Physics-Aware Truck and Drone Delivery Planning Using Optimization and Machine Learning

Yineng Sun[†], Armin Fügenschuh[‡], and Vikrant Vaze[†]

[†]Thayer School of Engineering, Dartmouth College, Hanover, NH, USA
[‡]Brandenburgische Technische Universität Cottbus-Senftenberg, Cottbus, Germany

1 Introduction

The last two decades have seen rapid development of the e-commerce business. According to the US Census Bureau, total e-commerce sales for 2023 were more than a trillion dollars, up 7.6% from 2022. Optimization of e-commerce merchandise delivery logistics is increasingly crucial to improving profits and customer experience. Given the large number of packages that need to be delivered in a timely and cost-effective manner to individual e-commerce customers, even small improvements in delivery solution strategies could have a significant bottom-line impact.

Last-mile merchandise are typically delivered using trucks, each operating in a predefined urban, suburban, or rural region. With recent technological developments in unmanned aerial vehicles ("drones"), researchers have proposed new last mile delivery schemes, in which delivery trucks are paired with drones to assist with package delivery. Murray and Chu (2015) defined a mathematical formulation for the flying sidekick traveling salesman problems (FSTSP) for this set-up. Many studies have focused on variations of the FSTSP setup. In particular, Agatz et al. (2018) defined and studied the traveling salesman problem with drones (TSP-D), where, for a single truck and a single drone, delivery assignment and vehicle routing decisions are jointly optimized. This modeling strategy differs from the arc-based modeling approach commonly used in FSTSP, subdividing the entire truck-and-drone tour into segments called operations.

2 The Problem

We consider a system of one truck and one drone to deliver packages to customers in minimum total travel time. Both vehicles begin at a depot with the drone resting on top of the truck and packages loaded in the truck. Each delivery location corresponds to a single package. If multiple customers share a location or if a single customer has multiple deliveries, then we redefine the set of packages as a single package per delivery location. The truck driver can park at the delivery locations to deliver packages & the drone can fly to serve other locations. The drone can take off and land on top of the truck at any time, even when the truck is in motion. This feature helps the drone save energy and time. The drone visits at most one delivery location in each operation.

To estimate drone travel times, delivery logistics optimization studies simplify drone trajectory details. Integrating the complexities of the physics-based modeling of drone trajectories (Schmidt and Fügenschuh, 2023) into the TSP-D model can lead to a more accurate but complicated model. However, it is unclear to what extent such integration could provide practical benefits. In this paper, we show that the prevalent simplifications of drone flight physics made to enable integrating truck and drone operations in parcel delivery service planning can lead to operational plans that can be highly suboptimal and often infeasible. We introduce a new two-stage optimization to address these limitations and improve the quality of the solution while ensuring the feasibility of the truck and drone travel routes. Our modeling approach allows planning for drone trajectories in which the drone can take off from or land on the truck anywhere along the truck path. In this regard, we expand the original concept of drone "operation" as defined by Agatz et al. (2018).

As shown in previous studies, the base models for the integrated planning of truck and drone routing logistics is a difficult combinatorial optimization problem, even when simplifying the drone physics considerations. Therefore, a naive direct inclusion of accurate drone travel trajectory calculations into the truck-and-drone route planning model is unlikely to be computationally tractable for realistic problem sizes. We propose a new physics-aware truck and drone delivery logistics planning model and an original heuristic solution approach that combines optimization and supervised machine learning to generate high-quality solutions in limited computational runtimes for realistically sized problem instances. In addition, we incorporate restricted airspace (RAS) constraints in drone trajectories for additional realism in our computational experiments.

3 Model Formulation

The tour planning problem decomposes into two sets of decisions. The upper level assigns delivery nodes to the two vehicles (truck and drone) and decides the sequence of visits to the delivery nodes for each vehicle. The lower level decides the drone trajectory for each flight. The objective is to minimize the total truck-and-drone tour time. We break ties by minimizing drone energy costs as a secondary objective. The upper-level non-linear optimization formulation is an extension of the model by Agatz et al. (2018). Because the operation time is a variable instead of a constant, we have a non-linear term in the objective function. Upper-level model also includes all constraints in the Agatz et al. (2018) model, ensuring that all delivery nodes are visited, each operation must cover at least one delivery node that is not in other operations, the flow balance of the truck-and-drone tour is maintained, and the tour starts and ends at the depot. The lower-level model builds on the drone trajectory planning model of Schmidt and Fügenschuh (2023). By coordinating with the truck's travel path and velocity when taking off and landing, the lower-level model optimizes the drone trajectory. It has the following constraints.

- Drone trajectories obey rules of physics for displacement, velocity & acceleration updates.
- Drone flying altitude is in one of the discretized altitude bands stipulated by local regulation.
- The drone speed and acceleration levels are within technical specifications at all times.
- The drone cannot fly into any of the restricted airspaces as per the local regulations.
- An operation ends when both the truck and the drone finish their operation assignments.
- When resting on the truck, location and speed of the drone are the same as that of the truck.
- Drone deliveries can occur only if the drone visits a customer node and reaches zero speed.

- Approximations leveraging l_∞ and l_1 norms are used to linearize l_2 norms of displacement, velocity & acceleration.

