Graph machine learning for activity scheduling: Location choice set formation

Luca Bataillard*[1] and Tim Hillel[1]

[1]Department of Civil, Environmental and Geomatic Engineering, University College London, UK

Short summary

This paper represents a first step towards a graph-based, econometric activity modelling framework, incorporating both location and spatial choices into the behavioural process. Current econometric activity based models typically focus predominantly on the temporal aspects of activity scheduling, leaving spatial decisions to agent-based microsimulation frameworks. Whilst recent work has established behavioural frameworks that incorporate location choice into the scheduling process, they have not yet been applied practically as (i) joint schedule optimisation across large numbers of candidate locations is computationally infeasible, and (ii) associated methods to establish individual location-choice sets for this context have not yet been developed. We address this gap by framing location choice set formation as a graph-based node classification problem. Using Graph Neural Networks (GNNs), we predict, for each individual, the probability of visiting different locations across a transport network graph enriched with point-of-interest, land-use, and individual socio-demographic features. We then demonstrate two different sampling strategies to generate choice set from the predicted node probabilities. We validate our approach on a synthetic case-study as well as real-world historic travel data from the Geneva public transport network, showing that GNNs can recover behaviourally plausible location choice sets while remaining computationally tractable.

Keywords: Activity-based modelling, location choice, graph machine learning

1 Introduction

Activity-based models of travel demand (ABM) assume that individuals schedule activities to meet their needs, and that travel demand emanates from these schedules (Bowman & Ben-Akiva, 2001). To construct daily schedules, ABMs require individuals to make choices across several dimensions, such as participation (whether to perform an activity), timing (start time and duration), location (where to go), and mode (how to get there). Recent econometric ABMs such as OASIS (Pougala et al., 2023) jointly optimise all choice dimensions. While this approach is behaviourally consistent, it is computationally demanding: the joint choice set grows as the product of the number of alternatives in each dimension, quickly becoming intractable. This phenomenon is primarily driven by the number of timing slots and candidate locations. In this work, we focus on the location dimension. In practice, location choice is often treated as a mechanistic process and deferred to the microsimulation stage, after schedules have been generated, thereby decoupling it from the joint optimisation. This disregards the influence of spatial factors, location and surrounding transport networks, on travel behaviour. Our method inverts this paradigm and produces a data-driven set of candidate locations to serve as a choice set in the joint optimization.

Location choice has long been studied as a stand-alone problem within discrete choice modelling. Residential location (Schirmer et al., 2014) and workplace location (Ho & Hensher, 2016) models represent individuals choosing between a finite set of zones, with utilities specified as functions of socio-demographics, accessibility and land use. To address the spatial correlation between location alternatives, researchers have used variants of the nested logit model (McFadden, 1978) or the spatially correlated logit model (Bhat & Guo, 2004). However, there has been limited focus on deeper integration of spatial choices in activity-based modelling pipelines.

Integrating location choice into ABM pipelines requires candidate sets of a tractable size: with large agent populations and many decisions per agent, evaluating all possible locations is computationally prohibitive. Several strategies have been proposed to reduce the size of location choice sets in ABMs: Pagliara & Timmermans (2009) categorise these into deterministic approaches, which filter candidates by activity type or maximum travel time, and stochastic approaches, which

sample alternatives uniformly or weighted by a measure of utility. Tsoleridis et al. (2022) integrate notions of activity-space with the stochastic approach by stratifying alternatives into three spatial strata constrained by time, space, and spatial awareness, then sampling from each stratum with probabilities proportional to observed choice frequencies. While this prioritises spatially relevant alternatives, the stratum boundaries still rest on strong behavioural assumptions. Machine learning (ML) methods have also been explored, though mainly for route and work destination choice. Yao & Bekhor (2022) train a variational autoencoder to generate the likely attributes of a route and retrieve the closest matching alternative, and Haghi et al. (2026) extend this to work destination choice. Both approaches are indirect: they generate attributes of plausible alternatives rather than identifying those alternatives in the network directly.

The highly stochastic nature of locations sets makes them hard to capture within econometric frameworks but, conversely, well suited to data-driven ML approaches. In particular, graph neural networks (GNN) are a natural fit for location choice set formation because the transport network already provides a graph structure over locations. GNNs process graph-structured data through iterative message passing, aggregating neighbourhood information to produce node, edge, or graph-level representations (Kipf & Welling, 2017; Hamilton et al., 2017), and have been applied broadly to mobility and transportation tasks such as traffic forecasting and urban planning (Rahmani et al., 2023). Most relevant to our work, Cheng et al. (2025) propose GNN-DCM for residential location choice, representing zones as nodes connected by spatial adjacency and training the GNN to output location utility, capturing spatial dependencies.

