

# UnderOneFacade: Worldwide Facade Semantic Segmentation Benchmark Dataset

Yi Wang<sup>1\*</sup>, Fan Wang<sup>1\*</sup>, Wanru Yang<sup>2\*</sup>, Prabin Gyawali<sup>1</sup>, Ziyang Xu<sup>1</sup>, Anna Klimkowska<sup>3</sup>, Yixiong Jing<sup>2</sup>, Filip Biljecki<sup>4</sup>, Christoph Holst<sup>1</sup>, Benjamin Busam<sup>1</sup>, Brian Sheil<sup>2</sup>, and Olaf Wysocki<sup>2</sup>


<sup>1</sup> Technical University of Munich, Munich, Germany

<sup>2</sup> CV4DT, University of Cambridge, Cambridge, UK

<sup>3</sup> University of Nottingham, Nottingham, UK

<sup>4</sup> National University of Singapore, Singapore

\* equal contribution



**Fig. 1:** UnderOneFacade: We present a centimeter-accurate and cross-continental facade point clouds, with fine-grained semantic segmentation of architectural elements, and hierarchical facade taxonomy enabling multi-level evaluation (LoFG2–LoFG3), totaling 2.7B annotated points.

**Abstract.** Globally consistent semantic digital twins require centimeter-accurate and geographically transferable 3D facade segmentation. However, progress in facade parsing is limited by the lack of large-scale, standardized benchmarks for evaluating cross-domain and cross-country generalization. Existing datasets are geographically narrow, sensor-specific, semantically inconsistent, or insufficiently precise. We introduce **UnderOneFacade**, the largest cross-continental 3D facade benchmark to

date, comprising centimeter-accurate point clouds with hierarchical, harmonized, and architecturally grounded semantic labels totaling 2.7 billion annotated points. Through a systematic evaluation of representative point-, graph- and transformer-based architectures, we show that current methods struggle to recognize fine-grained architectural elements and degrade significantly across geographic regions, with the best models achieving only up to 33 IoU on the fine-grained LoFG3 benchmark. By combining geometric precision with standardized semantics at unprecedented scale, UnderOneFacade establishes a rigorous benchmark for developing robust and transferable 3D segmentation models for and beyond facade understanding. The dataset, evaluation scripts, and pre-trained models are available here: [https://jiangyuanwangyi.github.io/UnderOneFacade\\_official/](https://jiangyuanwangyi.github.io/UnderOneFacade_official/)

**Keywords:** 3D Point Cloud Segmentation · Urban Scene Understanding · Cross-Domain Generalization

## 1 Introduction

Facade semantic segmentation has been studied since the early days of computer vision [41], owing to the structured yet intricate patterns present in building facades and their importance for urban scene understanding, large-scale mapping, and semantic digital twins. Accurate per-element labeling of facade components enables downstream applications such as architectural analysis, urban planning, energy modeling, and change detection [1, 10, 19, 28, 50]. Despite steady progress in 2D and 3D scene segmentation, facade understanding remains challenging due to its fine-grained structure, high intra-class variability, and strong dependence on precise geometric detail.

A key difficulty arises from the long-tail distribution of facade elements commonly observed in real-world data (e.g., ScanNet [5], A2D2 [15]). While dominant classes such as walls are well represented, semantically critical architectural elements appear far less frequently, including doors, windows, and balconies. This imbalance mirrors challenges observed across indoor and outdoor semantic segmentation benchmarks [5, 15, 18, 49, 52] and significantly limits the effectiveness of data-driven approaches. As a result, models often produce biased predictions favoring dominant classes while failing to recognize smaller and structurally intricate facade elements even within the same architectural style [27, 49].

These challenges are further exacerbated by the limited availability and diversity of annotated facade datasets. Existing benchmarks are typically limited to small-scale acquisitions covering individual buildings or homogeneous building blocks within narrow geographic regions (e.g., ZAHA [49], ArCH [27]). Consequently, current models tend to overfit dataset- and country-specific architectural styles and struggle to generalize across cities or countries with different construction practices and design principles.

A seemingly straightforward solution is to leverage widely available street-level imagery to expand facade annotations. However, such data cannot provide

the centimeter-level geometric accuracy of laser scanners required for precise architectural parsing. Facade segmentation is inherently accuracy-critical: even small geometric deviations can corrupt the delineation of fine structural elements and undermine semantic validity.

Beyond geometric fidelity, consistent semantic interpretation is equally important. Architecturally meaningful reasoning requires unified, standard-grounded class definitions rather than ad hoc, dataset-specific taxonomies. Without harmonized semantic structures transferable across cities and countries, benchmarking remains inconsistent and cross-domain evaluation unreliable.

The absence of large-scale 3D facade datasets that jointly ensure centimeter-grade geometric fidelity and standardized, architecturally grounded semantics therefore remains a core limitation of current benchmarks.

To address these limitations, we introduce **UnderOneFacade**, the largest cross-country 3D facade benchmark to date, comprising centimeter-accurate point clouds with unified semantic annotations grounded in architectural standards and totaling **2.7 billion annotated points**. The dataset spans multiple cities and architectural styles across different countries, enabling systematic evaluation of cross-country and cross-continent generalization. Our experiments show that even strong baseline architectures achieve only around **33 mIoU** on the most challenging LoFG3 benchmark, highlighting that fine-grained facade understanding remains far from solved. Beyond facade parsing itself, the proposed benchmark also provides a controlled setting for studying broader problems in 3D scene understanding, including cross-domain generalization across architectural styles, long-tailed recognition of structural elements, and multi-modal geometric-radiometric reasoning. Our contributions are summarized as follows:

- We introduce **UnderOneFacade**, the largest 3D facade segmentation benchmark with centimeter-level geometric accuracy and unified architectural semantics, comprising **2.7B annotated points**.
- We enable the first systematic evaluation of cross-country and cross-continent generalization in facade segmentation.
- We provide a comprehensive benchmark revealing the remaining challenges of fine-grained facade segmentation under realistic long-tailed distributions.

## 2 Related Works

**3D Semantic Segmentation Methods.** Early approaches imposed regular structure on unorganized point clouds (e.g., voxels), enabling convolutional processing but introducing discretization artifacts and limiting scalability [6, 24, 29, 37]. More recent methods process raw points directly. Point-based networks such as PointNet and PointNet++ [35, 36] introduced hierarchical aggregation to capture local geometric context, which is critical for thin and repetitive architectural structures. Subsequent work explored graph-based and transformer-based architectures to model local and global dependencies more effectively

[22, 34, 45, 47, 51, 53]. Some approaches additionally incorporate hierarchical spatial partitions (e.g., octrees) to balance geometric precision and computational efficiency [38, 54]. Despite these advances, accurately segmenting fine-grained facade elements remains challenging due to long-tail class distributions and the limited availability of large-scale facade datasets.

**3D Urban Semantic Segmentation Datasets.** Large-scale urban point cloud datasets have steadily expanded over the past decade, driven by advances in LiDAR sensing and 3D reconstruction. These benchmarks have enabled substantial progress in urban semantic segmentation and scene understanding. Representative datasets include Semantic3D [18], Paris-Lille-3D [39], and Toronto-3D [42]. However, these datasets primarily target coarse semantic categories such as buildings, roads, and vegetation. Fine-grained facade-level parsing is typically absent or treated marginally, limiting progress toward detailed architectural reasoning in 3D.

