Geospatial Data Pipelines for Urban Digital Twin Applications

Joanna Hobbins, Melisa Kopan, Nick Giannias
CAE
Canada
joanna.hobbins@cae.com, melisa.kopan@cae.com, nick.giannias@cae.com

Simon Merrick
CAE UK
United Kingdom
simon.merrick@cae.com

Sean Lilley, Shehzan Mohammed
Cesium
Philadelphia, PA
sean@cesium.com, shehzan@cesium.com

Ralph Coleman
Bluesky International Ltd.
United Kingdom
ralph.coleman@bluesky-world.com

ABSTRACT

Digital twins have become an important means of offering actionable insight into our world. The existing and potential applications are extremely broad in scope and drive a need for rapid creation of large scale digital replications of the real world. We organize our digital twin application around four conceptual categories or "layers": Physical, Human, Cognitive, and Resource. The Physical layer represents the world "as-built" and can include imagery, terrain and 3D models. The Human layer includes the population and their location, either through simulated pattern of life or real time geolocation feeds. The Cognitive layer encapsulates human awareness, behaviour, decision making, social connections and communication. The Resource layer contains the assets available for planning and simulation. This paper reports on the design, implementation and integration of the geospatial data used to create the Physical and Human layers of one such digital twin, the city of London. Our project made extensive use of both cloud technology and open standards from the Open Geospatial Consortium (OGC). We faced multiple challenges related to: (i) the conversion of large amounts of photogrammetry data for streaming using the OGC 3D Tiles standard (ii) the sourcing and processing of sufficiently detailed data for pattern of life and power network simulations and (iii) the performance challenges associated with using OGC WMS/WFS for real-time simulation. Viewed individually, none of the challenges we faced are unique in the geospatial and simulation communities. However, the workflows detailed in this paper have laid the foundation for patterns and approaches, including data variation points, to provide a method baseline to support the wider adoption of urban digital twin development.

ABOUT THE AUTHORS

Joanna Hobbins is a Subject Matter Expert in Geographic Information Systems (GIS) at CAE with over 20 years of experience and a focus in automating geospatial workflows. She holds a B.Sc. in Earth Science with a specialty in GIS from McGill University.

Melisa Kopan is Team Coordination and Geographic Information Systems (GIS) Lead at CAE where a background in Synthetic Environment Visual Database Design has been developed over her two years at the company. She holds a B.Sc. in Environmental Geoscience and Geomatics and an M.Sc. in Geography.

Nick Giannias is a Senior Technical Fellow working for the Advanced Technology and Innovation group at CAE and the Chief Architect for CAE's simulator products. With over 30 years of experience, he specializes in software development and real-time simulations for concept development, analysis, mission rehearsal, and training applications. Nick holds a bachelor's degree in Mechanical Engineering (Hons.) specializing in Aerospace from McGill University.

Simon Merrick is a Software Delivery Lead at CAE with 20 years of experience in software development spanning the training & simulation, defence, and GIS sectors. He currently leads the CAE team developing digital twin technology and holds a B.Sc (Hons) in Computer Science.

Sean Lilley leads the Cesium 3D team in developing software and open specifications for 3D geospatial. He is the co-editor of 3D Tiles and a contributor to glTF extensions. He holds a B.S.E in Digital Media Design from the University of Pennsylvania.

Shehzan Mohammed leads open source and outward-facing 3D engineering at Cesium. Shehzan is responsible for CesiumJS, Cesium for Unreal, and other runtime engines, along with growing the 3D geospatial ecosystem and its adoption. Shehzan holds a Master degree in Computer Graphics and Game Technology from University of Pennsylvania, where he also currently teaches GPU Programming and Architecture**.**

Ralph Coleman is Sales Director and Board Member at Bluesky International Ltd a leading aerial survey and geospatial information company headquartered in the UK. Ralph has a Geography degree and has been involved in the geospatial industry for over 20 years. Ralph has a passion for 3D data and in particular working with Bluesky's customers and partners to develop the concept of the Smart City and its societal benefits.

Geospatial Data Pipelines for Urban Digital Twin Applications

Joanna Hobbins, Melisa Kopan, Nick Giannias
CAE
Canada
joanna.hobbins@cae.com, melisa.kopan@cae.com, nick.giannias@cae.com

Simon Merrick
CAE UK
United Kingdom
simon.merrick@cae.com

Sean Lilley, Shehzan Mohammed
Cesium
Philadelphia, PA
sean@cesium.com, shehzan@cesium.com

Ralph Coleman
Bluesky International Ltd.
United Kingdom
ralph.coleman@bluesky-world.com

INTRODUCTION

Our first urban digital twin is of Greater London in the United Kingdom which covers an area over 600 square miles and a population of over eight million people. Our selected use case for this digital twin is one of "national resilience". We specifically elected to explore this potential application of urban digital twins with respect to a cyberattack on the power supply infrastructure within the city. The aim was to establish means and mechanisms for the provision of a real time understanding of:

- The state of critical national infrastructure impacted by such an event
- The second order impacts on the human population in terms of their movement and sentiment
- The location and status of relevant "resources", in this context specifically the police, fire, and ambulance services including both buildings and vehicles

Altogether these layers allow decisions makers to prioritize allotment of resources, particularly first responders, to higher risk situations.

