



This is a repository copy of *Dynamic multi-period vehicle routing with touting*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/204576/>

Version: Published Version

Article:

Keskin, M. orcid.org/0000-0001-6799-4917, Branke, J., Deineko, V. et al. (1 more author) (2023) Dynamic multi-period vehicle routing with touting. *European Journal of Operational Research*, 310 (1). pp. 168-184. ISSN 0377-2217

<https://doi.org/10.1016/j.ejor.2023.02.037>

Reuse

This article is distributed under the terms of the Creative Commons Attribution (CC BY) licence. This licence allows you to distribute, remix, tweak, and build upon the work, even commercially, as long as you credit the authors for the original work. More information and the full terms of the licence here:

<https://creativecommons.org/licenses/>

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



eprints@whiterose.ac.uk
<https://eprints.whiterose.ac.uk/>



Contents lists available at ScienceDirect

European Journal of Operational Research

journal homepage: www.elsevier.com/locate/ejor

Production, Manufacturing, Transportation and Logistics

Dynamic multi-period vehicle routing with touting

Merve Keskin^{a,*}, Juergen Branke^b, Vladimir Deineko^b, Arne K. Strauss^c^a Sheffield University Management School, Sheffield, S10 1FL, United Kingdom^b Warwick Business School, University of Warwick, Coventry, CV4 7AL, United Kingdom^c Mercator Endowed Chair of Demand Management & Sustainable Transport, WHU - Otto Beisheim School of Management, Vallendar 56179, Germany

ARTICLE INFO

Article history:

Received 29 October 2021

Accepted 25 February 2023

Available online 1 March 2023

Keywords:

Routing

Multi-period

Demand management

Touting

Metaheuristics

ABSTRACT

This paper introduces a dynamic multi-period vehicle routing problem with touting as demand management technique, where customers that have not yet placed an order can be actively encouraged to order a service sooner. Touting the right customers, such as those located nearby customers who already placed orders, allows for more efficient routes over time. However, it also increases the frequency of visits at such touted customers as they are serviced before they would normally require, which leads to smaller demand volumes per visit. To tackle this trade-off, we propose several strategies to decide which customers to tout and when, using the characteristics of the customers as well as the current plan at the time of touting. Specifically, using the demand and the location information, we approach the ones which are close to the current tour, relatively far from the depot and not likely to easily be covered in the near future. This information is then used as a part of different touting strategies, which are further embedded in a rolling-time horizon vehicle routing algorithm to address the multi-period nature of the problem. These different strategies are empirically compared in a simulation based on a real-world waste collection problem. We demonstrate that touting indeed allows to significantly reduce the travel distance in a dynamic vehicle routing problem.

© 2023 The Authors. Published by Elsevier B.V.

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

1. Introduction

In this paper, we consider a vehicle routing problem (VRP) with due dates and customer orders (e.g., for collections of goods) arriving dynamically over multiple periods. The addressed problem is motivated by a real-world problem faced by a waste collection company. Naturally it is more efficient to collect waste from customers located in the same region than from customers that are spread out. Therefore we investigate the option to actively influence demand by approaching additional customers that can be easily integrated into the current route and encouraging them to order sooner, a demand management technique known as *touting*. We assume that all customers are loyal and regular customers, whose demand distributions are known or can be estimated. Using these distributions, we predict the customers who are likely to place an order in the near future and tout these. Our problem is related to dynamic VRPs with stochastic customer arrivals, where the information of the demand distributions of the customers is available

and can be used for future planning. However, our approach is novel, as we examine the benefit of touting, i.e., interacting with customers and encouraging them to order earlier. The outcome of touting is determined by the customer, who may accept or reject the touting offer. In the former case, the customer is served on the next route, otherwise, there is no change with respect to the customer's status. So touting will not increase overall demand, just shift the timing of the demand. Because of this, touting may seem counter-productive, as it is likely to increase the number of required visits to a customer as customers are encouraged to order earlier, and thus smaller amounts. However, as we show in this paper, by specifically touting customers that fit well with the already acquired customer orders, touting can significantly reduce the distances traveled to service all customers. To the best of our knowledge, this is the first paper considering touting strategies in vehicle routing.

Note that the focus of this study is not on proposing a novel algorithm to solve dynamic VRPs, but on linking dynamic VRP solvers to the demand management technique of touting. We propose several heuristics to decide which customers should be touted, and when, in order to minimize overall distances traveled. The touting strategies are then combined with a vehicle routing algorithm operating on a rolling time horizon. We test the resulting

* Corresponding author.

E-mail addresses: m.keskin@sheffield.ac.uk (M. Keskin), juergen.branke@wbs.ac.uk (J. Branke), vladimir.deineko@wbs.ac.uk (V. Deineko), arne.strauss@whu.edu (A.K. Strauss).

algorithms in a simulation based on a real-world waste-collection problem of an industry partner that motivated our work.

We would like to stress that the idea of touting is not limited to waste collection; it can be beneficial in a wide variety of similar dynamic vehicle routing problems, involving collection, delivery, or service. For example, companies that offer preventive maintenance services may need to repeatedly serve customers by known deadlines and could contact some customers to check whether they would accept an early service. In delivery problems with deadlines, we likewise may want to pro-actively contact certain customers whom we believe are likely to be interested in a delivery (of food, for instance) and whose locations fit well with the current planned routes. In the following, we stick to the waste collection application to explain our approach.

The paper is organized as follows. We start with a review of related literature in Section 2, followed by a description of the considered problem in Section 3. Section 4 explains the proposed touting strategies as well as the methodology to solve the routing problem. The simulation model with empirical results of the different strategies are presented in Section 5. Finally, Section 6 summarizes the paper and suggests some avenues for future research.

2. Literature review

Vehicle routing is a very large and diverse research area (Braekers, Ramaekers, & Van Nieuwenhuysse, 2016), and even the class of dynamic vehicle routing problems would be too large to be covered here. The reader may refer to Pillac, Gendreau, Guéret, & Medaglia (2013), Ritzinger, Puchinger, & Hartl (2016), Ulmer, Goodson, Mattfeld, & Thomas (2020), or Soeffker, Ulmer, & Mattfeld (2022) for detailed reviews on dynamic VRP literature.

In this section, we will focus on the papers most relevant to our work, primarily on the class of multi-period vehicle routing problems (MPVRPs).

2.1. Nature of the demand - static or dynamic

In static problems, all data is available upfront. Although in our problem the customers arrive dynamically over time, we plan on a rolling horizon, and each sub-problem could be regarded as a static problem. In this sense, our problem is related to those static problems involving due dates. Athanasopoulos & Minis (2013) presented a general model for the MPVRP, where the tasks have time windows and allowed visit days. Each period has a time constraint and the objective is to minimize transportation cost of the entire horizon. Archetti, Jabali, & Speranza (2015) and Larrain, Coelho, Archetti, & Speranza (2019) considered customer release and due dates, and their objective is to minimize the distance traveled as well as the inventory holding cost at the depot until the goods are shipped to the customers.

In the dynamic MPVRP, new orders arrive dynamically over time and the initial plan may need to be revised to take into account the new information. In this case, the problem at hand changes due to the changed set of customers, and therefore, the solution found at the beginning of a period may not be the optimal solution of the new problem in the next period (Ozbaygin & Savelsbergh, 2019). Angelelli, Grazia Speranza, & Savelsbergh (2007) studied a variant in which a set of orders are revealed at the beginning of each time period and they have to be served either in that period or in the next one by a single uncapacitated vehicle. The plans are modified at the beginning of each period depending on the new orders. Wen, Cordeau, Laporte, & Larsen (2010) optimized the routes such that the total customer waiting time and travel time are minimized, and daily workload of the vehicles is balanced over the planning horizon. Cordeau, Dell'Amico, Falavigna, & Iori (2015) addressed an auto-carrier transportation

problem with heterogeneous fleet and introduced penalties for late deliveries. To address the dynamic nature of the problem, a rolling-horizon approach is used.

2.2. Plan update

While the above papers assume that the plan can only be changed once at the beginning of every period (as we do in our paper), some authors have considered the case where the plan may be changed as soon as a new order arrives, re-directing the vehicle to follow the updated plan. Ninikas, Athanasopoulos, Zeimpekis, & Minis (2014) allowed diversion from the current plan for new orders requesting urgent service, while due dates of regular orders still need to be obeyed. Their objective is minimizing overall transportation cost and maximizing the number of urgent customers covered over the horizon. Dayarian, Crainic, Gendreau, & Rei (2015) and Dayarian, Crainic, Gendreau, & Rei (2016) analyzed the case where the customers have stochastic demands. The initial plan is constructed using the probabilistic information. Then during the day, if the realized demand is higher than expected, vehicles are diverted to the depot to unload. Angelelli, Bianchessi, Mansini, & Speranza (2009) studied different strategies to solve a 2-period routing problem. At the beginning of each period, they route a fleet of uncapacitated vehicles using the known customers. Then, while the vehicles are traveling, new orders arrive and the routes are replanned. The new plan may postpone some customers originally scheduled today to the next day. The objective is to service all requests with minimum average cost. Ulmer, Soeffker, & Mattfeld (2018) studied the same problem by increasing the number of periods in the planning horizon. However, their objective is to maximize the number of same-day services.

2.3. Use of historical information

In some studies, knowledge on future demands is used in route planning. Subramanyam, Mufalli, Pinto, & Gounaris (2017) modeled the problem as a robust multi-stage optimization problem that hedges against customer order uncertainty, whose support functions are known from historical orders. The objective of the problem is minimizing the transportation cost over the planning horizon. Billing, Jaehn, & Wensing (2018) also used historical data to obtain probabilities that customers place an order at a period and used these to make decisions about whether existing orders should be served today or be postponed to a later period. Albareda-Sambola, Fernández, & Laporte (2014) modelled the dynamic VRP as a prize-collecting VRP by assigning prize measures to known customers using the information on future orders. They aim at routing the known customers such that the plan is also convenient for likely future requests. Ferrer & Alba (2019) considered a waste collection problem with prediction of the fill levels of the containers. The route planning is done based on these predictions.

2.4. Steering the demand in planning

In a MPVRP context, few papers consider demand management to decrease transportation costs. Estrada-Moreno, Savelsbergh, Juan, & Panadero (2019) analyzed effects of price discounts offered to customers to relax their preferred delivery day by one day, either to the day before or after the preferred day. The aim is to minimize total distribution cost and discounts paid over the planning horizon. Yildiz & Savelsbergh (2020) solved a simplified setting, where all the nodes are located on a line with the depot on one end. They considered discounts in exchange for delivery day flexibility, however assumed that they are accepted by the customers with a certain probability only. The customer set is static in

both studies. On the other hand, in attended home delivery literature, several demand management techniques are studied. In recent years, online grocery sales have increased rapidly, especially after the pandemic many people prefer their groceries delivered to their home instead of shopping in-store. Therefore, attended home delivery problems have become important. In this context, demand management techniques are extensively studied to steer the customers in a desired direction. One approach uses price incentives in the shape of discounts for some delivery slots to have more efficient plans (Yang & Strauss, 2017; Yang, Strauss, Currie, & Eglese, 2016). Another technique is to tailor the provided service or products according to the customers having the service. For example, some delivery slots may not be opened to a specific geographical area, some areas may be offered more slots, or the length of the slot for an area may be adjusted according to some features of the customers or of the area (Agatz, Campbell, Fleischmann, & Savelsbergh, 2011). Recently, Agatz, Fan, & Stam (2021) have analysed using *green labels* that are defined as 'environmentally friendly time slots' to motivate selection of specific delivery times.

