Generation of Interactive Knowledge Graphs to Enable Research of the Effects of Trauma Center Organization on Patient Outcomes

Emma Chappell[1], Justin Whorton[1[0009-0003-4268-6207]], Adria Abella Villafranca[1[0000-0003-1503-2163]], Reza Shahriari[2 [0009-0003-0196-2199]], Scott C. Brakenridge[3], Judy L. Bennett[1], Songthip Ounpraseuth[1], Eric D. Ragan[2 [0000-0002-7192-3457]], William R. Hogan[4[0000-0002-9881-1017]], Jonathan Bona[1[0000-0003-1402-9616]], Kevin W. Sexton[5[0000-0002-1460-9867]] and Mathias Brochhausen[1[0000-0003-1834-3856]]

[1] University of Arkansas for Medical Sciences, 4301 W. Markham, Little Rock, AR 72205, USA
[2] University of Florida, Gainesville, FL 32611, USA
[3] UW Harborview Medical Center, 325 9th Ave Main Hospital, Seattle, WA 98104, USA
[4] Medical College of Wisconsin, 8701 Watertown Plank Rd., Milwaukee, WI 53226, USA
[5] Vanderbilt University Medical Center, 1211 Medical Center Drive, Nashville, TN 37232, USA
mbrochhausen@uams.edu

Abstract. This research addresses a knowledge gap on the effect of trauma center organization on patient outcomes. To allow trauma care stakeholders to explore data from their center and identify potential opportunities to improve patient outcomes by adjusting organizational patterns, we are creating knowledge graphs from data on relevant parameters. These include data on organizational parameters of trauma centers, currently an understudied topic, and patient outcome data. We have recruited 42 trauma centers in the US and have collected data from those sites. Our knowledge graph uses an ontology to aid knowledge collection and data integration across medical institutions. We are presenting a novel tool, the Knowledge Path Explorer, that allows users to query and expand the knowledge graph at their own pace and based on their needs. This paper presents the basics of the Knowledge Path Explorer, the extension of the underlying ontology to capture trauma patient outcomes, and the data curation process underlying the knowledge graph.

Keywords: Knowledge Graph, Biomedical Ontologies, Visualization, Patient Outcomes, Trauma Care

1 Introduction

In the United States and in many other countries [1], trauma care is organized around nationally designated trauma centers as key components ranging from Levels 1 to Level 5 trauma centers. Level 1 trauma centers are academic medical centers with 24/7 comprehensive services and Level 5 trauma centers provide initial evaluation, stabilization and diagnosis. For the two highest levels of trauma center verification, Level 1 (L1) and Level 2 (L2), there is no difference in services provided, but L1 centers need to be medical centers with specific research and physician training requirements. Trauma center designation captures only a small subset of organizational parameters. Thus, the impact of organizational decisions on trauma patient outcomes goes unevaluated. Despite growing standardization of clinical trauma care, at L1 and L2 trauma centers there remains significant variability in patient outcomes across trauma centers on both levels [2, 3].

Research on the causes of this variability and the general impact of organizational parameters on patient outcomes in trauma care is hindered by the lack of data on organizational parameters [4]. There is also a lack of tools that allow trauma care stakeholders, such as trauma medical directors, trauma program managers, and institutional decision makers, to explore their organizational parameters against their patient outcomes. Thus, missing the opportunity for fast-track, local changes to improve performance.

In this paper, we describe recent work and improvement of the Knowledge Path Explorer, an interactive, online tool that allows trauma care stakeholders to explore a knowledge graph regarding their trauma center's organizational parameters and patient outcomes at their own pace and self-directed. The KPE relies on the Ontology for Organizational Structures of Trauma Centers and Trauma Systems (OOSTT) [5] to provide relevant classes for data instances as well as relations. This paper also specifies the extensions of OOSTT, expanding its coverage of trauma outcomes.

2 Study Description

The Trauma Institutional Priorities and Teams for Outcomes Efficacy (TIPTOE) project aims to close the knowledge gap between organizational structures of trauma centers and their patient outcomes [4]. Its aim is to empower trauma center stakeholders by linking semantically integrated data on organizational structures and patient outcomes in a Knowledge Graph. To allow the stakeholders to explore the data regarding their institution focusing on aspects of their institution they are interested in, we created the KPE, a novel graphical user interface that allows trauma center employees to explore the relations between their outcomes and their center's organizational parameters. This tool guides hypothesis generation of trauma center stakeholders and facilitates decision making regarding the best ways to evolve organizational parameters to improve patient outcomes.

