The Cartographer in the Loop: An Open-Source LiDAR-to-3D Workflow for Heat-Resilient Streetscape Planning

Mahsa Saharkhiz [1*], Siavash Karimi [2] and Mohamed Aly Etman [2]

[1] Department of Geography, College of Arts and Sciences, University at Buffalo, NY
[2] Department of Architecture, School of Architecture and Planning, University at Buffalo, NY
[*] mahsasah@buffalo.edu

Keywords: airborne LiDAR; 3D city models; individual tree detection; urban heat island; outdoor thermal comfort; open-source GIS; QGIS; human-in-the-loop geoprocessing

Research question and motivation

Extreme heat is among the deadliest weather hazards in temperate cities, and the mortality burden of non-optimal temperature is measurable across countries and climates (Gasparrini et al., 2015). Effective streetscape interventions, including street trees, shade structures, and ventilation corridors must respond to local three-dimensional conditions. The microclimate and outdoor thermal-comfort models used to evaluate these interventions require geometrically explicit 3D inputs, including tree locations and heights and every building footprints and roof heights (Lindberg et al., 2018).

Large metropolitan areas often possess such models. However, small and mid-sized municipalities often lack the resources needed to produce them. Proprietary photogrammetry, dedicated LiDAR acquisition, and manual modeling can place neighborhood-scale heat analysis beyond the reach of local planning offices.

Our research question is practical: Can the public airborne LiDAR available through the USGS 3D Elevation Program (3DEP), processed entirely in free and open-source software by a non-specialist office, yield streetscape geometry trustworthy enough for heat-resilience design, and what division of labor between automated safeguards and analyst judgement does that demand? If so, where must analyst judgment remain with the workflow to identify and correct predictable automated errors?

The workflow responds directly to the conference theme by treating human review not as a failure of automation, but as a designed component of reliable municipal geoprocessing.

Background

Canopy-height-model (CHM) individual tree detection by local-maxima search is mature, transparent, and fast, but its weaknesses are amplified in cities, where canopy is heterogeneous and intermingled with poles, signs, and rooftops; recent multi-city evaluation shows that even state-of-the-art large-area canopy-height products systematically misclassify urban canopy and mis-estimate tree heights, so locally derived airborne LiDAR remains necessary (Dong et al., 2025). Existing open routes to simulation-ready 3D city models stop short of what streetscape design needs: voxel frameworks assemble pre-computed global height products into gridded semantic models rather than deriving per-object geometry from local returns (Fujiwara et al., 2026), while per-building LiDAR pipelines typically depend on the supervised point-cloud classification that public tiles cannot reliably support (Park & Guldmann, 2019). On the footprint side, OpenStreetMap outlines are geometrically good where present but spatially uneven in completeness (Fan et al., 2014; Herfort et al., 2023). 3DEP itself is accurate and consistent across acquisitions (Stoker & Miller, 2022), yet it cannot be used off the shelf: classification labels are frequently untrustworthy, water voids interpolate into phantom canopy, and hard returns from street furniture masquerade as trees. The missing piece is a documented, low-barrier bridge from this flawed but abundant public resource to per-object 3D vector geometry and a thermal-comfort analysis.

Methods

The workflow runs entirely in QGIS with two plugins and short console scripts, in five stages: (1) terrain, surface, and canopy height models; (2) building footprints and roof heights; (3) individual tree detection; (4) cleaning and 3D scene assembly; and (5) a comparative thermal-comfort analysis. Every raster is built at a fixed 0.25 m cell size on an identical grid, and a single projected CRS is enforced per site, with arriving point clouds either relabeled (when coordinates are native but mislabeled) or reprojected exactly once at the raw-point stage (when the delivery is genuinely foreign, such as Web Mercator). The surface model is built with an inclusion-only filter—first returns from only the classes a tree can legitimately occupy—which sidesteps, rather than repairs, the unreliable classification of public tiles. Trees are detected by constrained local-maxima search on a building-masked CHM, filtered by a plausible height band whose upper cap is calibrated ecologically and verified by mapping the removed candidates, and by a canopy ‘crown test’ that rejects poles, masts, and signs. Footprints come from OpenStreetMap, on-screen digitizing against a LiDAR hill shade, or georeferenced authoritative CAD.

