

A bag-of-concepts model improves relation extraction in a narrow knowledge domain with limited data

Jiyu Chen, Karin Verspoor and Zenan Zhai

School of Computing and Information Systems

The University of Melbourne, Australia

jiyuc@student.unimelb.edu.au, {karin.verspoor, zenan.zhai}@unimelb.edu.au

Abstract

This paper focuses on a traditional relation extraction task in the context of limited annotated data and a narrow knowledge domain. We explore this task with a clinical corpus consisting of 200 breast cancer follow-up treatment letters in which 16 distinct types of relations are annotated. We experiment with an approach to extracting typed relations called window-bounded co-occurrence (WBC), which uses an adjustable context window around entity mentions of a relevant type, and compare its performance with a more typical intra-sentential co-occurrence baseline. We further introduce a new bag-of-concepts (BoC) approach to feature engineering based on the state-of-the-art word embeddings and word synonyms. We demonstrate the competitiveness of BoC by comparing with methods of higher complexity, and explore its effectiveness on this small dataset.

1 Introduction

Applying automatic relation extraction on small data sets in a narrow knowledge domain is challenging. Here, we consider the specific context of a small clinical corpus, in which we have a variety of relation types of interest but limited examples of each. Transformation of clinical texts into structured sets of relations can facilitate the exploration of clinical research questions such as the potential risks of treatments for patients with certain characteristics, but large-scale annotation of these data sets is notoriously difficult due to the sensitivity of the data and the need for specialized clinical knowledge.

Rule-based methods (Abacha and Zweigenbaum, 2011; Verspoor et al., 2016) typically determine whether a particular type of relation exists in a given text by leveraging the context in which key clinical entities are mentioned. For instance,

if two entities with type `TestName` and `TestResult`, respectively, are observed in a given sentence, it is likely that a relation of type `TestFinding` exists between them. However, construction of high-precision rules defining relevant contexts is time-consuming and expensive, requiring extensive effort from domain experts.

The state-of-the-art machine learning algorithms such as neural network models (Nguyen and Grishman, 2016; Ammar et al., 2017; Huang and Wang, 2017) may over-fit in performing relation extraction in this context, due to a limited quantity of training instances.

In this work, we experiment with two automatic approaches to semantic relation extraction applied to a small corpus consisting of breast cancer follow-up treatment letters (Pitson et al., 2017), comparing a simple rule-based co-occurrence approach to machine learning classifiers.

The first approach, simple co-occurrence (Verspoor et al., 2016), is based on the assumption that most relevant relations are intra-sentential, that is, the relation between a pair of named entities is expressed within the scope of a single sentence. However, some relations may be expressed across sentence boundaries, and thus a single sentence may not be the ideal choice of scope, as shown in prior work that considers inter-sentential relations (also known as non-sentence or cross-sentence relations) (Panyam et al., 2016; Peng et al., 2017). We extend the co-occurrence approach to allow explicit adjustment of context window size, from one to two sentences, a method called Window-Bounded Co-occurrence (WBC). The best window size for a given relation is identified by choosing the one which produces the highest score under F_1 -measure on a development set.

The second approach is based on supervised binary classification. We transform the multi-relation extraction task into several independent

binary tasks. We build on a bag-of-concepts (BoC) (Sahlgren and Cöster, 2004) approach which models the text in terms of phrases or pre-identified concepts, extending it with word embeddings and word synonyms. We compare two different pre-trained word embedding models, and a number of other model variations. We also explore grouping of synonyms into abstracted *concepts*.

We find that the intra-sentential rule-based approach outperforms the approach which allows for a larger context window. The supervised learning models outperforms rule-based approaches under F_1 measure, and their results show that models using BoC features outperform models with BoW, dependency parse, or sentence embedding features. We also show that SVM outperforms complex models such as a feed-forward ANN in our low resource scenario, with less tendency to over-fitting.

2 Background

At present, the two primary approaches to automatic relation extraction over biomedical corpora are rule-based approaches (Verspoor et al., 2016; Abacha and Zweigenbaum, 2011) and machine learning approaches based on learners such as logistic regression, support vector machines (SVM) (Panyam et al., 2016) and convolutional neural networks (CNN) (Nguyen and Verspoor, 2018) together with sophisticated feature engineering methods.

