From Feature Attribution to Spatial Information Structure: Interpreting Urban Pedestrian Flows with Explainable AI

Atakilti Kiros
Faculty of Engineering, Ariel University, Israel

Achituv Cohen
Faculty of Engineering, Ariel University, Israel

Yuri Ribakov
Faculty of Engineering, Ariel University, Israel

Israel Klein
Ariel University, Israel

Abstract

Machine-learning models can predict pedestrian flows with high accuracy, but they provide limited insight into how urban environments organize movement. This paper introduces an interpretive framework that links explainable artificial intelligence with Spatial Information Theory by conceptualizing feature attribution as evidence of spatial information structure – the relative salience of temporal, morphological, accessibility, and environmental dimensions in shaping pedestrian movement. Using pedestrian sensor data from Zurich and Berlin, we train comparable CatBoost models and analyze SHAP-based feature contributions. We examine whether different urban contexts encode movement in systematically distinct ways. The results reveal clear contrasts: Berlin is dominated by temporal dynamics and situational conditions, whereas Zurich shows stronger influence from urban morphology and accessibility. We interpret these differences as model-based evidence that cities organize movement-related information differently and demonstrate that a compact SHAP-based comparative design provides a practical methodological bridge between GeoAI and the study of spatial representation.

2012 ACM Subject Classification Computing methodologies → Machine learning; Information systems → Geographic information systems; Applied computing → Environmental sciences; Information systems → Spatial-temporal systems; Human-centered computing → Visualization techniques

Keywords and phrases GeoAI, explainable artificial intelligence, SHAP, pedestrian mobility, spatial information theory, urban analytics, feature attribution, spatial data science

Digital Object Identifier 10.4230/LIPIcs.COSIT.2026.30

Category Short Paper

1 Introduction

Understanding how urban space structures human movement remains a central challenge in geography, urban science, and Spatial Information Theory (SIT). While machine-learning models can accurately predict pedestrian flows using spatial, temporal, and environmental data, they offer limited insight into how movement is organized within urban environments. Recent advances in GeoAI demonstrate strong predictive capabilities for modeling complex urban dynamics, particularly through the integration of spatial, temporal, and multimodal data using machine learning and deep learning approaches [5, 7, 2]. Explainable artificial intelligence (XAI) provides a pathway to address this limitation by enabling interpretation of model behavior in geospatial contexts. Recent work has also explored explainable spatio-temporal models in urban contexts, highlighting the importance of interpretability for

17th International Conference on Spatial Information Theory (COSIT 2026).
Editors: Sabine Timpf, Gabriele Filomena, Armand Kapaj, Rui Zhu, Nicholas Giudice, and Ed Manley; Article No. 30; pp. 30:1–30:8
Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

understanding complex urban systems [1]. SHapley Additive exPlanations (SHAP) enables the quantification of feature contributions to model predictions [8], offering a means to examine how different dimensions of urban space relate to movement. However, the potential of XAI to contribute to SIT remains underexplored. From a SIT perspective, urban space can be understood as an information system that encodes relationships such as accessibility, morphology, and environmental context, which guide human behavior [4, 9, 6]. Despite this conceptual foundation, empirical methods for approximating such representations from behavioral data remain limited.

This paper addresses this gap by introducing the concept of spatial information structure, defined as the relative configuration of temporal, morphological, accessibility, and environmental dimensions that organize pedestrian movement. Rather than treating variables as independent predictors, we interpret them as components of a relational system through which urban space constrains and affords activity. We operate this concept using SHAP-based analysis of pedestrian flow data from two European cities, Zurich and Berlin. By comparing feature contributions across cities under a consistent modeling framework, we examine whether different urban contexts exhibit distinct configurations of spatial information.

2 Conceptual Framework: Spatial Information Structure

SIT conceptualizes urban environments as information-rich systems in which multiple dimensions such as built form, activity patterns, environmental conditions, and temporal rhythms jointly shape movement [3, 9, 6]. Building on this perspective, we use the term spatial information structure to refer to the relational organization of temporal, morphological, accessibility, and environmental dimensions within urban space. In this view, these dimensions are not treated as isolated predictors, but as interconnected components of a spatial system that may constrain and afford movement. Observed pedestrian flows are therefore considered not as the structure itself, but as a behavioral expression through which aspects of this underlying spatial information structure may be empirically approximated.

