



Invited Review

Stochastic dynamic vehicle routing in the light of prescriptive analytics: A review

Ninja Soeffker^a, Marlin W. Ulmer^b, Dirk C. Mattfeld^{c,*}^a University of Vienna, Department of Business Decisions and Analytics, Kolingasse 14-16, 1090 Vienna, Austria^b Technische Universität München, School of Management, Arcisstraße 21, 80333 München, Germany^c Technische Universität Braunschweig, Decision Support Group, Mühlenpfordtstraße 23, 38106 Braunschweig, Germany

ARTICLE INFO

Article history:

Received 20 March 2020

Accepted 6 July 2021

Available online 16 July 2021

Keywords:

Routing

Survey

Stochastic dynamic vehicle routing

Prescriptive analytics

ABSTRACT

Stochastic dynamic vehicle routing problems have become an essential part of logistics and mobility services. In such problems, a sequence of vehicle routing decisions has to be made in reaction and anticipation of newly revealed stochastic information. To this end, a variety of computational operations research methods has emerged in the literature, increasingly integrating potential future information in their decision making. The integration of information models into decision models via computational methods is known as prescriptive analytics, the most recent advance of business analytics. In this paper, we explore the existing work and future potential of prescriptive analytics for stochastic dynamic vehicle routing. We identify the characteristics of decision models and information models unique in stochastic dynamic vehicle routing and analyze how different methodology meets the characteristics' requirements. We use the insights to derive recommendations about promising methodology when approaching specific stochastic dynamic vehicle routing problems.

© 2021 The Author(s). Published by Elsevier B.V.

This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>)

1. Introduction

Logistics and mobility services play a major role in all our lives. Many of us use mobility services such as taxis, dial-a-ride, or ride-sharing regularly and whenever we want. We order food or goods online and expect near-instant gratification a short time after our order is placed. And we order services to our homes, technicians repairing our internet router or nurses caring for our health.

In all these cases, vehicle routing is required to fulfill the customer demand with fleet resources, usually in an urban environment. In many cases, the information in the three dimensions is not static and deterministic, but changes over time, e.g., more customers requesting service, streets congest unexpectedly, or even the fleet resources change, for example, due to crowdsourced drivers entering or leaving the system spontaneously. Such frequent, disruptive, and stochastic information changes force providers to dynamically renew their planning. However, mere reactions to information changes in one or more of the three dimensions are insufficient. There is limited benefit in changing a driver's route when the vehicle is already stuck in traffic, leaving

a ride-sharing vehicle idling in a far distant corner of the city until a new customer demands service, or calling drivers for support when lunch delivery orders are rolling in and no crowdsourced driver is around for delivery. Instead, flexible, anticipatory planning is required, preferably utilizing the vast amounts of data that companies have collected over the years. The incorporation of predictive information models into decision making is known as prescriptive analytics (PA), the most recent advance in the business analytics-sequence of descriptive, predictive, and prescriptive analytics (Lepenioti, Bousdekis, Apostolou, & Mentzas, 2020).

Business analytics (BA) has received tremendous attention in academia and business (Conboy, Mikalef, Dennehy, & Krogstie, 2020; Hindle, Kunc, Mortensen, Oztekin, & Vidgen, 2019; Holsapple, Lee-Post, & Pakath, 2014). BA operates in the intersection between computational technologies or computer science, decision making in management, and quantitative mathematical methods (compare Fig. 1 in Mortenson, Doherty, & Robinson, 2015). The goal of BA is to analyze available data, to derive predictive information models, and to use the models for decision making. The first step, descriptive analytics, derives insights and patterns from past observations (e.g., customer demand for delivery services). The second step, predictive analytics, uses the insights and patterns to develop an information model that allows predictions of future developments, for example, the likelihood customers request service dur-

* Corresponding author.

E-mail addresses: ninja.soeffker@univie.ac.at (N. Soeffker), marlin.ulmer@tum.de (M.W. Ulmer), d.mattfeld@tu-braunschweig.de (D.C. Mattfeld).

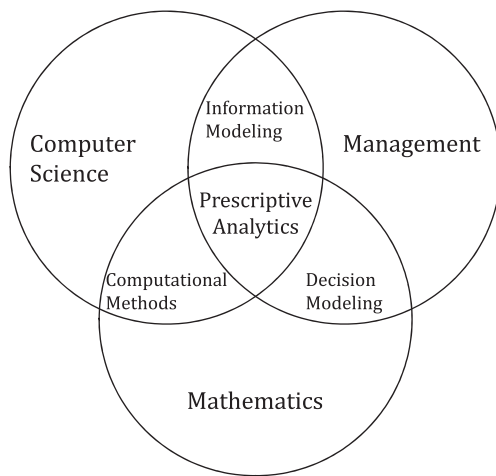


Fig. 1. Prescriptive analytics, adapted from Mortenson et al. (2015).

ing the day and over the service area. The final step, PA, integrates the information model into decision making, for example, evaluating assignment and routing decisions with respect to potential future demand. To define PA in detail, we specify the depiction of Mortenson et al. (2015) in Fig. 1, positioning PA in the interface of computer science, mathematics, and management. More specific, PA combines information modeling with decision modeling and computational methods in a holistic way. The information models capture potential information changes and are provided by predictive analytics, the decision models define the management problem via mathematical modeling, and the computational methods derive decisions with respect to decision and information models. The holistic view provided by prescriptive analytics allows to choose computational methods that fit the given decision and information models. That is, PA provides methods tailored to properties and structure of the application's information and decision models (Lustig, Dietrich, Johnson, & Dziekan, 2010).

While logistics and mobility services might benefit from PA tremendously, work on PA in stochastic dynamic vehicle routing problems (SDVRPs) is still very limited (Ulmer, Goodson, Mattfeld, & Thomas, 2020b). Solving the decision models of vehicle routing problems is by itself challenging, even without changes in the available information. The stochasticity and dynamism add to the existing challenges since decisions now have to consider future changes in the information and decision model. To cope with the challenges, researchers develop problem-specific computational methods based on intuition and experience. Consequently, there is no “standard” PA-procedure to approach SDVRPs, let alone, commercial software. The foundation for such standard procedures is to gather and classify the existing work and experience and to derive insights in the performance of computational methods with respect to decision and information models. Reviews provided by Psaraftis, Wen, & Kontovas (2016) and Ulmer (2017) present classifications of SDVRP literature with a focus on the considered problem and the applied computational methods. Powell (2019) presents a general overview on computational methods in stochastic optimization with a strong connection to the modeling of uncertainty. In our work, we aim to present a holistic PA-view that incorporates the interplay of problem and model characteristics with methodology characteristics. Thus, our paper discusses SDVRPs and their characteristics as presented in Ulmer (2017) with the view and language of Fig. 1 and the available models and methodology of Powell (2019). To this end, this paper takes the following steps, in accordance to Fig. 1:

1. It gives a short overview on the main sources of uncertainty in SDVRP and their characteristics in Section 2.

2. It presents information models, decision models, and their interrelation in Section 3.
3. It describes and classifies the existing prescriptive methodology with respect to their exploitation of the information model in Section 4.
4. It uses the insights of information modeling, decision modeling, and computational methods for a (qualitative) recommendation when selecting a prescriptive method for a specific problem in Section 5.

In this process, we suggest that there is no “dominant” method to approach every SDVRP. Instead, the suitability of a prescriptive method depends on the characteristics of information and decision model (compare, e.g., Spivey & Powell, 2004).

2. Stochastic dynamic vehicle routing problems

In this section, we briefly discuss the main dimensions of SDVRPs. We spend a particular focus on the uncertainty observed by the service providers which forms the basis for the first two BA-steps, descriptive and predictive analytics. We also introduce an example of a stochastic dynamic vehicle routing problem that we use throughout the article for illustration purposes.

2.1. Dimensions

Applications areas for SDVRPs are vast (see Appendix A.1, for an overview). In general, application areas of stochastic dynamic vehicle routing problems span the transportation of goods (Ichoua, Gendreau, & Potvin, 2000), the transportation of passengers (Psaraftis, 1980), as well as the conduction of services at customers' homes (Larsen, Madsen, & Solomon, 2002). Depending on the application, different objectives (e.g., cost, service rate) and restrictions are considered, for example, time windows (Gendreau, Guertin, Potvin, & Taillard, 1999), working hours (Angelelli, Bianchessi, Mansini, & Speranza, 2009), or capacities (Sáez, Cortés, & Núñez, 2008).

While the applications may differ, the problem settings share characteristics. Stochastic dynamic vehicle routing problems occur on an operational level in a preset urban environment where resources like drivers and vehicles of the fleet are set. Decisions are made about how to use the resources within the environment to satisfy customer demand which generates some reward for the service provider. The objective function is to maximize the reward generated by satisfying customer demand or to minimize the resource usage required to serve a given demand.

The three dimensions demand, resources, and environment are sketched in Fig. 2. The left box in Fig. 2 represents potential customer demand a service provider faces. Depending on the application area, customers ask for a transportation service, goods, or a service to be conducted at their homes. The right box in the figure depicts the resources of the service provider. These resources are typically vehicles of various types, like bikes or trucks, with different characteristics as well as the staff that operates the vehicles and conducts services. The environment, depicted in the central box in Fig. 2, consists of the road and logistics network as well as additional factors like weather conditions or traffic management decisions about traffic signal strategies or road closures.

All of the described dimensions may be subject to a change in information. This leads to several potential sources of stochasticity which are described in the following.

2.2. Uncertainty

As in static vehicle routing, decisions are made about the assignment of resources to demand and the implementation within




Fig. 2. Dimensions of vehicle routing problems faced by a service provider.

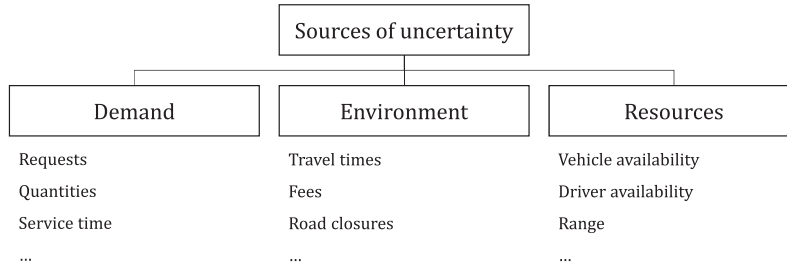


Fig. 3. Sources of uncertainty in stochastic dynamic vehicle routing.

the surrounding environment. However, as we observe uncertain changes in the information, decisions are not made once but are made dynamically over time. Fig. 3 depicts the dimensions, demand, resources, environment, and some typical uncertainties. This uncertainty can be modeled by means of an information model that collects observations of the uncertainty and unites them in a model. To derive an information model, the observations must be analyzed and structured first, a task of descriptive analytics (Lustig et al., 2010). Based on the structures and insights, information models can then be determined via predictive analytics (see Section 3.2). Such an information model is able to depict the dimensionality and heterogeneity of the uncertainty. Also, potential disruptivity and correlation can be described by means of an information model. In the following, we discuss typical sources of uncertainty in the different dimensions, their reasons and characteristics.

Uncertain demand. For the service provider, uncertainty about the customer demand often consists of when and where a customer requests, how much the customer requests, or how long the service will take. Uncertainty in demand may, for example, deal with the occurrence of customer demand for transportation, home services, or the delivery of goods. In the resulting problem settings, vehicles already serve customers while new customer requests arrive (Agussurja, Cheng, & Lau, 2019; Sheridan et al., 2013). This source of uncertainty is a rather common one in times of modern communication and may occur in the three application areas of goods (as in same-day delivery), passengers (as in dial-a-ride), and services (as in technician services). Requests can be rather heterogeneous or less heterogeneous both with respect to their location and request time. In the second typical form, the demand quantity of the customers, for example for gasoline, liquids, or other inventory, only becomes known when the driver arrives at the customer’s location. Here, uncertainty occurs in the form of unknown demand quantities (Novoa & Storer, 2009). This uncertainty can be encountered by service providers offering the transportation of goods or passengers. Uncertain service times especially occur in the case of services conducted at the customer’s home (Larsen et al., 2002).

Uncertain environment. Uncertainties may also affect the environment the service provider operates in. These uncertainties are usually independent of the application. Weather, traffic jams, and other events can impact whether a road can be used and how long it takes to traverse it, i.e., travel times (Schilde, Doerner, & Hartl, 2014). Since travel times are required both for every road

and every time, they are described using multiple dimensions. Especially with the rise of air pollution monitoring, traffic administrations may increase the price for using certain roads by means of tolls for all or certain vehicles. They may even block roads for certain types of vehicles due to high pollution levels (Köster, Ulmer, Mattfeld, & Hasle, 2018). In addition to uncertain travel times, roads may therefore also be affected by uncertain fees such as tolls or by sudden disruptive events like road closures.

Uncertain resources. For some providers, independent of the application area, the resources available to fulfill services are uncertain. Disruptive events like accidents or malfunctions are examples for a stochastic vehicle availability (Xiang, Chu, & Chen, 2008). The vehicle then has to be replaced or repaired. Also, drivers can be absent due to sick leave unexpectedly. Finally, self-employed drivers, also known as occasional or crowdsourced drivers, can decide on their own when they want to work and which jobs they want to fulfill (Dayarian & Savelsbergh, 2020). Stochastic driver availability is therefore another possible form of uncertain resources (Arslan, Agatz, Kroon, & Zuidwijk, 2019). Especially with the rise of electric vehicles, there exists uncertainty about the remaining range, which impacts the availability of the resource vehicle. Such an uncertainty severely impacts the ability to serve customer requests because the possible length of the tours is not known in advance.

