

A Comparison of Word-based and Context-based Representations for Classification Problems in Health Informatics

Aditya Joshi[♣], Sarvnaz Karimi[♣], Ross Sparks[♣], Cécile Paris[♣], C Raina MacIntyre[♣]

[♣]CSIRO Data61, Sydney, Australia

[♣]Kirby Institute, University of New South Wales, Sydney, Australia

{firstname.lastname}@csiro.au , r.macintyre@unsw.edu.au

Abstract

Distributed representations of text can be used as features when training a statistical classifier. These representations may be created as a composition of word vectors or as context-based sentence vectors. We compare the two kinds of representations (word versus context) for three classification problems: influenza infection classification, drug usage classification and personal health mention classification. For statistical classifiers trained for each of these problems, context-based representations based on ELMo, Universal Sentence Encoder, Neural-Net Language Model and FLAIR are better than Word2Vec, GloVe and the two adapted using the MESH ontology. There is an improvement of 2-4% in the accuracy when these context-based representations are used instead of word-based representations.

1 Introduction

Distributed representations (also known as ‘embeddings’) are dense, real-valued vectors that capture semantics of concepts (Mikolov et al., 2013). When learned from a large corpus, embeddings of related words are expected to be closer than those of unrelated words. When a statistical classifier is trained, distributed representations of textual units (such as sentences or documents) in the training set can be used as feature representations of the textual unit. This technique of statistical classification that uses embeddings as features has been shown to be useful for many Natural Language Processing (NLP) problems (Zhang et al., 2015; Joshi et al., 2016; Chou et al., 2016; Simova and Uszkoreit, 2017; Fu et al., 2016; Buscaldi and Priego, 2017) and biomedical NLP problems (Yadav et al., 2017; Kholghi et al., 2016). In this paper, we experiment with three classification problems in health informatics: influenza infection classification, drug usage classification and

personal health mention classification. We use statistical classifiers trained on tweet vectors as features. To compute a tweet vector, *i.e.*, a distributed representation for tweets, typical alternatives are: (a) tweet vector as a function of word embeddings of the content words¹ in the tweet; or, (b) a contextualised representation that computes sentence vectors using language models. The former considers meanings of words in isolation, while the latter takes into account the order of these words in addition to their meaning. We compare word-based and context-based representations for the three classification problems. This paper investigates the question:

‘When statistical classifiers are trained on vectors of tweets for health informatics, how should the vector be computed: using word-based representations that consider words in isolation or context-based representations that account for word order using language models?’

For these classification problems, we compare five approaches that use word-based representations with four approaches that use context-based representations.

2 Related Work

Distributed representations as features for statistical classification have been used for many NLP problems: semantic relation extraction (Hashimoto et al., 2015), sarcasm detection (Joshi et al., 2016), sentiment analysis (Zhang et al., 2015; Tkachenko et al., 2018), co-reference resolution (Simova and Uszkoreit, 2017), grammatical error correction (Chou et al., 2016), emotion intensity determination (Buscaldi

¹Content words refers to all words except stop words.

	Representation	Details
Word-based	A tweet vector is the average of the vectors of the content words in the tweet.	
	Word2Vec_PreTrain, GloVe_PreTrain	Vectors of the content words are obtained from pre-trained embeddings from Word2Vec & GloVe respectively.
	Word2Vec_SelfTrain	Vectors of the content words are based on embeddings learned from the training set, separately for each fold.
	Word2Vec_WithMeSH, Glove_WithMeSH	Vectors of the content words are pre-trained word embeddings from Word2Vec & GloVe (respectively) retrofitted using MeSH ontology.
Context-based	A tweet vector is obtained from a pre-trained language model that uses context.	
	ELMo, USE, NNLM, FLAIR	Context-based representations of tweets are obtained from pre-trained models of ELMo, USE, NNLM and FLAIR respectively. They account for relationship between words using language models.