**Table 1:** Urban point cloud benchmark datasets for semantic segmentation. Columns indicate multi-national scope, multi-sensor acquisition, and whether hierarchical semantic classes are provided.

| Name                        | Year | Sensor      | Real / Synth | #Cls | Facade classes? | #Facade-Labeled points | Multi-National? | Multi-Sensor? | Hierarchical classes? |
|-----------------------------|------|-------------|--------------|------|-----------------|------------------------|-----------------|---------------|-----------------------|
| Oakland 3D [30]             | 2009 | MLS         | real         | 44   | ~               | <0.1B                  | ×               | ×             | ×                     |
| ETH PRS [23]                | 2012 | TLS         | real         | 0    | ×               | ×                      | ×               | ×             | ×                     |
| Sydney Urban Objects [7]    | 2013 | MLS         | real         | 26   | ×               | ×                      | ×               | ×             | ×                     |
| Paris-rue-Madame [40]       | 2014 | MLS         | real         | 27   | ~               | <0.1B                  | ×               | ×             | ~                     |
| iQumulus [44]               | 2015 | MLS         | real         | 8    | ×               | ×                      | ×               | ×             | ×                     |
| TUM-MLS-2016 [55]           | 2016 | MLS         | real         | 9    | ×               | ×                      | ×               | ×             | ×                     |
| Semantic3D.net [18]         | 2017 | TLS         | real         | 9    | ×               | ×                      | ✓               | ×             | ×                     |
| Paris-Lille-3D [39]         | 2018 | MLS         | real         | 50   | ×               | ×                      | ×               | ×             | ×                     |
| SynthCity [17]              | 2019 | MLS         | synthetic    | 9    | ×               | ×                      | ×               | ×             | ×                     |
| A2D2 [15]                   | 2020 | MLS         | real         | 38   | ×               | ×                      | ×               | ×             | ×                     |
| ArCH [28]                   | 2020 | TLS/MLS/UAV | real         | 10   | ✓               | 0.1B                   | ×               | ✓             | ×                     |
| Toronto-3D [42]             | 2020 | MLS         | real         | 8    | ×               | ×                      | ×               | ×             | ×                     |
| Wbu-TLS [9]                 | 2020 | TLS         | real         | 0    | ×               | ×                      | ×               | ×             | ×                     |
| BIMAGE [2]                  | 2021 | MLS         | real         | 0    | ×               | ×                      | ×               | ×             | ×                     |
| KITTI-360 [25]              | 2021 | MLS         | real         | 19   | ×               | ×                      | ×               | ×             | ×                     |
| Paris-CARLA-3D [8]          | 2021 | MLS         | real/synth   | 23   | ×               | ×                      | ✓               | ×             | ×                     |
| TUM-FAÇADE [48]             | 2022 | MLS         | real         | 17   | ✓               | 0.1B                   | ×               | ×             | ×                     |
| HelixNet [26]               | 2022 | MLS         | real         | 9    | ×               | ×                      | ×               | ×             | ×                     |
| SUD [16]                    | 2023 | MLS         | real         | 8    | ×               | ×                      | ×               | ×             | ×                     |
| ZAHA [49]                   | 2025 | MLS         | real         | 15   | ✓               | 0.6B                   | ×               | ×             | ✓                     |
| Tuscany [33]                | 2025 | MLS         | real         | 10   | ✓               | <0.1B                  | ×               | ×             | ×                     |
| TrueCity [31]               | 2025 | MLS         | real/synth   | 12   | ✓               | 0.1B                   | ×               | ×             | ×                     |
| City-Facade [4]             | 2026 | MLS         | real         | 9    | ✓               | 0.2B                   | ×               | ×             | ×                     |
| <b>UnderOneFacade (Our)</b> | 2026 | TLS/MLS     | real         | 15   | ✓               | 2.7 B                  | ✓               | ✓             | ✓                     |

**Facade-Level 3D Datasets.** Motivated by applications such as digital twins and automated building reconstruction [32, 50], several datasets have begun to target facade-specific semantic segmentation with higher geometric fidelity and class granularity. Examples include ArCH [28], TUM-FAÇADE [48], and ZAHA [49]. Nevertheless, existing datasets remain limited in scale, architectural diversity, and geographic coverage. Their restricted size constrains deep model training and systematic evaluation. This bottleneck largely stems from the high cost of centimeter-accurate acquisition and the labor-intensive annotation of small, repetitive, and structurally complex facade elements.

**Semantic Consistency.** Beyond dataset scale, facade benchmarks also exhibit substantial variability in class definitions, often relying on ad hoc or project-specific taxonomies. Such inconsistency impedes cross-dataset evaluation and reproducibility. The ZAHA dataset introduces the Level-of-Facade-Granularity



(LoFG) concept, a hierarchical taxonomy grounded in architectural standards. While this represents an important step toward semantic harmonization, its realization within a single German city limits architectural variability and effectively validates the concept primarily on Central European typologies. These limitations highlight the need for a large-scale, centimeter-accurate, cross-country 3D facade benchmark with standardized and architecturally grounded semantic labels.

UnderOneFacade addresses this gap by providing the largest harmonized facade dataset to date, enabling systematic evaluation across countries and supporting the development of long-tail-robust and transferable 3D facade segmentation methods. In this sense, UnderOneFacade can be viewed as a cross-country extension of ZAHA: we inherit the LoFG taxonomy and the German subset from ZAHA, while contributing a  $4.5\times$  increase in scale, multi-sensor acquisition, and a cross-country and cross-continent evaluation protocol.

### 3 Dataset Design

UnderOneFacade prioritizes highest possible geometric accuracy by deploying laser scanners with up to centimeter local accuracy and globally georeferenced centimeter acquisition, increasing architecture pattern richness by acquiring cross-continent data, multi-modal scanner setups to ensure scanner-related patterns variability but still ensure even cm-grade accuracy, and ensuring high-definition labels by manual and hierarchical annotations.

Three representative regions are selected to capture architectural diversity across continents. The European subsets (UK and Germany) share historical influences but differ in facade materials, ornamentation, and structural composition (Victorian vs Haussmann). In contrast, the Singapore subset represents Southeast Asian urban architecture, combining traditional Malay and Chinese elements with modern high-density construction. This selection introduces substantial stylistic variation and enables systematic evaluation of cross-regional generalization in facade segmentation.

#### 3.1 Laser Scans

**UK.** A static terrestrial laser scanning (TLS) campaign was conducted in The Park residential district in Nottingham, United Kingdom, using the Leica RTC360 [14] scanner, which provides local accuracy at the centimeter-level and rapid full-dome capture. A total of 17 scanning stations were placed along the urban segment to ensure full facade coverage, reduce occlusions, and maintain sufficient overlap between neighboring scans. To ensure consistent global alignment, the survey was fully georeferenced using a network of ground control points (GCPs) measured with GNSS and RTK corrections, providing centimeter-level positional accuracy in the national coordinate system. Individual scans were initially aligned using the scanner’s visual-inertial system and subsequently refined through global registration. The final dataset captures dense facade geometries


with fine architectural detail, including window frames, balconies, cornices, and recessed entrances. Compared to mobile mapping data, the static multi-station acquisition significantly increases point density on vertical structures while reducing occlusions. Using the georeferenced data, individual buildings were extracted using OpenStreetMap (OSM) building footprints [11], expanded with a 5 m buffer to account for footprint inaccuracies. The resulting UK subset contains 997M points across approximately 100 facades. No additional noise filtering was applied in order to preserve raw acquisition fidelity.

