

Relational Verification, Feature Sensitivity, and Cross-City Transfer in Heterogeneous Graph Learning for Urban Functional Zone Classification

Original

Relational Verification, Feature Sensitivity, and Cross-City Transfer in Heterogeneous Graph Learning for Urban Functional Zone Classification / Vazirov, E., Apiletti, D.. - 3023:(2026), pp. 127-141. (International Conference on Artificial Intelligence for Digital Transformations (AIDT 2026) Baku, Azerbaijan 22/06/2026 - 24/06/2026) [10.1007/978-3-032-31319-5_9].

Availability:

This version is available at: 11583/3015702 since: 2026-09-18T18:39:25Z

Publisher:

Springer

Published

DOI:10.1007/978-3-032-31319-5_9

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)



Relational Verification, Feature Sensitivity, and Cross-City Transfer in Heterogeneous Graph Learning for Urban Functional Zone Classification

Etibar Vazirov^{1,2}(✉)  and Daniele Apiletti¹

¹ Politecnico di Torino, Turin, Italy

{etibar.vazirov,daniele.apiletti}@polito.it

² ADA University, Baku, Azerbaijan

Abstract. Urban functional zone classification is a key component of data-driven urban analytics, supporting applications in planning, mobility, and digital governance. Although point-of-interest (POI) composition is widely used to infer regional functionality, it remains unclear whether predictive performance reflects meaningful relational structure or primarily direct semantic signals derived from the same data source. This study introduces a heterogeneous graph learning framework that models Region–POI, POI–Category, and Region–Region interactions for functional zone classification. Comparable graph representations are constructed for three cities with diverse urban characteristics—Baku, Berlin, and Montevideo—using a uniform grid partition and OpenStreetMap-derived POI data. To enable rigorous evaluation, the analysis incorporates structured relation ablation, cross-city transfer, and a feature sensitivity setting where direct semantic region features are removed.

Experimental results show that conventional models and homogeneous graph baselines experience substantial performance degradation under feature removal, whereas the heterogeneous model maintains strong performance in both in-city and cross-city settings. Relation ablation further indicates that Region–POI interactions contribute most significantly to predictive performance, while spatial adjacency provides comparatively limited explanatory power.

These findings suggest that heterogeneous relational modeling captures meaningful structural dependencies in urban data and provides a more reliable representation of functional organization under varying feature and transfer conditions.

Keywords: Heterogeneous Graph Neural Networks · Urban Functional Zone Classification · Relational Modeling · Cross-City Transfer · Feature Sensitivity

1 Introduction

Urban functional zone classification aims to identify semantic region types such as residential, commercial, industrial, and mixed-use areas from heterogeneous

urban data. Accurate functional labeling supports data-driven planning, infrastructure optimization, and AI-enabled digital transformation. As urban systems increasingly rely on automated decision-support pipelines, understanding how semantic structure emerges from spatial and relational dependencies becomes a central research challenge. Existing approaches range from rule-based schemes to POI-driven machine learning and multimodal geospatial frameworks [9, 13, 19]. POI composition is widely adopted due to its availability and interpretability; however, most methods treat regions independently or rely on aggregated statistics, thereby neglecting structured relational dependencies inherent in urban systems. In addition, the reliance on POI-derived features for both labeling and prediction raises concerns about feature-label circularity, potentially overstating model performance. Graph-based modeling provides a principled alternative for capturing spatial interactions. Graph neural networks (GNNs) have shown strong performance in urban forecasting and infrastructure modeling [7, 10], while heterogeneous GNNs (HeteroGNNs) explicitly model multiple node and relation types, enabling richer representations of interactions among regions, POIs, categories, and spatial adjacencies [1]. Graph-based functional zoning and multimodal urban representation learning further demonstrate the benefits of relational modeling [3, 14]. Despite these advances, three challenges remain. First, the contribution of individual relation types is rarely quantified in a systematic, counterfactual manner; improvements are typically attributed to graph structure without isolating which dependencies drive performance gains. Second, cross-city transferability remains underexplored. Although transfer learning has been studied in traffic and spatiotemporal forecasting [8, 11, 18], domain shift effects in region-level functional classification remain insufficiently examined. Third, it remains unclear whether the predictive power of existing models reflects genuine relational learning or is dominated by direct semantic features derived from the same POI source. To address these limitations, a heterogeneous graph learning framework is introduced for region-level functional zone classification, modeling three relational components: region-POI containment, POI-category mapping, and region-region spatial adjacency. Using a uniform 750-m grid, comparable heterogeneous urban graphs are constructed for three cities with diverse urban characteristics, namely Baku, Berlin, and Montevideo, based on OpenStreetMap data accessed via OSMnx and Overpass. Beyond predictive evaluation, the analysis incorporates structured relation ablation to quantify the contribution of each relation type, cross-city transfer experiments to assess transferability under domain shift, and a feature sensitivity setting in which direct semantic region features are removed to evaluate reliance on POI-derived attributes.

