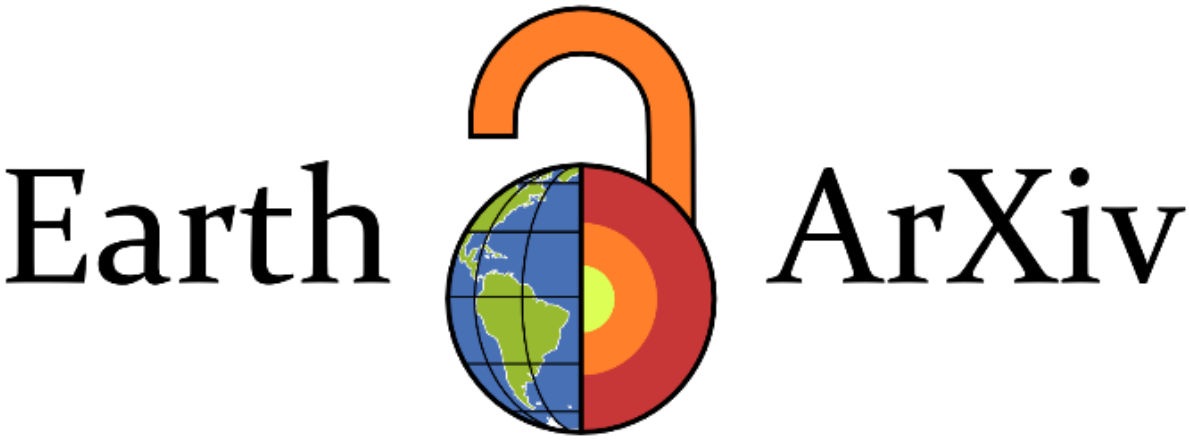

Pixel-Level Urban Housing-Price Mapping Based on AlphaEarth Foundations: Evidence from 36 Chinese Cities

Shuyang Hou[a,b], Haoyue Jiao[a]*, Ziqi Liu[a], Donghai Liu[a], Lutong Xie[a], Shaowen Wu[a], Guanyu Chen[a], Yaxian Qing[a], Qingyang Xu[a,c], Xuefeng Guan[a], Huayi Wu[a]

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

b. HUBEI LUOJIA LABORATORY, Wuhan, China

c. School of Computer Science and Technology, Zhejiang Normal University, Jinhua, China

*Corresponding authors:
Haoyue Jiao, email: haoyuejiao@whu.edu.cn

Abstract

This study develops and validates a multi-source modelling framework for continuous, pixel-level urban housing-price mapping using surface embeddings from AlphaEarth Foundations (AEF). Pixel-level labels calibrated against multi-source market data are constructed for 288 city–year samples across 36 Chinese cities (2017–2024), and AEF's 64-dimensional, 10 m annual surface embeddings are systematically evaluated across four training strategies, fourteen feature combinations, and four baseline models, with validation through PCA ablation, SHAP attribution, IDW comparison, spatial clustering, and temporal consistency analysis. AEF embeddings alone yield limited predictive power, with pixel-level R^2 of only 0.66 and 0.15 under random and spatial cross-validation, respectively; incorporating points of interest, nighttime light, population density, spatial structure, and remote-sensing indices raises these values to 0.84 and 0.49. Full-coverage prediction surfaces for 36 cities over eight years achieve a median Pearson r of 0.948; spatial clustering intensity closely matches the IDW benchmark (Moran's I = 0.865), and the median adjacent-year structural similarity index reaches 0.986. Analysis further reveals that AEF implicitly encodes substantial built-form information—explicit physical features such as building height and footprint area provide only marginal gains—yet price-relevant signals are distributed across the full 64-dimensional space without a stable cross-city subset, confirming that socioeconomic auxiliary factors remain indispensable for generalisation. Building on the continuous prediction surfaces, downstream applications—housing-price tier transitions, CBD premium gradients, and spatial redistribution during market corrections—show that surface foundation-model embeddings support not only housing-price prediction but also large-scale, continuous, and temporally consistent mapping and spatial analysis of urban economic variables.

Keywords: AlphaEarth Foundations; urban housing-price mapping; representation learning; foundation models; continuous socioeconomic variables; urban remote sensing

1. Introduction

Housing constitutes one of the most important stores of wealth for urban residents, and the spatial

distribution of housing prices encapsulates land value, public-service provision, employment opportunities, and socioeconomic inequality. In China, residential assets account for approximately 60% of household wealth[1, 2], making the timely characterisation of spatial housing-price patterns essential for urban planning, tax-base assessment, housing policy, and the monitoring of spatial inequality. However, sparse transaction records, the non-disclosure of historical listing prices, and the heavy reliance of community-level hedonic modelling on auxiliary variables[3] have long constrained the availability of housing-price data that is broadly covered, finely resolved, and temporally continuous.

Remote sensing offers a scalable data foundation for large-scale housing-price modelling. Prior studies have estimated residential value using high-resolution optical imagery, texture features, convolutional neural networks, and vision foundation models[4-7], yet most evaluations and outputs remain confined to the administrative-district, street, or neighbourhood scale, which fails to capture the continuous yet locally abrupt price gradients within cities. Housing prices carry clear economic meaning and are additionally shaped by micro-scale factors such as school-district quality[8], transport access[9, 10], commercial amenities[11], and landscape resources[12]; naive spatial aggregation therefore obscures the price stratification between adjacent neighbourhoods. Meanwhile, medium-resolution sensors such as nighttime-light and Landsat imagery have been used predominantly to predict composite socioeconomic indicators such as poverty and wealth[13, 14], leaving direct pixel-level modelling of continuous housing-price variables comparatively underexplored[15].

AlphaEarth Foundations (AEF) provides a new representational basis for addressing this gap. By integrating multi-source observations—optical, synthetic aperture radar (SAR), climate, and textual data—at approximately 10 m resolution, AEF employs self-supervised learning to generate 64-dimensional annual surface embeddings that offer globally consistent coverage, annual updates, and analysis-ready outputs[16]. Prior work has demonstrated that AEF generalises well across a broad spectrum of downstream Earth observation tasks, including crop mapping and agricultural monitoring[17], land-cover mapping and functional-zone classification[18], gridded population spatialization[19], slum detection and urban poverty mapping[20], and socioeconomic indicator estimation[21]; however, housing prices are jointly shaped by visible built form and unobservable socioeconomic processes. Signals such as building density and street texture may already be implicitly encoded by AEF, whereas factors such as school-district quality, commercial vibrancy, and transport accessibility likely require supplementary auxiliary variables. Whether AEF embeddings can directly support pixel-level housing-price regression, and how they complement socioeconomic auxiliary factors, therefore still requires systematic validation.

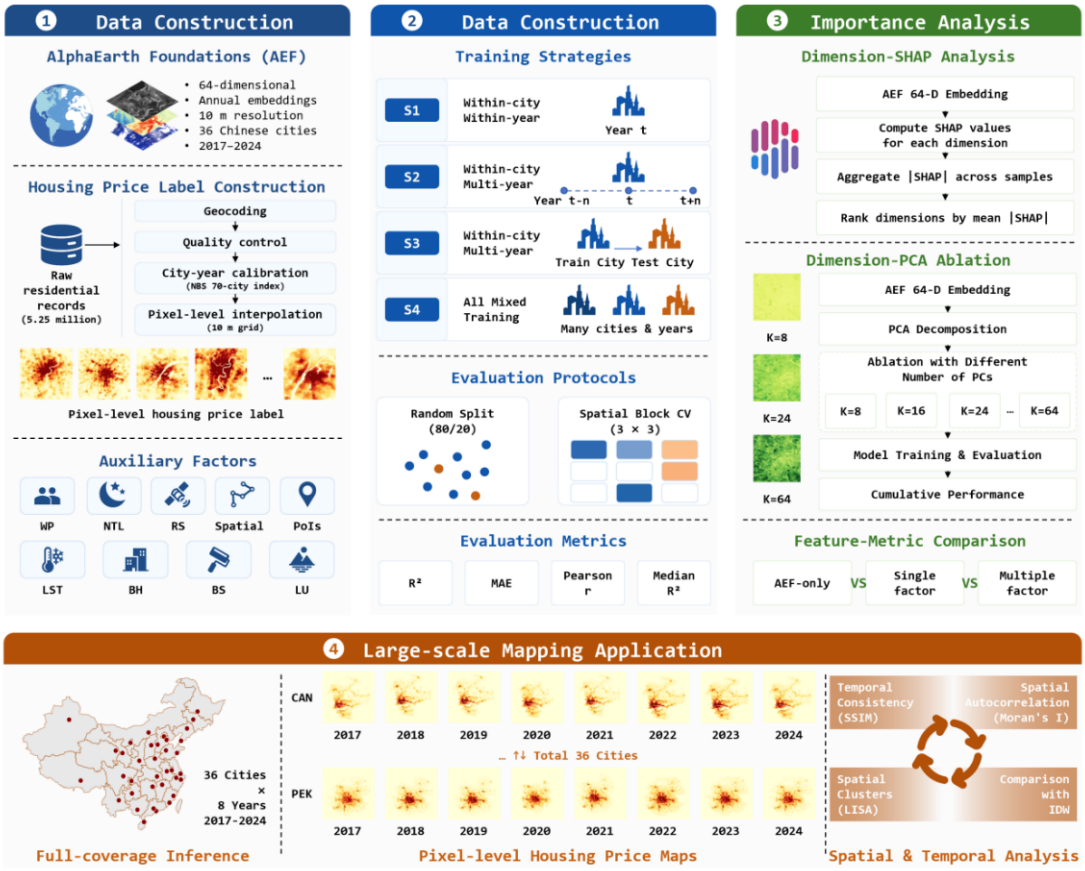

Figure 1. Overall workflow for pixel-level urban housing-price mapping from AEF embeddings, covering data construction, model training and evaluation, importance analysis, and large-scale mapping (map lines delineate study areas and do not necessarily depict accepted national boundaries).

This study focuses on 36 representative Chinese cities, comprising 288 city–year samples spanning 2017–2024. Pixel-level housing-price labels are constructed from publicly accessible secondary-housing listings on multiple major online real estate platforms in China, and calibrated against multi-source market data, including the National Bureau of Statistics' housing-price index for 70 large and medium-sized cities as well as supplementary data from commercial housing market research institutions. Centred on the workflow shown in **Figure 1**, this study addresses three core questions: first, whether AEF's 64-dimensional embeddings can support intra-city and cross-city housing-price prediction under different training strategies and evaluation protocols; second, how housing-price information is distributed within the AEF embedding space, and whether a stable, cross-city subset of key dimensions exists; and third, to what extent auxiliary factors—POI, nighttime light, population density, spatial structure, and remote-sensing indices—improve predictive accuracy and enable full-coverage, continuous, and temporally consistent urban housing-price mapping.

The main contributions of this study are threefold. First, it provides the first systematic evaluation of AEF's global surface embeddings for continuous urban housing-price regression, covering 288 city–year samples across 36 Chinese cities from 2017 to 2024, thereby clarifying both the capabilities and limitations of foundation-model embeddings in modelling continuous socioeconomic variables. Second, it constructs a pixel-level housing-price label system calibrated against multi-source market data, systematically quantifies the marginal contribution and synergistic effects of nine auxiliary

factors, and reveals the dimensional structure of AEF embeddings, their generalisation boundaries, and their complementarity with socioeconomic auxiliary factors. Third, it generates continuous housing-price surfaces for 36 cities using the best-performing model and validates their reliability in terms of spatial pattern, temporal consistency, and representative downstream applications, demonstrating that foundation-model embeddings can support large-scale, continuous, and high-resolution urban housing-price mapping.

The remainder of this paper is organised as follows. Section 2 introduces the study area and data; Section 3 describes the research methods; Sections 4 and 5 report the experimental results and full-coverage mapping; Section 6 discusses the main findings and limitations; and Section 7 concludes the paper.

2. Study Area and Data

2.1. Study area

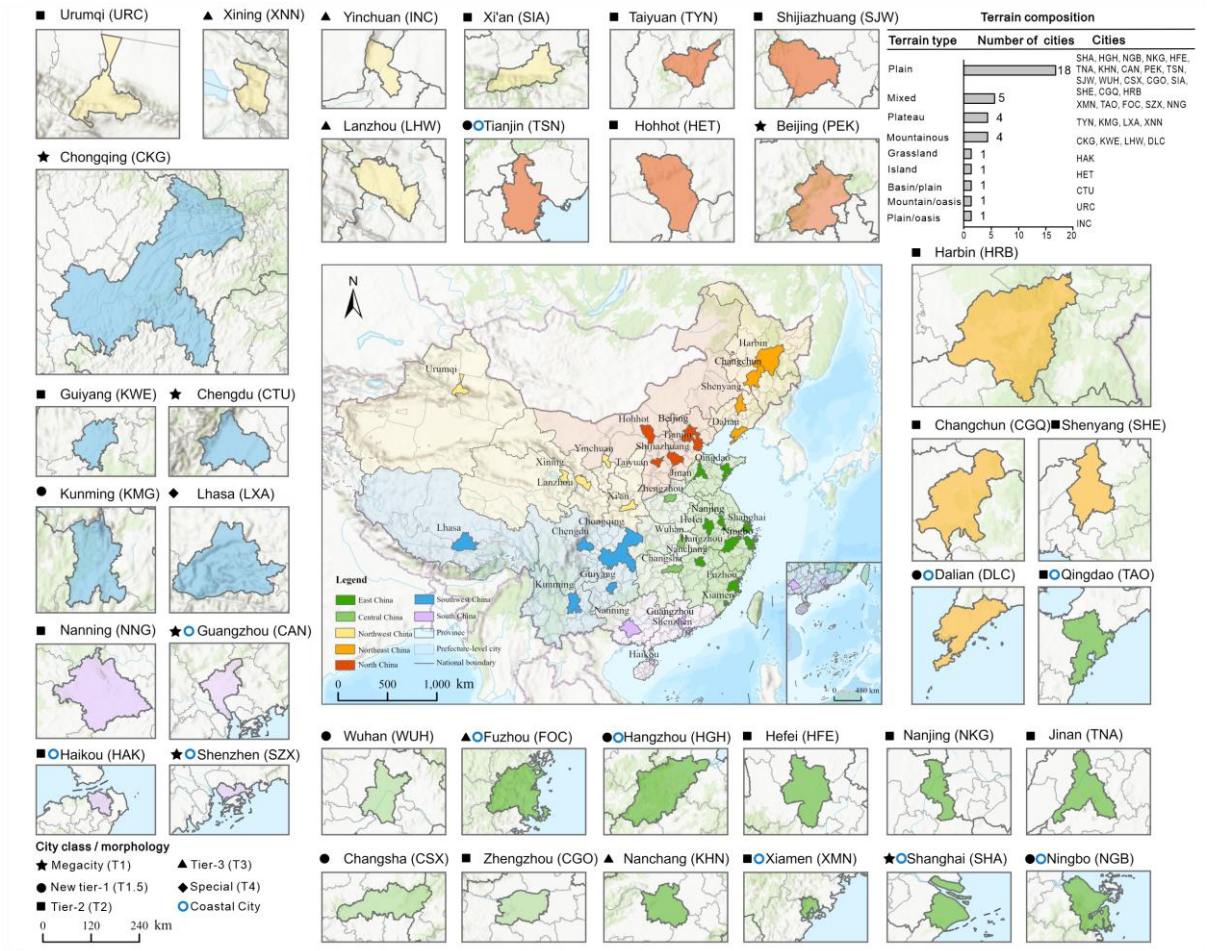

Figure 2. Spatial distribution, city tiers, coastal setting, and terrain composition of the 36 study cities in China.

