

Multi-Task Learning of Keyphrase Boundary Classification

Isabelle Augenstein *

Department of Computer Science
University College London
i.augenstein@ucl.ac.uk

Anders Søgaard *

Department of Computer Science
University of Copenhagen
soegaard@di.ku.dk

Abstract

Keyphrase boundary classification (KBC) is the task of detecting keyphrases in scientific articles and labelling them with respect to predefined types. Although important in practice, this task is so far underexplored, partly due to the lack of labelled data. To overcome this, we explore several auxiliary tasks, including semantic super-sense tagging and identification of multi-word expressions, and cast the task as a multi-task learning problem with deep recurrent neural networks. Our multi-task models perform significantly better than previous state of the art approaches on two scientific KBC datasets, particularly for long keyphrases.

1 Introduction

The scientific keyphrase boundary classification (KBC) task consists of a) determining keyphrase boundaries, and b) labelling keyphrases with their types according to a predefined schema. KBC is motivated by the need to efficiently search scientific literature, which can be summarised by their keyphrases. Several companies are working on keyphrase-based recommender systems for scientific literature or search interfaces where scientific articles decorate graphs, in which nodes are keyphrases. Such keyphrases must be dynamically retrieved from the articles, because important scientific concepts emerge on a daily basis, and the most recent concepts are typically the ones of interest to scientists.

KBC is not a common task in NLP, and there are only few small annotated datasets for inducing supervised KBC models, made available recently

(QasemiZadeh and Schumann, 2016; Augenstein et al., 2017). Typical KBC approaches therefore rely on hand-crafted gazetteers (Hasan and Ng, 2014) or reduce the task to extracting a list of keyphrases for each document (Kim et al., 2010) instead of identifying mentions of keyphrases in sentences. For related more common NLP tasks such as named entity recognition and identification of multi-word expressions, neural sequence labelling methods have been shown to be useful (Lample et al., 2016). In order to overcome the small data problem, we study using more widely available data for tasks related to KBC and exploit their synergies in a deep multi-task learning setup.

Multi-task learning has become popular within natural language processing and machine learning over the last few years; in particular, *hard* parameter sharing of hidden layers in deep learning models. This approach to multi-task learning has three advantages: a) It significantly reduces Rademacher complexity (Baxter, 2000; Maurer, 2007), i.e., the risk of over-fitting, b) it is space-efficient, reducing the number of parameters, and c) it is easy to implement.

This paper shows how hard parameter sharing can be used to improve gazetteer-free keyphrase boundary classification models, by exploiting different syntactically and semantically annotated corpora, as well as more readily available data such as hyperlinks.

Contributions We study the so far widely underexplored, though in practice important task of scientific keyphrase boundary classification, for which only a small amount of training data is available. We overcome this by identifying good auxiliary tasks and cast it as a multi-task learning problem. We evaluate our models across two new, manually annotated corpora of scientific articles and outperform single-task approaches by up

*Both authors contributed equally

to 9.64% F1, mostly due to better performance for long keyphrases.

2 Keyphrase Boundary Classification

Consider the following sentence from a scientific paper:

- (1) We find that simple interpolation methods, like log-linear and linear interpolation, improve the performance but fall short of the performance of an oracle.

This sentence occurs in the ACL RD-TEC 2.0 corpus. Here, *interpolation methods* and *log-linear and linear interpolation* are annotated as technical keyphrases, *performance* as a keyphrase related to measurements, and *oracle* is a keyphrase labelled as miscellaneous. Below, we are interested in predicting the boundaries *and* the types of all keyphrases.

3 Multi-Task Learning

Multi-task learning is an approach to learning, in which generalisation is improved by taking advantage of the inductive bias in training signals of related tasks. When abundant labelled data is available for an auxiliary task, but little data for the target task, multi-task learning can act as a form of semi-supervised learning combined with a distant supervision signal. Inducing a model from only the sparse target task data may lead to overfitting to random noise in the data, but relying on auxiliary data helps the model generalise, making it easier to abstract away from noise, as well as leveraging the marginal distribution of auxiliary input data. From a representation learning perspective, auxiliary tasks can be used to induce representations that may be beneficial for the target task. Caruana (1993) also suggests that the auxiliary task can help focus attention in the induction of the target task model. Finally, multi-task learning can be cast as a regulariser as studies show reductions in Rademacher complexity in multi-task architectures over single-task architectures (Baxter, 2000; Maurer, 2007).

