



GlobalBuildingAtlas: an open global and complete dataset of building polygons, heights and LoD1 3D models

Xiao Xiang Zhu^{1,2}, Sining Chen^{1,2,★}, Fahong Zhang^{1,★}, Yilei Shi¹, and Yuanyuan Wang¹

¹Technical University of Munich, Munich, Germany

²Munich Center for Machine Learning, Munich, Germany

★These authors contributed equally to this work.

Correspondence: Xiao Xiang Zhu (xiaoxiang.zhu@tum.de)

Received: 2 June 2025 – Discussion started: 8 July 2025

Revised: 8 September 2025 – Accepted: 12 October 2025 – Published: 1 December 2025

Abstract. We introduce GlobalBuildingAtlas, a publicly available dataset providing global and complete coverage of building polygons, heights and Level of Detail 1 (LoD1) 3D building models. This is the first open dataset to offer high quality, consistent, and complete building data in 2D and 3D form at the individual building level on a global scale. Towards this dataset, we developed machine learning-based pipelines to derive building polygons and heights (called GBA.Height) from global PlanetScope satellite data, respectively. Also a quality-based fusion strategy was employed to generate higher-quality polygons (called GBA.Polygon) based on existing open building polygons, including our own derived one. With more than 2.75 billion buildings worldwide, GBA.Polygon surpasses the most comprehensive database to date by more than 1 billion buildings. GBA.Height offers the most detailed and accurate global 3D building height maps to date, achieving a spatial resolution of $3\text{ m} \times 3\text{ m}$ – 30 times finer than previous global products (90 m), enabling a high-resolution and reliable analysis of building volumes at both local and global scales. Finally, we generated a global LoD1 building model (called GBA.LoD1) from the resulting GBA.Polygon and GBA.Height. GBA.LoD1 represents the first complete global LoD1 building models, including 2.68 billion building instances with predicted heights, i.e., with a height completeness of more than 97 %, achieving RMSEs ranging from 1.5 to 8.9 m across different continents. With its height accuracy, comprehensive global coverage and rich spatial details, GlobalBuildingAtlas offers novel insights on the status quo of global buildings, which unlocks unprecedented geospatial analysis possibilities, as showcased by a better illustration of where people live and a more comprehensive monitoring of the progress on the 11th Sustainable Development Goal of the United Nations. The code is publicly available at <https://github.com/zhu-xlab/GlobalBuildingAtlas> (last access: 1 November 2025). The GBA dataset described in this manuscript can be accessed on mediaTUM under <https://doi.org/10.14459/2025mp1782307> (Zhu et al., 2025b).

1 Introduction

Buildings anchor human life and define the form and function of urban environments. According to the United Nations (UN), over 50 % of the global population currently reside in cities and considering the ongoing urbanization, it is estimated that this ratio will reach nearly 70 % by 2050 (United Nations, 2024). While urbanization enables more people to

live in connected communities, it also introduces a range of challenges, including inequalities and urban poverty, inadequate transportation infrastructure, air pollution, and limited access to open public spaces. In response, the UN has established “sustainable cities and communities” as its 11th Sustainable Development Goal (SDG 11) in its 2030 Agenda.

To monitor the progress toward this goal, one of the key indicators being used is the ratio of the land consumption

rate to the population growth rate (Indicator 11.3.1, United Nations, 2025) – a ratio that measures how much land is being developed or consumed relative to the rate at which the population is growing. This key indicator is computed based on the built-up area. Although this metric provides some insights into the spatial distribution of the built environment, its reliance on two-dimensional measurements can lead to biased or misleading interpretations. Cities are inherently three-dimensional (3D), and ignoring vertical space overlooks crucial information about how urban space is actually used. For example, a densely populated informal settlement may have the same built-up area per capita as a well-designed urban neighborhood with multi-story buildings, despite their vastly different spatial, social and infrastructural conditions. Therefore, a more complete understanding of the actual urban form requires information on building heights, which would enable obtaining a volumetric perspective of the built environment. Such 3D insights are essential for urban planning, infrastructure management and policy-making – especially in resource-limited contexts where the strategic allocation of funding and intervention is critical.

Despite its importance, the availability of comprehensive building height data remains unevenly distributed across the globe. While developed countries often leverage advanced technologies to monitor urban development, many regions undergoing rapid change – particularly in the Global South – lack the observational infrastructure and technical capacity to do so. Notably, these are often the areas that are most vulnerable and most in need of accurate, timely data. A promising alternative lies in Earth observations from space, which can offer scalable and repeatable insights independent of local ground-based or aerial-based resources.

However, existing global-scale building products utilizing data collected from space fall short in several key aspects. First, true global coverage and completeness remain limited, with most datasets covering only continental or national extents. Second, many existing approaches yield only coarse, aggregate-level building representations, which are insufficient for applications requiring more detailed, building-level information. Third, the reliance on ancillary data – such as socioeconomic variables, like population – hinders the scalability and timeliness of such product, especially in fast-changing urban contexts where rapid updates are critical.

To address these limitations, we introduce GlobalBuildingAtlas (GBA), a global-scale dataset originating from a pipeline that relies exclusively on optical satellite imagery. GBA provides comprehensive building polygons, heights and level of detail 1 (LoD1) 3D building models. Our main contributions are

- *GBA.Polygon*. The first complete set of global building polygons consisting of 2.75 billion buildings, closing the current gap of more than 40 % global buildings that were previously not accounted for.

- *GBA.Height*. The most detailed and accurate global 3D building height map to date, achieving an unprecedented spatial resolution of $3\text{ m} \times 3\text{ m}$ – 30 times finer than previous global products (90 m), and enabling a high-resolution and reliable analysis of building volumes at both local and global scales, achieving root mean square errors (RMSEs) ranging from 38.0 to 580.0 m^3 per 100 m^2 across different continents.
- *GBA.LoD1*. The first complete global LoD1 building model, including 2.68 billion building instances with predicted heights, which achieves RMSEs ranging from 1.5 to 8.9 m across different continents.

2 Related Work

3D building information is crucial for many purposes, including urban planning. Such information captures structural attributes that extend beyond the ground footprint of the building, including the building height and volume. Table 1 summarizes existing mainstream products that primarily map 3D building information in two main aspects: raster-based and instance-level building representations.

2.1 Raster-based Building Representation Products

Raster-based building representations provide the statistics of buildings in regular grids. The products are delivered as rasters showing aggregated building heights or total building volumes in a grid format, typically spanning tens to hundreds of meters. The results are often estimated from diverse data sources using simple regression models. For example, World Settlement Footprint (WSF) 3D (Esch et al., 2020, 2022) quantifies average building heights, total volumes and building fractions at 90 m resolution. Global Human Settlement Layers (GHSL) (Pesaresi et al., 2024) model global building distributions, with GHS-BUILT-H (Pesaresi and Politis, 2023a) depicting building heights and GHS-BUILT-V (Pesaresi and Politis, 2023b) depicting volumes at 250 m resolution. Ma et al. (2024) provided global building heights at a spatial resolution of 150 m. Generally, these products are derived either from low-resolution data sources of multiple modalities, such as Sentinel-1 Synthetic Aperture Radar (SAR) images and Sentinel-2 multi-spectral images, or from global digital elevation model (DEM) sources at coarse resolution. To achieve higher resolution mapping, recently Google (Sirko et al., 2023) demonstrated the potential of optical satellite imagery for the rapid retrieval of building information by generating building presence and height data at a 4 m resolution using only Sentinel-2 images. However, this dataset is not globally available yet and is primarily limited to regions in the Global South.

In summary, raster-based products of this kind usually have a large coverage or even global coverage, but with the compromise of a lower quality and resolution, restricting

Table 1. Comparison of existing large-scale building map and building height products. BF: Building footprint; BH: Building height; B: billion.

Product	Data Structure	Data Sources	BF Coverage	BH Coverage	Resolution	Building Count/with Height	Year
WSF 3D	Raster	S1, S2, TanDEM-X	Global	Global	90 m	/	vary from 2011–2019
GHS-BUILT-H	Raster	5 Global DEM sources	Global	Global	250 m	/	vary from 2000–2011
GBH 2020	Raster	LS-8, S1, S2	Global	Global	150 m	/	vary from 2019–2021
Google 2.5D	Raster	S2	Global South	Global South	4 m	/	2016–2023 Yearly
GBA.Height (ours)	Raster	PlanetScope	Global	Global	3 m	/	2019
Microsoft	Vector	Multi-source VHR images	Global excl. China	NA, EU, OC	/	1.4 B/0.28 B	vary from 2014–2021
OSM	Vector	Human Annotation	Global in part	Global in part	/	0.45 B/0.02 B	since 2008
3D-GloBFP	Vector	10 different sources	Global	Global	/	1.7 B/1.7 B	vary from 2014–2021
GBA.LoD1 (ours)	Vector	4 BF sources, PlanetScope	Global	Global	/	2.75 B/2.68 B	2019

their use in applications requiring high granularity at the instance level.

