

Dependence Modeling in Ultra High Dimensions with Vine Copulas and the Graphical Lasso

Dominik Müller^{*†} Claudia Czado[†]

September 5, 2017

Abstract

To model high dimensional data, Gaussian methods are widely used since they remain tractable and yield parsimonious models by imposing strong assumptions on the data. Vine copulas are more flexible by combining arbitrary marginal distributions and (conditional) bivariate copulas. Yet, this adaptability is accompanied by sharply increasing computational effort as the dimension increases. The approach proposed in this paper overcomes this burden and makes the first step into ultra high dimensional non-Gaussian dependence modeling by using a divide-and-conquer approach. First, we apply Gaussian methods to split datasets into feasibly small subsets and second, apply parsimonious and flexible vine copulas thereon. Finally, we reconcile them into one joint model. We provide numerical results demonstrating the feasibility of our approach in moderate dimensions and showcase its ability to estimate ultra high dimensional non-Gaussian dependence models in thousands of dimensions.

1 Introduction

In many areas of scientific research but also in business applications, high dimensional problems arise. For example, if a financial institution owns stocks S_1, \dots, S_{100} and wants to calculate a *portfolio Value-at-Risk*, see e.g. McNeil et al. (2006), a joint dependence model in $d = 100$ dimensions is required. For large companies, one can easily imagine this number to increase into thousand for a single asset class and much higher when several asset classes are considered. Another active field of research in which dimensions grow rapidly is computational biology. For example in *gene expression data*, genes are measured simultaneously to make inference about dependence within the biological system or with respect to some disease, see e.g. Toh and Horimoto (2002). In this case, the number

^{*}Corresponding author

[†]Department of Mathematics, Technische Universität München, Boltzmannstraße 3, 85748 Garching, Germany. E-Mail: dominik.mueller@ma.tum.de, cczado@ma.tum.de.

of genes measured describes the dimension d and in recent applications can increase to several thousands. Similar to this, hundreds of *metabolites* in the human blood can be measured and analyzed for dependence, see e.g. Krumsiek et al. (2011), where $d = 151$. This example is also easily imaginable to be extended to thousands of different metabolites.

Among the most prominently used models to analyse such datasets are *Gaussian Graphical Models*, which are based on the assumption that the data originates from a multivariate Gaussian distribution in d dimensions. Neglecting the mean of the distribution, the problem remains to estimate a $d \times d$ covariance matrix Σ . A favourable property of the multivariate Gaussian distribution is that for $\Omega = \Sigma^{-1}$ and a zero entry of Ω , i.e. $\omega_{ij} = 0$, we have conditional independence of i and j given the rest $\{1, \dots, d\} \setminus \{i, j\}$. By drawing a graph with nodes $\{1, \dots, d\}$ and omitting an undirected edge (i, j) whenever $\omega_{ij} = 0$, we obtain a *graphical model* for the conditional independence in this distribution. There has been a considerable effort on how to estimate (sparse) inverse covariance matrices and thus, Gaussian graphical models. It started with *covariance selection* of Dempster (1972) to the current state of the art algorithm, the *graphical Lasso* (Friedman et al., 2008). The huge advantage of these methods is their computational tractability also in ultra high dimensions, i.e. several thousands of variables. The underlying assumption of Gaussianity is however very strict, by imposing also assumptions on the marginals. This has been relaxed by the so called *non-paranormal* Liu et al. (2012) where the marginal distributions must not necessarily be Gaussian. For the dependence part however, especially for financial data, the multivariate Gaussian is a too strong simplification. It does not allow for heavy tails as the Student's-t distribution or asymmetric dependence, thus more sophisticated models are required.

The main idea of the so called *pair copula construction (PCC)*, see Aas et al. (2009), is to couple $d(d-1)/2$ (conditional) bivariate distributions and d marginal distributions to obtain a joint d -dimensional distribution, called *vine copula*. The huge benefit is now that all involved distributions can be chosen arbitrarily and entirely independently from each other. However, this construction is not unique but is described by a graphical model, a *regular (or R-)vine* (Bedford and Cooke, 2001, 2002). Each possible R-vine constitutes a different construction of a d -dimensional distribution. As the number of possible R-vines grows super-exponentially in dimensions, a search for an optimal model in terms of a goodness of fit criteria such as *log-likelihood* is not feasible in any dimension, and heuristic algorithms as the one of Dißmann et al. (2013) come into place. Secondly, an R-vine model needs $d(d-1)/2$ pair copulas. Clearly, this is growing too fast in hundreds of dimensions, unspoken of thousands. Thus, this also demands sophisticated approaches in high dimensions to keep models parsimonious. The large number of applications with vine copulas in recent years shown in Aas (2016) can be easily thought to expand to hundreds and thousands of dimensions, which makes an extension of vine copulas in ultra high dimensions desirable.

Our contribution is twofold. First, we will show how well established Gaussian methods in high dimensions can be fundamental for clustering R-vines. Thus, we break a d -dimensional dependence model in multiple dependence sub-models with significantly

smaller dimensions. These sub-models are now tractable again. Secondly, within these sub-models, we use a refined algorithm to improve the accuracy of the standard search algorithm for vine copulas by Dißmann et al. (2013). Afterwards, the sub-models are recombined to obtain one joint parsimonious model in ultra high dimensions. We show that this is working well in moderate dimensions and outperforming previous methods in several hundreds of dimensions in computation time and goodness of fit. Going to ultra high dimensions, i. e. several thousands, it is to our knowledge the only feasible way to estimate vine copula models, and is actually doable in a comparably short amount of time. Finally, we will also demonstrate that non-Gaussian models give clearly a competitive edge on Gaussian models using real world financial data.

The paper is structured as follows. We will briefly introduce vine copulas in Section 2 and discuss current model selection methods. We recapitulate graphical models based on the multivariate Gaussian distribution in Section 3 where we focus on the graphical Lasso. Section 4 contains our *divide-and-conquer* approach where we split the R-vine selection into sub-problems according to a path of solutions of the graphical Lasso and solve each separately with increased accuracy. We sketch an algorithmic implementation and continue to numerical examples in Section 5. There, we show a simulation study in 85 dimensions to demonstrate the feasibility of our approach in moderate dimensions. Afterwards, we increase the dimension to over 1700 to demonstrate the high efficiency of our approach with respect to time consumption, at the same time outperforming standard methods in terms of penalized goodness of fit measures. We finally include an example in more than 2000 dimensions which demonstrates that a pure Gaussian fitting is too restrictive for real world datasets, and can be improved significantly by ultra high dimensional vine copulas.

2 Dependence Modeling with Vine Copulas

We will use the following conventions. Upper case letters X denote random variables, and lower case letters x their realizations. We use bold letters \mathbf{X} for random vectors and \mathbf{x} for the vector of realizations. Matrices M are identified by upper case letters. We write v_i for the i -th entry of the vector \mathbf{v} . Denote sub vectors of $\mathbf{x} = (x_1, \dots, x_d)^T$ by $\mathbf{x}_{D(e)} := (\mathbf{x}_j)_{j \in D(e)}$. When considering matrices, we denote $m_{i,j}$ the j -th entry in the i -th row of the matrix M . Additionally, we use the following data scales.

- (i) *x-scale*: the original scale of X_i , i.i.d., with density $f_i(x_i)$, $i = 1, \dots, d$,
- (ii) *u-scale* or *copula-scale*: $U_i = F_i(X_i)$, F_i the cdf of X_i and $U_i \sim \mathcal{U}[0, 1]$, $i = 1, \dots, d$,
- (iii) *z-scale*: $Z_i = \Phi^{-1}(U_i)$, Φ the cdf of $\mathcal{N}(0, 1)$ thus $Z_i \sim \mathcal{N}(0, 1)$, $i = 1, \dots, d$.

For a random vector $\mathbf{X} = (X_1, \dots, X_d)$ we denote the joint distribution function and density by F and $f(x_1, \dots, x_d) = \frac{\partial F}{\partial x_1 \dots \partial x_d}(x_1, \dots, x_d)$, respectively. To model F , we exploit the famous Theorem of Sklar (1959). It separates the marginal distributions from the joint distribution such that $F(x_1, \dots, x_d) = C(F_1(x_1), \dots, F_d(x_d))$, where C denotes a d -dimensional *copula*. If all marginal distributions F_i are continuous, C is unique. The

corresponding density f with respect to the copula is obtained by taking derivatives

$$f(x_1, \dots, x_d) = c(F_1(x_1), \dots, F_d(x_d)) \prod_{i=1}^d f_i(x_i), \quad (2.1)$$

with $c(x_1, \dots, x_d) = \frac{\partial C}{\partial x_1, \dots, \partial x_d}(x_1, \dots, x_d)$ the *copula density*. While there exist many model classes for the univariate marginal distributions, this is not the case for copulas in arbitrary dimensions. The most well known copula functions are the multivariate Gaussian, Student t , Archimedean and extreme value copulas. Using these copulas to model d -dimensional data comes along with several drawbacks such as lack of flexibility and computational issues. To overcome this, Joe (1996) constructed distributions in d dimensions modelled by $d(d-1)/2$ bivariate distributions. Extending his work, Aas et al. (2009) developed the *pair-copula-construction (PCC)*, which builds up a d -dimensional distribution using d marginals and $d(d-1)/2$ (conditional) bivariate copulas. These building blocks can be chosen entirely independent from each other and thus provide a very flexible modeling approach. For example, pair copulas with heavy tails or asymmetric dependence can be used. Yet, the construction of a d -dimensional distribution with $d(d-1)/2$ (conditional) pairs is not unique. More precisely, there exist an exponentially growing number of valid constructions, see Kurowicka and Joe (2011, p. 190). With the work of Bedford and Cooke (2001, 2002), introducing *regular vines*, a framework was developed which allowed to organize the possible constructions by *vine trees*. In total, $d-1$ of these trees are required to define a d -dimensional distribution and are given by $\mathcal{V} = (T_1, \dots, T_{d-1})$ such that

- (i) T_1 is a tree with nodes $V_1 = \{1, \dots, d\}$ and edges E_1 ,
- (ii) for $i \geq 2$, T_i is a tree with nodes $V_i = E_{i-1}$ and edges E_i ,
- (iii) if two nodes in T_{i+1} are joined by an edge, the corresponding edges in T_i must share a common node (proximity condition).

By (ii), edges become nodes and are connected with new edges recursively. For a node, e.g. $\{1, 2\} \in V_2$, we define the two nodes $\{1\}, \{2\} \in V_1$ of which the node in V_2 is combined, as *m-children*. For some node in T_k , define the *m-family* as the union of all its m-children and their m-children in trees T_1, \dots, T_{k-1} . Each edge in one of the R-vine trees consists of a bivariate *conditioned set* and a *conditioning set*, ranging from the empty set to a set containing $d-2$ variables. To specify how an edge represents a specific (conditional) pair, let the *complete union* of an edge e be $A_e := \{j \in V_1 | \exists e_1 \in E_1, \dots, e_{i-1} \in E_{i-1} : j \in e_1 \in \dots \in e_{i-1} \in e\}$. The conditioning set of an edge $e = \{a, b\}$ is then given by $D_e := A_a \cap A_b$. The conditioned set is given by $\mathcal{C}_e := \mathcal{C}_{e,a} \cup \mathcal{C}_{e,b}$ with $\mathcal{C}_{e,a} := A_a \setminus D_e$ and $\mathcal{C}_{e,b} := A_b \setminus D_e$. For all edges $e \in E_i$, $1 \leq i \leq d-1$, we define the set of bivariate copula densities by $\mathcal{B}(V) = \{c_{j(e), \ell(e); D(e)} | e \in E_i, 1 \leq i \leq d-1\}$ with the conditioned set $j(e), \ell(e)$ and the conditioning set $D(e)$. Hence, with the PCC,

Equation (2.1) can be written as

$$f(x_1, \dots, x_d) = \left(\prod_{i=1}^d f_i(x_i) \right) \times \left(\prod_{i=1}^{d-1} \prod_{e \in E_i} c_{j(e), \ell(e); D(e)} \left(F(x_{j(e)} | \mathbf{x}_{D(e)}), F(x_{\ell(e)} | \mathbf{x}_{D(e)}) \right) \right). \quad (2.2)$$

In (2.2), we implicitly took into account the *simplifying assumption*. It imposes that a two-dimensional conditional copula density, e. g.

$$c_{13;2} (F_{1|2}(x_1|x_2), F_{3|2}(x_3|x_2); x_2)$$

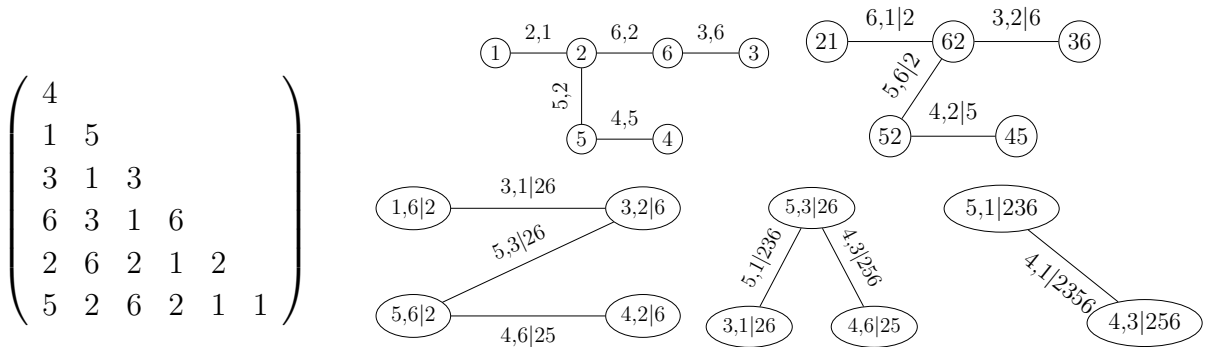
is independent of the conditioning value $X_2 = x_2$. A detailed discussion can be found in Stöber et al. (2013). We define the parameters of the bivariate copula densities $\mathcal{B}(V)$ by $\theta(\mathcal{B}(V))$. This determines the R-vine copula $(V, \mathcal{B}(V), \theta(\mathcal{B}(V)))$. An intuitive representation of vine copulas is given by lower triangular $d \times d$ matrices, see Dißmann et al. (2013). Such an R-vine matrix $M = (m_{i,j})_{i=1, \dots, d; j=1, \dots, d}$ has to satisfy three properties.