In the empirical implementation, features are grouped into a small set of conceptually distinct dimensions: temporal dynamics, urban morphology, accessibility, and environmental conditions. Temporal features capture cyclical patterns of activity; morphological features represent structural constraints of the built environment; accessibility features encode proximity and connectivity to services and transport; and environmental variables describe climatic and situational influences. These dimensions provide an operational approximation of how different types of urban information contribute to movement, consistent with data-driven approaches in urban analytics and GeoAI [3, 10]. Because the spatial information structure is not directly observable, it must be inferred indirectly from the relationship between environmental descriptors and observed behavior.

To support this inference, we use SHAP values to quantify the contribution of individual features to the predictions of the model [8]. Importantly, SHAP values are not interpreted as direct measurements of spatial structure but as model-based evidence that can be aggregated to approximate the relative salience of different information dimensions. This approach aligns with recent work in explainable GeoAI, which emphasizes interpreting model behavior to better understand spatial processes [1], rather than focusing solely on predictive performance in data-driven urban models [11].

This shifts the focus from predictive accuracy to identifying which dimensions of spatial information organize movement. Spatial information structures may vary across cities, requiring consistent models and features for comparison. It is defined as the relative importance of temporal, morphological, accessibility, and environmental contributions.

3 Study Design and Methods

The analysis is based on pedestrian sensor data from Zurich and Berlin collected throughout the full year of 2024. The dataset comprises approximately 60 sensors in Zurich and 33 sensors in Berlin, each providing hourly pedestrian counts across diverse urban contexts. Sensor data were obtained from municipal open data platforms. Environmental variables were derived from meteorological sources such as Meteostat, while urban morphology, accessibility, and network features were computed using OpenStreetMap (OSM) and remote sensing products such as Sentinel-2.

To ensure cross-city comparability, a harmonized set of 21 features was constructed using identical preprocessing pipelines. These features capture four conceptual dimensions: temporal dynamics, urban morphology, accessibility, and environmental conditions. Spatial variables were computed using standardized buffer-based and network-based methods, including density estimation, proximity measures, and centrality metrics. Temporal dynamics are represented as two cyclic features, Hour (cos) $= \cos(2\pi \cdot \text{hour}/24)$ and Hour (sin) $= \sin(2\pi \cdot \text{hour}/24)$, which preserve cyclic proximity between adjacent hours and allow the model to learn diurnal phase without a discontinuity at midnight. A detailed feature-level description, including data sources and operationalisation, is provided in Appendix A. To facilitate meaningful comparison, identical preprocessing procedures, feature definitions, and model settings were applied across both cities. This controlled design enables differences in feature contributions to be interpreted as differences in the organization of movement-relevant urban information rather than artifacts of model specification.

Pedestrian flows were modelled using CatBoost, with separate models trained for each city under identical configurations. We intentionally excluded autoregressive lag features to reduce simple temporal persistence effects that could dominate feature attribution and obscure relationships between movement and static urban context; this choice prioritises the interpretability of spatial and contextual predictors over maximal short-term forecasting skill. The modelling workflow consists of four stages: data integration and preprocessing; multimodal feature engineering; model training and validation with standardised splits; and post hoc interpretation using SHAP values.

Robustness was assessed across five random seeds and five independent 80/20 train–validation splits, with feature importance rankings aggregated to assess their stability across model runs.

4 Results

4.1 Model Performance

The predictive models achieve strong performance in both study areas, indicating that the selected feature set captures meaningful and spatially relevant patterns in pedestrian mobility.

As shown in Figure 1, the observed-vs.-predicted scatterplots demonstrate that the models explain a substantial proportion of the variance in pedestrian flows, with $R^2 = 0.76$ in Zurich and $R^2 = 0.80$ in Berlin. In both cities, the points cluster closely around the 1:1 reference line, showing that the models track the overall magnitude and distribution of pedestrian counts well, while still allowing for some dispersion at higher flow values. This pattern suggests that the combination of temporal, environmental, and urban form variables provides a robust basis for modeling pedestrian dynamics across distinct urban contexts.