We frame activity location choice set formation as a graph-based ML problem and propose a discriminative, fully data-driven, GNN-based method for generating location choice sets. We treat location choice set formation as a complex decision-making process that is difficult to capture with econometric methods alone, and instead learn a black-box model that reduces the candidate location set to a tractable size while preserving behavioural validity. We construct a graph in which nodes represent locations and edges encode spatial adjacency or transport links. Using observed trip diaries, we label each node as visited or not visited per individual, train a GNN to predict visit probabilities from network features, individual socio-demographics, and home location, then sample from these probabilities to obtain the choice set. In this paper, we make the following contributions:

1. We formulate a theory of individual location choice set formation and define choice set formation as a graph node classification problem on a transport network graph.
2. We present a GNN-based pipeline that produces individual-level location choice sets by sampling from predicted visit probabilities.
3. We report results on two case studies: (i) a synthetic dataset; and (ii) a real travel data from the Geneva public transport operator (TPG), showing that GNNs recover behaviourally plausible location choice sets at tractable scale.

This paper focuses on how to formulate activity location choice set formation as a graph-based problem. Whilst we do investigate multiple ML prediction approaches, the ML implementation is not the primary contribution. We therefore give less detail on it for conciseness, focusing instead on the problem formulation and results.

2 Methodology

We base our method on a behavioural representation of how individuals form location consideration sets, drawing on the theory of consideration sets presented by Shocker et al. (1991). Consider an individual who is planning when and where to carry out grocery shopping in their working day. In practice, rather than evaluating every possible supermarket in the city, individuals typically choose from one of a restricted set of specific, considered locations. For shopping, this might be their local convenience store, a larger supermarket on their commute from work, or a hypermarket outside the city. We call this reduced set of locations the *consideration set*. It is shaped by spatial factors such as distance or transport accessibility from home or other key locations (e.g. work, leisure, etc.), location-specific attributes such as product range, and price, as well as individual habits, social context, and past behaviour. We treat consideration sets as stable over the medium term, typically evolving when an individual goes through a major life change, such as moving house or changes in job or household composition.

Our method generates these location choice sets from past travel data, producing candidate sets that can serve as input to an ABM. While not explicitly covered in this paper, our intention is that the consideration sets generated by our method can serve as input for an ABM, as they require a set of candidate locations to consider when generating schedules. This paper represents the first step towards this integrated approach. We first define a graph formulation of travel diary data, then present a GNN approach to predicting visit probabilities from the graph formulated travel data, finally sampling from predicted visit probabilities to form the desired location choice sets.

Graph-based scheduling

We frame activity scheduling as a graph-based problem to model both temporal and spatial dimensions simultaneously. Consider a graph representation of a transport network in which *nodes* correspond to to physical locations (e.g. neighbourhoods, zones, or public transport stops) and *edges* represent some degree of transport connection between locations (e.g. spacial adjacency or direct public transport links). Scheduling choices then map naturally onto node or edge labelling problems: in this paper, we assign binary labels and/or probabilities for whether a node is included in the location choice set, but the same representation extends to more complex problems such as multinomial probabilities over different activity types, start-times, or durations of activities at each node.

We model the decision making process in three steps; starting from the full set of locations and narrowing down sequentially. The *network graph* (NG) represents all candidate locations in a region: nodes correspond to physical locations (neighbourhoods, zones, or public transport stops) and edges connect spacially adjacent nodes or nodes sharing a direct transport connection. Each node carries features that contain relevant information such as point of interest (POI) counts and land use composition. The NG is fully observed and fixed for a given region. When planning activities, an individual does not consider all nodes in the NG. Instead, they filter the NG down to a *consideration graph* (CG_i), a latent subgraph shaped by individual attributes (socio-demographics, home and work location) and network attributes (travel times, amenities). From the CG_i, the individual selects the locations that maximise their utility, yielding the *realised graph* (RG_i) of actually visited nodes. These three graphs satisfy the relation $\mathrm{RG}_i \subseteq \mathrm{CG}_i \subseteq \mathrm{NG}$

Within this framework, location choice set formation can be reduced to an individual-specific node classification problem over a shared NG[1]. In this method, we combine network information, common to all individuals, and individual-specific attributes. Accordingly, we distinguish three types of input: (i) *network-level* information, such as transport connections, land-use mix and POI counts by category; (ii) *individual-spatial-level* information, specific to both a location and an individual, such as home location and observed location visits; and (iii) non-spatial *individual-level* information, such as socio-demographics or car ownership. The NG is shared across individuals, but can be made individual-specific through addition of individual-level and individual-spatial-level features.