Verspoor et al. (2016) established a typical intra-sentential co-occurrence baseline with competitive performance comparing to a comprehensive machine learning-based system, PKDE4J (Song et al., 2015), on the extraction of relations between human genetic variants and disease on the Variome corpus (Verspoor et al., 2013). The sentential baseline is based on the assumption that the scope of relations is within one sentence, and further assumes that any pair of two entities mentioned in the same sentence and satisfying the type constraints of a given relation, expresses that relation. For example, if two entities with type TimeDescriptor and EndocrineTherapy respectively, the relation TherapyTiming will be extracted. The sentential co-occurrence baseline set a benchmark for relation extraction on Variome corpus.

Abacha and Zweigenbaum (2011) explored semantic rules for the extraction of relations between

medical entities on PubMed Central (PMC) articles using linguistic patterns. They provide an example of implementing an end-to-end relation extraction system, applying named entity recognition in the first stage, then followed by the stage of relation extraction. They define several relation patterns based on medical knowledge, and leveraging the dependency parse tree of sentences in which entities occur. However, the linguistic patterns and rules developed for their corpus likely are not directly applicable to our semantically distinct context of clinical letters.

Machine learning methods vary based on the choice of models and the features considered. In model selection, the multi-type relation extraction task can be assigned to several independent binary classifiers, each making the decision of whether a certain type of relation exist or not. A basic binary classifier such as logistic regression with ridge regularization is capable of performing relation extraction in this scenario. Panyam et al. (2016) used support vector machines (SVM) with a dependency graph kernel to perform relation extraction on two biomedical relation extraction tasks, showing competitive results. Brown Clustering (Brown et al., 1992) is a hierarchical approach to clustering words into classes through maximizing mutual information of bi-grams; it showed competitive performances in many NLP tasks (Turian et al., 2010). Nguyen and Verspoor (2018) implemented a method using character-based word embeddings which can capture unknown words within the context, coupled with CNN and LSTM neural network models. This approach obtained state-of-the-art performance in extracting chemical-disease relations on the BioCreative-V CDR corpus (Li et al., 2016).

For feature engineering, text features can generally be divided into the two categories of lexical features and syntactic features. Typical features used in other relation extraction tasks are summarized here.

- **Bag-of-words (BoW) features** based on white-space delimited tokens, are used in many tasks as a starting point.
- **Bag-of-concepts (BoC) features** (Sahlgren and Cöster, 2004) represent the text in terms of concepts, that is, phrases in the text that correspond to meaningful units. The current methods for generating concepts are

based on techniques such as mutual information (Sahlgren and Cöster, 2004), or through dictionary-based strategies (Funk et al., 2014). For clinical texts, the MetaMap tool (Aronson, 2001) is often used to recognize clinical concepts.

- **Syntactic features** take sentence structure into account. For example, RelEx (Fundel et al., 2006) uses dependency parse trees associated with small numbers of rules in extracting relations from MEDLINE abstract and reaches an overall 80% precision and recall. Approximate Sub-graph Matching (ASM) (Liu et al., 2013) enables sentences to be matched by considering the similarity of the structure of dependency parse subgraphs that connect relevant entities to subgraphs in the training data.
- **Word embeddings** aim to capture word semantics through lower-dimension projections of word contexts and can be used to find word synonyms by measuring cosine similarity between word vectors. There are two widely used approaches to train word embeddings, co-occurrence matrix based methods such as GloVe (Pennington et al., 2014), and learning-based methods using skip-grams (Mikolov et al., 2013b) and CBOW (Mikolov et al., 2013a).

3 Methods

We improve the approaches described above to achieve better efficiency in relation extraction in our context of a narrow knowledge domain with limited data, specific to cancer follow-up treatment.

3.1 Corpus

We consider a previously introduced corpus related to breast cancer follow-up treatment, randomly sampled and manually annotated by two physicians (Pitson et al., 2017). The corpus contains around 1000 sentences and 47,186 tokens. Despite its small size, the corpus is richly annotated with 16 medical named entity types and 16 types of semantic relations linking those entities with over 1,500 relation occurrences. Entities within the corpus are related to clinical therapies, temporal events, diseases, and so on. The annotation of clinical relations includes the associations

between identified entities. For example, in the context "She remains on Arimidex tablets.", "remains on" is a TimeDescriptor, and "Arimidex" is an EndocrineTherapy. The relation TherapyTiming holds between these two entities in this context (e.g., TherapyTiming(TimeDescriptor, EndocrineTherapy)). For conciseness, we use abbreviations to refer to the entity types; hence we will use TD to represent the entity type of TimeDescriptor. While the dataset hasn't been published yet, the full terminology list of entity types can be found in Pitson et al. (2017). To focus exclusively on the relation extraction task, we decouple the named entity recognition task from the relation extraction task by utilizing the gold standard entity annotations from the corpus.

