

Tiresias: Bilingual Question Answering over DBpedia Abstracts through Machine Translation and BERT

Michalis Mountantonakis^{1,2,*}, Michalis Bastakis^{2,†}, Loukas Mertzanis^{1,2,†} and Yannis Tzitzikas^{1,2}

¹Institute of Computer Science, FORTH, Heraklion, Greece

²Department of Computer Science, University of Crete, Heraklion, Greece

Abstract

There is a high need for providing multilingual solutions to Natural Language Processing tasks, such as for Question Answering, since a lot of people do not speak and read English. In this way, DBpedia abstracts, which are offered in several languages, can be exploited, since they cover valuable information which are never transformed to structured data (to triples), whereas for the same entity, the abstract in different languages can contain complementary information. For making it feasible to answer simple open domain factoid questions over DBpedia abstracts, we introduce *Tiresias*, a research prototype that supports bilingual Question Answering (QA) over DBpedia abstracts. In particular, it receives a question either in Greek or in English language, and by exploiting Named Entity Recognition models for recognizing the entity of the question, the DBpedia abstracts written in the mentioned languages for the identified entity, Machine Translation (MT) tools, and BERT QA models (pretrained in English corpus), it produces the final answer. Concerning the evaluation, we provide experimental results about the effectiveness and efficiency of MT and QA process by using a Greek evaluation collection, and we present statistics of DBpedia abstracts and several use cases with real examples.

Keywords

Multilingual QA, DBpedia Abstracts, BERT, Greek Language, Machine Translation, Deep Learning

1. Introduction

It is of primary importance to provide multilingual solutions to Natural Language Processing (NLP) tasks, such as for Question Answering (QA), given that a lot of people worldwide do not speak and read English, i.e., only 17% of people worldwide speak English¹. Even in popular tourist destinations, such as Greece, a high percentage of people cannot speak and read English, i.e., 49% in Greece². For aiding such tasks, we believe that DBpedia abstracts [1], which are offered in several languages [2] can be exploited. Indeed, they cover valuable information that is not always transformed to structured data, whereas for the same entity the abstract in different languages can contain complementary information. In particular, the DBpedia abstracts are derived from Wikipedia, where each page is offered in one or more languages, and for each

[†]These authors contributed equally.

✉ mountant@ics.forth.gr (M. Mountantonakis); mbastakis@gmail.com (M. Bastakis); mertzanis@ics.forth.gr (L. Mertzanis); tzitzik@ics.forth.gr (Y. Tzitzikas)



© 2022 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

CEUR Workshop Proceedings (CEUR-WS.org)

¹<https://www.babbel.com/en/magazine/how-many-people-speak-english-and-where-is-it-spoken>

²<https://doublespeakdojo.com/how-common-is-spoken-english-in-greece/>

language different information can be covered, especially for entities that are more related to a specific country (and language). As an example, see the DBpedia abstract for *Xanthippe*, the wife of Socrates: i) the following text “She was likely much younger than Socrates, perhaps by as much as 40 years”, which is part of both English and Greek abstract, has not been transformed to structured data (i.e., to RDF triples), whereas ii) the Greek abstract is quite larger comparing to the English one, and contains complementary information (that is not offered in English).

For supporting bilingual factoid Question Answering (QA) over DBpedia abstracts for questions like “How many years was Xanthippe younger than Socrates?”, we introduce *Tiresias*, a web application that follows a generic pipeline that can be adopted for any language, for answering such questions over Greek and English, by using as input the DBpedia abstracts in both languages. In particular, a) *Tiresias* receives a question in English or Greek, b) it recognizes the main entity in the question by using Named Entity Recognition (NER) tools, such as DBpedia Spotlight [3], c) it retrieves the abstract of the entity in both English and Greek, through a SPARQL query, d) it translates any Greek input (question, context) in English through Machine Translation (MT) tools (e.g., Bing), e) it produces the final answer in English through BERT QA models (pretrained in an English corpus) by using the DBpedia abstracts as context, and f) it can possibly translate back the answer to Greek. Moreover, it also supports QA by giving a specific context a priori, i.e., without using NER and DBpedia abstracts.