Here, we follow the probably most common approach to multi-task learning, known as *hard parameter sharing*. This was introduced in Caruana (1993) in the context of deep neural networks, in which hidden layers can be shared among tasks. We assume T different training set, D_1, \dots, D_T ,

where each D_t contains pairs of input-output sequences $(w_{1:n}, y_{1:n}^t)$, $w_i \in V$, $y_i^t \in L^t$. The input vocabulary V is shared across tasks, but the output vocabularies (tagset) L^t are task dependent. At each step in the training process we choose a random task t , followed by a random training instance $(w_{1:n}, y_{1:n}^t) \in D_t$. We use the tagger to predict the labels \hat{y}_i^t , suffer a loss with respect to the true labels y_i^t and update the model parameters. The parameters are trained jointly for a sentence, i.e. cross-entropy loss over each sentence is employed. Each task is associated with an independent classification function, but all tasks share the hidden layers. Note that for our experiments, we only consider one auxiliary task at a time.

4 Experiments

Experimental Setup We perform experiments for both keyphrase boundary identification (*unlabelled*), and keyphrase boundary identification and classification (*labelled*). Metrics measured are token-level precision, recall and F1, which are micro-average results across keyphrase types. Types are defined by the two datasets studied.

Auxiliary tasks We experiment with five auxiliary tasks: (1) syntactic chunking using annotations extracted from the English Penn Treebank, following Søgaard and Goldberg (2016); (2) frame target annotations from FrameNet 1.5 (corresponding to the target identification and classification tasks in Das et al. (2014)); (3) hyperlink prediction using the dataset from Spitzkovsky et al. (2010), (4) identification of multi-word expressions using the Streusle corpus (Schneider and Smith, 2015); and (5) semantic super-sense tagging using the Semcor dataset, following Johansson et al. (2014). We train our models on the main task with one auxiliary task at a time. Note that the datasets for the auxiliary tasks are not annotated with keyphrase boundary identification or classification labels.

Datasets We evaluate on the SemEval 2017 Task 10 dataset (Augenstein et al., 2017) and the the ACL RD-TEC 2.0 dataset (QasemiZadeh and Schumann, 2016). The SemEval 2017 dataset is annotated with three keyphrase types, the ACL RD-TEC dataset with seven. For the former, we test on the development portion of the dataset, as the test set is not released yet. We randomly split ACL RD-TEC into a training and test set, reserv-

	SemEval 2017 Task 10	ACL RD-TEC
Labels	Material, Process, Task	Technology and Method, Tool and Library, Language Resource, Language Resource Product, Measures and Measurements, Models, Other
Topics	Computer Science, Physics, Material Science	Natural Language Processing
Number all keyphrases	5730	2939
Proportion singleton keyphrases	31%	83%
Proportion single-word mentions	18%	23%
Proportion mentions with word length ≥ 2	82%	77%
Proportion mentions with word length ≥ 3	51%	33%
Proportion mentions with word length ≥ 5	22%	8%

Table 1: Characteristics of SemEval 2017 Task 10 and ACL-RD-TEC corpora, statistics of training sets

ing 1/3 for testing. Key dataset characteristics are summarised in Table 1. One important observation is that the SemEval 2017 dataset contains a significantly higher proportion of long keyphrases than the ACL dataset.

Models Our single- and multi-task networks are three-layer, bi-directional LSTMs (Graves and Schmidhuber, 2005) with pre-trained SENNA embeddings.¹ For the multi-task networks, we follow the training procedure outlined in Section 3. The dimensionality of the embeddings is 50, and we follow Søgaard and Goldberg (2016) in using the same dimensionality for the hidden layers. We add a dropout of 0.1 to the input and train these architectures with momentum SGD with initial learning rate of 0.001 and momentum of 0.9 for 10 epochs.