2.2 Instance-level Building Representation Products

Instance-level building representation maps delineate individual buildings and assign their corresponding 3D features. There already exist several 3D building model products at the instance level, albeit on a smaller scale (e.g., city or country). For instance, Peters et al. (2022) created 3D-BAG, the first open 3D building dataset based on Light Detection and Ranging (LiDAR) data covering the Netherlands; while the City of Helsinki (2017) released a high-quality 3D semantic city model derived from high-resolution aerial photographs. Similar models are available in other developed areas worldwide, e.g., Philadelphia by Pennsylvania State University (2015), and 260 cities in Japan by the Ministry of Land, Infrastructure, Transport and Tourism (MLIT) Japan (2024). Although they are derived from high-quality observations, and are capable of reconstructing detailed 3D building structures at Level of Detail 2 (LoD2) or higher, such products are not widely available globally due to their high operational costs and significant computational demands.

In addition to the aforementioned observation-based techniques, Volunteered Geographic Information (VGI) offers an alternative approach. While VGI – based on human annotations – is often recognized for its positional accuracy, it typically suffers from limited completeness, particularly with respect to building height information. Therefore, current mainstream products are primarily derived from aerial or satellite remote sensing data using algorithmic approaches. Microsoft (Microsoft, 2024) maintains a dataset of 1.4 billion building footprints, of which approximately 20 % include estimated building heights. Che et al. (2024) provided 3D-GloBFP – a 3D building footprint dataset encompassing 1.66 billion buildings. Despite achieving substantial coverage in terms of building footprints and height estimates, none of these datasets are derived from complete or globally consistent sources. Consequently, the total number of buildings represented falls significantly short of the United Nations’ estimate of approximately 4 billion buildings world-

wide (United Nations Human Settlement Program: UN Habitat, 2019).

2.3 Gaps in Existing Products

While raster-based building representation products prioritize broader coverage, they suffer from low resolution and quality limitations. Instance-level representation products, on the other hand, offer detailed information on individual buildings but are constrained by limited coverage and incomplete datasets. Moreover, the production of high-level building representation products often involves multiple data sources, complicating the process of updating these datasets efficiently.

To address these gaps, we propose a novel pipeline to generate global LoD1 building models based solely on PlanetScope images. Leveraging this approach, we generated 3 m-resolution raster height maps alongside globally consistent building polygons. By fusing the generated building polygons with existing open-source building footprints based on quality assessments, we constructed a global building footprint dataset optimized for the highest possible quality. Finally, we assigned building heights to every individual building polygon using our generated height maps, which led to the first complete global LoD1 building models.

3 Data Sources

This section introduces all the data sources that were used to develop the GlobalBuildingAtlas dataset.

3.1 Satellite Optical Images

We collected global-scale PlanetScope Surface Reflectance (PSR) imagery to support polygonal building mapping and height estimation. PSR data are orthorectified, multi-spectral satellite images provided by Planet’s satellite constellation. These images undergo atmospheric correction, ensuring that the reflectance values accurately represent the Earth’s surface properties. The PSR data we utilized include four spectral bands – three in the visible RGB range and one near-infrared (NIR) band, with a spatial resolution of approxi-

Table 2. Comparison of existing large-scale building footprint products. B: billion.

Product	Coverage	Number of Buildings	Year
Microsoft	Global excl. China	1.4 B	vary from 2014–2021
OSM	Global in part	0.45 B	since 2008
Open Buildings	Global South	1.8 B	until 2021
CLSM	East Asia	0.28 B	vary from 2020–2022
3D-GloBFP	Global	1.7 B	vary from 2014–2021
GBA.Polygon (ours)	Global	2.75 B	vary from 2014–2021




Figure 1. Distribution of city-scale regions of interest where training and test data were collected. The 3D and 2D training sets were used to train the building height estimation and building polygon generation pipelines, respectively. Basemap © CARTO © OpenStreetMap contributors 2025. Distributed under the Open Data Commons Open Database License (ODbL) v1.0.

mately 3 m. Additionally, the PlanetScope constellation offers a high temporal revisit frequency, capturing imagery of the same location up to daily, which is beneficial for monitoring dynamic urban environments.

3.2 Building Footprints

We utilized building footprint datasets from various sources (see Table 2), including OSM (OpenStreetMap contributors, 2025), Google Open buildings (Google Research, 2023), Microsoft Building Footprints (Microsoft, 2024), and CLSM (Shi et al., 2024). Since none of the above-mentioned footprints is complete, we also generated our own global building polygons from an updated version of GlobalBuildingMap (Zhu et al., 2024), with a novel building regularization and polygonization pipeline applied. The building mapping and regularization networks were trained on PSR imagery using manually filtered OpenStreetMap (OSM) annotations

with high completeness. The data distribution is illustrated in Fig. 1. To produce the final complete global building footprint dataset – GBA.Polygon – we integrated all available sources, including our own generated polygons, using a quality-based fusion approach.

3.3 LiDAR

As the most precise 3D measurement available, we used aerial LiDAR data covering a wide range of geolocations (see Fig. 1) in developed countries. A comprehensive list of the covered countries is provided in Appendix A. These data are publicly released by governments for open use. Due to the high operational cost of LiDAR observations, no such data are available in Africa. LiDAR observations were primarily acquired in 2019 to ensure temporal consistency with the PSR data. When 2019 acquisitions were not available, data from adjacent years were utilized. Since the coverage is

concentrated in developed regions, where building changes are generally limited, the temporal discrepancies between the PSR data and the 3D labels are expected to be negligible. We processed the LiDAR data into normalized digital surface models (nDSMs), which represent the heights of ground objects. These nDSMs served as reference data for training our neural networks.

4 Methodology

4.1 Overall Workflow

The overall workflow for generating the GlobalBuildingAtlas dataset is illustrated in Fig. 2. The workflow can be divided into 4 sections, including global data acquisition, global building polygon generation, global building height estimation and post processing, i.e., global LoD1 model generation.

4.2 Global Data Acquisition

This section details the acquisition, organization and pre-processing of PSR data to prepare them for the subsequent deep learning models and LoD1 model generation. We divided the Earth's surface into grid cells of $0.2^\circ \times 0.2^\circ$. Grids overlapping with built-up areas, as defined by the Global Urban Footprint (GUF) dataset (Esch et al., 2010, 2011, 2012, 2013), were selected as our areas of interest. For these regions, we acquired satellite imagery from 2019 that intersects with the selected grids, aiming to maximize coverage. In total, approximately 800 000 PlanetScope scenes were downloaded, each with a spatial resolution of 3 m and covering an area of about 287.5 km².

To ensure data quality, we applied a cloud filter that retained only scenes with less than 10 % cloud cover. In areas where suitable cloud-free imagery from 2019 was unavailable, the dataset was supplemented with images from 2018. For each $0.2^\circ \times 0.2^\circ$ grid cell, PSR imagery was mosaicked using a scene prioritization strategy guided by the corresponding Unusable Data Masks (UDMs), ensuring the selection of the clearest and most complete pixels across overlapping scenes.

4.3 Global Building Polygon Generation

This section outlines the global-scale building polygon generation process, as illustrated in the upper middle part of Fig. 2. We begin by describing the curation of the training data, followed by a detailed description of the building footprint extraction model, the regularization model, and the polygonization method used to generate the building polygons. Finally, we present the filtering process employed to remove false-positive polygons.

4.3.1 Data Curation

To construct a dataset for training the building polygon generation model, we sampled PSR image patches from 107 predefined regions of interest (RoIs) distributed globally. These RoIs included 92 densely built urban areas and 15 non-urban regions, such as ocean, forest, desert, and mountain environments, which served as negative samples.

PSR imagery overlapping with the RoIs was divided into patches of size 256×256 to ensure compatibility with the deep learning models, resulting in a total of 142 722 training samples. For each sampled PSR image patch, we obtained building polygons from OSM for regions outside of China. For Chinese cities, we used annotations provided by the dataset published by Cao and Huang (2021). All the acquired building polygon annotations were rasterized to the spatial resolution of 3 m to match the resolution of the PSR images.

4.3.2 Building Map Extraction

A building map extraction network was designed to map the input PSR images to binary building masks. This followed an encoder–decoder architecture based on UPerNet (Xiao et al., 2018), with ConvNeXt-Tiny (Liu et al., 2022) serving as the backbone. The network is supervised by cross-entropy loss. Following the practice in Zhao et al. (2017), an auxiliary loss is calculated using the third-layer features extracted from ConvNeXt backbone to enhance the feature representation. Before being fed into the network, image patches were up-scaled by a factor of 4 and cropped to a size of 512×512 . The network was trained for 160 000 iterations with a batch size of 8, using the AdamW optimizer (Loshchilov and Hutter, 2019). The global inference was built-up on the pipeline for generating GlobalBuildingMap, as described in Zhu et al. (2024).

