



Politecnico di Bari

Repository Istituzionale dei Prodotti della Ricerca del Politecnico di Bari

On the Relevance of Explanation for RDF Resources Similarity

This is a pre-print of the following article

Original Citation:

On the Relevance of Explanation for RDF Resources Similarity / Colucci, Simona; Donini, Francesco M.; Di Sciascio, Eugenio. - STAMPA. - 488:(2023), pp. 96-107. (Intervento presentato al convegno 3rd International workshop on Model-driven Organizational and Business Agility, MOBA 2023 tenutosi a Zaragoza, Spain nel June 12-13, 2023) [10.1007/978-3-031-45010-5_8].

Availability:

This version is available at <http://hdl.handle.net/11589/262080> since: 2023-11-17

Published version

DOI:10.1007/978-3-031-45010-5_8

Publisher: Springer

Terms of use:

(Article begins on next page)

On the relevance of explanation for RDF resources similarity

Simona Colucci¹, Francesco M. Donini², and Eugenio Di Sciascio¹

¹ Politecnico di Bari, Bari, Italy

`simona.colucci@poliba.it`, `eugenio.disciascio@poliba.it`

² Università degli Studi della Tuscia, Viterbo, Italy

`donini@unitus.it`

Abstract. Artificial Intelligence (AI) has been shown to productively affect organizational decision making, in terms of returned economic value. In particular, agile business may significantly benefit from the ability of AI systems to constantly pursue contextual knowledge awareness. Undoubtedly, a key added value of such systems is the ability to explain results. In fact, users are more inclined to trust and feel the accountability of systems, when the output is returned together with a human-readable explanation. Nevertheless, some of the information in an explanation might be *irrelevant* to users—despite its truthfulness. This paper discusses the relevance of explanation for resources similarity, provided by AI systems. In particular, the analysis focuses on one system based on Large Language Models (LLMs)—namely ChatGPT— and on one logic-based tool relying on the computation of the Least Common Subsumer in the Resource Description Framework (RDF). This discussion reveals the need for a formal distinction between relevant and irrelevant information, that we try to answer with a definition of relevance amenable to implementation.

Keywords: Explainable Artificial Intelligence (XAI), Explanation Relevance, Large Language Models (LLMs), Resource Description Framework (RDF), Least Common Subsumer (LCS)

1 Introduction

The role of Artificial intelligence (AI) in providing business value is nowadays universally recognized and has been widely investigated [10]. In particular, it has been discussed the opportunity of embedding AI techniques in processes traditionally performed by humans—like strategical and organizational decision making [24], recruitment [23], and corporate training [7]. A key feature of AI systems is their ability to constantly acquire new knowledge from the context, either by learning or by reasoning on formal models. This ability is crucial in supporting business organizations to rapidly adapt to changing conditions, or, in other words, to pursue *agility*.

The observations above motivate the investigation on a model of governance for businesses based on AI [19]. The work by Schneider *et al.* [19] also collects

main open challenges in this kind of governance. Among them, the authors refer to two undesirable features of AI output: i) it is often not understandable; ii) some results are beyond the control of an organization.

In fact, the ability of AI systems to provide their users an explanation for a given behavior is nowadays considered as an important feature [17], improving, among others, trustworthiness and accountability of the system. In particular, a *logic-based* approach has been advocated [12] to compute and validate explanations [9]. However, methodologies for computing explanations usually do not consider the user whom the explanation is given to—in Miller’s words [17, p.29], “an explanation is an interaction between two roles: explainer and explainee”. In this respect, the various parts an explanation can be made of should be *relevant* for the user—not just trivially true—what Miller summarizes as *epistemic relevance* in explanations [17, p.38].

In the realm of logic-based explanations, in this paper we focus on relevance in explaining similarity of RDF resources, *i.e.*, logical methods for constructing (relevant) explanations about why two or more RDF resources were declared similar—or equivalently, why they were put in the same cluster. In a recent paper [5], Colucci *et al.* explored the verbalization of logical explanations regarding why some tenders in a Public Procurement knowledge graph were clustered by standard algorithms, like k-Means. Colucci *et al.* made use of Least Common Subsumers (LCS) [4] of RDF resources to verbalize an explanation of the similarities among them. However, in the examples presented, pertinent phrases were mixed with general information. For instance [5, Fig.2], all recorded tenders were completed, so reporting the information of a “status complete”, or the information that all tenders referred to a classification schema called “Common Procurement Vocabulary (CPV)”, while being true, was not epistemically relevant. In contrast, the information that all tenders in the cluster referred to truck transportation was relevant in the general context of tenders.

