

Summary of the response letter

Title: Mapping Three Decades of Urban Growth in China: A 30 m Annual Building Height Dataset (1990–2019)

Dear Editor and Reviewers,

Thank you for allowing us to respond to the reviewers' valuable comments that help to significantly improve our analyses and the paper. We carefully and extensively revise our manuscript to account for the reviewers' recommended changes. First, the Datasets section is enhanced by providing more details on the grid-level building height definition and reference datasets. Second, the Method and validation section is strengthened by clarifying the CCDC-based annual reference data construction and reconstructed VV backscatter, and the building height estimation models are optimized using a ratio-controlled resampling strategy to improve the estimation of mid- and high-rise buildings. Third, the Results and discussion section is revised by further evaluating the spatial transferability and cross-city robustness of the models and validating temporal reliability through comparisons with stable areas and historical imagery; a new subsection on uncertainty analysis and limitations is also added. We believe the revised manuscript accounts for all the reviewers' comments, and it is significantly improved as a result. We provide detailed responses to all questions and recommendations from the editor and reviewers in the attached letter. The revised manuscripts with and without track changes are attached.

Sincerely yours,

Quanhua Dong

National Satellite Meteorological Center (National Center for Space Weather),

China Meteorological Administration, Beijing, 100081, China

dqh@pku.edu.cn

Reviewer #1

General comments and responses

This study develops a time-series 30 m building height dataset for China by integrating multi-source remote sensing data, representing a valuable contribution to large-scale mapping of urban vertical dynamics. The dataset has clear potential for urban studies and related applications. However, several issues related to data quality, robustness, and validation—particularly for high-rise buildings and regions with limited reference data—still need to be carefully addressed to strengthen the reliability and usability of the product. The following comments are provided for the authors' consideration.

We express our deep gratitude for the reviewer's positive feedback on our work and sincerely appreciate the constructive comments on highlighting important areas for improvement regarding data quality, robustness, and validation. In the following sections, we address each of the reviewer's points in detail, organizing our responses in a question (black)-response (blue)-revision (red) format to ensure clarity.

Specific comments and responses

1. In Eq. (6), averaging all 1 m pixels (900 pixels) within a 30 m grid mixes building and non-building pixels. Please clarify the rationale for this choice, rather than computing height only over building-covered pixels.

We appreciate the reviewer's professional comment regarding the calculation of building height in Eq. (6) of the original manuscript. Eq. (6) calculates the height value of each 30 m × 30 m grid by averaging the height values of the 900 underlying 1 m × 1 m pixels. For clarity, this terminology is used in responses to related questions: “grid” refers to a 30 m × 30 m grid, while “pixel” refers to a 1 m × 1 m pixel.

$$H^{30m} = \frac{\sum_i^{900} H_i^{1m}}{900} = \frac{\sum_j^{nbc} H_j^{1m}}{900} \quad (6) \text{ of the original manuscript}$$

In this equation, H_i^{1m} denotes the height value of each pixel within the grid, including both building and non-building pixels. For non-building pixels, H_i^{1m} is assigned a value of 0. Therefore, the numerator is equivalent to the sum of height values over only the building pixels, denoted as H_j^{1m} . Here, nbc represents the number of building-covered pixels among the 900 pixels within the grid. Importantly, the denominator is kept as 900 rather than nbc .

This design means that Eq. (6) does not calculate the mean height of building footprints alone. Instead, it derives an area-weighted grid-level building height that preserves both building height information and building coverage density within each grid. If the average is calculated only over building-covered pixels, grids with similar building heights but markedly different building coverage fractions could yield similar mean height values, despite substantial differences in land development intensity. For example, a grid containing sparse high-rise structures and another grid with extensive, contiguous high-rise coverage could show proximate average heights if only building-covered pixels are considered. In contrast, using all 900 pixels as the denominator incorporates the proportion of non-building pixels within the grid, enabling the derived height to more comprehensively represent the overall vertical development intensity of the built environment

at the grid scale.

Although the denominator of Eq. (6) is the total number of pixels within each grid, including both building and non-building pixels, the grids themselves are constrained using the Global Artificial Impervious Area (GAIA) dataset and building-related information. This masking procedure ensures that the retained grids mainly represent buildings and their surrounding built-up environment such as roads and plazas, rather than unrelated natural land-cover types such as vegetation, water bodies, or bare land. As a result, grids without any building information are excluded from the reference sample preparation. In addition, zero-height reference samples are excluded during model training to ensure that the training samples effectively characterize building-related areas.

Through the above processing, this study computes an area-weighted average building height at the grid level. This metric reflects the overall vertical development intensity per unit grid area across the built-up region, rather than merely the average height within building footprints, and therefore captures the combined effects of building height and building density within each grid.

Furthermore, calculating building height at the grid level by considering both buildings and associated built-up surfaces is consistent with established practices in large-scale urban morphology studies. For instance, Zhou et al. (2022) state in their methodology that

“...In this study, the urban built-up height (unit: meter) was calculated as the mean height of all areas within a 500 m × 500 m grid in the urban domain, including both buildings and non-buildings such as streets, parking lots, and green space...”

Similarly, Li et al. (2020) note that their building height

“...should be noted that the building height in this study specifically refers to the mean height within the 500 m grid, including buildings and non-buildings like streets and parking lots...”

Other recent high-impact works, such as Chen et al. (2025), also adopt the average RH96 value of all GEDI samples within a grid to represent building height without explicitly filtering for building-only footprints.

In addition, our grid-level calculation aligns with the Gross Building Height definition adopted in the Global Human Settlement Layer (GHSL) framework, which distinguishes between the Net and Gross metrics (Pesaresi et al., 2024). Net Building Height refers to the average height of buildings only, excluding built-up surfaces without buildings, whereas Gross Building Height averages buildings and associated built-up surfaces within the grid. Providing Net Building Height together with building fraction information can offer a more detailed characterization of urban form. However, simultaneously mapping long-term building height and sub-pixel building fractions across decadal scales remains challenging in terms of computational cost and data consistency.

Therefore, considering that grid-level building height is widely adopted in regional and global building height mapping to represent the overall vertical development intensity of the built environment within each grid, Eq. (6) is used to calculate the area-weighted building height.

Thanks again for the reviewer’s professional comments. Considering the importance of building height calculation in this study, further clarification is added in Section 2.6 of the revised manuscript:

“ ...

To generate rasterized reference data for 2019, vector footprints are converted to 1m resolution height raster to preserve height details. Then, these raster datasets are resampled to 30m, averaging the value of every 1m pixel within the 30mx30m grid (Zhou et al., 2022).

$$H^{30m} = \frac{\sum_i^{900} H_i^{1m}}{900} = \frac{\sum_j^{nbc} H_j^{1m}}{900} \quad (6)$$

where H_i^{1m} denotes the height value of each 1m sub-pixel. Non-building sub-pixels are assigned a value of 0, so the numerator is equivalent to the sum of height values over building-covered sub-pixels only, denoted as H_j^{1m} . Here, nbc represents the number of building-covered 1m sub-pixels among the 900 sub-pixels within the grid. The denominator is kept as 900 rather than nbc , because using nbc would produce a building-footprint-only mean height and could assign comparable values to grids with similar building heights but different building coverage fractions, thereby failing to capture variations in built-up intensity.

Although non-building 1m sub-pixels are included in this calculation, the 30m grids are constrained using the GAIA impervious surface mask and building information. As a result, grids without any building information are excluded during reference sample preparation. In addition, zero-height reference samples are omitted during model training to ensure that the training data effectively represent building-related areas. Accordingly, these 30 m grids serve as the basic units for all subsequent processing, model training, validation, and dataset generation, and correspond to the 30m pixels in the final building height dataset.

With these constraints, the derived height better captures the combined effects of building height and building coverage. This grid-level calculation is consistent with established practices in large-scale urban morphology studies (Chen et al., 2025; Li et al., 2020; Zhou et al., 2022) and aligns with the Gross Building Height concept in the GHSL framework (Pesaresi et al., 2024).

...

Reference:

- [1] Zhou, Y., Li, X., Chen, W., Meng, L., Wu, Q., Gong, P., and Seto, K. C.: Satellite mapping of urban built-up heights reveals extreme infrastructure gaps and inequalities in the Global South, *Proceedings of the National Academy of Sciences*, 119(46), e2214813119, <https://doi.org/10.1073/pnas.2214813119>, 2022
- [2] Li, X., Zhou, Y., Gong, P., Seto, K. C., and Clinton, N.: Developing a method to estimate building height from Sentinel-1 data, *Remote Sensing of Environment*, 240, 111705, <https://doi.org/10.1016/j.rse.2020.111705>, 2020
- [3] Chen, P., Huang, H., Qin, P., Liu, X., Wu, Z., Zhao, F., Liu, C., Wang, J., Li, Z., Cheng, X., and Gong, P.: Characterizing dynamics of built-up height in China from 2005 to 2020 based on GEDI, Landsat, and PALSAR data, *Remote Sensing of Environment*, 325, 114776, <https://doi.org/10.1016/j.rse.2025.114776>, 2025
- [4] Pesaresi, M., Schiavina, M., Politis, P., Freire, S., Krasnodębska, K., Uhl, J. H., Carioli, A., Corbane, C., Dijkstra, L., Florio, P., Friedrich, H. K., Gao, J., Leyk, S., Lu, L., Maffneni, L., Mari-Rivero, I., Melchiorri, M., Syrris, V., Van Den Hoek, J. and Kemper, T.: 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, 2024

2. The model performs well in cities with available building height samples (e.g., Beijing and Tianjin). However, in cities without reference height samples (Fig. 14), the heights of mid-rise and high-rise buildings are significantly underestimated. This indicates limited spatial transferability and cross-city robustness of the model. Please

address this limitation, and explicitly discuss the applicability and associated uncertainty in such regions.

We sincerely appreciate the reviewer's constructive feedback regarding the spatial transferability and cross-city robustness of our model. The underestimation observed in cities lacking reference samples, as illustrated in Fig. 14 in the original manuscript, is a critical point that requires further clarification and a more transparent discussion of uncertainty.

The observed underestimation in non-sampled regions is largely attributable to the inherent challenges of large-scale building height mapping in China. High-quality, wide-coverage ground truth data is often restricted to proprietary datasets from commercial platforms like Gaode or Baidu (Wu et al., 2023; Yan et al. 2024), which primarily focus on major cities. This uneven spatial distribution of available training data leads to a certain degree of overfitting to the specific urban morphologies and architectural styles of sampled metropolitan areas, thereby limiting the model's generalization when applied to cities with different developmental profiles.

In response to this issue, the model is first optimized using a ratio-controlled sampling strategy, as elaborated in our reply to Question 3. After optimization, the model is trained with improved representation of mid- and high-rise buildings, thereby reducing prediction inaccuracies for such structures and improving overall performance, as shown in Fig. 4 of this response letter.

To further evaluate the spatial transferability and cross-city robustness of the optimized model, we conduct an independent validation using sparse building height information collected from the Lianjia real-estate platform. After being matched and processed with CMAB building footprints (Zhang, Zhao and Long, 2025), the Lianjia data provide community-level building heights. Since these data are neither pixel-level nor footprint-level ground truth, they are not used for model training. Instead, they provide an independent out-of-distribution benchmark for evaluating model performance in extrapolated cities that are excluded from the reference samples.




Fig. 1. Accuracy assessment of model performance in cities excluded from the reference samples using Lianjia real-estate data.

The results in Fig. 1 of this response letter indicate that the optimized model maintains reasonable estimation performance in extrapolated cities, with an overall RMSE of 13.40 m. For representative cities, the RMSE values are mostly within 3–6 m in low-rise built-up areas, including Linyi, Puyang, Weifang, and Quzhou. For mid- and high-rise areas, the RMSE values range from 10 to 17 m in Zhenjiang, Jiangmen, Yancheng, and Xuchang.

