

# Learning Mobility Flows from Urban Features with Spatial Interaction Models and Neural Networks\*

Gevorg Yeghikyan  
*Scuola Normale Superiore*  
Pisa, Italy  
gevorg.yeghikyan@sns.it

Felix L. Opolka  
*University of Cambridge*  
Cambridge, UK  
flo23@cam.ac.uk

Mirco Nanni  
*ISTI-CNR*  
Pisa, Italy  
mirco.nanni@isti.cnr.it

Bruno Lepri  
*FBK*  
Trento, Italy  
lepri@fbk.eu

Pietro Liò  
*University of Cambridge*  
Cambridge, UK  
pl219@cam.ac.uk

**Abstract**—A fundamental problem of interest to policy makers, urban planners, and other stakeholders involved in urban development projects is assessing the impact of planning and construction activities on mobility flows. This is a challenging task due to the different spatial, temporal, social, and economic factors influencing urban mobility flows. These flows, along with the influencing factors, can be modelled as attributed graphs with both node and edge features characterising locations in a city and the various types of relationships between them. In this paper, we address the problem of assessing origin-destination (OD) car flows between a location of interest and every other location in a city, given their features and the structural characteristics of the graph. We propose three neural network architectures, including graph neural networks (GNN), and conduct a systematic comparison between the proposed methods and state-of-the-art spatial interaction models, their modifications, and machine learning approaches. The objective of the paper is to address the practical problem of estimating potential flow between an urban development project location and other locations in the city, where the features of the project location are known in advance. We evaluate the performance of the models on a regression task using a custom data set of attributed car OD flows in London. We also visualise the model performance by showing the spatial distribution of flow residuals across London.

**Index Terms**—urban mobility flows, spatial interaction models, graph neural networks, urban computing

## I. INTRODUCTION

Planning and managing city and transportation infrastructures requires understanding the relationship between urban mobility flows and spatial, structural, and socio-economic features associated with them. There exists extensive literature addressing this problem ranging from the classical gravity model and its modifications [1], [2] to the more recent spatial econometric interaction models [3] and the non-parametric radiation models [4] that attempt to characterise cross-sectional origin-destination (OD) flow matrices. Furthermore, various neural network-based models have been proposed for predicting temporal OD flow matrices [5], [6]. However, modelling OD flow matrices in their entirety, the mentioned works do not address the problem of assessing flows between a specific location and every other location in the city, given all other flows, other location characteristics, as well as information on the dyadic relations between those locations.

More specifically, the motivation for this task comes from a scenario in which it is necessary to assess the impact of an urban development project on the OD flows in and out of the project’s location. Examples of these motivating scenarios include retail location choice and consumer spatial behaviour prediction, which have been approached with the Huff model and its modifications [7]. These models, however, suffer from a series of drawbacks related mostly to overly restrictive assumptions. In this paper, we take a different approach and focus on the problem of evaluating OD flows in and out of a location of interest. By modelling urban flows as attributed graphs in which the nodes represent locations in a city (i.e. each node is described by a vector of features such as population density, Airbnb prices, available parking areas, etc.), and the edges represent the car flows between them (each one described by a vector of features such as road distance, average time required to travel, average speed, etc.), this project aims to offer an instrument for assessing flows between a specific location and all other locations in the city.

Since a rigorous experimental setting would have required difficult-to-obtain longitudinal data of OD flows *before* and *after* the completion of an urban development project, we set up a *quasi-experimental* setting. We randomly select locations in a city and the flows associated with them as a test set, and attempt to find a function that takes the urban features describing city locations and the remaining flows as input, and predicts the flows in the test set as output.

In sum, our paper makes the following contributions:

- We propose three neural network architectures for predicting car flows between a location of interest and every other location in a city. Two of the models use graph convolutional layers that pool information from geographical or topological neighbourhoods around relevant nodes to incorporate more information (Section V).
- We evaluate and compare our models on a custom dataset of aggregate OD car flows in London, containing node and edge features (Section VI).
- We show that the proposed neural network models outperform well-known spatial interaction and machine learning models. A comparison among neural network models reveals that graph convolutions do not substantially improve prediction performance on the formulated task (Sections IV, VI).

\*To appear in the Proceedings of 2020 IEEE International Conference on Smart Computing (SMARTCOMP 2020)

- We describe our custom dataset and make it publicly available along with the code for this study (Section III).

## II. RELATED WORK

The problem of estimating human flows between locations in a geographical space has been first addressed by [1] through a family of spatial interaction models and subsequently extended by [2]. Spatial interaction models, extensively used to estimate human mobility flows and trip demand between locations as a function of the location features, have become an acknowledged method for modelling geographical mobility in transportation planning [8], [9], commuting [10], and spatial economics [11]. The spatial interaction models are usually calibrated via an Ordinary Least Squares (OLS) regression, which assumes normally distributed data. However, OD flows are usually not distributed normally, are count data, and contain a large number of zero flows. This makes the setting incompatible with OLS estimation and requires either a Poisson model or, in the presence of over-dispersion, a Negative Binomial Regression (NB) model [12].

Another major concern in this modelling scenario are the complex interactions often caused by spatial dependencies and non-stationarity. The former arises from spill-over effects from a location to its neighbourhoods, while the latter is caused by the influence of independent variables varying across space. These issues have been addressed in literature by spatial autocorrelation and geographically weighted modelling techniques [13], [3], [14], [12].

Another approach within the spatial interaction modelling paradigm is the Huff model and its extensions [7]. Originally developed mainly for retail location choice and turnover prediction, they represent a probabilistic formulation of the gravity model. The Huff model considers OD flows as proportional to the relative attractiveness and accessibility of the destination compared to other competing destinations. The probability  $P_{ij}$  of a consumer at location  $i$  of choosing to shop at a retail location  $j$  is framed as:

$$P_{ij} = \frac{A_j^\alpha D_{ij}^{-\beta}}{\sum_{j=1}^n A_j^\alpha D_{ij}^{-\beta}}, \quad (1)$$

where  $A_j$  is a measure of attractiveness of retail location  $j$ , such as area or a linear combination of different features,  $D_{ij}$  is the distance between locations  $i$  and  $j$ ,  $\alpha$  and  $\beta$ , estimated from empirical observations, are attractiveness and distance decay parameters, respectively.

