

UW-BHI at MEDIQA 2019: An Analysis of Representation Methods for Medical Natural Language Inference

William R. Kearns[†], Wilson Lau[†], and Jason A. Thomas[†]

[†]Department of Biomedical Informatics and Medical Education, University of Washington
850 Republican Street
Seattle, WA
{kearnsw, wlau, thomasjt}@uw.edu

Abstract

Recent advances in distributed language modeling have led to large performance increases on a variety of natural language processing (NLP) tasks. However, it is not well understood how these methods may be augmented by knowledge-based approaches. This paper compares the performance and internal representation of an Enhanced Sequential Inference Model (ESIM) between three experimental conditions based on the representation method: Bidirectional Encoder Representations from Transformers (BERT), Embeddings of Semantic Predications (ESP), or Cui2Vec. The methods were evaluated on the Medical Natural Language Inference (MedNLI) sub-task of the MEDIQA 2019 shared task. This task relied heavily on semantic understanding and thus served as a suitable evaluation set for the comparison of these representation methods.

1 Introduction

This paper describes our approach to the Natural Language Inference (NLI) subtask of the MEDIQA 2019 shared task (Ben Abacha et al., 2019). As it is not yet clear the extent to which knowledge-based embeddings may provide task-specific improvement over recent advances in contextual embeddings, we provide an analysis of the differences in performance between these two methods. Additionally, it is not yet clear from the literature the extent to which information stored in contextual embeddings overlaps with that in knowledge-based embeddings for which we provide a preliminary analysis of the attention weights of models that use these two representation methods as input. We compare BERT fine-tuned to MIMIC-III (Johnson et al., 2016) and PubMed to Embeddings of Semantic Predications (ESP) trained on SemMedDB and a baseline that

uses Cui2Vec embeddings trained on clinical and biomedical text.

Two recent advances in the unsupervised modeling of natural language, Embeddings of Language Models (ELMo) (Peters et al., 2018) and Bidirectional Encoder Representations from Transformers (BERT) (Devlin et al., 2018), have led to drastic improvements across a variety of shared tasks. Both of these methods use transfer learning, a method whereby a multi-layered language model is first trained on a large unlabeled corpus. The weights of the model are then frozen and used as input to a task specific model (Peters et al., 2018; Devlin et al., 2018; Liu et al., 2019). This method is particularly well-suited for work in the medical domain where datasets tend to be relatively small due to the high cost of expert annotation.

However, whereas clinical free-text is difficult to access and share in bulk due to privacy concerns, the biomedical domain is characterized by a significant amount of manually-curated structured knowledge bases. The BioPortal repository currently hosts 773 different biomedical ontologies comprised of over 9.4 million classes. SemMedDB is a triple store that consists of over 94 million predications extracted from PubMed by SemRep, a semantic parser for biomedical text (Rindfleisch and Fiszman, 2003; Kilocoglu et al., 2012). These available resources make a strong case for the evaluation of knowledge-based methods for the Medical Natural Language Inference (MedNLI) task (Romanov and Shivade, 2018).

2 Related Work

In this section, we provide a brief overview of methods for distributional and frame-based semantic representation of natural language. For a more detailed synthesis, we refer the reader to the

review of Vector Space Models (VSMs) by Turney and Pantel (2010).

2.1 Distributional Semantics

The distributed representation of words has a long history in computational linguistics, beginning with latent semantic indexing (LSI) (Deerwester et al., 1990; Hofmann, 1999; Kanerva et al., 2000), maximum entropy methods (Berger et al., 1996), and latent Dirichlet allocation (LDA) (Blei et al., 2003). More recently, neural network methods have been applied to model natural language (Bengio et al., 2003; Weston et al., 2008; Turian et al., 2010). These methods have been broadly applied as a method of improving supervised model performance by learning word-level features from large unlabeled datasets with more recent work using either Word2Vec (Mikolov et al., 2013; Pavlopoulos et al., 2014) or GloVe (Pennington et al., 2014) embeddings. Recent work has learned a continuous representation of Unified Medical Language System (UMLS) (Aronson, 2006) concepts by applying the Word2Vec method to a large corpus of insurance claims, clinical notes, and biomedical text where UMLS concepts were replaced with their Concept Unique Identifiers (CUIs) (Beam et al., 2018).

Models that incorporate sub-word information are particularly useful in the medical domain for representing medical terminology and out-of-vocabulary terms common in clinical notes and consumer health questions (Romanov and Shivade, 2018). Most approaches use a temporal convolution over a sliding window of characters and have been shown to improve performance on a variety of tasks (Kim et al., 2015; Zhang et al., 2015; Seo et al., 2016; Bojanowski et al., 2017).

Embeddings from Language Models (ELMo) computes word representations using a bidirectional language model that consist of a character-level embedding layer followed by a deep bidirectional long short-term memory (LSTM) network (Peters et al., 2018). Bidirectional Encoder Representations from Transformers (BERT) replaces the each forward and backward LSTMs with a single Transformer that simultaneously computes attention in both the forward and backward directions and is regarded as the current state-of-the-art method for language representation (Vaswani et al., 2017; Devlin et al., 2018). This method additionally substitutes two new unsupervised train-

ing objectives in place of the classical language models, i.e., masked language modeling (MLM) and next sentence prediction (NSP). In the case of MLM, a percentage of the words in the corpus are replaced by a [MASK] token. The task is then for the system to predict the masked token. For NSP, the task is given two sentences, s_1 and s_2 , from a document to determine whether s_2 is the next sentence following s_1 .