This study selects 36 representative Chinese cities spanning seven major regions—East China, South China, North China, Central China, Southwest China, Northwest China, and Northeast

China—covering tier-1, new-tier-1, tier-2, tier-3, and special-category cities, and encompassing diverse geographical settings such as coastal, inland, mountainous, plateau, plain, and oasis environments. City selection followed four criteria: (1) inclusion in the National Bureau of Statistics' housing-price index for 70 large and medium-sized cities, to ensure reliable annual city-level price anchors; (2) balanced representation across economic development levels and city tiers; (3) coverage of diverse urban morphologies; and (4) complete availability of all auxiliary data across cities and years.

The study period spans 2017–2024, comprising eight years and 288 city–year samples; the spatial distribution and key attributes of the 36 study cities are illustrated in **Figure 2**.

2.2. Data and preprocessing

2.2.1. AlphaEarth Foundations

AlphaEarth Foundations (AEF) embeddings were obtained via Google Earth Engine, covering the full study extent of all 36 cities from 2017 to 2024. AEF provides globally consistent 64-dimensional annual surface embeddings at approximately 10 m spatial resolution that are directly usable for downstream analysis[16]. This study adopts the EPSG:4326 coordinate system, storing data in float32 format, with each city–year represented by a single 64-band raster.

2.2.2. Auxiliary factors

AEF embeddings are derived from multi-source remote-sensing data through self-supervised learning, and the extent to which they encode housing-price-relevant socioeconomic semantics remains to be verified. This study therefore introduces auxiliary factors alongside AEF and compares the predictive performance of "AEF alone" against "AEF plus auxiliary factors" to assess which information AEF already encodes implicitly and which semantic content still requires external supplementation.

Nine categories of auxiliary factors were constructed: population density (WP), nighttime light (NTL), remote-sensing indices (RS), spatial structure (Spatial)[22], points of interest (POI), land surface temperature (LST), building height (BH), building footprint area (BS), and land use and building age (LU). These factors were selected to balance theoretical interpretability with respect to housing-price formation mechanisms, public data availability, and complete spatiotemporal coverage across the 36 cities for 2017–2024[23-25].

The data sources, native resolutions, and processing pipelines for each auxiliary factor are summarised in **Table 1**. All datasets were temporally aligned to the corresponding year and spatially resampled to 10 m to match the AEF grid. Each auxiliary factor was ultimately stored as a pixel-wise feature and concatenated channel-wise with the AEF embedding to form the model input; the resulting dimensional composition is detailed in **Table 2**.

Table 1. Auxiliary factors: data sources and processing pipeline

Name	Data source	Original resolution	Processing pipeline
WP	WorldPop population density	100 m	Bilinear resampling to 10 m; yearly alignment
NTL	VIIRS DNB+DMSP-OLS	~500 m	Monthly-to-annual aggregation; outlier clipping; bilinear resampling
RS	Landsat 8+GHSL	10–30 m	Annual median compositing; NDVI/NDBI calculation; spatial alignment
Spatial	OpenStreetMap road network	Vector	Road graph construction; distance and density calculation; rasterization
POI	OpenStreetMap POIs	Vector	POI categorization; kernel density estimation; nearest-distance calculation
LST	MODIS MOD11A2	1 km	Annual averaging; bilinear resampling
BH	JRC GHS_BUILT_H	10 m	Missing-value filling; cross-year reuse
BS	JRC GHS_BUILT_S	10 m	Temporal mosaicking; raster alignment
LU	CLCD+GISA	30 m	Nearest-neighbour resampling; land-use encoding; first-built-year extraction

Following upstream processing, each factor was stored as a pixel-level feature column in the label CSV file to facilitate channel-wise concatenation with the AEF embedding. The dimensional composition and sub-feature semantics of each auxiliary factor at the model-input level are presented in **Table 2**.

Table 2. Input-level dimensional composition and sub-feature semantics of the nine auxiliary factors

Tag	Source	Dim	Sub-feature composition	Data type
F1	WP	1	Population density	Continuous
F2	NTL	18	VIIRS annual statistics (mean/median/max/min/std/p10/p90/cv/lit-ratio,9 dims) +seasonal contrast (1 dim) + 3×3 focal mean & contrast (2 dims) + stray light (1 dim) + DMSP-OLS 2013 baseline (2 dims) + growth since 2013 (3 dims)	Continuous
F3	RS	3	NDVI (1 dim) + NDBI (1 dim) + GHSL built-up density (1 dim)	Continuous (NDVI/NDBI $\in [-1,1]$; GHSL $\in [0,1]$)
F4	Spatial	9	Distance-to-center (1 dim) + distance-to-water (1 dim) + road density (arterial/sub-arterial/branch/total,4 dims) + distance-to-subway (1 dim) + subway stations within 1 km (1 dim) + distance-to-bus (1 dim)	Continuous
F5	POI	12	Kernel density for 12 categories: food, education, healthcare, finance, shopping, entertainment, public service, leisure, accommodation, religion, transport, office	Continuous
F6	POI	12	Nearest distance for the same 12 POI categories	Continuous
F7	LST	1	Annual mean land surface temperature(°C)	Continuous
F8	BH	1	Mean building height(m)	Continuous, right-skewed
F9	BS	1	Building surface area per cell(m²)	Continuous, right-skewed
F10	LU	3	CLCD 9-class land-use (1 dim) + GISA first-built year (1 dim) +built-by-2018 indicator (1 dim)	Categorical+continuous+binary

2.2.3. Housing-price labels

Community-level housing-price data were continuously collected by the research team between 2017 and 2024 from publicly available secondary-housing listings on multiple major online real estate platforms in China. The raw dataset covers the 36 study cities and comprises approximately 5.25 million community–year observations, including fields such as community name, address, monthly average listing price, unit-type distribution, and year of construction. To mitigate cross-platform discrepancies and the influence of anomalous listings, cross-platform price consistency was first checked for the same community and month, and records with a relative deviation exceeding 20% were discarded.

Because most historical records contain only textual addresses and some coordinate fields are missing or insufficiently precise, community addresses were converted to WGS-84 coordinates using the AutoNavi (Amap) Geocoding API[26], and a geocoding confidence score was constructed based on the match level, administrative-boundary consistency, and matching pattern:

$$\eta_i = w_L s_L(level_i) + w_A s_A(admin_i) + w_M s_M(match_i) \quad (1)$$

Here, s_L, s_A, and s_M denote the match-level, administrative-boundary consistency, and matching-pattern scores, respectively, with weights set to 0.5, 0.3, and 0.2. Low-confidence records with ($\eta_i < 0.8$) below the threshold were manually reviewed, and records with evident positional offsets were further corrected or removed based on reverse geocoding, study-area boundaries, and the positional consistency of same-named communities across their historical records.

To improve the cross-city and cross-year comparability of listing prices, this study constructs annual city-level price anchors from authoritative price data. Specifically, the National Bureau of Statistics' housing-price index for 70 large and medium-sized cities is combined with supplementary residential listing and transaction price data from commercial housing market research institutions, using a weighted average that prioritises the official government index (assigned a weight of 0.5, with the remaining weight distributed across the supplementary commercial sources), to obtain the annual city-level price benchmark ($A_{c,y}$). Community listing prices were then compared against the corresponding city–year anchor, and outliers deviating from the city's annual mean by more than 3σ, as well as records with a cross-platform relative deviation exceeding 20%, were removed.

For communities with intermittent missing years, missing values were imputed using a "city price anchor plus community fixed effect" approach: the community's average offset relative to the city's annual mean price was first estimated from its observed years to capture its stable premium or discount level, and this offset was then combined with the city price anchor for the missing year to back-calculate the community price. Communities with more than five consecutive missing years, or with observations in only a single year, were excluded rather than imputed. The resulting dataset comprises community-level housing-price labels for 36 cities over the eight years from 2017 to 2024; a summary is provided in **Table 3**.

Table 3. Overview of the community-level housing-price dataset

Dimension	Indicator	Value
Sample scale	Number of unique communities	1,514,257
	Final number of community–year records	5,482,360
	Average years covered per community	3.62
	Share of original valid observations	89.90%
	Share of model-imputed observations	10.10%
Spatial coverage	Number of cities	36
	Number of city–year samples	288
	Median number of communities in new-tier-1/tier-2/tier-3 cities	50,425 / 27,800 / 19,041
Price distribution	Median price across the full sample (yuan/m²)	22,675
	Price range (yuan/m²)	2,261—267,679
	Average price range in tier-1/tier-3 cities (yuan/m²)	33,800—70,500 / 7,800—12,500

Building on this dataset, community-level labels were further converted into 10 m pixel-level housing-price labels. Because the raw data contain only representative community points rather than true parcel boundaries, a distance-decay weighting method was adopted to generate continuous labels [27]. For each city (c), year (y), and pixel (π_k), the set of valid communities within radius (R) was identified ($N(\pi_k, y)$), and distance-based weights were constructed as follows:

$$\omega_{k,u} = \frac{\exp[-d(\pi_k,u)/\tau]}{\sum_{v\in N(\pi_k,y)} \exp[-d(\pi_k,v)/\tau]} \quad (2)$$

The pixel-level housing price is defined as the weighted average of neighbouring community

prices:

$$p_{\pi_k,y} = \sum_{u \in N(\pi_k,y)} \omega_{k,u} p_{u,y} \quad (3)$$

If no valid community lies within the search radius, the K nearest communities are used for supplementary estimation; if the nearest community still exceeds the maximum tolerance distance, no label is generated for that pixel. This study sets R = 1000 m, τ = 300 m, and K = 5, balancing local price variation against spatial continuity.

3. Methods

3.1. Baseline models

This study employs four widely used regression models as baselines, namely Ridge regression, HistGradientBoosting, Random Forest, and a PyTorch-based multilayer perceptron (TorchMLP)[28], to compare the housing-price predictive performance of linear, tree-ensemble, and neural-network models under AEF-embedding and auxiliary-factor inputs. All models are trained with the $y' = log(1 + y)$-transformed housing price as the supervision target; pixel-level metrics are computed on the log1p scale, while city–year aggregate metrics are computed on the back-transformed original price scale. Model hyperparameters are listed in **Table 4**.

Table 4. Baseline models and hyperparameters

Family	Model	Hyperparameters
Linear	Ridge	($\alpha = 1.0$)； Huber loss ($\delta = 10.0$)
Ensemble	HistGBT	max_depth=6；max_iter=200；learning_rate=0.1；Huber loss ($\delta = 10.0$)
Ensemble	RandomForest	n_estimators=200； max_depth=12； min_samples_leaf=5
Neural	TorchMLP	4-layer MLP((512\rightarrow 256\rightarrow 128\rightarrow 64))； Adam， lr=(1×10^{-3})； early stopping patience=20

3.2. Feature preprocessing

All auxiliary factors were projected to the WGS-84 coordinate system (EPSG:4326) and resampled to a 10 m grid to match the AEF embedding, using bilinear interpolation for continuous variables and nearest-neighbour interpolation for categorical variables. Long-tailed variables such as POI density and building footprint area were $log(1 + x)$-transformed to attenuate the influence of extreme values. The Ridge model was normalised using RobustScaler, tree-ensemble models used the original feature scale directly, and TorchMLP achieved implicit normalisation through batch normalisation (BatchNorm). AEF embeddings were input to the models using their original normalised values.

During full-coverage dense prediction, the output raster was aligned with the inverse-distance weighting (IDW) reference grid (EPSG:4326, approximately 100 m resolution), and features were extracted pixel-by-pixel on this grid to generate the prediction results.

3.3. Experimental strategies

3.3.1. Direct comparison strategies

Using only the 64-dimensional AEF embedding as input, we first compare how the spatiotemporal extent of the training data affects model performance. Four training strategies were designed: single-city single-year training (S1), temporal-expansion training (S2), cross-city transfer training (S3), and global pooled training (S4). All strategies are evaluated on the same target city–year and compared under an equal training-sample size, so as to prevent differences in sample size from confounding the assessment of model performance. Detailed definitions are provided in **Table 5**.

Table 5. Four training strategies

Strategy	Training set	Test set	Description
S1	Target city, target year	Target city, target year	Within-city, within-year baseline
S2	Target city, all years	Target city, target year	Temporal expansion
S3	Other cities, same year	Target city, target year	Zero-shot cross-city transfer
S4	All data except the target city–year	Target city, target year	Global mixed training

For fairness, the training-sample size for S2–S4 was aligned with that of S1, with samples drawn proportionally at random from the pixel counts of the source cities or years. All four baseline models were evaluated in parallel under S1–S4, and the random and spatial protocols were jointly used to select the optimal model $M^{\star}$ and training strategy $S^{\star}$, which together serve as the unified benchmark for the subsequent dimensional analysis, feature-combination experiments, and full-coverage mapping.