The problem is that all information present in an LCS of a set of resources S is “relevant” in the usual sense, in that it describes some characteristic which truly belongs to each single resource. However, the epistemically relevant information should be one that

1. *discriminates* resources in S with respect to other resources not in S (what has been also called *contrastive* explanation in other contexts), and
2. *is not already known* by the user—which appeals to the fact that the explanatory service should maintain a *model* of the user being informed.

In the following sections, we analyze some answers of the system by Colucci *et al.*, pointing out the irrelevant parts of the explanations it gives. We analyze also a state-of-the-art Large Language Model (LLM), namely ChatGPT, and find that LLMs too fail to make a distinction between relevant and irrelevant information in explaining the similarity of some items. To overcome such limitations, in Section 4 we attempt a logical definition of relevance addressing Points 1–2 above, which may lead us—in future research—to mechanize the computation of relevant explanations for similarity and clusterization. A final section concludes the paper.

2 Comparing LCS with LLM

LLMs have recently reached an indisputable fame and become rather pervasive in real life, mostly thanks to the public availability of ChatGPT³, a chatbot providing detailed and articulate answers to questions across different knowledge domains. In a nutshell, the term LLM denotes deep learning models with a couple of distinguishing characteristics: (i) their number of parameters is at least in the order of billions; (ii) they are general-purpose systems, trained for excellently performing a wide set of tasks.

ChatGPT implements as a Web API the Generative Pre-trained Transformer 3.5 (GPT-3.5)[3], a LLM developed by OpenAI to produce natural language text on the basis of deep learning techniques [18]. The answers returned by ChatGPT are quite impressive in terms of information richness and readability, as every user may verify by querying the chatbot.

As an example, by asking "What do the drugs "cefepime" and "ceftazidime" have in common?", ChatGPT⁴ returns the explanation below:

Cefepime and ceftazidime are both cephalosporin antibiotics, which means they belong to a class of medications that are used to treat bacterial infections. They both work by inhibiting the synthesis of the bacterial cell wall, which leads to the death of the bacteria and the resolution of the infection. Cefepime and ceftazidime are similar in that they both have a broad spectrum of activity and can be used to treat a wide range of infections caused by different types of bacteria. They are both active against many gram-positive and gram-negative bacteria, including some species of Pseudomonas and Escherichia coli. However, cefepime and ceftazidime differ in some respects. Cefepime has a more extended spectrum of activity than ceftazidime, meaning it is effective against a wider range of bacteria. Cefepime is also more resistant to certain types of beta-lactamases, enzymes produced by bacteria that can break down cephalosporin antibiotics and make them less effective. Ceftazidime is more often used to treat infections in the respiratory tract, urinary tract, and skin, while cefepime is more often used to treat infections in the blood and urinary tract.

This answer is learned by facts collected in the Web from published documents.

Nevertheless, knowledge is not always available in such a plain format: Linked Open Data (LOD) [20] also represents a valuable information source for answering users questions. In this context, knowledge is available in RDF syntax [11] and usually publicly accessible.

As an example, we extracted from Bioportal⁵ the explanation in Figure 1 for the similarity of the same two resources (through an independent implementation of the verbalization tool by Colucci *et al.* [5]).

The reader may notice that the returned answer is less human-readable than the one from ChatGPT, but rather detailed and informative.