- (i) $\{m_{d,i}, \dots, m_{i,i}\} \subset \{m_{d,j}, \dots, m_{j,j}\}$ for $1 \geq i \geq j \geq d$,
- (ii) $m_{i,i} \notin \{m_{i+1,i+1}, \dots, m_{d,i+1}\}$,
- (iii) for all $j = d-2, \dots, 1$, $i = j+1, \dots, d$, there exist (k, ℓ) with $k < j$ and $\ell < k$ such that

$$\begin{aligned} \{m_{i,j}, \{m_{d,j}, \dots, m_{i+1,j}\}\} &= \{m_{k,k}, \{m_{1,k}, \dots, m_{\ell,k}\}\} \text{ or} \\ \{m_{i,j}, \{m_{d,j}, \dots, m_{i+1,j}\}\} &= \{m_{\ell,k}, \{m_{1,k}, \dots, m_{\ell-1,k}, m_{k,k}\}\}. \end{aligned} \quad (2.3)$$

The last property reflects the proximity condition. We now give an example R-vine.

Example 2.1 (R-vine in 6 dimensions). *The R-vine tree sequence in Figure 1 is described by the R-vine matrix M as follows. Edges in T_1 are pairs of the main diagonal and the lowest row, e. g. (2,1), (6,2), (3,6). T_2 is given by the main diagonal and the second last row conditioned on the last row, e. g. 6,1|2; 3,2|6. Higher order trees are characterized similarly. M -children of e. g. 6,1|2 are (2,1) and (6,2) and its m -family comprises additionally 1, 2, 6. For a column p in M , only entries of the main-diagonal right of p , i. e. values in $(m_{p+1,p+1}, \dots, m_{d,d})$ are allowed and no entry occurs more than once in a column.*



R-vine matrix M .

Figure 1: R-vine trees T_1, T_2 (top), T_3, T_4, T_5 (bottom), left to right. We abbreviate $c_{j,\ell|D} := c_{j,\ell;D}(F(x_i|\mathbf{x}_D), F(x_j|\mathbf{x}_D))$ for the conditioning vector \mathbf{x}_D . Additionally, let $\mathbf{x} = (x_1, \dots, x_6)$ and $f_i := f_i(x_i)$. The density in (2.2) for this case becomes

$$\begin{aligned} f(\mathbf{x}) &= f_1 \times f_2 \times f_3 \times f_4 \times f_5 \times f_6 \times c_{2,1} \times c_{6,2} \times c_{3,6} \times c_{5,2} \times c_{4,5} \times c_{6,1|2} \times c_{3,2|6} \\ &\quad \times c_{5,6|2} \times c_{4,2|5} \times c_{3,1|26} \times c_{5,3|26} \times c_{4,6|25} \times c_{5,1|236} \times c_{4,3|256} \times c_{4,1|2356}. \end{aligned}$$

The pair copula families and parameters are also described by lower triangular family and parameter matrices $\Gamma = (\gamma_{i,j})_{i=1,\dots,d;j=1,\dots,d}$ and $P = (p_{i,j})_{i=1,\dots,d;j=1,\dots,d}$. Thus, the family and parameters of the edge $6, 1|2$ are given by $\gamma_{5,4}$ and $p_{5,4}$, respectively. When two-parametric pair copulas are considered, we use an additional parameter matrix P_2 .

2.1 Model Assessment

We consider an R-vine model in d dimensions with specification $\Theta = (V, \mathcal{B}(V), \theta(\mathcal{B}(V)))$. Additionally, assume we have n replications of d -dimensional data $(\mathbf{x}_1, \dots, \mathbf{x}_n)^T \in \mathbb{R}^{n \times d}$ with $\mathbf{x}_k \in \mathbb{R}^d$ for $k = 1, \dots, n$. Including the marginal distributions $f_i(x_i)$, $i = 1, \dots, d$, the log-likelihood for the specification Θ on the u -scale is

$$\mathcal{L}(\Theta, (\mathbf{x}_1, \dots, \mathbf{x}_n)) = \sum_{k=1}^n \left(\sum_{i=1}^d \log \left(f_i(x_{k,i}) \right) + \sum_{i=1}^{d-1} \sum_{e \in E_i} \log \left(c_{j(e), \ell(e); D(e)} \left(F(x_{k,j(e)} | \mathbf{x}_{k, D(e)}), F(x_{k, \ell(e)} | \mathbf{x}_{k, D(e)}) \right) \right) \right)$$

The log-likelihood will always increase whenever more parameters are included in a model. Thus, in high dimensional setups where the number of significant parameters grows slower than the total number of possible parameters, it is not feasible to use log-likelihood for model discrimination. Because of this, *penalized goodness of fit* measures as the (AIC) and the *Bayesian information criterion (BIC)* (Schwarz, 1978) were developed. For $n \geq 8$, BIC will penalize more than AIC. If the number of possible parameters in an R-vine $q(d) = 2 \times d(d-1)/2$ is greater or equal than the sample size and the model is comparably small, BIC is no longer consistent and will penalize too little. For these setups, we use the *generalized information criterion (GIC)*, see Fan and Tang (2013). More precisely, we have

$$\begin{aligned} AIC(\Theta, (\mathbf{x}_1, \dots, \mathbf{x}_n)) &= -2\mathcal{L}(\Theta, (\mathbf{x}_1, \dots, \mathbf{x}_n)) + 2p \\ BIC(\Theta, (\mathbf{x}_1, \dots, \mathbf{x}_n)) &= -2\mathcal{L}(\Theta, (\mathbf{x}_1, \dots, \mathbf{x}_n)) + \log(n)p, \\ GIC(\Theta, (\mathbf{x}_1, \dots, \mathbf{x}_n)) &= -2\mathcal{L}(\Theta, (\mathbf{x}_1, \dots, \mathbf{x}_n)) + \log(\log(n)) \log(p)p, \end{aligned} \tag{2.4}$$

where p equals the number of parameters in the model Θ . Finally note that the number of parameters correspond to the the number of one parametric copulas and two times the number of two parametric copulas in the model.

2.2 Model Selection

Since the space of all R-vine structures is too large to explore explicitly each model, the standard estimation method relies on using a search heuristic such as the *Dißmann algorithm* (Dißmann et al., 2013). Initially, for each pair $(j, \ell) \in \binom{d}{2}$, Kendall's τ of the pair (U_j, U_ℓ) is calculated. The intuition is that variable pairs with high dependence should contribute significantly to the model fit and should be included in the first trees. Since the T_1 must be a tree, the $d-1$ edges with highest sum of absolute value of Kendall's

τ are chosen based on a maximum spanning tree algorithm, e.g. Prim (1957). Afterwards, on the selected edges either maximum likelihood estimation for all desired pair copula types is performed or the corresponding copula parameters are estimated by inversion of the empirical Kendall's τ . The later is only possible for one-parametric pair copula families. From these estimators, *pseudo-observations* are generated. More precisely, assume we want to estimate the pair copula density $c_{j(e),\ell(e);D(e)}$. Then, we use the pseudo-observations generated from $\widehat{F}_{j|D}(x_{j(e)}|\mathbf{x}_{D(e)})$ and $\widehat{F}_{\ell|D}(x_{\ell(e)}|\mathbf{x}_{D(e)})$ to estimate pair copula families and parameters on this pair. The corresponding $\widehat{F}_{j|D}$ and $\widehat{F}_{\ell|D}$ are given explicitly by derivatives of the pair copula distribution functions in lower trees evaluated at the estimated parameters. After taking into account the proximity conditions, Kendall's τ is calculated on all admissible pairs of pseudo-observations and again, a maximum spanning tree is calculated. After $d - 1$ iterations, the R-vine structure is determined. This proceeding has some drawbacks. First, it is not ensured that for each tree the maximum spanning tree in terms of Kendall's τ leads to a structure with e.g. optimal log-likelihood. Second, as lower order trees influence higher order trees, sub-optimal choices lead to error-propagation. Finally, in each step a tree is fitted over all remaining (conditional) pairs, and clusters within dependence are not treated any different from structures with less dependence. Overall the effort is of order d^2 since $d(d - 1)/2$ pair copulas are estimated.

As mentioned, using Kendall's τ as edge weights is only heuristically backed by the goal to model the strongest dependency first. However, this approach is not ensured to optimize e.g. log-likelihood or AIC in a given tree and thus, also not in the entire R-vine. We can instead also estimate pair copula densities $c_{j,\ell}$ for each edge (j, ℓ) on all $d(d - 1)/2$ edges in the first tree and then calculate a maximum spanning tree with respect to an edge weight $\mu_{j,\ell}$, e.g. log-likelihood or AIC, based on the actual fit. Thus, this choice would optimize for e.g. log-likelihood or AIC in a given tree. For higher trees, this can be done similarly. This approach has firstly been discussed in Czado et al. (2013). However, it was considered having unacceptable computational complexity. We will however come back to it later for an improved version.

The algorithm proposed by Müller and Czado (2017a) uses graphical models, more precisely *directed acyclic graphs (DAGs)*, to find parsimonious structures and set the majority of pair copulas to the independence copula in larger datasets, which eases the computational effort. However, for both their and Dißmann's algorithm, more than 500 – 1000 dimensions are not solvable because of the quadratically increasing effort in terms of computation time and memory.

Another promising proposal has been made by Kraus and Czado (2017), explicitly searching for simplified R-vine structures. One can expect that the true structure of the R-vine can be correctly identified if the data is originating from a simplified R-vine. However, it works similar to Dißmann's algorithm and hence, has the same computational complexity. A different approach was proposed in Müller and Czado (2017b), where the Lasso (Tibshirani, 1994) is used to compute not only the R-vine structure, but also a regularization path of the R-vine matrix, allowing for flexible thresholding to reduce the model complexity based on a single structure. Even though the approach does not rely on maximum

spanning trees, a d -dimensional R-vine needs to be fitted and stored, which is accompanied by the same computational complexity. Large computational gains are however obtained by setting a majority of all pair copulas to independence copulas upfront.

There have been attempts to relate undirected graphical models to R-vines for structure selection. Hobæk Haff et al. (2016) showed that a k -truncated R-vine can be expressed as chordal graph with maximal clique size $k + 1$ and vice versa. However, this chordal graph needs to adhere to some other non-trivial properties. These are in practice not met when a graphical model is fitted. Thus, finding a sparse undirected graph and translate it into a sparse R-vine is hard and computationally infeasible. Kovács and Szántai (2016) propose an algorithm to calculate a k -truncated R-vine from a chordal graph with maximal clique size $k + 1$ and show that taking into account the aforementioned property only leads to a chordal graph with maximal clique size $k + 2$. The later can then be used to estimate a $k + 2$ truncated R-vine. However, the problem is only deferred. Where there exist several methods for finding sparse undirected graphical models in high dimensions, these will not be chordal in most cases. Finding a so called *chordal cover* with clique size at most $k + 1$ is known to be *NP-complete*, see Arnborg and Corneil (1987). Hence, for dimensions $d > 1000$ where k can not be assumed too small, the problem is also intractable.