While ELMo has been shown to outperform GloVe and Word2Vec on consumer health question answering (Kearns and Thomas, 2018), BERT has outperformed ELMo on various clinical tasks (Si et al., 2019) and has been fine-tuned and applied to the biomedical literature and clinical notes (Alsentzer et al., 2019; Huang et al., 2019; Si et al., 2019; Lee et al., 2019). BERT supports the transfer of a pretrained general purpose language model to a task-specific application through fine-tuning. The next sentence prediction objective in the pre-training process suggests this method would be inherently suitable for NLI. In addition, BERT utilizes character-based and WordPiece tokenization (Wu et al., 2016) to learn the morphological patterns among inflections. The subword segmentation such as `##nea` in the word *dyspnea* makes it capable to understand the context of an out-of-vocabulary word making it a particularly suitable representation for clinical text.

2.2 Frame-based Semantics

FrameNet is a database of sentence-level frame-based semantics that proposes human understanding of natural language is the result of frames in which certain roles are expected to be filled (Baker et al., 1998). For example, the predicate “*replace*” has at least two such roles, the thing being replaced and the new object. A sentence such as “*The table was replaced.*” raises the question “*With what was the table replaced?*”. Frame-based semantics is a popular approach for semantic role labeling (SRL) (Swayamdipta et al., 2018), question answering (QA) (Shen and Lapata, 2007; Roberts and Demner-fushman, 2016; He, 2015; Michael et al., 2018), and dialog systems (Larsson and Traum, 2000; Gupta et al., 2018).

Vector symbolic architectures (VSA) are an approach that seeks to represent semantic predications by applying binding operators that define a directional transformation between entities (Levy and Gayler, 2008). Early approaches in-

cluded binary spatter code (BSC) for encoding structured knowledge (Kanerva, 1996, 1997) and Holographic Embeddings that used circular convolution as a binding operator to improve the scalability of this approach to large knowledge graphs (Plate, 1995). The resurgence of neural network methods has focused attention on extending these methods as there is a growing interest in leveraging continuous representations of structured knowledge to improve performance on downstream applications.

Knowledge graph embeddings (KGE) are one approach that represents entities and their relationships as continuous vectors that are learned using TransE/R (Bordes and Weston, 2009), RESCAL (Nickel et al., 2011), or Holographic Embeddings (Plate, 1995; Nickel et al., 2015). Stanovsky et al. (2017) showed that RESCAL embeddings pre-trained on DbPedia improved performance on the task of adverse drug reaction labeling over a clinical Word2Vec model. RESCAL uses tensor products whose application to representation learning dates back to Smolensky (1986; 1990) that used the inner product and has recently been applied to the bAbI dataset (Smolensky et al., 2016; Weston et al., 2016). Embeddings of Semantic Predications (ESP) are a neural-probabilistic representational approach that uses VSA binding operations to encode structured relationships (Cohen and Widdows, 2017). The Embeddings Augmented by Random Permutations (EARP) used in this paper are a modified ESP approach that applies random permutations to the entity vectors during training and were shown to improve performance on the Bigger Analogy Test Set by up to 8% against a fastText baseline (Cohen and Widdows, 2018).

3 Methods

In this section, we provide details on the three representation methods used in this study, i.e. BERT, Cui2Vec, and ESP. We continue with a description of the inference model used in each experiment to predict the label for a given hypothesis/premise pair.

3.1 Representation Layer

There are many publicly available biomedical BERT embeddings which were initialized from the original BERT Base models. BioBERT was trained on PubMed Abstracts and PubMed Central

Full-text articles (Lee et al., 2019). In this study, we applied ClinicalBERT that was initialized from BioBERT and subsequently trained on all MIMIC-III notes (Alsentzer et al., 2019).

For Cui2Vec, we used the publicly available implementation from Beam et al. (2018) that was trained on a corpus consisting of 20 million clinical notes from a research hospital, 1.7 million full-text articles from PubMed, and an insurance claims database with 60 million members.

For ESP, we used a 500-dimensional model trained over SemMedDB using the recent Embeddings Augmented by Random Permutations (EARP) approach with a 10^{-7} sampling threshold for predications and a 10^{-5} sampling threshold for concepts excluding concepts that had a frequency greater than 10^6 (Cohen and Widdows, 2018).

To apply Cui2Vec and ESP, we first processed the MedNLI dataset (Romanov and Shivade, 2018) with MetaMap to normalize entities to their concept unique identifier (CUI) in the UMLS (Aronson, 2006). MetaMap takes text as input and applies biomedical and clinical entity recognition (ER), followed by word sense disambiguation (WSD) that links entities to their normalized concept unique identifiers (CUIs). Entities that mapped to a UMLS CUI were assigned a representation in Cui2Vec and ESP. Other tokens were assigned vector representations using fastText embeddings trained on MIMIC-III data (Bojanowski et al., 2017; Romanov and Shivade, 2018).

