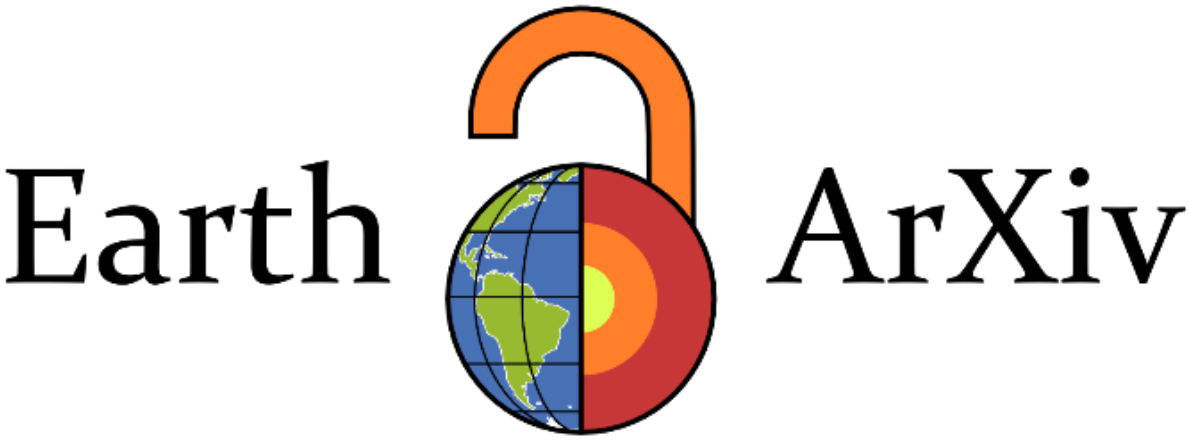

Can AlphaEarth Foundations Redefine the Paradigm of Gridded Population Mapping? A Systematic Evaluation across 18 Global Cities and Large-Scale Mapping Applications

Shuyang Hou[a], Haoyue Jiao[a], Qingyang Xu[a], Yaxian Qing[a], Ziqi Liu[a], Lutong Xie[a], Zhangyan Xu[a], Xuefeng Guan[a], Huayi Wu[a]*

*a. State Key Laboratory of Information Engineering in Surveying, Mapping, and Remote Sensing, Wuhan University, Wuhan, China

*Corresponding author: Huayi Wu, email: wuhuayi@whu.edu.cn

Abstract

Population mapping typically relies on census data and its update cycles, and is further constrained by manual feature engineering and limited cross-regional generalization. AlphaEarth Foundations (AEF) provides globally consistent, analysis-ready 64-dimensional annual surface embeddings, offering a new data foundation for reducing dependence on frequent census updates and enabling more transferable, data-driven mapping approaches. However, its suitability for modeling population distribution—a socio-economic variable that is not directly observable—remains insufficiently evaluated. This study investigates 18 cities distributed across six continents, integrating AEF data from 2017–2024 with WorldPop 100 m population grids from 2017–2020. We evaluate the performance of linear models, random forests, and deep learning models under four training strategies, four spatial scales, and four categories of auxiliary variables. Results show that AEF embeddings achieve strong explanatory power under random-split validation, with median R^2 values ranging from 0.82 to 0.96. However, spatial block cross-validation reveals a pronounced dependence on spatial autocorrelation, with R^2 decreasing by 0.46–0.67. Compared to cross-city training, single-city multi-year training yields more stable spatial generalization, while cross-city transfer is constrained by distributional shifts in AEF features. Among auxiliary variables, spatial structure factors provide the most substantial improvements in spatial generalization, whereas high-dimensional points-of-interest (POI) features may introduce redundancy and degrade performance in deep learning models. PCA and SHAP analyses further indicate that AEF dimensions exhibit a backbone–long-tail contribution pattern, where a small number of dominant dimensions coexist with many weak contributors; more complete utilization of the full 64-dimensional embedding is associated with improved spatial generalization. Based on the optimal combination of strategies, we generate annual population maps for 2021–2024 at 1 km resolution globally and 100 m resolution for the 18 cities. These results demonstrate that AEF-driven approaches can effectively complement temporal gaps in WorldPop and support the updating of high-resolution population distribution data.

Keywords: Population mapping; AlphaEarth Foundations; WorldPop; global mapping; multi-source data fusion modeling

1. Introduction

Population spatial patterns constitute a fundamental data basis for fine-grained urban governance,

disaster risk assessment, and public health response, and also provide essential support for multiple indicators of the United Nations Sustainable Development Goals (SDGs) (Leyk et al., 2019; Liu et al., 2024). Traditional population statistics are typically reported at the administrative-unit level, resulting in coarse spatial resolution that fails to capture intra-urban heterogeneity (Georganos et al., 2022; Liu et al., 2024). To address this limitation, products such as WorldPop, LandScan, and GHSL adopt dasymetric mapping approaches, redistributing population counts from census units to fine-resolution grids using multi-source ancillary variables, and achieving population distribution estimates at approximately 100 m resolution on a global scale (Dobson et al., 2000; Pesaresi et al., 2024; Tatem, 2017). However, this paradigm still faces several key limitations. First, population modeling generally relies on census or statistical data as calibration benchmarks, and the long update cycles of censuses may constrain the timeliness of population mapping (Neal et al., 2022; Thomson et al., 2020). For example, although WorldPop provides datasets for 2015–2030, only data up to 2020 are directly constrained by census inputs, while subsequent years are primarily based on model extrapolation, limiting their ability to capture short-term population dynamics driven by conflicts, disasters, or policy changes; similarly, GHS-POP is released at five-year intervals, with the current R2023A version providing estimates up to 2030. Second, improvements in mapping accuracy depend on the construction and integration of multiple remote sensing–derived variables, such as nighttime lights, land use, and building morphology (Pesaresi et al., 2024; Stevens et al., 2015; Zheng et al., 2023); this process typically involves complex data processing workflows and may encounter challenges in ensuring consistency and transferability across regions, while the introduction of deep learning methods is further constrained by the scarcity of pixel-level population labels required for supervision (Georganos et al., 2022; Metzger et al., 2024). Third, manual feature engineering relies heavily on prior knowledge, and spatial signals related to population distribution that are not explicitly defined may remain underutilized, thereby limiting the model's ability to capture complex spatial patterns (Tseng et al., 2023). Taken together, these limitations indicate that current approaches still have room for improvement in terms of update efficiency, feature representation completeness, and information utilization, motivating the exploration of unified representations capable of encoding diverse surface semantics, coupled with data-driven models to better extract latent signals underlying population spatial distributions.

In July 2025, Google DeepMind released AlphaEarth Foundations (AEF) along with its annual satellite embedding data product (Brown et al., 2025). At a global scale, AEF integrates optical imagery, multi-temporal synthetic aperture radar (SAR) observations, and spatiotemporal contextual information, mapping multi-source remote sensing data into surface representations that are consistent across sensors and time, and generating a 64-dimensional embedding vector for each surface unit at approximately 10 m resolution. Compared with traditional population mapping approaches that rely on manual feature engineering, AEF offers three potential advantages. First, AEF embeddings are independent of census updates, thereby potentially reducing reliance on frequently updated census data in population mapping. Second, multi-source remote sensing information is compressed into dense embeddings within a unified framework, which helps reduce the need for region-specific feature engineering workflows and improves the consistency of feature representation. Third, the 64-dimensional embeddings are learned from

surface observations through representation learning, enabling the capture of spatial information that is difficult to explicitly define using manually engineered features. However, whether the representational capacity of AEF can effectively support population spatial modeling remains to be systematically evaluated.

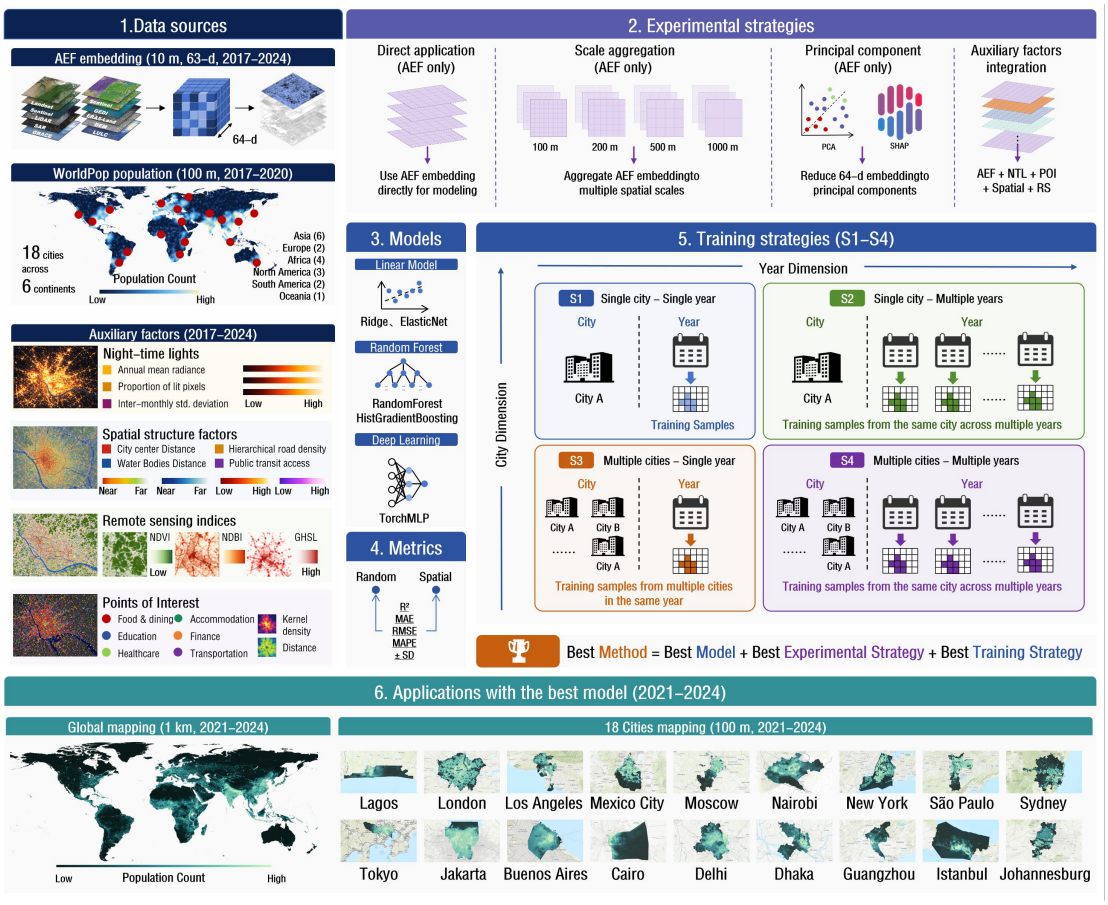

Fig. 1. Overall research framework. Part 1 lists the three data sources; Parts 2–4 specify the experimental strategies, regression models, and evaluation metrics; Part 5 defines the four training strategies S1–S4 over the city and year dimensions; Part 6 shows the mapping applications produced by the best-performing combination for 2021–2024.

Existing studies on AEF have primarily focused on tasks such as crop mapping (Fang et al., 2025; Ma et al., 2026; Zvonkov et al., 2025), land cover classification (Khan & Ahmad, 2025), and environmental variable estimation (Alvarez et al., 2025; Qu et al., 2026). These tasks typically exhibit relatively direct spectral, structural, or temporal correspondences with remote sensing observations, and are therefore well aligned with the surface representation design of AEF. In contrast, population distribution is not directly observable from remote sensing data, but rather emerges as a socio-economic spatial outcome shaped by multiple factors, including the built environment, urban functions, industrial structure, transportation accessibility, policy changes, and disruptive events (Wu et al., 2005). Consequently, the ability of AEF embeddings to represent population spatial structures, their generalization performance under different training strategies, and their interactions with auxiliary variables cannot be reliably inferred from the design principles alone. Against this backdrop, this study systematically evaluates the applicability of

AEF for population spatial modeling and further explores population mapping approaches that leverage AEF as the core representation, with the aim of supporting temporal updates of global population distribution data beyond 2020.

Against this backdrop, this study focuses on the applicability of AEF to population spatial modeling and its optimal usage, and formulates five core research questions:

- Q1 (Feasibility): Can the 64-dimensional surface embeddings of AEF effectively characterize urban population spatial distributions without incorporating any auxiliary features? How stable is their performance across different cities and training–testing evaluation settings?
- Q2 (Training strategies): How does the performance of AEF-based population modeling vary under four training strategies — single-city single-year (S1), single-city multi-year (S2), cross-city same-year (S3), and globally mixed (S4)? Which strategy best leverages the spatiotemporally aligned embeddings provided by AEF?
- Q3 (Scale adaptability): Does the performance of AEF-based population modeling remain consistent when scaling from 100 m to 1 km resolution? What are its performance limits and scale-dependent characteristics?
- Q4 (Dimensional structure): Do the 64 dimensions of AEF embeddings contain redundant information? Is there a subset of core dimensions that contributes disproportionately to population prediction, and is this contribution pattern consistent across cities?
- Q5 (Auxiliary gains): When integrating AEF with auxiliary features such as nighttime lights, remote sensing indices, spatial structure variables, and points-of-interest (POIs), what are the marginal gains associated with each feature type, and do these gains depend on the choice of training strategy?

To address the five research questions outlined above, this study integrates AEF embedding data (2017 – 2024) for 18 cities across six continents, WorldPop population data (2017 – 2020), and multiple auxiliary variables, including nighttime lights, POIs, remote sensing indices, and road network features, to systematically evaluate the applicability of AEF for population spatial modeling and identify optimal usage strategies. Based on the optimal model configurations, two mapping applications are conducted: (1) global population mapping at 1 km resolution for 2021–2024 to assess the extrapolation stability of the proposed framework at the global scale; and (2) population mapping at 100 m resolution for the 18 cities over 2021 – 2024. The overall study design and technical workflow are illustrated in Fig. 1.

The main contributions of this study are as follows:

- We develop a comprehensive experimental framework that integrates multiple model types, training strategies, spatial scales, and auxiliary variables, providing, to our knowledge, the most systematic empirical evaluation of AEF for population mapping as an indirectly coupled task.

- Through spatial block cross-validation, we quantify the dependence of AEF embeddings on spatial autocorrelation and demonstrate that random-split validation can substantially overestimate model generalization performance.
- By combining PCA and SHAP analyses, we identify a backbone – long-tail contribution structure in the AEF embeddings, showing that effective spatial generalization relies on leveraging a near-complete 64-dimensional representation; we further find that spatial structure variables, together with the single-city multi-year (S2) training strategy, form an effective combination for improving generalization performance.
- We show that spatial structure variables serve as a critical complement for enhancing generalization, whereas high-dimensional POI features may introduce negative gains in deep learning models, providing empirical evidence for both beneficial and detrimental auxiliary factors.
- Based on the optimal model configurations, we generate annual population maps for 2021–2024 at 1 km resolution globally and at 100 m resolution for the 18 cities, offering a practical data-driven solution to complement temporal gaps in WorldPop beyond 2020.

The remainder of this paper is organized as follows: Section 2 reviews related work on population mapping and AEF; Section 3 describes the study areas and datasets; Section 4 presents the modeling methods and experimental design; Section 5 reports the results; Section 6 demonstrates population mapping applications at 1 km resolution globally and 100 m resolution for the 18 cities; and Section 7 discusses the applicability and limitations of the proposed approach, and concludes with directions for future research.

2. Related Work

2.1. Population Spatial Distribution Mapping

Population spatial distribution mapping aims to transform population statistics or proxy observations into fine-resolution gridded population distributions. Early approaches were based on statistical interpolation and spatial redistribution. Tobler's (1979) pycnophylactic interpolation emphasized generating smooth continuous surfaces while preserving zonal population totals (Tobler, 1979), whereas Mennis (2003) introduced dasymetric mapping to constrain population redistribution using ancillary data such as land use (Mennis, 2003). These methods established the methodological foundation of population mapping, but typically relied on smoothing assumptions or coarse categorical partitions, limiting their ability to capture complex intra-urban heterogeneity.