2.5. Inventory routing problems

Our problem is conceptually also similar to the Stochastic Inventory Routing or Vendor Managed Inventory Routing Problems, where the decisions are centralized at the supplier and the customer demand is stochastic. Some papers in this area (Coelho, Cordeau, & Laporte, 2014; Markov, Bierlaire, Cordeau, Maknoon, & Varone, 2018; 2020) assume that the supplier can precisely monitor the inventory level at customer sites, and needs to ensure that customers do not have stock-outs. The aim is to minimize both the transportation and inventory holding costs. Other papers assume that the supplier has no knowledge of the inventory levels and will only observe the inventory level when arriving at the customer (Huang & Lin, 2010; Jaillet, Bard, Huang, & Dror, 2002, and Ketzenberg & Metters, 2020). Hence, the plans are based on expected values and recourse actions are defined if the plans do not meet the actual requirements of the customers. Our problem sits between the above two cases: the supplier cannot monitor the customers' inventory levels but has information on the historical demand, and as soon as a customer requests a collection, the amount to be collected is revealed by the customer, i.e., the demand becomes known. In addition, it is possible to call (tout) a customer, offer them a service, and, if the customer accepts the offer and places an order, again the amount to be collected becomes known.

Table 1 provides an overview of similarities and differences of our study compared to the existing literature in multi-period vehicle routing. In summary, our study extends the current literature of multi-period vehicle routing problem by introducing the concept of touting, i.e., actively approaching a customer with the purpose to elicit orders earlier, and combining this demand management technique with routing decisions.

3. Problem definition

Our work is motivated by the real-world application of a waste collection company in the UK. This section first illustrates the procedures of the company, then explains how we model them in our paper.

3.1. Real-life challenge

The company has a fleet of waste collection vehicles with different capacities. It services a stable set of customers, from whom they collect the waste products. Customers accumulate waste over time, i.e., on one day a high amount of waste may be produced, whereas on another day the waste production may be low. The

waste can be stored in tanks, and when the accumulated amount gets close to their storage capacity, they request a collection, specifying the amount of waste to be collected. The waste collection company promises service within a certain number of days. While the waste is being collected, there is a service time spent at the customer site. If the driver knows the client's premises, this reduces the time needed to collect the waste. For this reason, the company has decided to assign each driver to a designated service area, which allows an independent planning of each driver's route. That is, a multi-period single vehicle routing problem is addressed in each service area.

The company plans the routes of the vehicles on a daily basis. Throughout the day, the customer service department collects orders from the customers and at the end of the day, the routes of the collection vehicles for the following day are decided subject to the capacity of the vehicles, daily and weekly working time limits for the drivers, as well as due dates of the customers. The routes are constructed based on the customers who have already requested collection. As the customers are to be serviced within a certain number of days, the problem spans several days. Planning is done on a rolling horizon, i.e., the route for the following day is executed as planned, while there might be changes in the other routes as further requests are received. The stakeholders involved in this process are the customers, the customer service department and the route planners of the company as well as the drivers of the waste collection vehicles. The objective of the company is to service its customers with the least transportation cost, i.e., with the minimum total distance driven by the collection vehicles.

The company does not perform any forecasting on predicting the potential customers. This is done by the drivers based on their experiences. If they know that a customer has not requested a collection for some time and if they are servicing an area nearby that customer, then they ask the customer service department to communicate with that customer and check whether they need a collection. This action is called "touting". If the customer is happy with the collection, then they are added to the current plan and the driver collects their waste within the day. This way, the customers nearby those already scheduled for the current route can be serviced without a large detour.

3.2. Touting for the vehicle routing problem

Our paper focuses primarily on this touting aspect. Before finalizing the next day's tour, the company can attempt to elicit additional orders via touting. In the model, the drivers' experience on predicting the potential customers is replaced by a forecasting model. When a potential customer is approached and offered a collection on the next day, they will accept this offer with a probability depending on the current fill level of their storage tank. If the touted customer accepts the offer, they will be serviced on the next day's tour. Thus, customer requests are received during the day, while touting only takes place at the end of the day when there are no more requests. Figure 1 illustrates the sequence of decisions and the random events on a planning day, where the chart on the right shows the flow of events between 5pm and 6pm.

At 8am the company starts receiving orders from the customers via phone calls or emails. At the same time, the drivers execute their planned routes, which have been finalized on the day before, by visiting the customers in the planned order. At a cut-off time (5pm), the company starts planning and touting for the next day. Using the known requests, i.e., the orders received until 5pm as well as the unserved orders from previous days, a tentative routing plan is constructed. After that, a forecasting model is run to predict customers which are likely to request a service soon. If there are such customers, then they are added to the potential customers list. If this list contains some customers which could be

Table 1

Multi-period vehicle routing literature overview. It highlights the nature of the customers (static or dynamic), the nature of the demand (deterministic, i.e., amount revealed at time of order, or stochastic, i.e., amount revealed when vehicle arrives at customer location), whether i) the method used in the paper anticipates the future demand and uses this information in the optimization, ii) the customers have due dates, and iii) the plans are re-optimized according to the revealed information while the vehicle is on the road and finally the demand management technique if the paper includes one.

Papers	Set of customers		Demand		Anticipation of demand	Due dates	Re- dates opt.	Demand manag.
	static	dynamic	det.	stoc.				
Athanasopoulos & Minis (2013)	✓		✓					
Archetti et al. (2015)	✓		✓					
Larrain et al. (2019)	✓		✓					
Angelelli et al. (2007)		✓	✓					
Wen et al. (2010)		✓	✓					
Cordeau et al. (2015)		✓	✓					
Ninikas et al. (2014)		✓	✓		✓			
Dayarian et al. (2015)	✓			✓	✓			
Dayarian et al. (2016)	✓			✓	✓			
Angelelli et al. (2009)		✓	✓					
Ulmer et al. (2018)		✓			✓			
Subramanyam et al. (2017)		✓	✓		✓			
Billing et al. (2018)		✓	✓		✓			
Albareda-Sambola et al. (2014)		✓	✓		✓			
Ferrer & Alba (2019)		✓	✓		✓			
Estrada-Moreno et al. (2019)	✓		✓					discounts
Yildiz & Savelsbergh (2020)	✓		✓					discounts
Coelho et al. (2014)		✓		✓	✓			
Markov et al. (2018)		✓		✓	✓			
Markov, Bierlaire, Cordeau, Maknoon, & Varone (2020)		✓		✓	✓			
Jaillet et al. (2002)		✓		✓	✓			
Huang & Lin (2010)		✓		✓	✓			
Ketzenberg & Metters (2020)		✓	✓	✓	✓			
Our paper		✓	✓		✓			touting

added to the next day's plan making it more efficient, i.e., by collecting a large amount of waste without a long detour, then the most relevant of those is identified, and this customer is contacted and asked whether they would accept a collection on the next day, in other words, they are 'touted'. If the customer accepts having a collection on the next day, they are added to the tentative plan and removed from the potential customers list, otherwise they are also removed from this list to prevent further communication with the customer. After that, other suitable customers are touted, until there are no more such customers or no additional customer can be feasibly added to the existing plan. These events take place until 6pm, after which the route plans for the next day are finalized. Of course, the schedule used here is representative only, there may be different settings in different practices.

There are two types of decisions: selecting the customers to tout, and routing the current customers. Figure 2 illustrates a series of touting and routing decisions along with the changes in the system depending on these decisions and on the exogenous information. Assume that the vehicle capacity is 4, all customers have a unit demand, and at most one customer can be touted (this is for illustration purposes, there is no such restriction in the model). At the beginning (Fig. 2-1), i.e., at 5pm on day 1, there are 5 customers with known orders (white circles) and 3 potential customers, i.e., predicted to have a sufficient amount of waste through a forecasting model (grey circles). Using the known orders, an initial plan is constructed as illustrated in Fig. 2-1, i.e., customers 1, 2, 3, and 4 are planned to be serviced the next day, and customer 5 is left for the following day. Note that because of the vehicle's capacity restriction it cannot service the fifth customer. However, as long as their due dates are respected, the planner may want to postpone some customers to a later day if they require long detours to be serviced and it is expected that the vehicle may visit these areas in the future. The first decision is to tout customer 7, which changes the state of the system as shown in Fig. 2-2. If customer 7 accepts the touting offer, then the system evolves to a state as shown in Fig. 2-3.a, which is then followed by a routing decision

resulting in the tour for the following day as in Fig. 2-4.a. Note that customer 4 has been removed from the tour to be serviced on another day so that it is feasible to include customer 7. If customer 7 does not accept the touting offer, then the system evolves to a state as shown in Fig. 2-3.b. Another customer may be considered for touting, e.g. customer 6, which changes the state of the system as shown in Fig. 2-4.b. Their acceptance of the touting offer evolves the system into a state as shown in Fig. 2-5. After that, the routing decision is made as shown in Fig. 2-6, in which customer 1 is replaced by customer 6.

This is essentially a multi-period dynamic vehicle routing problem with capacity constraints, time constraints and due dates. The planning is done for a single vehicle, which belongs to the driver operating in the area under consideration. We assume that all customers are regular customers known to the company, as the number of new customers is negligible. Customers request a collection when the amount of their waste products reaches a threshold, and the company knows this amount only when the customer specifies it while placing the order. While it has been considered to install smart sensors at the clients' premises for monitoring their inventory levels, in practice this has been deemed too expensive and is therefore not done. Instead, historical data may be used to forecast a customer's inventory level.

Let us emphasize again that touting does not generate additional demand, it just nudges a customer to order earlier, and as a consequence, the amount to be collected is less than that if the company had just waited for the customer to place the order. While this implies more frequent visits to a customer picking up smaller amounts, and thus higher cost, it also allows to influence the timing of the order, thus opening the opportunity to save travel distance by visiting neighboring customers on the same day, or moving demand from high-demand periods to low-demand periods.

Let us return to the simple example from above, assuming the demand of each order is one, the capacity of the vehicle is four, and each customer must be served within two days. In Fig. 3 five

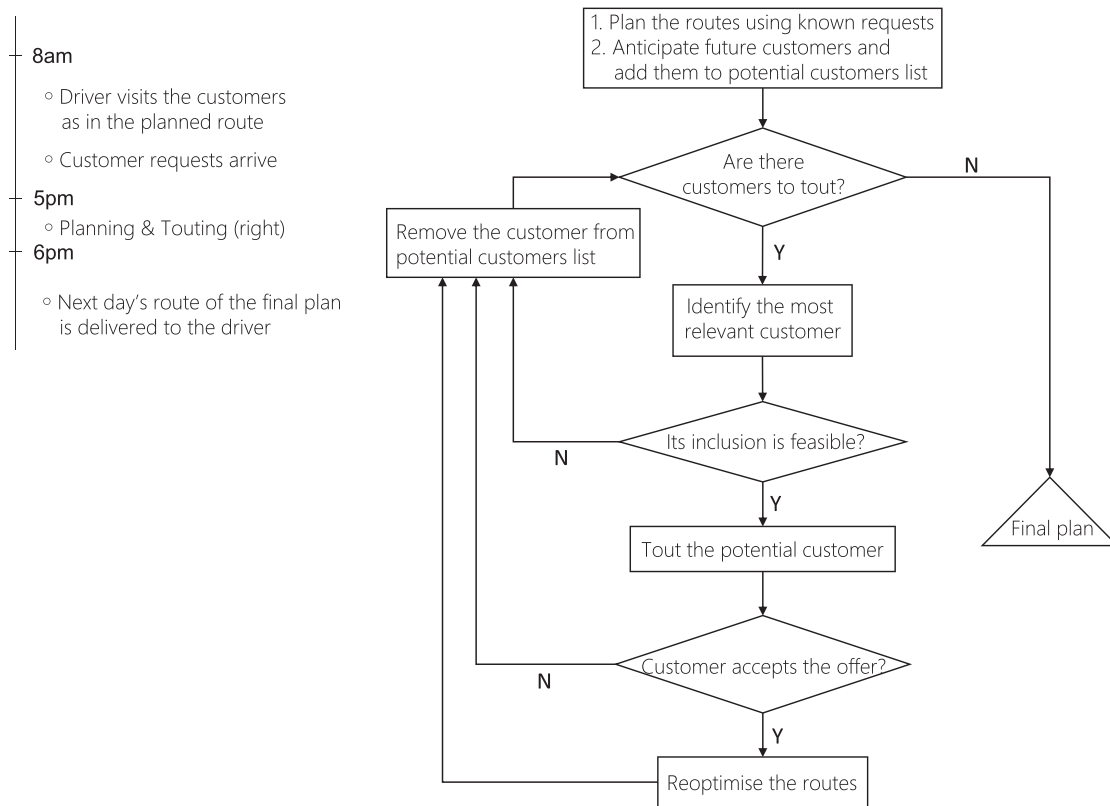


Fig. 1. Sequence of decisions and random events on a planning day.