Regarding the Research Questions, we would like to investigate a) how effective and efficient is the mentioned pipeline, and b) whether by using the proposed approach, we can answer questions that are expressed in English from a Greek text (i.e., abstract), and vice versa. Concerning our contribution, *Tiresias*, which is available in <https://demos.isl.ics.forth.gr/tiresias>, is the first online system offering open domain QA in Greek by using MT and BERT models, and QA in English by exploiting texts in Greek. The web application offers a rich configuration for performing bilingual QA over DBpedia abstracts, by combining MT and BERT models. Regarding the evaluation, we provide experimental results about the efficiency of MT and QA process, statistics of DBpedia abstracts and several use cases with real examples.

The rest of this paper is organized as: §2 discusses related work, §3 presents the process and web UI of *Tiresias*, §4 presents the experimental evaluation and §5 concludes the paper.

2. Related Work

The multilingualism in QA is one major challenge in the area of knowledge graphs and semantic web [4], and recently, there is a trend towards that direction. In particular, [5] use machine translation tools for evaluating QA over knowledge graphs for questions in languages that are unsupported by a KGQA system for 8 different languages. Moreover, [6] uses bilingual lexicon induction (BLI) methods and multilingual models for performing QA over DBpedia. As regards online QA systems, Platypus [7], is a multilingual QA that supports QA for 3 languages by using Wikidata [8], whereas QAnswer [9] provides multilingual QA by exploiting many knowledge bases. Finally, Wikidata is used by DeepPavlov [10] for QA in two languages.

Concerning QA approaches over knowledge bases that use also the textual description of entities, [11] proposed an approach for combining the textual description of entities for QA and triples from DBpedia, for providing a hybrid QA. Regarding QA systems in Greek language,

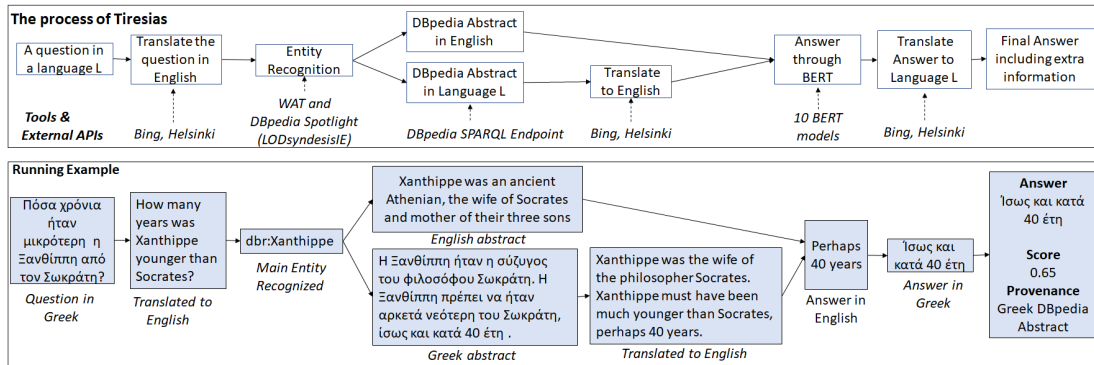


Figure 1: The process of *Tiresias* and an example with a simple factoid question in Greek language

APANTISIS [12] is a QA system that can be plugged in relational databases. Furthermore, there are multilingual evaluation collections over Knowledge graphs, e.g., QALD-9 [13], whereas QALD-7 [14] provides an evaluation collection for Hybrid QA systems in English. Moreover, [15] surveys approaches over QA and Linked Data, including hybrid QA tools that exploit textual information. Finally, concerning approaches offering multilingual QA over textual sources by using MT, [16] transforms both the question and context in English for offering QA for French and Japanese, whereas in [17] a new metric was proposed for evaluating such methods.

Novelty of Tiresias. To the best of our knowledge, *Tiresias* is the first online system offering open domain QA in Greek (by exploiting contexts in both English and Greek), whereas it is the first application offering QA in English by exploiting the Greek version of DBpedia abstracts. Finally, it is the first approach combining MT and BERT models for QA in Greek.

3. The Steps of *Tiresias*

Here, we present the steps of the proposed approach (see the upper side of Fig. 1), by also showing a running example (see the lower side of Fig. 1).

Input. *Tiresias* receives one or more simple factoid questions in a language L (i.e., Greek and English), and possibly a context text. When the context is not given, extra steps are needed for recognizing the question entity and retrieving the context, e.g., see Fig. 1.

