

# Semantic 3D City Interfaces – intelligent interactions on Dynamic Geospatial Knowledge Graphs

Arkadiusz Chadzynski<sup>1</sup>, Shiyong Li<sup>2</sup>, Ayda Grišiūtė<sup>2</sup>, Jefferson Chua<sup>3</sup>,  
Jingya Yan<sup>1</sup>, Huay Yi Tai<sup>1</sup>, Emily Lloyd<sup>1</sup>, Mehal Agarwal<sup>1</sup>,  
Markus Hofmeister<sup>3</sup>, Jethro Akroyd<sup>3,1</sup>, Pieter Herthogs<sup>2</sup>, Markus Kraft<sup>3,1,4,5</sup>

released: August 2, 2022

<sup>1</sup> CARES  
Cambridge Centre for Advanced  
Research and Education in Singapore  
1 Create Way  
CREATE Tower, #05-05  
Singapore, 138602

<sup>2</sup> Singapore-ETH Centre,  
Future Cities Lab Global Programme,  
CREATE campus,  
1 Create Way,  
#06-01 CREATE Tower,  
Singapore, 138602

<sup>3</sup> Department of Chemical Engineering  
and Biotechnology  
University of Cambridge  
Philippa Fawcett Drive  
Cambridge, CB3 0AS  
United Kingdom

<sup>4</sup> School of Chemical  
and Biomedical Engineering  
Nanyang Technological University  
62 Nanyang Drive  
Singapore, 637459

<sup>5</sup> The Alan Turing Institute  
2QR, John Dodson House  
96 Euston Road  
London, NW1 2DB  
United Kingdom

Preprint No. 297



---

*Keywords:* GeoWeb, Cognitive Architecture, Artificial Intelligence, Knowledge Graph, Automation, City Modelling, Geospatial, CityGML, Semantic Web, Ontology

**Edited by**

Computational Modelling Group  
Department of Chemical Engineering and Biotechnology  
University of Cambridge  
Philippa Fawcett Drive  
Cambridge, CB3 0AS  
United Kingdom

**E-Mail:** [mk306@cam.ac.uk](mailto:mk306@cam.ac.uk)

**World Wide Web:** <https://como.ceb.cam.ac.uk/>



## Abstract

This paper presents a system architecture and a set of interfaces that can build scalable information systems capable of large city modelling based on dynamic geospatial knowledge graphs to avoid pitfalls of Web 2.0 applications while blending artificial and human intelligence during the knowledge enhancement processes. We designed and developed a GeoSpatial Processor, an SQL2SPARQL Transformer, and a Geospatial Tiles Ordering tasks and integrated them into a City Export Agent to visualise and interact with city models on an augmented 3D web client. We designed a Thematic Surface Discovery Agent to automatically upgrade the model's level of detail to interact with thematic parts of city objects by other agents. We developed a City Information Agent to help retrieve contextual information, provide data concerning city regulations, and work with a City Energy Analyst Agent that automatically estimates the energy demands for city model members. We designed a Distance Agent to track the interactions with the model members on the web, calculate distances between objects of interest, and add new knowledge to the Cities Knowledge Graph. The logical foundations and CityGML-based conceptual schema used to describe cities in terms of the OntoCityGML ontology, together with the system of intelligent autonomous agents based on the JPS Agent Framework, make such systems capable of assessing and maintaining ground truths with certainty. This new era of GeoWeb 2.5 systems lowers the risk of deliberate misinformation within Geography Web Systems used for modelling critical infrastructures.



## Highlights

- GeoWeb 2.5 minimising deliberate misinformation risks while modeling critical infrastructures.
- Geographical Web Systems blending human and artificial intelligence during knowledge enhancement.
- Interactive interfaces to Dynamic Geospatial Knowledge Graphs.

# Contents

<b>1</b>	<b>Introduction</b>	<b>3</b>
<b>2</b>	<b>From GIS to GeoWeb 2.0 and Participatory Urban Design</b>	<b>5</b>
2.1	GIS, Big Data and City Modeling at Scale . . . . .	5
2.2	A "View From Everywhere" on the GeoWeb 2.0 . . . . .	13
<b>3</b>	<b>Enabling GeoWeb 2.5 on Dynamic Geospatial Knowledge Graphs</b>	<b>15</b>
3.1	Intelligent Autonomous Agency and Knowledge Enhancement . . . . .	15
3.1.1	Interactive Interfaces to the Thematic Surface Discovery Agent .	16
3.1.2	Interactive Interfaces to the City Information Agent . . . . .	19
3.1.3	Interactive Interfaces to the Distance Agent . . . . .	22
3.2	Ground Truth Maintenance on Dynamic Geospatial Knowledge Graphs .	23
<b>4</b>	<b>Conclusions and Future Work</b>	<b>25</b>
<b>A</b>	<b>Thematic Surface Discovery Agent - UML Activity Diagram</b>	<b>28</b>
<b>B</b>	<b>City Information Agent - UML Activity Diagram</b>	<b>29</b>
<b>C</b>	<b>City Energy Analyst Agent - UML Activity Diagram</b>	<b>30</b>
	<b>References</b>	<b>32</b>

# 1 Introduction

## General context and problem space

Dynamic geospatial knowledge graphs that make use of existing standards [62] ported to ontologies [10] demonstrate the capabilities of scalable information systems based on sustainable digitisation practices [70] and enable interoperability [7] between otherwise siloed datasets. As described by Chadzynski et al. [11] such systems, coupled with intelligent automation based on cognitive architecture [46], could aid efforts focused on finding answers to globally most pressing problems [68] and existential threats [55]. Both, the means and the ends, are recognised and agreed upon by the global governing bodies, such as the UN, G20 and the World Bank [62], as well as many individual countries around the world [1, 3, 14, 58, 72]. There is also a recognition that, apart from the just mentioned capabilities as well as the capability to conduct complex multi-factor optimisation scenarios at a city scale [60], urban modelling could serve as a tool for storytelling allowing decision-makers to communicate the motivations behind policy implementations to wider audiences [36]. This, in turn, requires the design and development of appropriate interfaces to such knowledge graphs.

Cities Knowledge Graph (CKG) [16], an active research project collaboratively worked on by the Cambridge Centre for Advanced Research and Education in Singapore (CARES) [15] and the Singapore-ETH Centre (SEC) [17], is an example of a dynamic geospatial knowledge graph based on sustainable digitisation practices [10]. As a knowledge graph built around the Semantic 3D City Database [10] that is a semantic equivalent of the 3D City DB [73], originally developed at the Technische Universität München (TUM) for relational geospatial databases [66], it is designed to produce and process multi dimensional representations of urban environments modelled in accordance with the CityGML 2.0 standard by the Open Geospatial Consortium (OGC) [35]. This knowledge persistence layer, developed to be compliant with the semantic web standards and recommendations provided by the W3C, is coupled with a system of intelligent autonomous agents [11] designed upon principles of a cognitive architecture [49].

CKG is a part of The World Avatar (TWA), a knowledge graph encompassing other micro and macro scale domains. This approach combines city-specific knowledge with knowledge from multiple heterogeneous domains and provides a comprehensive point of view applying to a feature under consideration. TWA has been already used for optimal site selection of modular nuclear power plants [20], simulations of chemical kinetic reaction mechanisms [24], quantum chemistry calculations [47], combustion chemistry [24] and power systems optimisation in eco-industrial parks [21]. Cities are known to be one of the largest contributors into the global CO<sub>2</sub> emissions due to their high energy demands, so far mostly supplied by fossil fuels. A thorough understanding of those macro structures taking into account multiple points of view could optimise those energy demands [53] and reduce their carbon footprint [23]. Knowledge graphs are designed to facilitate such interoperability as well as provide information that is usually inaccessible without combining multiple datasets.

## Human Computer Interaction on the GeoWeb 2.0

A combination of issues related to data integration and the ability to easily update parts

of existing models [8] as well as collaborative work on them by multiple people simultaneously has been an *Achilles' heel* of traditional GIS applications [2]. This motivated research [61] and development [12] of GeoWeb services and applications that align geographical modelling closer with the rest of the modern software. Although the GeoWeb moves GIS into a new architectural pattern where data, applications and services are more openly available and accessible even to *neo-geographers* [44] for content mashups, the service providers within the industry still emphasise the importance of standardisation to facilitate interoperability [19].

This new paradigm requires new interfaces allowing to enter new information, retrieve, remove or correct existing information to the GeoWeb systems on a very large scale. The interfaces should be easy enough to operate as well as accessible to interested participants exposing feature-rich information systems without requiring extensive training or expertise. Interfaces to the traditional GIS should be extended or modified to adapt them to this new set of requirements. "Where once a geographer would pack a map, compass and notepad, the NeoGeographer merely needs a mobile device to explore, navigate, document, capture and log the environment." [44]. Web 2.0 proved to be able to integrate more traditional information systems as well as enable them for wider participation in the past. GeoWeb interfaces make use of similar Human Computer Interaction (HCI) principles and port proven existing designs to work with the new type of content [65].

### Synthesis

TWA could be regarded as an example of a new paradigm for the GeoWeb information systems that go beyond Web 2.0. Being based on the Semantic 3D City Database [10] at its core, it ports existing GIS standards to a new graph database and takes advantage of the Open World Assumption (OWA), which is absent in the equivalent relational geospatial databases [66]. Coupling it with a system of intelligent autonomous agents based on cognitive architecture extends and scales existing geospatial data transformation tools. Its agents proved to be able to automatically create a semantic model of Berlin that consists of 419 909 661 Subject-Predicate-Object statements in total. Apart from the traditionally laborious and error-prone process of model creation, the agents also autonomously created a representation of it suitable for interaction with the model using web interfaces. Moreover, the agents demonstrated the ability to track user interactions with the model on the web, create new knowledge and also display it to the user autonomously [11].

Although there have been similar attempts to visualise geospatial knowledge using semantic web principles, they seem not to reach large city model scale [41]. Moreover, it is well recognised that Web 2.0 systems introduce a new set of problems and challenges to overcome. As pointed out by Goodchild [34], the risk of deliberate misinformation, which is prevalent on the other web domains, is inherited by GeoWeb 2.0. However, at the same time, systems that provide geospatial models are arguably more critical to the policymakers at the local as well as national governance levels. Therefore, it is even more important to minimise and mitigate such risks on the new GeoWeb that is designed to enable wider and non-expert participation to avoid potential deliberate policy misinfluence.

**The purpose of this paper** is to present multiple web interfaces to the 3D city objects of the TWA that allow for interactions blending artificial and human intelligence as well as to point out how dynamic geospatial knowledge graphs can minimise Web 2.0 prob-

lems and open a new era of GeoWeb 2.5 systems. The issues concerning more traditional Geospatial Information Systems (GIS) and the resulting motivation for the GeoWeb are elaborated on in the next section (section 2). It also shows how TWA already overcomes particular issues by providing a scalable way of interacting with large city models on the web. Section 3 focuses on the interfaces that allow TWA agents to autonomously enhance existing city models by upgrading their Levels Of Details (LODs) (subsection 3.1.1), interact with city objects and retrieve information from other domains that is applicable to them (subsection 3.1.2) as well as track user interactions with the city objects on the web to add new analytical knowledge (subsection 3.1.3). The advantages of knowledge graphs, and the way in which they can avoid pitfalls of Web 2.0 systems and mitigate the risks of potential misuse of the GeoWeb systems are presented in subsection 3.2. The last section 4 includes conclusions and potential directions of the future research and development for GeoWeb 2.5 based on such dynamic geospatial knowledge graphs.

## 2 From GIS to GeoWeb 2.0 and Participatory Urban Design

The majority of the issues related to GIS of the past revolve around their scalability. They come to light when one starts trying to make use of traditional systems for applications imposed by Industry 4.0 [52]. Because of the concurrent real-time data collection from multiple heterogeneous sources that is constantly transformed and analysed to fulfil its application demands, the era of Smart Cities is also the era of Big Data [2]. Such demands motivated the evolution of the traditional systems to Geospatial Web Services (GWS) based on open standards to ensure interoperability [22]. Such systems are also referred to as GeoWeb 2.0 or even as WikiGIS [65]. The Wiki aspect indicates that GWS, apart from being able to accommodate and integrate heterogeneous data generated by machines, also open those systems to collaboration between their human counterparts with a diverse range of skills and motivations. This shows that GWS have even the potential to realise the idea of participatory urban design and planning [59]. TWA, being an information system containing web interfaces, could be regarded as a GWS, realising some of the mentioned ideas. The next two subsections describe this aspect of it in more detail.