The contributions of this work are threefold:

1. A unified heterogeneous graph framework for region-level functional zone classification across multiple urban environments;
2. A structured relation ablation strategy for quantifying the contribution of heterogeneous relations;

3. A combined evaluation protocol including cross-city transfer and feature sensitivity analysis for assessing robustness and mitigating feature–label circularity.

2 Related Work

2.1 Urban Functional Zone Modeling

Urban functional zone identification is a central task in smart city analytics and digital urban planning. Early approaches rely on land-use records, remote sensing data, and statistical aggregation of point-of-interest (POI) distributions, while recent surveys highlight a transition toward data-driven and learning-based frameworks [9].

Recent studies increasingly incorporate multimodal geospatial data and graph-based representations to capture spatial dependencies beyond aggregated POI statistics. Multilevel graph neural networks have been applied to functional zone classification using heterogeneous spatial features [19], and multimodal graph fusion approaches integrate POIs, road networks, and auxiliary signals to improve predictive performance [14]. Dynamic and multi-temporal graph models further attempt to capture evolving urban functions [3]. Despite these advances, most approaches focus primarily on predictive accuracy, with limited analysis of how individual relational components contribute to performance.

2.2 Graph Modeling and Heterogeneous Urban Representations

Graph learning has become a key paradigm for modeling complex urban systems, including infrastructure networks and semantic spatial interactions [6, 15]. Heterogeneous graph neural networks (HeteroGNNs) extend classical GNNs by modeling multiple node and relation types, making them particularly suitable for urban environments characterized by structural and semantic heterogeneity [1].

Recent developments include spatial graph-based foundation models for large-scale urban representation learning [21] and domain-enhanced multimodal graph models for planning applications [16]. While these methods demonstrate strong representational capacity, they generally do not provide systematic mechanisms for interpreting or verifying the role of different relations within the graph.

2.3 Cross-City Transfer and Domain Shift

Transferability across cities remains a major challenge in urban AI systems. Prior work has extensively studied cross-city transfer in traffic and spatiotemporal prediction [8, 11, 20], where alignment-based and knowledge graph approaches are used to mitigate structural and distributional differences [18].

However, most existing studies focus on temporal forecasting rather than region-level semantic classification. Systematic cross-city evaluation for functional zone classification remains limited. Recent surveys on graph learning under distribution shift emphasize the importance of robustness and transferability capabilities in graph-based models [17], but empirical evidence at the level of regional functional semantics is still scarce.

2.4 Research Gap

Existing literature demonstrates the effectiveness of graph-based models for urban representation and highlights the importance of transfer learning under domain shift. However, two key limitations remain: (i) the lack of systematic verification of relational dependencies in heterogeneous urban graphs, and (ii) limited empirical evaluation of cross-city transferability for region-level functional classification.

This work addresses these limitations by combining structured relation ablation, cross-city transfer evaluation, and feature sensitivity analysis within a unified heterogeneous graph framework, enabling both relational verification and robust behavior under transfer assessment.