The relatively larger errors for taller buildings are consistent with patterns reported in existing large-scale building height studies. Che et al. (2024) show that high-rise buildings tend to contribute more strongly to regional errors; for example, in Africa, the RMSE for buildings above 50 m reaches 25.52 m, which is much higher than that for buildings below 20 m. Therefore, although underestimation still occurs in some high-rise areas, the extrapolation results remain comparable to existing findings and demonstrate the optimized model's cross-city applicability and generalization capability.

We thank the reviewer again for the constructive comment. To further evaluate the model's spatial transferability and cross-city robustness, we add a new Section 4.2.3, "Cross-city transferability assessment":

" ...

To further evaluate the spatial transferability and cross-city robustness of the models, an independent validation is conducted using sparse building height information collected from the Lianjia real-estate platform. After being matched and processed with CMAB building footprints (Zhang, Zhao, and Long, 2025), the Lianjia data provide community-level building heights. Since these data are neither pixel-level nor footprint-level ground truth, they are not used for model training. Instead, they provide an independent benchmark for evaluating model performance in extrapolated cities that are excluded from the reference samples.

The extrapolation validation is conducted using the XGB-S1S2 model trained with the 2019 reference samples, because this validation only requires a single-year cross-sectional assessment. The results show that this model maintains reasonable estimation performance in cities excluded from the reference samples, with an overall RMSE of 13.40 m. For representative cities, the RMSE values are mostly within 3–6 m in low-rise built-up areas, including Linyi, Puyang, Weifang, and Quzhou. For mid- and high-rise areas, the RMSE values range from 10 to 17 m in Jiangmen, Zhenjiang, Yancheng, and Xuchang.

The relatively larger errors for taller buildings are consistent with patterns reported in existing large-scale building height studies. Che et al. (2024) show that high-rise buildings tend to contribute more strongly to regional errors; for example, in Africa, the RMSE for buildings above 50 m reaches 25.52 m, which is much higher than that for buildings below 20 m. Therefore, although underestimation still occurs in some high-rise areas, the extrapolation results remain comparable to existing findings and demonstrate this model's cross-city applicability and generalization capability.

..."

We also clarify that while the model remains highly applicable for regional-scale urban analysis, caution should be exercised when interpreting absolute height values in regions far removed from the training data distribution. As for the current manuscript upon evaluating the inherent uncertainties in model predictions, a confidence layer has been developed to quantitatively assess prediction reliability. This uncertainty analysis is presented in a dedicated section, with detailed discussion available in the response to Question 9(2) and in the newly added Section 4.6. This

limitation will be a primary focus of our future work, where we aim to integrate self-supervised learning and multi-source remote sensing data to further enhance cross-city robustness.

Reference:

- [1] Wu, W. B., Ma, J., Banzhaf, E., Meadows, M. E., Yu, Z. W., Guo, F. X., Sengupta, D., Cai, X. X., and Zhao, B.: A first Chinese building height estimate at 10 m resolution (CNBH-10 m) using multi-source earth observations and machine learning, *Remote Sensing of Environment*, 291, 113578, <https://doi.org/10.1016/j.rse.2023.113578>, 2023
- [2] Yan, W., Wu, J., Zhang, C., Chen, X., Ren, J., Xiao, Z., Liao, Z., Laforтеzza, R., and Su, Y.: Developing an annual building volume dataset at 1-km resolution from 2001 to 2019 in China, *International Journal of Digital Earth*, 17(1), 2330690, <https://doi.org/10.1080/17538947.2024.2330690>, 2024
- [3] Che, Y., Li, X., Liu, X., Wang, Y., Liao, W., Zheng, X., Zhang, X., Xu, X., Shi, Q., Zhu, J., Yuan, H., and Dai, Y.: 3D-GloBFP: the first global three-dimensional building footprint dataset, *Earth Syst. Sci. Data*, 16, 5357–5374, <https://doi.org/10.5194/essd-16-5357-2024>, 2024
- [4] Zhang, Y., Zhao, H., and Long, Y.: CMAB: A Multi-Attribute Building Dataset of China, *Scientific Data*, 12(1), 430, <https://doi.org/10.1038/s41597-025-04730-5>, 2025

3. The model shows limited capability in predicting high-rise buildings. Please clarify the maximum building height represented in the training samples. Moreover, this limitation may be related to the uneven distribution of building height samples. In that case, more appropriate sample setting strategies or model optimization schemes could be considered. In addition, the height bins in Fig. 4 (e.g., “<10 m” and “>60 m”) should be more clearly defined (e.g., “3–10 m”, “60–X m”).

We would like to thank the reviewer for these insightful suggestions regarding the training sample distribution and the clarity of our data categorization. We have taken these comments seriously and have implemented a new ratio-controlled strategy to improve the model's performance on high-rise buildings.




Fig. 2. Maximum height in training samples each year

The maximum building height represented in the training samples is reflected in 0 of this response letter, with 139 m being the highest from 2003 to 2019, 118 m from 1990 to 2002, demonstrating

the existence of high-rise buildings in the samples. Although the Jinmao Tower broke China's building height record after 1999, reaching 420.5 m, the sample preparation procedure averages building height together with surrounding built-up surfaces, as clarified in response to Question 1. Therefore, the resulting value represents an area-weighted grid-level height rather than the height of the individual building itself, which lowers the maximum height value represented in the dataset. Based on the maximum building height represented in the reference samples, building heights are divided into four intervals: (0 m, 10 m], (10 m, 30 m], (30 m, 60 m], and (60 m, 140 m]. Section 4.1 of the manuscript is updated to report the maximum heights:

“...The highest sample is at 139 m from 2003 to 2019, and 118 m from 1990 to 2002, demonstrating the existence of high-rise buildings in the samples....”

Regarding the sample distribution, a significant imbalance indeed exists in our original dataset, where buildings under 10 m accounted for over 70% of the total samples. The initial attempt is to balance all height intervals equally through resampling; however, this strategy produces suboptimal results. This is likely because low-rise buildings constitute the dominant building type in the study area, and maintaining a higher proportion of low-rise samples is essential for preserving the model's prediction accuracy for these prevalent building types (Moraes, Campagnolo, and Cartano, 2024).

Consequently, a more nuanced ratio-controlled resampling scheme is adopted to balance minority-class enhancement and majority-class preservation. This approach is consistent with established practices for handling imbalanced samples in remote sensing and urban studies (Chawla et al., 2002; Naboureh et al., 2020). In the updated strategy, the sample size of each height interval above 10 m is capped at 1/10 of that of the (0 m, 10 m] category. This design largely preserves the dominance of low-rise buildings while increasing the relative proportions of the (30 m, 60 m] and (60 m, 140 m] high-rise groups. Specifically, down-sampling is applied to over-represented height bins, whereas oversampling is applied to under-represented bins. As the number of newly added pixels through oversampling exceeds the number of pixels removed through down-sampling, the total sample size increases, as shown in Fig. 3 of this response letter. Although the proportion of the (10 m, 30 m] interval decreases, the experiment shows no negative effect on model accuracy. This strategy maintains the prediction accuracy for predominant low-rise buildings and simultaneously enhances the representation of high-rise samples, thereby improving the model's ability to estimate taller buildings. A new Section 3.2 “Data Augmentation” is added to describe the process:

“...The original height distribution of reference samples is highly imbalanced, with low-rise buildings accounting for the majority. This imbalance may limit the model's ability to accurately estimate relatively high buildings. Although equal-proportion resampling can balance different height intervals, it reduces the representation of the original majority class and leads to a decline in prediction accuracy for the most prevalent building types (Moraes, Campagnolo, and Cartano, 2024). To improve the model's performance for high-rise buildings while preserving its accuracy for dominant low-rise structures, a ratio-controlled resampling strategy is adopted.

Specifically, building heights are divided into four intervals based on reference samples: (0 m, 10 m], (10 m, 30 m], (30 m, 60 m], and (60 m, 140 m]. The sample size of each height interval above 10 m is adjusted through down-sampling or oversampling and capped at one-tenth of that of the (0 m, 10 m]

category (Chawla et al., 2002; Naboureh et al., 2020). As the number of newly added pixels through oversampling exceeds the number of pixels removed through down-sampling, the total sample size increases. Although the proportion of the (10 m, 30 m] interval decreases, the experiment shows no negative effect on model accuracy. This strategy maintains the prediction accuracy for predominant low-rise buildings and simultaneously enhances the representation of high-rise samples, thereby improving the model's ability to estimate taller buildings....”

Based on the optimized sample distribution, the models are retrained and a new dataset is generated. The updated results show substantial improvements in accuracy. The updated results and regenerated dataset are further presented in our response to Question 4, where the accuracy assessment figures (i.e., Fig. 4 of this response letter) are also revised accordingly. Specifically, the RMSE values for 2019, 2017, and 2014 decrease from 6.37 m, 6.60 m, and 6.69 m to 3.90 m, 4.14 m, and 4.11 m, respectively; the MAE values decrease from 3.66 m, 3.79 m, and 3.80 m to 2.70 m, 2.86 m, and 2.83 m, respectively; and the R^2 values increase from 0.46, 0.43, and 0.42 to 0.80, 0.78, and 0.78, respectively. These results demonstrate the effectiveness of the ratio-controlled sampling strategy in improving model accuracy.

It should be clarified that the ratio-controlled resampling strategy is applied only during model training to mitigate sample imbalance. The accuracy assessment is conducted using the original reference samples, rather than the resampled training samples. This ensures that the reported metrics reflect the model's predictive performance on the actual reference data distribution and are not affected by the artificial sample distribution introduced during resampling (Elreedy et al., 2019).

Notably, these optimizations directly address the reviewer's concern regarding the model's capability to predict mid- and high-rise structures and further enhance the model's cross-city robustness and spatial transferability, as shown in the response to Question 2.




Fig. 3 (the revised Fig. 4 in the original manuscript). Number of reference samples and height proportion each year. (a) Number and proportion before resample; (b) Number and proportion after resample.

Regarding the height-bin notations, we sincerely appreciate the reviewer's suggestion to define them more clearly. The height-bin notations in Figs. 4 and 7 of the original manuscript are revised to (0 m, 10 m], (10 m, 30 m], (30 m, 60 m], and (60 m, 140 m]. This updated notation is more precise, as it explicitly defines the mathematical range of each height interval.

As explained in our response to Question 1 and 4, the area-weighted grid-level height is used as the reference height, which may result in height values between 0 m and 3 m, lower than the height

of a single floor. Since pixels without any building information are excluded during model training, the first height interval is denoted as (0 m, 10 m]. For the last interval, (60 m, 140 m], the upper bound of 140 m is selected because it is higher than any building height represented in the reference samples, as shown in Fig. 2 of this response letter.

References:

- [1] Moraes, D., Campagnolo, M. L. and Caetano, M.: Training data in satellite image classification for land cover mapping: a review. *European Journal of Remote Sensing*, 57(1), 2341414, 2024
- [2] Chawla, N. V., Bowyer, K. W., Hall, L. O. and Kegelmeyer, W. P.: SMOTE: synthetic minority over-sampling technique. *Journal of artificial intelligence research*, 16, 321-357, 2002
- [3] Naboureh, A., Li, A., Bian, J., Lei, G. and Amani, M.: A hybrid data balancing method for classification of imbalanced training data within google earth engine: Case studies from mountainous regions. *Remote Sensing*, 12(20), 3301, 2020
- [4] Elreedy, D. and Atiya, A. F.: A comprehensive analysis of synthetic minority oversampling technique (SMOTE) for handling class imbalance. *Information sciences*, 505, 32-64, 2019

4. In Fig. 7, the reference building heights are mainly concentrated within 0–8 m. Pixels with height values below 3 m are likely non-building pixels and could be excluded from validation. Moreover, reporting relative error metrics (e.g., rRMSE) would provide a more informative evaluation.

We appreciate the reviewer's valuable suggestions regarding the height thresholds and additional error metrics for the Fig. 7 of the original manuscript.