3.2 Method 1: Typed Sentential Co-occurrence

The simplest rule-based approach, given typed named entities, is to extract every pair of entities in a document that satisfies the type constraints of a relation. Such an approach yields high recall but poor precision, due to lack of use of context needed to ensure that a specific relation is expressed as holding between the entities. For example, in our data, only the semantic relation of Toxicity is defined as connecting a Therapy to a ClinicalFinding (expressing that a therapy was found to cause a specific toxic effect) and so it might seem reasonable to assume that a Toxicity event is being expressed in a document where both a Therapy and a ClinicalFinding are mentioned. However, the occurrence of these two entity types together in a document do not strictly indicate an occurrence of Toxicity relation between them; the entities may be connected to other mentioned entities via different relations. Hence assuming a Toxicity relation between them would result in a false positive.

We used intra-sentential constraints as described by (Verspoor et al., 2016) to improve precision by only considering that named entities that co-occur in the same sentence can have valid semantic relations. With this intra-sentential co-occurrence constraint, the relation extraction performance of the sentential baseline achieved strong recall and competitive precision, as well as reasonable overall F-score, on the Variome corpus (Verspoor et al., 2013).

We introduce the approach of Window-

Bounded Co-occurrence (WBC) to explore the impact of relaxing the constraint that the sentence boundary defines the scope of a relation. WBC defines a context window for a relation as the expansion to a base window of the sentence where an entity of the appropriate type occurs, and the adjustment of tokens beyond that sentence within which the related entity appears. We assume the occurrence of an inter-sentential relation relies on the distance of two entities, where distance is defined based on the number of tokens considered beyond the base sentence in which an entity occurs. We introduce a hyperparameter, ρ , to represent the distance of an entity pair in a context, namely the number of tokens allowed in exceeding the single base sentence. $\rho = 0$ denotes the entity pair is intra-sentential; $\rho = x$ denotes the second entity in a pair is x number of tokens away from the base window. If ρ is large than the length of the second context window, then the context window will set to the scope of two sentences by default.

Figure 1 shows an example when $\rho = 5$, WBC allows for the extraction of the semantic relation of TestToAssess across the base window containing the entity of ClinicalFinding, and into the expanded window encompassing the subsequent sentence and containing the TestName entity.




Figure 1: Example of length-awareness sliding window of WBC in TestToAssess relation case

3.3 Method 2: Supervised Binary Classification Approach

We adopt a traditional pipeline as the architecture of the relation extraction system. Each stage is introduced below.

3.3.1 Data Preparation

Considering the semantic variation in the texts, and the small number of examples, training several independent binary classifiers is more robust for mining individual type of semantic relation patterns. Therefore, we transform the original dataset into 16 independent subsets, by grouping

instances by their relation type. An instance consists of a typed entity pair, one or two sentence(s) with the relevant named entities inside as context, a label as an indication of relation occurrence, and a treatment letter id for mapping its position in the original dataset. We remove four relation types with fewer than 10 annotated instances, specifically the relations Intervention, EffectOf, RecurLink and GetOpinion.

We apply context selection for generating instances during data transformation. One instance represents the occurrence of a single semantic relation, containing one relevant entity pair. The context for each instance is the entire raw text of the sentence where the entity pair appears. In cases where more than one entity pair occurs in the same text context, we generate an independent instance for each entity pair. Where cross-sentence relations occur, we concatenate the two sentences containing the relevant entities into a single sentence, structurally indicating that the two sentences are related. In each instance, we replace the two named entity phrases with their type. In the case of overlapping named entities, such as where one named entity partially or completely collides with another entity, both types are retained, adjacent to each other in the text. Other entities mentioned in the context not relevant for the specific relation are left in their original textual form.

We use NLTK (Bird et al., 2009) to perform tokenization and lemmatization to normalize the representation of the text, and strip punctuation. We use the Snowball English stopword list¹ to remove stopwords.

Further details are presented in the feature engineering section below.

3.3.2 Feature Engineering

We implement a set of traditional semantic features and three main feature sets based on ASM, BoC, and sentence embeddings in the sections below.

- **Traditional Semantic Features**

The traditional semantic features includes bag-of-words (count-based), lemmas (base, uninflected form of a noun or verb), algebraic expressions, named entity type (derived from the

¹<http://anoncv.s.postgresql.org/cvsweb.cgi/pgsql/src/backend/snowball/stopwords/>

gold-standard), POS tags and dependency parse based on Stanford CoreNLP (Manning et al., 2014), and a transformation from dependency parse tree to graph using NetworkX (Hagberg et al., 2008) where edges are dependencies and nodes are tokens/labels.