Digital Twins

At its heart, a "digital twin" is a digital (typically virtual) replica of a real world subject. The subject may be singular or plural and may range from physical devices to people, processes or systems. There is no consensus as to the exact requirements or even definition of a digital twin with most authors providing their own take (Zheng, 2018). Digital twins, as referred to in this paper, serve as a foundational reference for the topics discussed herein; and enable us to share our views and learning experience in this area. The definition provided is not intended to be a normative or even necessarily a complete one.

The subjects of today's digital twins tend to be physical in nature, although it is conceivable that they need not be. As long as the subject is a real-world system, it could be an abstraction rather than a physical entity provided that it meets the defining characteristics of a digital twin as listed below.

The purpose of such replicas varies depending on the needs of the user, although one canonical theme is to support decision making. Such decisions can range from decisions based on supervising/observing the current state of the real world subject, through predicting its future state based on data collected from the digital twin, and potentially including modeled behaviour and/or external empirical data collections.

For these reasons, the following hold as key defining characteristics of a digital twin for this paper's purpose:

- A digital twin must replicate a real world subject.
- The replication must provide access to real, current, data from that subject.
- It is not necessary, perhaps even desirable, to replicate every possible aspect of the subject.

- The aspects (as well as the fidelity) of the replicated data may vary based on the specific purpose and use case of the digital twin.
- It is necessary that replicated aspects serve a real world use case or purpose, and that the selected aspects reflect the true state of the real world subject being twinned – whether that be the current or historical state.
- The replication is not just a “visualisation” of the subject but provides appropriate access to the relevant data “behind the scenes” in a manner that enables an understanding of replicated aspects.

The digital revolution and the global spread of the internet have given rise to such phenomena as "Data-Driven City" and "Smart City". A smart city can be defined as the adoption and application of mobile computing systems through practical data management networks amongst all components and layers of the city itself (Kirimtat et al, 2020). An urban digital twin is a replica of the physical built environment in an urban area such as a city. This replica includes aspects such as the city’s buildings, transport networks and infrastructure. It is reasonable to draw the link between smart cities and digital twins, and to anticipate that as each of these fields evolve, in adoption and sophistication, they will symbiotically benefit the other, enabling further enhancements and capabilities as a result.

Conceptual Layers

Whereas (Ikanov et al, 2020) discusses a framework for a city digital twin comprising of multiple digital twin elements feeding a principal digital twin, we organize our digital twin application around four conceptual categories or “layers”: Physical, Human, Cognitive and Resource. This model has been adopted in order to explore the following premise: the integrated replication of these particular aspects of a real-world operating theater ought to provide actionable insight that supports better understanding and decision making within our increasingly complex world. Further to this, we postulate that modern, digital age, technologies now enable the creation and real world use of such digital twins – providing a potentially transformational opportunity to governments and militaries in the way that they organize and operate. The separation of these four layers enables the encapsulation of different data, models and applications into distinct groupings while providing the opportunity to reason over, and computationally analyze the relationships and correlations between the data in each layer. Furthermore, we are able to leverage the layers separately in the different web and gaming engine viewports, optimizing them for both display and analysis.

For example, the **Physical layer** is perhaps the most intuitive to grasp as it comprises the replication of the physical environment including imagery, terrain, buildings, and infrastructure. It provides the canvas upon which the other layers can be overlaid and intuitively links the user to the real world context.

The **Human layer** covers the general population, including individuals’ identifying attributes and location within the real world. This enables the user and the system to reason over the demographics and movement of populations at scale. It also holds the potential to support automated reasoning over crowd density and pedestrian/traffic flows with respect to the physical environment as well as the population’s contact networks and digital footprints.

The **Cognitive layer** builds upon the Human layer, adding the potential to support the understanding, analysis and visualization of the digital communications, and associated sentiment of individuals and groups within the population.

The separation of the **Resource layer** is key to the user’s utility of the digital twin. It encapsulates the resources available to the user, whether or not they conceptually span the other layers. The user can select to view or hide the available data feed and analytics from the first three layers, and on top of this context may “drape” their view of the current resource laydown in order to understand the situation in context and/or to define and assess different potential courses of action and events.

TECHNOLOGY

A guiding principle of our digital twin design is the extensive use of open standards and open source software. To that end, we decided on the Open Geospatial Consortium (OGC) 3D Tiles format for 3D data (OGC, 2019) and the OGC Web Map Service (WMS) (OGC, 2014) and Web Feature Service (WFS) (OGC, 2006) protocols for raster and vector data. Our open source software stack included CesiumJS and Cesium for Unreal for visualization and Geoserver for OGC web services.

3D Tiles

3D Tiles is an OGC Community Standard for streaming massive 3D geospatial content across desktop, web, and mobile applications. Source geospatial data, such as massive 3D models derived from photogrammetry, is often terabytes or larger in size and is impractical to stream over a network or fit in system memory. To address this, 3D Tiles defines a spatial data structure that enables Hierarchical Level of Detail (HLOD) so that only visible tiles are streamed - and only those tiles which are most important for a given 3D view.

