A General Optimization Framework for Dynamic Two-Stage Order Fulfillment Problems

Gal Neria[1], Michal Tzur[1], Marlin W Ulmer[2]

[1] *Department of Industrial Engineering, Tel Aviv University, Tel Aviv, Israel*
[2] *Otto-von-Guericke Universität Magdeburg, Magdeburg, Germany*

Extended abstract submitted for presentation at the 12th Triennial Symposium on Transportation Analysis conference (TRISTAN XII)
June 22-27, 2025, Okinawa, Japan

March 31, 2025

Keywords: dynamic order fulfillment, real-time decision making, two stage optimization, machine learning.

1 INTRODUCTION

The rise of on-demand delivery services is transforming industries like warehousing, food services, and production, where timely delivery is critical for maintaining efficiency and customer satisfaction (Waßmuth *et al.*, 2023). This often means that capture and fulfillment phases overlap, i.e., orders come in at the same time as delivery takes place (Fleckenstein *et al.*, 2023). Success in these sectors hinges on seamless real-time coordination between two key stages: the first involves the in-house preparation, picking, or production of goods, and the second focuses on their timely delivery to the customers.

Several examples highlight the complexity of these challenges. In modern e-commerce systems, companies like Amazon or Walmart manage large warehouses where products are picked, packed, and dispatched to customers within a short time frame. Aligning order picking with delivery schedules is essential for ensuring consolidated, on-time shipments without delay and with limited routing cost. In meal delivery platforms, such as those operated by UberEats, kitchens must align the preparation of meals with the dispatch schedules of delivery drivers with the primary goal of on-time delivery without delay. Orders from multiple restaurants may need to be combined into a single delivery route, further complicating operations. Also, preparing meals too early risks freshness loss, while preparing them too late can cause delivery delays and dissatisfied customers. There is also an increasing development of on-demand delivery in production, e.g., for supplying construction sites with tailored materials right on time. Similarly, in the metalworking and packaging industries, aligning production with outbound distribution minimizes downtimes and ensures efficient resource utilization (Berghman *et al.*, 2023). The importance of synchronizing interdependent stages is also evident in other areas, e.g., in the preparation and distribution of healthcare products where pharmaceuticals, chemotherapy treatments, or radiopharmaceuticals require precise synchronization of production and distribution to ensure medicines reach patients without delays.

While these applications vary across industries, their general structure shows similarities. All of them consist of two strongly connected stages, a preparation stage and a delivery stage, with challenging combinatorial decisions to make in both stages in real-time. All of them have constraining first-stage resources and the cost mainly on the second stage, the delivery (the cost for delivery routing or the customer inconvenience with respect to expected time commitments). Finally, all of them require effective dynamic decision-making with respect to orders being placed sequentially over time. This means keeping both preparation and delivery resources flexible with respect to future orders and decisions.

In this work, we introduce the Dynamic Two-Stage Order Fulfillment Problem (DTS-OFP), a novel class that unifies traditionally distinct logistical challenges (e.g., warehousing, meal delivery, production) under a single cohesive framework. The significance of this class lies in its ability to model both preparation and delivery stages simultaneously, under dynamic and stochastic conditions — an area not sufficiently addressed in the literature yet.

For effective decision-making, we introduce the innovative Two-Stage Heuristic with Value Function Approximation (TSH-VFA) framework. The framework is grounded in an analytical theorem that demonstrates how, under certain conditions, optimizing the second stage, i.e., the stage contributing to the objective function, is sufficient for overall system performance. For this stage, we develop a Large Neighborhood Search (LNS) framework that searches for effective and efficient delivery routes. By focusing only on the second stage in LNS, the decision space is reduced and the search is more successful. At the same time, first-stage feasibility is considered by means of fast heuristic algorithms or, alternatively, machine learning predictions. In cases where feasibility can be ensured, a fast algorithm creates the full decision for both stages. This algorithm is designed in a way that it focuses on flexible decisions for both stages. The flexibility of a full decision is then determined via VFA, estimating the expected future cost when the decision is taken.

