

Towards User Friendly Medication Mapping Using Entity-Boosted Twin Modeling

Shaoqing Yuan¹ Parminder Bhatia² Busra Celikkaya² Haiyang Liu¹ Kyunghwan Choi¹

Abstract

Recent advancements in medical entity linking have been applied in the area of scientific literature and social media data. However, with the adoption of telemedicine and conversational agents such as Alexa in healthcare settings, medical name inference has become an important task. Medication name inference is the task of mapping user friendly medication names from a free-form text to a concept in a normalized medication list. This is challenging due to the differences in the use of medical terminology from health care professionals and user conversations coming from the lay public. We begin with mapping descriptive medication phrases (DMP) to standard medication names (SMN). Given the prescriptions of each patient, we want to provide them with the flexibility of referring to the medication in their preferred ways. We approach this as a ranking problem which maps SMN to DMP by ordering the list of medications in the patient’s prescription list obtained from pharmacies. Furthermore, we leveraged the output of intermediate layers and performed medication clustering. We present the Medication Inference Model (MIM) achieving state-of-the-art results. By incorporating medical entities based attention, we have obtained further improvement for ranking models.

1. Introduction

Medication names are extremely hard to pronounce for patients without a proper medical background. Thus, when interacting with Alexa on medication names, patients without this background may have many different ways to refer to a medication (e.g. Bumetanide can be referred to as Bumetanide (generic name), Bumex (brand name), high blood pressure pill (disease name)). On the other hand, pa-

¹Amazon Alexa, Seattle, Washington, USA ²Amazon AWS, Seattle, Washington, USA. Correspondence to: Shaoqing Yuan <shaoqiny@amazon.com>.

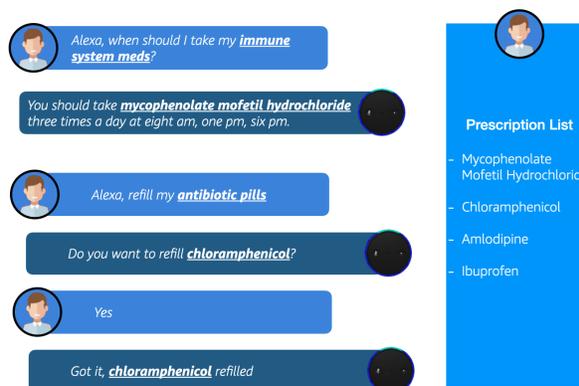


Figure 1. Different ways of how users interact with conversational agents for medical queries

tients with medical knowledge may use abbreviations or specialized ways to refer to medication names. For example, patients may use “immune meds” to refer to “mycophenolate mofetil hydrochloride” in their prescription list.

In this paper, we describe a new problem about finding the generic medication name (SMN: standardized medication name) based on a patient’s description (DMP: descriptive medication phrase) from a list of medications the patient is consuming. According to our internal user research, in the United States, patients with chronic diseases usually take around four to five medications daily. This problem is different from medical concept normalization (Zhu et al., 2019) which tries to map a health-related entity mention in a free-form text to a concept in a controlled vocabulary (Miftahutdinov & Tutubalina, 2019) which is a generic concept list rather than a patient specific prescription list and is generally much longer.

We structure this as a ranking problem. Here we rank all medications a patient is consuming based on the relationship with the patient’s description and the one ranked highest will be the inference result. We present a hard attention based entity boosted CNN architecture achieving 4% over earlier ranking methods.

Furthermore, the mapping between SMN and DMP contains the patient’s understanding of the medications, especially

from the usage perspective of the medications. Using latent output from our model, we build a medication clustering system which groups together medications with similar effects and disease treatments. The output is designed to aid physicians to consider other medications as a substitution for decreasing cost as well as helping patients distinguish medications that are similar in their impression but should, in reality, be used in different conditions. Moreover, with clustering patients will have an intuitive understanding of the relationship of the medications they are consuming. Our contributions are as follows:

- We present a medical entity boosted architecture, Medication Inference Model (MIM) achieving a 7% – 9% improvement over strong BERT baselines.
- We benchmark against state-of-the-art ranking architectures, demonstrating robustness of our work.
- We present medication clustering results which group together medications with similar effects and treat the similar diseases.

