

Off-line approximate dynamic programming for the vehicle routing problem with stochastic customers and demands via decentralized decision-making

M. Dastpak, F. Errico

G-2021-53

September 2021

La collection *Les Cahiers du GERAD* est constituée des travaux de recherche menés par nos membres. La plupart de ces documents de travail a été soumis à des revues avec comité de révision. Lorsqu'un document est accepté et publié, le pdf original est retiré si c'est nécessaire et un lien vers l'article publié est ajouté.

The series *Les Cahiers du GERAD* consists of working papers carried out by our members. Most of these pre-prints have been submitted to peer-reviewed journals. When accepted and published, if necessary, the original pdf is removed and a link to the published article is added.

Citation suggérée : M. Dastpak, F. Errico (Septembre 2021). Off-line approximate dynamic programming for the vehicle routing problem with stochastic customers and demands via decentralized decision-making, Rapport technique, Les Cahiers du GERAD G-2021-53, GERAD, HEC Montréal, Canada.

Suggested citation: M. Dastpak, F. Errico (September 2021). Off-line approximate dynamic programming for the vehicle routing problem with stochastic customers and demands via decentralized decision-making, Technical report, Les Cahiers du GERAD G-2021-53, GERAD, HEC Montréal, Canada.

Avant de citer ce rapport technique, veuillez visiter notre site Web (<https://www.gerad.ca/fr/papers/G-2021-53>) afin de mettre à jour vos données de référence, s'il a été publié dans une revue scientifique.

Before citing this technical report, please visit our website (<https://www.gerad.ca/en/papers/G-2021-53>) to update your reference data, if it has been published in a scientific journal.

La publication de ces rapports de recherche est rendue possible grâce au soutien de HEC Montréal, Polytechnique Montréal, Université McGill, Université du Québec à Montréal, ainsi que du Fonds de recherche du Québec – Nature et technologies.

The publication of these research reports is made possible thanks to the support of HEC Montréal, Polytechnique Montréal, McGill University, Université du Québec à Montréal, as well as the Fonds de recherche du Québec – Nature et technologies.

Dépôt légal – Bibliothèque et Archives nationales du Québec, 2021
– Bibliothèque et Archives Canada, 2021

Legal deposit – Bibliothèque et Archives nationales du Québec, 2021
– Library and Archives Canada, 2021

GERAD HEC Montréal
3000, chemin de la Côte-Sainte-Catherine
Montréal (Québec) Canada H3T 2A7

Tél. : 514 340-6053
Télec. : 514 340-5665
info@gerad.ca
www.gerad.ca

Off-line approximate dynamic programming for the vehicle routing problem with stochastic customers and demands via decentralized decision-making

Mohsen Dastpak ^{a, b, c}

Fausto Errico ^{a, b, c}

^a *Department de génie de la construction, École de technologie supérieure, Montréal (Qc), Canada, H3C 1K3*

^b *GERAD, Montréal (Qc), Canada, H3T 1J4*

^c *CIRRELT, Montréal (QC), Canada, H3C 3J7*

mohsen.dastpak@cirrelt.ca

fausto.errico@cirrelt.ca

September 2021
Les Cahiers du GERAD
G–2021–53

Copyright © 2021 GERAD, Dastpak, Errico

Les textes publiés dans la série des rapports de recherche *Les Cahiers du GERAD* n'engagent que la responsabilité de leurs auteurs. Les auteurs conservent leur droit d'auteur et leurs droits moraux sur leurs publications et les utilisateurs s'engagent à reconnaître et respecter les exigences légales associées à ces droits. Ainsi, les utilisateurs:

- Peuvent télécharger et imprimer une copie de toute publication du portail public aux fins d'étude ou de recherche privée;
- Ne peuvent pas distribuer le matériel ou l'utiliser pour une activité à but lucratif ou pour un gain commercial;
- Peuvent distribuer gratuitement l'URL identifiant la publication.

Si vous pensez que ce document enfreint le droit d'auteur, contactez-nous en fournissant des détails. Nous supprimerons immédiatement l'accès au travail et enquêterons sur votre demande.

The authors are exclusively responsible for the content of their research papers published in the series *Les Cahiers du GERAD*. Copyright and moral rights for the publications are retained by the authors and the users must commit themselves to recognize and abide the legal requirements associated with these rights. Thus, users:

- May download and print one copy of any publication from the public portal for the purpose of private study or research;
- May not further distribute the material or use it for any profit-making activity or commercial gain;
- May freely distribute the URL identifying the publication.

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Abstract : This paper studies a stochastic variant of the vehicle routing problem (VRP) where both customer locations and demands are uncertain. In particular, potential customers are not restricted to a predefined customer set but are continuously spatially distributed in a given service area. The objective is to maximize the served demands while fulfilling vehicle capacities and time restrictions. We call this problem the VRP with stochastic customers and demands (VRPSCD). For this problem, we first propose a Markov Decision Process (MDP) formulation representing the classical *centralized* decision-making perspective where one decision-maker establishes the routes of all vehicles. While the resulting formulation turns out to be intractable, it provides us with the ground to develop a new MDP formulation of the VRPSCD representing a *decentralized* decision-making framework, where vehicles autonomously establish their own routes. This new formulation allows us to develop several strategies to reduce the dimension of the state and action spaces, resulting in a considerably more tractable problem. We solve the decentralized problem via Reinforcement Learning, and in particular, we develop a Q-learning algorithm featuring state-of-the-art acceleration techniques such as Replay Memory and Double Q Network. Computational results show that our method considerably outperforms two commonly adopted benchmark policies (random and heuristic). Moreover, when comparing with existing literature, we show that our approach can compete with specialized methods developed for the particular case of the VRPSCD where customer locations and expected demands are known in advance. Finally, we show that the value functions and policies obtained by our algorithm can be easily embedded in Rollout algorithms, thus further improving their performances.

Keywords: Stochastic vehicle routing problem, approximate dynamic programming, Q-learning, decentralized decision-making

1 Introduction

Freight distribution plays a crucial role in the supply chain system for its significant contribution to the overall operating cost, as well as for its impact on the service level provided to customers (Ghiani, Laporte, and Musmanno 2004). The current growth of e-commerce, by drastically increasing the volume of home deliveries, mostly in the form of parcel distribution, is further contributing to the relevance and the challenges of the last-mile transportation distribution. Fortunately, new tools are available to planners nowadays. On the one hand, advances in Information Technology (IT) enable fully connected distribution systems where vehicle and customer statuses are updated in real-time and shared across the system. On the other hand, the rapidly approaching advent of Autonomous Vehicles (AVs) and Unmanned Aerial Vehicles not only allows the planner to overcome several operating constraints afflicting traditional transportation, such as driver shift schedules but also opens the door to entirely different management paradigms, where traditional centralized decision making may shift towards a more autonomous and decentralized approach. It is not a surprise, however, that these new opportunities come with a completely new set of planning challenges, thus calling for new streams of research (see, for example Narayanan, Chaniotakis, and Antoniou 2020).

One of the core decision problems in freight distribution is known under the name of the vehicle routing problem (VRP, Toth and Vigo 2014). In its simpler definition, the VRP consists in dispatching a fleet of vehicles, initially located at a depot, to serve the demand of a known set of customers in such a way that a given performance measure is optimized (typically operations costs). In general, several side constraints must be fulfilled, such as vehicle capacities, time restrictions, Etc. The VRP has been intensively studied in the last 50 years, and the related literature is extremely rich (Vidal, Laporte, and Matl 2020). Most of this vast literature focuses on deterministic and static versions of the problem. However, when dealing with freight distribution in real-life contexts, a portion of the data may be unknown at the moment of planning, thus calling for stochastic versions of the VRP (SVRP). The most studied sources of uncertainty in literature are customer demands, customer locations, and travel and service times. For a complete review on stochastic VRPs, we refer the reader to Gendreau, Jabali, and Rei (2014), Oyola, Arntzen, and Woodruff (2018).

This paper introduces a stochastic variant of the VRP where both customer locations and demands are uncertain. We call this problem the *Vehicle Routing Problem with Stochastic Customers and Demands* (VRPSCD). In our setting, the planner has no information about the exact number of customers, their position within the service area, and their expected demands before the day of operation. Similarly to the work of Gendreau, Laporte, and Séguin (1995, 1996), the stochastic information is revealed in two phases: 1) at the beginning of each operating day, customer locations and expected demands become known, and 2) upon the visit to a customer, its actual demand is observed. We do not assume the knowledge of the probability distribution underlying the customers' locations and demand realization processes. However, we do assume the availability of historical data in the form of customer demand realizations. The VRPSCD then consists in dispatching a homogeneous fleet of vehicles initially located at the depot to serve the demands of the realized set of customers within a given deadline. Vehicles are allowed to perform so-called *preventive restocking* (see, for example Louveaux and Salazar-González 2018), thus enabling the possibility to return to the depot before the vehicle capacity is filled. The objective is to maximize the total served demand.

The way the VRPSCD models the uncertainty, alongside the considered information realization process, finds application in several real-world problems. Companies offering courier services, for example, typically possess statistical information about the expected number of requests and volumes for a given geographical area. However, the actual customer locations are only revealed shortly before operations. The second layer of uncertainty in the VRPSCD is given by customer demands. Typically, the uncertainty in customer demands is due to factors such as inaccurate descriptions of the parcel provided by customers and the adoption of simplified models to estimate the actual space a given set of parcels occupies in the vehicle (see, for example Côté, Gendreau, and Potvin 2020, Ghosal and Wiesemann 2020). Further examples of similar information realization processes can be found

in contexts such as waste collection for industrial applications, beverage distribution for convenience stores, fuel distribution for gas stations, and others.

Several streams of literature are related to the VRPSCD. The first one focuses on variants of the VRP where customers are known in advance while their demand is stochastic. Within this stream, we can identify two main problem settings. In the first, vehicles must fully serve all customer demands. Due to the uncertainty on the latter, vehicles may have to return to the depot to reload and continue the service. The objective is to minimize the expected total costs, which includes returns to the depot (see, for example Bertsimas 1992, Laporte et al. 2002, Secomandi 2000, 2001, Teodorović and Pavković 1992, Dror, Laporte, and Trudeau 1989, Louveaux and Salazar-González 2018, Florio, Hartl, and Minner 2020). In the second stream, similarly to the VRPSCD, a deadline on the duration of the operations is imposed. As a result, not all customers may be served, and the objective is to maximize the total demand serviced within the deadline (see, for example Goodson, Ohlmann, and Thomas 2013, Goodson, Thomas, and Ohlmann 2016, Erera, Morales, and Savelsbergh 2010, Mendoza, Rousseau, and Villegas 2016).

In addition to stochastic demands, the VRPSCD also considers stochastic customers. Stochastic customers and demands have been studied in Gendreau, Laporte, and Séguin (1995, 1996). The authors consider a predefined customer set with known locations. Customers are characterized by a service request probability, together with their expected demands. Similar to the VRPSCD introduced in this paper, customers and demands are revealed in two phases. The first phase reveals the subset of customers requesting service. The actual demands are observed upon visiting the customers. It is worth noting that stochastic customers' presence has also been studied in several works focusing on the Traveling Salesman Problem (see, for example Campbell and Thomas 2008, Jaillet 1988).

From the modeling perspective, this paper can be seen as a generalization of the work by Gendreau, Laporte, and Séguin (1995, 1996) on the one hand, and the work by Goodson, Ohlmann, and Thomas (2013), Goodson, Thomas, and Ohlmann (2016) on the other. In fact, we generalize the problem setting in Gendreau, Laporte, and Séguin (1995, 1996) by considering customer locations to be continuously spatially distributed in the service area, rather than restricted to a predefined discrete set. On the other hand, the problem in Goodson, Ohlmann, and Thomas (2013), Goodson, Thomas, and Ohlmann (2016) can be seen as a particular case of our model where customer locations and their expected demands are known in advance.

We can broadly classify solution methods for stochastic VRPs into two main categories (Gendreau, Jabali, and Rei 2016): the a priori optimization and the re-optimization categories. By somewhat generalizing the definition provided by Gendreau, Jabali, and Rei (2016), we say that a method belongs to the a priori category when it establishes a set of complete routes *here and now*, before the stochastic information is revealed. This feature is common to several mathematical frameworks for stochastic VRPs, including two-stage stochastic programming, chance-constraint programming, robust optimization, and distributionally robust optimization. These routes *may* be modified, for example in two-stage stochastic programming, as the stochastic information is revealed, according to a set of predefined rules called *recourse* policies. On the other hand, we say that a method belongs to the re-optimization paradigm when it builds routes dynamically by iteratively leveraging the information revealed during operations, rather than relying (solely) on a priori routes. It is worth noting that, according to this definition, solution methods providing operating policies such as those obtained by stochastic dynamic programming and related approximation methods are re-optimization methods.

In our problem setting, given that the set of potential customers is not known in advance, nor is their location, the a priori approach cannot be straightforwardly applied. For this reason, we make the choice to address the VRPSCD via the re-optimization paradigm. More specifically, we first provide a complete model of the VRPSCD in terms of a Markov Decision Process (MDP). This model represents the classical centralized decision-making perspective, where the decision-maker has full information on the vehicles' and clients' statuses and establishes the routes for all vehicles. The resulting model is, however, intractable even for small problem instances. The main challenge comes

from the multi-vehicle nature of the VRPSCD and the consequent explosion in the dimension of the state and action spaces. Nonetheless, this formulation provides us with the ground to develop a new *decentralized* MPD formulation of the VRPSCD, which allows us to reduce the dimension of the state and action spaces drastically. One could think of the resulting formulation as representing a different decision-making framework where vehicles autonomously decide the next action to be taken based on the information currently available to them. As it will be clarified in Section 4, while this formulation can be cast in terms of autonomously regulated vehicles, the resulting decision policy can also be applied in a centralized operating fashion. The decentralized formulation also opens the door to two approximation strategies. The first is based on the observation that two vehicles will choose the same action if they are in the same situation. As a result, instead of searching for one policy for each vehicle, we rather search for a single policy that is the same for all vehicles. We will see in Section 4 that this is crucial for the success of the proposed solution algorithm. In the second approximation strategy, the decentralized method facilitates the selection of information that is more relevant to the decision-making of a single vehicle (for example, the set of closest customers).

To find the resulting operational policy, we resort to Reinforcement Learning (Watkins and Dayan 1992, Sutton and Barto 2018), and in particular, we implement a Q-learning algorithm (Watkins and Dayan 1992), which is an approximate and model-free form of the Value Iteration Algorithm for stochastic dynamic programming (Bertsekas 2010). It is an *off-line* method in the sense that it provides an operational policy that is computed entirely before the beginning of the operations. Q-learning enables us to estimate the value of each state-action pair via simulation. In its simplest form, states and actions are discrete, and state-action values are represented as a lookup table. However, this paper proposes a continuous state representation and incorporates a two-layer artificial neural network to approximate the Q values. Our implementation features two main key elements. First, the challenges of variable-sized customer sets are addressed by adopting a fixed-size vector based on a particular heatmap-style encoding technique. Second, given the proposed decentralized model and our choice to develop a single policy for all vehicles, our algorithm is only required to maintain a single set of Q values, which is iteratively updated according to the simulation data from all vehicles. Finally, we develop several state-of-the-art techniques to improve the overall performance of our algorithm, including Replay Memory (Mnih et al. 2015) and Double Q Network (Van Hasselt, Guez, and Silver 2016). We call the resulting algorithm DecQN.

We conduct a thorough computational analysis. First, given that, to the best of our knowledge, no instance set is available for the VRPSCD in the literature, we construct a new instance set. We then implement two benchmark policies to compare with the DecQN. Results show that the DecQN considerably outperforms the considered benchmark policies. To provide a comparison with existing literature, we trained our DecQN on instances of the VRPSD, which is a particular case of the VRPSCD where the customer set is not stochastic. Results show that even though our algorithm is not tailored to exploit the in-advance knowledge of customer locations, it is competitive with the current state-of-the-art algorithm for the VRPSD (Goodson, Thomas, and Ohlmann 2016). Further computational analysis is then geared towards testing how the training phase of the DecQN can be guided towards generalization over different instances and problem dimensions, such as duration limits, stochastic variability, Etc. Finally, in the last experiment, we show that the obtained policies can be easily employed in an *on-line* algorithm, such as the Rollout Algorithm, thus further improving their performances.

