

An Effective Approach to Biomedical Information Extraction with Limited
Training Data

by

Siddhartha Jonnalagadda

A Dissertation Presented in Partial Fulfillment
of the Requirements for the Degree
Doctor of Philosophy

Approved April 2011 by the
Graduate Supervisory Committee:

Graciela Gonzalez, Chair
Trevor Cohen
Robert Greenes
Douglas Fridsma

ARIZONA STATE UNIVERSITY

May 2011

ABSTRACT

In the current millennium, extensive use of computers and the internet caused an exponential increase in information. Few research areas are as important as information extraction, which primarily involves extracting concepts and the relations between them from free text. Limitations in the size of training data, lack of lexicons and lack of relationship patterns are major factors for poor performance in information extraction. This is because the training data cannot possibly contain all concepts and their synonyms; and it contains only limited examples of relationship patterns between concepts. Creating training data, lexicons and relationship patterns is expensive, especially in the biomedical domain (including clinical notes) because of the depth of domain knowledge required of the curators.

Dictionary-based approaches for concept extraction in this domain are not sufficient to effectively overcome the complexities that arise because of the descriptive nature of human languages. For example, there is a relatively higher amount of abbreviations (not all of them present in lexicons) compared to everyday English text. Sometimes abbreviations are modifiers of an adjective (e.g. CD4-negative) rather than nouns (and hence, not usually considered named entities). There are many chemical names with numbers, commas, hyphens and parentheses (e.g. t(3;3)(q21;q26)), which will be separated by most tokenizers. In addition, partial words are used in place of full words (e.g. up- and downregulate); and some of the words used are highly specialized for the domain. Clinical notes contain peculiar drug names, anatomical nomenclature, other specialized names and phrases that are not standard in everyday English or in published articles (e.g. "I shoulder inj"). State of the art concept extraction systems use machine learning algorithms

to overcome some of these challenges. However, they need a large annotated corpus for every concept class that needs to be extracted.

A novel natural language processing approach to minimize this limitation in concept extraction is proposed here using distributional semantics. Distributional semantics is an emerging field arising from the notion that the meaning or semantics of a piece of text (discourse) depends on the distribution of the elements of that discourse in relation to its surroundings. Distributional information from large unlabeled data is used to automatically create lexicons for the concepts to be tagged, clusters of contextually similar words, and thesauri of distributionally similar words. These automatically generated lexical resources are shown here to be more useful than manually created lexicons for extracting concepts from both literature and narratives. Further, machine learning features based on distributional semantics are shown to improve the accuracy of BANNER, and could be used in other machine learning systems such as cTakes to improve their performance.

In addition, in order to simplify the sentence patterns and facilitate association extraction, a new algorithm using a "shotgun" approach is proposed. The goal of sentence simplification has traditionally been to reduce the grammatical complexity of sentences while retaining the relevant information content and meaning to enable better readability for humans and enhanced processing by parsers. Sentence simplification is shown here to improve the performance of association extraction systems for both biomedical literature and clinical notes. It helps improve the accuracy of protein-protein interaction extraction from the literature and also improves

relationship extraction from clinical notes (such as between medical problems, tests and treatments).

Overall, the two main contributions of this work include the application of sentence simplification to association extraction as described above, and the use of distributional semantics for concept extraction. The proposed work on concept extraction amalgamates for the first time two diverse research areas –distributional semantics and information extraction. This approach renders all the advantages offered in other semi-supervised machine learning systems, and, unlike other proposed semi-supervised approaches, it can be used on top of different basic frameworks and algorithms.

DEDICATION

bījam ímām śarva-bhūtānām ´
viddhi pārtha sanātanam
buddhir buddhimatām asmi
tejas tejasvinām aham

This work belongs to He who is the original seed of all existences, the intelligence of the intelligent, and the brilliance of all brilliant men.

ACKNOWLEDGMENTS

I would like to express my gratitude to my thesis adviser Dr. Graciela Gonzalez for her guidance and patience throughout my degree. Dr. Gonzalez's creativity and inspiration in research serves as a role model in how good research should be done. I would also like to thank Drs. Robert Greenes, Douglas Fridsma and Trevor Cohen for their time in supervising this dissertation writing. In particular, I want to thank Dr. Trevor Cohen for his collaboration in the research. Dr. Greenes as the Chair of the department has been outstandingly supportive during my stay as a student. I am also happy to have had useful discussions about my research with (ex-)lab members – Dr. Luis Tari, Dr. Jorg Hakenberg, Robert Leaman, Ryan Sullivan, Ehsan Emadzadeh and Xiaoxiao Wang. I also acknowledge the support of the funding agencies – Science Foundation Arizona (CAA 0277-08, PI:Gonzalez) and National Library of Medicine (NLM-RDPS-10-121-KP, PI:Gonzalez). I also thank the researchers, developers and supporters of open source software such as Mallet, WEKA, Semantic Vectors, Stanford Parser, Link Grammar Parser that were useful for my research. I am grateful to Jonathon Cadillac Bates who assisted in editing and revising the dissertation.

I am fortunate to have my family and friends who have been with me through the good and rough times. Lastly, my words are not enough to express my gratitude for the almighty who stood with me in the toughest times and gave me the tough times to make me responsible.

TABLE OF CONTENTS

Chapter	Page
List of Tables.....	viii
List of Figures.....	x
Section A. Prologue.....	1
1 Introduction	2
1.1 Overview of the problem and approach.....	2
1.2 Specification of the problem and approach	9
2 Background	13
2.1 Natural Language Processing	14
2.2 Extracting concepts	19
2.2.1 Concept extraction from biomedical literature	20
2.2.2 Concept extraction from clinical notes.....	22
2.3 Extracting relationships.....	23
2.3.1 Protein-Protein interaction extraction	24
2.3.2 Relationship extraction from clinical notes	25
2.4 Distributional semantics	26
2.4.1 Random indexing.....	26
2.4.2 Syntagmatic and Paradigmatic relationships	28
2.4.3 Different models of context.....	29
2.4.4 Hyperspace analogue to language model (HAL)	30
2.4.5 Permutation-based sliding window model	32
2.5 Support Vector Machines (SVM)	33
2.6 K-Means Clustering.....	34
Section B. Named Entity Recognition with Limited Training Data	36
3 Distributional semantics aids clinical concept extraction	37

Chapter	Page
3.1 Methods	38
3.1.1 Choice of unlabeled data	41
3.1.2 Choice of the paradigmatic model.....	41
3.1.3 Automatic word categorization test.....	42
3.1.4 Adding distributional semantic features to machine learning system.....	44
3.2 Results.....	49
3.2.1 Automatic word categorization results.....	49
3.2.2 Concept extraction results	50
3.3 Discussion	53
3.3.1 Comparison with other systems	53
3.3.2 Analysis of the output	55
3.3.3 Limitations.....	55
3.4 Conclusion.....	55
4 Domain knowledge for concept extraction without dictionaries.....	57
4.1 Methods	58
4.1.1 Large unlabeled corpus	58
4.1.2 Directional model.....	59
4.1.3 SVM-based features.....	59
4.1.4 Adding clustering-based features.....	60
4.1.5 N-Nearest neighbors-based features.....	60
4.1.6 Lexicons created by humans	60
4.1.7 Sentence level features	61
4.1.8 Training corpus.....	61
4.1.9 Conditional Random Fields algorithm	62

Chapter	Page
4.2 Results	62
4.3 Conclusion	65
5 A distributional semantics approach to simultaneous recognition of multiple classes of named entities	66
Section C. Discovering More Patterns for Association Extraction	75
6 Dependency-based sentence simplification	76
6.1 Introduction	76
6.2 Corpora	77
6.3 Methods	77
6.3.1 Non-syntactic transformation	78
6.3.2 Syntactic transformation	79
6.3.3 Grammatical correctness of a sentence	81
6.3.4 Overview of rules for simplification	84
6.4 Impact of simplification on the PPI extraction	85
6.5 Conclusion	90
7 Parse tree-based sentence simplification.....	91
7.1 Introduction.....	91
7.2 Methods	93
7.2.1 Noun phrase replacement.....	94
7.2.3 Syntactic Simplification	95
7.3 PPI Extraction Evaluation	98
7.4 Results and Conclusion	99
8 Sentence simplification for clinical relation extraction	101
8.1 Machine learning features.....	102

Chapter	Page
8.2 Choice of baseline for testing the impact of sentence simplification .	
. 105	
8.4 Using sentence simplification as the post-processing step	106
8.5 Impact of sentence simplification	108
8.6 Error analysis.....	112
8.7 Conclusion.....	118
Section D. Overall Discussion and Conclusions	119
9 Discussion	120
9.1 Summary	120
9.2 Implications for the field.....	121
9.3 Limitations.....	124
9.4 Potential applications and Future work.....	125
10 Conclusion.....	129
References	131

LIST OF TABLES

Table		Page
2.1.	Example for HAL matrix computation	31
3.1.	List of features used	45
3.2.	Comparison of concept extraction performance with different models	51
3.3.	The accuracy of baseline vs. system with distributional semantic features	52
4.1.	Clinical concept extraction – comparison of different features	62
4.2.	Clinical concept extraction – comparison of SVM-based features with N-Nearest neighbor-based feature	63
4.3.	Clinical concept extraction – comparison of Clustering-based features with other distributional semantic features	63
4.4.	Protein tagging – Impact of distributional semantic features on BANNER	64
5.1.	SimFind Algorithm	69
5.2.	SimFind results for the GENIA entities	72
6.1.	Example of impact of BioSimplify in improving association extraction	88
6.2.	Results of PIE on selection from AIMed	92
7.1.	Syntactic simplification algorithm	95
7.2.	Syntactic simplification rules with explanation and examples ..	97
7.3.	Results for PIE I	101
7.4.	Results for PIE II	101
8.1.	List of features used for extracting clinical relations	105
8.2.	Performance of three baselines	109

Table	Page
8.3. Comparison of different teams	109
8.4. U Test results for clauses	119
8.5. U Test results for word length	119
8.6. U Test results for word length II	120
8.7. U Test results for clauses II	120
8.8. U Test results for average number of clauses per word	120
8.9. U Test results for number of concepts in a sentence	120

LIST OF FIGURES

Figure	Page
1.1. Overview of techniques used	9
2.1. Example for permuted vector computation	32
3.1. Overall architecture of the clinical concept extraction system	40
3.2. Number of dimensions vs. Performance	46
3.3. Number of seeds vs. Performance	46
3.4. Window radius in positional model vs. Performance	47
3.5. Window radius in directional model vs. Performance	48
3.6. Window radius in positional+basic model vs. Performance	48
3.7. Comparison of different models	49
3.8. Comparison of top clinical concept extraction systems	54
4.1. Overall architecture of the system	58
5.1. SimFind's architecture	68
5.2. Depiction of which entities cause confusion for each entity	75
6.1. Linkages after simplification of the original sentence	81
6.2. Link Grammar's output.....	84
6.3. Link Grammar's another output	84
6.4. Impact of sentence simplification on PPI extraction	90
8.1. Architecure for testing the impact of sentence simplification	110
8.2. iP/R curves of annotations with and without use of sentence simplification	113
8.3. F-score vs. Confidence with and without use of sentence simplification	114
8.4. F-score vs. Confidence with and without use of sentence simplification (for confidence > 0.4)	115

SECTION A: PROLOGUE

This section introduces a challenge in automatic extraction of information from text and the overall approach taken to overcome the challenge.

How does one overcome the inherent ambiguity present in human written text – be it in literature or in narratives, be it in semantics or syntax, be it in identifying concepts or expressing relations between two concepts in a sentence? This is the challenge.

Can a computer learn semantics of words to extract concepts such as proteins, drugs and diseases? Can one help by giving it some context? Can a computer parse the complex syntax of text well enough to extract the associations between the concepts in a sentence? Can one help it by paraphrasing, or simplifying the sentence a bit? An approach is presented to do these things.

This research involves in-depth application of natural language processing, machine learning and also distributional semantics. A background on these topics is presented in Chapter 2.

The methods and evaluations related to automatic incorporation of semantics for extracting concepts are presented in Section B. Section C describes the methods and evaluations related to simplification of syntax for association extraction. Finally, a summary of the entire research and the implications are discussed in Section D.

1. INTRODUCTION

Biomedical information extraction involves automatically extracting biomedical concepts such as genes, treatments, tests and medical problems from unstructured texts, such as clinical notes and biomedical literature, and also extracting the binary associations between them.

1.1 Overview of the problem and approach

Importance of information extraction from clinical narratives: The health care industry in the United States is one of the largest players in the national economy, with trillions of dollars of market share and more than a billion clinical documents per year (Heinze, Morsch, & Holbrook, 2001). Information from these documents could lead to improved health care outcomes by enabling secondary uses of the data such as tracking performance, optimizing resources, appraising treatments and alerting the community about potential post-marketing adverse drug effects. The unstructured format of the non-coded (free text) parts of clinical documents is a barrier for mining valuable data, so advanced natural language processing methods are required. This involves automatically extracting the different concepts and also the relations between the concepts.

Importance of information extraction from biomedical literature: The study of protein-protein interactions and other molecular events is a central tenet of modern translational and genomic research. Publications centering on reports of such atomic events abound. The manual extraction of these events from the literature currently consumes the time of trained curators that deposit them in databases such as DIP, MINT or IntAct. Manual curation, however, despite years of effort has only made a small dent (around 7%) into

the volume of publications believed to report protein-protein interactions (Chitta Baral et al., 2005). Thus, automatic extraction of such facts is a priority for biomedical text mining researchers. This involves first extracting individual concepts such as genes or proteins and then extracting the associations between them.

A limitation in concept extraction: Most basic systems that extract concepts and relations use rule- or pattern-based approaches with the help of lexicons; however, such approaches are not reliable in biomedical information extraction because of the complex ways used while referring to a concept or an association, as elaborated below.

In biomedical literature, the inter-annotator agreement rate for part of speech annotation is approximately 87% (Yuka Tateisi & Jun'ichi Tsujii, 2004). There is a relatively higher amount of usage of abbreviations (not all of them in lexicons) compared to everyday English text. Sometimes abbreviations are modifiers of an adjective (e.g. CD4-negative). Thus, not all abbreviations are nouns and hence named entities. There are many chemical names with numbers, commas, hyphens and parentheses (e.g. 1,25(OH)2D3, beta-(1,3)-glucan, t(3;3)(q21;q26)). Most tokenizers separate these terms. In addition, partial words are used in the place of full words (e.g. alpha- and beta-catenin, up- and downregulate, transcription factor(s)). Moreover, some of the words used are highly specialized for biomedical domain (e.g. p53, c-Abl).

Clinical notes share all the above attributes, thereby making extraction of concepts from both biomedical literature and clinical notes too difficult to be addressed by a simple rule-based or dictionary-based system. Clinical notes contain peculiar drug names, anatomical nomenclature, other specialized

names and phrases that are not standard in everyday English such as urinary incontinence, benign positional vertigo, l shoulder inj, po pain medications, a c5-6 acdf, st changes, resp status and o2 sats. There is also a high incidence of abbreviation usage and many of the abbreviations have a different meaning in other genres of English. For example: ASA (Acetyl Salicylic Acid, not as soon as), NAD (Nicotinamide Adenine Dinucleotide, here not no acute distress) and NC (No Change, not not clear).

Dictionary-based approaches are not sufficient to effectively overcome these complexities that arise because of the descriptive nature of human languages. Supervised machine learning based approaches offer a promising alternative. All supervised machine learning algorithms such as Conditional Random Fields (Lafferty, A. McCallum, & Pereira, 2001) require a training set labeled with concepts. The i2b2/VA NLP shared task corpus is such a corpus of clinical notes labeled with medical problems, tests and treatments. The GENIA corpus (J. D. Kim, Ohta, Y. Tateisi, & J. Tsujii, 2003) is a corpus of biomedical abstracts labeled with 46 biomedical concepts such as proteins, RNA and cell. Since such methods are statistical, a large corpus with as many relevant examples as possible yields an accurate system. However, because of privacy concerns, large corpora of clinical text are not available for research purposes, and thus the annotated sets that are publicly available are even smaller in size. The corpora available for certain semantic types and relations in biomedical literature and clinical notes are also limited because of the domain expertise and effort needed for annotation.

Approach to minimize the limitation in concept extraction: This work proposes to use unannotated data from biomedical texts to design a semi-

supervised machine learning approach for the purpose of extracting concepts. There is an earlier system in the biomedical domain owned by IBM Watson Research Center (Ando, 2007) that used a large amount of unannotated data for winning the BioCreative II gene mention shared task. That system used a computationally expensive machine learning algorithm called Alternating Structure Optimization (ASO). They start with 5 million PubMed abstracts and prune the sentences which do not contain words that appear less than 25 times. C. Liu & Ng (C. Liu & Ng, 2007), who applied ASO to Semantic Role Labeling states, "Some of our experiments are limited by the extensive computing resources required for a fuller exploration. However, we have been unable to use unlabeled data to improve the accuracy." This work is proposing the use of unannotated data through construction of a vector-based similarity model using Random Indexing which is much faster than previous methods (processing the entire Medline corpus takes around 30 minutes using an octa-core Xeon server and 16GB RAM). This approach is thus scalable to huge unannotated corpora and will promote widespread use of unannotated data for the task of clinical concept extraction.

A limitation in association extraction: Sentences in biomedical literature are significantly more complex than those in newspaper articles because of higher average sentence length (Jonnalagadda et al., 2009), higher perplexity measures (Elhadad, 2006), greater lexical density, and increased number of relative clauses and prepositional phrases (Gemoets, 2004). Clinical notes are "not well-written, noise prone, ungrammatical and with much cryptic content" (Y. Wang, 2009). Some of the sentences are too complex to be correctly interpreted by the association extraction system. Presented below are sentences from one of the clinical notes from the i2b2/VA shared task corpus

where a relation is not detected by a state of the art system. The concepts that are associated in the gold standard are highlighted.

*He 'd been having **lower abdominal pain** for approximately the past week , **a symptom** for which he 's been admitted in the past .*

*After discussing this with his PCP , Leon was clear that the patient had had **recurrent DVTs** and ultimately **a PE** and his PCP felt strongly that he required long-term anticoagulation .*

*11. Mesalamine 400 mg Tablet , Delayed Release (E.C.) Sig : Three (3) Tablet , Delayed Release (E.C.) PO TID (3 times a day) as needed for **ulcerative colitis** w/o recent **severe flares** .*

The baseline system was able to detect the associations when the above sentences were respectively simplified as below.

lower abdominal pain is a symptom .

*The patient had had **recurrent DVTs** and ultimately **a PE** .*

*Delayed Release (E.C.) Sig : Three Tablet Delayed Release (E.C.) PO TID as needed for **ulcerative colitis** w/o recent **severe flares** .*

Approach to minimize the limitation in association extraction:

Simplifying sentences and then using the simpler sentences for protein-protein interaction extraction might improve the overall accuracy. Impact on accuracy on extracting relationships in clinical narratives will also be evaluated. This result might also apply for general relationship extraction. If the assumption that sentence simplification aids relationship extraction is true

for both biomedical literature and clinical notes, it would encourage researchers in other fields to use sentence simplification to improve relationship extraction and similar tasks.

In what follows, these proposed approaches are formally hypothesized and the evaluation approaches are specified. The subsequent sections (B and C) detail the experiments.

Publications: This work resulted in the below peer-reviewed research papers.

1. Siddhartha Jonnalagadda, Graciela Gonzalez. BioSimplify: an open source sentence simplification engine to improve recall in automatic biomedical information extraction. In Annual Proceedings of AMIA 2010, Washington D.C., November 13-17, 2010
2. Jörg Hakenberg, Robert Leaman, Nguyen Ha Vo, Siddhartha Jonnalagadda, Ryan Sullivan, Christopher Miller, Luis Tari, Chitta Baral, Graciela Gonzalez. Efficient extraction of protein-protein interactions from full-text articles. IEEE/ACM TCBB. 2010
3. Siddhartha Jonnalagadda, Robert Leaman, Trevor Cohen and Graciela Gonzalez. A Distributional Semantics Approach to Simultaneous Recognition of Multiple Classes of Named Entities. CICLing, LNCS 6008. 2010
4. Siddhartha Jonnalagadda, Graciela Gonzalez. Sentence Simplification Aids Protein-Protein Interaction Extraction. The 3rd International

Symposium on Languages in Biology and Medicine, Jeju Island, South Korea, November 8-10, 2009

5. Siddhartha Jonnalagadda, Luis Tari, Jörg Hakenberg, Chitta Baral and Graciela Gonzalez. Towards Effective Sentence Simplification for Automatic Processing of Biomedical Text. In Proc. of the NAACL-HLT, Boulder, USA, June 2009

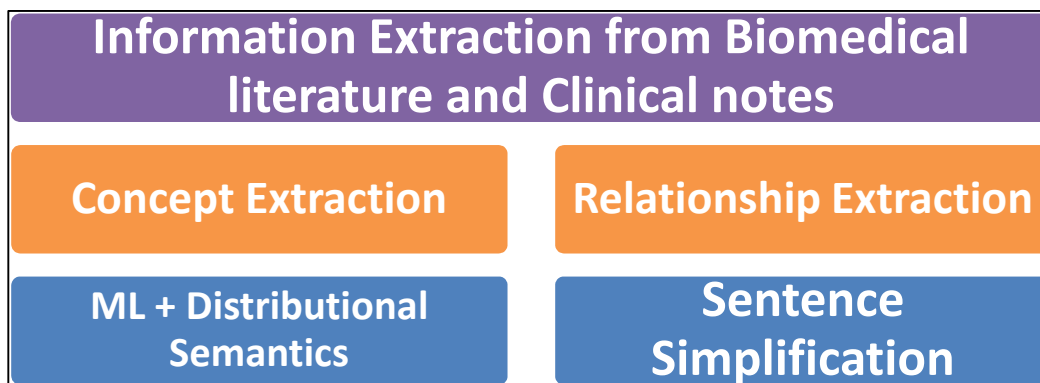
Software published. In addition to the research accomplishments, the software contributions achieved are included below.

