ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences, Volume XII-4/W1-2026
21st International 3D GeoInfo Conference 2026, 28–29 September 2026, Sofia, Bulgaria

A proposed framework for generating plausible cities

Oliver J. Post[1], Hugo Ledoux[1], Akshay Patil[1]

[1] Delft University of Technology, the Netherlands—olepost1@gmail.com [h.ledoux, a.l.patil]@tudelft.nl

Keywords: procedural city generation, urban morphology, 3D city models, typology, open-source

Abstract

Three-dimensional city models are widely used for simulation and analysis, yet their creation remains constrained by data availability, geometric errors, and the labour-intensive nature of manual modelling. Procedurally generated cities offer a practical alternative, but existing methods lack plausibility because they rely on generalised rule-based approaches rather than real urban data. In this paper, we propose the CityStack framework, a three-step pipeline for the analysis, encoding, and generation of plausible cities grounded in real-world morphological data. The framework decomposes urban form into hierarchical layers; in this paper, we focus on the minor road network and the building footprint layers. Using publicly available geospatial data from 43 cities worldwide, we introduce new morphological metrics for characterising street network patterns and apply a Gaussian Mixture Model to identify 13 global road typologies, alongside a supervised gradient boosting classifier that assigns one of 9 building footprint typologies. Urban form is encoded in a typology grid, a raster representation in which each 100 m $\times$ 100 m cell stores the dominant typology, serving as a compact, comparable intermediate format between analysis and generation. New city variants are synthesised by adapting typology grids via simulated annealing and generating roads and buildings from typology statistics. All components are implemented in an open-source command-line tool, making the full pipeline reproducible and accessible for a range of urban modelling applications.

1. Introduction

Three-dimensional city models offer significant value across multiple domains (Biljecki et al., 2015), and are often used as input for simulation, for instance, wind simulation (García-Sánchez et al., 2018; Patil and García-Sánchez, 2025), energy estimation (Kaden and Kolbe, 2013; León-Sánchez et al., 2021), noise modelling (Kumar et al., 2017; Hobeika et al., 2024), etc. However, due to difficulties in reconstructing 3D cities and in obtaining high-quality height datasets, the availability of existing models is somewhat limited.

Rather than reconstructing existing cities, we propose in this paper to generate procedurally *plausible* variations of existing cities, that is, cities that resemble the character of a real-life city while differing in their exact geometry and arrangement. In this paper, plausibility means morphological consistency with observed urban patterns: a generated city should preserve characteristic typology distributions and spatial organisation without reproducing the source city exactly. Doing this offers several advantages for research and practical use cases. High-quality synthetic 3D city models can be used to test the implementation of algorithms and file formats, to provide input for computational fluid dynamics (CFD) studies that are notoriously difficult to construct in practice (Pađen et al., 2024), and to serve as training data for machine learning models (Vargiemezis and Gorlé, 2025). It also circumvents privacy concerns associated with using real population data (Kim et al., 2018): a procedurally generated city, similar but not identical to a real one, allows population attributes to be detailed yet anonymous, for example in studies of economic segregation (Sarkar et al., 2024).

As further explained in Section 2, existing procedural generation methods for cities often lack *plausibility* (Groenewegen and Smelik, 2009) because they rely on generalised rule-based approaches instead of real city data, and often they focus only on one aspect (eg road network).

We address these limitations through a comprehensive urban morphological analysis and develop the *CityStack* framework to generate realistic city variations directly from the morphological patterns of existing urban environments. As shown in Figure 1, and further explained in Section 3, our framework abstracts and conceptualises cities as a hierarchy of layers, where lower layers constrain those above. Within this framework, we focus particularly on the more ambiguous layers, specifically the minor road networks and building structures, that significantly influence a city's visual character yet prove challenging to categorise (Oliveira, 2020). While existing city-generation methodologies that produce realistic city layouts focus solely on land use, our approach incorporates detailed morphological characteristics.

This paper is guided by two research questions: first, whether globally applicable and interpretable typologies for minor roads and building footprints can be derived from morphological data across many cities; and second, whether those typologies can be encoded and regenerated in a way that preserves the statistical structure of a reference city while producing a new layout. Our proposed framework is a pipeline composed of three steps: (1) analysis of the urban forms of current cities; (2) encoding of the form in a standardised format (which we call a typology grid), and (3) the generation of plausible cities by modifying the typology grids of a city. This framework is implemented in the open-source command-line tool (called `citypy` and available at `https://github.com/OliverJPost/CityStack`), which can reproduce all analysis and encoding results with simple commands, making urban morphology analysis accessible to researchers and practitioners without specialised expertise. Within the scope of this paper, success is measured by the interpretability and transferability of the typologies, the comparability of the typology grid, and the preservation of target typology distributions and landscape statistics in generated outputs. We have analysed 43 cities worldwide (including large ones like Paris, which required up

