Transportation Research Part E 142 (2020) 102049

Contents lists available at ScienceDirect

Transportation Research Part E

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

Optimization and analysis of a robot-assisted last mile delivery system

Michele D. Simoni[a,c,*], Erhan Kutanoglu[b], Christian G. Claudel[a]

[a] Dept. of Civil, Architectural, and Environmental Engineering, University of Texas at Austin, 301E Dean Keeton, St., Austin 78712, TX, USA
[b] Operations Research and Industrial Engineering, University of Texas at Austin, 204E Dean Keeton, St., Austin 78712, TX, USA
[c] Division of Systems Analysis and Economics, KTH Royal Institute of Technology, Stockholm, Sweden

ARTICLE INFO

Keywords:
Integrated truck-robot delivery
Traveling salesman with robot
City logistics
Last-mile delivery
Local search with adaptive perturbation

ABSTRACT

The last mile of freight distribution is a critical part of the supply chain because of its significant costs and customers' increasing expectations from e-commerce and same-day delivery services. Automated technologies in freight transportation represent an opportunity to develop more efficient systems characterized by the integration of different and complementary modes. In this study, we focus on the possibility of implementing an integrated truck-robot system for the last-mile delivery. This typology of problem shares similarities with truck-drone problems, although robots are characterized by much slower speeds and can perform several consecutive deliveries. Based on these particular features, a heuristic that efficiently identifies solutions based on initial truck tours and corresponding joint robot operations is presented. This solution approach leverages a special version of the "Weighted Interval Scheduling Problem," which allows for a very efficient Dynamic Programming solution. The developed solution approach is adopted to analyze the influence on efficiency of different features concerning the robot's design and operation, and the surrounding environment. The results show that robot-assisted last-mile delivery systems are quite efficient if robots are employed in heavily congested areas and appropriately retrofitted to accommodate several compartments in the robot's storage.

1. Introduction

The last mile of freight distribution represents the weakest link of the supply chain and it is a source of considerable congestion and pollution externalities (Rodrigue et al., 2009). The future increase of urbanization and the expansion of e-commerce will bring additional pressure for the development of effective, innovative "City Logistics" solutions (Savelsbergh and Van Woensel, 2016).

Technological advances in automation offer the opportunity to develop newer, more sustainable, and more efficient delivery systems. In particular, recent technological progress in automated vehicles (e.g. driverless vehicles, robots, unmanned aerial vehicles) is creating the ground for the development of innovative delivery models that could transform the landscape of last-mile delivery.

Automated modes, such as drones and robots, are transitioning from a purely conceptual phase to the actual prototyping and testing driven by key players in high-tech services and delivery market. Amazon and Google (Halzack, 2016; Nicas, 2018), and DHL and UPS (Hern, 2014; Desjardins, 2018), are all working on the possibility of replacing the more costly traditional truck-based delivery process with drones in areas characterized by low accessibility (due to geographical constraints) or long delivery times (due to traffic). In order to compensate fuel efficiency and size issues, drones can operate as an assisting mode to traditional freight

* Corresponding author.
E-mail address: micheles@kth.se (M.D. Simoni).

https://doi.org/10.1016/j.tre.2020.102049
Received 30 May 2019; Received in revised form 27 June 2020; Accepted 26 July 2020
1366-5545/

Fig. 1. Starship delivery robot. (Source: Mercedes-Benz, n.d.)

distribution modes and depart and return to delivery trucks (Joerss et al., 2016). Robots represent an interesting solution as well, particularly in urban environments characterized by high densities of stops and relatively short delivery distances. Robots' storage could be divided into different compartments so that each trip could serve more than a single customer at a time (although this feature is not currently present in real-world systems) and compensate for their lower speeds. Different companies are already testing robots for mail delivery (Bishop, 2016) and food deliveries (Burns, 2016; Coldewey, 2019). In 2016, the truck manufacturer Mercedes-Benz Vans started a partnership with Starship Technologies, the maker of six-wheeled delivery robots, to develop an integrated delivery service with trucks and robots (Daimler, 2017), as shown in Fig. 1. Here, the truck ferries robots that can swarm around the neighbourhood to perform the last few hundred meters of delivery. Robots' storage is locked and can be opened only by the customers with a code. For safety reasons, the robots, which move on sidewalks, travel at pedestrian speeds. Given the limited service range (2 miles) and load (about 20 kg), such integrated service seems particularly suitable for the last-mile distribution of small-sized items like parcels, groceries and food. In 2019, Amazon too has officially started field-testing its new robot-based delivery service, "Amazon Scout," for same-day package deliveries (Scott, 2019). It is unclear, though, whether these robots would be employed in connection with trucks or independently.

Although different companies are already considering the employment of delivery robots and running initial tests, little is known about the potential gains in delivery time or costs of carriers' operations. While several theoretical studies have focused on the efficiency of drones' delivery systems in terms of saved delivery times and costs (Murray and Chu, 2015; Carlsson and Song, 2017; Wang et al., 2017; Ha et al., 2018; Agatz et al. 2018), very few studies have focused on the specific case of robots. Although drones and sidewalk robots can be broadly included in the same category of "automated delivery modes", several differences can be identified. From a design perspective, robots are characterized by having higher capacities (20–30 kg vs 2–5 kg) and a potentially higher number of compartments (which would ultimately allow two or more consecutive stops in their routes). From an operational perspective, robots are characterized by having considerably lower speeds (5–10 kmph vs 50–100 kmph) and longer ranges (5–10 km vs 10–30 km). Based on these features, in the near future, robots seem to be more suitable for deliveries of low-value items (e.g. groceries, mail) in dense urban environments whereas drones could be more appropriate for deliveries of high priority valuable items (e.g. healthcare and fashion products) in remote or rural areas. Finally, robots and drones will probably differ from a regulatory framework as well since the second ones seem to face more safety concerns due to their risk of harming people and infrastructure. For this reason, drones may need additional licenses for the use of airspace. Along these lines, this study investigates the implementation of a coordinated truck-robot delivery service for the last-mile parcel delivery and its possible efficiency gains under different scenarios characterized by different traffic conditions, demand configurations, design and system operational features. The analyses are restricted to relatively small scenarios (neighborhood size with 50 stops) in order to focus on the influence of different factors on reliable scenarios.

The first contribution of this paper consists in the derivation of an optimization algorithm for the development of efficient "mixed" delivery routes where the original truck's route is modified to include robot trips (sub-tours). The problem is an extension of the traditional Traveling Salesman Problem (TSP) and shares similarities with some recent optimization approaches proposed for the drone-assisted delivery problem. The main features of the proposed approach consist in the development of initial truck routes and their partition into optimal robot sub-tours based on the concept of robot operations. The formulation is a special case of "Independent Set Problem" (a NP-hard problem) known as "Weighted Interval Scheduling Problem", which allows for an efficient Dynamic Programming-based solution algorithm. In addition, the proposed approach leverages specific features of the problem, such as the lower speed of the ancillary robot and the sequential nature of the problem to reduce the size of the optimization problem.

The second contribution of this paper consists in a detailed evaluation of the potential benefits of robot-assisted deliveries based on different deployment and environmental features. The influence of factors such as traffic speed and congestion, the capacity of robot's storage, and the length of drop-off operations is evaluated in terms of improved efficiency. The results of the analyses lead to a broader discussion on the opportunities and challenges for the implementation of integrated delivery service with trucks and robots.

In this paper, after a brief presentation of previous research on related problems, we provide a description of the robot-assisted truck delivery problem. We describe our proposed formulation and solution approach, and we investigate their performance. In the final part of the paper, we present the analysis of alternative scenarios characterized by different robots' operational features and deployment conditions, followed by a discussion on implementation and conclusions.

2. Related literature

The robot-assisted truck delivery problem studied here (hereinafter referred to as Traveling Salesman Problem with Robot or TSP-R) can be formulated similarly to other extensions of the "traditional" TSP characterized by the possibility of serving some of the customers with ancillary vehicles. Cuda et al. provide a relatively comprehensive overview (Cuda et al., 2015); their article is recommended for an in-depth literature review of two-echelon routing problems. However, the most relevant studies are summarized below.

The first extension of TSP to account for a combination of two modes in joint tours and sub-tours is the Truck and Trailer Routing Problem (TTRP). In this problem, a truck with a detachable trailer serves two different sets of customers: the first one that can be served by either a truck or truck-trailer, and the second one that can be served by a truck alone. The rationale behind this division was the presence of accessibility constraints in areas of the city that prevented the employment of large vehicles. In the TTRP, the solution that minimizes costs consists of a combination of the main route traveled by two joint vehicles and several truck sub-tours. Different heuristics have been adopted to solve this problem, including TABU search (Chao, 2002; Scheuerer, 2006), and Simulated Annealing (Lin et al., 2009). Villegas et al. (2011) and Villegas et al. (2013) propose metaheuristics based on a greedy randomized adaptive search procedure (GRASP) and use a "cluster-first, route-second" approach. Exacts approaches based on branching optimization techniques have been investigated by Drexl (2012) and Drexl (2014).

Another similar problem is the Traveling Salesman Problem with drones (TSP-D), where deliveries can be performed by means of an integrated system of trucks and drones.[1] Unlike customers in the TTRP problem, any customer can be served by a drone that is launched and picked up by a truck along its route. The drones typically have a flight range and capacity limited to one parcel (only one customer can be served per trip). Furthermore, drones' drop-off and rendezvous operations with the truck need to be coordinated. As a result, a considerable number of solutions involving different combinations of launch, pickup nodes and customers served might arise even for small problems involving a handful of customers. In the "flying sidekick traveling salesman problem," Murray and Chu (2015) formulate the problem for optimal truck and drone routes by means of a Mixed Integer Programming (MIP) formulation and propose different heuristic approaches to solve it. The same formulation has been adopted by Ha et al. (2018) who developed a heuristic GRASP based on a split procedure specifically tailored to the TSP-D. Wang et al. (2017) investigate the TSP-D from a worst-case point of view by identifying maximal potential savings (compared to the traditional truck-based delivery) from different delivery options. Other studies have focused on analytical approaches to identify the efficiency of such combined delivery systems (Carlsson and Song, 2017). Finally, Agatz et al. (2018) provide an alternative Integer Programming (IP) formulation of the problem based on the concept of "operations" that is adopted to optimally solve the TSP-D problem for instances up to 12 customers. In addition, they propose different heuristics based on the combination of local search and dynamic programming techniques. The number of studies on the topic of aerial drones and their applications has rapidly grown, and more than a hundred articles have been published between 2015 and 2017. For a recent up-do-date survey of this issue the reader is referred to Otto et al. (2018).

