Why Has LOD3 Not Scaled? From Reconstruction Algorithms to the Generation of 3D City Models

Mehmet Akif Ortak[1], Burak Bek[1], and Jörg Rainer Noennig[1]

[1] HafenCity University Hamburg, Henning-Voscherau-Platz 1, 20457 Hamburg, Germany - mehmet.ortak@hcu-hamburg.de

Keywords: CityGML, LOD3, Digital twins, 3D city models, Facade modeling

Abstract

Level of Detail 3 (LOD3) has been part of the CityGML standard for nearly two decades, yet it has never become a routine product in operational 3D city model production. While LOD2 building models are now widely generated and maintained by mapping agencies, LOD3 remains largely limited to research prototypes and pilot projects. This paper examines why this gap persists despite sustained progress in photogrammetry, computer vision, and AI-assisted facade reconstruction. We examine LOD3 as a production ecosystem problem involving application demand, industrial workflows, quality assurance, maintenance, and production effort, not as another reconstruction algorithm. We argue that semantically enriched LOD2 remains sufficient for many urban applications, whereas explicit facade geometry can add analytical value for applications such as material inventories, BIM–GIS integration, entrance-aware navigation, flood analysis, and radio propagation simulation in the Next Generation City Networking (NGCN) context. Building on this distinction, we present an industrial pilot in which AI-assisted window and door detection is integrated into an existing photogrammetric LOD2 production workflow and validated by experienced operators. The final pilot dataset contains 505 LOD3 buildings, produced at approximately 50–75% lower operator effort than fully manual photogrammetric restitution according to indicative production estimates rather than controlled benchmarks. We hypothesize that advances in reconstruction technology alone are insufficient to explain or overcome the limited adoption of LOD3, and that ecosystem-level constraints have become an important barrier alongside the reconstruction challenges that remain. The pilot demonstrates workflow integration and operational feasibility rather than city-, regional-, or national-scale scalability.

1. Introduction

Semantic 3D city models have become a core component of Urban Digital Twins, supporting applications from urban planning and infrastructure management to environmental simulation, telecommunications, and disaster response. CityGML has emerged as the de facto standard for semantic 3D city modeling, providing a common representation of geometry, semantics, topology, and appearance across application domains (Gröger et al., 2008, Kolbe, 2009, Gröger and Plümer, 2012).

Advances in airborne photogrammetry, laser scanning, and dense image matching have transformed Level of Detail 2 (LOD2) from a research topic into a routine geospatial product. National and municipal mapping programs now maintain LOD2 datasets at scale, supported by mature software ecosystems optimized for roofscapes, building envelopes, and terrain integration (Buyukdemircioglu and Kocaman, 2020, Wysocki et al., 2024, Landesbetrieb Geoinformation und Vermessung Hamburg, 2026). The situation is markedly different for Level of Detail 3 (LOD3). Although part of the CityGML standard since its initial release, operational LOD3 datasets remain rare. Photogrammetric reconstruction from oblique imagery, terrestrial laser scanning, procedural modeling (Müller et al., 2006), and AI-based facade parsing have all produced usable results for windows, doors, and balconies, yet these advances have not translated into routine city-scale production (Harshit et al., 2024, Ogawa et al., 2024, Hanke et al., 2025, Wysocki et al., 2022). As illustrated in Figure 1, LOD3 extends the simplified LOD2 building envelope with explicit facade and roof elements such as windows, doors, balconies, and chimneys.

Demand for facade-level information is increasing, driven by

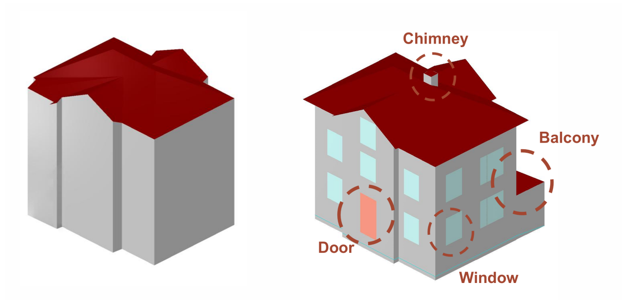

Figure 1. Comparison of LOD2 and LOD3 building representation.

material stock estimation, building renovation planning, urban energy modeling, radio propagation analysis in the Next Generation City Networking (NGCN) project, and high-fidelity urban simulation (Tanikawa and Hashimoto, 2009, Heisel and Rau-Oberhuber, 2020, Johari et al., 2022, Wang et al., 2025, Iye et al., 2025). At the same time, not every application requires full LOD3 geometry: prior work has shown that additional geometric detail does not always improve analytical results sufficiently to justify the extra production effort (Biljecki et al., 2016b). Identifying these boundaries is therefore as important as improving reconstruction methods.

This suggests that the remaining challenges extend beyond geometry reconstruction to production workflows, quality assurance, maintenance, specification ambiguity, benchmark availability, and economic sustainability (Biljecki et al., 2016a, Harshit et al., 2024). The gap between technical feasibility and operational production has received insufficient attention in the literature.

This paper addresses two related questions: why has LOD3 not emerged as an operational standard despite nearly two decades of research, and under what conditions does facade-level detail provide sufficient additional value over LOD2 to justify its higher production effort? Drawing on existing literature, industry collaboration, and practical experience gained during the development of an AI-assisted production workflow, the paper examines these questions from an operational and application-oriented perspective.