Table 1: Summary of the representations used in our experiments.

and Priego, 2017) and sentence similarity detection (Fu et al., 2016). In terms of the biomedical domain, word embedding-based features have been used for entity extraction in biomedical corpora (Yadav et al., 2017) or clinical information extraction (Kholghi et al., 2016). Several approaches for personal health mention classification have been reported (Aramaki et al., 2011; Lamb et al., 2013a; Yin et al., 2015). Aramaki et al. (2011) use bag-of-words as features for personal health mention classification. Lamb et al. (2013a) use linguistic features including coarse topic-based features, while Yin et al. (2015) use features based on parts-of-speech and dependencies for a statistical classifier. Feng et al. (2018) compare statistical classifiers with deep learning-based classifiers for personal health mention detection. In terms of detecting drug-related content in text, there has been work on detecting adverse drug reactions (Karimi et al., 2015). Nikfarjam et al. (2015) use word embedding clusters as features for adverse drug reaction detection.

3 Representations

A tweet vector is a distributed representation of a tweet, and is computed for every tweet in the training set. The tweet vector along with the output label is then used to train the statistical classification model. The intuition is that the tweet vector captures the semantics of the tweet and, as a result, can be effectively used for classification. To obtain tweet vectors, we experiment with two alter-

natives that have been used for several text classification problems in NLP: word-based representations and context-based representations. They are summarised in Table 1, and described in the following subsections.

3.1 Word-based Representations

A word-based representation of a tweet combines word embeddings of the content words in the tweet. We use the average of the word embeddings of content words in the tweet. Average of word embeddings have been used for different NLP tasks (De Boom et al., 2016; Yoon et al., 2018; Orasan, 2018; Komatsu et al., 2015; Ettinger et al., 2018). As in past work, words that were not learned in the embeddings are dropped during the computation of the tweet vector. We experiment with three kinds of word embeddings:

1. **Pre-trained Embeddings:** Denoted as *Word2Vec_PreTrained* and *GloVe_PreTrained* in Table 1, we use pre-trained embeddings of words learned from large text corpora: (A) Word2Vec by Mikolov et al. (2013): This has been pre-trained on a corpus of news articles with 300 million tokens, resulting in 300-dimensional vectors; (B) GloVe by Pennington et al. (2014): This has been pre-trained on a corpus of tweets with 27 billion tokens, resulting in 200-dimensional vectors.
2. **Embeddings Trained on The Training Split:** It may be argued that, since the pre-trained embeddings are learned from a cor-

Classification	# tweets (# true tweets)
IIC	9,006 (2,306)
DUC	13,409 (3,167)
PHMC	2,661 (1,304)

Table 2: Dataset statistics.

pus from an unrelated domain (news and general, in the case of Word2Vec and GloVe respectively), they may not capture the semantics of the domain of the specific classification problem. Therefore, we also use the Word2Vec Model available in the gensim library (Řehůřek and Sojka, 2010) to learn word embeddings from the documents. For each split, the corresponding training set is used to learn the embeddings. The embeddings are then used to compute the tweet vectors and train the classifier. We refer to these as *Word2Vec_SelfTrain*.

3. **Pre-trained embeddings retrofitted with medical ontologies:** Another alternative to adapt word embeddings for a classification problem is to use structured resources (such as ontologies) from a domain same as that of the classification problem. Faruqui et al. (2015) show that word embeddings can be retrofitted to capture relationships in an ontology. We use the Medical Subject Headings (MeSH) ontology (Nelson et al., 2001), maintained by the U.S. National Library of Medicine, which provides a hierarchically-organised terminology of medical concepts. Using the algorithm by Faruqui et al. (2015), we retrofit pre-trained embeddings from Word2Vec and GloVe, with the MeSH ontology. The retrofitted embeddings for Word2Vec and GloVe are referred to as *Word2Vec_WithMeSH*, and *GloVe_WithMeSH* respectively.

The three kinds of word-based representations result in five configurations: *Word2Vec_PreTrained*, *GloVe_PreTrained*, *Word2Vec_SelfTrain*, *Word2Vec_WithMeSH*, and *GloVe_WithMeSH*.

3.2 Context-based Representations

Context-based representations may use language models to generate vectors of sentences. Therefore, instead of learning vectors for individual

words in the sentence, they compute a vector for sentences on the whole, by taking into account the order of words and the set of co-occurring words.