Finally, there also exist Bayesian methods for R-vine estimation (Gruber and Czado, 2015b,a). However, these do require even more computational effort and are hence not feasible in more than $d > 20$ dimensions.

2.3 Model Simplification

All of the proposed approaches may be modified by testing each pair copula for the independence copula with density $c^\perp(u_1, u_2) = 1$. Thus, a type-1 error $\alpha \in (0, 1)$ is specified and each pair-copula is tested for the null hypothesis to be the independence copula. Only if this hypothesis can be rejected at the level $1 - \alpha$, an estimation of pair copula family and parameter is performed. Similarly, entire R-vine trees can be tested for only containing the independence copula. However, this also requires that an additional tree is fitted before it can be tested. Contrary to this, a *truncation level* $k \in \{1, \dots, d - 2\}$ can be specified upfront. By doing so, only the first k R-vine trees are estimated and independence is assumed for all higher trees. If a k -truncation is imposed, (2.2) becomes

$$f(x_1, \dots, x_d) = \left(\prod_{i=1}^d f_i(x_i) \right) \times \left(\prod_{i=1}^k \prod_{e \in E_i} c_{j(e), \ell(e); D(e)} \left(F(x_{j(e)} | x_{D(e)}), F(x_{\ell(e)} | x_{D(e)}) \right) \right).$$

In Example 2.1, a k -truncated R-vine is given by $c_{j, \ell; D} = c^\perp$ if $|D| \geq k$. More details can be found in Brechmann et al. (2012). Finally, testing for independence does not explicitly decrease the computational effort. However, a k truncation leads to only kd pair copulas, whereas identifying a sensible truncation level prior to estimation is hard.

3 Graphical Independence Models

To find structures for R-vines, we will use tools which have proven very successful in the Gaussian setting. More precisely, we will use models based on *undirected graphs*.

Most of the terminology follows Lauritzen (1996, pp. 4-7). Let $V \neq \emptyset$ be a finite set, the *node set* and let $E \subseteq \{(\alpha, \beta) \mid (\alpha, \beta) \in V \times V \text{ with } \alpha \neq \beta\}$ be the *edge set*. Thus, we obtain a graph $\mathcal{G} = (V, E)$ as pair of node set and edge set. We assume $(\alpha, \beta) \in E \Rightarrow (\beta, \alpha) \in E$, i. e. only *undirected edges* and hence an *undirected graph*. Define a *path* of length k from nodes α to β by a sequence of distinct nodes $\alpha = \alpha_0, \dots, \alpha_k = \beta$ such that $(\alpha_{i-1}, \alpha_i) \in E$ for $i = 1, \dots, k$. A *cycle* is defined as a path with $\alpha = \beta$. A graph $\mathcal{H} = (W, \mathcal{F})$ is a *subgraph* of $\mathcal{G} = (V, E)$ if $W \subseteq V$ and $\mathcal{F} \subseteq E$. We have an *induced subgraph* $\mathcal{H} = (W, \mathcal{F})$ if $W \subseteq V$ and $\mathcal{F} = \{(\alpha, \beta) \mid (\alpha, \beta) \in W \times W \text{ with } \alpha \neq \beta\} \cap E$, i. e. \mathcal{H} contains a subset of nodes of \mathcal{G} and all the edges existing between these nodes in \mathcal{G} . If a path from α to β exists for all $\alpha, \beta \in V$, we say that \mathcal{G} is *connected*. Whenever for a graph $\mathcal{G} = (V, E)$ we have that there exists a disjoint partition of $V = \bigcup_{i=1}^p V_i$ such that the p subgraphs \mathcal{H}_i induced by V_i for $i = 1, \dots, p$ are connected subgraphs, we speak of *connected components* of \mathcal{G} . For \mathcal{G} undirected, $\alpha, \beta \in V$, a set $S \subseteq V$ is said to be an (α, β) *separator* in \mathcal{G} if all paths from α to β intersect S . S is said to *separate* A from B in \mathcal{G} if it is an (α, β) separator in \mathcal{G} for every $\alpha \in A, \beta \in B$ and we denote it by $\alpha \perp \beta \mid S$ [\mathcal{G}] and $A \perp B \mid S$ [\mathcal{G}], respectively. We define the *adjacency matrix* $\Pi^{\mathcal{G}} = (\pi_{i,j}^{\mathcal{G}})_{i,j=1,\dots,d} \in \{0, 1\}^{d \times d}$ of a graph $\mathcal{G} = (V, E)$ such that $\pi_{i,j}^{\mathcal{G}} = \pi_{j,i}^{\mathcal{G}} = 1 \Leftrightarrow (i, j) \in E$ and 0 otherwise.

3.1 Probabilistic (Gaussian) Graphical Models

Graph theory and statistical models can be linked, obtaining *probabilistic graphical models*. More precisely, consider a random vector $\mathbf{X} = (X_1, \dots, X_d)$ and assume $\mathbf{X} \sim \mathcal{N}_d(\mathbf{0}, \Sigma)$, denoting a d -dimensional Gaussian distribution with probability density function

$$f(\mathbf{x}; \mathbf{0}, \Sigma) = (2\pi)^{-\frac{d}{2}} \det(\Sigma)^{-\frac{1}{2}} \exp\left(-\frac{1}{2} \mathbf{x}^T \Sigma^{-1} \mathbf{x}\right). \quad (3.1)$$

Let $i, j \in \{1, \dots, d\}$ with $i \neq j$ and $S \subseteq \{1, \dots, d\} \setminus \{i, j\}$. We denote that X_i is conditionally independent of X_j given \mathbf{X}_S by $X_i \perp\!\!\!\perp X_j \mid \mathbf{X}_S$. In the remaining, let $\mathcal{G} = (V = \{1, \dots, d\}, E)$. We say that \mathbf{X} is *Markov* with respect to \mathcal{G} when

$$X_j \perp\!\!\!\perp X_\ell \mid \{X_1, \dots, X_d\} \setminus \{X_j, X_\ell\} \Leftrightarrow (j, \ell) \notin E. \quad (3.2)$$

This means, each missing edge in the graph \mathcal{G} corresponds to random variables which are conditionally independent given the rest. The later is however also directly expressed via the inverse of the correlation matrix. More precisely, let $\Omega = \Sigma^{-1}$, then

$$\Omega_{j\ell} = 0 \Leftrightarrow X_j \perp\!\!\!\perp X_\ell \mid \{X_1, \dots, X_d\} \setminus \{X_j, X_\ell\}.$$

Thus, estimating the *sparsity pattern* of Ω is equivalent to estimating the graph \mathcal{G} . From the graph however, we can also extract more information. If the underlying distribution

is Gaussian, we consider the *global Markov property* such that for $S \subseteq \{1, \dots, d\} \setminus \{i, j\}$ we have that

$$X_j \perp\!\!\!\perp X_\ell \mid \mathbf{X}_S \Leftrightarrow j \perp \ell \mid S \ [\mathcal{G}], \quad (3.3)$$

which is a more favourable property than (3.2) and which we will exploit later. Next, we will focus on how to estimate such graphs, i. e. the matrix Ω , given data.

Example 3.1 (Graphical model). *Consider the graph in Figure 2. By virtue of the*




Figure 2: Graphical model in 5 dimensions.

graphical separation, we have e. g. $1, 4 \perp 5 \mid 2, 3 \ [\mathcal{G}] \Rightarrow 1, 4 \perp\!\!\!\perp 5 \mid 2, 3$.

3.2 Estimating Sparse Inverse Covariance Matrices with the Graphical Lasso

To estimate a sparse precision matrix Ω , we use the well-known *graphical Lasso*, see Friedman et al. (2008). Denote the sample covariance matrix by $S = X^T X/n \in \mathbb{R}^{d \times d}$ where $X = (x_{i,j})_{i=1, \dots, n, j=1, \dots, d} \in \mathbb{R}^{n \times d}$ is the observed and centred data matrix with $\mathbf{X} \sim \mathcal{N}_d(\mathbf{0}, \Sigma)$. Then, the graphical Lasso calculates a sparse undirected graphical model by finding a solution for Ω . Considering the logarithm of (3.1) and taking derivatives with respect to Σ we obtain as optimization problem

$$\max_{\Omega \in \mathbb{R}^{d \times d}} \log(\det(\Omega)) + \text{tr}(S\Omega). \quad (3.4)$$

with solution $\widehat{\Omega} = S^{-1}$. However, this will in general have no zero entries and hence, induce no sparse graph. Furthermore, in high dimensional data sets we often have $d > n$, which leads to a singular matrix S and the inverse of S does not exist. The graphical Lasso overcomes this by introducing a penalty in (3.4) and solving

$$\max_{\Omega \in \mathbb{R}^{d \times d}} \log(\det(\Omega)) + \text{tr}(S\Omega) + \lambda \sum_{i=1}^d \sum_{j=1}^d |\Omega_{ij}|. \quad (3.5)$$

depending on some *regularization parameter* $\lambda \geq 0$. For this optimization problem, many efficient numerical solvers also for thousands of dimensions exist. For $\lambda = 0$, there is no penalization and the solutions of (3.4) and (3.5) coincide. For fixed $\lambda > 0$, denote the solution of (3.5) by $\widehat{\Omega}^\lambda$ and define $\mathcal{G}^\lambda = (V, E^\lambda)$ by

$$(i, j) \in E^\lambda \Leftrightarrow \widehat{\Omega}_{ij}^\lambda \neq 0. \quad (3.6)$$

Varying $\lambda > 0$, we obtain a piecewise constant *solution path* of graphs with different levels of sparsity. Letting $\lambda \rightarrow 0$, the solutions \mathcal{G}^λ will become more and more dense. For $\lambda \rightarrow \infty$, the number of *connected components* will increase but their individual sizes decreases. In practice, the solution path is calculated along a vector $\boldsymbol{\lambda} = (\lambda_1, \dots, \lambda_J)$ with $\lambda_j > 0$. Several modifications and improvements for the graphical Lasso have been proposed. For instance, Meinshausen and Bühlmann (2006) show that neighbourhood selection consistently estimates the graph \mathcal{G} , however, no estimate $\widehat{\Omega}$ is obtained. Witten et al. (2011) demonstrate that the search for (3.5) can be carried out in terms of block-diagonal matrices, breaking apart the large problem into smaller ones.

Vital for the first part of our proposed approach is that the connected components with respect to some $\lambda > 0$ can also be calculated directly from the sample covariance matrix S . More precisely, consider a fixed $\lambda > 0$ and a solution $\widehat{\Omega}^\lambda$ of (3.5) with $\mathcal{G}^\lambda = (V, E^\lambda)$ as defined in (3.6). Thus, $\widehat{\Omega}^\lambda$ and equivalently E^λ induce a vertex partition

$$V = \bigcup_{i=1}^p V_i^\lambda,$$

where each V_i^λ is a connected component for $i = 1, \dots, p$, based on the edge set E^λ . Alternatively, define a graph $\mathcal{H} = (V, \mathcal{F}^\lambda)$ based on the sample covariance matrix S with edge set \mathcal{F}^λ and adjacency matrix $\Pi^\mathcal{H}$ such that

$$\pi_{i,j}^\mathcal{H} = 1 \Leftrightarrow (i, j) \in \mathcal{F}^\lambda \Leftrightarrow |S_{ij}| \geq \lambda, \quad (3.7)$$

This way of assigning the edge set is called *screening*. The graph \mathcal{H} has now, say q connected components we denote by W_i for $i = 1, \dots, q$ and consider the associated partition $V = \bigcup_{i=1}^q W_i$. It has now been shown by Mazumder and Hastie (2012), that $p = q$ and moreover, $V_i = W_i$ for all $i = 1, \dots, p$. This makes a decomposition of the entire graphical Lasso problem more tractable as we can split it into p parallel tasks which can be performed entirely independent from each other. Furthermore, we have a very easy screening rule for intractably high dimensional datasets to decompose their dependence behaviour in multiple smaller parts which now are tractable. Inside the connected components, we can then use the graphical Lasso to obtain non-dense graphs. We give a brief example.