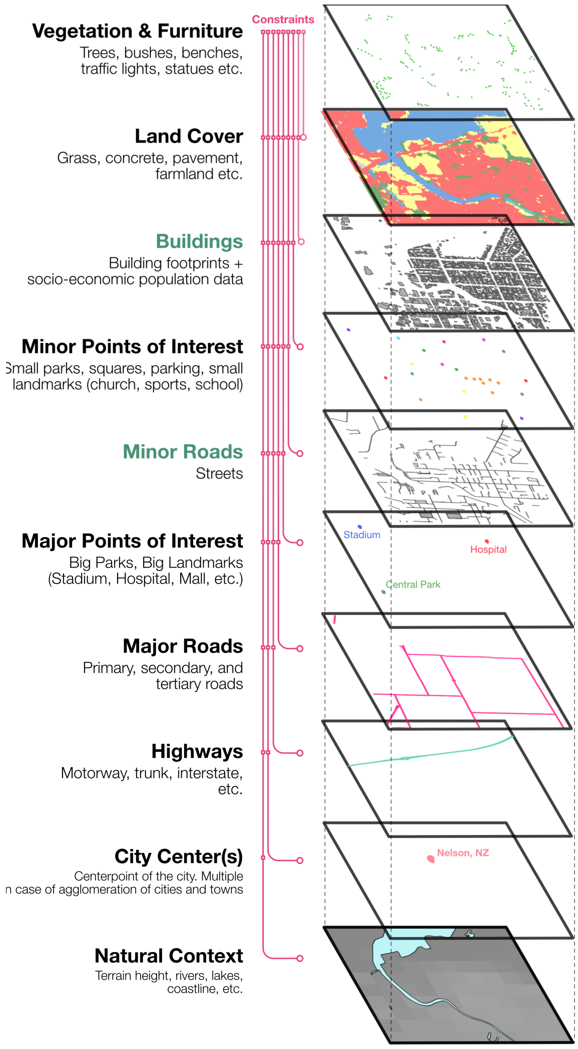

Figure 1. Schematic representation of the CityStack. Layers highlighted in green are investigated in this paper. Magenta lines indicate that lower layers in the stack constrain layers above.

to 100 GB of RAM), and we demonstrate in Sections 4 to 6 examples of cities generated with the framework.

The specific contributions of this paper are: (i) new morphological metrics for characterising street network patterns at the sub-city scale; (ii) a global set of 13 road network typologies and 9 building typologies derived from data covering 43 cities worldwide; (iii) a typology grid encoding that bridges morphological analysis and procedural generation; (iv) a modular generation pipeline that produces road networks and building footprints consistent with real urban morphology; and (v) an open-source implementation (available at `https://github.com/OliverJPost/CityStack`) that makes the full pipeline reproducible.

The remainder of the paper is structured as follows. Section 2 reviews related work on procedural city generation and urban morphology analysis. Section 3 presents the overview of the CityStack framework. Sections 4 through 6 describe the three framework steps. Section 7 presents a discussion and directions for future work.

2. Related work

2.1 Procedural generation of cities

Several methods for procedurally generating cities or their constituent layers exist. Parish and Müller (2001) adapts L-systems to grow road networks from image maps encoding terrain and population density; results are visually plausible but require manually crafted input maps that do not reflect real urban data. Chen et al. (2008) introduces tensor fields as a framework for interactive street layout design, yielding coherent orthogonal patterns but with limited applicability to organic urban fabrics. Aliaga et al. (2008) proposes using an existing road network as a template to extend road patterns locally; while the method produces plausible extensions, it does not support city-wide generation. For land use and city layout, Groenewegen and Smelik (2009) and Kim et al. (2018) discuss generation approaches, identifying plausibility as a core open problem: without grounding in real city data, generated layouts remain stylised rather than representative. A shared limitation across these methods is the reliance on hand-crafted rules that cannot systematically capture the diversity of urban forms observed worldwide. More recently, learning-based approaches such as StreetGAN (Hartmann et al., 2017) and the AETree model of Han et al. (2024) have shown that neural networks can synthesise road network patches or city object layouts from examples. However, these methods generally operate as end-to-end generators and offer less explicit control over layer-specific constraints, interpretability, and user adjustment than a typology-driven workflow.

2.2 Urban morphology analysis

Complementary to generation, a body of work focuses on the systematic analysis and classification of urban form. Marshall (2004) and Louf and Barthélemy (2014) classify road networks from geometric and topological properties; however, their typologies operate at the city scale and do not resolve the sub-city variation needed for per-neighbourhood generation. Badhrudeen et al. (2022) extends street pattern classification to a global context using geometric metrics, providing a basis for cross-city comparison. For buildings, Steiniger et al. (2008) and Wurm et al. (2016) define classification schemes grounded in morphological attributes, which form the basis for the supervised model used in this work. Fleischmann et al. (2022) demonstrates that unsupervised clustering applied to morphological metrics can identify contiguous urban tissue areas across multiple cities from open geospatial data; their approach, implemented in the `momepy` library, classifies all urban layers simultaneously. This conflates building and street pattern typologies into a combined representation, requiring a large number of clusters to describe all combinations, and does not provide a direct link to per-layer generation algorithms.

Our work addresses these limitations by analysing each layer separately using named and interpretable typologies, and by connecting the analysis directly to a layer-by-layer generation pipeline.

3. The CityStack framework

Our CityStack framework conceptualises and describes an existing city as a set of layers connected by constraints. Each layer describes a distinct aspect of urban form: from the natural context at the bottom to points of interest at the top. The key simplification is that constraints flow in a single direction from lower layers in the stack to higher ones, and only in that direction. Because of this single direction, city generation can proceed layer by layer from the bottom of the stack upward.

Our framework has three main steps:

1. **Analysis**: We capture the urban form of real-world cities, layer by layer, using publicly available data from OpenStreetMap[1].
2. **Encoding**: We represent the analysis results in a standardised template that we call *typology grids*. Those enable comparison across cities and serve as input for generation.
3. **Generation**: We produce new plausible cities, layer by layer, based on the encoded template. This is done using simulated annealing optimisation to create new typology grids that statistically *resemble* those of actual cities.

