Adversarial Self-Supervised Learning for Secure and Robust Urban Region Profiling

Weiliang Chen, Qianqian Ren, Yong Liu, Jianguo Sun, and Feng Lin, *Senior Member, IEEE*

***Abstract*—Urban region profiling is essential for forecasting and decision-making in dynamic and noisy urban environments. However, existing approaches struggle with adversarial attacks, data incompleteness, and security vulnerabilities, which undermine predictive accuracy and reliability. This paper introduces Enhanced Urban Region Profiling with Adversarial Self-Supervised Learning (EUPAS), a robust framework that integrates adversarial contrastive learning with self-supervised and supervised objectives. To fortify resilience against adversarial attacks and noisy data, we introduce perturbation augmentation, a trickster generator, and a deviation copy generator, which collectively enhance the robustness of learned embeddings. EUPAS significantly outperforms state-of-the-art models in forecasting tasks, including crime prediction, check-in prediction, and land usage classification, achieving up to 12.2% improvement in forecasting performance. Additionally, our model demonstrates superior resilience against transfer-based black-box and white-box attacks compared to baseline models. By addressing key security challenges in data-driven urban modeling, EUPAS provides a scalable and adversarially robust solution for smart city applications.**

***Index Terms*—Urban region profiling, adversarial contrastive learning, robust forecasting, adversarial attacks, smart city security.**

I. Introduction

URBAN region profiling plays a fundamental role in urban management, enabling critical tasks such as traffic forecasting, socio-demographic analysis, and crime prediction. The rapid expansion of urban sensor networks and the increasing availability of multi-source urban data (e.g., mobility records, Points of Interest (POIs), and socio-economic indicators) have significantly enhanced the potential for data-driven decision-making [1], [2], [3]. However, the high-dimensional, heterogeneous, and noisy nature of urban data, coupled with the increasing vulnerability of smart city systems to adversarial attacks, poses significant challenges for robust and secure urban forecasting.

Existing urban region embedding models have been widely applied in prediction tasks such as crime forecasting [4], socio-demographic analysis [5], and land usage classification [6]. Recently, deep learning-based approaches have shown promising results in urban region profiling. MV-PN [7] introduced a region embedding model that captures intra- and inter-regional similarities using POI networks and spatial autocorrelation layers. CGAL [8] extended this idea by incorporating collective adversarial training, while MVGRE [9] leveraged multi-view joint learning to enhance region representations. MGFN [10] focused on traffic pattern extraction but overlooked POI data, which is crucial for capturing regional functionalities. ROMER [11] improved urban region embeddings by modeling long-range dependencies through a global attention graph network, while HREP [12] introduced prefix-tuning to enhance adaptability in downstream tasks. Despite their advancements, these models struggle with suboptimal embeddings caused by inevitable noise and data incompleteness in urban region profiling. To address these issues, contrastive learning has emerged as a promising solution by learning robust representations without extensive labeled data [13], [14], [15].

However, contrastive learning in the context of urban region modeling still faces critical challenges. One major obstacle lies in the generation of semantically meaningful augmented samples for contrastive learning. Traditional strategies such as random cropping or feature masking are often unsuitable for urban settings, where they can disrupt key spatial relationships. For instance, as illustrated in Fig. 1, replacing a POI (e.g., a café) with a nearby location (e.g., a restaurant or shopping mall) can fundamentally alter the semantics of a mobility pattern, thereby misleading the model and degrading representation quality. These distortions hinder the effectiveness of contrastive training and reduce generalization performance in downstream tasks.

Furthermore, existing studies overlook a critical aspect in urban analytics: **security**. Graph Neural Networks (GNNs), which form the core of most region representation models, are inherently vulnerable to adversarial perturbations. Small but imperceptible changes in node features or graph structure often guided by gradient information can drastically change prediction outcomes [16]. This fragility poses a serious threat in high-stakes urban applications such as crime prediction or traffic forecasting, where model failures can lead to dangerous real-world consequences.

Given these limitations, there is a pressing need for more **robust and secure learning frameworks**. Adversarial contrastive learning has emerged as a promising paradigm in domains like computer vision [17], [18], [19] and natural language processing [20], where it has demonstrated improved

Received 4 February 2025; revised 7 July 2025; accepted 21 July 2025. Date of publication 30 July 2025; date of current version 7 August 2025. This work was supported by China Postdoctoral Science Foundation under Grant 2022M711088. The associate editor coordinating the review of this article and approving it for publication was Dr. Yan Wang. *(Corresponding authors: Qianqian Ren; Jianguo Sun.)*

Weiliang Chen, Qianqian Ren, and Yong Liu are with the Department of Computer Science and Technology, Heilongjiang University, Harbin 150080, China (e-mail: chanweiliang@s.hlju.edu.cn; renqianqian@hlju.edu.cn; liuyong123456@hlju.edu.cn).

Jianguo Sun is with Hangzhou Institute of Technology, Xidian University, Hangzhou 311231, China (e-mail: jgsun@xidian.edu.cn).

Feng Lin is with the Department of Computer Science and Technology, Zhejiang University, Hangzhou 310027, China (e-mail: flin@zju.edu.cn).

Digital Object Identifier 10.1109/TIFS.2025.3594165

(a) (b)

Fig. 1. Difficulties in data augmentation for urban data: (a) Implacability of geographic neighbor sets; (b) Replacing a POI (e.g., a café) with a nearby location (e.g., a restaurant or shopping center) fundamentally changes the semantics and purpose of the trip, distorting the model's learning process.

resistance to distributional shifts and malicious attacks. By integrating adversarial perturbations into the contrastive learning process, models can learn more resilient features that withstand both noise and targeted manipulation.

However, applying adversarial contrastive learning to urban region profiling presents new challenges. Urban data is inherently heterogeneous and multi-relational, consisting of mobility flows, POI distributions, and geographic structures. Ensuring semantic consistency across these modalities while preserving relational integrity under adversarial settings remains an open problem. Most existing frameworks are not designed to address such complex structural and semantic interactions, leading to poor robustness in noisy or adversarial environments.

To address these challenges, we propose EUPAS (Enhanced Urban Region Profiling with Adversarial Self-Supervised Learning), a novel framework that explicitly enhances security, robustness, and adaptive learning in urban profiling tasks. Our key innovations include:

(1) **Perturbation-Aware Data Augmentation.** To mitigate adversarial vulnerabilities, EUPAS introduces a perturbation augmentation module, which injects targeted adversarial noise during training to enhance model resilience against security threats.

(2) **Deviation Copy Generator for Semantic-Preserving Augmentation**. To address the limitations of conventional data augmentation, EUPAS employs a Deviation Copy Generator that constructs semantically meaningful positive samples, preventing erroneous augmentations that alter urban region semantics and avoiding distortions that could compromise model learning.

(3) **Trickster Generator for Hard Negative Sample Construction.** To improve the robustness of urban embeddings, EUPAS introduces the Trickster Generator, which not only automatically synthesizes more difficult negative samples, enhancing the model's ability to distinguish urban regions even under distribution shifts, but also ensures structural stability by preserving the relationships between nodes.

The contributions of this paper are summarized as follows.

- **General Aspect.** We propose **EUPAS**, the first security-aware adversarial self-supervised learning framework for urban region profiling. EUPAS addresses key challenges in adversarial robustness, data augmentation, and multi-source data fusion, enhancing the reliability of urban forecasting systems.
- **Methodologies.** We introduce a perturbation-aware augmentation strategy comprising two novel modules: the **Deviation Copy Generator**, which produces semantically consistent positive samples, and the **Trickster Generator**, which crafts hard negative samples. This design improves the model's robustness to adversarial attacks while preserving the semantic and structural integrity of urban data.
- **Experimental Evaluation.** Extensive experiments on two real-world urban datasets demonstrate that EUPAS achieves superior performance across tasks such as check-in prediction, land usage classification, and crime forecasting. It achieves up to a 12.2% improvement in prediction accuracy. In terms of adversarial robustness, EUPAS outperforms state-of-the-art models under both white-box and black-box attack settings, exhibiting stronger resilience and better transferability.

II. Related Work

A. Graph Neural Networks for Region Representation

GNNs have become foundational in urban graph modeling. RGCN [21] and HetGNN [22] extend Graph Convolutional Networks (GCNs) to heterogeneous settings, well-suited for urban data fusion. GraphSAGE [23] and GAT [24] further enable inductive learning and attention-based message passing. However, most GNNs presume reliable graph structures and are vulnerable to adversarial attacks involving subtle modifications to edges or features [25]. Defense approaches like AD-GCL [26] and RGCN [27] attempt to address this but are primarily designed for homogeneous graphs. In contrast, our approach considers adversarial resilience in noisy, heterogeneous urban graphs by designing two generative modules to produce hard positive and negative samples.

B. Region Representation Learning

With the increasing availability and scale of urban data sources, research on learning urban region representations has expanded, primarily into single and multiple data sources. Regional human mobility data is commonly used as a single data source to mine inter-regional correlations. For example, HDGE [4] constructs flow and spatial graphs to learn region embeddings through human mobility. ZE-Mob [6] extracts mobility patterns from taxi trajectories and learns region embeddings using source-destination co-occurrences. MGFN

[10] analyzes mobility data from different periods to build mobility patterns for region embedding. Methods integrating multiple urban data sources offer more comprehensive regional attributes, leading to richer region embeddings. MV-PN [7] constructs graph representations of inter-region human mobility data and intra-region POIs, using them as initial vectors in an AutoEncoder to learn final region embeddings. CGAL [8] enhances embedding learning by connecting multiple graphs through a network-based strategy. MVGRE [9] improves performance by building region-based graphs from various data sources and employing a multi-view fusion mechanism. Region2Vec [28] leverages knowledge graphs to explore global and local correlations in multi-source data, boosting region representation learning. ROMER [11] excels by capturing multi-view dependencies from diverse data sources, using global graph attention networks and a dual-stage fusion module. Recently, HREP [12] introduced prefix prompt learning from NLP to optimize and guide downstream tasks automatically.

Despite these advances, many of these methods assume clean data and lack resilience to adversarial perturbations. Though adversarial augmentation [29] and robust contrastive learning [30] have shown promise, their direct application in urban analytics remains limited. Our work addresses this gap by explicitly modeling structured adversarial signals in region-level contrastive learning.

C. Adversarial Contrastive Learning

Contrastive learning has shown strong potential in unsupervised graph representation [31], [32], but remains susceptible to adversarial perturbations [26], [33]. Adversarial training (AT) [34] has emerged as a reliable defense mechanism, training models directly on perturbed data to enhance robustness [35], [36].

Several prior works have explored the integration of AT into contrastive frameworks. RoCL [37] was among the first to introduce adversarial instance-level perturbations in self-supervised contrastive learning, crafting hard positives to strengthen feature discrimination without requiring labels. GASSL [16] extended these ideas to the graph domain, proposing automated adversarial views to avoid manual augmentation heuristics. In the vision domain, Chen et al. [17] studied the effect of pretraining under adversarial contrastive objectives and highlighted the impact on downstream robustness. Similarly, Fan et al. [19] examined robustness transferability from pretraining to finetuning and provided theoretical insights into when adversarial invariance is preserved.

Beyond vision and graphs, adversarial contrastive learning has also been applied to language models. ARE [20] demonstrated improved generalization and robustness bounds in pre-trained LLMs under adversarial noise. These developments collectively underscore the importance of contrastive supervision under adversarial conditions, motivating our tailored adversarial contrastive framework for multi-relation urban graphs. In contrast to prior work such as RoCL [37], GASSL [16], which focus on instance-level adversarial views in unimodal settings, our approach introduces relationally structured perturbations over multi-relational urban graphs, enabling semantically-aware positive and negative sample generation for spatial tasks.