To the best of our knowledge, there are only two published papers that have formally investigated the implementation of robot-assisted parcel deliveries (Boysen et al., 2018; Jennings and Figliozzi, 2019). In the first study, the proposed delivery system relies on small depots where the parcels can be transshipped from the truck to robots in charge of single-item last-mile deliveries. The authors investigate optimal scheduling procedures for the delivery truck and the robots' drop-offs along the truck route by adopting different MIP formulations. In the second one, the authors propose a continue approximation model to identify the deployment of several robots dropped off and picked up by a single truck at predetermined points. Their corresponding travel time, distance, and delivery savings are determined for a specific case study. Two conference papers have also investigated the possibility of deploying sidewalk robots in support of truck delivery (Poeting et al., 2019; Sonneberg et al., 2019). In the first study, the authors develop a system of micro-depots served by trucks that can be used as hubs for robot deliveries. In the second one, a Location Routing Problem involving delivery robots is investigated to determine the efficiency of this new technology and the influence of alternative compartment sizes.

In this study, the TSP-R is considered more as a particular case of the TSP-D since no additional infrastructure is involved and the "drop-off" and "pickup" of a single robot needs to be coordinated with the truck delivery stops. There are, however, different features of this problem that make it unique among its own kind. In Table 1 we provide a qualitative summary of the main differences among the TTRP, TSP-D, and TSP-R based on the available literature. First, since robots, unlike drones, are always characterized by lower speeds than trucks, they are suitable for deliveries only in particular situations: relatively small-scale areas characterized by a high density of stops. For this reason, in our analyses we focus on scenarios of 50 customers in a 6 km^2 area. The loading/unloading operations are performed while the truck is stopped for a delivery (the robot cannot re-join the truck while it is cruising). Thus, the efficiency of loading/unloading operations becomes critical in the TSP-R (whereas this aspect is relatively overlooked in the TTRP and TSP-D). Finally, thanks to the possibility of developing divisible storage, robots can perform more than one consecutive delivery. While the addition of two or more storage compartments does not yield significant changes in the mathematical formulation from the TSP-D, the corresponding growth of solutions makes the TSP-R computationally more challenging. As a result, the optimization approach presented in this study leverages specific features of the robot problem to derive an efficient and accurate heuristic.

[1] Different definitions of the similar versions of the same problem have been found in the literature. In this study we refer to the combined problem as TSP-D, as suggested by Agatz et al. (2018).

Table 1
Main differences between TTRP, TSP-D and TSP-R.

	TTRP	TSP-D	TSP-R
Type of area studied	Variable (depending on the application)*	Variable (depending on the application)**	Urban area. Downtown/neighborhood size
Drop-off/Pickup points	The truck waits for the trailer at the same node	The truck drops and picks up the drone in any point	The truck drops and picks up the robot at delivery stops
Accessibility constraints	Some customers can be served only by a truck and not a trailer	Drones can serve any node as long as in their service range (10 to 20 km round trip)	Robots can serve any node as long as in their service range (less than 4 km round trip)
Efficiency (compared to TSP)	–	Dependent on operational and design constraints	Dependent on operational and design constraints
"Supporting mode" tour	Several deliveries	One single delivery	One or more consecutive deliveries

* Typically involves distribution problems in cities where neighborhoods cannot be easily accessed by truck. See Derigs et al. (2013) for a review of applications.

** Ranging from delivery problems in cities with several deliveries to emergency operations in rural areas with low accessibility. See Otto et al. (2018) for a review of applications.

3. Problem formulation and heuristic solution approach

In this section, after presenting the integer programming formulation of the TSP-R, we present a description of our heuristic optimization approach. Such an approach is based on a "route-first cluster-second" method where, by means of a local search, truck routes are explored in order to integrate robot deliveries.

3.1. Integer programming formulation

The TSP-R described in the previous section is formally described here by means of an IP formulation. This formulation is a modification of the TSP-D by Agatz et al. (2018), which is based on the concept of "operation." Here, an operation consists of a sequence of nodes that can be served by the truck with a robot on board, or by a truck and robot that split at a departure node, serve nodes independently and rejoin at a pickup node. In particular, given a graph $G(V, E)$ where the truck starts and finishes its trip from/to an entry (or depot) node v_0 and serves customers corresponding to l nodes ($v_1, \cdots v_l$), our problem formulation is based on the following assumptions:

1. Each customer node can be visited by a truck or a robot. This is a realistic assumption if we allow customers to come at the building entrance for pickup.
2. Each robot has a maximum capacity of p parcels. As previously discussed, robots' storage could be retrofitted to accommodate two or more compartments.
3. Robot pickup and drop-off can occur at any node, but cannot occur between two stops. While a more flexible framework with additional dedicated stops for pickup and drop-off might be more beneficial, it might not be always feasible (e.g. lack of parking).
4. Robots can be dropped off only when the truck is ready to leave for the following customer node since it is a procedure supervised by the driver, which requires an amount of time t_{drop}. Picking-up (empty) robots also requires additional time (in addition to the wait) expressed as t_{pickup}.
5. Each delivery operation requires a homogenous period of time t_{del}. In reality this task is likely to follow a normal distribution, however this simplification allows us to focus on analysis of the design and operations of robots.
6. Orders do not have delivery time-windows, hence there are no time-constraints affecting the order of stops. This is a reasonable assumption when considering relatively small sets of customers in the same area (most likely characterized by the same time-window).

Truck nodes are visited by the truck alone. Robot nodes are visited by the robot alone. Combined nodes are nodes that are visited by the truck and robot together. An operation o is composed of: two combined nodes, called the start or drop-off node (s_o) and end or pick-up node (e_o); n_o robot nodes, $j_1, j_2, \cdots, j_{n_o} \in J_o$ such that $n_o \leq p$; and a nonnegative number m_o truck nodes, $h_1, h_1, \cdots, h_{m_o} \in H_o$. An operation may contain a combination of truck and robot nodes or only truck nodes (meaning that the two modes travel jointly). Two different types of operations are illustrated in Fig. 2 where the same configuration of four customers is served by either a "joint" truck and robot trip or by the two modes working in parallel. (We refer to Agatz et al. (2018) for a more detailed explanation of this concept).

Each operation o is associated with a certain cost c_o corresponding to the time required to go from the start node to the end node, by visiting robot nodes (if any) and truck nodes. The operation's cost can be derived as:

$$c_o = \max \left\{ \begin{array}{c} d'(s_o, j_1) + \sum_{i=1,..n_o-1} d'(j_i, j_{i+1}) + d'(j_{n_o}, e_o) + n_o \bullet t_{del} + \delta_o \bullet (t_{drop} + t_{pick}); \\ d(s_o, h_1) + \sum_{i=1,...m_o-1} d(h_i, h_{i+1}) + d(h_{m_o}, e_o) + (m_o + 1) \bullet t_{del} \end{array} \right\} \quad (1)$$

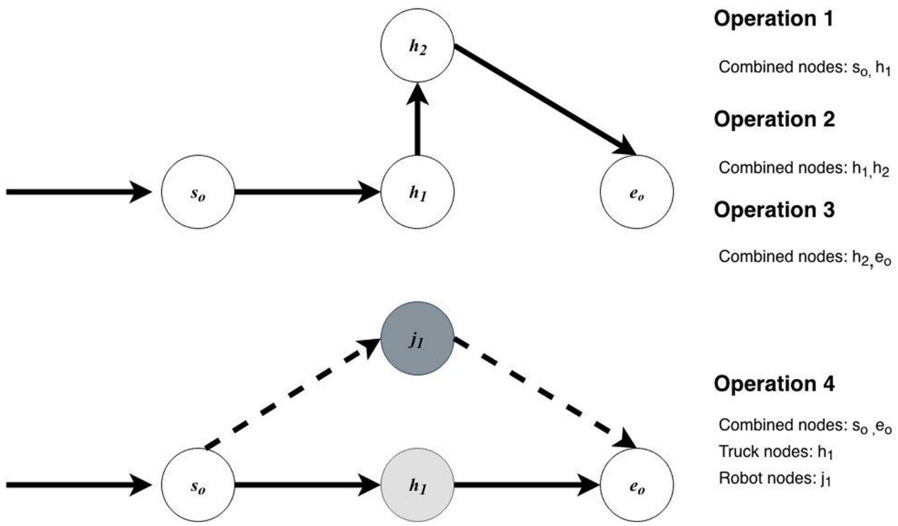

Fig. 2. Alternative types of operations within a sequence of four customers.

where $d(a, b)$ corresponds to the truck travel time from node a to node b, $d'(a, b)$ is the corresponding robot travel time for the same couple of nodes, and δ_o is 1 if operation o includes at least one robot note (meaning $n_o > 0$), 0 otherwise. Note that the operation can contain only truck nodes with $n_o = 0$. Since the robot can rejoin the truck while it is stopped for delivery, the operation cost corresponds to the minimum time for the two modes, including truck's stop at the pick-up node. A feasible operation should respect the capacity constraint.

The solution of the TSP-R could be expressed as the combination of different operations in the set of feasible ones O that minimizes the total time required to serve all the customer locations by either truck or robot. In essence, a mixed truck and robot route will be composed by a series of operations that connect each node at the minimum cost. Once all the feasible operations are created, the integrated routing problem can be formulated as the following IP problem:

$$\min \sum_{o \in O} c_o x_o \tag{2}$$

$$\textit{subject to} \quad \sum_{o \in O(v)} x_o \geq 1, \forall v \in V \tag{3}$$

$$\sum_{o \in O^-(v)} x_o \leq y_v, \forall v \in V \tag{4}$$

$$\sum_{o \in O^-(v)} x_o = \sum_{o \in O^+(v)} x_o, \forall v \in V \tag{5}$$

$$\sum_{o \in O^+(S)} x_o \geq y_v, \forall S \subset V\{v_0\}, v \in S \tag{6}$$

$$\sum_{o \in O^+(v_o)} x_o \geq 1 \tag{7}$$

$$y_{v_o} = 1 \tag{8}$$

$$x_o \in \{0, 1\}, \forall o \in O \tag{9}$$

$$y_v \in \{0, 1\}, \forall v \in V. \tag{10}$$

Here, x_o is a binary variable corresponding to 1 if operation o is chosen and 0 otherwise; y_v is a binary auxiliary variable corresponding to 1 if v is chosen as a start node in at least one operation and 0 otherwise; $O^-(v) \subset O$ represents the subsets of operations with v as a start node; $O^+(v) \subset O$ represents the subsets of operations with v as an end node, and $O(v) \subset O$ is the subset of operations that visit node v. The subset $S \subset V$ is introduced such that: $O^-(S)$ and $O^+(S)$ consist respectively of the subset of operations with start node and end node in S, and their corresponding end and start node in $V\backslash S$. Constraint (3) ensures that each node is covered by an operation. Constraint (4) ensures that at most one operation with start node v is chosen. Constraint (5) ensures that the chosen operations visit each node once and Constraints (6–7) ensure that truck-and-robot tours are connected by ensuring each customer node has at least one corresponding operation starting from and finishing there. Constraint (8) ensures that the delivery tour starts and ends at the given entry/exit node.