Configuration. The user selects the desired language, and which MT and BERT QA models to use. Below, we provide all the possible choices. In particular, the language can be either Greek or English, whereas for Machine Translation two options are offered: *Bing* and *Helsinki* [18]. Concerning BERT models for QA, 10 models are offered: 1) Deepset/roBERTA, 2) DistilBERT-cased, 3) BERT-Mask-uncased, 4) Deepset/Mask-uncased, 5) DistilBERT-uncased, 6) rsvp-ai/Bertserini, 7) deepset/MiniLM, 8) BioBERT, 9) deepset/BERT-cased, and 10) BERT-Mask-Cased.

Step A. Named Entity Recognition (NER). Given a question either in Greek or in English, we perform NER to identify the main entity of the question (the first recognized entity). Indeed we use a combination of WAT [19] and DBpedia Spotlight [3] tools through LODsynthesisIE [20], that offers NER for the English language and returns the DBpedia link of the identified entity. However, an extra step of translation is needed, when the question is given in Greek, e.g., in Fig. 1 we translated the question and we identified the entity “Xanthippe” and its DBpedia link.

Step B. Retrieval of Bilingual Context (DBpedia Abstract). Here, we send a SPARQL query to DBpedia endpoint for retrieving the abstracts in both languages, e.g., English and Greek. In the second case the abstract is translated to English for being used in the next step. By using only the English version, there would be no need for translation, however, this is not always the ideal case. Indeed, there are 3 abstract categories, a) the abstract exists in English but not in L , e.g., for the basketball player *Jason Tatum*, there is not available a Greek abstract, b) the abstract exists in L but not in English, i.e., many Wikipedia pages are not offered in English, e.g., see the Cretan village *Ano Asites*, and c) the abstract exists in both languages. In that case, *Tiresias* uses both abstracts (e.g., see Fig. 1), since they can contain complementary data, e.g., for *Xanthippe*, the Greek DBpedia abstract contains 568 words and the English one only 44 words. Thereby, a question written in English may be answerable only from a Greek context (and vice versa). More examples, showing the impact of using both languages, and statistics are given in Section 4. As we can see in Fig. 1, both abstracts were retrieved for the main entity.

Step C. QA over the (translated) context through BERT QA pretrained models in English. Since we want the process to be generic for covering more languages in the future, we decided to perform MT for translating the non-English context (i.e., the Greek) to English, and then to use BERT QA models, pretrained in English corpus for performing QA over the translated input, since they are quite effective [21]. Another option could be to use either i) a pretrained QA model in a language L or ii) a Multilingual BERT QA model. However, for i), they are not available in many languages, including Greek, and for ii), they do not cover all the languages, and can be slower due to their huge training set size, since they are usually trained in a very large corpus for covering many languages (see also Section 4).

Therefore, as a preprocessing step, we use the desired MT model for translating any Greek input in English. Afterwards, we just use the desired BERT QA model, which has been pretrained for the English language. In particular, any of the available models receives as input the given question and the combination of the two DBpedia abstracts and produces the answer with the highest confidence score in English language, e.g., the answer “perhaps 40 years” in Fig. 1.

Step D. Translate Back the Answer. The final step is to translate back the answer, if the initial question is not given in English, e.g., in Greek. In the end, the final answer, its confidence score, and extra information such as the provenance of the answer are returned to the user. E.g., see the final answer in Fig. 1, which is presented in Greek language.

3.1. The Website & Architecture of *Tiresias* Web Application

Tiresias is available in <https://demos.isl.ics.forth.gr/tiresias>, and it is quite responsive for several devices. It runs on a server with 4GB main memory, 2 cores and 500 GB disk space.

Examples and Screenshots. Some screenshots are shown in Figure 2, i.e., see the configuration, the question, i.e., “Which is the contract of G. Antetokounmpo”, and the corresponding answer, i.e., “\$228 million”. The real example in the upper and in the lower side is the same but for the two supported languages, i.e., English and Greek. In both cases, we used the Helsinki model for MT, and the RoBERTA model for QA, and we retrieved the correct answer. By clicking on the info button, we can retrieve more information about the answer: its provenance, which was the Greek DBpedia abstract in this example, its confidence score, a link to the corresponding DBpedia abstract, and its response time. As we can see it was faster for the English language,

