VOWLExplain: Knowledge Graph Visualization for Explainable Artificial Intelligence

Filipa Serrano, Susana Nunes and Catia Pesquita

LASIGE, Faculdade de Ciências, Universidade de Lisboa, Portugal

Abstract
Ontologies and Knowledge Graphs are a potential solution to the problem of lack of explainability in Artificial Intelligence, and are especially suited to explain how a given prediction fits with existing knowledge in a domain. Communicating these semantic explanations to end users in a correct, clear and trustworthy fashion is crucial to support the adoption of artificial intelligence in critical and complex domains such as healthcare. We developed VOWLExplain, a tool that supports the visualization of semantic post-hoc explanations for predictions made by AI black-box models. We performed a small-scale user study comparing text-based and graph-visualization based explanations in a case study for personalized medicine. The results highlighted the diversity of how users perceive explanations, and demonstrated that although users indicate a slight preference for graph based representations, they generally rate them as correct and as trustworthy as text-based explanations, but do consider them clearer.

Keywords
Explainable AI, Knowledge Graph, Visualization

1. Introduction

Artificial Intelligence (AI), specifically Machine Learning (ML) algorithms, have been gaining more importance due to the development of powerful models, such as Deep Learning (DL). There are several applications in which DL models are being used, with great potential and promising results [1].

However, the application of black-box AI in critical use cases is hindered by their lack of explainability. Black-box models are opaque models whose internal mechanism is unknown or uninterpretable to humans. Explainability, the ability of a user to understand, evaluate and eventually trust a specific prediction made by a machine learning model is essential for applying these models in sensitive fields, where decisions highly impact people's lives [2, 3].

The concept of explainable AI (XAI) is not new and has been used since the beginnings of artificial intelligence. There have been efforts to clearly define XAI terminology, distinguishing concepts such as transparency, interpretability and explainability [2]. Explainability approaches allow users to have a clearer understanding of why certain AI predictions were made, which may help increase their trust and acceptance in these predictions.

Explanations are usually divided into two categories: post-hoc and ante-hoc explanations. Ante-hoc systems are interpretable by design, which includes decision trees and linear regression.

VOILA! 2022: 7th International Workshop on Visualization and Interaction for Ontologies and Linked Data, October 23, 2022, Virtual Workshop, Hangzhou, China, Co-located with ISWC 2022.

Post-hoc systems, on the other hand, suggest possible explanations for specific predictions made by ML models, maintaining high fidelity to the original model and producing helpful explanations instead of trying to explain the original model itself [4]. An additional dimension of explainability is how the model's predictions fit with prior knowledge. This is especially relevant in areas where the body of knowledge cannot be fully understood by users due to its size and complexity, with a prime example being the clinical domain [5]. The need for knowledge representations to support AI in clinical and biomedical settings is recognized by the scientific community, but their effective use is still an open challenge [4].

Ontologies and Knowledge Graphs (KGs) afford a potential solution to this need [3]. KGs are graph-based representations of knowledge that use nodes to represent entities and edges to represent relations. The possible types of elements and relations can be described by an ontology [6]. The structured and connected form of modelling domains in an ontology allows a facilitated integration, as well as an extensive vocabulary and clear identifiers [7], which allows a shared common knowledge [8].

The abundance and diversity of ontologies and KGs in the biomedical and clinical domains are an opportunity that can be explored by XAI. One particular area where explanations are crucial is personalized medicine. Personalized medicine tries to answer the question of *"What is the right drug at the right dose for the right patient?"* by integrating and analysing very large volumes of diverse and heterogeneous data coming from a variety of sources and different scientific and clinical domains. Black-box algorithms are leading in the field [9], but their opaque nature is a recognized challenge. Moreover, a lack of consideration of users' expectations is among the chief reasons for the limited adoption of ML systems in critical and complex domains [10], so effective user interfaces are also a requirement [4]. If KGs are to be the backbone of an explanation for an AI prediction, then the visualization of the KG can be the communication means of such an explanation. Ontology and KG visualization are active research areas, with several existing tools, but the challenges of visualizing large graphs and adapting to specific use cases remain [11].