2.3. Example: dynamic vehicle routing with stochastic customer requests

In this following, we present an example that we use throughout the paper for the illustration of application, modeling, and methods. We consider a prominent problem from practice and literature, the dynamic vehicle routing problem with stochastic customer requests and therefore uncertainty in demand. The problem described here is an example of a broader class of problems that vary with respect to their focus (Thomas, 2007; Ulmer, Goodson, Mattfeld, & Hennig, 2019). The customer requests could ask for the conduction of a service or for the collection of a parcel. Some requests are known in advance, but other customers request service during the day. A single vehicle is available to serve customer requests in a given time horizon defined by the vehicle driver’s shift. When a new request occurs, the service provider needs to decide whether the new request is accepted or rejected, and in case of acceptance, how to update the vehicle’s ongoing route to include the

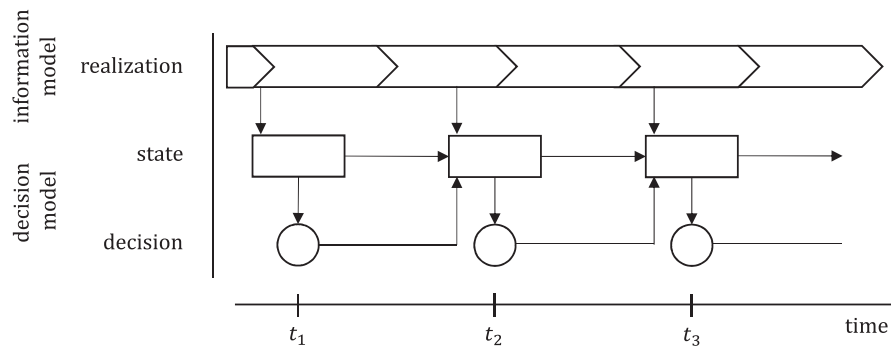


Fig. 4. Stochastic dynamic decision processes, adapted from Meisel (2011).

new request. The service provider aims at serving as many customers over the time horizon as possible.

The structure of the observed uncertainty depends on the application considered. If the service is the collection of express mail, many requests may accumulate in office districts and therefore mainly during office hours. If the service is maintenance or repair instead, some customers may be known beforehand and additional requests are likely to occur more evenly spread over the city and time of day. Thus, for the same problem type, heterogeneity and disruptivity of the uncertainty may vary significantly. In the next section, we will describe how this problem can be modeled and how such differences can be captured by the information models.

3. Modeling

In the following, we describe how we model stochastic dynamic vehicle routing problems as sequential decision processes. The modeling contains three steps, the transfer of the analyzed observed uncertainty to a predictive information model, the definition of the decision model, and the modeling of their interactions. We first describe the general functionality of sequential decision processes as an interaction of decision and information model. We then describe information model and decision model in detail, highlighting their characteristics for SDVRPs.

3.1. Sequential decision process

Stochastic dynamic decision problems can be modeled as sequential decision processes (Powell, 2011). A sequential decision process depicts a sequence of states where in each state a decision model instance is solved based on realizations of the information model. The functionality of a sequential decision process is shown in Fig. 4. We define the components in Fig. 4 as follows:

- The x-axis depicts the evolution of time where decisions are made at *decision points* $k = 1, \dots, K$ with respect to newly revealed information. Parameter K depicts the final decision point and may be a random variable.
- The *information model* Ω describes the uncertain changes over time. In a decision point k , a *realization* ω_k of the information model becomes known.
- The *states* s_k describe the information available to the decision maker when making a *decision*. Each state is an instance of the problem's *decision model*.
- The *decision* x_k represents a solution of the decision model instance. Such a decision usually contributes to the objective function either by providing a reward, $R(s_k, x_k)$, or causing costs (For the sake of simplicity, we focus on a maximization objective in this work). We note that the reward may be an expected value. The combination of decision x_k and information model

realization ω_{k+1} leads to a transition with transition function T to the new state $s_{k+1} = T(s_k, x_k, \omega_{k+1})$ at time t_{k+1} .

A solution of a sequential decision process is a policy π . A policy is a function that maps a state s_k to a decision $x_k = X^\pi(s_k)$. The goal for sequential decision processes is to find an optimal policy π^* maximizing the expected sum of rewards. An optimal policy satisfies the Bellman Equation in each state s_k , maximizing the sum of the immediate reward $R(s_k, x_k)$ and the expected sum over the future rewards given s_k and decision x_k and when applying the optimal policy π^* throughout:

$$X^{\pi^*}(s_k) = \arg \max_{x_k \in X(s_k)} \left\{ R(s_k, x_k) + \mathbb{E} \left[\sum_{j=k+1}^K R(s_j, X^{\pi^*}(s_j)) \mid (s_k, x_k) \right] \right\}. \quad (1)$$

The second term of the Bellman Equation is often called the *value* of either a state–decision pair, $V(s_k, x_k)$ or the post–decision state, $V(s_k^x)$ that follows deterministically when choosing decision x_k in state s_k . Knowing the value function V for every state–decision pair (or post–decision state) allows to make optimal decisions with the Bellman Equation.

The presented sequential decision process is generic and can be used in many different areas. As we will describe in the remainder of this section in detail, the sequential decision processes for SDVRPs show particular characteristics in decision and information models. The information model represents potential changes in the three aforementioned dimensions resources, demand, and environment. The decision model represents problem-specific variants of vehicle routing problems with decisions being (tentative) routing plans. For SDVRPs, there is usually a mixed-integer program that defines the decision model instance of a state s_k . Thus, Eq. (1) could also be written as:

$$\begin{aligned} & \max && R(s_k, x_k) + V(s_k^x) \\ & \text{subject to} && A(s_k)x_k \leq b(s_k), \\ & && x_k \in \{0, 1\}^{n(s_k)} \end{aligned} \quad (2)$$

The first line is the Bellman Equation. The second part spans the set of potential decisions via the decision variables and the constraints of the decision model instance, for example, routing constraints or capacity constraints (We note that dependent on the modeling style and the problem at hand, parts of the variables may be of real values). With respect to the sequential decision process model, for SDVRPs, both decisions and information model realizations change the system state and therefore the characteristics of $A(s_k)$ and $b(s_k)$. We will illustrate this with the example introduced in Section 2 later in this section.

While every SDVRP (and every sequential decision process) can be modeled as shown in Fig. 4, they differ in their information model and their decision model, i.e., the structure of $R(s_k, x_k)$, s_k , x_k and ω_k and consequently the development of $A(s_k)$ and $b(s_k)$.

For an in-depth overview on modeling sequential decision processes for SDVRPs, we refer to [Ulmer et al. \(2020b\)](#). In the following, we investigate information models and decision models for SDVRPs in more detail.

3.2. Information model

In this section, we describe how uncertainty can be modeled by means of an information model, a task of predictive analytics.¹ Since the uncertainty is “external”, it is also referred to as “exogenous information” ([Powell, 2019](#)) that becomes known over time. The information model Ω describes the stochastic information and its dimensions and ω_k describes the specific realization that becomes known in decision point k .

The information model in SDVRPs can represent uncertainty in the dimensions demand, the environment, or the resources. That is, information models can contain information for three dimensions. [Waller & Fawcett \(2013\)](#) also provide a list of several different logistics applications along with data that can be gathered. In the following, we discuss how the corresponding information can be modeled. We further illustrate that the information models differ in several characteristics such as *dimensionality*, *heterogeneity*, *disruptivity*, and *correlation*, always dependent on the underlying SDVRPs.

Stochastic demand. While stochastic customer requests may occur at any time point and at any location in the service area, certain patterns can be typically observed. For example, customer requests for personal goods or services occur rather in residential areas whereas requests for business-oriented services or goods also occur in industrial areas. Correlation can occur for applications such as food delivery, e.g., dependent on the weather. Customers also place more requests for a meal delivery at actual meal times resulting in heterogeneity in demand over time and space. Stochastic customer request locations are either modeled by a set of discrete locations ([Albareda-Sambola, Fernández, & Laporte, 2014](#)) or via a continuous distribution over the service area ([Angelelli et al., 2009](#)). The ratio between the number of stochastic requests and the overall number of requests can be described by the degree of dynamism (DOD, [Larsen et al., 2002](#)). This metric describes the effect of stochastic requests, that is, the disruptivity that can be associated with the stochasticity. With a small DOD, the main characteristics of [Eq. \(2\)](#) stay the same and decisions may only require slight adaptations. With a high DOD and many new customers in the system, $A(s_k)$ and $b(s_k)$ change substantially and this disruption likely renders the previous routing plan insufficient. Suitable distribution types for stochastic demand quantities depend on whether continuous items like liquids ([Hvattum, Løkketangen, & Laporte, 2007](#)) or discrete items like parcels ([Secomandi, 2001](#)) are considered. Besides the demand quantity, demand may contain additional dimensions such as type of service or required equipment, again modeled via discrete distributions ([Chen, Thomas, & Hewitt, 2017](#)). Service times usually depend on customer features, like medical conditions in health care applications or the requested service. For example, an installation may be faster than a repair. Also, the service times may be less variable for installations. Such stochastic demand quantities may change the $b(s_k)$ by enforcing additional trips to the depot and by using resources. Also, by specifying customer demands, they modify $A(s_k)$. For service times, an asymmetrical continuous distribution can capture the rare cases where the service time is extremely high ([Maxwell, Restrepo, Henderson, &](#)

[Topaloglu, 2010](#)). Dependent on the demand realizations, resources are consumed, changing the available resources in $b(s_k)$. Again, if the distributions are very volatile, this might lead to a disruption and the requirement to entirely replan the routing.

Stochastic environment. Travel times denote how long it takes to travel between locations. If they change, they modify $A(s_k)$. Travel times show spatial and temporal heterogeneity, for example, they differ in rush hours and in downtown or rural areas. Further, there may be disruptive binary events like heavy congestion, road closures, or traffic management intervention ([Köster et al., 2018](#)). Another challenge in stochastic travel times is the potential correlation between roads. If one road is congested, it is likely that adjacent roads show congestion as well ([Schilde et al., 2014](#)). Travel times are usually modeled via asymmetrical long-tail distributions. However, modeling correlation is challenging, for example, because roads may only be congested in one direction.

Stochastic resources. Distributions about vehicle availability or about drivers' absences need to depict very rare, disruptive events. Changes in the resources lead to changes in $b(s_k)$. In the case of crowdsourced drivers ([Arslan et al., 2019](#)), spatio-temporal distributions describe when and where drivers are available. Such a driver availability might show correlations with realized and expected demand and may also depend on the last assignment. A distribution depicting the range of vehicles may have to consider factors like weather, altitude differences, or the traffic situation.

Summary. Overall, the information models differ in a variety of factors. A suitable information model is necessary not only for the purposes of descriptive and predictive analytics, but also for prescriptive analytics as it defines the information available for the decision making. Some SDVRPs require information models of high dimensionality to capture the specifics of the revealed uncertainty. Other SDVRPs require information models that can capture the disruptivity of a few events. In some cases, the information model might require more detail because of spatial and temporal heterogeneity. Finally, for some SDVRPs, correlation is important and therefore must be captured by the information model.

3.3. Decision model

In a decision point of the sequential decision process, an instance of the decision model is solved. Here, information about all three dimensions is available and describes the decision model instance via s_k , $A(s_k)$, $b(s_k)$, x_k , $R(s_k, x_k)$, and $V(s_k^x)$. This contains, for example, information about demand in the form of customer locations or time windows, information about the environment, such as travel times, and information about the available resources, for example the vehicle locations. The information that describes the decision model instance is equivalent to the information captured in the state s_k . The decision model contains an objective, decision variables, and constraints. In this section, we describe the three components for SDVRPs and the underlying characteristics, the horizontal balance of the objective function, the complexity and restrictiveness of the constraints, and the dimensionality and range of the decision variables.

Objective function. The objective function generally focuses on either satisfying as much demand as possible with the given resources or minimize the used resources while satisfying all demand. In contrast to static VRPs, the objective function contains two parts, one representing the immediate reward $R(s_k, x_k)$, the other the expected future reward $V(s_k^x)$. The type of “reward” depends on the problem and comprises, for example, maximizing the number of served customers ([Meisel, Suppa, & Mattfeld, 2011](#)), maximizing the achieved revenue ([Albareda-Sambola et al., 2014](#)), minimizing travel plus service costs/times ([Arslan et al., 2019](#)), or minimizing customer inconveniences such as delays or travel time violations ([Bertsimas & Van Ryzin, 1991](#)). Dependent on the SD-

¹ We note that while we assume the information model and its interrelation with the decision making to be known, this is not always the case. [Bertsimas & Kallus \(2020\)](#) discuss decision making for problems with a single decision stage where the information model is not fully known, but auxiliary information is used for decision making instead.

VRPs, the balance between immediate reward and expected future reward varies. For example, in problems where demand occurs dynamically and over the entire time horizon, the immediate reward is usually significantly smaller than the expected future reward over the remainder of the time horizon (Angelelli et al., 2009). Other problems have a more equal balance of immediate and future rewards, for example, when time window violations are minimized (Schilde et al., 2014). Finally, for some problems, a major part of the objective value manifests already in the immediate reward when the first tentative route is planned, e.g. problems where travel times may change over time and the expected overall travel times are minimized (Köster et al., 2018). The initial reward represents the main amount of travel times needed. Further decisions may change the routes but the reward of these decisions is rather a delta value.

Decision variables. Decision variables x_k generally determine how the resources are used to satisfy demand in the environment. Decisions usually include an assignment of customers to vehicles and a tentative route for every vehicle. Furthermore, decisions may comprise offering customer a service (Klapp, Erera, & Toriello, 2018), pricing different services (Ulmer, 2020), determining inventory levels (Brinkmann, Ulmer, & Mattfeld, 2019), and charging (Kullman, Goodson, & Mendoza, 2021) or repositioning idling vehicles (Riley, Van Hentenryck, & Yuan, 2020). Decisions differ in their dimensionality and in their range over the time horizon. For some problems, decisions are “just” about routing the vehicle from customer to customer (Larsen et al., 2002). In other cases, decisions may comprise several additional dimensions, e.g., pickup and delivery, depot returns, charging, pricing, or inventory. Furthermore, for some SDVRPs such as dial-a-ride or taxi services, decisions only range over a short horizon of time. For others, the (tentative) plan of a decision may span the entire time horizon, e.g., in dynamic delivery routing with time windows.

