

Textual Entailment Graphs

Lili Kotlerman, Ido Dagan, Bernardo Magnini, Luisa Bentivogli

(Submitted to JNLE, under review)

Since the Textual Entailment paradigm of modelling semantic inference was first introduced by [1], it has become a notable concept in the field. Yet, although entailment has been successfully utilized in various NLP systems, until recently it had been applied only in a pairwise manner, to recognize the entailment relation between single pairs of elements. Recently, researchers started utilizing entailment to construct entailment graphs, where nodes represent language expressions and directed edges represent entailment between nodes. Berant et al. ([2]) proposed a global algorithm over entailment graphs with predicates at nodes, in order to improve acquisition of entailment rules between predicates, such as '*X marry Y*' \rightarrow '*Y is X's spouse*'. Mehdad et al. ([3]) built an entailment graph of token n-grams for topic labelling. Levy et al. ([4]) suggested an approach for organizing and consolidating open IE propositions, such as '*cure(aspirin, headache)*', using the notion of proposition entailment graphs.

In this work we suggest a novel type of graphs, namely Textual Entailment Graphs (TEG), where nodes represent complete natural language texts rather than single concepts (like in concept hierarchies), n-grams or reductive structures such as predicates and open-IE propositions. Given a set of texts (graph nodes), the task of constructing a textual entailment graph is to recognize all the entailments among the texts, i.e. deciding which directional edges connect which pairs of nodes. The main difference between this task and the Recognizing Textual Entailment (RTE) task is that the text pairs are not independent. The nodes in the graph are inter-connected via entailment edges, and since entailment is a transitive relation, the edges should not violate transitivity.

We formally define Textual Entailment Graphs and describe a complete methodology we developed for their construction. We introduce the first dataset we created for the task of TEG construction, which contains 29 consistent manually-annotated textual entailment graphs. We use this dataset to evaluate the performance of several state-of-the-art RTE engines for the task of automatic TEG construction and thus provide a number of baselines for this task. We show that organizing the texts in a TEG and enforcing transitivity results in improved pairwise entailment decisions.

Since our research was motivated by joint work with industrial partners in the text analytics area, we performed a user study which showed the usefulness of TEGs for text exploration. However, while our own motivation and dataset focus at text exploration setting, we suggest that textual entailment graphs might have different usages and suggest that automatic creation of such graphs is an interesting task for the community.

References:

- [1] Dagan, Ido and Oren Glickman. 2004. Probabilistic textual entailment: Generic applied modelling of language variability.
- [2] Berant, Jonathan, Ido Dagan, and Jacob Goldberger. 2011. Global learning of typed entailment rules. In Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies-Volume 1, pages 610–619. Association for Computational Linguistics.
- [3] Mehdad, Yashar, Giuseppe Carenini, Raymond T Ng, and Shafiq Joty. 2013. Towards topic labelling with phrase entailment and aggregation. In Proceedings of NAACL-HLT, pages 179–189.
- [4] Levy, Omer, Ido Dagan, and Jacob Goldberger. 2014. Focused Entailment Graphs for Open IE Propositions. In Proceedings of CoNLL.