We are then faced with two challenges: (1) Can KGs be used to craft semantic explanations of how a particular AI prediction fits with existing knowledge?; (2) Are KG-based visualizations an effective means to communicate such an explanation? In this work, we focus on the second challenge, building on the following definition of a semantic explanation: an explanation for a specific AI prediction for a given instance that corresponds to a subgraph extracted from the KG that includes a representation of the instance, the prediction and a path in the KG that connects both. This type of explanation requires that both the instance for which the prediction was made and the prediction itself to be encoded in the KG. To address this challenge, we require a tool that enables the visualization of a semantic explanation, and a user study that compares this visualization with other communication approaches, e.g, text-based.

The main goal of this work is to develop a visualization tool, that given a semantic explanation of how a particular AI prediction fits with prior knowledge represents this explanation in a visual manner. The main contributions of this work are: (1) the extension of the VOWL language to represent additional KG elements required for semantic explanations; (2) the adaptation of the WebVOWL tool to represent semantic explanations; (3) the development of representative semantic explanations to evaluate the tool; (4) the design of a user study; (5) a small-scale user study.

2. Related Work

This work is related to two complementary domains: the use of KGs for XAI and the visualization of KGs. Below we provide a brief overview of relevant works in these domains.

KGs for XAI We focus on the exploration of KGs for post-hoc explanations both in supervised learning and pattern mining settings. Explanations that use background knowledge are likely to be closer to human conceptualizations and thus more useful in applications.

Trepan Reloaded [12] has been a recent extension of Trepan [13], an algorithm that creates a decision tree that tries to replicate deep neural network model predictions and employs ontologies to select the most general concepts, determined through the hierarchy of the ontology, to then be used as tree nodes. The authors consider that the more general concepts will provide the most understandable explanation, which is a reasonable although semantically poor criterion.

Lécué and Wu [14] developed a method that uses ontologies to help explain predictions of classification models. It selects representative data points and their semantic context is then built by characterizing them with their respective concepts using an ontology. The concepts can then be divided into positive concepts, if they characterize points in a certain class, and negative concepts if they describe points in the opposite class. An algorithm is then used to select the most useful positive and negative concepts for explaining each class, which are preferably the more general ones. This results in a list of ordered informative explanations, which are based on the contrasting concepts of each class.

Ontologies can also help in filtering and organizing results of pattern mining techniques. Jay and D'Aquin (2013) developed a tool to interpret results obtained from data mining with the use of Linked Open Data (LOD) [15]. Their approach is applied to results from pattern mining techniques, which are sequential patterns regarding hospitalized patients' trajectories. This approach makes use of linked data to extract information about the result patterns and to organize them in a hierarchical way. The tool also allows the linkage of the patterns to their terminology, making their interpretation of patterns easier.

These representative works differ from ours in their definition of an explanation. In [12], an explanation is the model itself, which is built with input from an ontology. In [14], the explanations are the most common classes that represent positive and negative examples. In [15] the explanations are the semantic representations of the extracted patterns. In all works, the presentation of the explanations is addressed, however both [12] and [14] disregard the contexual and semantic properties of the ontology they explore to generate the explanations when presenting them to users. [12] presents the decision tree where although nodes correspond to ontology classes, their semantic properties are ignored, while [14] merely proposes to present a list of relevant classes, without providing any other semantic information. On the other hand, [15] explore the semantic properties of the linked data and ontologies they use, allowing patterns to be navigated according to the ontology hierarchy.