III. Preliminaries

In this section, we introduce key notations and formally define the urban region embedding problem. Consider a city divided into N non-overlapping regions. We construct a graph $\mathcal{G} = (\mathcal{V}, E)$, where $\mathcal{V}$ is the set of region nodes, each representing a unique urban area, and E is the set of edges capturing inter-regional relationships. We assume the existence of K distinct relation types ($K > 1$) between regions. The edge set associated with the k-th relation is denoted by E_k, where $1 \leq k \leq K$. For each relation, we define a corresponding subgraph $\mathcal{G}_k = (\mathcal{V}, E_k)$. These relational subgraphs are constructed based on human mobility patterns, POIs similarity, and geographic adjacency.

Definition 1 (Human Mobility): Human mobility is defined as a collection of trips between urban regions. Let $t = (r_o, r_t)$ represent a travel record, where r_o and r_t denote the origin and destination regions respectively, with $1 \leq o, t \leq N$. The set of all trip records is denoted as $\mathcal{T} = \{t_1, t_2, \ldots, t_{|M|}\}$, where M is the total number of trips.

Definition 2 (POI Information): The functional characteristics of each region are represented by its POIs. Let f be the number of POI categories. We denote the POI matrix as $\mathcal{P} = [p_1, p_2, \ldots, p_N] \in \mathbb{R}^{f \times N}$, where each column vector $p_i \in \mathbb{R}^f$ represents the POI distribution of region r_i.

Definition 3 (Geographic Neighbors): Geographic neighbor information encodes the spatial relationships between a region and its adjacent regions. We denote the geographic neighbors of each region as $\mathcal{N} = \{\mathcal{N}_1, \mathcal{N}_2, \ldots, \mathcal{N}_N\}$, where $\mathcal{N}_i = \{r_j \mid r_j \text{is adjacent to } r_i\}$ represents the set of geographic neighbors for region r_i.

Problem Definition: Given the graph $\mathcal{G}$, human mobility data $\mathcal{T}$, POI information $\mathcal{P}$, and geographic neighbors $\mathcal{N}$, our objective is to learn low-dimensional embeddings $\mathcal{E} = \{e_1, e_2, \ldots, e_N\}$, where each $e_i \in \mathbb{R}^d$ denotes the d-dimensional embedding of region $r_i \in \mathcal{V}$. These embeddings should effectively preserve critical information derived from mobility patterns, geographic relationships, and POI features. Such representations are expected to benefit various downstream urban computing tasks, including check-in prediction, land usage classification, and crime prediction.

IV. Methodology

The overall architecture of EUPAS is illustrated in Fig. 2. EUPAS is designed to improve the robustness, generalization, and security of urban region embedding.

Workflow Overview. EUPAS follows a structured pipeline to generate robust region embeddings. It begins with **Region Representation Learning**, which uses a multi-relational graph neural network to extract spatial-temporal features from urban data (e.g., POIs, mobility, geography). The **Perturbation Augmentation** module then introduces controlled adversarial noise to improve robustness. Next, the **Attentive Supervised Module** assigns dynamic attention to different data sources, emphasizing informative features and suppressing

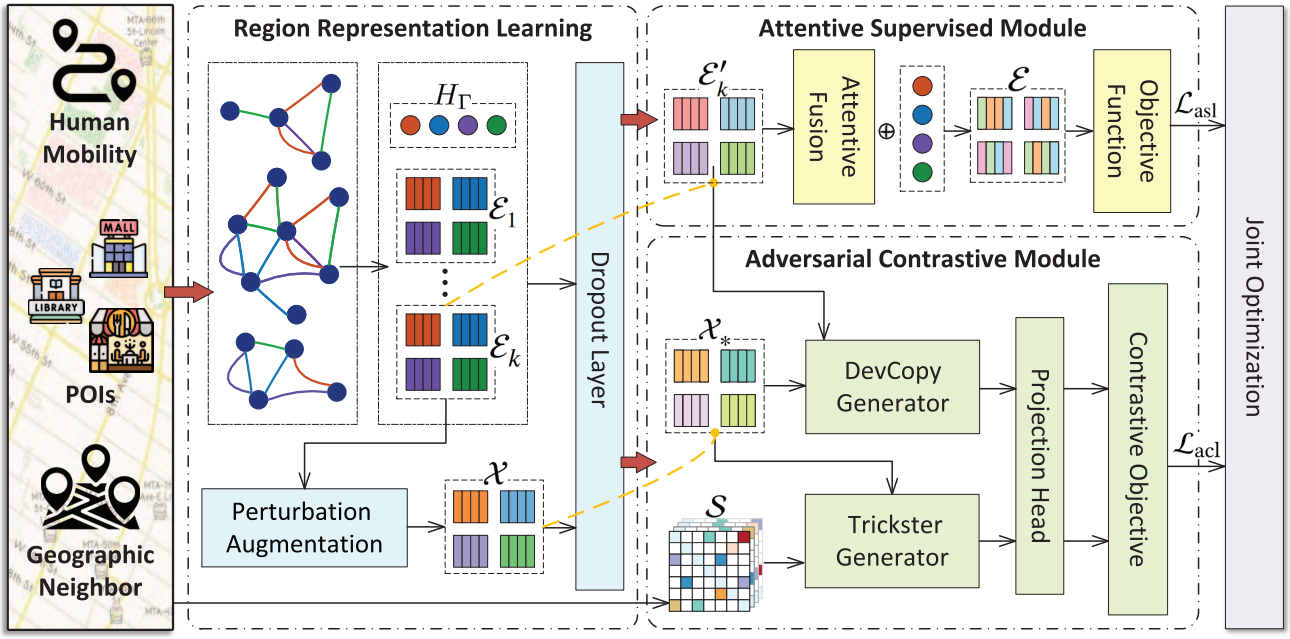

Fig. 2. The architecture of EUPAS integrates four key components: region representation learning, perturbation augmentation, attentive supervised module, and adversarial contrastive learning.

TABLE I
LIST OF IMPORTANT NOTATIONS

N	the number of regions
K	the number of relation types
$\mathcal{S}$	similarity matrix for different data sources
H_Γ	relation embedding representing inter-region dependencies
$\mathcal{E}_k$	relation-focused region embedding for region k
$\mathcal{E}'_k$	anchor embedding for region k
$\mathcal{E}$	final task-specific region embedding
$\mathcal{X}_*$	perturbation-augmented node embedding
$\mathcal{X}_*^-$	hard negative samples (generated by Trickster)
$\mathcal{X}_*^+$	strong positive samples (generated by DevCopy)
$\mathcal{L}_{\text{acl}}$	adversarial contrastive loss function
$\mathcal{L}_{\text{asl}}$	attentive supervised loss function
$\mathcal{L}_{\text{total}}$	total loss function
τ	temperature coefficient for contrastive loss scaling
α	balancing factor for positive and negative samples in contrastive loss
β	regularization factor for total loss
λ	adversarial perturbation coefficient

noise. Finally, **Adversarial Contrastive Learning** enhances representation discrimination via two strategies: the **Deviation Copy Generator (DevCopy)** produces semantically aligned but challenging positives, while the **Trickster Generator** generates structurally similar yet semantically distinct negatives. The key notations used in EUPAS are summarized in Table I.

A. Region Representation Learning

GCNs offer a natural message-passing mechanism that stabilizes representation learning over irregular and heterogeneous graph topologies. This property is especially useful in urban settings, where each region engages in diverse and sparse relational patterns (e.g., human mobility, geographic proximity, and POI similarity).

We first learn the embedding representations of regions under different relations. Let $\mathcal{E}_k^{(l)} = \left\{e_{1k}^{(l)}, e_{2k}^{(l)}, \ldots, e_{Nk}^{(l)}\right\}$ denote the region embeddings under the k-th relation at layer l, and $H_\Gamma^{(l)} = \left\{h_0^{(l)}, h_1^{(l)}, \cdots, h_K^{(l)}\right\}$ represent relation embeddings. Given initial node embeddings $\mathcal{E}_k^{(0)}$ and relation embeddings $H_\Gamma^{(0)}$, their updates at the l-th layer follow:

$$e_{uk}^{(l)} = \sigma\left(\sum_{\gamma\in\Gamma}\sum_{v\in\mathcal{N}_u^\gamma}\phi_{u,\gamma}\mathbf{W}^{(l)}(e_{vk}^{(l-1)} \circ h_k^{(l-1)})\right), \tag{1}$$

$$h_k^{(l)} = \mathbf{w}_k^{(l)} h_k^{(l-1)} + \mathbf{b}_k^{(l)}, \tag{2}$$

where $u, v \in \{1, 2, \ldots, N\}$ and $k \in \{0, 1, \ldots, K\}$. $\sigma(\cdot)$ denotes the LeakyReLU activation function, and $\circ$ represents the element-wise product. $\mathcal{N}_u^k$ is the set of neighbors of region u under relation k, and $\phi_{u,k} = \frac{1}{|\mathcal{N}_u^k|}$ serves as a normalization factor. $\mathbf{W}^{(l)}$, $\mathbf{w}_k^{(l)}$, and $\mathbf{b}_k^{(l)}$ are all learnable parameters. This formulation allows the model to adaptively integrate structural and semantic cues from heterogeneous relational views.

The element-wise interaction between $e_{vk}^{(l-1)}$ and $h_k^{(l-1)}$ provides an efficient and parameter-free mechanism to inject relational semantics into node updates. While simple, it preserves localized relational expressiveness and avoids over-parameterization, making it well-suited for urban graphs with sparse multi-relational structures (as verified in Section V-E).

To mitigate feature smoothing in GNNs, we adopt ResNet-based [38] aggregation:

$$\mathcal{E}_k^{(l+1)} = \mathcal{E}_k^{(l)} + \sigma\left(D^{-\frac{1}{2}}AD^{-\frac{1}{2}}\mathcal{E}_k^{(l)}\mathbf{W}^{(l)}\right), \tag{3}$$

where A is the adjacency matrix and D is the diagonal degree matrix.

To efficiently handle large-scale urban data, we store adjacency matrices $\mathbf{A}_k$ as sparse matrices, significantly reducing memory consumption by retaining only non-zero elements. Each region maintains only the top k-nearest neighbors based

on similarity measures of human mobility, POIs, and geographic distances.

B. Perturbation Augmentation for Robustness and Security

Urban region data is often noisy and sparse, making learned representations vulnerable to distortion, especially under adversarial perturbations. To address this, we introduce a perturbation-based augmentation that injects controlled noise into region embeddings, enhancing both robustness and generalization. Specifically, the region embeddings $\mathcal{E}_k$ for the k-th relation are modified through a controlled perturbation process:

$$\mathcal{X}_* = f\left(\mathcal{E}_k + \eta \Delta\boldsymbol{\theta}_{sp}\right), \quad \Delta\boldsymbol{\theta}_{sp} \sim \mathcal{N}(0, \boldsymbol{\sigma}), \tag{4}$$

where $\Delta\boldsymbol{\theta}_{sp}$ is a Gaussian noise perturbation with zero mean and standard deviation σ. The scalar η controls the perturbation intensity, while $f(\cdot)$ represents a dropout layer, which further enhances robustness by randomly masking features to prevent overfitting. The resulting perturbed embeddings, $\mathcal{X}_*$, ensure that the model remains robust even in dynamic, uncertain urban environments.