3D Tiles is also designed for runtime efficiency. Tiles are built on glTF (Khronos, 2021), the runtime asset format which maps to modern graphics APIs and includes extensions for mesh and texture compression such as Draco, Meshopt, and KTX2/BasisU. Metadata can be stored at multiple granularities within the tileset. Per-feature metadata such as building heights, traffic sign types, and vehicle information can drive simulations and actor behaviors. After source data has been processed into 3D Tiles, it can be streamed to any runtime engine that supports the 3D Tiles standard. The two engines used in this digital twin are CesiumJS and Cesium for Unreal.

CesiumJS is an open source JavaScript / WebGL engine for creating 3D globes on the web. It provides an extensive API for interacting with 3D Tiles including picking, styling, and querying metadata. Cesium for Unreal is an open source plugin that brings the same high-accuracy full-scale WGS84 globe to Unreal Engine. Cesium for Unreal is tightly integrated with Unreal Engine making it possible to visualize and interact with real-world content in the Unreal editor and at runtime. The plugin also has support for Unreal Engine physics, collisions, character interaction, and landscaping tools. Under the hood it uses Cesium Native as the core 3D Tiles streaming engine.

Geoserver

Geoserver is an Open Source Geospatial Foundation[1] project that serves geospatial data through OGC protocols. Of particular interest for our digital twin is its ability to connect to a PostGIS database and serve the data through both WMS and WFS protocols.

To increase response time GeoWebCache is used to pre-generate the WMS tiles to zoom level 15 (approximately 5m ground resolution) with the remaining levels generated on the fly when requested by a client. This compromise reduces the storage costs of the higher zoom levels, many of which will never or extremely rarely be accessed while building a cache of those that are. This threshold becomes increasingly important when the digital twin increased from city wide to province or even country wide systems. While a web client will predominantly use the WMS protocol to retrieve and render the data, the simulation component requires vector data and will use the WFS protocol to retrieve the data in CSV or GeoJSON formats.

DATA

Physical Layer

The principal component of the Physical layer is high resolution imagery. It provides a backdrop that a user readily recognizes and provides easy correlation between the digital twin and the real world. Traditionally 3D visualizations have the imagery draped over a terrain mesh and models of 3D structures like buildings and sometimes trees placed on it. Producing geo-specific 3D models that correspond to the actual building is time and labour expensive, the geometry and features are often simplified, and smaller structures may not be digitized at all. In more recent years region-scale LIDAR photogrammetry models are replacing these separate layers with a single high resolution surface mesh. See Figure 1 for an example of a building captured through photogrammetric scans vs traditional OGC CityGML LOD2 modeling.

The benefits of photogrammetry include fast capture and processing times to produce a textured 3D mesh output and complete correlation. The downside is the lack of feature specific attribution and no separation between terrain and

[1] Open Source Geospatial Foundation https://www.osgeo.org/

features. The London Digital Twin uses MetroVista[2] photogrammetry derived from 5cm aerial photography and 16 points per meter LIDAR (Figure 2).

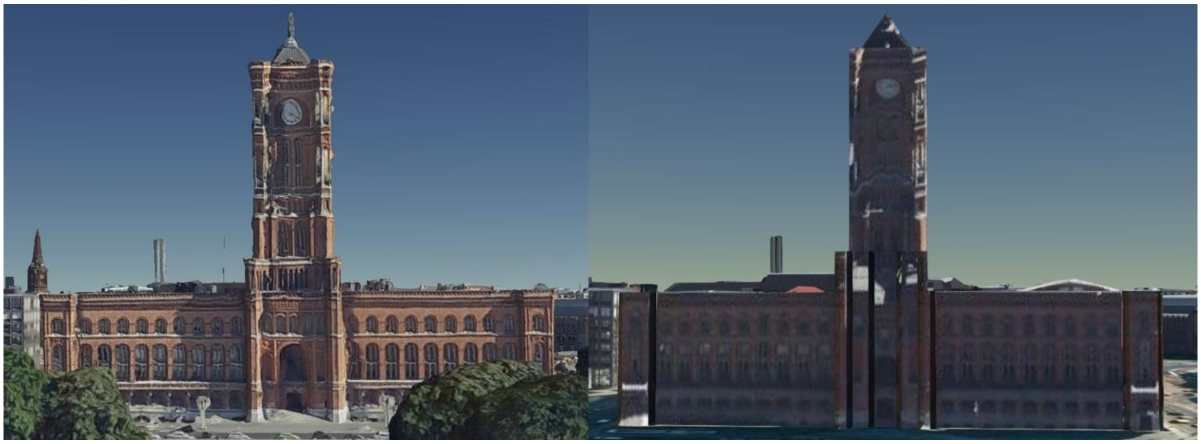

Figure 1 . Photogrammetry (left) vs LOD2 model (right)[3]

Figure 2. Metrovista photomesh of Tower Bridge and surrounds, London

Further context to the Physical layer is provided by topographic and landuse/landcover datasets that can be draped semi-transparently over the imagery. Ordnance Survey (OS) MasterMap[4] datasets including Topography Layer, Highways Network – Roads, and AddressBase were used in this project, though less detailed open source datasets are also available. Major infrastructure like power and transportation networks are also considered part of the Physical layer as they have significant impact on the landscape as well as relevance to the potential uses of urban digital twins. The power transmission network data was provided by National Grid[5] and the distribution networks by UK Power Networks[6] (UKPN), Scottish & Southern Electricity Networks[7] (SSEN), and OpenStreetMap[8] (OSM).