Baselines Our baselines are Finkel et al. (2005)² and Lample et al. (2016)³, in order to compare to a lexicalised and a state-of-the-art neural method. We use the implementations released by the authors and re-train models on our data.

5 Results and Analysis

Results for SemEval 2017 Task 10 corpus are presented in Table 2, and for the ACL RD-TEC corpus in Table 3. For the SemEval corpus, all five labelled multi-task learning models outperform both examples of previous work, as well as our single-task BiLSTM baseline, by some margin. For ACL RD-TEC, three of our five multi-task learning la-

belled models perform better than the single-task BiLSTM baseline.

On the SemEval corpus, the F1 error reduction of the best labelled model over the Stanford tagger is 9.64%. The lexicalised Finkel et al. (2005) model shows a surprisingly competitive performance on the ACL RD-TEC corpus, where it is only 2 points in F1 behind our best performing labelled model and on par with our best performing unlabelled model. Results with Lample et al. (2016), on the other hand, are lower than the Finkel et al. (2005) baseline. This might be due to the model having a large set of parameters to model state transitions which poses a difficulty for small training datasets.

Overall, multi-task models show bigger improvements over baselines for the SemEval corpus, and all models achieve better results on ACL RD-TEC. Statistics shown in Table 1 help to explain this. Most noticeably, the SemEval dataset contains a significantly higher proportion of long keyphrases than the ACL dataset. Interestingly, ACL RD-TEC contains a large proportion of keyphrases which only appear once in the training set (singletons), significantly fewer keyphrases and more keyphrase type, but that does not seem to impact results as much as a high proportion of long keyphrases.

All models struggle with semantically vague or broad keyphrases (e.g. ‘items’, ‘scope’, ‘key’) and long keyphrases, especially those containing clauses (e.g. ‘complete characterisation of the oxide particles’, ‘earley deduction proof procedure for definite clauses’). The multi-task models generally outperform the BiLSTM baseline for long phrases (e.g. ‘language-independent system for

¹<http://ronan.collobert.com/senna/>

²<http://nlp.stanford.edu/software/CRF-NER.shtml>

³<https://github.com/clab/stack-lstm-ner>

Method	Unlabelled			Labelled		
	Precision	Recall	F1	Precision	Recall	F1
Finkel et al. (2005)	77.89	50.27	61.10	49.90	27.97	35.85
Lample et al. (2016)	71.92	49.37	58.55	41.36	28.47	33.72
BiLSTM	81.58	57.86	67.71	45.80	32.48	38.01
BiLSTM + Chunking	82.88	52.08	63.96	55.54	34.90	42.86
BiLSTM + Framenet	77.86	56.05	65.18	54.04	38.91	45.24
BiLSTM + Hyperlinks	76.59	60.53	67.62	46.99	44.09	41.13
BiLSTM + Multi-word	74.80	70.18	72.42	46.99	44.09	45.49
BiLSTM + Super-sense	83.70	51.76	63.93	56.94	35.25	43.54

Table 2: Results for keyphrase boundary classification on the SemEval 2017 Task 10 corpus

Method	Unlabelled			Labelled		
	Precision	Recall	F1	Precision	Recall	F1
Finkel et al. (2005)	84.16	80.08	82.07	59.97	53.86	56.75
Lample et al. (2016)	65.60	86.06	74.45	31.30	41.07	35.53
BiLSTM	83.40	80.36	81.85	59.62	57.45	58.51
BiLSTM + Chunking	83.36	79.46	81.37	59.26	57.24	57.84
BiLSTM + Framenet	84.11	79.39	81.68	60.64	57.24	58.89
BiLSTM + Hyperlinks	83.94	79.12	81.46	60.18	56.73	58.40
BiLSTM + Multi-word	84.86	76.92	80.69	59.81	54.21	56.87
BiLSTM + Super-sense	84.67	78.29	81.36	61.35	56.73	58.95

Table 3: Results for keyphrase boundary classification on the ACL RD-TEC corpus

automatic discovery of text in parallel translation’, ‘honeycomb network of graphite bricks’). Being able to recognise long keyphrases correctly is part of the reason our multi-task models outperform the baselines, especially on the SemEval dataset, which contains many such long keyphrases.

