Supplementary Information

Hierarchical organization of pedestrian mobility is stable across cities

Fatemeh Nourmohammadi, Taha H. Rashidi, Meead Saberi

1 Overview of the inference framework

Pedestrian demand was inferred using an integrated framework that combines trip generation, destination choice, route assignment, and origin–destination (OD) demand calibration. The framework builds directly on previously developed component models for transferable trip generation [Nourmohammadi et al., 2025b], graph-structured destination choice [Nourmohammadi et al., 2025a], empirically estimated pedestrian route choice [Lilasathapornkit et al., 2026], and high-dimensional OD calibration [Osorio, 2019]. In the present study, these components are applied sequentially and without local re-estimation of the behavioral model parameters in the target city.

The framework proceeds in four stages. See Figure 1. First, transferable trip generation and destination choice models are applied to infer a prior origin–destination (OD) demand. Second, this demand is assigned to the pedestrian network using a probabilistic route choice model to obtain link-level flows. Third, the simulated flows are compared with observed pedestrian counts, and the OD matrix is updated through a regularized inverse calibration procedure. Finally, the calibrated demand is represented as a weighted directed graph and analyzed to determine whether calibration primarily affects flow magnitudes, the structure of spatial interactions, or both. The resulting networks are further analyzed in terms of their hierarchical organization, including the concentration of flows across nodes and the extent to which this structure differs from null expectations.