We have recruited 42 L1 and L2 trauma centers. Through these centers our project has access to data about 33092 trauma patients. To generate the KPE we use patient outcome data routinely collected by the trauma centers as part of the American College of Surgeons' Committee on Trauma's Trauma Quality Improvement Program (TQIP) [6]. Participating centers submit their TQIP data every 4 months to us in a process ensuring no PHI is transmitted. While we collect all patient outcome data in TQIP, our statistical analysis focuses on three outcome measures: mortality, length of stay, and major complications. In addition, each center answers a questionnaire regarding organizational parameters every 6 months. To harmonize incoming data and add background knowledge, we use the Ontology for Organizational Structures in Trauma Centers and Trauma Systems (OOSTT). This ontology is linked up to the organizational parameters questionnaire to help users provide correct and meaningful answers by providing definitions of elements of the questionnaire [4].

2 Material and Methods

2.1 Ontology Expansion Workflow

Two research team members with experience in trauma organization and trauma outcomes data reviewed data needs regarding trauma patient outcomes for our study that aim to identify correlations between organizational parameters and those outcomes. OOSTT already covered organizational parameters of trauma centers and trauma systems [4, 5], hence those were excluded. TQIP, the source of the patient outcomes data, uses the National Trauma Data Standard (NTDS) terms [7], which is organized into 11 topic areas. From those we identified five prioritization areas: *demographic information, injury information, diagnosis information, hospital information, hospital events*. In the prioritization areas we identified a total of 94 terms that were required to represent patient outcome data relevant to the study.

Our ontology development process followed the methods suggested by Arp et al. [8], which are consistent with BFO representational principles [9]. It used the OBO Foundry Principles [10, 11] and naming conventions for ontology terms [12]. The ontology was implemented in OWL2 [13] using the Protégé ontology editor [14]. To prevent the creation of terms already covered by OBO Foundry, we used the OntoBee, an ontology term lookup website [15]. When a term was found to already exist, we imported the term using the Minimum Information to Reference an External Ontology Term (MIREOT) method [16], using a Protégé MIREOT plugin [17].

Applying the ontology development methodologies all required terms were checked against pre-existing OBO Foundry ontologies, a suite of community developed biological and biomedical ontologies aiming to be interoperable with each other [11, 18]. The OBO Foundry was chosen to ensure that the ontological commitments across multiple diverse domains, such as medicine, administration, social and political entities remain consistent. If there was no pre-existing term in the OBO Foundry a textual definition and hierarchical placement was proposed. In an iterative process including experts in trauma care, healthcare organization, and biomedical ontologies, the definitions were optimized until agreement was reached. The terms were added into the ontology as new classes. The HermiT reasoner [19] was used to ensure the consistency of the ontology.

2.3 KPE Data Visualization

The front end of the KPE is an interactive web application accessible by web browser. The KPE allows the user to select a “starting point” for their exploration of the knowledge graph., letting them narrow down the graph to a more manageable focus. Currently, starting points include e.g., “Trauma Program”, “Organization”, and

“Trauma Medical Director”. The chosen starting point is displayed as the root node of the tree, and each predicate and object one step away from that subject are also displayed. When a node one step away from the root is clicked, the remaining choices are removed, and the clicked node now displays its predicates and objects. This leads to a path being displayed as nodes are clicked, leading from the starting point to the concept the user wants to explore.

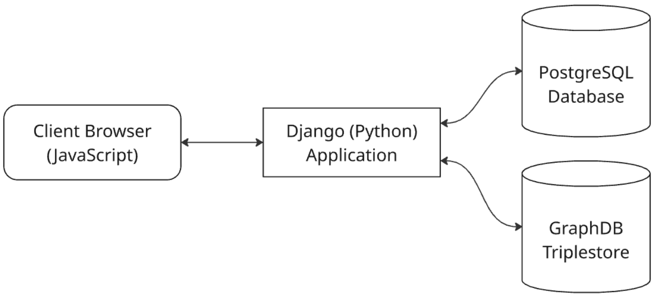

Fig. 1. Architecture of the Knowledge Path Explorer.