Graph formulation of location choice set formation

We construct the graph representation from spatial data and travel diaries. Consider a region with N locations and a travel diary recording the trips of I individuals. We build a static NG graph $\mathcal{G} = (\mathcal{N}, \mathcal{E}, X)$ where nodes $\mathcal{N}$ is the set of locations and edges $\mathcal{E}$ encode spatial adjacency or transport connections. Each node is associated with network-level features collected in a matrix $X \in \mathbb{R}^{N \times d_1}$.

Individual-level features are represented as a matrix $J \in \mathbb{R}^{I \times d_2}$, where row J_i corresponds to individual i. For each individual i, we also define individual-spatial-level features $S_i \in \mathbb{R}^{N \times d_3}$, where each row corresponds to a node and encodes features specific to the pair (i, n) such as an indicator for home location. Observed visits are represented by the binary vector $y_i \in \{0, 1\}^N$, where $y_{i,n} = 1$ if individual i visits node n.

For a given individual i, we construct an annotated graph

$$\mathcal{G}_i = (\mathcal{N}, \mathcal{E}, X_i, y_i) \tag{1}$$

[1]This approach assumes that choice set membership is a function of visitation probability, rather than the other way around (Manski & Lerman, 1977). We will revisit this assumption formally in future work when integrating with a scheduling model.

where the node feature matrix is defined as

$$X_i = [X \parallel S_i] \tag{2}$$

with $[\cdot \parallel \cdot]$ denoting column-wise concatenation. In this way, the shared NG structure is augmented with individual-specific node features X_i and global individual-level information J_i, and the prediction task reduces to node-wise classification with labels y_i.

GNNs for location choice set formation

We model the node-wise classification problem defined above using GNNs. GNNs operate by iterative message passing: at each layer, nodes aggregates feature vectors from their neighbours and updates their representations. After K layers, each node embedding captures information from its K-hop neighbourhood. This property suits our problem well, since the attractiveness of a location depends on the spatial context provided by nearby locations and transport connections.

We augment the network-level node features X with positional encodings: random walk encodings (Dwivedi et al., 2022) and Laplacian eigenvector encodings (Dwivedi et al., 2023). This provides the model with additional local and global graph structure beyond aggregated neighbourhood features.

Let f_θ denote the GNN. For each individual i, it models the conditional probability

$$\Pr(y_{i,n} | \mathcal{G}_i, J_i) \tag{3}$$

that individual i visits node n. The model takes as input the annotated graph $\mathcal{G}_i = (\mathcal{N}, \mathcal{E}, X_i)$ and the individual-level features J_i, which are incorporated as a graph feature that is concatenated to the final node embeddings.

The GNN operates in three stages: f_θ (i) applies a shared input multi-layer perceptron (MLP) to each node to transform node features; (ii) performs K rounds of graph convolution to propagate neighbourhood information; and (iii) applies a final node-level MLP with sigmoid activation to output per-node visit probabilities:

$$\hat{y}_{i,n} = f_\theta(\mathcal{G}_i, J_i). \tag{4}$$

We train the model using a binary cross entropy (BCE) loss function with added ℓ_1 regularisation to encourage sparsity.

Sampling choice sets from predicted probabilities

We investigate two methods to draw a choice set of m locations for individual i from the predicted probabilities $\hat{y}_i$. In *Poisson sampling*, each node n is included independently with probability $\hat{y}_{i,n}$, yielding a set whose expected size equals $\sum_n \hat{y}_{i,n}$; the size is thus stochastic and controlled implicitly via the predicted probabilities. In *probability-proportional-to-size (PPS) sampling*, we normalise $\hat{y}_i$ to a proper distribution and draw exactly m locations without replacement, giving a choice set of fixed, analyst-specified size. Poisson sampling is simpler and preserves the probabilistic interpretation of each node independently, whereas PPS sampling is preferable when a fixed choice set size is required by a downstream ABM. A shared limitation is that draws are made independently across nodes, ignoring spatial correlations between nearby locations. We plan to investigate more principled sampling approaches in future work.

3 Results and discussion

We perform two experiments to validate our framework. The synthetic experiment, where ground-truth consideration graphs CG_i are known by construction, accomplishes two goals: (i) assess whether the GNN recovers the correct data generation process; and (ii) investigate different GNN architectures in a controlled setting. The real-world experiment, as an early proof of concept, then applies the best-performing architecture from the synthetic experiment to historic travel data from TPG.

Synthetic experiment

This experiment aims to show that GNN can recover behaviourally plausible choice sets and capture spatial dependencies on a fixed graph with synthetic travel data, and to select the best architec-

ture before applying it to real data. By specifying the exact data generation process, including network, individual, and individual-spatial factors, we can measure how well each model recreates the underlying decision process and compare against baselines.