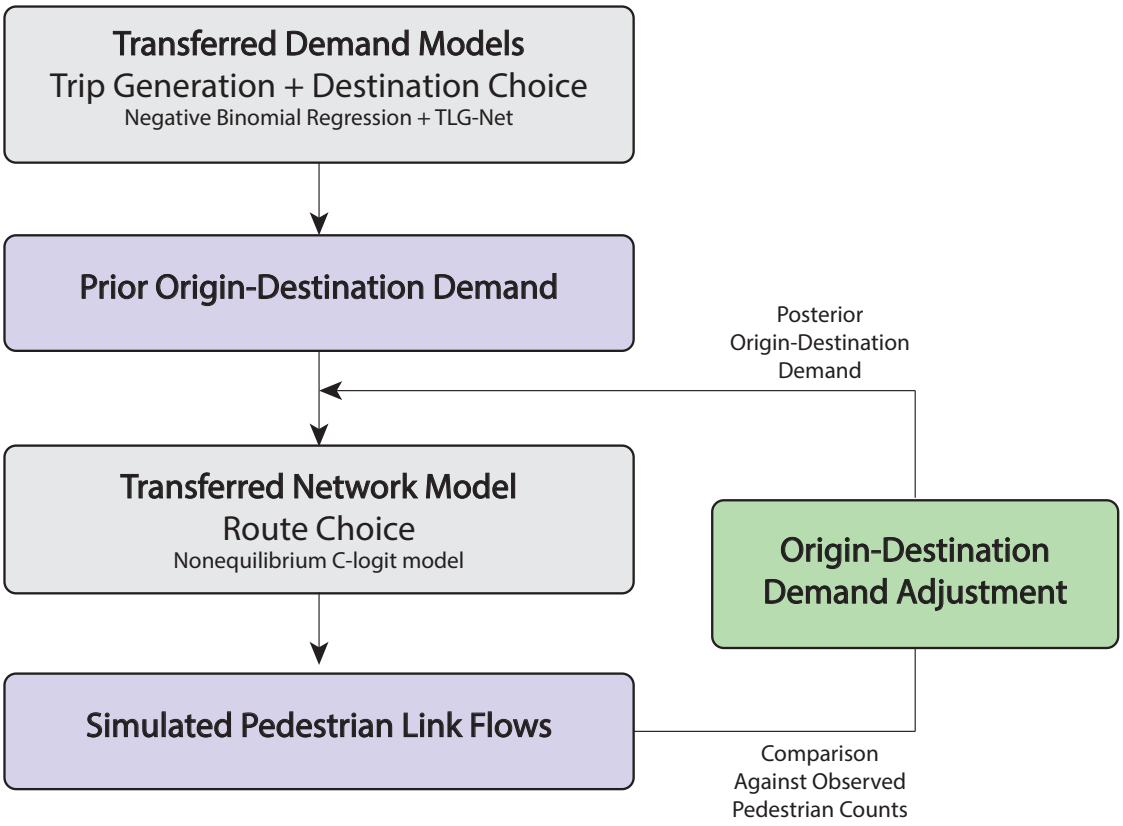

Figure 1: Overview of the pedestrian demand inference framework. Transferred trip generation and destination choice models are used to construct a prior origin–destination (OD) demand, which is assigned to the pedestrian network using a route choice model to generate link-level flows. These simulated flows are then compared with observed pedestrian counts, and the OD matrix is updated through a regularized calibration procedure. The resulting posterior OD demand is subsequently analyzed to evaluate the structural organization of pedestrian movement.

2 Detailed methods

Let o and d denote origin and destination zones, respectively. The set of candidate destinations for origin o is denoted by C_o, and the set of feasible routes between o and d is denoted by R_{od}. In the graph-based analysis, the number of nodes and edges are denoted by $n = |V|$ and $m = |E|$, respectively.

2.1 Trip generation

Pedestrian trip generation follows the transferable Negative Binomial formulation proposed by Nourmohammadi et al. [2025b]. Let T_i denote the number of pedestrian trips generated at zone i. Trip counts are modeled as

$$T_i \sim \text{NegBin}(\mu_i, \alpha), \tag{1}$$

where μ_i is the expected number of trips produced in zone i and $\alpha > 0$ is the dispersion parameter. The conditional mean is specified through the log-link

$$\mu_i = \exp\left(\beta_0 + \sum_m \beta_m x_{mi}\right), \tag{2}$$

where x_{mi} is the value of explanatory variable m for zone i, and β_m is the corresponding coefficient.

In this study, the trip generation coefficients estimated in the source city are transferred directly to the target city without re-estimation. The resulting zonal trip productions serve as the initial demand input to the framework.

2.2 Destination choice

Destination choice follows the graph-structured OD modeling approach of Nourmohammadi et al. [2025a]. For each origin o, the probability of selecting destination d is derived from a utility-based score s_{od} that captures destination attractiveness, travel impedance, and graph-based spatial features. Choice probabilities are computed using a sparse transformation:

$$p_{od} = \text{sparsemax}(s_{od}) = \max(s_{od} - \tau_o, 0), \tag{3}$$

where τ_o is a normalization threshold chosen so that the probabilities sum to one over the destination choice set C_o, that is, $\sum_{d \in C_o} p_{od} = 1$.

The OD flow from origin o to destination d is then given by

$$T_{od} = T_o \, p_{od}, \tag{4}$$

where T_o is the total number of trips generated at origin o. In the subsequent calibration stage, these OD flows are also denoted by q_{od}.

As in trip generation, the destination choice model is transferred to the target city without re-estimation, so that its output serves as a prior demand estimate rather than a locally fitted OD matrix.

2.3 Route choice and link-flow assignment

The inferred OD demand is assigned to the pedestrian network using the C-logit route choice model estimated by Lilasathapornkit et al. [2026]. For each OD pair (o, d), a finite set of feasible routes R_{od} is generated. The probability of selecting route $i \in R_{od}$ is

$$P_{od}(i) = \frac{\exp\left[\mu \left(V_{iod} - CF_{iod}\right)\right]}{\sum_{j \in R_{od}} \exp\left[\mu \left(V_{jod} - CF_{jod}\right)\right]}, \tag{5}$$

where V_{iod} is the perceived utility of route i, CF_{iod} is the commonality factor accounting for overlap among alternatives, and μ is the scale parameter.

Route utility is specified as

$$V_{iod} = \sum_k \beta_k z_{k,iod}, \tag{6}$$

where $z_{k,iod}$ denotes route attribute k for route i connecting origin o and destination d, and β_k is the corresponding coefficient. In this study, route attributes include route length, maximum gradient, number of turns, number of points of interest, Green View Index, and the number of crossings.