1. Updates to BANNER (<http://banner.sourceforge.net>) to use distributional semantic features
2. Minor feature updates to Semantic Vectors package (<http://code.google.com/p/semanticvectors>) to enable using start list.
3. SimFind, a rule-based concept extraction system using distributional semantics (http://www.public.asu.edu/~sjonnal3/SV_NER_src.zip)
4. A sentence simplification system using Link Grammar parser (<http://biosimplify.sourceforge.net>)
5. Another sentence simplification system using Phrase Structure Grammar parser (<http://biosimplify.sourceforge.net>)
6. A clinical relations extraction system (available as part of <http://biosimplify.sourceforge.net>)

1.2 Specification of the problem and approach

This work introduces the notions that distributional semantic methods can be applied for concept extraction in general, and also in biomedical¹ and clinical² domains; and that simplifying the biomedical and clinical sentences can improve extraction of relationships. Figure 1.1 presents the overview of this work.

Figure 1.1: Overview of techniques used



The overall goal is to extract concepts and relationships from biomedical literature and clinical notes. ML refers to a state-of-the-art machine learning system. "ML + Distributional Semantics" refers to using distributional semantics to improve ML. Sentence simplification is technique proposed to improve relationship extraction.

For concept extraction in any domain, the main challenge is the scarcity of annotated examples and that no such large corpus can practically be created without raising privacy concerns. The first hypothesis is that the **distributional information of terms in unannotated corpora can be used to compensate for the limited vocabulary present in a small**

¹ biomedical = scientific biomedical literature

² clinical = electronic medical record text

annotated corpus and allow more accurate concept recognition.

Existing state of the art machine learning systems in both the domains will be extended by adding distributional semantic features, respectively, to extract medical problems and treatments from clinical narratives and to extract proteins from biomedical literature text. Improvement in accuracy after adding distributional semantic features (using the i2b2/VA NLP shared task corpus as a gold standard for the clinical domain and BioCreative shared task corpus for the biomedical domain) would validate the first hypothesis.

For extracting relationships, a small training corpus means that grammatical variants of patterns that express relationships will be limited and might not contain the patterns present in the test set. The second hypothesis is that **through sentence simplification the variants in grammatical patterns present in the test data could be normalized in order to detect the corresponding relationships.** If relationship patterns obtained from the simplified sentences of a training set are used against sentences of a test set simplified to a certain level, then it could improve the accuracy of automatic relationship extraction. Improvement in accuracy after using sentence simplification (as evaluated on the i2b2/VA NLP shared task gold standard for clinical domain and the AIMED corpus for the biomedical literature domain) would validate this approach.

Hypotheses:

Hypothesis 1: Distributional information from unannotated corpora can be used to compensate for limited corpus vocabulary. If similarity derived from such distributional information is used, then it could improve the accuracy of automatic extraction of concepts from text.

Method 1: Adapt regular orthographic and linguistic features known to be useful for concept extraction to create a baseline system and create a new concept extraction system by including the distributional semantic features.

Evaluation 1: The concept extraction system will be tested for extracting gene mentions from biomedical literature and also for extracting medical problems, tests and treatments from clinical notes. The performance will be compared against the baseline system that is equivalent in all respects other than not using distributional semantic features. Improvement in performance in either domain would prove Hypothesis 1. Improvement in performance over both domains would validate the generalizability of Hypothesis 1.

Hypothesis 2: Sentence simplification can compensate for limited grammatical variation in a training corpus. Relationship patterns obtained from the training set might not capture all possible ways of expressing a relationship. If relationship patterns obtained from the training set are used against sentences of a test set simplified to a certain level, then it could improve the accuracy of automatic relationship extraction.

Method 2: Design a system that automatically creates the set of all simpler sentences that can be composed from the original sentence.

Method 3: Use sentence simplification as a post-processing step for relationship extraction in biomedical and clinical domains.

Evaluation 2: Method 3 will be tested for extracting protein-protein interactions from biomedical literature and also for extracting relationships between concepts from clinical text. The performance will be compared against a system that is equivalent in all respects other than not using sentence simplification features. Improvement in performance in either domain would prove Hypothesis 2. Improvement in performance over both domains would validate the generalizability of Hypothesis 2.

2. BACKGROUND

This dissertation is about a novel approach to the task of extracting concepts and relationships from text using distributional semantics and sentence simplification. Distributional semantics is an emerging field arising from the notion that the meaning or semantics of a piece of text (discourse) depends on the distribution of the elements of that discourse in relation to its surroundings. The goal of sentence simplification is to reduce the grammatical complexity of surroundings while retaining the relevant information content and meaning to enable enhanced processing by parsers, and better readability for humans. This research is primarily centered around concept extraction and association extraction, the fundamental blocks of text mining. However, it is easier to understand these with a general overview of natural language processing which is first revised. Sentence simplification, a novel method of syntactic processing of natural language text is also introduced in this section.

The subsequent sections introduce concept extraction and relation extraction in general, and also in the context of biomedical literature and clinical notes. In addition, related literature is also reviewed briefly.

We believe this is the first work that applies distributional semantics for the task of information extraction. Thus, distributional semantics is explained extensively in the next section.

The subsequent sections describe Support Vector Machine and K-Means algorithm which will be used in Chapter 4 to automatically classify and cluster the semantic vectors for the terms.

2.1 Natural Language Processing

Text mining involves information extraction which is the discovery by computer of new, previously unfound information, by automatically extracting information from different written resources (Hearst, 1999). Information extraction primarily constitutes concept extraction, also known as named entity recognition, and relation extraction, also known as association extraction. Natural Language Processing (NLP) is the art of text mining. NLP deals with written text at level of documents, words, grammar, meaning and context.

At the level of **documents**, the entire document is split into individual sentences. While simple rules such as based on punctuation marks and capital letters could help in achieving a reasonable accuracy, more accurate methods employ statistical techniques. For example, Tomanek, Wermter, & Hahn, 2007 (Tomanek, Wermter, & Hahn, 2007) demonstrate that the use of Conditional Random Fields machine learning substantially improves the accuracy.

At the level of **words**, the sentences need to first broken into individual tokens or words. This process is known as tokenization and is executed using rules such as those used by Apache Lucene³. However, there are also machine learning based tokenizers such as that implemented by OpenNLP⁴ using MaxEnt algorithm. Also at the level of **words**, the base form of each word could be found using a process known as lemmatization. The base form of the word, also known as lemma, is independent of its part of speech or a

³http://lucene.apache.org/java/2_3_0/api/org/apache/lucene/analysis/standard/StandardTokenizer.html

⁴ <http://opennlp.sourceforge.net>

morphological variation. Stem, usually used as an approximation for lemma, is that part of the word which is common to all inflected forms. For example, from "deduced", the lemma is "deduce", but the stem is "deduc", because there are words such as "deduction". Computational methods such as Porter Stemmer algorithm (Porter, 1980) are used to strip the suffixes and normalize the inflections automatically.

At the level of **grammar**, the part of speech of the word could be determined using statistical models such as Hidden Markov Model. Several open source tools such as OpenNLP, Lingpipe⁵ and Dragon Toolkit⁶ provide such models customized for specific genres of text. Most part of the speech taggers output the corresponding tags for each word as described in the Penn Tree Bank manual (Marcus, Santorini, & Marcinkiewicz, 1993). For example, the sentence - "A good correlation was found between the grade of Barrett's esophagus dysplasia and high p53 positivity." would be tagged as "A_DT good_JJ correlation_NN was_VBD found_VBN between_IN the_DT grade_NN of_IN Barrett's_NNP esophagus_NN dysplasia_NN and_CC high_JJ p53_NN positivity_NN ._.", where DT is determiner, JJ is adjective, NN is singular noun, VBD is past tense verb, VBN is past participle verb, IN is preposition, NNP is singular proper noun and CC is coordinating conjunction. For detailed information about the Penn Tree Bank format part of speech tags, refer to the documentation at (Marcus et al., 1993) or <ftp://ftp.cis.upenn.edu/pub/treebank/doc/tagguide.ps.gz>. Also at the level of **grammar** is phrase chunking and parsing.

⁵ <http://alias-i.com/lingpipe/demos/tutorial/posTags/read-me.html>

⁶ <http://dragon.ischool.drexel.edu/>

A phrase is a collection of consecutive words that are functional units of a clause and can be primarily classified as noun phrase, verb phrase, prepositional phrase and so on. Phrase chunking, also known as shallow parsing, is the process of identifying all the phrases that are not nested. For example, in the sentence – “John Smith will eat the beans.”, there is a noun phrase - *John Smith*, a verb phrase - *will eat* and a noun phrase - *the beans*.

Parsing is the process of determining the complete grammatical structure of a sentence with respect to a given formal grammar. Some parsers such as Stanford parser (Klein & C. D. Manning, 2003) utilize phrase structure grammar to represent the output. In phrase structure grammar, rules are represented by a tree whose nodes are the different phrases in the sentence and the edges indicate relationship between different phrases. For example, the sentence – “The strongest rain ever recorded in India shut down the financial hub of Mumbai, snapped communication lines, closed airports and forced thousands of people to sleep in their offices or walk home during the night, officials said today.” could be parsed as below:

```
(S
  (S
    (NP
      (NP (DT The) (JJS strongest) (NN rain))
      (VP
        (ADVP (RB ever))
        (VBN recorded)
        (PP (IN in)
          (NP (NNP India))))))
    (VP
      (VP (VBD shut)
        (PRT (RP down))
        (NP
          (NP (DT the) (JJ financial) (NN hub))
          (PP (IN of)
            (NP (NNP Mumbai))))))
      (, ,)
      (VP (VBD snapped)
        (NP (NN communication) (NNS lines)))
```

```

( , , )
(VP (VBD closed)
  (NP (NNS airports)))
(CC and)
(VP (VBD forced)
  (NP
    (NP (NNS thousands))
    (PP (IN of)
      (NP (NNS people)))))
(S
  (VP (TO to)
    (VP
      (VP (VB sleep)
        (PP (IN in)
          (NP (PRP$ their) (NNS offices))))
      (CC or)
      (VP (VB walk)
        (NP (NN home))
        (PP (IN during)
          (NP (DT the) (NN night))))))))))
( , , )
(NP (NNS officials))
(VP (VBD said)
  (NP-TMP (NN today)))
(. .))

```

The parsers representing the output using phrase structure grammar predominantly use Penn Tree Bank format. For detailed information about the Penn Tree Bank format phrase tags, refer to the documentation at (Marcus et al., 1993) or <ftp://ftp.cis.upenn.edu/pub/treebank/doc/manual/root.ps.gz>.

Some of the parsers, such as Link Grammar (Sleator, 1998) utilize dependency grammar where a dependency graph is used to represent the output. In a dependency graph, nodes are the different tokens in the sentence and the edges indicate relationship between the individual tokens. Link grammar has a specific dependency format; however, there are general dependency formats such as Stanford dependency format (Marneffe, MacCartney, & Christopher D. Manning, 2006). For the example sentence, dependencies in Stanford format would be:

```

det(rain-3, The-1)
amod(rain-3, strongest-2)

```

```

nsubj(shut-8, rain-3)
nsubj(snapped-16, rain-3)
nsubj(closed-20, rain-3)
nsubj(forced-23, rain-3)
advmod(recorded-5, ever-4)
partmod(rain-3, recorded-5)
prep_in(recorded-5, India-7)
ccomp(said-40, shut-8)
prt(shut-8, down-9)
det(hub-12, the-10)
amod(hub-12, financial-11)
dobj(shut-8, hub-12)
prep_of(hub-12, Mumbai-14)
conj_and(shut-8, snapped-16)
ccomp(said-40, snapped-16)
nn(lines-18, communication-17)
dobj(snapped-16, lines-18)
conj_and(shut-8, closed-20)
ccomp(said-40, closed-20)
dobj(closed-20, airports-21)
conj_and(shut-8, forced-23)
ccomp(said-40, forced-23)
dobj(forced-23, thousands-24)
prep_of(thousands-24, people-26)
aux(sleep-28, to-27)
xcomp(forced-23, sleep-28)
poss(offices-31, their-30)
prep_in(sleep-28, offices-31)
xcomp(forced-23, walk-33)
conj_or(sleep-28, walk-33)
dobj(walk-33, home-34)
det(night-37, the-36)
prep_during(walk-33, night-37)
nsubj(said-40, officials-39)
tmod(said-40, today-41)

```

Also at the level of **grammar**, sentences could be paraphrased or simplified.

The problem of simplifying long sentences in common English text has been studied before, notably by (Chandrasekar & Srinivas, 1997), (J. Carroll, Minnen, Canning, Devlin, & Tait, 1998) in their Practical Simplification of English text project or PSET and (Siddharthan, 2006). The goal of syntactic simplification is to reduce the grammatical complexity of a text while retaining the relevant information content and meaning to enable better processing by parsers, and better readability for humans. The work of (Chandrasekar & Srinivas, 1997) is aimed at improving processing by parsers, while that of

(Carroll et al, 1998) and (Siddharthan, 2006) seek to enhance readability for humans. In the former work, the goal was to preserve both semantic content and grammatical correctness, but not necessarily cohesiveness. The latter projects aim to create sentences that are shorter, grammatically correct, information-preserving and cohesive (property that the context of a discourse element can be inferred from its precedents). Sentence simplification is also applied in text summarization systems like SumBasic (Vanderwende, Suzuki, Brockett, & Nenkova, 2007) where the focus is to preserve only the important content. This approach, called sentence shortening or sentence compression, does not necessarily preserve semantic content.

At the level of **semantics**, distributional semantics (introduced subsequently) is helpful to automatically infer the similarity of words using unannotated text data. At the level of **pragmatics or context**, section names could provide that information.

2.2 Extracting concepts

The problem of extracting the relevant concepts automatically from text is also known as "Named Entity Recognition and Classification (NERC)", or "Named Entity Recognition (NER)". The problem of Named Entity Recognition and Classification was studied in Computer Science for almost two decades (Rau, Res, Center, & Schenectady, 1991) and there has been significant progress in the field. Earlier attempts were predominantly dictionary or rule-based systems; however, most of the modern systems use supervised machine learning where a system is trained to recognize named entity mentions in text based on specific (and numerous) features associated with the mentions that the system learns from annotated corpora. Thus, machine

learning based methods are dependent on the specific technique or implementation details and the features used for it.

In the former category, generative models (Naïve Bayes Classifier and Hidden Markov Models) and instance-based classifiers (Logistic Regression and Naïve Bayes Classifier) are less accurate for extracting concepts or named entities from text than sequence-based discriminative models like Conditional Random Fields (Minka, 2005; Sutton & A. McCallum, 2007). Advances in machine learning, such as active learning and semi-supervised learning, still need to be thoroughly explored for concept extraction.

Most of the contemporary high-performing tools use non-semantic features such as parts of speech, lemmata, regular expressions, prefixes and n-grams. The high computational cost associated with using deep syntactic and semantic features largely restricted the NERC systems to the orthographic, morphological and shallow syntactic features.

Another common limitation of NERC systems based on machine learning techniques, such as conditional random fields, is the significant computational needs when training on a large, rich corpus like GENIA. While such probabilistic graphical models have been used for multi-class NERC (J. R. Finkel & C. D. Manning, 2009a; A. McCallum & W. Li, 2003; Burr Settles, 2005; Song, E. Kim, G. G. Lee, & Yi, 2004), these are typically trained for less than six entities and are not particularly computationally efficient.

2.2.1 Concept extraction from biomedical literature

State of the art for NER in biomedical domain is at par with that in general domain. Semantic features of varying degrees of sophistication have been used previously in systems like ABNER (Burr Settles, 2005) and the joint

parser and NER tool developed in Stanford by (J. R. Finkel & C. D. Manning, 2009a).

ABNER, a pioneering system for Biomedical NERC using conditional random fields, uses list-look up techniques based on 17 dictionaries that map individual tokens to their semantic types. The dictionaries include those corresponding to Greek letters, amino acids, chemical elements, known viruses, genes, chromosome locations, proteins, cell lines, and abbreviations. However adding these semantic features to the existing word-level features actually had a deleterious effect of decreasing the f-measure by 0.3%.

Finkel's tool uses the distributional similarity model built by (Clark, 2000) to determine the cluster to which a particular token belongs to. The clusters were built apriori from the British National corpus and English Gigaword corpus. The major limitations of this approach are that Clark's model uses only the immediate adjacent tokens to calculate the distributional similarity and that the ambiguity in the semantic type of the token depending upon the larger context is not taken into consideration. It is also reported that because they were able to find only 200 clusters, it resulted in slower inference and no improvement in performance was reported.

On the other hand, most of the state-of-art NERC systems such as BANNER (Leaman & G. Gonzalez, 2008), that has an F-score accuracy of 82% for recognizing protein names, do not use any semantic features including distributional semantic features for want of evidence for scalability and impact on performance.

2.2.2 Concept extraction from clinical notes

The first application of information extraction in clinical domain can be attributed to (Hirschman, Grishman, & Sager, 1976) for converting a corpus of x-ray reports on patients with breast cancer into a structures database using a theory of sublanguage of grammars (Sager, 1975). Medical Language Extraction and Encoding System (MEDLEE (Carol Friedman, 1997)) automatically generates coded information from general clinical notes using rule-based approach in addition to finding modifiers.

MetaMap (Aronson, 2001), an initiative of NLM to map text to UMLS metathesaurus, uses large lexicon and it is still being actively used. A direct application of MetaMap for detecting medical problems (S. Meystre & P. J Haug, 2005) shows that it has an F-score accuracy of 75% for that task and a recent open-source tool HITEX (Zeng et al., 2006) uses MetaMap to map concepts to UMLS strings.

More recently, hospital systems are moving forward in developing open source clinical information extraction systems, such as cTakes (Guergana K Savova, Kipper-Schuler, Buntrock, & Chute, 2008), which uses naïve bayes classifier with lexical and syntactic features to achieve 56% strict F-score accuracy in identifying clinical named entities such as diseases, signs/symptoms, anatomical sites and procedures.

Some other clinical extraction projects include (Chuang, Carol Friedman, & Hripcsak, 2002; Fiszman, W. W. Chapman, Evans, & P J Haug, 1999; P J Haug et al., 1995; P J Haug, Ranum, & Frederick, 1990; Stephane Meystre & Peter J Haug, 2005; Stéphane Meystre & Peter J Haug, 2006; Schadow & McDonald, 2003; Taira & Soderland, 1999; Wicentowski & Sydes, 2008; Hua

Xu & Carol Friedman, 2003). However, as pointed out in a comprehensive review of clinical concept extraction (S. M. Meystre, G K Savova, Kipper-Schuler, & Hurdle, 2008), clinical NLP is currently lagging behind biomedical NLP because of insufficient experience of NLP researchers with clinical text and rarity of annotated corpora. For a systematic review on automatic clinical coding and concept classification systems, refer to (Stanfill, Williams, Fenton, Jenders, & Hersh, 2010)

2.3 Extracting relationships

The concepts extracted in the NER step are constructing blocks for extracting knowledge from text for useful applications like sentiment analysis and relationship extraction. Mathematically, obtaining more information involves finding n-tuples of extracted concepts satisfying a defined function. For example, according to (Hu & B. Liu, 2004), an opinion is the 5-tuple of the target object, feature of object, the opinion holder, the time when it is expressed, and the sentiment value or class of the opinion holder on the feature of the object at that time. According to (Bach & Badaskar, 2007), the problem of relationship extraction is to find the n-tuples for a relation $r(e_1, e_2, \dots, e_n)$. In the biomedical domain, relationship extraction is primarily studied as interaction between proteins. In clinical domain, relationship extraction involves finding associations between medical problems, tests and treatments. In general, relationship extraction is as straightforward as finding whether there is an association or not. There are often attempts to classify the kind of association. In the biomedical domain, the BioNLP shared task (J. D. Kim, Ohta, Pyysalo, Kano, & J. Tsujii, 2009) for event extraction attempted classifying association between proteins or genes as: 1) gene expression, 2) transcription, 3) protein catabolism, 4) localization, 5) binding, 6)

phosphorylation, 7) positive regulation and 8) negative regulation. In the clinical domain, the i2b2/VA NLP shared task⁷ is attempting to classify relationships between tests and problems as whether: 1) test reveals the medical problem, and 2) test conducted for the medical problem; and relationships between treatment and medical problem as whether: 1) treatment improves medical problem, 2) treatment worsens medical problem, 3) treatment causes medical problem, 4) treatment is administered for medical problem, or 5) treatment is not administered for medical problem. Relationship extraction from clinical notes is still in its beginning stages, unlike in the biomedical domain, where there are a number of well-established systems (Ahmed, Chidambaram, Davulcu, & Chitta Baral, 2005; Blaschke, 1999; Bunescu et al., 2005; Katrin Fundel, Kuffner, & Zimmer, 2007; Hakenberg et al., 2010; M. Huang, Zhu, & M. Li, 2006; Lawrence Hunter et al., 2008; Jang et al., 2006; S. Kim et al., 2008; Edward M. Marcotte, Xenarios, & Eisenberg, 2001; Ono, Hishigaki, Tanigami, & Takagi, 2001; Phuong, D. Lee, & LeeKwang, 2003; Schneider, Kaljurand, & Rinaldi, 2009; Temkin & Gilder, 2003; Thomas, Milward, Ouzounis, Pulman, & M. Carroll, 2000; L. Wong, 2001), with some already publicly available (M. Huang et al., 2006; S. Kim et al., 2008; Rebholz-Schuhmann, Arregui, Gaudan, Kirsch, & Jimeno, 2007).

