Modelling the performance of delivery vehicles across urban micro-regions to accelerate the transition to cargo-bike logistics

Max Schrader[0][*]
mcschrader@crimson.ua.edu

Navish Kumar[0][†]
navish@iitkgp.ac.in

Nicolas Collignon[‡]
Kale Collective

Esben Sørig[‡]
Kale Collective

Soonmyeong Yoon[‡]
Kale Collective

Akash Srivastava[§]
akashsri@mit.edu

Kai Xu[¶]
me@xuk.ai

Maria Astefanoaei[∥]
msia@itu.dk

Abstract

Light goods vehicles (LGV) used extensively in the last mile of delivery are one of the leading polluters in cities. Cargo-bike logistics has been put forward as a high impact candidate for replacing LGVs, with experts estimating over half of urban van deliveries being replaceable by cargo bikes, due to their faster speeds, shorter parking times and more efficient routes across cities.

By modelling the relative delivery performance of different vehicle types across urban micro-regions, machine learning can help operators evaluate the business and environmental impact of adding cargo-bikes to their fleets. In this paper, we introduce two datasets, and present initial progress in modelling urban delivery service time (e.g. cruising for parking, unloading, walking). Using Uber's H3 index to divide the cities into hexagonal cells, and aggregating OpenStreetMap tags for each cell, we show that urban context is a critical predictor of delivery performance.

1 Introduction

Today, transport accounts for around 30% of a city's carbon emissions, and is expected to grow at a faster rate than any other sector in the coming decade [1]. One of the drivers behind this acceleration is the increasing demand for faster logistics due to the boom of the e-commerce industry. By 2025, the number of packages delivered around the world is expected to climb to 200 billion, up from fewer than 90 billion in 2018 [2]. Given the current trends, it is estimated that both the amount of CO2 emissions and traffic congestion caused by urban last-mile deliveries will increase by 30% by 2030 (6 million tonnes CO2 compared to 2019) [3].

The single most challenging aspect for logistics operators in meeting this exceptional rise in demand comes down to the last mile of the delivery (with some reporting it accounts for more than 50% of their total delivery cost). Recently, cargo bike logistics has been put forward as a competitive, ultra-low emission, and environmentally sustainable candidate for replacing vans-only logistics fleets in cities. Scientific reports and pilot studies have been confirming the potential of cargo bike logistics

---

[0]These authors contributed equally to this work

[*]Department of Mechanical Engineering, University of Alabama, Tuscaloosa, Alabama, USA

[†]Department of Humanities and Social Sciences, IIT Kharagpur, Kharagpur, West Bengal 721302, India

[‡]Kale Collective, London, UK

[§]MIT-IBM Watson AI Lab, Cambridge, MA, USA

[¶]University of Edinburgh, Edinburgh, UK

[∥]Data-intensive Systems and Applications, IT University of Copenhagen, Copenhagen, Denmark

and their superior efficiency in dense urban areas [4, 5, 6]. As an example, a recent study found that 67% of daily van operations of a large logistics operator in Paris could be substituted by cargo-bikes at no extra cost [7].

The importance of data-driven decision-making is paramount in an industry known for high pressure and tight margins. However, data collection processes vary widely in maturity. Incumbent companies struggle from data quality issues across the logistics chain, which hampers their data science efforts, and little data around cargo-bike logistics exists.

Current tools used by operators fall short of capturing the complexity of driver behaviour, even less the varying vehicle efficiency in the real world. Experienced delivery drivers have tacit knowledge that help them efficiently navigate highly complex urban environments to serve customers. The theoretical routes generated by optimisation engines still diverge considerably from the complexities of the real wold; they fail to take into account factors such as time spent by delivery drivers searching for parking, or the time spent walking between the vehicle and multiple delivery locations [8], that both have significant impact on the delivery time and depend heavily on the type of vehicle used. As such, predicting the duration of delivery runs as well as estimating delivery times remains a considerable challenge for operators [9].

2 Modelling delivery performance of vehicles across urban micro-regions

Modelling the performance of delivery vehicles (here, vans and cargo-bikes) will help towards accurately representing their relative efficiency across different urban contexts. Operators are faced with complex trade-offs that need to consider to the quantity and size of deliveries, their density patterns across different urban areas, and the capacity and efficiency of vehicles in their fleets. We present initial results on predicting service time, which includes time searching for parking, unloading parcels, and walking to destinations.

The importance of service time for delivery performance. A 2018 study in London highlighted that walking can account for 62% of the total van round time [10] due to limited suitable kerbside parking space in urban centres. Similarly, studies found that delivery drivers spend on average 5.8 minutes and 24 minutes searching for each parking spot in Seattle and New York City, respectively [11, 12]. Despite the significant amount of time currently spent looking for parking, optimisation tools largely ignore the search time for parking in route planning [13]. Choosing when and where to park remains a decision made by the driver [8], which can have significant impact on their performance [14]. The nimbleness of cargo-bikes presents a competitive advantage by allowing riders to find parking much quicker, and closer to the final destination.