Link-level pedestrian volumes are obtained by summing the probabilities and OD flows across all routes that traverse each link.

The route choice model is not re-estimated in the target city. This preserves a controlled comparison across transfer experiments, so that differences in final link flows can be attributed to the transferred demand inputs and their subsequent calibration.

2.4 OD demand adjustment

Following initial demand estimation and route assignment, the prior OD matrix is updated using observed pedestrian counts. Let $\mathbf{q}^0 = \{q^0_{od}\}$ denote the prior OD demand obtained from the transferred trip generation and destination choice models, and let $\hat{\mathbf{q}} = \{\hat{q}_{od}\}$ denote the adjusted OD demand after calibration. Let y_i denote the observed pedestrian count on link i, and let $\hat{y}_i(\hat{\mathbf{q}})$ denote the simulated flow on link i obtained by assigning $\hat{\mathbf{q}}$ through the route choice model.

Following Osorio [2019], the calibration problem is formulated as the regularized optimization

$$\min_{\hat{\mathbf{q}}} \ \frac{1}{|I|}\sum_{i\in I}\left(y_i - \hat{y}_i(\hat{\mathbf{q}})\right)^2 + \delta\,\frac{1}{|Z|}\sum_{(o,d)\in Z}\left(q^0_{od} - \hat{q}_{od}\right)^2, \tag{7}$$

subject to

$$0 \le \hat{q}_{od} \le q^{\max}_{od}, \qquad \forall (o,d) \in Z, \tag{8}$$

where I is the set of observed links, Z is the set of OD pairs, δ is a regularization parameter, and $q^{\max}_{od}$ is an upper bound on the adjusted flow for OD pair (o,d).

The first term minimizes the discrepancy between observed and simulated link flows. The second term penalizes deviations from the transferred prior demand. This regularization is needed because the problem is underdetermined. The number of OD pairs is much larger than the number of observed link counts. In this context, the transferred demand model acts as prior information, and calibration adjusts this prior only to the extent needed to match observed counts. This formulation ensures that calibration adjusts flow magnitudes while preserving the structural organization of the transferred demand, enabling subsequent analysis of whether hierarchical patterns are retained after adjustment.

2.5 Graph-based diagnostics of demand structure

To distinguish changes in flow magnitude from changes in structural organization, OD demand is represented as a weighted directed graph $G = (V, E, W)$, where V is the set of zones, E is the set of directed OD connections, and $W = [w_{ij}]$ is the matrix of edge weights, with w_{ij} denoting the flow from node i to node j. The binary adjacency matrix is defined by $a_{ij} = 1$ if an edge from i to j exists and $a_{ij} = 0$ otherwise.

Several standard graph-theoretic metrics are used [Newman, 2010]. For directed graphs, the in-degree and out-degree of node i are defined as

$$k^{\text{in}}_i = \sum_j a_{ji}, \qquad k^{\text{out}}_i = \sum_j a_{ij}, \tag{9}$$

which quantify the number of incoming and outgoing connections, respectively. The intensity of interactions is captured using node strength,

$$s_i = \sum_j w_{ij} + \sum_j w_{ji}, \tag{10}$$

which represents the total incoming and outgoing flow associated with node i.

At the network level, density is computed as

$$D = \frac{m}{n(n-1)}, \tag{11}$$

where n is the number of nodes and m is the number of directed edges. Clustering is computed on the undirected projection of the graph. Let k_i denote the undirected degree of node i and let e_i denote the number of edges among its neighbors. The clustering coefficient is

$$C_i = \frac{2e_i}{k_i(k_i-1)}. \tag{12}$$

A weighted clustering coefficient is also reported:

$$C_i^w = \frac{1}{s_i(k_i-1)} \sum_{j,h} \frac{w_{ij}+w_{ih}}{2} a_{ij}a_{ih}a_{jh}. \tag{13}$$

The importance of a node in mediating flows is measured using betweenness centrality:

$$C_B(v) = \sum_{s \neq v \neq t} \frac{\sigma_{st}(v)}{\sigma_{st}}, \tag{14}$$

where σ_{st} is the total number of shortest paths between nodes s and t, and $\sigma_{st}(v)$ is the number of such paths passing through node v. Shortest-path distances are computed using inverse edge weights,

$$d_{ij} = \frac{1}{w_{ij}+\epsilon}, \tag{15}$$

where ϵ is a small positive constant introduced to avoid division by zero.