3.2 Inference Model

For all experiments, we used the AllenNLP implementation (Gardner et al., 2018) of the Enhanced Sequential Inference Model (ESIM) architecture (Chen et al., 2017). This model encodes the premise and hypothesis using a Bidirectional LSTM (BiLSTM) where at each time step the hidden state of the LSTMs are concatenated to represent its context. Local inference between the two sentences is then achieved by aligning the relevant information between words in the premise and hypothesis. This alignment based on soft attention is implemented by the inner product between the encoded premise and encoded hypothesis to produce an attention matrix (Figure 1 and 2). These attention values are used to create a weighted representation of both sentences. An enhanced representation of the premise is created by concatenating the encoded premise, the weighted hypoth-

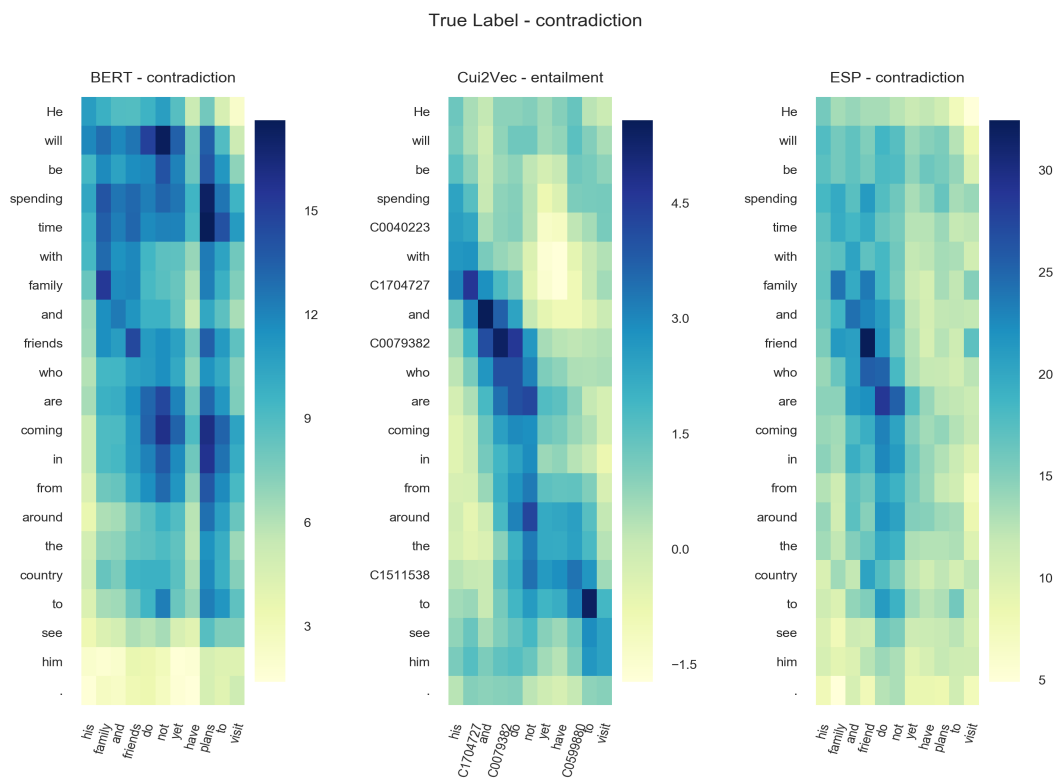


Figure 1: An example of a correct BERT prediction demonstrating its general domain coverage and contextual embedding. Premise: “He will be spending time with family and friends who are coming in from around the country to see him.” Hypothesis: “his family and friends do not yet have plans to visit.”

esis, the encoded premise minus the weighted hypothesis, and the element-wise multiplication of the encoded premise and the weighted hypothesis. The enhanced representation of the hypothesis is created similarly. This operation is expected to enhance the local inference information between elements in each sentence. This representation is then projected into the original dimension and fed into a second BiLSTM inference layer in order to capture inference composition sequentially. The resulting vector is then summarized by max and average pooling. These two pooled representations are concatenated and passed through a multi-layered perceptron followed by a sigmoid function to predict probabilities for each of the sentence labels, i.e. *entailment*, *contradiction*, and *neutral*.

4 Results

The ESIM model achieved an accuracy of 81.2%, 65.2%, and 77.8% for the MedNLI task using BERT, Cui2Vec, and ESP, respectively. Table 1 shows the number of correct predictions by each embedding type. The BERT model has the highest

accuracy on predicting *entailment* and *contradiction* labels, while the ESP model has the highest accuracy on predicting *neutral* labels. However, the difference is only significant in the case of *entailment*.

To evaluate the ability to set a predictive threshold for use in clinical applications, we sought to measure the certainty with which the model made its predictions. To achieve this goal, we used the predicted probabilities of each embedding type on their respective subset of correct predictions such that. We found the predicted probability of ESP to be much higher than the others as depicted in Figure 3. ESP’s minimum predicted probability as well as the variance of its distribution is the lowest among all embedding types.

4.1 Error Analysis

To examine the relationship between embedding prediction performance and hypothesis focus, we first annotated the test set for:

- hypothesis focus (e.g. *medications*, *procedures*, *symptoms*, etc.)