In this paper, we focus on the two layers that are most ambiguous to divide into distinct categories: (1) the *minor roads* layer, and (2) the *buildings* layer. This is because, in our opinion, these contribute most to the visual character of a city. Also, we restrict the analysis to 2D, simplifying complexity and enhancing data availability (but we discuss in Section 7 how 3D can, in the future, be incorporated in the framework to generate 3D models).

User input is an important aspect of city generation; we designed the typology-based approach to be intuitive, comprehensible, and adjustable.

4. Analysis of the urban fabric of cities

To develop a generative framework grounded in real-world data rather than generalised rules, we characterise the urban form of a large, geographically diverse set of cities using morphological metrics. As illustrated in Figure 1, the current CityStack implementation focuses on two layers: the minor road network and the building footprints. The analysis follows the programmatic approach of Fleischmann et al. (2022) but applies it independently to each layer, enabling separate typology sets and a direct connection to the per-layer generation stage.

4.1 Data acquisition

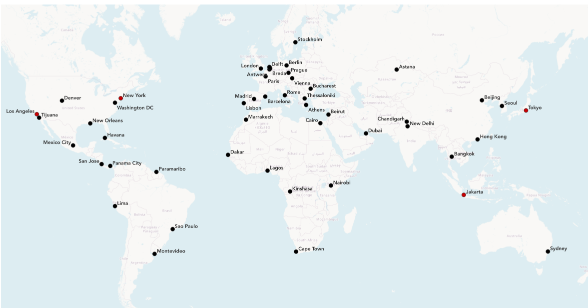

Figure 2. Cities considered for the building footprint and road network classification analysis. Black dots mark cities included in the analysis; red dots mark cities excluded due to analysis time constraints.

We assembled a geographically diverse dataset covering 43 cities worldwide (see Figure 2). Road network data were obtained from OpenStreetMap and simplified into four hierarchical levels: highway, major road, street, and path (based on the OSM tags). For building footprints, we used the Overture Maps Foundation dataset[2], which integrates OpenStreetMap, Google Open Buildings, Esri Community Maps, and Microsoft Open Building Footprints into a combined dataset of approximately 2.4 billion footprints, offering substantially better coverage in the Global South than OpenStreetMap alone. Since both sources lack height data for the majority of buildings, approximate heights were added from the GHSL Built-H dataset (Pesaresi, 2023), a global raster at 100 m resolution, as a building attribute. Seven cities were used as ground-truth datasets for supervised building classification; the remaining cities were classified using the trained model and visually verified against satellite imagery from Google Earth and Maxar Technologies.

4.2 Road network typologies

For road segments, we computed metrics capturing both individual segment properties and neighbourhood context. Several new metrics were introduced to address known limitations of existing approaches at the sub-city scale:

- *Continuity* measures how far one can travel in a straight direction without a significant turn, defined as the ratio of a road stretch length to the maximum stretch length (500 m).
- *Stretch curvilinearity* captures road curviness using a logarithmic adaptation of linearity applied to road stretches rather than individual sections, addressing known problems with section-based linearity (Fleischmann et al., 2022).
- *Right-neighbour distance and angle deviation* captures the spatial relationship between adjacent parallel streets.

Each metric was also computed as a weighted neighbourhood aggregate, so that each segment is described not only by its own properties but also by the character of its surroundings, following Fleischmann et al. (2022).

A Gaussian Mixture Model (Reynolds, 2009) was applied to this feature set across all 43 cities. The optimal number of clusters was determined at 13 by the elbow of the Bayesian Information Criterion curve. The resulting 13 road network typologies are shown in Figure 3 and include angular streets (c0), irregular grid (c1), interrupted grid (c2), highly continuous streets (c3), dense core or parking lot (c4), disconnected (c5), regular grid (c6), curvy roads (c7), structured angular (c8), spacious winding roads (c9), scattered and orthogonal roads (c10), and two for which no clear typology was identified (c11 and c12). The GMM approach has the advantage of operating on unseen cities without requiring cluster re-specification, and handles overlapping cluster boundaries naturally. A detailed description of all metrics and their derivation is given in Post (2024) (the MSc thesis used as basis for this paper).