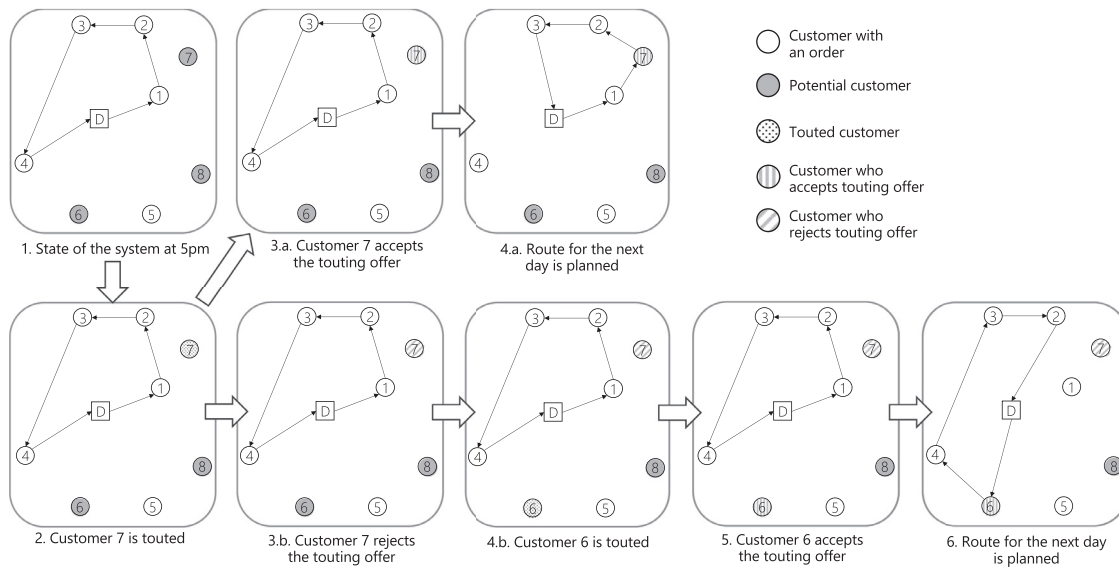


Fig. 2. Example of a series of touting and routing decisions.

customers (unfilled circles) are known on day 1 and the others will place their order on day 2. If touting is not used, the planner would route the corresponding known customers on each day, and obtain a result as depicted in Fig. 3-a. With touting, the planner may predict some customers who are likely to order soon, and tout for example customer 7, as their location is close to customers that would be visited tomorrow. If the customer accepts, they can then be serviced on the next day, and the resulting tours would look like in Fig. 3-b, with a significantly shorter overall distance.

In order to help clarify the structure of the decision making problem, we provide a dynamic programming formulation of the problem (even though one cannot directly solve it in this way). The details of the formulation can be found in Appendix D.

4. Rolling horizon route planning and demand management

Our paper focuses on integration of demand management with route planning via touting. The actual route planning algorithm used is secondary, but necessary to evaluate our strategies empirically.

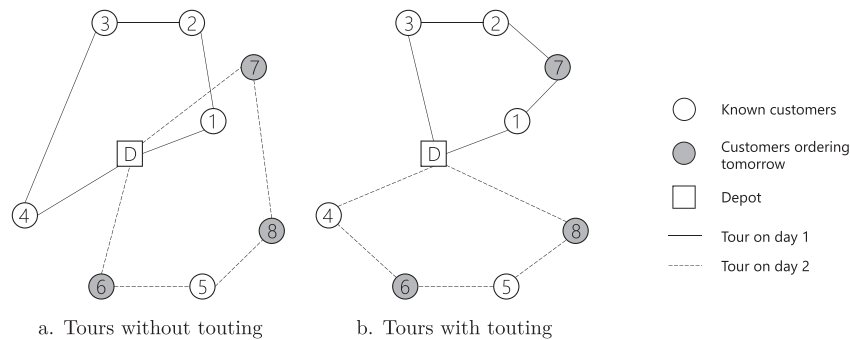


Fig. 3. Route plans with and without touting customers.

ically. Planning at the touting phase needs to have an anticipatory touting strategy for including suitable customers to the plan as well as a fast and high quality routing procedure so as to ensure that we arrive at efficient overall routes. Therefore, the proposed methodology is essentially a combination of routing and touting algorithms, which are applied to find the solution on a planning day. For the former, we use Large Neighborhood Search (LNS) to construct the routes the vehicle will be performing. It uses the information of the customers with existing collection requests and proposes a route plan to serve these customers. For the latter, we use a touting strategy, which will be elaborated in Section 4.2. It uses the information of all customers in the portfolio of the company, e.g., their history of collections to i) determine potential customers to tout, ii) and among them to select the most relevant one to approach. Within this strategy, we also utilize the routing algorithm to calculate how the solution would be revised in case the touted customer accepts the offer and is included in the solution. The combination of these algorithms continues to be used until there are no customers that can feasibly be added to the solution. This procedure is repeated on every planning day with the updated information. Thus, in this section we will describe our routing algorithm and the touting strategies.

4.1. Route planning

Given the dynamic nature of the problem, planning is done on a rolling horizon. That is, at the end of each day, we solve a multi-period VRP with all the order data available. On the next day, we execute the plan for this day, remove all serviced customers from the set of orders, and add any new orders that arrived during this day. Then we solve the next multi-period VRP with this new order data and the cycle repeats.

To solve the VRP of each day, we chose to use LNS. This was motivated by the fact that VRP is NP-hard and thus an exact method is computationally expensive, but also because for a dynamic problem, solving each sub-problem of the rolling horizon procedure exactly does not guarantee overall optimality anyway, as we will later demonstrate in Section 5.2. LNS starts by constructing an initial solution via cheapest insertion. The orders are sorted according to their due dates and the ones having the earliest due dates are considered first. Out of those, the customer that can be inserted with the least additional driving distance is inserted, and the procedure is repeated until all orders have been scheduled. This procedure ensures all capacity and time constraints are obeyed and creates new routes as needed. An overview of the algorithm is given in Algorithm 1, where d_{ij} stands for the driving distance between nodes i and j .

The improvement step follows the removal and repair heuristics introduced in Ropke & Pisinger (2006). In each improvement step, some customers are removed from the current solution either randomly, or based on worst marginal distance, proximity time, or

Algorithm 1 Initial Solution Construction.

Input: List of Unscheduled Customers

```

1: while Unscheduled Customers List not empty do
2:   Put the customers having the earliest due dates into the Priority List
3:   while Priority List not empty do
4:     for all customers in Priority List do
5:       for all possible insertion positions in the routes covering
           all time periods do
6:         if insertion of customer  $i$  between nodes  $j$  and  $k$  is
           feasible then
7:           Calculate the insertion cost as:
8:              $(-d_{jk} + d_{ji} + d_{ik})$ 
9:           end if
10:        end for
11:       end for
12:      Determine the customer with the least insertion cost, i.e.,
           customer  $i$ 
13:      Perform the cheapest insertion for customer  $i$ 
14:      Delete customer  $i$  from Priority List and Unscheduled Customers List
15:    end while
16:  end while

```

demand. They are then re-inserted into the partial solution using greedy and regret insertion heuristics. The details of the algorithm are outlined in Algorithm 2 in Appendix B.

4.2. Touting algorithm

Let us assume that an initial plan based on the orders received so far is given, i.e., via the routing algorithm discussed in the previous subsection, and that we have a given route to execute for the next day. We then attempt to make it more efficient by exploiting information (from a forecasting model) about a set of potential customers whom we could contact to elicit their business.

More specifically, our aim is to identify customers who are likely to request service in the near future and to tout the ones that would make the overall plan better. We assume that these potential customers who may require a service soon are available from a forecasting model. To determine a customer to tout among these potential customers, we may consider different criteria, such as required detour from the next day's route, its distance from the depot, or whether it is possible to cover that potential customer in the near future. Figure 4 illustrates these criteria using a solution with four scheduled customers (circles) and a depot (rectangle) as well as several potential customers, shown with grey circles. The values next to the arcs are the lengths of these arcs, which will be used to quantify the measures. In Fig. 4a, touting customer 5

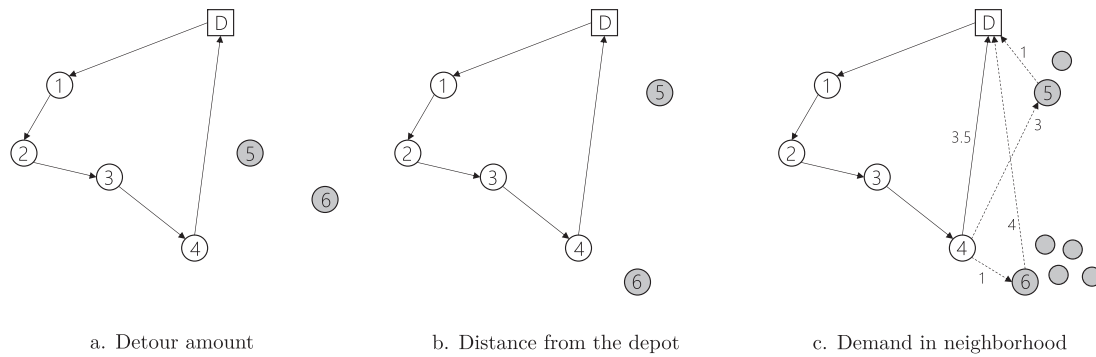


Fig. 4. Factors in selecting the most relevant customer to tout.

seems more advantageous than touting customer 6 in terms of the required detour to insert the customer. Figure 4b shows an example where it may be more beneficial to tout a customer far from the depot if the vehicle is traveling in their neighborhood, than a customer close to the depot, as the latter may be easily included in another tour. Finally, other potential customers in the neighborhood of a potential customer may also be taken into account to decide which customer to tout. If there are many potential customers in the neighboring area, then another vehicle may be sent to cover that area on another day. In Fig. 4c, Customer 6 has many neighbors, whereas Customer 5 has only one, which makes Customer 5 more relevant to tout.

4.2.1. Relevance measure

We now define a relevance measure for each potential customer that takes into account all three criteria mentioned above. Let N be the set of potential customers that may be added to the route, whose calculation is elaborated in Section 4.2.2. We index the depot as 0. For a potential customer $i \in N$, the detour amount is calculated by determining the cheapest insertion cost c_i , i.e., for each pair of consecutive nodes (j, k) in the next day's route with distance d_{jk} , the insertion cost is calculated as $c_i = \text{argmin}(d_{ji} + d_{ik} - d_{jk})$. For example, for customer 5 in Fig. 4c, the insertion cost is $c_5 = 0.5$ and it is $c_6 = 1.5$ for customer 6. Let l_i be the cost of exclusively covering customer i by a single vehicle, i.e., $l_i = d_{0i} + d_{i0}$. For customers 5 and 6 in Fig. 4c, these costs will be $l_5 = 2$ and $l_6 = 8$, respectively. Finally, we define a neighborhood demand measure p_i for customer i based on their neighbors at a distance of less than ρ . More formally, define $N_\rho(i) = \{j \in N | d_{ij} \leq \rho\}$. Let $Q_i = \sum_{j \in N_\rho(i) \cup \{i\}} q_j$ be the total demand of i and their neighbors, where q_i stands for the demand of customer i . We then define $p_i = Q_i / L$, $i \in N$, with L the load capacity of the vehicle, as a measure to assess whether it is worthwhile to drive to an area where customer i is located. To quantify this measure using customers 5 and 6 in Fig. 4c, let us assume that the potential customers have unit demand and the capacity of the vehicle is 5. Then the demand measures for customers 5 and 6 are calculated as $p_5 = 0.4$ and $p_6 = 1$, respectively. If p_i is high, it means that there is high demand in the area and the vehicle would go there anyway to service these customers. Then customer i can be covered along with other customers nearby by another vehicle in the near future, hence it does not necessarily need to be touted right away. However, if p_i is low, meaning that customer i is alone or has few demand around, then it is beneficial to include them in the next day's route.