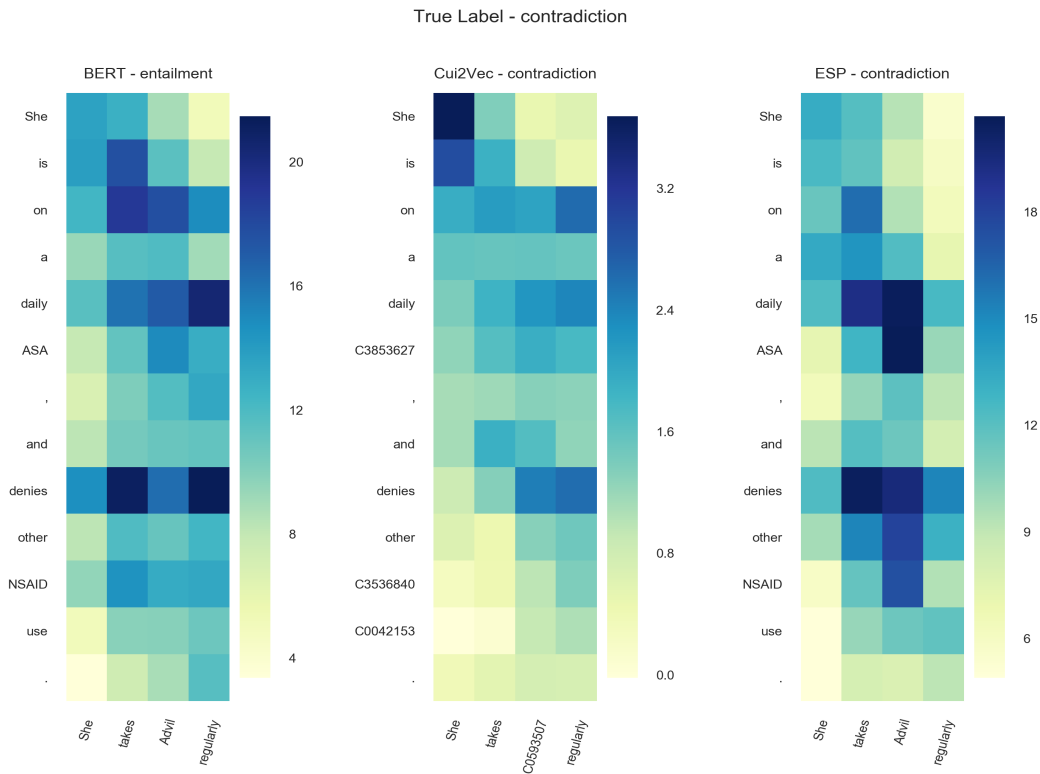


Figure 2: An example of a correct ESP prediction demonstrating its ability to associate Advil as a subclass of NSAIDs. Premise: “She is on a daily ASA, and denies other NSAID use.” Hypothesis: “She takes Advil regularly.”

Label	Embedding Type		
	BERT	Cui2Vec	ESP
Entailment	82.22% (n=111)	60.00% (n=81)	71.85% (n=97)
Contraction	88.15% (n=119)	74.81% (n=101)	87.41% (n=118)
Neutral	73.33% (n=99)	60.74% (n=82)	74.07% (n=100)

Table 1: Model accuracy for each label by embedding type.

- hypothesis tense (e.g. *past, current, future*)

4.1.1 Focus

A total of eleven, non-mutually exclusive hypothesis focus classes were arrived at by consensus of the three authors after an initial blinded round of annotation by two annotators. The remaining data was annotated by one of these annotators. We provide definitions of the classes and their overall counts in Table 2. The classes are: *State, Anatomy, Disease, Process, Temporal, Medication, Clinical Finding, Location, Lab/Imaging, Procedure, and Examination*.

We then performed Pearson’s chi-squared test with Yates’ continuity correction on 2x2 contingency tables for each embedding sentence pair prediction (correct or incorrect) with each hy-

pothesis focus (presence or absence) using the *chisq.test* function in R software and results reported in Table 3.

The only significant relationships between hypothesis focus and embedding accuracy were found between BERT and *Disease* (p-value = 0.01) and Cui2Vec and *Disease* (p-value = 0.01) through Pearson’s Chi-squared test with Yates’ continuity correction. Both embeddings achieved higher accuracy on sentence pairs with a hypothesis focus labeled *Disease* (BERT=90.4%; Cui2Vec=76.6%) than without (BERT=78.5%; Cui2Vec=61.7%).

4.1.2 Tense

Each hypothesis was annotated for tense into one of three mutually exclusive classes: *Past, Current,*