Along with traditional gravity methods, the Huff model and its variations have found their way to numerous applications including location selection of movie theaters [15], a university campus [16], or the analysis of spatial access to health care [17].

However, these models suffer from too restrictive assumptions such as considering the ratio of the probabilities of an individual selecting two alternatives as being unaffected by the introduction of a third alternative. Although the competing destinations model [18] has overcome this, it has the

disadvantage of considering either spatial agglomeration or competition effects, ignoring the fact that they can coexist in the same location. Even though a number of extensions to the Huff model and the gravity framework in general have been proposed to overcome spatial non-stationarity and to include a larger array of features affecting the flows [19], [20], this family of models, along with the non-parametric radiation and population-weighted opportunities model, have demonstrated to fall short of high predictive capacity particularly at the city scale [21], [22], [23].

More recently, machine learning, particularly a Random Forest approach, has shown promising results in reconstructing inter-city OD flow matrices [24]. However, its performance on intra-urban flow data remains to be tested.

Moreover, as already mentioned, the discussed models address the problem of modelling the OD flow matrix as a whole and have to be adapted to our specific task of estimating flows between a specific location and all other locations, given the other flows in the city, the location features, and the features describing the dyadic relations between them, respectively.

The problem of estimating OD flows has also been addressed with neural network methods [25]. As flows are most naturally modelled by graphs, most work has focused on the use of graph neural networks for flow estimation. An early neural network model for graph structured data has been suggested in [26]. Later work has specifically focused on generalising Convolutional Neural Networks from the domain of regular grids to the domain of irregular graphs [27], [28]. One of the most commonly used graph neural network models is the Graph Convolutional Neural Network (GCN) proposed in [29].

Graph neural networks have previously been applied to urban planning tasks. In [5], they have been used to predict the flow of bikes within a bike sharing system. Unlike our model, flows are modelled as node-level features, which requires a different neural network model and does not allow to predict flows between specific pairs of nodes. Although [30] uses graph neural networks to predict flows between parts of a city, their model operates on spatio-temporal data and focuses on the temporal aspect of the data. Beyond flow prediction, in [31], a graph neural network model has been proposed for building site selection. A broader overview of machine learning methods applied to the task of urban flow prediction is given in [32]. In this work, we define neural network models that make use of stationary node and edge features and compare different neural network architectures based on fully connected networks and graph neural networks.

## III. DATA DESCRIPTION

We publicly release<sup>1</sup> a custom dataset of aggregate origin-destination (OD) flows of private cars in London augmented with feature data describing city locations and dyadic relations between them. The workflow of building the dataset is as follows:

<sup>1</sup>Dataset will be released at <https://trackandknowproject.eu/file-repository/>. Code available at [github.com/FelixOpolka/Mobility-Flows-Neural-Networks](https://github.com/FelixOpolka/Mobility-Flows-Neural-Networks).

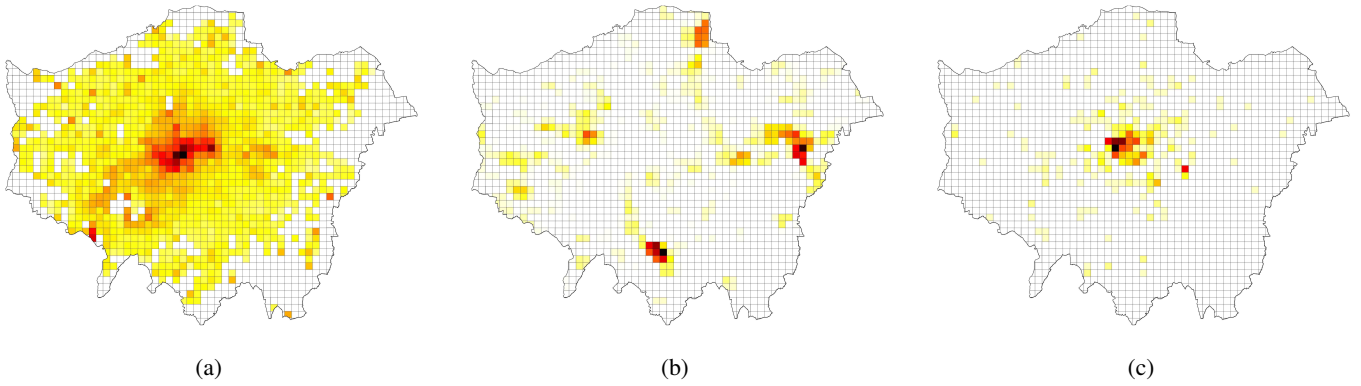


Fig. 1: Examples of node (cell) features (a) Average Airbnb listing prices (b) Proportion of grid cell area allotted to industrial activity (c) Number of museums and galleries per grid cell. Darker colours indicate higher values.

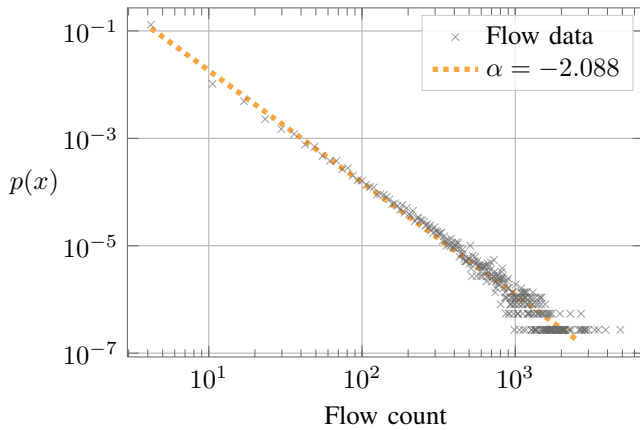


Fig. 2: Log-log plot of the probability distribution of the OD flows fitted with a power-law distribution  $p(x) \propto x^{-\alpha}$  with exponent of  $\alpha = -2.088$ .