[1] `https://www.openstreetmap.org/`

[2] `https://overturemaps.org/`

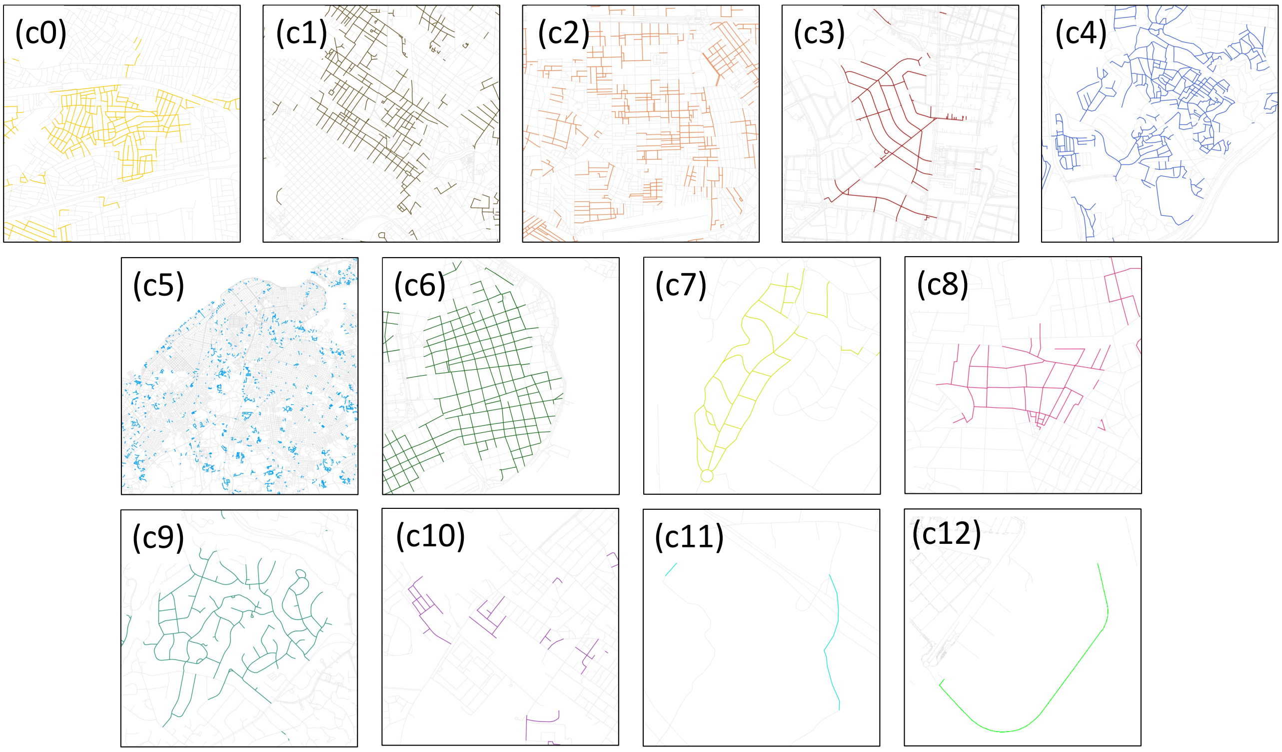

Figure 3. Representative examples of road network typologies for a selection of cities, from Cluster 0 (c0) to Cluster 12 (c12).

4.3 Building footprint typologies

For building footprints, the architectural and functional diversity of buildings makes it feasible to predefine a set of distinct typologies. Based on the classifications of Steiniger et al. (2008) and Wurm et al. (2016), refined through visual inspection of cities worldwide, we established nine typologies: detached housing, terraced housing, filled city block, perimeter city block, irregular blocks, apartment buildings, industrial buildings, complex buildings, and big commercial buildings. These are illustrated in Figure 4 for the city of Antwerp.

Figure 4. Building typology classification for the city of Antwerp, Belgium. Each colour corresponds to one of the nine typologies.

Each building was assigned a typology using a supervised gradient boosting classifier (Dorogush et al., 2018) trained on footprint attributes and neighbourhood context metrics: footprint area, shared wall ratio, land use category, covered area, GHSL height, and weighted neighbourhood aggregates of these features. Feature importance analysis shows that neighbourhood context dominates: 17 of the 20 most important features describe aggregated neighbourhood properties, confirming that building typology is determined primarily by surroundings rather than individual footprint geometry.

5. Encoding the typologies into a typology grid

5.1 The typology grid

The typology grid is a raster representation in which each cell stores a single typology value, defined as the dominant typology within that cell by intersection area. The default resolution is 100 × 100 m, which provides a balance between spatial detail and computational tractability. This encoding bridges the analysis and generation phases of the workflow: it is compact and comparable across cities, it is interpretable and adjustable by users, and it is compatible with patch-level landscape statistics (McGarigal and Marks, 1995), which provide quantitative descriptors of spatial pattern. A 3 × 3 mode kernel is applied as a post-processing step to reduce noise; across all 43 cities, this caused 17.5% of cells to change value, improving spatial coherence without loss of the overall typology pattern.

Figure 5 shows a typology grid encoding the building classification for the city of Antwerp. The various typologies are treated as globally universal and are therefore encoded independently for all cities, enabling direct cross-city comparison of urban character.

5.2 Typology templates and city comparison

In addition to the typology grid, each typology is described statistically in a typology template: an aggregation of numerical and categorical attributes across all instances of that typology across the full dataset. For example, a regular grid road typology template encodes a narrow distribution of intersection

https://doi.org/10.5194/isprs-annals-XII-4-W1-2026-277-2026 | © Author(s) 2026. CC BY 4.0 License.

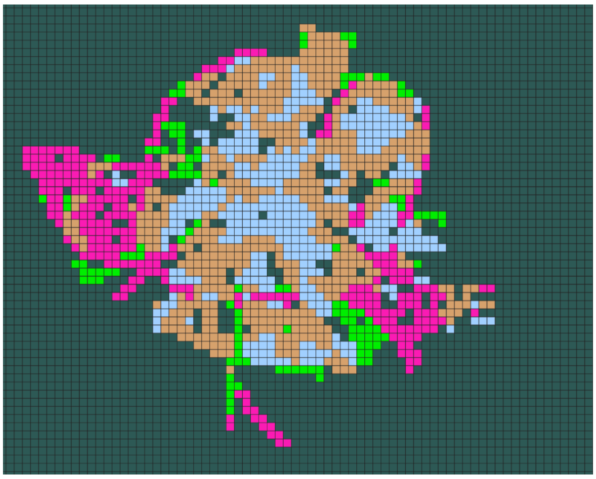

Figure 5. Typology grid encoding the building classification for the city of Antwerp, Belgium. Each colour represents one of the nine building typologies.

angles centred near 90 degrees, a regular distribution of segment lengths, and low curvature values. The templates carry cell-level, area-level, and city-level statistics, encoding both typical values and their variability. They serve as the primary source of parameters for all generation algorithms described in Section 6.