Constraints. Constraints in the form of $A(s_k)x_k \leq b(s_k)$ can be imposed on resources, demand, and the environment. Constraints on resources are, e.g., vehicle capacity, travel speed, skills, or area accessibility. Constraints on demand comprises for example the service itself, time windows, pickup before delivery, or demand volumes. Constraints on the environment capture the travel times, street capacities, or the availability of urban infrastructure such as parking or micro-hubs.

For different SDVRPs, the constraints are more or less complex and restrictive. For some problems, the decision space can be searched freely for a high quality decision. For other problems, the issue is rather the satisfiability of the constraints and finding a feasible decision of Eq. (2) at all. For example, a parcel pickup routing problem as presented in our example may only have constraints about the travel times and the final return to the depot (Soeffker, Ulmer, & Mattfeld, 2019). For other problems, constraints are more complex, e.g., for a dial-a-ride problem where vehicle capacity, pickup before delivery, and customer time windows have to be considered (Schilde et al., 2014). Finding a feasible solution is therefore significantly more challenging.

Summary. Overall, the decision models of different SDVRPs differ in objective function, decision variables, and constraints. For some problems, the major part of the objective value realizes at time the decision is made, for others, it is balanced between now and the future, and in some cases, the immediate reward is relatively small compared to the value of the future. Further, decision variables differ in their dimensionality and their range. In some cases, the decision only comprises the routing component. In other cases, additional dimensions have to be considered. Finally, the constraints differ in their complexity and their restrictiveness ranging from pure routing constraints to a large number of different, very restrictive constraints with respect to resources, demand, and environment.

3.4. Example: modeling the dynamic VRP with stochastic customer requests

In this section, we describe the sequential decision process, the information model, and the decision model for the example introduced in Section 2.3. In this problem, the information model Ω describes the probability distribution of customer requests over time and space. An information model realization ω_k in decision point k is a specific realization of this distribution and contains a new customer request with request time and location. As discussed in Section 2, the spatial distribution as well as the temporal distribution are likely to depend on the service offered. Thus, for the same problem type, heterogeneity and disruptivity of the information model may vary significantly.

The decision model is an orienteering problem (Vansteenwegen, Souffriau, & Van Oudheusden, 2011), an optimization problem in the form of Eq. (2) with an objective function in the form of the Bellman Equation maximizing the number of customers served. The decision model instance in a state s_k describes the specific optimization problem in a decision point, that is, the current vehicle location, the locations of customers that need to be visited, and the new request. The decision variables x_k describe a tentative routing plan. The routing plan may span the entire horizon to ensure feasibility, but it can be modified in the next decision point. The constraints $A(s_k)x_k \leq b(s_k)$ are standard VRP constraints limiting the time horizon, ensuring that all accepted requests are served, as well as the subtour elimination constraints. After choosing a decision, the transition to the next decision point occurs based on the state, the decision taken, and the next information model realization, that is, a new customer request. Since an acceptance requires resources, such a decision impacts the constraints of the new decision model instance.

As discussed, the distribution of the rewards for this problem also depends on the underlying application. If only a small number of additional requests can be served over the overall horizon as for the maintenance and repair case, then the immediate reward is relatively large in comparison to the expected future reward. If many requests occur as for the express mail case, the relationship between the immediate reward and the expected future rewards is less balanced.

4. Computational methods

Once problem, information model, and decision model are defined, we need to derive policies. As defined in Section 3, a policy assigns a decision to every instance of the decision model (or state). Because typically SDVRPs are too complex to be solved exactly, heuristic computational methods are applied. In this section, we discuss the general types of such methods.

All methods to solve SDVRPs prescribe decisions and therefore are by definition methods of prescriptive analytics (Lustig et al., 2010). However, they can be classified in their different exploitation of the information model. Methods may consider the information model in a descriptive way, just operating on the observed information model realization in a state. Methods may exploit the information model in a predictive way, analyzing and using the information model characteristics to derive decisions, e.g. via scenarios (Shmueli & Koppius, 2011). Finally, they can exploit the information model in a prescriptive way, not only analyzing the information model in isolation, but also its general interaction with the decision model (Bertsimas & Kallus, 2020). The latter two can further be differentiated in where the information model is considered; externally to derive a method (e.g., based on a-priori analytical considerations) or internally when applying the method (e.g. via sampling of realizations in a state). For our example problem, an internal consideration may sample future customer

Table 1
Method classification with respect to their exploitation of the information model.

		Information model use		
		None	External	Internal
Exploitation	Descriptive	Rolling horizon	–	–
	Predictive	–	I Policy function approximation Cost function approximation	IIa Lookahead Rollout
	Prescriptive	–	–	IIb Policy iteration Value function approximation

requests to evaluate current decisions while an external consideration may have derived a classification of customers a-priori and then uses this classification in a decision rule to decide about offering service or not. Based on the three ways of exploiting the information model plus the differentiation of external or internal usage, we can classify the method classes for SDVRPs in Table 1.

The first method class is rolling horizon reoptimization. Rolling horizon methods focus only on the current state information, in particular, the currently observed information model realization. Such methods solve the current decision model instance (based on the current information model realization) with the current reward as objective function (e.g., Branchini, Armentano, & Løkketangen (2009); Gendreau, Guertin, Potvin, & Séguin (2006); Gendreau et al. (1999); Ichoua et al. (2000)). For our example, a rolling horizon method would always accept a request if feasible. Rolling horizon methods assume that the observed information model realizations will remain unchanged in the future. They exploit the information model only in a descriptive way and therefore provide myopic, often inflexible decisions. Thus, we refrain from discussing this method class in detail.

Besides rolling horizon, we consider two categories with the respective method classes either considering the information (and decision) model externally before defining the policy (category I), or internally in the policy (category II). The second category is further differentiated in their exploitation of the information model, either predictive (IIa) or prescriptive (IIb).

Category I describes approaches where the problem and its functionality are analyzed, often in a simplified way. Based on the analysis, a policy is derived, e.g., hard coded decision rules such as waiting or threshold strategies. This policy is then applied in every state without additional considerations of future information model realizations or decision model instances. However, because the structure of both are usually the foundation of finding the policies, we see them exploiting the information model (at least) in a predictive way.

In category II, the information model is used internally in the method, usually via sampling realizations. Category IIa summarizes approaches that derive solutions considering information model realizations by “looking ahead” into the future. However, they integrate future decision making only to a very limited extent and do not analyze the general dependencies between information and decision model (i.e., predictive exploitation). Approaches from category IIb go one step further and consider the general interaction between information model realizations and decision making (i.e., prescriptive exploitation).

We illustrate the differences by extending Fig. 4 to Fig. 5. The additional boxes depict the basis for decision making in decision point t_1 . The smallest light gray box with a dotted line depicts approaches from category I, not integrating any future developments internally. The box in light gray dashed lines depicts approaches in category IIa, mainly focusing on future information model realizations. The largest box with solid light gray lines describes the information basis for approaches in category IIb that consider both information model realizations and future decision model instances.

In the following, we describe categories I, IIa, and IIb and how they consider information model and decision model along with examples. We discuss that the methods differ in the level of detail of information and decision model they capture as well as the length of the time horizon their considerations span. We also sketch corresponding methods for our SDVRP-example.

4.1. Category I: predictive and external exploitation of the information model

Category I contains two method classes, policy function approximation (PFA) and cost function approximation (CFA). Both are based on preliminary analyses of the problem, often with simplified assumptions about information and decision models (e.g., Thomas, 2007). The derived insights are then used to define the decision policy.

The first idea is to follow some analytically derived concept or common-sense of what “good” solutions look like. Such policies are called policy function approximations. While such strategies are very problem-specific, they all have in common that both current reward and future value are disregarded and a rule defines how to derive a decision for the current decision model instance. Examples are, amongst others, waiting strategies (Branke, Middendorf, Noeth, & Dessouky, 2005; Thomas, 2007), pre-defined policies that decide about the routing (Pavone, Bisnik, Frazzoli, & Isler, 2009), or methods that incentivize vehicles serving particular areas in the city (Ichoua, Gendreau, & Potvin, 2006; Van Hemert & La Poutré, 2004). The resulting decisions may yield some flexibility and anticipation of the future. However, the information model is not considered internally in every state. Further, limited effort is spent on searching the decision model instance in detail.

The second idea is to alleviate the myopic and inflexible decision making of the rolling horizon reoptimization. This is done by manipulating either reward function $R(s_k, x_k)$ or constraints via $A(s_k)$ and $b(s_k)$ to incentivize flexible decisions (Riley et al., 2020; Ulmer, Nowak, Mattfeld, & Kaminski, 2020a) and prohibit inflexible decisions (Al-Kanj, Nascimento, & Powell, 2020). A common example are safety buffers (Ulmer, Thomas, Campbell, & Woyak, 2021). Such methods are called cost function approximations. In CFAs, no additional information about the information model is used in a state. Instead, CFAs invest the computational resources into solving the current, but manipulated, decision model instance as effectively as possible. To this end, neighborhood search (Angelelli et al., 2009; Ulmer et al., 2020a; Ulmer et al., 2021) or metaheuristics (Ichoua et al., 2006) are applied.

Example: approaches in I for the dynamic VRP with stochastic customer requests For our example, a cost-function approximation may change the available time in $b(s_k)$ to integrate requests, for example, based on their location in the service area. A policy function approximation may define a detour-threshold and only accept requests if the increase in tour duration lies below this threshold.

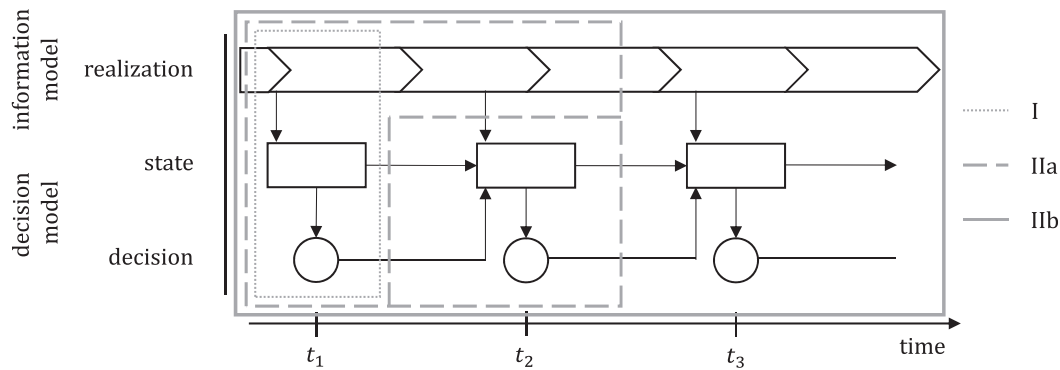


Fig. 5. Approaches with their use of information model and decision model.

4.2. Category IIa: predictive and internal exploitation of the information model

There is a group of methods that consider future realizations of the information model internally but without fully considering future decision making. Thus, they exploit the information model in a predictive way. This idea is depicted in Fig. 5 as the light gray boxes with the notation “IIa”. These boxes cover parts of the information model that lay in the future. Optionally, they may also cover parts of the decision model instances in the future in a very limited way. The methods have in common that they integrate sampled information model realizations in their decision making. In the following, we describe the two main strategies, scenario-based approaches and post-decision rollout algorithms in detail.

In scenario-based approaches, information model realizations are sampled to create a set of scenarios (like in Azi, Gendreau, & Potvin (2012); Bent & Van Hentenryck (2004); Ghiani, Manni, Quaranta, & Triki (2009); Hvattum, Løkketangen, & Laporte (2006); Schilde et al. (2014)). Such scenarios are static and deterministic, thus potential changes in information and decision model over time are ignored. Besides a few exceptions where a multi-stage decision problem is solved (Ferrucci, Bock, & Gendreau, 2013), the scenarios are solved individually, either via mixed-integer solvers, neighborhood searches, or metaheuristics. Based on the individual scenario solution, a decision is derived.

The most prominent scenario-based approach is the multiple-scenario approach (MSA), described in Algorithm 1. Given a state

ates m scenarios via function $Scenario(s_k, \omega^i)$ (e.g., by augmenting the set of current customers with the set of sampled customers). Then, for each scenario \hat{s} , the augmented decision model realization is solved with respect to the immediate reward with function $Decision(\hat{s})$, e.g. via mixed-integer solvers or metaheuristics. Then, the solutions are reduced to omit any augmented information (e.g., by skipping sampled customers in the routes). This results in a set X of m potential decisions for state s_k . To select one, the MSA draws on a consensus function $Consensus(x_i, X)$. This function measures the similarity between a decision x_i and the overall set of solutions X (e.g., the number of edges of decision x_i that are used in other decisions). The most similar decision x^* is then applied to state s_k .

Scenario-based approaches consider changes in the information model explicitly, but do not consider their dynamism over the time horizon. They also spend significant effort in solving the individual scenarios in detail. There is an alternative method that focuses more on the dynamism over the time horizon and less on a detailed optimization within the scenarios. We call the methods post-decision rollout algorithm (RA), compare Goodson, Thomas, & Ohlmann (2017), but similar concepts are known as discrete event simulation (Fishman, 2001) or Monte-Carlo-tree search (Browne et al., 2012). The idea of an RA is, in every state, to preselect a set of potential decisions, and approximate the value of each resulting post-decision state via simulation of information model realizations and decision model instances (Goodson, Ohlmann, & Thomas, 2013; Ulmer et al., 2019). The approximated value is then used in the Bellman Equation to select one of the decisions. Within the simulation, generally a base policy, a simple, runtime efficient decision rule, is applied to solve the simulated decision model instances (Novoa & Storer, 2009; Secomandi, 2001). Further, the simulation horizon is often limited to a number of steps because the simulations diverge increasingly from reality over the simulation horizon (Brinkmann et al., 2019).

