

# Making Urban Energy Use More Intelligible Using Semantic Digital Twins

Sander R. de Meij<sup>1,\*</sup>, Alex J.A. Donkers<sup>1</sup>, Dujuan Yang<sup>1</sup> and Matthijs Klepper<sup>2</sup>

<sup>1</sup>Eindhoven University of Technology (TU/e), Groene Loper 6, Eindhoven, 5412AZ, The Netherlands

<sup>2</sup>Royal KPN N.V., Wilhelminakade 123, 3072 AP, Rotterdam, The Netherlands

## Abstract

There is great potential in urban energy modeling for mitigating the effects of increasing energy consumption in cities. However, there is limited integration of traditional building information and urban data in general. Therefore, this project suggests a novel data integration structure, the Neighborhood Energy Ontology (NEO). This ontology aims to connect urban data from different domains and scales to provide more intelligible insight to the end user. In order to assist with this goal, a dashboard was created which allows the end-user to interact with the data and come to new insights. It is suggested that the created ontology, in combination with the dashboard, is a suitable proof-of-concept to show how semantic solutions can aid in improving the potential of urban energy modeling to mitigate the adverse effects of increasing urbanization.

## Keywords

Semantic digital twin, Linked Data, Urban energy usage, Neighborhood Energy Ontology, Dashboard

## 1. Introduction

Growing urban populations [1] cause increased energy consumption in cities. As a result, cities currently consume ‘two-thirds of primary energy resources and are responsible for more than 70% of Green House Gas emissions worldwide’ [2]. Buildings in these cities account for 40% of global energy consumption [3]. In order to address this issue, there is great potential in urban energy modeling which can result in increased energy use efficiency on an urban and building level [4]. However, according to Curry et al. [5], there is limited integration of traditional building information and other data, such as energy consumption. This limits possibilities for cross-domain monitoring, simulation, and interventions. More readily available information could indeed facilitate the identification of problems and solutions with respect to urban energy consumption [6]. Moreover, while many studies focus on integrating data on the building scale [7, 3, 5, 8], most goals for reducing energy use and Green House Gases are set on a national level, and most action is taken at the city scale [9]. Abbasabadi and Mehdi Ashayeri [10] describe the fundamentally different approaches in assessing urban energy use. On the highest level, methods can be divided into top-down and bottom-up approaches, where ‘the top-down approach focuses on the macro scale and treats the built environment as


*Proceedings LDAC2023 – 11th Linked Data in Architecture and Construction, June 15–16, 2023, Matera, Italy*

\*Corresponding author.

✉ s.r.d.meij@tue.nl (S. R. d. Meij); a.j.a.donkers@tue.nl (A. J.A. Donkers); d.yang@tue.nl (D. Yang); m.klepper@kpn.com (M. Klepper)



© 2023 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

 CEUR Workshop Proceedings (CEUR-WS.org)

a whole energy user, without taking into account individual end-users. It relies on historical aggregated energy data to understand energy consumption in cities. In contrast, the bottom-up approach takes a localized approach to studying energy use and considers urban attributes at the microscale of individual units, such as individual buildings or a collective set of buildings. This approach estimates energy use for individual end-users and extrapolates it into regional and national scales' [10]. The aim of this project is to adapt semantic web technologies to integrate cross-domain data and information on multiple scales, enabling both top-down and bottom-up urban energy assessment. Moreover, the project aims to visualize this integration in order to allow stakeholders to interpret and work with the data more intuitively. Placing this effort in the current digital twin paradigm, the definition provided by VanDerHorn and Mahadevan [11] should provide some insight, as they describe a digital twin as 'a virtual representation of a physical system (and its associated environment and processes) that is updated through the exchange of information between the physical and virtual systems.' The goal of this project is to take the initial steps toward such a digital twin. In order to show the potential of this approach, firstly, an ontological structure is defined (section 2). After this, a visual dashboard is presented that allows stakeholders to read and interact with the linked data. This dashboard is tested using the city of Eindhoven (the Netherlands) (section 3). This use case integrates publicly available data from multiple stakeholders.

