Technische Universität München

Fakultät für Mathematik

Constrained Clustering via Generalized Voronoi Diagrams

Fabian Klemm

Vollständiger Abdruck der von der Fakultät für Mathematik der Technischen Universität München zur Erlangung des akademischen Grades eines

Doktors der Naturwissenschaften (Dr. rer. nat.)

genehmigten Dissertation.

Vorsitzender:	Prof. Dr. Michael Ulbrich
Prüfer der Dissertation:	1. Prof. Dr. Peter Gritzmann
	2. Prof. Dr. Martin Henk
	3. Prof. Jesús A. De Loera, Ph.D.

Die Dissertation wurde am 03.06.2020 bei der Technischen Universität München eingereicht und durch die Fakultät für Mathematik am 13.10.2020 angenommen.

Abstract

We are interested in grouping data into a given amount of k clusters under constraints that fix the aggregated values of data features for each cluster. A typical type of constraints is fixing the weights of clusters. We furthermore assume the data space to be equipped with a meaningful similarity measure that allows to judge the quality of a clustering. Exemplary applications are districting problems such as trade area planning or electoral district design.

Voronoi diagrams and their generalizations have been a powerful tool in various streams of research in data analysis and optimization. We define a general notion of Voronoi diagrams in arbitrary spaces and discuss several particular types in our context. A direct correspondence between generalized Voronoi diagrams and constrained clusterings is recalled and generalized. Clusters of certain optimality w. r. t. the space's similarity measure can be embedded into the cells of generalized Voronoi diagrams and thus to some extent inherit their geometric properties. Further, we recall some theory of (semi) infinite linear programming and transfer the established correspondence from the finite setting to arbitrary compact Hausdorff spaces. Also, we recall and extend the existing theory for balanced-k-means clustering from both a theoretical and algorithmic point of view.

As prime application, we consider the problem of electoral district design. We demonstrate our proposed methodology using three different types of Voronoi diagrams for the data of the federal elections in Germany.

Zusammenfassung

Wir beschäftigen uns mit der Gruppierung von Daten in eine gegebene Anzahl von k Gruppen (Cluster), so dass Nebenbedingungen bezüglich aggregierter Datenmerkmale für jeden Cluster eingehalten werden. Das Festsetzen von Clustergewichten ist hier ein typisches Beispiel. Darüber hinaus nehmen wir an, dass im Datenraum ein aussagekräftiges Ähnlichkeitsmaß vorliegt, welches eine Beurteilung der Qualität eines Clusterings zulässt. Anwendungsbeispiele liefern Aufteilungsprobleme, etwa für Handelsgebiete oder Wahlkreise.

Voronoi Diagramme und deren Verallgemeinerungen sind ein mächtiges Werkzeug, das in verschiedenen Forschungsgebieten der Datenanalyse und Optimierung eingesetzt wird. Wir definieren einen allgemeinen Voronoi Diagramm Begriff in beliebigen Räumen und diskutieren einige wichtige Spezialfälle im Hinblick auf unsere Methodik. Wir wiederholen und verallgemeinern eine direkte Beziehung zwischen verallgemeinerten Voronoi Diagrammen und Clusterings unter Nebenbedingungen. Cluster, die bezüglich des gegebenen Ähnlichkeitsmaßes in gewissem Sinne optimal sind, können in die Zellen eines verallgemeinerten Voronoi Diagramms eingebettet werden und erben so einige ihrer geometrischen Eigenschaften. Darüber hinaus wiederholen wir einige Theorie von (semi) unendlich-dimensionalen linearen Programmen und übertragen die beschriebene Beziehung für endliche Probleme auf beliebige kompakte Hausdorffräume. Außerdem besprechen und erweitern wir Ergebnisse für k-means Clustering unter Gewichtsnebenbedingungen, sowohl in theoretischer als auch algorithmischer Hinsicht.

Als Vorzeigebeispiel betrachten wir das Problem der Wahlkreiseinteilung. Wir demonstrieren unseren Ansatz anhand drei verschiedener Typen von Voronoi Diagrammen für Daten zur Bundestagswahl in Deutschland.

Contents

Chapter 1

Introduction

1.1 Scope

Clustering describes the task of grouping data in a meaningful way (see [GMW07], [DHS00, Chapter 10] for comprehensive introductions). As we are living in the age of information, or — more likely — at the very beginning of it, clustering self-evidently plays an important role in the day-to-day handling of the accompanying data amounts.

In the context of "big data" and "artificial intelligence", clustering is also considered as a *unsupervised* learning process that assigns labels to unlabeled data ([DHS00; HTF09]). Data elements with identical labels then form the clusters of a given data set. Thus, one may also consider clustering to belong to the huge field of machine learning. The theory of clustering algorithms and cluster analysis is, of course, wide-spread and ongoing (cf. the surveys in [JMF99; Sax+17]).

In order to decide on the meaningfulness of a clustering, similarities of the elements in between a cluster or dissimilarities between the clusters can be evaluated ([GMW07, Section 6]). For the context of this thesis, we assume data to be given in a meaningful geometric space that allows to inherit those measures of similarity or dissimilarity.

Let us explain this by means of the famous k-Means clustering algorithm ([Bis06, Section 9.1]). Here, one aims to group points to clusters such that the sum of squared errors, i. e., the sum of squared distances of points in a cluster to the cluster's centroid, is minimized. It is then easy to see that in any optimum, points are assigned to the cluster which yields the closest centroid. This reveals an important underlying geometric structure. If we assign each point in Euclidean space to the closest among a finite set of reference points, we obtain a tessellation of the space whose structure is well-known as a *Voronoi diagram* ([AKL13]). As one major characteristic, Voronoi diagrams consist of convex cells. This already implies a lot of information on the clusters obtained from the k-Means algorithm and hence provides a meaningful understanding of their basic geometric properties.

In this thesis, we generally aim at clusterings that allow this kind of structural knowledge, i. e., that are induced by diagrams. In order to be more flexible, however, we consider generalized versions of Voronoi diagrams. In their essence, Voronoi diagrams assign points to clusters that minimize a certain dissimilarity measure. Here, classical Voronoi diagrams consider the Euclidean norm and distances to reference points. By allowing basically any kind of dissimilarity — for example, by using different metrics or transformations of distances — one obtains new classes of generalized Voronoi diagrams. As long as the chosen distance measures are meaningful, one can hope to obtain a meaningful understanding of the tessellation that results from a diagram and hence of the induced clustering.

Next, we are interested in finding special clusterings that are further constrained. There are several ways how to constrain a clustering. An overview is given by Basu, Davidson, and Wagstaff [BDW08]. Here, the classical type of constraints are on an *instance level*, i. e., they constrain a relation between pairs of data elements. This

includes *must-link* or *cannot-link* constraints that enforce or disallow to group certain elements in the data into the same cluster, for example. Other types of constraints deal with properties of clusters as a whole. An important example are *balancing constraints* that define eligible sizes for the clusters. In this work, we are interested in a similar class of constraints. For each data element, we assume a vector of numerical features to be known. Such a feature can be a weight or coordinates of an element. We then assume that the feature values aggregated over each cluster are known a priori, and only allow clusterings that fulfill those aggregated values. Thus, our class of allowed constraints in particular implies balancing constraints.

Districting problems yield a well-suited class of applications in this setting ([Kal15], [Río20]). Here, areas have to be subdivided into smaller districts. Typically, those districts should be created with respect to certain side conditions that result in constraints or optimization criteria. Examples are police districts ([Mit72]), trade areas ([BS97]), farmland consolidation ([BG04]), or electoral districts ([RSS13]), to name just a few. We will use the latter as a prime example to demonstrate the proposed theory. Here, a state area has to be divided into electoral districts that each elect their own democratic representatives. A major constraint is that those districts are equally balanced in order to obey the one-man-one-vote principle. However, the decision what a "good" or even "admissible" electoral district plan is, can be tricky. Electoral districts in general should be "compact" in order to allow a good representation of their citizens and in order to prevent oddly-shaped districts that root from a partisanship-driven district design - more famously known as *gerrymandering*. Another important constraint is contiguity of electoral districts. However, this as well can be ambiguous in case that districts are separated by water or natural reserves, for example. Also, electoral districts have to be re-designed regularly due to census developments. In order to be politically acceptable, this redistricting should be as continuous as possible. Thus, this application reveals a wide range of constraints and optimality criteria.

Here, the usage of Voronoi diagrams and their generalizations will help us to cope with this variety of demands. Depending on the focus of optimality and feasibility criteria, we may choose a diagram type that in general yields a dissection of the plane into cells of desirable geometric properties that come along with this focus.

This work is intended to present a general framework for constrained clustering problems of the proposed type. Generalizations of Voronoi diagrams will yield the basic structure that embeds the resulting clusterings. As it turns out, there is a direct correspondence between those diagrams and clusterings that appear as optimizers of certain linear programs. In particular, the data that is to be clustered is not restricted to consist of a finite number of points in Euclidean spaces. Instead, both the concept of generalized Voronoi diagrams as well as the proposed correspondence to optimal clusterings carries over to very general settings.

1.2 Outline and Contributions

This work aims to provide a self-contained description of a methodological usage of generalized Voronoi diagrams for constrained clustering problems as motivated above.

A major starting point has been the results on the relation between power diagrams and least-squares weight-balanced clusterings, which in particular have been successfully applied to the problem of farmland consolidation ([Bor10; BBG11; BBG14; BG04; BG12]). As one major contribution, Chapters 2 to 5 generalize both the setting as well as the results of this theory in manifold ways.

Chapter 2 will define the notion of constrained clusterings. Sets of clusterings that obey a suitable type of constraints yield constrained clustering polytopes. Here, the special case of weight-balanced clusterings yield transportation polytopes which are well-understood and whose properties have been successfully exploited in the aforementioned balanced clustering problems. A particular important result is the limitation of the number of (strictly) fractional assignments of data points to several clusters. Those appear as we (temporarily) have to allow for fractional assignments to clusters in order to exploit a linear programming duality in the following. While any vertex of a weight-balanced clustering polytope yields at most $k-1$ many such fractional assignments, a very similar result can be shown for our general clustering polytopes as introduced in this chapter. Besides, this chapter will introduce clustering graphs as a helpful tool for dealing with clusterings and as well recall the existing results for weight-balanced clusterings which to some extent can be generalized, too.

Chapter 3 then introduces the concept of generalized Voronoi diagrams. Of course, Voronoi diagrams and their generalizations are subject to a lot of research on their own. We will give a concise introduction, provide an abstract definition of generalized Voronoi diagrams for our context, and discuss some basic properties. We then discuss some basic types that generalize the concept of classical Voronoi diagrams in Euclidean space as well as so-called shortest-path diagrams over graphs. As of certain importance for the applications in our mind, we particularly discuss different parametrizations of power and anisotropic power diagrams. In particular, we provide results on the invariances of those parametrizations.

In Chapter 4, we then generalize the aforementioned correspondence between weight-balanced clusterings and power diagrams to our setting. We recall the existing findings in the literature. We then successively follow the steps of generalizations as existing in the literature in order to finally derive our — to the best of our knowledge — further generalized correspondence. This relates the constrained clusterings of Chapter 2 to the concept of generalized Voronoi diagrams from Chapter 3. By exploiting standard linear-programming duality theory, we show that specific faces of a constrained clustering polytope yield exactly those clusterings that may be derived from generalized Voronoi diagrams. More practically, we may consider an assignment of data points to clusters that is optimal w. r. t. the aggregated costs that result from distance functions associated

with each cluster. The corresponding dual solution then yields the linear factors for combining the constraint-defining feature maps. The latter yield correction terms for the cluster-associated distance measures such that the resulting generalized Voronoi diagram then induces the optimal clustering.

The results of Chapter 4 are based on basic linear programming theory. This particularly implies a setting of a finite amount of data points to be clustered. However, it is natural to ask to which extent this can be extended to a continuous setting. For example, instead of clustering a finite number of points in the plane one may way want to find a clustering of the plane itself. Instead of having weighted points one may have a continuous probability distribution over the plane. Chapter 5 is intended to break down the correspondence as established in the preceding chapter to its very core by considering a way more general setting. Here, we will assume any compact Hausdorff space equipped with some regular Borel measure to yield the dataset to be clustered. In order to keep this thesis self-contained, this chapter first recalls some necessary definitions from topological and functional analysis, and then gives an introduction to linear programming in locally convex spaces. We then carry our terminology and problem over to this more general setting. As the finite case for balanced clusterings results in transportation problems, we may then generalize our duality results from the previous section similar to existing findings on semi-infinite transportation problems. However, we will see that the correspondence between generalized Voronoi diagrams and optimal clusterings as proven for the finite case in Chapter 4 does not hold in general in this new setting. Still, we can show that by allowing an arbitrary small error in the cluster constraints, the correspondence can be re-established.

An important and already well-researched special case of our setting are least-squares balanced clusterings. As this yields the weight-constrained version of the famous k-Means clustering problem, we refer to it as Balanced k-Means clustering. Chapter 6 intends to provide a cohesive algorithmic treatment of this problem. We define two variants of the Balanced k-Means problem and provide several complexity results. We then outline an algorithmic approach to it. First, the Balanced k-Means-algorithm from the literature yields the constrained counterpart of the famous k-Means algorithm. Next, we recall and adapt existing theory on so-called gravity polytopes that allows to interpret the Balanced k-Means problem as an ellipsoidal norm-maximization problem. We then present a revised version of the Balanced k-Means-algorithm that guarantees convergence to a local optimum. Also, using our adapted theoretical results we are able to slightly improve its guaranteed running time.

Chapter 7 presents our prime application of the electoral district design problem by means of a case study for the federal elections in Germany. First, we shortly introduce and discuss the problem of electoral district design. We then formulate a general methodological framework that allows to apply several types of generalized Voronoi diagrams. By means of a toy example, we then discuss the applicability of power diagrams, shortest-path diagrams and anisotropic power diagrams. Here, the

choice of the diagram type can be made w. r. t. the desired focus of optimization. The shortest-path diagram approach is the only one that guarantees contiguous districts a priori. Finally, anisotropic power diagrams can be exploited in order to aim at a district design that is similar to an existing (but maybe infeasible) district design. We present detailed computational results. Those demonstrate that all approaches are capable of providing pleasant district designs. Also, they underline how the choice of the diagram type sets the focus w. r. t. the various optimization criteria.

1.3 Prerequisites and Notations

Besides a general mathematical background, this thesis assumes the reader to be familiar with some fundamental concepts related to the field of discrete optimization. This concerns basic terms from convex geometry such as polytopes, halfspaces and cones (see [Roc70]). Also, a fundamental knowledge of the theory of linear programming (cf. [Sch98; Van08]), discrete mathematics (such as graphs, see [Die17]) as well as combinatorial optimization (including basics of complexity theory, see, for example, [KV12; Pap82]) is assumed. Whenever ambiguous or beyond this expected scope, this thesis strives to provide the necessary definitions and terminology in order to yield a self-contained work.

Let us clarify some general notation used in this thesis. In addition, an overview at page 251 is to provide a look-up possibility of the most important notations and objects as used throughout this thesis.

The set $\mathbb{N}$ denotes the set of all non-negative integers including 0, while $\mathbb{N}_{>0}$ denotes the set of all strictly positive integer numbers. For an integer $k \in \mathbb{N}$, the set $[k] := \{1, \dots, k\}$ denotes the set of the first k integers, with the convention $[0] := \emptyset$. For a finite set S, we denote by $\binom{S}{k}$ the set of all subsets of S of cardinality k. The sets of all, all non-negative, and all strictly positive reals are denoted by $\mathbb{R}$, $\mathbb{R}_{\geq 0}$, and $\mathbb{R}_{>0}$, respectively.

Let $d \in \mathbb{N}_{>0}$ denote some dimension. The set $\mathcal{SO}(d)$ denotes the special orthogonal group of all orthogonal matrices of determinant 1 in $\mathbb{R}^{d\times d}$. The matrix $\mathrm{Id}_d \in \mathbb{R}^{d\times d}$ yields the identity matrix and for a vector $v \in \mathbb{R}^d$ the diagonal matrix with the components of v on its diagonal is denoted by $\mathrm{diag}(v)$. The vector $u^{(i)}$ will denote the ith standard unit vector (with the dimension d being clear by context). We write $\mathbb{1}^{(d)} := (1, \dots, 1)^\intercal \in \mathbb{R}^d$ for the vector with all components equal to 1 and $0^{(d)} := (0, \dots, 0)^\intercal \in \mathbb{R}^d$ for the vector with all components equal to 0 (we omit the dimension d if clear by context).

For a matrix $A \in \mathbb{R}^{n\times m}$, we denote by $A_{i,\cdot}$ its ith row, and by $A_{\cdot,j}$ its jth column. For some $S \subset [d]$, the vector $A_{\cdot,S} \in \mathbb{R}^{|S|}$ denotes the submatrix of A obtained by restricting A to the columns with index in S (with the analogous notation used for restricting rows). Labeled vectors in $\mathbb{R}^d$ will be indexed by super indices such as $x^{(i)}$

in order to prevent ambiguities with component references.

For a set $S \subset \mathbb{R}^d$, the sets $\operatorname{conv}(S)$, $\operatorname{lin}(S)$, and $\operatorname{aff}(S)$ denote the smallest convex set, linear subspace, and affine subspace that contain S, respectively. For some $a \in \mathbb{R}^d$ and $\beta \in \mathbb{R}$, the set $H^{\leq}_{(a,\beta)} := \left\{x \in \mathbb{R}^d : x^\intercal a \leq \beta\right\}$ denotes a closed halfspace. Analogously we define the closed or open halfspaces $H^{\geq}_{(a,\beta)}, H^{<}_{(a,\beta)}, H^{>}_{(a,\beta)}$, as well as the hyperplane $H_{(a,\beta)} := H^{=}_{(a,\beta)}$. A polyhedron is the intersection of a finite amount of closed halfspaces and called polytope if its furthermore bounded. We assume the reader for be familiar with the notion of faces of convex sets. For a convex set C, we denote by $\operatorname{ext}(C)$ the set of all extremal points, i. e., 0-faces of C (that we call vertices if C is a polyhedron).

1.4 Acknowledgments

A lot of the results of Chapter 7 were obtained in joint work with Prof. Dr. Peter Gritzmann and Prof. Dr. Andreas Brieden, and have been published in [BGK17].

Of course, the whole of this thesis has profited from the many rewarding conversations and discussions with my supervisor, Peter Gritzmann. Besides our scientific cooperation, I want to thank him for giving me opportunity to join his group, as well as his support, kindness, and advise throughout my years at *M9*, the chair for applied geometry and discrete mathematics at the TUM math department. Thank you, Peter, sincerely.

Life in a research group such as M9 comes along with the — maybe once-in-a-lifetime — opportunity of working with the brightest and kindest minds. My gratitude for this experience goes out to the current, former, and associated people of M9. Among all I would like to mention Katherina von Dichter, Raymond Hemmecke, Melanie Herzog, Wei Huang, Stefan Kobler, Nicole Megow, Bernardo Gonzáles Merino, Benedikt Plank, Kevin Schewior, Felix Schmiedl, Andreas S. Schulz, Ina Seidel, Matthias Silbernagel, Annika Treyer, Stefan Weltge, Mel Zürcher. I am grateful for our endless conversations during lunch, conferences, team events, teaching, or the kitchen — all of which I will miss, for sure.

I want to thank the colleagues with whom I shared an office: Tobias Windisch, for our joyful talks in the late office hours, and Anja Kirschbaum, for her empathy, kindness and support. I would like to thank Steffen Borgwardt for introducing me to the field of constrained clusterings and power diagrams in the very beginning. Special thanks go to René Brandenberg for taking care of everything concerning my position at M9 and our enjoyable teamwork in our common teaching units. I thank Michael Ritter for his advise, unconditional support, and for our priceless teaching experiences with *his* case studies of discrete optimization. I am grateful for my fellow students, colleagues, and friends Paul Stursberg and Wolfgang Riedl, who involuntarily had to serve as reference persons to me. I am thankful to Andreas Alpers for our many helpful discussions on anisotropic power diagrams in mixture with sharing his insights on team management in volleyball. I thank Andreas Brieden for our fruitful joint work that

was especially enlightened by his inexhaustible optimism. I wish to thank the people who particularly helped me in the final phase of the thesis, namely Carolin Bauerhenne and Maximilian Fiedler (also for our joyful joint work, of course). I am deeply in dept with Viviana Ghiglione for her thoughtfulness, invaluable support, honesty, caring, and serving the strongest coffee at just the right time.

It was my pleasure to serve as representative for the doctoral candidates at TUM mathematics for some years. I am glad to have shared discussions and events with the countless people from our department, the ISAM, the TUM Graduate School, and the TUM Graduate Council. Warm greetings go out to Felix Dietrich, Patrick Gonthar, Simon Plazotta, Katharina Schaar, and so many more. I want to thank the former and current ISAM coordinators Sonya Gzyl and Isabella Wiegand, as well as our dean's office referent Lydia Weber, who all have done a great job in alleviating the bureaucratic pressure from the doctoral candidates as much as possible.

I want to thank my family and my friends for their support throughout the years while not asking too many questions at the wrong times! Finally, there are no words to express my gratitude to my wife Silvia, for her endless patience and support.

Chapter 2

Constrained Clustering

In this chapter, we define the terminology of constrained clusterings for the scope of this thesis.

In order to so, Section 2.1 defines constrained clustering polytopes. Some elementary properties of those are then discussed in Section 2.2. Section 2.3 then introduces the notion of clustering graphs which will serve as a useful tool to work with clusterings later on.

2.1 Definition

The clustering problems we are interested in deal with the task of partitioning a set into a fixed number of k subsets. Those subsets — called *clusters* in our context — obey certain constraints. These constraints may, for example, be of the form of fixed weights of the resulting clusters w. r. t. a given weighting of the elements in our set of interest. Another example might be fixing the clusters' centroids if we are to partition points in $\mathbb{R}^d$. In general, we consider constraints that can be expressed as a fixed average or sum over the clusters. Also, we do allow elements to be fractionally assigned to several clusters. This gives rise to the following very broad definition.

Definition 2.1 (Constrained Fractional Clustering)
Let $n, k \in \mathbb{N}_{>0}$. Furthermore, let for every $i \in [k]$ a number $m_i \in \mathbb{N}$, a matrix $A_i \in \mathbb{R}^{m_i \times n}$, and a vector $b^{(i)} \in \mathbb{R}^{m_i}$ be given.

We set

$$T := T\left((A_i)_{i\in[k]}, (b^{(i)})_{i\in[k]}\right) := \{\xi = (\xi_{i,j})_{\substack{i\in[k],\\ j\in[n]}} \in \mathbb{R}^{k\times n} :$$

$$\xi^\intercal \mathbb{1}^{(k)} = \mathbb{1}^{(n)} \tag{2.0a}$$

$$A_i \xi_{i,\cdot}^\intercal = b^{(i)} \quad \forall i \in [k] \tag{2.0b}$$

$$\xi_{i,j} \geq 0 \quad \forall i \in [k], j \in [n]\}. \tag{2.0c}$$

Then we call $\xi \in T$ a $((A_i)_{i\in[k]}, (b^{(i)})_{i\in[k]})$-*constrained (fractional) clustering*. If $\xi \in T \cap \mathbb{Z}^{k\times n}$, we call the clustering *integer*.

For each $i \in [k]$ we set $C_i := (\xi_{i,1}, \ldots, \xi_{i,n})^\intercal \in \mathbb{R}^n$ and call it (the ith) *cluster*. Next, for each cluster C_i we denote by

$$\operatorname{supp}(C_i) := \{j \in [n] : \xi_{i,j} \neq 0\} \tag{2.1}$$

the *support* of the ith cluster.

We will call the case $A_1 = \ldots = A_k =: A \in \mathbb{R}^{m\times n}$ *homogeneously constrained*. Then we write $T := T\left(A, (b^{(i)})_{i\in[k]}\right)$ and call $\xi \in T$ a $(A, (b^{(i)})_{i\in[k]})$-constrained clustering.

The case $m = 1$ will be referred to as *single-constrained*. For the sake of a uniform notation, we define

$$T_{\text{uncstr.},k,n} := T\left((0^{(n)})^\intercal, (0)_{i\in[k]}\right) = \left\{\xi \in \mathbb{R}^{k\times n} : \xi^\intercal \mathbb{1}^{(k)} = \mathbb{1}^{(n)}, \xi \geq 0\right\}. \tag{2.2}$$

This is the n-fold product of $(n-1)$-simplices, and we will refer to $\xi \in T_{\text{uncstr.},k,n}$ as an *unconstrained* clustering.

Definition 2.1 defines a clustering $\xi = (\xi_{i,j})_{i,j}$ as a matrix in $\mathbb{R}^{k\times n}$. Using the Frobenius inner product the latter is a Hilbert space that can be identified with the Euclidean space $\mathbb{R}^{k\cdot n}$. Hence, for the sake of simpler denotation we will occasionally identify a clustering ξ with a vector in $\mathbb{R}^{k\cdot n}$ (whenever this is clear from the context). In particular, we identify the set T from Definition 2.1 with a polytope in $\mathbb{R}^{k\cdot n}$ and call it *(constrained) clustering polytope*.

Example 2.2.
Consider the set of points $X = \left\{x^{(1)}, \dots, x^{(6)}\right\}$ with $x^{(1)} := \begin{pmatrix}0\\2\end{pmatrix}$, $x^{(2)} := \begin{pmatrix}1\\3\end{pmatrix}$, $x^{(3)} := \begin{pmatrix}2\\2\end{pmatrix}$, $x^{(4)} := \begin{pmatrix}7/2\\2\end{pmatrix}$, $x^{(5)} := \begin{pmatrix}3\\1\end{pmatrix}$, and $x^{(6)} := \begin{pmatrix}3/2\\1/2\end{pmatrix}$. Furthermore, the points shall have weights $\omega_1 := 1$, $\omega_2 := 2$, $\omega_3 := 1$, $\omega_4 := 1$, $\omega_5 := 1$, and $\omega_6 := 1$.

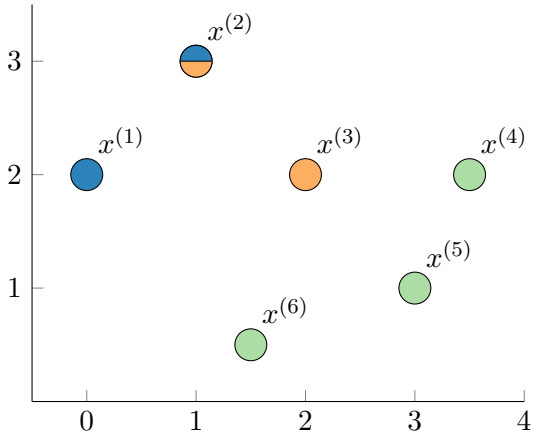

Figure 2.1: Exemplary constrained fractional clustering from Example 2.2.

Let us assume we are interested in finding a clustering of X into $k := 3$ clusters of prescribed weights $\kappa_1 := 2$, $\kappa_2 := 2$, and $\kappa_3 := 3$.

As all clusters have to obey the same type of constraint, we are in the homogeneously single-constrained case. Thus, we set $A := \begin{pmatrix}\omega_1 & \dots & \omega_6\end{pmatrix}$ and $b^{(i)} := \kappa_i$ for every $i \in [3]$ and obtain with $T := T\left(A, (b^{(i)})_{i\in[k]}\right)$ as defined in Definition 2.1

the polytope of all $\left(A, (b^{(i)})\right)$-constrained clusterings. Fig. 2.1 depicts the clustering $\xi = \begin{pmatrix} 1 & \frac{1}{2} & 0 & 0 & 0 & 0 \\ 0 & \frac{1}{2} & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 1 & 1 \end{pmatrix} \in T$ with colors blue/orange/green for clusters 1/2/3, respectively. In particular, this clustering is not integer as $x^{(2)}$ is split to equal parts between cluster 1 and 2. Here, we have $\mathrm{supp}(C_1) = \{1, 2\}$, $\mathrm{supp}(C_2) = \{2, 3\}$, and $\mathrm{supp}(C_3) = \{3, 4, 5\}$.

If it is clear from the context, we will refer to a cluster C_i simply by its index i. Also, we will refer to the elements that are to be clustered as *units*. Units can be points (such as in Example 2.2), the set of numbers $[n]$ or any other set of objects. Again, we will refer to the unit of index j simply as unit j, if the reference is clear.

Example 2.2 introduces a very important special class of homogeneously-constrained clustering polytopes that fall into the class of *transportation polytopes*. Those received major interest as description of the feasible region of the *Hitchcock Transportation Problem* or *Hitchcock-Koopmans Problem* in the 40s of the last century (cf. [Hit41],[Koo49], [KW68], [YKK84, Chapter 6]). Here, one searches for the optimal transport of a homogeneous good from sources with supply amounts $(a_j)_{j\in[n]}$ to destinations of demands $(b_i)_{i\in[k]}$, with costs linear in the transportation amount between each source-terminal pair. This may also be understood as a network flow problem in a complete bipartite graph between source and terminal nodes ([Van08, Chapter 15]). In a more general setting, the *Monge-Kantorovich transportation problem* seeks for an optimal probability measure for the product space of two measurable spaces with prescribed marginal distributions (see [Mon81] and [Kan60] for original publications as cited in [Kol+17]; see [Vil09] for a thorough treatment of optimal transport theory; further, application-oriented introductions can be found in [Kol+17] and [PC19]). The special case of both measurable spaces being finite then yields the previously described transportation problem.

Besides for classical transportation problems, transportation polytopes are also of interest in the field of statistics. Here, *contingency tables* describe the possible outcome from sampling multivariate discrete random variables with fixed marginal sums. In this context, transportation polytopes are, for example, researched w. r. t. the question of how much information may be drawn about individuals when knowing cumulated numbers (see, for example, [CCE85], [Cho+99], [BFH07], and the references therein). Consequently, transportation polytopes have been extensively studied and are still matter of current research (see [KW68], [Bol72], [YKK84, Chapters 6 - 8], [Bru06, Chapter 8], [DK14] for collections of the most relevant results).

In our setting in the homogeneously single-constrained case with only positive parameters, i. e., $m = 1$ and $A = (a_1, \ldots, a_n) \in \mathbb{R}_{>0}^{1\times n}$ and $b^{(i)} \in \mathbb{R}_{>0}$ for all $i \in [k]$, we obtain a classical transportation polytope by rescaling $T := T\left(A, (b^{(i)})_{i\in[k]}\right)$ as follows.

Let $D : \mathbb{R}^{k\times n} \to \mathbb{R}^{k\times n}$ be the scaling $(\xi_{i,j})_{i,j} \mapsto (\frac{1}{a_j}\xi_{i,j})_{i,j}$, then

$$D \cdot T = \left\{ \left(\frac{1}{a_j}\xi_{i,j}\right)_{i,j} \in \mathbb{R}^{k\times n} : \sum_{i\in[k]} \xi_{i,j} = 1 \ \forall j \in [n] \ \wedge \ \sum_{j\in[n]} a_j\xi_{i,j} = b^{(i)} \forall i \in [k] \right\}$$
$$= \left\{ \left(\hat{\xi}_{i,j}\right)_{i,j} \in \mathbb{R}^{k\times n} : \sum_{i\in[k]} \hat{\xi}_{i,j} = a_j \ \forall j \in [n] \ \wedge \ \sum_{j\in[n]} \hat{\xi}_{i,j} = b^{(i)} \forall i \in [k] \right\}$$

is a transportation polytope in the classical definition. Consequently, many of the following results are obtained using basic facts about transportation polytopes or contain them as special cases.

2.2 Properties of Clustering Polytopes

Let us discuss some basic properties of constrained clustering polytopes.

2.2.1 Feasibility

First, we will investigate the conditions under which a clustering polytope T is non-empty. We will restrict this to the homogeneously constrained case and consider some homogeneously constrained clustering polytope $T = T\left(A, (b^{(i)})_{i\in[k]}\right)$ in this section. We will furthermore denote $A = \begin{pmatrix} a^{(1)} & \dots & a^{(n)} \end{pmatrix} \in \mathbb{R}^{m\times n}$ and $b := \left((b^{(1)})^\intercal, \dots, (b^{(k)})^\intercal\right)^\intercal \in \mathbb{R}^{k\cdot m}$.

For the case of transportation polytopes, i. e., $m = 1$, it is well-known that $T\left(A, (b^{(i)})_{i\in[k]}\right)$ is non-empty if and only if $\sum_{j\in[n]} a^{(j)} = \sum_{i\in[k]} b^{(i)}$. However, this is not sufficient anymore as soon as the non-negativity of the $a^{(j)}$ and $b^{(i)}$ is dropped. For example, setting $A := (-1, 1)$ and $(b^{(1)}, b^{(2)}) := (-2, 2)$ obviously leads to T being empty.

While the general homogeneous case for $m = 1$ may be determined easily, we can only provide some necessary conditions on the choice of A and $(b^{(i)})_{i\in[k]}$ in higher dimensions.

Proposition 2.3
In the setting as above, the following holds:

a) In the case $m = 1$ it holds that $T \neq \emptyset$ if and only if

$$\sum_{\substack{j\in[n]:\\ a^{(j)}\leq 0}} a^{(j)} \leq \sum_{\substack{i\in[k]:\\ b^{(i)}\leq 0}} b^{(i)} \quad \text{and} \quad \sum_{i\in[k]} b^{(i)} = \sum_{j\in[n]} a^{(j)}.$$

b) In general, if $T \neq \emptyset$, then for all $u \in \mathbb{R}^m$ it follows that

$$\sum_{j\in[n]} u^\intercal a^{(j)} = \sum_{i\in[k]} u^\intercal b^{(i)} \quad \text{and} \quad \sum_{\substack{j\in[n]:\\ u^\intercal a^{(j)}\leq 0}} u^\intercal a^{(j)} \leq \sum_{\substack{i\in[k]:\\ u^\intercal b^{(i)}\leq 0}} u^\intercal b^{(i)}.$$

Proof. For part a, let $I^+ := \left\{i \in [k] : b^{(i)} \geq 0\right\}$, $I^- := \left\{i \in [k] : b^{(i)} < 0\right\}$, $J^+ := \left\{j \in [n] : a^{(j)} > 0\right\}$, and $J^- := \left\{j \in [n] : a^{(j)} < 0\right\}$. We may w.l.o.g. assume that $a^{(j)} \neq 0$ for all $j \in [n]$ (as those columns of A obviously do not affect feasibility) and hence $J^- \cup J^+ = [n]$. Next, we denote $b^+ := \sum_{j\in I^+} b^{(j)}$, $b^- := \sum_{i\in I^-} b^{(i)}$, $a^+ := \sum_{j\in J^+} a^{(j)}$, and $a^- := \sum_{j\in J^-} a^{(j)}$. Our assumption hence translates into

$$a^- \leq b^- \quad \text{and} \quad b^- + b^+ = a^- + a^+. \tag{2.3}$$

Now if Eq. (2.3) holds, we first set $r := a^+ - b^+ = b^- - a^- \geq 0$. Then, we define $\xi \in \mathbb{R}^{k\times n}$ as

$$\xi_{i,j} := \frac{-b^{(i)} + r/k}{-a^-} \quad \forall i \in I^-, j \in J^-, \quad \xi_{i,j} := \frac{r/k}{-a^-} \quad \forall i \in I^+, j \in J^-$$
$$\xi_{i,j} := \frac{b^{(i)} + r/k}{a^+} \quad \forall i \in I^+, j \in J^+, \quad \xi_{i,j} := \frac{r/k}{a^+} \quad \forall i \in I^-, j \in J^+.$$

Note that this is well-defined as a^- (a^+) being zero implies that J^- (J^+) is empty. We can now easily verify that for each $j \in J^-$ it holds that

$$\sum_{i\in[k]} \xi_{i,j} = \sum_{i\in I^-} \frac{-b^{(i)} + r/k}{-a^-} + \sum_{i\in I^+} \frac{r/k}{-a^-} = \frac{-b^- + k \cdot \frac{r}{k}}{-a^-} = 1$$

and for every $i \in I^-$ that

$$\sum_{j\in[n]} a^{(j)}\xi_{i,j} = \sum_{j\in J^-} a^{(j)} \cdot \frac{-b^{(i)} + r/k}{-a^-} + \sum_{j\in J^+} a^{(j)} \cdot \frac{r/k}{a^+} = b^{(i)} - r/k + r/k = b^{(i)}.$$

The analogous check for $j \in J^+$ and $i \in I^+$, respectively, together with the observation that $\xi \geq 0$ then yields that indeed $\xi \in T$ holds.

For the other implication, let $\xi \in T$. Then we get

$$\sum_{i\in[k]} b^{(i)} = \sum_{i\in[k]} \sum_{j\in[n]} a^{(j)}\xi_{i,j} = \sum_{j\in[n]} a^{(j)}(\sum_{i\in[k]} \xi_{i,j}) = \sum_{j\in[n]} a^{(j)}$$

as well as

$$\begin{aligned} b^- &= \sum_{i\in I^-} b^{(i)} = \sum_{i\in I^-}\sum_{j\in[n]} a^{(j)}\xi_{i,j} \geq \sum_{i\in I^-}\sum_{j\in J^-} a^{(j)}\xi_{i,j} \\ &= \sum_{j\in J^-} \underbrace{a^{(j)}}_{\leq 0} \cdot \underbrace{(\sum_{i\in I^-} \xi_{i,j})}_{\leq 1} \geq \sum_{j\in J^-} a^{(j)} = a^-, \end{aligned}$$

and so Eq. (2.3) holds.

For part b, note that for every $u \in \mathbb{R}^m$ it follows that

$$T\left(A, (b^{(i)})_{i\in[k]}\right) \subset T\left(u^\intercal A, (u^\intercal b^{(i)})_{i\in[k]}\right).$$

Hence, this follows from part a. □

In order to get a more general understanding, let us rewrite the set of all feasible right-hand sides b. If there exists $\xi \in T$, it follows that

$$\begin{pmatrix} b^{(1)} \\ \vdots \\ b^{(k)} \end{pmatrix} \overset{(2.0b)}{=} \begin{pmatrix} A\xi_{1,\cdot}^\intercal \\ \vdots \\ A\xi_{k,\cdot}^\intercal \end{pmatrix}$$

$$= \sum_{j\in[n]} \begin{pmatrix} \xi_{1,j}a^{(j)} \\ \vdots \\ \xi_{k,j}a^{(j)} \end{pmatrix} \overset{\substack{(2.0a),\\(2.0c)}}{\in} \sum_{j\in[n]} \operatorname{conv}\left(\left\{\begin{pmatrix} a^{(j)} \\ 0 \\ \vdots \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ a^{(j)} \\ 0 \\ \vdots \end{pmatrix}, \ldots, \begin{pmatrix} 0 \\ \vdots \\ 0 \\ a^{(j)} \end{pmatrix}\right\}\right).$$

Vice versa, we get that the coefficients of any convex combination in each of the convex hulls in the Minkowski sum above yield a feasible clustering $\xi \in T$ by definition. Hence, the set of right-hand sides b such that T is non-empty is given by

$$B := \operatorname{conv}\left(\sum_{j\in[n]} \left\{\begin{pmatrix} a^{(j)} \\ 0 \\ \vdots \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ a^{(j)} \\ 0 \\ \vdots \end{pmatrix}, \ldots, \begin{pmatrix} 0 \\ \vdots \\ 0 \\ a^{(j)} \end{pmatrix}\right\}\right). \tag{2.4}$$

Let $S^{(k-1)} := \left\{x \in \mathbb{R}^{k-1} : x \geq 0 \wedge (\mathbb{1}^{(k-1)})^\intercal x \leq 1\right\}$ denote the standard simplex in $\mathbb{R}^{k-1}$ and $\pi_{[(k-1)m]} : \mathbb{R}^{k\cdot m} \to \mathbb{R}^{(k-1)\cdot m}$ the projection of $\mathbb{R}^{k\cdot m}$ onto the first $(k-1)\cdot m$ coordinates. Identifying $\mathbb{R}^{(k-1)\cdot m}$ with $\mathbb{R}^{(k-1)\times m}$ we see that

$$\pi_{[(k-1)m]}(B) = \sum_{j\in[n]} S^{(k-1)}(a^{(j)})^\intercal.$$

Since $S^{(k-1)}$ is full-dimensional, it follows that $\dim(\pi_{[(k-1)m]}(B)) = (k-1) \cdot \operatorname{rank}(A)$. Furthermore, for $b \in B$ and $\xi \in T$ we have that $b^{(k)} = A\xi_{k,\cdot} = A(\mathbb{1}^{(n)} - \sum_{i\in[k-1]} \xi_{i,\cdot}) = \sum_{j\in[n]} a^{(j)} - \sum_{i\in[k-1]} b^{(i)}$.

Consequently, it holds that $\dim(B) = (k-1)\operatorname{rank}(A)$ and

$$\operatorname{aff}(B) = \left\{ \tilde{b} \in \mathbb{R}^{k\cdot m} : \sum_{i\in[k]} \tilde{b}^{(i)} = \sum_{j\in[n]} a^{(j)} \right\}.$$

Of course, B is invariant under a permutation of the k blocks of dimension m (or the rows in each matrix in B, when we again identify $\mathbb{R}^{k\cdot m}$ with $\mathbb{R}^{k\times m}$). Furthermore, with $\pi_{(i-1)m+[m]}$ denoting the projection of $\mathbb{R}^{k\cdot m}$ to the ith m-dimensional block, it holds that

$$\pi_{(i-1)m+[m]}(B) = \sum_{j\in[n]} [0,1]a^{(j)}.$$

This is, $\pi_{(i-1)m+[m]}(B)$ is the zonotope of the columns of A (cf. Fig. 2.2).

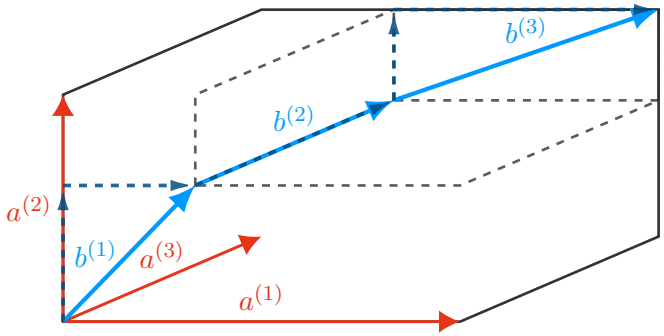

Figure 2.2: Exemplary zonotope $\pi_{(i-1)m+[m]}(B)$. The dashed zonotopes illustrate the feasible choices for $b^{(i)}$ for $i \in [k-1]$ after $b^{(1)}, \ldots, b^{(i-1)}$ have been fixed.

2.2.2 Results on Fractionality

Of course, allowing a fractional assignment of units to clusters only yields a relaxation in case that the model on hand actually requires binary assignments. However, using the structure of a clustering polytope T one can show that for extremal solutions the number of units that are non-integrally assigned can be reasonably bounded. For the single-homogeneously-constrained case, i. e., $m = 1$ and $T = T\left(A, (b^{(i)})_{i\in[k]}\right)$, it is well known that clusterings that are vertices of T may not assign more than $k-1$ units non-integrally (see [BG12; GLW97; Hoj96; KNS05; Mar81; SW77; ZS83]).

In order to show the direct generalization of this result to our setting, we first require the following technical lemma that yields the rank of the matrix of equality constraints in the definition of clustering polytopes.

Lemma 2.4
Consider the setting of Definition 2.1, i. e., let $A_i \in \mathbb{R}^{m_i \times n}$ *and* $b^{(i)} \in \mathbb{R}^{m_i}$ *for* $i \in [k]$ *be given and consider* $T := T\left((A_i)_{i\in[k]}, (b^{(i)})_{i\in[k]}\right)$. *Furthermore, let* $\operatorname{rank}(A_i) = m_i$ *hold for all* $i \in [k]$.

Let $\overline{A}$ *denote matrix that defines the left-hand sides of Equations* (2.0a) *and* (2.0b). *Then*

$$\operatorname{rank}(\overline{A}) = n + \sum_{i=1}^{k} m_i - \dim\left(\bigcap_{i=1}^{k} \operatorname{lin}(A_i^\intercal)\right).$$

Here, $\operatorname{lin}(A_i^\intercal)$ *denotes the linear space spanned by the rows of* A_i.

Proof. Considering T to be a subset of $\mathbb{R}^{k\cdot n}$ as described earlier, the left-hand side's matrix of the linear equality system given by the equations (2.0a) and (2.0b) can be written as

$$\overline{A} := \begin{pmatrix} \mathrm{Id}_n & \mathrm{Id}_n & \dots & \mathrm{Id}_n \\ A_1 & & & \\ & A_2 & & \\ & & \ddots & \\ & & & A_k \end{pmatrix}. \tag{2.5}$$

We can now perform elementary row operations on $\overline{A}$. For this purpose, we set $\mathcal{L}_i := \bigcap_{l=1}^{i} \operatorname{lin}(A_l^\intercal)$ for $i \in [k]$ and w.l.o.g. assume that for $i \geq 2$ there exist (possibly empty) matrices A_i^+, A_i^- such that $A_i = \begin{pmatrix} A_i^+ \\ A_i^- \end{pmatrix}$, $\operatorname{lin}((A_i^+)^\intercal) = \operatorname{lin}(A_i^\intercal) \cap \mathcal{L}_{i-1}^\perp$, and $\operatorname{lin}((A_i^-)^\intercal) = \operatorname{lin}(A_i^\intercal) \cap \mathcal{L}_{i-1} = \mathcal{L}_i$ (otherwise, we perform elementary row operations

on A_i). Gaussian elimination then yields

$$\begin{pmatrix} \mathrm{Id}_n & \mathrm{Id}_n & \dots & \mathrm{Id}_n \\ A_1 & & & \\ & A_2 & & \\ & & \ddots & \\ & & & A_k \end{pmatrix} \rightsquigarrow \begin{pmatrix} \mathrm{Id}_n & \mathrm{Id}_n & \dots & \mathrm{Id}_n \\ 0 & -A_1 & \dots & -A_1 \\ & A_2 & & \\ & & \ddots & \\ & & & A_k \end{pmatrix} \overset{\substack{\mathrm{lin}((A_2^-)^\intercal)\subset \\ \mathcal{L}_1 = \mathrm{lin}(A_1^\intercal)}}{\rightsquigarrow} \begin{pmatrix} \mathrm{Id}_n & \mathrm{Id}_n & \mathrm{Id}_n & \dots & \mathrm{Id}_n \\ & -A_1 & -A_1 & \dots & -A_1 \\ & A_2^+ & & & \\ & 0 & -A_2^- & \dots & -A_2^- \\ & & A_3 & & \\ & & & \ddots & \\ & & & & A_k \end{pmatrix}$$

$$\rightsquigarrow \cdots \rightsquigarrow \begin{pmatrix} \mathrm{Id}_n & \mathrm{Id}_n & \mathrm{Id}_n & \dots & \mathrm{Id}_n \\ & -A_1 & -A_1 & \dots & -A_1 \\ & A_2^+ & & & \\ & & -A_2^- & \dots & -A_2^- \\ & & A_3^+ & & \\ & & & \ddots & \\ & & & & A_k^+ \\ & & & & 0 \end{pmatrix}.$$

As all A_i have full (row) rank, it follows that

$$\mathrm{rank}(\overline{A}) = n + \sum_{i=1}^{k} m_i - \mathrm{rank}(A_k^-) = n + \sum_{i=1}^{k} m_i - \dim\left(\bigcap_{i=1}^{k} \mathrm{lin}(A_i^\intercal)\right). \qquad \square$$

In the particular case of homogeneously single-constrained clusterings, Lemma 2.4 states that the rank of the equality system defining T is $n + k - 1$, which is, of course, in consistency with the according results for transportation polytopes. For the latter, the relative interior consists of points with strictly positive components and thus their dimension is $n \cdot k - \mathrm{rank}(\overline{A}) = (n-1)(k-1)$ ([YKK84, Chapter 6, Proposition 1.1]). Note that — as we do not make any further assumptions on a clustering polytope's parameters — this may not be true in our setting, as the following short example shows.

Example 2.5
Let $n := k := 2$, $A := (b^{(1)}, b^{(2)}) := (-1, 1)$. Then $T := T\left(A, (b^{(i)})_{i\in[k]}\right) = \{\mathrm{Id}_2\}$ and thus $\dim(T) = 0 < 1 = (n-1)(k-1)$.

However, as $\mathrm{rank}(\overline{A})$ yields the size of any basis of a vertex of T, we make the following definition analogously to transportation polytopes:

Definition 2.6
Let $T = T\left((A_i)_{i\in[k]}, (b^{(i)})_{i\in[k]}\right)$ be a clustering polytope as in Definition 2.1 and ξ be a vertex of T.

Then ξ is called *non-degenerate* if and only if $\sum_{i=1}^{k} \mathrm{supp}(C_i) = n + \sum_{i=1}^{k} m_i - \dim\left(\bigcap_{i=1}^{k} \mathrm{lin}(A_i^\intercal)\right)$.

Otherwise, ξ is called *degenerate*.

Using Lemma 2.4 we may now bound the number of units that are non-integrally assigned by a constrained clustering.

Theorem 2.7
Consider the setting of Definition 2.1, i. e., let $A_i \in \mathbb{R}^{m_i \times n}$ *and* $b^{(i)} \in \mathbb{R}^{m_i}$ *for* $i \in [k]$ *be given, and consider* $T := T\left((A_i)_{i\in[k]}, (b^{(i)})_{i\in[k]}\right)$.

Furthermore, let $\operatorname{rank}(A_i) = m_i$ *hold for all* $i \in [k]$.

Then every vertex of T *yields at most*

$$\sum_{i=1}^{k} m_i - \dim\left(\bigcap_{i=1}^{k} \operatorname{lin}(A_i^\intercal)\right)$$

non-integral assignments.

Proof. With $\overline{A}$ as defined by Eq. (2.5), Lemma 2.4 yields

$$\operatorname{rank}(\overline{A}) = n + \sum_{i=1}^{k} m_i - \dim\left(\bigcap_{i=1}^{k} \operatorname{lin}(A_i^\intercal)\right).$$

Every vertex of T has (at most) $\operatorname{rank}(\overline{A})$-many non-zero components. As for every $j \in [n]$ the column $\xi_{\cdot,j}$ requires at least one positive entry due to Definition 2.1, at most $\sum_{i=1}^{k} m_i - \dim\left(\bigcap_{i=1}^{k} \operatorname{lin}(A_i^\intercal)\right)$ of that columns may contain more than one non-zero entry, which yields the claim. □

Note that assuming A_i to have full rank for every $i \in [k]$ is, of course, nothing else than demanding to drop (obviously) redundant constraints. Otherwise, we would have to replace m_i by $\operatorname{rank}(A_i)$ in the statement of Theorem 2.7.

In the homogeneously-constrained case the claim of Theorem 2.7 simplifies to the following corollary. In particular, this implies the well-known case for $m = 1$.

Corollary 2.8
Consider $T = T\left(A, (b^{(i)})_{i\in[k]}\right)$ *as defined in Definition 2.1 in the homogeneously-constrained case and assume* $\operatorname{rank}(A) = m$.

Then every vertex of T *yields at most*

$$m(k-1)$$

non-integral assignments.

Proof. This follows immediately from Theorem 2.7 as here $\dim\left(\bigcap_{i=1}^{k} \operatorname{lin}(A_i^\intercal)\right) = \dim\left(\operatorname{lin}(A^T)\right) = m$. □

Besides in the number of fractionally assigned units, we are furthermore interested in how those fractionally assigned units may be distributed among the clusters. The next result limits the number of units that are shared by two clusters in case that all constraints are of the same type.

Lemma 2.9
Consider $T = T\left(A, (b^{(i)})_{i\in[k]}\right)$ *as defined in Definition 2.1 in the homogeneously-constrained case. Furthermore, assume* $\operatorname{rank}(A) = m$.

Let $\xi = (C_1, \ldots, C_k)^\intercal$ *be a vertex of* T. *Then for every* $i, l \in [k]$, $i \neq l$, *we have*

$$|\operatorname{supp}(C_i) \cap \operatorname{supp}(C_l)| \leq m.$$

Proof. Let $S := \operatorname{supp}(C_i) \cap \operatorname{supp}(C_l)$ and assume $|S| > m$.

Then the $|S|$ columns of A indexed by S are obviously linearly dependent, i. e., there exists $\zeta \in \mathbb{R}^n \setminus \{0\}$ such that $\zeta_{[n]\setminus S} = 0$ and $A\zeta = 0$. Furthermore, we may assume $\|\zeta\|$ to be sufficiently small such that $0 \leq C_i \pm \zeta \leq 1$ as well as $0 \leq C_l \pm \zeta \leq 1$.

We define $\xi^{\pm} := \xi \pm \left(u^{(i)}\zeta^\intercal - u^{(l)}\zeta^\intercal\right)$ (with $u^{(r)} \in \mathbb{R}^k$ denoting the rth standard unit vector for $r \in [k]$).

Then it holds that

$$(\xi^{\pm})^\intercal \mathbb{1} = \xi^\intercal \mathbb{1} \pm \left(u^{(i)}\zeta^\intercal - u^{(l)}\zeta^\intercal\right)^\intercal \mathbb{1} = \mathbb{1} \pm \zeta\left((u^{(i)})^\intercal \mathbb{1} - (u^{(l)})^\intercal \mathbb{1}\right) = \mathbb{1}.$$

Altogether, this gives $\xi^{\pm} \in T$, a contradiction to ξ being extremal. □

2.3 Clustering Graphs

A natural and helpful tool for the analysis of (fractional) clusterings and relations between different clusters are graph representations.

We consider the complete bipartite graph with clusters and units as nodes. A fractional clustering $\xi \in T$ induces a subgraph by taking only the edges between a cluster node i and a unit node j whenever j lies in $\operatorname{supp}(C_i)$. Here, we will identify the node that is associated with a cluster C_i with the cluster itself. However, as a cluster C_i has been identified with the ith row of the clustering ξ, from a strictly formal perspective this causes issues in case of identical rows in ξ. In order to avoid the introduction of yet another notation and as this will not actually lead to ambiguities, we will ignore this issue.

Definition 2.10 (Clustering Graph)
Let $\xi = (C_1, \ldots, C_k)^\intercal \in \mathbb{R}^{k\times n}$ and set $\mathcal{C} := \{C_1, \ldots, C_k\}$. Then we call

$$G(\xi) := (\mathcal{C} \cup [n], E)$$

with $E := \{\{C_i, j\} : i \in [k], j \in [n] \ \wedge \ \xi_{i,j} \neq 0\}$ the *clustering graph* (of ξ).

In our context, we are particularly interested in a more careful analysis of clusters that share a fractionally assigned unit and thus are linked to each other in a certain sense. For this reason, we define the following multigraph (using the notation as, for example, introduced in [Die17, Chapter 1]) with nodes corresponding to clusters and edges for all shared units:

Definition 2.11 (Contracted Clustering Graph)
Let $\xi = (C_1, \dots, C_k)^\intercal \in \mathbb{R}^{k \times n}$ and $G(\xi) := (\mathcal{C} \cup [n], E)$ be the corresponding clustering graph. Then we call

$$G_{\mathcal{C}}(\xi) := (\mathcal{C}, E_{\mathcal{C}}, \nu)$$

with

$$E_{\mathcal{C}} := \left\{ (\{C_i, C_l\}, j) \in \binom{\mathcal{C}}{2} \times [n] : i \neq l \ \wedge \ \{C_i, j\} \in E \ \wedge \ \{C_l, j\} \in E \right\} \text{ and}$$

$$\nu : E_{\mathcal{C}} \to \binom{\mathcal{C}}{2}, \quad (\{C_i, C_l\}, j) \mapsto \{C_i, C_l\}$$

the *contracted clustering graph* (of ξ).

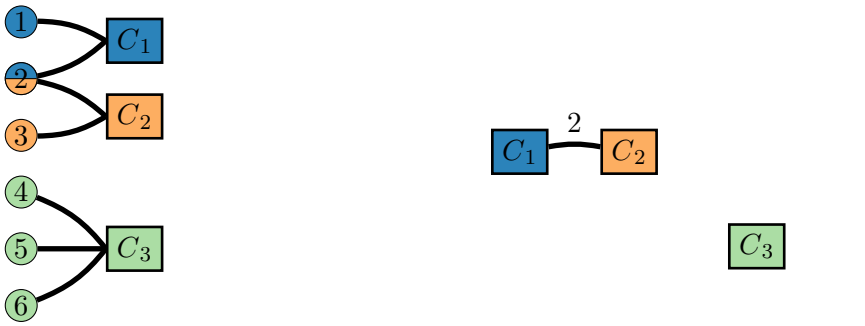

(a) Clustering graph $G(\xi)$. **(b)** Contracted clustering graph $G_{\mathcal{C}}(\xi)$.

Figure 2.3: Clustering graphs of Example 2.2.

Example 2.2 (continuing from p. 11).
For our running example and the clustering ξ as introduced before, we obtain

$$G_{\mathcal{C}}(\xi) = (\{C_1, C_2, C_3\}, \{(\{C_1, C_2\}, 2)\}, \nu)$$

with $\nu((\{C_1, C_2\}, 2)) = \{C_1, C_2\}$. Both the clustering graph as well as the contracted clustering graph are depicted in Fig. 2.3.

Note that we may not only consider clustering graphs for clusterings from a clustering polytope T, but allow arbitrary matrices $\xi \in \mathbb{R}^{k\times n}$. In particular, we will consider the clustering graphs that result from differences $\xi^{(2)} - \xi^{(1)}$ of clusterings $\xi^{(1)}$, $\xi^{(2)} \in T$. Those are then used to describe and understand the transition from one clustering $\xi^{(1)}$ to another clustering $\xi^{(2)}$.

For transportation polytopes it is known that for every extremal point the corresponding subgraph in the bipartite graph that connects sources and sinks is cycle-free (cf. [KW68, Theorem 4]). Also, there is a one-to-one correspondence between trees and bases of the transportation polytope ([YKK84, Chapter 8, Theorem 2.2]). In our setting, this translates into the following theorem:

Theorem 2.12
Consider $T = T\left(A, (b^{(i)})_{i\in[k]}\right)$ *as defined in Definition 2.1 in the homogeneously-constrained case for* $m = 1$ *and assume that* $A \in (\mathbb{R} \setminus \{0\})^{1\times n}$.

Then $\xi \in T$ *is a vertex of* T *if and only if* $G(\xi)$ *is acyclic.*

Moreover, $\xi \in T$ *is a non-degenerate vertex of* T *if and only if* $G(\xi)$ *is a tree.*

Before the proof of Theorem 2.12, let us consider our more general setting — at least for the homogeneously constrained case.

Here, being acyclic translates into a more general kind of connectivity property. A (multi)graph is said to be *l-edge-connected* if it stays connected under the removal of any l edges (cf. [Die17, Chapter 1]). The following result now yields that extremal points of the homogeneously-constrained clustering polytope T requires a certain connectivity of the corresponding contracted clustering graph.

Theorem 2.13
Consider $T = T\left(A, (b^{(i)})_{i\in[k]}\right)$ *as defined in Definition 2.1 in the homogeneously-constrained case.*

Assume $\operatorname{rank}(A) = m$ *and let* ξ *be a non-degenerate vertex of* T.

Then $G_{\mathcal{C}}(\xi)$ *is* m*-edge-connected.*

Proof. Assume $G_{\mathcal{C}}(\xi)$ is not m-edge-connected. Then there exists a cut $(S, \mathcal{C} \setminus S)$ of size at most $m-1$, i. e. $|\{e \in E : e \cap S \neq \emptyset \wedge e \cap (\mathcal{C} \setminus S) \neq \emptyset\}| \leq m-1$. W.l.o.g. assume $S = \{C_1, \ldots, C_l\}$ for some $1 \leq l < k$.

From the proof of Lemma 2.4 we know that the matrix $\overline{A}$, that yields the left-hand side of the equality system defining T, has rank $n + (k-1)m$. More precisely, we notice

that the last m rows are redundant so that the matrix

$$\hat{A} := \begin{pmatrix} \mathrm{Id}_n & \mathrm{Id}_n & \dots & \mathrm{Id}_n & \mathrm{Id}_n \\ A & & & & \\ & A & & & \\ & & \ddots & & \\ & & & A & 0 \end{pmatrix} \in \mathbb{R}^{(n+(k-1)m)\times k\cdot n}$$

yields the left-hand side of an irredundant representation of T.

Let $B \subset [k] \times [n]$ be a basis of ξ. This means $|B| = n + (k-1)m$, $\det(\hat{A}_{\cdot,B}) \neq 0$, and we can express ξ via $\xi_B = (\hat{A}_{\cdot,B})^{-1} \left((\mathbb{1}^{(n)})^\intercal, (b^{(1)})^\intercal, \dots, (b^{(k-1)})^\intercal \right)^\intercal$ and $\xi_{([k]\times[n])\setminus B} = 0$. Furthermore, as ξ is non-degenerate, it must moreover hold that $\xi_{i,j} > 0$ if and only if $(i,j) \in B$.

Now the cut on hand naturally partitions $B = B_1 \dot\cup B_2$ via $B_1 := B \cap [l] \times [n]$ and $B_2 := B \setminus B_1 = B \cap ([k] \setminus [l]) \times [n]$.

Furthermore, we may partition $[n] = N_1 \dot\cup N_2 \dot\cup N_3$ via

$$N_1 := \{j \in [n] : \exists i \in [l] : (i,j) \in B_1 \wedge \nexists i \in [k] \setminus [l] : (i,j) \in B_2\},$$
$$N_2 := \{j \in [n] : \exists i \in [l] : (i,j) \in B_1 \wedge \exists i \in [k] \setminus [l] : (i,j) \in B_2\},$$

and $N_3 := [n] \setminus (N_1 \cup N_2)$. In other words, N_1 yields the indices of units that are assigned only to clusters in S. Furthermore, N_2 yields the indices of units each assigned to at least one cluster in both S and $\mathcal{C} \setminus S$. Again, we may w.l.o.g. assume that $N_1 = [n_1]$ and $N_2 = [n_2] \setminus [n_1]$ (and thus $N_3 = [n] \setminus [n_2]$) for some $0 \leq n_1 < n_2 \leq n$. (Note that for the special case $n_1 = 0$ we assume $N_1 = \emptyset$).

The indices in N_2 are exactly those corresponding to edges in the $(S, \mathcal{C} \setminus S)$-cut of $G_{\mathcal{C}}(\xi)$. Thus, we have

$$|N_2| = n_2 - n_1 \leq m - 1.$$

By definition of B_1, N_1, and N_2, we have that $B_1 \subset [l] \times [n_2]$. Thus, we conclude that $\hat{A}_{\cdot,B_1}$ is a submatrix of

$$\begin{array}{c} \\ 1 \\ 2 \\ \vdots \\ l \end{array} \begin{array}{c} \begin{array}{cccc} 1 & 2 & \dots & l \end{array} \\ \begin{pmatrix} \mathrm{Id}_{n_2} & \mathrm{Id}_{n_2} & \dots & \mathrm{Id}_{n_2} \\ 0 & 0 & \dots & 0 \\ A_{\cdot,[n_2]} & & & \\ & A_{\cdot,[n_2]} & & \\ & & \ddots & \\ & & & A_{\cdot,[n_2]} \\ 0 & 0 & \dots & 0 \end{pmatrix} \end{array}$$

and thus $\operatorname{rank}(\hat{A}_{\cdot,B_1}) \leq n_2 + (l-1)\cdot m$ (by the same argumentation as for $\overline{A}$ in the proof of Lemma 2.4; this may be interpreted as considering the clustering instance reduced to units in $[n_2]$ and clusters in $[l]$).

Similarly, $\hat{A}_{\cdot,B_2}$ is a submatrix of

$$
\begin{array}{c}
\\
l{+}1 \\
l{+}2 \\
\vdots \\
k{-}1
\end{array}
\begin{array}{c}
\begin{array}{ccccc} l{+}1 & l{+}2 & \dots & k{-}1 & k \end{array} \\
\begin{pmatrix}
0 & 0 & \dots & 0 & 0 \\
\mathrm{Id}_{n-n_1} & \mathrm{Id}_{n-n_1} & \dots & \mathrm{Id}_{n-n_1} & \mathrm{Id}_{n-n_1} \\
0 & 0 & \dots & 0 & 0 \\
A_{\cdot,[n]\setminus[n_1]} & & & & \\
& A_{\cdot,[n]\setminus[n_1]} & & & \\
& & \ddots & & \\
& & & A_{\cdot,[n]\setminus[n_1]} & 0
\end{pmatrix}
\end{array}
$$

and so $\operatorname{rank}(\hat{A}_{\cdot,B_2}) \leq n - n_1 + (k-1-l)\cdot m$.

Together, this gives

$$
\begin{aligned}
\operatorname{rank}(\hat{A}_{\cdot,B}) &= \operatorname{rank}(\hat{A}_{\cdot,B_1}) + \operatorname{rank}(\hat{A}_{\cdot,B_2}) \\
&\leq n_2 + (l-1)\cdot m + n - n_1 + (k-1-l)\cdot m \\
&= n + (k-2)\cdot m + n_2 - n_1 \leq n + (k-2)m + m - 1 \\
&= n + (k-1)\cdot m - 1,
\end{aligned}
$$

which is a contradiction to $\det(\hat{A}_{\cdot,B}) \neq 0$. □

Theorem 2.13 yields a minimal connectivity of clusterings that are extremal w. r. t. T. Of course, we also would like to approach the reverse implication.

Theorem 2.14
Consider $T = T\left(A, (b^{(i)})_{i\in[k]}\right)$ as defined in Definition 2.1 in the homogeneously-constrained case.

Assume $\operatorname{rank}(A_{\cdot,S}) = m$ for every $S \in \binom{[n]}{m}$ and let $\xi \in T$.

If $G_{\mathcal{C}}(\xi)$ does not contain a $(m+1)$-edge-connected subgraph, then ξ is a vertex.

Proof. Assume ξ is not a vertex of T. Then there exists $\zeta \in \mathbb{R}^{k\times n} \setminus \{0\}$ such that $\xi \pm \zeta \in T$. The latter is equivalent to $\xi \pm \zeta \geq 0$, $\zeta^\intercal \mathbb{1}^{(k)} = 0$ and $A\zeta^\intercal = 0$.

In particular, we can observe that $G_{\mathcal{C}}(\zeta) \subset G_{\mathcal{C}}(\xi)$, as for all $i \in [k], j \in [n]$ we must have that $0 \leq \xi_{i,j} \pm \zeta_{i,j} \leq 1$ and thus $\zeta_{i,j} \neq 0$ implies $\xi_{i,j} > 0$.

By constraint (2.0a) in the definition of T, it follows that $\zeta^\intercal \mathbb{1}^{(k)} = 0$ and so for every $i \in [k], j \in [n]$ with $\zeta_{i,j} \neq 0$ there must exist a $l \in [k] \setminus \{i\}$ with $\zeta_{l,j} \neq 0$. Hence, together with $\zeta \neq 0$ this means $G_{\mathcal{C}}(\zeta)$ is not empty.

We may further assume ζ to be chosen such that there is no subset of rows $R \subsetneq [k]$ such that for $\tilde{\zeta} := \mathrm{diag}(\mathbb{1}_R^{(k)})\zeta$ — i. e., the matrix in $\mathbb{R}^{k\times n}$ that equals ζ in the rows with index in R and is 0 otherwise — it holds that $\tilde{\zeta} \neq 0$ and $\tilde{\zeta}^\intercal \mathbb{1}^{(k)} = 0$. Otherwise, as then $\xi \pm \tilde{\zeta} \in T$ holds as obviously $\xi \pm \tilde{\zeta} \geq 0$ as well as $A\tilde{\zeta}^\intercal = 0$, we could proceed with $\tilde{\zeta}$ instead of ζ.

From our assumption it follows that we can find a partition $S_1 \dot\cup S_2 = \mathcal{C}$ such that S_1 and S_2 yield a cut of $G_\mathcal{C}(\zeta)$ of size at most m but greater than 0. (For doing so, consider a non-trivial connected component of $G_\mathcal{C}(\xi)$ and use that it is not $(m+1)$-edge-connected.)

Let $U := \{j \in [n] : \exists C_i \in S_1, C_l \in S_2 : \zeta_{i,j} \neq 0 \wedge \zeta_{l,j} \neq 0\}$, i. e., the indices of those units that correspond to edges in this cut. Thus, we get $1 \leq |U| \leq m$. By definition of the cut, we may further obtain a partition of the remaining edges and thus unit indices such that

$$\begin{aligned}
&[n] := M_1 \dot\cup U \dot\cup M_2,\\
&j \in M_1 \Rightarrow \zeta_{i,j} = 0 \ \forall C_i \in S_2, \text{ and}\\
&j \in M_2 \Rightarrow \zeta_{i,j} = 0 \ \forall C_i \in S_1
\end{aligned}$$

hold. In other words, we may assume the following block-structure of ζ:

$$\zeta \cong \begin{array}{c} \\ S_1 \\ S_2 \end{array} \begin{array}{c} \begin{array}{ccc} M_1 & U & M_2 \end{array} \\ \begin{pmatrix} * & * & 0 \\ 0 & * & * \end{pmatrix} \end{array}$$

From $\zeta^\intercal \mathbb{1}^{(k)} = 0$ we conclude

$$\sum_{C_i \in S_1} (\zeta_{i,\cdot})^\intercal = \begin{array}{c} M_1 \\ U \\ M_2 \end{array} \begin{pmatrix} 0 \\ u \\ 0 \end{pmatrix} = -\sum_{C_i \in S_2} (\zeta_{i,\cdot})^\intercal$$

for some $u \in \mathbb{R}^{|U|}$. By our assumption on the choice of ζ, we can further conclude $u \neq 0$. From $A\zeta^\intercal = 0$ it follows that

$$0 = \sum_{C_i \in S_1} (A\zeta^\intercal)_{\cdot,i} = A\Big(\sum_{C_i \in S_1} (\zeta_{i,\cdot})^\intercal\Big) = (A_{\cdot,U})u.$$

However, this contradicts the at most m columns of $A_{\cdot,U}$ that are linearly independent.□

Note that Lemma 2.9 is actually also a corollary of Theorem 2.14.

The restriction in Theorem 2.14 to constraint matrices A that only have linearly independent columns may be understood as the units that are represented by the

columns in A being in a *general position*. In particular, equal columns may be understood as equivalent units that can be arbitrarily interchanged.

We can observe by the following example that, in general, the reverse of Theorem 2.14 is not true.

Example 2.15
For $n := 6$ we consider points $x^{(j)} \in \mathbb{R}$ defined by $x^{(j)} := j$ for $j \in [6]$. Furthermore, we give each point $x^{(j)}$ the weight 1. We would like to (fractionally) cluster those points into $k := 4$ clusters such that each cluster has a total weight $3/2$ and such that clusters 1, 2, and 3 have their centroid in 3, i. e., $\frac{2}{3}\sum_{j=1}^{6} \xi_{i,j} x_j = 3$ for $i \in [3]$, and the 4th cluster has the centroid 5.

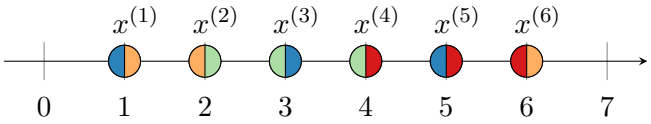

Figure 2.4: Illustration of Example 2.15.

Thus, we are in a homogeneously-constrained case with $m = 2$ and set $A := \begin{pmatrix} 1 & 1 & 1 & 1 & 1 & 1 \\ 1 & 2 & 3 & 4 & 5 & 6 \end{pmatrix}$, $b^{(i)} := \begin{pmatrix} 3/2 \\ 9/2 \end{pmatrix}$ for $i \in [3]$, and $b^{(4)} := \begin{pmatrix} 3/2 \\ 15/2 \end{pmatrix}$.
Figure 2.4 illustrates this situation and the fractional clustering

$$\xi^* = \begin{pmatrix} 1/2 & 1/2 & 0 & 0 & 0 & 1/2 \\ 1/2 & 0 & 1/2 & 0 & 1/2 & 0 \\ 0 & 1/2 & 1/2 & 1/2 & 0 & 0 \\ 0 & 0 & 0 & 1/2 & 1/2 & 1/2 \end{pmatrix}.$$

Indeed, $\xi^* \in T\left(A, (b^{(i)})_{i\in[4]}\right)$ holds. Moreover, from Lemma 2.4 we can conclude that $\operatorname{rank}(\overline{A}) = 6+4\cdot 2-2 = 12$ with $\overline{A}$ as defined by Eq. (2.5). Hence, a basis of any vertex of T has cardinality $4 \cdot 6 - 12 = 12$. We observe that ξ^* is indeed a non-degenerate vertex of T with according basis $B = \{(1,1),(1,3),(1,5),(2,1),(2,2),(2,6),(3,2),(3,3),(3,4),(4,4),(4,5),(4,6)\}$. However, the resulting clustering graph is complete (as simple graph), i. e., $G_{\mathcal{C}}(\xi^*) \cong K_4$, and thus 3-edge-connected (cf. Fig. 2.5).

Still, for the case of homogeneously single-constrained clusterings, Theorem 2.12 states that the reverse of Theorem 2.14 indeed holds.

Here, the resulting transportation polytopes are very well-understood. In particular, let us recall that edges of those polytopes can be identified with cycles in the clustering graph $G(\xi)$. In the setting of Theorem 2.12, consider a clustering $\xi \in T$ and let

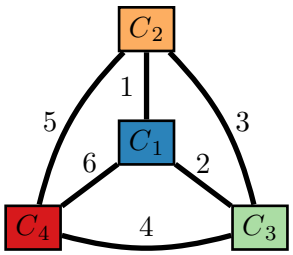

Figure 2.5: Contracted clustering graph of Example 2.15.

$C_{i_1}, j_1, \ldots, j_{r-1}, C_r, j_r, C_{i_{r+1}} = C_{i_1}$ be a cycle in $G(\xi)$. Then, we define the *cyclic exchange* $\zeta \in \mathbb{R}^{k \times n}$ via

$$\zeta_{i_l,j_l} := -(A_{1,j_l})^{-1} \quad \text{and} \quad \zeta_{i_{l+1},j_l} := (A_{1,j_l})^{-1} \tag{2.6}$$

for $l = 1, \ldots, r$ and 0 in all remaining components. Note that by construction it holds that $\zeta^\intercal \mathbb{1}^{(k)} = 0^{(n)}$ as well as $\zeta A^\intercal = 0^{(k)}$. Vice versa, for any $\tilde{\zeta} \in \mathbb{R}^{k \times n} \setminus \{0\}$ such that $G\left(\tilde{\zeta}\right)$ is the same cycle and $\zeta^\intercal \mathbb{1}^{(k)} = 0^{(n)}$ as well as $\zeta A^\intercal = 0^{(k)}$ hold, we can deduce that $\tilde{\zeta}$ equals the corresponding cyclic exchange defined by Eq. (2.6) up to scaling.

Let us now briefly supplement the proof of Theorem 2.12.

Proof (Theorem 2.12). From Theorem 2.14 we get that if $G(\xi)$ is acyclic and therefore is $G_{\mathcal{C}}(\xi)$, then ξ is extremal.

For the reverse implication, assume that there exists a cycle in $G(\xi)$ and let ζ be the corresponding cyclic exchange. Then there exists $\epsilon > 0$ such that $\xi \pm \epsilon\zeta \geq 0$ and hence by construction of ζ it follows that $\xi \pm \epsilon\zeta \in T$, so ξ is not extremal.

The second claim then is a direct consequence of Theorem 2.13. □

With Theorem 2.12, we may now rephrase the following well-known characterization of edges of transportation polytopes in our context.

Lemma 2.16

In the setting of Theorem 2.12, let $\xi^{(1)} \neq \xi^{(2)} \in T$ be two vertices.

Then $\xi^{(1)}, \xi^{(2)}$ are neighbors if and only if $G\left(\xi^{(1)}\right) \cup G\left(\xi^{(2)}\right)$ contains exactly one cycle, namely $G\left(\xi^{(2)} - \xi^{(1)}\right)$.

Proof. Let us first make an observation. Consider any $\zeta \in \mathbb{R}^{k \times n}$ with $\zeta^\intercal \mathbb{1}^{(k)} = 0^{(n)}$ and $\zeta A^\intercal = 0^{(k)}$. Those two equalities imply $deg_{()}G(\zeta)(j) \neq 1$ for every $j \in [n]$ and $deg_{()}G(\zeta)(C_i) \neq 1$ for every $i \in [k]$, respectively. Consequently, $G(\zeta)$ is a union of (maybe not edge-disjoint) cycles. In particular, $G\left(\xi^{(2)} - \xi^{(1)}\right)$ is a union of cycles.

Set $\tilde{\xi} := \frac{1}{2}(\xi^{(1)} + \xi^{(2)})$. Observe that by construction

$$G\left(\xi^{(2)} - \xi^{(1)}\right) \subset G\left(\tilde{\xi}\right) = G\left(\xi^{(1)}\right) \cup G\left(\xi^{(2)}\right).$$

By Theorem 2.12, $G\left(\tilde{\xi}\right)$ contains a cycle. Let ζ be the corresponding cyclic exchange.

Now first assume that $G\left(\tilde{\xi}\right)$ contains another cycle in addition to $G(\zeta)$. Let $\tilde{\zeta}$ be the corresponding cyclic exchange. Then there exists $\epsilon > 0$ such that $\tilde{\xi} \pm \epsilon\tilde{\zeta} \in T$. As ζ and $\tilde{\zeta}$ are linearly independent (due to their different support), this implies that the dimension of the face of $\tilde{\xi}$ is at least 2, so $\xi^{(1)}$ and $\xi^{(2)}$ are not neighbors. As $G\left(\xi^{(2)} - \xi^{(1)}\right)$ is a union of cycles and not empty, it follows that $G\left(\xi^{(2)} - \xi^{(1)}\right) = G(\zeta)$.

Next, assume that $G\left(\xi^{(1)}\right) \cup G\left(\xi^{(2)}\right)$ only contains the cycle $G(\zeta)$. As $G\left(\xi^{(2)} - \xi^{(1)}\right)$ is a union of cycles, it follows that $G(\zeta) = G\left(\xi^{(2)} - \xi^{(1)}\right)$. Now consider any $\tilde{\zeta} \neq 0$ such that $\tilde{\xi} \pm \tilde{\zeta} \in T$. It follows that $G\left(\tilde{\zeta}\right)$ is a union of cycles contained in $G\left(\tilde{\xi}\right)$ and hence $G\left(\tilde{\zeta}\right) = G\left(\xi^{(2)} - \xi^{(1)}\right)$. This implies $\tilde{\zeta} = \lambda\left(\xi^{(2)} - \xi^{(1)}\right)$ for some $\lambda \in \mathbb{R} \setminus \{0\}$. Thus the dimension of the face of $\tilde{\xi}$ is 1 and $\xi^{(1)}$ and $\xi^{(2)}$ are neighbors . □

A useful corollary is that cyclic exchanges suffice to traverse between any two points in a homogeneously single-constrained clustering polytope:

Corollary 2.17
In the setting of Theorem 2.12, let $\xi^{(1)}, \xi^{(2)} \in T$.

Then there exist $r \in \mathbb{N}$, cyclic exchanges $\zeta^{(1)}, \ldots, \zeta^{(r)} \in \mathbb{R}^{k \times n}$ (as defined by Eq. (2.6)), and $\lambda_i \in \mathbb{R}$ for $i \in [r]$ such that

$$\xi^{(1)} - \xi^{(2)} = \sum_{i \in [r]} \lambda_i \zeta^{(i)}$$

Moreover, the cyclic exchanges can be chosen such that $G\left(\zeta^{(i)}\right) \subset G\left(\xi^{(2)} - \xi^{(1)}\right)$ for every $i \in [r]$.

Proof. The first statement follows directly from Lemma 2.16 which yields that every edge direction of T is collinear to a cyclic exchange. For the second statement, note that it is sufficient to consider edge directions of the face of $\frac{1}{2}(\xi^{(1)} + \xi^{(2)})$. As all extremal clusterings of that face assign units to clusters equally if $\xi^{(1)}$ and $\xi^{(2)}$ do, it follows that any of those cyclic exchanges can only use edges in $G\left(\xi^{(2)} - \xi^{(1)}\right)$. □

2.4 Summary & Conclusion

This chapter mainly defined the notion of constrained (fractional) clusterings. In particular, we allow constraints that determine the cluster aggregates of numerical values associated with the points to be clustered. Here, a typical example is a weighted point set and constraints that determine the resulting cluster weights.

We then discussed some basic properties of the underlying clustering polytopes in Section 2.2. Most important for our scope, we can bound the number of units that are assigned to multiple clusters by a vertex, which generalizes a well-known result for transportation polytopes. This is of major interest for applications in which the possibility of fractional clusterings is a relaxation of binary constraints. Here, the achieved bound can be exploited to bound the resulting relaxation error. Of course, many issues regarding clustering polytopes remain open. For example, an improvement of Proposition 2.3 that yields sufficient conditions for the general feasibility of a clustering polytope might be of interest.

Finally, Section 2.3 introduced and discussed two types of clustering graphs that will yield a useful tool in the following chapters when dealing with the relations between clusters that share commonly assigned units.

Chapter 3

Generalized Voronoi Diagrams

3.1 Motivation

Voronoi diagrams are a simple and yet powerful tool that has been applied successfully in various fields — both of theoretical and practical nature. The supposedly earliest known literature, in which they appear explicitly, is the work of René Descartes in the 17th century. He used a Voronoi diagram-like structure to decompose the solar system. Their name roots back to Georgy Feodoseevich Voronoy, who in 1908 gave a formal treatment of their concept (apparently, this is only one of the many English translations of his Russian name that can be found in the literature). Even before Voronoy, in the middle of the 19th century, Carl Friedrich Gauß and Gustav Lejeune Dirichlet formalized this kind of diagram when studying quadratic forms (we refer to [Oka+00] and [AKL13] for the claimed dates and the original references).

Since then, Voronoi diagrams and strongly related concepts have been appearing in various domains of science, such as biology, chemistry, computer science, crystallography, geography, physics — to list just a few of the more than twenty fields mentioned in the introduction of [Oka+00]. As a particular milestone in computational geometry, Shamos and Hoey [SH75] presented an efficient construction algorithm for Voronoi diagrams. They used this in order to design superior algorithms for classical geometric problems such as the Euclidean minimum spanning tree or finding a smallest enclosing circle. This might be considered a starting signal for the widely spread usage of Voronoi diagrams in the field of computational geometry and computer science ([Oka+00, Chapter 1.2]).

One drawback of this scattered field of applications is that many concepts have been developed independently. This results in different denotations for identical objects such as cells or regions, sites or generators. Voronoi diagrams themselves have been known under many different names, such as Dirichlet Tessellations, Wirkungsbereiche [domain of actions], Thiessen polygons or Wigner-Seitz regions, to list just a few (cf. [Now33], [Aur91]). Also, identical terms are used with different notions. For example, the term Voronoi diagram itself might — among others — refer to a (polyhedral) cell complex of its faces ([Aur87a]), the union of its edges ([AKL13]), or to the collection of its (closed) cells (as we are going to do in Section 3.2).

Aurenhammer, Klein, and Lee [AKL13] provide a comprehensive treatment of Voronoi diagrams and their generalizations. They present a thorough study of the structural properties of the various kinds of Voronoi diagrams, computational concepts for their construction, and geometric applications. Another important book of reference is provided by Okabe et al. [Oka+00]. They as well give a comprehensive treatment of the several generalizations of Voronoi diagrams from an application-driven perspective. Furthermore, they provide some historical background. For further, less extensive surveys on Voronoi diagrams we refer to [Aur91], [For95], [AK00], and [For04].

There are several reasons to consider Voronoi diagrams and their generalizations when clustering data. Here, let us discuss two main perspectives in the following.

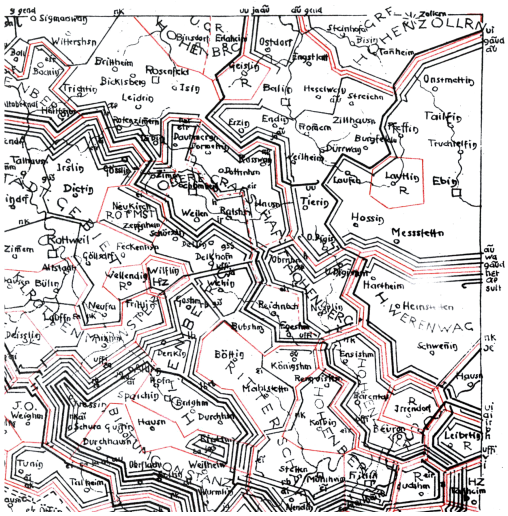

Figure 3.1: Extract of a map by Karl Haag from 1928 illustrating linguistic boundaries in a south German region, taken from [Haa29, p. 3]. Regions result from unions of cells from the Voronoi diagram with municipal sites. Red lines depict political borders (with multiplicity indicating political changes over time). Black lines indicate phonetic changes in the local dialect, with the thickness of lines indicating the intensity of the linguistic differences.

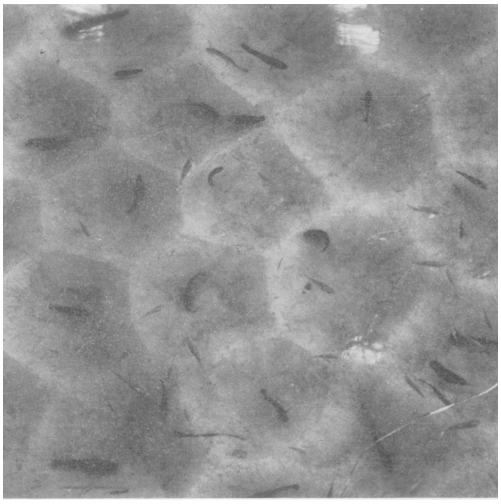

Figure 3.2: Extract of a photograph showing the breeding territories of African mouthbreeding cichlid fish, taken from [Bar74, Plate XIV, Fig. 1]. Males dig pits into sandy ground that mark their territories. In this experimental setting, Barlow [Bar74] succeeded in providing empirical evidence for the theory of regular hexagonal animal territories in homogeneous areas. This, of course, results in the Voronoi diagram resulting from a densest sphere packing in the plane.

From an Application's Point of View

First, their simple structure does not only seem natural, but can indeed be observed in nature itself. Let us provide a few examples (many more may be found in [Oka+00, Chapter 7]).

Voronoi diagrams play an important role in the field of crystallography. Here, denoted by *domains of actions*, Voronoi diagrams w. r. t. sites arranged on lattices have been researched in order to characterize regular crystalline structures (see, for example, [Nig27], [Now76]).

In biology, Voronoi diagrams explain the regular patterns of skin cells of mammals as the result of their regeneration process ([Sai82]). In an experiment researching breeding territories of the African mouthbreeding cichlid fish, Barlow [Bar74] was able to confirm the theory that animal territories in an ideal uniform setting should result in a regular hexagonal tiling (see Fig. 3.2). In the areas of geography and social sciences, the *central place theory* tries to explain frequency, sizes, and locations of urban

settlements in rural areas (cf. the review in [Kin84]). Here, a basic model assumes differences in consumption costs to only result from different transportation costs. This immediately results in Voronoi regions as trade areas of urban centers. Other models try to measure the influence of each urban center by indexes obtained, for example, from their respective number of telephone connections, shops, or population sizes. This results in trade areas that are weighted generalizations of Voronoi diagrams.

Another example stems from the rather surprising field of linguistics. In 1898, the German linguist Karl Haag started to publish maps of linguistic boundaries that depict the changes of dialect in some German regions (cf. [Haa98], [Haa29], and the review in [HW82]). By depicting "Grenzen zwischen Nachbargemeinden […] durch die Mittellote auf der Distanz" [boundaries between neighboring municipalities as the middle perpendiculars of their distance] ([Haa98, p. 7]), he as well discovered the structure of Voronoi diagrams on his own. Furthermore, Haag depicted both the political as well as the phonetic boundaries in this way. He justified this as follows: "Die politischen Gebiete nach dem thatsächlichen Umfang der Gemeindemarkungen […] abzugrenzen hat hier keinen Sinn, da die Zusammengehörigkeit und die Entfernung allein in Betracht kommt" [To demarcate the political regions by their actual municipal extent is not reasonable here, as only affiliation and distance are of interest.] ([Haa98, p. 7]). Figure 3.1 yields an extract of one of Haag's maps. In conclusion, not only is the simple structure of Voronoi diagrams sufficiently rich in order to contain all required information — they may sometimes even be superior, as in this example of transitions of dialect, where their resulting territorial structure is allegedly more accurate than any politically drawn boundaries.

From an Optimization Point of View

A second reason for us to be interested in Voronoi diagrams is the fact that they (and their generalizations) naturally appear as an underlying structure for optimal solutions in various optimization problems.

As already mentioned, Shamos and Hoey [SH75] pointed out that geometric problems such as the nearest neighbor problem, finding an Euclidean minimum spanning tree, or finding the smallest circle enclosing a set of points, can be efficiently solved by computing the Voronoi diagram that reveals their solutions. In the context of clusterings, the quality of clusters is often measured by a function that depends only on the distances of points in a cluster to each other or to some cluster representative (cf. [JMF99], [AKL13, Section 8.4]). In many of those cases, an optimal clustering will be induced by a Voronoi diagram (when using the Euclidean distance) or a generalization. Let us provide some examples (without any claim of completion, of course).

The method of *k-means* clustering as introduced by MacQueen [Mac67] aims to find an (unconstrained) clustering that minimizes the within-cluster sum of squares (or average of cluster variances, also referred to as *squared error* or *moment of inertia*).

More precisely, for a set of points $X := \left\{x^{(1)}, \ldots, x^{(n)}\right\} \subset \mathbb{R}^d$ and a weight function $\omega : X \to \mathbb{R}_{>0}$, one wants to find a partition $X = \dot{\bigcup}_{i\in[k]} X_i$ such that

$$\sum_{i\in[k]} \sum_{x\in X_i} \omega(x) \left\|x - c^{(i)}\right\|_2^2$$

with $c^{(i)} := \sum_{x\in X_i} \frac{\omega(x)}{\omega(X_i)} \cdot x$ and $\omega(X_i) := \sum_{x\in X_i} \omega(x)$ for $i \in [k]$ is minimized. One can easily verify that any optimal partition must result from the Voronoi diagram obtained by the optimal partition's centroids as sites. In fact, the famous k-means algorithm as introduced in [Llo82] (thus also referred to as *Lloyd's method*) iteratively assigns each point to the cluster that corresponds to the cell of the Voronoi diagram given by the last iteration's centroid that contains it. The k-means objective might be the "most intuitive and frequently used criterion function in partitional clustering" ([JMF99, p. 278]). Whenever data can be embedded into the Euclidean space in a meaningful way, the variance of each cluster yields a natural criterion of how well it can be represented by a single point or, more precisely, its mean. Due to its outlier-sensitivity, this criterion is also well-suited in order to measure the consolidation (or informal "compactness") of clusters.

Similarly, in the case of *k-median* clustering one wants to minimize

$$\sum_{i\in[k]} \sum_{x\in X_i} \omega(x) \left\|x - s^{(i)}\right\|_2$$

over the space of all clusterings $X_1, \ldots, X_k$ and sites $s^{(1)}, \ldots, s^{(k)} \in \mathbb{R}^d$. Here, any optimizer will always be induced by the Voronoi diagram defined by the clusters' geometric medians (or Fermat-Weber points, cf. [Bri95]).

Inaba, Katoh, and Imai [IKI94] furthermore show that any clustering minimizing the all-pairs sum of squared errors, i. e., $\sum_{i\in[k]} \sum_{x,y\in X_i} \|x - y\|_2^2$, results from a generalized Voronoi diagram which we will denote as an anisotropic power diagram (see Section 3.3.3).

Capoyleas, Rote, and Woeginger [CRW91] show that for any objective $F(r(X_1), \ldots, r(X_k))$ with $F : \mathbb{R}^k \to \mathbb{R}$ being a monotone increasing function (such as the sum or the maximum, for example) and $r(X_i)$ being the Euclidean radius of the cluster X_i, there exists an optimal clustering that is induced by a power diagram (cf. Section 3.3.2). In case of $F = \max$, this yields the *k-center* problem. In this special case, it can be easily seen that the clustering induced by the Voronoi diagram obtained from the k optimal centers will be optimal.

Beyond that, the geometric properties of Voronoi diagrams and their generalizations themselves can be of particular interest in applications. This is the case in our prime example of electoral district design that we will treat in Chapter 7. Here, consolidation of districts is a major optimization criterion. Furthermore, electoral law typically

requires (some sort of) contiguity of districts. We will see that optimally consolidated districts will be embedded into the cells of certain Voronoi diagram types. To some extent, this will then imply the required contiguity a priori, without formalizing it as an explicit constraint in the model.

3.2 Definition

As already pointed out in the previous section, many different generalizations of Voronoi diagrams have been defined in the literature, sometimes multiple times but independently of each other, sometimes with different names or deviations in the precise notions. In this work, we will use a very broad definition similar to the one by Edelsbrunner and Seidel [ES86]. Their intent was to provide a definition that answers the question "What constitutes a VoD [Voronoi Diagram] in its most general form?" ([ES86, p. 26]). For our purposes, a Voronoi diagram is to yield a dissection of some arbitrary space $\mathcal{X}$ into a predefined finite number of (not necessarily disjoint) cells. For each cell, we assume a *dissimilarity measure* in form of an arbitrary real-valued function to be given. While a classical Voronoi diagram in Euclidean space assigns each point to the closest among the cells' sites, we assign each point to its least dissimilar cell as determined by the cells' individual dissimilarity measures.

Definition 3.1
Let $\mathcal{X}$ be a set, $k \in \mathbb{N}$ and functions $f_i : \mathcal{X} \to \mathbb{R}$ for $i \in [k]$ be given. We set

$$P_i := \{x \in \mathcal{X} : f_i(x) \leq f_l(x) \ \forall l \in [k]\} \tag{3.1}$$

and call it (the ith) *Voronoi cell* and the tuple

$$\mathcal{P} := (P_1, \ldots, P_k)$$

generalized Voronoi diagram (w. r. t. $(f_i)_{i\in[k]}$).

For any $i \neq l \in [k]$, we furthermore call $B_{i,l} := \{x \in \mathcal{X} : f_i(x) = f_l(x)\}$ the *bisector* of the cells i and l, and $H_{i,l} := \{x \in \mathcal{X} : f_i(x) \leq f_l(x)\}$ the *dominance region* of i over l, so that $P_i = \bigcap_{l\in[k]} H_{i,l}$.

Note that this definition differs from the one provided in [ES86]. There, for each $I \subset [k]$ the corresponding Voronoi cell is a set of points for which the functions f_i for *all* $i \in I$ yield the minimum value among all functions f_i for $i \in [k]$. Thus, they obtain an actual partition of the underlying space $\mathcal{X}$. In particular, if all differences of functions $f_i - f_l$ for $i, l \in [k]$ are affine (which is the case for the classical Voronoi diagram), this partition yields a polyhedral cell complex (cf. [Aur87b]). However, for our purposes, the less granular variant of Definition 3.1 suffices.

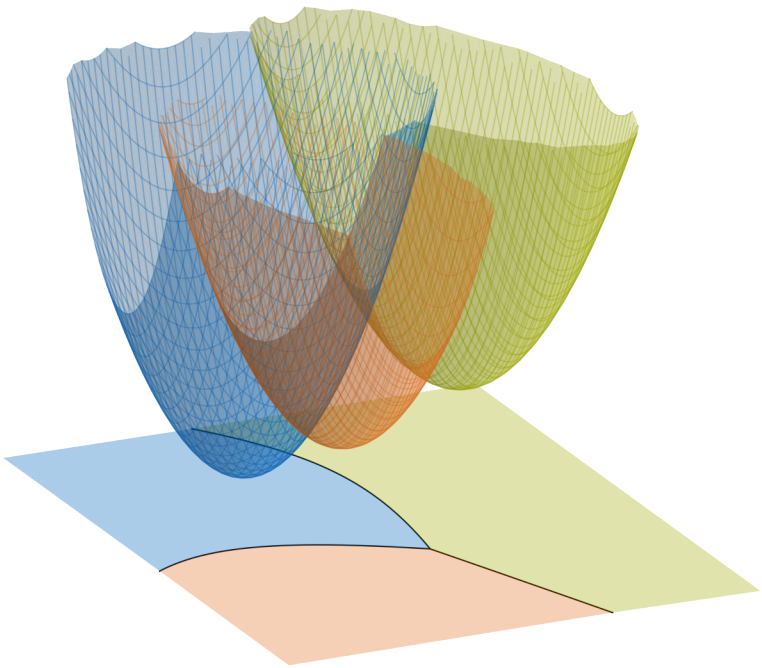

Figure 3.3: Voronoi surface in $\mathbb{R}^3$ of a two-dimensional generalized Voronoi diagram. The meshed surfaces depict the graphs of the Voronoi functions f_i (here: paraboloids, cf. Section 3.3.3).