3 Methodology

3.1 Problem Formulation and Graph Construction

Let $\mathcal{C} = \{c_1, \dots, c_M\}$ denote the set of cities under study. For each city $c \in \mathcal{C}$, the urban area is partitioned into a uniform grid of 750-m spatial regions $\mathcal{R}_c = \{r_1, \dots, r_N\}$. Each region is associated with a functional label $y_i \in \mathcal{Y}$.

$$\mathcal{Y} = \{\text{Residential, Commercial, Industrial, Leisure, Mixed}\}. \quad (1)$$

The task is to learn a function

$$f_\theta : \mathcal{R}_c \rightarrow \mathcal{Y}, \quad (2)$$

that predicts region-level functional zones from multi-relational urban structure.

For each city, a heterogeneous graph is constructed as

$$\mathcal{G}_c = (\mathcal{V}_c, \mathcal{E}_c), \quad (3)$$

where

$$\mathcal{V}_c = \mathcal{R}_c \cup \mathcal{P}_c \cup \mathcal{K}_c$$

contains region nodes, POI nodes, and semantic category nodes. The edge set consists of three relation types: region–POI containment, POI–category assignment, and region–region spatial adjacency. POIs are obtained from OpenStreetMap data through OSMnx and Overpass, mapped to grid cells using geometric containment, and grouped into semantic super-categories. Region adjacency is defined by shared grid boundaries in projected coordinate space.

This graph design separates spatial proximity from semantic composition, allowing relation-aware learning instead of collapsing all interactions into a homogeneous graph. Heterogeneous graph neural networks are well suited to this setting because they model multiple node and edge types through relation-specific transformations [1]. The resulting schema is shown in Fig. 1.




Fig. 1. Schema of the heterogeneous urban graph. Region nodes connect to POIs through containment relations, POIs link to semantic categories, and neighboring regions are connected through spatial adjacency.

3.2 Functional Label Definition and Feature Settings

City-scale authoritative functional labels are often incomplete, unavailable, or inconsistent across urban contexts. To ensure reproducibility and cross-city comparability, functional labels are derived from POI composition within each spatial region.

Let $\mathcal{P}(r_i)$ denote the set of POIs contained in region r_i , and let $\mathcal{P}_k(r_i)$ denote the subset belonging to semantic super-category k . The normalized composition vector of region r_i is defined as

$$\mathbf{c}_i = \left[\frac{|\mathcal{P}_1(r_i)|}{|\mathcal{P}(r_i)|}, \dots, \frac{|\mathcal{P}_K(r_i)|}{|\mathcal{P}(r_i)|} \right], \quad (4)$$

where K is the number of super-categories.

Functional labels are assigned using a dominance-based rule:

$$y_i = \arg \max_{k \in \mathcal{Y}} \mathbf{c}_{i,k}, \quad (5)$$

subject to a fixed dominance threshold applied consistently across all cities. Regions without a sufficiently dominant category are assigned to the *Mixed* class.

This strategy follows the common use of POI distributions as proxies for regional semantics [9, 13], while enabling a controlled and consistent evaluation across cities despite differences in label distribution.

To assess the impact of direct semantic features, two region-feature settings are considered. The *FULL* setting includes direct semantic composition features such as category counts and aggregate diversity indicators. The *NODIRECT* setting removes direct category-count features from region nodes and retains only indirect structural summaries, including total POI count, diversity, active category count, maximum category ratio, second-largest category ratio, and dominance gap.

The heterogeneous graph structure, including POI and category nodes and their relations, is kept unchanged. This setup allows isolating whether predictive performance is driven by explicit semantic attributes or by relational structure encoded in the graph.

3.3 Heterogeneous Message Passing Model

Given $\mathcal{G}_c$, node representations are updated using relation-specific message passing. Let $\mathcal{R}$ denote the set of relation types. At layer ℓ , the embedding of node v is updated as

$$\mathbf{h}_v^{(\ell+1)} = \sigma \left(\sum_{r \in \mathcal{R}} \sum_{u \in \mathcal{N}_r(v)} \frac{1}{|\mathcal{N}_r(v)|} \mathbf{W}_r^{(\ell)} \mathbf{h}_u^{(\ell)} \right), \quad (6)$$

where $\mathcal{N}_r(v)$ is the relation-specific neighborhood of v , $\mathbf{W}_r^{(\ell)}$ is a learnable transformation for relation r , and $\sigma(\cdot)$ is a nonlinear activation function. Mean aggregation is used to reduce sensitivity to node degree.