2.3.1 Protein-Protein interaction extraction

The study of protein-protein interactions and other molecular events is a central tenet of modern translational and genomic research. Publications centering on reports of such atomic events abound, and their manual extraction from the literature occupies many trained curators that deposit

⁷ <https://www.i2b2.org/NLP/Relations/>

them in databases such as DIP, MINT, or IntAct. Manual curation, however, despite years of effort, has only made a small dent (calculated at around 7%) into the volume of publications believed to report protein-protein interactions. Automatic extraction of such facts is a foremost priority for biomedical text mining researchers; however, the accuracy of the systems is still inadequate. The BioCreative II effort (Krallinger, Leitner, & Valencia, 2007) to compare the PPI extraction tools of 16 international teams revealed that the best system had an F-score of only 0.30. Although this is the performance for the combined task of relationship extraction and gene normalization, the fact remains that there is still much to be done in this area. While the different reported methods for interaction were reported that use syntactic trees, counting co-occurrences of proteins, lemmatization and so on, the best systems from BioCreative II and BioCreative II.5, respectively (W. Baumgartner et al., 2008) and (Hakenberg et al., 2010) use pattern matching after de-identification of certain words or word patterns to improve matching.

2.3.2 Relationship extraction from clinical notes

Unlike in biomedical domain, relationship extraction is still in beginning stages in the clinical domain with only a couple planned for general use (Angus Roberts, Robert Gaizauskas, Mark Hepple, & Yikun Guo, 2008; Uzuner, Mailoa, Ryan, & Sibanda, 2010). One of the most sophisticated systems is ONYX (Christensen, Harkema, P. J. Haug, Irwin, & W. W. Chapman, 2009) which uses probabilistic context free grammar to extract surface-of-part relationship and location-of-condition relationship with an ERROR of 70%. Most existing systems for finding relationships between clinical concepts use measures of simple co-occurrence (E. S. Chen, G. Hripcsak, H. Xu, Markatou, & C. Friedman, 2008), or variations of it like chi-square (X. Wang, Chused,

Elhadad, C. Friedman, & Markatou, 2008) and point wise mutual information (X. Wang, G. Hripcsak, & C. Friedman, 2009). Some of the recent work is using machine learning techniques (A. Roberts, R. Gaizauskas, M. Hepple, & Y. Guo, 2008) and dependency trees (X. Zhou, Han, Chankai, Prestrud, & Brooks, 2006) to extract relationship between an entity and the modifiers around (not necessarily a specific entity-type) it.

2.4 Distributional semantics

Distributional semantics is an emerging area of research arising from the notion that the semantics of a piece of text (discourse) can be inferred from the distribution of the elements of that discourse in relation to their surroundings. Methods of distributional semantics derive measures of semantic relatedness between terms and text passages from large bodies of unannotated natural language text (for a review, see (T. Cohen & Widdows, 2009)). These measures of relatedness have been shown to correlate with human estimates of relatedness; however, little is known about how best to apply them to support structured prediction tasks (such as parsing), or sequence labeling tasks such as named entity recognition (Turian, Opérationnelle, Ratinov, & Bengio, 2010).

2.4.1 Random indexing

Geometric models of distributional semantics represent each term as a vector in high-dimensional space. Each dimension of this space correspond to a context in the corpus, such as (co-occurrence near) another term or (co-occurrence within) a document. Consequently, a vocabulary of terms is represented as a term-by-context matrix. However, since distributional semantics models constructed based on millions of documents and/or millions of terms would be unmanageable in size, distributional models approaching

corpora of this magnitude tend to reduce dimensionality. Traditional dimensionality reduction techniques, such as Singular Value Decomposition (SVD), are computationally expensive (the commonly utilized algorithm for SVD is cubic in complexity (Trefethen & Bau, 1997)). Recently, Random Indexing (Kanerva, Kristofersson, & Holst, 2000) emerged as a promising alternative to the use of SVD for the dimension reduction step in the generation of term-by-context vectors. Random Indexing and other similar methods are motivated by the Johnson–Lindenstrauss Lemma (W. B. Johnson & Lindenstrauss, 1984) which states that the distance between points in a vector space will be approximately preserved if they are projected into a reduced-dimensional subspace of sufficient dimensionality. Random Indexing scales at a rate that is linear to the size of the data, and has the added advantage that it is not necessary to represent the term-document or term-term matrix in memory: dimension reduction occurs “on-the-fly” as each new context is encountered. This is accomplished by assigning to each document (in term-document models) or term (in sliding window models) a sparse high-dimensional (on the order of 1000) *elemental vector*, a vector comprising of mostly zero elements with a small number (on the order of 10) set to either +1 or -1. These non-zero elements are determined at random, and on account of the sparseness of the vectors this results in a set of vectors that are highly likely to be orthogonal or close-to-orthogonal to one another (that is to say, they are likely to have few non-zero dimensions in common). *Semantic vectors* for each term are then generated as the normalized vector sum of the vectors representation the contexts in which they have occurred.

The major advantages of Random Indexing over established methods employing SVD for dimension reduction are scalability and the capacity for

incremental updates. Hence, for a small size of corpus like i2b2/VA corpus, there are supervised dimensionality reduction techniques such as LDA (Linear Discriminant Analysis) that use SVD computation and could replace Random Indexing for designing kernels based on training set annotations. LDA, while having the limitation of being inapplicable for reducing dimensionality in unlabeled data, is widely applied in NER before and it will be unsurprising if kernels built using LDA perform better than the kernels built using Random Indexing that is an unsupervised dimensionality reduction method and does not exploit the labels of the data. Random indexing is more suitable when applied to a huge unlabeled corpus such as hundreds and thousands of clinical notes or clinical abstracts.

2.4.2 Syntagmatic and Paradigmatic relationships

Recent work (notably (Sahlgren, Holst, & Kanerva, 2008)) in distributional semantics explored the differences between relations extracted depending on the type of context used to construct a model. Sahlgren and other authors (Rapp, 2002) distinguish between two types of relationships between terms captured by distributional models, which they designate as “syntagmatic” and “paradigmatic” relations, terminology derived from the work of Swiss linguist Ferdinand de Saussure (Saussure et al., 1922). If two terms co-occur significantly in the same discourse, they are said to be in syntagmatic relationship. Examples include term pairs such as Barack and Obama, p53 and tumor, APOE and AD, and poliomyelitis and leg. If two terms can substitute for each other in a sentence (i.e. they occur in similar local contexts throughout the corpus), they are said to be in a paradigmatic relationship. For example, (President) Bush and (President) Obama, p53 (gene) and gata1 (gene), AD and SDAT (synonyms), and poliomyelitis and

polio (also synonyms). Since terms in paradigmatic relationship generally do not occur together in the same context, extracting such a relationship typically requires 2nd order (inter-document) analysis, while a 1st order (intra-document) analysis is sufficient to extract syntagmatic relationships. The NER task involves finding words that could conceivably replace the token to be labeled without disturbing syntactic structure, i.e., finding words that are paradigmatically related. In scientific language, domain semantics also determines which terms could replace one another (Z. S. Harris, 2002).

2.4.3 Different models of context

Distributional models derive representations for terms such that terms occurring in similar contexts across a set of documents in a corpus have similar representations. However, the definition of what constitutes a context differs across models. For example, Latent Semantic Analysis (LSA) uses an entire document as a context. In LSA, one builds a term-document matrix and find the semantic representation for words and documents through singular value decomposition (SVD). In contrast, other models such as the Hyperspace Analogue to Language (Lund & Burgess, 1996) use as a context a sliding window that is moved through the text corpus to generate a term-term matrix, T , where $T[i, j]$ is the number of times the word representing the j^{th} column appears in the vicinity of the word representing the i^{th} column. Two words are said to be in vicinity of each other if and only if the number of words separating them is less than an integer parameter known as the sliding window radius. Hence, when a small sliding window is used, this model captures paradigmatic relationships, as two words that can substitute for one another will frequently occur within similar (if not identical) narrow sliding windows. Sahlgren argues that using a small sliding-window rather than an

entire document as a context is better suited to extracting paradigmatic relations, and supports this argument with empirical results. For example, narrow sliding windows are shown to preferentially generate associations between terms of the same part of speech, and between synonymous terms, when compared to term-document based models or wider sliding windows (Sahlgren, 2006).

2.4.4 Hyperspace analogue to language model (HAL)

HAL (Lund & Burgess, 1996) is an extension of the basic sliding window model where one also factors in: a) the direction in which a word occurs with respect to the other, and b) the actual distance between the two words. The cooccurrence strength is weighted inversely according to the number of words separating them. The direction is taken care of by having two columns for each word, with one column representing the number of weighted occurrences to the left and the other column representing the number of weighted occurrences to the right. Table 2.1 shows an example matrix computed for "she was increasingly using her right side" using a window width of six words. Although HAL captures the direction of the words and the distance between them, it does not precisely capture the word order.

The elemental vectors for each term were created using Random Indexing. A random permutation function (π) is generated. A permutation function is a function that maps each element in a collection to a different element or the same element if no two elements are mapped to the same element. This permutation function, π , is applied to each elemental vector to obtain a permutation.

Table 2.1: Example for HAL matrix computation

	she	was	increasingly	using	her	right	side	she'	was'	increasingly'	using'	her'	right'	side'
she	0	6	5	4	3	2	1	0	0	0	0	0	0	0
was	0	0	6	5	4	3	2	6	0	0	0	0	0	0
increasingly	0	0	0	6	5	4	3	5	6	0	0	0	0	0
using	0	0	0	0	6	5	4	4	5	6	0	0	0	0
her	0	0	0	0	0	6	5	3	4	5	6	0	0	0
right	0	0	0	0	0	0	6	2	3	4	5	6	0	0
side	0	0	0	0	0	0	0	1	2	3	4	5	6	0

The numbers in first 7 columns represent the word in the column to the right of the word in the row. The numbers in last 7 columns represent the word in the column to the right of the word in the row. The closer the words, the higher the weight.

To build the contextual vectors for each occurrence of the term, the weighted original elemental vectors of each term occurring to the right in the sliding window are added to the weighted permuted elemental vectors of each term occurring to the left in the sliding window. The term-term matrix such as the one in Table 2.1 is used for this calculation. For example, the contextual vector for the word "using" in the sliding window is calculated as follows:

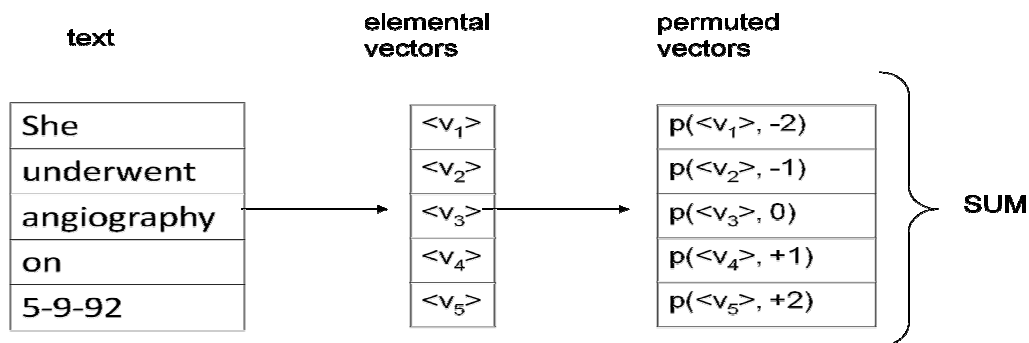
$$c(\text{using}) = 6 * e(\text{her}) + 5 * e(\text{right}) + 4 * (side) + 4 * \pi(e(\text{she})) + 5 * \pi(e(\text{was})) + 6 * \pi(e(\text{increasingly})),$$

where $c(t)$ is the contextual vector for term t in that sliding window, $e(t)$ is the elemental vector of term t and $\pi(u)$ is the permuted vector of the vector u .

The semantic vector for a term is obtained by adding the contextual vectors obtained at each occurrence of that term.

2.4.5 Permutation-based sliding window model

Figure 2.1: Example for permuted vector computation



Semantic term vector for "angiography" = sum (permuted vectors)

$v_1, v_2, v_3, v_4,$ and v_5 respectively are the randomly generated vectors for She, underwent, angiography, on and 5-9-92. $p(v,i)$ would be the vector obtained by shuffling the vector by $|i|$ positions to the right or left depending on the sign of i . The final semantic vector for the term is then generated by adding the vectors generated at its each occurrence.

Sahlgren's permutation-based method (Sahlgren et al., 2008) encodes word-order, thus accounting for the sequential structure of language. The order of the word signifies the grammatical role and thus the meaning of the word. This method is an alternative implementation to the convolution and superposition operations used by BEAGLE (Jones & Mewhort, 2007) (based on Tony Plate's Holographic Reduced Representation (HRR) (Plate, 2002)) to encode word-order information in word spaces. Sahlgren's method captures word information by permutation of vector coordinates, which is a computationally light alternative to BEAGLE's convolution operation. Nevertheless, similar performance improvements have recently been achieved using HRR in the frequency domain (De Vine & Bruza, 2010). To achieve this,

Sahlgren et al. first randomly generate sparse high-dimensional elemental vectors for each term, and use permutation, specifically shuffling of coordinates (shifting of all of the non-zero values of a sparse elemental vector to the left or right according to the relative position of terms) to replace the convolution operator. In this way, a different close-to-orthogonal elemental vector is generated for each term depending on its position within the sliding window. A semantic term vector for each term is then generated as the linear sum of the permuted elemental vectors for each term co-occurring with this term in a sliding window. An example is shown in Figure 2.1. This permutation function is reversible, allowing for construction of order-based queries (such as a query to find terms occurring one position to the right of the term "president").

2.5 Support Vector Machines (SVM)

It was shown previously that distributional semantic measures help in text categorization (Sahlgren & Cöster, 2004), where SVM was used as the machine learning system for benchmarking. This work will be using Conditional Random Fields (CRF) (Lafferty et al., 2001) as the machine learning system as a baseline and SVM will be used for generating features after creating boundaries using the concepts labeled in the training corpus as support vectors. These boundaries help us create the lexicons of concepts. SVM (Cortes & Vapnik, 1995) is designed to draw hyper-planes separating two class regions such that they have a maximum margin of separation. Thus, it is useful for drawing boundaries for regions that are desirable in an N-dimensional space. The optimal hyper planes are achieved by solving the Quadratic Programming equation. There are many implementations to solve the problem, such as Platt's Sequential Minimal Optimization (SMO) (Platt,

1998) algorithm, for fast training of SVMs using sequential minimum optimization and Fan's algorithm (Fan, P. H. Chen, & C. J. Lin, 2005), that uses second order information to achieve fast convergence.

SVM can also be extended to multi-class classification by reducing it to multiple binary classification problems. The outputs from the individual classifiers are ensembled to predict the final class. Most common ways to do this is: a) one of the labels to the rest (one-versus-all) or b) between every pair of classes (one-versus-one). In the first case, there are N binary classifications. The label that has the highest output function is assigned. In the second case, there are $N(N-1)/2$ binary classifications. The label that is assigned in the most classifications is finally assigned.

2.6 K-Means Clustering

Most of the work in concept extraction that uses distributional features use clustering techniques such as Brown's clusters and Clark's clusters (Clark, 2000). A drawback with such methods is that it is not possible to assign cluster to a newly encountered word. Clustering is a kind of unsupervised learning that assigns class labels to data without using prior training information. Algorithms that start with the individual elements and arrive finally at clusters are called agglomerative algorithms, while those that start with all the elements together and eventually arrive at the clusters are called divisive algorithms. They can also be classified as partitional (those that partition the elements) and hierarchical (those that construct the hierarchy of the elements).

K-means, an agglomerative partitional algorithm, groups elements into pre-specified number of clusters. The goal is achieved by minimizing the

aggregate distances of the elements from the centroid. The K means to be found are randomly assigned first and the clusters are then found by assigning elements to the closest centroid. The new centroids are found from the clusters and the process is repeated till the same centroids are found in successive iterations. This algorithm might get stuck at local optima, a known limitation with all algorithms which have multiple local optima. This problem is overcome by running multiple iterations and choosing the most optimal solution. This algorithm will be used to create distributionally similar words.

SECTION B: NAMED ENTITY RECOGNITION WITH LIMITED TRAINING DATA

This section shows how a computer can learn semantics of words to extract concepts such as proteins, drugs and diseases. This is achieved by giving it some context obtained by computing similarity between the words using their distribution in the respective wordspaces. These measures of relatedness have been shown to correlate well with human estimates of relatedness; however, little is known about how best to apply them to support structured prediction tasks (such as parsing), or sequence labeling tasks such as named entity recognition (Turian et al., 2010).

Chapter 3 presents a new method that utilizes distributional semantics to transcend limitations imposed by small training corpora for concept extraction. It also demonstrates how to effectively derive semantics for a particular domain such as clinical narratives. Different distributional semantic models are considered and optimal parameter values are calculated.

Chapter 4 demonstrates the adaptability of this approach of imparting semantic knowledge by using it to improve extraction of gene mentions from the literature.

Chapter 5 shows how distributional semantics can be used to solve multi-label classification by automatically extracting around 50 nested classes from GENIA, a biomedical NER corpus.

3. DISTRIBUTIONAL SEMANTICS AIDS CLINICAL CONCEPT EXTRACTION

Extracting concepts (such as drugs and diagnoses) from clinical narratives constitutes a basic enabling technology that is a prerequisite to unlocking the knowledge within to support more advanced reasoning applications. For example, diagnosis explanation, disease progression modeling and intelligent analysis of the effectiveness of treatment. The recent release of annotated training sets of de-identified clinical notes has contributed to the development and refinement of methods to achieve this end. However, as the annotation process is human-intensive, these training data are limited in their breadth and scope presenting a limitation for systems using supervised machine learning. A novel natural language processing approach is proposed to transcend this limitation by automatically extracting concepts from clinical notes using a combination of supervised machine learning and an unsupervised approach based on distributional statistics.

A sliding-window based variant of the Random Indexing approach is selected as a means to estimate the relatedness between terms in the corpus. This model is used to identify additional features for machine learning. A sequential discriminative classifier is implemented using Conditional Random Fields algorithm for extracting the medical problems, treatments and tests from clinical notes. The additional features added to each term are the N terms that are most similar to it, as measured using the distributional approach.

The evaluation is performed using the i2b2/VA concept extraction corpus. It is observed that incorporating features based on the distributional statistics of

terms across a large unlabeled corpus significantly aids concept extraction. The micro-averaged f-measure for exact match increased from 80.31% to 81.75% and the micro-averaged inexact f-measure increased from 89.73% to 90.94%. The significance of the improvement is 100.0% as measured using bootstrap resampling technique.

Distributional semantic features significantly improve the performance of concept extraction.

3.1 Methods

Harris' sublanguage theory (Z. Harris, 1991) suggests that the constraints of language in a specialized domain include domain semantics in addition to purely grammatical constraints, suggesting the applicability of a semantic grammar in which domain semantics are an integral component of syntax. In highly constrained genres of texts like biomedical literature and clinical notes (C. Friedman, Kra, & Rzhetsky, 2002), there are inherent inequalities of the likelihood of certain patterns of words occurring. As a corollary, high probability combinations convey implicit information. For example, sentences such as *Patient X received drug Y* are more common than sentences such as *Doctor X received drug Y*. Thus, a sentence pattern *A received B* could be decoded to assume that A might refer to patient and B might refer to treatment. Thus, the concept extraction task can be seen as finding a class of terms that could conceivably replace the token to be labeled without disturbing the surrounding syntactic structure. Consequently, substituting an antibiotic for a protein would still disturb syntax even though they are both nouns. Hence, the concept extraction task relates to the finding of terms that are paradigmatically related.