³ <https://chat.openai.com/>

⁴ Queried on December, 28th 2022

⁵ <https://bioportal.bioontology.org/>

- 1)They are all Organic Chemical that is Chemical Viewed Structurally ;
- 2)They are all Antibiotic that is Pharmacologic Substance
- 3)They are all C preparations that is Classification
- 4)They are all C preparations that is Drug Products by Generic Ingredient Combinations
- 5)They all has_ingredient something that is Organic Chemical that is Chemical Viewed Structurally ;
- 6)They all has_ingredient something that is Antibiotic that is Pharmacologic Substance ;
- 7)They all has_physiologic_effect Decreased Cell Wall Synthesis & Repair that is Organ or Tissue Function ;
- 8)They all has_physiologic_effect Decreased Cell Wall Synthesis & Repair that is Cell Wall Alteration;
- 9)They all may_treat Serratia Infections that is Enterobacteriaceae Infections;
- 10)They all may_treat something that is Enterobacteriaceae Infections;
- 11)They all may_treat Urinary Tract Infections that is Infection;
- 12)They all may_treat something that is Infection;
- 13)They all may_treat Acinetobacter Infections that is Moraxellaceae Infections ;
- 14)They all may_treat Escherichia coli Infections that is Enterobacteriaceae Infections ;
- 15)They all may_treat something that is Enterobacteriaceae Infections;
- 16)They all may_treat Neutropenia that is Agranulocytosi;
- 17)They all may_treat Pneumonia, Bacterial that is Bacterial Infections ;
- 18)They all may_treat something that is Bacterial Infections ;
- 19)They all may_treat Haemophilus Infections that is Pasteurellaceae Infections ;
- 20)They all may_treat Streptococcal Infections that is Gram-Positive Bacterial Infections ;
- 21)They all may_treat something that is Gram-Positive Bacterial Infections ;
- 22)They all may_treat Proteus Infections that is Enterobacteriaceae Infection;
- 23)They all may_treat Fever that is Finding ;
- 24)They all may_treat Fever that is Body Temperature Changes ;
- 25)They all may_treat Fever induced_by PLAGUE VACCINE INJ ;
- 26)They all may_treat Sepsis that is Infection ;
- 27)They all may_treat Pseudomonas Infections that is Gram-Negative Bacterial Infections ;
- 28)They all may_treat Klebsiella Infections that is Enterobacteriaceae Infections ;
- 29)They all may_treat Bone Diseases, Infectious that is Bone Diseases ;
- 30)They all may_treat Skin Diseases, Bacterial that is Bacterial Infections;
- 31)They all are contraindicated with Drug Hypersensitivity that is Hypersensitivity ;
- 32)They all has_mechanism_of_action Enzyme Inhibitors that is Molecular Function ;
- 33)They all share has_mechanism_of_action Enzyme Inhibitors that is Enzyme Interactions ;

Fig. 1. Explanation of the commonalities between "cefepime" (<http://purl.bioontology.org/ontology/NDFRT/N0000022054>) and "ceftazidime" (<http://purl.bioontology.org/ontology/NDFRT/N0000145931>) computed through an independent implementation of the verbalization tool by Colucci *et al.* [5]

Unfortunately, if we query ChatGPT about knowledge available *only* on the Web of Data, returned answers look like not exploiting RDF knowledge content.

For instance, TheyBuyForYou (TBFY) knowledge graph [21] includes an ontology for procurement data, based on the Open Contracting Data Standard (OCDS)[22]. Figure 2 shows the answer by ChatGPT, when asked for comparing two referenced contracting processes:

`http://data.tbfy.eu/contractingProcess/ocds-0c46vo-0133-026258-2019` and
`http://data.tbfy.eu/contractingProcess/ocds-0c46vo-0001-76e76119-992d-40ef-8444-7b020809ff81` .



explain similarity of "http://data.tbfy.eu/contractingProcess/ocds-0c46vo-0001-76e76119-992d-40ef-8444-7b020809ff81" and "http://data.tbfy.eu/contractingProcess/ocds-0c46vo-0133-026258-2019"



The URLs "<http://data.tbfy.eu/contractingProcess/ocds-0c46vo-0001-76e76119-992d-40ef-8444-7b020809ff81>" and "<http://data.tbfy.eu/contractingProcess/ocds-0c46vo-0133-026258-2019>" are both links to procurement information on the TheyBuyForYou website. They both represent a specific procurement process and they are similar in that they both contain the same prefix "<http://data.tbfy.eu/contractingProcess>" which suggests that they are both part of the same dataset of public procurement information.

The similarities don't end there, they both contain the prefix "ocds-0c46vo" which is a standard format for Open Contracting Data Standard (OCDS) and it's used to identify procurement information. The rest of the string of characters in the URL are unique identifiers that are specific to each procurement.

It can be inferred that these two procurement processes are related in some way and they were released by the same organization, but without more information on the procurement it's hard to know more about the similarity.

