

Neural Networks for Joint Sentence Classification in Medical Paper Abstracts

Franck Deroncourt*
MIT
francky@mit.edu

Ji Young Lee*
MIT
jjylee@mit.edu

Peter Szolovits
MIT
psz@mit.edu

Abstract

Existing models based on artificial neural networks (ANNs) for sentence classification often do not incorporate the context in which sentences appear, and classify sentences individually. However, traditional sentence classification approaches have been shown to greatly benefit from jointly classifying subsequent sentences, such as with conditional random fields. In this work, we present an ANN architecture that combines the effectiveness of typical ANN models to classify sentences in isolation, with the strength of structured prediction. Our model outperforms the state-of-the-art results on two different datasets for sequential sentence classification in medical abstracts.

1 Introduction

Over 50 million scholarly articles have been published (Jinha, 2010), and the number of articles published every year keeps increasing (Druss and Marcus, 2005; Larsen and Von Ins, 2010). Approximately half of them are biomedical papers. While this repository of human knowledge abounds with useful information that may unlock new, promising research directions or provide conclusive evidence about phenomena, it has become increasingly difficult to take advantage of all available information due to its sheer amount. Therefore, a technology that can assist a user to quickly locate the information of interest is highly desired, as it may reduce the time required to locate relevant information.

When researchers search for previous literature, for example, they often skim through abstracts in order to quickly check whether the papers match

the criteria of interest. This process is easier when abstracts are *structured*, i.e., the text in an abstract is divided into semantic headings such as objective, method, result, and conclusion. However, a significant portion of published paper abstracts is *unstructured*, which makes it more difficult to quickly access the information of interest. Therefore, classifying each sentence of an abstract to an appropriate heading can significantly reduce time to locate the desired information.

We call this the *sequential sentence classification task*, in order to distinguish it from general text classification or sentence classification that does not have any context. Besides aiding humans, this task may also be useful for automatic text summarization, information extraction, and information retrieval.

In this paper, we present a system based on ANNs for the sequential sentence classification task. Our model makes use of both token and character embeddings for classifying sentences, and has a sequence optimization layer that is learned jointly with other components of the model. We evaluate our model on the NICTA-PIBOSO dataset as well as a new dataset we compiled based on the PubMed database.

2 Related Work

Existing systems for sequential sentence classification are mostly based on naive Bayes (Ruch et al., 2007; Huang et al., 2013), support vector machine (McKnight and Srinivasan, 2003; Yamamoto and Takagi, 2005; Hirohata et al., 2008; Yamamoto and Takagi, 2005), Hidden Markov models (Lin et al., 2006), and conditional random fields (CRFs) (Kim et al., 2011; Hassanzadeh et al., 2014; Hirohata et al., 2008). They often require numerous hand-engineered features based on lexical (bag-of-words, n-grams, dic-

* These authors contributed equally to this work.

tionaries, cue words), semantic (synonyms, hyponyms), structural (part-of-speech tags, headings), and sequential (sentenced position, surrounding features) information.

On the other hand, recent approaches to natural language processing (NLP) based on artificial neural networks (ANNs) do not require manual features, as they are trained to automatically learn features based on word as well as character embeddings. Moreover, ANN-based models have achieved state-of-the-art results on various NLP tasks, including the most relevant task of text classification (Socher et al., 2013; Kim, 2014; Kalchbrenner et al., 2014; Zhang et al., 2015; Conneau et al., 2016; Xiao and Cho, 2016; dos Santos and Gatti, 2014). For text classification, many ANN models use word embeddings (Socher et al., 2013; Kim, 2014; Kalchbrenner et al., 2014; Gehrmann et al., 2017), and most recent works are based on character embeddings (Zhang et al., 2015; Conneau et al., 2016; Xiao and Cho, 2016). Approaches combining word and character embeddings have also been explored (dos Santos and Gatti, 2014; Deroncourt et al., 2016).

However, most existing works using ANNs for short-text classification do not use any context. This is in contrast with sequential sentence classification, where each sentence in a text is classified taking into account its context, i.e. the surrounding sentences and possibly the whole text. One exception is a recent work on dialog act classification (Lee and Deroncourt, 2016), where each utterance in a dialog is classified into its dialog act, but only the preceding utterances were used, as the system was designed with real-time applications in mind.