The procedure of an RA is depicted in Algorithm 2. Inputs are the state s_k , the information model realization sequences $\{\omega^1, \dots, \omega^m\}$, the horizon limit L , and the base policy π_b . The RA generates a set of candidate decisions for state s_k with function $Decisions(s_k)$ (e.g., by applying a set of runtime-efficient policy-function approximations). For each of the candidate decisions, the RA generates the corresponding post-decision state, with function $PDS(s_k, x)$ and then approximates the value \hat{V} for each post-decision state. For approximation, m simulation runs are conducted that consist of sequences of information model realizations and decision model instances. Each simulation run i terminates either when L decision points were simulated or when the process reached termination state s_K . Else, the transition function T with current post-decision state s^x and sampled realization ω^i generates a new state. In every state s of the simulation, the basis policy

Algorithm 1: Multiple scenario approach.

```

Input: State  $s_k$ , Information Model Realization Sequences  $\{\omega^1, \dots, \omega^m\}$ 
Output: Decision  $x^*$ 
2 // Generate Scenario Solutions
3  $X \leftarrow \emptyset$  // Set of Scenario Solutions
4 for  $i = 1, \dots, m$  do
5    $\hat{s} \leftarrow Scenario(s_k, \omega^i)$  // Generate Scenario  $i$ 
6    $\hat{x}_i \leftarrow Decision(\hat{s})$  // Solve Decision Model for Scenario  $i$ 
7    $X \leftarrow X \cup Reduce(\hat{x}_i, s_k)$  // Add Decision Reduced to  $s_k$ 
8 end
10 // Selection
11  $M^* \leftarrow 0$  // Consensus Measure
12 for all  $x_i \in X$  do
13   if  $(Consensus(x_i, X) \geq M^*)$  // Highest Consensus Value
14   then
15      $x^* \leftarrow x_i$  // Best Found Solution
16      $M^* \leftarrow Consensus(x_i, X)$  // Highest Found Consensus Value
17   end
18 end
19 return  $x^*$ 

```

s_k and a set of m sampled information model realization sequences $\{\omega^1, \dots, \omega^m\}$ (e.g., future customer requests), the algorithm gener-

Algorithm 2: Post-decision rollout algorithm.

Input: State s_k , Information Model Realization Sequences $\{\omega^1, \dots, \omega^m\}$, Limit L , Base Policy π_b
Output: Decision x^*

```

1  $X \leftarrow \text{Decisions}(s_k)$  // Generate Set of Decisions
2 for all  $x \in X$  do
3    $s_k^x \leftarrow \text{PDS}(s_k, x)$  // Generate Post-Decision State
4    $\hat{V}(s^x) \leftarrow 0$  // Initialize Value Function
5   // Simulation
6   for  $(i = 1, \dots, m)$  //  $m$  Simulation Runs per Post-Decision State
7   do
8      $s^x \leftarrow s_k^x, l \leftarrow 0$ 
9     while  $\text{AND}((s^x \neq s_k), (l \leq L))$  // Termination of Run
10    do
11       $s \leftarrow T(s^x, \omega^i)$  // Transition Function
12       $x \leftarrow X^{\pi_b}(s)$  // Solve Decision Model via Base Policy
13       $\hat{V}(s_k^x) \leftarrow \hat{V}(s_k^x) + \frac{1}{m}R(s, x)$  // Update Value Function
14       $s^x \leftarrow \text{PDS}(s, x)$ 
15       $l \leftarrow l + 1$ 
16    end
17  end
18 end
19 end
20 // Selection
21  $R^* \leftarrow -\text{bigM}$ 
22 for all  $x \in X$  do
23   if  $(R(s_k, x) + \hat{V}(\text{PDS}(s_k, x)) \geq R^*)$  // Bellman Equation
24   then
25      $x^* \leftarrow x$  // Best Found Solution
26      $R^* \leftarrow R(s_k, x) + \hat{V}(\text{PDS}(s_k, x))$  // Best Found Value
27   end
28 end
29 end
30 return  $x^*$ 

```

π_b is used to derive decisions $X^{\pi_b}(s)$. The corresponding reward is added to the approximated value \hat{V} proportionately. After the values for all post-decision state candidates are approximated, the Bellman Equation is used to select a decision x^* .

In essence, MSAs and RAs incorporate future information model realizations in much detail but only for a very limited horizon. Further, while searching the current decision model instance receives significant attention, future decision making is either ignored entirely or simplified substantially. Due to the solution effort of the scenarios, MSAs and RAs require a high computational effort in a decision point.

Example: approaches in IIa for the dynamic VRP with stochastic customer requests In an MSA for the example from Section 2.3, possible future requests are sampled and both the actual new request and the sampled artificial ones are added to the current orienteering problem. This problem (including potential future requests) is solved and the sampled requests are removed from the solutions. From the possible routing plans, the MSA chooses the one that is most similar to the others. The new request is accepted in case it is part of this solution. While the MSA can search large solution spaces, it is limited with respect to the consideration of future dynamics. For an RA, first, a small set of possible decisions is selected. For example, only the acceptance part of the decision is considered and the planned route is determined by an insertion procedure. Then, for the two decisions, future demand is simulated and as long it is feasible integrated via an insertion procedure. When comparing the two procedures, we observe that the MSA spends more effort on searching the routing solution space while the RA focuses more on dynamism. This observation is common for the two types of scenario-based methods.

4.3. Category IIb: prescriptive and internal exploitation of the information model

There are methods that consider future changes in both information model realizations and decision model (and their interac-

tions) in their decision making. Thus, they exploit the information model in a prescriptive way. The methods in this category are depicted in Fig. 5 by the box with solid lines surrounding the entire chart. The corresponding methods fall in the broader category of reinforcement learning, e.g., value function approximation (for example in Schmid, 2012), policy iteration (as in Secomandi, 2000), or Q-learning (Chen, Ulmer, & Thomas, 2021). The approaches approximate a policy or the value function by iteratively simulating through the time horizon. Because the repeated simulation requires substantial computation time, it is usually done in an offline manner, that is, in advance of the actual decision making. Once the offline learning is complete, the learned policy can be applied without additional, runtime-expensive simulations as long as the information model remains valid.

Within the offline simulations, information model realizations are sampled and every decision model instance is solved based on the learned experience of previous simulations. After each simulation run, the learned experience is updated in terms of the observed states, decisions, and rewards. This procedure alleviates the curse of dimensionality in the information space (Powell, 2011). For the decision space, a pre-selection of decision candidates is required, as for the RA. The state space is aggregated to a set of features and the values for the features are stored, either in groups or in functional form, in an approximation architecture.

In the following, we focus on value function approximation as a representative of offline learning approaches. We first present the algorithmic procedure and then discuss the forms of aggregation applied to SDVRPs.

4.3.1. General procedure

The procedure of value function approximation (VFA) is depicted in Algorithm 3. Inputs are a (large) set of information

Algorithm 3: Value function approximation.

Input: Information Model Realization Sequences $\omega^1, \dots, \omega^N$, Initial Values \hat{V}
Output: Values \hat{V}

```

1 // Initialization
2  $\mathcal{O} \leftarrow \emptyset$  // Observations
3 for  $(j = 1, \dots, N)$  // Simulation of  $N$  Realizations (Learning)
4 do
5    $\mathcal{O}^j \leftarrow \emptyset$  // Observations in run  $j$ 
6    $s^x \leftarrow s_k^0$  // State Initialization for the next Realization
7   while  $(s^x \neq s_k)$  // Simulation of one Realization
8   do
9      $s \leftarrow T(s^x, \omega_j)$  // Next State
10     $x_k \leftarrow \text{Decision}(s, \hat{V}, \mathcal{O})$  // Solve Decision Model
11     $\mathcal{O}^j \leftarrow \mathcal{O}^j \cup \{\text{Aggregate}(\text{PDS}(s, x_k))\}$  // Save Aggregated Observation
12     $\mathcal{R} \leftarrow \mathcal{R} \cup \{R(s, x_k)\}$  // Save Reward
13     $s^x \leftarrow \text{PDS}(s, x_k)$ 
14  end
15  // Update Values
16   $\mathcal{O} \leftarrow \text{UpdateObservations}(\mathcal{O}^j, \mathcal{R})$ 
17   $\hat{V} \leftarrow \text{UpdateValues}(\hat{V}, \mathcal{O})$ 
18 end
19 return  $\hat{V}$ 

```

model realization sequences $\omega^1, \dots, \omega^N$ and an initial value function \hat{V} . Initially, the set of stored observations \mathcal{O} is empty. The VFA iteratively simulates the N realization sequences to learn the value function. Every simulation run starts in the initial state s_k^0 . Then, a sequence of information model realizations and decision model instances is simulated. The decisions are derived based on the approximate value function and the already encountered observations via function $\text{Decision}(s, \hat{V}, \mathcal{O})$. This function can use the Bellman Equation or exploration methods that aim on encountering new observations. We note that only a small set of candidate decisions is searched rather than the entire decision space. After a

decision is selected, the post-decision state $PDS(s, x_k)$ and the reward accumulated up to this point \mathcal{R} is stored. This storage happens in an aggregated way (see Section 4.3.2). After each simulation run, the set of observations \mathcal{O} is updated with respect to the observed post-decision states \mathcal{O}^j and rewards \mathcal{R} via function *UpdateObservations*. Then, the value function \hat{V} is updated with respect to \mathcal{O} in function *UpdateValues*. The specific update procedure depends on the VFA-tuning and the underlying approximation architecture.

4.3.2. Aggregation

For SDVRPs, the post-decision state space is vast and each post-decision state consists of a large number of dimensions, for example, time of day, vehicle locations and planned routes, or customer locations, potentially with additional information. As discussed earlier, post-decision states are aggregated to a set of features. Features are chosen that are assumed to have a meaningful relationship with the values. Examples are the current time in the decision horizon, the available capacity of vehicles, or information about how many customers still have to be visited. The resulting vector of features is then mapped to a value. The mapping depends on the approximation architecture used (Bertsekas & Tsitsiklis, 1996), for example, a lookup table, linear functions, or neural nets. In the following, we analyze feature selection and approximation architecture for SDVRPs.

Feature selection. In Section 2.1, we discussed the three dimensions where uncertainty can occur, demand, resources, and environment. Consequently, features should reflect the state of the three dimensions, especially in case of uncertainty. While some features specifically refer to one of the three dimensions, other features may combine information from multiple dimensions.

Until now, there is no work applying a solution approach from IIb to a problem with stochasticity in the resources or in the environment. Features relating to the environment are therefore not considered in the literature yet. However, since the demand has an effect on the resources, the resources are often captured by the features as well.

In literature on *stochastic service times* (Maxwell et al., 2010; Schmid, 2012), resource-features consider the vehicle locations, the available vehicles, as well as information on the available vehicles in the future. Information about the demand is depicted by features describing customer orders that need to be served, customer demand that cannot be reached, as well as expected future demand.

In literature considering *stochastic demand quantities* (Pandelis, Kyriakidis, & Dimitrakos, 2012; Secomandi, 2000; Secomandi & Margot, 2009), information about the available resources is depicted by the current vehicle location and the remaining capacity. Information about the demand is represented via the service-status of the customers and about the expected future demand.

When considering *stochastic customer requests*, information about the resources is often depicted in terms of the remaining time capacity (slack) (Meisel et al., 2011; Ulmer, Mattfeld, & Köster, 2018a), the currently available vehicles (Agussurja et al., 2019), the vehicle locations (Thomas & White, 2004), and the vehicles that are available in the future (Maxwell et al., 2010). The demand is depicted by means of the service-status of the potential customers in the system (Meisel et al., 2011; Thomas & White, 2004) or features describing the customer orders waiting to be served (Schmid, 2012). More detailed information and a table summarizing the features applied in the literature can be found in A.2.

In general, while the features represent important information about vehicles and customers, they lack detail, for example, spatial customer information (Ulmer et al., 2018a) or the distribution of vehicles within the city (Al-Kanj et al., 2020).

Approximation architecture. After choosing features, the mapping between the aggregated states and the values needs to be determined. This mapping therefore builds the bridge between the features and the values as depicted in Fig. 6. For this mapping, different approximation architectures can be chosen. For the sake of simplicity, the one depicted in Fig. 6 is a lookup table.

Possible ways to depict such a relationship between features and values are either parametric or non-parametric dependencies. Non-parametric architectures can be, for example, lookup tables (Thomas & White, 2004; Ulmer, Soeffker, & Mattfeld, 2018b) and their variants, decision trees, or state representatives (Agussurja et al., 2019; Soeffker et al., 2019). Non-parametric architectures have the advantage that they can capture complex value function structures, often observed in stochastic dynamic vehicle routing. However, they do not scale well and require many simulation runs to learn (Ulmer & Thomas, 2020). In the case of parametric dependencies between states and values, the type of relationship has to be defined, possible options here are linear (Meisel et al., 2011; Secomandi, 2000), weighted combinations of basis functions (Maxwell et al., 2010; Schmid, 2012), or, theoretically, nonlinear dependencies. Parametric architectures allow a faster approximation, but often fail to fully capture the structure of the value function (Ulmer & Thomas, 2020). There is also an increasing amount of work using neural net architectures (Chen et al., 2021; Joe & Lau, 2020). Neural nets can be seen as a hybrid class between non-parametric and parametric, because they are tuned via parameters while the value function does not necessarily follow any functional form. However, training them often requires a large number of simulation runs and the interpretation of the outcome is generally difficult. More information about the approximation architectures used in literature can be found in A.3.

