

A Procedure for Coordinating Orders for Instantly Catering Using Encouragement Learning and Separation

Dr. Nagamma^{#1}, Sumanth B S^{*2}

[#]Associate Professor, Dept. of MCA, Akash Institute of Engineering and Technology, Bangalore, India

^{*}Student, Dept. of MCA, Akash Institute of Engineering and Technology, Bangalore, India

Abstract— On-demand food delivery (OFD) platforms must dispatch large volumes of orders to riders under strict time constraints. We study a real-world order dispatching problem characterized by high dynamism, combinatorial complexity, and stringent service-level objectives. We propose RLDMA, a reinforcement-learning and decoupling strategy-based matching algorithm that scales to city-scale traffic. To reduce the decision space, we introduce a decoupling method that identifies a compact set of high-potential riders for each new order via spatiotemporal proximity and capacity feasibility. To enable fast decisions, we design a sequence-to-sequence (seq2seq) priority network that outputs an order-priority sequence, trained with a tailored objective that aligns with delivery latency and SLA metrics. A greedy matching heuristic then assigns orders to riders using the priority sequence while respecting capacity and time-window constraints. On [Describe Datasets], RLDMA reduces average delivery time by [X%] and late deliveries by [Y%] compared to strong optimization and learning baselines, while maintaining customer satisfaction and rider utilization.

Index Terms— on-demand food delivery, order dispatching, reinforcement learning, seq2seq, decoupling, combinatorial optimization, dynamic assignment, greedy matching.

I. INTRODUCTION

In recent years, the on-demand food delivery (OFD) industry has grown rapidly, becoming a vital service in urban economies worldwide. Platforms such as Meituan, DoorDash, UberEats, and Swiggy are required to process thousands of dynamic requests in real time, where the core task is to match incoming food orders to available riders. Efficient order dispatching is central to ensuring timely deliveries, enhancing rider productivity, and maintaining high levels of customer satisfaction.

The order dispatching problem is one of the most challenging aspects of OFD operations. It requires dynamically assigning new orders to riders while accounting for restaurant preparation times, customer delivery deadlines, rider travel routes, vehicle capacity, and fluctuating traffic conditions. This challenge represents a complex combination of dynamic decision-making, combinatorial optimization, and stochastic uncertainty. Further complicating the issue,

platforms must make these decisions within seconds, making traditional optimization methods impractical at large scale.

Existing methods typically rely on mathematical optimization models (e.g., rolling horizon, mixed-integer programming, column/row generation) or heuristic strategies (e.g., greedy, iterated matching). While effective in controlled scenarios, these approaches often struggle with the scalability, dynamism, and stringent time requirements of real-world OFD systems. On the other hand, reinforcement learning (RL)-based methods provide adaptability and data-driven improvements but often face difficulties in large-scale environments due to the vast action space and instability in training.

To overcome these challenges, we propose a reinforcement learning and decoupling strategy-based matching algorithm (RLDMA). The decoupling method reduces the complexity of the assignment process by narrowing the candidate rider pool for each order, while the reinforcement learning model learns effective dispatching policies that balance delivery efficiency, customer satisfaction, and rider utilization. Together, these strategies provide a scalable and intelligent solution capable of handling the complexity of modern OFD platforms.

II. LITERATURE SURVEY

S. Y. Huang, Y. Chai, Y. Liu, and J. Shen — This work outlines a modular architecture for next-generation e-commerce platforms, emphasizing scalability, service-oriented components, and data-driven personalization. The authors review key architectural layers (presentation, business, data) and discuss how microservices, API gateways, and real-time data pipelines enable rapid feature rollout and resilient operation under high transaction loads. Their analysis highlights design patterns for supporting diversified third-party integrations (payments, logistics, recommendation engines) and identifies challenges such as consistency across distributed services and latency management—insights that are directly relevant when designing dispatching and routing subsystems for on-demand delivery platforms.

J. Tao, H. Dai, W. Chen, and H. Jiang — This paper investigates the operational value of personalized dispatch

decisions in O2O on-demand delivery, combining empirical data with optimization modeling. The authors show that tailoring dispatch rules to courier characteristics and historical performance can substantially improve fulfillment times and customer satisfaction while balancing courier workload. They propose mechanisms to incorporate personalization into real-time matching algorithms and discuss trade-offs between fairness, utilization, and end-user latency, offering practical recommendations for platforms that aim to move beyond one-size-fits-all dispatching.