C. Attentive Supervised Module

To integrate representations across multiple relations, we introduce an attentive supervised module that assigns weights to different relation types. This mechanism allows the model to focus on the most informative relational signals, guided by a supervised loss that improves embedding quality.

First, we define the semantic fusion coefficient α_k for the k-th relation as follows:

$$\alpha_k = \frac{1}{|N|} \sum_{j=1}^{|N|} \mathbf{q}^\top \cdot \sigma\left(\mathbf{W} e_{jk} + \mathbf{b}\right), \tag{5}$$

where $\mathbf{q}$ is the attention vector, and $\sigma(\cdot)$ represents the LeakyReLU activation function. The parameters $\mathbf{q}$, $\mathbf{W}$, and $\mathbf{b}$ are shared across all region embeddings, ensuring a consistent projection into the same space for the computation of α_k.

Then, using the learned coefficients α_k, the final region representation integrates all relations as:

$$\mathcal{E} = \sum_{k=1}^{K} \text{softmax}\left(\alpha_k\right) \cdot \mathcal{E}'_k \cdot h_k, \tag{6}$$

where h_k is the corresponding relation embedding, and $\mathcal{E}'_k$ represents the region embeddings under the k-th relation after applying a dropout layer to $\mathcal{E}_k$. This aggregation ensures that the representation is both comprehensive and focused on significant relationships.

Inspired by [12], we design a unified loss function to optimize the module, defined as:

$$\begin{aligned}\mathcal{L}_{\text{asl}} = &\sum_{i=1}^{N} \max\left\{\left\|e_i - e_i^+\right\|_2 - \left\|e_i - e_i^-\right\|_2, 0\right\} \\ &+ \sum_{(r_i, r_j) \in \mathcal{M}} \left(\log \frac{\hat{d}_o(r_j \mid r_i)}{p_o(r_j \mid r_i)} + \log \frac{\hat{d}_t(r_i \mid r_j)}{p_t(r_i \mid r_j)}\right) \\ &+ \sum_{i=1}^{N} \sum_{j=1}^{N} \left[\mathcal{S}_p^{ij} - (e_{ip})^\top e_{jp}\right]^2, \end{aligned} \tag{7}$$

where e_i^+ and e_i^- are the positive and negative geographic neighbors of region r_i, respectively. e_{ip} and e_{jp} represent the embeddings of regions r_i and r_j under the POI relation. $\mathcal{S}_p$ is the POI similarity matrix from [9]. $p_o(r_j \mid r_i)$ and $p_t(r_i \mid r_j)$ are the original mobility distributions. $\hat{d}_o(r_j \mid r_i)$ and $\hat{d}_t(r_i \mid r_j)$ are the reconstructed origin and target distributions, computed as:

$$\hat{d}_o(r_j \mid r_i) = \frac{\exp(e_i^{s\top} e_j^t)}{\sum_j \exp(e_i^{s\top} e_j^t)}, \quad \hat{d}_t(r_i \mid r_j) = \frac{\exp(e_j^{t\top} e_i^s)}{\sum_i \exp(e_i^{t\top} e_i^s)}. \tag{8}$$

This loss function balances the supervised learning objectives across multiple relationships, ensuring the embeddings capture the most relevant features while maintaining consistency with the observed data distributions.

D. Adversarial Contrastive Module

To overcome the intrinsic separability of region embeddings, which often limits contrastive learning effectiveness, we propose two key modules: the **Deviation Copy Generator (DevCopy)** and the **Trickster Generator (Trickster)**, as shown in Fig. 3. These modules generate adversarially enhanced positive and negative samples, encouraging the model to learn more discriminative and robust region representations.

1) Deviation Copy Generator: DevCopy addresses a central challenge in contrastive learning: generating positive samples that are both semantically consistent and sufficiently challenging. Common augmentation methods (e.g., masking or cropping) often distort spatial semantics in urban contexts, leading to unreliable embeddings. In contrast, DevCopy introduces harder positives while preserving semantic integrity. It achieves this through a two-stage adversarial process:

Directional Semantic Adjustment: A perturbation ψ is applied to the original embedding $\mathcal{E}'_k$ to introduce controlled deviation:

$$\check{\mathcal{X}} = \mathcal{E}'_k - \psi \frac{\boldsymbol{v}}{\|\boldsymbol{v}\|_2}, \quad \boldsymbol{v} = \nabla_{\mathcal{E}'_k} \mathcal{L}_{\text{cl}}^+. \tag{9}$$

Here, $\mathcal{L}_{\text{cl}}^+$ is the positive contrastive loss formally defined in Eq. (14). $\boldsymbol{v}$ guides the perturbation direction along the gradient of the contrastive loss, ensuring the modified sample remains informative and semantically aligned.

Semantic Regularization via KL Divergence: To ensure semantic consistency during adversarial perturbation, we constrain the perturbed embedding to yield similar semantic predictions as the original under each relation. Specifically, we define the semantic prediction distribution as:

$$\mathcal{L}_{\text{KL}} = \sum_{k=1}^{K} D_{\text{KL}}\left(P\left(\mathcal{S}_k \mid \check{\mathcal{X}}\right) \| P\left(\mathcal{S}_k \mid \mathcal{E}'_k\right)\right). \tag{10}$$

where $P(\mathcal{S}_k \mid \cdot)$ denotes the predicted semantic distribution under the k-th relation. This encourages the perturbed embedding to preserve semantic consistency with the original sample while increasing its representational difficulty in feature space.

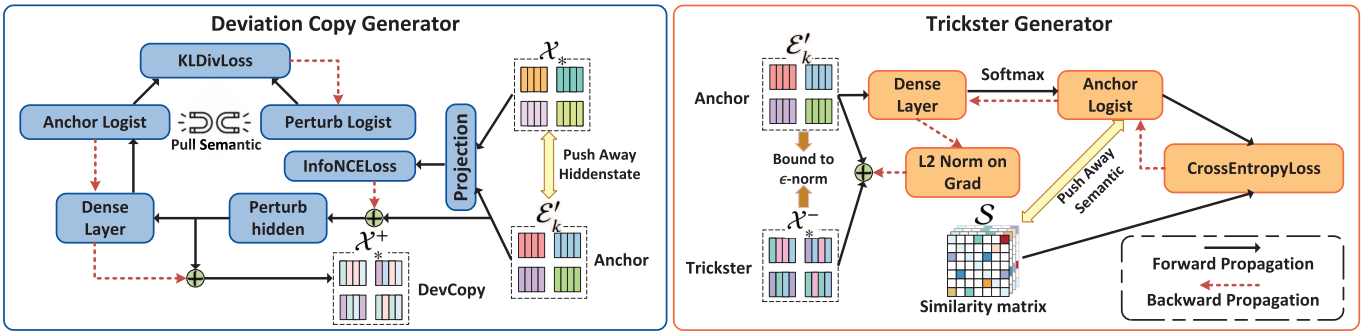

Fig. 3. Workflow diagram for the deviation copy and trickster generators.

Final Positive Sample Refinement: After semantic regularization, the refined hard positive sample $\mathcal{X}_*^+$ is obtained by applying an additional adversarial perturbation:

$$\mathcal{X}_*^+ = \check{\mathcal{X}} - \psi \frac{\omega}{\|\omega\|_2}, \quad \omega = \nabla_{\check{\mathcal{X}}} \mathcal{L}_{\mathrm{KL}}. \tag{11}$$

This two-stage optimization produces hard positives that preserve semantic meaning while increasing the model's sensitivity to fine-grained variations in urban region representations.

2) Trickster Generator: The Trickster Generator aims to create **hard negative samples** that are close in structure yet semantically divergent from the anchor. Traditional contrastive learning often relies on randomly selected or weakly perturbed negatives, which limits representation discriminability. To address this, we design a gradient-guided adversarial process that generates informative negatives while maintaining structural proximity.

Adversarial Hard Negative Generation: Given the clean region embedding $\mathcal{E}'_k$, we generate a hard negative by introducing a small perturbation δ constrained under $\|\delta\|_2 \le \epsilon$. The objective is to minimize semantic consistency, defined by the prediction distribution $P(\mathcal{S}_k \mid \cdot)$:

$$\mathcal{L}_{\mathrm{neg}} = -\sum_{k=1}^{K} \log P\left(\mathcal{S}_k \mid \mathcal{E}'_k + \delta\right) \tag{12}$$

Gradient-Based Perturbation: We instantiate δ using normalized gradient ascent:

$$\mathcal{X}_*^- = \mathcal{E}'_k + \epsilon \frac{\boldsymbol{g}}{\|\boldsymbol{g}\|_2}, \quad \boldsymbol{g} = \nabla_{\mathcal{E}'_k} \mathcal{L}_{\mathrm{neg}} \tag{13}$$

This perturbation maximizes semantic discrepancy while preserving structural coherence, enabling the model to learn more robust and discriminative region boundary.

3) Integration and Task Relevance: To ensure effectiveness in downstream tasks, we apply separate linear projections to $\mathcal{X}_*^+$ and $\mathcal{X}_*^-$ to align them with specific urban forecasting objectives, as inspired by [39]. These modules significantly enhance embedding robustness in applications such as traffic forecasting, where fine-grained spatial distinctions are critical for accurate predictions.

Algorithm Analysis. The computational complexity of the Deviation Copy Generator and Trickster Generator is as follows. For the Deviation Copy Generator, two gradient computations are required: $\boldsymbol{v} = \nabla_{\mathcal{E}'_k} \mathcal{L}_{\mathrm{cl}^+}$ and $\boldsymbol{\omega} = \nabla_{\check{\mathcal{X}}} \mathcal{L}_{\mathrm{KL}}$, each with complexity $\mathcal{O}(d)$, where d is the embedding dimension. Additionally, computing the KL divergence across K relations incurs $\mathcal{O}(Kd)$, leading to a total per-sample complexity of $\mathcal{O}(Kd)$. For the Trickster Generator, the gradient computation $\nabla_{\mathcal{E}'_k} \mathcal{L}_{\mathrm{neg}}$ and normalization together cost $\mathcal{O}(d)$. When both modules are applied to N samples, the overall complexities become $\mathcal{O}(NKd)$ and $\mathcal{O}(Nd)$ for the Deviation Copy and Trickster modules, respectively. The dominant term depends on the number of relations K and embedding dimension d, with the DevCopy module being more computationally intensive as K increases.

4) Contrastive Optimization for Robust Augmentation: To enhance the robustness and generalizability of contrastive learning in noisy urban environments, we design a dual-objective InfoNCE-based loss [40], which separately models positive and negative contrastive signals. The goal is to

Algorithm 1 Deviation Copy Generator (DevCopy)

Input: Anchor embedding $\mathcal{E}'_k$, contrastive loss $\mathcal{L}_{\mathrm{cl}^+}$, perturbation step size ψ.

Output: Hard positive sample $\mathcal{X}_*^+$.

Stage 1: Initial Perturbation

Compute gradient of contrastive loss:

$$\boldsymbol{v} = \nabla_{\mathcal{E}'_k} \mathcal{L}_{\mathrm{cl}^+}$$

Apply perturbation to the anchor embedding:

$$\check{\mathcal{X}} = \mathcal{E}'_k - \psi \frac{\boldsymbol{v}}{\|\boldsymbol{v}\|_2}$$

Stage 2: Semantic Refinement

Minimize KL divergence between $\check{\mathcal{X}}$ and $\mathcal{E}'_k$:

$$\mathcal{L}_{KL} = \sum_{k=1}^{K} D_{KL}\left(P\left(\mathcal{S}_k \mid \check{\mathcal{X}}\right) \| P\left(\mathcal{S}_k \mid \mathcal{E}'_k\right)\right)$$

Compute gradient of KL loss:

$$\boldsymbol{\omega} = \nabla_{\check{\mathcal{X}}} \mathcal{L}_{KL}$$

Refine the positive sample:

$$\mathcal{X}_*^+ = \check{\mathcal{X}} - \psi \frac{\boldsymbol{\omega}}{\|\boldsymbol{\omega}\|_2}$$

return $\mathcal{X}_*^+$

Algorithm 2 Trickster Generator

Input: Clean region embedding $\mathcal{E}'_k$, semantic similarity matrix $\mathcal{S}_k$, perturbation budget ϵ.