As Agatz et al. (2018) notice, it is possible to reduce the total number of operations by selecting the least expensive one among those with the same start node, end node, and equal sets of truck nodes and/or robot nodes (if any), but served in different orders. However, the number of operations would still grow exponentially, and the computational time required to optimally solve problems of any practical size (more than 8–10 customers) would be prohibitive. For this reason, in this study, we develop a heuristic for TSP-R (Section 3.2), based on the similar concept of "robot operation" for given truck tours. Such an approach not only significantly simplifies the original problem, but also allows for a very efficient computational approach.

3.2. Heuristic solution approach

The proposed heuristic is based on the observation that, for a given truck-only route, finding the "optimal set of robot operations" that improves such a route corresponds to a particular case of the Maximum Weighted Independent Set Problem known as Weighted Interval Scheduling Problem. Such a problem can be solved efficiently and optimally with a dynamic programming algorithm. Hence, a "route-first cluster-second" method is adopted, where an initial truck route is created and then the corresponding optimal combination of robot operations is obtained with dynamic programming (Section 3.2.1). A Local Search with Adaptive Perturbation (LS-AP) where each solution corresponds to a different initial truck route is performed to iteratively improve the TSP-R solutions (Section 3.2.2).

3.2.1. Optimal robot deployment for any given truck tour

Given a truck tour, a "robot operation" consists of a robot sub-tour and it is represented by a sequence of nodes that include a drop-off node where the robot departs from the truck (v_s) and a pick-up node where the robot rejoins the truck (v_e). A robot operation can contain from 3 up to ($p + 2$) nodes (according to the robot capacity constraint p). If the stops in the original truck tour can be indexed according to their order as: $\{0, 1, 2, \cdots, n, n+1\}$, then a robot operation k can be defined as a subset R such that: $v_s \in \{0, 1, \cdots, n-1\}$ and $v_e \in \{1, \cdots, n+1\}$ and $v_s \leq v_e$. Based on this definition, each operation can be identified as a "forward" operation or "backward" operation. A forward operation consists of a subset R where all the served customers' indices in the original truck route sequence are lower than the pick-up node. Vice versa, a backward operation consists of a subset R where at least one served customer's index in the original truck route sequence is higher than the pick-up node.

A robot operation essentially entails that one or more of the original stops truck route can now be served by a robot. Depending on the choice of the drop-off and pick-up nodes, the truck can perform several stops while detached from the robot. Given this definition, for each robot operation it is possible to identify the corresponding "savings" s_k deriving from employing the robot between node v_s and v_e:

$$s_k = c(H[v_s, v_e]) - c_k + e_k \tag{11}$$

where $c(H[v_s, v_e])$ corresponds to the original truck tour cost between the drop-off and pick-up nodes (including all nodes visited in between), c_k corresponds to the robot operation's cost and it is derived from the first argument of the max in Eq. (1), and e_k corresponds to the "extra savings" derived from serving some customers after the pickup node in case of a backward operation. Indeed, in some particular situations, if no served node has index between v_s and v_e, the operation would yield negative savings as the original truck route between these two nodes would be unaltered. Such "extra savings" can be derived as follows:

$$e_k = c(H[v_e, v_0]) - c(H'[v_e, v_0]) \tag{12}$$

where $c(H'[v_e, v_0])$ corresponds to the updated cost of the "new" truck route between the pickup node and the exit node with the same robot customers removed. In a nutshell, based on this formulation the overall savings of a specific robot operation depends on the original sequence of nodes visited by the truck, the drop-off and pickup nodes, and the replaced original truck nodes by robot nodes.

Based on this formulation of robot operations, the "TSP-R sub-problem" for a given truck route, could be straightforwardly reformulated as a Maximum Weight Independent Set Problem (MWISP) where the objective is to maximize total savings by selecting the optimal combination of robot operations:

$$\max \sum_{k \in K} s_k x_k \tag{13}$$

$$\textit{subject to} \ \sum_i a_{ik} x_i \leq M \bullet (1 - x_k), \ \forall \ k \in K \tag{14}$$

$$x_k = \{0, 1\}, \ \forall \ k \in K \tag{15}$$

where the variable x_k indicates whether an operation k with savings s_k is chosen, and constraint (14) is introduced to avoid "conflicting" operations to be chosen together, and it is based on a "preprocessed" $K \times K$ size conflict matrix A, where K is the number of all potential bot operations. The values of the conflict matrix, a_{ik} are equal to 1 if operation i and operation k are conflicting with each other, and 0 otherwise.

Operation i and operation k are considered conflicting if:

- They have the same drop-off node, i.e., $v_{e,i} = v_{e,k}$
- They have the same pickup node, i.e., $v_{s,i} = v_{s,k}$

- They share any served node: $R_i \cap R_k = \emptyset \backslash \{v_{s,i}, v_{e,i}\}$
- They overlap such that $v_{e,i} < v_{e,k}$ and $v_{e,i} < v_{s,k}$

Based on this definition and on the fact that each solution is derived from an initial truck route, the connectivity is guaranteed. In order to reduce the size of the optimization problem, it is possible to exclude some operations from the set of potential solutions without losing the conditions for optimality (Algorithm 1). First, robot operations that yield negative savings can be discarded since they would imply an inefficient use of the robot in comparison to just performing deliveries with the truck (for the specific configuration envisaged by the operation). Second, it is possible to identify several "dominated" robot operations. For example, among robot operations with same drop-off, pickup and served nodes, only the one yielding the highest savings can be considered. In addition, among operations with the same drop-off and pickup nodes, it is possible to consider only the one with the highest savings. Indeed, regardless of the customers served and their serving order, between the drop-off and pickup node it is not possible to perform more than one forward operation. Furthermore, operations are "not affecting" other preceding or following operations. Hence, for the principle of optimality, given a set of robot operations with the same drop-off and pickup node, "dominated" ones can be discarded from the solution space.

Another possibility of reduction of bot operations comes directly from the operational features of the TSP-R. For any given truck tour, it is possible to exclude *a priori* some robot operations, based on the fact that the robot is slower than the truck. Given a potential drop-off node, if the time required by the robot to reach and serve a potential customer is higher than the cost of completing the remaining tour by truck, then such robot operation would yield negative savings, and can hence be discarded from the solution set. For each node visited in the original truck route (H), it is possible to identify a "Reachable Set" (S) of customers/nodes that can potentially determine positive savings. Considering the Reachable Set of customers allows to accelerate the creation of robot operations to different extents, depending on the speed ratio between the two modes, the area served and configuration of customer nodes in the area.

Algorithm 1. (*Pseudocode to derive robot operations*).

```
INPUT: truck tour H
  OUTPUT: robot operation set K
  FOR each customer h in H\{n,n + 1} (following the order of customers served):
    Generate Reachable Set S
    FOR each customer s in S:
      Generate operations sets P_i given by permutation of j served nodes of size i different from the given drop-off h (v_s) and pick-up s (v_e):
      P_i = (S, i) such that 1 ≤ i ≤ p (capacity constraint) and j ≠ h ≠ k for any j in the permutation set P_i
      Set s_max = 0
      FOR each operation j in P_i:
        Derive savings s_j by means of Eq.11
        IF s_j is positive:
          IF s_j > s_max:
            Add operation to the robot operation set K
        END IF
      END FOR
    END FOR
  END FOR
```

Given the particular "sequential" nature of the MWISP[2] it is possible to relate it to the more specific Weighted Interval Scheduling Problem that can be solved very efficiently by means of dynamic programming (DP).

Each operation k with savings s_k can be considered as an interval characterized by a start index $t_k = v_s$ and a finish index $f_k = v_e$ in case of forward operation or $f_k = \max\{v_e, \max(v_j)\}$ (where $\max(v_j)$ corresponds to the served node with the highest index) in case of backward operation. Two operations (or intervals) i and k are "compatible" or non-conflicting if $t_i \geq f_k$ or vice versa, meaning that they do not overlap. Based on that, it is possible to adopt the DP approach proposed by Kleinberg and Tardos (2006) to solve this type of problem. Given the set of K robot operations, the algorithm consists of the following steps:

1. Sort the operations such that $f_1 \leq f_2 \leq \cdots \leq f_K$
2. For each operation k calculate i_k, which corresponds to the largest positional index i such that i and k are disjoint (highest index of compatible prior operation)
3. For each operation k calculate the optimal solution to the problem, OPT(k), which corresponds to:

$$OPT(k) = max\{s_k + OPT(i_k), OPT(k-1)\}.$$

In order to avoid recursion, it is possible to store the values in global array N of optimal solutions to sub-problems and adopt the following iterative algorithm (Algorithm 2):

[2] The initial truck route can be expressed as a sequence of customers served in chronological order.

