Distant Supervision for Temporal Resolution

Jeniya Tabassum
Computer Science and Engineering
Ohio State University
bintejafar.1@osu.edu

Alan Ritter
Computer Science and Engineering
Ohio State University
ritter.1492@osu.edu

Temporal expression resolution in newswire text is a long-studied problem; traditional approaches have relied mostly on hand-written rules (Mani and Wilson, 2000; Chang and Manning, 2012). Recently, there has been increasing interest in learning to resolve time-expressions directly from data (Angeli et al., 2012; Lee et al., 2014), however previous methods are limited by their reliance on small hand-annotated datasets. In this work we investigate the feasibility of learning robust date resolvers by scaling up to learning from massive datasets through distant supervision.

Most previous work on distant supervision has focused on the task of relation extraction (Mintz et al., 2009). As far as we are aware, no previous work has applied distant supervision to resolving temporal information contained in text. Existing models of temporal resolution make use of hand-crafted rules and supervised learning.

One motivation for the proposed approach is to make temporal resolution domain independent. Since distant supervision does not rely on manually labeled data in a fixed domain, it is able to scale up to large unlabeled corpora. Our proposed approach has the potential to learn abbreviations and other phenomenons that are common in social media but not present in well-edited newswire text.

The assumption behind the proposed distant supervision model is that tweets mentioning an event will often contain temporal expressions referring to the event's date. Given a set of events represented as tuples (e, d) consisting of a key entity, e, and calendar date d, we identify tweets written near the time of the event which contain mentions of the entity. Each tweet is then associated with a set of temporal labels derived from the event's date and also the creation date of the tweet. Orthographic and contextual features (Mintz et al., 2009) of selected tweets are extracted which we use as the input to the temporal classifier. The classifier outputs the temporal tags of the tweets using a conditionally trained latent variable model of distant supervision. We adopt the MultiR (Hoffmann et al., 2011) and DNMAR models (Ritter et al., 2013) in the context of extracting temporal relations. Figure-1 presents the graphical model structure of our Temporal Tagger. The extracted temporal tags are used as an input to a separate date resolution classifier along with the tweet and its creation date, which finally predicts the calendar date of the event from the tweet.

To evaluate the performance of our proposed approach, we have conducted experiments over an annotated twitter corpus consisting of tweets from years 2011 and 2012 and found that the DNMAR model can extract the temporal tags from the tweets with an F-score of 58.3% where the MultiR model extracts tags with a F-score of 48.6%. Using the tags form the DNMAR model and creation dates of the tweet, the date classifier can extract the event date with an F-score of 79.9%.

In future work, we would like to apply the proposed system over a human-annotated corpus and evaluate its performance.

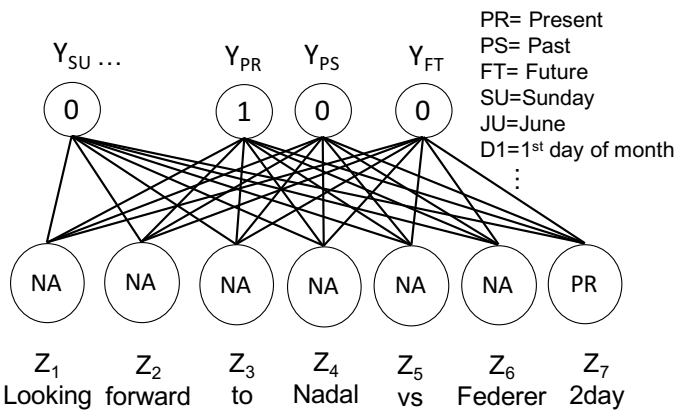

Figure 1: Our approach is capable of learning word-level temporal tags given only sentence-level labels.

References

Gabor Angeli, Christopher D. Manning, and Daniel Jurafsky. 2012. Parsing time: Learning to interpret time expressions. In *NAACL-HLT*.

Angel X Chang and Christopher D Manning. 2012. Sutime: A library for recognizing and normalizing time expressions. In *LREC*.

Raphael Hoffmann, Congle Zhang, Xiao Ling, Luke Zettlemoyer, and Daniel S. Weld. 2011. Knowledge-based weak supervision for information extraction of overlapping relations.

Kenton Lee, Yoav Artzi, Jesse Dodge, and Luke Zettlemoyer. 2014. Context-dependent semantic parsing for time expressions. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*.

Inderjeet Mani and George Wilson. 2000. Robust temporal processing of news. In *Proceedings of the 38th Annual Meeting on Association for Computational Linguistics*.

Mike Mintz, Steven Bills, Rion Snow, and Dan Jurafsky. 2009. Distant supervision for relation extraction without labeled data.

Alan Ritter, Luke Zettlemoyer, and Oren Etzioni. 2013. Modeling missing data in distant supervision for information extraction.