2. Task Definition

Each example is represented as a tuple $(Q, P_1, P_2, \dots, P_n, Y)$, where $Q = (q_1, q_2, \dots, q_{l_q})$ is a DMP, with a length l_q , $P_i = (p_1, p_2, \dots, p_{l_{p_i}})$ is a SMN, with a length l_{p_i} , and $Y = (y_1, y_2, \dots, y_n)$ is the label representing the relationship between Q and P_1, P_2, \dots, P_n . Y and P have the same length. $y_i = 1$ if P_i is the generic medication name that Q is referring to, 0 otherwise.

Y	$Q = \{q_1(\text{high}), q_2(\text{blood}), q_3(\text{pressure})\}$
$y_1 = 0$	$P_1 = \{p_1(\text{morphine}), p_2(\text{suppository})\}$
$y_2 = 1$	$P_2 = \{p_1(\text{hydrochlorothiazide})\}$
$y_3 = 0$	$P_3 = \{\dots\}$
\dots	\dots

It is possible that among P_1, P_2, \dots, P_n , more than one medications may be referred to by Q . Thus, $\sum_{i=0}^n y_i = m$ where m is the number of medications in P_1, P_2, \dots, P_n that could be referred by Q . Ideally, we should make it possible that for the estimated \hat{Y} , $\sum_{i=0}^n \hat{y}_i > 1$. In this paper, however, we assume $\sum_{i=0}^n y_i = 1$.

The clustering task is defined as grouping medications across the prescriptions of different patients. i.e., we assign each medication in (P_1, P_2, \dots, P_N) to a group according to the DMP Q associated with them.

3. Method

Instead of comparing n medications in each sample, we begin with two. Each sample consists of (Q, P_1, P_2, Y) and the model must distinguish which SMN in P_1, P_2 the patient is referring to by Q . To simplify, we assume only one of

(P_1, P_2) could be referred by Q . In practice, when there are n (where $n > 2$) medications in a patient’s prescription we run the model on all the combinations of the medications and rank them accordingly.

3.1. Entity Boosted Two-Tower Neural Network

Motivated by the facial recognition problem, where the models evaluate the similarity between images of faces (Chopra et al., 2005), we apply a Two-tower neural network to our problem. We regard the DMP as the query of each medicine and the SMNs as the medication candidates. The purpose is to match the correct SMNs to the DMP provided.

Descriptions can often be verbose and can contain a large amount of noise. To improve the robustness and reduce noise, we have incorporated medical entity based hard attention (Luong et al., 2015) using Amazon Web Services Comprehend Medical (CM) (Bhatia et al., 2019) which is a natural language processing service to perform entity and relation extraction.

For each instance of data, we use the generic name as the SMN in our model, and generate DMPs from free-form text data that describe the usage of each medication in patient friendly terms. To reduce the noise, we feed the description to CM to extract entities. CM is able to extract relevant medical information from unstructured text and classify extracted entities into five categories and 28 types. In this work we use the entities marked with types “dx name” (diagnostic indicator), “treatment name”, “system organ site”, “swap”, “generic name”, “procedure name”, “brand name”, “test name” as DMPs.

3.1.1. MEDICATION INFERENCE MODEL

Figure 2 outlines the Medication Inference Model (MIM) network. In this model, there are two different sets of unshared embedding weights, where one is used to embed SMNs and the other DMP. We use convolutional (CNN) layers (Kim, 2014) followed by pooling on top of the embedding layer to get a vector representation of SMNs and DNP separately. We use cosine distance to measure the separation between the two SMNs and DMP.

3.1.2. BERT-BASED MODEL

It is natural to leverage models which prove to be successful in solving question-answering problems to process our task. Here, DMP are regarded as the patient’s query, with the n SMNs as the answer candidates.

We concatenate the DMP output with different SMNs separately and combine them into a BERT (Devlin et al., 2018) based multi-choice model (Guo et al., 2019). The vector representations of the [CLS] tokens are used to represent the combinations of DMP with each SMN, and are fed into a