- **ASM features**

The classical ASM measurement was developed by Liu et al. (2013), and was later extended to kernel method by Panyam et al. (2016). The ASM kernel was applied to the chemical induced disease (CID) task (Wei et al., 2015) and SeedeV shared tasks (Chaix et al., 2016). The performance of ASM significantly depends on the result of POS-tagging and dependency tree parsing. All nodes are normalized to their lemmas. Here, Stanford CoreNLP (Manning et al., 2014) is used for POS tagging and dependency tree parsing of the text. The context is split into sentences before dependency parsing is applied on individual sentences.

We produce the ASM features following Panyam et al. (2016). Where the context includes two sentences, a dummy root node is introduced to connect the root nodes of two dependency parse trees. Figure 2 shows an example. After pre-processing, the dependency tree structure is transformed into a graph where nodes are lemmas with their POS tags, edges are dependencies across lemmas within sentence. Then, a shortest path algorithm is applied on the dependency graph to generate flat features. In cases where sentences are very long, processing time is unacceptable, and no ASM features are generated.

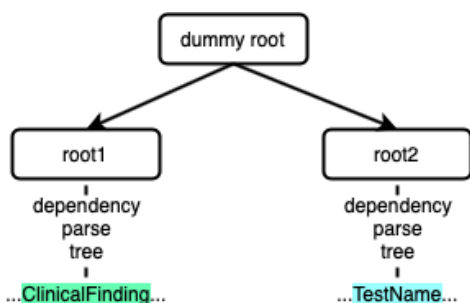


Figure 2: Illustration of concatenation of two sentence parsing results using dummy root in TestToAssess relation setting

- **Bag-of-Concepts features**

Word embeddings are used to capture word

similarity based on shared surrounding context. The size of the surrounding context, known as window size control, varies the representation of word embeddings from more semantic (shorter window size) to more syntactic (longer window size). Synonyms can be identified by identifying two words with similar embeddings, based on cosine similarity measurement. Overfitting can occur for word embeddings, where a training corpus is not large enough or a corpus is limited to a narrow domain of knowledge. Therefore, instead of training word embeddings on our corpus, we use two publicly available pre-trained word embeddings, GloVe (Pennington et al., 2014) and a Wikipedia-PubMed-PMC embedding (Moen and Ananiadou, 2013) to capture more clinically relevant vocabulary. The vocabulary of word embeddings denotes the total number of words that are represented. In our experiment, only the top 20,000 most frequent lemmas are selected. Gensim (Rehurek and Sojka, 2010) is used to find the synonyms of a lemma from the vocabulary by measuring similarity between GloVe word vectors.

We then implement an algorithm for building BoC. Using a word2concept algorithm (see Equation 1), we map a lemma (key) to a concept (value) based on the embedding of the lemma expressed as $E(lemma)$ and the similarity threshold expressed as μ as a tunable hyperparameter.

$$CONCEPT = f(E(lemma), \mu) \quad (1)$$

The algorithm starts by extracting BoW features for each generated instance after data preparation process into a list L . Then, starting from the first lemma w_1 from L , we retrieve its embedding $x_{w_1} = E(w_1)$. We then retrieve a new lemma w_i and its embedding x_{w_i} from the vocabulary V , and calculate the similarity score $S = \cos(x_{w_1}, x_{w_i})$. If $S \geq \mu$, create mappings between $w_1 \rightarrow concept_1$ and $w_i \rightarrow concept_1$. If no w_i satisfies the condition of $S \geq \mu$, then w_1 will be kept in its original form. Next, we move to the second word w_2 in L , check whether w_2 has already been mapped to a concept $concept_*$, and if so, directly create the mapping $w_2 \rightarrow concept_*$. Otherwise, we iterate.

Note that in this model, named entities will effectively be treated as out-of-vocabulary terms,

since they have been mapped to and replaced with the names of the relevant entity types (e.g., “TestName” which is not a token that would be expected to be represented in any pre-trained word embedding model).

- **Sentence embedding features**

Apart from being a tool for finding word synonyms and generating BoC data representation, word embeddings can also be used to obtain sentence embeddings through weighted average pooling. If S denotes a sentence, $E(S)$ denotes the embedding of sentence S . $E(w_1)$ denotes the embedding of the first word w_1 of sentence S , $Score(w_1)$ denotes the TF-IDF score of word w_1 , the sentence embeddings based on weighted average pooling can be expressed as Equation 2. We calculate TF-IDF scores of each word using the original documents, as each instance has an index to its original document id. Out-of-vocabulary tokens and gold standard named entities will be ignored. However, entity information has been considered during that data preparation stage, because the generation of instances takes gold standard entities into account.

$$E(S) = \frac{1}{n} \sum_{i=1}^n E(w_i) \times Score(w_i) \quad (2)$$

3.4 Supervised Learning Models

We build individual binary classifiers for each relation type. We introduce the SVM classifiers and Feed-forward ANN models briefly here.