Also, not surprisingly, even this definition does not cover every type of Voronoi diagram. Another, quite common, generalization of Voronoi diagrams are diagrams of *higher order*. Here, a point is assigned to the p closest (or least dissimilar) cells for some fixed $p \leq k - 1$ (cf. [ES86], [AKL13, Section 6.5], [Oka+00, Section 3.2]). Further, dynamic Voronoi diagrams with the cells' dissimilarity measures changing over time have been considered (cf. [Oka+00, Section 3.9]). Although not regarded any further in this thesis, a natural question that arises is how those could be utilized in our context of constrained clusterings.

A helpful understanding of Definition 3.1 is to interpret a generalized Voronoi diagram as a projection of the *lower envelope* or *minimization diagram* of the graphs of its defining functions, i. e.,

$$\left\{(x, y) \in \mathcal{X} \times \mathbb{R} : y = \min_{i \in [k]} f_i(x)\right\}.$$

Consequently, Huttenlocher, Kedem, and Sharir [HKS93] call this graph the *Voronoi surface* (and describe how the combination of several Voronoi surfaces can be used to compute the Hausdorff distance under translation of two finite sets). Figure 3.3 depicts

such a surface for the case of the functions f_i being quadratic polynomials (which we will investigate further in Section 3.3.3). The complexity of lower envelopes has been studied in several settings, which immediately yields bounds on the complexity of generalized Voronoi diagrams. Here, by complexity we mean the number of all connected components of cells of all dimensions, with a cell defined in accordance with [ES86] as mentioned above. For the case of $\mathcal{X} = \mathbb{R}^d$ and the functions f_i being polynomials of fixed degree, it is shown in [Sha94] that the complexity of the resulting diagram is $\mathcal{O}\left(k^{d+\epsilon}\right)$ for any fixed $\epsilon > 0$ (assuming also the dimension to be fixed; this is based on the according result for $d = 2$ in [HS94]). Moreover, Agarwal, Aronov, and Sharir [AAS97] provide a randomized algorithm with an expected running time in $\mathcal{O}\left(k^{d+\epsilon}\right)$ for any fixed $\epsilon > 0$ that computes all vertices, edges and two-dimensional cells in this setting. They also present an extended version that yields the full incidence structure for $d = 3$. The important special case of all f_i being affine linear, i. e., the case that the graphs of the functions f_i are (non-vertical) hyperplanes, is well-understood and will lead us to *power diagrams* in Section 3.3.2. Lower envelopes mostly appear in the context of *arrangements* that are the decompositions of $\mathbb{R}^d$ into connected cells obtained from a collection of hypersurfaces or surface patches. In particular, arrangements are useful when considering generalized Voronoi diagrams of higher order (cf. Edelsbrunner and Seidel [ES86]). A survey on general arrangements including lower envelopes is provided by Agarwal and Sharir [AS00].

It is clear from Definition 3.1 that the functions that define a generalized Voronoi diagram are anything but unique. For a given diagram $\mathcal{P}$, any functions $(f_i)_{i\in[k]}$ which preserve in each point $x \in \mathcal{X}$ the indices of its minimizing functions, i. e., $\operatorname{argmin}\{f_i(x) : i \in [k]\} = \{i \in [k] : x \in P_i\}$, are valid. In particular, given a diagram w. r. t. functions $(f_i)_{i\in[k]}$, any transformation of those functions that strictly preserves their point-wise order will yield the same diagram. Let us store a simple but useful consequence of this observation in the following lemma.

Lemma 3.2
Let $\mathcal{P}$ be a generalized Voronoi diagram w. r. t. functions $(f_i)_{i\in[k]}$ as defined in Definition 3.1, and let $a : \mathcal{X} \to \mathbb{R}_{>0}$ and $b : \mathcal{X} \to \mathbb{R}$ be two further functions. Then with $\tilde{\mathcal{P}}$ being the generalized Voronoi diagram w. r. t. $(a \cdot f_i + b)_{i\in[k]}$, it holds that $\mathcal{P} = \tilde{\mathcal{P}}$.

As a consequence, we could, for example, without loss of generality demand $f_k = 0$ in Definition 3.1.

However, either way Definition 3.1 does not provide any meaningful diagram as long as the class of functions $(f_i)_{i\in[k]}$ is not suitably restricted. In fact, any collection of sets $P_i \subset \mathcal{X}$ for $i \in [k]$ with $\bigcup_{i\in[k]} P_i = \mathcal{X}$ yields a generalized Voronoi diagram as defined in Definition 3.1. For example, it suffices to set $f_i := -\mathbb{1}_{P_i}$ to be the negative respective indicator function for every $i \in [k]$.

For our applications, we will typically restrict the functions f_i to be of the form

$$f_i(x) := h(d_i(x, s^{(i)})) + \mu_i \tag{3.2}$$

for $x \in \mathcal{X}$. Here, d_i is supposed to be a (possibly asymmetric or pseudo-) metric that may depend on i. For each $i \in [k]$, $s^{(i)} \in \mathcal{X}$ yields a cell's reference point, which is referred to as *site*. Next, $h : \mathbb{R} \to \mathbb{R}$ shall be a common transformation function and $\mu_i \in \mathbb{R}$ some further additive weight for every $i \in [k]$.

In many applications of our interest, geometric characteristics such as the connectivity of cells are of major importance. The following lemma ensures a strongly related property in general metric spaces if the class of allowed diagrams as proposed by Eq. (3.2) is restricted further.

Lemma 3.3
Let d be a (possibly non-symmetric) metric on $\mathcal{X}$. Let $\mathcal{P}$ be a generalized Voronoi diagram w. r. t. functions $(f_i)_{i\in[k]}$ with

$$f_i(x) = d(x, s^{(i)}) + \mu_i \tag{3.3}$$

with $s^{(i)} \in \mathcal{X}, \mu_i \in \mathbb{R}$ for all $x \in \mathcal{X}, i \in [k]$. Then for every fixed $i \in [k]$ and $x \in P_i$ it holds that

$$\left\{y \in \mathcal{X} : d(x, s^{(i)}) = d(y, s^{(i)}) + d(x, y)\right\} \subset P_i. \tag{3.4}$$

Proof. Let $y \in \mathcal{X}$ with $d(x, s^{(i)}) = d(y, s^{(i)}) + d(x, y)$ and $l \in [k]$. Then

$$\begin{aligned} f_i(y) = d(y, s^{(i)}) + \mu_i &= -d(x, y) + d(x, s^{(i)}) + \mu_i \\ &\overset{x\in P_i}{\leq} -d(x, y) + d(x, s^{(l)}) + \mu_l \overset{\Delta\text{-inequ.}}{\leq} d(y, s^{(l)}) + \mu_l = f_l(y), \end{aligned}$$

and hence $y \in P_l$. □

We will refer to property (3.4) as *site-star-shaped*. For example, if $\mathcal{X}$ is a (possibly asymmetric) Minkowski space and d its norm-induced metric, then for any generalized Voronoi diagram w. r. t. functions of the form (3.3), the cells are star-shaped in the sense that for any $x \in P_i$, it follows that $[s^{(i)}, x] \subset P_i$.

A similar result can be found in [AKL13]. The authors consider the connectivity of cells w. r. t. functions $f_i(x) = d(x, s^{(i)})$. Here, d is a so-called *nice* metric in the plane with some favorable properties concerning its relation to the Euclidean topology and the resulting bisector curves (cf. [AKL13, Lemma 7.4]). We may adapt some of their ideas that use a fundamental result by Menger [Men28] to the setting of Eq. (3.3). Menger calls a set $\mathcal{X}$ equipped with a metric d a *convex metric space* if

$$\forall a \neq b \in \mathcal{X}\ \exists c \in \mathcal{X} \setminus \{a, b\} : d(a, b) = d(a, c) + d(c, b). \tag{3.5}$$

In a convex metric space $\mathcal{X}$, we call a set $A \subset \mathcal{X}$ *connected*, if for any two points $a, b \in A$ there exists a continuous curve $\gamma : [0,1] \to A$ with $\gamma(0) = a$ and $\gamma(1) = b$. If furthermore $\mathcal{X}$ is complete (w. r. t. the d-induced convergence of Cauchy sequences), then for all points $a, b \in \mathcal{X}$ there exists a *geodesic path* from a to c, i. e., a continuous curve γ as above such that for all $\lambda_1, \lambda_2 \in [0,1]$ it holds that $d(\gamma(\lambda_1), \gamma(\lambda_2)) = |\lambda_1 - \lambda_2| d(a,c)$ (cf. [Men28, p. 89]). As points along such a path by definition satisfy Eq. (3.5), we can set $a := s^{(i)}$ and $c := x$ in Lemma 3.3 and deduce the following:

Corollary 3.4
Consider the setting of Lemma 3.3. Assume furthermore $\mathcal{X}$ to be a complete metric space that satisfies Eq. (3.5)*. Then for every $i \in [k]$ the Voronoi cell P_i is connected.*

Furthermore, Lemma 3.3 implies that non-empty cells always contain their respective sites:

Corollary 3.5
Consider the same setting as in Lemma 3.3. Then for every $i \in [k]$ it holds that

$$P_i \neq \emptyset \Rightarrow s^{(i)} \in P_i.$$

From an application point of view, the bisectors $B_{i,l}$ of a generalized Voronoi diagram often imply an unpleasant uncertainty, as elements in bisectors cannot be unambiguously mapped to a cell. For generalized Voronoi diagrams in $\mathbb{R}^d$ with defining functions of the form of (3.2), the following lemma guarantees that bisectors do not intersect the interior of cells if the metrics d_i are induced by strictly convex norms. Note that this result and its proof have been published in [BGK17, Lemma 6] and have only been slightly adapted to the notation of this thesis.

Lemma 3.6 ([BGK17, Lemma 6])
Let $\mathcal{X} = \mathbb{R}^d$ and functions f_i for $i \in [k]$ in the form of Eq. (3.2) *be given with the further assumptions that every metric d_i is induced by a strictly convex norm, the function h is continuous and injective, and the sites are pairwise distinct, i. e., $s^{(i)} \neq s^{(l)}$ for $i \neq l \in [k]$.*

Then for the generalized Voronoi diagram $\mathcal{P} = (P_1, \ldots, P_k)$ w. r. t. $(f_i)_{i \in [k]}$ it holds that $\operatorname{int}(P_i) \cap \operatorname{int}(P_l) = \emptyset$ for every $i \neq l \in [k]$.

Proof. Let $U_1, U_2 \subset \mathbb{R}^n$ be the unit balls of the norms that induce d_1 and d_2, respectively. Furthermore, denote by $\|\cdot\|_{U_i}$, $i = 1, 2$, the corresponding norms, i. e., $d_i(0,x) = \|x\|_{U_i} = \min\{\rho \geq 0 : x \in \rho U_i\}$ for $x \in \mathbb{R}^d$, $i = 1, 2$.

Suppose that there exists $x^{(0)} \in \mathbb{R}^d$ and $\delta > 0$ such that $x^{(0)} + \delta \mathbb{B}_2 \subset B_{1,2}$ where $\mathbb{B}_2$

is the Euclidean unit ball. This means that

$$h(\left\|x - s^{(1)}\right\|_{U_1}) - h(\left\|x - s^{(2)}\right\|_{U_2}) = \mu_2 - \mu_1 \tag{3.6}$$

holds for all $x \in x^{(0)} + \delta\mathbb{B}_2$. We may w.l.o.g. assume that $s^{(1)}, s^{(2)}$ and $x^{(0)}$ are affinely independent.

Next, let $a \in \mathbb{R}^d \setminus \{0\}$ such that

$$H^{\leq}_{a,a^\intercal x^{(0)}} := \{x \in \mathbb{R}^d : a^\intercal x \leq a^\intercal x^{(0)}\}$$

is a halfspace that supports $s^{(1)} + \left\|x^{(0)} - s^{(1)}\right\|_{U_1} U_1$ in $x^{(0)}$.

If $H^{\leq}_{a,a^\intercal x^{(0)}}$ does not support $s^{(2)} + \left\|x^{(0)} - s^{(2)}\right\|_{U_2} U_2$ in $x^{(0)}$, it follows that there exists $z \in x^{(0)} + \delta\mathbb{B}_2$ with $\left\|z - s^{(1)}\right\|_{U_1} = \left\|x^{(0)} - s^{(1)}\right\|_{U_1}$ and $\left\|z - s^{(2)}\right\|_{U_2} \neq \left\|x^{(0)} - s^{(2)}\right\|_{U_2}$. As h is injective, this implies that Eq. (3.6) does not hold for z, a contradiction. Hence, $H^{\leq}_{a,a^\intercal x^{(0)}}$ must support $s^{(2)} + \left\|x^{(0)} - s^{(2)}\right\|_{U_2} U_2$ in $x^{(0)}$.

Furthermore, by continuity there exist $\lambda > 1$ and $\nu \in \mathbb{R}$ such that for $x^{(1)} := s^{(1)} + \lambda(x^{(0)} - s^{(1)})$ and $x^{(2)} := s^{(2)} + \nu(x^{(0)} - s^{(2)})$ it holds that $x^{(2)} \in \operatorname{int}(x^{(0)} + \delta\mathbb{B}_2)$ and $\left\|x^{(1)} - s^{(1)}\right\|_{U_1} = \left\|x^{(2)} - s^{(1)}\right\|_{U_1}$. Furthermore, due to the affine independence of $s^{(1)}, s^{(2)}$ and $x^{(0)}$ we have that $x^{(1)} \neq x^{(2)}$.

As $H^{\leq}_{a,a^\intercal x^{(0)}}$ supports $s^{(1)} + \left\|x^{(0)} - s^{(1)}\right\|_{U_1} U_1$ in $x^{(0)}$, it follows that $H^{\leq}_{a,a^\intercal x^{(1)}}$ supports $s^{(1)} + \left\|x^{(1)} - s^{(1)}\right\|_{U_1} U_1$ in $s^{(1)} + \frac{\left\|x^{(1)} - s^{(1)}\right\|_{U_1}}{\left\|x^{(0)} - s^{(1)}\right\|_{U_1}}(x^{(0)} - s^{(1)}) = x^{(1)}$. Analogously, $H^{\leq}_{a,a^\intercal x^{(2)}}$ supports $s^{(2)} + \left\|x^{(2)} - s^{(2)}\right\|_{U_2} U_2$ in $x^{(2)}$. By the same argumentation as before, we see that $H^{\leq}_{a,a^\intercal x^{(2)}}$ must also support $s^{(1)} + \left\|x^{(2)} - s^{(1)}\right\|_{U_1} U_1 = s^{(1)} + \left\|x^{(1)} - s^{(1)}\right\|_{U_1} U_1$ in $x^{(2)}$ (as otherwise we find a point that contradicts Eq. (3.6)).

Hence, $s^{(1)} + \left\|x^{(1)} - s^{(1)}\right\|_{U_1} U_1$ is supported in $x^{(1)}$ and $x^{(2)}$ by the halfspaces $H^{\leq}_{a,a^\intercal x^{(1)}}$ and $H^{\leq}_{a,a^\intercal x^{(2)}}$, respectively. This contradicts the strict convexity of U_1. □

Many of the generalizations considered in the literature match the form of Eq. (3.2). Since Shamos and Hoey [SH75] presented an efficient construction algorithm for planar Voronoi Diagrams (i. e., the case $\mathcal{X} = \mathbb{R}^2$, $h = \mathrm{id}$, $\mu_i = 0$ and d_i being the Euclidean metric for every $i \in [k]$), many similar construction algorithms for generalized Voronoi diagrams in the plane have followed. Hwang [Hwa79] as well as Lee and Wong [LW80] consider the case of d_i being the L_1 or L_∞ metric. Lee and Drysdale, III [LD81] use line segments as sites. (Those can be brought to the form of Eq. (3.2) by choosing $s^{(i)}$ as the line segments middle points. Further, choose the metric d_i to be induced by the norm

of the unit ball obtained as the Minkowski sum of the line segment and the Euclidean unit ball.) Chew and Dyrsdale [CD85] then study the construction of Voronoi diagrams in planar (possibly asymmetric) Minkowski spaces, i. e., for $\mathcal{X} = \mathbb{R}^2, h = \mathrm{id}, \mu_i = 0$, and with d_i being induced by a (cell-independent, possibly asymmetric) Minkowski norm. Ash and Bolker [AB86] as well consider the planar case $\mathcal{X} = \mathbb{R}^2$ and provide characterizations of diagrams for $d(x, y) := \|x - y\|_2$ and $h(\cdot) = (\cdot)^2$ or $h = \mathrm{id}$ as well as $d_i(x, y) := \alpha_i \|x - y\|_2$ for $\alpha_i > 0$, $h = \mathrm{id}$ and $\mu_i = 0$. Icking et al. [Ick+01] research the bisectors, and in particular their number of connected components, for generalized Voronoi diagrams in the plane for $h = \mathrm{id}$ and $d_i := \left\|x - s^{(i)}\right\|_{B_i}$, with $\|\cdot\|_{B_i}$ being a (possibly asymmetric) norm defined by some convex (not necessarily symmetric) unit ball B_i for each $i \in [k]$.

This, of course, only yields a glimpse of the many variants of Voronoi diagrams researched in the literature. Still, it underlines that the class of functions of type (3.2) seems accurate in order to produce a useful theory. We will stick to the very general form of Definition 3.1 in order to define the major relations between constrained clusterings and generalized Voronoi diagrams. For all applications, however, we will consider diagrams of type (3.2).

3.3 Basic Types

Let us introduce some classes of generalized Voronoi diagrams that are important in our context and discuss the properties that are most relevant to us. For more comprehensive surveys of generalizations we refer to [AK00, Chapters 6 – 7] and [Oka+00, Chapter 3].

3.3.1 Additively Weighted Voronoi Diagrams

A first generalization of classical Voronoi diagrams is achieved by the introduction of additive weights for each cell. In the schema of Eq. (3.2), this means that we consider diagrams with $\mathcal{X} := \mathbb{R}^d$, $h := \mathrm{id}$, and d_i being the Euclidean metric for every $i \in [k]$. Thus, we obtain defining functions of the form

$$f_i(x) = \left\|x - s^{(i)}\right\|_2 + \mu_i \tag{3.7}$$

for sites $s^{(i)} \in \mathbb{R}^d$ and additive weights $\mu_i \in \mathbb{R}$ for $i \in [k]$. We will call a generalized Voronoi diagram of the form (3.7) an *additively weighted Voronoi diagram*.

Additively weighted Voronoi diagrams have occurred in the literature multiple times. In particular, in the context of mineralogy they are known as *Johnson-Mehl-Model* ([JM39]). Here, the diagram is considered to be the result of a growth process. Seeds of cells are generated at different points in time and then start growing radially at constant rate. The value $f_i(x)$ hence yields the point in time at which the ith cell,

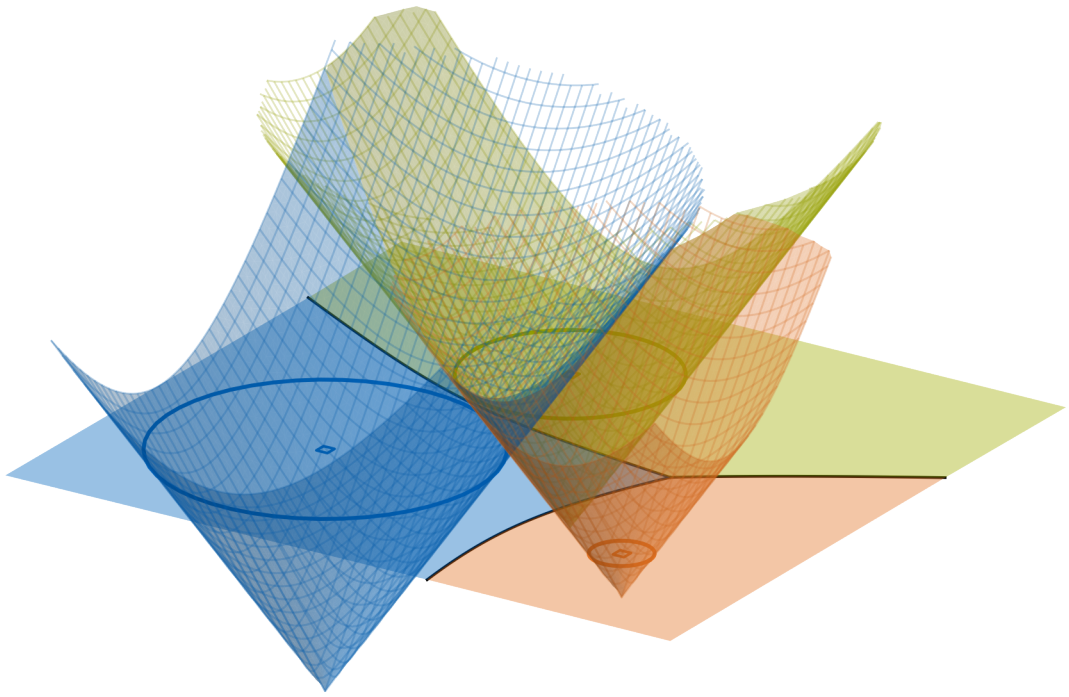

Figure 3.4: Voronoi surface of an additively weighted Voronoi diagram.

which started growing in $s^{(i)}$ at time μ_i, reaches point x. The growing process in a certain direction stops whenever another cell collides. As an example, Johnson and Mehl [JM39] describe the process of freezing metal in which frozen nuclei are formed in a stochastic manner and then — in an ideal isothermal setting — grow radially.

Ash and Bolker [AB86], who called them *hyperbolic tessellations*, state necessary and sufficient conditions to decide whether a given tessellation in the plane is an additively weighted Voronoi diagram. Several algorithms for constructing additively weighted Voronoi diagrams in the plane have been proposed resulting in a complexity of $\mathcal{O}(k \log(k))$ ([LD81] , [Sha85], and [For87]). It has been shown that single cells of additively weighted Voronoi diagrams are combinatorically equivalent to the intersection of power diagrams (see Section 3.3.2) and the unit sphere, which yields a construction algorithm for a whole additively weighted Voronoi diagram in $\mathbb{R}^d$ of complexity $\mathcal{O}\left(k^2 \log(k) + k^{\lceil \frac{d}{2} \rceil + 1}\right)$ ([BD05; BK03]).

Let us clarify some characteristics of the cells of an additively weighted Voronoi diagram. Due to Lemma 3.2 we may w.l.o.g. assume that $\mu_i < 0$ for every $i \in [k]$ (as we may otherwise add a sufficiently negative constant b to every function f_i without changing the diagram). This allows us to associate every cell of the resulting diagram with an Euclidean sphere with center $s^{(i)}$ and radius $\sigma_i := -\mu_i$. The distance

measure associated with the ith cell is hence the *signed* Euclidean distance to the cell's associated sphere, where "signed" means that points contained in the corresponding ball are assigned the negative distance.

For each $i \in [k]$, the graph of the function f_i is hence

$$\begin{pmatrix} s^{(i)} \\ -\sigma_i \end{pmatrix} + [0,\infty) \cdot \sigma_i \cdot \left(\mathbb{S}^{d-1} \times \{1\} \right) \subset \mathbb{R}^{d+1},$$

i. e., the conical (hyper)surface with apex $((s^{(i)})^\intercal, -\sigma_i)^\intercal$ through the cell's associated sphere embedded into the hyperplane $x_{d+1} = 0$. Consequently, a bisector between two cells is obtained as the projection of the intersection of the two associated cones, which yields a sheet of a hyperboloid in $\mathbb{R}^d$. Figure 3.4 depicts this for an example of $k = 3$ sites in $\mathbb{R}^2$. The resulting diagram is drawn by its embedding into the x_1x_3-plane, i. e., $\{x \in \mathbb{R}^3 : x_3 = 0\}$.

Let us obtain a better understanding of how the choice of the sites and sphere radii influences the bisectors and hence the shape of the diagram cells. For this purpose, let two sites $s^{(1)}$ and $s^{(2)}$ in $\mathbb{R}^d$ as well as radii $\sigma_1, \sigma_2 \in \mathbb{R}_{>0}$ be given. We may w.l.o.g. assume that $\sigma_1 \geq \sigma_2$, otherwise we relabel. We may furthermore w.l.o.g. assume that $s^{(1)} + s^{(2)} = 0$ and that $s^{(2)} - s^{(1)} = \left\| s^{(2)} - s^{(1)} \right\|_2 u^{(1)}$, otherwise we apply a translation and rotation. Next, we assume $s^{(1)} \neq s^{(2)}$, as otherwise either $P_2 = \emptyset$ or $P_2 = P_1$ holds (depending on $\sigma_1 > \sigma_2$ or $\sigma_1 = \sigma_2$). The case $\sigma_1 = \sigma_2$ yields the bisector of a classical Voronoi diagram, i. e., the hyperplane $\left\{ x \in \mathbb{R}^d : x^\intercal (s^{(2)} - s^{(1)}) = \frac{1}{2} \left(\left\| s^{(2)} \right\|_2^2 - \left\| s^{(1)} \right\|_2^2 \right) \right\}$. Thus, we finally assume $\sigma_1 > \sigma_2$. Figure 3.5 depicts such a situation in the plane.

The bisector $B_{1,2}$ is given by all points $x \in \mathbb{R}^d$ that satisfy $f_1(x) = f_2(x)$, i. e.,

$$\left\| x - s^{(1)} \right\|_2 - \sigma_1 = \left\| x - s^{(2)} \right\|_2 - \sigma_2 \tag{3.8}$$

$$\Leftrightarrow \left\| x - s^{(1)} \right\|_2 - \left\| x - s^{(2)} \right\|_2 = \sigma_1 - \sigma_2. \tag{3.9}$$

Equation (3.9) immediately reads as the set of all points that have a constant difference of distances to $s^{(1)}$ and $s^{(2)}$, respectively, namely $\sigma_1 - \sigma_2$. Thus, the sites $s^{(1)}$ and $s^{(2)}$ are the foci of the desired hyperboloid (if existent). For the sake of simplicity, let us introduce the two abbreviations

$$\Delta s := \left\| s^{(1)} - s^{(2)} \right\|_2 > 0 \quad \text{and} \quad \Delta\sigma := \sigma_1 - \sigma_2 > 0.$$

We now rewrite Eq. (3.8) further using our assumptions $s^{(1)} = -s^{(2)} = -\frac{1}{2}\Delta s \cdot u^{(1)}$

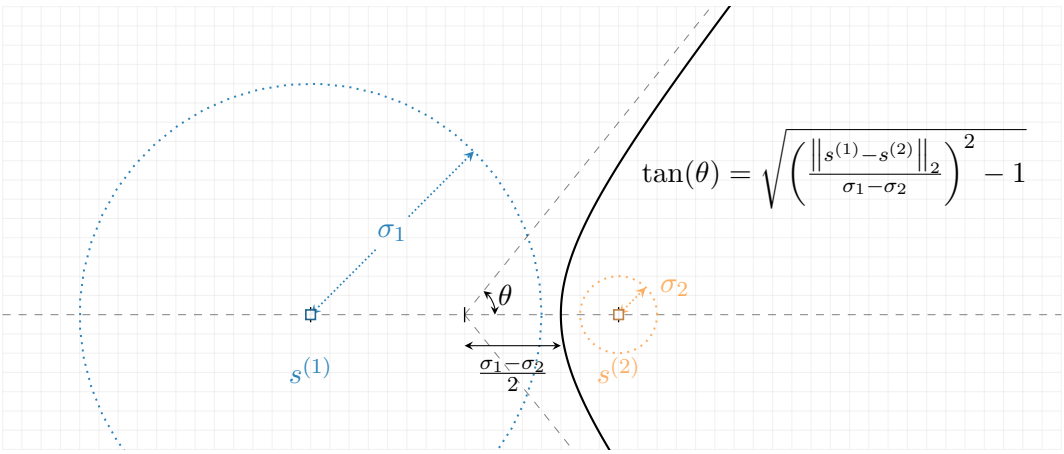

Figure 3.5: Determination of the bisector, i. e., a hyperbola, between two cells of an additively weighted Voronoi diagram in the plane for two sites $s^{(1)}$ and $s^{(2)}$ with sphere radii $\sigma_1 > \sigma_2$.

as well as $\Delta\sigma > 0$:

$$\begin{aligned}
&\left\|x - s^{(1)}\right\|_2 = \left\|x - s^{(2)}\right\|_2 + \Delta\sigma && \left|(\cdot)^2\right. \\
\Leftrightarrow\ & 4x^\intercal s^{(2)} - \Delta\sigma^2 = 2\Delta\sigma \left\|x - s^{(2)}\right\|_2 && \left|(\cdot)^2\right. \\
\Leftrightarrow\ & 0 = x^\intercal x - \frac{4(x^\intercal s^{(2)})^2}{\Delta\sigma^2} - \frac{1}{4}\Delta\sigma + \left\|s^{(2)}\right\|_2^2 \quad \wedge \quad x^\intercal s^{(2)} \geq \frac{1}{4}\Delta\sigma^2 \\
\Leftrightarrow\ & \left(1 - \frac{\Delta s^2}{\Delta\sigma^2}\right) x_1^2 + \sum_{i\geq 2} x_i^2 = \frac{1}{4}(\Delta\sigma^2 - \Delta s^2) \quad \wedge \quad x_1 \geq \frac{1}{2}\frac{\Delta\sigma^2}{\Delta s}
\end{aligned} \tag{3.10}$$

We may distinguish now three final cases. First, assume $\Delta\sigma = \Delta s$. Then Eq. (3.10) becomes

$$\sum_{i\geq 2} x_i^2 = 0 \quad \wedge \quad x_1 \geq \frac{1}{2}\Delta\sigma. \tag{3.11}$$

In this (degenerated) case, the bisector is the half-ray $s^{(2)} + [0, \infty) \cdot u^{(1)}$. From a more geometric perspective, in this case the sphere belonging to the second cluster is contained in the (convex hull of) the first one, with a touching point in $s^{(2)} + \sigma_2 u^{(1)}$. Consequently, the only points having equal distance to both spheres are given by the corresponding cone of outer normals translated into $s^{(2)}$, which yields the calculated half-ray. Figure 3.6 depicts this case.

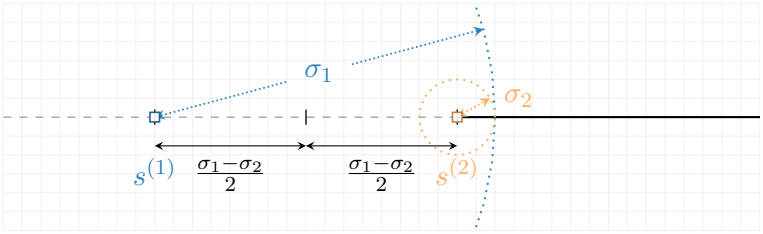

Figure 3.6: Situation of Fig. 3.5 for the degenerated case $\sigma_1 - \sigma_2 = \left\|s^{(1)} - s^{(2)}\right\|_2$. The resulting bisector is given by the black-drawn half-ray.

If $\Delta\sigma \neq \Delta s$, we may simplify Eq. (3.10) further to

$$\left(\frac{4}{\Delta\sigma^2}\right) x_1^2 + \frac{4}{\Delta\sigma^2 - \Delta s^2} \sum_{i \geq 2} x_i^2 = 1 \quad \wedge \quad x_1 \geq \frac{1}{2} \frac{\Delta\sigma^2}{\Delta s} \tag{3.12}$$

Let us briefly discuss the second case $\Delta\sigma > \Delta s$. Geometrically, this means that the interior of the (convex hull of the) first sphere contains the second one, i. e., $s^{(2)} + \sigma_2 \mathbb{B}_2^d \subset \operatorname{int}\left(s^{(1)} + \sigma_1 \mathbb{B}_2^d\right)$. From this perspective, it is immediately clear that the signed distance of any point to the first sphere will always be strictly less than the signed distance to the second sphere. Thus, the second cell and in particular the bisector of interest are empty. In view of Eq. (3.12), the equality yields an ellipsoid with axes aligned to the coordinate system and a first elliptic radius of $\frac{1}{2}\Delta s < \frac{1}{2}\Delta s \cdot \frac{\Delta s}{\Delta\sigma}$. Thus, any point on that ellipsoid violates the inequality in (3.12).

Finally, let us consider the case $\Delta\sigma < \Delta s$. This holds whenever the spheres are not contained in (the convex hulls of) one another. Then, for $d \geq 3$ the bisector is one sheet of the hyperboloid obtained as the (hyper)surface of revolution when rotating a hyperbola around the x_1-axis in all remaining dimensions (i. e., every intersection of this hypersurface with a hyperplane perpendicular to the x_1-axis is either empty or a dilatation of $\mathbb{S}^{d-2}$). For $d = 2$ there is no dimension for rotation left, so the bisector is just the according branch of the hyperbola. The main characteristics of this hyperbola are readily given by Eq. (3.12). In particular, we obtain the vertex $\frac{\Delta\sigma}{2} \cdot u^{(1)}$ and an asymptotic slope $\sqrt{\frac{\frac{1}{4}(\Delta s^2 - \Delta\sigma^2)}{\frac{1}{4}\Delta\sigma^2}} = \sqrt{\left(\frac{\Delta s}{\Delta\sigma}\right)^2 - 1}$ (cf. Fig. 3.5). For $\Delta s \gg \Delta\sigma$, i. e., when the two spheres are very distant, this slope tends to infinity and thus the bisector tends to the bisector of the classical Voronoi diagram in form of a hyperplane as discussed for the case $\Delta\sigma = 0$. For $\Delta s \to \Delta\sigma$ the bisector tends to the degenerated case of Fig. 3.6.

With this understanding of bisectors we may state some basic facts about the geometric characteristics of cells of additively weighted Voronoi diagrams. Those can also be found in this or very similar forms in [AB86; Sha85].

Proposition 3.7
Let $\mathcal{P} = (P_1, \dots, P_k)$ be an additively Voronoi diagram in $\mathbb{R}^d$ w. r. t. functions $(f_i)_{i\in[k]}$ given by $f_i(x) = \left\|x - s^{(i)}\right\|_2 - \sigma_i$ for $i \in [k]$ with $s^{(i)} \in \mathbb{R}^d$, $\sigma_i \in \mathbb{R}_{>0}$ and such that $f_i \neq f_l$ for all $i \neq l \in [k]$. Fix some $i \in [k]$. Then the following holds:

i) $\exists l \in [k] : s^{(i)} + \sigma_i \mathbb{B}_2^d \subset \operatorname{int}\left(s^{(l)} + \sigma_l \mathbb{B}_2^d\right) \Rightarrow P_i = \emptyset$

ii) $P_i \neq \emptyset \Rightarrow s^{(i)} \in P_i$

iii) P_i *is star-shaped w. r. t.* $s^{(i)}$*, i. e.,* $x \in P_i \Rightarrow [s^{(i)}, x] \subset P_i$

iv) $\left(s^{(i)} + \sigma_i \mathbb{B}_2^d\right) \setminus \left(s^{(l)} + \sigma_l \mathbb{B}_2^d\right) \neq \emptyset \ \forall l \in [k] \setminus \{i\} \Leftrightarrow s^{(i)} \in \operatorname{int}(P_i) \neq \emptyset$

Proof. First note that the assumptions of the f_i being pairwise unequal is equivalent to the assumption of the corresponding spheres being pairwise distinct. This is a very mild assumption as cells of identical functions would coincide, of course.

The prerequisite of part i is the case if and only if $\left\|s^{(l)} - s^{(i)}\right\|_2 < \sigma_l - \sigma_i$, so this immediately follows from the observations above. Parts ii and iii are directly provided by Corollary 3.5 and Lemma 3.3, respectively.

For part iv, it is sufficient to show that for any $l \in [k] \setminus \{i\}$ we have

$$\left(s^{(i)} + \sigma_i \mathbb{B}_2^d\right) \setminus \left(s^{(l)} + \sigma_l \mathbb{B}_2^d\right) \neq \emptyset \ \Leftrightarrow \ s^{(i)} \in \operatorname{int}(H_{i,l}) \ \Leftrightarrow \ \operatorname{int}(H_{i,l}) \neq \emptyset.$$

Here, the latter equivalence immediately follows from our observations above. However, we have seen that $\operatorname{int}(H_{i,l}) = \emptyset$ holds if and only if either we are in the degenerated case of $s^{(i)} + \sigma_i \mathbb{B}_2^d$ touching $s^{(l)} + \sigma_l \mathbb{B}_2^d$ from the inside in a single point (as $s^{(i)} + \sigma_i \mathbb{B}_2^d \neq s^{(l)} + \sigma_l \mathbb{B}_2^d$ by assumption) or $s^{(i)} + \sigma_i \mathbb{B}_2^d \subset \operatorname{int}(s^{(l)} + \sigma_l \mathbb{B}_2^d)$. This is the case if and only if $\left(s^{(i)} + \sigma_i \mathbb{B}_2^d\right) \setminus \left(s^{(l)} + \sigma_l \mathbb{B}_2^d\right) = \emptyset$ holds. □

3.3.2 Power Diagrams

The next generalization we want to consider are *power diagrams*. Due to their favorable geometric properties, power diagrams are of particular importance in our context of constrained clustering (cf. Chapter 6). From a rather theoretical point of view, they are also of importance as many types of generalized Voronoi diagrams correspond to transformations of power diagrams (cf. [BD05; BK03; BWY06; Wor08]; see also Section 3.3.3).

Besides that, they have several applications such as the illumination of balls, organic compounds, or muscle fibres (to name just a few; see [Aur87a] and [Oka+00, Section 3.1.6] for a more detailed list and references). Aurenhammer [Aur87a] provides a comprehensive treatment of power diagrams, further details and references can also be found in [AKL13, Section 6.2] and [Oka+00, Section 3.1.4].

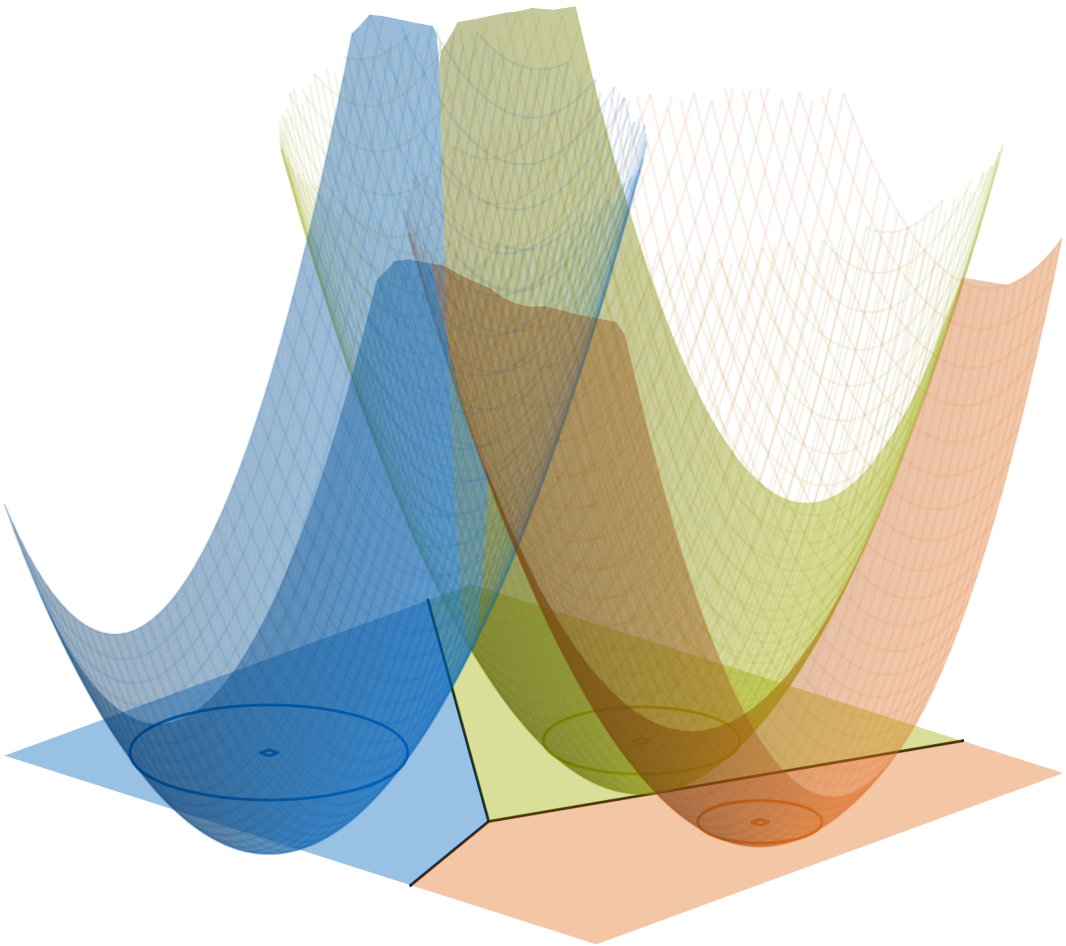

Figure 3.7: Voronoi surface of a power diagram.

Definition and Basic Properties

We again consider diagrams in d-dimensional Euclidean space, but this time with squared distances. This means, we set $h = (\cdot)^2$ in the schema of Eq. (3.2). Thus, we consider functions

$$f_i(x) = \left\| x - s^{(i)} \right\|_2^2 + \mu_i \tag{3.13}$$

for sites $s^{(i)} \in \mathbb{R}^d$ and additive weights $\mu_i \in \mathbb{R}$ for $i \in [k]$, and call the resulting generalized Voronoi diagram $\mathcal{P}$ w. r. t. $(f_i)_{i\in[k]}$ a *power diagram*. We will denote $\mathcal{P} = \mathrm{PD}\left(\left(s^{(i)}, \mu_i\right)_{i\in[k]}\right)$.

While the term "power diagram" is in coherence with most of the literature known to

us, they are also known as *sectional Dirichlet tessellation* ([AB86]), or *Voronoi diagram in the Laguerre geometry* ([IIM85]).

Again, we can w.l.o.g. assume that $\mu_i < 0$ for every $i \in [k]$. Under this assumption, we can again associate each cell for $i \in [k]$ with a sphere centered in $s^{(i)}$ of radius $\sigma_i := \sqrt{-\mu_i}$. For a point outside this sphere, the value $f_i(x)$ is then the squared distance of the point x and the intersection point of a tangent of that sphere through x. Functions of the type (3.13) are also referred to as *power functions* (cf. [Aur87a]). For any point $x \in \mathbb{R}^d$ the value $\left\|x - s^{(i)}\right\|_2^2 + \mu_i$ is then called the *power* of x w. r. t. the ith sphere.

The graph of a function f_i is consequently a circular paraboloid with apex in $\left(s^{(i)}, \mu_i\right)$ that intersects the hyperplane $x_{d+1} = 0$ in the embedded sphere associated with i, i. e., $s^{(i)} + \sigma_i \mathbb{S}^{d-1}$. Figure 3.7 depicts the graphs of a power diagram of 3 cells in the plane and the resulting Voronoi surface.

For $i \neq l \in [k]$ the bisector $B_{i,l}$ is given by

$$\begin{aligned}
&\left\|x - s^{(i)}\right\|_2^2 + \mu_i = \left\|x - s^{(l)}\right\|_2^2 + \mu_l \\
\Leftrightarrow\ &\|x\|_2^2 - 2x^\intercal s^{(i)} + \left\|s^{(i)}\right\|_2^2 + \mu_i = \|x\|_2^2 - 2x^\intercal s^{(l)} + \left\|s^{(l)}\right\|_2^2 + \mu_l \\
\Leftrightarrow\ &- 2x^\intercal s^{(i)} + \left\|s^{(i)}\right\|_2^2 + \mu_i = -2x^\intercal s^{(l)} + \left\|s^{(l)}\right\|_2^2 + \mu_l \qquad (3.14) \\
\Leftrightarrow\ &x^\intercal \left(2s^{(l)} - 2s^{(i)}\right) + \left\|s^{(i)}\right\|_2^2 + \mu_i - \left\|s^{(l)}\right\|_2^2 - \mu_l = 0 \qquad (3.15)
\end{aligned}$$

Thus, the bisector $B_{i,l}$ is the hyperplane with outer normal $s^{(l)} - s^{(i)}$ that lies midway between $s^{(l)}$ and $s^{(i)}$ shifted by $\frac{\mu_l - \mu_i}{2\left\|s^{(l)} - s^{(i)}\right\|_2}$ towards $s^{(l)}$, cf. Fig. 3.8.

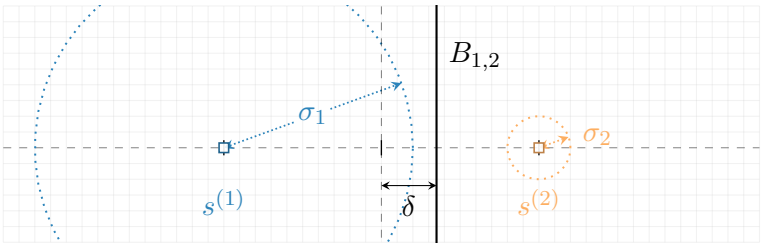

Figure 3.8: Bisector $B_{1,2}$ of a power diagram with two spheres centered in $s^{(1)}$ and $s^{(2)}$ and radii σ_1 and σ_2, respectively. The bisector is the midway perpendicular hyperplane between $s^{(1)}$ and $s^{(2)}$ shifted by $\delta = \frac{\sigma_1^2 - \sigma_2^2}{2\left\|s^{(1)} - s^{(2)}\right\|_2}$ towards $s^{(2)}$.

From this observation it already follows that every cell of a power diagram is a polyhedron. However, the following parametrization of power diagrams might yield a more straightforward view on the structure of power diagrams.

Affine Parametrization

We reconsider Eq. (3.14) and define a transformation function

$$\Phi : \begin{cases} \mathbb{R}^d \times \mathbb{R} & \to \mathbb{R}^d \times \mathbb{R} \\ (s, \mu) & \mapsto \Phi(s, \mu) := (-2s, \|s\|_2^2 + \mu). \end{cases} \tag{3.16}$$

We set $(a_i, \alpha_i) := \Phi\left(s^{(i)}, \mu_i\right)$ and

$$\hat{f}_i(x) := {a_i}^\intercal x + \alpha_i \tag{3.17}$$

for every $i \in [k]$. From comparison with Eq. (3.14) we see that the generalized Voronoi diagram $\mathcal{P}$ w. r. t. $\left(\hat{f}_i\right)_{i\in[k]}$ is equal to the one w. r. t. the functions given by Eq. (3.13). Obviously, Φ is a bijection with

$$\Phi^{-1}(a, \alpha) = \left(-\frac{1}{2}a, \alpha - \frac{1}{4}\|a\|_2^2\right) \tag{3.18}$$

for $(a, \alpha) \in \mathbb{R}^d \times \mathbb{R}$. Hence, both parametrizations are equivalent (this is also observed in [Wor08] and [Bor15]). Alternatively, Lemma 3.2 with the term $b(x) := x^2$ being added to f_i or subtracted from $\hat{f}_i$ for every $i \in [k]$, respectively, yields this observation, too.

We will refer to the parametrization (3.17) as *affine* and denote it

$$\mathcal{P} = \mathrm{PD}_{\mathrm{aff}}\left((a_i, \alpha_i)_{i\in[k]}\right).$$

The first parametrization as given by Eq. (3.13) will be called *spherical*.

The affine representation of power diagrams simplifies some insights into their basic properties. In particular, it establishes the following correspondence between power diagrams in $\mathbb{R}^d$ and polyhedra in $\mathbb{R}^{d+1}$. Figure 3.9 depicts the Voronoi surface of the same diagram as in Fig. 3.7 but in affine parametrization. This is a lower envelope obtained from hyperplanes in $\mathbb{R}^{d+1}$. We define the polyhedron

$$\overline{P} := \left\{ \begin{pmatrix} x \\ x_{d+1} \end{pmatrix} \in \mathbb{R}^{d+1} : x_{d+1} \leq {a_i}^\intercal x + \alpha_i \forall i \in [k] \right\} = \bigcap_{i\in[k]} H^{\leq}_{\left(\binom{-a_i}{1}, \alpha_i\right)}. \tag{3.19}$$

Then the Voronoi surface is the union of all facets of $\overline{P}$ and every non-empty cell P_i is the orthogonal projection of the facet of $\overline{P}$ with outer normal $\begin{pmatrix} -a_i \\ 1 \end{pmatrix}$ onto $\mathbb{R}^d$. Moreover, this implies that a power diagram yields a cell complex that is up to projection identical to the boundary complex of $\overline{P}$ (as defined in [Grü03], for example). Vice versa, from Eq. (3.19) we see that for any polyhedron in $\mathbb{R}^{d+1}$, its upper surface, i. e., the union of

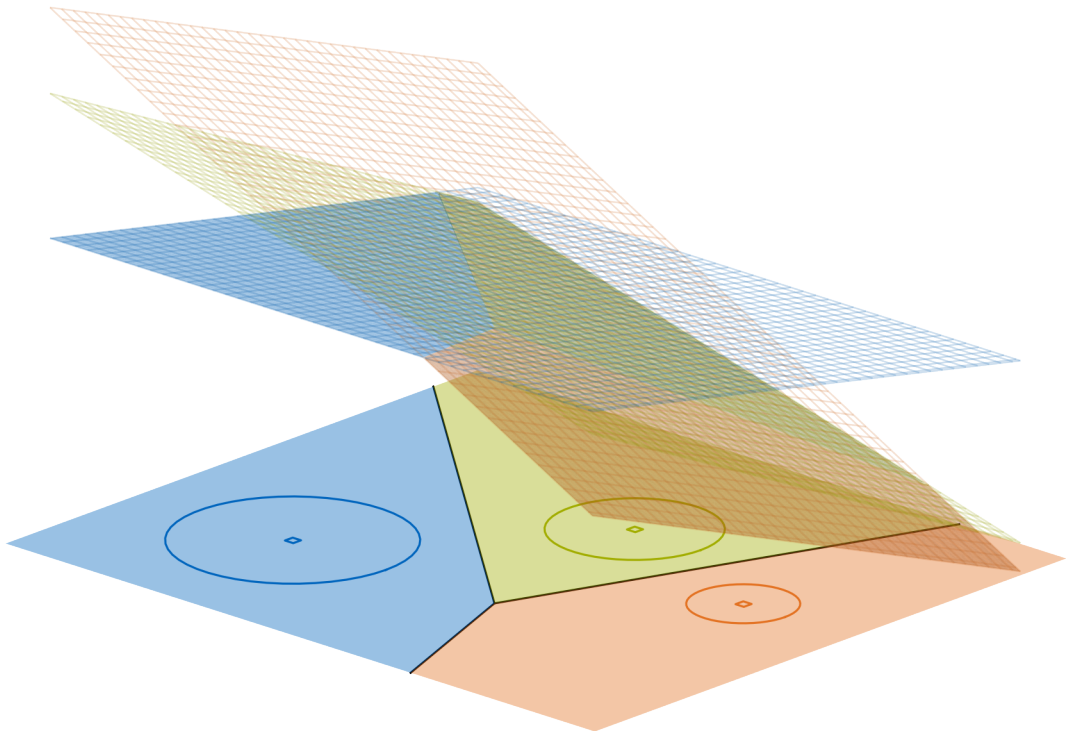

Figure 3.9: Voronoi surface of the power diagram in Fig. 3.7 but w. r. t. the affine representation of Eq. (3.17). The surface has been vertically shifted and scaled for the sake of a better visualization.

all facets with outer normals with positive $(d+1)$st coordinate, can be identified with a power diagram in $\mathbb{R}^d$.

As a byproduct, one obtains an algorithm to construct power diagrams in general dimension by constructing the corresponding upper surfaces of $\overline{P}$ using polarity as proposed by Aurenhammer [Aur87b]. Let us briefly describe its outline. In order to obtain an irredundant representation of $\overline{P}$, we can consider the cone

$$K := \left\{ \begin{pmatrix} x \\ x_{d+1} \\ x_{d+2} \end{pmatrix} \in \mathbb{R}^{d+1} : \; -{a_i}^\intercal x + x_{d+1} - \alpha_i x_{d+2} \leq 0 \; \forall i \in [k] \right\}.$$

This means, $\overline{P} \times \{1\} = K \cap \{x_{d+2} = 1\}$. With

$$D := \operatorname{conv}\left(\left\{ \begin{pmatrix} -a_i \\ -\alpha_i \end{pmatrix} : i \in [k] \right\} \right)$$

it follows that $K = \bigcap_{\binom{a}{\alpha} \in D} H^{\leq}_{((a^\intercal,1,\alpha)^\intercal,0)}$. Basic arguments from convex geometry then

yield that $H_{(\binom{x}{1},\beta)} \cap D$ is a facet of D if and only if $\begin{pmatrix} x \\ \beta \\ 1 \end{pmatrix}$ is a vertex of $K \cap \{x_{d+2} = 1\}$ and thus if and only if $\begin{pmatrix} x \\ \beta \end{pmatrix}$ is a vertex of $\overline{P}$. Hence, computing a power diagram, i. e. the full associate cell complex, can be done by computing the incidence structure of D and then only considering facets that have an outer normal pointing upwards, i. e., with positive $(d+1)$st component.

Due to this correspondence, one furthermore obtains sharp bounds on the complexity of the cell complex associated with power diagrams. Applying the upper bound theorem by McMullen [McM70] on the maximum number of j-faces of a $(d+1)$-polytope with a given number of vertices to D, one obtains sharp upper bounds for the number of j-faces of (the cell complex of) a power diagram, which for $0 \leq j \leq d-1$ is $\mathcal{O}\left(n^{\lceil \frac{d}{2} \rceil}\right)$ (see [Aur87a, Theorem 1]).

Note, however, that not every d-dimensional cell complex can be derived as the projection of a $(d+1)$-polyhedron and hence not every generalized Voronoi diagram with polyhedral cells is a power diagram (see [Aur87b] for a counterexample in dimension $d = 2$).

Invariances of Parametrizations

Due to Lemma 3.2 it is obvious that the choice of parameters for a power diagram in neither spherical (Eq. (3.13)) nor affine (Eq. (3.17)) parametrization is unique. For the affine parametrization of (3.17), it follows that the power diagram is invariant to a dilatation and translation of the parameters $(a_i, \alpha_i)_{i \in [k]}$:

Theorem 3.8
Let $d \geq 2$ and $\mathcal{P} = (P_1, \ldots, P_k)$ with $k \geq d+1$ be a power diagram such that every cell P_i has a vertex (i. e., every cell is non-empty and line-free), and every vertex is contained in exactly $d+1$ cells. Let $((a_i, \alpha_i))_{i \in [k]} \in (\mathbb{R}^d \times \mathbb{R})^k$ such that $\mathcal{P} = \mathrm{PD}_{\mathrm{aff}}\left((a_i, \alpha_i)_{i \in [k]}\right)$.

Then for $((b_i, \beta_i))_{i \in [k]} \in (\mathbb{R}^d \times \mathbb{R})^k$ it holds that $\mathcal{P} = \mathrm{PD}_{\mathrm{aff}}\left((b_i, \beta_i)_{i \in [k]}\right)$ if and only if there exist $\lambda > 0$, $\Delta a \in \mathbb{R}^d$ and $\Delta\alpha \in \mathbb{R}$ such that

$$\begin{pmatrix} b_i \\ \beta_i \end{pmatrix} = \lambda \cdot \begin{pmatrix} a_i \\ \alpha_i \end{pmatrix} + \begin{pmatrix} \Delta a \\ \Delta\alpha \end{pmatrix} \quad \text{for every } i \in [k]. \tag{3.20}$$

Proof. The second implication immediately follows from Lemma 3.2.

Now, let $((a_i, \alpha_i))_{i \in [k]}$ and $((b_i, \beta_i))_{i \in [k]}$ both be generators for $\mathcal{P}$. W.l.o.g. assume that the cells with indices in $[d+1]$ intersect in a vertex $v \in \mathbb{R}^d$.

Consider the corresponding Voronoi surface in $\mathbb{R}^{d+1}$. Then $\begin{pmatrix} v \\ {a_1}^\intercal v + \alpha_1 \end{pmatrix}$ is a regular vertex of the polyhedron $\overline{P}$ defined in Eq. (3.19). Thus, the outer normals $\begin{pmatrix} -a_i \\ 1 \end{pmatrix}$ for $i \in [d+1]$ are linearly independent.

Also, w.l.o.g. assume that $a_1 = 0$ and $\alpha_1 = 0$. Otherwise, first consider $((\tilde{a_i}, \tilde{\alpha_i}))_{i\in[k]}$ defined via $\tilde{a_i} := a_i - a_1$ and $\tilde{\alpha_i} := \alpha_i - \alpha_1$, which results in the same diagram due to Lemma 3.2. Then obtain $\tilde{\lambda}$, $\Delta\tilde{a}$, $\Delta\tilde{\alpha}$, and set $\lambda := \tilde{\lambda}$ and $(\Delta a, \Delta\alpha) := (\Delta\tilde{a}, \Delta\tilde{\alpha}) - \tilde{\lambda}(a_1, \alpha_1)$, which yields the desired result.

With $a_1 = 0$ and $\begin{pmatrix} a_i \\ -1 \end{pmatrix}$ for $i \in [d+1]$ being linearly independent, it follows that $a_2, \ldots, a_{d+1}$ are linearly independent (in $\mathbb{R}^d$), too.

Also, as v is a regular vertex, for every $i, l \in [d+1]$, $i \neq l$, it follows that $P_i \cap P_i$ is a $(d-1)$-dimensional facet of both P_i and P_l. That facet is contained in the hyperplane $H_{(a_i - a_l, -\alpha_i + \alpha_l)}$. As this must hold for the second set of generators, too, we can conclude that $H^{\leq}_{(a_i - a_l, -\alpha_i + \alpha_l)} = H^{\leq}_{(b_i - b_l, -\beta_i + \beta_l)}$ for every $i, l \in [d+1], i \neq l$.

Thus, we can conclude that there exist $\lambda_i > 0$ for $i \in [d+1]$ such that

$$\begin{pmatrix} b_i - b_{i+1} \\ -\beta_i + \beta_{i+1} \end{pmatrix} = \lambda_i \begin{pmatrix} a_i - a_{i+1} \\ -\alpha_i + \alpha_{i+1} \end{pmatrix} \forall i \in [d] \quad \wedge \quad \begin{pmatrix} b_{d+1} - b_1 \\ -\beta_{d+1} + \beta_1 \end{pmatrix} = \lambda_{d+1} \begin{pmatrix} a_{d+1} - a_1 \\ -\alpha_{d+1} + \alpha_1 \end{pmatrix}. \tag{3.21}$$

Thus, we get

$$\begin{aligned} 0 &= \sum_{i\in[d]} (b_i - b_{i+1}) \ + \ b_{d+1} - b_1 = \sum_{i\in[d]} \lambda_i (a_i - a_{i+1}) \ + \ \lambda_{d+1}(a_{d+1} - a_1) \\ &\overset{a_1 = 0}{=} \sum_{i\in[d+1]\setminus\{1\}} (\lambda_i - \lambda_{i-1}) a_i. \end{aligned}$$

Due to the linear independence of $a_2, \ldots, a_{d+1}$, we can conclude $\lambda_i = \lambda_l$ for all $i, l \in [d+1]$. Therefore, we define

$$\lambda := \lambda_1, \text{ and} \tag{3.22}$$

$$\begin{pmatrix} \Delta a \\ \Delta\alpha \end{pmatrix} := \begin{pmatrix} b_1 \\ \beta_1 \end{pmatrix} - \lambda \begin{pmatrix} a_1 \\ \alpha_1 \end{pmatrix}. \tag{3.23}$$

Together with Eq. (3.21) this inductively gives

$$\begin{aligned} \begin{pmatrix} b_i \\ \beta_i \end{pmatrix} &= \begin{pmatrix} b_{i-1} \\ \beta_{i-1} \end{pmatrix} + \begin{pmatrix} b_i - b_{i-1} \\ \beta_i - \beta_{i-1} \end{pmatrix} = \lambda \begin{pmatrix} a_{i-1} \\ \alpha_{i-1} \end{pmatrix} + \begin{pmatrix} \Delta a \\ \Delta\alpha \end{pmatrix} + \lambda \begin{pmatrix} a_i - a_{i-1} \\ \alpha_i - \alpha_{i-1} \end{pmatrix} \\ &= \lambda \begin{pmatrix} a_i \\ \alpha_i \end{pmatrix} + \begin{pmatrix} \Delta a \\ \Delta\alpha \end{pmatrix} \end{aligned}$$

for every $i \in [d+1] \setminus \{1\}$.

We may do so for every vertex of the diagram. If two vertices are contained in a common bisector, it follows from Eqs. (3.21) and (3.22) that the resulting value for the scaling factor λ obtained for each of the two vertices must be identical. Due to Eq. (3.23), the same holds for the resulting translation $\begin{pmatrix}\Delta a \\ \Delta\alpha\end{pmatrix}$. Now by assumption, every cell of the diagram has a vertex. The graph with a node for every cell and an edge between cells $i \neq l \in [k]$ whenever $P_i \cap P_l \neq \emptyset$ is connected (note that we can identify this graph with the 1-skeleton of the polar of $\overline{P}$). Therefore, Eq. (3.20) holds.□

We can easily translate Theorem 3.8 into the spherical parametrization. Here, we still have the freedom of a dilatation and a translation, but obtain further correction terms for the additive weights.

Corollary 3.9
Let $d \geq 2$ and $\mathcal{P} = (P_1, \ldots, P_k)$ with $k \geq d+1$ be a power diagram such that every cell P_i has a vertex and every vertex is contained in exactly $d+1$ cells for every $i \in [k]$. Let $\left((s^{(i)}, \mu_i)\right)_{i\in[k]} \in (\mathbb{R}^d \times \mathbb{R})^k$ such that $\mathcal{P} = \mathrm{PD}\left(\left(s^{(i)}, \mu_i\right)_{i\in[k]}\right)$.

Then for $\left((\hat{s}^{(i)}, \hat{\mu}_i)\right)_{i\in[k]} \in (\mathbb{R}^d \times \mathbb{R})^k$ it holds that $\mathcal{P} = \mathrm{PD}\left(\left((\hat{s}^{(i)}, \hat{\mu}_i)\right)_{i\in[k]}\right)$ if and only if there exist $\lambda > 0$, $\Delta s \in \mathbb{R}^d$ and $\Delta\mu \in \mathbb{R}$ such that for every $i \in [k]$ it holds that

$$\hat{s}^{(i)} = \lambda s^{(i)} + \Delta s \tag{3.24}$$

and

$$\hat{\mu}_i = \lambda\mu_i \ + \ \Delta\mu \ + \ \lambda\left\|s^{(i)}\right\|_2^2 - \left\|\lambda s^{(i)} + \Delta s\right\|_2^2. \tag{3.25}$$

Proof. With Φ as defined by Eq. (3.16) it holds that

$$\mathrm{PD}\left(\left((s^{(i)}, \mu_i)\right)_{i\in[k]}\right) = \mathrm{PD}\left(\left((\hat{s}^{(i)}, \hat{\mu}_i)\right)_{i\in[k]}\right)$$

if and only if $\mathrm{PD}_{\mathrm{aff}}\left((a_i, \alpha_i)_{i\in[k]}\right) = \mathrm{PD}_{\mathrm{aff}}\left((b_i, \beta_i)_{i\in[k]}\right)$ with

$$(a_i, \alpha_i) := \Phi(s^{(i)}, \mu_i) \text{ and } (b_i, \beta_i) := \Phi(\hat{s}^{(i)}, \hat{\mu}_i)$$

for every $i \in [k]$. By Theorem 3.8 this is equivalent to the existence of $\lambda > 0$, $\Delta a \in \mathbb{R}^d$ and $\Delta\alpha \in \mathbb{R}$ such that $b_i = \lambda a_i + \Delta a$ and $\beta_i = \lambda\alpha_i + \Delta\alpha$ for every i. Now fix $i \in [k]$.

Then

$$\begin{aligned}\left(\hat{s}^{(i)}, \hat{\mu}_i\right) &= \Phi^{-1}\left(\lambda a_i + \Delta a,\ \lambda\alpha_i + \Delta\alpha\right) \\ &\overset{(3.16)}{=} \Phi^{-1}\left(\left(\lambda(-2s^{(i)}) + \Delta a\right), \left(\lambda(\left\|s^{(i)}\right\|_2^2 + \mu_i) + \Delta\alpha\right)\right) \\ &\overset{(3.18)}{=} \left(\lambda s^{(i)} - \frac{1}{2}\Delta a,\ \lambda(\left\|s^{(i)}\right\|_2^2 + \mu_i) + \Delta\alpha - \frac{1}{4}\left\|-2\lambda s^{(i)} + \Delta a\right\|_2^2\right).\end{aligned}$$

With $\Delta s := -\frac{1}{2}\Delta a$ and $\Delta\mu := \Delta\alpha$ this yields the claim. □

Let us briefly discuss the necessity of the regularity assumptions in Theorem 3.8 and Corollary 3.9.

First, by a simple continuity argument it is clear that any empty cell must be excluded. Assume a power diagram given in affine representation, i. e., $\mathcal{P} = \mathrm{PD}_{\mathrm{aff}}\left((a_i, \alpha_i)_{i\in[k]}\right)$. If the ith cell is empty, this implies that its corresponding hyperplane $H_{\left(\binom{-a_i}{1}, \alpha_i\right)}$ in $\mathbb{R}^{d+1}$ can be strongly separated from the polyhedron $\overline{P}$ as defined in Eq. (3.19). Thus, any sufficiently small disturbance of (a_i, α_i) will not change $\overline{P}$ and hence the power diagram.

Let us consider two small examples explaining the necessity of a dimension greater than one, line-free cells and the regularity of vertices. For the sake of a better illustration, we consider those in the non-affine parametrization by spheres.

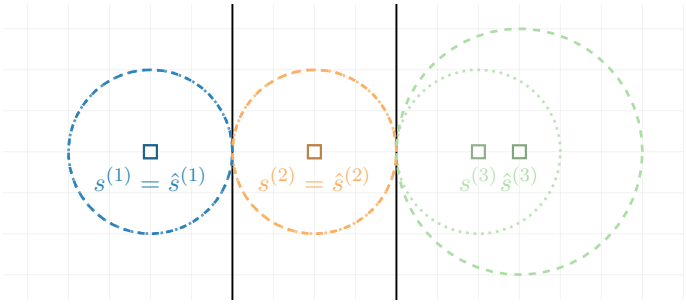

Figure 3.10: Illustration of Example 3.10.

Example 3.10

We consider a two-dimensional power diagram $\mathcal{P}$ of $k = 3$ cells and two parametrizations $\mathcal{P} = \mathrm{PD}\left(\left(s^{(i)}, \mu_i\right)_{i\in[k]}\right) = \mathrm{PD}\left(\left(\hat{s}^{(i)}, \hat{\mu}_i\right)_{i\in[k]}\right)$ with sites

$$s^{(1)} := \hat{s}^{(1)} := \begin{pmatrix}0\\0\end{pmatrix}, s^{(2)} := \hat{s}^{(2)} := \begin{pmatrix}2\\0\end{pmatrix}, \text{ and } s^{(3)} := \begin{pmatrix}4\\0\end{pmatrix}, \hat{s}^{(3)} := \begin{pmatrix}\frac{9}{2}\\0\end{pmatrix},$$

and radii σ_i (recall that $\mu_i = -\sigma_i^2$)

$$\sigma_1 := \tilde{\sigma}_1 := 1,\ \sigma_2 := \hat{\sigma}_2 := 1, \text{ and } \sigma_3 = 1, \hat{\sigma}_3 := \frac{3}{2}.$$

Figure 3.10 depicts the resulting diagram together with both generating sets of spheres.

As all sites are contained in the x_1-axis, it is clear that all bisectors are parallels of the x_2-axis and hence the diagram consists of vertical stripes. Also, from Fig. 3.10 one directly sees that $B_{1,2} = \{x \in \mathbb{R}^2 : x_1 = 1\}$ and $B_{2,3} = \{x \in \mathbb{R}^2 : x_1 = 3\}$, as the corresponding pairs of spheres touch each other in $(1, 0)^\intercal$ and $(3, 0)^\intercal$, respectively. This yields that the corresponding power functions have value 0 in those points.

Now $s^{(1)} = \hat{s}^{(1)}$ and $s^{(2)} = \hat{s}^{(2)}$ in Eq. (3.24) gives $\lambda = 1$ and $\Delta s = 0$. As $s^{(3)} \neq \hat{s}^{(3)}$, however, we see that the claim of Corollary 3.9 does not hold.

As all sites are collinear, this example implies the one-dimensional case by considering the intersection of the diagram with the x_1-axis. Although in this case the cells are obviously line-free, the proof of Theorem 3.8 fails as vertices are not connected by the pairwise intersection of cells.

Second, we consider a similarly simple example that reveals the necessity of the non-degeneracy assumption for vertices.

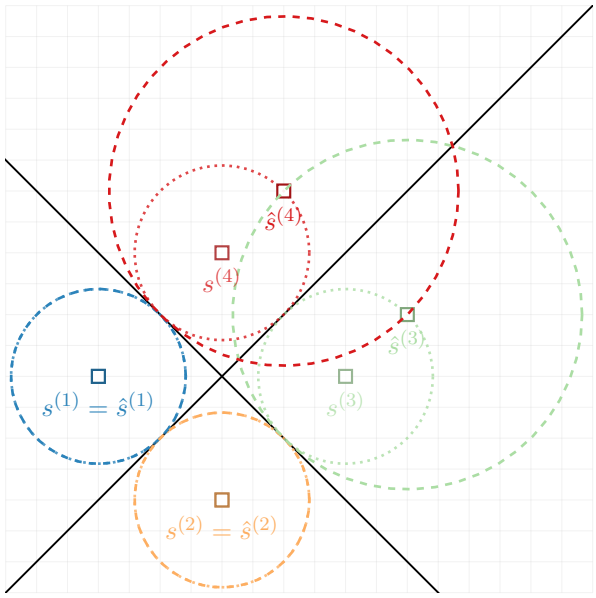

Figure 3.11: Illustration of Example 3.11.

Example 3.11

We consider a power diagram $\mathcal{P}$ in the plane with $k = 4$ cells given by two parametriza-

tions $\mathcal{P} = \mathrm{PD}\left(\left(s^{(i)}, \mu_i\right)_{i\in[k]}\right) = \mathrm{PD}\left(\left(\hat{s}^{(i)}, \hat{\mu}_i\right)_{i\in[k]}\right)$ with sites

$$s^{(1)} := \hat{s}^{(1)} := \begin{pmatrix} -2 \\ 0 \end{pmatrix}, \qquad s^{(2)} := \hat{s}^{(2)} := \begin{pmatrix} 0 \\ -2 \end{pmatrix},$$
$$s^{(3)} := \begin{pmatrix} 2 \\ 0 \end{pmatrix}, \hat{s}^{(3)} := \begin{pmatrix} 3 \\ 1 \end{pmatrix}, \qquad s^{(4)} := \begin{pmatrix} 0 \\ 2 \end{pmatrix}, \hat{s}^{(4)} := \begin{pmatrix} 1 \\ 3 \end{pmatrix},$$

and radii

$$\sigma_1 := \sigma_2 := \sigma_3 := \sigma_4 := \sqrt{2}, \text{ and } \hat{\sigma}_1 := \hat{\sigma}_2 = \sqrt{2}, \ \hat{\sigma}_3 := \hat{\sigma}_4 := 2\sqrt{2}.$$

Figure 3.11 illustrates the resulting diagram. Again, correctness of both parametrizations can be deferred from the intersection points of the spheres, that must be contained in the corresponding bisectors. In particular, $(0, 0)^\intercal$ is a vertex of all four cells and hence degenerated.

Once more, $s^{(1)} = \hat{s}^{(1)}$ and $s^{(2)} = \hat{s}^{(2)}$ would yield $\lambda = 1$ and $\Delta s = 0$ in Eq. (3.24). Hence, as $s^{(3)} \neq \hat{s}^{(3)}$, we again see that Eq. (3.24) does not hold for this degenerated case.

From the proof of Theorem 3.8 it is clear, however, that this assumption could be relaxed by only requiring that the graph that consists of all *regular* vertices and edges for incident $(d-1)$-faces of cells is connected.

3.3.3 Anisotropic Power Diagrams

Our next step in generalizing Voronoi diagrams is the introduction of individual norms for each of the cells. There are several practical motivations for doing so. One may be that the Euclidean distance does not adequately model the distances of a real world application and hence should be locally adjusted. In another situation, one might want to enforce a certain orientation of the cells and hence evaluate distances individually.

In the following, we will assume for each $i \in [k]$ the distance measure d_i to be the distance induced by an ellipsoidal norm. As a consequence, the distance of a point from a site depends on the direction, hence we are talking about *anisotropic* diagrams.

We set $\mathcal{X} := \mathbb{R}^d$ and define $f_i : \mathcal{X} \to \mathbb{R}$ via

$$f_i(x) := (x - s^{(i)})^\intercal \Sigma_i (x - s^{(i)}) + \mu_i \tag{3.26}$$

for symmetric matrices $\Sigma_i \in \mathbb{R}^{d\times d}_{\mathrm{sym}}$, sites $s^{(i)} \in \mathbb{R}^d$, and additive weights $\mu_i \in \mathbb{R}$, for every $i \in [k]$. We call the resulting generalized Voronoi diagram $\mathcal{P}$ w. r. t. $(f_i)_{i\in[k]}$ *anisotropic power diagram* and denote it $\mathcal{P} = \mathrm{APD}\left(\left(\Sigma_i, s^{(i)}, \mu_i\right)_{i\in[k]}\right)$. Figure 3.12 depicts the Voronoi surface for an exemplary diagram of 3 cells in the plane.

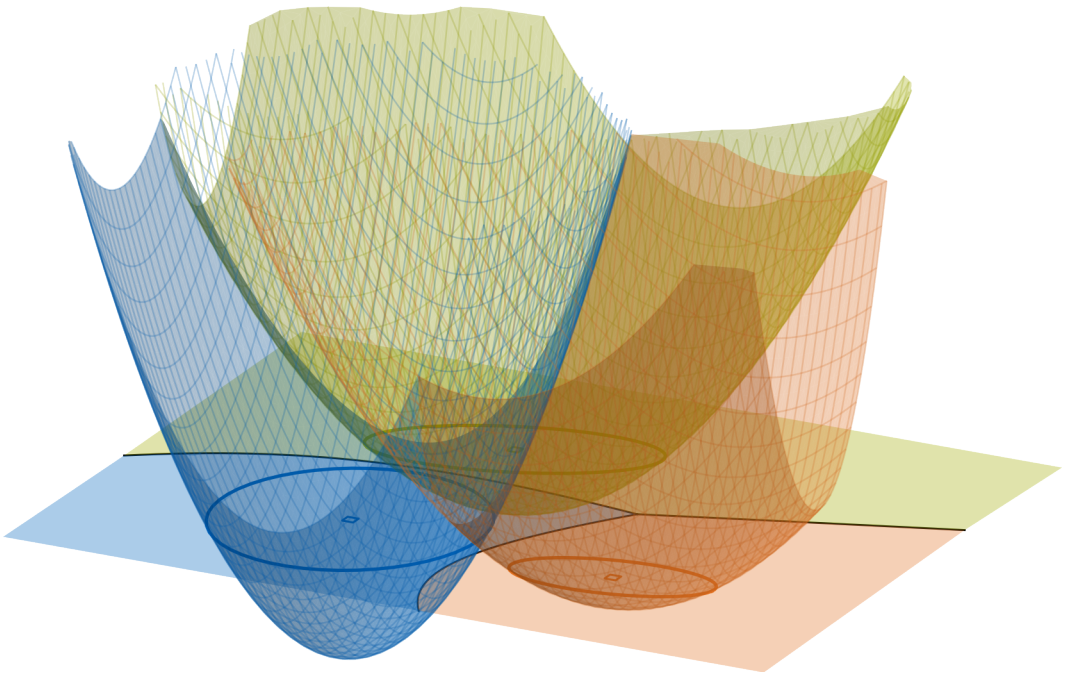

Figure 3.12: Voronoi surface of an anisotropic power diagram.

If Σ_i is positive definite and $\mu_i < 0$, we may again set $\sigma_i := \sqrt{-\mu_i}$ and obtain

$$f_i(x) := \left\| x - s^{(i)} \right\|_{\Sigma_i}^2 - \sigma_i^2. \tag{3.27}$$

Here $\|\cdot\|_{\Sigma_i}$ denotes the norm with the ellipsoidal unit ball

$$\mathbb{B}_{\Sigma_i} := \left\{ x \in \mathbb{R}^d : x^\intercal \Sigma_i x \leq 1 \right\} = \Sigma_i^{-\frac{1}{2}} \cdot \mathbb{B}_2. \tag{3.28}$$

We may then associate the ellipsoid

$$\mathbb{E}_i := s^{(i)} + \left\{ x \in \mathbb{R}^d : x^\intercal \Sigma_i x = \sigma_i^2 \right\} = s^{(i)} + \sigma_i \Sigma_i^{-\frac{1}{2}} \cdot \mathbb{S}^{d-1} \tag{3.29}$$

with the ith cell. We will see shortly that the additional assumptions of positive definite matrices and negative additive weights are not restrictive.

We call a point $v \in \mathbb{R}^d$ a *vertex* of the anisotropic power diagram if there exists an open set $U \subset \mathbb{R}^d$ and cell indices $I \subset [k]$ such that $\{v\} = \bigcap_{i \in I} P_i \cap U$. We say that a vertex is *regular* if it is the intersection of exactly $d+1$ cells (i. e., $|I| = d+1$ and I is uniquely determined).

For $\emptyset \neq \mathcal{I} \subset [k]$, we call $\mathcal{F} := \bigcap_{i \in \mathcal{I}} P_i$ a *n-face* of the diagram, if $\mathcal{F} \setminus \bigcup_{i \in [k] : \mathcal{F} \not\subset P_i} P_i$ is a (possibly disconnected) n-manifold (in $\mathbb{R}^d$). Note that in our definition faces do

not need to be connected. In particular, a 0-face may consist of several vertices, see Fig. 3.13 for an example.

We call the diagram *vertex-connected* if every cell contains at least one vertex and the graph with the diagram's vertices as nodes and with edges connecting nodes whose vertices belong to at least two common cells is connected.

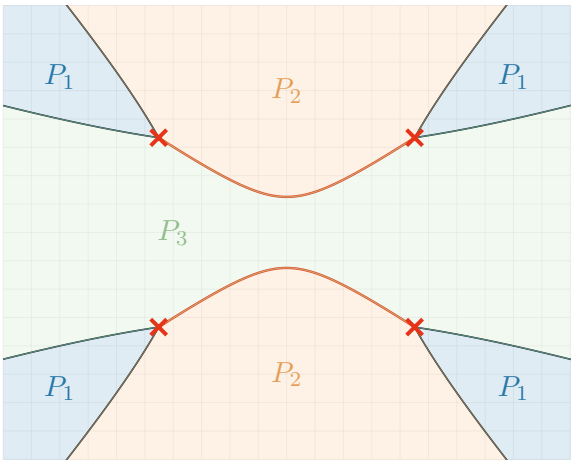

Figure 3.13: Examples of faces of an anisotropic power diagram. The red bold lines yield the 1-face $P_2 \cap P_3$. The points marked by crosses yield the 0-face $P_1 \cap P_2 \cap P_3$ which consists of four vertices.

Anisotropic power diagrams have appeared in the literature in various settings. Labelle and Shewchuk [LS03] introduced anisotropic Voronoi Diagrams, i. e., anisotropic power diagrams without additive weights, motivated by the application of mesh generation. Here, the anisotropy is used in order to compensate for stretched or elongated areas that suggest the usage of meshes with accordingly asymmetric triangles. In particular, they describe sufficient conditions for the positioning of sites in the plane such that the resulting cells are star-shaped. In the same context, Canas and Gortler [CG11] describe sufficient conditions w. r. t. the distance of sites such that anisotropic power diagrams provide connected cells. In [Alp+15] they are introduced as *generalized power diagrams* for the purpose of grain reconstruction. Boissonnat, Wormser, and Yvinec [BWY06] and Wormser [Wor08] discuss *curved* Voronoi diagrams and in particular describe the relation of power diagrams and anisotropic power diagrams (which will be discussed shortly) as well as different possible parametrizations. A special case of anisotropic power diagrams that we do not discuss in more detail are *Möbius diagrams* ([BK03; BWY06]). Those are obtained for the case that all ellipsoidal unit balls are in fact scaled Euclidean balls.

Quadratic Parametrization

Similar to power diagrams, it will be useful to introduce a second parametrization of anisotropic power diagrams. This time, we define the parameter transformation function

$$\Phi : \begin{cases} \mathbb{R}^{d\times d}_{\text{sym}} \times \mathbb{R}^d \times \mathbb{R} & \to \mathbb{R}^{d\times d}_{\text{sym}} \times \mathbb{R}^d \times \mathbb{R} \\ (\Sigma, s, \mu) & \mapsto \Phi(\Sigma, s, \mu) := (\Sigma, -2\Sigma s, s^\intercal \Sigma s + \mu). \end{cases} \tag{3.30}$$

For every $i \in [k]$ we set $(A_i, a_i, \alpha_i) := \Phi\left(\Sigma_i, s^{(i)}, \mu_i\right)$, and with

$$\hat{f}_i(x) := x^\intercal A_i x + a_i^\intercal x + \alpha_i \tag{3.31}$$

it trivially holds that $\hat{f}_i = f_i$ (with f_i as defined by Eq. (3.26)). So, the resulting generalized Voronoi diagrams are clearly identical.

Obviously, Φ is only a bijection when restricted to regular matrices in its first argument. However, given $\mathcal{P} = \text{APD}_{\text{quad}}\left((A_i, a_i, \alpha_i)_{i\in[k]}\right)$ we can always find $E \in \mathbb{R}^{d\times d}_{\text{sym}}$ such that $A_i + E$ is regular for every $i \in [k]$. However, then once more Lemma 3.2 yields that $\text{APD}_{\text{quad}}\left((A_i + E, a_i, \alpha_i)_{i\in[k]}\right) = \text{APD}_{\text{quad}}\left((A_i, a_i, \alpha_i)_{i\in[k]}\right)$ as the associated functions differ by a term $x^\intercal E x$. Hence, we can conclude that the class of all generalized Voronoi diagrams defined by functions of type (3.31) is identical to the class of anisotropic power diagrams (as also noted in [BWY06]).

We will refer to the parametrization (3.31) as *quadratic* and denote the resulting anisotropic power diagram $\mathcal{P} = \text{APD}_{\text{quad}}\left((A_i, a_i, \alpha_i)_{i\in[k]}\right)$. If not clear from the context, we will refer to the parametrization (3.27) as *ellipsoidal*.

Bisectors

In order to understand the characteristics of anisotropic power diagrams, let us consider a bisector $B_{i,l}$ for $i \neq l \in [k]$. It holds that $x \in B_{i,l}$ if and only if

$$\begin{aligned} &(x - s^{(i)})^\intercal \Sigma_i (x - s^{(i)}) + \mu_i = (x - s^{(l)})^\intercal \Sigma_l (x - s^{(l)}) + \mu_l \\ \Leftrightarrow &x^\intercal(\Sigma_i - \Sigma_l)x - 2(\Sigma_i s^{(i)} - \Sigma_l s^{(l)})^\intercal x + (s^{(i)})^\intercal \Sigma_i s^{(i)} + \mu_i - (s^{(l)})^\intercal \Sigma_l s^{(l)} - \mu_l = 0. \end{aligned} \tag{3.32}$$

The latter equation yields a polynomial of degree at most 2 and so $B_{i,l}$ is a (possibly degenerate) quadric hypersurface (see, for example, [Aud03, Chapter 6] for a general introduction).

From an application point of view, the ellipsoidal parametrization (3.27) might give a better understanding for the interdependency of the diagram cells. Hence, let us discuss how the relations of the norm-associated ellipsoids affect the characteristics of the cells.

In order to do so, we can normalize Eq. (3.32) in a common way. As $\Sigma_i - \Sigma_l \in \mathbb{R}^{d\times d}_{\text{sym}}$, there exists an orthonormal matrix $Q = \begin{pmatrix} Q_R & Q_N \end{pmatrix} \in \mathcal{SO}(d)$ with $R \subset [d]$ and $|R| = \operatorname{rank}(\Sigma_i - \Sigma_l)$, $N = [d] \setminus R$, and $\delta_r \in \mathbb{R} \setminus \{0\}$ for $r \in R$ such that with $D_R := \operatorname{diag}(\delta_1, \ldots, \delta_{|R|}) \in \mathbb{R}^{|R|\times|R|}$ it holds that $\Sigma_i - \Sigma_l = Q_R D_R Q_R^T$. We then obtain the pseudo inverse of $\Sigma_i - \Sigma_l$ as $(\Sigma_i - \Sigma_l)^+ = Q_R D_R^{-1} Q_R^\intercal$. With

$$c := (\Sigma_i - \Sigma_l)^+(\Sigma_i s^{(i)} - \Sigma_l s^{(l)}), \tag{3.33}$$

$$b_N := -2Q_N^\intercal(\Sigma_i s^{(i)} - \Sigma_l s^{(l)}), \text{ and} \tag{3.34}$$

$$\begin{aligned}\gamma := {} & (s^{(i)})^\intercal \Sigma_i s^{(i)} + \mu_i - (s^{(l)})^\intercal \Sigma_l s^{(l)} - \mu_l \\ & - (\Sigma_i s^{(i)} - \Sigma_l s^{(l)})^\intercal (\Sigma_i - \Sigma_l)^+ (\Sigma_i s^{(i)} - \Sigma_l s^{(l)})\end{aligned} \tag{3.35}$$

we can define normalized coordinates via

$$y := \begin{pmatrix} y_R \\ y_N \end{pmatrix} := \begin{pmatrix} Q_R^\intercal(x - c) \\ Q_N^\intercal x \end{pmatrix} \tag{3.36}$$

and get that Eq. (3.32) is equivalent to

$$\sum_{r\in R} \delta_r y_r^2 + b_N^\intercal y_N + \gamma = 0. \tag{3.37}$$

From Eq. (3.32) we already see that the corresponding quadratic form and hence the characterization of the quadric is determined by the difference of the norm-defining matrices $\Sigma_i - \Sigma_l$. From a rather geometric perspective, this means that a bisector's main characteristic can be deferred from the unit balls of the corresponding ellipsoidal norms. Let us collect some basic facts on that dependency in the following proposition.

Proposition 3.12
Consider the setting above. Then:

i) For $r \in R$ and $q^{(r)} := Q_{\cdot,r}$, it holds that $\left\|q^{(r)}\right\|_{\Sigma_i}^2 - \left\|q^{(r)}\right\|_{\Sigma_l}^2 = \delta_r$.

ii) It holds that $\delta_r > 0$ for all $r \in R$ if and only $\mathbb{B}_{\Sigma_i} \subset \mathbb{B}_{\Sigma_l}$. Furthermore, it holds that $\delta_r > 0$ for all $r \in R$ and $R = [d]$ if and only if $\mathbb{B}_{\Sigma_i} \subset \operatorname{int}(\mathbb{B}_{\Sigma_l})$.

iii) Let $U := \operatorname{lin}(\{Q_{\cdot,t} : t \in N\})$. Then $x \in U \setminus \{0\}$ if and only if $\mathbb{B}_{\Sigma_i}$ and $\mathbb{B}_{\Sigma_l}$ touch in $(\|x\|_{\Sigma_i})^{-1}x = (\|x\|_{\Sigma_l})^{-1}x$.

iv) It holds that $N = [d]$ if and only if $\mathbb{B}_{\Sigma_i} = \mathbb{B}_{\Sigma_l}$.

Proof. Part i follows directly from the definition of Q_R as $(\Sigma_i - \Sigma_r)q^{(r)} = \delta_r q^{(r)}$ and hence $\|q_r\|_{\Sigma_i}^2 - \left\|q^{(r)}\right\|_{\Sigma_l}^2 = (q^{(r)})^\intercal \Sigma_i q^{(r)} - (q^{(r)})^\intercal \Sigma_l q^{(r)} = \delta_r (q^{(r)})^\intercal q^{(r)} = \delta_r$.

Now the first statement of part ii is a direct consequence of part i, as $\mathbb{B}_{\Sigma_i} \subset \mathbb{B}_{\Sigma_l}$ is equivalent to $\|x\|_{\Sigma_i} \geq \|x\|_{\Sigma_l}$ for all $x \in \mathbb{R}^d$ and a symmetric matrix is positive semi-definite if and only if all eigenvalues are non-negative. As $\mathbb{B}_{\Sigma_i} \subset \mathbb{B}_{\Sigma_l}$ but $\mathbb{B}_{\Sigma_i} \not\subset \operatorname{int}(\mathbb{B}_{\Sigma_l})$ implies that there exists a point x with $\|x\|_{\Sigma_i} = \|x\|_{\Sigma_l}$ and thus that $\Sigma_i - \Sigma_l$ is not positive-definite, the second statement of part ii follows as well.

By definition, $\mathbb{B}_{\Sigma_i}$ and $\mathbb{B}_{\Sigma_l}$ touch in some point $\hat{x} \in \mathbb{R}^d$ if and only if $1 = \|\hat{x}\|_{\Sigma_i} = \|\hat{x}\|_{\Sigma_l}$ and $\Sigma_i\hat{x} = \Sigma_l\hat{x}$. The latter is equivalent to $\hat{x} \in \ker(\Sigma_i - \Sigma_l)$ and $\|\hat{x}\|_{\Sigma_i} = 1$. This yields part iii and immediately implies part iv. □

Note that if $N = [d]$ holds for all pairs of cells, i. e., if all ellipsoidal norms are identical, the diagram is in fact a power diagram.

Exemplary Cases in the Plane

Let us illustrate the observations above for the simple case of $k = 2$ cells in the plane. In particular, we aim to get an understanding of how the choice of parameters in the ellipsoidal parametrization determines the cells' characteristics.

First, we consider the case that neither of the two ellipsoidal unit balls are contained in one another. An example is illustrated in Fig. 3.14. This case implies $R = \{1, 2\}$ and $\delta_1 \cdot \delta_2 < 0$. Consequently, we see from Eq. (3.37) that the bisector $B_{1,2}$ is a pair of hyperbolas if $\gamma \neq 0$ and a pair of crossing lines otherwise. Up to translation, the asymptotes of the hyperbolas are solely determined by the choices of Σ_1 and Σ_2. For the diagram in Fig. 3.14b, this results in the cell P_1 (blue) to consist of two connected components. In order to grasp the impact of changing sites, assume all parameters to be fixed except the site $s^{(2)}$. The bisector is point symmetric w. r. t. c which is affine in $s^{(2)}$. We observe that γ in dependency of the site $s^{(2)}$ is yet another quadratic function. In Fig. 3.14, the red-dashed line depicts one of the pair of hyperbolas that yields the choices of $s^{(2)}$ such that $\gamma = 0$. In particular, when $s^{(2)}$ crosses this line, the cells swap their connectivity property (Figs. 3.14c and 3.14d). Finally, as γ changes linearly in the additive parameters μ_1 and μ_2, they determine the connectivity of the cells, too.

Next, we consider the case that one ellipsoidal unit ball is strictly contained in the other, i. e., we assume $\mathbb{B}_{\Sigma_2} \subset \operatorname{int}(\mathbb{B}_{\Sigma_1})$. This means we are in the setting of part ii of Proposition 3.12. Consequently, the bisector $B_{1,2}$ is either an ellipse or empty, depending on the sign of γ. Figure 3.15 illustrates this situation. Again, the choice of Σ_1 and Σ_2 already determines both orientation and ratio of the axes of this ellipse. The choice of the sites then again determines its center c as given by Eq. (3.33). Again, we may interpret γ as a quadratic function in $s^{(2)}$. This time, the set of sites such that $\gamma = 0$ yields yet another ellipse[1] (or may be empty). This is again drawn as red dashed line in Fig. 3.15b. Any choice of $s^{(2)}$ inside this ellipse leads to an empty cell P_2. Due

[1] In order to see this, note that by elementary linear algebra it holds that $\Sigma_1 \succ 0$, $\Sigma_2 \succ 0$ and $\Sigma_1 - \Sigma_2 \succ 0$ imply $\Sigma_2 - \Sigma_2(\Sigma_1 - \Sigma_2)^{-1}\Sigma_2 \succ 0$.

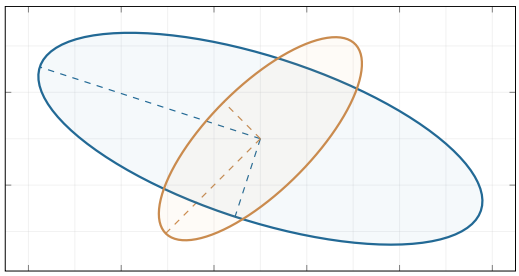

(a) Ellipsoidal unit balls $\mathbb{B}_{\Sigma_1}$ (blue) and $\mathbb{B}_{\Sigma_2}$ (orange).

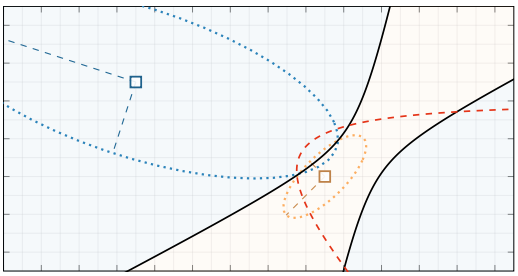

(b) Diagram with connected cell P_2.

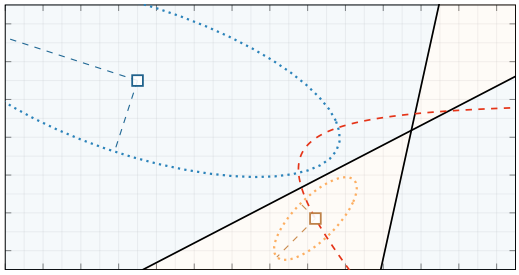

(c) Degenerated case of crossing lines.

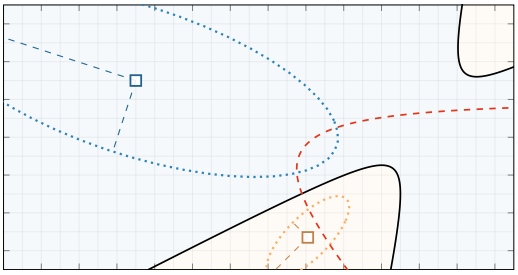

(d) Diagram with connected cell P_1.

Figure 3.14: Anisotropic power diagram of two cells in the plane, with no ellipsoidal unit ball contained in the other. The dotted ellipses in Figs. 3.14b to 3.14d yield the associated ellipsoids $\mathbb{E}_1$ and $\mathbb{E}_2$ (Eq. (3.29)). The red dashed line illustrates the set $\gamma = 0$ when interpreting γ as a function in $s^{(2)}$.

to the linear term $\mu_2 - \mu_1$ in the definition of γ, we can conclude that for $\mu_1 - \mu_2$ large, the ellipse $B_{1,2}$ and thus the cell P_2 are scaled by an asymptotic factor of $\sqrt{\mu_1 - \mu_2}$.