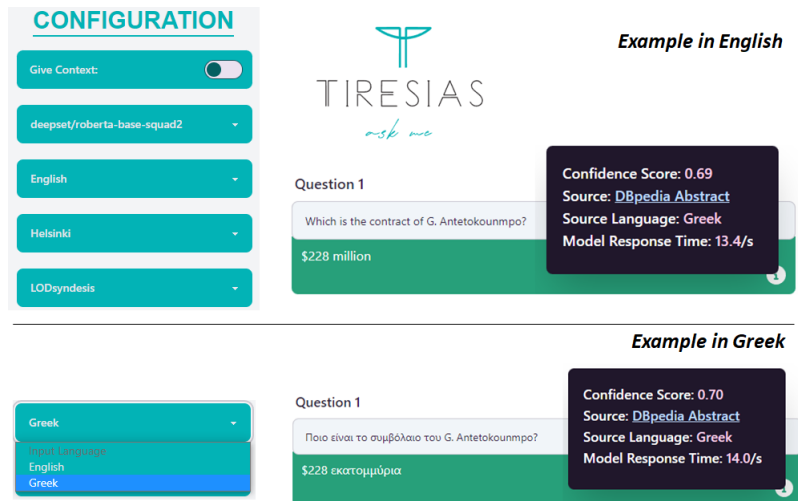


Figure 2: Screenshots- A real example of *Tiresias* in Greek and English language

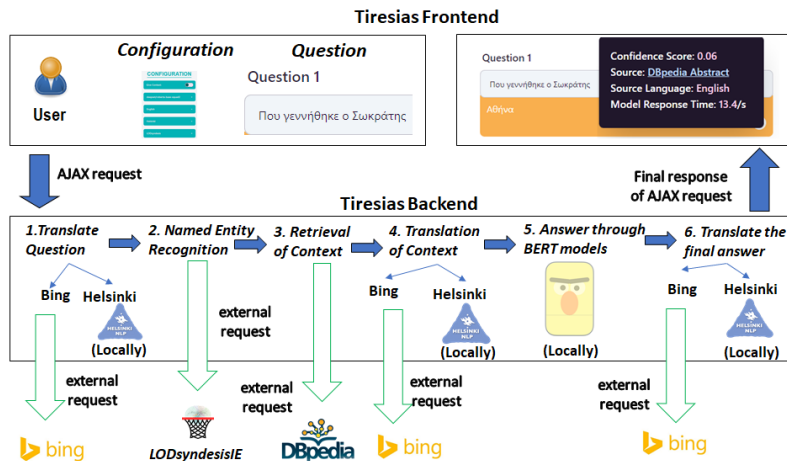


Figure 3: The architecture of *Tiresias* Web App by having as input a Greek question

since some translations (i.e., of the question and answer) were not required.

Additional Functionality. *Tiresias* can also receive as input the context, e.g., in Greek or English. In such a case, steps A and B are omitted, i.e., there is no need to retrieve a context from DBpedia. Moreover, it offers an error handling mechanism by showing several messages for informing the users for any possible error, and buttons for providing a fast configuration. Finally, several running examples are offered for both Greek and English language.

Architecture. Regarding its architecture, it is depicted in Fig. 3, e.g., for the case of having a Greek question as an input. As we can see after receiving the question through an AJAX request, *Tiresias* sends several requests to external APIs, such as to Bing (in case we do not use Helsinki) to LODsyndesisIE for retrieving the main entity, to DBpedia SPARQL endpoint for retrieving the abstracts and again to Bing for translating the Greek context and the answers. On the contrary, the Helsinki model and the BERT QA models have been locally downloaded, therefore we do not send external requests for these models. After completing all the steps,

the backend communicates with the frontend for sending the final response (through an AJAX request) and for presenting the final answer to the user.

Github link for Code and Tutorial Video. Concerning its code (mainly written in python and javascript) and guidelines for running *Tiresias*, they are available in <https://github.com/mbastakis/Tiresias>, and a tutorial video can be accessed through <https://youtu.be/eNiD7Dvco6M>.

Limitations and Possible Extensions. Concerning its limitations, since it sends several requests to external APIs, it can be affected when they do not work properly, e.g., when they are out-of-service. Moreover, some external APIs, such as Bing, offer a limited number of requests per day. Concerning possible extensions, since the selected MT tools and the DBpedia abstracts support hundreds of languages, more languages can be added. Moreover, *Tiresias* has been designed for enabling the easy addition of new NER, MT and BERT QA models.