4.3.3 Building Map Regularization

The building map extraction network described in Sect. 4.3.2 generates binary building masks from input satellite images by applying a threshold of 0.5 to the predicted probability maps. However, the resulting masks are often noisy, with adjacent buildings merging due to the limited resolution of the PSR images, which are typically insufficient to capture fine building details. To address this, we trained a building map regularization network to refine the generated masks to improve their accuracy and delineation.

During training, we generated two binary masks, $\hat{\mathbf{M}}$ and $\mathbf{M}$, using building polygon annotations, treating them as the network's input and output, respectively. $\hat{\mathbf{M}}$ was created by rasterizing the building polygon annotations to match the resolution of the corresponding images, with random noise added to the vertices and edges of the polygons. In contrast, $\mathbf{M}$ was generated similarly, but without any added noise. By




Figure 2. Workflow of the proposed pipeline. The basemap in the “Building Model” image is derived from © OpenStreetMap contributors 2025. Distributed under the Open Data Commons Open Database License (ODbL) v1.0. © Google and the Google logo are trademarks of Google LLC. © Microsoft and the Microsoft logo are trademarks of Microsoft Corporation.

learning to predict $\mathbf{M}$ from $\hat{\mathbf{M}}$, the network was able to de-noise the generated binary masks, producing regularized and more accurate building mask representations.

Since both networks followed an encoder-decoder architecture, the architecture and training details of the building map regularization network were identical to those of the building map extraction network described in Sect. 4.3.2.

4.3.4 Building Polygonization and Simplification

After the regularized binary building maps were generated, the subsequent building polygonization and simplification algorithms mapped them to polygonal representations. Specifically, the binary masks were first converted into dense vector representations using a contour tracing-like algorithm provided by the GDAL library (The GDAL Development Team, 2020). Then, a polygon simplification algorithm, as described in Zhang et al. (2025a), was applied to simplify the polygonal representation.

4.3.5 False Positives Filtering

Since the available training data do not sufficiently cover the diverse terrains and land cover types worldwide, and the quality of satellite data can vary by region, the generated building polygons may include false positives, particularly in areas with cloud cover, snow, water bodies, forests, and similar features. To filter out these false positives, we implemented a filtering strategy based on the global land cover map product World Cover (Zanaga et al., 2021). We began by dilating the built-up mask of the World Cover map with a 250 m window size, and then filter out any generated polygons that fell outside the masked area.

4.4 Global Building Height Estimation

The proposed pipeline acquires 3D building information exclusively from optical satellite imagery, ensuring broad accessibility and ease of data acquisition. We utilized deep learning methods that are particularly well-suited for this task, as they can effectively leverage large-scale data while offering flexibility and scalability – key advantages for global geospatial applications. This section details the height information retrieval process, illustrated in the lower middle part of Fig. 2. We first describe the curation of the training data and then elaborate on the monocular height estimation process.

4.4.1 Data Curation

The dataset used for training the monocular height estimation model consisted of samples from 168 city-scale RoIs worldwide (see Fig. 1). These regions were primarily located in North America, Europe and Oceania, where high-quality 3D observations are more readily available. The 3D data originated from LiDAR point clouds released by governments, which were processed into height maps in raster format – specifically, normalized digital surface models (nDSMs) that match the resolution of PSR imagery. The nDSMs were cropped into patches of size 256×256 , yielding a total of 231 656 samples.

4.4.2 Monocular Height Estimation

We trained an HTC-DC Net (Chen et al., 2023) to predict heights from single images, following a classification-regression paradigm. The model consisted of an EfficientNet-B5 (Tan and Le, 2020) backbone for feature extraction, a classification module, and a hybrid

regression process. The classification module employed a vision transformer encoder (Dosovitskiy et al., 2020) to capture relationships between local and global features, dynamically determining bin edges and their corresponding probabilities. The classification output is treated as a distribution, which was refined through the hybrid regression process to obtain the final height prediction. The network was trained for 150 epochs with a batch size of 8, using the AdamW optimizer (Loshchilov and Hutter, 2019).

4.4.3 Uncertainty Quantification

To assess the uncertainty of the predicted height values, we employed the test-time augmentation (TTA) technique during the inference process. Inference was conducted on mosaicked PSR imagery on a $0.2^\circ \times 0.2^\circ$ scale using a sliding window approach. The window slides with a stride of 128 pixels, allowing up to 4 predictions per pixel. The variance across these 4 predictions serves as the uncertainty measure for height prediction.

4.5 Global LoD1 Building Model Generation

The generation of the global LOD1 building model relies on the availability of high-quality global building polygons and the corresponding building heights. Given the limited availability of accurate and detailed global building height datasets, GBA.Height represents the most suitable option for this purpose. However, to leverage existing large-scale building footprint datasets – such as the vector sources listed in Table 2, which were derived from higher-quality input data – we introduce a quality-guided fusion strategy. This approach, detailed in Sect. 4.5.1, enables the integration of multiple building polygon sources. The final global LoD1 building model was then constructed based on this fused footprints and our generated GBA.Height, as described in Sect. 4.5.2.

4.5.1 Quality-guided Building Polygon Fusion

As shown in Table 2, different instance-level building footprint products exhibit varying degrees of completeness across different regions. Even in areas where multiple building footprint products are available, their quality differs significantly. To leverage the strengths of all available datasets, we performed a quality-guided polygon fusion process to produce a final global building footprint dataset – GBA.Polygon.

The fusion processing was carried out within each administration boundary as defined by the GADM dataset (GADM, 2025). The process begins with selecting a base layer, i.e., the primary source, which should represent the highest-quality building footprint at a given location. Zhu et al. (2024) compared various building footprint datasets across different continents and concluded that OpenStreetMap (OSM) is the most suitable base layer for all continents except

South America and Africa. In these two regions, Google's Open Building dataset outperforms others due to its well-delineated and superior preservation of details.

However, while these datasets may provide the best quality at a large scale, concerns about their completeness remain. To address this limitation, we introduced a secondary source to complement the primary dataset with region-specific coverage.

To address this, we evaluated other sources against the base layer to identify the most suitable secondary source, with the ranking determined using a combined metric of recall and area gain. Recall measures the building areas in the primary source that are covered by a given secondary source, while area gain quantifies the additional building area introduced by the secondary source. The source with the highest combined metric was selected as the secondary source, ensuring that it aligned well with the primary source while maximizing the inclusion of new building footprints.

Finally, the selected primary and secondary sources were merged at the instance level. All buildings from the primary source were retained, and any additional buildings from the secondary source were incorporated into the final dataset. This approach ensures that the final polygon set integrates the strengths of the two most reliable data sources, resulting in a more complete and accurate global building footprint product.

4.5.2 LoD1 Building Model Generation

To generate the final LoD1 building models, the predicted height maps were integrated with the fused building footprints to assign a height value to each building instance. This height value was determined by selecting the maximum height within each building's footprint. The variance at the location where the maximum value was taken served as the uncertainty measure for the corresponding building instance.

5 Results

5.1 Dataset Overview

Our resulting GlobalBuildingAtlas dataset consists of global building polygons (GBA.Polygon), building height maps (GBA.Height) and LoD1 models (GBA.LoD1). Figure 3 gives an overview on the GlobalBuildingAtlas dataset. It illustrates the global distribution of building volumes (in the middle) as well as statistics on the building counts, areas and volumes and associated errors by continent (top), and showcases the resulting building height map and LoD1 model from selected areas in North America (middle-left) and Asia (middle-right), respectively. The bottom subfigures offer a closer look at the building volume distribution in 10 representative cities across the globe, using the same colormap as the global map for consistency. The total building count across the globe is 2.75 billion, corresponding to a total




Figure 3. Overview of our resulting GlobalBuildingAtlas dataset, consisting of global building polygons (GBA.Polygon), building height map (GBA.Height) and LoD1 models (GBA.LoD1). The top section presents the dataset statistics by continent, including the number of buildings, total building area and volume, as well as height RMSE and volume RMSE across the test cities. The middle section displays the global distribution of building volume, computed over $480\text{ m} \times 480\text{ m}$ grid cells. The bottom section offers a closer look at the building volume distribution in 10 representative cities, using the same colormap as the global map for consistency. LoD1 model visualization base images © Google Maps.

building area of $506.64\text{ billion m}^2$ (square meters) and a total building volume of 2.85 trillion m^3 (cubic meters). This suggests that the actual global building count is lower than the UN's estimation of 4 billion. As anticipated, building volumes are concentrated in metropolitan regions, with notable clusters in East Asia, Europe, and North America.

The global distribution of building-related metrics by continent reveals significant geographic variations in terms of quantity and scale.

Asia leads in nearly all metrics, with an estimated 1.22 billion buildings, the highest among all the continents. This reflects both the continent's large population and the extensive urban development across countries such as China, India, and those in Southeast Asia. Africa follows with approximately 540 million buildings, surpassing Europe (403 mil-

lion) and North America (295 million) in terms of building count. South America contributes around 264 million buildings, while Oceania has the lowest count at 14 million buildings.