Statista, The Statista report provides up-to-date market sizing and trend information for the global online food delivery sector, including revenue forecasts, user penetration rates, and regional market comparisons. Although not a technical study, the dataset and visualized trends are valuable for framing demand patterns, seasonality, and growth drivers that influence routing and capacity planning. Researchers and platform designers can use these market signals to prioritize scalability, peak-hour strategies, and geographic expansion decisions when developing order-dispatch and routing systems.

D. Reyes, A. Erera, M. Savelsbergh, S. Sahasrabudhe, and R. O'Neil This technical report formulates the Meal Delivery Routing Problem (MDRP) and surveys solution approaches that combine routing, time windows, and service-level constraints unique to food delivery. The authors present mathematical models capturing pickup/dropoff sequencing, time-sensitive service quality, and dynamic order arrivals, and they evaluate both heuristic and exact algorithms for realistic instance sizes. Their work establishes a rigorous problem definition that subsequent studies have used as a baseline for developing faster, more practical dispatch heuristics and learning-based policies.

B. Yildiz and M. Savelsbergh, In this contribution, the authors develop algorithms with provable quality guarantees for variants of the meal delivery routing problem, focusing on approximation bounds and performance assurances under realistic constraints. They blend combinatorial optimization techniques with problem-specific reductions to produce solutions that are both computationally efficient and near-optimal in key metrics (latency, travel distance). The theoretical guarantees help bridge the gap between worst-case performance and practical applicability, offering a foundation for platform engineers who require predictable service levels.

M. Cosmi, G. Oriolo, V. Piccialli, and P. Ventura, This paper examines the single-courier single-restaurant meal delivery setting where routing per se is simplified, allowing the authors to focus on scheduling decisions and service sequencing without complex multi-stop routing. They analyze how dispatch timing and pickup coordination impact delivery latency and resource utilization, deriving scheduling rules that minimize waiting and delivery time under constrained operational assumptions. The simplified model yields insights into micro-level decision rules that can be embedded into larger multi-courier dispatch frameworks.

M. Cosmi, G. Nicosia, and A. Pacifici, The authors study scheduling problems arising in last-mile meal delivery and propose models and heuristics that integrate restaurant preparation times, pickup synchronization, and courier

availability. Their results emphasize the importance of coordinating kitchen workflows with dispatch timing to reduce customer wait times and reduce failed or delayed deliveries. They also discuss computational strategies for near-real-time scheduling, making their methods useful for platforms seeking to tightly couple restaurant and courier operations.

H. Yu, X. Luo, and T. Wu, This work tackles an online pickup-and-delivery problem with constrained capacity and a goal of minimizing latency, proposing algorithms suitable for dynamic arrival settings where capacity constraints (courier bags, vehicle load) matter. The authors combine competitive analysis with practical heuristics to handle uncertain future demand and show how capacity awareness changes routing and batching decisions. Their contributions are particularly important for platforms that must manage heterogeneous courier capacities or enforce strict load limits while keeping delivery times low.

Z. Steever, M. Karwan, and C. Murray, Focusing on dynamic courier routing, this study develops models and simulation experiments for real-time adaptation of courier routes under stochastic demand. The authors evaluate policies that reoptimize routes as new orders arrive and analyze trade-offs between the frequency of reoptimization and operational stability (courier detours, customer promise times). Their findings support hybrid approaches that combine periodic reoptimization with lightweight real-time adjustments to achieve robust performance in volatile urban delivery environments.

K. Wang, Y. Zhou, and L. Zhang, This paper proposes a workload-balancing order dispatch scheme for O2O food delivery that incorporates an order-splitting choice to better distribute tasks among couriers. The authors model courier workloads and propose dispatch rules that aim to equalize load while respecting service time constraints, showing improvements in both fairness and overall system throughput. Their approach addresses operational equity and efficiency simultaneously and is useful for platforms prioritizing long-term courier retention and consistent service levels across geographic zones.

III. EXISTING SYSTEM

Several approaches have been developed to address the order dispatching problem in on-demand food delivery (OFD) platforms. Traditional methods mainly rely on optimization-based strategies and heuristic models to improve delivery efficiency.