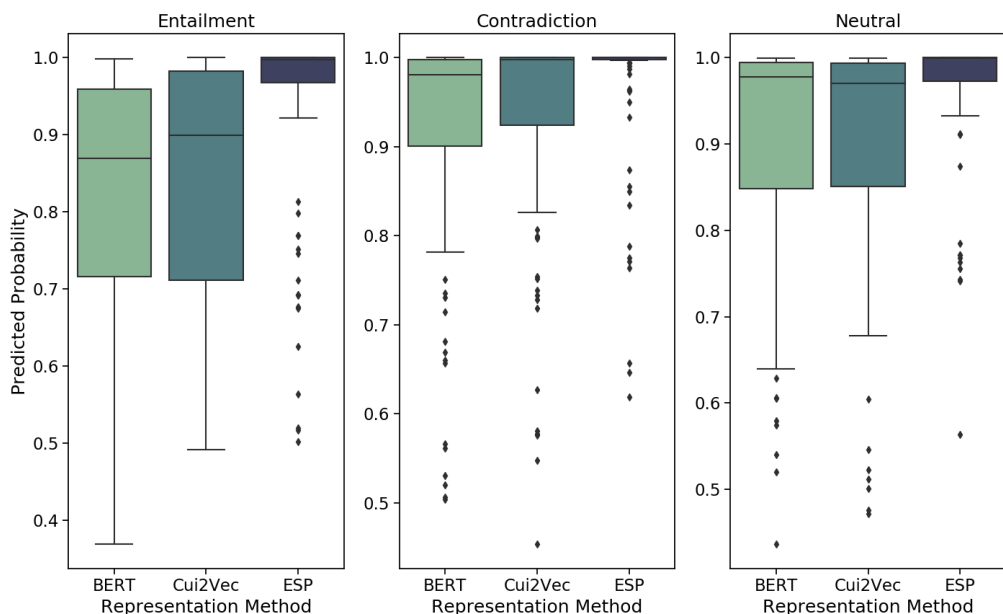


Figure 3: Distribution of predicted probability of the gold label from the subset of correct predictions for each representation method.

Hypothesis Focus	Definition	Count(%)
State	Patient state or symptoms (e.g. “...has high blood pressure...”)	251 (62.0)
Anatomy	Specific body part referenced (e.g. “... has back pain”)	115 (28.4)
Disease	Similar to state, but a defined disease (e.g. “...has Diabetes”)	95 (23.5)
Process	Events like transfers, family visiting, scheduling, or vague references to interventions (e.g. “...received medical attention”)	52 (12.8)
Temporal	Reference to time (e.g. “...initial blood pressure was low”) besides tense or history	51 (12.6)
Medication	Any reference to medication (e.g. “antibiotics”, “fluids”, “oxygen”, “IV”) including administration and patient habits	32 (7.9)
Clinical Finding	Results of an exam, lab/image, procedure, or a diagnosis	28 (6.9)
Location	Specific physical location specified (e.g. “...discharged home”)	28 (6.9)
Lab/Imaging	Laboratory tests or imaging (e.g. <i>histology</i> , <i>CBC</i> , <i>CT scan</i>)	24 (5.9)
Procedure	Physical procedure besides Lab/Image or exam (e.g. “intubation”, “surgery”, “biopsies”)	14 (3.5)
Examination	Physical examination or explicit use of the word exam(ination)	3 (0.7)

Table 2: Hypothesis foci definitions, examples, and count for all 405 hypotheses in the test set.

and *Future*. Test set hypotheses were predominantly *Current* (n=273; 67.4%) or *Past* (n=131; 32.3%) tense. Only one hypothesis (0.2%) was *Future* tense. A subset (n=22; 7.9%) of the *Current* tense hypotheses explicitly described patient history (e.g. “The patient has a history of PE”).

5 Discussion

Our preliminary analysis, identified several patterns from the attention heatmaps that differentiated the three representation methods. We describe two here and provide the entire set of attention matrices along with supplemental analysis on Github ¹.

¹<https://kearnsw.github.io/MEDIQA-2019/>

Focus	BERT			Embedding Type Cui2Vec			ESP		
	(+)	(-)	p-value	(+)	(-)	p-value	(+)	(-)	p-Value
Anatomy	93	22	1	73	42	0.74	90	25	0.99
Clinical Finding	24	4	0.71	16	12	0.47	24	4	0.42
Disease	85	9	0.01	72	22	0.01	78	16	0.21
Examination	3	0	0.93	2	1	0.58	3	0	0.82
Lab/Imaging	30	7	1	22	15	0.55	31	6	0.48
Location	21	7	0.53	14	14	0.12	19	9	0.28
Medication	27	5	0.81	24	8	0.30	28	4	0.25
Procedure	12	2	0.93	7	7	0.35	11	3	1
Process	41	11	0.78	35	17	0.85	40	12	1
State	198	53	0.16	158	93	0.27	191	60	0.36
Temporal	38	12	0.41	37	13	0.22	41	9	0.56

Table 3: Results from chi-squared (with Yates’ continuity correction) test of correct(+) and incorrect(-) predictions by embedding and hypothesis focus type.

The coverage of entities and their associations was characteristic of BERT predictions (Figure 1). BERT associated “*spending time*” with “*plans*” in addition to the lexical overlap of the word “*family*” which is attended by each experimental condition in this example. All three embeddings identified the contradictory significance of the word “*not*” in the hypothesis. However, BERT associated it with both spans “*will be*” and “*are coming*” in the premise, which led to the correct prediction. Cui2Vec over-attended the lexical match of the words “*and*”, “*to*” and “*C0079382*”, which led to the wrong prediction.

The ESP model recognized hierarchical relationships between entities, e.g. “*Advil*” and “*NSAIDs*” (Figure 2). In this example, the ESP approach attends to the daily use of “*ASA*” (acetyl-salicylic acid), i.e. aspirin, and the patient denying the use of “*other NSAIDs*”. This pattern was recognized multiple times in our analysis and provides a strong example of how continuous representations of biomedical ontologies may be used to augment contextual representations.