We experiment with four deep contextualised vectors: (A) **Embeddings from Language Models (ELMo)** by Peters et al. (2018): ELMo uses character-based word representations and bidirectional LSTMs. The pre-trained model computes a contextualised vector of 1024 dimensions. ELMo is available in the Tensorflow Hub², a repository of machine learning modules; (B) **Universal Sentence Encoder (USE)** by Cer et al. (2018): The encoder uses a Transformer architecture that uses attention mechanism to incorporate information about the order and the collection of words (Vaswani et al., 2017). The pre-trained model of USE that returns a vector of 512 dimensions is also available on Tensorflow Hub; (C) **Neural-Net Language Model (NNLM)** by Bengio et al. (2003): The model simultaneously learns representations of words and probability functions for word sequences, allowing it to capture semantics of a sentence. We use a pre-trained model available on Tensorflow Hub, that is trained on the English Google News 200B corpus, and computes a vector of 128 dimensions; (D) **FLAIR** by Akbik et al. (2018): This library by Zalando research³ uses character-level language models to learn contextualised representations. We use the pooling option to create sentence vectors. This is a concatenation of GloVe embeddings and the forward/backward language model. The resultant is a vector of 4196 dimensions.

Table 1 refers to the four configurations as *ELMo*, *USE*, *NNLM* and *FLAIR* respectively.

4 Experiment Setup

We conduct our experiments on three boolean classification problems in health informatics: (A) **Influenza Infection Classification (IIC)**: The goal is to predict if a tweet reports an influenza infection (*‘I have been coughing all day’*, for example) or describes information about influenza (*‘flu outbreaks are common in this month of the year’*, for example). We use the dataset presented in Lamb et al. (2013b); (B) **Drug Usage Classification (DUC)**: The objective here is to

²<https://www.tensorflow.org/hub/>; Accessed on 3rd June, 2019.

³<https://github.com/zalandoresearch/flair>; Accessed on 3rd June, 2019.

	# dim.	IIC	DUC	PHMC
(A) <i>Word-based Representations</i>				
Word2Vec_PreTrain	300	0.8106 (σ : 0.024)	0.7417 (σ : 0.153)	0.7632 (σ : 0.037)
GloVe_PreTrain	200	0.7996 (σ : 0.015)	0.7549 (σ : 0.120)	0.7765 (σ : 0.033)
Word2Vec_SelfTrain	300	0.5099 (σ : 0.001)	0.7450 (σ : 0.028)	0.7418 (σ : 0.003)
Word2Vec_WithMeSH	300	0.6944 (σ : 0.021)	0.7450 (σ : 0.046)	0.7427 (σ : 0.050)
GloVe_WithMeSH	200	0.7264 (σ : 0.017)	0.7635 (σ : 0.030)	0.7425 (σ : 0.010)
(B) <i>Context-based Representations</i>				
ELMo	1024	0.8010 (σ : 0.021)	0.7724 (σ : 0.090)	0.7814 (σ : 0.02)
USE	512	0.8164 (σ : 0.008)	0.7790 (σ : 0.100)	0.8155 (σ : 0.030)
NNLM	128	0.8520 (σ : 0.006)	0.7610 (σ : 0.070)	0.7495 (σ : 0.020)
FLAIR	4196	0.8000 (σ : 0.021)	0.7667 (σ : 0.116)	0.7896 (σ : 0.031)

Table 3: Comparison of five word-based representations with four context-based representations; Average accuracy with standard deviation (σ) indicated in brackets.

detect whether or not a tweet describes the usage of a medicinal drug (*‘I took some painkillers this morning’*, for example). We use the dataset provided by Jiang et al. (2016); (C) **Personal Health Mention classification (PHMC)**: A personal health mention is a person’s report about their illness. We use the dataset provided by Robinson et al. (2015). For example *‘I have been sick for a week now’* is a personal health mention while *‘Rollercoasters can make you sick’* is not. It must be noted that IIC involves influenza while the PHMC dataset covers a set of illnesses as described later.

The datasets for each of the classification problems consist of tweets that have been manually annotated as reported in the corresponding papers. The statistics of these datasets are shown in Table 2. The values in brackets indicate the number of true tweets (*i.e.*, tweets that have been labeled as true), since these are boolean classification problems. For details on inter-annotator agreement and the annotation techniques, we refer the reader to the original papers. Based on sentence vectors obtained using either word-based or context-based representations, we train logistic regression with default parameters available as a part of the Liblinear package (Fan et al., 2008). We report five-fold cross-validation results for our experiments. Each fold is created using stratified k-fold sampling available in scikit-learn⁴.