[2] Metrovista https://www.blueskymapshop.com/products/metrovista
[3] Berlin 3D, https://www.businesslocationcenter.de/berlin3d-downloadportal/
[4] OS MasterMap https://www.ordnancesurvey.co.uk/business-government/products
[5] National Grid https://www.nationalgrid.com/electricity-transmission/network-and-infrastructure/
[6] UK Power Networks https://www.ukpowernetworks.co.uk
[7] Scottish and Southern Electricity Networks https://www.ssen.co.uk/our-services/tools-and-maps/
[8] OpenStreetMap https://www.openstreetmap.org/

Human Layer

The major component of the Human layer is demographic data. Census data[9] is used to generate a geotypical population down to the individual entity. As the 2021 census results are expected to be published in October 2022, the 2011 census was processed as a placeholder. For a given census output area counts of gender and age are preserved, though the combinations on an individual level will not necessarily be true to life. These individuals are assigned home and work locations based on address point data. Depending on the specific context of the digital twin, the availability of data and on nations' applicable laws and regulations, this population dataset could be replaced by a true database of individuals and/or live location feeds. Key socioeconomic indicators from the census like ethnicity, income and employment level are also displayed as thematic maps.

Political elections results[10] are also of interest as party affiliation can have an impact on understanding underlying context and reasons for the human responses in a given area as well as the simulated outcome of various courses of actions where applicable. Simplistically, the same action in a right wing dominated area might be seen as positive whereas a left wing area might react negatively and vice versa. Various administrative boundaries, including special zones of regulatory bodies are important for any decision making, as these affect who is in charge and which agencies need to be consulted.

Cognitive Layer

The link between social and other media content (such as news outlets) with the Human and Physical layer data can provide a basis to geolocate both content and related population sentiment. However, there are currently no geospatial components integrated into the London digital twin Cognitive layer and so discussion of it is out of scope for this paper. See Social Media Synthesis using AI for Decision Support (Harris et al, 2022) for details on this layer.

Resource Layer

The Resource layer contains the locations of emergency services like hospitals, police and fire stations. Asset availability, such as manpower, number of free hospital beds, fire engines, etc., could be projected/theoretical or come from real-time feeds. For the London digital twin location data for these service points came from the OS AddressBase dataset, while asset information came from a variety of datasets in the London Datastore[11], UK Open data[12], and National Health Service (NHS) Statistics[13] data portals.

DATA PROCESSING PIPELINES

Photogrammetric Mesh

A key input to mesh generation is high resolution aerial imagery from a multi-head digital sensor. The sensor comprises of a nadir camera and four oblique cameras, each camera generating a 151 MPixel multispectral image. The sensor is flown along with a GNSS receiver and Inertial Navigation System (INS), which produce an accurate trajectory from which initial exterior orientations for each image exposure can be extracted. The appearance of the aerial images is important, and a radiometric correction is applied to the imagery to create as balanced an appearance as possible across the project area.

Once acquired and balanced, the images need to be accurately tied together and connected to a control network; this is achieved with aerial triangulation. This task generates a set of tie points - locations which are precisely measured on all overlapping aerial images, both nadir and obliques, and this helps define the relative orientation from image to image. In addition to the tie point measurements, precisely surveyed ground control points are measured on all overlapping images. A bundle adjustment is performed, which takes all the tie point measurements across the images,

[9] UK Census 2011 https://www.nomisweb.co.uk/census/2011
[10] 2019 General Election: https://commonslibrary.parliament.uk/research-briefings/cbp-8749/
[11] London Datastore https://data.london.gov.uk/
[12] UK Open data https://data.gov.uk/
[13] National Health Service https://www.england.nhs.uk/statistics/

the ground control point image measurements, the control points surveyed coordinates as well as the initial exterior orientations of each image, and computes the best fit amongst all the inputs. The main output of the bundle adjustment is a refined set of exterior orientations for every aerial image. These parameters, along with precise camera calibration parameters, feed into the mesh generation process.

The mesh generation can be broken down into several key steps (Figure 3): point cloud generation, mesh model generation and finally texturing. Input parameters are the aerial nadir and oblique images, camera calibration parameters and exterior orientations generated from the aerial triangulation bundle adjustment.