This paper makes three contributions. First, it provides an ecosystem-level synthesis of the factors limiting the operational adoption of LOD3, formalized as the LOD3 Gap: a dependency chain linking application demand, facade acquisition, reconstruction, human validation, quality assurance, maintenance, production effort, and demonstrated application value (Section 3.3). Second, it introduces an application-oriented distinction between cases where semantically enriched LOD2 is sufficient and those where explicit facade geometry creates additional analytical value. Third, it documents a hybrid AI-assisted production workflow that integrates facade-level objects into an existing photogrammetric CityGML pipeline. Together, these contributions support the central perspective of the paper: reconstruction capability has advanced considerably, but the surrounding production ecosystem has not matured at the same pace.

1.1 Research Approach

This conceptual and perspective paper is grounded in production practice rather than a controlled empirical study. We combine three complementary evidence streams: a synthesis of literature on CityGML, LOD3 reconstruction, and facade-dependent applications; exploratory interactions with industry and practitioner communities; and an industrial LOD3 pilot conducted in an operational production environment. We synthesize these sources to identify recurring technical, organizational, workflow, and production constraints, which we organize along the production chain formalized in Section 3.3. Practitioner interactions were not conducted as a structured survey, and the study does not claim statistical generalizability. The pilot demonstrates operational integration and feasibility rather than a controlled benchmark of scalability.

2. State of the Art: The LOD3 Ecosystem

2.1 Evolution of CityGML Levels of Detail

The Level of Detail concept has been a defining component of CityGML since the first OGC specification was published in 2008, combining geometry, semantics, and intended application into a unified multi-scale framework ranging from regional footprints (LOD0) to architecturally detailed building shells (LOD3) and interior representations (LOD4) (Gröger et al., 2008, Kolbe, 2009). Practical adoption exposed conceptual limitations in this hierarchy, and a more flexible multi-representation concept was later proposed (Löwner et al., 2016). Biljecki et al. (2016a) showed that the original definitions mixed geometric and semantic complexity in ways that produced ambiguity between neighboring levels, while a complementary study showed that higher geometric detail does not uniformly improve analytical outcomes: the benefit of moving from LOD2 to LOD3 is application-dependent rather than universal (Biljecki et al., 2016b). CityGML 3.0 responded by removing LOD4, treating interior modeling as an orthogonal dimension (Kutzner et al., 2020, Open Geospatial Consortium, 2021). However, while the conceptual model matured, the production challenges specific to LOD3 remained largely unaddressed.

2.2 Existing LOD3 Production Approaches

Manual and operator-driven workflows represent the earliest operational route to LOD3 production. The resulting datasets were produced through manual modeling in CAD and SketchUp-based environments with CityGML export plugins, using terrestrial photography and point clouds as geometric reference (Chair of Geoinformatics, TUM, 2023, Wysocki et al., 2024). These projects confirm that geometrically and semantically compliant LOD3 models can be generated manually. Semi-automated alternatives still require substantial human intervention for cross-registration, topology checks, and schema conversion (Harshit et al., 2024). Their fundamental limitation is throughput: production chains that depend on per-building editing or validation remain economically difficult to extend to routine city-scale maintenance.

Airborne photogrammetry and oblique imagery extend the acquisition strategies used for conventional city-scale mapping. Standard nadir airborne acquisition, which underpins LOD2 production at national scale, provides insufficient facade coverage for LOD3. Oblique aerial imaging systems, which simultaneously acquire nadir and lateral views, substantially improve facade visibility and have enabled semi-automated extraction of windows and facade openings (Tuttas and Stilla, 2013, Haala and Kada, 2010). However, mutual occlusion in dense urban environments, limited ground-floor coverage constrained by flight altitude, and the absence of integrated pipelines from oblique imagery to CityGML-compliant LOD3 remain unresolved. A key limitation is that most operational airborne acquisition campaigns continue to be planned around LOD2 requirements rather than facade reconstruction, limiting the facade completeness of the resulting datasets.

Street-level acquisition and terrestrial approaches address this facade visibility problem more directly. Mobile laser scanning (MLS) and terrestrial imagery provide substantially richer facade detail than airborne sensors and have become the dominant acquisition strategy in research-oriented LOD3 reconstruction. Several research groups have demonstrated pipelines from point clouds and street-level imagery that address facade segmentation, opening detection, and CityGML export (Wysocki et al., 2022, Pantoja-Rosero et al., 2022, Ogawa et al., 2024). The TUM2TWIN initiative has produced one of the most thoroughly documented openly accessible LOD3 benchmark datasets available to the research community (Wysocki et al., 2026). The operational limitations of terrestrial approaches are well established: rear and courtyard facades remain inaccessible without supplementary acquisition; urban canyons create persistent occlusion; and terrestrial campaigns require separate mobilization from the airborne workflows used for LOD2, substantially increasing per-building cost.

Recent open-data reviews make this imbalance measurable. Wysocki et al. (2024) inventoried 27 open semantic 3D city model collections comprising an estimated 216.5 million building models. Yet only 98 LOD3 buildings were reported across the surveyed European datasets, compared with 851 LOD3 buildings in Japan. The result is not a general shortage of 3D city models, but a shortage of facade-grade, ground-truth LOD3 examples suitable for training and benchmarking automated reconstruction.