4. Experimental Evaluation & Use Cases

In this section we provide an experimental evaluation for several components of *Tiresias*, by mainly focusing on the effectiveness/efficiency of the combination of Machine Translation and BERT QA models (pretrained in English), whereas we provide some statistics for DBpedia abstracts and several use cases for indicating the impact of using DBpedia abstracts.

Evaluation Collection for QA over Greek Language. We have created an evaluation collection, called *GreekTexts*. We collected manually 20 texts for QA, derived from a Greek text bank (<https://www.greek-language.gr/certification/dbs/teachers/index.html>), and we have written for each of the 20 texts, 10 questions and their corresponding golden answer, i.e., in total 200 questions and answers. The texts cover a variety of subjects and each text contains on average 398.5 words, each question 7.9 words and each answer 5.6 words. An alternative choice could be to use the collection XQuAD [22], which is a translated version of SQuAD that covers Greek language, however, the translated Greek version of its texts was not so accurate. However, we plan also to perform an evaluation with this collection in the future.

Evaluation Metrics and Problems. The translation can result to several problematic cases, e.g., words/phrases can be replaced by synonyms tenses/suffixes can be changed, especially for complex languages, like Greek, which uses a very extensive inflection. It means that the same word can be represented with many different suffixes (very common case for the Greek language), i.e., for denoting tenses, genders, singular and plural, and others. Therefore, in such cases it can be infeasible to evaluate the results even through F1score. For this reason, for providing accurate results we decided to annotate manually the predicted answers, i.e., in total 4,200 answers. We divided them in 3 categories by checking the golden answer: a) *Correct*, if the predicted answer has exactly the same meaning with the golden one, b) *Partially Correct*, if the predicted answer covers either a part or the whole golden answer, however it contains additional information that are irrelevant, and c) *Wrong*, if both the string representation and the meaning of the predicted answer is totally irrelevant compared to the golden answer.

4.1. Experimental Results over *Greek Texts* collection

Concerning the experimental results, we have used 20 combinations of MT/BERT QA models, i.e., for each of the 10 BERT/QA models of Section 3, whereas we use both Bing and Helsinki

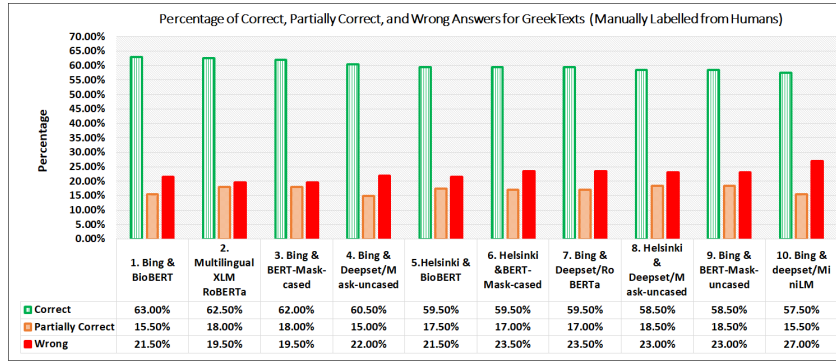


Figure 4: The top-10 models with the highest percentage of correct answers (Human Labelling)

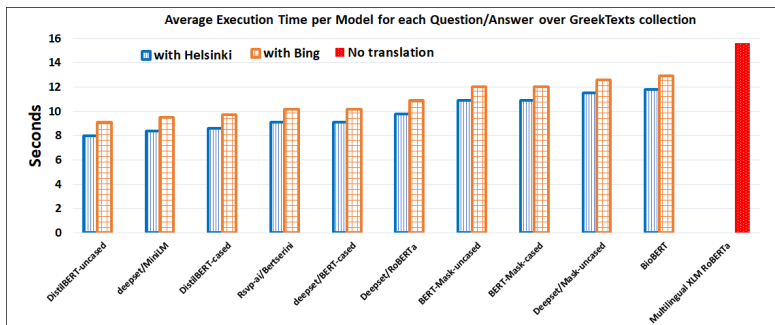


Figure 5: Efficiency of the different QA models for the GreekTexts collection

translation models. Moreover, we compared these models with a Multilingual BERT QA model, called XLM-RoBERTa, which has been pretrained in several languages including Greek.