To summarize the intensity and heterogeneity of OD flows, edge-weight statistics are also reported:

$$\bar{w} = \frac{1}{m} \sum_{(i,j)\in E} w_{ij}, \qquad \sigma_w = \sqrt{\frac{1}{m} \sum_{(i,j)\in E} (w_{ij}-\bar{w})^2}, \qquad w_{\max} = \max_{(i,j)\in E} w_{ij}. \tag{16}$$

These measures allow a direct comparison of demand structure before and after calibration. Changes in topology indicate changes in the organization of spatial interactions, whereas changes in strength and edge-weight distributions indicate changes in flow magnitude and concentration.

In addition to these standard metrics, the hierarchical organization of flows is characterized using concentration measures, including the share of total flow captured by the top fraction of nodes, and by comparisons with null models in which flows are randomly redistributed across the network.

2.5.1 Null model construction

We construct null models in which flows are randomly redistributed across the network while preserving selected structural properties. Two types of null models are considered. In the first, edge weights are randomly permuted across existing origin–destination connections, preserving the network topology while removing any systematic association between connectivity and flow intensity. In the second, destination nodes are randomly reassigned for each origin, preserving the set of edge weights but altering the spatial structure of origin–destination interactions.

For each network, multiple realizations of both null models are generated. Flow concentration metrics, including the share of total flow captured by the top fraction of nodes, are computed for each realization. These distributions are then compared with the corresponding metrics from the observed networks to evaluate whether the hierarchical structure exceeds what would be expected under random allocation of flows.

3 Data

3.1 Study areas

The analysis includes two source cities and two target cities. See Figure 1. Brisbane, Australia, at the Statistical Area Level 1 (SA1), and Seattle, USA, at the Census Tract (CT) level, are used as source cities for estimating transferable demand models. Sydney, Australia, at the SA1 level, and New York City, USA, at the CT level, are used as target cities for transfer and validation. The Brisbane-trained demand model is transferred only to Sydney. It is not transferred to New York City because the SA1 spatial aggregation is substantially finer than the CT level and because New York City is much larger than Brisbane, making direct transfer likely to under-represent demand. The total number of SA1s in Sydney is 786, and the total number of CTs in New York City is 2,322.