### 2.1 GIS, Big Data and City Modeling at Scale

At the very high level, TWA architecture could be compared to the architectures following the Model-View-Controller (MVC) design pattern [29]. City objects are modelled within the Semantic 3D City Database that is a graph equivalent of the 3D City DB. The objects are described in a form of quads [63] that assign every statement to a named graph that is an equivalent to a corresponding original relational database table [10]. It allows mapping the semantically stored objects to software models in Java by using an Object Graph Mapper (OGM) engine developed for this purpose within the TWA. It could be roughly compared to the existing Object Relational Mapper (ORM) engines commonly used with object orientated programming languages [51]. The cognitive agents layer could be com-

pared to the controller layer in the MVC. TWA agents are based on the cognitive architecture of the JPS Agent framework [11]. They operate on instances of the city objects in Java that are mapped to appropriate graphs. While agents perform intelligent operations on the instances, objects' instantiation, persistence of updates, additions and deletions are taken care of by the OGM layer automatically. All the changes to the model as a whole as well as to the individual city object members could be immediately reflected on the TWA view layer thanks to its reuse of the augmented 3D web client, based on Cesium virtual globe framework [9] and previously worked on at TUM [13]. The view specific data is exported in the Keyhole Markup Language (KML) [57] format by the City Export Agent (CExA) that is capable of generating it for the whole large model at once or for individual model members, depending upon the request received from the other agents. The parts of the model are loaded dynamically on the client because of the adopted geospatially ordered data tiling strategy. This set of features put together also differentiates the TWA from the original TUM stack by making it capable of reflecting changes to the model on the web interface immediately. The original stack requires laborious manual work using desktop tools to achieve the same end result [66]. Eliminating manual steps from the process makes TWA a fully dynamic GWS solution that also retains other advantages of knowledge graphs at the same time.

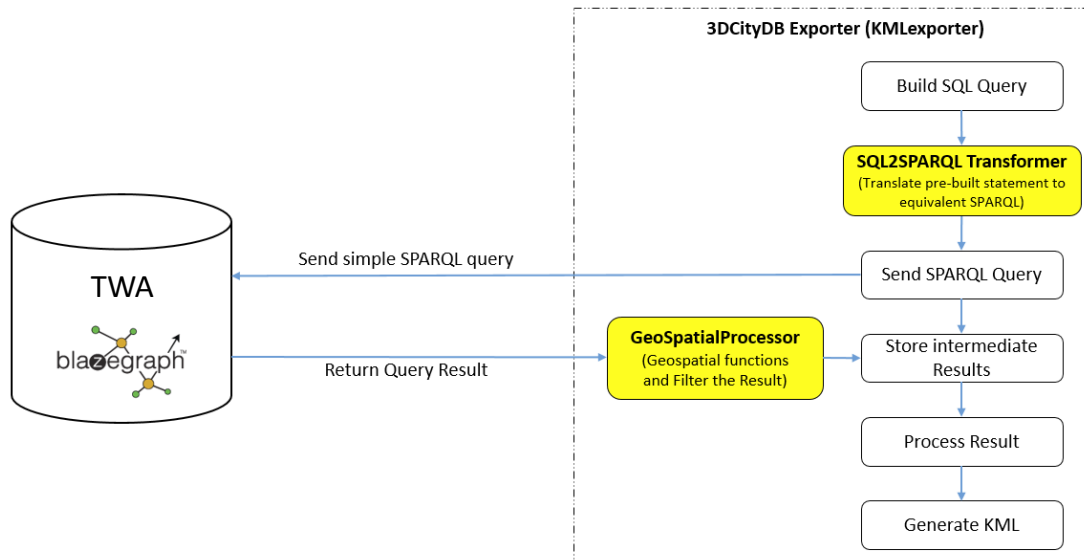
Cesium framework is widely used to visualise 3D geospatial data and develop web-based 3D geospatial applications in a virtual globe. It provides a GWS user ability to freely move around in the virtual environment by zooming in and out, changing the viewing angle and position. It supports many open formats to visualise 3D data, including CityGML, GeoJSON, KML/COLLADA and so on. 3D Tiles within Cesium are designed for optimised streaming and rendering large volumes of 3D geospatial data. It is an open specification for sharing, visualising, fusing, and interacting with massive heterogeneous 3D geospatial models across different platforms [32], which is adopted by OGC as a community standard. TWA transforms the semantic 3D building representation from the knowledge graph data model to 3D Tiles in KML format based on the hierarchy level of detail of 3D Tiles. It is relatively easy to implement 3D visualisation and develop interactive 3D geospatial applications using CesiumJS. However, it supports only two commonly used coordinate systems, one is the WGS84 geographic coordinate system and the other is the Cartesian space Cartesian coordinate system. Hence, for the 3D visualisation, all data needs to have converted its geographic coordinate system to WGS84.

The framework offers an off-the-shelf web visualisation component supporting open standards without requiring to build a custom 3D rendering engine [61]. Visualisation data preparation strategy in a form of geospatially ordered and dynamically loaded tiles allows to visualise cities of the size of Berlin in the TWA at LOD2 at the present moment [18]. The framework has also been used for other thematic 3D visualisations on the web [31] as well as for visualisations of dynamic data sources, such as floods [48] and energy flow simulations [54] which are of a particular interest when city models are combined with models from other domains within TWA.

The KML visualisation data is prepared by the CExA, implemented as one of the TWA agents using the JPS Agent framework. It embedded the augmented version of the Importer/Exporter tool (ImpExp) so that it can export the city model from the semantic data store (e.g., Blazegraph™) by sending an HTTP request with required inputs. After the



validation of the inputs and preparation of the configuration file, the ImpExp code is used to export the city model and save it into KML files which are added as a data layer for visualisation on the 3D Web Map Client. The connectivity to the semantic database has been augmented using Jena JDBC, a SPARQL over JDBC driver framework. Moreover, the ImpExp needed the development of two additional components to make it work with semantic database: SQL2SPARQL Transformer and GeoSpatial Processor. Their placement within the overall architecture is illustrated in Figure 1.



**Figure 1:** Additional components, not present in the original TUM tool, needed to export web visualisation data in KML format from a raw semantic data store.

The SQL2SPARQL Transformer translates the SQL statements to equivalent SPARQL statements according to the OntoCityGML schema. Some SQL statements in the original ImpExp make use of the built-in geospatial functions provided by the PostGIS database. However, those functions are not present in the current Blazegraph™ version, such as *ST\_TRANSFORM* for coordinate transformation, *ST\_Area* for calculating the area covered by polygon, *ST\_IsValid* for checking valid geometry etc. To overcome the limitation of the current Blazegraph™, a GeoSpatialProcessor class is implemented to post-process the query results in a similar manner to provide the same accurate information for generating KML files that are used for the visualisation of the city model. Another limitation of current Blazegraph™ is limited computational power for large data storage. To overcome these limitations, some of the complex queries need to be factorised and executed one by one so that part of the computational load can be taken off the database. This implementation can be found in the StatementTransformer class using Java implementation.

To query the corresponding surface geometry of the buildings, a sequence of different queries are executed to retrieve the desired geometry information. One of the SQL examples is presented in the Listing 1:

```

1 SELECT ST_Union(get_valid_area.simple_geom) FROM (

```

```

2  SELECT * FROM (
3      SELECT * FROM (
4          SELECT ST_Force2D(sg.geometry) AS simple_geom
5          FROM citydb.SURFACE_GEOMETRY sg
6          WHERE sg.root_id IN (
7              SELECT b.lod2_multi_surface_id FROM citydb.BUILDING b
8              WHERE b.id = ? AND b.lod2_multi_surface_id IS NOT NULL
9              UNION
10             SELECT b.lod2_solid_id FROM citydb.BUILDING b
11             WHERE b.id = ? AND b.lod2_solid_id IS NOT NULL
12             UNION
13             SELECT ts.lod2_multi_surface_id FROM citydb.
14                 THEMATIC_SURFACE ts
15                 WHERE ts.building_id = ? AND ts.lod2_multi_surface_id IS
16                     NOT NULL)
17             AND sg.geometry IS NOT NULL
18         ) AS get_geoms WHERE ST_IsValid(get_geoms.simple_geom)='TRUE'
19     ) AS get_valid_geoms WHERE ST_Area(ST_Transform(get_valid_geoms.
20         simple_geom,4326)::geography, true) >0.001
21 ) AS get_valid_area

```

**Listing 1:** A nested SQL query used by ImpExp tool to retrieve the surface geometry. The question mark ? is the input variable placeholder and is replaced by an actual value during the evaluation

A nested query is a query within another SQL query and embedded within the WHERE clause. The results of the subquery are used in the main query as a condition to further restrict the data to be retrieved. The database engine evaluates a nested query from inner to outer, which means the most inner query is evaluated first (e.g., from lines 7 to 14 in the Listing 1). As the whole statement in the Listing 1 contains geospatial functions that are not present in the current Blazegraph™ version, a complete translation does not exist. However, part of this statement can be translated. The Listing 2 is the corresponding translation of the section from lines 4 to 15 in Listing 1.

```

1  PREFIX ocgml: <http://www.theworldavatar.com/ontology/ontocitygml/
2      citieskg/0ntoCityGML.owl#>
3  BASE <http://127.0.0.1:9999/blazegraph/namespace/berlin/sparql/>
4  SELECT ?geometry
5  FROM <surfacegeometry/>
6  WHERE { ?id ocgml:rootId ?rootId ;
7      ocgml:GeometryType ?geometry .
8  FILTER (!isBlank(?geometry))
9  { SELECT (?lod2MultiSurfaceId AS ?rootId)
10     WHERE
11     {GRAPH <building/>
12     { ?id ocgml:buildingId ?building_id ;
13         ocgml:lod2MultiSurfaceId ?lod2MultiSurfaceId
14         FILTER (!isBlank(?lod2MultiSurfaceId)) }}

```

```

14     }
15     UNION
16     { SELECT (?lod2SolidId AS ?rootId)
17         WHERE
18             {GRAPH <building/>
19                 { ?id ocgml:buildingId ?building_id ;
20                   ocgml:lod2SolidId ?lod2SolidId
21                   FILTER (!isBlank(?lod2SolidId)) }}
22         }
23     UNION
24     { SELECT (?lod2MultiSurfaceId AS ?rootId)
25         WHERE
26             {GRAPH <thematicsurface/>
27                 { ?id ocgml:buildingId ?building_id ;
28                   ocgml:lod2MultiSurfaceId ?lod2MultiSurfaceId
29                   FILTER (!isBlank(?lod2MultiSurfaceId)) }}
30         }
31     }

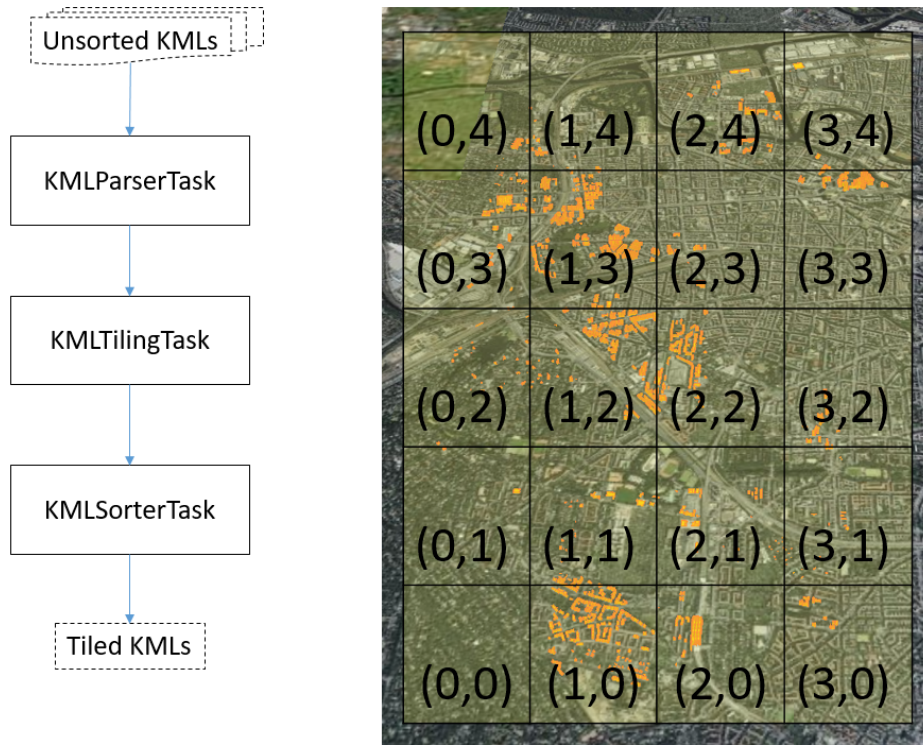
```

**Listing 2:** Equivalent SPARQL statement translated from previous SQL statement (lines 4 to 15 in Listing 1)

The Listing 2 is composed of a nested query with two stages. The first stage is a composition of three independent queries via a UNION operator where their results are merged and used in the WHERE clause of the outer query. The second stage is the outer query which queries the geometry information based on the WHERE condition of stage 1. One of the current limitations in Blazegraph™ is that its computational load increases with the increasing size of the data storage in the database. In general, this type of query can be executed seamlessly on a smaller database. However, it cannot be executed within a reasonable time range in TWA that combines city data with data representing other domains. To work with TWA, additional optimisation of the software architecture is required to achieve an acceptable execution time. One of the implemented optimisations is to factorize the complex original SPARQL statement into multiple simple queries that are executed concurrently. Returned results are collected and merged by the StatementTransformer class within augmented ImpExp to reduce computational load on the database engine. This architecture is illustrated in Figure 1.

Some of the required SQL statements in the Listing 1 make use of built-in geospatial functions provided by PostGIS, but these are not present in the current version of Blazegraph™ - the backend for the Semantic 3D City Database of the TWA. The GeoSpatialProcessor class is implemented to provide the missing geospatial functionalities such as *ST\_Transform*, *ST\_Area*, *ST\_IsValid* using Geospatial Data Abstraction Library (GDAL) [30]. As illustrated in the Listing 1, these functions are predominantly used to filter intermediate query results which can be used for the outer queries. The export process incorporating the StatementTransformer class and GeoSpatialProcessor class should provide identical information to the next process for generating KML files as the original implementation with PostGIS and SQL.