In a meta-analysis, we specify and implement TSH-VFA for three selected problems, derived from the three major areas of warehousing, meal delivery, and production, each coming with its own, domain-specific constraints and objectives. For each problem, we specify the required components of TSH-VFA, the LNS operators, the feasibility check, and the algorithm to transfer an LNS solution for the second stage to a full problem decision. For the VFA, we rely on the work of Neria *et al.* (2024). We compare our method to several benchmark policies from the literature. We show that our method provides the best performance by far and that a targeted search on the second stage is superior to alternatives searching the entire decision space at once. For each component of TSH-VFA, we further present an alternative to illustrate their value and functionality: We present an LNS with standard operators and show the value of problem-specific operators. For feasibility, we present a prototypical check based on supervised learning (SL) related to van der Hagen *et al.* (2024). To the best of our knowledge, this is the first time that SL is used in an LNS to predict feasibility. We show that it speeds up runtime without significant loss in quality. For the transfer to a full solution, we present an alternative, myopic algorithm to highlight that while optimization should focus on the second stage, flexibility must be guaranteed over both stages.

2 Problem Definition

In a typical DTS-OFP setting, orders arrive dynamically over time and must go through a preparation stage by resources (such as machines or pickers) before being delivered to customers. The preparation and delivery stages are highly interdependent: each order must complete preparation before delivery can commence, and resources and vehicles must remain flexible for future orders where other synchronization constraints might also apply. To model the DTS-OFP, we use a Markov Decision Process (MDP). Decisions are made at key points when new orders arrive, and the goal is to update the schedule for resources and vehicle trips in response to the new information.

The system's state, S_k, includes information about the orders, resource schedules, and vehicle schedules. Each decision, x_k, updates this state while satisfying feasibility constraints, such as ensuring that resources complete order preparation before dispatch and vehicles meet capacity limits. A policy, denoted by $\pi \in \Pi$, is a mapping that assigns a decision $\pi(S_k) = x_k$ to each state S_k. Starting from an initial state S_0, the goal is to determine an optimal policy π^* that minimizes the expected total cost. This objective is equivalent to minimizing the expected sum of immediate costs across all decision points, beginning from S_0 and consistently applying policy π throughout the process:

$$\min_{\pi \in \Pi} E\left[\sum_{k=0}^{K} C(S_k, \pi(S_k)) \mid S_0\right], \tag{1}$$

where $C(S_k, x_k)$ is the variant specific immediate cost associated with state S_k, and decision x_k.

We model three specific variants within the DTS-OFP framework. The Warehousing Variant (following D'Haen *et al.* (2024)) in which orders are picked from a warehouse and delivered under a combined travel-time and delay-minimizing objective. The challenge lies in coordinating picking resources with delivery schedules to reduce delays and travel costs. The Meal Delivery Variant (following Neria *et al.* (2024)) in which orders from ghost kitchens must meet freshness constraints during delivery. This variant introduces synchronization challenges due to freshness requirements, necessitating fast coordination between meal preparation and dispatch to minimize tardiness. Finally, the Production Variant (following Wu *et al.* (2022)) involves manufacturing orders with setup times between tasks. The goal is to minimize violations of delivery time windows, considering varying product setups and time-window constraints.

3 METHODOLOGY

The TSH-VFA framework provides a structured approach to addressing DTS-OFP challenges by decomposing the problem into two optimization stages.This decoupled approach allows for efficient exploration of delivery scheduling options by searching through partial trip decisions that include trip sequences and departure times rather than searching the entire decision space. A feasibility check ensures that the preparation stage can accommodate the selected delivery sequence. If a sequence is feasible, the framework generates a full decision that includes both preparation and delivery stages. Our method leverages a neural network-based VFA to estimate the long-term cost implications of each decision. The VFA guides the search process by providing fast estimations of future costs, which enables the framework to select decisions that maintain system flexibility. We next detail the components of the TSH-VFA.

Based on a theoretical result that shows that it is sufficient to optimize the second stage, an LNS explores potential second-stage decisions to improve solution quality by modifying trips and their departures sequence. It modifies the order assignments within trips and generates potential delivery sequences while ensuring that each trip is independently feasible. Then, another algorithm refines the partial trips created by LNS by determining specific vehicle assignments and departure times. It also verifies feasibility using fast heuristic checks. An additional algorithm assigns preparation resources to orders based on the generated schedule. If a feasible preparation sequence can be arranged, it completes the full decision. Finally, a VFA evaluates each decision by estimating its future impact, allowing for anticipatory optimization that maintains system adaptability.