We generate a fixed random graph of N nodes, assigning each a scalar feature $X_n \in \mathbb{R}$ drawn from a uniform distribution. For each synthetic individual, we draw an individual feature $J_i \in \mathbb{R}$ from a uniform distribution and sample a home location h_i uniformly at random. We then define an individual-specific utility function over nodes combining node features X_n, individual attributes J_i, home node feature $S_i = X^1_{h_i} \in \mathbb{R}$, and a distance penalty from the home location:

$$\eta_{i,n} = \beta_1 J_i (X_n - \bar{X_n}) - \beta_2 \left| X_n - S_i \right| - \beta_3 J_i \text{dist}(n, h_i) - \alpha \tag{5}$$

where $\beta_1, \beta_2, \beta_3, \alpha \in \mathbb{R}$ are fixed parameters. A node is labelled as visited, i.e.$y_{i,n} = 1$, if its utility plus a Gumbel noise draw exceeds zero, yielding a sparse, spatially structured set of visited locations that mirrors realistic travel behaviour. Figure 1 illustrates this process for an example individual.

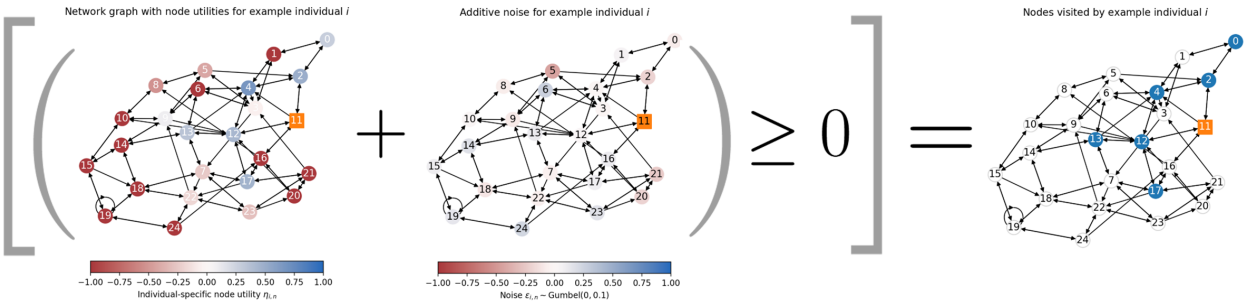


Figure 1: Synthetic data generation. Each node's individual-specific utility $\eta_{i,n}$ is computed from network features, individual attributes, and distance from the home node (orange square). A node is labelled as visited (blue) if its utility plus an independent Gumbel noise draw is non-negative.

We compare three model classes: (i) an MLP with access only to node features X_n and individual features J_i, with no home-location or graph structure information (*MLP-base*); (ii) an MLP augmented with all components of the utility function, serving as a best-case oracle baseline (*MLP-full*); and (iii) GNN architectures that receive node features, individual features, and home indicators, and learn graph structure through message passing (*GNN*). We evaluate GNNs with varying numbers of GCN (Kipf & Welling, 2017) and GAT (Veličković et al., 2018) convolution layers. Table 1 shows that certain GNN architectures match the node classification performance of *MLP-full*, demonstrating that message passing recovers spatial dependencies in location choice sets without an explicitly specified utility function. We find that the best-performing GNNs use a number of message-passing layers equal to the graph diameter and that GAT convolutions consistently outperform GCN convolutions. Additionally, residual connections between convolution layers and a skip connection from initial node features to the final node MLP head prove essential to matching baseline performance. These findings are consistent with recent findings in the literature from Luo et al. (2025). Figure 2 shows the predicted visit probabilities from GATSkipRes alongside the resulting Poisson and PPS choice sets for the same individual.

Table 1: BCE of tested GNN architectures and baselines on the synthetic experiment, lower is better. For GNN models, we report the best-performing number of convolutional layers.

Model	**Conv. type**	**Num. conv. layers**	**Residual conn.**	**Skip conn.**	**BCE ↓**
MLP-full	–	–	–	–	**0.1108**
GATSkipRes	GAT	5	Yes	Yes	**0.1151**
GATSkip	GAT	5	No	Yes	0.1157
GCNSkipRes	GCN	5	Yes	Yes	0.1218
GCNSkip	GCN	5	No	Yes	0.1318
GCN	GCN	4	Yes	No	0.1832
MLP-base	–	–	–	–	0.2075

