GUIDES – Geospatial Urban Infrastructure Data Engineering Solutions (Demo Paper)

Booma Sowkarthiga Balasubramani¹, Omar Belingheri^{1,2}, Eric S. Boria¹, Isabel F. Cruz¹, Sybil Derrible¹, Michael D. Siciliano¹

¹University of Illinois at Chicago, USA

²University of Milan Bicocca, Italy

bbalas3@uic.edu, o.belingheri@campus.unimib.it, {eboria2, ifcruz, derrible, siciliano}@uic.edu

ABSTRACT

As the underground infrastructure systems of cities age, maintenance and repair become an increasing concern. Cities face difficulties in planning maintenance, predicting and responding to infrastructure related issues, and in realizing their vision to be a smart city due to their incomplete understanding of the existing state of the infrastructure. Only few cities have accurate and complete digital information on their underground infrastructure (e.g., electricity, water, natural gas) systems, which poses problems to those planning and performing construction projects. To address these issues, we introduce *GUIDES* as a new data conversion and management framework for urban underground infrastructure systems that enable city administrators, workers, and contractors along with the general public and other users to query digitized and integrated data to make smarter decisions. This demo paper presents the *GUIDES* architecture and describes two of its central components: (i) mapping of underground infrastructure systems, and (ii) integration of heterogeneous geospatial data.

CCS CONCEPTS

- **Information systems** → **Geographic information systems**;
- **Computing methodologies** → *Ontology engineering*;

KEYWORDS

GIS, Geospatial data, Data integration, Ontology, Data analytics, Data management, Urban underground infrastructure, Smart cities

1 INTRODUCTION AND MOTIVATION

Complex geo-analytical applications require the integration of multiple cross-domain geospatial datasets, such as soil types, underground water pipes, and traffic conditions, which change with respect to time and space for effective spatial decision-making (e.g., identification of the most effective sites where to repair a water leakage). *Geospatial data integration* involves combining two or more geospatial datasets from different sources to facilitate analysis, reasoning, querying, and data visualization.

Significant opportunities for smarter data management of urban infrastructure systems are on the rise, as many US cities are moving towards the vision of “smart cities”, creating open data

portals that enable city administrators and residents to explore urban data and perform predictive analyses. Despite the availability of a tremendous volume of available data on cities, the lack of accurate geospatial data for underground infrastructure systems remains a problem. The need to address the poor state of the existing underground infrastructure is a strong rationale to develop such data management systems. For example, New York City has over 6,800 miles of water mains whose average age is 69 years. Over two thirds of them are made of materials susceptible to internal corrosion and prone to leakage, leading to 400 water main breaks in 2013 alone. Cross-domain querying is vital for an effective infrastructure maintenance (e.g., to locate pipes that need to be replaced in order of priority while coordinating across agencies to perform road excavation at the same time), and reiterate the need for integrating multiple heterogeneous geospatial datasets, thereby facilitating queries such as *retrieve all the components from the multiple thematic layers (e.g., census, water pipes, road network)* in a given region, and *how many low-income families will be affected by the burst of a given water main*. Such queries are complex to process due to various kinds of heterogeneities associated with them. Therefore, traditional ontology matching techniques, and the statistical and geospatial data processing tools (e.g., QGIS, ArcGIS) are insufficient to handle such queries.

Data come from various sources, they possess differences in format, representation, context, tools, traits, structure, events, data models, spatio-temporal resolution, data collection and storage techniques, and the relationship between various system properties in a given region. Also, data are most often erroneous, incomplete, and inconsistent, leading to uncertainty. All these factors affect a data conversion and management framework, resulting in imprecise results when the data are analyzed.

GUIDES aims to enable a wide variety of users to explore and query underground infrastructure systems and analyze the impacts of disruptions in these systems (e.g., traffic conditions due to a water main break, malaria incidence in a county due to wastewater leakage), while addressing several technical challenges associated with achieving this vision, and protecting sensitive data simultaneously. In this paper, we describe *GUIDES*, a novel framework to map and query urban underground infrastructure systems. We intend to demo the mapping, pre-processing, and part of the integration process of *GUIDES*, using the infrastructure data of the University of Illinois at Chicago (UIC) campus.