3 Model

In the following, we denote scalars in italic lowercase (e.g., k , b_f), vectors in bold lowercase (e.g., \mathbf{s} , \mathbf{x}_i), and matrices in italic uppercase (e.g., W_f) symbols. We use the colon notations $x_{i:j}$ and $\mathbf{v}_{i:j}$ to denote the sequences of scalars (x_i, x_{i+1}, \dots, x_j), and vectors ($\mathbf{v}_i, \mathbf{v}_{i+1}, \dots, \mathbf{v}_j$), respectively.

3.1 ANN model

Our ANN model (Figure 1) consists of three components: a hybrid token embedding layer, a sentence label prediction layer, and a label sequence optimization layer.

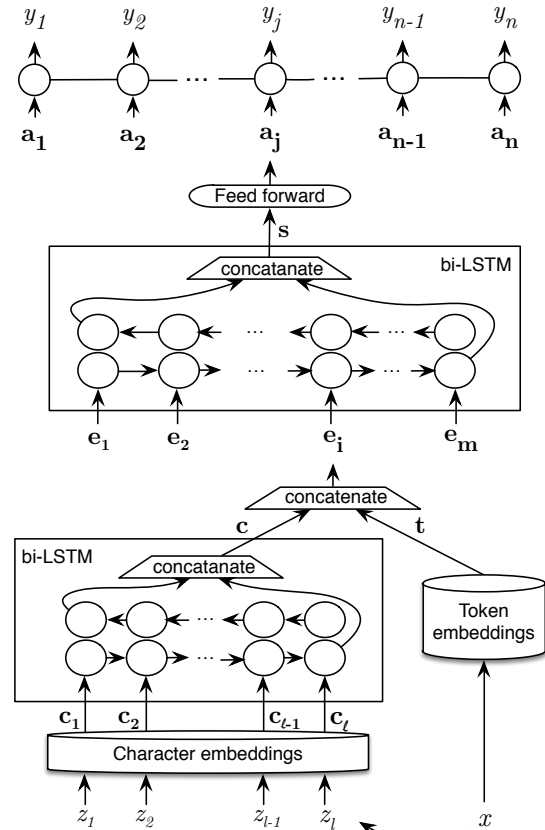


Figure 1: ANN model for sequential sentence classification. x : token, t : token embeddings (300), z_i : i^{th} character of x , c_i : character embeddings (25), c : character-based token embeddings (50), e_i : hybrid token embeddings (350), s : sentence vector (200), \mathbf{a}_j : sentence label vector (number of classes), y_j : sentence label. The numbers in parenthesis indicate the dimension of the vectors. Token embeddings are initialized with GloVe (Pennington et al., 2014) embeddings pretrained on Wikipedia and Gigaword 5 (Parker et al., 2011). Replacing LSTMs with convolutional neural networks did not improve the results: we therefore use LSTMs.

3.1.1 Hybrid token embedding layer

The hybrid token embedding layer takes a token as an input and outputs its vector representation utilizing both the token embeddings and as well as the character embeddings.

Token embeddings are a direct mapping $\mathcal{V}_T(\cdot)$ from token to vector, which can be pre-trained on large unlabeled datasets using programs such as word2vec (Mikolov et al., 2013b; Mikolov et al., 2013a; Mikolov et al., 2013c) or GloVe (Pennington et al., 2014). Character embeddings are also defined in an analogous manner, as a direct mapping $\mathcal{V}_C(\cdot)$ from character to vector.

Let $z_{1:l}$ be the sequence of characters that comprise a token x . Each character z_i is first mapped to its embedding $c_i = \mathcal{V}_C(z_i)$, and the resulting sequence $c_{1:l}$ is input to a bidirectional LSTM, which outputs the character-based token embedding c .

The output \mathbf{e} of the hybrid token embedding layer for the token x is the concatenation of the character-based token embedding \mathbf{c} and the token embedding $\mathbf{t} = \mathcal{V}_T(x)$.

3.1.2 Sentence label prediction layer

Let $x_{1:m}$ be the sequence of tokens in a given sentence, and $\mathbf{e}_{1:m}$ be the corresponding embedding output from the hybrid token embedding layer. The sentence label prediction layer takes as input the sequence of vectors $\mathbf{e}_{1:m}$, and outputs \mathbf{a} , where the k^{th} element of \mathbf{a} , denoted $\mathbf{a}[k]$, reflects the probability that the given sentence has label k .

To achieve this, the sequence $\mathbf{e}_{1:m}$ is first input to a bidirectional LSTM, which outputs the vector representation \mathbf{s} of the given sentence. The vector \mathbf{s} is subsequently input to a feedforward neural network with one hidden layer, which outputs the corresponding probability vector \mathbf{a} .