**Germany.** The German subset is based on the ZAHA dataset [49], which we integrate into the UnderOneFacade benchmark. ZAHA builds upon the publicly available TUM-MLS-2016 dataset [55], acquired in the Munich downtown area using the Mobile Distributed Situation Awareness (MODISSA) platform. The platform is equipped with two Velodyne HDL-64E LiDAR sensors mounted at the front of the vehicle and two Velodyne VLP-16 sensors at the rear, providing dense multi-view urban coverage. Georeferencing is ensured through an inertial navigation system combined with RTK corrections from the German satellite positioning service (SAPOS) [3, 55]. The dataset captures diverse architectural styles ranging from late 19th century historic buildings to modern urban developments, including residential, commercial, and cultural heritage structures. Following the same OSM-based extraction procedure as used for the UK subset, facade point clouds were clipped per building footprint with a 5 m buffer. The resulting German subset contains 601M points across approximately 66 facades. No additional noise removal was applied.

**Singapore.** For the Singapore subset, data were collected using the Leica BLK360 [12] laser scanner, which captures high-density point clouds with centimeter-level accuracy. Scanning was performed with the scanner mounted to a backpack and from street-level positions along major urban corridors, carefully selected to cover a broad range of architectural styles from modern high-rise buildings to colonial-era structures. Individual scans were registered using geometric feature alignment and subsequently referenced to geographic coordinates using OpenStreetMap (OSM) data. The registered point clouds were merged into a unified coordinate frame, producing a dense and geometrically consistent dataset suitable for facade-level analysis. As with the other subsets, buildings were extracted using OSM footprints with a 5 m buffer. The Singapore subset contains 1.117B points across approximately 200 facades. No additional noise filtering was applied.

### 3.2 Hierarchical Semantic Classes Annotations


**Hierarchical Classes.** High-fidelity 3D facade semantic segmentation requires both hierarchical class definitions and cross-country consistency. To address the lack of benchmarks with sufficient facade variability and standardized semantics, we adopt the Level of Facade Generalization (LoFG) concept [49] to unify semantic labels across diverse architectural contexts. LoFG defines three levels of detail (LoFG1, LoFG2, and LoFG3) allowing precise formulation of segmen-



tation tasks, supporting transfer learning and label propagation, and enabling consistent evaluation across datasets and regions.

In this work, we cross-country-adopted LoFG to provide 15 detailed classes at LoFG3, aggregated into 5 coarser classes at LoFG2, and a single abstract class at LoFG1. The adaptation ensures that the taxonomy remains valid across countries with distinct architectural styles while adhering to international urban modeling and architectural standards such as CityGML, IFC, and the Art and Architecture Thesaurus (AAT) [20,21,28]. LoFG3 includes fine-grained elements such as wall, balcony, molding, decorative elements, stairs, columns, arches, windows, blinds, doors, roof, ground, terrain, interior, and other (mainly noise). LoFG2 merges these into broader functional categories of structural (wall, columns, arch, balcony, stairs), opening (blinds, door, window), decoration (molding, deco), floor (ground surface, terrain), and other elements (roof, other, interior) classes, while LoFG1 abstracts the entire facade.

**Annotation Process.** All annotations were performed manually by trained annotators with expertise in architecture and geospatial modeling, following the architectural standards described earlier. Labeling was carried out in professional large-scale point cloud software, Cyclone 3DR [13], enabling precise interaction with dense point clouds. Annotating a single facade required approximately 15 hours, resulting in roughly 5,500 hours of manual annotation for the full dataset of about 366 facades. All annotations were further harmonized and cross-checked through annotator consensus, with a correction rate of around 10%, occurring mainly in ambiguous categories (e.g., *molding* vs. *decoration*) as defined in our annotation guidelines (Suppl. Sec. 9).



**Fig. 2:** Comparison of per-point class distributions of UnderOneFacade (2.7BN points) and ZAHA (0.6 BN). Cross-country aggregation yields a substantially heavier-tailed and structurally more diverse distribution.

**Data Splits.** Each country subset is divided into training, validation, and test splits following a 70/20/10 ratio based on the total number of points. Splitting is performed per building, ensuring that complete facades belong to only one subset and that the splits are spatially disjoint. It was also ensured to preserve both dominant and underrepresented semantic classes across all subsets. For reference, mean facade heights are 8.2/10.6/12.1 m, while mean point densities are 11,483/28,713/33,199 pts/m<sup>2</sup> before and 575/427/2,446 pts/m<sup>2</sup> after voxelization (Singapore/UK/Germany, respectively).

**UnderOneFacade Benchmark Challenges** UnderOneFacade is designed to evaluate facade segmentation along three benchmark axes: (a) cross-country domain shifts, (b) hierarchical semantic granularity, and (c) realistic long-tailed class distributions. As shown in Fig. 2, aggregating multiple countries results in a markedly heavier-tailed distribution compared to single-country datasets, opening new challenges. By combining centimeter-accurate geometry, harmonized LoFG semantic labels, and architectural diversity across continents, the benchmark enables systematic evaluation of segmentation robustness beyond single-country datasets and can serve as a generic long-tail benchmark beyond facade segmentation. These three principles define the core evaluation axes of UnderOneFacade. In the following section, we detail the experimental protocol designed to systematically assess architectural sensitivity, cross-domain generalization, and long-tail robustness across representative point cloud segmentation architectures.

**Release.** The dataset is released under an open data license (CC BY 4.0, Creative Commons Attribution), hosted on a cloud service (i.e., Google Drive <https://drive.google.com/drive/folders/1Yzz7PmyeK1qe0tkTFCfkbw7IEHXcMJo8?usp=sharing>) to allow unrestricted community use; the project page with additional details: [https://jiangyuanwangyi.github.io/UnderOneFacade\\_official/](https://jiangyuanwangyi.github.io/UnderOneFacade_official/). No ethical approval was required for this study as it did not involve human or animal subjects.


## 4 Experimental Setup

**Baselines.** We select representative baselines from the major families of point cloud networks to reflect different underlying working principles. Namely, point-based: PointNet++ [36], KPConv [43]; graph-based: DGCNN [34]; and transformer-based approaches: OctFormer [45], Point Transformer (abbrv. as PTV1) [53], and Point Transformer v3 (abbrv. as PTV3) [46]. This selection allows us to systematically analyze how performance varies depending on architectural design and feature aggregation strategy. The chosen methods are arguably among the most widely used models in the literature and serve as strong reference points across numerous point cloud benchmarks and methods [15, 42, 49, 52]. Moreover, they are prevalent in related tasks such as facade segmentation and large-scale urban scene understanding [28, 49, 50], making them particularly suitable for comparison in our setting. To isolate the geometric difficulty of facade segmentation, the primary benchmark uses only xyz coordinates as input (radiometric features impact is analyzed separately in Table 5 and accompanying paragraph).

**Hyperparameters.** For a fair and controlled comparison, we evaluate all baseline methods under a unified training protocol inspired by ZAHA [49] (e.g., a fixed 100 training epochs), rather than performing per-method hyperparameter tuning. To analyze the impact of facade architecture and acquisition on various methods under identical budgets, we intentionally used unweighted cross-entropy. And to maintain experimental consistency, we train each method for the same number of epochs and apply identical data preprocessing and uniform point sampling strategies across all networks. We also deploy standard metrics to analyze the performance on the two chosen semantic granularities: Overall Accuracy (OA), Precision (P), Recall (R), F1 score (F1), and Intersection over Union (IoU). More implementation details in Suppl.