6 Limitations

The results presented in this paper compare a single model for each representation method fine-tuned to the development set. However, it is well known that the weights of the same model may vary slightly between training runs. Therefore, a more comprehensive approach would be to present the average attention weights across mul-

iple training runs and to examine the weights at each attention layer of the models which we leave for future work.

7 Conclusion

We have presented our analysis of representation methods on the MedNLI task as evaluated during the MEDIQA 2019 shared task. We found that BERT embeddings fine-tuned using PubMed and MIMIC-III outperformed both Cui2Vec and ESP methods. However, we found that ESP had the lowest variance and highest predictive certainty, which may be useful in determining a minimum threshold for clinical decision support systems. *Disease* was the only hypothesis focus to show a significant positive relationship with embedding prediction accuracy. This association was present for BERT and Cui2Vec embeddings - but not ESP. Overall, contradiction was the easiest label to predict for all three embeddings, which may be the result of an annotation artifact where contradiction pairs had higher lexical overlap often differentiated by explicit negation. However, overfitting on the negation can lead to lower accuracy on other entailment labels. Further, our preliminary results indicate that recognition of hierarchical relationships is characteristic of ESP suggesting that they can be used to augment contextual embeddings which, in turn, would contribute lexical coverage including sub-word information. We propose combining these methods in future work.

Acknowledgments

We would like to acknowledge Trevor Cohen for sharing the Embeddings of Semantic Predications used in this study. Author Jason A. Thomas' work was supported, in part, by the National Library of Medicine (NLM) Training Grant T15LM007442. This work was facilitated, in part, through the use of the advanced computational, storage, and networking infrastructure managed by the Research Computing Club at the University of Washington and funded by an STF award.

References

- Emily Alsentzer, John R. Murphy, Willie Boag, Weihung Weng, Di Jin, Tristan Naumann, and Matthew B. A. McDermott. 2019. [Publicly available clinical BERT embeddings](#). *CoRR*, abs/1904.03323.
- Alan R. Aronson. 2006. [Metamap: Mapping text to the umls metathesaurus](#). *Bethesda MD NLM NIH DHHS*, pages 1–26.
- Collin F. Baker, Charles J. Fillmore, and John B. Lowe. 1998. [The berkeley framenet project](#). In *Proceedings of the 36th Annual Meeting of the Association for Computational Linguistics and 17th International Conference on Computational Linguistics - Volume 1*, ACL '98/COLING '98, pages 86–90, Stroudsburg, PA, USA. Association for Computational Linguistics. <https://doi.org/10.3115/980845.980860>.
- Andrew L. Beam, Benjamin Kompa, Inbar Fried, Nathan P. Palmer, Xu Shi, Tianxi Cai, and Isaac S. Kohane. 2018. [Clinical concept embeddings learned from massive sources of medical data](#). *CoRR*, abs/1804.01486.
- Asma Ben Abacha, Chaitanya Shivade, and Dina Demner-Fushman. 2019. Overview of the mediq 2019 shared task on textual inference, question entailment and question answering. In *Proceedings of the BioNLP 2019 workshop, Florence, Italy, August 1, 2019*. Association for Computational Linguistics.
- Yoshua Bengio, Réjean Ducharme, Pascal Vincent, and Christian Janvin. 2003. [A neural probabilistic language model](#). *J. Mach. Learn. Res.*, 3:1137–1155.
- Adam L. Berger, Vincent J. Della Pietra, and Stephen A. Della Pietra. 1996. [A maximum entropy approach to natural language processing](#). *Comput. Linguist.*, 22(1):39–71.
- David M. Blei, Andrew Y. Ng, and Michael I. Jordan. 2003. [Latent dirichlet allocation](#). *J. Mach. Learn. Res.*, 3:993–1022.
- Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2017. [Enriching word vectors with subword information](#). *Transactions of the Association for Computational Linguistics*, 5:135–146. https://doi.org/10.1162/tacl_a_00051.
- Antoine Bordes and Jason Weston. 2009. [Learning Structured Embeddings of Knowledge Bases](#). *Artificial Intelligence*, (Bengio):301–306. <https://doi.org/10.1016/j.procs.2017.05.045>.
- Qian Chen, Xiaodan Zhu, Zhen-Hua Ling, Si Wei, Hui Jiang, and Diana Inkpen. 2017. [Enhanced LSTM for natural language inference](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1657–1668, Vancouver, Canada. Association for Computational Linguistics. <https://doi.org/10.18653/v1/P17-1152>.
- Trevor Cohen and Dominic Widdows. 2017. [Embedding of semantic predications](#). *Journal of Biomedical Informatics*, 68:150–166. <https://doi.org/10.1016/j.jbi.2017.03.003>.
- Trevor Cohen and Dominic Widdows. 2018. [Bringing order to neural word embeddings with embeddings augmented by random permutations \(EARP\)](#). In *Proceedings of the 22nd Conference on Computational Natural Language Learning*, pages 465–475, Brussels, Belgium. Association for Computational Linguistics.
- Scott C. Deerwester, Susan T. Dumais, Thomas K. Landauer, George W. Furnas, and Richard A. Harshman. 1990. Indexing by latent semantic analysis. *JASIS*, 41:391–407.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. [BERT: pre-training of deep bidirectional transformers for language understanding](#). *CoRR*, abs/1810.04805.
- Matt Gardner, Joel Grus, Mark Neumann, Oyvind Tafjord, Pradeep Dasigi, Nelson F. Liu, Matthew Peters, Michael Schmitz, and Luke Zettlemoyer. 2018. [AllenNLP: A deep semantic natural language processing platform](#). In *Proceedings of Workshop for NLP Open Source Software (NLP-OSS)*, pages 1–6, Melbourne, Australia. Association for Computational Linguistics.
- Sonal Gupta, Rushin Shah, Mrinal Mohit, Anuj Kumar, and Mike Lewis. 2018. [Semantic parsing for task oriented dialog using hierarchical representations](#). *CoRR*, abs/1810.07942.
- Luheng He. 2015. [Question-Answer Driven Semantic Role Labeling : Using Natural Language to Annotate Natural Language](#). *Emnlp2015*, (September):643–653. <https://doi.org/10.18653/v1/D15-1076>.