## 2. Method

Before describing the newly proposed ontological structure named Neighborhood Energy Ontology (NEO), existing ontologies should be considered. The study by De Nicola and Villani [12] gives a preliminary overview of available ontologies related to several urban topics. Regarding energy use, they identify several ontologies, however, they are unsuitable for the purposes of this research. They lack the possibility to describe urban data on a neighborhood level, as they mostly relate to other urban units like microgrids [13], houses [14] or appliances [15]. Moreover, the authors give an overview of available ontologies describing urban systems, which are not aligned with the specific goal of this project as they mostly describe specific urban infrastructure or three-dimensional geospatial objects. Besides these ontologies, SAREF4CITY [16] can be considered. This ontology focuses on extending SAREF [17] in order to create a common core of general concepts for smart cities and data-oriented to the Internet of Things (IoT) field [16] and describes a data structure that allows for the description of several city objects, their geographical definition and corresponding measurements. While this ontology does reflect the core idea of NEO, it is deemed unsuitable for the purposes of this project, as it is more aimed at structuring IoT data and Key Performance Indicator (KPI) measurements. Moreover, NEO tries to capture urban data, without relying on a geographical definition as they are considered hard to work with and impractical for practical implementation. However, alignment between SAREF4CITY and NEO could be achieved (see section 4). Secondly, the Energy Management and Key Performance Indicator (EM-KPI) ontology can be considered. This ontology is created to describe the relationship between the master data sources for identifying energy performance problems and key areas for improvement and to help energy managers make informed decisions regarding energy efficiency measures [18]. Again, this ontology has similar goals and structure

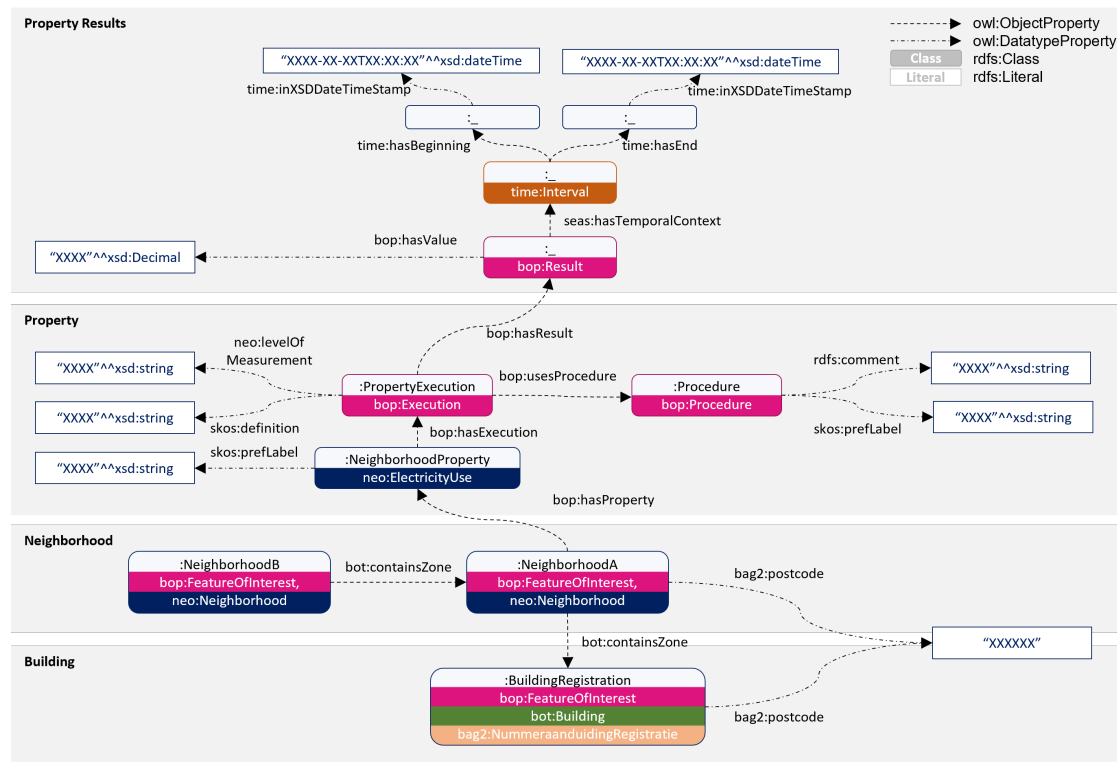
to NEO, however, is deemed unsuitable for the purposes of this project as it has an extensive focus on KPI measurement, energy systems, and building aspects, while NEO is focused on purely urban data. Moreover, as will be explained below, the central concepts of NEO are the Building Performance Ontology (BOP, <https://w3id.org/bop>) and Building Topology Ontology (BOT, <https://w3c-lbd-cg.github.io/bot/>). These ontological structures are considered to be able to describe building-level information and data in high detail and are therefore extended in this project. As the previously described ontologies are not connected to either BOT or BOP, they are not reused in this project. Therefore, a new ontological structure has been created named Neighborhood Energy Ontology (NEO), which reuses and extends multiple existing data structures. NEO tries to achieve the previously described goals by defining ‘neighborhoods’ as urban areas, which can contain other urban areas of a different (smaller) scale. These neighborhoods are linked to certain properties that are attributable to these areas, following a similar structure as defined in the Building Performance Ontology. In this project, neighborhoods are considered a ‘bop:FeatureOfInterest’ and therefore can be associated with a ‘bop:Property’. This structure is given in Figure 1 (namespaces are defined in Table 1). In this overview, it is shown how neighborhoods can contain other neighborhoods, of a different scale. NEO can therefore describe data on multiple levels, where a contained neighborhood might be assumed to inherit the properties described by the containing neighborhood. Moreover, multiple properties can be attributed to a neighborhood, which might come from different domains. Therefore, a more holistic description of urban data can be given (examples of different domains are given in section 3 below). In order to create a high-level structure of these properties, the property structure in Figure 2 is adopted. This figure shows that a ‘neo:Neighborhood’ is a sub-class of a ‘bot:Zone’, which allows for the previously described relationship where neighborhoods (zones) can contain other neighborhoods (zones). Moreover, neighborhoods can thus contain buildings (bot:Building) as shown in Figure 1. The anchor point for building information in NEO is the existing data structure of the Cadastre, Land Registry and Mapping Agency of the Netherlands. Their Key Register for Addresses and Buildings (Dutch acronym: BAG) is published as linked data and knowledge related to buildings, public spaces, cities is captured in the BAG2 ontology. The building registration is reused in this project (Figure 1). This information is captured in the ‘Basisregistratie Adressen en Gegevens’ ontology (BAG2 [19]), which is reused as the basis for capturing data on buildings in this project (Figure 1, ‘Building’). While this ontology entails a larger structure, only the building registration aspect is shown in this Figure as it is the core element used in this project.