- **Support Vector Machine** We select the general SVM model (Hearst, 1998) and kernels provided by `scikit-learn` (Pedregosa et al., 2011). For SVM kernels, we integrate ASM kernel as part of feature engineering, to avoid colliding with the use of the linear kernel and the RBF (Radial Basis Function) kernel.
- **Feed-forward ANN** We use Keras (Chollet et al., 2015) to construct a simple feed-forward neural network model with two fully connected layers as shown in Figure 3.

The dimension of each input and the number of hidden units in each layer is the same as the dimension of feature vectors under each type of relation. The activation function for the first dense layer is ReLU, and for the second dense layer is softmax.




Figure 3: Architecture of FNN

3.5 Experiment Design

Considering the dataset is small, we split the transformed dataset into three independent combinations of training, development, and test sets with the ratio of 6:1:3.

In training stage, we train independent models for each specific relation type. We use cross-validation and grid search to tune the hyperparameters of the classifiers.

In prediction stage, the decision of applying sentence-bounded or window-bounded approach is made by setting the size of sliding window. Setting the window size to 0 will apply the sentence-bounded co-occurrence constraint. We choose window sizes of 0, 5, 10 to explore the value of additional context. In supervised binary classification approach, utilizing a similarity threshold of 1 leads to strict use of BoW (word) features, while relaxing the similarity threshold μ of 0.9, 0.8 will generate BoC (concepts). We compare the influence of different word embeddings in generating BoC based on their relation extraction performance on the test set.

Both rule-based approach and supervised binary classification approach will make predictions on the same test set, which allows empirical comparison between rule-based and machine learning approaches.

We compare the impact of increasing data size for BoW and BoC by sub-sampling the training set into nine instance-incremental and non-overlapping sub-sets (combining them into sets representing 10% to 90% of the original training set) and visualize the performance variation. We explore whether word embeddings as a medium for generating BoC are more effective than the direct use of sentence embeddings in cases where the dataset is small and knowledge domain is restricted to the specific domain of breast cancer

treatment. We also explore the combination of BoC and sentence embeddings feature, in order to investigate how best to make integration of them

Finally, we explore the possibility of applying more complex models for analysis of our small and specific knowledge domain dataset. We start from simple linear models including logistic regression and lin-SVM, then apply rbf-SVM, feed-forward ANN on Keras with 32 batch per time, and epochs of 2, 10, 100, and 500 for each relation type.

3.6 Evaluation

We evaluate both overall and per relation type performance. Evaluation of the two approaches is performed on the same test set derived from the data preparation stage. In addition, considering the dataset is small, we calculate the mean score from three independent combinations of training, development and test sets. We evaluate results using micro F_1 -measure since the number of positives and negatives were highly imbalanced across all relation types. We finally evaluate the performance growth over nine sub-sampled training sets of increasing size.

4 Result and Discussion

We present experiment results with micro F_1 measurement. After experimenting with different similarity threshold values to generate BoC features, the best performance is achieved when the threshold is set to 0.9 (results not shown). Word embeddings derived from Wiki-PubMed-PMC outperform GloVe-based embeddings (Table 1). The models using BoC outperform models using BoW as well as ASM features.

As shown in Table 2, the intra-sentential co-occurrence baseline outperforms other approaches which allow boundary expansion. This is because a majority of relations in the corpus are intra-sentential.

A visualization of the growth in performance for both BoW and BoC-based models as training set size increases, over 12 relation types, based on micro F_1 , is shown in Figure 4. The results of BoC in this figure is collected from $\text{lin-SVM}_{\text{Wiki-PubMed-PMC}, \mu = 0.9}$, and BoW is collected from lin-SVM . We find BoC tends to outperform BoW with only a small number of training instances, and also performs better than BoW with incremental training instance.




Figure 4: Performance variation of BoW and BoC($\text{Wiki-PubMed-PMC}, \mu = 0.9$) with increasing data fractions, under linSVM , F_1 measure.

The reason that Wikipedia-PubMed-PMC embeddings (Moen and Ananiadou, 2013) outperforms GloVe (Mikolov et al., 2013a) in the extraction of most relation types (Table 1) is because its training corpus has a more similar domain and vocabulary as our dataset. Therefore, it leads to more relevant models of the distributional semantics of words. On the other hand, the GloVe embeddings are derived from a more general corpus; thus the semantics of domain specific terms in our dataset are not captured.