AI-assisted reconstruction and hybrid workflows now form the most active area of methodological development. Recent deep learning methods have produced the most significant advances in automated LOD3 reconstruction. Conflict-map methods combine MLS point clouds with low-level 3D building priors to identify facade openings, while end-to-end image-based frameworks match omnidirectional street imagery with building data to create high-detail 3D models (Hanke et al., 2025, Ogawa et al., 2024). These methods show that automated facade enrichment is technically feasible, but they also depend on controlled acquisition, accurate data alignment, and validation-ready input models (Harshit et al., 2024, Hanke et al., 2025). A broader trend toward hybrid production workflows has also emerged, combining automated facade detection with operator validation rather than pursuing full automation. However, AI-based methods have been evaluated primarily on small or curated benchmark datasets; generalization to heterogeneous building stock and variable production-scale data quality remains an open challenge (Wysocki et al., 2022, Wysocki et al., 2024). No hybrid workflow has yet been demonstrated in routine city-scale deployment.

Table 1 compares LOD3 approaches by their fit with production environments that already maintain city-scale LOD2 models. The scale potential reported there expresses an expectation derived from the reviewed literature rather than a demonstrated result; city-scale deployment has not been demonstrated for any of these pathways in the present study. Routine LOD3 production may nevertheless be more likely to emerge from targeted extensions of existing photogrammetric workflows than from isolated reconstruction pipelines.

Approach	Main operational characteristic	Scale potential
Single-building manual modeling	Accurate geometry and full semantic control, but per-building manual work limits use to landmarks and small pilot areas.	Low
Photogrammetric restitution within the LOD2 production line	Uses established orientation, restitution, editing, quality assurance, and CityGML export environments, but manual facade restitution remains operator-intensive.	Medium
AI-supported photogrammetric production workflow	Adds AI suggestions for repetitive facade elements to the existing photogrammetric line. Operators validate, correct, and handle complex cases while preserving production control and semantic quality assurance.	Medium–High

Table 1. Operational comparison of LOD3 production pathways. Scale potential is indicative and based on the reviewed literature.

2.3 Critical Synthesis

The literature reviewed above shows that the individual technical components required for LOD3 reconstruction, including point cloud processing, facade segmentation, opening detection, semantic classification, and CityGML encoding, have each advanced substantially across multiple research groups. Acquisition conditions, occlusion, and complex facade elements nevertheless continue to constrain automated methods, while the organizational and economic factors that determine whether a feasible reconstruction method becomes an operational product have received comparatively little attention (Harshit et al., 2024, Ogawa et al., 2024, Hanke et al., 2025). The following section therefore develops the hypothesis that ecosystem-level constraints have become an important barrier alongside the reconstruction challenges that remain.

3. Why Did LOD3 Fail to Emerge?

Substantial technical foundations for LOD3 reconstruction are in place, yet LOD3 has not evolved into an operational production standard comparable to LOD2. Technical barriers clearly persist, particularly for complex facade elements, acquisition conditions, occlusion, and validation. Our synthesis of practitioner interactions and pilot experience suggests that ecosystem-level constraints are also an important part of the adoption bottleneck. We state this as a hypothesis rather than a demonstrated result: advances in reconstruction technology alone appear insufficient to explain or overcome the limited adoption of LOD3.

3.1 Demand and Business Incentives

The first barrier is economic: LOD3 production involves materially more operator effort than established LOD2 workflows (indicative operator-effort estimates are reported in Table 2). Existing LOD3 datasets have mostly been developed within research projects, pilot implementations, or application-specific proof-of-concepts. Projects such as Ingolstadt (SAVe/SAVeNoW) and the TUM campus in Munich (TUM2TWIN) show that LOD3 production is technically feasible (Chair of Geoinformatics, TUM, 2023, Wysocki et al., 2026). However, few have evolved into continuously maintained workflows covering entire cities or regions.

During exploratory discussions with software vendors, mapping agencies, and industrial partners in 2024–2025, one remark came up again and again:

> *“Nobody has asked us for LOD3 before.”*

Consistent with the research approach described in Section 1.1, we treat this remark as a practitioner signal rather than as survey evidence. It nevertheless points to a market condition that the published record confirms. Unlike LOD2, which became a standard deliverable for municipalities and national mapping agencies, LOD3 has rarely appeared as an explicit contractual requirement. Software vendors therefore had little incentive to optimize production pipelines beyond research-oriented implementations.

Published examples show the same pattern: open LOD3 benchmarks and recent production studies remain concentrated in campus- or neighborhood-scale test areas rather than in continuously maintained national datasets, whereas LOD2 production is embedded in municipal and national mapping practice (Wysocki et al., 2022, Harshit et al., 2024, Ogawa et al., 2024, Hanke et al., 2025). This asymmetry reflects where procurement has historically been concentrated. As long as LOD2 remained sufficient for mapping, planning, visualization, and asset management, municipalities had little reason to procure more detailed production workflows (Biljecki et al., 2016b), and LOD3 did not enter the reinforcing cycle of demand, investment, and declining production effort formalized in Section 3.3.

Emerging applications in material stock estimation, renovation planning, NGCN applications, and urban simulation may create the economic incentives needed for long-term investment in operational LOD3 production.

3.2 Labor-Intensive Production Workflows

Beyond weak demand, a second constraint is that LOD3 production remains labor-intensive even where reconstruction methods are technically available. Artificial intelligence has significantly improved facade interpretation, but production remains fundamentally human-in-the-loop (Harshit et al., 2024, Hanke et al., 2025); Section 7 reports indicative operator-effort estimates for this division of labor.

A representative workflow begins with conventional photogrammetric reconstruction, producing a geometrically consistent LOD2 building model from aerial imagery and dense point clouds. AI-based object detection is subsequently applied to identify facade openings, primarily windows and doors, which are then integrated into the reconstruction workflow (Ogawa et al., 2024, Hanke et al., 2025). Experienced operators review and correct these detections before export to CityGML; more complex architectural elements are handled through hybrid production steps, as detailed in Section 6. AI therefore assists operators rather than replacing them.