2 FRAMEWORK

This section introduces the *GUIDES* framework (Figure 1) and briefly describes its components.

2.1 Mapping

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

SIGSPATIAL '17, November 7–10, 2017, Los Angeles Area, CA, USA

© 2017 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-5490-5/17/11

<https://doi.org/10.1145/3139958.3139968>

Figure 1: The GUIDES framework.

2.1.1 Data Sources and Data Providers. Big-data driven decision making in smart city applications requires the integration of diverse map-based data sources, many of which are non-standardized. Standardization of data sources and coordination among *data providers*, such as municipalities and service providers can improve the accuracy of data that is being centrally integrated. While some municipalities are working to verify map accuracy through on-field inspection and real-time sensor information, the accuracy of such geospatial data still remains problematic. An initial challenge of the *GUIDES* framework is to create accurate GIS-based representations from existing legacy sources to enable the mapping of multiple thematic layers (e.g., buildings layer and water pipes layer).

2.1.2 Mapping & Pre-processing. Mapping deals with the conversion of data from one or more non-standardized sources into a single standardized format. Legacy data formats lack geographical information and often contain all the relevant information in one single source. Dimensions, for example, are often shown directly on the engineering drawing (e.g., CAD) as opposed to being an attribute of a piece of the infrastructure. Pre-processing algorithms that can automatically detect and solve these issues are critical. *GUIDES* follows a three-step approach for pre-processing. First, a set of rules is developed based on domain knowledge to identify errors (e.g., two overlapping or co-located points should be flagged as they may be a single point). The second step is to generate new variables (e.g., based on network properties) that can be used to further identify errors (e.g., a water valve should have at least two connections). At the same time, we also incorporate GIS features to test whether a point is located within a polygon or not. After highlighting misplaced or missing elements, the third step is to suggest the correct configuration, for which, we develop algorithms and leverage the information present in other infrastructure systems. For instance, given that most underground infrastructure systems are buried under roads, road data can be used to suggest where missing infrastructure should be located. Once complete, all errors and added infrastructure elements can be flagged until they are validated manually during maintenance or new construction.

2.2 Geospatial Data Integration

Data may be collected with different spatial and temporal resolutions, update frequencies, and geometry types [7] with heterogeneity across dimension, location, scale and source. To address these challenges, *GUIDES* uses two kinds of ontologies: (i) a set of domain ontologies; and (ii) a spatio-temporal ontology. The domain

ontology deals with instances related to a specific domain (e.g., water pipes) in the GIS database or relevant external data sources (e.g., census or economic data), whereas the spatio-temporal ontology consists only of the spatial (e.g., urban spatial hierarchy) and temporal (e.g., aggregation of monthly series to annual levels) hierarchies and their corresponding instances. Data integration is then carried out by performing *instance matching*, which enables combining the datasets based on the similarity between their spatio-temporal components by matching their corresponding domain ontology with the spatio-temporal ontology.

2.3 Geospatial Data Analytics & Visualization

The *analytics* module incorporates geostatistical models and spatio-temporal processing mechanisms which enable precise predictions of values for geospatial entities, and quantification of uncertainty. For example, this module applies the spatial function *contains* to identify whether a census block contains a broken water pipe when computing the number of low-income families affected by a water main break in a given region. The *visualization* module consists of a map-based interface for data exploration and comparison of various geostatistical models. Infrastructure elements can be displayed simultaneously for a given spatial entity (e.g., a street with several infrastructure elements including water pipes and buildings) to facilitate better decision making and data exploration with focus (e.g., details on an area where a water leakage is being repaired) and context (e.g., a sketch of other infrastructure elements around the focus area) at the same time.