The relevance measure of a potential customer i is then calculated as $r_i = \alpha \frac{c_i}{M} - \beta \frac{l_i}{I} + \gamma \frac{p_i}{P}$ where α , β and γ are the weights assigned to each criterion, and M , I , and P are used to normalize the different components. They represent the maximum detour

amount to include a customer, the maximum distance to cover a customer with an individual route, and the maximum of the neighborhood demand measures over all eligible customers for touting, respectively. The total distance to cover a customer can be calculated at the beginning of the planning horizon and it does not change. However, the other two values need to be updated in each relevance measure calculation whenever a new maximum is achieved. To set the initial values for the first calculation, we run the simulation once and feed the final maximum values attained at the end of the simulation. The lower the relevance value, the more beneficial to tout that customer since we want to include those customers with a) small detour amount from the next day's tour, b) long distance from the depot, and c) little demand nearby.

4.2.2. Customers considered for touting

Touting customers when they have only accumulated a very small amount of waste will not only annoy the customer, but also lead to very frequent collections of tiny amounts. On the other hand, only touting customers which are predicted to order soon anyway may severely restrict the choice of customers to tout. In the touting algorithm, we restrict touting to customers whose predicted fill level is at least 50% of their tank capacity. Section 5.4.2 will examine the algorithm's sensitivity to this choice.

4.2.3. Waiting vs. touting

Obviously touting only makes sense if the route planned for the next day still has sufficient capacity to incorporate additional customers. On the other hand, if there is little demand and there are no customers who have to be serviced on the next day because of their due date, then it may be more beneficial to simply wait for new orders to arrive and not send the vehicle out at all. Consequently, we consider touting only when the utilization of next day's vehicle is under a threshold, $\Delta\%$ of vehicle capacity, the vehicle must be dispatched because at least one customer must be served the next day, and serving all these urgent customers still allows for some slack in time to potentially serve others.

While we apply this waiting rule in combination with all touting heuristics, waiting is of course also possible if touting is not used. This strategy is called *Wait-if-Possible*, where if the vehicle utilization of the next day's route is under $\Delta\%$ of its capacity and there are no customers due next day. We check whether it is possible to shift the customers in the initial plan to the following day without deteriorating the objective function. If this is possible, the vehicle is not dispatched, hoping that next day we receive orders from conveniently located customers such that we can come up with a more efficient plan. The scheduled jobs remain in the set of open orders. De Bruecker, Beliën, De Boeck, De Jaeger, & Demeulemeester (2018) also stated that instead of having multiple half days work, it may be more advantageous to have a complete day off. For example, the drivers might use these off days for training purposes.

4.2.4. Touting heuristics

Having defined the background information, we now propose the following heuristics to tout customers:

- *Tout using Distance per Litre*: This heuristic touts customers in order of the benefit obtained on the total distance per litre of waste collected after their inclusion, given that the customer is inserted at the position requiring the smallest additional detour.
- *Tout using Relevance Measures*: This approach first touts the customer which has the lowest relevance measure.

It is possible to insert more than one additional customer via touting. We thus also test whether it is beneficial to re-optimize the VRP after the inclusion of every customer, and before the next touting is attempted. More specifically, we distinguish between

- *Tout without Re-Optimization*: Here, we only add customers to the next day's route without re-optimizing all tours.
- *Tout & Re-Optimize*: This strategy allows re-optimizing the routing decisions after the addition of the touted customer in the next day's route.

In all cases, touting continues until either there are no further potential customers or it is not feasible to add another customer to the next day's route. Furthermore, a customer is only touted as long as their insertion does not violate time and capacity feasibility of the solution and if and only if the solution with the additional customer is better than the current solution, i.e., it has a smaller value for total distance per litre of collected waste. Note that, the routing problem solved on each day has the objective of minimization of total distance traveled by the vehicles as it considers a static set of customers at the time of planning. However, if additional customers are included in the tours through touting, then the total demand changes compared to the case where there are no additional customers. Therefore, comparing the solutions to these different sets of customers through their total distances is not fair. Thus, while comparing different strategies, we consider total distance driven per litre of waste collected. Nevertheless, since the amount of waste collected is constant except for end-of-problem effects, these two objectives are equal in the long-run.

5. Experimental study

In this section, we first introduce the instances we use in the experiments. Then, we show that for a dynamic multi-period VRP, solving each period's subproblem with an exact method is no better than the LNS heuristic we employ in this paper. Finally, we perform simulation studies to investigate the proposed touting strategies. The heuristics as well as the simulations were implemented in Java and all computations were performed on a PC equipped with an Intel Core i7 2.60GHz processor and 32 GB of RAM.

In LNS, the maximum number of iterations is set to 4000, and the removal rules select half of the customers to be removed from the current solution. Regarding the percentage of customers removed from the solution, in several papers studying LNS, the percentage of customers to be removed from the solution is randomly selected from the interval [10%–40%]. However, our initial experiments showed that 50% provides better results, which has been observed also in the literature. For example, Liu, Tao, & Xie (2019) conducted experiments with this value but decided to use a lower ratio because of the longer computational times of the former. As we are not constrained with the time because of the small number of customers to be routed on a day, we decided to continue with 50% for more promising results.

5.1. Problem instances

In the experiments, we use real world data from a waste collection company operating in the UK. The company has several depot




Fig. 5. Locations of the customers and the depot for two drivers.

locations across the country and each depot has dedicated drivers who cover different geographical areas. Customers' waste is accumulated in tanks at a random rate which depends on the size of their business. When the accumulated waste is close to the maximal tank capacity, the customer requests a collection. They also state the lead time they require, typically seven working days beginning from the day of request.

As only the data from one depot and its two drivers was made available to us, in the experiments, we use this depot of the company and the historical data for these drivers who operate for that depot. The dataset covers three months of waste collections of these two drivers. Each collection order includes the customer's name, postcode, date when the waste collection order has been placed, and the amount of waste requested to be collected. Customer locations are shown in Fig. 5 where the depot is represented by a red house. Total numbers of collections in the instances are 273 and 260, whereas the numbers of unique customers in the network are 142 and 125 for drivers 1 and 2, respectively. The dataset, which includes these historical orders as well as the problem parameters, such as vehicle capacity and time limits is available at Mendeley Data (Branke, Deineko, & Strauss, 2023).

The service time at a customer location is substantial and cannot be ignored. It consists of a constant setup time and a variable collection time. Setup time includes preparing and dismantling the collection tools before and after the collection, respectively, whereas the collection time is the time spent on actually removing the waste from the tank and is proportional to the amount of waste collected. Since the company does not have records of the times the drivers spend on customer sites, we estimate these service times using linear regression based on data from different drivers' operations. We collected 430 data points and using these values, we found the service time can be approximated by $s = 12.06 + 0.01565w$, where w is the amount of waste collected in litres and s is the service time in minutes.

We use this historical data of the company directly in the experiments presented in Section 5.2, whereas in Section 5.3, we use simulated data, which is generated using the properties of this data.

5.2. Comparing exact and heuristic VRP solvers

Since the requests arrive continuously over time and the problem is solved on a rolling time horizon, an optimal plan for a particular planning horizon may become sub-optimal after arrival of new orders, i.e., solving each period's VRP to optimality does not

Table 2
Comparison of results of solution approaches using exact and heuristic algorithms.

Subproblems are solved	Driver 1		Driver 2	
	exactly	heuristically	exactly	heuristically
Total distance (km)	13,327	12,943	9334	9274
Average lead time (days)	2.52	2.52	3.14	3.16
Average capacity utilization	62.24%	62.78%	74.28%	74.37%
Number of routes	55.0	54.5	56	55.8
Average number of customers served per route	4.96	5.01	4.66	4.65
Computational time (sec)	13,205	28.92	51,058	51.64

guarantee that the overall solution is optimal. This is illustrated through a toy example in Appendix C. For this reason, and because of the computational time required by exact solvers, we use the LNS described in Section 4 in our experiments. However, in order to judge the quality of this heuristic, we compare it with an exact solver. Note that each method is applied to solve the problem for a given single stage on a rolling horizon, rather than solving the whole dynamic problem for the entire planning horizon. The mathematical model can be found in Appendix A, and has been solved using CPLEX.

Table 2 presents the results of both approaches tested on two real-world order datasets. As the heuristic has a randomized component, results are averaged over 30 runs, while the mathematical model is deterministic and solved only once. Capacity utilization is calculated by dividing the amount of waste collected by the vehicle capacity. Lead time shows the average number of days between the day a customer requests a collection and the day it is served. The computational time is the total time to solve the problems over a 3-month horizon. Interestingly for these problem instances, the approach using the LNS heuristic to solve the every-day routing problem is able to find solutions with even smaller total distance than if the exact method is used to solve each day’s VRP. While this may seem surprising, this is due to the loss of overall optimality of the exact method when applied on a rolling horizon. In fact, optimal solutions are often more brittle to changes than heuristically generated ones. These findings have been previously pointed out also in other studies (Brinkmann, Ulmer, & Mattfeld, 2020; Powell, Towns, & Marar, 2000). We take these results as confirmation that the chosen LNS heuristic is fit for purpose.

5.3. A simulator for demand management

In the previous subsection we used the real-world order data. However, touting changes the time of ordering due to some touted customers ordering sooner than expected, and thus, in order to test and compare different touting strategies, we needed to create a simulator based on the real-world data that would allow the touting algorithm to interact with customers.

Therefore, we first create simulated data based on the real instances introduced in Section 5.1. To achieve this, we keep the customer locations the same, as visualized in Fig. 5. However, rather than using a historical set of orders as in the previous subsection, we assume that customers accumulate waste over time, with the amount accumulated each day following a normal distribution with mean μ and variance σ^2 . The parameters of the waste accumulation distributions are derived from real data on the order amounts and the times between consecutive orders using maximum likelihood estimation, and are assumed fixed throughout the simulation. The storage capacity of each customer is set to the maximum amount collected from the corresponding customer. For each customer, the initial amount of waste at the start of the simulation is chosen uniformly at random between 0% and 95% of its capacity. Then, on each planning day, a random value, which represents the waste generated by that customer on that day, is gen-

Table 3
Distance per litre values for different relevance measure parameter settings.

Parameter Values	$\alpha, \beta, \gamma = 1/3$	$\alpha = 1$	$\beta = 1$	$\gamma = 1$
Driver 1	0.02046	0.02061	0.02085	0.02090
Driver 2	0.01362	0.01374	0.01369	0.01387

erated according to the customer’s demand distribution and its accumulated waste is increased accordingly.

The customers request service when their accumulated waste reaches 90% of their storage capacity. When touted before they would usually order, they will agree to a collection with probability $P_{accept} = \frac{w}{0.9W}$ where w is the amount of waste accumulated and W is the customer’s storage capacity. In other words, their probability of accepting a collection increases linearly with the amount of waste accumulated. The neighborhood threshold used in the relevance measure calculation, ρ , is set to 25 km. The threshold for the utilization of next day’s vehicle, Δ , is set to 90%.

The dataset including the distribution parameters of each customer’s waste accumulation, tank storage capacities as well as the distance and travel time matrices is available at Mendeley data (Branke, Deineko, & Strauss, 2023). The simulation horizon has been set to 240 days, corresponding to one business year.