Properties of the neighborhood are linked to a bop:Execution and bop:Procedure. As a single property can be measured (and therefore defined) through multiple methods, this structure allows for these differences to occur in the data. For example, an area’s population can be measured by multiple methods, and will likely deviate across different datasets. These measurements (of type bop:Execution) are linked to the same property but describe different values measured by different procedures. Conversely, different properties (e.g. the populations of multiple areas) might apply similar measurement procedures. Lastly, each execution of a property can have one or multiple results which provide a value (Figure 1, ‘Property Results’). Moreover, each result is assigned a specific time interval that denotes its temporal relevance. Thus it can be said that the result is only relevant within the time interval attributed to it. To summarize, neighborhoods of varying scales may possess multiple properties which can be measured or assessed through

**Table 1**

Prefixes and corresponding namespaces used in Neighborhood Energy Ontology

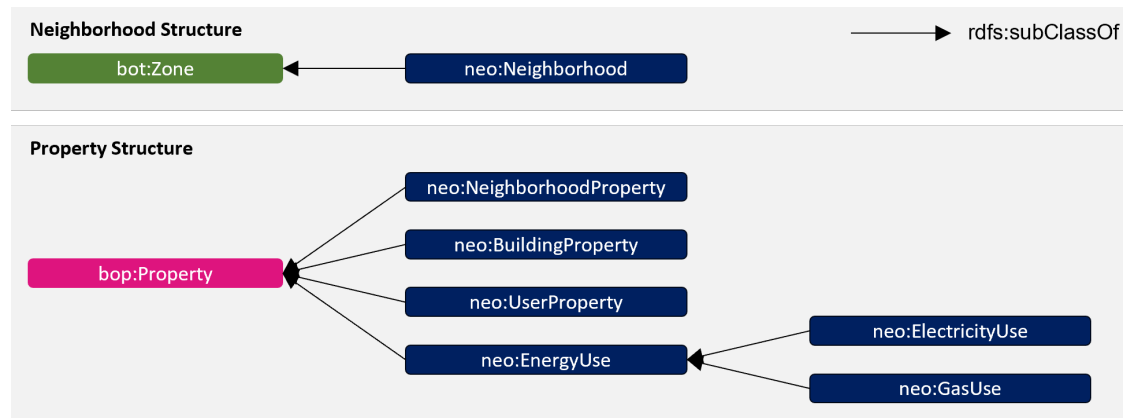
Prefix	Namespace	Color Representation
bag2	https://bag2.basisregistraties.overheid.nl/bag/def/	Orange
bot	https://w3id.org/bot#	Green
bop	https://w3id.org/bop#	Pink
time	http://www.w3.org/2006/time#	Brown
neo	https://sanderdemeij.github.io/neo/	Blue
skos	http://www.w3.org/2004/02/skos/core#	-
rdfs	http://www.w3.org/2000/01/rdf-schema#	-
xsd	http://www.w3.org/2001/XMLSchema#	-
seas	https://w3id.org/seas/EvaluationOntology#	-



**Figure 1:** Overview of data integration method

multiple executions with different procedures. These executions may yield multiple results, each with a designated time interval that reflects its temporal relevance.