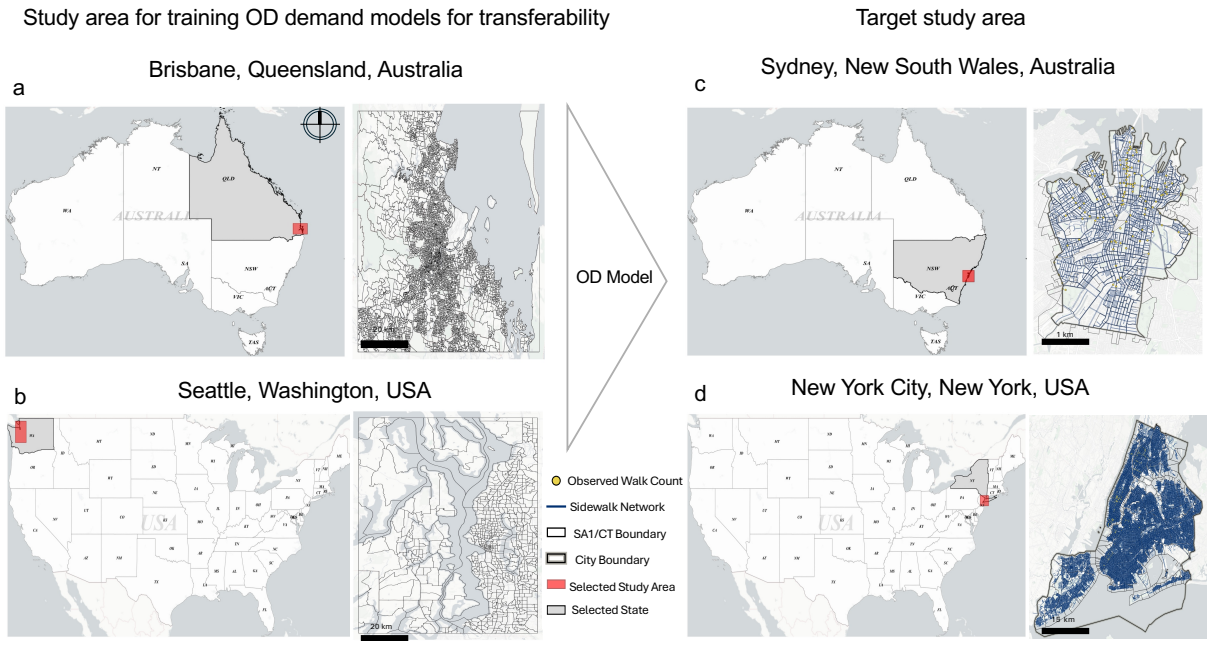

Figure 2: Study areas and data sources used in the analysis. Brisbane and Seattle serve as source cities for estimating transferable demand models, while Sydney and New York City are target cities for transfer and validation. The figure shows the spatial extent of each study area, the underlying pedestrian network, and the locations of observed pedestrian counts used for calibration. Differences in spatial scale and network structure across cities motivate the evaluation of transferability and the analysis of hierarchical flow organization.

3.2 Demand input data

Demand input data are used to estimate trip generation and trip distribution in the source cities and to apply the transferred models in the target cities.

Household travel survey data. Household travel survey data from Brisbane (Queensland Household Travel Survey, 2019–2022) and Seattle (2017–2023) are used to estimate the transferable demand models. Household travel survey data are not available for Sydney at the required spatial resolution. As a result, no local trip generation model is estimated for Sydney, and all demand inference relies on models transferred from the source cities. In Seattle, pedestrian trip data are available at the CT level, with 417 CTs retained after preprocessing. In Brisbane, more disaggregated trip records ($n = 5{,}103$) are aggregated to the SA1 level. For both source cities, the weighted number of weekday pedestrian trips per day is extracted, including walking-only trips and the walking components of multimodal trips.

Census data. Census data from 2021 are used to obtain socio-demographic variables not available in the travel surveys. For Sydney and Brisbane, SA1-level variables include population, vehicle ownership, median household income, and average household age. For Seattle and New York City, population data are obtained at the CT level from U.S. Census sources.

Land-use data. Land-use data are aggregated to the spatial units used for demand modeling. Categories include industrial, commercial, hospital, education, residential, primary production, parkland, transportation, water, and other uses.

Network connectivity variables. Street-network characteristics are extracted from OpenStreetMap at the SA1 and CT levels. These include intersection count, number of nodes, edge density, and the ratio of nodes to links. These variables capture local connectivity and structural complexity [Tian and Ewing, 2017].

Public transport accessibility. Public transport accessibility is represented by the number of stations within each spatial unit. For Australian cities, station counts are derived from GTFS feeds. For Seattle, station counts are derived using OpenStreetMap.

3.3 Pedestrian network and route attributes

Additional data are used to support route choice and link-flow assignment in the target cities.

Pedestrian network. The sidewalk network for Sydney is obtained from Footpath AI and contains 10,312 nodes and 13,301 edges. The sidewalk network for New York City is obtained from the city-scale pedestrian flow study used in the original paper and contains 188,983 nodes and 315,577 edges.

Digital elevation model. Elevation data are used to derive slope and gradient. For Sydney, a LiDAR-derived DEM at 5 m resolution is obtained from Geoscience Australia's ELVIS system. For New York City, DEM data at 1-foot resolution are used.

Green View Index. Google Street View imagery is used to compute the Green View Index, which measures the proportion of visible greenery along streets [Li et al., 2015].