- 1) The urban territory has been subdivided into  $n$  Cartesian grid cells of size  $500 \times 500$  m, and each such quadratic cell is considered a node in the graph.
- 2) The GPS trajectories of around 10000 cars spanning a period of one year, provided by a car insurance company for research purposes, have been superimposed on the grid, and trip origins and destinations have been extracted (Figure 3a).
- 3) The OD network has been built from the extracted origin-destination pairs by aggregating the flow counts over a year (Figure 3b). Since the aggregation spans such a long time period, the OD matrix is approximately symmetric, and thus has been converted into a symmetric matrix by averaging the matrix with its transpose.
- 4) The node features have been built by engineering 35 features from various open sources [33], [34], [35] and from the GPS data. These features include population density, average Airbnb prices, parking areas, areas covered by residential buildings, number of restaurants, bars, banks, museums, road network density, average

radius of gyration, etc. per cell. Examples of node features and their spatial distribution are visualised in Figure 1.

- 5) Similarly, the edge features encode information on 12 dyadic relations such as network distance, average time, average speed, temporal correlation between car incidence in cells, public transport connections, etc. The detailed attribute description is provided with the dataset.

#### IV. PROBLEM STATEMENT

In this section, we describe the problem we are addressing and state definitions of important terms.

We define a **weighted attributed graph**  $G = (\mathcal{V}, \mathcal{E}, \mathbf{W}, \mathbf{X}^v, \mathbf{X}^e)$  with feature information associated with both nodes and edges. More specifically,  $\mathcal{V}$  is the set of  $n$  nodes, and  $\mathcal{E} = \{e_{ij} = (i, j) : i, j \in \mathcal{V}\}$  represents the set of  $m$  edges in graph  $G$ . Furthermore,  $\mathbf{W} \in \mathbb{R}^{n \times n}$  is the weighted adjacency matrix, essentially the OD matrix, with  $\mathbf{W}_{ij} \geq 0 \forall i, j \in \mathcal{V}$  corresponding to the flow between cells  $i$  and  $j$ . Additionally, we denote the node feature matrix as  $\mathbf{X}^v \in \mathbb{R}^{n \times p}$ , where  $p$  is the number of node features. The edge feature matrix, on the other hand, is denoted as  $\mathbf{X}^e \in \mathbb{R}^{m \times k}$ , where  $k$  is the number of edge features.

The urban mobility flow network  $T$  is a weighted undirected attributed graph whose nodes are  $500 \times 500$  m city grid cells, and the edges are the aggregate flows between them. The nodes and edges are additionally augmented by feature vectors described in detail in Section III. Furthermore, each edge  $e_{ij}$  in the urban mobility flow network  $T$  is associated with a target (or ground truth) flow  $w_{ij}$ , which is the corresponding entry in the weighted adjacency matrix  $\mathbf{W}$  of  $T$ . It represents the aggregate mobility flow between cell (node)  $i$  and cell (node)  $j$  in the network.

In our prediction setting, we are given the urban mobility flow network  $T = (\mathcal{V}, \mathcal{E}, \mathbf{W}, \mathbf{X}^v, \mathbf{X}^e)$  and a node of interest  $i$  for which the target flows  $W_{i1}, \dots, W_{in}$  are unknown. Hence, we aim to learn a mapping  $f : \{\mathcal{V}, \mathcal{E}, \mathbf{W}, \mathbf{X}^v, \mathbf{X}^e\} \rightarrow \mathbb{R}^n$  from the urban mobility flow network to the missing flows, i.e.  $[W_{i1}, \dots, W_{in}] = f(i, \mathbf{W}, \mathbf{X}^v, \mathbf{X}^e) \forall i \in \mathcal{V}$ . In other words, the

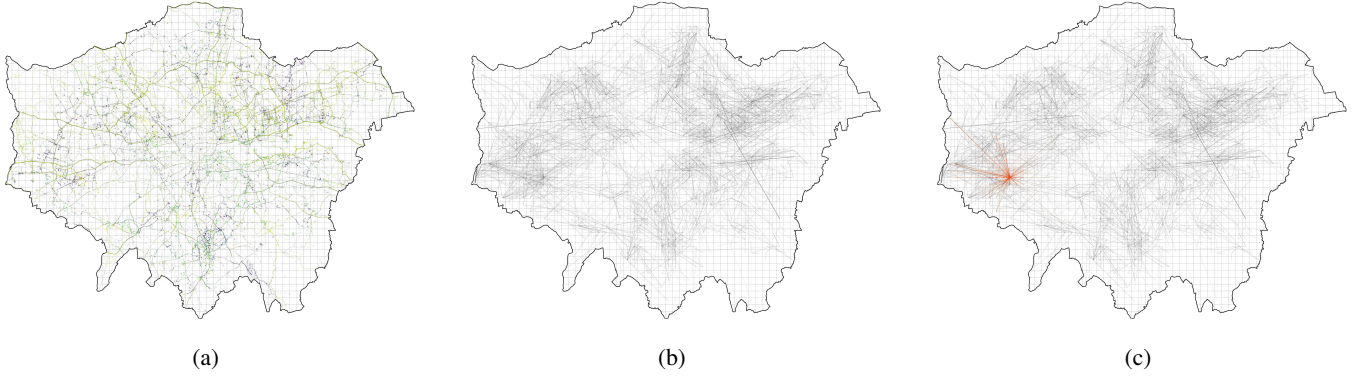


Fig. 3: (a) Car GPS trajectories over grid cells in London. (b) Origin-Destination (*OD*) flow network in London. (c) Target flows between a node of interest and every other node.

aim is to predict the missing target flows (Figure 3c), given the features of node  $i$  and the rest of the graph.

## V. METHODOLOGY

In the following, we describe three neural network models that are trained to predict the unknown flows in the urban mobility flow network  $T$ . When a model makes a prediction for the flow associated with an edge going from a node of interest to another node in the graph, it can use all node and edge features in the graph, as these features are available even for nodes of interest, i.e. sites of prospective urban development projects. Furthermore, it may use the ground truth flows for edges that are not connected to a node of interest. In a practical situation, this corresponds to the flows between *existing* locations in the city for which flow information is therefore available.