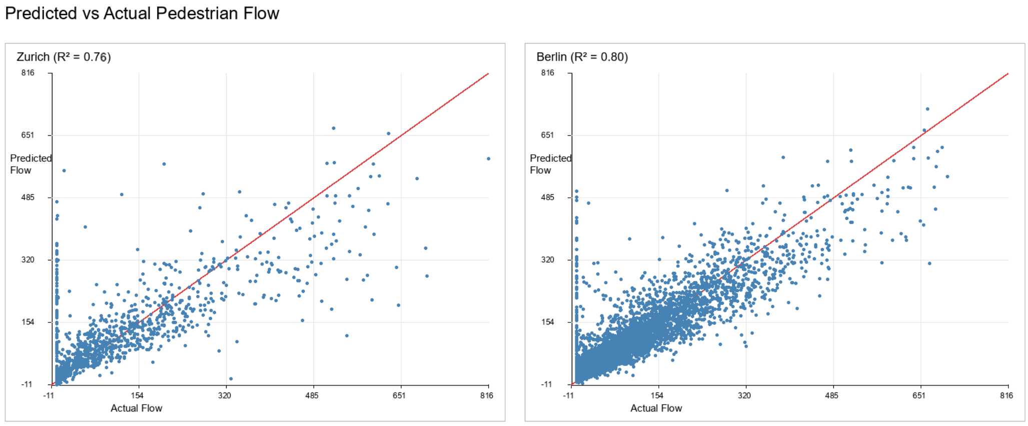

Figure 1 Predicted versus observed pedestrian flows in Zurich and Berlin, showing strong model performance ($R^2 = 0.76$ and 0.80).

4.2 Global Feature Contributions

Across both cities, temporal variables dominate the model, with cyclical hour-of-day features exhibiting the highest contributions. In both Zurich and Berlin, hour (cos) and hour (sin) are the most influential predictors, indicating that pedestrian flows are primarily structured by consistent diurnal rhythms. Beyond these shared temporal effects, non-temporal contributions differ systematically across the two urban contexts.

Table 1 Top five feature contributions based on mean absolute SHAP values and normalized importance.

Zurich			*Berlin*		
Feature	SHAP	Normalized	Feature	SHAP	Normalized
Hour (cos)	11.36	1.00	Hour (cos)	39.77	1.00
Facade continuity	9.14	0.81	Hour (sin)	27.13	0.68
Hour (sin)	8.67	0.76	Temperature	25.57	0.64
Built volume density	8.06	0.71	Street enclosure	23.68	0.60
Transit stop distance	6.65	0.59	Weekend	15.81	0.40

As shown in Table 1, Zurich is characterized by stronger morphological and accessibility-related contributions, particularly facade continuity (0.81), built volume density (0.71), and transit stop distance (0.59). In contrast, Berlin exhibits greater influence from temperature (0.64) and street enclosure (0.60), while accessibility plays a comparatively smaller role.

These patterns indicate that, although temporal dynamics provide a common structural basis, the relative salience of non-temporal features varies across cities. From a Spatial Information Theory perspective, this suggests that pedestrian movement reflects city-specific spatial information structures, with Zurich more strongly shaped by built form and accessibility, and Berlin by environmental and enclosure conditions.

4.3 Local Contribution Patterns

The SHAP beeswarm plots figure 2 provide a more detailed view of how influential features contribute to pedestrian flow predictions in each city. As shown in the global importance analysis, the temporal variables have the strongest impact in both Zurich and Berlin.

This pattern is particularly clear in Berlin, where higher values of Hour (cos) are generally associated with negative SHAP values and lower values with positive SHAP values, reflecting the cyclical nature of daily pedestrian rhythms.

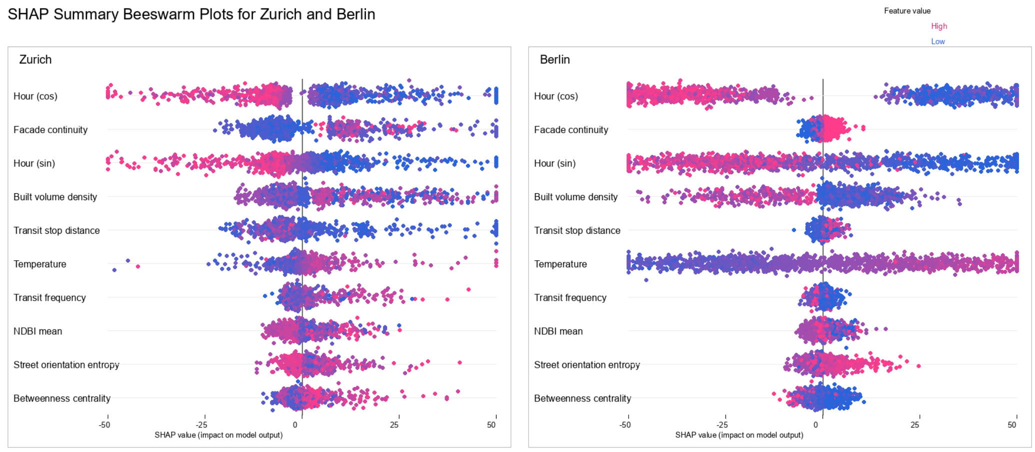

Figure 2 SHAP beeswarm plots for Zurich and Berlin illustrating the distribution and direction of feature effects across observations.