Visualization of ontologies and KGs The majority of tools to visualize ontologies employ two-dimensional node-link visualizations with a focus on class hierarchies and are rarely use case oriented [11]. 10 out of the 33 tools surveyed by [11] are plugins for the popular ontology editor Protégé [16]. Protégé affords a visualization of an ontology as an indented list, but a variety of plugins cover other layouts, such as trees and graphs (e.g., OWLViz [17], OntoGraf [18]). Other popular tools are browser-based. Ontodia [19] supports quick visualization of RDF datasets and OWL ontologies on the Web. WebVOWL [20] uses VOWL (Visual Notation for Ontologies) to support web-based ontology visualization aiming at a better and more intuitive experience for the user. VOWLMap [21] targets the visualization of ontology alignments. Graph databases that include visualization interfaces can also be used to visualize KGs, such as Gruff for AllegroGraph [1].

Orthogonally, a recent study [22] performed a comparative evaluation of state-of-the-art linked data visualization tools based on a number of use cases including the ability to *visualize the paths that connect different instances.* Only one of ten tools was able to accomplish this use case [23], however it failed on other relevant use cases such as visualizing the information related to a class or a property.

There is not one-size-fits all solution to the problem of ontology and KG visualization and it is clear that different use cases demand different visualization and interaction techniques. For the specific use case of semantic explanation visualization no existing tool is capable of answering all requirements.

3. Methodology

Figure 1 represents two semantic explanations that illustrate the fit between prior knowledge encoded in the KG and the AI prediction. In this case, the instance is John Doe and the predicted drug to treat this patient's disease is Sunitinib, an antineoplastic agent, and they are connected through two paths that link the semantic representation of John Doe to the semantic representation of Sunitinib. These paths provide two possible explanations of why Sunitinib was predicted for this patient, one of them more generic (grey) and one more specific (colors). The generic explanation states that John Doe has a mutation MET T540 that is related to renal cell carcinoma, which is a type of cancer and cancers can be treated by the antineoplastic agents, of which Sunitinib is an example. The more specific explanation declares that the patient has a specific mutation MET T540 that promotes the transcription of the MET gene that is related to tyrosine kinase activity which is inhibited by Sunitinib. Both explanations are valid, but the specific one provides more information to understand the possible link between a patient feature (the mutation) and the drug effect.

To visualize semantic explanations we need to fulfil the following requirements:

1. load a semantic explanation, i.e., a KG subgraph
2. visualize the instance and its properties, i.e, the KG individual for whom the AI prediction was made
3. visualize the predicted class and its properties, i.e., the KG class that represented the predicted class
4. visualize the path between instance and predicted class composed by individuals, classes and properties

[1]https://allegrograph.com/products/gruff/

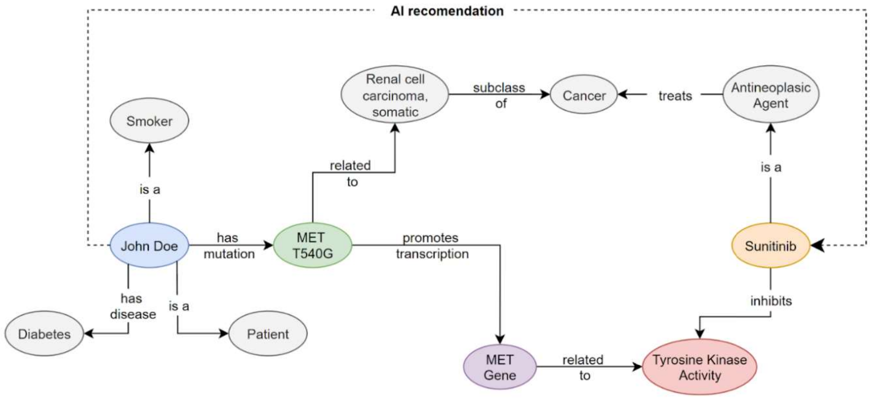

Figure 1: Example of two semantic explanations for the recommendation of sunitinib to patient John Doe. A more specific semantic explanation is represented in color and a more generic one represented in grey.

5. expand the neighbourhood of nodes in the path to include neighbouring regions of the KG

The visualization of ontologies can be supported by visual languages, such as VOWL [24] and its associated visualization tool WebVOWL [25]. In this work, we extended VOWL to support the visualization of individuals and adapted WebVOWL to the visualization of semantic explanations.