The export process generates KML files that can be added as layers to the augmented 3D web client for visualisation. However, the platform is unable to display large 3D models as a whole due to its current limitations. A geospatially ordered tiling strategy is hence adopted as a post-processing step in the CExA to reorganise all exported KML files into geospatially ordered tiles. Each tile consists of a new KML file containing information about buildings physically located within the area defined by the tile. This process allows large models exported by the CExA to utilise the existing feature in the augmented 3D web client to dynamically load and unload tiles based on the camera view.



**Figure 2:** The flowchart on the left describes reorganisation steps, that transform the exported KML files into spatial tiles. The graph on the right describes the tiling algorithms. The *KMLTilingTask* assigns each building into a tile based on the tile location and its boundaries (the grid in the foreground). The (X, Y) value on the grid indicates the tile location.

The tiling process is divided into three tasks executed in the following order – the *KMLParserTask*, *KMLTilingTask*, and *KMLSortingTask*. The *KMLParserTask* is responsible for parsing all exported KML files to identify all building objects to be processed. The NASA WorldWindJava library [71] is used to read and recognise the KML structure of the files parsed, allowing for the identification of each building object as a separate entity and extracting information from KML tags. For each building object, the coordinates of its envelope are read from the exported KML file and used to calculate the envelope’s centroid of the building object. During this process, the task also evaluates the extent of the bounding box that encloses all building objects. The *KMLParserTask* produces a sum-

mary of all building objects recorded in comma-separated values (CSV) format, which is passed to the KMLTilingTask. For each building object, this includes its gmlId, envelope, envelope centroid, as well as the exported KML file name where information of the building object is stored.

The KMLTilingTask is responsible for assigning each building object to its respective tile. Using the extent of the bounding box and a predefined length of each tile, the number of tiles and their individual boundaries are determined. A building object is assigned to a tile if its envelope centroid is located within the boundaries of the tile. The tile number of each building object is appended to the summary in CSV format and saved to a new file. After this task, the CSV content is reorganised based on the tile assignment.

Following the assignment of all building objects, the KMLSorterTask works to generate a new KML file for each tile. For each tile, the KMLSorterTask extracts the relevant building object information from the associated exported KML file, and inserts it into a newly created tile KML file. This read and write process uses the Java API for the KML library [45] to preserve the KML structure when extracting building information and to create new KML files.

The tiling process requires a lot of I/O operations which can slow down the program when the unsorted KML files are very large. The CSV summary file created during the KMLParserTask can decrease the search time and reduce the redundant I/O operations during the KMLSortingTask. Initially, while KMLSortingTask is generating the spatially organised tiles, the program extracts the relevant building object information from the initial KML files. In the worst case, one would have to search for a specific building each time among all the unsorted KML files. Using a hash table to store the temporary data reduces searching time and avoid redundant I/O operations. A hash table is known as a data structure that can map keys to values. This data structure can help to store temporary data in the computer memory.

The CSV summary file created during KMLParserTask records the associated KML file-name for each building object. Two hash tables are introduced during KMLSortingTask for the program optimisation: the first one is a building hash table that stores the building id and its geometry information, the table content is growing when an unsorted KML file is parsed into this hash map. The second hash map is called the file status hash table which records the file name and its respective status. If a file is read, the object information of all buildings within the file is added into the building hash map.

While the KMLSortingTask is writing a KML file for a tile, for each building in this tile, the program firstly gets the respective file location from the CSV summary file and checks the file status in the hash table. If the status of the relevant file is false, meaning that it has not been read, the program adds the unsorted KML file into the existing building hash map and updates the file status. After the relevant building information is written into a KML file, its record is removed from the building hash table to prevent the hash table from growing too large. Described optimisation can prevent an unsorted KML file from being read more than once as each of these files can be very large. Reading large files multiple times would increase the total processing time.

The web interface to the semantic representation of Berlin in the TWA presented at the Figure 3, showcases its capabilities to handle the so-called *five V* problems in the smart



**Figure 3:** *Visualisation of Berlin on the Dynamic Knowledge Graph architecture of the Cities Knowledge Graph, part of the World Avatar Project. KML tiles for visualisation are dynamically loaded based on the current scope. Typical problems found in traditional GIS systems that relate to model updates have been solved. Partial model updates are possible and, thanks to the design of the City Export Agent, reflecting changes on the user interface is not a problem anymore.*

city data management [2]. As described by Chadzynski et al. [10], its underlying Semantic 3D City Database uses Blazegraph™ as a data store. Because of that, this technology supports up to 50 billion edges on a single machine, and it is fully compliant with the semantic web standards; TWA is the only knowledge graph of this kind that is capable of handling Volumes of data needed to store models of large cities, so far described in the literature on the subject matter. Such models undoubtedly have a Value from the perspective of city planning [69] that can address issues mentioned in the introduction (section 1) to the present paper. When built environment representation is put in the context of TWA, the large Variety of data and different domain models could be combined with the data stored solely for city modelling. This also means that the system is designed to perform complex analytics operations over such a huge amount of heterogeneous data and this way addresses problems of Velocity and Variability. Moreover, dynamic data source coupled with the web interface capable of reflecting changes immediately in the TWA eliminates the need for the whole model being generated again every time there is an update on the model, noticed by Buyukdemircioglu and Kocaman [8] during the attempt

to work on the CityGML 2.0 model for the whole country in the past.

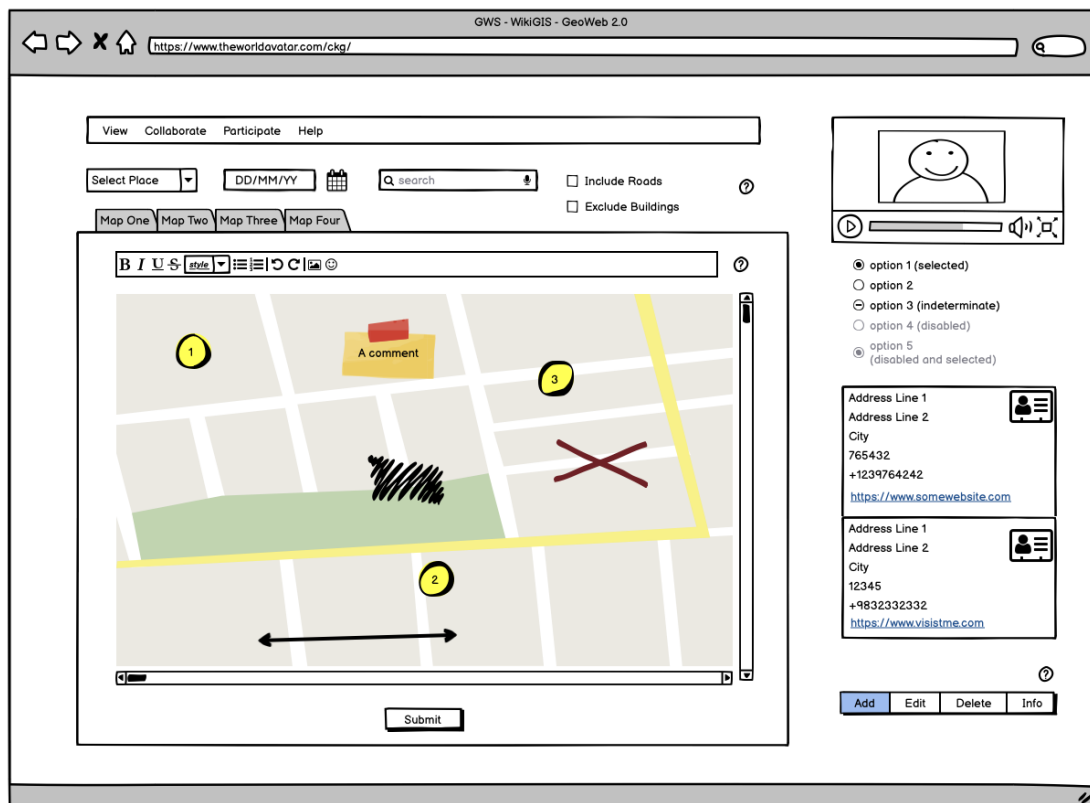
## 2.2 A "View From Everywhere" on the GeoWeb 2.0

Although typical WikiGIS user interfaces are not yet implemented in the TWA, it is worth mentioning that the system would be capable of supporting participatory GIS applications that are the subject of the most recent research in the domain of urban planning [5]. The idea of collaborative GWS has got many advantages in terms of cost and time savings required for data acquisition and preparation of it to be ready for web interactions. However, the development of such systems involves certain risks that are also elaborated on in the present subsection. Knowledge graph specific solutions to those problems, absent in other systems, are presented in section 3.2.

Ideas of WikiGis, GeoWeb 2.0, neogeography and participatory urban design are closely intertwined. Neogeography - a geography without geographers - is only possible with web, mobile and virtual reality-based interfaces enabling community engagement and participation in urban planning. This is also one of the reasons for referring to it as *wikification of GIS* that comes from recognition of the importance of narratives in this domain [28]. Urban simulation models can be used as policy narratives that have the power to persuade and congeal public opinion. Narratives do not have to be factual but can be regarded as 'what if' scenarios that show interconnections between various parts of the model as well as associated processes that occur between them. "The purpose of the narrative is therefore to make the connections to ordinary experience through ordinary language and peel away the presuppositions so as to engage as well as to persuade" [36]. Within TWA, the Parallel World Framework has been already used in the past to create similar simulations altering electrical grid of Jurong Island in Singapore by adding renewable energy sources and optimising the infrastructure for carbon tax [23]. Therefore, narratives could also be regarded as alternative worlds or versions of reality that communicate a story.

The idea of *wikification* of knowledge also entails a shift from a "top-down to bottom-up" approach in terms of knowledge acquisition and generation. Enabling neogeographers' participation in those processes significantly shortens times and reduces costs required for content preparation in non-participatory GIS. In the GeoWeb 2.0 systems knowledge is produced voluntarily and not as requested to be produced by field experts. This way ordinary citizens and their communities become the knowledge stakeholders instead of large institutional or private producers [65]. This shift requires an appropriate set of interfaces in GWS allowing for such content production by individuals. A wireframe that extends the current TWA web interface to allow for such participation is presented in Figure 4. Apart from existing elements and controls for viewing time versioned city models and their various aspects as either actual stages of a particular city or as a narratives illustrating some changes, the system would need to include elements and controls allowing new web stakeholders to create their own geo-narratives and mash them up with custom types of content.

Some researchers, like Foth et al. [27], even argue that only enabling such interfaces would allow the current most pressing problems that involve cities to be addressed. That



**Figure 4:** A wireframe for a WikiGIS user interface. It allows the creation of custom maps with multiple participants able to collaborate. It is possible to enrich geographical information by mashing it up with other multimedia content.

includes climate change as well as social inequality problems that would help get closer to achieving sustainable development goals [67] while planning cities in a modern way, by producing a comprehensive "view from everywhere" and giving a voice even to previously marginalised groups of stakeholders. According to Hudson-Smith et al. [44], the ability of many to engage and interact is the key feature that defines Web 2.0. Controls and interface elements from it should be added to the GWS' in order to make them capable of supporting this functionality as well. This type of web of ordinary web stakeholder empowerment is also seen as a critical success factor of the GeoWeb 2.0 systems [59].

Despite technology obstacles that generally require a higher level of programming knowledge to get started with web-mapping in contrast to 'mash ups' that are built around the more accessible programming interfaces of Web 2.0 tools, neogeographers still show a strong preference for using such toolset over more traditional ones. As noticed by Hall et al. [37] during field experiments, while the participants had the option of recording their assets on supplied paper maps, most opted to use the web tools to select existing features, create new features, make associated comments, and use the software to chat with other participants about features of mutual interest. However, this shortening of 'time to market' of geospatial content on the web 2.0 comes with its own costs. Credibility assessment of crowd-sourced information could be very complex and not always conclusive [25]. Such



systems, similarly to the other Web 2.0 systems, are also open to deliberate misinformation. Although the geo-web stakeholders already provide informal community-based data quality control, the robust mechanisms needed to detect and remove errors, and to build the same level of trust and assurance that national mapping agencies have traditionally enjoyed are absent in such systems so far [34]. At the same time, geographic information is critical as far as local and country-level governance processes are concerned. Opening it to deliberate misinfluencing in systems not offering the mentioned mechanisms poses high levels of misgoverning risks.

### **3 Enabling GeoWeb 2.5 on Dynamic Geospatial Knowledge Graphs**


Whenever a particular narrative reaches its goal of persuading or justifying the implementation of certain policies, it results in some changes to the actual city landscape. Such changed landscapes are then reflected in the original models and form a part of the new base model for the next cycle of narratives that may result in similar changes as well. Because of that, some authors describe digital models that facilitate such recursive interactions using a metaphor of ‘digital mirrors’ [43]. When looked at as narratives, they provide insights into some potential futures. When considered as base models, they reflect the current state of affairs. TWA, as a system that is built around principles of cognitive architecture (apart from a knowledge base that forms a set of models represented in a semantic form), contains a layer of intelligent autonomous agents that work with such models by creating them, analysing, enhancing and discovering new knowledge. Therefore, in such systems, the recursion of changes is also a result of interactions between human and artificial intelligence. Such a blend of actors has also got some properties specific to knowledge graph architecture that help with mitigation of the mentioned GeoWeb 2.0 risks. GWS built on such architectures blending human and artificial intelligence to solve Web 2.0 issues may be already regarded as GeoWeb 2.5 systems.