In this project, data from multiple domains is being collected to provide a more intelligible overview for end-users. However, these datasets which describe data about the same urban areas (neighborhoods) are often stored in separate data silos with different data owners, leading to a lack of cooperation and connection. To connect these datasets, a concrete sample of the



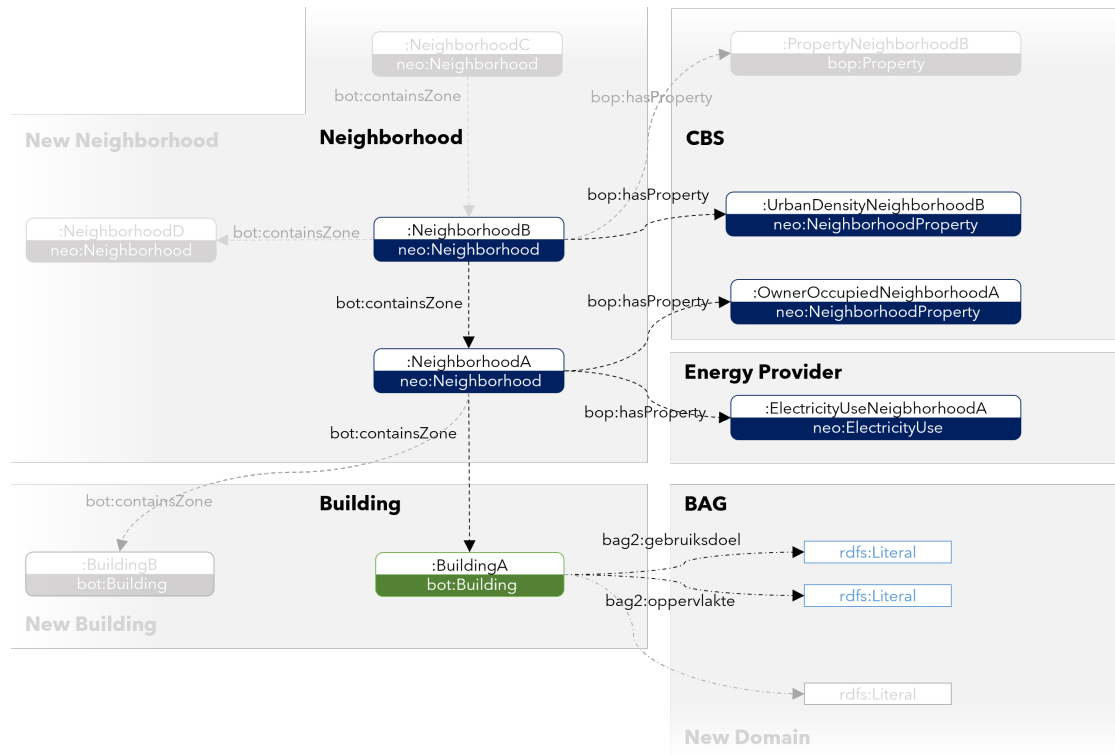
**Figure 2:** Class structure of Neighborhood Energy Ontology (NEO)

integrated data mentioned above is in Figure 3. This figure illustrates how data from multiple domains is connected on multiple scale levels. Similarly, a sample of this data in Turtle format is shown in Figure 4. While Figure 3 is a visual representation, 4 is a machine-readable example of the data which is used in section 3. By adopting this method, the data is made accessible for further analysis and interpretation by different stakeholders (an example of this will be given in section 3.3). Moreover, in future development, different scales and domains can be added easily (which will be compatible with the designed dashboard).

## 3. Results

### 3.1. Dashboard Design

In order to achieve the goal of more insightful urban energy use data, a web-based viewer is created that visualizes the integrated data from multiple stakeholders. The back end of the viewer consists of two databases. First, a graph database (Ontotext GraphDB [20]) is used to store the linked data. This project incorporates data from four main sources: (1) energy use data obtained from a Dutch energy provider [21]; (2) socio-economic data collected from several sources of the Dutch Central Bureau of Statistics (CBS) [22]; (3) energy label data procured from the Dutch Enterprise Agency (Dutch acronym: RVO) [23], and (4) building-related data gathered from the BAG data [19]. A Python converter has been developed to convert tabular data into RDF data. The converter can integrate data from multiple datasets as long as the tabular data has a postal code identifier. It is a helpful addition to the dashboard described above, which can accommodate any data adhering to the data structure outlined in section 2. The aforementioned data is used as an example in this project, but the converter allows for the easy extension of new datasets in the future. The geometric city information is typically unsuitable for graph databases, which is why Cesium Ion [24] is used to store and stream the geometric data. This data consists of both 2D and 3D geometries. The dashboard, a web application in JavaScript, visualizes the geometric information. This data is linked with the linked data using GUIDs. The web application lowers the entry barrier for end-users and can be easily built upon.