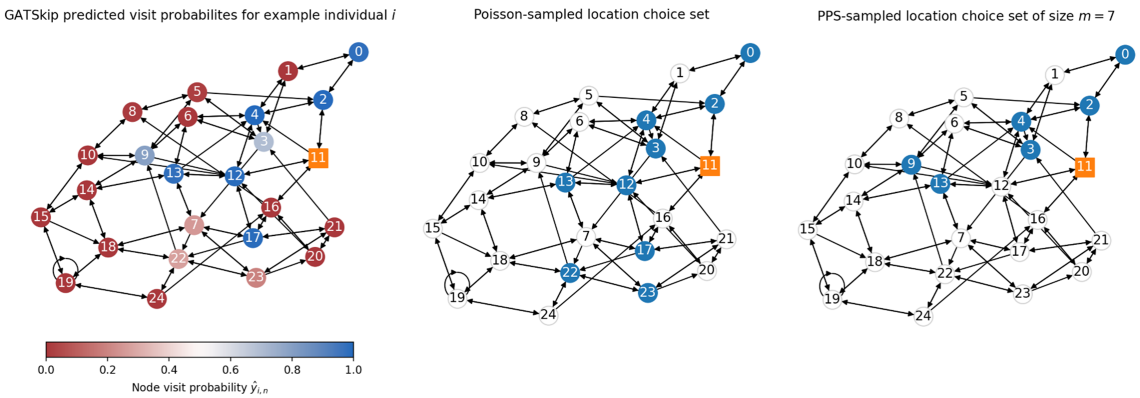


Figure 2: GATSkipRes predicted visit probabilities for the example individual from Figure 1 (left), and the resulting location choice sets under Poisson sampling (centre) and PPS sampling (right).

Geneva transport experiment

We apply our method to real-world travel diary data from TPG as a proof of concept, demonstrating that the graph construction yields meaningful predictions on real data. We show that GNN models outperform baseline models that do not use information from surrounding neighbourhoods, which implies that the graph formulation of location choice set formation is worthwhile and that future work should investigate model architectures, richer graph formulations, and integration into an ABM pipeline.

The TPG travel diary covers 37K trips made by 16K individuals across 470 contiguous statistical zones (SITG, 2026) which we treat as nodes in NG. Each individual visits between 2 and 4 unique nodes. The diary includes trip purposes, from which we infer home locations, but no socio-demographic or other individual-level information.

We build the network graph using the city2graph library (Sato, 2025): nodes are connected by edges encoding spatial adjacency[2]. We enrich node features with subsector area, population and job densities from the Geneva statistical office (SITG, 2025), POI counts and land-use densities from Overture Maps (Overture Maps Foundation, 2026; OpenStreetMap contributors, 2026), and positional encodings. Distances between subsector centroids are added as edge features. We construct individual-specific graph copies by appending a home-location indicator to the node features and setting node visit indicators as labels. Figure 3 shows the graph structure and displays the visited locations of an example individual in the training set.

We train the best-performing architecture from the synthetic experiment, GATSkipRes, with 128-dimensional hidden layers, a 1-layer input MLP, 8 GAT convolution layers, and a 3-layer MLP classification head. As a baseline, we train a 3-layer node-level MLP that receives the same node features but no neighbourhood information. Both models are trained with BCE loss and L1 regularisation over 50 epochs with dropout probability 0.2. Since visited nodes are a small minority, we weight positive examples more strongly in the BCE loss. We compare the ML models against four simple probabilistic baselines computed from training set frequencies: (i) *Uniform*, which assigns a 0.5 probability to every node; (ii) *GlobalFrequency*, which outputs the global visit rate; (iii) *NodeFrequency*, which outputs the per-node visit frequency; and (iv) *HomeConditioned*, which outputs per-node visit frequency conditioned on the home node.

Table 2 shows that GATSkipRes achieves the best BCE, Precision@5, and Recall@5 on the test set, outperforming all probabilistic baselines by a substantial margin. The gap between GATSkipRes and MLP is consistent but small –BCE of 0.1538 vs. 0.1579, Precision@5 of 0.2234 vs. 0.2225– which we attribute primarily to the limited individual-level information available in this dataset. With only home location as individual context, the GNN has little individual-specific signal to propagate through the graph; we expect the gap to widen once richer features are available. Figure 4 illustrates this qualitatively: GATSkipRes assigns higher visit probabilities to nodes surrounding the home location, whereas MLP does not, confirming that the GNN leverages spatial correlations that the MLP cannot capture.

[2]Edges encoding public transport connections will be added in future steps.

Figure 5 shows location choice sets produced by Poisson and PPS sampling from the GATSkipRes predicted probabilities. A formal evaluation of both sampling procedures is left for future work, in particular once integrated into an ABM pipeline.

Table 2: Benchmark of GATSkipRes against MLP and probabilistic baselines on the TPG dataset. Models are evaluated on the test set using weighted BCE (lower is better) and Precision and Recall at $k = 5$ (higher is better).