Points of interest. Points of interest are extracted from OpenStreetMap and aggregated along routes. These include commercial, recreational, and service-related destinations.

4 Additional results

4.1 Transferred trip generation

All types of pedestrian trips are modeled at the SA1 level in Sydney and the CT level in New York City. The Seattle-trained trip generation model is transferred to both Sydney and New York City, while the Brisbane-trained model is transferred to Sydney only.

In Sydney, the Seattle-based model produces 807,172 trips, and the Brisbane-based model produces 239,672 trips. In New York City, the Seattle-based model produces 1,680,944 trips. For context, aggregated survey-based estimates indicate approximately 785,750 pedestrian trips per day in the City of Sydney, while New York City records nearly 14 million pedestrian trips per day. In total, the transferred destination choice models generate 6,069 origin–destination (OD) pairs for the City of Sydney and 42,883 OD pairs for New York City. These OD structures reflect the spatial interaction patterns inferred from the source models. The estimated coefficients for the two transferable trip generation models are reported in Table 1.

The distributions of trip generation input variables in the source and target cities are used as a transferability condition, ensuring that key explanatory variables are sufficiently comparable across cities [Nourmohammadi et al., 2025b].

Table 1: Estimated coefficients for transferred trip generation models trained in Brisbane and Seattle.

Variable	Brisbane, Australia		Seattle, USA	
	β	p-value	β	p-value
Constant	4.618	0.000	6.411	0.000
Population	0.134	0.006	0.182	0.000
Number of stations	0.113	0.000	0.043	0.002
Weekly household income	-0.115	0.000	-0.102	0.000
% commercial land use	0.111	0.000	0.365	0.000
Number of vehicles	-0.158	0.001		
Household median age	-0.104	0.000		
Intersection count	0.031	0.020		
Street node average	0.089	0.000		
Links / nodes ratio	0.015	0.011		
% residential land use	0.016	0.010		
% education land use	0.070	0.005		
Pseudo R^2	0.09		0.115	
RMSE	287		1,347	

4.2 Route choice, assignment, and validation

For each origin–destination pair, up to five feasible routes are generated using the BFS-LE algorithm. Among these alternatives, the shortest path is treated as the reference route.

The estimated route choice parameters are reported in Table 2.

Table 2: C-logit route choice model estimation results in Sydney.

Variable (unit)	β	p-value
Length (per 1 km)	-0.50	0.050
Number of turns	-0.11	0.000
Number of crossings (per 100 crossings)	0.01	0.091
Maximum gradient (percentage along route)	0.16	0.000
Number of POIs (per 10 points along route)	0.01	0.231
Green View Index (average percentage)	0.29	0.121
Commonality factor	0.65	0.000

A sensitivity analysis of the regularization parameter indicated that $\delta = 0.001$ provided the best performance among the tested values 0.1, 0.01, and 0. Figure 3 presents the calibration performance across different values of δ. The sensitivity analysis indicates that the impact of the regularization term is not pronounced; however, using a small positive value yields slightly better performance than setting it to zero.