**Figure 3:** Overview of a sample of the integrated data in visual and code form

The dashboard serves two main functions, visualization, and querying. End-users are able to build visual queries, just like the common web shop filter bars (see Figure 5). These queries are transformed into SPARQL. Expert users can choose to type SPARQL queries manually. The results of these queries will be visualized in the 3D map, in a table, and in a graph. The 3D map is created using the Cesium package [25], and the graph is constructed using vis.js [26]. All these items are interactive so that if a user clicks on a certain neighborhood in the map, on a row in the table, or on an element in the graph, the SPARQL query will be automatically updated and new results will pop up. Users are therefore not limited to querying functionalities but can explore the available data via different intuitive methods. These functionalities are shown in Figure 5, where the map is shown in the top left, while the query is shown in the top right. The table and graph are shown in the bottom left and right respectively. The remainder of this section will show these elements in more detail and show the different forms they can take (which correspond to different functionalities).

As has been mentioned, the map is created using the Cesium package [25], which allows for a multitude of geospatial representations. Cesium was chosen specifically in order to be able to visualize 3D shapes on the map. The main map displays the city under investigation (Eindhoven, the Netherlands) to the end-user, including all 6-digit postal code areas in this city. The user is able to zoom, pan, and rotate the view to investigate all aspects of the city. Moreover, within the map, the user has the ability to click on the postal code areas, which gives more information on

```

<http://example.org/NeighborhoodB> a neo:Neighborhood,
    bop:FeatureOfInterest ;
    bag2:postcode "XXXXB"^^xsd:string ;
    bop:hasProperty <http://example.org/UrbanDensityNeighborhoodB> ;
    bot:containsZone <http://example.org/NeighborhoodA>.

<http://example.org/NeighborhoodA> a neo:Neighborhood,
    bop:FeatureOfInterest ;
    bag2:postcode "XXXXAA"^^xsd:string ;
    bop:hasProperty <http://example.org/ElectricityUseNeighborhoodA> ,
        <http://example.org/OwnerOccupiedNeighborhoodA> ;
    bot:containsZone <https://bag2.basisregistraties.overheid.nl/bag/id/registratie/BuildingA>.

<https://bag2.basisregistraties.overheid.nl/bag/id/registratie/BuildingA> a bop:FeatureOfInterest,
    bot:Building ;
    bag2:gebruiksdoel "XXXX"^^xsd:string ;
    bag2:oppervlakte "XX"^^xsd:string .

<http://example.org/UrbanDensityNeighborhoodB> a neo:NeighborhoodProperty ;
    skos:prefLabel "urbanDensity"^^xsd:string ;
    bop:hasExecution <http://example.org/UrbanDensityNeighborhoodBProcedureA>.

<http://example.org/UrbanDensityNeighborhoodBProcedureA> neo:levelOfMeasurement "XXXX"^^xsd:string ;
    skos:definition "XXXXXXXXXXXX"^^xsd:string ;
    bop:hasResult [ a bop:Result ;
        bop:hasSimpleUnit "XXXX"^^xsd:string ;
        bop:hasValue 1000.0 ;
        seas:hasTemporalContext <http://example.org/UrbanDensityTemporalContextA> ],
        [ a bop:Result ;
        bop:hasSimpleUnit "XXXX"^^xsd:string ;
        bop:hasValue 1000.0 ;
        seas:hasTemporalContext <http://example.org/UrbanDensityTemporalContextB> ] ;
    bop:usesProcedure <http://example.org/ProcedureA> .

<http://example.org/UrbanDensityTemporalContextA> a time:Interval ;
    time:hasBeginning [ time:inXSDdateTimeStamp "XXXX-XX-XXTXX:XX:XX"^^xsd:dateTimeStamp ] ;
    time:hasEnd [ time:inXSDdateTimeStamp "XXXX-XX-XXTXX:XX:XX"^^xsd:dateTimeStamp ] .

<http://example.org/ProcedureA> a bop:Procedure ;
    skos:prefLabel "XXXX"^^xsd:string ;
    rdfs:comment "XXXXXXXXXXXX"^^xsd:string .

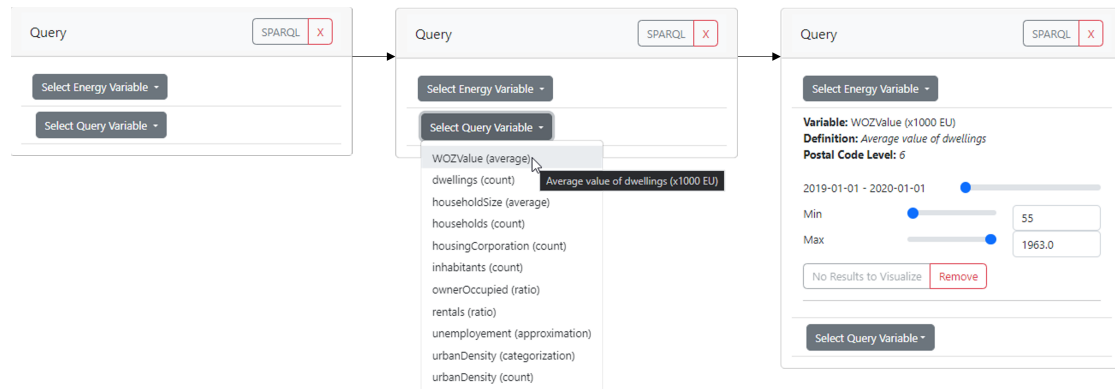
```