As final case, we assume that $\mathbb{B}_{\Sigma_2} \subsetneq \mathbb{B}_{\Sigma_1}$ but $\mathbb{B}_{\Sigma_1} \cap \mathbb{B}_{\Sigma_2} \neq \emptyset$, i. e., the two unit balls touch, as depicted in Fig. 3.16a. This means $\operatorname{rank}(\Sigma_1 - \Sigma_2) = 1$. In particular, Fig. 3.16a furthermore depicts the null space U as used in part iii of Proposition 3.12. Consequently, the resulting normalized form of Eq. (3.37) contains a linear term y_N. Again, assume all parameters to be fixed except $s^{(2)}$. This time, the red dashed lines in Figs. 3.16b to 3.16d mark the choices of $s^{(2)}$ that result in $b_N = 0$. As

$$b_N = -2Q_N^\intercal(\Sigma_1 s^{(1)} - \Sigma_2 s^{(2)}) = -2\underbrace{Q_N^\intercal(\Sigma_1 - \Sigma_2)}_{=0} s^{(1)} - 2Q_N^\intercal\Sigma_2(s^{(1)} - s^{(2)}),$$

this yields a line through $s^{(1)}$. Which side of this line $s^{(2)}$ is placed then determines the orientation of the parabola that is $B_{1,2}$. In the degenerated case that $s^{(2)}$ lies on this line, the bisector consists of two parallel lines and the cell P_1 is disconnected (cf. Figs. 3.16b to 3.16d). With respect to the additive parameters, the term $\mu_1 - \mu_2$ determines the vertical position along the $Q_R = (q_1)$ direction in the non-degenerated

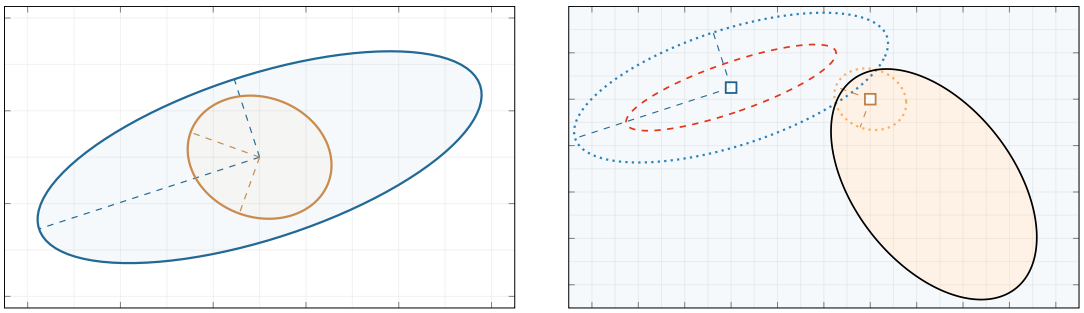

(a) Ellipsoidal unit balls $\mathbb{B}_{\Sigma_1}$ (blue) and $\mathbb{B}_{\Sigma_2}$ (orange).

(b) Resulting anisotropic power diagram with ellipsoidal bisector. The dotted ellipses mark $\mathbb{E}_i$, $i = 1, 2$ (Eq. (3.29)). The red-dashed line marks the choices of $s^{(2)}$ resulting in $\gamma = 0$.

Figure 3.15: Anisotropic power diagram of two cells in the plane, with one ellipsoidal unit ball strictly contained in the other, i. e., $\mathbb{B}_{\Sigma_2} \subset \text{int}(\mathbb{B}_{\Sigma_1})$.

case. In the degenerated case, it determines the distance of the two parallel lines (which might be smaller than 0 causing P_2 to be empty).

Parametrization Invariances

As we have already seen for the exemplary cases in the plane, the choice of parameters for an anisotropic power diagram is not unique. By stating a few lemmas, we prepare a theorem that will narrow down the invariances in the parametrization of anisotropic power diagrams. First, regular vertices guarantee the affine independence of the parameters of the associated cells.

Lemma 3.13
Let $\mathcal{P} = (P_1, \dots, P_k) = \text{APD}_{\text{quad}}\left((A_i, a_i, \alpha_i)_{i\in[k]}\right)$ *be an anisotropic power diagram and* $v \in \mathbb{R}^d$ *be a regular vertex. Let* $I \subset [k]$ *such that* $\{v\} = \bigcap_{i\in I} P_i$. *Then* (A_i, a_i, α_i) *for* $i \in I$ *are affinely independent (when interpreted as vectors in* $\mathbb{R}^{d^2+d+1}$*).*

Proof. As v is regular, we may w.l.o.g. assume $I = [d+1]$. Next, we may furthermore w.l.o.g. assume $(A_{d+1}, a_{d+1}, \alpha_{d+1}) = 0$ (otherwise we set $\left(\hat{A}_i, \hat{a}_i, \hat{\alpha}_i\right) := (A_i, a_i, \alpha_i) - \left(\hat{A}_{d+1}, \hat{a}_{d+1}, \hat{\alpha}_{d+1}\right)$ for $i \in [d+1]$ and obtain the same diagram due to Lemma 3.2, which, of course, does not affect affine independence). Thus, we need to show the linear independence of the first d vectors. Assume this is not the case and once more w.l.o.g. that there are $\lambda_i \in \mathbb{R}$ for $i \in [d-1]$ such that $(A_d, a_d, \alpha_d) = \sum_{i\in[d-1]} \lambda_i (A_i, a_i, \alpha_i)$.

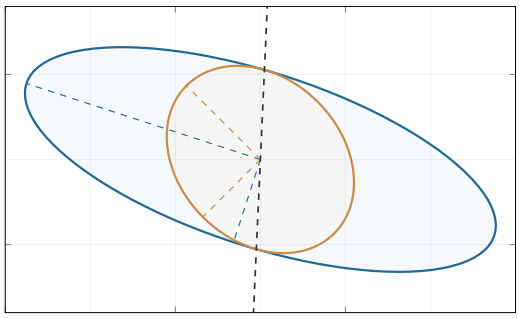

(a) Ellipsoidal unit balls $\mathbb{B}_{\Sigma_1}$ (blue) and $\mathbb{B}_{\Sigma_2}$ (orange). The gray dashed line marks the subspace U as defined in part iii of Proposition 3.12.

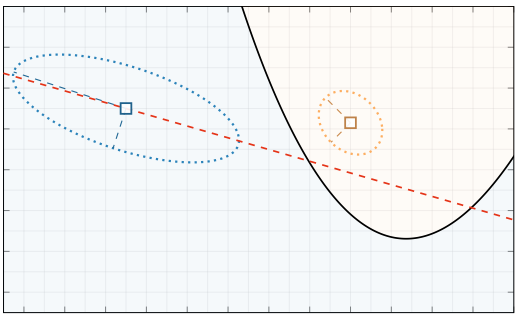

(b) Diagram with a parabola as bisector.

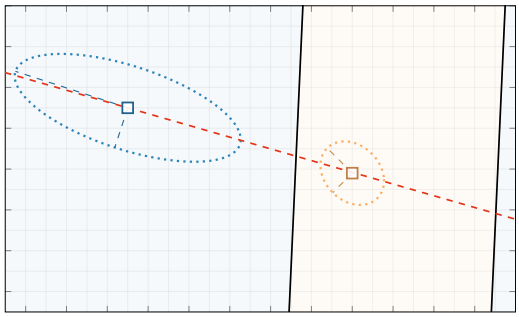

(c) Diagram with the bisector consisting of two parallel lines.

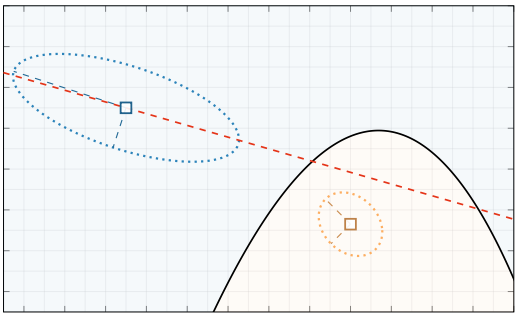

(d) Diagram with a parabola as bisector.

Figure 3.16: Anisotropic power diagram of two cells in the plane, with the two distinct ellipsoidal unit balls touching each other. Dotted ellipses mark $\mathbb{E}_1$ (blue) and $\mathbb{E}_2$ (orange) as defined by Eq. (3.29). The red dashed line marks the choices for $s^{(2)}$ such that $b_N = 0$ (cf. Eq. (3.34)).

Then

$$\bigcap_{i\in[d-1]} B_{i,d+1} = \left\{x \in \mathbb{R}^d : x^\intercal A_i x + a_i^\intercal x + \alpha_i = 0 \ \forall i \in [d-1]\right\}$$
$$\subset \left\{x \in \mathbb{R}^d : x^\intercal \Big(\sum_{i\in[d-1]} \lambda_i A_i\Big)x + \sum_{i\in[d-1]} \lambda_i a_i{}^\intercal x + \sum_{i\in[d-1]} \lambda_i \alpha_i = 0\right\} = B_{d,d+1}$$

This immediately gives $\bigcap_{i\in[d-1]} P_i \cap P_{d+1} \subset B_{d,d+1} \cap P_{d+1}$. Now observe that $B_{d,d+1} \cap P_{d+1} = P_d \cap P_{d+1}$. ("$\supset$" is clear by definition. Then note that $x \in B_{d,d+1} \cap P_{d+1}$ implies $f_d(x) = f_{d+1}(x) = \min_{i\in[k]} f_i(x)$ and hence $x \in P_d$).

Thus, we get $\bigcap_{i\in[d-1]} P_i \cap P_{d+1} \subset P_d \cap P_{d+1}$ and hence $\bigcap_{i\in[d+1]} P_i = \bigcap_{i\in[d-1]} P_i \cap P_{d+1} = \{v\}$, which contradicts the regularity of v. □

Furthermore, in order to determine the degrees of freedom in the parametrization of anisotropic power diagrams, we require a few technical results on real quadrics. In particular, we need a little terminology from algebraic geometry concerning quadrics or – more generally – affine varieties. Here, we follow the terminology of [CLO15]. For polynomials $p_1, \dots, p_n \in \mathbb{R}[x_1, \dots, x_d]$ the set

$$\mathcal{V}(p_1, \dots, p_n) := \left\{ x \in \mathbb{R}^d : p_i(x) = 0 \ \forall i \in [n] \right\}$$

is called an *affine variety*.

We will have to exclude degenerated cases of anisotropic power diagrams that yield bisectors which are (locally) contained in hyperplanes. We will call a set $A \subset \mathbb{R}^d$ *locally flat* if there exists an open set $U \subset \mathbb{R}^d$ and a hyperplane $H \subset \mathbb{R}^d$ such that $A \cap U = H \cap U \neq \emptyset$.

The next two lemmas are of fundamental nature and could be established using Hilbert's Nullstellensatz ([CLO15, Chapter 4, Theorem 2]) over an algebraically closed field. However, as we only consider the non-algebraically closed field $\mathbb{R}$ and do not intend to make use of deeper theoretical algebraic geometry in our context, we provide fundamental proofs on our own.

By the first lemma, we will ensure that whenever bisectors are not locally flat, the corresponding polynomials are not reducible.

Lemma 3.14
Let $p \in \mathbb{R}[x_1, \dots, x_d]$ be a polynomial with $deg(()p) = 2$. Then p is reducible if and only if $\mathcal{V}(p)$ is locally flat.

Proof. If p is reducible, it is the product of two affine functions. Thus, $\mathcal{V}(p)$ is the union of two hyperplanes, which is locally flat.

If p is locally flat, there exists a hyperplane $H \subset \mathbb{R}^d$ and an open set $U \subset \mathbb{R}^d$ such that $\emptyset \neq \mathcal{V}(p) \cap U = H \cap U$. We can w.l.o.g. assume that $H = \left\{ x \in \mathbb{R}^d : x_1 = 0 \right\}$ (otherwise, we transform coordinates). Let $g \in \mathbb{R}[x_1, \dots, x_d]$ and $h \in \mathbb{R}[x_1, \dots, x_{d-1}]$ such that $p(x) = x_1 g(x) + h(\overline{x})$ for $x =: (x_1, \overline{x}^\intercal)^\intercal \in \mathbb{R}^d$. As $p|_{H \cap U} = 0$, it follows that h is zero on the projection of $H \cap U$ onto the first $d-1$ coordinates $\overline{U} := \pi_{[d-1]}(H \cap U) = \pi_{[d-1]}(U)$. Since this is an open set in $\mathbb{R}^{d-1}$, it follows that $h = 0$. Thus, p is reducible. □

The next lemma states that the quadratic polynomial that is given by the bisector equation (3.32) is already uniquely determined up to scaling as soon as some small excerpt of the resulting bisector is known and the bisector is non-degenerated, i. e., not the union of two hyperplanes. We mention that a version of the following statement for general complex varieties can be found in [Ken77, Theorem 2.11].

Lemma 3.15
Let $d \geq 2$, $p, \tilde{p} \in \mathbb{R}[x_1, \ldots, x_d]$, p irreducible with $deg(()p) = 2$, $deg(()\tilde{p}) \leq 2$, and $U \subset \mathbb{R}^d$ open such that $\mathcal{V}(p) \cap U$ is $(d-1)$-manifold in $\mathbb{R}^d$ and

$$\mathcal{V}(p) \cap U = \mathcal{V}(\tilde{p}) \cap U.$$

Then there exists $\lambda \neq 0$ such that $\tilde{p} = \lambda p$.

Proof. We may w.l.o.g. assume that p is in normal form (otherwise we may consider an affine transformation of $\mathbb{R}^d$). This means, that for a suitable $r \in [d]$ we set $R := [r]$ and $N := [d] \setminus R$ and there are $\delta_i \in \mathbb{R} \setminus \{0\}$ for $i \in R$, $a_N \in \mathbb{R}^{|N|}$ and $\alpha \in \mathbb{R}$ such that

$$p(x) = \sum_{i \in [r]} \delta_i x_i^2 + a_N^\intercal x_N + \alpha.$$

Here, we denote $x = (x_R^\intercal, x_N^\intercal)^\intercal \in \mathbb{R}^d$. Similarly, we will denote $x = (x_1, \overline{x}^\intercal)^\intercal$ with $\overline{x} \in \mathbb{R}^{d-1}$.

As we are only interested in the equality of p and $\tilde{p}$ up to scaling, we may w.l.o.g. assume $\delta_1 = -1$ (otherwise consider $-\frac{1}{\delta_1} p$) and define $q \in \mathbb{R}[x_2, \ldots, x_d]$ via $q(\overline{x}) := \sum_{2 \leq i \leq r} \delta_i x_i^2 + a_N^\intercal x_N + \alpha$.

For all $x \in \mathcal{V}(p) \cap U$ it then holds that

$$x_1^2 = q(\overline{x}).$$

As $\mathcal{V}(p) \cap U$ is $(d-1)$-manifold and p is irreducible, we can conclude from Lemma 3.14 that $U \cap \mathcal{V}(p)$ is full-dimensional. Hence, there exists $x^{(0)} \in U \cap \mathcal{V}(p)$ such that $(x^{(0)})_1 \neq 0$. Furthermore, we may assume that $\operatorname{sign}(x_1) = \operatorname{sign}(x_1^{(0)})$ as well as $q(\overline{x}) > 0$ for all $x = (x_1, \overline{x}^\intercal)^\intercal \in U$.

Consider the projection

$$\overline{U} := \pi_{[d]\setminus\{1\}}(U \cap \mathcal{V}(p)) := \left\{ \overline{x} \in \mathbb{R}^{d-1} : \exists x_1 \in \mathbb{R} : (x_1, \overline{x}^\intercal)^\intercal \in U \cap \mathcal{V}(p) \right\} \subset \mathbb{R}^{d-1}.$$

Then

$$U \cap \mathcal{V}(p) = \left\{ \left(\operatorname{sign}(x_1^{(0)}) \cdot \sqrt{q(\overline{x})},\ \overline{x}^\intercal \right)^\intercal : \overline{x} \in \overline{U} \right\}.$$

In particular, we can assume that $\overline{U}$ is open (in $\mathbb{R}^{d-1}$) as $\pi_{[d]\setminus\{1\}}(U)$ must contain a neighborhood of $\overline{x}^{(0)}$ in which $\left(\operatorname{sign}(x_1^{(0)}) \cdot \sqrt{q(\overline{x})}, \overline{x}^\intercal \right)^\intercal \in U$.

Now, let $\lambda \in \mathbb{R}$ and $g, h \in \mathbb{R}[x_1, \ldots, x_{d-1}]$ with $deg(()g) \leq 1$, $deg(()h) \leq 2$ such that

$$\tilde{p}(x) = \lambda x_1^2 + x_1 g(\overline{x}) + h(\overline{x})$$

for all $x \in \mathbb{R}^d$.

As $U \cap \mathcal{V}(p) = U \cap \mathcal{V}(\tilde{p})$ it follows for all $\overline{x} \in \overline{U}$ that

$$\begin{aligned} 0 &= \lambda q(\overline{x}) + \operatorname{sign}(x_1^{(0)}) \cdot \sqrt{q(\overline{x})} \cdot g(\overline{x}) + h(\overline{x}) \\ \Rightarrow & q(\overline{x}) g(\overline{x})^2 = (\lambda q(\overline{x}) + h(\overline{x}))^2. \end{aligned}$$

As $\overline{U}$ is open and a polynomial is uniquely defined by its restriction to any open set, this implies

$$q \cdot g^2 = (\lambda q + h)^2.$$

If $g \neq 0$, then the latter implies $g|(\lambda f + h)$ and thus with $t := (\lambda q + h)/g \in \mathbb{R}[x_1, \ldots, x_{d-1}]$ it holds that $q = t^2$. However, then $p(x) = t(\overline{x})^2 - x_1^2 = (t(\overline{x}) - x_1)(t(\overline{x}) + x_1)$, so p is reducible, a contradiction. Hence, we obtain $g = 0$ and thus

$$\lambda q = -h.$$

If $\lambda = 0$, it follows that $h = 0$ and hence $\tilde{p} = 0$, a contradiction to $U \cap \mathcal{V}(\tilde{p}) = U \cap \mathcal{V}(p) \neq U$ as $p \neq 0$. We conclude $q = -\frac{1}{\lambda} h$ and so for any $x \in \mathbb{R}^d$

$$\tilde{p}(x) = \lambda x_1^2 + h(\overline{x}) = -\lambda(-x_1^2 - \frac{1}{\lambda} h(\overline{x})) = -\lambda(-x_1^2 + q(\overline{x})) = -\lambda p(x). \qquad \square$$

By the help of Lemma 3.15 we can now state and prove the analogue of Theorem 3.8 for the quadratic parametrization of anisotropic power diagrams.

Theorem 3.16

Let $d \geq 2$ and $\mathcal{P} = (P_1, \ldots, P_k)$ with $k \geq d+1$ be an anisotropic power diagram, and let $A_i \in \mathbb{R}^{d \times d}_{\text{sym}}, a_i \in \mathbb{R}^d, \alpha_i \in \mathbb{R}$ for $i \in [k]$ such that $\mathcal{P} = \text{APD}_{\text{quad}}\left((A_i, a_i, \alpha_i)_{i \in [k]}\right)$. Let furthermore $B_i \in \mathbb{R}^{d \times d}_{\text{sym}}, b_i \in \mathbb{R}^d, \beta_i \in \mathbb{R}$ for $i \in [k]$.

If there exist $\lambda > 0$ and $\Delta A \in \mathbb{R}^{d \times d}_{\text{sym}}, \Delta a \in \mathbb{R}^d, \Delta\alpha \in \mathbb{R}$ such that

$$(B_i, b_i, \beta_i) = \lambda\,(A_i, a_i, \alpha_i) + (\Delta A, \Delta a, \Delta\alpha)$$

for every $i \in [k]$, then $\text{APD}_{\text{quad}}\left((A_i, a_i, \alpha_i)_{i \in [k]}\right) = \text{APD}_{\text{quad}}\left((B_i, b_i, \beta_i)_{i \in [k]}\right)$.

If for any $i \neq l \in [k]$ with $P_i \cap P_l \neq \emptyset$ it holds that $P_i \cap P_l$ is a non locally flat $(d-1)$-face, every vertex is regular and the diagram vertex-connected, then the reverse is also true.

Proof. Once more, the first implication immediately follows from Lemma 3.2 by setting $a(x) := \lambda$ and $b(x) := x^\intercal \Delta A x + \Delta a^\intercal x + \Delta\alpha$ for $x \in \mathbb{R}^d$.

Now assume $\mathcal{P} = \mathrm{APD}_{\mathrm{quad}}\left((A_i, a_i, \alpha_i)_{i\in[k]}\right) = \mathrm{APD}_{\mathrm{quad}}\left((B_i, b_i, \beta_i)_{i\in[k]}\right)$ and that $\mathrm{int}(P_i) \neq \emptyset$ for every $i \in [k]$ and $P_i \cap P_l$ is not locally flat for any $i \neq l \in [k]$ with $P_i \cap P_l \neq \emptyset$.

Due to the first implication, we can w.l.o.g. assume that $(A_1, a_1, \alpha_1) = 0$.

Define the polynomials $p_i, q_i \in \mathbb{R}[x_1, \ldots, x_d]$ for $i \in [k]$ via $p_i(x) := x^\intercal A_i x + a_i{}^\intercal x + \alpha_i$ and $q_i(x) := x^\intercal B_i x + b_i{}^\intercal x + \beta_i$.

By assumption, P_1 contains a vertex v. Due to regularity, we may assume that $\{v\} = \bigcap_{i\in[d+1]} P_i \cap U_v$ for some open set $U_v \subset \mathbb{R}^d$. Also by regularity, it holds that $v \notin P_i$ for any $i \in [k] \setminus [d+1]$ and hence we may choose U_v sufficiently small such that $U_v \cap \bigcup_{i\in[k]\setminus[d+1]} P_i = \emptyset$.

Let $i \neq l \in [d+1]$. By assumption, $P_i \cap P_l$ is a $(d-1)$-face of the given diagram. Furthermore, we claim that there is no $r \in [d+1]$ such that $P_i \cap P_l \cap U_v \subset P_r$. If that was the case, then $\{v\} = \bigcap_{i\in[d+1]} P_i \cap U_v \supset \bigcap_{i\in[d+1]\setminus\{r\}} P_i \cap U_v$, which would contradict the regularity of v. Hence, the definition of a $(d-1)$-face implies that $(P_i \cap P_l \cap U_v) \setminus \bigcup_{r\in[k]\setminus\{i,l\}} P_r$ is a $(d-1)$-manifold and in particular non-empty. Hence, there exists $x \in P_i \cap P_l \setminus \bigcup_{r\in[k]\setminus\{i,l\}} P_r$. Thus, we can find an open neighborhood U_x of x such that $P_i \cap P_l \cap U_x = B_{i,l} \cap U_x = \mathcal{V}(p_i - p_l) \cap U_x = \mathcal{V}(q_i - q_l) \cap U_x$. As $P_i \cap P_l$ is furthermore non-flat, Lemma 3.14 implies that $p_i - p_l$ is irreducible. Then Lemma 3.15 states that there exists a $\lambda_{i,l} \neq 0$ such that

$$p_i - p_l = \lambda_{i,l}(q_i - q_l). \tag{3.38}$$

Next, we claim that $\lambda_{i,l} > 0$. If $\lambda_{i,l} < 0$, then $p_i(x) - p_l(x) < 0$ implies $q_i(x) - q_l(x) > 0$, which implies an empty symmetric difference of the cells, i. e., $P_i \Delta P_l = \emptyset$. However, this means $P_i = P_l$, which again contradicts the regularity of v.

As Eq. (3.38) holds for arbitrary $i \neq l \in [d+1]$, we can proceed as in the proof of Theorem 3.8 and consider the sum

$$\begin{aligned} 0 &= \sum_{i=1}^{d}(p_i - p_{i+1}) + (p_{d+1} - p_1) = \sum_{i=1}^{d} \lambda_{i,i+1}(q_i - q_{i+1}) + \lambda_{d+1,1}(q_{d+1} - q_1) \\ &= (\lambda_{1,2} - \lambda_{d+1,1})q_1 + \sum_{i=2}^{d+1}(\lambda_{i,i+1} - \lambda_{i-1,i})q_i + (\lambda_{d+1,1} - \lambda_{d,d+1})q_{d+1}. \end{aligned}$$

Due to the regularity of v, Lemma 3.13 states that the coefficient vectors of the polynomials q_i for $i \in [d+1]$ are affinely independent. Observe that the coefficients of the polynomials in the last equation sum up to zero, hence affine independence implies

$$\lambda_{1,2} = \lambda_{2,3} = \ldots = \lambda_{d+1,1} =: \lambda.$$

As we may reorder cells, this gives $\lambda = \lambda_{i,l}$ for all $i \neq l \in [d+1]$. We then define

$$(\Delta A, \Delta a, \Delta\alpha) := (B_1, b_1, \beta_1) - \lambda\,(A_1, a_1, \alpha_1)$$

and can (analogously to the proof of Theorem 3.8) verify that Eq. (3.38) then yields

$$(B_i, b_i, \beta_i) = \lambda\,(A_i, a_i, \alpha_i) + (\Delta A, \Delta a, \Delta\alpha) \tag{3.39}$$

for all $i \in [d+1]$. As the diagram is vertex-connected and each pair of connected vertices shares two common cells, we can conclude that Eq. (3.39) holds for all $i \in [k]$.□

As we did for power diagrams, we obtain the invariances for the parametrization of anisotropic power diagrams by ellipsoids as a corollary.

Corollary 3.17
Let $d \geq 2$ and $\mathcal{P} = (P_1, \ldots, P_k)$ with $k \geq d+1$ be an anisotropic power diagram, and let $\Sigma_i \in \mathbb{R}^{d\times d}_{\mathrm{sym},>0}$, $s^{(i)} \in \mathbb{R}^d$, $\mu_i \in \mathbb{R}$ for $i \in [k]$ such that $\mathcal{P} = \mathrm{APD}\left(\left(\Sigma_i, s^{(i)}, \mu_i\right)_{i\in[k]}\right)$. Let furthermore $\hat{\Sigma}_i \in \mathbb{R}^{d\times d}_{\mathrm{sym},>0}$, $\hat{s}^{(i)} \in \mathbb{R}^d$, $\hat{\mu}_i \in \mathbb{R}$ for $i \in [k]$.

If there exist $\lambda > 0$ and $\Delta\Sigma \in \mathbb{R}^{d\times d}_{\mathrm{sym}}$, $\Delta s \in \mathbb{R}^d$, and $\Delta\mu \in \mathbb{R}$, such that

$$\hat{\Sigma}_i = \lambda\Sigma_i + \Delta\Sigma \tag{3.40}$$

$$\hat{s}^{(i)} = \lambda\left(\lambda\Sigma_i + \Delta\Sigma\right)^{-1}\Sigma_i\left(s^{(i)} + \Sigma_i^{-1}\Delta s\right) \tag{3.41}$$

$$\hat{\mu}_i = \lambda\mu_i + \lambda\left(s^{(i)\intercal}\Sigma_i s^{(i)} - (\Sigma_i s^{(i)} + \Delta s)^\intercal\lambda\left(\lambda\Sigma_i + \Delta\Sigma\right)^{-1}(\Sigma_i s^{(i)} + \Delta s)\right) + \Delta\mu \tag{3.42}$$

for every $i \in [k]$, then $\mathrm{APD}\left(\left(\Sigma_i, s^{(i)}, \mu_i\right)_{i\in[k]}\right) = \mathrm{APD}\left(\left(\hat{\Sigma}_i, \hat{s}^{(i)}, \hat{\mu}_i\right)_{i\in[k]}\right)$.

If for any $i \neq l \in [k]$ with $P_i \cap P_l \neq \emptyset$ it holds that $P_i \cap P_l$ is a non locally flat $(d-1)$-face, every vertex is regular and the diagram vertex-connected, then the reverse is also true.

Proof. With the parameter transformation function Φ as defined in Eq. (3.30) we get

$$\Phi(\Sigma_i, s^{(i)}, \mu_i) = \left(\Sigma_i, -2\Sigma_i s^{(i)}, s^{(i)\intercal}\Sigma_i s^{(i)} + \mu_i\right).$$

For the first implication, let $\hat{\Sigma}_i, \hat{s}^{(i)}, \hat{\mu}_i$ for $i \in [k]$ be given as in Eqs. (3.40) to (3.42). We can reformulate Eqs. (3.41) and (3.42) as

$$\hat{s}^{(i)} = \lambda\hat{\Sigma}_i^{-1}(\Sigma_i s^{(i)} + \Delta s) \qquad \text{and}$$

$$\hat{\mu}_i = \lambda\left(\mu_i + (s^{(i)})^\intercal\Sigma_i s^{(i)}\right) - (\hat{s}^{(i)})^\intercal\hat{\Sigma}_i\hat{s}^{(i)} + \Delta\mu.$$

With $(\Delta A, \Delta a, \Delta\alpha) := (\Delta\Sigma, -2\lambda\Delta s, \Delta\mu)$ we then get for all $i \in [k]$ that

$$\begin{aligned}\Phi(\hat{\Sigma}_i, \hat{s}^{(i)}, \tilde{\mu}_i) &= \left(\hat{\Sigma}_i,\ -2\hat{\Sigma}_i\hat{s}^{(i)},\ (\hat{s}^{(i)})^\intercal\hat{\Sigma}_i\hat{s}^{(i)} + \hat{\mu}_i\right)\\ &= \left(\lambda\Sigma_i + \Delta\Sigma,\ -2\lambda(\Sigma_i + \Delta s),\ \lambda\left(\mu_i + (s^{(i)})^\intercal\Sigma_i s^{(i)}\right) + \Delta\mu\right)\\ &= \lambda\Phi(\Sigma_i, s^{(i)}, \mu_i) + (\Delta A, \Delta a, \Delta\alpha).\end{aligned}$$

Hence, we obtain from Theorem 3.16 that

$$\begin{aligned}\mathrm{APD}\left(\left(\Sigma_i, s^{(i)}, \mu_i\right)_{i\in[k]}\right) &= \mathrm{APD}_{\mathrm{quad}}\left(\left(\Phi(\Sigma_i, s^{(i)}, \mu_i)\right)_{i\in[k]}\right)\\ = \mathrm{APD}_{\mathrm{quad}}\left(\left(\Phi(\hat{\Sigma}_i, \hat{s}^{(i)}, \hat{\mu}_i)\right)_{i\in[k]}\right) &= \mathrm{APD}\left(\left(\hat{\Sigma}_i, \hat{s}^{(i)}, \hat{\mu}_i\right)_{i\in[k]}\right)\end{aligned}$$

holds.

For the reverse implication, assume that the further assumptions as well as

$$\mathrm{APD}\left(\left(\Sigma_i, s^{(i)}, \mu_i\right)_{i\in[k]}\right) = \mathrm{APD}\left(\left(\hat{\Sigma}_i, \hat{s}^{(i)}, \hat{\mu}_i\right)_{i\in[k]}\right)$$

hold. Consequently, we have that

$$\mathrm{APD}_{\mathrm{quad}}\left(\left(\Phi(\Sigma_i, s^{(i)}, \mu_i)\right)_{i\in[k]}\right) = \mathrm{APD}_{\mathrm{quad}}\left(\left(\Phi(\hat{\Sigma}_i, \hat{s}^{(i)}, \hat{\mu}_i)\right)_{i\in[k]}\right).$$

Thus, Theorem 3.16 yields the existence of $\lambda > 0$, $\Delta A \in \mathbb{R}^{d\times d}_{\mathrm{sym}}$, $\Delta a \in \mathbb{R}^d$ and $\Delta\alpha \in \mathbb{R}$ such that $\Phi(\Sigma_i, s^{(i)}, \mu_i) = \lambda\Phi(\hat{\Sigma}_i, \hat{s}^{(i)}, \hat{\mu}_i) + (\Delta A, \Delta a, \Delta\alpha)$ holds for all $i \in [k]$. With $(\Delta\Sigma, \Delta s, \Delta\mu) := \left(\Delta A, \frac{-1}{2\lambda}\Delta a, \Delta\alpha\right)$ and using that $\Phi|_{\mathbb{R}^{d\times d}_{\mathrm{sym},>0}\times\mathbb{R}^d\times\mathbb{R}}$ is invertible, we then get the desired result from the calculations above. □

Comparing the prerequisites for the reverse implications of Theorem 3.8 and Theorem 3.16, we notice anisotropic power diagrams to be slightly more demanding. In both cases, however, we require the existence and regularity of vertices.

Regarding the additional requirements, let us first briefly discuss the restriction to non-flat bisectors. If this is not the case, then the polynomial that defines a bisector $B_{i,l}$ either has degree one or must be reducible into the product of two polynomials of degree one. Consequently, the bisector consists of a single or the union of two hyperplanes. Thus, the "known" subset $P_i \cap P_l \subset B_{i,l}$ might miss the information of one of those hyperplanes.

The following example is to demonstrate the necessity of the vertex-connectivity as required in Theorem 3.16.

Example 3.18
We consider the anisotropic power diagram $\mathrm{APD}\left(\left(\Sigma_i, s^{(i)}, \mu_i\right)_{i\in[k]}\right)$ for $k = 5$ cells as depicted in Fig. 3.17. The figure illustrates the ellipsoidal parametrization of the diagram. Here, the norm associated with the first cell is simply the Euclidean norm, while the remaining 4 cells are all equipped with the same norm whose unit ball is a shrunk Euclidean ball. Thus, we obtain the spheres $\mathbb{E}_i = \left\{x \in \mathbb{R}^d : \left\|x - s^{(i)}\right\|_{\Sigma_i} = \sigma_i\right\}$ (with $\sigma_i = \sqrt{-\mu_i}$) drawn as dotted lines. This qualitative description suffices for the scope of this example.

The resulting diagram consists of 5 connected cells and has in total 4 non-degenerate vertices. However, it holds that $(P_2 \cup P_3) \cap (P_4 \cup P_5) = \emptyset$. Hence, the vertices in $P_1 \cap P_2 \cap P_3$ and $P_1 \cap P_4 \cap P_5$, respectively, are not connected.

We are now interested in alternative parameters that yield this diagram. In order to do so, let us temporarily consider the two diagrams that only consist of the cells with indices in either $\{1, 2, 3\}$ or $\{1, 4, 5\}$, respectively. Those diagrams would each meet the prerequisites of Corollary 3.17. Hence, we can conclude that for any alternative parametrization of the full diagram, there must exist respective scaling and shift parameters as stated in the corollary for each of those two reduced diagrams individually. However, as this example shows, those are not necessarily equal.

In order to see this, we choose an arbitrary $\lambda > 0$. We then set $\Delta\Sigma := \frac{1-\lambda}{\lambda}\Sigma_1$, $\Delta s := \frac{1-\lambda}{\lambda}\Sigma_1 s^{(1)}$, $\Delta\mu := (1-\lambda)(\mu_1 - (s^{(1)})^\intercal \Sigma_1 s^{(1)})$, and define $\left(\hat{\Sigma}_i, \hat{s}^{(i)}, \hat{\mu}_i\right)$ via Eqs. (3.40) to (3.42) for $i \in [3]$. For $i = 4, 5$ we set $\left(\hat{\Sigma}_i, \hat{s}^{(i)}, \hat{\mu}_i\right) := \left(\Sigma_i, s^{(i)}, \mu_i\right)$.

Now observe that by construction $\left(\hat{\Sigma}_1, \hat{s}^{(1)}, \hat{\mu}_1\right) = \left(\Sigma_1, s^{(1)}, \mu_1\right)$ holds. Consequently, we have found an alternative parametrization for the whole diagram that is not of the form presented in Corollary 3.17. Figure 3.17 also depicts the resulting adapted spheres $\hat{\mathbb{E}}_2$ and $\hat{\mathbb{E}}_3$ for some $\lambda > 1$.

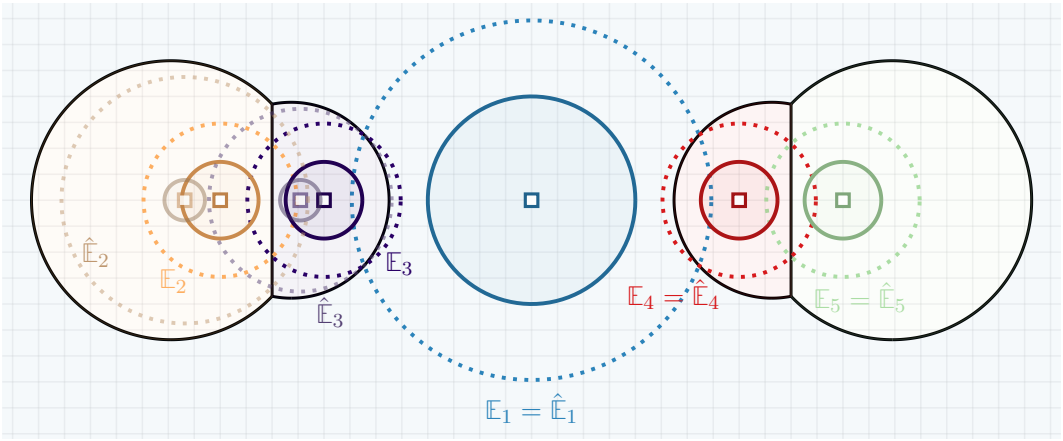

Figure 3.17: Anisotropic power diagram that is not vertex-connected. The solid colored circles depict the unit balls of the respective norms, the dotted circles the associated ellipsoids $\mathbb{E}_i = \left\{x \in \mathbb{R}^d : \left\|x - s^{(i)}\right\|_{\Sigma_i} = \sigma_i\right\}$, cf. Example 3.18.

Number of Connected Components

A cell of an anisotropic power diagram may consist of several connected components. In order to give a bound on the number of those, we need the following simple fact about quadratic inequalities.

Lemma 3.19
Let $P \dot\cup N \dot\cup Z = [d]$, $\delta_i > 0$ for $i \in P \cup N$, $b_Z \in \mathbb{R}^{|Z|}$, and $\gamma \in \mathbb{R}$. Set

$$C := \left\{ y \in \mathbb{R}^d : \sum_{i \in P} \delta_i y_i^2 - \sum_{i \in N} \delta_i y_i^2 + b_Z^\intercal y_Z + \gamma \geq 0 \right\}.$$

If $\gamma < 0$, either $Z = \emptyset$ or $b_Z = 0$, and $|P| = 1$, then C consists of exactly two connected components that are strongly separated by the hyperplane $\left\{ y \in \mathbb{R}^d : y_p = 0 \right\}$ with $\{p\} := P$. Otherwise, C has a single connected component.

Proof. Let $y, \hat y \in C$.

As first case, we consider $b_Z \neq 0$ (and in particular $Z \neq \emptyset$). For $\tau \in [0,1]$ we set

$$y(\tau) := (y_N(\tau), y_P(\tau), y_Z(\tau)) := (\sqrt{1-\tau}\ y_N, \sqrt{1-\tau}\ y_P, (1-\tau) y_Z - \tau \frac{\gamma}{b_Z^\intercal b_Z} b_Z).$$

Then

$$\sum_{i \in P} \delta_i y_i(\tau)^2 - \sum_{i \in N} \delta_i y_i(\tau)^2 + b_Z^\intercal y_Z(\tau) + \gamma = (1-\tau)(\sum_{i \in P} \delta_i y_i^2 - \sum_{i \in N} \delta_i y_i^2 + b_Z^\intercal y_Z + \gamma) \geq 0$$

and hence $y(\tau) \in C$ holds for all $\tau \in [0,1]$. We define $\hat y(\tau)$ analogously and have $y(1) = \hat y(1) = (0, 0, -\frac{\gamma}{b_Z^\intercal b_Z} b_Z)$, so y and $\hat y$ are connected in C.

Next, we consider the case $b_Z = 0$ or $Z = \emptyset$. In both situations we can disregard the components in Z. Thus, we w.l.o.g. assume $Z = \emptyset$.

For $\gamma \geq 0$, we may again define $y(\tau) := (\sqrt{1-t}\ y_N, \sqrt{1-t}\ y_P)$ and get

$$\sum_{i \in P} \delta_i y_i(\tau)^2 - \sum_{i \in N} \delta_i y_i(\tau)^2 + \gamma = (1-\tau)(\sum_{i \in P} \delta_i y_i^2 - \sum_{i \in N} \delta_i y_i^2 + \gamma) + \tau\gamma \geq 0,$$

which yields connectivity. Now assume $\gamma < 0$. We set

$$y(\tau) := (\sqrt{1-\tau} y_N, \sqrt{1 - \tau \frac{\sum_{i \in N} \delta_i y_i^2}{\sum_{i \in P} \delta_i y_i^2}}\, y_P)$$

for $\tau \in [0,1]$ (which is well-defined as $\sum_{i \in P} \delta_i y_i^2 \geq -\gamma + \sum_{i \in N} \delta_i y_i^2 > \sum_{i \in N} \delta_i y_i^2 \geq 0$). Then

$$\sum_{i \in P} \delta_i y_i(\tau)^2 - \sum_{i \in N} \delta_i y_i(\tau)^2 + \gamma = \sum_{i \in P} \delta_i y_i(\tau)^2 - \tau \sum_{i \in N} \delta_i y_i^2 - (1-\tau) \sum_{i \in N} \delta_i y_i^2 + \gamma \geq 0$$

and $y(1) = (\sqrt{1 - \frac{\sum_{i\in N}\delta_i y_i^2}{\sum_{i\in P}\delta_i y_i^2}}\, y_P, 0)$ hold. Thus, we can w.l.o.g. assume $y_N = 0$ as well as $\hat{y}_N = 0$ right away (otherwise we first "traverse" to $y(1)$ and $\hat{y}(1)$, respectively). With the ellipsoid $E := -\gamma \cdot (\operatorname{diag}((\delta_i^{-\frac{1}{2}})_{i\in P})\mathbb{B}_2^{|P|} \subset \mathbb{R}^{|P|}$ it holds that $\sum_{i\in P}\delta_i z_i^2 + \gamma \geq 0$ for all $z_P \in \mathbb{R}^{|P|} \setminus \operatorname{int}(E)$. We conclude $y, \hat{y} \in \mathbb{R}^{|P|} \times \{0\} \setminus (\operatorname{int}(E) \times \{0\}) \subset C$. If $|P| \geq 2$ holds, this yields connectivity.

As final case, we assume $|P| = 1$, $\gamma < 0$, and $Z = \emptyset$ (or $b_Z = 0$). W.l.o.g. let $P = \{1\}$. Then for all $y \in C$ it follows that $\delta_1 y_1^2 - \sum_{i\in N}\delta_i y_i(\tau)^2 + \gamma \geq 0$ and thus $|y_1| \geq \sqrt{\frac{-\gamma}{\delta_1}}$. We conclude that $C \subset \mathbb{R}^d \setminus \left\{y \in \mathbb{R}^d : -\sqrt{\frac{-\gamma}{\delta_1}} < y_1 < \sqrt{\frac{-\gamma}{\delta_1}}\right\}$ which yields the desired separation. Furthermore, $\pm(\sqrt{\frac{-\gamma}{\delta_1}}, 0, 0) \in C$ holds, so it follows that C has exactly two connected components. □

We can now formulate the following sharp bound on the number of connected components of a single cell in an anisotropic power diagram.

Theorem 3.20
Let $\mathcal{P}$ be an anisotropic power diagram. Then a cell of the diagram consists of at most

$$\sum_{i=0}^{d} \binom{k-1}{i} \tag{3.43}$$

connected components. This bound is sharp.

Proof. We show that there can be as many connected components of a cell as there are d-faces of a hyperplane arrangement in $\mathbb{R}^d$.

Consider the cell P_1 of $\mathcal{P}$. For every $i \in [k] \setminus \{1\}$ the dominance region of 1 over i is given by a quadratic inequality which can be brought to the form of Eq. (3.37) by means of a translation $t \in \mathbb{R}^d$ and a rotation $Q \in \mathbb{R}^{d\times d}$. From Lemma 3.19 it then follows that the dominance region of 1 over i consists of at most 2 connected components, which are separated by the hyperplane $H_i := \left\{x : (u^{(p)})^\intercal Qx = (u^{(p)})^\intercal Qt\right\}$ for a suitable $p \in [d]$.

Now consider the d-cells of the hyperplane arrangement from the at most $k-1$ hyperplanes obtained this way for $i = 2, \ldots, k$. Consequently, each of those cells can contain at most a single connected component of P_1. As is well-known (see, for example, [EOS86, Lemma 2.5]), an arrangement of n hyperplanes in $\mathbb{R}^d$ has at most $\sum_{i=0}^{d}\binom{n}{i}$ many d-cells. This bound is tight whenever the hyperplanes are in general position, i. e., any intersection of l of those hyperplanes is a $(d-l)$-dimensional affine subspace.

With $n := k-1$ this already yields the desired upper bound. In order to show that this bound is tight, let $H_i := H_{(b^{(i)},\beta_i)}$ for $i = 2, \ldots, k$ be some $k-1$ hyperplanes in

general position. For $i = 2, \dots, k$, w.l.o.g. assume $\left\|b^{(i)}\right\|_2 = 1$ and let $Q_i \in \mathbb{R}^{d\times d}$ be an orthonormal matrix with $b^{(i)}$ as the first column. Furthermore, let $\epsilon, \delta > 0$ and define the diagonal matrix $D := \operatorname{diag}(1, -\epsilon, \dots, -\epsilon) \in \mathbb{R}^d$.

We define an anisotropic power diagram $\mathrm{APD}_{\mathrm{quad}}\left((A_i, a_i, \alpha_i)_{i\in[k]}\right)$ in quadratic representation via $A_1 := 0$, $a_1 = 0$, $\alpha_1 = 0$ and $A_i := Q_i D Q_i^\intercal$, $a_i := -2\beta_i (b^{(i)})^\intercal Q_i D Q_i^\intercal$, $\alpha_i := (\beta_i b^{(i)})^\intercal Q_i D Q_i^\intercal (\beta_i b^{(i)}) - \delta = \beta_i^2 - \delta$, for $i = 2, \dots, k$.

Then the dominance region of 1 over i for $i \in [k] \setminus \{1\}$ is given by all $x \in \mathbb{R}^d$ that satisfy

$$\begin{aligned}&(x - \beta_i b^{(i)}) Q_i D Q_i^\intercal (x - \beta_i b^{(i)}) \geq \delta\\ \Leftrightarrow &(x^\intercal b^{(i)} - \beta_i)^2 \geq \delta + \epsilon(\|x\|_2^2 - ((b^{(i)})^\intercal x)^2).\end{aligned}$$

For $\epsilon \to 0$ this set converges (pointwise) to the set of all points of distance at least $\sqrt{\delta}$ from H_i. Letting furthermore $\delta \to 0$, this means that P_1 converges to the d-cells of the hyperplane arrangement given by $H_2, \dots, H_k$. □

Figure 3.18 illustrates the final construction of the proof of Theorem 3.20 in the plane for $k = 4$. Note that we could, of course, have set $\epsilon = 0$ right away. However, the construction shows that no degeneracy is necessary in order to achieve the bound given by Theorem 3.20.

Relation to Power Diagrams

As a final thought on anisotropic power diagrams, let us describe a further relevant relation to power diagrams. As also pointed out in [Wor08], anisotropic power diagrams can be constructed as certain projections of power diagrams in higher dimensional space.

Let an anisotropic power diagram $\mathcal{P} := (P_1, \dots, P_k) := \mathrm{APD}_{\mathrm{quad}}\left((A_i, a_i, \alpha_i)_{i\in[k]}\right)$ in quadratic representation be given. We define the mapping

$$\varphi : \begin{cases}\mathbb{R}^d & \to \mathbb{R}^{d+\frac{d(d+1)}{2}}\\ x & \mapsto \varphi(x) := \left(x_1, \dots, x_d, x_1^2, x_1x_2, \dots, x_{d-1}x_d, x_d^2\right)^\intercal.\end{cases} \tag{3.44}$$

We then consider the power diagram $\overline{\mathcal{P}} := (\overline{P_1}, \dots, \overline{P_k}) := \mathrm{PD}_{\mathrm{aff}}\left((\overline{a}_i, \alpha_i)_{i\in[k]}\right)$ in $\mathbb{R}^{d+\frac{d(d+1)}{2}}$ with

$$\overline{a}_i := (a_i^\intercal, (A_i)_{1,1}, 2(A_i)_{1,2}, \dots, 2(A_i)_{d-1,d}, (A_i)_{d,d})^\intercal$$

for $i \in [k]$. Let $(f_i)_{i\in[k]}$ and $(\overline{f_i})_{i\in[k]}$ denote the defining functions of $\mathcal{P}$ and $\overline{\mathcal{P}}$, respectively. Then by definition (cf. Eqs. (3.17) and (3.31)) it holds that $f_i(x) =$

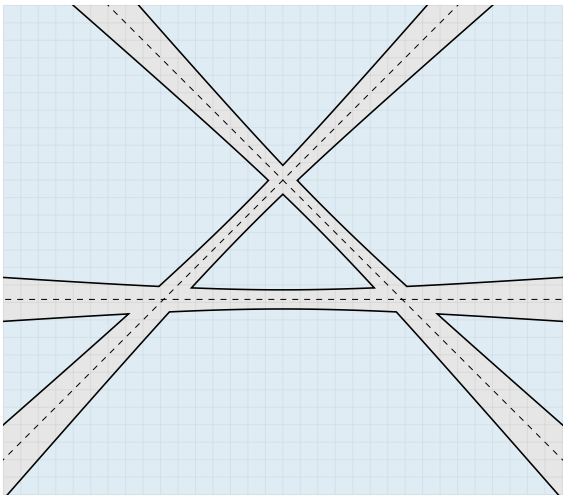

Figure 3.18: Anisotropic power diagram cell in the plane with a maximum number of connected components, as described in the proof of Theorem 3.20. The dashed lines depict the hyperplanes that separate the regions of dominance of the cell that is depicted in blue over each other cell in the diagram. The gray area depicts the union of the three remaining cells. Thus, $k = 4$ and Eq. (3.43) yields a maximum number of 7 connected components. This is the case for the cell depicted in blue.

$\overline{f_i}(\varphi(x))$ for all $x \in \mathbb{R}^d$ and $i \in [k]$. Consequently, we obtain the cells of the anisotropic power diagram $\mathcal{P}$ as the projection of the cells of the power diagram $\overline{\mathcal{P}}$ intersected with the image of φ, i. e., for every $i \in [k]$ it holds that

$$P_i = \pi_{[d]}(\overline{P_i} \cap \varphi(\mathbb{R}^d)).$$

3.3.4 Shortest-Path Diagrams

So far, all considered diagram types have lived in $\mathcal{X} = \mathbb{R}^d$. However, Definition 3.1 is not restricted to this setting. In particular, we can apply the concept of generalized Voronoi diagrams to a discrete or even finite setting. This may be given in form of a graph structure.

Variants of Voronoi diagrams on graphs are also discussed in [AKL13, Section 9.4.1] and [Oka+00, Section 3.8]. According to [AKL13], Voronoi diagrams on graphs are explicitly introduced for the first time by Mehlhorn [Meh88], who presents an approximation algorithm for the minimal Steiner tree problem that exploits the Voronoi diagram with the terminal nodes as sites. Most similar to our setting, Okabe et al. [Oka+08] consider "generalized network diagrams". They present several variants of diagrams on graphs, including (additively) weighted network Voronoi diagrams that coincide with our definition below.

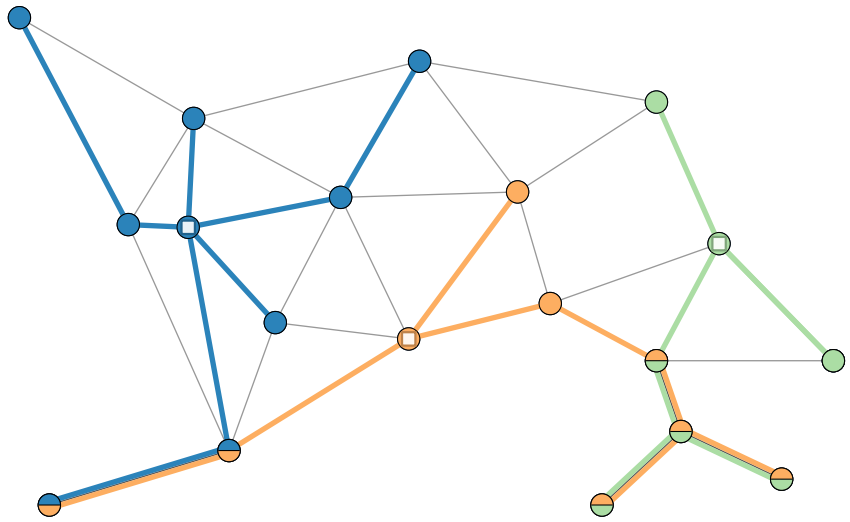

Figure 3.19: Exemplary shortest-path diagram.

Let us now assume $\mathcal{X}$ to be a finite set and $G = (\mathcal{X}, E, c)$ to be a connected, edge-weighted graph with $E \subset \binom{\mathcal{X}}{2}$ and $c : E \to \mathbb{R}_{\geq 0}$. Note that while we restrict ourselves to undirected graphs for the sake of simplicity, all results naturally carry over to directed ones, too. Let d_G denote the metric over $\mathcal{X}$ induced by G, i. e., $d_G(v, w)$ denotes the length (w. r. t. c) of a shortest v-w-path in G for $v, w \in \mathcal{X}$.

For sites $s^{(1)}, \ldots, s^{(k)} \in \mathcal{X}$ and weights $\mu_1, \ldots, \mu_k \in \mathbb{R}$ we then set

$$f_i(x) := d_G(s^{(i)}, x) + \mu_i \tag{3.45}$$

for $x \in \mathcal{X}$ and $i \in [k]$. Note that this is in accordance with our proposed framework for diagram functions given by Eq. (3.2) for $h = \mathrm{id}$ (we will review this choice also later in Section 7.3.2). We then call the diagram $\mathcal{P}$ w. r. t. the functions $(f_i)_{i \in [k]}$ a *shortest-path diagram*.

Figure 3.19 depicts an exemplary shortest-path diagram. It shows a graph that consists of 19 nodes and 35 edges whose weights are determined by the Euclidean distance of the nodes embedded in the plane as depicted. Nodes are colored according to the diagram cells they are contained in. As always, the sites are marked by squares. The colored edges illustrate the shortest paths of nodes to their assigned sites.

From Lemma 3.3 we immediately get that the cells of a shortest-path diagram are site-star-shaped (which is also observed in [Oka+00, Section 3.8]). In particular, this implies that subgraphs induced by the diagram cells, i. e., $G[P_i]$, $i \in [k]$, are connected.

Let us consider the structure of the diagram cells and in particular the resulting bisectors a little bit closer. For any site $s^{(i)}$, we can consider the tree T_i that consists of all nodes $\mathcal{X}$ but only edges in E that are contained on a shortest path from $s^{(i)}$ to

one of the edge's nodes. Whenever shortest paths are ambiguous, we could use some arbitrary tie-breaker rule so that T_i actually is a tree. However, in order to avoid technicalities, we will rather assume shortest paths to be unique (which is reasonable for the applications that we have in mind). We can furthermore root T_i in $s^{(i)}$ in order to obtain a partial order on $\mathcal{X}$. We define this order as $v \succeq_{T_i} w$ for $v, w \in \mathcal{X}$ whenever w lies on the shortest $s^{(i)}$-v-path in G.

From Eq. (3.4) in Lemma 3.3 we see that the subgraph of the trees T_i that is induced by the respective cell P_i, i. e., $T_i[P_i]$, is still connected. In Fig. 3.19, the colored edges depict those subtrees. The diagram cells are hence "spanned" by the shortest-path trees, which gives rise to the diagram name.

Concerning bisectors, we can deduce that whenever a node v lies in a bisector $B_{i,l}$, this must be the case for all common successors of v in T_i and T_l:

Lemma 3.21
Let G be a graph as introduced above. Assume the shortest path between any pair of nodes in G (w. r. t. c) to be unique. Let $\mathcal{P} = (P_1, \dots, P_k)$ be a shortest-path diagram over G and $v \in P_i \cap P_l$ for some $i, l \in [k]$.

Then $\{w \in P_i : w \succeq_{T_i} v\} = \{w \in P_l : w \succeq_{T_l} v\}$.

Proof. Let $w \in P_i$ such that w is a direct successor in T_i of v, i. e., $d_G(s^{(i)}, w) = d_G(s^{(i)}, v) + c(\{v, w\})$.

As $v \in P_i \cap P_k$ it holds that $f_i(v) + \mu_i = f_l(v) + \mu_l$ and hence $d_G(s^{(i)}, v) + \mu_i = d_G(s^{(l)}, v) + \mu_k$

In particular, $f_l(v) + \mu_l = d_G(s^{(l)}, w) + \mu_l \leq d_G(s^{(l)}, v) + c(\{v, w\}) + \mu_l = d_G(s^{(i)}, v) + c(\{v, w\}) + \mu_i = f_i(w) + \mu_i$. As $w \in P_i$, the latter must hold with equality and in particular implies $w \in P_l$ and that w must be a direct successor of v in T_l, too (as shortest paths are unique).

As this argument can be successively applied to all successors of v w. r. t. T_i, the claim follows. □

Consequently, we conclude that cells of shortest-path diagrams intersect in subtrees of the underlying shortest-path trees.

3.4 Summary & Conclusion

Generalized Voronoi diagrams yield a simple yet powerful tool that we are going to exploit in the remainder of this thesis.

In this chapter, Section 3.1 first gave a self-containing short introduction that was to demonstrate the versatility and capability of the concept of Voronoi diagrams.

Section 3.2 then formally introduced the terminology of generalized Voronoi diagrams for the scope of this thesis and stated some fundamental characteristics. In particular,

we will focus on diagrams defined w. r. t. functions in the form of Eq. (3.2) in order to derive meaningful diagram types.

Starting from classical Voronoi diagrams, Section 3.3 then introduced some basic types in increasing degree of generalization. Besides discussing some basic properties, we focused on the impact of the choice of parameters to the geometric properties of the resulting diagrams. Here, we believe that a proper understanding of those dependencies is of major benefit for any practical application. In particular, we have seen different parametrizations for both power and anisotropic power diagrams. While the one parametrization might give a simpler understanding of the impact of a parameter choice, the other one can be simpler to be handled in order to obtain theoretical claims. Also, invariances of those parametrizations were discussed. Of course, those are of interest as soon as the question how to determine the parameters that yield a "suitable" diagram will appear later on. We closed by the introduction of shortest-path diagrams, which due to their discrete nature might stick out. However, the possibility of working not only in Euclidean or related spaces but also on more general structures such as graphs once more demonstrates the broadness of the concept of generalized Voronoi diagrams.

Chapter 4

Correspondence between Diagrams and Clusterings: Finite Space

In Chapter 2 we defined the notion of constrained clusterings and in Chapter 3 generalized Voronoi diagrams were introduced. For our applications of interest, we aim to find constrained clusterings that furthermore have some desirable geometric or topological properties. In this chapter, we see that this can be achieved due to a strong correspondence between certain constrained clusterings and generalized Voronoi diagrams of prescribed type. More precisely, we will see that certain faces of the underlying clustering polytopes contain exactly those constrained clusterings that are induced by certain generalized Voronoi diagrams.

We will provide the definitions that yield the terminology of the relations between diagrams and clusterings in Section 4.1. Section 4.2 then gives an overview of the existing results on that relations. In the hope of providing a best-possible understanding, the remaining sections will then successively develop the desired correspondence with increasing generality. Section 4.3 will do so for power diagrams and least-squares clusterings with size constraints. Section 4.4 will then consider weight constraints and general diagram types. Finally, Section 4.5 will establish the correspondence for arbitrarily constrained clusterings as defined in Chapter 2.

4.1 Definition

First, let us properly define what is meant by a generalized Voronoi diagram to "induce" a clustering. The following definition is strongly oriented at the terminology as introduced in [BG12; BGK17].

Definition 4.1
Let $\mathcal{X}$ be a set and $\mathcal{P} = (P_1, \dots, P_k)$ be a generalized Voronoi diagram over $\mathcal{X}$. Let $X = \left\{x^{(1)}, \dots, x^{(n)}\right\} \subset \mathcal{X}$ be a finite set of n distinct elements in $\mathcal{X}$, and $T = T\left((A_i)_{i\in[k]}, (b^{(i)})_{i\in[k]}\right) \subset \mathbb{R}^{k\times n}$ a clustering polytope. Let $\xi \in T$.

We say $\mathcal{P}$ is *feasible* for ξ, if and only if

$$\left[\xi_{i,j} > 0 \Rightarrow x^{(j)} \in P_i\right] \text{ for all } i \in [k], j \in [n]. \tag{4.1}$$

Moreover, $\mathcal{P}$ *supports* ξ, if and only if

$$\left[\xi_{i,j} > 0 \Leftrightarrow x^{(j)} \in P_i\right] \text{ for all } i \in [k], j \in [n]. \tag{4.2}$$

As an alternative to the notation above, observe that $\mathcal{P}$ is feasible for ξ if and only if $\left\{x^{(j)} : j \in \operatorname{supp}(C_i)\right\} \subset P_i$ and $\mathcal{P}$ supports ξ if and only if $\left\{x^{(j)} : j \in \operatorname{supp}(C_i)\right\} = X \cap P_i$ for all $i \in [k]$.

Example 4.2
Let us briefly illustrate the terminology introduced by Definition 4.1 with the help of the example depicted in Fig. 4.1. Here, a fractional clustering of $n = 6$ points into $k = 3$ clusters is illustrated. The color of each of the points drawn as circles shall indicate their assignment. All points are integrally assigned to a single cluster with the exception of the point $(1, 2)^\intercal$ that is fractionally assigned to both the blue and the green cluster. Both Figs. 4.1a and 4.1b depict the same clustering.

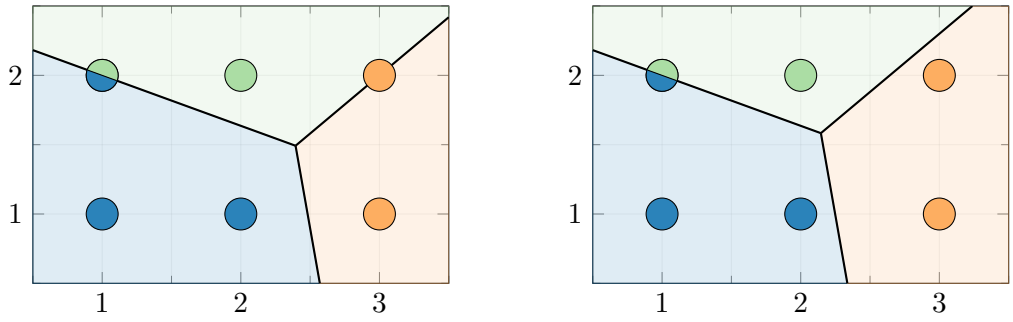

(a) Feasible diagram, not supporting.

(b) Supporting diagram.

Figure 4.1: Examples for feasible and supporting diagrams.

Furthermore, each of the two figures depict (different) power diagrams whose regions are indicated by the colored fillings and are bounded by the bold black lines.

Both diagrams are feasible for the given clustering as every point is contained in every region it is assigned to. However, only the diagram of Fig. 4.1b supports the clustering, as here no region contains any point that is not assigned to it by at least some positive fraction. By contrast, the diagram in Fig. 4.1a does not support the clustering. Here, the point $(3, 2)^\intercal$ lies in the intersection of the orange and green region but is fully assigned to the orange cluster.

One reason for the distinction between feasible and supporting diagrams is to avoid situations of ambiguity. While a feasible diagram only constraints the supports of clusters to be contained in the respective regions, a supporting diagram provides full information on them. Hence, if one is only provided the information that a diagram is feasible, the number of clusterings — even when being further restricted with respect to size or other constraints — that can be considered to be "induced" by that diagram may be large. Thus, we would like to avoid situations in which points may accidentally lie in a region without being assigned to it.

Example 4.3
Let us illustrate this by means of a second example. Figure 4.2 depicts the generalized Voronoi diagram w. r. t. functions $f_i(x) := \left\| x - s^{(i)} \right\|_1$ for $i = 1, 2$ and sites $s^{(1)} =$

$(19/8, 19/8)^\intercal$ and $s^{(2)} = (21/8, 21/8)^\intercal$ (depicted as white squares).

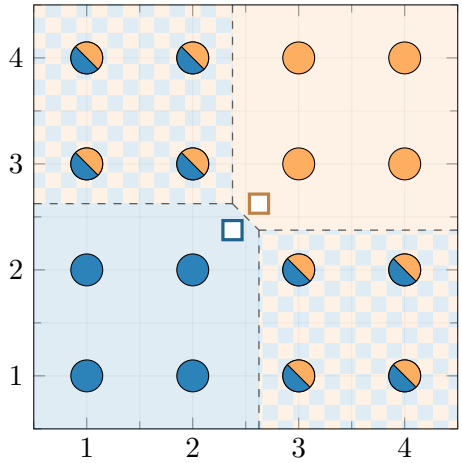

Figure 4.2: Diagram w. r. t. functions $f_i(x) = \left\|x - s^{(i)}\right\|_1$ and a supported clustering (see Example 4.3).

The sets $P_1 \setminus P_2$ and $P_2 \setminus P_1$ are filled by the respective colors blue and orange, while the intersection $P_1 \cap P_2$ is indicated by the blue-orange checkered pattern.

In addition, we consider the unit set $X = \{(x, y)^\intercal : x \in [4], y \in [4]\}$. Now for any clustering that assigns the four points in each the bottom left and top right quarter to the blue and orange cluster, respectively, and arbitrarily assigns the other 8 remaining points, the given diagram is feasible. On the other side, in order for the diagram to support a clustering, it must assign each point in the top left and bottom right quarter with some positive fraction to each of the two clusters. The latter case is depicted in Fig. 4.2.

As indicated by Example 4.3, the distinction between feasible and supporting diagrams becomes particularly crucial whenever we are considering diagrams with region intersections that may possibly contain a significant fraction of the data points to be clustered.

4.2 A Short Literature Review

In the following sections we establish a correspondence between constrained clusterings as defined in Chapter 2 and generalized Voronoi diagrams as defined in Chapter 3. Various variants and flavors of this relation have been presented in the literature. Let us give an informal overview of the ones most relevant to our approach.

Barnes, Hoffman, and Rothblum [BHR92] considered the partitioning of (unweighted)

point sets in $\mathbb{R}^d$. They introduce bounded-shape partition polytopes which consist of the vectors obtained by summing over all points for each cluster while the sizes of clusters are restricted. Here, they particularly observe that vertices of this polytope correspond to clusterings with disjoint convex hulls, which implies the existence of power diagrams that contain the clusters in the cell interiors.

Aurenhammer, Hoffmann, and Aronov [AHA98] then established that a clustering is optimal w. r. t. a least-square assignment to given sites under weight constraints if and only if there exists a power diagram that is feasible for it. In contrast to our approach, they prove this relation using mainly geometric arguments instead of an (explicit) optimization duality relation.

By similar arguments, Aronov, Carmi, and Katz [ACK09] established the analogue relation for additively weighted Voronoi diagrams and weight-constrained clusterings minimizing average Euclidean distances.

The relation of power diagrams and least-square assignments was then also established by Brieden and Gritzmann [BG12] by means of a linear programming duality that we consider, too. Here, in contrast to Aurenhammer, Hoffmann, and Aronov [AHA98] and our setting, also ranges of cluster weights are allowed. In particular, they also distinguish between feasible and supporting diagrams. Furthermore, they research gravity polytopes which are related to the bounded-shape partition polytopes and will be discussed in more detail in Chapter 6. The relation of power diagrams and weight-constrained least-squares clusterings will be treated in Section 4.3.

Geiß et al. [Gei+14] then established the more general correspondence between weight-constrained clusterings and generalized Voronoi diagrams with additive correction terms as we will see in Section 4.4. Note that they consider continuous clusterings of a bounded region in $\mathbb{R}^d$ such that clusters yield a partition into sets of prescribed Lebesgue measures. Similar to [AHA98], their proof relies only on geometric arguments. Quite the same result but using the same linear programming duality as we are going to do in the following was then established by Carlsson, Carlsson, and Devulapalli [CCD16] (cf. also [Río20, Section 4]). They discretize their continuous setting in order to use classical linear programming theory in their proof.

Also, it is worth mentioning that Schröder [Sch01] observed that clusters occurring in a Lagrangian-relaxation approach to an integer, weight-constrained clustering problem are contained in sets that from our point of view are the cells of generalized Voronoi diagrams with additive correction terms (cf. [Sch01, Chapter 7]; see also Section 7.3.2). In particular, he concludes those cells to be convex polygons for the case of squared Euclidean distances, i. e., power diagrams ([Sch01, Proposition 7.16]).

To the best of our knowledge, the relation of generally constrained clusterings and generalized Voronoi diagrams as presented in Section 4.5 has not occurred in the literature. While this is an admittedly straightforward generalization w. r. t. the proof, it significantly widens the theoretical understanding and the practical possibilities (as demonstrated by the example of moment-constrained clusterings in Section 4.5.3).

Also, the careful distinction between feasible and supporting diagrams (as in [BG12] for power diagrams) will be of practical relevance for us (for example, for the usage of shortest-path diagrams in Section 7.3.2).

4.3 Power Diagrams and Size-Constrained Clusterings

Eventually, we will establish the correspondence between generalized Voronoi diagrams and constrained clusterings as a linear programming duality result. However, before we state the general result, let us derive it in a more natural way.

In order to do so, let n distinct points $X = \left\{x^{(1)}, \dots, x^{(n)}\right\} \subset \mathbb{R}^d$ in some dimension d be given. Assume we are interested in finding a (fractional) clustering into k clusters of these points such that each cluster is of prescribed size $\kappa_i \in \mathbb{N}_{>0}$ for $i \in [k]$. More precisely, this means we accept all clusterings in the clustering polytope $T = T\left(A, (b^{(i)})_{i\in[k]}\right)$ for $A = \begin{pmatrix}1 & \dots & 1\end{pmatrix} \in \mathbb{R}^{1\times n}$ and $b^{(i)} = \kappa_i$ for $i \in [k]$.

We wonder whether there exists a clustering $\xi \in T$ together with some power diagram $\mathcal{P} = \mathrm{PD}\left(\left(s^{(i)}, \mu_i\right)_{i\in[k]}\right)$ for suitably chosen sites $s^{(i)} \in \mathbb{R}^d$ and parameters $\mu_i \in \mathbb{R}$, $i \in [k]$, such that this diagram $\mathcal{P}$ is feasible or, better yet, supports ξ. Next, we follow the argumentation in [BGK17].

Due to Eq. (4.1) in Definition 4.1, a diagram $\mathcal{P}$ is feasible for ξ whenever $\xi_{i,j} > 0$ implies that $x^{(j)} \in P_i$ for every $i \in [k]$ and $j \in [n]$. By definition of a cell P_i in Definition 3.1, we may rephrase Eq. (4.1) as

$$\left[\xi_{i,j} > 0 \Rightarrow \left\|x^{(j)} - s^{(i)}\right\|_2^2 + \mu_i = \min_{l\in[k]}\left(\left\|x^{(j)} - s^{(l)}\right\|_2^2 + \mu_l\right)\right] \forall i \in [k], j \in [n]. \quad (4.3)$$

Now for any $\xi \in T$ and $s^{(i)} \in \mathbb{R}^d, \mu_i \in \mathbb{R}$ for $i \in [k]$, it holds that $\xi_{i,j} \geq 0$ and trivially $\left\|x^{(j)} - s^{(i)}\right\|_2^2 + \mu_i \geq \min_{l\in[k]}\left(\left\|x^{(j)} - s^{(l)}\right\|_2^2 + \mu_l\right)$. Hence,

$$\xi_{i,j} \cdot \left(\left\|x^{(j)} - s^{(i)}\right\|_2^2 + \mu_i - \min_{l\in[k]}\left(\left\|x^{(j)} - s^{(l)}\right\|_2^2 + \mu_l\right)\right) \geq 0 \ \forall i \in [k], j \in [n] \quad (4.4)$$

holds independently of the feasibility of $\mathcal{P}$ for ξ.

Using Eq. (4.4) we can further reformulate Eq. (4.3) as

$$\begin{aligned} &\xi_{i,j} \cdot \left(\left\|x^{(j)} - s^{(i)}\right\|_2^2 + \mu_i - \min_{l\in[k]}\left(\left\|x^{(j)} - s^{(l)}\right\|_2^2 + \mu_l\right)\right) = 0 \ \forall i \in [k], j \in [n] \\ \Leftrightarrow &\sum_{i\in[k]}\sum_{j\in[n]} \xi_{i,j} \cdot \left(\left\|x^{(j)} - s^{(i)}\right\|_2^2 + \mu_i - \min_{l\in[k]}\left(\left\|x^{(j)} - s^{(l)}\right\|_2^2 + \mu_l\right)\right) = 0. \end{aligned} \quad (4.5)$$

In particular, due to Eq. (4.4) we see that Eq. (4.5) holds with "$\geq$" for any choice of $\xi \in T$ and diagram parameters $(s^{(i)}, \mu_i)_{i\in[k]}$. Thus, the task of finding a clustering and corresponding feasible diagram is equivalent to minimizing the left-hand side of Eq. (4.5) over all clusterings and diagram parameters. Then, whenever the optimal value 0 can be achieved, a clustering with corresponding feasible diagram exists and is provided by an optimizer.

At first glance it is not clear whether the left-hand side of Eq. (4.4), which is a non-convex function in ξ, $(s^{(i)})_{i\in[k]}$ and $(\mu_i)_{i\in[k]}$, can be minimized easily. However, we can simplify Eq. (4.4) in two steps.

First, as $\xi \in T$, we know that for all $j \in [n]$ it holds that

$$\sum_{i\in[k]} \xi_{i,j} = 1 \tag{4.6}$$

as well as for all $i \in [k]$ it holds that

$$\sum_{j\in[n]} \xi_{i,j} = \kappa_i. \tag{4.7}$$

We may now use Eq. (4.6) to simplify

$$\sum_{i\in[k]} \sum_{j\in[n]} \xi_{i,j} \cdot \min_{l\in[k]} \left(\left\| x^{(j)} - s^{(l)} \right\|_2^2 + \mu_l \right) = \sum_{j\in[n]} \min_{l\in[k]} \left(\left\| x^{(j)} - s^{(l)} \right\|_2^2 + \mu_l \right) \tag{4.8}$$

and Eq. (4.7) to simplify

$$\sum_{i\in[k]} \sum_{j\in[n]} \xi_{i,j}\mu_i = \sum_{i\in[k]} \kappa_i\mu_i. \tag{4.9}$$

Second, after plugging Eqs. (4.8) and (4.9) into Eq. (4.5), we may rearrange terms in order to separate all the terms containing ξ from the terms that contain components of $(\mu_i)_{i\in[k]}$. Doing so, we get that Eq. (4.5) is equivalent to

$$\sum_{i\in[k]} \sum_{j\in[n]} \xi_{i,j} \cdot \left\| x^{(j)} - s^{(i)} \right\|_2^2 = \sum_{j\in[n]} \min_{l\in[k]} \left(\left\| x^{(j)} - s^{(l)} \right\|_2^2 + \mu_l \right) - \sum_{i\in[k]} \kappa_i\mu_i. \tag{4.10}$$

Let us now consider both sides of Eq. (4.10) independently. Recall that due to Eq. (4.4) it holds that Eq. (4.5) and hence in particular Eq. (4.10) hold with "$\geq$" for any clustering and diagram. Thus, we obtain a minimization problem for the left-hand side and a maximization problem for the right-hand side.

Moreover, as those problems are coupled by the power diagram parameters $(s^{(i)})_{i\in[k]}$, let us for a moment assume those to be fixed to some arbitrary values.

Consequently, minimization of the left-hand side over $\xi \in T$ immediately yields the following linear program:

$$\min_{\xi \in \mathbb{R}^{k \times n}} \sum_{i \in [k]} \sum_{j \in [n]} \left\| x^{(j)} - s^{(i)} \right\|_2^2 \cdot \xi_{i,j} \tag{4.11}$$

$$\text{s.t.} \sum_{i \in [k]} \xi_{i,j} = 1 \qquad \forall j \in [n] \tag{4.11a}$$

$$\sum_{j \in [n]} \xi_{i,j} = \kappa_i \qquad \forall i \in [k] \tag{4.11b}$$

$$\xi \geq 0 \tag{4.11c}$$

Maximization over $(\mu_i)_{i \in [k]}$ of the right-hand side as stated by Eq. (4.10) yields an unconstrained maximization problem with a piecewise linear, concave objective. (The latter can be easily verified, as the pointwise minimum over a set of affine linear functions yields a concave, piecewise linear function.) By consideration of the epigraphs that correspond to the convex terms $-\min_{l \in [k]} \left(\left\| x^{(j)} - s^{(l)} \right\|_2^2 + \mu_l \right)$ we can easily rewrite this as a linear program as well. In order to do so, we introduce for every $j \in [n]$ an auxiliary variable η_j such that

$$\begin{aligned} & \eta_j \geq -\min_{l \in [k]} \left(\left\| x^{(j)} - s^{(l)} \right\|_2^2 + \mu_l \right) \\ \Leftrightarrow \qquad & -\eta_j \leq \left\| x^{(j)} - s^{(l)} \right\|_2^2 + \mu_l \qquad \forall l \in [k]. \end{aligned}$$

Thus, the problem of maximizing the right-hand side of (4.10) is equivalent to the linear program

$$\max_{\substack{(\mu_i)_i \in \mathbb{R}^k, \\ (\eta_j)_j \in \mathbb{R}^n}} -\sum_{j \in [n]} \eta_j - \sum_{i \in [k]} \kappa_i \mu_i \tag{4.12}$$

$$\text{s.t.} \quad -\eta_j - \mu_i \leq \left\| x^{(j)} - s^{(i)} \right\|_2^2 \quad \forall j \in [n], i \in [k]. \tag{4.12a}$$

However, we may now recognize that the linear programs (4.11) and (4.12) are in fact dual to each other. Here, $(-\eta_j)_{j \in [n]}$ yields the dual variables corresponding to constraints (4.11a) and $(-\mu_i)_{i \in [k]}$ yields the dual variables corresponding to constraints (4.11b).

Under the assumption of $T \neq \emptyset$, the program (4.11) is feasible and bounded. As a consequence, linear programming duality implies that, indeed, both programs share the same optimal objective value. This means, however, that any optimal primal-dual pair $(\xi, (\mu_i)_{i \in [k]}, (\eta_j)_{j \in [n]})$ fulfills Eq. (4.5), so the resulting diagram is feasible.

In particular, this does not depend on the choice of the sites $(s^{(i)})_{i \in [k]}$. Consequently, we conclude that for *any* choice of sites and cluster weights $(\kappa_i)_{i \in [k]}$ such that $\sum_{i \in [k]} \kappa_i =$

n (which implies $T \neq \emptyset$, cf. Proposition 2.3), we can always find a clustering $\xi \in T$ and additive weights $(\mu_i)_{i\in[k]}$, such that the resulting power diagram is feasible for ξ.

By construction, in any optimum of the program (4.12)

$$\eta_j = -\min_{l\in[k]} \left(\left\| x^{(j)} - s^{(l)} \right\|_2^2 + \mu_l \right) \forall j \in [n]. \tag{4.13}$$

holds. With this, Eq. (4.5) reads as

$$\sum_{i\in[k]} \sum_{j\in[n]} \xi_{i,j} \cdot \left(\left\| x^{(j)} - s^{(i)} \right\|_2^2 + \mu_i + \eta_j \right) = 0. \tag{4.14}$$

Hence, we can recognize Eq. (4.5) to be nothing else than the complementary slackness condition of the dual programs (4.11) and (4.12).

Next, one may ask whether this is still possible if we do not only head for a feasible but a supporting power diagram. In total analogy to Eq. (4.3), we may rephrase the definition of $\mathcal{P}$ supporting ξ as

$$\left[\xi_{i,j} > 0 \Leftrightarrow \left\| x^{(j)} - s^{(i)} \right\|_2^2 + \mu_i = \min_{l\in[k]} \left(\left\| x^{(j)} - s^{(l)} \right\|_2^2 + \mu_l \right) \right] \forall i \in [k], j \in [n]. \tag{4.15}$$

Now it follows that Eq. (4.15) is equivalent to Eq. (4.5) with the additional requirement that for each addend exactly one of the two factors equals zero. With the interpretation of Eq. (4.5) as a complementary slackness condition, this additional requirement directly translates into *strict* complementary slackness. From a more geometric point of view, this means that any primal dual pair of the programs (4.11) and (4.12) such that each solution lies in the respective relative interior of the optimal faces, results in a clustering together with a *supporting* power diagram.

Finally, another canonical question is whether *all* pairs of clusterings and corresponding feasible or supporting power diagrams can be derived this way. Let a clustering $\xi \in T$ and a power diagram with parameters $((s^{(i)})_{i\in[k]}, (\mu_i)_{i\in[k]})$ be given. Now define the auxiliary variables $(\eta_j)_{j\in[n]}$ via Eq. (4.13) and thus obtain a pair of primal and dual feasible solutions. By definition, the diagram is feasible for the clustering ξ if and only if Eq. (4.5) holds. By choice of $(\eta_j)_{j\in[n]}$, the latter is equivalent to Eq. (4.14), which is the complementarity condition for the two dual programs. Analogously, a supporting diagram yields a primal-dual pair with strict complementarity.

Let us summarize those observations in the following proposition:

Proposition 4.4
Let $x^{(j)} \in \mathbb{R}^d$ for $j \in [n]$, $s^{(i)} \in \mathbb{R}^d$, $\mu_i \in \mathbb{R}$, and $\kappa_i \in \mathbb{R}_{\geq 0}$ for $i \in [k]$ such that $\sum_{i\in[k]} \kappa_i = n$. Furthermore, let $\xi \in T\left((\mathbb{1}^{(n)})^\intercal, (\kappa_i)_{i\in[k]}\right)$. Define $(\eta_j)_{j\in[n]}$ via Eq. (4.13), so that $\left((\mu_i)_{i\in[k]}, (\eta_j)_{j\in[n]}\right)$ is feasible for the linear program (4.12).

Then the power diagram $\mathrm{PD}\left(\left(s^{(i)}, \mu_i\right)_{i\in[k]}\right)$ *is feasible for* ξ *if and only if* $\left(\xi, \left((\mu_i)_{i\in[k]}, (\eta_j)_{j\in[n]}\right)\right)$ *is a primal-dual optimizer of the programs* (4.11) *and* (4.12).

Furthermore, the power diagram w. r. t. $((s^{(i)})_{i\in[k]}, (\mu_i)_{i\in[k]})$ *supports* ξ *if and only if* $\left(\xi, \left((\mu_i)_{i\in[k]}, (\eta_j)_{j\in[n]}\right)\right)$ *satisfies the complementary slackness conditions of the programs* (4.11) *and* (4.12) *strictly.*

The main claim of Proposition 4.4, which sets power diagrams in correspondence to optimizers of the program (4.11), was (to best of our knowledge) first observed in [AHA98]. The formulation as a primal-dual relation was then established in [BG12].

From Proposition 4.4 we can conclude that the class of power diagrams is sufficiently large so that for any size constraints, there exists a diagram together with a feasible clustering. More precisely, it yields a one-to-one correspondence between feasible diagrams and certain faces of the underlying constrained clustering polytope T. A choice of sites $(s^{(i)})_{i\in[k]}$ determines the face of T that is the optimizer w. r. t. the direction $\left(\left\|x^{(j)} - s^{(i)}\right\|_2^2\right)_{i\in[k],j\in[n]}$. The dual variables that belong to the cluster constraints (4.11b) then yield an additive correction term μ_i for each of the power diagram's functions.

With this understanding in mind, we may now generalize this relation by reducing the argumentation above to its very core. As a very first step, we may recognize that no particular characteristic of the term $\left\|x^{(j)} - s^{(i)}\right\|_2^2$ was exhausted. In fact, let *any* functions $f_i : \mathbb{R}^d \to \mathbb{R}$ for $i \in [k]$ be given. If we now substitute the terms $\left\|x^{(j)} - s^{(i)}\right\|_2^2$ by $f_i(x^{(j)})$ in the previous section's argumentation, we immediately get that there exist additive weights $(\mu_i)_{i\in[k]}$ and a clustering $\xi \in T$ such that the generalized Voronoi diagram w. r. t. functions $\hat{f}_i := f_i + \mu_i$, $i \in [k]$, is feasible for ξ.

For another perspective, consider any face of the size-constrained clustering polytope T, i. e., a set of optimizers for a direction $(\gamma_{i,j})_{i\in[k],j\in[n]} \in \mathbb{R}^{k\times n}$. Then for any functions $f_i : \mathbb{R}^d \to \mathbb{R}$, $i \in [k]$, such that $f_i(x^{(j)}) = -\gamma_{i,j}$, there exist weights $(\mu_i)_{i\in[k]}$ such that the generalized Voronoi Diagram w. r. t. $\hat{f}_i := f_i + \mu_i$ is feasible for any clustering in that face, and even supporting if the clustering is taken from the relative interior.

4.4 Weight-Constrained Clusterings and Generalized Voronoi Diagrams

A next natural adaption is to introduce weights for each data point. Again, let n distinct points $X = \left\{x^{(1)}, \ldots, x^{(n)}\right\} \subset \mathbb{R}^d$ for some dimension d and functions $f_i : \mathbb{R}^d \to \mathbb{R}$ for $i \in [k]$ be given. Those may be thought of as individual distance measures for

each cluster. In order to equip the points with weights, let $\omega_j \in \mathbb{R}_{>0}$ for each $j \in [n]$ be given. Finally, let target cluster weights $\kappa_i \in \mathbb{R}_{>0}$ for $i \in [k]$ be given such that $\sum_{i\in[k]} \kappa_i = \sum_{j\in[n]} \omega_j$.

Again, we want to find a homogeneously single-constrained clustering, but this time in the polytope

$$T_{\kappa,\omega} := T\left(\begin{pmatrix}\omega_1 & \dots & \omega_n\end{pmatrix}, (\kappa_i)_{i\in[k]}\right). \tag{4.16}$$

If all weights are integers, this can be — ignoring any computational complexity issues — easily reduced to the unweighted case by simply introducing ω_j unweighted copies of each data point $x^{(j)}$. As a result of the previous section, by solving the program (4.11) we obtain a clustering (of the copied points) together with weights $(\mu_i)_{i\in[k]}$, such that the generalized Voronoi diagram w. r. t. $\hat{f}_i := f_i + \mu_i$ for $i \in [k]$ is feasible or even supporting if our solution obeys strict complementarity. By recombining the copies of each point we may then obtain a clustering in T so that each point $x^{(j)}$ is fractionally assigned to the ith cluster by the sum of fractional assignments of its copies divided by ω_j. Obviously, the originally obtained diagram is then still feasible or supporting for this clustering.

As we will work with weighted point sets throughout the remainder of this thesis, let us introduce some simplifying notation. Interpreting $(\omega_j)_{j\in[n]}$ as a measure on X, we will write

$$\omega(B) := \sum_{x^{(j)}\in X} \omega_j$$

for any $B \subset X$. For a clustering $\xi = (C_1, \dots, C_k)^\intercal \in T_{\kappa,\omega}$ we denote by

$$\omega(C_i) := \sum_{j\in[n]} \xi_{i,j}\omega_j \tag{4.17}$$

the weight of the ith cluster for $i \in [k]$.

As can be seen by just reproducing each step in the previous section, instead of splitting each point up to smaller fractions in order to recompose it afterwards, we could have just added a weighting factor to each of its occurrences in the derivation of the desired diagram.

Hence, we adapt the program (4.11) to

$$\min_{\xi\in\mathbb{R}^{k\times n}} \sum_{i\in[k]}\sum_{j\in[n]} \omega_j f_i(x^{(j)}) \cdot \xi_{i,j} \tag{4.18}$$

$$\text{s.t.} \quad \sum_{i\in[k]} \xi_{i,j} = 1 \quad \forall j \in [n] \tag{4.18a}$$

$$\sum_{j\in[n]} \omega_j\xi_{i,j} = \kappa_i \quad \forall i \in [k] \tag{4.18b}$$

$$\xi \geq 0. \tag{4.18c}$$

In analogy to Proposition 4.4, we can state the correspondence between optimizers of (4.18) and additively weighted generalized Voronoi diagrams by means of the following proposition. Here, the correspondence between weight-constrained clusterings and generalized Voronoi diagrams with additive correction terms was stated similarly in both [Gei+14] and [CCD16], with the latter exploiting the same linear programming duality.

Proposition 4.5
Let $x^{(j)} \in \mathbb{R}^d$ for $j \in [n]$. For each $i \in [k]$ let $\kappa_i \in \mathbb{R}_{\geq 0}$ and a function $f_i : \mathbb{R}^d \to \mathbb{R}$ be given. For each $j \in [n]$ let $\omega_j \in \mathbb{R}_{>0}$ such that $\sum_{j\in[n]} \omega_j = \sum_{i\in[k]} \kappa_i$ holds.

Let $\xi \in T_{\kappa,\omega}$.

Then ξ is an optimizer (in the relative interior of the optimal face) of the program (4.18), if and only if there exist $\mu_i \in \mathbb{R}$ for $i \in [k]$ such that the generalized Voronoi diagram w. r. t. $(f_i + \mu_i)_{i\in[k]}$ is feasible for (and supports) ξ.

Let us illustrate the results of Section 4.3 and Section 4.4 by means of a small example.

Example 4.6
Figure 4.3 depicts a Voronoi diagram in the plane w. r. t. $k = 4$ sites $(s^{(i)})_{i\in[k]}$ drawn as white squares with colored frame. There are $n = 25$ points drawn as disks. Each point $x^{(j)}$ has an associated weight $\omega_j \in \{1, 2, 3\}$. The weights are indicated by the disks' radii. As there are no points in the intersections of diagram cells, it uniquely induces the clustering as indicated by the points' colors. The resulting cluster weights are

$$\begin{pmatrix} \text{blue} & \text{orange} & \text{green} & \text{red} \\ 8 & 13 & 12 & 14 \end{pmatrix}.$$

Now, let us assume new desired cluster weights

$$\begin{pmatrix} \text{blue} & \text{orange} & \text{green} & \text{red} \\ 10 & 10 & 4 & 23 \end{pmatrix}$$

to be given. From Proposition 4.5 we know that there must exist a power diagram and a clustering $\xi^* \in T_{\kappa,\omega}$ (see Eq. (4.16)) such that the diagram supports ξ^*. (Note that $T_{\kappa,\omega} \neq \emptyset$ as by construction $\sum_{i\in[k]} \kappa_i = \sum_{j\in[n]} \omega_j$.)

Indeed, solving (4.18) with $f_i(x) := \left\|x - s^{(i)}\right\|_2^2$ for $i \in [k]$ yields a unique optimizer ξ^* and a corresponding power diagram with generating functions $\hat{f}_i(x) = \left\|x - s^{(i)}\right\|_2^2 + \mu_i$ with weights $(\mu_i)_{i\in[k]}$ (rounded to 4 decimal digits) given as

$$\begin{pmatrix} \mu_{\text{blue}} & \mu_{\text{orange}} & \mu_{\text{green}} & \mu_{\text{red}} \\ 0.0000 & 2.9239 & 4.3111 & -9.8454 \end{pmatrix}.$$

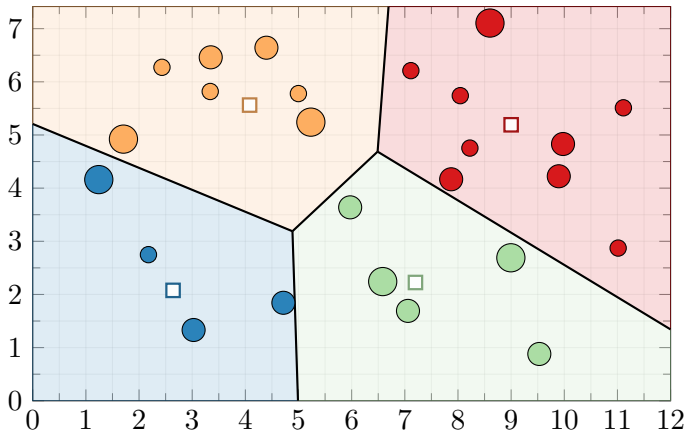

Figure 4.3: Clustering induced by the Voronoi diagram in Example 4.6.

Figure 4.4 depicts ξ^* together with the resulting supporting power diagram. In accordance with Corollary 2.8, there are in total 3 points that are non-integrally assigned, which are split among the 4 clusters as follows:

$x^{(j)}$	$(6.6, 2.2)^\intercal$	$(5.2, 5.2)^\intercal$	$(1.7, 4.9)$
blue	0	0	$2/3$
orange	0	$2/3$	$1/3$
green	$2/3$	0	0
red	$1/3$	$1/3$	0

4.5 General Correspondence

So far it has been shown that in order to fulfill size constraints of clusters, any choice of distance functions for the clusters may be adapted by additive shifts in order to obtain a clustering with a corresponding feasible or even supporting diagram. In the following, we will generalize this result for clusterings from any constrained clustering polytope as in Definition 2.1.

4.5.1 Constrained Clusterings

In this general setting, we assume that points $x^{(j)}$ for $j \in [n]$ in some arbitrary space $\mathcal{X}$ shall be clustered into k clusters. Note that in the following neither the data points $x^{(j)}$ nor their containing space $\mathcal{X}$ are of any theoretical interest themselves, as only their mappings by arbitrary distance measures or feature maps will be regarded.

Once more, we assume functions $f_i : \mathcal{X} \to \mathbb{R}$ for $i \in [k]$ to be given. Again, we may think of each f_i as a kind of distance measure for the ith cluster.

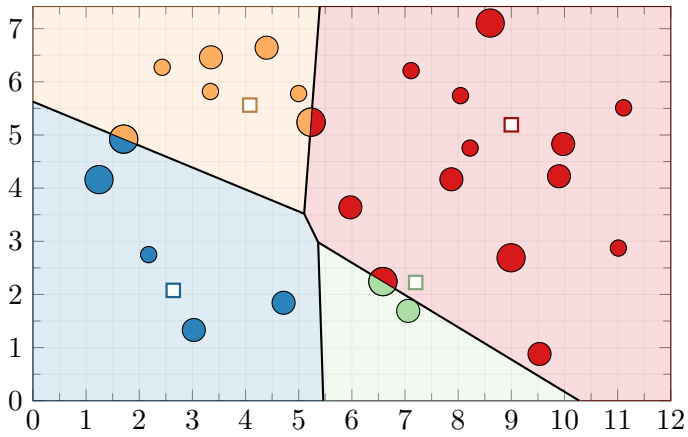

Figure 4.4: Power diagram resulting from enforcing cluster weights in Example 4.6.

Each cluster may be constrained w. r. t. the sums of quantitative features that are presumably derived from the data points $x^{(j)}$. In order to do so, we assume feature maps $\Phi_i : \mathcal{X} \to \mathbb{R}^{m_i}$ and target aggregates $b^{(i)} \in \mathbb{R}^{m_i}$ for each cluster $i \in [k]$ to be given. We set $A_i := \begin{pmatrix} \Phi_i(x^{(1)}) & \dots & \Phi_i(x^{(n)}) \end{pmatrix} \in \mathbb{R}^{m_i \times n}$ for each $i \in [k]$ and define the clustering polytope $T := T\left((A_i)_{i\in[k]}, (b^{(i)})_{i\in[k]}\right)$.

In analogy to Sections 4.3 and 4.4, we now consider the linear program that minimizes the aggregated assignment costs in terms of the distance measures $(f_i)_{i\in[k]}$ over T:

$$\min_{\xi \in \mathbb{R}^{k\times n}} \sum_{i\in[k]} \sum_{j\in[n]} f_i(x^{(j)}) \cdot \xi_{i,j} \tag{4.19}$$

$$\text{s.t.} \quad \sum_{i\in[k]} \xi_{i,j} = 1 \qquad \forall j \in [n] \tag{4.19a}$$

$$\sum_{j\in[n]} \Phi_i(x^{(j)})\xi_{i,j} = b^{(i)} \qquad \forall i \in [k] \tag{4.19b}$$

$$\xi \geq 0. \tag{4.19c}$$

In order to derive a corresponding dual program, we denote the dual variables that correspond to the constraints (4.19a) by $(\eta_j)_{j\in[n]} \in \mathbb{R}^n$ and the (vectors of) dual variables that correspond to the constraints in (4.19b) by $y^{(i)} \in \mathbb{R}^{m_i}$ for each $i \in [k]$. For the sake of readability later on, let us furthermore switch the signs of the (sign-unrestricted) dual variables $y^{(i)}$. We hence obtain the following dual of (4.19a):

$$\max_{\substack{y^{(i)}\in\mathbb{R}^k, i\in[k]\\ \eta_j\in\mathbb{R}, j\in[n]}} \sum_{j\in[n]} \eta_j - \sum_{i\in[k]} (b^{(i)})^\intercal y^{(i)} \tag{4.20}$$

$$\text{s.t.} \quad \eta_j - (\Phi_i(x^{(j)}))^\intercal y^{(i)} \quad \leq \quad f_i(x^{(j)}) \quad \forall j \in [n], i \in [k]. \tag{4.20a}$$

The complementary slackness condition for (4.19) and (4.20) then reads as

$$\sum_{i\in[k]} \sum_{j\in[n]} \xi_{i,j} \cdot \left(f_i(x^{(j)}) + (\Phi_i(x^{(j)}))^\intercal y^{(i)} - \eta_j \right) = 0. \tag{4.21}$$

As we did in Section 4.3, we will now find that for every optimal pair of (4.19) and (4.20) there is a generalized Voronoi diagram that is feasible or even supports the optimum of (4.19).