To summarize, the contributions of the present work are as follows:

- We introduce the VRPSCD, which is a variant of the VRP with stochastic customers and demands. Our model generalizes existing works by allowing potential customers to be continuously distributed in a given service area instead of being constrained to a predefined set of customers.
- For this new problem, we provide a formulation based on MDPs, representing the traditional centralized decision-making framework. Given the fact that we are dealing with multiple vehicles, the resulting problem is intractable even for small problem instances.

- We then provide a decentralized MDP-based formulation of the VRPSCD. Among the main contributions of this paper is that this new decentralized formulation allows us not only to decompose the problem by vehicle, but also to search for a single policy that is the same for all vehicles. Another contribution descending from the decentralized formulation is that it enables us to develop computationally efficient state and action space aggregation strategies.
- We solve the resulting problem via a state-of-the-art Q-learning algorithm, called DecQN, featuring recent accelerating strategies such as Replay Memory (Mnih et al. 2015) and Double Q Networks (Van Hasselt, Guez, and Silver 2016). It is worth noting that the learning nature of the DecQN, together with the decentralized formulation, is a key factor for the efficiency of the overall algorithm because simulated experiences can be shared among vehicles.
- We perform an extensive computational experiment. Results show that the DecQN clearly outperforms two benchmark policies for VRPSCD. Furthermore, when tested on instances of the VRPSD, DecQN is competitive with state-of-the-art methods specifically designed for that problem. Moreover, we show that the training phase of the DecQN can be generalized to tackle problem instances varying in terms of several dimensions, such as duration limits and stochastic variability. Finally, we show that the obtained off-line policies are suitable to be used as base policies in on-line algorithms, such as the Rollout Algorithm, thus further improving their performances.

The paper is organized as follows. We begin with a detailed literature review in Section 2. In Section 3 we describe the problem in detail by introducing the notations and providing the first centralized MDP formulation. We then present the decentralized approach in Section 4 and describe the proposed Q-learning algorithm in Section 5. In Section 6 the experimental results are presented. Finally, we conclude this research in Section 7.

2 MDP-based solution methods for stochastic and dynamic VRPs

This section reviews works addressing stochastic or dynamic versions of the VRP and adopting the re-optimization paradigm according to the provided definition, i.e., methods building solutions dynamically by iteratively leveraging the information revealed during operations. More specifically, given the existing links with our approach, we focus on MDP-based solution methods, including approximate dynamic programming, reinforcement learning, and multi-agent reinforcement learning. For a more general review on re-optimization methods, the reader is referred to Psaraftis, Wen, and Kontovas (2016), Ritzinger, Puchinger, and Hartl (2016).

Dynamic programming (DP)(Bellman 1958) provides a solid theoretical framework and a family of algorithms to solve decision problems stated in terms of MDPs. Unfortunately, classical DP algorithms, such as Value Iteration and Policy Iteration, tend to be affected by the so-called curses of dimensionality (Powell 2011); thus, their applicability is limited to relatively small-sized decision problems. Furthermore, these algorithms assume the knowledge of the probability distribution governing the evolution of the system at hand, although this assumption rarely holds in real-life problems. Several approaches have been proposed in the literature to overcome the above difficulties; some limit the state-space exploration to a subset of the possible system states (Secomandi and Margot 2009), others provide approximated versions of the classical DP algorithms, e.g., the Approximate Value Iteration such as the Q-learning (Watkins and Dayan 1992), or the Approximate Policy Iteration (API) such as Policy Gradients (Williams 1992, Sutton et al. 2000). Generally speaking, depending on how the computational effort is distributed, whether during or before operations, we can classify these methods into on-line and off-line algorithms, respectively.

The Rollout Algorithm (RA) is one of the most common on-line algorithms and adopts a base policy and simulation over a scenario sample to approximate the value function (Bertsekas, Tsitsiklis, and Wu 1997). Secomandi (2001) introduced an RA for the single-vehicle VRPSD, where the vehicle must serve all customers and minimize the total travel time. The same problem was addressed by Novoa and

Storer (2009). They improved the work of Secomandi (2001) by employing a different base policy and using a two-step, rather than the classical one-step, rollout. Fan, Wang, and Ning (2006) addressed the VRPSD with multiple vehicles. They statically partitioned the customer set according to a geographical criterion and applied the method by Secomandi (2001) on each partition. Goodson, Ohlmann, and Thomas (2013), Goodson, Thomas, and Ohlmann (2016) extended this problem by incorporating a duration limit within which all operations must be completed. Unlike previous works, the objective function is to maximize the total served demand. Goodson, Ohlmann, and Thomas (2013) proposed a fixed-route heuristic comprising a neighborhood search heuristic to generate alternative routes and an approximate method to estimate the value of each route. They embedded the fixed-route heuristic into an RA. Inspired by the decomposition proposed by Fan, Wang, and Ning (2006), they presented a dynamic decomposition method to dynamically re-partition customers among vehicles during operations. In particular, customers are assigned to the same partition if they are visited by the same route according to the results of the fixed-route heuristic. Goodson, Thomas, and Ohlmann (2016) propose a new procedure, called restocking fixed-route heuristic, to estimate the expected value of a fixed route. Using an auxiliary dynamic program, they provide route evaluations both with or without preventive restocking. Similar to Goodson, Ohlmann, and Thomas (2013), they could not solve instances with more than 25 customers and used the dynamic decomposition method, enabling them to handle instances with up to 100 customers.

In addition to the RA, other on-line methods have been proposed in the literature. For the single-vehicle VRPSD, Secomandi and Margot (2009) presented the Partial Re-optimization method, where a subset of states is heuristically selected and evaluated by backward DP. For the VRPSD with two vehicles, Zhu et al. (2014) introduced the Paired Cooperative Re-optimization method inspired by the Partial Re-optimization approach proposed by Secomandi and Margot (2009). The authors formulated the problem as a bilevel MDP and suggested a heuristic procedure to approximate the value function.

We can find several off-line methods that have been successfully applied to SVRPs (Powell and Topaloglu 2006, Godfrey and Powell 2002). For the single-vehicle VRPSD, Secomandi (2000) proposed an API method and compared it with an RA. For the single-vehicle VRP with stochastic customers, Nazari et al. (2018) presented an Actor-Critic method. The authors assume that potential customers' locations are unknown, and customers can require service at any point in the service area. The authors encode the stochastic customer set via an Attention mechanism to obtain a representation in terms of a vector whose size is invariant with the number of customers. Stochastic customer sets have also been addressed in the context of the TSP. Dai et al. (2017) adopt structure2vec for graph embedding, Kool, van Hoof, and Welling (2018) propose an attention encoder, while Miki, Yamamoto, and Ebara (2018) implement a Convolutional Neural Network.

The literature on re-optimization methods for VRPSD with multiple vehicles is relatively scarce. In particular, to the best of our knowledge, no off-line re-optimization method has been proposed for this problem in the literature. However, for the multi-vehicle VRP with deterministic demands and stochastic customer realizations, a multi-agent attention model has been proposed in Zhang et al. (2020). The model is trained according to the policy gradient method and provides a policy that generates a set of routes.

Multi-agent RL-based algorithms have also been introduced for the Dial-a-Ride Problem with multiple vehicles and stochastic orders (Qin et al. 2020, Tang et al. 2019, Kullman et al. 2020, Holler et al. 2019). Qin et al. (2020) and Tang et al. (2019) implemented Q values in the form of one Deep Q Network (DQN) for each vehicle and used a central combinatorial optimization problem as a coordinator to assign orders to vehicles. Kullman et al. (2020) adopted an attention encoder-decoder as the central coordinator and trained the model with Actor-Critic. For a similar problem, Holler et al. (2019) compared Actor-Critic and DQN methods without observing significant performance differences.

Summarizing the above literature review, we can say that re-optimization algorithms seem to be promising for the solution of stochastic and dynamic versions of the VRP and have received increased

attention, especially during the last decade. However, to the best of our knowledge, no re-optimization method exists in the literature for the VRPSCD introduced in this paper. The VRPSCD generalizes several existing problems, such as those introduced in Gendreau, Laporte, and Séguin (1995, 1996) and Goodson, Ohlmann, and Thomas (2013), Goodson, Thomas, and Ohlmann (2016), and it is an extremely challenging problem. The fact that the set of potential customers is not known in advance makes it difficult to adopt route-based policies. On the other hand, the fact that the number of customers is stochastic requires particular techniques to enable fixed-size vector representations to be suitably adopted in a Reinforcement Learning algorithm. We argue that the decentralized formulation proposed in this paper, together with the adoption of the heatmap-style representation of the stochastic customer set, are original contributions to the literature and essential components for the success of the proposed approach.

3 The centralized VRPSCD

In Section 3.1 we formally define the VRPSCD and introduce some notation. The centralized MDP formulation is provided in Section 3.2.

3.1 Problem definition

We are given a service area and a depot located in the service area where a set \mathcal{V} of m homogeneous vehicles with capacity Q is initially placed. Customers may request service from any point in the service area. The set of customers requesting service \mathcal{C}_0 is revealed at the beginning of the operation period (e.g., every morning). We denote n the number of realized customers ($n = |\mathcal{C}_0|$). Each customer $c \in \mathcal{C}_0$ is characterized by a location l_c , and a probabilistic demand whose expected value \bar{d}_c is assumed to be known. Accordingly, each customer is defined by a tuple of (l_c, \bar{d}_c) . The customer's actual demand \hat{d}_c is only realized upon the first visit. We assume that the customer's demand is splittable, meaning that if a vehicle is unable to fully serve the customer, it will do so to the fullest extent possible. The rest of the demand may be served by the same or a different vehicle at a later moment. The time taken to travel between locations a and b is assumed to be deterministic and is denoted $\tau_{a,b}$. Vehicles begin their trips at the depot, serve a certain subset of customers, and end their routes at the depot. Furthermore, vehicles must complete their operations before a given duration limit L . Vehicles may return to the depot during their trip for restocking operations after leaving a customer and before visiting the next one, even if the vehicle capacity has not been filled yet. Such a policy is known in the literature as *restocking*. The objective is to maximize the expected total served demand.

3.2 The centralized formulation

In this section, we provide a centralized MDP formulation of the VRPSCD. We use the term centralized to refer to the classical decision-making approach, where routes are built by one central decision-maker. An MDP formulation generally consists of a *state* which represents the current status of the system, a control *action*, and a function describing the transition of the system to a new state once an action is taken and the exogenous information is revealed.

In our centralized MDP formulation, the *decision epoch* is any point in time at which either new information is revealed, such as at the beginning of the horizon, when the set of customers is revealed, and during operations, when at least one vehicle arrives at a customer, or when a vehicle arrives at the depot. At each decision epoch k , the status of the system is captured and represented by a state s_k . For the VRPSCD, we represent the state of the system as follows:

$$s_k = ([l_c, h_c, \bar{d}_c, \hat{d}_c]_{c \in \mathcal{C}_0}, [l_v, a_v, q_v]_{v \in \mathcal{V}}, t_k), \quad (1)$$

where the first two groups of components refer to the state of customers and vehicles, respectively, and the last component indicates the time.

Note that although the customer and vehicle state components are naturally indexed by the decision epoch k , we have simplified the notation by omitting the index k . The state of each customer $c \in \mathcal{C}_0$ is represented as a tuple of $(l_c, h_c, \bar{d}_c, \hat{d}_c)$. The first two components indicate the location and customer's availability, while the last two components are the customer's expected demand and the customer's unserved demand. The customer's availability h_c takes value 1 if the customer is available to be served and 0 otherwise. The customer's unserved demand indicates the volume of demand that is observed but not served yet. For each vehicle $v \in \mathcal{V}$, a tuple of (l_v, a_v, q_v) forms its state, where l_v , a_v , and q_v are the destination location, arrival time at destination, and the current available capacity, respectively. The time of the system at decision epoch k is denoted as t_k .

An m -dimensional vector of $x_k = (x_k^1, \dots, x_k^m)$ denotes the action set, where x_k^v indicates the next location vehicle v travels to. We define $\bar{\mathcal{V}}_k = \{v \in \mathcal{V} | a_v = t_k\}$ the subset of vehicles that arrive at their destinations and are available to take action at decision epoch k . We call $\bar{\mathcal{V}}_k$ the set of active vehicles. At decision epoch k , the action x_k belongs to the action space $A(s_k)$ defined as:

$$A(s_k) = \{x_k \in \{\mathcal{C}_0 \cup \{depot\}\}^m : \quad (2)$$

$$x_k^v = l_v, \quad \forall v \in \mathcal{V} \setminus \bar{\mathcal{V}}_k, \quad (3)$$

$$x_k^v = depot, \quad \forall \{v \in \bar{\mathcal{V}}_k | q_v = 0\}, \quad (4)$$

$$x_k^v \neq x_k^{v'}, \quad \forall \{v, v' \in \mathcal{V} | v \neq v', x_k^{v'} \neq depot\}, \quad (5)$$

$$x_k^v \in J(s_k, v) \cup depot, \quad \forall v \in \bar{\mathcal{V}}_k, \quad (6)$$

$$x_k^v \neq depot, \quad \forall \{v \in \bar{\mathcal{V}}_k | J(s_k, v) \setminus X_k^{-v} \neq \emptyset, l_v = depot\}, \quad (7)$$

where $X_k^{-v} = \{x_k^1, \dots, x_k^{v-1}, x_k^{v+1}, \dots, x_k^m\}$ denotes the set of actions at time k for all vehicles except vehicle v , and $J(s_k, v)$ the set of reachable customers from vehicle v that have not been served before:

$$J(s_k, v) = \{c \in \mathcal{C}_0 | h_c = 1, \tau_{v,c} + \tau_{c,depot} \leq L - t_k\}. \quad (8)$$

In the defined action space $A(s_k)$, Condition (3) enforces the fact that vehicles cannot be relocated when assigned to a destination. Condition (4) implies that vehicles with no available capacity should return to the depot to replenish. Based on Condition (5), no two vehicles can choose the same destination unless the destination is the depot. Also, Condition (6) obliges vehicles to choose actions from the set of reachable customers $J(s_k, v)$ and the depot. Finally, in Condition (7), waiting in the depot is only allowed when there is no feasible customer to visit.

The information about customers' presence and demand volumes is revealed in two phases. At the first decision epoch ($k = 0$), no customer is realized yet. Consequently,

$$s_0 = ([], [(depot, 0, Q)]_{v \in \mathcal{V}}, t_0).$$

At this point, a dummy action $x_0 = (depot)_{v \in \mathcal{V}}$ transfers the system to the next decision epoch ($k = 1$) where the system time and the vehicle states remain the same; however, a scenario of realized locations and expected demands populates the set of customers \mathcal{C}_0 :

$$s_1 = ([(l_c, 1, \bar{d}_c, -1)]_{c \in \mathcal{C}_0}, [(depot, 0, Q)]_{v \in \mathcal{V}}, t_1).$$

For each realized customer $c \in \mathcal{C}_0$, we set \hat{d}_c to -1 at the beginning and will update it upon its first visit. When \hat{d}_c is zero, it means that the customer is fully served and no longer needs any service. Also, h_c is initially set to one for all customers; however, if customer c is already fully served or a vehicle is on its way to serve it, h_c takes the value of zero.

It is worth noting that the way the state s_0 is initialized highlights one of the main features of the VRPSCD. In this MDP formulation, we are looking to develop a policy that accounts for all possible customer realizations and can tackle those realizations that have never been seen before. Accordingly, the system starts with no customers, and a random set of customers will be revealed in decision epoch

1. From this decision epoch on ($k \geq 1$), \mathcal{C}_0 is assumed to be fixed. Taking the action x_k from the set of possible actions in the state s_k denoted as $A(s_k)$ transits the system to the next state s_{k+1} . In addition to the current state and the taken action, the next state s_{k+1} is also dependent on the exogenous information w_{k+1} .

$$s_{k+1} = S^M(s_k, x_k, w_{k+1}).$$

We let w be a sample scenario of customer demands and $w_k \subset w$ be the set of realized demand volumes of those customers whose actual demands are observed at decision epoch k . To elaborate on the transition function $S^M(\cdot)$, we divide it into two steps. In the first step, the post-decision state s_k^x is the state immediately after taking action x_k , but before revealing the exogenous information w_{k+1} :

$$h_{x_k^v} = 0, \quad \forall \{v \in \bar{\mathcal{V}}_k | x_k^v \neq depot\}. \quad (9)$$

$$a_v = t_k + \tau_{v, x_k^v}, \quad \forall v \in \bar{\mathcal{V}}_k. \quad (10)$$

$$l_v = l_{x_k^v}, \quad \forall v \in \bar{\mathcal{V}}_k. \quad (11)$$

$$t_{k+1} = \min_{v \in \mathcal{V}} a_v. \quad (12)$$

Equation (9) flags the chosen customers as unavailable. The arrival time as well as the location of active vehicles are updated using Equations (10) and (11). Finally, Equation (12) progresses the time to the next decision epoch, where the exogenous information w_{k+1} will be revealed. All the other components are left unmodified.