**Figure 4:** Sample of integrated data in Turtle format

the selected neighborhood. Firstly, an automatic query is generated which queries all variables associated with this postal code area. This allows the user to investigate the extent of the data available for this area, and get an initial sense of the area. Secondly, a 3D visualization of the buildings within the area is shown. These buildings are part of the 3D BAG dataset [27] and have Level of Detail (LoD) 2. The buildings are available for the entire city, however, buildings are only shown based on the selected postal code area, due to load time restrictions. The map can be considered the core of the dashboard. Moreover, all other functionalities will interact with the map in some way in order to provide the end user with a clear overview of the data. While the map can be considered the core element of the dashboard, the query functionality adds the first layer of interactive capabilities. As is shown in Figure 5, the user has the option to select energy use variables, as well as other variables, in order to construct a query. These variables are automatically generated based on the available instances of bop:Property in the graphs. Moreover, the user is able to select multiple variables in order to construct more complex queries (section 3.2) and visualize the results of the query on the map. Lastly, the user is able to visualize the data in tabular and graph form. The table represents data in common form and allows the user to select individual neighborhoods for further investigation. Moreover, the graph visualizes the structure of the data to the user. Here, the user can investigate the data on







**Figure 6:** Process of constructing a query using the dashboard

existing questions or formulate new questions which the user was unable to form before.

### 3.2. Querying

The structure for constructing a query is provided in Figure 6. Firstly, the user can select one of the available properties (with different execution methods). This list of properties is created dynamically based on the available data. Secondly, the user selects a property to query, after which the property is represented correctly. The representation method is based on the level of measurement indicated for the execution method (Figure 1, 'Property') and a high-level check of the found values. Based on whether the data is numerical, categorical, or nominal, a different user interface is generated to build the visual query method. In the example provided in Figure 6, the data is numerical, and therefore the user is represented with means to query this data accordingly. Moreover, the user is able to select multiple properties to run a complex query. Meaning that individual SPARQL queries are constructed and executed, which all return the correct postal codes and their values. When multiple variables are queried the overlap between the found postal codes is established, and these postal codes are shown on the map. When the user chooses to visualize the results, the found values for that variable are converted to five bins, each represented by a color (as shown in Figure 5). This allows a more intuitive way to explore the spread of the found values over the city. The user can switch between the visualization of all the found variables, however, the visualization of multiple variables at once is not available as the interpretation of this visualization would be too complex.

### 3.3. Generating Urban Energy Insights

In order to show the potential of the data structure and the dashboard, a relevant use case is explored. In order to cope with the increasing energy prices, the Dutch government has initiated a maximum price for energy [28]. Part of this policy is the fact that energy price is only limited when energy use remains under a certain limit annually. For gas, this limit is set at 1200 m<sup>3</sup>. Given this policy, the municipality (or another stakeholder concerned with urban policy) might be interested to know where in the city this consumption is exceeded. Given

the provided data, the end-user can start an initial investigation of the data by querying for all postal codes where the average gas use exceeds these limits, as is done in Figure 5. Moreover, it is suggested that persons who rent their homes are more vulnerable to energy poverty and therefore increasing energy prices [29, 30]. Therefore, the query is extended to areas with a high percentage of ‘renters’. Lastly, in order to exclude areas where few dwellings are present (and are more likely to include industry), a query is included to search for areas with more than 50 dwellings. The results of this query are shown in Figure 5 and indeed show that there are multiple areas in the city under investigation that might be relevant for further investigation. For example the larger area in the south of the city, as well as the cluster of areas in the north-west. The end-user might want to go into more detail on why gas use is relatively high in these areas, however, this dashboard has provided a first step in the process. As Figure 3 shows, data from multiple stakeholders is (and can) be used in this use case, while the possibility exists for further expansion of the data in further conducted research into these areas of interest. The graph function shown in Figure 5 allows the user to see which data is already available for the areas under investigation. As an example, the end-user might collect more detailed building information in these areas to further the investigation. While this type of data collection is infeasible for the entire city, it is possible for only a few areas. What should also be noted is the fact that a multitude of areas are marked gray, indicating that no data is available for these areas for one of the variables as indicated by the provider of the data. As the provider of the data has indicated the data as unavailable, it is unclear whether the data does not exist, is zero, or is kept secret (for instance for privacy reasons). During the addition of the variables, it became clear that these missing values are attributable to the data on rented dwellings. While this might limit the validity of the analysis, it might be useful information to the end-user and indicate several steps to be taken in order to do a more thorough investigation of this problem.