3.3.2. Feature importance and principal component analysis

Building on $S^{\star}$ and $M^{\star}$, this section analyses the internal structure of the AEF embedding alone. First, the SHapley Additive exPlanations (SHAP) attribution method is used to compute the relative contribution of each of the 64 AEF dimensions to housing-price prediction[29, 30], and the stability of key dimensions is examined via the rank correlation of importance across city–years. Rank consistency is quantified using Kendall's τ coefficient[31]:

$$\tau = \frac{C - D}{\frac{1}{2}n(n-1)} \tag{4}$$

where n = 64, and C and D denote the number of concordant and discordant pairs between the two rankings, respectively. The closer τ is to 1, the more consistent the key dimensions relied upon across different city–years.

Second, an ablation based on principal component analysis (PCA)[32] was conducted on the AEF embedding, compressing the 64-dimensional representation to $k \in 4,8,16,32,48,64$ dimensions and retraining the model under the same training strategy. This experiment determines whether housing-price prediction depends on a small number of leading principal components or requires a near-complete AEF representation space; PCA performance is further compared against the cumulative explained-variance ratio to distinguish statistical variance retention from the retention of task-relevant information.

3.3.3. Feature combinations

To evaluate the marginal contribution of auxiliary factors, this study progressively concatenates semantically distinct factors onto the AEF embedding. Let the AEF embedding of the i-th pixel be $a_i \in \mathbb{R}^{64}$, and let the k-th auxiliary factor be $f_i^{(k)}$; the input corresponding to feature combination C is then given by:

$$x_i(C) = a_i \,\|_{k \in C}\, f_i^{(k)} \tag{5}$$

where $\|$ denotes vector concatenation. A total of 14 feature combinations (C0–C13) were constructed. C0 serves as the pure-AEF baseline; C1–C9 each add a single auxiliary factor to evaluate its independent marginal contribution; and C10–C13 are multi-factor combinations used to examine the overall gain, synergistic effects, and redundancy arising from jointly inputting multiple sources of information. Detailed configurations are provided in **Table 6**. Specifically, C10 combines WP, NTL, RS, Spatial, and POI to form a classic remote-sensing–socioeconomic combination; C11 adds LST to C10; C12 further adds BH, BS, and LU; and C13 incorporates all nine auxiliary factors.

Table 6. Feature combinations

Code	Composition	Total dim
C0	AEF(baseline)	64
C1	AEF+WP	65
C2	AEF+NTL	82
C3	AEF+RS	67
C4	AEF+Spatial	73
C5	AEF+POI	88
C6	AEF+LST	65
C7	AEF+BH	65
C8	AEF+BS	65
C9	AEF+LU	67
C10	AEF+WP+NTL+RS+Spatial+POI	119
C11	C10+LST	120
C12	C10+BH+BS+LU	124
C13	AEF+all nine auxiliary factors	125

3.4. Evaluation metrics

Model performance is evaluated at four levels: pixel-level accuracy, city–year aggregate accuracy, spatial structure, and temporal consistency. Pixel-level regression metrics include the coefficient of determination R^2, the mean absolute error (MAE), and the Pearson correlation coefficient, all computed on the $log(1+y)$ price scale to remain consistent with the model's training objective. City–year aggregate accuracy is instead computed on the back-transformed, original price scale: for each city c and year y, the median pixel-level predicted price $P_{c,y}$ is compared against the

city-level annual price anchor $A_{c,y}$ obtained from multi-source calibration to compute the city–year aggregate R^2.

Spatial structure is evaluated using global Moran's I, residual Moran's I, and the classification agreement of local indicators of spatial association (LISA)[33]. Global Moran's I quantifies the spatial-clustering intensity of the predicted price surface, while residual Moran's I is computed from the difference between the machine-learning prediction surface and the IDW reference surface to determine whether their discrepancy retains systematic spatial structure. Spatial weights are constructed using a Gaussian kernel with a 1000 m bandwidth, with a maximum of 5,000 randomly sampled points per city–year to constrain computational cost. Temporal consistency is assessed using the structural similarity index (SSIM) between prediction surfaces of adjacent years[34]. To emphasise spatial pattern rather than year-on-year changes in overall price level, each prediction surface is z-score standardised before computing SSIM; where raster extents differ between adjacent years, the later year is resampled to the grid of the earlier year.

All pixel-level regression experiments adopt two evaluation protocols: random and spatial. The random protocol randomly splits each target city–year into training and test sets at an 80:20 ratio and can be regarded as an optimistic upper bound on model performance. The spatial protocol partitions the target area into a 3 × 3 grid of nine spatial blocks, using one block as the test set and the remaining eight as the training set in turn, and takes the median of the nine folds as the reported metric. This protocol reduces train–test leakage caused by spatial autocorrelation[35], more closely approximating the model's generalisation performance on unseen areas. The difference between the two protocols $\Delta R^2 = R^2_{random} - R^2_{spatial}$ is used to diagnose the degree of spatial-autocorrelation leakage.

4. Results

This chapter reports experiments based on 288 city–year samples across the 36 cities over eight years. The primary metric is computed on the $log(1+y)$ scale, namely the pixel-level R^2; the random protocol serves as an optimistic upper bound, and the spatial protocol provides a conservative estimate of spatial generalisation. MAE and Pearson's r serve as auxiliary metrics, while city–year aggregate metrics and full-coverage mapping validation are reported in Section 5.

4.1. Training strategy and model selection

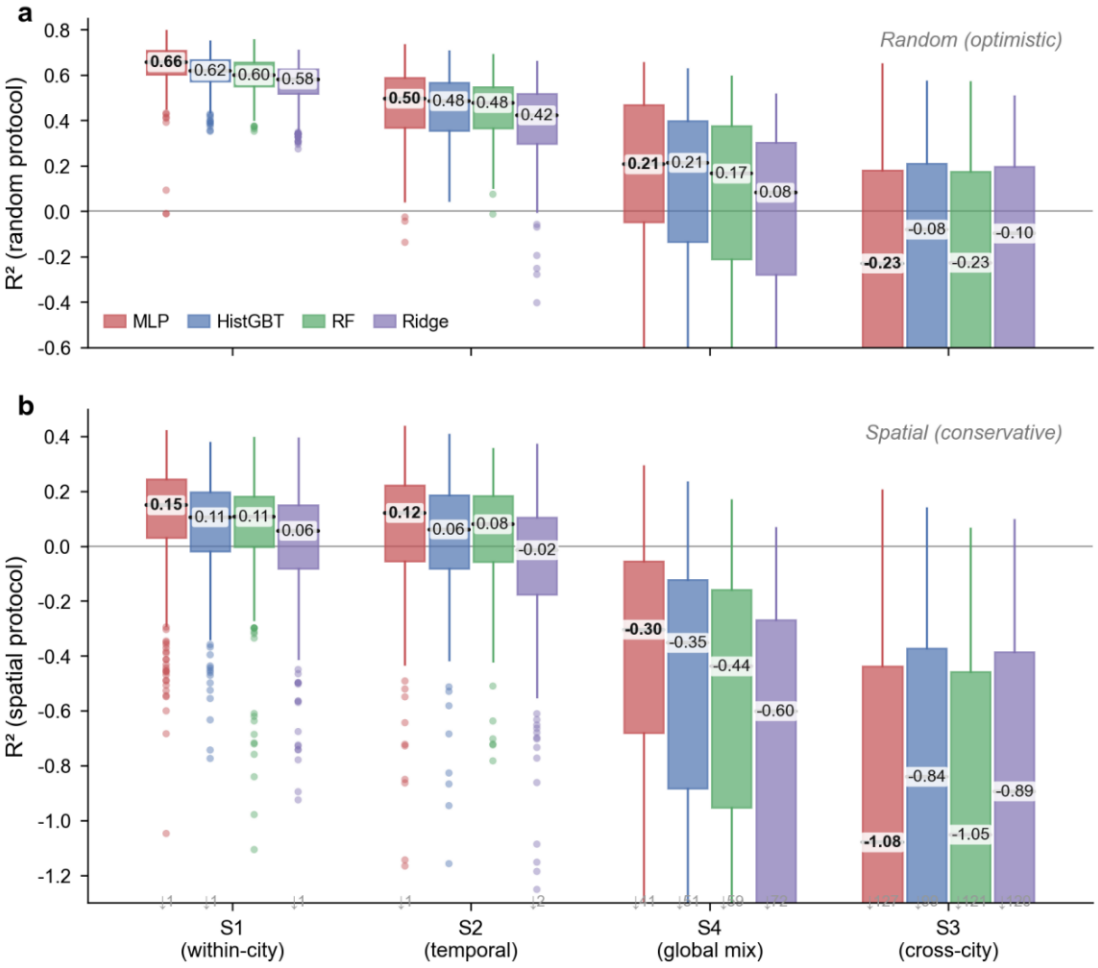

Figure 3. Pixel-level R² (log1p scale) across 288 city–years for four training strategies(S1–S4) and four baseline models. (a) Random 80/20 split (optimistic upper bound). (b) Spatial 3×3 block cross-validation (conservative lower bound). Box: interquartile range; line: median (value annotated with white background; MLP in bold); whiskers:1.5×IQR; dots: outliers beyond whiskers. Values below−1.3 in(b)are clipped, with counts noted at bottom.

Figure 3 shows the distribution of pixel-level R² across the four training strategies and four baseline models under the random and spatial protocols. A substantial gap exists between the two protocols: for S1 × MLP, R^2 falls from 0.66 under random to 0.15 under spatial, indicating that random splitting is strongly affected by spatial autocorrelation and tends to overestimate model performance. Under the more stringent spatial protocol, S1 and S2 retain positive predictive power, with median MLP R^2 values of 0.15 and 0.12, respectively; S4 falls to negative values and S3 performs worst, indicating that pure AEF embeddings struggle to achieve stable cross-city transfer. At the model level, MLP performs best overall across strategies, suggesting that the housing-price signal within the 64-dimensional AEF embedding depends primarily on nonlinear combinations rather than a single dimension or a linear relationship.

Table 7. Median Pearson r (log1p scale) across 288 city–years for four strategies × four models.

Strategy	Model	Pearson_random	Pearson_spatial	Δ (leakage)
S1	**MLP**	**0.811**	**0.555**	**0.256**
S1	HistGBT	0.788	0.548	0.24
S1	RF	0.775	0.54	0.235
S1	Ridge	0.763	0.53	0.233
S2	**MLP**	**0.759**	**0.576**	**0.183**
S2	HistGBT	0.741	0.555	0.186
S2	RF	0.728	0.545	0.182
S2	Ridge	0.726	0.534	0.192
S4	**MLP**	**0.667**	**0.457**	**0.211**
S4	HistGBT	0.658	0.46	0.199
S4	RF	0.649	0.457	0.192
S4	Ridge	0.634	0.449	0.185
S3	**MLP**	**0.642**	**0.431**	**0.212**
S3	HistGBT	0.628	0.434	0.194
S3	RF	0.606	0.423	0.183
S3	Ridge	0.606	0.432	0.174

Table 7 further shows that the decline in Pearson's r across strategies is smaller than that of R^2: under the spatial protocol, Pearson's r for S1 × MLP is 0.555, while S3 × MLP still reaches 0.431. This indicates that AEF can partially capture the relative ranking of high and low prices within a city even in cross-city settings, but struggles to independently calibrate absolute price levels across different cities. This study therefore adopts S1 as the training strategy and MLP as the baseline model for subsequent experiments, analysing the AEF dimensional structure and the gains from auxiliary factors on this basis.

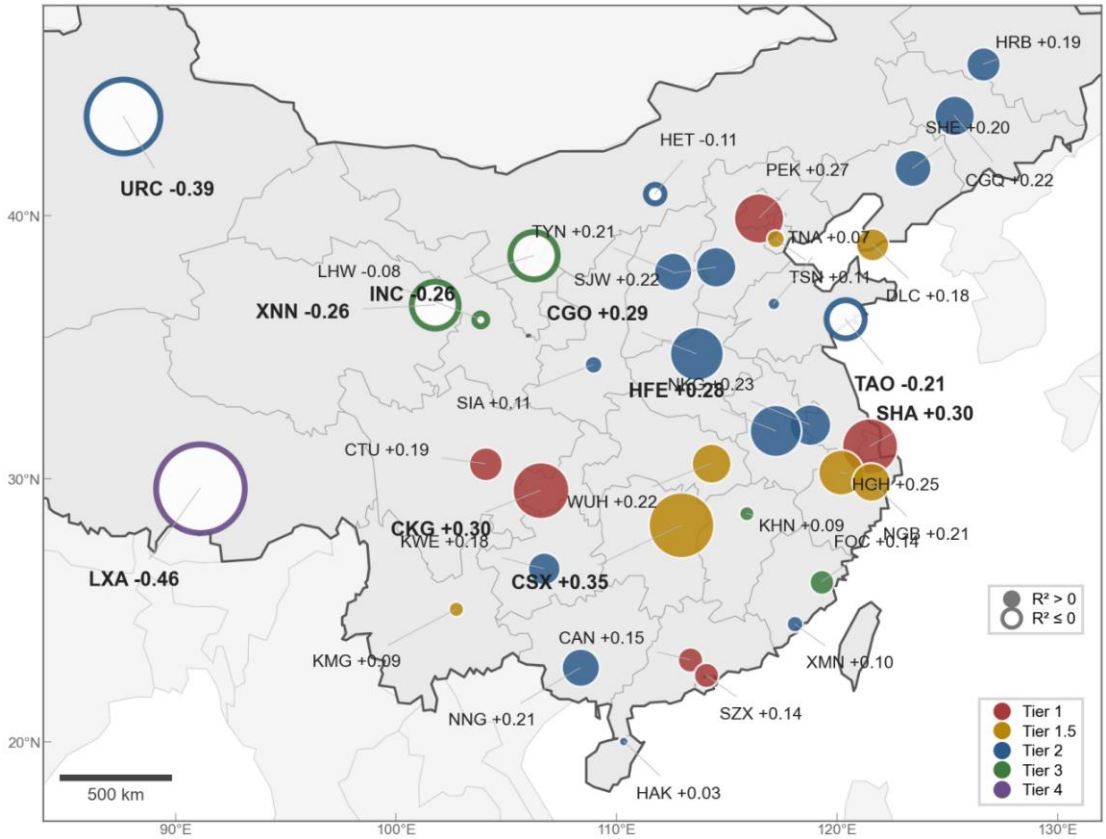

Figure 4. Median pixel-level R² (log1p scale, spatial protocol) for 36 cities under S1×MLP. Bubble area is proportional to $|R^2|$ (power-scaled); filled circles denote $R^2 > 0$, open circles $R^2 \leq 0$. Colors indicate city tier (map lines delineate study areas and do not necessarily depict accepted national boundaries).