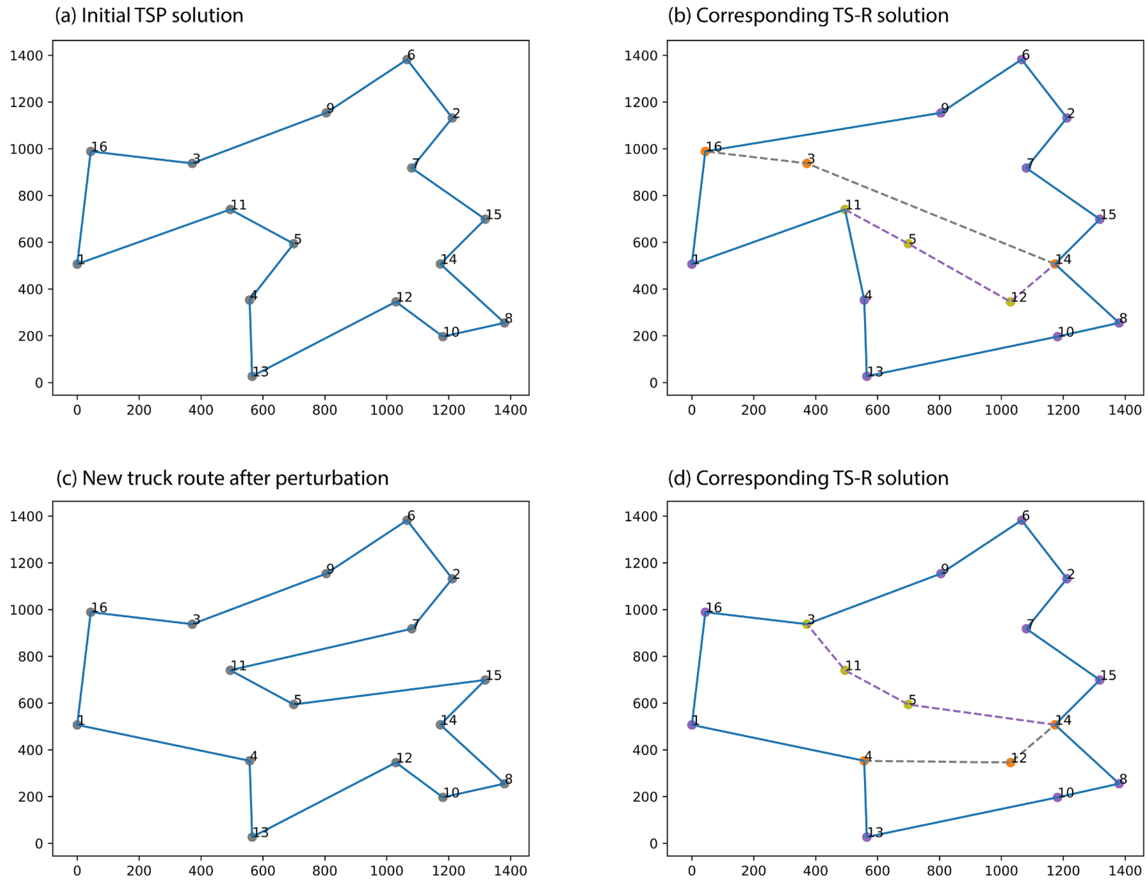

Fig. 3. Example of proposed ILS approach for a route of 15 customers (truck route in solid line and robot routes in dashed line).

Algorithm 2. (*Algorithm to compute optimal robot solutions for a given truck route*).

```
Set N[-1] = 0
    Set N[0] = 0
    FOR k = 1,…,K:
      N[k] = max{s_k + N[i[k]], N[k − 1]}
```

4. Compute the solution by "tracing back" through array N

The algorithm has a complexity $O(K\log K)$ and yields optimal solutions. We refer to Kleinberg and Tardos (2006) for formal proofs of the optimality of the algorithm.

In order to further reduce the computational efforts for problems involving a large number of customers (e.g. 50 or more), the original truck route could be divided in smaller chunks where the same approach is adopted for each of them independently. This approach would however introduce a loss of efficiency as it would miss operations involving nodes from different segments.

3.2.2. *Iterated local search*

As already discussed in other studies (Cuda et al., 2015), starting with the optimal TSP tour does not necessarily yield the optimal solution to the two-echelon or two-mode routing problems. However, it is also expected that TSP-R solutions derived from the optimal truck tour are typically better than the majority of solutions obtained from random tours. For this reason, in the first step of the algorithm a TSP-R solution is derived from the best truck route and performing a systematic local search with the goal of identifying better combined truck and robot solutions by modifying the original truck tour. In Fig. 3, the evolution of the solution in the Iterated Local Search (ILS) for a configuration of 15 customers is shown: first, an initial truck-bot configuration is derived from the TSP solution; then after the truck route is modified, an improved truck-bot configuration is identified.

In order to find a suitable balance between an intensified and diversified search ("exploitation vs. exploration") we adopt an adaptive perturbation strategy where two different types of perturbations are applied depending on the current state of search (Benlic and Hao, 2013). Its algorithmic steps are illustrated in Algorithm 3. "Directed perturbations" are performed at early steps after a local improvement is obtained and they consist of the "first" improving 2-opt swap move identified for a random customer in the truck tour sequence. In case no improving swap move could be identified for such randomly selected customer, other customers are tried until a threshold number of attempts is reached. With increasing number of local searches performed without improvement, diversification grows by applying "random perturbations" with higher probability. Random perturbations are characterized by a variable number of swap moves selected uniformly at random. The number of moves also depends on the state of the search and increases with consecutive local searches without improvements. The probability of applying directed perturbations P is derived as follows:

$$P = \max\left(P_0, e^{-\frac{i_{wi}}{T}}\right) \tag{16}$$

where P_0 corresponds to a minimum given threshold for selecting a directed perturbation, i_{wi} corresponds to the number of local search iterations without improvement, and T a given threshold (input parameter).

Algorithm 3. (*Pseudocode for Local Search with Adaptive Perturbation***).**

INPUT: Initial (optimal) truck tour H, max total local searchesn_l, maximum local searches without improvement n_{wi}
$H \rightarrow S(H)$ by Algorithm (1–3) ***%Derive TSP-R Solutions***
$S \rightarrow S_{best}$ ***%Add solution to the best solution***
Set $i = 0$ and $i_{wi} = 0$ ***%Initialize counter total iterations and iterations without improvement***
WHILE $i < n_l$ and $i_{wi} < n_{wi}$:
Derive P with Eq. (16) ***%Derive probability of directed perturbation/randomized perturbation***
IF $P < $ random(0;1):
$\varphi_d(H) \rightarrow H'$ ***%Apply directed perturbation (first improving 2-opt swap move)***
ELSE:
$\varphi_r(H) \rightarrow H'$ ***%Apply randomized perturbation***
END IF
$H' \rightarrow S'(H')$ ***%Derive new TSP-R solution***
IF$S' < S$:
$S \rightarrow S_{best}$ ***%Add new solution to the best solution***
$i_{wi} = 0$
ELSE:
$i_{wi} = i_{wi} + 1$
END IF
$i = i + 1$
END WHILE

4. Performance of the heuristic approach

The quality of the proposed heuristic approach is tested for randomly generated instances with different features. Depending on the size of the instance considered, the results of the heuristic are compared to the solution of the corresponding integer formulation solved with commercial software or with a theoretical lower bound.

instance

4.1. Instances

For the experiments, several instances are randomly derived by varying: the number of customers (8, 15, 30, 50 and 100), the size of served area (2 km^2, 4 km^2, 9 km^2, 12 km^2, and 15 km^2), and the geographical distribution of customers (minimum distance between customers ranging from 100 m to 500 m) (Table 2). In total, 68 instances are generated.

Considering the nature of the problem investigated (the last-mile delivery issue for specific neighborhoods or downtown areas), it is reasonable to limit ourselves to relatively small-scale problems (up to 100 customers, although solving optimally a TSP-R involving even smaller (say, 12 customers) by means of state-of-the-art IP solvers would already require hundreds of thousands of variables and lead to already prohibitive computation times). For these experiments we assume the robot speed $v_r = 1.0$ m/s (slightly lower than an adult pedestrian's average speed of 1.4 m/s (Knoblauch et al., 1996)), and the truck speed $v_t = 6m/s$ (in line with the average traffic speeds identified in different urban settings across the world (Uber Movement, 2019)). Each delivery requires a stop (t_{del}) of 180 s (3 min). These values are in line with Conway et al. (2016) and Allen et al. (2017). Finally, we assume the average pickup/drop-off (launch) time $t_{drop} = 60s$.

4.2. Performance analysis

In order to investigate the quality of the proposed heuristic, we compare its results with those obtained by solving the problem according to the original IP formulation using the commercial optimization solver GUROBI. Preliminary tests have shown that, with a laptop machine (i7-7700HQ, 2.80 GHz), it is not possible to solve optimally instances larger than 8 customers in less than a few hours

Table 2
Overview of the instances.

Instances	Customers	Served Area (km^2)	Customers Spacing (min distance in m)
1–4	8	2	100
5–8	8	2	200
9–12	8	4	200
13–16	8	9	200
17–20	15	2	200
21–24	15	4	200
25–28	15	4	400
29–32	15	9	400
33–36	30	4	200
36–40	30	4	400
41–44	30	6	200
45–48	30	9	400
49–52	50	6	200
53–56	50	6	300
57–60	50	12	300
61–64	100	9	300
65–68	100	15	500

Table 3
Performance of the proposed heuristic in comparison with the optimal solution.