## 4. Discussion

Urban energy modeling has great potential to support energy efficiency in buildings and the urban environment. However, the lack of data integration across different scales and domains has limited this potential. As an example, the paper by Ali et al. [6] describes a seven-step process ‘to analyze and visualize large dataset and extract meaningful information from the data’. To address this issue, and make the extraction of meaningful information much easier, this project has proposed a new data structure that enables the integration of multiple relevant data sets into urban energy modeling. Moreover, we proposed an interactive visual dashboard, which allows end-users to investigate the data in more detail and gain insights that were previously difficult to obtain. As the proposed data structure is only demonstrated in the context of urban energy use, more research is needed to test its suitability for other urban properties. In addition, the data currently used is rather static data (yearly or monthly). More tests are needed to integrate dynamic data, such as real-time sensor data. Here, alignment with currently existing ontologies like SAREF4CITY and EM-KPI (section 2) should be investigated, as they show potential for including this type of data. Additionally, the data structure was designed for Dutch urban systems, which defines neighborhood by postal code areas. The effectiveness of the data structure in more complex urban structures needs to be tested further. The dashboard

presented in this project serves as a proof of concept. It is suggested that this is a good example of using the proposed data structure. However, more research is required to fully explore the possibilities of semantic digital twins for urban policy testing, by integrating additional data sets such as mobility, traffic, and stakeholder-specific use cases. Referring back to the definition of a digital twin provided in section 1 by VanDerHorn and Mahadevan [11], it is suggested that the presented dashboard is a first step towards an urban digital twin, using semantic web technologies. The definition provided that a digital twin consists of three elements: a physical reality, a virtual representation, and an interconnection that exchanges information between the physical reality and virtual representation [11]. While the main focus of the project has been on the second of these elements, it can be argued that the third element needs further development (as has been discussed with the further exploration of dynamic data integration). However, VanDerHorn and Mahadevan [11] argue that update frequency and manner in which a digital twin should be updated is arguable. The authors state that this does not have to be online and by frequent sensor updates, for instance, but can depend on the specific use case. Considering the urban nature of this project, it could therefore be argued that the proposed dashboard is indeed an adequate first step towards a semantic urban digital twin.

## 5. Conclusion

In conclusion, this project has demonstrated the potential of data integration with the semantic modeling method for urban energy analysis and planning. By collecting data from multiple sources and proposing a new ontology (NEO), the project has shown how previously disconnected data sets can be linked and analyzed in a more meaningful way. The visual and interactive semantic dashboard developed in this project allows end users to explore and gain insights into urban energy use and related socio-economic factors. However, the project also highlights several limitations and challenges, such as the need to further testing and refinement of the data structure for different urban systems and the inclusion of real-time or near-real-time data in the data structure. Nevertheless, the project provides a useful example of the possibility of semantic digital twins for urban energy modeling and suggests future directions for research and development in this area. Overall, the project contributes to the broader goal of achieving a more sustainable and efficient urban energy system, which is crucial for improving urban livability.

## 6. Acknowledgments

The authors would like to gratefully acknowledge the support from Eindhoven University Technology and the funding by Smart One W&I TKI KPN flagship.

## References

- [1] United Nations Department of Economics and Social Affairs, World Population Prospects 2019, 2019. URL: <https://population.un.org/wpp/>.