Figure 4 shows the spatial distribution of R^2 for S1 × MLP across the 36 cities. The predictive power of pure AEF varies markedly across cities: most cities exhibit positive R^2, but overall values remain low, with Changsha reaching the highest value (0.35). Tier-1 cities such as Shenzhen and Guangzhou are not markedly superior to some tier-2 cities, indicating that AEF performance does not simply increase with city tier. Cities with negative values are concentrated mainly in the Northwest and plateau regions, such as Lhasa, Urumqi, Xining, and Yinchuan, suggesting that surface embeddings alone struggle to establish stable price discrimination in cities with smaller built-up areas, limited morphological variation, or weak price variability. Overall, AEF provides some housing-price-relevant representation, but using it alone remains insufficient to support robust pixel-level housing-price mapping.

4.2. Dimensional structure of AEF

This section analyses the effective dimensionality, importance distribution, and cross-city stability of the 64-dimensional embedding under $S^\star = S1$, $M^\star = MLP$, and pure AEF input (C0).

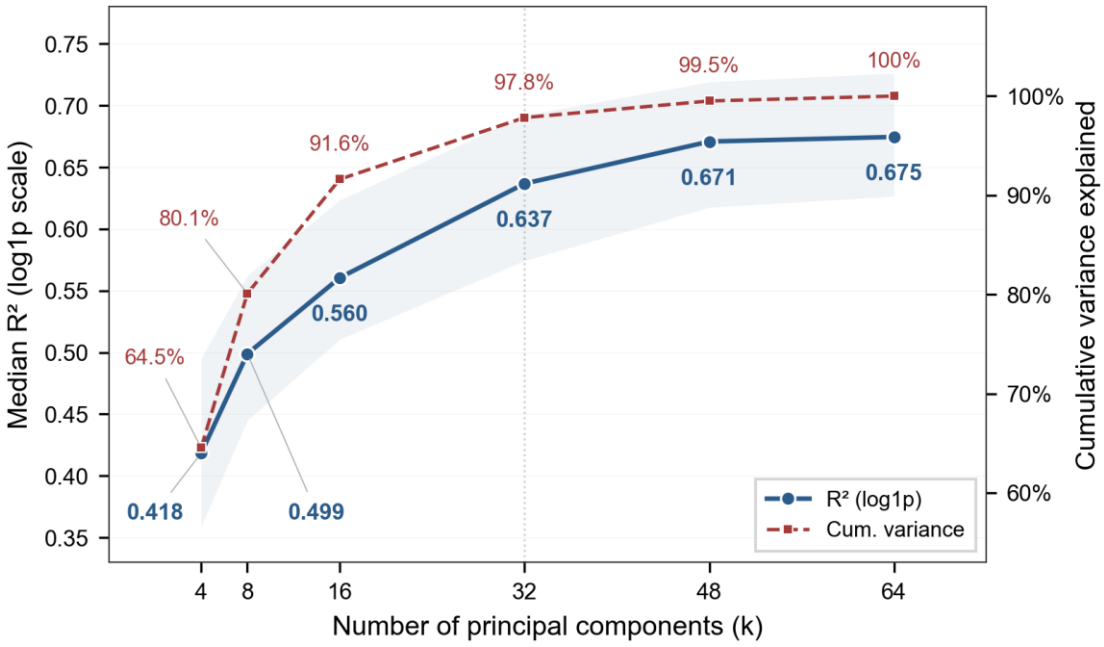

Figure 5. PCA ablation on AEF 64-dimensional embeddings under S1×MLP (random protocol). Left axis: median R^2 (log1p scale) across 288 city–years; shaded band shows interquartile range. Right axis: cumulative variance explained. Dashed vertical line marks k = 32.

Figure 5 shows that model performance under PCA ablation improves as the number of principal components increases: median R^2 is 0.418 at k = 4, rising to 0.637 at k = 32 and 0.675 at k = 64. Although the first 32 components already explain 97.8% of the statistical variance, full-dimensional performance is not fully recovered, indicating that housing-price-relevant information is not confined to high-variance components alone but is partly distributed across tail dimensions. Performance plateaus beyond k = 48, suggesting that housing-price prediction requires a near-complete AEF representation space.

The SHAP results further support this conclusion. **Figure 6** shows that dim_32, dim_37, dim_46, dim_10, and dim_44 are the dimensions with the highest global contributions, although overall contribution is not concentrated in a small number of dimensions but is instead dispersed across multiple embedding channels. More importantly, a clear "red–blue separation" structure is observed across many dimensions, whereby high values (red) and low values (blue) correspond to positive and negative contributions to the model output, respectively; however, the ordering of this positive–negative effect is not consistent across dimensions—some dimensions exhibit a monotonic "high value → positive contribution" relationship, while others show a "mid-range reversal" or bidirectional effect. This structural heterogeneity indicates that different AEF dimensions do not act on housing price through a unified mechanism but instead contribute to prediction through distinct nonlinear mappings.

At the city level, the spatial distribution of key dimensions exhibits pronounced city specificity. In **Figure 7**, each row represents the normalised 64-dimensional SHAP importance for one city; if AEF contained universal, cross-city "backbone dimensions" for housing price, this would manifest as several columns remaining consistently highlighted across most cities, with the Top-3 markers concentrated on similar dimensions. Instead, the white markers in the figure are highly dispersed across all 64 dimensions, with no stable columnar clustering emerging across regions or city tiers,

indicating that different cities rely on different combinations of AEF dimensions. **Figure 8** further confirms this from the perspective of rank correlation: across 41,328 city–year pairs, the median Kendall's τ is −0.001, with the distribution centred close to zero; the histograms for within-city cross-year pairs and cross-city pairs overlap substantially, and no discernible regional or city-tier block structure emerges in the 36 × 36 city-pairing matrix. This indicates that AEF dimension importance lacks not only cross-city stability but also a stable ranking of key dimensions across years within the same city.

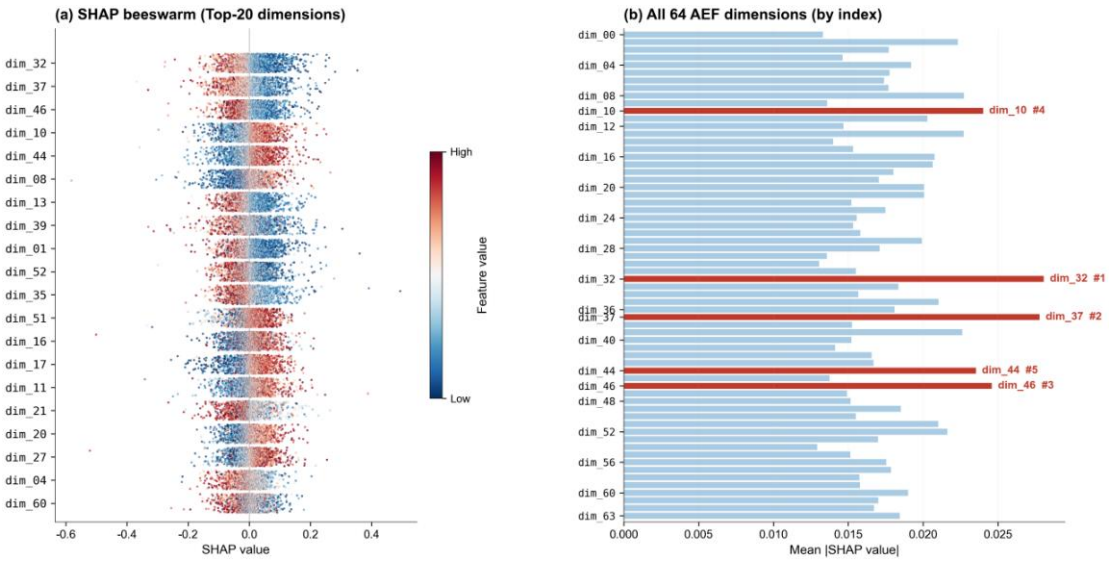

Figure 6. SHAP analysis of AEF embedding dimensions under S1×MLP across 288 city–years. (a) Beeswarm plot of the Top-20 dimensions ranked by global mean |SHAP value| (5,000 random samples); color encodes normalized feature value. (b) Mean |SHAP value| for all 64 dimensions arranged by index; Top-5 highlighted in red with global rank annotated.

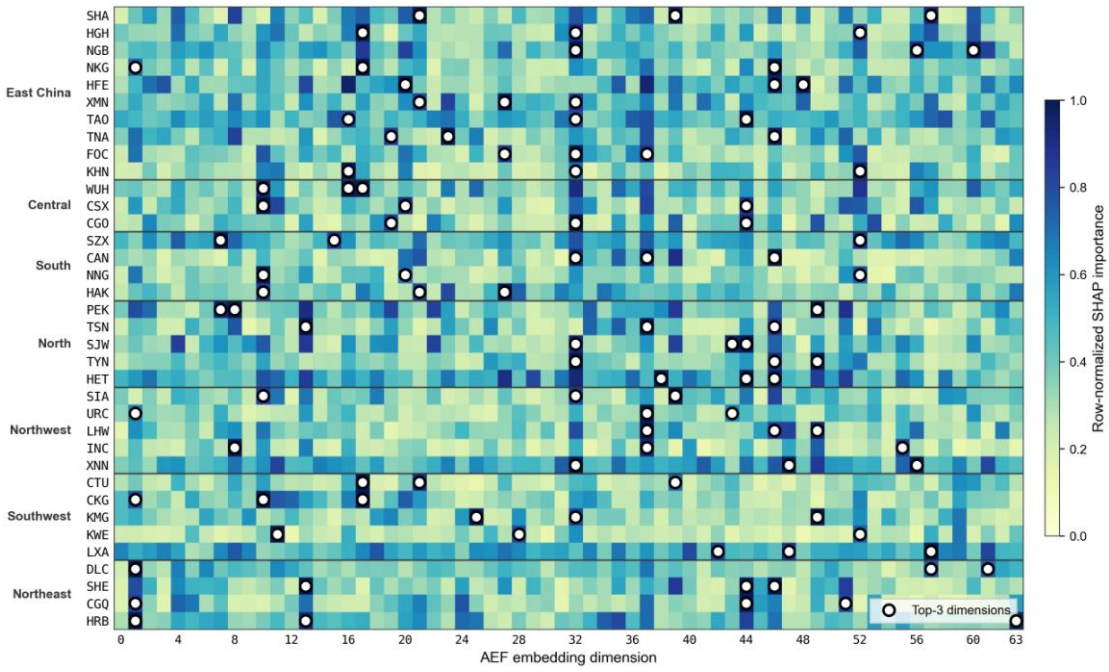

Figure 7. Row-normalized SHAP importance heatmap for 36 cities (rows) across 64 AEF dimensions (columns). Each row is scaled to its own maximum. White circles mark each city's

Top-3 dimensions. Cities grouped by region; black lines denote region boundaries.

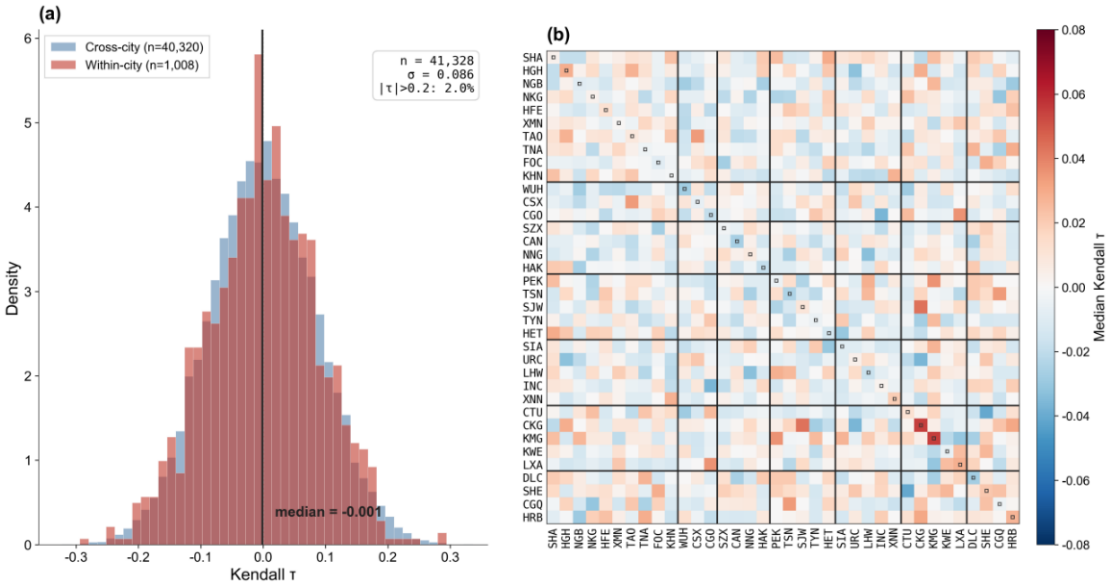

Figure 8. Cross-city stability of SHAP dimension rankings. (a)Distribution of pairwise Kendall τ across 41,328 city–year pairs; blue: cross-city, red: within-city cross-year. (b)Median Kendall τ between each pair of 36 cities (averaged over all year combinations); cities grouped by region. Color centered at zero (±0.08 range); diagonal squares mark within-city entries.

4.3. Effects of stacking auxiliary factors

This section evaluates the gains in housing-price prediction from the nine auxiliary factors and their combinations, building on $S^{\star} = S1$, $M^{\star} = MLP$, and the full 64-dimensional AEF embedding.