3.1. Extending VOWL

The VOWL notation was extended to represent new elements required for the visualization of individuals and their properties, as presented in Table 1. We defined the representation of Named Individuals and their relations (both the "Instance of" relation to their corresponding class, as well as the object properties that connect the individuals between themselves).

3.2. VOWLExplain

VOWLExplain was developed over WebVOWL. Previous works have already demonstrated the potential to adapt WebVOWL to develop new visualization tools adapted to specific tasks [21].

WebVOWL takes as input a JSON file with the desired ontology. This JSON file has a specific format that describes the different elements to be interpreted and represented by WebVOWL and is generated by the OWL2VOWL tool. However, this tool does not allow the representation of individuals, so we developed a tool to process the KG subgraph and generate a JSON representation that follows the structure of WebVOWL but contains the extensions required to represent individuals. We then modified the WebVOWL code to include the representation of the new VOWL elements. The first part of this adaptation included recognizing and processing these elements from the JSON file. Then, we guaranteed their accurate representation, with the new VOWL notation, by creating the new graphical elements

Table 1
Extension of the VOWL notation.

VOWL Element	Notation
Individual	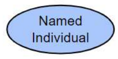
Instance of	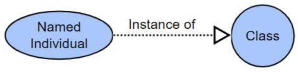
Object Property	
Object Property and Inverse Property	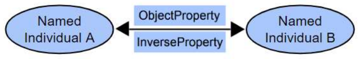

in the VOWLExplain code. The addition of the new VOWL elements also included the adaptation of all of the useful original features of WebVOWL for each element, such as moving, selecting, and showing the details in the lateral menu: Name, Type, other characteristics, and domain and range (in case of relations). The overall appearance and functionalities remained the same, with the addition of the new VOWL elements, as well as a new feature for collapsing and expanding the neighboorhood of nodes, in order to facilitate the visualization of the explanation paths.

3.3. Evaluation strategy

The evaluation of VOWLExplain was grounded in the specific case of personalized medicine for renal cancer[2].

Ontologies, Data and Explanations To evaluate VOWLExplain we built a KG based on a network of aligned ontologies and simulated drug recommendations based on data describing real patients and the drugs that were used for their clinical case.

The ontology network comprises a set of 28 biomedical ontologies aligned to each other to build the semantic backbone of the KG [26, 27]. The ontologies cover a wide range of domains, including clinical data, clinical trial data and 'omics data, such as immunopeptidomics and transcriptomics and proteomics.

The patient data (clinical features, gene mutations and administered drug) was obtained from The Cancer Genome Atlas (TCGA), which contains rich metadata, such as the clinical characterizations of patients, and transcriptomics data from the work by Braun *et al.* [28], which describes gene activity and mutations in renal cancer.

We developed semantic explanations for the drug recommendations for patients by creating paths in the KG between patient's gene mutations or clinical characteristics and the recommended drug using Protégé [16]. We created six semantic explanations, four which represented specific explanations where the mechanism of a genetic mutation and the effect of a drug are

[2]in the context of European Commission funded KATY project https://katy-project.eu/