The second step of the transition function happens at the beginning of a new decision epoch ($k+1$), when the uncertain demands are revealed, by transiting from the post-decision state s_k^x to the next state s_{k+1} . Let $\tilde{\mathcal{C}}_{k+1} = \{c \in \mathcal{C}_0 | \exists v \in \bar{\mathcal{V}}_{k+1} \wedge l_c = l_v\}$ be the set of customers that are being served at decision epoch $k+1$. If the actual demand of a customer in $\tilde{\mathcal{C}}_{k+1}$ is not yet realized ($\hat{d}_c = -1$), its value will be set to the observed demand w_{k+1}^c .

$$\hat{d}_c = w_{k+1}^c, \quad \forall \{c \in \tilde{\mathcal{C}}_{k+1} | \hat{d}_c = -1\}. \quad (13)$$

Let η_v^{k+1} be the demand volume that vehicle $v \in \bar{\mathcal{V}}_{k+1}$ serves at decision epoch $k+1$, defined as:

$$\eta_v^{k+1} = \min\{\hat{d}_{c_v}, q_v\}, \quad \forall \{v \in \bar{\mathcal{V}}_{k+1} | l_v \neq depot\}, \quad (14)$$

where c_v refers to a customer $c \in \tilde{\mathcal{C}}_{k+1}$ that is served by vehicle v (i.e. $l_c = l_v$). The unserved demands of customers in $\tilde{\mathcal{C}}_{k+1}$ and the available capacity of vehicles in $\bar{\mathcal{V}}_{k+1}$ are then updated as follows:

$$\hat{d}_{c_v} = \hat{d}_{c_v} - \eta_v^{k+1}, \quad \forall \{v \in \bar{\mathcal{V}}_{k+1} | l_v \neq depot\}. \quad (15)$$

$$q_v = \begin{cases} q_v - \eta_v^{k+1} & \text{if } l_v \neq depot, \\ Q & \text{otherwise.} \end{cases} \quad \forall v \in \bar{\mathcal{V}}_{k+1}. \quad (16)$$

Also, customers in $\tilde{\mathcal{C}}_{k+1}$ whose demands are not fully served will be available again to be served later:

$$h_c = 1, \quad \forall \{c \in \tilde{\mathcal{C}}_{k+1} | \hat{d}_c > 0\}. \quad (17)$$

These two steps are illustrated in Figure 1 using an example of 4 customers and 2 vehicles. We define the terminal decision epoch K as a point of time at which all vehicles are located at the depot, and vehicles have no other choice than to stay in the depot ($A(s_K) = \{(depot)_{v \in \mathcal{V}}\}$).

The reward function $R(s_k, x_k, w)$ is defined as the demands to be served by taking action x_k in state s_k , based on the exogenous information w :

$$R(s_k, x_k, w) = \sum_{\{v \in \bar{\mathcal{V}} | x_k^v \neq depot\}} \min\{q_v, d_{x_k^v}^w\} \quad (18)$$

$$d_c^w = \begin{cases} w_k^c & \text{if } \hat{d}_c = -1, \\ \hat{d}_c & \text{otherwise.} \end{cases} \quad \forall c \in \mathcal{C}_0, \quad (19)$$

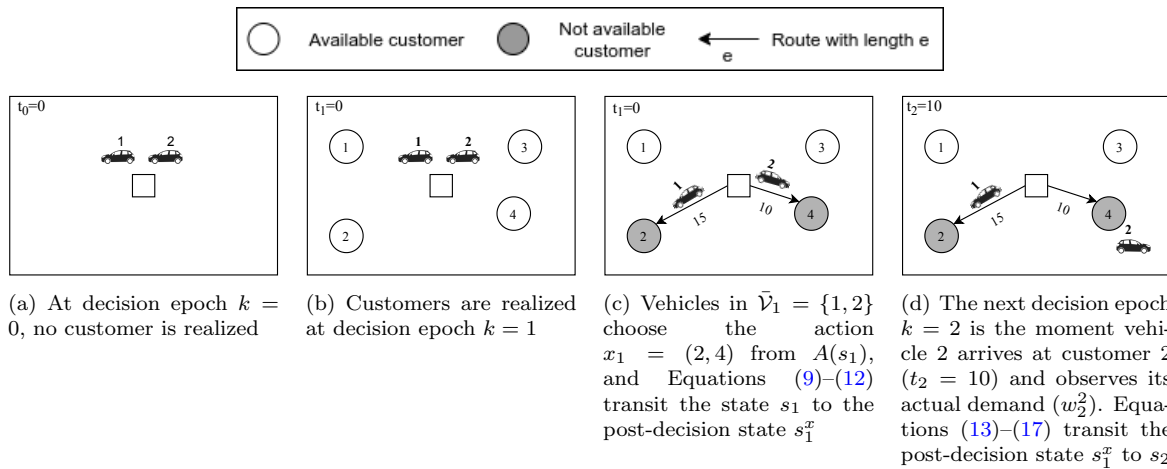


Figure 1: The transition function in a centralized MDP formulation

where $w_{\bar{k}}^c$ is the demand of customer c observed when the vehicle v visits the customer c at the decision epoch \bar{k} . The decision epoch \bar{k} refers to the moment when the vehicle $v \in \bar{\mathcal{V}}$, which has selected customer x_k^v at the decision epoch k , arrives at its destination:

$$\bar{k} = \{k' \in [0, \dots, K] | t_{k'} = t_k + \tau_{v, x_k^v}\}.$$

For an active vehicle that decides in k , the decision epoch \bar{k} can be further than $k + 1$ in time, and other vehicles may decide while it is en route. However, it is worth noting that the reward for its action at decision epoch k is solely determined by its state, action, and the delayed observed demands, not by the other vehicles.

In this formulation, the value of being in state s_k is defined as $V(s_k)$. This function indicates the expected cumulative reward the system can collect from state s_k onward:

$$V(s_k) = \max_{x_k \in A(s_k)} \mathbb{E}_w [R(s_k, x_k, w) + \gamma V(s_{k+1})], \quad (20)$$

where γ is a discount factor and we remind that $s_{k+1} = S^M(s_k, x_k, w_{k+1})$.

There are some challenges making this formulation intractable for the VRPSCD. First, the dimension of the state representation is dependent on the cardinality of customers, which is stochastic. Consequently, the size of the vector s_k would be variable and dependent on the number of realized customers, preventing the use of many solution methods that assume a fixed-size state. For example, no linear regression method can be implemented to approximate the value function in Equation (20), when the size of the input (state vector) is variable. Furthermore, the variable number of customers combined with the continuous representation of their state greatly expands the state space. On the other side, having multiple vehicles makes the size of the action space, defined in Equations (2)–(7), excessively large. Denoting \underline{n} the number of customers available to be served at decision epoch k , the size of the action space will be $\frac{\underline{n}!}{(\underline{n}-m)!}$. For example, for 8 vehicles and 50 customers, 2×10^{13} various actions are available to be taken. Even for mid-size problems, handling such a large action set is computationally prohibitive.

The following section explains how the decentralized formulation allows us to cope with the challenges mentioned above.

4 The decentralized VRPSCD

The main idea behind decentralizing the VRPSCD is to allow vehicles to establish the next customer to be visited based on partial observations of the surrounding environment. In our approach, we enforce

the restriction that only one vehicle, called the *active vehicle*, is allowed to make a decision at a given decision epoch. In particular, the vehicle makes a decision based upon the observation of the system state from *its point of view*. While in this setting, vehicles decide individually, the objective function is to maximize the demands collectively served by all vehicles. It is worth noting that although it is suggestive to think of the decentralized formulation in terms of self-deciding AVs, the resulting policy can also be operated centrally. In fact, given the full knowledge of the system, and according to the current vehicle states, the central dispatcher can compute the observed state relative to each vehicle and apply the decentralized policy to decide which is the next customer to be visited.

Decentralizing the problem has several advantages: first of all, since the decision is restricted to one vehicle only, the size of the action space will be drastically reduced to \underline{n} (from potentially up to $\frac{n}{(n-m)!}$), where \underline{n} is the number of available customers in the current decision epoch. Second, it allows us to represent the system state based on the point of view of the active vehicle by ignoring the less relevant information for the active vehicle. This enables us to reduce the size of the state space hugely. Third, given the assumption that the operating fleet is homogeneous, and the fact that two identical vehicles would make the same action if the observed states are identical, the proposed approach creates one policy only, shared by all vehicles, instead of m policies.

Starting from the centralized formulation, we show how to obtain the decentralized formulation in Section 4.1. In Section 4.2, we provide a detailed description of the vehicle observation function and its encoding. Finally, in Section 4.3, we discuss cooperation issues and show how they are addressed in our decentralized approach.

4.1 The decentralized formulation

To model the decentralized VRPSCD, we define two levels of state representation. The first level is the *global state*. The global state s_k is the state that we presented in the centralized MDP formulation, and it is responsible for keeping track of the evolution of the entire system. In the second level, we introduce the observation function $O(s_k, \bar{v})$, which takes as input the global state s_k and the current active vehicle \bar{v} , and returns $o_{k,\bar{v}}$, which is the state as observed by vehicle \bar{v} :

$$o_{k,\bar{v}} = O(s_k, \bar{v}). \quad (21)$$

The observation $o_{k,\bar{v}}$ is comprised of four parts:

$$o_{k,\bar{v}} = [F_k | H_k | G_k | t_k], \quad (22)$$

where F_k is the state of a subset of customers that are judged more promising for \bar{v} , H_k is an aggregated overview of all customers in the service area, G_k is the state of vehicles, and t_k is the time at decision epoch k . Details of the observation function and the observed state will be provided in Section 4.2; however, the main point is that the size of the observed state is significantly smaller than the size of the global state.

In our approach, only one vehicle at a time is allowed to make a decision. If multiple vehicles are potentially active at decision epoch k , one is randomly chosen as the active vehicle \bar{v} . Therefore, the set of active vehicles \bar{V}_k reduces to one active vehicle \bar{v} .

In addition to the state representation, we redefine the action and the action space specifically for the active vehicle as $x_k = (x_k^{\bar{v}})$ and $A(s_k, \bar{v})$, respectively. We have:

$$A(s_k, \bar{v}) = \{x_k \in \{J(s_k, \bar{v}) \cup \{\text{depot}\}\} : \quad (23)$$

$$x_k = \text{depot} \quad \text{if } q_{\bar{v}} = 0, \quad (24)$$

$$x_k \neq \text{depot} \quad \text{if } J(s_k, \bar{v}) \neq \emptyset, l_{\bar{v}} = \text{depot}\}, \quad (25)$$

where Condition (23) obliges x_k to be a customer from the set of feasible customers $J(s_k, \bar{v})$ or the depot, Condition (24) forces the active vehicle to travel to the depot when it has no remaining capacity, and Condition (25) does not allow the vehicle to stay at the depot if there are reachable customers to be served.

As mentioned earlier, we use the global state s_k to keep track of the system dynamics. Therefore, the state transition function $S^M(\cdot)$ and its associated equations (Equations (9)-(17)) are still applicable here, given that $\bar{\mathcal{V}}_k = \{\bar{v}\}$. Figure 2 illustrates the two levels of state representation as well as their interactions for an example of two vehicles. At each decision epoch k , if the current location of the active vehicle \bar{v} is the depot, its capacity is replenished (Equation (16)); otherwise, the actual demand of the customer $c \in \tilde{\mathcal{C}}_k$ is realized (Equation (13)) and the active vehicle \bar{v} serves it to the fullest extent. At this point, the available capacity of \bar{v} (Equation (16)) as well as the unserved demand of the customer $c \in \tilde{\mathcal{C}}_k$ is updated (Equations (15) and (17)). Then, the active vehicle \bar{v} observes $o_{k,\bar{v}} = O(s_k, \bar{v})$ and chooses an action x_k from the action set $A(s_k, \bar{v})$ to travel. Equations (9)-(11) transit the global state s_k to the post-decision state s_k^x and Equation (12) updates the time of the next decision epoch. The operation terminates at decision epoch K where all vehicles have returned to the depot and there is no customer to be served or no time to serve any additional customers.

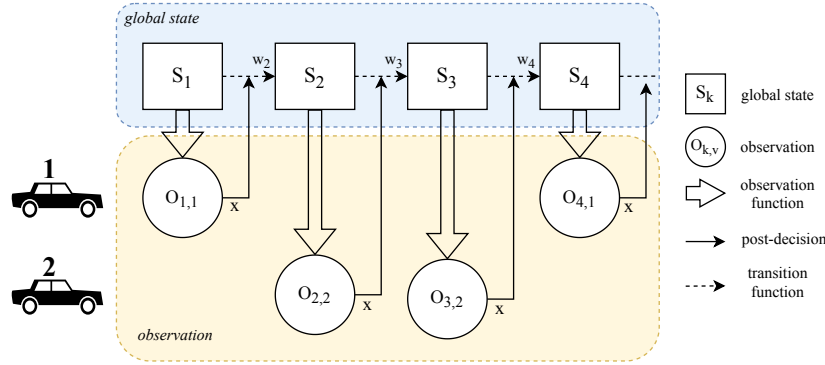


Figure 2: A schema of system dynamics in decentralized approach for an example of two vehicles

In our approach, we look for a policy π that indicates the best location to travel to for the active vehicle \bar{v} in the feasible action space $A(s_k, \bar{v})$, when the observation of the global state is $o_{k,\bar{v}}$. The resulting policy selects actions such that the total served demand by all vehicles is maximized. The optimal policy π^* can be obtained by solving the following equations for each observation $o_{k,\bar{v}}$ at decision epoch k :

$$V(o_{k,\bar{v}}) = \max_{x_k \in A(s_k, \bar{v})} \mathbb{E}_w[R(s_k, x_k, w) + \gamma V(O(s_{k+1}, \bar{v}'))], \quad (26)$$

$$\pi^*(o_{k,\bar{v}}) = \arg \max_{x_k \in A(s_k, \bar{v})} \mathbb{E}_w[R(s_k, x_k, w) + \gamma V(O(s_{k+1}, \bar{v}'))], \quad (27)$$

where \bar{v}' is the active vehicle in decision epoch $k+1$ and $V(o_{k,\bar{v}})$ is the expected reward-to-go starting from observation $o_{k,\bar{v}}$ onward. In the terminal global state s_K , there is no customer available to be served. Hence:

$$V(O(s_K, v)) = 0, \quad \forall v \in \mathcal{V}.$$

4.2 The observation function and its representation

The primary purpose of the observation function is to suitably select the most relevant information for the active vehicle and aggregate the remaining information. More specifically, we focus on the customers' state, which forms the most considerable portion of the information in the state representation s_k . For this purpose, we identify a subset of customers $\mathcal{C}_{\bar{v}} \subset \mathcal{C}_0$ that are judged more promising for the active vehicle and call it the set of *target customers*. Therefore, instead of keeping the in-detail information of all customers, the observation function only maintains the comprehensive information of target customers and aggregates the rest.

