

# BioNLP-OST 2019 RDoC Tasks: Multi-grain Neural Relevance Ranking Using Topics and Attention Based Query-Document-Sentence Interactions

\*Yatin Chaudhary<sup>1,2</sup>, \*Pankaj Gupta<sup>1,2</sup>, Hinrich Schütze<sup>2</sup>

<sup>1</sup>Corporate Technology, Machine-Intelligence (MIC-DE), Siemens AG Munich, Germany

<sup>2</sup>CIS, University of Munich (LMU) Munich, Germany

{yatin.chaudhary, pankaj.gupta}@siemens.com

## Abstract

This paper presents our system details and results of participation in the RDoC Tasks of BioNLP-OST 2019. Research Domain Criteria (RDoC) construct is a multi-dimensional and broad framework to describe mental health disorders by combining knowledge from *genomics* to *behaviour*. Non-availability of RDoC labelled dataset and tedious labelling process hinders the use of RDoC framework to reach its full potential in Biomedical research community and Healthcare industry. Therefore, *Task-1* aims at retrieval and ranking of PubMed abstracts relevant to a given RDoC construct and *Task-2* aims at extraction of the most relevant sentence from a given PubMed abstract. We investigate (1) attention based supervised neural topic model and SVM for retrieval and ranking of PubMed abstracts and, further utilize BM25 and other relevance measures for re-ranking, (2) supervised and unsupervised sentence ranking models utilizing *multi-view* representations comprising of *query-aware attention-based sentence representation (QAR)*, bag-of-words (BoW) and TF-IDF. Our best systems achieved 1st rank and scored 0.86 mAP and 0.58 macro average accuracy in Task-1 and Task-2 respectively.

## 1 Introduction

The scientific research output of the biomedical community is becoming more sub-domain specialized and increasing at a faster pace. Most of the biomedical domain knowledge is in the form of unstructured text data. Natural Language Processing (NLP) techniques such as relation extraction and information retrieval have enabled us to effectively mine relevant information from a large corpus. These techniques have significantly reduced the time and effort required for knowledge min-

ing and information extraction from past scientific studies and electronic health reports (EHR).

Information Retrieval (IR) is the process of retrieving relevant information from an unstructured text corpus, which satisfies a given query/requirement, for example Google search, email search, database search etc. This is generally achieved by converting the query and the document collection into an external representation which by preserving the important semantical information can reduce the IR processing time. This external representation can be generated using either statistical approach i.e., word counts or distributed semantical approach i.e., word embeddings. Therefore, there is a motivation to develop such IR system which can understand the specialized sub-domain language and domain-specific jargon of biomedical domain and assist researchers and medical professionals by effectively and efficiently retrieving most relevant information given a query.

RDoC Tasks aims at exploring information retrieval (IR) and information extraction (IE) tasks on selected abstracts from PubMed dataset. While Task-1 aims to rank abstracts i.e., *coarse granularity*, Task-2 aims to rank sentences i.e., *fine granularity* and hence the term *multi-grain*. An RDoC construct combines information from multiple sources like genomics, symptoms, behaviour etc. and therefore, is a much broader way of describing mental health disorders than symptoms based approach. Table 1 shows the association between PubMed abstracts and RDoC constructs depending on the semantic knowledge of the highlighted content words. Both of these tasks aim in the direction of *ease of accessibility* of PubMed abstracts labelled with diverse RDoC constructs so that this information can reach its full potential and can be of help to biomedical researchers and healthcare professionals.

\* : Equal Contribution

PMID	RDoC Construct	PubMed Abstract
14998902	Acute_Threat_Fear	<i>Title:</i> Mother lowers <b>glucocorticoid</b> levels of preweaning rats after <b>acute threat</b> . <i>Abstract:</i> Exposure to a <b>deadly threat</b> , an adult male rat, induced the release of <b>corticosterone</b> in 14-day-old rat pups. The <b>endocrine stress response</b> was decreased when the pups were reunited with their mother immediately after exposure. These findings demonstrate that social variables can reduce the consequences of an <b>aversive experience</b> .
21950094	Sleep_Wakefulness	<i>Title:</i> Central mechanisms of <b>sleep-wakefulness cycle</b> <i>Abstract:</i> Brief anatomical, <b>physiological and neurochemical</b> basics of the <b>regulation of wakefulness, slow wave (NREM) sleep and paradoxical (REM) sleep</b> are regarded as representing by the end of the first decade of the second millennium.

Table 1: RDoC construct - This table shows two PubMed abstracts labelled with two different RDoC construct and PubMed ID (PMID). Highlighted words (blue and red) in each abstract shows content words which together provide the semantic understanding of the corresponding RDoC constructs.

## 2 Task Description and Contributions

**RDoc-IR Task-1:** The task aims at retrieving and ranking the PubMed abstracts (within each of the eight clusters) that are relevant for the RDoC construct (i.e, a query) related to the cluster in the abstract appears. The training data consists of abstracts (title + sentences) each annotated with one or more RDoC constructs. Test data consists of abstracts without annotation and the goal is to submit a ranked lists of relevant articles for each medical domain RDoC construct.

**RDoc-IE Task-2** The task aims at extracting the most relevant sentence from each PubMed abstract for the corresponding RDoC construct. The input consists of an abstract (title  $t$  and sentences  $s$ ) for an RDoC construct  $q$ . The training data consists of abstracts each annotated with one RDoC construct and the most relevant sentence. Test data contains abstracts relevant for RDoC constructs and the goal is to submit a list of predicted most relevant sentence for each abstract.

**Our Contributions:** Following are our multi-fold contributions in this paper:

(1) **RDoC-IR Task-1:** We perform document (or abstract) ranking in two steps, first using supervised neural topic model and SVM. Moreover, we have introduced attentions in supervised neural topic model, along with pre-trained word embeddings from several sources. Then, we re-rank documents using BM25 and similarity scores between query and query-aware attention-based document representation.

Comparing with other participating systems in the shared task, our submission is ranked 1<sup>st</sup> with a mAP score of 0.86.

(2) **RDoC-IE Task-2:** We have addressed the

sentence ranking task by introducing unsupervised and supervised sentence ranking schemes. Moreover, we have employed multi-view representations consisting of bag-of-words, TF-IDF and query-aware attention-based sentence representation via enhanced query-sentence interactions. We have also investigated relevance of title with the sentences and coined ways to incorporate both query-sentence and title-sentence relevance scores in ranking sentences with an abstract.