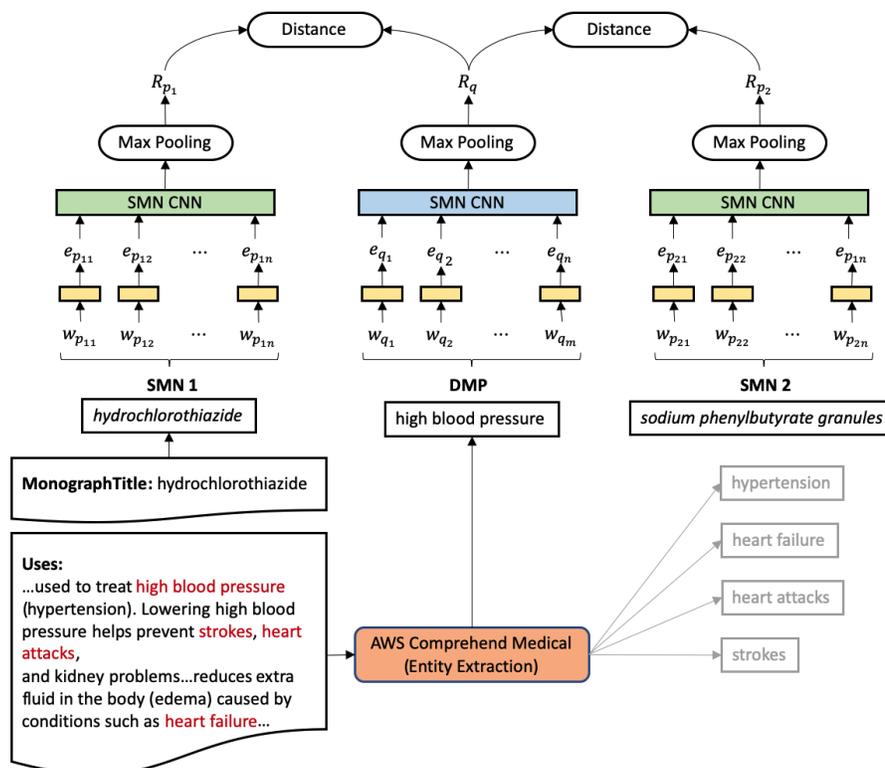


Figure 2. Medication Inference Model structure

fully connected layer. Finally we use hinge loss for ranking to compare the scale values with ground truth.

4. Experiments

4.1. Dataset and Evaluation Metrics

Synthetic Data Set The training and test dataset are generated from 2,683 medication descriptions from the FDB¹ PEM (patient education module) dataset. FDB stands for First DataBank which is a known drug database and medical device database provider. Each PEM file contains a patient facing medication description including medication generic name, uses, warnings side effects etc.².

The SMNs are collected from the generic name section and DMPs are generated from the "USES" section of the PEM files using CM as described in Section 3.1.1. To evaluate the effect of the entity extraction component, we generate another DMP set by randomly drawing n-grams (where $n = [1, 2, 3, 4, 5]$) from the "USES" section of the PEM files as a replacement for CM.

¹<https://www.fdbhealth.com/>

²CommonNames, Warning, Uses, HowToUse, SideEffects, Precautions, DrugInteractions, Overdose, Notes, MissedDose, Storage, MedicAlert

Next, we use the SMNs and DMPs collected to generate our training and test sets. we generate each instance starting with a DMP according to following steps.

1. For each DMP, we generate a positive SMNs set which consist of SMNs extracted from the same PEM file where the DMP is extracted from. It is possible that one DMP may have multiple positive SMNs if the DMP is a very general phrase. For example, the DMP "high blood pressure" may have multiple SMNs since many medications can be used to treat hypertension.
2. For each DMP, we also generate a negative SMN set. The negative SMNs are all the medications covered by the PEM files excluding the positive SMNs identified above and should follow the constraint that the entities extracted from the "USES" section of the negative SMNs' PEM files should have no overlap with that of the DMPs.
3. Each instance in the training and validation data set consists of 1 positive SMN and $n - 1$ negative SMNs randomly selected from the SMN sets described above. The label of each instance is indicating which SMN is positive.

For the training and validation splits, n is set to 2 in step 3

Model	Number of candidates							
	Baseline				Entity-based Attention			
	2	3	4	5	2	3	4	5
NormalBert	69.7	54.7	46.6	42.7	79.7	72.7	62.4	62.5
BioBert	73.4	59.8	51.9	46.4	81.5	74.8	66.2	66.4
ClinicalBert	73.5	57.8	51.9	47.4	79.5	74.4	64.8	62.7
ARC-1	65.2	46.2	37.5	33.4	76.8	66.8	58.6	59.4
ARC-2	64.0	48.1	39.1	36.5	75.7	65.7	62.1	57.0
ConvKNRM	65.4	52.0	42.6	38.8	76.7	66.6	60.7	57.3
MatchLSTM	64.9	50.8	40.9	35.4	82.4	70.9	62.6	58.3
MatchPyramid	59.5	45.0	36.5	33.5	74.0	58.0	58.0	54.1
MIM	73.9	57.9	48.9	49.7	87.7	80.8	78.9	76.9

Table 1. Synthetic test set: For each model we report top1 accuracy including and excluding entity-based attention.