Example 3.2 (Screening). *Assume we have a dataset on the z-scale in 6 dimensions with the following empirical covariance matrix*

$$S = \begin{pmatrix} 1.0000 & 0.2058 & 0.1794 & 0.7340 & 0.7298 & 0.7167 \\ 0.2058 & 1.0000 & 0.3212 & 0.2643 & 0.3158 & 0.2848 \\ 0.1794 & 0.3212 & 1.0000 & 0.1895 & 0.2105 & 0.2327 \\ 0.7340 & 0.2643 & 0.1895 & 1.0000 & 0.9606 & 0.9089 \\ 0.7298 & 0.3158 & 0.2105 & 0.9606 & 1.0000 & 0.9378 \\ 0.7167 & 0.2848 & 0.2327 & 0.9089 & 0.9378 & 1.0000 \end{pmatrix}.$$

We use the *huge* R-package (Zhao et al., 2015) to calculate a sequence of $J = 4$ values for λ given by

$$\boldsymbol{\lambda} = (0.9607, 0.7438, 0.3452, 0.2070).$$

Solving (3.5) for these values we obtain graphical models $\mathcal{G}^{\lambda_1}, \dots, \mathcal{G}^{\lambda_4}$ as shown in Figure 3. Initially, there are only isolated nodes and hence $|V| = 6$ connected components. In the




Figure 3: Example 3.2: Sequence of estimated graphical models $\mathcal{G}^{\lambda_1}, \dots, \mathcal{G}^{\lambda_4}$.

second and third graph, we have one connected component of size 3 together with isolated nodes and size 4 with isolated nodes. In the fourth graph, only one connected components of size $|V| = 6$ exists. If we consider $\lambda_2 = 0.7438$ and apply the screening rule (3.7), we obtain the following adjacency matrix based on S

$$\Pi^{\mathcal{G}^{\lambda_2}} = \begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 1 & 0 & 1 \\ 0 & 0 & 0 & 1 & 1 & 0 \end{pmatrix}.$$

which constitutes the same connected components as in the corresponding graph of the graphical Lasso solution path. The adjacency matrix Π indicates that the nodes 4, 5, 6 are all connected to each other. However, the result by Mazumder and Hastie (2012) only assures that the connected components of the screening rule and the graphical Lasso are the same for a specific $\lambda > 0$, but no knowledge about the edges in the connected components is can be drawn from the screening rule.

4 Selection of High Dimensional R-vines with Graphical Models

Recall that our goal is ultimately to estimate R-vines in ultra high dimensions. Thus instead of estimating an R-vine as described in Section 2 on data in d dimensions, we first cluster the dataset into a partition and then perform estimation on the elements of

the partition. The partition is provided by the methods presented in Section 3, i. e. the graphical Lasso. Inside these clusters, we can then estimate R-vines in smaller dimensions with improved accuracy as a consequence of exploiting the global Markov property, (3.3).

4.1 Clustering High Dimensional Data

Assume that we are given data $X \in \mathbb{R}^{n \times d}$, where $d \approx 1000$ or more. We will consider a sequence of J disjoint partitions of $V = \{1, \dots, d\}$ into p_j connected components $V^j := \bigcup_{i=1}^{p_j} V_i^j$, $j = 1, \dots, J$. This is expressed by a sequence of graphical models

$$\mathcal{G}_1 = \left(V^1 = \bigcup_{i=1}^{p_1} V_i^1, \bigcup_{i=1}^{p_1} E_i^1 \right), \dots, \mathcal{G}_J = \left(V^J = \bigcup_{i=1}^{p_J} V_i^J, \bigcup_{i=1}^{p_J} E_i^J \right). \quad (4.1)$$

In most practical applications, $J = 15$ or $J = 30$, see Zhao et al. (2015). If partition V^j is only a single connected component, we have $p_j = 1$. The sequence \mathcal{G}_j for $j = 1, \dots, J$ can be identified as solution path corresponding to the graphical Lasso for J different penalization values of $\lambda > 0$. To identify the size of connected components in these graphs, define for each partition V^j , $j = 1, \dots, J$,

$$\delta_j = \max_{i=1, \dots, p_j} |V_i^j|. \quad (4.2)$$

Instead of considering the entire dataset in d dimensions, we consider subsets of lower dimensions on the connected component with maximum dimension $\delta_j < d$ dimensions. In practical applications, we will have some *threshold dimension* $0 < d_T < d$ and calculate the solution path of the graphical Lasso for a sequence $\{\lambda_1, \dots, \lambda_J\}$ based on the screening property (3.7). This works very fast and we can select the corresponding graphical model and associated partition V^T by

$$T = \arg \max_{j=1, \dots, J} \delta_j \text{ such that } \delta_j \leq d_T. \quad (4.3)$$

Finally, we denote the chosen partition T by V^T and the corresponding graph by $\mathcal{G}_T = (V^T = \bigcup_{i=1}^{p_T} V_i^T, \bigcup_{i=1}^{p_T} E_i^T)$.

Example 4.1 (Example 3.2 cont.). *Consider the sequence of graphs in Example 3.2. Using the notation of (4.1) and defining $\mathcal{G}_j = \mathcal{G}^{\lambda_j}$, we have $J = 4$ and p_j, δ_j for $j = 1, \dots, 4$ as follows: Assume $d_T = 4$, then $T = \arg \max_{j=1, \dots, 4} \delta_j$ such that $\delta_j \leq 4$, thus*

j	1	2	3	4
p_j	6	4	3	1
δ_j	1	3	4	6

Table 1: Example 4.1: Number of connected components p_j and maximal component sizes δ_j for graphs in Figure 3.

$T = 3$. Hence, the graph \mathcal{G}_3 is selected with partition $V^3 = \{1, 4, 5, 6\} \cup \{2\} \cup \{3\}$ with $p_T = 3$ and $\delta_T = 4$.

Now, we consider the problem of estimating *sub-R-vines* on these connected components induced by a partition V^T . Thus, with respect to this partition we estimate R-vines on the connected components of \mathcal{G}_T , i. e. the elements of the partition V_i^T , $i = 1, \dots, p_T$ with at most dimension δ_T . Each of these *sub-R-vines* is then stored in an R-vine matrix and the corresponding matrices can be combined non-uniquely to an R-vine matrix of dimension $d \times d$. This is however an incomplete R-vine matrix as it does not contain information how the connected components are connected to each other. Additionally, connected components of size 1, i. e. isolated nodes are not yet included. Both the missing connections and the isolated nodes can however be easily connected afterwards as we will show in a subsequent example. For this, we also introduce the *fill-level* $k_F \geq 0$ which determines how many R-vine trees *outside* the connected components should be estimated.

Example 4.2 (Toy example). Assume we have the graphical model \mathcal{G}_T as in Figure 4. Given this graphical model, we estimate sub-R-Vines on the components $\{1, 4, 6\}$ and




Figure 4: Example 4.2: Graphical model \mathcal{G}_T in 6 dimensions.

$\{2, 5\}$, respectively. These two components give rise to the following R-vine matrices, where we assume that the estimate in M_1 is optimal with respect to some edge weight.

$$M_1 = \begin{pmatrix} 1 & & \\ 6 & 4 & \\ 4 & 6 & 6 \end{pmatrix} \quad M_2 = \begin{pmatrix} 2 & \\ 5 & 5 \end{pmatrix}.$$

Together with the isolated node 3, these two R-vines can be arranged into a joint R-vine on 6 dimensions described by the matrix M_{123} .

$$M_{123} = \begin{pmatrix} 3 & & & & & \\ & 2 & & & & \\ & & 5 & & & \\ & & & 1 & & \\ \hline & & & 6 & 4 & \\ \square_1 & 5 & \square_2 & 4 & 6 & 6 \end{pmatrix}, \quad M'_{123} = \begin{pmatrix} 3 & & & & & \\ & 2 & & & & \\ & & 5 & & & \\ & & & 1 & & \\ \hline \triangle_1 & \triangle_2 & \triangle_3 & 6 & 4 & \\ 2 & 5 & 1 & 4 & 6 & 6 \end{pmatrix}.$$

This would correspond to the graphical model in \mathcal{G}_T . Now, we can connect the connected components with each other in the first k_F trees, i. e. to the fill-level. This works as following for the example of $k_F = 2$. The entries in M_{123} marked by \square describe the pair copulas between the connected components and are chosen from $\square_1 \in \{2, 5, 1, 4, 6\}$ and $\square_2 \in \{1, 4, 6\}$. For example, $\square_2 \notin \{3, 2\}$ since only diagonal entries from the right of the

this increases the required effort for an R-vine tree sequence significantly. As we have $\nu - 1$ trees on $\nu - i + 1$ nodes for $i = 2, \dots, \nu - 2$ and consider in the worst case all possible pairs in each tree, this sums up to

$$\sum_{i=1}^{\nu-2} \frac{(\nu - i + 1)(\nu - i)}{2} = \frac{\nu^3 - \nu}{6} - 1, \quad (4.4)$$

pair copulas, where the leading term can only be bounded from above by order ν^3 . Note that this is a worst case complexity since in higher trees, the proximity condition can exclude some edges, and hence, no pair copula needs to be estimated for these. In the particular case of a *D-vine*, where the first tree is a path, i. e. each node except for the two end nodes has degree 2, the remaining R-vine tree sequence is completely determined and the effort collapses to order $\nu(\nu - 1)/2$. Whereas in the case of a *C-vine*, where each tree is a star, the upper bound in (4.4) is attained. Recall that the method of Dißmann et al. (2013), only $\nu(\nu - 1)/2$ pair copulas have to be estimated. Hence, for general R-vines, estimation of ν^3 pair copulas is not admissible even if $\nu \ll d$. However, we can take into account the (conditional) independence information by the graph \mathcal{H} to set a huge number of pair copulas to independence copulas upfront, leading to a significant decrease of computational effort.

- Recall that \mathcal{F} is the edge set of \mathcal{H} and define $\mathcal{H}_0 = (W, W \times W)$, i. e. a full graph on W and assign pair copulas and weights

$$c_{j,\ell} = 1 \Leftrightarrow \mu_{j,\ell} = 0 \Leftrightarrow (i, j) \notin \mathcal{F}. \quad (4.5)$$

For all remaining pairs $(j, \ell) \in W \times W$, we perform maximum likelihood estimation on the pair copulas and obtain weights $\mu_{j,\ell}$ based on the actual fit. Thus, we have weights for all edges and can now calculate a maximum spanning tree $T_1 = (W, E_1)$.

- We define $T_k = (E_{k-1}, E_{k-1} \times E_{k-1})$ for $k = 2, \dots, \nu - 1$, i. e. a full graph on E_{k-1} . We remove all edges not allowed by the proximity condition. For the remaining edges, we assign pair copulas and weights

$$c_{j,\ell|\mathbf{D}} = 1 \Leftrightarrow \mu_{j,\ell|\mathbf{D}} = 0 \Leftrightarrow j \perp \ell \mid \mathbf{D} [\mathcal{H}]. \quad (4.6)$$

For all remaining pairs $(j, \ell) \in W \times W$ we perform maximum likelihood estimation on the pair copulas and obtain weights $\mu_{j,\ell}$ based on the actual fit. Thus, we have weights for all edges and can now calculate a maximum spanning tree T_k for $k = 2, \dots, \nu - 1$.

If the connected component \mathcal{H} is not too dense, i. e. $|\mathcal{F}| \sim \nu$ instead of $|\mathcal{F}| \sim \nu^2$, overall about ν^2 pair copulas are estimated for the entire R-vine in ν dimensions. From our point of view, it is more beneficial to neglect the edges which are not chosen anyway because of (conditional) independence and perform a more thorough analysis on the remaining edges. The algorithm of Dißmann might miss important edges, especially if all estimated Kendall's τ values are similar. Furthermore, our approach fosters sparsity by setting

pair copulas to independence straight away. This is not possible in Dißmann’s algorithm, except for truncation, which however has to be specified upfront in a very inflexible manner. We will show that our approach is able to capture significantly more dependence compared to Dißmann’s. An example illustrating this approach finalizes the section.

Example 4.3 (Example 4.1 cont.). *Assume the graphical model \mathcal{H} as in Figure 5 (left). Using our approach, we fit pair copulas $c_{13}, c_{14}, c_{23}, c_{34}$ and consider the corresponding*




Figure 5: Example 4.3: Graphical model \mathcal{H} (left) and first R-vine tree T_1 (right).

goodness of fit values $\mu_{13}, \mu_{14}, \mu_{23}, \mu_{34}$, given for example by the AIC. The other missing edges are neglected, i. e. no pair copula is estimated. Assume additionally that the optimal choice with respect to AIC is given by the tree given in the right panel of Figure 5. Then, the second R-vine tree T_2 may contain for example the possible edge $c_{24|3}$. However, we can see from the left panel of Figure 5 that $2 \perp 4 \mid 3$ [\mathcal{H}] and thus, the pair copula $c_{24|3}$ is set to the independence copula upfront in T_2 .