Figure 1: The open-source QGIS workflow and its analyst-in-the-loop checkpoints. Canopy-height-model detection of trees and buildings (top left) is cross-checked by eye against date-matched aerial imagery (remaining panels), confirming that detections fall on canopy rather than roofs or pavement and exposing the recurring hard cases — canopy overhanging low roofs, and small buildings partly hidden beneath tree crowns.

Roof heights illustrate the analyst-in-the-loop design most sharply. The per-footprint CHM maximum—the obvious extrusion height—is a single cell, and ranking footprints by it exposes three recurring patterns: genuinely tall buildings; real but slender rooftop structures (steeples, stacks, masts); and overhanging street-tree canopy inflating one- and two-story buildings to tree height. The workflow therefore extrudes on a robust statistic—the median of the CHM inside each footprint shrunk inward by 1.5 m, with a fallback chain for small footprints and a per-building provenance flag—in line with the robust per-building statistics used for city-scale morphology extraction from aerial LiDAR (Bonczak & Kontokosta, 2019), while the ranked audit remains a designed human checkpoint: automation proposes every height, the analyst inspects the extremes, and the one case no zonal statistic can resolve (a small structure entirely beneath canopy) is flagged for individual review rather than silently accepted. At North Tonawanda such canopy-swallowed buildings were frequent enough in tree-lined blocks that this manual pass over the largest corrected heights, within the area selected for thermal analysis, is treated as required: the statistic repaired thousands of heights automatically, and the analyst repaired the residual handful. The assembled scene drives a comparative thermal-comfort index read as spatial pattern and scenario difference rather than absolute prediction.

Findings

The workflow was exercised on three New York sites spanning two IECC climate zones and two UTM zones. At Saranac Lake, an Adirondack village, it recovered 620 buildings and 8,931 height-tagged trees; the CHM maximum of 44.26 m is a plausible tallest-tree value for the region. At North Tonawanda, a city of about 30,600, the same procedure scaled without modification to 13,812 buildings and 45,668 trees—several times the City's inventory of 12,927 public street and park trees (City of North Tonawanda, 2022b), the excess being exactly the private and backyard canopy a right-of-way inventory cannot see. The roof-height audit there quantified why robust statistics matter in a tree-lined city: of 13,127 footprints within LiDAR coverage (from 23,263 candidates across the acquisition window), 3,546—27%—carried a zonal maximum more than 5 m above their buffered-median height. The fifteen largest maxima resolved into exactly the anticipated patterns: two genuinely tall buildings (a 28.9 m residential tower; a 27.1 m industrial block), a church steeple (31.8 m over an 11.7 m nave), two stepped industrial complexes, and ten low-rise buildings whose 29–38 m maxima were overhanging street trees above 4.9–9.5 m roofs. At the University at Buffalo North Campus, detection was validated two ways, with a cautionary result: against a 2018 CAD site drawing—stale and contaminated with non-tree symbols—F1 was 0.38, but against verified source-epoch imagery on ten one-hectare tiles the same detections scored precision 0.74, recall 0.89, F1 = 0.81. The gap measures the reference, not the detector—a warning for anyone benchmarking detection against class-based municipal inventories. Geometric accuracy is sub-meter where independently checkable: the tallest campus building is recovered at 45.37 m against 45.7 m in Google Earth, and six repeated-design buildings agree within 0.28 m.