2.1 Representing urban context at the micro-region level

Dividing cities in micro-regions. To explore the role of urban context in service time prediction we divide the greater delivery area into micro-regions. For this, we turn to Uber's Hexagonal Hierarchical Spatial Index (H3) [15]. The H3 spatial index divides the Earth's sphere into hexagonal cells at different levels of granularity, providing us with a method to tessellate the delivery areas in a way that is agnostic of the city or country. We use a H3 resolution of 9 (average edge length of 174.4 meters), which was found to capture several city blocks effectively [16].

Using OpenStreetMap tags as descriptions of urban micro-regions. To describe the characteristics of urban micro-regions, we rely on OpenStreetMap (OSM) tags [17]. Each tag is comprised of sub-tags in the form of *key=value* pairs (e.g. highway=footway). We aggregate these OSM tags on the H3 cell level by counting the tagged-objects to include information that may be predictive about service time such as land-use, road types, density of apartment buildings, traffic lights, parking zones, businesses, and more. Theses tags come to represent the urban characteristics of each H3 cell (or the urban context of a delivery). Embeddings using this type of representation have been shown to exhibit strong semantically interpretable characteristics of urban space [16]. We use OSMNx [18] to extract the OSM data and consider additional tags to the ones in [16] that may be directly relevant to urban logistics.

2.2 Van and cargo-bike delivery datasets

Analysis of the Amazon van-delivery dataset. Amazon has recently open-sourced a large dataset of van-deliveries for their Last Mile Routing Research Challenge [9]. It is one of the largest and most up-to-date van delivery datasets openly available, spanning five U.S. cities. The 9,184 routes in the dataset provide good coverage of city topography, including high-density multi-family deliveries, lower density suburban deliveries, and deliveries to businesses and business centres. In the dataset, each stop has an associated planned service time, which encapsulates the time required to find parking and make the delivery. The importance of service time in a driver's day is reflected in the data, as they spend 56% of their day either looking for parking or walking to make deliveries, averaging 1.8 minutes of service time per stop [9]. The distribution of service times is heavy-tailed with stops in dense urban areas having a high likelihood of lasting longer than 10 minutes (see Figure 1).

Pedal Me and Urbike cargo-bike delivery datasets. Pedal Me (London, UK) and Urbike (Brussels, Belgium) are two cargo-bike operators, with fleet sizes of around 100 and 30 bikes respectively. We are currently in the process of designing a dataset based on GPS traces, as well as pick-up and drop-off locations that will enable a comparison of delivery efficiency with vans across urban areas.

2.3 Preliminary results on the Amazon dataset

To understand which urban features impact delivery service time, we perform a regression analysis on H3 cell average service time using NGBoost [19] and cell tag counts as features. NGBoost predicts parametric conditional probability distributions. We use NGBoost to predict the log-distribution (mean and SD) for service times within the 1197 hexagons (and 35442 deliveries) conditioned on the tag count features (we consider 674 features). Following a 5-fold validation on the data, we obtain a mean NLL of 3.51 ($SD = 1.23$) and mean R^2 value of 0.55 ($SD = 0.08$). We find that the most important features are coherent with characteristics of urban space one would expect to influence service time, such as the presence of parking, the type and density of buildings, or the type of roads.
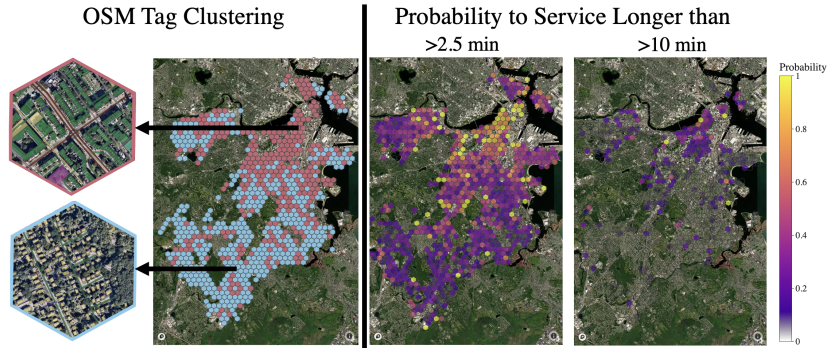

Figure 1: Clustering cells based on their OSM tags in the left-most image. The most representative cells for each of the clusters are inflated on the left, with the OSM geometries overlaid on top.

In Figure 1, we perform agglomerative clustering with complete linkage and cosine distance [20] on the significant OSM tags for the city of Boston, Massachusetts. For simplicity, two clusters are created. The clusters approximately split the residential (blue) and urban (red) areas. The residential cluster has a median service time of 95 seconds, with the urban cluster having a service time of 147 seconds ($t(13052) = -31.02, p < 0.001$, after log-transforming the data). The difference in average service times points to the pivotal role that urban context plays in delivery.