#### **3.1 Intelligent Autonomous Agency and Knowledge Enhancement**

The layer of intelligent autonomous agents of the TWA is built upon the cognitive architecture of the JPS Agent Framework, previously described by Chadzynski et al. [11]. The agents operate on memories, and their elements are stored in the form of subject-predicate-object triples in the knowledge graph. Those memories form larger structures, such as city models. The framework offers a security mechanism in the form of syntactic validation, preventing invalid agent inputs’ from entering the downstream data processing and knowledge graph. The framework is built around a microservices architecture and a minimal services model, making the components independent of each other and replaceable with minimal impact on the whole system. The agents operating on the knowledge graph can create, analyse, and, enhance existing knowledge without human involvement or supervision in those processes. Cognitive agents, similarly to any other intelligent beings, typically demonstrate the following capabilities: recognition, decision making,

choice, monitoring, execution, action, interaction, communication, as well as acquisition, representation, refinement and organisation of knowledge. As demonstrated in the next three subsections, web interfaces of the TWA allow for the blending of human and artificial intelligence while working on city modelling and knowledge enhancement.

### 3.1.1 Interactive Interfaces to the Thematic Surface Discovery Agent



**Figure 5:** *CityGML Level of Detail (LOD) upgrade performed by the Thematic Surface Discovery Agent at the city level (the bottom part showing Pirmasens in Germany) and the appropriate knowledge enhancement of a single building before and after the transformation (the top part, from left to right). The agent discovered walls, roofs and ground surfaces from the set of surface geometries and assigned them to appropriate thematic surface categories in the knowledge graph. The thematic surfaces are assigned different colours on the user interface and can be interacted with separately, whereas before the enhancement, it was only possible to interact with the whole buildings and not any of their parts.*

The objective of the Thematic Surface Discovery Agent (TSDA) is to upgrade buildings of CityGML LOD1 representation to LOD2 representation by converting the geometry tree describing the exterior shell to an array of CityGML BoundarySurface entities describing the exact same exterior shell geometry, but semantically differentiated into the different surface types. The polygon classification is a result of applying algorithm 1 to a set of CityGML SurfaceGeometry objects with no thematic information. This process is referred to as ‘thematicisation’. Specifically, the TSDA classifies the polygons in the direct surface geometry trees of buildings into wall, roof and ground polygons. It transforms their surface geometry trees to an OntoCityGML-compliant wall surface-, roof surface-

and ground surface-based hierarchy while preserving as much of the original tree structure as possible. The scope of valid inputs are buildings with top-down topography, i.e. theoretically encodable in heightmaps of rooftop (and base) elevation, with no overhangs or interior geometry and well-conditioned winding order, i.e. polygon coordinate sequences consistently indicate the direction of face normals by a right- or left-hand rule. The TSDA is nevertheless executable on buildings of different levels of detail that do not satisfy the topographical requirement but may result in errors. However, this may still be useful for some applications which do not require high precision, e.g. estimating the total area of roof. Detailed implementation details of the agent are included in the Appendix A.

---

**Algorithm 1:** Surface theme identification algorithm for a single building.

---

**Input:** Array of surface polygons of a building,  $S$ .

**Output:** Array of surface-theme pairs,  $R$ .

```

1 begin
2   for  $s_i \in S$  do
3     Project the polygon of  $s_i$  to the Universal Transverse Mercator projection.
4     Calculate the vector area  $\vec{s}_i$  of the UTM polygon by the shoelace formula in each of
       the  $xy$ ,  $xz$  and  $yz$  planes.
5     Calculate the normal  $\hat{n}_i$  by normalising  $\vec{s}_i$ .
6     if  $|\hat{n}_z| < \text{tolerance}$  then
7       Append ( $s_i$ , “wall”) to  $R$ .
8     else if  $\hat{n}_z > 0$  then
9       Append ( $s_i$ , “roof”) to  $R$ .
10    else
11      Append ( $s_i$ , “ground”) to  $R$ .
12    Calculate the average centroid  $r$  of “roof” surfaces.
13    Calculate the average centroid  $g$  of “ground” surfaces.
14    if  $r$  is defined and  $g$  is defined and  $r > g$  then
15      Switch all “roof” entries in  $R$  to “ground” and vice versa.
16  return  $R$ .
```

---

To evaluate the performance of the TSDA in identifying the thematic surfaces of city objects, the agent was validated using the Semantic 3D City Database of Berlin, described by Chadzynski et al. [10, 11]. The process resulted in the report of the accuracy with which the TSDA identifies the roof, wall and ground type bottom-level thematic geometries linked to the thematic surfaces. The validation of the TSDA was performed independently on each of the 12 districts in Berlin namely Friedrichshain-Kreuzberg, Lichtenberg, Charlottenburg-Wilmersdorf, Marzahn-Hellersdorf, Neukoelln, Spandau, Pankow, Tempelhof-Schoeneberg, Steglitz-Zehlendorf, Reinickendorf, Mitte and Treptow-Koepenick.

A namespace called ‘berlin’ was created in the Blazegraph™ workbench, which serves as the triple store. The namespace was created in the quads mode with the geospatial feature enabled. For each district, the data in n-quads format was uploaded to this triple store. A SPARQL update was executed to insert the Coordinate Reference System (CRS) information in the namespace. This was followed by 2 SPARQL updates which served to fix any missing links (if present) in the building and thematicsurface named graphs to link the buildings and thematic surfaces with the actual bottom-level thematic geometries respectively.

For each building in each district, all the thematic surfaces linked to it have an attribute

called `objectId` which can assume three possible values: 33, 34 or 35 representing roof, wall, and ground type thematic surfaces respectively. This information serves as the ground truth.

For the validation purpose, the TSDA is run in the VALIDATE mode, whereby it uses the polygon data of each bottom-level thematic geometry and classifies it into one of the three categories: roof, wall, or ground. The classifications are added into the surfacegeometry named graph as `rdfs:comment` tags of “roof”, “wall”, or “ground”.

A SPARQL query is executed to compare the bottom-level thematic geometry type (roof, wall, or ground) identified by the TSDA with the ground truth `objectId` (33, 34, or 35) associated with those surface geometries. The SPARQL query used to carry out this validation is included in the listing 3.

```

1 PREFIX ocgml: <http://www.theworldavatar.com/ontology/ontocitygml/
   citieskg/OntoCityGML.owl#>
2 BASE <http://127.0.0.1:9999/blazegraph/namespace/berlin/sparql/>
3 SELECT ?srfc_class ?comment (COUNT(*) AS ?count)
4 WHERE {
5   GRAPH <surfacegeometry/>
6   { ?surfaceGeometry ocgml:cityObjectId ?thematicSurface;
7     ocgml:GeometryType ?geometryType.
8     OPTIONAL{?surfaceGeometry rdfs:comment ?comment}
9     FILTER(!ISBLANK(?geometryType)) }
10  GRAPH <thematicsurface/>
11  { ?thematicSurface ocgml:objectId ?srfc_class;
12    ocgml:buildingId ?bldg. }
13  GRAPH <building/>
14  { ?bldg ocgml:objectId ?bldg_class}
15 }
16 GROUP BY ?srfc_class ?comment

```

**Listing 3:** *Thematic surface discovery validation SPARQL query.*

For each district, the results from the above query are tabulated and the combined accuracy of the TSDA on all 12 districts is computed. The correct classification according to the Semantic 3D City Database is shown in Table 1. The combined validation results for all 12 districts in Berlin are shown in Table 2.

**Table 1:** *Correct thematic surface classification according to the Semantic 3D City database.*

surface_class	comment
33	roof
34	wall
35	ground

The TSDA has been applied to transform a LOD1 model of Pirmasens in Germany, described in OntoCityGML, to LOD2. Results of identifying roofs, walls and ground surfaces are presented in Figure 5. The dataset was initially prepared by extruding ground

**Table 2:** *TSDA Validation Results on Berlin dataset consisting of 539 274 buildings, 9 558 218 surface geometries and 2 936 408 thematic surfaces in total.*

surface_class	comment	count	TSD Agent Classification	
33	roof	1457294	correct	99.88%
33	wall	1017	incorrect	0.07%
33	ground	795	incorrect	0.05%
34	wall	4316401	correct	99.99%
34	ground	145	incorrect	0.00%
34	roof	153	incorrect	0.00%
35	ground	701019	correct	99.99%
35	roof	19	incorrect	0.00%
35	wall	24	incorrect	0.00%

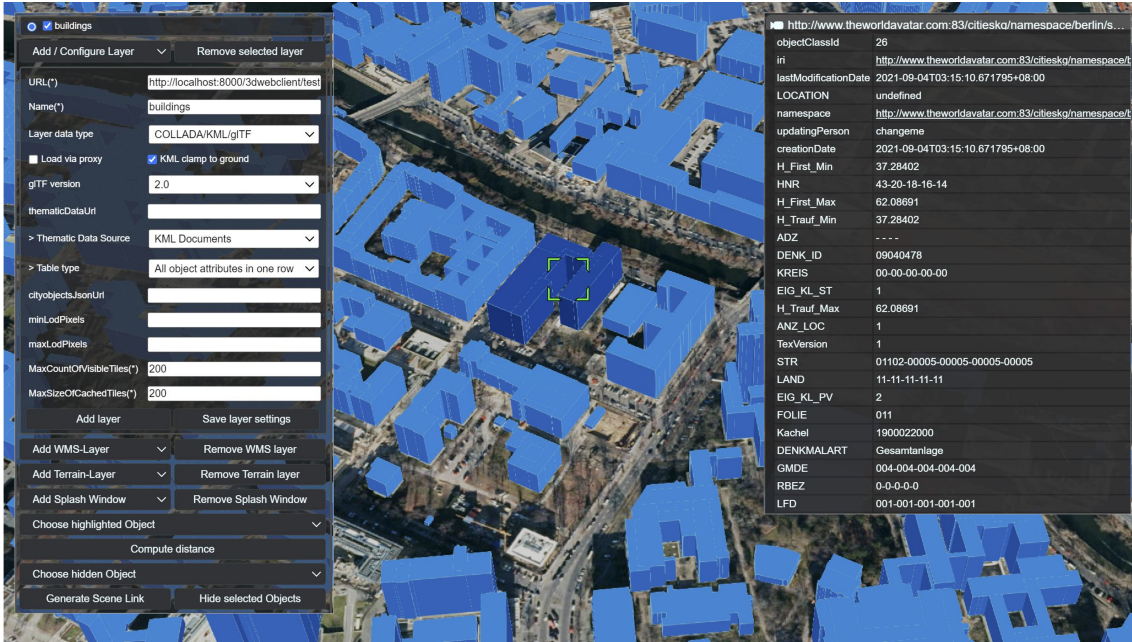
surfaces into buildings in LOD1. In this level of detail, buildings are represented as solids without any thematic information. After transforming the model to the KML format by CExA and visualising it on the augmented 3D web map client in TWA, it is possible to interact only with the whole buildings, by retrieving information out of the TWA that is applicable to them. All buildings are presented as blue geometrical 3D figures of various shape complexity. After the agent broke down such solids into sets of polygons and discovered their spatial orientation, it assigned individual surface polygons into three thematic groups and added thematic surface-specific models to the knowledge graph. When the enhanced dataset was exported into the KML, the individual thematic surfaces started being displayed in separate colours - walls in grey, ground surfaces in green and roof surfaces in red. A richer set of interactions became possible as well. Clicking on individual thematic surfaces makes it possible to retrieve information that is applicable to those surfaces specifically, apart from information concerning the whole building. This may become important in assessing the energy efficiency of the buildings, by examining the thermal insulation properties of walls as well as the solar potential of roofs, etc.

### 3.1.2 Interactive Interfaces to the City Information Agent

City Information Agent (CIA) facilitates contextual interactions with city objects. The agent serves as a general-purpose semantic information retrieval mediator on its own as well as by involving other intelligent agents. In essence, the agent finds, retrieves and provides the required information to the user interface of Web-Map-Client (WMC) upon detecting an interaction event with a particular city object.

The CIA utilises an earlier presented OGM engine to interact with structured graph data. OGM represents OntoCityGML ontology-based knowledge graph classes as Java models and ontological properties or relations as model fields. These models and fields are populated with object-related information from the knowledge graph.

Upon interaction with the city object on the WMC interface, CIA receives an HTTP POST request containing unique object IRI and if provided, context-related information. Then, the CIA creates a JAVA city object class instance using the OGM engine. The created



**Figure 6:** *The City Information Agent (CIA) facilitates contextual interactions with city objects. The agent finds and provides the required information to the user interface upon detecting an interaction event with a particular city object. It contacts any other agents to compile an answer that consists of statements about the city object of interest relevant to the specified context. I.e. it can provide information about urban regulations or energy usage projections that apply to the city object interacted with.*

city object model is populated with data from the knowledge graph using the *pullAll()* function. By executing this function, the graph database is queried for all quads or triples containing the city object model instance’s IRI as the subject or object. Each row of the query response is processed, and the value and datatype are injected into the city object model instance’s corresponding field. The city object model is packaged into the HTTP response body that is displayed on the interface.