4.3 Implementation of the Algorithm

We combine the two previous steps in one algorithm, allowing to estimate high dimensional vine copulas based on a clustering by the graphical Lasso. For this, we need to specify a threshold dimension $d_T < d$ and a fill-level $k_F \geq 0$ describing until which tree we will estimate pair-copulas *outside* the connected components. The later is beneficial since we only use the connected components to break our original problem into tractable sub-problems. However, we assume there is dependence outside the connected components. From a computational point of view, estimating within the connected components of dimension at most $d_T < d$ first and then connect these components afterwards is much more beneficial than estimating an R-vine on d dimensions. We use the **huge** R-package, see Zhao et al. (2015) to generate high dimensional undirected graphical models and use the default settings there. Since we normally operate on *copula-data* (U_1, \dots, U_n) , i. e. data with uniform marginals, we transform our observations to the *z-scale*, i. e. consider $(Z_1 = \Phi^{-1}(U_1), \dots, Z_d = \Phi^{-1}(U_d)) \sim \mathcal{N}(0, \Sigma)$ in a Gaussian set up. The only change we perform in the default setting of **huge** is that we always want to obtain a sequence of 30 graphs, this is regulated by the number of $\lambda_j \geq 0$ to be evaluated. In terms of the previous notation, we set $J = 30$. Our algorithm then selects a partition such that the maximal component size is less or equal d_T as shown in Section 4.1 and then performs on each of the components an improved R-vine selection based on Section 4.2. The edge metric μ we use is the AIC of the associated pair copula term in our considerations. Afterwards,

we combine these *sub-R-vines* into one *joint R-vine* matrix on which we operate further. Finally, we estimate the pair copulas in the first $k_F \geq 0$ trees of the joint R-vine. From this point on, we only operate on the R-vine matrix M . Thus, we have to take into account that in a certain column j of M , only values of the main diagonal right of j , i. e. in $M_{j+1,j+1}, \dots, M_{d,d}$ can occur. Thus, for a given entry in $M_{i,j}$, we check which entries in $M_{j+1,j+1}, \dots, M_{d,d}$ are valid according to the proximity condition. For those, we fit pair copulas and select the best choice according to some metric μ , e. g. $\mu = \text{AIC}$. The entire algorithm, which we will refer to as *RVineClusterSelect*, is given Appendix, Section A.

5 Numerical Examples

To show the feasibility of our approach, we present different numerical examples.

- (i) A simulation study of several sparse dependence models with respect to a subset in $d = 85$ dimensions of the S&P100 index constituents. We compare our approach to the Dißmann algorithm with respect to goodness of fit and computation time.
- (ii) A runtime comparison with Dißmann's in up to $d = 1750$ dimensions.
- (iii) An ultra high-dimensional data application involving $d = 2131$ stocks from different sectors and multiple geographies. We compare our approach to Gaussian models.

5.1 Preprocessing of Data

Because of the wide availability, we apply our methods to high dimensional financial daily data. More precisely, we consider closing prices of shares adjusted for dividends and splits S_i^j with $i = 1, \dots, n$ observations for $j = 1, \dots, d$ shares. In the next step, we calculate daily log returns by $R_i^j = \log(S_i^j/S_{i-1}^j)$ for $i = 2, \dots, n$. These log returns are then filtered for trend and seasonality, i. e. idiosyncratic behaviour. We use ARMA-GARCH(p, q) models with $(p, q) \in \{0, 1\} \times \{0, 1\}$, i. e. four different specifications and allow for residuals distributed according to three different distributions, Gaussian, Student's-t or skewed Student's-t. In total we consider $4 \times 3 = 12$ models which we fit for each marginal time series and choose the best in terms of log-likelihood. Next, we compute standardized residuals $x_{ij} = (R_i^j - \widehat{R}_i^j) / \widehat{\sigma}_j$ for $j = 1, \dots, d$ and $i = 1, \dots, n$ based on the estimated time series models, where $\widehat{\sigma}_j^2, j = 1, \dots, d$ is the estimated variance of the error distribution. Finally, we calculate the *empirical cumulative distribution function* \widehat{F}_j of x_{1j}, \dots, x_{nj} for $j = 1, \dots, d$ to obtain *copula data* $U_{ij} = \widehat{F}_j(x_{ij})$ for $i = 1, \dots, n$ and $j = 1, \dots, d$, on which we work on. The data was acquired using the `quantmod` R-package (Ryan and Ulrich, 2017) and the marginal time series models were fitted using the `rugarch` R-package (Ghalanos, 2015).

5.2 Simulation Study

First of all, we want to see whether our approach generates feasible models in moderate dimensions, e. g. $d \sim 100$. There, also the standard algorithm for vine copula estimation of Dißmann et al. (2013) works well. However, as our approach is especially targeted to sparse scenarios, we want to account for this and design three scenarios. We obtained stock quotes for $d = 85$ stocks in the S&P100 index from 01.01.2013 to 31.12.2016 and process them as previously described. The remaining 15 stocks are removed since they dropped out of the index by the end of the observation period. We fit three different vine copula models V_1, V_2, V_3 to the data allowing for all parametric pair copula families implemented in the `VineCopula` R-package of Schepsmeier et al. (2017), however imposing a 2, 5 and 10-truncation. Additionally, we perform an independence test at significance level $\alpha = 0.05$ to pair copulas to the independence copula. Clearly, the 2-truncation is the most parsimonious model, however, also the 5 and 10-truncation are quite reduced models compared to a full model with 84 trees. From these fitted models, we simulate 50 replications with $n = 1000$ observations each for which we compare our method with the algorithm of Dißmann. We use runs with threshold dimensions $d_T = 25, 50, 75$, see (4.3) and $k_F = \lceil \log(d) \rceil = \lceil \log(85) \rceil = 5$. This choice worked very well in the numerical experiments we carried out. Before evaluating the corresponding models, we consider the maximal component sizes δ_T^i and the number of connected components p_T^i of the chosen partitions T_i for all 50 replications $i = 1, \dots, 50$. There, we observe that for higher threshold dimension, the number of the connected components decrease while the maximal component sizes naturally increases, see Figure 6. Next, we present the results




Figure 6: Scenario V_1 , 2-truncation: Distribution of maximal component sizes δ_T^i (left) and distribution of the number of connected components p_T^i (right) for each of the 50 replications, $i = 1, \dots, 50$ for different threshold dimensions $d_T = 25, 50, 75$.

with respect to goodness of fit, number of parameters and computation time for the 2-

truncation in Figure 7 and defer the remaining scenarios to the Appendix, Section B as the results behave quite similarly. We denote the model from which we simulate as true model. First of all, we see that Dißmann’s algorithm obtains highest log-likelihood in all

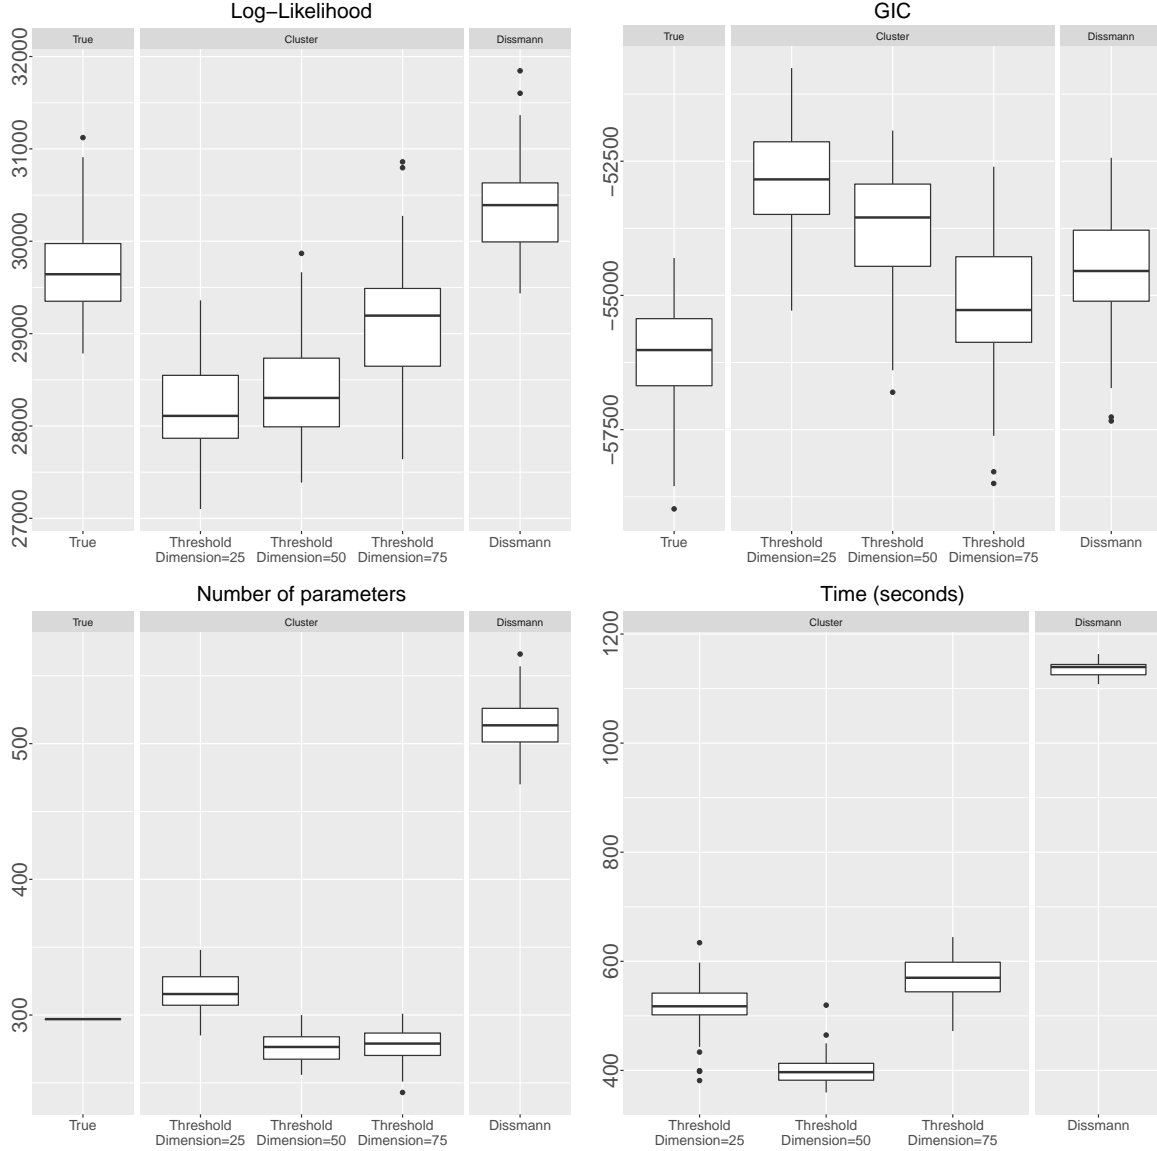


Figure 7: Scenario V_1 , 2-truncation: Comparison of `RVineClusterSelect` algorithm with threshold dimension $d_T = 25, 50, 75$ and Dißmann’s algorithm on *u-scale*: log-likelihood (upper left), GIC (upper right), number of parameters (lower left) and computation time (lower right).

scenarios. However, it also tends to overfit as it has almost twice as much parameters as the true model. Our approach captures the number of parameters better and this is accompanied with lower, i. e. better GIC, see (2.4), for the threshold dimension $d_T = 75$. We also see significantly lower computation times. Generally, our approach needs far less parameters but selects the most important bivariate dependencies to yield parsimonious models.