Model	**BCE** $\downarrow$	**Precision@5** $\uparrow$	**Recall@5** $\uparrow$
Probabilistic baselines			
Uniform	0.7402	0.0057	0.0114
GlobalFrequency	0.3875	0.0057	0.0114
NodeFrequency	0.3418	0.0547	0.1070
HomeConditioned	0.4637	0.2133	0.4405
ML models			
MLP	0.1579	0.2225	0.4614
GATSkipRes	**0.1538**	**0.2234**	**0.4632**

4 Conclusions

We have taken a first step towards a graph-based, data-driven framework for integrating location choice into econometric ABMs. We reframe activity scheduling as a graph problem, representing the transport network as nodes and edges, and casting location choice set formation as individual-specific node classification. We formalise the problem via three nested graphs, the network graph, the latent consideration graph, and the realised graph, and train a GNN to identify candidate locations from network structure and individual context without explicitly specified utilities. Experiments on both synthetic and real-world data confirm that the method produces behaviourally plausible choice sets: on the synthetic experiment, GATSkipRes matches oracle performance, and on the Geneva dataset it outperforms all ML and probabilistic baselines.

Several directions remain open for future work. Next steps include: (i) validating the approach on additional real-world datasets that include household and individual-level socio-demographics, to assess how much individual context improves choice set quality; (ii) enriching the graph with heterogeneous node and edge types, additional POI and land-use features, and richer transport network attributes, and conducting a thorough architecture comparison on real-world data; (iii) developing a systematic evaluation framework benchmarking our method against existing deterministic and stochastic choice set formation approaches; (iv) integrating the generated location choice sets into a downstream ABM to measure the end-to-end impact on schedule quality; (v) investigating more advanced generative ML approaches and improving the sampling procedure; and (vi) expanding this method beyond location choice sets to include timings and activity types, moving towards an integrated spatio-temporal modelling framework.

Acknowledgements

We thank TPG for providing the travel data used in this study.

References

Bhat, C. R., & Guo, J. (2004). A mixed spatially correlated logit model: formulation and application to residential choice modeling. *Transportation Research Part B: Methodological*, *38*(2), 147–168.

Bowman, J., & Ben-Akiva, M. (2001). Activity-based disaggregate travel demand model system with activity schedules. *Transportation Research Part A: Policy and Practice*, *35*(1), 1-28. Retrieved from `https://www.sciencedirect.com/science/article/pii/S0965856499000439` doi: https://doi.org/10.1016/S0965-8564(99)00043-9

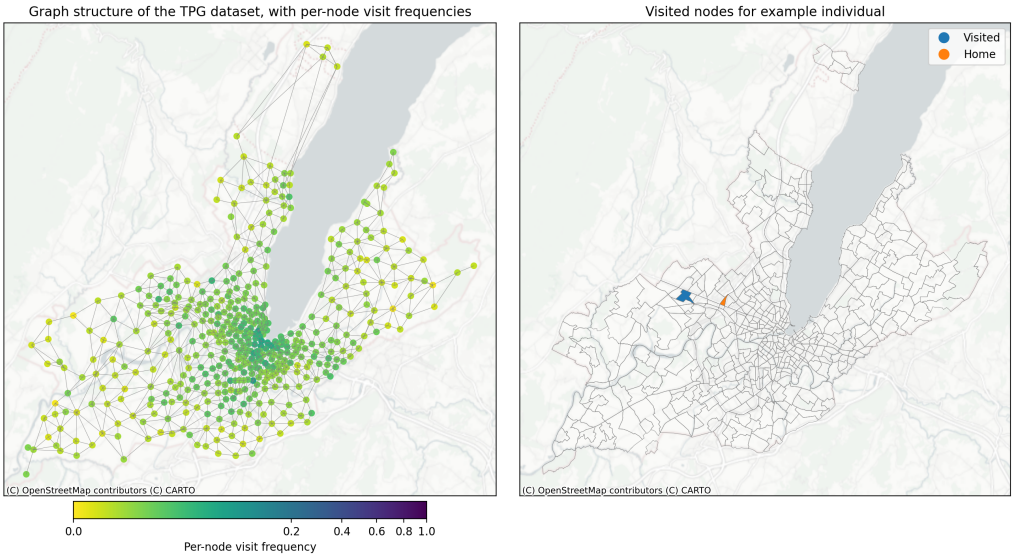


Figure 3: Left: graph structure of the TPG dataset, with per-node visit frequencies computed on the training set. Right: visited and home locations for an example individual in the test set.