The model uses stacked heterogeneous convolution layers with relation-wise transformations, ReLU activations, and dropout regularization. All node types are projected into a shared latent space, while only region embeddings are used for classification. After the final layer, the predicted class distribution for region r_i is computed as

$$\hat{\mathbf{y}}_i = \text{softmax} \left(\mathbf{W}_o \mathbf{h}_{r_i}^{(L)} \right), \quad (7)$$

where $\mathbf{W}_o$ is the output projection matrix. The parameters are optimized using cross-entropy loss over labeled training regions:

$$\mathcal{L} = - \sum_{r_i \in \mathcal{R}_c^{\text{train}}} \sum_{k=1}^{|\mathcal{Y}|} y_{ik} \log \hat{y}_{ik}. \quad (8)$$

3.4 Evaluation Protocol

The evaluation consists of three complementary components. First, in-city classification is performed independently for each city using stratified splits, and

results are reported as mean and standard deviation over five random seeds. Second, structured relation ablation is used to quantify the importance of each relation type. For a relation r , a modified graph $\mathcal{G}_c^{-r}$ is created by removing all edges of that relation while keeping all other graph components fixed. The relation contribution is measured as

$$\Delta_r = \text{MacroF1}(\mathcal{G}_c) - \text{MacroF1}(\mathcal{G}_c^{-r}). \quad (9)$$

Third, cross-city transfer is used to assess transferability under domain shift. Given a source city c_s and target city c_t , the model is trained on $\mathcal{G}_{c_s}$ and evaluated on $\mathcal{G}_{c_t}$ without fine-tuning:

$$\theta^* = \arg \min_{\theta} \mathcal{L}_{c_s}(\theta), \quad \text{evaluate on } \mathcal{G}_{c_t}. \quad (10)$$

This setting is important because graph learning models may be sensitive to distribution shifts across domains [17]. Both FULL and NODIRECT settings are evaluated to distinguish ordinary predictive performance from robustness to direct semantic feature removal.

4 Experimental Setup

4.1 Datasets and Graph Construction

Experiments are conducted on three structurally diverse urban environments: Baku (Azerbaijan), Berlin (Germany), and Montevideo (Uruguay). These cities differ in terms of geographic location, scale, and urban organization, providing a controlled yet diverse setting for evaluating transferability under heterogeneous urban distributions. For each city, the urban space is partitioned into a uniform grid of 750-m regions, and inactive cells without POIs are removed.

The resulting heterogeneous graphs contain:

- **Baku:** 1097 regions and 12,566 POIs,
- **Berlin:** 1323 regions and 66,485 POIs,
- **Montevideo:** 579 regions and 9,222 POIs.

Each graph includes three relation types: region–POI containment, POI–category semantic mapping, and region–region spatial adjacency.

To explicitly control potential circularity between POI-based labels and features, two feature configurations are evaluated:

- **FULL:** includes category-level POI composition features,
- **NODIRECT:** excludes direct category features and retains only aggregate statistics (e.g., diversity, dominance, and POI counts).

This design enables systematic evaluation of whether models rely on direct feature-label correspondence or exploit relational structure.

4.2 Baselines and Training Protocol

The proposed heterogeneous GNN is compared against both non-graph and graph-based baselines to provide a controlled evaluation across different modeling paradigms:

- **LGBM**: a feature-based model using region-level statistics,
- **GCN**: a homogeneous graph model using only region adjacency,
- **GIN**: a structure-aware graph baseline with enhanced neighborhood aggregation.

These baselines are selected to isolate the contribution of relational modeling by contrasting feature-only learning, homogeneous graph learning, and more expressive graph-based aggregation. For the heterogeneous model, hyperparameters are selected via grid search based on validation Macro-F1. The search space is summarized in Table 1.

Table 1. Hyperparameter search space for the heterogeneous GNN.

Hyperparameter	Search Values
Hidden dimension (d)	{32, 64}
Dropout rate	{0.3, 0.5}
Learning rate	{0.005, 0.01}
Number of layers (L)	2
Optimizer	Adam

The best configuration is selected on the validation split and re-evaluated across multiple random seeds to ensure stability. The search space is intentionally kept compact to ensure fair comparison and avoid excessive hyperparameter tuning.