Data acquisition further reinforces this challenge. Most airborne mapping campaigns have historically been designed for LOD2 production, where nadir imagery provides sufficient roof coverage. LOD3 reconstruction, however, depends on substantially improved facade visibility obtained through carefully planned oblique imagery (Tuttas and Stilla, 2013, Haala and Kada, 2010). Existing oblique acquisition campaigns may already provide useful facade information, but flight planning is still rarely optimized specifically for facade reconstruction. Acquisition strategy therefore deserves stronger consideration where LOD3 is an explicit objective, instead of treating facade modeling only as a downstream enhancement of LOD2-oriented data collection.

3.3 Formalizing the LOD3 Gap

Significant progress has been achieved in many individual components, but they have largely evolved independently rather than as parts of an integrated production system. We formalize the *LOD3 Gap* as a dependency chain:

Application Demand → Facade Acquisition → Facade Reconstruction → Human Validation → Quality Assurance (QA) → Maintenance → Production Effort → Application Value.

Constraints at any of these stages can limit operational production regardless of progress elsewhere in the chain. The relationship is also recursive: demonstrated application value can stimulate future demand and investment, while weak demand limits investment, keeps production effort high, restricts LOD3 availability, and leaves fewer opportunities to demonstrate application value. Research has concentrated heavily on reconstruction, whereas the surrounding acquisition, validation, maintenance, and production conditions have received less systematic attention. In this framing, the LOD3 Gap describes the disconnect across the production and adoption chain rather than the absence of a single technology.

4. When Is LOD3 Really Necessary?

The preceding sections established that LOD3 production has remained technically feasible but economically marginal. The question that follows is more precise: for which applications does facade-level geometry create measurable analytical value that enriched LOD2 cannot provide? A useful organizing principle emerges from the application evidence: applications that treat buildings as statistical objects generally work well with enriched LOD2, whereas applications that interact directly with the building envelope may benefit from explicit facade geometry (Biljecki et al., 2016b). The distinction is not categorical. Required detail depends on analytical scale, accuracy requirements, and modeling approach. We therefore use a narrower criterion: LOD3 is warranted where explicit facade geometry adds analytical value that enriched LOD2 cannot provide, rather than whenever facade-related information is required.

4.1 Applications Where Enriched LOD2 Is Sufficient

For the majority of established urban modeling workflows, standard or semantically enriched LOD2 provides sufficient geometric fidelity at substantially lower production effort. Large-scale solar potential assessment, planning visualization, population density estimation, noise propagation modeling, and broad-scale energy demand aggregation operate at a spatial resolution where prismatic volumes with attributed properties are adequate (Biljecki et al., 2015, Biljecki et al., 2016b). In many energy simulation workflows, attaching window-to-wall ratios and construction-year classes as thematic properties to LOD2 surfaces captures the dominant sources of thermal variation without altering the underlying geometry (Johari et al., 2022, Xu et al., 2023).

4.2 Applications Requiring Explicit Facade Geometry

The common characteristic of the applications examined here is that they interact directly with the building envelope rather than with the building as a single volumetric object. For such applications, the question is not simply whether opening information is present, but whether geometric representation confers analytical advantages over attribute-based approximations.

This distinction first became apparent within the CircBoost project, where CityGML-derived material stock estimates were validated against BIM-derived reference models. For regional material cadastres, LOD2 volumes and archetype-based material intensities are often sufficient because most building mass is held in foundations, floors, and structural walls (Tanikawa and Hashimoto, 2009). The required level of detail changes once individual building components become the unit of analysis. LOD3 becomes relevant when the task shifts from aggregate mass estimation to object-level circular-economy tracking, where windows, doors, and facade components must remain distinct assets across BIM–GIS exchange (Heisel and Rau-Oberhuber, 2020, Honic et al., 2019). In the project that motivated this work, LOD3 was selected for this reason: explicit facade geometry provides a city-model representation substantially closer to the BIM-derived building-envelope information against which material stock calculations are compared. High-resolution thermal simulation faces a related constraint: peropening geometry drives solar gain and fabric heat loss calculations at the individual facade zone level in ways that attributed approximations cannot replicate (Johari et al., 2022, Xu et al., 2023).

A comparable dependency arises in flood inundation modeling and building-level damage assessment. Buildings are treated not as sealed volumes but as objects whose vulnerability depends on the relation between water levels and building-specific geometry. BIM–GIS flood assessment therefore benefits from detailed building geometry and 3D visualisation at the individual-building scale (Amirebrahimi et al., 2016). Entrance-aware routing illustrates the boundary of this criterion. Entrance locations can be stored as point attributes on an LOD2 model, which may be sufficient for many routing tasks. Explicit facade geometry adds value where the geometric relation between the opening, the building envelope, and the surrounding model must itself be evaluated, for example in building-level flood exposure. The analytical object is then no longer the building shell in general but the access interface. This is evident in last-mile logistics, where the operational destination is no longer a building centroid but a specific doorway connected to the street network. Commercial geocoding and navigation platforms now expose entrance coordinates as routable destinations, illustrating a shift from building-level to entrance-level location services (Mapbox, 2026).

A further potential dependency is being investigated in urban communication-network research within the NGCN project. At millimeter-wave frequencies used in 5G and emerging 6G systems, the physical scale of facade features, including window positions, structural reveals, and cladding material transitions, approaches the signal wavelength, suggesting an increasing geometric sensitivity in ray-tracing propagation models (Wang et al., 2025). Recent 5G mobility emulation work based on Project PLATEAU, OpenAirInterface, and Sionna RT illustrates the conversion bottleneck: CityGML is used as geometric input, but the propagation pipeline operates on exported mesh geometry with separately assigned electromagnetic materials rather than on the original semantic city model (Iye et al., 2025). We do not claim this as an established necessity; we identify it as a research direction whose practical value is still being quantified.