Comparing with other participating systems in the shared task, our submission is ranked 1<sup>st</sup> with a macro average accuracy of 0.58. Our code is available at [https://github.com/YatinChaudhary/RDoC\\_Task](https://github.com/YatinChaudhary/RDoC_Task).

## 3 Methodology

In this section, we first describe representing a query, sentence and document using local and distributed representation schemes. We further describe enhanced query-document (query-title and query-content) and query-sentence interactions to compute query-aware document or sentence representations for Task-1 and Task-2, respectively. Finally, we discuss the application of supervised neural topic modeling in ranking documents for task 1 and introduce unsupervised and supervised sentence rankers for Task-2.

### 3.1 Query, Sentence and Document Vectors

In this paper, we deal with texts of different lengths in form of query, sentence and document. In this section, we describe the way we represent the different texts.

**Bag-of-words (BoW) and Term frequency-inverse document frequency (TF-IDF):** We use two the local representation schemes: BoW and

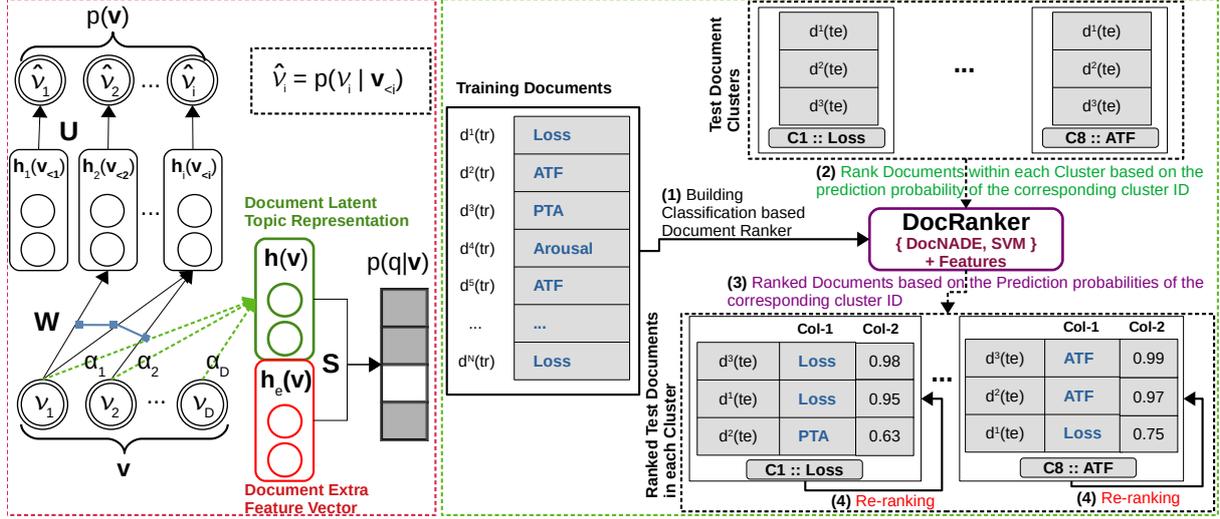


Figure 1: (Left) DocNADE Topic Model: Blue colored lines signify parameter sharing,  $\alpha$  attention weights are used to compute *latent document representation*  $h(v)$ ; (Right) RDoC Task-1 system architecture, where the numbered arrow (1) denotes the flow. **Col-1** indicates “predicted label” by the DocRanker and **Col-2** indicates “prediction probability” ( $p(q|v)$ ). “Features” inside DocRanker indicates *FastText* and *word2vec* pretrained embeddings.

TF-IDF (Manning et al., 2008) to compute sentence/document vectors.

**Embedding Sum Representation (ESR):** Word embeddings (Mikolov et al., 2013; Pennington et al., 2014) have been successfully used in computing distributed representation of text snippets (short or long). In ESR scheme, we employ the pre-trained word embeddings from FastText (Bojanowski et al., 2017) and word2vec (Mikolov et al., 2013). To represent a text (query, sentence or document), we compute the sum of (pre-trained) word vectors of each word in the text. E.g., ESR for a document  $d$  with  $D$  words can be computed as:

$ESR(d) = \tilde{\mathbf{d}} = \sum_{i=1}^D \mathbf{e}(d_i)$  where,  $\mathbf{e} \in \mathbb{R}^E$  is the pre-trained embedding vector of dimension  $E$  for the word  $d_i$ .

**Query-aware Attention-based Representation (QAR) for Documents and Sentences:** Unlike ESR, we reward the maximum matches between a query and document by computing density of matches between them, similar to McDonald et al. (2018). In doing so, we introduce a weighted sum of word vectors from pre-trained embeddings and therefore, incorporate importance/attention of certain words in document (or sentence) that appear in the query text.

For an enhanced query-aware attention based document (or sentence) representation, we first compute an histogram  $\mathbf{a}_i(d) \in \mathbb{R}^D$  of attention weights for each word  $k$  in the document  $d$  (or

sentence  $s$ ) relative to the  $i$ th query word  $q_i$ , using cosine similarity:

$$\mathbf{a}_i(d) = [a_{i,k}]_{k=1}^D \text{ where, } a_{i,k} = \frac{\mathbf{e}(q_i)^T \mathbf{e}(d_k)}{\|\mathbf{e}(q_i)\| \|\mathbf{e}(d_k)\|}$$

for each  $k$ th word in the document  $d$ . Here,  $\mathbf{e}(w)$  refers to an embedding vector of the word  $w$ .

We then compute an query-aware attention-based representation  $\Phi_i(d)$  of document  $d$  from the viewpoint of  $i$ th query word by summing the word vectors of the document, weighted by their attention scores  $\mathbf{a}_i(d)$ :

$$\Phi_i(d) = \sum_{k=1}^D a_{i,k}(d) \mathbf{e}(d_k) = \mathbf{a}_i(d) \odot [\mathbf{e}(d_k)]_{k=1}^D$$

where  $\odot$  is an element-wise multiplication operator.