The following is a key theorem that exactly expresses this correspondence between primal-dual solution pairs of (4.19) and (4.20) with pairs of clusterings and corresponding feasible or supporting diagrams. More informally, while the primal feasible region of (4.19) by definition yields all clusterings obeying the given constraints, the feasible region of the dual program (4.20) projected to the $(y^{(i)})_{i\in[k]}$ components may be identified with the set of different generalized Voronoi diagrams that are obtained by adapting the distance functions f_i with linear combinations of the cluster constraints' feature mappings Φ_i. Here, the dual values $(y^{(i)})_{i\in[k]}$ yield the coefficients of those linear combinations. A diagram obtained from such adapted functions is then feasible (and supporting) for a clustering from the constrained clustering polytope, if and only if (strict) complementary slackness holds.

Theorem 4.7
Let $n \in \mathbb{N}$, $\mathcal{X}$ be some arbitrary space and $x^{(j)} \in \mathcal{X}$ for $j \in [n]$. Let $k \in \mathbb{N}$ and for every $i \in [k]$ let functions $f_i : \mathcal{X} \to \mathbb{R}$ and $\Phi_i : \mathcal{X} \to \mathbb{R}^{m_i}$ for $m_i \in \mathbb{N}$ as well as $b^{(i)} \in \mathbb{R}^{m_i}$ be given.

Define the constrained clustering polytope

$$T := T\left(\left(\begin{pmatrix} \Phi_i(x^{(1)}) & \dots & \Phi_i(x^{(n)}) \end{pmatrix} \right)_{i\in[k]}, \left(b^{(i)} \right)_{i\in[k]} \right) \tag{4.22}$$

and assume $\xi \in T$.

Finally, let $y^{(i)} \in \mathbb{R}^{m_i}$ for $i \in [k]$ and $\mathcal{P}$ be the generalized Voronoi diagram w. r. t. $\left(f_i + (y^{(i)})^\intercal \Phi_i \right)_{i\in[k]}$. For every $j \in [n]$, set

$$\eta_j := \min_{i\in[k]} \left(f_i(x^{(j)}) + \left(y^{(i)} \right)^\intercal \Phi_i(x^{(j)}) \right). \tag{4.23}$$

Then $\left((\eta_j)_{j\in[n]}, (y^{(i)})_{i\in[k]} \right)$ is feasible for (4.20) and the following equivalences hold:

$\mathcal{P}$ *is feasible for* ξ	$\Leftrightarrow$	ξ *and* $\Big((\eta_j)_{j\in[n]}, (y^{(i)})_{i\in[k]}\Big)$ *satisfy the complementary slackness condition* (4.21)
$\mathcal{P}$ *supports* ξ	$\Leftrightarrow$	ξ *and* $\Big((\eta_j)_{j\in[n]}, (y^{(i)})_{i\in[k]}\Big)$ *satisfy* (4.21) *with strict complementarity*

Proof. First, feasibility of $\Big((\eta_j)_{j\in[n]}, (y^{(i)})_{i\in[k]}\Big)$ for (4.20) follows immediately from the definition of $(\eta_j)_{j\in[n]}$ in Eq. (4.23).

In analogy to Section 4.3 we achieve the first equivalence quite immediately from the definition of the regions of $\mathcal{P}$:

$$
\begin{aligned}
&\mathcal{P} \text{ is feasible for } \xi \\
&\overset{(4.1)}{\Leftrightarrow} \Big[\xi_{i,j} > 0 \Rightarrow x^{(j)} \in P_i\Big] && \forall i \in [k], j \in [n] \\
&\overset{(3.1)}{\Leftrightarrow} \Big[\xi_{i,j} > 0 \Rightarrow f_i(x^{(j)}) + (y^{(i)})^\intercal \Phi_i(x^{(j)}) \\
&\qquad\qquad = \min_{l\in[k]} \Big(f_i(x^{(j)}) + (y^{(l)})^\intercal \Phi_l(x^{(j)})\Big)\Big] && \forall i \in [k], j \in [n] \\
&\overset{(4.23)}{\Leftrightarrow} \Big[\xi_{i,j} > 0 \Rightarrow f_i(x^{(j)}) + (y^{(i)})^\intercal \Phi_i(x^{(j)}) = \eta_j\Big] && \forall i \in [k], j \in [n] \\
&\overset{\substack{(4.19c),\\(4.20a)}}{\Leftrightarrow} \sum_{i\in[k], j\in[n]} \xi_{i,j} \cdot \Big(f_i(x^{(j)}) + (y^{(i)})^\intercal \Phi_i(x^{(j)}) - \eta_j\Big) = 0
\end{aligned}
$$

Next, Eq. (4.21) is satisfied with *strict* complementarity, if and only if

$$\Big[\xi_{i,j} > 0 \Leftrightarrow f_i(x^{(j)}) + (y^{(i)})^\intercal \Phi_i(x^{(j)}) - \eta_j = 0\Big] \qquad \forall i \in [k], j \in [n].$$

As this is (by the same argumentation as ahead) equivalent to

$$\Big[\xi_{i,j} > 0 \Leftrightarrow x^{(j)} \in P_i\Big] \qquad \forall i \in [k], j \in [n],$$

we obtain the second equivalence. □

The next two corollaries yield an interpretation of Theorem 4.7 as follows: Consider the constrained clustering polytope T obtained via feature maps Φ_i as in Eq. (4.22) and distance functions f_i for each cluster. Then the clusterings that allow a feasible (or supporting) generalized Voronoi diagram obtained from feature-map-adjusted distance functions are exactly those that lie in the (relative interior of the) optimal face of T w. r. t. the objective of the program (4.19). Also, whenever we are given a generalized Voronoi diagram that is feasible (or supports) a clustering from any clustering polytope, we know that this clustering must lie in the (relative interior of the) face with its outer normal determined by the diagram's defining functions.

Corollary 4.8
Let $\mathcal{X}$ be some arbitrary space and $x^{(j)} \in \mathcal{X}$ for $j \in [n]$. For $i \in [k]$, let functions $f_i : \mathcal{X} \to \mathbb{R}$ and $\Phi_i : \mathcal{X} \to \mathbb{R}^{m_i}$ for $m_i \in \mathbb{N}$, as well as $b^{(i)} \in \mathbb{R}^{m_i}$ be given. Let T be the clustering polytope as defined by Eq. (4.22) and $\xi \in T$.

Finally, let F and D denote the optimal faces of the linear programs (4.19) and (4.20), respectively. Let $(\eta_j)_{j\in[n]}$ be given by Eq. (4.23).

Then $F \neq \emptyset$, $D \neq \emptyset$ and there exists $S \subset \prod_{i\in[k]} \mathbb{R}^{m_i}$ such that $D = \left\{(\eta_j)_{j\in[n]}\right\} \times S$. Furthermore, the following holds:

i) $\xi \in F \quad \Leftrightarrow$ *For all $(y^{(i)})_{i\in[k]} \in S$, the generalized Voronoi diagram w. r. t. $\left(f_i + (y^{(i)})^\intercal \Phi_i\right)_{i\in[k]}$ is feasible for ξ.*

ii) $\xi \in \operatorname{relint}(F) \quad \Leftrightarrow$ *For all $(y^{(i)})_{i\in[k]} \in \operatorname{relint}(S)$, the generalized Voronoi diagram w. r. t. $\left(f_i + (y^{(i)})^\intercal \Phi_i\right)_{i\in[k]}$ supports ξ.*

Proof. As $T \neq \emptyset$ by assumption and bounded by definition, we immediately get $F, D \neq \emptyset$. From the constraints (4.20a) together with the objective of the linear program (4.20), it follows immediately that Eq. (4.23) must hold in any optimum of (4.20). Hence, the set S as described must exist.

By Theorem 4.7, it follows that $\mathcal{P}$ is feasible for ξ if and only if ξ and $\left((\eta_j)_{j\in[n]}, (0)_{i\in[k]}\right)$ obey the complementary slackness condition (4.21). As this is a primal and dual feasible pair by construction, this is equivalent to both being optimal. This implies i).

Next, strict complementarity for any feasible pair of solutions for a pair of primal-dual linear programs holds if and only if both solutions lie in the relative interior of the respective optimal faces (cf. [Van08, Section 5] or [Roo92, Corollary 2.1]). Thus, we may further conclude from Theorem 4.7 that $\mathcal{P}$ supporting ξ is equivalent to both ξ and $\left((\eta_j)_{j\in[n]}, (0)_{i\in[k]}\right)$ lying in the relative interior of the primal and dual optimal face, respectively. Thus, we can also conclude ii). □

Corollary 4.9
Let $\mathcal{X}$ be some arbitrary space and $x^{(j)} \in \mathcal{X}$ for $j \in [n]$ pairwise distinct.

Let $T = T\left((A_i)_{i\in[k]}, (b^{(i)})_{i\in[k]}\right) \subset \mathbb{R}^{k\times n}$ be a clustering polytope, $\xi^* \in T$ some clustering and $\mathcal{P}$ be a generalized Voronoi diagram w. r. t. functions $f_i : \mathcal{X} \to \mathbb{R}$, $i \in [k]$.

Let F be the optimal face of the linear program

$$\min_{\xi\in T} \sum_{i\in[k], j\in[n]} f_i(x^{(j)})\xi_{i,j}. \tag{4.24}$$

Then the following holds:

i) $\mathcal{P}$ *is feasible for* $\xi^* \Rightarrow \quad \xi^* \in F$

ii) $\mathcal{P}$ *supports* $\xi^* \quad \Rightarrow \quad \xi^* \in \operatorname{relint}(F)$

Proof. For every $i \in [k]$ we may write $A_i =: \left(a^{(i,1)}, \ldots, a^{(i,n)}\right) \in \mathbb{R}^{m_i \times n}$ and define a function $\Phi_i : \mathcal{X} \to \mathbb{R}^{m_i}$ via $\Phi_i(x^{(j)}) := a^{(i,n)}$ for all $j \in [n]$ and an arbitrary extension on $\mathcal{X}$. The Φ_i are well-defined as we require the $x^{(j)}$ to be pairwise distinct.

Then the programs (4.24) and (4.19) are in fact identical. We set $y^{(i)} := 0 \in \mathbb{R}^{m_i}$ for $i \in [k]$ and define $(\eta_j)_{j\in[n]}$ as in Eq. (4.23).

Hence, by Theorem 4.7 it follows that $\mathcal{P}$ is feasible for ξ^* if and only if ξ^* and $\left((\eta_j)_{j\in[n]}, (0)_{i\in[k]}\right)$ obey the complementary slackness condition (4.21). Thus, feasibility of $\mathcal{P}$ implies optimality of ξ^*. If $\mathcal{P}$ is furthermore supporting, Theorem 4.7 yields strict complementarity which furthermore implies $\xi^* \in \operatorname{relint}(F)$. □

4.5.2 Constrained Clusterings with Point Weights

In Section 4.4 we have already treated clusterings of weighted point sets. Weight-constraints are typical for the sort of clustering applications we are interested in. However, the notation in the previous Section 4.5.1 might not immediately reveal how additional weights of the data points $x^{(j)}$ could be considered and which type of constraints easily allow the combination with weight-constraints.

Hence, let us consider the same setting as in Section 4.5.1, but furthermore assume point weights $\omega_j \in \mathbb{R}_{>0}$ for $j \in [n]$ to be given. Now instead of constraints that sum up the feature vectors Φ_i, we consider constraints that sum up the weighted feature vectors, i. e., we consider the clustering polytope

$$T := T\left(\left(\begin{pmatrix}\omega_1 \cdot \Phi_i(x^{(1)}) & \ldots & \omega_n \cdot \Phi_i(x^{(n)})\end{pmatrix}\right)_{i\in[k]}, \left(b^{(i)}\right)_{i\in[k]}\right). \tag{4.25}$$

In the same manner, we adapt the objective of the linear program (4.19) by a weighted version and thus obtain the adapted program

$$\min_{\xi\in\mathbb{R}^{k\times n}} \sum_{i\in[k]}\sum_{j\in[n]} \omega_j \cdot f_i(x^{(j)}) \cdot \xi_{i,j} \tag{4.26}$$

$$\text{s.t.} \quad \sum_{i\in[k]} \xi_{i,j} = 1 \qquad \forall j \in [n] \tag{4.26a}$$

$$\sum_{j\in[n]} \omega_j \cdot \Phi_i(x^{(j)}) \cdot \xi_{i,j} = b^{(i)} \quad \forall i \in [k] \tag{4.26b}$$

$$\xi \geq 0. \tag{4.26c}$$

as well as its dual

$$\max_{\substack{y^{(i)}\in\mathbb{R}^k, i\in[k]\\ \eta_j\in\mathbb{R}, j\in[n]}} \sum_{j\in[n]} \eta_j - \sum_{i\in[k]} (b^{(i)})^\intercal y^{(i)} \tag{4.27}$$

$$\text{s.t.} \quad \eta_j - (\omega_j \Phi_i(x^{(j)})^\intercal y^{(i)} \leq \omega_j \cdot f_i(x^{(j)}) \quad \forall j \in [n], i \in [k]. \tag{4.27a}$$

We may now state a point-weighted version of Corollary 4.8. Note that the only difference is the replacement of both the clustering polytope of interest as well as the underlying linear program by their weighted versions.

Corollary 4.10
Let $\mathcal{X}$ be some arbitrary space, $x^{(j)} \in \mathcal{X}$ for $j \in [n]$ pairwise distinct, and weights $\omega_j \in \mathbb{R}_{>0}$ for $j \in [n]$ be given. Let $k \in \mathbb{N}$ and for every $i \in [k]$, let functions $f_i : \mathcal{X} \to \mathbb{R}$ and $\Phi_i : \mathcal{X} \to \mathbb{R}^{m_i}$ for $m_i \in \mathbb{N}$, as well as $b^{(i)} \in \mathbb{R}^{m_i}$ be given. Let T be the clustering polytope given by Eq. (4.25) and $\xi \in T$.

Finally, let F and D denote the optimal faces of the linear programs (4.26) and (4.27), respectively. Set $\eta_j := \omega(x^{(j)}) \cdot \min_{i\in[k]} \left(f_i(x^{(j)}) + \left(y^{(i)}\right)^\intercal \Phi_i(x^{(j)}) \right)$ for all $j \in [n]$.

Then $F \neq \emptyset$, $D \neq \emptyset$ and there exists $S \subset \prod_{i\in[k]} \mathbb{R}^{m_i}$ such that $D = \{(\eta_j)_j\} \times S$. Furthermore, the following holds:

i) $\xi \in F$ $\Leftrightarrow$ *For all $(y^{(i)})_i \in S$, the generalized Voronoi diagram w. r. t. $\left(f_i + (y^{(i)})^\intercal \Phi_i\right)_{i\in[k]}$ is feasible for ξ.*

ii) $\xi \in \operatorname{relint}(F)$ $\Leftrightarrow$ *For all $(y^{(i)})_i \in \operatorname{relint}(S)$, the generalized Voronoi diagram w. r. t. $\left(f_i + (y^{(i)})^\intercal \Phi_i\right)_{i\in[k]}$ supports ξ.*

Proof. This is analogous to Corollary 4.8. Let $\omega : \mathcal{X} \to \mathbb{R}_{>0}$ with $\omega(x^{(j)}) = \omega_j$ for $j \in [n]$ be some arbitrary extension of the given weights to $\mathcal{X}$. For $i \in [k]$, define $\hat{f}_i := \omega \cdot f_i$ as well as $\hat{\Phi}_i := \omega \cdot \Phi_i$ and inject those into Corollary 4.8. Now equivalence follows by the simple observation that a generalized Voronoi diagram w. r. t. functions $\left(f_i + (y^{(i)})^\intercal \Phi_i\right)_{i\in[k]}$ is equal to the diagram w. r. t. the functions $\left(\omega \cdot (f_i + (y^{(i)})^\intercal \Phi_i)\right)_{i\in[k]}$ by Lemma 3.2. □

4.5.3 Moment-Constrained Clusterings

Let us now demonstrate the practical relevance of this general relation by means of the class of generalized Voronoi diagrams obtained from polynomials of bounded degree. As in Section 4.5.2, we assume a set of n distinct points $X = \left\{x^{(1)}, \ldots, x^{(n)}\right\} \subset \mathbb{R}^d$ and weights $\omega_j \in \mathbb{R}_{>0}$ for $j \in [n]$ to be given.

Furthermore, assume $q \in \mathbb{N}_{>0}$ to be given. As the common feature map $\Phi =: \Phi_i$ for all clusters $i \in [k]$, we map $x \in \mathbb{R}^d$ to all monomials of degree at most q, i. e.,

$$\Phi : \begin{cases} \mathbb{R}^d & \to \mathbb{R}^{\binom{q+d}{d}} \\ x & \mapsto (x^\alpha)_{\alpha \in \mathbb{N}^d : |\alpha| \leq q} \end{cases}. \tag{4.28}$$

Here, we use the multi-index notation $x^\alpha = \prod_{i=1}^d x_i^{\alpha_i}$ for $\alpha = (\alpha_1, \dots, \alpha_d)^\intercal \in \mathbb{N}^d$ with $|\alpha| := \sum_{i=1}^d \alpha_i$.

Now consider the homogeneously constrained clustering polytope T as defined by Eq. (4.25) for given right-hand sides $b^{(i)} \in \mathbb{R}^{\binom{q+d}{d}}$, $i \in [k]$.

We may interpret T as follows: For a clustering $\xi = (C_1, \dots, C_k)^\intercal \in T$, a cluster C_i can be identified with the probability distribution on X that maps $x^{(j)} \mapsto \frac{\omega_j}{\omega(C_i)} \xi_{i,j}$ for $j \in [n]$. As $\xi \in T$, the (statistical) *moments* up to degree q of this distribution are determined by the according right-hand side $b^{(i)}$ by definition. More formally, if the d-variate random variable $\overline{X}$ is distributed according to such a distribution, this fixes $\mathbb{E}\left[\overline{X}^\alpha\right]$ for any index vector $\alpha \in \mathbb{N}^d$ with $|\alpha| \leq q$. In particular, this implies the mean (i. e., centroid) and, if $q \geq 2$, the covariance of clusters. In other words, this determines the first q coefficients of the moment-generating function's Taylor series for this distribution (cf. [DS12, Section 4.4]).

We can now formulate Corollary 4.8 for this setting. Basically, this yields that fixing moments up to degree q results in correction terms in form of polynomials of degree at most q.

Corollary 4.11
Let $q \in \mathbb{N}$, $x^{(j)} \in \mathbb{R}^d$ and $\omega_j \in \mathbb{R}_{>0}$ for $j \in [n]$ be given. Further, let $k \in \mathbb{N}_{>0}$, functions $f_i : \mathbb{R}^d \to \mathbb{R}$, and vectors $b^{(i)} \in \mathbb{R}^{\binom{q+d}{d}}$ for $i \in [k]$ be given. With Φ as given by Eq. (4.28), define the homogeneously-constrained clustering polytope

$$T := T\left(\begin{pmatrix} \omega_1 \Phi(x^{(1)}) & \dots & \omega_n \Phi(x^{(n)}) \end{pmatrix}, \left(b^{(i)}\right)_{i \in [k]}\right) \tag{4.29}$$

and let $\xi \in T$.

Then ξ is an optimizer (in the relative interior of the optimal face) of the according linear program (4.26) if and only if there exist polynomials $p_i \in \mathbb{R}[x_1, \dots, x_d]$ with $deg(()p_i) \leq q$ for $i \in [k]$ such that the generalized Voronoi diagram w. r. t. $(f_i + p_i)_{i \in [k]}$ is feasible for (and supports) ξ.

Proof. This follows immediately from Corollary 4.8 as the definition of Φ yields that the resulting correction terms $(y^{(i)})^\intercal \Phi$ are polynomials of degree less than or equal to q. □

Implicitly, this result inherits a certain invariance w. r. t. the choice of the functions $(f_i)_{i\in[k]}$. For the current setting, Corollary 4.11 implies that there is invariance w. r. t. the space of polynomials up to degree q:

Corollary 4.12
In the same setting as in Corollary 4.11, let $\xi \in T$ be an optimizer of the program (4.19) *and let polynomials $p_i \in \mathbb{R}[x_1, \dots, x_d]$ with $deg(()p_i) \leq q$ for $i \in [k]$ be given.*

Then ξ is also an optimizer of the program (4.26) *when replacing the functions f_i by $f_i + p_i$ for $i \in [k]$.*

Proof. By Corollary 4.11, there exist polynomials $\hat{p}_i$ for $i \in [k]$ such that the diagram w. r. t. $(f_i+\hat{p}_i)_{i\in[k]}$ is feasible for ξ. With $\tilde{p}_i := \hat{p}_i - p_i$ for $i \in [k]$, this means the diagram w. r. t. $(f_i + p_i + \tilde{p}_i)_{i\in[k]}$ is feasible for ξ and so the claim follows by Corollary 4.11, too. □

Let us break this down onto a practical level. Theoretically, Corollary 4.11 states that the class of diagrams that result from the polynomials up to degree q is rich enough in order to fix the first q moments of clusters. However, from Corollary 4.12 we learn that we are at the same time invariant to shifts by polynomials of degree q. If we choose each function f_i as a polynomial of degree q, the resulting clustering is the same as setting $f_i = 0$ for every $i \in [k]$. This yields the degenerated case of every point lying in any diagram cell and thus a fully fractional clustering (in a supporting case).

However, we may patch this drawback by adding another degree: We choose as class of diagram functions polynomials of degree $q + 1$. If we furthermore assume that pairwise differences $f_i - f_l$ for $i \neq l \in [k]$ are polynomials of (exact) degree $q + 1$ as well, we can conclude that we will end up with a non-degenerate diagram that supports the resulting clustering.

We have already seen examples for this in the previous sections. Power diagrams as introduced in Section 3.3.2 have been shown to allow any weight-constrained clusterings in Section 4.3. This is in accordance with the observations above, as we know that power diagrams (in affine parametrization) are in fact generated by polynomials of degree 1 and hence allow fixing the 0th moment. Note that for power diagrams in spherical parametrization the differences of the resulting functions are in fact affine. Similarly, anisotropic power diagrams have been shown to be parametrized by polynomials of degree 2 (cf. Section 3.3.3). Here, if the quadratic terms do not vanish in the differences $f_i - f_l$, i. e., if the local norms are chosen pairwise distinct, we see that anisotropic power diagrams allow constraining both the weight as well as the first moments, i. e., the centroids, of clusters.

Furthermore, we see that due to Corollary 4.12 the initial choice of the additive parameters μ_i for both power and anisotropic power diagrams is irrelevant if weights of clusters are constrained. If furthermore the centroids are constrained in the case of

anisotropic power diagrams, also the choice of sites will be dominated by the resulting correction terms.

Let us close this chapter with an example that demonstrates the fixation of centroids by using anisotropic power diagrams, as well as how the distinction between feasible and supporting diagrams in Corollary 4.9 comes into play here.

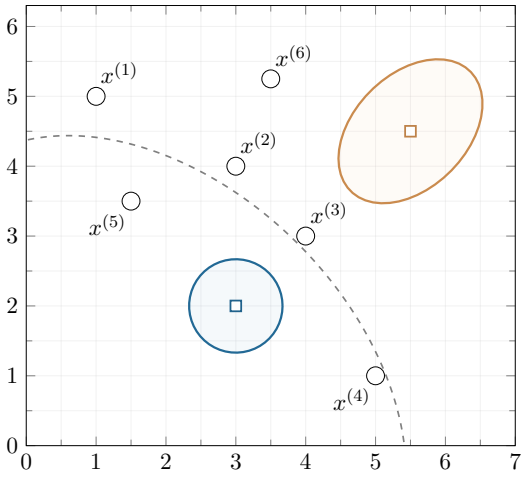

Figure 4.5: Basic setting in Example 4.13.

Example 4.13
Consider the setting as depicted in Fig. 4.5. Let the set of 6 points $X = \left\{x^{(1)}, \dots, x^{(6)}\right\}$ in $\mathbb{R}^2$ be given as drawn in the figure. Furthermore, we assume two ellipsoidal norms with unit balls as depicted by the blue and orange ellipses to be given. The gray-dashed line in Fig. 4.5 depicts the bisector of the anisotropic power diagram w. r. t. those norms, sites as given by the blue and orange framed squares in the figure, and no additive weights (i. e., $\mu_1 = \mu_2 = 0$).

Assume all points to have the same unit weight, i. e., $\omega_j = 1$ for $j \in [n]$. We would now like to obtain a clustering into $k = 2$ clusters that is constrained to cluster weights $\kappa_1 = \frac{7}{2}$ and $\kappa_2 = \frac{5}{2}$ as well as centroids $c^{(1)} = (3, 3)^\intercal$ and $c^{(2)} = (3, 9/2)^\intercal$, respectively.

In order to do so, we define Φ as in Eq. (4.28), functions f_i according to Eq. (3.26), and solve the resulting linear program (4.26). We then obtain a clustering and an anisotropic power diagram according to Corollary 4.11.

The obtained clustering is

$$\xi = \begin{array}{c|cccccc} & x^{(1)} & x^{(2)} & x^{(3)} & x^{(4)} & x^{(5)} & x^{(6)} \\ \hline C_1 & 1/2 & 1/2 & 1/2 & 1 & 1 & 0 \\ C_2 & 1/2 & 1/2 & 1/2 & 0 & 0 & 1 \end{array}.$$

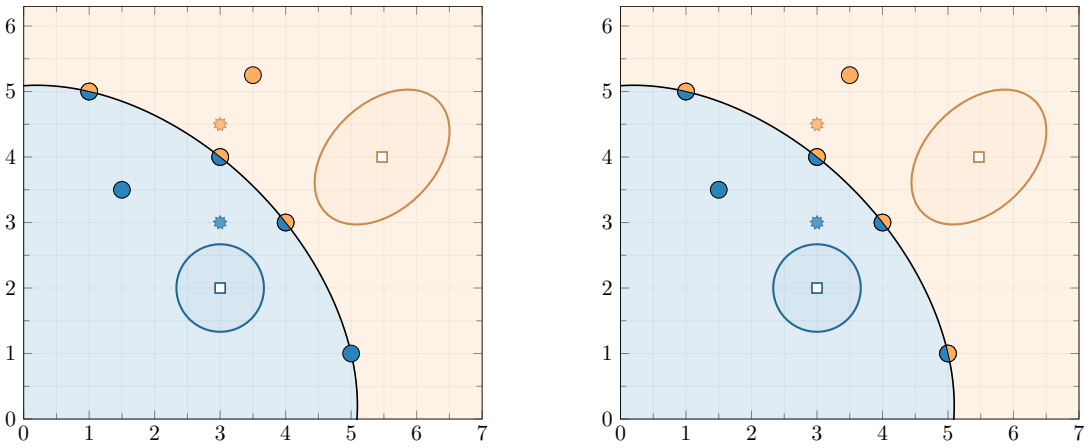

(a) Extremal clustering.

(b) Clustering in the relative interior of the optimal face.

Figure 4.6: Setting in Example 4.13 with non-unique optimal clustering.

For the corresponding anisotropic power diagram we get the sites and additive weights (up to some numerical precision)

$$s^{(1)} = \begin{pmatrix} 3 \\ 2 \end{pmatrix}, \quad \mu_1 = 0, \text{ and } \quad s^{(2)} = \begin{pmatrix} 5.4724 \\ 4 \end{pmatrix}, \quad \mu_2 = 2.2190.$$

Here, $s^{(i)}$ as well as μ_i for $i \in [2]$ are readily obtained from the dual solution $(y^{(i)})_{i\in[2]}$ of the dual program (4.27) using the transformation from quadratic to ellipsoidal parametrization for anisotropic power diagrams (cf. Section 3.3.3). Figure 4.6a illustrates this result. Here, the jagged points mark the centroids of the clusters $c(C_1)$ (blue) and $c(C_2)$ (orange).

As the diagram is obviously feasible for ξ, Corollary 4.10 ensures us that $(\xi, (y^{(i)})_{i\in[k]})$ is indeed an optimal primal dual pair for the programs (4.26) and (4.27).

We can deduce that both the primal and the dual solutions are extremal. From Theorem 2.14 we immediately can conclude this for ξ: The corresponding clustering graph connects the two clusters by exactly 3 edges and thus does not contain a 4-connected subgraph. Here, the rank condition of Theorem 2.14 is obeyed as each subset of 3 points in X is affinely independent (i. e., no three points lie on a line). For the dual, we may fix $y^{(1)} = 0$ (as constraints associated with one cluster are redundant in a homogeneously constrained clustering polytope; cf. also Lemma 2.4). Then the fractionally assigned points $x^{(1)}, x^{(2)}, x^{(3)}$ uniquely determine $y^{(2)}$ (again due to affine independence).

However, due to $x^{(4)} \in P_1 \cap P_2$, the diagram does not support ξ. We already argued

that the dual solution (with $y^{(1)}$ fixed) is unique, so it must be contained in the relative interior of the respective optimal face. Thus, Corollary 4.9 yields that this cannot be the case for the primal solution.

Indeed, as the four points on the diagram's bisector are (minimally) affinely dependent, this means that the constraint matrix $A_{\cdot,[4]}$ has a null space of dimension one. Adding this null space direction to C_1 with appropriate sign (w. r. t. the component referring to $x^{(4)}$) and subtracting it from C_2, we can move along the one-dimensional optimal face of the clustering polytope T. Doing so, we obtain the second vertex of this face as $\xi' := \begin{pmatrix} C_1' & C_2' \end{pmatrix}^\intercal$ with $C_1' := \begin{pmatrix} \frac{5}{8} & 0 & 1 & \frac{7}{8} & 1 & 0 \end{pmatrix}^\intercal$ and $C_2' := \mathbb{1}^{(6)} - C_1'$. Consequently, $\frac{1}{2}(\xi + \xi')$ is a clustering in the relative interior of the optimal face and hence must be supported by the diagram $\mathcal{P}$. This is depicted in Fig. 4.6b.

Let us slightly adapt the example in order to make the optimal dual solution non-unique. We consider a third clustering

$$\xi'' = \begin{array}{c|cccccc} & x^{(1)} & x^{(2)} & x^{(3)} & x^{(4)} & x^{(5)} & x^{(6)} \\ \hline C_1'' & 1/2 & 1/2 & 1 & 1 & 1 & 0 \\ C_2'' & 1/2 & 1/2 & 0 & 0 & 0 & 1 \end{array}.$$

This results in cluster weights $\kappa_1'' = 4$, $\kappa_2'' = 2$ and centroids $c(C_1'') = (25/8, 3)^\intercal$, $c(C_2'') = (11/4, 39/8)^\intercal$. Figure 4.7a depicts this situation. As the diagram $\mathcal{P}$ is still feasible for ξ'', we can once more deduce optimality in the program (4.26) for an underlying adapted clustering polytope T''. Again, the diagram is obviously not supporting, but this time due to $x^{(3)}$ as well as $x^{(4)}$. By the same argumentation as before, ξ'' is an extremal point of T''. This time, however, it is the unique optimizer of the program (4.26). We could verify this by checking that the only direction in the kernel of $A_{\cdot,[4]}$ has different signs in the components corresponding to $x^{(3)}$ and $x^{(4)}$. However, this also becomes evident from the optimal dual solutions. Here, we find that the kernel of $A_{\cdot,[2]}^\intercal$ yields a direction $d \in \mathbb{R}^3 \setminus \{0\}$ such that $d^\intercal \left(A_{\cdot,\{3,4,5\}}\right) < 0$ and $d^\intercal \left(A_{\cdot,\{6\}}\right) > 0$. Thus, $y^{(2)} + \lambda d$, $\lambda \in \mathbb{R}_{\geq 0}$, yields an extreme ray that is the optimal dual face (assuming $y^{(1)} = 0$ to be fixed as before). Figure 4.7b depicts some of the resulting diagrams for increasing λ (drawn as dashed lines with the color fading out as λ grows) and corresponding sites (also fading out with increasing λ). We may interpret this as follows: The clustering ξ'' is supported by the power diagram with the bisector given as the affine hull of $x^{(1)}$ and $x^{(2)}$. As the blue ball is fully contained in the orange ellipse, we know from Section 3.3.3 that the anisotropic power diagram resulting from any choice of sites and additive weights will always result in the blue cell being an ellipse and the orange cell its surrounding. As the site $s^{(2)}$ is pushed away along $s^{(2)} + \lambda \cdot \Sigma_2^{-1} d$ with μ_2 adapted accordingly, this ellipse grows such that the resulting diagram converges to a power diagram (pointwise everywhere except for the middle point between $x^{(1)}$ and $x^{(2)}$).

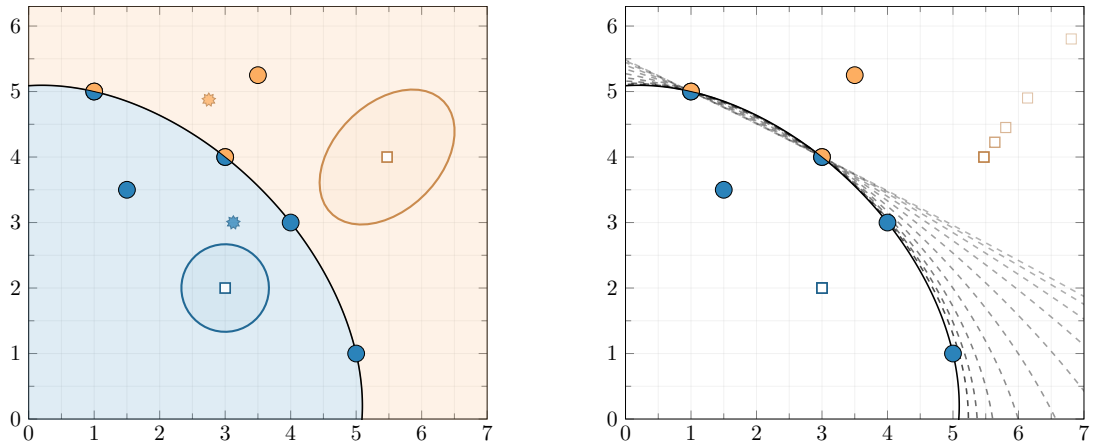

(a) Clustering ξ'' and the diagram obtained from the extremal dual solution.

(b) Diagrams obtained from the relative interior of the dual optimal face and corresponding sites.

Figure 4.7: Setting in Example 4.13 with non-unique diagram.

4.6 Summary & Conclusion

In this chapter we established the main correspondence between constrained clusterings and generalized Voronoi diagrams. This correspondence yields the main tool in our approach to constrained clustering.

We started by a precise definition of the desired relation between fractional clusterings and generalized Voronoi diagrams in Section 4.1. Here, the feasibility of a diagram yields that a unit assigned to a cluster also must lie in the respective cell. The supporting property then strengthens this relation by also demanding the reverse implication. Note that this distinction will be of practical impact in the remaining chapters.

This relation is not fundamentally new. A short review on the existing results was given in Section 4.2. In particular, the correspondence between power diagrams and weight-balanced clusterings is well-researched and will also play a major role in the remainder of this thesis.

The core relation can be understood as a linear programming duality as expressed in our core Theorem 4.7. In particular, the feasibility of a diagram can be formalized as a complementarity condition for the underlying dual pair of linear programs. The supporting property then furthermore requires this complementarity to be strict. As this can be traced back to the facets and their relative interior of the clustering polytope, it implies a geometric understanding of the proposed correspondence.

In order to provide a more plausible approach to the essence of this dual relation, we developed it step by step in Sections 4.3 to 4.5 (basically following the historical

order from Section 4.2). For the final correspondence, we assume feature maps of points to be given that determine the aggregation constraints of the clusters. To the best of our knowledge, the correspondence in this generality has not been stated yet. In order to stress the practical impact of this general version, we closed this chapter by the discussion of moment-constrained clusterings in Section 4.5.3. Here, we saw that the constraining of moments of clusters results in polynomial correction terms of according degree. In particular, this reveals the increased capabilities of anisotropic power diagrams compared to power diagrams. While the latter "only" allow weight constraints, the former may further be exploited to constrain centroids of clusters. However, this comes at the price of the more complex nature of anisotropic power diagrams as we had seen before in Chapter 3. Of course, there are many other kinds of constraints of practical relevance. In fact, any integral evaluations of functions over clusters are possible. This includes, for example, the Fourier coefficients of functions with support restricted to the clusters. As long as the resulting correction terms are conform with the chosen diagram class, this can offer a very broad scope of applications.

In the next chapter we will continue to generalize the established correspondence w. r. t. the space that is to be clustered. While this might be more of theoretical rather than practical interest, it will stress the simple yet powerful nature of this correspondence.

Chapter 5

Correspondence between Diagrams and Clusterings: General Space

We would like to extend our results to constrained clusterings of more general — in particular non-discrete — spaces. However, so far we have considered the finite case in which n units have to be (fractionally) assigned to k clusters. In particular, if we are interested in partitioning some subset of $\mathbb{R}^d$, we aim at a clustering of not n but infinitely many points that might furthermore be weighted by some continuous measure.

With the results of Chapter 4 in mind, a natural approach is to consider a series of refining discretizations and thus obtain continuous versions of the results for the discrete setting. For weight-constrained clusterings as considered in Section 4.4, such an approach is presented in [CCD16].

We will attempt a more general approach in this chapter. This leads us from an optimization problem over a constrained clustering polytope in $\mathbb{R}^{k\times n}$ to problems described in infinite-dimensional spaces that, for example, yield probability distributions over the plane. For weight-constrained clusterings, we end up with a semi-infinite transportation problem. Several authors have considered the problem of obtaining an optimal partition of a subset of $\mathbb{R}^d$ into sets of prescribed Lebesgue measures. Aurenhammer, Hoffmann, and Aronov [AHA98] showed the correspondence between least-square partitions and power diagrams as depicted in Section 4.3 for such a continuous setting by geometric arguments. Similarly, Geiß et al. [Gei+14] obtained a correspondence between generalized Voronoi diagrams and weight-constrained clusterings as treated in Section 4.4. Alternatively, a duality theory for the underlying semi-infinite transportation problem has been developed that implies the same results ([CR72a; CR72b; HK93; KY82; LH76; Tod78]). Linear programs in general topological vector spaces were taken into consideration very soon after the classical theory of linear programming had been developed (cf. [Duf56], [Hur14]). A very comprehensive introduction to linear programs in general topological vector spaces as well as a duality theory for semi-infinite transportation problems can be found in [AN87]. We will mainly follow this approach in order to convey our results of Section 4.5 into this very general setting.

However, this requires various aspects from topological vector spaces, functional analysis and measure theory. In order to provide a self-containing treatment and to prevent ambiguous terminology, the following Section 5.1 will state the necessary definitions and results. We will then formulate the constrained clustering problem in this more general setting in Section 5.2. Section 5.3 then argues to what extent the presented problems are solvable and strong duality can be re-established. While the desired correspondence between constrained clusterings and generalized Voronoi diagrams can not be fully established in general, Section 5.4 then shows how this can still be derived in an approximate manner.

5.1 Preliminaries

5.1.1 Definitions

Let us first recall and clarify some basic definitions from topology and functional analysis. Here, we mainly follow the terminology of [Kan60, Section III.3].

Topological Vector Spaces

A vector space equipped with some topology will be called *topological vector space* whenever both addition and scalar multiplication are continuous w. r. t. that topology. In particular, this implies that both translations and scalings of open sets are again open. To be more precise, let us state an equivalent definition via a basis of 0-neighborhoods as provided in [Kan60, Theorem 1]. We call a set V a *neighborhood* of x (or x-neighborhood) whenever V contains an open set that contains x.

Definition 5.1 (Topological Vector Space, cf. [Kan60])
Let X be a $\mathbb{R}$-vector space equipped with some topology. We say X is a *topological vector space* if there exists a *basis of* 0*-neighborhoods* $\mathcal{B} \subset 2^X$ (i. e., such that a subset U of X is open if and only for every $x \in U$ there exists $V \in \mathcal{B}$ with $x + V \subset U$) with the following properties:

- For any $V_1, V_2 \in \mathcal{B}$ there exists $V_3 \in \mathcal{B}$ such that $V_3 \subset V_1 \cap V_2$.
- Every $V \in \mathcal{B}$ is *balanced*, i. e., for all $x \in V$ it follows that $[-x, x] \subset V$.
- Every $V \in \mathcal{B}$ is *absorbent*, i. e., for all $x \in X$ there exists $\lambda > 0$ such that $[-\lambda x, \lambda x] \subset V$.
- For every $V \in \mathcal{B}$ there exists $U \in \mathcal{B}$ such that $U + U \subset V$.

Furthermore, we will require the considered topological vector spaces to be locally convex, in particular, as this will allow for separation theorems.

Definition 5.2 (Locally Convex Space, cf. [Kan60])
Let X be a $\mathbb{R}$-vector space equipped with some topology. We say X is a *locally convex space*, if

- X is a topological vector space,
- X is Hausdorff (or separated), i. e., for $x, y \in X$, $x \neq y$, there exist disjoint neighborhoods V_x and V_y of x and y, respectively, and
- X has a basis of convex 0-neighborhoods.

For any vector space X, we write X^+ for its algebraic dual. If X is a vector space, equipped with a topology $\mathcal{T}$, we write $(X, \mathcal{T})^*$ for its (topological) dual, or X^* whenever $\mathcal{T}$ is clear from the context.

The theory of linear programming in finite dimensions, of course, crucially relies on the separability of convex sets. Fortunately, locally convex spaces provide us with this tool, too. For any $f \in X^*$ we call $H_{(f,\beta)} := \{x \in X : f(x) = \beta\}$ a closed *hyperplane* and $H^{\leq}_{(f,\beta)} := \{x \in X : f(x) \leq \beta\}$ and $H^{<}_{(f,\beta)} := \{x \in X : f(x) < \beta\}$ a closed and open *halfspace*, respectively. Two non-empty sets $A, B \subset X$ are now *(strictly) separable* if there exists a hyperplane $H_{(f,\beta)}$ such that $A \subset H^{\leq}_{(f,\beta)}$ ($A \subset H^{<}_{(f,\beta)}$) and $B \subset H^{\geq}_{(f,\beta)}$ ($B \subset H^{>}_{(f,\beta)}$), respectively, hold.

The geometrical form of the famous Hahn-Banach theorem now yields separability of convex sets in this very general setting (cf. [SW99, Chapter II, Section 9]).

Theorem 5.3 ([SW99, Chapter II, Section 9, Second Separation Theorem]) *Let X be a locally convex space and $A, B \subset X$ two non-empty, disjoint convex sets such that A is closed and B is compact. Then A and B are strictly separable.*

As singletons are compact, we obtain the classical characterization of closed convex sets as a corollary:

Corollary 5.4 ([SW99, Chapter II, Section 9.2, Corollary 2]) *Let X be a locally convex space and $C \subset X, C \neq \emptyset$, be convex and closed. Let $\mathcal{H} := \left\{H^{\leq}_{(f,\alpha)} : C \subset H^{\leq}_{(f,\alpha)}\right\}$ be the set of all closed halfspaces containing C. Then*

$$C = \bigcap_{H^{\leq}_{(f,\alpha)} \in \mathcal{H}} H^{\leq}_{(f,\alpha)}.$$

Dual Pairings of Spaces

Next, in order to establish a duality theory, we require the notion of dual parings of spaces. Here, we follow [AN87, Section 3.2], [SW99, Chapter IV] and [Osb14, Section 3.6].

Let X be a vector space over $\mathbb{R}$. A subset Y of X^+ is *total* on X, if for all $x \in X \setminus \{0\}$ there is a $f \in Y$ such that $f(x) \neq 0$.

If Y is a total linear subspace of X^+, then $\langle X, Y\rangle$ is a *dual pairing*. In particular, this implies that $\langle Y, X\rangle$ is a dual pairing, too, by embedding X into Y^+ via the evaluation map that maps $x \mapsto g_x$ with $g_x(f) := f(x)$ for all $f \in Y$. In case of a topological vector space, $\langle X, X^*\rangle$ is such a pairing. A topology $\mathcal{T}$ on X is called *compatible* with the dual pairing $\langle X, Y\rangle$, whenever $(X, \mathcal{T})^* = Y$.

For any $A \subset X$, the *polar* of A is $A^\circ := \{f \in Y : |f(x)| \leq 1 \ \forall x \in A\}$. In this context,

the concept of polarity allows the definitions of several useful topologies (cf. [Osb14, Section 3.6]).

The *weak topology*, denoted $\sigma(X, Y)$, can be obtained via the basis of 0-neighborhoods given by the polars of all finite sets in Y, i. e., $\{A^\circ : A \subset Y, |A| < \infty\}$. In particular X equipped with the topology $\sigma(X, Y)$ is a locally convex vector space. For any topology that is compatible with $\langle X, Y\rangle$ and any $f \in Y$ it must hold that $\{f\}^\circ = f^{-1}([-1, 1])$ is a 0-neighborhood. Thus, it can be followed that $\sigma(X, Y)$ is the coarsest topology that is compatible with $\langle X, Y\rangle$ ([Osb14, p. 71]).

As a consequence of the Mackey-Arens theorem, there also exists a finest topology, called the *Mackey topology* and denoted by $\tau(X, Y)$, that is compatible with $\langle X, Y\rangle$ (cf. [SW99, Section 3.2]). It can be obtained via the basis of 0-neighborhoods $\{A^\circ : A \subset Y, A$ is $\sigma(Y, X)$-compact and absolutely convex$\}$. (Here, a set C is absolutely convex if for any $x_1, x_2 \in C$ and $\lambda_1, \lambda_2 \in \mathbb{R}$ with $|\lambda_1| + |\lambda_2| = 1$ it follows that $\lambda_1 x_1 + \lambda_2 x_2 \in C$.) Again, this makes $(X, \tau(X, Y))$ a locally convex space.

As a consequence of Corollary 5.4, we obtain the following standard result:

Corollary 5.5
Let $\langle X, Y\rangle$ be a dual pairing, $\mathcal{T}$ a locally convex, compatible topology , and $C \subset X$ be convex. Then the $\mathcal{T}$-closure of C is identical to the $\sigma(X, Y)$-closure of C.

Proof. This follows immediately from Corollary 5.4 as by compatibility the set of closed halfspaces that contain C (and thus its closure) are the same for both topologies. □

Corollary 5.5 will turn out to be very useful in the context of Slater-like duality conditions, as we will consider ($\sigma(X, Y)$-)closures of sets that have an empty $\sigma(X, Y)$, but non-empty $\tau(X, Y)$-interior. In particular, a locally convex space X such that its topology is the Mackey topology $\tau(X, X^*)$ is called a *Mackey space*. Important for us, metrizable and thus in particular normed spaces are Mackey spaces ([SW99, Section 3.4]).

The following theorem of Alaoglu-Bourbaki will help us to state solvability of our optimization problems, as it provides a connection of compactness to the boundedness of sets in the dual of normed spaces.

Theorem 5.6 (Theorem of Alaoglu-Bourbaki, [Kan60, Chapter III, Section 3, Theorem 7])
If X is a locally convex space, then the polar U° of each neighborhood of zero $U \subset X$ is $\sigma(X^*, X)$-compact.

For normed spaces, this yields following corollary:

Corollary 5.7
Let X be a normed space, then X^* is a normed space w. r. t. the operator norm and the unit ball $B^* = \{f \in X^* : |f(x)| \leq \|x\| \ \forall x \in X\}$ is $\sigma(X^*, X)$-compact.

Proof. As every linear functional on X is continuous (w. r. t. the norm-induced topology) if and only if it is bounded ([KA82, Section V.1, Theorem 1]), X^* together with the operator norm is indeed a normed space. Now observe that $B^* = B^\circ$ with B being the unit ball in X. So the second claim follows immediately from Theorem 5.6.□

5.1.2 Linear Programming in Paired Spaces

We are now able to establish a basic linear programming theory in paired spaces. We follow the notation and definitions of [AN87, Section 3.2].

In order to do so, let $\langle X, Y\rangle$ and $\langle Z, W\rangle$ be two dual pairings of spaces. Furthermore, let $c \in Y$, $b \in Z$ and $A : X \to Z$ be a continuous linear mapping (continuous w. r. t. the weak topologies $\sigma(X, Y)$ and $\sigma(Z, W)$).

While X will function as the *primal* (variable) space, the space Z will be referred to as the *constraint* space in the following.

Partial Order

Next, we equip our primal space with a partial order. In order to do so, we assume $P \subset X$ to be a convex cone. This defines a partial order "$\geq$" on X via

$$a \geq b :\Leftrightarrow a - b \in P$$

for all $a, b \in X$. In particular, $P = \{x \in X : x \geq 0\}$ holds by that definition. P also yields a partial order on the dual space Y via its dual

$$P^* := -P^\circ = \{y \in Y : \langle x, y\rangle \geq 0 \ \forall x \in P\}.$$

However, in order to stress the role of P, we will continue to write $x \in P$ instead of $x \geq 0$ in the following (as well as $y \in P^*$ instead of $y \geq 0$).

Linear Program and its Dual

We define a *linear program* in the given setting as

$$\min_{x \in X} \langle c, x\rangle \tag{5.1}$$
$$\text{s.t.} \quad Ax = b \tag{5.1a}$$
$$x \in P. \tag{5.1b}$$

The adjoint of A, i. e., $A^* : W \to Y$, maps $w \in W$ to the linear functional in Y, denoted A^*w, that is given by $x \mapsto \langle x, A^*w\rangle := \langle Ax, w\rangle$ (well defined as A is continuous w. r. t. the weak topology, cf. [SW99, Chapter IV, Section 2]).

We define the program

$$\max_{w\in W} \langle w, b\rangle \tag{5.2}$$
$$\text{s.t.} \quad c - A^*w \quad \in \quad P^* \tag{5.2a}$$

and call it the *dual* of (5.1).

In accordance with [AN87] we call the problem (5.1) *consistent*, whenever there exists a feasible solution, i. e., there exists $x \in P$ with $Ax = b$. The *value* of the program (5.1) is defined as

$$\text{value}((5.1)) := \inf\{\langle x, c\rangle : x \in P, Ax = b\}. \tag{5.3}$$

In particular, the program (5.1) is consistent if and only if $\text{value}((5.1)) < \infty$. We call the program (5.1) *solvable*, if there exists $x \in P$ with $Ax = b$ and $\langle x, c\rangle = \text{value}((5.1))$.

Note that the program (5.2) can be easily brought to the form of the program (5.1). In order to do so, we introduce slack variables $z \in Y$, replace the maximization expression by the equivalent $-\min_{w\in W, z\in Y}\langle w, -b\rangle$, and rewrite (5.2a) as $A^*w + z = c, (w, z) \in W \times P^*$. In particular, this implies that all terminology for programs in the form of (5.1) translates to programs in the form of (5.2) in a canonical way. In particular,

$$\text{value}((5.2)) = \sup\{\langle w, b\rangle : w \in W, c - A^*w \in P^*\}. \tag{5.4}$$

As a first observation, note that whenever P is closed, i. e., $P = (P^*)^*$, then the program (5.1) is the dual program of (5.2), too. In this case, $(W \times P^*)^* = \{0\} \times P$ as well as $A^{**} = A$ hold (both as W is total), so the latter follows from some standard reformulations.

Duality

In general, there is no strong duality between the programs (5.1) and (5.2). However, based on Theorem 5.3, the following generalized version of the Farkas-Minkowski Lemma establishes a slightly weaker relation. To the best of our knowledge, this was first established in similar versions in [Duf56], [Hur14] and [Kre61].

Theorem 5.8 (Farkas-Minkowski)
Let $\langle X, Y\rangle, \langle Z, W\rangle$ be dual pairings and $A : X \to Z$ a $\sigma(X, Y)$-$\sigma(Z, W)$-continuous linear mapping and $b \in Z$.

Set

$$Z_A := \{b \in Z : \exists x \in X : Ax = b, x \in P\} \qquad \text{and}$$
$$V_A := \{b \in Z : \forall w \in W : A^*w \in P^* \Rightarrow \langle b, w\rangle \geq 0\}.$$

Then

$$V_A = \mathrm{cl}_{\sigma(Z,W)}(Z_A).$$

Proof. We first show $Z_A \subset V_A$: Let $b \in Z_A$ and $w \in W$ with $A^*w \in P^*$. Then there exists $x \in P$ with $Ax = b$ and thus

$$\langle b, w\rangle = \langle Ax, w\rangle = \langle x, A^*w\rangle \overset{A^*w\in P^*, x\in P}{\geq} 0,$$

so $b \in V_A$ follows.

Next, we show that V_A is $\sigma(Z, W)$-closed. This follows immediately as $V_A = \bigcap_{w\in W: A^*w\in P^*} H^{\geq}_{(w,0)}$ and thus V_A is an intersection of weakly closed halfspaces. This implies $\mathrm{cl}_{\sigma(Z,W)}(Z_A) \subset V_A$.

For the final inclusion, assume there exists $b_0 \in V_A \setminus \mathrm{cl}_{\sigma(Z,W)}(Z_A)$. As $\{b_0\}$ is compact, Theorem 5.3 yields that there exists $w_0 \in W$ that separates $\{b_0\}$ and $\mathrm{cl}_{\sigma(Z,W)}(Z_A)$ strictly, i. e., there exists $\alpha \in \mathbb{R}$ such that $\langle b_0, w_0\rangle < \alpha$ and $\mathrm{cl}_{\sigma(Z,W)}(Z_A) \subset H^{\geq}_{(w_0,\alpha)}$. As $\mathrm{cl}_{\sigma(Z,W)}(Z_A)$ is a cone, it must hold that $\alpha = 0$ and so $\langle b_0, w_0\rangle < 0$. However, for any $x \in P$ it follows that $\langle x, A^*w_0\rangle = \langle Ax, w_0\rangle \overset{Ax\in Z_A}{\geq} 0$. Hence, $A^*w_0 \in P^*$ and so $b_0 \in V_A$ implies $\langle b_0, w_0\rangle \geq 0$, a contradiction. □

Note that Theorem 5.8 implies the classical Farkas-Minkowski Lemma: For $n, m \in \mathbb{N} \setminus \{0\}$, consider $X = Y = \mathbb{R}^n$, $Z = W = \mathbb{R}^m$ and a linear mapping identified with the matrix $A \in \mathbb{R}^{m\times n}$. Then Theorem 5.8 states that $b \in \mathbb{R}^m$ lies in the positive hull of the columns of A (which is closed as finitely generated) if and only if for all $w \in \mathbb{R}^m$ with $A^\intercal w \geq 0$ it follows that $b^\intercal w \geq 0$.

Let us compare Theorem 5.8 to its classical counterpart which can be applied in order to prove strong linear programming duality in finite dimensions. Thus, we define

$$Z_{A,c} := \left\{(\tilde{b}, \alpha) \in Z \times \mathbb{R} : \exists x \in X : \langle x, c\rangle \leq \alpha, Ax = \tilde{b}, x \in P\right\} \tag{5.5}$$

to be the cone of feasible right-hand sides and attainable objective values. This particularly implies

$$\mathrm{value}((5.1)) = \inf\{\alpha \in \mathbb{R} : (b, \alpha) \in Z_{A,c}\}. \tag{5.6}$$

With the classical proof of linear programming duality in mind, Theorem 5.8 motivates the idea that strong duality between the programs (5.1) and (5.2) will hold if $Z_{A,c}$ is

closed. As this is not the case in general, this motivates to define the *subvalue* of the program (5.1) as

$$\text{subvalue}((5.1)) := \inf\{\alpha \in \mathbb{R} : (b, \alpha) \in \text{cl}(Z_{A,c})\}. \tag{5.7}$$

Here, the closure is taken w. r. t. the product of the $\sigma(Z, W)$ topology and the standard topology in $\mathbb{R}$. We call the program (5.1) *subconsistent* whenever $\text{subvalue}((5.1)) < \infty$ holds.

The generalized Farkas-Minkowski-Lemma may now be used to prove the following theorem, which firstly states the weak duality between the programs (5.1) and (5.2) and secondly yields a strong duality result between the subvalue of (5.1) and the value of (5.2) (cf. [AN87, Theorem 3.5, Theorem 3.3]).

Theorem 5.9
It holds that

$$\text{value}((5.1)) \geq \text{subvalue}((5.1)) \geq \text{value}((5.2)) \tag{5.8}$$

and if not both $\text{subvalue}((5.1)) = \infty$ *and* $\text{value}((5.2)) = -\infty$ *it holds that*

$$\text{subvalue}((5.1)) = \text{value}((5.2)). \tag{5.9}$$

Proof. In order to apply Theorem 5.8, we set

$$V_{A,c} := \{(b, \alpha) \in Z \times \mathbb{R} : \forall w \in W, \beta \in \mathbb{R}_{\geq 0} : A^* w + c\beta \in P^* \Rightarrow \langle b, w\rangle + \alpha\beta \geq 0\}. \tag{5.10}$$

We may write $Z_{A,c} = \left\{(b, \alpha) \in Z \times \mathbb{R} : \exists x \in X, \beta \in \mathbb{R} : \tilde{A}(x, \beta) = \tilde{b}, x \in P, \beta \geq 0\right\}$ with $\tilde{A}(x, \beta) := \begin{pmatrix} \langle x, c\rangle + \beta \\ Ax \end{pmatrix}$ and $\tilde{b} := \begin{pmatrix} \alpha \\ b \end{pmatrix}$ and obtain $\tilde{A}^*(w, \beta) = A^* w + c\beta$ as well as $(P \times \mathbb{R}_{\geq 0})^* = P^* \times \mathbb{R}_{\geq 0}$.

Therefore, Theorem 5.8 yields $V_{A,c} = \text{cl}(Z_{A,c})$.

Then we get

$$\begin{aligned}
\text{value}((5.2)) &= \sup\{\alpha \in \mathbb{R} : \exists w \in W : -A^* w + c \in P^*, \alpha = \langle b, w\rangle\} \\
&= \inf \underbrace{\{\alpha \in \mathbb{R} : \forall w \in W : A^* w + c \in P^* \Rightarrow \langle b, w\rangle + \alpha \geq 0\}}_{=:B_D} \\
&\leq \inf \underbrace{\{\alpha \in \mathbb{R} : \forall w \in W, \beta \geq 0 : A^* w + \beta c \in P^* \Rightarrow \langle b, w\rangle + \alpha\beta \geq 0\}}_{=:B_P} \\
&= \inf\{\alpha \in \mathbb{R} : (b, \alpha) \in V_{A,c}\} = \inf\{\alpha \in \mathbb{R} : (b, \alpha) \in \text{cl}(Z_{A,c})\} \\
&= \text{subvalue}((5.1)).
\end{aligned}$$

Here, note that the first equality is particularly true for the case that the problem (5.24) is inconsistent as $\sup(\emptyset) = -\infty$ holds by convention. The inequality follows as the choice $\beta = 1$ in the definition of B_P yields all constraints on α in the definition of B_D.

The first inequality in Eq. (5.8) is clear by the definitions of value and subvalue, so Eq. (5.8) is shown.

Now assume $\inf(B_D) < \inf(B_P)$. Then there exist $\alpha_0 \in B_D \setminus B_P$ and consequently $w_0 \in W, \beta_0 \geq 0$ such that $A^*w_0 + c\beta_0 \in P^*$ and $\langle b, w_0\rangle + \alpha_0\beta_0 < 0$.

If $\beta_0 > 0$ we get $A^*(\frac{1}{\beta_0}w_0) + c \in P^*$ and $\langle b, \frac{1}{\beta_0}w_0\rangle + \alpha_0 < 0$, a contradiction to $\alpha_0 \in B_D$.

If $\beta_0 = 0$, this means $A^*w_0 \in P^*$ and $\langle b, w_0\rangle < 0$. As this particularly implies $\langle b, w_0\rangle + \beta_0\alpha < 0$ for all $\alpha \in \mathbb{R}$, it follows that $B_P = \emptyset$ and thus subvalue((5.1)) $= \infty$. Then by assumption value((5.2)) $> -\infty$ holds, i. e., the program (5.2) must be consistent. Thus, there exists $w_1 \in W$ with $-A^*w_1 + c \in P^*$ and therefore $-A^*(w_1 - \lambda w_0) + c \in P^*$ holds for all $\lambda > 0$. However, from $\langle b, w_1 - \lambda w_0\rangle \to \infty$ (for $\lambda \to \infty$), we get value((5.2)) $= \infty$, a contradiction to value((5.2)) $= \inf(B_D) < \inf(B_P)$. This completes the proof. □

Most importantly, Theorem 5.9 implies that whenever $Z_{A,c} = \mathrm{cl}(Z_{A,c})$, strong duality between (5.1) and (5.2) follows, as then value((5.1)) = subvalue((5.1)) holds.

For the dual program (5.2) we may set

$$Y_{A^*,b} = \{(\tilde{c}, \delta) \in Y \times \mathbb{R} : \exists w \in W : \langle b, w\rangle \geq \delta, \tilde{c} - A^*w \in P^*\} \tag{5.11}$$

and obtain

$$\text{subvalue((5.2))} = \sup\{\delta : (c, \delta) \in \mathrm{cl}(Y_{A^*,b})\}. \tag{5.12}$$

In the case of P being closed, i. e., when the program (5.1) is the dual of (5.2), Theorem 5.9 immediately yields the following corollary.

Corollary 5.10
If P is closed, it holds that

$$\text{value((5.2))} \leq \text{subvalue((5.2))} \leq \text{value((5.1))}$$

and if furthermore not both subvalue((5.2)) $= -\infty$ *and* value((5.1)) $= \infty$ *it holds that*

$$\text{value((5.2))} \leq \text{subvalue((5.2))} = \text{value((5.1))}.$$

If $Z_{A,c}$ is not closed, the value and subvalue of the program (5.1) are not equal in general (see [AN87] for examples) and strong duality does not hold. However,

several conditions have been found which yield closure of $Z_{A,c}$ (cf., for example,[Kre61], [Fan65], [Yam68], [NY79], [AN87]). The following Slater-like condition was established in [Kre61] (however, with an incomplete proof that is corrected in [Yam68]). We provide the statement in our setting and a proof that combines ideas from [Kre61] and [AN87].

Theorem 5.11
Let P be closed and let $w_0 \in W$ such that $c - A^*w_0$ lies in the $\tau(Y,X)$-interior of P^*. Then $Z_{A,c}$ is closed.

Proof. Let $(b_0, \alpha_0) \in V_{A,c} = \mathrm{cl}(Z_{A,c})$ as defined by Eq. (5.10) and assume $(b_0, \alpha_0) \notin Z_{A,c}$. Here, note that it does not matter whether we take the closure w. r. t. the Mackey topology or the weak topology due to Corollary 5.5. Hence, in the following we will always assume to work in the Mackey topology.

We define the auxiliary cone

$$\hat{Y} := \{(\tilde{c}, \delta) \in Y \times \mathbb{R} : \exists w \in W, \beta \geq 0 : \tilde{c} + A^*w + \beta c \in P^*, \delta - \langle b_0, w\rangle - \beta\alpha_0 = 0\}.$$

With the mapping $\tilde{A} : W \times Y \times \mathbb{R} \to Y \times \mathbb{R}$ given by

$$\tilde{A}(w, y, \beta) := (y - A^*w - \beta c, \langle b_0, w\rangle + \beta\alpha_0)$$

and the cone $\tilde{P} := W \times P^* \times \mathbb{R}_{\geq 0}$, it then holds that

$$\hat{Y} = \Big\{(\tilde{c}, \delta) \in Y \times \mathbb{R} : \exists (w, y, \beta) \in \tilde{P} : \tilde{A}(w, y, \beta) = (\tilde{c}, \delta)\Big\}.$$

Furthermore, with $\tilde{P}^* = \{0\} \times P \times \mathbb{R}_{\geq 0}$ (using that W is total as well as P is closed and hence $P^{**} = P$) we get

$$\begin{aligned}
\hat{V} &:= \Big\{(\tilde{c}, \delta) \in Y \times \mathbb{R} : \forall (x, \gamma) \in X \times \mathbb{R} : \tilde{A}^*(x, \gamma) \in \tilde{P}^* \Rightarrow \langle x, \tilde{c}\rangle + \gamma\delta \geq 0\Big\} \\
&= \{(\tilde{c}, \delta) \in Y \times \mathbb{R} : \forall (x, \gamma) \in X \times \mathbb{R} : \\
&\qquad\qquad - Ax + \gamma b_0 = 0 \wedge x \in P \wedge -\langle x, c\rangle + \gamma\alpha_0 \geq 0 \Rightarrow \langle x, \tilde{c}\rangle + \gamma\delta \geq 0\} \\
&= \{(\tilde{c}, \delta) \in Y \times \mathbb{R} : \forall x \in P, \gamma \in \mathbb{R} : Ax = \gamma b_0 \wedge \langle x, c\rangle \leq \gamma\alpha_0 \Rightarrow \langle x, \tilde{c}\rangle + \gamma\delta \geq 0\}.
\end{aligned}$$

Now, Theorem 5.8 yields $\mathrm{cl}(\hat{Y}) = \hat{V}$.

Assume there exist $x \in P$ and $\gamma > 0$ with $Ax = \gamma b_0$ and $\langle x, c\rangle \leq \gamma\alpha_0$. Then we set $\tilde{x} := \frac{1}{\gamma}x \in P$ and obtain $A\tilde{x} = b_0$ and $\langle \tilde{x}, c\rangle \leq \alpha_0$, and thus $(b_0, \alpha_0) \in Z_{A,c}$, which contradicts the assumption. Consequently, we get

$$\hat{V} = \{(\tilde{c}, \delta) \in Y \times \mathbb{R} : \forall x \in P, \gamma \leq 0 : Ax = \gamma b_0, \langle x, c\rangle \leq \gamma\alpha_0 \Rightarrow \langle x, \tilde{c}\rangle + \gamma\delta \geq 0\}.$$

In particular, we can easily check that $(0, -1) \in \hat{V} = \mathrm{cl}(\hat{Y})$.

By assumption there exist $w_0 \in W$ and a 0-neighborhood $U \subset Y$ such that $c - A^* w_0 + U \subset P^*$. As Z is a topological vector space, we have that $\frac{1}{n}U$ is a 0-neighborhood for every $n \geq 1$. We may furthermore w.l.o.g. assume that U is balanced (otherwise, we replace it by an element of a fundamental 0-neighborhood that is contained in U).

As $(0, -1) \in \mathrm{cl}(\hat{Y})$, for every $n \geq 1$ there exist $w^{(n)} \in W$, $y^{(n)} \in Y$ and $\beta^{(n)} \in \mathbb{R}_{\geq 0}$ such that

$$y^{(n)} \in \frac{1}{n}U, \;\; y^{(n)} + A^* w^{(n)} + \beta^{(n)} c \in P^*, \text{ and } \;\; \langle b_0, w^{(n)} \rangle + \beta^{(n)} \alpha_0 \in (-1 - \frac{1}{n}, -1 + \frac{1}{n}).$$

We will now achieve a contradiction by slightly moving the elements $(w^{(n)}, \beta^{(n)})$ towards $(-w_0, 1)$ and thus obtain a feasible constraint in the definition of $V_{A,c}$ in Eq. (5.10). Hence, we set for every $n \geq 1$

$$\hat{w}^{(n)} := \frac{n}{n+1} w^{(n)} + \frac{1}{n+1}(-w_0) \quad \text{and} \quad \hat{\beta}^{(n)} := \frac{n}{n+1}\beta^{(n)} + \frac{1}{n+1}.$$

This yields

$$\langle b_0, \hat{w}^{(n)} \rangle + \hat{\beta}^{(n)} \alpha_0 = \frac{n}{n+1} \underbrace{\left(\langle b_0, w^{(n)} \rangle + \beta^{(n)} \alpha_0 \right)}_{\to -1 \;\; (n \to \infty)} + \frac{1}{n+1} \left(\langle b_0, -w_0 \rangle + \alpha_0 \right) \overset{n \to \infty}{\longrightarrow} -1.$$

Furthermore, it holds that

$$\begin{aligned} A^* \hat{w}^{(n)} + \hat{\beta}^{(n)} c &= \frac{n}{n+1} \left(A^* \hat{w}^{(n)} + \hat{\beta}^{(n)} c \right) + \frac{1}{n+1} \left(c - A^* w_0 \right) \\ &= \frac{n}{n+1} \underbrace{\left(A^* \hat{w}^{(n)} + \hat{\beta}^{(n)} c + y^{(n)} \right)}_{\in P^*} + \frac{1}{n+1} \underbrace{\left(c - A^* w_0 - n y^{(n)} \right)}_{\in c - A^* w_0 - U \subset P^*} \in P^*. \end{aligned}$$

Here, we used $U = -U$ as U is balanced by assumption.

However, as $(b_0, \alpha_0) \in V_{A,c}$ this implies $\langle b_0, \hat{w}^{(n)} \rangle + \alpha_0 \hat{\beta}^{(n)} \geq 0$ for every n, a contradiction to the limit above. □

We mention that Anderson and Nash [AN87] provide a similar statement showing that value((5.2)) = subvalue((5.2)) (without implying closure of either $Z_{A,c}$ or $Y_{A^*,b}$) under the same conditions except the requirement of P being closed ([AN87, Theorem 3.11, Theorem 3.13]).

5.1.3 Duality of Measures and Continuous Functions

As a final ingredient, we need to prepare the choice of topological spaces that we will eventually use to bring the program (4.26) into a program in the style of (5.1).

In the discrete setting, each cluster of a fractional clusterings can be understood as a conditional discrete probability distribution on the points to be clustered. In a more general setting, a cluster will thus be a measure on some given space. In order to prepare for this, let us recall a version of the famous Riesz Representation theorem that allows to identify the set of continuous functions on a compact Hausdorff space with the dual of the space of signed measures.

For the remainder of this section, let X be a compact Hausdorff space (which the reader may think of as the space containing the points to be clustered). As several formulations of the Riesz Representation theorem exist, we will make precise the one we will refer to. In particular, definitions vary slightly in the literature (e. g., [RF10] requires a Borel measure of a compact set to be finite while [Coh13] does not), so we will state the definitions that are crucial for our context.

Definition 5.12 (Signed Regular Borel Measure, cf. [Coh13, Section 7.2], [RF10, Section 21.3])
Let X be a Hausdorff space and $\mathcal{B}(X)$ be the Borel σ-algebra.

A measure ω on $(X, \mathcal{B}(X))$ with $\omega(K) < \infty$ for every compact set $K \subset X$ is called *Borel measure*.

A Borel measure ω is *regular* if for all $A \in \mathcal{A}$ it holds that

$$\omega(A) = \inf\{\omega(U) : A \subset U, U \text{ open}\}$$

and for each open subset $U \subset X$ it holds that

$$\omega(U) = \sup\{\omega(K) : K \subset U, K \text{ compact}\}.$$

A regular Borel measure is also called *Radon measure*.

If X is compact, we say that ω is a *signed regular Borel measure* (*signed Radon measure*) if there are regular Borel measures ω^+, ω^- on X, and $X^+, X^- \in \mathcal{B}(X)$ with $X^+ \cap X^- = \emptyset$, $X = X^+ \cup X^-$, such that $\omega^+(X^-) = 0$, $\omega^-(X^+) = 0$ and $\omega = \omega^+ - \omega^-$.

Note that the restriction to compact spaces for a signed measure will be sufficient for our context and avoids dealing with a more general definition of signed measures on non-compact spaces.

We will denote

$$M_r(X) := \{\omega : \omega \text{ is a signed regular Borel measure on } X\}.$$

Then $M_r(X)$ is a vector space, and equipped with the norm of total variation

$$\|\omega\|_{\text{var}} := \omega^+(X) + \omega^-(X)$$

with ω^+ and ω^- as in Definition 5.12 (uniquely determined by the Jordan Decomposition theorem, cf. [RF10, Section 17]), this is a normed vector space (and thus in particular a locally convex space w. r. t. the norm-induced topology).

The second space we consider is that of real-valued continuous functions over X, i. e.,

$$C(X) := \Big\{ f \in \mathbb{R}^X : f \text{ continuous} \Big\}.$$

Equipped with maximum-norm

$$\|f\|_{\max} := \max_{x \in X} |f(x)|$$

it is a normed vector space, too.

Note that every continuous function is Borel-measurable and in particular

$$\langle f, \omega \rangle := \int_X f(x) d(\omega(x))$$

is well-defined.

Theorem 5.13 (Riesz Representation Theorem for the Dual of $C(X)$ ([RF10, Section 21.5]))
Let X be a compact Hausdorff space.

Then the mapping $\psi : M_r(X) \to C(X)^*$ that maps ω to the linear functional $\psi_\omega(\cdot) :=< \cdot, \omega >$ is a linear isometric isomorphism of $M_r(X)$ onto $C(X)^*$.

Here, $C(X)^*$ is a normed space using the operator norm (cf. Corollary 5.5). In other words, we may identify the dual space of $C(X)$ with $M_r(X)$. In particular, as the described isomorphism is isometric for any $\omega \in M_r(X)$ it holds that

$$\|\omega\|_{\text{var}} = \sup_{f \in C(X) : \|f\|_{\max} \leq 1} \left| \int_X f(x) d(\omega(x)) \right|.$$

Also, this implies that $\langle M_r(X), C(X) \rangle$ is a dual pairing.

Later on, we will define clusters as non-negative measures. Thus, let

$$P := \{ \omega \in M_r(X) : \omega(A) \geq 0 \ \forall A \in \mathcal{B}(X) \} \tag{5.13}$$

be the cone of all non-negative signed regular Borel measures.

Then the dual cone of P is given by all non-negative continuous functions:

Lemma 5.14
With P as defined by Eq. (5.13) it holds that $P^* = \{ f \in C(X) : f(x) \geq 0 \ \forall x \in X \}$.

Proof. The inclusion "$\supset$" is clear from the definition of P^*. The other inclusion follows as for every $f \in P^*$ and $x \in X$, the Dirac measure δ_x is a regular Borel measure and $\int_X f(\tilde{x}) d\delta_x(\tilde{x}) = f(x)$. □

Furthermore, it holds that P is closed (and hence in particular $P^{**} = P$). In order to provide a proof for this, we use the following result:

Theorem 5.15 (Lusin's Theorem [Coh13, Theorem 7.4.3])
Let X be a locally compact Hausdorff space, ω a regular Borel measure and $f : X \to \mathbb{R}$ measurable.

For $A \in \mathcal{B}(X) \setminus \{\emptyset\}$ with $\omega(A) < \infty$ and $\epsilon > 0$ there exists a compact subset $K \subset A$ such that $\omega(A \setminus K) < \epsilon$ and such that $f|_K$ is continuous.

Furthermore, there exists a compactly supported continuous function $g : X \to \mathbb{R}$ such that $g|_K = f|_K$ and $\sup\{|g(x)| : x \in X\} \leq \sup\{|f(x)| : x \in X\}$.

We conclude the closure of P:

Corollary 5.16
Let X be a compact Hausdorff space and P as given by Eq. (5.13).

Then P is closed (w. r. t. the $\sigma(M_r(X), C(X))$ topology).

Proof. Let $\omega \in M_r(X) \setminus P$. Then there are regular Borel-measures ω^+, ω^- of disjoint support such that $\omega = \omega^+ - \omega^-$.

By assumption there exists a measurable set A such that $\omega^-(A) > 0$ and $\omega^+(A) = 0$. We approximate the indicator function $\mathbb{1}_A$. Using Lusin's Theorem (Theorem 5.15) applied to the regular Borel measure $\omega^+ + \omega^-$, we can find a compact set $K \subset A$ such that $\omega^-(X \setminus K) + \omega^+(X \setminus K) < \frac{\omega^-(A)}{3}$ and a continuous function g such that $g|_K = (\mathbb{1}_A)|_K$ and $\|g\|_\infty \leq 1$. Now we may define $\tilde{g}(x) := \max\{g(x), 0\}$, so $\tilde{g} \in P^*$ (as given by Lemma 5.14), i. e., $\tilde{g}$ is a non-negative continuous function. Then

$$\begin{aligned}\langle \tilde{g}, \omega \rangle &= \langle \mathbb{1}_{X \setminus K} \tilde{g}, \omega \rangle + \langle \mathbb{1}_K \tilde{g}, \omega \rangle \leq \omega^+(X \setminus K) + \omega^-(X \setminus K) + \langle \mathbb{1}_K, \omega \rangle \\ &\leq \omega^+(X \setminus K) + \omega^-(X \setminus K) - \omega^-(A) \\ &< \frac{2}{3}\omega^-(A) - \omega^-(A) < 0.\end{aligned}$$

Consequently, $\tilde{g}$ separates ω strictly from the cone

$$P^{**} = \{\tilde{\omega} \in M_r(X) : \langle f, \tilde{\omega} \rangle \geq 0 \forall f \in P^*\}.$$

Hence, $P^{**} \subset P$ and so $P^{**} = P$. In particular, this implies that P is closed. □

5.2 Constrained Clusterings in General Space

We can now transfer our previous theory into this more general setting. For the remainder of this chapter, we assume X to be a compact Hausdorff space. This is the space we assume to be clustered. Furthermore, we assume $\omega_0 \in M_r(X)$ to be a regular Borel measure on X.

As pointed out in [AN87] for semi-infinite transportation problems, this setting is crucial. Depending on that choice, consistency, solvability and the relation of the resulting linear programs are affected. In particular, as $\langle M_r(X), C(X)\rangle$ is a dual pairing, this dictates the type of cluster constraints in the following.

5.2.1 Definitions

First, let us transfer the definitions and notations of Chapter 2 into this (more general) setting. Here, a clustering is now given as a tuple of regular Borel measures that adds up to ω_0. We then constrain clusterings w. r. t. the integral values of continuous feature maps.

Definition 5.17 (Fractional Clustering (in a compact Hausdorff space))
Let X be a compact Hausdorff space and ω_0 a regular Borel measure on X. Let $k \in \mathbb{N} \setminus \{0\}$ and for every $i \in [k]$ let $m_i \in \mathbb{N}$ and functions $\Phi_{i,l} \in C(X)$ for $l \in [m_i]$ as well as $b^{(i)} \in \mathbb{R}^{m_i}$ be given.

We set

$$
\begin{aligned}
T &:= T\left(\omega_0, (\Phi_{i,l})_{l\in[m_i],i\in[k]}, (b^{(i)})_{i\in[k]}\right) && \text{(5.13a)}\\
&:= \{(\xi_i)_{i\in[k]} \in M_r(X)^k : \\
&\qquad \sum_{i\in[k]} \xi_i \quad = \quad \omega_0 && \text{(5.13b)}\\
&\qquad \langle \phi_{i,l}, \xi_i\rangle \quad = \quad b_l^{(i)} \qquad \forall i \in [k], l \in [m_i]\}. && \text{(5.13c)}
\end{aligned}
$$

Then we call $(\xi_i)_{i\in[k]} \in T$ *a* $(\omega_0, (\Phi_{i,l})_{l\in[m_i],i\in[k]}, (b^{(i)})_{i\in[k]})$*-constrained (fractional) clustering.*

For each $i \in [k]$ we call the measure ξ_i (the ith) *cluster*. For each cluster ξ_i its support is given (in the usual way) by

$$
\operatorname{supp}(\xi_i) := X \setminus \bigcup_{U \text{open},\ \xi_i(U)=0} U. \tag{5.14}
$$

If the measures ξ_i are mutually singular, i. e., there exist pairwise disjoint measurable sets $X_1, \ldots, X_k$, such that $\xi_i(X_i) = \xi_i(X)$ for all $i \in [k]$, we call the clustering *integer*.

Next, the terminology of feasible and supporting generalized Voronoi Diagrams may be transferred with almost no adaption.

Definition 5.18
Let X be a compact Hausdorff space, $\mathcal{P} = (P_1, \dots, P_k)$ be a generalized Voronoi Diagram in X, and a set of constrained clusterings $T = T\left(\omega_0, (\Phi_{i,l})_{l\in[m_i],i\in[k]}, (b^{(i)})_{i\in[k]}\right)$ be given. Let $\xi \in T$.

We say $\mathcal{P}$ is *feasible* for ξ, if and only if

$$\operatorname{supp}(\xi_i) \subset P_i \ \forall i \in [k] \tag{5.15}$$

and $\mathcal{P}$ *supports* ξ, if and only if

$$\operatorname{supp}(\xi_i) = P_i \cap \operatorname{supp}(\omega_0) \ \forall i \in [k]. \tag{5.16}$$

5.2.2 Program Formulations

We can now transfer the program (4.26) and its dual to this setting. In order to do so, let $T = T\left(\omega_0, (\Phi_{i,l})_{l\in[m_i],i\in[k]}, (b^{(i)})_{i\in[k]}\right)$ as in Definition 5.17 be given. Furthermore, let functions $f_i \in C(X)$ for $i \in [k]$ be given.

We obtain the following primal problem:

$$\min_{(\xi_i)_i \in M_r(X)^k} \sum_{i=1}^{k} \langle f_i, \xi_i \rangle \tag{5.17}$$

$$\text{s.t.} \quad \sum_{i\in[k]} \xi_i = \omega_0 \tag{5.17a}$$

$$\langle \phi_{i,l}, \xi_i \rangle = b_l^{(i)} \quad \forall i \in [k], l \in [m_i] \tag{5.17b}$$

$$\xi_i \geq 0 \quad \forall i \in [k] \tag{5.17c}$$

Let us provide some sanity checks: Our primal solution space is $M_r(X)^k$. Note that we can identify its dual space by the linear functionals given as sum of functionals from the dual of $M_r(X)$, i. e., $C(X)$ (cf. [SW99, Chapter I, Section 2]).

The left-hand sides of the constraints (5.17a) (which yield mappings $M_r(X)^k \to M_r(X)$) are indeed linear and continuous (w. r. t. the box topology obtained from $\sigma(M_r(X), C(X))$), as $(M_r(X), \sigma(M_r(X), C(X)))$ is a topological vector space. The constraints (5.17b) are linear and continuous (w. r. t. the weak topology), too, as we assume $\phi_{i,l} \in C(X)$ for every $i \in [k], l \in [m_i]$. Note that constraint (5.17c) implies that the associated positive cone is given by $(P)^k$ with P as given by Eq. (5.13).

We conclude that the program (5.17) is indeed of the form (5.1).

Next, we formulate the dual of (5.17).

The constraint space (i. e., the possible right hand sides) is $M_r(X) \times \mathbb{R}^{m_1} \times \ldots \times \mathbb{R}^{m_k}$. Consequently, the variable space of the dual of (5.17) is $C(X) \times \mathbb{R}^{m_1} \times \ldots \times \mathbb{R}^{m_k}$.

Hence, we obtain the dual program

$$\max_{\eta \in C(X), y^{(i)} \in \mathbb{R}^{m_i}, i \in [k]} \sum_{i=1}^{k} (y^{(i)})^\intercal b^{(i)} + \langle \eta, \omega_0 \rangle \tag{5.18}$$

$$\text{s.t.} \qquad \eta + (y^{(i)})^\intercal \phi_i \quad \leq \quad f_i \quad \forall i \in [k] \tag{5.18a}$$

Here, $\phi_i := (\phi_{i,l})_{l \in [m_i]}$ is to denote a vector of functions. This yields a continuous function $X \to \mathbb{R}^{m_i}$ for every $i \in [k]$. Recall that from Theorem 5.13 we know that $C(X)^*$ can be identified with $M_r(X)$. Furthermore, from Corollary 5.16 we know that P is closed, or — equivalently — $P^{**} = P$. In particular, this implies that the program (5.17) is the dual program of (5.18).

5.3 Solvability and Duality

Next, let us discuss solvability as well as the duality properties of the programs (5.17) and (5.18).

We apply several results of Section 5.1 in order to obtain the following theorem:

Theorem 5.19
Assume problem (5.17) *to be consistent. Then the following holds:*

i) The program (5.17) *is solvable.*

ii) The program (5.18) *is consistent with finite value.*

iii) There is no duality gap, i. e., value((5.17)) = value((5.18)).

Proof. This proof partially follows [AN87, Section 5.5].

For part i, let $(\xi_i)_{i \in [k]}$ be feasible for the program (5.17). Then for any $i \in [k]$ it follows that

$$\sup_{f \in C(X): \|f\|_\infty \leq 1} \left| \int_X f(x) d\xi_i(x) \right| = \|\xi_i\|_{\text{var}} = \xi_i(X) \leq \omega_0(X) < \infty.$$

By Corollary 5.7, this implies that all feasible clusters ξ_i are contained in a $\sigma(M_r(X), C(X))$-compact set. As the objective is (by choice) $\sigma(M_r(X), C(X))$-continuous, this implies solvability.

For part ii, a feasible solution is obtained by setting $y^{(i)} := 0$ for $i \in [k]$ and η to be constant with

$$\eta := \min_{i \in [k], \tilde{x} \in X} f_i(\tilde{x}) - 1 \tag{5.19}$$

(which is well-defined as we assume X to be compact and $f_i \in C(X)$ for every $i \in [k]$). By weak duality and as we assume the program (5.17) to be consistent, the optimal value of the program (5.18) must be finite.

For part iii, recall that the cone of regular Borel measures on X is $\sigma(M_r(X), C(X))$-closed by Corollary 5.16. Thus, it is sufficient to show value((5.18)) = subvalue((5.18)) due to Corollary 5.10.

As $C(X)$ is a normed space, it is a Mackey space with the Mackey topology being the norm induced topology. With $y^{(i)} = 0$ for $i \in [k]$ and η defined by Eq. (5.19), we get that $f_i(x) - \eta(x) - (y^{(i)})^\intercal \phi_i(x) \geq 1$ for all $x \in X$. Hence, we can apply Theorem 5.11 which yields the claim. □

A crucial missing point in Theorem 5.19 is the solvability of the program (5.18). However, as the following example shows, this cannot be guaranteed in general for our current setting.

Example 5.20
Let $X = [0,1] \subset \mathbb{R}$ and $k = 2$. Let $f_1, f_2 \in C([0,1])$ be constant functions with $f_1 = 0$ and $f_2 = 1$.

Furthermore, let $m_1 := m_2 := 1$, set $\phi_{1,1} := \phi_{2,1} := \phi :=$ id, i. e., $\phi(x) = x$ for all $x \in [0,1]$, and $b^{(1)} := 0$, $b^{(2)} := 1/2$.

Finally, denote by $\lambda_{[0,1]}$ the Lebesgue measure on the interval $[0,1]$. We set $\omega_0 := \lambda_{[0,1]}$.

With $\int_{[0,1]} f_1 d\xi_1(x) = 0$ and $\int_{[0,1]} f_2 d\xi_2(x) = \xi_2([0,1])$, problem (5.17) simplifies to

$$\min_{\xi_1,\xi_2 \in M_r([0,1])} \xi_2([0,1]) \tag{5.20}$$

$$\text{s.t.} \quad \xi_1 + \xi_2 = \lambda_{[0,1]} \tag{5.20a}$$

$$\int_{[0,1]} x \; d\xi_1(x) = 0 \tag{5.20b}$$

$$\int_{[0,1]} x \; d\xi_2(x) = \frac{1}{2} \tag{5.20c}$$

$$\xi_1, \xi_2 \geq 0 \quad . \tag{5.20d}$$

From constraint (5.20b) it follows that no feasible ξ_1 can have a non-zero absolutely continuous part. Thus, we can deduce from constraint (5.20a) that $\xi_1 = 0$ and $\xi_2 = \lambda_{[0,1]}$ is the only possible solution, which is consistent with the (redundant) constraint (5.20c). So, the program (5.20) is solvable with value 1.

However, we can easily see that for no $y^{(1)}, y^{(2)} \in \mathbb{R}$, the generalized Voronoi diagram w. r. t. the functions $\hat{f}_i = f_i - y^{(i)}\phi_{i,1}$ will be feasible for (ξ_1, ξ_2). As depicted in Fig. 5.1, only in the limit, i. e., either for $y^{(1)} \to -\infty$ or $y^{(2)} \to \infty$, the resulting diagrams become feasible.

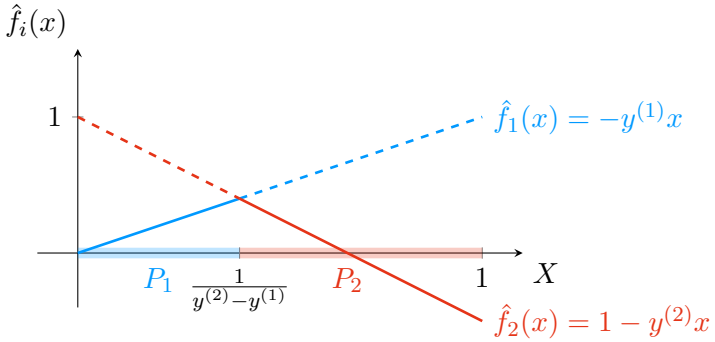

Figure 5.1: Exemplary generalized Voronoi diagram functions in Example 5.20 for $y^{(1)} = -1$ and $y^{(2)} = \frac{3}{2}$.

This also becomes clear when considering the dual. For any $\eta \in C([0,1])$ it holds that $\int_{[0,1]} \eta(x) d\omega_0(x) = \int_0^1 \eta(x) d\lambda(x)$. Hence, for this instance problem (5.18) becomes

$$\max_{\eta \in C([0,1]), y^{(1)}, y^{(2)} \in \mathbb{R}} \frac{1}{2} y^{(2)} + \int_0^1 \eta(x) d\lambda(x) \tag{5.21}$$

$$\text{s.t.} \quad \eta(x) + y^{(1)}x \leq 0 \quad \forall x \in [0,1] \tag{5.21a}$$

$$\eta(x) + y^{(2)}x \leq 1 \quad \forall x \in [0,1]. \tag{5.21b}$$

We may assume that any feasible point of the program (5.21) fulfills for every $x \in [0,1]$

$$\eta(x) = \min\left\{-y^{(1)}x, 1 - y^{(2)}x\right\} = \begin{cases} -y^{(1)}x & , y^{(2)} - y^{(1)} \leq 1 \\ -y^{(1)}x & , y^{(2)} - y^{(1)} > 1 \wedge x \leq \frac{1}{y^{(2)}-y^{(1)}} \\ 1 - y^{(2)}x & , y^{(2)} - y^{(1)} > 1 \wedge x \geq \frac{1}{y^{(2)}-y^{(1)}} \end{cases}. \tag{5.22}$$

This is feasible as it is continuous and obeys constraints (5.21a) and (5.21b) for any choice of $(y^{(i)})_{i \in [k]}$. Consequently, the objective of the program (5.21) reduces to

$$\frac{1}{2} y^{(2)} + \int_0^1 \eta(x) d\lambda(x) = \begin{cases} \frac{1}{2}(y^{(2)} - y^{(1)}) & , y^{(2)} - y^{(1)} \leq 1 \\ 1 - \frac{1}{2(y^{(2)}-y^{(1)})} & , y^{(2)} - y^{(1)} \geq 1 \end{cases}$$

Thus, (5.21) has the optimal value 1 which, for example, is the limit objective value of the (feasible) points $y^{(2)} = n, y^{(1)} = 0$ and η given by Eq. (5.22) for $n \to \infty$. However, the program (5.21) is obviously not solvable.

5.4 Approximability

The conclusion of Example 5.20 is rather disappointing as, of course, we would like to obtain a result in analogy to the finite case. Namely, that the dual problem (5.18) would yield us the parameters for a generalized Voronoi diagram that would be feasible or even support the clustering obtained from the program (5.17). However, this is not possible in general as Example 5.20 has demonstrated.

Still, Example 5.20 does offer some further insights. In this example, the greater we choose the dual $y^{(2)}$, the closer our resulting Voronoi diagram fits the optimal clustering. Also, we know that if the projection of the feasible region of (5.18) to the space of the y variables was bounded, then the program (5.18) would be solvable by the same argumentation as for the program (5.17) and using that η can be determined explicitly via $(y^{(i)})_{i\in[k]}$.

For some fixed $\tau > 0$, we consider the following adapted version of the program (5.17):

$$\min_{(\xi_i)_i \in M_r(X)^k, \rho\in\mathbb{R}} \frac{1}{\tau}\sum_{i=1}^{k}\langle f_i, \xi_i\rangle \; + \rho \tag{5.23}$$

$$\text{s.t.} \quad \sum_{i\in[k]} \xi_i = \omega_0 \tag{5.23a}$$

$$\langle \phi_{i,l}, \xi_i\rangle + \rho \geq b_l^{(i)} \quad \forall i \in [k], l \in [m_i] \tag{5.23b}$$

$$\langle \phi_{i,l}, \xi_i\rangle - \rho \leq b_l^{(i)} \quad \forall i \in [k], l \in [m_i] \tag{5.23c}$$

$$\xi_i \geq 0 \qquad \forall i \in [k] \tag{5.23d}$$

$$\rho \geq 0 \tag{5.23e}$$

Here, an error variable ρ has been added which yields the maximum error from the former equality constraint (5.17b). Penalizing this maximum error by a factor τ in the objective of (5.23) then leads to a limitation (in the dual) ℓ_1-norm in the resulting dual problem. Here, we skip the introduction of dual variables $y_l^{(i,+)}$ and $y_l^{(i,-)}$ for the constraints (5.23b) and (5.23c), respectively. Instead, we directly replace them by $y_l^{(i)} := y_l^{(i,+)} - y_l^{(i,-)}$ and recognize $y_l^{(i,+)} + y_l^{(i,-)} = \left|y_l^{(i)}\right|$. This yields the following variant of the dual of the program (5.23):

$$\max_{\eta\in C(X),y^{(i)}\in\mathbb{R}^{m_i},i\in[k]} \sum_{i=1}^{k}(y^{(i)})^\intercal b^{(i)} + \langle \eta, \omega_0\rangle \tag{5.24}$$

$$\text{s.t.} \quad \eta + (y^{(i)})^\intercal \phi_i \leq \frac{1}{\tau} f_i \qquad \forall i \in [k] \tag{5.24a}$$

$$\sum_{i=1}^{k} \left\| y^{(i)} \right\|_1 \leq 1 \tag{5.24b}$$

In this relaxed setting, we can formulate a stronger version of Theorem 5.19.

Theorem 5.21
It holds that:

i) The program (5.23) *is solvable.*

ii) The program (5.24) *is solvable.*

iii) There is no duality gap, i. e., value((5.23)) = value((5.24)).

Proof. Concerning part i, we first note that the program (5.23) is consistent. For this, we set $\xi_1 := \omega_0$, $\xi_i := 0$ for $i = 2, \ldots, k$, and $\rho := \max_{i\in[k],l\in[m_i]} \left| b_l^{(i)} - \langle \phi_{i,l}, \xi_i \rangle \right|$. We may then follow the exactly same argumentation as in Theorem 5.19 part i and conclude solvability, too.

For part ii, we once more may w.l.o.g. assume due to $\omega_0 \geq 0$ that for any solution of the program (5.24) it holds that $\eta(x) = \min_{i\in[k]} \frac{1}{\tau} f_i(x) - (y^{(i)})^\intercal \phi_i(x)$ for all $x \in X$. Thus, we can reformulate the program (5.24) as follows:

$$\max_{y^{(i)}\in\mathbb{R}^{m_i},i\in[k]} \sum_{i=1}^{k}(y^{(i)})^\intercal b^{(i)} + \langle \min_{i\in[k]}(\frac{1}{\tau} f_i - (y^{(i)})^\intercal \phi_i), \omega_0\rangle \tag{5.25}$$

$$\text{s.t.} \quad \sum_{i=1}^{k} \left\| y^{(i)} \right\|_1 \leq 1 \tag{5.25a}$$

This is a finite-dimensional optimization problem over the ℓ_1-unit ball in $\mathbb{R}^{\sum_{i\in[k]} m_i}$ with a continuous objective function. The latter follows, for example, via bounded convergence: The functions f_i and ϕ_i are continuous and hence bounded over the compact set X. Therefore, $\frac{1}{\tau} f_i(x) - (y^{(i)})^\intercal \phi_i(x)$ can be uniformly bounded in x and $y^{(i)}$ for every $i \in [k]$. Hence, the program (5.25) is solvable and therefore is the program (5.24), too.