## 5 Results and Discussion

**Overall Performance on Hierarchical Semantic.** Overall, the results show that fine-grained facade segmentation remains far from solved (Tab. 2 and Fig. 3). Even the strongest model reaches only about 33 mIoU on the combined dataset (DGCNN: 33.4), highlighting the difficulty of long-tailed facade semantics. Interestingly, the graph-based DGCNN consistently outperforms several more recent transformer-based architectures (e.g., PTV1: 15.1 and PTV3: 27.3 mIoU overall), suggesting that current transformer designs may not yet effectively capture the geometric structure of facade point clouds. While several architectures achieve high overall accuracy on individual datasets, their mIoU values remain limited, indicating that dominant classes such as *wall* or *roof* are segmented reliably whereas rare facade components remain challenging. This observation is further confirmed by the per-class analysis in Table 3. While aggregating classes into the coarser LoFG2 hierarchy slightly increases scores for dominant categories, it does not resolve the underlying segmentation difficulty.



**Fig. 3:** Qualitative facade segmentation results on UnderOneFacade on the LoFG3: We visualize predictions of representative architectures across scenes from different countries. Rows show example facades, while columns correspond to different segmentation models. Despite correct segmentation of dominant structures such as *walls* and *roofs*, models struggle to consistently recognize fine-grained facade elements, including *windows*, *doors*, and decorative components. These examples illustrate the challenges posed by long-tailed semantics and cross-country architectural variability.

**Table 2:** Accuracy across outdoor long-tail datasets and architectures on the most challenging LoFG3. The Overall column reports models trained on the full UnderOneFacade dataset, while others solely within countries.

| Family      | Method     | UK   |      | Singapore |      | Germany |      | Overall |      |
|-------------|------------|------|------|-----------|------|---------|------|---------|------|
|             |            | OA   | mIoU | OA        | mIoU | OA      | mIoU | OA      | mIoU |
| Point       | PointNet++ | 68.5 | 22.1 | 59.6      | 29.8 | 66.4    | 25.6 | 60.9    | 29.7 |
|             | KPConv     | 83.3 | 23.4 | 41.5      | 11.2 | 55.8    | 11.2 | 49.0    | 15.0 |
| Graph       | DGCNN      | 77.8 | 28.3 | 62.2      | 31.3 | 71.1    | 33.4 | 62.7    | 33.4 |
| Transformer | PTv1       | 79.0 | 31.0 | 44.7      | 19.8 | 75.0    | 41.6 | 48.8    | 15.1 |
|             | PTv3       | 79.0 | 26.5 | 58.5      | 29.6 | 56.8    | 15.7 | 58.3    | 27.3 |
|             | OctFormer  | 75.5 | 27.7 | 61.5      | 34.6 | 44.9    | 13.5 | 57.9    | 27.3 |

**Table 3:** Per-class F1 scores for complete UnderOneFacade (%). Results are reported for two semantic granularity levels: LoFG3 (fine-grained) and LoFG2 (coarser). Methods are grouped by architectural family.

| Class                                  | Point-based |        | Graph-based | Transformer-based |             |             |
|----------------------------------------|-------------|--------|-------------|-------------------|-------------|-------------|
|                                        | PointNet++  | KPConv | DGCNN       | PT                | PTv3        | OctFormer   |
| <b>LoFG3 (fine facade semantics)</b>   |             |        |             |                   |             |             |
| wall                                   | 70.5        | 61.2   | 71.7        | 60.5              | 65.2        | <b>71.9</b> |
| window                                 | <b>37.2</b> | 0.0    | 36.9        | 14.7              | 5.1         | 28.1        |
| door                                   | 13.9        | 0.0    | <b>26.7</b> | 0.2               | 16.4        | 10.3        |
| balcony                                | 19.3        | 0.0    | <b>46.7</b> | 0.0               | 6.4         | 43.4        |
| molding                                | 46.8        | 31.0   | <b>48.4</b> | 15.4              | 8.7         | 41.7        |
| deco                                   | <b>10.0</b> | 0.0    | 9.2         | 1.0               | 9.2         | 9.6         |
| column                                 | <b>45.7</b> | 0.0    | 43.7        | 1.1               | 43.3        | 29.5        |
| arch                                   | 42.6        | 8.4    | 41.0        | 2.5               | <b>52.8</b> | 31.5        |
| stairs                                 | 6.5         | 0.0    | <b>9.1</b>  | 1.1               | 0.7         | 1.3         |
| ground surface                         | 54.1        | 48.3   | 54.8        | 20.2              | <b>57.9</b> | 51.6        |
| terrain                                | 70.4        | 44.6   | 73.8        | 64.6              | <b>75.0</b> | 67.7        |
| roof                                   | 68.8        | 46.0   | <b>72.6</b> | 45.8              | 69.3        | 70.3        |
| blinds                                 | 19.5        | 0.0    | <b>28.6</b> | 0.4               | 12.8        | 23.3        |
| interior                               | 61.9        | 35.0   | 69.9        | 53.5              | <b>74.6</b> | 53.6        |
| other                                  | 70.7        | 63.0   | <b>73.1</b> | 55.0              | 69.5        | 63.4        |
| <b>LoFG2 (coarse facade semantics)</b> |             |        |             |                   |             |             |
| floor                                  | <b>93.3</b> | 66.7   | 72.2        | 85.8              | 57.9        | 69.7        |
| decoration                             | <b>45.1</b> | 0.0    | 45.1        | 4.1               | 9.2         | 36.0        |
| structural                             | <b>73.1</b> | 5.7    | 47.7        | 64.2              | 65.2        | 44.3        |
| opening                                | 44.3        | 83.8   | <b>92.1</b> | 20.1              | 16.4        | 87.8        |
| other                                  | 75.4        | 66.6   | <b>75.5</b> | 67.0              | 69.5        | 71.7        |

For example, the *floor* class reaches high performance (e.g., 93.3 F1 for PointNet++), whereas structurally meaningful categories remain challenging. The *decoration* class shows large variability across architectures (45.1 for PointNet++ vs. 4.1 for PT), and even the aggregated *opening* class reveals strong architectural sensitivity, with transformer-based models performing substantially worse (PT: 20.1, PTv3: 16.4) than DGCNN (92.1). Similar patterns are observed in the fine-grained LoFG3 classes such as *window*, *door*, and *molding*. Overall, semantic aggregation alone does not overcome the long-tailed distribution and geometric complexity of facade elements.