Next, we compute density of matches between several words in query and the document by summing each of the attention histograms  $\mathbf{a}_i$  for all the query terms  $i$ . Therefore, the *query-aware document representation* for a document (or sentence) relative to all query words in  $q$  is given by:

$$QAR(d) = \Phi_q(d) = \sum_i^{|q|} \Phi_i(d) \quad (1)$$

Similarly, a *query-aware sentence representation*  $\Phi_q(s)$  and *query-aware title representation*  $\Phi_q(t)$  can be computed for the sentence  $s$  and document title  $t$ , respectively.

For query representation, we use ESR scheme as  $\tilde{\mathbf{q}} = \sum_{i=1}^{|q|} \mathbf{e}(w_i)$ .

Figure 2 illustrates the computation of *query-aware attention-based sentence representation*.

### 3.2 Document Neural Topic Models

Topic models (TMs) (Blei et al., 2003) have shown to capture thematic structures, i.e., topics appearing within the document collection. Beyond interpretability, topic models can extract latent document representation that is used to perform document retrieval. Recently, Gupta et al. (2019a) and Gupta et al. (2019b) have shown that the neural network-based topic models (NTM) outperform LDA-based topic models (Blei et al., 2003; Srivastava and Sutton, 2017) in terms of generalization, interpretability and document retrieval.

In order to perform document classification and retrieval, we have employed supervised version of neural topic model with extra features and further introduced word-level attention in a neural topic model, i.e. in DocNADE (Larochelle and Lauly, 2012; Gupta et al., 2019a).

**Supervised NTM (SupDocNADE):** Document Neural Autoregressive Distribution Estimator (DocNADE) is a neural network based topic model that works on bag-of-words (BoW) representation to model a document collection in a language modeling fashion.

Consider a document  $d$ , represented as  $\mathbf{v} = [v_1, \dots, v_i, \dots, v_D]$  of size  $D$ , where  $v_i \in \{1, \dots, Z\}$  is the index of  $i$ th word in the vocabulary and  $Z$  is the vocabulary size. DocNADE models the joint distribution  $p(\mathbf{v})$  of document  $\mathbf{v}$  by decomposing  $p(\mathbf{v})$  into autoregressive conditional of each word  $v_i$  in the document, i.e.,  $p(\mathbf{v}) = \sum_{i=1}^D p(v_i | \mathbf{v}_{<i})$ , where  $\mathbf{v}_{<i} \in \{v_1, \dots, v_{i-1}\}$ .

As shown in Figure 1 (left), DocNADE computes each autoregressive conditional  $p(v_i | \mathbf{v}_{<i})$  using a feed forward neural network for  $i \in \{1, \dots, D\}$  as,

$$p(v_i = w | \mathbf{v}_{<i}) = \frac{\exp(b_w + \mathbf{U}_{w,:} \mathbf{h}(\mathbf{v}_{<i}))}{\sum_{w'} \exp(b_{w'} + \mathbf{U}_{w',:} \mathbf{h}(\mathbf{v}_{<i}))}$$

$$\mathbf{h}_i(\mathbf{v}_{<i}) = f(\mathbf{c} + \sum_{j<i} \mathbf{W}_{:,v_j})$$

where,  $f(\cdot)$  is a non-linear activation function,  $\mathbf{W} \in \mathbb{R}^{H \times Z}$  and  $\mathbf{U} \in \mathbb{R}^{Z \times H}$  are encoding and decoding matrices,  $\mathbf{c} \in \mathbb{R}^H$  and  $\mathbf{b} \in \mathbb{R}^Z$  are encoding and decoding biases,  $H$  is the number of units in latent representation  $\mathbf{h}_i(\mathbf{v}_{<i})$ . Here,

$\mathbf{h}_i(\mathbf{v}_{<i})$  contains information of words preceding the word  $v_i$ . For a document  $\mathbf{v}$ , the log-likelihood  $\mathcal{L}(\mathbf{v})$  and latent representation  $\mathbf{h}(\mathbf{v})$  are given as,

$$\mathcal{L}^{unsup}(\mathbf{v}) = \sum_{i=1}^D \log p(v_i | \mathbf{v}_{<i}) \quad (2)$$

$$\mathbf{h}(\mathbf{v}) = f(\mathbf{c} + \sum_{i=1}^D \mathbf{W}_{:,v_i}) \quad (3)$$

Here,  $\mathcal{L}(\mathbf{v})$  is used to optimize the topic model in unsupervised fashion and  $\mathbf{h}(\mathbf{v})$  encodes the topic proportion. See Gupta et al. (2019a) for further details on training unsupervised DocNADE.

Here, we extend the unsupervised version to DocNADE with a hybrid cost  $\mathcal{L}^{hybrid}(\mathbf{v})$ , consisting of a (supervised) discriminative training cost  $p(y = q | \mathbf{v})$  along with an unsupervised generative cost  $p(\mathbf{v})$  for a given query  $q$  and associated document  $\mathbf{v}$ :

$$\mathcal{L}^{hybrid}(\mathbf{v}) = \mathcal{L}^{sup}(\mathbf{v}) + \lambda \cdot \mathcal{L}^{unsup}(\mathbf{v}) \quad (4)$$

where  $\lambda \in [0, 1]$ . The supervised cost is given by:

$$\mathcal{L}^{sup}(\mathbf{v}) = p(y = q | \mathbf{v}) = \text{softmax}(\mathbf{d} + \mathbf{S} \mathbf{h}(\mathbf{v}))$$

Here,  $\mathbf{S} \in \mathbb{R}^{L \times H}$  and  $\mathbf{d} \in \mathbb{R}^L$  are output matrix and bias,  $L$  is the total number of unique RDoC constructs (i.e., unique query labels).

**Supervised Attention-based NTM (a-SupDocNADE):** Observe in equation 3 that the DocNADE computes document representation  $\mathbf{h}(v)$  via aggregation of word embedding vectors without considering attention over certain words. However, certain content words own high important, especially in classification task. Therefore, we have introduced attention-based embedding aggregation in supDocNADE (Figure 1, left):

$$\mathbf{h}(\mathbf{v}) = f(\mathbf{c} + \sum_{i=1}^D \alpha_i \mathbf{W}_{:,v_i}) \quad (5)$$

Here,  $\alpha_i$  is an attention score of each word  $i$  in the document  $\mathbf{v}$ , learned via supervised training.