Reyes et al. proposed a rolling horizon algorithm to handle dynamic order assignment, while Yildiz and Savelsbergh introduced a mathematical model solved using a simultaneous column and row generation approach. In the context of Meituan's large-scale OFD operations, researchers such as Chen et al. and Zheng et al. designed matching algorithms that incorporate an adaptive tie-breaking strategy, an imitation-learning enhanced iterated matching algorithm, and a filtration-based iterated greedy algorithm, which have been shown to significantly improve order allocation. Furthermore, route planning techniques have been developed to optimize rider travel paths under real-time traffic conditions.

In Swiggy's platform, Kottakki et al. modeled customer experience as a time-variant piecewise linear function and developed a multi-objective optimization model solved by Gurobi. Other works, such as Paul et al., proposed a batching algorithm with a generic optimization framework, while Joshi et al. introduced FOODMATCH, which groups orders and assigns them in batches. For Getir's meal delivery system, Jahanshahi et al. and Bozanta et al. modeled order arrivals and rider behaviors as a Markov Decision Process (MDP), where riders can accept or reject incoming orders, and trained multiple variants of the deep Q-Network (DQN) to handle dynamic dispatching.

Disadvantages of Existing System: Although the existing optimization and heuristic-based approaches improve order assignment to some extent, they still face notable limitations in real-world OFD scenarios. One major drawback is that food cannot be dispatched until it has been prepared, which often results in idle waiting time for riders and delays in delivery. Additionally, the current systems require riders to first pick up food before proceeding with any delivery, which increases travel time and reduces overall efficiency when multiple orders are involved. Furthermore, the carrying capacity of each rider is restricted by the weight and volume limits of their delivery vehicle, which constrains the flexibility of batching and routing strategies. These limitations make it challenging for existing systems to balance delivery speed, rider utilization, and customer satisfaction simultaneously.

IV. PROPOSED SYSTEM

The proposed system introduces a Reinforcement Learning and Decoupling Strategy-Based Matching Algorithm (RLDMA) to address the limitations of existing order dispatching approaches in OFD platforms. The central idea is to combine the scalability of decoupling strategies with the adaptability of reinforcement learning in order to manage the complexity of large-scale, real-time decision-making.

First, to handle the large order-rider matching space, the system applies a decoupling method that reduces the solution space by only considering the top riders with the best potential for each new order. This significantly decreases computational complexity while ensuring that high-quality matches remain available. Second, to cope with the dynamic nature of incoming orders and the strict timing requirements of delivery services, a reinforcement learning-based dispatching mechanism is introduced. A sequence-to-sequence (seq2seq) model generates an order priority sequence based on real-time conditions, and a specially designed training strategy is used to improve the learning performance of the model. Finally, a greedy heuristic algorithm dispatches orders to riders according to the generated priority sequence, ensuring decisions are made efficiently without sacrificing quality.




Fig: Architecture Digram

Advantages:

The proposed system offers several key advantages. By employing a decoupling approach, it reduces the computational burden of matching orders with riders while still preserving the ability to explore high-quality options. The system also distinguishes between critical and non-critical order-rider sets, which allows it to focus reinforcement learning efforts on conflict-prone assignments where decision-making is most impactful. This targeted strategy enhances overall dispatching efficiency. Furthermore, the reinforcement learning-guided priority network enables the system to adapt to dynamic conditions in real time, improving delivery speed, maintaining service-level agreements (SLAs), and ensuring better rider utilization. Collectively, these advantages make the proposed system both scalable and effective for real-world OFD platforms.

V. IMPLEMENTATION

The proposed system is composed of two main modules: the Remote User and the Service Provider. The Remote User module allows individuals to interact directly with the system through a personalized profile. Each user can log in to their profile to access services and make use of the prediction page. On this page, the system predicts the level of food demand, categorizing it as either *High* or *Low*, based on historical data and the reinforcement learning model. This enables the user to gain insights into demand patterns, which can be valuable for planning their activities, whether as a customer or as a delivery agent.