Despite Africa's relatively high number of buildings, its total building area is significantly lower, amounting to only 38 billion m², in contrast to Asia's 218 billion m², Europe's 138 billion m², and North America's 79 billion m². South America and Oceania contribute 28 and 5 billion m², respectively. This disparity indicates that African buildings tend to be smaller in size, possibly reflecting the dominance of low-rise or informal structures – an expected pattern in less developed or peri-urban areas.

A similar pattern can be observed with the building volume statistics. Asia again ranks highest, with a cumulative building volume of approximately 1.272 trillion m³, followed by Europe (763 billion m³) and North America (538 billion m³). Africa, despite its large number of structures, contributes only 117 billion m³, further underscoring the prevalence of small-scale or single-story buildings. South America and Oceania contribute 123 billion m³ and 26 billion m³, respectively.

The building height and volume estimation accuracy of the GlobalBuildingAtlas dataset across continents is mainly influenced by several factors: (1) variations in building morphology, (2) the representation of regional characteristics in the training data, and (3) the comprehensiveness and quality of the validation data. The global average height RMSE of GBA.LoD1 stands at 5.5 m. The lowest error is observed in Oceania (1.5 m), followed by Europe (4.1 m), North America (5.3 m), and Asia (5.9 m). South America, however, exhibits a significantly higher height RMSE of 8.9 m, suggesting a greater uncertainty or variation in building morphology, possibly due to a mixture of informal and vertical developments in urban regions. The global average volume RMSE of GBA.LoD1 is approximately 152.6 m³ per 100 m². South America again shows the highest error at 586.8 m³ per 100 m², which may be attributed to compounding uncertainties in both the building footprint and height estimation. Asia follows with 247.7 m³ per 100 m², then Europe (150.9 m³ per 100 m²), North America (135.5 m³ per 100 m²), and Oceania with the lowest error at 46.8 m³ per 100 m². Furthermore, GBA.Height enhances building volume estimation in South America and Asia, achieving RMSE values of 580.0 and 38.0 m³ per 100 m², respectively.

In summary, the statistics reflect significant continental differences in the scale of the built environment and estimation accuracy. While Asia dominates in terms of the building count, area, and volume, Africa's high building count but low aggregate volume and area underscores the heavy presence of small-scale and potentially informal constructions. Conversely, Europe and Oceania demonstrate relatively lower estimation errors, suggesting a higher consistency in terms of the building structures and their better representation in the

training data. South America, despite its moderate building count, emerges as a region with notably high estimation errors, indicating potential challenges in urban morphology or limitations in the training data.

These findings provide valuable insights for global-scale urban modeling, infrastructure planning, and geospatial analytics.

5.2 Comparison with Existing Products

The released GlobalBuildingAtlas dataset provides comprehensive information for both raster-based and instance-level building representation. For comparison, we compared it against three raster-based products, namely GHS-BUILT-H (Pesaresi and Politis, 2023b), WSF-3D (Esch et al., 2022), and Google 2.5D (Sirko et al., 2023), as well as two instance-level products, namely Microsoft (Microsoft, 2024) and 3D-GloBFP (Che et al., 2024).

5.2.1 Validation Dataset

We collected a test dataset (see Fig. 1) comprising LoD1 building data from government sources across 28 cities in Asia, Europe, North America, South America, and Oceania to evaluate the performance of the various building height products. Note, owing to the lack of reference data in Africa, it was not possible to conduct a quantitative evaluation in that region.

5.2.2 Metrics

The evaluation covered assessments of both 2D building footprint and 3D building information. The quality of 2D building footprints was quantified using the following metrics.

- Intersection over union (IoU). IoU was used to evaluate both the raster-based and instance-level building products. For evaluating raster-based building map, we rasterized the reference polygons to 3 m resolution, resizing the building map to the same resolution and then calculated the IoU of the building area. For the instance-level building products, the IoU was calculated at the polygon level.
- Average precision of building polygons at an IoU threshold of 0.5 (AP_{50}^{poly}). To evaluate the quality of building polygons with vectorized building data, we adopted the standard average precision (AP) metric, which is widely used in object detection and instance segmentation (Everingham et al., 2010). Specifically, we computed the AP for building polygons using an IoU threshold of 0.5, denoted as AP_{50}^{poly} , following the classic evaluation protocol.

- Average recall of building polygons at an IoU threshold of 0.5 (AR_{50}^{poly}). In addition to evaluating the average precision (AR_{50}^{poly}), we also computed the average recall (AR) to assess how comprehensive the dataset's ability to comprehensively capture building instances is. Specifically, AR_{50}^{poly} was determined to measure the recall performance of detected building polygons at an IoU threshold of 0.5, following the standard evaluation protocol.
- *N*-ratio of the detected building count. To evaluate the accuracy of the different instance-level building products, we employed the *N*-ratio metric, defined as the ratio between the number of detected buildings and the corresponding ground truth count. This metric provides a straightforward measure of how well the predicted building counts align with reference data. An *N*-ratio closer to 1 indicates better performance.

The assessment of the 3D building information included estimated building volume and height accuracy, corresponding to the instance-level and raster-based building representations, respectively.

- Root mean square error ($RMSE_{BV}$) and mean absolute error (MAE_{BV}) of the estimated building volume. The ground truth building models and the instance-level building datasets were rasterized. The RMSE and MAE were computed on a per-pixel basis with a $1\text{ m} \times 1\text{ m}$ grid resolution.
- RMSE ($RMSE_{BH}$) and MAE (MAE_{BH}) of the estimated building heights. Building instances in the ground truth and in the various datasets were matched according to the maximum overlapping area. The RMSE and MAE of the building height were then computed based on these matched pairs.
- Building height completeness (Comp.). This metric represents the proportion of ground truth buildings (in numbers) for which valid predictions are available. A prediction is considered valid if it exceeds a minimum building height threshold, which was set to 1 m in this study.

5.2.3 Quantitative Results

We present the quantitative comparison results of various building height products in Table 3. To provide a quick visual summary, the key metrics and characteristics are also illustrated in Fig. 4.

For metrics related to 2D building footprints – including AP_{50}^{poly} , AR_{50}^{poly} , *N*-ratio, and IoU – our building polygon dataset GBA.Polygon consistently achieves the best performance across most regions, with the exception of South American cities. This demonstrates the effectiveness of our

proposed quality-guided polygon fusion strategy in leveraging existing multi-source vectorized building footprint datasets.

In terms of building height-related metrics, both our raster-format height product GBA.Height and vector-format LoD1 model GBA.LoD1 achieve the highest completeness scores across all categories. This suggests that the generated height map provides the most comprehensive and meaningful coverage of building heights over the reference polygons. For building volume estimation, as measured by $RMSE_{BV}$ and MAE_{BV} , our GBA.LoD1 outperforms all other compared methods across every continent except South America, confirming its suitability for city-scale volume estimation. Lastly, in terms of building height accuracy – measured by $RMSE_{BH}$ and MAE_{BH} – our GBA.LoD1 delivers the best results in North America and Oceania, and remains competitive or only slightly behind other methods in the remaining regions, confirming its consistent high-quality across the globe.

Due to the lack of reference data in Africa, the results could not be directly evaluated. However, the satisfactory performance observed in South America – where both training and test data are similarly limited and building morphology and regional characteristics differ significantly from those of more developed regions – suggests that the results in Africa would likely be acceptable.

5.2.4 Qualitative Results

To provide an intuitive understanding of the performance of various large-scale building height products, we visualize the results from the different methods on selected test cities across multiple continents: Portland, USA for North America, Medellín, Colombia for South America, Bordeaux, France for Europe, Launceston, Australia for Oceania, and Wakayama, Japan for Asia. As shown in Fig. 5, the raster-based products, such as GHS-BUILT-H (Pesaresi et al., 2024; Pesaresi and Politis, 2023a) and WSF 3D (Esch et al., 2020, 2022), provide coarse, city-scale building height estimates but fail to capture the detailed building geometries – a critical limitation for many downstream applications. Google 2.5D (Sirko et al., 2023) produces accurate building heights in South American cities, like Medellín, but its geographic coverage is limited primarily to the Global South. Microsoft's LoD1 model offers detailed polygon-level height estimations with good spatial resolution. However, it lacks coverage in Asia and South America and tends to underestimate heights in European and Asian cities, such as Bordeaux and Wakayama. 3D-GloBFP (Che et al., 2024) provides global but incomplete coverage, yet tends to overestimate building heights, especially in Medellín and Bordeaux. In contrast, our building height map GBA.Height and LoD1 model GBA.LoD1 achieve accurate estimations in cities across North America, Europe, and Oceania, including Portland, Bordeaux, and Launceston. While some under-

Table 3. Per-continent evaluation results for various global building height products using the test dataset. “AS”, “EU”, “NA”, “OC”, and “SA” represent Asia, Europe, North America, Oceania, and South America, respectively. The products are grouped based on whether they provide vectorized building footprints. Metrics of the building volume (BV) are reported in the units of m³ per 100 m²; metrics of the building heights (BH) are reported in m. Best metrics in each group are shown in bold.