2.4 Query & Update

The *query* module allows a wide range of geospatial queries for any spatial entity (e.g., census block, street, or a drawn extent) selected by users, whereas the *update* module enables users to add, remove or modify the infrastructure elements in a dataset. Both the *query* and *update* modules restrict their allowed operations, depending on the category of the end-users (e.g., administrators, residents, maintenance crews) and their particular data needs and authorized level of access. For example, the general public should not be aware of the underground infrastructure data that are deemed sensitive, and hence are denied access to those data.

3 DEMONSTRATION SCENARIOS

This section demonstrates how the *GUIDES* framework enables pre-processing and ontology-based data integration mechanisms for urban infrastructure data, using the *Water Pipes* and *Buildings* maps of the UIC campus. These maps were initially in DWG (AutoCAD¹ drawing format), but were converted to shapefile format, and visualized using QGIS.² The original data contained several errors and inconsistencies, and the conversion process generated several errors as well. The maps are transformed into a list of nodes and edges using a Python script [6], splitting the edges at the intersections with nodes.

3.1 Water Pipes Map Pre-processing

This subsection details how *GUIDES* facilitates identification and correction of errors in geospatial datasets.

3.1.1 Fixing Duplicate Nodes. In Figure 2b, although the feature highlighted in red appears to be a single node, the corresponding

¹<https://www.autodesk.com/products/autocad/>

²<http://qgis.org/>

Figure 2: Errors in water pipes map.

feature table in Figure 2c shows that it is in fact two nodes with two different IDs. That is, the two nodes are separate features within the same layer and there is no edge connecting them. Such scenarios are common and pose obvious issues, even when the most basic operations on the network are performed. For example, trying to find a path that goes through the edges in Figure 2b will fail, simply because the overlapping nodes are not connected. To resolve this, a Python script involving the *GDAL*³ and *NetworkX*⁴ libraries is run to remove one node for each pair of such duplicate nodes and connect its edges to the other copy of the node.

3.1.2 Differentiating Infrastructure Elements and AutoCAD Symbols. Figure 2a is an example of a circle representing a manhole. By zooming in, we can see that only one of the three nodes is actually connected to the circle edge. The circle was deleted and replaced with a new node at its center, with proper connections to the other nodes on either side. A field named *Is_manhole* with value 1 for this node, is added in the attribute table of the map, so that the information is kept intact, even though the circle is removed.

Figure 3: (a) Streets (black lines) and water pipes (blue lines); (b) After the removal of some pipes; (c) Missing pipes inferred by the algorithm; (d) A false positive (edge AB) was avoided, thanks to the streets layer.

3.1.3 Context-aware Pre-processing. The *Buildings* layer was used to further identify errors in the *Water Pipes* map. For example, intuitively, a water pipe should either end in a building, or be connected to other water pipes. Otherwise, it is reasonable to assume that there is an error that should be flagged for correction. Such cases can be identified by finding the nodes with degree 1 (end nodes) in the *Water Pipes* layer. This hypothesis has been confirmed by our experiments with synthetic map layers for *Water Pipes* and *Streets* (Figure 3). After the random removal of water pipes (Figure 3b), the algorithm suggested proper corrections to restore the initial map (Figure 3c). In doing so, using the constraints

enforced by the *Streets* layer (water pipes normally run underneath streets) has proven to be fundamental in reducing the number of false positives (incorrectly added pipes), raising the precision from 59% to 93%. Figure 3d shows an example of a pipe whose incorrect addition has been avoided with the help of these constraints.

Applying this hypothesis to the UIC datasets, we should also ensure that the end nodes within the perimeter of a building are not flagged, which is essentially a *point-in-polygon* problem [9]. To resolve this, we use the *GDAL* Python library, which, given a point (a node in the water pipes) and a polygon (a building), checks whether the point falls within the area of the polygon.