6 Related Work

Multi-Task Learning Hard sharing of all hidden layers was introduced in Caruana (1993), and popularised in NLP by Collobert et al. (2011a). Several variants have been introduced, including hard sharing of selected layers (Søgaard and Goldberg, 2016) and sharing of parts (subspaces) of layers (Liu et al., 2015). Søgaard and Goldberg (2016) show that hard parameter sharing is an effective regulariser, also on heterogeneous tasks such as the ones considered here. Hard parameter sharing has been studied for several tasks, including CCG super tagging (Søgaard and Goldberg, 2016), text normalisation (Bollman and Søgaard, 2016), neural machine translation (Dong et al., 2015; Luong et al., 2016), and super-sense tagging (Martínez Alonso and Plank, 2017). Sharing of information can further be achieved by extending LSTMs with an external memory shared

across tasks (Liu et al., 2016). A further instance of multi-task learning is to optimise a supervised training objective jointly with an unsupervised training objective, as shown in Yu et al. (2016) for natural language generation and auto-encoding, and in Rei (2017) for different sequence labelling tasks and language modelling.

Boundary Classification KBC is very similar to named entity recognition (NER), though arguably harder. Deep neural networks have been applied to NER in Collobert et al. (2011b); Lample et al. (2016). Other successful methods rely on conditional random fields, thereby modelling the probability of each output label conditioned on the label at the previous time step. Lample et al. (2016), currently state-of-the-art for NER, stack CRFs on top of recurrent neural networks. We leave exploring such models in combination with multi-task learning for future work.

Keyphrase detection methods specific to the scientific domain often use keyphrase gazetteers as features or exploit citation graphs (Hasan and Ng, 2014). However, previous methods relied on corpora annotated for type-level identification, not for mention-level identification (Kim et al., 2010; Sterckx et al., 2016). While most applications

rely on extracting keyphrases (as types), this has the unfortunate consequence that previous work ignores acronyms and other short-hand forms referring to methods, metrics, etc. Further, relying on gazetteers makes overfitting likely, obtaining lower scores on out-of-gazetteer keyphrases.

7 Conclusions and Future Work

We present a new state of the art for keyphrase boundary classification, using data from related, auxiliary tasks; in particular, super-sense tagging and identification of multi-word expressions. Deep multi-task learning improves significantly on previous approaches to KBC, with error reductions of up to 9.64%, mostly due to better identification and labelling of long keyphrases.

In future work, we want to explore alternative multi-task learning regimes to hard parameter sharing and experiment with additional auxiliary tasks. The auxiliary tasks considered here are standard NLP tasks, hyperlink prediction aside. Other tasks may be more directly relevant such as predicting the layout of calls for papers for scientific conferences, or predicting hashtags in tweets by scientists, since both data sources contain scientific keyphrases.

Acknowledgments

We would like to thank Elsevier for supporting this work.