Continent	Dataset	Vector	AP ₅₀ ^{poly}	AR ₅₀ ^{poly}	<i>N</i> -ratio	IoU	RMSE _{BV}	MAE _{BV}	RMSE _{BH}	MAE _{BH}	Comp.
AS	GHS-BUILT-H	<i>x</i>	–	–	–	13.3	322.6	113.5	5.7	3.8	0.95
	WSF 3D	<i>x</i>	–	–	–	16.4	268.9	52.6	6.1	4.4	0.98
	Google 2.5D	<i>x</i>	–	–	–	–	–	–	–	–	–
	GBA.Height (ours)	<i>x</i>	–	–	–	–	247.7	41.4	5.8	4.4	0.99
	Microsoft	✓	6.5	17.8	0.53	23.4	–	–	–	–	–
	3D-GloBFP	✓	8.7	24.0	0.76	32.0	321.7	49.7	6.0	3.2	0.67
	GBA.LoD1 (ours)	✓	27.9	46.8	1.04	44.7	235.0	41.3	5.9	4.6	0.76
EU	GHS-BUILT-H	<i>x</i>	–	–	–	10.2	206.4	60.4	7.1	5.0	0.74
	WSF 3D	<i>x</i>	–	–	–	12.8	198.6	40.0	5.6	3.7	0.96
	Google 2.5D	<i>x</i>	–	–	–	–	–	–	–	–	–
	GBA.Height (ours)	<i>x</i>	–	–	–	–	150.9	24.3	5.6	3.5	0.99
	Microsoft	✓	13.9	28.1	0.65	52.6	168.4	27.4	3.9	2.8	0.55
	3D-GloBFP	✓	5.1	12.0	0.31	25.1	257.1	47.2	6.2	4.6	0.52
	GBA.LoD1 (ours)	✓	46.0	69.3	1.19	67.2	139.4	22.7	4.1	3.0	0.83
NA	GHS-BUILT-H	<i>x</i>	–	–	–	10.3	185.1	44.3	8.8	4.0	0.93
	WSF 3D	<i>x</i>	–	–	–	12.7	173.4	23.8	5.9	2.9	0.99
	Google 2.5D	<i>x</i>	–	–	–	–	–	–	–	–	–
	GBA.Height (ours)	<i>x</i>	–	–	–	–	144.1	16.6	6.4	2.4	0.99
	Microsoft	✓	37.8	58.0	0.94	58.5	149.3	16.1	6.4	2.6	0.81
	3D-GloBFP	✓	41.3	55.5	0.77	63.2	153.9	18.3	8.2	6.7	0.90
	GBA.LoD1 (ours)	✓	54.1	72.4	1.09	66.8	135.5	15.1	5.3	2.1	0.96
SA	GHS-BUILT-H	<i>x</i>	–	–	–	25.1	560.2	265.4	11.3	6.5	0.98
	WSF 3D	<i>x</i>	–	–	–	29.4	566.6	198.7	13.8	7.4	0.99
	Google 2.5D	<i>x</i>	–	–	–	46.6	519.7	188.4	8.6	5.1	0.97
	GBA.Height (ours)	<i>x</i>	–	–	–	–	580.0	171.5	10.1	5.7	1.00
	Microsoft	✓	0.0	0.6	0.10	42.2	–	–	–	–	–
	3D-GloBFP	✓	0.1	1.2	0.14	43.4	1708.5	697.3	29.2	24.5	0.89
	GBA.LoD1 (ours)	✓	0.5	4.8	0.69	37.6	586.8	180.6	8.9	5.0	0.87
OC	GHS-BUILT-H	<i>x</i>	–	–	–	3.9	61.4	11.6	1.9	1.4	0.85
	WSF 3D	<i>x</i>	–	–	–	8.6	42.7	3.9	2.1	1.4	0.97
	Google 2.5D	<i>x</i>	–	–	–	–	–	–	–	–	–
	GBA.Height (ours)	<i>x</i>	–	–	–	–	38.0	3.6	1.5	1.0	0.99
	Microsoft	✓	18.9	46.1	1.15	37.5	46.7	3.9	1.5	1.0	0.57
	3D-GloBFP	✓	19.2	45.9	1.14	38.7	52.2	4.8	1.8	1.1	0.84
	GBA.LoD1 (ours)	✓	36.7	56.0	1.16	39.8	46.8	4.2	1.5	1.0	0.87

estimation is still observed in Medellín and Wakayama, our models strike a significantly better overall balance between resolution, completeness, and accuracy compared to existing solutions.

5.3 Validation of the UN's Building Count Estimation

The UN has estimated that there are approximately 4 billion buildings worldwide (United Nations Human Settlement Program: UN Habitat, 2019), a figure that exceeds the number of building instances in any existing datasets, including our GBA. However, by using the *N*-ratio evaluation

metric shown in Table 3, which reflects the completeness of GBA.Polygon, we provide an alternative estimate of the global building count based on our product to validate the UN's estimate.

Initially, *N*-ratios were computed continentally and globally against reference data, and ranged from 0.69 in South America to 1.19 in Europe, with a global average of 1.03. Using these ratios, we estimated the total number of buildings per continent by scaling the known building counts accordingly. For Africa, where no *N*-ratio is available, we computed




Figure 4. Overview of the continental comparison results across large-scale building height products. In the resolution dimension, the products were first scored based on whether they provide vectorized building footprints. Raster-based products were then further scored according to their spatial resolution.

three estimates using the minimum, maximum, and global average N -ratios from other continents.

Aggregate projections across continents yield a global estimate of 2.71 billion buildings, while the estimated bounds range from 2.64 billion to 2.97 billion buildings. This analysis suggests that the UN's estimate may exceed actual global building counts.

5.4 Strengths and Limitations

Based on the preceding analysis of our dataset, its key strengths and limitations can be summarized as follows.

Strengths

- GBA.Polygon represents the most comprehensive dataset of global building polygons to date, comprising approximately 2.75 billion buildings. This achievement could be primarily attributed to two key factors: (1) the development of a multi-source building polygon fusion strategy that effectively integrates all existing large-scale polygon datasets; and (2) the implementation of a global building polygon generation pipeline based on PlanetScope satellite imagery, which addresses existing data gaps by producing previously unavailable building footprints, albeit with limited quality in some regions constrained by a lower spatial resolution of the input satellite data.
- GBA.Height represents the first high-resolution building height map available at the global scale. In contrast to existing global building height datasets (Pesaresi and Politis, 2023a; Ma et al., 2024; Esch et al., 2022), which offer coarse spatial resolutions ranging from 90 to 150 m, our approach can achieve a resolution of 3 m – representing at least a 30-fold improvement. This substantial enhancement enables more detailed and accurate representation of urban structures, facilitating a broader range of downstream geospatial applications. Furthermore, height estimation in GBA.Height is derived exclusively from a single and affordable data source, namely PlanetScope satellite imagery, instead of very high-resolution aerial data. Owing to the reduced data dependency data cost and the high temporal revisit frequency of PlanetScope, our approach supports rapid and scalable updates of global building height maps.
- GBA.LoD1 constitutes the most comprehensive global LoD1 building model to date, encompassing approximately 2.68 billion buildings and filling a gap of over 1 billion previously unrepresented structures worldwide. Compared to existing LoD1 datasets, GBA.LoD1 offers globally consistent and high-accuracy building height estimates, with RMSEs ranging from 1.5 to 8.9 m across different geographic regions.




Figure 5. Visual comparison of existing building height products on the test cities Portland (North America), Medellín (South America), Bordeaux (Europe), Launceston (Oceania) and Wakayama (Asia). Satellite images © Google Maps.

Limitations

- Limited availability of height data in Africa for training and validation. Due to the scarcity of accurate reference data in Africa, our building height estimation model was neither trained nor validated on samples from this region. Consequently, the model may be subject to domain shift, potentially limiting its generalizability and accuracy when applied to urban areas in African contexts. The limitations in available training data could be potentially addressed by applying weakly supervised learning, i.e. training dedicated models by considering all existing building height data as weak labels. However, community efforts will still be needed to curate the validation dataset.
- Underestimation of the heights in high-rise building areas. As shown in the visual results in Fig. 5, our monocular height estimation model tends to underestimate building heights in some South American and Asian cities. This may be attributed to the model being trained to minimize pixel-level height errors rather than building instance-level errors, leading to more conservative predictions in unfamiliar high-rise urban scenarios.

6 Applications and Enrichment

In this section, we demonstrate two applications of the global LoD1 building models. The first examines the correlation between building volume and population, while the second evaluates the utility of building volume information for computing SDG indicators.