Model	Upper limit for number of candidates				
	2	3	4	5	10
BioBert	83.10	73.50	67.00	61.8	52.70
MIM	83.50	74.50	69.50	63.60	53.80

Table 2. Real test set: For each model we report top1 accuracy.

above, which means there are two SMNs in each instance. The training and validation data set contains 680K instances and 70% of them are used for training and 30% for validation and testing.

For testing purposes, we generated four synthetic test sets with n in step 3 set to 2, 3, 4, 5 separately to simulate the real situations where patients with chronic disease in the U.S. usually have four to five medications in their prescription list at a time.

Real Data Set The real data set is generated based on 251 prescriptions collected from the i2b2 data set³ which contains the de-identified patient discharge summaries. Internal human annotators generate DMPs for each medication in the prescriptions. It is observed that in a real prescription, multiple medications may serve the same purpose and a general DMP could be used to refer to multiple medications in a prescription. In our current experiment, we assume the ground truth of each DMP is only the medication used to generate the DMP in a prescription. In this way, we will get the lower bound of the performance of the models. For testing purpose, we limited the number of medications in each prescription to be 10, 5, 4, 3, 2 respectively. For the test set with 10 as max number of medications, we go through all the 251 prescriptions and only select the prescriptions that has less than 10 medications into our test set. We randomly truncate the prescriptions in the 10 medication test set to 5, 4, 3, 2 medications as the other test sets. Further more, in order to evaluate the situations where one DMP may refer to multiple SMNs in a prescription, the annotators are currently working on labeling all the SMNs that a DMP could refer to in a prescription and if the model outputs one of the medications in the ground truth SMNs, the test sample will

³<https://portal.dbmi.hms.harvard.edu/projects/n2c2-nlp/>

be marked as success in future experiments.

We report accuracy as the main evaluation metric, i.e., the correctness of selecting the positive SMN from n SMNs. When evaluating on the test data, the model goes through all pairwise combinations of the SMNs and ranks all the SMNs accordingly.

4.2. Experimental Details

For the CNN-based model, we test multiple word embedding models including 200-dimensional BioWordVec (Zhang et al., 2019; Chen et al., 2018) and 300-dimensional FastText word embeddings (Bojanowski et al., 2017) trained with 3,466 articles from the Mayo Clinic. The two dimensional CNN layer consisted of 200 filters with window size 2, strip as 1 and no regularization. Batch size is set to 150 and we observed model convergence after six epochs. For the pre-trained language model, we leverage Clinical BERT (Alsentzer et al., 2019), BioBERT (Lee et al., 2019), and original BERT models (Devlin et al., 2018). We used the default settings for all BERT models as provided by Devlin et al. (2018). Batch size is set to 32, learning rate is set to 5×10^{-5} and dropout rate is set to 0.2. We observed the model converged after 10 epochs. We trained and evaluated all the models using a Tesla V100 GPU.

4.2.1. BASELINES

When evaluating the performance of our model, we compare the medication name inference performance with baseline models listed below.

- **ARC-I (Hu et al., 2014)**: ARC-I finds the representation of each sentence with CNN layers, and then compares the representation for the two sentences with a multi-layer perceptron (MLP).
- **ARC-II (Hu et al., 2014)**: ARC-II improves based on ARC-I by calculating the interaction features between sentences with CNN.
- **ConvKNRM (Dai et al., 2018)**: Conv-KNRM uses

CNN to represent n-grams of various lengths and soft matches them in a unified embedding space. The n-gram soft matches are then utilized by the kernel pooling and a fully connected layer to generate the final ranking score.

- **MatchLSTM (Wang & Jiang, 2016):** The matchLSTM sequentially aggregates the matching of the attention-weighted question to each token of the answer and uses the aggregated matching result to make a final prediction.
- **MatchPyramid (Pang et al., 2016):** MatchPyramid generates a matching matrix which represents the similarity between mention and candidate and then apply CNN layers on top of the matrix followed by a MLP layer to calculate the similarity score.