Fig. 2. Chat, dated January 16, 2023, about the similarity of two referenced contracting processes in TBFY: `http://data.tbfy.eu/contractingProcess/ocds-0c46vo-0133-026258-2019` and `http://data.tbfy.eu/contractingProcess/ocds-0c46vo-0001-76e76119-992d-40ef-8444-7b020809ff81`

By carefully inspecting the answer in Figure 2, the reader may notice that only generic information about the knowledge graph structure is returned. The chatbot ends up to admit the hardness of providing an explanation for similarity.

Yet, the verbalization tool by Colucci *et al.* [5] is able to provide the explanation in Figure 3 for the similarity of the same two referenced contracting processes, by exploiting only the knowledge in TBFY.

The explanation in Figure 3 is built on the computation of a Common Subsumer (CS) of the two contracting processes, as detailed in the work by Colucci

The resources in analysis present the following properties in common:

- 1) They all have a release referencing some resource
 - which has publisher name "Open Opps"
 - and has publisher schema "Companies House"
 - and has publisher web page "https://openopps.com"

- 2) They all present a tender referencing some resource
 - which require a specific item(s) referencing some resource
 - which has classification code "34134000 (Flatbed and Tipper trucks)"
 - and has classification schema "Common Procurement Vocabulary (CPV)"

 - and has tender status "complete"

Fig. 3. Verbal explanation for the similarity of two referenced resources: <http://data.tbfy.eu/contractingProcess/ocds-0c46vo-0133-026258-2019> and <http://data.tbfy.eu/contractingProcess/ocds-0c46vo-0001-76e76119-992d-40ef-8444-7b020809ff81>

et al. [5]. Figure 4 shows the CS knowledge graph, as an evidence of the existing commonalities between the two contracting processes.

The reader may agree on the greater informativeness of the explanation in Figure 3 w.r.t. the one returned by ChatGPT in Figure 2.

3 A critical view on relevance in explanations

In the previous section, we showed by example that the informative richness of explanation services depends on the underlying knowledge model. In particular, when the domain is modelled in RDF, a LCS-based approach seems to be able to return richer explanations than ChatGPT. On the contrary, when knowledge is available in web documents, LLMs-based tools reach an impressive informative level.

By the way, none of the above-mentioned approaches copes with the problem of providing explanations which are *epistemically relevant* to the querying entity. Indeed, often explanations include obvious information, describing a common knowledge in the specific query domain.

To support our thesis, we refer to examples both in ChatGPT and in the verbalization tool by Colucci *et al.* [5]. Recalling Point 2 in the Introduction, the irrelevance of an explanation may depend on the knowledge of the user to be informed: the answer should provide information that the user does not already know.

First, we go back to ChatGPT answer to the question: "What do the drugs "cefepime" and "ceftazidime" have in common?". The first sentence in the response includes a definition for antibiotics (*.....antibiotics, which means they belong to a class of medications that are used to treat bacterial infections*) that, although true, is obvious in the medical domain and known by every physician. An explanation targeted to the medical field should omit such an information.




Fig. 4. A Common Subsumer (CS) of the two referenced resources: <http://data.tbfy.eu/contractingProcess/ocds-0c46vo-0133-026258-2019> and <http://data.tbfy.eu/contractingProcess/ocds-0c46vo-0001-76e76119-992d-40ef-8444-7b020809ff81>

Also by looking at the explanation in Figure 1, the reader may notice that at least lines 12 and 18 are completely irrelevant to physicians, as well as the second part of lines 24, that includes a definition for "fever".

Let us consider now the answer in Figure 2; the second paragraph includes a reference to a common prefix:

...they both contain the same prefix "http://data.tbfy.eu/contractingProcess" which suggests that they are both part of the same dataset of public procurement information.

which is irrelevant to every user able to refer to specific TBFY URIs in the query. The answer to such a query could omit the information about common prefix without any informative loss.

For further examples, we apply the LCS-based verbalization tool to clustering results in TBFY. In particular, all contracting processes released on January,30th 2019 have been clustered with K-means [13] algorithm⁶ and the smallest cluster has been explained in terms of commonalities (on the basis of the LCS of all items in it).

The resulting explanation is given in Figure 5.

The resources in analysis present the following properties in common:

```
1) They all have a release referencing some resource
   which has publisher name "Open Opps"
   and has publisher schema "Companies House"
   and has release publisher "TICON UK LIMITED"
   and has publisher web page "https://openopps.com"
   and has release date "30 January 2019"
```

Fig. 5. Explanation (obtained by using the Common Subsumer technology) of the commonalities in the smallest cluster returned by clustering with k-Means all contracting processes released on January 30, 2019.