In Zurich, non-temporal variables such as facade continuity, built volume density, temperature, and transit stop distance also show substantial contribution ranges, indicating that they play an important role in the model. However, their effects are less clearly directional, suggesting more context-dependent relationships. In Berlin, temperature shows a much clearer pattern, with lower temperatures contributing negatively and higher temperatures contributing positively to the model output. Built volume density also appears to have a clearer directional effect, while facade continuity and transit stop distance play a comparatively smaller role than in Zurich.

These results suggest that although temporal rhythms dominate in both cities, the structure of non-temporal influences differ meaningfully between them.

4.4 Comparative Structures

Although the feature-level analysis in Section 4.2 highlights differences in the importance of individual predictors, the normalized contributions shown in Figure 3 reveal higher-level structural patterns. Temporal variables dominate in both cities, yet the relative importance of non-temporal features differs systematically.

Zurich exhibits a more distributed configuration of influence across morphological and accessibility dimensions, whereas Berlin shows a more concentrated reliance on temporal and environmental signals. This suggests that pedestrian movement in Zurich is shaped by a more heterogeneous spatial information structure, while in Berlin it is governed by a narrower set of dominant factors. From a SIT perspective, these patterns indicate that urban environments may encode behaviorally relevant information in fundamentally different ways across cities.

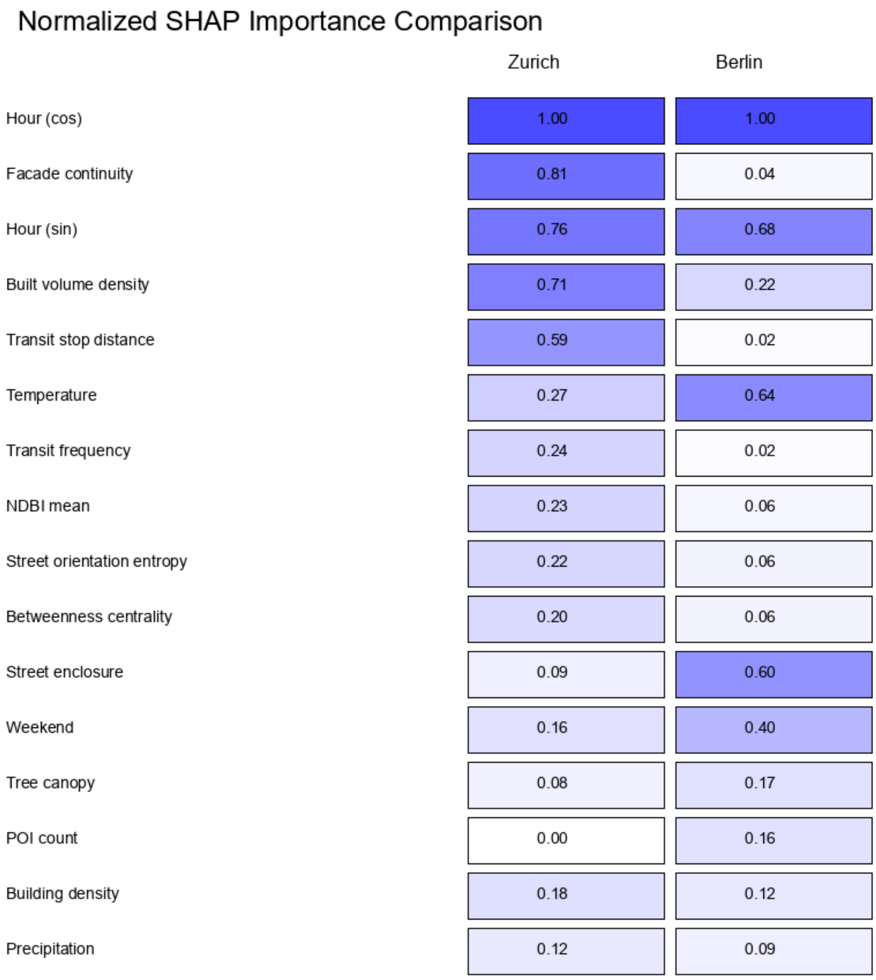

Figure 3 Normalized SHAP importance profiles for Zurich and Berlin.