Output: Adversarial hard negative sample $\mathcal{X}^-_*$.

Compute adversarial loss to minimize semantic consistency:

$$\mathcal{L}_{\text{neg}} = -\sum_{k=1}^{K} \log P\left(\mathcal{S}_k \mid \mathcal{E}'_k + \delta\right)$$

Compute gradient of the loss with respect to $\mathcal{E}'_k$:

$$\boldsymbol{g} = \nabla_{\mathcal{E}'_k} \mathcal{L}_{\text{neg}}$$

Generate adversarial hard negative sample via normalized gradient ascent:

$$\mathcal{X}^-_* = \mathcal{E}'_k + \epsilon \cdot \frac{\boldsymbol{g}}{\|\boldsymbol{g}\|_2}$$

return $\mathcal{X}^-_*$

maximize the agreement between anchor embeddings and their corresponding DevCopy samples, while minimizing similarity with adversarially crafted Trickster negatives.

The positive contrastive loss is defined as:

$$\mathcal{L}_{\text{cl}^+} = -\sum_{k=1}^{K} \log \frac{\varphi\left(\mathcal{E}'_k, \mathcal{X}^+_*\right)}{\sum_{\boldsymbol{e}'_x \in Z'} \varphi\left(\mathcal{E}'_k, \boldsymbol{e}'_x\right)}, \tag{14}$$

where $\varphi(\cdot) = \exp(\text{cosine}(\cdot)/\tau)$ is a temperature-scaled cosine similarity, and $Z' = Z \cup \left\{\mathcal{X}^-_*\right\}$ denotes the extended negative sample pool.

Similarly, the negative contrastive loss is formulated as:

$$\mathcal{L}_{\text{cl}^-} = -\sum_{k=1}^{K} \log \frac{\varphi\left(\mathcal{E}'_k, \mathcal{X}_*\right)}{\sum_{\boldsymbol{e}'_x \in Z'} \varphi\left(\mathcal{E}'_k, \boldsymbol{e}'_x\right)}, \tag{15}$$

where $\mathcal{X}_*$ is perturbation-augmented node embedding. We use different positive samples in the contrastive loss formulation to avoid overfitting to a single type of positive example. This strategy not only enhances the model's ability to distinguish between positive and negative samples but also encourages it to learn more nuanced representations that capture diverse positive semantic relationships. By ensuring that the model is exposed to multiple forms of positive signals, we promote a more generalizable learning process.

As a result, this approach increases the model's robustness to the inherent noise and heterogeneity of urban data, ensuring that it can handle a variety of urban contexts more effectively. This is especially important in noisy, real-world urban environments where data from multiple sources often interact in complex and unpredictable ways.

To balance the influence of positive and negative signals, we define the adversarial contrastive loss as:

$$\mathcal{L}_{\text{acl}} = \alpha \mathcal{L}_{\text{cl}^-} + (1-\alpha)\mathcal{L}_{\text{cl}^+}, \tag{16}$$

where $\alpha \in [0, 1]$ is a tunable hyperparameter that controls the relative emphasis on repulsive (negative) versus attractive (positive) forces.

TABLE II
DATA DESCRIPTION OF EXPERIMENTED DATASETS

Dataset	City	Number of Records
Human Mobility	NYC	10 million trips per month
	Chicago	2.4 million trips per month
POI and Check-ins	NYC	20,000 POI records
	Chicago	112,000 POI records
Crime Data	NYC	Over 40,000 records annually
	Chicago	Over 32,000 records annually

E. Joint Optimization

We integrate the adversarial contrastive loss $\mathcal{L}_{\text{acl}}$ with the attentive supervised learning objective $\mathcal{L}_{\text{asl}}$ to jointly optimize semantic fidelity and structural robustness. To prevent overfitting and promote generalization, we further apply ℓ_2 regularization to all learnable parameters Θ. The full objective is expressed as:

$$\mathcal{L}_{\text{total}} = \beta \mathcal{L}_{\text{asl}} + (1-\beta)\mathcal{L}_{\text{acl}} + \mu\|\Theta\|^2, \tag{17}$$

where $\beta \in [0, 1]$ governs the trade-off between supervised and self-supervised components, and μ denotes the regularization coefficient.

V. EXPERIMENTS

In this section, we present an extensive set of experiments designed to evaluate the effectiveness, security, and reliability of the proposed model.

A. Datasets

We evaluate our framework on three tasks: crime prediction, land usage classification, and check-in prediction using urban datasets from New York City (NYC) and Chicago. Regions are defined based on census block boundaries and street-level data, following [1], [4], and [9]. Specifically, NYC is divided into 180 regions and Chicago into 234.[1] **Human mobility data** is derived from NYC yellow taxi records[2] and Chicago trip data,[3] capturing approximately 10 million and 2.4 million trips per month, respectively. **POI and check-in data** is collected via the Foursquare API,[4] comprising about 20,000 POI records for NYC and 112,000 for Chicago. Each record includes the venue name, category, check-in count, and visitor statistics, used to characterize regional functions. **Crime data** is sourced from the NYC Open Data and Chicago Data Portal, providing over 40,000 and 32,000 annual crime records, respectively. Each entry includes the time, location, and type of offense.

To avoid temporal leakage in time series tasks (check-in and crime prediction), we use a chronological data split: the first 42 weeks (80%) as training and the final 10 weeks (20%) for testing. This mirrors real-world forecasting scenarios where models are trained on past data to predict future events. A summary of dataset statistics is shown in Table II.

[1] http://www.census.gov
[2] https://opendata.cityofnewyork.us/
[3] https://data.cityofchicago.org/
[4] https://location.foursquare.com/developer/

B. Experimental Settings

EUPAS is implemented in PyTorch. All experiments are conducted on an Intel(R) Xeon(R) E5-2680 v4 CPU with an NVIDIA GeForce RTX 3090 Ti-24G GPU. The dimension of our model is 144. The model is optimized using the Adam optimizer with a learning rate of 0.001. The heterogeneous GCNs component contain 3 layers. Key hyperparameters are set as follows: standard deviation $\sigma = 0.01$, scale factor $\eta = 1$, perturbation threshold $\epsilon = 1$, and weighting factors $\alpha = 0.50$, $\beta = 0.15$, and temperature $\tau = 4$.

All hyperparameters are selected via grid search on the validation set of the Manhattan dataset. The final selected values are then applied to both Manhattan and Chicago without further tuning to evaluate generalization across cities. For instance, τ is searched from $\{2, 4, 6\}$, while α and β are selected from $[0.1, 0.9]$ in increments of 0.1. We observe that EUPAS maintains stable performance across a wide range of hyperparameter settings, suggesting that the model does not heavily rely on fine-tuning. For fair comparison, all baseline models are tuned according to their original configurations or publicly recommended settings.

C. Baselines and Evaluation Metrics

We compare EUPAS with four categories of baseline methods:

1) Shallow Embedding Methods:

- **LINE** [41]: Captures first- and second-order proximities via edge-wise loss.
- **node2vec** [42]: Learns embeddings using biased random walks and the Skip-Gram model.
- **GAE** [43]: Encodes nodes via a graph autoencoder optimized for reconstruction.

2) Graph Neural Network (GNN) Methods:

- **GCN** [44]: Aggregates neighborhood features through spectral convolutions.
- **GraphSAGE** [45]: Learns inductive embeddings via sampled neighborhood aggregation.
- **GAT** [24]: Applies attention weights to neighboring nodes during message passing.

3) Multi-View and Spatial Graph Models:

- **POI** [9]: Generates region vectors based on POI-TFIDF spatial semantics.
- **HDGE** [4]: Integrates mobility paths and spatial relations through hybrid graphs.
- **ZE-Mob** [6]: Utilizes region co-occurrence in mobility trips for embedding.
- **MV-PN** [7]: Builds multi-view POI graphs for regional representation.
- **CGAL** [8]: Combines POI and mobility views using adversarial graph alignment.
- **MVGRE** [9]: Employs cross-view attention for mobility and POI integration.
- **MGFN** [10]: Fuses multi-mobility graphs using a multi-level attention module.

4) State-of-the-Art Urban Embedding Approaches:

- **ROMER** [11]: Captures multi-view global dependencies via graph attention and staged fusion.
- **HREP** [12]: Adopts prefix-tuning to guide region embeddings in downstream urban tasks.

Evaluation Metrics: For predictive tasks (crime and check-in), we adopt Mean Absolute Error (MAE), Root Mean Squared Error (RMSE), and R^2. For clustering (land usage), we report Normalized Mutual Information (NMI) and Adjusted Rand Index (ARI).

D. Main Results

1) Check-In and Crime Prediction: The check-in and crime prediction tasks aim to evaluate how well the learned region embeddings capture spatiotemporal dynamics and semantic information useful for forecasting future regional activity. Specifically, the goal is to predict the number of check-in events (e.g., at restaurants, shops) or the number of reported crimes in each region over the course of one year using the learned region embeddings.

We evaluate the predictive effectiveness of region embeddings on weekly check-in and crime counts across Manhattan and Chicago. Each city's dataset spans 52 weeks with **weekly granularity**, forming 52 time steps per region. Following a chronological split, the first 42 weeks (80%) are used for training and the remaining 10 weeks (20%) for testing, yielding 7,560 training and 1,800 testing samples per task per city.

Table III reports the results. In Manhattan, EUPAS achieves notable improvements over the strongest baseline (HREP) on check-in prediction: 6.9% (MAE), 2.9% (RMSE), and 2.7% (R^2). On crime prediction, the gains are even larger: 10.8% (MAE), 8.5% (RMSE), and 7.4% (R^2). Similar performance improvements are observed in Chicago, confirming the generalizability of EUPAS across cities with distinct spatial and socio-economic structures.

Classical methods such as LINE, node2vec, and GAE underperform due to their limited capacity to model complex spatial dependencies and vulnerability to noise. Multi-view approaches (HDGE, ZE-Mob, MV-PN) improve expressiveness but remain sensitive to sparsity and perturbation. Advanced graph-based models (MVGRE, MGFN, ROMER, HREP) incorporate attention and fusion mechanisms, yet still struggle under incomplete or noisy data.

EUPAS surpasses these baselines by jointly modeling semantic-preserving perturbations and adversarial contrastive signals. Its deviation-based augmentation and trickster-driven hard negative generation enhance the robustness of region embeddings. The consistent performance across both cities underscores EUPAS's resilience and scalability in real-world urban analytics.