Customers in the set of target customers should have certain characteristics. For a target customer $c \in \mathcal{C}_{\bar{v}}$, the simplest desirable attributes are: 1) requesting large demand volumes (d_c), and 2) being

close to the active vehicle ($\tau_{c,\bar{v}}$). In our problem, customers with higher demands are more promising because of their potential contribution to the objective function. On the other hand, visiting customers closer to the active vehicle is more favorable due to the duration limit. Therefore, we define a new measure to identify the target customers. Accordingly, we choose \tilde{n} customers with higher $\rho_{c,\bar{v}}$:

$$\rho_{c,\bar{v}} = \frac{\min(\tilde{d}_c, q_{\bar{v}})}{\tau_{c,\bar{v}}},$$

$$\tilde{d}_c = \begin{cases} \bar{d}_c & \hat{d} = -1 \\ \hat{d}_c & \text{otherwise} \end{cases}, \quad (28)$$

where \tilde{n} is a hyper-parameter. This measure normalizes the customers' demands by their distance to the active vehicle, and also accounts for the available capacity. For example, for two customers $c_1, c_2 \in \mathcal{C}_0$, with $\tau_{c_1,\bar{v}} < \tau_{c_2,\bar{v}}$ and both demands greater than the available capacity ($d_{c_1} \geq q_{\bar{v}}, d_{c_2} \geq q_{\bar{v}}$), the closer one to \bar{v} (i.e., customer c_1) is more appealing ($\rho_{c_1,\bar{v}} > \rho_{c_2,\bar{v}}$), even if its demand is much smaller than the other customer ($d_{c_1} < d_{c_2}$). The effectiveness of selecting target customers according to their score $\rho_{c,\bar{v}}$ was validated through preliminary computational tests.

In order to keep in-detail information of target customers, we represent each customer using a set of features as $(l_c, \tau_{c,\bar{v}}, \tau_{c,d}, \tilde{d}_c, \min(\tilde{d}_c, q_{\bar{v}}), \mu_c)$. In this feature set, l_c refers to the customer's location, $\tau_{c,\bar{v}}$ is the travel time between the customer and the active vehicle, $\tau_{c,d}$ is the travel time between the customer and the depot, \tilde{d}_c is the customer's unserved demand (or expected demand, if not realized yet), $\min(\tilde{d}_c, q_{\bar{v}})$ shows the portion of customer's demand that can be served by the active vehicle, and μ_c indicates whether the actual demand is realized or not. The vector of target customers is denoted F_k :

$$F_k = [(l_c, \tau_{c,\bar{v}}, \tau_{c,d}, \tilde{d}_c, \min(\tilde{d}_c, q_{\bar{v}}), \mu_c)]_{c \in \mathcal{C}_{\bar{v}}}$$

Less relevant data may now be aggregated. In particular, we seek an aggregated representation of the remaining customers to provide the active vehicle with an overview of the service area. The challenge here is that the stochastic customer set necessitates an aggregating technique that handles a variable number of inputs while delivering a fixed-size output. To this purpose, we propose a heatmap-style approach to encode all customers in the service area as a fixed-size vector H_k . In particular, we divide the service area into a set of square-shaped *partitions*. We denote P the set of partitions and $\mathcal{C}_p \subset \mathcal{C}_0$ the subset of customers that are located in partition $p \in P$. We characterize each partition by two values (ξ_p^c, ξ_p^d) representing the number of customers and the total expected demands in that partition, respectively.

$$\xi_p^c = |\mathcal{C}_p|, \quad \xi_p^d = \sum_{c \in \mathcal{C}_p} \tilde{d}_c$$

We define $H_k = [(\xi_p^c, \xi_p^d)]_{p \in P}$ the fixed-size vector representing the aggregated information of all customers in the service area. In addition to H_k and F_k , the state of the vehicles G_k forms the third segment of the observation representation. We define $G_k = [(l_v, a_v, q_v)]_{v \in \mathcal{V}}$, which has the same representation used in the centralized state.

The main difference between s_k and $o_{k,\bar{v}}$ is the customers' state, with a variable size, replaced by a fixed-size vector $[F_k, H_k]$. The observation representation has two advantages: first, it overcomes the problem of a variable-sized state representation that was previously addressed. Second, depending on the hyper-parameter \tilde{n} and the number of partitions $|P|$, the dimension of the observation can be much lower than the state representation. In particular, the customers' state in s_k has the dimension of $4 * |\mathcal{C}_0|$, while the dimension for $[F_k, H_k]$ is $(6 * \tilde{n} + 2 * |P|)$. As a numerical example, if we assume the average number of customers to be 50, the dimension of the customers' state will be on average 200. However, this value for $o_{k,\bar{v}}$ will be 110 ($60 + 50$), where \tilde{n} and $|P|$ are assumed to be 10 and 25, respectively.

In addition to making the observation representation considerably smaller, the idea of identifying and focusing on target customers enables us to reduce the dimension of the active vehicle's action space

by limiting its actions to target customers. Therefore, in the definition of $J(s_k, \bar{v})$ (Equation (8)), which is used to construct the active vehicle's action space $A(s_k, \bar{v})$, we substitute \mathcal{C}_0 with $\mathcal{C}_{\bar{v}}$. As a result, the size of the action space is reduced to $|\mathcal{C}_{\bar{v}}| + 1$, where $|\mathcal{C}_{\bar{v}}| \leq \tilde{n}$.

4.3 Cooperation issues

In a decentralized setting, non-cooperating vehicles will generally provide suboptimal policies for the VRPSCD. This is illustrated in Figure 3. The example considers 4 customers ($\mathcal{C}_0 = \{1, 2, 3, 4\}$) and 2 vehicles ($Q = 100, L = 30$). Figure 3(a) shows the position of customers and the depot, the deterministic demands ($\{10, 10, 1, 1\}$), and the travel times. Notice that the duration limit is such that vehicles cannot serve customers to the left and to the right of the depot within the time limit. Non-cooperating vehicles will try to increase their total served demand regardless of (or even competing with) other vehicles. In the example of the figure, one vehicle will travel to customer 1, and the other vehicle to customer 2, with 20 total served demand (Figure 3(b)). However, by considering collaboration among vehicles (Figure 3(c)), one vehicle would serve the customers to the left, while the other would serve the customers to the right, with total served demand 22. In this example, the cooperating strategy was globally advantageous; however, one of the vehicles had to "sacrifice" its own reward to let the other vehicle collect a higher reward.




Figure 3: An example of collaboration in VRP

Similar cooperation issues have been considered, for example, in Cooperative Multi-Agent System (CMAS) (Stone and Veloso 2000). In CMASs, a set of agents (corresponding to vehicles in our case) share the environment and act independently to reach a given goal (Panait and Luke 2005). The approach aims to find a policy (or policies) for agents in order to maximize a given global objective function. Although an agent does not decide for other agents, its action may significantly affect their decision strategy. More specifically, the value of each action for an agent depends on what the other agents choose. Consequently, similarly to what we discussed for the decentralized VRPSCD, the collaboration of agents is crucial to achieving an optimal decision policy. In CMAS, the collaboration is addressed by a central coordination mechanism. The role of this mechanism is typically to take into account the action of *all* agents in evaluating their value. For example, the Q-learning algorithm considers the central coordination by defining the Q values as the value of taking a *joint* action x_k in state s_k (not the action of the individual agent, x_k^v). Similarly, in approximate policy iteration methods, such as the policy gradient, the value function employed to iteratively evaluate policies in the training phase plays the role of the central coordination mechanism. This value function, which is usually defined as in Equation (20), estimates the cumulative rewards all agents may collect from the current state.

In contrast to CMASs, where a central coordination mechanism plays a crucial role in enabling collaboration between agents, our decentralized approach proposed in this study eliminates the necessity of a central coordinator. In fact, the decentralized approach exploits the special characteristic of the VRPSCD that forces vehicles to decide sequentially (one per decision epoch). Consequently, the reward function and the transition function will be constructed based on its action ($x_k^{\bar{v}}$), independently

from the other vehicles. This results in a value function only dependent on the action of the active vehicle ($x_k^{\bar{v}}$), not on the joint action (x_k).

5 Solution method

The VRPSCD, in both centralized and decentralized formulations, has a considerably large state and action space. The consequent so-called curse of dimensionality make it hard to solve the VRPSCD by conventional stochastic dynamic programming. Specialized approximate methods have been proposed in the literature to overcome the difficulties arising in these cases (Powell 2011). In this paper, we consider the Q-learning algorithm. Our preliminary experiments suggest that Q-learning is not able to solve the centralized formulation of the VRPSCD due to its huge state-action space. For this reason, in this section, we only report on our implementation of the decentralized formulation. In particular, in Section 5.1 we give a general overview of the Q-learning algorithm. In Section 5.2 we provide the details of our implementation of the decentralized VRPSCD, which features a neural network to approximate the Q-values, as well as several techniques to accelerate the convergence of the algorithm, such as Replay Memory and Double Q Networks.

5.1 An overview of Q-learning algorithm

Q-learning is a model-free algorithm based on approximate value iteration in which an agent learns how to behave in order to maximize the collected reward through a sequence of interactions with the environment, resulting in a policy that satisfies Equation (27). More specifically, this method iteratively updates the value of making the action x_k in state s_k , known as the Q factor of that state-action pair ($Q(s_k, x_k)$). In fact, it is possible to rewrite the Bellman equation (defined in Equation (26)) in terms of Q factors as follows:

$$V(s_k) = \max_{x_k \in A(s_k)} Q(s_k, x_k), \quad (29)$$

where:

$$Q(s_k, x_k) = R(s_k, x_k, w) + \gamma \max_{x_{k+1} \in A(s_{k+1})} Q(s_{k+1}, x_{k+1}). \quad (30)$$

The reader is referred to (Powell 2011) for a complete overview of the algorithm. Algorithm 1 illustrates a generic Q-learning framework. To summarize, this algorithm simulates the problem for several trials ($Trials_MAX$) over different sample scenarios of stochastic components of the problem. The simulation follows the so-called ϵ -greedy policy, which enables random exploration of the action space. The level of exploration is typically larger in the initial phases of the algorithm when actions are mostly chosen randomly. As the algorithm proceeds, actions with higher estimated values are generally preferred. Q-learning exploits the interactions between the agent and the environment to update the Q values (the so-called *experiences*). An experience is defined as a tuple of (s_k, x_k, r_k, s_{k+1}) , describing interaction of the agent with the environment consisting of the agent observation of the state s_k , its action x_k , the received reward r_k (in VRPSCD, $r_k = R(s_k, x_k, w)$) and the observation of the next state s_{k+1} . For each state-action pair experienced, we estimate their Q value as $\hat{Q}(s_k, x_k)$ using the following equation:

$$\hat{Q}(s_k, x_k) = r_k + \gamma \max_{x_{k+1} \in A(s_{k+1})} Q(s_{k+1}, x_{k+1}), \quad (31)$$

where γ is the discount factor. Then, we update the $Q(s_k, x_k)$ according to the following equation:

$$Q(s_k, x_k) \leftarrow Q(s_k, x_k) + \alpha [\hat{Q}(s_k, x_k) - Q(s_k, x_k)], \quad (32)$$

where α is the step size (learning rate).

Algorithm 1 returns a set of Q values that can be used as the decision policy. In particular, the optimal decision policy will be:

$$\pi^*(s_k) = \arg \max_{x_k \in A(s_k)} Q(s_k, x_k). \quad (33)$$

Algorithm 1: Generic Q-learning Algorithm

```

Initialize: set Q values to 0
while trials < Trials_MAX do
  Choose a sample scenario w
  Observe state s0
  Set k = 0, simulate = True
  while simulate do
    Take action:  $x_k = \begin{cases} \text{random action in } A(s_k) & \epsilon \\ \arg \max_{x' \in A(s_k)} Q(s_k, x') & 1 - \epsilon \end{cases}$ 
    Transit to  $s_{k+1} = S^M(s_k, x_k, w_{k+1})$  and observe the reward  $r_k$ 
    Estimate and update new Q values:
    •  $\hat{Q}(s_k, x_k) = \begin{cases} r_k & A(s_{k+1}) = \emptyset \\ r_k + \gamma \max_{x' \in A(s_{k+1})} Q(s_{k+1}, x') & \text{otherwise} \end{cases}$ 
    • Set  $Q(s_k, x_k) = Q(s_k, x_k) + \alpha[\hat{Q}(s_k, x_k) - Q(s_k, x_k)]$ 
    Set k = k + 1
    if  $A(s_k) = \emptyset$  then
      | Set simulate = False
    end
  end
  Set trials = trials + 1
  Decay the exploration rate  $\epsilon$ 
end
return Q values

```

5.2 The DecQN algorithm

As previously mentioned, our preliminary implementation of the Q-learning algorithm could not solve the centralized formulation of the VRPSCD. For the Q-learning to be successful, in fact, sufficient exploration of state-action pairs is required, which is impractical for very large state-action spaces, such as the one of the centralized formulation. The decentralized formulation proposed in Section 4, on the other hand, has a considerably smaller state-action space, and it is more suitable to be addressed by Q-learning. For this purpose, by using the proposed observation function $O(s_k, \bar{v})$, Equations (31) and (32) can be rewritten as follows:

$$\hat{Q}(o_{k,\bar{v}}, x_k) = R(s_k, x_k, w) + \gamma \max_{x_{k+1} \in A(s_{k+1}, \bar{v}')} Q(O(s_{k+1}, \bar{v}'), x_{k+1}), \quad (34)$$

$$Q(o_{k,\bar{v}}, x_k) \leftarrow Q(o_{k,\bar{v}}, x_k) + \alpha * [\hat{Q}(o_{k,\bar{v}}, x_k) - Q(o_{k,\bar{v}}, x_k)], \quad (35)$$

where \bar{v}' is the active vehicle at decision epoch $k + 1$.

In this study, we implement the Q values in the form of an artificial neural network instead of the traditional look-up tables. This method is called Q-network in the literature. The Q-network has the purpose of returning the approximate Q value for a given observation-action pair $(o_{k,\bar{v}}, x_k)$. Generally speaking, for computational efficiency, it is desirable to design the Q-network in such a way that only one forward call is sufficient to return the Q values for all $(|\mathcal{C}_{\bar{v}}| + 1)$ actions in $A(s_k, \bar{v})$. To this purpose, we consider a neural network that returns $\tilde{n} + 1$ Q values in every call. Notice that, under some circumstances, the number of target customers can be less than \tilde{n} , making the cardinality of $A(s_k, \bar{v})$ variable, but restricted to $\tilde{n} + 1$. Therefore, in the vector of $\tilde{n} + 1$ Q values approximated by the Q-network, we associate the first $|\mathcal{C}_{\bar{v}}|$ values to the Q value of target customers $\mathcal{C}_{\bar{v}}$ and assign the last one to the depot. Q values not associated with target customers nor the depot will be set to 0.

We implement a neural network with two hidden layers. This network gets the observation representation $(o_{k,\bar{v}})$ as the input layer and returns approximate Q values in the output layer. We denote h^{in} and h^{out} , respectively, the size of the input and output layers, where $h^{in} = 6\tilde{n} + 2|P| + 3m + 1$ and

$h^{out} = \tilde{n} + 1$. The hidden layers are assumed to be fully connected dense layers linked by an activation function (a rectified linear unit, ReLU). The size of the hidden layers is chosen based on empirical tests. In particular, we set their size to $\frac{2}{3}(h^{in} - h^{out}) + h^{out}$ and $\frac{1}{3}(h^{in} - h^{out}) + h^{out}$, respectively. Denoting by θ the trainable weights of the proposed neural network, the Q value of an observation-action pair is represented as $Q(o_{k,\bar{v}}, x_k, \theta)$. It is worth noting that any equation involving $Q(o_{k,\bar{v}}, x_k)$ will also apply to $Q(o_{k,\bar{v}}, x_k, \theta)$.

To train the proposed network, we aim at minimizing the so-called loss function. We define this function based on the difference between the new estimation of the values $\hat{Q}(o_{k,\bar{v}}, x_k, \theta)$ and the current estimation $Q(o_{k,\bar{v}}, x_k, \theta)$, denoted as Δ (i.e., $\Delta = \hat{Q}(o_{k,\bar{v}}, x_k, \theta) - Q(o_{k,\bar{v}}, x_k, \theta)$). Among several loss functions used in the literature, we found that the Huber loss function performs better than the commonly used functions such as Mean Squared Error (MSE) and Mean Absolute Error (MAE). The Huber loss acts as an MSE for small Δ , while it works like an MAE for larger differences. The Huber loss is defined as:

$$Huber(\Delta) = \begin{cases} \frac{1}{2}\Delta^2 & \text{for } |\Delta| \leq \delta \\ \delta(|\Delta| - \frac{1}{2}\delta) & \text{otherwise} \end{cases} \quad (36)$$

where δ is a control parameter. We obtained the best results by setting $\delta = 5$ and used the Adam optimizer (Kingma and Ba 2014) to update the network in order to minimize the loss function.