The concentration of samples around 0–8 m in the density scatter plot shown in Fig. 4 of this response letter, corresponding to the revised Fig. 7 in the original manuscript, is mainly caused by the following three factors:

First, the 30 m grid height in this study represents an area-weighted height, rather than the absolute height of individual buildings. In other words, the grid-level height is jointly affected by both building height and building coverage fraction. When buildings occupy only a small proportion of a grid, the area-weighted height can still be relatively low even if real buildings are present within the grid. Therefore, the concentration around 0–8 m also reflects the prevalence of low-rise buildings and areas with low building coverage in the reference samples.

Second, the reference samples themselves show a clear concentration in the low-height range. As shown in the sample distribution analysis in our response to Question 3, 30 m grids with heights in the interval of (0 m, 10 m] account for the vast majority of the reference samples. This indicates that low-rise buildings, sparsely built-up areas, and areas with low vertical development intensity are the dominant building-height types in the study area. Therefore, the concentration of samples in the low-height range in the accuracy assessment figure is primarily determined by the height distribution of the reference samples.

Third, the visualization mechanism of the density scatter plot further highlights areas with high sample frequency. The color intensity in the density scatter plot reflects the local concentration of sample points. Since the number of samples in the (0 m, 10 m] interval is much larger than that in other height intervals, this range shows a higher point density in the density plot, making the samples visually more concentrated around 0–8 m.

Overall, the concentration of samples in the low-height range in the density plot is a reasonable

phenomenon jointly caused by the area-weighted height definition, the reference sample distribution, and the visualization mechanism. It does not indicate that the model can only predict low-height buildings, nor does it mean that mid- or high-rise building samples are excluded from the validation.

To avoid this ambiguity, we add the following clarification in Section 4.2.1 of the revised manuscript:

“...The concentration of samples in the low-height range in the density plot is reasonable. It is mainly caused by three factors: the area-weighted definition of grid-level height; the reference sample distribution, which reflects the dominance of low-rise buildings, sparsely built-up areas, and areas with low vertical development intensity in China; and the visualization mechanism of the density scatter plot. This phenomenon does not indicate that the model can only predict low-height buildings, nor does it mean that mid- or high-rise building samples are excluded from the validation....”

In our dataset, grids with height values greater than 0m but lower than 3m generally represent low area-weighted grid-level building height, where the summed height of building pixels is averaged over all pixels within the grid. As mentioned in our response to Question 1, a grid containing sparse high-rise structures and another grid with intensive, contiguous high-rise coverage could show similar mean heights if only building-covered pixels are considered, despite substantial differences in land development intensity. Thus, grids with height values greater than 0m but lower than 3m usually indicate that buildings occupy only a limited proportion of the grid area, rather than the absence of buildings. They may correspond to low-rise structures, transitional urban fringes, or sparsely built-up areas that are physically present in real urban environments. Excluding these values would lead to an incomplete assessment of the model's performance in complex, heterogeneous urban settings. By retaining these 30m pixels/grids, the validation provides a more transparent and comprehensive evaluation of the dataset's ability to represent the full spectrum of urban vertical forms. Therefore, after careful consideration, 30m pixels with height values greater than 0m but lower than 3m are retained in the validation process. However, we agree that using 0m in the legend may cause misunderstanding. To avoid this ambiguity, we add the following clarification in Section 4.2.1 of the revised manuscript:

“...Since each 30m grid represents the area-weighted height of buildings and surrounding built-up surfaces, height values greater than 0m but lower than 3m may occur when buildings occupy only a limited proportion of the grid area, rather than indicating the absence of buildings. Accordingly, for visualization purposes only, the starting points of both the X- and Y-axes are set to 0m....”

Following the reviewer's suggestion, relative Root Mean Square Error (rRMSE) is incorporated into the accuracy assessment, as shown in Fig. 4 of this response letter. The rRMSE normalizes RMSE by the mean reference height, making it useful for comparing model performance across different urban densities and height ranges. The rRMSE values for 2019, 2017, and 2014 are 0.61, 0.63, and 0.63, respectively. We also update Fig. S1 in the Supplementary Material by adding the rRMSE values for each mapped year.

The rRMSE metric is incorporated, and its formula is added to Section 3.4 (former Section 3.3) of the revised manuscript:

“...Model performance is quantified by three metrics: Root Mean Square Error (RMSE, Eq. (8)), Mean Absolute Error (MAE, Eq. (9)), coefficient of determination (R^2 , Eq. (10)), and relative Root Mean Square Error (rRMSE, Eq. (11))...

$$rRMSE = \frac{RMSE}{\frac{\sum_{i=1}^n H_{i,ref}}{n}} \quad (11)$$

...”

All accuracy metrics reported in the manuscript are updated accordingly. For example, Section 4.2.1 is revised as follows:

“...with RMSE values of 3.90 m, 4.14 m, and 4.11 m; MAE values of 2.70 m, 2.86 m, and 2.83 m; R^2 values of 0.80, 0.78, and 0.78; and rRMSE values of 0.61, 0.63 and 0.63, respectively. Linear regression slopes between reference and estimated heights are 0.77, 0.75, and 0.75, with intercepts of 0.64 m, 0.71 m, and 0.71 m...”




Fig. 4 (the revised Fig. 7 in the original manuscript). Accuracy of XGBoost-based models in their initial years after resampling.

5. In building change areas, the accuracy of estimated height changes should be quantitatively assessed to evaluate whether the model can reliably capture building height dynamics. In stable areas, the temporal consistency of predicted building heights should be quantitatively examined to assess prediction stability.

We appreciate the reviewer’s suggestion to quantitatively assess the dataset performance in both building change areas and stable areas. This is indeed crucial for evaluating whether the generated long-term building height dataset can reliably characterize spatiotemporal urban dynamics.

To validate the reliability of the long-term building height dataset in characterizing building change areas, both quantitative comparisons and historical remote sensing images are used. Fig. 13

and Fig. 14 in the original manuscript are updated as Fig. 5 and Fig. 6 in this response letter, respectively. These two figures are designed to assess three typical urban change contexts as described in Section 3.1.1 of the manuscript, namely persistent built-up areas, complete demolition areas, and new construction areas, in both sampled and non-sampled cities.

Specifically, Fig. 5 in this response letter (Fig. 13 in the original manuscript) focuses on representative areas in sampled cities, including Tianjin, Beijing, and Wuhan. For these cities, predicted heights, reference heights, and historical remote sensing images are jointly compared, allowing a quantitative assessment of the mapped building height dynamics. Fig. 6 in this response letter (Fig. 14 in the original manuscript) focuses on representative areas in non-sampled cities, including Handan, Jining, and Anyang. Since reference height data are unavailable in these non-sampled cities, the quality of the extrapolated building height results is mainly examined using historical remote sensing images as visual evidence of construction, demolition, and redevelopment processes.

It should be further clarified that, due to the lack of long-term records of historical building heights, the annual reference building heights used in sampled cities are not direct observations for all historical building pixels. As described in Section 3.1.1 of the manuscript, they are derived from persistent built-up areas identified based on the 2019 reference data and CCDC breakpoint information. Specifically, for a given year, only pixels with no breakpoint between that year and 2019 are kept as reference data for that year. Therefore, the annual reference data mainly represent areas without clear land-use conversion or building change between the target year and 2019, and do not include pixels that experience complete demolition or substantial change during this period. This explains the local differences between predicted and reference heights in demolition areas such as Jiaohuachang (Fig. 5b in this response letter).

For sampled cities, the predicted heights generally show trends that are consistent with the reference heights and agree well with historical remote sensing images. In persistent built-up and new construction areas, such as Drum Tower in Tianjin (Fig. 5a in this response letter) and Yijiangyuan in Wuhan (Fig. 5c in this response letter), the dataset effectively captures building height dynamics. Taking Yijiangyuan as an example, the reference heights in 2000, 2005, 2013, and 2019 are 5.19 m, 5.26 m, 5.69 m, and 6.00 m, respectively, while the corresponding predicted heights are 5.15 m, 5.30 m, 5.72 m, and 6.11 m, with errors all below 0.11 m. This indicates good consistency between the predicted height changes and the reference data for sampled cities.

For the complete demolition case in the sampled cities, Jiaohuachang in Beijing is selected (Fig. 5b in this response letter). The coke plant is clearly visible in the 2005 historical remote sensing image but experiences demolition between 2005 and 2019. Therefore, this area does not satisfy the “no breakpoint between the target year and 2019” condition and, in principle, is not retained as a 2005 reference sample. In other words, the 2005 reference height cannot fully represent the true building height before demolition. As a result, a difference is observed between the predicted height and reference height in 2005: the predicted height is 5.71 m, whereas the reference height is 3.52 m. This difference mainly results from the applicability limitations of the annual reference data generation method, rather than from model bias. Moreover, historical remote sensing images in 2005, 2010, 2015, and 2018 confirm the demolition process, and the estimated height decreases from 5.71 m to 4.04 m, reflecting the transition from relatively high built-up structures to demolition and redevelopment.

For non-sampled cities, Fig. 6 in this response letter further examines the extrapolated building

height results of the dataset under the same three urban change contexts. Although reference height data are unavailable, the mapped height changes are broadly consistent with the construction and demolition processes visible in historical remote sensing images. In Handan Congtai Park (Fig. 6a in this response letter), the area is dominated by bungalows in 2006 and is gradually transformed into clusters of multi-story apartments, which become prominent in the 2013 and 2019 imagery. Correspondingly, the estimated height increases from 5.90 m to 7.40 m, then to 8.04 m, and 8.68 m, indicating that the dataset captures the redevelopment process from low-rise buildings to higher residential structures. In Jining Shiliying Village (Fig. 6b in this response letter), the settlement progressively sinks into a subsidence lake, leading to the abandonment and demolition of houses and the near disappearance of the village by 2019. This process is reflected by a gradual decrease in estimated height from 4.05 m to 3.91 m, then to 3.74 m, and 3.53 m. In Anyang Yiyuan (Fig. 6c in this response letter), new developments are built on previously undeveloped land and expand continuously between 2005 and 2019. The estimated height correspondingly increases from 4.91 m to 5.60 m, 7.22 m, and 8.63 m, showing that the dataset reflects the progressive construction process in this non-sampled city.

Overall, the supplementary analysis demonstrates that the generated dataset can reasonably characterize long-term building height dynamics across persistent built-up, complete demolition, and new construction areas. In sampled cities, the predicted height changes are supported by both reference data and historical imagery. The local differences between predicted and reference heights in demolition areas mainly arise from the applicability limitations of the annual reference data generation method. In non-sampled cities, although reference height data are unavailable, the extrapolated temporal trends are consistent with visually identifiable urban change processes in historical remote sensing images.

To evaluate prediction stability in unchanged areas, we identify a set of stable areas where building footprints and heights remain unchanged throughout the study period. These areas are defined as pixels with no change year detected by CCDC from 1990 to 2019. As shown in Fig. 7 of this response letter, the estimated mean building heights for these stable areas remain highly consistent from 1990 to 2019. The highest predicted mean height is 7.13 m in 2011, while the lowest is 6.95 m in 1992, corresponding to a fluctuation of less than 0.2 m. The red dotted line represents the mean reference height of the stable areas in 1990, which is 5.58 m and is close to the predicted heights. The orange bars, representing the difference between the predicted height in each year and that in 1990, remain small, with the largest deviation being 0.17 m in 2011. Overall, these results demonstrate that the dataset does not introduce artificial temporal noise and maintains high temporal stability in areas without physical urban change.




Fig. 5 (the revised Fig. 13 in the original manuscript). Quantitative comparison of urban areas in sampled cities. (a) Drum Tower 435 region, Tianjin; (b) Jiaohuachang region, Beijing; (c) Yijiangyuan region, Wuhan. Remote sensing images are from © Google.




Fig. 6 (the revised Fig. 14 in the original manuscript). Comparison between mapped height changes and historical remote sensing images in non-sampled cities. (a) Congtai Park region, Handan; (b) Shiliying region, Jining; (c) Yiyuan region, Anyang. Remote sensing images are from © Google.




Fig. 7. Predicted height and reference height in stable areas.