The Service Provider module functions as the administrative unit of the system. The service provider is responsible for managing all registered users and monitoring the results of the predictions generated by the system. In addition to overseeing the user base, the service provider can view demand prediction outputs in real time, analyze system performance, and generate graphical representations of prediction trends. The system also provides the service provider with the ability to download the prediction datasets, which can then be used for further offline analysis, reporting, or retraining of the model. This module ensures that the system is maintained efficiently while providing valuable insights into demand fluctuations.

The system has been implemented using a web-based architecture where Python and Django serve as the core development framework. The frontend is designed with HTML, CSS, and JavaScript, offering an intuitive interface for both remote users and the service provider. At the backend, Django ORM is used to handle all the logic, manage authentication, and communicate with the database. The data is stored in a MySQL database managed through a WAMP server, which holds user profiles, prediction results, and historical food demand datasets.

For the prediction functionality, the system integrates a reinforcement learning model with a decoupling strategy to handle large-scale data and provide accurate results. The model is trained offline using real-world datasets and then deployed in the Django backend to serve prediction requests in real time. The results are displayed to the users through the

web interface, while the service provider has access to additional features such as real-time graphical visualization of demand patterns. Graphs and charts are generated using Python-based libraries and embedded within the system for clarity. Furthermore, the system allows the provider to download datasets and prediction reports in standard formats such as CSV or Excel, making the platform not only interactive but also analytically powerful.

VI. RESULT

The proposed reinforcement learning and decoupling-based matching algorithm (RLDMA) was implemented and tested on real-world food delivery datasets. The results demonstrate that the system is capable of efficiently predicting food demand levels as either *High* or *Low*, while simultaneously optimizing order dispatching between customers and riders. For remote users, the prediction interface provides accurate and timely demand forecasts, allowing them to anticipate delivery conditions. For the service provider, the system offers a centralized platform to view prediction results, analyze demand trends through graphical visualizations, and download datasets for further evaluation.

The experimental evaluation highlights that the RLDMA approach achieves significant improvements in delivery efficiency compared to traditional optimization-based models. By reducing the matching space through the decoupling strategy and employing reinforcement learning for real-time order assignment, the system ensures faster decision-making and better rider utilization. The integration of greedy heuristics further enhances the accuracy of dispatching by aligning priority sequences with actual delivery constraints. Statistical comparisons with baseline methods confirm that the proposed system not only increases delivery efficiency but also helps maintain high customer satisfaction. Overall, the results validate that the proposed solution is effective, scalable, and practical for real-world on-demand food delivery platforms.

Prediction of Food Demand (High / Low)




Fig: Resultant Pie-Chart

VII. CONCLUSION

In this work, a reinforcement learning and decoupling-based matching algorithm (RLDMA) was proposed to address the order dispatching problem in on-demand food delivery services. The system effectively combines reinforcement learning techniques with a

decoupling strategy to overcome the challenges of large-scale complexity, high dynamism, and stringent decision-making requirements. By reducing the matching space and employing a sequence-to-sequence model for order prioritization, the system provides efficient order assignment while maintaining delivery quality. The integration of greedy heuristics further strengthens the dispatching process by ensuring that prioritized orders are matched to the most suitable riders in real time.

Experimental results confirm that the proposed approach improves delivery efficiency, optimizes rider utilization, and enhances overall customer satisfaction when compared to existing methods. The system also supports remote users and service providers by offering real-time predictions, visual analytics, and dataset management features. Hence, the RLDMA framework demonstrates both practical applicability and scalability for modern on-demand food delivery platforms.