Algorithm 2 summarizes our DecQN algorithm. Similar to what was suggested by Mnih et al. (2015), our preliminary results showed that using the target network (Van Hasselt, Guez, and Silver 2016) and the replay memory help the model achieve better convergence. Having a target network (parameterized by $\bar{\theta}$ in Algorithm 2) allows the model to learn from a stable target. The target network $\bar{\theta}$ is periodically copied from the primary network θ every β_d trials. In the replay memory, experiences are stored in a First-In-First-Out list, denoted as B , with a fixed size of N . At every decision epoch, with a probability of β_t , a random batch of experiences (\tilde{B}) is sampled from this list to update the weights (θ) in the neural network. Several advantages have been mentioned for using the replay memory (Mnih et al. 2015) such as breaking the correlation between consecutive experiences and giving more chances to those experiences that happen more. The exploration rate (ϵ) and the learning rate (α) are linearly decayed during the training phase. The decaying rate is decided based on the length of the training process.

As previously noticed, one of the main features of the proposed approach is that we only develop a single set of Q values, resulting in a policy shared among all vehicles. In particular, the experiences of all vehicles obtained via simulation are utilized for training a single set of Q values. Building a single policy offers several benefits: first, it allows vehicles to exchange their experiences (i.e., interactions with the environment) under a single policy. Also, it makes the development of the policy more sample efficient. In particular, the experiences of all m vehicles are gathered to construct one policy, which minimizes the total required simulated samples, rather than utilizing the experiences of each vehicle to establish its policy. Finally, maintaining and storing one policy is more manageable than maintaining m policies.

6 Experimental results

In this section, we present our computational experience. First, since no instance set is available in the literature, we construct a new set of instances for the VRPSCD. We then design four main experiments. The first one is aimed at evaluating the overall performance of the DecQN algorithm. To this purpose, we first implement two additional benchmark policies, i.e., the random policy and a heuristic policy. We then perform an extensive computational experiment to compare the three algorithms. The purpose of the second experiment is to compare the DecQN with the existing literature. The challenge is that no direct comparison is possible, given that, to the best of our knowledge, the VRPSCD has not been addressed before. The closest problem in the literature we are aware of is the VRPSD, for which

Algorithm 2: The DecQN

```

Initialize the Q-network and the target Q-network with random weights  $\theta, \bar{\theta}$ 
Initialize the replay memory  $B$  with the size of  $N$ 
Initialize  $Trials\_MAX$ ,  $\beta_t$ ,  $\beta_d$  as the training length, the probability for updating the network, the number of
  trials between two substitution of the target network, respectively
while  $trials < Trials\_MAX$  do
  Choose a sample scenario  $w$  and observe a new set of customers  $\mathcal{C}_0$ 
  Set  $t_0 = 0$ ,  $k = 0$ ,  $\bar{v}_{prev} = \emptyset$ ,  $Buff = [\emptyset]_{v \in \mathcal{V}}$ , and  $simulate = True$ 
  while  $simulate$  do
    Set  $\bar{v}$  = a random  $v \in \bar{\mathcal{V}}$ , where  $\bar{\mathcal{V}} = \{v \in \mathcal{V} | a_v = t_k\}$ 
    if  $k > 0$  then
      Transit from  $s_{k-1}^x$  to  $s_k$  and observe the reward  $r$ 
       $Buff[\bar{v}_{prev}] = (s_{k-1}, \bar{v}_{prev}, x_{k-1}, s_k, \bar{v})$ 
      if  $Buff[\bar{v}] \neq \emptyset$  then
         $(s, v, x, s', v') \leftarrow Buff[\bar{v}]$ 
        Add  $(s, v, x, r, s', v')$  to  $B$ 
      end
    end
    Observe  $o_{k,\bar{v}} = O(s_k, \bar{v})$ 
    Take action:  $x_k = \begin{cases} \text{random action in } A(s_k, \bar{v}) & rand[0, 1] < \epsilon \\ \arg \max_{x'} Q(o_{k,\bar{v}}, x', \theta) & \text{otherwise} \end{cases}$ 
    Transit from  $s_k$  to  $s_k^x$ 
    if  $rand[0, 1] < \beta_t$  then
      Estimate and update new Q values:
      

- Sample a batch of experiences  $\tilde{B}$  from  $B$
- Compute the estimation  $\hat{Q}(O(s, v), x, \theta) = \begin{cases} r & A(s_k, v) = \{depot\} \wedge l_v = depot \\ r + \gamma \max_{x'} Q(O(s', v'), x', \bar{\theta}) & \text{otherwise} \end{cases}$ ,  

 $\forall (s, v, x, r, s', v') \in \tilde{B}$
- Compute the loss function:  

 $\Phi = \mathbb{E}_{\tilde{B}} [Huber(Q(O(s, v), x, \theta) - \hat{Q}(O(s, v), x, \theta))]$
- Update the network weights  $\theta$ :  $\theta = \theta - \alpha \nabla_{\theta} \Phi$

end
    Set  $k = k + 1$ , and  $\bar{v}_{prev} = \bar{v}$ 
    if  $A(s_k, \bar{v}) = \emptyset \wedge l_v = depot, \forall v \in \mathcal{V}$  then
      Set  $simulate = False$ 
    end
  end
  Set  $trials = trials + 1$ 
  Every  $\beta_d$  trials, set  $\bar{\theta} = \theta$ 
  Decay the learning rate  $\alpha$  and the exploration rate  $\epsilon$ 
end
return Q values

```

Goodson, Thomas, and Ohlmann (2016) developed the current state-of-the-art method. As mentioned earlier, the VRPSD can be seen as a particular case of the VRPSCD where both customer locations and expected demands are known in advance. Even though our algorithm does not take advantage of the specific characteristics of the problem, it turns out, quite surprisingly, that our algorithm is competitive and sometimes outperforms the results in Goodson, Thomas, and Ohlmann (2016). The third experiment can be seen as an extension of the second one, where we investigate the possibility of modifying the training phase of the DecQN to develop a policy able to address a larger range of instances aggregated over specific parameters, such as the duration limit and stochastic variability. Finally, the last experiment provides a simple demonstration of the fact that the policies provided by our algorithm can be easily embedded as base policies in rollout algorithms.

Regarding the hyper-parameters in Algorithm 2, we consider a minibatch of 32 experiences ($|\tilde{B}| = 32$) uniformly sampled from the replay memory with $N = 50000$, where $\beta_t = 0.05$. We also replace the

target network every 1000 trials ($\beta_d = 1000$). The hyper-parameters related to the maximum number of trials ($Trials_MAX$) and the decaying factor for the learning rate (α) and the exploration rate (ϵ) will be specified for each experiment later. All the procedures have been implemented in python and executed on 3.6 GHz Intel Xeon processors with 64 GB RAM (no GPUs are used).

The experiments mentioned above are described in Sections 6.1 to 6.4 in detail. Also, a discussion on computational time is provided in Section 6.5.

6.1 Performance of DecQN in VRPSCD

In this section, we compare the performance of the proposed DecQN with the random policy and a heuristic policy. Given that the VRPSCD is new in the literature, we generate a new set of instances in Section 6.1.1. We then describe the details of the benchmarks in Section 6.1.2. We report about the computational experiment and results in Section 6.1.3.

6.1.1 VRPSCD instances

As custom when building computational experiments in learning-based methods, we need to provide two types of input data, one to be used in the training phase of the algorithm, the other in the testing phase. Elements of these two sets are samples of the problem at hand. Although training and testing sets are distinct, they are typically generated by sampling from the same distributions.

In the context of the VRPSCD, we define an instance i as the tuple $i = (\mathcal{A}, \Psi_{n_z}, \Psi_l, \Psi_{\bar{d}}, \Psi_{\hat{d}}, m, Q, L)$, where \mathcal{A} denotes the geography of the service area, which is assumed to be partitioned into a set of zones Z , and $\Psi_{n_z}, \Psi_l, \Psi_{\bar{d}}, \Psi_{\hat{d}}$ denote the probability distribution functions of the number of customers n_z per zone $z \in Z$ (customer density), customer locations, expected demands, and actual demands, respectively. Furthermore, m, Q, L indicate the number of vehicles, their capacity, and the duration limit, respectively. By varying the service area \mathcal{A} , the distributions $\Psi_{n_z}, \Psi_l, \Psi_{\bar{d}}, \Psi_{\hat{d}}$, and the values of parameters m, Q , and L , one can obtain different instances. As it will be clear in the detailed description below, this experiment considers three different distributions for the customer density Ψ_{n_z} , and three different values of the vehicle capacity Q , resulting in a set I of 9 instances. Furthermore, the distribution Ψ_{n_z} of the customer density uniquely determines the values of the number of vehicles m and the time limit L . Given an instance i , a realization \hat{i} (also referred to as sample, or scenario) can be obtained by sampling from distributions $\Psi_{n_z}, \Psi_l, \Psi_{\bar{d}}, \Psi_{\hat{d}}$. The set of all realizations \hat{i} of i is denoted $\hat{I}(i)$. It follows a detailed description of the method to generate problem instances and realizations.

We consider a squared-shape service area \mathcal{A} of dimension 100×100 with a depot located at the center ($depot = (50, 50)$). Customers can potentially be located at any position inside the service area. However, in real-life applications, certain zones of the service area never generate requests, such as non-residential areas. Therefore, for a given service area partitioned into a set of 5×5 square-shaped zones Z , we randomly pick a subset \mathcal{Z} of *active* zones, in which customers must be located. In our experiment, we consider 15 active zones out of 25 (i.e., 60% of the zones are active). Figure 4 shows the layout of the considered region.

To generate the customer locations and demands, we make the assumption that active zones are homogeneous in the sense that the customer density Ψ_{n_z} and distributions $\Psi_l, \Psi_{\bar{d}}, \Psi_{\hat{d}}$ do not depend on the specific zone. The customer density Ψ_{n_z} specifies the probability distribution of the number n_z of requests for each active zone. We consider three possible levels of customer density $\mathcal{D} = \{\text{Low}, \text{Moderate}, \text{High}\}$. In particular, n_z takes values $\{0, 1, 2, 3\}$ for Low, $\{2, 3, 4, 5\}$ for Moderate, and $\{4, 5, 6, 7\}$ for High customer densities, with probabilities $\{0.1, 0.4, 0.4, 0.1\}$, respectively. The corresponding expected number of customers at each zone \bar{n}_z is 1.5, 3.5, and 5.5, respectively. The total expected number of customers \bar{n} for a given Ψ_{n_z} is calculated as $\bar{n} = \mathbb{E}[\sum_{z \in \mathcal{Z}} n_z] = |\mathcal{Z}| * \bar{n}_z$, resulting in 23, 53, and 83 customers for Low, Moderate, and High customer density, respectively. Once n_z is sampled according to the distribution Ψ_{n_z} for each zone $z \in \mathcal{Z}$, the customer locations are




Figure 4: The layout of active zones

generated following the distribution Ψ_l . In this experiment, we choose Ψ_l to be the uniform distribution function over (x,y)-coordinates inside a given zone. The above process allows us to sample the initial set of customers \mathcal{C}_0 .

Given a customer $c \in \mathcal{C}_0$, the expected demand \bar{d}_c and actual demand \hat{d}_c are sampled according to distributions $\Psi_{\bar{d}}$ and $\Psi_{\hat{d}}$, respectively. These distributions are chosen following the method proposed in Gendreau, Laporte, and Séguin (1995). In particular, $\Psi_{\bar{d}}$ is the discrete uniform distribution over the values $\{5, 10, 15\}$. Once the value \bar{d}_c has been sampled, the actual demands follow the distribution $\Psi_{\hat{d}}$ which is the discrete uniform distribution on values $\{\bar{d}_c - 5, \dots, \bar{d}_c, \dots, \bar{d}_c + 5\}$. In order to avoid the no-demands case when $\bar{d}_c = 5$, we use the discrete uniform distribution function on values $\{\bar{d}_c - 4, \dots, \bar{d}_c, \dots, \bar{d}_c + 4\}$.

We determine the number of vehicles m and the time limit L uniquely as functions of the distribution Ψ_{n_z} by linking them through the concept of the filling rate. The filling rate f indicates the portion of all demands that can be served by all vehicles without replenishment:

$$f = \frac{\mathbb{E}[\sum_{c \in \mathcal{C}_0} \bar{d}_c]}{m * Q}.$$

In the equation above, $\mathbb{E}[\sum_{c \in \mathcal{C}_0} \bar{d}_c]$ is the expected total demand and can be approximated by $\bar{n}_z * |\mathcal{Z}| * \bar{\bar{d}}$, where $\bar{\bar{d}}$ is the average expected demand of a customer. Since the distribution function $\Psi_{\bar{d}}$ assigns expected demands 5, 10, or 15 to customers uniformly (with a probability of $\frac{1}{3}$), the average expected demand of a customer ($\bar{\bar{d}}$) will be equal to 10 ($\frac{5+10+15}{3}$) in all instances. By fixing the filling rate and the vehicle capacity to 1 and 75, respectively, the number of vehicles for each level of customer density Ψ_{n_z} can then be obtained by the definition of the filling rate:

$$m = \left\lceil \frac{\mathbb{E}[\sum_{c \in \mathcal{C}_0} \bar{d}_c]}{f * Q} \right\rceil = \left\lceil \frac{\bar{n}_z * |\mathcal{Z}| * \bar{\bar{d}}}{f * Q} \right\rceil = \left\lceil \frac{\bar{n}_z * 15 * 10}{1 * 75} \right\rceil = \lceil 2 * \bar{n}_z \rceil. \quad (37)$$

In this way, the number of required vehicles for each class of customer density is 3, 7, and 11, respectively. Notice that, according to the expression $\mathbb{E}[\sum_{c \in \mathcal{C}_0} \bar{d}_c]$, the expected total demand at each instance will be 230, 530, and 830 for Low, Moderate, and High customer density, respectively.

To determine the duration limit L for a given instance with the customer density distribution Ψ_{n_z} and its associated m , we utilize the average travel time of vehicles acquired by solving a set of 250,000 realizations of that instance via the Greedy policy (described in Section 6.1.2). In this process, we assumed that $L = \infty$, $Q = 75$, and realizations are randomly sampled from distributions $\Psi_{n_z}, \Psi_l, \Psi_{\bar{d}}, \Psi_{\hat{d}}$. To make the duration limit constraining, we let the duration limit be the average travel time of all vehicles multiplied by 0.75. The resulting duration limits for Low, Moderate, and High customer densities are 221.47, 195.54, and 187.29, respectively. To complete the description of the instance generation, we need to determine the values Q . According to Equation (37), the number of vehicles m is determined such that when $Q = 75$, vehicles can serve the total expected demand

without replenishing at the depot. Therefore, for a given duration limit, while using smaller capacities (< 75) may increase the frequency of restocking operations, which results in serving fewer customers in a limited time, considering larger vehicle capacities (> 75) may not necessarily affect the performance of vehicles. Therefore, in this experiment, we consider three possible values $\mathcal{Q} = \{25, 50, 75\}$.

To summarize, the instance set I is obtained by varying the customer density distribution function $\Psi_{n_z} \in \mathcal{D} = \{\text{Low, Moderate, High}\}$ and the vehicle capacity $Q \in \mathcal{Q} = \{25, 50, 75\}$, resulting in 9 different instances. Fixing the customer density distribution also uniquely determines the number of vehicles m and the time limit L . The rest of the instance-defining elements, namely the service area \mathcal{A} and the distributions $\Psi_l, \Psi_{\bar{d}}, \Psi_{\hat{d}}$ are kept fixed throughout the experiment. In other words an instance in this experiment is completely identified by the pair (Ψ_{n_z}, Q) . Finally, realizations \hat{i} of an instance i are obtained by sampling from distributions $\Psi_{n_z}, \Psi_l, \Psi_{\bar{d}}, \Psi_{\hat{d}}$.