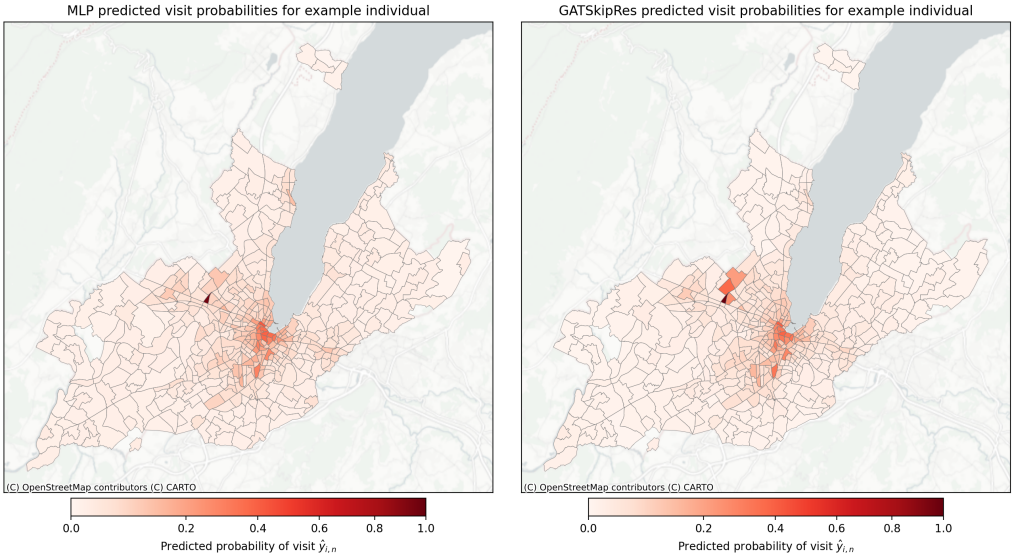


Figure 4: Comparison of predictions made by the MLP and GATSkipRes models for the example individual in Figure 3.

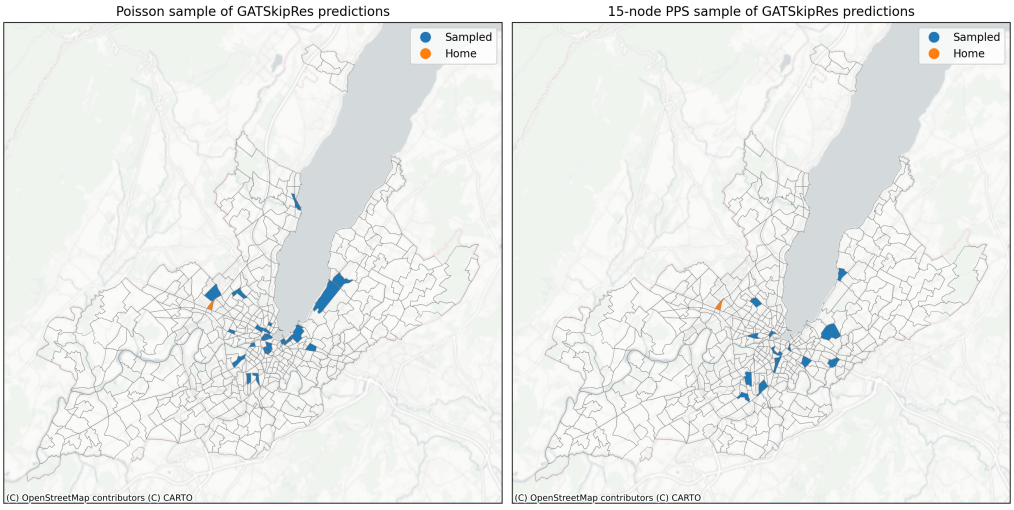


Figure 5: Demonstration of location choice set formation using the Poisson and PPS sampling procedures on the GATSkipRes predictions from Figure 4.

Cheng, Z., Hu, L., Bu, Y., Zhou, Y., & Wang, S. (2025). *Graph neural networks for residential location choice: connection to classical logit models.* Retrieved from https://arxiv.org/abs/2507.21334

Dwivedi, V. P., Joshi, C. K., Luu, A. T., Laurent, T., Bengio, Y., & Bresson, X. (2023). Benchmarking graph neural networks. *Journal of Machine Learning Research*, *24*(43), 1–48.

Dwivedi, V. P., Luu, A. T., Laurent, T., Bengio, Y., & Bresson, X. (2022). Graph neural networks with learnable structural and positional representations. In *International conference on learning representations.*

Haghi, H., Patterson, Z., & Farooq, B. (2026). Choice set generation in work destination choice modeling with variational autoencoders. *Transportation Research Record*, 03611981251409202. Retrieved from https://doi.org/10.1177/03611981251409202 doi: 10.1177/03611981251409202

Hamilton, W., Ying, Z., & Leskovec, J. (2017). Inductive representation learning on large graphs. In I. Guyon et al. (Eds.), *Advances in neural information processing systems* (Vol. 30). Curran Associates, Inc. Retrieved from https://proceedings.neurips.cc/paper_files/paper/2017/file/5dd9db5e033da9c6fb5ba83c7a7ebea9-Paper.pdf

Ho, C. Q., & Hensher, D. A. (2016). A workplace choice model accounting for spatial competition and agglomeration effects. *Journal of Transport Geography*, *51*, 193-203. Retrieved from https://www.sciencedirect.com/science/article/pii/S0966692316000065 doi: https://doi.org/10.1016/j.jtrangeo.2016.01.005

Kipf, T. N., & Welling, M. (2017). Semi-supervised classification with graph convolutional networks. In *International conference on learning representations.* Retrieved from https://openreview.net/forum?id=SJU4ayYgl

Luo, Y., Shi, L., & Wu, X.-M. (2025). Can classic gnns be strong baselines for graph-level tasks? simple architectures meet excellence. In *International conference on machine learning* (pp. 41290–41310).