Example: approaches in IIb for the dynamic VRP with stochastic customer requests Ulmer et al. (2018a) presents a VFA for our example. Decisions are reduced to acceptance of requests, routing is conducted by a runtime-efficient insertion procedure. Post-decision states are aggregated to two features, point of time and slack. Slack indicates the free time available to serve additional requests, once all existing customers are served. Both features indicate changes in the value function. With increasing point of time, the expected future demand decreases, as does the value function. With decreasing slack, the possibility to integrate future demand decreases, as does the value function. As approximation architecture, a lookup table is chosen because the functional dependencies between time, slack and value function are complex, especially in case of spatial heterogeneity in the information model.

5. Recommendation

In the previous sections, we illustrated that every SDVRP follows the same structure, a sequential process with iterations of information model realizations and decision model instances. However, they differ substantially in the characteristics of information and decision model. We further illustrated the available prescriptive analytics (PA) methodology, their functionality, and their strengths and weaknesses. In this section, we combine the insights of models and methodology to derive recommendations about how to approach a specific SDVRP with prescriptive analytics. We first give an anecdotal evidence and then a general recommendation.

5.1. The impact of information model heterogeneity

In the following, we show how the heterogeneity within the information model impacts the performance of different PA-methods for the stochastic dynamic vehicle routing problem with stochastic




Fig. 6. Depicting the relationship between features and values by means of an approximation architecture.




Fig. 7. Relative performance of an MSA compared to a VFA for varying heterogeneity in the information model, Ulmer et al. (2019).

requests described in the example in Section 2.3 (details on problem, methods, and results can be found in Ulmer et al., 2019). We recall that decisions are made about accepting and routing customer requests with the objective to maximize the expected number of served customers per day.

The tested PA-methods are a value function approximation (from category IIb) and a variant of the multiple-scenario approach (from category IIa) as described in the examples in Sections 4.2 and 4.3. The VFA aggregates post-decision states to two features, the point of time and slack. The MSA samples a set of near-future requests and evaluates decisions with respect to the number of sampled requests that can be served. Thus, both methods treat the information model very differently. The VFA incorporates the information model highly aggregated without explicit consideration of spatial information such as vehicle or customer locations. However, due to the frequent simulations of entire days, the VFA captures longer term effects of decisions and many different (potentially disruptive) realizations of the information model. The MSA considers the available and sampled information in full detail but only a few realizations and over a very limited time horizon.

Ulmer et al. (2019) compare the two methods for varying heterogeneity in the information model, namely, the spatial distribution of the customer requests. They generate instances stepwise shifting from a fully homogeneous spatial distribution where customers are uniformly distributed over the city to a very heterogeneous distribution where customers accumulate in three compact clusters. Both methods are applied to the different instances and the relative performance of the MSA compared to the VFA is calculated. The results are depicted in Fig. 7. The x-axis depicts the heterogeneity of the instances from fully homogeneous (0%) to very heterogeneous (100%). The y-axis shows the relative performance of the MSA. Values below 1 indicate that the VFA performs better, values above 1 indicate that the MSA provides better results.

We observe a nearly linear dependency between heterogeneity and relative performance. The VFA is superior for homogeneous instances and the MSA is superior for instances with significant heterogeneity when the more detailed consideration of the information model becomes important. This experiment highlights that

there is not one dominant PA-method for the problem but that the suitability of the method depends on the heterogeneity in the information model and how well the method can depict such heterogeneity in detail. While this observation is rather anecdotal evidence, we will use it as motivation for the following general considerations.

5.2. Generalization

In this paper, we illustrated that information models and decision models differ in several characteristics for different SDVRPs. We also illustrated by means of examples that PA-methods consider decision and information model in their decision making in different level of detail and with a different time horizon. We now connect models to methodology by analyzing when detail and horizon in information and decision model are important and how they are treated by the different methods.

Information models can be characterized with respect to dimensionality, heterogeneity, correlation, and disruptivity. *Dimensionality* represents the number of different information dimensions a realization reveals. *Heterogeneity* indicates the differences in the realized information, e.g., in time and space. *Correlation* indicates that realized information might be correlated, e.g., again over time and space. And *disruptivity* means that two subsequent information model realizations differ substantially. As motivated by the example of the previous section, the information model characteristics require different treatment. This treatment should consider the level of detail spent on the information model realization and the time horizon over that future realizations are integrated in decision making. Detail is required in case of high dimensionality and heterogeneity as well as correlation. A longer time horizon is required in case of high disruptivity to anticipate potentially severe changes in the information model realizations.

Decision models differ in their objective function, their decision variables, and their constraints. The objective functions differ in how immediate reward and future value are *balanced*. Decision variables differ in their *dimensionality* and their *time range*. The constraints result in different *complexity* and *restrictiveness*, and therefore, the potential of finding feasible solutions fast. As for the

Table 2
Overview methodology.

Category	Method	Information model		Decision model	
		Detail	Horizon	Detail	Horizon
I	Rolling-Horizon Reoptimization	+++	–	+++	–
	Cost Function Approximation	++	+	++	+
	Policy Function Approximation	–	++	–	++
IIa	Multiple Scenario Approach	++	+	++	+
	Rollout Algorithm	+	++	++	++
IIb	Value Function Approximation	+	+++	+	+++

information model, the different characteristics require a different treatment of detail and horizon of the decision model. If the current reward is relatively marginal and a large amount of reward manifests in the future value, based on future decision making, a longer horizon is required. If decisions have a large number of dimensions, a more detailed consideration of the decision model is needed. In case decisions span a larger time range into the future, a longer horizon is required as well. Decision models with very complex and restrictive constraints also require a detailed consideration, e.g., just to find a feasible solution.

We established that the requirement of considering detail and time horizon varies depending on the characteristics of information and decision model. In Section 4, we illustrated that methods differ in which level of detail and time horizon they capture. We present an (qualitative) assessment of the methods’ capabilities in Table 2. The table depicts the three method categories with the methods described in this paper. It further depicts a method’s ability to capture detail and horizon of information model and decision model. We differentiate between no capability (“–”), slight capability (“+”), good capability (“++”), and excellent capability (“+++”). We note that the assessment is not validated by quantitative measures but is based on conceptual considerations and experience described earlier in this paper.

Rolling-horizon reoptimization can fully capture the details of current information model realizations and decision model instances, but ignores changes over the time horizon. The cost function approximation can capture decision model and information model now and, implicitly based on analytical considerations, in the near future. Thus, the overall level of detail is relatively high but the horizon comparably short. The policy function approximation disregards detail in information and decision model and selects a decision fast and straightforwardly. However, because the decision making is based on analytical or practical considerations, such decisions capture effects of information and decision model over a longer horizon. Similar to the cost-function approximation, the multiple scenario approach considers decision and information model now and in the near future in detail, but disregards longer term effects. The rollout algorithm evaluates only a few candidate decisions of the current decision model instance, but with explicit integration of future changes in information and decision model over a (limited) time horizon. Finally, the value function approximation captures general, long term horizon effects due to the manifold simulation runs. However, the information is highly aggregated and decision making is reduced to a set of few candidate decisions. Thus, the level of detail is rather limited.

When comparing the methods in Table 2, we observe that there is no dominant method. Instead, the method selection should be made based on the characteristics of information and decision model. If they require consideration of detail, e.g. because of dimensionality, heterogeneity, correlation, or complex constraints, methods such as cost function approximation or multiple scenario approach are promising choices. If they require consideration of longer term effects, e.g., because of disruptions or because the future reward is substantially higher than the immediate rewards,

methods such as policy function approximation or value function are likely the better choice. If they are mixed, a rollout algorithm might provide a sufficient consideration of detail and horizon while the other methods fall short in one of them. Essentially, when approaching a new SDVRP, we recommend to first analyze the characteristics of information and decision model, then determine the requirements of detail and horizon based on the characteristics, and finally select a method that matches the requirements with the its strengths.

6. Outlook

In this paper, we have analyzed the work on stochastic dynamic vehicle routing problems (SDVRPs) in the light of prescriptive analytics (PA). PA for SDVRPs combines decision modeling with information modeling. A SDVRP can be seen as a mutually dependent sequence of iterations between information model realizations and decision model instances. Information model and decision model implement the characteristics of the specific SDVRP. We have concluded that dependent on the characteristics, different consideration should be given to the level of detail and time horizon of information and decision model. We have further presented a classification of the SDVRP-methodology with respect to prescriptive analytics, i.e., the treatment of information model and decision model. We have finally given a guideline to the methods’ suitability with respect to the characteristics of information and decision model.

There are plenty of opportunities for future work in the field of SDVRPs, e.g., due to the vast number of new e-commerce and urban mobility applications. Working on these problems will provide important methodological insights and managerial impact. From a methodological view, it might also be valuable to consider problems with an incomplete or inaccurate information model, but with auxiliary information that can be used instead (Bertsimas & Kallus, 2020). We see another important avenue in validating and refining the proposed interdependencies between characteristics of information and decision model and the performance of the computational methods. First, the qualitatively derived characteristics of information and decision model require quantitative measures to differentiate problems and instances. Second, the conceptual evaluation of the methods in Table 2 requires quantitative validation. Because this requires significant conceptual work in combination with the development and testing of different methods on SDVRPs with different characteristics, this will undergo cumulative research of the community.

Acknowledgment

The authors like to thank the associate editor and the three anonymous reviewers for the constructive and thorough remarks. They were very valuable in revising the manuscript. Marlin Ulmer’s work is funded by the Emmy Noether Programme of the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation), project 444657906. This support is gratefully acknowledged.

Appendix A.

In the appendix, we provide further information on different application areas in Section A.1, the choice of features in Section A.2, on approximation architectures in Section A.3, and on the modeling of uncertainties in Section A.4. We provide a summarizing list of literature classified according to the proposed features in Section A.5.

A1. Application areas

In the following, we give a brief overview on application areas where stochastic dynamic problems emerge. We classify those areas in transportation of goods, passengers, and services at the customer’s location. While in all three areas customer demand is fulfilled by the provider’s resources, they differ in the fulfillment process and the customers’ involvement. Thus, they often differ in restrictions and goal. In Fig. A.1, these three broad application areas along with some exemplary applications are depicted. In the following, we discuss the three application areas in more detail and highlight some special characteristics inherent to the application areas.

Transportation of goods. The most prominent applications of stochastic dynamic vehicle routing problems dealing with transporting goods are courier services (Ichoua et al., 2000), food delivery, or same-day delivery (Voccia, Campbell, & Thomas, 2019).

In such settings, a fleet of vehicles serves customers in a service area by either picking up or delivering goods. While the goods are moved through time and space, the customer is only involved in the first or final step of the fulfillment process, the pickup or delivery. For such problems, goal and restrictions spread along the pickup and delivery process, e.g., time windows (Gendreau et al., 1999), working hours (Angelelli et al., 2009), or vehicle capacity (Van Hemert & La Poutre, 2004). In the case of delivery problems (Secomandi, 2000), the depot has to be visited before the customer visit to load the customer’s good onto the vehicle. And in the case of the pickup and delivery problem (Mitrović-Minić, Krishnamurti, & Laporte, 2004), the pickup location has to be visited before the delivery location.

Passenger transportation. The second application area is the group of passenger transportation problems, for example, dial-a-ride problems (Psaraftis, 1980), emergency transportation of patients (Schmid, 2012), or ride-sharing (Agussurja et al., 2019). Usually, they are pickup and delivery problems where customers have to be driven from an origin to a destination within a short amount of time. Thus, the customer is involved in nearly the entire service process. The transportation of passengers poses additional restrictions on the problem settings. Customers have to be picked up before dropping them off at their destination and time windows may be to be considered. Also, vehicles can only carry a small number of people at the same time (imposing capacity restrictions, Sáez et al., 2008) and detours, both spatial and temporal, have

to be kept small to avoid customer inconvenience (Schilde et al., 2014). In emergency transport applications, profitability is only of secondary interest as fast response times and safety are the main goals and serving customers is mandatory. Since only individual requests are served at a time and those have to be processed as fast as possible, decisions are rather made about how to allocate vehicles within the service area than about routing plans (Maxwell et al., 2010).

Services. Finally, the provision of services such as technician services (Larsen et al., 2002) or health care services (Demirbilek, Branke, & Strauss, 2021) at a customer’s home is becoming more and more relevant. In contrast to transportation services, a major part of the fulfillment process takes place at the customer’s location and the service itself is often more complex. Here, the qualification of the vehicle driver has to match the service to be provided (Demirbilek et al., 2021). Healthcare services usually define non-urgent problems where patients have to be visited regularly (Demirbilek et al., 2021). Customer time windows and driver-customer consistency might be considered to provide regularity to patients (Song, Ulmer, Thomas, & Wallace, 2020). Especially in health care services, service times might be rather long and may depend on the driver’s familiarity with the patient (Ulmer et al., 2020a). In emergency problem settings, patients sometimes are not only transported, but are also treated medically (Maxwell et al., 2010; Schmid, 2012). Due to the connection with the patient transport, capacity restrictions have to be considered.

A2. Feature selection

In the following, we analyze the literature with respect to how the features relate to the source of uncertainty. We are not aware of such methods in literature considering stochastic travel times or stochastic resources and therefore only consider stochastic service times, stochastic demand quantities, and stochastic requests. All of these are sources of uncertainty concerning the customer demand which utilizes the resources in different ways. In the following and in Fig. A.1, we describe which features are chosen by literature for the different sources of uncertainty.