4.3 Evaluation Strategy

Performance is evaluated using Accuracy and Macro-F1, with Macro-F1 serving as the primary metric due to class imbalance.

$$\text{MacroF1} = \frac{1}{K} \sum_{k=1}^K \text{F1}_k, \quad (11)$$

where F1_k denotes the F1-score for class k .

To ensure robustness, all experiments are repeated across multiple random seeds and reported as mean $\pm$ standard deviation. Beyond standard predictive evaluation, four complementary analysis protocols are employed:

- **Feature decoupling (FULL vs NODIRECT)** to assess dependence on direct POI-derived features,
- **Structured relation ablation** to quantify the contribution of each edge type,
- **Cross-city transfer** to evaluate transferability under domain shift,
- **Deletion-based fidelity** to measure the influence of local POI structures on model predictions.

For cross-city transfer, models are trained on all labeled regions of a source city and directly evaluated on a target city without fine-tuning, providing a strict assessment of domain transferability.

5 Results and Analysis

This section presents a comprehensive evaluation of the proposed heterogeneous graph framework, focusing on relational robustness, structural verification, and transferability under domain shift. All experiments are conducted over five independent random seeds, and results are reported as mean and standard deviation to ensure statistical reliability.

5.1 Robustness to Feature Removal: FULL vs NODIRECT

To assess the dependence of the model on direct region-level features, a controlled experiment is conducted by comparing two configurations: (i) FULL, where all region features are used, and (ii) NODIRECT, where direct region features are removed and only relational information is retained.




Fig. 2. Performance comparison between FULL and NODIRECT configurations across three cities.

As shown in Fig. 2, the proposed HeteroGNN maintains strong performance even in the absence of direct region features. In Baku and Berlin, the performance degradation remains limited, indicating that relational structure alone provides sufficient information for accurate classification.

In contrast, homogeneous baselines such as GCN and GIN exhibit substantial performance drops under the NODIRECT setting, suggesting that they rely heavily on explicit node features and are less capable of exploiting multi-relational dependencies.

These findings provide empirical evidence that the proposed model does not depend solely on aggregated region-level statistics, but instead leverages structured relational signals, mitigating concerns regarding feature-label circularity.

5.2 Structured Relation Ablation

To quantify the contribution of each relation type, structured relation ablation is performed as defined in Eq. (9). Each relation type is removed independently, and the resulting degradation in macro-F1 is measured.




Fig. 3. Relation ablation results across three cities.

Figure 3 shows that removing Region–POI edges consistently leads to the largest performance drop across all cities. This confirms that semantic composition, captured through POI aggregation, is the dominant factor in functional zone classification. The removal of Region–Region adjacency results in comparatively smaller degradation, indicating that spatial proximity provides complementary contextual information rather than acting as a primary signal. The POI–Category relation contributes moderately, improving representation stability but not driving predictions independently. Importantly, the observed patterns remain consistent across all five random seeds and across different urban environments, supporting the robustness of the relational interpretation.

5.3 Cross-City Transfer Under Domain Shift

To evaluate transferability, models are trained on one city and evaluated on another without fine-tuning. The results are presented in Fig. 4 for FULL and NODIRECT settings.




Fig. 4. Cross-city transfer performance under domain shift.

A noticeable performance reduction is observed compared to in-city results, indicating the presence of domain shift across urban environments. However, the proposed HeteroGNN consistently outperforms the LGBM baseline in transfer scenarios, suggesting that relational representations capture partially transferable structural patterns.

The transfer behavior is asymmetric, with certain city pairs yielding stronger results than others, reflecting differences in urban composition and data distribution.

Homogeneous graph baselines (GCN and GIN) are not included in this evaluation due to their significantly lower in-city performance and unstable behavior under domain shift observed in preliminary experiments. As such, they do not provide meaningful reference points for transfer analysis.

It is important to note that these results should be interpreted as evidence of transferability rather than definitive generalization. Given the limited number of cities, the findings indicate robustness under domain shift, while the proposed framework remains extensible to additional urban environments for broader validation.