With the advancement of remote sensing technologies and the increasing availability of global geospatial data, global population products such as LandScan (Dobson et al., 2000), GHSL (Pesaresi et al., 2024), and WorldPop (Stevens et al., 2015; Tatem, 2017) have been developed. These products generally integrate multiple ancillary variables—including nighttime lights, land cover, road networks, building morphology, and topography—and apply statistical or machine learning models to disaggregate population data into grid cells, substantially improving the spatial detail of population representations. In recent years, mobile phone data (Deville et al., 2014), social media data (Xu et al., 2021), and point-of-interest (POI) data (Yang et al., 2019; Zhang & Zhao, 2024) have also been incorporated to complement the limitations of traditional remote

sensing variables in capturing dynamic population activities. However, such data are often constrained by privacy concerns, platform biases, spatial coverage, and cross-regional consistency, making them difficult to use for globally consistent population mapping.

2.2. Remote Sensing Foundation Representation Models and Applications of AEF

Early remote sensing modeling largely followed a task-driven paradigm, in which separate models were developed for specific applications such as land cover classification, crop identification, or environmental variable estimation (Huo et al., 2025; Xiao et al., 2025). With the rapid growth of multi-source remote sensing data, this paradigm has gradually shifted toward a "pretrained general surface representation followed by downstream task adaptation" framework (Mendieta et al., 2023). Tile2Vec learns semantic embeddings of the Earth's surface based on spatial neighborhood consistency (Jean et al., 2018), while SeCo (Manas et al., 2021) extracts temporally stable features through consistency constraints across multi-temporal imagery. More recent approaches, such as Presto (Tseng et al., 2023) and TESSERA (Feng et al., 2025), further explore unified representations from multi-source time series data. Although these studies have advanced foundational representation learning in remote sensing, most are still delivered as models or tools, requiring users to perform data acquisition, preprocessing, and feature extraction in practical applications.

AEF extends this paradigm by transforming foundational representations from model-based tools into directly usable annual embedding data products (Brown et al., 2025). Developed by Google DeepMind, AEF integrates multi-source Earth observation data along with spatial, temporal, and observational context to generate globally consistent 64-dimensional annual surface embeddings at approximately 10 m resolution, providing standardized, analysis-ready inputs for downstream mapping tasks. Publicly available information indicates that the AEF dataset spans 2017–2025, and its official documentation highlights its general-purpose representation capability across tasks such as land use and land cover mapping, crop identification, and species distribution modeling.

Existing evaluations of AEF have primarily focused on surface or environmental tasks that are more directly linked to remote sensing observations, including land cover classification (Benavides-Martinez et al., 2026), crop mapping (Ma et al., 2026), yield estimation (Rahman, 2026), and species distribution modeling (Fang et al., 2026). In these cases, the target variables typically exhibit relatively clear spectral, structural, or temporal relationships with remote sensing signals, making them well suited to benefit from AEF's multi-source surface representations. In contrast, population distribution is not directly observable from remote sensing data but emerges as a socio-spatial outcome shaped by the built environment, economic activity, social structure, and policy factors (Song et al., 2024). As a result, whether AEF embeddings can effectively capture spatial variations in population density, how well they generalize across cities, scales, and training strategies, and whether traditional auxiliary variables remain necessary all remain open questions. These gaps constitute the primary motivation of this study.

3. Study Areas and Data

3.1. Study Areas

This study selects 18 cities distributed across six continents as the study areas. The selection follows three main criteria: (1) geographic coverage, ensuring representation across Asia, Europe, Africa, North America, South America, and Oceania, including cities such as Tokyo, Guangzhou, Delhi, Moscow, Istanbul, London, Cairo, Lagos, Nairobi, New York, Los Angeles, Mexico City, São Paulo, Buenos Aires, and Sydney; (2) diversity in city size and development stage, covering a range of population scales, levels of urbanization, income levels, and built environment types to evaluate model performance across heterogeneous urban contexts; and (3) consistency with WorldPop data availability, prioritizing cities located in countries with stable national-level WorldPop coverage to ensure the availability and temporal consistency of training labels. Detailed information on each city is provided in Appendix A, and their geographic distribution is shown in Fig. 2.

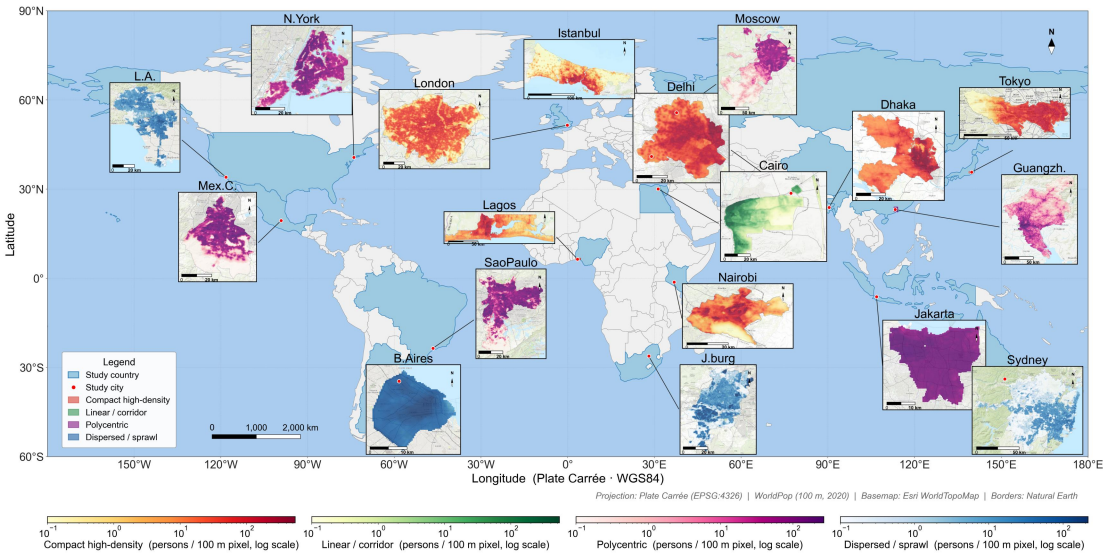

Fig. 2. Geographic distribution and population density patterns of the 18 study cities. Insets show 2020 WorldPop population density (100 m, log scale, persons/pixel), color-coded by urban morphology: compact high-density (red), linear/corridor (green), polycentric (purple), and dispersed/sprawl (blue). The sample spans six continents and four morphology classes.

3.2. Data

3.2.1. AEF

AEF data for eight years (2017–2024) were acquired via the Google Earth Engine (GEE) platform, covering 18 cities worldwide. A total of 144 valid files were obtained, with an overall data volume of approximately 38 GB. For each city, AEF data are stored as 64-band GeoTIFF files in the corresponding UTM coordinate system, with a spatial resolution of approximately 10 m and individual file sizes ranging from 0.13 to 0.37 GB.

3.2.2. WorldPop Population Grids

WorldPop employs a random forest – based modeling framework that integrates multi-source ancillary variables, including census data, nighttime lights, land cover, and road networks, to generate annual population grids at 100 m resolution (unit: persons per pixel) through a dasymetric redistribution approach. It is one of the most widely used global population datasets with extensive coverage and high spatial resolution. In this study, WorldPop data are used as reference labels for model training and accuracy assessment. The publicly available dataset extends to 2020; therefore, data from 2017 – 2020, temporally aligned with AEF, are used for training and evaluation.

3.2.3. Auxiliary Variables

Four categories of auxiliary variables are incorporated to complement aspects of urban socio-spatial structure that may not be fully captured by AEF embeddings. Their marginal contributions are evaluated through feature combination experiments. Nighttime lights (NTL) are derived from VIIRS/DNB monthly composites to generate annual statistical features, supplemented by DMSP-OLS 2013 stable lights as a historical baseline to quantify expansion dynamics, resulting in 20 feature layers that jointly characterize urban economic activity intensity and temporal variation. Remote sensing indices are computed from Landsat 8 annual composites, including NDVI and NDBI, along with GHSL built-up density, forming three feature layers to represent vegetation cover and impervious surface distribution. Spatial structure variables are derived from historical OpenStreetMap (OSM) data, including distance to city center, distance to water bodies, hierarchical road network density, and public transport accessibility, yielding 10 feature layers that describe urban centrality gradients and transportation structure. POIs are extracted from OSM, with kernel density and nearest-distance metrics calculated for 12 functional categories, resulting in 24 feature layers that capture the spatial concentration of urban service functions. Detailed definitions of all variables are provided in Appendix B.

3.3. Data Preprocessing

For spatial alignment, all datasets are projected to the corresponding UTM coordinate system of each city and resampled to a 100 m grid. WorldPop grids are reprojected and resampled using bilinear interpolation, and all auxiliary variables are processed accordingly to ensure pixel-level alignment. The original AEF data, with a spatial resolution of approximately 10 m, are aggregated to 100 m by applying mean pooling to each of the 64 bands, producing a 64-dimensional embedding vector for each 100 m grid cell.

For spatial masking, all data layers are clipped to city administrative boundaries, retaining only valid pixels within each boundary to avoid interference from sparsely populated peripheral areas and to focus the analysis on intra-urban population structure.

For feature normalization, linear models and neural networks, which are sensitive to feature scale, are standardized using RobustScaler based on the median and interquartile range to reduce the influence of outliers, whereas tree-based ensemble models, which rely on feature ordering and split thresholds, are trained using raw inputs without scaling.

Given the long-tailed distribution of population density, a logarithmic transformation is applied to

the target variable for all models:

$$y' = \log(1 + y) \tag{1}$$

Where y denotes the population count per pixel. This transformation compresses high-value ranges and reduces skewness, improving model performance in low- and medium-density areas. Model predictions are converted back to the original scale during evaluation using.

$$y = exp(y') - 1 \tag{2}$$

4. Methods

4.1. Models and Training Parameter Selection

The functional relationship between AEF embeddings and population spatial distribution is not known a priori and may involve linear responses, nonlinear transformations, and higher-order feature interactions. Therefore, this study constructs a multi-paradigm comparison framework comprising linear models, ensemble learning models, and neural network models to characterize the explanatory capacity of AEF representations for population mapping under different hypothesis spaces.

4.1.1. Models

Linear and regularized regression models. Ridge regression and ElasticNet are selected as linear baselines. Ridge regression mitigates multicollinearity among high-dimensional features through an L2 penalty, whereas ElasticNet combines L1 and L2 regularization to balance feature selection and parameter stability.

Ensemble learning and tree-based models. Random Forest and Histogram-based Gradient Boosting (HistGBT) are used as representative tree-based models. Random Forest reduces variance through bootstrap sampling and random feature selection, while HistGBT fits residuals sequentially through gradient boosting, providing a favorable trade-off between nonlinear modeling capacity and computational efficiency. Both models make no assumptions about feature scaling and are therefore trained using raw feature inputs.

Deep neural network model. A multilayer perceptron (TorchMLP) is implemented in PyTorch, incorporating batch normalization and dropout to reduce overfitting, together with the Adam optimizer, cosine annealing learning-rate scheduling, and early stopping. CNNs and Transformers are not included in this study. These architectures rely on spatial neighborhood windows and positional encoding, respectively, and are designed for spatially continuous two-dimensional imagery. In contrast, the input used here consists of independent pixel-level embedding vectors, with spatial topology not explicitly encoded in the input; therefore, the inductive biases of CNNs and Transformers are unlikely to be fully exploited.

4.1.2. Hyperparameter Selection

For the linear and tree-based ensemble models,unified configurations are determined within reasonable ranges based on prior studies and preliminary experiments,without conducting large-scale grid search.The regularization coefficients of Ridge regression and ElasticNet are set to

1.0 and 0.01,respectively.For the tree-based models,the maximum depth is set within 6–12,the learning rate is set to 0.05,and subsampling is used to enhance generalization.Detailed settings are provided in Table C1 of Appendix C.

The multilayer perceptron is more sensitive to hyperparameter choices;therefore,its configuration is determined through multi-round iterative pruning.In the early stage,a broad search is conducted using lightweight sampling,with approximately 2×10^4 pixels per pair,covering hidden-layer architecture,initial learning rate $\{2 \times 10^{-4}, 5 \times 10^{-4}, 1 \times 10^{-3}\}$,batch size {2048,4096,8192} ,and dropout rate {0.1,0.2,0.3} .Values that consistently underperform the baseline are removed in successive rounds. In the later stage, the retained configurations are re-evaluated using all available pixels, and the mean test-set R2R^2R2 across two independent random seeds is used as the primary selection criterion, where R2R^2R2 denotes the coefficient of determination as defined in Section 4.2. The final configuration is reported in Table C2 of Appendix C.

4.2. Evaluation Metrics

Model performance is evaluated from three perspectives: goodness of fit, absolute error, and relative error. Fit-related metrics are computed in the log-transformed space, whereas error-based metrics are calculated on the original population scale after inverse transformation.

Goodness of fit is assessed using the coefficient of determination (R^2)and the Pearson correlation coefficient (r):

$$R^2 = 1 - \frac{\sum_{i=1}^{n} (y_i - \hat{y}_i)^2}{\sum_{i=1}^{n} (y_i - \bar{y})^2} \tag{3}$$

$$r = \frac{\sum_{i=1}^{n} (y_i - \bar{y})(\hat{y}_i - \bar{\hat{y}})}{\sqrt{\sum_{i=1}^{n} (y_i - \bar{y})^2} \sqrt{\sum_{i=1}^{n} (\hat{y}_i - \bar{\hat{y}})^2}} \tag{4}$$

where y_i and $\hat{y}_i$ denote the observed and predicted values in log space,respectively,and$\bar{y}$is the mean of the observed values.

Absolute errors are measured using mean absolute error (MAE)and root mean squared error (RMSE), both computed on the original population scale:

$$MAE = \frac{1}{n} \sum_{i=1}^{n} \left| \tilde{y}_i - \tilde{\hat{y}}_i \right| \tag{5}$$

$$RMSE = \sqrt{\frac{1}{n} \sum_{i=1}^{n} (\tilde{y}_i - \tilde{\hat{y}}_i)^2} \tag{6}$$

Relative error is evaluated using mean absolute percentage error (MAPE), calculated only for

pixels with$\tilde{y}_i > 10^{-6}$:

$$MAPE = \frac{1}{n}\sum_{i=1}^{n}\left|\frac{\tilde{y}_i - \tilde{\hat{y}}_i}{\tilde{y}_i}\right| \times 100\% \quad (7)$$

For cross-validation experiments, metrics are averaged across folds, and the standard deviation of R^2 is used to quantify model stability.

4.3. Experimental Design

This section evaluates the capability of AEF for population spatial modeling through four progressive levels: (1) baseline modeling performance of AEF embeddings without introducing any auxiliary variables, assessed under four training strategies (Section 4.3.1); (2) extension of the optimal strategy to multiple spatial resolutions to examine scale adaptability (Section 4.3.2); (3) analysis of the internal structure of the 64-dimensional embeddings using PCA and SHAP (Section 4.3.3); and (4) incorporation of multi-source auxiliary variables under the four training strategies to evaluate their marginal contributions (Section 4.3.4).

4.3.1. Direct Application Performance

The four training strategies progressively expand the spatiotemporal coverage of the training data. S1 represents local training within the same city and year(baseline); S2 extends along the temporal dimension to multiple years within the same city; S3 expands along the spatial dimension to multiple cities within the same year; and S4 extends across both temporal and spatial dimensions simultaneously. Each strategy is evaluated under two modes:random and spatial.In the random mode,pixels from the target city–year are randomly split into training and test sets with an 80/20 ratio.In the spatial mode,valid pixels are partitioned into spatial blocks based on a 3×3 grid.Specifically,an enclosing rectangle is constructed from the spatial extent of valid pixels for each city and divided evenly into 3×3 blocks;pixels falling outside the boundary are assigned to the nearest block, and edge blocks with very few pixels (<1%of total pixels) are merged with adjacent blocks.