2) Land Usage Classification: To evaluate the semantic quality of region embeddings, we perform land usage classification via unsupervised clustering. Land usage classification helps assess how well the learned embeddings represent functional and spatial aspects of urban regions, which is crucial for urban region profiling tasks. We apply K-means clustering to the learned embeddings, with the number of clusters set to $k = 12$ for both Manhattan and Chicago, corresponding to real-world zoning systems (12 Community Board zones in

TABLE III

PERFORMANCE COMPARISON ON URBAN PREDICTION TASKS. BOLD INDICATES THE BEST PERFORMANCE, UNDERLINE INDICATES THE SECOND-BEST

Models	Crime Prediction						Land Usage Classification				Check-in Prediction					
	NYC			Chicago			NYC		Chicago		NYC			Chicago		
	MAE	RMSE	R^2	MAE	RMSE	R^2	NMI	ARI	NMI	ARI	MAE	RMSE	R^2	MAE	RMSE	R^2
Shallow Embedding Models																
LINE	117.53	152.43	0.06	334.6	396.29	0.04	0.17	0.01	0.07	0.005	564.59	853.82	0.08	1128.1	1733.1	0.03
GAE	96.55	133.10	0.19	289.27	390.45	0.20	0.47	0.23	0.17	0.05	498.23	803.34	0.09	1107.84	1691.72	0.05
node2vec	75.09	104.97	0.49	150.1	209.93	0.32	0.58	0.35	0.21	0.10	372.83	609.47	0.44	744.1	1080.3	0.12
GNNs																
GCN	96.21	134.89	0.19	279.37	389.71	0.21	0.48	0.25	0.18	0.06	489.12	765.23	0.13	1033.65	1547.21	0.08
GAT	92.67	131.23	0.22	279.19	378.23	0.34	0.35	0.28	0.13	0.02	465.21	721.86	0.42	862.35	1168.43	0.07
GraphSage	93.58	133.11	0.21	281.41	385.92	0.26	0.29	0.19	0.09	0.01	478.00	751.22	0.23	902.82	1307.86	0.03
Multi-view / Mobility Models																
POI	94.71	129.01	0.24	288.77	379.04	0.25	0.31	0.16	0.10	0.01	482.12	568.21	0.39	1021.91	1439.02	0.04
HDGE	72.65	96.36	0.58	145.2	208.78	0.45	0.59	0.29	0.22	0.08	399.28	536.27	0.57	799.7	1240.1	0.29
ZE-Mob	101.98	132.16	0.20	304.4	396.1	0.14	0.61	0.39	0.25	0.14	360.71	592.92	0.47	796.6	1071.9	0.2
MV-PN	92.30	123.96	0.30	276.5	372.9	0.35	0.38	0.16	0.15	0.03	476.14	784.25	0.08	888.7	1241.6	0.02
CGAL	71.82	95.49	0.59	149.3	192.16	0.52	0.64	0.40	0.27	0.18	335.58	529.18	0.59	679.5	991.2	0.24
MVGRE	69.28	96.51	0.57	142.9	189.6	0.57	0.78	0.61	0.37	0.24	312.63	513.02	0.61	553.21	875.8	0.31
MGFN	70.21	89.60	0.63	139.5	174.4	0.64	0.76	0.58	0.35	0.22	292.60	451.76	0.69	511.27	851.63	0.35
Recent State-of-the-art Models																
ROMER	67.17	86.46	0.68	141.3	187.9	0.61	0.81	0.67	0.41	0.28	285.44	433.96	0.73	501.9	821.21	0.38
HREP	65.66	84.59	0.68	130.8	178.2	0.65	0.80	0.65	0.39	0.26	270.28	406.53	0.75	482.18	801.4	0.41
EUPAS	**58.56**	**77.41**	**0.73**	**117.6**	**160.9**	**0.71**	**0.84**	**0.69**	**0.46**	**0.31**	**251.70**	**394.68**	**0.77**	**468.83**	**769.41**	**0.44**
Improvements	**10.8%**	**8.5%**	**7.4%**	**10.1%**	**7.7%**	**8.5%**	**3.7%**	**3.0%**	**12.2%**	**10.7%**	**6.9%**	**2.9%**	**2.7%**	**2.8%**	**4.0%**	**7.3%**

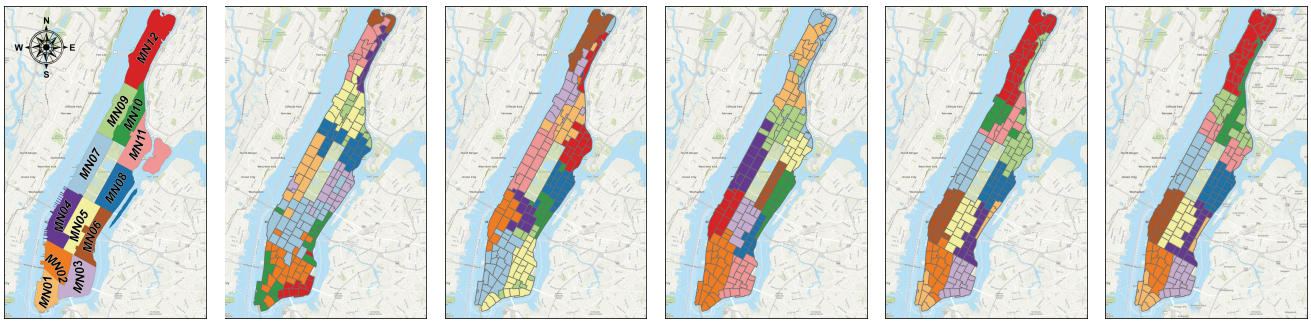

(a) Districts (b) HDGE (c) ZE-Mob (d) MVGRE (e) ROMER (f) Ours

Fig. 4. Comparison of Manhattan districts and region clusters obtained under different baseline methods.

Manhattan and functional land use categories in Chicago, as defined by the Chicago Metropolitan Planning Agency, e.g., residential, commercial, industrial).

As shown in Table III, EUPAS consistently outperforms all baselines. In Manhattan, it improves over ROMER by 4.9% (NMI) and 5.6% (ARI), indicating superior semantic alignment. In the more heterogeneous Chicago dataset, EUPAS achieves even greater improvements-12.2% (NMI) and 10.7% (ARI), demonstrating its strong ability to generalize to more complex, less structured urban layouts. Models such as ZE-Mob and MVGRE struggle with ambiguous or mixed-use regions, leading to unclear cluster boundaries. In contrast, EUPAS benefits from perturbation-aware contrastive learning and high-level semantic guidance, which enable it to generate more coherent region embeddings that better align with real land use categories.

Fig. 4 provides a visual comparison of clustering results in Manhattan. Regions within the same cluster are colored identically, and EUPAS demonstrates significantly higher alignment with real land use boundaries, showcasing its ability to capture both functional and spatial semantics despite data noise or sparsity.

E. Ablation Study

We conduct ablation experiments on all three downstream tasks to evaluate the contribution of key components in EUPAS. The following variants are considered:

w/o Spatial Augmentation: Replaces the spatial perturbation mechanism with standard augmentation to generate positive samples.

w/o Supervised: Removes the attentive supervised module.

w/o Self-Supervised: Removes the adversarial contrastive module.

w/o Trickster: Excludes the adversarially generated hard negatives $\mathcal{X}_*^-$.

w/o DevCopy: Excludes the adversarially generated hard positives $\mathcal{X}_*^+$.

As shown in Fig. 5, each module contributes meaningfully to the overall performance. Removing spatial perturbation

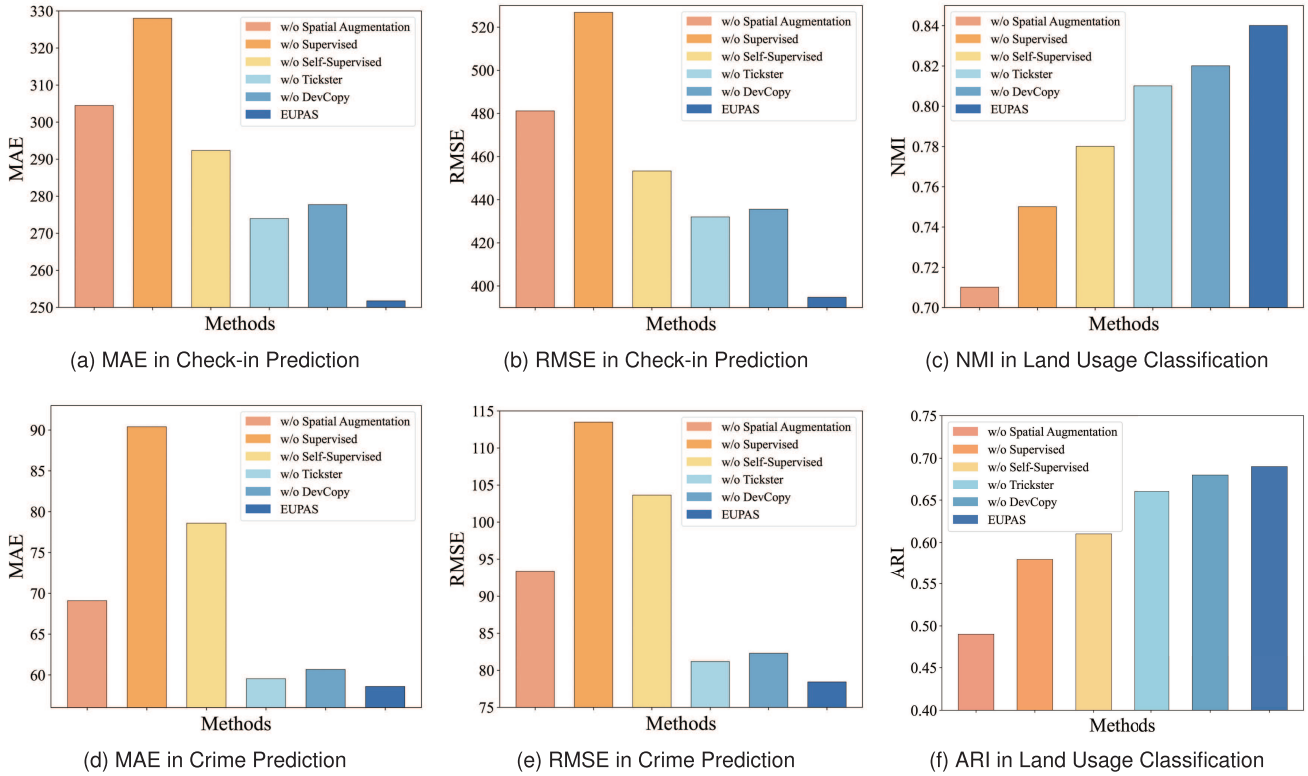

(a) MAE in Check-in Prediction (b) RMSE in Check-in Prediction (c) NMI in Land Usage Classification

(d) MAE in Crime Prediction (e) RMSE in Crime Prediction (f) ARI in Land Usage Classification

Fig. 5. Ablation studies for three tasks on NYC dataset.

TABLE IV
PERFORMANCE OF DIFFERENT INTERACTION FUNCTIONS

Interaction Function	Check-in (MAE)	Robust Performace (MAE)	Land Usage Classification (NMI)	Robust Performace (NMI)	Inference time
Hadamard (ours)	251.7	286.31	0.84	0.74	3.812
Concat + MLP	270.6	301.42	0.83	0.73	3.951
Bilinear	255.9	292.59	0.78	0.67	3.294

notably degrades accuracy, indicating the importance of structured augmentation in capturing urban spatial semantics. The absence of the self-supervised module causes the most severe drop, underscoring its critical role in improving robustness against noisy and incomplete data.

Furthermore, eliminating either the Trickster or DevCopy module leads to noticeable performance declines, demonstrating the effectiveness of generating hard negatives and hard positives for enhancing representation discrimination. Meanwhile, excluding the supervised module results in unstable clustering and reduced predictive accuracy, highlighting the necessity of hybrid supervision.