⁴<https://scikit-learn.org/stable/>; Accessed on 3rd June, 2019.

5 Results

We first present a quantitative evaluation to compare the two types of representations. Following that, we analyse sources of errors.

5.1 Quantitative Evaluation

We compare word-based and context-based representations for the three classification problems in Table 3. Accuracy is computed as the proportion of correctly classified instances. The table contains the average accuracy values with standard deviation values shown in parentheses. The table is divided into two parts. Part (A) corresponds to experiments using word-based representations, while Part (B) corresponds to those using context-based representations. In general, context-based representations result in an improvement in the three classification problems as compared to word-based representations. For IIC, the best word-based representation is when pre-trained Word2Vec embeddings (*Word2Vec_PreTrain*) of content words are averaged to generate the tweet vector. The accuracy in this case is 0.8106. In contrast, the best performing context-based representation is NNLM (0.8520). This is an improvement of 4% points. Similarly, tweet vectors created using USE result in an accuracy of 0.7790 for DUC and 0.8155 for PHMC. This is an improvement of 2-4% points each over the word-based representations for these two classification problems as well. In addition, for pre-trained embeddings (Word2Vec and GloVe) retrofitted with a medical ontology (MeSH), we observe a degrada-

	1 st -person men- tions		Present Partici- ple	
	Word	Context	Word	Context
IIC	58.2	41.0	79.6	72.5
DUC	66.4	54.75	33.0	40.75
PHMC	64.8	37.5	61.6	40.0

Table 4: Average number of instances (out of 100 randomly sampled mis-classified instances) containing first-person mentions and present participle form for the three classification problems and two types of representations.

tion in the accuracy for IIC and PHMC, as compared to without retrofitting. There is an improvement of 1% point in the case of DUC. Similarly, learning the embeddings on the specific training corpus does not work well. It leads to a degradation as compared to pre-trained embeddings. This could happen because pre-trained embeddings are trained on much larger corpora than our training datasets, thereby capturing semantics more effectively than the *Word2Vec_SelfTrain* variant.

5.2 Qualitative Evaluation

For a qualitative comparison of the two representations, we analyse 100 randomly sampled instances that are mis-classified by each classifier. While these instances need not be the same for each classifier, the trends in the errors show where one kind of representation scores over the other. We compared linguistic properties of these mis-classified instances, such as the person, tense and number. Table 4 shows two linguistic properties where we observed the most variation: first-person mentions and the use of present participles. The two properties are important in terms of the semantics of the three classification problems. First-person mentions are useful indicators to identify if the speaker has influenza, took a drug or reported a personal health mention. Similarly, present participle forms of verbs appear in situations where a person has had an infection or taken a drug. For ‘Word’, the average is over the five representations, while for ‘Context’, the average is over the four context-based representations. In the case of IIC, an average of 58.2 mis-classified instances from word-based representations contained first person mentions. The corresponding number for context-based representations was 41. For PHMC, the averages are 64.8 (word-based) and 37.5 (context-

based). The difference is not as high in the case of DUC (66.4 and 54.75 respectively). Differences are observed in the case of present participle in mis-classified instances. However, in the case of DUC, errors from context-based representations contain more average number of present participles (40.75) than word-based representations (33).

6 Conclusions

In this paper, we show that context-based representations are a better choice than word-based representations to create tweet vectors for classification problems in health informatics. We experiment with three such problems: influenza infection classification, drug usage classification and personal health mention classification, and compare word-based representations with context-based representations as features for a statistical classifier. For word-based representations, we consider pre-trained embeddings of Word2Vec and GloVe, embeddings trained on the training split, and the pre-trained embeddings of Word2Vec and GloVe retrofitted to a medical ontology. For context-based representations, we consider ELMo, USE, NNLM and FLAIR. For the three problems, the highest accuracy is obtained using context-based representations. In comparison with pre-trained embeddings, the improvement in classification is approximately 4% for influenza infection classification, 2% for drug usage classification and 4% for personal health mention classification. Embeddings trained on the training corpus or retrofitted on the ontology perform worse than those pre-trained on a large corpus.