Table 8. Marginal contribution of individual auxiliary factors to AEF-based housing price prediction (S1×MLP, spatial protocol). ΔR² denotes the gain relative to C0 (AEF only).

Code	Factor	Dim	R² (spatial)	ΔR²	R² (random)	MAE (spatial)	Pearson (spatial)
C0	AEF only	64	0.149	—	0.656	0.219	0.555
C5	+ POI	88	0.449	0.3	0.812	0.169	0.723
C4	+ Spatial	73	0.31	0.161	0.754	0.198	0.643
C2	+ NTL	82	0.234	0.085	0.717	0.206	0.604
C1	+ WP	65	0.208	0.059	0.674	0.212	0.586
C6	+ LST	65	0.186	0.036	0.671	0.215	0.573
C7	+ BH	65	0.182	0.033	0.67	0.215	0.58
C3	+ RS	67	0.171	0.022	0.662	0.218	0.564
C9	+ LU	67	0.163	0.014	0.662	0.216	0.566
C8	+ BS	65	0.157	0.008	0.661	0.218	0.568

Table 8 shows that every individual auxiliary factor improves spatial R^2, although the magnitude of the contribution varies markedly. POI yields the largest gain (+0.300), followed by Spatial (+0.161),

indicating that public services, commercial amenities, and transport accessibility represent key semantics that AEF alone cannot fully capture. NTL and WP form a second tier, improving performance by +0.085 and +0.059, respectively. In contrast, the gains from LST, BH, RS, LU, and BS do not exceed +0.036; the comparatively weak contribution of explicit physical built-form features such as building height, building footprint area, and land use suggests that AEF already implicitly encodes a substantial share of built-environment information.

Table 9. Performance of stacked feature combinations (S1×MLP). C10 achieves the highest spatial R^2 among all configurations; adding further factors yields no improvement. ΔR^2 values are computed on the spatial protocol.

Code	Composition	Dim	R^2 (random)	R^2 (spatial)	ΔR^2 vs C0	ΔR^2 vs C10
C0	AEF only	64	0.656	0.149	—	−0.341
C10	AEF + WP + NTL + RS + Spatial + POI	119	0.844	0.49	0.341	—
C11	C10 + LST	120	0.838	0.481	0.332	−0.009
C12	C10 + BH + BS + LU	124	0.837	0.473	0.324	−0.017
C13	AEF + all 9 factors	125	0.842	0.48	0.331	−0.010

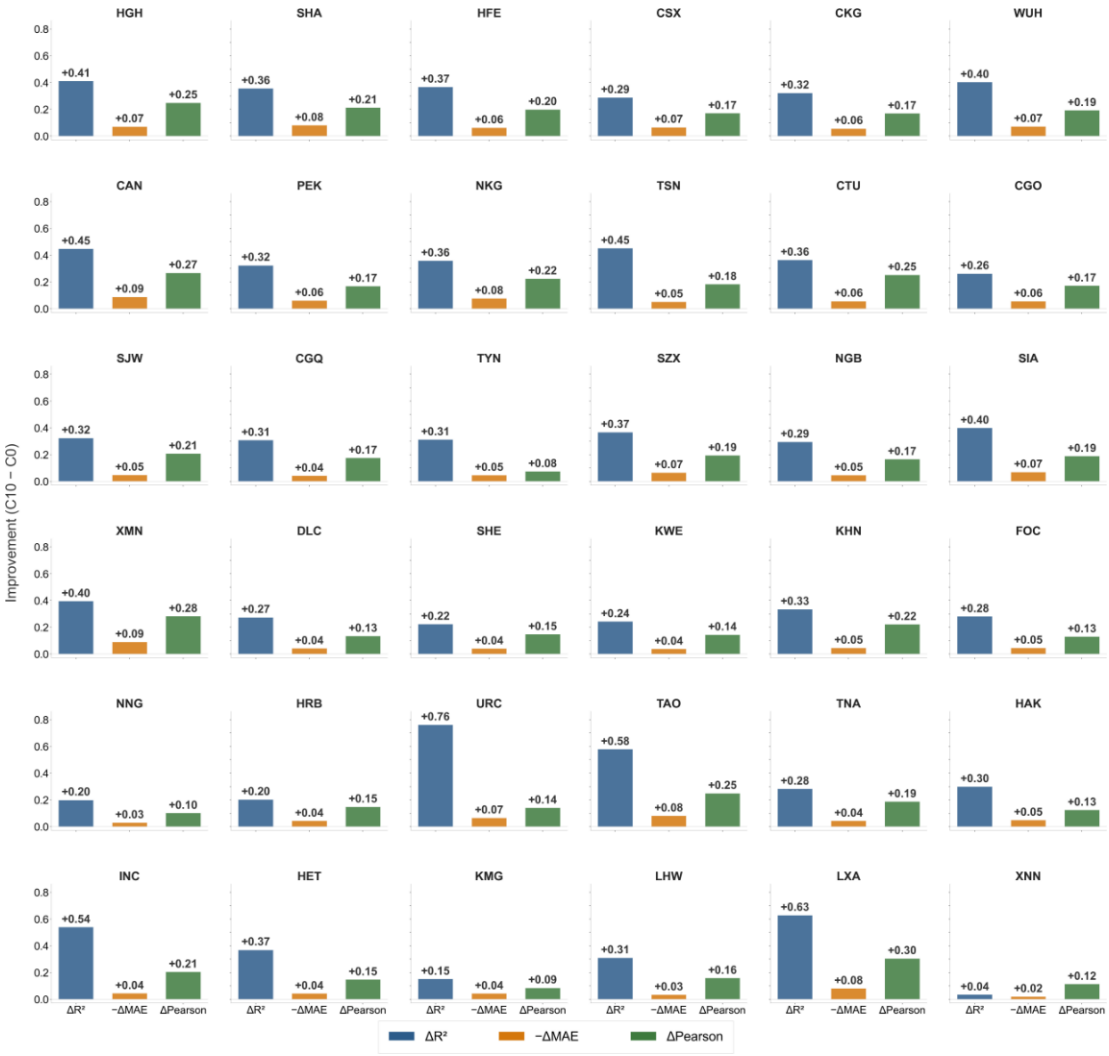

Figure 9. Per-city improvement from C0 (AEF only) to C10 (AEF + Classic 5) across three metrics under S1×MLP spatial protocol. Each panel shows one of 36 cities, sorted by C10 R^2 descending. Blue: ΔR^2; orange: −ΔMAE (positive indicates MAE reduction); green: $\Delta Pearson\, r$. All metrics are computed in log1p scale.

Table 9 further compares the performance of multi-factor combinations. C10, composed of WP, NTL, RS, Spatial, and POI, achieves the best results, with spatial R^2 rising from 0.149 (C0) to 0.490 and random R^2 rising from 0.656 to 0.844. Adding LST, building physical factors, or all nine auxiliary factors yields no further improvement; spatial R^2 in fact declines slightly, indicating that C10 is already close to information saturation and that additional variables mainly introduce redundancy or noise. C10 is therefore adopted as the optimal feature configuration for the subsequent full-coverage inference.

Figure 9 shows the city-by-city improvement from C0 to C10. R^2 and Pearson's r increase, and MAE decreases, in every city, indicating that the gain from auxiliary factors is universal. The cities with the largest gains are mostly those where pure AEF performs poorly, such as Urumqi, Lhasa, Qingdao, and Yinchuan, where ΔR^2 exceeds 0.5; cities that already showed reasonable predictive power under C0 show comparatively smaller marginal improvements. This result indicates that auxiliary factors are particularly effective at compensating for AEF's representational shortfalls in areas with limited morphological variation or invisible socioeconomic signals.

5. Full-Coverage Inference and Spatial Validation

5.1. Full-coverage dense mapping

Based on the optimal configuration ($S^\star = S1,\ M^\star = MLP,\ C^\star = C10$), full-coverage dense inference was performed for all 288 city–year units across the 36 cities and eight years. For each city–year, the model was trained on that city's valid anchors for the corresponding year, and C10 features were extracted on a roughly 100 m grid consistent with the IDW reference raster to generate a continuous housing-price prediction surface.

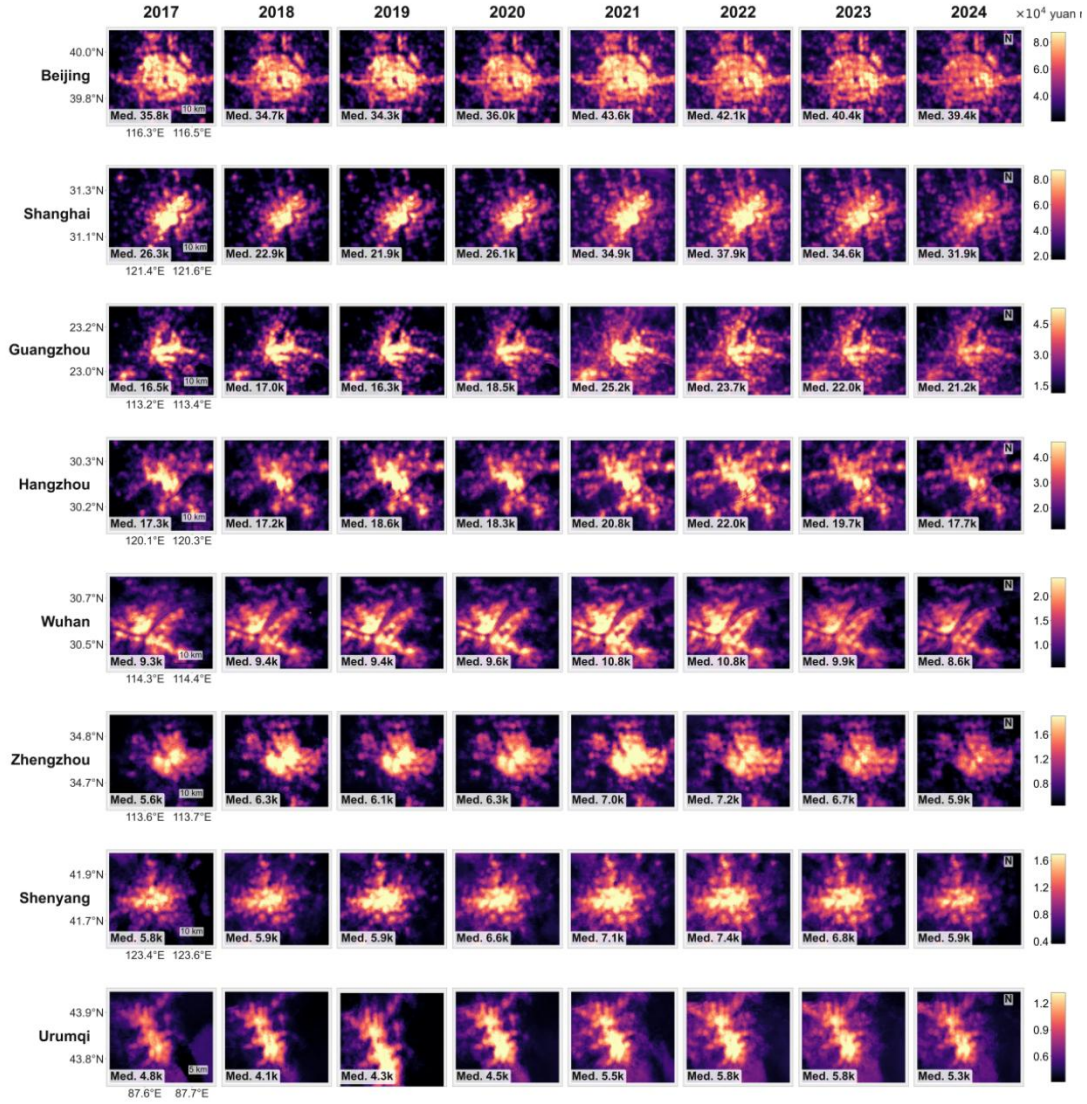

Figure 10. Pixel-level predicted housing prices for eight representative cities from 2017 to 2024, estimated using S1×MLP with C10 features at approximately 100 m resolution. Rows correspond to cities and columns to years. Colours show predicted price using city-specific 2nd–98th percentile normalization and are therefore comparable across years within each city row. Colour bars are reported in $\times 10^4$ yuan m^{-2}. Panel labels indicate annual median predicted price in $\times 10^4$ yuan m^{-2}. Grey pixels indicate non-built-up or no-prediction areas.

Figure 10 presents the full-coverage prediction results for eight representative cities from 2017 to 2024. Predicted prices maintain a clear gradient across cities: the annual median price is approximately 35,000–44,000 yuan m^{-2} in Beijing and 22,000–39,000 yuan m^{-2} in Shanghai, whereas Zhengzhou, Shenyang, and Urumqi mostly fall within 4,000–8,000 yuan m^{-2}, a near order-of-magnitude gap between the highest- and lowest-priced cities. Temporally, most cities exhibit a common cyclical pattern: relatively stable prices in 2017–2019, a marked upturn in 2020–2021, and a gradual decline thereafter. For example, Beijing's median price rose from 34,300 yuan m^{-2} in 2019 to 43,600 yuan m^{-2} in 2021 before falling to 39,400 yuan m^{-2} in 2024, while Shanghai's rose from 21,900 yuan m^{-2} in 2019 to 37,900 yuan m^{-2} in 2022 before declining to 31,900 yuan m^{-2} in 2024. Spatially, the location of each city's high-price core remains largely stable over the eight years,

indicating that the prediction surfaces primarily capture the persistent structure of intra-urban price gradients, with year-to-year variation reflecting overall shifts or corrections in price level rather than a restructuring of spatial pattern.

It should be noted that the approximately 100 m output resolution smooths local price discontinuities to some extent, so that the sharp boundary between a central-business-district (CBD) core and its surroundings may appear as a gradual transition; localised noise patches are also present in a few city–years, potentially related to AEF coverage quality or year-to-year fluctuations in auxiliary factors such as POI. Nevertheless, **Figure 10** demonstrates overall that the AEF + C10 model can generate temporally continuous, spatially stable urban housing-price prediction surfaces.