6.1.2 Benchmarks

We compare DecQN with two policies that are capable of handling stochastic customers with stochastic demands in real-time. These policies cannot handle the preemptive restocking, however. Therefore, the action set is redefined as:

$$A(s_k, \bar{v}) = \begin{cases} \{\text{depot}\} & q_{\bar{v}} = 0 \parallel J(s_k, \bar{v}) = \emptyset \\ J(s_k, \bar{v}) & \text{otherwise.} \end{cases} \quad (38)$$

The first benchmark is the *Random Policy*. In this policy, an action is randomly picked from the action set $A(s_k, \bar{v})$. This policy is important because it is exactly the policy that our model starts with when the exploration rate in the ϵ -greedy method is 1. Therefore, the improvement of our model over the Random policy demonstrates the effect of the learning. The performance of this policy will be reported as RP.

The second benchmark is the *Greedy Policy*. Given that the objective function in the VRPSCD is to maximize the expected total served demand, this policy selects customers according to the potential reward in terms of served demands. In particular, at each decision epoch the Greedy policy chooses the customer $c \in A(s_k, \bar{v})$ with the highest \tilde{d}_c . We recall that \tilde{d}_c was defined in Equation (28) and gives the expected or the remaining demand of customer c . If several customers have the same \tilde{d}_c , this policy selects the closer one. The results of the Greedy policy are reported as GP.

6.1.3 Results and discussion

In this experiment, for each instance $i \in I$, we generated the training set by sampling 5 million realizations \hat{i} from $\hat{I}(i)$. We then trained the DecQN on these realizations. The exploration rate decayed from 1.0 to 0.1 in the first 300K trials. The learning rate also decreased linearly from 10^{-3} to 10^{-4} in the first 2 million trials. In order to test the trained policies, we generate a new set of realizations, called the test set. For each instance $i \in I$, the test set is also sampled from $\hat{I}(i)$. More specifically, we generate 500 customer realizations and 500 demand realizations for each instance, resulting in 250,000 scenarios. It is important to notice that although the training and the test sets are sampled from the same distributions, they are disjoint.

We then apply DecQN, Random, and Greedy policy to the test set and compare the performances. The aggregated results are reported in Table 1. In this table, the first four columns indicate the characteristics of the instance. As an upper bound, the expected total demand at each instance is reported in column $\mathbb{E} \sum d_c$. The results of DecQN, along with the results of the two benchmarks, are reported in columns DecQN, RP, and GP, respectively. The column $\% \mathbb{E} \sum d_c$ shows the ratio of demands served by DecQN in each instance. Note that the reported results for each instance are averaged over 250,000 tests.

Table 1: Results of DecQN compared with Random and Greedy policies

Ψ_{n_z}	Q	m	L	$\mathbb{E} \sum d_c$	DecQN	% $\mathbb{E} \sum d_c$	RP	GP	%Imp. RP	%Imp. GP
L	25	3	221.47	230	159.12	69.18%	99.97	143.00	59.18%	11.28%
	50				193.11	83.96%	120.56	171.37	60.18%	12.69%
	75				210.10	91.35%	128.90	192.36	63.00%	9.23%
Avg:								60.79%	11.06%	
M	25	7	195.54	530	360.93	68.10%	217.95	321.54	65.60%	12.25%
	50				452.67	85.41%	264.50	417.96	71.14%	8.30%
	75				501.24	94.57%	269.57	460.03	85.94%	8.96%
Avg:								74.23%	9.84%	
H	25	11	187.29	830	552.54	66.57%	334.03	502.53	65.41%	9.95%
	50				703.04	84.70%	402.77	664.81	74.55%	5.75%
	75				778.66	93.81%	412.20	738.23	88.91%	5.48%
Avg:								76.29%	7.06%	
Total Avg:								70.43%	9.76%	

We define the measure %Imp. X to evaluate the improvement percentage of our method (DecQN) over another method X , computed as:

$$\% \text{Imp. } X = \frac{\text{DecQN} - X}{X} \times 100$$

In Table 1, columns %Imp. RP and %Imp. GP refer to the percentage improvement of our DecQN over the Random and Greedy policies. Notice that in this table, customer densities L, M, and H are acronyms for Low, Moderate, and High customer densities, respectively.

Averaged results on all instances show that DecQN outperforms the Random and Greedy policies by 70.43% and 9.76%, respectively. The significant improvement over the Random policy demonstrates the effect of the learning. According to Table 1, %Imp. GP falls from 11.06% to 7.06% as the customer density increases. To explain this behavior, we argue that as the customer density increases, the instance becomes relatively simpler for the Greedy policy. To demonstrate this, we need to analyze the sensitivity of GP to variations in customer density. However, since the value of the served demands by the Greedy policy depends on the specification of the solved instance, we must normalize it by removing its reliance on customer density. To this purpose, we divide the GP by $\mathbb{E} \sum d_c$. Notice that the value of $\mathbb{E} \sum d_c$ solely depends on the level of Ψ_{n_z} ; hence, $\% \frac{\text{GP}}{\sum d_c}$ gives us a normalized measure that can be used to assess the performance of the Greedy policy with respect to the variations in customer density. This value can also be viewed as the ratio of demands served by the Greedy policy. Table 2 (on the left) reports computed values of $\% \frac{\text{GP}}{\sum d_c}$ for different instances. Accordingly, as the customer density rises, the Greedy policy can serve a more significant ratio of demands. More specifically, the average ratio of demands served by the Greedy policy increases from 73.44% to 76.53%, when the customer density increases from Low to High. It can also be intuitively inferred that as customers in the service area become dense, the decision problem for each vehicle will be more of a customer selection problem than a routing problem where a long-term collective reward plays a decisive role. As a result, for policies such as the Greedy, where the routing factor is ignored, a more densely populated service area would be preferred than one with fewer customers. On the other hand, Table 2 (on the right), which reports the ratio of demands served by the DecQN, demonstrates that the performance of the DecQN remains approximately constant as the customer density varies. Following what was discussed earlier, we may attribute this result to the fact that, unlike the Greedy policy, the DecQN accounts for routing.

To analyze how the vehicle capacity impacts results, we study the sensitivity of %Imp. GP to the changes in the vehicle capacity. It can be observed in Table 1 that the %Imp. GP is higher when Q is smaller. For example, in instances with High customer density, the %Imp. GP declines from 9.95%

Table 2: The ratio of demands served by DecQN and Greedy policy

	$\% \frac{GP}{\mathbb{E} \sum d_c}$	Q					$\% \frac{DecQN}{\mathbb{E} \sum d_c}$	Q			
		25	50	75	Avg			25	50	75	Avg
Ψ_{n_z}	L	62.17%	74.51%	83.63%	73.44%	Ψ_{n_z}	L	69.18%	83.96%	91.35%	81.50%
	M	60.67%	78.86%	86.80%	75.44%		M	68.10%	85.41%	94.57%	82.69%
	H	60.55%	80.10%	88.94%	76.53%		H	66.57%	84.70%	93.81%	81.70%
Avg:		61.13%	77.82%	86.46%	75.14%	Avg:		67.95%	84.69%	93.25%	81.96%

to 5.48% when the vehicle capacity increases from 25 to 75. A possible reason explaining this result is that, according to Table 2, as Q rises, the percentage of demands served by both policies approach to serve almost all demands (i.e., from 67.95% and 61.13% to 93.25% and 86.46%, respectively). Thus there is a lower margin of improvement in higher vehicle capacities, which reduces the potential difference between their performances.

Another significant finding from Table 2 (on the right) is that for a given Q , the ratio of demands served by DecQN ($\% \frac{DecQN}{\mathbb{E} \sum d_c}$) remains almost constant as the customer density varies. In particular, it can be seen that the value of $\% \frac{DecQN}{\mathbb{E} \sum d_c}$, averaged across different values of Q , remains around $\approx 82\%$ for each level of customer density. This finding is noteworthy because it implies that our method scales effectively and does not suffer from the curses of dimensionality that arise in ADP-based methods when the number of customers or/and vehicles grows.

6.2 Performance of DecQN in VRPSD

The VRPSD can be seen as a particular version of the VRPSCD, where the set of customers, including their locations and expected demands, is known in advance. Even though the DecQN, being designed to handle stochastic customer sets, does not exploit the specific characteristics of having previous knowledge about fixed customers, we show in this experiment that DecQN can compete with the solution method proposed by Goodson, Thomas, and Ohlmann (2016), which is the best-performing benchmark specialized for VRPSDs with duration limits and multiple vehicles.

In Section 6.2.1, we define the instance set and describe how we generate the instance realizations. For each instance, we train a model for 3 million trials. Similarly to the first experiment, we decay the exploration rate and the learning rate from 1.0 and 10^{-3} to 0.1 and 10^{-4} , respectively, in the first 300K and 1.5 million trials. We describe the benchmark in Section 6.2.2 and discuss the results in Section 6.2.3.

6.2.1 VRPSD instances

The VRPSD, as a particular version of the VRPSCD, assumes that the number of customers, their locations and their expected demands are given in advance. Therefore, the distribution functions $\Psi_{n_z}, \Psi_l, \Psi_{\bar{d}}$ in Section 6.1.1 do not play a role in this experiment. The remaining instance-defining components $(\mathcal{A}, \Psi_{\bar{d}}, m, Q, L)$ are defined according to the procedure used in Goodson, Thomas, and Ohlmann (2016). The authors used the *R101* and *C101* instances of Solomon (Solomon 1987) and took the first n customers to generate the set \mathcal{C}_0 with $n = 25, 50, 75$, and 100. Concerning expected demands, the authors used the deterministic demands provided by the Solomon instances. To determine the number of vehicles m and the duration limit L , Goodson, Thomas, and Ohlmann (2016) solved each instance as a classical VRP using a heuristic with a vehicle capacity of 100. The authors used the resulting number of routes as the number of required vehicles and multiplied the length of the longest route by 0.75, 1.25, and 1.75, constructing the set of three possible duration limits, denoted as $\mathcal{L} = \{\text{Short, Medium, Long}\}$. They also varied the vehicle capacity Q to assume values $\{25, 50, 75\}$.

Concerning the distribution functions of the customer demands, they characterized $\Psi_{\bar{d}}$ by an instance identifier U , denoted as the stochastic variability of the demands. The stochastic variability U may have a value in $\mathcal{U} = \{\text{Low}, \text{Moderate}, \text{High}\}$. In particular, $\Psi_{\bar{d}}$ takes the distribution function of $\{\bar{d}/2, \bar{d}, 3\bar{d}/2\}$ with probabilities $(0.05, 0.9, 0.05)$ for Low, $\{0, \bar{d}/2, \bar{d}, 3\bar{d}/2, 2\bar{d}\}$ with probabilities $(0.05, 0.15, 0.6, 0.15, 0.05)$ for Moderate, and $\{0, \bar{d}/2, \bar{d}, 3\bar{d}/2, 2\bar{d}\}$ with a uniform probability for High stochastic variability, respectively. For details, the reader may refer to Goodson, Thomas, and Ohlmann (2016). In consequence, the authors conduct experiments on 216 instances, comprised of 8 different sets of customers (i.e., four levels for the number of customers $\{25, 50, 75, 100\} \times$ two classes of instances $\{\text{R101}, \text{C101}\}$). In particular, for each set of customers with the same locations and expected demands, they did tests on 27 distinct instances.

Testing the DecQN on the full set of instances is time-consuming and beyond the scope of this paper. Therefore, we choose a subset of relatively challenging instances. In particular, we focus on instances with $n = 75$ constructed from R101 Solomon instances. In these instances, the number of vehicles m is 11. We also restrict this experiment to instances with $Q = 50$. Therefore, in this experiment, we identify a VRPSD instance $i \in I$ as a pair (L, U) with $L \in \mathcal{L}$ and $U \in \mathcal{U}$. The set $I = \mathcal{L} \times \mathcal{U}$ of instances contains nine elements. Notice that in the selected instances, duration limits Short, Medium, and Long are 103.05, 171.75, and 240.45, respectively.

6.2.2 Benchmarks

Goodson, Thomas, and Ohlmann (2016) proposed an RA based on a restocking fixed-route policy. The RA is a form of forward-dynamic programming that can be viewed as a one-step policy iteration (Bertsekas 2013). At each decision epoch, they generate a set of fixed routes using a local search heuristic. Each fixed route is then evaluated using a reward-to-go function for a set of sample demand realizations. Finally, the fixed route with the highest reward-to-go is picked, and the next move on that route is taken. To account for the preemptive actions, they solve an auxiliary dynamic program to evaluate the reward-to-go of following a fixed route with and without performing the restocking move. Since the size of the auxiliary dynamic program is highly dependent on the size of the problem, their proposed method cannot handle instances with more than 25 customers. Therefore, they also proposed a dynamic decomposition method to partition customers between vehicles, enabling them to handle instances of up to 100 customers.

Goodson, Thomas, and Ohlmann (2016) obtain the initial fixed route by solving the VRPSD for a given instance $i \in I$ using a Simulated Annealing method. They showed that the high-quality initial fixed route plays a vital role in this method. According to their findings, the RA that starts with a high-quality initial solution outperforms the one that starts with a low-quality (randomly produced) fixed route by an average of 11.2%. We refer to the benchmark policy as GHQ when the algorithm starts with a high-quality fixed route and denote it GLQ when it begins with a low-quality fixed route.

6.2.3 Results and discussion

To test the obtained policies, we construct a set of realizations for each instance $i \in I$, denoted as $\hat{I}(i)$, by generating a set of 500 demand realizations for customers in the VRPSD. Table 3 shows the results of our method for the selected instances along with the benchmarks. In this table, the first two columns indicate the characteristics of the instance. The total served demand by DecQN averaged over 500 realizations ($\hat{I}(i)$) for each instance $i \in I$ is reported as DecQN. Column $(\% \mathbb{E} \sum d_c)$ shows the ratio of demands that DecQN serves. Note that the total expected demand for each instance is 1079. The results of the benchmarks are shown in columns GLQ and GHQ. Columns %Imp. GLQ and %Imp. GHQ display the percentage improvement of our method over GLQ and GHQ, respectively. Notice that in Table 3 duration limits S, M, and L, respectively refer to Short, Medium, and Long. Also, stochastic variability L, M, and H are abbreviations for Low, Moderate, and High, respectively.

Table 3: Results of DecQN compared with Goodson, Thomas, and Ohlmann (2016)

L	U	DecQN	$\% \mathbb{E} \sum d_c$	GLQ	GHQ	%Imp. GLQ	%Imp. GHQ
S	L	793.45	73.54%	738.25	833.95	7.48%	-4.86%
	M	766.97	71.08%	718.45	794.68	6.75%	-3.49%
	H	745.76	69.12%	704.93	760.43	5.79%	-1.93%
Avg:			71.25%			6.67%	-3.43%
M	L	1053.27	97.62%	1019.21	1077.53	3.34%	-2.25%
	M	1036.22	96.04%	1002.8	1066.94	3.33%	-2.88%
	H	1021.5	94.67%	977.11	1043.61	4.54%	-2.12%
Avg:			96.1%			3.74%	-2.42%
L	L	1079.31	100.03%	1078.32	1078.29	0.09%	0.09%
	M	1080	100.09%	1076.31	1076.89	0.34%	0.29%
	H	1079.1	100.0%	1072.28	1075.28	0.64%	0.36%
Avg:			100.04%			0.36%	0.25%
Total Avg:			89.13%			3.59%	-1.86%

Overall, results show that DecQN outperforms the GLQ on average by 3.59%, while GHQ outperforms DecQN by 1.86%. Comparing the percentage improvements, averaged for each level of duration limit, demonstrates the considerable impact of the duration limit on the gap between solution methods. In particular, when the duration limit becomes longer, the absolute value of %Imp. GLQ and %Imp. GHQ fall from 6.67% and 3.43% to 0.36% and 0.25%, respectively. This reduction in the absolute gaps may be explained by observing that as the two methods become capable of serving almost all demands in the VRPSCD (e.g., by extending the duration limit), the range of potential difference between their performance becomes narrower. This logic is also supported by results reported in the fourth column ($\% \mathbb{E} \sum d_c$) of Table 3. Accordingly, the DecQN serves almost all demands (at least 94.67%) when the duration limit is Medium or Long.