6.1 Correlation with the Population

6.1.1 Fine-scale Continental Analysis: Europe as a Case Study

We investigated the relationship between population distribution and building volume across a $1\text{ km} \times 1\text{ km}$ grid over the 27 member states of the European Union (EU). The population data used in this analysis were sourced from the EU's official population grid dataset in year 2021 (Eurostat, 2021). As can be seen in Fig. 6, we performed a logarithmic regression between the population count and building volume for each grid cell, both at the EU-wide level and for individual member states. To assess the significance of the relationship, we computed both the Pearson correlation coefficient (r) and the Spearman rank correlation coefficient (ρ), which measure linear and monotonic association, respectively.

In general, population exhibits a strong positive correlation with building volume: areas with higher populations tend to be associated with buildings of greater volume. While the slope of the regression line reflects the average building volume per capita, the correlation coefficient provides

insights into the equity of the distribution. Higher correlation values suggest a more uniform allocation of building volume relative to population across spatial units, indicating a closer alignment between the built infrastructure and population density.

As illustrated in Fig. 7, a considerable disparity emerges when comparing the building volume per capita and correlation coefficients across the 27 EU countries. More than half of the member states fall below the EU-wide average in terms of building volume per capita. Finland, which ranks highest in this metric, possesses 6 times the per-capita building volume of Greece, which ranks lowest. The disparities are somewhat less pronounced in the correlation coefficients, with most countries exhibiting relatively strong correlations, with only seven countries falling significantly below the EU-wide correlation benchmark.

Notably, some of the countries with the highest building volume per capita, such as Finland and Estonia, exhibit a particularly imbalanced allocation of building volume with respect to their population, as evidenced by their lower correlation coefficients. This suggests that while they have abundant building stock, its distribution is less aligned with their population density.

6.1.2 Coarse-scale Global Analysis

Similar conclusions can be drawn from a coarse, country-level analysis conducted at the global scale. We performed a logarithmic regression between the total population of each country or territory (The World Bank, 2019b) in year 2019 and the corresponding total building volume derived from our dataset. As can be seen in Fig. 8, the regression again indicates a strong positive correlation between the population and building volume.

However, the analysis also reveals a much greater imbalance in building volume per capita across countries. Countries with the largest building volumes per capita are predominantly located in Europe, while those with the smallest values are primarily in Africa. As an example, Finland continues to exhibit the highest building volume per capita, whereas in Niger, Africa, individuals have access to only 0.36 % of that value. This figure is also 27 times lower than the global average, highlighting significant disparities in built infrastructure worldwide.

6.2 Potential SDG Indicator

To monitor progress toward SDG 11 – *Make cities and human settlements inclusive, safe, resilient, and sustainable* – several indicators have been defined (United Nations, 2025). Among these, Indicator 11.3.1 – *Ratio of land consumption rate to population growth rate* – plays a crucial role (United Nations Human Settlements Programme (UN-Habitat), 2021). The indicator is computed by deriving the built-up area change and population change. Due to the ab-




Figure 6. Regression analysis between population and building volume were conducted for the entire EU as well as for each of the 27 EU member states individually. The first graph shows the regression analysis of the EU as a whole, followed by the 27 member countries sorted in descending order based on the harmonic metric that averages the Pearson (r) and Spearman (ρ) correlation coefficients (Year 2021).

sence of time series data, we were unable to directly compute the indicator to substantiate the rationale for employing volume-based indicators, though generating time series data using our developed pipeline is relatively straightforward. As an alternative approach, we examined the use of volume-based indicators by considering the secondary indicator defined in the indicator metadata (United Nations Human Settlements Programme (UN-Habitat), 2021) – *built-up area per capita*. In this context, we substituted built-up area with built-up volume to assess whether the revised indicator could provide a more accurate reflection of the urban development status.

Sustainable development is influenced not only by factors related to the built-up environment but also by a wide range of other social, economic, and environmental variables. Nevertheless, correlation analysis between built-up environment indicators and development-related measures can provide insights into the relative effectiveness of different indicators. In this study, we used gross domestic product (GDP) per capita in year 2019 (The World Bank, 2019a) as a proxy for development status and compared its correlation with both the building area per capita and building volume per capita. The per-capita building metrics were computed using population data for the year 2019 (The World Bank, 2019b).

As Fig. 9 shows, GDP per capita exhibits a strong positive correlation with both building area per capita and building volume per capita, with correlation coefficients of 0.76

and 0.85, respectively. The significantly higher correlation coefficient between building volume per capita and GDP per capita suggests building volume per capita might be a more suitable measure.

To further validate this observation, we assessed the alignment between GDP per capita and building-based indicators – namely, building volume per capita and building area per capita – by randomly selecting pairs of countries and comparing their respective rankings. For each pair, we examined whether the ranking based on the indicator agreed with the ranking based on GDP per capita. As shown in Table 4, among 20 301 country or territory pairs sampled from a total of 202 countries and territories, using building volume per capita instead of building area per capita resulted in 788 additional ranking agreements. This substitution increased the agreement rate from 79.6 % (area-based) to 83.5 % (volume-based), demonstrating the superior discriminative power of building volume per capita in capturing differences in economic development.

In conclusion, volume-based indicators demonstrate greater potential than area-based metrics for assessing progress toward the SDGs related to sustainable urban development.




Figure 7. Building volume per capita and the harmonized correlation coefficients for the 27 EU member states and the EU as a whole. The harmonized correlation coefficients were computed as an average of the Pearson (r) and Spearman (ρ) correlation coefficients (Year 2021).

Table 4. Comparison of the ranking accuracy for GDP per capita estimation using building volume per capita and building area per capita (Year 2019). Best metrics are shown in bold.

Indicator	Agreement Between Indicators	Only One Indicator Correct	Total Correct Rankings	Ranking Accuracy [%]
Volume per Capita	15 724	1229	16 953	83.5
Area per Capita	15 724	441	16 165	79.6

7 Data availability

The GBA dataset described in this manuscript can be accessed on mediaTUM under <https://doi.org/10.14459/2025mp1782307> (Zhu et al., 2025b). The dataset comprises building footprints from OSM and Microsoft, both provided under the Open Database License (ODbL). This portion of the dataset can be accessed at <https://doi.org/10.57967/hf/6771> (Zhu et al., 2025a).

8 Code availability

All code is available on GitHub under an MIT license with the Commons Clause: <https://github.com/zhu-xlab/GlobalBuildingAtlas> (last access: 31 May 2025; <https://doi.org/10.5281/zenodo.17543406>, Zhang et al., 2025b). The Commons Clause restricts the use of this software for commercial purposes. The repository includes the full development code for the GBA product, as well as scripts for reproducing the figures in this manuscript. A portal for interactive preview of the dataset is available at the same URL.




Figure 8. Regression analysis of population and building volume across all countries and territories globally with Pearson (r) and Spearman (ρ) correlation coefficients. Top 10 and bottom 10 countries or territories are displayed by building volume per capita (Year 2019).




Figure 9. Correlations between GDP per capita and two built environment indicators: building volume per capita and building area per capita. Pearson (r) and Spearman (ρ) correlation coefficients are reported (Year 2019).

9 Conclusions

In this study, we introduce GlobalBuildingAtlas, an open global dataset comprising building polygons, heights, and LoD1 3D building models. Our motivation stems from wanting to fill the significant gaps in existing datasets, particularly regarding the lack of comprehensive 3D information at

an individual building level, which limits applications in both large-scale urban environments and detailed assessments.

We developed a streamlined pipeline that utilizes only optical satellite imagery for LoD1 3D building modeling. This approach facilitates rapid updates for time-series monitoring and enables a wide range of applications. The pipeline consists of two main components: building polygon generation and building height estimation. The building polygon gen-

eration process addresses gaps in current building footprint datasets by supplementing missing building instances. Building height estimation produces pixel-wise height maps at a resolution of 3 m, extending applications beyond the built environment alone. The resulting LoD1 3D building model constitutes the most comprehensive and accurate dataset to date, encompassing 2.75 billion buildings, with 97.7 % of them having height estimates. Our height estimation achieves high accuracy, with RMSEs ranging from 1.5 to 8.9 m. Additionally, the resulting building volumes demonstrates significant advantages over existing products, with RMSEs ranging from 46.8 to 586.8 m³ per 100 m².

We further explored the relationship between population and building volume, revealing a strong correlation. Our analysis highlighted the substantial disparities in building volume per capita on both continental and global scales. Moreover, we propose using building volume-based indicators to assess progress toward the SDGs, as these indicators better reflected development status compared to building area-based metrics in our analysis.

Appendix A: Additional Information of the 3D Training Set

The 3D training set introduced in Sect. 3.3 comprises a total of 168 cities (Table A1). The majority of samples originate from Europe (109 cities), followed by North America (39) and Oceania (17). In contrast, Asia and South America contribute only two and one cities, respectively, while no publicly available LiDAR data were identified for Africa that could be used to supervise monocular height estimation models. In total, the training set contains 187 239 paired samples of PSR imagery and corresponding nDSM patches, each with a spatial size of 256 pixel × 256 pixel at 3 m resolution, corresponding to a ground coverage of approximately 110 438 km².