The first neural network architecture is a fully connected neural network operating on the features of the target edge and the features of its two incident nodes. More specifically, when predicting the flow for target edge  $e_{ij}$ , we concatenate the node features  $\mathbf{x}_i^v$  and  $\mathbf{x}_j^v$  for incident node features, as well as the corresponding edge features  $\mathbf{x}_{ij}^e$ . The concatenated vector

$$\bar{\mathbf{x}} = [\mathbf{x}_i^v, \mathbf{x}_{ij}^e, \mathbf{x}_j^v] \quad (2)$$

is passed into a fully connected neural network with ReLU-non-linearities, defined as  $\text{ReLU}(z_j) = \max(0, z_j)$ , where  $z_j$  is the  $j^{\text{th}}$  output of the linear transformation. Each fully connected layer is followed by batch normalisation [36] and dropout [37] to counter overfitting. We refer to this model as *FCNN*.

The second model builds upon the *FCNN* model through the additional use of graph convolutions to generate embeddings of node neighbourhoods. We use a graph convolutional neural network (GCN) [29] to generate node embeddings  $\mathbf{h}_i, \mathbf{h}_j$  for the two nodes incident to the target edge  $e_{ij}$ . GCN layers extend fully-connected layers with an additional neighbourhood aggregation step before the non-linearity. The layer applies a linear transformation to all node features  $\mathbf{h}_i^{(l-1)}$  in the graph and then, for each node, computes a weighted average of the

resulting representations at the central node and in the 1-hop neighbourhood of the central node:

$$\mathbf{z}_i^{(l)} = \sum_{j \in \mathcal{N}(i) \cup \{i\}} \frac{1}{\sqrt{(d_i + 1)(d_j + 1)}} \mathbf{h}_j^{(l-1)} \Theta, \quad (3)$$

where  $\Theta \in \mathbb{R}^{D^{(l-1)} \times D^{(l)}}$  is a learned weight matrix,  $\mathcal{N}(i)$  refers to the 1-hop neighbourhood of node  $i$ , and  $d_i$  denotes the degree of node  $i$ . This aggregation scheme is followed by a non-linearity and can be written more compactly using matrix multiplication as

$$\mathbf{H}^{(l)} = \text{ReLU}(\tilde{\mathbf{D}}^{-\frac{1}{2}} \tilde{\mathbf{W}} \tilde{\mathbf{D}}^{-\frac{1}{2}} \mathbf{H}^{(l-1)} \Theta). \quad (4)$$

where  $\tilde{\mathbf{W}} = \mathbf{W} + \mathbf{I}$  and  $\tilde{\mathbf{D}}$  is the degree matrix of  $\tilde{\mathbf{W}}$ . Equation 4 defines a graph convolutional layer and multiple such layers can be stacked to form a multi-layer graph neural network. A GNN with  $k$  layers allows us to compute embeddings encoding node feature information from within a  $k$ -hop neighbourhood.

For the second model, we apply multiple graph convolutions as defined above on the flow-weighted geographical adjacency matrix  $\mathbf{W}^{\text{geo}}$  where  $W_{ij}^{\text{geo}}$  is non-zero if and only if node  $i$  is in the geographical neighbourhood of node  $j$  and  $W_{ij}^{\text{geo}} = W_{ij}$ , i.e. the flow between  $i$  and  $j$ . The resulting node embeddings  $\mathbf{h}_i, \mathbf{h}_j \in \mathbb{R}^D$  for the two nodes incident to edge  $e_{ij}$  are added to the representation of  $\bar{\mathbf{x}}$  (see Equation 2 after the first fully connected layer):

$$\mathbf{h}_{ij}^{(1)} = \phi_1 \text{FCN}(\bar{\mathbf{x}}) + \phi_2 [\text{GNN}(\mathbf{x}_i) + \text{GNN}(\mathbf{x}_j)], \quad (5)$$

where  $\phi_1, \phi_2$  are learned weighting coefficients. We note that both mentions of  $\text{GNN}(\cdot)$  refer to the same sequence of graph convolutional layers. We then feed  $\mathbf{h}_{ij}^{(1)}$  into a number of fully connected layers, again with dropout and batch normalisation, such that the resulting model contains the same number of fully connected layers as the *FCNN* model. We call the resulting model *GNN-geo*.

Finally, we evaluate a third model, denoted by *GNN-flow*, which is equivalent to *GNN-geo* except graph convolutions are performed using the flow-based adjacency matrix  $\mathbf{W}^{\text{flow}} =$




Fig. 4: Overview of the neural network model architectures. When predicting the flow for edge  $e_{ij}$ , all three models concatenate the corresponding edge features  $\mathbf{x}_{ij}^e$ , and the node features  $\mathbf{x}_i^v, \mathbf{x}_j^v$  of the incident nodes. The resulting vector is fed into a single fully connected layer. In case of the GNN-based models *GNN-geo* and *GNN-flow*, the network also perform graph convolutions on the neighbourhoods of  $v_i$  and  $v_j$  and computes a weighted sum of both neighbourhood embeddings and the edge embedding. A further set of fully connected layers maps the sum to the predicted flow  $\hat{y}_{ij}$ . The *FCNN* model skips the addition step and does not perform graph convolutions.

$\mathbf{W}$ , where  $W_{ij}^{flow}$  is the flow between  $i$  and  $j$ . Hence, the adjacency matrix used by *GNN-flow* will contain additional edges to those used by *GNN-geo*. A visualisation of the model architectures is given in Figure 4.

The graph based models *GNN-geo* and *GNN-flow* require flow information for the adjacency matrices. While this is readily available for edges between two regular nodes, we have to approximate flow between a regular node  $i$  and a node of interest  $j$ . This is done by taking the average of the flows from node  $i$  to each node in the neighbourhood of  $j$ , i.e.

$$\tilde{W}_{ij} = \frac{1}{|\mathcal{N}(j)|} \sum_{k \in \mathcal{N}(j)} W_{ik}. \quad (6)$$

We note that even though the *FCNN* does not use graph convolutions and hence does not qualify as a common graph neural network, it does use graph structure information by concatenating specifically the features  $\mathbf{x}_i^v, \mathbf{x}_j^v$  of the nodes incident to the target edge  $e_{ij}$ .