Populating city object model instances with values from the knowledge graph may be performed recursively. While non-recursive pulls straightforwardly populate the specified model, the recursive pulls proceed to execute more recursive pulls on model objects linked in the parent level pull. According to the CityGML schema, the city object is linked to *genericAttribute* and *externalReference* cityGML classes. Therefore, when creating a city object Java model counterpart, by setting recursion depth, the *pullAll()* can be used to populate the created corresponding linked object model fields with query results.

CIA also recognises the ‘context’ parameter in the received HTTP request. The parameter typically contains another agent IRI and request body that should be passed. Using this information, the CIA contacts any other agents to compile an answer that consists of statements about the city object of interest relevant to the specified context and display it on the web interface. For instance, it can provide information about urban regulations or energy usage projections that apply to a particular city object. The form in which such

information is currently displayed is presented in Figure 6.

CIA supports synthesis processes in participatory city planning and mediates information between knowledge graphs and end-users. By retrieving all related information about a particular city object, the agent automates otherwise labour-intensive tasks, enabling more resources for design explorations, like public engagements or workshops. Further, individual unique use cases or research questions can be targeted using the CIA context parameter, determining the agent's communication with other agents for retrieving relevant information, e.g. factual as well as simulated with other software. Thus, the synthesis capabilities of the CIA, when integrated together with the visual WMC interface, remove cognitive loads that end users typically experience when confronted with multi-source, unstructured urban data. Thus, inspections of geospatial and related semantic data and engagement with such data may be performed with greater ease, catering to non-expert end-users. Detailed implementation details of the agent are included in the Appendix B. The following two distinctive CIA applications illustrate the agent's synthesis support capabilities.

In the first application, the CIA is used to return site regulation data for plots across Singapore. It creates a JAVA model for the interacted city object - a building plot represented as a 2D polygon - and populates it with retrieved information from the knowledge graph. In such a way, plot zoning type data and Gross Plot Ratio (GPR) are visualised on the web interface. Further planning operations can be performed on it, e.g. querying allowable land uses or programs for the selected city object or comparing estimated Gross Floor Area (GFA) with retrieved GPR.

The second use case makes use of the context parameter to contact the City Energy Analyst (CEA) agent. The CEA agent is used to automate the workflow of running the CEA software [26]. The agent has a run endpoint that, on receiving a request with a building IRI in the body, retrieves the required input data using queries sent to the TWA and then runs a predetermined CEA workflow on the building. Key outputs returned by the CEA are represented ontologically and added to the TWA through automatic SPARQL updates. This removes the requirement for any domain knowledge in how to operate the CEA manually, as well as to understand the output results. The output data added are the building energy demand (grid consumption, electricity consumption, heating consumption and cooling consumption) and the potential solar energy availability of the building (solar panel area, solar panel energy supply).

The CIA can display this information if the CEA Agent query endpoint is added as a context parameter. This means the CIA sends an HTTP request to the CEA agent's query endpoint to retrieve the energy profile information of the building specified by the IRI given in the body of the HTTP request. The returned energy information is then displayed on the web interface alongside the city object information. This brings data from multiple sources together in one access point for a user to get detailed information about buildings. Detailed implementation details of the agent are included in the Appendix C.



**Figure 7:** *The Distance Agent automatically calculates distances between city object representations, which were interacted with on the web map client. It dynamically manifests the acquired knowledge by displaying the learned information about spatial relationships through connection lines and distance values whenever it is ready.*

### 3.1.3 Interactive Interfaces to the Distance Agent

The Distance Agent (DA) is an example of intelligent autonomous agents tracking interactions on the city models and their parts presented on the TWA's web interface. Similarly to the other JPS agents deployed as web services, while operating in the listening mode, the agent waits to receive HTTP POST requests about the city objects interacted with on the web. As presented in the sections 3.1.1 and 3.1.2, web users of the TWA can interact with the whole buildings or their individual thematic parts to retrieve more specific contextual information that applies to a particular object of interest in a given situation. Upon clicking on those objects, their IRIs are sent over to DA in JSON format. This operational mode allows the agent to make sense of its environment and react accordingly to the detected events and processes.

The agent performs syntactic input validation for cybersecurity purposes by checking whether all the IRIs in the request payload are well-formed and conform to the HTTP specification. It responds with the HTTP error 400 to the invalid inputs and halts any further data processing. Valid inputs trigger the agent to operate in the distance calculation mode.

To calculate distances between objects of interest, the agent uses the OGM engine for object instantiation and data retrieval from the knowledge graph. The objects are retrieved using the technology-agnostic Access Agent and Store Router regardless of which data stores and namespaces their target IRIs resolve to. This allows the agent to operate on a highly distributed system without the need to know specific data storage locations for the city objects. Before calculating the distance between particular objects of interest, the agent checks if an existing distance relation is already stored in the knowledge graph.



If the query result is not empty, the distance is returned, and the agent gets back to the listening mode of operation. Otherwise, the agent computes centroids of the objects in the first step. To do that, it also makes sure that the object's envelopes that specify their bounding boxes are encoded within the same coordinate reference system. The agent also performs the coordinate reference systems' conversions in case of finding any discrepancies. After computing centroids of envelopes described in the uniform coordinate reference system, the agent computes distance and adds appropriate new statements to the knowledge graph, describing distance in terms of Units of Measure Ontology (OM) [64]. As depicted in Figure 7, the distances are also returned in the agent's HTTP response to the TWA web interface and displayed there immediately. A more detailed description of the implementation can be found in [11].

The agent operates autonomously and independently from all the other TWA agents. It is only dependent on the Access Agent that provides a technology-agnostic way of accessing the knowledge graph in TWA. However, it can be used for decision support use cases while working with them in tandem. The agent is capable of calculating distances between buildings, and in this way could be used for similar studies to those performed by Huang et al. [42] for urban bicycling suitability but on a larger scale. This way, it would be possible to conduct multi-factor optimisations for less fossil fuel-intensive transportation systems on large city models. The agent could also work with models enhanced by TSDA and assess distances between roofs of buildings to minimise the waste of materials required for solar panels deployment as a potential use case; especially when they are interacted with by means of the CIA working in the context of the CEA, described in the section 3.1.2. When the CIA works in the context of regulations applying to plots, the agent could aid optimal site selection for usages that require distance from others (i.e. chemical processing plants and housing estates) and so on. Such cases show how intelligent autonomous agents of the TWA could lead to adding new knowledge as well as enhancing existing knowledge about cities concurrently with interactions on its web interfaces while performing intelligent operations and analyses by their human counterparts.

### **3.2 Ground Truth Maintenance on Dynamic Geospatial Knowledge Graphs**

Semantic Web Technologies (SWTs) are generally regarded as an essential enabler for Web 3.0. This new Web is envisioned around the ideas of decentralisation, web stakeholder data ownership, privacy, and data security [50]. Blockchain technologies are seen as facilitators for the decentralisation of the Web as well as, because of their origins in cryptography, technologies enhancing security. TWA and the JPS Agent Framework are designed as decentralised systems at the core. In particular, the framework is built upon the idea of a minimal service model and microservices architecture [11]. Moreover, the Agent Composition Framework of TWA [75] is designed to work with a blockchain-based agent marketplace [74]. Both are accessible by interfaces integrated into the JPS Base Library that TWA agents extend and make use of. Compliance with W3C standards makes TWA storage layer a knowledge graph based on the SWTs. This gives TWA good foundations to evolve into the new era of web information systems in the future. TWA contains city modelling subsystem dedicated to representing the domain of cities in a similar form

to other domain representations - by making use of SWTs. Because of using Blazegraph™ triple store and being configured for geospatial search, the TWA is a dynamic geospatial knowledge graph fully compliant with SWTs [10].

The adopted system architecture strongly positions the TWA to be ready for the era of GeoWeb 3.0 systems at some point in the future too. However, as was elaborated more in section 2, GWS systems have not yet fully evolved into GeoWeb 2.0 systems. Moreover, enabling Web 2.0 features on them carries high levels of risk for critical infrastructures to be exposed to malicious activities. Embracing graph structure in SWTs makes them naturally suitable for content mashups that are going to remain an essential part of the next generation of the Web as well [38]. As such, the TWA is a GeoWeb 2.0 ready GWS that only requires slight extensions to its web interfaces to facilitate ordinary web stakeholder participation in geo-content creation, editing and sharing.

OntoCityGML [10] is a transformation of the CityGML 2.0 schema into an ontology serialised in the Web Ontology Language (OWL) and provides a vocabulary and a set of relationships between vocabulary items for representations of city objects in the TWA. The ontology as a conceptual schema for city modelling serves as a Tbox [4] for the Semantic 3D City Database. OWL is a language based on description logic [40] and could be used with automated reasoning technologies. The OntoCityGML ontology implements 344 classes, 272 object properties and 78 datatype properties. It also contains 3,363 implemented axioms. The ontology has been evaluated according to the criteria found in the literature on the subject matter [39]. It passes consistency, coherence, accuracy, conciseness and completeness tests using Protégé [56] ontology editor and Hermit reasoner [33].

Such strong foundations in logic already give the TWA models an advantage over other modelling techniques without schemas checked this way when considering potential GeoWeb 2.0 problems. A consistent, coherent, accurate, concise and complete vocabulary and a set of rules already provide a language for geographical model descriptions that minimise risks of deliberate misinformation due to the obscurity of language used for such descriptions otherwise. All elements of city models are stored as subject-predicate-object statements in the knowledge graph. They form an Abox described in terms of the OntoCityGML Tbox. A particular city model can be regarded as a logical conjunction of all such statements and also assessed as true or false. The truth of such models can be maintained automatically. Apart from JPS Agents layer validation mechanisms enforcing the correctness of data, Blazegraph™ triple store can also be configured for the automated truth maintenance [6] that checks the model whenever any new statements are added, deleted or altered by performing inferencing on the Abox automatically. Whereas web interfaces of the TWA allow for information retrieval, analysis and enhancement by blending artificial and human intelligence, the underlying knowledge graph backend can take care of the truth of the model resulting from such actions. Those architectural properties of the TWA provide certainty of the model assessments. In contrast, existing machine learning techniques, such as spam filtering or sentiment analysis that are used in misinformation battles on Web 2.0 are based on probabilistic algorithms. Therefore, TWA can be regarded as a GeoWeb 2.5 GWS that is capable of avoiding commonly recognised pitfalls of Web 2.0 by maintaining the truth of geographical models automatically and with certainty.

## 4 Conclusions and Future Work

Virtual city models can be regarded as narratives useful to communicate policies as well as to persuade residents to accept changes at the regional and national levels. As cities are one of the biggest contributors to global CO<sub>2</sub> emissions, such models could aid in solving the most pressing problems of the present time. Recent research shows that to address those issues, interfaces to the models should be enabled for public participation and content creation, apart from offering view-only access. This shift is similar to the one that occurred while moving to Web 2.0 in the past. However, the dawn of the new era of GWS, sometimes also referred to as WikiGIS, is not without concerns of opening those systems to already well-known risks. At the same time, geospatial content is critical to the national and regional governing bodies. The majority of existing techniques offer only probabilistic algorithms to protect the content from misinformation, falsehoods and, in turn, policy misinfluencing. This paper presented one of the ways to avoid GWS designs opening geographical content to such malicious activities and making them capable of offering interactive interfaces to the city models.

A high-level overview of the TWA architecture, capable of solving the *five-V* problems in smart city data management, was presented in the section 2. At first, subsection 2.1 demonstrated how interactive web interfaces to large city models are created by the CEA from the raw triples stored within the Semantic 3D City Database. The dynamism of the augmented 3D web client was achieved by implementing a geospatially ordered tiling strategy that allows KML replacement on demand and also reflects any changes to the underlying model. The next subsection (2.2) presented how existing web interfaces to the TWA could be extended to make it act as a GeoWeb 2.0 system. The risks of opening the system to interactive public participation and geospatial content creation without additional safeguards have been highlighted as well. The way of avoiding such risks within dynamic geospatial knowledge graphs architecture that makes them GeoWeb 2.5 systems has been presented in the section 3. Subsection 3.1 presents three agents built using the JPS Agent Framework based on a cognitive architecture and shows their role in knowledge enhancement that occurs as a result of human and artificial intelligence with TWA web interfaces. TSDA, presented in subsection 3.1.1, is capable of autonomously upgrading levels of detail in city models by analysing the spatial orientation of surface geometries that belong to their elements. Such changes allow for interactions with individual parts of the buildings on the web and more fine-grained analyses. The CIA, elaborated on in subsection 3.1.2, can retrieve contextual information that relates to the city objects and view it using web interface elements and controls. It can act with the whole objects or their parts after the models are upgraded by the TSDA. The DA, presented in subsection 3.1.3, tracks interactions with the city objects on the web. This includes objects that are interacted with while the information is provided by the CIA. The agent calculates the distances between those objects and adds new knowledge to the TWA. The last subsection (3.1) elaborates on how the dynamic spatial knowledge graph architecture based on semantic web standards and recommendations by the W3C, as well as the web ontology equivalent of the CityGML 2.0 encoding standard, can protect from deliberate misinformation related to the geospatial content in GWS.