Appendix B: Building Volume Distribution by Country or Territory

Figure B1 illustrates the distribution of building volumes by country or territory. The results reveal that China holds the largest global share, accounting for 24.8 % of total building volume, followed by the USA at 15.4 %. These findings are consistent with the prior study (Che et al., 2024). Notably, due to improved building data completeness, higher proportions of building volumes for Russia, India, and Brazil are uncovered. Considering the high population count and land area in these countries and territories, the finding seems to be reasonable. Conversely, building volumes for other countries and territories are estimated to constitute a smaller percentage compared to the earlier research, which aligns more closely with the expected patterns.

Table A1. List of Countries with Numbers of City-scale Regions in the 3D Training Set

<i>Europe</i>	109
Austria	5
Belgium	8
Denmark	5
Estonia	2
France	3
Germany	25
Ireland	1
Latvia	2
Luxembourg	1
Netherlands	7
Norway	8
Poland	1
Slovakia	1
Slovenia	2
Spain	18
Switzerland	3
United Kingdom	17
<i>North America</i>	39
Canada	19
Mexico	1
United States	19
<i>Asia</i>	2
China	1
Philippines	1
<i>Oceania</i>	17
Australia	11
New Zealand	6
<i>South America</i>	1
Brazil	1
<i>Africa</i>	0
TOTAL	168




Figure B1. Building volume distribution by country or territory. Top: spatial distribution of building volumes; Bottom: Top 20 countries or territories with most building volumes. The colorbars are consistent.

Appendix C: Contribution of Different Data Sources in Quality-guided Building Polygon Fusion

Figure C1 illustrates the contributions of various building data sources to GBA.Polygon. In terms of building count, Open Buildings (Google Research, 2023) is the largest contributor in Asia, Africa, and South America, with a total of 1.62 billion buildings. OpenStreetMap (OSM) provides 0.49 billion buildings, showing more complete coverage in Europe and North America, but relatively sparse representation in Africa and South America. Microsoft's building footprints account for 0.43 billion buildings across multiple continents, and serve as a key supplement to OSM in Europe. CLSM (Shi et al., 2024) and our own generated building polygons primarily address the gaps in Asia, contributing 0.07 and 0.14 billion buildings, respectively.

When comparing building area contributions, the proportion from Open Buildings decreases across continents, suggesting that its average building size is smaller. In contrast, our generated building polygons contribute more area relative to their building count, likely due to their derivation from lower-resolution satellite imagery, which can lead to the merging of adjacent buildings into single polygon instances. Despite this limitation, our building polygons significantly complement existing datasets, enabling a more complete and globally consistent coverage of building footprints.

Figure C2 presents example polygons generated through the quality-guided polygon fusion process described in Sect. 4.5.1. By integrating additional building footprint datasets, the fusion process substantially enhances the overall completeness and coverage of the final building polygons.




Figure C1. Contribution of various building polygon sources by continent in the quality-guided polygon fusion process. The number of buildings and total building area are attributed to each source. “PSR-derived (ours)” refers to the building footprint polygons generated from PSR as described in Sect. 4.3.




Figure C2. Examples of fused building polygons derived from multiple footprint sources. “PSR-derived (ours)” refers to the building footprint polygons generated from PSR as described in Sect. 4.3. OSM data come from © OpenStreetMap and are distributed under the Open Data Commons Open Database License (ODbL) v1.0. Satellite images © Google Maps.

Author contributions. Conceptualization: XXZ; methodology: XXZ, SC, FZ, YS; software: SC, FZ, YS; results validation: XXZ, SC, FZ; analysis: XXZ, SC, FZ; data collection: FZ, SC, XXZ; writing, and original draft preparation: XXZ, SC, FZ; paper revision: XXZ, SC, FZ, YW; visualization: SC, FZ, XXZ, YW; funding acquisition: XXZ; project administration: XXZ; resources: XXZ.

Competing interests. The contact author has declared that none of the authors has any competing interests.

Disclaimer. Publisher's note: Copernicus Publications remains neutral with regard to jurisdictional claims made in the text, published maps, institutional affiliations, or any other geographical representation in this paper. While Copernicus Publications makes every effort to include appropriate place names, the final responsibility lies with the authors. Views expressed in the text are those of the authors and do not necessarily reflect the views of the publisher.

Acknowledgements. Source of the Planet data used in the publication: Planet Labs Inc. This work is jointly supported by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No. (ERC-2016-StG-714087), Acronym: *So2Sat*), by the German Federal Ministry for Economic Affairs and Climate Action in the framework of the "national center of excellence ML4Earth" (grant number: 50EE2201C), by the German Research Foundation (DFG GZ: ZH 498/18-1; Project number: 519016653), by the Excellence Strategy of the Federal Government and the Länder through the TUM Innovation Network EarthCare, by the European Union (project number: 101091374) and by Munich Center for Machine Learning. The authors gratefully acknowledge the computational and data resources provided by the Leibniz Supercomputing Centre (<https://www.lrz.de>, last access: 1 November 2025).

Financial support. This research has been supported by the European Research Council, HORIZON EUROPE European Research Council (grant no. ERC-2016- StG-714087), the Bundesministerium für Wirtschaft und Klimaschutz (grant no. 50EE2201C), the Deutsche Forschungsgemeinschaft (grant no. ZH 498/18-1), the Technische Universität München, Graduate School, Technische Universität München (grant no. EarthCare), the European Commission, HORIZON EUROPE Framework Programme (grant no. 101091374), and the Munich Center for Machine Learning.

Review statement. This paper was edited by Yuyu Zhou and reviewed by two anonymous referees.

References

- Cao, Y. and Huang, X.: A deep learning method for building height estimation using high-resolution multi-view imagery over urban areas: A case study of 42 Chinese cities, *Remote Sensing of Environment*, 264, 112590, <https://doi.org/10.1016/j.rse.2021.112590>, 2021.
- Che, Y., Li, X., Liu, X., Wang, Y., Liao, W., Zheng, X., Zhang, X., Xu, X., Shi, Q., Zhu, J., Zhang, H., 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.
- Chen, S., Shi, Y., Xiong, Z., and Zhu, X. X.: HTC-DC Net: Monocular Height Estimation From Single Remote Sensing Images, *IEEE Transactions on Geoscience and Remote Sensing*, 61, 1–18, <https://doi.org/10.1109/TGRS.2023.3321255>, 2023.
- City of Helsinki: Helsinki 3D City Model, [3D model], https://hri.fi/data/en_GB/dataset/helsingin-3d-kaupunkimalli (last access: 1 November 2025), 2017.
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., and Houslsby, N.: An image is worth 16 × 16 words: Transformers for image recognition at scale, *arXiv [preprint]*, <https://doi.org/10.48550/arXiv.2010.11929>, 2020.
- Esch, T., Thiel, M., Schenk, A., Roth, A., Müller, A., and Dech, S.: Delineation of Urban Footprints From TerraSAR-X Data by Analyzing Speckle Characteristics and Intensity Information, *IEEE Transactions on Geoscience and Remote Sensing*, 48, 905–916, <https://doi.org/10.1109/TGRS.2009.2037144>, 2010.
- Esch, T., Schenk, A., Ullmann, T., Thiel, M., Roth, A., and Dech, S.: Characterization of Land Cover Types in TerraSAR-X Images by Combined Analysis of Speckle Statistics and Intensity Information, *IEEE Transactions on Geoscience and Remote Sensing*, 49, 1911–1925, <https://doi.org/10.1109/TGRS.2010.2091644>, 2011.
- Esch, T., Taubenböck, H., Roth, A., Heldens, W., Felbier, A., Thiel, M., Schmidt, M., Müller, A., and Dech, S.: TanDEM-X Mission – New Perspectives for the Inventory and Monitoring of Global Settlement Patterns, *Journal of Applied Remote Sensing*, 6, 061702, <https://doi.org/10.1117/1.JRS.6.061702>, 2012.
- Esch, T., Marconcini, M., Felbier, A., Roth, A., Heldens, W., Huber, M., Schwinger, M., Taubenböck, H., Müller, A., and Dech, S.: Urban Footprint Processor – Fully Automated Processing Chain Generating Settlement Masks from Global Data of the TanDEM-X Mission, *IEEE Geoscience and Remote Sensing Letters*, 10, 1617–1621, <https://doi.org/10.1109/LGRS.2013.2272953>, 2013.
- Esch, T., Zeidler, J., Palacios-Lopez, D., Marconcini, M., Roth, A., Mönks, M., Leutner, B., Brzoska, E., Metz-Marconcini, A., Bachofer, F., Loekken, S., and Dech, S.: Towards a Large-Scale 3D Modeling of the Built Environment—Joint Analysis of TanDEM-X, Sentinel-2 and Open Street Map Data, *Remote Sensing*, 12, 2391, <https://doi.org/10.3390/rs12152391>, 2020.
- Esch, T., Brzoska, E., Dech, S., Leutner, B., Palacios-Lopez, D., Metz-Marconcini, A., Marconcini, M., Roth, A., and Zeidler, J.: World Settlement Footprint 3D – A first three-dimensional survey of the global building stock, *Remote Sensing of Environment*, 270, 112877, <https://doi.org/10.1016/j.rse.2021.112877>, 2022.

