

Towards Open-Domain Topic Classification

Hantian Ding¹, Jinrui Yang^{1,2}, Yuqian Deng¹, Hongming Zhang¹, Dan Roth¹

¹University of Pennsylvania, ²University of Melbourne

{hantian2, jinruiy, yuqiand, hzhangal, danroth}@seas.upenn.edu

Abstract

We introduce an open-domain topic classification system that accepts user-defined taxonomy in real time. Users will be able to classify a text snippet with respect to any candidate labels they want, and get instant response from our web interface. To obtain such flexibility, we build the backend model in a zero-shot way. By training on a new dataset constructed from Wikipedia, our label-aware text classifier can effectively utilize implicit knowledge in the pretrained language model to handle labels it has never seen before. We evaluate our model across four datasets from various domains with different label sets. Experiments show that the model significantly improves over existing zero-shot baselines in open-domain scenarios, and performs competitively with weakly-supervised models trained on in-domain data.¹²

1 Introduction

Text classification is a fundamental natural language processing problem, with one of its major applications in topic labeling (Lang, 1995; Wang and Manning, 2012). Over the past decades, supervised classification models have achieved great success in closed-domain tasks with large-scale annotated datasets (Zhang et al., 2015; Tang et al., 2015; Yang et al., 2016). However, they are no longer effective in open-domain scenarios where the taxonomy is unbounded. Retraining the model for every new label set often incurs prohibitively high cost in the sense of both annotation and computation. By contrast, having one classifier that is flexible with unlimited labels can save such tremendous efforts while keeping the solution simple. Therefore, in this work, we build a system for open-domain topic classification that can classify a given text snippet into any categories defined by users.

¹Interactive online demo at https://cogcomp.seas.upenn.edu/page/demo_view/ZeroShotTC

²Code and data available at http://cogcomp.org/page/publication_view/980

At the core of our system is a zero-shot text classification model. While supervised models are typically insensitive to class names, a zero-shot model is usually label-aware, meaning that it can understand label semantics directly from the name or definition of the label, without accessing any annotated examples. Our model **TE-Wiki** combines a **Textual Entailment (TE)** formulation with Wikipedia finetuning. Specifically, we construct a new dataset that contains three million article-category pairs from Wikipedia’s subcategory graph, and finetune a pretrained language model (e.g. BERT) to predict the entailment relations between articles and their associated categories. We simulate the diversity in open-domain classification with the wide coverage of Wikipedia, while preserving label-awareness through an entailment framework.

In our benchmarking experiments, **TE-Wiki** outperforms all previous zero-shot methods on four benchmarks from different domains. It also shows competitive performance against weakly-supervised models trained on in-domain data. By learning from Wikipedia, our method does not require any data that is specifically collected from the evaluation domains. On the other hand, since our model is label-aware, it can flexibly classify text pieces into any labels outside Wikipedia.

Finally, we compare our system against humans for further insights. We show that even humans are sometimes confused by ambiguous labels through a crowdsourcing study, which explains the performance gap between open-domain and supervised classification. The gap is reduced significantly when label meanings are clear and well aligned with the semantics of text. We also use an example to illustrate the negative effect of a bad label name. Through the analysis, we demonstrate the importance of choosing proper label names in open-domain classification.