4 RESULTS

Our computational study uses three real world datasets from Meituan (2024), Iowa City restaurants, and ORTEC (2022). We test the TSH-VFA framework on each domain, evaluating its performance against four benchmark methods: FIFO (first-in-first-out), Integrated (optimizing both stages simultaneously, D'Haen *et al.* (2024)), AI (optimizing both stages simultaneously and using VFA, Neria *et al.* (2024)), and TSH-G (equivalent to TSH-VFA without VFA).

The results show that TSH-VFA consistently outperforms these benchmarks, achieving significant improvements in minimizing costs across all problem variants. TSH-VFA's ability to balance efficiency in the delivery stage while ensuring preparation feasibility is evident across all settings. In fact, TSH-VFA delivers the largest improvements over FIFO, which lacks the coordinated optimization of both stages, followed by substantial improvements over Integrated and AI. The superiority of TSH-G over Integrated demonstrates already the effectiveness of focusing on the second-stage optimization. The TSH-VFA improvement over TSH-G, which lacks VFA, demonstrates the importance of anticipatory decision-making in dynamic stochastic systems.

Each problem variant provides insights into how TSH-VFA adapts to different settings: In the Warehousing Variant we observe that TSH-VFA effectively reduces travel and delay costs by balancing the importance of timely delivery against the cost of vehicle routing. In the Meal Delivery Variant TSH-VFA's adaptive scheduling achieves up to 60% improvement in delay while meeting freshness constraints over benchmarks, particularly in large instances. This variant highlights the framework's capability to synchronize preparation and delivery. For the Production Variant the framework significantly reduces delivery time-window violations and minimizes average setup times in production, enabling more efficient resource utilization in manufacturing.

We also made methodological analysis. For example, we introduced a novel machine learning-based feasibility filter to reduce computational time. The filter identified infeasible solutions early, reducing runtime by over 80% without sacrificing solution quality. This demonstrates the potential of ML-based filters for streamlining large-scale scheduling problems.

References

Berghman, Lotte, Kergosien, Yannick, & Billaut, Jean-Charles. 2023. A review on integrated scheduling and outbound vehicle routing problems. *European Journal of Operational Research*, **311**(1), 1–23.

D'Haen, Ruben, Ramaekers, Katrien, Archetti, Claudia, & Braekers, Kris. 2024. Integrating order picking and vehicle routing decisions in a dynamic e-commerce setting. *Computers & Operations Research*, 106762.

Fleckenstein, David, Klein, Robert, & Steinhardt, Claudius. 2023. Recent advances in integrating demand management and vehicle routing: A methodological review. *European Journal of Operational Research*, **306**(2), 499–518.

Meituan. 2024. *TSL-Meituan*. INFORMS TSL Data-Driven Research Challenge.

Neria, Gal, Hildebrandt, Florentin D, Tzur, Michal, & Ulmer, Marlin W. 2024. The restaurant meal delivery problem with ghost kitchens. *Transportation Science, forthcoming.*

ORTEC. 2022. *EURO NEURIPS VRP 2022 Quickstart Instances*. `https://github.com/ortec/euro-neurips-vrp-2022-quickstart/tree/main/instances`. Accessed: July 31, 2024.

van der Hagen, Liana, Agatz, Niels, Spliet, Remy, Visser, Thomas R, & Kok, Leendert. 2024. Machine learning–based feasibility checks for dynamic time slot management. *Transportation Science*, **58**(1), 94–109.

Waßmuth, Katrin, Köhler, Charlotte, Agatz, Niels, & Fleischmann, Moritz. 2023. Demand management for attended home delivery—A literature review. *European Journal of Operational Research*, **311**(3), 801–815.

Wu, Gen-Han, Cheng, Chen-Yang, Pourhejazy, Pourya, & Fang, Bai-Lyn. 2022. Variable neighborhood-based cuckoo search for production routing with time window and setup times. *Applied Soft Computing*, **125**, 109191.