Stochastic service times. Work on stochastic service times is presented in Maxwell et al. (2010) and Schmid (2012). In these articles, stochastic service times are combined with stochastic customer requests and decisions are made about how to redeploy ambulances after serving a customer. Stochastic service times impact the time when (and where) vehicles will be available again for the service provider. Thus, information about this time and the locations is required. Because customer requests are uncertain, information about demand is required as well, and how well this demand can be served by the fleet. Both articles describe the spatial and temporal utilization of the ambulances by means of direct or indirect information about the locations of the vehicles and the remaining customer requests that still need to be served.

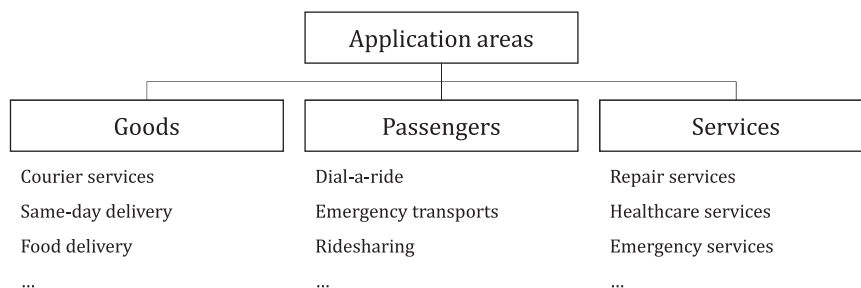


Fig. A1. Application areas within stochastic dynamic vehicle routing.

Table A1

Features within the dimensions resources and demand used when considering stochastic demand quantity (upper part) or stochastic requests (lower part).

	Resources					Demand			
	Slack	Vehicle location	Available vehicles	Vehicles available in future	Remaining capacity	Status of all customers	Orders waiting to be served	Demand not reachable	Expected future demand
Secomandi (2000)		✓			✓	✓			✓
Secomandi & Margot (2009)		✓			✓	✓			
Pandelis et al. (2012)		✓			✓	(✓)			
Thomas & White (2004)		✓				✓			
Maxwell et al. (2010)			(✓)	(✓)			✓	✓	✓
Meisel et al. (2011)	✓								✓
Schmid (2012)		(✓)	✓				✓		
Ulmer et al. (2018a)	✓								(✓)
Ulmer et al. (2018b)	✓								(✓)
Agussurja et al. (2019)			✓				✓		
Chen et al. (2021)	(✓)						(✓)		(✓)
Soeffker et al. (2019)	✓								(✓)
Ulmer et al. (2019)	✓								(✓)
Van Heeswijk, Mes, & Schutten (2019)			✓				✓		
Al-Kanj et al. (2020)		✓			✓				(✓)
Joe & Lau (2020)						(✓)	(✓)		(✓)
Ulmer (2020)	✓					(✓)	(✓)		(✓)
Ulmer & Thomas (2020)	✓				✓				(✓)

Maxwell et al. (2010) combines the two aspects of resources and demand in the following features. These are

- the number of customers that have been assigned to a vehicle and cannot be served on time (1),
- the numbers of requests that are located such that currently no vehicle is in reach (2) and is on its way into the area (3),
- the numbers of requests that cannot be served because vehicles within reach are already occupied (4) and vehicles that will be in reach soon are already occupied (5).

While Maxwell et al. (2010) focus on specific requests and consider vehicles implicitly, Schmid (2012) selects features reflecting the vehicles in the service area and the time horizon. To this end, the service area is divided into subareas and the time horizon into intervals of equal length. Then, the number of idle vehicles and the number of pending requests are counted per subarea of the service area within a certain time interval.

Table A.1 lists all features used in the literature considering stochastic service times, stochastic demand quantities, or stochastic requests. Since the articles considering stochastic service times also considers stochastic requests, they are listed in this lower section of the table together with the other articles considering stochastic requests. Articles considering stochastic demand quantities are listed in the upper half of the table. We combine the literature as many features describing the resources or the demand overlap. As mentioned, since the environment is not considered stochastic, it is not described by means of features.

Stochastic demand quantities. Work considering stochastic demand quantities is considered in Secomandi (2000), Secomandi & Margot (2009), and Pandelis et al. (2012). If demand quantities at customer locations are stochastic, it is unclear if a vehicle in a specific location with a remaining fill level of goods can fully serve the next customer’s demand or has to return to the depot. This means that it is unclear how much of the resources travel time and goods are required to serve all customer demand. Thus, information about the currently remaining capacity is mandatory to depict the currently available resources. This feature is incorporated in all considered articles. Also, information about the current vehicle locations is required as it also describes the resources of the

provider. This feature is also depicted in all articles. In the case of anticipatory decision making that allows returns to the depot at any point in the tour, information about the customer locations that still have to be visited is relevant as well because it indirectly describes how much demand is still unsatisfied. Since the customers are known in these articles, the articles consider the status of all customers (served, partially served, not served yet). While Secomandi (2000) and Secomandi & Margot (2009) consider this information explicitly, it is considered implicitly in Pandelis et al. (2012) by means of the current location and a predefined tour. In addition to the remaining locations, Secomandi (2000) also considers the expected demand of the remaining customers.

Thus, in all articles, information about the resources in the form of the vehicle’s current location and the remaining capacity as well as some information about the demand in the form of the set of (remaining) customers is utilized. The work on stochastic demand quantities along with the used features is listed in Table A.1.

Stochastic requests. A larger number of articles considers stochastic customer requests. In the case of stochastic customer requests, the resource that can be spent is usually the time of the drivers. This resource is used to serve the occurring demand, that is, the customer requests. Therefore, information about the available resource driver time is necessary.

- The tour already planned in combination with the current time provides information about the resources that are already planned and about the resources that are still free. The remaining time is also called “slack”, this feature is used in the literature.
- Another possible feature depicting the use of this resource is the current vehicle location. However, this feature is only rarely used in literature.
- Alternatively, the aggregated information about the vehicles available currently or in the future can be used. While information about the available vehicles is rather common, the future availability is only implicitly considered by Maxwell et al. (2010).

Also, information about the existing and expected future demand for the resource driver time is required.

Table A2
Approximation architectures used.

	Parametric		Non-parametric		Neural nets
	Linear	Weighted comb.	Lookup table	Representatives	
Thomas & White (2004)			✓		
Secomandi & Margot (2009)			✓		
Maxwell et al. (2010)		✓			
Secomandi (2000)	✓				
Meisel et al. (2011)	✓				
Schmid (2012)		✓			
Ulmer et al. (2018a)			✓		
Ulmer et al. (2018b)			✓		
Agussurja et al. (2019)				✓	
Chen et al. (2021)					✓
Soeffker et al. (2019)				✓	
Al-Kanj et al. (2020)			✓		
Joe & Lau (2020)					✓
Ulmer (2020)	✓		✓		
Ulmer & Thomas (2020)	✓		✓		

Table A3
Modeling stochastic travel times.

Work	Discrete	Continuous	Correlated	Time-dependent
Haghani & Jung (2005)		✓		✓
Chen, Hsueh, & Chang (2006)		✓		✓
Potvin, Xu, & Benyahia (2006)		✓		✓
Pureza & Laporte (2008)	✓			✓
Xiang et al. (2008)		✓		✓
Lorini, Potvin, & Zufferey (2011)		✓		✓
Ghannadpour, Noori, & Tavakkoli-Moghaddam (2013)		✓		
Ferrucci & Bock (2014)		✓		
Schilde et al. (2014)		✓	✓	✓
Köster, Ulmer, & Mattfeld (2015)		✓	(✓)	✓
Kim, Ong, Cheong, & Tan (2016)		✓		✓
Köster et al. (2018)	✓		✓	

Table A4
Modeling stochastic service times.

Work	Discrete	Continuous	Symmetric	Asymmetric
Bertsimas & Van Ryzin (1991)		✓		
Bertsimas & Van Ryzin (1993a)		✓		
Bertsimas & Van Ryzin (1993b)		✓		
Papastavrou (1996)		✓		✓
Tassioulas (1996)		✓		
Larsen et al. (2002)		✓		✓
Xiang et al. (2008)		✓		
Maxwell et al. (2010)		✓		✓
Smith, Pavone, Bullo, & Frazzoli (2010)		✓		
Pavone, Frazzoli, & Bullo (2011)		✓		
Schmid (2012)		✓		✓
Zhang, Ohlmann, & Thomas (2018)		✓		✓
Ulmer et al. (2021)		✓		✓

Table A5
Modeling stochastic demand quantities.

Work	Discrete	Continuous
Secomandi (2000)	✓	
Secomandi (2001)	✓	
Hvattum et al. (2007)		✓
Novoa & Storer (2009)	✓	
Secomandi & Margot (2009)	✓	
Pandelis et al. (2012)	✓	
Goodson et al. (2013)	✓	
Goodson, Thomas, & Ohlmann (2016)	✓	
Sarasola, Doerner, Schmid, & Alba (2016)	✓	
Brinkmann et al. (2019)	✓	

Table A6
Modeling stochastic customer requests.

Work	Discrete	Continuous		Time-dependent
		Uniform	Other	
Psarafitis (1980)		✓		
Powell, Sheffi, Nickerson, Butterbaugh, & Atherton (1988)	✓			
Bertsimas & Van Ryzin (1991)		✓		
Bertsimas & Van Ryzin (1993a)		✓		
Bertsimas & Van Ryzin (1993b)			✓	
Papastavrou (1996)		✓		
Tassiulas (1996)			✓	
Savelsbergh & Sol (1998)		✓		
Gendreau et al. (1999)		✓	✓	
Swihart & Papastavrou (1999)		✓		
Ichoua et al. (2000)		✓	✓	
Mahmassani, Kim, & Jaillet (2000)		✓		
Larsen et al. (2002)		✓		
Bent & Van Hentenryck (2003)		✓		
Bent & Van Hentenryck (2004)			✓	
Mitrović-Minić et al. (2004)		✓		
Thomas & White (2004)	✓			✓
Van Hemert & La Poutré (2004)			✓	✓
Yang, Jaillet, & Mahmassani (2004)		✓		
Branke et al. (2005)		✓		
Haghani & Jung (2005)	✓			
Chen & Xu (2006)		✓	✓	
Chen et al. (2006)		✓	✓	
Gendreau et al. (2006)	✓			✓
Hvattum et al. (2006)	✓			✓
Ichoua et al. (2006)			✓	✓
Potvin et al. (2006)		✓	✓	
Bent & Van Hentenryck (2007)			✓	
Hanshar & Ombuki-Berman (2007)		✓	✓	
Hvattum et al. (2007)			✓	
Thomas (2007)	✓			
Pavone, Frazzoli, & Bullo (2007)			✓	
Pureza & Laporte (2008)		✓	✓	
Sáez et al. (2008)			✓	
Xiang et al. (2008)		✓		
Angelelli et al. (2009)		✓	✓	
Branchini et al. (2009)			✓	✓
Ghiani et al. (2009)			✓	
Pavone et al. (2009)		✓		
Maxwell et al. (2010)			✓	✓
Smith et al. (2010)		✓		
Wen, Cordeau, Laporte, & Larsen (2010)			✓	
Lorini et al. (2011)		✓	✓	
Meisel et al. (2011)			✓	
Pavone et al. (2011)			✓	
Azi et al. (2012)		✓		
Schmid (2012)			✓	✓
Ferrucci et al. (2013)			✓	✓
Ghannadpour et al. (2013)		✓	✓	
Sheridan et al. (2013)		✓	✓	✓
Albareda-Sambola et al. (2014)	✓			
Ferrucci & Bock (2014)		✓		
Schilde et al. (2014)	✓			
Sarasola et al. (2016)		✓		
Ulmer, Mattfeld, Hennig, & Goodson (2016)		✓	✓	
Billing, Jaehn, & Wensing (2018)	✓			✓
Klapp et al. (2018)			✓	
Ulmer et al. (2018a)		✓	✓	
Ulmer et al. (2018b)		✓	✓	
Zou & Dessouky (2018)			✓	✓
Agussurja et al. (2019)		✓		
Arslan et al. (2019)		✓	✓	
Brinkmann et al. (2019)	✓			✓
Chen et al. (2021)			✓	✓
Soeffker et al. (2019)		✓		
Ulmer et al. (2019)		✓	✓	
Van Heeswijk et al. (2019)			✓	
Voccia et al. (2019)		✓	✓	
Al-Kanj et al. (2020)			✓	✓
Joe & Lau (2020)	✓			
Riley et al. (2020)	✓			
Ulmer (2020)		✓	✓	
Ulmer & Thomas (2020)			✓	
Ulmer et al. (2020a)	✓			
Demirbilek et al. (2021)			✓	
Ulmer et al. (2021)	✓			

Table A7
Summary.