Beyond simulation and analysis, LOD3 also occupies a structurally distinct role in BIM–GIS integration. The exterior shell of a BIM model maps most naturally onto a CityGML LOD3 representation, and as building permit workflows digitize and public authorities increasingly receive BIM deliverables, LOD3 supports data exchange and improves the fidelity with which project-level information can be ingested into operational city models (Donkers et al., 2016, Noardo et al., 2020). This does not imply perfect interoperability, but it positions LOD3 as the most coherent transition layer between the two domains.

4.3 Implications for LOD3 Production

Because the applications that genuinely require explicit facade geometry form an identifiable subset rather than the majority of urban modeling tasks, a blanket commitment to LOD3 at municipal or national scale is difficult to justify on analytical grounds alone. LOD3 should instead be deployed selectively where facade geometry produces demonstrable analytical value, and such application-driven deployments provide a realistic pathway for expanding the production ecosystem.

5. From Research Demand to Operational Production

The pathway examined in this study did not begin with a new reconstruction algorithm. It began with a concrete application that exposed the limitations of existing LOD2 data. A material cadastre study comparing CityGML with BIM required close semantic and geometric correspondence between the two representations. Among the available CityGML representations, LOD3 provided the most suitable approximation of the BIM building envelope (Donkers et al., 2016, Noardo et al., 2020). Producing complete LOD3 models manually, however, was not viable for city-scale applications (Buyukdemircioglu and Kocaman, 2020, Harshit et al., 2024).

This requirement brought researchers, public authorities, and industrial software developers into close collaboration, shifting the discussion from isolated geometric reconstruction toward the extension of existing production workflows. LOD3 was thus treated not as a general-purpose upgrade of city models but as a specific production capability supporting emerging application domains.

Hamburg provided an important operational reference through its long-running municipal CityGML production environment, maintained by the State Office for Geoinformation and Surveying (LGV) together with industrial software providers (Landesbetrieb Geoinformation und Vermessung Hamburg, 2026). This existing production context provided the basis for collaboration with UVM Systems, whose software already supports photogrammetric CityGML production. Building on this experience, the LOD3 extension was subsequently developed and tested for the Belgrade pilot area. Our role was to formulate the application requirements and the required facade-level information, while UVM Systems implemented the software extensions and carried out the production.

The Belgrade pilot therefore examined whether facade-level LOD3 capabilities could be added to an established photogrammetric production chain covering image orientation, photogrammetric restitution, quality assurance, interactive editing, and CityGML export.

Development therefore focused on missing facade components rather than complete reconstruction. Roof superstructures were already reconstructed operationally, while AI-assisted methods targeted repetitive windows and doors (Ogawa et al., 2024, Hanke et al., 2025).

The resulting workflow combines AI with established photogrammetric practice instead of replacing it; the division of labor is described in Section 6. These observations come from a single production environment and are not statistically generalizable, but the pathway may be relevant to other cities that already maintain photogrammetric LOD2 workflows.

6. Hybrid LOD3 Production Pipeline

The hybrid workflow extends an existing operational LOD2 pipeline without replacing it: we defined the production requirements, and UVM Systems integrated the required facade-level capabilities into its established photogrammetric pipeline (Figure 2).

https://doi.org/10.5194/isprs-archives-L-4-W2-2026-155-2026 | © Author(s) 2026. CC BY 4.0 License.

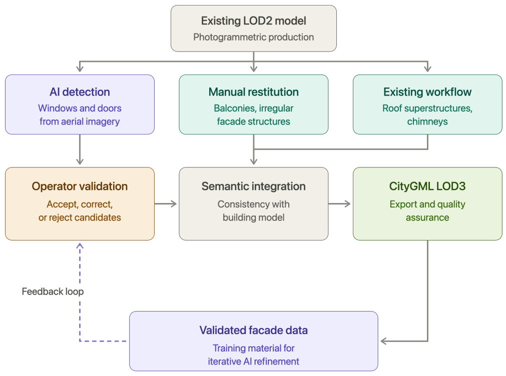

Figure 2. Hybrid LOD3 production workflow used in this study. AI proposes window and door candidates, operators validate them, complex facade structures remain operator-led, and validated facade data supports iterative AI refinement.

The unchanged LOD2 pipeline covers aerial image processing, orientation, stereo restitution, semantic reconstruction, quality assurance, and CityGML export. The extension adds facade-level semantic objects to this production environment. Mature CityGML delivery chains exist for LOD2, while facade-level enrichment remains a separate production challenge (Landesbetrieb Geoinformation und Vermessung Hamburg, 2026, Harshit et al., 2024).

The workflow employs AI-assisted methods to detect repetitive facade elements, primarily windows and doors, from aerial imagery. These detections do not generate the final LOD3 model directly; they enter the existing editing environment as candidate semantic objects associated with the corresponding building surfaces. Experienced operators subsequently validate the proposed detections, correct ambiguous cases, and ensure semantic and geometric consistency with the underlying building model (Ogawa et al., 2024, Hanke et al., 2025). Operators still reconstruct architectural structures with greater geometric variability, particularly balconies, using conventional photogrammetric restitution, while roof superstructures such as chimneys remain part of the established production workflow.

The result is a human-in-the-loop workflow: automation supports repetitive patterns, while expert interpretation handles complex elements and final quality assurance.