- Thomas Hofmann. 1999. **Probabilistic latent semantic indexing**. In *Proceedings of the 22Nd Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '99*, pages 50–57, New York, NY, USA. ACM. <https://doi.org/10.1145/312624.312649>.
- Kexin Huang, Jaan Altsosaar, and Rajesh Ranganath. 2019. **ClinicalBERT: Modeling Clinical Notes and Predicting Hospital Readmission**. *arXiv:1904.05342 [cs]*. ArXiv: 1904.05342.
- Alistair E W Johnson, Tom J Pollard, Lu Shen, Liwei H Lehman, Mengling Feng, Mohammad Ghassemi, Benjamin Moody, Peter Szolovits, Leo Anthony Celi, and Roger G Mark. 2016. **MIMIC-III, a freely accessible critical care database**. pages 1–9.
- Pentti Kanerva. 1996. **Binary spatter-coding of ordered k-tuples**. In *Proceedings of the 1996 International Conference on Artificial Neural Networks, ICANN 96*, pages 869–873, London, UK, UK. Springer-Verlag.
- Pentti Kanerva. 1997. **Fully distributed representation**. In *In Proceedings Real World Computing Symposium (Report TR-96001)*, pages 358–365.
- Pentti Kanerva, Jan Kristoferson, and Anders Holst. 2000. **Random indexing of text samples for latent semantic analysis**. In *In Proceedings of the 22nd Annual Conference of the Cognitive Science Society*, pages 103–6. Erlbaum.
- William R Kearns and Jason A Thomas. 2018. **Resource and response type classification for consumer health question answering**. *AMIA Annual Symposium proceedings. AMIA Symposium, 2018*:634–643.
- Halil Kilicoglu, Dongwook Shin, Marcelo Fiszman, Graciela Rosemblat, and Thomas C. Rindflesch. 2012. **SemMedDB: A PubMed-scale repository of biomedical semantic predications**. *Bioinformatics*, 28(23):3158–3160. <https://doi.org/10.1093/bioinformatics/bts591>.
- Yoon Kim, Yacine Jernite, David Sontag, and Alexander M. Rush. 2015. **Character-aware neural language models**. *CoRR*, abs/1508.06615.
- Staffan Larsson and David R. Traum. 2000. **Information state and dialogue management in the trindi dialogue move engine toolkit**. *Nat. Lang. Eng.*, 6(3-4):323–340. <https://doi.org/10.1017/S1351324900002539>.
- Jinhyuk Lee, Wonjin Yoon, Sungdong Kim, Donghyeon Kim, Sunkyu Kim, Chan Ho So, and Jaewoo Kang. 2019. **BioBERT: a pre-trained biomedical language representation model for biomedical text mining**. *arXiv:1901.08746 [cs]*.
- Simon D. Levy and Ross Gayler. 2008. **Vector symbolic architectures: A new building material for artificial general intelligence**. In *Proceedings of the 2008 Conference on Artificial General Intelligence 2008: Proceedings of the First AGI Conference*, pages 414–418, Amsterdam, The Netherlands, The Netherlands. IOS Press.
- Xiaodong Liu, Pengcheng He, Weizhu Chen, and Jianfeng Gao. 2019. **Multi-task deep neural networks for natural language understanding**. *CoRR*, abs/1901.11504.
- Julian Michael, Gabriel Stanovsky, Luheng He, Ido Dagan, and Luke Zettlemoyer. 2018. **Crowdsourcing question-answer meaning representations**. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 560–568, New Orleans, Louisiana. Association for Computational Linguistics. <https://doi.org/10.18653/v1/N18-2089>.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. **Distributed representations of words and phrases and their compositionality**. In *Proceedings of the 26th International Conference on Neural Information Processing Systems - Volume 2, NIPS'13*, pages 3111–3119, USA. Curran Associates Inc.
- Maximilian Nickel, Lorenzo Rosasco, and Tomaso A. Poggio. 2015. **Holographic embeddings of knowledge graphs**. *CoRR*, abs/1510.04935.
- Maximilian Nickel, Volker Tresp, and Hans-Peter Kriegel. 2011. **A three-way model for collective learning on multi-relational data**. In *Proceedings of the 28th International Conference on International Conference on Machine Learning, ICML'11*, pages 809–816, USA. Omnipress.
- Ioannis Pavlopoulos, Aris Kosmopoulos, and Ion Androutsopoulos. 2014. **Continuous Space Word Vectors Obtained by Applying Word2Vec to Abstracts of Biomedical Articles**. <http://bioasq.lip6.fr/info/BioASQword2vec/>.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. **Glove: Global vectors for word representation**. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, Doha, Qatar. Association for Computational Linguistics. <https://doi.org/10.3115/v1/D14-1162>.
- Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. **Deep contextualized word representations**. *CoRR*, abs/1802.05365.
- T. A. Plate. 1995. **Holographic reduced representations**. *IEEE Transactions on Neural Networks*, 6(3):623–641. <https://doi.org/10.1109/72.377968>.