Finally, part iii follows almost exactly as part iii of Theorem 5.21. We set $y^{(i)} := 0$ for all $i \in [k]$ and $\eta := \min_{i\in[k],\tilde{x}\in X} \frac{1}{\tau} f_i(\tilde{x}) - 1$. By doing so, we obtain a point in the (Mackey) interior of the dual positive cone and can again apply Theorem 5.11. □

Theorem 5.21 states a strong duality between the programs (5.23) and (5.24). Thus, we may again defer the feasibility of a generalized Voronoi Diagram as the complementarity property of the optimal primal and dual solution. However, as Example 5.20 showed that this may not be possible for solutions that meet the clustering constraints exactly, we have to allow an arbitrarily small error.

Theorem 5.22
Let X be a compact Hausdorff space and ω_0 a regular Borel measure on X. Let $k \in \mathbb{N} \setminus \{0\}$ and for every $i \in [k]$ let $m_i \in \mathbb{N}$ and functions $f_i \in C(X)$ and $\Phi_{i,l} \in C(X)$, $l \in [m_i]$, as well as $b^{(i)} \in \mathbb{R}^{m_i}$ be given. Next, set

$$\begin{aligned}\rho^* \quad := \min \Big\{ \rho \in \mathbb{R}_{\geq 0} : \exists (e_l^{(i)})_{i\in[k],l\in[m_i]} \in \mathbb{R}^{\sum_{i\in[k]} m_i}, \|e\|_\infty \leq \rho, \\ T\left(\omega_0, (\Phi_{i,l})_{l\in[m_i],i\in[k]}, (b^{(i)} + e^{(i)})_{i\in[k]}\right) \neq \emptyset \Big\}.\end{aligned} \tag{5.26}$$

Then for every $\epsilon > 0$, there exist $(e_l^{(i)})_{i\in[k],l\in[m_i]} \in \mathbb{R}^{\sum_{i\in[k]} m_i}$ with $\|e\|_\infty < \rho^* + \epsilon$, a clustering $\xi \in T\left(\omega_0, (\Phi_{i,l})_{l\in[m_i],i\in[k]}, (b^{(i)} + e^{(i)})_{i\in[k]}\right)$, and $(y_l^{(i)})_{i\in[k],l\in[m_i]} \in \mathbb{R}^{\sum_{i\in[k]} m_i}$ such that the generalized Voronoi Diagram w. r. t. functions $\left(f_i + (y^{(i)})^\intercal \Phi_i\right)_{i\in[k]}$ is feasible for ξ.

Proof. First note that ρ^* is actually well-defined. This follows once more due to the compactness of the space of positive regular Borel measures bounded by ω_0, or by applying Theorem 5.21 part i to the program (5.23) with $f_i = 0$ for every $i \in [k]$.

Now, for any $\tau > 0$, let ξ^τ, ρ^τ be an optimizer of the resulting program (5.23) and $(\eta_\tau, (y^{(i,\tau)})_{i\in[k]})$ be an optimizer of (5.24). As for every $i \in [k]$ it holds that $0 \leq \xi_i^\tau \leq \omega_0$, we get that

$$\left|\frac{1}{\tau}\langle f_i, \xi_i^\tau\rangle\right| \leq \frac{1}{\tau}\int_X \|f_i\|_\infty \, d\omega_0(x) \to 0 \quad (\tau \to \infty).$$

By definition of ρ^* it holds that $\rho^\tau \geq \rho^*$ and that there exists

$$\xi^* \in T\left(\omega_0, (\Phi_{i,l})_{l\in[m_i],i\in[k]}, (b^{(i)} + e^{(i)})_{i\in[k]}\right)$$

with $\|e\|_\infty = \rho^*$. In particular, (ξ^*, ρ^*) is feasible for (5.23). In total, it holds that

$$\frac{1}{\tau}\sum_{i=1}^k \langle f_i, \xi_i^\tau\rangle + \rho^* \leq \frac{1}{\tau}\sum_{i=1}^k \langle f_i, \xi_i^\tau\rangle + \rho^\tau \leq \frac{1}{\tau}\sum_{i=1}^k \langle f_i, \xi_i^*\rangle + \rho^*$$

and so it follows that $\rho^\tau \to \rho^*$ for $\tau \to \infty$.

For the remainder of this proof, assume τ to be sufficiently large such that $\rho^\tau < \rho^* + \epsilon$. Again, due to optimality we may w.l.o.g. assume $\eta_\tau(x) = \min_{i\in[k]} \frac{1}{\tau} f_i(x) - (y^{(i)})^\intercal \Phi_i(x)$. As strong duality by Theorem 5.21 part iii holds, we get

$$\begin{aligned}
&\frac{1}{\tau}\sum_{i=1}^{k}\langle f_i, \xi_i^\tau\rangle + \rho^\tau = \langle \eta_\tau, \omega_0\rangle + \sum_{i\in[k]} (y^{(i,\tau)})^\intercal b^{(i)} \\
\Leftrightarrow &\sum_{i\in[k]} \int_X \frac{1}{\tau} f_i(x)\ d\xi_i(x) + \rho^\tau = \sum_{i\in[k]} \int_X \min_{l\in[k]} \left(\frac{1}{\tau} f_l(x) - (y^{(l,\tau)})^\intercal \Phi_l(x)\right)\ d\xi_i(x) \\
&\qquad + \sum_{i\in[k]} (y^{(i,\tau)})^\intercal \Big(\int_X \Phi_i(x)\ d\xi_i^\tau(x) + b^{(i)} - \int_X \Phi_i(x)\ d\xi_i^\tau(x)\Big) \\
\Leftrightarrow &\underbrace{\sum_{i\in[k]} \int_X \frac{1}{\tau} f_i(x) - (y^{(i,\tau)})^\intercal \Phi_i(x) - \min_{l\in[k]} \left(\frac{1}{\tau} f_l(x) - (y^{(l,\tau)})^\intercal \Phi_l(x)\right)\ d\xi_i^\tau(x)}_{=:\circledast} \\
&\qquad + \underbrace{\rho^\tau - \sum_{i\in[k]} (y^{(i,\tau)})^\intercal \left(b^{(i)} - \int_X \Phi_i(x)\ d\xi_i^\tau(x)\right)}_{=:\circledast\circledast} = 0.
\end{aligned}$$

As $\xi_i^\tau \geq 0$ for every $i \in [k]$ it holds that $\circledast \geq 0$ and as

$$\begin{aligned}
\sum_{i\in[k]} (y^{(i,\tau)})^\intercal \left(b^{(i)} - \int_X \Phi_i\ d\xi_i^\tau(x)\right) &\leq \sum_{i\in[k]} \left\|y^{(i,\tau)}\right\|_1 \left\|b^{(i)} - \int_X \Phi_i(x)\ d\xi_i^\tau(x)\right\|_\infty \\
&\leq \sum_{i\in[k]} \left\|y^{(i,\tau)}\right\|_1 \rho^\tau \leq \rho^\tau
\end{aligned}$$

we get $\circledast\circledast \geq 0$, too. In particular, this implies $\circledast = 0$ and thus

$$\int_X \frac{1}{\tau} f_i(x) - (y^{(i,\tau)})^\intercal \Phi_i(x) - \min_{l\in[k]} \left(\frac{1}{\tau} f_l(x) - (y^{(l,\tau)})^\intercal \Phi_l(x)\right)\ d\xi_i^\tau(x) = 0$$

for every $i \in [k]$.

For any $x_0 \in \operatorname{supp}(\xi_i^\tau)$, this implies that $i \in \operatorname{argmin}_{l\in[k]} \left(\frac{1}{\tau} f_l(x_0) - (y^{(l,\tau)})^\intercal \Phi_l(x_0)\right)$: Otherwise, by continuity there exist a $\delta > 0$ and a neighborhood V_{x_0} of x_0 such that $\frac{1}{\tau} f_i(x) - (y^{(i,\tau)})^\intercal \Phi_i(x) - \min_{l\in[k]} \left(\frac{1}{\tau} f_l(x) - (y^{(l,\tau)})^\intercal \Phi_l(x)\right) \geq \delta$ for all $x \in V_{x_0}$. As $x_0 \in \operatorname{supp}(\xi_i^\tau)$, it must hold that $\xi_i^\tau(V_{x_0}) > 0$. However, then it holds that

$$\circledast \geq \int_{V_{x_0}} \frac{1}{\tau} f_i(x) - (y^{(i,\tau)})^\intercal \Phi_i(x) - \min_{l\in[k]} \left(\frac{1}{\tau} f_l(x) - (y^{(l,\tau)})^\intercal \Phi_l(x)\right) \geq \delta \xi_i^\tau(V_{x_0}) > 0,$$

a contradiction.

For every $i \in [k]$, we now set $y^{(i)} := -\tau \cdot y^{(i,\tau)}$ and thus have $\operatorname{supp}(\xi_i^\tau) \subset P_i$ with $(P_1, \dots, P_k)$ being the generalized Voronoi Diagram w. r. t. the functions $\left(f_i + (y^{(i)})^\intercal \Phi_i\right)_{i\in[k]}$. □

Example 5.20 (continuing from p. 125).
We may revisit our previous Example 5.20 in this approximate setting. For any $\tau > 0$, we obtain a new primal program

$$\min_{\xi_1,\xi_2 \in M_r([0,1]),\rho} \frac{1}{\tau}\xi_2([0,1]) + \rho \tag{5.27}$$

$$\text{s.t.} \quad \xi_1 + \xi_2 = \lambda_{[0,1]} \tag{5.27a}$$

$$-\rho \leq \int_{[0,1]} x \; d\xi_1(x) \leq \rho \tag{5.27b}$$

$$\frac{1}{2} - \rho \leq \int_{[0,1]} x \; d\xi_2(x) \leq \frac{1}{2} + \rho \tag{5.27c}$$

$$\xi_1,\ \xi_2,\ \rho \geq 0 \tag{5.27d}$$

and a corresponding dual program

$$\max_{\eta \in C([0,1]), y^{(1)}, y^{(2)} \in \mathbb{R}} \frac{1}{2}y^{(2)} + \int_0^1 \eta(x) d\lambda(x) \tag{5.28}$$

$$\text{s.t.} \quad \eta(x) + y^{(1)}x \leq 0 \qquad \forall x \in [0,1] \tag{5.28a}$$

$$\eta(x) + y^{(2)}x \leq \frac{1}{\tau} \qquad \forall x \in [0,1] \tag{5.28b}$$

$$\left|y^{(1)}\right| + \left|y^{(2)}\right| \leq 1. \tag{5.28c}$$

Let us again first consider a solution $(\eta, y^{(1)}, y^{(2)})$ of the dual program (5.28). Here, we can assume for any $x \in [0,1]$ that

$$\eta(x) = \min\left\{-y^{(1)}x, \frac{1}{\tau} - y^{(2)}x\right\} = \begin{cases} y^{(1)}x & , y^{(2)} - y^{(1)} \leq \frac{1}{\tau} \vee x \leq \frac{1}{\tau(y^{(2)} - y^{(1)})} \\ 1 - y^{(2)}x & , y^{(2)} - y^{(1)} \geq \frac{1}{\tau} \wedge x \geq \frac{1}{\tau(y^{(2)} - y^{(1)})} \end{cases}$$

holds and obtain a corresponding objective value

$$\frac{1}{2}y^{(2)} + \int_0^1 \eta(x) d\lambda(x) = \begin{cases} \frac{1}{2}(y^{(2)} - y^{(1)}) & , y^{(2)} - y^{(1)} \leq \frac{1}{\tau} \\ \frac{1}{\tau} - \frac{1}{2\tau^2(y^{(2)} - y^{(1)})} & , y^{(2)} - y^{(1)} \geq \frac{1}{\tau} \end{cases}.$$

By constraint (5.28c), it must hold that $y^{(2)} - y^{(1)} \leq 1$. We see that — independently of the choice of τ — an optimizer must fulfill $y^{(2)} - y^{(1)} = 1$. Thus, we obtain for an optimal solution $(\eta^*, y^{(1,*)}, y^{(2,*)})$ a resulting objective value of

$$\frac{1}{2}y^{(2,*)} + \int_0^1 \eta^*(x) d\lambda(x) = \begin{cases} \frac{1}{2} & \tau \leq 1 \\ \frac{1}{\tau} - \frac{1}{2\tau^2} & \tau \geq 1 \end{cases}.$$

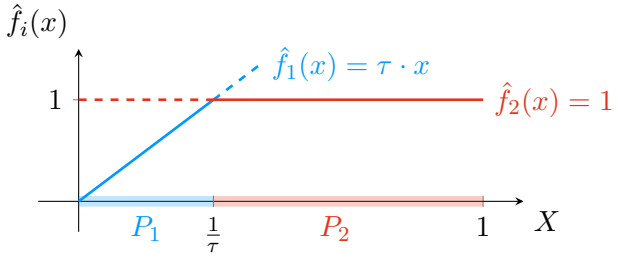

Figure 5.2: Optimal solution and corresponding diagram of program (5.27) in Example 5.20 for $y^{(1,*)} = -1$ and $y^{(2,*)} = 0$.

Let now $\mathcal{P} = (P_1, P_2)$ be the generalized Voronoi Diagram w. r. t. the functions $\hat{f}_1(x) := -\tau \cdot y^{(1,*)} \cdot x$ and $\hat{f}_2(x) := 1 - \tau \cdot y^{(2,*)} \cdot x$ for $x \in [0,1]$. Then

$$\begin{aligned} P_1 &= \left\{x \in [0,1] : -\tau \cdot y^{(1,*)} \cdot x \leq 1 - \tau \cdot y^{(2,*)} \cdot x\right\} \\ &= \left\{x \in [0,1] : x(y^{(2,*)} - y^{(1,*)}) \leq \frac{1}{\tau}\right\} = \left[0, \min\left\{\frac{1}{\tau}, 1\right\}\right] \end{aligned}$$

and, analogously, $P_2 = [\min\{\frac{1}{\tau}, 1\}, 1]$. From Theorem 5.21 we know that $\mathcal{P}$ must be feasible for any optimizer of (5.27). Thus, we deduce that the unique optimizer ξ^* of the program (5.27) is given by

$$\xi_1^* = \lambda\left(\left[0, \min\left\{\frac{1}{\tau}, 1\right\}\right]\right) \quad \text{and} \quad \xi_2^* = \lambda\left(\left[\min\left\{\frac{1}{\tau}, 1\right\}, 1\right]\right).$$

This yields $\rho^* = \max\left\{\int_0^{\min\{\frac{1}{\tau},1\}} x\, dx, \frac{1}{2} - \int_{\min\{\frac{1}{\tau},1\}}^1 x\, dx\right\} = \min\left\{\frac{1}{2\tau^2}, \frac{1}{2}\right\}$. (Note that taking the maximum of the two constraint violations is actually not necessary as the constraints (5.27a) and (5.27b) imply the constraint (5.27c).) We may now verify our calculations by computing the resulting objective value of (5.27). This is

$$\frac{1}{\tau}\xi_2^*([0,1]) + \rho^* = \frac{1}{\tau}\left(1 - \min\left\{\frac{1}{\tau}, 1\right\}\right) + \min\left\{\frac{1}{2\tau^2}, \frac{1}{2}\right\} = \begin{cases} \frac{1}{2} & , 0 < \tau \leq 1 \\ \frac{1}{\tau} - \frac{1}{2\tau^2} & , \tau \geq 1 \end{cases},$$

which, of course, is consistent with our obtained dual objective value. Figure 5.2 depicts this situation (for $\tau > 1$).

5.5 Summary & Conclusion

This chapter considered the correspondence between constrained clusterings and generalized Voronoi diagrams as established in Chapter 4 for clusterings over (almost)

arbitrary Hausdorff spaces. In particular, this accounts for clusterings of a whole subset of $\mathbb{R}^d$ (instead of a finite point set). Although the results in this chapter are of mainly theoretical interest, they reveal the core essence of the proposed correspondence.

As this, of course, required a lot prior knowledge w. r. t. algebraic spaces, topology and measure theory, an introduction to a linear programming theory over general Hausdorff spaces was given in Section 5.1. Here, we mainly followed the theory of [AN87], but focused particularly on duality as it yields the essence of the relation between clusterings and diagrams. Section 5.2 then transferred our previous terminology to this more general setting and in particular defined clusterings as tuples of measures.

Having established the preliminaries, the desired relation could once more be established rather easily. However, we failed to do so in Section 5.3 as solvability of the dual program and hence the existence of suitable parameters for a feasible diagram cannot be guaranteed in general. Section 5.4 demonstrated how this can be fixed by allowing an arbitrarily small error w. r. t. the cluster constraints. Note that this result in particular carries over to our existing theory of the previous chapter in case that a clustering polytope is empty.

We restricted the established correspondence in Theorem 5.22 to the feasibility of a diagram for the scope of this thesis. Of course, the questions in which cases a supporting property can be established and how this again translates into a kind of strict complementarity arise. As mentioned in the beginning of this chapter, at least for clusterings over $\mathbb{R}^d$ and using a Lebesgue measure as ω_0, one may also obtain similar results using refining discretizations (as sketched by [CCD16]). Both [CCD16] as well as [Gei+14] establish the correspondence between weight-constrained clusterings and additively weighted generalized Voronoi diagrams for this setting. However, both require that the bisectors of the diagrams are null sets. Under that assumption, feasibility and the supporting property of a diagram collapse (up to sets of measure zero).

Chapter 6

Balanced k-Means Clustering

Power diagrams have an exposed position in the context of geometric clustering as they come along with some favorable properties, are well-understood and have been successfully applied in various areas ([Aur87a], [BG12]; cf. also Section 3.3.2). In Section 4.3 they have been our prime example for the correspondence between constrained clusterings and generalized Voronoi diagrams.

In unconstrained clustering, classical Voronoi diagrams yield least-squares assignments. Finding the best choice of sites then is the well-known k-Means problem. Analogously, power diagrams lead to least-squares assignments under weight-constraints. The problem of which sites to choose then states the Balanced k-Means problem ([BBG17]).

In our prime application of electoral district design in Chapter 7, we will as well make extensive use of this rich theory. Therefore, this chapter is to collect, combine and extend the existing results on the Balanced k-Means problem. Section 6.1 will discuss the Balanced k-Means problem itself, while Section 6.2 will then discuss algorithmic approaches to it.

6.1 The Balanced k-Means Problem

In this section, we will see that constrained clustering via power diagrams can be understood as the weight-constrained version of the famous k-Means clustering, which will be called the Balanced k-Means problem.

For the remainder of this section, we consider the same setting as in Section 4.4. This means, we assume a set $X = \left\{x^{(1)}, \dots, x^{(n)}\right\}$ of n distinct points in $\mathbb{R}^d$ with weights $\omega_j \in \mathbb{R}_{>0}$ for $j \in [n]$ to be given. Furthermore, we assume cluster weights $\kappa_i \in \mathbb{R}_{>0}$ for $i \in [k]$ with $\sum_{i\in[k]} \kappa_i = \sum_{j\in[n]} \omega_j$ to be given. We are then interested in clusterings from the resulting clustering polytope $\xi \in T_{\kappa,\omega}$ as defined by Eq. (4.16).

The k-Means problem is to find an (unconstrained) clustering that minimizes the total squared error when each cluster is represented by a single point (cf. [JMF99; Wu12]). Of course, each representative of a cluster then must coincide with the centroid in an optimum. In our context, the total squared error is also referred to as *moment of inertia* ([LGC13]). More formally, for a clustering $\xi \in T_{\kappa,\omega}$, we denote by

$$c\left(C_i\right) := \sum_{j\in[n]} \frac{\omega_j}{\omega(C_i)} x^{(j)} \tag{6.1}$$

the centroid of the ith cluster. Then the moment of inertia of the clustering ξ is given by

$$\sum_{i\in[k]} \sum_{j\in[n]} \xi_{i,j}\omega_j \left\|x^{(j)} - c\left(C_i\right)\right\|_2^2. \tag{6.2}$$

6.1.1 Squared Error Analysis

As k-MEANS clustering plays an important role in the field of data analysis, its objective may be interpreted from a statistical point of view (cf. [DHS00, Section 10.2], for example). First of all, this shall give a better understanding of the objects that we will deal with in the following. Second, it prepares some formal dependencies that we can utilize later on.

Let $\overline{X}$ and $\overline{C}$ be random vectors supposed to yield a sampled unit and cluster, respectively. The units' weights now naturally define a probability distribution of $\overline{X}$ via

$$\mathbb{P}\left(\overline{X} = x^{(j)}\right) := \frac{\omega_j}{\omega(X)}.$$

Similarly, the cluster weights provide a distribution for $\overline{C}$ by

$$\mathbb{P}\left(\overline{C} = C_i\right) := \frac{\kappa_i}{\omega(X)}$$

for $i \in [k]$. A (fixed) clustering $\xi = (C_1, \ldots, C_k)^\intercal \in T_{\text{uncstr.},k,n}$ then defines for each cluster the conditional probabilities

$$\mathbb{P}\left(\overline{C} = C_i | \overline{X} = x^{(j)}\right) = \xi_{i,j}.$$

Let $c(X) := \sum_{j\in[n]} \frac{\omega_j}{\omega(X)} x^{(j)} = \mathbb{E}\left[\overline{X}\right]$ denote the centroid of all units. Similarly, let $c\left(\overline{C}\right)$ denote the random variable that yields the centroid of the random cluster $\overline{C}$, i. e., $c\left(\overline{C}\right) := \mathbb{E}\left[\overline{X}|\overline{C}\right]$. In particular, this implies

$$\mathbb{E}\left[c\left(\overline{C}\right)\right] = \sum_{i\in[k]} \mathbb{P}\left(\overline{C} = C_i\right) \mathbb{E}\left[\overline{X}|\overline{C} = C_i\right] = \mathbb{E}\left[\overline{X}\right].$$

The overall variance in the data is then

$$\begin{aligned}
\operatorname{var}\left(\overline{X}\right) &= \sum_{j\in[n]} \mathbb{P}\left(\overline{X} = x^{(j)}\right) \left\|x^{(j)} - c(X)\right\|_2^2 \qquad (6.3)\\
&= \sum_{i\in[k]} \mathbb{P}\left(\overline{C} = C_i\right) \sum_{j\in[n]} \mathbb{P}\left(\overline{X} = x^{(j)} | \overline{C} = C_i\right) \left\|x^{(j)} - c(X)\right\|_2^2 \\
&= \sum_{i\in[k]} \mathbb{P}\left(\overline{C} = C_i\right) \left(\sum_{j\in[n]} \mathbb{P}\left(\overline{X} = x^{(j)} | \overline{C} = C_i\right) \left\|x^{(j)} - c(C_i)\right\|_2^2 + \|c(C_i) - c(X)\|_2^2 \right) \\
&= \mathbb{E}\left[\operatorname{var}\left(\overline{X}|\overline{C}\right)\right] + \operatorname{var}\left(c\left(\overline{C}\right)\right) \\
&= \frac{1}{\omega(X)} \sum_{i\in[k]} \sum_{j\in[n]} \xi_{i,j}\omega_j \left\|x^{(j)} - c(C_i)\right\|_2^2 + \frac{1}{\omega(X)} \sum_{i\in[k]} \omega(C_i) \|c(C_i) - c(X)\|_2^2 . \qquad (6.4)
\end{aligned}$$

In other words, the total variance of the data splits into the expected *inner-cluster variance* $\mathbb{E}\left[\mathrm{var}\left(\overline{X}|\overline{C}\right)\right]$ and the *centroid variance* $\mathrm{var}\left(c\left(\overline{C}\right)\right)$. The k-Means clustering problem asks to find an (unconstrained) clustering that minimizes the expected inner-cluster variance. Due to Eq. (6.4), this is equivalent to maximizing the centroid variance. The latter is then the variance that is *explained* by the model, while the former remains as the unexplained, undesirable distortion. In our setting, the minimization of the moment of inertia means to maximize the explained cluster variance under the constraint of balanced clusterings.

There is also a very helpful geometrical interpretation of the centroid variance. However, in order to derive this later on, let us remain in the given stochastic model for one further step. Here, we may exploit the basic fact that for any two identically and independently distributed random vectors $\overline{A}, \overline{B}$ with existing first and second moment it holds that

$$\begin{aligned}\mathrm{var}\left(\overline{A}\right) &= \mathbb{E}\left[\left\|\overline{A}\right\|_2^2\right] - \left\|\mathbb{E}\left[\overline{A}\right]\right\|_2^2 = \frac{1}{2}\left(\mathbb{E}\left[\left\|\overline{A}\right\|_2^2 + \left\|\overline{B}\right\|_2^2\right] - 2\mathbb{E}\left[\overline{A}\right]^\intercal \mathbb{E}\left[\overline{B}\right]\right) \\ &= \frac{1}{2}\mathbb{E}\left[\left\|\overline{A} - \overline{B}\right\|_2^2\right].\end{aligned}$$

First, this yields that the expected inner-cluster variance can be rewritten as

$$\begin{aligned}\mathbb{E}\left[\mathrm{var}\left(\overline{X}|\overline{C}\right)\right] &= \frac{1}{\omega(X)} \sum_{i\in[k]} \sum_{j\in[n]} \xi_{i,j}\omega_j \left\|x^{(j)} - c\left(C_i\right)\right\|_2^2 \\ &= \frac{1}{2\omega(X)} \sum_{i\in[k]} \frac{1}{\omega(C_i)} \sum_{a,b\in[n]} \xi_{i,a}\omega_a\xi_{i,b}\omega_b \left\|x^{(a)} - x^{(b)}\right\|_2^2, \end{aligned} \tag{6.5}$$

which can be read as the expected pairwise squared distances of all units that are assigned to the same cluster.

Analogously, the centroid variance can be read as

$$\mathrm{var}\left(c\left(\overline{C}\right)\right) = \sum_{i\in[k]} \frac{\omega(C_i)}{\omega(X)} \left\|c\left(C_i\right) - c\left(X\right)\right\|_2^2 = \sum_{i<l\in[k]} \frac{\omega(C_i)\omega(C_l)}{\omega(X)^2} \left\|c\left(C_i\right) - c\left(C_l\right)\right\|_2^2. \tag{6.6}$$

So, in order to minimize the moment of inertia, we want to find a clustering such that the resulting cluster centroids are pairwise "pushed away from each other" as far as possible. At the same time, we note that

$$\begin{aligned}\mathrm{var}\left(c\left(\overline{C}\right)\right) &= \mathbb{E}\left[\left\|c\left(\overline{C}\right)\right\|_2^2\right] - \left\|\mathbb{E}\left[c\left(\overline{C}\right)\right]\right\|_2^2 = \mathbb{E}\left[\left\|c\left(\overline{C}\right)\right\|_2^2\right] - \left\|\mathbb{E}\left[\overline{X}\right]\right\|_2^2 \\ &= \frac{1}{\omega(X)} \sum_{i\in[k]} \omega(C_i) \left\|c\left(C_i\right)\right\|_2^2 \; - \left\|\mathbb{E}\left[\overline{X}\right]\right\|_2^2 \end{aligned} \tag{6.7}$$

holds. As $\left\|\mathbb{E}\left[\overline{X}\right]\right\|_2^2$ does not depend on the underlying clustering, the maximization of the cluster variance is equivalent to the maximization of the weighted sum of squared centroid norms.

6.1.2 Problem Definition and Complexity

In analogy to the famous k-Means problem, let us formally define the Balanced k-Means problem.

Problem Balanced k-Means:

Input:	$\omega_1, \ldots, \omega_n \in \mathbb{N}_{>0}$, $x^{(1)}, \ldots, x^{(n)} \in \mathbb{Q}^d$, $\kappa_1, \ldots, \kappa_k \in \mathbb{N}_{>0}$
Question:	Determine

$$\min \sum_{i\in[k]} \sum_{j\in[n]} \xi_{i,j}\omega_j \left\|x^{(j)} - s^{(i)}\right\|_2^2$$
$$\xi \in T_{\kappa,\omega}$$
$$\left(s^{(i)}\right)_{i\in[k]} \in (\mathbb{R}^d)^k.$$

In the definition of Balanced k-Means we do allow to choose the sites $\left(s^{(i)}\right)_{i\in[k]}$ independently of the clustering. However, we see that

$$\sum_{i\in[k]} \sum_{j\in[n]} \xi_{i,j}\omega_j \left\|x^{(j)} - s^{(i)}\right\|_2^2 = \sum_{i\in[k]} \left(\sum_{j\in[n]} \xi_{i,j}\omega_j \left\|x^{(j)} - c\left(C_i\right)\right\|_2^2 + \kappa_i \left\|s^{(i)} - c\left(C_i\right)\right\|_2^2 \right) \tag{6.8}$$

holds, which easily follows by reformulations in analogy to the previous section. Thus, for any optimum (of the optimization variant) of Balanced k-Means, the sites are identical to the cluster centroids (and hence the objective value equal to the moment of inertia as defined in Eq. (6.2)).

Let us analyze the complexity of Balanced k-Means. The famous k-Means problem has been shown to be $\mathcal{APX}$-hard ([Awa+15], [LSW17]). Furthermore, it is known that k-Means remains $\mathcal{NP}$-hard even in the plane ([MNV09]). (For a clarification on the used terms for approximation complexity we refer to [KV12, Chapters 15-16] and [Aus+99].) We may start our analysis by the observation that the decision variant of Balanced k-Means is $\mathcal{NP}$-complete. While this will follow as well from later results, an independent proof by a reduction of a variant of Exact Cover by 3-Sets is provided in Appendix A.

Proposition 6.1
The decision variant of Balanced k-Means *is $\mathcal{NP}$-complete.*

Proof. Containment in $\mathcal{NP}$ follows as we can restrict the sites to the centroids that result from a clustering. The latter on the other hand can be chosen extremal w. r. t. T, hence we obtain a rational certificate. Hardness will follow from Theorems 6.3 and 6.4. □

Unfortunately, we have not been able to transfer the inapproximability results of k-Means from the literature to Balanced k-Means. However, following Eq. (6.4) we can formulate an alternative variant:

Problem Balanced Max Centroid Variance:

Input:	$\omega_1, \dots, \omega_n \in \mathbb{N}_{>0}$, $x^{(1)}, \dots, x^{(n)} \in \mathbb{Q}^d$, $\kappa_1, \dots, \kappa_k \in \mathbb{N}_{>0}$
Question:	Determine

$$\max \sum_{i \in [k]} \kappa_i \left\| c\left(C_i\right) - c\left(X\right) \right\|_2^2 \quad \text{s.t.} \quad \xi \in T.$$

In other words, while we do not know about the approximability of the expected inner-cluster-variance, we will show inapproximability for the centroid variance that is optimized by Balanced Max Centroid Variance. We will shortly follow the established theory in [BG12] for a practical approach to the usage of power diagrams in our setting. This will reduce the Balanced k-Means problem via the Balanced Max Centroid Variance formulation to a certain norm maximization problem. For now, we can already exploit this connection in order to deduce inapproximability.

In order to do so, we will exploit a hardness result on Euclidean norm maximization over full-dimensional parallelotopes.

Problem $[-1, 1]$-ParMax:

Input:	$v^{(1)}, \dots, v^{(d)} \in \mathbb{Q}^d$ linearly independent
Question:	Determine

$$\max \left\{ \left\| v \right\|_2^2 : v \in \sum_{i \in [d]} [-1, 1] v^{(i)} \right\}.$$

Brieden, Gritzmann, and Klee [BGK00] showed the following inapproximability result on $[-1, 1]$-ParMax. In the following statement, the *ratio* of an approximation algorithm for a minimization problem refers to the quotient of the algorithm's output value and

the actual optimal value in the worst-case (so that by definition this is a value greater or equal than 1).

Theorem 6.2 ([BGK00, Theorem 2.1])
The problem $[-1,1]$-ParMax *is* $\mathcal{APX}$*-complete and there exists a polynomial-time approximation algorithm with ratio* $7/3$*. Unless* $\mathcal{P} = \mathcal{NP}$*, there does not exist a polynomial-time approximation algorithm of ratio* 1.090*.*

With this result we are able to show the following:

Theorem 6.3
The problem Balanced Max Centroid Variance *is* $\mathcal{APX}$*-hard, even when reduced to the case* $k = 2$*,* $\omega \equiv 1$ *and* $\kappa_1 = \kappa_2 = \frac{n}{2}$*.*

Proof. Let $v^{(1)}, \dots, v^{(d)} \in \mathbb{Q}^d$ define an instance of $[-1,1]$-ParMax and set $P := \sum_{i\in[d]}[-1,1]v^{(i)}$. We define an auxiliary polytope in $\mathbb{Q}^{2d}$ via

$$\tilde{P} := \Big\{ x \in \mathbb{Q}^{2d} : \begin{array}{l} x = \sum_{i\in[d]} \lambda_i \begin{pmatrix} v^{(i)} \\ 0 \end{pmatrix} + \sum_{i\in[d]} \lambda_{d+i} \begin{pmatrix} 0 \\ v^{(i)} \end{pmatrix}, \\ \lambda \in \mathbb{B}_\infty^{2d}, \quad \lambda^\intercal \mathbb{1} = 0 \end{array} \Big\}. \tag{6.9}$$

We will first show

$$\max_{x\in P} \|x\|_2^2 = \frac{1}{2} \max_{\tilde{x}\in\tilde{P}} \|\tilde{x}\|_2^2 . \tag{6.10}$$

For any $\lambda \in [0,1]^d$ and $x := \sum_{i\in[d]} \lambda_i v^{(i)}$ we can define $\tilde{\lambda} := (\lambda^\intercal, -\lambda^\intercal)^\intercal$. Thus, $\tilde{\lambda} \in \mathbb{B}_\infty^{2d}$, $\tilde{\lambda}^\intercal \mathbb{1} = 0$, and

$$\left\| \sum_{i\in[d]} \tilde{\lambda}_i \begin{pmatrix} v^{(i)} \\ 0 \end{pmatrix} + \sum_{i\in[d]} \tilde{\lambda}_{d+i} \begin{pmatrix} 0 \\ v^{(i)} \end{pmatrix} \right\|_2^2 = \left\| \sum_{i\in[d]} \tilde{\lambda}_i v^{(i)} \right\|_2^2 + \left\| \sum_{i\in[d]} \tilde{\lambda}_{d+i} v^{(i)} \right\|_2^2 = 2\|x\|_2^2$$

hold. Vice versa, let $\tilde{\lambda} \in \mathbb{B}_\infty^{2d}$ with $\tilde{\lambda}^\intercal \mathbb{1} = 0$ and obtain the corresponding $\tilde{x} \in \tilde{P}$ according to Eq. (6.9). With $\tilde{x} =: ((x^{(1)})^\intercal, (x^{(2)})^\intercal)^\intercal$ it follows that $x^{(1)}, x^{(2)} \in P$ as well as $\max\left\{ \left\|x^{(1)}\right\|_2^2, \left\|x^{(2)}\right\|_2^2 \right\} \geq \frac{1}{2}\|\tilde{x}\|_2^2$.

We now obtain an instance of Balanced Max Centroid Variance in dimension $2d$ via $k := 2$, $n := 2d$, $\kappa_1 := \kappa_2 := \frac{n}{2}$, $\omega_j := 1$ for $j \in [n]$ and $x^{(j)} := \begin{pmatrix} v^{(j)} \\ 0 \end{pmatrix}$, $x^{(d+j)} := \begin{pmatrix} 0 \\ v^{(j)} \end{pmatrix}$ for all $j \in [d]$. Let $T := T\left((\omega_1, \dots, \omega_n), (\kappa_i)_{i\in[k]}\right)$ be the corresponding

clustering polytope. We claim that

$$\frac{n^2}{4} \max_{\xi \in T} \|c(C_1) - c(C_2)\|_2^2 = \max_{\tilde{x} \in \tilde{P}} \|\tilde{x}\|_2^2 . \tag{6.11}$$

For any $\xi \in T$, we can set $\lambda := \xi_{1,\cdot} - \xi_{2,\cdot}$. Then both $\lambda \in \mathbb{B}_\infty^{2d}$ and $\lambda^\intercal \mathbb{1} = \sum_{j \in [n]} \xi_{1,j} - \xi_{2,j} = \frac{n}{2} - \frac{n}{2} = 0$ hold. For the corresponding $\tilde{x} := \sum_{j \in [2d]} \lambda_j x^{(j)} \in \tilde{P}$ (as in Eq. (6.9)) it follows that

$$\|\tilde{x}\|_2^2 = \left\| \sum_{j \in [2d]} \xi_{1,j} x^{(j)} - \sum_{j \in [2d]} \xi_{2,j} x^{(j)} \right\|_2^2 = \frac{n^2}{4} \|c(C_1) - c(C_2)\|_2^2 . \tag{6.12}$$

On the other hand, for any $\lambda \in \mathbb{B}_\infty^{2d}$ with $\lambda^\intercal \mathbb{1} = 0$ we can define $\xi_{1,\cdot} := \frac{1}{2}(\mathbb{1}^{(n)} + \lambda)$ as well as $\xi_{2,\cdot} := \frac{1}{2}(\mathbb{1}^{(n)} - \lambda)$. We obtain $\xi \geq 0$, $\xi_{1,\cdot} + \xi_{2,\cdot} = \mathbb{1}^{(n)}$ as well as $\xi_{i,\cdot}^\intercal \mathbb{1}^{(n)} = \frac{1}{2}((\mathbb{1}^{(n)})^\intercal \mathbb{1}^{(n)} \pm \lambda^\intercal \mathbb{1}^{(n)}) = \frac{n}{2} = \kappa_i$ for $i \in [2]$, so $\xi \in T$ holds. With

$$\sum_{j \in [2d]} \xi_{1,j} x^{(j)} - \sum_{j \in [2d]} \xi_{2,j} x^{(j)} = \sum_{j \in [2d]} \frac{1}{2}(1 - \lambda_j) x^{(j)} - \sum_{j \in [2d]} \frac{1}{2}(1 + \lambda_j) = \sum_{j \in [2d]} \lambda_j x^{(j)}$$

we see that (6.12) holds and so (6.11) follows. From Eqs. (6.6), (6.10) and (6.11), we get that

$$\max_{x \in P} \|x\|_2^2 = \frac{n^2}{8} \max_{\xi \in T} \|c(C_1) - c(C_2)\|_2^2 = \frac{n}{2} \max_{\xi \in T} \sum_{i \in [k]} \kappa_i \|c(C_i) - c(X)\|_2^2$$

and so this reduction in fact preserves any approximation ratio. Hence, we can conclude both claims from Theorem 6.2. □

Of course, from Theorem 6.3 we can deduce Proposition 6.1. However, we are particularly interested in the planar case for our applications of interest. In order to show hardness for this case, we can indeed apply the proof of Mahajan, Nimbhorkar, and Varadarajan [MNV09] for showing hardness of the planar k-Means problem.

Theorem 6.4
Balanced k-Means *is $\mathcal{NP}$-complete even when restricted to the plane, i. e., $d = 2$.*

Proof. Let us roughly describe the basic idea of the proof in [MNV09] and how this fits our setting. Here, hardness is shown by a reduction from Planar-3SAT. This variant of 3SAT only considers instances that yield a planar graph with variables and clauses as nodes. The edges of this graph form a cycle of all variable nodes as well as connect all variables to clauses containing an associated literal.

Using this planar embedding the authors then construct another graph with an embedding in the plane. This graph then consists of disjoint even cycles for each variable and is embedded in the plane such that the nodes lie on a grid. Furthermore, the node placement, the grid distance, as well as the choice of the parameter k are made such that an optimal k-MEANS clustering will produce a matching of nodes along those cycles. Which of the two possible matchings is chosen corresponds to the variable assignment for 3SAT. Finally, the cycles are arranged such that for each clause another point can be placed in a particular way. For each literal in a clause, this point then lies close to a pair of variable nodes such that a clustering of this pair fits the type of the literal. An optimal k-MEANS clustering will then still partition the variable node cycles into pairs of neighbors, while each of the clause points will be assigned to one of those pairs. Only if the 3SAT instance is a YES-instance, every clause will possess at least one close pair of variable nodes that is contained in one cluster.

Important for our setting, the sizes of the clusters in an optimal clustering are known a priori. Here, note that in order to enforce optimal clusterings that produce valid variable assignments, the variable points are weighted by a factor M, while the clause points have weight 1. (More precisely, the authors in [MNV09] consider an unweighted setting and thus create M copies of each of those points.) With $s_1, \dots, s_r \in 2\mathbb{N}_{>0}$ denoting the number of nodes on each of those cycles and m being the number of clauses, we get a total number of clusters $k = \frac{1}{2}\sum_{i\in[r]} s_i$. We furthermore know that any optimal k-MEANS clustering will have $k - m$ clusters of size $2M$ and m clusters of size $2M + 1$, so we can equivalently consider the according instance of BALANCED k-MEANS, which then yields the claim. □

Unfortunately, crucial for our applications is the equally balanced case in the plane. While the construction in the reduction in Theorem 6.4 de facto leads to asymptotically equally balanced clusters (of sizes $2M$ and $2M + 1$, where M can be chosen arbitrarily large), the proof from [MNV09] does not seem to be easily adaptable to us. Thus, this remains an open question.

6.2 Algorithmic Approaches

Despite the inapproximability results of the previous section, the BALANCED k-MEANS problem turns out to be practically manageable for many applications of our interest. This section is to summarize and adapt the existing algorithmic approaches (here, we mainly focus on [BBG17; BG04; BG10; BG12]).

6.2.1 A Balanced-k-Means Algorithm

A first natural approach is to adapt the famous k-MEANS *algorithm* (also known as Lloyd's method as first introduced in [Llo82]). This has been done by Borgwardt,

Brieden, and Gritzmann [BBG17], who defined and analyzed the (weight-)balanced k-means algorithm. Note that in [BBG17] the authors consider a more general setting as their variant allows lower and upper bounds for the weights of the resulting clusters. The Balanced k-Means-algorithm is given by Algorithm 1.

Algorithm 1: Balanced k-Means Algorithm

Input: $x^{(1)}, \dots, x^{(n)} \in \mathbb{R}^d$,
$\omega_1, \dots, \omega_n \in \mathbb{R}_{>0}, \kappa_1, \dots, \kappa_k \in \mathbb{R}_{>0}, s^{(1)}, \dots, s^{(k)} \in \mathbb{R}^d$

1 **Repeat**
2 Obtain an optimal solution $\xi = (C_1, \dots, C_k)^\intercal$ of (4.18) w. r. t. $f_i(x) = \left\|x - s^{(i)}\right\|_2^2$, $i = 1, \dots, k$, and according to the remaining input.
3 **if** $c(C_i) = s^{(i)} \quad \forall i \in [k]$ **then**
4 **return** ξ, $s^{(1)}, \dots, s^{(k)}$
5 **forall** $i \in [k]$ **do** $s^{(i)} \leftarrow c(C_i)$

In order to analyze the algorithm, the terminology for the relation of power diagrams and clusterings can be extended as follows: Given a clustering ξ, a unit set X, and a power diagram $\mathcal{P}$ w. r. t. functions $f_i(x) = \left\|x - s^{(i)}\right\|_2^2 + \mu_i$ for $i \in [k]$, the diagram is called *centroidal* if $c(C_i) = s^{(i)}$ holds for every $i \in [k]$. We have already concluded from Eq. (6.8) that in every optimum of Balanced k-Means the sites must coincide with the cluster centroids. Hence, in any optimum the corresponding feasible power diagram must in fact be centroidal.

Consequently, if this is not the case, the update step of Algorithm 1 (line 5) adjusts the sites by the current cluster centroids. From Eq. (6.8) we immediately see that this yields a sequence of (fractional) clusterings that is strictly decreasing w. r. t. the Balanced k-Means objective. By exploiting the relation to power diagrams, Borgwardt, Brieden, and Gritzmann [BBG17] furthermore show that the algorithm will terminate in $n^{\mathcal{O}(dk)}$ iterations. As is the case for the classical k-Means algorithm, Algorithm 1 does in general not lead to a global optimum of Balanced k-Means and in particular strongly depends on the choice of the start sites provided by the input. We will consider its convergence properties more closely later on.

6.2.2 Norm Maximization over Gravity Polytopes

While in Section 6.1 we saw that the search for a global optimum is indeed a hard problem, it has been demonstrated both theoretically ([BG10; BG12]) and practically by means of the problem of farmland consolidation ([BBG14; BG04]) that it can still be reasonably tackled in an approximate manner. Here, we will focus on the Balanced

Max Centroid Variance variant. Following Eq. (6.6), we see that the objective of Balanced Max Centroid Variance can be rewritten as

$$\sum_{i\in[k]} \kappa_i \left\| c\left(C_i\right) - c\left(X\right)\right\|_2^2 = \frac{1}{\omega(X)} \sum_{i<l\in[k]} \kappa_i\kappa_l \left\| c\left(C_i\right) - c\left(C_l\right)\right\|_2^2 . \tag{6.13}$$

Informally speaking, this may be read as the optimal clustering pushing the cluster centroids as much apart as possible. Brieden and Gritzmann [BG10] provide a thorough investigation of *clustering bodies*. In general, given norms $\|\cdot\|_\diamond$ and $\|\cdot\|_*$ in $\mathbb{R}^{\frac{k(k-1)}{2}}$ and $\mathbb{R}^d$, respectively, those describe the level sets

$$C := \left\{(c_1^\intercal, \dots, c_k^\intercal)^\intercal \in \mathbb{R}^{dk} : \|(\|c_1 - c_2\|_*, \|c_1 - c_3\|_*, \dots, \|c_{k-1} - c_k\|_*)\|_\diamond \le 1\right\}.$$

In [BG10], clustering bodies for various choices of $\|\cdot\|_\diamond$ and $\|\cdot\|_*$ are studied, in particular with regard to their approximability by polyhedra. Observe that the level sets of the objective (6.13) are up to taking the square root identical to the clustering body with $\|\cdot\|_* := \|\cdot\|_2$ being the Euclidean norm and $\|\cdot\|_\diamond$ being the ellipsoidal norm with the ellipsoid $\operatorname{diag}\left(\left(\sqrt{\frac{\omega(X)}{\kappa_i\kappa_l}}\right)_{i<l\in[k]}\right) \mathbb{B}_2^{\frac{k(k-1)}{2}}$ as unit ball.

For the fully Euclidean case, i. e., $\|\cdot\|_\diamond = \|\cdot\|_2$ and $\|\cdot\|_* = \|\cdot\|_2$, it follows that C has the lineality space $L = \left\{(c_i)_{i\in[k]} \in \mathbb{R}^{dk} : c_i = c_l \ \forall i, l \in [k]\right\}$. Furthermore, they obtain $C = \frac{1}{\sqrt{k}}\mathbb{B}_2^{dk} \cap L^\perp + L$ (cf. [BG10, Lemma 2.3, Theorem 4.3]). Thus, a polyhedral approximation of the $d(k-1)$-dimensional Euclidean ball (embedded in the subspace $L^\perp = \left\{(c_i)_{i\in[k]} \in \mathbb{R}^{dk} : \sum_{i\in[k]} c_i = 0\right\}$) will suffice in order to approximate the clustering body.

There is another way to recognize Balanced Max Centroid Variance as a norm maximization problem, as is also described in [BG12]. Note that for the purpose of solving Balanced Max Centroid Variance (or Balanced k-Means) we can assume the units to be centralized, i. e., $\sum_{j\in[n]} \omega_j x^{(j)} = 0$ or, equivalently, $c\left(X\right) = 0$. Otherwise, we could commonly translate all unit points (and thus all cluster centroids) without consequence for the considered problems. Also, note that a common translation of all units results in a translation of the cluster centroids along the gravity body's lineality space L. For the objective of Balanced Max Centroid Variance this assumption yields

$$\sum_{i\in[k]} \kappa_i \left\| c\left(C_i\right) - c\left(X\right)\right\|_2^2 = \sum_{i\in[k]} \kappa_i \left\| c\left(C_i\right)\right\|_2^2 . \tag{6.14}$$

Note that this is also, of course, in accordance with Eq. (6.7), as here we are assuming $\mathbb{E}\left[\overline{X}\right] = 0$.

As is clear from Eq. (6.14), our objective is only sensitive to changes of the cluster centroids. Thus, following [BG12] we define the *gravity polytope*

$$G := \left\{ (c(C_1)^\intercal, \dots, c(C_k)^\intercal)^\intercal \in \mathbb{R}^{kd} : \exists \xi = (C_1, \dots, C_k)^\intercal \in T_{\kappa,\omega} \right\} = A \cdot T_{\kappa,\omega}, \quad (6.15)$$

where $A : \mathbb{R}^{k \times n} \to \mathbb{R}^{kd}$ is the linear function that maps each cluster C_i to its resulting centroid $c(C_i) = \sum_{j \in [n]} \xi_{i,j} \frac{\omega_j}{\kappa_i} x^{(j)}$. Note that in [BG12] more general gravity *bodies* are considered as the authors include the possibility for cluster weights κ_i not to be fixed but to range in certain bounds. In particular, this breaks the linear relationship between the polytope of feasible clusterings and the set of resulting centroids. Also related to that, Hwang, Onn, and Rothblum [HOR98] define *single shaped partition polytopes* which result from the vector sums of size-constrained clusters of (unweighted) points and hence yield dilatations of the gravity polytopes as defined in (6.15). More closely, *mean* partition polytopes as defined in [CHR06] are identical to gravity polytopes as defined here but for unweighted points. Gravity polytopes for unweighted points are also researched by Borgwardt [Bor10] with the focus on the strict separation properties of their vertices as well as their edge-structure.

We observe that the gravity polytope lives in a $(k-1)d$ dimensional space:

Lemma 6.5
Assume $\dim(X) = d$. *For the gravity polytope* G *as defined by Eq.* (6.15) *it holds that* $\dim(G) = d(k-1)$ *and* $\operatorname{aff}(G) = \left\{ (c^{(i)})_{i \in [k]} \in \mathbb{R}^{dk} : \sum_{i \in [k]} \kappa_i c^{(i)} = \sum_{j \in [n]} \omega_j x^{(j)} \right\}$.

Proof. Set $L_G := \left\{ (c^{(i)})_{i \in [k]} \in \mathbb{R}^{dk} : \sum_{i \in [k]} \kappa_i c^{(i)} = \sum_{j \in [n]} \omega_j x^{(j)} \right\}$. It is clear by definition that $G \subset L_G$ as well as $\dim(L_G) = d(k-1)$ hold, so we only need to show $\dim(G) \geq d(k-1)$.

Define the clustering $\xi^{(0)} \in T_{\kappa,\omega}$ via $\xi^{(0)}_{i,j} := \frac{\kappa_i}{\omega(X)}$ for $i \in [k], j \in [n]$. It is easy to check that $\xi^{(0)} \in T_{\kappa,\omega}$ and, as $\xi^{(0)} > 0$, in particular $\xi^{(0)} \in \operatorname{relint}(T_{\kappa,\omega})$. Thus, with $c^{(0)} = A\xi^{(0)}$ it follows that $c^{(0)} = (c(X), \dots, c(X)) \in \operatorname{relint}(G)$. Now, let $(v^{(i)})_{i \in [k-1]} \in \mathbb{R}^{d(k-1)}$, set $v^{(k)} := -\frac{1}{\kappa_k} \sum_{i \in [k-1]} \kappa_i v^{(i)}$, and $v := (v^{(0)}, \dots, v^{(k)})$. Thus, $c^{(0)} + v \in L_G$ holds.

As $\dim(X) = d$, we may w.l.o.g. assume $\dim(\left\{ x^{(1)}, \dots, x^{(d+1)} \right\}) = d$.

As $\xi^{(0)} > 0$, it holds that the clustering graph $G\left(\xi^{(0)}\right)$ is complete (as a bipartite graph between the cluster and unit nodes). This enables us to construct cyclic exchanges that yield the desired perturbation of the cluster centroids. For this purpose, let $\zeta^{(i,j)}$ for $i \in [k-1]$ and $j \in [d]$ denote the cyclic exchange associated with the cycle $C_k, j, C_i, d+1, C_k$ (as defined by Eq. (2.6)).

Fix some $i \in [k-1]$. By construction, we get $(A\zeta^{(i,j)})_{i,\cdot} = \frac{1}{\kappa_i}\left(x^{(j)} - x^{(d+1)}\right)$, $(A\zeta^{(i,j)})_{k,\cdot} = \frac{1}{\kappa_k}\left(x^{(d+1)} - x^{(j)}\right)$, and $(A\zeta^{(i,j)})_{l,\cdot} = 0$ for all $l \in [k-1] \setminus \{i\}$.

Next, due to the assumed affine independence, we can choose $\lambda_j^{(i)} \in \mathbb{R}$ for $j \in [d]$, such that $\kappa_i v^{(i)} = \sum_{j\in[d]} \lambda_j^{(i)}(x^{(j)} - x^{(d+1)})$. This yields $A(\sum_{j\in[d]} \lambda_j^{(i)}\zeta^{(i,j)}) = \left(0,\dots,0,v^{(i)},0,\dots,0,-\frac{\kappa_i}{\kappa_k}v^{(i)}\right)$. So, in total

$$A\left(\sum_{i\in[k-1]}\sum_{j\in[d]}\lambda_j^{(i)}\zeta^{(i,j)}\right) = \left(v^{(1)},v^{(2)},\dots,v^{(k-1)},-\frac{1}{\kappa_k}\sum_{i\in[k-1]}\kappa_i v^{(i)}\right) = v.$$

Hence, for $\epsilon > 0$ sufficiently small, we have $\xi^{(0)} + \epsilon\sum_{i\in[k-1]}\sum_{j\in[d]}\lambda_j^{(i)}\zeta^{(i,j)} \in T_{\kappa,\omega}$ and hence $c^{(0)} + \epsilon v = A(\xi^{(0)} + \epsilon\sum_{i\in[k-1]}\sum_{j\in[d]}\lambda_j^{(i)}\zeta^{(i,j)}) \in G$, which shows the claim. □

Under the assumption of $c(X) = 0$, we can deduce from Eq. (6.14) that Balanced Max Centroid Variance reduces to a norm maximization problem over the gravity polytope, i. e.,

$$\max_{c=(c^{(i)})_{i\in[k]}\in G}\sum_{i\in[k]}\kappa_i\left\|c^{(i)}\right\|_2^2. \tag{6.16}$$

Together with Lemma 6.5, this means that problem (6.16) asks to maximize an ellipsoidal norm in the $d(k-1)$-dimensional subspace $\operatorname{aff}(G) = \operatorname{lin}(G)$ (as $c(X) = 0$) with the ellipsoid $\operatorname{diag}\left(\kappa_1^{-\frac{1}{2}},\dots,\kappa_k^{-\frac{1}{2}}\right)\mathbb{B}_2^{kd} \cap \operatorname{lin}(G)$ as its unit ball.

Now let us merge the observations above and our previous theory of the correspondence between power diagrams and weight-balanced clusterings. As a first step towards an algorithmic approach for the problem (6.16), we may first consider the faces of G. We denote a face of G w. r. t. an outer normal $d \in \mathbb{R}^{dk}$ by

$$F(G,d) := \{c \in G : d^\intercal c \geq d^\intercal \hat{c}\ \forall \hat{c} \in G\} \tag{6.17}$$

and faces of $T_{\kappa,\omega}$ accordingly. Note that due to $G = AT_{\kappa,\omega}$, we can obtain any face of G as the image of the face of $T_{\kappa,\omega}$ with outer-normal $A^\intercal d$, i. e., $F(G,d) = AF(T_{\kappa,\omega}, A^\intercal d)$. Observe that for any vector $d = (d^{(i)})_{i\in[k]} \in \mathbb{R}^{dk}$ it holds that

$$A^\intercal d = \left(\frac{\omega_j}{\kappa_i}(x^{(j)})^\intercal d^{(i)}\right)_{\substack{i\in[k],\\ j\in[n]}} \in \mathbb{R}^{k\times n}.$$

Now let $(a_1{}^\intercal,\dots,a_k{}^\intercal)^\intercal \in \mathbb{R}^{dk}$. For reasons to become clear shortly, we are interested in the face $F(G,(-\kappa_1 a_1{}^\intercal,\dots,-\kappa_k a_k{}^\intercal)^\intercal)$. We obtain this as the image

$$F(G,(-\kappa_1 a_1{}^\intercal,\dots,-\kappa_k a_k{}^\intercal)^\intercal) = AF\left(T_{\kappa,\omega},\left(-\omega_j a_i{}^\intercal x^{(j)}\right)_{i\in[k],j\in[n]}\right). \tag{6.18}$$

The latter face is the set of optimizers of

$$\min_{\xi \in T_{\kappa,\omega}} \sum_{i\in[k]} \sum_{j\in[n]} \xi_{i,j}\omega_j \cdot {a_i}^\intercal x^{(j)}. \tag{6.19}$$

We observe that the linear program (6.19) is just of the form (4.18) but w. r. t. linear functions $f_i(x) := {a_i}^\intercal x$. Following Proposition 4.5, this means that for optimal solutions of (6.19), there are feasible (or even supporting) generalized Voronoi diagrams w. r. t. affine functions $f_i(x) := {a_i}^\intercal x + \mu_i$. This, of course, yields a power diagram in affine representation as introduced in Section 3.3.2. An alternative way to observe this is to take any $\xi \in T_{\kappa,\omega}$ and rewrite the Balanced k-Means objective as

$$\begin{aligned}\sum_{i\in[k]} \sum_{j\in[n]} \xi_{i,j}\omega_j \left\|x^{(j)} - s^{(i)}\right\|_2^2 &= \sum_{i\in[k]} \sum_{j\in[n]} \xi_{i,j}\omega_j (\left\|x^{(j)}\right\|_2^2 - 2{(s^{(i)})}^\intercal x^{(j)} + \left\|s^{(i)}\right\|_2^2) \\ &= \sum_{j\in[n]} \omega_j \left\|x^{(j)}\right\|_2^2 + \sum_{i\in[k]} \kappa_i \left\|s^{(i)}\right\|_2^2 + \sum_{i\in[k]} \sum_{j\in[n]} \xi_{i,j}\omega_j (-2s^{(i)})^\intercal x^{(j)}. \end{aligned} \tag{6.20}$$

Hence, the considered optimal face of the program (4.18) equals that of (6.19) for $a_i = -2s^{(i)}$ for $i \in [k]$, which is in accordance with Section 3.3.2.

In other words, every face of the gravity polytope corresponds to exactly those clusterings that allow a feasible power diagram whose sites are determined from the corresponding outer normal. The following theorem precises this relation between the faces of G and power diagrams. In particular, it shows that there is a one-to-one correspondence between vertices of the gravity polytope and vertices of the clustering polytope that allow a supporting power diagram. The latter statement which is part iii of the following theorem combines [BG12, Theorem 2.1] and [BG12, Theorem 4.7]. The uniqueness of clusterings that have a centroid vector which is extremal may also be deduced from the according result for single shaped partition polytopes in [HOR98, Theorem 3.1] and has also been shown in [Bor10, Lemma 2.12].

Theorem 6.6
Let $a_i \in \mathbb{R}^d$ for $i \in [k]$.
The following holds:

i) Let $\xi \in T_{\kappa,\omega}$. Then there exists $(\mu_i)_{i\in[k]} \in \mathbb{R}^k$ such that the power diagram $\mathrm{PD}_{\mathrm{aff}}\left((a_i, \mu_i)_{i\in[k]}\right)$ is feasible for ξ if and only if $A\xi \in F\left(G, (-\kappa_1 {a_1}^\intercal, \ldots, -\kappa_k {a_k}^\intercal)^\intercal\right)$.

ii) Let $c = \left(c^{(i)}\right)_{i\in[k]} \in G$. Then there exists a clustering $\xi \in T_{\kappa,\omega}$ with $c = A\xi$ and $(\mu_i)_{i\in[k]} \in \mathbb{R}^k$ such that the power diagram $\mathrm{PD}_{\mathrm{aff}}\left((a_i, \mu_i)_{i\in[k]}\right)$ supports ξ if and only if $c \in \mathrm{relint}\left(F\left(G, (-\kappa_1 {a_1}^\intercal, \ldots, -\kappa_k {a_k}^\intercal)^\intercal\right)\right)$.

iii) The vector c is a vertex of G if and only if there exists a unique clustering $\xi \in T_{\kappa,\omega}$ such that $c = A\xi$, ξ is a vertex of $T_{\kappa,\omega}$, and there exists a power diagram that supports ξ.

Proof. Part i immediately follows from Proposition 4.5 as the problem (6.19) fits the form of (4.18) (cf. the discussion above).

Part ii follows almost analogously. Let $d := (-\kappa_1 a_1{}^\intercal, \dots, -\kappa_k a_k{}^\intercal)^\intercal$. From Eq. (6.18) it particularly follows that

$$\operatorname{relint}(F(G,d)) = A \cdot \operatorname{relint}(F(T_{\kappa,\omega}, A^\intercal d)). \tag{6.21}$$

Here, we use that taking the relative interior and applying a linear transformation commutes for convex sets in general (see, for example, [HL93, Chapter 3, Proposition 2.1.12]). Consequently, for every $c \in \operatorname{relint}(F(G,d))$, there exists a clustering $\xi \in \operatorname{relint}(F(T_{\kappa,\omega}, A^\intercal d)$ which then yields the desired claim in connection with Proposition 4.5.

For part iii, first assume that c is a vertex of G. Let $d \in \mathbb{R}^{dk}$ such that $F(G,d) = \{c\}$ (i. e., d is in the interior of the normal cone of G in c).

Now let $\xi^{(1)}, \xi^{(2)} \in T_{\kappa,\omega}$ such that $A\xi^{(1)} = A\xi^{(2)} = c$. Furthermore, set $\xi^* := \frac{1}{2}(\xi^{(1)} + \xi^{(2)})$. From Corollary 2.17 we know that $\xi^{(2)} - \xi^{(1)}$ is a linear combination of cyclic exchanges whose corresponding cycles are all contained in $G\left(\xi^{(2)} - \xi^{(1)}\right) \subset G(\xi^*)$. Let ζ be such a cyclic exchange. As $G(\zeta) \subset G(\xi^*)$, there exists $\epsilon > 0$ such that $\xi^* \pm \epsilon\zeta \in T_{\kappa,\omega}$. Consequently, $c \pm \epsilon A\zeta = A(\xi^* \pm \epsilon\zeta) \in G$. As c is extremal, it follows that $A\zeta = 0$. Let j^+, C_i, j^- be a subsequent tuple of a unit, cluster, and unit node along the cycle $G(\zeta)$. By construction, $(A\zeta)_{i,\cdot} = x^{(j^+)} - x^{(j^-)}$. Thus, $x^{(j^+)} = x^{(j^-)}$, which contradicts our assumption of unit points being pairwise distinct. Consequently, there does not exist such a cyclic exchange and hence $\xi^{(1)} = \xi^* = \xi^{(2)}$ is the unique clustering with centroid vector c.

Now for any $\xi \in F(T_{\kappa,\omega}, A^\intercal d)$ it holds by definition that $d^\intercal A\xi \geq d^\intercal A\xi^* = d^\intercal c$. Thus, $A\xi \in F(G,d) = \{c\}$. As ξ^* is unique, this means $F(T_{\kappa,\omega}, A^\intercal d) = \{\xi^*\}$, so ξ^* is a vertex of $T_{\kappa,\omega}$. The existence of a supporting power diagram then follows from part ii.

For the reverse implication of part iii, let ξ be a vertex of $T_{\kappa,\omega}$ and the power diagram $\mathrm{PD}_{\mathrm{aff}}\left((a_i,\mu_i)_{i\in[k]}\right)$ supporting ξ. Then Proposition 4.5 in connection with (6.19) yields $\xi \in \operatorname{relint}(F(T_{\kappa,\omega}, A^\intercal d))$ with $d := (-\kappa_1 a_1{}^\intercal, \dots, -\kappa_k a_k{}^\intercal)^\intercal$. As ξ is a vertex, this implies $F(T_{\kappa,\omega}, A^\intercal d) = \{\xi\}$. Then Eq. (6.21) yields that $\operatorname{relint}(F(G,d)) = \{A\xi\}$, and so $A\xi$ is a vertex of G. □

Note that Theorem 6.6 also applies to the gravity polytope itself as a non-proper face. Our definition of power diagrams does not exclude the parameters that are all 0. This yields a (degenerated) power diagram whose cells all coincide. The diagram indeed supports all clusterings that assign each point to every cluster and whose

resulting centroid vectors give the relative interior of the gravity polytope (according to Theorem 6.6 part ii).

Also note that in [BHR92] it was shown that extremal points of single shaped partition polytopes (and hence scaled gravity polytopes for unweighted points) lead to clusters with disjoint convex hulls. This is, of course, in accordance with the findings here, as the unweighted case leads to integer extremal clusterings. Thus, a supporting power diagram must contain the clusters' points in the interior of the corresponding cells which implies that the convex hulls are indeed disjoint. This was also demonstrated in [Bor10, Theorem 2.52].

The relation of faces of the gravity polytope G to power diagrams furthermore allows us to bound their number. In order to do so, we exploit a result from [BBG17]. Here, the authors bound the number of iterations of the Balanced k-Means-algorithm as they consider the set of all *power patterns* given as

$$\begin{aligned} PP(X) := \Big\{ (X \cap P_i)_{i\in[k]} : \forall i \in [k] \exists s^{(i)} \in \mathbb{R}^d, \mu_i \in \mathbb{R} : \\ (P_1, \dots, P_k) = \mathrm{PD}\left((s^{(i)}, \mu_i)_{i\in[k]} \right) \Big\}. \end{aligned} \tag{6.22}$$

In [BBG17], the cardinality of $PP(X)$ is bounded. Let us briefly describe the key idea to do so. Consider some power diagram $(P_1, \dots, P_k) = \mathrm{PD}_{\mathrm{aff}}\left((a_i, \alpha_i)_{i\in[k]}\right)$. Then in order to decide for a unit $x^{(j)} \in X$ whether it is contained in the dominance region $H_{i,l}$ for $i < l \in [k]$, one has to evaluate the sign of $a_i{}^\intercal x^{(j)} + \alpha_i - a_l{}^\intercal x^{(j)} + \alpha_l$. This is linear in the power diagram parameters $(a_i, \alpha_i)_{i\in[k]}$. In total, there are $n\binom{k}{2}$ such linear expressions. A sign-pattern of those linear expressions is now given by the signs in $\{-1, 0, 1\}$ of their evaluations for some parameter vector. A result from [War68], that allows to bound the number of occurring sign patterns for any set of polynomial expressions, is then adapted. This yields the following theorem[1]:

Theorem 6.7 ([BBG17, Theorem 2])
It holds that

$$|PP(X)| \leq \left(\frac{4e \cdot (k-1)n}{d} \right)^{(d+1)k-1}.$$

In [BBG17], the power patterns are then set into relation to certain vertices of $T_{\kappa,\omega}$. Using a similar argumentation, we can relate power patterns to those faces of $T_{\kappa,\omega}$ that yield the faces of G and hence bound their number.

[1]Note that in the original theorem in [BBG17] the bound is given as $\left(\frac{8e\cdot(k-1)n}{d}\right)^{(d+1)k-1}$. Here, a factor of 2 in the inner expression results from the fact that in [BBG17] the spherical representation of power diagrams is considered and hence a quadratic (instead of linear) expression is obtained. However, it can be easily checked that this does not require any further adaption of the proof.

Theorem 6.8
It holds that

$$\left|\left\{F\left(G,d\right) : d \in \mathbb{R}^{dk}\right\}\right| \leq \left(\frac{4e \cdot (k-1)n}{d}\right)^{(d+1)k-1}.$$

Proof. We show the claim by arguing that every face of G can be associated with a distinct power pattern, so the claim follows from Theorem 6.7.

Let $d = ((d^{(1)})^\intercal, \ldots, (d^{(k)})^\intercal)^\intercal \in \mathbb{R}^{dk}$. By definition of $T_{\kappa,\omega}$ as the intersection of an affine subspace with the non-negative orthant, it is clear that all $\xi \in \operatorname{relint}(F\left(T_{\kappa,\omega}, A^\intercal d\right))$ have the same support vector $(\operatorname{supp}(C_i))_{i\in[k]}$, as well as that the relative interiors of distinct faces of $T_{\kappa,\omega}$ have distinct support vectors. Further, there exists a power diagram $\mathcal{P} = (P_1, \ldots, P_k)$ that supports ξ, which by definition yields $P_i \cap X = \left\{x^{(j)} \in X : j \in \operatorname{supp}(C_i)\right\}$ for all $i \in [k]$. Thus, every face $F\left(T_{\kappa,\omega}, A^\intercal d\right)$ can be associated with a distinct power pattern. As every face of the gravity polytope G is obtained as $F\left(G,d\right) = AF\left(T_{\kappa,\omega}, A^\intercal d\right)$, the claim follows. □

Next, we are interested in (local) optima of the problem (6.16). As we are maximizing a strictly convex function over a polytope, we know that every local optimum will be a vertex of G, which according to Theorem 6.6 corresponds to a vertex of $T_{\kappa,\omega}$. The following theorem concludes that a local optimizer is achieved for any clustering that is extremal w. r. t. $T_{\kappa,\omega}$ and that allows a supporting centroidal diagram. Note that this result coincides with [BG12, Theorem 2.4] (up to minimal changes as we do allow degenerate diagrams).

Theorem 6.9
Let $c = \left(c^{(i)}\right)_{i\in[k]} \in G$.
Then the following holds:

i) The vector c is a local optimizer of (6.16) *if and only if* $F\left(G, (\kappa_i c^{(i)})_{i\in[k]}\right) = \{c\}$.

ii) The vector c is a local optimizer of (6.16) *if and only if there exists a vertex $\xi \in T_{\kappa,\omega}$ with $c = A\xi$ and a power diagram that supports ξ and is centroidal (w. r. t. ξ).*

Proof. As problem (6.16) asks for maximizing a strictly convex smooth function over a polytope, we know that $c \in G$ is a local maximum of (6.16) if and only if the gradient of this function at c lies in the interior of the normal cone in c. With $g : \mathbb{R}^{dk} \to \mathbb{R}$, $g(c^{(1)}, \ldots, c^{(k)}) := \sum_{i\in[k]} \left\|c^{(i)}\right\|_2^2$, we have that $\nabla g(c^{(1)}, \ldots, c^{(k)}) = (2\kappa_i c^{(i)})_{i\in[k]}$ holds.

Thus, $c = \left(c^{(i)}\right)_{i\in[k]}$ is a local maximum of (6.16) if and only if $\nabla g(c) \in \operatorname{int}(N_G(c))$ and hence $F\left(G, (\kappa_1 c^{(1)}, \ldots, \kappa_k c^{(k)})\right) = \{c\}$.

Part ii now follows from part i in combination with parts ii and iii of Theorem 6.6. Assume $F\left(G, (\kappa_i c^{(i)})_{i\in[k]}\right) = \{c\}$. Then by part ii of Theorem 6.6, as $c \in \operatorname{relint}\left(F\left(G, (2\kappa_1 c^{(1)}, \ldots, 2\kappa_k c^{(k)})\right)\right)$, there is a clustering with centroids c and a supporting power diagram $\mathrm{PD}_{\mathrm{aff}}\left(\left(-2c^{(i)}, \mu_i\right)_{i\in[k]}\right)$. Recalling the relation between affine and spherical representations of power diagrams (see Section 3.3.2), we find $\tilde{\mu}_i \in \mathbb{R}$ for $i \in [k]$ such that $\mathrm{PD}_{\mathrm{aff}}\left(\left(-2c^{(i)}, \mu_i\right)_{i\in[k]}\right) = \mathrm{PD}\left(\left(c^{(i)}, \tilde{\mu}_i\right)_{i\in[k]}\right)$. Hence, this power diagram is in fact centroidal. From Theorem 6.6 part iii we get that the clustering must be a vertex of $T_{\kappa,\omega}$.

Vice versa, if there exists a vertex $\xi \in T_{\kappa,\omega}$ with corresponding centroid vector c and a centroidal power diagram that supports ξ, Theorem 6.6 part iii yields that c is a vertex, and Theorem 6.6 part ii (using the same relation as above) that $c \in \operatorname{relint}(F\left(G, (\kappa_i c^{(i)})_{i\in[k]}\right)$. So $F\left(G, (\kappa_i c^{(i)})_{i\in[k]}\right) = \{c\}$. □

Concerning our original Balanced k-Means problem, we get the analogous result:

Theorem 6.10
Let sites $(s^{(i)})_{i\in[k]} \in \mathbb{R}^{dk}$ and a clustering $\xi = (C_1, \ldots, C_k)^\intercal \in T_{\kappa,\omega}$ be given. Then the following statements are equivalent:

i) The pair $((s^{(i)})_{i\in[k]}, \xi)$ is a local optimum of Balanced k-Means.

ii) The clustering ξ is the unique optimum of (6.19) *for $a_i := -c(C_i) = -s^{(i)}$ for $i \in [k]$.*

iii) The clustering ξ is a vertex of $T_{\kappa,\omega}$ and there exists a power diagram with sites $(s^{(i)})_{i\in[k]}$ that supports ξ and is centroidal.

Proof. First, it is clear that all three statements imply $s^{(i)} = c(C_i)$ for every $i \in [k]$ (in particular recall Eq. (6.8)).

Second, let us reassure that we may also w.l.o.g. assume that

$$c(X) = \sum_{j\in[n]} \frac{\omega_j}{\omega(X)} x^{(j)} = 0$$

(or otherwise shift all points uniformly by $-c(X)$). As we assume $s^{(i)} = c(C_i)$ for $i \in [k]$, the objective of Balanced k-Means is invariant to uniform shifts of the units (and hence the cluster centroids). Concerning the second statement and problem (6.19),

we can observe that with $e := ((e^{(0)})^\intercal, \dots, (e^{(0)})^\intercal)^\intercal$ for any $e^{(0)} \in \mathbb{R}^d$ and $\xi \in T_{\kappa,\omega}$, it holds that

$$e^\intercal A\xi = \sum_{i\in[k]} \sum_{j\in[n]} \xi_{i,j}\omega_j (e^{(0)})^\intercal x^{(j)} = (e^{(0)})^\intercal (\sum_{j\in[n]} \omega_j x^{(j)})$$

and so for any $d \in \mathbb{R}^{dk}$, it follows that $F\left(T_{\kappa,\omega}, A^\intercal d\right) = F\left(T_{\kappa,\omega}, A^\intercal(d+e)\right)$. This implies that the program (6.19) is invariant w. r. t. common shifts of the vectors a_i. Hence, a common shift of the centroids does not affect the second statement, too. For the third statement, this holds true, too, as the cells of any centroidal power diagram shift uniformly if the sites do.

Next, following the preceding section we know that $\left((c\left(C_i\right))_{i\in[k]}, \xi\right)$ for $\xi = (C_1, \dots, C_k)^\intercal \in T_{\kappa,\omega}$ is a local optimum of Balanced k-Means if and only if ξ is a local optimum of Balanced Max Centroid Variance. Under the made assumption the latter is equivalent to $c := \left(c\left(C_i\right)\right)_{i\in[k]}$ being a local optimum of (6.16).

According to Theorem 6.9 part i this is equivalent to

$$\{c\} = F\left(G, \left(\kappa_i c\left(C_i\right)\right)_{i\in[k]}\right). \tag{6.23}$$

If Eq. (6.23) holds, then c is a vertex of G. Thus, Eq. (6.18) together with Theorem 6.6 part iii yield that Eq. (6.23) holds if and only if $F\left(T_{\kappa,\omega}, \left(-\omega_j c\left(C_i\right)^\intercal x^{(j)}\right)_{i\in[k],j\in[n]}\right) = \{\xi\}$, which is exactly part ii.

Finally, part iii readily follows from Theorem 6.9 part ii. □

The following examples stress that both the *supporting* property as well as the claim for an *extremal* solution in part ii of Theorem 6.9 are necessary. Note that those translate into the *uniqueness* requirement for clusterings in Theorem 6.10.

The first example shows the possibility of a clustering whose centroid vector is a vertex of G but not a local optimum of the problem (6.16), although there exists a feasible centroidal power diagram.

Example 6.11
We consider the setting as illustrated in Fig. 6.1. Here, we set

$$x^{(1)} := \begin{pmatrix} -1 \\ -1 \end{pmatrix}, x^{(2)} := \begin{pmatrix} -1 \\ 0 \end{pmatrix}, x^{(3)} := \begin{pmatrix} 1 \\ 0 \end{pmatrix}, x^{(4)} := \begin{pmatrix} 1 \\ 1 \end{pmatrix},$$

$\omega_j := 1$ for $j \in [4]$, and $\kappa_1 := \kappa_2 := 2$.

Let $\xi^{(0)}$ be the clustering as depicted in Fig. 6.1a, i. e., assigning units $x^{(1)}$ and $x^{(3)}$ to cluster 1 and $x^{(2)}$ and $x^{(4)}$ to cluster 2. Then $\xi^{(0)}$ is extremal (as it is integer) with

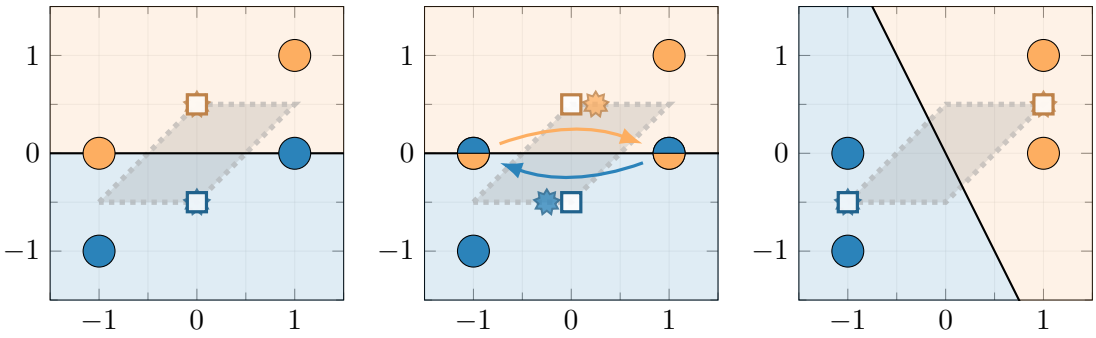

(a) Centroidal and extremal, but no local optimum.

(b) Improving movement of cluster centroids.

(c) Global optimum.

Figure 6.1: Illustrations for Example 6.11. Squares depict the sites of the power diagrams, while the jagged points mark the resulting cluster centroids. The gray-shaped area illustrates the projection of the gravity polytope to the first d coordinates.

centroids

$$c(C_1) = \begin{pmatrix} 0 \\ -1/2 \end{pmatrix}, \quad c(C_2) = \begin{pmatrix} 0 \\ 1/2 \end{pmatrix},$$

and the power diagram $\mathrm{PD}\left(\left(c\left({C_1}^{(0)}\right), 0\right), \left(c\left({C_2}^{(0)}\right), 0\right)\right)$ is feasible for $\xi^{(0)}$ and centroidal.

However, $c^{(0)} := \left(c\left({C_1}^{(0)}\right), c\left({C_2}^{(0)}\right)\right)$ is *not* a local optimum of the problem (6.16): Consider the cyclic exchange ζ resulting form the cycle $C_1, 3, C_2, 2, C_1$ (as defined by Eq. (2.6)). Then for $0 \leq \epsilon \leq 1$ the clustering $\xi^{(\epsilon)} := \xi + \epsilon\zeta$ has centroids

$$c\left({C_1}^{(\epsilon)}\right) = \begin{pmatrix} -\epsilon \\ -1/2 \end{pmatrix} \text{ and } c\left({C_2}^{(\epsilon)}\right) = \begin{pmatrix} \epsilon \\ 1/2 \end{pmatrix},$$

as illustrated in Fig. 6.1b. Then $c^{(\epsilon)} := \left(c\left({C_1}^{(\epsilon)}\right), c\left({C_2}^{(\epsilon)}\right)\right)$ results in an objective value $1 + 4\epsilon^2$ for problem (6.16), which for $\epsilon = 1$ results in the global optimum (cf. Fig. 6.1c). Thus, ξ is *not* a local optimizer. Crucial here, it can be easily verified that $F\left(G, (\kappa_i c^{(i)})_{i\in[k]}\right) = \mathrm{conv}\left(\left\{c^{(0)}, c^{(1)}\right\}\right)$ is an edge of G. Figure 6.1 also depicts in gray the projection of G onto the first two coordinates (i. e., one of the two centroids), which preserves the faces of G (note that $\dim(G) = 2$). In particular, $c^{(0)}$ is indeed a vertex of G. A (non-centroidal) power diagram that supports ξ in accordance with Theorem 6.6 part iii is, for example, given by $\mathrm{PD}\left(\left(\begin{pmatrix} 1/2 \\ -1/2 \end{pmatrix}, 0\right), \left(\begin{pmatrix} -1/2 \\ +1/2 \end{pmatrix}, 0\right)\right)$ (which simply results in a counter-clockwise rotation of the bisector of the power diagram in Fig. 6.1a).

Similarly, we can easily find an example in which there exists a centroidal power diagram that is even supporting, but the corresponding centroids do not yield a local optimum of (6.16) as the face $F\left(G, (\kappa_i c^{(i)})_{i\in[k]}\right)$ is not a vertex.

Example 6.12
We consider the situation as depicted by Fig. 6.2 which is similar to that of Example 6.11.

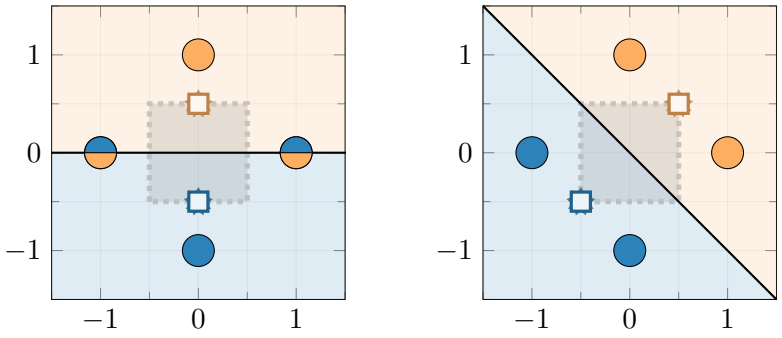

(a) Centroidal but not extremal.

(b) Optimal clustering.

Figure 6.2: Illustration for Example 6.12, depictions as in Fig. 6.1.

This time, we have

$$x^{(1)} := \begin{pmatrix} 0 \\ -1 \end{pmatrix}, x^{(2)} := \begin{pmatrix} -1 \\ 0 \end{pmatrix}, x^{(3)} := \begin{pmatrix} 1 \\ 0 \end{pmatrix}, x^{(4)} := \begin{pmatrix} 0 \\ 1 \end{pmatrix},$$

while still $\omega_j := 1$ for $j \in [4]$, and $\kappa_1 := \kappa_2 := 2$.

Figure 6.2a depicts the clustering $\xi^{(0)}$ that integrally assigns units $x^{(1)}$ and $x^{(4)}$ to clusters 1 and 2, respectively, while units $x^{(2)}$ and $x^{(3)}$ are split equally. Thus, the centroids are

$$c\left(C_1\right) = \begin{pmatrix} 0 \\ -1/2 \end{pmatrix} \text{ and } c\left(C_2\right) = \begin{pmatrix} 0 \\ 1/2 \end{pmatrix}.$$

The power diagram $\mathrm{PD}\left((c\left({C_1}^{(0)}\right), 0), (c\left({C_2}^{(0)}\right), 0)\right)$ both supports $\xi^{(0)}$ and is centroidal.

However, as indicated by Fig. 6.2a, we can see that

$$F\left(G, (\kappa_i c^{(i)})_{i\in[k]}\right) = \operatorname{conv}\left(\left\{\left(\begin{pmatrix} -1/2 \\ -1/2 \end{pmatrix}, \begin{pmatrix} 1/2 \\ 1/2 \end{pmatrix}\right), \left(\begin{pmatrix} 1/2 \\ -1/2 \end{pmatrix}, \begin{pmatrix} -1/2 \\ 1/2 \end{pmatrix}\right)\right\}\right).$$

So following Theorem 6.9, this cannot be a local optimum. In fact, any vertex of G is a global optimum in this example (cf. Fig. 6.2b).

Examples 6.11 and 6.12 also point out an issue with the Balanced k-Means-algorithm as stated by Algorithm 1: In the given form, degenerate cases such as given in the examples might prevent the algorithm from convergence to a local optimum.

However, from Theorem 6.10 we know that this can be prevented as long as we require the solutions of the problem (4.18) to be unique and supported by a centroidal power diagram. By its definition, the Balanced k-Means-algorithm does not stop before a centroidal power diagram has been found. We know further that if we always take a solution from the relative interior of the optimal face of (4.18), we enforce the supporting property. Now the following lemma implies that if we have obtained a supporting and centroidal power diagram but the clustering is not unique, we can improve in any direction on the current face of the gravity polytope. Note that this situation appears in Example 6.12.

Lemma 6.13
Let $c = (c^{(i)})_{i\in[k]} \in G$ and $\xi \in T_{\kappa,\omega}$ such that $c = A\xi$ and $\{c\} \subsetneq F := F\left(G, (\kappa_i c^{(i)})_{i\in[k]}\right)$.

Then for any centroid vector $\hat{c} \in F \setminus \{c\}$ and clustering $\hat{\xi} \in T_{\kappa,\omega}$ with $\hat{c} = A\hat{\xi}$, it holds that

$$\sum_{i\in[k]}\sum_{j\in[n]} \hat{\xi}_{i,j}\omega_j \left\|x^{(j)} - \hat{c}^{(i)}\right\|_2^2 < \sum_{i\in[k]}\sum_{j\in[n]} \xi_{i,j}\omega_j \left\|x^{(j)} - c^{(i)}\right\|_2^2.$$

Proof. As $c, \hat{c} \in F$ it follows by Eqs. (6.18) and (6.20) that

$$\sum_{i\in[k]}\sum_{j\in[n]} \hat{\xi}_{i,j}\omega_j \left\|x^{(j)} - c^{(i)}\right\|_2^2 = \sum_{i\in[k]}\sum_{j\in[n]} \xi_{i,j}\omega_j \left\|x^{(j)} - c^{(i)}\right\|_2^2.$$

From Eq. (6.8) we get

$$\sum_{i\in[k]}\sum_{j\in[n]} \hat{\xi}_{i,j}\omega_j \left\|x^{(j)} - c^{(i)}\right\|_2^2 - \sum_{i\in[k]}\sum_{j\in[n]} \hat{\xi}_{i,j}\omega_j \left\|x^{(j)} - \hat{c}^{(i)}\right\|_2^2 = \sum_{i\in[k]} \kappa_i \left\|c^{(i)} - \hat{c}^{(i)}\right\|_2^2 > 0,$$

which in conclusion yields the claim. □

6.2.3 A revised Balanced k-Means Algorithm

We may now use the latter findings in order to obtain a revised version of the Balanced k-Means-algorithm in form of Algorithm 2.

For the revised Balanced k-Means-algorithm as stated by Algorithm 2 we obtain the following convergence result. This extends the analogous result in [BBG17] as Algorithm 2 guarantees convergence to a local optimum. Of course, due to the degenerate nature of the examples in which Algorithm 1 may not reach a local optimum, this is more of theoretical than of practical interest. Also, we can slightly improve the bound for the maximum number of iterations from $\left(\frac{40e\cdot k^2 n}{d}\right)^{(d+1)k-1}$ to $\left(\frac{4e\cdot(k-1)n}{d}\right)^{(d+1)k-1}$.

Algorithm 2: Balanced k-Means Algorithm (Revised)

Input: $x^{(1)}, \dots, x^{(n)} \in \mathbb{R}^d$ pairwise distinct,
$\omega_1, \dots, \omega_n \in \mathbb{R}_{>0}, \kappa_1, \dots, \kappa_k \in \mathbb{R}_{>0}, s^{(1)}, \dots, s^{(k)} \in \mathbb{R}^d$

1 **Repeat**
2 Obtain an optimal solution $\xi = (C_1, \dots, C_k)^\intercal$ from the *relative interior* of the optimal face of (4.18) w. r. t. $f_i(x) = \left\|x - s^{(i)}\right\|_2^2$, $i \in [k]$, and according to the remaining input.
3 **if** $c(C_i) = s^{(i)} \quad \forall i \in [k]$ **then**
4 **if** *ξ is the unique solution of* (4.18) **then**
5 **return** ξ, $s^{(1)}, \dots, s^{(k)}$
6 **else**
7 Determine a cycle in $G(\xi)$ and let ζ be the corresponding cyclic exchange.
8 $\lambda \leftarrow \min\{\omega_j \xi_{i,j} : i \in [k], j \in [n] \text{ with } \zeta_{i,j} \neq 0\}$
9 $\xi \leftarrow \xi + \lambda\zeta$
10 **forall** $i \in [k]$ **do** $s^{(i)} \leftarrow c(C_i)$

Theorem 6.14
Algorithm 2 terminates in a local optimum of Balanced k-Means *after at most* $\left(\frac{4e\cdot(k-1)n}{d}\right)^{(d+1)k-1}$ *iterations.*

Proof. We first assure that Algorithm 2 yields a strictly decreasing sequence w. r. t. the objective of Balanced k-Means.

Let $\xi \in T_{\kappa,\omega}$ be the solution obtained in line 2 and set $d := (\kappa_1(s^{(1)})^\intercal, \dots, \kappa_k(s^{(k)})^\intercal)^\intercal$. By updating the clustering this way but keeping the sites, we get a solution that is as least as good as the previous iteration's one.

Now assume that the centroidal property check in line 3 returns true but the solution is not returned as it is not unique. As we choose ξ from the relative interior of the optimal face, this is equivalent to the clustering ξ not being a vertex of $T_{\kappa,\omega}$. According to Theorem 2.12, this means that the clustering graph $G(\xi)$ does indeed contain a cycle, so the cyclic exchange ζ is well-defined. Furthermore, due to $G(\zeta) \subset G(\xi)$ we obtain $\lambda > 0$ in line 8. As by definition $\zeta_{i,j} \in \left\{0, \pm\frac{1}{\omega_j}\right\}$ for all $i \in [k]$ and $j \in [n]$, it follows that $\xi + \lambda\zeta \geq 0$ and in particular $\xi + \lambda\zeta \in T_{\kappa,\omega}$. In particular, once more as $G(\zeta) \subset G(\xi)$ and as $\xi \in \operatorname{relint}(F(T_{\kappa,\omega}, A^\intercal d))$, we get $\xi + \lambda\zeta \in F(T_{\kappa,\omega}, A^\intercal d)$. However, as noticed before, the assumption of pairwise distinct points implies that the cyclic exchange must result in a shift of centroids, i. e., $A\zeta \neq 0$. Thus, the updated clustering

obtained in line 9 is still an optimizer of (4.18) but with centroids that are not identical to the current sites.

Hence, whenever line 10 is reached, the update of the sites will strictly improve the objective value of Balanced k-Means according to Eq. (6.8).

Recall that the objective value of Balanced k-Means only depends on the resulting centroids (cf. Eq. (6.4) and the derivation of Balanced Max Centroid Variance). This implies that no centroid vector $c = (c(C_i))_{i\in[k]} = A\xi$ for a clustering $\xi \in T_{\kappa,\omega}$ can appear in more than one iteration. From Eq. (6.20) we know that the set of optimal solutions of problem (4.18) is $F(T_{\kappa,\omega}, A^\intercal d)$ with $d = (\kappa_i s^{(i)})_{i\in[k]}$. Thus, line 2 may select any clustering $\xi \in \operatorname{relint}(F(T_{\kappa,\omega}, A^\intercal d))$ and thus a centroid vector $A\xi \in \operatorname{relint}(F(G,d)) = A \cdot \operatorname{relint}(F(T_{\kappa,\omega}, A^\intercal d))$ (cf. Eq. (6.21)). As no centroid vector may appear twice, this implies that the gravity polytope's face $F(G,d)$ of the centroid vector must be distinct for each iteration, too. So, Theorem 6.8 implies that there can be at most $\left(\frac{4e\cdot(k-1)n}{d}\right)^{(d+1)k-1}$ iterations.

At termination, it holds that ξ is the unique optimizer of (4.18) for $s^{(i)} = c(C_i)$ for $i \in [k]$ and thus the unique optimizer of problem (6.19) for $a_i = -c(C_i)$, $i \in [k]$ (cf. (6.20)). Hence, Theorem 6.10 states that ξ is a local optimum. □

The revised Balanced k-Means-algorithm offers a very practical possibility to obtain a local optimum. We have seen in Section 6.1.2 that Balanced k-Means is in fact a hard problem and its counterpart Balanced Max Centroid Variance even $\mathcal{APX}$-hard. However, the latter takes the form of a norm maximization problem over the gravity polytope as stated by Eq. (6.16).

We are given an implicit representation of the gravity polytope G as a linear transformation of the clustering polytope $T_{\kappa,\omega}$. In this view, the linear program (6.19) can be understood as a *(strong) optimization oracle* for G. This means, given any vector $d \in \mathbb{R}^{dk}$ we can solve (6.19) in order to either decide that $G = \emptyset$ or to find $c^* \in G$ such that c^* maximizes $d^\intercal c$ over G. As sufficient for this scope which shall only provide a brief review of existing results, we disregard the issue of numerical errors and assume an infinite-precision arithmetic for the following complexity statements. We refer to [GLS93] and [Pap94, Section 14.3] for a comprehensive treatment of oracle-based algorithms and their running time analysis.

The polynomial-time approximability of radii maximization, including the Euclidean norm, over convex and compact sets (short: *bodies*) that are represented by an oracle has been thoroughly researched ([Bri+98], [Bri+01]). Here, a polynomial-time, or oracle-polynomial-time, algorithm can be understood as a polynomial-time algorithm (in the classical sense of a Turing machine model) that can further call the provided oracle at only the computational cost of writing and reading its corresponding in- and outputs, respectively (cf. [GLS93, Section 1.2]). Brieden et al. [Bri+01] consider both deterministic and randomized approximation algorithms. They show that for

both deterministic as well as randomized[2] algorithms, any oracle-polynomial time algorithm that maximizes the Euclidean norm over an oracle-represented body must have an accuracy in $\Omega(\sqrt{\frac{log(d)}{d}})$ with d being the dimension ([Bri+98, Theorem 3.2, Theorem 4.2]). Here, an algorithm approximates the Euclidean maximum norm with an accuracy $r \in [0,1]$ if for the algorithm result $\phi_{\mathcal{A}}(K)$ for a body K it holds that $r \cdot \max_{x \in K} \|x\|_2 \leq \phi_{\mathcal{A}}(K) \leq \max_{x \in K} \|x\|_2$. At the same time, they provide both deterministic and randomized algorithms that do solve the problem with an accuracy in $\mathcal{O}\left(\sqrt{\frac{log(d)}{d}}\right)$ ([Bri+01, Theorem 3.2.1, Theorem 2.3]). However, it is pointed out in [Bri+01] that although randomization does not help to increase the asymptotic accuracy, their proposed randomized algorithm is superior w. r. t. the estimated running time bounds.

The randomized algorithm now obtains a polyhedral approximation of the Euclidean unit ball from a uniform sample of extreme points of its polar, i. e., the unit sphere. Let us briefly explain the key idea and sketch its complexity analysis. Here, one crucial figure is the *s-cap ratio*, $\tau(d,s)$.

For $s \in (0,1)$, a *s-cap* is given as the intersection of the sphere with a half-space at distance s from the origin, i. e., $C(s,a) := \mathbb{S}^d \cap H^{\geq}_{(a,s)}$ for $a \in \mathbb{S}^{d-1}$. For any dimension d, $\tau(d,s)$ then denotes the ratio of surface volumes of an s-cap to the whole sphere, i. e., $\tau(d,s) := \frac{\mathrm{vol}_{d-1}(C(s,a))}{\mathrm{vol}_{d-1}\left(\mathbb{S}^{d-1}\right)}$ for any $a \in \mathbb{S}^{d-1}$. Observe that for all $a \in C(s, \frac{x}{\|x\|_2})$, it holds by definition that $a^\intercal x = \|x\|_2 a^\intercal \frac{x}{\|x\|_2} \geq s \|x\|_2$. Now let S be a set of uniformly and independently distributed samples from the unit sphere. Let further $x^* \in K$ be a norm-maximizer of K. Then it can be shown that if $|S| \geq 2\lceil \tau(d,s) \rceil$, with probability greater than $6/7$ it holds that $C(s, \frac{x^*}{\|x^*\|_2}) \cap S \neq \emptyset$ and thus an accuracy of s is obtained ([Bri+98, Theorem 4.2]). Furthermore, for $s > \frac{2}{\sqrt{d}}$ it holds that $2s\sqrt{d}(1-s^2)^{\frac{1-d}{2}} < \tau(d,s) < 10s\sqrt{d}(1-s^2)^{\frac{1-d}{2}}$ ([Bri+98, Lemma 5.1]). Basic calculus then yields $\tau\left(d, \alpha\sqrt{\frac{log(d)}{d}}\right) = \Theta(\sqrt{\log(d)} d^{\frac{\alpha^2}{2}})$ for any $\alpha > 0$ (which is $\mathcal{O}(d)$ for any $\alpha < \sqrt{2}$). Note that it can also be shown that this bound is best-possible, i. e., any randomized algorithm of accuracy s must make $\Omega(\tau(d,s))$ oracle calls in the worst case ([Bri+98, Theorem 4.2]).