**Cross-Country Architectural Variability.** We analyze how architectural differences across countries affect facade segmentation, where Table 4 provides class-level comparisons between the most distinct UK and Singapore subsets. At the fine-grained LoFG3 level, several classes exhibit substantial performance drops in Singapore, particularly *wall* (e.g., PT: 83.7→43.8, −39.9) and *terrain*

**Table 4:** Per-class F1 comparison between UK and Singapore datasets (%). Each cell shows  $UK / SG$  and the difference  $\Delta$  (SG–UK): red  $\rightarrow$  Singapore harder, green  $\rightarrow$  Singapore easier.

| Class                           | Point-based |             | Graph-based<br>DGCNN | Transformer-based |             |             |
|---------------------------------|-------------|-------------|----------------------|-------------------|-------------|-------------|
|                                 | PointNet++  | KPConv      |                      | PT                | PTv3        | OctFormer   |
| LoFG3 (fine facade semantics)   |             |             |                      |                   |             |             |
| wall                            | 72.5 / 54.4 | 81.1 / 49.5 | 79.5 / 58.0          | 83.7 / 43.8       | 81.7 / 50.8 | 81.1 / 64.4 |
|                                 | -18.1       | -31.6       | -21.5                | -39.9             | -30.9       | -16.7       |
| window                          | 13.7 / 3.1  | 0.0 / 0.0   | 27.9 / 2.6           | 23.1 / 5.2        | 10.9 / 0.0  | 19.6 / 9.1  |
|                                 | -10.6       | 0.0         | -25.3                | -17.9             | -10.9       | -10.5       |
| door                            | 5.7 / 0.6   | 0.0 / 0.0   | 10.3 / 0.3           | 21.7 / 0.0        | 0.0 / 0.0   | 21.8 / 17.6 |
|                                 | -5.1        | 0.0         | -10.0                | -21.7             | 0.0         | -4.2        |
| molding                         | 23.6 / 60.2 | 31.7 / 18.5 | 40.7 / 63.3          | 46.3 / 22.9       | 41.8 / 41.5 | 29.0 / 56.3 |
|                                 | +36.6       | -13.2       | +22.6                | -23.4             | -0.3        | +27.3       |
| column                          | 2.8 / 59.6  | 0.0 / 4.5   | 14.1 / 62.9          | 18.9 / 30.7       | 0.0 / 62.1  | 0.0 / 65.4  |
|                                 | +56.8       | +4.5        | +48.8                | +11.8             | +62.1       | +65.4       |
| arch                            | 0.0 / 62.0  | 0.0 / 23.8  | 0.0 / 67.3           | 0.0 / 57.6        | 0.0 / 76.2  | 0.0 / 80.0  |
|                                 | +62.0       | +23.8       | +67.3                | +57.6             | +76.2       | +80.0       |
| terrain                         | 95.4 / 62.4 | 97.5 / 55.7 | 98.2 / 64.0          | 97.6 / 57.9       | 97.2 / 57.5 | 96.3 / 70.4 |
|                                 | -33.0       | -41.8       | -34.2                | -39.7             | -39.7       | -25.9       |
| roof                            | 77.3 / 75.2 | 87.8 / 7.6  | 84.7 / 80.4          | 87.0 / 64.4       | 88.7 / 83.7 | 87.2 / 80.5 |
|                                 | -2.1        | -80.2       | -4.3                 | -22.6             | -5.0        | -6.7        |
| interior                        | 23.2 / 40.0 | 0.0 / 0.8   | 49.7 / 36.6          | 55.4 / 62.1       | 47.7 / 69.6 | 35.0 / 50.9 |
|                                 | +16.8       | +0.8        | -13.1                | +6.7              | +21.9       | +15.9       |
| LoFG2 (coarse facade semantics) |             |             |                      |                   |             |             |
| floor                           | 71.9 / 94.8 | 83.0 / 93.0 | 79.9 / 79.9          | 82.1 / 88.5       | 82.0 / 96.1 | 79.8 / 97.0 |
|                                 | +22.9       | +10.0       | 0.0                  | +6.4              | +14.1       | +17.2       |
| decoration                      | 24.8 / 60.5 | 0.0 / 30.2  | 28.0 / 51.1          | 36.5 / 36.6       | 33.8 / 40.1 | 20.3 / 49.4 |
|                                 | +35.7       | +30.2       | +23.1                | +0.1              | +6.3        | +29.1       |
| structural                      | 21.5 / 61.7 | 15.3 / 47.7 | 40.6 / 66.3          | 42.1 / 52.4       | 36.9 / 60.3 | 34.0 / 65.5 |
|                                 | +40.2       | +32.4       | +25.7                | +10.3             | +23.4       | +31.5       |
| opening                         | 94.9 / 23.8 | 97.7 / 0.0  | 97.6 / 20.2          | 97.7 / 6.8        | 97.5 / 17.2 | 95.9 / 35.6 |
|                                 | -71.1       | -97.7       | -77.4                | -90.9             | -80.3       | -60.3       |
| other                           | 82.9 / 72.8 | 89.6 / 66.6 | 88.1 / 67.5          | 88.9 / 55.5       | 90.0 / 78.5 | 86.5 / 79.8 |
|                                 | -10.1       | -23.0       | -20.6                | -33.4             | -11.5       | -6.7        |

(97.6 $\rightarrow$ 57.9, -39.7), indicating a clear architectural domain shift between European and Southeast Asian scenes. Similarly, detailed facade elements such as *window* degrade strongly (DGCNN: 27.9 $\rightarrow$ 2.6). Conversely, elements such as *column* (2.8 $\rightarrow$ 59.6) and *arch* (0.0 $\rightarrow$ 62.0) appear predominantly in Singapore, resulting in large positive gains. At the coarser LoFG2 level, most classes remain comparatively stable across countries, suggesting improved robustness under hierarchical semantic aggregation. However, the *opening* class exhibits severe degradation across architectures (e.g., PT: 97.7 $\rightarrow$ 6.8, -90.9; DGCNN: 97.6 $\rightarrow$ 20.2), highlighting persistent difficulties in intricate facade openings of Southeast Asian buildings compared to the repetitive patterns of UK Victorian architecture.

**Impact of Radiometric Features.** We analyze the impact of radiometric information by comparing models trained with geometric features only (xyz) against models using additional radiometric attributes, including projected RGB (rgb) values and laser intensity (i). Table 5 shows that incorporating radiometric features can improve performance, increasing mIoU by +7.3 for PointNet++ and +8.1 for PT. The improvements are most pronounced for facade elements with weak geometric signatures, such as *window*, *door*, and *molding*, where gains exceed +30 F1 for windows and +24 F1 for doors, likely due to consistent color



**Table 5:** Impact of radiometric features (RGBI) on the UK LoFG3 test set.  $\Delta$  indicates the performance change when adding RGBI features ( $+rgbi - xyz$ ).