4.5 Robustness of Feature Contribution Patterns

Across both cities, temporal variables dominate the model, with cyclical hour-of-day features exhibiting the highest contributions. In both Zurich and Berlin, hour (cos) and hour (sin) are the most influential predictors, indicating that pedestrian flows are primarily structured by consistent diurnal rhythms. Beyond these shared temporal effects, non-temporal contributions differ systematically across the two urban contexts.

These patterns indicate that, while temporal dynamics provide a common baseline, the importance of non-temporal features varies across cities. From a SIT perspective, this suggests that pedestrian movement reflects city-specific spatial information structures, with Zurich shaped more by built form and accessibility, and Berlin by environmental and enclosure conditions.

5 Discussion and Conclusion

The results indicate that while temporal dynamics provide a common structural basis for pedestrian flows in both cities, the relative importance of non-temporal features varies in ways that reflect distinct urban configurations. Zurich exhibits a more distributed contribution structure, with stronger influence from morphological and accessibility-related variables such

Table 2 Robustness of normalized SHAP feature contributions across repeated model runs.

	Zurich					*Berlin*			
Feature	Mean	SD	Rank	Range	Feature	Mean	SD	Rank	Range
Hour (cos)	1.00	0.00	1.0	1–1	Hour (cos)	1.00	0.00	1.0	1–1
Facade continuity	0.81	0.04	2.2	2–3	Hour (sin)	0.68	0.03	2.1	2–3
Hour (sin)	0.76	0.03	2.8	2–4	Temperature	0.64	0.05	3.0	2–4
Built volume density	0.71	0.05	4.0	3–5	Street enclosure	0.60	0.04	4.0	3–5
Transit stop distance	0.59	0.06	5.1	4–6	Weekend	0.40	0.05	5.2	4–6

as facade continuity, built volume density, and transit stop distance. In contrast, Berlin shows a more concentrated pattern in which environmental and enclosure-related factors, particularly temperature and street enclosure, are more prominent. These differences suggest that, beyond shared temporal rhythms, cities encode movement-relevant information through distinct configurations of urban form, accessibility, and environmental context.

These findings should not be interpreted as evidence of direct causal relationships, but rather as model-based approximations derived from SHAP. From a Spatial Information Theory perspective, their significance lies in revealing how different dimensions of urban space become salient within the model's learned representation of movement. Rather than identifying independent determinants of pedestrian activity, the analysis highlights how multiple dimensions jointly organize behaviorally meaningful information. In this sense, the observed contribution patterns can be understood as an empirical approximation of a city's spatial information structure. This interpretation remains subject to limitations. The inferred structures are derived from observational data and may reflect confounding factors, feature dependencies, or sensor distribution effects. The results should therefore be regarded as exploratory rather than definitive. Nevertheless, they provide a theoretically grounded and empirically interpretable basis for comparing how different urban environments organize movement-related information.

In summary, this paper introduced spatial information structure as a framework linking explainable GeoAI with Spatial Information Theory. The comparative analysis of Zurich and Berlin demonstrates that, while temporal dynamics form a shared foundation, non-temporal contributions vary systematically across cities, reflecting distinct informational configurations. By interpreting feature contributions as components of a relational system, the study moves beyond prediction toward a more interpretable understanding of how urban environments encode behaviorally relevant information. The comparison should also be interpreted in light of substantial differences in city size, population, and urban structure between Zurich and Berlin. Rather than comparing equivalent urban systems, the analysis examines whether a common modeling framework can reveal distinct configurations of movement-relevant information across contrasting urban contexts. A key direction for future work is to analyze the spatio-temporal structure of prediction errors to better understand the limits of the inferred spatial information structure. Identifying where and when the model fails may reveal missing contextual or event-driven factors and support the integration of context-aware variables, improving interpretability and generalization.

References

1. A. Alamelh. Explain-city: Explainable spatio-temporal deep learning for urban resilience and risk governance. *IEEE Access*, 13:181989–182004, 2025. doi:10.1109/ACCESS.2025.3623066.
2. S. Avram. Advancing urban analytics: Geoai applications in spatial decision-making and sustainable cities. *Urban Science*, 10, 2026. doi:10.3390/urbansci10030148.