Instance	τ	σ_1	$\gamma_{1,w}$	$\gamma_{1,b}$	$\Delta\bar{\gamma}_1$	σ_2	$\gamma_{2,w}$	$\gamma_{2,b}$	$\Delta\bar{\gamma}_2$	σ_3	$\gamma_{3,w}$	$\gamma_{3,b}$	$\Delta\bar{\gamma}_3$
1	2158.5	2001.8	2001.8	2001.8	**0.00**	1829.8	1843.8	1829.8	**0.31**	1829.8	1848.9	1829.8	**0.35**
2	2081	1873.0	1879.3	1873.0	**0.03**	1741.5	1741.5	1741.5	**0.00**	1741.5	1741.5	1741.5	**0.00**
3	2127.6	1935.5	1935.5	1935.5	**0.00**	1841.0	1841.0	1841.0	**0.00**	1841.0	1841.0	1841.0	**0.00**
4	2127.6	1997.1	1997.1	1997.1	**0.00**	1845.0	1880.8	1845.0	**0.19**	1845.0	1845.0	1845.0	**0.00**
5	2085	1863.6	1863.6	1863.6	**0.00**	1747.3	1747.3	1747.3	**0.00**	1747.3	1747.3	1747.3	**0.00**
6	2261.6	2062.3	2069.1	2062.3	**0.03**	1933.3	1933.3	1933.3	**0.00**	1933.3	1933.3	1933.3	**0.00**
7	2116.6	1949.3	1949.3	1949.3	**0.00**	1769.1	1769.1	1769.1	**0.00**	1769.1	1769.1	1769.1	**0.00**
8	2133.3	1893.1	1893.1	1893.1	**0.00**	1851.5	1851.5	1851.5	**0.00**	1851.5	1858.0	1851.5	**0.21**
9	2410.3	2258.8	2258.8	2258.8	**0.00**	2060.5	2060.5	2060.5	**0.00**	2060.5	2060.5	2060.5	**0.00**
10	2541.1	2257.6	2294.8	2257.6	**0.11**	2240.0	2240.0	2240.0	**0.00**	2240.0	2240.0	2240.0	**0.00**
11	2516.3	2387.1	2395.3	2387.1	**0.07**	2303.3	2303.3	2303.3	**0.00**	2303.3	2303.3	2303.3	**0.00**
12	2264.6	2112.1	2112.1	2112.1	**0.00**	1912.1	1931.6	1912.1	**0.20**	1912.1	1912.1	1912.1	**0.00**
13	2807.1	2678.3	2678.3	2678.3	**0.00**	2520.3	2520.3	2520.3	**0.00**	2520.3	2520.3	2520.3	**0.00**
14	2723.1	2426.6	2426.6	2426.6	**0.00**	2366.6	2366.6	2366.6	**0.00**	2366.6	2366.6	2366.6	**0.00**
15	2884.3	2719.1	2719.1	2719.1	**0.00**	2536.6	2536.6	2536.6	**0.00**	2536.6	2536.6	2536.6	**0.00**
16	2928	2691.9	2691.9	2691.9	**0.00**	2691.9	2691.9	2691.9	**0.00**	2691.9	2691.9	2691.9	**0.00**

(similarly to previous studies on the TSP-D). Hence, for practical reasons, we focus here on solving optimally only the first 16 instances (Instance 1–16). In Table 3, the performance of the heuristic is reported for different levels of the robot's available storage capacity (varying p to maximum 1, 2, or 3 customers served consecutively) in comparison to the optimal solutions. The computational effort required to solve the problems to optimality for each level of capacity corresponds respectively to: 190 s, 570 s, and 1660 s (on average). The average computation time for the heuristic, given a maximum threshold of 1000 local searches (truck route iterations), varies between 1.1 and 14.1 s depending on the level of the robot's capacity. We indicate the original minimum truck cost (obtained by solving the corresponding truck-only TSP problem optimally) with τ. The optimal solution values of the TSP-R for three different levels of capacity consisting of 1, 2 or 3 maximum items (allowing for more consecutive deliveries per robot trip) are indicated respectively with σ_1, σ_2, and σ_3. Since the iterated local search is probabilistic, we report both the best- and worst-case performances of the heuristic out of 100 runs for each instance to give a more complete analysis of the heuristic performance. The worst and best solutions obtained from 100 runs of the proposed heuristic are indicated respectively with γ_w and γ_b. The average percentage gaps between the optimal solution and the results from the heuristic runs are indicated respectively with$\Delta\bar{\gamma}$.

The results shown in Table 3 indicate that the proposed heuristic, for relatively small instances, achieves results within a 0–0.2% gap from the optimal solution of the TSP-R on average. In all the analyzed instances, the heuristic approach was able to obtain, at least once (in some cases in all runs), the optimal solution. It is interesting to see that the quality of the heuristic seems independent from the maximum capacity of the robots, as the average gap does not increase for increasing levels of capacity. It is also interesting to see that, given the relatively small number of customers (8), and speed ratio between robot and truck (1 to 6), both in the optimal TSP-R and heuristic solutions, the capacity of robots is not fully utilized for capacity higher than 2.

For larger instances (Instances 17–68), the results of the heuristic (for single robot capacity) are compared to the solution achieved with a regular TSP. We summarize in Fig. 4a-b the results of the experiments for larger instances (Instance 17–68), with the standard parameters used in Section 4.1 (v_r = 1m/s, v_t = 6m/s, t_{stop} = 180s) and different n_l, n_{wi} parameters for the Local Search in relation to the size of the problem (respectively 500 and 100 for instances with 15 customers, 200 and 50 for instances with 30

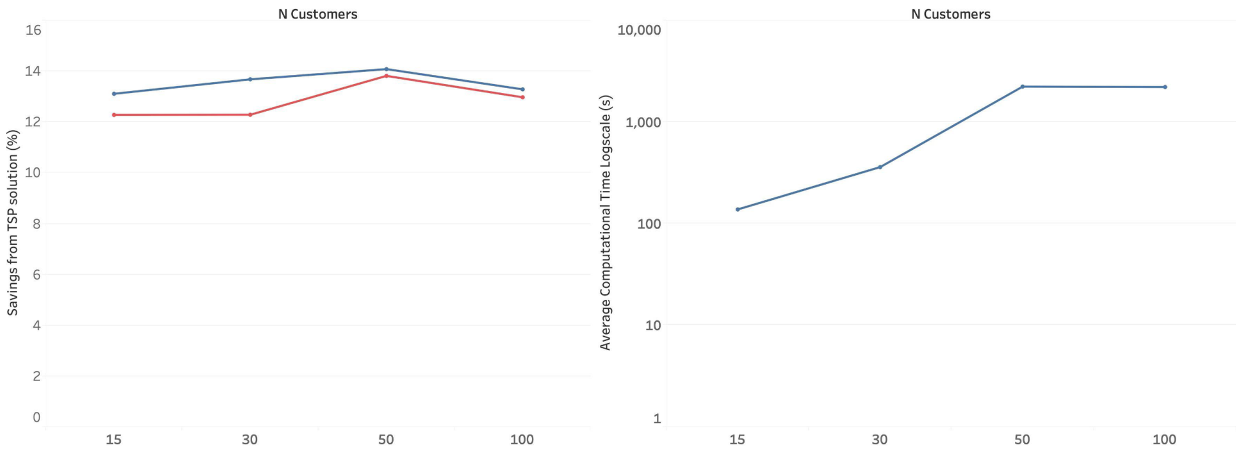

Fig. 4. Maximum savings (blue) and minimum savings (red) achieved for different instance sizes (a) and corresponding average computational times (b).

customers, 50 and 10 for instances with 50 customers, and 20 and 5 for instances with 100 customers). The full results of the experiments are reported in Appendix A. The results obtained for larger instances are in line with those of smaller instances. Although it is not possible to directly compare the heuristic solutions with the optimal ones, the achieved savings (in comparison to the TSP solution) are consistent for increasing sizes of the problem and always within a 1–2% range from those obtained for Instances 1–16. As expected, the average savings achieved with the heuristic vary depending on the typology of scenario investigated. For relatively dense configurations (i.e., 15 or more customers within an area of 2 km^2, 30 customers or more within an area of 4 km^2) the heuristic solutions achieve savings between 14% and 17%. However, for sparser configurations (i.e., 15 customers within an area of 4 km^2, 30 customers an area of 9 km^2, and 50 customers or more for an area of 12 km^2 or more), the savings would decrease to 10–13%. Lower spacing among customers (and consequently the higher presence of customer clusters) seems to favor robot utilization and more efficient integrated routes (as shown in Fig. 5).

Finally, although the obtained results are characterized by a decreasing number of maximum local searches in order to manage the computational effort, the solution quality does not seem to be significantly affected. In future research, it would be interesting to investigate the performance of the heuristic for higher thresholds of local search with more powerful computational resources.

Overall, the results of the different tests show that the proposed heuristic approach yields to solutions reasonably close to the optimal ones, with a high level of confidence particularly in scenarios characterized by relatively small areas and high customer density.

5. Savings and factors of influence

In this section, different factors influencing the overall performance of truck-robot delivery systems are investigated: the speed of traffic and the truck speed, the robot's storage capacity, and the duration of drop-off operation. Based on the results, some considerations on practical implications for the implementation of robot-assisted delivery systems are made.

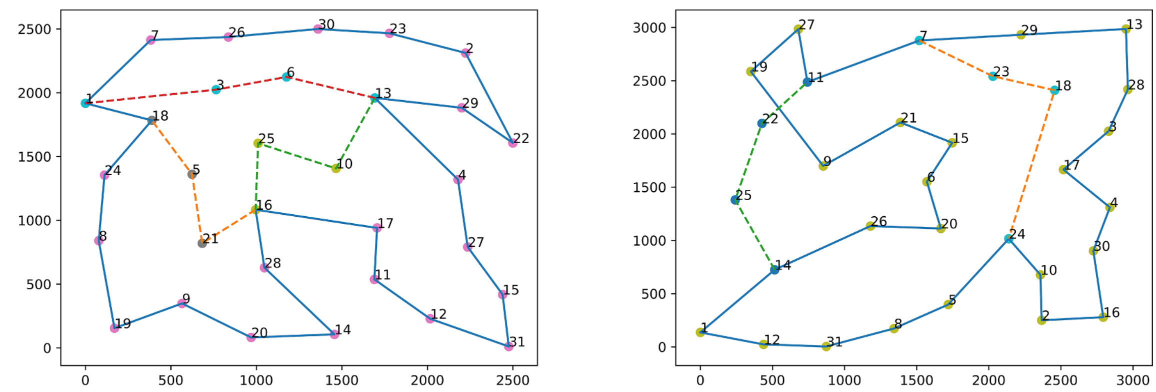

Fig. 5. TS-R solution for different levels of customer density: (a) Instance 35: 30 customers in 4 km^2 and (b) Instance 45: 30 customers in 9 km^2 (truck route in solid line and robot routes in dashed line).

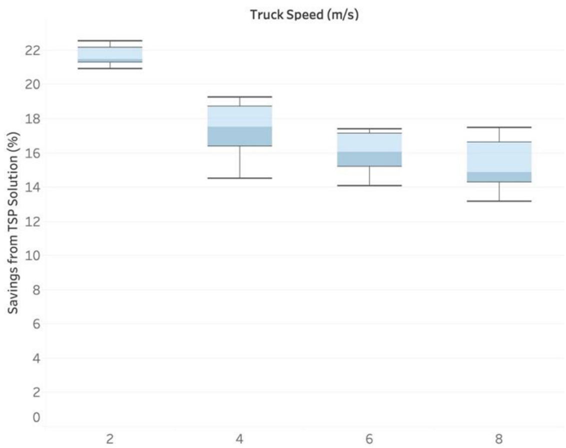

Fig. 6. Analysis of average savings from the TSP-R with different truck (or traffic) speeds, [6]Each box plot shows the first quartile, median, third quartile, and lowest/highest datum within 1.5 IQR of the lower/higher quartile, out of 10 instances.