5.4. Analysis of demand management strategies

Here, we summarize the results for the touting strategies presented in Section 4.2.4. We will start with an analysis of the weights in the calculation of our relevance measure. We usually assume that only customers with a predicted amount of accumulated waste that is greater than half of their storage capacity are considered for touting, although we will vary this threshold in Section 5.4.2. Finally, we demonstrate the advantage of touting by comparing it with approaches that don’t use touting. The key objective to minimize is the distance per litre collected, as the total volume collected depends on the touting strategy and thus distance alone is not a suitable objective.

5.4.1. Relevance measure parameters

The relevance measure proposed in Section 4.2.1 is a linear combination of three criteria. To better understand the importance of the different criteria, we run an experiment using the *Tout using Relevance Measures & Re-Optimize* approach where the relevance is calculated using only one of the criteria, or equal weighting of the different criteria. In Table 3, we report the average of the distance traveled per litre of waste collected over the planning horizon for both drivers. The results show that using an equal weighting in the calculation leads to a lower average distance per litre of waste collected for both drivers, thus gives better results. One-sided paired t-tests show that the equal weighting of the three criteria is significantly better at 0.05 level than any of the individual criteria for all cases and for both drivers. Therefore, we use equal weighting in the following experiments.

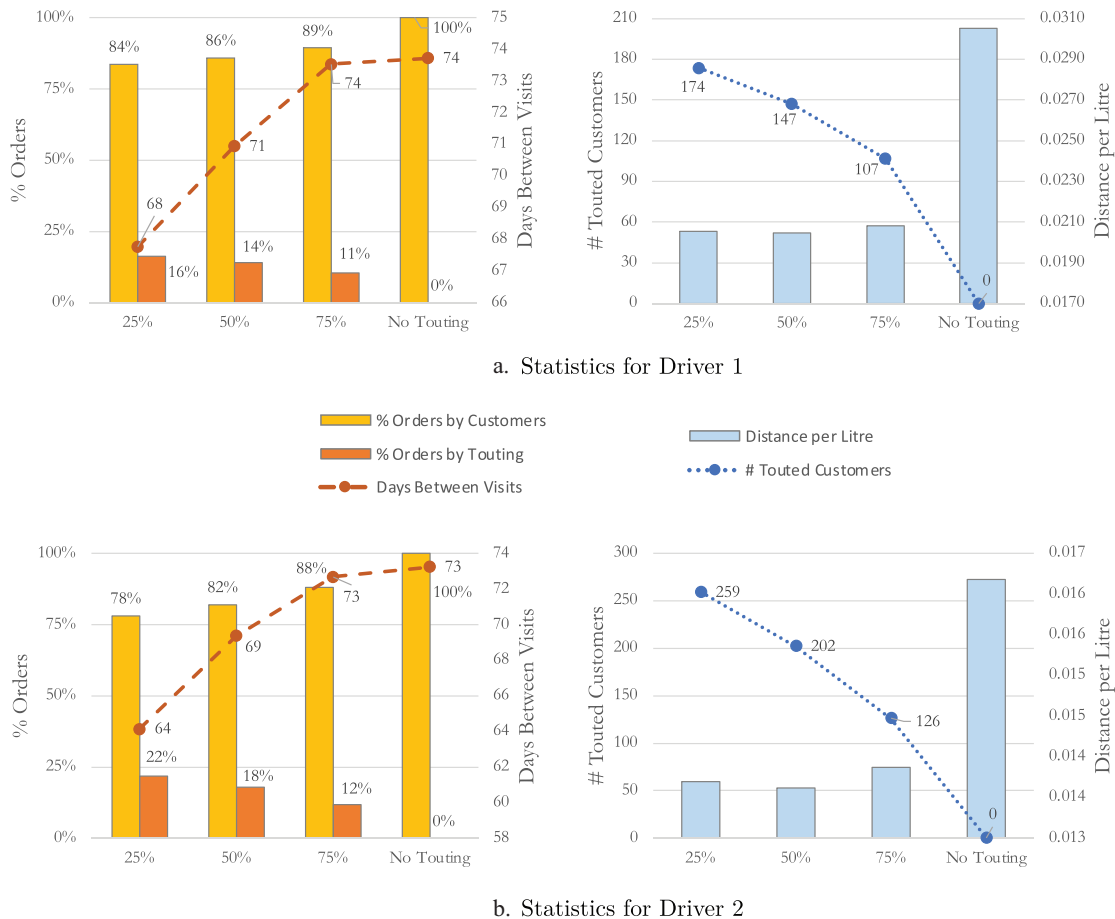


Fig. 6. Statistics for different accumulation levels.

5.4.2. Accumulation threshold

We use a minimum accumulation threshold to determine which customers may be considered for touting. The default setting in this paper is 50%. We increase it to 75% and decrease it to 25%, then we compare the performances of all values for *Tout using Relevance Measures & Re-Optimize* strategy.

In Fig. 6, we report some statistics emerging from the experiments with these accumulation levels. According to these results, while the accumulation threshold decreases, we observe an increasing trend in the number of touted customers because the set of potential customers that are considered in touting gets larger. Therefore, there are more customers to tout and it is easier to find more appropriate customers from a larger set. In parallel, the average number of days between two visits decreases since the customers are visited more frequently due to more toutings. This is also reflected in the % times the customers order by themselves or are touted, with the first percentage increasing while the threshold is increasing. As a result, when the accumulation threshold is set to 25%, visiting customers more frequently slightly decreases the efficiency with a higher total distance traveled per litre of waste collected compared to 50% setting. On the other hand, in the 75% setting, since the set of potential customers is smaller, fewer customers can be approached. This decreases the advantage of touting by not being able to reach the appropriate customers, which results in higher distance traveled per collected amount. We again perform one-sided paired *t*-tests in order to validate the results. The test statistics show that the differences between 25% and 50% as well as between 50% and 75% are significant. Therefore, we con-

clude that 50% is an appropriate threshold for the accumulation level to determine the potential customers to tout.

5.4.3. Benefit of touting strategies

Based on the above analysis, we now test different touting strategies for both drivers using the accumulation threshold of 50% and equal weighting for the relevance parameters ($\alpha = \beta = \gamma = 1/3$). Table 4 summarizes the average results of 100 simulation runs. The results in the *No-Touting* column are obtained without applying any strategies, i.e., they belong to the solutions of the rolling horizon route planning heuristic only. Column 3 summarizes the results obtained by the *Wait-if-Possible* strategy without touting. The results of *Tout using Distance per Litre*, abbreviated by *dist./lt.*, and *Tout using Relevance Measures*, abbreviated by *relevance*, are grouped according to whether they are applied along with re-optimization or not. *Tout without Re-Optimization* strategies are given in the fourth and fifth columns, whereas the last two columns summarize the results of *Tout & Re-Optimize* strategies.

We report the total distance traveled throughout the planning horizon, total number of vehicles dispatched, which is the number of routes generated, average capacity utilization of the vehicle, average lead time (in days) for the customers served, i.e., the number of days between the day a customer requests a collection and the day they have been serviced, the number of touted customers, the number of customer visits during the planning horizon, the average of the distance traveled per litre of waste collected over the planning horizon, and the computational times in seconds required to simulate the entire year. We then compare the performances

Table 4
Results of different strategies.

Strategies		Wait-		No Re-Optimization		Re-Optimization	
		No-Touting	if-Possible	Dist./Lt.	Relevance	Dist./Lt.	Relevance
Driver 1	Total distance	50,289	35,500	34,167	34,181	35,559	33,931
	Number of routes	233	159	149	150	150	149
	Capacity utilization	59%	86%	93%	92%	92%	93%
	Average lead time	2.41	4.10	3.81	3.75	3.80	3.76
	# Touted customers	-	-	125	149	112	147
	# Customer visits	815	815	854	863	850	861
	Distance / litre collected	0.03050	0.02153	0.02060	0.02060	0.02146	0.02046
	Imp. wrt. <i>Wait-if-Possible</i>	-41.7%	-	4.3%	4.3%	0.3%	5.0%
	<i>No-Touting</i>	-	29.4%	32.5%	32.5%	29.7%	32.9%
	Computational time	106.7	107.3	123.5	132.3	214.7	230.4
Driver 2	Total distance	37,712	33,387	32,228	32,030	33,150	31,926
	Number of routes	240	217	208	209	208	208
	Capacity utilization	81%	90%	94%	94%	94%	94%
	Average lead time	3.59	4.73	4.22	4.17	4.22	4.16
	# Touted customers	-	-	165	193	168	202
	# Customer visits	879	879	935	945	936	950
	Distance / litre collected	0.01618	0.01432	0.01375	0.01367	0.01414	0.01362
	Imp. wrt. <i>Wait-if-Possible</i>	-13.0%	-	4.0%	4.5%	1.3%	4.9%
	<i>No-Touting</i>	-	11.5%	15.0%	15.5%	12.6%	15.8%
	Computational time	104.2	123.6	134.7	160.1	310.4	322.4

of different touting strategies with those of *No-Touting* and *Wait-if-Possible* strategies and report the improvements with respect to these strategies.

The results reveal that the default rolling horizon planning that solves a VRP each day and sends out a vehicle if there is an open customer order is a rather inefficient strategy. The simple *Wait-if-Possible* strategy that delays sending out the truck until there is either sufficient demand or it has to be sent out because of an imminent due date is dramatically more efficient (29.4% for Driver 1, 11.5% for Driver 2). This makes sense because occasional waiting means the pool of waiting orders is larger, allowing to construct more efficient tours. On top of this, touting is able to further improve efficiency by around 4.5%. The differences between the touting strategies are relatively small, except a generally poor performance of the distance per litre priority rule together with re-optimization. Without re-optimization, the distance per litre and relevance heuristics perform comparable for Driver 1, while the relevance heuristic is substantially better for Driver 2. We thus conclude that overall, the relevance measure with re-optimization is the most sensible choice.

Further observations complement the analysis. The number of tours is the greatest for *No-Touting*, since the truck is always sent out if there is an open customer order. The policy *Wait-if-Possible* follows *No-Touting* and the touting strategies have very close numbers for both drivers. As a result, the utilization of the vehicles is higher when touting is performed. One may argue that touting will not only be more efficient and cost effective, but also has advantages for the drivers, as the number of days a driver goes out to visit customers is smaller. This may create an opportunity for training for the drivers, or simply to take some days off.

Average lead time is the shortest in *No-Touting* since all customers are serviced as soon as possible. On the other hand, *Wait-if-Possible* has the longest average lead time, as the customers may be kept unserved until their due dates if there are few orders to fill the vehicle. Touting strategies lead to smaller average lead times than those in *Wait-if-Possible* case, as the touted customers are served one day after they are approached, which provides the smallest possible lead time.

When the routes are re-optimized after the addition of the touted customers, the existing customers may be redistributed to other routes such that the total distance is lower compared to the

initial solution. Hence, a better solution is obtained for this set of customers. However, as new customers arrive on subsequent days, the previously shifted customers may prevent obtaining a good solution. It may be even harder to tout more customers. This is observed when the customers to tout are selected using the distance per litre criterion, which is a myopic approach. However, when the relevance measure is used as criterion, then re-optimisation is beneficial in the long-term and further increases the efficiency. Note that this benefit is achieved at the expense of increased computational times (74% and 101% higher compared to the no re-optimization strategy for Driver 1 and 2). It is also observed that the number of customer visits during the planning horizon is higher when touting is used, but this is outweighed by the more efficient routes.