Following the reviewer’s suggestion, Section 4.4.3 “Comparison with historical remote sensing images” is expanded, reorganized, and retitled as “Comparison with stable areas and historical imagery for temporal reliability validation”. This section not only retains the comparison with historical remote sensing images, but also adds a temporal stability assessment in unchanged areas and validation results for building change areas in both sampled and non-sampled cities. Considering that the figure numbering in the manuscript may be adjusted in subsequent revisions, the figure numbers used in the revised manuscript text below are referenced according to the numbering in this response letter.

“ ...

Section 4.4.3 Comparison with stable areas and historical imagery for temporal reliability validation

To evaluate the temporal stability of the generated building height dataset, stable areas are identified where building footprints and heights remain unchanged throughout the study period. These areas are defined as pixels with no change year detected by CCDC from 1990 to 2019. As shown in Fig. 15, the estimated mean building heights in these stable areas remain highly consistent over time. The highest mean height is 7.13 m in 2011, while the lowest is 6.95 m in 1992, corresponding to a fluctuation of less than 0.2 m. The mean reference height in 1990 is 5.58 m, shown by the red dotted line, and is close to the estimated heights. The orange bars show annual differences from the 1990 estimated height, with a maximum deviation of 0.17 m in 2011. These results indicate that the dataset does not introduce artificial temporal noise and maintains high temporal stability in areas without physical urban change.

To validate the reliability of the long-term building height dataset in building change areas, Fig. 16 and Fig. 17 examine three typical urban change contexts described in Section 3.1.1, including persistent built-up, complete demolition, and new construction areas. For sampled cities, predicted heights, reference heights, and historical remote sensing images are jointly compared; for non-sampled cities, where reference height data are unavailable, historical remote sensing images are used as visual evidence for the extrapolated results.

In sampled cities, the predicted heights generally agree well with both reference heights and historical imagery. For persistent built-up areas and new construction areas, such as Drum Tower in Tianjin (Fig. 16a) and Yijiangyuan in

Wuhan (Fig. 16c), the dataset effectively captures building height dynamics. Taking Yijiangyuan as an example, the reference heights in 2000, 2005, 2013, and 2019 are 5.19 m, 5.26 m, 5.69 m, and 6.00 m, respectively, while the corresponding predicted heights are 5.15 m, 5.30 m, 5.72 m, and 6.11 m, with errors all below 0.11 m. For the complete demolition case, Jiaohuachang in Beijing (Fig. 16b) is selected. The estimated height decreases from 5.71 m to 4.04 m, consistent with the demolition and redevelopment process visible in historical imagery. As described in Section 3.1.1, annual reference heights are derived from persistent built-up areas identified using the 2019 reference data and CCDC breakpoint information, rather than direct observations of all historical building pixels. Therefore, the 2005 difference between the predicted and reference heights, 5.71 m versus 3.52 m, mainly arises from the applicability limitation of the annual reference data generation method in demolition areas. This difference reflects the limited representativeness of the reference height for demolished buildings, rather than model bias or a data error in the generated dataset.

In non-sampled cities, mapped temporal trends are broadly consistent with visually identifiable urban changes. The estimated height at Handan Congtai Park (Fig. 17a) increases from 5.90 m to 7.40 m, 8.04 m, and 8.68 m, reflecting redevelopment from bungalows to multi-story apartments. At Jining Shiliying Village (Fig. 17b), the estimated height decreases from 4.05 m to 3.91 m, 3.74 m, and 3.53 m, consistent with settlement abandonment and demolition caused by subsidence. For Anyang Yiyuan (Fig. 17c), the estimated height increases from 4.91 m to 5.60 m, 7.22 m, and 8.63 m, reflecting progressive new construction on previously undeveloped land. These results indicate that the generated dataset can reasonably characterize long-term building height dynamics across persistent built-up, demolition, and new construction areas.

...”

6. For dataset comparison, it would be more appropriate to focus on comparable 30 m resolution multi-temporal or time-series building height datasets, such as He et al. (2023) and Chen et al. (2025).

We thank the reviewer for suggesting a comparison with other high-quality time-series building height datasets. Following this suggestion, two representative 30m resolution multi-temporal/long-term building height datasets are incorporated into comparative analysis: the 1985–2010 dataset of He et al. (2023) and the 2005\2010\2015\2020 dataset by Chen et al. (2025). Our dataset is compared with these multi-temporal/long-term datasets at both the nationwide scale and the city scale.

Since different datasets vary substantially in reference data sources and building height inference methods, making strict pixel-level quantitative cross-validation not fully appropriate. For example, our study uses vector-based building heights as the source of reference data; He et al. (2023) adopt a sample-free method based on DSM neighborhood analysis; and Chen et al. (2025) mainly extract building heights from GEDI spaceborne lidar observations. Therefore, we focus on comparing temporal height trends and spatial pattern consistency to more reasonably evaluate differences among these datasets.

First, at the nationwide scale, the four datasets all show relatively stable temporal trends in building height. During 1990–2019, the nationwide mean height of our dataset is close to that of He et al. (2023), with values of approximately 5.17 m and 5.18 m, respectively, whereas Chen et al. (2025) reports a relatively higher mean height of approximately 10.44 m. The CMAB dataset reports even higher values, exceeding 14 m. These differences in building height values are mainly related

to differences in sample coverage and reference data sources among the datasets. For example, CMAB focuses more on urban center built-up areas, whereas the other long-term datasets have broader coverage, including rural areas, townships, and low-density built-up areas, resulting in relatively lower overall mean heights. Despite systematic height differences among the datasets, they all show relatively stable temporal trends at the nationwide scale, indicating that our dataset reasonably reflects the stable characteristics of long-term building height changes in China. This stability also suggests that China's urban development maintains a basic balance between horizontal expansion and vertical growth.




Fig. 8 (the revised Fig. 12 in the original manuscript). Temporal change in annual average height in three long-term building height datasets.

Second, at the urban local scale, Beijing and Guangzhou are selected as two representative cities for detailed comparison. To further examine the ability of different datasets to characterize local building height changes, we select Zhongguancun in Beijing and Liwan in Guangzhou as case study areas with the aid of historical remote sensing images (Fig. 9 in this response letter).

(1) As shown by the historical remote sensing images in Fig.9a in this response letter, Zhongguancun in Beijing experiences obvious urban renewal and demolition–reconstruction processes around 2005, making it suitable for evaluating the ability of datasets to capture dynamic changes in building height. The results show that the dataset of He et al. (2023) fails to sufficiently reflect the demolition and reconstruction process in this area, whereas the dataset of Chen et al. (2025) shows a certain temporal drifting phenomenon, namely unreasonable fluctuations in the heights of some stable buildings across different time steps. In contrast, our dataset better reflects the building height growth and renewal process in the local area.

(2) Liwan District in Guangzhou (Fig. 9b in this response letter) has remained generally stable since 2005, making it suitable for testing the temporal consistency of datasets in stable built-up areas. The comparison shows that the dataset of Chen et al. (2025) also exhibits a certain temporal drifting phenomenon in this area, whereas our dataset maintains better height continuity and temporal consistency in stable built-up areas.

Overall, the supplementary comparison results show that, compared with other long-term datasets, our dataset maintains reasonable long-term trends at the nationwide scale, better reflects height changes in local renewal areas at the urban scale, and preserves good temporal consistency in stable built-up areas.




Fig. 9. Comparison of long-term building height datasets in different cities. (a) Beijing; (b) Guangzhou.

Following the reviewer’s suggestion, Section 4.4.2 “Comparison with long-term building height products” is revised as follows. Considering that the figure numbering in the manuscript may be adjusted in subsequent revisions, the figure numbers used in the revised manuscript text below are referenced according to the numbering in this response letter.

“ ...

Section 4.4.2 Comparison with long-term building height products

To provide a comprehensive comparison with existing long-term building height products, three representative datasets are selected: the 1985–2010 dataset of He et al. (2023), the 2005/2010/2015/2020 dataset of Chen et al. (2025), and the China Multi-Attribute Building (CMAB) dataset (Zhang, Zhao, and Long, 2025). Since He’s dataset and CMAB record building construction completion years, the average building height for each reference year is calculated using buildings constructed before that year. For our dataset and Chen et al. (2025), direct annual height averaging is applied.

At the nationwide scale, Fig. 13 compares annual mean height trends among these datasets. During 1990–2019, our dataset and He et al. (2023) show similar nationwide mean heights of approximately 5.17 m and 5.18 m, respectively, whereas Chen et al. (2025) reports a higher mean height of approximately 10.44 m, and CMAB exceeds 14 m. These differences are mainly related to sample coverage and reference data sources. For example, CMAB focuses more on urban center built-up areas, while the other long-term datasets cover broader regions, including rural areas, townships, and low-density built-up areas. Despite these systematic height differences, all datasets show relatively stable temporal trends at the nationwide scale, suggesting that our dataset reasonably reflects the long-term stability of building height changes in China.

.....

At the urban local scale, Beijing and Guangzhou are selected for detailed comparison (Fig. 14). In Zhongguancun, Beijing, where obvious urban renewal and demolition–reconstruction occur around 2005. He et al. (2023) does not sufficiently reflect the local redevelopment process, while Chen et al. (2025) shows temporal drifting with unreasonable height fluctuations for some stable buildings. Our dataset better reflects local building height growth and renewal.

In Liwan District, Guangzhou, which has remained generally stable since 2005, Chen et al. (2025) also shows temporal drifting, whereas our dataset maintains better height continuity and temporal consistency.

...”

References:

- [1] He, T., Wang, K., Xiao, W., Xu, S., Li, M., Yang, R., and Yue, W.: Global 30 meters spatiotemporal 3D urban expansion dataset from 1990 to 2010, *Scientific data*, 10(1), 321, <https://doi.org/10.1038/s41597-023-02240-w>, 2023
- [2] Chen, P., Huang, H., Qin, P., Liu, X., Wu, Z., Zhao, F., Liu, C., Wang, J., Li, Z., Cheng, X., and Gong, P.: Characterizing dynamics of built-up height in China from 2005 to 2020 based on GEDI, Landsat, and PALSAR data, *Remote Sensing of Environment*, 325, 114776, <https://doi.org/10.1016/j.rse.2025.114776>, 2025

7. Building height results are visualized using both classified and continuous color schemes, with inconsistent value ranges across figures (Figure 11, Figure 13, Figure 14, Figure 15). Adopting a unified classified visualization scheme would improve clarity and consistency.

We thank the reviewer for the suggestion regarding the unification of building height legends and color schemes. We agree that clear and consistent legends are helpful for improving the readability and comparability of the figures.

It should be noted that different classified or continuous color schemes are used in the original manuscript mainly because the figures differ in spatial scale, study object, and visualization purpose. Some figures are designed for the nationwide scale or large municipal administrative regions, which include many low-rise buildings, townships, and rural built-up areas, where building heights are generally low. Other figures are focused on community-scale areas or typical building change areas, where local height differences and temporal change patterns are emphasized. If exactly the same color range and classification intervals are adopted for all figures, subtle height differences in low-rise areas may be compressed, or high-rise areas may become color-saturated, thereby weakening the representation of local details.

Following the reviewer’s suggestion and considering the updated dataset, the legends and color schemes of the relevant figures are optimized. For figures requiring cross-region or cross-temporal comparison, more consistent classified color schemes and height intervals are adopted. For figures where local details need to be highlighted, the height ranges and classification standards are explicitly marked in the legends to avoid ambiguity.

Specifically, Fig. 11 in the original manuscript, corresponding to Fig. 10 in this response letter, covers large municipal regions and rural built-up areas with generally low building heights. Therefore, the legend range is set to 0–30 m to better highlight differences in low-rise buildings and urban–rural transitional areas. Fig. 13 (Fig. 5 in this response letter) and Fig. 14 (Fig. 6 in this response letter) mainly show community-scale or local building change areas; therefore, classification schemes suitable for local interpretation are retained, with the corresponding height ranges clearly labelled. For dense urban core areas with more high-rise structures, the legend range of Fig. 15c–e in the original manuscript (Fig. 11c–e in this response letter) is set to 0–150 m to avoid truncation or color saturation. In addition, a truncated classification scheme is adopted in Fig. 15a–b in the original manuscript (Fig. 11a–b in this response letter) to enhance nationwide visual contrast

and avoid excessive homogeneity caused by the dominance of low-rise buildings.