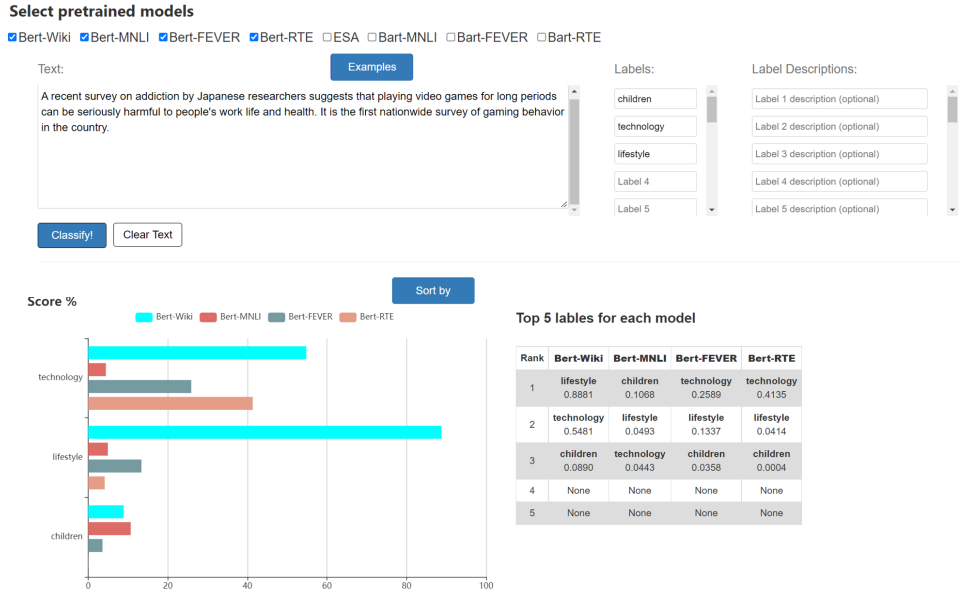


Figure 1: An overview of our open-domain topic classification system. Users can choose multiple models (*top*), and define their own text input and candidate labels (*middle*). Prediction results from different models are displayed in the bar chart and the table (*bottom*).

2 Related Work

Open-domain zero-shot text classification was first studied in the NLP domain in (Chang et al., 2008) (under the name “dataless classification”) as a method that classifies solely based on a general knowledge source and does not require any in-domain data, whether labeled or not. It was proposed to embed both the text and labels into the same semantic space, via Explicit Semantic Analysis, or ESA (Gabrilovich and Markovitch, 2007), and pick the label with the highest relevance score. This idea was further extended to hierarchical (Song and Roth, 2014) and cross-lingual (Song et al., 2016) text classification. Later on, (Yin et al., 2019) called this protocol “label fully unseen” and proposed an entailment approach to transfer knowledge from textual entailment to text classification. It formulates an n -class classification problem as n binary entailment problems by converting labels into hypotheses and the text into the premise, and selects the premise-hypothesis pair with highest entailment score. More recently, another concurrent work (Chu et al., 2021) proposed to explore resources from Wikipedia for zero-shot text classification, but with a different formulation.

There are many other methods that also require less labeling than supervised classification, though in slightly different settings. For example, previous works have explored to generalize from a set of known classes (with annotation) to unknown classes (without annotation) using word embed-

dings of label names (Pushp and Srivastava, 2017; Xia et al., 2018; Liu et al., 2019), class correlation on knowledge graphs (Rios and Kavuluru, 2018; Zhang et al., 2019), or joint embeddings of documents and labels (Nam et al., 2016). Besides, Weakly supervised approaches (Mekala and Shang, 2020; Meng et al., 2020) learn from an unlabeled, but in-domain training set. Given a set of predefined labels, a label-aware knowledge mining step is first applied to find class-specific indicators from the corpus, followed by another self training step to further enhance the model by propagating the knowledge to the whole corpus. However, none of these approaches are suitable for building an open-domain classification system. They either require domain-specific annotation or knowing test labels beforehand.

3 System Description

We present details about our open-domain topic classification system, starting with an overview of our web interface, followed by the backend model.

3.1 User Interface

Figure 1 is a snapshot of our online demo. The system is supported by multiple backend models for test and comparison. Among them, “Bert-Wiki”, corresponding to **TE-Wiki** in this paper, is the best-performing one in our evaluation. After selecting the model(s), users can create their own taxonomy in the “Labels” column, and input the text snippet.