Another perspective from the customers' standpoint is that they are served slightly more often with touting (see Fig. 6). For the instance for Driver 1, the average number of days between two consecutive collections is 71 and 74 when touting is used and not used, respectively. These values are 69 and 73 for the instance for Driver 2. These differences are small and are not expected to make a noticeable difference in the experience of the customers. On the other hand, for the instances for Driver 1 and Driver 2, when touting is performed, 14% and 18% of the time, the customers are served even before they need to place an order via touting, whereas when touting is not performed, they need to place the orders by themselves all the time. According to our partner company, some customers do not always pay attention to how much waste has been collected in their containers and sometimes this leads to overflowing. This is not only environmentally costly, but also requires that an urgent collection must be arranged for this customer. Although we have not included these kinds of uncertainties in our simulation, customers are probably happy to be approached and have their waste collected before an overflow happens.

We perform additional experiments using the *Tout using Relevance Measures & Re-Optimize* strategy by removing the probability that a touted customer accepts the offer, i.e., they always accept being served on the next day. For both drivers, as expected, the numbers of touted customers and overall customer visits increase and the average lead time as well as the number of days between two consecutive visits decreases. However, vehicle utilization and number of routes remain unchanged. For the distance traveled per

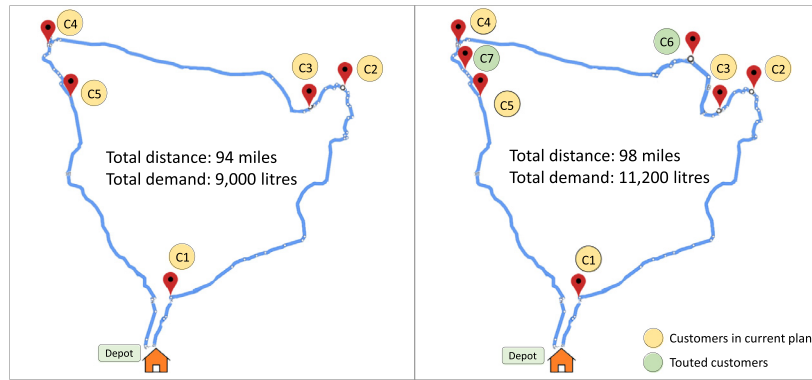


Fig. 7. Route plans before and after touting potential customers.

Table 5
Constructive heuristic results of different strategies.

Strategies		Wait- / No Re-Optimization / Re-Optimization					
		No-Touting	if-Possible	No Re-Optimization		Re-Optimization	
				Dist./Lt.	Relevance	Dist./Lt.	Relevance
Driver 1	Total distance	53,342	40,233	39,931	39,873	39,432	39,024
	Number of routes	234	148	146	146	146	146
	Capacity utilization	59%	93%	94%	94%	94%	94%
	Average lead time	2.16	4.46	4.37	4.35	4.30	4.25
	# Touted customers	-	-	31	32	46	53
	# Customer visits	815	815	828	829	831	836
	Distance / litre collected	0.03235	0.02440	0.02411	0.02410	0.02384	0.02357
	Imp. wrt. Wait-if-Possible	-32.6%	-	1.2%	1.2%	2.3%	3.4%
	Imp. wrt. No-Touting	-	24.6%	25.5%	25.5%	26.3%	27.1%
	Computational time	0.03	0.47	0.52	0.53	1.25	1.32
Driver 2	Total distance	46,378	41,494	40,713	40,508	39,459	38,848
	Number of routes	239	207	205	205	205	205
	Capacity utilization	81%	94%	95%	95%	95%	95%
	Average lead time	2.83	5.02	4.78	4.76	4.73	4.65
	# Touted customers	-	-	73	82	94	112
	# Customer visits	879	879	908	912	913	921
	Distance / litre collected	0.01989	0.01780	0.01740	0.01732	0.01686	0.01659
	Imp. wrt. Wait-if-Possible	-11.8%	-	2.2%	2.7%	5.3%	6.8%
	Imp. wrt. No-Touting	-	10.5%	12.5%	13.0%	15.2%	16.6%
	Computational time	0.02	0.91	1.08	1.13	1.67	1.76

litre of waste collected, although it has slightly increased and decreased for Driver 1 and 2, respectively, paired t-tests show that the differences are not significant.

Figure 7 illustrates a simulation day in which touting is performed using *Tout using Relevance Measures & Re-Optimize* strategy. The depot is represented by an orange house and the customer set in the current plan involves customers 1–5. The planned route for the next day is depicted on the left with a total distance of 94 miles and 9000 litres of waste collected. Customers 6 and 7, which are shown in the right figure, with demands of 1200 and 1000 litres of waste are found as potential customers. Because their addition is feasible in terms of capacity and time, they are touted and included in the current plan. As a result, the vehicle could collect 2200 litres of additional waste with only 4 miles of detour. Distance traveled per litre collected has decreased from 0.0104 to 0.00875.

Finally, in order to demonstrate that the benefits of touting are independent of the algorithm used to solve the routing problem, we conduct another set of simulation experiments. Here, we only use the construction heuristic to obtain a solution, we do not improve it by means of LNS. The details of this heuristic are outlined in Algorithm 1. Similarly, in *Tout & Re-Optimize* strategies, while re-optimizing the solution after considering a touted customer, only the initial solution is constructed, it is not further im-

proved. Table 5 presents the results for Drivers 1 and 2. The relative findings are analogous to those obtained with the more sophisticated heuristic.

Since these results are obtained by using only the initial solution construction heuristic, it is expected that re-optimization results in better solutions. Compared to using the LNS heuristic, of course, the solution quality is on average 13.3% and 23.6% worse in terms of total distance per litre collected for Driver 1 and Driver 2, respectively. In addition, average lead times are longer in this approach, except for the *no-touting* strategy. This is due to the nature of the construction heuristic, which tries to fill the next day's vehicle. Since more jobs are scheduled in the first vehicle, the average lead time decreases. While the relative performances of different strategies are similar to that of the heuristic with improvement component, *Tout using Distance per Litre & Re-Optimize* strategy provides a better result compared to the strategies where re-optimization is not used. Overall, these results demonstrate that touting relevant customers improves the efficiency in the long-run, independent of the method to solve the VRP.

5.4.4. Experiments on additional instances

As our industrial partner supplied the data of two drivers that are discussed above, to further validate the performance of the proposed strategies, we conduct additional experiments on in-

Table 6
Results of different strategies for C101 instance.

Strategies	No-Touting	Wait- if-Possible	No Re-Optimization		Re-Optimization		
			Dist./Lt.	Relevance	Dist./Lt.	Relevance	
C101	Total distance	21,869	16,282	16,246	16,264	16,364	16,133
	Number of routes	224	165	161	161	161	161
	Capacity utilization	67.9%	92.5%	95.7%	95.6%	95.7%	95.9%
	# Touted customers	-	-	47	49	46	52
	# Customer visits	609	609	633	636	634	637
	Distance / litre collected	1.79440	1.33382	1.32003	1.32044	1.32779	1.30905
	Imp. wrt. Wait-if-Possible	-34.5%	-	1.0%	1.0%	0.5%	1.9%
	Imp. wrt. No-Touting	0.0%	25.7%	26.4%	26.4%	26.0%	27.0%
R101	Total distance	21,339	15,542	15,515	15,502	15,598	15,487
	Number of routes	232	172	168	169	168	168
	Capacity utilization	69.8%	94.3%	96.9%	96.7%	97.0%	97.0%
	# Touted customers	-	-	45	47	42	51
	# Customer visits	789	789	814	815	812	817
	Distance / litre collected	1.65046	1.19930	1.18937	1.18784	1.19503	1.18605
	Imp. wrt. Wait-if-Possible	-37.6%	-	0.8%	1.0%	0.4%	1.1%
	Imp. wrt. No-Touting	0.0%	27.3%	27.9%	28.0%	27.6%	28.1%
RC101	Total distance	26,492	19,316	19,287	19,186	19,342	19,194
	Number of routes	230	169	165	165	165	165
	Capacity utilization	68.5%	93.4%	96.1%	96.0%	96.1%	96.3%
	# Touted customers	-	-	39	40	37	45
	# Customer visits	724	724	747	747	745	750
	Distance / litre collected	2.10621	1.53282	1.51796	1.51068	1.52220	1.50967
	Imp. wrt. Wait-if-Possible	-37.4%	-	1.0%	1.4%	0.7%	1.5%
	Imp. wrt. No-Touting	0.0%	27.2%	27.9%	28.3%	27.7%	28.3%

stances we generate based on the well-known VRPTW instances of Solomon (Solomon, 1987). We select one instance from each problem category, where the customers' locations are randomly distributed across the network (R), or they are clustered (C) or half of them is clustered and the other half is randomly distributed (RC), specifically we use instances C101, R101, and RC101. As the data is designed for a static VRP, we make some additions to adapt it to our problem. First, we assume that a customer's demand value in the Solomon instance is the tank capacity of the customer in the revised instance for our problem. Then to simulate the random accumulation process, we generate mean and variance values based on one of the drivers' data. Using the relationship between the capacity and the mean and the variance in that data, we generate the mean and variance values in Solomon's instances.¹ Then we set the maximum lead time to 10 days to obtain feasible schedules. We apply the simulation process as discussed in Section 5.3 and the average results of 100 simulation runs are presented in Table 6

We observe that the results for all three spatial distributions are consistent with the ones obtained using the company's dataset. Touting strategies bring substantial improvement relative to the No-Touting strategy, and using the relevance measure works better than using distance/litre.

There are also improvements with respect to Wait-if-Possible strategy with again Touting & Re-Optimize using Relevance Measures bringing the highest improvement.

6. Conclusion

We have presented a dynamic multi-period vehicle routing problem as faced by a UK waste collection company. It is solved on a rolling-horizon with a Large Neighborhood Search to solve each individual VRP. The paper introduces the idea to integrate demand management into the tour planning via a method called *touting*.

¹ This dataset is included in the Mendeley Data (Keskin, Branke, Deineko and Strauss, 2023).

Touting consists of contacting a customer expected to order a collection soon, and nudge them to place their order now. While this means smaller amounts are collected and the customers need to be visited more frequently, it opens up the opportunity for the company to visit a customer when they are nearby anyway, potentially reducing the overall distance traveled.

We have proposed different strategies for touting potential future customers. Using real world data from waste collection industry, we have shown that touting appropriate future customers on relevant days may save a considerable amount of fuel due to decreased total distance traveled per litre of waste collected. Furthermore, the number of tours required is shown to be smaller if the touting strategies are followed, which may provide additional monetary benefits from the vehicle acquisition costs. Although these advantages are shown using a real-world waste collection problem, the idea of exploiting knowledge of demand and active management of that demand in the form of touting can be applied to many routing problems in which customers arrive dynamically, including maintenance routing - as visiting the customers a few days before their deadline would not make much difference - and routing random and subscribed customers, for example Amazon could contact subscribed customers to offer service before the scheduled dates if their routes are nearby customers on the tour.

This research opens up several directions for future research. First, in this study, the decisions are made at the end of each time period. Future work may attempt to make decisions more dynamically within the day, and the routes of the vehicles may be updated according to the new information, i.e., each time a new order is received, or when a customer accepts the touting offer. This will make not only routing, but also touting dynamic, i.e., customers to be touted change according to the new set of customers and the current tour. Second, we have considered a single-vehicle routing problem as our industry partner uses fixed service regions per driver for operational reasons. While we anticipate that results carry over to VRPs with multiple vehicles, this should be demonstrated. Third, the most successful strategy for touting was the proposed relevance measure that is a linear combination of three criteria with equal weighting. Additional criteria such as the

customer’s own anticipated demand could be integrated, and the weighting of the criteria could be more sophisticated. Finally, another future work may include incentivising customers by offering discounts at the time of touting to increase the probability that they accept the offer.

Acknowledgment

We thank our industrial partner Slicker Recycling for their valuable insights and for providing real-world data. We gratefully acknowledge funding by Slicker Recycling and Innovate UK through grant KTP11213.