After revision, the building height visualization schemes in Fig. 11, Fig. 13, Fig. 14, and Fig. 15 of the original manuscript (Fig. 10, Fig. 5, Fig. 6, and Fig. 11 in this response letter, respectively) are more standardized, with clearer legends and better balance between cross-figure comparability and local detail representation.




Fig. 10 (the revised Fig. 11 in the original manuscript). Comparison of building height datasets in different cities in 2019. (a) Beijing; (b) Shanghai; (c) Chengdu; (d) Guangzhou.




Fig. 11 (the revised Fig. 15 in the original manuscript). Spatial distribution and temporal progression of building heights. (a, b) Spatial distribution of building heights in 1990 and 2019; (c-e) Temporal progression of building height in Beijing, Shanghai and Shenzhen.

8. The comparison presented in Fig. 10 raises several concerns.

We would like to express our gratitude to the reviewer for the meticulous examination of Fig. 10 in the original manuscript, corresponding to Fig. 12 in this response letter. The concerns raised are highly valuable and have guided us in performing a more rigorous re-validation of our comparative analysis.

(1) Different building height products adopt different height definitions (e.g., average height, area-weighted average height), and these definitional differences should be explicitly considered. In addition, it is worth noting that some gridded building height products may include only building information, which appears to differ from the

definition adopted in this study. As a result, direct resampling and comparison across products may not be strictly comparable.




Fig. 12 (the revised Fig. 10 in the original manuscript). Comparison of reference height, our dataset and other datasets in Beijing in 2019. (a) Height distribution of different datasets; (b) Scatter plot of our dataset and reference heights; (c) Scatter plot of WSF 3D (Esch et al., 2022) and reference heights; (d) Scatter plot of CNBH-10m (Wu et al., 2023) and reference heights; (e) Scatter plot of Ma's dataset (Ma et al., 2024) and reference heights.

We thank the reviewer for pointing out the differences in height definitions among different building height products. We fully agree with this comment. Different products may adopt different height definitions, such as average height calculated only over building-covered areas, area-weighted average height that incorporates building coverage fraction, or different treatments of non-building built-up surfaces. If these definitional differences are ignored, direct resampling and comparison across products may lead to limited comparability and affect the fairness and interpretability of the accuracy assessment.

To address this issue, no existing building height product is modified, and the comparison is not based on simply resampling all products to the same spatial scale and directly comparing their absolute height values. Instead, for each evaluated product, corresponding reference data are generated to match its spatial resolution, spatial reference, and height definition. Specifically, the same set of Baidu building height vectors is first converted into a 1m resolution height raster to ensure consistency in the original reference data source. Then, according to the spatial resolution, spatial reference, and height definition of each product, the corresponding aggregation method is applied to generate matched reference data. In this way, consistency in scale, spatial location, and height definition between each evaluated product and its corresponding reference data is ensured as

far as possible.

Specifically, as shown in Table 1 in this response letter, for the datasets of Wu et al. (2023) and Ma et al. (2024), their products mainly represent the height of building-covered areas and exclude non-building areas. Therefore, when the corresponding reference data are generated, only building-covered pixels are averaged and then aggregated to the spatial resolutions of the respective products, namely 10m and 150m. This ensures consistency between the reference height and their “building-only” definition. In contrast, WSF 3D (Esch et al., 2022) derives building height from volume data, and its definition is closer to an area-weighted height that integrates buildings and built-up surfaces, which is more consistent with the grid-level height concept adopted in this study. Therefore, when WSF 3D is evaluated, the corresponding reference data are generated by averaging all pixels within each 90m grid, including non-building built-up surfaces surrounding buildings, to ensure consistency with its volume-based derivation logic.

Therefore, the comparison is not intended to directly compare absolute height differences among different building height products. Instead, while the original products are kept unchanged, each product is evaluated against its definition-matched reference data using accuracy metrics such as RMSE, MAE, rRMSE, and R². Through this treatment, comparison bias caused by differences in height definition, spatial resolution, and spatial reference is reduced as much as possible within a more consistent and fair evaluation framework.

Based on the revised comparison results, our dataset achieves the lowest RMSE (4.58 m), rRMSE (0.48), and MAE (3.13 m), as well as a relatively high R² (0.77) in Beijing in 2019. This indicates that, after differences in height definitions among products are explicitly considered and matched reference data are used for evaluation, our dataset still shows good accuracy performance.

Table 1. Resolution and height calculation method for each dataset compared in Fig. 12 in this response letter, BA_i stands for building area with an i m resolution pixel.

Dataset	CNBH-10m (Wu et al., 2023)	Ma et al., 2024	WSF 3D (Esch et al., 2022)	Our dataset
Resolution	10m	150m	90m	30m
Height definitions	Building only	Building only	Building and built-up surface	Building and built-up surface
Height calculation	$\frac{\sum_i^{BA_{10}} H_i^{1m}}{BA_{10}}$	$\frac{\sum_i^{BA_{150}} H_i^{1m}}{BA_{150}}$	$\frac{\sum_i^{8100} H_i^{1m}}{8100}$	$\frac{\sum_i^{900} H_i^{1m}}{900}$

We thank the reviewer again for the constructive comment. Following the reviewer’s suggestion, Section 4.4.1 is revised as follows:

“...

For comparison with single-year building height datasets, Beijing in 2019 is selected as the study area because of its complex urban morphology and heterogeneous building height distribution.

...

To ensure a fair assessment across products with different height definitions, the original building height products are kept unchanged, and definition-matched reference data are generated from the same Baidu building height vectors. The vectors are first converted into a 1m height raster and then aggregated according to each product’s resolution, spatial reference, and height definition. For Wu et al. (2023) and Ma et al. (2024), only building-covered

pixels are averaged and aggregated to 10m and 150m, respectively; for WSF 3D (Esch et al., 2022), all pixels within each 90m grid are averaged to match its area-weighted, volume-derived height definition. Therefore, each product is evaluated against its matched reference data using consistent accuracy metrics, rather than by directly comparing absolute height values across products with different definitions. Notably, all accuracy metrics are calculated using the full validation dataset; the 45 m axis limit is used only for visualization to better show the dominant data distribution.

...

In Beijing in 2019, our dataset achieves the best overall accuracy, with the lowest RMSE (4.58 m), rRMSE (0.48), and MAE (3.13 m), as well as a relatively high R^2 (0.77) (Fig. 11b). CNBH-10m, WSF 3D and Ma's datasets expose visible biases, with RMSE differing from 6.50m to 14.39m and MAE differing from 4.43m to 9.71m (Fig. 11c-e).

....”

(2) The reported RMSE of 748.79 m is physically implausible and requires clarification.

We sincerely apologize for the physically implausible RMSE value (748.79 m) reported in the initial submission. This error occurred during the coding phase, where NaN values in the reference dataset are erroneously included in the numerical calculation. We have conducted a thorough review of our validation scripts and corrected this oversight. The updated RMSE for Ma's dataset is now 11.82 m, which is consistent with the validation results reported in recent literature (e.g., Fig. 15 in Chen et al., 2025). We deeply regret this technical error and have implemented stricter data-cleaning protocols for all reported metrics.

(3) Please note that it is unclear whether the same reference building height data were used consistently across Fig. 10b–e. Reference heights above 20 m appear in some panels (b–d) but not in others (e), and zero-height pixels are shown for Ma's dataset, although such values do not appear in Ma's original released data.

We thank the reviewer for pointing out this issue. As clarified in our response to Question 8(1), the original reference data source is consistent across Fig. 10b–e in the original manuscript, corresponding to Fig. 12b–e in this response letter. The same set of Baidu building height vectors is first converted into a 1m resolution height raster, and then aggregated according to the spatial resolution, spatial reference, and height definition of each evaluated product to generate matched reference data. Therefore, the reference data in Fig. 10b–e (Fig. 12b–e in this response letter) are derived from the same original source, and the apparent differences among panels are mainly caused by the previous data process, rather than by inconsistent reference data sources. We have carefully checked and updated the corresponding comparison results. As shown in the revised Fig. 10e (Fig. 12e in this response letter), the reference height distribution for Ma's dataset now correctly includes values above 20 m, and the estimated heights are consistent with the characteristics of Ma's original released data.

(4) It is unclear whether this part of the validation is limited to buildings with heights below 45 m.

We sincerely appreciate the reviewer's comment. We clarify that this part of the validation is not

limited to buildings below 45 m. All accuracy metrics are calculated using the full validation dataset, including buildings higher than 45 m. The axes are capped at 45 m only for visualization purposes in the figure. This is because the majority of data points are concentrated between 0 and 40 m; displaying the full height range would compress most points into a small corner of the plot and make the main data distribution difficult to inspect. As shown in the density plot in the revised Fig. 10 (Fig. 12 in this response letter), the 45 m axis limit provides a better visualization balance and clearer presentation of the dominant data distribution, but it does not affect the validation samples or the reported accuracy metrics. The following clarification is added in Section 4.4.1 of the revised manuscript:

“...Notably, all accuracy metrics are calculated using the full validation dataset; the 45 m axis limit is used only for visualization to better show the dominant data distribution....”

9. Overall, Section 4.4 requires careful re-examination.

(1) More detailed descriptions of the accuracy assessment and dataset comparison procedures are required, for instance, whether the validation is conducted at the pixel level or using object-based approaches based on building footprints (with height).

We thank the reviewer for suggesting a more detailed description of the accuracy assessment and dataset comparison procedures. We fully agree that clarifying whether validation is conducted at the pixel level or through an object-based approach based on building footprints is important for ensuring the transparency, reproducibility, and interpretability of the dataset evaluation.

Following the reviewer’s suggestion, more detailed descriptions of the accuracy assessment and dataset comparison procedures are added to the revised manuscript, especially regarding data processing, scale matching, and statistical methods used in different validation scenarios.

Specifically, in Section 4.4.1 “Comparison with single-year building height datasets”, following our response to Question 8, the pixel-level validation procedure for comparison with single-year building height datasets is further clarified. This section, based on Fig. 12 in this response letter, clarifies that the original products are kept unchanged, while the corresponding reference data are matched and aggregated according to each product’s spatial resolution, spatial reference, and height definition. Each product is then evaluated at the pixel level against its definition-matched reference data, rather than through a direct comparison of absolute height values among products with different definitions.

In Section 4.4.2 “Comparison with long-term building height products”, following our response to Question 6, the processing procedure for comparing long-term building height products is further supplemented. Based on Fig. 8 and Fig. 9 in this response letter, this section clarifies that validation is mainly based on pixel-level or grid-level statistics. It is also clarified that the building footprints in CMAB are used only to calculate the nationwide average building height in Fig. 8 in this response letter.

In Section 4.4.3 “Comparison with stable areas and historical imagery for temporal reliability validation”, following our response to Question 5, the temporal reliability validation procedure for the dataset is further clarified. This section, based on Fig. 5, Fig. 6, and Fig. 7 in this response letter, explains that the validation of stable areas, building change areas in sampled cities, and extrapolated change areas in non-sampled cities is also mainly based on pixel-level or grid-level results. Stable

areas are identified using pixels with no change year detected by CCDC, and their multi-year estimated height stability is then assessed. Building change areas are evaluated by comparing pixel-level height change trends with reference data or historical remote sensing images.

Overall, the revised manuscript more clearly distinguishes the data processing methods and statistical scales used in different validation scenarios. The validation and comparison in this study are mainly based on pixel-level/grid-level data processing and statistical assessment. Building footprint data are used only in specific comparisons as spatial statistical boundaries or reference data sources, and do not constitute object-based validation. We thank the reviewer again, as these clarifications help better explain the validation scale and processing workflow and improve the transparency and reproducibility of the assessment results.

(2) If possible, providing spatial uncertainty or confidence layers would substantially enhance the usability of the dataset.