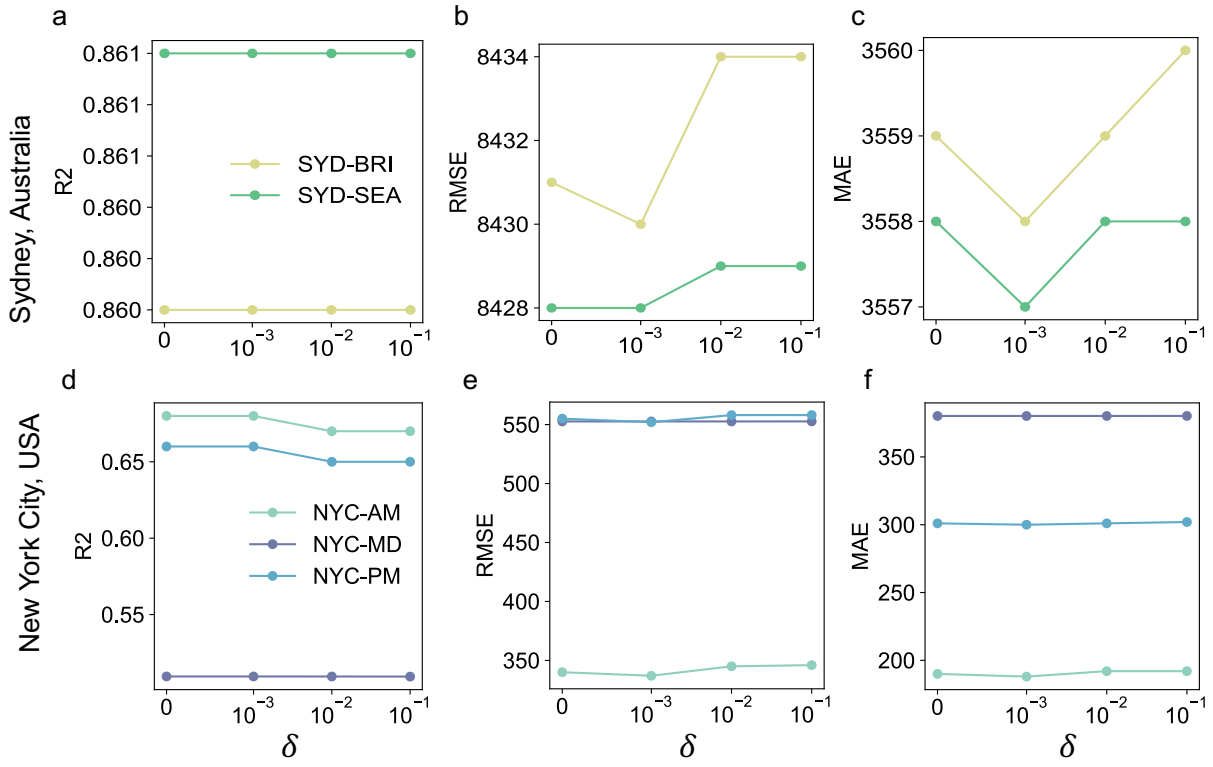

Figure 3: Sensitivity analysis of model performance across different regularisation values (δ). Panels (a)–(c) show R^2, RMSE, and MAE for Sydney using the Brisbane and Seattle models. Panels (d)–(f) present R^2, RMSE, and MAE for New York City using the Seattle model across AM, midday, and PM peak periods.

In Sydney, 115 pedestrian count locations are available. The total prior demand before calibration is 239,672 trips under the Brisbane-based trip generation model and 807,172 trips under

the Seattle-based model. In New York City, the total prior daily demand is 1,680,944 trips. Observed pedestrian count data in New York City are available for three weekday periods: the AM peak (8:00–9:00 a.m.), representing 6% of daily trips; the midday (MD) peak (12:30–1:00 p.m.), representing 9%; and the PM peak (5:00–6:00 p.m.), also representing 9%. Accordingly, daily OD demand is scaled to these proportions, and three separate OD calibration procedures are performed. This corresponds to 100,856 trips for the AM peak and 151,284 trips for both the MD and PM peaks before OD adjustment. After calibration, these values increase to 255,887 in the AM peak, 225,775 in the MD peak, and 354,138 in the PM peak.

The pedestrian network used for validation in Sydney contains 27,233 links, corresponding to 54,577 bidirectional links. In New York City, the network consists of 315,577 links, corresponding to 631,152 bidirectional links.

In Sydney, both transferred models achieve strong agreement with observed daily pedestrian counts after OD calibration. The Brisbane-based model achieves $R^2 = 0.86$, RMSE = 8,427, and MAE = 3,556, while the Seattle-based model achieves $R^2 = 0.86$, RMSE = 8,430, and MAE = 3,558. In New York City, the Seattle-trained demand model achieves $R^2 = 0.68$, RMSE = 337, and MAE = 0.52 in the AM peak; $R^2 = 0.52$, RMSE = 548, and MAE = 376 in the MD peak; and $R^2 = 0.66$, RMSE = 552, and MAE = 300 in the PM peak.