Interaction Mechanism. To investigate the effectiveness of the element-wise product used to model interactions between node and relation embeddings (Eq. (1)), we compare it against two alternatives:

- **Concatenation**: $[e_{vk}^{(l-1)}; h_k^{(l-1)}]$ followed by a linear layer.
- **Bilinear Pooling**: $e_{vk}^{\top} \mathbf{W}_b h_k$ as a learnable bilinear interaction.

As shown in Table IV, although the Hadamard product slightly lags behind Bilinear in terms of inference time, it outperforms both alternatives in terms of check-in prediction accuracy (MAE) and robustness under adversarial attacks (PGD). Given its superior accuracy and effectiveness in large-scale urban modeling, we conclude that the Hadamard product is the optimal choice despite the minor trade-off in inference time.

F. Hyperparameter Sensitivity

We evaluate the sensitivity of EUPAS to key hyperparameters by varying one parameter at a time while fixing others to their default settings in the NYC dataset. The results are presented in Fig. 6, covering three critical hyperparameters: the temperature τ, contrastive weight α, and the balancing coefficient β.

Effect of τ. This parameter controls the sharpness of similarity distributions in the contrastive loss (Equations 14, 15). We vary τ across $\{0.1, 0.5, 1, 2, 3, 4, 5, 6\}$ and observe that smaller values ($\leq$ 3) tend to destabilize training by overemphasizing hard negatives in the similarity distribution. Conversely, large values (>4) overly smooth the similarity space, reducing discrimination power. The optimal setting, $\tau = 4$, provides a good balance between gradient sharpness and embedding robustness.

Effect of α. This coefficient weights the contributions of positive and negative sample losses in adversarial

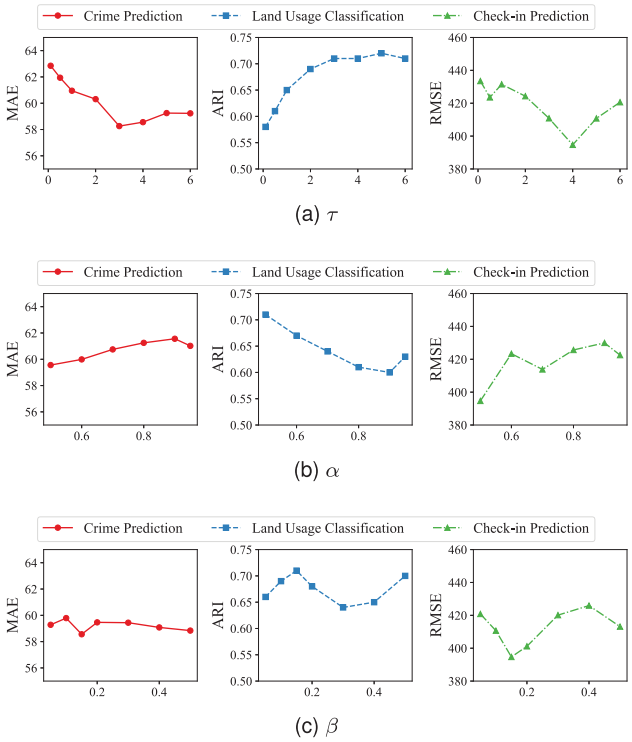

Fig. 6. Impact of τ, α and β to our model.

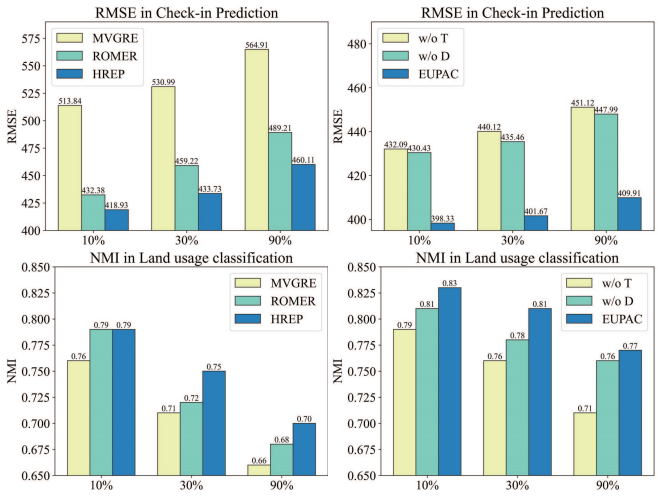

Fig. 7. Performance changes as the training data qualities decrease.

contrastive learning (Equation 16). We test values from $\{0.5, 0.6, 0.7, 0.8, 0.9, 0.95\}$ and find that $\alpha = 0.5$ performs best, suggesting that equal emphasis on positive and negative examples leads to more stable optimization and improved contrastive discrimination.

Effect of β. This parameter controls the contribution of the self-supervised loss in the total objective (Equation 17). We explore $\beta \in \{0.05, 0.1, 0.15, 0.2, 0.3, 0.4, 0.5\}$ and observe that $\beta = 0.15$ offers optimal performance. This reflects the complementary nature of supervised and self-supervised learning in our framework, with supervised learning contributing more strongly in noisy, real-world settings.

G. Adversarial Robustness Evaluation

We comprehensively evaluate the robustness of the proposed EUPAS model under both natural noise and adversarial perturbations. The evaluation consists of three perspectives: resilience to random noise, robustness against white-box adversarial attacks, and generalization under black-box transfer attacks.

1) Robustness to Realistic Data Noise: To simulate natural disturbances in urban sensing data (e.g., mobility counts, sensor fluctuations), we inject Poisson noise [46] into different portions of the input. Poisson noise is particularly suited for count-based or discrete event modeling, common in traffic and check-in datasets. Specifically, we perturb 10%, 30%, and 90% of the input embeddings using Poisson-distributed values (mean noise level =1). As visualized in Fig. 7, EUPAS consistently outperforms all baselines across noise levels. While performance naturally degrades as noise intensity increases, EUPAS shows significantly higher tolerance than alternatives, especially in downstream prediction tasks.

TABLE V
WHITE-BOX PGD ATTACK RESULTS FOR EUPAS AND HREP MODELS

Model	Task	Metric	Original	Robust Performance
EUPAS	Crime Prediction	MAE	58.56	64.11
		RMSE	77.41	84.80
	Land Usage Classification	NMI	0.84	0.74
		ARI	0.69	0.53
	Check-in Prediction	MAE	251.70	286.31
		RMSE	394.68	449.21
HREP	Crime Prediction	MAE	65.66	80.45
		RMSE	84.59	107.25
	Land Usage Classification	NMI	0.80	0.62
		ARI	0.65	0.45
	Check-in Prediction	MAE	270.28	331.45
		RMSE	406.53	502.15

2) Resistance to White-Box Adversarial Attacks: In this experiment, we assess the vulnerability of both EUPAS and HREP models under a white-box attack scenario using PGD, a widely used gradient-based adversarial attack method. PGD iteratively optimizes the adversarial perturbations based on the model's gradients, making it effective in evaluating the robustness of deep learning models. The attack parameters are configured with a maximum perturbation of $\lambda = 0.03$, a step size of $\alpha = 0.01$, and 20 iterations, applying an ℓ_2 norm constraint. The evaluation is conducted both on the original data and the adversarial samples generated by PGD.

Table V presents the results of the PGD attack on both models, showing the performance degradation for each model under adversarial perturbation. The results indicate that while both models experience a decrease in performance under PGD, EUPAS maintains a higher level of robustness with lower MAE and RMSE values compared to HREP. This demonstrates that adversarial contrastive learning, the core technique in EUPAS, effectively enhances the model's resilience against adversarial attacks, consistently outperforming the strongest baseline model.

3) Robustness Under Transfer-Based Black-Box Attacks: To further evaluate the robustness of HREP and EUPAS under

TABLE VI

PERFORMANCE OF EUPAS AND HREP MODELS AGAINST BLACK-BOX ATTACKS ON THE NYC DATASET. EACH COLUMN DENOTES THE ATTACK STRENGTH WITH $\lambda = 0.02$ AND $\lambda = 0.1$. WE GENERATE PGD ATTACK EXAMPLES (PGD) AND EXPECTATION OF TRANSFORMATION ATTACK (EOT) FROM EUPAS AND HREP MODELS. EACH ROW SHOWS THE PERFORMANCE OF THE TARGET MODEL TRAINED WITH ℓ_2

		NYC							
Target \ Source		$\lambda = 0.02$				$\lambda = 0.10$			
		ROMER	HREP	EUPAS(PGD)	EUPAS(EoT)	ROMER	HREP	EUPAS(PGD)	EUPAS(EoT)
Crime prediction (MAE)	ROMER	-	72.81	79.71	89.02	-	84.58	88.06	90.49
	HREP	67.69	-	77.60	85.14	70.21	-	85.11	88.91
	EUPAS	63.51	66.36	-	-	69.55	80.93	-	-
Land usage classification (NMI)	ROMER	-	0.80	0.78	0.72	-	0.72	0.67	0.70
	HREP	0.79	-	0.77	0.70	0.70	-	0.69	0.67
	EUPAS	0.83	0.82	-	-	0.79	0.76	-	-

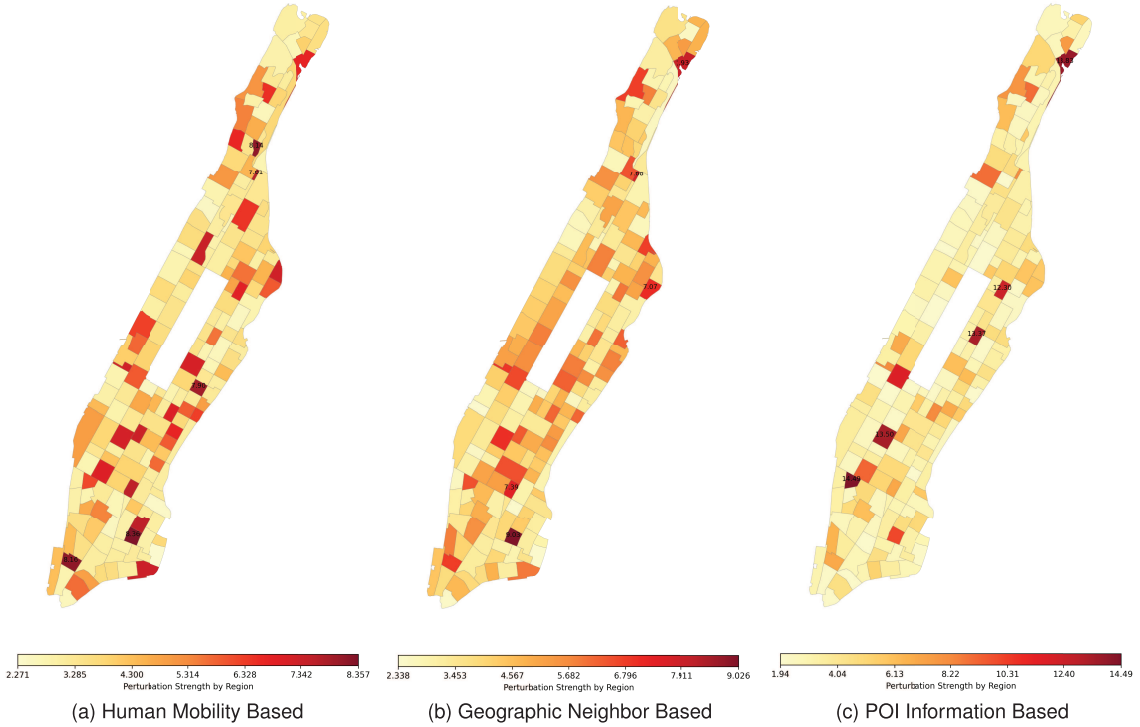

(a) Human Mobility Based (b) Geographic Neighbor Based (c) POI Information Based

Fig. 8. Spatial perturbation strength heatmaps in Manhattan derived from embeddings generated using different input data sources.

realistic attack scenarios, we conduct transfer-based black-box attacks on the NYC dataset. In this setting, adversarial examples are crafted using a source model and then transferred to a target model for evaluation, simulating cases where the attacker has no access to the internal parameters of the victim model. We also generate black-box adversarial examples with EUPAS by attacking the EUPAS with a linear layer using the PGD attack (EUPAS (PGD)), and the EUPAS with a projector using the instance-wise attack (EUPAS (EoT)).