Given the substantial variation in the number of valid pixels across cities (ranging from 2×10^4 to 7×10^5), strategies S2–S4 apply equal-size sampling to control the training set size and eliminate the influence of sample size differences on cross-strategy comparisons. Cross-validation is performed using spatial blocks as fold units, and final performance is reported as the average across folds. The test set partition is fixed prior to experimentation and shared across S1–S4, such that differences between strategies arise solely from variations in the spatiotemporal extent of the training data.

To further control for the influence of data volume, the total training sample size in S2 – S4 is strictly constrained to match that of S1. When additional data sources are introduced, samples from each source are drawn proportionally to their original pixel counts, jointly filling the same training quota as S1. The composition of training datasets under the four strategies is summarized in Table 1.

Table 1. Training and test set construction for the four training strategies (S1–S4) under random and spatial evaluation.

Strategy	Expansion Dimension	Test Set	Random Training Set	Spatial Training Set
S1	Baseline	City C, Year Y	Remaining 80% of pixels from City C, Year Y	Remaining spatial blocks of City C, Year Y
S2	+Temporal	City C, Year Y	Proportional sampling from (remaining 80% of C–Y + all pixels from other years of C), total size equal to remaining 80% of C–Y	Proportional sampling from (complement of current fold in C–Y + all pixels from other years of C), total size equal to complement of current fold in C–Y
S3	+Spatial	City C, Year Y	Proportional sampling from (remaining 80% of C–Y + all pixels from Year Y in other cities), total size equal to remaining 80% of C–Y	Proportional sampling from (complement of current fold in C–Y + all pixels from Year Y in other cities), total size equal to complement of current fold in C–Y
S4	+Temporal + Spatial	City C, Year Y	Proportional sampling from (remaining 80% of C–Y + all pixels from other years of C + all pixels from all years in other cities), total size equal to remaining 80% of C–Y	Proportional sampling from (complement of current fold in C–Y + all pixels from other years of C + all pixels from all years in other cities), total size equal to complement of current fold in C–Y

4.3.2. Scale Adaptation

This experiment evaluates the consistency of AEF-based population modeling performance across spatial scales ranging from 100 m to 1 km. WorldPop labels at 100 m and AEF embeddings are aggregated to 200 m, 500 m, and 1000 m resolutions using spatial summation and mean pooling, respectively. At each target resolution, training and test datasets are constructed independently, while keeping the model architecture and training configuration unchanged, and the full evaluation framework described in Section 4.3.1 is repeated. To examine whether potential performance degradation of TorchMLP at coarser scales arises from a mismatch between model capacity and sample size rather than inherent limitations of neural network approaches, an additional variant, TorchMLP-Adaptive, is introduced as a control. This model dynamically adjusts network capacity according to the effective training sample size at each scale, while keeping all other hyperparameters identical to the original TorchMLP; detailed configurations are provided in Table C3 of Appendix C. By comparing the performance differences between the original TorchMLP and its adaptive counterpart across scales, we assess whether scale sensitivity can be mitigated through capacity adjustment, thereby providing a more comprehensive characterization of the limits of AEF-based population modeling across spatial resolutions.

4.3.3. Feature Importance Analysis

- Principal Component Analysis (PCA). The 64-dimensional AEF embeddings are reduced to $k \in \{8,16,24,32,38,48,56,64\}$ principal components using PCA. Under the optimal training strategy, models are retrained and evaluated for all 18 city–year combinations. The saturation behavior of both random and spatial R^2 with respect to k is analyzed, together with the cumulative explained variance curve, to determine an empirical upper bound on the effective dimensionality of AEF representations.

- SHAP analysis. Under the S1 and S2 strategies, models are retrained for all 18 city–year combinations, and mean absolute SHAP values for each dimension are computed immediately after training. Linear models use LinearSHAP on the full test set, while tree-based models use TreeSHAP with 2,000 test samples. For TorchMLP, a three-level fallback strategy is adopted: DeepExplainer is used preferentially (500 test samples, 100 background samples); if it fails, GradientExplainer is applied; if both fail, KernelSHAP is used (200 test samples, 50 background samples, 50 perturbations). The ranking consistency of high-contribution dimensions across the 18 cities is quantified, and statistical significance of cross-city consensus is evaluated using 1,000 permutation tests, allowing differentiation between globally consistent signals and location-specific patterns. The PCA and SHAP results are interpreted jointly to identify the effective dimensionality of AEF representations from complementary perspectives of information concentration and feature contribution.

4.3.4. Effects of Auxiliary Variables

This section systematically evaluates the marginal contributions of four categories of auxiliary variables to AEF-based population modeling: nighttime lights (NTL), remote sensing indices (RS-Index), spatial structure variables (Spatial), and POI; detailed descriptions of these variables are provided in Section 3.4 and Appendix B. The experiments are based on the best-performing model under the S1 strategy. Auxiliary variables are directly concatenated with the 64-dimensional AEF embeddings to form extended input features. Using the AEF-only configuration (C0) as the baseline, each category of auxiliary variables is introduced individually (C1 – C4) to assess its independent contribution, followed by a full combination scheme (C5). The configurations are summarized in Table 2.

Table 2. Auxiliary factor combinations (C0–C5) for evaluating marginal gains over the AEF baseline.

Scheme	Feature Combination	Dimensionality
C0	AEF (baseline)	64
C1	AEF + NTL	84
C2	AEF + RS-Index	67
C3	AEF + Spatial	74
C4	AEF + POI	88
C5	AEF + NTL + RS-Index + Spatial + POI (full)	121

4.4. Experimental Setup

All experiments are conducted on a local high-performance workstation. The hardware and software configurations are summarized in Table 3.

Table 3. Hardware and software configuration of the experimental environment.

Category	Configuration
CPU	Multi-core processor (≥16 cores)
GPU	NVIDIA GeForce RTX 4090 × 2 (24 GB VRAM each)
Memory	128 GB RAM
Storage	SSD (≥2 TB)
Operating System	Ubuntu 20.04 LTS
Programming Language	Python 3.10
Deep Learning Framework	PyTorch 2.4.1 (CUDA 12.1)

5. Results

5.1. Direct Application Performance

Under the random 80/20 split, the performance of the five models across five evaluation metrics for 2017 – 2020 is shown in Fig. 3. In terms of median performance, TorchMLP consistently achieves the best results across all years and metrics (R^2 mean: 0.952–0.956; MAE: 8.89–9.43; RMSE: 23.39 – 24.74), whereas ElasticNet performs the worst (R^2 median: 0.821 – 0.828). Nevertheless, all models achieve median R^2 values above 0.82, with the minimum city-level R^2 remaining above 0.60, indicating that the surface information encoded in AEF embeddings can be effectively utilized across all three modeling paradigms.

In terms of variability, TorchMLP exhibits the narrowest interquartile range across all metrics, indicating the highest cross-city consistency, while ElasticNet shows the widest distribution, reflecting greater variability in performance across cities. Regarding extreme values, Tokyo consistently achieves the highest city-level R^2 across all models (reaching 0.990 with TorchMLP), whereas Jakarta consistently records the lowest values (0.604 – 0.637). The relative performance gaps between cities remain stable across models, suggesting that urban spatial structure is a primary determinant of prediction difficulty.

Dhaka, characterized by a dominance of high-density pixels, exhibits the largest MAE and RMSE values (with RMSE up to 154), while Sydney, with a large proportion of near-zero-density pixels, shows extremely high MAPE values (6,800 – 7,900%) but the lowest absolute errors (RMSE ≈ 6.6). These results highlight the inherent limitations of relying on a single error metric in cities with highly skewed population distributions.

Across years, the median R^2 values of all models vary by less than 0.008, indicating strong temporal stability in the feature representations encoded by AEF over the study period.

Fig. 3. Per-city performance of the five models across 2017–2020 under random 80/20 evaluation. Rows correspond to the five evaluation metrics (R^2, Pearson r, MAE, RMSE, MAPE); columns correspond to the four years. Within each subplot, boxplots aggregate the 18 city-level scores per model: the model with the best median is highlighted in red, the worst in green, and the remaining three in blue. Light-blue dots show individual cities, with the top and bottom cities labeled (Tok = Tokyo, Jak = Jakarta, NYC = New York, Dha = Dhaka, Syd = Sydney, LA = Los

Angeles). The legend in each subplot reports the best (▲) and worst (▼) median values. Triangular markers at panel edges denote outliers clipped to the visible range.

However, the random-split results are subject to a systematic risk of overestimation. Urban population density exhibits strong spatial autocorrelation; under random splitting, spatial neighbors of test pixels are highly likely to appear in the training set, allowing models to achieve high R^2 by exploiting local spatial similarity rather than genuine feature-based generalization. The spatial block cross-validation results confirm this concern (Fig. 4): all five models show substantial declines in R^2 under spatial CV (Δ= -0.46 to -0.67), with highly consistent reductions in Pearson's r (Δ= -0.15 to -0.17). This indicates that limited generalization across spatial blocks is a common constraint of the AEF feature system rather than a model-specific issue. The magnitude of degradation is negatively associated with model capacity: TorchMLP shows the smallest decline (ΔR^2= -0.46), whereas the linear model exhibits the largest decline (ElasticNet: −0.67), suggesting that higher-capacity models can extract more spatially transferable feature relationships from the 64-dimensional embeddings. In terms of MAPE, Random Forest shows the largest increase (+103.5), followed by HistGBT (+86.8) and TorchMLP (+78.0), whereas the linear models show smaller increases, indicating that nonlinear models are more prone to overfitting local density patterns under random splitting. For RMSE, linear and tree-based ensemble models show lower values under spatial CV than under random splitting ($\Delta \approx$ -8 to -10), whereas TorchMLP shows an increase (+5.2). This divergence arises because weaker models already struggle to fit high-density pixels under random splitting, and the more balanced density distribution of spatial CV test blocks narrows RMSE; in contrast, TorchMLP can accurately capture high-density cores under random splitting, but its predictions deteriorate once spatial separation is imposed, leading to higher RMSE. In terms of stability, R^2_{std} decreases monotonically with model capacity (Ridge: 0.644; ElasticNet: 0.622; Random Forest: 0.479; HistGBT: 0.435; TorchMLP: 0.278), indicating that stronger models produce more consistent predictions across spatial blocks.

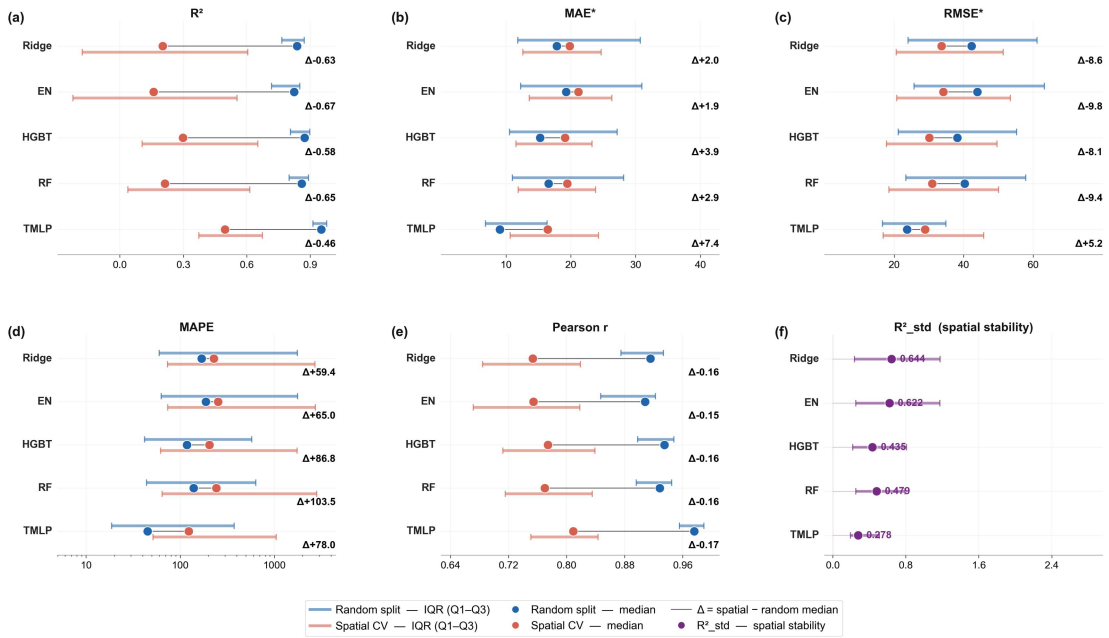

Fig. 4. Performance shift from random 80/20 split to spatial block cross-validation across the five models. Each panel reports one evaluation metric across the 18 cities: (a) R^2, (b) MAE, (c) RMSE, (d) MAPE, (e) Pearson r, and (f) cross-fold standard deviation of R^2 (R^2_std) under spatial CV as a stability indicator. Blue and red bars denote the interquartile range (Q1–Q3) and dots the median under random split and spatial CV, respectively; the right-side annotation Δ gives the median shift (spatial − random). MAE and RMSE are computed in the original population scale (asterisks), MAPE on a logarithmic axis. A negative ΔR^2 indicates degradation under spatial isolation; the magnitude of $|\Delta R^2|$ inversely correlates with model capacity.

To further examine cross-city heterogeneity in performance degradation, Fig. 5 presents a scatter matrix of random-split R^2 versus spatial CV R^2 for the 18 cities. Using the median values of both metrics as reference thresholds, cities are categorized into four spatial generalization types. The HH group (Tokyo, Sydney, São Paulo, Mexico City, Lagos, Istanbul; 6 cities) exhibits above-median performance in both metrics, indicating that AEF embeddings perform consistently under both random and spatial evaluations. The LH group (London, New York, Delhi, Los Angeles; 4 cities) shows below-median random R^2 but relatively high spatial R^2, suggesting that models retain reasonable generalization under spatial separation. The HL group (Cairo, Nairobi, Moscow, Dhaka, Guangzhou; 5 cities) displays above-median random R^2 but substantially lower spatial R^2, representing the most pronounced cases of spatial autocorrelation dependence. The LL group (Jakarta, Buenos Aires, Johannesburg; 3 cities) falls below the median in both metrics, indicating limited overall representational capacity of AEF embeddings for these cities. Notably, Cairo is the only city where the correlation between random and spatial R^2 is not statistically significant ($r = 0.12$, $p = 0.619$), implying a complete decoupling between high random performance and true spatial generalization, and representing an extreme case of spatial autocorrelation dependence. In contrast, Buenos Aires exhibits negative random R^2 values under several models, indicating a failure of AEF embeddings to provide even a baseline level of representation for population distribution in this city.

Fig. 6 presents spatial CV results for four representative cities—Tokyo (HH), London (LH), Cairo (HL), and Jakarta (LL)—showing predicted population density, absolute error, and relative error maps for the best-performing linear, tree-based ensemble, and deep learning models. Under spatial CV, TorchMLP achieves R^2 values of 0.716, 0.839, −13.389, and 0.411 for the four cities, respectively, closely aligning with the previously defined typology. Three common patterns emerge across cities. First, absolute error maps consistently show systematic underestimation in high-density core areas, reflected by dominant blue tones, indicating a general limitation of AEF embeddings under spatial separation. Second, the spatial distribution of errors is highly consistent across model types, suggesting that error patterns are primarily driven by underlying population structure rather than model choice. Third, relative errors in low-density suburban areas are extremely high and must be interpreted in conjunction with absolute errors.