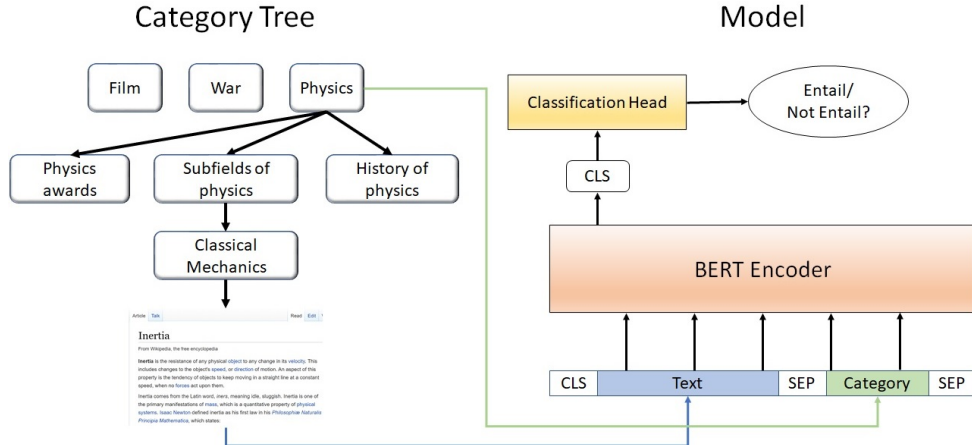


Figure 2: An overview of our proposed **TE-Wiki**. *Left*: the data collection process. For each of the top-level categories, we run DFS to find its descendant categories as well as their member articles. These articles are paired with the root category for model input. *Right*: the model architecture. We use BERT for sequence classification. The article text is concatenated with the category name to feed into a BERT encoder. The classification head takes the output embedding of the "[CLS]" token to classify the input text-category pair.

The system will then classify the text with the user-defined taxonomy. Results are presented in two formats: a bar chart and a ranking table. The table on the right provides a clear view of rankings by each model, while the bar chart on the left is useful to compare the scale of the scores from different models for different labels. These scores, ranging from 0 to 1, are probabilities of the label being relevant to the text, which we will explain further in the next section.

Consider the example in Figure 1. The input text is most relevant to *lifestyle*, somewhat relevant to *technology*, and irrelevant to *children*, which aligns with the prediction of our “Bert-Wiki” model.

3.2 TE-Wiki

We now describe our best performing model **TE-Wiki**. Previous work (Yin et al., 2019) has demonstrated that an n -way classification problem can be converted into n binary entailment problems. Specifically, we can use the text as the premise, and candidate labels as the hypotheses, to generate n statements “[Text] **Entails** [Label $_i$]” for $i \in [n]$. The motivation is that classification is essentially a special kind of entailment. Suppose we want to classify a document into 3 classes: *politics*, *business*, *sports*. We can ask three binary questions: “Is it about politics?”, “Is it about business?”, “Is it about sports?”. By doing so, the model is no longer constrained to a fixed label set, as we can always ask more questions to handle new labels.

With the above framework, it is straightforward to train a model on an entailment dataset (e.g. MNL (Williams et al., 2018), FEVER (Thorne

et al., 2018), RTE (Dagan et al., 2005; Wang et al., 2019).) and use it for classification. However, this may not be the optimal choice as topic classification only focuses on high-level concepts, while textual entailment has a much wider scope and involves many other aspects (e.g., see (Dagan et al., 2013)). Therefore, we propose to construct a new dataset from Wikipedia with articles as premises, and categories as hypotheses. Our desired training pair should meet the following two criteria:

1. The hypothesis is consistent with the premise, i.e. the categorization is correct.
2. The hypothesis should be abstract and concise to reflect the high-level idea of the premise, rather than focus on certain details.

Directly using all the categories associated with an article satisfies the first criterion, but fails with the second, as some of them do not represent the article well. For example, the page *Bill Gates* is assigned *Category: Cornell family*, which is correct about the person but probably not a suitable label for the whole article. To resolve the issue, we instead use higher-level categories on Wikipedia’s subcategory graph to yield better hypotheses.