As shown in Table VI, EUPAS consistently generates stronger black-box attacks compared to HREP and ROMER. When EUPAS is used as the source model, it significantly degrades the performance of all target models. For example, under the EoT attack with $\lambda = 0.10$, the MAE of HREP increases from 67.69 to 88.91 in crime prediction, and the NMI of ROMER drops from 0.80 to 0.67 in land usage classification. These large margins indicate high transferability of adversarial examples from EUPAS. In contrast, attacks transferred from HREP and ROMER are relatively weaker. Across both tasks, adversarial examples generated from HREP result in smaller MAE increases and less NMI degradation on EUPAS. This asymmetry highlights EUPAS's dual advantage: it is more effective as an attacker and more robust as a defender. These results validate that EUPAS not only benefits from stronger adversarial representation learning, but also acts as a more potent adversarial generator, affirming its utility in real-world threat scenarios where transferability is critical.

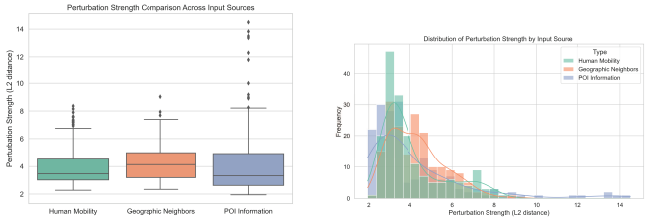

(a) Perturbation strengths across input types. (b) Distribution of perturbation strengths (L2 distance).

Fig. 9. Statistical analysis of perturbation strengths derived from different semantic sources.

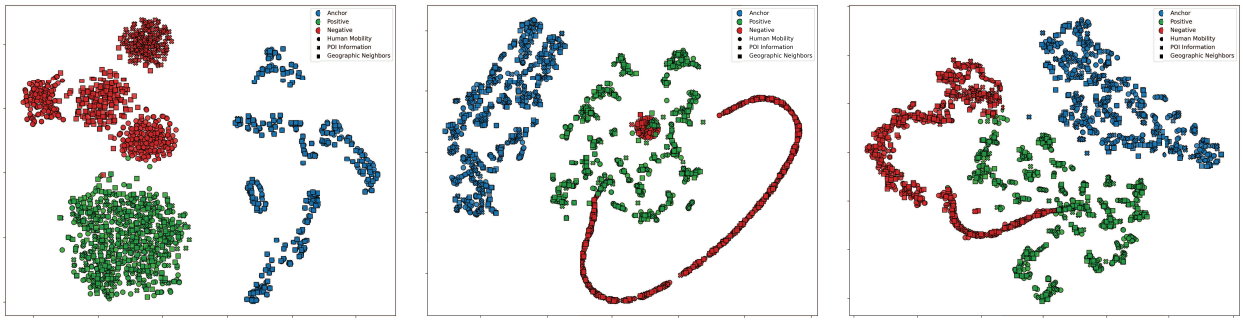

(a) Initial Stage (b) Mid Stage (c) Final Stage

Fig. 10. t-SNE visualizations of anchors, positives, and adversarial negatives with different embedding sources across stages of training.

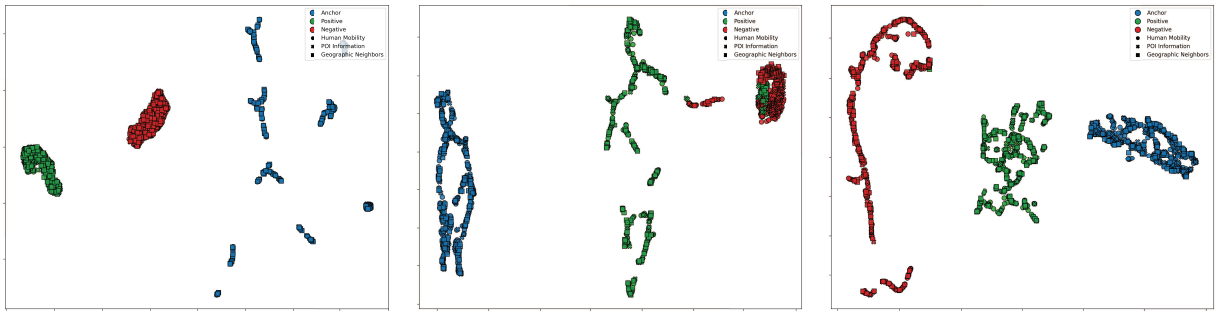

(a) Initial Stage (b) Mid Stage (c) Final Stage

Fig. 11. UMAP visualizations of anchors, positives, and adversarial negatives with different embedding sources across stages of training.

H. Regional Interpretability of Adversarial Perturbations

To provide spatial interpretability and better understand the mechanism behind adversarial perturbations, we conduct a region-level analysis of perturbation strengths across Manhattan using embeddings derived from three semantic sources: Human Mobility, Geographic Neighbors, and POI Information. As visualized in Fig. 8, the spatial heatmaps reveal distinct perturbation patterns under different inputs. Human Mobility and POI-based embeddings show concentrated perturbations in high-activity areas such as Midtown and Downtown, reflecting their greater semantic complexity and susceptibility to adversarial disruption. In contrast, the Geographic Neighbor-based map displays a more uniform distribution, suggesting less regional sensitivity under this relation.

To quantify these differences, Fig. 9a provides a boxplot of perturbation magnitudes (ℓ_2 distance). POI embeddings exhibit higher variability and a larger median perturbation, indicating greater adversarial vulnerability. This pattern is further confirmed by the density distribution in Fig. 9b, where POI-based embeddings show a longer tail and heavier skew, revealing a higher frequency of extreme perturbations in specific urban regions.

I. Contrastive Embedding Evolution With DevCopy and Trickster

To better understand the role of the proposed DevCopy and Trickster modules, we qualitatively visualize how they shape the contrastive embedding space. While our methodology conceptually distinguishes DevCopy as generating semantically aligned hard positives and Trickster as producing semantically divergent but locally proximate negatives, these effects are further validated through empirical visualization in Fig. 10 and Fig. 11.

We use both t-SNE and UMAP to track the evolution of anchor, positive, and adversarial negative embeddings at different training stages (initial, mid, final). t-SNE excels at preserving local structures and revealing data clusters, while UMAP better captures both local and global structures, maintaining the overall data geometry and large-scale relationships. In early training, all samples are loosely distributed with little semantic structure. As training progresses, DevCopy samples move closer to anchor clusters, enhancing intra-class compactness. Meanwhile, Trickster samples remain nearby in Euclidean space but form distinct manifolds, reflecting their role as hard negatives that challenge semantic consistency.

These dynamics confirm the intended behavior: DevCopy helps the model learn to tolerate local variations without losing semantic alignment, while Trickster introduces adversarially misleading yet spatially close samples to improve contrastive discrimination. By the final stage, embeddings form well-separated, semantically meaningful clusters, demonstrating that our generators not only provide hard samples but also induce a more structured and interpretable feature space.

J. Computational Efficiency

To comprehensively evaluate the computational efficiency of the EUPAS model, we conducted ablation experiments on two city datasets of different scales: NYC and Chicago. The training and inference time comparisons are reported in Table VII. The results show that EUPAS exhibits a clear advantage in training time over previous methods, particularly considering the complexity of the model, while maintaining high efficiency.

TABLE VII
COMPARISON OF TRAINING TIME & INFERENCE TIME

City	Task	Model	Time (s/100epoch)	
			Training	Inference
NYC	Crime Prediction	MVGRE	33.46	1.207
		ROMER	31.38	1.232
		HREP	8.143	1.987
		w/o Self-Supervised	1.621	1.368
		w/o Tickster	3.703	3.276
		w/o DevCopy	2.713	2.233
		EUPAS (Ours)	4.024	3.812
Chicago	Land Usage Classification	MVGRE	48.76	1.842
		ROMER	46.11	1.787
		HREP	11.64	2.723
		w/o Self-Supervised	2.23	2.035
		w/o Tickster	4.88	4.641
		w/o DevCopy	3.99	3.258
		EUPAS (Ours)	5.84	4.982

To the best of our knowledge, we are the first to incorporate security concerns in the field of urban area analysis. The Trickster Generator and Deviation Copy modules we propose, while adding some overhead to inference time, are specifically designed to enhance predictive accuracy while addressing security-related challenges. This trade-off allows EUPAS to not only surpass traditional methods in accuracy but also make a significant contribution to ensuring data privacy and security.

VI. CONCLUSION

In this paper, we propose EUPAS, a robust and efficient framework for urban region representation learning, addressing critical challenges such as noise, data incompleteness, and semantic biases, which are central concerns in the field of secure and trustworthy data-driven modeling. Our framework combines a joint attentive supervised and adversarial contrastive learning approach, which ensures reliable and resilient performance in urban tasks like check-in prediction, crime prediction, and land usage classification. By introducing innovative components such as perturbation augmentation and adversarial contrastive modules, EUPAS effectively mitigates the negative impact of noise and data incompleteness, providing a more robust solution for urban data analysis. Our experiments on two structurally distinct cities, Manhattan and Chicago, demonstrate that EUPAS generalizes well across different urban environments. The model maintains high performance in all tasks despite variations in city morphology and data distribution, confirming its potential for broader deployment in smart city applications.

Looking forward, we aim to extend EUPAS's capabilities further by focusing on enhancing its security features. Specifically, we plan to incorporate stronger adversarial defenses to better withstand sophisticated attacks and explore the integration of differential privacy techniques to safeguard sensitive urban data.

REFERENCES

[1] Q. Zhang, C. Huang, L. Xia, Z. Wang, Z. Li, and S. Yiu, "Automated spatio-temporal graph contrastive learning," in *Proc. ACM Web Conf.*, Apr. 2023, pp. 295–305, doi: 10.1145/3543507.3583304.

[2] X. Yang, S. He, K. G. Shin, M. Tabatabaie, and J. Dai, "Cross-modality and equity-aware graph pooling fusion: A bike mobility prediction study," *IEEE Trans. Big Data*, vol. 11, no. 1, pp. 286–302, Feb. 2025.

[3] L. Gong et al., "Contrastive pre-training with adversarial perturbations for check-in sequence representation learning," in *Proc. 37th AAAI Conf. Artif. Intell. 35th Conf. Innov. Appl. Artif. Intell. 13th Symp. Educ. Adv. Artif. Intell.*, vol. 37, 2023, pp. 4276–4283, doi: 10.1609/aaai.v37i4.25546.

[4] H. Wang and Z. Li, "Region representation learning via mobility flow," in *Proc. ACM Conf. Inf. Knowl. Manage.*, Nov. 2017, pp. 237–246.

[5] N. Jean, S. Wang, A. Samar, G. Azzari, D. B. Lobell, and S. Ermon, "Tile2Vec: Unsupervised representation learning for spatially distributed data," in *Proc. AAAI Conf. Artif. Intell.*, May 2018, pp. 3967–3974.