Figure 1 shows the architecture of the KPE. Data from the questionnaire on organizational parameters is saved in a PostgreSQL database. The database contains a mapping to define RDF triples from questionnaire data, and a Python script uses the mapping to generate one or more RDF triples per question. This approach allows us to change the RDF mapping and then regenerate the triples from the same survey responses. The RDF triples are stored in a GraphDB [20] triple store, with each survey being saved in its own graph. In addition, the triple store contains triples generated from the TQIP data of participating institutions including aggregate analysis. GraphDB comes with an inbuilt reasoner using the RDFS-Plus ruleset. OOSTT is also stored as its own named graph in the triple store. Another Python script is used to materialize the inferences on each survey's graph. The range on a representative subset of graphs of participating institutions is in the range of 200 to 300 triples. With the inferences materialized the numbers range between 1000 and 1300. This materialization allows the KPE front end to load the relevant information quickly. The KPE uses Django (Python) as a middle layer to query the triple store using SPARQL and serve HTML pages to the client browser. These pages use D3.js [21] to render the knowledge graph on the client side.

Table 1. New classes added to OOSTT to represent patient outcomes

ventilator associated pneumonia	abbreviated injury scale code	total length ventilator use	surgical site infection
superficial incisional surgical site infection	sex information content entity	organ/space surgical site infection	deep surgical site infection

occupational information content entity	ethnics category information content entity	deployed airbag information content entity	occupational industry information content entity
child protective device	minor administrative subdivision	catheter related urinary tract infection	racial category information content entity
bloodstream infection	age at injury	municipality	trauma incident
abbreviated injury scale version identifier	work related injury information content entity	central line associated bloodstream infection	cardiac arrest with cardiopulmonary resuscitation
international classification of diseases tenth revision place of occurrence external cause code	international classification of diseases tenth revision additional external cause code	international classification of diseases tenth revision primary external cause code	

3 Results

3.1 Ontology Updates

Based on the needs assessment and analysis of terms required from NTDS, we have created three new releases of OOSTT. The latest one can be found here: http://purl.obolibrary.org/obo/oostt.owl.

The updates provide the necessary representations to capture patient outcome parameters from TQIP and to integrate those into the Knowledge Graph and make them available for display them in the KPE. A total of 27 classes were created de novo, including new International Resource Identifiers (IRIs) and textual definitions (Table 1). Table 2 shows source ontologies for classes that have been individually imported from pre-existing ontologies and the number of classes that have been imported from those.

In addition to the changes based on adding patient outcome-related representations from NTDS, the latest release version also incorporates fixes to issues identified in the ontology review by a research team member with experience in biomedical ontologies, who is not otherwise involved in the ontology development.

3.2 KPE Updates

The expansion of OOSTT enabled the addition of patient outcome related data from the TQIP reports of the participating institutions. This expands the KG to include patient outcomes and will allow the next development step of adding relationships between organizational structures and patient outcomes as the statistical results of the study become available. Figure 2 shows examples of two patient outcome parameters added to the KG and how they are displayed in the KPE. The code for the KPE can be found here: https://github.com/tiptoe-trauma/knowledgePath.

Table 2. Source ontologies for terms imported using the MIREOT method

Ontology acronym	Ontology name	Number of classes
Apollo_SV	Apollo Structured Vocabulary	1
GEO	Geographical Entity Ontology	2
HP	Human Phenotype Ontology	4
MAXO	Medical Action Ontology	7
MONDO	Mondo Disease Ontology	11
NCIT	National Cancer Institute Thesaurus	2
OBI	Ontology for Biomedical Investigations	2
OMRSE	Ontology for Modeling and Representation of Social Entities	1
UBERON	Uberon multi-species anatomy ontology	4

4 Conclusions

From the results presented here, we conclude that our technology stack allows us to generate a knowledge graph of organizational parameters and patient outcomes for trauma centers. We also demonstrated that we can render the knowledge graph in an interactive tool allowing stakeholder to explore the graph at their own pace.

5 Discussion and Next Steps

Some terms were overly specific, e.g., “injury incident data”. Composite terms like that can be broken down into more generic components and represented by RDF patterns using more generic OWL classes [22]. In our case, for example, “injury incident date” was broken up into “trauma incident” and “date.” This way, the generic term “trauma incident” could be connected to date, time, zip code, etc., thus, making ontology representation more efficiently and semantically clear.