By observation, most lemmas that map into concepts are digits, time stamps, common verbs and medical terminologies. For example, in $\text{TimeStamp}(\text{TD}, \text{TP})$, drugs with similar effects such as “letrozole” and “anastrozole” map into “ CONCEPT_3 ”; all single digits, ranges from 0-9 map into “ CONCEPT_{26} ”; year tags such as “2011, 2009” map into “ CONCEPT_{56} ”. We consider the cause of differences between the BoW and BoC representations. In the normalization process for BoW, stop-words collected from general knowledge domain are not well-suited to the knowledge domain for our specific task, while limited data does not allow the construction of an appropriate stop-words list from the data. In contrast, BoC models normalize the differences between individual words with shared meaning. While intuitively this should support an improvement over BoW models, we find BoW outperforms BoC when extracting certain relations such as $\text{TestTiming}(\text{TN}, \text{TD})$, $\text{TestToAssess}(\text{TN}, \text{TR})$ with gold standard named entities as arguments. The reason is that synonyms may express slightly different meanings; concept mapping discards such differences and leads to information loss, potentially causing more mis-classifications. The deci-

Feature	LR			SVM			ANN		
	P	R	F ₁	P	R	F ₁	P	R	F ₁
+BoW	0.93	0.91	0.92	0.94	0.92	0.93	0.91	0.91	0.91
+BoC (Wiki-PubMed-PMC)	0.94	0.92	0.93	0.94	0.92	0.93	0.91	0.91	0.91
+BoC (GloVe)	0.93	0.92	0.92	0.94	0.92	0.93	0.91	0.91	0.91
+ASM	0.90	0.85	0.88	0.90	0.86	0.88	0.89	0.89	0.89
+Sentence Embeddings(SEs)	0.89	0.89	0.89	0.90	0.86	0.88	0.88	0.88	0.88
+BoC(Wiki-PubMed-PMC)+SEs	0.92	0.92	0.92	0.94	0.92	0.93	0.91	0.91	0.91

Table 1: Performance of supervised learning models with different features.

Relation type	Count	Intra-sentential co-occ.			BoC(Wiki-PubMed-PMC)		
		$\rho = 0$	$\rho = 5$	$\rho = 10$	LR	SVM	ANN
TherapyTiming(TP,TD)	428	0.84	0.59	0.47	0.78	0.81	0.78
NextReview(Followup,TP)	164	0.90	0.83	0.63	0.86	0.88	0.84
Toxicity(TP,CF/TR)	163	0.91	0.77	0.55	0.85	0.86	0.86
TestTiming(TN,TD/TP)	184	0.90	0.81	0.42	0.96	0.97	0.95
TestFinding(TN,TR)	136	0.76	0.60	0.44	0.82	0.79	0.78
Threat(O,CF/TR)	32	0.85	0.69	0.54	0.95	0.95	0.92
Intervention(TP,YR)	5	0.88	0.65	0.47	-	-	-
EffectOf(Com,CF)	3	0.92	0.62	0.23	-	-	-
Severity(CF,CS)	75	0.61	0.53	0.47	0.52	0.55	0.51
RecurLink(YR,YR/CF)	7	1.0	1.0	0.64	-	-	-
RecurInfer(NR/YR,TR)	51	0.97	0.69	0.43	0.99	0.99	0.98
GetOpinion(Referral,CF/other)	4	0.75	0.75	0.5	-	-	-
Context(Dis,DisCont)	40	0.70	0.63	0.53	0.60	0.41	0.57
TestToAssess(TN,CF/TR)	36	0.76	0.66	0.36	0.92	0.92	0.91
TimeStamp(TD,TP)	221	0.88	0.83	0.50	0.86	0.85	0.83
TimeLink(TP,TP)	20	0.92	0.85	0.45	0.91	0.92	0.90
Overall	1569	0.90	0.73	0.45	0.92	0.93	0.91

Table 2: F₁ score results per relation type of the best performing models.

sion of whether to use the BoC or BoW will depend on the characteristics of particular relation types.

Table 1 also shows the combination feature of BoC and sentence embeddings outperforms sentence embeddings alone, but do not exceed the upper boundary of BoC feature, in which again demonstrating the competitiveness of BoC feature.

Since this corpus is much smaller than other narrow knowledge domain corpus such as CID (Wei et al., 2015) and Seedev shared task (Chaix et al., 2016), the training instances are not enough for the learners to generalize well using syntactic representation. Therefore, the models using ASM kernel (Panyam et al., 2016) do not outperform the simple linear classifiers.