3.1.3 Label sequence optimization layer

The label sequence optimization layer takes the sequence of probability vectors $\mathbf{a}_{1:n}$ from the label prediction layer as input, and outputs a sequence of labels $y_{1:n}$, where y_i is the label assigned to the token x_i .

In order to model dependencies between subsequent labels, we incorporate a matrix T that contains the transition probabilities between two subsequent labels; we define $T[i, j]$ as the probability that a token with label i is followed by a token with the label j . The score of a label sequence $y_{1:n}$ is defined as the sum of the probabilities of individual labels and the transition probabilities:

$$s(y_{1:n}) = \sum_{i=1}^n \mathbf{a}_i[y_i] + \sum_{i=2}^n T[y_{i-1}, y_i].$$

These scores can be turned into probabilities of the label sequences by taking a softmax function over all possible label sequences:

$$p(\hat{y}_{1:n}) = \frac{e^{s(\hat{y}_{1:n})}}{\sum_{y_{1:n} \in Y^n} e^{s(y_{1:n})}}$$

with Y being the set of all possible labels. During the training phase, the objective is to maximize the log probability of the gold label sequence. In the testing phase, given an input sequence of tokens, the corresponding sequence of predicted labels is chosen as the one that maximizes the score.

Computing the denominator $\sum_{y \in Y^n} e^{s(y_{1:n})}$ can be done in $O(n|C|^2)$ time using dynamic

programming (where $|C|$ denotes the number of classes), as demonstrated below. Let $A_{(n, y_n)}$ be the log of the sum of the scores of all the sequence of length n the last label of which is y_n . Then:

$$\begin{aligned} A_{(n, y_n)} &\stackrel{\text{def.}}{=} \log \left(\sum_{y_{1:(n-1)} \in Y^{n-1}} e^{s(y_{1:n})} \right) \\ &= \log \left(\sum_{y_{1:(n-1)} \in Y^{n-1}} e^{s(y_{1:(n-1)}) + T(y_{n-1}, y_n) + a_n(y_n)} \right) \\ &= \log \left(\sum_{y_{n-1} \in Y} \left(\sum_{y_{1:(n-2)} \in Y^{n-2}} e^{s(y_{1:(n-1)})} \right) e^{T(y_{n-1}, y_n) + a_n(y_n)} \right) \\ &= \log \left(\sum_{y_{n-1} \in Y} e^{A_{(n-1, y_{n-1})}} e^{T(y_{n-1}, y_n) + a_n(y_n)} \right) \end{aligned}$$

Since $A_{(n, y_n)}$ can be computed in $\Theta(|C|)$ time given $\{A_{(n-1, y_{n-1})} | y_{n-1} \in Y\}$, computing $\{A_{(n, y_n)} | y_n \in Y\}$ takes $\Theta(|C|^2)$ time given $\{A_{(n-1, y_{n-1})} | y_{n-1} \in Y\}$. Consequently, computing $\{A_{(n, y_n)} | y_n \in Y\}$ takes $O(n|C|^2)$ time.

4 Experiments

4.1 Datasets

We evaluate our model on the sentence classification task using the following two medical abstract datasets, where each sentence of the abstract is annotated with one label. Table 1 presents statistics on each dataset.

NICTA-PIBOSO This dataset was introduced in (Kim et al., 2011) and was the basis of the ALTA 2012 Shared Task (Amini et al., 2012).

PubMed 20k RCT This corpus was introduced in (Deroncourt et al., 2017)¹. It is based on the PubMed database of biomedical literature and uses 5 sentence labels: objectives, background, methods, results and conclusions

Dataset	$ C $	$ V $	Train	Validation	Test
PubMed	5	68k	15k (195k)	2.5k (33k)	2.5k (33k)
NICTA	6	17k	722 (8k)	77 (0.9k)	200 (2k)

Table 1: Dataset overview. $|C|$ denotes the number of classes, $|V|$ the vocabulary size. For the train, validation and test sets, we indicate the number of abstracts followed by the number of sentences in parentheses.