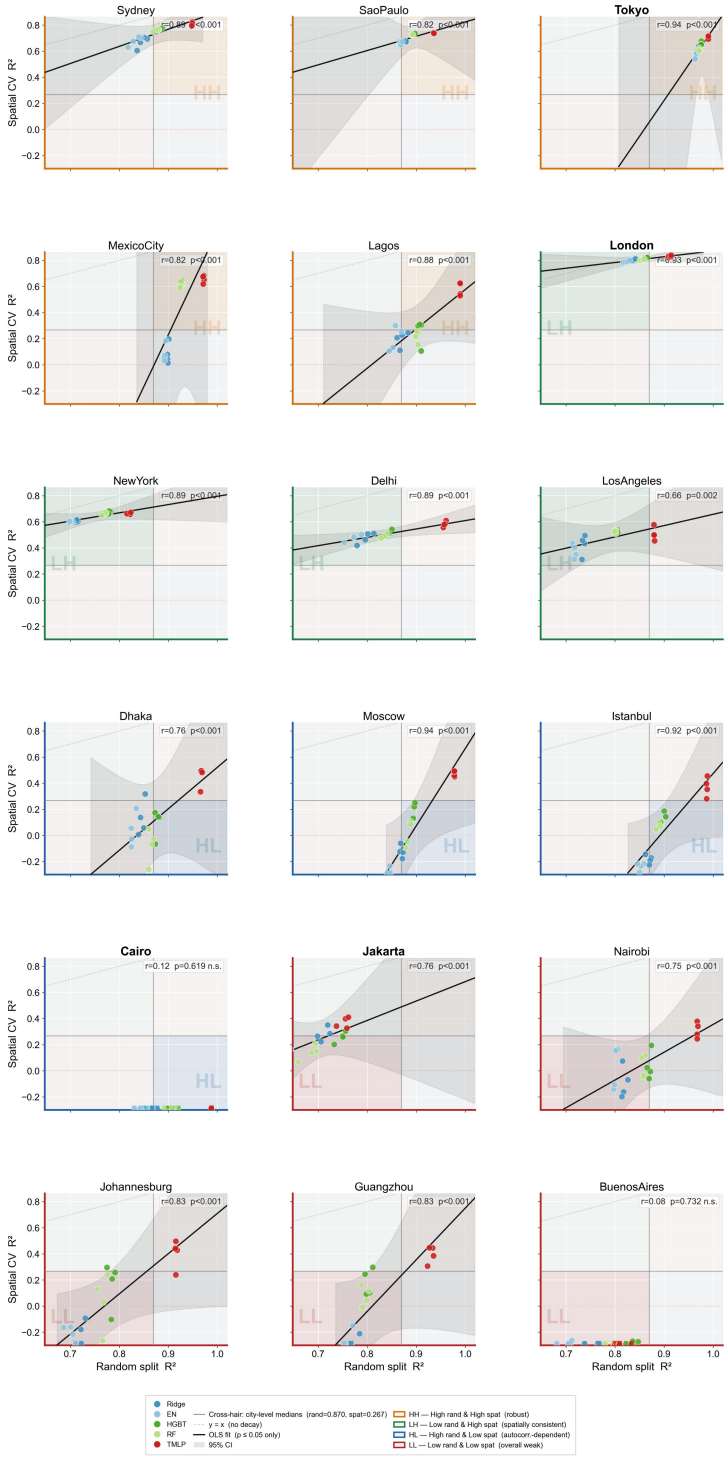

Fig. 5. Per-city scatter of random split R² versus spatial CV R² across the five models, classified into four spatial generalization types. Each panel shows one city, with five colored dots per year-model combination (Ridge, EN, HGBT, RF, TMLP across 2017–2020). The two grey cross-hair lines mark the cross-city medians of random and spatial R² (rand = 0.870, spat = 0.267), partitioning each panel into four quadrants in analogy to Moran's local indicators of spatial association: HH (high–high, robust; orange frame), LH (low–high, spatially consistent; green frame), HL (high–low, autocorrelation-dependent; blue frame), and LL (low–low, overall weak;

red frame). Frame color denotes the quadrant in which the city's centroid falls. Panels are ordered by quadrant (HH → LH → HL → LL) and by descending spatial R^2 within each quadrant. Bold panel titles mark the four representative cities (Tokyo, London, Cairo, Jakarta) used for case-study mapping in Fig. 6. Black solid line and grey band give the OLS fit and 95 % CI, shown only when $p \leq 0.05$; the grey dashed diagonal denotes y = x (no decay between the two evaluation modes). Downward triangles at the bottom edge indicate spatial R^2 values clipped below the y-axis range.

Inter-city differences correspond closely to their respective population spatial structures. In Tokyo, population density follows a concentric gradient, allowing models to learn stable density–feature relationships, with underestimation limited to a small number of extreme high-density pixels. London exhibits a mosaic pattern, where high-density patches are embedded within a medium- to low-density background; dense inner-city areas, particularly historical districts, are consistently underestimated across all models. In Cairo, population is highly concentrated along narrow corridors on both sides of the Nile River, with near-zero values elsewhere; under spatial separation, model performance collapses, with R^2 dropping to extremely low values, rendering both absolute and relative error metrics unreliable. In contrast, Jakarta displays a dispersed and fragmented population pattern without clear centers of concentration; the resulting error maps exhibit fine-grained alternations of over- and underestimation, indicating that models fail to establish stable density–feature relationships.

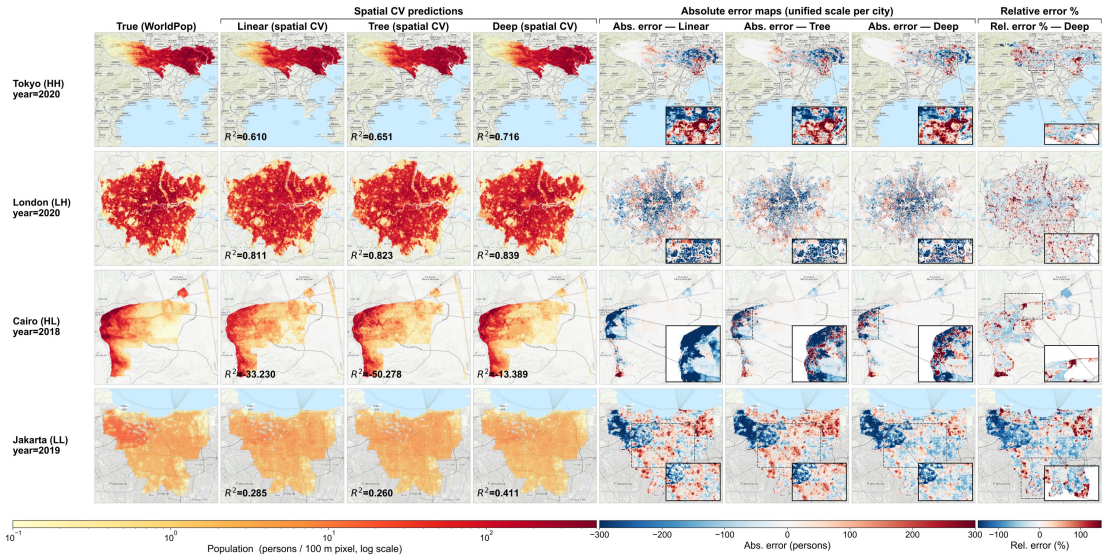

Fig. 6. Spatial CV prediction and error maps for the four representative cities of HH/LH/HL/LL types. Rows correspond to the four representative cities of the spatial generalization types defined in Fig. 5: Tokyo (HH, 2020), London (LH, 2020), Cairo (HL, 2018), and Jakarta (LL, 2019). Columns 1–4 show the WorldPop reference and the spatial CV predictions of the best linear (Ridge or EN), tree-ensemble (HGBT or RF), and deep (TorchMLP) models, all rendered on a shared logarithmic population colormap (persons per 100 m pixel) with the panel-level R^2 annotated. Columns 5–7 show absolute error maps (predicted − reference) of the three models on a unified diverging scale per city — blue indicates underestimation, red overestimation. Column 8 shows the relative error of the deep model. The inset in each error panel

zooms into the urban core to reveal fine-grained patterns obscured at the city scale.

In the training strategy comparison experiment (Fig. 7), S2 (single-city multi-year) exhibits a consistent trade-off pattern across most cities: spatial R^2 improves relative to S1, while random R^2 shows a slight decrease, indicating that the incorporation of historical data provides a beneficial regularization effect. In contrast, when cross-city data are introduced in S3 and S4, both metrics decline simultaneously in the majority of cities, accompanied by a substantial increase in inter-fold variance. This suggests the presence of systematic distributional shifts in AEF feature values across cities, whereby the heterogeneity introduced by cross-city training degrades both predictive accuracy and model stability.

Exceptions are observed in Moscow and Buenos Aires, where S2 results in simultaneous declines in both metrics compared to S1, and in Nairobi, where S3 and S4 fail to outperform the S1 baseline under spatial evaluation. These cases indicate that, for cities characterized by highly corridor-constrained population structures or where the representational capacity of AEF is already limited, incorporating additional temporal or cross-city data may dilute the model's ability to capture local spatial patterns. Overall, temporal expansion (S2) emerges as a relatively robust strategy for performance improvement, whereas the fundamental limitation of cross-city generalization lies in systematic distributional shifts of AEF features across different cities.

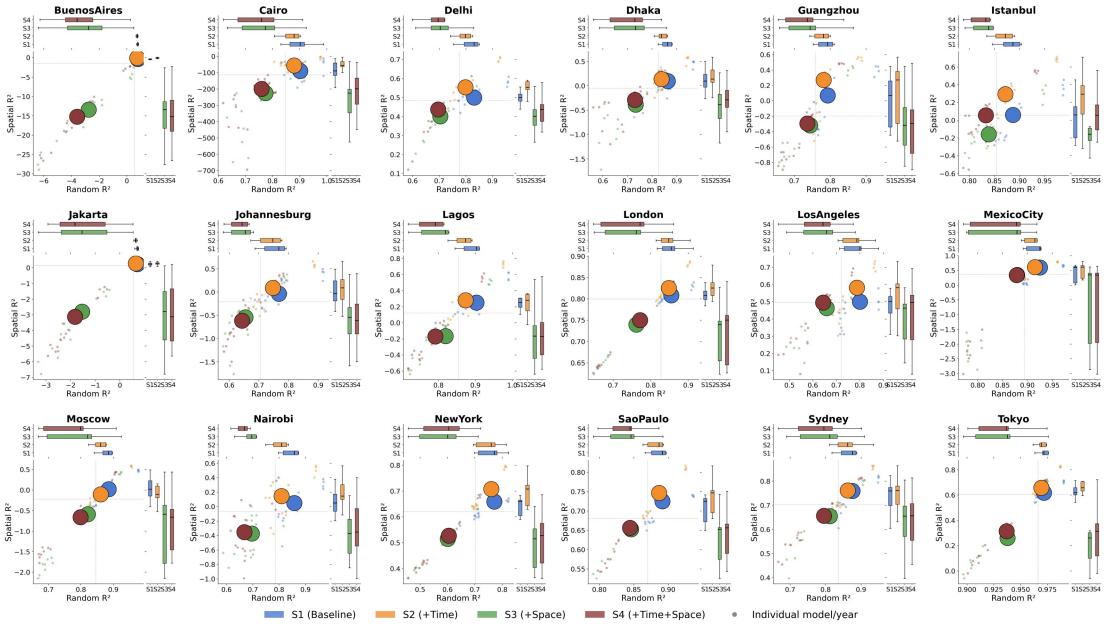

Fig. 7. Per-city comparison of the four training strategies (S1–S4) under random and spatial evaluation. Each panel shows one city. Within each panel, the four large dots represent the four training strategies — S1 (baseline, single city–single year, blue), S2 (+Time, single city–multiple years, orange), S3 (+Space, multiple cities–single year, green), and S4 (+Time+Space, multiple cities–multiple years, dark red) — positioned by their median random R^2 (x-axis) and median spatial R^2 (y-axis) across all model–year combinations. Small grey dots show individual model–year scores under each strategy. Marginal boxplots on the top and right axes summarize the dispersion of the four strategies along the random and spatial dimensions, respectively. Grey cross-hair lines mark the panel-level medians of the two metrics. A strategy point in the

upper-right region indicates joint gains in both random fitting and spatial generalization; movement toward the upper-left or lower-right reflects a trade-off between the two. Note that x- and y-axis ranges are panel-specific to preserve resolution within each city; absolute values are not directly comparable across panels.

5.2. Scale Adaptation

Using S1 as the baseline, this section evaluates the performance of five model types across four spatial resolutions from 100 m to 1000 m (Fig. 8). Linear and tree-based models (Ridge, ElasticNet, Random Forest, HistGBT) exhibit largely stable R^2 values during scale aggregation, with differences across resolutions typically within 0.05, indicating strong robustness to spatial scaling. In contrast, TorchMLP shows a markedly different pattern: under random splitting, performance slightly improves with coarser resolution, whereas under spatial CV, performance degrades substantially. In many cities (e.g., Buenos Aires, Cairo, Johannesburg, Moscow, New York, São Paulo, Tokyo), spatial CV R2R^2R2 drops into negative values at 500 m or 1000 m. This degradation is driven by the scale dependence of sample size. At 100 m resolution, the number of valid pixels per city ranges from 2×10^4 to 7×10^5; after aggregation to 1000 m, the smallest cities (e.g., Buenos Aires) retain only about 240 pixels. Such sample scarcity directly limits the fitting capacity of high-capacity neural networks, whereas linear and tree-based models, with fewer parameters and stronger inductive biases, are less affected.

The results of TorchMLP-Adaptive further support this explanation (Fig. 8). Comparing the TMLP-A row with the original TorchMLP row in the heatmap shows that, at 500 m and 1000 m resolutions in most cities, the adaptive version maintains performance comparable to linear and tree-based models (ΔR^2 close to zero or slightly negative), whereas the original TorchMLP exhibits widespread and pronounced performance degradation ($\Delta R^2 <- 0.3$). This comparison indicates that the scale sensitivity of TorchMLP primarily arises from a mismatch between model capacity and sample size, rather than an inherent limitation of neural network approaches at coarser scales. When model capacity is adaptively adjusted to match the available sample size, neural networks can retain performance comparable to that of linear and tree-based models.

The optimal spatial resolution varies substantially across cities (Fig. 9, annotated as Best R / Best S). Cities with larger pixel counts, such as Guangzhou, Sydney, and Cairo, exhibit performance peaks or plateaus at intermediate resolutions ($200 – 500$ m), reflecting gains from improved signal-to-noise ratios at moderate aggregation levels. In contrast, cities with smaller sample sizes, such as Buenos Aires and Jakarta, consistently achieve their best performance at 100 m, as further aggregation reduces the effective sample size and leads to rapid performance deterioration.

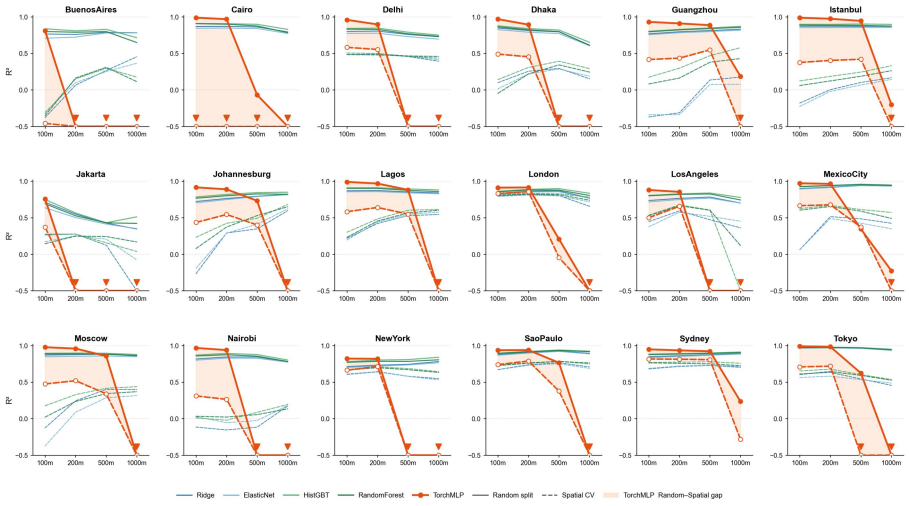

Fig. 8. Scale dependence of five models (100–1000 m) across cities under random and spatial evaluation. Solid vs. dashed lines denote random split and spatial CV R^2. TorchMLP (thick orange) shows a widening random–spatial gap at coarser resolutions. Triangles indicate clipped spatial CV R^2 (< -0.5). Performance degradation is mainly due to reduced sample size (~10^4 – 10^5 to ~240), not scale (Fig. 9).