As the results of applying the co-occurrence baseline ($\rho = 0$) shows (Table 2), the semantic relations in this data are strongly concentrated

within a sentence boundary, especially for the relation of RecurLink, with an F₁ of 1.0. The machine learning approaches based on BoC lexical features effectively complement the deficiency of cross-sentence relation extraction.

Lin-SVM outperforms other classifiers in extracting most relations. The feed-forward ANN displays significant over-fitting across all relation types, as the performance decreases when increasing the training epochs. Specifically, with only two training epochs, the performance of ANN is still slightly worse than lin-SVM. The result of lin-SVM present its robustness of avoiding over-fitting compares to feed-forward ANN with BoW, BoC, ASM flat features and sentence embeddings.

5 Conclusion and Future Work

We proposed two ways to perform relation extraction for a narrow knowledge domain, with only small available data set. We implemented a window-based context approach and experiment with determining the best context size for the relation extraction in the rule-based settings. The typical sentential co-occurrence baseline is competitive when most relations are intra-sentential. We implemented a BoC feature engineering method, by leveraging word embeddings as a tool for finding word synonyms and mapping them to concepts. BoC feature outperforms BoW, ASM syntactic feature and sentence embeddings derived by weighted average pooling across word embeddings in small dataset with respect to its significant improvements in micro F_1 score. In addition, it would be expected to show competitive results on other relation extraction tasks where it is useful to generalize specific tokens such as digits or time stamps.

We also explored the performance of models with different level of complexity, such as logistic regression, lin-SVM, rbf-SVM, and a simple feed-forward ANN. The results highlight that strategies to avoid over-fitting must be considered since the number of training instances is limited.

In future work, we will explore a number of directions, including some unsupervised learning approaches. We will test the performance of BoC on other corpora, to explore BoC vs. BoW as a baseline data representation. We will address comparisons between BoC and other word clustering methods such as Brown Clustering (Brown et al., 1992). Finally, the integration of current named entity recognition tools and end-to-end relation extraction, to remove the reliance on gold standard named entity annotations, will also be explored.

References

- Asma Ben Abacha and Pierre Zweigenbaum. 2011. Automatic extraction of semantic relations between medical entities: a rule based approach. *Journal of biomedical semantics*, 2(5):S4.
- Waleed Ammar, Matthew Peters, Chandra Bhagavatula, and Russell Power. 2017. The AI2 system at SemEval-2017 Task 10 (ScienceIE): semi-supervised end-to-end entity and relation extraction. In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*, pages 592–596.
- Alan R Aronson. 2001. Effective mapping of biomedical text to the UMLS Metathesaurus: the MetaMap program. In *Proceedings of the AMIA Symposium*, page 17. American Medical Informatics Association.
- Steven Bird, Ewan Klein, and Edward Loper. 2009. *Natural language processing with Python: analyzing text with the natural language toolkit*. ” O’Reilly Media, Inc.”.
- Peter F Brown, Peter V Desouza, Robert L Mercer, Vincent J Della Pietra, and Jenifer C Lai. 1992. Class-based n-gram models of natural language. *Computational linguistics*, 18(4):467–479.
- Estelle Chaix, Bertrand Dubreucq, Abdelhak Fatihi, Dialekti Valsamou, Robert Bossy, Mouhamadou Ba, Louise Deléger, Pierre Zweigenbaum, Philippe Bessieres, Loic Lepiniec, et al. 2016. Overview of the Regulatory Network of Plant Seed Development (SeeDev) Task at the BioNLP Shared Task 2016. In *Proceedings of the 4th BioNLP Shared Task Workshop*, pages 1–11.
- François Chollet et al. 2015. Keras. <https://keras.io>.
- Katrin Fundel, Robert Küffner, and Ralf Zimmer. 2006. RelExRelation extraction using dependency parse trees. *Bioinformatics*, 23(3):365–371.
- Christopher Funk, William Baumgartner, Benjamin Garcia, Christophe Roeder, Michael Bada, K Cohen, Lawrence Hunter, and Karin Verspoor. 2014. Large-scale biomedical concept recognition: an evaluation of current automatic annotators and their parameters. *BMC Bioinformatics*, 15(1):59.
- Aric Hagberg, Pieter Swart, and Daniel S Chult. 2008. Exploring network structure, dynamics, and function using NetworkX. Technical report, Los Alamos National Lab.(LANL), Los Alamos, NM (United States).
- Marti A. Hearst. 1998. Support Vector Machines. *IEEE Intelligent Systems*, 13(4):18–28.
- Yi Yao Huang and William Yang Wang. 2017. Deep Residual Learning for Weakly-Supervised Relation Extraction. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, EMNLP 2017, Copenhagen, Denmark, September 9-11, 2017*, pages 1803–1807.
- Jiao Li, Yueping Sun, Robin J Johnson, Daniela Sciaky, Chih-Hsuan Wei, Robert Leaman, Allan Peter Davis, Carolyn J Mattingly, Thomas C Wieggers, and Zhiyong Lu. 2016. Biocreative v cdr task corpus: a resource for chemical disease relation extraction. *Database*, 2016.