Appendix A. Mathematical model for the static VRP at each time period

Here, we present the mathematical formulation for the routing subproblem on a day within the planning horizon. $H = \{1, 2, \dots, l\}$ stands for the planning horizon for the day the problem is solved and is a subset of the whole horizon. Let $C = \{1, \dots, n\}$ be the set of known customers, which have not yet been serviced at the beginning of the planning day. A depot is located at vertex 0. d_{ij} and t_{ij} stand for the distance and travel time between vertices i and j . The vehicle has a load capacity L . Let the amount of goods to be collected from customer i , the due date of the order, and the time spent while collecting the goods at the customer site be q_i , d_i , and s_i respectively. Hence, the last period in the planning horizon, l , is determined by the latest due date of the customers. Each day $t \in H$ has a time limit, i.e., the vehicle should be back at the depot before T_{max}^t . Furthermore, there is a limit T_{max} on the total time spent over the planning horizon, e.g., weekly working time limit according to the hours of service regulations. The decision variable x_{ijt} is 1 if and only if arc (i, j) is traversed on day t . u_{ijt} and τ_t track the tank load of the vehicle upon traversing arc (i, j) , and total time spent on day t , respectively. The mathematical model is formulated as follows:

$$\text{minimize } \sum_{t \in H} \sum_{i, j \in C} d_{ij} x_{ijt} \tag{A.1}$$

subject to

$$\sum_{j \in C} \sum_{t=1}^{d_i} x_{ijt} = 1 \quad i \in C \tag{A.2}$$

$$\sum_{j \in C} x_{ijt} = \sum_{j \in C} x_{jit} \quad i \in C, t \in \{1 \dots d_i\} \tag{A.3}$$

$$\sum_{j \in C} u_{ijt} - \sum_{j \in C} u_{jit} = q_i \sum_{j \in C} x_{ijt} \quad i \in C, t \in \{1 \dots d_i\} \tag{A.4}$$

$$u_{ijt} \leq L x_{ijt} \quad i, j \in C, t \in H \tag{A.5}$$

$$\sum_{j \in C} u_{j0t} - \sum_{j \in C} u_{0jt} = \sum_{i, j \in C} q_i x_{ijt} \quad t \in H \tag{A.6}$$

$$\sum_{i, j \in C} (t_{ij} + s_i) x_{ijt} \leq \tau_t \quad t \in H \tag{A.7}$$

$$\tau_t \leq T_{max}^t \quad t \in H \tag{A.8}$$

$$\sum_{t \in H} \tau_t \leq T_{max} \tag{A.9}$$

$$x_{ijt} \in \{0, 1\} \quad i, j \in C, t \in H \tag{A.10}$$

$$u_{ijt}, \tau_t \geq 0 \quad i, j \in C, t \in H \tag{A.11}$$

The objective function, (A.1) minimizes the total distance traveled over the remaining planning periods. Constraints (A.2) ensure that each customer is serviced in one of the periods until their due date, whereas Constraints (A.3) establish the flow conservation. Constraints (A.4) - (A.6) track the load of the vehicle and ensure that the tank capacity is not exceeded. Constraints (A.7) calculate arrival time at the depot after completion of the route and Constraints (A.8) make sure that it does not exceed the time limit in each period. Total time limit over all periods is satisfied by Constraint (A.9). Finally, (A.10) - (A.11) define domains of the decision variables.

Appendix B. Framework for solving the vehicle routing problem

In this appendix, the details of the algorithm, which we use to solve the multi-period vehicle routing problem throughout the simulation, are presented in Algorithm 2, where x stands for a so-

Algorithm 2 An LNS Framework to Solve Multi-Period VRP.

```

1: Construct an initial solution,  $x_{init}$ 
2:  $x_{best}, x_{previous} \leftarrow x_{init}$ 
3: while  $iter \leq$  maximum number of iterations do
4:   Select a removal rule randomly
5:   Determine  $q\%$  of the scheduled customers according to the removal rule
6:   Remove the determined customers from the current solution
7:   Select an insertion rule randomly
8:   Insert the removed customers into the solution according to the insertion rule, obtain a new solution,  $x_{current}$ 
9:   if  $f(x_{current}) < f(x_{previous})$  then
10:      $x_{previous} \leftarrow x_{current}$ 
11:   end if
12:   if  $f(x_{current}) < f(x_{best})$  then
13:      $x_{best} \leftarrow x_{current}$ 
14:   end if
15: end while

```

lution, which is a set of routes belonging to different time periods and executed by the same vehicle, while $f(x)$ stands for the total length of the routes in solution x . $x_{current}$, $x_{previous}$, and x_{best} correspond to the solution obtained in the current iteration, the incumbent solution the algorithm has had in the previous iteration, and the overall best solution found until that iteration, respectively.

Appendix C. Comparison of solutions on a rolling horizon

This appendix illustrates the fact that for the problems where the customers arrive dynamically over time, the optimal solution obtained using the customers known on one day may not be necessarily optimal when applied on a rolling horizon. Let us consider the example in Fig. 2 without customer 4 to make the calculations simpler. On the first day, customers 1, 2, 3, and 5 are known, and they are optimally served as in Fig. C.8-a. Total distance of this tour is 19.12, which is calculated assuming unit length for each arc and Euclidean distances. Then the next day three more customers arrive and are served as shown in Fig. C.8-b. Total distance of this tour is 16.45. On the contrary, if only 3 customers are served on the first day as shown in Fig. C.8-c, and one of them is left to the following day, then the total distances traveled on these two days would become 12.80 and 16.53, respectively, which decreases the total distance of two days’ tours from 35.57 to 29.33.

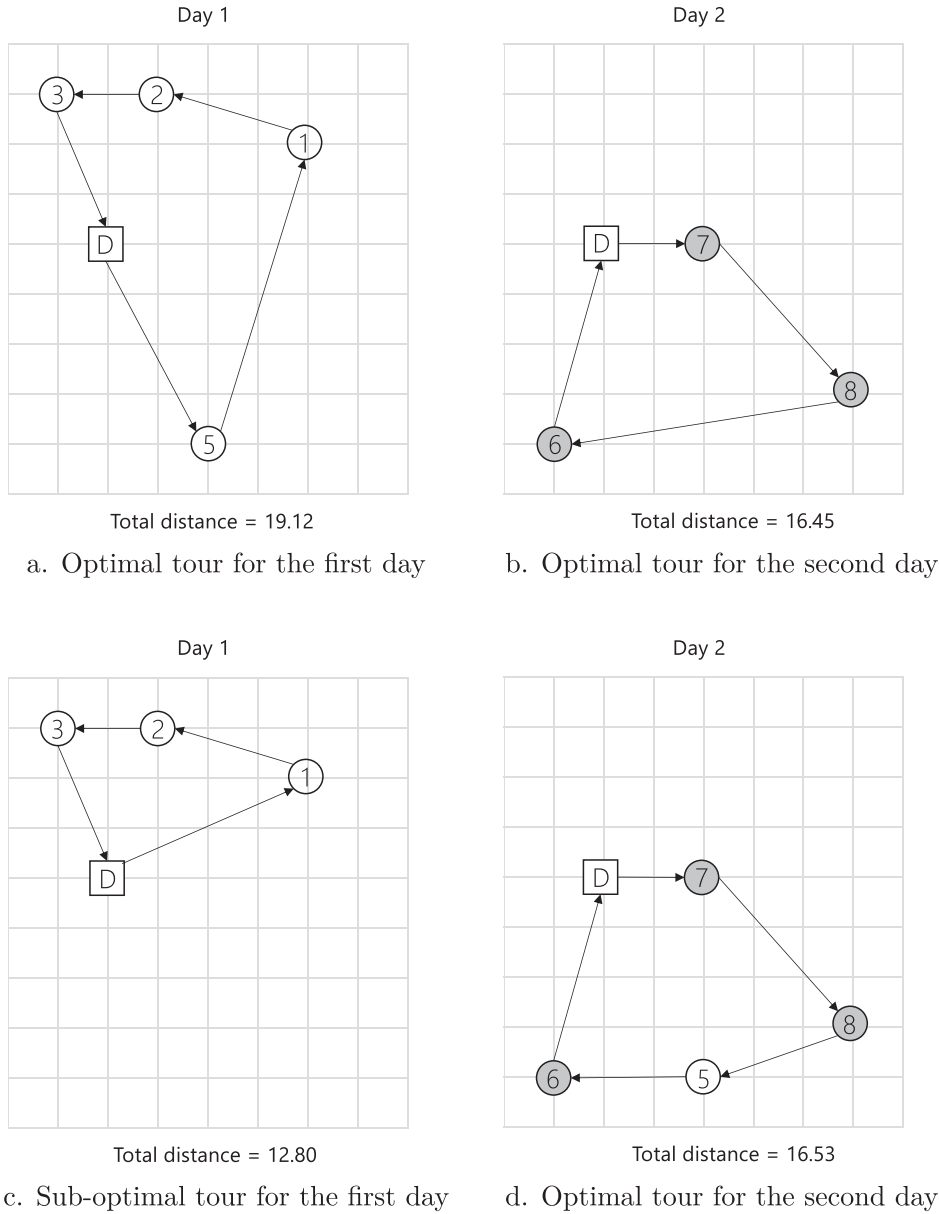


Fig. C.8. Illustration of different solutions on a rolling horizon setting.

Appendix D. Markov decision process representation

This appendix provides a dynamic programming formulation of the problem. To that end, let us define the state, decisions, rewards, exogenous information and the state transitions of the Markov Decision Process that underpins our problem. With this terminology, we then can formally state its dynamic programming formulation.

D1. Stages, states and decisions

We consider a finite planning horizon with a discrete time grid of decision epochs (stages). Let us assume that the customer arrivals occur within the day exogenously and this random information regarding demand collected over the day becomes available at the end of the day. Let us further assume that touting takes place in the next m epochs after the customer arrivals are complete. Then the decision space consists of decision epochs $k \in \{e_1^1, e_2^1, \dots, e_m^1, e_{m+1}^1, e_1^2, e_2^2, \dots, e_m^2, e_{m+1}^2, \dots, K\}$, where e_b^a is the b^{th} decision epoch to determine which customer to tout on day a , and

finally there is a decision epoch to decide the tours after there are no further customers to tout. On the first day, we observe demand that has arrived over the day in one batch at the beginning of period e_1^1 , and proceed on this basis with sequential touting decisions. The next m epochs represent the time interval over which touting decisions are being made and we observe the customers' responses to them; and the routing decision takes place in the last epoch, e_{m+1}^1 . Therefore, we decide whether to tout a customer (and if so, whom – expressed by a decision variable x_k^t), and then observe whether the touted customer responds positively to the offer. At the respective last epoch e_{m+1}^a , for each day a , we observe no further new information but schedule customers to be served on the next day based on all requests (including touted ones) received to date, expressed by a decision vector x_k^t .

At stage k , the system is in a state $S_k = (R_k = (q_k^c, d_k^c, h_k^c)_{c \in C}, T_{a(k)})$, where C stands for the set of customers. The array R_k contains all yet unscheduled customer requests, consisting of q_k^c as the requested quantity to be collected from customer c , with $q_k^c > 0$ if there is an outstanding order, and

0 otherwise. Next, d_k^c corresponds to the deadline of the collection request if the order is outstanding, or has no meaning if $q_k^c = 0$. Note that this notation allows us to also handle the case where orders remain unfulfilled beyond their deadline. The array h_k^c represents customer c 's collection history, consisting of collection dates and collection amounts, which we use to predict the fill level at each customer and subsequently decide whether a customer would be touted or not. Recall that we assume a finite and known customer population C ; thus their locations are not changing from one state to another and we do not need to explicitly carry the locations of the customers in the state variable. Finally, $T_{a(k)}$ represents the set of touted customers on the corresponding day $a(k)$ to which period k belongs, up until period k . We need to keep track of touted customers to prevent contacting them again. Of course, this information is also added to the history arrays h , but separating out which customers have been touted already makes the model easier to read.