Terms regarding race, ethnicity, and sex have recently been updated in the Ontology for Modeling and Representation of Social Entities (OMRSE) [23]. To unify these terms across the OBO Foundry and prevent duplication of representation, we will import those terms from OMRSE and obsolete the relevant terms in our ontology.

In addition, we are working on developing rules using the OOSTT ontology allowing replacing patterns of multiple nodes and edges with an abbreviated presentation. This will allow a principled and reproducible simplification of the knowledge graph visualized by the KPE and make the KPE more user-friendly to trauma care stakeholders.

The next step of development for the KPE is the integration of the relationships between trauma patient outcomes and the organizational parameters of the trauma centers. We will use the outcomes of our statistical data analysis to represent those in our Knowledge Graph.

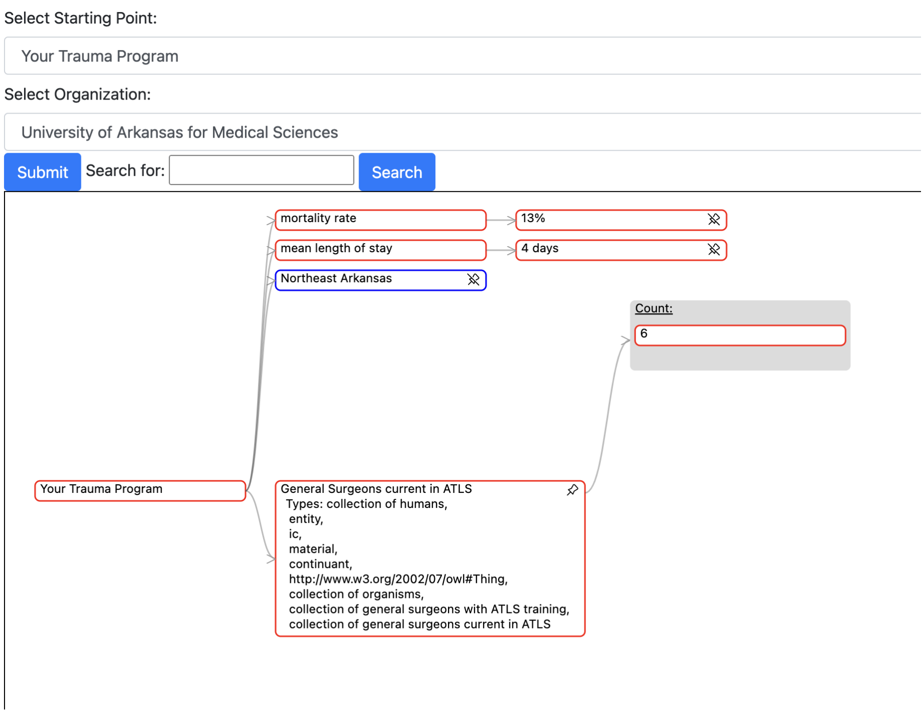

Fig. 2. Screenshot from the KPE showing organizational parameters (General Surgeons current in ATLS) and two key patient outcomes metrics (mortality rate, mean length of stay).

Acknowledgments

Research reported in this publication was supported by the National Institute of General Medical Sciences of the National Institutes of Health under award number R01GM111324.

References

1. Dijkink, S., Nederpelt, C.J., Krijnen, P., Velmahos, G.C., Schipper, I.B.: Trauma systems around the world: A systematic overview. J Trauma Acute Care Surg. 83(5), 917-925 (2017). doi: 10.1097/TA.0000000000001633.
2. Shafi, S., Stewart, R.M., Nathens, A.B., Friese, R.S., Frankel, H., Gentilello, L.M.: Significant variations in Mortality occur at similarly designated trauma centers. JAMA Surgery 144(1), 64-68 (2008). doi:10.1001/archsurg.2008.509.
3. Cudnik, M.T., Sayre, M.R., Hiestand, B., Steinberg, S.M.: Are all trauma centers created equally? A statewide analysis. Academic Emergency Medicine 17(7), 701-708 (2010). doi.org/10.1111/j.1553-2712.2010.00786.x.
4. Brochhausen, M., Ball, J.W., Sanddal, N.D., Dodd, J., Braun, N., Bost, S., et al.: Collecting data on organizational structures of trauma centers: the CAFÉ web service. Trauma Surg Acute Care Open 5(1):e000473 (2020). doi: 10.1136/tsaco-2020-000473.