- Thomas C Rindflesch and Marcelo Fiszman. 2003. The interaction of domain knowledge and linguistic structure in natural language processing: interpreting hypernymic propositions in biomedical text. *Journal of Biomedical Informatics*, 36(6):462–477. <https://doi.org/10.1016/j.jbi.2003.11.003>.
- Kirk Roberts and Dina Demner-fushman. 2016. [Annotating Logical Forms for EHR Questions](#). In *Proceedings of the 10th International Conference on Language Resources and Evaluation*, Section 3, pages 3772–3778.
- Alexey Romanov and Chaitanya Shivade. 2018. [Lessons from natural language inference in the clinical domain](#). *CoRR*, abs/1808.06752.
- Min Joon Seo, Aniruddha Kembhavi, Ali Farhadi, and Hannaneh Hajishirzi. 2016. [Bidirectional attention flow for machine comprehension](#). *CoRR*, abs/1611.01603.
- Dan Shen and Mirella Lapata. 2007. [Using semantic roles to improve question answering](#). In *Proceedings of the 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP-CoNLL)*, page 12–21.
- Yuqi Si, Jingqi Wang, Hua Xu, and Kirk Roberts. 2019. [Enhancing Clinical Concept Extraction with Contextual Embedding](#). *arXiv:1902.08691 [cs]*. ArXiv: 1902.08691.
- P. Smolensky. 1986. [Parallel distributed processing: Explorations in the microstructure of cognition, vol. 1](#). chapter Information Processing in Dynamical Systems: Foundations of Harmony Theory, pages 194–281. MIT Press, Cambridge, MA, USA.
- P. Smolensky. 1990. [Tensor product variable binding and the representation of symbolic structures in connectionist systems](#). *Artif. Intell.*, 46(1-2):159–216. [https://doi.org/10.1016/0004-3702\(90\)90007-M](https://doi.org/10.1016/0004-3702(90)90007-M).
- Paul Smolensky, Moontae Lee, Xiaodong He, Wen tau Yih, Jianfeng Gao, and Li Deng. 2016. Basic reasoning with tensor product representations. *CoRR*, abs/1601.02745.
- Gabriel Stanovsky, Daniel Gruhl, and Pablo Mendes. 2017. [Recognizing mentions of adverse drug reaction in social media using knowledge-infused recurrent models](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, pages 142–151, Valencia, Spain. Association for Computational Linguistics.
- Swabha Swayamdipta, Sam Thomson, Kenton Lee, Luke Zettlemoyer, Chris Dyer, and Noah A. Smith. 2018. [Syntactic scaffolds for semantic structures](#). *CoRR*, abs/1808.10485.
- Joseph Turian, Lev-Arie Ratinov, and Yoshua Bengio. 2010. [Word representations: A simple and general method for semi-supervised learning](#). In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*, pages 384–394, Uppsala, Sweden. Association for Computational Linguistics.
- Peter D. Turney and Patrick Pantel. 2010. [From frequency to meaning: Vector space models of semantics](#). *Journal of Artificial Intelligence Research*, 37:141–188. <https://doi.org/10.1613/jair.2934>.
- 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 I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 5998–6008. Curran Associates, Inc.
- Jason Weston, Antoine Bordes, Sumit Chopra, and Tomas Mikolov. 2016. [Towards ai-complete question answering: A set of prerequisite toy tasks](#). *CoRR*, abs/1502.05698.
- Jason Weston, Frédéric Ratle, and Ronan Collobert. 2008. [Deep learning via semi-supervised embedding](#). In *Proceedings of the 25th International Conference on Machine Learning, ICML '08*, pages 1168–1175, New York, NY, USA. ACM. <https://doi.org/10.1145/1390156.1390303>.
- Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V. Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, Jeff Klingner, Apurva Shah, Melvin Johnson, Xiaobing Liu, Łukasz Kaiser, Stephan Gouws, Yoshikiyo Kato, Taku Kudo, Hideto Kazawa, Keith Stevens, George Kurian, Nishant Patil, Wei Wang, Cliff Young, Jason Smith, Jason Riesa, Alex Rudnick, Oriol Vinyals, Greg Corrado, Macduff Hughes, and Jeffrey Dean. 2016. [Google’s Neural Machine Translation System: Bridging the Gap between Human and Machine Translation](#). *arXiv:1609.08144 [cs]*. ArXiv: 1609.08144.
- Xiang Zhang, Junbo Jake Zhao, and Yann LeCun. 2015. [Character-level convolutional networks for text classification](#). *CoRR*, abs/1509.01626.