D2. Rewards, exogenous information and state transitions

Our overall objective is to minimize the expected routing cost across all days in the planning horizon. Routing costs are incurred only at the end-of-day stages $k \in E = \{e_{m+1}^a : a \in A\}$, where A is the set of all days in the planning horizon. These costs, denoted by $f(S_k, x_k^r)$, pertain to the final schedule for the next day given requests in state S_k and routing decision x_k^r . During the day we only collect orders and this exogenous information becomes available at the beginning of each day in period e_1^a of day a in the form of an array R_k^{new} that contains quantities and deadlines of all new requests. We assume that the underpinning distribution is known such that the expectation over R_k^{new} is well-defined. For the first decision epoch $k = e_1^a$ of day a , we start in state $S_k = (R_{k-1} \cup R_k^{\text{new}}, T_a = \emptyset)$.

Subsequently, in each decision epoch e_i^a , $i \in \{2, \dots, m\}$, we decide on which customer to tout (if any), and then observe whether the touted customer will accept the offer as well as what quantity to be collected. If customer c accepts the touting offer, then its information is added to the state as $R_k^c = (q_k^c, d_k^c)$; otherwise, if customer c does not accept or if we do not tout in the first place, then $R_k^c = \emptyset$ for all $c \in C$. The set of customers touted so far, T_a , is expanded such that it includes customer c and this new demand information is added to the request history of customer c , h^c . We assume that a touted request must always be scheduled for the next day. Therefore, for the touted customers who have accepted collection, d_k^c is set to the next day. The state transition is written as $R_{k+1} = R_k \cup R_k^c$.

At the end of a day ($k \in E$), we schedule to serve certain customer requests for the next day, which then are removed from the pool of orders to be served by setting the corresponding indicator $q_k^c = 0$, for all scheduled customers $c \in C$, and by adding the collection information to the corresponding collection history h^c . Assuming $R(c_k)$ denotes the customers scheduled for the next day, we use the shorthand notation $R_{k+1} = R_k \setminus R(c_k)$ to represent this transition.

D3. Dynamic programming formulation

With this notation, we now can express the problem as a dynamic program. Let $V_k(S_k)$ be the value function at stage k and state S_k ; more specifically, it is the minimal expected cost from stage k until the end of the time horizon K and is given by:

$$V_k(S_k) = \begin{cases} \mathbb{E}_{R_k^{\text{new}}} [V_{k+1}(R_k \cup R_k^{\text{new}}, T_a = \emptyset)] & k = e_1^a \quad \forall a \in A, \\ \min_{c \in C \setminus \{c' \in C : q_{c'} > 0\} \cup T_a \cup \{0\}} \mathbb{E}_{R_k^c} [V_{k+1}(R_k \cup R_k^c, T_a \cup \{c\})] & k \in \{e_2^a, \dots, e_m^a\} \quad \forall a \in A, \\ \min_{x_k^r} f(S_k, x_k^r) + V_{k+1}(S_k \setminus S(x_k^r)) & \forall k \in E. \end{cases}$$

which is a combination of different cases depending on the type of the state. The first case corresponds to the initial states at the end of a day, after the customer arrivals are complete and the information on these new customers is available. The second case is for the states where touting is performed, i.e., the customers to tout are determined and their responses to the touting offers are observed. The minimization is over the touted customers, however it also includes the case in which we do not tout any customers, which is represented by 0. Finally, the last case stands for the last states of a day, in which the routing decision is taken based on the available information on unscheduled customers, and routing costs based on the function $f(S_k, x_k^r)$ are incurred. The boundary condition is given by $V_{K+1}(S_k) = 0$ for all states S_k (since we simply assume that there are no further orders coming in on the final day in the planning horizon).

Clearly, this dynamic program is intractable due to its large state space and the fact that it involves vehicle routing problems in each stage $k \in E$. Therefore, in Section 4, we present a heuristic approach for tackling this problem. This heuristic approach constructs routes for currently known customers, then uses different ways to decide which (if any) customer to tout. In the Markov decision model, routing decisions are made only in the final periods $k \in E$ since it is not possible to calculate the value function of the DP formulation in reasonable time for realistic instances of the problem, whereas in our heuristic we progressively add touted customers to the routes.

References

Agatz, N., Campbell, A., Fleischmann, M., & Savelsbergh, M. (2011). Time slot management in attended home delivery. *Transportation Science*, 45(3), 435–449.

Agatz, N., Fan, Y., & Stam, D. (2021). The impact of green labels on time slot choice and operational sustainability. *Production and Operations Management*, 30(7), 2285–2303.

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

Angelelli, E., Bianchessi, N., Mansini, R., & Speranza, M. G. (2009). Short term strategies for a dynamic multi-period routing problem. *Transportation Research Part C: Emerging Technologies*, 17(2), 106–119.

Angelelli, E., Grazia Speranza, M., & Savelsbergh, M. W. (2007). Competitive analysis for dynamic multiperiod uncapacitated routing problems. *Networks: An International Journal*, 49(4), 308–317.

Archetti, C., Jabali, O., & Speranza, M. G. (2015). Multi-period vehicle routing problem with due dates. *Computers & Operations Research*, 61, 122–134.

Athanasopoulos, T., & Minis, I. (2013). Efficient techniques for the multi-period vehicle routing problem with time windows within a branch and price framework. *Annals of Operations Research*, 206(1), 1–22.

Billing, C., Jaehn, F., & Wensing, T. (2018). A multiperiod auto-carrier transportation problem with probabilistic future demands. *Journal of Business Economics*, 88(7–8), 1009–1028.

Braekers, K., Ramaekers, K., & Van Nieuwenhuysse, I. (2016). The vehicle routing problem: State of the art classification and review. *Computers & Industrial Engineering*, 99, 300–313.

Brinkmann, J., Ulmer, M. W., & Mattfeld, D. C. (2020). The multi-vehicle stochastic-dynamic inventory routing problem for bike sharing systems. *Business Research*, 13(1), 69–92.

Coelho, L. C., Cordeau, J.-F., & Laporte, G. (2014). Heuristics for dynamic and stochastic inventory-routing. *Computers & Operations Research*, 52, 55–67.

Cordeau, J.-F., Dell'Amico, M., Falavigna, S., & Iori, M. (2015). A rolling horizon algorithm for auto-carrier transportation. *Transportation Research Part B: Methodological*, 76, 68–80.

Dayarian, I., Crainic, T. G., Gendreau, M., & Rei, W. (2015). A branch-and-price approach for a multi-period vehicle routing problem. *Computers & Operations Research*, 55, 167–184.

Dayarian, I., Crainic, T. G., Gendreau, M., & Rei, W. (2016). An adaptive large-neighborhood search heuristic for a multi-period vehicle routing problem. *Transportation Research Part E: Logistics and Transportation Review*, 95, 95–123.

De Bruecker, P., Beliën, J., De Boeck, L., De Jaeger, S., & Demeulemeester, E. (2018). A model enhancement approach for optimizing the integrated shift scheduling and vehicle routing problem in waste collection. *European Journal of Operational Research*, 266(1), 278–290.

Estrada-Moreno, A., Savelsbergh, M., Juan, A. A., & Panadero, J. (2019). Biased-randomized iterated local search for a multiperiod vehicle routing problem with price discounts for delivery flexibility. *International Transactions in Operational Research*, 26(4), 1293–1314.

- Ferrer, J., & Alba, E. (2019). Bin-ct: Urban waste collection based on predicting the container fill level. *Bio Systems*, 186, 103962.
- Huang, S.-H., & Lin, P.-C. (2010). A modified ant colony optimization algorithm for multi-item inventory routing problems with demand uncertainty. *Transportation Research Part E: Logistics and Transportation Review*, 46(5), 598–611.
- Jaillet, P., Bard, J. F., Huang, L., & Dror, M. (2002). Delivery cost approximations for inventory routing problems in a rolling horizon framework. *Transportation Science*, 36(3), 292–300.
- Keskin, M., Branke, J., Deineko, V., & Strauss, A.K. (2023). Dataset for a Dynamic Multi-Period Vehicle Routing Problem. doi:10.17632/cbkzp5b8hb.1.
- Ketzenberg, M. E., & Metters, R. D. (2020). Adapting operations to new information technology: A failed internet of things application. *Omega*, 92, 102152.
- Larrein, H., Coelho, L. C., Archetti, C., & Speranza, M. G. (2019). Exact solution methods for the multi-period vehicle routing problem with due dates. *Computers & Operations Research*, 110, 148–158.
- Liu, R., Tao, Y., & Xie, X. (2019). An adaptive large neighborhood search heuristic for the vehicle routing problem with time windows and synchronized visits. *Computers & Operations Research*, 101, 250–262.
- Markov, I., Bierlaire, M., Cordeau, J.-F., Maknoon, Y., & Varone, S. (2018). A unified framework for rich routing problems with stochastic demands. *Transportation Research Part B: Methodological*, 114, 213–240.
- Markov, I., Bierlaire, M., Cordeau, J.-F., Maknoon, Y., & Varone, S. (2020). Waste collection inventory routing with non-stationary stochastic demands. *Computers & Operations Research*, 113, 104798.
- Ninikas, G., Athanassopoulos, T., Zaimpekis, V., & Minis, I. (2014). Integrated planning in hybrid courier operations. *The International Journal of Logistics Management*, 25(3), 611–634.
- Ozbaygin, G., & Savelsbergh, M. (2019). An iterative re-optimization framework for the dynamic vehicle routing problem with roaming delivery locations. *Transportation Research Part B: Methodological*, 128, 207–235.
- Pillac, V., Gendreau, M., Guéret, C., & Medaglia, A. L. (2013). A review of dynamic vehicle routing problems. *European Journal of Operational Research*, 225(1), 1–11.
- Powell, W. B., Towns, M. T., & Marar, A. (2000). On the value of optimal myopic solutions for dynamic routing and scheduling problems in the presence of user noncompliance. *Transportation Science*, 34(1), 67–85.
- Ritzinger, U., Puchinger, J., & Hartl, R. F. (2016). A survey on dynamic and stochastic vehicle routing problems. *International Journal of Production Research*, 54(1), 215–231.
- Ropke, S., & Pisinger, D. (2006). An adaptive large neighborhood search heuristic for the pickup and delivery problem with time windows. *Transportation Science*, 40(4), 455–472.
- Soeffker, N., Ulmer, M. W., & Mattfeld, D. C. (2022). Stochastic dynamic vehicle routing in the light of prescriptive analytics: A review. *European Journal of Operational Research*, 298(3), 801–820.
- Solomon, M. M. (1987). Algorithms for the vehicle routing and scheduling problems with time window constraints. *Operations Research*, 35(2), 254–265.
- Subramanyam, A., Mufalli, F., Laínez-Aguirre, J. M., Pinto, J. M., & Gounaris, C. E. (2021). Robust multiperiod vehicle routing under customer order uncertainty. *Operations Research*, 69(1), 30–60.
- Ulmer, M. W., Goodson, J. C., Mattfeld, D. C., & Thomas, B. W. (2020). On modeling stochastic dynamic vehicle routing problems. *EURO Journal on Transportation and Logistics*, 9(2), 100008.
- Ulmer, M. W., Soeffker, N., & Mattfeld, D. C. (2018). Value function approximation for dynamic multi-period vehicle routing. *European Journal of Operational Research*, 269(3), 883–899.
- Wen, M., Cordeau, J.-F., Laporte, G., & Larsen, J. (2010). The dynamic multi-period vehicle routing problem. *Computers & Operations Research*, 37(9), 1615–1623.
- Yang, X., & Strauss, A. K. (2017). An approximate dynamic programming approach to attended home delivery management. *European Journal of Operational Research*, 263(3), 935–945.
- Yang, X., Strauss, A. K., Currie, C., & Eglese, R. (2016). Choice-based demand management and vehicle routing in e-fulfillment. *Transportation Science*, 50(2), 473–488.
- Yildiz, B., & Savelsbergh, M. (2020). Pricing for delivery time flexibility. *Transportation Research Part B: Methodological*, 133, 230–256.