Using Multiple RDF Knowledge Graphs for Enriching ChatGPT Responses*

Michalis Mountantonakis ✉ and Yannis Tzitzikas

Institute of Computer Science - FORTH-ICS, Greece
and Computer Science Department - University of Crete, Greece
{mountant, tzitzik}@ics.forth.gr

Abstract. There is a recent trend for using the novel Artificial Intelligence ChatGPT chatbox, which provides detailed responses and articulate answers across many domains of knowledge. However, in many cases it returns plausible-sounding but incorrect or inaccurate responses, whereas it does not provide evidence. Therefore, any user has to further search for checking the accuracy of the answer or/and for finding more information about the entities of the response. At the same time there is a high proliferation of RDF Knowledge Graphs (KGs) over any real domain, that offer high quality structured data. For enabling the combination of ChatGPT and RDF KGs, we present a research prototype, called GPT•LODS, which is able to enrich any ChatGPT response with more information from hundreds of RDF KGs. In particular, it identifies and annotates each entity of the response with statistics and hyperlinks to LODsyndesis KG (which contains integrated data from 400 RDF KGs and over 412 million entities). In this way, it is feasible to enrich the content of entities and to perform fact checking and validation for the facts of the response at real time.
URL: https://demos.isl.ics.forth.gr/GPToLODS/Annot_Enrichment
Demo Video: https://youtu.be/H30bSv9NfUw.

Keywords: ChatGPT, LOD, RDF, Knowledge Graphs, Annotation

1 Introduction

ChatGPT is a novel Artificial Intelligence (AI) chatbox (https://chat.openai.com/), which is built on GPT-3.5 and GPT-4 families of large language models (LLMs) [2], and provides detailed responses and human-like answers across many domains of knowledge. However, it has not been designed to store or retrieve facts, e.g., like a relational database or a Knowledge Graph (KG). For this reason, in many cases it returns plausible-sounding but incorrect or inaccurate responses [3]. Thereby, it is hard for the user to check the validity of the answers returned

* **This is a preprint of the demo paper: Michalis Mountantonakis and Yannis Tzitzikas, "Using Multiple RDF Knowledge Graphs for Enriching ChatGPT Responses". Accepted for publication in ECML PKDD 2023 Demos Track.**

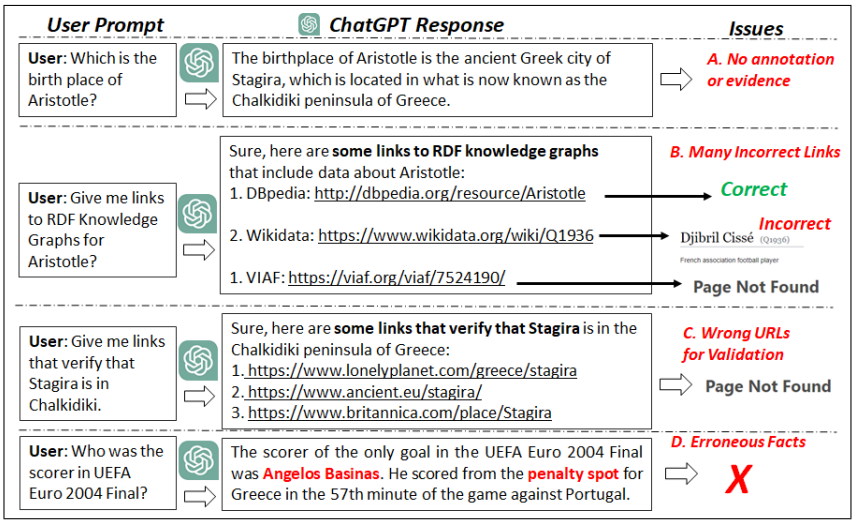

Fig. 1. ChatGPT issues through real conversations (March 27, 2023)

by ChatGPT. Sometimes the responses contain entities that do not exist, URLs that are wrong, facts that cannot be verified, outdated data, and many others. On the contrary, there are available numerous RDF KGs [4] (e.g., DBpedia [5], Wikidata [17], YAGO [15], etc), that provides high quality structured data (that are updated at least periodically), by using Linked Data techniques [1].

Concerning some key issues of ChatGPT (that can be improved by using KGs), Fig. 1 shows a real conversation with ChatGPT (in March 27, 2023). First, we asked about the birth place of Aristotle (issue A), and the output was a plain text, without evidence or annotations. Afterwards, we desired to find available RDF links for Aristotle (issue B), and only one of them was correct, whereas the URI of Wikidata refers to a completely different entity. Then, we asked about sources verifying that Stagira is located in Chalkidiki, and all the returned URIs were invalid (issue C). In the last case, we asked a question from a different domain, i.e., "Who scored the goal in UEFA Euro 2004 Final?", and we retrieved erroneous facts (issue D); its response was "Angelos Basinas from the penalty", however the correct answer is "Angelos Charisteas with a header".