The typology grid representation enables direct quantitative comparison of cities. Figure 2 shows the global distribution of cities used in this study. The typology distributions reveal characteristic differences between cities: some are dominated by organic street patterns, others by regular grids; some have predominantly detached housing, others are dominated by city blocks or apartments.

6. Generation of plausible cities

To generate a new plausible city from the analysis encoded in the layers, the CityStack framework proceeds layer by layer from the bottom of the stack upward, with each layer constrained by the layers beneath it. Since different layers encode distinct aspects of urban form, this approach does not prescribe the generation algorithm for any given layer, making the framework modular: any algorithm that accepts a typology grid and associated constraints as input and produces geometrically valid output can be incorporated.

For each layer, the generation proceeds in two steps:

1. A new typology grid is generated that preserves the statistical properties of the template typology grid for the corresponding city.

2. Individual geometric elements are generated by algorithms that adapt their output according to the typology of the intersecting typology grid cell.

6.1 Typology grid generation via simulated annealing

We use simulated annealing (Aarts and van Laarhoven, 1989) to generate new typology grids. The algorithm initialises a grid with the same typology distribution as the template and iteratively proposes random swaps of two cells with different typology values. Each swap is accepted or rejected based on a temperature-dependent criterion, allowing the search to escape local optima. The objective function measures the similarity between the generated grid and the template using patch area, core area index, and shape index from landscape statistics (McGarigal and Marks, 1995), binned by distance to the city centre to preserve the characteristic centre-to-periphery gradient found in most cities.

Using the city of Breda (the Netherlands) at one-ninth resolution (300 m cells), the annealing process converges after approximately 1.3 million iterations (not shown here), reaching 95% similarity to the input typology grid. The generated grid reproduces the spatial organisation of typologies while placing them in a new geometric configuration, producing a statistically consistent city variant that is distinct from the original layout. At the typology-grid level, this preservation of typology distributions and distance-binned landscape statistics is our operational criterion for plausibility.

6.2 Road network generation

The road network forms the lowest layer in the current framework. We adapt the L-system approach of Parish and Müller (2001) to use local typology statistics as input parameters, replacing the manually crafted image maps of the original method. The algorithm grows the road network segment by segment; at each step, the length, angle, and branching probability of the next segment are sampled from the typology template statistics of the local typology grid cell. The road network is additionally constrained by terrain, water bodies, and previously generated segments (but this is future work as it is not yet implemented). This data-driven parameterisation allows the same algorithm to generate a dense orthogonal grid where the typology grid specifies a regular pattern, and an organic winding network where it specifies curved or angular typologies. Once generated, the road network is polygonised into blocks and insets to derive the buildable polygons used by the building layer. In this way, the geometry of each road segment constrains the frontage, access, and maximum area available to buildings in adjacent cells.

The generation process is illustrated in Figure 6, showing an initial road segment growing into a full road network constrained by the underlying typology grid. The approach is computationally efficient, generating 300,000 road segments in seconds.

6.3 Building footprint generation

The building footprint layer is generated using five algorithms, each mapped to one or more of the nine typologies, as detailed in Table 1. Method 1 recursively splits an inset city-block polygon using object-aligned bounding boxes until a minimum area threshold is reached, with split ratios and thresholds drawn from typology statistics. Method 2 samples real building footprints from the template city and distributes them along block roads according to typology statistics (used for detached housing, apartments, industrial, complex, and big commercial buildings, where footprint shape diversity is highest). Method 3 adapts Method 1 by randomly discarding a fraction of the resulting buildings and irregularly insetting the remainder, producing irregular block patterns. Method 4 creates a courtyard block using nested insets of the original block polygon. Method 5 sweeps a row of identical house units at a fixed interval along selected street segments to produce terraced housing. All five methods operate inside the buildable polygons defined by the generated road network. In Method 2 in particular, sampled