5.1. Influence of speed ratio between truck and robot

The influence of the speed ratio between the two modes, which can reflect the effect of traffic congestion and given infrastructure (e.g. traffic lights and street configuration) on the efficiency of the robot-assisted delivery system is tested by means of several simulations assuming constant robot speed and variable truck speed (Fig. 6). Several randomized delivery scenarios (10) characterized by different configurations of 50 customers in an area of 6 km^2 are analyzed for a given robot speed of 1 m/s and truck speeds ranging from 2 m/s to 8 m/s (respectively corresponding to a travel cost ratio $\alpha = 0.5$ and $\alpha = 0.125$). These scenarios could represent a 3–4 h portion of alternative truck routes taking place in different types of neighborhood (e.g., commercial vs. residential) or different times of the day (e.g. peak and off-peak hours) or even different downtown speeds in certain cities. Such values are consistent with the average downtown last-mile speeds[3] measured in several cities' central business districts (Reed and Kidd, 2019). For example, 2018 INRIX's Global Traffic Scorecard reported average values around 4–5 m/s (8–10 mph) for cities like London, Paris, and San Francisco. As expected, the most significant savings (over the corresponding TSP solution) occur in situations characterized by high levels of congestion where the truck is only two or four times faster than the robot. In these two cases the savings are respectively around 21% and 17%. For higher truck (and traffic) speeds, the savings obtained from the usage of robots decrease to values around 15%. Interestingly, above a certain truck speed (6 m/s or higher), the amount of savings stabilizes around 15%. This phenomenon can be explained by the possibility of reducing the total delivery times thanks to the parallelization of delivery operations, which takes 180 s. Indeed, under these assumptions, it is still possible to perform at least one robot operation per twelve stops to serve one or two customers.

Although it is unlikely that motorized traffic would have speeds lower than 2–3 m/s in most of the urban settings worldwide (Reed and Kidd, 2019), it is possible that some limited areas would experience temporary increases in travel times due to accidents or special events. In these types of situations, robot-assisted deliveries could become particularly beneficial to avoid clusters of congestion and reduce the overall delivery time. In Table 4, the savings (in comparison with the truck-only TSP solution) for alternative levels of congested stops (where traffic speed is equal to 1 m/s in the first/last 200 m surrounding the stop) and general truck speeds are shown for a scenario of 50 customers. For scenarios with relatively limited congested areas (e.g., 10 or 20% of congested stops), the obtained savings are already considerably higher than the corresponding ones with homogenous general truck speeds (e.g., for 30% of congested stops, the savings almost double for any initial truck speed). Interestingly, for the same levels of congested stops, similar savings can be achieved regardless of the initial truck speed (already for relatively …). This result can be explained by the possibility of the joint truck-robot delivery system to efficiently adapt and modify the route to minimize the delivery time. For example, as shown in Fig. 7, for the same configuration of 50 stops (included the congested ones) and two different average truck speeds (4 m/s and 6 m/s), the robot operations' drop-off and pick-up points, and the truck's sequence of stops change leading to similar levels of savings (respectively 23.5% and 23.0% in comparison to regular TSP solution). For increasing speeds, the truck would modify its route such that the robot could serve several congested areas in a single operation.

5.2. Influence of robot's storage capacity

The possibility of performing several consecutive deliveries with the robot, thanks to its different storage compartments, represents one of the main advantages of this delivery system (especially in comparison to drones). In order to quantify the influence of

[3] Last mile speed is defined here as: "The speed at which a driver can expect to travel one mile into the central business district during peak hours". (Reed and Kidd, 2019)

Table 4
Savings (%) corresponding to alternative levels of congested stops for given average traffic speeds.

General truck speed (m/s)			
Congested stops	**4**	**6**	**8**
0%	18.2	16.6	14.6
10%	22.7	22.8	22.9
20%	25.7	26.1	26.4
30%	28.0	28.2	27.9

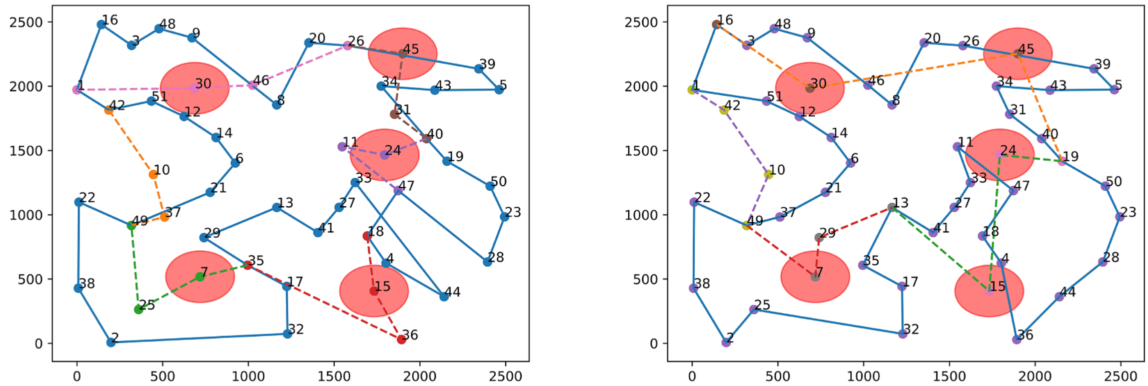

Fig. 7. TSP-R solutions with congested stops (in red circles) for alternative truck speeds at (a) 4 m/s and (b) 2 m/s (truck route in solid line and robot routes in dashed line)

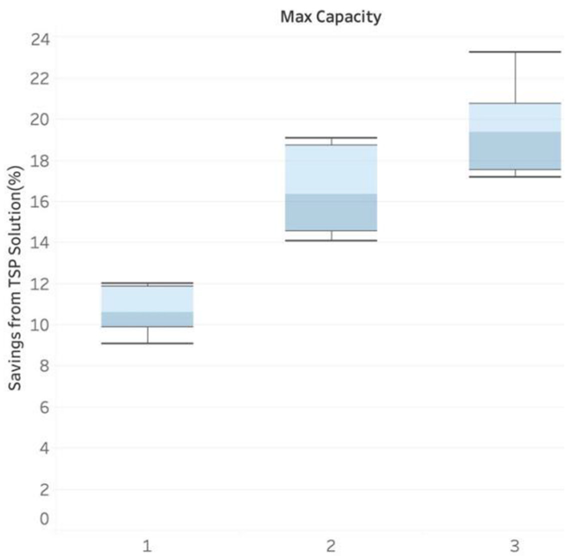

Fig. 8. Savings per different levels of robot capacity (maximum consecutive stops per operation).

this feature, 10 randomized scenarios with the same characteristics as earlier are investigated for three levels of available robot capacity, or maximum consecutive stops per robot operation (1, 2 or 3). As shown in Fig. 8, the savings (in comparison to the original TSP solution) significantly increase from 10.7%, corresponding to the single-stop robot operation solutions, to 16.5% and 19.6% corresponding respectively to 2- and 3-stop robot operation solutions. Increasing the number of compartments from one to three can hence double the efficiency of robot delivery systems. It is also interesting to see that, for increasing levels of capacity, the majority of robot operations utilize the added capacity. In scenarios with two and three maximum stops per robot operation, almost 93% and 68% of all robot operations are performed at capacity (Fig. 9). Given the relatively low speed of the robot that can serve only a few customers, allowing multiple consecutive stops allows for longer operations and more efficient solutions.

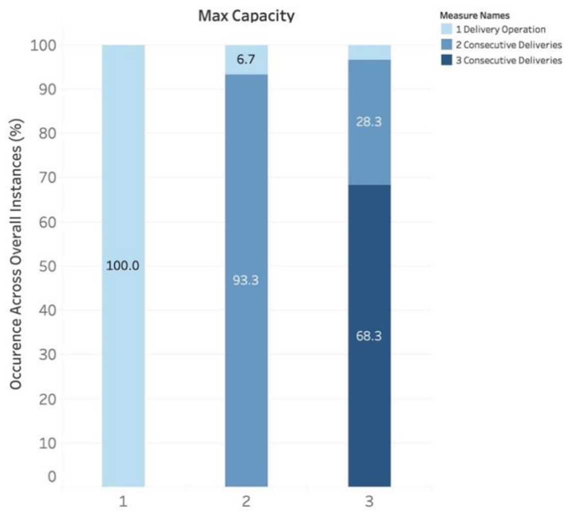

Fig. 9. Capacity usage per different levels of robot capacity (maximum consecutive stops per operation).

5.3. Influence of pickup and drop-off operations length

Another possible factor of influence on the overall efficiency of integrated truck-robot delivery systems corresponds to the length of drop-off operations. Depending on the level of automation of the truck-robot interface and the possibility of loading the robots' compartments with items before the actual launch, this process could take from a few seconds to a few minutes. In Fig. 10, the impacts of three alternative drop-off durations, ranging from 60 s to 3 min, are tested for 50 randomized scenarios with the same characteristics as before (Original instances have $t_{drop} = 60$ s). The results show that the overall efficiency of the system could decrease by around 5% when the length of such operations increases from 60 s to 180 s.

5.4. Influence of robot's operation range

Another important factor influencing the overall efficiency of integrated delivery systems corresponds to the maximum range of operation delivery robots. Given robots' reliance on rechargeable batteries, the presence of several physical obstacles in urban

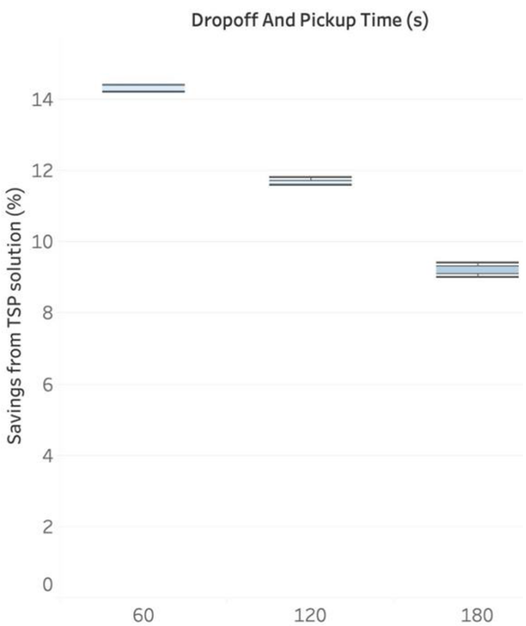

Fig. 10. Savings per different durations of drop-off operations.