2.4 Future work

Our ongoing work is focused on comparing the relative delivery performance of cargo-bikes and vans across types of urban micro-regions, considering both service and navigation times. We are exploring the use of embeddings for generalisable representations of urban context and plan on releasing datasets annotated with urban characteristics based on OSM tags. We believe this will unlock opportunities for the machine learning community to accelerate the decarbonisation of urban freight and contribute to healthier cities.

Acknowledgments and Disclosure of Funding

References

[1] Felix Creutzig, Patrick Jochem, Oreane Y Edelenbosch, Linus Mattauch, Detlef P van Vuuren, David McCollum, and Jan Minx. Transport: A roadblock to climate change mitigation? *Science*, 350(6263):911–912, 2015.

[2] Statista. ecommerce report. Technical report, Statista Digital Market Outlook, 2021.

[3] T Deloison, E Hannon, A Huber, B Heid, Ch Klink, R Sahay, and C Wolff. The future of the last-mile ecosystem. In *World Econ. Forum*, volume 1, pages 1–28, 2020.

[4] Johannes Gruber, Alexander Kihm, and Barbara Lenz. A new vehicle for urban freight? an ex-ante evaluation of electric cargo bikes in courier services. *Research in Transportation Business & Management*, 11:53–62, 2014.

[5] Manali Sheth, Polina Butrina, Anne Goodchild, and Edward McCormack. Measuring delivery route cost trade-offs between electric-assist cargo bicycles and delivery trucks in dense urban areas. *European transport research review*, 11(1):1–12, 2019.

[6] Walther Ploos van Amstel, Susanne Balm, Jos Warmerdam, Martin Boerema, Martijn Altenburg, Frank Rieck, and Toin Peters. City logistics: Light and electric. *Amsterdam University of Applied Sciences*, 2018.

[7] Antoine Robichet, Patrick Nierat, and Francois Combes. First and last miles by cargo bikes: Ecological commitment or economically feasible? the case of a parcel service company in paris. *Transportation Research Record*, page 03611981221086632, 2022.

[8] Nils Boysen, Stefan Fedtke, and Stefan Schwerdfeger. Last-mile delivery concepts: a survey from an operational research perspective. *Or Spectrum*, 43(1):1–58, 2021.

[9] Daniel Merchán, Jatin Arora, Julian Pachon, Karthik Konduri, Matthias Winkenbach, Steven Parks, and Joseph Noszek. 2021 amazon last mile routing research challenge: Data set. *Transportation Science*, 2022.

[10] Julian Allen, Maja Piecyk, Marzena Piotrowska, Fraser McLeod, Thomas Cherrett, Karen Ghali, Thuba Nguyen, Tolga Bektas, Oliver Bates, Adrian Friday, et al. Understanding the impact of e-commerce on last-mile light goods vehicle activity in urban areas: The case of london. *Transportation Research Part D: Transport and Environment*, 61:325–338, 2018.

[11] Giacomo Dalla Chiara and Anne Goodchild. Do commercial vehicles cruise for parking? empirical evidence from seattle. *Transport Policy*, 97:26–36, 2020.

[12] José Holguín-Veras, Johanna Amaya, Trilce Encarnacion, Sofia Kyle, and Jeffrey Wojtowicz. Impacts of freight parking policies in urban areas: The case of new york city. 2016.

[13] Sara Reed, Ann Melissa Campbell, and Barrett W Thomas. Does parking matter? the impact of search time for parking on last-mile delivery optimization. *arXiv preprint arXiv:2107.06788*, 2021.

[14] Oliver Bates, Adrian Friday, Julian Allen, Tom Cherrett, Fraser McLeod, Tolga Bektas, ThuBa Nguyen, Maja Piecyk, Marzena Piotrowska, Sarah Wise, et al. Transforming last-mile logistics: Opportunities for more sustainable deliveries. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, pages 1–14, 2018.

[15] H3: Uber's hexagonal hierarchical spatial index, Jun 2018.

[16] Szymon Woźniak and Piotr Szymański. Hex2vec – Context-Aware Embedding H3 Hexagons with OpenStreetMap Tags. In *Proceedings of the 4th ACM SIGSPATIAL International Workshop on AI for Geographic Knowledge Discovery*, pages 61–71, November 2021.

[17] OpenStreetMap contributors. Planet dump retrieved from https://planet.osm.org . `https://www.openstreetmap.org`, 2017.

[18] Geoff Boeing. OSMnx: New methods for acquiring, constructing, analyzing, and visualizing complex street networks. *Computers, Environment and Urban Systems*, 65:126–139, 2017.

[19] Tony Duan, Avati Anand, Daisy Yi Ding, Khanh K Thai, Sanjay Basu, Andrew Ng, and Alejandro Schuler. Ngboost: Natural gradient boosting for probabilistic prediction. In *International Conference on Machine Learning*, pages 2690–2700. PMLR, 2020.

[20] Daniel Müllner. Modern hierarchical, agglomerative clustering algorithms. *arXiv preprint arXiv:1109.2378*, 2011.