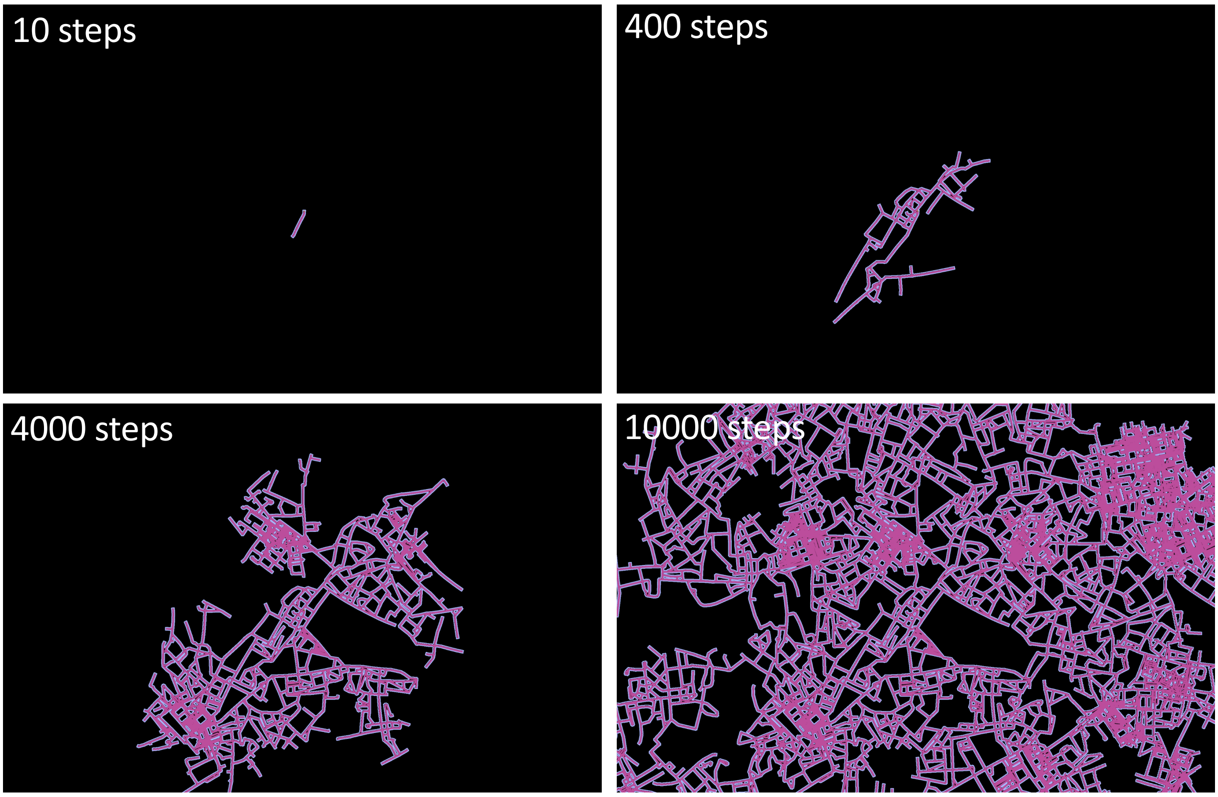

Figure 6. Street generation process using the adapted L-systems approach.

footprints are first aligned to the local street frontage and are rejected whenever they intersect the road network, the block boundary, or a previously accepted footprint; sampling then continues until the typology-specific density target is reached. This prevents overlap while allowing complex real-world footprint shapes to be reused.

	Method				
Typology	1	2	3	4	5
Detached Housing		×			
Terraced Housing					×
Filled City Blocks	×				
Perimeter City Blocks				×	
Irregular City Blocks			×		
Apartments		×			
Industrial Buildings		×			
Complex Buildings		×			
Big Commercial Buildings		×			

Table 1. Building footprint generation method used for each building typology class.

A visual comparison of all five building generation methods is provided in Figure 7.

7. Discussion and future work

This paper has demonstrated that urban morphological patterns can be systematically captured, encoded, and used to generate plausible city variations through a programmatic approach. Rather than simply recreating existing cities, our framework enables a data-driven morphological characterisation and supports the generation of new urban environments that maintain the morphological character of reference cities while adapting to different contexts. Within the scope of this paper, plausibility is interpreted narrowly: generated outputs should preserve the typology distributions, spatial gradients, and local morphological statistics observed in a reference city, while differing in their exact geometry. Analysis shows that unsupervised clustering can derive a global set of road pattern typologies, that supervised classification can assign building typologies, and that the typology grid provides a compact, comparable, and adjustable encoding linking analysis and generation.

This evidence should nevertheless be interpreted with care. The current evaluation is strongest at the typology and typology-grid levels: we show interpretable typologies, transferable encodings, and simulated annealing results that preserve target landscape statistics. We do not yet provide a full task-based evaluation, nor a quantitative comparison between final generated geometries and real cities across all morphology measures. Likewise, the supervised building classification currently relies on seven ground-truth cities and visual verification elsewhere, so a fuller accuracy assessment remains necessary.

The framework also inherits limitations from the source data. OpenStreetMap road data and derived urban attributes remain spatially uneven in completeness and semantic consistency (Biljecki et al., 2023), and although Overture improves building coverage, under-represented urban contexts may still be assigned to the nearest available typology rather than a truly local one. In addition, the 100 m typology grid resolution is a compromise: finer grids retain more local variation but increase noise and computation, while coarser grids improve robustness at the cost of morphological detail.

Computational scalability is another practical limitation. The current analysis pipeline processes large cities largely monolithically; for Paris, this required up to 100 GB of RAM. Future work should therefore investigate tiled or out-of-core pro-