For concept extraction using supervised machine learning methods, the main challenges are the scarcity of annotated examples, the level of expertise required to generate these examples, and the fact that no such large corpus of clinical text can practically be shared without raising privacy concerns. The proposed hypotheses was that information derived from the distribution of terms in a larger, unannotated corpus can be used to compensate for the limited vocabulary present in a small annotated corpus and allow more accurate concept recognition. To evaluate this hypothesis, distributional semantic features are added to commonly used features to extract medical problems and treatments from clinical narratives. Improvement in accuracy after adding distributional semantic features (using i2b2/VA NLP shared task corpus as gold standard) would validate the utility of these additional features.

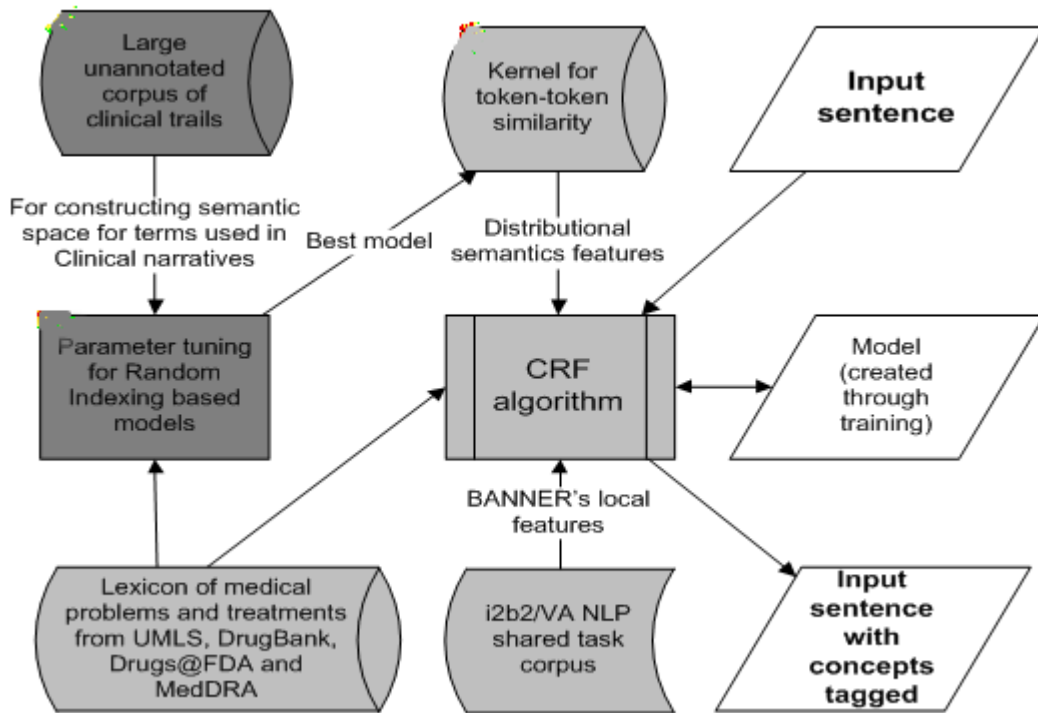
Building a sliding-window based model involves:

- Constructing elemental term vectors of pre-determined dimension N and seed S , where $N-2*S$ dimensions are zeroes, S dimensions are $+1$ s, and S dimensions are -1 s. To ensure that these elemental vectors have a high probability of being mutually orthogonal, or close-to-orthogonal, $S \ll N$.
- Computing the semantic vector representation of the term based on the terms surrounding each occurrence of the term.

This process is performed using the open source Semantic Vectors package (Widdows & T. Cohen, 2010), which supports the generation of basic,

direction-aware and permutation-based sliding window models using Random Indexing for dimension reduction.

Figure 3.1: Overall architecture of the clinical concept extraction system



The design of the system to identify treatment, test and medical problem entities includes components that are needed only offline (dark gray), both offline and online (light gray) and only online (white).

Features from an existing system BANNER (Leaman & G. Gonzalez, 2008) are adapted to create the baseline. BANNER uses the Conditional Random Fields algorithm implemented in MALLET (McCallum, 2002). Any concept extraction system using similar algorithm can easily adapt the proposed set-up. Our system will be trained on i2b2/VA NLP training corpus using a) sentence-level features; b) lexicon from UMLS, DrugBank, Drugs@FDA and

MedDRA; and c) distributional semantic features based on the large unannotated corpus. Creating the distributional semantic features involves tuning the parameters and finding the optimal model among different distributional semantic models. The best model would be used to create a kernel for adding distributional semantic features. After the systems are trained and tuned offline and a CRF machine learning model is created, processing a sentence from a clinical record involves the same set of NLP features used for training.

3.1.1 Choice of unlabeled data

All the documents (349 annotated clinical notes and 827 unannotated clinical notes) provided as part of the i2b2/VA training set and all the MEDLINE abstracts that are indexed as pertaining to "clinical trials" are used as unlabeled data. There are 447k MEDLINE abstracts indexed as pertaining to "clinical trials".

3.1.2 Choice of the paradigmatic model

An inverted index of the chosen unlabeled data and Semantic Vectors software (Widdows & T. Cohen, 2010) is used to find the vector model of each term in the abstracts that appears at least twice. The different paradigmatic model algorithms considered are the positional model, the directional model and the positional+basic model. The directional model approximates the HAL model by permuting the terms occurring before the target term in one specific way and the terms occurring after the target term in another specific way. The positional+basic model is the combination of the positional and basic sliding window models where the semantic vector of a term is the vector sum of the corresponding semantic vectors in the positional model and the basic model.

The parameters of the vector model that could be customized are: a) dimensions of the final vector space, b) seed length, c) half-window size of the models. The following methodology was used for deciding the optimal parameter values:

- Experiment with dimensions of sizes 100, 500, 1000, 1500 and 2000; and chose the optimal number of dimensions.
- Experiment with seeds of sizes 5, 10, 15, and 20; and chose the optimal seed length.
- Experiment with half-windows of sizes 1, 2, 3, 4, 5, 6, 7, 8, 9 and 10 separately for each paradigmatic model algorithm; and chose the optimal half-window size.

For efficient testing, the below automatic word categorization test before annotating all the clinical concepts was first performed.

3.1.3 Automatic word categorization test

The automatic word categorization test is designed to test the ability to correctly categorize terms. To model paradigmatic relationships, the model that performed the best in the experiments in the automatic word categorization test is chosen. This test is similar to part-of-speech test performed by (Sahlgren, 2006) to study the differences between different wordspace models. Sahlgren studied the extent to which different distributional models retrieved nearest neighboring terms that shared a part of speech with a cue term. He used two settings which he referred to as strict and lax. Both of them are variants of k-nearest neighbor algorithm – the first

one has $k=1$ and the second one has $k=10$ – where all the neighbors are treated equally.

All the UMLS single word phrases are divided into 4 categories: problem, treatment, test and none. A term is problem if its UMLS semantic type is pathologic functions, disease or syndrome, mental or behavioral dysfunction, cell or molecular dysfunction, congenital abnormality, acquired abnormality, injury or poisoning, anatomic abnormality, neoplastic process, virus/bacterium, sign or symptom. A term is treatment if its UMLS semantic type is therapeutic or preventive procedure, medical device, steroid, pharmacologic substance, biomedical or dental material, antibiotic, clinical drug, or drug delivery device. A term is test if its semantic type is laboratory procedure or diagnostic procedure. These definitions are taken from the 2010 i2b2/va concept extraction corpus annotation guidelines:

<https://www.i2b2.org/NLP/Relations/assets/Concept%20Annotation%20Guideline.pdf>

This test is performed for different paradigmatic models proposed in the previous section. For each model, 10 nearest neighbors are found for each term. The term itself is removed from the nearest neighbor list. Each of the 4 categories had a weight equal to the sum of the cosines of the neighbors that belong to the category. The category with the highest score is assigned to the term. This automatically determined word category is compared to the actual word category. For example, thoracotomy has 6 treatment neighbors ($.79 + .86 + .82 + .82 + .79 + .81 = 4.89$) and 3 test neighbors ($.82 + .80 + .82=2.44$). Hence, thoracotomy is correctly classified as treatment. Thus, it is known whether the assignment is a true positive, false positive, false

negative, or true negative with respect to each category. In each case, the individual recall and precision as well as micro-averaged recall, precision, F-score and accuracy was recorded. This automatic word categorization test is also a variant of k-nearest neighbor algorithm with k=9.

3.1.4 Adding distributional semantic features to machine learning system

The kernel K over the terms in the i2b2/VA corpus is then constructed, where

$$K(w_1, w_2) = \begin{cases} \text{cosine of the semantic vector representations of } w_1 \text{ and } w_2, \\ \text{if both terms exist in the } \frac{i2b2}{VA} \text{ corpus} \\ \\ \text{zero,} \\ \text{if either of the terms do not exist in the corpus} \end{cases}$$

The kernel is used to automatically build a thesaurus of terms. Each entry in the thesaurus consists of a token from the i2b2/VA NLP corpus and N most similar terms (also from the i2b2/VA corpus) based on the distributional semantics. Computing the kernel (K) scales linear in the number of dimensions of each vector and quadratic in the number of terms in the i2b2/VA corpus. Computing the thesaurus (with a pre-determined number of similar terms from the kernel) scales linear to the number of terms. On the whole, it scales linearly with respect to the number of dimensions of each vector and quadratic in the number of term vectors. Using the pre-computed kernel instead of directly computing thesaurus online saves computing the cosines (O(N) time complexity) during the construction of present thesaurus and more importantly for multiple values of parameter N, the number of similar terms.


The 1st order Conditional Random Fields (CRF) algorithm, as implemented by MALLET, is used. The time complexity of CRF algorithm is $O(L^2 * N * M * F * I)$,

where L is the number of labels, N is the number sequences (sentences), M is the average length of the sequences, and F is the average number of the features, and I the number of iterations. It is observed (Leaman & G. Gonzalez, 2008) that the accuracy is almost the same for all label types such as – IO, IOB and IOBEW, where I stands for labeling a token to be Inside, O for Outside, B for Beginning, E for End and W for Within. The IO notation is chosen for labeling to minimize time complexity. Thus, 4 labels - Iproblem, Itest, Itreatment and O are used since 3 concepts are annotated. In addition to all the features used in a state of the art Biomedical NER system known as BANNER (Leaman & G. Gonzalez, 2008), three additional feature types based on: 1) thesaurus, 2) vector representation of the token, and 3) dictionaries are added (see Table 3.1).

Table 3.1: List of features used


Feature name	Type	Description
Dictionary	Semantic	UMLS, DrugBank, Drugs@FDA and MedDRA
Distributional	Semantic	Distributional thesaurus and Dimensions of terms (word embeddings)
Section	Pragmatic	Name of the section in which the sentence appears
Part of speech	Syntactic	Part of speech of the token in the sentence
Others	Lexical	Lower case token, Lemma, Prefixes, Suffixes, n-grams, Matching patterns such as beginning with a capital, etc.

Figure 3.2: Number of dimensions vs. Performance




The performance initially increases and remains almost constant after 1000 dimensions

Figure 3.3: Number of seeds vs. Performance




The performance does not change with the number of seeds used in the model

Figure 3.4: Window radius in Positional model vs. Performance




The performance is the highest at the window radius of 2 corroborating with results from similar experiments by Sahlgren

Figure 3.5: Window radius in directional model vs. Performance




The performance is the highest at the window radius of 6.

Figure 3.6: Window radius in Positional+Basic model vs. Performance



The performance is the highest at the window radius of 4.

Figure 3.7: Comparison of different models



Positional+Basic and Directional model perform slightly better than the Positional model

3.2 Results

3.2.1 Automatic word categorization results

Figures 3.2 – 3.7 show how the experiment is conducted. Both F-score ($\frac{2*TP}{(2*TP+FP+FN)}$) and accuracy ($\frac{(TP+TN)}{(TP+TN+FP+FN)}$) is measured while fine-tuning each parameter. The advantage of using F-score is that in the NER (concept extraction) task, true positives are more important than true negatives. The advantage of using accuracy is that it is invariant to label switching which means the value does not change when a label is changed from positive to negative. Thus, the goal was to optimize both F-score and accuracy. There were situations when one of the two was almost constant and the other varied. In such situations, the setting where the latter's value is higher is chosen. The positional vector model is used as baseline to decide the optimal number of dimensions and seed length which are the parameters for the Random Indexing algorithm. These parameters are

general to the application of Random Indexing for the clinical text. Once these parameters were fine-tuned, they were used for benchmarking different sliding window algorithms amongst each other.

The accuracy and F-score increased rapidly in the beginning as the number of dimensions increased. However, the accuracy and F-score beyond 1000 dimensions were almost constant. Since time complexity is linearly proportionate to the number of dimensions, 2000 dimensions were used. The accuracy and F-score were constant as the number of seeds was increased. Therefore, 5 seeds were used.

The window radius is then changed for the three models – the positional permutation vector model, the directional permutation vector model, and the positional+basic model. As expected, the performance varies differently in each case with increase in the window radius. The positional model had the highest accuracy and F-score at a window radius of 2. This corroborates similar experiments by Sahlgren (Sahlgren, 2006; Sahlgren et al., 2008). For directional model, the performance increased till the window radius was increased to 6 and it decreased after that. The positional+basic model performed better than the positional model and the best accuracy and F-score was achieved at the window radius of 4.

3.2.2 Concept extraction results

Two of the parameters involved in use of the most optimum model found in the previous section are: a) the number of most similar words to consider, and b) the minimum similarity between the similar word and the original token. The optimal number of most similar words to use was found to be 20. With a threshold on the cosine similarity between the similar words and the

original token, the performance was slightly decreasing. Consequently, the threshold was set to 0, i.e., all the 20 nearest words were used irrespective of their cosine value. Table 3.2 compares the results with using different models with these settings.

Table 3.2: Comparison of concept extraction performance with different models

#	Model	R_{exact}	P_{exact}	F_{exact}	R_{inexact}	P_{inexact}	F_{inexact}
1	Non Distributional semantic features	<i>78.10</i>	<i>82.65</i>	<i>80.31</i>	<i>88.66</i>	<i>90.84</i>	<i>89.74</i>
2	1 + Distributional Semantic features using i2b2 corpus	78.71	83.20	80.89	89.02	91.27	90.13
3	2 + Positional model based features using Medline Clinical Trials	79.85	83.74	81.75	89.84	91.61	90.72
4	2 + Directional model based features using Medline Clinical Trials	79.92	83.50	81.67	90.26	91.63	90.94
5	2 + Positional+basic model based features using Medline Clinical Trials	79.81	83.55	81.64	90.05	91.63	90.83

The highest scores are in bold and the lowest scores are in italics. R_{exact} is the micro-averaged recall for exact match. P_{exact} is the micro-averaged precision for exact match. F_{exact} is the micro-averaged F-score for exact match. R_{inexact} is the micro-averaged recall for inexact match. P_{inexact} is the micro-averaged precision for inexact match. F_{inexact} is the micro-averaged F-score for inexact match.

The results for the system trained on the competition training corpus and tested on the testing corpus, but without any distributional semantic features and that of the best system with distributional semantic features (direction-based, 2000 dimensions, 5 seeds and window radius = 6) are in the below Table 3.3.

Table 3.3: The accuracy of baseline (subscript 1) vs. system with distributional semantic features (subscript 2)

Type	R ₁	R ₂	P ₁	P ₂	F ₁	F ₂
Concept exact span	80.4	82.0	85.1	85.6	82.7	83.7
Class exact span	78.1	79.9	82.7	83.5	80.3	81.7
Problem exact span	80.0	81.9	83.6	84.7	81.8	83.2
Treatment exact span	80.2	81.6	85.8	86.1	82.9	83.8
Test exact span	81.3	82.4	86.6	86.5	83.4	84.4
Problem matching class	78.6	80.5	81.6	83.0	80.0	81.7
Treatment matching class	77.0	79.1	82.9	83.4	79.8	81.2
Test matching class	78.6	80.0	84.1	84.4	81.3	82.1
Concept inexact span	89.0	90.5	94.2	94.6	91.6	92.5
Class inexact span	88.9	90.3	90.8	91.6	89.7	90.9
Problem inexact span	89.7	91.5	93.8	94.7	91.7	93.1
Treatment inexact span	88.3	89.6	94.4	94.3	91.2	91.9
Test inexact span	88.8	90.1	94.7	94.8	91.6	92.4
Problem inexact span matching class	87.6	89.5	90.9	92.2	89.2	90.8
Treatment inexact span matching class	84.2	86.2	90.6	91.0	87.3	88.5
Test inexact span matching class	85.1	86.7	91.0	91.5	87.9	89.1

R=Recall, P=Precision, F=F-score.

It is encouraging to see that addition of distributional semantic features increases both the recall and precision in all cases. Bootstrap Resampling (Noreen, 1989) with 1000 repetitions on the test corpus showed the improvement after adding the distributional semantic features is highly

significant (confidence=100%). However, the improvement because of adding local distributional semantic based features after adding distributional semantic features from clinical trials was insignificant (confidence=56.6%). Hence, it might also be concluded that addition of distributional semantic features using a large unannotated corpus is sufficient and need not be supplemented by distributional semantic features from smaller corpora. It was also found that using the dimensions of vector representations as tokens decreases the accuracy of the system. This could be due to dimensionality curse which exponentially increases the size of the training set needed based on the number of features used.


3.3 Discussion

3.3.1 Comparison with other systems

The results section showed that the system significantly improves with respect to the accuracy of extracting concepts after adding distributional semantic features. However, the practical significance of the improvement is best determined by how much it contributes to improving state of the art. Thus, the performance of the system was compared against the top systems in the i2b2/VA concept extraction task 2010 in which 22 international teams have submitted multiple runs as shown in Figure 3.8. Our baseline system (without distributional semantic features) ranked 7th both in F-scores measured using exact match and inexact match. It is to be noted that best system uses the relatively novel technique of using semi-HMM trained through Collin's structured perceptron. Systems with ranks 2-5 use proprietary (non-public) components. System 6 uses Conditional Random Fields algorithm similar to us. After adding distributional semantic features, the F-score as measured by inexact match is better than systems 3-6. The F-

score as measured by exact match is almost the same as systems 4-5 and better than system 6. This suggests that adding distributional semantic features to a supervised machine learning system significantly improves the state of art.

Figure 3.8: Comparison of top clinical concept extraction systems



NRC = National Research Council Canada. deBruijn et al.; VU = Vanderbilt Univ. Xu et al.; Erasmus = Erasmus Univ. Kors et al.; SCAI = SCAI. Hofmann et al.; Sydney = Univ. Sydney. Patrick et al.; George = Georgetown Univ. Liu et al.; **ASUbefore = local baseline; ASUafter = local baseline after adding distributional semantic features.**

In addition, this system is the only one among this list that adds the semi-supervised component on the top of an existing supervised system. Our semi-supervised component (distributional semantic features) can be adapted with other machine learning algorithms such as Hidden Markov Models (used in the

best system submission for i2b2/VA concept extraction task) and Naïve Bayes (used in cTakes system) to achieve incremental improvements. Our method is adaptable to other domains within Biomedical Informatics such as biomedical literature and health news and also to other domains. Another important feature of the method is that it allows secondary use of clinical notes.

3.3.2 Analysis of the output

The output of the baseline is compared with the output of the final version after adding distributional semantic features. The final version has more true positives [2202 vs. 1382] (hence, less false negatives) and less false positives [2535 vs. 2803]. All these contributed to an increase in both precision and recall. This phenomenon is found for all the concepts aggregately and also for each concept separately. The annotations that are different in both versions are analyzed. There are, of course, more new terms in the true positives [360 vs. 283] and true negatives [496 vs. 439] newly found in the final version. This corroborates the hypothesis that the performance increases because of being able to give a semantic representation for the newly found terms, and also because of a more accurate semantic representation of the existing terms. It was also found that approximately half of extra true positives found [1068 out of 2202] contained terms that were predicted to belong to the class based on the weighted K-NN algorithm (refer to "Automatic word categorization" subsection in the methods section). This correlation is interesting, especially because the distributional semantic features are only a fraction of all the features used.

3.3.3 Limitations

The unlabeled data currently used are clinical trials from Medline. The improvement might have been higher if one used clinical notes from hospital

records since that would also capture semantics of non standard words used in the clinical notes. Future work should include comparison of the features with other models such as Collobert and Weston's word embeddings (Collobert & Weston, 2008), Brown's word clusters (Brown, Desouza, Mercer, Pietra, & Lai, 1992) and Clark's word clusters (Clark, 2000).

3.4 Conclusion

In sum, the results indicate that distributional semantic features aid clinical concept extraction. The next step would be to use clinical records as unlabeled data to expect more increase in accuracy, since there would be more words of the same type as problem, treatment and test, and similar context around different words.

4. DOMAIN KNOWLEDGE FOR CONCEPT EXTRACTION WITHOUT DICTIONARIES

One of the most time-consuming tasks of an NLP researcher who is trying to adapt their concept extraction system for a different domain is to create (compile if such a resource is available) and curate the lexicons. The other time consuming tasks are: a) annotating a corpus and b) feature selection. Annotation of a corpus is necessary for any supervised machine learning system and active learning (B. Settles, 2004) might decrease the amount of corpus that needs to be annotated. Although feature selection is helpful in decreasing the number of useful features, it was not found to be helpful in increasing accuracy (Klinger & Friedrich, 2009) with state of the art algorithms such as Conditional Random Fields algorithm.