Effectiveness. Fig. 4 shows the top-10 models according to the percentage of their correct answers. As we can see, the best combination uses the Bing translator, i.e., the combination of “Bing/BioBERT” answered correctly 63% of answers, partially correctly 15.5% and wrongly 21.5%. For this dataset, it outperformed the multilingual BERT QA model (which has been also pretrained in Greek), although for the latter we did not perform any translation. Moreover, some other models using the Bing translator, such as “BERT-Mask-cased” and “Deepset/Mask-uncased” were quite effective. Generally, we obtained better results by using Bing, however, Helsinki translator was quite effective in many cases.

Efficiency. Here, we present results about the efficiency of performing MT and QA with BERT, by using the *GreekTexts* collection. Fig. 5 shows the efficiency for each of the BERT QA model for both Helsinki and Bing MT models, starting from the fastest one. We needed from 8 to 13 seconds, whereas by using the Helsinki translator (we do not send an external request), we achieved lower execution times for each BERT model. In comparison to the multilingual “XLM RoBERTa” model, we achieved lower execution times for each of the combination of MT/BERT QA English models, mainly since the multilingual model has been pretrained in a huge corpus (with many languages). Indeed, even in the worst case the combination of MT/BERT QA model was 2.7 seconds faster. Concerning the execution time of the translation steps (they are included in the total execution time of Fig. 5), we needed on average 1.6 seconds for the translation of the context, question, and answer by using Bing, and only 0.5 seconds by using Helsinki.

Category	Greek Abstract	English Abstract	Combined (Greek + English)
Median Words	73.0	84.0	189.0
Average Words	115.4	121.8	237.2
Max Words	1676.0	1078.0	1995.0
% of Abstracts with words \leq 50	36.9%	29.5%	5.7%
% of Abstracts with 50<words \leq 100	24.1%	27.9%	18.2%
% of Abstracts with 100<words \leq 150	12.7%	15.0%	16.2%
% of Abstracts with 150<words \leq 200	5.9%	9.9%	12.3%
% of Abstracts with words>200	20.4%	17.7%	47.5%

Table 1

Statistics for DBpedia abstracts for 7,500 entities having both an English and a Greek abstract

4.2. Evaluation of the Other Components of Tiresias and More Results.

Concerning the effectiveness of the other components, experimental results about LODsyndesisIE are presented in [20], whereas in the webpage of each BERT QA model, one can find more results about their effectiveness through F1score and Exact Match metrics. As regards their efficiency, the NER process of LODsyndesisIE and the retrieval of DBpedia abstracts, they are quite fast, i.e., usually less than 1 second is needed for retrieving this information. Finally, more experimental results and the evaluation texts are available in <https://github.com/mbastakis/Tiresias>.

4.3. Statistics of DBpedia Abstracts in Greek and English

Since the context is derived from DBpedia abstracts, we present some relevant statistics. First, there are 61,159 DBpedia abstracts offered both in Greek and English. Table 1 shows some statistics by randomly using 7,500 entities having abstracts in both languages. On average, for these entities the English abstracts contain more words than the Greek ones, i.e., 121.8 versus 115.4. However, the Greek abstracts contain more words in 44% of cases, and the English abstracts in 56% of cases. By combining both abstracts, the average number of words is 237.2, i.e., smaller than the average number of words in the texts of *GreekTexts* collection, whereas only a small percentage of entities have a long abstract (over 200 words), e.g., even in the combined case, less than a half of entities. Thereby, we expect that for entities having abstracts with fewer words than the used collection, we can achieve lower execution times (i.e., than those of Fig. 5),

4.4. Use Cases of Tiresias based on Real Examples

Here, we provide some use cases of *Tiresias*, with real examples that are presented in Table 2.

•**UC1. Bilingual QA DBpedia Entities.** The proposed pipeline can be exploited for multilingual QA over any DBpedia entity, by automatically retrieving its DBpedia abstracts as a context. Concerning the current status of *Tiresias*, it can be used for answering plain open domain factoid questions for both English and Greek language, by exploiting the corresponding language versions of DBpedia abstracts. Therefore, even questions written in English whose answer is included only inside a Greek abstract (and vice versa), are feasible to be answered through *Tiresias*. Such real examples, where the answer is included only inside either the Greek or the English abstract, are presented in the examples with IDs 1-2 in Table 2. Moreover, the third example includes a question that can be answered from both English and Greek abstract.