All models output the flow corresponding to the target edge  $e_{ij}$  and are trained to minimise the mean squared error between the predicted and the actual flow. More details on the experimental setup are provided in Section VI-C.

## VI. EXPERIMENTS

We evaluate the described model on the London dataset described in Section III. In the following, we describe the goodness-of-fit metrics we use to measure model performance, the baseline methods we compare our models to, and the experimental setup.

### A. Goodness-of-fit measures

*Mean absolute error (MAE)*. Let  $\hat{y}_{ij}$  be the predicted flow between  $i$  and  $j$ ,  $y_{ij}$  be the ground truth flow, then

$$\text{MAE} = \frac{1}{|\mathcal{E}|} \sum_i \sum_j |y_{ij} - \hat{y}_{ij}|. \quad (7)$$

*Binned MAE*. Due to the highly skewed distribution of the flow data, the vast majority of flows have a small flow count, with only a handful of flows with a very large flow value (see Figure 2). Because of this, the total MAE will be biased downwards. To account for this, we additionally measure the MAE of all models within 4 bins with the following boundaries:  $0 \leq 10.0 \leq 100.0 \leq 1000.0 \leq 10000.0$ , corresponding to  $\text{MAE}_0, \text{MAE}_1, \text{MAE}_2, \text{MAE}_3$ , respectively. Finally, we define the MAE bin mean as

$$\text{Bin mean MAE} = \frac{\text{MAE}_0 + \text{MAE}_1 + \text{MAE}_2 + \text{MAE}_3}{4}, \quad (8)$$

where  $\text{MAE}_i$  refers to MAE of the  $i^{\text{th}}$  bin.

*Mean absolute percentage error (MAPE)*. To display the model accuracy with respect to the ground-truth flow values, we further use the mean absolute percentage error, defined as

$$\text{MAPE} = 100 \times \frac{1}{|\mathcal{E}|} \sum_i \sum_j \left| \frac{y_{ij} - \hat{y}_{ij}}{y_{ij}} \right|, \quad (9)$$

*Sorensen similarity index*. We use a modified version of the Sorensen similarity index (SSI), which has been extensively used in spatial interaction modelling [23], [38], and is defined as

$$\text{SSI} = \frac{1}{|\mathcal{E}|} \sum_i \sum_j \frac{2 \min(y_{ij}, \hat{y}_{ij})}{y_{ij} + \hat{y}_{ij}}, \quad (10)$$

and takes on values between 0 and 1, with values closer to 1 denoting a better fit.

*Common part of commuters.* Further, we use a similar metric, the common part of commuters, used specifically for mobility OD flow networks [38]:

$$CPC = \frac{2 \sum_{i,j=1}^n \min(y_{ij}, \hat{y}_{ij})}{\sum_{i,j=1}^n y_{ij} + \sum_{i,j=1}^n \hat{y}_{ij}}. \quad (11)$$

This measure takes on the value 0, when the flows in the two networks completely differ, and 1, when they are in perfect agreement.

*Common part of links.* Finally, to measure the degree to which the topological structure of the original network has been reconstructed, we use the common part of links (CPL) [39] defined as

$$CPL = \frac{2 \sum_{i,j=1}^n \mathbb{1}_{y_{ij}>0} \cdot \mathbb{1}_{\hat{y}_{ij}>0}}{\sum_{i,j=1}^n \mathbb{1}_{y_{ij}>0} + \sum_{i,j=1}^n \mathbb{1}_{\hat{y}_{ij}>0}}, \quad (12)$$

where  $\mathbb{1}_A$  is the indicator function of condition  $A$ . The common part of links shows the proportion of links between the observed and predicted networks such that  $y_{ij} > 0$  and  $\hat{y}_{ij} > 0$ . It takes on the value zero if the two networks have no common links and one if the networks are topologically equivalent.

### B. Baseline models

In this study, we compare the proposed model to the following baselines, using the same experimental setup for all models:

- **Doubly constrained gravity model (DC-GM):** The classical gravity model with a power law decay has several formulations with respect to preserving the total in- nor out-flows during model calibration: unconstrained, origin-constrained, destination-constrained, and doubly constrained. Here we take the latter.
- **Huff model:** A probabilistic formulation of the gravity model described in Section II.
- **Poisson regression:** An instance of the Generalized Linear Modelling framework, in which the dependent variable, being count data, is assumed to be drawn from a Poisson distribution.
- **Negative Binomial regression (NB):** A generalization of the Poisson regression in which the restrictive assumption that the mean and the variance of the dependent variable are equal is loosened.
- **Spatial Autoregressive Model (SAM):** An extension to the Generalized Linear Modelling framework by accounting for spatial dependence among the flows by using spatial lags represented by spatial weight matrices built from observed data [3].
- **Generalised hypergeometric ensemble multilayer network regression (gHypE):** This recent random graph approach [40] provides a statistical ensemble of all possible flow networks under the constraints of preserving in- and out-flows from each node, as well as respecting pairwise flow propensities of nodes. The multilayer network regression considers these propensities as latent variables,

inferred from the edge features describing the dyadic relations between city locations. As opposed to conventional regression methods, this method intrinsically respects the network constraints.

- **Random Forest regression (RF):** We follow the approach proposed in [24] aimed at predicting inter-city mobility flows with a set of attributes describing each city. We adapt the same approach to our problem of intra-city flow prediction. Following the described method, we use a Random Forest approach with eXtreme Gradient Boosting (*XGBoost*) [41] through 5-fold cross-validation, model and feature selection, and hyperparameter tuning.