Thus, lexicons are primary resources for domain adaptation in addition to an annotated corpus. It is found that adding such domain specific information improves the performance of named entity recognition, although in a few systems an increase in performance was not found (Burr Settles, 2005; Guodong Zhou, Dan Shen, Jie Zhang, Jian Su, & S. Tan, 2005). The focus of this chapter is to explore the possibility of creating lexical resources automatically for the purpose of incorporating domain knowledge using distributional information inferred from unlabeled corpora. To this end three types of lexical resources are considered: a) lexicons for the concepts to be tagged, b) clusters of distributionally similar words, and c) thesaurus of distributionally similar words. These resources would be constructed using high dimensional vector space created from large text from that domain, such as Medline abstracts for genes and unlabeled clinical notes from hospitals for clinical concepts.

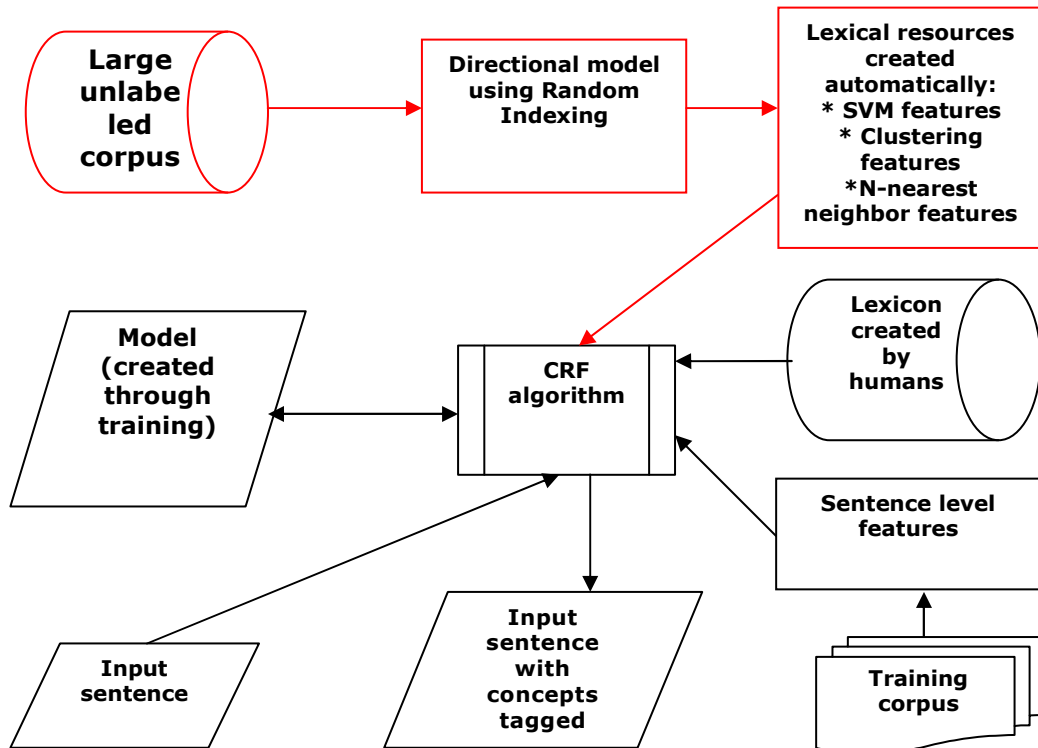
4.1 Methods

The architecture of the system is in Figure 4.1. This design of the concept extraction system is to identify treatment, test and medical problem entities in clinical notes and identify proteins in biomedical literature. Our system is trained using: a) sentence-level features using training corpus, b) lexicon created, compiled and curated by humans for each domain, and c) distributional semantic features based on a large unannotated corpus of domain-relevant text. After the systems are trained, a CRF-based machine learning model is created to process input sentences using the same set of NLP features. The output will be the sentence with the concepts tagged.

4.1.1 Large unlabeled corpus

For clinical concept extraction, 447K Medline abstracts that are indexed as

Figure 4.1: Overall architecture of the system



The design of the system to identify concepts using machine learning and distributional semantics. Components related to distributional semantics are outlined in red.

pertaining to clinical trials are used as the unlabeled corpus. For protein concept extraction, all the Medline abstracts of the 2008 baseline⁸ are used as the large unlabeled corpus.

4.1.2 Directional model

Directional model takes into account the direction in which a word occurs with respect to another by having two columns for each word. One column represents the number of occurrences to the left and the other column represents the number of occurrences to the right. The semantic vector for a term is obtained by adding the contextual vectors obtained at each occurrence of the term. The process of Random Indexing and directional model takes around a minute for thousands of documents and a few hours for the entire PubMed.

4.1.3 SVM-based features

SVM algorithm (Platt, 1998) implemented in WEKA (Witten, Frank, & Hall, 2011) is used to construct the support vectors. The respective training corpus is used to obtain the terms and their labels. For example, in clinical concept extraction, each term can either belong to one or more of the classes: problem, treatment, test or neither of these. The objective in creating the categories for different terms is to obtain samples of regions in the distributional hyper-space that contain terms from the desired (problem, treatment, test and neither) semantic types so that the boundaries can be drawn using SVM. Each term is labeled as Iproblem, Itest, Itreatment or Inone. To remove ambiguity, terms that belong to more than one category are neglected. Each term has a representation in the distributional hyper-

⁸

http://www.nlm.nih.gov/archive//20090811/bsd/licensee/2008_stats/2008_LO.html

space of 2000 dimensions. The parameter of each dimension is a feature (or attribute) for the SVM. Six binary SVM classifiers are generated - one for each pair of the category that can be used to predict the class of any term. During the execution of the training and testing phase of the machine learning algorithm, the class predicted by the SVM classifiers for each term is used as a feature for that term.

The features for protein tagging are constructed similarly.

4.1.4 Adding clustering-based features

K-means clustering algorithm implemented in WEKA is used to group the terms in the training corpus into 200 clusters using distributional semantic vectors. The cluster identifier assigned to the target term is used as a feature for the Conditional Random Fields based system for concept extraction. This feature is similar to the Clark's automatically created clusters (Clark, 2000) used by (J. R. Finkel & C. D. Manning, 2009b), where the same number of clusters are used.

4.1.5 N-Nearest neighbors-based features

Cosine similarity of vectors is used to find the 20-nearest terms for each term. These nearest terms are used as features for the respective target term.

4.1.6 Lexicons created by humans

For the clinical concept extraction, UMLS, DrugBank, Drugs@FDA and MedDRA are used to create dictionaries for medical problems, treatments and tests. The guidelines of the i2b2/VA NLP concept extraction task (<https://www.i2b2.org/NLP/Relations/assets/Concept%20Annotation%20Guideline.pdf>) are followed to identify the corresponding UMLS semantic types for

each of the 3 concepts. The other three resources are used to add more terms to the existing lexicon. In an extensive evaluation on the nature of the resources by (Gurulingappa, Klinger, Hofmann-Apitius, & Fluck, 2010), UMLS and MedDRA were found to be the best resources for extracting information about medical problems among several other resources.

For the protein tagging, the 344K word lexicon compiled and curated for BANNER (Leaman & G. Gonzalez, 2008), one of the best protein tagging system (Kabiljo, Clegg, & Shepherd, 2009) was used. This lexicon can be downloaded from <https://banner.svn.sourceforge.net/svnroot/banner/trunk/dict/single.txt>. This lexicon was constructed using the BioCreative II gene normalization training set (Morgan et al., 2008).

4.1.7 Sentence level features

Several local orthographic and linguistic features such as lower-case token, lemma, prefixes, suffixes, n-grams, matching patterns such as beginning with a capital and part of speech were also used.

4.1.8 Training corpus

For clinical concept extraction, the fourth i2b2/VA NLP shared task corpus (<https://www.i2b2.org/NLP/Relations>) for extracting concepts of the classes - problems, treatments and tests is used. The corpus contains 349 clinical notes as training data and 477 clinical notes as testing data.

For protein tagging, the BioCreative II Gene Mention task (Wilbur, L. Smith, & Tanabe, 2007) corpus (http://biocreative.sourceforge.net/biocreative_2_gm.html) is used. The corpus contains 15000 training set sentences and 5000 testing set sentences.

4.1.9 Conditional Random Fields algorithm

BANNER that internally uses Conditional Random Fields (CRF) algorithm as implemented by MALLET (AK McCallum, 2002) is being used. The time complexity of CRF algorithm is $O(L^2 * N * M * F * I)$, where L is the number of labels, N is the number sequences (sentences), M is the average length of the sequences, F is the average number of the features and I the number of iterations. The number of iterations is much higher during the training than during testing. 4 labels were used for the clinical concept extraction - Iproblem, Itest, Itreatment and O, respectively for inside a problem, test, treatment and none. For protein tagging, the default IOB notation (I=inside, O=outside, B=beginning) was used.

4.2 Results

Table 4.1: Clinical concept extraction – comparison of different features

Setting	Exact F	Inexact F	Exact Increase	Inexact Increase
ASU_Basic	77.9	88.2		
ASU_noDict	79.4	89.2	1.5	1
ASU_Dict	80.3	89.7	2.4	1.5
ASU_noDict+DistSem	81.4	90.8	3.5	2.6
ASU_Dict+DistSem	81.9	91	4	2.8

Table 4.1 shows the F-score for exact match increased over the baseline (ASU_Basic - which uses only the words themselves as features) after adding different local features (ASU_noDict) by 1.5%. After dictionary features are added (ASU_Dict), it increased by 2.4%. On the other hand, if only distributional semantic features were added without using dictionary (ASU_noDict+DistSem), it would have increased by 3.5%. It increases only by 0.5% more, if the dictionary features were also used along with

distributional semantic features (ASU_Dict+DistSem). The F-score for inexact match follows a similar pattern.

Table 4.2: Clinical concept extraction – comparison of SVM-based features with N-Nearest Neighbors-based features

Setting	Exact F	Inexact F	Exact Increase	Inexact Increase
ASU_Dict	80.3	89.7		
ASU_Dict+SVM	80.6	90	0.3	0.3
ASU_Dict+NN	81.7	90.9	1.4	1.2
ASU_Dict+NN+SVM	81.9	91	1.6	1.3

The F-score for exact match increased by 0.3% after adding SVM-based features, whereas it increased by 1.4% after adding N-nearest neighbor-based features. The F-score slightly increased further with the use of both features. The F-score for inexact match follows a similar pattern.

Table 4.3: Clinical concept extraction – Comparison of Clustering-based features with other distributional semantic features

Setting	Exact F	Inexact F	Exact Increase	Inexact Increase
ASU_Dict	80.3	89.7		
ASU_Dict+CL	80.8	90.1	0.5	0.4
ASU_Dict+NN+SVM	81.9	91	1.6	1.3
ASU_Dict+NN+SVM+CL	81.7	90.9	1.4	1.2

The F-score for exact match increased by 0.5% after adding clustering-based features, whereas it increased by 1.6% after adding N-nearest neighbor-based and SVM-based features. The F-score slightly decreased with the use of both of the features. The F-score for inexact match follows a similar pattern.

The evaluations for clinical concept extraction reveal that the distributional semantic features are better than traditional dictionary features. The accuracy further increases when both manually created dictionaries and distributional semantic feature types are used. This shows distributional semantic features

not only supplement manually developed dictionaries, but also complement them. Further, the N-nearest neighbor features are better than SVM-based features and Clustering-based features.

Table 4.4: Protein tagging - Impact of distributional semantic features on BANNER

Rank	Setting	Precision	Recall	F-score	Significance
1	rank 1 system	88.48	85.97	87.21	6-11
2	rank 2 system	89.30	84.49	86.83	8-11
3	BANNER_Dict + DistSem	88.25	85.12	86.66	8-11
4	rank 3 system	84.93	88.28	86.57	8-11
5	BANNER_noDict+DistSem	87.95	85.06	86.48	10-11
6	rank 4 system	87.27	85.41	86.33	10-11
7	rank 5 system	85.77	86.80	86.28	10-11
8	rank 6 system	82.71	89.32	85.89	10-11
9	BANNER_Dict	86.41	84.55	85.47	-
10	rank 7 system	86.97	82.55	84.70	-
11	BANNER_noDict	85.63	83.10	84.35	-

In Table 4.4, the performance of BANNER with distributional semantic features (3) and without distributional semantic features (9) is compared with the top ranking systems in the most recent gene mention task at the BioCreative shared tasks. Each system has an F-score that has a statistically significant comparison ($p < 0.05$) with the teams indicated in the Significance column. The significance is estimated using Table 1 in Biocreative II gene mention task report (Wilbur et al., 2007). The performance of BANNER with distributional semantic features and no dictionary features is better than BANNER with dictionary features and no distributional semantic features. This

demonstrates again that distributional semantic features (that are generated automatically) are more useful than dictionary features (that are usually compiled and curated manually).

The improvement after adding the distributional semantic features to BANNER is significant. The absolute ranking of BANNER with respect to other systems in the competition improves from 8 to 3. However, the rank 1 system does not have significantly better F-score than BANNER with distributional semantic features. Distributional semantic features are again more useful than dictionary features. The advantage of the features is that they are independent of the machine learning system used and can be used to further improve the performance of forthcoming algorithms. For example, the variants of CRF algorithm proposed in (C.-N. Hsu et al., 2008; Y. Li, H. Lin, & Yang, 2009).

4.3 Conclusion

Our evaluations using clinical notes and biomedical literature validate that distributional semantic features are useful to automatically obtain domain information irrespective of the domain and can reduce the need to create, compile and curate dictionaries.

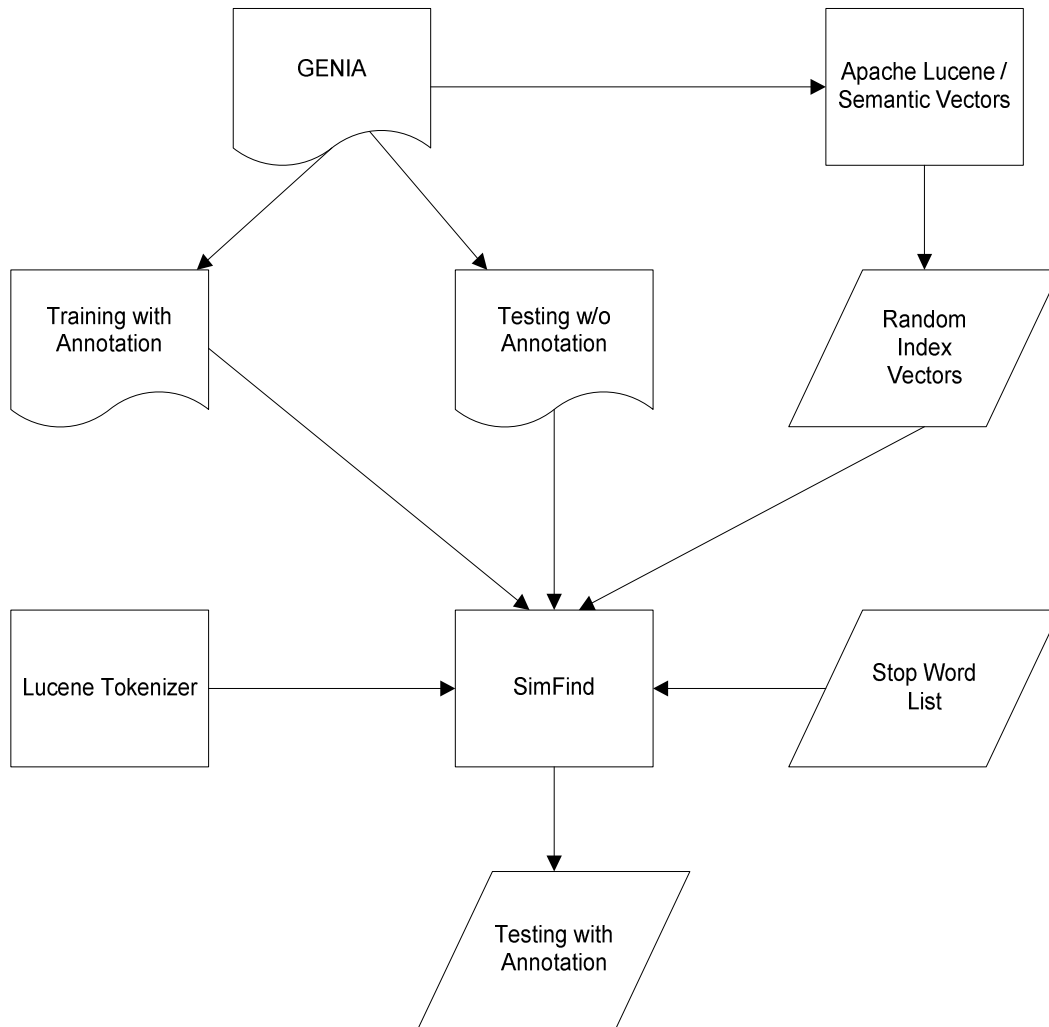
5. A DISTRIBUTIONAL SEMANTICS APPROACH TO SIMULTANEOUS RECOGNITION OF MULTIPLE CLASSES OF NAMED ENTITIES

Recent advances in distributional semantics allow us to efficiently create paradigmatic models that encode word order. Sahlgren *et. al.*'s permutation-based variant of the Random Indexing model was used to create a system to simultaneously recognize multiple entity classes mentioned in natural language, which is validated on the GENIA corpus which has annotations for 46 biomedical entity classes and supports nested entities. Using distributional semantic features only, it achieves an overall micro-averaged F-measure of 67.3% based on fragment matching with performance ranging from 7.4% for "DNA substructure" to 80.7% for "Bioentity".

The architecture is a 2-stage pipeline as shown in Figure 5.1. The entire corpus is broken into more than 18000 documents, each of which contains a unique sentence of the GENIA corpus. A Lucene index is built for this set of documents. The term and document vectors are built using the Semantic Vectors package.

The corpus is divided into two halves – one half is the training set and the other half is the test set. The Lucene tokenizer breaks the sentence into tokens and the SimFind algorithm is used to find the token in the training set that is most similar to the target token. The entity class of the similar token is then assigned to the target token. SimFind therefore takes into consideration the surrounding context when determining the semantic type of each token while previous methods considered the semantic type of the token independent of the context.

Figure 5.1: SimFind's system architecture



SimFind is a rule-based system that uses distributional semantic information to annotate biomedical concepts in text.

In this research, the estimates of similarity provided by Random Indexing were utilized for two purposes. Firstly, as token labels are context-dependent, the 100 most similar sentences from the training set that are similar to the vector sum of the terms belonging to the target sentence were found. Next, the first token from the similar sentences that is same as the target token or similar to it was found. Thus, the SimFind algorithm takes into account all the other tokens present in the sentence and it also does not assume that the

target token is present in the training set. The pseudo code for the algorithm is explained in Table 5.1.

Table 5.1: SimFind Algorithm

<pre> SimFind(targetToken, Line){ List simSentences = getSimilarSentences(Line,100); List goldenTokenLabel = getTokenLabels(simSentences); STEP1: FOREACH (goldenTokenLabel) IF (goldenTokenLabel has targetToken as token) return goldenTokenLabel; STEP2: IF (token IN STOPLIST) return <token,NONE>; terms = 1; STEP3: terms *= 10; <equivTokens,simIndex>=getSimWords(targetTok en,terms); FOREACH (equivToken) FOREACH (goldenTokenLabel) IF (goldenTokenLabel has targetToken as token) return goldenTokenLabel; IF (simIndex>0.5) goto STEP3; return <token, NONE>; } </pre>	<p>The SimFind function is the core method which retrieves the sentences which share the same context as the target sentence and for each token in the target sentence. The algorithm first checks for the earliest appearance of the target token in the set of similar sentences arranged in the order of similarity. The next step should be to search for the presence of the tokens similar to the target tokens. However, to minimize the total time taken, eliminate the tokens which appear too frequently in common English and hence are highly unlikely to be part of a biomedical entity.</p>
<pre> getSimilarSentences(line, numberOfResults){ break line into tokens using Lucene Tokenizer; form query vector by computing the sum of tokens; search for similar documents in Random Index Vectors; set the number of results to be numberOfResults; listOfSimilarSentences are the sentences from the training set which correspond to these documents; return listOfSimilarSentences; } </pre>	<p>The getSimilarSentences function is responsible for finding the specified number of sentences from the training set that are similar to the vector sum of the terms belonging to a given sentence in the test set.</p>
<pre> getSimWords(targetToken, count){ form query vector as targetToken; search for similar terms in Random Index Vectors; </pre>	<p>The getSimWords function is responsible for fetching the tokens in the corpus similar to</p>

<pre> set the number of results to be count; return list of similar terms; } </pre>	a given token.
<pre> getTokenLabels(simSentences){ for each token in simSentences find the label from the xml annotation and add <token, label> to listOfTokenLabels; return listOfTokenLabels; } </pre>	The getTokenLabels function is used to get the semantic type of the tokens in an annotated sentence.