In conclusion, calling the optimization oracle for K for $\mathcal{O}\left(\sqrt{\log(d)d}\right)$ many uniform samples from $\mathbb{S}^{d-1}$ yields the desired approximation (with an arbitrarily high probability).

We may now apply this directly to the problem (6.16). Of course, maximizing an ellipsoidal norm is equivalent to the maximization of the Euclidean norm in an

[2] A randomized algorithm must solve the problem (within the given accuracy) with a strictly positive probability $p \in (0,1]$. This probability can then be boosted to reach an arbitrary probability p^* by simply repeating the algorithm $\frac{\log(1-p^*)}{\log(1-p)}$ many times (cf. [Pap94, Chapter 11]).

accordingly transformed space. Also, due to Lemma 6.5 and using the assumption $c(X) = 0$, the gravity polytope lives in the $k(d-1)$-dimensional linear subspace

$$L := \left\{ (c^{(i)})_{i\in[k]} \in \mathbb{R}^{dk} : \sum_{i\in[k]} \kappa_i c^{(i)} = 0 \right\}.$$

With the diagonal matrix $D := \operatorname{diag}(\kappa_1 \mathrm{Id}_d, \dots, \kappa_k \mathrm{Id}_d) \in \mathbb{R}^{dk\times dk}$ let $U \in \mathbb{R}^{dk\times d(k-1)}$ be a matrix whose columns are an orthonormal basis of $D^{\frac{1}{2}}L$. Finally, set

$$\epsilon := 1 - \alpha\sqrt{\frac{\log(d(k-1))}{d(k-1)}}$$

for some constant $\alpha > 0$ and let S be a set of uniform samples from $\mathbb{S}^{d(k-1)}$ sufficiently large to obtain an accuracy of $(1-\epsilon)$ with the desired probability as depicted before. Then

$$\begin{aligned}
\max_{c\in G} \sqrt{c^\intercal Dc} &= \max_{x\in D^{\frac{1}{2}}G, a\in\mathbb{S}^{dk-1}} a^\intercal x = \max_{x\in D^{\frac{1}{2}}G, a\in\mathbb{S}^{dk-1}\cap D^{\frac{1}{2}}L} a^\intercal x = \max_{x\in D^{\frac{1}{2}}G, a\in\mathbb{S}^{d(k-1)-1}} a^\intercal U^\intercal x \\
&= \max_{x\in U^\intercal D^{\frac{1}{2}}G, a\in\mathbb{S}^{d(k-1)-1}} a^\intercal x \geq (1-\epsilon) \max_{x\in U^\intercal D^{\frac{1}{2}}G, a\in S} a^\intercal x \\
&= (1-\epsilon) \max_{c\in G, a\in D^{-\frac{1}{2}}US} a^\intercal Dc = (1-\epsilon) \max_{\substack{\xi\in T_{\kappa,\omega}, \\ (a_i)_{i\in[k]}\in D^{-\frac{1}{2}}US}} \sum_{i\in[k]}\sum_{j\in[n]} \xi_{i,j}\omega_j \cdot a_i{}^\intercal x^{(j)}.
\end{aligned}$$

The final expression can be computed by solving (6.19) for each $(a_i)_{i\in[k]}$ in $D^{-\frac{1}{2}}US$. Doing so, we obtain a randomized $\tilde{\epsilon}$-approximation algorithm for Balanced Max Centroid Variance for $\tilde{\epsilon} := 2\epsilon - \epsilon^2$ (as our objective takes the square of the norm). For our original Balanced k-Means problem and without the assumption of centralized units, this suggests to sample sites $(s^{(i)})_{i\in[k]}$ uniformly from the ellipsoid $c(X) + D^{-\frac{1}{2}}U\mathbb{S}^{d(k-1)} \subset \operatorname{aff}(G)$ (recall that the transformation from affine to spherical power diagram parameters implies $a_i = -2s^{(i)}$ for $i \in [k]$). We may then solve the program (4.18) for every sample of sites. Of course, this can then be combined with the Balanced k-Means-algorithm for every sample in order to reach a local optimum.

As stressed in [BG12], while the actual approximation bound is, of course, a rather limiting result, norm maximization over the gravity polytope still comes with the advantage of a rather low-dimensional space of dimension $(k-1)d$. In particular, this is independent of the number of units n. Furthermore, for many of the applications of our interest the dimension d is in $\{2, 3\}$. For our main application example of electoral district design this indeed will turn out to be a very manageable problem (cf. Section 7.3.1).

6.3 Summary & Conclusion

This chapter treated the Balanced k-Means problem as a constrained version of the famous k-Means problem. Here, we recalled and summarized results from the rich existing theory of the application of power diagrams in our context.

Section 6.1 formally introduced the Balanced k-Means problem, discussed its objective and formulated the alternative Balanced Max Centroid Variance problem. In particular, we contributed some complexity results by showing that Balanced k-Means is $\mathcal{NP}$-complete even in the plane and that Balanced Max Centroid Variance is even $\mathcal{APX}$-hard.

Section 6.2 then discussed algorithmic approaches to the Balanced k-Means problem. First, we recalled the natural adaption of the famous k-Means problem. We then recalled and slightly extended results on gravity polytopes. We saw that the Balanced Max Centroid Variance problem yields a norm maximization problem over the latter. In particular, local optima can be characterized by extremal clusterings that allow a supporting centroidal power diagram (as was originally shown in [BG12]). We then used those results in order to obtain a revised version of the Balanced k-Means-algorithm which guarantees convergence to a local optimum. Finally, we recalled some results on the approximation of Euclidean norm maximization which in combination with the Balanced k-Means-algorithm yields the algorithmic approach to power diagram clustering that (among others) will be exploited in the following chapter.

Chapter 7

Electoral District Design

In this chapter, we want to demonstrate the power and versatility of our geometric clustering approach by means of the problem of electoral district design. Here, municipalities of a state have to be grouped to electoral districts. One of the dominating constraints is to achieve balanced districts, i. e., districts that contain more or less the same amount of eligible voters. However, other requirements w. r. t. shape or other characteristics of the districts yield further yet more elusive constraints. We will see that using different classes of generalized Voronoi diagrams depending on the individual emphasis of those requirements can result in very pleasing results. Note that parts of this chapter have been developed and published in [BGK17] in collaboration with Prof. Dr. Andreas Brieden and Prof. Dr. Peter Gritzmann.

Section 7.1 will introduce the problem of electoral district design. Section 7.2 will then outline the general methodology which will be used to apply our theory of geometric clustering by generalized Voronoi diagrams to this problem. Next, Section 7.3 describes how this applies to selected classes of diagrams. Finally, Section 7.4 describes the implementation and results of those approaches for the example of federal elections in Germany.

7.1 The Problem of Electoral District Design

7.1.1 Problem Description

Let us first describe the problem of electoral district design from the application point of view before providing a formalization in Section 7.1.4.

Typically, representative democratic systems require the subdivision of the state area into electoral districts. Depending on the democratic system, this subdivision may have crucial impact on the political outcome of an election. As a well-known example, the expression *gerrymandering* has been established to describe the intentional shaping of electoral districts in order to influence an election outcome. Its name roots back to the U.S. state governor Elbridge Gerry, who in 1812 made heavily use of this practice and as a side effect created a district in salamander shape ([Wika]). This, of course, mainly affects plurality voting systems. For the example of Germany, the so-called "overhang seats" (Überhangmandate) used to open the possibility of gerrymandering until an electoral law reformation in 2013 ([SeiBPB], [DeBu]). Indeed, there had been allegations that certain districts in the eastern part of Germany had been designed to disfavor a particular party ([Eis+01], [EG02]). However, even without intentionally "gerrymandered" districts, a proper district design is indispensable both due to the juridical requirements defined by electoral law as well as to cope with the intention of districts to reflect a fair representation of their citizens.

Here, a first crucial aspect is to ensure a balancing w. r. t. the district sizes in order to obey the "one man, one vote" principle. However, this implies that the natural demographic changes over time result in the need for a regular review and possible

adaption of a state's electoral districts.

For the example of Germany, the federal electoral law ([BWG]) states that the deviation of district sizes from the average should not exceed 15%. Furthermore, a deviation exceeding 25% enforces a redistricting. In comparison, the United States Supreme Court ruled in 1964 that congressional districts must be equal in size "as nearly as is practical" ([USS64] as cited in [Lev08, p. 44]). For U.S. state legislative districts a deviation of 10% is acceptable ([Lev08]).

Prior to the question about the range of feasible balancing deviations is the question how to actually compute those. How to determine the actual size of a district depends, for example, on whether you only count eligible voters or imply people without a voting right (such as children). Also, the choice of the data sources such as statistical census data or official electoral registers is obviously crucial. For the example of Germany, the Federal Constitutional Court clarified in 2012 that only eligible voters are to be considered ([Bun12]).

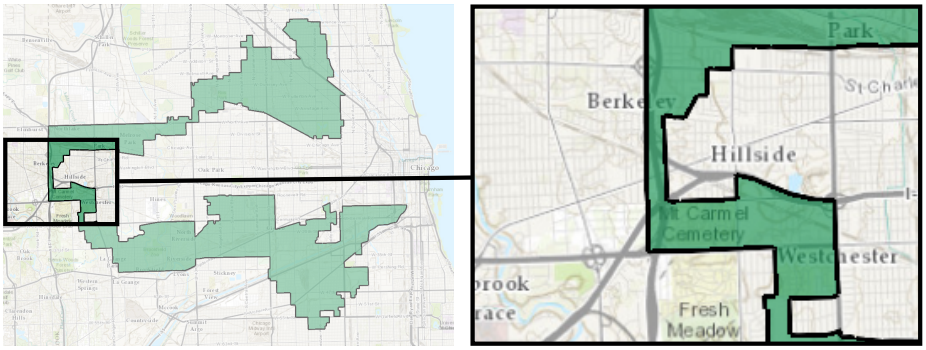

Figure 7.1: The 4th congressional district of Illinois (since 2013, [Wikb]). The northern and southern parts are only connected in the west by the Interstate 294. The underlying image is an extract from [Wik14].

A next common restriction on electoral districts is contiguity. For the example of Germany, this is explicitly demanded by the federal election law ([BWG, §3 (1)]). However, this again can be ambiguous. For example, one has to decide whether water or unincorporated areas such as natural reserves can separate a district. Another arguable example is given by the 4th U.S. congressional district of Illinois. This district consists of a northern and southern part that are only connected by an interstate highway (cf. [Ing]), see Fig. 7.1.

Another typical requirement for electoral districts is "compactness". This criterion seems reasonable as an electoral district forms a geographic unit with a common democratic representation and thus should be consolidated in some sense. Furthermore, a common hope of advocates of this criterion is the prevention of artificial district shapes that root from political partisanship (cf., e. g., [Alt98]). Horn, Hampton, and

Vandenberg [HHV93, p. 103] consider the "standard [to require compactness] as an essential preventive to gerrymandering". However, "beyond that I-know-it-when-I-see-it definition, there is little agreement about when a district is compact" [Lev08, p. 51].

Many authors have discussed how to mathematically model the juridical claim for compactness of districts. The surveys in [You88], [Nie+90] and [HHV93] categorize and compare manifold approaches to this problem. The proposed measures for compactness consider a district's area, perimeter, length and/or width (in various directions), or the degree of convexity (e.g., by comparison to its convex hull), to name just a few. Another class of measures relate to the moment of inertia (e. g. [Hes+65], [Kai66]) that measures a squared error to the center of gravity of a district (cf. Section 6.1).

For each of the measures proposed therein, [You88] provides an example for which it seems to fail to caption the juridical understanding of "compactness". Thus, the authors conclude that a "reliance on formulas has the semblance, but not the substance, of justice" [You88, p. 113]. To overcome that, both [Nie+90] and [HHV93] conclude that multiple measures should be used simultaneously.

Due to those manifold, sometimes ambiguous, and often even arguable criteria for the evaluation of an electoral district design, it is clearly not feasible to provide a single algorithmic approach to the problem that yields satisfactory results w. r. t. all requirements for any general setting. However, we will show how the geometric clustering methods proposed in the earlier chapters yield a powerful and flexible toolbox that indeed can be adjusted and applied in order to provide satisfying results.

7.1.2 Literature

The problem of electoral district design has been approached in manifold ways and is, of course, very similar to various districting-like problems. For a more complete overview, we refer to the surveys by [Tas11], [RSS13], [Kal15] and [Río20] (the latter two treating districting problems in general).

In the following, we list some of the approaches most related to ours in order to point out the overlaps and differences. Indeed, some of the proposed approaches turn out to be special cases of our approach using generalized Voronoi diagrams, e. g., as they implicitly or explicitly use a fixed diagram type. However, to the best of our knowledge, the general framework of intentionally applying generalized Voronoi diagrams and exploration of the resulting flexibility has first been proposed in [BGK17].

Many approach strategies with no or only very little methodological relation to ours have been proposed, too. Those include the use of enumerative techniques (e. g., [GN70], [Nyg88]), multi-kernel growth (e. g., [Bod73]), successive dichotomies ([KNS05]), meta-heuristics (e. g., [Alt98], [BLP05], [RS08], [RF09]), or column generation ([MJN98]), to name just a few.

Hess et al. [Hes+65] are considered the first (according to [RSS13]) to provide an optimization model for electoral district design. They already model the problem as

a transportation problem with binary variables and use squared Euclidean distances to some fixed sites as costs. They furthermore describe an improving procedure very similar to the BALANCED k-MEANS procedure as proposed in Chapter 6 that we will exploit in Section 7.3.1. Hence, they aim to minimize the moment of inertia. However, they neither provide any claims about the characteristics of the resulting districts, nor a convergence result for their proposed algorithm.

While focusing rather on the design of service territories than electoral districts, Segal and Weinberger [SW77] use shortest-path distances defined via a contiguity graph in order to solve a minimum cost flow problem to generate their districts. However, they aim at basic solutions which lead to possibly non-contiguous districts, as we will discuss in Section 7.3.2. Zoltners and Sinha [ZS83] follow a similar approach by also using shortest-path distances in a binary model. They furthermore introduce precedence constraints in order to ensure contiguity. They solve their model by a Lagrangian procedure which is, naturally, similar to our linear programming duality approach (see also Section 7.3.2).

Li and Wang [LW07] describe a non-convex quadratic programming model that uses shortest-path distances over the contiguity graph and aims to minimize the sum of intra-cluster distances in districts. However, beyond showing practical results for the example of congressional districts for the State of New York, they do not discuss their model any further.

Marlin [Mar81] considers districting problems in a more general sense and formulates several transportation problems with varying objectives including squared Euclidean distances (as in [Hes+65]) but with no further theoretical results on the characteristics of the resulting districts.

Hojati [Hoj96] extends the approach of [Hes+65] by replacing the balanced k-means procedure for improving the choice of sites with a Lagrangian relaxation approach. He furthermore states the *split resolution problem* in order to optimally round the fractionally assigned units, which he proves to be $\mathcal{NP}$-hard.

Schröder [Sch01] revises the Lagrangian approach from [ZS83]. Here, he furthermore observes that the usage of squared Euclidean distances leads to convex cells that contain the districts. Furthermore, as in [ZS83], he considers shortest-path distances and discusses the resulting contiguity of clusters (see also Section 7.3.2). In the major part of his work, however, he states different variants of the split resolution problem and discusses their complexity (which we will discuss in more detail in Section 7.2.2).

George, Lamar, and Wallace [GLW97] solve a maximum flow problem that also resembles the transportation problem from [Hes+65]. They implement various kinds of distance functions and imply the districts' balancing constraints by means of penalty terms.

Not surprisingly, Voronoi diagrams or generalizations have also explicitly appeared in the context of electoral district design. In her master's thesis, Miller [Mil07] describes a heuristic approach that aims to create centroidal Voronoi diagrams that yield the

electoral district design. However, she does not consider any balancing constraints. Fryer and Holden [FH11] already exploit centroidal power diagrams for electoral districting. They recognize those to be optimal w. r. t. the moment of inertia and use a gradient descent procedure based on the results of [AHA98]. This procedure is then combined with an iterative search for sites similar to the balanced k-means approach. As outlined in Chapter 6, centroidal power diagrams have also been successfully applied to the problem of farmland consolidation (see [BG04], [BBG11], [BBG14], [BG12]). Here, the problem is to redistribute farmland among farmers in order to obtain more consolidated and hence more efficiently farmable areas. Of course, this redistribution has to be done in a balanced way, so the task is very similar to the electoral districting problem. We already outlined the underlying theory in Chapter 6 and will apply it in Section 7.3.1.

Ricca, Scozzari, and Simeone [RSS08] use multiplicatively weighted Voronoi diagrams over a contiguity graph and describe a heuristic that iteratively adapts the multiplicative weights in order to obtain balanced districts. They claim that their procedure produces contiguous districts as well as provide theoretical convergence results for their procedure. However, they report a "lack of population equality" ([RSS08, p.1476]) in their results.

Our practical results consider the example of Germany. This has also been treated by Goderbauer [God16] (see also [God14]). He describes a heuristic divide and conquer approach that yields refining partitions along three administrative levels of Germany.

Our general framework now demonstrates a relation between many of the listed approaches. While many authors (explicitly or implicitly) model the electoral district design problem as a transportation problem, we are not aware of the explicit exploit of its strong relation to diagrams. By reconsidering some of the mentioned approaches in the context of our methodology, we are able to provide useful insights on the algorithmic approach and resulting solutions. On the other hand, we naturally rediscover some independent results. One example is the limited number of fractionally assigned components in any transportation-based approach (cf. Section 2.2.2). For the electoral districting problem or related applications, this has been mentioned in [SW77], [Mar81], [ZS83], [Hoj96], [GLW97], [KNS05], and [BG12] (without claiming completeness).

7.1.3 A Toy Example

Before we provide a formal problem description, let us prepare a toy example that will be useful to illustrate some difficulties in formally stating and solving the electoral district design problem.

Consider Fig. 7.2a that depicts a fictive state that is subdivided into a total of 42 municipal areas. These have to be grouped to form a given amount of, let's say, $k = 4$ electoral districts. For the sake of simplicity, let us assume that each of the municipalities is inhabited by a single voter. Consequently, every district should contain (approximately) $10^1/_2$ voters.

It seems fairly accurate to consider a state territory to be a subset of the (flat) Euclidean plane. However, as Fig. 7.2a illustrates, this subset may not be convex and may even contain holes (caused by lakes or city states, for example).

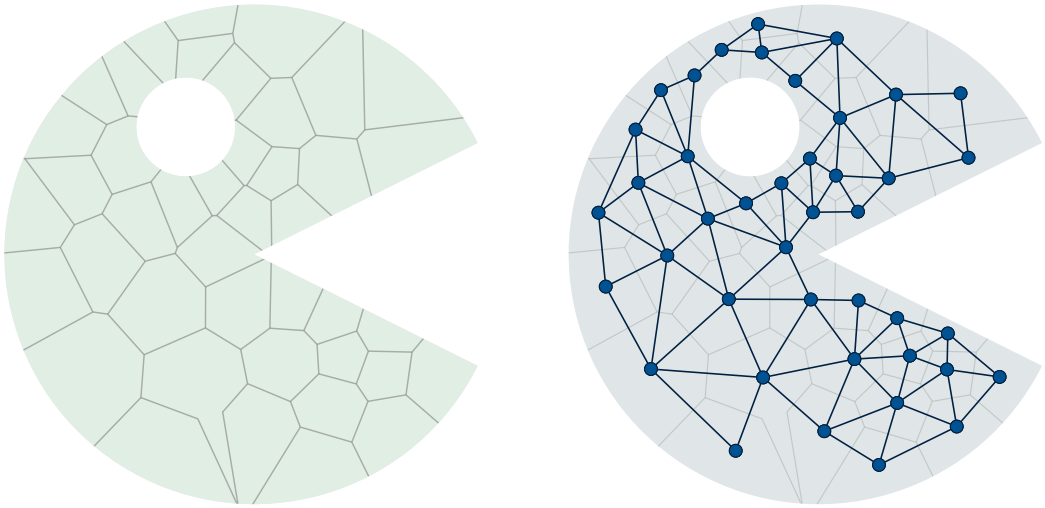

(a) Exemplary state territory with municipal areas that shall be grouped into districts.

(b) Resulting unit points (as centroids of the areas) and the contiguity graph obtained from municipalities with common border.

Figure 7.2: Toy example that will be used to discuss the different diagram approaches in the context of electoral district design.

As is the case for basically all geometrically driven approaches from the literature listed in Section 7.1.2, we will reduce each municipality to a single point in the plane (for example, given by its territory's centroid). This, of course, implies that we lose some information, in particular about the neighborhood of the corresponding municipality that defines contiguity. However, as pointed out in Section 7.1.1, even without this reduction, the definition of contiguity can be ambiguous. For example, municipalities that are separated by water and hence do not share a common boundary may be connected via a ferry service and thus should be considered contiguous from the application point of view. Consequently, in order to be able to decide on the contiguity of a district plan, we can assume to be provided with a *contiguity graph* that encodes this information. Figure 7.2b depicts the resulting instance for our toy example. Each area is reduced to its centroid, and edges between the resulting centroids state contiguity of the corresponding municipalities.

As we intend to apply the methodology of generalized Voronoi diagrams to the electoral district design problem, we have to decide on suitable distance functions that

define the class of resulting diagrams and their geometric characteristics. The usage of Euclidean distances seems accurate as many distance-oriented compactness measures (cf. Section 7.1.1) are in fact Euclidean. However, the usage of the Euclidean metric may fail when it comes to holes or non-convex notches in the underlying territory. In those cases, replacing the Euclidean by geodesic distances can be a solution. Also other distances, such as traveling times between the municipality centers, yield plausible and interpretable measures.

7.1.4 Problem Formulation

Now, let us provide the missing formal problem description. We assume to be given a set of n *units* that represent the administrative entities (such as municipalities) that have to be grouped to districts. As argued in the preceding section, we assume those units to be given by pairwise distinct points in the plane, so we denote by $X = \left\{x^{(1)}, \dots, x^{(n)}\right\} \subset \mathbb{R}^2$ the set units. Each unit $x^{(j)}$ is associated with a weight $\omega_j \in \mathbb{N}$. Furthermore, we assume a contiguity graph $G = (X, E, c)$ to be given. Here, we assume that $\left\{x^{(a)}, x^{(b)}\right\} \in E$ for $a \neq b \in [n]$ if and only if the units $x^{(a)}$ are $x^{(b)}$ are considered to be neighbors (i. e., the union of the two associated administrative entities yields a contiguous area). Furthermore, the edge weight function $c : E \to \mathbb{R}_{\geq 0}$ shall represent distances between units (as used in a shortest-path diagram approach).

The n units are to be partitioned into a given number of k districts. Thus, we may encode a (possibly infeasible) electoral district plan as an integer unconstrained clustering $\xi \in (C_1, \dots, C_k)^\intercal \in T_{\text{uncstr.},k,n} \cap \{0,1\}^{k \times n}$ (cf. Section 2.1). Recall that for $i \in [k]$ we denote by $\omega(C_i) = \sum_{j \in [n]} \xi_{i,j} \omega_j$ the weight of cluster i (see Section 4.4).

Now every district should be of approximately equal population. Here, we denote the targeted average district weight as

$$\overline{\kappa} := \frac{1}{k} \sum_{j \in [n]} \omega_{j.}$$

We assume to be given a relative error tolerance $\delta \in [0, 1]$ that defines the maximum absolute relative deviation of a district's total weight from the average such that a district can still be considered legally feasible.

In order to formulate the electoral district design problem as an optimization problem, let $g : T_{\text{uncstr.},k,n} \to \mathbb{R}$ be a function such that $g(\xi)$ evaluates the quality of a district plan ξ. However, as pointed out in Section 7.1.1, the evaluation of an electoral district plan strongly depends on the democratic system and legal boundaries. Even given a particular instance, this can be highly discussed and may not result in a single-criterion optimization problem (cf. [You88], [Nie+90]). This thesis does not aim at answering the question about which optimization criterion might or might not be superior to others. Instead, we focus on how our methodology of generalized Voronoi diagrams

can be applied to the electoral district design problem in general and how the choice of the diagram classes can be utilized in order to stress certain optimization criteria. Thus, let us for a moment assume some objective function g to be given but not further characterized.

We may then formulate the Electoral District Design problem (EDD) as follows:

$$\max_{\xi \in \mathbb{R}^{k \times n}} \; g(\xi) \tag{7.1}$$

$$\text{s.t.} \quad \xi \in T_{\text{uncstr.},k,n} \tag{7.1a}$$

$$\xi \in \{0,1\}^{k \times n} \tag{7.1b}$$

$$\omega(C_i) \in [(1-\delta)\overline{\kappa}, (1+\delta)\overline{\kappa}] \quad \forall i \in [k] \tag{7.1c}$$

$$G[\text{supp}(C_i)] \text{ connected} \quad \forall i \in [k] \tag{7.1d}$$

Complexity

Due to its general structure it is not surprising that the EDD states a hard problem, even when we drop the unspecified objective g and only focus on the remaining feasibility problem. Indeed, several well-known problems can be reduced to sub-problems of the EDD that result from restricted classes of instances. As just stated, we may only consider the corresponding feasibility problem (e. g., by setting $g = 0$). We may drop the balancing constraints (7.1c) if we only allow instances with a δ that is sufficiently large. Next, we may drop the contiguity constraints (7.1d) by demanding that the graph G is complete. Furthermore, we might incorporate those relaxed constraints in the objective g instead.

Altman [Alt97] provides $\mathcal{NP}$-hardness results for several such sub-problems. Let us briefly describe those and some others of interest in our context (cf. [Alt97, Table 3], see also [RSS08] and [God16] for further variants).

It is straightforward to check that already the feasibility problem associated with the EDD is $\mathcal{NP}$-hard. For $k = 2$, $\delta = 0$, and relaxed contiguity constraints, i. e., for G complete, it coincides with the well-known Partition problem ([GJ79, p. 47]).

The incorporation of the contiguity constraints (7.1d) introduces some further complexity. For any fixed $k \geq 3$, $\delta = 0$ and $\omega_j = 1$ for all $j \in [n]$, the problem reduces to finding a balanced partition of the nodes of G such that each resulting subgraph is connected. This is shown to be $\mathcal{NP}$-hard even if G is bipartite by Dyer and Frieze [DF85, Theorem 2.2]. De Simone et al. [De +90] show that (the decision variant of) finding a partition of a node-weighted graph into k connected subgraphs, such that the sum of deviations from the average partition size is minimized, is $\mathcal{NP}$-hard, even if the graph is restricted to be a tree with at most one node of degree greater or equal than 3. Chataigner, Salgado, and Wakabayashi [CSW07] furthermore show that for any fixed $k \geq 2$ the same task under maximizing the weight of the smallest partition

for a k-connected graph is $\mathcal{NP}$-hard even in the strong sense.

If we accept the moment of inertia as measure for compactness (cf. Section 6.1) and relax all constraints, we have the classical k-Means problem which is hard even in the plane ([MNV09]). Similarly, other compactness measures under relaxed constraints lead to classical hard problems, such as k-Center ([MS84]) when minimizing the maximum circumradius of districts, or k-Median when minimizing the sum of distances to a district center ([KH79], [MS84]).

For a slightly different problem setting, Puppe and Tasnádi [PT09] show that the problem of gerrymandering is $\mathcal{NP}$-hard. They are given a set system of the voters whose voting decision is assumed to be known and search for a subsystem that yields a partition of all voters. They then want to maximize the number of districts that a favored party wins and show hardness by a reduction of Set Packing ([GJ79, SP3, p. 221]).

7.2 Districting by Generalized Voronoi Diagrams

In the following, we outline how the theory of Chapter 4 shall be exploited for the EDD and depict the general methodology in Section 7.2.1. Here, the split resolution problem occurs as a subproblem and is discussed in more detail in Section 7.2.2.

7.2.1 General Methodology

Figure 7.3 is to illustrate the basic schema that we will adapt for the particular approaches.

Choice of Diagram Class and Parameters

For any approach, we consider a particular class of generalized Voronoi diagrams. Recall from Section 3.2 that we typically consider generator functions in the form of Eq. (3.2), i. e., given by a transformed and additively shifted (local) distance measure to some reference site. We have seen in Section 3.3 how the choice of the underlying metrics and transformation results in different classes of generalized Voronoi diagrams. For such a class, we may choose a parametrization $\mathcal{F} : \mathcal{S} \to \left(\mathbb{R}^X\right)^k$ that maps each choice of parameters $S \in \mathcal{S}$ to functions $f_i : X \to \mathbb{R}$ for $i \in [k]$. Here, the parameter set $\mathcal{S}$ typically consists of the choice of the sites $(s^{(i)})_{i\in[k]}$ that serve as cluster representatives. However, it may, for example, as well include the choice of the parameters that determine the local metrics $(d_i)_{i\in[k]}$.

With a diagram class fixed, we have to choose parameters S that determine a generalized Voronoi diagram. Those may not yet lead to a feasible districting plan, but can be considered to determine the eventual clustering's main structure. In which way, by what intention and how well this can be done, strongly depends on the chosen

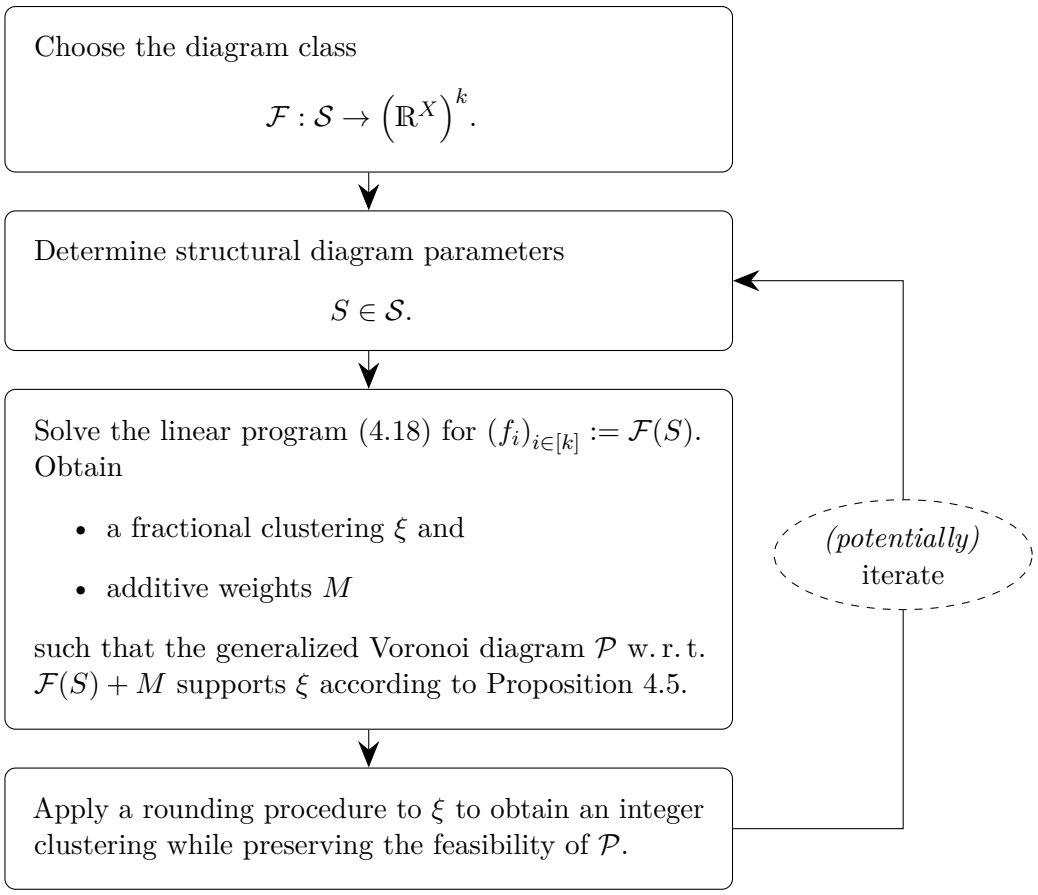

Figure 7.3: General methodological schema applied in the different approaches.

diagram class. We will describe different possibilities in our proposed approaches in Section 7.3.

Constrained Clustering Subroutine

As a subroutine, we then solve the linear program (4.18). Here, we inject the functions $(f_i)_{i\in[k]} = \mathcal{F}(S)$ into the objective, use the instance's unit weights $(\omega_j)_{j\in[n]}$ as such, and set cluster weights $\kappa_i := \overline{\kappa}$ for $i \in [k]$. Doing so, we obtain an optimal fractional clustering ξ. Following Proposition 4.5, the corresponding dual solution then yields weights such that the generalized Voronoi diagram w. r. t. functions $\mathcal{F}(S)$ disturbed by those weights is feasible for ξ (or even supporting, if ξ is chosen from the relative interior of the optimal face of (4.18)). To be more precise, let $M = (\mu_i)_{i\in[k]}$ be the part of an optimal dual solution that corresponds to the balancing constraints (4.18b).

Then the functions $\mathcal{F}(S) + M = (f_i + \mu_i)_{i\in[k]}$ yield the desired diagram.

Rounding Procedure

Assume that we have solved the program (4.18) and hence obtained a clustering $\xi = (C_1, \dots, C_k)^\intercal$ together with a feasible (or even supporting) generalized Voronoi diagram $\mathcal{P} = (P_1, \dots, P_k)$. Of course, solving (4.18) implies that we have relaxed the integrality constraints (7.1b) as we allow fractional clusterings. Consequently, we have to apply some rounding procedure in order to get an actual solution for the districting problem.

Of course, we would like to do so without losing the diagram's feasibility. Recall that the feasibility of a diagram is equivalent to the inclusions $\operatorname{supp}(C_i) \subset \left\{j \in [n] : x^{(j)} \in P_i\right\}$ for all $i \in [k]$ and that the diagram is supporting if and only if equality holds. Hence, we would like to choose an integer clustering $\hat{\xi}$ from

$$R(\xi) := \left\{\hat{\xi} \in \{0,1\}^{k\times n} : \mathbb{1}^{(k)}\hat{\xi}^\intercal = \mathbb{1}^{(n)} \ \wedge \ \operatorname{supp}(\hat{C}_i) \subset \operatorname{supp}(C_i) \ \forall i \in [k]\right\}. \tag{7.2}$$

By definition, the diagram $\mathcal{P}$ will be feasible for any $\hat{\xi} \in R(\xi)$. However, we will surely lose its potential supporting property if $\hat{\xi} \neq \xi$, as any formerly fractionally assigned unit must lie in all cells that correspond to its assigned clusters.

Note that for the case $\omega_j = 1$ for all $j \in [n]$ and $k|n$, every extremal solution of the transportation problem (4.18) is in fact integral (cf. [KW68, Corollary 1]). As mentioned in Section 7.1.2, the general case still yields that any extremal solution provides a fractional clustering with at most $k-1$ non-integrally assigned units. Again, this is due to the nature of the underlying transportation polytope (cf. [KW68, Corollary 3]). In our context, this is also provided by Corollary 2.8 as we are in the homogeneously single-constrained case. Furthermore, Theorem 2.12 tells us that the clustering graph (see Section 2.3) corresponding to any extremal solution of (4.18) is acyclic and furthermore a tree in the non-degenerated case. The tree structure can now be exploited in a rounding procedure that minimizes the resulting maximum absolute deviation of rounded cluster weights from the desired average value $\overline{\kappa}$, while not losing the feasibility of the diagram $\mathcal{P}$. Assuming that we are given an extremal solution of (4.18), we have to decide for "only" up to $k-1$ fractionally assigned units which cluster they should be assigned to. This issue has been given the name Split Resolution problem in the literature ([Kal15]). We will discuss this in more detail in Section 7.2.2.

However, we will also see that we are not always in the comfortable situation in which the optimum of (4.18) is unique and hence an extremal solution of (4.18) is in the relative interior of the optimal face. The following continuation of Example 4.3 will illustrate such a situation. This example also stresses the necessity to generate a supporting diagram in the second step of our methodology. For a diagram type as given in the example, the ambiguity of the optimum of (4.18) leads to a situation

in which the diagram is still feasible for any rounded solution from $R(\xi)$. However, pairwise cell intersections $P_i \cap P_l$ (and thus $\operatorname{supp}(C_i) \cap \operatorname{supp}(C_l)$) can contain multiple units. Therefore, no geometric claims about how those units are clustered in a rounded solution can be made. In particular, this can harm the contiguity of the resulting clusters. In such a case, a more careful rounding procedure is required. We will face this issue for shortest-path diagrams in Section 7.3.2.

Example 4.3 (continuing from p. 83).
Let us reconsider Example 4.3 from Section 4.1. Here, we considered a generalized Voronoi diagram w. r. t. functions $f_i(x) := \left\| x - s^{(i)} \right\|_1$ for $i = 1, 2$. This diagram supported a fractional clustering of sixteen points, see Fig. 4.2. We may easily interpret this example as an EDD instance. First, let us assume that each point has a weight $\omega_j = 1$ for $j \in [16]$. Thus, each cluster is to be assigned a weight of $\overline{\kappa} = 8$. For the depicted clustering of Fig. 4.2 — let us denote it ξ — we may furthermore assume that each of the four points in the top right and bottom left corner, respectively, are integrally assigned to their respective clusters, as depicted. The remaining 8 points shall each be equally assigned to both clusters by the same fraction $\frac{1}{2}$.

Hence, the clustering ξ is not only feasible for the resulting problem (4.18), but due to Proposition 4.5 even optimal. For the corresponding dual weights, we can thus conclude $M = 0$. Even more, as the diagram supports ξ, it must be contained in the relative interior of the optimal face of (4.18).

However, by construction we are in fact in the totally unimodular case, i. e., all extremal solutions of the program (4.18) will be integral. More precisely, we can conclude that $R(\xi)$ equals the set of all optimal extremal solutions.

As the considered functions f_i are of the form (3.3), Lemma 3.3 yields that every cell of the diagram is site-star-shaped, i. e., the line from any point of a cell to the corresponding site is contained in the respective cell. From a practical point of view, this more or less guarantees contiguity (cf. the issues discussed in Section 7.1.3), but only in case that the diagram supports the clustering. However, we lose the supporting property when choosing a $\hat{\xi} \in R(\xi)$. Figure 7.4 depicts two possible choices for $\hat{\xi} \in R(\xi)$. While the diagram is feasible for both choices (and non-supporting for either one), Fig. 7.4a clearly fails w. r. t. contiguity, in contrast to Fig. 7.4b.

For an extremal solution, on the other hand, Lemma 2.9 guarantees that $|\operatorname{supp}(C_i) \cap \operatorname{supp}(C_l)| \leq 1$ holds for any two clusters $C_i \neq C_l$, so this effect does not appear even after rounding.

The following theorem shows that for a wide class of generalized Voronoi diagrams, we can fairly assume the solution of (4.18) to be unique. Note that this can be regarded as a generalization of [BG12, Lemma 4.1] and has been published with almost identical form and proof in [BGK17].

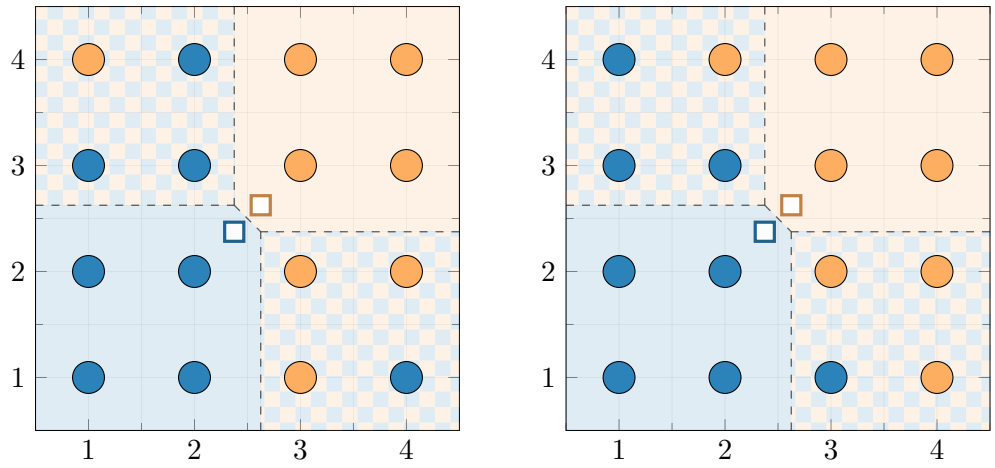

(a) Losing contiguity. **(b)** Preserving contiguity.

Figure 7.4: Rounded solutions in Example 4.3.

Theorem 7.1 ([BGK17, Theorem 7])
Consider the linear program (4.18) *for* $\omega_j > 0$ *for* $j \in [n]$, $\kappa_i \in \mathbb{R}_{\geq 0}$ *for* $i \in [k]$ *such that* $\sum_{i\in[k]} \kappa_i = \sum_{j\in[n]} \omega_j$, *and functions* $f_i : \mathbb{R}^d \to \mathbb{R}$ *for* $i \in [k]$ *in the form of Eq.* (3.2). *Assume further that every metric* d_i *in* (3.2) *is induced by a strictly convex norm, the function* h *is continuous and injective, and that the sites are pairwise distinct, i. e.,* $s^{(i)} \neq s^{(l)}$ *for* $i \neq l \in [k]$.

Then for every $\epsilon > 0$ *there exists* $\left(\hat{s}^{(i)}\right)_{i\in[k]} \in (\mathbb{R}^d)^k$ *with* $\sum_{i=1}^k \left\|s^{(i)} - \hat{s}^{(i)}\right\| < \epsilon$ *and* $\delta > 0$ *such that for all* $\left(\tilde{s}^{(i)}\right)_{i\in[k]} \in \left(\hat{s}^{(i)}\right)_{i\in[k]} + \delta\mathbb{B}_2^{d\times k}$, *it holds that with* $\tilde{f}_i(x) := h(d_i(x, \tilde{s}^{(i)})) + \mu_i$ *for* $i \in [k]$, *the linear program* (4.18) *with objective* $\sum_{i\in[k]} \sum_{j\in[n]} \omega_j \tilde{f}_i(x^{(j)}) \cdot \xi_{i,j}$ *has a unique optimizer.*

Proof. Suppose that the solution of (4.18) with objective $\sum_{i\in[k]} \sum_{j\in[n]} \omega_j f_i(x^{(j)}) \cdot \xi_{i,j}$ is not unique. Let $T = T_{\kappa,\omega}$ be the clustering polytope that is the feasible region of (4.18) as given by Eq. (4.16).

Let $\xi \in T$ be an optimal solution in the relative interior of the optimal face of (4.18).

By Theorem 2.12, the corresponding clustering graph $G(\xi)$ contains a cycle with corresponding cyclic exchange ζ as defined in Eq. (2.6). W.l.o.g. let this cycle be $C_1, 1, C_2, \ldots, C_r, r, C_1$. By optimality of ξ, it follows that

$$0 = \sum_{i\in[k]} \sum_{j\in[n]} \zeta_{i,j}\omega_j f_i(x^{(j)}) = -f_1(x^{(1)}) + \sum_{l=2}^{r} \left(f_l(x^{(l-1)}) - f_l(x^{(l)})\right) + f_1(x^{(r)})$$

and thus

$$f_1(x^{(1)}) - f_1(x^{(r)}) = \sum_{l=2}^{r} f_l(x^{(l-1)}) - f_l(x^{(l)}) =: \alpha.$$

After inserting the definition of f_1 from Eq. (3.2) this reads as

$$h(d_1(s^{(1)}, x^{(1)})) - h(d_1(s^{(1)}, x^{(r)})) = \alpha. \tag{7.3}$$

Note that α as defined above in particular does not depend on $s^{(1)}$. We can thus interpret the set of all sites $s^{(1)} \in \mathbb{R}^d$ that obey Eq. (7.3) as the bisector of the generalized Voronoi diagram w. r. t. functions $\hat{f}_1(x) := h(d_1(x^{(1)}, x)$ and $\hat{f}_2(x) := h(d_1(x^{(r)}, x)$. By Lemma 3.6, this bisector has an empty interior. Also, due to the given assumptions, the left hand side of Eq. (7.3) is continuous in $s^{(1)}$. Hence, we can choose $\hat{s}^{(1)} \in \mathbb{R}^d$ with strictly positive distance to this bisector but sufficiently close to $s^{(1)}$. As we can do so for any cycle and there are only finitely many (in the complete bipartite graph of clusters and units), the claim follows. □

In other words, Theorem 7.1 states that for distance measures that are obtained from strictly convex norms and reasonable transformation functions, an arbitrarily small perturbation of the sites suffices in order to obtain sites in the interior of the set of sites with a unique optimum.

Iterating

After rounding, i. e., having obtained some integer clustering, we may then iterate this process. Here, the (re-)selection of the structural parameters in general can be made w. r. t. several criteria. First of all, we may not have obtained a feasible solution for the EDD yet. For the balancing constraints (7.1c), we will see in Section 7.2.2 that we may efficiently perform the rounding w. r. t. the resulting maximum deviation. However, the best guarantee to be given here will be for a relative error of $\frac{\max_{j\in[n]} \omega_j}{\overline{\kappa}}$. In case of fairly heavily weighted units, we may hence violate the constraints (7.1c) and so we must consider this in an adaption of the structural parameters S.

The contiguity constraints (7.1d) are only considered indirectly. Here, we assume the class of generalized Voronoi diagrams $\mathcal{F}$ to be chosen such that feasibility of the obtained diagram will yield contiguity of the resulting districts. However, in particular for approaches based on Euclidean distances, this can occasionally fail due to the non-euclidean underlying data (cf. Section 7.1.3).

Finally, both the choice of the diagram class $\mathcal{F}$ and the selection of structural parameters S should incorporate the (yet unspecified) objective of the EDD.

7.2.2 The Split Resolution Problem

As our methodology aims at exploiting the LP-duality based correspondence between generalized Voronoi diagrams and balanced clusterings, this comes at the price of intermediate non-integral solutions. As already stated, this issue can be reasonably controlled, provided an extremal solution of (4.18) is given. We will assume ξ to be such an extremal solution for the remainder of this section.

In order to preserve feasibility, we have to find an integer clustering $\xi^* \in R(\xi)$. While we assume that the feasibility of the diagram suffices to obey the contiguity constraints (7.1d), the rounding must consider the balancing constraints (7.1c). This suggests to minimize the resulting maximum deviation after rounding. However, depending on the actual instance, other types of objective functions can be suitable, too. In accordance with the literature ([KNS05; Sch01]), we define the following variants of the Split Resolution problem. First, one may minimize a p-norm of the vector of resulting deviations. This yields the p−Split Resolution problem:

Problem p-Split Resolution:

Input:	$\delta \in \mathbb{Q}$, $\omega_1, \ldots, \omega_n \in \mathbb{N}$, $\xi \in \text{ext}\left(T\left((\omega_1, \ldots, \omega_n), (\overline{\kappa})_{i\in[k]}\right)\right)$
Question:	Does there exist $\hat{\xi} = \left(\hat{C}_1, \ldots, \hat{C}_k\right) \in R(\xi)$ such that $\sum_{i\in[k]} \left\lvert\overline{\kappa} - \omega(\hat{C}_i)\right\rvert^p \leq \delta$?

Another interesting variant has been formulated in [Hoj96]. Here, given individual maximum deviation bounds for each cluster, the goal is to find a clustering that minimizes the number of fractional assignments that is necessary to obey those bounds:

Problem min-fract-Split Resolution:

Input:	$m \in \mathbb{N}$, $\delta_1, \ldots, \delta_k \in \mathbb{N}$, $\omega_1, \ldots, \omega_n \in \mathbb{N}$, $\xi \in \text{ext}\left(T\left((\omega_1, \ldots, \omega_n), (\overline{\kappa})_{i\in[k]}\right)\right)$
Question:	Does there exist $\hat{\xi} = \left(\hat{C}_1, \ldots, \hat{C}_k\right)^\intercal \in T_{\text{uncstr.},k,n}$ such that $\text{supp}(\hat{C}_i) \subset \text{supp}(C_i)\ \forall i \in [k]$, $\lvert\overline{\kappa} - \omega(\hat{C}_i)\rvert \leq \delta_i\ \forall i \in [k]$, and $\sum_{j\in[n]} deg_{()}G\left(\hat{\xi}\right)(j) - n \leq m$?

Here, the expression $\sum_{j\in[n]} deg(_)G\left(\hat{\xi}\right)(j) - n$ counts the number of fractional assignments (with $deg(_)G\left(\hat{\xi}\right)(j)$ being the degree of node j in the clustering graph $G\left(\hat{\xi}\right)$). There is a certain practical relevance to min-fract-Split Resolution. If the administrative units that shall form a state's electoral districts are individually too big in order to obtain a sufficiently balanced districting plan, there might be no other choice than to split some of those (for example, this can be the case for greater cities). In this case, min-fract-Split Resolution answers the question how many of those splits have to be made.

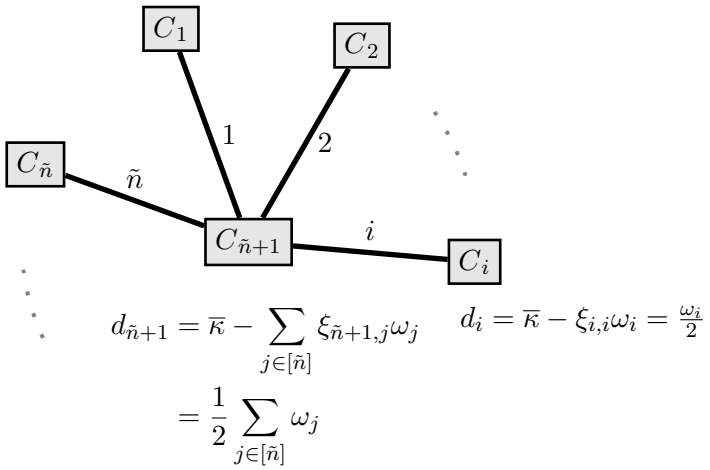

Figure 7.5: Construction for the reduction of Partition to the variants of Split Resolution (oriented at [Hoj96] and [Sch01]).

The problems min-fract-Split Resolution as well as p-Split Resolution for $1 \leq p < \infty$ have been shown to be $\mathcal{NP}$-hard by [Hoj96] and [Sch01], respectively (cf. also [KNS05]). In both cases, a very similar reduction of Partition has been applied. Let us shortly discuss this reduction. We assume an instance of partition to be given by weights $\omega_1, \ldots, \omega_{\tilde{n}} \in \mathbb{N}_{>0}$. We now construct an instance of EDD with $k = \tilde{n} + 1$ districts and $n = 2\tilde{n}$ units. For $j \in [\tilde{n}]$, we give the jth unit a weight ω_j and the $(j + \tilde{n})$th unit a weight of $\frac{1}{2}\sum_{r\in[\tilde{n}]\setminus j} \omega_r$. In particular, this yields an average cluster weight of $\overline{\kappa} = \frac{1}{2}\sum_{j\in[\tilde{n}]} \omega_j$. Then, an extremal feasible clustering is constructed by assigning each unit $j \in [\tilde{n}]$ equally to the jth and $(\tilde{n}+1)$st cluster, as well as every unit $j \in [2\tilde{n}] \setminus [\tilde{n}]$ with full extent to the $(j - \tilde{n})$th cluster. This results in the contracted clustering graph as depicted in Fig. 7.5. Each cluster $j \in [\tilde{n}]$ hence has a "demand" of $d_j = \frac{1}{2}\omega_j$ w. r. t. the fractionally assigned units, while for the $(\tilde{n} + 1)$st cluster all units have been fractionally assigned, so its demand is $\frac{1}{2}\sum_{j\in[\tilde{n}]} \omega_j$.

Following [Hoj96], we obtain an instance of min-fract-Split Resolution by setting $\delta_j := \lceil \frac{1}{2}\omega_j \rceil$ for every $j \in [\tilde{n}]$ and $\delta_{\tilde{n}+1} := 0$. This way, the first $\tilde{n}$ clusters actually do not

constrain us. Thus, we set $m = 0$ and have that this is a YES-instance of min-fract-SPLIT RESOLUTION if and only if the given PARTITION instance is a YES-instance.

Following [Sch01], we obtain an instance of p-SPLIT RESOLUTION by setting $\delta := \sum_{j\in[\tilde{n}]}(\frac{1}{2}\omega_j)^p$. Here, the first $\tilde{n}$ clusters will always create an error of $\sum_{j\in[\tilde{n}]}(\frac{1}{2}\omega_j)^p = \delta$ by construction. The $(\tilde{n}+1)$st cluster, however, will contribute 0 if and only if the given PARTITION instance is a YES-instance.

Of course, the formulation of min-fract-SPLIT RESOLUTION suffers from the possibility of individual error bounds δ_i (which is highly made use of in the reduction). To the best of our knowledge, the complexity of this problem when restricted to $\delta_1 = \ldots = \delta_k$ is not known.

On the other hand, Schröder [Sch01] shows that the case $p = \infty$, i. e., to minimize the maximum occurring deviation, is polynomially solvable. In his work, he considers several variants of more general balanced tree decomposition problems and deduces the polynomial solvability of $\infty-$SPLIT RESOLUTION as a special case (cf. [Sch01, Algorithmus 5.6, Korollar 7.10]). He also provides an adapted version that allows an improved running time (cf. [Sch01, Algorithmus 7.2, Satz 7.15]). He is able to do so by considering subproblems in form of approximate subset sum problems with precedence constraints. Those subproblems are solved by an adaption of classical dynamic programming schemes whose general pseudo-polynomial running time can be resolved as the chosen instances reveal to be sufficiently small. The final algorithm then proceeds by successively contracting parts of its input graph that yield a validly balanced partial decomposition within a prescribed maximum deviation. Finally, he applies a bisection procedure in order to determine the actual optimal value.

Due to its generality (even of the adapted version), Schröder [Sch01] does not provide a comprehensive version of his algorithm. In the following, we will present our own adaption of the algorithm in [Sch01]. Besides a self-containing formulation of the algorithm for our problem, this allows us to alter it slightly and obtain a different complexity result. The resulting complexity of the algorithm presented in [Sch01] is $\mathcal{O}\left(\delta_{\mathcal{C}}^3 k^2 \log\left(\max_{j\in[n]}\omega_j + \delta_X |J_f|\right)\right)$ ([Sch01, Satz 7.15]). Here, $\delta_{\mathcal{C}}$ and δ_X denote the maximal degrees of cluster and unit nodes in $G(\xi)$, respectively, and $J_f \subset [n]$ the set of all units that are fractionally assigned. By contrast, we can provide a bound of $\mathcal{O}\left(k^3 \log\left(\max_{j\in[n]}\omega_j\right) + kn\right)$ for our algorithm. This will be achieved by a careful running time analysis as well as by exploiting a different complexity result for the following subproblem.

An approximate variant of SUBSET SUM will appear as a basic subproblem in our algorithm. While we use the notation of [Sch01], note that we consider a less general form by enforcing non-negative weights as well as a symmetric error.

Problem Weak Subset Sum:

Input:	$\delta \in \mathbb{N}_{>0}$, $\kappa \in \mathbb{Z}$, $\omega_1, \ldots, \omega_n \in \mathbb{N}$
Question:	Does there exist $J \subset [n]$ such that $\left\lvert \sum_{j \in J} \omega_j - \kappa \right\rvert \leq \delta$?

Typically, problems in the style of Weak Subset Sum can be solved in a pseudo-polynomial time using dynamic programming. Kellerer, Pferschy, and Speranza [KPS97] provide a FPTAS for the Subset Sum problem. While at its core this implements classical dynamical programming approaches, they provide a rather sophisticated method that solves a knapsack-like problem and can be applied to Weak Subset Sum in order to obtain the following running time.

Theorem 7.2
For an instance $(\delta; \kappa; \omega_1, \ldots, \omega_n)$ of Weak Subset Sum*, we can compute a valid index set $J \subset [n]$ that solves the instance or the conclude that such a set does not exist in $\mathcal{O}\left(\min\left\{n \cdot \frac{\kappa}{\delta}, n + \left(\frac{\kappa}{\delta}\right)^2 \log\left(\frac{\kappa}{\delta}\right)\right\}\right)$.*

Proof. The algorithm proposed in [KPS97] solves the knapsack problem

$$\max\left\{\sum_{j \in J} \omega_j : J \subset [n], \sum_{j \in J} \omega_j \leq c\right\}$$

for the same input values as in our setting as well as $c \in \mathbb{N}$ with accuracy ϵ in $\mathcal{O}\left(\min\left\{\frac{n}{\epsilon}, n + (\frac{1}{\epsilon})^2 \log(\frac{1}{\epsilon})\right\}\right)$. After excluding the trivial case $\kappa \leq \delta$ we can set $c := \kappa + \delta$ and $\epsilon := \frac{2\delta}{\kappa + \delta}$ to obtain the statement (cf. also [KPS97, Theorem 7, Theorem 8]). □

Note that Schröder [Sch01] points out that the procedure from [KPS97] is unsuitable as it may not be adapted to his (more general) context which comes with certain precedence constraints (cf. [Sch01, Section 5.1.1.2]). Instead, he proposes a self-developed dynamic programming procedure which does not preserve the running time bounds of [KPS97]. However, in our situation we can exploit the simple structure of extremal clustering graphs and thus are indeed able to apply the results from [KPS97].

Description of the Algorithm

We will now outline our version of the algorithm by Schröder [Sch01].

For the sake of convenience, we cut the algorithm into three parts: Algorithm 3 splits the problem into the connected components of the contracted clustering graph $G_{\mathcal{C}}(\xi)$. Algorithm 4 describes a forward-recursion for each component. Finally, Algorithm 5 describes how an optimal rounding is determined locally. Note that all algorithms as presented here only yield the objective value but not an optimal solution. However, it

should be clear from the description that a solution can be obtained by storing and combining the solutions of the appearing subproblems in a straightforward way without any harm to the running time. In particular, due to its recursive nature, the procedure computes optimal solution values for certain subgraphs of $G_{\mathcal{C}}(\xi)$. As this requires the introduction of quite a lot of notation, Fig. 7.6 is meant to illustrate those and their referenced parts of $G_{\mathcal{C}}(\xi)$.

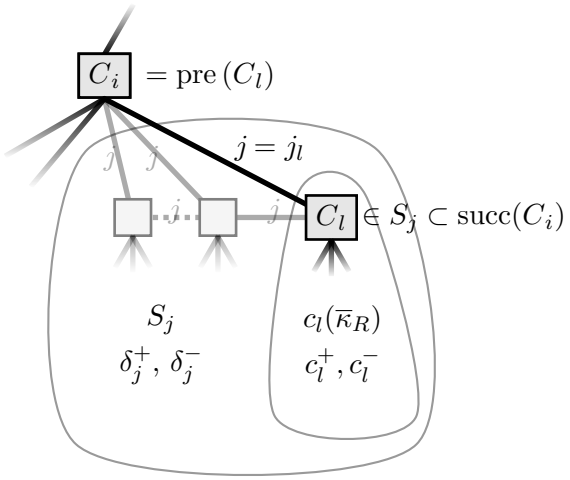

Figure 7.6: Illustration of the components of $G_{\mathcal{C}}(\xi)$ referenced by the definitions in Section 7.2.2.

We assume to be given an extremal solution $\xi \in \mathrm{ext}(T\left((\omega_1, \ldots, \omega_n), (\overline{\kappa})_{i\in[k]}\right))$. The algorithm exploits the tree and tree-like structures of the clustering graphs $G(\xi) = (\mathcal{C} \cup [n], E)$ and $G_{\mathcal{C}}(\xi) = (\mathcal{C}, E_{\mathcal{C}}, \nu)$ (cf. Definitions 2.10 and 2.11). From Theorem 2.12 we know that $G(\xi)$ is a forest. Together with the definition of the contracted clustering graph $G_{\mathcal{C}}(\xi)$, this implies that the only cycles in $G_{\mathcal{C}}(\xi)$ appear in the case that a unit is fractionally assigned to more than two clusters. In this case, the clusters sharing this unit form a clique. If we contracted those cliques, we would again obtain a forest. As a consequence, shortest paths between any two nodes in any connected component of $G_{\mathcal{C}}(\xi)$ are unique. Thus, for any connected component of $G_{\mathcal{C}}(\xi)$ we can choose any contained node C_r and *root* this component in C_r. As for trees in Section 3.3.4, we obtain a partial order on this component via

$$C_i \succeq_{G_{\mathcal{C}}} (\succ_{G_{\mathcal{C}}}) C_l \quad :\Leftrightarrow \qquad C_l \text{ lies on the shortest } C_i, C_r\text{-path (and } C_i \neq C_l)$$

for every pair of nodes C_i, $C_l \in \mathcal{C}$ that are contained in it. Due to the uniqueness of shortest paths, every node C_i except the root node has a unique *predecessor* $\mathrm{pre}(C_i)$

with $\mathrm{dist}_{G_{\mathcal{C}}(\xi)}(C_r, \mathrm{pre}(C_i)) = \mathrm{dist}_{G_{\mathcal{C}}(\xi)}(C_r, C_i) - 1$ and a set of *successors*

$$\mathrm{succ}(C_i) := \Big\{ C_l \in \mathcal{C} : \mathrm{dist}_{G_{\mathcal{C}}(\xi)}(C_r, C_l) = \mathrm{dist}_{G_{\mathcal{C}}(\xi)}(C_r, C_i) + 1 \Big\}.$$

Clearly, for any node C_i, the induced subgraph $G_{\mathcal{C}}(\xi)[\{C_l \in \mathcal{C} : C_l \succeq_{G_{\mathcal{C}}} C_i\}]$ is again rooted in C_i preserving the same order.

Algorithm 3: Solver for ∞-Split Resolution

Input: Extremal clustering $\xi \in \mathrm{ext}(T\Big((\omega_1, \ldots, \omega_n), (\overline{\kappa})_{i\in[k]}\Big))$, $k \in \mathbb{N}$, $(\omega_j)_{j\in[n]} \in \mathbb{N}^n_{>0}$

1 Compute $G_{\mathcal{C}}(\xi)$ and root the connected components (in arbitrary root nodes).
2 $\overline{\kappa} \leftarrow \frac{1}{k}\sum_{j\in[n]} \omega_j$
3 **foreach** $C_i : C_i$ *is the root of a connected component of* $G_{\mathcal{C}}(\xi)$ **do**
4 $\quad \delta_i \leftarrow$ *Algorithm 4* $\Big(\xi; G_{\mathcal{C}}(\xi); (\omega_j)_{j\in[n]}; C_i; \overline{\kappa}; \overline{\kappa}\Big)$
5 **return** $\max\{\delta_i : C_i$ is the root of a connected component of $G_{\mathcal{C}}(\xi)\}$

The proposed algorithm now first roots $G_{\mathcal{C}}(\xi)$ by choosing an arbitrary root for any connected component and then solves the $\infty-$Split Resolution for each component independently. This framework is given by Algorithm 3.

Next, Algorithm 4 performs a forward-recursion. It starts from the root node of a connected component of $G_{\mathcal{C}}(\xi)$ and recursively solves a variant of the $\infty-$Split Resolution for each subgraph $G_{\mathcal{C}}(\xi)[\{C_l : C_l \succeq_{G_{\mathcal{C}}} C_i\}]$ for an input node C_i. As a further input, it receives a target weight $\overline{\kappa}_R$. Let

$$j_l := \begin{cases} j \text{ such that } (\{C_l, \mathrm{pre}(C_l)\}, j) \in E & , \text{ if } C_l \text{ is not a root node} \\ -1 & , \text{ otherwise} \end{cases} \tag{7.4}$$

be the unit index associated with the edge to the (unique) predecessor of a cluster C_l if it is not a designated root. Algorithm 4 now computes the minimal maximum weight deviation of the clusters "below" C_i when rounding the fractional assignments. Here, every cluster except C_i should ideally be assigned a weight of $\overline{\kappa}$. The target weight for the cluster C_i is given by $\overline{\kappa}_R$ and we assume the j_ith unit (if it exists) *not* to be assigned to C_i. We can formalize this value by

$$c_i(\overline{\kappa}_R) := \min \Bigg\{ \max\Bigg\{ \Big\| \Big(\overline{\kappa} - \omega(\hat{C}_l)\Big)_{C_l \succ_{G_{\mathcal{C}}} C_i} \Big\|_\infty, \Big|\overline{\kappa}_R - \omega(\hat{C}_i)\Big| \Bigg\} : \hat{\xi} = \Big(\hat{C}_1, \ldots, \hat{C}_k\Big)^{\mathsf{T}} \in R(\xi) \ \wedge \ j_i \notin \mathrm{supp}(\hat{C}_i) \Bigg\}. \tag{7.5}$$

Algorithm 4: Solver for ∞-Split Resolution on connected subgraphs

Input: Extremal clustering $\xi \in \mathrm{ext}(T\left((\omega_1,\ldots,\omega_n),(\overline{\kappa})_{i\in[k]}\right))$, rooted contracted clustering graph $G_{\mathcal{C}}(\xi) = (\mathcal{C}, E_{\mathcal{C}}, \nu)$, $(\omega_j)_{j\in[n]} \in \mathbb{N}^n_{>0}$, $C_i \in \mathcal{C}$, $\overline{\kappa}, \overline{\kappa}_R \in \mathbb{Q}$

1 **foreach** $C_l \in \mathrm{succ}(C_i)$ **do**

2 $\quad j_l \leftarrow j \in [n]$ such that$(\{i,l\}, j) \in E_{\mathcal{C}}$

3 $\quad c_l^- \leftarrow$ *Algorithm 4* $\left(\xi; G_{\mathcal{C}}(\xi); (\omega_j)_{j\in[n]}; C_l; \overline{\kappa}; \overline{\kappa}\right)$

4 $\quad c_l^+ \leftarrow$ *Algorithm 4* $\left(\xi; G_{\mathcal{C}}(\xi); (\omega_j)_{j\in[n]}; C_l; \overline{\kappa}; \overline{\kappa} - \omega_{j_l}\right)$

5 $J \leftarrow \{j \in [n] : (\{C_i, C_l\}, j) \in E_{\mathcal{C}}, C_l \in \mathrm{succ}(C_i)\}$

6 **foreach** $j \in J$ **do**

7 $\quad \delta_j^- \leftarrow \max\{c_r^- : (\{C_i, C_l\}, j) \in E_{\mathcal{C}}, C_l \in \mathrm{succ}(C_i)\}$

8 $\quad \delta_j^+ \leftarrow \min\left\{ \max\left\{c_l^+\right\} \cup \left\{c_r^- : (\{C_i, C_r\}, j) \in E_{\mathcal{C}}, C_r \in \mathrm{succ}(C_i) \setminus \{C_l\}\right\} : (\{C_i, C_l\}, j) \in E_{\mathcal{C}}, C_l \in \mathrm{succ}(C_i)\right\}$

9 **return** *Algorithm 5* $\left(\overline{\kappa}_R - \sum_{j:\xi_{i,j}=1} \omega_j; J; (\omega_j)_{j\in J}; (\delta_j^-)_{j\in J}; (\delta_j^+)_{j\in J}\right)$

Note that, in particular, for every root node C_r the value $c_r(\overline{\kappa})$ is the optimal value of $\infty-$Split Resolution when reduced to the connected component of C_r.

For each cluster C_i that is not a root node, we are interested in two particular values for $\overline{\kappa}_R$: First for the assumption that we will not, and second for the assumption that we will assign the unit j_i to C_i. Let us define those values as

$$c_i^- := c_i(\overline{\kappa}) = \min\left\{ \left\| \left(\overline{\kappa} - \omega(\hat{C}_l)\right)_{C_l \succeq_{G_{\mathcal{C}}} C_i} \right\|_\infty : \hat{\xi} = (\hat{C}_l)_{l\in[k]} \in R(\xi) \wedge \hat{\xi}_{i,j_i} = 0 \right\} \quad (7.6)$$

and

$$c_i^+ := c_i(\overline{\kappa} - \omega_{j_i}) = \min\left\{ \left\| \left(\overline{\kappa} - \omega(\hat{C}_l)\right)_{C_l \succeq_{G_{\mathcal{C}}} C_i} \right\|_\infty : \hat{\xi} = (\hat{C}_l)_{l\in[k]} \in R(\xi) \wedge \hat{\xi}_{i,j_i} = 1 \right\}, \quad (7.7)$$

respectively.

We will show that lines 1 to 4 of Algorithm 4 compute these values for any successor of a cluster C_i. This allows us to compute $c_i(\overline{\kappa}_R)$. First, we determine for any unit that is fractionally assigned to C_i and one of its successors, the minimum deviation that we

obtain from rounding the fractional assignments "strictly below" C_i — depending on whether we do or do not assign this unit to C_i. We denote the set of available units by J (cf. line 5). For every unit $j \in J$ let

$$S_j := \{C_t \in \mathcal{C} : \exists C_l \in \mathrm{succ}(C_i) : C_t \succeq_{G_\mathcal{C}} C_l \wedge (\{C_i, C_l\}, j) \in E_\mathcal{C}\} \tag{7.8}$$

be the set of all clusters that succeed cluster C_i via an edge labeled by the unit j. Then

$$\delta_j^- := \min\left\{ \left\| \left(\overline{\kappa} - \omega(\hat{C}_l)\right)_{C_l \in S_j} \right\|_\infty : \hat{\xi} = \left(\hat{C}_1, \ldots, \hat{C}_k\right)^\intercal \in R(\xi) \ \wedge \ \hat{\xi}_{i,j} = 1 \right\} \tag{7.9}$$

and

$$\delta_j^+ := \min\left\{ \left\| \left(\overline{\kappa} - \omega(\hat{C}_l)\right)_{C_l \in S_j} \right\|_\infty : \hat{\xi} = \left(\hat{C}_1, \ldots, \hat{C}_k\right)^\intercal \in R(\xi) \ \wedge \ \hat{\xi}_{i,j} = 0 \right\} \tag{7.10}$$

formalize the described minimum deviation values. We will show that lines 6 to 8 of Algorithm 4 correctly compute those values for every $j \in J$.

Next, we pass these values to Algorithm 5, which then computes $c_i(\overline{\kappa}_R)$ via a bisection procedure. Here, from each available unit $j \in J$ we know that we cannot do any better then $\min\left\{\delta_j^-, \delta_j^+\right\}$, as we either will or will not assign j to one of the successors of C_i. This gives the lower bound δ_L as defined in line 3. On the other hand, we will see that the error will not exceed the maximum error computed so far as well as the maximum weight of the available units. We denote this bound by δ_U (line 4). Now assume some fixed value $\delta \in \mathbb{N}$ with $\delta_L \leq \delta \leq \delta_U$ to be given for which we want to check whether $c_i(\overline{\kappa}_R) \leq \delta$. Only for units $j \in J$ such that $\delta \geq \max\left\{\delta_j^-, \delta_j^+\right\}$ we can actually decide whether or not to assign them to C_i if we want to obey the bound δ. We denote those by J_δ (cf. line 7). By the definitions of δ_j^- and δ_j^+, any unit $j \in J$ with $\delta < \delta_j^-$ must be assigned to one of the successors of C_i and any unit $j \in J$ with $\delta < \delta_j^+$ must be assigned to C_i. We thus obtain a remaining weight demand κ_δ by subtracting the weights of the latter units from the provided demand weight $\hat{\kappa}$ (cf. line 9). If J_δ is not empty, we solve the resulting Weak Subset Sum instance, otherwise, we just have to check whether $\kappa_\delta \leq \delta$. We may then update the lower or upper bounds accordingly.

Algorithm Analysis

First, let us prove the correctness and running time of Algorithm 5. Here, similar to [Sch01] we particularly exploit that the occurring instances of Weak Subset Sum are sufficiently small in order to be solvable in polynomial time. As a byproduct, we obtain a valid upper bound for $\infty-$Split Resolution. Schröder [Sch01] obtains the same bound in a similar manner ([Sch01, Satz 7.13]), while another independent (and simpler) constructive proof can be found in [BGK17, Theorem 5].

Algorithm 5: Optimal subgraph rounding

Input: $\hat{\kappa} \in \mathbb{Q}$, $J \subset [n]$, $\omega_j \in \mathbb{N}_{>0}$, $\delta_j^+, \delta_j^- \in \mathbb{N}$ for $j \in J$

if $J = \emptyset$ **then**
  return $|\hat{\kappa}|$
$\delta_L \leftarrow \max\left\{\min\left\{\delta_j^+, \delta_j^-\right\} : j \in J\right\}$
$\delta_U \leftarrow \max \bigcup_{j \in J}\left\{\omega_j, \delta_j^+, \delta_j^-\right\}$
repeat
  $\delta \leftarrow \lfloor \frac{1}{2}(\delta_L + \delta_U) \rfloor$
  $J_\delta \leftarrow \left\{j \in J : \delta \geq \max\left\{\delta_j^+, \delta_j^-\right\}\right\}$
  $J_\delta^+ \leftarrow \left\{j \in J : \delta < \delta_j^+\right\}$
  $\kappa_\delta \leftarrow \hat{\kappa} - \sum_{j \in J_\delta^+} \omega_j$
  if $J_\delta \neq \emptyset$ **then**
    if Weak Subset Sum $(\delta; \kappa_\delta; (\omega_j)_{j \in J_\delta})$ *is a* yes-*instance* **then**
      $\delta_U \leftarrow \delta$
    else
      $\delta_L \leftarrow \delta$
  else
    if $|\kappa_\delta| \leq \delta$ **then**
      $\delta_U \leftarrow \delta$
    else
      $\delta_L \leftarrow \delta$
until $\delta_L \geq \delta_U - 1$
return δ_U

Lemma 7.3
Assume a valid input of Algorithm 3 to be given.

Let $i \in [k]$, $\overline{\kappa}_R \in \begin{cases} \{\overline{\kappa}, \overline{\kappa} - \omega_{j_i}\} \text{ if } j_i \geq 0 \\ \{\overline{\kappa}\} \qquad\quad \text{, otherwise} \end{cases}$ *(with* j_i *as defined by Eq.* (7.4)*),* $J = \{j \in [n] : (\{i, l\}, j) \in E, C_l \in \mathrm{succ}(C_i)\}$*, and let* δ_j^+, δ_j^- *for* $j \in J$ *as defined by Eqs.* (7.9) *and* (7.10)*.*

Then Algorithm $5\left(\overline{\kappa}_R - \sum_{j:\xi_{i,j}=1} \omega_j; J; (\omega_j)_{j \in J}, (\delta_j^-)_{j \in J}; (\delta_j^+)_{j \in J}\right)$ *correctly computes* $c_i(\overline{\kappa}_R)$ *as defined by Eq.* (7.5) *in* $\mathcal{O}\left(k^2 \log(\max_{j \in [n]} \omega_j)\right)$*.*

In particular, $c_i(\overline{\kappa}_R) \leq \max_{j \in [n]} \omega_j$*.*

Proof. First, assume that C_i is maximal w. r. t. "$\succ_{G_\mathcal{C}}$" (i. e., C_i is a leaf in $G_\mathcal{C}(\xi)$).

Then it follows that a single unit, namely j_i, has been fractionally assigned to C_i by ξ and thus $J = \emptyset$. Thus, the resulting error is either $|\overline{\kappa} - \sum_{j:\xi_{i,j}=1} \omega_j|$ or $|\overline{\kappa} - \sum_{j:\xi_{i,j}=1} \omega_j - \omega_{j_i}|$, depending on whether the unit j_i is or is not assigned to C_i, respectively. This is correctly returned in line 2. Furthermore, $\overline{\kappa} = \sum_{j:\xi_{i,j}=1} \omega_j + \xi_{i,j_i}\omega_{i_j}$ holds, so it follows that $c_i(\overline{\kappa}_R) \leq \omega_{j_i}$.

For the case $J \neq \emptyset$, we can inductively assume that $\max\left\{\delta_j^-, \delta_j^+\right\} \leq \max_{j\in[n]} \omega_j$. As argued before, it holds that $c_i(\overline{\kappa}_R) \geq \min\left\{\delta_j^-, \delta_j^+\right\}$ for all $j \in J$ by definition. Next, we conclude $c_i(\overline{\kappa}_R) \leq \max \bigcup_{j\in J}\left\{\omega_j, \delta_j^+, \delta_j^-\right\} \leq \max_{j\in[n]} \omega_j$ (the latter by induction). This follows as we can set $J^* := \left\{j^* \in J : \sum_{j\in[n]:\xi_{i,j}=1} \omega_j + \sum_{j\in J: j\leq j^*} \omega_j \leq \overline{\kappa}\right\}$ and obtain a rounded clustering by assigning all units in J^* to C_i and all units in $J \setminus J^*$ to one of the successors of C_i to which the respective unit has been fractionally assigned. For all clusters C_l with $C_l \succ_{G_\mathcal{C}} C_i$, this yields an error of at most $\max \bigcup_{j\in J}\left\{\delta_j^+, \delta_j^-\right\}$ by definition of δ_j^+ and δ_j^-, as well as an error of at most $\max_{j\in[n]} \omega_j$ for cluster C_i by definition of J^*. Consequently, we have that $c_i(\overline{\kappa}_R) \leq \delta_U$ as defined in line 4, so this is a valid upper bound.

We already argued that only for units in J_δ as defined in line 7 we can actually decide whether or not to assign them to C_i, as well as that all units in J_δ^+ as given in line 8 must be assigned to C_i. Thus, lines 10 to 19 correctly detect whether $c_i(\overline{\kappa}_R) \geq \delta$ or $c_i(\overline{\kappa}_R) \leq \delta$ and update the bounds accordingly. Hence, this bisection procedure correctly computes $c_i(\overline{\kappa}_R)$.

It remains to show the running time. As $\delta_U - \delta_L \leq \max_{j\in[n]} \omega_j$ holds, the bisection procedure in form of the **repeat** block requires $\mathcal{O}\left(\log(\max_{j\in[n]} \omega_j)\right)$ iterations. Lines 6 to 9 can be performed in $\mathcal{O}(k)$ as $|J_\delta| \leq k$ (as there are at most $k-1$ fractionally assigned units).

For the decision on Weak Subset Sum we can assume that $0 < \kappa_\delta \leq \sum_{j\in J_\delta} \omega_j$ holds. This can be checked in $\mathcal{O}(k)$ (as otherwise Weak Subset Sum can be trivially decided). Now let $j \in J_\delta$. In particular, this implies $\delta \geq \delta_j^-$. Furthermore, let $C_l \in \mathrm{succ}(C_i)$ such that j is fractionally assigned to C_l. It follows that $c_l(\overline{\kappa} - \omega_j) \leq \delta$ and $c_l(\overline{\kappa}) \leq \delta$. This means that there are two rounded assignments for the clusters C_t with $C_t \succeq_{G_\mathcal{C}} C_l$ — one excluding and one including unit j — that both result in a maximal error of at most δ. In the latter case, an additional weight of ω_j has to be distributed among the clusters, each of which can additionally take a weight of at most 2δ in order to stay feasible. Thus, by the pigeonhole principle it follows that

$$\omega_j \leq |\{C_t : C_t \succeq_{G_\mathcal{C}} C_l\}| \cdot 2\delta.$$

Thus, we can conclude

$$\frac{\kappa_\delta}{\delta} \leq \frac{\sum_{j\in J_\delta} \omega_j}{\delta} \leq \frac{\sum_{C_l\in\mathrm{succ}(C_i)} |\{C_t : C_t \succeq_{G_\mathcal{C}} C_l\}| \cdot 2\delta}{\delta} \leq 2k.$$

Together with $|J_\delta| \leq k$, Theorem 7.2 then yields that we can solve the resulting instance of Weak Subset Sum in $\mathcal{O}(k^2)$.

Altogether, this gives the claimed running time in $\mathcal{O}\left(k^2 \log(\max_{j\in[n]} \omega_j)\right)$. □

Theorem 7.4
Algorithm 3 correctly computes the optimal value of $\infty-$Split Resolution *in* $\mathcal{O}\left(nk + k^3 \log(\max_{j\in[n]} \omega_j)\right)$.

Proof. In order to show correctness, it suffices to show that Algorithm 4 correctly returns $c_i(\overline{\kappa}_R)$. Then Algorithm 3 correctly computes the maximal error $c_i(\overline{\kappa})$ for each connected component of $G_{\mathcal{C}}(\xi)$ independently and returns the maximum. Correctness of Algorithm 4 follows almost immediately from Lemma 7.3. If the input cluster C_i in Algorithm 4 is a leaf, $J = \emptyset$ follows and hence the call of Algorithm 5 correctly returns $c_i(\overline{\kappa}_R)$. If C_i is not a leaf, we may inductively assume that the recursive call of Algorithm 4 correctly returns c_l^+ and c_l^- as defined by Eqs. (7.6) and (7.7) for any $C_l \in \mathrm{succ}(C_i)$ in lines 3 and 4.

Thus, we only have to check that the values δ_j^- and δ_j^+ for $j \in J$ as computed in lines 7 and 8 are consistent with Eqs. (7.9) and (7.10). By definition, δ_j^- is to yield the weight deviation of clusters $C_l \succ_{G_{\mathcal{C}}} C_i$ in case j is assigned to C_i. For any successor C_l, the value c_l^- yields the weight deviation among clusters $C_t \succ_{G_{\mathcal{C}}} C_l$ if j is not assigned to C_l, so the maximum among those values yields δ_j^-. Similarly, δ_j^+ is defined to be the weight deviation if j is assigned to *some* successor C_l. For each C_l, the resulting value is the maximum of c_l^+ and the values c_t^- for all other successors $C_t \in \mathrm{succ}(C_i) \setminus \{C_l\}$. Taking the minimum over all clusters C_l then yields δ_j^+, as is done in line 8.

Concerning the running time, we can clearly compute and root $G_{\mathcal{C}}(\xi)$ in the given bounds. Now for each designated root node of a connected component of $G_{\mathcal{C}}(\xi)$ Algorithm 4 is called once, and for each other node exactly twice in lines 3 and 4 of Algorithm 4. Hence, there are $\mathcal{O}(k)$ many calls to Algorithm 4. With respect to lines 5 to 8, note that the sets J are disjoint for each C_i in the input of Algorithm 4 and in union yield the set of all fractionally assigned units, of which there are at most $k-1$ many. Every iteration of the **for**-loop in lines 6 to 8 can be done in $\mathcal{O}(k^2)$, so overall lines 5 to 8 run in $\mathcal{O}(k^3)$. Finally, the sum in the input for Algorithm 5 in line 9 of Algorithm 4 will take $\mathcal{O}(n)$ flops. Lemma 7.3 yields a running time of $\mathcal{O}\left(k^2 \log(\max_{j\in[n]} \omega_j)\right)$ for Algorithm 5. Hence, in total the calls to Algorithm 4 can be done in $\mathcal{O}\left(nk + k^3 \log(\max_{j\in[n]} \omega_j)\right)$. □

We may now incorporate these results on $\infty-$Split Resolution into our schema for the EDD and obtain the following result.

Corollary 7.5
Let an instance of EDD *and functions* $f_i : X \to \mathbb{R}$ *for* $i \in [k]$ *be given.*

Furthermore, assume all input values to be rational, i. e., $X = \left\{x^{(1)}, \dots x^{(n)}\right\} \subset \mathbb{Q}^2$ *and* $f_i(X) \subset \mathbb{Q}$ *for all* $i \in [k]$*. Then we can in polynomial time compute an integer clustering* $\xi = (C_1, \dots, C_k)^\intercal \in T_{uncstr.,k,n} \cap \{0,1\}^{k \times n}$ *such that* $|\omega(C_i) - \overline{\kappa}| \leq \max_{j \in [n]} \omega_j$ *as well as* $M = (\mu_i)_{i \in [k]} \in \mathbb{R}^k$*, such that the generalized Voronoi diagram w. r. t.* $(f_i + \mu_i)_{i \in [k]}$ *is feasible for* ξ*.*

Proof. Following the schema of Fig. 7.3 this can be deduced by combining Proposition 4.5 with Theorem 7.4 and Lemma 7.3. □

Obviously, Corollary 7.5 is too weak in order to provide a solution that satisfies the balancing constraint (7.1c) for any relative error $\delta < \frac{\max_{j \in [n]} \omega_j}{\overline{\kappa}}$. Yet we will see that for our real-world instances and used approaches, the schematic of Fig. 7.3 will result in very pleasant results w. r. t. the resulting balancing errors (cf. Section 7.4).

7.3 Approaches

This section outlines how we apply the general methodology from Section 7.2 for particular choices of generalized Voronoi diagrams. As pointed out in Section 7.1.1, the search for a suitable methodology that finds an optimal electoral district design is particularly tricky as it is preceded by the issue of how to actually define the notion of optimality of a district design. Hence, we apply different exemplary classes in order to demonstrate how the choice of a diagram type can be utilized in order to take the individual needs and optimization criteria into consideration. In particular, this choice determines the geometric characteristics of the resulting diagrams.

We will consider power diagrams in Section 7.3.1, shortest-path diagrams in Section 7.3.2 and anisotropic power diagrams in Section 7.3.3.

7.3.1 Power Diagrams

The first class we want to consider are power diagrams as introduced in Section 3.3.2. As described in Section 7.1.2, power diagrams already occurred in some attempts to the EDD (in particular in [FH11], [Sch01]), as well as in related applications. In particular, we will apply the theory as recalled in Chapter 6.

We also note that in some existing approaches for electoral district design, power diagrams even occurred without being explicitly recognized. Basically, all approaches that eventually solve a transportation problem in the form of the program (4.18) with

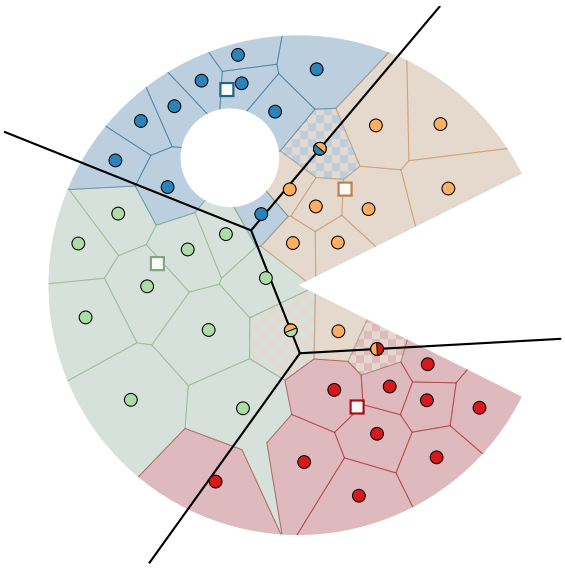

Figure 7.7: Clustering of our toy example that is supported by a power diagram.

objective

$$\sum_{i\in[k]}\sum_{j\in[n]} \xi_{i,j}\omega_j \left\|x^{(j)} - s^{(i)}\right\|_2^2 \tag{7.11}$$

for sites $s^{(i)} \in \mathbb{R}^2, i \in [k]$, yield a feasible power diagram for every optimal solution (as, for example, done in [Hes+65] or [Kai66]). As the moment of inertia is a respected measure for the consolidation degree of districts, this is not a coincidence, of course.

In Chapter 6 the problem of minimizing the moment of inertia was introduced as the Balanced k-Means problem. We have seen a one-to-one relation between faces of the gravity polytope and clusterings that allow a supporting power diagram and characterized the local optima of the Balanced k-Means problem by extremal clusterings supported by centroidal power diagrams.

Suitability of Power Diagrams for the EDD

If one accepts the moment of inertia as a measure for the compactness of a district plan, power diagrams yield the most compact districts due to the correspondence depicted in Chapter 6.

With respect to the weight-balancing constraints (7.1c), we can reasonably assume the solutions of the underlying linear programs to be unique following Theorem 7.1. Hence, we know from Section 7.2.2 that we can efficiently solve the resulting rounding problem. However, we obtain no better a priori guarantee for the resulting relative

error δ than $\max_{j\in[n]} \frac{\omega_j}{\kappa}$ from Lemma 7.3. For our considered data set, this bound does not provide a sufficient guarantee a priori, as large municipalities do occur. However, we will also see that this problem can be treated very well.

Next, let us discuss the contiguity constraints (7.1d). As we know from Section 3.3.2, power diagrams are the only generalized Voronoi diagrams with linear bisectors and thus the only diagrams that yield convex cells. However, here we are facing two issues.

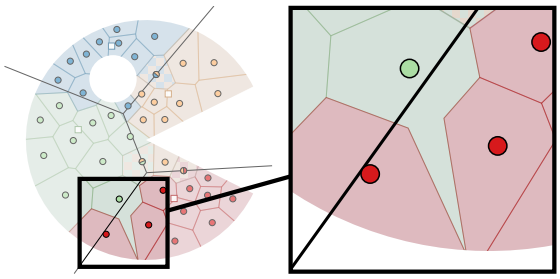

Figure 7.8: Non-contiguous clusters as a consequence of the point-representation of units.

First of all, we have assumed for our data that each municipality is represented by a point in the plane, e. g., determined by its territory's center of gravity. On the other hand, contiguity is determined by the contiguity graph which, for example, may be derived by testing the municipalities' areas for common borders. Thus, even if a feasible power diagram guarantees the convex hull of the representatives of a district to be contained in its corresponding cell, this may not imply contiguity. Figure 7.8 highlights this effect in our toy example. Here, the green unit separates the two red ones. However, as we will also see for our real-life data, this may be regarded as an artifact due to rather oddly-shaped units.

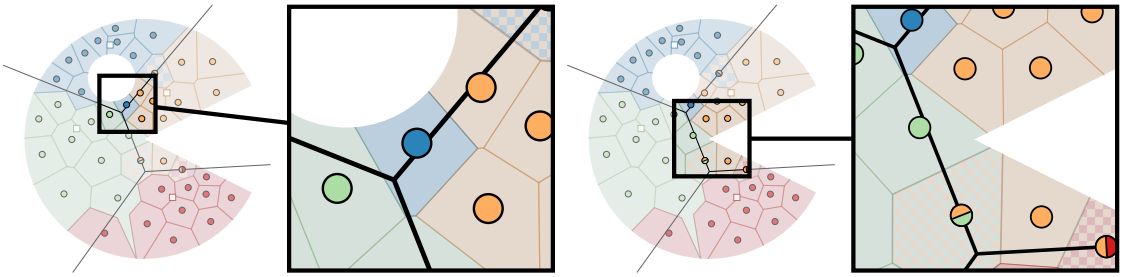

(a) Blue unit isolated due to a hole.

(b) An indentation of the area's boundary splits the orange cluster.

Figure 7.9: Examples for the non-convexity of the units' underlying region leading to non-contiguous clusters.

Second, the non-convexity of the underlying state may be the cause for non-contiguous clusters, too. Figure 7.9 depicts such effects in our toy example. As a power diagram

is blind for both holes (as in Fig. 7.9a) as well as indentations in the state's border (Fig. 7.9b), those can lead to non-contiguity. We will see that those kind of artifacts do appear in our real-world application, although rather rarely. Still, in order to enforce contiguous districts, a manual post-processing is necessary.

Handling of Non-Contiguous Districts

Power diagrams rely on Euclidean distances while the effects described above root from information that is intrinsic to the shapes of the municipalities' and the state's areas. In order to incorporate this information, we might add further constraints to the underlying linear program (4.18). Of course, this will have a malicious impact on our desired primal-dual relation of optimal clusterings and power diagrams, which we shall briefly discuss next.

First, we observe that representing each municipality by its center of gravity is in fact the natural choice w. r. t. power diagrams. A canonical idea how to incorporate the shape information of municipality areas into the model is to represent each municipality not by one but multiple points. Those are then linked via *must-link* constraints. Let $x^{(j)}$ be some unit with weight ω_j and let us further assume that $A_j \subset \mathbb{R}^2$ is its associated territory. We may further assume ω to be a measure on $\mathbb{R}^2$ (i. e., the population distribution) such that $\omega_j = \omega(A_j)$ holds for $j \in [n]$. In order to pay regard to the shape of the municipality, we may partition the municipality area $A_j = \bigcup_{r\in[t]} A_{j_r}$ such that $\omega(A_j) = \sum_{r\in[t]} \omega(\hat{A}_{j_r})$. We may then replace each unit $x^{(j)}$ in our model by units $\hat{x}^{(j_1)}, \dots, \hat{x}^{(j_t)} \in \mathbb{R}^2$ that represent the respective splits and have weights $\hat{\omega}_{j_r} := \omega(A_{j_r})$ for $r \in [t]$. Let $\hat{\xi}$ denote the assignment variables of the program (4.18) after this adaption. As we (eventually) must assign the municipalities to districts as a whole, this suggests to add the must-link constraints

$$\hat{\xi}_{i,j_a} = \hat{\xi}_{i,j_b} \ \forall i \in [k], a, b \in [t]. \tag{7.12}$$

Now recall a basic fact from the theory of linear programming: Whenever columns in a primal problem are set equal, this causes the corresponding constraints in the dual to be replaced by their sum[1]. This means that for a fixed $i \in [k]$, the dual constraints (4.27a) corresponding to the cluster index i and the unit indices $j_1, \dots, j_t$ are replaced

[1]Consider a linear program $\min_{x,y} c^\intercal x + d^\intercal y$ s.t. $Ax + By = b, x, y \geq 0$ with corresponding dual $\max_u b^\intercal u$ s.t. $A^\intercal u \leq c, B^\intercal u \leq d$ (with parameters A, B, c, d, b from appropriate spaces). Adding the constraint $x - y = 0$ to the primal alters the dual to $\max_{u,v} b^\intercal u$ s.t. $A^\intercal u + v \leq c, B^\intercal u - v \leq d$ which can be simplified to $\max_u b^\intercal u$ s.t. $A^\intercal u + B^\intercal u \leq c + d$ (as we may set $v = c - A^\intercal u$).

by

$$\sum_{r\in[t]} \hat{\eta}_{j_r} - \underbrace{\sum_{r\in[t]} \hat{\omega}_{j_r}}_{=\omega_j} \mu_i \leq \underbrace{\sum_{r\in[t]} \hat{\omega}_{j_r} f_i(\hat{x}^{(j)})}_{=\omega_j \sum_{r\in[t]} \frac{\hat{\omega}_{j_r}}{\omega_j} f_i(\hat{x}^{(j)})} . \tag{7.13}$$

Now let us consider the case of power diagrams in affine representation, i. e., $f_i(x) = a_i{}^\intercal x + \alpha_i$ for $i \in [k]$ and suitable parameters (cf. Section 3.3.2). Furthermore, we can substitute $\eta_j := \sum_{r\in[t]} \hat{\eta}_{j_r}$ in the dual program (4.27). Then Eq. (7.13) becomes

$$\eta_j - \omega_j \mu_i \leq \omega_j \left(a_i{}^\intercal \left(\sum_{r\in[t]} \frac{\hat{\omega}_{j_r}}{\omega_j} \hat{x}^{(j_r)} \right) + \alpha_i \right). \tag{7.14}$$

This is exactly the dual constraint we obtain when choosing $x^{(j)} = \sum_{r\in[t]} \frac{\hat{\omega}_{j_r}}{\omega_j} \hat{x}^{(j_r)}$ as the representative for A_j right away. Thus, we conclude that the usage of must-link constraints in combination with power diagrams is equivalent to replacing the linked points by their convex combination according to their weights. In particular, if we consider arbitrarily small partitions of A_j and assume $\hat{x}^{(j_r)} \in A_{j_r}$ for $r \in [t]$, this yields $x^{(j)} = \frac{1}{\omega_j} \int_{A_j} x\omega(dx)$, i. e., the municipality's centroid, in the limit.