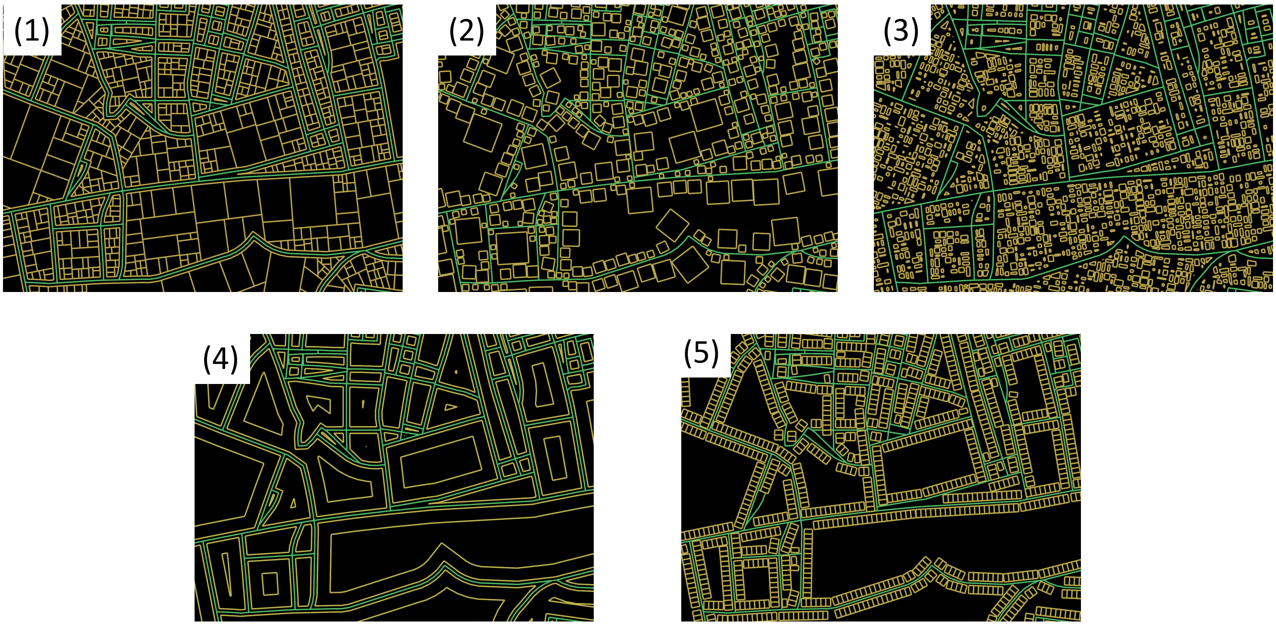

Figure 7. Five building footprint generation methods within the CityStack framework; green lines indicate street network segments and yellow lines indicate building footprints.

cessing, caching of intermediate metrics, and more aggressive parallelisation of neighbourhood aggregation so that the framework remains usable on standard research hardware.

The CityStack framework is explicitly designed for extensibility. Future research can add more layers to the stack, working upward from the currently modelled road and building layers: the *major roads* and *highways* layers can use the same analysis approach as the minor roads layer; the *vegetation and street furniture* layer can be added above buildings; and *land use* and *points of interest* layers already have relatively unambiguous typologies that can be encoded directly from OpenStreetMap. The *natural context* layer (terrain and water), although not used for analysis in this work, is already present in the dataset for future use.

The quality of road pattern typologies can be improved by refining the clustering methodology, trying alternative metrics or aggregation strategies, or attempting supervised learning on an expert-created global set of typologies. The integration of deep learning, as demonstrated by Wang et al. (2024) for combined urban tissue classification, is a promising direction. For typology grid generation, machine learning models such as convolutional auto-encoders, autoregressive transformers, or diffusion models conditioned on distance to the city centre could provide an alternative to simulated annealing while preserving centre-to-periphery gradients.

The most important avenue for future work is the extension of CityStack to full 3D city models, for instance, stored in the CityJSON (Ledoux et al., 2019) or CityGML-XML formats (Gröger and Plümer, 2012; OGC, 2023). Our current pipeline already provides the building footprint geometry and approximate height from the GHSL dataset. The next natural step is therefore to extrude building footprints to their assigned height and assign roof shapes based on typology statistics, producing models conforming to the CityGML standard at LoD 1 or LoD 2 (Biljecki et al., 2016), directly usable in downstream applications such as solar irradiation estimation, noise propagation modelling, and CFD wind simulation (Biljecki et al., 2015). Different building typologies will require different 3D generation strategies, and the typology templates already encode much of the statistical information needed to parameterise them. Beyond geometry, semantic attributes such as land use, function, and demographic variables could also be assigned based on typology and position within the stack, enabling synthetic cities for simulation and privacy-sensitive applications.

Acknowledgments

AP would like to acknowledge that this work was supported by the RefMap. This project has received funding from the European Union's HORIZON Research and Innovation Programme under Grant Agreement number 101096698. OP, HL, and AP would like to thank Claudiu Forgaci for their insights and discussions during the preparation of this manuscript. All authors acknowledge the use of a Large Language Model (LLM) in the preparation of this manuscript. The authors have thoroughly reviewed both the manuscript and the associated code following the use of the LLM and accept full responsibility for the accuracy and integrity of the content presented.

References

Aarts, E. H. L., van Laarhoven, P. J. M., 1989. Simulated annealing: An introduction. *Statistica Neerlandica*, 43(1), 31–52.

Aliaga, D. G., Vanegas, C. A., Benes, B., 2008. Interactive example-based urban layout synthesis. *ACM SIGGRAPH Asia 2008 papers*, SIGGRAPH '08, ACM, 1–10.

Badhrudeen, M., Derrible, S., Verma, T., Kermanshah, A., Furno, A., 2022. A geometric classification of world urban road networks. *Urban Science*, 6(1), 11.

Biljecki, F., Chow, Y. S., Lee, K., 2023. Quality of crowdsourced geospatial building information: A global assessment of OpenStreetMap attributes. *Building and Environment*, 237, 110295.

Biljecki, F., Ledoux, H., Stoter, J., 2016. Generation of multi-LoD 3D city model in CityGML with the procedural modelling engine Random3DCity. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, IV-4-W1, 51–59.