5. Results and Discussion

Table 1 provides the accuracy results for each model we experimented with on the synthetic test data set. Number of candidates represent test data sets with 2, 3, 4, 5 medications in each test instance, as described in Section 4.1. We report test results for each model with and without AWS Comprehend Medical as “Entity-based Attention” and “Baseline” columns.

Table 1 demonstrates the robustness of MIM. We observe that MIM and BERT based models outperform current state-of-the-art models such as ARC-I and MatchPyramid across a different number of candidates. Table 2 further compares the performance of two best performing models on real test set.

MIM outperforms BERT based models with a 7-9% improvement in accuracy. We believe the major reason for this is that MIM, by encoding SMN and DMP separately, is able to encode the representation in a more robust way in comparison to BERT based models which concatenate the representations together using a special separator token. We also observe performance variation of BERT models based on their pre-training. We found that domain specific pre-training helps, giving 2-3% improvement when compared to the baseline BERT.

We observe the entity-boosted description gives robust results across all the model settings achieving significant improvement in accuracy over non-entity based models. This alleviates the problem of noise in the lengthy descriptions.

Furthermore we see that our MIM model, with a relatively simpler CNN encoder as well as separate encoders for SMN and DMP, has the distinct advantage of generating inference results with low latency. This is ideal for real-time industrial settings. According to our experiments, the average latency

for the MIM model for five medications is 10ms, while, compared against BERT at 89ms.

5.1. Medication Clustering Result

We apply k -nearest neighbor (KNN) clustering based on the CNN max pooling output from the Two-tower neural network. The generic names of 2,683 medications are represented by vectors of 200 dimension. The number of clusters of KNN is determined by a Silhouette analysis (Rousseeuw, 1987) with result given in Figure 3. The Silhouette analysis shows that clustering performs better when number of classes is 31.

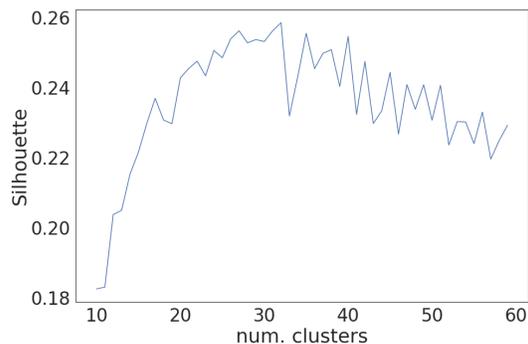


Figure 3. Average Silhouette value for different number of clusters

Figure 4 shows the t-SNE visualization of the result where 2,683 medications are grouped into 31 clusters. The figure illustrates that medications with same effects, treating same disease or having similar drug types are mapped close to each other.

Table 3 shows three examples of the clustering result. The nearby SMNs are sampled from the same cluster which example SMNs belong to and ranked according to their distance to the given SMN. The nature of the problem enables and requires the inference model to group SMNs based on multiple dimensions. We list examples from three dimensions and it is very natural for users to refer to their medications by diagnosis, disease symptoms and drug type. For example, coughing is a common symptom for multiple diseases including the common cold, pulmonary diseases such as pneumonia, and even from seasonal allergies. In column (b), the model is able to cluster medications that could relieve cough symptom of different underlying causes, for example, promethazine and antihistamine are used to treat allergies whereas zanamivir is used to treat and prevent flu.

6. Related Work

Earlier work on medical concept normalization (Zhu et al., 2019) relied on lexicon based string matching and dictio-

nary lookup to map limited number of variations of text to a pre-defined medical vocabulary (Aronson, 2001; Brennan & Aronson, 2003). (Leaman et al., 2013) introduced DNrom as the first pairwise learning ranking model that compares associations between mentions and entities of various disease. (Limsopatham & Collier, 2016; Lee et al., 2017) then further leveraged deep learning models, convolutional neural network (Limsopatham & Collier, 2016) and recurrent neural network models (Belousov et al., 2017) trained on large corpus of medical articles etc.. Currently, researchers enhanced the deep learning based model with different model structure to incorporate context information, better process out of vocabulary (OOV) words and take advantages of interaction features from different semantic levels (Luo et al., 2018; Miftahutdinov & Tutubalina, 2019; Niu et al., 2019).