Additionally, we incorporate extra word features, such as pre-trained word embeddings from several sources: FastText ( $\mathbf{E}^{fast}$ ) (Bojanowski et al., 2017) and word2vec ( $\mathbf{E}^{word2vec}$ ) (Mikolov et al., 2013). We introduce these features by concatenating  $\mathbf{h}_e(\mathbf{v})$  with  $\mathbf{h}(\mathbf{v})$  in the supervised portion of the a-supDocNADE model, as

$$\mathbf{h}_e(\mathbf{v}) = f\left(\mathbf{c} + \sum_{i=1}^D \alpha_i (\mathbf{E}_{:,v_i}^{fast} + \mathbf{E}_{:,v_i}^{word2vec})\right) \quad (6)$$

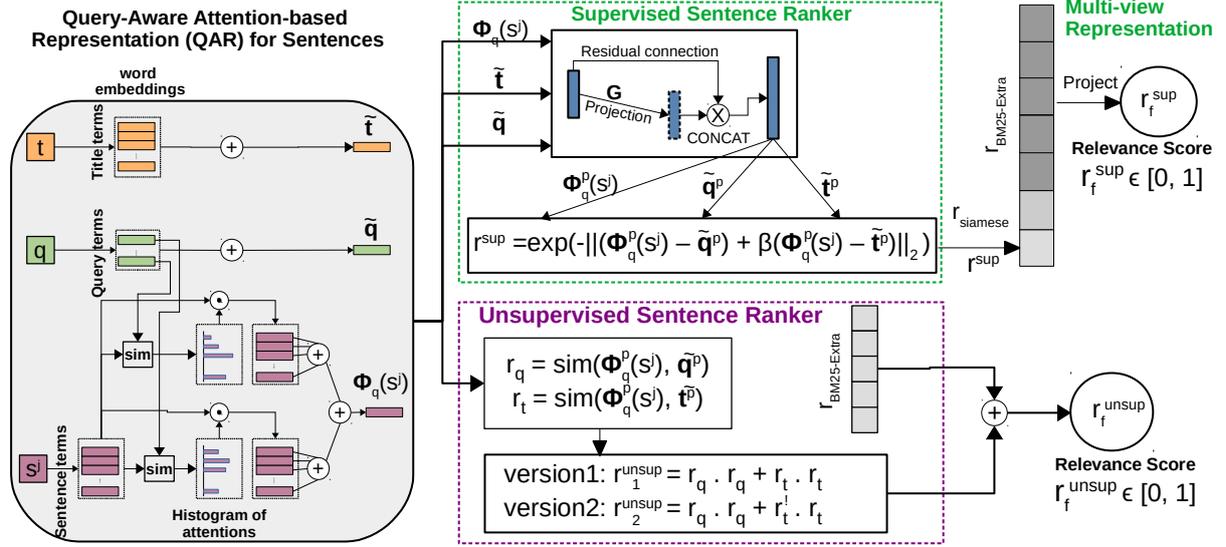


Figure 2: RDoC Task-2 System Architecture for Supervised and Unsupervised sentence ranking, consisting of: Query-aware Representation, Supervised and Unsupervised sentence rankers for computing the final relevance scores  $r_f^{sup}$  and  $r_f^{unsup}$ , respectively. Here, *sim* refers to cosine similarity.

Therefore, the classification portion of *a-supDocNADE* with additional features is given by:

$$p(q|\mathbf{v}) = \text{softmax}(\mathbf{d} + \mathbf{S}' \cdot \text{concat}(\mathbf{h}(\mathbf{v}), \mathbf{h}_e(\mathbf{v})))$$

where,  $\mathbf{S}' \in \mathbb{R}^{H' \times L}$  and  $H' = H + E^{fast} + E^{word2vec}$ .

### 3.3 Traditional Methods for IR

**BM25:** A ranking function proposed by [Robertson and Zaragoza \(2009\)](#) is used to estimate the relevance of a document for a given query.

**BM25-Extra:** The relevance score of BM-25 is combined with four extra features: (1) percentage of query words with exact match in the document, (2) percentage of query words bigrams matched in the document, (3) IDF weighted document vector for feature #1, and (4) IDF weighted document vector for feature #2. Therefore, BM25-Extra returns a vector of 5 scores.

### 3.4 System Description for RDoC Task-1

RDoC Task-1 aims at retrieving and ranking of PubMed abstracts (title and content) that are relevant for 8 RDoC constructs. Participants are provided with 8 clusters, each with a RDoC construct label and required to rank abstracts within each cluster based on their relevance to the corresponding cluster label. Each cluster contains abstracts relevant to its RDoC construct, while some (or most) of the abstracts are noisy in the sense that

they belong to a different RDoC construct. Ideally, the participants are required to rank abstracts in each of the clusters by determining their relevance with the RDoC construct of the cluster in which they appear.

To address the RDoC Task-1, we learn a mapping function between latent representation  $\mathbf{h}(\mathbf{v})$  of a document (i.e., abstract)  $\mathbf{v}$  and its RDoC construct, i.e., query words  $q$  in a supervised fashion. In doing so, we have employed supervised classifiers, especially supervised neural topic model **a-supDocNADE** (section 3.2) for document ranking. We treat  $q$  as label and maximize  $p(q|\mathbf{v})$  leading to maximize  $\mathcal{L}^{hybrid}(\mathbf{v})$  in *a-supDocNADE* model.

As demonstrated in Figure 1 (right), we perform document ranking in two steps:

(1) **Document Relevance Ranking:** We build a supervised classifier using all the training documents and their corresponding labels (RDoC constructs), provided with the training set. At the test time, we compute prediction probability score  $p(CID = q|\mathbf{v}^{test}(CID))$  of the label=CID for each test document  $\mathbf{v}^{test}(CID)$  in the cluster, CID. This prediction probability (or confidence score) is treated as a relevance score of the document for the RDoC construct of the cluster. Figure 1(right) shows that we perform document ranking using the probability scores (col-2) of the RDoC construct (e.g. *loss*) within the cluster  $C1$ . Observe that a test document with least confidence

for a cluster are ranked lower within the cluster and thus, improving mean average precision (mAP). Additionally, we also show the predicted RDoC construct in col-1 by the supervised classifier.