| Metric/Class   | PointNet++ |       |          | PT   |       |          | DGCNN |       |          |
|----------------|------------|-------|----------|------|-------|----------|-------|-------|----------|
|                | xyz        | +rgbi | $\Delta$ | xyz  | +rgbi | $\Delta$ | xyz   | +rgbi | $\Delta$ |
| OA             | 68.5       | 73.9  | +5.4     | 79.0 | 83.9  | +4.9     | 77.8  | 76.5  | -1.3     |
| $\mu P$        | 28.9       | 39.0  | +10.1    | 40.0 | 49.7  | +9.7     | 40.5  | 43.1  | +2.6     |
| $\mu R$        | 27.8       | 37.7  | +9.9     | 38.1 | 47.2  | +9.1     | 33.4  | 37.6  | +4.2     |
| $\mu F1$       | 27.9       | 38.0  | +10.1    | 38.5 | 47.9  | +9.4     | 34.9  | 38.4  | +3.5     |
| $\mu IoU$      | 22.1       | 29.4  | +7.3     | 31.0 | 39.1  | +8.1     | 28.3  | 30.2  | +1.9     |
| wall           | 72.5       | 76.2  | +3.7     | 83.7 | 86.1  | +2.4     | 79.5  | 78.9  | -0.6     |
| window         | 13.7       | 48.8  | +35.1    | 23.1 | 54.5  | +31.4    | 27.9  | 38.2  | +10.3    |
| door           | 5.7        | 30.2  | +24.5    | 21.7 | 44.8  | +23.1    | 10.3  | 43.0  | +32.7    |
| balcony        | 0.0        | 0.0   | 0.0      | 0.0  | 0.0   | 0.0      | 0.0   | 0.0   | 0.0      |
| molding        | 23.6       | 41.5  | +17.9    | 46.3 | 60.9  | +14.6    | 40.7  | 47.1  | +6.4     |
| deco           | 12.1       | 18.1  | +6.0     | 31.6 | 43.6  | +12.0    | 20.5  | 34.2  | +13.7    |
| column         | 2.8        | 15.4  | +12.6    | 18.9 | 33.9  | +15.0    | 14.1  | 24.1  | +10.0    |
| arch           | 0.0        | 0.0   | 0.0      | 0.0  | 0.0   | 0.0      | 0.0   | 0.0   | 0.0      |
| stairs         | 0.0        | 0.0   | 0.0      | 0.0  | 0.0   | 0.0      | 0.0   | 0.0   | 0.0      |
| ground surface | 0.0        | 0.0   | 0.0      | 0.0  | 0.0   | 0.0      | 0.0   | 0.0   | 0.0      |
| terrain        | 95.4       | 89.1  | -6.3     | 97.6 | 97.8  | +0.2     | 98.2  | 96.6  | -1.6     |
| roof           | 77.3       | 82.5  | +5.2     | 87.0 | 92.7  | +5.7     | 84.7  | 85.5  | +0.8     |
| blinds         | 8.3        | 17.0  | +8.7     | 19.3 | 36.9  | +17.6    | 6.0   | 0.3   | -5.7     |
| interior       | 23.2       | 64.4  | +41.2    | 55.4 | 74.8  | +19.4    | 49.7  | 39.8  | -9.9     |
| other          | 83.4       | 87.1  | +3.7     | 93.1 | 93.2  | +0.1     | 92.3  | 87.8  | -4.5     |

patterns across samples (e.g., white or brown doors and window frames). In contrast, large planar structures such as *terrain* remain largely unaffected. However, radiometric cues are not universally beneficial. For several classes, adding RGBI features leads to performance degradation, suggesting that models may learn dataset-specific radiometric shortcuts rather than robust geometric representations (e.g., OA -1.3 for DGCNN). Overall, these results highlight that radiometric information can complement geometric features for small facade elements, but may also introduce appearance biases that limit robustness. Additional results are provided in the Suppl.

**Cross-Country Generalization and Distribution Shift.** We evaluate cross-continental generalization under a zero-shot transfer protocol. Models are trained on one continental subset and evaluated directly on the other without any target-domain adaptation. Two transfer directions are considered: Europe  $\rightarrow$  Asia (training on UK + ZAHA, testing on Singapore) and Asia  $\rightarrow$  Europe. Both training sets contain comparable data volumes of roughly 1.2B points for Asia vs 1.5B points for Europe to ensure that observed differences reflect distributional rather than scale effects. Results are reported at both LoFG2 and LoFG3 levels.

Table 6 reveals a clear asymmetry in transfer performance. At LoFG2, Asia $\rightarrow$ Europe achieves a higher peak OA (65.8%) than Europe $\rightarrow$ Asia (60.9%), while mIoU remains similar (40.5% vs. 40.8%). The asymmetry becomes more pronounced at LoFG3, where Asia $\rightarrow$ Europe consistently outperforms Europe $\rightarrow$ Asia across all methods, reaching 51.9% OA and 23.6 mIoU compared to 42.6% OA and 14.8 mIoU. Class distribution analysis provides a data-grounded explanation for this effect. The European training set is highly concentrated in dominant classes such as *wall* and *other*, which together account for 58.3% of all points, while the Singapore set shows more balanced coverage across

**Table 6:** Cross-continental generalization (zero-shot) on LoFG2 and LoFG3 using xyz coordinates as input features (Europe: UK + ZAHA; Asia: Singapore), where a clear diagonal trend is distinguishable.

| Family            | Method     | Europe $\rightarrow$ Asia<br>(LoFG2) |             | Europe $\rightarrow$ Asia<br>(LoFG3) |             | Asia $\rightarrow$ Europe<br>(LoFG2) |             | Asia $\rightarrow$ Europe<br>(LoFG3) |             |
|-------------------|------------|--------------------------------------|-------------|--------------------------------------|-------------|--------------------------------------|-------------|--------------------------------------|-------------|
|                   |            | OA                                   | mIoU        | OA                                   | mIoU        | OA                                   | mIoU        | OA                                   | mIoU        |
|                   |            |                                      |             |                                      |             |                                      |             |                                      |             |
| Point-based       | PointNet++ | 56.4                                 | 33.9        | 40.5                                 | 14.0        | <b>65.8</b>                          | 38.8        | <b>51.9</b>                          | <b>23.6</b> |
|                   | KPConv     | 48.6                                 | 25.7        | 32.1                                 | 6.6         | 48.9                                 | 28.8        | 33.9                                 | 9.4         |
| Graph-based       | DGCNN      | 56.4                                 | 38.3        | <b>42.6</b>                          | <b>14.8</b> | 63.9                                 | <b>40.5</b> | 44.8                                 | 17.4        |
| Transformer-based | OctFormer  | 59.2                                 | 37.0        | 39.0                                 | 11.5        | 62.3                                 | 40.0        | 50.5                                 | 16.2        |
|                   | PTv1       | 55.1                                 | 34.2        | 33.8                                 | 14.0        | 58.1                                 | 38.1        | 41.9                                 | 20.4        |
|                   | PTv3       | <b>60.9</b>                          | <b>40.8</b> | 41.5                                 | 11.0        | 53.4                                 | 34.2        | 49.0                                 | 15.7        |

facade categories. Consequently, the Asian training data provides stronger supervision for rare classes, which improves transfer performance at the fine-grained LoFG3 level.

These results highlight that balanced semantic coverage in the source domain plays a key role in cross-continental generalization and underline the importance of geographically diverse datasets for robust facade understanding.

## 5.1 Limitations and Future Work

UnderOneFacade combines data from multiple sensing configurations, including vehicle-mounted mobile laser scanners, backpack systems, and static terrestrial scanners. While this increases acquisition diversity and realism, it makes it difficult to isolate the influence of individual sensor modalities without controlled acquisitions of the same environment. We acknowledge that the observed cross-country and cross-continental domain shifts likely arise from a combination of factors, including sensor type, point density, acquisition geometry, and architectural style, rather than architecture alone. UnderOneFacade intentionally reflects these entangled real-world conditions through its per-country (sensor-homogeneous) and overall (sensor-heterogeneous) splits. We provide partial disentanglement via (i) uniform 5 cm training-voxelization to reduce density differences across subsets, and (ii) the radiometric feature analysis in Table 5 and Suppl., which isolates radiometric shifts between the UK and Singapore subsets. Furthermore, the dataset preserves the naturally occurring long-tailed distribution of facade elements. Although this reflects realistic urban scenes, it results in limited training samples for some rare classes. The current benchmark covers three cities, which, despite their architectural diversity, cannot fully represent the global variability of all facade structures. Nevertheless, the proposed hierarchical taxonomy and annotation protocol are designed to be extensible, enabling future integration of additional cities and sensing setups into the benchmark.