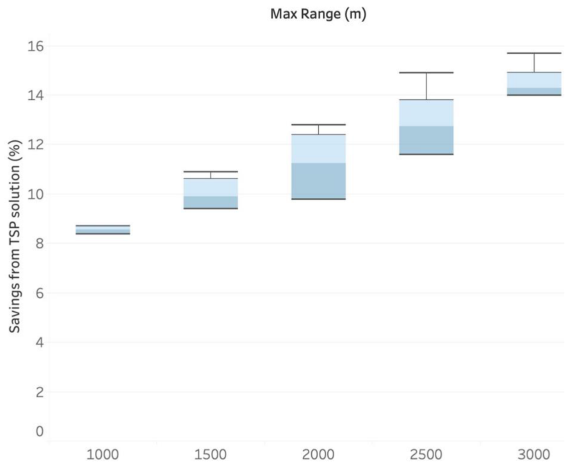

Fig. 11. Savings per different maximum distanced allowed for drop-off operations.

environments and general safety concerns, robots' trips distance might be limited to a few hundred meters. It emerges from Fig. 11, that imposing different maximum distances per bot operation has a considerable impact on the overall efficiency of the system. In particular, the savings achievable by robots increases steadily when the maximum threshold for operation distance grows from 1000 m (8.7% savings) to 3000 m (14.6% savings). The gain per 500 m step increase varies between 1.36% and 1.66%. For thresholds above 3000 m, the achievable savings are in line with those obtained in the baseline configuration (i.e. around 19–20 percent for a truck/robot speed ratio of 6).

5.5. Discussion

The results from the previous analyses allow some useful considerations regarding the implementation of robot-assisted delivery systems in scenarios characterized by a few dozens of customers, where one robot can depart from the truck for one or more consecutive deliveries. Depending on several factors concerning the operational and design characteristics of robots, and the surrounding environment, such integrated systems could provide significant travel time savings (compared to the traditional delivery method by truck).

Given the relatively low speed of the robot (1 m/s or 4 km/h), the ideal deployment scenario consists of a limited area (downtown or neighborhood size), characterized by a high level of traffic congestion (with average speed below 4 m/s or 15 km/h), and with a dense customer configuration (10 or more customers per square kilometer). Here, the savings achieved with respect to the truck-only tour would be around 20 percent. These kinds of conditions could realistically occur in densely populated urban areas during the most congested hours of the day. On the other hand, navigation in very crowded areas represents a significant challenge for robots. The presence of small areas of congestion around a limited number of customers also represents a valuable opportunity for the deployment of the robot, which can increase the overall delivery efficiency up to 30 percent. This type of scenario often occurs as a result of traffic accidents, temporary bottlenecks and special events.

Design and operation features of the robot, such as its storage capacity and launch time, play a significant role in the overall system efficiency as well. The possibility of dividing the storage into multiple compartments for multiple consecutive deliveries allows considerable efficiency gains (from 10 to 17 percent). Since presently tested robots by private companies appear to have only one single storage compartment, their current design might have to be modified to accommodate divisible storage. The maximum traveling distance of the robot also plays a significant impact on the performance of the systems: a relatively low threshold of 1000–2000 m can curb the overall savings (over traditional truck systems) by 6–8 percent. This is an important factor to consider in view of the battery life of these devices and the possibility of recharging them in a timely fashion. For example, the longer ranges might involve higher battery utilization and ultimately longer re-charging times. Finally, the efficiency of robot drop-off operations can affect the overall performance of the system with differences of about 6 percent between 1-minute and 3-minute deliveries. The length of this process will strongly depend on the level of automation and integration between the truck and robot. For example, the possibility of loading the robot while the truck is *en route* might considerably save time. The results of these analyses are partially in line with those from previous studies. Jennings and Figliozzi (2019) identified higher values for a slightly different system (involving several robots dropped off and picked up at predetermined locations) and by using a continuous approximation method. Although Boysen et al. (2018) focus on lateness of deliveries as a performance measure, a relation can be seen between the positive effect of higher density of drop-off points in their study and the overall customers' density in ours.

Compared to drone-assisted deliveries, robot-assisted systems are in general less performing due to their relatively low speeds and restriction to operate on the same road infrastructure network of trucks. Regardless of the adopted heuristic approach to solve the TSP-D problem, drone-based systems have been proved to achieve 20 percent or more savings over the traditional TSP tours when drones are as fast as the truck and have long fly ranges (above 30 min) (Murray and Chu, 2015). Nevertheless, when considering distance range and time endurance limitations, together with other constraints, such as no-fly zones, the gap between becomes much narrower. In particular, denser and congested delivery areas of limited size offer the opportunity to robot-assisted systems to considerably improve their delivery performance with gains up to 30 percent. Furthermore, thanks to their fewer related safety and privacy issues, as compared to drones, robot-assisted deliveries might become a more viable delivery option, at least in urban settings. Retrofitting robots' storage with several compartments could play a crucial role in the overall efficiency of the system and does not seem particularly compelling. More challenging issues would entail the navigation (including real-time synchronization) and automation of the drop-off and pick-up process.

6. Conclusions

Delivery robots represent a novel opportunity for enhancing last-mile delivery in urban settings. The goal of robot-assisted truck delivery systems is to partially replace truck routes and increase the speed of the last-mile delivery process. Although several leaders in the supply-chain field, together with robotics start-ups and automakers, have been exploring the feasibility of such a solution, currently very little is known about its potential efficiency and implications. In this study, we investigate a robot-assisted truck delivery system where a single robot could depart from the truck to perform one or more deliveries by traveling on sidewalks.

First, the TSP with robot is defined and compared to similar problems in the literature, such as the truck-and-trailer problem and the TSP with drones. Since the problem is NP-hard, it can be solved exactly as a MILP only for small instances. For this reason, an efficient and accurate heuristic approach based on the concept of robot operations is proposed. The proposed approach adopts iterated local search with an adaptive perturbation (LS-AP) scheme to explore possible truck routes, which can be modified to replace some stops with robot deliveries. This sub-problem, which can be formulated as a special case of the Maximum Weighted Independent Set Problem, can be optimally solved by means of an efficient dynamic programming algorithm. The performance of the heuristic is evaluated for different instances by using the corresponding optimal solution and by checking the consistency of the results across different sizes of the problem (up to 100 customers). For small instances, the heuristic achieves high quality results (with gaps typically below 1 percent) with significant computational time savings.

Then, a systematic analysis of the system's efficiency based on different features of the robot and its operating environment is performed. The results indicate that the extent of achievable travel-time savings is not straightforward and strongly depends on factors such as the speed ratio between truck and robot, the capacity of the robot's storage, and the configuration of customers. Despite their low travel speeds, robots can yield considerable (time) efficiency gains when performing several consecutive deliveries and in the presence of traffic congestion. Interestingly, robot-assisted deliveries are particularly beneficial in the presence of a small portion of customers located in heavily congested areas. The maximum distance that can be covered by the robot is an important factor to consider for the overall efficiency, with significant limitations for ranges below 1 km in scenarios with 10 customers or more per km^2. To a minor extent, the duration of drop-off can affect the overall efficiency as well.

In comparison to drone-assisted deliveries (with drones with large flying ranges and speeds twice or three times that of trucks), robot-assisted ones are characterized by lower savings. However, when employing only a single drone with greater constraints in terms of speed and flying range, the benefits of the two systems are comparable (see Murray and Chu, 2015; Agatz et al., 2018). This is an interesting outcome to keep in mind when making considerations about commercial implementation of last-mile delivery services based on automated technologies.

The analyses performed in this study address a considerable set of questions, while leaving others for future research. First, the investigations were limited to relatively small scenarios (considering a few dozens of customers) given the small-scale application of this delivery solution. Some of the assumptions, such as the homogeneity of delivery operation length and could also be relaxed in order to increase the realism of the analyses in future studies. Second, in this study, travel time was used as a main performance indicator, while other types of costs, such as operation and maintenance costs (which depend on insurance, labor, and fuel), were not considered. The effects of some types of constraints concerning the accessibility of customers and energy constraints, which were not included in this first study, could be explicitly included in future studies. Since these issues, along with safety and regulation challenges (e.g. capability to operate in crowded environments) will likely play a role in the future adoption of this emerging technology, it would be interesting to address them in future research. Finally, from an optimization perspective, a natural extension of this work would consist of the development of heuristic approaches for the deployment of several robots in support of a single truck. Increasing the number of ancillary robots would significantly increase the complexity of the problem from both a practical and theoretical perspective (at least with the approach proposed in this study). An alternative deployment for multiple robots would consist of selecting fixed drop-off/pick-up points where the robots can swarm from and return to the truck.

CRediT authorship contribution statement

Michele D. Simoni: Conceptualization, Methodology. **Erhan Kutanoglu:** Conceptualization, Supervision. **Christian G. Claudel:** Conceptualization, Supervision.

Appendix A

See Appendix Table 5.

Table 5
Analysis of instances (results for Instance 17–68).