Significance and feedback sought

The contribution is a documented, auditable, fully open-source bridge from free national LiDAR to design-ready 3D streetscape geometry—one whose reliability comes not from removing the human but from placing designed checkpoints where automation is known to fail: the top-of-table roof audit, the mapped dropped set of the height cap, the identify-tool spot checks after every raster step. That division of labor is, we argue, the practical shape of 'the cartographer in the loop' for municipal geoprocessing. The work is in progress toward a journal submission; at CaGIS we particularly seek feedback on the design of a calibrated height-error field campaign, on coupling the comparative comfort index to SOLWEIG/UMEP for physical cross-checking, and on cartographic strategies for communicating per-object uncertainty to planning audiences.

References

Bonczak, B., & Kontokosta, C. E. (2019). Large-scale parameterization of 3D building morphology in complex urban landscapes using aerial LiDAR and city administrative data. Computers, Environment and Urban Systems, 73, 126–142. https://doi.org/10.1016/j.compenvurbsys.2018.09.004

City of North Tonawanda. (2022). North Tonawanda's community forest management plan (Prepared by PlanIT Geo, Inc.). https://planitgeo.com/library/community-forest-management-plan-north-tonawanda-ny/

Dong, H., Xu, Z., Wu, J., Lan, T., Wang, L., Shao, G., & Tang, L. (2025). High-precision airborne LiDAR remains essential for urban forestry: Revealing the limitations of recent large-scale canopy height products in urban contexts. International Journal of Applied Earth Observation and Geoinformation, 143, 104791. https://doi.org/10.1016/j.jag.2025.104791

Fan, H., Zipf, A., Fu, Q., & Neis, P. (2014). Quality assessment for building footprints data on OpenStreetMap. International Journal of Geographical Information Science, 28(4), 700–719. https://doi.org/10.1080/13658816.2013.867495

Fujiwara, K., Tsurumi, R., Kiyono, T., Fan, Z., Liang, X., Lei, B., Yap, W., Ito, K., & Biljecki, F. (2026). VoxCity: A seamless framework for open geospatial data integration, grid-based semantic 3D city model generation, and urban environment simulation. Computers, Environment and Urban Systems, 123, 102366. https://doi.org/10.1016/j.compenvurbsys.2025.102366

Gasparrini, A., Guo, Y., Hashizume, M., Lavigne, E., Zanobetti, A., Schwartz, J., Tobias, A., Tong, S., Rocklöv, J., Forsberg, B., Leone, M., De Sario, M., Bell, M. L., Guo, Y.-L. L., Wu, C.-F., Kan, H., Yi, S.-M., de Sousa Zanotti Stagliorio Coelho, M., Saldiva, P. H. N., … Armstrong, B. (2015). Mortality risk attributable to high and low ambient temperature: A multicountry observational study. The Lancet, 386(9991), 369–375. https://doi.org/10.1016/S0140-6736(14)62114-0

Herfort, B., Lautenbach, S., Porto de Albuquerque, J., Anderson, J., & Zipf, A. (2023). A spatio-temporal analysis investigating completeness and inequalities of global urban building data in OpenStreetMap. Nature Communications, 14, 3985. https://doi.org/10.1038/s41467-023-39698-6

Lindberg, F., Grimmond, C. S. B., Gabey, A., Huang, B., Kent, C. W., Sun, T., Theeuwes, N. E., Järvi, L., Ward, H. C., Capel-Timms, I., Chang, Y., Jonsson, P., Krave, N., Liu, D., Meyer, D., Olofson, K. F. G., Tan, J., Wästberg, D., Xue, L., & Zhang, Z. (2018). Urban Multi-scale Environmental Predictor (UMEP): An integrated tool for city-based climate services. Environmental Modelling & Software, 99, 70–87. https://doi.org/10.1016/j.envsoft.2017.09.020

Park, Y., & Guldmann, J.-M. (2019). Creating 3D city models with building footprints and LIDAR point cloud classification: A machine learning approach. Computers, Environment and Urban Systems, 75, 76–89. https://doi.org/10.1016/j.compenvurbsys.2019.01.004

Stoker, J., & Miller, B. (2022). The accuracy and consistency of 3D Elevation Program data: A systematic analysis. Remote Sensing, 14(4), 940. https://doi.org/10.3390/rs14040940