We fully agree with the reviewer that providing a spatial uncertainty or confidence layer is essential for enhancing the usability and transparency of the dataset. In response to this suggestion, a per-pixel confidence layer is generated for the building height products based on the Mahalanobis distance (MD) in the feature space.

The MD-based confidence layer is used to describe the relative reliability of building height estimates by measuring the distance between each predicted pixel and the multivariate distribution of the training samples. Specifically, the Mahalanobis distance is calculated as:

$$D_M(x) = \sqrt{(x - \mu)^T \Sigma^{-1} (x - \mu)} \quad , \quad (1)$$

where $D_M(x)$ denotes the Mahalanobis distance from a given pixel x to the distribution, μ is the mean vector of the training samples, and Σ^{-1} is the inverse covariance matrix of the training samples.

A lower MD value indicates that the pixel is closer to the feature distribution of the training samples and therefore has higher confidence in the building height estimate. Conversely, a higher MD value suggests that the pixel is farther from the training distribution, indicating higher uncertainty or lower confidence (Lee et al., 2018). To make the uncertainty information easier to interpret, D_M^2 is further converted into a 0–1 confidence score using the survival function of the chi-squared distribution (Etherington, 2019; Frost, 2020):

$$D_M^2 \sim \chi^2(k) \quad (2)$$

$$\text{confidence score} = S(D_M^2) \quad (3)$$

where k denotes the degrees of freedom of the chi-squared distribution, corresponding to the number of feature variables used in each year, and $S(\cdot)$ is the survival function of the chi-squared distribution. The confidence score indicates how consistent a pixel is with the training sample distribution, with larger values representing pixels closer to the distribution center. To further generate a binary usability layer U , a practical threshold α is adopted, commonly set to 0.05 or 0.01 (Etherington, 2019). Pixels with confidence scores lower than α are more likely to fall outside the sample ellipsoid and are therefore regarded as potential outliers. In this study, $\alpha = 0.01$ is used to generate the binary confidence layer:

$$U = \begin{cases} 1 & \text{confidence score} \geq \alpha \\ 0 & \text{confidence score} < \alpha \end{cases} \quad (4)$$

where $U = 1$ indicates pixels within the acceptable range of the training feature distribution, and $U = 0$ indicates pixels that may be outliers or extrapolated samples and should be used with greater caution.

To evaluate the usability of the dataset, binary confidence layers are generated for representative large cities, including Beijing, Shanghai, and Chengdu. As shown in the binary confidence layers from 1990 to 2019 (Fig. 13 in this response letter), most building areas in these cities fall within the 99% confidence range, indicating that their feature distributions are close to the training samples and that the corresponding building height estimates are generally reliable.

Beijing (Fig. 13a in this response letter) and Shanghai (Fig. 13b in this response letter) show good spatial continuity in 1990, 1995, 2000, 2005, 2010, 2015, and 2019, with only a small number of sparsely distributed low-confidence patches. Chengdu (Fig. 13c in this response letter) also shows a high proportion of building areas within the 99% confidence range in most years, although a relatively concentrated low-confidence area appears in 1995. This pattern may be related to the smaller early built-up extent, differences in remote sensing feature quality, or local urban morphology that differs from the training sample distribution. Overall, these results demonstrate the reliability of the dataset across representative Chinese cities.




Fig. 13. Binary confidence layers for Beijing, Shanghai, and Chengdu from 1990 to 2019.

Following the reviewer's suggestion, a new Section 4.6, "Uncertainty analysis and Limitations", is added to the revised manuscript to further discuss the reliability, uncertainty, and usability of the dataset.

Three main aspects are described in the new Section 4.6. First, the spatial uncertainty of building height estimates is discussed based on the confidence layers. Second, the potential underestimation in high-rise building areas is clarified and identified as one of the directions for future improvement. Third, the spatial transferability and cross-city robustness of the model are further discussed using the validation results in non-sampled cities, as supplemented in our response to Question 2.

Through these revisions, the reliability, potential uncertainty, and usability of the dataset are more

clearly described in the revised manuscript. This helps users apply the dataset more appropriately across different urban contexts and building height types.

“... ”

Section 4.6 Uncertainty analysis and Limitations

Although the generated long-term building height dataset shows good overall accuracy and temporal consistency, several limitations and uncertainty sources should still be considered to support appropriate data use. Spatial uncertainty varies across regions because the reliability of building height estimates depends on the similarity between predicted pixels and the training samples in the feature space. Areas whose feature distributions deviate from the training samples may have higher uncertainty. This sample-representation issue is also reflected in some high-rise areas, where underestimation remains mainly due to the limited availability of high-rise reference samples.

In addition to feature-space similarity, uncertainty in the building height estimates may also arise from reconstructed VV, CCDC-based change detections, and the use of geographic coordinates as spatial predictors. The joint use of reconstructed VV and spectral and topographic variables suppresses random reconstruction noise during building height prediction. Although localized inaccuracies may appear as discrete low-confidence areas in heterogeneous landscapes, they do not distort the overall spatial patterns of the final 30 m building height maps. Uncertainty may also be higher in the early period because of the relative sparsity of CCDC-based change detections and the constrained availability of historical Landsat observations during the 1990s, which result in relatively lower transfer accuracy of reference data. The pre-2000 estimates still demonstrate high reliability for capturing regional-scale urban expansion, provincial building volume trends, and metropolitan-level development trajectories, although caution is advised when conducting fine-grained, pixel-level, year-to-year causal interpretations. Geographic coordinates help capture regional spatial trends and broad-scale variations in urban morphology, but may also introduce risks of spatial overfitting. The training dataset covers most of China’s topographical and climatic regions, allowing the models to capture the major patterns of urban morphology. Spatial transferability remains challenging in cities excluded from the reference samples, especially where local zoning or building morphology differs substantially from the training data.

To provide additional reliability information, a per-pixel confidence layer is generated using the Mahalanobis distance (MD), which measures the distance between each predicted pixel and the multivariate distribution of the training samples (Lee et al., 2018):

$$D_M(x) = \sqrt{(x - \mu)^T \Sigma^{-1} (x - \mu)} \quad (12)$$

where $D_M(x)$ denotes the Mahalanobis distance of pixel x , μ is the mean vector of the training samples, and Σ^{-1} is the inverse matrix. A lower MD value indicates that the pixel is closer to the training sample distribution and has higher confidence, whereas a higher MD value indicates greater uncertainty.

For easier interpretation, D_M^2 is converted into a 0–1 confidence score using the survival function of the chi-squared distribution. A threshold of $\alpha=0.01$ is then used to generate a binary confidence layer, where pixels within the 99% confidence range are regarded as reliable, while pixels outside this range are considered potential outliers or extrapolated samples and should be used with caution (Etherington, 2019; Frost, 2020).

The binary confidence layer is further examined in representative cities, including Beijing, Shanghai, and Chengdu (Fig. 20). From 1990 to 2019, most

building areas in these cities fall within the 99% confidence range, indicating that their feature distributions are generally close to the training samples and that the corresponding building height estimates are reliable. Beijing and Shanghai show good spatial continuity across all selected years, with only sparsely distributed low-confidence patches. Chengdu also shows a high proportion of reliable building areas in most years, although a relatively concentrated low-confidence area appears in 1995, possibly related to early built-up extent, remote sensing feature quality, or local morphology differences. These results demonstrate the reliability of the dataset across representative Chinese cities. Remaining uncertainties suggest that further refinement could be beneficial for improving the dataset across different urban contexts.

....”

References:

- [1] Lee, K., Lee, K., Lee, H., and Shin, J.: A simple unified framework for detecting out-of-distribution samples and adversarial attacks, *Advances in neural information processing systems*, 31, 2018.
- [2] Etherington, T. R.: Mahalanobis distances and ecological niche modelling: correcting a chi-squared probability error, *PeerJ*, 7, e6678, <https://doi.org/10.7717/peerj.6678>, 2019
- [3] Frost, H. R.: Variance-adjusted Mahalanobis (VAM): a fast and accurate method for cell-specific gene set scoring, *Nucleic acids research*, 48, e94, <https://doi.org/10.1093/nar/gkaa582>, 2020

(3) Please carefully check the reference formatting. The citation of Che et al. (2024) should refer to the final published version, rather than the Earth System Science Data Discussions version.

We appreciate the reviewer’s careful attention to the reference formatting. The manuscript is thoroughly checked, and the two citations, Che et al. and Zhu et al., are updated as requested:

“...Che, Y., Li, X., Liu, X., Wang, Y., Liao, W., Zheng, X., Zhang, X., Xu, X., Shi, Q., Zhu, J., Yuan, H., and Dai, Y.: 3D-GloBFP: the first global three-dimensional building footprint dataset, *Earth Syst. Sci. Data*, 16, 5357–5374, <https://doi.org/10.5194/essd-16-5357-2024>, 2024. ...”

“...Zhu, X. X., Chen, S., Zhang, F., Shi, Y., and Wang, Y.: GlobalBuildingAtlas: an open global and complete dataset of building polygons, heights and LoD1 3D models, *Earth Syst. Sci. Data*, 17, 6647–6668, <https://doi.org/10.5194/essd-17-6647-2025>, 2025. ...”

A comprehensive check is also performed of all other references to ensure they align with the journal’s formatting requirements and reflect the most recent published versions available.

Reviewer #2

General comments and responses

Overall, this is a valuable and timely study. Long-term, annual building-height reconstruction remains an important gap in urban remote sensing, especially at fine spatial resolution and national scale. The proposed dataset has clear potential to support studies of three-dimensional urban growth and renewal in China. Nevertheless, several methodological and interpretive issues still deserve further clarification. Addressing these points would improve the robustness, transparency, and broader usability of the dataset.

We would like to express our sincere gratitude to the editor and the reviewers for their valuable time, constructive comments, and thoughtful suggestions on our manuscript. In the following sections, we address each of the reviewer's points in detail, organizing our responses in a question (black)-response (blue)-revision (red) format to ensure clarity.

General comments and responses

1. On Page 9, Lines 205–210 and Page 10, Lines 217–219, the manuscript states that reconstructed 1 km VV backscatter is used as an independent variable for XGB-ReVV during 1990–2014. Since this VV layer is reconstructed from Landsat-derived variables and terrain/DSM information, its uncertainty may be propagated into the subsequent 30 m building-height prediction. Although the authors report satisfactory performance of the reconstruction model in the Beijing–Tianjin–Hebei region, it remains unclear whether this performance can be generalized to the whole of China, particularly across regions with different urban forms, climate conditions, and surface characteristics. I suggest that the authors add a short discussion on the robustness and potential error propagation of the reconstructed VV variable, and clarify the extent to which the regional validation supports its nationwide application.

We sincerely thank the reviewer for this constructive suggestion. It is agreed that the nationwide robustness of the 1 km VV reconstruction and its potential error propagation need to be evaluated to ensure the reliability of building height predictions during 1990–2014. To verify whether the 1 km VV reconstruction model is generalizable beyond the Beijing–Tianjin–Hebei region, validation is conducted across 18 cities representing 9 distinct ecological and geographic regions across China, using the actual 2019 Sentinel-1 VV observations as the reference. These cities are characterized by highly diverse urban forms, local climates, and topographic conditions. The reconstruction accuracy metrics are calculated based on the actual 2019 Sentinel-1 VV observations. The results summarized in Table 1 demonstrate the high robustness of the reconstructed 1 km VV at the national scale. Regarding the metrics for different regions, the lowest RMSE is obtained in Lanzhou, with an RMSE of 1.57 dB and an R^2 of 0.90, whereas the highest RMSE is observed in Guangzhou, with an RMSE of 2.45 dB and an R^2 of 0.84. This indicates that although the VV regression model is generally reliable, its predictive ability varies across regions. The national pooled RMSE is 1.97 dB, which is slightly larger than the 1.5 dB RMSE reported for the Beijing–Tianjin–Hebei region in the original article, with an rRMSE of 36.7%, a mean absolute error (MAE) of 1.34 dB, and an R^2 of 0.89. These results confirm that the reconstruction model can be safely applied at the national scale.