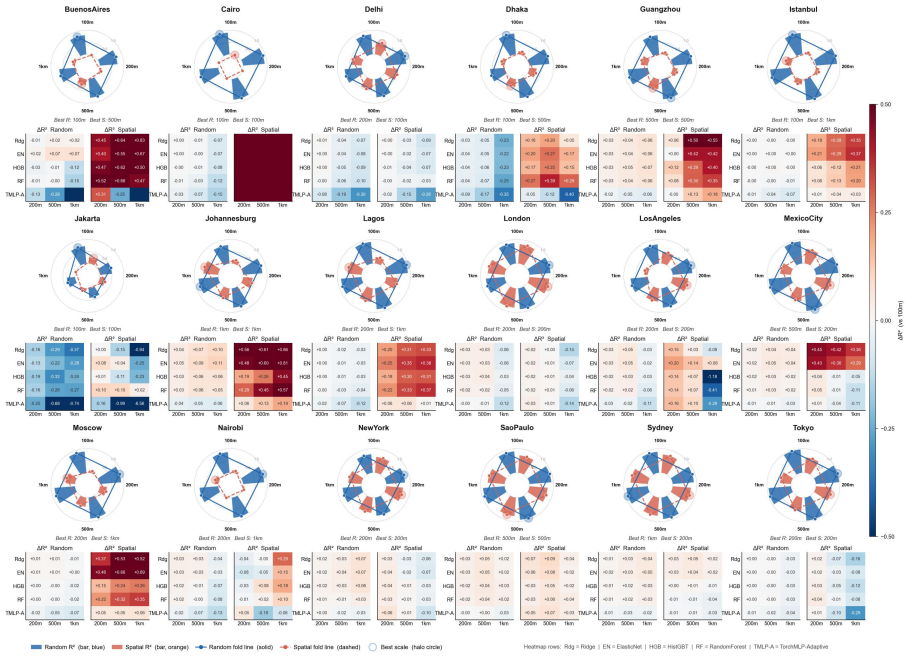

Fig. 9. Per-city scale-resolved performance and capacity-adaptive verification across the five models plus TorchMLP-Adaptive. Each city is shown by a paired radar–heatmap module. The radar chart (top) plots the four spatial resolutions as radial axes (100 m, 200 m, 500 m, 1000 m); blue and orange shaded polygons depict random R^2 and spatial R^2 of all six models, with solid and dashed connecting lines tracing the across-resolution trajectory of each model. The halo circle marks the resolution at which the city achieves its best random ("Best R") or best spatial ("Best S") R^2, annotated above the heatmap. The heatmap (bottom) reports ΔR^2 relative to the 100 m baseline at the three coarser scales (200 m, 500 m, 1000 m) for both random and spatial evaluation; rows

are the six models — Ridge, EN, HGBT, RF, TorchMLP, and TorchMLP-Adaptive (TMLP-A). Comparing the TMLP and TMLP-A rows isolates the contribution of capacity adaptation: TMLP-A largely eliminates the deep-blue degradation blocks seen in the original TMLP row at 500 m and 1000 m, confirming that the coarse-scale failure of TorchMLP stems from a mismatch between fixed model capacity and shrinking sample size rather than an intrinsic limitation of the deep-learning approach.

5.3. Feature Importance Analysis

This section examines the internal structure of the 64-dimensional AEF embeddings from two complementary perspectives—PCA and SHAP. All analyses are conducted under both S1 and S2 strategies, with all five model types retrained, to ensure that the findings are not dependent on a specific strategy or model.

Projecting the AEF embeddings onto $k \in \{8,16,24,32,38,48,56,64\}$ principal components and evaluating relative to the full $k=64$ baseline reveals markedly different saturation behaviors for random and spatial R^2 (Fig. 10a). Across the five models, the median value of k^* required to reach 99% saturation under random evaluation is 38 (range: $32-64$), whereas under spatial evaluation the median increases to 64 (range: 38–64). For Ridge, Random Forest, and TorchMLP, the spatial performance gap remains significantly positive even at k=56. These results indicate that within-city random fitting requires only approximately half of the principal components, whereas generalization across spatial blocks approaches the full 64-dimensional representation. At k=32, the cumulative explained variance already reaches 98.5%, yet the spatial R^2 of TorchMLP recovers only 84.7% of the baseline, revealing a clear decoupling between variance explained and predictive performance. This suggests that the key signals supporting cross-spatial generalization are distributed within low-variance principal components.

Aggregating the top-8 SHAP contributions from each model and performing cross-model voting (Fig. 10b) reveals a backbone–long-tail structure. Dimensions 16 and 36 are consistently identified by all five models, dimensions {32, 35, 56} by four models, and {1, 7} by three models. However, the top 10 dimensions collectively account for only about 30% of total SHAP contribution, and nearly the top 53 dimensions are required to reach 90% coverage. The apparent discrepancy between PCA ("k=32 is sufficient") and SHAP ("top-53 needed for 90% contribution") is not contradictory. PCA represents orthogonal linear combinations of all 64 dimensions, allowing low-dimensional projections to reconstruct most variance, whereas SHAP quantifies the conditional contribution of original features to prediction, which inherently follows a more long-tailed distribution.

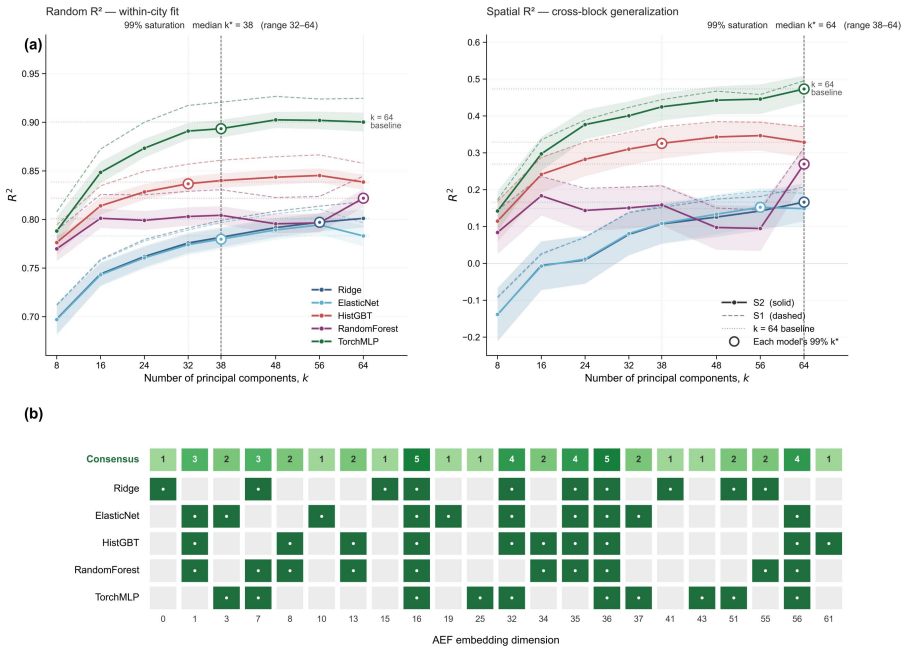

Fig. 10. Effective dimensionality of AEF embeddings: PCA saturation (a) and cross-model SHAP consensus (b). (a) R^2 as a function of principal component number k, under random split (left) and spatial CV (right). Solid lines: S2; dashed: S1; shaded bands: ±1 SD across 18 city-year combinations. Open circles mark each model's k* (smallest k reaching 99 % of the k = 64 baseline); cross-model median k* annotated above. (b) Top-8 SHAP-contributing dimensions per model; the "Consensus" row counts endorsing models per dimension (only dimensions endorsed by ≥1 model shown). Dim 16 and 36 attain full 5-model consensus.

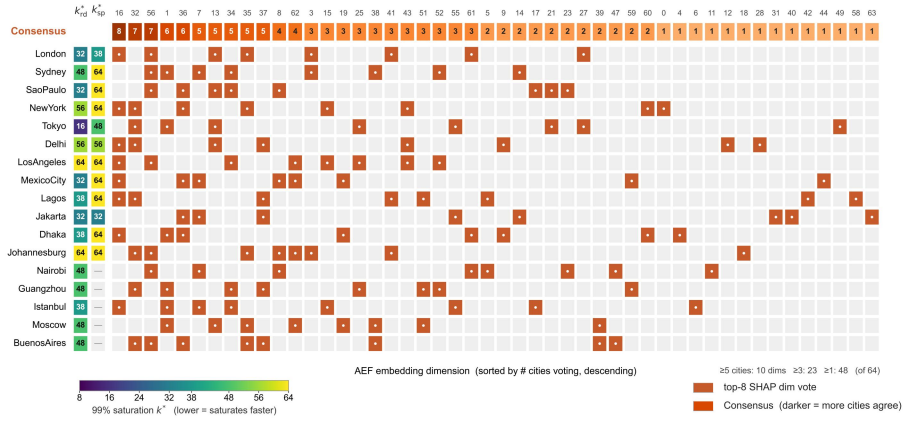

Fig. 11. Per-city decomposition of top-8 SHAP-contributing AEF dimensions across 17 cities. Filled cells indicate that the corresponding AEF dimension ranks within a city's top 8 by mean absolute SHAP value (averaged over the five models). Cities are sorted by descending consensus strength; dimensions are sorted by the number of cities endorsing them (top "Consensus" row, darker = more cities). The two leftmost columns report each city's k_rd and k_sp (99 % saturation under random and spatial CV); "—" denotes failure to saturate at k = 64. Cairo is excluded due to spatial CV collapse (n = 17).

City-level decomposition of the top-8 SHAP voting results (Fig. 11; n=17 cities, with Cairo excluded because spatial CV failed completely) shows that the overall backbone structure is partly stable but also partly disrupted at the city level. The stable pattern is evident in the following dimensions: dim 16 ranks first with consensus across 8 cities, dim 1 and dim 32 each appear in 7 cities, dim 36 and dim 7 each appear in 6 cities, and dims {13, 34, 35, 37, 8} each appear in 5 cities, yielding 10 dimensions with consensus across at least 5 cities. This set largely overlaps with the cross-model backbone set {1, 7, 16, 32, 35, 36, 56}, indicating that these dimensions are consistently identified as top-8 contributors within most cities rather than being artifacts driven by a few dominant cases.

Substantial heterogeneity is also observed. The consensus set expands to 23 dimensions when the threshold is relaxed to at least 3 cities, and to 48 dimensions when dimensions appearing in at least 1 city are included, covering 75% of all AEF dimensions and revealing a pronounced long-tail pattern. Several cities exhibit top-8 dimensions that deviate markedly from the overall backbone: Jakarta uniquely selects dims {40, 42, 44, 63}, Delhi uniquely selects dims {12, 18, 28}, Mexico City uniquely selects dims {31, 59}, and New York uniquely selects dims {39, 47}. These dimensions have low overall importance but enter the top-8 list in specific cities. The city-level distribution of k^* further confirms this heterogeneity: k^*_{rd} spans the full range from 16 in Tokyo to 64 in Los Angeles and Johannesburg, while for k^*_{sp}, five cities—Nairobi, Guangzhou, Istanbul, Moscow, and Buenos Aires—fail to reach 99% saturation before k=64 and are marked as "—", whereas Jakarta reaches saturation at only k^*_{sp}=32. These results indicate fundamental differences in how individual cities depend on the dimensional structure of AEF embeddings.

5.4. Effects of Auxiliary Variables

Fig. 12 presents the marginal contributions of four categories of auxiliary variables to population spatial modeling performance across the five models under S1–S4 training strategies. Panel (a) shows absolute R^2 values under different feature combinations, with the dashed line indicating the AEF-only baseline (C0); Panel (b) summarizes the median ΔR^2 for C1-C5 under the S1 strategy using a heatmap. Cairo is excluded from aggregated statistics due to complete model collapse under spatial CV (with R^2 as low as -694) and is instead marked separately with an asterisk.

The introduction of auxiliary variables consistently yields positive gains for linear and tree-based models (Ridge, ElasticNet, HistGBT, Random Forest) under both random and spatial evaluation modes. In contrast, TorchMLP exhibits a selective response: the inclusion of POI features (+POI) and the full combination (+All) results in negative gains under random evaluation (ΔR^2= -0.054 and -0.033, respectively), and also under spatial evaluation (+POI: -0.125; +All: -0.075). This indicates that the addition of high-dimensional socio-functional features may introduce noise when AEF embeddings already provide strong representational capacity, thereby degrading generalization stability.

Auxiliary variables contribute more substantially to spatial generalization than to random fitting. In Panel (b), ΔR^2 values in the spatial column are consistently larger than those in the random column, with +Spatial achieving the largest gains across all models (Ridge: +0.111; ElasticNet: +0.101; HistGBT: +0.299; Random Forest: +0.319; TorchMLP: +0.074). For several models, the

inclusion of spatial structure variables elevates spatial R^2 from negative or near-zero values to clearly positive levels—an improvement not attainable under the C0 baseline.

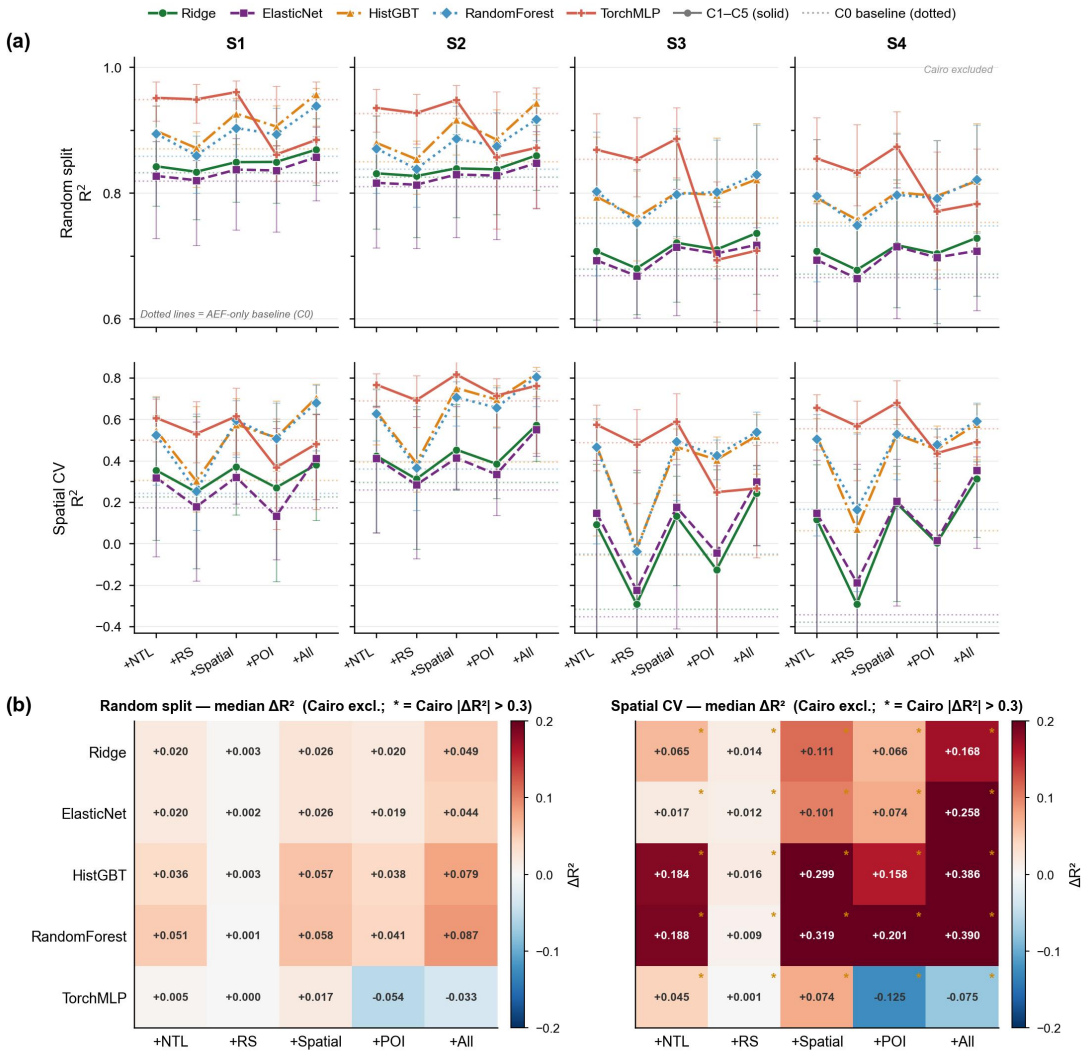

Fig. 12. Marginal gains of four auxiliary factor groups over the AEF baseline across the four training strategies. (a) Absolute R^2 of the five models under five auxiliary factor combinations (C1: +NTL; C2: +RS; C3: +Spatial; C4: +POI; C5: +All), evaluated under random split (top) and spatial CV (bottom) for the four training strategies S1–S4. Dotted horizontal lines mark each model's AEF-only baseline (C0); error bars give ±1 SD across cities. (b) Median ΔR^2 relative to C0 under S1, summarized as heatmaps. Cairo is excluded from medians due to spatial CV collapse (R^2 as low as −694); cells where Cairo's individual $|\Delta R^2|$ exceeds 0.3 are flagged with an asterisk.