### C. Experimental setup

For training and evaluating the three proposed models, we divide the dataset into a training, validation, and test set of edges. The subsets contain 70%, 10%, and 20% of the edges respectively. To construct the test set, we randomly select nodes in the graph and add their incident edges to the test set. We ensure that an equal number of edges fall in each of the four bins split by flow magnitude. Hence, once a bin is full, no more edges are added to the test set that would fall into this bin. We use the same procedure to construct the validation set. Nodes in the validation and test set are considered nodes of interest, while nodes in the training set are considered regular nodes.

We train all models on the same training set. To address the imbalance between flows of different magnitude, we resample the data such that each bin contains the same number of samples. We perform hyperparameter search to determine the optimal dimension of the intermediate representations, i.e. the outputs of the GCN and fully connected layers, the dropout rate, and the number of fully connected and GCN layers. We select models based on the bin mean MAE (see Equation 8) achieved on the validation set. The selected models have a total of four fully connected layers. The GNN-based models use a single GCN layer. We use a dimensionality of 32 for intermediate representations and the dropout rate is set to 0.5.

We train for a total of 110 epochs using the Adam optimiser [42] with a batch size of 256 and a learning rate of 0.01. We reduce the learning rate by a factor of ten after 50 epochs and every 15 epochs after that. We stop training early once the performance of the model does no longer improve in terms of bin mean MAE on the validation set.

We have also experimented with using different types of graph neural network layers including GAT layers [43], GIN layers [44], and Jumping Knowledge layers [45]. We did not find these layers to improve performance on the validation data set and hence preferred the conceptually simpler GCN layers.

## VII. RESULTS

We compare our models to the baseline ones in terms of MAE in Table I. We find that all three neural network models outperform all the spatial interaction models (*DC-GM*, *Huff*, *Poisson*, *NB*, *SAM*) as well as *gHypE* and *XGBoost* in terms of total MAE by a large margin. Crucially, the MAEs per bin



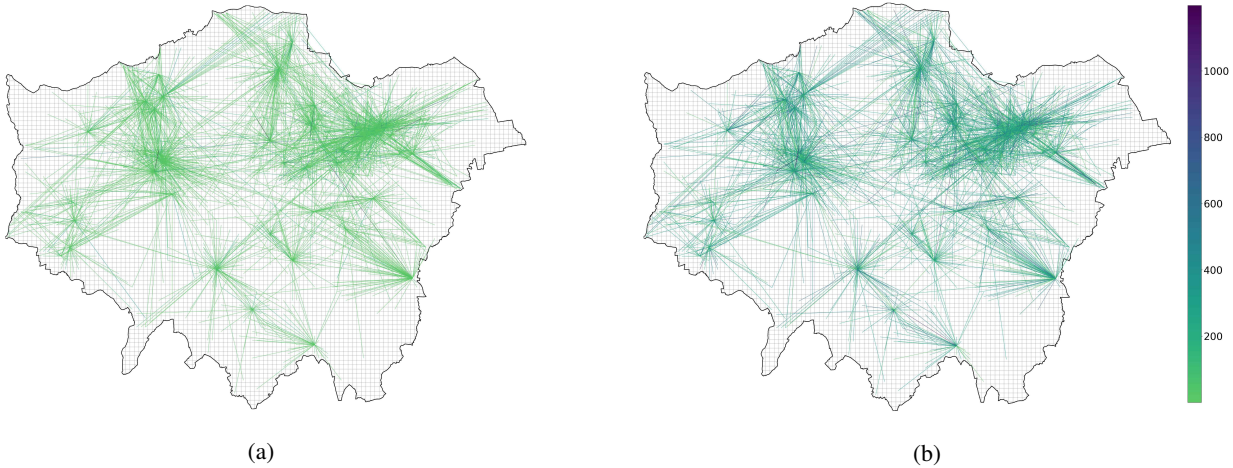


Fig. 5: MAE residuals of flows associated with test nodes (a) *GNN-geo*. (b) *XGBoost*.

MAE	Total	[0; 10)	[10; 10 <sup>2</sup> )	[10 <sup>2</sup> ; 10 <sup>3</sup> )	[10 <sup>3</sup> ; 10 <sup>4</sup> )	bin mean
<b>DC-GM</b>	167.58	64.88	170.45	881.98	2176.35	823.42
<b>Huff</b>	122.89	48.21	99.86	511.41	1476.72	534.05
<b>Poisson</b>	106.74	40.69	88.56	475.23	1261.41	466.47
<b>NB</b>	92.62	33.02	76.96	431.44	1087.12	407.14
<b>SAM</b>	75.09	19.31	61.53	395.01	989.30	366.29
<b>gHypE</b>	58.11	9.02	53.10	346.96	832.26	310.34
<b>XGBoost</b>	31.59 ± 5.88	2.61 ± 0.89	45.12 ± 11.06	228.96 ± 39.96	549.83 ± 84.79	206.63 ± 34.18
<b>FCNN</b>	12.55 ± 0.91	0.33 ± 0.08	28.97 ± 4.93	161.12 ± 22.36	408.88 ± 36.59	149.82 ± 13.65
<b>GNN-geo</b>	13.34 ± 2.51	0.52 ± 0.40	31.63 ± 9.68	161.32 ± 9.09	422.04 ± 25.70	153.88 ± 9.74
<b>GNN-flow</b>	15.35 ± 4.23	0.63 ± 0.62	38.66 ± 16.65	170.06 ± 17.41	458.05 ± 64.56	166.85 ± 16.39

TABLE I: Comparison of model performance in terms of mean absolute error grouped by flow magnitude.

	SSI	MAPE [10 <sup>3</sup> ; 10 <sup>4</sup> )	CPL	CPC
<b>DC-GM</b>	0.39	162.59	0.38	0.49
<b>Huff</b>	0.48	106.91	0.56	0.54
<b>Poisson</b>	0.46	102.10	0.57	0.54
<b>NB</b>	0.54	91.03	0.62	0.56
<b>SAM</b>	0.59	66.65	0.68	0.58
<b>gHypE</b>	0.62	52.99	0.79	0.60
<b>XGBoost</b>	0.67 ± 0.02	40.90 ± 5.85	0.86 ± 0.02	0.61 ± 0.01
<b>FCNN</b>	0.71 ± 0.00	27.16 ± 2.23	1.0 ± 0.00	0.69 ± 0.01
<b>GNN-geo</b>	0.70 ± 0.01	27.06 ± 1.65	1.0 ± 0.00	0.68 ± 0.04
<b>GNN-flow</b>	0.71 ± 0.02	30.67 ± 4.18	1.0 ± 0.01	0.65 ± 0.05

TABLE II: Comparison of model performance in terms of MAPE, SSI, CPL, and CPC.

reveal that the neural network models achieve high accuracy across bins relative to the magnitude of flows, hence the neural network does not only perform well on small flows, which are highly overrepresented in the dataset.