Table 1 The VV regression model accuracy in different cities across China in 2019

Geographic Regions	Representative Cities	Number of Pixels	RMSE (dB)	MAE (dB)	rRMSE	R²
Northern China Plain	Beijing	645	1.74	1.17	34.6%	0.90
	Shijiazhuang	398	1.91	1.29	27.4%	0.89
Northeast Plain	Harbin	451	1.92	1.35	28.9%	0.89
	Changchun	422	1.77	1.21	22.0%	0.87
Northwest Arid Zone	Lanzhou	160	1.57	1.26	34.6%	0.90
	Yinchuan	120	1.98	1.37	31.0%	0.85
Yangtze River Delta	Shanghai	548	1.98	1.29	28.6%	0.90
	Suzhou	463	1.97	1.26	34.1%	0.90
Yangtze River Basin	Wuhan	380	2.16	1.31	35.0%	0.89
	Changsha	320	2.34	1.41	38.1%	0.89
Sichuan Basin	Chengdu	493	1.97	1.29	38.1%	0.89
	Chongqing	179	1.74	1.17	30.7%	0.93
Southwest Highland	Kunming	130	1.68	1.18	40.1%	0.90
	Guiyang	250	1.81	1.31	31.3%	0.89
Southeast Coastal	Fuzhou	172	2.37	1.61	49.3%	0.82
	Xiamen	262	1.76	1.19	29.5%	0.92
Pearl River Delta	Guangzhou	267	2.45	1.35	56.1%	0.84
	Shenzhen	259	2.05	1.34	43.6%	0.86
National Pooled	All 59 cities	13997	1.97	1.34	36.7%	0.89

In addition to the nationwide validation results described above, the reconstructed VV backscatter is not used as an isolated predictor in the 30 m building-height model. Instead, it is integrated with the 30 m Landsat/Sentinel spectral bands and topographic variables to provide a complementary macro-scale structural constraint. In this framework, the 1 km reconstructed VV is introduced as a macro-scale spatial prior because its coarse spatial resolution mainly reflects regional built-up density and structural volume. By contrast, the 30 m Landsat/Sentinel spectral bands and topographic variables provide finer-scale spatial information, allowing local spatial and temporal details, such as building edges and height changes, to be captured (Frantz et al., 2021; Li et al., 2020; Dong et al., 2024). Therefore, potential local regression errors in the reconstructed VV are prevented from directly distorting the final 30 m building height maps through its integration with other remote-sensing variables. Although some false regressions may still exist and are reflected by the black squares in the confidence layer shown in Question 9(2) in the response letter to Reviewer #1 (Fig. 13 in this response letter), the overall spatial pattern of the final 30 m building height maps is preserved. Moreover, higher accuracy is achieved when the regressed VV is used as an input variable, as demonstrated in Fig. 9 of the original manuscript (Fig. 14 in this response letter).




Fig. 14. RMSE differences between models with and without reconstructed VV. Cropped from Fig. 9 of the original manuscript.

We again thank the reviewer for this valuable suggestion. Section 3.1.2 of the manuscript is updated to include a brief discussion on the reliability of the VV regression:

“ ...

The regression framework is further extended nationwide in this study, and the regressed VV backscatter is utilized as an independent variable in subsequent XGBoost-based models for building height prediction. The nationwide robustness of the 1 km reconstructed VV backscatter is confirmed through validation across diverse geographic regions characterized by different urban forms, local climates, and topographic conditions. Detailed accuracy metrics are provided in Supplementary Table S6.

...”

A detailed discussion is added in Section 4.6 of the manuscript to clarify the data uncertainty potentially introduced by the VV reconstruction:

“ ...

The joint use of reconstructed VV and spectral and topographic variables suppresses random reconstruction noise during building height prediction. Although localized inaccuracies may appear as discrete low-confidence areas in heterogeneous landscapes, they do not distort the overall spatial patterns of the final 30 m building height maps.

...”

References:

- [1] Frantz, D., Schug, F., Okujeni, A., Navacchi, C., Wagner, W., van der Linden, S., and Hostert, P.: National-scale mapping of building height using Sentinel-1 and Sentinel-2 time series, *Remote Sensing of Environment*, 252, 112128, <https://doi.org/10.1016/j.rse.2020.112128>, 2021
- [2] Li, X., Zhou, Y., Gong, P., Seto, K. C., and Clinton, N.: Developing a method to estimate building height from Sentinel-1 data, *Remote Sensing of Environment*, 240, 111705, <https://doi.org/10.1016/j.rse.2020.111705>, 2020
- [3] Dong, B., Zheng Q., Lin Y., Chen B., Ye Z., Huang C., Tong C., Li S., Deng J., and Wang, K.: Integrating physical model-based features and spatial contextual information to estimate building height in complex urban areas, *International Journal of Applied Earth Observation and Geoinformation*, 126, 103625, <https://doi.org/10.1016/j.jag.2023.103625>, 2024

2. On Page 7, Lines 148–151, the 2019 reference height data are derived from Baidu Maps building footprints, where floor counts are converted to height using a simple assumption of 3 m per floor. Although this approach is supported by Liu et al. (2021), the cited validation was based on a limited building-height assessment of only 519 buildings in Shenyang, with a reported mean height deviation of approximately 1 m and an accuracy of 86.8%. This raises the question of whether the same level of accuracy can be assumed for nationwide and multi-temporal building-height mapping. In addition, Section 4.1 would benefit from a clearer explanation of how the 2019 reference heights are transferred backward to 1990. In particular, if a pixel experienced in-situ height changes over the past three decades, such as demolition followed by reconstruction of a taller building, it would be helpful to clarify how such cases are handled during annual reference sample generation.

We thank the reviewer for these comments. Further scientific justification for the 3 m floor height assumption is provided, and the mathematical logic of our backward reference transfer mechanism is detailed in the response below.

Although the 86.8% accuracy is initially reported in Shenyang, the 3 m-per-floor conversion factor is widely validated and adopted as a highly stable vertical conversion parameter in national-scale urban morphology studies across China (Zhang, Zhao and Long, 2025; Wu et al., 2023). Importantly, the dataset is generated at a 30 m grid resolution, where each pixel represents an aggregated measure of building height derived from all sub-pixel structures within the corresponding cell. At this spatial scale, fine-grained architectural variations, such as high-ceiling spaces, rooftop decorations, and floor-level heterogeneity, are inherently smoothed through spatial aggregation, consistent with the scale-dependent nature of remote sensing observations (Woodcock and Strahler, 1987; Zhang et al., 2017). Consequently, the 3 m-per-floor assumption is used as an effective first-order approximation for large-scale gridded urban morphology representation, rather than as an exact building-level physical parameter.

To prevent historical reference labels from being contaminated by urban renewal such as in-situ demolition and taller reconstruction, the Continuous Change Detection and Classification (CCDC) algorithm is utilized as a temporal filter. For any target historical year t ($1990 \leq t \leq 2018$) and pixel (x, y) , the annual reference height $H_t(x, y)$ is determined by:

$$H_t(x, y) = \begin{cases} H_{2019}(x, y) & \text{if } T_{break}(x, y) = 0 \text{ or } T_{break}(x, y) < t \\ Excluded & \text{otherwise} \end{cases} \quad (1)$$

where $T_{break}(x, y)$ is the latest breakpoint year detected by the CCDC algorithm for that pixel, with 0 meaning no breakpoint detected. This method guarantees that the 2019 reference data set is only projected backward to year t if the land surface has remained spectrally stable and undisturbed since year t . If any spectral breakpoint is detected within the window $[t, 2019]$, the pixel is completely excluded from the training-validation pool for year t . This strictly prevents label mismatch or label pollution caused by urban renewal, ensuring high purity in our historical reference samples.

We thank the reviewer again for this valuable suggestion. Section 2.6 of the original manuscript is modified to include the following content:

“ ...

The 3 m-per-floor conversion is a statistically robust approximation for grid-level aggregation supported by different researches, although being unable to fully capture localized floor-height variations in specialized structures (Zhang, Zhao and Long, 2025; Wu et al., 2023).

...”

Section 3.1.1 of the original manuscript is updated, including the $H_t(x, y)$ formula to explain the annual reference extraction process using a more systematic approach:

“ ...

Given a target historical year t ($1990 \leq t \leq 2018$) and pixel (x, y) , the annual reference height $H_t(x, y)$ is determined by:

$$H_t(x, y) = \begin{cases} H_{2019}(x, y) & \text{if } T_{break}(x, y) = 0 \text{ or } T_{break}(x, y) < t \\ \text{Excluded} & \text{otherwise} \end{cases}$$

where $T_{break}(x, y)$ is the latest breakpoint year detected by the CCDC algorithm for that pixel, with 0 meaning no breakpoint detected. This method guarantees that the 2019 reference height is only projected backward to year t if the land surface has remained spectrally stable and undisturbed since year t . If any spectral breakpoint is detected within the window $[t, 2019]$, the pixel is completely excluded from the training-validation pool for year t .

...”

In addition, a comprehensive uncertainty analysis is added in Section 4.6, which quantifies and discusses potential error sources including those stemming from reference data (Reviewer #2, Question 2), CCDC-based temporal extraction (Reviewer #2, Question 3), and the use of spatial coordinates (Reviewer #2, Question 5). Furthermore, as detailed in our response to Reviewer #1 in Question 9, confidence layers are produced to represent these uncertainties spatially. These layers, along with a full methodological discussion, have been incorporated into the final manuscript to provide a transparent assessment of our model's predictive reliability.

References:

- [1] Zhang, Y., Zhao, H., and Long, Y.: CMAB: A Multi-Attribute Building Dataset of China, Scientific Data, 12(1), 430, <https://doi.org/10.1038/s41597-025-04730-5>, 2025
- [2] Wu, W. B., Ma, J., Banzhaf, E., Meadows, M. E., Yu, Z. W., Guo, F. X., Sengupta, D., Cai, X. X., and Zhao, B.: A first Chinese building height estimate at 10 m resolution (CNBH-10 m) using multi-source earth observations and machine learning, Remote Sensing of Environment, 291, 113578, <https://doi.org/10.1016/j.rse.2023.113578>, 2023
- [3] Kleman, J. and Fagerlund, E.: Influence of different nitrogen and irrigation treatments on the spectral reflectance of barley, Remote Sensing of Environment, 21.1, 1-14, [https://doi.org/10.1016/0034-4257\(87\)90002-2](https://doi.org/10.1016/0034-4257(87)90002-2), 1987
- [4] Zhang, L., Weng, Q. and Shao, Z.: An evaluation of monthly impervious surface dynamics by fusing Landsat and MODIS time series in the Pearl River Delta, China, from 2000 to 2015, Remote sensing of environment, 201, 99-114, <https://doi.org/10.1016/j.rse.2017.08.036>, 2017

3. On Page 9, Lines 195–198, the authors state that applying the 1999 CCDC mask uniformly to 1990–1999 has “negligible effects” on reference data accuracy. However, Page 13, Lines 271–274 show that the validation accuracy of annual reference samples remains only around 65%–70% during 1990–2002. Therefore, I suggest replacing “negligible effects” with a more cautious statement, such as “limited but non-negligible uncertainty may remain for the pre-2000 period.” The authors may also clarify in Section 4.1 or Section 6 that early-year estimates are more suitable for regional-scale or trend-level analyses, rather than for strong interpretation of pixel-level inter-annual changes.

We completely agree with the reviewer's insightful comment and sincerely appreciate the suggestion reviewer has provided. Acknowledging early-year uncertainties in a transparent manner improves the scientific rigor of our dataset. Section 3.1.1 of the manuscript is updated, replacing "negligible effects" with "limited but non-negligible uncertainty may remain for the pre-2000 period."

“ ...
which may introduce minor yet non-negligible uncertainty to the pre-2000
period, keeping the overall accuracy around 65%.
...”