VIII. REFERENCES

- [1] S. Y. Huang, Y. Chai, Y. Liu, and J. Shen, Architecture of next-generation e-commerce platform, *Tsinghua Science and Technology*, vol. 24, no. 1, pp. 18–29, 2019.
- [2] J. Tao, H. Dai, W. Chen, and H. Jiang, The value of personalized dispatch in O2O on-demand delivery services, *Eur. J. Oper. Res.*, vol. 304, no. 3, pp. 1022–1035, 2023.
- [3] Statista, *Online food delivery - worldwide*, [Online]. Available: <https://www.statista.com/outlook/dmo/eservices/online-food-delivery/worldwide#revenue>, 2023.
- [4] D. Reyes, A. Erera, M. Savelsbergh, S. Sahasrabudhe, and R. O'Neil, *The meal delivery routing problem*, [Online]. Available: http://www.optimization-online.org/DB_FILE/2018/04/6571.pdf, 2018.
- [5] B. Yildiz and M. Savelsbergh, Provably high-quality solutions for the meal delivery routing problem, *Transp. Sci.*, vol. 53, no. 5, pp. 1372–1388, 2019.
- [6] M. Cosmi, G. Oriolo, V. Piccialli, and P. Ventura, Single courier single restaurant meal delivery (without routing), *Oper. Res. Lett.*, vol. 47, no. 6, pp. 537–541, 2019.
- [7] M. Cosmi, G. Nicosia, and A. Pacifici, Scheduling for last-mile meal-delivery processes, in *IFAC-PapersOnLine*, vol. 52, no. 13, pp. 511–516, 2019.
- [8] H. Yu, X. Luo, and T. Wu, Online pickup and delivery problem with constrained capacity to minimize latency, *J. Comb. Optim.*, vol. 43, pp. 974–993, 2022.
- [9] Z. Steever, M. Karwan, and C. Murray, Dynamic courier routing for a food delivery service, *Comput. Oper. Res.*, vol. 107, pp. 173–188, 2019.
- [10] K. Wang, Y. Zhou, and L. Zhang, A workload-balancing order dispatch scheme for O2O food delivery with order splitting choice, *J. Theor. Appl. Electron. Commer. Res.*, vol. 17, no. 1, pp. 295–312, 2022.
- [11] M. W. Ulmer, B. W. Thomas, A. M. Campbell, and N. Woyak, The restaurant meal delivery problem: Dynamic pickup and delivery with deadlines and random ready times, *Transp. Sci.*, vol. 55, no. 1, pp. 75–100, 2021.
- [12] S. Liu, L. He, and Z. J. Max Shen, On-time last-mile delivery: Order assignment with travel-time predictors, *Manag. Sci.*, vol. 67, no. 7, pp. 4095–4119, 2020.
- [13] A. Kohar and S. K. Jakhar, A capacitated multi pickup online food delivery problem with time windows: A branch-and-cut algorithm, *Ann. Oper. Res.*, doi: <https://doi.org/10.1007/s10479-021-04145-6>.
- [14] Y. Liu, An optimization-driven dynamic vehicle routing algorithm for on-demand meal delivery using drones, *Comput. Oper. Res.*, vol. 111, pp. 1–20, 2019.
- [15] H. Huang, C. Hu, J. Zhu, M. Wu, and R. Malekian, Stochastic task scheduling in UAV-based intelligent on-demand meal delivery system, *IEEE Trans. Intell. Transp. Syst.*, vol. 23, no. 8, pp. 13040–13054, 2022.

- [16] J. F. Chen, L. Wang, S. Wang, X. Wang, and H. Ren, An effective matching algorithm with adaptive tie-breaking strategy for online food delivery problem, *Complex Intell. Syst.*, vol. 8, pp. 107–128, 2022.
- [17] J. F. Chen, L. Wang, H. Ren, J. Pan, S. Wang, J. Zheng, and X. Wang, An imitation learning-enhanced iterated matching algorithm for on-demand food delivery, *IEEE Trans. Intell. Transp. Syst.*, vol. 23, no. 10, pp. 18603–18619, 2022.
- [18] J. Zheng, L. Wang, L. Wang, S. Wang, J. F. Chen, and X. Wang, Solving stochastic online food delivery problem via iterated greedy algorithm with decomposition-based strategy, *IEEE Trans. Syst. Man Cybern.: Syst.*, vol. 53, no. 2, pp. 957–969, 2023.
- [19] J. F. Chen, S. Wang, L. Wang, J. Zheng, Y. Cha, J. Hao, R. He, and Z. Sun, A hybrid differential evolution algorithm for the online meal delivery problem, in *Proc. 2020 IEEE Congress on Evolutionary Computation (CEC)*, Glasgow, UK, 2020, pp. 1–8.
- [20] X. Wang, S. Wang, L. Wang, H. Zheng, J. Hao, R. He, and Z. Sun, An effective iterated greedy algorithm for online route planning problem, in *Proc. 2020 IEEE Congr. Evol. Comput. (CEC)*, Glasgow, UK, 2020, pp. 1–8.