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

Chen, G., Esch, G., Wonka, P., Müller, P., Zhang, E., 2008. Interactive procedural street modeling. *ACM SIGGRAPH 2008 papers*, SIGGRAPH '08, ACM, 1–10.

Dorogush, A. V., Ershov, V., Gulin, A., 2018. CatBoost: Gradient boosting with categorical features support. `https://arxiv.org/abs/1810.11363`.

Fleischmann, M., Feliciotti, A., Romice, O., Porta, S., 2022. Methodological foundation of a numerical taxonomy of urban form. *Environment and Planning B: Urban Analytics and City Science*, 49(4), 1283–1299.

García-Sánchez, C., Van Beeck, J., Gorlé, C., 2018. Predictive large eddy simulations for urban flows: Challenges and opportunities. *Building and Environment*, 139, 146–156.

Groenewegen, S. A., Smelik, R. M., 2009. Procedural city layout generation based on urban land use models. *Proceedings of Eurographics 2009: Short Papers*.

Gröger, G., Plümer, L., 2012. CityGML—Interoperable semantic 3D city models. *ISPRS Journal of Photogrammetry and Remote Sensing*, 71, 12–33.

Han, W., Wen, C., Chok, L., Tan, Y. L., Chan, S. L., Zhao, H., Feng, C., 2024. Autoencoding tree for city generation and applications. *ISPRS Journal of Photogrammetry and Remote Sensing*, 208, 176–189.

Hartmann, S., Weinmann, M., Wessel, R., Klein, R., 2017. StreetGAN: Towards road network synthesis with generative adversarial networks.

Hobeika, N., van Rijssel, L., Prusti, M., Dinklo, C., Giannelli, D., Dukai, B., Kok, A., van Loon, R., Nota, R., Stoter, J., 2024. Automated noise modelling using a triangulated terrain model. *Geo-Spatial Information Science*, 27(6), 1893–1913.

Kaden, R., Kolbe, T. H., 2013. City-wide total energy demand estimation of buildings using semantic 3D city models and statistical data. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 2, 163–171.

Kim, J.-S., Kavak, H., Crooks, A., 2018. Procedural city generation beyond game development. *SIGSPATIAL Special*, 10(2), 34–41.

Kumar, K., Ledoux, H., Commandeur, T. J. F., Stoter, J. E., 2017. Modelling urban noise in CityGML ADE: Case of the Netherlands. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 4, 73–81.

Ledoux, H., Arroyo Ohori, K., Kumar, K., Dukai, B., Labetski, A., Vitalis, S., 2019. CityJSON: A compact and easy-to-use encoding of the CityGML data model. *Open Geospatial Data, Software and Standards*, 4(1).

León-Sánchez, C., Giannelli, D., Agugiaro, G., Stoter, J., 2021. Testing the new 3D bag dataset for energy demand estimation of residential buildings. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 46, 69–76.

Louf, R., Barthélemy, M., 2014. A typology of street patterns. *Journal of the Royal Society Interface*, 11, 20140924.

Marshall, S., 2004. *Streets & patterns*. 1st ed edn, Taylor & Francis.

McGarigal, K., Marks, B. J., 1995. FRAGSTATS: Spatial pattern analysis program for quantifying landscape structure. Technical Report PNW-GTR-351, U.S. Department of Agriculture, Forest Service, Pacific Northwest Research Station.

OGC, 2023. OGC city geography markup language (CityGML) part 2: GML encoding standard. Open Geospatial Consortium Inc. Document 21-006r2, version 3.0.

Oliveira, V., 2020. *Urban morphology*. Springer.

Pađen, I., Peters, R., García-Sánchez, C., Ledoux, H., 2024. Automatic high-detailed building reconstruction workflow for urban microscale simulations. *Building and Environment*, 265, 111978.

Parish, Y. I. H., Müller, P., 2001. Procedural modeling of cities. *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, SIGGRAPH01, ACM, 301–308.

Patil, A., García-Sánchez, C., 2025. Quantifying the impact of urban geometric detail for urban air mobility risk forecasting. *Sustainable Cities and Society*, 106750.

Pesaresi, M., 2023. GHS-BUILT-H R2023A - GHS building height, derived from AW3D30, SRTM30, and Sentinel2 composite (2018). European Commission, Joint Research Centre (JRC).

Post, O. J., 2024. The City Stack: A morphology-based city analysis and generation framework. Master's thesis, Delft University of Technology.

Reynolds, D., 2009. Gaussian mixture models. S. Z. Li, A. Jain (eds), *Encyclopedia of biometrics*, Springer US, 659–663.

Sarkar, S., Cottineau-Mugadza, C., Wolf, L. J., 2024. Spatial inequalities and cities: A review. *Environment and Planning B: Urban Analytics and City Science*, 51(7), 1391–1407.

Steiniger, S., Lange, T., Burghardt, D., Weibel, R., 2008. An approach for the classification of urban building structures based on discriminant analysis techniques. *Transactions in GIS*, 12(1), 31–59.

Vargiemezis, T., Gorlé, C., 2025. From large-eddy simulations to deep learning: A U-net model for fast urban canopy flow predictions. *Sustainable Cities and Society*, 107005.

Wang, J., Huang, W., Biljecki, F., 2024. Learning visual features from figure-ground maps for urban morphology discovery. *Computers, Environment and Urban Systems*, 109, 102076.

Wurm, M., Schmitt, A., Taubenböck, H., 2016. Building types' classification using shape-based features and linear discriminant functions. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 9(5), 1901–1912.