The other canonical way of further constraining the clustering is to exclude or to enforce the assignment of units to clusters. Assume we exclude the assignment of the jth unit to the ith cluster, i. e., we introduce the constraint

$$\xi_{i,j} = 0. \tag{7.15}$$

In other words, we drop the associated column in the program (4.26) and hence the associated constraint (4.27a) in the program (4.27). Assume the clustering ξ^* and the diagram $\mathcal{P} = (P_1, \dots, P_k)$ w. r. t. functions $(f_l)_{l\in[k]}$ to be obtained from a primal-dual pair of (4.26) and (4.27) as proposed in (the proof of) Corollary 4.10. Also, assume that the jth unit is assigned to the rth cluster by some positive fraction, i. e., $\xi^*_{r,j} > 0$. It is easy to verify that the remaining complementarity conditions now only imply $x^{(j)} \in P_i \cup P_r$. More precisely, if we drop the ith cell and consider the diagram $\tilde{\mathcal{P}} = \left(\tilde{P}_l\right)_{l\in[k]\setminus\{i\}}$ w. r. t. functions $(f_l)_{l\in[k]\setminus\{i\}}$, then complementarity still guarantees $x^{(j)} \in \tilde{P}_r \subset P_r \cup P_i$. This means that we might lose feasibility of the diagram, however, only w. r. t. the cluster from which the unit has been excluded.

Note that this argument is repeatable. So if we introduce constraints (7.15) for some $j \in [n]$ and all $i \in \mathcal{I}$ for some $\mathcal{I} \subset [k]$, the desired relation is weakened further to $x^{(j)} \in \bigcup_{i\in\mathcal{I}} P_i \cup P_r$.

Alternatively, we may enforce the assignment of the jth unit to the ith cluster by setting

$$\xi_{i,j} = 1. \tag{7.16}$$

Of course, due to constraint (4.26a), this is equivalent to the introduction of constraints of type (7.15) for all cluster indices $r \in [k] \setminus \{i\}$. So, according to the previous observations, we simply lose any guaranteed relation between the resulting diagram and the unit $x^{(j)}$. This, of course, becomes also clear as instead of constraint (7.16) we can just solve the program (4.26) for the reduced unit set $X \setminus \left\{x^{(j)}\right\}$ and with a target weight $\kappa_i := \overline{\kappa} - \omega_j$ for the ith cluster.

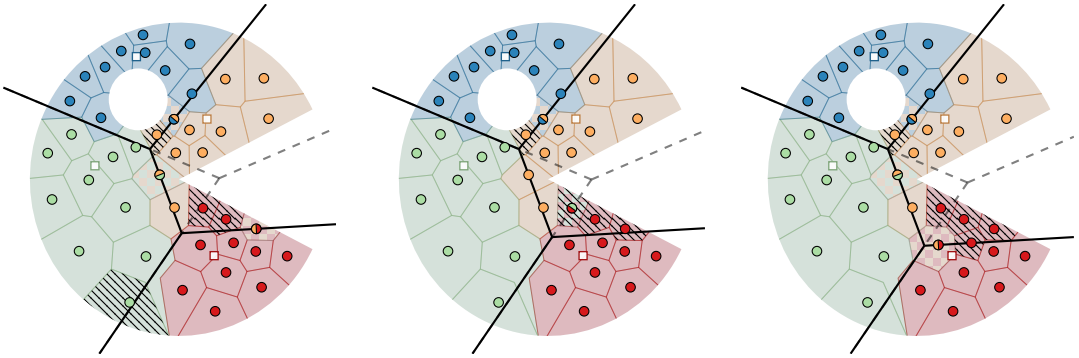

(a) Marked units are excluded from the clusters that they were assigned to in Fig. 7.7. The gray dashed lines indicate the bisectors that result from dropping the orange cluster.

(b) Result after adding one further constraints in comparison to Fig. 7.10a. Also, the constraint for the unit that is now correctly contained in the green cluster can be dropped.

(c) Result after adding two further constraints in comparison to Fig. 7.10b. The unit that was fractionally assigned to both the green and the red cluster is now further excluded from the former, as well as another red unit from the orange cluster.

Figure 7.10: Power diagram clustering of the toy example with additional exclusion constraints of type (7.15). Units that are affected by such constraints are marked by a black hatching.

Despite the conclusion that adding constraints of type (7.15) or (7.16) practically undermines the desired relation between diagrams and clusterings, this seems still practically feasible as long as only a few boundary units of clusters are concerned. In particular, the constraints of type (7.15) can be considered "less harmful" to our theory. Hence, we will apply a manual post-processing in which we successively add constraints for boundary units that are disconnected from their assigned cluster. Figure 7.10 depicts this post-processing for our toy example. Here, multiple iterations may be necessary in order to establish contiguity (up to the issue of fractionally assigned units). However, we will see that for our real-word data the necessary efforts for this kind of post-processing are within very reasonable bounds.

Resulting Method

Let us now fit the theory of Chapter 6 together with the preceding thoughts into our general methodology of Section 7.2.1 as depicted in Fig. 7.3.

For a fixed number N of iterations we sample sites from a uniform distribution over $\mathbb{S}^{dk-1} \cap \left\{c = (c^{(i)})_{i\in[k]} \in \mathbb{R}^{dk} : \sum_{i\in[k]} c^{(i)} = \sum_{j\in[n]} \frac{\omega_j}{\omega(X)} x^{(j)}\right\}$. Our experiments on our given data suggest that reasonable numbers (such as $N = 1000$) can be assumed to suffice (cf. Section 7.4.3).

For each of those sites, we then perform the (revised) Balanced k-Means-algorithm (Algorithm 2) as proposed in Chapter 6. Let a clustering $\xi \in T_{\kappa,\omega}$ and the sites $s^{(i)} = c(C_i)$ for $i \in [k]$ be the output of such a run. The additive weights $(\mu_i)_{i\in[k]}$ (obtained from the dual solution corresponding to constraints (4.18b), cf. Proposition 4.5) then yield the power diagram $\mathrm{PD}\left(\left(s^{(i)}, \mu_i\right)_{i\in[k]}\right)$ which supports the clustering ξ. Recall that from Theorem 6.10 we know that ξ is a vertex of $T_{\kappa,\omega}$. Thus, we may apply the split resolution algorithm (Algorithm 3) from Section 7.2.2. If a better solution w. r. t. the Balanced k-Means objective has been found, we update the stored best known solution.

At this point, we may furthermore reject solutions that disqualify for other reasons. In particular, we may reject solutions that yield too large balancing errors after an integral solution has been obtained. We will see in Section 7.4.3 to which degree this impacts the achieved objective value.

For the best found solution, we apply the post-processing in order to establish full contiguity w. r. t. the contiguity graph G as described before. Thus, we successively introduce constraints of types (7.15) and (7.16) to the problem (4.18) (using the sites of the best known solution). Note that this has to be done in a manual manner.

7.3.2 Shortest-Path Diagrams

Contiguity of districts is a natural constraint of electoral district design. However, we have already described situations in which a district's contiguity may be ambiguous even when the geometry of the unit areas is fully taken into consideration. Moreover, we have also seen for power diagrams that a diagram type, that relies on the representation of units by points in the Euclidean plane, may fail to produce contiguous districts. Clearly, unless the information that is encoded in the contiguity graph G is incorporated into the diagram-defining functions, contiguity cannot be guaranteed a priori in general.

For many instances, the necessity of a manageable manual post-processing as we have seen for the example of power diagrams might be well-acceptable. However, in particular in situations when transparency and independence of the district design process is of greatest importance, a human-controlled post-processing can be unacceptable. In particular, in a situation as in the United States, where there are cases in which

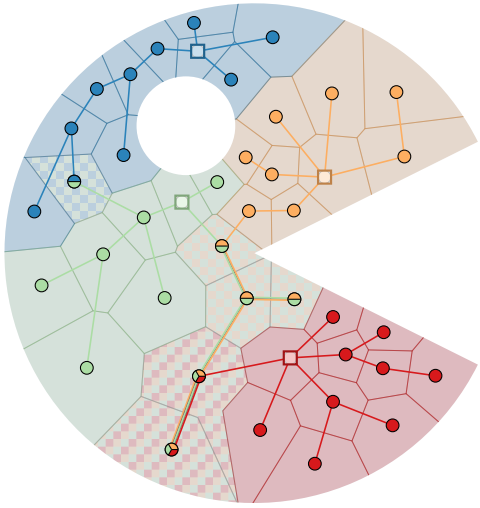

Figure 7.11: Exemplary shortest path diagram for our toy example.

the district design process is allegedly driven by political intentions, the desire for a fully-automated process arises (cf. [PP91], [Alt97]).

The canonical way to include the inherent information of the contiguity graph G into our context of generalized Voronoi diagrams is to use diagram-defining functions that are deduced from G. Consequently, we apply shortest-path diagrams as introduced in Section 3.3.4 over the contiguity graph G. We know that a shortest-path diagram is site-star-shaped. In particular, we have seen that the cells are in fact subtrees of the shortest-path trees w. r. t. the sites. Consequently, any clustering that is supported by a shortest-path diagram will yield star-shaped and hence contiguous clusters.

Related to our approach, shortest-path distances have also been considered by Zoltners and Sinha [ZS83] for the application of sales territory alignment. They propose a Lagrangian dual approach. From our point of view, their model optimizes shortest-path distances over a single-constrained clustering polytope while including binary constraints. This may be identified as an integer version of the program (4.18). They then relax constraints (4.18b) and perform a sub-gradient procedure to optimize the Lagrange multipliers (i. e., the additive terms μ_i in our context). Here, they observe that in each iteration units are assigned to the clusters with smallest adjusted objective, which, of course, is the clustering induced by the generalized Voronoi diagram resulting from the current Lagrangian multipliers. In particular, they conclude contiguity of the resulting clusters. However, they do not discuss the case of those assignment being ambiguous (i. e., the case of non-empty cell-intersections). Schröder [Sch01] also recalls the approach of [ZS83] in his work and in particular points out this issue (cf. [Sch01,

Chapters 7-8]). We will see shortly that for shortest-path diagrams the non-uniqueness of optimal solutions is indeed a crucial issue. Here, our LP-duality based relation of diagrams and clusterings will help to get a more precise understanding.

Choice of Structural Parameters - Sites

In the following, we will walk through the different steps of applying the schematic of Fig. 7.3 to shortest-path diagrams and summarize the resulting procedure afterwards.

The only structural parameters to choose are the sites. However, we may note at this point that — depending on the considered application — further parametrizations are possible. For example, one might further consider a tweak of edge weights in order to control the resulting shortest-path trees. An artificial increase of edge weights could be utilized in order to pay regard to further barriers such as borders of administrative levels that should be regarded in the optimization process.

In contrast to power diagrams, the choice of sites is limited to the unit set X. In the unconstrained case, this is also referred to as *k-medoid* clustering (cf. [KR05, Chapter 2]). A classical algorithm to approach this is the "partitioning around medoids" — short *PAM* — algorithm. In its originally proposed form in [KR87], the algorithm works by selecting k random sites (or medoids) and assigning all remaining points to their respective closest site. Then, all single swaps of sites and non-sites are tested for improving the overall costs. As long as improving swaps can be found, the best one is performed and the procedure iterates. As we have seen for the Balanced k-Means-algorithm in comparison with the (unconstrained) k-Means algorithm, this concept canonically carries over to our constrained setting. Several variants and improvements on the PAM algorithm have been proposed (see, for example, [PJ09; ZC05]). However, as sufficient for our purposes, we will use a very simple adaption of it that solves program (4.18) as a subroutine and is conform with our schema of Fig. 7.3. As we will see that rounding errors become more problematic with shortest-path diagrams, we will search for swaps that improve on the resulting balancing error instead of the resulting assignment costs given by the objective of (4.18).

Significance of the Supporting Property

We have seen in Section 3.3.4 that bisectors of shortest-path diagrams consist of subtrees of the intersections of the shortest-path trees associated with the cells. More precisely, whenever a unit lies in the intersection of two cells, all successors of the unit that are common to both associated shortest-path trees are contained in the corresponding bisector, too.

It is important to notice that the feasibility of a shortest-path diagram is insufficient in order to guarantee the contiguity of the resulting clusters. Assume that a unit that is contained in the intersection of two cells is integrally assigned to one of those. The

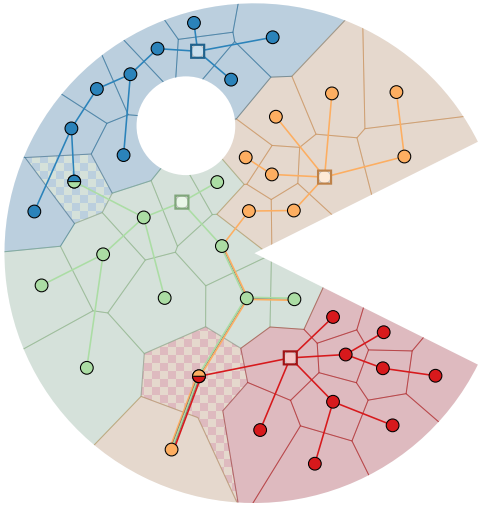

Figure 7.12: Extremal fractional clustering in our toy example for the case of shortest-path diagrams.

common successors of that unit along the associated shortest-path trees might then also be feasibly assigned to either of the two clusters without being necessarily connected to their assigned sites.

As an example, Fig. 7.12 depicts a clustering for our toy example that is an extremal solution of the program (4.18) using shortest-path distances. While the diagram (depicted in form of the trees T_i as introduced in Section 3.3.4) is indeed feasible, it is not supporting and the clustering is non-contiguous.

Recall that for the case of power diagrams, we know that the Balanced k-Means-algorithm will terminate in an extremal solution that is the unique optimum w. r. t. the final sites. Also, by Lemma 3.6, we know that for power diagrams in general we can reasonably assume the optimum to be unique and thus supported by the resulting diagrams. However, this does not hold true for shortest-path diagrams. Whenever shortest-path trees overlap in the described way, the additive correction terms that root from the balancing constraints can force the whole overlapping subtree to be contained in the intersection of the corresponding shortest-path diagram cells (cf., for example, the orange and green colored trees in Fig. 7.12). Any fractional assignment of points in such an overlap that results in the same overall weight apportionment as any optimal solution is hence optimal, too. Figure 7.11 depicts a solution in the relative interior of the optimal face of the same instance of (4.18) together with the same diagram as in Fig. 7.12. Here, however, the diagram is supporting.

So, two issues arise as depicted in the following.

Precedence Constraints

First, we should ensure that our underlying algorithm that solves the program (4.18) selects a solution that is actually lying in the relative interior of the optimal face in order to obtain a supporting diagram that guarantees contiguity. However, for our purposes, we can even slightly relax this. As our main intention is the contiguity of clusters, we can explicitly incorporate the star-shaped property of clusters as additional constraints in the program as follows:

Assume a clustering to be given that is supported by a shortest-path diagram. Recall the definition of the shortest-path trees T_i in Section 3.3.4. Now for every $i \in [k]$, we know that the subgraph of G induced by the ith cell, i. e., $G[P_i]$, is a subtree of T_i. Thus, for any unit assigned to the cluster C_i that is not a site, its predecessor in the rooted tree T_i must be assigned to this cluster, too. Let

$$\mathrm{pre}_{T_i}(j) := r \in [n] \text{ such that } (x^{(r)}, x^{(j)}) \in T_i$$

denote the predecessor (index) of $x^{(j)}$ in the rooted tree T_i for every $j \in [n]$ with $x^{(j)} \neq s^{(i)}$. We can now introduce precedence constraints

$$\xi_{i,\mathrm{pre}_{T_i}(j)} \geq \xi_{i,j} \quad \forall j \in [n] : x^{(j)} \neq s^{(i)} \ \forall i \in [k]. \tag{7.17}$$

Let us formally state the feasibility of those constraints in our context. Here, in order to apply Lemma 3.21, we will assume that shortest paths in the contiguity graph G are unique. For the real-word instances we have in mind, the edge weights are typically determined by some sort of distance measure such as Euclidean distances or traveling times. As those can be naturally assumed to imply unique shortest paths, this is a fairly mild assumption.

Theorem 7.6
Consider the program (4.18) *w. r. t. functions* $(f_i)_{i\in[k]}$ *as given by Eq.* (3.45)*. Furthermore, assume shortest paths between nodes in the contiguity graph* G *to be unique.*

Then there exists a solution in the relative interior of the optimal face that fulfills the constraints (7.17)*.*

Proof. We will show this by induction over the number of units n.

For $n = 1$ there is nothing to show.

Now for any $n \geq 2$, let ξ^* be any optimal solution in the relative interior of the optimal face of (4.18). Let $\mathcal{P} = (P_1, \ldots, P_k)$ be a corresponding supporting shortest-path diagram w. r. t. functions $(f_i)_{i\in[k]}$.

For any node $x^{(j)}$, we denote by

$$I(x^{(j)}) := \left\{ i \in [k] : x^{(j)} \in P_i \right\}$$

the set of all cluster indices such that $x^{(j)}$ lies in the respective cells. Note that the supporting property states $\xi^*_{i,j} > 0$ if and only if $i \in I(x^{(j)})$.

Now we select a unit $x^{(j)} \in X$ such that $x^{(j)}$ is a leaf of $T_i[P_i]$ for every $i \in I(x^{(j)})$. In fact, any leaf of $T_i[P_i]$ for any $i \in [k]$ has this property due to Lemma 3.21, hence this is well-defined. W.l.o.g. let $x^{(n)}$ be such a unit.

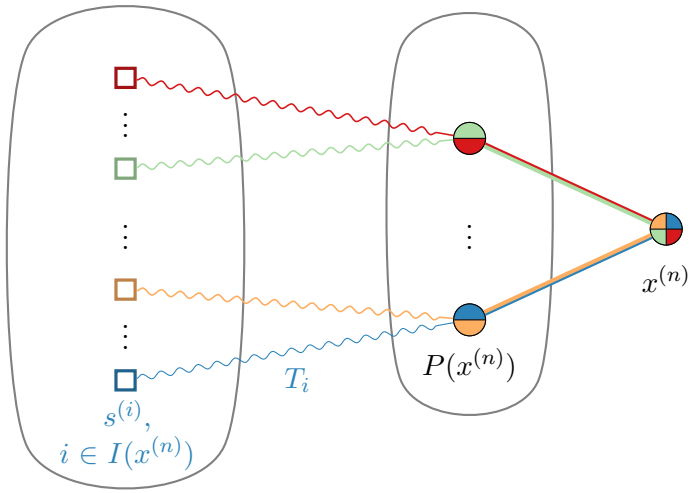

Figure 7.13: Exemplary illustration of the situation in the proof of Theorem 7.6 (here for the case that no sites coincide with $x^{(n)}$, i. e., $S(x^{(n)}) = \emptyset$).

We set

$$P(x^{(n)}) := \left\{ x^{(\mathrm{pre}_{T_i}(n))} : i \in I(x^{(n)}) \wedge x^{(n)} \neq s^{(i)} \right\} \tag{7.18}$$

to be the set of all predecessor nodes of $x^{(n)}$ in the trees $T_i[P_i]$ for $i \in I(x^{(n)})$. Also, let

$$S(x^{(n)}) := \left\{ i \in I(x^{(n)}) : x^{(n)} = s^{(i)} \right\} \tag{7.19}$$

be the set of cluster indices such that $x^{(n)}$ coincides with the corresponding sites. The choice of $x^{(n)}$ implies $P_i = \left\{ x^{(n)} \right\}$ for all $i \in S(x^{(n)})$, so the set $S(x^{(n)})$ describes the set of clusters that only consist of $x^{(n)}$. Figure 7.13 illustrates the given situation (with $S(x^{(n)}) = \emptyset$ for the sake of presentability).

We claim that the set of clusters that contain $x^{(n)}$ can be partitioned into the clusters containing each of its predecessor nodes $P(x^{(n)})$ and those consisting only of $x^{(n)}$, i. e.,

$$I(x^{(n)}) = \dot{\bigcup}_{x^{(j)} \in P(x^{(n)})} I(x^{(j)}) \ \dot{\cup}\ S(x^{(n)}). \tag{7.20}$$

Here, "$\subset$" follows directly from the definitions (7.18) and (7.19), while "$\supset$" follows again from Lemma 3.21 (and the definition of $S(x^{(n)})$ in (7.19)). Also from Lemma 3.21 we see that the sets $I(x^{(j)})$ for $x^{(j)} \in P(x^{(n)})$ are disjoint: For any $i \in I(x^{(j)})$, there exists some $l \in I(x^{(j)})$ such that $x^{(j)} = x^{(\mathrm{pre}_{T_l}(n))}$ by definition. Lemma 3.21 then implies $x^{(j)} \succeq_{T_i} x^{(n)}$ and so $x^{(j)}$ must be the unique predecessor of $x^{(n)}$ in T_i.

Now, we adapt the given instance by removing $x^{(n)}$, dropping the clusters in $S(x^{(n)})$ and increasing the weights of predecessor units in $P(x^{(n)})$. W.l.o.g. assume that $S(x^{(n)}) = [k] \setminus [\hat{k}]$ holds for some $\hat{k} \in [k]$. For all $x^{(j)} \in P(x^{(n)})$, set

$$\hat{\omega}_j := \omega_j + \sum_{i \in I(x^{(j)})} \xi^*_{i,n} \omega_n$$

and $\hat{\omega}_j := \omega_j$ for all remaining $j \in \left\{ j \in [n-1] : x^{(j)} \notin P(x^{(n)}) \right\}$. We then set

$$\hat{\xi}_{i,j} := \frac{\xi^*_{i,j}\omega_j + \xi^*_{i,n}\omega_n}{\hat{\omega}_j}$$

for all j such that $x^{(j)} \in P(x^{(n)})$ and $i \in I(x^{(j)})$ as well as $\hat{\xi}_{i,j} := \xi^*_{i,j}$ for all remaining pairs of $i \in [\hat{k}]$ and $j \in [n-1]$.

It is now easy to verify that $\hat{\xi} \in T_{\hat{\kappa},\hat{\omega}}$ (as defined by Eq. (4.16) for cluster weights $(\kappa_i)_{i\in[\hat{k}]}$ and unit weights $(\hat{\omega}_j)_{j\in[n-1]}$) holds. Moreover, it holds that the diagram $\left(P_1 \setminus \left\{x^{(n)}\right\}, \ldots, P_{\hat{k}} \setminus \left\{x^{(n)}\right\}\right)$ obtained from $(f_i|_{X\setminus\{x^{(n)}\}})_{i\in[\hat{k}]}$ supports $\hat{\xi}$, so it must hold that $\hat{\xi}$ lies in the relative interior of the optimal face of (4.18) w. r. t. the adapted instance. Thus, by induction we can w.l.o.g. assume $\hat{\xi}$ to obey (7.17). (Otherwise, select an according $\hat{\xi}$. Note that this would not change the supports of $\hat{\xi}$.)

We can now reverse our previous adaption. We define $\xi^{**} \in \mathbb{R}^{k\times n}$ via

$$\xi^{**}_{i,n} := \hat{\xi}_{i,\mathrm{pre}_{T_i}(n)} \cdot \left(\sum_{l \in I(x^{(\mathrm{pre}_{T_i}(n))})} \xi^*_{l,n} \right)$$

for all $i \in I(x^{(n)}) \setminus S(x^{(n)})$. For all $i \in S(x^{(n)}) = [k] \setminus [\hat{k}]$ we set $\xi^{**}_{i,n} := \xi^*_{i,n}$ and for all $i \in [k] \setminus I(x^{(n)})$ we set $\xi^{**}_{i,n} := 0$. Finally, for all $i \in [k], j \in [n-1]$ we set

$$\xi^{**}_{i,j} := \begin{cases} \hat{\xi}_{i,j} & \text{if } i \in [\hat{k}] \\ 0 & \text{otherwise} \end{cases}.$$

Let us check feasibility, i. e., $\xi^{**} \in T_{\kappa,\omega}$. We get

$$\begin{aligned}
\sum_{i\in[k]} \xi^{**}_{i,n} &= \sum_{i\in I(x^{(n)})\setminus S(x^{(n)})} \xi^{**}_{i,n} + \sum_{i\in S(x^{(n)})} \xi^{**}_{i,n} \\
&= \sum_{i\in I(x^{(n)})\setminus S(x^{(n)})} \hat{\xi}_{i,\mathrm{pre}_{T_i}(n)} \cdot \left(\sum_{l\in I(x^{(\mathrm{pre}_{T_i}(n))})} \xi^{*}_{l,n} \right) + \sum_{i\in S(x^{(n)})} \xi^{*}_{i,n} \\
&\overset{\substack{(7.18),\\(7.20)}}{=} \sum_{x^{(j)}\in P(x^{(n)})} \sum_{i\in I(x^{(j)})} \hat{\xi}_{i,j} \cdot \left(\sum_{l\in I(x^{(j)})} \xi^{*}_{l,n} \right) + \sum_{i\in S(x^{(n)})} \xi^{*}_{i,n} \\
&= \sum_{\substack{x^{(j)}\in P(x^{(n)}),\\ l\in I(x^{(j)})}} \xi^{*}_{l,n} + \sum_{i\in S(x^{(n)})} \xi^{*}_{i,n} \overset{(7.20)}{=} \sum_{i\in I(x^{(n)})\setminus S(x^{(n)})} \xi^{*}_{i,n} + \sum_{i\in S(x^{(n)})} \xi^{*}_{i,n} = 1.
\end{aligned}$$

We conclude that the assignment constraints (2.0a) in the definition of $T_{\kappa,\omega}$ are fulfilled (as for the remaining $j \in [n-1]$ this is inherited from $\hat{\xi}$). Now let us check the balancing constraints (2.0b). For $i \in S(x^{(n)})$ those are fulfilled directly by construction. Let $i \in [k] \setminus S(x^{(n)})$. We claim that the weight of the units in $P(x^{(n)}) \cup \left\{x^{(n)}\right\}$ assigned by ξ^* to the ith cluster is the same as the adapted weight of units in $P(x^{(n)})$ assigned to the ith cluster by $\hat{\xi}$. As other unit assignments are "untouched" by construction, this then shows that constraint (2.0b) is obeyed. If $i \notin I(x^{(n)})$, no weight of those units is assigned (due to (7.20) and as both clusterings have the same unit supports by construction). It holds (once more by (7.20)) that there exists a unique unit $x^{(j)} \in P(x^{(n)})$ with $j \in \mathrm{supp}(C_i{}^*)$. In particular, $x^{(j)} = \mathrm{pre}_{T_i}(x^{(n)})$ must hold and we get

$$\xi^{**}_{i,j}\omega_j + \xi^{**}_{i,n}\omega_n = \hat{\xi}_{i,j}\omega_j + \hat{\xi}_{i,j} \left(\sum_{l\in I(x^{(j)})} \xi^{*}_{l,n} \right) \omega_n = \hat{\xi}_{i,j}\hat{\omega}_j.$$

We conclude $\xi^{**} \in T_{\kappa,\omega}$.

As by construction for all $x^{(j)} \in P(x^{(n)})$ it holds that $\xi^{**}_{i,j} \geq \xi^{**}_{i,n}$, the constraints (7.17) are fulfilled. Also by construction, ξ^{**} has the same cluster supports as the clustering ξ^*. So we deduce that it lies in the relative interior of the optimal face of (4.18). □

Note that the precedence constraints (7.17) are not sufficient for the solution to lie in the optimal face of the clustering polytope $T_{\kappa,\omega}$. An example is provided in Fig. 7.14. However, they obviously do guarantee the contiguity of clusters which in the context of our application might be thought of as "sufficiently supporting".

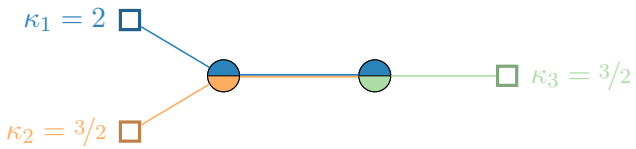

Figure 7.14: Clustering that obeys constraints (7.17), but that is not supported by an optimal shortest-path diagram. Edges depict the trees $T_i[P_i]$ that span the diagram's cells. All units are assumed to have weight 1. Both fractionally assigned units are equally assigned to the indicated clusters. Thus, the clustering does obey constraints (7.17). However, it is not supported by the diagram.

Zoltners and Sinha [ZS83] also introduce precedence constraints similar to (7.17). They conclude those to be already implied by any assignment occurring in their Lagrangian-relaxation, which is a clustering induced by a shortest-path diagram. However, this is clearly not the case for non-trivial intersections of diagram cells (as, for example, in Fig. 7.12). The statement is corrected by Schröder [Sch01, Satz 8.1], who demands uniqueness of the solution. This can be translated into empty cell-intersections in our context. Theorem 7.6 and the succeeding observations can be viewed as a clarification of this issue.

Iterative Rounding Procedure

The second problem that arises from the described nature of the bisectors of shortest-path diagrams is the consequence of greater balancing errors. As we cannot reasonably assume that our clustering solutions are extremal, we are likely to obtain more than the desired $k - 1$ fractionally assigned units. Also, as described in Section 7.2.1, in general we lose the supporting property of a diagram when rounding a fractional clustering. However, if we want to preserve the star-shapedness of the clusters, the subtrees that constitute the diagram's bisectors will have to be assigned uniformly to a single one of their fractionally assigned clusters.

For our application, we tackle this issue in two ways. First, we apply an iterative rounding procedure that does not enforce the assignment of bisector subtrees as a whole. Of course, star-shapedness of clusters is not necessary in order to achieve contiguity of clusters. Assume a fractional clustering that is supported by a shortest-path diagram to be given. For a unit that is fractionally assigned to several clusters, there may exist several paths in the contiguity graph G from this unit to units in the integrally assigned parts of those clusters.

For our toy example solution of Fig. 7.11, we see that indeed there is no necessity to assign all units that are fractionally assigned to the orange, green, or red cluster as a whole to one of them. In fact, most units share an edge with a unit that is integrally assigned to the green cluster.

Algorithm 6: Iterative assignment procedure for shortest path trees.

Input: Shortest-path clustering instance (X, κ, ω, G), sites $s^{(1)}, \dots, s^{(k)} \in X$, clustering $\xi = (C_1, \dots, C_k)^\intercal$ that is optimal for (4.18) with additional constraints (7.17).

Output: Adapted clustering ξ.

repeat

$\hat{X}_f \leftarrow \left\{x^{(j)} : \exists i \in [k] : \xi_{i,j} \in (0,1)\right\} \setminus \left\{s^{(1)}, \dots, s^{(k)}\right\}$

for $i \in [k]$ **do** $S_i \leftarrow \left\{s^{(i)}\right\} \cup \left\{x^{(j)} \in X : \xi_{i,j} = 1\right\}$

$\hat{X} \leftarrow \left\{s^{(1)}, \dots, s^{(k)}\right\} \cup \hat{X}_f =: \left\{x^{(j_1)}, \dots, x^{(j_{\hat{n}})}\right\}$

for $r \in [\hat{n}]$ **do** $\hat{\omega}_r \leftarrow \begin{cases} \omega_{j_r} & \text{, if } x^{(j_r)} \in \hat{X}_f \\ \sum_{x^{(j)} \in S_i} \omega_j & \text{, if } x^{(j_r)} = s^{(i)} \end{cases}$

$\hat{E} \leftarrow E \cap \binom{\hat{X}}{2} \cup \left\{\left\{s^{(i)}, b\right\} : i \in [k], a \in S_i, b \in \hat{X}_f, \{a,b\} \in E\right\}$
$\cup \left\{\left\{s^{(i)}, s^{(l)}\right\} : i \neq l \in [k], a \in S_i, b \in S_l, \{a,b\} \in E\right\}$

for $\{a, b\} \in \hat{E}$ **do**

$\hat{c}(\{a,b\}) \leftarrow \begin{cases} c(\{a,b\}) & \text{, if } \{a,b\} \subset \hat{X}_f \\ \min\{c(\{c,b\}) : c \in S_i\} & \text{, if } a = s^{(i)}, b \in \hat{X}_f \text{ for } i \in [k] \\ \min\{c(\{c,d\}) : c \in S_i, d \in S_l\} & \text{, if } a = s^{(i)}, b = s^{(l)} \text{ for } i \neq l \in [k] \end{cases}$

$\hat{G} \leftarrow \left(\hat{X}, \hat{E}, \hat{c}\right)$

$\hat{\xi} \leftarrow$ Optimizer of (4.18) for $(\hat{X}, \kappa, \hat{\omega}, \hat{G})$ with

$$f_i(x^{(j_r)}) = \begin{cases} \operatorname{dist}_{\hat{G}[\operatorname{supp}(C_i)]}\left(s^{(i)}, x^{(j_r)}\right) & \text{, if } x^{(j_r)} \in \operatorname{supp}(C_i) \\ \infty & \text{, otherwise.} \end{cases}$$

for $i \in [k]$, $x^{(j_r)} \in \hat{X}$ and with adapted precedence constraints

$$\hat{\xi}_{i,\operatorname{pre}_{\hat{T}_i}(j_r)} \geq \hat{\xi}_{i,r} \qquad \forall r \in [\hat{n}] : x^{(j_r)} \in \operatorname{supp}(C_i) \setminus \left\{s^{(i)}\right\} \; \forall i \in [k], \tag{7.21}$$

where $\hat{T}_i$ denotes the shortest-path tree rooted in $s^{(i)}$ in $\hat{G}[\operatorname{supp}(C_i)]$.

for $i \in [k], j \in [n]$ **do** $\xi_{i,j} \leftarrow \begin{cases} \hat{\xi}_{i,j_r}, \text{ if } x^{(j)} \in \hat{X}_f \text{ with } j = j_r, r \in [\hat{n}] \\ \hat{\xi}_{i,j_r}, \text{ if } x^{(j)} \in S_i \text{ with } s^{(i)} = x^{(j_r)}, r \in [\hat{n}] \end{cases}$

until $|\hat{X}|$ *has not changed from the previous iteration.*

return ξ

In order to pay regard to this observation in our rounding procedure, we proceed as follows. Consider an optimal solution ξ of the program (4.18) that fulfills the precedence constraints (7.17). Recall that we aim to construct some $\hat{\xi} \in R(\xi)$ as defined by Eq. (7.2). We now iteratively construct sequences of adapted shortest-path clustering instances. The procedure is formally given by Algorithm 6. In every iteration, we contract all currently integrally assigned nodes of a cluster to a single one (cf. line 3). In particular, every contracted node must contain the respective cluster's site. Thus, sites serve as representatives for contracted nodes. All remaining nodes are collected in the unit set $\hat{X}_f$, which together with the sites forms the adapted unit set (line 4). In the contracted graph, two nodes share an edge whenever any nodes from the respective contracted node sets share an edge in G (line 6). We weight those edges by the minimum weight (i. e., distance) among their corresponding contracted edges (line 8). We then re-solve the problem on this adapted instance. However, in order to ensure that we obtain a new solution in $R(\xi)$, we restrict the shortest-path distances to the cluster support-induced subtrees (line 10). The assignment costs of initially fractionally assigned units are now determined by their shortest-path distances to the contracted site nodes. In particular, this impacts the overlapping of the clusters' shortest-path trees. Thus, one can reasonably hope for less fractionally assigned units. This is now iterated until the clustering does not change after re-solving the problem.

Let us formally state and prove correctness of Algorithm 6.

Theorem 7.7
Given a valid input, Algorithm 6 terminates after at most n iterations with a clustering ξ that is optimal for the program (4.18) *(for the original instance) and such that $G[\mathrm{supp}(C_i)]$ is connected for every $i \in [k]$.*

Proof. First, let us ensure that the program (4.18) as stated in line 10 is feasible with a finite objective value. The current clustering ξ canonically defines a solution for the adapted instances by assigning all units of a contracted node to the corresponding cluster. Furthermore, this solution must be feasible and of finite objective value as distances as well as precedence constraints are determined along shortest paths in the cluster supports.

Vice versa, every clustering obtained in line 11 is feasible for the original version of (4.18) (without precedence constraints). Furthermore, by construction the cluster supports of the final solution are subsets of the respective cluster supports of the input clustering, hence the final clustering must be an optimizer of (4.18).

Contiguity of the resulting clusters is immediately implied by the (adapted) precedence constraints (7.21) and the definition of the contracted graph $\hat{G}$.

Once a unit is integrally assigned to a cluster, it will be contracted with all further integrally assigned nodes of that cluster in line 3. Due to the choice of the functions f_i in line 10, such a node cannot be assigned to a different cluster any more. Thus, the set

of integrally assigned units must increase in every iteration, otherwise the algorithm will terminate. This is the case after at most n iterations. □

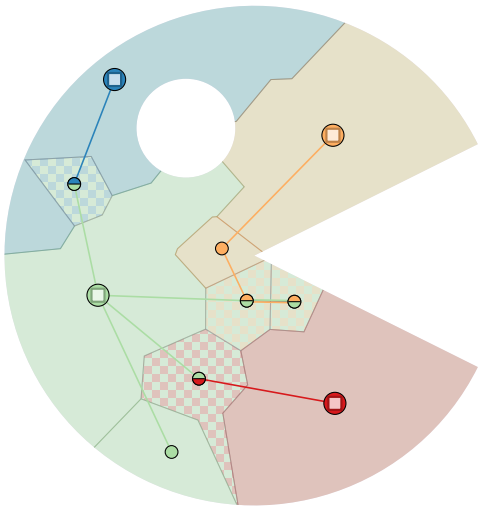

Figure 7.15: Clustering after the first iteration of Algorithm 6.

Let us illustrate Algorithm 6 by means of our toy example. Figure 7.11 depicts a clustering that is optimal for the program (4.18) and that obeys the precedence constraints (7.17) (and that moreover lies in the relative interior of the optimal face as the diagram is supporting). The integrally-assigned units are now contracted and new distances according to Algorithm 6 line 8 are determined. The problem is now re-solved as stated in line 10. Figure 7.15 depicts the resulting solution. This is then repeated a second time, however, without leading to any further integral assignments and hence the algorithm terminates (cf. Fig. 7.16). Doing so, the number of fractional assignments has been reduced from 6 to 4 while all resulting clusters are indeed contiguous.

For the final solution, a rounding procedure is then applied that minimizes the maximum balancing error while obeying the precedence constraints (w. r. t. the final graph in Algorithm 6). As for our application the number of remaining fractional assignments is fairly small, we do so by means of the following simple integer linear program (7.22). Let ξ^* be an optimal shortest-path clustering obtained after applying Algorithm 6. Only for the sake of notation simplicity, we assume here ξ^* to be identical to the initial optimal clustering that is the input for Algorithm 6 (i. e., no contractions

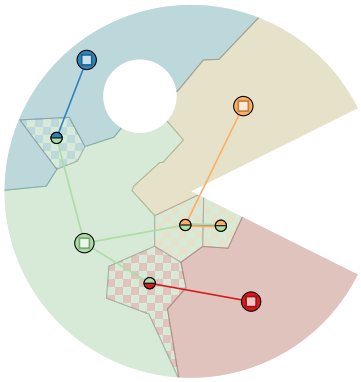

Figure 7.16: Clustering after the second (and final) iteration of Algorithm 6.

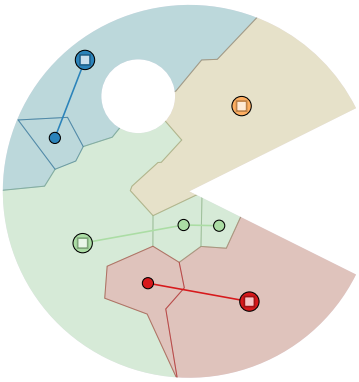

Figure 7.17: Rounded clustering after applying the ILP (7.22).

have been performed).

$$\min_{\delta\in\mathbb{R},\xi\in\mathbb{R}^{k\times n}} \delta \tag{7.22}$$

$$\text{s.t.} \quad \delta \geq \left|\kappa_i - \sum_{j\in[n]} \xi_{i,j}\omega_j\right| \qquad \forall i \in [k] \quad (7.22a)$$

$$\sum_{i\in[k]} \xi_{i,j} = 1 \qquad \forall j \in [n] \quad (7.22b)$$

$$\xi_{i,\mathrm{pre}_{T_i}(j)} \geq \xi_{i,j} \qquad \forall j \in [n] : x^{(j)} \neq s^{(i)}\ \forall i \in [k] \quad (7.22c)$$

$$\xi_{i,j} = 0 \qquad \forall i \in [k], j \in [n] : \xi^*_{i,j} = 0 \quad (7.22d)$$

$$\xi_{i,j} \in \{0,1\} \qquad \forall i \in [k], j \in [n] \quad (7.22e)$$

Note that constraint (7.22a) can, of course, be linearized in the usual way. Also, presumably most columns can be dropped due to (7.22d), as only units that have been fractionally assigned must be re-assigned.

In our toy example, Fig. 7.17 depicts the final clustering after solving the ILP (7.22). Here, a maximum absolute balancing error of 0.5 is obtained, which in fact is best possible (as we are assuming a unit weight on every of the 42 units).

Application of the GVD Schema

Finally, let us summarize how the general schema as depicted in Section 7.2.1 applies to shortest-path diagrams.

As we have seen for our toy example, the iterative procedure in form of Algorithm 6 can help in order to decrease fractionality. However, the resulting number of fractionally assigned units may still be arbitrarily large. For example, consider an instance of two sites that are connected to all remaining units via a single bottleneck unit. Thus, we decided to optimize our choice of sites w. r. t. the resulting balancing error. Of course, other objectives or mixed forms could be thought of, too. Here, we use a simple local search procedure that can be interpreted as a PAM-like method (as described before).

We start with any choice of sites. For example, those can be generated at random. In our application, we will use sites that are closest to the final sites of our power diagram approach.

In every iteration, each of the sites is then exchanged by any non-site unit. For practical reasons, here we reduce the set of exchange candidates to a certain number of closest units of that site w. r. t. G. If there are any improving exchanges w. r. t. our objective criterion, i. e., the resulting rounding error, we choose the best-performing exchange and continue. Algorithm 7 depicts this procedure which yields the proposed realization of the schema in Fig. 7.3 for shortest-path diagrams.

Algorithm 7: Overall shortest-path diagram procedure as applied to the electoral districting design problem.

Input: Shortest-path instance (X, κ, ω, G), start sites $S = \left\{s^{(1)}, \dots, s^{(k)}\right\} \subset X$, parameter $N \in \mathbb{N}_{>0}$

Output: Integral clustering ξ^*

```
Repeat
    δ ← Optimal value of ILP (7.22) after solving (4.18) with functions as given
      by (3.45) for sites S as well as precedence constraints (7.17) and then
      applying Algorithm 6.
    for i ∈ [k] do
        foreach ŝ^(i) from the N closest neighbors of s^(i) in G do
            Ŝ ← S \ {s^(i)} ∪ {ŝ^(i)}
            δ̂ ← Optimal value of ILP (7.22) after solving (4.18) with functions as
              given by (3.45) for sites Ŝ as well as precedence constraints (7.17)
              and then applying Algorithm 6.
    if δ̂ < δ then
        δ ← δ̂
        S ← Ŝ for Ŝ that yielded δ̂.
    else
        return Clustering obtained in the ILP (7.22) that yielded δ.
```

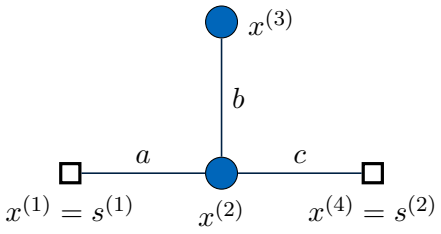

Figure 7.18: Example in proof of Theorem 7.8.

On the Theoretical Limitations of Shortest-Path Diagrams

While shortest-path diagrams come with the great benefit of a-priori-guaranteed contiguity, this comes at the price of less consolidated clusters. In particular, this is the case as the underlying optimization problem is less outlier-sensitive (as, for example, the squared distances in the case of power diagrams).

One natural approach here is to use modified functions f_i in the form of Eq. (3.2) that still rely on the metric induced by the contiguity graph.

However, the following result states that this is — at least in the context of generalized Voronoi diagrams — not feasible. To be more precise, for any transformation function h in Eq. (3.2) that is not affine, clusters may be disconnected despite being supported by the resulting generalized Voronoi diagram. The following result and its proof have been adapted from [BGK17, Theorem 11].

Theorem 7.8 (Variant of [BGK17, Theorem 11])
Let (X, d_G) be the metric space induced by G and let $\mathcal{P} = (P_1, \dots, P_k)$ be the generalized Voronoi diagram w. r. t. functions (f_i) as given by Eq. (3.2) for the metric d_G, a function $h : \mathbb{R} \to \mathbb{R}$, and $\mu_i \in \mathbb{R}$, $i \in [k]$.

If $h(x) = \alpha \cdot x + \beta$ for some $\alpha \in \mathbb{R}_{\geq 0}, \beta \in \mathbb{R}$ then $\mathcal{P}$ is site-star-shaped.

If h is any continuous function but not of the above type, this is not true in general.

Proof. If $\alpha > 0$, then the first claim follows directly from Lemma 3.3 together with Lemma 3.2, as $\alpha \cdot d_G$ again yields a metric on X. The degenerate case $\alpha = 0$ yields cells that are either empty or equal to X. Hence, the first claim follows.

For the second claim, let some continuous function $h : \mathbb{R}_{\geq 0} \to \mathbb{R}$ be given. Consider a graph G as illustrated in Fig. 7.18 with $X = \{x^{(1)}, x^{(2)}, x^{(3)}, x^{(4)}\}$, edges $E = \left\{\{x^{(1)}, x^{(2)}\}, \{x^{(2)}, x^{(3)}\}, \{x^{(2)}, x^{(4)}\}\right\}$ and edge weights $c(\left\{x^{(1)}, x^{(2)}\right\}) = a$, $c(\left\{x^{(2)}, x^{(3)}\right\})) = b$ and $c(\left\{x^{(2)}, x^{(4)}\right\}) = c$ for some $a, b, c \in \mathbb{R}_{>0}$.

We set $\mu_1 := 0$ and $\mu_2 := h(d_G(s^{(1)}, x^{(3)})) - h(d_G(s^{(2)}, x^{(3)}))$. Thus, $x^{(3)} \in P_1 \cap P_2$ holds by construction. If the diagram $\mathcal{P}$ is site-star-shaped, this implies $x^{(2)} \in P_1 \cap P_2$ and hence $h(d_G(s^{(1)}, x^{(2)})) = h(d_G(s^{(2)}, x^{(2)})) + \mu_2$. Together, this gives

$h(d_G(s^{(1)}, x^{(3)})) - h(d_G(s^{(2)}, x^{(3)})) = h(d_G(s^{(1)}, x^{(2)})) - h(d_G(s^{(2)}, x^{(2)}))$. In terms of a, b, and c and using the shortest-path distances this yields

$$h(a+b) - h(b+c) = h(a) - h(c).$$

Set $\tilde{h} := h - h(0)$. By taking the limit $c \to 0$ and using the continuity of h we get

$$\tilde{h}(a+b) = \tilde{h}(a) + \tilde{h}(b).$$

Since $a, b \in \mathbb{R}_{>0}$ are arbitrary and $\tilde{h}$ is continuous, it readily follows that $\tilde{h}$ is linear on $\mathbb{R}_{\geq 0}$ and thus h is affine. Finally, let us ensure that $\alpha \geq 0$ is required in general. Consider the example $X = \{x^{(1)}, x^{(2)}\}$, $s^{(i)} = x^{(i)}$ for $i = 1, 2$ and a single edge $\left\{x^{(1)}, x^{(2)}\right\}$ of arbitrary positive length. Set $\mu_1 := \mu_2 := 0$. Then for $\alpha < 0$ we get $P_1 = \left\{x^{(2)}\right\}$ and $P_2 = \left\{x^{(1)}\right\}$, so $\mathcal{P}$ is not site-star-shaped. □

At first glance, Theorem 7.8 suggests that using outlier-sensitive transformations (such as squaring distances) is unsuitable if star-shapedness (as a warrant for contiguity) should be guaranteed. However, this only affects the "strict" usage of shortest-path-diagrams. For example, one could still attempt a hybrid approach of power diagrams and shortest-path diagrams which uses squared Euclidean distances as proximity measure while enforcing continuity by means of precedence constraints (7.17). While this would not yield clusterings that are supported by a generalized Voronoi diagram, this could help to overcome some of the drawbacks of shortest-path-diagrams w. r. t. consolidation.

7.3.3 Anisotropic Power Diagrams

As a third approach, we apply anisotropic power diagrams to the electoral district design problem.

While our previous approaches show very promising techniques in order to produce well-suited district plans, they disregard a maybe unwritten but rather natural requirement: continuity. Redistricting has to be applied regularly due to census changes. However, it is plausible to demand this adaption to be done with preferably as little changes to the existing districts as necessary. Thus, one might want to pay regard to the "difference" between the newly generated and the currently existing district plan in an optimization procedure.

In order to do so in our context, we make use of anisotropic power diagrams. As mentioned before, those have been used successfully in the context of grain reconstruction ([Alp+15]). Here, the core idea is to resemble a tiling of a certain area into grains by means of an anisotropic power diagram. We can adapt the same methodology to our application.

Choice of Local Norms

Let $\xi^{(0)} \in T_{\text{uncstr.},k,n}$ be an existing clustering obtained from a current district design.

As in [Alp+15], we now want to determine a local norm for every district that inherits some information on its shape. If a district is elongated in a certain direction, the local norm should be chosen such that distances in that elongated direction are relatively decreased compared to other directions. From the point of view of a single district, we would like to change our normed space such that the original district is uniformly spread into all directions. A typical tool in order to perform such a data transformation is, of course, *principal component analysis* (PCA). Let $c_i^{(0)} := c\left(C_i^{(0)}\right) \in \mathbb{R}^2$ denote the centroid of the ith cluster of $\xi^{(0)}$ for $i \in [k]$.

For every $i \in [k]$, the covariance matrix of the ith cluster is given by

$$\operatorname{cov}(C_i^{(0)}) = \sum_{j\in[n]} \xi_{i,j}^{(0)} \frac{\omega_j}{\omega(C_i^{(0)})} \left(x^{(j)} - c_i^{(0)}\right)\left(x^{(j)} - c_i^{(0)}\right)^{\mathsf{T}}. \tag{7.23}$$

We then set

$$\Sigma_i := \operatorname{cov}(C_i^{(0)})^{-1} \tag{7.24}$$

for $i \in [k]$.

For every district $i \in [k]$, we can now consider the Minkowski space $(\mathbb{R}^d, \Sigma_i^{-\frac{1}{2}}\mathbb{B}_2^d)$ equipped with the ellipsoid $\Sigma_i^{-\frac{1}{2}}\mathbb{B}_2^d$ as the unit ball. Let us provide a quick sanity check why this choice of a space could be useful in our context. We can map this space into the standard Euclidean space via $x \mapsto \hat{x} := \Sigma_i^{\frac{1}{2}} x$. By construction, the covariance of the district under this mapping is

$$\begin{aligned}
&\sum_{j\in[n]} \xi_{i,j}^{(0)} \frac{\omega_j}{\omega(C_i^{(0)})} \left(\hat{x}^{(j)} - \hat{c}_i^{(0)}\right)\left(\hat{x}^{(j)} - \hat{c}_i^{(0)}\right)^{\mathsf{T}} \\
&= \sum_{j\in[n]} \xi_{i,j}^{(0)} \frac{\omega_j}{\omega(C_i^{(0)})} \Sigma_i^{\frac{1}{2}} \left(x^{(j)} - c_i^{(0)}\right)\left(x^{(j)} - c_i^{(0)}\right)^{\mathsf{T}} \Sigma_i^{\frac{1}{2}} = \mathrm{Id}_d.
\end{aligned} \tag{7.25}$$

So indeed, the district may be considered as equally spread in all directions in this space.

Consequently, we choose the centroids as sites, i. e., $s^{(i)} := c\left(C_i^{(0)}\right)$ for $i \in [k|$, and the inverted covariance Σ_i for $i \in [k]$ as norm-defining matrices for an ellipsoidal parametrization of the desired anisotropic power diagram. Thus, the pair $\left((\Sigma_i)_{i\in[k]}, \left(c_i^{(0)}\right)_{i\in[k]}\right)$ yields the structural parameter S in the context of our methodological schema from Section 7.2.1. For every cluster $i \in [k]$, we thus set

$$f_i(x) := \left(x - c_i^{(0)}\right)^{\mathsf{T}} \Sigma_i \left(x - c_i^{(0)}\right) = \left\|x - c_i^{(0)}\right\|_{\Sigma_i}^2 \tag{7.26}$$

as distance functions for the problem (4.18).

Note that by construction, for every $i \in [k]$ we get

$$\sum_{j\in[n]} \xi_{i,j}^{(0)} \omega_j \left\| x - c_i^{(0)} \right\|_{\Sigma_i}^2 = \sum_{j\in[n]} \xi_{i,j}^{(0)} \omega_j^{\mathsf{T}} \Sigma_i \left(x - c_i^{(0)} \right) \left(x - c_i^{(0)} \right)^{\mathsf{T}} \overset{(7.25)}{=} d \cdot \omega(C_i^{(0)}).$$

This means, the local norms have been chosen such that the original clustering yields normalized costs of $d \cdot \omega(X)$.

Post-Processing

Similar to the case of power diagrams in Section 7.3.1, there is no guarantee of contiguity of districts. Even worse, for the case of anisotropic power diagrams we know that the diagram cells in the plane might be non-contiguous themselves (cf. Section 3.3.3).

Thus, in order to ensure contiguity, we apply the same post-processing steps as we did with power diagrams by the introduction of constraints of type (7.15) and (7.16) and resolving the problem (4.18) (see Section 7.3.1).

Again, we will see that this turns out to be very manageable for our real-world application.

Overall Workflow

Let us shortly summarize our proposed overall workflow for the usage of anisotropic power diagrams for the electoral district design problem.

Given district plan — or clustering — we define the functions $(f_i)_{i\in[k]}$ via a principal component analysis as given by Eqs. (7.24) and (7.26). We then solve the program (4.18) and obtain a fractional clustering together with additive weights $M = (\mu_i)_{i\in[k]}$ such that the anisotropic power diagram w. r. t. functions $(f_i + \mu_i)_{i\in[k]}$ supports the clustering. If there are non-contiguous districts, we perform a post-processing as we do for the case of power diagrams (see Section 7.3.1). Recall that this preserves the feasibility of the diagram. Finally, we apply a rounding procedure. Due to Theorem 7.1, we can reasonably assume that the optimum of the problem (4.18) is unique and hence we can solve a $\infty-$Split Resolution as described in Section 7.2.2.

Figure 7.19 depicts some (more or less) arbitrary clustering for our toy example. Figure 7.20 then shows the resulting clustering and supporting anisotropic power diagram obtained from the program (4.18). In particular, it depicts the ellipses that are the unit balls of the respective local norms.

Comparison of District Plans

In order to evaluate and compare the different approaches in regard of our main intention for using anisotropic power diagrams, we need a suitable measure. For our

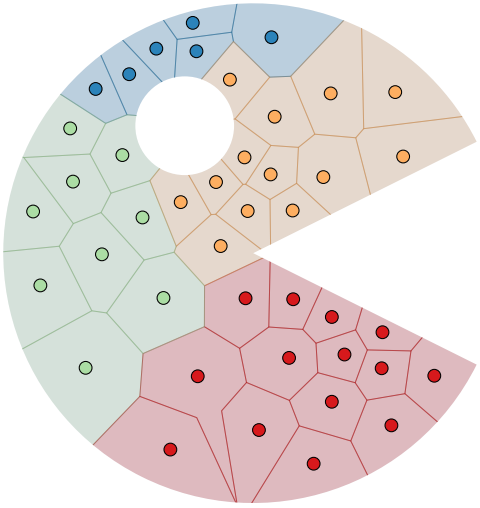

Figure 7.19: Assumed initial districts for the anisotropic power diagram approach.

purposes, we want to consider two district plans to be similar if pairs of voters who are assigned to a common district in one plan are not separated in the other. With $\xi^{(0)}$ still being our reference clustering, let $\xi^* \in \{0,1\}^{k \times n}$ denote a newly generated one that we assume to be integer. In total, there are $\sum_{i \in [k]} \binom{\omega(C_i^{(0)})}{2}$ pairs of voters that are assigned to the same district by $\xi^{(0)}$. Let

$$A(\xi^{(0)}, \xi^*) := \left\{ \{j,r\} \in \binom{[n]}{2} : \exists i \in [k] : \xi_{i,j}^{(0)} = \xi_{i,r}^{(0)} = 1 \ \wedge \ \forall i \in [k] : \xi_{i,j}^* \cdot \xi_{i,r}^* = 0 \right\}$$

denote the index set of all pairs of units that are assigned to a common cluster by $\xi^{(0)}$ but to different clusters by ξ^*. Then

$$\Delta(\xi^{(0)}, \xi^*) := \frac{1}{\sum_{i \in [k]} \binom{\omega(C_i^{(0)})}{2}} \sum_{\{j,r\} \in A(\xi^{(0)}, \xi^*)} \omega_j \cdot \omega_r \tag{7.27}$$

gives the ratio of voter pairs that clustering ξ^* separates in comparison to $\xi^{(0)}$.

7.4 Results

We demonstrate the proposed methodology by the example of the federal elections in Germany. Sections 7.4.1 and 7.4.2 will describe the dataset and test environment, Sections 7.4.3 to 7.4.5 describe the implementations of the different used diagram types.

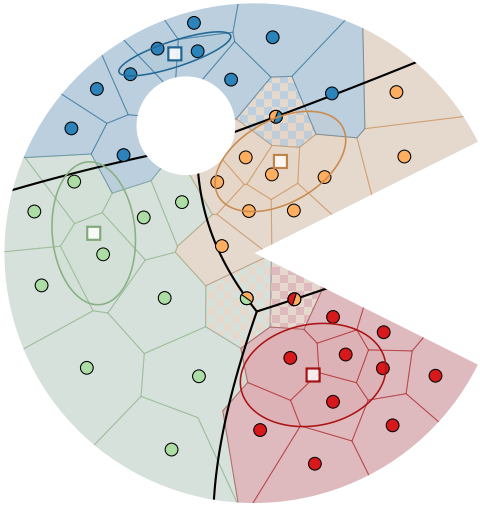

Figure 7.20: Anisotropic power diagram clustering of the toy example. The colored ellipses depict the respective cluster norms' unit balls that result from the principal component analysis of the clusters in Fig. 7.19.

Results and figures are then provided and compared in Section 7.4.6. More detailed result tables can be found in Appendix B.

7.4.1 Dataset

The German electoral law ([BWG]) states that 299 electoral districts are to be apportioned to the 16 federal states. This apportionment is done following the Sainte-Laguë method ([Wikc]). Electoral districts have to obey the federal state borders and must be contiguous. As far as possible, administrative borders ("Kreise", "Gemeinden") shall be obeyed. Thus, in our context each federal state yields an electoral district design instance on its own. By law, each district's size should not deviate from the federal average by more than 15%; a deviation exceeding 25% enforces a redesign of districts. According to a ruling by the Federal Constitutional Court, the district's size is determined by its number of eligible voters ([Bun12]).

We mention that in particular the number of electoral districts is currently politically discussed in Germany ([SpOn19; Zei19]). However, neither the number nor the way of apportionment of districts to states is part of this work. Also, we do not consider the administrative levels above municipalities, i. e., counties ("Kreise"), as their total number of 489 for the whole of Germany ([SABL19]) is, of course, way too narrowing in order to build 299 districts. Also, the (political) evaluation of which county may or may not be split is beyond the demonstrative purpose of this work.

We use the data from the federal German elections held on September 24th 2017. Note that analogous results on the database of the federal elections of September 22nd 2013 have been published in [BGK17].

The data of the electoral districts of 2017 has been taken from [Bwl16], municipality sizes, i. e., number of eligible voters, root from [SABL19, Table 14111-01-03-5]. The geographical data has been taken from [BKG19, Dataset VG250-EW, 1.1.2019][2]. Geographic shapes are mapped to the plane by the EPSG 25832 spatial reference system ([Wikd]).

Overall, there are 11 091 municipalities with an average amount of 5 556 eligible voters per municipality. There are a total of 61 688 482 eligible voters, so that the 299 districts yield a desired size of 206 316 voters per district. However, as the 16 federal states are to be treated independently, the average district sizes of each state actually ranges between 194 316 (Saarland) and 237 076 (Bremen).

Next, several larger municipalities formed several districts on their own in the 2017 election. For example, the city of Munich is split into 4 districts. As this results in large units that a priori yield districts on their own, we did not include them in our dataset. More precisely, we excluded all municipalities that are at least 5% greater than the desired district size. Also, the cities Berlin, Bremen and Hamburg, that are federal states on their own, were excluded for the same reason.

Overall, this reduces our data set to a total of 13 instances that in summary have 266 remaining districts. Table B.1 in the appendix lists more detailed information on the original districts of 2017, while Table B.2 yields the key figures for our resulting 13 clustering instances.

Figure 7.21 depicts the districts of the 2017 election. Here and in all further illustrations of district plans, we use a total of 6 colors to color the districts. Colors are chosen such that neighboring districts are colored differently. Furthermore, by mapping the districts of each of the generated clusterings to the original districts of 2017 and coloring mapped districts identically, the illustrations are sought to be comparable more easily. Municipalities that are excluded for the reasons described above are depicted by a gray-shaded filling.

For each municipality, the centroid of its territory was determined as a representative in the unit set X. A contiguity graph G was derived by connecting units whose territories share a common border. Some additional edges were introduced in order to connect island units to the mainland by considering ferry connections. Euclidean distances of the unit points were then used as edge lengths. Figure 7.22 depicts the resulting graphs for all 13 instances.

[2]Both data from [SABL19] as well as [BKG19] are licensed under the license “Data licence Germany – attribution – version 2.0”, see [GovD]. All geographical material in this section is ©GeoBasis-DE / BKG (2019).

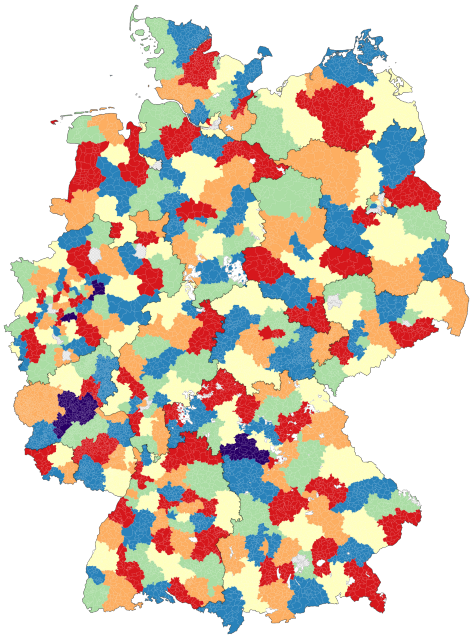

Figure 7.21: Districts of the 2017 federal elections in Germany.

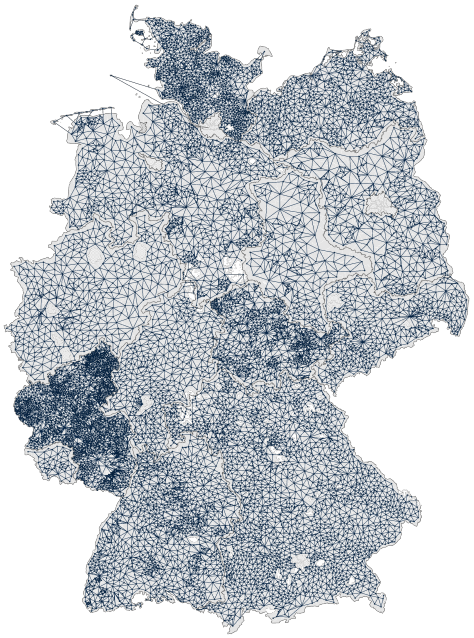

Figure 7.22: Unit centroids and contiguity graphs of the 13 instances.

7.4.2 Testing Environment and Implementation Notes

The runs were performed in Linux environments with slightly varying specification. However, our application is not sensitive to running times, so we did not perform any exact measurements w. r. t. those. From our observations, all approaches should be executable on a reasonably modern desktop computer at reasonable running times (i. e., at most a few hours per instance and approach).

The main algorithms were implemented in Java ([Orac]). For solving the underlying linear and integer linear programs, the solvers FICO Xpress ([FICO]) as well as Gurobi ([Guro]) were used.

The open-source software QGIS ([QGIS]) was used for graphics generation as well as for the interactive post-processing with power diagrams and anisotropic power diagrams. Here, we used the QGIS Python interface in order to inject our own Python ([Pyth]) code that performs the computations of further graphics such as diagram lines, interacts with the Java backend and enables a semi-automated generation of figures. Data storage and exchange between the several application levels was done in a PostgreSQL database ([PSQL]) equipped with the PostGIS extension ([PGIS]) for handling geographic data.

7.4.3 Power Diagrams

Let us discuss the results for power diagrams following the methodology of Section 7.3.1. Figure 7.23 depicts the result for the state of Bavaria.

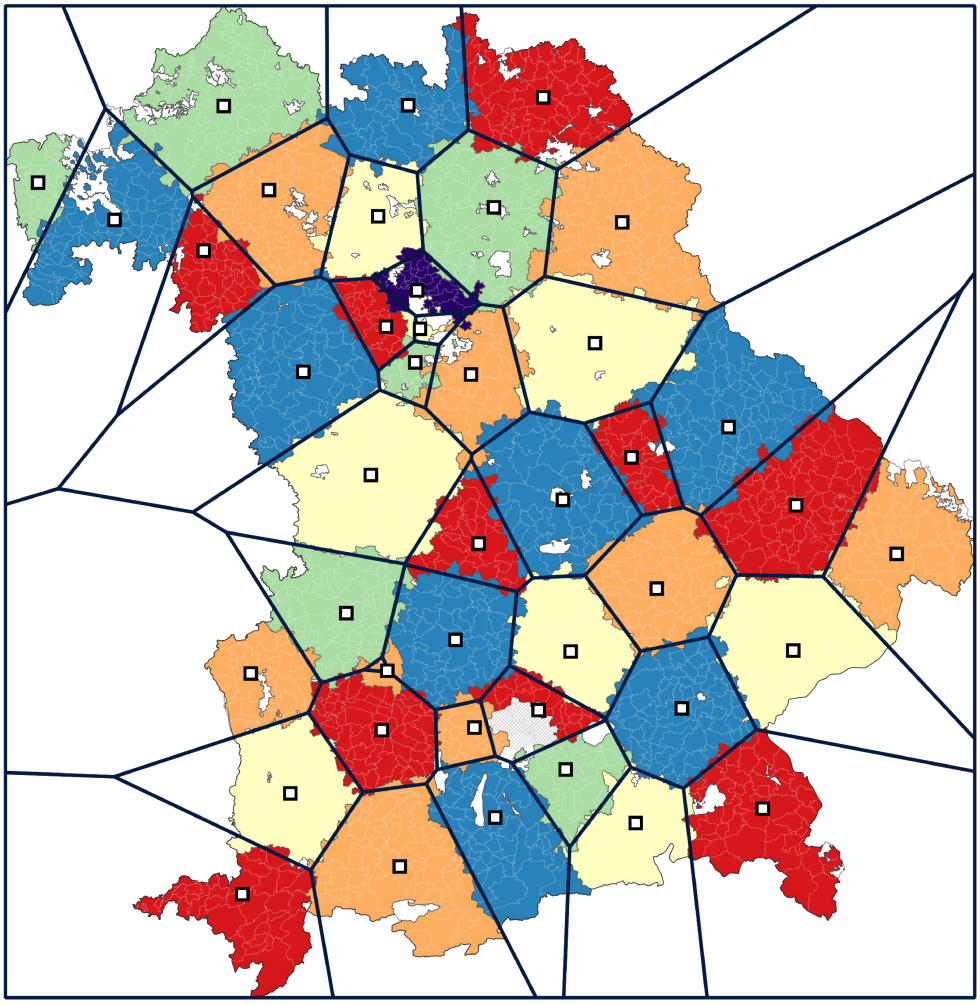

Figure 7.23: Power diagram clustering result for the state of Bavaria.

Necessary Number of Iterations

First, we discuss what number N of samples of start sites for the Balanced k-Means-algorithm may be reasonable w. r. t. our data.

In order to do so, an assumably large number of 50 000 samples of sites for each of the 13 state instances was generated and used as start sites for a Balanced k-Means run.

For each run, we then considered an estimated approximation ratio by taking the ratio of the resulting Balanced k-Means objective value and the best solution found over all samples. As it turns out, the final objective values after running the Balanced k-Means-algorithm are mostly in close range to the best found solution for all instances in our dataset.

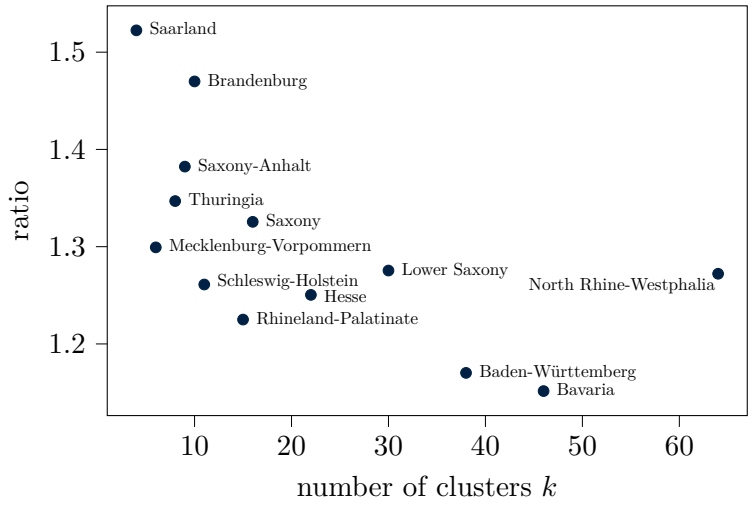

Figure 7.24: For each problem instance, the ratio between the worst and the best obtained objective values of 50 000 Balanced k-Means runs in relation to the number of clusters.

Figure 7.24 depicts for each of the thirteen instances the ratio of the worst and best found solution over all 50 000 samples. Here, we observe that for our given data this ratio is at worst 1.52 for the state of Saarland. The latter is also the smallest among all considered states, while the ratio of all states with at most 20 districts never exceeds 1.3. Going more into detail, we are interested in the resulting empirical distributions of the ratio between each and the best found objective value. Figure 7.25 shows the histograms of this ratio for the five largest states, while Fig. 7.26 yields the histograms for the remaining eight states.

Here, the histograms for smaller states reveal a less continuous behavior compared with the greater states. This is also in accordance with the observation that for smaller

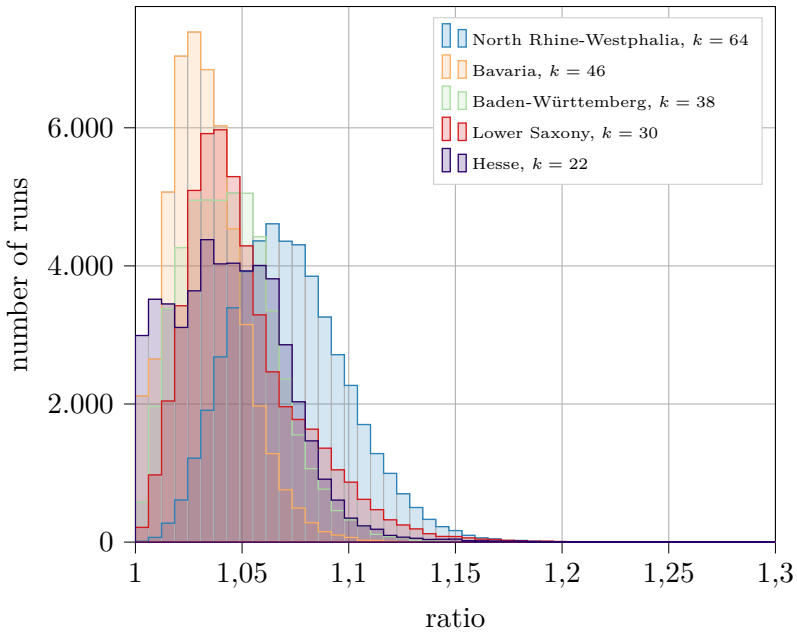

Figure 7.25: Histograms of the ratios between the individual and the best found objective values over the 50 000 Balanced k-Means runs. Here, for the five largest federal states.

states less distinct solutions were found. While for the greatest state of North Rhine-Westphalia no solution in our test run was obtained twice, for the "medium"-sized state of Saxony 4 665, and for the smallest state of Saarland only 51 distinct solutions occurred.

Most important for our approach, for all instances a significant amount of samples lies very close to the best found solution. In order to quantify this, let $\alpha > 1$ and consider S to be a random sample of structural parameters, i. e., sites, as input for the Balanced k-Means-algorithm. Let ξ^S denote the resulting clustering and for $\xi \in T_{\kappa,\omega}$ let $g(\xi)$ denote the resulting Balanced k-Means objective value (i. e., the moment of inertia, cf. Section 6.1.2). With ξ^* being the optimal clustering, we would like to know the probability

$$p^* := \mathbb{P}\left(\frac{g(\xi^S)}{g(\xi^*)} \leq \alpha\right) \tag{7.28}$$

of the random sample yielding an objective value within a ratio of α to the optimal solution. Of course, we do not know ξ^*. Hence, we replace it by $\hat{\xi}^*$ to be defined as the best solution found in our empirical runs. However, our observations, in particular of Figs. 7.25 and 7.26, suggest that we can reasonably assume $g(\xi^*) \approx g(\hat{\xi}^*)$. We can

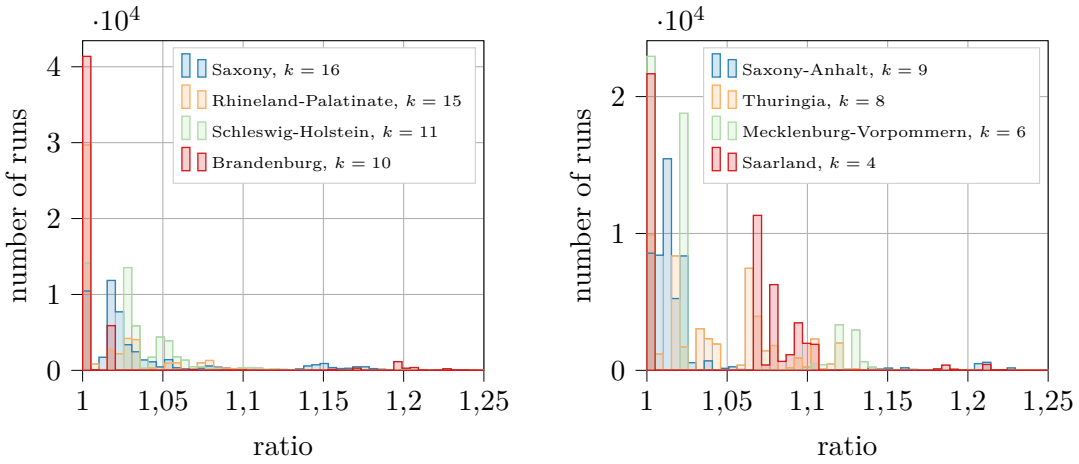

Figure 7.26: Histograms of the objective ratios as in Fig. 7.25 for the remaining states. Here, a few outliers (at most 0.32% per state) exceeding a ratio of 1.25 have been dropped for the sake of presentability (cf. also the worst-case values in Fig. 7.24).

then obtain a confidence-interval for $\mathbb{P}\left(\frac{g(\xi^S)}{g(\xi^*)} \leq \alpha\right)$ (i. e., for the success probability of a Bernoulli distribution). Following [DCB01] we use the Agresti-–Coull interval as estimator. We then consider the lower bound of this interval in order to compute the number of samples necessary to achieve a ratio of α at some boosted target probability.

For our setting, we choose $\alpha := 1.01$ and estimate the success probability at a 99.9% confidence level. We then compute the number of runs that are necessary to have a 99.9% likelihood (at this level of confidence) of obtaining a 1.01-approximation (of the best solution found among all samples). Figure 7.27 depicts the resulting estimates for the required numbers of runs compared to the instance's number of clusters. Here, we see that for all except one instance, less than $N = 524$ runs (obtained for the state of Lower Saxony) suffice. An outlier is given by the state of North Rhine-Westphalia with 14 856 runs. However, note that even for this state, 1000 iterations would suffice for a ratio of $\alpha = 1.02$ with the same probability and confidence. On average, the Balanced k-Means-algorithm took 20.6 iterations (with the averages of the individual instances weighted by their number of clusters). Thus, for our given instances, a value of $N = 1000$ results in solving roughly $2 \cdot 10^4$ transportation problems of the form (4.18), which is not an actual computational challenge on a modern machine.

Of course, our data is too limited in order to allow a generalization. Still, it indicates that a very reasonable number of iterations may suffice in order to obtain a satisfying approximation. In particular this holds true for the greater instances in our data, maybe excluding North Rhine-Westphalia. Note that the latter instance also considerably differs from the other instances w. r. t. the number of units per cluster as it is the

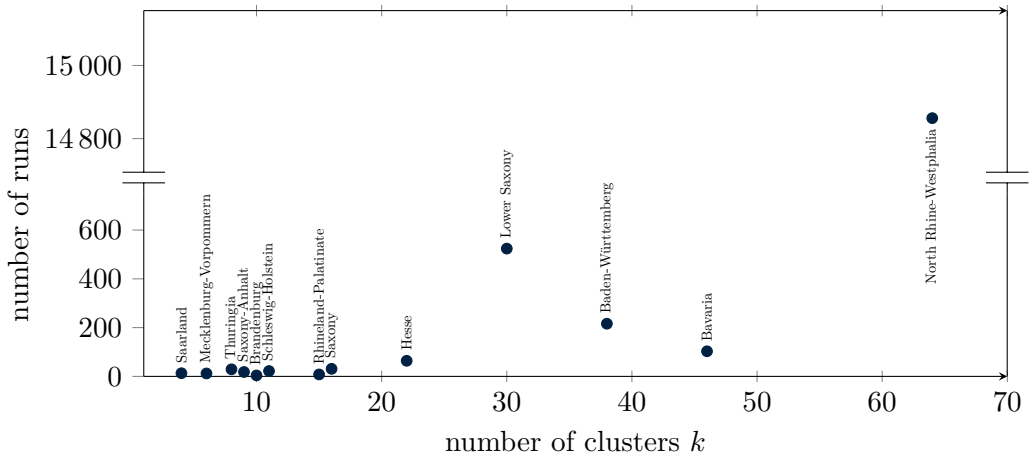

Figure 7.27: Estimated number of runs necessary to obtain a 1.01-approximation ratio (w. r. t. the best solution found) with a probability of 99.9% at a 99.9% confidence level.

densest-populated state in our dataset (cf. Table B.2).

Let us furthermore attempt to illustrate the practical meaning of accepting an approximation ratio of $\alpha = 1.01$. Here, let us assume that the optimal clustering ξ^* is integral and for each $j \in [n]$ let $i(j)$ be the unique cluster index that the jth unit is assigned to. Observe that if we replaced each unit representative $x^{(j)}$ by a translation $x^{(j)} + (\sqrt{\alpha} - 1)\left(x^{(j)} - c^{(i(j))}\right)$ while keeping the clustering ξ^* as well as the centroids (by construction) fixed, the Balanced k-Means objective value would increase by a factor of α. In our data, this translation is on average of length 68.8m. By comparison, the average maximum distance of a representative $x^{(j)}$ to the respective unit's border measures 8.25km.

Balancing Behavior

Using the procedure as described so far we obtained an average relative balancing error after applying the rounding procedure of 5.2%. Moreover, the maximum error for the individual instances ranged between 7.3% and 45.7% due to some greater cities being fractionally assigned. However, in the preceding section we observed that a considerably great ratio of solutions obtained from running the Balanced k-Means-algorithm for a set of randomly sampled sites yields satisfying values w. r. t. the moment of inertia. Thus, we decided to tweak the procedure by rejecting any result whose maximum balancing error after applying the rounding procedure exceeded a predefined threshold.

For all but three states in our dataset we set this threshold to 5%. However, for the states of Saxony, Saxony-Anhalt and North Rhine-Westphalia, we had to increase this threshold to 8%, 12.5%, and 15%, respectively, as in fact no feasible solution within

the threshold of 5% could be found for those instances.

This, of course, weakened the results w. r. t. the BALANCED k-MEANS objective. Here, we observed an average relative increase of the BALANCED k-MEANS objective by 2.5%. For the state of Saxony-Anhalt, a particularly high increase by 23.6% occurred (while all remaining states yielded increases of less than 5%).

Post-Processing

Next, we applied the post-processing in order to ensure contiguity of districts. Here, we were successful in establishing contiguity for all our instances. Overall, 19 (out of 266) clusters were not contiguous and we injected 30 constraints of type (7.15) as well as 8 constraints of type (7.16). This affected 36 units, i. e., 0.3% of all units.

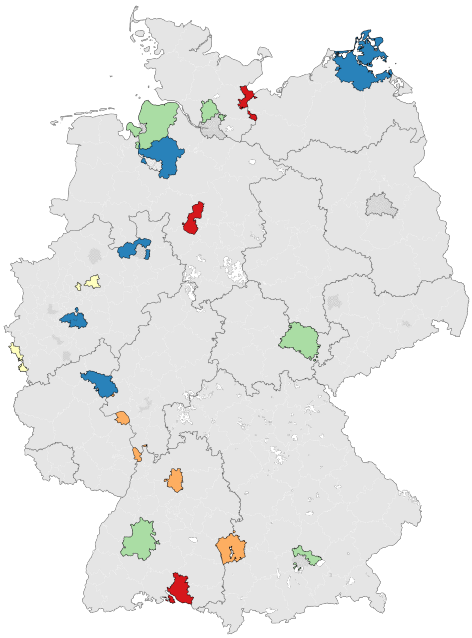

Figure 7.28: Non-contiguous districts (colored) before post-processing. Over all instances, 19 districts were non-contiguous.

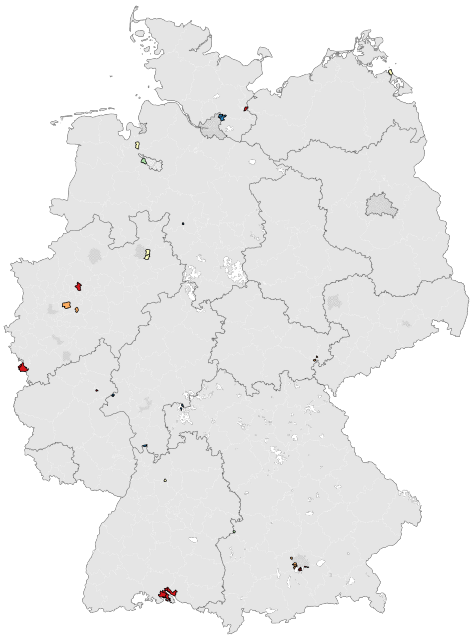

Figure 7.29: Units affected by the power diagram post-processing (colored, 36 in total).

Figure 7.28 depicts the non-contiguous clusters after the initial run, while Fig. 7.29 depicts all units that were affected by the additionally introduced constraints. By using the powerful open-source software QGIS ([QGIS]) that beyond the image-processing of geographic maps offers the capability of injecting individual source code, this process

turned out to be fairly convenient.

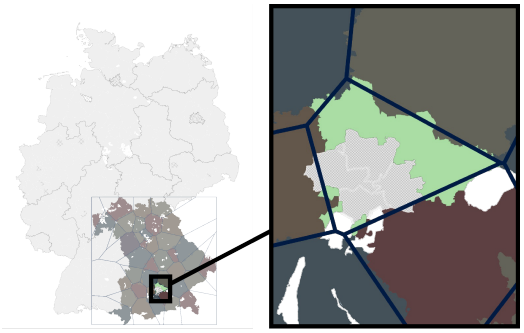

Figure 7.30: Non-contiguous district around the city of Munich. The gray-shaded areas depict the excluded districts that belong to Munich.

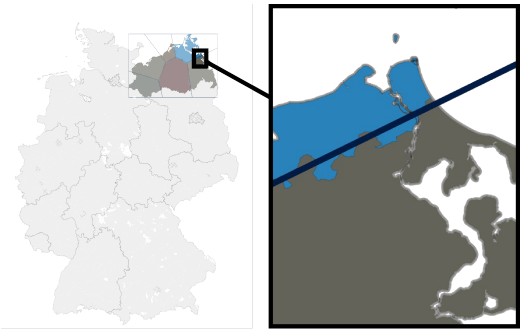

Figure 7.31: Non-contiguous district around the city of Peenemünde, that belongs to the northern part of the island Usedom.

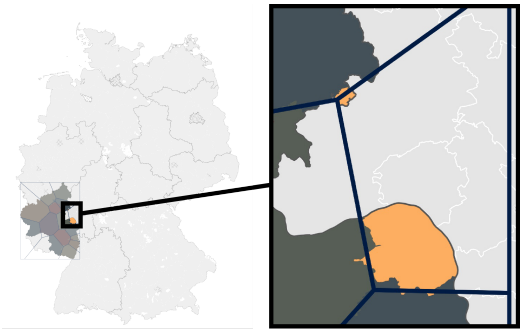

Figure 7.32: Non-contiguous district resulting from an indentation in the eastern border of Rhineland-Palatinate.

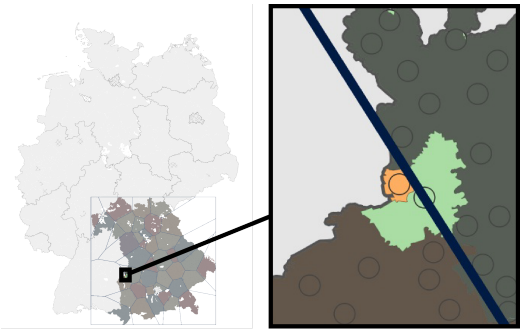

Figure 7.33: Non-contiguous district in Bavaria which is caused by the representation of unit areas by their centroids (depicted as circles).

Still, artifacts as described in Section 7.3.1 do occur in our real-world data. Figures 7.30 and 7.31 depict examples of "holes" separating a district in analogy to Fig. 7.9a. Here, Fig. 7.30 is due to four excluded districts of the city of Munich, while in Fig. 7.31 the town of Peenemünde, that belongs to the most northern part of the island Usedom, is separated from the western mainland. Artifacts in analogy to Fig. 7.9b can be observed, too. An example is depicted in Fig. 7.32, where an indentation in the eastern border of the state of Rhineland-Palatinate results in a non-contiguous district. Also, non-contiguity due to the representation of units by single points as we illustrated in Fig. 7.8 can be observed. An example is given by Fig. 7.33 that depicts how the area of the Bavarian town Gundelfingen an der Donau encloses the municipality of

Bächingen which is hence separated from its assigned district.

The post-processing also has some impact on the resulting Balanced k-Means objective values as well as on the balancing error. Here, the results did not differ notably w. r. t. the moment of inertia (with an average relative change of 0.4%). However, w. r. t. the balancing error, there was a slight increase of the observed maximum absolute deviation per state instance by 1.14 percentage points.

7.4.4 Shortest-Path Diagrams

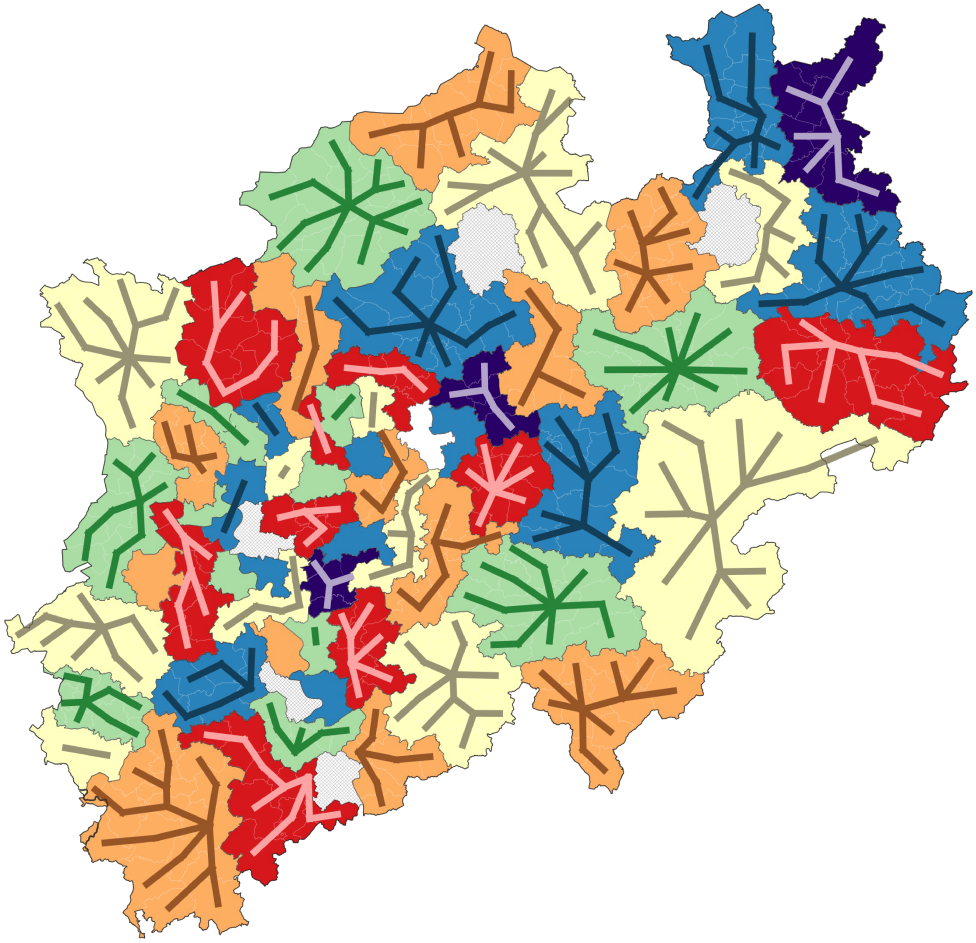

Figure 7.34: Shortest-path diagram clustering result for the state of North Rhine-Westphalia.

Next, we quickly discuss the application of shortest-paths to our data following Section 7.3.2. Figure 7.34 depicts an exemplary result for the state of North Rhine-

Westphalia. In particular, it also shows the shortest-path trees that represent the corresponding shortest-path diagram cells.

We implemented the local search procedure including the iterative rounding procedure as given by Algorithm 7. Here, we chose a parameter $N = 50$ for the neighborhood sizes. As initial sites, we chose the units closest to the sites resulting from the power diagram approach of Section 7.4.3. For all instances in our data set, this led to total running times of less than an hour in our testing environment.

7.4.5 Anisotropic Power Diagrams

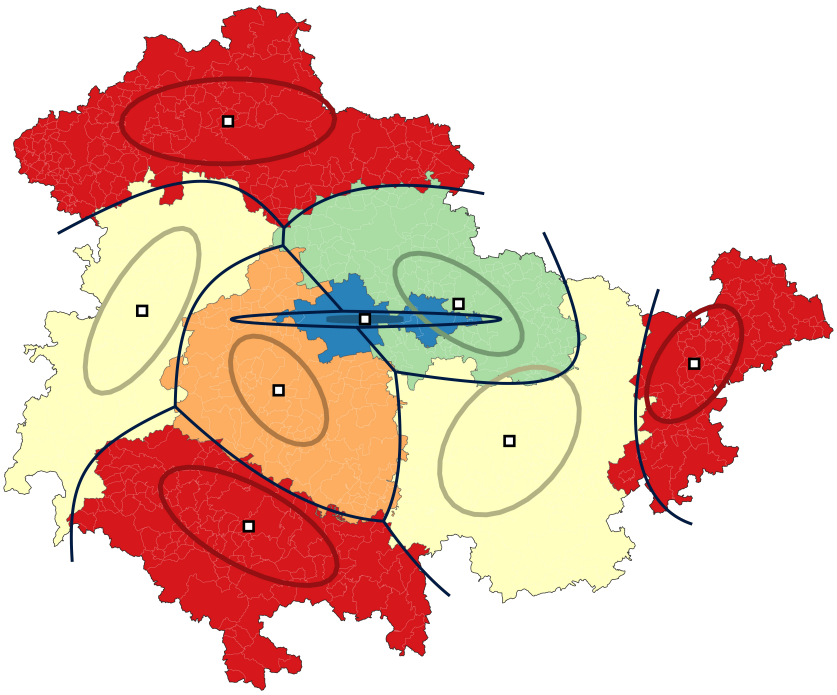

Figure 7.35: Resulting district plan and anisotropic power diagram for the state of Thuringa. The colored ellipses depict the unit balls of the associated local norms.

The anisotropic power diagram approach as described in Section 7.3.3 was implemented in a straightforward way.

One issue, however, is the existence of districts that consist of a single or two units in the original clustering. Here, the resulting covariance matrix is of rank 0 or 1, respectively, and so a corresponding norm cannot be determined in the proposed way. For the case of singletons, the unit balls were set to Euclidean balls. Here, the radius

was determined as a certain fraction of the minimum among all occurring axis lengths of the remaining ellipses. For the case of districts containing exactly two units, the zero eigenvalue of the covariance matrix was replaced by a fraction of the other, non-zero eigenvalue in order to mimic a "very small" variation in that direction.

Figure 7.35 shows an exemplary result for the state of Thuringa. For comparison, Fig. 7.36 depicts the original districts of the 2017 elections. We will discuss and compare the results in more detail in Section 7.4.6.

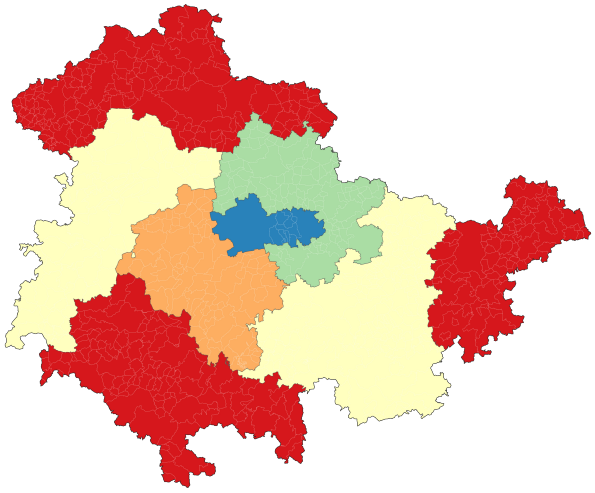

Figure 7.36: Districts of the state of Thuringa in the 2017 German elections.

Post-Processing

We applied the same post-processing in order to ensure contiguity as we did for the case of power diagrams.

In accordance with the fact that anisotropic power diagrams can be non-contiguous themselves, we indeed observed a higher number of non-contiguous districts compared to the power diagram approach. One example for such a non-contiguous district is illustrated in Fig. 7.37. Here, Fig. 7.37a depicts the original districts of the 2017 elections. Two districts and their corresponding ellipses that yield their norms' unit balls are highlighted. Their elongated shape and almost orthogonal orientation then causes one of the two corresponding diagram cells to be non-contiguous (Fig. 7.37b).

However, the number of non-contiguous clusters as well as the number of constraints necessary to resolve this still turned out to be very manageable.

In total, 31 (out of 266) clusters were not contiguous. We added 54 constraints of type (7.15) as well as 7 constraints of type (7.16). Overall, 59 units, i. e., 0.5%

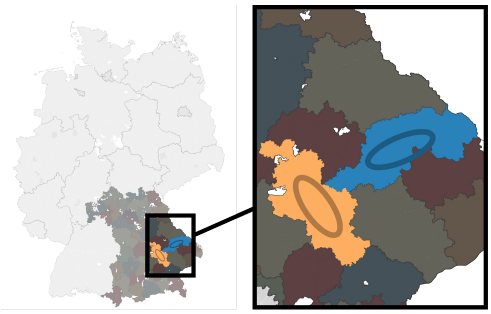

(a) Districts of 2017 and resulting ellipses for the two highlighted districts.

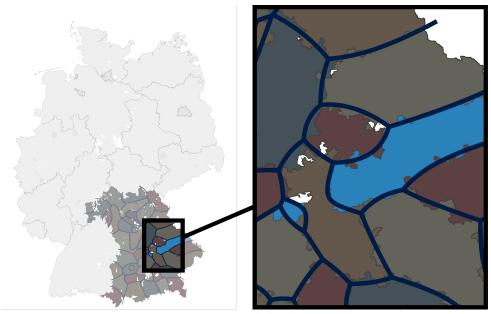

(b) Non-contiguous district in the anisotropic power diagram solution before post-processing.

Figure 7.37: Non-contiguous district resulting from non-contiguous diagram cells in the anisotropic power diagram approach for the state of Bavaria.