The strategy is designed to evolve rather than to eliminate manual work immediately. Operational production continuously generates validated facade geometries through routine quality control, creating reference data for later model refinement (Wysocki et al., 2022). This instantiates the feedback relation introduced in Section 3.3: application demand justifies investment, production generates validated data, and improved detectors reduce operator effort in later projects, while the photogrammetric backbone remains in place.

7. Operational Pilot Demonstration

We piloted the workflow, in cooperation with UVM Systems, in an operational production environment in Belgrade, Serbia, as part of the EU Horizon CircBoost project. The pilot area covers approximately 3 km^2 and includes around 600 predominantly residential buildings with different heights and facade typologies (Figure 3). Of these, 505 buildings were delivered as CityGML LOD3 in the final pilot dataset. UVM Systems first reconstructed the area through its established LOD2 production chain using combined aerial and mobile point-cloud acquisition. The pilot tested integration into an industrial CityGML production workflow, not the accuracy of a standalone AI method.

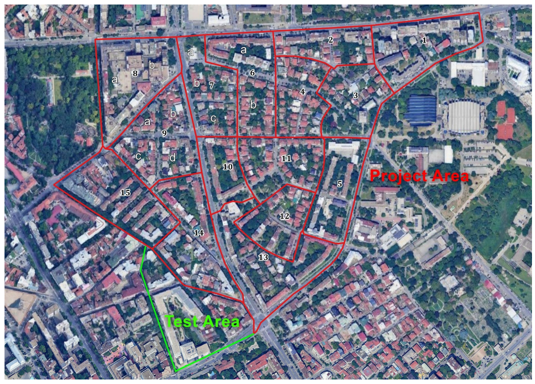

Figure 3. Belgrade pilot area in Palilula municipality. The red boundary indicates the project area and the green boundary indicates the test area used for the operational demonstration.

UVM Systems implemented the LOD3 extension within its CityGRID production environment, following the facade-level requirements defined in the project. AI-assisted methods generated candidate window and door objects on building facades, which experienced operators reviewed and corrected. Operators integrated accepted facade objects into the semantic building model and exported them as CityGML LOD3.

The pilot did not require replacement of the existing production chain. AI reduced repetitive editing for regular facade elements, while operators retained responsibility for correction, quality control, and complex objects. This division of labor matters because current automated LOD3 methods still struggle when acquisition, registration, and validation must be handled at city scale (Ogawa et al., 2024, Harshit et al., 2024).

The pilot also confirmed the limits of the approach. Occluded openings, irregular facade details, and locally specific architectural forms still require manual interpretation. Its practical value lies in adding facade-level semantics where they can be produced efficiently, without disrupting the existing LOD2 production chain.

Table 2 summarizes indicative operator-effort estimates provided by the industry partner from its production experience. Relative to fully manual photogrammetric LOD3 restitution, the hybrid AI-assisted workflow corresponds to approximately 50–75% lower operator effort.

These estimates were not obtained through a controlled time-and-motion study and should be interpreted as indicative production estimates rather than benchmark measurements. They vary with building complexity, input data quality, and operator experience and reflect predominantly residential buildings in the pilot area. In this human-in-the-loop workflow, throughput depends on the combined time for AI proposal, operator verification, and correction; standalone detector metrics (precision, recall, acceptance rate, and false-negative rate) were not recorded as controlled experimental measurements. The pilot therefore demonstrates workflow integration and operational feasibility, but not city-, regional-, or national-scale scalability.

Workflow	Estimated operator effort
Established LOD2 production	2–4 min/building
Earlier manual SketchUp-style LOD3	3–5 h/building
Manual photogrammetric LOD3 restitution	∼60 min/building
Hybrid AI-assisted LOD3 workflow	15–30 min/building

Table 2. Indicative operator-effort estimates provided by the industry partner based on production experience. LOD2 production and earlier manual SketchUp-style LOD3 modeling are included as reference baselines only.

8. Conclusions

This paper examines the conditions required for operational LOD3 production rather than another reconstruction method (Biljecki et al., 2016a, Harshit et al., 2024, Ogawa et al., 2024, Hanke et al., 2025).

We draw three conclusions. First, we hypothesize that reconstruction capability alone is insufficient to explain the limited adoption of LOD3: ecosystem-level constraints have become an important barrier alongside the reconstruction challenges that remain, as captured by the LOD3 Gap framework of Section 3.3. Second, LOD3 should be produced selectively where explicit facade geometry creates demonstrable analytical value, rather than treated as the default representation for all city models (Biljecki et al., 2016b). Third, AI can be introduced into existing municipal production environments through a hybrid workflow that combines automated facade detection with photogrammetric expertise and established quality-control procedures (Harshit et al., 2024, Ogawa et al., 2024, Hanke et al., 2025).

The Belgrade pilot illustrates this perspective (Figure 4): 505 LOD3 buildings were delivered through an existing industrial workflow at approximately 50–75% lower operator effort than fully manual photogrammetric restitution, demonstrating workflow integration and operational feasibility rather than scalability.

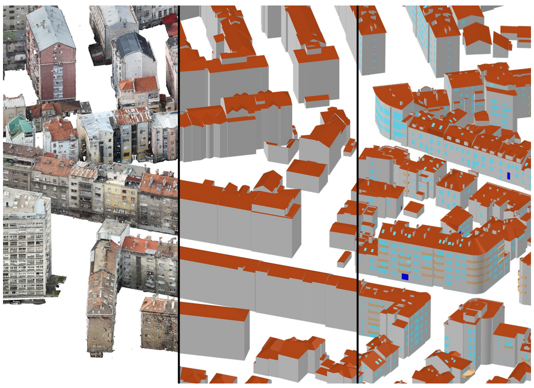

Figure 4. From measured point-cloud data to an operational LOD3 city model with explicit semantic facade and roof objects. The resulting model retains facade and roof objects as semantic components.