Furthermore, a clear guidance for data users is added in the revised Section 4.6. It is clarified that while the early-year estimates (1990–1999) are highly reliable for capturing regional-scale urban expansion, provincial building volume trends, and metropolitan-level development trajectories, users should exercise caution when conducting fine-grained, pixel-level, year-to-year causal interpretations in the 1990s due to the relative scarcity of CCDC change detections and historical Landsat observations in that era:

“ ...
Uncertainty may also be higher in the early period because of the relative
sparsity of CCDC-based change detections and the constrained availability of
historical Landsat observations during the 1990s, which result in relatively lower
transfer accuracy of reference data. The pre-2000 estimates still demonstrate
high reliability for capturing regional-scale urban expansion, provincial building
volume trends, and metropolitan-level development trajectories, although
caution is advised when conducting fine-grained, pixel-level, year-to-year causal
interpretations.
...”

4. On Page 10–11, Lines 225–228, the manuscript states that the annual reference height samples are randomly split into 90% training and 10% testing sets. Since both subsets are derived from the same CCDC-filtered and 2019-reference-based sample generation procedure, this evaluation is better interpreted as testing the consistency of model fitting against CCDC-derived reference labels, rather than as a fully independent validation of spatial or temporal generalization. I suggest that the authors make this distinction clearer when presenting the accuracy results, so that the reported RMSE, MAE, and R^2 are not over-interpreted as completely independent validation metrics.

We sincerely appreciate this insightful methodological observation provided by the reviewer, which encourages further refinement of the validation discussion. It is clarified that each year is treated as an independent entity in our modeling framework. Specifically, the training and validation sets for a given year t are derived exclusively from the CCDC-filtered reference pool for that year, ensuring that no temporal data leakage occurs between years. Although the reference labels originate from the same long-term CCDC-processed dataset, the labels for each year t are independently generated and are used only to train and validate the model created specifically for that corresponding year, rather than to fit a single model using all reference data generated across years. Consequently, the annual accuracy metrics, including RMSE, MAE, and R^2 , reflect the model performance on year-specific independent data, as described in Section 3.3 of the original manuscript:

"XGBoost models are adopted due to their demonstrated computational efficiency and predictive accuracy in prior architectural height regression studies (Che et al., 2024; Stipek et al., 2024). For each year, an independent XGBoost model is trained with year-specific variables extracted from remote sensing data, paired with corresponding reference building heights derived from CCDC masks. This approach maximizes the utility of annual data characteristics in the estimation process."

Furthermore, independent validations of spatial and temporal generalization, in addition to the annual accuracy metrics, are performed and demonstrated in the original manuscript. For spatial independent validation, the annual estimates are validated using a completely independent external database of housing transaction records from Lianjia, which contains real community construction years, as shown in Fig. 1 in the Response to Reviewer #1. Every city validated using the Lianjia dataset is excluded from the sampled cities, ensuring minimum correlation with the training set. In this way, the spatial transferability of the models is verified.

For temporal independent validation, historical visual audits are conducted in multiple non-sampled cities that are completely excluded from the training dataset, such as Handan, Jinjing, and Anyang, as shown in Fig. 6 in the Response to Reviewer #1 (the revised Fig. 14 of the original manuscript). These audits demonstrate the temporal accuracy of the models. As shown in Fig. 4 in the response letter, the dataset matches the historical construction timeline and reflects the consistent dynamics of urban development.

Finally, the results are benchmarked against entirely independent external datasets that do not share our reference data or modeling framework for cross-dataset validation, including WSF 3D, Ma's LiDAR-based height dataset and Wu's CNBH-10m as shown in Fig. 12 in the Response to Reviewer #1 (the revised Fig. 10 of the original manuscript, updated due to computational error). This cross-dataset validation strategy ensures the independence of our model evaluation and underscores the robustness of the product when compared against datasets derived from different sources.

We thank the reviewer for this valuable comment. According to the reviewer's suggestion, the year-specific nature of the models is further clarified in the modified manuscript. For example, Section 4.2.1 is revised as follows:

"...

The accuracy of the three XGBoost-based models is further evaluated for three representative years: XGB-ReVV in 2014, XGB-S1 in 2017, and XGB-S1S2 in 2019 (Fig. 7).

...”

5. Page 16, Lines 335–342 show that latitude and longitude are the most important predictors, even exceeding most optical and SAR variables. Meanwhile, Page 15–16, Lines 318–326 report clear regional discrepancies in model performance, especially poorer performance in parts of southern China and the Pearl River Delta. This suggests that the model may partly rely on regional urban-form priors rather than only pixel-level remote-sensing signals. I recommend that the authors add a short discussion on whether the strong importance of geographic coordinates helps capture spatial heterogeneity, or whether it may introduce risks when extrapolating to regions or cities with limited representation in the training samples.

We thank the reviewer for raising this key geographic modeling question. In national-scale modeling, latitude and longitude serve as spatial trend surface predictors that capture macro-geographic variations in urban morphology (Wu et al., 2023; Chen and Zhao, 2022; Meyer et al., 2019). In China, building forms and zoning regulations are heavily influenced by geographic locations. For example, northern cities have strict regulations for building spacing and heights to ensure winter solar access, resulting in lower building density and regular heights, whereas southern and coastal cities exhibit high-density, compact, and high-rise vertical patterns. (Yang, Xuan and Zhou, 2022; Wang et al., 2024). Introducing coordinates allows the XGBoost model to capture these regional urban-form priors and resolve spatial non-stationarity.

While these coordinates effectively account for spatial variations, their impact on temporal dynamics remains neutral. Specifically, since latitude and longitude are static variables over time, they cannot explain any temporal variations. Thus, all predicted annual height growth, demolition, and reconstruction dynamics over the 30-year period are strictly driven by the time-varying, physical remote sensing features, which are Landsat and Sentinel spectral composites and radar backscatter. This prevents coordinates from causing temporal extrapolation failures (Meyer et al., 2019)

Regarding the lower accuracy in parts of southern China including the Pearl River Delta, apart from coordinate overfitting, natural environmental physical noise is usually an inevitable cause. Frequent cloud cover and heavy precipitation in Southern China degrade the quality of annual optical composites (Zhang and Weng, 2016). Moreover, southern cities are characterized by dense sub-tropical evergreen vegetation, which interacts with high-density vertical structures, causing severe radar double-bounce and volume scattering noise in 30 m pixels (Li et al., 2020).

Upon model training, while highly autocorrelated coordinates can lead to overfitting, this risk is effectively controlled in our model, as presented in Fig. 2 in the Response to Reviewer #1. Our dataset is trained on approximately 12 million pixels across 59 geographically dispersed cities, which ensures a highly continuous and representative coordinate space.

To further evaluate the influence of geographic coordinates and mitigate potential concerns regarding spatial overfitting, an additional ablation experiment is conducted. A separate model is trained for the year 2019 using an identical training protocol, but excluding latitude and longitude as input predictors. The accuracy metrics of two models are presented in Table 2.

Table 2 Accuracy of model w/ and w/o coordinates in 2019

XGB-S1S2 in 2019	RMSE (m)	MAE (m)	rRMSE	R ²
------------------	----------	---------	-------	----------------

w/ coords	3.90	2.70	61.0%	0.80
w/o coords	4.06	2.88	63.2%	0.78

The results indicated that while the model excluding coordinates maintained consistent predictive trends, the inclusion of geographic coordinates yielded superior performance across all evaluated metrics, having the model with coordinates acquiring a RMSE of 3.90m, MAE of 2.70m and R^2 of 0.80, the model without coordinates falling behind with a RMSE of 4.06m, MAE of 2.88m and R^2 of 0.78. This improvement confirms that geographic coordinates effectively capture regional spatial heterogeneity that are not fully represented by remote-sensing signals alone. Geographic coordinates function as a supportive regional prior, thereby reinforcing the model's reliability across diverse urban landscapes.

To clarify the role of geographic coordinates and address potential concerns regarding spatial overfitting, Section 4.3 has been revised, explicitly distinguishing between the role of coordinates as macro-scale spatial priors and the reliance on time-varying remote sensing signals for detecting annual changes, providing a more robust validation of our model's spatial-temporal consistency:

“ ...

Geographic coordinates serve as essential spatial trend surface predictors that allow the XGBoost model to resolve spatial non-stationarity across China's diverse urban morphologies. By functioning as supportive regional priors, these coordinates capture macro-scale variations that are not fully represented by remote-sensing signals alone (Chen and Zhao, 2022). On the other hand, as these variables are static over time, they remain decoupled from inter-annual temporal dynamics, ensuring that all detected height changes are driven solely by time-varying physical features (Meyer et al., 2019)."

...”

Section 4.6 is updated as well to further discuss the uncertainty coordinates may introduce to the dataset:

“ ...

Geographic coordinates help capture regional spatial trends and broad-scale variations in urban morphology, but may also introduce risks of spatial overfitting. The training dataset covers most of China's topographical and climatic regions, allowing the models to capture the major patterns of urban morphology. Spatial transferability remains challenging in cities excluded from the reference samples, especially where local zoning or building morphology differs substantially from the training data.

...”

References:

- [1] Wu, W. B., Ma, J., Banzhaf, E., Meadows, M. E., Yu, Z. W., Guo, F. X., Sengupta, D., Cai, X. X., and Zhao, B.: A first Chinese building height estimate at 10 m resolution (CNBH-10 m) using multi-source earth observations and machine learning, Remote Sensing of Environment, 291, 113578, <https://doi.org/10.1016/j.rse.2023.113578>, 2023
- [2] Chen, Y. and Zhao, S.: A building height dataset across China in 2017 estimated by the spatially-informed approach, Scientific Data, 9.1, 76, <https://doi.org/10.1038/s41597-022-01192-x>, 2022

- [3] Meyer, H., Reudenbach, C., Wöllauer, S. and Nauss, T.: Importance of spatial predictor variable selection in machine learning applications—Moving from data reproduction to spatial prediction, *Ecological Modelling*, 411, 108815, <https://doi.org/10.1016/j.ecolmodel.2019.108815>, 2019
- [4] Yang, G., Xuan, Y. and Zhou Z.: Influence of building density on outdoor thermal environment of residential area in cities with different climatic zones in China—Taking Guangzhou, Wuhan, Beijing, and Harbin as examples, *Buildings*, 12.3, 370, <https://doi.org/10.3390/buildings12030370>, 2022
- [5] Wang, Y., Sun, G., Wu, Y. and Rosenberg, M. W.: Urban 3D building morphology and energy consumption: empirical evidence from 53 cities in China, *Scientific Reports*, 14.1, 12887, <https://doi.org/10.1038/s41598-024-63698-1>, 2024
- [6] Zhang, L. and Weng, Q.: Annual dynamics of impervious surface in the Pearl River Delta, China, from 1988 to 2013, using time series Landsat imagery, *ISPRS Journal of Photogrammetry and Remote Sensing*, 113, 86-96, <https://doi.org/10.1016/j.isprsjprs.2016.01.003>, 2016
- [7] Li, M., Koks, E., Taubenböck, H., and van Vliet, J.: Continental-scale mapping and analysis of 3D building structure, *Remote Sensing of Environment*, 245, 111859, <https://doi.org/10.1016/j.rse.2020.111859>, 2020

6. Several numerical values should be checked for consistency. The Abstract reports an RMSE range of 5.96–6.69 m, whereas Page 13, Lines 283–284 report 5.94–6.69 m, and Page 29, Lines 490–491 again report 5.96–6.69 m. Similarly, Page 27, Line 460 reports the 2019 national building volume as 884.69 km³, while Page 27, Lines 463–464 report 884.49 km³. I suggest that the authors harmonize these values throughout the manuscript to avoid giving the impression of inconsistent data versions or calculation procedures.

We apologize for these typographical inconsistencies. These discrepancies are remnants of earlier dataset versions and intermediate model runs. The entire manuscript is thoroughly audited, and all numerical statistics are harmonized:

- The national RMSE range is strictly standardized to 3.20-4.27 m in all sections of the manuscript. This value range is changed because new models are trained based on the adopted new sampling method, resulting in updated RMSE values.
- The 2019 national building volume is strictly standardized to 882.43 km³ across all pages.

All other figures, tables, and supplementary materials are double-checked to ensure complete numerical consistency throughout the revised manuscript.