The last sentence in Figure 5 states that all contracting process in the cluster have been released on January 30, 2019, which is a feature common to all items in the clustered dataset. This causes the information to be irrelevant to the audience. In other words, when we search for commonalities in a subset of resources, all features that are also common to larger sets should be not put in evidence in the explanation.

All observations above ask for a formal definition of explanation *relevance* in all its facets. Thus, we try to formalize the main aspects of such definitions in the next section.

⁶ The implementation at <https://scikit-learn.org/stable/modules/generated/sklearn.cluster.KMeans.html> has been used

4 Defining relevance in CS-based explanations

There have been several attempts to define relevance in symbolic Artificial Intelligence [16, 14]. However, such attempts concentrate on *relevance for reasoning*, a concept which is tightly coupled with independence [15]—a formula ψ is irrelevant for another formula ϕ if the interpretation of ϕ is independent from the interpretation of ψ . The concept of relevance was also studied in the Information Retrieval (IR) research field [2]. However, relevance in IR is about document relevance, as it aims “. . . to retrieve all the relevant documents [and] at the same time retrieving as few of the non-relevant as possible.” [25, p.6]. This aim leads to definitions that may either involve a subjective judgement about the retrieval task and the user’s needs [2], or a probabilistic analysis of relevance based on documents descriptions [26]. In both cases, we did not find connections to the concept of relevance in explaining similarity.

In our case, we are interested on the relevance of communicating *parts* of an LCS, when both the LCS itself, and the background knowledge about a user, are represented in RDF. To the best of our knowledge, definitions in the field of IR do not cope with RDF/RDFS; thus, we provide new, RDF-specific definitions for relevance. In our past research, we already dealt with problem of providing informative (L)CS [8], but w.r.t. Description Logics (DLs) [1].

We concentrate first on *relevance in a context*, meaning that some characteristics, that are in common to a given set of resources S , may not be relevant to communicate because they are common also to larger sets of resources $T \supset S$, while others do. We introduce this kind of relevance with an example based on the previous sections.

Example 1. Suppose T is a set of resources representing all European contracting processes of a given day of the year, say, January 30, 2019, and $S = \{r_1, r_2\} \subset T$ is a pair of such resources that an algorithm like k-Means, while clustering T , puts in the same cluster. Clearly r_1 and r_2 share the same release date “2019-01-30”; however, *in the larger context of T* which these resources were drawn from, this information is irrelevant. If instead r_1 and r_2 were drawn from the set $T' \supset T$ of *all* tenders in the whole year 2019, being released in the same date would become relevant.

Intuitively, a common characteristic P of a set of resources S is relevant in the context of a larger set T , when P is *not* common to the larger set T . In the following, formal definition, we refer to the notion of LCS as defined by Colucci *et al.* [4].

Definition 1. [*Relevance w.r.t. a larger set*] Let T, S be sets of RDF resources, with $S \subset T$, and for each resource $r \in T$, a rooted graph $\langle r, T \rangle$ can be computed. Let $L = \langle x, T \rangle$ the rooted graph representing the Least Common Subsumer of all resources in S , and let $P \subseteq L$ a path in L starting from x .

We say that P is relevant in T if $LCS(T) \not\subseteq P$.

The above definition of relevance may model *contrastive* explanations, in the following sense: given a resource r which has *not* been put into a cluster S of

resources, a contrastive explanation about “Why resources in S were put in the same cluster, while r was not?” can be formed by taking the LCS of S , and the LCS of $T = S \cup \{r\}$, and finding a characteristic which is in the former but not in the latter.

A different notion of relevance comes up when user’s knowledge—at least that part that can be expressed in RDF—can be taken into account. Again, we introduce this aspect with another example:

Example 2. When identifying the common characteristics of two antibiotics in a drugs databank [6], a common characteristics that is—correctly—found is that they both “are used in bacterial infections”. This information may be useful for a generic user, while for a physician it would be trivially irrelevant. In this case, the knowledge that the system implicitly attributes to the user comes into play. The system should be able to distinguish at least between a generic user and a physician, and choose epistemically relevant information accordingly.