(2) **Document Relevance Re-ranking:** Secondly, we re-ranked each document  $\mathbf{v}$  (title+abstract) within each cluster (with label  $q$ ) using *unsupervised ranking*, where the relevance scores are computed as: (a) **reRank(BM25-Extra)**: sum each of the 5 relevance scores to get the final relevance, and (b) **reRank(QAR)**: cosine-similarity(QAR( $\mathbf{v}$ ),  $\tilde{\mathbf{q}}$ ).

### 3.5 System Description for RDoC Task-2

The RDoC Task-2 aims at extracting the most relevant sentence from each of the PubMed abstract for the corresponding RDoC construct. Each abstract consists of title  $t$  and sentences  $s$  with an RDoC construct  $q$ .

To address RDoc Task-2, we first compute *multi-view* representation: BoW, TF-IDF and QAR (i.e.,  $\Phi_q(s^j)$ ) for each sentence  $s^j$  in an abstract  $d$ . On other hand, we compute ESR representation for RDoC construct (query  $q$ ) and title  $t$  of the abstract  $d$  to obtain  $\tilde{\mathbf{q}}$  and  $\tilde{\mathbf{t}}$ , respectively. Figure 2 and section 3.1 describe the computation of these representations. We then use the representations ( $\Phi_q(s^j)$ ,  $\tilde{\mathbf{t}}$  and  $\tilde{\mathbf{q}}$ ) to compute a relevance scores of a sentence  $s_j$  relative to  $q$  and/or  $t$  via *unsupervised* and *supervised* ranking schemes, discussed in the following section.

#### 3.5.1 Unsupervised Sentence Ranker

As shown in Figure 2, we first extract representations:  $\Phi_q(s^j)$ ,  $\tilde{\mathbf{t}}$  and  $\tilde{\mathbf{q}}$  for the sentence  $s^j$  query  $q$  and title  $t$ . During ranking sentences within an abstract for the given RDoC construct  $q$ , we also consider title  $t$  in computing the relevance score for each sentence relative to  $q$  and  $t$ . It is inspired from the fact that the title often contains relevant terms (or words) appearing in sentence(s) of the document (or abstract). On top, we observe that  $q$  is a very short text and non-descriptive, leading to minimal text overlap with  $s$ .

We compute two relevance scores:  $r_q$  and  $r_t$  for a sentence  $s^j$  with respect to a query  $q$  and title  $t$ , respectively.

$$r_q = \text{sim}(\tilde{\mathbf{q}}, \Phi_q(s^j)) \text{ and } r_t = \text{sim}(\tilde{\mathbf{t}}, \Phi_q(s^j))$$

Now, we devise two ways to combine the rele-

vance scores  $r_q$  and  $r_t$  in unsupervised paradigm:

$$\text{version1: } r_1^{\text{unsup}} = r_q \cdot r_q + r_t \cdot r_t$$

Observe that the relevance scores are weighted by itself. However, the task-2 expects a higher importance to the relevance score  $r_q$  over  $r_t$ . Therefore, we coin the following weighting scheme to give higher importance to  $r_q$  only if it is higher than  $r_t$  otherwise we compute a weight factor  $r'_t$  for  $r_t$ .

$$\text{version2: } r_2^{\text{unsup}} = r_q \cdot r_q + r'_t \cdot r_t$$

where  $r'_t$  is compute as:

$$r'_t = (r_t > r_q) |r_t - r_q|$$

The relevance score  $r_2^{\text{unsup}}$  is effective in ranking sentences when a query and sentence does not overlap. In such a scenario, a sentence is scored by title, penalized by a factor of  $|r_t - r_q|$ .

At the end, we obtain a final relevance score  $r_f^{\text{unsup}}$  for a sentence  $s^j$  by summing the relevance scores of BM25-Extra and  $r_1^{\text{unsup}}$  or  $r_2^{\text{unsup}}$ .

#### 3.5.2 Supervised Sentence Ranker

Beyond unsupervised ranking, we further investigate sentence ranking in supervised paradigm by introducing a distance metric between the query (or title) and sentence vectors.

Figure 2 describes the computation of relevance score for a sentence  $s^j$  using a supervised sentence ranker scheme. Like the unsupervised ranker (section 3.5.1), the supervised ranker also employs vector representations:  $\Phi_q(s^j)$ ,  $\tilde{\mathbf{t}}$  and  $\tilde{\mathbf{q}}$ . Using the projection matrix  $\mathbf{G}$ , we then apply a projection to each of the representation to obtain  $\Phi_q^p(s^j)$ ,  $\tilde{\mathbf{t}}^p$  and  $\tilde{\mathbf{q}}^p$ . Here, the operator  $\otimes$  performs concatenation of the projected vector with its input via residual connection. Next, we apply a Manhattan distance metric to compute similarity (or relevance) scores, following Gupta et al. (2018):

$$r^{\text{sup}} = \exp\left(-\|(\Phi_q^p(s^j), \tilde{\mathbf{q}}^p) + \beta (\Phi_q^p(s^j), \tilde{\mathbf{t}}^p)\|_2\right)$$

where  $\beta \in [0, 1]$  controls the relevance of title, determined by cross-validation. A final relevance score  $r_f^{\text{sup}} \in [0, 1]$  is computed by feeding a vector [ $r^{\text{sup}}$ ,  $r_{\text{siamese}}^{\text{sup}}$ , BM25-extra] into a supervised linear regression, which is trained end-to-end by minimizing mean squared error between the  $r_f^{\text{sup}}$  and  $\{0, 1\}$ , i.e., 1 when the sentence  $s^j$  is relevant to query  $q$ . Here,  $r_{\text{siamese}}^{\text{sup}}$  refers to a relevance

	L1	L2	L3	L4	L5	L6	L7	L8	Total
All data	39	38	47	21	28	27	48	18	266
Train set	31	30	37	16	22	21	38	14	209
Dev set	8	8	10	5	6	6	10	4	57
Test set (Task1)	79	108	123	144	138	139	122	146	999
Test set (Task2)	19	26	30	35	34	34	30	36	244

Table 2: Data statistics - # of PubMed abstracts belonging to each RDoC construct in different data partitions. (L1: “Acute Threat Fear”; L2: “Arousal”; L3: “Circadian Rhythms”; L4: “Frustrative Nonreward”; L5: “Loss”; L6: “Potential Threat Anxiety”; L7: “Sleep Wakefulness”; L8: “Sustained Threat”)

