Opponent's Review of Diploma Thesis

Title of the thesis:
Unsupervised Open Information Extraction with Large Language Models

Author:
Amrita Harikrishnan Nair

Task specification

The submitted thesis focuses on Open Information Extraction (OpenIE), especially on systems that work in an unsupervised manner using Large Language Models (LLMs). The author studies in detail one existing system, namely *DeepEx*, which automatically extracts entities and their relationships from large textual corpora in the form of structured triples consisting of a pair of noun phrases connected with a predicate. *DeepEx* employs the LLM technology to generate candidate triples, and then uses a "contrastive ranking model" to identify triples that explicitly convey the relational information between the noun phrases.

In this thesis the student proposes an improvement of the OpenIE output generated by *DeepEx*, and experiments with another, newly proposed and simpler ranking model, which is based on linguistic acceptability and trained on relatively small amounts of sample data. She evaluates the proposed system thoroughly on several evaluation benchmarks and in several different settings. The main goal of the thesis was to investigate the possible conjunction of the OpenIE performance and the linguistic acceptability of the generated triples, which have been generally considered as two disparate fields so far. She also analyses the specific contributions of linguistic acceptability measures to the overall performance in the OpenIE task.

Thesis structure

The text of the thesis is well stuctured in a standard way. Three main parts of the text can be recognised. The first introductory part gives a brief introduction to the field, and some necessary background (in Introduction, and in Chapters 1 and 2, about 15 pages). Motivation of the work is described first, then a brief info on the LLM technology and the notion of linguistic acceptability is given, and finally a brief explanation of *DeepEx* system. Given this background, the author summarises the objectives of her thesis at the very end of Chapter 2.

The second main part is given in Chapter 3 (about 17 pages). Here the author describes the used datasets, and details on linguistic acceptability computing, and then she explains how the proposed system is implemented, and the design of her experiments.

Then the third main part of the thesis gives experimental results, its analysis, error analysis, discussion, and conclusions. Here she thoroughly describes the experiments and the results, and discusses the differences in the different models' outputs. Her main experiments to evaluate the new ranking method are described in Chapter 4.1, and are conducted for five different benchmarks. Moreover, in Chapter 4.2 she conducts an ablation study to identify the

components of the new system that impact the OpenIE performance, and also, another short study on the connection between OpenIE and semantic role labelling. She also gives some ideas about limitations of her approach, and on future work. This is covered in Chapter 4 and in short Conclusion (about 15 pages in total).

Formal aspects and evaluation

The submitted thesis fulfills all general formal requirements. It is written in good English and all the text is well comprehensible. Minor grammatical mistakes and typos are not too frequent. The list of the used bibliography consists of around 70 items. The lists of Tables and Figures are added at the end.

In the thesis, the author convincingly demonstrated good knowledge of the field and practical ability to perform a large number of non-trivial experiments. The resulting data are presented in numerous tables with rich commentaries.

To conclude my review, I appreciate that the student has successfully contributed to the OpenIE research area by performing a systematic and quite extensive series of experiments with ranking OpenIE triples. In my view she did a lot of work, and the work has been both done and presented thoroughly. Her results seem to be "promising across benchmarks and datasets, despite using a very small amount of training data" (see her Conclusion).

I definitely recommend the submitted thesis for the defense and I suggest accepting this work as a diploma thesis.

September 3, 2024

RNDr. Martin Holub, Ph.D.
Institute of Formal and Applied Linguistics
Charles University in Prague