The overview of **TE-Wiki** is illustrated in Figure 2. Specifically, we start with a set of 700 top-level categories from Wikipedia’s overview page³ as roots. For each of them, we run a depth-first search (DFS) to find its subcategories. In our experiment, we set the max depth to 2 to ensure the subcategories found are strongly affiliated with the

³<https://en.wikipedia.org/wiki/Wikipedia:Contents/Categories>

Algorithm 1: Collect training data

Input : Top-level category set S ,
Wikipedia subcategory graph \mathcal{G} ,
max search depth $r = 2$;
Initialize $d(x, c) = \infty$ for any article $x \in X$ and
 $c \in S$. $M = \{\}$;
for c in S **do**
 $T = \text{DFS}(c, \mathcal{G}, r)$;
 for t in T .nodes **do**
 for x in t .articles **do**
 $d(x, c) = \min\{d(x, c), 1 + \text{depth}(t)\}$;
 end
 end
end
for x in X **do**
 if $\min_{c \in S} d(x, c) < \infty$ **then**
 $P = \text{argmin}_{c \in S} d(x, c)$;
 for c in P **do**
 Add $(x, c, 1)$ to M ;
 end
 Sample c' from $S - P$;
 Add $(x, c', -1)$ to M ;
 end
end
Output : M

root. We collect all member articles of categories in the DFS tree, including both leaves and internal nodes, and pair them with the root to construct positive examples. In case an article can be reached from multiple root categories, we only pair it with the root(s) that has the smallest tree distance to the article to ensure supervision quality. Then for each article, we randomly choose a different category to construct a negative example. While we have tried more sophisticated negative sampling strategies with the aim to confuse the model, none of them makes a significant improvement. Thus, we keep to this simple version. The final training set $D = \{(x_i, c_i, p_i)_{i=1}^n\}$ consists of 3-tuples such that x_i is a Wikipedia article, c_i is the corresponding high-level category name, and $p_i \in \{+1, -1\}$ is the label. The procedure for constructing the training set is summarized in Algorithm 1.

We then fine-tune the pre-trained BERT model (Devlin et al., 2019) with the collected dataset. Given a tuple (x_i, c_i, p_i) , the concatenation of x_i and c_i is passed to a BERT encoder, followed by a classification head to predict whether the article x_i belongs to the category c_i . During test, (i) for the single-labeled case, we pick the label with the highest predicted probability, (ii) for the multi-labeled case, we pick all labels predicted as positive (i.e. probability > 0.5). We do not use any hypothesis template to convert label names into sentences as in (Yin et al., 2019), for consistency with training.

Dataset	#Classes	#samples
Yahoo (Zhang et al., 2015)	10	100,000
Situation (Mayhew et al., 2019)	12	3525
AG News (Zhang et al., 2015)	4	7,600
DBpedia (Lehmann et al., 2015)	14	70,000

Table 1: Dataset Statistics.

4 Evaluation

We evaluate all the backend models of our system on four classification benchmarks to compare their performance. We also compare them against weakly-supervised and supervised models to quantify how much we can achieve without any domain-specific training data.

4.1 Experiment setup

Datasets: We summarize all test datasets in Table 1. For *Yahoo! Answers*, we use the reorganized train/test split by (Yin et al., 2019). All datasets are in English. Among the four, *Situation Typing* is a multi-labeled dataset with imbalanced classes, for which we report the weighted average of per-class F1 score. We refer readers to (Yin et al., 2019) for the class distribution statistics. The other three are single-labeled and class-balanced, and we report the classification accuracy.

Models: Apart from **TE-Wiki**, we run five zero-shot models for open-domain evaluation, as well as a weakly-supervised and a supervised model for close-domain comparison.

- **Word2Vec** (Mikolov et al., 2013): To measure cosine similarity between the embedding vectors of text and label.
- **ESA** (Chang et al., 2008): Same as above, except using embeddings in Wikipedia title space
- **TE-MNLI, TE-FEVER, TE-RTE** (Yin et al., 2019): Textual entailment models by finetuning BERT on MNLI, FEVER, and RTE respectively.⁴
- **LOTClass** (Meng et al., 2020): A weakly-supervised method that learns label information from unlabeled, but in-domain training data.
- **BERT** (Devlin et al., 2019): We finetune a supervised BERT on training data for each dataset.