Instance	Average solution (γ)	Average savings (%)	Max savings (%)	Min savings (%)	Computation time (s)
17	3375.5	10.7	13.2	8.8	129.2
18	3045.5	17.6	17.6	17.6	122.7
19	3007.5	16.1	16.6	15.2	114.8
20	3125.2	16.1	16.4	15.3	138.7
21	3520.9	12.9	13.8	11.9	129.8
22	3570.3	10.8	10.8	10.8	118.3
23	3380.1	12.4	12.7	12.3	377.9
24	3348.2	11.4	11.4	11.4	126.9
25	3673.4	12.5	12.8	11.8	142.5
26	3618.5	13.0	13.0	13.0	109.0
27	3627.2	12.4	12.4	12.4	121.7
28	3690.8	11.9	12.7	11.1	115.7
29	4228.0	11.5	12.4	10.9	114.8
30	4348.5	10.9	10.9	10.9	131.1
31	4118.8	11.9	11.9	11.9	116.4
32	4256.3	11.2	11.2	11.2	102.8
33	6123.2	14.6	15.1	13.5	101.3
34	6134.6	14.4	16.1	13.1	131.7
35	6096.3	14.6	14.8	13.5	125.7
36	6036.8	14.7	15.6	14.2	674.6
37	6946.2	11.9	13.6	11.2	412.8
38	6817.6	13.5	15.2	12.1	410.0
39	6880.7	12.3	12.8	11.4	458.3
40	6979.2	11.2	11.2	11.2	417.1
41	6313.6	14.3	16.1	12.9	548.8
42	6619.1	11.3	11.6	10.9	327.3
43	6354.3	15	15	15	288.7
44	6632.6	13.3	13.8	12.7	558.7
45	7127.5	12.3	13.2	12	406.7
46	7066.8	12.6	13.1	12.4	225.5
47	7325.8	10.7	10.7	10.7	452.6
48	7359.0	10.3	11	9.8	230.6
49	9828.9	16.4	16.6	16.3	2250.0
50	9854.9	14.4	14.7	14.2	2968.1
51	9964.5	15.4	15.4	15.3	2550.7
52	9960.0	14.8	15	14.7	2188.0
53	10630.0	14.5	14.5	14.5	1714.5
54	10627.3	14.6	14.7	14.5	2010.6
55	10451.2	15.3	15.3	15.3	1218.8
56	10449.8	15.1	15.1	15.1	2324.9
57	11377.3	12.3	12.5	12.2	2572.7
58	11369.8	12.5	12.9	12.3	1992.2
59	11410.3	11.7	11.8	11.6	2562.2
60	11491.3	10.1	10.5	9.8	2443.0
61	20813.9	14.5	14.5	14.5	3015.1
62	21110.9	13.1	13.2	13	1581.3
63	20965.1	13.4	13.8	13.1	2027.2
64	21076.0	13.3	13.7	13	2667.6
65	21644.2	13.2	13.4	13.1	2980.2
66	21953.5	12.1	12.1	12.1	1811.2
67	24261.0	13.3	13.5	13.1	2032.4
68	21920.3	12.0	12.1	11.9	1574.7

Appendix B. Supplementary material

Supplementary data to this article can be found online at https://doi.org/10.1016/j.tre.2020.102049.

References

Agatz, N., Bouman, P., Schmidt, M., 2018. Optimization approaches for the traveling salesman problem with drone. Transport. Sci. 52 (4), 965–981.
Allen, J., Piecyk, M., Piotrowska, M., McLeod, F., Cherrett, T., Ghali, K., Wise, S., 2017. Understanding the impact of e-commerce on last-mile light goods vehicle activity in urban areas: the case of London. Transport. Res. D: Transp. Environ. 61, 325–338.
Benlic, U., Hao, J.K., 2013. A study of adaptive perturbation strategy for iterated local search. In: European Conference on Evolutionary Computation in Combinatorial Optimization (pp. 61-72). Springer, Berlin, Heidelberg.

Bishop, C., 2016. Swiss Post trials robot parcel deliveries in Bern. The Local. Retrieved at: https://www.thelocal.ch/20160823/swiss-post-trials-robot-parcel-deliveries-in-bern.

Boysen, N., Schwerdfeger, S., Weidinger, F., 2018. Scheduling last-mile deliveries with truck-based autonomous robots. Eur. J. Oper. Res. 271 (3), 1085–1099.

Burns, J., 2016. Domino's Pizza Robot Making Deliveries In Australia. Forbes. Retrieved from: https://www.forbes.com/sites/janetwburns/2016/03/18/dominos-pizza-robot-is-making-deliveries-in-australia/#2528686b7d59.

Carlsson, J.G., Song, S., 2017. Coordinated logistics with a truck and a drone. Manage. Sci. 64 (9), 4052–4069.

Chao, I.M., 2002. A tabu search method for the truck and trailer routing problem. Comput. Oper. Res. 29 (1), 33–51.

Coldewey D. (April, 2019). Kiwi's food delivery bots are rolling out to 12 more colleges. TechCrunch. Retrieved from: https://techcrunch.com/2019/04/25/kiwis-food-delivery-bots-are-rolling-out-to-12-new-colleges/.

Conway, A., Wang, X., Chen, Q., Schmid, J.,2016. Freight Costs at the Curbside. Final Report. Retrieved from: http://www.utrc2.org/sites/default/files/Final-Report-Freight-Costs-at-Curbside.pdf (Accessed November, 2018).

Cuda, R., Guastaroba, G., Speranza, M.G., 2015. A survey on two-echelon routing problems. Comput. Oper. Res. 55, 185–199.

Daimler, 2017. Vans & Robots. Small delivery robots out of the Sprinter. Retrieved from: https://www.mercedes-benz.com/en/mercedes-benz/vehicles/transporter/vans-robots-small-delivery-robots-out-of-the-sprinter/.

Derigs, U., Pullmann, M., Vogel, U., 2013. Truck and trailer routing—problems, heuristics and computational experience. Comput. Oper. Res. 40 (2), 536–546.

Desjardins, J. (March, 2018). Amazon and UPS are betting big on drone delivery. Business Insider. Retrieved from: http://www.businessinsider.com/amazon-and-ups-are-betting-big-on-drone-delivery-2018-3.

Drexl, M., 2012. Synchronization in vehicle routing—a survey of VRPs with multiple synchronization constraints. Transport. Sci. 46 (3), 297–316.

Drexl, M., 2014. Branch-and-cut algorithms for the vehicle routing problem with trailers and transshipments. Networks 63 (1), 119–133.

Ha, Q.M., Deville, Y., Pham, Q.D., Hà, M.H., 2018. On the min-cost traveling salesman problem with drone. Transport. Res. Part C: Emerg. Technol. 86, 597–621.

Halzack, S., 2016. Amazon makes its first drone delivery to a real customer. The Washington Post. Retrieved at: https://www.washingtonpost.com/news/business/wp/2016/12/14/amazon-makes-its-first-drone-delivery-to-a-real-customer/?utm_term=.c4c7027b4021.

Hern, A., 2014. DHL launches first commercial drone 'parcelcopter' delivery service. The Guardian. Retrieved at: https://www.theguardian.com/technology/2014/sep/25/german-dhl-launches-first-commercial-drone-delivery-service.

Jennings, D., Figliozzi, M., 2019. Study of Sidewalk Autonomous Delivery Robots and Their Potential Impacts on Freight Efficiency and Travel. Transp. Res. Rec.

Joerss M., Schroder J., Neuhaus F., Klink, C., Mann, F. (2016). Parcel delivery. The future of last mile. McKinsey&Company. September, 2016. Retrieved on May, 2018 from: https://www.mckinsey.com/~/media/mckinsey/industries/travel%20transport%20and%20logistics/our%20insights/how%20customer%20demands%20are%20reshaping%20last%20mile%20delivery/parcel_delivery_the_future_of_last_mile.ashx.

Kleinberg, J., Tardos, E., 2006. Algorithm design. Pearson.

Knoblauch, R., Pietrucha, M., Nitzburg, M., 1996. Field studies of pedestrian walking speed and start-up time. Transport. Res. Rec.: J. Transport. Res. Board 1538, 27–38.

Lin, S.W., Vincent, F.Y., Chou, S.Y., 2009. Solving the truck and trailer routing problem based on a simulated annealing heuristic. Comput. Oper. Res. 36 (5), 1683–1692.

Mercedes-Benz (n.d.). "Vans, robots, and the future of delivery." Retrieved on November, 2019 from: https://www.mercedes-benz.co.nz/vans/en/mercedes-benz-vans/love-your-work/future-of-delivery.

Murray, C.C., Chu, A.G., 2015. The flying sidekick traveling salesman problem: optimization of drone-assisted parcel delivery. Transport. Res. C: Emerg. Technol. 54, 86–109.

Nicas, J., 2018. Google's Parent Births New Businesses: Balloons and Drones. The New York Times. Retrieved on August, 2018 from: https://www.nytimes.com/2018/07/11/technology/google-drones-internet-balloons.html.

Otto, A., Agatz, N., Campbell, J., Golden, B., Pesch, E., 2018. Optimization approaches for civil applications of unmanned aerial vehicles (UAVs) or aerial drones: a survey. Networks 72 (4), 411–458.

Poeting, M., Schaudt, S., Clausen, U., 2019. Simulation of an optimized last-mile parcel delivery network involving delivery robots. In: Interdisciplinary Conference on Production, Logistics and Traffic, Springer, Cham, pp. 1–19.

Reed, T., Kidd, J., 2019. Global Traffic Scorecard. INRIX Research. Retrieved at: http://inrix.com/scorecard/.

Rodrigue, J.P., Comtois, C., Slack, B., 2009. The "Last Mile" in Freight Distribution. In: The Geography of Transport Systems, (2nd ed.). Routledge, pp. 212.

Savelsbergh, M., Van Woensel, T., 2016. 50th anniversary invited article – city logistics: challenges and opportunities. Transportation Science 50 (2), 579–590.

Scheuerer, S., 2006. A tabu search heuristic for the truck and trailer routing problem. Comput. Oper. Res. 33 (4), 894–909.

Sonneberg, M.O., Leyerer, M., Kleinschmidt, A., Knigge, F., Breitner, M.H., 2019. Autonomous Unmanned Ground Vehicles for Urban Logistics: Optimization of Last Mile Delivery Operations. In: Proceedings of the 52nd Hawaii International Conference on System Sciences.

Scott, S., 2019. Meet Scout: Field testing a new delivery system with Amazon Scout. The Amazon Blog. Retrieved at: https://blog.aboutamazon.com/transportation/meet-scout?linkCode=w50&tag=w050b-20&imprToken=ZeMIelf0qQWgVzt0u3nPig&slotNum=0.

Uber Movements, 2019. Uber Technologies, Inc. Data retrieved on 05-20-19 at: https://movement.uber.com.

Villegas, J.G., Prins, C., Prodhon, C., Medaglia, A.L., Velasco, N., 2011. A GRASP with evolutionary path relinking for the truck and trailer routing problem. Comput. Oper. Res. 38 (9), 1319–1334.

Villegas, J.G., Prins, C., Prodhon, C., Medaglia, A.L., Velasco, N., 2013. A matheuristic for the truck and trailer routing problem. Eur. J. Oper. Res. 230 (2), 231–244.

Wang, X., Poikonen, S., Golden, B., 2017. The vehicle routing problem with drones: several worst-case results. Optimiz. Lett. 11 (4), 679–697.