Review of a Doctoral Thesis at FIT BUT

Doctoral thesis (title of the thesis):

AUTOMATED FACTOID QUESTION ANSWERING AND FACT-CHECKING IN NATURAL LANGUAGE

Name of the doctoral student (name and surname):

MARTIN FAJČÍK

Name and institution of the reviewer (full name of the reviewer, full name and country of the institution):

James Thorne, Korea Advanced Instituted of Science and Technology (KAIST), South Korea

---

Please state your opinion on the following aspects of (I) the candidate's doctoral thesis and (II) the candidate's overall achievements, and (III) state your conclusion (a minimum of approx. 300 characters for each item? <u>point</u> below is recommended):

I. Doctoral Thesis

<u>Appropriateness and Relevance</u>

The area addressed by the thesis is appropriate to the field of Natural Language Processing (NLP). The thesis addresses relevant and contemporary problems in this field, specifically focusing on three core contributions: extractive question answering, fact verification, and information retrieval (supporting both previous contributions). The problems addressed by this thesis are relevant and pertinent problems. Martin's contributions not only support development of models with higher accuracy, but also support improved runtime efficiency and explainability of NLP systems. The relevance of this thesis is validated through a collection of publications accepted at well-reputed peer-reviewed conference and workshop venues.

<u>A summary of the Contributions of the Thesis</u>

The thesis aims to develop training and inference methodologies of common NLP models to enhance the performance, efficiency, and interpretability of neural NLP systems. The thesis achieves its goal. The main contributions are listed as follows:

1) Alternative probabilistic formulation for extractive question answering:

    The thesis introduces a joint inference approach, considering the relationship between span start and span end classification tasks, in extractive question answering. This formulation aids improving the accuracy and modelling assumptions of answer extraction in QA.

2) Hybrid Architecture for Open-Domain Question Answering:

    Another contribution of this thesis is the development of a hybrid architecture that integrates abstractive and extractive question answering methods for open-domain question answering tasks. This architecture has become a well-cited standard baseline in the field, facilitating model

comparisons and benchmarking in open-domain QA. The approach uses a fusion mechanism to aggregate the scores from abstractive and extractive readers.

3) Study on Indexes in Information Retrieval:

The thesis includes an interesting take on the use of large indexes used in inner-product-based information retrieval. It demonstrates the feasibility of pruning these indexes to reduce their memory footprint, thereby enhancing the efficiency of information retrieval processes. This also demonstrates an information partition in current open-domain QA formulations. This could inspire future work on long tail datasets and building in additional robustness, or in performance optimization of models when applied to real-world settings.

4) Interpretable Automated Fact Verification:

A further study in the thesis uses the attention mechanism of a fact verification classifier system to add interpretability. This approach adds a mechanism to transfer sentence or passage level supervision signals to induce token level highlights.

Novelty and Significance:

The PhD thesis exhibits a suitable level of novelty and significance.

(1) Introducing the new extractive QA objective that explicitly models the interdependency between span start and end tasks presents a modest extension of existing approaches implicitly modelling this through neural architecture.

(2) The Hybrid QA Architecture combining abstractive and extractive methods is an interesting approach that bridges two predominant approaches in the field. The method introduced in this thesis has become a standardised baseline for the task. Furthermore, this approach has a higher potential for real-world impact as the fusion of the methods in retrieval-based question answering may yield more robust systems that can be applied to everyday settings leveraging the advantages of both methods.

(3) Interpretable Fact Verification: The Claim Dissector system achieved high accuracy for claim veracity prediction, compared to other leader board systems. This approach propagates the attention scores as token-level probabilities. Employing attention mechanisms for interpretability in fact verification is an innovative step towards yielding explanations from models. However, I would expect the student to add additional discussion to existing work on the use of attention in explanation (e.g. Jain and Wallace 2019).

However, for Chapter 6 (Ensembling Models for Stance Classification), an aspect that is overlooked in the thesis is the efficiency (runtime and parameter count) of creating such large ensembles of models.

Evaluation of the Formal Aspects of the Thesis:

The thesis has a suitable narrative arc: building on work in QA, to open-domain QA, to indexes, to applications to fact verification. The referencing and citation style is appropriate for this discipline. The writing style of the thesis is acceptable. However, there are small number of instances where the student uses a less formal tone may need adjustment to maintain the formal and scholarly tone expected of a doctoral thesis.

Quality of Publications

The student reports five publications in appropriate NLP conference and workshop venues. The student has two papers under findings of EMNLP and ACL which are international flagship conference venues for natural language processing. The workshop papers are similarly collocated with appropriate conferences. This level of publication is suitable for work in natural language processing.

II. Candidate's Overall Achievements

Overall R&D Activities Evaluation:

The student's thesis presents an appropriate level of breadth and depth of knowledge on the topics of core topics relating to open-domain natural language processing. The contributions in the document exhibit a suitable level of creativity and originality, building on existing works in the field to introduce new findings, build more accurate, robust and explainable NLP systems. Through the thesis, and through examination of the students publications outside of the thesis, there is suitable demonstration of scholarly contribution.

Assessment of Other Candidate Characteristics (optional):

N/A

III. Conclusion

The conclusion should contain an explicit statement saying whether, in your opinion, the doctoral thesis and the student´s achievements until now meet the generally accepted requirements for the award of an academic degree (in accordance with Section 47 of Act No. 111/1998 Coll., on higher education institution).*

* Short overview of both the Act and corresponding internal BUT regulations is enclosed.

Based on the examination of the thesis document, the student's CV, and publication history, it is the opinion of the examiner that the thesis contribution and the student's achievements meet generally accepted requirements for the award of the degree of PhD and should proceed for examination through public defense. Please refer to the attached annotated version of the thesis PDF file for suggested minor amendments to aid clarity.

Remote evaluation (Seoul, South Korea) 04.01.2024

Signature of the reviewer