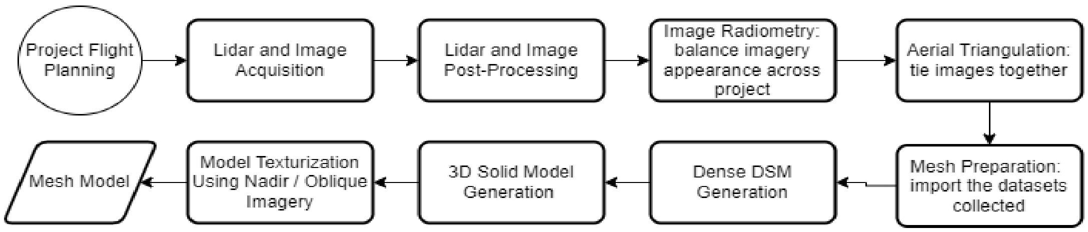

Figure 3. Photomesh Generation

The first stage is the computation of a dense digital surface model (DSM) point cloud based on the aerial images; this can be augmented with a LIDAR point cloud. The LIDAR can help to fill in potential void areas of the photogrammetric DSM and help define the mesh model better in narrow streets as well as around and under trees. The mesh model is generated using the dense DSM, using a technique like the 'marching cubes' algorithm (Lorensen & Cline, 1987). The final task is the texturing of the mesh model, each mesh polygon using the optimum 'look' to either a nadir or oblique image for its texture.

The 3D Tiles format was one that we had not worked with before. Testing was done to work out the process for moving from the photogrammetric mesh to one capable of being ingested by the cesium model tiler. Considerations included moving from the native coordinate system (British National Grid) to a local cartesian coordinate system based on an Earth-centered, Earth-fixed ECEF reference frame, calculating the origin of the 3D Tileset for position in latitude, longitude and ellipsoidal height and switching axis directions.

The tileset's use of ellipsoidal height (Figure 4) allows the data to be rendered efficiently on the GPU without a runtime coordinate conversion step. Additionally, grid shift files such as EGM2008 that are normally required for vertical coordinate transformations do not need to be packaged with the runtime engine.

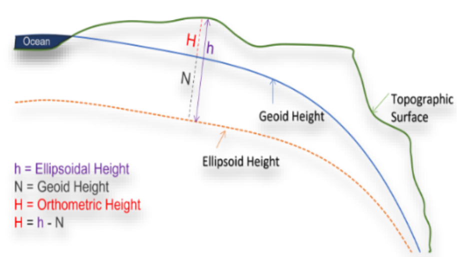

Figure 4. Ellipsoidal vs Orthometric Height

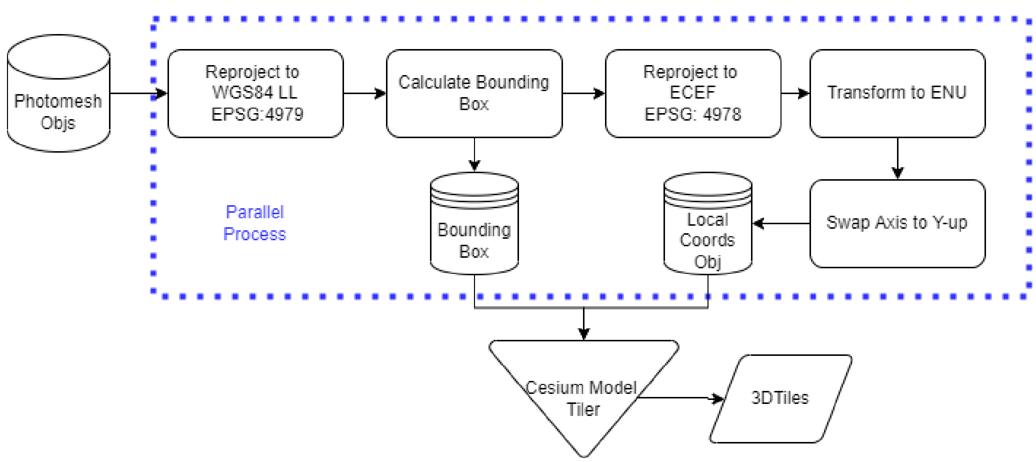

Figure 5. Photomesh obj to 3D Tiles processing

A commercial tool, FME, was used to rapidly prototype the conversions and establish the workflow. .NET scripts using GDAL were written to parallel process the data in the cloud. (Figure 5). Scaling up from the sample area introduced a new issue; when the size of input textures exceeded 1.5 GB it resulted in a heap out of memory error. Increasing the memory allocation did not resolve this error, so we opted to process the data in smaller batches. Processing 9GB of 5cm photogrammetry data took 3200 CPU hours, but only 130 hours of actual time.

Terrain Clamping

The previously mentioned lack of separation between terrain and surface features in the photogrammetry models lead to some undesired behaviour in the Unreal engine where the pedestrians moving along the sidewalk would fly into tree canopies or overhanging balconies, traverse them, and then drop back down to the ground. A separate LIDAR digital terrain model was proposed as a terrain clamping source, but this proved insufficient for activities like climbing exterior steps, where the pedestrians were occasionally knee deep into the stairwell or floating slightly above the steps. Instead, we opted to process the photogrammetric model to remove all surface features, leaving a “bare earth” mesh for clamping.

This was accomplished by first removing all faces overlapping in xy, keeping only the one with the lowest Z-value. Then vertical faces were removed based on the Z-normal, keeping only those with a value greater than 0.8. The isolated clusters of mesh faces were then compared to the digital terrain model to determine if they represented the top of a surface feature such as flat roofs of buildings or cars or were ground features like steps. The unwanted clusters were discarded and then the resulting holes were interpolated. Once again, the recipe (Figure 6) was developed using FME followed by a C# script written to allow parallel processing of the full dataset.