5.2. Spatial validation

Figure 11 presents the reference price surface, the C10 prediction surface, and the spatial residuals for 36 cities in 2020. Overall, the prediction surfaces successfully reproduce the intra-urban price tiers and the spatial ranking of high and low prices. Pixel-level metrics support this finding: across the 36 cities, Pearson's r ranges from 0.864 to 0.964 (median 0.948), and R^2 ranges from 0.380 to 0.919 (median 0.825; **Table 10**). Cities with the strongest performance—including Taiyuan, Urumqi, Haikou, and Hohhot—all achieve R^2 above 0.89, whereas cities such as Qingdao, Chengdu, Chongqing, and Ningbo perform comparatively worse, mainly owing to steep price gradients or complex urban spatial morphology. The residual surfaces generally exhibit a mean-reversion pattern, with underestimation in central areas and overestimation in peripheral areas, indicating that the model recovers the dominant spatial pattern but remains comparatively conservative in capturing the magnitude of extreme high- and low-price areas.

At the city–year aggregate level, the Pearson's r between the median predicted price and the multi-source calibrated price anchor across the 288 city–years reaches 0.988, with R^2 of 0.617, indicating that the model preserves the overall price ranking and temporal trends across cities and years. This result shows that the C10 dense prediction surfaces capture not only intra-urban spatial structure but also remain consistent with city-level market benchmarks.

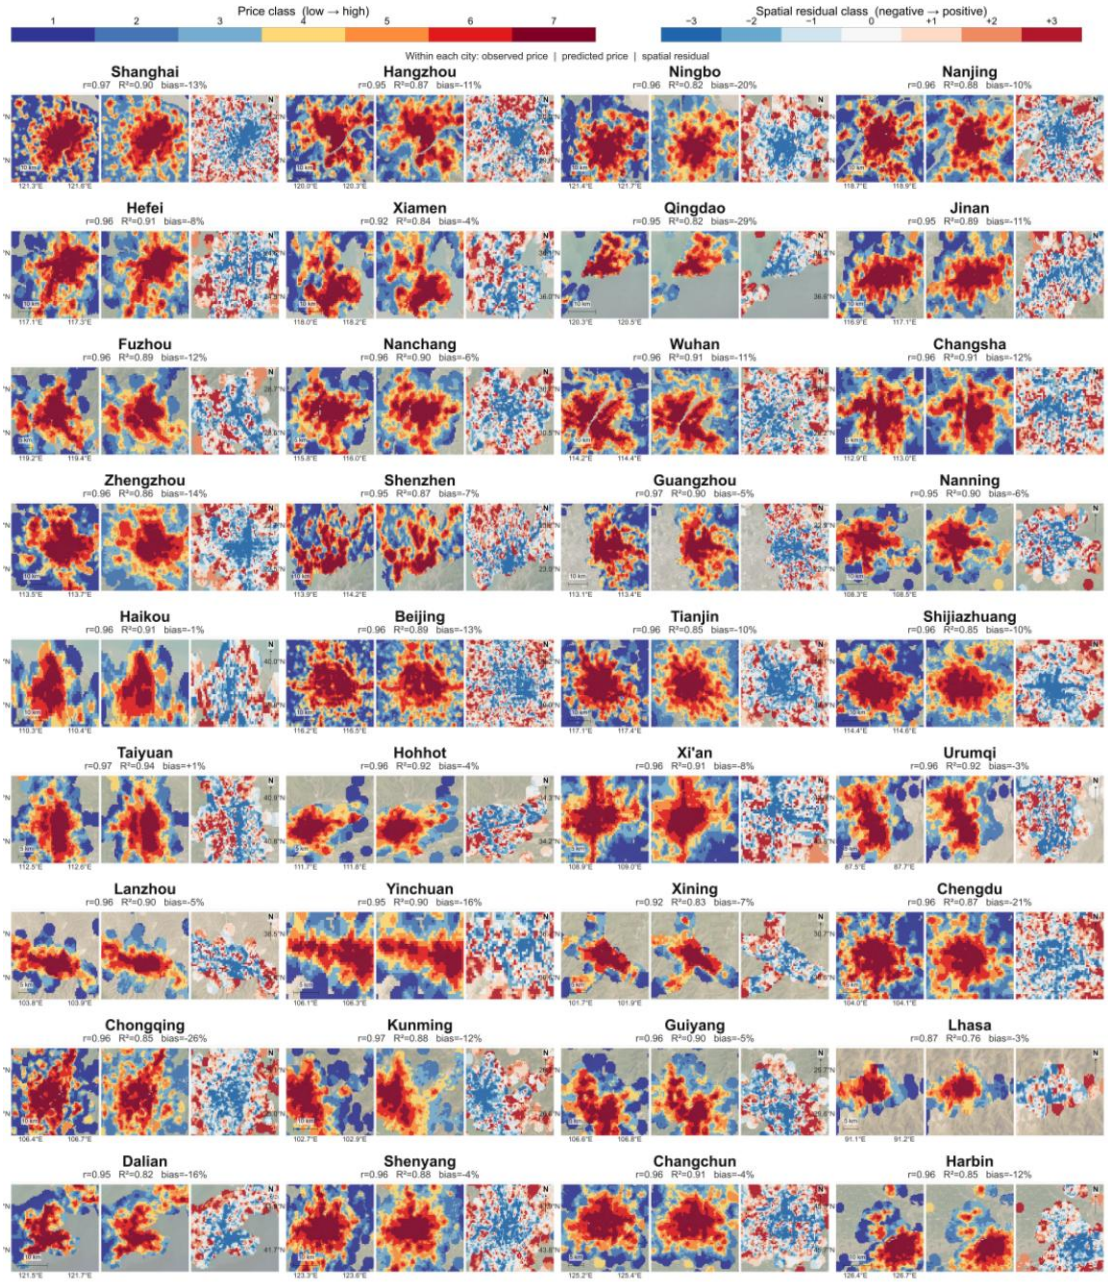

Figure 11. Pixel-level housing price mapping across 36 Chinese cities (2020). Each city panel shows three 0.5 km grid maps: observed reference surface (left), bias-corrected C10 dense prediction (centre), and spatial residual (right). Colour scales follow quantile-based classification within each city. Road network contours and ESRI satellite imagery are overlaid for geographic reference. City-level Pearson r, R^2, and original bias are annotated above each panel.

Local spatial-clustering results further confirm the structural reliability of the prediction surfaces. The median cell-wise LISA classification agreement between the C10 prediction surface and the IDW reference surface is 84.8% (IQR: 82.4%–86.3%), and global Moran's I is also highly consistent between the two (0.865 vs. 0.859). **Figure 12** shows that, in the six cities with the highest agreement, the High–High hotspots and Low–Low coldspots of the prediction surface largely coincide with those of the IDW reference, with discrepancies concentrated mainly in the high–low price transition zones. The median residual Moran's I is 0.493, indicating that the discrepancy between the two surfaces is not random noise but stems primarily from the model's systematic compression of high-price core

areas and low-price peripheral areas. Overall, the C10 prediction surfaces maintain a high degree of consistency with the reference data in terms of spatial ranking, clustering structure, and city-level trends, making them a viable alternative to IDW interpolation and capable of providing continuous estimates in sample-sparse areas.

Table 10. Per-city spatial agreement between C10 dense prediction and reference surface (2020, pixel-level, original scale)

City	Pearson r	R^2	City	Pearson r	R^2	City	Pearson r	R^2
Shanghai	0.964	0.841	Wuhan	0.952	0.846	Beijing	0.951	0.802
Kunming	0.962	0.83	Changsha	0.943	0.803	Tianjin	0.954	0.802
Taiyuan	0.96	0.919	Hangzhou	0.944	0.798	Shijiazhuang	0.947	0.802
Urumqi	0.954	0.904	Nanjing	0.945	0.828	Xi'an	0.944	0.843
Guangzhou	0.956	0.879	Jinan	0.937	0.797	Hefei	0.95	0.857
Zhengzhou	0.956	0.789	Fuzhou	0.946	0.81	Shenzhen	0.937	0.822
Changchun	0.95	0.882	Harbin	0.951	0.795	Xining	0.901	0.754
Haikou	0.947	0.892	Nanchang	0.95	0.874	Xiamen	0.909	0.814
Hohhot	0.951	0.896	Nanning	0.945	0.875	Lhasa	0.864	0.736
Shenyang	0.948	0.862	Lanzhou	0.945	0.868	Dalian	0.941	0.705
Guiyang	0.947	0.869	Yinchuan	0.942	0.665	Ningbo	0.945	0.629
Chongqing	0.948	0.625	Chengdu	0.94	0.547	Qingdao	0.943	0.38

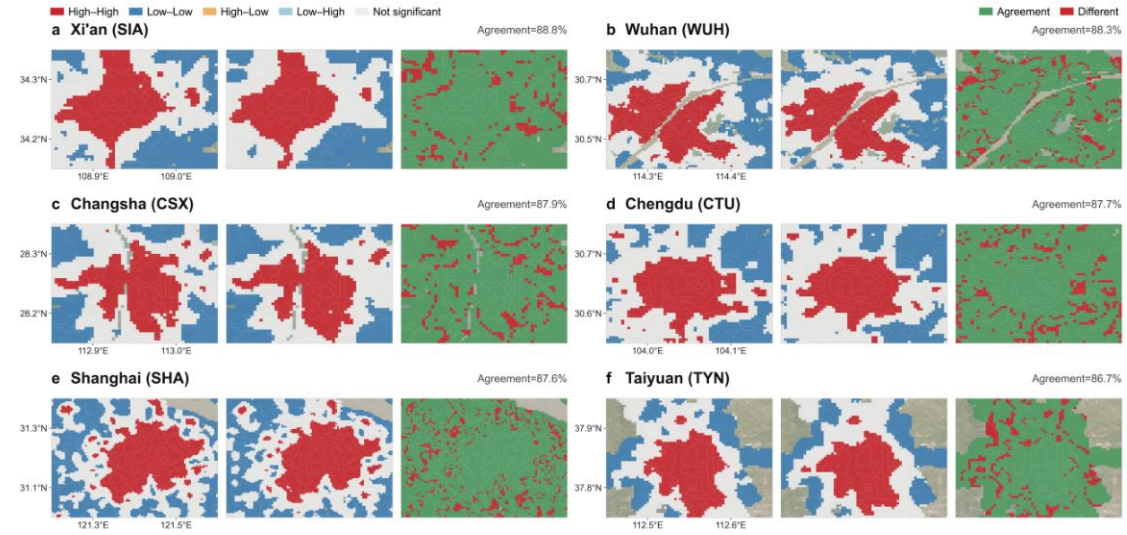

Figure 12. Spatial agreement between observed and predicted Local Moran's I (LISA) cluster patterns in six representative cities (2020). Local Moran's I (LISA) cluster maps derived from the IDW interpolation (left) and the model prediction (middle), with the corresponding spatial agreement maps (right), for the six cities exhibiting the highest agreement in 2020. High–High, Low–Low, High–Low, and Low–High clusters are shown in red, blue, orange, and light blue, respectively; non-significant locations are shown in grey. Agreement maps indicate matching LISA cluster classifications in green and mismatched classifications in red. The Esri World Gray Canvas basemap is included for geographic reference.

5.3. Temporal consistency validation

To examine whether the dense prediction surfaces are suitable for long-term monitoring, structural SSIM was computed between C10 prediction surfaces of adjacent years for the 36 cities. Each prediction surface was z-score standardised before computation to remove changes in overall price level and isolate spatial pattern alone, yielding 252 adjacent-year pairs in total.

Figure 13 shows that the prediction surfaces exhibit high temporal stability: the median SSIM across the 252 adjacent-year pairs is 0.986 (IQR: 0.977–0.992), with the vast majority exceeding 0.95. The median SSIM for each of the seven adjacent-year intervals remains no lower than 0.977; even during price upswings such as 2019–2020 and 2020–2021, spatial structure remains stable, indicating that year-to-year change is expressed mainly as an adjustment in overall price level rather than a restructuring of the intra-urban high–low price pattern. At the city level, SSIM remains above 0.98 for most cities, with the exception of a small number of cities such as Shenzhen and Xiamen. Combined with the spatial-validation results in Section 5.2, the AEF + C10 prediction surfaces exhibit both spatial-pattern reliability and temporal continuity, supporting their use for long-term urban housing-price monitoring.

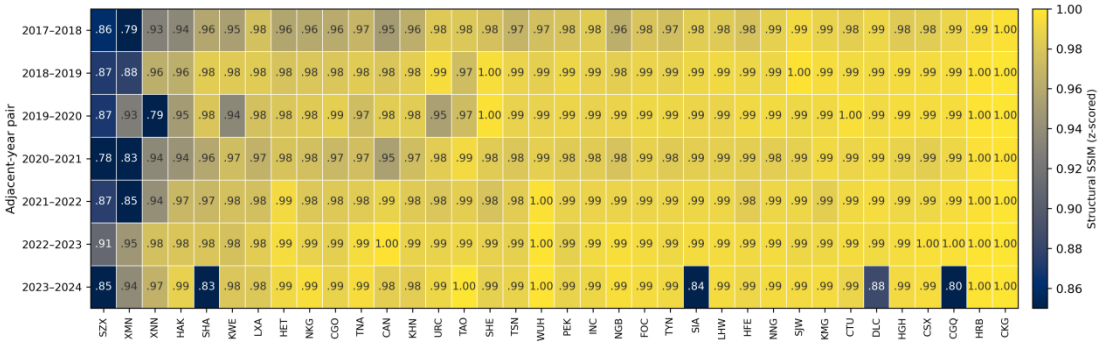

Figure 13. Year-to-year structural stability of C10 dense predictions across 36 cities. Each cell shows the structural SSIM between adjacent-year predicted price surfaces for one city (columns, ordered by median SSIM) and one year pair (rows); SSIM is computed after per-surface z-score normalization. Colour encodes SSIM from low (dark blue) to high (yellow).