The evidence suggests that bottlenecks extend beyond geometry reconstruction to the production and adoption chain, while substantial reconstruction challenges remain. Future work should test larger production campaigns, quantify detector performance under controlled conditions, and evaluate whether operator-effort reductions persist across building typologies and acquisition conditions.

Acknowledgements

This research originated in the European Union's Horizon Europe programme under Grant Agreement No. 101082068 (CircBoost), where the need for facade-level semantic 3D city models emerged in the context of circular economy and material stock estimation, and was further advanced through the project *Next Generation City Networking* (Grant No. 19DZ24004) at the Hanseatic Wireless Innovation Competence Center (HAW-ICC), funded by the German Federal Ministry of Transport through the German Center for Future Mobility (DZM).

The authors used OpenAI's ChatGPT and Codex to support language editing and drafting, and take full responsibility for the final content.

References

Amirebrahimi, S., Rajabifard, A., Mendis, P., and Ngo, T., 2016. A BIM-GIS integration method in support of the assessment and 3D visualisation of flood damage to a building. *J. Spat. Sci.*, 61(2), 317–350. doi:10.1080/14498596.2016.1189365.

Biljecki, F., Stoter, J., Ledoux, H., Zlatanova, S., and Çöltekin, A., 2015. Applications of 3D City Models: State of the Art Review. *ISPRS Int. J. Geo-Inf.*, 4(4), 2842–2889. doi:10.3390/ijgi4042842.

Biljecki, F., Ledoux, H., and Stoter, J., 2016a. An Improved LOD Specification for 3D Building Models. *Comput. Environ. Urban Syst.*, 59, 25–37. doi:10.1016/j.compenvurbsys.2016.04.005.

Biljecki, F., Ledoux, H., Stoter, J., and Vosselman, G., 2016b. The variants of an LOD of a 3D building model and their influence on spatial analyses. *ISPRS J. Photogramm. Remote Sens.*, 116, 42–54. doi:10.1016/j.isprsjprs.2016.03.003.

Buyukdemircioglu, M. and Kocaman, S., 2020. Reconstruction and Efficient Visualization of Heterogeneous 3D City Models. *Remote Sens.*, 12(13), 2128. doi:10.3390/rs12132128.

Chair of Geoinformatics, TUM, 2023. LoD3 modeling based on mobile laser scans. Creating CityGML Datasets documentation, Technical University of Munich. https://creating-citygml-datasets.readthedocs.io/en/latest/creation-guidelines/lod3-models-based-on-point-clouds.html (accessed 14 July 2026).

Donkers, S., Ledoux, H., Zhao, J., and Stoter, J., 2016. Automatic conversion of IFC datasets to geometrically and semantically correct CityGML LOD3 buildings. *Trans. GIS*, 20(4), 547–569. doi:10.1111/tgis.12162.

Gröger, G., Kolbe, T. H., Czerwinski, A., and Nagel, C. (Eds.), 2008. OpenGIS City Geography Markup Language (CityGML) Encoding Standard, Version 1.0.0. Open Geospatial Consortium, OGC Doc No. 08-007r1.

Gröger, G. and Plümer, L., 2012. CityGML: Interoperable Semantic 3D City Models. *Comput. Geosci.*, 49, 136–145. doi:10.1016/j.cageo.2012.07.011.

Haala, N. and Kada, M., 2010. An update on automatic 3D building reconstruction. *ISPRS J. Photogramm. Remote Sens.*, 65(6), 570–580. doi:10.1016/j.isprsjprs.2010.09.006.

Hanke, F., Bieringer, A., Wysocki, O., and Jutzi, B., 2025. CM2LoD3: Reconstructing LoD3 Building Models Using Semantic Conflict Maps. *ISPRS Ann. Photogramm. Remote Sens. Spatial Inf. Sci.*, X-4/W6-2025, 81–88. doi:10.5194/isprs-annals-X-4-W6-2025-81-2025.

Harshit, Chaurasia, P., Zlatanova, S., and Jain, K., 2024. Low-Cost Data, High-Quality Models: A Semi-Automated Approach to LOD3 Creation. *ISPRS Int. J. Geo-Inf.*, 13(4), 119. doi:10.3390/ijgi13040119.

Heisel, F. and Rau-Oberhuber, S., 2020. Calculation and evaluation of circularity indicators for the built environment using the case studies of UMAR and Madaster. *J. Clean. Prod.*, 243, 118482. doi:10.1016/j.jclepro.2019.118482.

Honic, M., Kovacic, I., and Rechberger, H., 2019. Improving the recycling potential of buildings through Material Passports (MP): An Austrian case study. *J. Clean. Prod.*, 217, 787–797. doi:10.1016/j.jclepro.2019.01.212.

Iye, T., Sakamoto, M., Takaya, S., Sato, E., Susukida, Y., Nagaoka, Y., Maruta, K., and Nakazato, J., 2025. Open Wireless Digital Twin: End-to-End 5G Mobility Emulation With OpenAirInterface and Ray Tracing. *IEEE Access*. doi:10.1109/ACCESS.2025.3619105.

Johari, F., Munkhammar, J., Shadram, F., and Widén, J., 2022. Evaluation of simplified building energy models for urban-scale energy analysis of buildings. *Build. Environ.*, 211, 108684. doi:10.1016/j.buildenv.2021.108684.