Apart from reusing existing components, this work required the design and development

of the following novel elements:

- A Geospatial Processor and SQL2SPARQL Transformer integrated into the City Export Agent allowing the creation of KML files suitable for web visualisation of the large city models by exporting data from raw triples stored in the Semantic 3D City Database.
- A task to geospatially order KML tiles, integrated into the City Export Agent (CExA), allowing the rearrangement of unordered KML files for dynamic loading on the large city models in an ordered manner.
- An augmented 3D web client that is capable of working with dynamic geospatial knowledge graphs and detecting information required for context processing from IRIs.
- An Object Graph Mapping engine that can automatically instantiate model elements from the knowledge graph into Java objects as well as persist any changes to them resulting from interactions via the TWA interfaces.
- A Thematic Surface Discovery Agent that autonomously upgrades levels of details of buildings by surface geometry analysis and adds thematic information to the knowledge graph.
- A City Information Agent that can retrieve contextual information related to the city objects from the TWA and present it on the web.
- A City Energy Analyst Agent (CEA) that automates the creation of building energy demand projections and works with the CIA to present it on the web, when required.

The presented system architecture and interfaces already allow the TWA to avoid typical Web 2.0 risks. However, to get closer to the GeoWeb 3.0, the current system foundations need to be strengthened even further. Apart from logically evaluated Tboxes that provide languages describing cities and related information, there is a necessity to add components and processes that are capable of similar Abox evaluations performed at scale. The JPS Agent Framework, together with the Agent Composition Framework, provide means for the development of agents responsible for ground truth maintenance by automating existing inferencing engines and autonomously working with the knowledge graph in the background. Making those agents work on a blockchain-based agent marketplace within TWA would also allow for agent-based knowledge provenance tracking, independently of the model related knowledge persistence layer. This future work, moving dynamic geospatial knowledge graphs even beyond GeoWeb 2.5, could be regarded as the next brick in paving the road toward self-sustainable knowledge graphs.

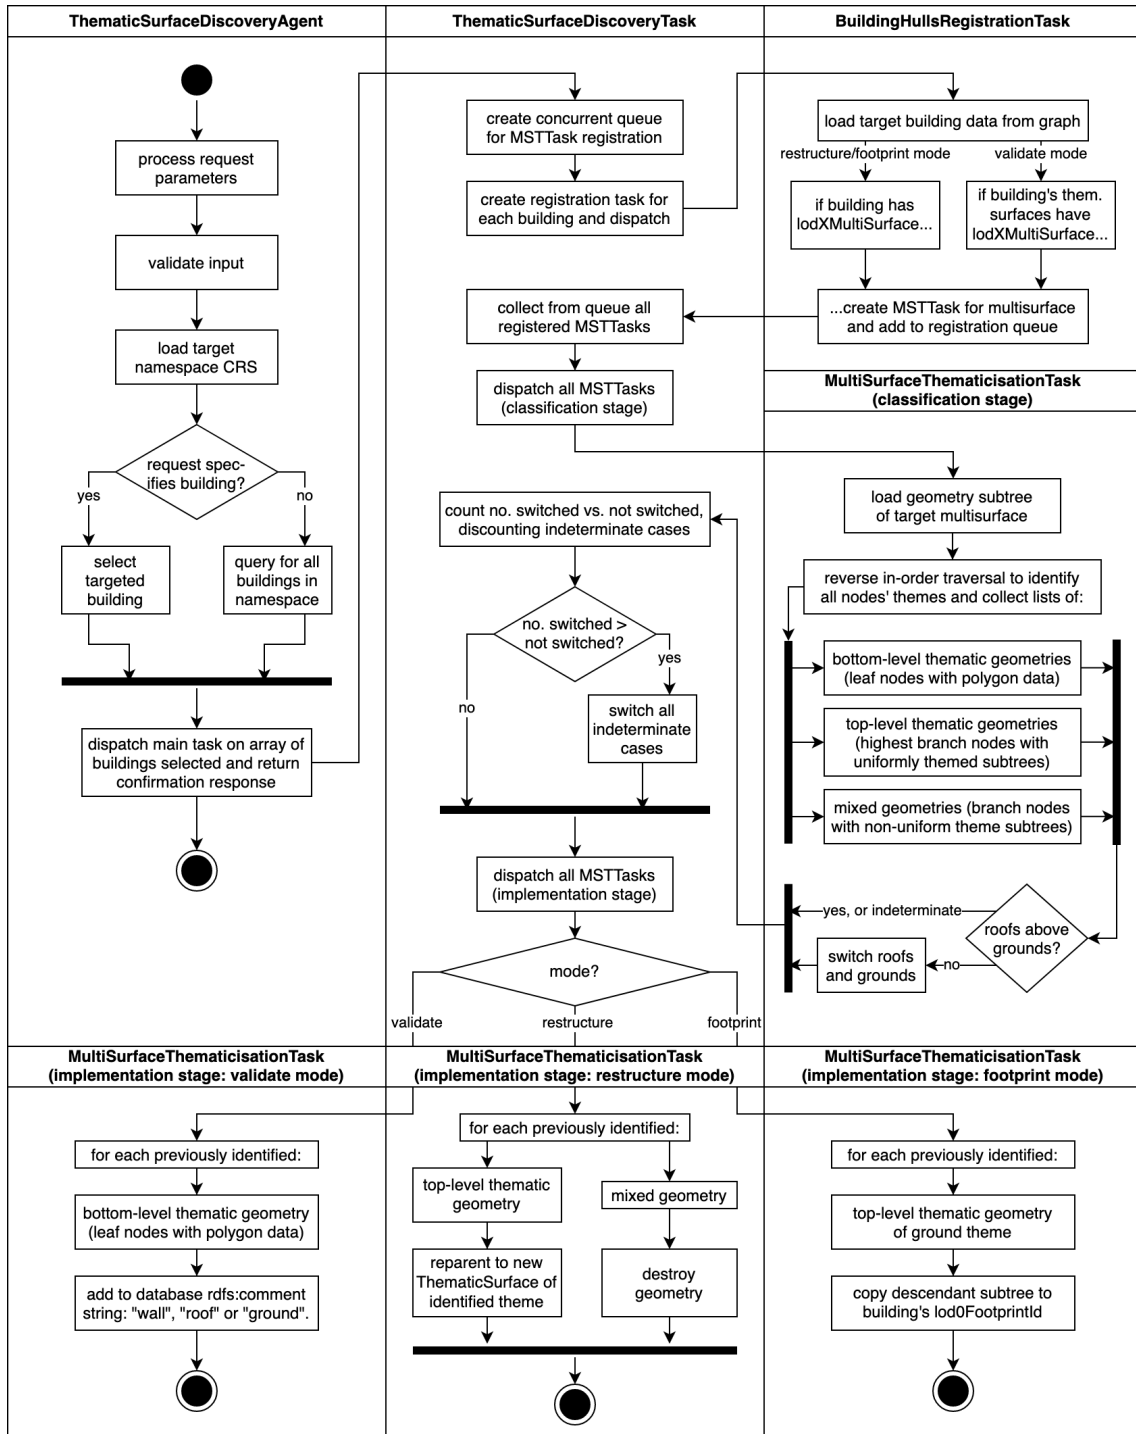
## Acknowledgements

This research is supported by the National Research Foundation, Prime Minister's Office, Singapore under its Campus for Research Excellence and Technological Enterprise

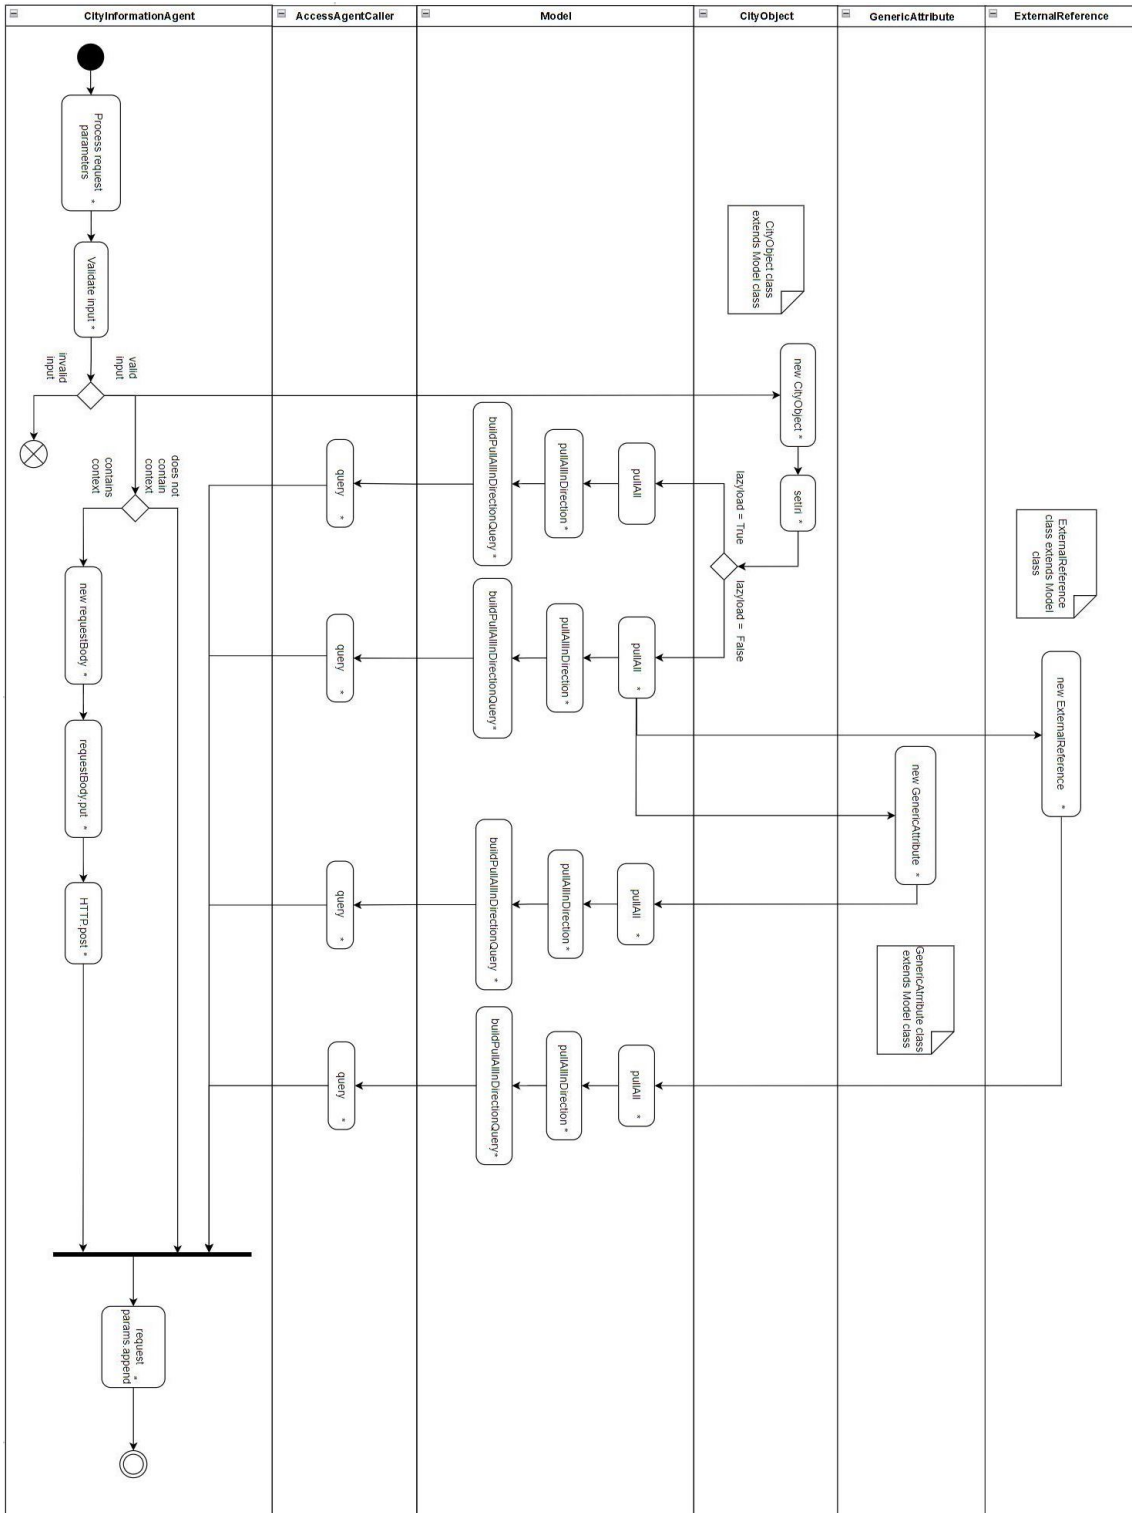
(CREATE) programme. Markus Kraft gratefully acknowledges the support of the Alexander von Humboldt foundation.

The research was conducted as part of an Intra-CREATE collaborative project involving CARES (Cambridge Centre for Advanced Research and Education in Singapore), which is University of Cambridge's presence in Singapore, and Future Cities Laboratory at the Singapore-ETH Centre, which was established collaboratively between ETH Zurich and the National Research Foundation Singapore. The authors would like to thank Andrew Breeson for help with proofreading and language correction of the paper.