Intuitively, a characteristic P is relevant for a user u when it is not part of u ’s prior knowledge; a formal definition can be made as follows:

Definition 2. [*Relevance w.r.t. a user’s knowledge*] Let S be a set of RDF resources, and for each resource $r \in S$, a rooted graph $\langle r, S \rangle$ can be computed. Let $L = \langle x, S \rangle$ the rooted graph representing the Least Common Subsumer of all resources in S , and let $P \subseteq L$ a path in L starting from x . Moreover, let K_u be the RDF-graph representing the knowledge (expressible in RDF) of a given user u .

We say that P is relevant for u if $K_u \not\supseteq P$.

Observe that it is not necessary to elicit K_u from the user through a long—and presumably tedious—knowledge elicitation process. As a first attempt, the system may use as K_u a domain ontology, whose knowledge can be commonly attributed to specific user categories, such as physicians or procurement brokers.

5 Conclusion

In this paper, we analyzed the explanations about the similarity of two or more resources, drawn from the Business domain of Public Procurement, and the biological domain of drugs. The compared explanations were given by the RDF-logic-based system of Colucci *et al.* [5] and the LLM-based system of ChatGPT. We highlighted that both approaches mix relevant and irrelevant information, where relevance can be evaluated based on (i) either a larger context of resources, (ii) or prior user’s expert knowledge. Finally, for the case of knowledge expressed in RDF, we provided two objective, logical definitions of relevance of information in such explanations, one for each of the above characterizations of relevance.

Regarding future work, we plan to implement the definitions of relevance given in the previous section by adding an RDF Reasoner in the pipeline building the linguistic realization of the LCS. To exploit Definition 1, we plan to compute

the LCS of a larger set of resources, and check for non-implication. In this way, our tool for providing explanations may avoid verbalizing trivial information about commonalities that are also common to supersets. Regarding Definition 2, we do not plan for the moment to try to model knowledge of single users—this would be too time-consuming. Instead, we plan to use RDF ontologies for specific domains, like those in Bioportal⁷, which can model common knowledge of users expert of the domain.

6 Acknowledgements

Projects Regione Lazio-DTC/“SanLo” (CUP F85F21001090003) and Ministero dello Sviluppo Economico/“Casa delle Tecnologie Emergenti dell’area metropolitana di Bari: Bari Open Innovation Hub” (CUP J99J19000300003) partially supported this work.

References

1. Baader, F., Calvanese, D., McGuinness, D., Patel-Schneider, P., Nardi, D.: The description logic handbook: Theory, implementation and applications. Cambridge university press (2003)
2. Borlund, P.: The concept of relevance in ir. *Journal of the American Society for Information Science and Technology* 54(10), 913–925 (2003), <https://onlinelibrary.wiley.com/doi/abs/10.1002/asi.10286>
3. Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J.D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al.: Language models are few-shot learners. *Advances in neural information processing systems* 33, 1877–1901 (2020)
4. Colucci, S., Donini, F., Giannini, S., Di Sciascio, E.: Defining and computing least common subsumers in RDF. *Web Semantics: Science, Services and Agents on the World Wide Web* 39, 62 – 80 (2016)
5. Colucci, S., Donini, F.M., Iurilli, N., Sciascio, E.D.: A business intelligence tool for explaining similarity. In: Babkin, E., Barjis, J., Malyzhenkov, P., Merunka, V. (eds.) *Model-Driven Organizational and Business Agility - Second International Workshop, MOBA 2022, Leuven, Belgium, June 6-7, 2022, Revised Selected Papers. Lecture Notes in Business Information Processing*, vol. 457, pp. 50–64. Springer (2022), https://doi.org/10.1007/978-3-031-17728-6_5
6. Colucci, S., Donini, F.M., Sciascio, E.D.: Logical comparison over RDF resources in bio-informatics. *J. Biomed. Informatics* 76, 87–101 (2017), <https://doi.org/10.1016/j.jbi.2017.11.004>
7. Colucci, S., Noia, T.D., Sciascio, E.D., Donini, F.M., Ragone, A.: Semantic-based skill management for automated task assignment and courseware composition. *J. Univers. Comput. Sci.* 13(9), 1184–1212 (2007), <https://doi.org/10.3217/jucs-013-09-1184>
8. Colucci, S., Tinelli, E., Sciascio, E.D., Donini, F.M.: Automating competence management through non-standard reasoning. *Eng. Appl. Artif. Intell.* 24(8), 1368–1384 (2011), <https://doi.org/10.1016/j.engappai.2011.05.015>