The objective of this paper is to aid the above issues, however, there are several challenges that should be tackled. Indeed, it is quite difficult to check the validity of entities and URIs, since it requires access to numerous KGs, sources and resources in general. For tacking these challenges, we demonstrate the research prototype `GPT•LODS`, that enables the user to make a question, and instead of getting the raw answer from ChatGPT, it retrieves the ChatGPT response annotated with the identified entities, and with relevant data about these entities (URIs, facts and KGs), For making this feasible at real time, we exploit LODsyndesis suite of services [9,11] including an Information Extraction (IE) service [12]. These services rely on LODsyndesis KG; a large KG, equipped with special indexes and algorithms, that has integrated 2 billion facts for 412 million entities, from 400 real RDF KGs. In this way, the user is able to speed

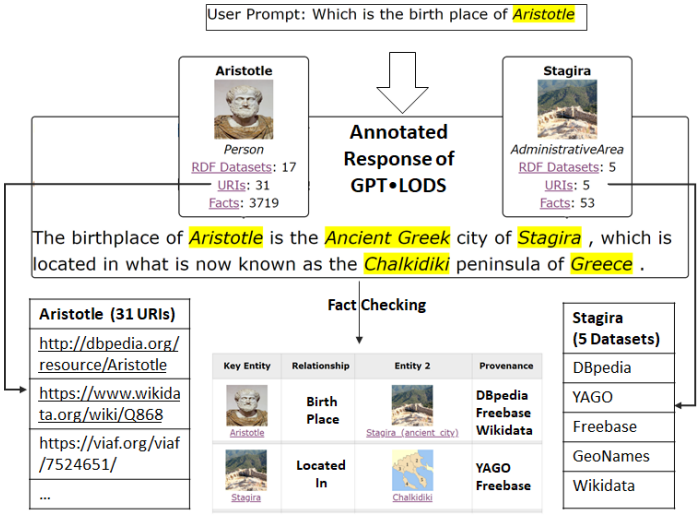

Fig. 2. Screenshots from GPT•LODS research prototype

up the validation of the ChatGPT response and to retrieve more information. Concerning the novelty, to the best of our knowledge it is the first system offering annotation and linking of the ChatGPT responses to hundreds of KGs.

The rest of this demo paper is organized as follows; 2 describes the related work, 3 presents the process and the use cases and 4 concludes the paper.

2 Related Work

First, concerning ChatGPT and KGs, [13] provides a comparison for the Question Answering task and they concluded that ChatGPT can have high precision in general knowledge domains but very low scores in unseen domains, compared to a KG-based approach. Regarding IE tools for Entity Recognition, there are available approaches from several areas, i.e., from Natural Language Processing (NLP)[6], from KGs [7,14,8], and from Neural Networks [16]. However, these tools link the entities to a single KG, and for this reason GPT•LODS uses the machinery of LODsyndesisIE [12], which combines tools from NLP and KG, and links the entities to 400 RDF KGs. Regarding the novelty, to the best of our knowledge there is no other related system that annotates the response of ChatGPT and provides links and services using hundreds of KGs at real time.

3 The Process of GPT•LODS and Use Cases

First, the name of the prototype GPT•LODS comes from the mathematical notation for function composition, i.e. $(GPT \bullet LODS)(x) = LODS(GPT(x))$, where $LODS$ comes from LODSyndesis [11]. Concerning the process (see Fig. 2), the user submits through the web application a question (in english), GPT•LODS sends

the question to the ChatGPT API and retrieves the response. Afterwards, it applies the machinery of LODsyndesisIE (described in [12]) for recognizing the entities of the response, i.e., it combines widely used Entity Recognition tools (i.e., DBpedia Spotlight [7], WAT [14] and Stanford CoreNLP[6]) for recognizing the entities. The next step is to further process the response for creating the annotations and for adding statistics and links to LODsyndesis KG for each recognized entity, and finally it returns the annotated response to the user. Fig. 2 shows that by clicking on each entity, one can see its name, image, type and statistics, such as the number of its RDF KGs, URIs and facts that occur in LODsyndesis. By clicking on the links the user can browse (or download) all this data (e.g., the list of all the URIs of each entity). In addition, many other services are offered, including a fact checking service (see Fig. 2) that shows all the relations between any pair of entities of the response (for fact validation). Finally, `GPT•LODS` can be easily extended for annotating and enriching the response of any new LLM in case of offering an analogous API as ChatGPT.

Scenario & Use Cases. Below, we present a scenario which will be demonstrated in the conference. The target audience can be any researcher of AI area, since `GPT•LODS` (which is accessible in `https://demos.isl.ics.forth.gr/GPToLODS`, including a link to github) is a research prototype that combines tools and techniques of several AI components, including LLMs, NLP, and Knowledge Representation and Reasoning (Linked Data and KGs). The scenario is about the user questions (and needs) shown in Fig. 1, i.e., starting with the question "Which is the birth Place of Aristotle". For tackling the needs of this scenario, the results include: the annotated entities, related information (identifiers, facts and datasets) about these entities and fact validation. The scenario can be accessed in a video (`https://youtu.be/H30bSv9NfUw`), that presents the issues of Fig. 1 and how they can be solved through the use cases.