4.2 Training

The model is trained using stochastic gradient descent, updating all parameters, i.e., token embed-

¹The dataset can be found online at <https://github.com/Franck-Deroncourt/pubmed-rct>

Model	PubMed 20k	NICTA
LR	83.1	71.6
Forward ANN	86.1	75.1
CRF	89.5	81.2
Best published	–	82.0
Our model	90.0	82.7

Table 2: F1-scores on the test set with several baselines, the best published method (Lui, 2012) from the literature, and our model. Since PubMed 20k RCT was introduced in this work, there is no previously published method for this dataset. The presented results for the ANN-based models are the F1-scores on the test set of the run with the highest F1-score on the validation set.

dings, character embeddings, parameters of bidirectional LSTMs, and transition probabilities, at each gradient step. For regularization, dropout is applied to the character-enhanced token embeddings before the label prediction layer. We selected the hyperparameters manually, though we could have used some hyperparameter optimization techniques (Bergstra et al., 2011; Deroncourt and Lee, 2016).

5 Results and Discussion

Table 2 compares our model against several baselines as well as the best performing model (Lui, 2012) in the ALTA 2012 Shared Task, in which 8 competing research teams participated to build the most accurate classifier for the NICTA-PIBOSO corpus.

The first baseline (LR) is a classifier based on logistic regression using n-gram features extracted from the current sentence: it does not use any information from the surrounding sentences. The second baseline (Forward ANN) uses the model presented in (Lee and Deroncourt, 2016): it computes sentence embeddings for each sentence, then classifies the current sentence given a few preceding sentence embeddings as well as the current sentence embedding. The third baseline (CRF) is a CRF that uses n-grams as features: each output variable of the CRF corresponds to a label for a sentence, and the sequence the CRF considers is the entire abstract. The CRF baseline therefore uses both preceding and succeeding sentences when classifying the current sentence. Lastly, the model presented in (Lui, 2012) developed a new approach called feature stacking, which is a meta-learner that combines multiple feature sets, and is the best performing system on NICTA-PIBOSO published in the literature.

Model	PubMed 20k	NICTA
Full model	89.9	82.7
- character emb	89.7	82.7
- pre-train	88.7	78.0
- token emb	88.9	77.0
- seq opt	85.0	72.8

Table 3: Ablation analysis. F1-scores are reported. “- character emb” is our model using only token embeddings, without character-based token embeddings. “- pre-train” is our model where token embeddings are initialized with random values instead of pre-trained embeddings. “- token emb” is our model using only character-based token embeddings, without token embeddings. “- seq opt” is our model without the label sequence optimization layer.




Figure 2: Transition matrix learned on PubMed 20k RCT. The rows represent the label of the previous sentence, the columns represent the label of the current sentence.

The LR system performs honorably on PubMed 20k RCT (F1-score: 83.1), but quite poorly on NICTA-PIBOSO (F1-score: 71.6): this suggests that using the surrounding sentences may be more important in NICTA-PIBOSO than in PubMed 20k RCT.

The Forward ANN system performs better than the LR system, and worse than the CRF: this is expected, as the Forward ANN system only uses the information from the preceding sentences but do not use any information from the succeeding sentences, unlike the CRF.

Our model performs better than the CRF system and the (Lui, 2012) system. We hypothesize that the following four factors give an edge to our model:

No human-engineered features: Unlike most other systems, our model does not rely on any human-engineered features.

No n-grams: While other systems heavily relies on n-grams, our model maps each token to a token embedding, and feeds it as an input to an RNN. This helps combat data scarcity, as for example “chronic tendonitis” and “chronic tendinitis” are

Sentence	Predicted	Actual
This study investigated whether oxytocin can affect attentional bias in social anxiety.	Background	Methods
The biological mechanisms by which oxytocin may be exerting these effects are discussed .	Conclusions	Results
Leuprolide pharmacokinetics were characterized for 11.25 and 30 mg 3-month depot injections.	Conclusions	Results
While, 6%HES 130/0.4 (free flex 6%HES 130/0.4, Fresenius Kabi) infusion was different [...]	Results	Methods
Arterial and central venous blood gas analyses were performed every 20 minutes [...]	Results	Methods
Cytokine responses accompanying [...] immunotherapy [...] have not previously been reported.	Background	Objectives

Table 4: Examples of prediction errors of our model on PubMed 20k RCT. The “predicted” column indicates the label predicted by our model for a given sentence. Our model takes into account all the sentences present in the abstract in which the classified sentence appears. The “actual” column indicates the gold label of the sentence.

	PubMed 20k RCT			
	Precision	Recall	F1-score	Support
Background	71.8	88.2	79.1	3621
Conclusion	93.5	92.9	93.2	4571
Methods	93.7	96.2	94.9	9897
Objectives	78.2	48.1	59.6	2333
Results	94.8	93.1	93.9	9713
Total	90.1	89.9	90.0	30135

Table 5: Results for each class obtained by our model on PubMed 20k RCT.

two different bigrams, but share the same meaning, and their token embeddings should therefore be very similar.