Implementation: We finetune the bert-base-uncased model on the Wikipedia article-category dataset to train **TE-Wiki**. We removed 26 categories whose name starts with "List of" from the 700 top-level categories, resulting in 674 categories

⁴In experiments, we always use bert-base-uncased.

Supervision Type	Methods	Yahoo	Situation	AG News	DBPedia
Zero-shot	Word2Vec (Mikolov et al., 2013)	35.7	15.6	71.1	69.7
	ESA (Chang et al., 2008)	40.4	30.2	71.1	64.7
	TE-MNLI (Yin et al., 2019)	37.9	15.4	68.8	55.3
	TE-FEVER (Yin et al., 2019)	40.1	21.0	78.0	73.0
	TE-RTE (Yin et al., 2019)	43.8	37.2	60.5	65.9
	TE-Wiki	57.3	41.7	79.6	90.2
Weakly-supervised	LOTClass (Meng et al., 2020)	54.7	N/A	86.4	91.1
Supervised	BERT (Devlin et al., 2019)	75.3	58.0	94.4	99.3

Table 2: Test results of all methods on four datasets. Compared with Word2Vec and ESA, ESA-WikiCate is overall the best among the three embedding-based methods. TE-WikiCate outperforms all other zero-shot methods across all four datasets, and performs competitively against the weakly-supervised LOTClass.

as hypotheses and 1,367,784 articles as premises. The final training set contains 3,387,028 article-category pairs. We set the max sequence length to be 128 tokens and the training batch size to be 64. The model is optimized with AdamW (Loshchilov and Hutter, 2019) with initial learning rate as $5e-5$. Since we do not have a development set in the zero-shot setting, we train the model for 1500 steps to prevent overfitting. For all zero-shot methods, we train once and evaluate the model on all test datasets. For supervised and weakly-supervised methods, we train a different model for each different dataset.

4.2 Result Analysis

The main results are presented in Table 2. We observe that **TE-Wiki** performs the best among all zero-shot methods on all four datasets with different labels and from different domains, demonstrating its effectiveness as an open-domain classifier. It also performs closely with the weakly-supervised **LOTClass** which is trained on in-domain data with known taxonomy, showing that an open-domain zero-shot model can achieve similar accuracy as those domain-specific classifiers. In particular, **TE-Wiki** outperforms **LOTClass** on *Yahoo*, whose training set contains quite a few ambiguous examples. These examples can have negative impact on self-training. On the other hand, our zero-shot model does not rely on any domain-specific data, making it more robust against imperfect data.

It is possible that some of the testing labels also appear in the Wikipedia categories used for training.⁵ To ensure the quality and fairness of our zero-shot evaluation, we remove the overlapping categories from Wikipedia training data, and retrain the **TE-Wiki** model for each test set. Specifically, we normalize labels and categories by their lowercased, lemmatized names, and perform a token-

⁵Throughout this subsection, we use the word "category" for the training set and "label" for the testing sets.

based matching. We report in Table 3 the performance before and after deduplication. We find that deduplication has little or even positive influence on performance, which shows that **TE-Wiki** does not rely on seeing test labels during training. In particular, the performance on *Yahoo* gets improved with deduplication. We suspect that exact match between training and testing labels can lead to overfitting, since the same label may have different meanings under different context. Notice that this study is only for justifying our zero-shot evaluation. For real-world applications, excluding overlapping categories is neither necessary nor feasible as users do not know the test labels beforehand in zero-shot scenarios.