- [2] N. Abbasabadi, M. Ashayeri, R. Azari, B. Stephens, M. Heidarinejad, An integrated data-driven framework for urban energy use modeling (UEUM), *Applied Energy* 253 (2019). doi:10.1016/j.apenergy.2019.113550.
- [3] E. Corry, P. Pauwels, S. Hu, M. Keane, J. O'Donnell, A performance assessment ontology for the environmental and energy management of buildings, *Automation in Construction* 57 (2015) 249–259. doi:10.1016/j.autcon.2015.05.002.
- [4] U. Ali, M. H. Shamsi, C. Hoare, E. Mangina, J. O'Donnell, Review of urban building energy modeling (UBEM) approaches, methods and tools using qualitative and quantitative analysis, 2021. doi:10.1016/j.enbuild.2021.111073.
- [5] E. Curry, J. O'Donnell, E. Corry, S. Hasan, M. Keane, S. O'Riain, Linking building data in the cloud: Integrating cross-domain building data using linked data, *Advanced Engineering Informatics* 27 (2013) 206–219. doi:10.1016/j.aei.2012.10.003.
- [6] U. Ali, M. H. Shamsi, M. Bohacek, K. Purcell, C. Hoare, E. Mangina, J. O'Donnell, A data-driven approach for multi-scale GIS-based building energy modeling for analysis, planning and support decision making, *Applied Energy* 279 (2020). doi:10.1016/j.apenergy.2020.115834.
- [7] U. Ali, M. H. Shamsi, F. Alshehri, E. Mangina, J. O'Donnell, Application of intelligent algorithms for residential building energy performance rating prediction, in: *Building Simulation Conference Proceedings*, volume 5, International Building Performance Simulation Association, 2019, pp. 3177–3184. doi:10.26868/25222708.2019.210232.
- [8] H. E. Degha, F. Z. Laallam, B. Said, Intelligent context-awareness system for energy efficiency in smart building based on ontology, *Sustainable Computing: Informatics and Systems* 21 (2019) 212–233. doi:10.1016/j.suscom.2019.01.013.
- [9] W. Li, Y. Zhou, K. Cetin, J. Eom, Y. Wang, G. Chen, X. Zhang, Modeling urban building energy use: A review of modeling approaches and procedures, 2017. doi:10.1016/j.energy.2017.11.071.
- [10] N. Abbasabadi, J. K. Mehdi Ashayeri, Urban energy use modeling methods and tools: A review and an outlook, 2019. doi:10.1016/j.buildenv.2019.106270.
- [11] E. VanDerHorn, S. Mahadevan, Digital Twin: Generalization, characterization and implementation, *Decision Support Systems* 145 (2021). doi:10.1016/j.dss.2021.113524.
- [12] A. De Nicola, M. L. Villani, Smart city ontologies and their applications: A systematic literature review, 2021. doi:10.3390/su13105578.
- [13] S. Chun, J. Jung, X. Jin, S. Seo, K. H. Lee, Designing an integrated knowledge graph for smart energy services, *Journal of Supercomputing* 76 (2020) 8058–8085. doi:10.1007/s11227-018-2672-3.
- [14] C. Reinisch, M. J. Kofler, F. Iglesias, W. Kastner, Thinkhome energy efficiency in future smart homes, *Eurasip Journal on Embedded Systems* 2011 (2011). doi:10.1155/2011/104617.
- [15] L. Daniele, SAREF4ENER: an extension of SAREF for the energy domain created in collaboration with Energy@Home and EEBus associations, 2020. URL: <https://saref.etsi.org/saref4ener/v1.1.2/>.
- [16] M. Poveda-Villalon, R. Garcia-Castro, P. Espinoza-Arias, SAREF extension for Smart City, 2020. URL: <https://saref.etsi.org/saref4city/v1.1.2/#references>.
- [17] L. Daniele, F. den Hartog, J. Roes, Created in Close Interaction with the Industry: The Smart Appliances REFERENCE (SAREF) Ontology, in: *Lecture Notes in Business*

- Information Processing, volume 225, Springer Verlag, 2015, pp. 100–112. doi:10.1007/978-3-319-21545-7{\\_}9.
- [18] Y. Li, R. García-Castro, N. Mihindukulasooriya, J. O’Donnell, S. Vega-Sánchez, Enhancing energy management at district and building levels via an EM-KPI ontology, *Automation in Construction* 99 (2019) 152–167. doi:10.1016/j.autcon.2018.12.010.
- [19] Kadaster, Basisregistratie Adressen en Gebouwen, 2021. URL: <https://bag2.basisregistraties.overheid.nl/bag/def/>.
- [20] Ontotext, GraphDB by Ontotext, 2023. URL: <https://graphdb.ontotext.com/>.
- [21] Enexis Netbeheer, Open Data, 2023. URL: <https://www.enexis.nl/over-ons/open-data#datasets>.
- [22] CBS, Kerncijfers per postcode, 2019. URL: <https://www.cbs.nl/nl-nl/dossier/nederland-regionaal/geografische-data/gegevens-per-postcode>.
- [23] Rijksdienst voor Ondernemend Nederland (RVO), Openbare data energielabels, 2022. URL: <https://www.ep-online.nl/PublicData>.
- [24] Cesium, Cesium Ion, 2023. URL: <https://cesium.com/platform/cesium-ion/>.
- [25] Cesium, Cesium - The Platform for 3D Geospatial, 2022. URL: <https://cesium.com/>.
- [26] vis.js, Vis.js community edition, 2023. URL: <https://visjs.org/>.
- [27] R. Peters, B. Dukai, S. Vitalis, J. Van Liempt, J. Stoter, Automated 3D reconstruction of LoD2 and LoD1 models for all 10 million buildings of the Netherlands, Technical Report, Delft University of Technology, 2021.
- [28] Rijksoverheid, Prijsplafond voor gas, stroom en stadsverwarming, 2022. URL: <https://www.rijksoverheid.nl/onderwerpen/koopkracht/plannen-kabinet-met-prijsplafond-voor-gas-en-elektriciteit>.
- [29] B. Mashhoodi, Who is more dependent on gas consumption? Income, gender, age, and urbanity impacts, *Applied Geography* 137 (2021). doi:10.1016/j.apgeog.2021.102602.
- [30] F. Belaïd, Exposure and risk to fuel poverty in France: Examining the extent of the fuel precariousness and its salient determinants, *Energy Policy* 114 (2018) 189–200. doi:10.1016/j.enpol.2017.12.005.