The *GDAL* library allows for the creation of multipolygons, which are objects that can contain several polygons. With this feature, one object contained the polygons of all the buildings, instead of having one object (a polygon) for each building. The point-in-polygon check was then performed with the multipolygon in one iteration over the nodes, instead of using two iterations to check if any of the nodes (1st iteration) are in any of the buildings/polygons (2nd iteration). The solution with two iterations results in a much faster computation compared to the one with single iteration, and was therefore chosen for the final implementation.

Figure 4: (a) Buildings (purple lines) & corresponding polygons (green areas); (b) Sample water pipes & buildings layer.

From Figure 4a, we can see that the polygons (green areas) do not cover all of the buildings (purple lines) that the map contains because of the map inconsistencies (e.g., broken edges and detached nodes), which make it impossible for *GUIDES* to build all the polygons properly. Therefore, this layer needs to be pre-processed to remove impurities and connect nodes that define the boundaries of a building, which we do by testing whether a *building* node is on the edge of a full polygon or not, and if not, it is connected to the closest node and flagged. Figure 4b shows a water pipe entering a building. Although node A has degree 1, it will not be flagged as it lies within the area of the building. Implementation of machine learning algorithms (e.g., SVM) to identify inconsistencies in the data and suggest correct configurations is underway.

3.2 Ontologies in Geospatial Data Integration

Figure 5: (a) Domain ontology (e.g., water pipes) - Partial; (b) Spatio-temporal ontology - Partial.

The integration component makes use of the pre-processed data and aids in resolving queries on location- and time-specific data.

³<http://www.gdal.org/>

⁴<https://github.com/networkx/networkx>

For example, given a region: (a) retrieve water pipes and buildings information; (b) retrieve all the components of the multi-layer network (road network, water pipes, rail network, and so on). Queries such as (a) facilitates the identification of potential spots where a water pipe is to be laid, when a new building is constructed. These queries are also particularly difficult to process, because of the heterogeneity of the spatial regions associated with the different datasets. For example, different infrastructure systems may belong to different spatial entities. Similarly, temporal queries also require the matching of heterogeneous data, mostly due to different temporal resolutions and update frequencies.

Figure 6: Geospatial data integration - An example.

To retrieve all components of a multilayer network for a given region (Figure 6), we need to consider the different spatial and temporal resolutions they exhibit (e.g., road networks running across different cities, and water pipes managed individually by each city). To resolve such queries, we perform ontology-based geospatial data integration. *GUIDES* encompasses a domain ontology for each dataset (e.g., water pipes as in Figure 5a), and a generic spatio-temporal ontology (Figure 5b), constructed using Protégé.⁵ Once the query is issued, the spatial and temporal components for the query are identified and their corresponding super- and sub-classes in the spatio-temporal ontology are obtained. We then retrieve the mappings (consisting of these super- and sub-classes) already obtained by matching the instances in the corresponding domain ontology with the instances of the spatio-temporal ontology, based only on the spatial and temporal components, using the AgreementMakerLight (AML) [5] framework. Spatio-temporal functions such as *within*, *crosses*, are used to obtain query results only for the region and time selected.

4 RELATED WORK

GIVA [4], an interactive map-based application, facilitates integration of data from multiple datasets, for a given region and a time interval. *GUIDES* adds on to the capabilities of GIVA, in terms of mapping and context-aware pre-processing, use of external data sources, and the mechanism for data integration, focusing on the urban and underground infrastructure domains. The City of Chicago's OpenGrid [10], a map-based open-source platform, supports advanced queries to identify and monitor incidents across the city. However, it only accepts queries on limited datasets and does not support data integration, nor cross-domain querying, but can be extended to perform predictive analytics on urban data [1]. SocialGlass [8] is a web-based system for visual exploration of large-scale and heterogeneous urban data, but it focuses on events and

not the urban infrastructure. Chang et al. [3] propose a model for visualization of urban relationships using data aggregation techniques. Their model does not support geospatial data integration. The framework of Beck et al. [2] integrates utility data using lightweight ontologies, but requires major changes when a new dataset needs to be integrated. In conclusion, real-life scenarios are more complex than previous work can handle, thus reinforcing the need for a new framework like *GUIDES*.