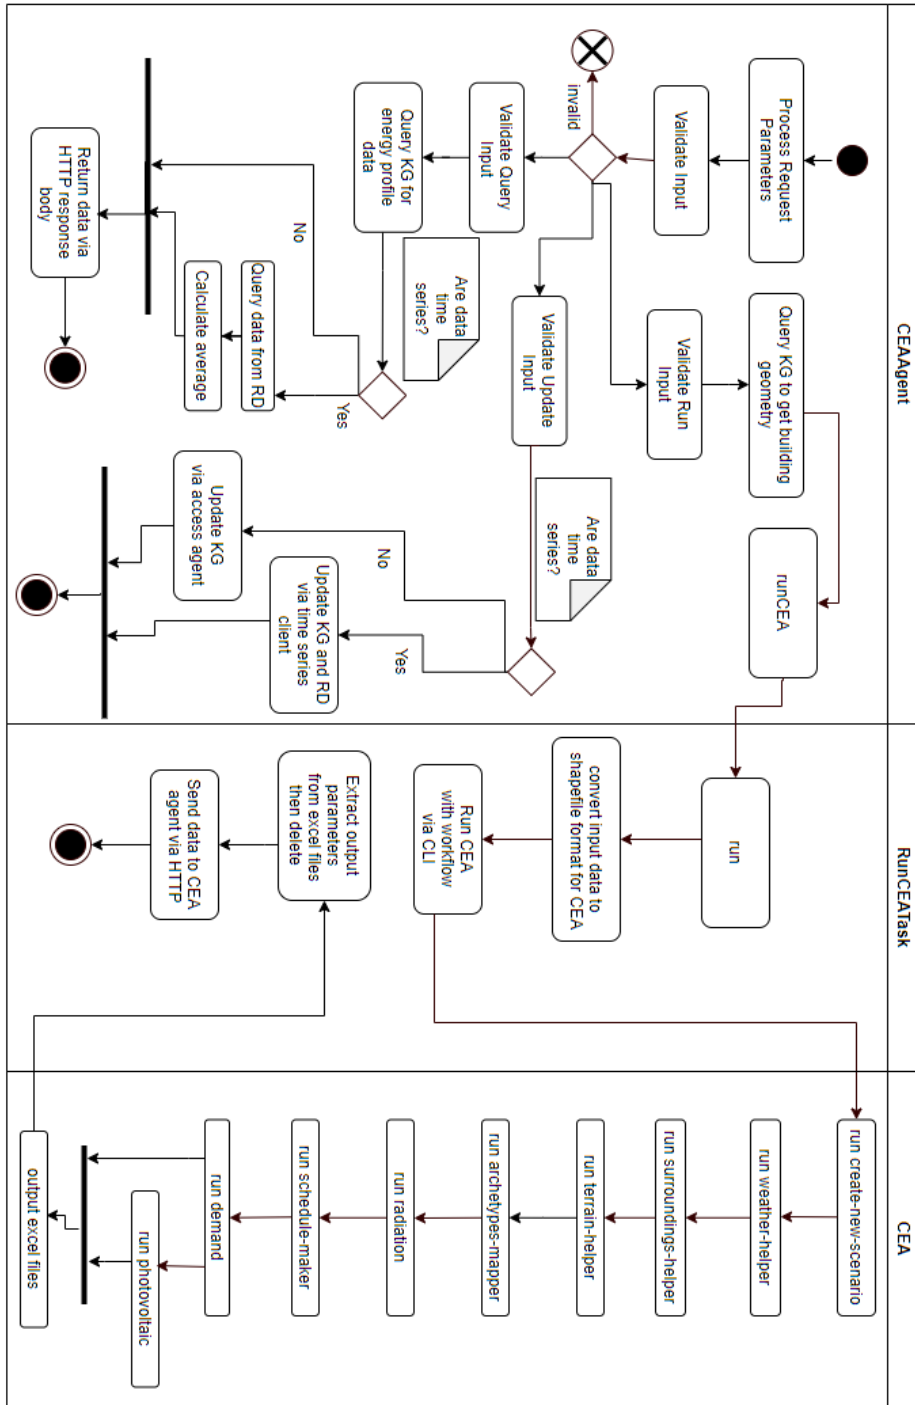
# A Thematic Surface Discovery Agent - UML Activity Diagram



# B City Information Agent - UML Activity Diagram



# C City Energy Analyst Agent - UML Activity Diagram





## List of abbreviations

<b>CARES</b>	<b>C</b> ambridge <b>C</b> entre for <b>A</b> dvanced <b>R</b> esearch and <b>E</b> ducation in <b>S</b> ingapore
<b>CEA</b>	<b>C</b> ity <b>E</b> nergy <b>A</b> nalyst
<b>CExA</b>	<b>C</b> ity <b>E</b> xport <b>A</b> gent
<b>CIA</b>	<b>C</b> ity <b>I</b> nformation <b>A</b> gent
<b>CIM</b>	<b>C</b> ity <b>I</b> nformation <b>M</b> odel
<b>CWA</b>	<b>C</b> losed <b>W</b> orld <b>A</b> ssumption
<b>DA</b>	<b>D</b> istance <b>A</b> gent
<b>DL</b>	<b>D</b> escription <b>L</b> ogic
<b>DT</b>	<b>D</b> igital <b>T</b> win
<b>EIP</b>	<b>E</b> co- <b>I</b> ndustrial <b>P</b> ark
<b>ETH</b>	<b>E</b> idgenössische <b>T</b> echnische <b>H</b> ochschule <b>Z</b> ürich
<b>GML</b>	<b>G</b> eography <b>M</b> arkup <b>L</b> anguage
<b>HTTP</b>	<b>H</b> ypertext <b>T</b> ransfer <b>P</b> rotocol
<b>IRI</b>	<b>I</b> nternationalized <b>R</b> esource <b>I</b> dentifier
<b>JDBC</b>	<b>J</b> ava <b>D</b> atabase <b>C</b> onnectivity
<b>JPS</b>	<b>J</b> - <b>P</b> ark <b>S</b> imulator
<b>JSON</b>	<b>J</b> avaScript <b>O</b> bject <b>N</b> otation
<b>LODx</b>	<b>L</b> evel <b>O</b> f <b>D</b> etail <b>x</b>
<b>OWL</b>	<b>W</b> eb <b>O</b> ntology <b>L</b> anguage
<b>OBDA</b>	<b>O</b> ntology- <b>B</b> ased <b>D</b> ata <b>A</b> ccess
<b>OGC</b>	<b>O</b> pen <b>G</b> eospatial <b>C</b> onsortium
<b>OWA</b>	<b>O</b> pen <b>W</b> orld <b>A</b> ssumption
<b>RDF</b>	<b>R</b> esource <b>D</b> escription <b>F</b> ramework
<b>RDBMS</b>	<b>R</b> elational <b>D</b> atabase <b>M</b> anagement <b>S</b> ystem
<b>SEC</b>	<b>S</b> ingapore- <b>E</b> TH <b>C</b> entre
<b>SQL</b>	<b>S</b> tructured <b>Q</b> uery <b>L</b> anguage
<b>SPARQL</b>	<b>S</b> PARQL <b>P</b> rotocol and <b>R</b> DF <b>Q</b> uery <b>L</b> anguage
<b>TSDA</b>	<b>T</b> hematic <b>S</b> urface <b>D</b> iscovery <b>A</b> gent
<b>TUM</b>	<b>T</b> echnische <b>U</b> niversität <b>M</b> ünchen
<b>TWA</b>	<b>T</b> he <b>W</b> orld <b>A</b> vatar
<b>URL</b>	<b>U</b> niform <b>R</b> esource <b>L</b> ocator
<b>URI</b>	<b>U</b> niform <b>R</b> esource <b>I</b> dentifier
<b>W3C</b>	<b>W</b> orld <b>W</b> ide <b>W</b> eb <b>C</b> onsortium
<b>WKT</b>	<b>W</b> ell- <b>K</b> nown <b>T</b> ext
<b>XML</b>	<b>E</b> xtensible <b>M</b> arkup <b>L</b> anguage

## References

- [1] K. Al Kalbani and A. B. Abdul Rahman. Appropriateness of using CityGML standard version 2.0 for developing 3D city model in Oman. In M. Ben Ahmed, İ. Rakıp Karas, D. Santos, O. Sergeyeva, and A. A. Boudhir, editors, *Innovations in Smart Cities Applications Volume 4*, pages 1332–1343, Cham, 2021. Springer International Publishing. ISBN 978-3-030-66840-2.
- [2] M. Amović, M. Govedarica, A. Radulović, and I. Janković. Big data in smart city: Management challenges. *Applied Sciences*, 11(10), 2021. doi:10.3390/app11104557.
- [3] S. Ates Aydar, J. Stoter, H. Ledoux, E. Demir Ozbek, and T. Yomralioglu. Establishing a national 3D geo-data model for building data compliant to CityGML: Case of Turkey. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, XLI-B2:79–86, 2016. doi:10.5194/isprs-archives-XLI-B2-79-2016.
- [4] F. Baader, I. Horrocks, C. Lutz, and U. Sattler. *An Introduction to Description Logic*. Cambridge University Press, 2017. doi:10.1017/9781139025355.
- [5] M. Batty and W. Yang. A digital future for planning, February 2022. URL <https://digital4planning.com/wp-content/uploads/2022/02/A-Digital-Future-for-Planning-Full-Report-Web.pdf>. Accessed 19th April, 2022.
- [6] B. Bebee. Inference and truth maintenance, 2020. URL <https://github.com/blazegraph/database/wiki/InferenceAndTruthMaintenance>. Accessed 6th May, 2022.
- [7] A. Buccella and A. Cechich. Towards integration of geographic information systems. *Electronic Notes in Theoretical Computer Science*, 168:45–59, 2007. doi:10.1016/j.entcs.2006.08.023. Proceedings of the Second International Workshop on Views on Designing Complex Architectures (VODCA 2006).
- [8] M. Buyukdemircioglu and S. Kocaman. Reconstruction and efficient visualization of heterogeneous 3D city models. *Remote Sensing*, 12(13), 2020. doi:10.3390/rs12132128. URL <https://www.mdpi.com/2072-4292/12/13/2128>.
- [9] Cesium, 2022. URL <https://www.cesium.com>. Accessed 07th April, 2022.
- [10] A. Chadzynski, N. Krdzavac, F. Farazi, M. Q. Lim, S. Li, A. Grisiute, P. Herthogs, A. von Richthofen, S. Cairns, and M. Kraft. Semantic 3D city database — an enabler for a dynamic geospatial knowledge graph. *Energy and AI*, 6:100106, 2021. doi:10.1016/j.egyai.2021.100106.
- [11] A. Chadzynski, S. Li, A. Grisiute, F. Farazi, C. Lindberg, S. Mosbach, P. Herthogs, and M. Kraft. Semantic 3d city agents—an intelligent automation for dynamic geospatial knowledge graphs. *Energy and AI*, 8:100137, 2022. doi:10.1016/j.egyai.2022.100137.

- [12] K. Chaturvedi, Z. Yao, and T. H. Kolbe. Web-based exploration of and interaction with large and deeply structured semantic 3D city models using HTML5 and WebGL. In T. P. Kersten, editor, *Bridging Scales - Skalenübergreifende Nah- und Fernerkundungsmethoden*, 35. *Wissenschaftlich-Technische Jahrestagung der DGPF*, volume 24, Köln, 2015. Deutsche Gesellschaft für Photogrammetrie, Fernerkundung und Geoinformation e.V., Deutsche Gesellschaft für Photogrammetrie, Fernerkundung und Geoinformation e.V.
- [13] K. Chaturvedi, Z. Yao, and T. H. Kolbe. Web-based exploration of and interaction with large and deeply structured semantic 3d city models using html5 and we-bgl. In T. P. Kersten, editor, *Bridging Scales - Skalenübergreifende Nah- und Fernerkundungsmethoden*, 35. *Wissenschaftlich-Technische Jahrestagung der DGPF*, volume 24, Köln, 2015. Deutsche Gesellschaft für Photogrammetrie, Fernerkundung und Geoinformation e.V., Deutsche Gesellschaft für Photogrammetrie, Fernerkundung und Geoinformation e.V.
- [14] H.-c. Chiang. Data modelling and application of 3D cadastre in Taiwan. In *Proceedings of the 3rd International Workshop on 3D Cadastres: Developments and Practices, Shenzhen, China*, 2012.
- [15] Cities Knowledge Graph - CARES, 2021. URL <https://www.cares.cam.ac.uk/research/cities/>. Accessed 15th April, 2021.
- [16] Cities Knowledge Graph - ResearchGate, 2021. URL <https://www.researchgate.net/project/Cities-Knowledge-Graph>. Accessed 15th April, 2021.
- [17] Cities Knowledge Graph - SEC, 2021. URL <https://fcl.ethz.ch/research/research-projects/cities-knowledge-graph.html>. Accessed 15th April, 2021.
- [18] CKG Berlin, 2022. URL <http://www.theworldavatar.com:83/citieskgweb/index.html?city=berlin>. Accessed 07th April, 2022.
- [19] J. Dangermond. GIS and the GeoWeb, 2008. URL <https://www.esri.com/news/arcnews/summer08articles/gis-and-geoweb.html>. Accessed 30th March, 2022.
- [20] A. Devanand, M. Kraft, and I. A. Karimi. Optimal site selection for modular nuclear power plants. *Comput. Chem. Eng.*, 125:339–350, 2019. doi:10.1016/j.compchemeng.2019.03.024.
- [21] A. Devanand, G. Karmakar, N. Krdzavac, R. Rigo-Mariani, E. Y. S. Foo, I. A. Karimi, and M. Kraft. OntoPowSys: A power system ontology for cross domain interactions in an eco industrial park. *Energy and AI*, 1:100008, 2020. doi:10.1016/j.egyai.2020.100008.
- [22] C. Dietz. Geospatial web services, open standards, and advances in interoperability: A selected, annotated bibliography. *ala map and geography round table.*, 2010. URL <https://hdl.handle.net/1969.1/129189>. Accessed 30th March, 2022.