While these observations are based on statistical classifiers, the corresponding benefit of context-based representations on neural architectures can be validated as a future work. In addition, while we average the word vectors to obtain tweet vectors, other options for tweet vector computation can be considered for word-based representations. In terms of the dataset, the comparison should be validated for text forms other than tweets, such as medical records. Medical records are expected to have typical challenges such as the use of abbreviations and domain-specific phrases that may not have been learned in pre-trained embeddings.

Acknowledgment

The authors would like to thank the anonymous reviewers for their helpful comments.

References

- Alan Akbik, Duncan Blythe, and Roland Vollgraf. 2018. Contextual string embeddings for sequence labeling. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 1638–1649.
- Eiji Aramaki, Sachiko Maskawa, and Mizuki Morita. 2011. Twitter catches the flu: Detecting influenza epidemics using Twitter. In *Empirical Methods in Natural Language Processing*, pages 1568–1576, Edinburgh, UK. Association for Computational Linguistics.
- Yoshua Bengio, Réjean Ducharme, Pascal Vincent, and Christian Jauvin. 2003. A neural probabilistic language model. *Journal of machine learning research*, 3(Feb):1137–1155.
- Davide Buscaldi and Belem Priego. 2017. Lipn-uam at emoint-2017: combination of lexicon-based features and sentence-level vector representations for emotion intensity determination. In *8th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*, pages 255–258, Copenhagen, Denmark.
- Daniel Cer, Yinfei Yang, Sheng-yi Kong, Nan Hua, Nicole Limtiaco, Rhomni St John, Noah Constant, Mario Guajardo-Cespedes, Steve Yuan, Chris Tar, et al. 2018. Universal sentence encoder. *arXiv preprint arXiv:1803.11175*.
- Wei-Chieh Chou, Chin-Kui Lin, Yuan-Fu Liao, and Yih-Ru Wang. 2016. Word order sensitive embedding features/conditional random field-based chinese grammatical error detection. In *3rd Workshop on Natural Language Processing Techniques for Educational Applications*, pages 73–81, Osaka, Japan. The COLING 2016 Organizing Committee.
- Cedric De Boom, Steven Van Canneyt, Thomas De-meester, and Bart Dhoedt. 2016. Representation learning for very short texts using weighted word embedding aggregation. *Pattern Recognition Letters*, 80:150–156.
- Allyson Ettinger, Ahmed Elgohary, Colin Phillips, and Philip Resnik. 2018. Assessing composition in sentence vector representations. In *27th International Conference on Computational Linguistics*, pages 1790–1801, Santa Fe, New Mexico.
- Rong-En Fan, Kai-Wei Chang, Cho-Jui Hsieh, Xiang-Rui Wang, and Chih-Jen Lin. 2008. Liblinear: A library for large linear classification. *Journal of machine learning research*, 9:1871–1874.
- Manaal Faruqui, Jesse Dodge, Sujay Kumar Jauhar, Chris Dyer, Eduard Hovy, and Noah Smith. 2015. Retrofitting word vectors to semantic lexicons. In *Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1606–1615.
- Shichao Feng, Keyuan Jiang, Jiyun Li, Ricardo A Calix, and Matrika Gupta. 2018. Detecting personal experience tweets for health surveillance using unsupervised feature learning and recurrent neural networks. In *Workshops at the 32nd AAAI Conference on Artificial Intelligence*, pages 425–430, New Orleans, Louisiana.
- Cheng Fu, Bo An, Xianpei Han, and Le Sun. 2016. Is-cas_nlp at semeval-2016 task 1: Sentence similarity based on support vector regression using multiple features. In *10th International Workshop on Semantic Evaluation*, pages 645–649, San Diego, California.
- Kazuma Hashimoto, Pontus Stenetorp, Makoto Miwa, and Yoshimasa Tsuruoka. 2015. Task-oriented learning of word embeddings for semantic relation classification. *arXiv preprint arXiv:1503.00095*.
- Keyuan Jiang, Ricardo Calix, and Matrika Gupta. 2016. Construction of a personal experience tweet corpus for health surveillance. In *ACL Workshop on biomedical natural language processing*, pages 128–135, Berlin, Germany.
- Aditya Joshi, Vaibhav Tripathi, Kevin Patel, Pushpak Bhattacharyya, and Mark Carman. 2016. Are word embedding-based features useful for sarcasm detection? In *Empirical Methods in Natural Language Processing*, pages 1006–1011, Austin, Texas.
- Sarvnaz Karimi, Chen Wang, Alejandro Metke-Jimenez, Raj Gaire, and Cecile Paris. 2015. Text and data mining techniques in adverse drug reaction detection. *ACM Computing Surveys*, 47(4):56.
- Mahnoosh Kholghi, Lance De Vine, Laurianne Sitbon, Guido Zuccon, and Anthony Nguyen. 2016. The benefits of word embeddings features for active learning in clinical information extraction. In *Australasian Language Technology Association Workshop*, pages 25–34, Melbourne, Australia.
- Hiroya Komatsu, Ran Tian, Naoaki Okazaki, and Kentaro Inui. 2015. Reducing lexical features in parsing by word embeddings. In *Pacific Asia Conference on Language, Information and Computation*, pages 106–113, Shanghai, China.
- Alex Lamb, Michael J Paul, and Mark Dredze. 2013a. Separating fact from fear: Tracking flu infections on twitter. In *North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 789–795, Atlanta, Georgia.
- Alex Lamb, Michael J. Paul, and Mark Dredze. 2013b. Separating fact from fear: Tracking flu infections on twitter. In *North American Chapter of the Association for Computational Linguistics: Human Language Technologies*.

- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pages 3111–3119.
- Stuart J Nelson, W Douglas Johnston, and Betsy L Humphreys. 2001. Relationships in medical subject headings (mesh). In *Relationships in the Organization of Knowledge*, pages 171–184.
- Azadeh Nikfarjam, Abeed Sarker, Karen Oconnor, Rachel Ginn, and Graciela Gonzalez. 2015. Pharmacovigilance from social media: mining adverse drug reaction mentions using sequence labeling with word embedding cluster features. *Journal of the American Medical Informatics Association*, 22(3):671–681.
- Constantin Orasan. 2018. Aggressive language identification using word embeddings and sentiment features. In *1st Workshop on Trolling, Aggression and Cyberbullying*, pages 113–119, Santa Fe, New Mexico.
- Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. Glove: Global vectors for word representation. In *Empirical Methods in Natural Language Processing*, pages 1532–1543.
- Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. In *North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2227–2237, New Orleans, Louisiana.
- Radim Řehůřek and Petr Sojka. 2010. Software Framework for Topic Modelling with Large Corpora. In *LREC Workshop on New Challenges for NLP Frameworks*, pages 45–50, Valletta, Malta. <http://is.muni.cz/publication/884893/en>.
- Bella Robinson, Ross Sparks, Robert Power, and Mark Cameron. 2015. Social media monitoring for health indicators. In *International Congress on Modelling and Simulation*, Gold Coast, Australia.
- Iliana Simova and Hans Uszkoreit. 2017. Word embeddings as features for supervised coreference resolution. In *Recent Advances in Natural Language Processing*, pages 686–693, Varna, Bulgaria.
- Maksim Tkachenko, Chong Cher Chia, and Hady Lauw. 2018. Searching for the x-factor: Exploring corpus subjectivity for word embeddings. In *Annual Meeting of the Association for Computational Linguistics*, pages 1212–1221, Melbourne, Australia.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems*, pages 5998–6008.
- Shweta Yadav, Asif Ekbal, Sriparna Saha, and Pushpak Bhattacharyya. 2017. Entity extraction in biomedical corpora: An approach to evaluate word embedding features with pso based feature selection. In *European Chapter of the Association for Computational Linguistics*, pages 1159–1170.
- Zhijun Yin, Daniel Fabbri, S Trent Rosenbloom, and Bradley Malin. 2015. A scalable framework to detect personal health mentions on twitter. *Journal of medical Internet research*, 17(6).
- Su-Youn Yoon, Anastassia Loukina, Chong Min Lee, Matthew Mulholland, Xinhao Wang, and Ikkyu Choi. 2018. Word-embedding based content features for automated oral proficiency scoring. In *3rd Workshop on Semantic Deep Learning*, pages 12–22, Santa Fe, New Mexico.
- Zhijia Zhang, Guoshun Wu, and Man Lan. 2015. Ecnu: Multi-level sentiment analysis on twitter using traditional linguistic features and word embedding features. In *9th International Workshop on Semantic Evaluation*, pages 561–567, Denver, Colorado.