From the perspective of training strategies, the gain patterns of auxiliary variables under S1 and S2 are broadly consistent, although S2 exhibits lower inter-fold variance under spatial evaluation, suggesting that the incorporation of historical data provides an effective regularization effect. Under S3 and S4, spatial R^2 values under the C0 baseline drop substantially into negative ranges; while the addition of +Spatial partially restores spatial generalization, the improvement is

insufficient to offset the degradation caused by cross-city feature distribution shifts. This indicates that auxiliary variables cannot fundamentally resolve the bottleneck of cross-city generalization.

Overall, the combination of the S2 training strategy, +Spatial auxiliary variables, and the TorchMLP model constitutes the optimal modeling configuration. Although the relative gain of TorchMLP from +Spatial (+0.074) is smaller than that of tree-based models (+0.299 to +0.319), it achieves the highest absolute spatial R^2, maintaining strong random performance while delivering the best cross-spatial generalization. This configuration is adopted as the baseline for the global 1 km and 18-city 100 m population mapping applications presented in Section 6.

6. Mapping Applications

6.1. Global 1 km Mapping

Extending the optimal modeling configuration identified in Section 5.4 (TorchMLP + S2 multi-year training strategy + C3 spatial structure variables) from the city scale to global-scale mapping requires the adoption of a unified spatial reference framework based on EPSG:4326 latitude–longitude grids, along with global preprocessing of AEF embeddings and spatial structure variables. For AEF data, annual embeddings for 2017–2024 covering global land areas ($-60°$ to $85°$ latitude) are exported from the GEE platform in longitudinal tiles, with ocean pixels masked using MODIS land cover products and resampled to 1 km resolution. For spatial structure variables, urban center pixels are identified based on the GHSL global human settlement model using 2020 as a reference, and nearest-neighbor distances are computed; combined with global water body and coastline data from OpenStreetMap, this yields two additional distance layers (water and coastline), resulting in three static distance features. In addition, global historical OSM PBF data are used to derive annual road network density at three hierarchical levels (arterial, collector, local), along with total density, forming four dynamic features. In total, seven bands of global spatial structure variables are constructed. Model training follows the optimal configuration described in Section 5.4, using global land pixels from $2017-2020$ together with WorldPop population labels resampled to 1 km resolution.

To evaluate global-scale generalization performance, back-testing is conducted using WorldPop data for 2017–2020 as reference. Tile-level R^2 is computed in log space, and tiles with fewer than 10,000 valid pixels are excluded. A population-weighted average R^2 is then calculated using WorldPop population totals as weights. The weighted R^2 values are 0.846 (2017), 0.849 (2018), 0.849 (2019), and 0.847 (2020), with effective population coverage reaching 98.4% in all years. Tiles with negative R^2 are primarily located in extremely low-density regions such as the Sahara and Arctic, with negligible impact on overall population coverage. These results indicate that the model achieves stable predictive performance across populated regions globally, supporting its use for extrapolation. Based on this, forward prediction is applied to AEF embeddings and corresponding spatial structure variables for $2021-2024$, generating annual global population maps at 1 km resolution without retraining (Fig. 13).

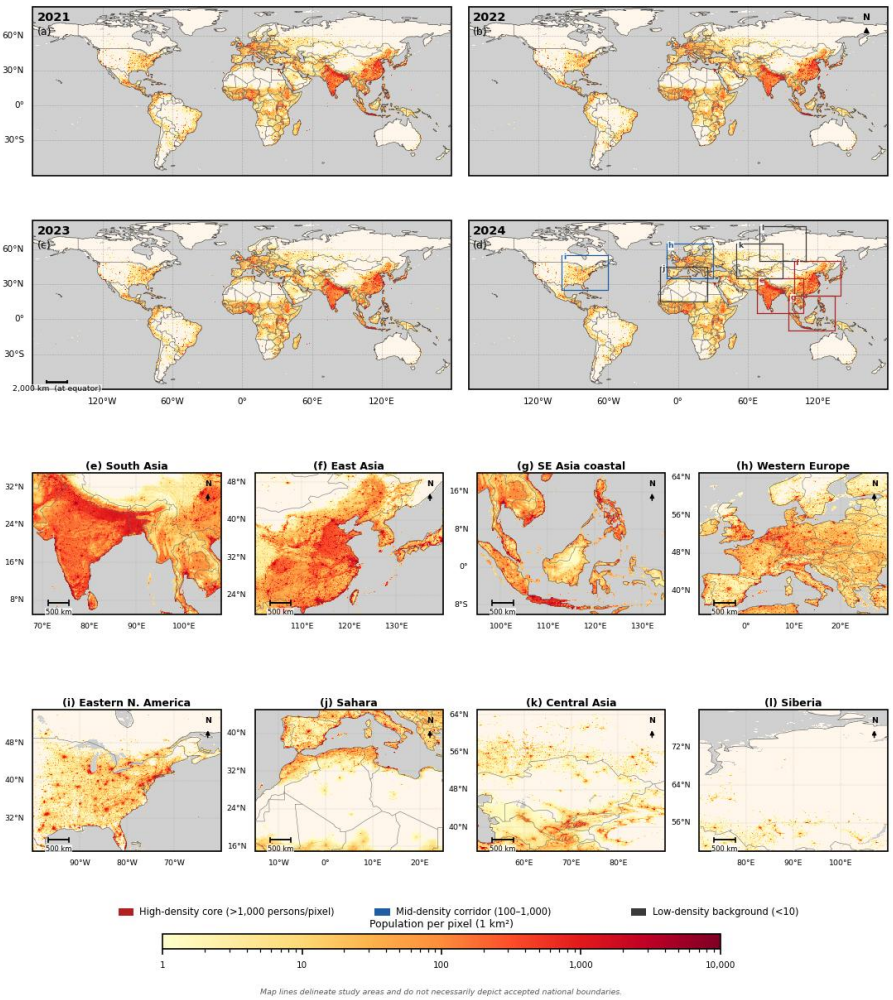

Fig. 13. Global 1 km annual population spatial distribution for 2021–2024 derived from the optimal modelling combination (TorchMLP + S2 + C3). Panels (a–d) show global density maps on a shared logarithmic colormap (persons per 1 km² pixel); panel (d) overlays colored bounding rectangles marking eight representative regional insets classified by density tier — red: high-density core (>1,000 persons/pixel), blue: mid-density corridor (100–1,000), dark grey: low-density background (<10). Panels (e–l) show 2024 regional insets at a uniform 40°×30° extent, each with a 500 km scale bar and north arrow. The shared color bar uses a log scale (1–10,000); map lines do not necessarily depict accepted national boundaries.

The resulting maps exhibit strong agreement with known global population patterns: the highest-density cores (>1000 persons per pixel) are observed in South Asia (India, Bangladesh), East Asia (eastern China, Japan), and Southeast Asian coastal urban agglomerations; continuous medium-to-high density corridors (100–1000 persons per pixel) are evident in Western Europe and the eastern United States; and extremely low-density regions (<10 persons per pixel) persist across the Sahara, Central Asian arid zones, and Siberia. The spatial structure of the maps remains highly consistent across the four years, with minimal visible interannual differences, indicating that no

systematic distributional drift is introduced during temporal extrapolation. This further supports the feasibility of the proposed framework for filling post-2020 temporal gaps in WorldPop. On this basis, interannual difference analysis is conducted for the 2021 – 2024 products to further characterize the dynamic evolution of global population spatial patterns over the study period.

Using 2021 as the baseline, the global population spatial changes captured by the model for 2022–2024 are shown in Fig. 14. Overall, most land areas exhibit light colors (indicating near-zero change rates), suggesting that population spatial patterns remain largely stable across years and that no systematic drift is introduced during temporal extrapolation. In contrast, pronounced red and blue signals are concentrated in specific hotspot regions, which constitute the primary focus of the following analysis.

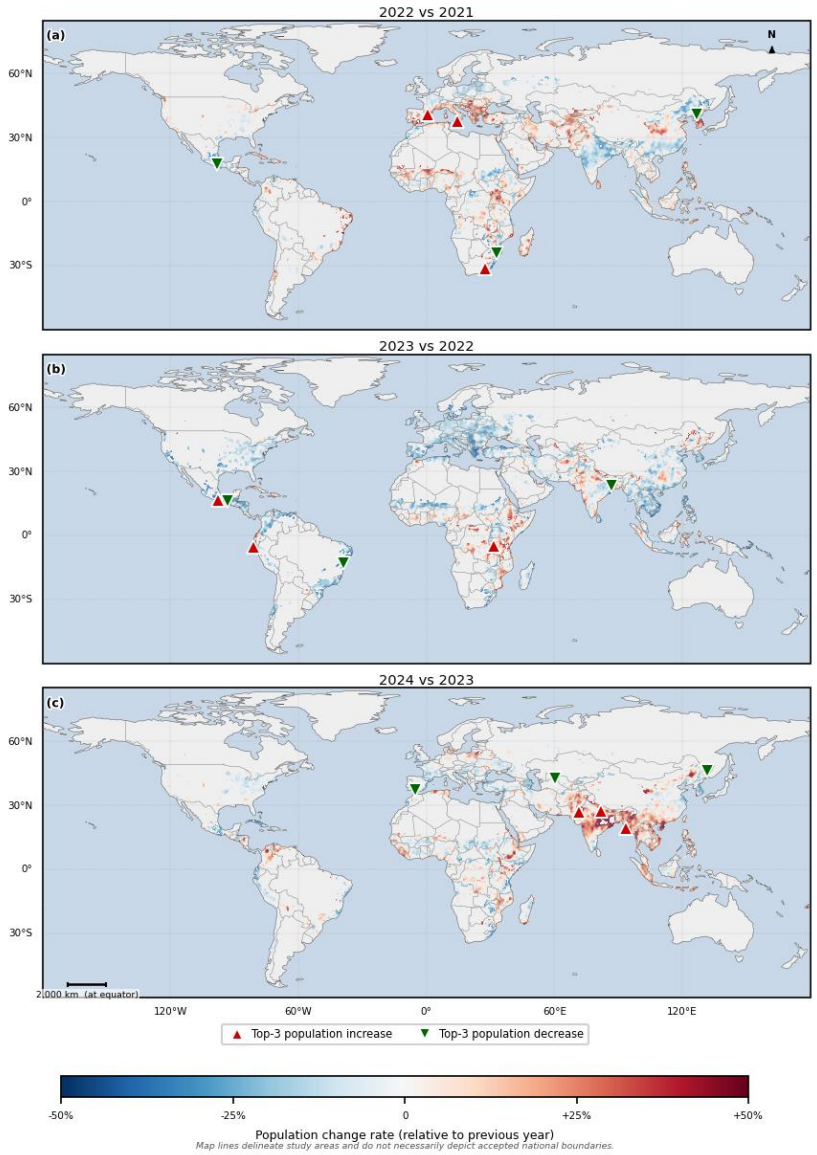

Fig. 14. Year-over-year global population change rate maps for 2022–2024 relative to the preceding year, derived from the 2021–2024 mapping products. Panels (a–c) show pixel-level

population change rate (relative to the previous year) on a diverging colormap (blue: decrease; red: increase; range −50% to +50%); red upward triangles and green downward triangles mark the top-3 increase and decrease hotspots per panel, respectively. The shared color bar (bottom) indicates population change rate relative to the previous year; scale bar (2,000 km at equator) and north arrow are shown in panel (c) and (a), respectively. Map lines do not necessarily depict accepted national boundaries.

In 2022, growth hotspots are concentrated in the transitional zone between the Middle East and North Africa (notably around Turkey and Syria), consistent with ongoing refugee resettlement and cross-border population movements in the region. Declining signals are observed in East Asia (Japan and the Korean Peninsula) and Central America; the former aligns with long-term structural population decline, while the latter corresponds to sustained northward migration leading to local population outflows. Extensive blue patterns across Russia and Siberia may reflect population mobilization and urban – rural redistribution following the outbreak of the Russia – Ukraine conflict.

In 2023, growth signals shift toward northwestern South America (along the Colombia – Peru Amazon border), consistent with expanding agricultural frontiers and increasing informal settlement activities. The decline pattern in Central America persists, while new decline hotspots emerge in northwestern South Asia (border regions of Pakistan and India), potentially linked to delayed population displacement following the severe flooding in Pakistan in 2022. North America shows widespread decline signals during the same period, possibly associated with the post-pandemic reversal of population flows back to urban areas following earlier suburbanization trends.

In 2024, South Asia and the Middle East form the most prominent growth hotspots during the study period, with extensive red signals across Pakistan, northwestern India, and Gulf countries. These patterns may reflect the combined effects of population displacement associated with the escalation of the Gaza conflict and continued labor inflows into Gulf economies. Localized decline signals appear in parts of Central and East Asia, potentially related to secondary migration of populations leaving Russia. Eastern coastal China exhibits noticeable growth signals in 2024, which may be associated with the post-pandemic recovery of built-environment activity and accelerated urbanization, although the underlying drivers require further investigation.

Overall, the spatiotemporal patterns from 2022 to 2024 indicate that the Middle East–South Asia arc represents the most dynamically evolving region in global population redistribution during this period, shaped by a combination of geopolitical conflicts, climate-related disasters, and economic attraction. It should be noted that these inferred changes are derived from interannual variations in AEF surface embeddings; direct validation against census or mobile data remains to be conducted. Therefore, the above interpretations should be regarded as exploratory rather than definitive conclusions.

6.2. 100 m Mapping for Major Cities

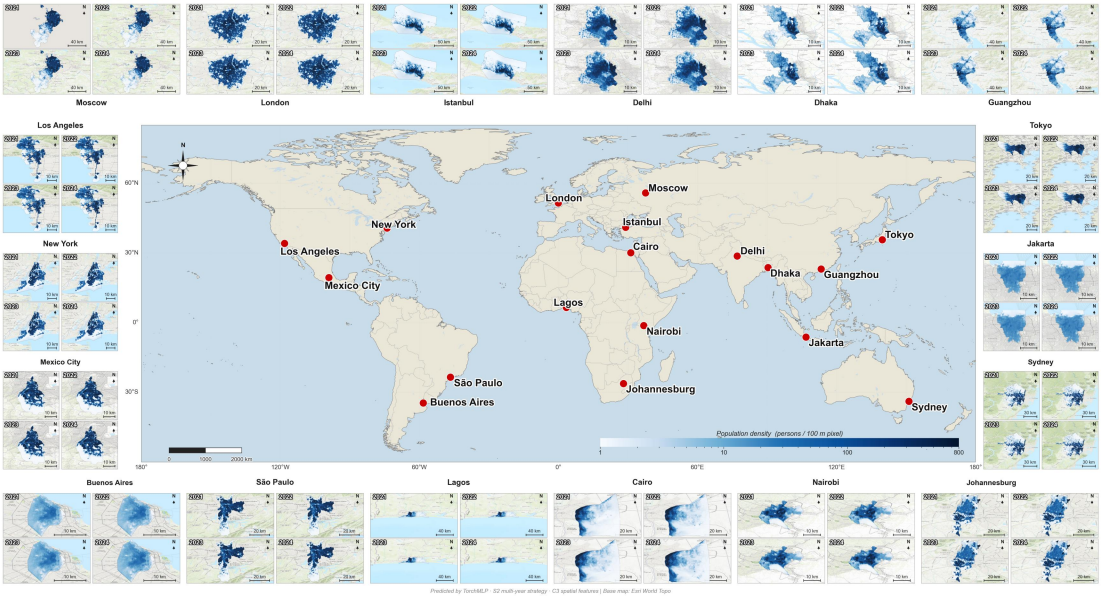

Fig. 15. Annual 100 m population density mapping for 18 study cities over 2021–2024, derived from the optimal modelling combination (TorchMLP + S2 + C3). The central panel shows the geographic distribution of all 18 cities on a world basemap; surrounding insets display the four-year prediction sequence (2021–2024, left to right) for each city, rendered on a shared linear colormap (population density in persons per 100 m pixel, range 10–600). Each inset includes a scale bar and north arrow. Predicted by TorchMLP + S2 multi-year strategy + C3 spatial features; basemap: Esri World Topo.