ID	UC	Question (Greek version)	<i>Tiresias</i> Answer (Greek version)	Explanation
1	UC1 UC2	When was the last eruption of Santorini Volcano? (Πότε έγινε η τελευταία έκρηξη στο ηφαιστειο στη Σαντορίνη)	Answer: 1950 • By using Bing/RoBERTa	This answer exists only in Greek abstract
2	UC1 UC2	How many kittens a female Cat usually has? (Πόσα μικρά έχει μια Γάτα συνήθως)	Answer: two to five (δύο έως πέντε) • By using Helsinki/BERT-Mask-cased	This answer exists only in English abstract
3	UC1 UC2	How many people visited Heraklion in 2018? (Πόσοι άνθρωποι επισκέφτηκαν το Ηράκλειο το 2018)	Answer: 3.4 Million (3.4 εκατομμύρια) • By using Helsinki/BERT-Mask-cased	This answer exists in both Greek and English abstract
4	UC3	Where was 1976 Olympic Games located? (Που έγιναν οι Ολυμπιακοί Αγώνες του 1976)	Answer 1: Toronto (Τορόντο) • By using Helsinki/DistilBERT Answer 2: Montreal (Μόντρεαλ) • By using Helsinki/RoBERTa	Different Answers from different models (a wrong and a correct one)
5	UC3	Who was the scorer in Euro 2004 final? (Ποιος έβαλε το γκολ στον τελικό του Euro 2004)	Answer 1: Angel Christas • By using Helsinki/RoBERTa Answer 2: Angelos Charisteas (Άγγελος Χαριστέας) • By using Bing/RoBERTa	A bad and a correct translation of the Named Entity of the answer

Table 2

Real examples from *Tiresias* for the Use Cases, showing the impact of DBpedia Abstracts

- **UC2. Multilingual Hybrid QA systems.** The process of *Tiresias* can be used from multilingual hybrid QA systems that also use templates of SPARQL queries, i.e., by combining both structured data and DBpedia abstracts. Indeed, the answer of examples 1-3 exist only inside the abstracts and not in separate triples. On the contrary, the triples of an entity usually cover complementary information, e.g., information about “the major” and “people that were born in” Santorini island exists in separate triples but not in the English and Greek DBpedia abstracts. Therefore, it would be interesting to combine all this information for hybrid QA.

- **UC3. Comparison of MT and QA models.** First, the process of *Tiresias* can be even used as a baseline for Multilingual QA, e.g., for comparing its performance against single-language or multilingual pretrained BERT models. Moreover, through *Tiresias*, one can compare the performance of different combinations of MT and QA models, e.g., for evaluating the QA process over Greek or English texts. For instance, in the example 4 of Table 2, we can see a real case where two models provided a different answer for the same question, a wrong (Toronto) and the correct one (Montreal). Moreover, in the example 5, the two translation models provided a different translation for the entity of the answer, i.e., “Angelos Charisteas”. In particular, Bing provided the correct answer and the Helsinki model failed to correctly translate the final answer.

5. Conclusion

We presented *Tiresias*, a research prototype that combines popular NER, MT and BERT QA models (pretrained in English), for offering bilingual QA in Greek and English over DBpedia abstracts. We described all its steps and its architecture, and we provided measurements about its effectiveness and efficiency for a Greek evaluation collection with 20 texts and 200 questions and we presented use cases based on real examples. The human labelled results showed that the combination of MT/BERT QA models can be exploited for QA in Greek, e.g., the best

combination predicted correctly 63% of the answers. As a future work, we plan to a) evaluate *Tiresias* with specialized multilingual collections (e.g., by also constructing a new multilingual collection based on DBpedia abstracts), b) offer a REST API, c) extend *Tiresias* by supporting more languages and complex questions with many entities, d) provide a hybrid QA model, and e) use textual resources from other KBs, e.g. LODsyndesis [23] and Wikidata.