We also observe that there is no clear difference in the performance between the three neural network based models. Surprisingly, the graph neural networks (*GNN-geo*, *GNN-flow*) do not outperform the fully connected neural network *FCNN*. This indicates that node neighbourhood information does not result in stronger predictive performance for this dataset and prediction task. We stress, however, that while *FCNN* is not a

graph neural network, it does use graph structural information by concatenating edge features with features of incident nodes. Furthermore, previous work on mobility flow prediction has omitted an explicit comparison of GNNs to fully connected neural networks, hence it remains unclear whether GNNs offer a predictive advantage in the urban mobility setting.

Finally, we compare the neural network models to the baselines in terms of SSI, MAPE of the largest bin, CPL, and CPC. These results also confirm that the neural network models find a better fit to the data compared to the state-of-the-art.

To further illustrate the effectiveness of the GNN models, we represent the MAE residuals on the London diagrammatic maps in Figure 5. These representations show the difference between predicted and ground-truth flows between the locations in the test set. We compare the state-of-the-art *XGBoost* model with our *GNN-flow* model and observe that the latter results in spatially smoother residuals.

## VIII. CONCLUSION

In this paper, we formulated and addressed the problem of learning urban mobility flows between a location of interest and every other location in the city, given the array of socio-economic and structural features describing each location and the pairwise dyadic relations between them. We proposed

three novel neural network architectures, using fully connected and graph convolutional layers, and compared them to a set of strong baseline models. We find that the neural network models achieve state-of-the-art performance and outperform the baselines by a large margin.

In fulfilment of the stated objective, our work has direct utility to urban planners and policy makers in offering a technique for assessing mobility flows between an urban development project location and other locations in the city.

#### ACKNOWLEDGMENT

F.L.O acknowledges funding from the Huawei Hisilicon Studentship at the Department of Computer Science and Technology of the University of Cambridge. This work is partially funded by the EU H2020 programme under Grant Agreement No. 780754, “Track & Know”.