The optimal modeling configuration identified in Section 5.4 (TorchMLP + S2 multi-year training + C3 spatial structure variables) is transferred to the extrapolative mapping task with two adjustments: the training samples are expanded to include AEF embeddings and corresponding WorldPop labels for all 18 cities from 2017 to 2020; during inference, forward prediction is directly applied to annual AEF embeddings and synchronized C3 variables for 2021–2024 without retraining. This workflow does not rely on high-barrier data sources such as census data, mobile phone records, or commercial POIs. For any city, once AEF embeddings, administrative boundaries, and OSM-derived C3 features are available, the trained weights can be reused for population mapping, demonstrating a lightweight and transferable paradigm enabled by the analysis-ready nature of AEF. Based on this workflow, 72 inference runs are completed for the 18 cities over 2021–2024, with the mapping results shown in Fig. 15. The model produces population estimates that are spatially consistent with actual urban structures across cities with different geographic locations, size classes, and development stages, while the four annual predictions show smooth temporal transitions in most cities, supporting the feasibility of the optimal configuration for both interannual extrapolation and cross-regional application.

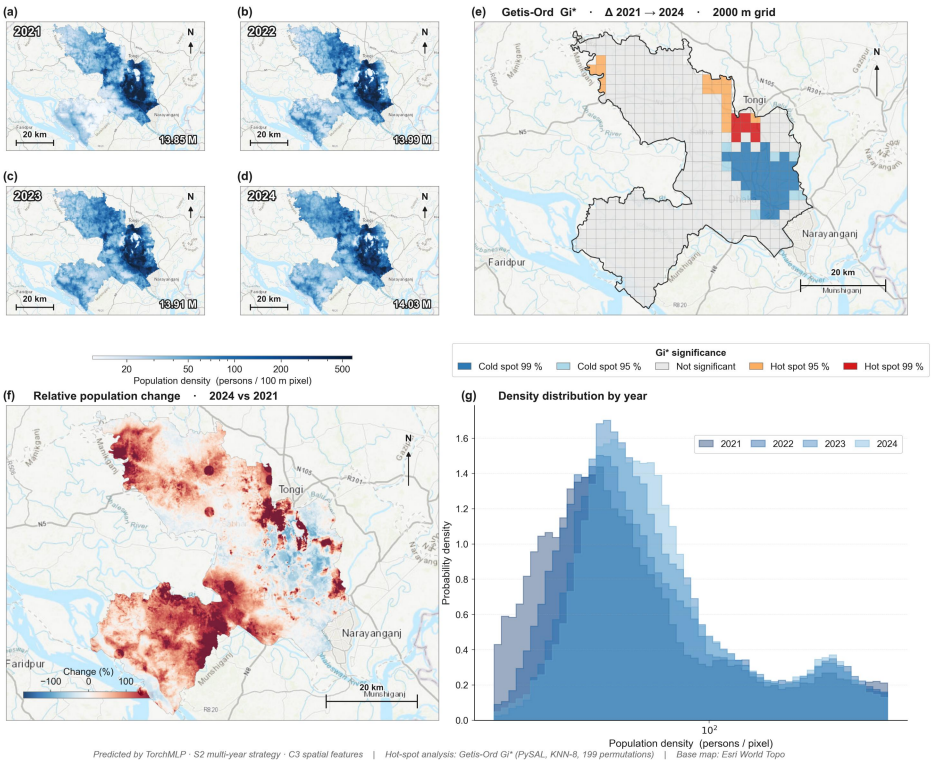

Fig. 16. Population spatial dynamics of Dhaka, 2021–2024, derived from the optimal modelling combination (TorchMLP + S2 + C3). Panels (a–d) show annual 100 m population density estimates for 2021–2024 on a shared logarithmic colormap (persons per 100 m pixel); total city population (M) is annotated at the lower-right of each panel. Panel (e) shows Getis-Ord Gi* hotspot analysis of cumulative population change (2021→2024) at 2 km grid resolution (PySAL, KNN-8, 199 permutations), with significance levels classified into five tiers (Cold/Hot spot at 99% and 95%, and Not significant). Panel (f) maps pixel-level relative population change rate between 2024 and 2021 on a diverging colormap (range −100% to +100%). Panel (g) shows log-scale probability density histograms of population density for all four years, with overlapping distributions indicating temporal stability. Basemap: Esri World Topo.

To further assess the ability of the optimal model to characterize fine-scale urban population dynamics, Dhaka is selected as a representative case for detailed analysis. Dhaka is one of the most densely populated megacities in the world, with an estimated urban population of 13.85 million in 2021. Its spatial structure is characterized by a dense urban core, river constraints, and newly developed peripheral settlements, posing substantial challenges for both local prediction accuracy and interannual consistency. Fig. 16(a)–(d) shows annual 100 m population density estimates for 2021–2024. The total urban population increases from 13.85 million to 14.03 million, corresponding to a compound annual growth rate of +0.42% per year, while the four annual predictions maintain smooth temporal transitions. Fig. 16(e) presents a Getis–Ord Gi analysis of population change from 2021 to 2024 at the 2 km grid scale (PySAL, eight nearest neighbors, 199 conditional randomizations), identifying a significant growth hotspot northeast of Tongi (99% hotspot), a surrounding contiguous 95% hotspot zone, and extensive significant cold spots along the southeastern river corridor (99% cold spots). Fig. 16(f) shows pixel-level population change rates in 2024 relative to 2021, revealing a pattern of “central contraction and peripheral

expansion.” Together with Fig. 16(e), this provides methodological cross-validation: the former identifies statistically significant zones of change, whereas the latter quantifies the magnitude and direction of change. Fig. 16(g) presents log-density histograms of population density across the four years; the four curves largely overlap, with their peaks stably anchored around the 10210^2102 level, indicating that the model does not introduce systematic distributional drift over time. These results demonstrate that the optimal modeling configuration can effectively capture interannual spatial redistribution signals, supporting its applicability for fine-scale urban population monitoring and spatiotemporal dynamics analysis.

7. Discussion

7.1. Main Findings

- Feasibility. The 64-dimensional AEF embeddings demonstrate strong explanatory capacity for population spatial patterns under random evaluation, with TorchMLP achieving mean R^2 values of $0.952-0.956$. This confirms that multi-source remote sensing information, when encoded within a unified representation framework, can effectively capture built-environment characteristics and population agglomeration patterns. However, spatial block cross-validation reveals substantial spatial autocorrelation dependence across all models, with ΔR^2 ranging from -0.46 to -0.67. This indicates that the population prediction ability of AEF embeddings depends to a considerable extent on the spatial proximity between training and test samples, rather than on genuine generalization to urban population structure. This finding highlights the systematic risk of using random-split results to assess the generalization capacity of mapping models and has broader methodological implications for remote sensing modeling practices.

- Training strategies. Single-city multi-year training (S2) provides the most robust pathway for improving spatial generalization. The inclusion of historical data acts as an effective regularizer, improving spatial CV performance without substantially compromising random accuracy. In contrast, the failure of cross-city strategies (S3 and S4) reveals systematic distributional shifts in AEF features across cities. Although AEF is designed for global consistency, its embeddings still contain pronounced region-specific information at the urban scale, creating fundamental obstacles for cross-city knowledge transfer. This finding challenges the assumption of “global universality” often associated with foundation representation models and suggests that future research should further explore domain adaptation strategies.

- Scale adaptability. Linear and tree-based ensemble models exhibit strong robustness across spatial scales, whereas the degradation of TorchMLP at coarser resolutions is shown to arise from a mismatch between model capacity and sample size rather than from an inherent limitation of the method itself. The TorchMLP-Adaptive experiment effectively separates the contributions of these two factors. This result provides practical guidance for model selection under different application scenarios: TorchMLP is preferable for fine-resolution tasks with sufficient pixel samples, whereas capacity-adaptive configurations or tree-based ensemble models are more robust choices for coarser-scale or small-area applications.

- Dimensional structure. PCA and SHAP analyses jointly reveal a backbone–long-tail structure in AEF embeddings. A small number of dimensions, such as dim 16 and dim 36, form stable cross-city and cross-model consensus, yet spatial generalization requires a near-complete 64-dimensional representation, and the dependency structure varies fundamentally across cities. This finding indicates that AEF is not redundantly designed: each dimension may carry non-substitutable local signals under different cities and tasks. Therefore, dimensional compression for computational efficiency is likely to come at the cost of reduced spatial generalization.

- Auxiliary variables. Spatial structure variables (C3) provide substantially greater benefits for spatial generalization than other auxiliary factor types. Their inclusion raises spatial R^2 from negative or near-zero values to positive ranges for multiple models—an improvement that cannot be achieved using AEF alone. This indirectly confirms the inherently indirect coupling between population distribution and remote sensing observations: surface appearance information encoded by AEF must be coupled with explicit urban spatial organization to establish stable population density mappings. The negative gains from high-dimensional POI features in TorchMLP further reveal the generalization risks of deep models when exposed to sparse, high-dimensional socio-functional data, providing important negative evidence for auxiliary feature selection.

7.2. Limitations

This study has three main limitations.

- **Label quality.** The training labels used in this study are derived from WorldPop estimates rather than directly observed population data. Therefore, the upper bound of the learned mapping is constrained by the varying quality of WorldPop in regions with weak census systems, and evaluation metrics in some areas may not fully reflect the actual reliability of the mapping products. In addition, the native resolution of WorldPop (~930 m) is coarser than the modeling scale used in this study (100 m). Resampling WorldPop to a 100 m grid is therefore effectively an interpolation toward finer scales, and the resulting pixel-level labels inevitably contain spatial smoothing effects, further constraining the credible upper bound of accuracy assessment.

- **Model generalization.** Although the combination of the S2 strategy and C3 variables substantially improves spatial generalization, predictions still nearly collapse under spatial separation in cities with strongly corridor-constrained population structures, such as Cairo. This indicates clear generalization boundaries when the model encounters urban population structures that differ markedly from those represented in the training data. The stratified design of the training set also leads to systematic overestimation in extremely low-density global regions. As a result, some hotspot signals in the interannual change analysis may arise from prediction jumps in low-density areas rather than actual population displacement. Although extreme anomalies are mitigated through base-population filtering and quantile truncation, this issue has not been fundamentally resolved at the methodological level.

- **Validation and extrapolation.** For global 1 km mapping, the training labels and back-testing reference data are derived from the same source, resulting in an in-sample evaluation rather than an independent test. The 2021–2024 extrapolated products lack cross-validation against contemporaneous independent reference data, and their absolute accuracy remains to be assessed using external sources such as mobile phone data, census records, or nighttime lights. In the experimental design, the temporal composition of the S2 training set, which mixes multiple years, may introduce both temporal information gains and data-volume gains, making their respective contributions difficult to separate; future work could use leave-one-year-out temporal blocking experiments to clarify this issue. More fundamentally, training on 2017 – 2020 data and extrapolating to 2021 – 2024 implicitly assumes that the relationship between AEF embeddings and population density remains stable over this period. However, disruptive events such as the COVID-19 pandemic, geopolitical conflicts, and climate-related disasters may cause structural breaks in population–surface relationships in specific regions, exceeding the representational capacity of the model.

8. Conclusions

This study provides a systematic evaluation of the applicability of AEF for population spatial modeling as an indirectly coupled task. The main conclusions are as follows. First, AEF embeddings exhibit strong explanatory power for population spatial patterns under controlled evaluation settings; however, their generalization performance is jointly constrained by spatial autocorrelation dependence and cross-city feature distribution shifts. Second, the combination of the S2 training strategy, C3 spatial structure variables, and the TorchMLP model represents the most effective modeling configuration identified in this study. Third, under this configuration, global annual population maps at 1 km resolution for 2021–2024 can be generated through stable temporal extrapolation without retraining, providing a feasible solution to fill the temporal gaps in WorldPop after 2020.

Future work can be advanced along several directions. At the methodological level, incorporating domain adaptation techniques or city-level embedding normalization to mitigate cross-city feature shifts is a priority for improving cross-regional generalization. Integrating fine-grained urban morphological data, such as building height and functional zoning, as auxiliary variables may further compensate for the limitations of AEF in indirectly coupled tasks. At the validation level, independent accuracy assessment using mobile phone signaling data, time-series nighttime lights, or census microdata is essential to enhance the reliability of extrapolated products. At the application level, the annual update cycle of AEF products offers opportunities for near real-time population monitoring; integrating the proposed framework with time-sensitive applications such as disaster response and epidemic modeling would provide a meaningful test of its practical value.

Disclosure statement

No potential conflict of interest was reported by the authors.

Funding

The work was supported by National Natural Science Foundation of China (no. 452B2013).

ORCID

Shuyang Hou: http://orcid.org/0009-0000-6984-9959

Huayi Wu: http://orcid.org/0000-0003-3971-0512

Funding

The work was supported by the National Natural Science Foundation of China [Grant number 452B2013].

Acknowledgements

This work was supported by the National Natural Science Foundation of China under Grant No. 452B2013, awarded to Shuyang Hou.

References

Alvarez, C. I., Ulloa Vaca, C. A., & Echeverria Llumipanta, N. A. (2025). Machine learning for urban air quality prediction using Google AlphaEarth Foundations satellite embeddings: a case study of Quito, Ecuador. *Remote Sensing*, *17*(20), 3472.

Benavides-Martinez, I. F., Guthrie, J., Arias, J. E., Garces-Gomez, Y. A., Guzman-Alvis, A. I., Portilla-Cabrera, C. V., Mondal, S., Allyn, A. J., & Ganguly, A. R. (2026). What on Earth is AlphaEarth? Hierarchical structure and functional interpretability for global land cover. *arXiv preprint arXiv:2603.16911*.

Brown, C. F., Kazmierski, M. R., Pasquarella, V. J., Rucklidge, W. J., Samsikova, M., Zhang, C., Shelhamer, E., Lahera, E., Wiles, O., & Ilyushchenko, S. (2025). Alphaearth foundations: An embedding field model for accurate and efficient global mapping from sparse label data. *arXiv preprint arXiv:2507.22291*.

Deville, P., Linard, C., Martin, S., Gilbert, M., Stevens, F. R., Gaughan, A. E., Blondel, V. D., & Tatem, A. J. (2014). Dynamic population mapping using mobile phone data. *Proceedings of the National Academy of Sciences*, *111*(45), 15888-15893.

Dobson, J. E., Bright, E. A., Coleman, P. R., Durfee, R. C., & Worley, B. A. (2000). LandScan: a global population database for estimating populations at risk. *Photogrammetric engineering and remote sensing*, *66*(7), 849-857.

Fang, C., Buhendwa, B. M., Riaz, F., & Wang, J. (2026). Performance Assessment and Application of AlphaEarth Embedding In Air Quality-Case of PM2. 5 in Changchun. *Remote Sensing Applications: Society and Environment*, 101989.

Fang, J., Wu, M., Zhang, Z., & Luo, W. (2025). Leveraging AlphaEarth Foundations Embeddings for High-Accuracy County-Scale Corn and Soybean Yield Estimation. *Authorea Preprints*.