5.4 Deletion-Based Fidelity Analysis

To further examine model behavior, deletion-based fidelity analysis is conducted by removing the top- K most influential POIs for each region based on gradient-based importance scores.

As shown in Fig. 5, prediction confidence decreases as increasingly important POIs are removed, indicating that the model relies on semantically meaningful




Fig. 5. Deletion-based fidelity across cities.

relational signals. The effect is more pronounced in larger and denser cities such as Berlin, while more moderate variations are observed in Montevideo.

These results provide additional evidence that the model captures meaningful dependencies rather than relying on spurious correlations or noise in the data.

6 Discussion

The results highlight several important aspects of heterogeneous relational modeling for urban functional zone classification.

First, semantic containment between regions and POIs emerges as the dominant signal. Structured ablation consistently shows that removing Region–POI relations leads to the largest performance degradation, while spatial adjacency contributes comparatively less. This indicates that functional zoning is primarily driven by semantic composition, with spatial proximity providing complementary context rather than a primary signal.

Second, cross-city experiments reveal clear distributional differences between urban environments. Although performance decreases under transfer, heterogeneous representations remain more stable than simpler baselines when partial structural alignment exists. These findings suggest that the proposed framework captures transferable relational patterns to a limited extent, while still reflecting city-specific characteristics.

Third, deletion-based fidelity analysis confirms that predictions are driven by semantically meaningful signals. The consistent confidence drop observed after removing influential POIs indicates that the model relies on structured semantic information rather than spurious correlations. Moreover, the stability of this behavior across cities supports the robustness of the learned representations.

From a practical perspective, the framework enables data-driven functional zoning using publicly available POI data, making it applicable in settings where official annotations are incomplete or unavailable. At the same time, the results

should be interpreted with appropriate caution, as the evaluation is limited to a small number of cities and relies on POI-based proxy labels.

Finally, the use of publicly available POI data introduces considerations related to data quality, potential biases, and responsible use. While no personal or sensitive individual-level information is involved, POI datasets may reflect uneven spatial coverage or socio-economic biases inherent in volunteered geographic information. These factors can influence model behavior and should be considered when interpreting results or deploying such systems in real-world applications.

7 Conclusion and Future Work

This paper presented a heterogeneous graph neural network framework for region-level urban functional zone classification, explicitly modeling Region–POI, POI–Category, and Region–Region relations within a unified structure. By preserving relational dependencies during message passing, the approach captures semantic patterns that are not accessible through purely spatial or tabular representations.

Experimental results across multiple cities show that heterogeneous modeling consistently outperforms homogeneous graph baselines and remains competitive with strong tabular methods. Structured relation ablation indicates that Region–POI containment is the most influential component, while spatial adjacency provides complementary contextual information. Cross-city transfer experiments reveal measurable performance degradation under distribution shift, suggesting that learned representations are partially transferable but still influenced by city-specific characteristics.

These findings highlight the importance of explicitly modeling semantic composition in urban functional analysis, while also emphasizing the challenges associated with transferring models across heterogeneous urban environments. The reported results should be interpreted as evidence of transferability rather than full generalization, given the limited number of cities considered.

Future work will investigate domain adaptation and cross-city alignment strategies to improve robustness under distribution shift. Extending the framework to a broader set of cities and incorporating additional data sources, such as temporal and multimodal urban signals, may further enhance scalability and applicability.

References

1. Bing, R., Yuan, G., Zhu, M., Meng, F., Ma, H., Qiao, S.: Heterogeneous graph neural networks analysis: a survey of techniques, evaluations and applications. *Artif. Intell. Rev.* **56**(8), 8003–8042 (2023)
2. Chen, Y., Yang, J., Yang, R., Xiao, X., Xia, J.C.: Contribution of urban functional zones to the spatial distribution of urban thermal environment. *Build. Environ.* **216**, 109000 (2022)