	Problem domain	Uncertainty	Solution approach
Psarafitis (1980)	p	r	-
Powell et al. (1988)	g	r	-
Bertsimas & Van Ryzin (1991)	s	r, s	I
Bertsimas & Van Ryzin (1993a)	s	r, s	I
Bertsimas & Van Ryzin (1993b)	s	r, s	I
Papastavrou (1996)	s	r, s	-
Tassioulas (1996)	s	r, s	I
Savelsbergh & Sol (1998)	g	r	-
Gendreau et al. (1999)	g	r	-
Swihart & Papastavrou (1999)	p	r	I
Ichoua et al. (2000)	g	r	I
Mahmassani et al. (2000)	g	r	-
Secomandi (2000)	g	d	IIa, IIb
Secomandi (2001)	g	d	IIa
Larsen et al. (2002)	s	r, s	-
Bent & Van Hentenryck (2003)	g	r	IIa
Bent & Van Hentenryck (2004)	g	r	IIa
Mitrović-Minić et al. (2004)	g	r	I
Thomas & White (2004)	g	r	IIb
Van Hemert & La Poutré (2004)	g	r	I
Yang et al. (2004)	g	r	-
Branke et al. (2005)	-	t	I
Haghani & Jung (2005)	g	r, t	-
Chen & Xu (2006)	g	r	-
Chen et al. (2006)	g	r, t	I
Gendreau et al. (2006)	g	r	-
Hvattum et al. (2006)	g	r	IIa
Ichoua et al. (2006)	g	r	IIa
Potvin et al. (2006)	g	r, t	I
Bent & Van Hentenryck (2007)	g	r	IIa
Hanshar & Ombuki-Berman (2007)	g	r	-
Hvattum et al. (2007)	g	r, d	IIa
Pavone et al. (2007)	s	r	I
Thomas (2007)	g	r	IIa
Pureza & Laporte (2008)	g	r, t	I
Sáez et al. (2008)	p	r	IIa
Xiang et al. (2008)	p	r, t, s, v	-
Angelelli et al. (2009)	g	r	I
Branchini et al. (2009)	g	r	I
Ghiani et al. (2009)	g	r	IIa
Novoa & Storer (2009)	g	d	IIa
Pavone et al. (2009)	s	r	I
Secomandi & Margot (2009)	g	d	IIb
Maxwell et al. (2010)	p, s	r, s	IIb
Smith et al. (2010)	s	r, s	I
Wen et al. (2010)	g	r	-
Lorini et al. (2011)	g	r, t	-
Meisel et al. (2011)	g	r	IIb
Pavone et al. (2011)	s	r, s	I
Azi et al. (2012)	g	r	IIa
Pandelis et al. (2012)	g	d	IIb
Schmid (2012)	p, s	r, s	IIb
Ferrucci et al. (2013)	g	r	IIa
Ghannadpour et al. (2013)	g	r, t	-
Goodson et al. (2013)	g	d	IIa
Sheridan et al. (2013)	p	r	I
Albareda-Sambola et al. (2014)	g	r	IIa
Ferrucci & Bock (2014)	g	r, t, v	I
Schilde et al. (2014)	p	r, t	IIa
Goodson et al. (2016)	g	d	IIa
Köster et al. (2015)	g	t	-
Kim et al. (2016)	g	t	IIa
Sarasola et al. (2016)	g	r, d	IIa
Ulmer et al. (2016)	g	r	IIa
Billing et al. (2018)	g	r	IIa
Klapp et al. (2018)	g	r	IIa
Köster et al. (2018)	g	t	IIa
Ulmer et al. (2018a)	g	r	IIb
Ulmer et al. (2018b)	g	r	IIb
Zhang et al. (2018)	s	s	IIa
Zou & Dessouky (2018)	g	r	IIa
Agussurja et al. (2019)	p	r	IIb
Arslan et al. (2019)	g	r, dr	-
Brinkmann et al. (2019)	g	d, r	IIa
Chen et al. (2021)	g	r	IIb
Soeffker et al. (2019)	g	r	IIb

(continued on next page)

Table A7 (continued)

	Problem domain	Uncertainty	Solution approach
Ulmer et al. (2019)	g	r	I1b
Van Heeswijk et al. (2019)	g	r	I1b
Voccia et al. (2019)	g	r	I1a
Al-Kanj et al. (2020)	p	r	I1b
Joe & Lau (2020)	g	r	I1b
Riley et al. (2020)	p	r	I1a
Ulmer (2020)	g	r	I1b
Ulmer & Thomas (2020)	g	r	I1b
Ulmer et al. (2020a)	g	r	I
Demirbilek et al. (2021)	s	r	I1a
Ulmer et al. (2021)	g	r, s	I

- Since the requests are not known in advance, listing status information about all customers is not applicable.
- Instead, information about the requests waiting to be served is considered as this depicts which demand has to be served.
- As discussed above, Maxwell et al. (2010) also use information about demand that cannot be served.
- In addition to known customer requests, many articles incorporate the expected future demand. The current time implicitly provides information about how much demand for the driver time is likely to occur in the remaining time horizon, articles considering the future demand in this form are denoted by a “(✓)” in Table A.1.

It can be seen that all articles consider information about the resources as well as about the demand. In the lower part of Table A.1, the literature is listed with the features used.

A3. Approximation architectures

Table A.2 lists approximation architectures used in the literature on stochastic dynamic vehicle routing and shows where they are used in the literature. We observe a broader range of architectures.

A4. Information model characteristics

In this section, we list the literature according to their source of uncertainty and depict how the information about this uncertainty was modeled. Table A.3 describes whether travel times are modeled being discrete or continuous. Also, it is denoted whether correlation between arcs close to each other and time-dependency are modeled. Table A.4 depicts how service times are modeled in the considered articles. It specifically describes whether the service times are modeled as being discrete or continuous and whether they are symmetric or asymmetric. In Table A.5, the literature considering stochastic demand quantities is listed and it is denoted whether the demand quantities are considered to be discrete or continuous. Table A.6 finally depicts how stochastic customer requests are modeled. Here, it is distinguished whether the requests are considered to be coming from discrete locations or from a continuous plane. For a continuous modeling, it is further separated whether the locations are distributed uniformly or in a different way. Also, we denote whether a time-dependency is considered.

A5. Summary

In this section, the resulting overview about the considered literature is provided in Table A.7. For each article, the problem domain, the considered uncertainty, and the solution approach are classified. The problem domain corresponds to one of the application areas goods (“g”), passengers (“p”), or services (“s”). The uncertainty considered can be any combination of demand quantity (“d”), requests (“r”), service time (“s”), travel time (“t”), vehicle

availability (“v”), or driver availability (“dr”). The classification of the solution approaches follows the categories described in the article (“-” for rolling-horizon reoptimization, and categories “I”, “I1a”, or “I1b”).

References

Agussurja, L., Cheng, S.-F., & Lau, H. C. (2019). A state aggregation approach for stochastic multiperiod last-mile ride-sharing problems. *Transportation Science*, 53, 148–166.

Al-Kanj, L., Nascimento, J., & Powell, W. B. (2020). Approximate dynamic programming for planning a ride-hailing system using autonomous fleets of electric vehicles. *European Journal of Operational Research*, 284, 1088–1106.

Albareda-Sambola, M., Fernández, E., & Laporte, G. (2014). The dynamic multiperiod vehicle routing problem with probabilistic information. *Computers and Operations Research*, 48, 31–39.

Angeles, E., Bianchessi, N., Mansini, R., & Speranza, M. G. (2009). Short term strategies for a dynamic multi-period routing problem. *Transportation Research Part C: Emerging Technologies*, 17, 106–119.

Arslan, A. M., Agatz, N., Kroon, L., & Zuidwijk, R. (2019). Crowdsourced delivery a dynamic pickup and delivery problem with ad hoc drivers. *Transportation Science*, 53, 222–235.

Azi, N., Gendreau, M., & Potvin, J. Y. (2012). A dynamic vehicle routing problem with multiple delivery routes. *Annals of Operations Research*, 199, 103–112.

Bent, R., & Van Hentenryck, P. (2003). Dynamic vehicle routing with stochastic requests. In *IJCAI* (pp. 1362–1363).

Bent, R., & Van Hentenryck, P. (2007). Waiting and relocation strategies in online stochastic vehicle routing. In *IJCAI* (pp. 1816–1821).

Bent, R. W., & Van Hentenryck, P. (2004). Scenario-based planning for partially dynamic vehicle routing with stochastic customers. *Operations Research*, 52, 977–987.

Bertsekas, D. P., & Tsitsiklis, J. N. (1996). *Neuro-dynamic programming*. Athena Scientific, Belmont.

Bertsimas, D., & Kallus, N. (2020). From predictive to prescriptive analytics. *Management Science*, 66, 1025–1044.

Bertsimas, D. J., & Van Ryzin, G. (1991). A stochastic and dynamic vehicle routing problem in the euclidean plane. *Operations Research*, 39, 601–615.

Bertsimas, D. J., & Van Ryzin, G. (1993a). Stochastic and dynamic vehicle routing in the euclidean plane with multiple capacitated vehicles. *Operations Research*, 41, 60–76.

Bertsimas, D. J., & Van Ryzin, G. (1993b). Stochastic and dynamic vehicle routing with general demand and interarrival time distributions. *Advances in Applied Probability*, 25, 947–978.

Billing, C., Jaehn, F., & Wensing, T. (2018). A multiperiod auto-carrier transportation problem with probabilistic future demands. *Journal of Business Economics*, 88, 1009–1028.

Branchini, R. M., Armentano, V. A., & Løkketangen, A. (2009). Adaptive granular local search heuristic for a dynamic vehicle routing problem. *Computers and Operations Research*, 36, 2955–2968.

Branke, J., Middendorf, M., Noeth, G., & Dessouky, M. (2005). Waiting strategies for dynamic vehicle routing. *Transportation Science*, 39, 298–312.

Brinkmann, J., Ulmer, M. W., & Mattfeld, D. C. (2019). Dynamic lookahead policies for stochastic-dynamic inventory routing in bike sharing systems. *Computers and Operations Research*, 106, 260–279.

Browne, C. B., Powley, E., Whitehouse, D., Lucas, S. M., Cowling, P. I., Rohlfshagen, P., et al. (2012). A survey of monte carlo tree search methods. *IEEE Transactions on Computational Intelligence and AI in Games*, 4, 1–43.

Chen, H.-K., Hsueh, C.-F., & Chang, M. S. (2006). The real-time time-dependent vehicle routing problem. *Transportation Research Part E: Logistics and Transportation Review*, 42, 383–408.

Chen, X., Thomas, B. W., & Hewitt, M. (2017). Multi-period technician scheduling with experience-based service times and stochastic customers. *Computers and Operations Research*, 82, 1–14.

Chen, X., Ulmer, M. W., & Thomas, B. W. (2021). Deep q-learning for same-day delivery with vehicles and drones. *European Journal of Operational Research*.