With the success of deep learning, many neural network based models have been proposed for semantic matching, and document ranking. Models such as ARC-I (Hu et al., 2014) first compute the representation of the two sentences, and then compute their relevance. Semantic/text matching techniques fits well to solve the medical concept normalization problems when the number of candidates is limited. As listed in (Guo et al., 2019), recently researchers have focused on developing deep learning models to solve document retrieval, question answering, conversational response ranking, and paraphrase identification (Guo et al., 2019) problems and introduced state-of-the-art models such as ARC-I (Hu et al., 2014), ARC-II (Hu et al., 2014), ConvKNRM (Dai et al., 2018), MatchLSTM (Wang & Jiang, 2016), MatchPyramid (Pang et al., 2016), Bert (Devlin et al., 2018).

In recent years, natural language processing (NLP) techniques have demonstrated increasing effectiveness in clinical text mining (Bhatia et al., 2019) (Bhatia et al., 2019). Electronic health record (EHR) narratives, e.g., discharge summaries and progress notes contain a wealth of medically relevant information such as diagnosis information and adverse drug events. Automatic extraction of such information and representation of clinical knowledge in standardized formats (Singh & Bhatia, 2019) could be employed for a variety of purposes such as clinical event surveillance, decision support (Jin et al., 2018), pharmacovigilance, and drug efficacy studies.

This paper describes a problem that is a combination of the medical concept normalization and semantic matching problem using medical entity based hard attention. The nature of the problem presented in this paper requires the solution be able to extract informations from short phrases with limited context information.

7. Conclusion and Future Work

In this paper, we introduce a new problem common in the development of medication voice interaction products. We evaluate the accuracy of different solutions and show that our entity boosted MIM outperform baseline models. The specialty of this problem is that the context information is very limited when compared against other NLP tasks and the short length of the phrases prevent us from leveraging other advanced techniques that rely on words relationship in a phrase. The evaluation result also show that the problem prefers simple model structure. Since the phrases structure is very simple, the quality of word embeddings is more important in this problem and keeping the embedding weight unchanged is important when the training data is not sufficient enough to enhance the relationship between words either due to the nature of the data or small sample sizes.

We also observe the discrepancy between synthetic collected datasets from real patients. For example, the combinations of the medicine on synthetic prescriptions may not be valid from a practitioner's or patient's perspective. We plan to further validate our model on real patient data to increase practicality. Finally on top of comparing and evaluating on two medications samples, we plan to experiment with more medications in each sample in training to closer mimic real world scenarios.

References

- Alsentzer, E., Murphy, J., Boag, W., Weng, W.-H., Jin, D., Naumann, T., and McDermott, M. Publicly available clinical BERT embeddings. In *Proceedings of the 2nd Clinical Natural Language Processing Workshop*, pp. 72–78, Minneapolis, Minnesota, USA, June 2019. Association for Computational Linguistics. doi: 10.18653/v1/W19-1909. URL <https://www.aclweb.org/anthology/W19-1909>.
- Aronson, A. R. Effective mapping of biomedical text to the umls metathesaurus: the metamap program. In *Proceedings of the AMIA Symposium*, pp. 17. American Medical Informatics Association, 2001.
- Belousov, M., Dixon, W., and Nenadic, G. Using an ensemble of generalised linear and deep learning models in the smm4h 2017 medical concept normalisation task. In *Proceedings of the Second Workshop on Social Media Mining for Health Applications (SMM4H)*. Health Language Processing Laboratory, 2017.
- Bhatia, P., Arumae, K., and Celikkaya, E. B. Dynamic transfer learning for named entity recognition. In *International Workshop on Health Intelligence*, pp. 69–81. Springer, 2019.
- Bhatia, P., Celikkaya, B., Khalilia, M., and Senthivel, S.