The algorithm used the list of 421 stop words created by Fox (Fox, 1989) from the Brown corpus to improve the efficiency of SimFind. These stop words were selected to be maximally efficient and effective in filtering the semantically neutral words. There are several options for the labeling model. The simplest is the IO model, which indicates whether the token is **i**nside an entity or **o**utside an entity, which is the model employed for this work. Another possible model is IOB, where each token is labeled to be either **b**eginning of an entity, **i**nside an entity, or **o**utside an entity. There are also systems using IOBEW model which in addition label for the **e**nd of the entity and one-**w**ord entity. In the recent evaluation of BANNER (Leaman & G. Gonzalez, 2008), an NERC tool, which used a corpus annotated with biomedical entities for recognizing gene entities, the difference between the performances of these three labeling models was found to be less than 1%. Each token can belong to multiple semantic types as GENIA annotates nested entities. Since there are 36 entity classes at leaf level (J. D. Kim et al., 2003), there are 2^{36} possible types of labels with the IO model.

NERC systems are typically evaluated using exact matching, which requires that both the left and right boundary match exactly. For many applications, however, determining the exact boundary is not necessary and it is sufficient to determine whether the sentence contains an entity of the specified type or

not and its approximate location. Thus, recently more realistic matching techniques like core-term matching and fragment matching are becoming prominent (Tsai et al., 2006). In core-term matching, the system's annotated named entity must contain a core term of the named entity in the gold standard. This requires that every annotation in the corpus should also mention which is the core-term. In a corpus like GENIA with around 100,000 entities, this would require an excessive amount of annotation resources. In fragment matching, each token is treated separately. This provides a measure of how much fraction of the entity is matched and is thus more realistic than conventional exact matching and loose partial matching.

Since it is shown that 5x2 validation is statistically more powerful than 10x1 validation (Dietterich, 1998), 5x2 validation was used. Table 5.2 presents the precision, recall and F-score measures achieved by the system on all the entities annotated in the GENIA corpus except the biologically irrelevant entities like Protein N/A, DNA N/A, and those with insufficient data. It also provides the count of true positives, false positives, and false negatives in each case. For most of the entities, this work is one of the first to use GENIA for evaluation. Hence the results also serve as comparison for all NERC systems that would be evaluated using GENIA corpus. In addition, for each entity the F-score for a hypothetical system that randomly assigns a positive or negative in the ratio of the number of actual true or false cases, respectively was calculated. If a corpus has t tokens belonging to a particular entity class and f tokens not belonging to that entity class, a system which randomly assigns tokens to that class in proportion to the known proportion of positives and negatives would result in both precision and recall approximating $t/(t+f)$. The F-score of the random system would therefore

also be approximately $t/(t+f)$, which serves as a quantitative estimate the difficulty of NERC task for a specific entity class. This quantity is labeled Random F-score in Table 5.2.

Table 5.2: SimFind results for the GENIA entities

<u>Entity</u>	<u>Precision (%)</u>	<u>Recall (%)</u>	<u>F score</u>	<u>TP</u>	<u>FP</u>	<u>FN</u>	<u>Random F-score</u>
Bio-entity	78.9	82.5	80.7	60479	16131	12868	26.22
Substance	77.0	79.6	78.3	46587	13796	11976	20.92
Organic compound	77.0	79.5	78.2	46244	13792	11944	20.82
Compound	77.0	79.5	78.2	46382	13822	11980	20.82
Amino acid	69.4	71.1	70.3	27331	11917	11091	13.69
Protein	69.2	71.0	70.1	26692	11864	10890	13.37
Lipid	66.1	67.0	66.5	1243	637	618	0.66
Virus	65.6	67.3	66.4	1641	862	797	0.86
Source	61.4	66.2	63.7	10434	6554	5326	5.62
Atom	62.0	60.2	61.1	150	92	99	0.10
Nucleotide	57.0	64.4	60.5	114	86	63	0.05
Other organic compound	62.1	58.9	60.5	2105	1285	1470	1.28
Protein molecule	60.9	59.4	60.1	13194	8456	9014	7.87
Organism	59.6	58.8	59.2	2085	1412	1460	1.23
Amino acid monomer	61.2	53.1	56.9	256	162	225	0.15
Mono Cell	69.4	45.9	55.3	100	44	118	0.10
Inorganic	57.1	53.3	55.1	97	73	85	0.66
Natural source	52.0	57.6	54.6	6017	5560	4427	3.76
Carbohydrate	63.2	45.7	53.1	43	25	51	0.05
Nucleic acid	51.0	54.2	52.6	9181	8803	7752	6.05
DNA	48.3	52.6	50.4	7829	8366	7051	5.31
DNA domain or region	44.4	48.5	46.4	5889	7362	6253	4.35
Cell type	42.7	50.7	46.3	3046	4089	2968	2.14
Cell line	44.0	44.9	44.5	2375	3022	2912	1.87
Artificial source	43.9	44.3	44.1	2442	3118	3074	1.98
RNA	47.0	41.2	43.9	707	797	1011	0.61
Body part	39.6	45.0	42.1	148	226	181	0.10
Other name	42.6	40.0	41.3	11591	15645	17367	10.31
Protein domain or region	41.9	38.8	40.2	606	842	958	0.56
Protein complex	40.4	40.1	40.2	1509	2226	2256	1.33

Protein family or group	34.0	39.8	36.7	3761	7289	5697	3.36
Peptide	41.9	32.7	36.7	149	207	307	0.15
RNA molecule	36.5	36.7	36.6	453	783	777	0.40
Multi Cell	36.5	34.7	35.6	315	547	593	0.30
Polynucleotide	44.9	27.0	33.7	62	76	168	0.10
Protein subunit	31.2	31.1	31.2	379	834	838	0.40
DNA molecule	24.6	22.6	23.6	174	533	597	0.30
Tissue	22.8	23.7	23.3	151	510	486	0.20
RNA family or group	28.3	15.7	20.2	67	170	360	0.15
Protein substructure	12.2	16.5	14.0	21	151	106	0.05
DNA family or group	12.8	14.5	13.6	270	1844	1588	0.66
DNA substructure	6.1	9.3	7.4	11	170	107	0.05
Overall Score	66.3	68.4	67.3	34233 0	17418 0	15790 9	

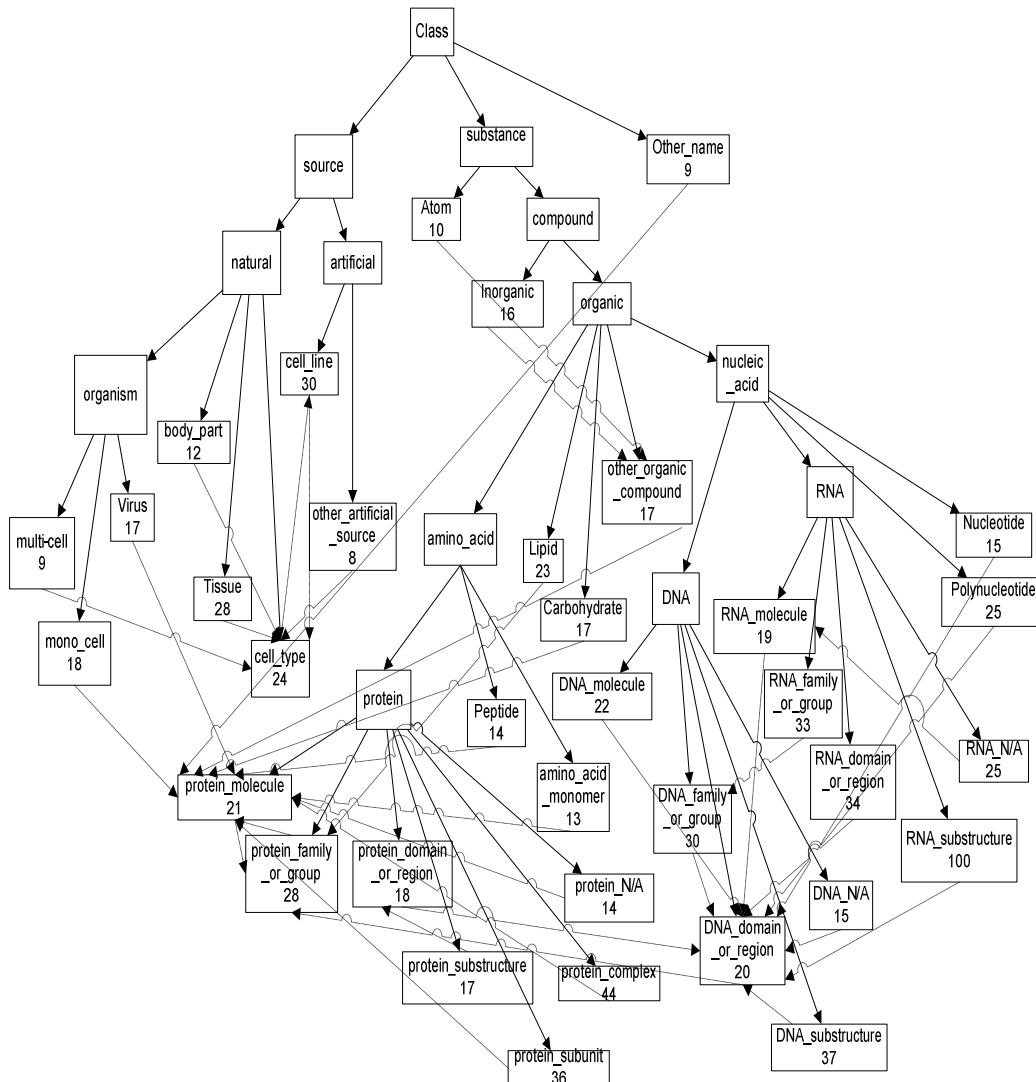
The entities in Table 5.2 are arranged in descending order of their F-scores based on the system. It is encouraging to see that more than half of the entity classes have an F-score greater than 50% just based on distributional semantic features and also the huge differences between F-score and Random F-score. The system also has a considerable good overall micro-averaged F-score of 67.3% which is calculated by adding the respective true positives, false positives and false negatives of each entity class. It took around 5 minutes to build the semantic vectors from the documents belonging to the GENIA corpus and around 3 hours to produce results for the testing set which constitutes more than 9000 sentences. This suggests that this framework is scalable and could have significant impact on the precision and recall of a more complex system.

There have been several attempts (Byrne, 2007; J. R. Finkel & C. D. Manning, 2009b; Gu, 2006; Marquez, Villarejo, Marti, & Taule, 2007; D. Shen, J. Zhang, G. Zhou, J. Su, & C. L. Tan, 2003; G. D. Zhou, 2006; G. Zhou, J.

Zhang, J. Su, D. Shen, & C. Tan, 2004) using machine learning to find nested entities in text with many entities like GENIA corpus. These systems limit themselves to work for less than six entities at a time due to computational cost. Since this framework also recognizes nested entities, it can be used to provide features that can be quickly calculated and can replace the features with slower inference. Figure 5.2 attempted to analyze the errors made by the system by characterizing the confusion between the entity classes. An entity class A is said to have confused entity class B, if and only if either at least one of the false positives of B actually belongs to A or at least one of the false negatives of B was considered by the system to belong to A. The confusion percentage of entity class A relative to entity class B can be defined as the percentage of times A confuses B for a given corpus and a given cross-fold validation. Such knowledge helps us in discovering, refining or validating relationship between entity classes and creating more meaningful ontologies. Information on which entity classes damage the results of the target entity class will be valuable in creating more efficient and powerful rules or features. For example: 34% of the mistakes in classifying "RNA domain or region" were caused because of "DNA domain or region"; 44% of the mistakes caused in classifying "Protein complex" were caused by "Protein molecule"; and 23% of the mistakes caused in classifying "Lipids" were caused by "Protein molecule". In a significant number of cases, most of the confusions were caused by the immediate siblings as would be expected, but there were many exceptions. For example: "RNA domain or region" with "DNA domain or region"; "Lipids" with "Protein molecule"; and "DNA domain or region" with "Protein family or group". This reflects both the ambiguity inherent in natural language and also the fact that while the GENIA ontology reflects a consideration of the major

properties of an entity, the local context of a mention may be more indicative of a single property that may be shared with entities which are otherwise significantly different.

Figure 5.2: Depiction of which entities cause confusion for each entity



Each dotted arrow shows which biologically-relevant leaf-level entity class (at head of the arrow) causes most confusion for each leaf-level entity class (at the tail of the arrow) with the corresponding confusion percentage below its name

SECTION C: DISCOVERING MORE PATTERNS FOR ASSOCIATION EXTRACTION

A computer should parse the complex syntax of text well enough to extract the associations between the concepts in a sentence. One can help it by paraphrasing, or simplifying the sentence.

In this section, a method to use link grammar parser's dependency output to simplify sentences for the purpose of improving extraction of protein-protein interaction from literature is first presented. In the second chapter of the section, an improvement in the method to use parse trees (or constituency trees), a more common way of expressing syntax for sentence simplification is described. In addition, a shotgun approach of using simplified sentences for relationship extraction is validated.

In the third chapter of the section, a machine learning based method to extract relations between medical problems, tests and treatments in clinical notes is presented. The impact of sentence simplification for extracting these associations is determined.

6. DEPENDENCY-BASED SENTENCE SIMPLIFICATION

Accurate systems for extracting Protein-Protein Interactions (PPIs) automatically from biomedical articles can help accelerate biomedical research. Biomedical Informatics re-searchers are collaborating to provide meta-services and advance the state-of-art in PPI extraction. One problem often neglected by current Natural Language Processing systems is the characteristic complexity of the sentences in biomedical literature. This chapter reports on the impact that automatic simplification of sentences has on the performance of a state of the art PPI extraction system, showing a substantial improvement in recall (8%) when the sentence simplification method is applied, without significant impact to precision.

6.1 Introduction

One of the challenges in parsing biomedical text is that it is significantly more complex than articles in typical English text. Different analysis show other problematic characteristics, including inconsistent use of nouns and partial words (Yuka Tateisi & Jun'ichi Tsujii, 2004), higher perplexity measures (Elhadad, 2006), greater lexical density, plus increased number of relative clauses and prepositional phrases (Gemoets, Rosemblat, Tse, & Logan, 2004), all of which correlate with diminished comprehension and higher text difficulty. These characteristics also lead to performance problems in terms of computation time and accuracy for parsers that are trained on common English text corpus.

Three categories of sentences were identified: 1) normal English sentences, like in Newswire text, 2) normal biomedical English sentences – those sentences which can be parsed without a problem by Link Grammar, and 3)

complex biomedical English sentences – those sentences which cannot be parsed by Link Grammar. Aside from the known characteristics mentioned before, sentences in the third group tended to be longer (18% of them had more than 50 words, while only 8% of those in group 2 and 2% of those in group 1 did). It has been observed that parsers perform well with sentences of reduced length (Chandrasekar & Srinivas, 1997; Siddharthan, 2006).

6.2 Corpora

There are currently many publicly available corpora of biomedical texts, the most popular among them being BioInfer, GENIA, AIMed, HPRD 50, IEPA, LLL and BioCreative I task corpus. Among these corpora, BioInfer includes the most comprehensive collection of sentences and careful annotation for links of natural parser, in both the Stanford and Link Grammar schemes. Therefore, BioInfer corpus, version 1.1.0 (Pyysalo, Salakoski, Aubin, & Nazarenko, 2007), containing 1100 sentences is chosen for evaluating the effectiveness of the simplification method on the performance of syntactic parsers. The AIMed corpus contains annotation for 197 abstracts which were identified by the Database of Interacting Proteins (DIP) to have PPIs and 29 more which do not have PPIs. In each sentence, all the proteins and all the pair-wise interactions among them are annotated. AIMed is publicly available at <ftp://ftp.cs.utexas.edu/pub/mooney/biodata/interactions.tar.gz>. AIMed corpus is used for evaluating the impact of sentence simplification for PPI extraction.

6.3 Methods

The method includes non-syntactic and syntactic transformations, detailed next.

6.3.1 Non-syntactic transformation

There are three steps in the approach: 1. preprocessing through removal of spurious phrases, 2. replacement of gene names, 3. replacement of noun phrases.

To improve the correctness of the parsing, each biomedical sentence is first preprocessed to remove phrases that are not essential to the sentence. This includes removal of *section indicators*, which are phrases that specify the name of the section at the beginning of the sentence, plus the removal of phrases in parentheses (such as citations and numbering in lists). Also, partially hyphenated words are transformed by combining with the nearest word that follows or precedes the partial hyphenated word to make a meaningful word. For instance, the phrase "alpha- and beta-catenin" is transformed into "alpha-catenin and beta-catenin".

Occurrences of multi-word technical terms and entity names involved in biomedical processes are common in biomedical text. Such terms are not likely to appear in the dictionary of a parser (perplexity is high), and will force it to use morpho-guessing and unknown word guessing. This is time consuming and prone to error. Thus, unlike typical text simplification that emphasizes syntactic transformation of sentences, the approach utilizes a named entity recognition engine, BANNER (Leaman & G. Gonzalez, 2008), to replace multi-word gene names with single-word placeholders. Replacement of gene names with single elements is not enough, however, and grammatical category (i.e. singular or plural) of the element has to be considered. Lingpipe shallow parser (Alias-i, 2006) for biomedical text identifies noun phrases and replaces them with single elements. A single element is considered singular when the following verb indicates a third-person singular verb or the

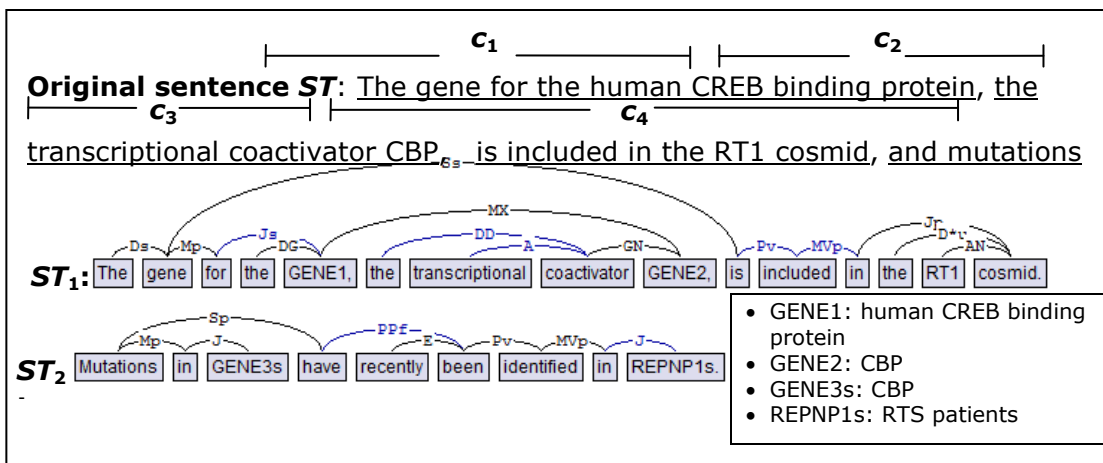
determiner preceded by the element is either "a" or "an". Otherwise it is considered as plural and an "s" is attached to the end of the element.

6.3.2 Syntactic transformation

The problem of simplifying long sentences in common English text has been studied before, notably by Chandrasekar & Srinivas (1997) and Siddharthan (2006). However, the techniques used in these studies might not totally solve the issue of parsing biomedical sentences. For example, using Siddharthan's approach, the biological finding "The Huntington's disease protein interacts with p53 and CREB-binding protein and represses transcription", and assuming multi-word nouns such as "CREB-binding protein" do not present a problem, would be simplified to:

"The Huntington's disease protein interacts with p53. The Huntington's disease protein interacts with CREB-binding protein. The Huntington's disease protein represses transcription."

Figure 6.1: Linkages after simplification of the original sentence



Our method transforms it to "*GENE1 interacts with GENE2 and GENE3 and represses transcription.*" Both decrease the average sentence length, but the earlier distorts the biological meaning (since the Huntington's disease protein might not repress transcription on its own), while the latter signifies it.

While replacement of gene names and noun phrases can reduce the sentence length, there are cases when the sentences are still too complex to be parsed efficiently. The example in Figure 6.1 illustrates a simple algorithm that utilizes *linkages* (specific grammatical relationships between pairs of words in a sentence) of the Link Grammar parser (Sleator, 1998) and punctuations for splitting sentences into clauses. Each linkage has a primary *link type* in CAPITAL followed by secondary link type in short. The intuition behind the algorithm is to try to identify independent clauses from complex sentences. The first step is to split the sentence ST into clauses c_1 , c_2 , c_3 and c_4 based on commas. c_1 is parsed using the Link Grammar parser, but c_1 cannot be a sentence as there is no "S" link in the linkage of c_1 . c_2 is then attached to c_1 and the linkage of " c_1, c_2 " does not contain a "S" link as well. " c_1, c_2, c_3 ." is recognized as a sentence, since the linkage contains an "S" link, indicating that it is a sentence, as well as the linkage of c_4 . So the algorithm returns two sentences ST_1 and ST_2 for ST .