- Haibin Liu, Lawrence Hunter, Vlado Kešelj, and Karin Verspoor. 2013. Approximate subgraph matching-based literature mining for biomedical events and relations. *PLoS one*, 8(4):e60954.
- Christopher Manning, Mihai Surdeanu, John Bauer, Jenny Finkel, Steven Bethard, and David McClosky. 2014. The Stanford CoreNLP natural language processing toolkit. In *Proceedings of 52nd annual meeting of the association for computational linguistics: system demonstrations*, pages 55–60.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013a. Efficient estimation of word representations in vector space. *CoRR*, abs/1301.3781.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013b. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pages 3111–3119.
- SPFGH Moen and Tapio Salakoski2 Sophia Ananiadou. 2013. Distributional semantics resources for biomedical text processing. In *Proceedings of the 5th International Symposium on Languages in Biology and Medicine, Tokyo, Japan*, pages 39–43.
- Dat Quoc Nguyen and Karin Verspoor. 2018. Convolutional neural networks for chemical-disease relation extraction are improved with character-based word embeddings. In *Proceedings of the BioNLP 2018 workshop*, pages 129–136. Association for Computational Linguistics.
- Thien Huu Nguyen and Ralph Grishman. 2016. Combining Neural Networks and Log-linear Models to Improve Relation Extraction. In *Proceedings of IJCAI Workshop on Deep Learning for Artificial Intelligence (DLAI)*.
- Nagesh C Panyam, Karin Verspoor, Trevor Cohn, and Rao Kotagiri. 2016. ASM Kernel: Graph Kernel using Approximate Subgraph Matching for Relation Extraction. In *Proceedings of the Australasian Language Technology Association Workshop 2016*, pages 65–73.
- F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830.
- Nanyun Peng, Hoifung Poon, Chris Quirk, Kristina Toutanova, and Wen-tau Yih. 2017. Cross-Sentence N-ary Relation Extraction with Graph LSTMs. *Transactions of the Association of Computational Linguistics*, 5(1):101–115.
- 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.
- Graham Pitson, Patricia Banks, Lawrence Cavedon, and Karin Verspoor. 2017. Developing a Manually Annotated Corpus of Clinical Letters for Breast Cancer Patients on Routine Follow-Up. *Studies in health technology and informatics*, 235:196–200.
- Radim Rehurek and Petr Sojka. 2010. Software framework for topic modelling with large corpora. In *In Proceedings of the LREC 2010 Workshop on New Challenges for NLP Frameworks*. Citeseer.
- Magnus Sahlgren and Rickard Cöster. 2004. Using bag-of-concepts to improve the performance of support vector machines in text categorization. In *Proceedings of the 20th international conference on Computational Linguistics*, page 487. Association for Computational Linguistics.
- Min Song, Won Chul Kim, Dahee Lee, Go Eun Heo, and Keun Young Kang. 2015. PKDE4J: Entity and relation extraction for public knowledge discovery. *Journal of Biomedical Informatics*, 57:320 – 332.
- Joseph Turian, Lev 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. Association for Computational Linguistics.
- Karin Verspoor, Antonio Jimeno Yepes, Lawrence Cavedon, Tara McIntosh, Asha Herten-Crabb, Zoë Thomas, and John-Paul Plazzer. 2013. Annotating the biomedical literature for the human variome. *Database*, 2013.
- Karin M Verspoor, Go Eun Heo, Keun Young Kang, and Min Song. 2016. Establishing a baseline for literature mining human genetic variants and their relationships to disease cohorts. *BMC medical informatics and decision making*, 16(1):68.
- Chih-Hsuan Wei, Yifan Peng, Robert Leaman, Allan Peter Davis, Carolyn J Mattingly, Jiao Li, Thomas C Wieggers, and Zhiyong Lu. 2015. Overview of the BioCreative V chemical disease relation (CDR) task. In *Proceedings of the fifth BioCreative challenge evaluation workshop*, pages 154–166.