Kolbe, T. H., 2009. Representing and Exchanging 3D City Models with CityGML. In: Lee, J. and Zlatanova, S. (Eds.), *3D Geo-Information Sciences*, Springer, Berlin, Heidelberg, 15–31. doi:10.1007/978-3-540-87395-2_2.

Kutzner, T., Chaturvedi, K., and Kolbe, T. H., 2020. CityGML 3.0: New Functions Open Up New Applications. *PFG – J. Photogramm. Remote Sens. Geoinf. Sci.*, 88(1), 43–61. doi:10.1007/s41064-020-00095-z.

Landesbetrieb Geoinformation und Vermessung Hamburg, 2026. 3D-Gebäudemodell LoD2-DE Hamburg. MetaVer metadata record, last updated 19 May 2026. https://metaver.de/trefferanzeige (accessed 14 July 2026).

Löwner, M.-O., Gröger, G., Benner, J., Biljecki, F., and Nagel, C., 2016. Proposal for a New LOD and Multi-Representation Concept for CityGML. *ISPRS Ann. Photogramm. Remote Sens. Spat. Inf. Sci.*, IV-2/W1, 3–12. doi:10.5194/isprs-annals-IV-2-W1-3-2016.

Mapbox, 2026. Mapbox brings doorway-level accuracy to delivery, logistics, and ride-hailing. Press release, 3 February 2026. https://www.mapbox.com/press-releases/mapbox-brings-doorway-level-accuracy-to-delivery-logistics-and-ride-hailing (accessed 14 July 2026).

Müller, P., Wonka, P., Haegler, S., Ulmer, A., and Van Gool, L., 2006. Procedural modeling of buildings. *ACM Trans. Graph.*, 25(3), 614–623. doi:10.1145/1141911.1141931.

Noardo, F., Harrie, L., Arroyo Ohori, K., Biljecki, F., Ellul, C., Krijnen, T., Eriksson, H., Guler, D., Hintz, D., Jadidi, M. A., Pla, M., Sanchez, S., Soini, V.-P., Stouffs, R., Tekavec, J., and Stoter, J., 2020. The GeoBIM benchmark project: Status and first results. *ISPRS Ann. Photogramm. Remote Sens. Spat. Inf. Sci.*, V-4-2020, 101–108. doi:10.5194/isprs-annals-V-4-2020-101-2020.

Ogawa, Y., Nakamura, R., Sato, G., Maeda, H., and Sekimoto, Y., 2024. End-to-end framework for the automatic matching of omnidirectional street images and building data and the creation of 3D building models. *Remote Sens.*, 16(11), 1858. doi:10.3390/rs16111858.

Open Geospatial Consortium, 2021. OGC City Geography Markup Language (CityGML) 3.0 Part 1: Conceptual Model Standard. OGC Doc No. 20-010r2. https://docs.ogc.org/is/20-010r2/20-010r2.html.

Pantoja-Rosero, B. G., Achanta, R., Kozinski, M., Fua, P., Perez-Cruz, F., and Beyer, K., 2022. Generating LoD3 building models from structure-from-motion and semantic segmentation. *Autom. Constr.*, 141, 104430. doi:10.1016/j.autcon.2022.104430.

Tanikawa, H. and Hashimoto, S., 2009. Urban stock over time: Spatial material stock analysis using 4d-GIS. *Build. Res. Inf.*, 37(5–6), 483–502. doi:10.1080/09613210903169394.

Tuttas, S. and Stilla, U., 2013. Reconstruction of façades in point clouds from multi aspect oblique ALS. *ISPRS Ann. Photogramm. Remote Sens. Spat. Inf. Sci.*, II-3/W3, 91–96. doi:10.5194/isprsannals-II-3-W3-91-2013.

Wang, B., Jiang, Z., Yang, Z., Huang, C., Shen, Y., Jiang, S., Zhu, C., Yang, Z., Jin, R., Zhang, Z., Muhaidat, S., and Debbah, M., 2025. DeepTelecom: A Digital-Twin Deep Learning Dataset for Channel and MIMO Applications. arXiv:2508.14507.

Wysocki, O., Hoegner, L., and Stilla, U., 2022. TUM-FAÇADE: Reviewing and enriching point cloud benchmarks for façade segmentation. *ISPRS Arch. Photogramm. Remote Sens. Spat. Inf. Sci.*, XLVI-2/W1-2022, 529–536. doi:10.5194/isprs-archives-XLVI-2-W1-2022-529-2022.

Wysocki, O., Schwab, B., Beil, C., Holst, C., and Kolbe, T. H., 2024. Reviewing Open Data Semantic 3D City Models to Develop Novel 3D Reconstruction Methods. *ISPRS Arch. Photogramm. Remote Sens. Spat. Inf. Sci.*, XLVIII-4-2024, 493–500. doi:10.5194/isprs-archives-XLVIII-4-2024-493-2024.

Wysocki, O., Schwab, B., Biswanath, M. K. et al., 2026. TUM2TWIN: Introducing the large-scale multimodal urban digital twin benchmark dataset. *ISPRS J. Photogramm. Remote Sens.*, 232, 810–830. doi:10.1016/j.isprsjprs.2025.12.013.

Xu, H., Kim, J. I., Zhang, L., and Chen, J., 2023. LOD2 for energy simulation (LOD2ES) for CityGML: A novel level of details model for IFC-based building features extraction and energy simulation. *J. Build. Eng.*, 78, 107715. doi:10.1016/j.jobe.2023.107715.