4.3 Early stopping and knowledge transfer

To study the convergence of our model, we sample a small dev set of 1000 examples from *Yahoo*'s original validation set. During training, we find that with 25 steps the **TE-Wiki** model already achieves a reasonably good performance on the dev set. Further training for longer steps yields some, but not significant gains. Since the model has only seen $25 \times 64 = 1600$ examples at that point, there is little chance for the model to acquire label specific knowledge with such a small amount of data. Hence, we believe that during the early steps, the model actually learns "*what topic classification is about*", while the knowledge specific to different labels has already been implicitly stored in the pretrained BERT encoder. The category prediction task takes a minor role in transferring world knowledge. Rather, it teaches the model how to use existing knowledge to make a good inference.

5 Importance of label names

Since zero-shot classifiers understand a label by its name, the quality of label names can be an important performance bottleneck in designing open-domain text classification systems. To study this, We con-

	Yahoo	Situation	AG News	DBPedia
TE-Wiki	57.3	41.7	79.6	90.2
- Overlapping categories	59.4(+2.1)	38.8(-2.9)	79.7(+0.1)	88.9(-1.3)
Removed training categories	15	2	4	10
Overlapping test examples (%)	100.0	24.0	100.0	50.0

Table 3: Performance before and after removing the overlapping categories, as well as their difference. We also show the number of removed categories, and the percentage of test documents that belong to the overlapping labels.

	TE-Wiki	Human	Supervised
Yahoo	58.9	64.8	77.4
Yahoo-5	82.1(+23.2)	88.1(+23.3)	89.4(+12.0)
AG News	79.0	80.2	94.7
AG News-5	86.4(+7.4)	90.8(+10.6)	96.4(+1.7)

Table 4: Classification accuracy on crowdsourcing datasets. Yahoo-5 and AG News-5 count only examples for which all five workers choose the same label.

duct crowdsourcing surveys on subsets of *Yahoo* and *AG News*. For each dataset, we randomly sample 1,000 documents while preserving class balance. Every document is independently annotated by five workers. In the survey question, we only provide the document to be classified and names of candidate labels, without giving workers examples for each class. We consider an example to be correctly classified by humans only if at least three workers choose the gold label. Details about the survey are in Appendix.

We summarize the results in Table 4. Row 1&3 are classification accuracy on the whole crowdsourcing datasets, and row 2&4 are on subsets of examples where all 5 workers choose the same label. We observe that when including all examples, both **TE-Wiki** and humans perform much worse than the supervised method. The supervised approach has the advantage that it learns data-specific features to resolve ambiguity among different classes. On the other hand, humans only make judgements based on their understanding of the labels and a stand-alone test document, and so does our zero-shot algorithm. Ideally, this task should not be difficult for humans as long as the labels properly describe the text topics. However, in some cases the labels could be ambiguous and confusing. Figure 3 shows an example of a bad label name leading to a mistake. The word “Reference” in the correct label actually means “quoting other people’s words”. However, it is hard for an ordinary person to understand the meaning without any example as illustration. 4 out of 5 annotators instead chose “Entertainment & Music” due to the movie “Star Wars”. By contrast, the supervised model has no difficulty in making the correct decision because it has seen plenty of quotation examples during




Figure 3: An example with a bad label name. Annotators are confused by the word “Reference”.

training and can easily capture the useful pattern like “Who said XXX”. The main reason for humans’ confusion here is that the label name does not directly reflect the semantics of the text. A better description of the class should be provided for classification without examples.

We also calculate the accuracy on examples where all 5 workers agree, as in row 2&4 in Table 4. We believe the high inter-annotator agreement here indicates a better alignment between the semantics of text and label. We find a significant improvement of human performance on these less ambiguous cases. The same happens to our zero-shot model, but the supervised method benefits much less. Consequently, the performance gap between humans and the supervised model is also getting closer, which demonstrates that ambiguous labels have a strongly negative impact on classification. Therefore, we believe picking good labels is crucial for open-domain topic classification.