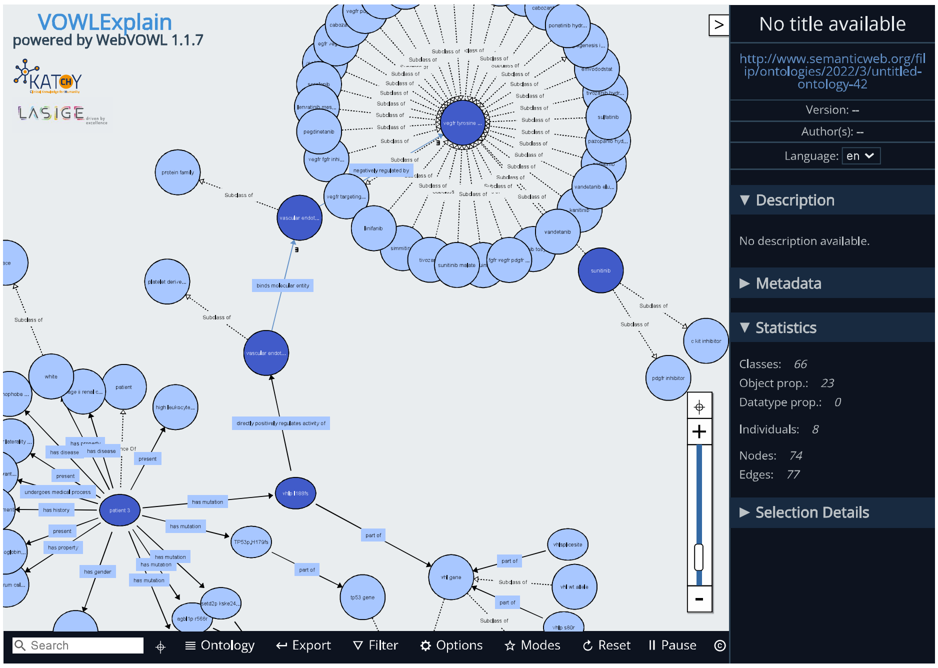

Figure 2: Example of semantic explanation loaded in VOWLExplain. The explanation shows a path (represented in dark blue) that connects the patient (patient 3) to the recommended drug (sunitinib), as well as some neighbors that provide context to the explanation (represented in light blue).

represented (an example of a specific graph-based explanation is presented in Figure 2) and two based on the generic anti-cancer effect of drugs (example of a generic text explanation is presented in Figure 3). Text explanations are simple transformations of class and property names in the explanation path into grammatically correct sentences.

Patient 2 has the disease Metastatic Renal Cell Carcinoma. Nivolumab is capable of inhibiting or preventing the pathological process of Metastatic Renal Cell Carcinoma.

Figure 3: Example of a text explanation presented in the user study. This is a generic explanation.

Ontologies, Data and Explanations We performed a preliminary user study, to gather feedback from a small pool of users, before embarking on a large-scale study. The goal of the user study was to evaluate the usefulness of visual semantic explanations and is based on comparing textual representations, handcrafted based on the semantic explanations and graph-based representations of semantic explanations using VOWLExplain.

We recruited four users with a background in health informatics. The study was both observational (online video call) and questionnaire based. The evaluation was task-based: users were given information about a patient and its corresponding AI recommendation (Figure 4) and then asked a number of questions about either a textual or a graph-based semantic explanation (SEQ) for the given patient and prediction:

Patient 3	
Clinical Data	Male Birth Year:1935 White Number of packs smoked a year: 48 Renal cell carcinoma, chromophobe type No neoadjuvant therapy given No prior treatment High Leukocyte count Low hemoglobin and serum calcium levels Stage II renal cell cancer Tumor Laterality Left Diagnosed at 71
Mutated Genes	VHL gene AGBL1 gene DST gene SETD2 gene TP53 gene
Prediction	Sunitinib

Figure 4: Example of table presented in the user studies with characterization of a patient and its corresponding AI recommendation.

(SEQ1) Rate the explanation in terms of Correctness (I don't know or 1-Not at all to 4-Completely)

(SEQ2) Rate the explanation in terms of Clarity (I don't know or 1-Not at all to 4-Completely)

(SEQ3) Rate the explanation in terms of Trustworthiness (I don't know or 1-Not at all to 4-Completely)

(SEQ4) How would you improve this particular explanation? (free text - optional)

After rating six different semantic explanations (three text and three graph-based), users were also asked the following general questions (GQ):

(GQ1) Do you think explanations, either graph or text based, are useful? (1-Not useful 5-Very useful)

(GQ2) Which explanations do you prefer? (1-Graph, 3 corresponds to no preference, 5-Text)

(GQ3) Adding context to the explanations (neighborhood in VOWLExplain) is useful? (1-Not useful 5-Very useful)

(GQ4) Any suggestions or comments to improve the graph-based explanations? (free text - optional)

(GQ5) Any suggestions or comments to improve the VOWLExplain tool? (free text - optional)

(GQ6) Any suggestions or comments to improve the textual explanations? (free text - optional)

(GQ7) Any suggestions or comments to improve the explanations, overall? (free text - optional)

Users' screens were recorded while using VOWLExplain to elucidate which features were used. Half the users were first presented with textual explanations followed by graph explanations, and the other half vice-versa. Users were never shown the same explanation in both forms.