[6] Z. Yao, Y. Fu, B. Liu, W. Hu, and H. Xiong, "Representing urban functions through zone embedding with human mobility patterns," in *Proc. 27th Int. Joint Conf. Artif. Intell.*, Jul. 2018, pp. 3919–3925.

[7] Y. Fu, P. Wang, J. Du, L. Wu, and X. Li, "Efficient region embedding with multi-view spatial networks: A perspective of locality-constrained spatial autocorrelations," in *Proc. AAAI Conf. Artif. Intell.*, vol. 33, no. 1, Jul. 2019, pp. 906–913.

[8] Y. Zhang, Y. Fu, P. Wang, X. Li, and Y. Zheng, "Unifying inter-region autocorrelation and intra-region structures for spatial embedding via collective adversarial learning," in *Proc. 25th ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, Jul. 2019, pp. 1700–1708.

[9] M. Zhang, T. Li, Y. Li, and P. Hui, "Multi-view joint graph representation learning for urban region embedding," in *Proc. 29th Int. Joint Conf. Artif. Intell.*, Jul. 2020, pp. 4431–4437.

[10] S. Wu et al., "Multi-graph fusion networks for urban region embedding," in *Proc. 31st Int. Joint Conf. Artif. Intell.*, Jul. 2022, pp. 2312–2318.

[11] W. Chan and Q. Ren, "Region-wise attentive multi-view representation learning for urban region embedding," in *Proc. 32nd ACM Int. Conf. Inf. Knowl. Manage.*, Oct. 2023, pp. 3763–3767, doi: 10.1145/3583780.3615194.

[12] S. Zhou, D. He, L. Chen, S. Shang, and P. Han, "Heterogeneous region embedding with prompt learning," in *Proc. AAAI Conf. Artif. Intell.*, Jun. 2023, vol. 37, no. 4, pp. 4981–4989.

[13] Z. Wang, W. Zhang, W. Bao, F. Long, and C. Yuan, "Adaptive contrastive learning for learning robust representations under label noise," in *Proc. 31st ACM Int. Conf. Multimedia*, Oct. 2023, pp. 4917–4927, doi: 10.1145/3581783.3612491.

[14] M. Abdelfattah, M. Hassan, and A. Alahi, "MaskCLR: Attention-guided contrastive learning for robust action representation learning," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2024, pp. 18678–18687.

[15] Z. Li, W. Huang, K. Zhao, M. Yang, Y. Gong, and M. Chen, "Urban region embedding via multi-view contrastive prediction," in *Proc. AAAI Conf. Artif. Intell.*, Mar. 2024, vol. 38, no. 8, pp. 8724–8732.

[16] L. Yang, L. Zhang, and W. Yang, "Graph adversarial self-supervised learning," in *Proc. Adv. Neural Inf. Process. Syst.*, M. Ranzato, A. Beygelzimer, Y. Dauphin, P. Liang, and J. W. Vaughan, Eds., Red Hook, NY, USA: Curran Associates, vol. 34, 2021, pp. 14887–14899. [Online]. Available: https://proceedings.neurips.cc/paper_files/paper/2021/file/7d3010c11d08cf990b7614d2c2ca9098-Paper.PDF

[17] T. Chen, S. Liu, S. Chang, Y. Cheng, L. Amini, and Z. Wang, "Adversarial robustness: From self-supervised pre-training to fine-tuning," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2020, pp. 696–705.

[18] M. Kim, J. Tack, and S. J. Hwang, "Adversarial self-supervised contrastive learning," 2020, *arXiv:2006.07589*.

[19] L. Fan, S. Liu, P.-Y. Chen, G. Zhang, and C. Gan, "When does contrastive learning preserve adversarial robustness from pretraining to finetuning?," 2021, *arXiv:2111.01124*.

[20] Y. Zhang, Z. Wei, J. Sun, and M. Sun, "Adversarial representation engineering: A general model editing framework for large language models," 2024, *arXiv:2404.13752*.

[21] M. Schlichtkrull, T. N. Kipf, P. Bloem, R. van den Berg, I. Titov, and M. Welling, "Modeling relational data with graph convolutional networks," in *Proc. Eur. Semantic Web Conf.*, 2018, pp. 593–607.

[22] C. Zhang, D. Song, C. Huang, A. Swami, and N. V. Chawla, "Heterogeneous graph neural network," in *Proc. 25th ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, 2019, pp. 793–803.

[23] W. L. Hamilton, R. Ying, and J. Leskovec, "Inductive representation learning on large graphs," 2018, *arXiv:1706.02216*.

[24] P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Liò, and Y. Bengio, "Graph attention networks," 2017, *arXiv:1710.10903*.

[25] W. Jin et al., "Adversarial attacks and defenses on graphs," *ACM SIGKDD Explor. Newslett.*, vol. 22, no. 2, pp. 19–34, Jan. 2021.

[26] S. Suresh, L. Pan, C. Hao, and J. Neville, "Adversarial graph augmentation to improve graph contrastive learning," in *Proc. Adv. Neural Inf. Process. Syst.*, Nov. 2021, pp. 15920–15933.
[27] D. Zhu, Z. Zhang, P. Cui, and W. Zhu, "Robust graph convolutional networks against adversarial attacks," in *Proc. ACM SIGKDD Int. Conf. Knowl. Discov. Data Min.*, 2019, pp. 1399–1407, doi: 10.1145/3292500.3330851.
[28] Y. Luo, F.-L. Chung, and K. Chen, "Urban region profiling via multi-graph representation learning," in *Proc. 31st ACM Int. Conf. Inf. Knowl. Manage.*, Oct. 2022, pp. 4294–4298, doi: 10.1145/3511808.3557720.
[29] L. Zhao, T. Liu, X. Peng, and D. Metaxas, "Maximum-entropy adversarial data augmentation for improved generalization and robustness," in *Proc. 34th Int. Conf. Neural Inf. Process. Syst.*, 2020, pp. 14435–14447.
[30] Z. Jiang, T. Chen, T. Chen, and Z. Wang, "Robust pre-training by adversarial contrastive learning," in *Proc. Adv. Neural Inf. Process. Syst. (NIPS)*, 2020, pp. 16199–16210.
[31] K. He, H. Fan, Y. Wu, S. Xie, and R. Girshick, "Momentum contrast for unsupervised visual representation learning," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2020, pp. 9726–9735.
[32] J.-B. Grill et al., "Bootstrap your own latent: A new approach to self-supervised learning," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 33, 2020, pp. 21271–21284.
[33] D. Xu, W. Cheng, D. Luo, H. Chen, and X. Zhang, "InfoGCL: Information-aware graph contrastive learning," 2021, *arXiv:2110.15438*.
[34] A. Madry, A. Makelov, L. Schmidt, D. Tsipras, and A. Vladu, "Towards deep learning models resistant to adversarial attacks," 2017, *arXiv:1706.06083*.
[35] Z. Allen-Zhu and Y. Li, "Feature purification: How adversarial training performs robust deep learning," in *Proc. IEEE 62nd Annu. Symp. Found. Comput. Sci. (FOCS)*, Los Alamitos, CA, USA, Feb. 2022, pp. 977–988.
[36] Z. Wei, Y. Wang, Y. Guo, and Y. Wang, "CFA: Class-wise calibrated fair adversarial training," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2023, pp. 8193–8201.
[37] M. Kim, J. Tack, and S. J. Hwang, "Adversarial self-supervised contrastive learning," in *Proc. Adv. Neural Inf. Process. Syst. (NeurIPS)*, 2020, pp. 1–11.
[38] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Apr. 2016, pp. 770–778.
[39] Y. Xie, Z. Xu, J. Zhang, Z. Wang, and S. Ji, "Self-supervised learning of graph neural networks: A unified review," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 45, no. 2, pp. 2412–2429, Feb. 2023.
[40] A. van den Oord, Y. Li, and O. Vinyals, "Representation learning with contrastive predictive coding," 2018, *arXiv:1807.03748*.
[41] J. Tang, M. Qu, M. Wang, M. Zhang, J. Yan, and Q. Mei, "Line: Large-scale information network embedding," in *Proc. 24th Int. Conf. World Wide Web*, 2015, pp. 1067–1077, doi: 10.1145/2736277.2741093.
[42] A. Grover and J. Leskovec, "Node2Vec: Scalable feature learning for networks," in *Proc. 22nd ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, Jul. 2016.
[43] T. N. Kipf and M. Welling, "Variational graph auto-encoders," 2016, *arXiv:1611.07308*.
[44] T. N. Kipf and M. Welling, "Semi-supervised classification with graph convolutional networks," 2016, *arXiv:1609.02907*.
[45] W. Hamilton, Z. Ying, and J. Leskovec, "Inductive representation learning on large graphs," in *Proc. Adv. Neural Inf. Process. Syst.*, I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, Eds., Red Hook, NY, USA: Curran Associates, vol. 30, 2017, pp. 1–11.
[46] A. Papoulis and S. Pillai, *Probability, Random Variables, and Stochastic Processes*. New York, NY, USA: McGraw-Hill, 2002.

Weiliang Chen received the B.S. degree in software engineering and the M.S. degree in computer science and technology from Heilongjiang University, Harbin, China, in 2022 and 2025, respectively.

His research interests include deep learning, urban data mining, and trustworthy artificial intelligence.

Qianqian Ren received the Ph.D. degree in computer science and technology from Harbin Institute of Technology, Harbin, China, in 2012. She was a Visiting Scholar at the University of Alberta, Edmonton, AB, Canada, from 2018 to 2019. She is currently a Professor with the School of Computer Science and Technology, Heilongjiang University, Harbin. Her research interests include graph learning, spatio-temporal data mining, and recommendation systems.

Yong Liu received the B.S. and M.S. degrees in computer science and technology from Heilongjiang University, China, in 1998 and 2001, respectively, and the Ph.D. degree from Harbin Institute of Technology, China, in 2010. He is currently a Professor with the School of Computer Science and Technology, Heilongjiang University. He has authored or co-authored over 50 publications. His research interests include graph learning and social network analysis.

Jianguo Sun is currently a Professor with Xidian University, China, and the Ph.D. Advisor. He serves as the Director for Zhejiang Engineering Research Center for Data Security Governance and the Technical Lead for the National Engineering Laboratory for Big Data System Computing Technology. He is also a Visiting Research Fellow at the Institute of Information Engineering, Chinese Academy of Sciences. His research interests include data and security governance, intelligent software technologies, and multi-modal computing.

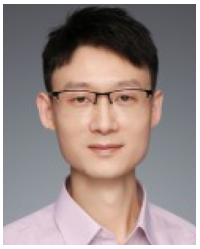

Feng Lin (Senior Member, IEEE) received the Ph.D. degree from the Department of Electrical and Computer Engineering, Tennessee Technological University, Cookeville, TN, USA, in 2015. He was an Assistant Professor with the University of Colorado Denver, Denver, CO, USA; a Research Scientist with the State University of New York (SUNY) at Buffalo, Buffalo, NY, USA; and an Engineer with Alcatel-Lucent (currently, Nokia). He is currently a Professor with the School of Cyber Science and Technology, College of Computer Science and Technology, Zhejiang University, China. His current research interests include mobile sensing, the Internet of Things security, biometrics, AI security, and the IoT applications. He was a recipient of the Best Paper Award from ACM MobiSys'20, IEEE Globecom'19, and IEEE BHI'17; the Best Demo Award from ACM HotMobile'18; and the First Prize Design Award from the 2016 International 3D Printing Competition.