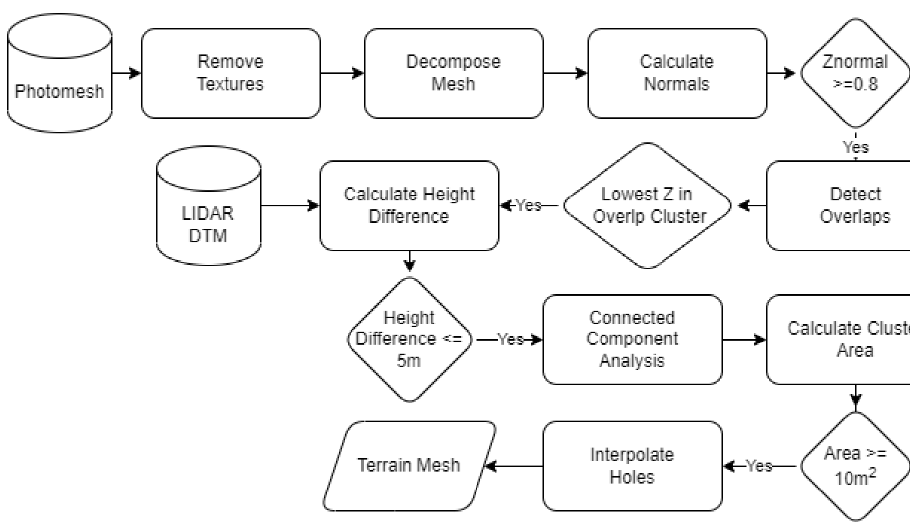

Figure 6. Bare earth processing diagram

Population

As discussed, the individual population dataset was derived from the UK 2011 census. The pattern of life simulation requires origin and destination points for the individuals. The Ordnance Survey (OS) AddressBase data contains a classification type for each location, using FME these classifications were mapped to residential (origin) and various destinations such as workplace, school, and retail. Then each individual was semi-randomly assigned locations, taking into account broad constraints such as the size of the business, the average household size, and the number of workers/commuters per area from the census data. Further details on both the population generation and full scale pattern of life can be found in Large Scale Pattern of Life Simulation for Real-Time Applications (Giannias et al, 2022).

Walkable Areas

The process of creating walkable area inputs, from which to derive a navigational mesh, was more involved. The OS MasterMap Topography Layer contains detailed polygons with feature rich attribution, in addition to the pavement (sidewalk) classification, other pedestrian friendly features including pathways, footbridges, grass surfaces (parks) and steps. Filtering to just these features showed two distinct issues, the most significant being that there was no data to indicate crosswalks, and secondly that driveways that crossed the pavement were part of the main road polygons. This leads to hundreds of isolated areas with no way to cross the street from one block to another (Figure 7).

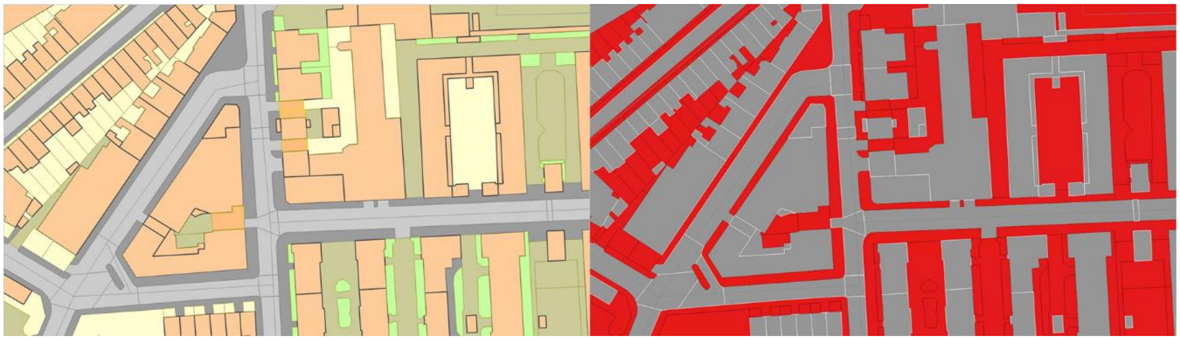

Figure 7. OS Topographic Data (left) and Walkable features in red (right)

The driveway gaps were filled in by triangulating between the pavement polygons with a maximum length of 10m. The resulting triangles were then filtered to keep only those that a) intersected a road polygon and b) did not intersect the road centerline. Polygons that were less than 3m^2 were discarded (Figure 8).

Figure 8. Driveway triangulation

To create crosswalks, we first used the OS MasterMap Highways data to identify nodes that correspond to a crossable intersection, as opposed to pseudo nodes, cul-de-sacs, or traffic islands. Once these nodes were identified, the corresponding road polygons were converted to linears and broken into segments based on angle. The road segment that crossed the road centerline was considered a candidate for a crosswalk. The candidate crosswalk was extended slightly in both directions and checked to see if it intersected 2 pavement polygons at both ends. If so, the line was buffered on the side away from the intersection node to create a 3m wide crosswalk (Figure 9).