Regarding the impact of the stochastic variability on the performances, Table 3 shows a considerable reduction in the absolute gaps (%Imp. GLQ and %Imp. GHQ) as U increases from Low to High, in instances with Low duration limit. This behavior can be explained by looking into the difference between GLQ and GHQ. Figure 5 illustrates the trend of served demands by DecQN, GLQ, and GHQ in different levels of duration limit with respect to changes in stochastic variability. In this figure, results for different duration limits are identified by an indication appended to the policy name. For example, DecQN-S refers to the results of DecQN in instances with a Short duration limit. As shown in this figure, the increase in U narrows the gap between GLQ and GHQ when the duration limit is Low. This finding implies that as the stochastic variability increases, having a high-quality initial solution becomes less important. This figure also explains why the gap between GLQ and GHQ, and consequently %Imp. GLQ and %Imp. GHQ, do not follow a descending trend (as in instances with Short L) when U rises. As can be seen, the amount of served demands in instances with Medium and Long duration limit are very close to the total expected demand (i.e., 1079), making the performance of policies bounded. Hence the gap between GLQ and GHQ cannot follow the same trend as in instances with a Short duration limit. Another important finding that can be concluded from this figure is that while the performance of GHQ, when U increases from Low to High in instances with Short L , drops by 8.82% (from 833.95 to 760.43), the performance of our DecQN decreases only by 6.01% (from 793.45 to 745.76). It can be deduced that our method is relatively better at dealing with stochastic variability than GHQ. This finding is important because our method does not presume any distribution function for stochastic components.

To conclude, it is worth noting that the DecQN is analogous to the GLQ because they both start from a random initial policy and solution, respectively. Results showed that the DecQN method could outperform GLQ by 3.59% on average. On the other hand, the average gap between DecQN and GHQ, where a high-quality initial solution is provided, is -1.86%. However, results showed that the

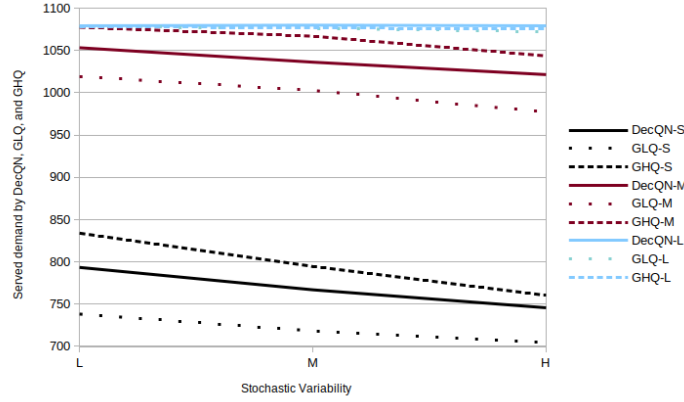


Figure 5: Performance of DecQN, GLQ, and GHQ with respect to the stochastic variability

advantages of starting from a high-quality initial solution become less pronounced for higher values of stochastic variability.

6.3 Generalized DecQN

The previous experiment showed that DecQN could compete with the best-performing method in the literature to solve a VRPSD. However, we had to train a new policy for each instance $i \in I$. In fact, a trained policy for an instance i is not optimized to solve another instance i' . In this experiment, we further explore the flexibility of the DecQN in terms of generalizing it to handle a set of instances. More specifically, we investigate the possibility of training one policy to solve a group of VRPSD instances I' , where $I' \subset I$. This also has some practical advantages, as it would be beneficial for a manager to apply the same policy regardless of, for example, the duration limit.

Using the instances defined in Section 6.2.1 (represented as $I = \mathcal{L} \times \mathcal{U}$), we carry out this experiment in three steps. In the first step, we aggregate instances with respect to the duration limit L . Accordingly, we train three policies, each for a level of $U \in \mathcal{U}$. More specifically, in instances aggregated over L , the duration limit follows a uniform distribution function to pick a value from \mathcal{L} . Similarly, we generalize the method over the stochastic variability U in the second step. For this purpose, for each level of $L \in \mathcal{L}$, we train a policy that can handle different stochastic variabilities. In the last step, we generalize the DecQN over both components. In this setting, a single policy is trained to solve all nine instances. To provide enough information for the policy, we append the value of the component on which we are aggregating to the observation representation. For example, for aggregating over the duration limit, we pass L as an extra component into the Equation (22). We have:

$$o_{k,\bar{v}} = [F_k H_k | G_k | t_k | L]$$

Once we obtain a policy, it is tested on all instances considered in the training phase. Observe that the considered training and testing realization sets are disjoint. For the first step, Table 4 reports the results of the trained policies (DecQN_L) and compares them with those trained in experiment 2 (DecQN) and two benchmark policies (GLQ and GHQ). In this table, the first two columns indicate the characteristics of the test instances. Results show that the aggregation over the duration limit performs on average 1.86% and 3.67% worse than the individual policies (DecQN) and the GHQ, respectively. However, it still outperforms the GLQ by 1.62%. It can also be seen that the gap between DecQN_L and DecQN enhances when the duration limit gets longer. Also, the aggregation works relatively better in higher stochastic variability.

Table 5 reports the results of the second step (DecQN_U), where instances are aggregated with respect to the stochastic variability. Similarly to the previous analysis, we compare the results with

Table 4: Results of DecQN generalized over the duration limit L

U	L	DecQN $_L$	DecQN	GLQ	GHQ	%Imp. DecQN	%Imp. GLQ	%Imp. GHQ
L	S	765.63	793.45	738.25	833.95	-3.51%	3.71%	-8.19%
	M	1027.03	1053.27	1019.21	1077.53	-2.49%	0.77%	-4.69%
	L	1078.96	1079.31	1078.32	1078.29	-0.03%	0.06%	0.06%
Avg:						-2.01%	1.51%	-4.27%
M	S	737.13	766.97	718.45	794.68	-3.89%	2.6%	-7.24%
	M	1017.7	1036.22	1002.8	1066.94	-1.79%	1.49%	-4.62%
	L	1078.22	1080	1076.31	1076.89	-0.16%	0.18%	0.12%
Avg:						-1.95%	1.42%	-3.91%
H	S	722.39	745.76	704.93	760.43	-3.13%	2.48%	-5.00%
	M	1002.76	1021.5	977.11	1043.61	-1.83%	2.63%	-3.91%
	L	1079.77	1079.1	1072.28	1075.28	0.06%	0.7%	0.42%
Avg:						-1.64%	1.94%	-2.83%
Total Avg:						-1.86%	1.62%	-3.67%

the individual (DecQN) and benchmark (GLQ and GHQ) policies. Results show that our method can handle a broader range of instances at the expense of performing on average 0.33% worse than the individual policies. Interestingly, it can be seen that the generalized DecQN outperforms the individual policies in instances with Medium and Long duration limits on average by 0.13% and 0.03%, respectively. Although the generalized model has a gap of -2.18% with the GHQ, it still outperforms the GLQ considerably (3.23%).

Table 5: Results of DecQN generalized over the stochastic variability U

L	U	DecQN $_U$	DecQN	GLQ	GHQ	%Imp. DecQN	%Imp. GLQ	%Imp. GHQ
S	L	773.54	793.45	738.25	833.95	-2.51%	4.78%	-7.24%
	M	760.22	766.97	718.45	794.68	-0.88%	5.81%	-4.34%
	H	745.32	745.76	704.93	760.43	-0.06%	5.73%	-1.99%
Avg:						-1.15%	5.44%	-4.52%
M	L	1051.71	1053.27	1019.21	1077.53	-0.15%	3.19%	-2.40%
	M	1038.64	1036.22	1002.8	1066.94	0.23%	3.57%	-2.65%
	H	1024.55	1021.5	977.11	1043.61	0.30%	4.86%	-1.83%
Avg:						0.13%	3.87%	-2.29%
L	L	1078.68	1079.31	1078.32	1078.29	-0.06%	0.03%	0.04%
	M	1080.28	1080	1076.31	1076.89	0.03%	0.37%	0.31%
	H	1080.3	1079.1	1072.28	1075.28	0.11%	0.75%	0.47%
Avg:						0.03%	0.38%	0.27%
Total Avg:						-0.33%	3.23%	-2.18%

Finally, Table 6 reports the results of testing the single trained policy in the third step (DecQN $_{All}$) on all nine instances. Similar to Tables 4 and 5, we compare the performance of the generalized DecQN with the individual and benchmark policies. Results show that we can train a single policy to solve a set of different instances concerning U and L at the expense of reducing on average 1.72% of the performance. Confronting these results with those obtained in Tables 4 and 5, we can deduce that most of the performance deterioration comes from the generalization over the time limit dimension L . This experiment implies that, in addition to stochastic customers and demands that are uncertain in advance, the generalized DecQN can also handle other unknown instance components (e.g., the duration limit) that only become evident during the operation period.

Table 6: Results of DecQN generalized over both U and L

L	U	DecQN _{All}	DecQN	GLQ	GHQ	%Imp. DecQN	%Imp. GLQ	%Imp. GHQ
S	L	755.88	793.45	738.25	833.95	-4.74%	2.39%	-9.36%
	M	743.02	766.97	718.45	794.68	-3.12%	3.42%	-6.50%
	H	731.63	745.76	704.93	760.43	-1.89%	3.79%	-3.79%
Avg:						-3.25%	3.32%	-6.55%
M	L	1026.51	1053.27	1019.21	1077.53	-2.54%	0.72%	-4.73%
	M	1019.22	1036.22	1002.8	1066.94	-1.64%	1.64%	-4.47%
	H	1009.27	1021.5	977.11	1043.61	-1.20%	3.29%	-3.29%
Avg:						-1.79%	1.88%	-4.17%
L	L	1078.05	1079.31	1078.32	1078.29	-0.12%	-0.03%	-0.02%
	M	1078.72	1080	1076.31	1076.89	-0.12%	0.22%	0.17%
	H	1077.89	1079.1	1072.28	1075.28	-0.11%	0.52%	0.24%
Avg:						-0.12%	0.24%	0.13%
Total Avg:						-1.72%	1.81%	-3.53%

6.4 On-line-off-line DecQN

In the previous experiments, we investigated the performance of the trained Q-network as a policy. More specifically, for a given Q-network that provides Q values for observation-action pairs, the policy applied the action with the maximum Q value. This section investigates the use of the trained Q-network as a reward-to-go approximation function in an on-line RA. In Section 6.4.1, the proposed RA is explained in detail. We use instances and trained policies as illustrated in Section 6.1. Finally, we compare the performance of the proposed on-line-off-line DecQN with the purely off-line method in Section 6.4.2.

6.4.1 Proposed rollout algorithm

The Rollout Algorithm is an on-line method in the form of forward-dynamic programming that can be viewed as a one-step policy iteration (Bertsekas 2013). Bertsekas and Tsitsiklis (1996) showed that if we roll out a known heuristic policy (the so-called base policy) to approximate reward-to-go values, the resulting policy will be improved over the base policy. In practice, the rollout algorithm estimates the reward-to-go value of an available action at the current decision epoch by simulating the base policy after choosing that action. Interested readers may refer to a survey by Bertsekas (2013) for a detailed discussion about RAs.

Similar to what defined in Equation (20), we let $V^\pi(s_k)$ be the expected reward that can be obtained by starting from the state s_k and following the base policy π . Accordingly, we define the rollout policy Π as:

$$\Pi(s_k) = \arg \max_{x_k \in A(s_k)} \left[R(s_k, x_k) + \mathbb{E}_{w \in W} [\gamma V^\pi(s_{k+1})] \right]. \quad (39)$$

In our case, the base policy is provided by the Q values trained by the DecQN method in the previous experiment. Therefore, we substitute the approximate value function $V^\pi(s_{k+1})$ according to Equation (29). Also, following the decentralized formulation, we replace s_k and $A(s_k)$ with $o_{k,\bar{v}}$ and $A(s_k, \bar{v})$, respectively. We rewrite the rollout policy Π as:

$$\Pi(o_{k,\bar{v}}) = \arg \max_{x_k \in A(s_k, \bar{v})} \left[R(s_k, x_k) + \mathbb{E}_{w \in W} \left[\gamma \max_{x' \in A(s_{k+1}, \bar{v}')} Q(o_{k+1, \bar{v}'}, x') \right] \right], \quad (40)$$

where \bar{v}' represents the active vehicle in decision epoch $k+1$.

Note that, since the set of customers is realized at the beginning of the operation, the customer demands are the only source of uncertainty in the equation above. Therefore, the expectation function

$\mathbb{E}(\cdot)$ is computed over a finite set of sample scenarios for the realization of the stochastic demands (W). It is important to note that, according to Equation (18), which defines the reward function for DecQN, the actual reward is realized with a delay. However, it is not feasible to wait to receive the actual reward through on-line methods, such as RA. Therefore, instead of using the reward function as in Equation (18), we define it as follows:

$$R(s_k, x_k) = \begin{cases} \tilde{d}_c & x_k \text{ is a } c \in \mathcal{C}_0 \\ 0 & \text{otherwise} \end{cases}, \quad (41)$$

where \tilde{d}_c is the unserved demand of customer c , if its actual demand is already realized; otherwise, it is the expected demand (Equation (28)).

The RA proceeds at decision epoch k by simulating the system for each possible action $x_k \in A(s_k, \bar{v})$ for one step and observe the immediate reward $R(s_k, x_k)$ as well as the next observation $o_{k+1, \bar{v}'}$. In the next observation, the maximum Q value of available actions is considered as the approximate reward expected to be obtained onward, denoted as \bar{f} . The expected reward-to-go of each action is computed for a set of demand realizations (W). The action with the maximum \bar{f} will be selected as the best action to take (x^*) at decision epoch k .

Given the delayed reward mechanism, we decided to simulate the problem for two steps for each action (instead of the traditional one-step) in our implementation. Therefore, the implemented Equation (40) has been modified as:

$$\Pi(o_k, \bar{v}) = \arg \max_{x_k \in A(s_k, \bar{v})} \mathbb{E}_{w \in W} [R(s_k, x_k) + \gamma R(s_{k+1}, x_{k+1}) + \gamma^2 \max_{x_{k+2} \in A(s_{k+2}, \bar{v}'')} Q(o_{k+2, \bar{v}'}, x_{k+2})], \quad (42)$$

where x_{k+1} is the action in decision epoch $k+1$ according to the base policy π (i.e., $x_{k+1} = \arg \max_{x' \in A(s_{k+1}, \bar{v}')} Q(o_{k+1, \bar{v}'}, x')$) and \bar{v}'' is the active vehicle in decision epoch $k+2$. Algorithm 3 illustrates our proposed RA for each decision epoch k .

6.4.2 Results and discussion

In Section 6.1, we tested each trained policy for 250,000 realizations. However, because the rollout algorithm requires substantial on-line computing, testing across the same realization set would be computationally expensive. Therefore, we test the proposed on-line-off-line DecQN for each instance $i \in I$, with $I = \mathcal{D} \times \mathcal{Q}$, on $\hat{I}(i)$ with 2,500 realizations (50 customer realizations \times 50 demand realizations). For this experiment, we set the size of the sample scenarios for the rollout algorithm ($|W|$) to 50.

Table 7 reports the results of the on-line-off-line DecQN (RA-DecQN) and compares them with the performance of the off-line policy (DecQN). The average improvement of the rollout algorithm over the off-line DecQN is seen in column %Imp. DecQN. According to the results, the proposed rollout outperforms the off-line version by an average of 1.46% in all cases. The average improvement of the rollout algorithm over the off-line DecQN is shown in Table 8. This table helps us to analyze how changes in customer density and vehicle capacity affect %Imp. DecQN. Accordingly, as the customer density decreases from High to Low, the gap widens significantly. Specifically, the average improvement over the off-line DecQN in instances with high customer density is 0.93%, whereas this value in low-density instances is 2.43%. It is essential to notice that, as discussed in the first experiment, the importance of a good routing strategy is magnified in instances with a lower customer density. This finding shows that the on-line-off-line method improves the off-line policy in terms of routing aspects. Furthermore, the sensitivity analysis over the changes in the vehicle capacity also confirms this finding. According to Table 8, the improvement percentage considerably increases from 0.63% to 2.01% when the vehicle capacity changes from 25 to 75. It is worth to notice that as the vehicle capacity becomes larger, vehicles are allowed to perform longer routes which highlights the necessity of an effective routing strategy.