- [23] A. Eibeck, A. Chadzynski, M. Q. Lim, L. K. Aditya, L. Ong, A. Devanand, G. Karmakar, S. Mosbach, R. Lau, I. A. Karimi, E. Y. S. Foo, and M. Kraft. A parallel world framework for scenario analysis in knowledge graphs. *Data-Centric Engineering*, 1:e6, 2020. doi:10.1017/dce.2020.6.
- [24] F. Farazi, J. Akroyd, S. Mosbach, P. Buerger, D. Nurkowski, M. Salamanca, and M. Kraft. OntoKin: An ontology for chemical kinetic reaction mechanisms. *J. Chem. Inf. Model.*, 60(1):108–120, 2020. doi:10.1021/acs.jcim.9b00960.
- [25] A. J. Flanagan and M. J. Metzger. The credibility of volunteered geographic information. *GeoJournal*, 72(3):137–148, Aug 2008. doi:10.1007/s10708-008-9188-y.
- [26] J. A. Fonseca, T.-A. Nguyen, A. Schlueter, and F. Marechal. City energy analyst (CEA): Integrated framework for analysis and optimization of building energy systems in neighborhoods and city districts. *Energy and Buildings*, 113:202–226, Feb. 2016. doi:10.1016/j.enbuild.2015.11.055.
- [27] M. Foth, N. Odendaal, and G. Hearn. The view from everywhere: Towards an epistemology for urbanites, 2007. URL <https://eprints.qut.edu.au/9149/1/9149.pdf>. Accessed 22nd April, 2022.
- [28] M. Foth, B. Bajracharya, R. Brown, and G. Hearn. The Second Life of urban planning? Using NeoGeography tools for community engagement. *Journal of Location Based Services*, 3(2):97–117, 2009. doi:10.1080/17489720903150016.
- [29] M. Fowler. *Patterns of Enterprise Application Architecture*. Addison-Wesley Professional, 2002. ISBN 0321127420.
- [30] GDAL/OGR, 2022. URL <https://gdal.org/download.html#current-releases>. Accessed 25th April, 2022.
- [31] M. Gede. Using Cesium for 3D thematic visualisations on the web. *Proceedings of the ICA*, 1:45, 2018. doi:10.5194/ica-proc-1-45-2018.
- [32] G. Getz. 3D tiles format specification, Jan 2019. URL <http://www.opengis.net/doc/CS/3DTiles/1.0>.
- [33] B. Glimm, I. Horrocks, B. Motik, G. Stoilos, and Z. Wang. Hermit: An OWL 2 reasoner. *Journal of Automated Reasoning*, 53(3):245–269, 2014. doi:10.1007/s10817-014-9305-1.
- [34] M. F. Goodchild. Citizens as voluntary sensors: Spatial data infrastructure in the world of web 2.0. *Int. J. Spatial Data Infrastructures Res.*, 2:24–32, 2007.
- [35] G. Gröger, T. H. Kolbe, C. Nagel, and K. H. Häfele. OGC city geography markup language (CityGML) en-coding standard, 2012. URL <https://www.ogc.org/standards/citygml>. Accessed 11th March, 2021.
- [36] S. Guhathakurta. Urban modeling as storytelling: Using simulation models as a narrative. *Environment and Planning B: Planning and Design*, 29(6):895–911, 2002. doi:10.1068/b12857.

- [37] G. B. Hall, R. Chipeniuk, R. D. Feick, M. G. Leahy, and V. Deparday. Community-based production of geographic information using open source software and web 2.0. *International Journal of Geographical Information Science*, 24(5):761–781, 2010. doi:10.1080/13658810903213288.
- [38] J. Hendler. Web 3.0 emerging. *Computer*, 42(1):111–113, 2009. doi:10.1109/MC.2009.30.
- [39] H. Hlomani and D. Stacey. Approaches, methods, metrics, measures, and subjectivity in ontology evaluation: A survey. *Sem. Web J.*, 1:1–11, 2014. URL <http://www.semantic-web-journal.net/content/approaches-methods-metrics-measures-and-subjectivity-ontology-evaluation-survey>. Accessed February 7th, 2019.
- [40] I. Horrocks. OWL: A description logic based ontology language. In P. van Beek, editor, *Principles and Practice of Constraint Programming - CP 2005*, pages 5–8, Berlin, Heidelberg, 2005. Springer Berlin Heidelberg. ISBN 978-3-540-32050-0.
- [41] W. Huang and L. Harrie. Towards knowledge-based geovisualisation using semantic web technologies: a knowledge representation approach coupling ontologies and rules. *International Journal of Digital Earth*, 13(9):976–997, 2020. doi:10.1080/17538947.2019.1604835.
- [42] W. Huang, K. Kazemzadeh, A. Mansourian, and L. Harrie. Towards knowledge-based geospatial data integration and visualization: A case of visualizing urban bicycling suitability. *IEEE Access*, 8:85473–85489, 2020. doi:10.1109/ACCESS.2020.2992023.
- [43] A. Hudson-Smith, R. Milton, J. Dearden, and M. Batty. Virtual cities: Digital mirrors into a recursive world, 2007. CASA Working Paper 125, available at <https://www.ucl.ac.uk/bartlett/casa/publications/2007/dec/casa-working-paper-125>.
- [44] A. Hudson-Smith, A. Crooks, M. Gibin, R. Milton, and M. Batty. Neogeography and web 2.0: concepts, tools and applications. *Journal of Location Based Services*, 3(2):118–145, 2009. doi:10.1080/17489720902950366.
- [45] Java API for KML - Micromata Labs, 2022. URL <https://labs.micromata.de/java-api-for-kml/>. Accessed 28th January, 2022.
- [46] I. Kotseruba and J. K. Tsotsos. 40 years of cognitive architectures: core cognitive abilities and practical applications. *Artificial Intelligence Review*, 53(1):17–94, Jan 2020. doi:10.1007/s10462-018-9646-y.
- [47] N. Krdzavac, S. Mosbach, D. Nurkowski, P. Buerger, J. Akroyd, J. Martin, A. Menon, and M. Kraft. An ontology and semantic web service for quantum chemistry calculations. *J. Chem. Inf. Model.*, 59(7):3154–3165, 2019. doi:10.1021/acs.jcim.9b00227.

- [48] K. Kumar, H. Ledoux, and J. Stoter. Dynamic 3D visualization of floods: Case of the Netherlands. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, XLII-4/W10:83–87, 2018. doi:10.5194/isprs-archives-XLII-4-W10-83-2018.
- [49] P. Langley, J. E. Laird, and S. Rogers. Cognitive architectures: Research issues and challenges. *Cognitive Systems Research*, 10(2):141–160, 2009. doi:10.1016/j.cogsys.2006.07.004.
- [50] M. Leppington. What is web 3.0? the fundamental guide to the new era of digital data ownership, 2022. URL <https://elastos.info/what-is-web-3-0/>. Accessed 05th May, 2022.
- [51] P. C. Linskey and M. Prud’hommeaux. An in-depth look at the architecture of an object/relational mapper. In *Proceedings of the 2007 ACM SIGMOD International Conference on Management of Data*, SIGMOD ’07, page 889–894, New York, NY, USA, 2007. Association for Computing Machinery. ISBN 9781595936868. doi:10.1145/1247480.1247581.
- [52] Y. Lu. Industry 4.0: A survey on technologies, applications and open research issues. *Journal of Industrial Information Integration*, 6:1–10, 2017. doi:10.1016/j.jii.2017.04.005.
- [53] A. Malhotra, M. Shamovich, J. Frisch, and C. van Treeck. Urban energy simulations using open CityGML models: A comparative analysis. *Energy and Buildings*, 255: 111658, 2022. doi:10.1016/j.enbuild.2021.111658.
- [54] B. Mao, Y. Ban, and B. Laumert. Dynamic online 3D visualization framework for real-time energy simulation based on 3D tiles. *ISPRS International Journal of Geo-Information*, 9(3), 2020. doi:10.3390/ijgi9030166.
- [55] Ministry of the Environment and Water Resources. Zero waste masterplan Singapore, 2019. URL <https://www.towardszerowaste.gov.sg/images/zero-waste-masterplan.pdf>. Accessed 20th September, 2021.
- [56] M. Musen. The Protégé project: A look back and a look forward. *AI Matters*, 1(4), 2015. doi:10.1145/2557001.25757003.
- [57] OGC - KML Standard, 2022. URL <https://www.ogc.org/standards/kml>. Accessed 07th April, 2022.
- [58] OGC Asia Forum - ShowCases, 2021. URL <http://www.ogc.org.tw/en-US/ShowCases>. Accessed 06th October, 2021.
- [59] B. Pak and J. Verbeke. Geoweb 2.0 for participatory urban design: Affordances and critical success factors. *International Journal of Architectural Computing*, 12(3): 283–305, 2014. doi:10.1260/1478-0771.12.3.283.
- [60] A. Perera, K. Javanroodi, and V. M. Nik. Climate resilient interconnected infrastructure: Co-optimization of energy systems and urban morphology. *Applied Energy*, 285:116430, 2021. doi:10.1016/j.apenergy.2020.116430.

- [61] F. Prandi, F. Devigili, M. Soave, U. Di Staso, and R. De Amicis. 3D web visualization of huge CityGML models. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, XL-3/W3:601–605, 2015. doi:10.5194/isprsarchives-XL-3-W3-601-2015.
- [62] W. J. Radermacher. *Official Statistics 4.0: The Era of Digitisation and Globalisation*, pages 119–156. Springer International Publishing, Cham, 2020. ISBN 978-3-030-31492-7. doi:10.1007/978-3-030-31492-7\_4.
- [63] RDF 1.1 N-Quads, 2014. URL <https://www.w3.org/TR/n-quads/>. Accessed 07th April, 2022.
- [64] H. Rijgersberg, M. Van Assem, and J. Top. Ontology of units of measure and related concepts. *Semantic Web*, 4(1):3–13, 2013.
- [65] S. Roche, B. Mericskay, W. Batita, M. Bach, and M. Rondeau. WikiGIS basic concepts: Web 2.0 for geospatial collaboration. *Future Internet*, 4(1):265–284, 2012. doi:10.3390/fi4010265.
- [66] A. Stadler, C. Nagel, G. König, and T. H. Kolbe. *Making Interoperability Persistent: A 3D Geo Database Based on CityGML*, pages 175–192. Springer Berlin Heidelberg, Berlin, Heidelberg, 2009. ISBN 978-3-540-87395-2. doi:10.1007/978-3-540-87395-2\_11.
- [67] UN - Sustainable Development Goals, 2022. URL <https://sdgs.un.org/goals>. Accessed 26th April, 2022.
- [68] United Nations Security Council. Climate change ‘biggest threat modern humans have ever faced’, world-renowned naturalist tells security council, calls for greater global cooperation, 2021. URL <https://www.un.org/press/en/2021/sc14445.doc.htm>. Accessed 20th September, 2021.
- [69] A. von Richthofen, P. Herthogs, M. Kraft, and S. Cairns. Semantic City Planning Systems (SCPS): A literature review. *Journal of Planning Literature*, 2022. doi:10.1177/08854122211068526. In press.
- [70] U. Winkelhake. *Roadmap for Sustainable Digitisation*, pages 127–178. Springer International Publishing, Cham, 2018. ISBN 978-3-319-71610-7. doi:10.1007/978-3-319-71610-7\_6.
- [71] WorldWind Java/NASA WorldWind, 2022. URL <https://worldwind.arc.nasa.gov/java/>. Accessed 28th January, 2022.
- [72] O. Wysocki. Awesome CityGML, 2021. URL <https://github.com/0lo0cki/awesome-citygml/>. Accessed 17th September, 2021.
- [73] Z. Yao, C. Nagel, F. Kunde, G. Hudra, P. Willkomm, A. Donaubaauer, T. Adolphi, and T. H. Kolbe. 3DCityDB - a 3D geodatabase solution for the management, analysis, and visualization of semantic 3D city models based on CityGML. *Open Geospatial Data, Software and Standards*, 3(1):5, May 2018. doi:10.1186/s40965-018-0046-7.

- [74] X. Zhou and M. Kraft. Blockchain technology in the chemical industry. *Annual Review of Chemical and Biomolecular Engineering*, 13(1):null, 2022. doi:10.1146/annurev-chembioeng-092120-022935.
- [75] X. Zhou, A. Eibeck, M. Q. Lim, N. B. Krdzavac, and M. Kraft. An agent composition framework for the J-Park Simulator - a knowledge graph for the process industry. *Computers & Chemical Engineering*, 130:106577, 2019. doi:10.1016/j.compchemeng.2019.106577.