⁷ <https://bioportal.bioontology.org/ontologies>

9. Cooper, M.C., Marques-Silva, J.: Tractability of explaining classifier decisions. *Artificial Intelligence* 316, 103841 (2023), <https://www.sciencedirect.com/science/article/pii/S0004370222001813>
10. Enholm, I.M., Papagiannidis, E., Mikalef, P., Krogstie, J.: Artificial intelligence and business value: A literature review. *Information Systems Frontiers* 24(5), 1709–1734 (2022)
11. Hayes, P., Patel-Schneider, P.F.: RDF 1.1 semantics, W3C recommendation (2014), <http://www.w3.org/TR/2014/REC-rdf11-mt-20140225/>
12. Ignatiev, A.: Towards trustable explainable AI. In: *IJCAI*. pp. 5154–5158 (2020)
13. Jin, X., Han, J.: K-Means Clustering, pp. 563–564. Springer US, Boston, MA (2010), https://doi.org/10.1007/978-0-387-30164-8_425
14. Lakemeyer, G.: Relevance from an epistemic perspective. *Artif. Intell.* 97(1-2), 137–167 (1997), [https://doi.org/10.1016/S0004-3702\(97\)00038-6](https://doi.org/10.1016/S0004-3702(97)00038-6)
15. Lang, J., Liberatore, P., Marquis, P.: Propositional independence: Formula-variable independence and forgetting. *J. Artif. Intell. Res.* 18, 391–443 (2003), <https://doi.org/10.1613/jair.1113>
16. Levy, A.Y., Fikes, R., Sagiv, Y.: Speeding up inferences using relevance reasoning: A formalism and algorithms. *Artif. Intell.* 97(1-2), 83–136 (1997), [https://doi.org/10.1016/S0004-3702\(97\)00049-0](https://doi.org/10.1016/S0004-3702(97)00049-0)
17. Miller, T.: Explanation in artificial intelligence: Insights from the social sciences. *Artificial Intelligence* 267, 1 – 38 (2019), <http://www.sciencedirect.com/science/article/pii/S0004370218305988>
18. OpenAI: ChatGPT: Optimizing language models for dialogue. Available (as of April 18, 2023) on the Web, <https://web.archive.org/web/20221130180912/https://openai.com/blog/chatgpt/>
19. Schneider, J., Abraham, R., Meske, C., Brocke, J.V.: Artificial intelligence governance for businesses. *Information Systems Management* 0(0), 1–21 (2022)
20. Shadbolt, N., Hall, W., Berners-Lee, T.: The semantic web revisited. *Intelligent Systems, IEEE* 21(3), 96–101 (2006)
21. Soylyu, A., Corcho, O., Elvesater, B., Badenes-Olmedo, C., Blount, T., Yedro Martinez, F., Kovacic, M., Posinkovic, M., Makgill, I., Taggart, C., Simperl, E., Lech, T.C., Roman, D.: TheyBuyForYou platform and knowledge graph: Expanding horizons in public procurement with open linked data. *Semantic Web* 13(2), 265–291 (2022)
22. Soylyu, A., Elvesæter, B., Turk, P., Roman, D., Corcho, O., Simperl, E., Konstantinidis, G., Lech, T.C.: Towards an ontology for public procurement based on the open contracting data standard. p. 230–237. Springer-Verlag, Berlin, Heidelberg (2019)
23. Tinelli, E., Cascone, A., Ruta, M., Noia, T.D., Sciascio, E.D., Donini, F.M.: I.M.P.A.K.T.: an innovative semantic-based skill management system exploiting standard SQL. In: Cordeiro, J., Filipe, J. (eds.) *ICEIS 2009 - Proceedings of the 11th International Conference on Enterprise Information Systems*, Volume AIDSS, Milan, Italy, May 6-10, 2009. pp. 224–229 (2009)
24. Trunk, A., Birkel, H., Hartmann, E.: On the current state of combining human and artificial intelligence for strategic organizational decision making. *Business Research* 13(3), 875–919 (2020)
25. Van Rijsbergen, C.: *Information Retrieval* (Book 2nd ed). Butterworth-Heinemann Newton, MA, USA (1979)
26. Van Rijsbergen, C.: *Information retrieval: theory and practice*. In: *Proceedings of the joint IBM/University of Newcastle upon Tyne seminar on data base systems*. vol. 79 (1979)