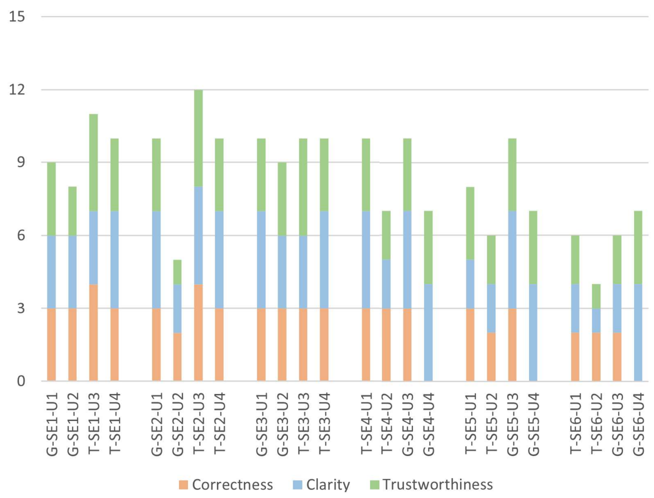

Figure 5: User (U) ratings for each semantic explanation (SE) in graph (G) or text (T) format, in terms of correctness, clarity and trustworthiness

4. Results and Discussion

4.1. VOWLExplain

Figure 2 depicts a semantic explanation loaded in VOWLExplain. It presents the path that connects the patient to the predicted drug, as well as the neighborhood of this path, for more context.

4.2. Preliminary User Studies

All users rated explanations as very useful (rating=5), regardless of type, all showed a slight preference for graph explanations (rating=2), and rated the context of graph explanations as useful (rating=4) or very useful (rating=5). However, the specific ratings on Correctness, Clarity and Trustworthiness reflect different opinions (see Figure 5). This preliminary user study was helpful in understanding the diversity of how users perceive explanations. For instance, U4 struggles with rating the correctness of all graph-based explanations (rating=I don't know), but shows no hesitancy in rating textual explanations. Moreover, U4 rates all explanations regardless of type with the same level of clarity and trustworthiness. On the other hand, U2 rates graph-based explanations generally higher in terms of clarity and trustworthiness, except for the generic explanation (G-SE2) which is rated considerably lower on par with the generic textual explanation (T-SE6). U1 rates graph-based explanations generally higher, while U3 rates the textual ones higher. Looking at the profiling information, these preferences make sense, since U3 rates their knowledge of KGs as *Novice* while U1 rate themselves as *Competent*.

In Table 2, we can see the median of all the answers regarding the Correctness, Clarity and Trustworthiness of the explanations. The scores for Correctness and Trustworthiness are equivalent for both text and graph visualization based explanations, but in Clarity the graph

Table 2
Median scores for correctness, clarity and trustworthiness of explanations.

Explanation Type	Correctness	Clarity	Trustworthiness
Graph Visualization	3	4	3
Text	3	3	3

explanations received a higher median score. It is possible the perceived increase in clarity comes with the additional context that the KG visualization affords.

Users provided some suggestions for improvements, such as fitting the entire explanation path on the screen and marking the instance and the prediction nodes with a different color.