Model	Feature	Classification Accuracy	Ranking mAP
SVM	BoW	0.947	<b>0.992</b>
Re-ranking	reRank #1	-	<b>0.992</b>
with	reRank #2	-	<b>0.992</b>
cluster label	reRank #3	-	<b>0.992</b>
a-supDocNADE	Random init	0.912	0.930
	+ FastText	0.947	0.949
	+ BioNLP	0.965	0.983
Re-ranking	reRank #1	-	0.985
with	reRank #2	-	<b>0.994</b>
cluster label	reRank #3	-	<b>0.994</b>

Table 3: RDoC Task-1 results (on development set): Classification accuracy and mean Average Precision (mAP) of a-supDocNADE and SVM models. Each model’s *classification accuracy* and *ranking mAP* (using prediction probabilities) are shown together. Furthermore, each model’s ranked clusters are re-ranked using different re-ranking algorithms. Best mAP score for each model is marked in **bold**. (*reRank #1*: “reRank(BM25-Extra)”; *reRank #2*: “reRank(QAR)”; *reRank #3*: “reRank(BM25-Extra) + reRank(QAR)”)

score computed between  $q$  and  $s^j$  via Siamese-LSTM (Gupta et al., 2018).

To perform sentence ranking within an abstract for a given RDoC construct  $q$ , the relevance score  $r_f^{sup}$  (or  $r_f^{unsup}$ ) is computed for all the sentences and a sentence with the highest score is extracted.

## 4 Experiments and Results

### 4.1 Data Statistics and Experimental Setup

**Dataset Description:** Dataset for RDoC Tasks contains a total of 266 PubMed abstracts labelled with 8 RDoC constructs in a single label fashion. Number of abstracts for each RDoC construct is described in Table 2, where *first* row describes the statistics for all abstracts and *second & third* row shows the split of those abstracts into *training* and *development* sets maintaining a 80-20 ratio for each RDoC construct. For Task-1, each PubMed

abstract contains its associated title, PubMed ID (PMID) and label (RDoC construct). In addition for Task-2, each PubMed abstract also contains a list of most relevant sentences from that abstract. Final evaluation test data for Task-1 & Task-2 contains 999 & 244 abstracts respectively.

We use “RegexTokenizer” from scikit-learn to tokenize abstracts and lower-cased all tokens. After this, we remove those tokens which occur in less than 3 abstracts and also remove stopwords (using nltk). For computing BM25-Extra relevance score, we use unprocessed raw text of sentences and titles.

**Experimental Setup:** As the training dataset labelled with RDoC constructs is very small, we use an external source of semantical knowledge by incorporating pretrained distributional word embeddings (Zhang et al., 2019) from FastText model (Bojanowski et al., 2017) trained on the entire corpus of PubMed and MIMIC III Clinical notes (Johnson et al., 2016). Similarly, we also use pretrained word embeddings (Moen and Ananiadou, 2013) from word2vec model (Mikolov et al., 2013) trained on PubMed and PMC abstracts. We create 3 folds\* of train/dev splits for cross-validation.

**RDoC Task-1:** For DocNADE topic model, we use latent representation of size 50. We use pretrained FastText embeddings of size 300 and pretrained word2vec embeddings of size 200. For SVM, we use Bag-of-words (BoW) representation of abstracts with radial basis kernel function. PubMed abstracts are provided in eight different clusters, one for each RDoC construct, for final test set evaluation.

**RDoC Task-2:** We use pretrained FastText embeddings to compute *query-aware sentence representation* of a sentence ( $\Phi_q(s^j)$ ), title ( $\tilde{t}$ ) and query ( $\tilde{q}$ ) representations. We also train Replicated-Siamese-LSTM (Gupta et al., 2018) model with input as sentence and query pair i.e.,  $(s^j, q)$  and label as 1 if  $s^j$  is relevant otherwise 0. We use  $\beta \in \{0, 1\}$ .

### 4.2 Results: RDoC Task-1

Table 3 shows the performance of supervised Document Ranker models i.e, **a-supDocNADE** and **SVM**, for Task-1. *SVM* achieves a classification accuracy of 0.947 and mean average precision

\*we only report results on  $fold_1$  because of best scores on partial test dataset

Ranking (with Prediction Probability)			Re-ranking (with BM25-Extra)	
PMID	Pred Prob	Gold Label	PMID	Gold Label
22906122	0.90	PTA	26005838	PTA
24286750	0.77	PTA	22906122	PTA
17598732	0.61	PTA	28828218	PTA
26005838	0.56	PTA	26773206	PTA
28316567	0.46	Loss	24286750	PTA
28828218	0.45	PTA	17598732	PTA
26773206	0.41	PTA	28316567	Loss

Table 4: RDoC Task-1 analysis: Ranking of PubMed abstracts within “Potential Threat Anxiety (PTA)” cluster using supervised prediction probabilities ( $p(q|\mathbf{v})$ ). It shows that an intruder/noisy abstract (Gold Label: *Loss*) is assigned higher probability than the abstracts with same Gold Label as the cluster. But, using re-ranking with BM25-Extra ( $reRank(BM25-Extra)$ ) relevance score assigns lowest relevance to the intruder abstract.

(mAP) of 0.992 by ranking the abstracts in their respective clusters using the supervised prediction probabilities ( $p(q|\mathbf{v})$ ). After that, we use three different relevance scores: (1)  $reRank(BM25-Extra)$ , (2)  $reRank(QAR)$  and (3)  $reRank(BM25-Extra) + reRank(QAR)$ , for re-ranking of the abstracts in their respective clusters. It is to be noted that the *ranking mAP* of the clusters using prediction probabilities is already the best possible i.e., the intruder abstracts (abstracts with different label (RDoC construct) than the cluster label) are at the bottom of the ranked clusters. Therefore, re-ranking of these clusters would not achieve a better score. Similarly, we train *a-supDocNADE* model with three different settings: (1) random weight initialization, (2) incorporating FastText embeddings ( $\mathbf{h}_e(\mathbf{v})$ ) and (3) incorporating FastText and word2vec embeddings ( $\mathbf{h}_e(\mathbf{v})$ ). By using the pretrained embeddings, the classification accuracy increases from 0.912 to 0.965, this shows that distributional pretrained embeddings carry significant semantic knowledge. Furthermore, re-ranking using  $reRank(BM25-Extra)$  and  $reRank(QAR)$  further results in the improvement of mAP score (0.994 vs 0.983) by shifting the intruder documents at the bottom of each impure cluster.