5.4. Downstream application cases

5.4.1. Temporal analysis of housing prices

Figure 14 illustrates the application of continuous housing-price prediction surfaces to the spatiotemporal analysis of housing-market evolution. The price index for all six representative cities peaks around 2021 before entering a correction phase, although the magnitude of the correction differs markedly: Shenzhen and Guangzhou exhibit the strongest fluctuations, while Beijing, Wuhan, and Changsha are comparatively stable, indicating that different cities respond inconsistently to the same market cycle. The 2021–2024 tier-transition matrix shows that most spatial units retain their original price tier, with transitions occurring mainly between adjacent tiers, indicating that the overall urban housing-price pattern remains stable during the correction period while changing only incrementally at the local level. Spatially, the persistently high-price zones in Beijing, Shanghai, Guangzhou, Wuhan, and Changsha are concentrated mainly in the urban core, whereas correction and upgrade zones are located more towards the periphery; Shenzhen instead exhibits a more fragmented

transition pattern, reflecting its stronger spatial heterogeneity. This case demonstrates that continuous prediction surfaces can support not only accuracy evaluation but also the monitoring of housing-market evolution, spatial risk identification, and policy assessment.

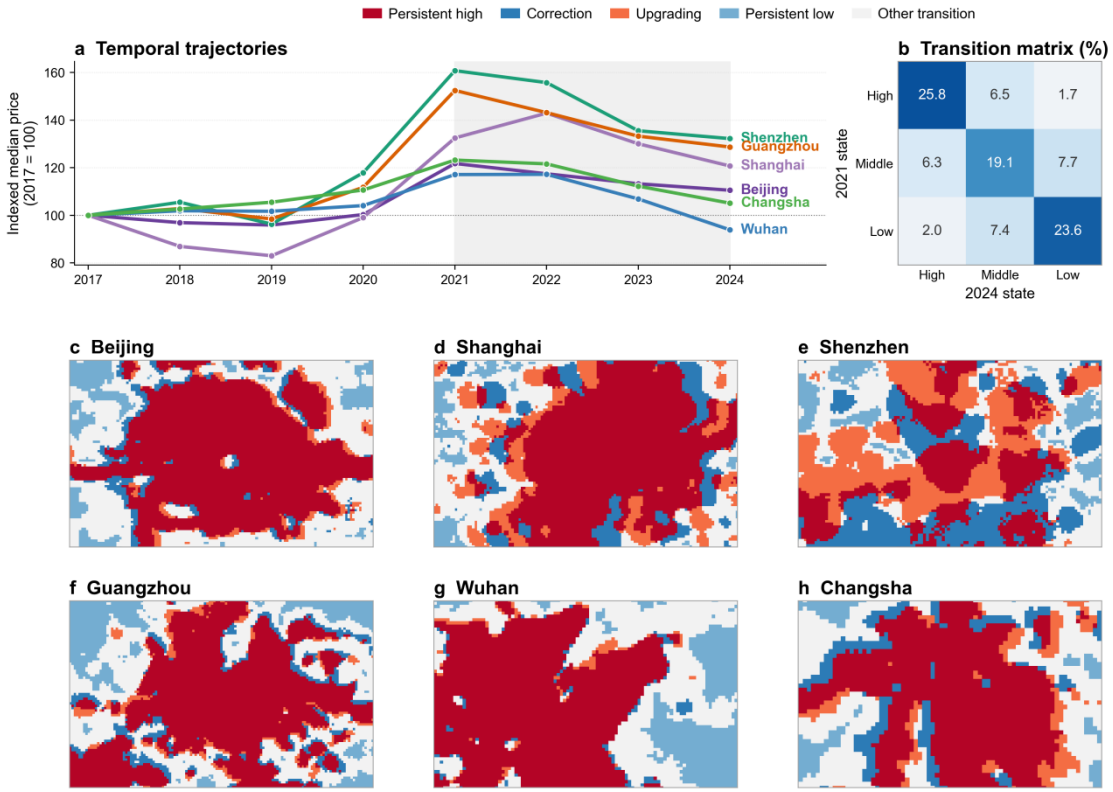

Figure 14. Spatiotemporal regime analysis of predicted housing prices, demonstrating a representative downstream application of the proposed framework by integrating long-term housing-price prediction with spatiotemporal regime analysis. Rather than focusing solely on prediction accuracy, this visualization illustrates how continuous annual prediction surfaces can be transformed into interpretable temporal trajectories and spatial transition patterns, thereby supporting the identification of long-term market dynamics.

5.4.2. Temporal evolution of the CBD premium gradient

To illustrate the value of continuous housing-price prediction surfaces for economic-geographic analysis, this study further examines the spatiotemporal characteristics of price variation with distance from the city centre, using the CBD price gradient as an example. **Figure 15** shows that the 2020 prediction surfaces recover the typical centre–periphery price structure: Beijing and Shanghai exhibit pronounced high-price cores with multi-centre expansion zones, Wuhan's high-price zone is relatively contained, and Urumqi shows an overall weak price gradient.

Regression of median price against CBD-distance rings shows a significant logarithmic distance-decay relationship in all four cities (**Figure 16**). The 2024 rent-gradient elasticity β ranges from −0.035 to −0.057, with model fit R^2 of 0.88–0.99, all statistically significant, indicating that the prediction surfaces capture genuine spatial price gradients consistent with urban-economic theory rather than local noise. The curve shapes for 2017, 2020, and 2024 are broadly similar, indicating strong temporal stability of the intra-urban price gradient.

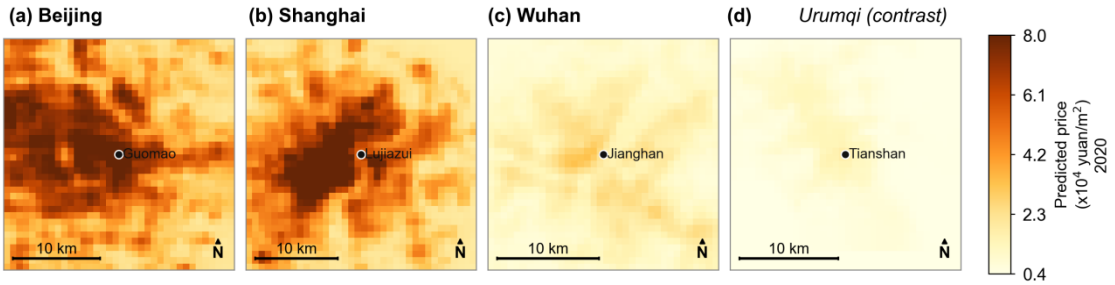

Figure 15. Predicted intra-urban price surfaces (2020) recover the polycentric rent structure of Chinese cities. Pixel-level predictions from the AEF-C10 model for four representative cities, each centred on its principal CBD and rendered at the city's own spatial extent (scale bars therefore differ); colour denotes predicted price on a shared scale. Beijing and Shanghai display pronounced high-value cores radiating into multi-nodal plateaus, whereas Wuhan and the monocentric contrast city Urumqi remain comparatively flat.

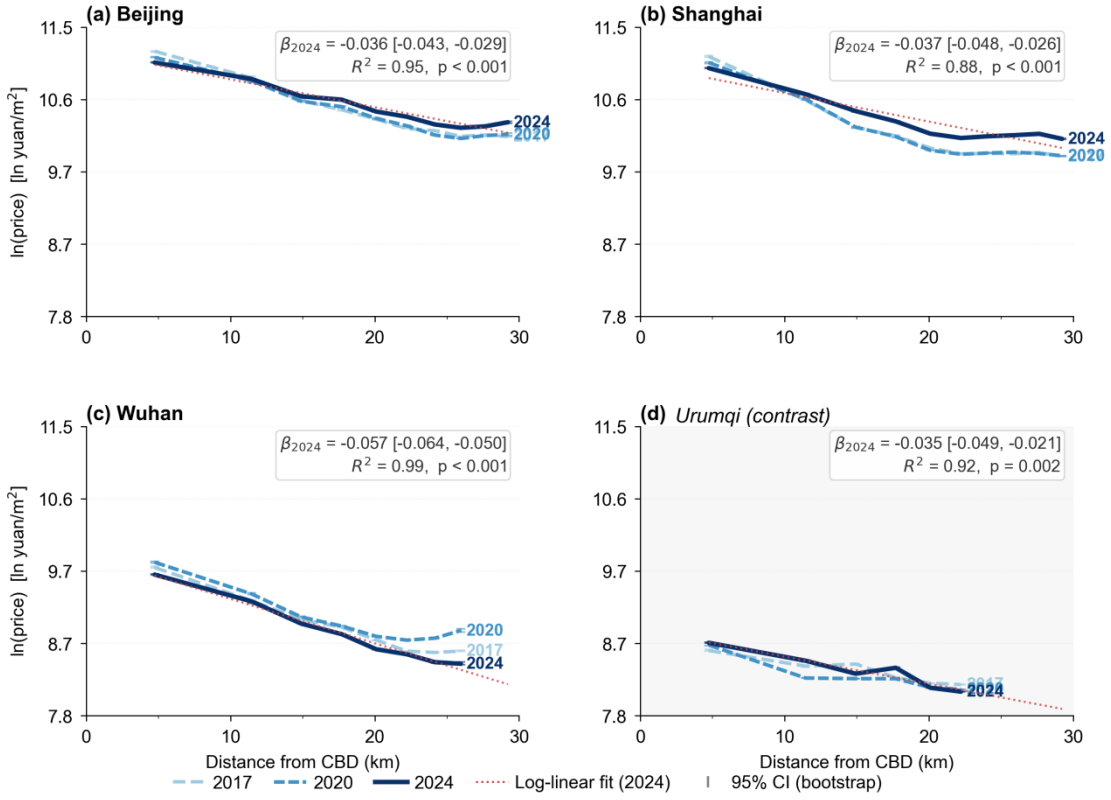

Figure 16. Log-linear distance-decay gradients are stable across cities and summarised by a single rent-gradient elasticity β. Median predicted price (ln scale, shared y-axis) against distance from the CBD for 2017, 2020 and 2024; dotted lines give the 2024 log-linear fit with bootstrap 95% confidence intervals. Estimated gradients (β_{2024}) range from −0.035 to −0.057 ln-units km^{-1} (all $p <$ 0.01; R^2 = 0.88–0.99), confirming a robust and cross-city comparable premium gradient even in the monocentric contrast case.

Figure 17 further shows that the price decline during the 2021–2024 correction is markedly centripetal. The largest losses occur near the CBDs of Beijing and Shanghai, at approximately −1.2 and −1.1 × 10^4 yuan m^{-2}, respectively, attenuating with increasing distance; Wuhan's decline is comparatively mild, and Urumqi remains nearly stable. This indicates that the market downturn was

not spatially uniform but was concentrated mainly in core areas that had previously commanded higher prices, thereby flattening the intra-urban price gradient. This case demonstrates that continuous housing-price prediction surfaces can support distance-decay, gradient-evolution, and spatial-redistribution analyses that are difficult to conduct with traditional discrete samples.

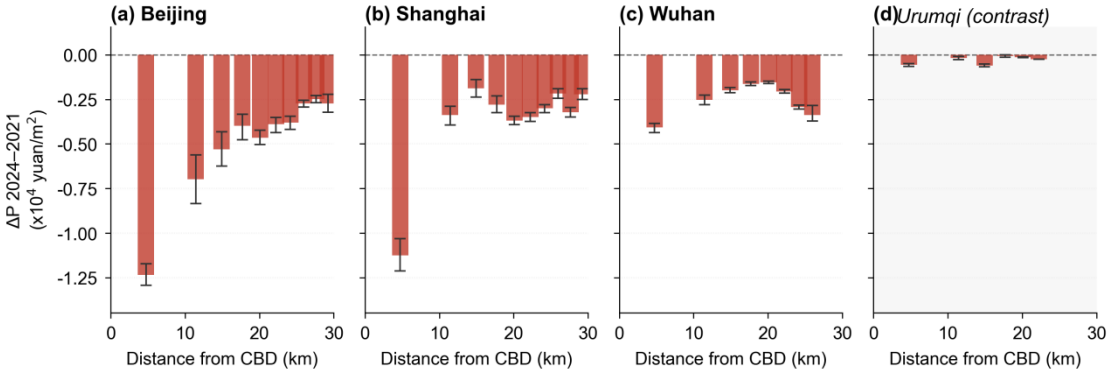

Figure 17. The 2021–2024 correction was spatially regressive, concentrating losses near the urban core. Bootstrap median change in predicted price (ΔP, 2024 minus 2021) by distance ring, with 95% confidence intervals; negative values denote decline, y-axis shared. Beijing and Shanghai register the steepest near-CBD losses (peak ΔP ≈ −1.2 and −1.1 $\times 10^4$ yuan m^{-2}), attenuating with distance, while Wuhan declines mildly and Urumqi remains essentially flat.

5.4.3. Spatial patterns of intra-urban price change during the peak-to-trough period

Figure 18 further presents the pixel-level price changes during the 2021–2024 correction for nine representative cities. Overall, the magnitude of price decline correlates with a city's baseline price level: the median declines in Beijing, Guangzhou, and Shanghai are −0.60, −0.45, and −0.35 × 10^4 yuan m^{-2}, respectively, with 97%, 95%, and 83% of pixels declining, whereas the median change in Chengdu and Urumqi is close to zero, indicating limited spatial variation in cities with a low baseline price or a weaker correction.

Pronounced intra-urban divergence is also evident. Beijing, Guangzhou, Wuhan, Shenyang, and Kunming are dominated by widespread declines, whereas Shanghai, Chengdu, and Urumqi exhibit more localised increases or resilient areas. Some cities show a clear spatial directionality—for instance, Xi'an's northern new district and its old town display contrasting patterns, and isolated peripheral areas of Shanghai rise against the overall trend. The histogram results further indicate marked differences in the distribution of price change across cities: Wuhan declines almost uniformly city-wide, whereas Chengdu's distribution is more balanced around zero. This case demonstrates that pixel-level prediction surfaces can decompose macro-level market corrections down to the intra-urban submarket scale, providing spatial evidence for identifying local risk zones and resilient areas.