References

- [1] J. Lehmann, et al., DBpedia—a large-scale, multilingual knowledge base extracted from wikipedia, *Semantic web* 6 (2015) 167–195.
- [2] M. Brümmer, M. Dojchinovski, S. Hellmann, DBpedia abstracts: a large-scale, open, multilingual nlp training corpus, in: *LREC*, 2016, pp. 3339–3343.
- [3] P. N. Mendes, M. Jakob, A. García-Silva, C. Bizer, DBpedia spotlight: shedding light on the web of documents, in: *Proceedings of conference on semantic systems*, 2011, pp. 1–8.
- [4] K. Höffner, S. Walter, E. Marx, R. Usbeck, J. Lehmann, A.-C. Ngonga Ngomo, Survey on challenges of question answering in the semantic web, *Semantic Web* 8 (2017) 895–920.
- [5] A. Perevalov, et al., Can machine translation be a reasonable alternative for multilingual question answering systems over knowledge graphs?, in: *Proceedings of the ACM WebConf*, 2022, pp. 977–986.
- [6] Y. Zhou, et al., Improving zero-shot cross-lingual transfer for multilingual question answering over knowledge graph, in: *NAACL*, 2021, pp. 5822–5834.
- [7] T. Tanon, M. D. de Assunção, E. Caron, F. Suchanek, Platypus—A Multilingual Question Answering Platform for Wikidata, Ph.D. thesis, LIP-ENS Lyon, 2018.
- [8] D. Vrandečić, M. Krötzsch, Wikidata: a free collaborative knowledge base, *Communications of the ACM* 57 (2014) 78–85.
- [9] D. Diefenbach, et al., QAnswer: A question answering prototype bridging the gap between a considerable part of the lod cloud and end-users, in: *The World Wide Web Conference*, 2019, pp. 3507–3510.
- [10] M. Burtsev, et al., Deepavlov: Open-source library for dialogue systems, in: *Proceedings of ACL 2018, System Demonstrations*, 2018, pp. 122–127.
- [11] C. Nikas, P. Fafalios, Y. Tzitzikas, Open domain question answering over knowledge graphs using keyword search, answer type prediction, SPARQL and pre-trained neural models, in: *ISWC*, Springer, 2021, pp. 235–251.
- [12] E. Marakakis, H. Kondylakis, P. Aris, APANTISIS: A greek question-answering system for knowledge-base exploration, in: *Strategic Innovative Marketing*, Springer, 2017, pp. 501–510.
- [13] A. Perevalov, et al., Qald-9-plus: A multilingual dataset for question answering over DBpedia and Wikidata translated by native speakers, in: *2022 IEEE 16th ICSC*, IEEE, 2022, pp. 229–234.
- [14] R. Usbeck, et al., Qald-7, in: *Semantic web evaluation challenge*, Springer, 2017, pp. 59–69.
- [15] E. Dimitrakis, K. Sgontzos, Y. Tzitzikas, A survey on question answering systems over linked data and documents, *Journal of intelligent information systems* 55 (2020) 233–259.
- [16] A. Asai, A. Eriguchi, K. Hashimoto, Y. Tsuruoka, Multilingual extractive reading comprehension by runtime machine translation, *arXiv preprint arXiv:1809.03275* (2018).
- [17] M. Krubiński, E. Ghadery, M. F. Moens, P. Pecina, Just ask! evaluating machine translation by asking and answering questions, in: *Proceedings of the Sixth Conference on Machine Translation*, 2021, pp. 495–506.
- [18] J. Tiedemann, S. Thottingal, Opus-mt—building open translation services for the world, in: *Proceedings of the 22nd Annual Conference of the European Association for Machine Translation*, 2020.
- [19] F. Piccinno, P. Ferragina, From TagME to WAT: a new entity annotator, in: *Proceedings of the first international workshop on Entity recognition & disambiguation*, 2014, pp. 55–62.
- [20] M. Mountantonakis, Y. Tzitzikas, Linking entities from text to hundreds of RDF datasets for enabling large scale entity enrichment, *Knowledge* 2 (2021) 1–25.
- [21] C. Zeng, et al., A survey on machine reading comprehension—tasks, evaluation metrics and benchmark datasets, *Applied Sciences* 10 (2020) 7640.
- [22] M. Artetxe, S. Ruder, D. Yogatama, On the cross-lingual transferability of monolingual representations, *CoRR abs/1910.11856* (2019). [arXiv:1910.11856](https://arxiv.org/abs/1910.11856).
- [23] M. Mountantonakis, Y. Tzitzikas, LODsyndesis: global scale knowledge services, *Heritage* 1 (2018) 23.