### 4.3 Analysis: RDoC Task-1

Table 4 shows an impure “Potential Threat Anxiety” cluster of abstracts containing an intruder abstract with label (RDoC construct) “Loss”. When this cluster is ranked on the basis of predic-

Model	Feature	Recall	F1	Macro-Average Accuracy
Unsupervised	$reRank(BM25-Extra)$ [#1]	0.316	0.387	0.631
	version1 [#2]	<b>0.351</b>	<b>0.412</b>	<b>0.701</b>
	version2 [#3]	0.263	0.345	0.526
Supervised	$r_f^{sup}(\beta = 0)$ [#4]	<b>0.386</b>	<b>0.436</b>	<b>0.772</b>
	$r_f^{sup}(\beta = 1)$ [#5]	0.368	0.424	0.737
Ensemble	{#1, #2, #4}	<b>0.395</b>	<b>0.441</b>	<b>0.789</b>
	{#1, #3, #4}	0.316	0.387	0.631
	{#2, #4, #5}	<b>0.395</b>	<b>0.441</b>	<b>0.789</b>
	{#1, #2, #3, #4, #5}	0.368	0.424	0.737

Table 5: RDoC Task-2 results (on development set): Performance of unsupervised and supervised sentence rankers (Figure 2) under different configurations. Best scores for each model is marked in **bold**.

tion probabilities ( $p(q|\mathbf{v})$ ), then “Loss” abstract is ranked third from the bottom and it degrades the mAP score of the retrieval system. But after re-ranking this cluster using  $reRank(BM25-Extra)$  relevance score, the “Loss” abstract is ranked at the bottom, thus maximizing the mAP score. Therefore, re-ranking with BM25-Extra on top of ranking with  $p(q|\mathbf{v})$  is, evidently, a robust abstract/document ranking technique.

### 4.4 Results: RDoC Task-2

Table 5 shows results for Task-2 using *three* unsupervised and *two* supervised sentence ranker models. For unsupervised model, using  $reRank(BM25-Extra)$  relevance score between a query ( $q$ ), label (RDoC construct) of an abstract, and all the sentences ( $s^j$ ) in an abstract, we get an macro-average accuracy (MAA) of 0.631. However, using *version1* and *version2* models (see Fig 2), we achieve a MAA score of 0.701 and 0.526 respectively. Higher accuracy of *version1* model suggests that title ( $t$ ) of an abstract also contains the essential information regarding the most relevant sentence. For supervised model, we get an MAA score of 0.772 and 0.737 by setting  $\beta = 0$  & 1 in supervised relevance score ( $r_f^{sup}$ ) equation in section 3.5.2. Hence, for supervised sentence ranker model, title ( $t$ ) is playing a negative influence in correctly identifying the relevance ( $r_f^{sup}$ ) of different sentences. Furthermore, we combine the knowledge of unsupervised and supervised sentence rankers by creating multiple ensembles (majority voting) of the predictions from different models. We achieve the highest MAA score of 0.789 by combining the predictions of (1)  $reRank(BM25-Extra)$ , (2) *version1*, and (3)  $r_f^{sup}$  with  $\beta = 0$ . Notice that all the proposed supervised and unsupervised sentence ranking mod-

Team	RDoC Task-1 (Official Results)										RDoC Task-2 (Official Results)							
	L1	L2	L3	L4	L5	L6	L7	L8	mAP	L1	L2	L3	L4	L5	L6	L7	L8	MAA
<i>MIC-CIS</i>	0.85	0.92	<b>1.00</b>	<b>0.73</b>	<b>0.78</b>	<b>0.94</b>	<b>1.00</b>	0.63	<b>0.86</b>	<b>0.74</b>	<b>0.73</b>	0.47	0.37	<b>0.74</b>	<b>0.59</b>	<b>0.60</b>	0.42	<b>0.58</b>
<i>Javad Rafiei Asl</i>	<b>0.89</b>	<b>0.93</b>	<b>1.00</b>	0.69	0.74	0.87	<b>1.00</b>	<b>0.64</b>	0.85	0.68	0.62	0.47	0.34	0.56	0.32	0.50	0.36	0.48
<i>Ramya Tekumalla</i>	0.83	0.91	<b>1.00</b>	0.67	0.71	0.89	<b>1.00</b>	<b>0.64</b>	0.83	0.37	0.12	0.10	0.11	0.26	0.15	0.17	0.14	0.18
<i>Daniel Laden</i>	0.67	0.84	<b>1.00</b>	0.61	0.61	0.81	<b>0.98</b>	0.41	0.74	-	-	-	-	-	-	-	-	-
<i>Shyaman Jayasundara</i>	-	-	-	-	-	-	-	-	-	0.47	0.42	0.60	0.29	0.62	0.38	0.57	<b>0.47</b>	0.48
<i>Fei Li</i>	-	-	-	-	-	-	-	-	-	0.58	0.46	<b>0.70</b>	<b>0.43</b>	0.26	0.41	0.33	<b>0.47</b>	0.46

Table 6: RDoC Tasks official results - performance on test set of different competing systems. Best score in each column is marked in **bold**. (Refer to Table 2 for header notations) (mAP: “Mean Average Precision”; MAA: Macro-Average Accuracy)

PubMed Abstract (PMID: “23386529”; RDoC construct: “Loss”)		
Most Relevant Sentence (using reRank (BM25-Extra))	Sentence ID	Gold Label
Nurses are expected to care for grieving women and families suffering from perinatal loss.	#1	Not relevant
Most Relevant Sentence (using Ensemble {#1, #2, #4})	-	-
We found that nurses experience a grieving process similar to those directly suffering from perinatal loss.	#6	Relevant

Table 7: RDoC Task-2 analysis: This table shows that the most relevant sentence predicted using reRank (BM25-Extra) is actually not a relevant sentence, but Ensemble {#1, #2, #4} (Table 5) predicts the correct sentence as the most relevant.

els (except [#3]) outperform traditional ranking models, e.g., *reRank(BM25-Extra)* in terms of query-document relevance score.