References

- Isabelle Augenstein, Mrinal Das, Sebastian Riedel, Lakshmi Vikraman, and Andrew McCallum. 2017. SemEval-2017 Task 10 : Extracting Keyphrases and Relations from Scientific Publications. In *Proceedings of SemEval, to appear*.
- Jonathan Baxter. 2000. A model of inductive bias learning. *Journal of Artificial Intelligence Research* 12:149–198.
- Marcel Bollman and Anders Søgaard. 2016. Improving historical spelling normalization with bi-directional LSTMs and multi-task learning. In *Proceedings of COLING*.
- Rich Caruana. 1993. Multitask Learning: A Knowledge-Based Source of Inductive Bias. In *Proceedings of ICML*.
- Ronan Collobert, Jason Weston, Léon Bottou, Michael Karlen, Koray Kavukcuoglu, and Pavel Kuksa. 2011a. Natural language processing (almost) from scratch. *The Journal of Machine Learning Research* 12:2493–2537.
- Ronan Collobert, Jason Weston, Léon Bottou, Michael Karlen, Koray Kavukcuoglu, and Pavel Kuksa. 2011b. Natural language processing (almost) from scratch. *The Journal of Machine Learning Research* 12:2493–2537.
- Dipanjan Das, Desai Chen, Andre Martins, Nathan Schneider, and Noah Smith. 2014. Frame-semantic parsing. *Computational linguistics* 40(1):9–56.
- Daxiang Dong, Hua Wu, Wei He, Dianhai Yu, and Haifeng Wang. 2015. Multi-Task Learning for Multiple Language Translation. In *Proceedings of ACL*.
- Jenny Finkel, Trond Grenager, and Christopher Manning. 2005. Incorporating non-local information into information extraction systems by Gibbs sampling. In *Proceedings of ACL*.
- Alex Graves and Jürgen Schmidhuber. 2005. Frame-wise Phoneme Classification with Bidirectional LSTM and other Neural Network Architectures. *Neural Networks* 18(5):602–610.
- Kazi Saidul Hasan and Vincent Ng. 2014. Automatic Keyphrase Extraction: A Survey of the State of the Art. In *Proceedings of ACL*.
- Anders Johannsen, Dirk Hovy, Héctor Martínez, Barbara Plank, and Anders Søgaard. 2014. More or less supervised supersense tagging of Twitter. In *Proceedings of *SEM*.
- Su Nam Kim, Olena Medelyan, Min-Yen Kan, and Timothy Baldwin. 2010. SemEval-2010 Task 5 : Automatic Keyphrase Extraction from Scientific Articles. In *Proceedings of SemEval*.
- Guillaume Lample, Miguel Ballesteros, Sandeep Subramanian, Kazuya Kawakami, and Chris Dyer. 2016. Neural Architectures for Named Entity Recognition. In *Proceedings of NAACL-HLT*. pages 260–270.
- Pengfei Liu, Shafiq Joty, and Helen Meng. 2015. Fine-grained Opinion Mining with Recurrent Neural Networks and Word Embeddings. In *Proceedings of EMNLP*.
- Pengfei Liu, Xipeng Qiu, and Xuanjing Huang. 2016. Deep Multi-Task Learning with Shared Memory for Text Classification. In *Proceedings of EMNLP*.
- Minh-Thang Luong, Quoc V. Le, Ilya Sutskever, Oriol Vinyals, and Lukasz Kaiser. 2016. Multi-task Sequence to Sequence Learning. In *Proceedings of ICLR*.
- Héctor Martínez Alonso and Barbara Plank. 2017. When is multitask learning effective? Semantic sequence prediction under varying data conditions. In *Proceedings of EACL*.
- Andreas Maurer. 2007. Bounds for Linear Multi Task Learning. *Journal of Machine Learning Research* 7:117–139.

- Behrang QasemiZadeh and Anne-Kathrin Schumann. 2016. The ACL RD-TEC 2.0: A Language Resource for Evaluating Term Extraction and Entity Recognition Methods. In *Proceedings of LREC*.
- Marek Rei. 2017. Semi-supervised Multitask Learning for Sequence Labeling. In *Proceedings of ACL, to appear*.
- Nathan Schneider and Noah Smith. 2015. A Corpus and Model Integrating Multiword Expressions and Supersenses. *Proceedings of NAACL-HLT*.
- Anders Søgaard and Yoav Goldberg. 2016. Deep multi-task learning with low level tasks supervised at lower layers. In *Proceedings of ACL*.
- Valentin Spitkovsky, Daniel Jurafsky, and Hiyan Alshawi. 2010. Profiting from Mark-Up: Hyper-Text Annotations for Guided Parsing. In *Proceedings of ACL*.
- Lucas Sterckx, Cornelia Caragea, Thomas Demeester, and Chris Develder. 2016. Supervised Keyphrase Extraction as Positive Unlabeled Learning. In *Proceedings of EMNLP*.
- Lei Yu, Jan Buys, and Phil Blunsom. 2016. Online Segment to Segment Neural Transduction. In *Proceedings of EMNLP*.