5.3 Runtime and Goodness of Fit Analysis on Real World Data

We obtained data from $d = 1757$ stock listed companies from the entire globe in the financial services industry. The data contains $n = 470$ trading days, i. e. about two years of data. Thus, we have a data matrix $U \in [0, 1]^{n \times d}$ and split it into 18 nested subsets $U_\ell \in [0, 1]^{n \times d_\ell}$ with $d_\ell \in \{50, 150, 250, \dots, 1750\}$ by only considering the first d_ℓ columns. Our goal is to compare our approach to Dißmann’s method considering the computation times with respect to the dimension. We expect that our approach is on the one hand side much faster, and on the other hand side more precise because of using a more accurate estimation approach as introduced in Section 4.2. The question arises, which maximal component size d_ℓ^T to choose for our algorithm and which fill level k_F . For the maximal component size, we evaluate a grid of 9 graphs with the maximal component sizes given Table 2. Here, we refer to the maximal component size of graph \mathcal{G}_j^ℓ estimated on the data matrix U_ℓ by δ_j^ℓ , $\ell = 1, \dots, 18$ and $j = 2, \dots, 9$. Note that we do not select a threshold dimension upfront but take it as generated by the estimation process of the `huge` package. The corresponding values of λ calculated for each of the 18 different data matrices behave

Data matrix	Dimension d_ℓ	δ_2^ℓ	δ_3^ℓ	δ_4^ℓ	δ_5^ℓ	δ_6^ℓ	δ_7^ℓ	δ_8^ℓ	δ_9^ℓ
U_1	50	3	3	4	8	10	15	18	22
U_2	150	9	17	20	36	46	53	57	63
U_3	250	18	29	36	49	61	75	84	100
U_4	350	26	40	52	71	90	105	121	172
U_5	450	24	34	64	86	107	134	152	205
U_6	550	17	38	67	102	127	148	175	203
U_7	650	20	43	84	111	154	181	210	307
U_8	750	20	43	105	144	187	212	248	359
U_9	850	21	55	75	160	197	239	278	406
U_{10}	950	22	59	82	176	220	265	307	452
U_{11}	1050	12	50	88	106	224	254	435	494
U_{12}	1150	12	54	93	112	238	270	458	530
U_{13}	1250	12	54	93	112	238	270	470	545
U_{14}	1350	12	54	92	113	239	272	483	562
U_{15}	1450	3	49	81	108	224	267	465	570
U_{16}	1550	3	49	81	108	228	271	477	618
U_{17}	1650	3	49	81	108	228	272	490	654
U_{18}	1750	3	49	81	108	228	272	439	680

Table 2: Maximal component sizes $\delta_2^\ell, \dots, \delta_9^\ell$ of graphs $\mathcal{G}_2^\ell, \dots, \mathcal{G}_9^\ell$, for data matrices U_ℓ , see (4.1) for each dimension d_ℓ , $\ell = 1, \dots, 18$.

quite similar, see Figure 8. We omit the first graph \mathcal{G}_1 since it is always an empty graph, and hence, there are no connected components and only isolated nodes. For the fit based on the associated graphs, we evaluate the fill levels $k_F = 0, 1, 5$. Note however that we have to be cautious here. If $k_F = 1$, by our algorithm, for the first tree, the worst case




Figure 8: Comparison of $\lambda_1, \dots, \lambda_{18}$ for U_1, \dots, U_{18} as calculated by the `huge` package for path length $J = 30$.

effort can only be bounded by d^2 as no independence information is used, which may take some time. It is hence more computationally feasible to only fit for example one-parametric pair copula families *outside* the connected components to ease computational effort. Our findings are given in Figure 9, describing the connection between dimension of the sample and computation time, log-likelihood, BIC and GIC. Here, we denote our algorithm by `Cluster-KF` with the corresponding fill-levels $k_F = 0, 1, 5$, dashed, dotted and dot-dashed, respectively. Note that Dißmann’s algorithm in the BIC plot is below all other plots.

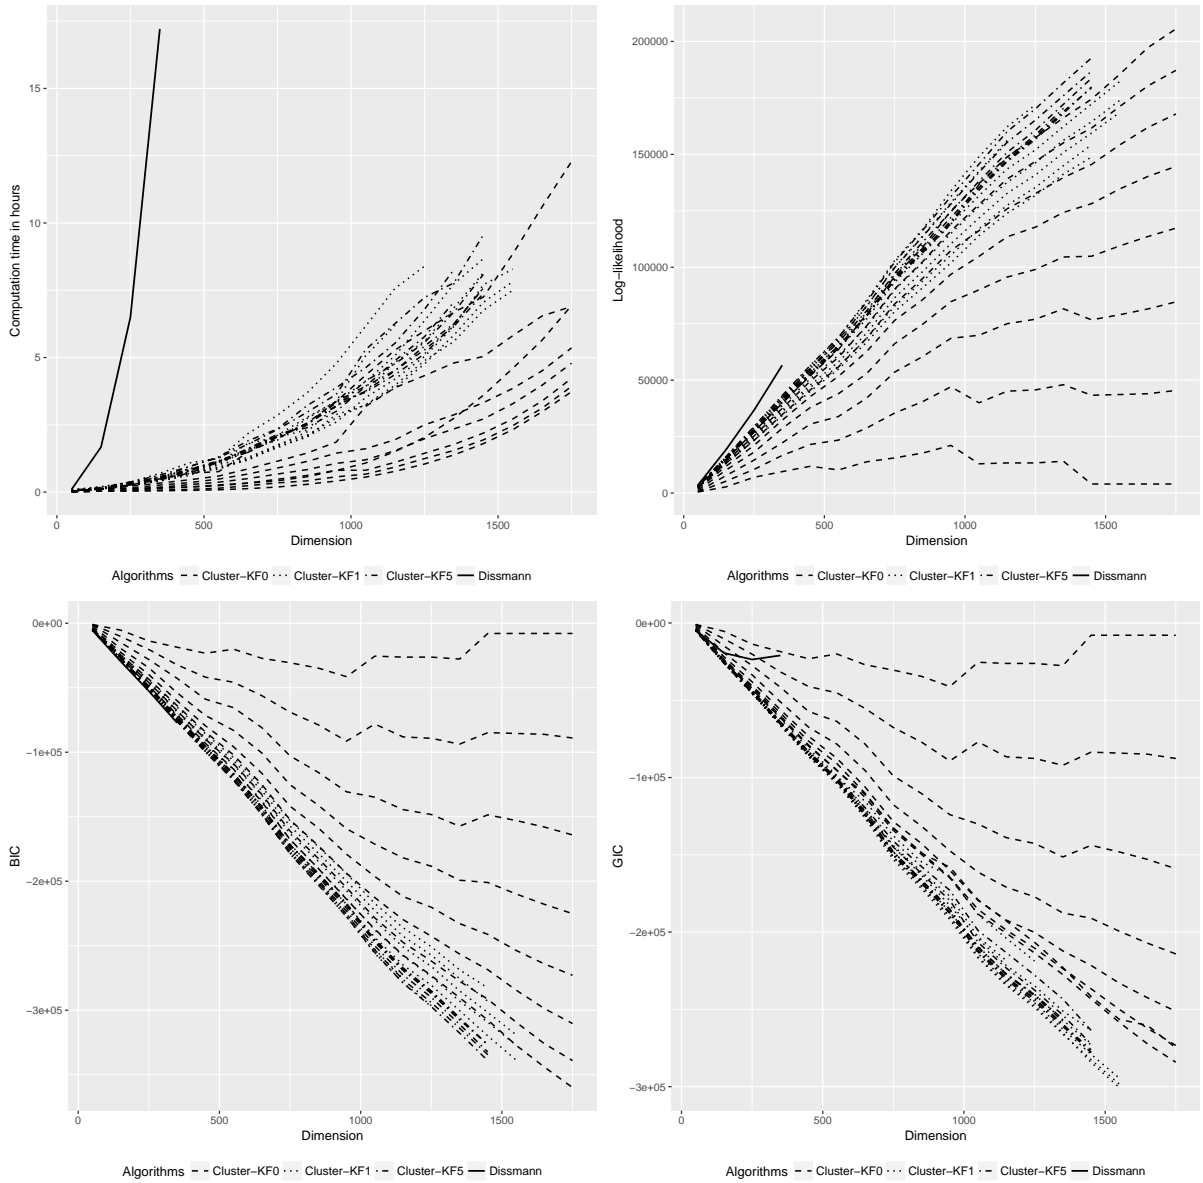


Figure 9: Comparison of `RVineClusterSelect` algorithm for $\mathcal{G}_2^\ell, \dots, \mathcal{G}_9^\ell$ with $k_F = 0, 1, 5$ and Dißmann's algorithm on d dimensions, $d = 50, 150, \dots, 1750$: computation time (upper left), log-likelihood (upper right), BIC (lower left), GIC (lower right).

In total, we have one line for Dißmann's algorithm and eight lines for our approach with different maximal cluster sizes as displayed in Table 2. We perform all the calculations on a Linux Cluster with 32 cores. We see that for computation time, Dißmann's algorithm computation time grows very rapidly once dimension exceeds $d > 100$. There are only 4 data points since for the fifth estimation with $d = 450$, we could not fit any model since we got a time out after 6 days. Additionally, memory consumption becomes also a burden in these dimensions. For our proposed approach, the slope is much less steep and fitting a model in the full size of $d = 1750$ only takes about 4 to 5 hours on the Linux cluster with $k_F = 0$. We note that a full maximum likelihood fitting for each pair copula was carried out and not just a mere estimation by inversion of the empirical Kendall's, which

is faster. We report however, that also Dißmann’s algorithm with this faster estimation is not feasible once the dimension exceeds about $d = 400$. Even though log-likelihood is larger for the standard algorithm, we attain equally good results in terms of BIC and even better results in terms of GIC. This is especially beneficial since we expect the overall dependence to decrease as dimension increases. This is because we expect that most intra-sectoral and intra-geographical dependencies are modelled first, and afterwards, we have more conditional independence. Thus, the number of parameters is expected not to grow as much as the dimension, expressed by lower GIC which penalizes complex models more than BIC. For $k_F = 1$ and $k_F = 5$, some data points in high dimensions are missing because of out-of-memory problems, which is related to many pair copula fits for higher fill levels in dimensions $d > 1500$. However, we see that even though computation time is increasing, it is acceptable given the large benefits in terms of log-likelihood and GIC.

5.4 Data Application

If we want to compare the Gaussian modeling approach with the non-Gaussian R-vine approach, we need to set the fill-level $k_F = 0$, i. e. only model the edges present in the Gaussian model by pair copulas and not include additional edges between the connected components. However, we will also include a fit using fill-level $k_F = 1$ in our considerations later on. Secondly, we have to assume the same marginal distributions. This time, we obtained data on $d = 2131$ stocks with $n = 999$ observations. The stocks are based in the US (1866), Sweden (154) and Australia (111). The following industry sectors are covered, see Table 3. Define S_i the i -th stock for $i = 1, \dots, 2131$, denote the industry

ID	Sector Description	Total	USA	Sweden	Australia
1	Materials and Energy	218	209	3	6
2	Industry Conglomerates	4	4	0	0
3	Consumer Staples	205	189	12	4
4	Financial Services and Real Estate	524	449	49	26
5	Healthcare and Chemicals	232	182	26	24
6	Manufacturing, Industrials and Defense	195	173	12	10
7	Business Management and Services	325	295	21	9
8	IT, Telecommunication and Software	360	301	28	31
9	Utilities	68	64	3	1
		2131	1866	154	111

Table 3: Distribution of stocks over industry sectors and geographies in data application.

sector assignment by

$$\mathcal{I}_\ell : \{i : \text{Stock } S_i \text{ belongs to industry } \ell\} \text{ for } \ell = 1, \dots, 9,$$

and the node set $V = \{1, \dots, 2131\}$ for all graphical models $\mathcal{G}_1, \dots, \mathcal{G}_{30}$. We estimate a path of 30 graphical Lasso solutions $\mathcal{G}_1, \dots, \mathcal{G}_{30}$ calculated by the huge R-package without setting a threshold dimension. To obtain also the log-likelihoods, we evaluate the

corresponding covariance matrices of the solutions. We report the first ten corresponding maximum connected component sizes of $\mathcal{G}_1, \dots, \mathcal{G}_{10}$ in Table 4.

δ_1	δ_2	δ_3	δ_4	δ_5	δ_6	δ_7	δ_8	δ_9	δ_{10}
1	2	42	84	128	215	453	804	1027	1231

Table 4: Maximal connected component sizes for $\mathcal{G}_1, \dots, \mathcal{G}_{10}$.