Manski, C. F., & Lerman, S. R. (1977). The estimation of choice probabilities from choice based samples. *Econometrica*, *45*(8), 1977–1988. Retrieved 2026-03-25, from http://www.jstor.org/stable/1914121

McFadden, D. (1978). Modeling the choice of residential location. *Transportation Research Record*(673).

OpenStreetMap contributors. (2026). *Openstreetmap.* https://www.openstreetmap.org. (Accessed: 2026-03-20. Data available under the Open Database License (ODbL) v1.0 at https://www.openstreetmap.org/copyright)

Overture Maps Foundation. (2026). *Overture Maps Open Map Data – Place and land use data.* https://www.overturemaps.org. (Retrieved: 2026-03-20 from s3://overturemaps-us-west-2/release/2026-03-18.0, Licensed under CDLA Permissive v2.0)

Pagliara, F., & Timmermans, H. (2009). Choice set generation in spatial contexts: a review. *Transportation Letters*, *1*(3), 181–196. Retrieved from https://doi.org/10.3328/TL.2009.01.03.181-196 doi: 10.3328/TL.2009.01.03.181-196

Pougala, J., Hillel, T., & Bierlaire, M. (2023). Oasis: Optimisation-based activity scheduling with integrated simultaneous choice dimensions. *Transportation Research Part C: Emerging Technologies*, *155*, 104291. Retrieved from https://www.sciencedirect.com/science/article/pii/S0968090X23002802 doi: https://doi.org/10.1016/j.trc.2023.104291

Rahmani, S., Baghbani, A., Bouguila, N., & Patterson, Z. (2023). Graph neural networks for intelligent transportation systems: A survey. *IEEE Transactions on Intelligent Transportation Systems*, *24*(8), 8846-8885. doi: 10.1109/TITS.2023.3257759

Sato, Y. (2025). *City2graph: Transform geospatial relations into graphs for spatial network analysis and graph neural networks.* Retrieved from https://github.com/c2g-dev/city2graph doi: 10.5281/zenodo.15858845

Schirmer, P. M., Van Eggermond, M. A., & Axhausen, K. W. (2014). The role of location in residential location choice models: a review of literature. *Journal of Transport and Land Use*, *7*(2), 3–21.

Shocker, A. D., Ben-Akiva, M., Boccara, B., & Nedungadi, P. (1991). Consideration set influences on consumer decision-making and choice: Issues, models, and suggestions. *Marketing Letters*, *2*(3), 181–197. Retrieved 2026-03-12, from http://www.jstor.org/stable/40216215

SITG. (2025). *Agglo – carroyage de population et emplois.* https://sitg.ge.ch/donnees/agglo-carreau-200. (Identifier: AGGLO_CARREAU_200. Published 2025-06-12. Freely accessible. Sources: STATPOP/STATENT (OFS, 2019), FiLoSoFi (INSEE, 2019), SIRENE (2020). Système d'information du territoire à Genève (SITG). Accessed: 2026-03-20)

SITG. (2026). *Sous-secteurs statistiques (GIREC).* https://sitg.ge.ch/donnees/geo-girec. (Identifier: GEO_GIREC. Published 2026-01-31. Freely accessible. Système d'information du territoire à Genève (SITG). Accessed: 2026-03-20)

Tsoleridis, P., Choudhury, C. F., & Hess, S. (2022). Utilising activity space concepts to sampling of alternatives for mode and destination choice modelling of discretionary activities. *Journal of Choice Modelling*, *42*, 100336. Retrieved from https://www.sciencedirect.com/science/article/pii/S1755534521000695 doi: https://doi.org/10.1016/j.jocm.2021.100336

Veličković, P., Cucurull, G., Casanova, A., Romero, A., Liò, P., & Bengio, Y. (2018). Graph attention networks. In *International conference on learning representations.*

Yao, R., & Bekhor, S. (2022). A variational autoencoder approach for choice set generation and implicit perception of alternatives in choice modeling. *Transportation Research Part B: Methodological*, *158*, 273-294. Retrieved from https://www.sciencedirect.com/science/article/pii/S019126152200042X doi: https://doi.org/10.1016/j.trb.2022.02.015