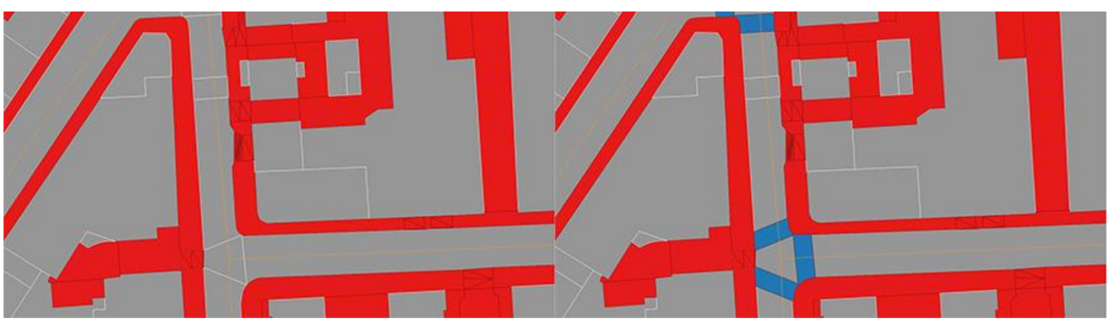

Figure 9. Walkable Areas before (left) and after (right) crosswalk generation

When the prototype area was expanded to the whole of London, a further source of disconnection from underpasses was discovered. The areal topographic layer is seamless, with no overlapping polygons, so both roads and pavements are interrupted by road and railroad bridges. Fortunately, the linear topographic layer includes boundary and obstacle features that can be used to construct the majority of these missing underpasses (Figure 10).

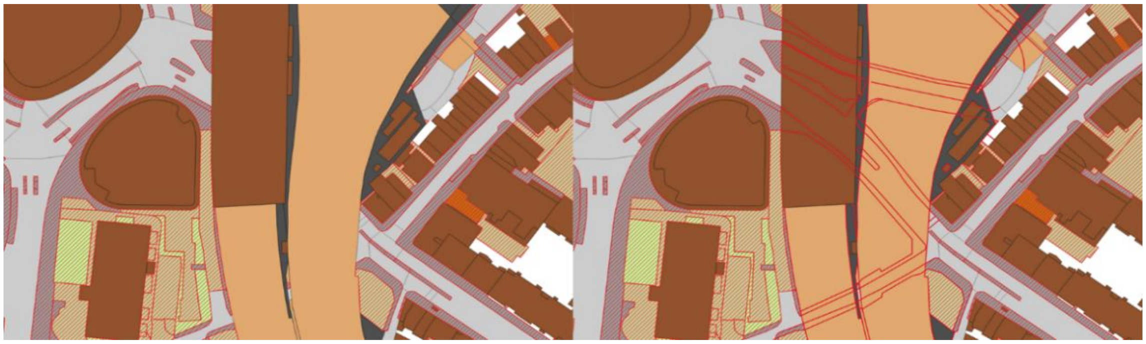

Figure 10. Areal topographic data (left) and topographic linears for potential underpasses (right)

The full set of walkable areas including the generated crosswalks and driveways and underpasses was put through a connected component analysis and isolated clusters were either corrected to connect them to the rest of the network or removed. A grassed or paved road median for instance would be removed unless it formed part of a crosswalk.

As the origin and destination points were generally inside buildings and the navigational software we were using at the time did not have the ability to snap a point to the navigation mesh, we created "walkways", polygons from the origin/destination points to the nearest walkable area. The software has since been updated to allow for snapping. The creation of walkways is still a valid concept when the origin/destination points are static or a subset of static data, as it will reduce the run time calculations and improve efficiency.

Lastly the data was reprojected to LL-WGS84 and donuts were cut into regular polygons due to another limitation of the software. Cutting the donuts after the reprojections prevents any disconnections in the resulting dataset due to floating point precision issues during the reprojection (Figure 11).

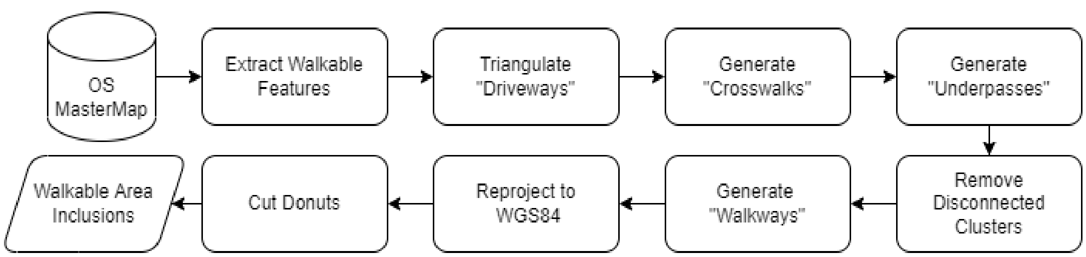

Figure 11. High level workflows for walkable areas

Power Network

The cyberattack scenario for our London digital twin required identifying zones affected when any given power infrastructure (generator, substation) is targeted and taken offline. In the ideal situation we would have full network information down to which structures are connected to which transformers. In reality we had limited data from four different sources, with entirely different schemas and levels of details. In the worst instance we had substation locations with voltage information, but no actual cable data for the distribution network.