The remaining 20 cluster sizes for $\mathcal{G}_{11}, \dots, \mathcal{G}_{30}$ are significantly higher than 1200 and thus not considered for our method, as these subsets are too large to fit an R-vine model onto them. In fact, we only fit R-vines on the graphs $\mathcal{G}_2, \dots, \mathcal{G}_9$ since the first partition with $\delta_1 = 1$ is an empty graph bearing no information and the largest \mathcal{G}_{10} with $\delta_{10} = 1231$ is also too large. Hence, we consider 8 graphical models $\mathcal{G}_2, \dots, \mathcal{G}_9$ in the following. An interesting property we can observe is the *industry sector concentration* within each connected component of the graphical models $\mathcal{G}_2, \dots, \mathcal{G}_9$. In the notation of Section 4.1, consider graphical models $\mathcal{G}_j = (V^j = \bigcup_{i=1}^{p_j} V_i^j, \bigcup_{i=1}^{p_j} E_i^j)$ with for $j = 2, \dots, 9$. Let us fix j and consider the k -th component V_k^j of V^j . Then, denote $|V_k^j| = \nu_k^j$, i.e. ν_k^j stocks are contained in this connected component. Recall that each node represents a stock associated to one industry sector. We define the industry sector occurring most often in the component V_k^j by

$$b_k^j = \arg \max_{\ell \in \{1, \dots, 9\}} \sum_{i \in V_k^j} \mathbb{1}_{\{i \in \mathcal{I}_\ell\}}(i).$$

A natural measure for the sector concentration in V_k^j is given by

$$\varrho_k^j = b_k^j / \nu_k^j. \quad (5.1)$$

In other words, we count the occurrence of all different sectors within each connected component and divide the number of the most often occurring sector per connected component by the total number of nodes in this connected component. We do this for all connected components within each of the graphs in the sequence $\mathcal{G}_2, \dots, \mathcal{G}_9$. The results, see Table 5, demonstrate a very high sector concentration in the connected components over the entire sequence of considered graphs. We see first of all, that a large portion of connected components have sector concentration $\varrho_j = 1$, second and third column of Table 5. Thus, at least 80% of the connected components are dominated by a single sector. Additionally, we see from the last column that also the connected components where more than one sector are present, have quite a high sector concentration. This backs our assumption that the graphical Lasso works very well to isolate highly dependent subsets from each other. The motivation for this is the assumption that especially in high dimensions, there exist *clusters of dependence*. As we have e.g. geographical or industry-sectoral dependency in a high dimensional stocks dataset, we can expect the model to make use of *conditional independences* or, in other words, *sparsity*. The idea is that after all intra-geographical or

\mathcal{G}_j	p_j	$ \{k \in 1, \dots, p_j : \varrho_k^j = 1\} $	$ \{k \in 1, \dots, p_j : \varrho_k^j = 1\} $ $/p_j$	mean of $\{\varrho_k^j : \varrho_k^j \neq 1\}$
\mathcal{G}_2	3	3	1	-
\mathcal{G}_3	20	19	0.95	0.83
\mathcal{G}_4	42	39	0.93	0.64
\mathcal{G}_5	67	62	0.93	0.74
\mathcal{G}_6	69	60	0.87	0.66
\mathcal{G}_7	61	46	0.75	0.69
\mathcal{G}_8	47	41	0.87	0.55
\mathcal{G}_9	31	26	0.84	0.58

Table 5: Number of connected components p_j in \mathcal{G}_j for $j = 2, \dots, 9$, number of connected components $k = 1, \dots, p_j$ with sector concentration $\varrho_k^j = 1$, percentage of connected components with sector concentration $\varrho_k^j = 1$ compared to all connected components and mean of the sector concentrations ϱ_k^j over the remaining connected components where sector concentration $\varrho_k^j \neq 1$.

sectoral dependencies are described, cross geographical or sectoral dependencies are weak and can be neglected outside the connected components. The same is often monitored for biological data, where only a small set of genes is affected by each other.

Finally, we compare the Gaussian model and the R-vine model in terms of log-likelihood and GIC on the z -scale for corresponding numbers of parameters, see Figure 10. In the Gaussian model, the number of parameters is equal to the numbers of edges in the graph whereas for the R-vine, the number of parameters is equal to the number of pair copula parameters. In both cases, we add d parameters for the estimated variance of the marginal distributions. This is because we need to add Gaussian marginals with the same variance as estimated by the graphical Lasso estimate to our R-vine model to make it comparable to the Gaussian model estimated with the graphical Lasso. We clearly see that the flexibility of the R-vine compared to the Gaussian model leads to significant out-performance with respect to log-likelihood and even more with respect to GIC. This stems from the fact that the R-vine gains much more exploratory power with adding additional parameters. This also true for the corresponding model with $k_F = 1$. The most parsimonious model in terms of GIC is the R-vine given by \mathcal{G}_8 . This is also similar for the Gaussian model.

In terms of pair copulas, the most prominently present families in \mathcal{G}_8 are *Frank-copulas* (5437, 39.9 %) and *Student's-t copulas* (1889, 13.8 %) before *Gaussian* (1581, 11.6 %), which is a clear indicator of non-Gaussianity and also tail dependence. When considering the Student's-t copulas in the R-vine and their parameters, i. e. the degrees of freedom df , we have that for lower df , we have more tail dependence and for higher $df > 30$, the distribution becomes quite similar to the Gaussian distribution. We consider boxplots of the degrees of freedom over the R-vines computed with the graphs $\mathcal{G}_3, \dots, \mathcal{G}_9$, see Figure 11. Note that in the R-vine computed based on \mathcal{G}_2 we do not have any Student's-t copulas. Since most of the degrees of freedom vary around 10 to 15, we clearly monitor and model heavy tailed data, which is quite important considering financial returns. If this property




Figure 10: Comparison of clustered R-vines with fill-levels $k_F = 0, 1$ and Gaussian Graphical Model fitted with the Graphical Lasso $\mathcal{G}_1, \dots, \mathcal{G}_9$: log-likelihood (left), GIC (lower right).

is not adequately taken into account, risk models are deemed to fail in times of heavy market turmoil when assets become highly correlated, leading to a significant reduction of diversification when it is most needed. Finally, our models based on the graphs $\mathcal{G}_2, \dots, \mathcal{G}_9$ took between 12 hours and 2 days for estimation on a Linux Cluster with 32 cores. The Gaussian graphical model needs only several minutes for estimation in these dimensions. It is worth noting that other methods for estimating R-vines as the one of Dißmann et al. (2013) failed in these dimensions because of memory consumption or time out.

6 Discussion

We developed a method to estimate ultra high dimensional vine copulas using a novel divide-and-conquer approach inspired by the graphical Lasso. The main idea is to exploit conditional independence for isolated consideration of sub problems of significantly lower dimension compared to the original data. Furthermore, we proposed using graphical independence models also in frameworks of moderate dimension to increase the estimation accuracy of standard vine copula estimation procedures significantly. In all the scenarios, our approach works very well in terms of computation times and penalized goodness of fit measures as GIC, which targets especially sparse datasets. At last, we showed that our approach is of several magnitudes faster than competing methods, however allowing for non-Gaussian dependence. We see further improvement for our contribution with respect to the selection of the corresponding solution of the graphical Lasso, which is the basis for the vine copula model. Currently, we evaluate all these solutions over a grid of penalty values. However, also using criteria as *StARS* (Liu et al., 2010) seems feasible and




Figure 11: Distribution of the degrees of freedom as parameter for the Student’s-t copula.

worth exploring. Additionally, a more sophisticated value of the currently implemented fill level $k_F = \lceil \log(d) \rceil$ may be considered, since the current choice is only a rule of thumb motivated by numerical experiments such as presented in Section 5. Finally, a method also worth exploring is to use an *iterated divide and conquer* approach where larger connected components as in $\mathcal{G}_{10}, \dots, \mathcal{G}_{30}$ of the data application can again be divided until all they are all below a given machine-induced threshold. We will consider these aspects as well as the search for applications to non-financial ultra high dimensional datasets in the future.

Acknowledgement

The first author is thankful for a research stipend of the Technische Universität München. The second author is supported by the German Research foundation (DFG grant CZ 86/4-1). Numerical computations were performed on a Linux cluster supported by DFG grant INST 95/919-1 FUGG. The *VineCopula* R-package (Schepsmeier et al., 2017), on which our code is based, is greatly acknowledged.

References

- Aas, K. (2016). Pair-copula constructions for financial applications: A review. *Econometrics* 4(4).
- Aas, K., C. Czado, A. Frigessi, and H. Bakken (2009). Pair-copula constructions of multiple dependence. *Insurance, Mathematics and Economics* 44, 182–198.
- Arnborg, S. and P. A. Corneil, Derek G. (1987). Complexity of finding embeddings in a k-tree. *SIAM Journal on Algebraic and Discrete Methods* 8(2), 277–284.

- Bedford, T. and R. Cooke (2001). Probability density decomposition for conditionally dependent random variables modeled by vines. *Annals of Mathematics and Artificial Intelligence* 32, 245–268.
- Bedford, T. and R. Cooke (2002). Vines - a new graphical model for dependent random variables. *The Annals of Statistics* 30(4), 1031–1068.
- Brechmann, E., C. Czado, and K. Aas (2012). Truncated regular vines in high dimensions with application to financial data. *Canadian Journal of Statistics* 40, 68–85.
- Czado, C., S. Jeske, and M. Hofmann (2013). Selection strategies for regular vine copulae. *Journal de la Société Française de Statistique* 154, 174–191.
- Dempster, A. P. (1972). Covariance selection. *Biometrics* 28(1), 157–175.
- Dißmann, J., E. Brechmann, C. Czado, and D. Kurowicka (2013). Selecting and estimating regular vine copulae and application to financial returns. *Computational Statistics and Data Analysis* 52(1), 52–59.
- Fan, Y. and C. Y. Tang (2013). Tuning parameter selection in high dimensional penalized likelihood. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 75(3), 531–552.
- Friedman, J., T. Hastie, and R. Tibshirani (2008). Sparse inverse covariance estimation with the graphical lasso. *Biostatistics* 9(3), 432.
- Ghalanos, A. (2015). *rugarch: Univariate GARCH models*. R package version 1.3-6.
- Gruber, L. and C. Czado (2015a). Bayesian model selection of regular vine copulas. *Preprint*.
- Gruber, L. and C. Czado (2015b). Sequential bayesian model selection of regular vine copulas. *Bayesian Analysis* 10, 937–963.
- Hobæk Haff, I., K. Aas, A. Frigessi, and V. L. Graziani (2016). Structure learning in Bayesian Networks using regular vines. *Computational Statistics and Data Analysis* 101, 186–208.
- Joe, H. (1996). Families of m -variate distributions with given margins and $m(m - 1)/2$ bivariate dependence parameters. In L. Rüschendorf, B. Schweizer, and M. D. Taylor (Eds.), *Distributions with fixed marginals and related topics*, pp. 120–141. Hayward: Institute of Mathematical Statistics.
- Kovács, E. and T. Szántai (2016). On the connection between cherry-tree copulas and truncated R-vine copulas. *arXiv preprint arXiv:1604.03269*.
- Kraus, D. and C. Czado (2017). Growing simplified vine copula trees: improving dißmann’s algorithm. *arXiv preprint arXiv:1703.05203*.

- Krumsiek, J., K. Suhre, T. Illig, J. Adamski, and F. J. Theis (2011). Gaussian graphical modeling reconstructs pathway reactions from high-throughput metabolomics data. *BMC systems biology* 5(1), 21.
- Kurowicka, D. and H. Joe (2011). *Dependence Modeling - Handbook on Vine Copulae*. Singapore: World Scientific Publishing Co.
- Lauritzen, S. L. (1996). *Graphical Models* (1st ed.). Oxford, England: University Press.
- Liu, H., F. Han, M. Yuan, J. Lafferty, and L. Wasserman (2012, 08). High-dimensional semiparametric gaussian copula graphical models. *Ann. Statist.* 40(4), 2293–2326.
- Liu, H., K. Roeder, and L. Wasserman (2010). Stability approach to regularization selection (stars) for high dimensional graphical models. In *Proceedings of the 23rd International Conference on Neural Information Processing Systems, NIPS'10, USA*, pp. 1432–1440. Curran Associates Inc.
- Mazumder, R. and T. Hastie (2012). Exact Covariance Thresholding into Connected Components for Large Scale Graphical Lasso. *Journal of Machine Learning Research* 13, 723–736.
- McNeil, A. J., R. Frey, and P. Embrechts (2006). *Quantitative Risk Management: Concepts, Techniques and Tools*. Princeton University Press.
- Meinshausen, N. and P. Bühlmann (2006, 06). High-dimensional graphs and variable selection with the Lasso. *The Annals of Statistics* 34(3), 1436–1462.
- Müller, D. and C. Czado (2017a). Representing Sparse Gaussian DAGs as Sparse R-vines Allowing for Non-Gaussian Dependence. *To appear in the Journal of Computational and Graphical Statistics*.
- Müller, D. and C. Czado (2017b). Selection of Sparse Vine Copulas in High Dimensions with the Lasso. *arXiv preprint arXiv:1705.05877*.
- Prim, R. C. (1957). Shortest connection networks and some generalizations. *Bell System Technical Journal* 36, 1389–1401.
- Ryan, J. A. and J. M. Ulrich (2017). *quantmod: Quantitative Financial Modelling Framework*. R package version 0.4-10.
- Schepsmeier, U., J. Stöber, E. C. Brechmann, B. Graeler, T. Nagler, and T. Erhardt (2017). *VineCopula: Statistical Inference of Vine Copulas*. R package version 2.1.2.
- Schwarz, G. (1978). Estimating the dimension of a model. *The Annals of Statistics* 6(2), 461–464.
- Sklar, A. (1959). Fonctions de répartition à n dimensions et leurs marges. *Publ. Inst. Stat. Univ. Paris* 8, 229–231.

- Stöber, J., H. Joe, and C. Czado (2013). Simplified pair copula constructions-limitations and extensions. *Journal of Multivariate Analysis* 119(0), 101 – 118.
- Tibshirani, R. (1994). Regression Shrinkage and Selection Via the Lasso. *Journal of the Royal Statistical Society, Series B* 58, 267–288.
- Toh, H. and K. Horimoto (2002). Inference of a genetic network by a combined approach of cluster analysis and graphical gaussian modeling. *Bioinformatics* 18(2), 287–297.
- Witten, D. M., J. H. Friedman, and S. Noah (2011). New Insights and Faster Computations for the Graphical Lasso. *Journal of Computational and Graphical Statistics* 20(4), 892–900.
- Zhao, T., X. Li, H. Liu, K. Roeder, J. Lafferty, and L. Wasserman (2015). *huge: High-Dimensional Undirected Graph Estimation*. R package version 1.2.7.