#### REFERENCES

- [1] A. G. Wilson, “A family of spatial interaction models, and associated developments,” *Environment and Planning A*, vol. 3, no. 1, pp. 1–32, 1971.
- [2] A. S. Fotheringham and M. E. O’Kelly, *Spatial interaction models: formulations and applications*, vol. 1. Kluwer Academic Publishers Dordrecht, 1989.
- [3] J. P. LeSage and R. K. Pace, “Spatial econometric modeling of origin-destination flows,” *Journal of Regional Science*, vol. 48, no. 5, pp. 941–967, 2008.
- [4] F. Simini, M. C. González, A. Maritan, and A.-L. Barabási, “A universal model for mobility and migration patterns,” *Nature*, vol. 484, no. 7392, p. 96, 2012.
- [5] D. Chai, L. Wang, and Q. Yang, “Bike flow prediction with multi-graph convolutional networks,” in *Proceedings of the 26th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*, pp. 397–400, ACM, 2018.
- [6] F. Toqué, E. Côme, M. K. El Mahrsi, and L. Oukhellou, “Forecasting dynamic public transport origin-destination matrices with long-short term memory recurrent neural networks,” in *2016 IEEE 19th International Conference on Intelligent Transportation Systems (ITSC)*, pp. 1071–1076, IEEE, 2016.
- [7] D. L. Huff, “A probabilistic analysis of shopping center trade areas,” *Land economics*, vol. 39, no. 1, pp. 81–90, 1963.
- [8] S. Erlander and N. F. Stewart, *The gravity model in transportation analysis: theory and extensions*, vol. 3. Vsp, 1990.
- [9] J. de Dios Ortuzar and L. G. Willumsen, *Modelling transport*. John Wiley & sons, 2011.
- [10] D. P. McArthur, G. Kleppe, I. Thorsen, and J. Ubøe, “The spatial transferability of parameters in a gravity model of commuting flows,” *Journal of Transport Geography*, vol. 19, no. 4, pp. 596–605, 2011.
- [11] R. Patuelli, A. Reggiani, S. P. Gorman, P. Nijkamp, and F.-J. Bade, “Network analysis of commuting flows: A comparative static approach to german data,” *Networks and Spatial Economics*, vol. 7, no. 4, pp. 315–331, 2007.
- [12] L. Zhang, J. Cheng, and C. Jin, “Spatial interaction modeling of od flow data: Comparing geographically weighted negative binomial regression (gwnbr) and ols (gwolsr),” *ISPRS International Journal of Geo-Information*, vol. 8, no. 5, p. 220, 2019.
- [13] A. S. Fotheringham, C. Brunsdon, and M. Charlton, *Geographically weighted regression: the analysis of spatially varying relationships*. John Wiley & Sons, 2003.
- [14] A. R. da Silva and T. C. V. Rodrigues, “Geographically weighted negative binomial regression incorporating overdispersion,” *Statistics and Computing*, vol. 24, no. 5, pp. 769–783, 2014.
- [15] P. Davis, “Spatial competition in retail markets: movie theaters,” *The RAND Journal of Economics*, vol. 37, no. 4, pp. 964–982, 2006.
- [16] G. Bruno and G. Improta, “Using gravity models for the evaluation of new university site locations: A case study,” *Computers & Operations Research*, vol. 35, no. 2, pp. 436–444, 2008.
- [17] N. Wan, B. Zou, and T. Sternberg, “A three-step floating catchment area method for analyzing spatial access to health services,” *International Journal of Geographical Information Science*, vol. 26, no. 6, pp. 1073–1089, 2012.
- [18] A. S. Fotheringham, “A new set of spatial-interaction models: the theory of competing destinations,” *Environment and Planning A: Economy and Space*, vol. 15, no. 1, pp. 15–36, 1983.
- [19] M. De Beule, D. Van den Poel, and N. Van de Weghe, “An extended huff-model for robustly benchmarking and predicting retail network performance,” *Applied Geography*, vol. 46, pp. 80–89, 2014.
- [20] Y. Li and L. Liu, “Assessing the impact of retail location on store performance: A comparison of wal-mart and kmart stores in cincinnati,” *Applied Geography*, vol. 32, no. 2, pp. 591–600, 2012.
- [21] A. P. Masucci, J. Serras, A. Johansson, and M. Batty, “Gravity versus radiation models: On the importance of scale and heterogeneity in commuting flows,” *Physical Review E*, vol. 88, no. 2, p. 022812, 2013.
- [22] X. Liang, J. Zhao, L. Dong, and K. Xu, “Unraveling the origin of exponential law in intra-urban human mobility,” *Scientific reports*, vol. 3, p. 2983, 2013.
- [23] X.-Y. Yan, C. Zhao, Y. Fan, Z. Di, and W.-X. Wang, “Universal predictability of mobility patterns in cities,” *Journal of The Royal Society Interface*, vol. 11, no. 100, p. 20140834, 2014.
- [24] G. Spadon, A. C. de Carvalho, J. F. Rodrigues-Jr, and L. G. Alves, “Reconstructing commuters network using machine learning and urban indicators,” *Scientific reports*, vol. 9, no. 1, pp. 1–13, 2019.
- [25] M. Lorenzo and M. Matteo, “Od matrices network estimation from link counts by neural networks,” *Journal of Transportation Systems Engineering and Information Technology*, vol. 13, no. 4, pp. 84–92, 2013.
- [26] F. Scarselli, M. Gori, A. C. Tsoi, M. Hagenbuchner, and G. Monfardini, “The graph neural network model,” *Trans. Neur. Netw.*, vol. 20, no. 1, p. 6180, 2009.
- [27] J. Bruna, W. Zaremba, A. Szlam, and Y. Lecun, “Spectral networks and locally connected networks on graphs,” in *International Conference on Learning Representations (ICLR2014), CBLS, April 2014*, 2014.
- [28] M. Defferrard, X. Bresson, and P. Vandergheynst, “Convolutional neural networks on graphs with fast localized spectral filtering,” in *Proceedings of the 30th International Conference on Neural Information Processing Systems, NIPS16, (Red Hook, NY, USA)*, p. 38443852, Curran Associates Inc., 2016.
- [29] T. N. Kipf and M. Welling, “Semi-Supervised Classification with Graph Convolutional Networks,” in *Proceedings of the 5th International Conference on Learning Representations, ICLR ’17*, 2017.
- [30] X. Wang, Z. Zhou, F. Xiao, K. Xing, Z. Yang, Y. Liu, and C. Peng, “Spatio-temporal analysis and prediction of cellular traffic in metropolis,” *IEEE Transactions on Mobile Computing*, vol. 18, no. 9, pp. 2190–2202, 2019.
- [31] D. Zhu and Y. Liu, “Modelling spatial patterns using graph convolutional networks,” in *10th International Conference on Geographic Information Science (GIScience 2018)*, 2018.
- [32] P. Xie, T. Li, J. Liu, D. Shengdong, Y. Xin, and J. Zhang, “Urban flows prediction from spatial-temporal data using machine learning: A survey,” 2019. arXiv preprint arXiv: 1908.10218.
- [33] OpenStreetMap contributors, “Planet dump retrieved from <https://planet.osm.org>.” <https://www.openstreetmap.org>, 2017.
- [34] Murray Cox, “Inside Airbnb retrieved from <http://insideairbnb.com/get-the-data.html>.” <http://insideairbnb.com/get-the-data.html>, 2019.
- [35] Transport for London, “Open data retrieved from <https://tfl.gov.uk/info-for/open-data-users/>.” <https://tfl.gov.uk/info-for/open-data-users/>, 2019.
- [36] S. Ioffe and C. Szegedy, “Batch normalization: Accelerating deep network training by reducing internal covariate shift,” in *Proceedings of the 32nd International Conference on International Conference on Machine Learning - Volume 37, ICML15*, p. 448456, JMLR.org, 2015.
- [37] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, “Dropout: A simple way to prevent neural networks from overfitting,” *J. Mach. Learn. Res.*, vol. 15, no. 1, p. 19291958, 2014.
- [38] M. Lenormand, S. Huet, F. Gargiulo, and G. Deffuant, “A universal model of commuting networks,” *PloS one*, vol. 7, no. 10, 2012.
- [39] M. Lenormand, A. Bassolas, and J. J. Ramasco, “Systematic comparison of trip distribution laws and models,” *Journal of Transport Geography*, vol. 51, pp. 158–169, 2016.
- [40] G. Casiraghi, “Multiplex network regression: How do relations drive interactions?,” *arXiv preprint arXiv:1702.02048*, 2017.



- [41] T. Chen and C. Guestrin, "Xgboost: A scalable tree boosting system," in *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*, pp. 785–794, 2016.
- [42] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings* (Y. Bengio and Y. LeCun, eds.), 2015.
- [43] P. Velickovi, G. Cucurull, A. Casanova, A. Romero, P. Li, and Y. Bengio, "Graph attention networks," in *International Conference on Learning Representations*, 2018.
- [44] K. Xu, W. Hu, J. Leskovec, and S. Jegelka, "How powerful are graph neural networks?," in *International Conference on Learning Representations*, 2019.
- [45] K. Xu, C. Li, Y. Tian, T. Sonobe, K.-i. Kawarabayashi, and S. Jegelka, "Representation learning on graphs with jumping knowledge networks," in *Proceedings of the 35th International Conference on Machine Learning* (J. Dy and A. Krause, eds.), vol. 80 of *Proceedings of Machine Learning Research*, pp. 5453–5462, 2018.