Use Case 1. Annotation, Evidence and Linking. This refers to the issues A and B of Figure 1, i.e., by having the annotation of the entities of the response, we are able to find more information (links, datasets and facts) for the entities of the response. Moreover, we can have access to the correct URIs for each entity, e.g., in Fig. 2 we retrieved the correct Wikidata and VIAF URIs for Aristotle (e.g., `https://www.wikidata.org/wiki/Q868`), and in total 31 URIs.

Use Case 2. Fact Validation and Correct Answer. In Fig. 1 we can see that a ChatGPT response can either provide wrong links for validation or even erroneous facts (issues C,D). In the first case, `GPT•LODS` can be used for fact validation, e.g., in Fig. 2 it verified 2 facts of the response from popular KGs, like DBpedia and Wikidata. Regarding issue D, `GPT•LODS` will not find the erroneous fact in LODsyndesis KG. For finding the correct answer, the user can further browse all the facts of an entity (by clicking on the corresponding link).

Use Case 3. Dataset Discovery and Enrichment. The user can discover all the datasets of each entity (e.g., see the lower right part of Fig. 2), and all (or a part of) the facts of that entity in the KGs that are included in LODsyndesis. This can be useful for enriching the available content of each entity, e.g., for creating an application, a data warehouse [10], for performing an analysis, etc.

4 Concluding Remarks

In this paper, we presented the research prototype GPT•LODS, which enables the real time annotation and linking of a ChatGPT response to hundreds of RDF KGs, the enrichment of its entities and the validation of its facts. As a future work, we plan to improve the GUI, and the fact checking service by performing relation extraction, to offer a REST API and to support multilinguality.

References

1. Bizer, C., Heath, T., Berners-Lee, T.: Linked data: The story so far. In: Semantic services, interoperability and web applications: emerging concepts, pp. 205–227. IGI global (2011)
2. 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)
3. van Dis, E.A., Bollen, J., Zuidema, W., van Rooij, R., Bockting, C.L.: ChatGPT: five priorities for research. Nature **614**(7947), 224–226 (2023)
4. Hogan, A., Blomqvist, E., Cochez, M., d'Amato, C., Melo, G.d., Gutierrez, C., Kirrane, S., Gayo, J.E.L., Navigli, R., Neumaier, S., et al.: Knowledge graphs. ACM Computing Surveys (CSUR) **54**(4), 1–37 (2021)
5. Lehmann, J., Isele, R., Jakob, M., Jentzsch, A., Kontokostas, D., et al.: DBpedia–a large-scale, multilingual knowledge base extracted from wikipedia. Semantic web **6**(2), 167–195 (2015)
6. Manning, C.D., Surdeanu, M., Bauer, J., Finkel, J.R., Bethard, S., McClosky, D.: The Stanford CoreNLP natural language processing toolkit. In: Proceedings of 52nd annual meeting of the association for computational linguistics: system demonstrations. pp. 55–60 (2014)
7. Mendes, P.N., Jakob, M., García-Silva, A., Bizer, C.: DBpedia spotlight: shedding light on the web of documents. In: Proceedings of the 7th international conference on semantic systems. pp. 1–8 (2011)
8. Moro, A., Cecconi, F., Navigli, R.: Multilingual word sense disambiguation and entity linking for everybody. In: ISWC (Posters & Demos). pp. 25–28. Citeseer (2014)
9. Mountantonakis, M.: Services for Connecting and Integrating Big Numbers of Linked Datasets, vol. 50. IOS Press (2021)
10. Mountantonakis, M., Tzitzikas, Y.: Large-scale semantic integration of linked data: A survey. ACM Computing Surveys (CSUR) **52**(5), 1–40 (2019)
11. Mountantonakis, M., Tzitzikas, Y.: Content-based union and complement metrics for dataset search over RDF knowledge graphs. ACM JDIQ **12**(2), 1–31 (2020)
12. Mountantonakis, M., Tzitzikas, Y.: Linking entities from text to hundreds of RDF datasets for enabling large scale entity enrichment. Knowledge **2**(1), 1–25 (2022)
13. Omar, R., Mangukiya, O., Kalnis, P., Mansour, E.: ChatGPT versus traditional question answering for knowledge graphs: Current status and future directions towards knowledge graph chatbots. arXiv preprint arXiv:2302.06466 (2023)
14. Piccinno, F., Ferragina, P.: From tagme to wat: a new entity annotator. In: Proceedings of the first international workshop on Entity recognition & disambiguation. pp. 55–62 (2014)

15. Rebele, T., Suchanek, F., Hoffart, J., Biega, J., Kuzey, E., Weikum, G.: YAGO: A multilingual knowledge base from wikipedia, wordnet, and geonames. In: The Semantic Web–ISWC 2016: 15th International Semantic Web Conference, Kobe, Japan, October 17–21, 2016, Proceedings, Part II 15. pp. 177–185. Springer (2016)
16. Van Hulst, J.M., Hasibi, F., Dercksen, K., Balog, K., de Vries, A.P.: Rel: An entity linker standing on the shoulders of giants. In: Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval. pp. 2197–2200 (2020)
17. Vrandečić, D., Krötzsch, M.: Wikidata: a free collaborative knowledgebase. Communications of the ACM **57**(10), 78–85 (2014)