5. Utecht, J., Ball, J., Bowman, S.M., Dodd, J., Judkins, J., Maxson R.T., et al. Development and validation of a controlled Vocabulary: An OWL representation of organizational structures of trauma centers and trauma systems. Stud Health Technol Inform. 264, 403-407 (2019). doi: 10.3233/SHTI190252.
6. Hemmila, M.R., Cain-Nielsen, A.H., Jakubus, J.L., Mikhail, J.N., Dimick, J.B.: Association of hospital participation in a regional trauma quality improvement collaborative with patient outcomes. JAMA Surg. 153(8), 747-756 (2018). doi: 10.1001/jamasurg.2018.0985.
7. National Trauma Data Standard (NTDS) Data Dictionary. https://www.facs.org/quality-programs/trauma/quality/national-trauma-data-bank/national-trauma-data-standard, last accessed 2025/07/28.
8. Arp, R., Smith, B., Spear A.D.: Building Ontologies with Basic Formal Ontology. MIT Press, Cambridge/Mas. (2015).
9. ISO/IEC 21838-2:2021Information technology — Top-level ontologies (TLO)Part 2: Basic Formal Ontology (BFO). https://www.iso.org/standard/74572.html, last accessed 2025/07/28.
10. OBO Foundry Principles. https://obofoundry.org/principles/fp-000-summary.html, last accessed 2025/07/28.
11. Smith, B., Ashburner, M., Rosse, C., Bard, J., Bug, W., Ceusters, W., et al.: The Obo Foundry: coordinated evolution of ontologies to support biomedical data integration. Nat Biotechnol 25(11), 1251-1255 (2007). DOI: 10.1038/nbt1346.
12. Schober, D., Smith, B., Lewis, S.E., Kusnierczyk, W., Lomax, J., Mungall, C., et al.: Survey-based naming conventions for use in OBO Foundry ontology development. BMC Bioinformatics, 10:125. (2009). doi: 10.1186/1471-2105-10-125.
13. OWL 2 Web Ontology Language Document Overview (Second Edition), https://www.w3.org/TR/owl2-overview, last accessed 2025/07/28.
14. Musen, M.A., Protégé Team: The Protégé project: A look back and a look forward. AI Matters. Association of Computing Machinery Specific Interest Group in Artificial Intelligence, 1(4) (2015). DOI: 10.1145/2557001.25757003.
15. Ong, E., Xiang, Z., Zhao, B., Liu, Y., Lin, Y., Zheng, J., et al.: Ontobee: A linked ontology data server to support ontology term dereferencing, linkage, query, and integration. Nucleic Acid Research 45(D1):D347-D352. (2017). DOI: 10.1093/nar/gkw918.
16. Courtot, M., Gibson, F., Ruttenberg, A.: MIREOT: the minimum information to reference an external ontology term. Applied Ontology, 6(1), 23-33, (2003). doi.org/10.3233/AO-2011-0087.
17. Chen, H.J., Crow, W.A., Hall, R., Liu, J., Pendurthi, T., Schmidt, T., et al.: Simplifying MIREOT: a Protégé plugin. CEUR WS 914, 25-28 (2012).
18. Jackson, R., Matenzoglu, N., Overton, J.A., Vita, R., Balhoff, J.P., Buttigieg, P.L., et al.: OBO Foundry in 2021: operationalizing open data principles to evaluate ontologies. Database (Oxford), 2021:baav069 (2021). doi: 10.1093/database/baab069.
19. Glimm, B., Horrocks, I., Motik, B., Stoilos, G., Wang, Z.: HermiT: an OWL 2 Reasoner. Journal of Automated Reasoning, 53, 245-69, (2014). doi.org/10.1007/s10817-014-9305-1.
20. GraphDB, https://graphdb.ontotext.com/documentation/11.0/index.html, last accessed 2025/07/29.
21. D3, https://d3js.org, last accessed 2025/07/29.
22. Bona, J.P., Nolan, T.S., Brochhausen, M.: Ontology-enhanced representations of non-imaging data in the Cancer Imaging Archive. CEUR WS 2285, 37 (2018).
23. Hicks, A., Hanna, J., Welch, D., Brochhausen, M., Hogan, W.R.: The ontology of medically related social entities. J Biomed Semantics, 12(7), 47 (2016). doi: 10.1186/s13326-016-0087-8.