- Eurostat: Population distribution – GEOSTAT, <https://ec.europa.eu/eurostat/web/gisco/geodata/population-distribution/population-grids> (last access: 24 November 2025), 2021.
- Everingham, M., Van Gool, L., Williams, C. K., Winn, J., and Zisserman, A.: The pascal visual object classes (voc) challenge, *International journal of computer vision*, *International Journal of Computer Vision*, 88, 303–338, 2010.
- GADM: Global Administrative Areas (GADM), <https://gadm.org/data.html> (last access: 23 April 2025), 2025.
- Google Research: Google Open Buildings Dataset, available under the Creative Commons Attribution 4.0 International License, <https://sites.research.google/open-buildings/> (last access: 30 April 2025), 2023.
- Liu, Z., Mao, H., Wu, C.-Y., Feichtenhofer, C., Darrell, T., and Xie, S.: A convnet for the 2020s, in: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 11976–11986, <https://doi.org/10.48550/arXiv.2201.03545>, 2022.
- Loshchilov, I. and Hutter, F.: Decoupled weight decay regularization, *arXiv [preprint]*, <https://doi.org/10.48550/arXiv.1711.05101>, 2019.
- Ma, X., Zheng, G., Xu, C., Moskal, L. M., Gong, P., Guo, Q., Huang, H., Li, X., Liang, X., Pang, Y., Wang, C., Xie, H., Yu, B., Zhao, B., and Zhou, Y.: A global product of 150-m urban building height based on spaceborne lidar, *Scientific Data*, 11, 1387, <https://doi.org/10.1038/s41597-024-04237-5>, 2024.
- Microsoft: Microsoft Global ML Building Footprints, available under the Open Data Commons Open Database License (ODbL), GitHub [data set], <https://github.com/microsoft/GlobalMLBuildingFootprints> (last access: 30 April 2025), 2024.
- Ministry of Land, Infrastructure, Transport and Tourism (MLIT) Japan: PLATEAU 3D City Model, [3D city model], <https://www.geospatial.jp/ckan/dataset/plateau> (last access: 1 November 2025), 2024.
- OpenStreetMap contributors: OpenStreetMap, available under the Open Database License, <https://www.openstreetmap.org> (last access: 30 April 2025), 2025.
- Pennsylvania State University: Philadelphia 3D City Model, [3D city model], <https://www.pasda.psu.edu/uci/DataSummary.aspx?dataset=7146> (last access: 1 November 2025), 2015.
- Pesaresi, M. and Politis, P.: GHS-BUILT-H R2023A – GHS building height, derived from AW3D30, SRTM30, and Sentinel-2 composite (2018) [data set], <https://doi.org/10.2905/85005901-3A49-48DD-9D19-6261354F56FE>, 2023a.
- Pesaresi, M. and Politis, P.: GHS-BUILT-V R2023A – GHS built-up volume grids derived from joint assessment of Sentinel-2, Landsat, and global DEM data, multitemporal (1975–2030), European Commission, Joint Research Centre (JRC) [data set], <https://doi.org/10.2905/AB2F107A-03CD-47A3-85E5-139D8EC63283>, 2023b.
- 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., Maffenini, 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, <https://publications.jrc.ec.europa.eu/repository/handle/JRC136539> (last access: 24 November 2025), 2024.
- Peters, R., Dukai, B., Vitalis, S., van Liempt, J., and Stoter, J.: Automated 3D reconstruction of LoD2 and LoD1 models for all 10 million buildings of the Netherlands, *Photogrammetric Engineering & Remote Sensing*, 88, 3, <https://doi.org/10.14358/PERS.21-00032R2>, 2022.
- Shi, Q., Zhu, J., Liu, Z., Guo, H., Gao, S., Liu, M., Liu, Z., and Liu, X.: The Last Puzzle of Global Building Footprints – Mapping 280 Million Buildings in East Asia Based on VHR Images, *Journal of Remote Sensing*, 4, 0138, <https://doi.org/10.34133/remotesensing.0138>, 2024.
- Sirko, W., Brempong, E. A., Marcos, J. T. C., Annkah, A., Korme, A., Hassen, M. A., Sapkota, K., Shekel, T., Diack, A., Nevo, S., Hickey, J., and Quinn, J. A.: High-Resolution Building and Road Detection from Sentinel-2, *arXiv [preprint]*, <https://doi.org/10.48550/arXiv.2310.11622>, 2023.
- Tan, M. and Le, Q. V.: EfficientNet: Rethinking model scaling for convolutional neural networks, *arXiv [preprint]*, <https://doi.org/10.48550/arXiv.1905.11946>, 2020.
- The GDAL Development Team: GDAL – Geospatial Data Abstraction Library, <https://gdal.org> (last access: 2 April 2023), 2020.
- The World Bank: GDP per capita (current US\$), The World Bank [data set], <https://data.worldbank.org/indicator/NY.GDP.PCAP> (last access: 16 April 2025), 2019a.
- The World Bank: Population, total – World Bank Data, The World Bank [data set], <https://data.worldbank.org/indicator/SP.POP.TOTL> (last access: 16 April 2025), 2019b.
- United Nations: The Sustainable Development Goals Report 2024, <https://unstats.un.org/sdgs/report/2024/The-Sustainable-Development-Goals-Report-2024.pdf> (last access: 8 April 2025), 2024.
- United Nations: Global Indicator Framework for the Sustainable Development Goals and Targets of the 2030 Agenda for Sustainable Development, <https://unstats.un.org/sdgs/indicators/Global-Indicator-Framework-after-2025-review-English.pdf> (last access: 8 April 2025), 2025.
- United Nations Human Settlement Program: UN Habitat: Urban Indicators Database, <https://data.unhabitat.org> (last access: 1 November 2025), 2019.
- United Nations Human Settlements Programme (UN-Habitat): Metadata for Indicator 11.3.1: Ratio of Land Consumption Rate to Population Growth Rate, Tech. rep., United Nations Statistics Division, <https://unstats.un.org/sdgs/metadata/files/Metadata-11-03-01.pdf> (last access: 16 April 2025), 2021.
- Xiao, T., Liu, Y., Zhou, B., Jiang, Y., and Sun, J.: Unified perceptual parsing for scene understanding, *arXiv [preprint]*, <https://doi.org/10.48550/arXiv.1807.10221>, 2018.
- Zanaga, D., Van De Kerchove, R., De Keersmaecker, W., Souverijns, N., Brockmann, C., Quast, R., Wevers, J., Grosu, A., Paccini, A., Vergnaud, S., Cartus, O., Santoro, M., Fritz, S., Georgieva, I., Lesiv, M., Carter, S., Herold, M., Li, L., Tsendbazar, N., Ramoino, F., and Arino, O.: ESA WorldCover 10m 2020 v100, Zenodo [data set], <https://doi.org/10.5281/zenodo.5571936>, 2021.
- Zhang, F., Shi, Y., and Zhu, X. X.: Global Collinearity-aware Polygonizer for Polygonal Building Mapping in Remote Sensing, *arXiv [preprint]*, <https://doi.org/10.48550/arXiv.2505.01385>, 2025a.

- Zhang, F., Chen, S., and Wang, Y.: zhu-xlab/GlobalBuildingAtlas: v1.0.0 (v1.0.0), Zenodo [code], <https://doi.org/10.5281/zenodo.17543406>, 2025b.
- Zhao, H., Shi, J., Qi, X., Wang, X., and Jia, J.: Pyramid scene parsing network, in: Proceedings of the IEEE conference on computer vision and pattern recognition, 2881–2890, <https://doi.org/10.1109/CVPR.2017.660>, 2017.
- Zhu, X. X., Li, Q., Shi, Y., Wang, Y., Stewart, A., and Prexl, J.: GlobalBuildingMap – Unveiling the Mystery of Global Buildings, arXiv [preprint], <https://doi.org/10.48550/arXiv.2404.13911>, 2024.
- Zhu, X. X., Chen, S., Zhang, F., Shi, Y., and Wang, Y.: zhu-xlab/GBA.ODbLPolygon, Hugging Face [data set], <https://doi.org/10.57967/hf/6771>, 2025a.
- Zhu, X. X., Chen, S., Zhang, F., Shi, Y., and Wang, Y.: Global-BuildingAtlas: An Open Global and Complete Dataset of Building Polygons, Heights and LoD1 3D Models, TUM [data set], <https://doi.org/10.14459/2025mp1782307>, 2025b.