## 6 Conclusion

We present **UnderOneFacade**, the largest cross-country 3D facade segmentation benchmark to date, comprising centimeter-accurate point clouds with hierarchical semantic annotations totaling 2.7B labeled points. The dataset combines multiple sensing modalities and architectural styles across countries, producing realistic long-tailed distributions and substantial geometric diversity. Our experiments show that the benchmark is highly challenging for current methods. Notably, transformer-based models do not consistently outperform earlier architectures, with DGCNN achieving the best overall performance, while still the highest IoU reaches only 33%. Radiometric features offer limited benefit and can even degrade performance, indicating reliance on dataset-specific appearance shortcuts. Cross-country evaluation further reveals strong architectural domain shifts, with several models degrading by more than 30 F1 between European and Asian scenes. These results highlight the limited robustness of existing models to architectural variability and long-tailed facade semantics. We believe UnderOneFacade provides an important step toward more robust facade understanding and high definition urban digital twins and generic large-scale semantic segmentation challenges.

## Acknowledgements

UnderOneFacade is the result of a collaborative effort spanning several years, from approximately 2023 to 2026, building upon the earlier development of the ZAHA dataset. We gratefully acknowledge the support of the TUM Global Incentive Fund, which enabled a research stay at the National University of Singapore (NUS) and the acquisition of the Singapore component of the dataset.

We sincerely thank the researchers of the Urban Analytics Lab at NUS for their support during the Singapore data acquisition campaign, and acknowledge the ZAHA team and its annotators for providing the hierarchical foundation that inspired and enabled this work. We are especially grateful to Clara Våth for her meticulous annotation of the Singapore and UK subsets.

Finally, we thank our collaborators at the University of Nottingham for their hospitality and support during the early stages of this project, and the TUM Graduate School which made the visit in Nottingham possible.

This benchmark represents the combined efforts of many researchers, annotators, and collaborators across multiple institutions and countries. We are delighted to see this long-term collaborative effort culminate in the presentation of UnderOneFacade at ECCV 2026.

## References

1. Biljecki, F., Stoter, J., Ledoux, H., Zlatanova, S., Çöltekin, A.: Applications of 3D city models: State of the art review. *ISPRS International Journal of Geo-Information* **4**(4), 2842–2889 (2015)

2. Blaser, S., Meyer, J., Nebiker, S.: Open urban and forest datasets from a high-performance mobile mapping backpack – a contribution for advancing the creation of digital city twins. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences* **XLIII-B1-2021**, 125–131 (2021)
3. Borgmann, B., Schatz, V., Kieritz, H., Scherer-Klößling, C., Hebel, M., Arens, M.: Data processing and recording using a versatile multi-sensor vehicle. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences* **IV-1**, 21–28 (2018)
4. Chen, Y., Li, J., Han, T., Feng, H., Chen, J., Wang, C.: City-facade: A city-level large-scale point cloud building facade dataset for semantic & instance segmentation. *ISPRS Journal of Photogrammetry and Remote Sensing* **232**, 675–688 (2026)
5. Dai, A., Chang, A.X., Savva, M., Halber, M., Funkhouser, T., Nießner, M.: Scannet: Richly-annotated 3d reconstructions of indoor scenes. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 5828–5839 (2017)
6. Dai, A., Nießner, M.: 3dmv: Joint 3d-multi-view prediction for 3d semantic scene segmentation. *Proceedings of the European Conference on Computer Vision (ECCV)* pp. 452–468 (2018)
7. De Deuge, M., Quadros, A., Hung, C., Douillard, B.: Unsupervised feature learning for classification of outdoor 3d scans. *Australasian Conference on Robotics and Automation* **2**(1) (2013)
8. Deschaud, J.E., Duque, D., Richa, J.P., Velasco-Forero, S., Marcotegui, B., Goulette, F.: Paris-CARLA-3D: A real and synthetic outdoor point cloud dataset for challenging tasks in 3D mapping. *Remote Sensing* **13**(22), 4713 (2021)
9. Dong, Z., Liang, F., Yang, B., Xu, Y., Zang, Y., Li, J., Wang, Y., Dai, W., Fan, H., Hyypä, J., et al.: Registration of large-scale terrestrial laser scanner point clouds: A review and benchmark. *ISPRS Journal of Photogrammetry and Remote Sensing* **163**, 327–342 (2020)
10. Gehrung, J.A.: Change detection in point clouds of urban street spaces using fuzzy spatial reasoning. Ph.D. thesis, Technical University of Munich, Munich, Germany (2022)
11. Geofabrik: Openstreetmap data extracts. <https://download.geofabrik.de/> (2020), Accessed: 2020-10-01
12. Geosystems, L.: Leica blk arc. <https://leica-geosystems.com/en-gb/products/laser-scanners/scanners/leica-blk-arc> (2026), Accessed: 2026-01-30
13. Geosystems, L.: Leica cyclone 3dr. <https://leica-geosystems.com/en-gb/products/laser-scanners/software/leica-cyclone/leica-cyclone-3dr> (2026), Accessed: 2026-01-30
14. Geosystems, L.: Leica rtc360 3d laser scanner. <https://leica-geosystems.com/products/laser-scanners/scanners/leica-rtc360> (2026), Accessed: 2026-01-30
15. Geyer, J., Kassahun, Y., Mahmudi, M., Ricou, X., Durgesh, R., Chung, A.S., Hauswald, L., Pham, V.H., Mühlegg, M., Dorn, S., et al.: A2D2: Audi autonomous driving dataset. *arXiv preprint arXiv:2004.06320* (2020)
16. González-Collazo, S.M., Balado, J., Garrido, I., Grandío, J., Rashdi, R., Tsiranidou, E., del Río-Barral, P., Rúa, E., Puente, I., Lorenzo, H.: Santiago urban dataset sud: Combination of handheld and mobile laser scanning point clouds. *Expert Systems with Applications* **238**, 121842 (2024)
17. Griffiths, D., Boehm, J.: SynthCity: A large scale synthetic point cloud. *arXiv preprint arXiv:1907.04758* (2019)
18. Hackel, T., Savinov, N., Ladicky, L., Wegner, J.D., Schindler, K., Pollefeys, M.: Semantic3d.net: A new large-scale point cloud classification benchmark. *arXiv preprint arXiv:1704.03847* (2017)