#### 4.5 Analysis: RDoC Task-2

Table 7 shows that the most relevant sentence predicted by *reRank(BM25-Extra)* is actually a non-relevant sentence. But an ensemble of predictions from both unsupervised and supervised ranker models correctly predicts the relevant sentence. This suggests that complementary knowledge of different models is able to capture the relevance of sentences on different scales and majority voting among them is, evidently, a robust sentence ranking technique.

#### 4.6 Results: RDoC Task 1 & 2 on Test set

Table 6 shows the final evaluation scores of different competing systems for both the RDoC Task-1 & Task-2 on final test set. Observe that our submission (MIC-CIS) scored a mAP score of 0.86 and MAA of 0.58 in Task-1 and Task-2, respectively. Notice that we outperform the second best

system by 20.83% (0.58 vs 0.48) margin in Task2.

## 5 Conclusion

In conclusion, both supervised neural topic model and SVM can effectively perform ranking of PubMed abstracts in a given cluster based on the prediction probabilities. However, a further re-ranking using *BM25-Extra* or *query-aware sentence representation (QAR)* has proven to maximize the mAP score by correctly assigning the lowest relevance score to the intruder abstracts. Also, unsupervised and supervised sentence ranker models using query-title-sentence interactions outperform the traditional BM25-Extra based ranking model by a significant margin.

In future, we would like to introduce complementary feature representation via hidden vectors of LSTM jointly with topic models and would like to further investigate the interpretability (Gupta et al., 2015; Gupta and Schütze, 2018) of the proposed neural ranking models in the sense that one can extract salient patterns determining relationship between query and text. Another promising direction would be introduce abstract information, such as part-of-speech and named entity tags (Lample et al., 2016; Gupta et al., 2016) to augment information retrieval (IR).

## Acknowledgment

This research was supported by Bundeswirtschaftsministerium (bmwi.de), grant 01MD19003E (PLASS (plass.io)) at Siemens AG - CT Machine Intelligence, Munich Germany.

## References

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. *TACL*, 5:135–146.
- Pankaj Gupta, Bernt Andrassy, and Hinrich Schütze. 2018. Replicated siamese LSTM in ticketing system for similarity learning and retrieval in asymmetric texts. In *Proceedings of the Third Workshop on Semantic Deep Learning, SemDeep@COLING 2018, Santa Fe, New Mexico, USA, August 20, 2018*, pages 1–11.
- Pankaj Gupta, Yatin Chaudhary, Florian Buettner, and Hinrich Schütze. 2019a. Document informed neural autoregressive topic models with distributional prior. In *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, The Thirty-First Innovative Applications of Artificial Intelligence Conference, IAAI 2019, The Ninth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2019, Honolulu, Hawaii, USA, January 27 - February 1, 2019.*, pages 6505–6512.
- Pankaj Gupta, Yatin Chaudhary, Florian Buettner, and Hinrich Schütze. 2019b. Texttvec: Deep contextualized neural autoregressive topic models of language with distributed compositional prior. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*.
- Pankaj Gupta, Thomas Runkler, Heike Adel, Bernt Andrassy, Hans-Georg Zimmermann, and Hinrich Schütze. 2015. *Deep Learning Methods for the Extraction of Relations in Natural Language Text*. Master Thesis.
- Pankaj Gupta and Hinrich Schütze. 2018. LISA: explaining recurrent neural network judgments via layer-wise semantic accumulation and example to pattern transformation. In *Proceedings of the Workshop: Analyzing and Interpreting Neural Networks for NLP, BlackboxNLP@EMNLP 2018, Brussels, Belgium, November 1, 2018*, pages 154–164.
- Pankaj Gupta, Hinrich Schütze, and Bernt Andrassy. 2016. Table filling multi-task recurrent neural network for joint entity and relation extraction. In *COLING 2016, 26th International Conference on Computational Linguistics, Proceedings of the Conference: Technical Papers, December 11-16, 2016, Osaka, Japan*, pages 2537–2547.
- Alistair EW Johnson, Tom J Pollard, Lu Shen, H Lehman Li-wei, Mengling Feng, Mohammad Ghassemi, Benjamin Moody, Peter Szolovits, Leo Anthony Celi, and Roger G Mark. 2016. Mimic-iii, a freely accessible critical care database. *Scientific data*, 3:160035.
- Guillaume Lample, Miguel Ballesteros, Sandeep Subramanian, Kazuya Kawakami, and Chris Dyer. 2016. Neural architectures for named entity recognition. In *NAACL HLT 2016, The 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, San Diego California, USA, June 12-17, 2016*, pages 260–270.
- Hugo Larochelle and Stanislas Lauly. 2012. A neural autoregressive topic model. In *Advances in Neural Information Processing Systems 25: 26th Annual Conference on Neural Information Processing Systems 2012. Proceedings of a meeting held December 3-6, 2012, Lake Tahoe, Nevada, United States.*, pages 2717–2725.
- Christopher D. Manning, Prabhakar Raghavan, and Hinrich Schütze. 2008. *Introduction to information retrieval*. Cambridge University Press.
- Ryan McDonald, George Brokos, and Ion Androutsopoulos. 2018. Deep relevance ranking using enhanced document-query interactions. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018*, pages 1849–1860.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Gregory S. Corrado, and Jeffrey Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Advances in Neural Information Processing Systems 26: 27th Annual Conference on Neural Information Processing Systems 2013. Proceedings of a meeting held December 5-8, 2013, Lake Tahoe, Nevada, United States.*, pages 3111–3119.
- SPFGH Moen and Tapio Salakoski Sophia Ananiadou. 2013. Distributional semantics resources for biomedical text processing. *Proceedings of LBM*, pages 39–44.
- Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing, EMNLP 2014, October 25-29, 2014, Doha, Qatar, A meeting of SIGDAT, a Special Interest Group of the ACL*, pages 1532–1543.
- Stephen E. Robertson and Hugo Zaragoza. 2009. The probabilistic relevance framework: BM25 and beyond. *Foundations and Trends in Information Retrieval*, 3(4):333–389.
- Akash Srivastava and Charles A. Sutton. 2017. Autoencoding variational inference for topic models. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*.
- Yijia Zhang, Qingyu Chen, Zhihao Yang, Hongfei Lin, and Zhiyong Lu. 2019. Biwordvec, improving biomedical word embeddings with subword information and mesh. *Scientific data*, 6(1):52.