Algorithm 3: Proposed Rollout Algorithm

```

input : Q values,  $s_k$ ,  $\bar{v}$ , W
output:  $x^*$ 

1  $f^* \leftarrow 0$ 
2 build the action set  $A(s_k, \bar{v})$ 
3 for  $x_k \in A(s_k, \bar{v})$  do
4    $\bar{f} \leftarrow 0$ 
5    $o_{k, \bar{v}} \leftarrow O(s_k, \bar{v})$ 
6   take action  $x_k$  and receive the reward  $r_k = R(s_k, x_k)$ 
7   for  $w \in W$  do
8     state transition  $k$  to  $k+1$ :
9      $s_{k+1} \leftarrow S^M(s_k, x_k, w)$ 
10     $\bar{v}' \leftarrow \{v \in \mathcal{V} | a_v = \min_{v \in \mathcal{V}} a_v\}$ 
11     $o_{k+1, \bar{v}'} \leftarrow O(s_{k+1}, \bar{v}')$ 
12    take action  $x_{k+1} = \arg \max_{x'} Q(o_{k+1, \bar{v}'}, x')$  and receive  $r'_k = R(s_{k+1}, x_{k+1})$ 
13    state transition  $k+1$  to  $k+2$ :
14     $s_{k+2} \leftarrow S^M(s_{k+1}, x_{k+1}, w)$ 
15     $\bar{v}'' \leftarrow \{v \in \mathcal{V} | a_v = \min_{v \in \mathcal{V}} a_v\}$ 
16     $o_{k+2, \bar{v}''} \leftarrow O(s_{k+2}, \bar{v}'')$ 
17     $f \leftarrow [r_k + \gamma r'_k + \gamma^2 \max_{x_{k+2} \in A(s_{k+2}, \bar{v}'')} Q(o_{k+2, \bar{v}''}, x_{k+2})]$ 
18     $\bar{f} \leftarrow \bar{f} + f$ 
19  end
20   $\bar{f} \leftarrow \bar{f} / |W|$ 
21  if  $\bar{f} > f^*$  then
22     $x^* \leftarrow x_k$ 
23     $f^* \leftarrow \bar{f}$ 
24  end
25 end

```

In Section 6.1, we demonstrated that the average improvement over GP drops from 11.16% to 7.89%, when the vehicle capacity increases from 25 to 75. However, this experiment shows that the on-line-off-line DecQN partly compensates for this reduction.

Table 7: Results of on-line-off-line DecQN compared with off-line DecQN

Ψ_{n_z}	Q	RA-DecQN	DecQN	%Imp. DecQN
L	25	164.34	161.34	1.86%
	50	202.72	196.87	2.97%
	75	217.08	211.85	2.47%
M	25	358.34	358.22	0.03%
	50	461.00	455.40	1.23%
	75	510.47	501.45	1.80%
H	25	551.71	551.66	0.01%
	50	708.65	701.44	1.03%
	75	790.34	776.71	1.75%
		Avg:	1.46%	

6.5 Computation time

The computation time is an important factor in optimization problems. In fact, real-world problems have very limited computational budget. Unlike on-line methods, where the majority of the computation is done during the process, off-line policies have the advantage of being developed in advance and used whenever needed. The computation time for using the off-line policy to obtain an action is negligible once it is trained. Hence, the computational time of training the off-line policies is usually

Table 8: Results of on-line-off-line DecQN compared with off-line DecQN - Averages

%Imp. DecQN		Q			Avg
		25	50	75	
Ψ_{n_z}	L	1.86%	2.97%	2.47%	2.43%
	M	0.03%	1.23%	1.80%	1.02%
	H	0.01%	1.03%	1.75%	0.93%
Avg:		0.63%	1.74%	2.01%	1.46%

the main concern. Therefore, we only report the approximate times for training the off-line policies. Depending on the number of trials, the training of each model took an average of 24 to 36 CPU hours.

Regarding the computation time in on-line-off-line DecQN, Table 9 illustrates the average CPU time for each decision epoch. Although, as expected, the computation time grows as the number of customers increases, the on-line-off-line DecQN can decide in 0.21 seconds, even in the most complex problems. Therefore, the proposed rollout algorithm can also be regarded as an efficient on-line method.

Table 9: Computation time of on-line-off-line DecQN per decision epoch

	Ψ_{n_z}		
	L	M	H
CPU time (s)	0.10	0.17	0.21

7 Conclusion

In this study, we focused on a version of the VRP in which customer locations and demands are stochastic, and a fleet of capacitated vehicles must complete the service within a specified time limit. The goal is to maximize the total served demand. For this problem, we provided two formulations based on Markov Decision Processes. The first formulation describes the traditional centralized decision-making perspective, while the second models a decentralized decision-making framework where vehicles autonomously establish their routes according to the observed state of the system. This second formulation enabled us to drastically reduce the dimension of the state and the action spaces, thus resulting in a more tractable problem. Variable-sized stochastic customer sets have been addressed via a new heatmap-style representation of the customer demands. This approach, in particular, partitions the stochastic customers geographically and encodes them as a fixed-size vector. We then solved the resulting problem by developing a Q-learning algorithm. We implemented a two-layer artificial neural network to approximate state-action Q factors. Computational results showed that our method significantly outperforms two benchmark policies. Furthermore, we evaluated the DecQN on the VRPSD, which is a particular case of VRPSCD, and showed that our method could compete with the state-of-the-art solution method specialized for that problem. As a further investigation into the capabilities of our framework, we tested and showed that the DecQN could be generalized over several problem components, such as the stochastic variability of demands. Finally, in our last experiment, we demonstrated that the obtained policies and value function can be employed as a base policy and a reward-to-go estimator in an online rollout method.

In respect to possible future lines of research, we observe that the DecQN can efficiently handle problems with stochastic customer sets. It is therefore promising to investigate how the DecQN could be extended to tackle dynamic versions of the VRP, where customers arise dynamically during the execution period. Another potential avenue is to investigate employing deep neural networks to construct a more sophisticated observation function, to better identify relevant information for the active vehicle.

References

- Bellman R, 1958 On a routing problem. *Quarterly of Applied Mathematics* 16(1):87–90.
- Bertsekas DP, 2010 *Dynamic programming and optimal control* (3rd edition) (Athena Scientific).
- Bertsekas DP, 2013 Rollout algorithms for discrete optimization: A survey. Pardalos PM, , Du DZ, , Graham RL, eds., *Handbook of Combinatorial Optimization*, 2989–3013 (New York, NY: Springer New York).
- Bertsekas DP, Tsitsiklis JN, 1996 *Neuro-dynamic programming* (Athena Scientific).
- Bertsekas DP, Tsitsiklis JN, Wu C, 1997 Rollout algorithms for combinatorial optimization. *Journal of Heuristics* 3(3):245–262.
- Bertsimas DJ, 1992 A vehicle routing problem with stochastic demand. *Operations Research* 40(3):574–585.
- Campbell AM, Thomas BW, 2008 Probabilistic traveling salesman problem with deadlines. *Transportation Science* 42(1):1–21.
- Côté JF, Gendreau M, Potvin JY, 2020 The vehicle routing problem with stochastic two-dimensional items. *Transportation Science* 54(2):453–469.
- Dai H, Khalil EB, Zhang Y, Dilkina B, Song L, 2017 Learning combinatorial optimization algorithms over graphs. *Advances in Neural Information Processing Systems 2017-Decem*:6349–6359.
- Dror M, Laporte G, Trudeau P, 1989 Vehicle routing with stochastic demands. Properties and solution frameworks. *Transportation Science* 23(3):166–176.
- Erera AL, Morales JC, Savelsbergh M, 2010 The vehicle routing problem with stochastic demand and duration constraints. *Transportation Science* 44(4):474–492.
- Fan J, Wang X, Ning H, 2006 A multiple vehicles routing problem algorithm with stochastic demand. 2006 6th World Congress on Intelligent Control and Automation, volume 1, 1688–1692 (IEEE).
- Florio AM, Hartl RF, Minner S, 2020 Optimal a priori tour and restocking policy for the single-vehicle routing problem with stochastic demands. *European Journal of Operational Research* 285(1):172–182.
- Gendreau M, Jabali O, Rei W, 2014 Chapter 8: Stochastic vehicle routing problems. *Vehicle Routing*, 213–239 (Philadelphia, PA: Society for Industrial and Applied Mathematics).
- Gendreau M, Jabali O, Rei W, 2016 50th Anniversary invited article future research directions in stochastic vehicle routing. *Transportation Science* 50(4):1163–1173.
- Gendreau M, Laporte G, Séguin R, 1995 An exact algorithm for the vehicle routing problem with stochastic demands and customers. *Transportation Science* 29(2):143–155.
- Gendreau M, Laporte G, Séguin R, 1996 A tabu search heuristic for the vehicle routing problem with stochastic demands and customers. *Operations Research* 44(3):469–477.
- Ghiani G, Laporte G, Musmanno R, 2004 *Introduction to logistics systems planning and control* (John Wiley & Sons, Inc.).
- Ghosal S, Wiesemann W, 2020 The distributionally robust chance-constrained vehicle routing problem. *Operations Research* 68(3):716–732.
- Godfrey GA, Powell WB, 2002 An adaptive dynamic programming algorithm for dynamic fleet management, I: single period travel times. *Transportation Science* 36(1):21–39.
- Goodson JC, Ohlmann JW, Thomas BW, 2013 Rollout policies for dynamic solutions to the multivehicle routing problem with stochastic demand and duration limits. *Operations Research* 61(1):138–154.
- Goodson JC, Thomas BW, Ohlmann JW, 2016 Restocking-based rollout policies for the vehicle routing problem with stochastic demand and duration limits. *Transportation Science* 50(2):591–607.
- Holler J, Vuorio R, Qin Z, Tang X, Jiao Y, Jin T, Singh S, Wang C, Ye J, 2019 Deep reinforcement learning for multi-driver vehicle dispatching and repositioning problem. 2019 IEEE International Conference on Data Mining (ICDM), volume 2019-Novem, 1090–1095 (IEEE).
- Jaillet P, 1988 A priori solution of a traveling salesman problem in which a random subset of the customers are visited. *Operations Research* 36(6):929–936.
- Kingma DP, Ba J, 2014 Adam: A method for stochastic optimization. 3rd International Conference on Learning Representations, ICLR 2015 (International Conference on Learning Representations, ICLR).
- Kool W, van Hoof H, Welling M, 2018 Attention, learn to solve routing problems! 7th International Conference on Learning Representations, ICLR 2019 .
- Kullman ND, Cousineau M, Goodson J, Mendoza JE, 2020 Dynamic Ridehailing with Electric Vehicles. *Transportation Science* In press.

- Laporte G, Louveaux FV, Hamme LV, van Hamme L, Hamme LV, 2002 An integer L-shaped algorithm for the capacitated vehicle routing problem with stochastic demands. *Operations Research* 50(3):415–423.
- Louveaux FV, Salazar-González JJ, 2018 Exact approach for the vehicle routing problem with stochastic demands and preventive returns. *Transportation Science* 52(6):1463–1478.
- Mendoza JE, Rousseau LM, Villegas JG, 2016 A hybrid metaheuristic for the vehicle routing problem with stochastic demand and duration constraints. *Journal of Heuristics* 22(4):539–566.
- Miki S, Yamamoto D, Ebara H, 2018 Applying deep learning and reinforcement learning to traveling salesman problem. 2018 International Conference on Computing, Electronics & Communications Engineering (iCCECE), 65–70 (IEEE).
- Mnih V, Kavukcuoglu K, Silver D, Rusu AA, Veness J, Bellemare MG, Graves A, Riedmiller M, Fidjeland AK, Ostrovski G, Petersen S, Beattie C, Sadik A, Antonoglou I, King H, Kumaran D, Wierstra D, Legg S, Hassabis D, 2015 Human-level control through deep reinforcement learning. *Nature* 518(7540):529–533.
- Narayanan S, Chaniotakis E, Antoniou C, 2020 Shared autonomous vehicle services: A comprehensive review. *Transportation Research Part C: Emerging Technologies* 111:255–293.
- Nazari M, Oroojlooy A, Takáč M, Snyder LV, Takáč M, 2018 Deep reinforcement learning for solving the vehicle routing problem. *Advances in Neural Information Processing Systems* 2018-Decem:9839–9849.
- Novoa C, Storer R, 2009 An approximate dynamic programming approach for the vehicle routing problem with stochastic demands. *European Journal of Operational Research* 196(2):509–515.
- Oyola J, Arntzen H, Woodruff DL, 2018 The stochastic vehicle routing problem, a literature review, part I: models. *EURO Journal on Transportation and Logistics* 7(3):193–221.
- Panait L, Luke S, 2005 Cooperative multi-agent learning: the state of the art. *Autonomous Agents and Multi-Agent Systems* 11(3):387–434.
- Powell WB, 2011 *Approximate Dynamic Programming*. Wiley Series in Probability and Statistics (Hoboken, NJ, USA: John Wiley & Sons, Inc.).
- Powell WB, Topaloglu H, 2006 Approximate dynamic programming for large-scale resource allocation problems. *Models, Methods, and Applications for Innovative Decision Making*, 123–147 (INFORMS).
- Psaraftis HN, Wen M, Kontovas CA, 2016 Dynamic vehicle routing problems: Three decades and counting. *Networks* 67(1):3–31.
- Qin ZT, Tang X, Jiao Y, Zhang F, Xu Z, Zhu H, Ye J, 2020 Ride-hailing order dispatching at DiDi via reinforcement learning. *INFORMS Journal on Applied Analytics* 50(5):272–286.
- Ritzinger U, Puchinger J, Hartl RF, 2016 A survey on dynamic and stochastic vehicle routing problems. *International Journal of Production Research* 54(1):215–231.
- Secomandi N, 2000 Comparing neuro-dynamic programming algorithms for the vehicle routing problem with stochastic demands. *Computers & Operations Research* 27(11-12):1201–1225.
- Secomandi N, 2001 A rollout policy for the vehicle routing problem with stochastic demands. *Operations Research* 49(5):796–802.
- Secomandi N, Margot F, 2009 Reoptimization approaches for the vehicle-routing problem with stochastic demands. *Operations Research* 57(1):214–230.
- Solomon MM, 1987 Algorithms for the vehicle routing and scheduling problems with time window constraints. *Operations Research* 35(2):254–265.
- Stone P, Veloso M, 2000 Multiagent systems: A survey from a machine learning perspective. *Autonomous Robots* 8(3):345–383.
- Sutton RS, Barto AG, 2018 *Reinforcement learning: An introduction* (MIT Press).
- Sutton RS, McAllester D, Singh S, Mansour Y, 2000 Policy gradient methods for reinforcement learning with function approximation. *Advances in Neural Information Processing Systems*, 1057–1063.
- Tang X, Qin ZT, Zhang F, Wang Z, Xu Z, Ma Y, Zhu H, Ye J, 2019 A deep value-network based approach for multi-driver order dispatching. *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, volume 11, 1780–1790 (New York, NY, USA: ACM).
- Teodorović D, Pavković G, 1992 A simulated annealing technique approach to the vehicle routing problem in the case of stochastic demand. *Transportation Planning and Technology* 16(4):261–273.
- Toth P, Vigo D, 2014 *Vehicle Routing* (Philadelphia, PA: Society for Industrial and Applied Mathematics).
- Van Hasselt H, Guez A, Silver D, 2016 Deep reinforcement learning with double Q-learning. *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 30.

- Vidal T, Laporte G, Matl P, 2020 A concise guide to existing and emerging vehicle routing problem variants. *European Journal of Operational Research* 286(2):401–416.
- Watkins CJCH, Dayan P, 1992 Q-learning. *Machine Learning* 8(3–4):279–292.
- Williams RJ, 1992 Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning* 8(3-4):229–256.
- Zhang K, He F, Zhang Z, Lin X, Li M, 2020 Multi-vehicle routing problems with soft time windows: A multi-agent reinforcement learning approach. *Transportation Research Part C: Emerging Technologies* 121:102861.
- Zhu L, Rousseau LM, Rei W, Li B, 2014 Paired cooperative reoptimization strategy for the vehicle routing problem with stochastic demands. *Computers and Operations Research* 50:1–13.