This adhoc approach to syntactic simplification is a first step, but it did not have a method to determine whether a simplified sentence is grammatically correct or not. So, the subsequent sections first address the problem of accurately determining the grammatical correctness of a sentence and then describe a rule-based approach for splitting sentences.

6.3.3 Grammatical correctness of a sentence

For the sentence "He is an amazing.", the Stanford parser gives a complete constituent tree as if it were a grammatically correct sentence, while the Link Grammar parser (Sleator, 1998) correctly gives an incomplete linkage leaving behind "an", thus saying that "He is amazing." is the closest grammatically correct sentence. However, Link Grammar gives too many false negatives. In fact, Link Grammar does not find a complete linkage for more than 33% of the 1100 sentences in BioInfer corpus, all of which are grammatically correct (Jonnalagadda et al., 2009).

The statistical parsers generally give more false-positives than dependency parsers as they try to make a 'best guess' at a badly-constructed sentence, given that robustness in the face of poor grammar is more desirable. Most of the evaluation of the performance of parsers (for example, (Clegg & Shepherd, 2007)) is done using a corpus of grammatically correct English sentences. Those results are not relevant in the context where the primary task is to determine the correctness of a sentence. (Foster, 2007) uses an automatically generated Tree bank of grammatically incorrect sentences, but rightly notes that the correct combination of grammatically correct and incorrect training data for a statistical parser has yet to be created. The statistical parsers trained using only grammatically correct sentences will not be able to catch errors not found in the training data.

Siddharthan (Siddharthan, 2006) comments that it is difficult to automatically determine the grammatical correctness of a sentence; so he evaluates them manually. However, it will be very useful for the task of simplification to have an efficient method to automatically compare the syntactical soundness of the simplified parts to the original part before committing the simplification. This

section proposes using the combination of the number of null links and disjunct cost from the cost vector returned by the Link Grammar (version: 4.5.3) to get an estimate of the grammatical correctness.

Figure 6.2: Link Grammar's output

```
linkparser> This is an amazing.
No complete linkages found.
+++Time                                0.00 seconds (0.01 total)
Found 2 linkages (2 had no P.P. violations) at null count 1
  Linkage 1, cost vector = (UNUSED=1 DIS=0 AND=0 LEN=5)

  +-----Xp-----+
  +---Wd---+Ss*h+---Paf---+
  |         |         |         |
LEFT-WALL this.p is.v [an] amazing.a .
```

Figure 6.3: Link Grammar's another output

```
linkparser> This is an dangerously.
No complete linkages found.
+++Time                                0.00 seconds (0.01 total)
Found 1 linkage (1 had no P.P. violations) at null count 1
  Unique linkage, cost vector = (UNUSED=1 DIS=2 AND=0 LEN=5)

  +-----Xp-----+
  +---Wd---+Ss*h+---MVA---+
  |         |         |         |
LEFT-WALL this.p is.v [an] dangerously .
```

The vector "UNUSED" and "DIS" in the cost vector indicate the number of null links and disjunct cost, respectively. Figure 6.2 is the output for the sentence "This is an amazing." Here "an" from the original sentence is removed to suggest the closest grammatically correct sentence.

Figure 6.3 is the output for the sentence "This is an dangerously." The null count for both the sentences is 1, but the suggestion for the second sentence – "This is dangerously." is less grammatically correct than the suggestion for

the first sentence – “This is amazing.” This difference in grammatical correctness is reflected by the disjunct cost or DIS vector; the DIS vector for the second sentence is 2 while that for the first sentence is 0. Disjunct cost, also called Connector cost, represents the level of inappropriateness (caused by using less frequent rules) in the linkage.

Every sentence can be uniquely associated with the 2-tuple of null count and disjunct cost. It is reasonable to assume that a null count (which represents unwanted words) needs more attention than the disjunct cost (which represents less likely linkages). Since null counts and disjunct costs are typically less than 10 (i.e., one-digit numbers), for the purpose of easy comparison and for capturing the 2-tuples in one dimension, a new cost vector GRAM is defined to be equal to $10 \times \text{UNUSED} + \text{DIS}$. It is an easy proof that GRAM value is equivalent to the 2-tuple of null count and disjunct cost, under the assumption that the disjunct cost of the corresponding collection of sentences is not more than 10.

Any syntactic simplification will be approved only if the resulting sentences are collectively at least as grammatically correct as the original sentence alone, i.e., the sum of GRAM values of the parts should be less than or equal to the GRAM value of the original sentence. For example, the GRAM value of the sentence – “These effects were associated with significantly lower blood pressure, though within the normal range, in captopril-treated versus control animals.” is 20 (UNUSED = 2, DIS = 0; skipped “though” and “versus”). One suggestion from BioSimplify for splitting into the two sentences is: “These effects were associated with significantly lower blood pressure.” and “Within the normal range, in captopril-treated versus control animals.”, whose

respective GRAM values are 0 (UNUSED =0, DIS =0) and 22 (UNUSED =2, DIS =2). Since the sum of GRAM values of the parts (22) is more than the GRAM value of the original sentence (20), this suggestion is rejected. This approximates how a human would discern, because the second sentence was not grammatically correct. The second suggestion from BioSimplify was to split it into the two sentences – “These effects were associated with significantly lower blood pressure in captopril-treated versus control animals.” and “Significantly lower blood pressure is though within the normal range.”, whose respective GRAM values are 10 (UNUSED =1, DIS =0; skipped “versus”) and 10 (UNUSED =1, DIS =0; skipped “though”). Since the sum of GRAM values of the parts (20) is same as the GRAM value of the original sentence (20), this suggestion is accepted. In fact, if the second sentence was instead “This blood pressure is still within the normal range,” the GRAM value of the constituents would be lesser than that of the original sentence. This would also be possible in the future with more advanced implementations for resolving relative and appositive clauses in the syntactic simplification.

6.3.4 Overview of rules for simplification

This work implemented the rules for prefix subordination, infix subordination and if-then coordination (details in Siddharthan, 2003). These rules were also adapted recently by SimText (Ong, et al., 2008), a text simplification system for improving the readability of medical literature, but without a mechanism to judge the grammatical correctness. The notion of the GRAM value (as described in the prior section) as a way to automatically judge the syntactical soundness of the simplified parts as compared to the original sentence was also added. There are seven rules in total – three for conjunction and two each for relative clauses and apposition.

The referring expression (Siddharthan, 2003) for the relative clause is determined using the "MX" link from the link grammar output and that for appositive clauses is determined using the "R" link from the link grammar output. In addition, all the describing phrases that occasionally occur at the beginning of the sentence like the underlined phrases in: "These results suggest that affixin is involved in reorganization of subsarcolemmal cytoskeletal actin by activation of Rac1 through alpha and betaPIXs in skeletal muscle." and "As reported previously, alphaPIX was specifically co-immunoprecipitated by anti-affixin and anti-betaPIX antibodies." are removed.

6.4 Impact of simplification on the PPI extraction

The BioCreative Meta Server (BCMS) platform (Leitner et al., 2008) is a framework provide meta-services for information extraction in molecular biology. There are 12 research labs providing BCMS servers, but currently (as the outcome of BioCreative II) the publicly available servers only give information on whether the abstract with a given PubMed ID contains at least one PPI or not. This section is studying whether simplification of a sentence helps PPI extraction systems, so it was more appropriate to use a tool which operates on single sentences. For this reason, and the fact that it uses parse tree information, PIE (Kim et al., 2008) was selected as an ideal tool for evaluating BioSimplify. PIE is available as an online web service that can test any sentence(s) for presence of PPIs, not just a PubMed abstract. More information about the usage of PIE is available at <http://bi.snu.ac.kr/pie>. PIE returns positive or negative for each sentence depending on whether or not it detects a PPI in it.

Table 6.1 shows some examples where BioSimplify simplified a sentence and helped correct the output of PIE. The sentences in which PIE reported a PPI are marked in **bold**.

Table 6.1: Examples of impact of BioSimplify in improving association extraction

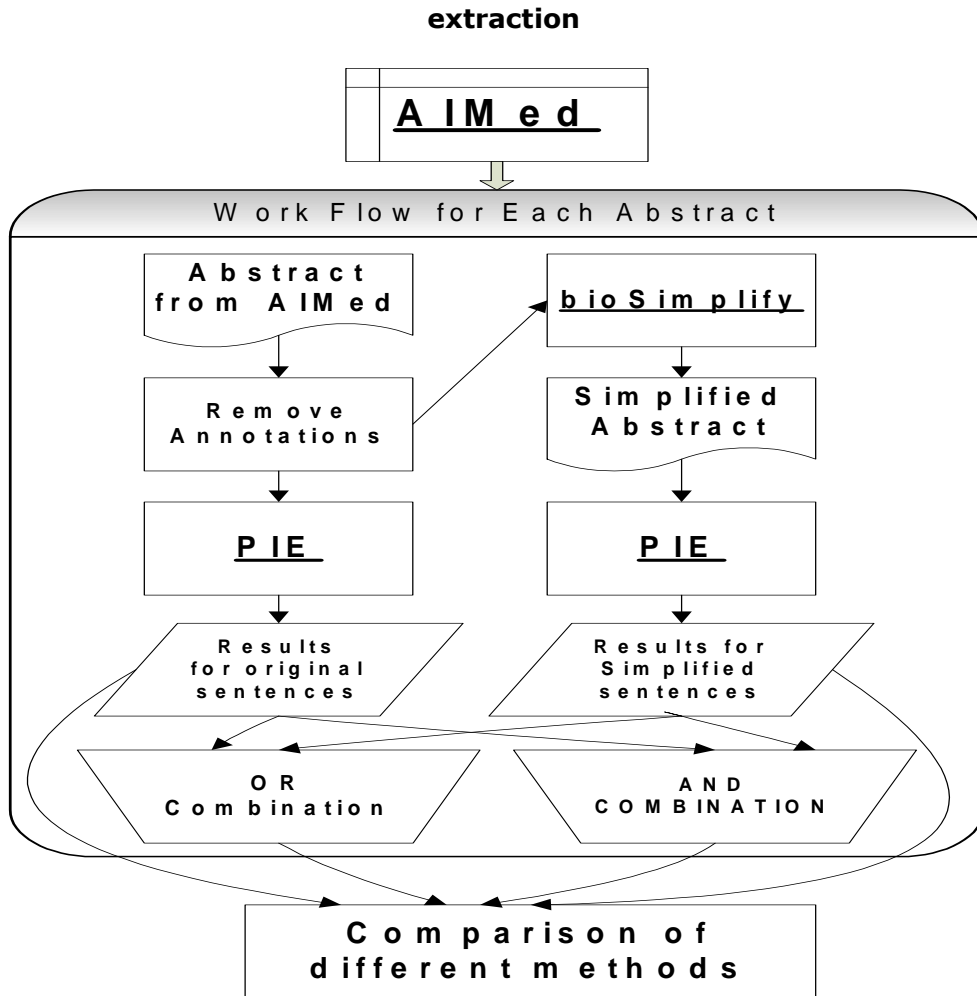
Original Sentence	Simplified Sentence	Comment
Unlike human IL - 6 (which uses many hydrophilic residues), the viral cytokine largely uses hydrophobic amino acids to contact gp130, which enhances the complementarity of the viral IL - 6 - gp130 binding interfaces.	Unlike human IL - 6, the viral cytokine largely uses acids to contact gp130. gp130 enhances the complementarity of the viral IL - 6 - gp130 binding interfaces.	False negative converted to True positive. BioSimplify has split the sentences into two parts (in addition to preprocessing and replacing noun phrases) and made it easy for PIE to identify the PPI.
LEC also induced calcium mobilization, but marginal chemotaxis via CCR5.	LEC also induced calcium mobilization, but chemotaxis via CCR5.	Replacing noun phrase with single word helps to concentrate on PPI indicating words and structure
The sequences that confer on FGF - 7 its specific binding to KGFR have not been identified.	The sequences confer on FGF - 7 its specific binding to KGFR . The sequences have not been identified.	Splitting sentences again uncovers the PPI indicated in part of the sentence.
It has been shown that LIGHT triggers apoptosis of various tumor cells including HT29 cells that express both lymphotoxin beta receptor (LTbetaR) and HVEM / TR2 receptors.	LIGHT triggers apoptosis of cells including cells. Cells express both lymphotoxin beta receptor and HVEM / TR2 receptors.	Noun phrase replacement, preprocessing and sentence splitting together separate the PPI containing part of the sentence for easy identification
Thus, Phe93 and Phe205 are important binding determinants for both EPO and EMP1, even though these ligands share no sequence or structural homology, suggesting that these residues may	Thus, Phe93 and Phe205 are determinants for both EPO and EMP1. These ligands share no sequence or structural homology, suggesting that these residues may represent	Splitting sentences again uncovers the PPI indicating part of the sentence.

represent a minimum epitope on the EPOR for productive ligand binding.	a minimum epitope on the EPOR for binding.	
We have crystallized a complex between human FGF1 and a two - domain extracellular fragment of human FGFR2.	We have crystallized a complex between human FGF1 and a two - domain fragment of human FGFR2.	Even a little amount of simplification sometimes highly influences the PPI extraction.
The structural arrangement in the active site is consistent with a mostly associative mechanism of phosphoryl transfer and provides an explanation for the activation of Ras by glycine - 12 and glutamine - 61 mutations.	The structural arrangement in the active site is consistent with a mostly associative mechanism of transfer. Arrangement provides an explanation for the activation of Ras by glycine - 12 and glutamine - mutations.	False positive converted to true negative. BioSimplify transformed the sentence and helped PIE in deciding that there is actually no PPI in the sentence, though there seems to be one because of the complexity of the sentence.

Figure 6.4 explains the evaluation steps. Considering a sentence as a positive if PIE detects a PPI in either the original sentence or in its simplified version should, in theory, lead to a system with higher recall without negatively affecting precision. For explanation, let us assume this set-up. BioSimplify transforms each sentence into at least one sentence that could be different from the original. Suppose it transforms a sentence A into A1,...,An for n>0. There are four possible outcomes: a) A is falsely assigned positive by PIE: Irrespective of whether A1,...,An are falsely assigned positive, there is no change in the precision and recall. b) A is correctly assigned negative: It is highly unlikely that the system which performed well with A, commits a mistake on the simpler parts – A1,...,An; so the precision is less likely to decrease and the recall would be the same irrespective of the performance of the system on A1,...,An. c) A is correctly assigned positive: Irrespective of the

performance of the system on A_1, \dots, A_n , there is no change in the precision and recall. d) A is falsely assigned negative: Though the system failed to identify PPI in A, the simplified segments A_1, \dots, A_n could allow a PPI extraction system to identify a PPI.

Figure 6.4: Evaluating impact of sentence simplification on PPI



Thus, addressing the false negatives (case d) without increasing false positives (case b) would increase both the recall and precision of the system. Cases a) and c) do not affect the system's performance. Case b) presents a slight likelihood for a decrease in precision (where the simplified sentence

triggers a false positive), but because this (in theory) happens sparsely - in 2% of the cases while processing the AIMed corpus - and there is a high likelihood of increase in recall and precision by resolving false negatives (case d), overall the system would have a higher recall with almost the same precision when processing simplified sentences. This set-up is referred to as **OR combination**. An alternate set-up where a sentence is considered positive only when PIE identifies a PPI in both the original sentence and the corresponding simplified sentence would be referred to as **AND combination**.

PIE was used to test for the presence of PPIs in 942 sentences before and after simplification. PIE reports whether a sentence with a potential PPI has a high probability of having a PPI or just a moderate probability. Any sentence in which PIE identifies a PPI with at least a moderate probability is considered as positive. These sentences are from 76 PubMed abstracts in AIMed with ids between 8816798 and 11470772 (this selection was based on DIP), and 14 PubMed abstracts chosen with PubMed ids between 11780382 and 11790884 for negative examples of interaction. Overall, out of the 942 sentences in these abstracts, 270 contain PPI(s). Each abstract was processed as illustrated in Figure 6.4. The aggregate results of PIE on all the sentences and their simplified counterparts are presented in Table 6.2.

Table 6.2: Results of PIE on selection from AIMed

Category	Recall	Precision	F-score
Before simplification	53%	49%	51%
After simplification	55%	52%	53%
AND combination	47%	53%	50%
OR combination	61%	48%	54%

6.5 Conclusion

The results of evaluation and error analysis allow us to conclude that BioSimplify, although still needing improvements, leads to improved PPI extraction results using PIE, which already uses syntactic information from parse trees. The results indicate that a system for sentence simplification used as a preprocessing step for natural language processing could improve the accuracy of PPI extraction process and other association extraction processes.

7. PARSE TREE-BASED SENTENCE SIMPLIFICATION

The last chapter presented a method of sentence simplification that is dependent on the output of the Link Grammar in a very specific dependency format. BioSimplify is optimized for accurate and flexible biomedical information extraction, and thus uses the output from parsers in the standard Penn Tree Bank (PTB) format. This is a reasonable choice considering that there has been significant increase in the accuracy of parsing biomedical text from 80% in 2005 (Lease & Charniak, 2005), to 84% in 2008 (McClosky & Charniak, 2008), and to 88% in 2009⁹ measured according to F-score and it facilitates inclusion of the simplification algorithm in different NLP pipelines. This choice also decouples the simplification process from the parsing step, allowing it to be done separately or to take advantage of the availability of collections of pre-processed sentences like the NLP web service provided by NCIBI (Ade, Wright, & Jagadish, 2009) or the PTDB (L. Tari et al., 2009) database of Arizona State University. This chapter uses Penn trees obtained from McClosky parser (McClosky & Charniak, 2008) which has an F-score of 88% for parsing biomedical text. BioSimplify can also be used with Penn trees produced from other parsers like Stanford parser and Link Grammar that can produce PTB-style (Marcus et al., 1993) output and also from Penn trees created apriori. Our goal is to produce all possible grammatically correct simplified sentences, assuming the available Penn tree is completely accurate.

7.1 Introduction

Explaining the limitation of the bag of words model, linguist Zellig Harris pointed out (Z. Harris, 1954): "language is not merely a bag of words but a tool with particular properties ... The linguist's work is precisely to discover

⁹ <http://www.cs.brown.edu/~dmcc/biomedical.html>

these properties.” When discourse analysts try to discover these properties, they usually break the sentence into simpler clauses. It should be noted however that a single independent clause or simple sentence as defined by (Quirk, Greenbaum, Leech, & Svartvik, 1985), may still be complex. Consider for example the following sentence, which even if one were to break it into simpler clauses, it would still be a complex: *We have identified a new TNF-related ligand, designated human GITR ligand (hGITRL), and its human receptor (hGITR), an ortholog of the recently discovered murine glucocorticoid-induced TNFR-related (mGITR) protein.* (Gee, 1999) advises that critically analyzing the discourse involves separating and unpacking clauses from sentences and phrases to understand all the perspectives. Automatically creating this set of simplified sentences for the purpose of information extraction on biomedical text is the subject of this paper. While existing NLP methods for information extraction already use grammatical information of text in the form of features like POS tags, parse trees and dependencies – informally known as “bag of NLP”, the usual focus is to choose one optimal such parsing for further processing. Our approach is rather a “shotgun” approach: use grammatical information in elemental chunks that can then be combined and recombined to generate many sentences from one (different perspectives) in order to maximize the likelihood that an automatic extraction engine can find in one (or several) of them the information contained in the original sentence. Thus, BioSimplify outputs the set of all sentences it can generate from the original sentence such that they are: 1) implied by the original sentence, 2) grammatically correct and 3) shorter than the original sentence.

7.2 Methods

While trying to understand how to break a sentence into simpler parts, this work focused on how sentences grow, and devised methods to undo the expansion. (Halliday, 1985) states that there are three ways for such expansion: 1) elaborating an existing basic structure, 2) extending it by addition or replacement and 3) enhancing its environment. These basic guidelines were used to design rules (listed in Table 7.2) for creating simpler sentences out of a complex one. The rest of this section presents the Noun Phrase Replacement module and the Syntactic Simplification algorithm along with the rules (Tables 7.1 and 7.2).

Table 7.1: Syntactic Simplification algorithm.