5 CONCLUSIONS AND FUTURE WORK

In this paper, we introduced *GUIDES*, a data conversion and management framework, which supports ontology-based data integration, querying, analytics, and visualization of heterogeneous geospatial datasets, focusing on the urban infrastructure domain. The framework also supports several types of users such as administrator, planner, maintenance crew, and the general public, with various levels of access. We highlighted the key architectural elements and their capabilities to handle several challenges associated with geospatial data. Given the novelty of *GUIDES* and the complexity of the problems this framework handles, there is a great potential for its expansion to ensure the highest level of usability and interoperability, cross-jurisdictional and inter-organizational collaboration, and workflow optimizations for crews. Opportunities for integration of the *GUIDES* framework with open data exploration platforms such as OpenGrid, will also be explored.

ACKNOWLEDGMENTS

We thank Roberto Tamassia and Goce Trajcevski for many helpful discussions. This work was partially supported by NSF awards CNS-1646395, III-1618126, CCF-1331800, and III-1213013 and by a Bill & Melinda Gates Foundation Grand Challenges Explorations grant. Supplementing the permissions stated in the first page, this work is licensed under a Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>).

REFERENCES

- [1] BALASUBRAMANI, B. S., SHIVAPRABHU, V. R., KRISHNAMURTHY, S., CRUZ, I. F., AND MALIK, T. Ontology-based Urban Data Exploration. In *Proceedings of the 2nd ACM SIGSPATIAL Workshop on Smart Cities and Urban Analytics* (2016), pp. 10:1–10:8.
- [2] BECK, A. R., FU, G., COHN, A. G., BENNETT, B., AND STELL, J. G. A Framework for Utility Data Integration in the UK. In *Proceedings of the Urban Data Management Society Symposium* (2007), Taylor & Francis, pp. 261–276.
- [3] CHANG, R., WESSEL, G., KOSARA, R., SAUDA, E., AND RIBARSKY, W. Legible Cities: Focus-dependent Multi-resolution Visualization of Urban Relationships. *IEEE Transactions on Visualization and Computer Graphics* 13, 6 (2007), 1169–1175.
- [4] CRUZ, I. F., GANESH, V. R., CALETTI, C., AND REDDY, P. GIVA: A Semantic Framework for Geospatial and Temporal Data Integration, Visualization, and Analytics. In *Proceedings of the 21st ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems* (2013), ACM, pp. 544–547.
- [5] FARIA, D., PESQUITA, C., SANTOS, E., PALMONARI, M., CRUZ, I. F., AND COUTO, F. M. The AgreementMakerLight Ontology Matching System. In *International Conference on Ontologies, DataBases, and Applications of Semantics (ODBASE)* (2013), Springer, pp. 527–541.
- [6] KARDUNI, A., KERMANSHAH, A., AND DERRIBLE, S. A Protocol to Convert Spatial Polyline Data to Network Formats and Applications to World Urban Road Networks. *Scientific Data* 3 (2016).
- [7] LE, Y. Challenges in Data Integration for Spatiotemporal Analysis. *Journal of Map & Geography Libraries* 8, 1 (2012), 58–67.
- [8] PSYLIDIS, A., BOZZON, A., BOCCONI, S., AND BOLIVAR, C. T. A Platform for Urban Analytics and Semantic Data Integration in City Planning. In *International Conference on Computer-Aided Architectural Design Futures* (2015), Springer, pp. 21–36.
- [9] SHARMA, A. K., AND GILL, S. K. Methods to Define a Single Point in the Polygon. In *International Journal of Computer Science and Information Technologies (IJCSIT)* (2014), vol. 5, CiteSeer, pp. 3429–3430.
- [10] TUECKE, S., FOSTER, I., AND KESSELMAN, C. The OpenGrid Services Architecture. *The Grid: Blueprint for a New Computing Infrastructure* (2004), 215–242.

⁵<http://protege.stanford.edu/>