Structured prediction: The labels for all sentences in an abstract are predicted jointly, which improves the coherency between the predicted labels in a given abstract. The ablation analysis presented in Table 3 shows that the sequence optimization layer is the most important component of the ANN model.

Joint learning: Our model learned the features and token embeddings jointly with the sequence optimization.

The sequence information is mostly contained in the transition matrix. Figure 2 presents an example of transition matrix after the model has been trained on PubMed 20k RCT. We can see that it effectively reflects transitions between different labels. For example, it learned that the first sentence of an abstract is most likely to be either discussing objective (0.23) or background (0.26). By the same token, a sentence pertaining to the methods is typically followed by a sentence pertaining to the methods (0.25) or the results (0.17).

Tables 5 and 6 detail the result of our model for each label in PubMed 20k RCT. The main difficulty the classifier has is distinguishing background sentences from objective sentences. In particular, a third of the objective sentences are incor-

	Backg.	Concl.	Methods	Obj.	Res.
Background	3193	28	116	277	7
Conclusions	55	4248	7	0	261
Methods	78	36	9523	35	225
Objectives	1112	1	95	1122	3
Results	11	232	426	1	9043

Table 6: Confusion matrix on PubMed 20k RCT obtained with our model. Rows correspond to actual labels, and columns correspond to predicted the labels. For example, 116 background sentences were predicted as method.

rectly classified as background, which causes the recall for objectives and the precision for background to be low. The classifier has also some difficulty in distinguishing method sentences from result sentences.

Table 4 presents a few examples of prediction errors. Our error analysis suggests that a fair number of sentence labels are debatable. For example, the sentence “We conducted a randomized study comparing strategies X and Y.” belongs to the background according to the gold target, but most humans would classify it as an objective.

6 Conclusions

In this article we have presented an ANN architecture to classify sentences that appear in sequence. We demonstrate that jointly predicting the classes of all sentences in a given text improves the quality of the predictions. Our model outperforms the state-of-the-art results on two datasets for sentence classification in medical abstracts.

Acknowledgments

The authors thank the anonymous reviewers for their insightful comments. The project was supported by Philips Research. The content is solely the responsibility of the authors and does not necessarily represent the official views of Philips Research.