To harmonize the data the substations were split into 3 categories: high voltage (HV), extra high voltage (EHV) and transmission with consistent attributes and projection. An influence zone was created for each high voltage substation, either through Voronoi polygons of the substations themselves, or in the case where distribution cables existed through the aggregation of structures connected to that substation. In the latter case a preprocessing step was required to create a topologically connected network as the substations were polygons and the cable linears ended inside the polygon but did not actually connect to one another. To address this, nodes were generated for power stations and substations and the cables were snapped to those nodes.

Each HV substation was then associated to an EHV substation either through connected cables where available, or by proximity where not. The EHV influence zone was created by dissolving all the associated HV substations. This same process was then repeated to associate EHV substations to Transmission substations, and Transmission to Power Generation Plants.

The disruption of any level of power infrastructure would cause a cascade of failures through the child substations and then highlight the complete area affected. This is an oversimplification of the electrical grid and does not consider the multiple redundancies in a real life system. Future work is required to take into account the maximum power output/throughput of the various stations and how the power requirements are load balanced across the remaining infrastructure when one or more pieces are taken offline.

CONCLUSION

There exists a variety of relevant data, formats, and open standards to support the development of urban digital twins. In this paper, we addressed some of these sources which built the four conceptual layers upon which our digital twin was designed. Even within the Greater London Area data availability and quality was inconsistent. Any large scale digital twin will have to be flexible enough to allow for these differences, as restricting the framework to just those datasets available nationally, or even regionally, will greatly reduce the effectiveness of the digital twin for decision making. At the same time, the geospatial pipelines should be designed with an eye to producing as close to standardized outputs as possible without locking into a specific schema that may be ideal for one location and very flawed for a different location.

It cannot be assumed that data of a quality sufficient for a digital twin is immediately available for any urban area globally. The first step for any digital twin implementation should be an investigation of the available data in terms of quality, including not only scale, accuracy, and attribution, but also currency. A highly detailed dataset that is 30 years old is generally not as useful as a less detailed dataset that is only a year old. The authority of a dataset should also be taken into consideration. While crowd-sourced datasets like OpenStreetMap can adapt to changes far more rapidly than an official source, which generally has a fixed yearly or quarterly update schedule, they often have little quality control and are vulnerable to bad actors deliberately compromising the integrity of the data. The criteria for which datasets to include in the digital twin should be established based on the specific use cases.

The workflows detailed in this paper have laid the foundation for patterns and approaches, including data variation points, to provide a method baseline to support the wider adoption of urban digital twin development. We have shown that it is possible to have simultaneous use of 3DTiles in a web browser and game engine, and correlate high fidelity interactive visualisation with the underlying GIS data. By leveraging open-source data and applying the methodologies developed in the London digital twin, we are well on the way to developing a country scale digital twin for our next project, Estonia.

REFERENCES

Giannias, N., Harris, E. P., Geraerts, R., Herfst, S., Hobson-Sayers, A., & Thorpe, A. (2022). Large Scale Pattern of Life Simulation for Real-Time Applications. *Interservice/Industry Training, Simulation and Education Conference, 22235*.

Harris, E. P., Kaur, J., Chaouachi, M., Durocher, M., Tiwari, R., Emirov, A., & Giannias, N. (2022). Social Media Synthesis using AI for Decision Support. *Interservice/Industry Training, Simulation and Education Conference, 22175.*

Ivanov, S., Nikolskaya, K., Radchenko, G., Sokolinsky, L., & Zymbler, M. (2020). Digital Twin of City: Concept Overview. *Global Smart Industry Conference*, 178-186.

Khronos 3D Formats Working Group, The. (2021). glTF™ 2.0 Specification. https://www.khronos.org/registry/glTF/specs/2.0/glTF-2.0.html

Kirimtat, A., Krejcar, O., Kertesz, A., & Tasgetiren, M. F. (2020). Future Trends and Current State of Smart City Concepts: A Survey. *IEEE Access* (8), 86448-86467.

Lorensen, W. E., & Cline, H. E. (1987). Marching cubes: A high resolution 3d surface construction algorithm. *ACM Computer Graphics. 21 (4)*,163–169.

Open Geospatial Consortium Inc. (OGC). (2006). OpenGIS Web Map Server Implementation Specification. OGC 06-042. https://www.ogc.org/standards/wms

Open Geospatial Consortium Inc. (OGC). (2014). OpenGIS Web Feature Service 2.0 Interface Standard. OGC 09-025r2. https://www.ogc.org/standards/wfs

Open Geospatial Consortium (OGC). (2019). 3D Tiles Specification 1.0. OGC 18-053r2. https://docs.ogc.org/cs/18-053r2/18-053r2.html

Shahat E., Hyun C. T., & Yeom C. (2021). City Digital Twin Potentials: A Review and Research Agenda. *Sustainability 13 (6)*, 3386.

Zheng, Y., Yang, S. and Cheng, H. (2019) An application framework of digital twin and its case study. *Journal of Ambient Intelligence and Humanized Computing (10)*, 1141–1153.