- Chen, Z.-L., & Xu, H. (2006). Dynamic column generation for dynamic vehicle routing with time windows. *Transportation Science*, 40, 74–88.
- Conboy, K., Mikalef, P., Dennehy, D., & Krogstie, J. (2020). Using business analytics to enhance dynamic capabilities in operations research: A case analysis and research agenda. *European Journal of Operational Research*, 281, 656–672.
- Dayarian, I., & Savelsbergh, M. (2020). Crowdsourcing and same-day delivery: Employing in-store customers to deliver online orders. *Production and Operations Management*, 29, 2153–2174.
- Demirbilek, M., Branke, J., & Strauss, A. K. (2021). Home healthcare routing and scheduling of multiple nurses in a dynamic environment. *Flexible Services and Manufacturing Journal*, 33, 253–280.
- Ferrucci, F., & Bock, S. (2014). Real-time control of express pickup and delivery processes in a dynamic environment. *Transportation Research Part B: Methodological*, 63, 1–14.
- Ferrucci, F., Bock, S., & Gendreau, M. (2013). A pro-active real-time control approach for dynamic vehicle routing problems dealing with the delivery of urgent goods. *European Journal of Operational Research*, 225, 130–141.
- Fishman, G. S. (2001). *Discrete-event simulation: Modeling, programming, and analysis*. Springer Science & Business Media, New York.
- Gendreau, M., Guertin, F., Potvin, J.-Y., & Séguin, R. (2006). Neighborhood search heuristics for a dynamic vehicle dispatching problem with pick-ups and deliveries. *Transportation Research Part C: Emerging Technologies*, 14, 157–174.
- Gendreau, M., Guertin, F., Potvin, J.-Y., & Taillard, E. (1999). Parallel tabu search for real-time vehicle routing and dispatching. *Transportation Science*, 33, 381–390.
- Ghannadpour, S. F., Noori, S., & Tavakkoli-Moghaddam, R. (2013). Multiobjective dynamic vehicle routing problem with fuzzy travel times and customers satisfaction in supply chain management. *IEEE Transactions on Engineering Management*, 60, 777–790.
- Ghiani, G., Manni, E., Quaranta, A., & Triki, C. (2009). Anticipatory algorithms for same-day courier dispatching. *Transportation Research Part E: Logistics and Transportation Review*, 45, 96–106.
- Goodson, J. C., Ohlmann, J. W., & Thomas, B. W. (2013). Rollout policies for dynamic solutions to the multivehicle routing problem with stochastic demand and duration limits. *Operations Research*, 61, 138–154.
- Goodson, J. C., Thomas, B. W., & Ohlmann, J. W. (2016). Restocking-based rollout policies for the vehicle routing problem with stochastic demand and duration limits. *Transportation Science*, 50, 591–607.
- Goodson, J. C., Thomas, B. W., & Ohlmann, J. W. (2017). A rollout algorithm framework for heuristic solutions to finite-horizon stochastic dynamic programs. *European Journal of Operational Research*, 258, 216–229.
- Haghani, A., & Jung, S. (2005). A dynamic vehicle routing problem with time-dependent travel times. *Computers and Operations Research*, 32, 2959–2986.
- Hanshar, F. T., & Ombuki-Berman, B. M. (2007). Dynamic vehicle routing using genetic algorithms. *Applied Intelligence*, 27, 89–99.
- Hindle, G., Kunc, M., Mortensen, M., Oztekin, A., & Vidgen, R. (2019). Business analytics: Defining the field and identifying a research agenda. *European Journal of Operational Research*, 281, 483–490.
- Holsapple, C., Lee-Post, A., & Pakath, R. (2014). A unified foundation for business analytics. *Decision Support Systems*, 64, 130–141.
- Hvattum, L. M., Løkketangen, A., & Laporte, G. (2006). Solving a dynamic and stochastic vehicle routing problem with a sample scenario hedging heuristic. *Transportation Science*, 40, 421–438.
- Hvattum, L. M., Løkketangen, A., & Laporte, G. (2007). A branch-and-regret heuristic for stochastic and dynamic vehicle routing problems. *Networks*, 49, 330–340.
- Ichoua, S., Gendreau, M., & Potvin, J. Y. (2000). Diversion issues in real-time vehicle dispatching. *Transportation Science*, 34, 426–438.
- Ichoua, S., Gendreau, M., & Potvin, J. Y. (2006). Exploiting knowledge about future demands for real-time vehicle dispatching. *Transportation Science*, 40, 211–225.
- Joe, W., & Lau, H. C. (2020). Deep reinforcement learning approach to solve dynamic vehicle routing problem with stochastic customers. In J. C. Beck, O. Buffet, J. Hoffmann, E. Karpas, & S. Sohrabi (Eds.), *Proceedings of the international conference on automated planning and scheduling: vol. 30* (pp. 394–402).
- Kim, G., Ong, Y. S., Cheong, T., & Tan, P. S. (2016). Solving the dynamic vehicle routing problem under traffic congestion. *IEEE Transactions on Intelligent Transportation Systems*, 17, 2367–2380.
- Klapp, M. A., Erera, A. L., & Toriello, A. (2018). The one-dimensional dynamic dispatch waves problem. *Transportation Science*, 52, 402–415.
- Köster, F., Ulmer, M. W., & Mattfeld, D. C. (2015). Cooperative traffic control management for city logistic routing. *Transportation Research Procedia*, 10, 673–682.
- Köster, F., Ulmer, M. W., Mattfeld, D. C., & Hasle, G. (2018). Anticipating emission-sensitive traffic management strategies for dynamic delivery routing. *Transportation Research Part D: Transport and Environment*, 62, 345–361.
- Kullman, N. D., Goodson, J. C., & Mendoza, J. E. (2021). Electric vehicle routing with public charging stations. *Transportation Science*, 55, 637–659.
- Larsen, A., Madsen, O., & Solomon, M. (2002). Partially dynamic vehicle routing models and algorithms. *Journal of the Operational Research Society*, 53, 637–646.
- Lepenioti, K., Bousdekis, A., Apostolou, D., & Mentzas, G. (2020). Prescriptive analytics: Literature review and research challenges. *International Journal of Information Management*, 50, 57–70.
- Lorini, S., Potvin, J.-Y., & Zufferey, N. (2011). Online vehicle routing and scheduling with dynamic travel times. *Computers and Operations Research*, 38, 1086–1090.
- Lustig, I., Dietrich, B., Johnson, C., & Dziekan, C. (2010). The analytics journey. *Analytics Magazine*, 3, 11–13.
- Mahmassani, H. S., Kim, Y., & Jaillet, P. (2000). Local optimization approaches to solve dynamic commercial fleet management problems. *Transportation Research Record*, 1733, 71–79.
- Maxwell, M. S., Restrepo, M., Henderson, S. G., & Topaloglu, H. (2010). Approximate dynamic programming for ambulance redeployment. *INFORMS Journal on Computing*, 22, 266–281.
- Meisel, S. (2011). Anticipatory optimization for dynamic decision making. *Operations Research/Computer Science Interfaces Series*: 51. Springer, New York.
- Meisel, S., Suppa, U., & Mattfeld, D. (2011). Serving multiple urban areas with stochastic customer requests. In H.-J. Krewski, B. Scholz-Reiter, & K. D. Thoben (Eds.), *Dynamics in logistics* (pp. 59–68). Springer, Berlin Heidelberg.
- Mitrović-Minić, S., Krishnamurti, R., & Laporte, G. (2004). Double-horizon based heuristics for the dynamic pickup and delivery problem with time windows. *Transportation Research Part B: Methodological*, 38, 669–685.
- Mortenson, M. J., Doherty, N. F., & Robinson, S. (2015). Operational research from taylorism to terabytes: A research agenda for the analytics age. *European Journal of Operational Research*, 241, 583–595.
- Novoa, C., & Storer, R. (2009). An approximate dynamic programming approach for the vehicle routing problem with stochastic demands. *European Journal of Operational Research*, 196, 509–515.
- Pandelis, D. G., Kyriakidis, E. G., & Dimitrakos, T. D. (2012). Single vehicle routing problems with a predefined customer sequence, compartmentalized load and stochastic demands. *European Journal of Operational Research*, 217, 324–332.
- Papastavrou, J. D. (1996). A stochastic and dynamic routing policy using branching processes with state dependent immigration. *European Journal of Operational Research*, 95, 167–177.
- Pavone, M., Bisnik, N., Frazzoli, E., & Isler, V. (2009). A stochastic and dynamic vehicle routing problem with time windows and customer impatience. *Mobile Networks and Applications*, 14, 350–364.
- Pavone, M., Frazzoli, E., & Bullo, F. (2007). Decentralized algorithms for stochastic and dynamic vehicle routing with general demand distribution. In *2007 46th IEEE conference on decision and control* (pp. 4869–4874). IEEE.
- Pavone, M., Frazzoli, E., & Bullo, F. (2011). Adaptive and distributed algorithms for vehicle routing in a stochastic and dynamic environment. *IEEE Transactions on Automatic Control*, 56, 1259–1274.
- Potvin, J.-Y., Xu, Y., & Benyahia, I. (2006). Vehicle routing and scheduling with dynamic travel times. *Computers and Operations Research*, 33, 1129–1137.
- Powell, W. B. (2011). *Approximate dynamic programming: Solving the curses of dimensionality*. Wiley Series in Probability and Statistics: 842. New York: John Wiley & Sons.
- Powell, W. B. (2019). A unified framework for stochastic optimization. *European Journal of Operational Research*, 275, 795–821.
- Powell, W. B., Sheffi, Y., Nickerson, K. S., Butterbaugh, K., & Atherton, S. (1988). Maximizing profits for north american van lines' truckload division: A new framework for pricing and operations. *Interfaces*, 18, 21–41.
- Psaraftis, H. N. (1980). A dynamic programming solution to the single vehicle many-to-many immediate request dial-a-ride problem. *Transportation Science*, 14, 130–154.
- Psaraftis, H. N., Wen, M., & Kontovas, C. A. (2016). Dynamic vehicle routing problems: Three decades and counting. *Networks*, 67, 3–31.
- Pureza, V., & Laporte, G. (2008). Waiting and buffering strategies for the dynamic pickup and delivery problem with time windows. *INFOR: Information Systems and Operational Research*, 46, 165–175.
- Riley, C., Van Hentenryck, P., & Yuan, E. (2020). Real-time dispatching of large-scale ride-sharing systems: Integrating optimization, machine learning, and model predictive control. arXiv preprint arXiv:2003.10942
- Sáez, D., Cortés, C. E., & Núñez, A. (2008). Hybrid adaptive predictive control for the multi-vehicle dynamic pick-up and delivery problem based on genetic algorithms and fuzzy clustering. *Computers and Operations Research*, 35, 3412–3438.
- Sarasola, B., Doerner, K. F., Schmid, V., & Alba, E. (2016). Variable neighborhood search for the stochastic and dynamic vehicle routing problem. *Annals of Operations Research*, 236, 425–461.
- Savelsbergh, M., & Sol, M. (1998). Drive: Dynamic routing of independent vehicles. *Operations Research*, 46, 474–490.
- Schilde, M., Doerner, K. F., & Hartl, R. F. (2014). Integrating stochastic time-dependent travel speed in solution methods for the dynamic dial-a-ride problem. *European Journal of Operational Research*, 238, 18–30.
- Schmid, V. (2012). Solving the dynamic ambulance relocation and dispatching problem using approximate dynamic programming. *European Journal of Operational Research*, 219, 611–621.
- Secomandi, N. (2000). Comparing neuro-dynamic programming algorithms for the vehicle routing problem with stochastic demands. *Computers and Operations Research*, 27, 1201–1225.
- Secomandi, N. (2001). A rollout policy for the vehicle routing problem with stochastic demands. *Operations Research*, 49, 796–802.
- Secomandi, N., & Margot, F. (2009). Reoptimization approaches for the vehicle-routing problem with stochastic demands. *Operations Research*, 57, 214–230.
- Sheridan, P. K., Gluck, E., Guan, Q., Pickles, T., Balcioglu, B., & Benhabib, B. (2013). The dynamic nearest neighbor policy for the multi-vehicle pick-up and delivery problem. *Transportation Research Part A: Policy and Practice*, 49, 178–194.
- Shmueli, G., & Koppius, O. R. (2011). Predictive analytics in information systems research. *MIS Quarterly*, 35, 553–572.
- Smith, S. L., Pavone, M., Bullo, F., & Frazzoli, E. (2010). Dynamic vehicle routing with priority classes of stochastic demands. *SIAM Journal on Control and Optimization*, 48, 3224–3245.
- Soeffker, N., Ulmer, M. W., & Mattfeld, D. C. (2019). Adaptive state space partitioning for dynamic decision processes. *Business & Information Systems Engineering*, 61, 261–275.

- Song, Y., Ulmer, M. W., Thomas, B. W., & Wallace, S. W. (2020). Building trust in home services stochastic team-orienting with consistency constraints. *Transportation Science*, 54, 823–838.
- Spivey, M. Z., & Powell, W. B. (2004). The dynamic assignment problem. *Transportation Science*, 38, 399–419.
- Swihart, M. R., & Papastavrou, J. D. (1999). A stochastic and dynamic model for the single-vehicle pick-up and delivery problem. *European Journal of Operational Research*, 114, 447–464.
- Tassiulas, L. (1996). Adaptive routing on the plane. *Operations Research*, 44, 823–832.
- Thomas, B. W. (2007). Waiting strategies for anticipating service requests from known customer locations. *Transportation Science*, 41, 319–331.
- Thomas, B. W., & White, C. C., III (2004). Anticipatory route selection. *Transportation Science*, 38, 473–487.
- Ulmer, M., Nowak, M., Mattfeld, D., & Kaminski, B. (2020a). Binary driver-customer familiarity in service routing. *European Journal of Operational Research*, 286, 477–493.
- Ulmer, M. W. (2017). *Approximate dynamic programming for dynamic vehicle routing volume 61*. Springer, Cham.
- Ulmer, M. W. (2020). Dynamic pricing and routing for same-day delivery. *Transportation Science*, 54, 1016–1033.
- Ulmer, M. W., Goodson, J. C., Mattfeld, D. C., & Hennig, M. (2019). Offline–online approximate dynamic programming for dynamic vehicle routing with stochastic requests. *Transportation Science*, 53, 185–202.
- Ulmer, M. W., Goodson, J. C., Mattfeld, D. C., & Thomas, B. W. (2020b). On modeling stochastic dynamic vehicle routing problems. *EURO Journal on Transportation and Logistics*, 9, 100008.
- Ulmer, M. W., Mattfeld, D. C., Hennig, M., & Goodson, J. C. (2016). A rollout algorithm for vehicle routing with stochastic customer requests. In D. Mattfeld, T. Spengler, J. Brinkmann, & M. Grunewald (Eds.), *Logistics management* (pp. 217–227). Springer, Cham.
- Ulmer, M. W., Mattfeld, D. C., & Köster, F. (2018a). Budgeting time for dynamic vehicle routing with stochastic customer requests. *Transportation Science*, 52, 20–37.
- Ulmer, M. W., Soeffker, N., & Mattfeld, D. C. (2018b). Value function approximation for dynamic multi-period vehicle routing. *European Journal of Operational Research*, 269, 883–899.
- Ulmer, M. W., & Thomas, B. W. (2020). Meso-parametric value function approximation for dynamic customer acceptances in delivery routing. *European Journal of Operational Research*, 285, 183–195.
- Ulmer, M. W., Thomas, B. W., Campbell, A. M., & Woyak, N. (2021). The restaurant meal delivery problem: Dynamic pickup and delivery with deadlines and random ready times. *Transportation Science*, 55, 75–100.
- Van Heeswijk, W. J. A., Mes, M. R. K., & Schutten, J. M. J. (2019). The delivery dispatching problem with time windows for urban consolidation centers. *Transportation Science*, 53, 203–221.
- Van Hemert, J. I., & La Poutré, J. A. (2004). Dynamic routing problems with fruitful regions: Models and evolutionary computation. In X. Yao, E. K. Burke, J. e. A. Lozano, J. Smith, J. J. Merelo-Guervos, J. A. Bullinaria, ... H. P. Schwefel (Eds.), *International conference on parallel problem solving from nature* (pp. 692–701). Berlin Heidelberg: Springer.
- Vansteenwegen, P., Souffriau, W., & Van Oudheusden, D. (2011). The orienteering problem: A survey. *European Journal of Operational Research*, 209, 1–10.
- Voccia, S. A., Campbell, A. M., & Thomas, B. W. (2019). The same-day delivery problem for online purchases. *Transportation Science*, 53, 167–184.
- Waller, M. A., & Fawcett, S. E. (2013). Data science, predictive analytics, and big data: A revolution that will transform supply chain design and management. *Journal of Business Logistics*, 34, 77–84.
- Wen, M., Cordeau, J.-F., Laporte, G., & Larsen, J. (2010). The dynamic multi-period vehicle routing problem. *Computers and Operations Research*, 37, 1615–1623.
- Xiang, Z., Chu, C., & Chen, H. (2008). The study of a dynamic dial-a-ride problem under time-dependent and stochastic environments. *European Journal of Operational Research*, 185, 534–551.
- Yang, J., Jaillet, P., & Mahmassani, H. (2004). Real-time multivehicle truckload pickup and delivery problems. *Transportation Science*, 38, 135–148.
- Zhang, S., Ohlmann, J. W., & Thomas, B. W. (2018). Dynamic orienteering on a network of queues. *Transportation Science*, 52, 691–706.
- Zou, H., & Dessouky, M. M. (2018). A look-ahead partial routing framework for the stochastic and dynamic vehicle routing problem. *Journal on Vehicle Routing Algorithms*, 1, 73–88.