5. Conclusions

In recent years, ontologies and KGs have been proposed as a fundamental piece of the XAI puzzle. In complex and critical domains, such as healthcare, they are widely recognized as essential [29]. This work presented VOWLExplain, a tool for visualizing semantic explanations for AI predictions that are based on elucidating how a prediction fits with existing knowledge encoded in the KG. A small-scale user study comparing text representations of semantic explanations with graph-visualization representations revealed a diversity of user perceptions, and although users stated a preference for the graph-based visualization, they did not rate them as more correct or trustworthy than text based ones. They did rate them as generally more clear. The user study also highlighted some limitations of VOWLExplain, including the lack of distinct representations for instance and prediction. In future work, we will address user suggestions and also integrate text explanations into VOWLExplain by exploring tools that translate OWL constructs into natural language such as NaturalOWL [30]. We believe presenting both types of explanation will make the tool more versatile and easier to pick up for users without familiarity with ontologies or KGs. We will also conduct a larger-scale user study, with users recruited from both clinical, biomedical and health informatics backgrounds.

Data and Source Code Availability

Data, code, video tutorial and user study form are openly available at: https://github.com/liseda-lab/VOWLExplain

Acknowledgements

This work was partially supported by the KATY project funded the European Union's Horizon 2020 research and innovation program (GA 101017453). It was also supported by FCT through the LASIGE Research Unit (UIDB/00408/2020 and UIDP/00408/2020). The authors are also grateful to Daniel Faria for the discussions and advice on transforming OWL into a JSON representation.

References

[1] S. Min, B. Lee, S. Yoon, Deep learning in bioinformatics, Briefings in bioinformatics 18 (2017) 851–869. doi:10.1093/bib/bbw068. arXiv:1603.06430.

[2] M. A. Clinciu, H. F. Hastie, A survey of explainable AI terminology, NL4XAI 2019 - 1st Workshop on Interactive Natural Language Technology for Explainable Artificial Intelligence, Proceedings of the Workshop (2019) 8–13. doi:10.18653/v1/w19-8403.

[3] S. Chari, D. M. Gruen, O. Seneviratne, D. L. Mcguinness, Foundations of Explainable Knowledge-Enabled Systems (2020). arXiv:2003.07520v1.

[4] A. Holzinger, C. Biemann, C. S. Pattichis, D. B. Kell, What do we need to build explainable AI systems for the medical domain? (2017) 1–28. URL: http://arxiv.org/abs/1712.09923. arXiv:1712.09923.

[5] K. N. Vokinger, S. Feuerriegel, A. S. Kesselheim, Mitigating bias in machine learning for medicine, Communications medicine 1 (2021) 1–3.

[6] Heiko Paulheim, Knowledge Graph Refinement: A Survey of Approaches and Evaluation Methods, Semantic Web 8 (2016) 489–508.

[7] J. J. Cimino, X. Zhu, The practical impact of ontologies on biomedical informatics., Yearbook of medical informatics (2006) 124–135. doi:10.1055/s-0038-1638470.

[8] R. Hoehndorf, P. N. Schofield, G. V. Gkoutos, The role of ontologies in biological and biomedical research: A functional perspective, Briefings in Bioinformatics 16 (2015) 1069–1080. doi:10.1093/bib/bbv011.

[9] N. J. Schork, Artificial intelligence and personalized medicine, in: Precision medicine in cancer therapy, Springer, 2019, pp. 265–283.

[10] C. J. Cai, E. Reif, N. Hegde, J. Hipp, B. Kim, D. Smilkov, M. Wattenberg, F. Viegas, G. S. Corrado, M. C. Stumpe, et al., Human-centered tools for coping with imperfect algorithms during medical decision-making, in: Proceedings of the 2019 chi conference on human factors in computing systems, 2019, pp. 1–14.

[11] M. Dudáš, S. Lohmann, V. Svátek, D. Pavlov, Ontology visualization methods and tools: a survey of the state of the art, The Knowledge Engineering Review 33 (2018).

[12] R. Confalonieri, T. Weyde, T. R. Besold, F. Moscoso del Prado Martín, Using ontologies to enhance human understandability of global post-hoc explanations of black-box models, Artificial Intelligence 296 (2021) 103471. URL: https://doi.org/10.1016/j.artint.2021.103471. doi:10.1016/j.artint.2021.103471.