- Comprehend medical: A named entity recognition and relationship extraction web service. In *2019 18th IEEE International Conference On Machine Learning And Applications (ICMLA)*, pp. 1844–1851, 2019.
- Bojanowski, P., Grave, E., Joulin, A., and Mikolov, T. Enriching word vectors with subword information. *Transactions of the Association for Computational Linguistics*, 5: 135–146, 2017. ISSN 2307-387X.
- Brennan, P. F. and Aronson, A. R. Towards linking patients and clinical information: detecting umls concepts in e-mail. *Journal of biomedical informatics*, 36(4-5):334–341, 2003.
- Chen, Q., Peng, Y., and Lu, Z. Biosentvec: creating sentence embeddings for biomedical texts. *CoRR*, abs/1810.09302, 2018. URL <http://arxiv.org/abs/1810.09302>.
- Chopra, S., Hadsell, R., and LeCun, Y. Learning a similarity metric discriminatively, with application to face verification. In *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)*, volume 1, pp. 539–546. IEEE, 2005.
- Dai, Z., Xiong, C., Callan, J., and Liu, Z. Convolutional neural networks for soft-matching n-grams in ad-hoc search. In *Proceedings of the eleventh ACM international conference on web search and data mining*, pp. 126–134, 2018.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- Guo, J., Fan, Y., Ji, X., and Cheng, X. Matchzoo: A learning, practicing, and developing system for neural text matching. In *Proceedings of the 42Nd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR'19*, pp. 1297–1300, New York, NY, USA, 2019. ACM. ISBN 978-1-4503-6172-9. doi: 10.1145/3331184.3331403. URL <http://doi.acm.org/10.1145/3331184.3331403>.
- Hu, B., Lu, Z., Li, H., and Chen, Q. Convolutional neural network architectures for matching natural language sentences. In *Advances in neural information processing systems*, pp. 2042–2050, 2014.
- Jin, M., Bahadori, M. T., Colak, A., Bhatia, P., Celikkaya, B., Bhakta, R., Senthivel, S., Khalilia, M., Navarro, D., Zhang, B., et al. Improving hospital mortality prediction with medical named entities and multimodal learning. *arXiv preprint arXiv:1811.12276*, 2018.
- Kim, Y. Convolutional neural networks for sentence classification. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 1746–1751, Doha, Qatar, October 2014. Association for Computational Linguistics. doi: 10.3115/v1/D14-1181. URL <https://www.aclweb.org/anthology/D14-1181>.
- Leaman, R., Islamaj Doğan, R., and Lu, Z. Dnorm: disease name normalization with pairwise learning to rank. *Bioinformatics*, 29(22):2909–2917, 2013.
- Lee, J., Yoon, W., Kim, S., Kim, D., Kim, S., So, C. H., and Kang, J. BioBERT: a pre-trained biomedical language representation model for biomedical text mining. *Bioinformatics*, 09 2019. ISSN 1367-4803. doi: 10.1093/bioinformatics/btz682. URL <https://doi.org/10.1093/bioinformatics/btz682>.
- Lee, K., Hasan, S. A., Farri, O., Choudhary, A., and Agrawal, A. Medical concept normalization for online user-generated texts. In *2017 IEEE International Conference on Healthcare Informatics (ICHI)*, pp. 462–469. IEEE, 2017.
- Limsopatham, N. and Collier, N. Normalising medical concepts in social media texts by learning semantic representation. Association for Computational Linguistics, 2016.
- Luo, Y., Song, G., Li, P., and Qi, Z. Multi-task medical concept normalization using multi-view convolutional neural network. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- Luong, M.-T., Pham, H., and Manning, C. D. Effective approaches to attention-based neural machine translation. *arXiv preprint arXiv:1508.04025*, 2015.
- Miftahutdinov, Z. and Tutubalina, E. Deep neural models for medical concept normalization in user-generated texts. *arXiv preprint arXiv:1907.07972*, 2019.
- Niu, J., Yang, Y., Zhang, S., Sun, Z., and Zhang, W. Multi-task character-level attentional networks for medical concept normalization. *Neural Processing Letters*, 49(3): 1239–1256, 2019.
- Pang, L., Lan, Y., Guo, J., Xu, J., Wan, S., and Cheng, X. Text matching as image recognition. In *Thirtieth AAAI Conference on Artificial Intelligence*, 2016.
- Rousseeuw, P. J. Silhouettes: a graphical aid to the interpretation and validation of cluster analysis. *Journal of computational and applied mathematics*, 20:53–65, 1987.
- Singh, G. and Bhatia, P. Relation extraction using explicit context conditioning. *arXiv preprint arXiv:1902.09271*, 2019.

Wang, S. and Jiang, J. Machine comprehension using match-lstm and answer pointer. *arXiv preprint arXiv:1608.07905*, 2016.

Zhang, Y., Chen, Q., Yang, Z., Lin, H., and Lu, Z. Biowordvec, improving biomedical word embeddings with subword information and mesh. *Scientific data*, 6(1):52, 2019.

Zhu, M., Celikkaya, B., Bhatia, P., and Reddy, C. K. Latte: Latent type modeling for biomedical entity linking. *arXiv preprint arXiv:1911.09787*, 2019.