Real-time Demand Forecasting for an Urban Delivery Platform

Alexander Hess[a,1], Stefan Spinler[a,1], Matthias Winkenbach[b,1]

[a] *WHU - Otto Beisheim School of Management, Burgplatz 2, 56179 Vallendar, Germany*
[b] *Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, MA 02139, United States*

Abstract

Meal delivery platforms like Uber Eats shape the landscape in cities around the world. This paper addresses forecasting demand into the short-term future. We propose an approach incorporating both classical forecasting and machine learning methods. Model evaluation and selection is adapted to demand typical for such a platform (i.e., intermittent with a double-seasonal pattern). The results of an empirical study with a European meal delivery service show that machine learning models become competitive once the average daily demand passes a threshold. As a main contribution, the paper explains how a forecasting system must be set up to enable predictive routing.

Keywords: demand forecasting, intermittent demand, machine learning, urban delivery platform

---

[1]Emails: alexander.hess@whu.edu, stefan.spinler@whu.edu, mwinkenb@mit.edu

1. Introduction

In recent years, many meal delivery platform providers (e.g., Uber Eats, GrubHub, DoorDash, Deliveroo) with different kinds of business models have entered the markets in cities around the world. A study by [5] estimates the global market size to surpass 20 billion Dollars by 2025. A common feature of these platforms is that they do not operate kitchens but focus on marketing their partner restaurants' meals, unifying all order related processes in simple smartphone apps, and managing the delivery via a fleet of either employees or crowd-sourced sub-contractors.

Various kind of urban delivery platforms (UDP) have received attention in recent scholarly publications. [6] look into heuristics to simultaneously optimize courier scheduling and routing in general, while [8] do so for the popular dial-a-ride problem and [9] investigate the effect of different fulfillment strategies in the context of urban meal delivery. [4] and [1] focus their research on the routing aspect, which is commonly modeled as a so-called vehicle routing problem (VRP).

Not covered in the recent literature is research focusing on the demand forecasting problem a UDP faces. Due to the customers' fragmented locations and the majority of the orders occurring ad-hoc for immediate delivery in the case of a meal delivery platform, forecasting demand for the near future (i.e., several hours) and distinct locations of the city in real-time is an essential factor in achieving timely fulfillment. In general, demand forecasting is a well-researched discipline with a decades-long history in scholarly journals as summarized, for example, by [3]. Even some meal delivery platforms themselves publish their practices: For example, [2] provide a general overview of supply and demand forecasting at Uber and benchmarks of the methods used while [7] investigate how extreme events can be incorporated.

The conditions such platforms face are not limited to meal delivery: Any entity that performs ad-hoc requested point-to-point transportation at scale in an urban area benefits from a robust forecasting system. Examples include ride-hailing, such as the original Uber offering, or bicycle courier services. The common characteristics are:

- **Geospatial Slicing**: Forecasts for distinct parts of a city in parallel
- **Temporal Slicing**: Forecasts on a sub-daily basis (e.g., 60-minute windows)

- **Order Sparsity**: The historical order time series exhibit an intermittent pattern
- **Double Seasonality**: Demand varies with the day of the week and the time of day

Whereas the first two points can be assumed to vary with the concrete application's requirements, it is the last two that pose challenges for forecasting a platform's demand: Intermittent demand (i.e., many observations in the historic order time series exhibit no demand at all) renders most of the commonly applied error metrics useless. Moreover, many of the established forecasting methods can only handle a single and often low seasonality (i.e., repeated regular pattern), if at all.

In this paper, we develop a rigorous methodology as to how to build and evaluate a robust forecasting system for an urban delivery platform (UDP) that offers ad-hoc point-to-point transportation of any kind. We implement such a system with a broad set of commonly used forecasting methods. We not only apply established (i.e., "classical") time series methods but also machine learning (ML) models that have gained traction in recent years due to advancements in computing power and availability of larger amounts of data. In that regard, the classical methods serve as benchmarks for the ML methods. Our system is trained on and evaluated with a dataset obtained from an undisclosed industry partner that, during the timeframe of our study, was active in several European countries and, in particular, in France. Its primary business strategy is the delivery of meals from upper-class restaurants to customers in their home or work places via bicycles. In this empirical study, we identify the best-performing methods. Thus, we answer the following research questions:

Q1: Which forecasting methods work best under what circumstances?
Q2: How do classical forecasting methods compare with ML models?
Q3: How does the forecast accuracy change with more historic data available?
Q4: Can real-time information on demand be exploited?
Q5: Can external data (e.g., weather data) improve the forecast accuracy?

To the best of our knowledge, no such study has yet been published in a scholarly journal.

The subsequent Section 2 reviews the literature on the forecasting methods included in the system. Section 3 introduces our forecasting system, and Section 4 discusses the results obtained in the empirical study. Lastly, Section 5 summarizes our findings and concludes with an outlook on further research opportunities.

2. Literature Review

3. Model Formulation

4. Empirical Study: A Meal Delivery Platform in Europe

5. Conclusion

References

[1] Alcaraz, J.J., Caballero-Arnaldos, L., Vales-Alonso, J., 2019. Rich vehicle routing problem with last-mile outsourcing decisions. Transportation Research Part E: Logistics and Transportation Review 129, 263–286.

[2] Bell, F., Smyl, S., 2018. Forecasting at uber: An introduction. `https://eng.uber.com/forecasting-introduction/`. Accessed: 2020-10-01.

[3] De Gooijer, J., Hyndman, R., 2006. 25 years of time series forecasting. International Journal of Forecasting 22, 443–473.

[4] Ehmke, J.F., Campbell, A.M., Thomas, B.W., 2018. Optimizing for total costs in vehicle routing in urban areas. Transportation Research Part E: Logistics and Transportation Review 116, 242–265.

[5] Hirschberg, C., Rajko, A., Schumacher, T., Wrulich, M., 2016. Mckinsey: The changing market for food delivery. `https://www.mckinsey.com/industries/high-tech/our-insights/the-changing-market-for-food-delivery`. Accessed: 2020-10-01.

[6] Hou, L., Li, D., Zhang, D., 2018. Ride-matching and routing optimisation: Models and a large neighbourhood search heuristic. Transportation Research Part E: Logistics and Transportation Review 118, 143–162.

[7] Laptev, N., Smyl, S., Shanmugam, S., 2017. Engineering extreme event forecasting at uber with recurrent neural networks. `https://eng.uber.com/neural-networks/`. Accessed: 2020-10-01.

[8] Masmoudi, M.A., Hosny, M., Demir, E., Genikomsakis, K.N., Cheikhrouhou, N., 2018. The dial-a-ride problem with electric vehicles and battery swapping stations. Transportation research part E: logistics and transportation review 118, 392–420.

[9] Wang, Z., 2018. Delivering meals for multiple suppliers: Exclusive or sharing logistics service. Transportation Research Part E: Logistics and Transportation Review 118, 496–512.