6 Conclusion

We introduce a system for open-domain topic classification. The system allows users to define customized taxonomy and classify text with respect to that taxonomy at real time, without changing the underlying model. To build a powerful model, we propose to utilize Wikipedia articles and categories and adopt an entailment framework for zero-shot learning. The resulting **TE-Wiki** outperforms all existing zero-shot baselines in open-domain evaluations. Finally, we demonstrate the importance of choosing proper label names in open-domain topic classification through a crowdsourcing study.

Acknowledgements

This work was supported by Contracts FA8750-19-2-1004 and FA8750-19-2-0201 with the US Defense Advanced Research Projects Agency (DARPA). Approved for Public Release, Distribution Unlimited. The views expressed are those of the authors and do not reflect the official policy or position of the Department of Defense or the U.S. Government.

References

- Ming-Wei Chang, Lev Ratinov, Dan Roth, and Vivek Srikumar. 2008. [Importance of Semantic Representation: Dataless Classification](#). In *Proc. of the Conference on Artificial Intelligence (AAAI)*.
- Zweiwei Chu, Karl Stratos, and Kevin Gimpel. 2021. [NAT-CAT: Weakly supervised text classification with naturally annotated resources](#). In *3rd Conference on Automated Knowledge Base Construction*.
- I. Dagan, O. Glickman, and B. Magnini. 2005. The pascal recognising textual entailment challenge. In *Proceedings of PASCAL first Workshop on Recognising Textual Entailment*.
- Ido Dagan, Dan Roth, Mark Sammons, and Fabio Massimo Zanzoto. 2013. Recognizing Textual Entailment: Models and Applications.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- E. Gabrilovich and S. Markovitch. 2007. Computing semantic relatedness using wikipedia-based explicit semantic analysis. In *Proc. of the International Joint Conference on Artificial Intelligence (IJCAI)*.
- Ken Lang. 1995. NewsWeeder: learning to filter net-news. In *Proc. of the International Conference on Machine Learning (ICML)*.
- Jens Lehmann, Robert Isele, Max Jakob, Anja Jentzsch, Dimitris Kontokostas, Pablo N. Mendes, Sebastian Hellmann, Mohamed Morsey, Patrick van Kleef, Sören Auer, and Christian Bizer. 2015. [Dbpedia - A large-scale, multilingual knowledge base extracted from wikipedia](#). *Semantic Web*, 6(2):167–195.
- Han Liu, Xiaotong Zhang, Lu Fan, Xuandi Fu, Qimai Li, Xiao-Ming Wu, and Albert Y.S. Lam. 2019. [Reconstructing capsule networks for zero-shot intent classification](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4799–4809, Hong Kong, China. Association for Computational Linguistics.
- Ilya Loshchilov and Frank Hutter. 2019. [Decoupled weight decay regularization](#). In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.
- Stephen Mayhew, Tatiana Tsygankova, Francesca Marini, Zihan Wang, Jane Lee, Xiaodong Yu, Xingyu Fu, Weijia Shi, Zian Zhao, Wenpeng Yin, Karthikeyan K. Jamaal Hay, Michael Shur, Jennifer Sheffield, and Dan Roth. 2019. [University of Pennsylvania LoReHLT 2019 Submission](#). Technical report.
- Dheeraj Mekala and Jingbo Shang. 2020. [Contextualized weak supervision for text classification](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 323–333, Online. Association for Computational Linguistics.
- Yu Meng, Yunyi Zhang, Jiaxin Huang, Chenyan Xiong, Heng Ji, Chao Zhang, and Jiawei Han. 2020. [Text classification using label names only: A language model self-training approach](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9006–9017, Online. Association for Computational Linguistics.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Gregory S. Corrado, and Jeffrey Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Advances in Neural Information Processing Systems 26: 27th Annual Conference on Neural Information Processing Systems 2013. Proceedings of a meeting held December 5-8, 2013, Lake Tahoe, Nevada, United States*.
- Jinseok Nam, Eneldo Loza Mencía, and Johannes Fürnkranz. 2016. [All-in text: Learning document, label, and word representations jointly](#). In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, February 12-17, 2016, Phoenix, Arizona, USA*, pages 1948–1954. AAAI Press.
- Pushpankar Kumar Pushp and Muktabh Mayank Srivastava. 2017. [Train once, test anywhere: Zero-shot learning for text classification](#). *CoRR*, abs/1712.05972.
- Anthony Rios and Ramakanth Kavuluru. 2018. [Few-shot and zero-shot multi-label learning for structured label spaces](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3132–3142, Brussels, Belgium. Association for Computational Linguistics.
- Yangqiu Song and Dan Roth. 2014. [On Dataless Hierarchical Text Classification](#). In *Proc. of the Conference on Artificial Intelligence (AAAI)*.