4.3 Graph-based comparison before and after OD adjustment

Tables 3 and 4 report the full set of graph metrics before and after calibration for Sydney and New York City. These values are used to assess changes in both structural and weighted properties of the inferred demand networks. These additional graph metrics are reported for completeness and to confirm that changes during calibration are primarily reflected in weighted properties rather than network topology.

Table 3: Graph feature metrics before and after demand calibration in Sydney.

	Seattle trained demand model		Brisbane trained demand model	
Graph Feature Metric	q_0	$\hat{q}$	q_0	$\hat{q}$
Nodes	1,179	1,179	1,179	1,179
Edges	11,528	11,528	11,528	11,528
Density	0.008	0.008	0.008	0.008
Average in degree	9.778	9.778	9.778	9.778
Average out degree	9.778	9.778	9.778	9.778
Average strength	123.232	1,783	36.591	1,785.9
Standard deviation strength	293.750	8,609	49.574	8,619
Average clustering undirected	0.259	0.259	0.259	0.259
Average clustering weighted	0.000	0.000	0.001	0.000
Mean betweenness	0.004	0.003	0.003	0.003
Max betweenness	0.114	0.190	0.081	0.177
Standard deviation betweenness	0.014	0.017	0.008	0.015
Mean edge weight	6.302	91.221	1.871	91.325
Standard deviation edge weight	51.027	1,024	4.549	1,029
Max edge weight	2,882	46,334	105	46,324

Table 4: Graph feature metrics before and after demand calibration in New York City.

Graph Feature Metric	AM Peak		MD Peak		PM Peak	
	q_0	$\hat{q}$	q_0	$\hat{q}$	q_0	$\hat{q}$
Nodes	4,774	4,774	4,774	4,774	4,774	4,774
Edges	40,635	40,635	40,635	40,635	40,635	40,635
Density	0.002	0.002	0.002	0.002	0.002	0.002
Average in degree	8.512	8.512	8.512	8.512	8.512	8.512
Average out degree	8.512	8.512	8.512	8.512	8.512	8.512
Average strength	42.252	170.097	63.379	124.452	63.379	230.276
Standard deviation strength	98.558	898.012	147.837	778.947	147.837	1,214
Average clustering undirected	0.234	0.234	0.234	0.234	0.234	0.234
Average clustering weighted	0.000	0.000	0.000	0.000	0.000	0.000
Mean betweenness	0.003	0.003	0.003	0.003	0.003	0.003
Max betweenness	0.126	0.202	0.126	0.184	0.126	0.187
Standard deviation betweenness	0.010	0.015	0.010	0.014	0.010	0.015
Mean edge weight	2.482	9.992	3.723	7.311	3.723	13.527
Standard deviation edge weight	13.092	97.783	19.638	76.336	19.638	119.626
Max edge weight	957	5,104	1,436	4,069	1,436	4,797

References

Xiaojiang Li et al. Assessing street-level urban greenery using google street view. *Urban Forestry & Urban Greening*, 14:675–685, 2015.

Tanapon Lilasathapornkit, Fatemeh Nourmohammadi, and Meead Saberi. Understanding walking route choice preferences and pedestrian network flows: From individual trajectories to city-scale patterns using empirical data from sydney. *Transportation Research Part A*, 208:104985, 2026.

Mark E. J. Newman. *Networks: An Introduction*. Oxford University Press, 2010.

Fatemeh Nourmohammadi, Taha H. Rashidi, and Meead Saberi. Graph-structured gravity model enhances transferable pedestrian flow prediction, 2025a. Preprint (Version 1) available at Research Square.

Fatemeh Nourmohammadi, Taha H. Rashidi, and Meead Saberi. Spatial transferability of pedestrian trip generation models. *Transportation Research Part A*, 199:104618, 2025b.

Carolina Osorio. High-dimensional origin–destination demand calibration. *Transportation Research Part B*, 124:18–43, 2019.

Guang Tian and Reid Ewing. Walk trip generation model. *Transportation Research Part D*, 52: 340–353, 2017.