3 Michael Batty. *The New Science of Cities*. MIT Press, 2013.

4 Max J. Egenhofer and David M. Mark. Naive geography. In Andrew U. Frank and Werner Kuhn, editors, *Spatial Information Theory: A Theoretical Basis for GIS*, pages 1–15. Springer, Berlin, 1995. doi:10.1007/3-540-60392-1_1.

5 Krzysztof Janowicz, Song Gao, Grant McKenzie, Yingjie Hu, and Budhendra Bhaduri. Geoai: Spatially explicit artificial intelligence techniques for geographic knowledge discovery and beyond. *International Journal of Geographical Information Science*, 34:625–636, 2020. doi:10.1080/13658816.2019.1684500.

6 Werner Kuhn. Core concepts of spatial information for transdisciplinary research. *International Journal of Geographical Information Science*, 26:2267–2276, 2012. doi:10.1080/13658816.2012.722637.

7 Peng Liu, Yu Zhang, and Filip Biljecki. Explainable spatially explicit geospatial artificial intelligence in urban analytics. *Environment and Planning B: Urban Analytics and City Science*, 51:1104–1123, 2024. doi:10.1177/23998083231204689.

8 Scott M. Lundberg and Su-In Lee. A unified approach to interpreting model predictions. In *Advances in Neural Information Processing Systems*, volume 30, 2017.

9 Daniel R. Montello. Cognitive map-design research in the twentieth century: Theoretical and empirical approaches. *Cartography and Geographic Information Science*, 31(4):283–304, 2004.

10 Fivos Papadimitriou. *Spatial Artificial Intelligence*. SpringerBriefs in Applied Sciences and Technology. Springer, Cham, 2025. doi:10.1007/978-3-031-82136-3.

11 Liudmila Prokhorenkova, Gleb Gusev, Aleksandr Vorobev, Anna Veronika Dorogush, and Andrey Gulin. Catboost: Unbiased boosting with categorical features. In *Advances in Neural Information Processing Systems*, volume 31, 2018.

A Appendix A: Feature Definition and Data Sources

Table 3 Feature taxonomy, data sources, and operationalisation of the 21 predictors.

Category	**Feature**	**Data Source**	**Operationalization**
Temporal Dynamics			
Temporal	hour_cos	Sensor timestamp	Cosine encoding of hour-of-day
Temporal	hour_sin	Sensor timestamp	Sine encoding of hour-of-day
Temporal	is_weekend	Calendar data	Binary weekend indicator
Environmental / Weather			
Environment	temp	Meteostat / NOAA	Hourly air temperature
Environment	precip	Meteostat / NOAA	Hourly precipitation intensity
Urban Morphology (Built Form)			
Built form	built_volume_density_200m	OpenStreetMap	Built volume within 200 m buffer
Built form	building_density_200m	OpenStreetMap	Building footprint density (200 m)
Built form	facade_continuity_proxy	OpenStreetMap	Proxy of street wall continuity
Built form	street_enclosure_index_proxy	OpenStreetMap	Enclosure ratio (height/width)
Built form	street_canyon_ratio_mean_200m	OpenStreetMap	Mean height-to-width ratio (200 m)
Urban Green / Environmental Context			
Green context	sensor_canopy_pct	Remote sensing	Tree canopy percentage
Green context	sensor_ndvi_mean	Sentinel-2	Mean NDVI (buffer)
Green context	ndbi_mean_s2	Sentinel-2	Mean built-up index
Accessibility / Transport			
Accessibility	transit_stop_distance	OpenStreetMap	Distance to nearest transit stop
Accessibility	transit_frequency_index	GTFS / OSM	Frequency-weighted accessibility index
Accessibility	bus_stop_density_400m	OpenStreetMap	Bus stop density (400 m)
Urban Activity / Land Use			
Land use	poi_count_300m	OpenStreetMap	POI count within 300 m
Land use	access_to_services_index	OpenStreetMap	Composite service accessibility index
Street Network Structure			
Network	intersection_density_300m	OpenStreetMap	Intersection density (300 m)
Network	betweenness_centrality	OpenStreetMap	Network centrality measure
Network	street_orientation_entropy	OpenStreetMap	Orientation diversity entropy