- Yangqiu Song, Shyam Upadhyay, Haoruo Peng, and Dan Roth. 2016. [Cross-lingual Dataless Classification for Many Languages](#). In *Proc. of the International Joint Conference on Artificial Intelligence (IJCAI)*.
- Duyu Tang, Bing Qin, and Ting Liu. 2015. [Document modeling with gated recurrent neural network for sentiment classification](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1422–1432, Lisbon, Portugal. Association for Computational Linguistics.
- James Thorne, Andreas Vlachos, Christos Christodoulopoulos, and Arpit Mittal. 2018. [FEVER: a large-scale dataset for fact extraction and VERification](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 809–819, New Orleans, Louisiana. Association for Computational Linguistics.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. 2019. [GLUE: A multi-task benchmark and analysis platform for natural language understanding](#). In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.
- Sida Wang and Christopher Manning. 2012. [Baselines and bigrams: Simple, good sentiment and topic classification](#). In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 90–94, Jeju Island, Korea. Association for Computational Linguistics.
- Adina Williams, Nikita Nangia, and Samuel Bowman. 2018. [A broad-coverage challenge corpus for sentence understanding through inference](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1112–1122, New Orleans, Louisiana. Association for Computational Linguistics.
- Congying Xia, Chenwei Zhang, Xiaohui Yan, Yi Chang, and Philip Yu. 2018. [Zero-shot user intent detection via capsule neural networks](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3090–3099, Brussels, Belgium. Association for Computational Linguistics.
- Zichao Yang, Diyi Yang, Chris Dyer, Xiaodong He, Alex Smola, and Eduard Hovy. 2016. [Hierarchical attention networks for document classification](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1480–1489, San Diego, California. Association for Computational Linguistics.
- Wenpeng Yin, Jamaal Hay, and Dan Roth. 2019. [Benchmarking Zero-shot Text Classification: Datasets, Evaluation, and Entailment Approach](#). In *Proc. of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Jingqing Zhang, Piyawat Lertvittayakumjorn, and Yike Guo. 2019. [Integrating semantic knowledge to tackle zero-shot text classification](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 1031–1040, Minneapolis, Minnesota. Association for Computational Linguistics.
- Xiang Zhang, Junbo Jake Zhao, and Yann LeCun. 2015. [Character-level convolutional networks for text classification](#). In *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*, pages 649–657.

A Crowdsourcing Setup

We conduct crowdsourcing annotations for 1000 documents sampled from the *Yahoo! Answers* dataset and another 1000 from the *AG News* on Amazon Mechanical Turk (AMTurk). Both crowdsourcing subsets preserve the class-balance as in the original datasets. We avoid using long documents so that each document contains no more than 512 characters. The 1000 samples are split into 40 assignments, each containing 25 examples. We request 5 AMTurk workers for multiple-choice questions on each assignment. In order to ensure the response quality, we use anchor examples and gold annotations from the original datasets to filter out low-quality answers. Specifically, in each assignment we insert two anchor examples that we believe are easy enough for workers to choose the correct answer as long as they pay attention. We reject a submission if a worker’s classification accuracy against gold annotations is below 30%, or both anchor examples are wrongly classified. With a small initial pilot, we estimate the average working time for labeling 25 examples to be 22 minutes, and we set the pay rate to be \$1.5 per assignment for each valid submission. The overall cost is \$300 for 200 valid submissions for each dataset.