References

- Iman Amini, David Martinez, and Diego Molla. 2012. Overview of the ALTA 2012 Shared Task. In *Australasian Language Technology Association Workshop 2012*, volume 7, page 124.
- James Bergstra, Rémi Bardenet, Yoshua Bengio, and Balázs Kégl. 2011. Algorithms for hyper-parameter optimization. In J. Shawe-Taylor, R. S. Zemel, P. L. Bartlett, F. Pereira, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 24*, pages 2546–2554. Curran Associates, Inc.
- Alexis Conneau, Holger Schwenk, Loïc Barrault, and Yann Lecun. 2016. Very deep convolutional networks for natural language processing. *arXiv preprint arXiv:1606.01781*.
- Franck Deroncourt and Ji Young Lee. 2016. Optimizing neural network hyperparameters with gaussian processes for dialog act classification. *IEEE Spoken Language Technology*.
- Franck Deroncourt, Ji Young Lee, Ozlem Uzuner, and Peter Szolovits. 2016. De-identification of patient notes with recurrent neural networks. *Journal of the American Medical Informatics Association (JAMIA)*.
- Franck Deroncourt, Ji Young Lee, and Peter Szolovits. 2017. PubMed 200k RCT: a dataset for sentence classification in medical paper abstracts. *arXiv preprint arXiv:1703*.
- Cícero Nogueira dos Santos and Maira Gatti. 2014. Deep convolutional neural networks for sentiment analysis of short texts. In *International Conference on Computational Linguistics (COLING)*, pages 69–78.
- Benjamin G Druss and Steven C Marcus. 2005. Growth and decentralization of the medical literature: implications for evidence-based medicine. *Journal of the Medical Library Association*, 93(4):499.
- Sebastian Gehrmann, Yeran Li, Franck Deroncourt, Eric T. Carlson, Joy T. Wu, Jonathan Welt, David W. Grant, Patrick D. Tyler, and Leo A. Celi. 2017. Comparing rule-based and deep learning models for patient phenotyping. *arXiv preprint arXiv:1703*.
- Hamed Hassanzadeh, Tudor Groza, and Jane Hunter. 2014. Identifying scientific artefacts in biomedical literature: The evidence based medicine use case. *Journal of biomedical informatics*, 49:159–170.
- Kenji Hirohata, Naoaki Okazaki, Sophia Ananiadou, Mitsuru Ishizuka, and Manchester Interdisciplinary Biocentre. 2008. Identifying sections in scientific abstracts using conditional random fields. In *International Joint Conference on Natural Language Processing (IJCNLP)*, pages 381–388.
- Ke-Chun Huang, I-Jen Chiang, Furen Xiao, Chun-Chih Liao, Charles Chih-Ho Liu, and Jau-Min Wong. 2013. PICO element detection in medical text without metadata: Are first sentences enough? *Journal of biomedical informatics*, 46(5):940–946.
- Arif E Jinha. 2010. Article 50 million: an estimate of the number of scholarly articles in existence. *Learned Publishing*, 23(3):258–263.
- Nal Kalchbrenner, Edward Grefenstette, and Phil Blunsom. 2014. A convolutional neural network for modelling sentences. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics*. Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics.
- Su Nam Kim, David Martinez, Lawrence Cavedon, and Lars Yencken. 2011. Automatic classification of sentences to support evidence based medicine. *BioMed Central (BMC) Bioinformatics*, 12(2):1.
- Yoon Kim. 2014. Convolutional neural networks for sentence classification. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1746–1751. Association for Computational Linguistics (ACL).
- Peder Olesen Larsen and Markus Von Ins. 2010. The rate of growth in scientific publication and the decline in coverage provided by science citation index. *Scientometrics*, 84(3):575–603.
- Ji Young Lee and Franck Deroncourt. 2016. Sequential short-text classification with recurrent and convolutional neural networks. In *Human Language Technologies 2016: The Conference of the North American Chapter of the Association for Computational Linguistics, NAACL HLT*.
- Jimmy Lin, Damianos Karakos, Dina Demner-Fushman, and Sanjeev Khudanpur. 2006. Generative content models for structural analysis of medical abstracts. *BioNLP06 Linking Natural Language Processing and Biology: Towards Deeper Biological Literature Analysis*, 6:65–72.
- Marco Lui. 2012. Feature stacking for sentence classification in evidence-based medicine. In *Australasian Language Technology Workshop 2012: ALTA Shared Task*, page 134.
- Larry McKnight and Padmini Srinivasan. 2003. Categorization of sentence types in medical abstracts. In *American Medical Informatics Association (AMIA)*.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013a. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013b. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pages 3111–3119.

- Tomas Mikolov, Wen-tau Yih, and Geoffrey Zweig. 2013c. Linguistic regularities in continuous space word representations. In *HLT-NAACL*, pages 746–751.
- Robert Parker, David Graff, Junbo Kong, Ke Chen, and Kazuaki Maeda. 2011. English Gigaword fifth edition. Technical report, Linguistic Data Consortium, Philadelphia.
- Jeffrey Pennington, Richard Socher, and Christopher D Manning. 2014. GloVe: global vectors for word representation. *Proceedings of the Empirical Methods in Natural Language Processing (EMNLP 2014)*, 12:1532–1543.
- Patrick Ruch, Celia Boyer, Christine Chichester, Imad Tbahriti, Antoine Geissbühler, Paul Fabry, Julien Gobeill, Violaine Pillet, Dietrich Rebholz-Schuhmann, Christian Lovis, et al. 2007. Using argumentation to extract key sentences from biomedical abstracts. *International journal of medical informatics*, 76(2):195–200.
- Richard Socher, Alex Perelygin, Jean Y Wu, Jason Chuang, Christopher D Manning, Andrew Y Ng, and Christopher Potts. 2013. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the conference on empirical methods in natural language processing (EMNLP)*, volume 1631, page 1642. Citeseer.
- Yijun Xiao and Kyunghyun Cho. 2016. Efficient character-level document classification by combining convolution and recurrent layers. *arXiv preprint arXiv:1602.00367*.
- Yasunori Yamamoto and Toshihisa Takagi. 2005. A sentence classification system for multi biomedical literature summarization. In *21st International Conference on Data Engineering Workshops (ICDEW'05)*, pages 1163–1163. IEEE.
- Xiang Zhang, Junbo Zhao, and Yann LeCun. 2015. Character-level convolutional networks for text classification. In *Advances in Neural Information Processing Systems (NIPS)*, pages 649–657.