4 Solution Approach

If the time and drone energy consumption for each operation is known beforehand, then one can directly solve the TSP-D model by Agatz et al. (2018). However, an accurate estimation of the drone travel time requires extensive computations to solve the lower-level formulation for each candidate operation, and the number of such operations grows exponentially with the instance size, making it unrealistic to perform a complete enumeration for practically sized instances. Instead, we propose a new neural network-based approach to provide accurate estimates of the time and drone energy consumption for each operation. Kundu et al. (2022) proposed a polynomial time split algorithm combined with local search techniques to solve the TSP-D problem and showed that overall approach is highly accurate and computationally scalable to realistic problem instances. Their approach relies on Euclidean travel distances, which we propose to replace with a machine-learning trained neural network (NN) predictor. Our NN model outputs a drone trajectory time estimate given a set of three node locations: a start node, a delivery node, and an end node, as well as the coordinates of the corners of restricted airspaces. Once the aforementioned modified heuristic generates the optimal set of operations, the detailed lower-level model provides a feasible 4D (x, y, z, and time) trajectory plan, drone travel time, and drone energy consumption cost for each operation.The input layer of the NN model uses features that represent drone travel locations within an operation, including the starting, delivery, and ending points. Specifically, we used the x and y coordinates of each of these three locations as features. The NN model's output is the drone travel time for an operation. We train the NN models on simulation-generated offline training data from a drone-only model for various (start, delivery, end) node triplets with and without RAS regions separately before applying them.

5 Results and Conclusion

We conduct extensive computational experiments on two sets of instances to test the performance of our approach. The first set (reported in Table 1) does not have any RAS regions, whereas the second (not reported here due to space restrictions) considers the existence of six RAS regions. For each set, we compare total travel time and drone energy consumption obtained by our approach (denoted by P) with that obtained by the benchmark approach (denoted by K) based on simplified Euclidean distance-based drone travel times. We also report results for a third approach (denoted by MK) which uses a constant multiplicative factor (calibrated using offline training data) to scale the Euclidean distance to estimate the drone travel time. We vary the number of delivery nodes from 10 to 200 and the average truck travel speed, measured in km per hour (kph), from 20 to 80. For each combination, we randomly generated 100 instances and report the average results in Table 1. Each value in this table reports the percentage saving in total travel time using Method P compared to Method K. To show how our approach performs in a real-world setting, we additionally applied it to four major city centers in the U.S. for a typical daily e-commerce delivery instance. Table 2 reports the savings in both total travel time and drone energy consumption, achieved by our Method P compared to Method K and Method MK. In all cases, we report the actual travel times and the actual values of drone energy consumption calculated using the trajectory model based on physics.

Speed \ Nodes	10	20	50	75	100	175	250
20	0.20%	0.48%	0.64%	1.07%	1.85%	4.11%	7.58%
30	-0.20%	1.41%	3.92%	5.71%	8.40%	15.15%	21.88%
40	2.19%	4.10%	11.16%	15.61%	19.10%	27.58%	34.28%
50	3.94%	9.86%	18.40%	23.35%	27.16%	34.89%	40.66%
60	5.89%	10.91%	19.68%	24.41%	28.30%	34.97%	40.07%
70	5.39%	9.48%	15.54%	18.99%	21.91%	27.12%	31.63%
80	3.79%	7.54%	13.67%	15.57%	18.24%	24.08%	28.32%

Table 1: Total Travel Time Reduction in Synthetic Instances: Method P vs. Method K

City Center	TTT (vs. K)	TTT (vs. MK)	DEC (vs. K)	DEC (vs. MK)
Los Angeles	4.38%	3.83%	14.21%	13.72%
Boston	3.72%	2.53%	13.73%	12.72%
Chicago	3.69%	2.41%	14.01%	12.85%
Philadelphia	3.15%	2.51%	11.58%	10.51%

Table 2: Total Travel Time (TTT) and Drone Energy Consumption (DEC) Saving for Method P

Method P consistently outperforms Method K in all but one of the 49 combinations in Table 1, with an average TTT savings of 15% (range 0%-41%). Method K also provides average DEC savings of 30% (range 2%-58%) [not shown here]. The relative benefits of Method P increase significantly with the number of customer nodes and also with the average truck speeds. For the four US city centers, Table 2 shows that Method K leads to TTT savings of 3.15%-4.38% compared to Method K (and 2.41%-3.83% compared to Method MK) and DEC savings of 11.58%-14.21% compared to Method K (and 10.51%-13.72% compared to Method MK). All models are solved in under a minute of runtime consistent with practical requirements. Thus, our neural network-based solution approach solves our physics-aware truck-and-drone delivery logistics planning model efficiently providing substantial and consistent travel time savings.

References

Agatz, N., Bouman, P., and Schmidt, M. (2018). Optimization approaches for the traveling salesman problem with drone. *Transportation Science*, 52(4):965–981.

Kundu, A., Escobar, R. G., and Matis, T. I. (2022). An efficient routing heuristic for a drone-assisted delivery problem. *IMA Journal of Management Mathematics*, 33(4):583–601.

Murray, C. C. and Chu, A. G. (2015). The flying sidekick traveling salesman problem: Optimization of drone-assisted parcel delivery. *Trans. Res. Part C: Emerging Tech.*, 54:86–109.

Schmidt, J. and Fügenschuh, A. (2023). A two-time-level model for mission and flight planning of an inhomogeneous fleet of unmanned aerial vehicles. *Comp, Opt. & App.*, 85(1):293–335.