A Algorithm RVineClusterSelect

```

input : Data  $\mathbf{X} \in \mathbb{R}^{n \times d}$ ,  $d_T \leq d$ ,  $k_F < d$ .
output: R-Vine in  $d$  dimensions.

1 calculate  $(\mathcal{G}_1 = (\bigcup_{i=1}^{p_1} V_i^1, \bigcup_{i=1}^{p_1} E_i^1), \dots, \mathcal{G}_J = (\bigcup_{i=1}^{p_J} V_i^J, \bigcup_{i=1}^{p_J} E_i^J))$ ;
2 select  $\mathcal{G}_T$  such that  $T = \arg \max_{j=1, \dots, J} \delta_j \leq d_T$  with  $\delta_j = \max_{i=1, \dots, p_i} |V_i^j|$ ;
3 for  $i = 1$  to  $p_T$  do
4   | set  $\nu_i = |V_i^T|$ ; // dimension of connected component i
5   | for  $(j, \ell) \in V_i^T$  do
6   |   | if  $(j, \ell) \notin E_i^T$  then
7   |   |   | set  $c_{j, \ell} = 1$ ;
8   |   |   | else
9   |   |   |   | estimate pair-copula  $c_{j, \ell}$ ;
10  |   |   | end
11  | end
12  | calculate weights  $\mu(j, \ell) = \mu(c_{j, \ell})$ ; // e. g. Log-Lik., AIC, BIC
13  | calculate maximum spanning tree  $T_1 = (V_1 = V_i, E_1)$  w.r.t.  $\mu$ ;
14  | for  $k = 2$  to  $\nu_i - 1$  do
15  |   | set  $V_k = E_{k-1}$  and  $E_k$  admissible edges by proximity condition;
16  |   | for  $(j, \ell | \mathbf{D}) \in E_k$  do
17  |   |   | if  $j \perp \ell | \mathbf{D} [\mathcal{G}_T]$  then
18  |   |   |   | set  $c_{j, \ell | \mathbf{D}} = 1$ ;
19  |   |   |   | else
20  |   |   |   |   | estimate pair-copula  $c_{j, \ell | \mathbf{D}}$ ;
21  |   |   |   | end
22  |   |   | end
23  |   |   | calculate weights  $\mu(j, \ell | \mathbf{D}) = \mu(c_{j, \ell | \mathbf{D}})$ ;
24  |   |   | calculate maximum spanning tree  $T_k = (V_k, E_k)$  w.r.t.  $\mu$ ;
25  |   | end
26 end
27 create R-vine matrix  $M$  in  $d$  dimensions combining sub-R-vines;
28 for  $k = 1$  to  $k_F$  do
29 | estimate pair-copulas between sub-R-vines in tree  $k$  of in R-vine matrix  $M$ ;
30 end

```

Algorithm 1: RVineClusterSelect: Selection of an R-vine in d dimensions.

B Additional Results for Simulation Study

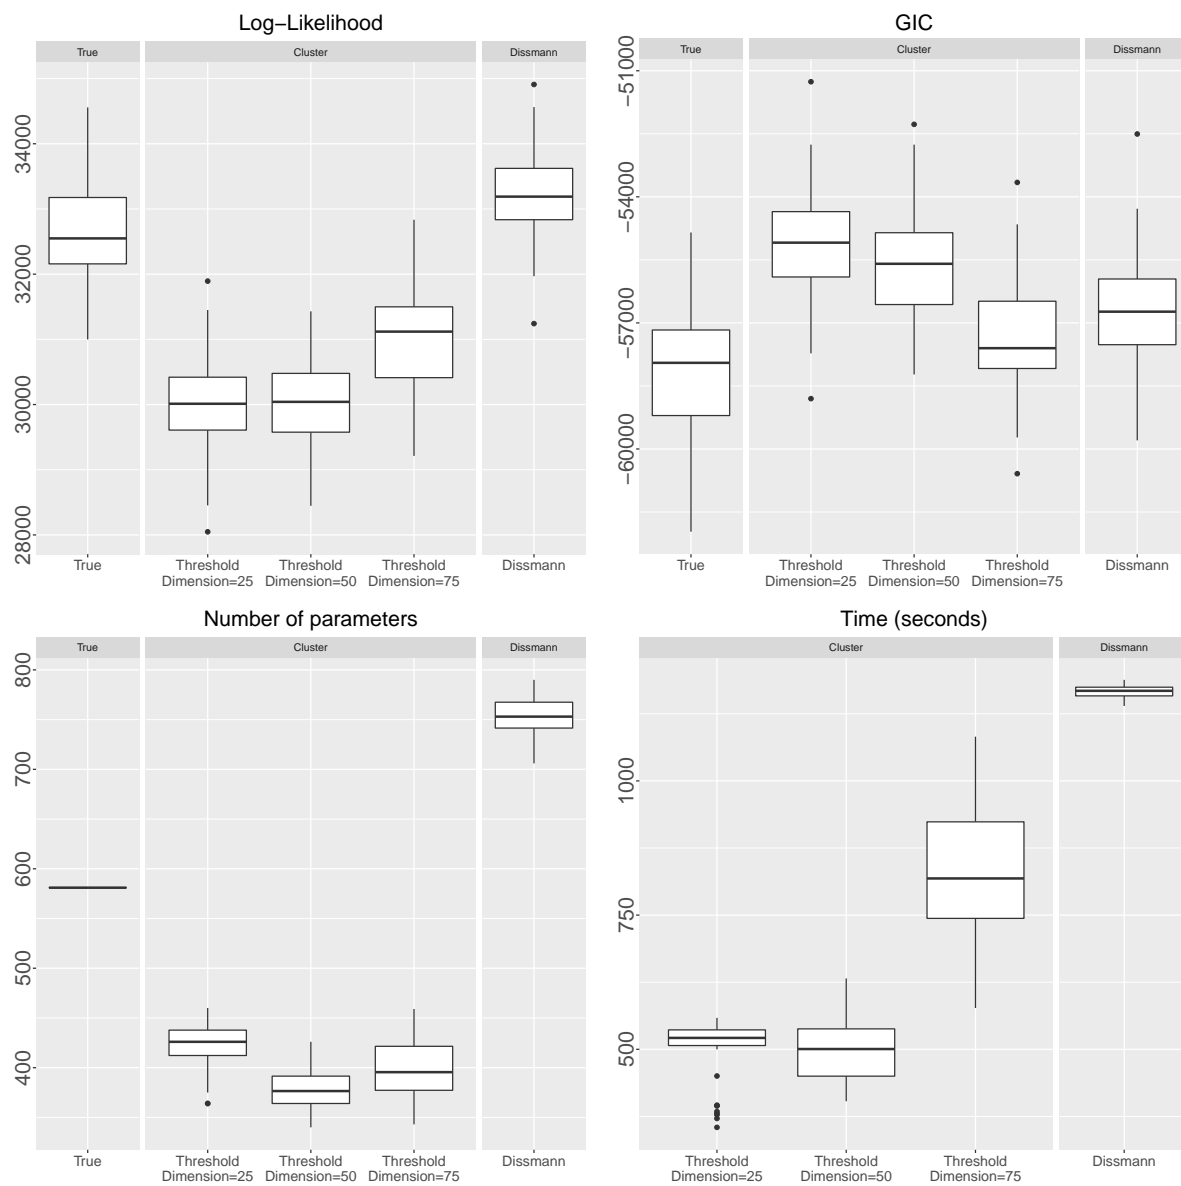


Figure 12: V_2 : Comparison of `RVineClusterSelect` algorithm with threshold dimension $d_T = 25, 50, 75$ and Dissmann's algorithm on u -scale: log-likelihood (upper left), GIC (upper right), number of parameters (lower left) and computation time (lower right).

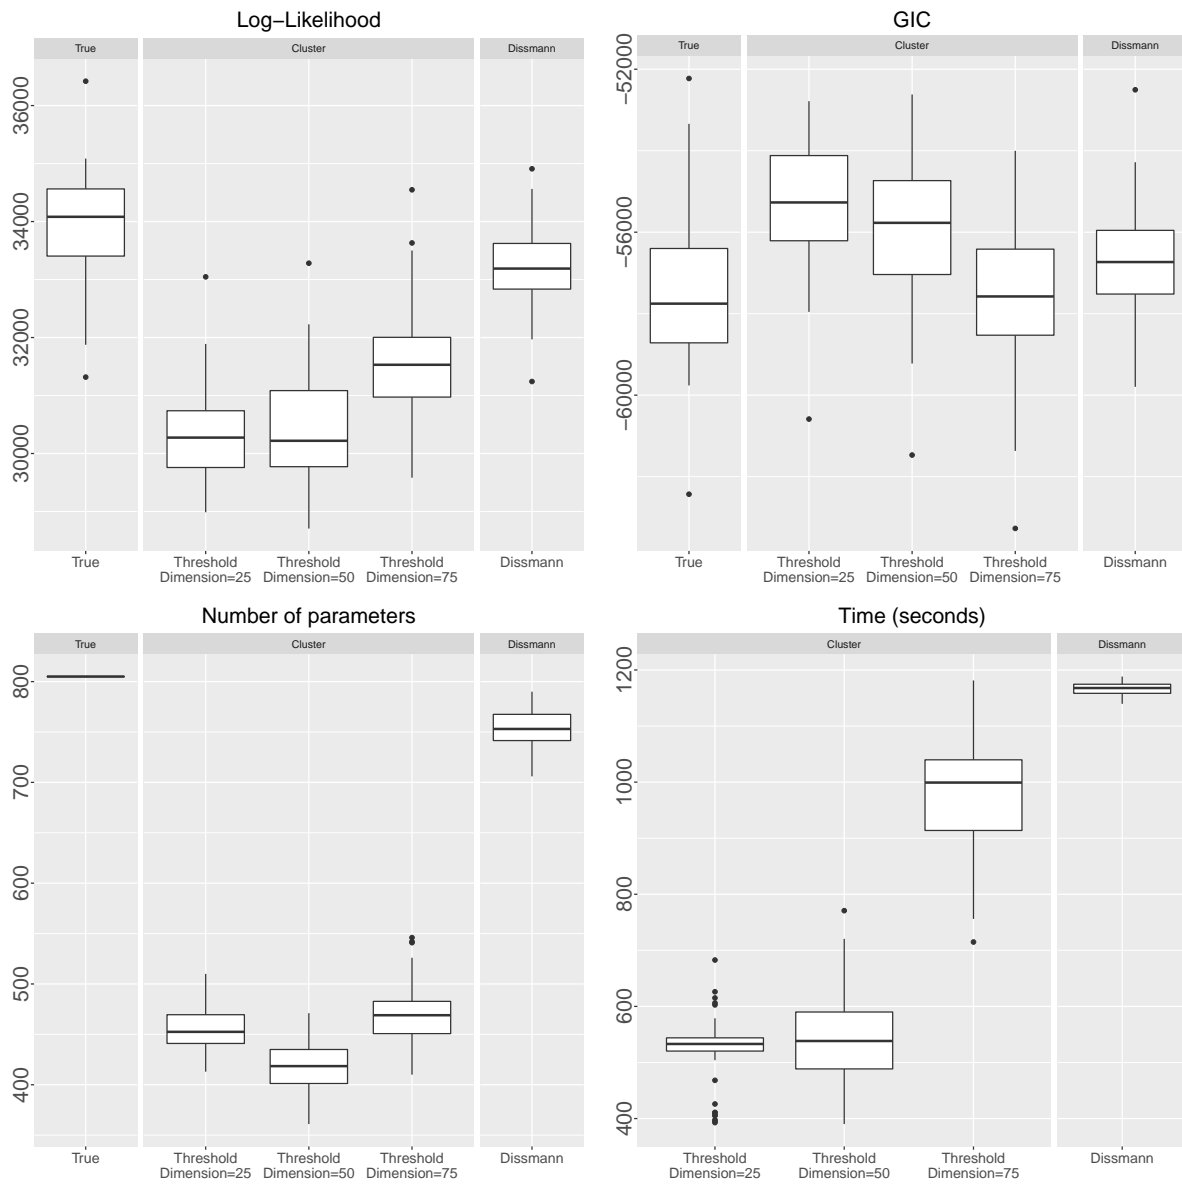


Figure 13: V_3 : Comparison of `RVineClusterSelect` algorithm with threshold dimension $d_T = 25, 50, 75$ and Dissmann's algorithm on u -scale: log-likelihood (upper left), GIC (upper right), number of parameters (lower left) and computation time (lower right).