[13] M. W. Craven, J. W. Shavlik, Extracting Tree-Structured Representations of Trained Networks, Advances in Neural Information Processing Systems 8 (1996) 24–30. URL: citeseer.ist.psu.edu/craven96extracting.html.

[14] F. Lecue, J. Wu, Semantic Explanations of Predictions (2018). URL: http://arxiv.org/abs/1805.10587. arXiv:1805.10587.

[15] N. Jay, M. D'Aquin, Linked data and online classifications to organise mined patterns in patient data., AMIA ... Annual Symposium proceedings / AMIA Symposium. AMIA Symposium 2013 (2013) 681–690.

[16] N. F. Noy, R. W. Fergerson, M. A. Musen, The knowledge model of protégé-2000: Combining interoperability and flexibility, Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics) 1937 (2000) 17–32. doi:10.1007/3-540-39967-4_2.

[17] M. Horridge, Owlviz, Available on: http://protegewiki. stanford. edu/wiki/OWLViz (2010).

[18] S. Falconer, Ontograf, Protégé Wiki (2010).

[19] D. Mouromtsev, D. S. Pavlov, Y. Emelyanov, A. V. Morozov, D. S. Razdyakonov, M. Galkin, The Simple Web-based Tool for Visualization and Sharing of Semantic Data and Ontologies., in: ISWC (Posters & Demos), 2015.

[20] S. Lohmann, S. Negru, F. Haag, T. Ertl, Visualizing ontologies with VOWL, Semantic Web 7 (2016) 399–419. doi:10.3233/SW-150200.

[21] A. Guerreiro, C. Pesquita, D. Faria, VOWLMap: Graph-based Ontology Alignment Visualization and Editing ?, CEUR Workshop Proceedings 2980 (2021) 1–12.

[22] F. Desimoni, N. Bikakis, L. Po, G. Papastefanatos, A comparative study of state-of-the-art linked data visualization tools, in: 5th International Workshop on Visualization and Interaction for Ontologies and Linked Data, VOILA 2020, volume 2778, CEUR-WS, 2020, pp. 1–13.

[23] P. Heim, S. Hellmann, J. Lehmann, S. Lohmann, T. Stegemann, Relfinder: Revealing relationships in rdf knowledge bases, in: International Conference on Semantic and Digital Media Technologies, Springer, 2009, pp. 182–187.

[24] S. Lohmann, S. Negru, F. Haag, T. Ertl, VOWL 2: User-oriented visualization of ontologies, in: International Conference on Knowledge Engineering and Knowledge Management, Springer, 2014, pp. 266–281.

[25] S. Lohmann, V. Link, E. Marbach, S. Negru, Webvowl: Web-based visualization of ontologies, in: International Conference on Knowledge Engineering and Knowledge Management, Springer, 2014, pp. 154–158.

[26] M. C. Silva, D. Faria, C. Pesquita, Integrating Knowledge Graphs for Explainable Artificial Intelligence in Biomedicine (2021). arXiv:2003.07523.

[27] M. C. Silva, D. Faria, C. Pesquita, Matching multiple ontologies to build a knowledge graph for personalized medicine, in: European Semantic Web Conference, Springer, 2022, pp. 461–477.

[28] D. A. Braun, Y. Hou, Z. Bakouny, M. Ficial, M. Sant'Angelo, J. Forman, P. Ross-Macdonald, A. C. Berger, O. A. Jegede, L. Elagina, et al., Interplay of somatic alterations and immune infiltration modulates response to pd-1 blockade in advanced clear cell renal cell carcinoma, Nature medicine 26 (2020) 909–918.

[29] C. Pesquita, Towards semantic integration for explainable artificial intelligence in the biomedical domain., in: HealthInf Conference - BIOSTEC 2021, volume 5, 2021, pp. 747–753.

[30] I. Androutsopoulos, G. Lampouras, D. Galanis, Generating natural language descriptions from owl ontologies: the naturalowl system, Journal of Artificial Intelligence Research 48 (2013) 671–715.