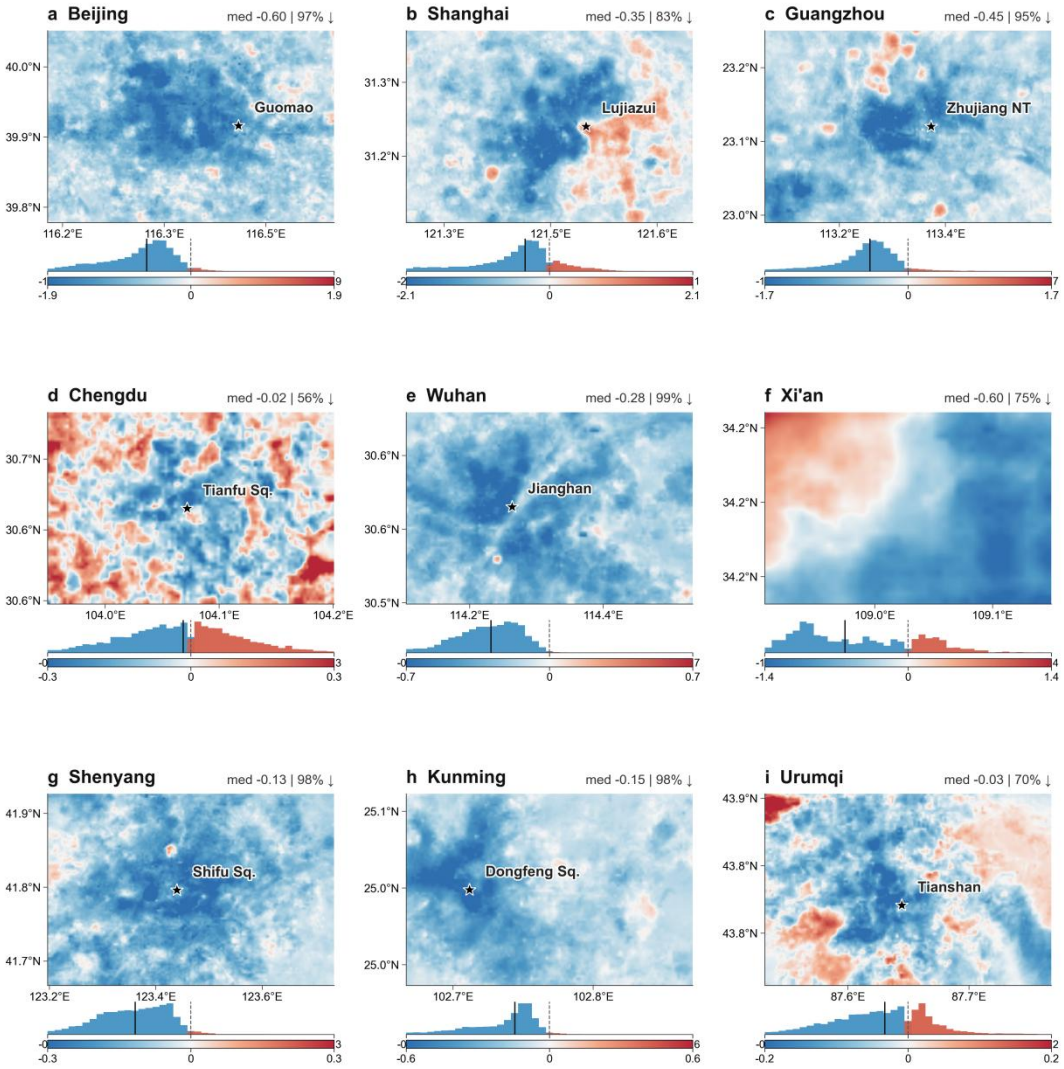

Figure 18. Pixel-level predicted price change ($\Delta P = P_{2024} - P_{2021}$) for nine representative cities during the 2021–2024 market correction. Blue denotes decline and red denotes increase, with city-specific symmetric colour scales set at the 98th percentile of $|\Delta P|$ ($\times 10^4$ yuan m^{-2}). Histograms show pixel-level ΔP distributions; solid lines mark the median, dashed lines mark zero. Black stars indicate CBD locations. Annotations report median ΔP and the share of declining pixels.

6. Discussion

6.1. Main findings

This study shows that AEF embeddings capture housing-price-relevant built-environment information, but their predictive power is limited when used alone, declining markedly under spatial cross-validation and cross-city transfer in particular. This indicates that housing prices are determined not solely by visible surface form but are also shaped by socioeconomic factors such as public-service provision, commercial vibrancy, and transport accessibility. The auxiliary-factor stacking experiments further confirm this: POI and spatial-structure variables contribute the most, whereas explicit physical

built-form features such as building height, building footprint area, and land use yield smaller gains, indicating that AEF already implicitly encodes substantial built-form information but still inadequately represents socioeconomic semantics.

Analysis of AEF's internal structure shows that housing-price information is not concentrated in a small number of stable dimensions but is instead distributed across nearly the full 64-dimensional embedding space. PCA ablation indicates that low-variance tail dimensions still carry effective housing-price signal, while SHAP and Kendall's τ results show that the key dimensions relied upon vary across cities and years without stable consistency. This indicates that AEF's representation for the housing-price task exhibits pronounced city specificity, which also explains the comparatively weak cross-city transfer ability of pure AEF.

The full-coverage inference results further demonstrate that combining AEF with auxiliary factors can generate continuous, stable urban housing-price prediction surfaces. The predicted surfaces show high consistency with the IDW reference surfaces in terms of spatial-clustering structure (Moran's I $\approx$ 0.86), and the median structural SSIM between adjacent years reaches 0.986, indicating that the model can both recover intra-urban price tiers and maintain multi-year temporal consistency. AEF + C10 therefore represents a viable technical pathway for large-scale housing-price mapping, particularly in settings where traditional interpolation methods struggle to generate continuous estimates in sample-sparse areas.

6.2. Limitations

This study has three main limitations. First, the housing-price labels are based primarily on publicly listed prices; although calibrated and quality-controlled using multiple sources, a discrepancy between listing prices and actual transaction prices may remain. Second, the model inputs are drawn mainly from remote-sensing data and publicly available auxiliary factors and do not yet incorporate finer-grained socioeconomic variables such as school-district quality, income level, or rail-transit service intensity, leaving the characterisation of complex housing-price mechanisms incomplete. Third, the study is limited to 36 Chinese cities, and the model's applicability to other countries, market institutions, and urban morphologies still requires further verification. Future work could incorporate transaction data, multimodal foundation models, and broader cross-regional experiments to improve the model's generalisability and practical value[36].

7. Conclusion

Using housing-price data from 36 Chinese cities over 2017–2024, this study evaluates the applicability of AEF embeddings to pixel-level urban housing-price mapping. The results show that pure AEF provides some housing-price-relevant representation but has limited predictive power and cross-city generalisation, whereas the C10 combination—incorporating POI, nighttime light, population density, spatial structure, and remote-sensing indices—substantially improves model performance. AEF already implicitly encodes substantial built-form information, but housing-price-relevant semantics are city-specific and require the joint expression of a near-complete embedding space and socioeconomic auxiliary factors. The continuous housing-price prediction surfaces generated for 36 cities using the best-performing model are validated in terms of both spatial clustering and temporal consistency, demonstrating that surface foundation-model embeddings can

offer a new technical pathway for large-scale, continuous mapping of urban economic variables.

References

[1] He Z, Ye J, Shi X. Housing wealth and household consumption in urban China [J]. Urban studies, 2020, 57(8): 1714-1732.

[2] Zou Y, Zhong N, Chen H, et al. Housing burden across Chinese cities: Spatio-temporal patterns and influential factors [J]. Applied Geography, 2024, 166: 103259.

[3] Wei C, Fu M, Wang L, et al. The research development of hedonic price model-based real estate appraisal in the era of big data [J]. Land, 2022, 11(3): 334.

[4] Yao Y, Zhang J, Hong Y, et al. Mapping fine-scale urban housing prices by fusing remotely sensed imagery and social media data [J]. Transactions in GIS, 2018, 22(2): 561-581.

[5] Chen M, Liu Y, Arribas-Bel D, et al. Assessing the value of user-generated images of urban surroundings for house price estimation [J]. Landscape and Urban Planning, 2022, 226: 104486.

[6] Fu X, Jia T, Zhang X, et al. Do street-level scene perceptions affect housing prices in Chinese megacities? An analysis using open access datasets and deep learning [J]. PloS one, 2019, 14(5): e0217505.

[7] Gong W, Srivastava U, Wang Y, et al. Earth Embeddings Reveal Diverse Urban Signals from Space [J]. arXiv preprint arXiv:260403456, 2026.

[8] Wen H, Xiao Y, Zhang L. School district, education quality, and housing price: Evidence from a natural experiment in Hangzhou, China [J]. Cities, 2017, 66: 72-80.

[9] He D, Sun G, Li L, et al. New metro and housing price and rent premiums: A natural experiment in China [J]. Urban Studies, 2024, 61(7): 1371-1392.

[10] Yang L, Chen Y, Xu N, et al. Place-varying impacts of urban rail transit on property prices in Shenzhen, China: Insights for value capture [J]. Sustainable Cities and Society, 2020, 58: 102140.

[11] Qu Z, Wu Y, Liu M. Multidimensional Neighborhood Amenity Characteristics and Housing Prices in Urban China [J]. Applied Spatial Analysis and Policy, 2025, 18(2): 68.

[12] Wu Y, Liu X, Li M, et al. The nonlinear impact of urban blue-green space landscape pattern on residential housing price in a new first-tier city of China [J]. Discover Cities, 2024, 1(1): 30.

[13] Jean N, Burke M, Xie M, et al. Combining satellite imagery and machine learning to predict poverty [J]. Science, 2016, 353(6301): 790-794.

[14] Levin N, Kyba C C, Zhang Q, et al. Remote sensing of night lights: A review and an outlook for the future [J]. Remote Sensing of Environment, 2020, 237: 111443.

[15] Li F, Yigitcanlar T, Nepal M, et al. Machine learning and remote sensing integration for leveraging urban sustainability: A review and framework [J]. Sustainable Cities and Society, 2023, 96: 104653.

[16] Brown C F, Kazmierski M R, Pasquarella V J, et al. Alphaearth foundations: An embedding field model for accurate and efficient global mapping from sparse label data [J]. arXiv preprint arXiv:250722291, 2025.

[17] Ma Y, Shen Y, Swatantran A, et al. Harvesting AlphaEarth: Benchmarking the geospatial foundation model for agricultural downstream tasks [J]. International Journal of Applied Earth Observation and Geoinformation, 2026, 149: 105258.

[18] Hou S, Jiao H, Liu Z, et al. AlphaEarth Foundations (AEF) in Earth Observation: A Systematic

Review of Applications and Practices [J]. 2026.
[19] Hou S, Jiao H, Xu Q, et al. Can AlphaEarth Foundations Redefine the Paradigm of Gridded Population Mapping? A Systematic Evaluation across 18 Global Cities and Large-Scale Mapping Applications [J]. 2026.
[20] Hou S, Liu Z, Jiao H, et al. Slum Detection and Density Mapping with AlphaEarth Foundations: A Representation Learning Evaluation Across 12 Global Cities [J]. arXiv preprint arXiv:260510029, 2026.
[21] Hou S, Liu Z, Jiao H, et al. AEF-Econ: Toward Plug-and-Play Socioeconomic Foundation Embeddings from AlphaEarth for Urban Remote Sensing [J]. arXiv preprint arXiv:260620697, 2026.
[22] Yu H, Yang J, Li T, et al. Morphological and functional polycentric structure assessment of megacity: An integrated approach with spatial distribution and interaction [J]. Sustainable Cities and Society, 2022, 80: 103800.
[23] Lloyd C T, Sorichetta A, Tatem A J. High resolution global gridded data for use in population studies [J]. Scientific data, 2017, 4(1): 170001.
[24] Oostwegel L J, Schorlemmer D, Guéguen P. From Footprints to Functions: A Comprehensive Global and Semantic Building Footprint Dataset [J]. Scientific Data, 2025, 12(1): 1699.
[25] Haklay M. How good is volunteered geographical information? A comparative study of OpenStreetMap and Ordnance Survey datasets [J]. Environment and planning B: Planning and design, 2010, 37(4): 682-703.
[26] Wu H, Jiao H, Hou S, et al. GeoColab: an LLM-based multi-agent collaborative framework for geospatial code generation [J]. International Journal of Digital Earth, 2025, 18(2): 2569405.
[27] Kim J, Lee Y, Lee M-H, et al. A comparative study of machine learning and spatial interpolation methods for predicting house prices [J]. Sustainability, 2022, 14(15): 9056.
[28] Soltani A, Heydari M, Aghaei F, et al. Housing price prediction incorporating spatio-temporal dependency into machine learning algorithms [J]. Cities, 2022, 131: 103941.
[29] Li Z. Extracting spatial effects from machine learning model using local interpretation method: An example of SHAP and XGBoost [J]. Computers, Environment and Urban Systems, 2022, 96: 101845.
[30] Wu H, Jiao H, Hou S, et al. GeoRRDI: An explainable and multi-source data-driven framework for rural revitalization assessment with spatial heterogeneity consideration [J]. Ecological Indicators, 2025, 181: 114387.
[31] Mirza F K, Öğrenci A S, Bilge A H, et al. Evaluating model concordance in a multi-model binary scoring framework: feature-specific and sampling-driven agreement analyses using Kendall's tau rank correlation for credit risk assessment [J]. Neural Computing and Applications, 2026, 38(7): 231.
[32] Benavides-Martinez I F, Guthrie J, Arias J E, et al. What on Earth is AlphaEarth? Hierarchical structure and functional interpretability for global land cover [J]. arXiv preprint arXiv:260316911, 2026.
[33] Feng Y, Wang Y. Spatiotemporal Evolution and Driving Forces of Housing Price Differentiation in Qingdao, China: Insights from LISA Path and GTWR Models [J]. Buildings, 2025, 15(16): 2941.
[34] Xiong S, Zhang X, Lei Y, et al. Time-series China urban land use mapping (2016–2022): an approach for achieving spatial-consistency and semantic-transition rationality in temporal

domain [J]. Remote Sensing of Environment, 2024, 312: 114344.
[35] Roberts D R, Bahn V, Ciuti S, et al. Cross-validation strategies for data with temporal, spatial, hierarchical, or phylogenetic structure [J]. Ecography, 2017, 40(8): 913-929.
[36] Hou S, Jiao H, Liang J, et al. GeoCogent: an LLM-based agent for geospatial code generation [J]. International Journal of Geographical Information Science, 2026, 40(4): 1073-1106.