3. Ding, C., Tang, J., Zhang, M., Wu, H., Wang, B.: Identifying dynamic functions of urban areas with a multi-temporal collaborative graph neural network model. *Int. J. Geogr. Inf. Sci.* **40**(1), 78–111 (2026)
4. Fan, Y., et al.: Towards predicting urban land use changes: a dynamic graph alignment perspective. *ACM Trans. Intell. Syst. Technol.* **16**(2), 1–24 (2025)
5. Geng, X.: Multi-source cross-city transfer learning for region-based spatial distribution prediction on ride-hailing demand. Hong Kong University of Science and Technology (Hong Kong) (2024)
6. Javaherian Pour, E., Atazadeh, B., Rajabifard, A., Sabri, S.: Developing a CityGML-based graph data model for utility infrastructure in smart cities. *ISPRS Ann. Photogrammetry Remote Sens. Spat. Inf. Sci.*, 405–412 (2025)
7. Jin, G., et al.: Spatio-temporal graph neural networks for predictive learning in urban computing: a survey. *IEEE Trans. Knowl. Data Eng.* **36**(10), 5388–5408 (2023)
8. Jin, Y., Chen, K., Yang, Q.: Selective cross-city transfer learning for traffic prediction via source city region re-weighting. In: *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pp. 731–741 (2022)
9. Liu, B., Deng, Y., Li, M., Yang, J., Liu, T.: Classification schemes and identification methods for urban functional zone: a review of recent papers. *Appl. Sci.* **11**(21), 9968 (2021)
10. Liu, T., Meidani, H.: End-to-end heterogeneous graph neural networks for traffic assignment. *Transp. Res. Part C: Emerg. Technol.* **165**, 104695 (2024)
11. Lone, S.N., Abouelela, M., Yu, T., Kim, J., Antoniou, C.: Transfer learning in spatiotemporal graph neural networks for traffic forecasting: taxonomy, analysis, and future directions. *Anal. Future Dir.* (2025)
12. Lu, K., Zhao, X., Li, M., Zhou, Y., Zhang, B.: An urban change detection method based on multimodal data and knowledge graph technology. *Int. J. Digit. Earth* **18**(2), 2564902 (2025)
13. Luo, G., Ye, J., Wang, J., Wei, Y.: Urban functional zone classification based on poi data and machine learning. *Sustainability* **15**(5), 4631 (2023)
14. Tao, Y., et al.: A graph-based multimodal data fusion framework for identifying urban functional zone. *Int. J. Appl. Earth Obs. Geoinf.* **136**, 104353 (2025)
15. Trudelle, C., Claramunt, C.: A graph-based modelling approach for the representation and analysis of urban conflicts. *Comput. Environ. Urban Syst.* **114**, 102201 (2024)
16. Wang, C., Wang, Y., Zhao, W.: Domain-graph enhanced multi-task multimodal large models for urban planning. *Indoor Built Environ.* **34**(9), 1636–1652 (2025)
17. Wu, M., et al.: Graph learning under distribution shifts: a comprehensive survey on domain adaptation, out-of-distribution, and continual learning (2024). arXiv preprint [arXiv:2402.16374](https://arxiv.org/abs/2402.16374)
18. Zhang, S., et al.: Deep transfer learning for city-scale cellular traffic generation through urban knowledge graph. In: *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pp. 4842–4851 (2023)
19. Zhang, T., et al.: Classifying urban functional zones via a multilevel graph neural network and multimodal geospatial data. *IEEE J. Sel. Top. Appl. Earth Obs. Remote Sens.* (2025)

20. Zhang, Y., Wang, X., Yu, X., Sun, Z., Wang, K., Wang, Y.: Drawing informative gradients from sources: a one-stage transfer learning framework for cross-city spatiotemporal forecasting. In: Proceedings of the AAAI Conference on Artificial Intelligence, vol. 39, pp. 1147–1155 (2025)
21. Zhang, Z., Xie, M., Balsebre, P., Huang, W., Luo, S., Cong, G.: UrbanMFM: spatial graph-based multiscale foundation models for learning generalized urban representation. *IEEE Trans. Knowl. Data Eng.* (2026)

Open Access This chapter is licensed under the terms of the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License (<http://creativecommons.org/licenses/by-nc-nd/4.0/>), which permits any noncommercial use, sharing, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if you modified the licensed material. You do not have permission under this license to share adapted material derived from this chapter or parts of it.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