19. Kolbe, T.H., Donaubauer, A.: Semantic 3D city modeling and BIM. In: Shi, W., Goodchild, M.F., Batty, M., Kwan, M.P., Zhang, A. (eds.) *Urban Informatics*. pp. 609–636. Springer Singapore, Singapore (2021)
20. Kutzner, T., Chaturvedi, K., Kolbe, T.H.: CityGML 3.0: New functions open up new applications. *PFG – Journal of Photogrammetry, Remote Sensing and Geoinformation Science* **88**(1), 1–19 (2020)
21. Laakso, M., Kiviniemi, A.: The IFC standard: A review of history, development, and standardization, information technology. *ITcon* **17**(9) (2012)
22. Lai, X., Liu, J., Jiang, L., Wang, L., Zhao, H., Liu, S., Qi, X., Jia, J.: Stratified transformer for 3d point cloud segmentation. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 8500–8509 (2022)
23. Lande, M.B.: Automatic registration of partially overlapping terrestrial laser scanner point clouds. [https://prs.igp.ethz.ch/research/completed\\_projects/automatic\\_registration\\_of\\_point\\_clouds.html](https://prs.igp.ethz.ch/research/completed_projects/automatic_registration_of_point_clouds.html) (2012), Accessed: 2020-10-30
24. Li, Y., Bu, R., Sun, M., Wu, W., Di, X., Chen, B.: Pointcnn: Convolution on x-transformed points. *Advances in Neural Information Processing System (NeurIPS)* **31** (2018)
25. Liao, Y., Xie, J., Geiger, A.: KITTI-360: A novel dataset and benchmarks for urban scene understanding in 2D and 3D. *arXiv preprint arXiv:2109.13410* (2021)
26. Loiseau, R., Aubry, M., Landrieu, L.: Online segmentation of lidar sequences: Dataset and algorithm. In: Avidan, S., Brostow, G., Cissé, M., Farinella, G.M., Hassner, T. (eds.) *Computer Vision – ECCV 2022*. pp. 301–317. Springer Nature Switzerland, Cham (2022)
27. Matrone, F., Lingua, A., Pierdicca, R., Malinverni, E.S., Paolanti, M., Grilli, E., Remondino, F., Murtiyoso, A., Landes, T.: A benchmark for large-scale heritage point cloud semantic segmentation. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences* **XLIII-B2-2020**, 1419–1426 (2020)
28. Matrone, F., Grilli, E., Martini, M., Paolanti, M., Pierdicca, R., Remondino, F.: Comparing machine and deep learning methods for large 3D heritage semantic segmentation. *ISPRS International Journal of Geo-Information* **9**(9), 535 (2020)
29. Maturana, D., Scherer, S.: Voxnet: A 3d convolutional neural network for real-time object recognition. *International conference on intelligent robots and systems (IROS)* pp. 922–928 (2015)
30. Munoz, D., Bagnell, J.A.D., Vandapel, N., Hebert, M.: Contextual classification with functional Max-Margin Markov networks. *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* pp. 975 – 982 (June 2009)
31. Nguyen, D., Lai, Y.L., Zhang, Q., Gyawali, P., Wang, F., Schwab, B., Wysocki, O., Kolbe, T.H.: Truecity: Real and simulated urban data for cross-domain 3d scene understanding. In: *Thirteenth International Conference on 3D Vision*
32. Pantoja-Rosero, B.G., Achanta, R., Kozinski, M., Fua, P., Perez-Cruz, F., Beyer, K.: Generating LoD3 building models from structure-from-motion and semantic segmentation. *Automation in Construction* **141**, 104430 (2022)
33. Pellis, E., Masiero, A., Betti, M., Tucci, G., Grussenmeyer, P.: A photogrammetric image-point dataset for the semantic segmentation of heritage buildings. *Data in Brief* **60**, 111661 (2025). <https://doi.org/https://doi.org/10.1016/j.dib.2025.111661>, <https://www.sciencedirect.com/science/article/pii/S2352340925003919>
34. Phan, A.V., Le Nguyen, M., Nguyen, Y.L.H., Bui, L.T.: Dgcnn: A convolutional neural network over large-scale labeled graphs. *Neural Networks* **108**, 533–543 (2018)

35. Qi, C.R., Su, H., Mo, K., Guibas, L.J.: Pointnet: Deep learning on point sets for 3d classification and segmentation. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 652–660 (2017)
36. Qi, C.R., Yi, L., Su, H., Guibas, L.J.: Pointnet++: Deep hierarchical feature learning on point sets in a metric space. *Advances in neural information processing systems* **30** (2017)
37. Riegler, G., Osman Ulusoy, A., Geiger, A.: Octnet: Learning deep 3d representations at high resolutions. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 3577–3586 (2017)
38. Robert, D., Raguet, H., Landrieu, L.: Efficient 3d semantic segmentation with superpoint transformer. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 17195–17204 (2023)
39. Roynard, X., Deschaud, J.E., Goulette, F.: Paris-Lille-3D: A large and high-quality ground-truth urban point cloud dataset for automatic segmentation and classification. *The International Journal of Robotics Research* **37**(6), 545–557 (2018)
40. Serna, A., Marcotegui, B., Goulette, F., Deschaud, J.E.: Paris-rue-Madame database: As 3D mobile laser scanner dataset for benchmarking urban detection, segmentation and classification methods. In: Proceedings of the International Conference on Pattern Recognition Applications and Methods. ACM, Angers, France, 6–8 March. pp. 819–824 (2014)
41. Szeliski, R.: Computer vision: Algorithms and applications. Springer Science & Business Media (2010)
42. Tan, W., Qin, N., Ma, L., Li, Y., Du, J., Cai, G., Yang, K., Li, J.: Toronto-3D: A large-scale mobile LiDAR dataset for semantic segmentation of urban roadways. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW). pp. 202–203 (2020)
43. Thomas, H., Qi, C.R., Deschaud, J.E., Marcotegui, B., Goulette, F., Guibas, L.J.: Kpconv: Flexible and deformable convolution for point clouds. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 6411–6420 (2019)
44. Vallet, B., Brédif, M., Serna, A., Marcotegui, B., Paparoditis, N.: TerraMobilita/iQmulus urban point cloud analysis benchmark. *Computers & Graphics* **49**, 126–133 (2015)
45. Wang, P.S.: Octformer: Octree-based transformers for 3d point clouds. *ACM Transactions on Graphics (TOG)* **42**(4), 1–11 (2023)
46. Wu, X., Jiang, L., Wang, P.S., Liu, Z., Liu, X., Qiao, Y., Ouyang, W., He, T., Zhao, H.: Point transformer v3: Simpler faster stronger. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 4840–4851 (2024)
47. Wu, X., Lao, Y., Jiang, L., Liu, X., Zhao, H.: Point transformer v2: Grouped vector attention and partition-based pooling. *Advances in Neural Information Processing Systems* **35**, 33330–33342 (2022)
48. Wysocki, O., Hoegner, L., Stilla, U.: TUM-FAÇADE: Reviewing and enriching point cloud benchmarks for façade segmentation. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences* **XLVI-2/W1-2022**, 529–536 (2022)
49. Wysocki, O., Tan, Y., Froech, T., Xia, Y., Wysocki, M., Hoegner, L., Cremers, D., Holst, C.: ZAHA: Introducing the Level of Facade Generalization and the Large-Scale Point Cloud Facade Semantic Segmentation Benchmark Dataset. In: Proceedings of the Winter Conference on Applications of Computer Vision (WACV). pp. 7637–7647 (February 2025)

50. Wysocki, O., Xia, Y., Wysocki, M., Grilli, E., Hoegner, L., Cremers, D., Stilla, U.: Scan2LoD3: Reconstructing semantic 3D building models at LoD3 using ray casting and Bayesian networks. *IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)* pp. 6547–6557 (2023)
51. Yang, Y.Q., Guo, Y.X., Xiong, J.Y., Liu, Y., Pan, H., Wang, P.S., Tong, X., Guo, B.: Swin3d: A pretrained transformer backbone for 3d indoor scene understanding. *Computational Visual Media* **11**(1), 83–101 (2025)
52. Yeshwanth, C., Liu, Y.C., Nießner, M., Dai, A.: ScanNet++: A high-fidelity dataset of 3d indoor scenes. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 12–22 (2023)
53. Zhao, H., Jiang, L., Jia, J., Torr, P.H., Koltun, V.: Point transformer. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 16259–16268 (2021)
54. Zhou, H., Feng, Y., Fang, M., Wei, M., Qin, J., Lu, T.: Adaptive graph convolution for point cloud analysis. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 4965–4974 (2021)
55. Zhu, J., Gehring, J., Huang, R., Borgmann, B., Sun, Z., Hoegner, L., Hebel, M., Xu, Y., Stilla, U.: TUM-MLS-2016: An annotated mobile LiDAR dataset of the TUM City Campus for semantic point cloud interpretation in urban areas. *Remote Sensing* **12**(11), 1875 (2020)