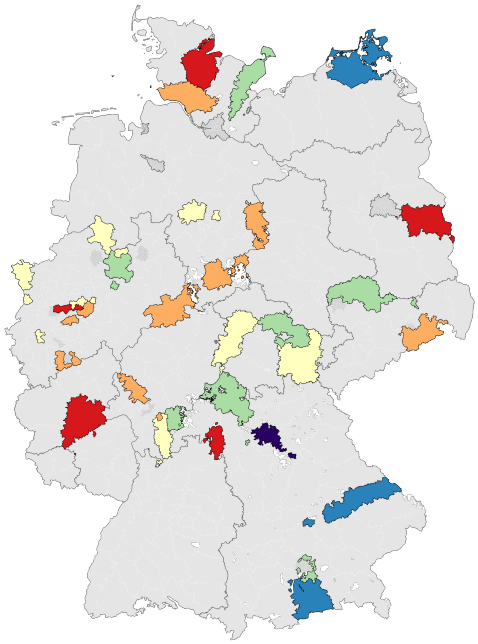

Figure 7.38: Non-contiguous districts (colored) in the anisotropic power diagram approach before post-processing. Overall, 31 districts are non-contiguous.

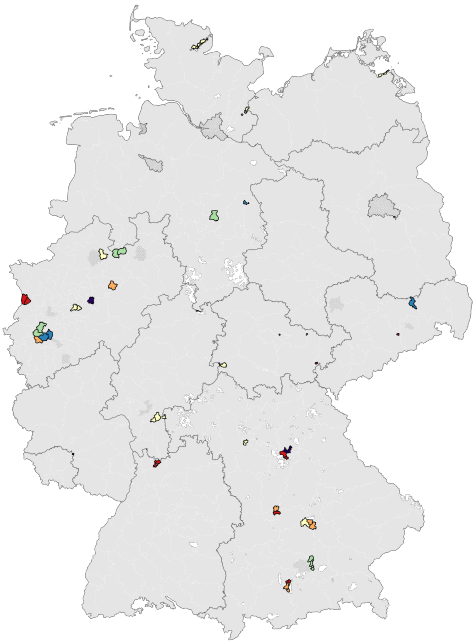

Figure 7.39: Units affected by the anisotropic power diagram post-processing (colored, 59 in total).

	Deviation (Germany)		Deviation (State)	
	∅	max	∅	max
Districts 2017	9.46%	24.89%	9.22%	24.79%
Power Diagrams	2.75%	14.37%	2.18%	13.97%
Shortest-Path Diagrams	2.01%	8.62%	1.28%	8.88%
Anisotropic Power Diagrams	3.23%	21.70%	2.68%	21.33%

Table 7.1: Absolute district size deviations relative to the overall German and individual state average district sizes, respectively. Average and maximum values are taken over all 13 state instances.

of all units were affected. Figure 7.38 depicts the non-contiguous districts of all 13 instances after the initial run. Figure 7.39 then depicts all units that were affected by the post-processing.

7.4.6 Evaluation & Comparison

In the following, we evaluate and compare the results for our three applied approaches.

First of all, let us state that all approaches finally led to contiguous district plans.

Balancing

Next, we consider how well the resulting district plans are balanced. In order to do so, we consider the absolute relative deviations of district sizes from the actual average district size. Recall that there are two different average district sizes we may refer to: First, the average district size of each instance, i. e., an average district size on state level. Second, we consider the average district size for the whole of Germany (cf. also Table B.2 for the precise values).

Table 7.1 shows the resulting deviations, both w. r. t. state as well as federal level, for all approaches summarized over all instances. Table B.3 in the appendix lists those values in more detail for all 13 instances. Figure 7.40 illustrates the resulting deviations (w. r. t. the federal level) for the whole of Germany.

In all our approaches, we can see a significant drop of the average deviation from 9.46% for the 2017 districts to values between 2.01% and 3.23% for our results (w. r. t. the federal level). Both the power diagram as well as the shortest-path diagram approach furthermore result in maximum deviations of less than 15%. Of course, this is not surprising, as for power diagrams we restricted solutions that resulted in too big deviations (cf. Section 7.4.3). The implemented local search method in the shortest-path approach even uses the resulting deviation as main objective.

The greatest deviations occur for the state of North Rhine-Westphalia. As this state has the by-far greatest average unit size, this results in greater rounding errors. Still,

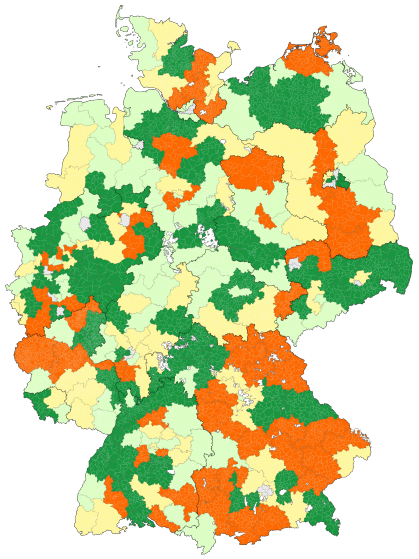

(a) Districts of 2017.

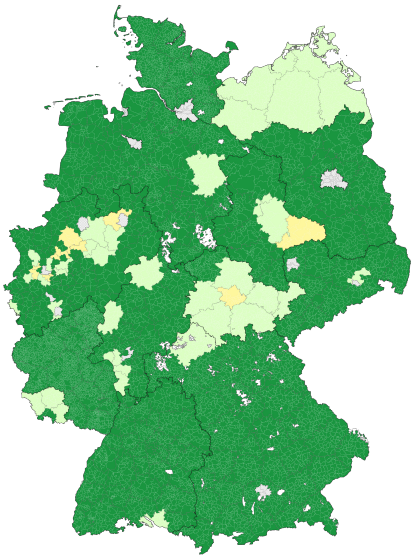

(b) Power Diagrams.

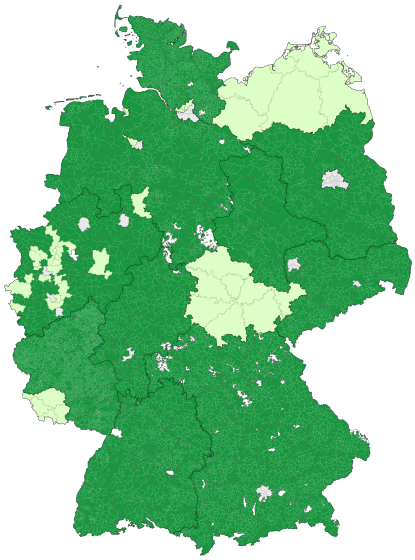

(c) Shortest-Path Diagrams.

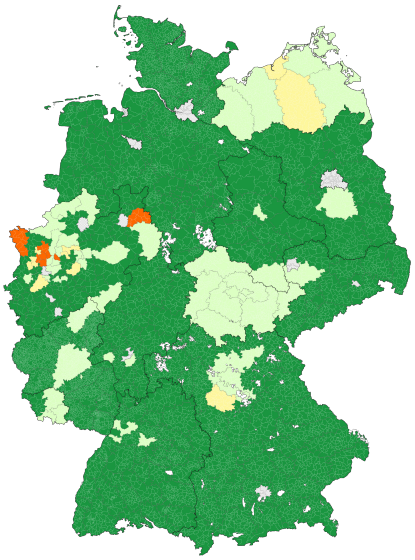

(d) Anisotropic Power Diagrams.

(e) Color legend.

Figure 7.40: Illustrations of the cluster size deviations for the different approaches as well as the original districts of 2017. See Table B.3 for the detailed results.

the local search routine in the shortest-path diagram approach results in an acceptable maximum deviation of 8.62%.

Consolidation

Table 7.2 shows the relative difference in the moment of inertia, i. e., the Balanced k-Means objective, between the three applied approaches and the districts of the 2017 elections. We denote this by ΔMoI. Of course, the power diagram approach dominates this criterion by construction. It yields an average decrease of 12.89%, while the shortest-path diagram approach still results in a decrease of 8.13% and the anisotropic power diagram approach results in an increase of 3.70%.

Table B.4 in the appendix shows the results in more detail. Here, we see that there are some exceptions in which the power diagram approach fails to yield the most-consolidated districts. While both rounding and post-processing can weaken this result, for our data this can be traced back to the rejection of solutions that would have resulted in greater balancing errors. For example, for the state of Saxony-Anhalt we observed that the final power diagram result has a moment of inertia that is 24.2% greater than the best-found but rejected solution.

	Power Diagrams	Shortest-Path Diagrams	Anisotropic Power Diagrams
ΔMoI	-12.89%	-8.13%	3.70%

Table 7.2: Relative differences in the moment of inertia (as given by Eq. (6.2)) between the clustering results and the actual districts of the 2017 elections. The values are averages over all 13 state instances weighted by their respective population sizes.

Similarity to Former Districts

Finally, let us consider the similarity of the newly created districts to the districts of 2017. We do so by measuring the ratio of separated pairs of voters as defined by Eq. (7.27). Here, both the power diagram and shortest-path diagram approach result in a ratio of 38.5% separated voter pairs (cf. Table 7.3). By contrast, only 20.0% of voter pairs are split by the anisotropic power diagram approach. Thus, we may conclude that this procedure provides districts that are more similar to the 2017 ones as intended. In the appendix, Table B.5 shows the ratios of separated pairs for all instances individually.

	Power Diagrams	Shortest-Path Diagrams	Anisotropic Power Diagrams
$\Delta(\xi^{(0)}, \xi^*)$	38.53%	38.46%	19.98%

Table 7.3: Ratio of separated pairs of voters in the different clustering approaches in comparison with the 2017 districts. The ratios are computed as defined by Eq. (7.27) for the combination of all 13 instances.

7.5 Summary & Conclusion

This chapter presented our prime application of electoral district design to demonstrate the practical relevance of our suggested constrained clustering approach.

Section 7.1 introduced the problem and listed related approaches from the literature. While, of course, the usage of generalized Voronoi diagrams for particular types is not new to this problem, our general framework offers the possibility of a targeted exploitation of the established diagram correspondence. In particular, this enables to deal with the manifold juridical and political requirements on an algorithmic approach to the electoral district design problem.

Section 7.2 then described a general framework in which we then applied the established theory to this problem. As we must consider a relaxation of binary constraints and allow fractional clusterings, this inevitably implies the usage of some rounding procedure. Therefore, we recalled the p-Split Resolution problem and in particular reviewed and adapted the algorithmic approach from [Sch01] for our setting (which yielded a slight improvement in the running time compared to [Sch01]).

Using our toy example, Section 7.3 then described the algorithmic implementations for the three diagram types, while Section 7.4 then presented and compared numerical results.

Here, we saw how the chosen diagram type determines the optimization focus of the approach. Power diagrams yielded the most consolidated districts, anisotropic power diagrams established a kind of continuity and shortest-path diagrams yielded a-priori-guaranteed contiguity and the best balancing results (as here, we optimized over the latter).

In concordance with recited opinions from the literature, we conclude that there does not exist a best-suited approach for electoral district design. However, this even more suggests to exploit the versatility of generalized Voronoi diagrams as well as their natural interpretability. Many adaptions or hybrid approaches could be thought of, too. For example, one could manipulate distances in a shortest-path approach in order to pay regard to administrative boundaries. One may use the local norms of anisotropic power diagrams in order to pay regard to the non-convex nature of state territories. One may also imply contiguity constraints in a power diagram approach and discuss to which extent this undermines the established diagram correspondence.

The example of electoral district design hence demonstrates how our proposed methodology does not provide a superior algorithm to a specific problem, but a very powerful and flexible toolbox for various settings.

Appendix A

Addendum Balanced k-Means Clustering

One can show the hardness of the Balanced k-Means problem, i. e., Proposition 6.1, also independent of the Balanced Max Centroid Variance problem by a reduction of a special case of Exact Cover By 3-sets. According to [PY82], the following problem is $\mathcal{NP}$-complete:

Problem Exact Cover by 3-Sets* (X3C*) :

Input:	Set X with $\|X\| = 3q$, family of 3-sets $S_1, \ldots, S_{3q} \in \binom{X}{3}$ such that for each $x \in X$ it holds that $\|\{S_i : x \in S_i\}\| = 3$.
Question:	Is there a subfamily $S_{i_1}, \ldots, S_{i_q}$ such that $\bigcup_{l=1}^{q} S_{i_l} = X$?

Now Proposition 6.1 can be proven via a reduction of X3C*.

Proof (Proposition 6.1). Let an instance of X3C* be given by a set $\mathcal{Y} = \{y_1, \ldots, y_{3q}\}$ and a family of 3-sets $S_1, \ldots, S_{3q} \in \binom{\mathcal{Y}}{3}$.

From this instance, we obtain an instance of Balanced k-Means for $k := 3q, d := 3q + \binom{3q}{2}$ and $n := 3q + 3q$. Next, all point weights are set to 1, i. e., $\omega_j := 1$ for $j \in [n]$. The cluster weights are set to $\kappa_1 := \ldots := \kappa_q := 4$ and $\kappa_{q+1} := \ldots := \kappa_{3q} := 1$.

We will partition the components of the Balanced k-Means instance's points into two parts: One associated with elements of $\mathcal{Y}$ and one associated with pairs of 3-sets $\{S_i, S_j\}$. For the sake of readability, we introduce the following notations: For a vector $x \in \mathbb{Q}^{3q+\binom{3q}{2}}$ denote $x = (x^{(\mathcal{Y})}, x^{(\mathcal{S})})$ for $x^{(\mathcal{Y})} \in \mathbb{Q}^{3q}$ and $x^{(\mathcal{S})} \in \mathbb{Q}^{\binom{3q}{2}}$ (and assume vectors being given in a row-representation). Next, components of $x^{\mathcal{Y}}$ and $x^{(\mathcal{S})}$ will be denoted by $x^{(\mathcal{Y})} = (x_1^{(\mathcal{Y})}, \ldots, x_{3q}^{(\mathcal{Y})})$ and $x^{(\mathcal{S})} = (x_{\{1,2\}}^{(\mathcal{S})}, \ldots, x_{\{3q-1,3q\}}^{(\mathcal{S})})$, respectively. Finally, $u^{(\mathcal{Y},i)} \in \{0,1\}^{3q}$, $i = 1, \ldots, 3q$ and $u^{(\mathcal{S},\{i,l\})} \in \{0,1\}^{\binom{3q}{2}}$, $\{i,l\} \in \binom{[3q]}{2}$ shall denote the respective standard basis vectors.

The family of points of the Balanced k-Means instance is now defined in the following way: For $j = 1, \ldots, 3q$, define

$$x^{(y_j)} = (x^{(\mathcal{Y},y_j,)}, x^{(\mathcal{S},y_j)}) := (u^{(\mathcal{Y},j)}, 0).$$

Next, for $i = 1, \ldots, 3q$, set

$$x^{(S_i)} = (x^{(\mathcal{Y},S_i)}, x^{(\mathcal{S},S_i)}) := \left(\sum_{y_j \in S_i} u^{(\mathcal{Y},j)}, \sum_{l \in [3q]\setminus\{i\}} (-1)^{l-i} M u^{(\mathcal{S},\{i,l\})} \right).$$

for some $M > 0$ to be specified later.

Due to $\omega \equiv 1$ we can assume an optimal solution of this Balanced k-Means instance to be integral (as in this case the system that defines the clustering polytope $T_{\kappa,\omega}$ is totally unimodular). Let $\xi \in \{0,1\}^{k \times m}$ be an optimal solution. This implies clusters $1, \ldots, q$ to contain exactly 4 points and clusters $q+1, \ldots, 3q$ to contain exactly 1 point.

Let $c^{(i)} = (c^{(\mathcal{Y},i)}, c^{(\mathcal{S},i)})$ denote the centroid of the ith cluster. Note that the clustering costs can be split into a $\mathcal{Y}$- and a $\mathcal{S}$-part:

$$\sum_{i=1}^{3q} \sum_{x \in \mathrm{supp}(C_i)} \left\| x - c^{(i)} \right\|^2 = \sum_{i=1}^{3q} \sum_{x \in \mathrm{supp}(C_i)} \left\| x^{(\mathcal{Y})} - c^{(\mathcal{Y},i)} \right\|^2 + \sum_{i=1}^{3q} \sum_{x \in \mathrm{supp}(C_i)} \left\| x^{(\mathcal{S})} - c^{(\mathcal{S},i)} \right\|^2$$

Let us now first show the following auxiliary claim: For $M > (24q)^{\frac{1}{2}}$, each cluster of ξ contains exactly one point $x^{(S_i)}$, i. e.,

$$\left| \mathrm{supp}(C_i) \cap \{x^{(S_1)}, \ldots, x^{(S_{3q})}\} \right| = 1, \forall\ i = 1, \ldots, 3q. \tag{A.1}$$

For the $\mathcal{Y}$-part of the points of cluster C_i, $i \in [q]$, observe that there are at most 12 components that are non-zero for at least one point in $\mathrm{supp}(C_i)$ and all values are in $\{0,1\}$. As clusters $q+1, \ldots, 3q$ consist of a single point each, they cause no costs. Thus, we have that $\sum_{i=1}^{3q} \sum_{x \in \mathrm{supp}(C_i)} \left\| x^{(\mathcal{Y})} - c^{(\mathcal{Y},i)} \right\|^2 < 12q$.

Now w.l.o.g. let $i = 1$. We will simply consider every case of

$$\rho := \left| \mathrm{supp}(C_1) \cap \{x^{(S_1)}, \ldots, x^{(S_{3q})}\} \right|.$$

The case $\rho = 0$ trivially results in $\mathcal{S}$-costs $\sum_{x \in \mathrm{supp}(C_1)} \left\| x^{(\mathcal{S})} - c^{(\mathcal{S},i)} \right\|^2 = 0$.

In the case $\rho = 1$ we can w.l.o.g. assume $\mathrm{supp}(C_1) = \{x^{(S_1)}, x^{(y_1)}, x^{(y_2)}, x^{(y_3)}\}$. It follows that $c^{(\mathcal{S},1)} = \frac{M}{4} \sum_{l \in [3q]\setminus\{1\}} \pm u^{(\mathcal{S},\{1,l\})}$. Thus

$$\begin{aligned} \sum_{x \in \mathrm{supp}(C_1)} \left\| x^{(\mathcal{S})} - c^{(\mathcal{S},i)} \right\|^2 &= \sum_{j=1}^{3} \left\| x^{(\mathcal{S},y_j)} - c^{(\mathcal{S},1)} \right\|^2 + \left\| x^{(\mathcal{S},S_1)} - c^{(\mathcal{S},1)} \right\|^2 \\ &= 3 \cdot (3q-1)(\frac{1}{4}M)^2 + (3q-1)(\frac{3}{4}M)^2 \\ &= \ldots = (3q-1)\frac{3}{4}M^2 =: \alpha. \end{aligned}$$

For the case $\rho = 2$, w.l.o.g. assume $\text{supp}(C_1) = \{x^{(S_1)}, x^{(S_2)}, x^{(y_1)}, x^{(y_2)}\}$.
It follows that $c^{(\mathcal{S},1)} = \frac{M}{4} \sum_{l \in [3q] \setminus \{1,2\}} \pm u^{(\mathcal{S},\{1,l\})} \pm u^{(\mathcal{S},\{2,l\})}$. Then,

$$\begin{aligned}
&\sum_{x \in \text{supp}(C_1)} \left\| x^{(\mathcal{S})} - c^{(\mathcal{S},i)} \right\|^2 \\
&\qquad = \sum_{j=1}^{2} \left\| x^{(\mathcal{S},y_j)} - c^{(\mathcal{S},1)} \right\|^2 \quad + \sum_{i=1}^{2} \left\| x^{(\mathcal{S},S_i)} - c^{(\mathcal{S},1)} \right\|^2 \\
&\qquad = 2 \cdot (3q-2) \cdot 2(\tfrac{1}{4}M)^2 \qquad +2 \cdot (M^2 + (3q-2)((\tfrac{1}{4}M)^2 + (\tfrac{3}{4}M)^2)) \\
&\qquad = \ldots = 2 \cdot (3q-1)\tfrac{3}{4}M^2 + \tfrac{1}{2}M^2 = 2 \cdot \alpha + \tfrac{1}{2}M^2.
\end{aligned}$$

For the case $\rho = 3$, we may w.l.o.g. assume $\text{supp}(C_1) = \{x^{(S_1)}, x^{(S_2)}, x^{(S_3)}, x^{(y_1)}\}$.
Then $c^{(\mathcal{S},1)} = \frac{M}{4} \sum_{l \in [3q] \setminus \{1,2,3\}} \pm u^{\mathcal{S}}_{\{1,l\}} \pm u^{\mathcal{S}}_{\{2,l\}} \pm u^{\mathcal{S}}_{\{3,l\}}$ and

$$\begin{aligned}
&\sum_{x \in \text{supp}(C_1)} \left\| x^{(\mathcal{S})} - c^{(\mathcal{S},i)} \right\|^2 \\
&\qquad = \left\| x^{(\mathcal{S},y_1)} - c^{(\mathcal{S},1)} \right\|^2 \quad + \sum_{i=1}^{3} \left\| x^{(\mathcal{S},S_i)} - c^{(\mathcal{S},1)} \right\|^2 \\
&\qquad = (3q-3) \cdot 3(\tfrac{1}{4}M)^2 \quad +3 \cdot (2M^2 + (3q-3)(2(\tfrac{1}{4}M)^2 + (\tfrac{3}{4}M)^2)) \\
&\qquad = \ldots = 3 \cdot (3q-1)\tfrac{3}{4}M^2 + \tfrac{3}{2}M^2 = 3 \cdot \alpha + \tfrac{3}{2}M^2.
\end{aligned}$$

For $\rho = 4$, we assume w.l.o.g. $\text{supp}(C_1) = \{x^{(S_1)}, x^{(S_2)}, x^{(S_3)}, x^{(S_4)}\}$.
Then $c^{(\mathcal{S},1)} = \frac{M}{4} \sum_{l \in [3q] \setminus \{1,2,3,4\}} \pm u^{\mathcal{S}}_{\{1,l\}} \pm u^{\mathcal{S}}_{\{2,l\}} \pm u^{(\mathcal{S},\{3,l\})} \pm u^{(\mathcal{S},\{4,l\})}$ and

$$\begin{aligned}
\sum_{x \in \text{supp}(C_1)} \left\| x^{(\mathcal{S})} - c^{(\mathcal{S},i)} \right\|^2 \quad &= \sum_{i=1}^{4} \left\| x^{(\mathcal{S},S_i)} - c^{(\mathcal{S},1)} \right\|^2 \\
&= 4 \cdot (3M^2 + (3q-4)(3(\tfrac{1}{4}M)^2 + (\tfrac{3}{4}M)^2)) \\
&= \ldots = 4 \cdot (3q-1)\tfrac{3}{4}M^2 + 3M^2 = 4 \cdot \alpha + 3M^2.
\end{aligned}$$

Hence, if a clustering obeys the auxiliary claim (A.1), it causes costs of at most $12q + q \cdot \alpha$. However, as there are as many clusters as 3-sets, each cluster containing more than one x_{S_i} results in additional $\mathcal{S}$-costs of at least $\frac{1}{2}M^2$. Thus, the auxiliary claim must be fulfilled if $\frac{1}{2}M^2 > 12q$.

Choosing M accordingly and thus with (A.1) fulfilled, we have that the optimal costs are given as

$$\sum_{i=1}^{3q} \sum_{x \in \text{supp}(C_i)} \left\| x - c^{(i)} \right\|^2 = \sum_{i=1}^{q} \sum_{x \in \text{supp}(C_i)} \left\| x^{(\mathcal{Y})} - c^{(\mathcal{Y},i)} \right\|^2 + q \cdot \alpha.$$

We now again w.l.o.g. consider cluster 1 and assume

$$\text{supp}(C_1) = \{x^{(S_1)}, x^{(y_1)}, x^{(y_2)}, x^{(y_2)}\}.$$

It is now easy to check that this cluster has $\mathcal{Y}$-costs

$$\sum_{x \in \text{supp}(C_1)} \left\| x^{(\mathcal{Y})} - c^{(\mathcal{Y},1)} \right\|^2 = \frac{1}{2} \left(9 - |S_1 \cap \{y_1, y_2, y_3\}|\right)$$

Thus, a yes-instance of X3C* will result in costs $3q + q\alpha$, while a no-instance must have costs of at least $3q + q\alpha + \frac{1}{2}$. This proves the claim. □

Appendix B

Statistics and Results for the German Elections Data

State	#Districts	#Municipalities	#Eligible Voters	∅District Size	Absolute Relative Balancing Error State-Level		Germany	
					max	∅	max	∅
Baden-Württemberg	38	1 102	7 732 597	2.035E+05	20.73%	8.32%	20.01%	8.07%
Bavaria	46	2 060	9 522 371	2.070E+05	23.42%	13.05%	23.83%	13.17%
Berlin	12	12	2 503 070	2.086E+05	13.65%	6.68%	14.91%	6.49%
Brandenburg	10	417	2 051 559	2.052E+05	21.23%	14.54%	20.55%	14.49%
Bremen	2	3	474 151	2.371E+05	6.50%	6.50%	22.38%	15.40%
Hamburg	6	6	1 296 656	2.161E+05	13.96%	8.20%	18.04%	10.57%
Hesse	22	427	4 408 986	2.004E+05	21.06%	9.08%	17.60%	9.08%
Lower Saxony	30	946	6 124 582	2.042E+05	17.66%	7.96%	18.52%	8.01%
Mecklenburg-Vorpommern	6	753	1 324 614	2.208E+05	9.11%	4.20%	16.76%	7.35%
North Rhine-Westphalia	64	407	13 174 577	2.059E+05	22.18%	8.41%	22.35%	8.38%
Rhineland-Palatinate	15	2 305	3 080 588	2.054E+05	21.24%	9.86%	20.69%	9.74%
Saarland	4	52	777 264	1.943E+05	8.67%	4.77%	13.98%	5.82%
Saxony	16	424	3 329 550	2.081E+05	21.23%	6.80%	20.55%	6.99%
Saxony-Anhalt	9	218	1 854 891	2.061E+05	19.64%	8.81%	19.72%	8.76%
Schleswig-Holstein	11	1 110	2 266 012	2.060E+05	20.05%	12.23%	19.86%	12.23%
Thuringia	8	849	1 767 014	2.209E+05	10.20%	4.98%	16.68%	8.68%
Germany	299	11 091	61 688 482	2.063E+05	23.42%	9.14%	23.83%	9.42%

Table B.1: Statistics of our data for the 2017 German elections for all 16 federal states.

State	# Districts	# Units	# Eligible Voters	∅ District Size	∅ Unit Size
Baden-Württemberg	38	1 102	7 732 597	2.035E+05	7.017E+03
Bavaria	42	2 056	8 598 759	2.047E+05	4.182E+03
Brandenburg	10	417	2 051 559	2.052E+05	4.920E+03
Hesse	21	426	4 185 044	1.993E+05	9.824E+03
Mecklenburg-Vorpommern	6	753	1 324 614	2.208E+05	1.759E+03
Lower Saxony	30	946	6 124 582	2.042E+05	6.474E+03
North Rhine-Westphalia	59	402	12 015 375	2.037E+05	2.989E+04
Rhineland-Palatinate	15	2 305	3 080 588	2.054E+05	1.336E+03
Saarland	4	52	777 264	1.943E+05	1.495E+04
Saxony	13	421	2 648 542	2.037E+05	6.291E+03
Saxony-Anhalt	9	218	1 854 891	2.061E+05	8.509E+03
Schleswig-Holstein	11	1 110	2 266 012	2.060E+05	2.041E+03
Thuringia	8	849	1 767 014	2.209E+05	2.081E+03

Table B.2: Key figures for the 13 instances resulting from our dataset of the 2017 federal elections in Germany. These instances root from the 16 states (cf. Table B.1) dropping all units that are 5% larger than the average district size (w. r. t. the state average) as well as all city states, cf. Section 7.4.1.

	Districts 2017				Power Diagrams				Shortest-Path Diagrams				Anisotropic Power Diagrams			
	Dev. (State)		Dev. (Ger)		Dev. (State)		Dev. (Ger)		Dev. (State)		Dev. (Ger)		Dev. (State)		Dev. (Ger)	
	max	∅	max	∅	max	∅	max	∅	max	∅	max	∅	max	∅	max	∅
Schleswig-Holstein	20.05%	12.23%	20.88%	12.21%	4.20%	2.09%	4.45%	2.06%	0.19%	0.08%	0.84%	0.70%	2.69%	1.05%	3.39%	1.16%
Lower Saxony	17.66%	7.96%	17.83%	7.97%	5.76%	1.76%	5.52%	1.82%	2.75%	1.23%	2.95%	1.26%	5.72%	1.28%	5.48%	1.33%
North Rhine-Westphalia	22.53%	8.01%	21.89%	7.94%	13.97%	4.56%	14.37%	4.55%	8.88%	3.59%	8.62%	3.59%	21.33%	6.13%	21.70%	6.06%
Hesse	21.74%	8.92%	18.60%	9.06%	4.44%	1.76%	5.90%	2.97%	1.81%	0.90%	4.14%	2.57%	4.39%	1.22%	6.87%	2.68%
Rhineland-Palatinate	21.24%	9.86%	21.71%	9.96%	0.71%	0.25%	1.09%	0.41%	0.09%	0.05%	0.47%	0.39%	9.33%	2.00%	9.74%	2.10%
Baden-Württemberg	20.73%	8.32%	20.09%	8.20%	6.24%	1.57%	5.65%	1.69%	2.40%	1.12%	2.94%	1.22%	6.18%	1.49%	6.70%	1.61%
Bavaria	24.79%	13.48%	24.89%	13.50%	4.64%	0.96%	4.58%	0.95%	1.02%	0.36%	1.02%	0.37%	14.88%	2.06%	14.95%	2.05%
Saarland	8.67%	4.77%	13.25%	5.47%	2.87%	1.74%	7.39%	4.99%	0.42%	0.30%	5.30%	5.01%	4.05%	2.05%	7.42%	4.98%
Brandenburg	21.23%	14.54%	21.57%	14.56%	4.18%	2.01%	4.46%	1.97%	0.91%	0.47%	1.19%	0.46%	5.38%	2.42%	5.66%	2.43%
Mecklenburg-Vorpommern	9.11%	4.20%	17.75%	8.19%	0.20%	0.11%	8.07%	7.90%	0.08%	0.05%	8.00%	7.92%	11.56%	3.88%	14.22%	9.60%
Saxony	19.54%	6.26%	19.88%	6.24%	7.63%	1.99%	7.17%	2.14%	0.70%	0.45%	1.04%	0.52%	5.57%	1.36%	5.97%	1.40%
Saxony-Anhalt	19.64%	8.81%	19.04%	9.19%	12.23%	3.10%	11.59%	3.36%	0.36%	0.18%	1.07%	0.75%	5.81%	2.70%	6.58%	2.77%
Thuringia	10.20%	4.98%	17.67%	9.22%	4.78%	1.90%	12.70%	8.02%	0.09%	0.05%	8.07%	7.97%	1.34%	0.63%	9.40%	7.95%
Germany	24.79%	9.22%	24.89%	9.46%	13.97%	2.18%	14.37%	2.75%	8.88%	1.28%	8.62%	2.01%	21.33%	2.68%	21.70%	3.23%

Table B.3: Detailed relative absolute deviations of the district sizes of the several approaches. Deviations are given from both the average district size for the individual states as well as the overall German average district size.

	Districts 2017	Power Diagrams		Shortest-Path Diagrams		Anisotropic Power Diagrams	
	MoI	MoI	ΔMoI	MoI	ΔMoI	MoI	ΔMoI
Baden-Württemberg	1.292E+15	1.119E+15	-13.45%	1.141E+15	-11.74%	1.335E+15	3.33%
Bavaria	2.623E+15	2.230E+15	-14.99%	2.307E+15	-12.04%	2.847E+15	8.52%
Brandenburg	1.069E+15	9.480E+14	-11.32%	1.046E+15	-2.11%	1.096E+15	2.52%
Hesse	7.787E+14	6.719E+14	-13.71%	6.785E+14	-12.87%	7.749E+14	-0.49%
Lower Saxony	2.005E+15	1.566E+15	-21.92%	1.623E+15	-19.07%	1.972E+15	-1.68%
Mecklenburg-Vorpommern	1.010E+15	8.927E+14	-11.61%	9.063E+14	-10.27%	9.423E+14	-6.70%
North Rhine-Westphalia	1.183E+15	1.071E+15	-9.45%	1.145E+15	-3.14%	1.264E+15	6.85%
Rhineland-Palatinate	7.312E+14	6.274E+14	-14.19%	7.206E+14	-1.44%	7.424E+14	1.54%
Saarland	7.675E+13	7.377E+13	-3.88%	8.629E+13	12.43%	7.569E+13	-1.38%
Saxony	6.616E+14	5.771E+14	-12.77%	6.631E+14	0.23%	6.907E+14	4.39%
Saxony-Anhalt	7.026E+14	7.620E+14	8.45%	7.690E+14	9.45%	7.259E+14	3.31%
Schleswig-Holstein	7.113E+14	6.517E+14	-8.38%	6.501E+14	-8.61%	7.655E+14	7.61%
Thuringia	7.828E+14	6.001E+14	-23.34%	6.698E+14	-14.43%	7.723E+14	-1.34%
Germany	1.363E+15	1.170E+15	-12.89%	1.225E+15	-8.13%	1.423E+15	3.70%

Table B.4: Resulting moment of inertia (MoI) values (as defined by Eq. (6.2)) for the different approaches and instances. The ΔMoI columns yield the relative difference to the values of the 2017 election districts.

	Power Diagrams $\Delta(\xi^{(0)}, \xi^*)$	Shortest-Path Diagrams $\Delta(\xi^{(0)}, \xi^*)$	Anisotropic Power Diagrams $\Delta(\xi^{(0)}, \xi^*)$
Baden-Württemberg	40.06%	40.81%	18.69%
Bavaria	44.48%	40.75%	31.54%
Brandenburg	29.06%	35.95%	20.25%
Hesse	39.04%	39.34%	17.78%
Lower Saxony	40.94%	41.73%	16.71%
Mecklenburg-Vorpommern	24.28%	23.38%	19.69%
North Rhine-Westphalia	37.20%	36.54%	17.78%
Rhineland-Palatinate	38.81%	49.83%	21.80%
Saarland	4.50%	21.75%	10.39%
Saxony	31.69%	31.25%	14.29%
Saxony-Anhalt	37.16%	28.13%	16.79%
Schleswig-Holstein	43.42%	43.48%	19.17%
Thuringia	43.61%	37.02%	11.70%
Germany	38.53%	38.46%	19.98%

Table B.5: Ratio of separated pairs of voters in the different clustering approaches in comparison with the 2017 districts.

List of Algorithms

List of Figures

Except for those with an explicit source citation (precisely, Figs. 3.1, 3.2 and 7.1), all figures in this thesis were created by the author. The geographic map data of Germany, that was used to depict the results in Chapter 7, is ©GeoBasis-DE / BKG (2019) (see [BKG19]).

Most diagrams and clustering illustrations have been created using the wonderful *TikZ* package ([TikZ]). The three-dimensional diagram surfaces in Chapter 3 were created with the help of the excellent *matplotlib* Python module ([Hun07]) and its *mplot3d* toolkit ([Mpl3d]). Geographic illustrations were created using the powerful QGIS software ([QGIS]). The histograms in Chapter 7 were created by combining matplotlib ([Hun07]) with the *tikzplotlib* Python module ([Sch19]).

List of Tables

Notations

General Mathematical Objects

$[n] := \{1, \dots, n\}$	Set of numbers 1 to n.
$\binom{S}{k} := \{U \subset S : \lvert U\rvert = k\}$	Set of all subsets of S of cardinality k.
$\mathrm{lin}(A)$ for $A \in \mathbb{R}^{n\times m}$	Subspace generated by the columns of A.
$u^{(i)}$	The ith standard unit vector (with the dimension implied by the context).
$\mathrm{Id}_d \in \mathbb{R}^{d\times d}$	Identity matrix in dimension d.
$\mathbb{R}^{d\times d}_{\mathrm{sym}}$	Symmetric matrices in $\mathbb{R}^{d\times d}$.
$\mathbb{R}^{d\times d}_{\mathrm{sym},>0}$	Positive definite symmetric matrices in $\mathbb{R}^{d\times d}$.
$\mathbb{N}$, $\mathbb{N}_{>0}$	Non-negative integers (including 0) and strictly positive integers, respectively.
$\mathbb{R}_{>0}$, $\mathbb{R}_{\geq 0}$	Strictly positive / Non-negative reals, respectively.
π_S	Orthogonal projection to the coordinates indexed by $S \subset [d]$ (with d being the original dimension), i. e., $\pi_S(x) = \sum_{i\in S} x^\intercal u^{(i)}$ for $x \in \mathbb{R}^d$.
$\mathbb{B}^d_2$	Euclidean unit ball in $\mathbb{R}^d$.
$\mathbb{S}^{d-1}$	Euclidean unit sphere in $\mathbb{R}^d$.
$\lVert\cdot\rVert_2$	Euclidean norm.
$\lVert\cdot\rVert_\Sigma$	Ellipsoidal norm w. r. t. $\Sigma \in \mathbb{R}^{d\times d}$ pos. def., defined via $\lVert x\rVert_\Sigma = \sqrt{x^\intercal \Sigma x}$ (see Section 3.3.3).
$\mathcal{SO}(d)$	Special orthogonal group in dimension d.
$H_{(a,\beta)}$, $H^{\leq}_{(a,\beta)}$, $H^{<}_{(a,\beta)}$	Hyperplane, closed, and open halfspace, respectively, with outer normal a and level β (see Section 1.3).
$F(C, d)$	Face of a convex set $C \subset \mathbb{R}^d$ w. r. t. an outer normal $d \in \mathbb{R}^d$ (see Eq. (6.17)).

Vector and Matrix Indexing

$\mathbb{1}^{(d)}$	Vector $(1,\dots,1)^\intercal \in \mathbb{R}^d$.
$0^{(d)}$	Vector $(0,\dots,0)^\intercal \in \mathbb{R}^d$.
$\mathbb{1}_S^{(d)}$ for $S \subset [d]$	Indicator vector of indices in S, i. e. $\mathbb{1}_S^{(d)} := \sum_{i\in S} u^{(i)}$.
$\mathbb{1}_A$ for $A \subset \mathbb{R}^d$	Indicator function of the a set $A \subset \mathbb{R}^d$, i. e., $\mathbb{1}_A(x) = 1 \Leftrightarrow x \in A$.
$A_{i,\cdot}$	The ith row of matrix A.
$A_{\cdot,j}$	The jth column of matrix A.
$A_{\cdot,S}$	Submatrix of A of all columns in S.
$\operatorname{diag}(v)$	Diagonal matrix in $\mathbb{R}^{d\times d}$ with diagonal entries given by $v \in \mathbb{R}^d$.

Constrained Clusterings

$\xi = (C_1,\dots,C_k)^\intercal \in \mathbb{R}^{k\times n}$	Clustering (cf. Section 2.1)
$T\left((A_i)_{i\in[k]}, (b^{(i)})_{i\in[k]}\right)$	Clustering polytope with constraint matrices $(A_i)_{i\in[k]}$ and right-hand sides $(b^{(i)})_{i\in[k]}$ (Definition 2.1 see also explanations in Section 2.1)
$T_{\text{uncstr.},k,n}$	Unconstrained clustering polytope, see Eq. (2.2).
$T_{\kappa,\omega}$	Weight-balanced clustering polytope, see Eq. (4.16).
$\operatorname{supp}(C_i)$	Support of the ith cluster, see Eq. (2.1).
$\omega(C_i)$	Weight of the ith cluster, i. e., $\sum_{j\in[n]} \xi_{i,j}\omega_j$, assuming weights $(\omega_j)_{j\in[n]}$ to be given, see Eq. (4.17).
$G(\xi)$	Clustering graph, see Definition 2.10.
$G_{\mathcal{C}}(\xi)$	Contracted clustering graph, , see Definition 2.11.

Generalized Voronoi Diagrams

$\mathcal{P} = (P_1,\dots,P_k)$	Notation for a generalized Voronoi diagram (as used throughout this thesis), see Definition 3.1).
$B_{i,l}$	Bisector of the cells i and l (see Definition 3.1).
$H_{i,l}$	Dominance region of cell i over l (see Definition 3.1).
$\operatorname{PD}\left(\left(s^{(i)},\mu_i\right)_{i\in[k]}\right)$	Power diagram in spherical representation (cf. Section 3.3.2).
$\operatorname{PD}_{\text{aff}}\left((a_i,\alpha_i)_{i\in[k]}\right)$	Power diagram in affine representation (cf. Section 3.3.2).
$\operatorname{APD}\left(\left(\Sigma_i, s^{(i)},\mu_i\right)_{i\in[k]}\right)$	Anisotropic power diagram in ellipsoidal representation (cf. Section 3.3.3).
$\operatorname{APD}_{\text{quad}}\left((A_i,a_i,\alpha_i)_{i\in[k]}\right)$	Anisotropic power diagram in quadratic representation (cf. Section 3.3.3).

Bibliography

[AAS97] P. K. Agarwal, B. Aronov, and M. Sharir. "Computing Envelopes in Four Dimensions with Applications". In: *SIAM Journal on Computing* 26.6 (1997), pp. 1714–1732. DOI: 10.1137/S0097539794265724.

[AS00] P. K. Agarwal and M. Sharir. "Arrangements and Their Applications". In: *Handbook of Computational Geometry*. Elsevier, 2000, pp. 49–119. DOI: 10.1016/B978-044482537-7/50003-6.

[Alp+15] A. Alpers, A. Brieden, P. Gritzmann, A. Lyckegaard, and H. F. Poulsen. "Generalized balanced power diagrams for 3D representations of polycrystals". In: *Philosophical Magazine* 95.9 (2015), pp. 1016–1028. DOI: 10.1080/14786435.2015.1015469.

[Alt98] M. Altman. "Modeling the effect of mandatory district compactness on partisan gerrymanders". In: *Political Geography* 17.8 (1998), pp. 989–1012. ISSN: 0962-6298. DOI: 10.1016/S0962-6298(98)00015-8.

[Alt97] M. Altman. "The computational complexity of automated redistricting : Is automation the answer ?" In: *Rutgers computer & technology law journal* 23 (1997), pp. 81–142.

[AN87] E. J. Anderson and P. Nash. *Linear Programming in Infinite-Dimensional Spaces: Theory and Applications*. Wiley-Interscience Series in Discrete Mathematics and Optimization Series. Wiley, 1987. ISBN: 978-0-608-00992-6.

[ACK09] B. Aronov, P. Carmi, and M. J. Katz. "Minimum-Cost Load-Balancing Partitions". In: *Algorithmica* 54.3 (July 2009), pp. 318–336. ISSN: 0178-4617. DOI: 10.1007/s00453-007-9125-3.

[AB86] P. F. Ash and E. D. Bolker. "Generalized Dirichlet tessellations". In: *Geometriae Dedicata* 20.2 (Apr. 1986), pp. 209–243. ISSN: 0046-5755. DOI: 10.1007/BF00164401.

[Aud03] M. Audin. *Geometry*. Universitext. Springer Berlin Heidelberg, 2003. ISBN: 978-3-540-43498-6. DOI: 10.1007/978-3-642-56127-6.

[Aur87a] F. Aurenhammer. "Power Diagrams: Properties, Algorithms and Applications". In: *SIAM Journal on Computing* 16.1 (Feb. 1987), pp. 78–96. ISSN: 0097-5397. DOI: 10.1137/0216006.

[AHA98] F. Aurenhammer, F. Hoffmann, and B. Aronov. "Minkowski-Type Theorems and Least-Squares Clustering". In: *Algorithmica* 20.1 (Jan. 1998), pp. 61–76. ISSN: 1432-0541. DOI: 10.1007/PL00009187.

[Aur87b] F. Aurenhammer. "A criterion for the affine equivalence of cell complexes in R^d and convex polyhedra in R^{d+1}". In: *Discrete & Computational Geometry* 2.1 (Mar. 1987), pp. 49–64. ISSN: 0179-5376. DOI: 10.1007/BF02187870.

[Aur91] F. Aurenhammer. "Voronoi diagrams – a survey of a fundamental geometric data structure". In: *ACM Computing Surveys (CSUR)* 23.3 (Sept. 1991), pp. 345–405. ISSN: 0360-0300. DOI: 10.1145/116873.116880.

[AK00] F. Aurenhammer and R. Klein. "Voronoi Diagrams". In: *Handbook of Computational Geometry*. Ed. by J. Sack and G. Urrutia. North-Holland, 2000, pp. 201–290. ISBN: 978-0-444-82537-7. DOI: 10.1016/B978-044482537-7/50006-1.

[AKL13] F. Aurenhammer, R. Klein, and D.-T. Lee. *Voronoi Diagrams and Delaunay Triangulations*. World Scientific, 2013. ISBN: 978-9814447638. DOI: 10.1142/8685.

[Aus+99] G. Ausiello, A. Marchetti-Spaccamela, P. Crescenzi, G. Gambosi, M. Protasi, and V. Kann. *Complexity and Approximation: Combinatorial Optimization Problems and Their Approximability Properties*. Springer Berlin Heidelberg, 1999. ISBN: 978-3-642-63581-6. DOI: 10.1007/978-3-642-58412-1.

[Awa+15] P. Awasthi, M. Charikar, R. Krishnaswamy, and A. K. Sinop. "The Hardness of Approximation of Euclidean k-Means". In: *31st International Symposium on Computational Geometry (SoCG 2015)*. Ed. by L. Arge and J. Pach. Vol. 34. Leibniz International Proceedings in Informatics (LIPIcs). Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2015, pp. 754–767. ISBN: 978-3-939897-83-5. DOI: 10.4230/LIPIcs.SOCG.2015.754.

[BLP05] F. Bação, V. Lobo, and M. Painho. "Applying genetic algorithms to zone design". In: *Soft Computing* 9 (2005), pp. 341–348. ISSN: 1432-7643. DOI: 10.1007/s00500-004-0413-4.

[Bar74] G. W. Barlow. "Hexagonal territories". In: *Animal Behaviour* 22 (Nov. 1974), 876–IN1. ISSN: 0003-3472. DOI: 10.1016/0003-3472(74)90010-4.

[BHR92] E. R. Barnes, A. J. Hoffman, and U. G. Rothblum. "Optimal partitions having disjoint convex and conic hulls". In: *Mathematical Programming* 54.1-3 (Feb. 1992), pp. 69–86. ISSN: 0025-5610. DOI: 10.1007/BF01586042.

[BDW08] S. Basu, I. Davidson, and K. Wagstaff. *Constrained Clustering: Advances in Algorithms, Theory, and Applications*. Chapman and Hall/CRC, Aug. 1, 2008. 472 pp. ISBN: 1584889969.

[Bis06] C. M. Bishop. *Pattern Recognition and Machine Learning*. Information Science and Statistics. Springer New York, Jan. 2006. ISBN: 978-0-387-31073-2.

[BFH07] Y. M. Bishop, S. E. Fienberg, and P. W. Holland. *Discrete Multivariate Analysis Theory and Practice*. Springer New York, 2007. ISBN: 978-0-387-72805-6. DOI: 10.1007/978-0-387-72806-3.

[Bod73] L. D. Bodin. "A DISTRICTING EXPERIMENT WITH A CLUSTERING ALGORITHM". In: *Annals of the New York Academy of Sciences* 219.1 (Nov. 1973), pp. 209–214. ISSN: 0077-8923. DOI: 10.1111/j.1749-6632.1973.tb41400.x.

[BD05] J.-D. Boissonnat and C. Delage. "Convex Hull and Voronoi Diagram of Additively Weighted Points". In: *Algorithms–ESA 2005*. Springer Berlin Heidelberg, 2005, pp. 367–378. ISBN: 978-3-540-29118-3. DOI: 10.1007/11561071_34.

[BK03] J.-D. Boissonnat and M. I. Karavelas. "On the Combinatorial Complexity of Euclidean Voronoi Cells and Convex Hulls of d-Dimensional Spheres". In: *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*. SODA '03. Baltimore, Maryland: Society for Industrial and Applied Mathematics, 2003, pp. 305–312. ISBN: 0-8987-1538-5.

[BWY06] J.-D. Boissonnat, C. Wormser, and M. Yvinec. "Curved Voronoi Diagrams". In: *Effective Computational Geometry for Curves and Surfaces*. Springer Berlin Heidelberg, 2006, pp. 67–116. ISBN: 978-3-540-33258-9. DOI: 10.1007/978-3-540-33259-6_2.

[Bol72] E. D. Bolker. "Transportation polytopes". In: *Journal of Combinatorial Theory, Series B* 13.3 (Dec. 1972), pp. 251–262. ISSN: 0095-8956. DOI: 10.1016/0095-8956(72)90060-3.

[BS97] B. Boots and R. South. "Modeling retail trade areas using higher-order, multiplicatively weighted voronoi diagrams". In: *Journal of Retailing* 73.4 (Dec. 1997), pp. 519–536. ISSN: 00224359. DOI: 10.1016/S0022-4359(97)90033-6.

[Bor15] S. Borgwardt. "On Soft Power Diagrams". In: *Journal of Mathematical Modelling and Algorithms in Operations Research* 14.2 (June 2015), pp. 173–196. ISSN: 2214-2487. DOI: 10.1007/s10852-014-9263-y.

[Bor10] S. Borgwardt. "A Combinatorial Optimization Approach to Constrained Clustering". PhD thesis. Technical University of Munich, 2010.

[BBG17] S. Borgwardt, A. Brieden, and P. Gritzmann. "An LP-based k-means algorithm for balancing weighted point sets". In: *European Journal of Operational Research* 263.2 (Dec. 2017), pp. 349–355. ISSN: 0377-2217. DOI: 10.1016/j.ejor.2017.04.054.

[BBG11] S. Borgwardt, A. Brieden, and P. Gritzmann. "Constrained minimum-k-Star Clustering and its application to the consolidation of farmland". In: *Operational Research* (2011). DOI: 10.1007/s12351-009-0041-y.

[BBG14] S. Borgwardt, A. Brieden, and P. Gritzmann. "Geometric Clustering for the Consolidation of Farmland and Woodland". In: *Mathematical Intelligencer* 36.2 (2014), pp. 37–44. ISSN: 0343-6993. DOI: 10.1007/s00283-014-9448-2.

[BG04] A. Brieden and P. Gritzmann. "A Quadratic Optimization Model for the Consolidation of Farmland by Means of Lend-Lease Agreements". In: *Operations Research Proceedings 2003*. Ed. by D. Ahr, R. Fahrion, M. Oswald, and G. Reinelt. Vol. 2003. Operations Research Proceedings. Springer Berlin Heidelberg, 2004, pp. 324–331. ISBN: 978-3-540-21445-8. DOI: 10.1007/978-3-642-17022-5_42.

[BG10] A. Brieden and P. Gritzmann. "On Clustering Bodies: Geometry and Polyhedral Approximation". In: *Discrete & Computational Geometry* 44.3 (2010), pp. 508–534. ISSN: 0179-5376. DOI: 10.1007/s00454-009-9226-7.

[BG12] A. Brieden and P. Gritzmann. "On Optimal Weighted Balanced Clusterings: Gravity Bodies and Power Diagrams". In: *SIAM Journal on Discrete Mathematics* 26.2 (2012), pp. 415–434. DOI: 10.1137/110832707.

[Bri+98] A. Brieden, P. Gritzmann, R. Kannan, V. Klee, L. Lovasz, and M. Simonovits. "Approximation of diameters: randomization doesn't help". In: *Proceedings 39th Annual Symposium on Foundations of Computer Science (Cat. No.98CB36280)*. IEEE Comput. Soc, 1998, pp. 244–251. DOI: 10.1109/SFCS.1998.743451.

[Bri+01] A. Brieden, P. Gritzmann, R. Kannan, V. Klee, L. Lovász, and M. Simonovits. "Deterministic and randomized polynomial-time approximation of radii". In: *Mathematika* 48.1-2 (Dec. 2001), pp. 63–105. ISSN: 0025-5793. DOI: 10.1112/S0025579300014364.

[BGK00] A. Brieden, P. Gritzmann, and V. Klee. "Inapproximability of some Geometric and Quadratic Optimization Problems". In: *Nonconvex Optimization and Its Applications*. Ed. by P. M. Pardalos. Vol. 42. Approximation and Complexity in Numerical Optimization. Springer Boston MA, 2000, pp. 96–115. ISBN: 978-1-4419-4829-8. DOI: 10.1007/978-1-4757-3145-3_7.

[BGK17] A. Brieden, P. Gritzmann, and F. Klemm. “Constrained clustering via diagrams: A unified theory and its application to electoral district design”. In: *European Journal of Operational Research* 263.1 (2017), pp. 18–34. ISSN: 0377-2217. DOI: 10.1016/j.ejor.2017.04.018.

[Bri95] J. Brimberg. “The Fermat—Weber location problem revisited”. In: *Mathematical Programming* 71.1 (Nov. 1995), pp. 71–76. ISSN: 0025-5610. DOI: 10.1007/BF01592245.

[Bru06] R. A. Brualdi. *Combinatorial Matrix Classes*. Encyclopedia of Mathematics and its Applications. Cambridge University Press, 2006. ISBN: 978-0-511-72118-2. DOI: 10.1017/CBO9780511721182.

[BKG19] Bundesamt für Kartographie und Geodäsie. *Verwaltungsgebiete 1:250 000 mit Einwohnerzahlen (Ebenen), Stand 31.12. (VG250-EW 31.12.)* URL: https://gdz.bkg.bund.de/index.php/default/digitale-geodaten/verwaltungsgebiete.html (visited on 08/30/2019).

[Bun12] Bundesverfassungsgericht der Bundesrepublik Deutschland. *Pressemitteilung Nr. 12/2012 vom 22. Februar 2012, Beschluss vom 31. Januar 2012, 2 BvC 3/11*. 2012.

[BWG] *Bundeswahlgesetz (Federal Elections Act)*. Version as promulgated on 23 July 1993 (Federal Law Gazette I pp. 1288–1594), last amended by Article 9 of the Ordinance of 31 August 2015 (Federal Law Gazette I p. 1474). 1993.

[CG11] G. D. Canas and S. J. Gortler. “Orphan-Free Anisotropic Voronoi Diagrams”. In: *Discrete & Computational Geometry* 46.3 (Oct. 2011), pp. 526–541. ISSN: 0179-5376. DOI: 10.1007/s00454-011-9372-6.

[CRW91] V. Capoyleas, G. Rote, and G. Woeginger. “Geometric clusterings”. In: *Journal of Algorithms* 12.2 (June 1991), pp. 341–356. ISSN: 0196-6774. DOI: 10.1016/0196-6774(91)90007-L.

[CCD16] J. G. Carlsson, E. Carlsson, and R. Devulapalli. “Shadow Prices in Territory Division”. In: *Networks and Spatial Economics* 16.3 (Sept. 2016), pp. 893–931. ISSN: 1566-113X. DOI: 10.1007/s11067-015-9303-9.

[CCE85] B. D. Causey, L. H. Cox, and L. R. Ernst. “Applications of Transportation Theory to Statistical Problems”. In: *Journal of the American Statistical Association* 80.392 (Dec. 1985), pp. 903–909. ISSN: 0162-1459. DOI: 10.1080/01621459.1985.10478201.

[CHR06] F.-H. H. Chang, F. K. Hwang, and U. G. Rothblum. “The Mean-Partition Problem”. In: *Journal of Global Optimization* 36.1 (Sept. 2006), pp. 21–31. ISSN: 0925-5001. DOI: 10.1007/s10898-006-9025-0.

[CSW07] F. Chataigner, L. R. Salgado, and Y. Wakabayashi. "Approximation and Inapproximability Results on Balanced Connected Partitions of Graphs". In: *Discrete Mathematics and Theoretical Computer Science* 9 (2007), pp. 177–192.

[CD85] L. P. Chew and R. L. Dyrsdale. "Voronoi diagrams based on convex distance functions". In: *Proceedings of the First Annual Symposium on Computational Geometry*. Vol. 75. SCG '85. ACM Press, 1985, pp. 235–244. ISBN: 0-8979-1163-6. DOI: 10.1145/323233.323264.

[Cho+99] S. D. Chowdhury, G. T. Duncan, R. Krishnan, S. F. Roehrig, and S. Mukherjee. "Disclosure Detection in Multivariate Categorical Databases: Auditing Confidentiality Protection Through Two New Matrix Operators". In: *Management Science* 45.12 (Dec. 1999), pp. 1710–1723. DOI: 10.1287/mnsc.45.12.1710.

[Coh13] D. L. Cohn. *Measure Theory*. Birkhäuser Advanced Texts Basler Lehrbücher. Springer New York, 2013. ISBN: 978-1-4614-6955-1. DOI: 10.1007/978-1-4614-6956-8.

[CR72a] H. W. Corley and S. D. Roberts. "A Partitioning Problem with Applications in Regional Design". In: *Operations Research* 20.5 (Oct. 1972), pp. 1010–1019. ISSN: 0030-364X. DOI: 10.1287/opre.20.5.1010.

[CR72b] H. W. Corley and S. D. Roberts. "Duality Relationships for a Partitioning Problem". In: *SIAM Journal on Applied Mathematics* 23.4 (Dec. 1972), pp. 490–494. ISSN: 0036-1399. DOI: 10.1137/0123052.

[CLO15] D. A. Cox, J. Little, and D. O'Shea. *Ideals, Varieties, and Algorithms*. Undergraduate Texts in Mathematics. Springer Cham, 2015. ISBN: 978-3-319-16720-6. DOI: 10.1007/978-3-319-16721-3.

[DCB01] A. DasGupta, T. T. Cai, and L. D. Brown. "Interval Estimation for a Binomial Proportion". In: *Statistical Science* 16.2 (May 2001), pp. 101–133. ISSN: 0883-4237. DOI: 10.1214/ss/1009213286.

[DK14] J. A. De Loera and E. D. Kim. "Combinatorics and Geometry of Transportation Polytopes: An Update". In: *AMS Special Session on Discrete Geometry and Algebraic Combinatorics*. Vol. 625. Contemporary Mathematics. American Mathematical Society, 2014, pp. 37–76. DOI: 10.1090/conm/625/12491.

[De +90] C. De Simone, M. Lucertini, S. Pallottino, and B. Simeone. "Fair dissections of spiders, worms, and caterpillars". In: *Networks* 20.3 (May 1990), pp. 323–344. ISSN: 1097-0037. DOI: 10.1002/net.3230200305.

[DS12] M. H. DeGroot and M. J. Schervish. *Probability and Statistics*. 4th Edition. Addison-Wesley, 2012. ISBN: 978-0-321-50046-5.

[Bwl16] Der Bundeswahlleiter, Statistisches Bundesamt, Wiesbaden 2016. *Wahlkreise und zugeordnete Gemeinden*. URL: https://www.bundeswahlleiter.de/bundestagswahlen/2017/wahlkreiseinteilung/downloads.html (visited on 08/01/2019).

[DeBu] Deutscher Bundestag. *Überhangmandate*. URL: https://www.bundestag.de/services/glossar/glossar/U/ueberh_mandate-245552 (visited on 07/19/2019).

[Die17] R. Diestel. *Graph Theory*. Fifth Edition. Vol. 173. Graduate Texts in Mathematics. Springer Berlin Heidelberg, 2017. ISBN: 978-3-662-53621-6. DOI: 10.1007/978-3-662-53622-3.

[DHS00] R. O. Duda, P. E. Hart, and D. G. Stork. *Pattern Classification*. John Wiley & Sons, 2000. ISBN: 978-0-471-05669-0.

[Duf56] R. Duffin. "Infinite Programs". In: *Linear Inequalities and Related Systems*. Ed. by H. W. Kuhn and A. W. Tucker. Vol. 38. Annals of Mathematics Studies. Princeton University Press, 1956, pp. 157–170. ISBN: 978-0-691-07999-8.

[DF85] M. Dyer and A. Frieze. "On the complexity of partitioning graphs into connected subgraphs". In: *Discrete Applied Mathematics* 10.2 (Feb. 1985), pp. 139–153. ISSN: 0166-218X. DOI: 10.1016/0166-218X(85)90008-3.

[EOS86] H. Edelsbrunner, J. O'Rourke, and R. Seidel. "Constructing Arrangements of Lines and Hyperplanes with Applications". In: *SIAM Journal on Computing* 15.2 (May 1986), pp. 341–363. ISSN: 0097-5397. DOI: 10.1137/0215024.

[ES86] H. Edelsbrunner and R. Seidel. "Voronoi diagrams and arrangements". In: *Discrete & Computational Geometry* 1.1 (Mar. 1986), pp. 25–44. ISSN: 0179-5376. DOI: 10.1007/BF02187681.

[EG02] S. Eisel and J. Graf. *Bundestagswahl 2002 – Die umstrittenen Wahlkreise*. Arbeitspapier 55. Konrad-Adenauer-Stiftung, Rathausallee 12, 53757 Sankt Augustin: Konrad-Adenauer-Stiftung e.V., Jan. 2002.

[Eis+01] S. Eisel, M. Schmitz, E. Giebel-Felten, R. T. Baus, and V. Neu. *Bundestagswahlen 2002 – Kandidatenentwicklung und Personalwechsel*. Arbeitspapier 50. Konrad-Adenauer-Stiftung, Rathausallee 12, 53757 Sankt Augustin: Konrad-Adenauer-Stiftung e.V., Nov. 2001.

[FICO] Fair Isaac Corporation. *FICO® Xpress Solver, Version 8.6*. URL: https://www.fico.com/de/products/fico-xpress-solver (visited on 12/01/2019).

[Fan65] K. Fan. "A generalization of the Alaoglu-Bourbaki theorem and its applications". In: *Mathematische Zeitschrift* 88.1 (Dec. 1965), pp. 48–60. ISSN: 0025-5874. DOI: 10.1007/BF01112692.

[For87] S. Fortune. "A sweepline algorithm for Voronoi diagrams". In: *Algorithmica* 2.1-4 (Nov. 1987), pp. 153–174. ISSN: 0178-4617. DOI: 10.1007/BF01840357.

[For95] S. Fortune. "Voronoi Diagrams and Delaunay Triangulations". In: *Computing in Euclidean Geometry*. Ed. by D.-Z. Du and F. Hwang. 2nd Edition. Vol. 4. Lecture Notes Series on Computing. World Scientific, Jan. 1995, pp. 225–265. DOI: 10.1142/9789812831699_0007.

[For04] S. Fortune. "Voronoi diagrams and Delaunay triangulations". In: *Handbook of Discrete and Computational Geometry*. Ed. by J. E. Goodman and J. O'Rourke. 2nd edition. Discrete Mathematics and its Applications. Chapman and Hall, 2004, pp. 513–528. ISBN: 978-1-58488-301-2.

[FH11] R. G. Fryer and R. Holden. "Measuring the Compactness of Political Districting Plans". In: *Journal of Law and Economics* 54.3 (Aug. 2011), pp. 493–535. ISSN: 0022-2186. DOI: 10.1086/661511.

[GMW07] G. Gan, C. Ma, and J. Wu. *Data Clustering: Theory, Algorithms, and Applications*. Society for Industrial and Applied Mathematics, July 11, 2007. 466 pp. ISBN: 978-0-89871-623-8. DOI: 10.1137/1.9780898718348. eprint: https://epubs.siam.org/doi/pdf/10.1137/1.9780898718348.

[GJ79] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. New York: W. H. Freeman & Co, 1979. ISBN: 0-7167-1044-7.

[GN70] R. S. Garfinkel and G. L. Nemhauser. "Optimal Political Districting by Implicit Enumeration Techniques". In: *Management Science* 16.8 (Apr. 1970), B–495–B–508. ISSN: 0025-1909. DOI: 10.1287/mnsc.16.8.B495.

[Gei+14] D. Geiß, R. Klein, R. Penninger, and G. Rote. "Reprint of: Optimally solving a transportation problem using Voronoi diagrams". In: *Computational Geometry* 47.3, Part B (Apr. 2014), pp. 499–506. ISSN: 0925-7721. DOI: 10.1016/j.comgeo.2013.11.003.

[GLW97] J. A. George, B. W. Lamar, and C. A. Wallace. "Political district determination using large-scale network optimization". In: *Socio-Economic Planning Sciences* 31.1 (Mar. 1997), pp. 11–28. DOI: 10.1016/S0038-0121(96)00016-X.

[God14] S. Goderbauer. "Optimierte Einteilung der Wahlkreise für die Deutsche Bundestagswahl". In: *OR News* 52 (Nov. 2014), pp. 19–21.

[God16] S. Goderbauer. *Mathematische Optimierung der Wahlkreiseinteilung für die Deutsche Bundestagswahl*. Springer Wiesbaden, 2016. ISBN: 978-3-658-15048-8. DOI: 10.1007/978-3-658-15049-5.

[GLS93] M. Grötschel, L. Lovász, and A. Schrijver. *Geometric Algorithms and Combinatorial Optimization*. Vol. 2. Algorithms and Combinatorics. Springer Berlin Heidelberg, 1993. ISBN: 978-3-642-78242-8. DOI: 10.1007/978-3-642-78240-4.

[Grü03] B. Grünbaum. *Convex Polytopes*. Ed. by V. Kaibel, V. Klee, and G. M. Ziegler. Vol. 221. Graduate Texts in Mathematics. Springer New York, 2003. ISBN: 978-0-387-40409-7. DOI: 10.1007/978-1-4613-0019-9.

[Guro] Gurobi. *Gurobi Optimizer, Version 8.1*. URL: https://www.gurobi.com/ (visited on 12/01/2019).

[Haa98] K. Haag. *Die Mundarten des oberen Neckar- und Donaulandes (Schwäbisch-alemannisches grenzgebiet: Baarmundarten)*. Reutlingen: Buchdruckerei E. Hutzler, 1898.

[Haa29] K. Haag. "Sprachwandel im Lichte der Mundartgrenzen". In: *Teuthonista* Jahrg. 6.H. 1 (1929), pp. 1–35.

[HS94] D. Halperin and M. Sharir. "New bounds for lower envelopes in three dimensions, with applications to visibility in terrains". In: *Discrete & Computational Geometry* 12 (Sept. 1994), pp. 313–326. ISSN: 0179-5376. DOI: 10.1007/BF02574383.

[HW82] H. Händler and H. E. Wiegand. "Das Konzept der Isoglosse: methodische und terminologische Probleme". In: *Dialektologie, Part 1*. Ed. by W. Besch, U. Knoop, W. Putschke, and H. E. Wiegand. Walter de Gruyter, Jan. 1982, pp. 501–527. ISBN: 978-3-110-20338-7. DOI: 10.1515/9783110059779.1.4.501.

[HTF09] T. Hastie, R. Tibshirani, and J. Friedman. *The Elements of Statistical Learning*. Second Edition. Springer Series in Statistics. New York: Springer New York, Feb. 1, 2009. 745 pp. ISBN: 978-0-387-84857-0. DOI: 10.1007/978-0-387-84858-7.

[Hes+65] S. W. Hess, J. B. Weaver, H. J. Siegfeldt, J. N. Whelan, and P. A. Zitlau. "Nonpartisan Political Redistricting by Computer". In: *Operations Research* 13.6 (Dec. 1965), pp. 998–1006. ISSN: 0030-364X. DOI: 10.1287/opre.13.6.998.

[HK93] R. Hettich and K. O. Kortanek. "Semi-Infinite Programming: Theory, Methods, and Applications". In: *SIAM Review* 35.3 (Sept. 1993), pp. 380–429. ISSN: 0036-1445. DOI: 10.1137/1035089.

[HL93] J.-B. Hiriart-Urruty and C. Lemaréchal. *Convex Analysis and Minimization Algorithms I*. Vol. 305. Grundlehren der mathematischen Wissenschaften. Springer Berlin Heidelberg, 1993. ISBN: 978-3-642-08161-3. DOI: 10.1007/978-3-662-02796-7.

[Hit41] F. L. Hitchcock. “The Distribution of a Product from Several Sources to Numerous Localities”. In: *Journal of Mathematics and Physics* 20.1-4 (Apr. 1941), pp. 224–230. DOI: 10.1002/sapm1941201224.

[Hoj96] M. Hojati. “Optimal political districting”. In: *Computers & Operations Research* 23.12 (Dec. 1996), pp. 1147–1161. ISSN: 0305-0548. DOI: 10.1016/S0305-0548(96)00029-9.

[HHV93] D. L. Horn, C. R. Hampton, and A. J. Vandenberg. “Practical application of district compactness”. In: *Political Geography* 12.2 (Mar. 1993), pp. 103–120. ISSN: 0962-6298. DOI: 10.1016/0962-6298(93)90031-2.

[Hun07] J. D. Hunter. “Matplotlib: A 2D graphics environment”. In: *Computing in Science & Engineering* 9.3 (2007), pp. 90–95. DOI: 10.1109/MCSE.2007.55.

[Mpl3d] J. Hunter, D. Dale, E. Firing, M. Droettboom, and the Matplotlib development team. *Matplotlib mplot3d Toolkit*. URL: https://matplotlib.org/3.1.1/tutorials/toolkits/mplot3d.html (visited on 02/10/2020).

[Hur14] L. Hurwicz. “Programming in Linear Spaces”. In: *Traces and Emergence of Nonlinear Programming*. Ed. by G. Giorgi and T. Kjeldsen. Birkhäuser, Basel, 2014, pp. 131–195. ISBN: 978-3-0348-0438-7. DOI: 10.1007/978-3-0348-0439-4_8.

[HKS93] D. P. Huttenlocher, K. Kedem, and M. Sharir. “The upper envelope of voronoi surfaces and its applications”. In: *Discrete & Computational Geometry* 9.3 (Mar. 1993), pp. 267–291. ISSN: 0179-5376. DOI: 10.1007/BF02189323.

[Hwa79] F. K. Hwang. “An O(n log n) Algorithm for Rectilinear Minimal Spanning Trees”. In: *Journal of the ACM* 26.2 (Apr. 1979), pp. 177–182. ISSN: 0004-5411. DOI: 10.1145/322123.322124.

[HOR98] F. K. Hwang, S. Onn, and U. G. Rothblum. “Representations and characterizations of vertices of bounded-shape partition polytopes”. In: *Linear Algebra and its Applications* 278.1-3 (July 1998), pp. 263–284. ISSN: 0024-3795. DOI: 10.1016/S0024-3795(97)10092-1.

[Ick+01] C. Icking, R. Klein, L. Ma, S. Nickel, and A. Weißler. “On bisectors for different distance functions”. In: *Discrete Applied Mathematics* 109.1-2 (Apr. 2001), pp. 139–161. ISSN: 0166-218X. DOI: 10.1016/S0166-218X(00)00238-9.

[IIM85] H. Imai, M. Iri, and K. Murota. “Voronoi diagram in the Laguerre geometry and its applications”. In: *SIAM Journal on Computing* 14.1 (Feb. 1985), pp. 93–105. ISSN: 0097-5397. DOI: 10.1137/0214006.

[IKI94] M. Inaba, N. Katoh, and H. Imai. “Applications of weighted Voronoi diagrams and randomization to variance-based k -clustering”. In: *Proceedings of the tenth annual symposium on Computational geometry - SCG '94*. ACM Press, June 1994, pp. 332–339. ISBN: 0-8979-1648-4. DOI: 10.1145/177424.178042.

[Ing] C. Ingraham, The Washington Post. *America's most gerrymandered congressional districts*. URL: https://www.washingtonpost.com/news/wonk/wp/2014/05/15/americas-most-gerrymandered-congressional-districts (visited on 07/10/2019).

[JMF99] A. K. Jain, M. N. Murty, and P. J. Flynn. “Data clustering: a review”. In: *ACM Computing Surveys* 31.3 (Sept. 1999), pp. 264–323. ISSN: 0360-0300. DOI: 10.1145/331499.331504.

[JM39] W. A. Johnson and R. F. Mehl. “Reaction kinetics in processes of nucleation and growth”. In: *Transactions of the Metallurgical Society of AIME* 135 (1939), pp. 416–442.

[Kai66] H. F. Kaiser. “An Objective Method for Establishing Legislative Districts”. In: *Midwest Journal of Political Science* 10.2 (May 1966), pp. 200–213. ISSN: 0026-3397. DOI: 10.2307/2109148.

[Kal15] J. Kalcsics. “Districting Problems”. In: *Location Science*. Ed. by G. Laporte, S. Nickel, and F. S. da Gama. Springer Cham, 2015, pp. 595–622. ISBN: 978-3-319-13110-8. DOI: 10.1007/978-3-319-13111-5_23.

[KNS05] J. Kalcsics, S. Nickel, and M. Schröder. “Towards a unified territorial design approach — Applications, algorithms and GIS integration”. In: *Top* 13.71 (June 2005), pp. 1–56. ISSN: 1134-5764. DOI: 10.1007/BF02578982.

[Kan60] L. V. Kantorovich. “Mathematical Methods of Organizing and Planning Production”. In: *Management Science* 6.4 (July 1960), pp. 366–422. DOI: 10.1287/mnsc.6.4.366.

[KA82] L. V. Kantorovich and G. P. Akilov. *Functional Analysis*. Elsevier, 1982. ISBN: 9780080230368. DOI: 10.1016/C2013-0-03044-7.

[KH79] O. Kariv and S. L. Hakimi. “An Algorithmic Approach to Network Location Problems. II: The p -Medians”. In: *SIAM Journal on Applied Mathematics* 37.3 (Dec. 1979), pp. 539–560. ISSN: 0036-1399. DOI: 10.1137/0137041.

[KR05] L. Kaufman and P. Rousseeuw. *Finding Groups in Data: An Introduction to Cluster Analysis*. Wiley Series in Probability and Statistics. John Wiley & Sons, 2005. ISBN: 978-0-470-31748-8.

[KR87] L. Kaufman and P. J. Rousseeuw. *Clustering by means of Medoids*. Tech. rep. 3. Delft University of Technology, 1987.

[KPS97] H. Kellerer, U. Pferschy, and M. G. Speranza. "An efficient approximation scheme for the subset-sum problem". In: *Algorithms and Computation*. Vol. 1350. Lecture Notes in Computer Science. Springer Berlin Heidelberg, 1997, pp. 394–403. ISBN: 3-5406-3890-3. DOI: 10.1007/3-540-63890-3_42.

[Ken77] K. Kendig. *Elementary Algebraic Geometry*. Vol. 44. Graduate Texts in Mathematics 2. Springer New York, 1977. ISBN: 978-1-4615-6901-5. DOI: 10.1007/978-1-4615-6899-5.

[Kin84] L. J. King. *Central place theory*. Ed. by G. I. Thrall. Scientific geography series. SAGE Publications, 1984. ISBN: 978-0-803-92324-9.

[KW68] V. Klee and C. Witzgall. "Facets and Vertices of Transportation Polytopes". In: *Mathematics of the Decision Sciences, Part 1*. Ed. by G. B. Dantzig and D. Cwinorr. American Mathematical Society, 1968, pp. 257–282.

[Kol+17] S. Kolouri, S. R. Park, M. Thorpe, D. Slepcev, and G. K. Rohde. "Optimal Mass Transport: Signal processing and machine-learning applications". In: *IEEE Signal Processing Magazine* 34.4 (July 2017), pp. 43–59. DOI: 10.1109/MSP.2017.2695801.

[Koo49] T. C. Koopmans. "Optimum Utilization of the Transportation System". In: *Econometrica* 17 (July 1949), pp. 136–146. ISSN: 0012-9682. DOI: 10.2307/1907301.

[KY82] K. O. Kortanek and M. Yamasaki. "Semi-infinite transportation problems". In: *Journal of Mathematical Analysis and Applications* 88.2 (Aug. 1982), pp. 555–565. ISSN: 0022-247X. DOI: 10.1016/0022-247X(82)90214-1.

[KV12] B. Korte and J. Vygen. *Combinatorial Optimization: Theory and Algorithms*. Fifth Edition. Vol. 21. Algorithms and Combinatorics. Springer Berlin Heidelberg, 2012. ISBN: 978-3-642-24487-2. DOI: 10.1007/978-3-642-24488-9.

[Kre61] K. S. Kretschmer. "Programmes in Paired Spaces". In: *Canadian Journal of Mathematics* 13 (Jan. 1961), pp. 221–238. ISSN: 1496-4279. DOI: 10.4153/CJM-1961-019-2.

[LS03] F. Labelle and J. R. Shewchuk. "Anisotropic voronoi diagrams and guaranteed-quality anisotropic mesh generation". In: *Proceedings of the nineteenth conference on Computational geometry - SCG '03*. ACM Press, June 2003, pp. 191–200. ISBN: 978-1-581-13663-0. DOI: 10.1145/777819.777822.

[LD81] D. T. Lee and R. L. Drysdale, III. "Generalization of Voronoi Diagrams in the Plane". In: *SIAM Journal on Computing* 10.1 (Feb. 1981), pp. 73–87. ISSN: 0097-5397. DOI: 10.1137/0210006.

[LW80] D. T. Lee and C. K. Wong. “Voronoi Diagrams in L_1 (L_∞ Metrics with 2-Dimensional Storage Applications”. In: *SIAM Journal on Computing* 9.1 (Feb. 1980), pp. 200–211. ISSN: 0097-5397. DOI: 10.1137/0209017.

[LSW17] E. Lee, M. Schmidt, and J. Wright. “Improved and simplified inapproximability for k-means”. In: *Information Processing Letters* 120 (Apr. 2017), pp. 40–43. ISSN: 0020-0190. DOI: 10.1016/j.ipl.2016.11.009.

[Lev08] J. Levitt. “A Citizen’s Guide to Redistricting”. In: *SSRN Electronic Journal* (2008). ISSN: 1556-5068. DOI: 10.2139/ssrn.1647221.

[LGC13] W. Li, M. F. Goodchild, and R. Church. “An efficient measure of compactness for two-dimensional shapes and its application in regionalization problems”. In: *International Journal of Geographical Information Science* 27.6 (June 2013), pp. 1227–1250. ISSN: 1365-8816. DOI: 10.1080/13658816.2012.752093.

[LW07] Z. Li and Y. Wang. “A Quadratic Programming Model for Political Districting Problem”. In: *The First International Symposium on Optimization and Systems Biology (OSB’07)*. Ed. by D.-Z. Du and X.-S. Zhang. Vol. 7. Lecture Notes in Operations Research. Beijing, Aug. 2007, pp. 427–435.

[Llo82] S. Lloyd. “Least squares quantization in PCM”. In: *IEEE Transactions on Information Theory* 28.2 (Mar. 1982), pp. 129–137. ISSN: 0018-9448. DOI: 10.1109/TIT.1982.1056489.

[LH76] T. J. Lowe and A. P. Hurter. “The Generalized Market Area Problem”. In: *Management Science* 22.10 (June 1976), pp. 1105–1115. ISSN: 0025-1909. DOI: 10.1287/mnsc.22.10.1105.

[Mac67] J. MacQueen. “Some Methods for classification and Analysis of Multivariate Observations”. In: *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability*. Vol. Volume 1: Statistics. 1967, pp. 281–297.

[MNV09] M. Mahajan, P. Nimbhorkar, and K. Varadarajan. “The Planar k-Means Problem is NP-Hard”. In: *WALCOM: Algorithms and Computation*. Ed. by U. R. Das S. Vol. 5431. Springer Berlin Heidelberg, 2009, pp. 274–285. ISBN: 978-3-642-00201-4. DOI: 10.1007/978-3-642-00202-1_24.

[Mar81] P. G. Marlin. “Application of the transportation model to a large-scale “Districting” problem”. In: *Computers & Operations Research* 8.2 (Jan. 1981), pp. 83–96. DOI: 10.1016/0305-0548(81)90036-8.

[McM70] P. McMullen. “The maximum numbers of faces of a convex polytope”. In: *Mathematika* 17.2 (Dec. 1970), pp. 179–184. ISSN: 0025-5793. DOI: 10.1112/S0025579300002850.

[MS84] N. Megiddo and K. J. Supowit. "On the Complexity of Some Common Geometric Location Problems". In: *SIAM Journal on Computing* 13.1 (Feb. 1984), pp. 182–196. ISSN: 0097-5397. DOI: 10.1137/0213014.

[Meh88] K. Mehlhorn. "A faster approximation algorithm for the Steiner problem in graphs". In: *Information Processing Letters* 27.3 (Mar. 1988), pp. 125–128. ISSN: 0020-0190. DOI: 10.1016/0020-0190(88)90066-X.

[MJN98] A. Mehrotra, E. L. Johnson, and G. L. Nemhauser. "An Optimization Based Heuristic for Political Districting". In: *Management Science* 44.8 (Aug. 1998), pp. 1100–1114. ISSN: 0025-1909. DOI: 10.1287/mnsc.44.8.1100.

[Men28] K. Menger. "Untersuchungen über allgemeine Metrik". In: *Mathematische Annalen* 100 (1928), pp. 75–163. ISSN: 0025-5831.

[Mil07] S. Miller. "The problem of redistricting: the use of centroidal Voronoi diagrams to build unbiasedcongressional districts". MA thesis. Withmann College, USA, 2007.

[Mit72] P. S. Mitchell. "Optimal Selection of Police Patrol Beats". In: *The Journal of Criminal Law, Criminology, and Police Science* 63.4 (Dec. 1972), p. 577. ISSN: 00220205. DOI: 10.2307/1141814.

[Mon81] G. Monge. "Mémoire sur la théorie des déblais et des remblais". In: *Histoire de l'Académie Royale des Sciences de Paris* (1781).

[NY79] T. Nakamura and M. Yamasaki. "Sufficient conditions for duality theorems in infinite linear programming problems". In: *Hiroshima Mathematical Journal* 9.2 (1979), pp. 323–334. DOI: 10.32917/hmj/1206134890.

[Nie+90] R. G. Niemi, B. Grofman, C. Carlucci, and T. Hofeller. "Measuring Compactness and the Role of a Compactness Standard in a Test for Partisan and Racial Gerrymandering". In: *The Journal of Politics* 52.4 (Nov. 1990), pp. 1155–1181. ISSN: 0022-3816. DOI: 10.2307/2131686.

[Nig27] P. Niggli. "XXIV. Die topologische Strukturanalyse. I." In: *Zeitschrift für Kristallographie - Crystalline Materials* 65.1-6 (Jan. 1927), pp. 391–415. ISSN: 2196-7105. DOI: 10.1524/zkri.1927.65.1.391.

[Now33] W. Nowacki. "Der Begriff „Voronoischer Bereich"". In: *Zeitschrift für Kristallographie - Crystalline Materials* 85.1-6 (Jan. 1933), pp. 331–332. ISSN: 2196-7105. DOI: 10.1524/zkri.1933.85.1.331.

[Now76] W. Nowacki. "Über allgemeine Eigenschaften von Wirkungsbereichen". In: *Zeitschrift für Kristallographie - Crystalline Materials* 143.1-6 (Jan. 1976), pp. 360–386. ISSN: 2196-7105. DOI: 10.1524/zkri.1976.143.jg.360.

[Nyg88] B. Nygreen. “European assembly constituencies for wales - comparing of methods for solving a political districting problem”. In: *Mathematical Programming* 42 (Apr. 1988), pp. 159–169. ISSN: 0025-5610. DOI: 10.1007/BF01589400.

[Oka+00] A. Okabe, B. Boots, K. Sugihara, S. N. Chiu, and D. G. Kendall. *Spatial Tessellations: Concepts and Applications of Voronoi Diagrams*. Second Edition. Wiley Series in Probability and Statistics. Hoboken, NJ, USA: John Wiley & Sons, Inc., July 2000. ISBN: 978-0-470-31701-3. DOI: 10.1002/9780470317013.

[Oka+08] A. Okabe, T. Satoh, T. Furuta, A. Suzuki, and K. Okano. “Generalized network Voronoi diagrams: Concepts, computational methods, and applications”. In: *International Journal of Geographical Information Science* 22.9 (Sept. 2008), pp. 965–994. ISSN: 1365-8816. DOI: 10.1080/13658810701587891.

[Orac] Oracle. *Java Platform, Standard Edition Documentation, Version 11*. Sept. 2018.

[Osb14] M. S. Osborne. *Locally Convex Spaces*. Vol. 269. Graduate Texts in Mathematics. Springer Cham, 2014. ISBN: 978-3-319-02044-0. DOI: 10.1007/978-3-319-02045-7.

[Pap82] C. Papadimitriou. *Combinatorial Optimization: Algorithms and Complexity*. Englewood Cliffs, N.J: Prentice Hall, 1982. ISBN: 978-0-13-152462-0.

[Pap94] C. H. Papadimitriou. *Computational Complexity*. Addison-Wesley, 1994, pp. I–XV, 1–523. ISBN: 978-0-201-53082-7.

[PY82] C. H. Papadimitriou and M. Yannakakis. “The complexity of restricted spanning tree problems”. In: *Journal of the ACM* 29.2 (Apr. 1982), pp. 285–309. ISSN: 0004-5411. DOI: 10.1145/322307.322309.

[PJ09] H.-S. Park and C.-H. Jun. “A simple and fast algorithm for K-medoids clustering”. In: *Expert Systems with Applications* 36.2 (Mar. 2009), pp. 3336–3341. ISSN: 0957-4174. DOI: 10.1016/j.eswa.2008.01.039.

[PC19] G. Peyré and M. Cuturi. *Computational Optimal Transport: With Applications to Data Science*. Now Publishers, 2019. ISBN: 978-1-680-83550-2. DOI: 10.1561/2200000073.

[PP91] D. D. Polsby and R. Popper. “The Third Criterion: Compactness as a Procedural Safeguard against Partisan Gerrymandering”. In: *Yale Law & Policy Review* 9.2 (1991). DOI: 10.2139/ssrn.2936284.

[PT09] C. Puppe and A. Tasnádi. "Optimal redistricting under geographical constraints: Why "pack and crack" does not work". In: *Economics Letters* 105.1 (Oct. 2009), pp. 93–96. ISSN: 0165-1765. DOI: 10.1016/j.econlet.2009.06.008.

[Pyth] Python Software Foundation. *Python Language Reference, Version 3.7*. June 2018.

[QGIS] QGIS Development Team. *QGIS Geographic Information System*. Open Source Geospatial Foundation. 2015. URL: http://qgis.osgeo.org (visited on 12/01/2019).

[RSS13] F. Ricca, A. Scozzari, and B. Simeone. "Political Districting: from classical models to recent approaches". In: *Annals of Operations Research* 204.1 (Jan. 2013), pp. 271–299. ISSN: 0254-5330. DOI: 10.1007/s10479-012-1267-2.

[RSS08] F. Ricca, A. Scozzari, and B. Simeone. "Weighted Voronoi region algorithms for political districting". In: *Mathematical and Computer Modelling* 48.9-10 (Nov. 2008), pp. 1468–1477. ISSN: 0895-7177. DOI: 10.1016/j.mcm.2008.05.041.

[RS08] F. Ricca and B. Simeone. "Local search algorithms for political districting". In: *European Journal of Operational Research* 189.3 (Sept. 2008), pp. 1409–1426. ISSN: 0377-2217. DOI: 10.1016/j.ejor.2006.08.065.

[RF09] R. Z. Ríos-Mercado and E. Fernández. "A reactive GRASP for a commercial territory design problem with multiple balancing requirements". In: *Computers & Operations Research* 36.3 (Mar. 2009), pp. 755–776. ISSN: 0305-0548. DOI: 10.1016/j.cor.2007.10.024.

[Río20] R. Z. Ríos-Mercado, ed. *Optimal Districting and Territory Design*. Springer International Publishing, 2020. DOI: 10.1007/978-3-030-34312-5.

[Roc70] R. T. Rockafellar. *Convex Analysis*. Princeton University Press, Dec. 1970. DOI: 10.1515/9781400873173.

[Roo92] C. Roos. "Interior point approach to linear programming: theory, algorithms & parametric analysis". In: *Topics in Engineering Mathematics*. Vol. 81. Mathematics and Its Applications. Springer Netherlands, 1992, pp. 181–216. ISBN: 978-94-010-4800-2. DOI: 10.1007/978-94-011-1814-9_7.

[RF10] H. Royden and P. M. Fitzpatrick. *Real analysis*. Fourth Edition. Pearson, 2010. ISBN: 978-0-135-11355-4.

[Sai82] N. Saitô. "Asymptotic regular pattern of epidermal cells in mammalian skin". In: *Journal of Theoretical Biology* 95.3 (Apr. 1982), pp. 591–599. ISSN: 0022-5193. DOI: 10.1016/0022-5193(82)90036-4.

[Sax+17] A. Saxena, M. Prasad, A. Gupta, N. Bharill, O. P. Patel, A. Tiwari, M. J. Er, W. Ding, and C.-T. Lin. "A review of clustering techniques and developments". In: *Neurocomputing* 267 (Dec. 2017), pp. 664–681. DOI: 10.1016/j.neucom.2017.06.053.

[SW99] H. H. Schaefer and M. P. Wolff. *Topological Vector Spaces*. Vol. 3. Graduate Texts in Mathematics. Springer New York, 1999. ISBN: 9781-4612-7155-0. DOI: 10.1007/978-1-4612-1468-7.

[Sch19] N. Schlömer. *tikzplotlib v0.9.0*. Nov. 2019. DOI: 10.5281/zenodo.3839605. URL: https://github.com/nschloe/tikzplotlib.

[Sch98] A. Schrijver. *Theory of Linear and Integer Programming*. Wiley, 1998. 484 pp. ISBN: 978-0-471-98232-6.

[Sch01] M. Schröder. "Gebiete optimal aufteilen". PhD thesis. Universität Karlsruhe, 2001.

[SW77] M. Segal and D. B. Weinberger. "Turfing". In: *Operations Research* 25.3 (June 1977), pp. 367–386. ISSN: 0030-364X. DOI: 10.1287/opre.25.3.367.

[SeiBPB] C. Seils, Bundeszentrale für politische Bildung. *Das neue Wahlrecht und die Krux mit den Überhangmandaten*. URL: http://www.bpb.de/politik/wahlen/bundestagswahlen/163311/das-neue-wahlrecht (visited on 07/19/2019).

[GovD] Senatskanzlei, Geschäfts- und Koordinierungsstelle GovData, Hamburg. *https://www.govdata.de/dl-de/by-2-0*. URL: https://www.govdata.de/dl-de/by-2-0 (visited on 08/30/2019).

[SH75] M. I. Shamos and D. Hoey. "Closest-point problems". In: *16th Annual Symposium on Foundations of Computer Science (sfcs 1975)*. IEEE, Oct. 1975, pp. 151–162. DOI: 10.1109/SFCS.1975.8.

[Sha94] M. Sharir. "Almost tight upper bounds for lower envelopes in higher dimensions". In: *Discrete & Computational Geometry* 12.3 (Sept. 1994), pp. 327–345. DOI: 10.1007/BF02574384.

[Sha85] M. Sharir. "Intersection and Closest-Pair Problems for a Set of Planar Discs". In: *SIAM Journal on Computing* 14.2 (May 1985), pp. 448–468. ISSN: 0097-5397. DOI: 10.1137/0214034.

[SpOn19] Spiegel Online. *Wie die Opposition den Bundestag verkleinern will*. URL: https://www.spiegel.de/politik/deutschland/wahlrecht-fdp-gruene-und-linken-machen-vorschlag-fuer-kleineren-bundestag-a-1291153.html (visited on 10/11/2019).

[SABL19] Statistische Ämter des Bundes und der Länder. *Regionaldatenbank Deutschland*. URL: https://www.regionalstatistik.de/ (visited on 11/01/2019).

[TikZ] T. Tantau. *The TikZ and PGF Packages. Manual for version 3.1.5*. Lübeck, Dec. 2019.

[Tas11] A. Tasnádi. "The political districting problem: A survey". In: *Society and Economy* 33.3 (Dec. 2011), pp. 543–554. ISSN: 1588-9726. DOI: 10.1556/SocEc.2011.0001.

[PGIS] The PostGIS Development Group. *PostGIS 2.5.4dev Manual*. URL: https://postgis.net/ (visited on 12/01/2019).

[PSQL] The PostgreSQL Global Development Group. *PostgreSQL 11.6 Documentation*. Nov. 2019.

[Tod78] M. J. Todd. "Solving the Generalized Market Area Problem". In: *Management Science* 24.14 (Oct. 1978), pp. 1549–1554. ISSN: 0025-1909. DOI: 10.1287/mnsc.24.14.1549.

[USS64] U.S. Supreme Court. *Wesberry v. Sanders*. 376 U.S. 1, 7-8. Feb. 1964.

[Van08] R. J. Vanderbei. *Linear Programming*. Third Edition. Vol. 114. International Series in Operations Research & Management Science. Springer Boston MA, 2008. ISBN: 978-0-387-74387-5. DOI: 10.1007/978-0-387-74388-2.

[Vil09] C. Villani. *Optimal Transport: Old and New*. Vol. 338. Grundlehren der mathematischen Wissenschaften. Springer Berlin Heidelberg, 2009. ISBN: 978-3-540-71049-3. DOI: 10.1007/978-3-540-71050-9.

[War68] H. E. Warren. "Lower bounds for approximation by nonlinear manifolds". In: *Transactions of the American Mathematical Society* 133 (1968), pp. 167–178. ISSN: 0002-9947. DOI: 10.2307/1994937.

[Wik14] Wikimedia Commons. *Illinois US Congressional District 4 (since 2013)*. 2014. URL: https://commons.wikimedia.org/wiki/File:Illinois_US_Congressional_District_4_(since_2013).tif (visited on 07/10/2019).

[Wika] Wikipedia. *Gerrymandering — Wikipedia, The Free Encyclopedia*. accessed 13-January-2016. URL: https://en.wikipedia.org/wiki/Gerrymandering (visited on 07/17/2019).

[Wikb] Wikipedia. *Illinois's 4th congressional district — Wikipedia, The Free Encyclopedia*. accessed 13-January-2016. URL: https://en.wikipedia.org/wiki/Illinois%27s_4th_congressional_district (visited on 07/17/2019).

[Wikc] Wikipedia. *Webster/Sainte-Laguë method — Wikipedia, The Free Encyclopedia*. accessed 13-January-2016. URL: https://en.wikipedia.org/w/index.php?title=Webster/Sainte-Lagu%C3%AB_method&oldid=920212859 (visited on 08/23/2019).

[Wikd] Wikipedia. *World Geodetic System — Wikipedia, The Free Encyclopedia*. accessed 13-January-2016. URL: https://en.wikipedia.org/w/index.php?title=World_Geodetic_System&oldid=696220732 (visited on 07/17/2019).

[Wor08] C. Wormser. "Generalized Voronoi Diagrams and Applications". PhD thesis. Université Nice Sophia Antipolis, 2008.

[Wu12] J. Wu. "Cluster Analysis and K-means Clustering: An Introduction". In: *Advances in K-means Clustering*. Springer Theses. Springer Berlin Heidelberg, 2012, pp. 1–16. ISBN: 978-3-642-29806-6. DOI: 10.1007/978-3-642-29807-3_1.

[Yam68] M. Yamasaki. "Duality theorems in mathematical programmings and their applications." In: *Journal of Science of the Hiroshima University, Series A-I (Mathematics)* 32.2 (1968), pp. 331–356. DOI: 10.32917/hmj/1206138657.

[YKK84] V. Yemelichev, M. Kovalev, and M. Kravtsov. *Polytopes, Graphs and Optimisation*. Cambridge University Press, 1984. ISBN: 978-0-521-25597-4.

[You88] H. P. Young. "Measuring the Compactness of Legislative Districts". In: *Legislative Studies Quarterly* 13.1 (Feb. 1988), pp. 105–115. ISSN: 0362-9805. DOI: 10.2307/439947.

[Zei19] Zeit Online. *Thomas Oppermann will Wahlrechtsreform noch dieses Jahr*. URL: https://www.zeit.de/politik/deutschland/2019-09/bundestag-thomas-oppermann-einigung-wahlrechtsreform-verkleinerung (visited on 09/21/2019).

[ZC05] Q. Zhang and I. Couloigner. "A New and Efficient K-Medoid Algorithm for Spatial Clustering". In: *Computational Science and Its Applications – ICCSA 2005*. Vol. 3482. Lecture Notes in Computer Science. Springer Berlin Heidelberg, 2005, pp. 181–189. ISBN: 978-3-540-25862-9. DOI: 10.1007/11424857_20.

[ZS83] A. A. Zoltners and P. Sinha. "Sales Territory Alignment: A Review and Model". In: *Management Science* 29.11 (Nov. 1983), pp. 1237–1256. ISSN: 0025-1909.