Feng, Z., Atzberger, C., Jaffer, S., Knezevic, J., Sormunen, S., Young, R., Lisaius, M. C., Immitzer, M., Jackson, T., & Ball, J. (2025). Tessera: Temporal embeddings of surface spectra for earth representation and analysis. *arXiv preprint arXiv:2506.20380*.

Georganos, S., Hafner, S., Kuffer, M., Linard, C., & Ban, Y. (2022). A census from heaven: Unraveling the potential of deep learning and Earth Observation for intra-urban population mapping in data scarce environments. *International Journal of Applied Earth Observation and Geoinformation*, *114*, 103013.

Huo, C., Chen, K., Zhang, S., Wang, Z., Yan, H., Shen, J., Hong, Y., Qi, G., Fang, H., & Wang, Z. (2025). When remote sensing meets foundation model: A survey and beyond. *Remote Sensing*, *17*(2), 179.

Jean, N., Wang, S., Azzari, G., Lobell, D., & Ermon, S. (2018). Tile2vec: Unsupervised representation learning for remote sensing data. *arXiv preprint arXiv:1805.02855*.

Khan, H., & Ahmad, A. (2025). Evaluating AlphaEarth Foundation Embeddings for Pixel-and Object-Based Land Cover Classification in Google Earth Engine.

Leyk, S., Gaughan, A. E., Adamo, S. B., De Sherbinin, A., Balk, D., Freire, S., Rose, A., Stevens, F. R., Blankespoor, B., & Frye, C. (2019). The spatial allocation of population: a review of large-scale gridded population data products and their fitness for use. *Earth System Science Data*, *11*(3), 1385-1409.

Liu, L., Cao, X., Li, S., & Jie, N. (2024). A 31-year (1990–2020) global gridded population dataset generated by cluster analysis and statistical learning. *Scientific Data*, *11*(1), 124.

Ma, Y., Shen, Y., Swatantran, A., & Lobell, D. B. (2026). Harvesting AlphaEarth: Benchmarking the geospatial foundation model for agricultural downstream tasks. *International Journal of Applied Earth*

Observation and Geoinformation, *149*, 105258.

Manas, O., Lacoste, A., Giró-i-Nieto, X., Vazquez, D., & Rodriguez, P. (2021). Seasonal contrast: Unsupervised pre-training from uncurated remote sensing data. Proceedings of the IEEE/CVF international conference on computer vision,

Mendieta, M., Han, B., Shi, X., Zhu, Y., & Chen, C. (2023). Towards geospatial foundation models via continual pretraining. Proceedings of the IEEE/CVF International Conference on Computer Vision,

Mennis, J. (2003). Generating surface models of population using dasymetric mapping. *The Professional Geographer*, *55*(1), 31-42.

Metzger, N., Daudt, R. C., Tuia, D., & Schindler, K. (2024). High-resolution population maps derived from sentinel-1 and sentinel-2. *Remote Sensing of Environment*, *314*, 114383.

Neal, I., Seth, S., Watmough, G., & Diallo, M. S. (2022). Census-independent population estimation using representation learning. *Scientific Reports*, *12*(1), 5185.

Pesaresi, M., Schiavina, M., Politis, P., Freire, S., Krasnodębska, K., Uhl, J. H., Carioli, A., Corbane, C., Dijkstra, L., & Florio, P. (2024). Advances on the Global Human Settlement Layer by joint assessment of Earth Observation and population survey data. *International Journal of Digital Earth*, *17*(1), 2390454.

Qu, P., Ouyang, W., Zhang, C., Chai, Y., Xu, S., Ye, L., Piao, Y., Zhang, M., & Lu, H. (2026). Utilizing Earth Foundation Models to Enhance the Simulation Performance of Hydrological Models with AlphaEarth Embeddings. *arXiv preprint arXiv:2601.01558*.

Rahman, M. (2026). Physically Interpretable AlphaEarth Foundation Model Embeddings Enable LLM-Based Land Surface Intelligence. *arXiv preprint arXiv:2602.10354*.

Song, Y., Wu, S., Chen, B., & Bell, M. L. (2024). Unraveling near real-time spatial dynamics of population using geographical ensemble learning. *International Journal of Applied Earth Observation and Geoinformation*, *130*, 103882.

Stevens, F. R., Gaughan, A. E., Linard, C., & Tatem, A. J. (2015). Disaggregating census data for population mapping using random forests with remotely-sensed and ancillary data. *PloS one*, *10*(2), e0107042.

Tatem, A. J. (2017). WorldPop, open data for spatial demography. *Scientific Data*, *4*(1), 170004.

Thomson, D. R., Rhoda, D. A., Tatem, A. J., & Castro, M. C. (2020). Gridded population survey sampling: a systematic scoping review of the field and strategic research agenda. *International journal of health geographics*, *19*(1), 34.

Tobler, W. R. (1979). Smooth pycnophylactic interpolation for geographical regions. *Journal of the American Statistical Association*, *74*(367), 519-530.

Tseng, G., Cartuyvels, R., Zvonkov, I., Purohit, M., Rolnick, D., & Kerner, H. (2023). Lightweight, pre-trained transformers for remote sensing timeseries. *arXiv preprint arXiv:2304.14065*.

Wu, S.-s., Qiu, X., & Wang, L. (2005). Population estimation methods in GIS and remote sensing: A review. *GIScience & Remote Sensing*, *42*(1), 80-96.

Xiao, A., Xuan, W., Wang, J., Huang, J., Tao, D., Lu, S., & Yokoya, N. (2025). Foundation models for remote sensing and earth observation: A survey. *IEEE Geoscience and Remote Sensing Magazine*.

Xu, Y., Song, Y., Cai, J., & Zhu, H. (2021). Population mapping in China with Tencent social user and remote sensing data. *Applied Geography*, *130*, 102450.

Yang, X., Ye, T., Zhao, N., Chen, Q., Yue, W., Qi, J., Zeng, B., & Jia, P. (2019). Population mapping with multisensor remote sensing images and point-of-interest data. *Remote Sensing*, *11*(5), 574.

Zhang, J., & Zhao, X. (2024). Using POI and multisource satellite datasets for mainland China's

population spatialization and spatiotemporal changes based on regional heterogeneity. *Science of The Total Environment*, *912*, 169499.

Zheng, Q., Seto, K. C., Zhou, Y., You, S., & Weng, Q. (2023). Nighttime light remote sensing for urban applications: Progress, challenges, and prospects. *ISPRS Journal of Photogrammetry and Remote Sensing*, *202*, 125-141.

Zvonkov, I., Tseng, G., Becker-Reshef, I., & Kerner, H. (2025). Cropland Mapping using Geospatial Embeddings. *arXiv preprint arXiv:2511.02923*.\

Appendix A. Basic Information of Study Cities

This study includes 18 cities worldwide as case studies for AEF-based population spatial modeling, covering Asia, Europe, Africa, North America, South America, and Oceania, with a temporal span from 2017 to 2024. Detailed information on each city—including continent, country, longitude and latitude range, and selection rationale—is provided in Table A.

Table A. Basic information of study cities

ID	Continent	City	Country	Longitude Range	Latitude Range	Selection Rationale
1	Asia	Tokyo	Japan	139.40°–139.98°E	35.45°–35.93°N	Developed East Asian megacity with high-density and well-structured spatial patterns
2		Jakarta	Indonesia	106.61°–107.08°E	6.45°–5.98°S	Largest Southeast Asian city, representative of informal urban expansion
3		Delhi	India	76.96°–77.50°E	28.42°–28.90°N	Extremely high-density South Asian city with strong internal heterogeneity
4		Guangzhou	China	113.00°–113.52°E	22.89°–23.37°N	Representative Chinese city with rapid urbanization
5		Dhaka	Bangladesh	90.14°–90.65°E	23.48°–23.96°N	One of the world's highest-density cities, typical low-income high-density context
6	Europe	Moscow	Russia	37.19°–38.04°E	55.52°–55.99°N	High-latitude megacity with legacy Soviet planning structure
7		Istanbul	Turkey	28.65°–29.28°E	40.77°–41.25°N	Transcontinental city, representative of developing Europe
8		London	United Kingdom	0.55°W–0.32°E	51.28°–51.76°N	Mature polycentric Western European city with high-quality data
9	North America	New York	United States	74.23°–73.61°W	40.46°–40.93°N	Typical high-density compact metropolitan area in North America
10		Los Angeles	United States	118.69°–118.12°W	33.88°–34.35°N	Low-density sprawled city, contrasting with New York
11		Mexico City	Mexico	99.38°–98.89°W	19.20°–19.66°N	Megacity in developing North America, representative of Latin America
12	South America	São Paulo	Brazil	46.89°–46.38°W	23.79°–23.31°S	Largest South American city with strong socioeconomic heterogeneity
13		Buenos Aires	Argentina	58.67°–58.10°W	34.84°–34.36°S	Mature metropolitan area, complementary

ID	Continent	City	Country	Longitude Range	Latitude Range	Selection Rationale
						to São Paulo in development stage
14		Cairo	Egypt	30.97°–31.51°E	29.82°–30.29°N	Compact urban form in arid North Africa
15		Lagos	Nigeria	3.17°–3.63°E	6.22°–6.68°N	Rapidly growing West African city dominated by informal expansion
16	Africa	Nairobi	Kenya	36.58°–37.20°E	1.55°S–0.95°S	Representative East African city with mixed formal–informal structure
17		Johannesburg	South Africa	27.85°–28.45°E	26.45°–25.90°S	City with pronounced spatial segregation shaped by historical legacy
18	Oceania	Sydney	Australia	150.65°–151.32°E	34.18°–33.59°S	Developed, low-density suburbanized city and the only Oceania representative

Appendix B. Summary of Auxiliary Variables

A total of four categories of auxiliary variables are constructed in this study, including nighttime lights, remote sensing indices, spatial structure variables, and POIs, comprising 57 feature layers in total. These variables are complementary in terms of data source, temporal dynamics, and spatial semantics. Nighttime lights and remote sensing indices are both derived from satellite observations: the former captures temporal dynamics of urban activity intensity, while the latter reflects physical properties of land surface cover. Spatial structure variables and POIs are extracted from historical OpenStreetMap (OSM) data: the former characterize macro-scale urban spatial organization (e.g., centrality, geographic constraints, and transport infrastructure), whereas the latter describe fine-scale functional distributions and service accessibility. All variables are generated annually and aligned at the pixel level with AEF embeddings. Detailed descriptions are provided in Table B.

Table B. Summary of multi-source auxiliary variables

Group	Feature Name	Description	Data Source	Temporal Type	Layers
Nighttime Lights (NTL)	viirs_mean	Annual mean radiance	VIIRS/DNB VCMSLCFG	Dynamic (annual)	1
	viirs_median	Annual median radiance		Dynamic	1
	viirs_max	Annual maximum radiance		Dynamic	1
	viirs_min	Annual minimum radiance		Dynamic	1
	viirs_std	Monthly standard deviation		Dynamic	1
	viirs_p10 / viirs_p90	10th/90th percentile radiance		Dynamic	2
	viirs_cv	Coefficient of variation (monthly variability)		Dynamic	1

Group	Feature Name	Description	Data Source	Temporal Type	Layers
	viirs_lit_fraction	Fraction of lit pixels (annual mean)		Dynamic	1
	viirs_summer / viirs_winter	Seasonal mean (summer/winter)		Dynamic	2
	viirs_seasonal_diff	Seasonal difference (summer–winter)		Dynamic	1
	viirs_cny / viirs_cny_diff	Chinese New Year mean and deviation		Dynamic	2
	viirs_focal_mean	3×3 neighborhood mean		Dynamic	1
	viirs_local_contrast	Local contrast (pixel − neighborhood mean)		Dynamic	1
	viirs_sl_mean	Annual mean with straylight correction	VIIRS/DNB VCMCFG	Dynamic	1
	viirs_straylight_contrib	Straylight contribution (VCMCFG − VCMSLCFG)		Dynamic	1
	dmsp_2013_stable	Stable lights (2013 baseline)	NOAA DMSP-OLS	Static	1
	ntl_growth_vs_dmsp2013	VIIRS-based growth relative to DMSP	VIIRS/DMSP	Dynamic	1

Group	Feature Name	Description	Data Source	Temporal Type	Layers
Subtotal					**20**
Remote Sensing Indices	NDVI	Normalized Difference Vegetation Index	Landsat 8 (GEE annual composite)	Dynamic	1
	NDBI	Normalized Difference Built-up Index		Dynamic	1
	GHSL built-up density	Built-up area density (interpolated 2015/2020)	JRC/GHSL P2023A	Quasi-static	1
Subtotal					**3**
Spatial Structure Variables	dist_center_main	Distance to primary urban center (km)	OSM historical data	Quasi-static	1
	dist_center_secondary	Distance to nearest secondary center (km)		Quasi-static	1
	dist_water	Distance to nearest water body (km)		Quasi-static	1
	road_arterial	Density of arterial roads (km/km²)		Dynamic	1
	road_collector	Density of collector roads (km/km²)		Dynamic	1

Group	Feature Name	Description	Data Source	Temporal Type	Layers
	road_local	Density of local roads (km/km²)		Dynamic	1
	road_density	Total road network density (km/km²)		Dynamic	1
	dist_metro	Distance to nearest metro station (km)		Dynamic	1
	metro_density	Number of metro stations within 1 km buffer		Dynamic	1
	dist_bus	Distance to nearest bus stop (km)		Dynamic	1
Subtotal					**10**
POI	poi_food_density / _dist	Kernel density / nearest distance (food)	OSM historical data	Dynamic	2
	poi_education_density / _dist	Education POI density / distance		Dynamic	2
	poi_healthcare_density / _dist	Healthcare POI density / distance		Dynamic	2
	poi_finance_density / _dist	Finance POI density / distance		Dynamic	2
	poi_shopping_density / _dist	Retail POI density / distance		Dynamic	2

Group	Feature Name	Description	Data Source	Temporal Type	Layers
	poi_entertainment_density / _dist	Entertainment POI density / distance		Dynamic	2
	poi_public_service_density / _dist	Public service POI density / distance		Dynamic	2
	poi_leisure_density / _dist	Leisure POI density / distance		Dynamic	2
	poi_accommodation_density / _dist	Accommodation POI density / distance		Dynamic	2
	poi_religion_density / _dist	Religious POI density / distance		Dynamic	2
	poi_transport_service_density / _dist	Transport service POI density / distance		Dynamic	2
	poi_office_density / _dist	Office POI density / distance		Dynamic	2
Subtotal					**24**
Total					**57**

Appendix C

Table C1. Hyperparameter settings for traditional machine learning models

Model	Hyperparameters
Ridge	α=1.0
ElasticNet	α=0.01；l1_ratio=0.5；max_iter=2000
HistGBT	max_iter=200；max_depth=6；learning_rate=0.05；min_samples_leaf=20
RandomForest	n_estimators=200；max_depth=12；min_samples_leaf=5；max_features=√d

Table C2. Optimal hyperparameter configuration for TorchMLP

Item	Value
Hidden layers	(512, 256, 128, 64)
Block structure	Linear → BatchNorm → ReLU → Dropout
Dropout	0.2
Optimizer	Adam (weight_decay = 1×10^{-4})
Loss	MSE
Initial learning rate	5×10^{-4}
LR scheduler	CosineAnnealingLR (T_max = epochs)
Batch size	4096
Max epochs	300
Early stopping	patience = 20
Validation split	15% (from training set)
Feature scaling	RobustScaler
Framework	PyTorch + CUDA

Table C3. TorchMLP-Adaptive configurations under different sample sizes

Tier	Training set size	Hidden layers	Batch size	Typical scale coverage
Large	n>50,000	(512,256,128,64)	4,096	100 m (mid–large cities)
Medium	10,000<n≤50,000	(256,128,64)	n/8	200 m (most cities)
Small	2,000<n≤10,000	(128,64,32)	n/8	500 m (most cities),200 m (small cities)
Tiny	500<n≤2,000	(64,32)	n/8	1,000 m (mid–large cities)
Minimal	n≤500	(32,64)	n/8	1,000 m (small cities)