<p>synSimp(t), where t is the Penn tree of the given sentence:</p> <ol style="list-style-type: none">1) Initialize <u>simpTrees</u>, the ordered set containing the Penn trees of all simplified sentences, with the Penn tree of the original sentence 2) FOR EACH subtree of t traversed in the order of depth-first traversal<ul style="list-style-type: none">- perform necessary simplifications at that node which are the simplifications that need not be repeated for all the parents to this node 3) Add the present tree to <u>simpTrees</u> 4) FOR EACH unprocessed tree in <u>simpTrees</u><ul style="list-style-type: none">- FOR EACH subtree of t traversed in the order of depth-first traversal- perform the simplifications for this node- add new trees in <u>simpTrees</u> if applicable 5) return the sentences represented by the trees in <u>simpTrees</u>
--

7.2.1 Noun phrase replacement

A noun phrase in English consists of an optional determinative, an optional premodifier, a mandatory head noun, and an optional post-modifier (Quirk et al., 1985). Noun Phrase chunkers usually return the noun phrases of the smallest length, excluding postmodifiers. Hence, the last word of an identified noun phrase is the head noun. Previously, Hakenberg et al. (Hakenberg et al., 2010) introduced the preprocessing step of replacing noun phrases in the sentences with their head noun. However, removal of the optional determinative makes the sentence grammatically incorrect, while removal of the premodifier still gives a grammatically correct sentence. So, in this work, only the premodifiers are removed. For example, the noun phrase "*the recently discovered murine glucocorticoid*" is replaced with "*the glucocorticoid*". The part-of-speech (POS) tags in the sentence are identified using Lingpipe¹⁰, one of the most widely used POS taggers trained on GENIA biomedical corpus¹¹. The OpenNLP maximum entropy method¹² was then used to identify the noun phrases in the sentence. The other chunkers that could be considered are GATE chunker, GENIA Tagger, Lingpipe and Yamacha. For noun phrase chunking, GENIA Tagger and OpenNLP perform the best (F-score of 90% on GENIA Corpus), but OpenNLP is more usable in the system workflow as it is written in Java, while the GENIA Tagger is written in C++. To remove the premodifiers, all the tokens other than the head noun and the starting determinative or numeral (if they exist) are removed from the noun phrases.

¹⁰ <http://alias-i.com/lingpipe>

¹¹ <http://www-tsuji.is.s.utokyo.ac.jp/~genia/topics/Corpus/>

¹² <http://maxent.sourceforge.net>

7.2.3 Syntactic Simplification

While POS taggers only indicate the grammatical role of a particular word, a parse tree represents the syntactic structure of the whole sentence, giving complete details on how the words in it are related to each other. Sentence simplification systems (Chandrasekar & Srinivas, 1997; S. Jonnalagadda & G. Gonzalez, 2009; Vanderwende et al., 2007) usually have parsers as integral part of their algorithm, while there are few systems (Siddharthan, 2006) that use only POS. The latter aim for fast simplification at the point of application, while the former give a higher importance to accuracy of the output. BioSimplify is optimized for accurate and flexible biomedical information extraction, and thus uses the output from parsers.

Table 7.2: Syntactic simplification rules with explanation and examples

Rule	Explanation	Example	Result
S ~ {S}. S contains NP S contains VP	Adds all simple sentences (phrases present in a larger sentence, but in themselves are a grammatical sentence)	In differentiating C2C12 cells, E2F complexes switch and DNA synthesis in response to serum are prevented when MyoD DNA binding activity and the cdk5 inhibitor MyoD downstream effector p21 are induced.	MyoD DNA binding activity and the cdk5 inhibitor MyoD downstream effector p21 are induced.
NP[NP1 VP1*] ~ [NP1] {NP1 "can be" VP1} *VP1 starts with a gerund, present participle or past participle	NP Postmodification by verb phrase (separate the verb phrase that provides extra information about the noun phrase, but is not related to the whole sentence)	The cloning of members of these gene families and the identification of the protein-interaction motifs found within their gene products has initiated the molecular identity of factors (TRADD, FADD/MORT, RIP, FLICE/MACH, and TRAFs) associated with both of the p60 and p80 forms of the TNF receptor and with other members of the TNF receptor superfamily.	The cloning of members of these gene families and the identification of the protein-interaction motifs has initiated the molecular identity of factors... The protein-interaction motifs can be found within their gene products.
NP[NP1 ADJP1] ~ [NP1] {NP1 "can be" ADJP1 }	NP Postmodification by adjective phrase (similar to above)	Src homology domain-2 (SH2)/SH3 domain - can be containing adapters such as Grb2, Crk, and Crk-L, which interact with guanine nucleotide exchange factors specific for the Ras family.	... interact with guanine nucleotide exchange factors. Guanine nucleotide exchange factors can be specific for the Ras family.
NP[NP1 PP] ~ [NP1]	NP Postmodification by prepositional phrase is removed	To explore the role of the different domains of the betaL subunit in IFNalpha signaling, we coexpressed wild-type alpha subunit and truncated forms of	To explore the role in IFNalpha signaling, we coexpressed ...

		the betaL chain in L-929 cells.	
VP[MD VP1 , S*] ~ [MD VP1] *S contains VP and not NP	VP Postmodification by verb phrase is removed	T lymphocytes can be activated normally in response to either stimulus, demonstrating that the effects of the inactive CaMKIV on activation are reversible.	T lymphocytes can be activated normally in response to either stimulus.
VP[... , PP] ~ }, PP{* *Terminal prepositional phrase and preceding comma are removed from verb phrase	VP Postmodification by prepositional phrase is removed	Because cell lines can lose their differentiated phenotype in culture across passages, documentation of gene expression must be determined for passage populations, for us to have knowledge of cell behavior in vitro.	Because cell lines can lose their differentiated phenotype in culture across passages, documentation of gene expression must be determined for passage populations.
NP[NP1 PRN] ~ [NP1] [PRN - LRB - RRB]* *The left and right brackets are removed from the parenthetical	Handling abbreviations: Replace with two sentences- one with abbreviation removed, the other with NP replaced by abbreviation	Coexpression of the alpha and betaL subunits of the human interferon alpha (IFNalpha) receptor is required for the induction of an antiviral state by human IFNalpha.	Coexpression of the alpha and betaL subunits of the human interferon alpha receptor is ... Coexpression of the alpha and betaL subunits of the human IFNalpha is ...
NP[NP : S*] ~ [S*] *S contains VP or NP	Section indicator is removed	OBJECTIVE: To investigate the relationship between the expression of Th1/Th2 type cytokines and the effect of interferon-alpha therapy.	To investigate the relationship between the expression of Th1/...
S[S1* , NP VP] ~ [NP VP] *S1 does not contain both NP and VP	Content clause is removed	To characterize these pathways, we focused on changes in the cyclin-dependent kinase inhibitors and their binding partners that underlie the cell cycle arrest at senescence.	We focused on changes in the cyclin-dependent kinase inhibitors and their ...
NP[NP SBAR] ~ [NP], {SBAR -WHNP + NP}* *Wh-NP in the relative clause is replaced by NP from main clause	Relative Clause of this type is separated from the original sentence	To characterize these pathways, we focused on changes in the cyclin-dependent kinase inhibitors and their binding partners that underlie the cell cycle arrest at senescence.	...changes in the cyclin-dependent kinase inhibitors and their binding partners. The cyclin-dependent kinase inhibitors and their binding partners underlie...
VP[... , SBAR...] ~ }, SBAR{	Relative Clause of this type is removed	As [Ca2+]o increased, [Ca2+]i rapidly increased, as monitored by fluorometry.	As [Ca2+]o increased, [Ca2+]i rapidly increased.
VP , CC VP2] ~ [VP1] [VP2] PP [PP1, CC PP2] ~ [PP1] [PP2] ADJP[ADJP1 , CC ADJP2] ~ [ADJP1] [ADJP2]	Coordination of verb phrases, or prepositional phrase, or adjective phrase is removed by separating the coordinates into different sentences	These mechanisms must be understood in order to prevent, or combat, the emergence of a virulent, multidrug-resistant form of the bacillus that would be uncontrollable by means of today's treatment strategies.	These mechanisms must be understood in order to prevent, the emergence of a virulent, multidrug ... These mechanisms must be understood in order to combat , the ...

Column 1 is the mathematical notation for the rule and Column 2 is the verbal explanation of it. Readers interested in adapting the rules in their own environment might be more interested in Column1 than others.

The algorithm has a time complexity of $O(n^2 \cdot R)$, where n is the number of tokens in the sentence and R is the number of simplification rules. The average time complexity is, however, $O(n \log(n) \cdot R)$. One of the features of BioSimplify is avoidance of domain-specific rules. For example, the system does not replace entity names (like genes) with shorter alternatives as is done with the noun phrases in the present version and with the gene names in the earlier version (S. Jonnalagadda & G. Gonzalez, 2009). The system also avoided hard-coding the words in the rules created to split sentences with relative clauses. These measures enhance the domain adaptability of the system. Table 7.2 explains some of the simplification rules used. There are around 40 syntactic simplification rules available with the source code and documentation. The format for the rules is $A[A_1 A_2 \dots A_n] \sim [B_1 \dots B_p] \{C_1 \dots C_q\} \{D_1 \dots D_r\}$, where the variables are non-terminals in the parse tree. $[A_1 A_2 \dots A_n]$ are the children of A in the original tree. $[B_1 \dots B_p]$ are preserved from the original tree, $\{C_1 \dots C_q\}$ are removed from the original tree and added as a separate tree, and $\{D_1 \dots D_r\}$ are removed from the original tree. To make sure that the rules are optimized for biomedical sentences, the researchers manually examined each sentence in GENIA biomedical corpus and designed rules that would create the largest possible “bag of simplified sentences” based on Halliday’s formalism for sentence simplification. Some rules are classified as necessary (Table 7.2), and are executed only once since that transformations based on those rules do not decrease the semantic content of the original sentence. However, some transformations, like the postmodification of noun phrases by prepositional phrases, often destroy the semantic content. For such transformations, the algorithm ensures that both the simplified sentences and the sentence from which they are derived are

preserved. Our rule-set consists of all the rules present in the most comprehensive rule-based system currently known (Siddharthan, 2006) – albeit using a different notation based on phrase structures – and includes many additional rules. Siddharthan’s system (Siddharthan, 2006) handles coordination only at the clause level, while this system handles it also at the phrase level. We also remove section indicators, content clauses, and postmodifiers of phrases not handled before. Premodification of noun phrases is handled at NP replacement stage. We currently do not handle pronoun resolution.

7.3 PPI Extraction Evaluation

For the purpose of evaluating the impact of sentence simplification, AIMed corpus and PIE (S. Kim et al., 2008) (a machine learning approach available as a web service that uses the parse tree information from the Collins statistical parser as its key component) were used. PIE returns two kinds of results – one with a high precision, which is called *tight PIE*; and the other with low precision, which is called *light PIE*. We also compare the present version of BioSimplify with the older version (S. Jonnalagadda & G. Gonzalez, 2009) which is limited in its functionality because it only implements the rules described by (Siddharthan, 2006). The present version which has an average time complexity of $O(n \log(n) * R)$ is faster than the older version which has a time complexity of $O(n^3 * R)$, where n is the number of tokens in the sentence and R is the number of rules. The older version has domain specific optimizations (like replacing the gene names with single-word identifiers), which were not used in the newer version for portability.

7.4 Results and Conclusion

Table 7.3: Results for applying Protein-Interaction Extraction system PIE (high precision setting) to sentences from AIMed, as compared to the same to sentences processed with BioSimplify.

	Precision	Recall	F-score
Original sentences	53	47	50
Sentences simplified by older(S. Jonnalagadda & G. Gonzalez, 2009) BioSimplify	55	44	49
Sentences simplified by current BioSimplify	49	67	57

Table 7.4: Results for low precision setting PIE

	Precision	Recall	F-score
Original sentences	46	58	51
Sentences simplified by older(S. Jonnalagadda & G. Gonzalez, 2009) BioSimplify	51	64	57
Sentences simplified by current BioSimplify	46	82	60

We used PIE to test for the presence of PPIs before and after simplification in both the versions. The test set is from 18 PubMed abstracts in AIMed with ids between 9121766 and 9427624. Overall, out of the 189 sentences in these abstracts, 63 contain PPI(s). The aggregate results of PIE are presented in Tables 7.3 and 7.4. Precision, Recall, and F-score assume conventional meaning. Using BioSimplify improved the performance of light PIE by 9% in F-score and in recall by 24%. It also enabled an improvement in f- score by 7% and in recall by 20% on tight PIE. These improvements are statistically significant based on the two-tailed paired t-test on the outputs (PPI present/absent) of PIE before and after simplification ($p=1.3 \times 10^{-5}$ for light PIE and $p=7.2 \times 10^{-8}$). Overall, the present version of BioSimplify performs much better than the older version. This is because using the “shotgun” model for simplification (many simpler sentences) instead of the original sentences improves the chances that the PPI engine will detect a relationship

hidden in complex syntax. The precision of the present system is slightly lower than that of the older system because of the exhaustive simplification. This was not the case with the old system where the rules were based on Siddharthan's work and did not stress as much exhaustiveness of simplification. In situations where precision is much more important than recall, one could use only a subset of the rules that are empirically found to be more precise. The slight loss in precision could also be attributed to the removal of domain-specific features like replacing all gene names with unique identifiers.

Three judges evaluated the precision and recall of the BioSimplify system itself by reading each simplified sentence produced. The evaluation criteria (grammatical correctness) and this post-hoc evaluation using judges follow the same model and rationale described by (Siddharthan, 2006). We used 404 sentences from AIMed for this evaluation, for an estimated precision of 90%. Since the evaluation for grammatical correctness is done post-hoc by humans, the recall could be overestimated, as it is cognitively difficult to think of all possible grammatically correct sentences. Our judges found less than 1% new simpler sentences that were not produced by the system. The corpus of 4511 biomedical sentences (out of which 2017 contain PPIs) produced from the original 404 is available at <https://biosimplify.sourceforge.net>. These sentences were artificially created by BioSimplify from AIMed corpus and are manually annotated to indicate the presence or absence of PPI(s).

8. SENTENCE SIMPLIFICATION FOR CLINICAL RELATION EXTRACTION

In narrative clinical records, relations may exist between medical problems and treatments, medical problems and tests, and between medical problems and other medical problems. This chapter explains the approach to extracting these relations defined in the i2b2/VA 2010 shared task:

1. Treatment improves a medical condition.
2. Treatment worsens a medical condition.
3. Treatment causes a medical problem.
4. Treatment is administered for a medical problem.
5. Treatment is not administered for a medical problem.
6. Test reveals a medical problem.
7. Test is conducted to reveal a medical problem.
8. Medical problem indicates another medical problem.

The machine learning features of the system designed are based on those proposed by (Uzuner et al., 2010) for a similar relation extraction task in clinical domain. It uses MaxEnt classifier as implemented by MALLET to annotate the relations found between medical problems, tests and treatments in clinical notes. To test the impact of sentence simplification on this task, BioSimplify developed by (Siddhartha Jonnalagadda & Graciela Gonzalez, 2010) is used. The i2b2/VA shared task organizers (<https://www.i2b2.org/NLP/Relations/>) have annotated a training set and test

set for this specific relation extraction task to be used in the 2010 competition. This corpus, containing 349 training set documents and 477 testing set documents will be used to test the impact of sentence simplification. For evaluation purpose, it is assumed that the correct concepts are available for the system from the annotators and the task is to extract the relations between the concepts.

8.1 Machine learning features

The baseline system is a machine learning system with different natural language based features at the level of words, syntax, semantics or pragmatics. Most of the features are based on those proposed in (Uzuner et al., 2010), however the feature extraction system is implemented originally in Java and is available as an open source in the BioSimplify project page¹³ along with the code for sentence simplification. The table below lists all the features that are implemented.

Table 8.1: List of features used for extracting clinical relations

Feature name	Feature Type
Relative ordering of the candidate concepts	Lexical
Distance between the candidate concepts	Lexical
Presence of intervening concepts	Lexical
Tokens in concepts	Lexical
Concepts themselves	Lexical
Lexical unigrams, bigrams and trigrams around the concepts	Lexical
Tokens between concepts	Lexical
Verb phrases	Syntactic
Head words of the concepts	Syntactic
Tokens that are connected to concepts in dependency graph	Syntactic
Link path between the headwords of the concepts	Syntactic
Matches to manually created relationship patterns	Syntactic
Name of the section	Pragmatic

¹³ <http://biosimplify.sourceforge.net>

In order to extract the features in Table 8.1, the system tokenized each sentence, found the phrase chunks using the OpenNLP chunker¹⁴ and constructed a Penn Treebank format parse tree using the Stanford parser (Klein & C. D. Manning, 2003).

It was found that using the information about the shortest path between the candidate concepts in the dependency graph of a sentence in literature helps in improving over-all accuracy (Ding, Berleant, J. Xu, & Fulmer, 2003) of relationship extraction. Hence, Uzuner et al. (Uzuner et al., 2010) used link path features from Link Grammar for extracting relations in clinical domain. However, Link Grammar is a rule-based system built for grammatical English sentences. As pointed out previously (S. Jonnalagadda & G. Gonzalez, 2009), it would not produce reasonable output for sentences that are not grammatically correct. It is well established that most of the sentences in clinical notes are grammatically incorrect sentences or phrases. Thus, the Stanford parser which is a statistical model is used, since it outputs the most probable parse tree given the tokens in the sentence. Work by Chapman et al. is currently in progress to create parse trees for sentences in clinical notes as part of the AMIA NLP working group annotation grant and the Sharp grant. It is possible to retrain the statistical model with the annotations specific to clinical notes to improve the performance.

The parse tree which is in PTB format is converted to a dependency graph following Stanford format designed by Marneffe et al. (Marneffe et al., 2006). Dijkstra's algorithm (Dijkstra, 1959) to find the shortest path between nodes in a graph is used to find the link path. The time complexity for this step is

¹⁴ <http://opennlp.sourceforge.net>

$O(N^2)$ (which is less than $O(N^3)$, the time complexity for constructing the parse tree), where N is the number of tokens in the sentence.

In addition, Smith-Waterman local alignment algorithm (T. F. Smith & Waterman, 1981) is implemented to find if the candidate relation instance is similar in pattern to the patterns created manually from the true relations instances. To create the patterns, each sentence is reduced to the shortest, continuous snippet that contains both concepts from the ground truth as well as some keywords that indicate an interaction. To preserve only the keywords, the system removed all the stop words (Fox, 1989) in between the concepts. For example, the interaction where the treatment improved medical condition - *hypertension was controlled on hydrochlorothiazide* - is converted to the template "PROBLEM_ controlled TREATMENT_" so that the pattern matches other sentences or snippets such as, "*She has an elevated cholesterol controlled with Zocor.*" All the annotated concepts were replaced with the corresponding entity tags. For example, in the pattern "TEST_ revealed PROBLEM with PROBLEM and PROBLEM_", the candidate concepts are represented by the first and last words suffixed by an underscore. The concepts that were not part of the relation were not suffixed by an underscore. The snippets thus automatically created were manually examined. Where ever insufficient information was contained to infer the relation they were corresponding to, the snippets were expanded using the least number of words from the right or left to have sufficient information for inferring the relation. The result of the match of the testing instance to each of the pattern is a feature. The different patterns that were used is presented in

<https://sourceforge.net/projects/biosimplify/files/ClinicalRelationPatterns.txt>.

8.2 Choice of baseline for testing the impact of sentence simplification

Three different baselines were created. The best of the three that also compares to state of the art systems will be chosen to test the impact of sentence simplification.

The first baseline is a rule-based system that uses approximate matching based on local alignment with the clinical relation patterns. The threshold for the approximate match is fine tuned using "leave one out" validation. In the leave one out validation, each sample is separately used for testing while the rest of the samples are used for training. This is a special case K-fold cross-validation, where K is the number of samples. Multiple thresholds of approximate match (measured using Smith-Waterman distance) starting from 0.5 to 1.0 in increments of 0.05 are used for testing. It is found that a threshold of 0.85 yields the best F-score with leave one out validation on the training set.

The second baseline is the machine learning system which uses all of the features discussed in section 8.1 above. Different machine learning algorithms such as MaxEnt, AdaBoost, C45, Decision Tree, Naïve Bayes and Winnow as implemented in Mallet (AK McCallum, 2002) are tested using 10-fold cross validation on the training set. MaxEnt performed significantly better than the other algorithms and hence it was chosen to be used in the final system.

The third baseline is also a MaxEnt-based machine learning system that uses all of the features discussed in section 8.1 above except the features using approximate pattern match.

Table 8.2: Performance of the three baselines at i2b2/VA shared task

#	Baseline name	Recall	Precision	F-score
1	Pattern matching	0.3614	0.3449	0.3530
2	Machine learning with pattern feature	0.7496	0.6513	0.6970
3	Machine learning without pattern feature	0.7518	0.6488	0.6965

As shown in the Table, baseline 2, the MaxEnt-based machine learning system that uses all of the features discussed in section 8.2 has the best performance. Table 8.3 compares this system with the Top-10 systems submitted to the i2b2/VA relation extraction shared task.

Table 8.3: Comparison of different teams at i2b2/VA shared task

Rank	Team	Recall	Precision	F-score
1	UTDallas, Roberts et al.	0.7534	0.7204	0.7365
2	National Research Council Canada, deBruijn et al.	0.6933	0.7738	0.7313
3	LIMSI, Grouin et al.	0.7076	0.7110	0.7093
4	Univ. Sydney, Patrick et al.	0.6751	0.7307	0.7018
5	Arizona State Univ., Gonzalez et al.	0.7496	0.6513	0.6970
6	Salt Lake City VA, Divita et al.	0.6510	0.7459	0.6952
7	TU Budapest, Solt et al.	0.7067	0.6384	0.6708
8	NLM, Aronson et al.	0.6207	0.7183	0.6660
9	Brandeis Univ., Anick et al.	0.6606	0.6654	0.6630
10	OHSU VA	0.6401	0.6723	0.6558

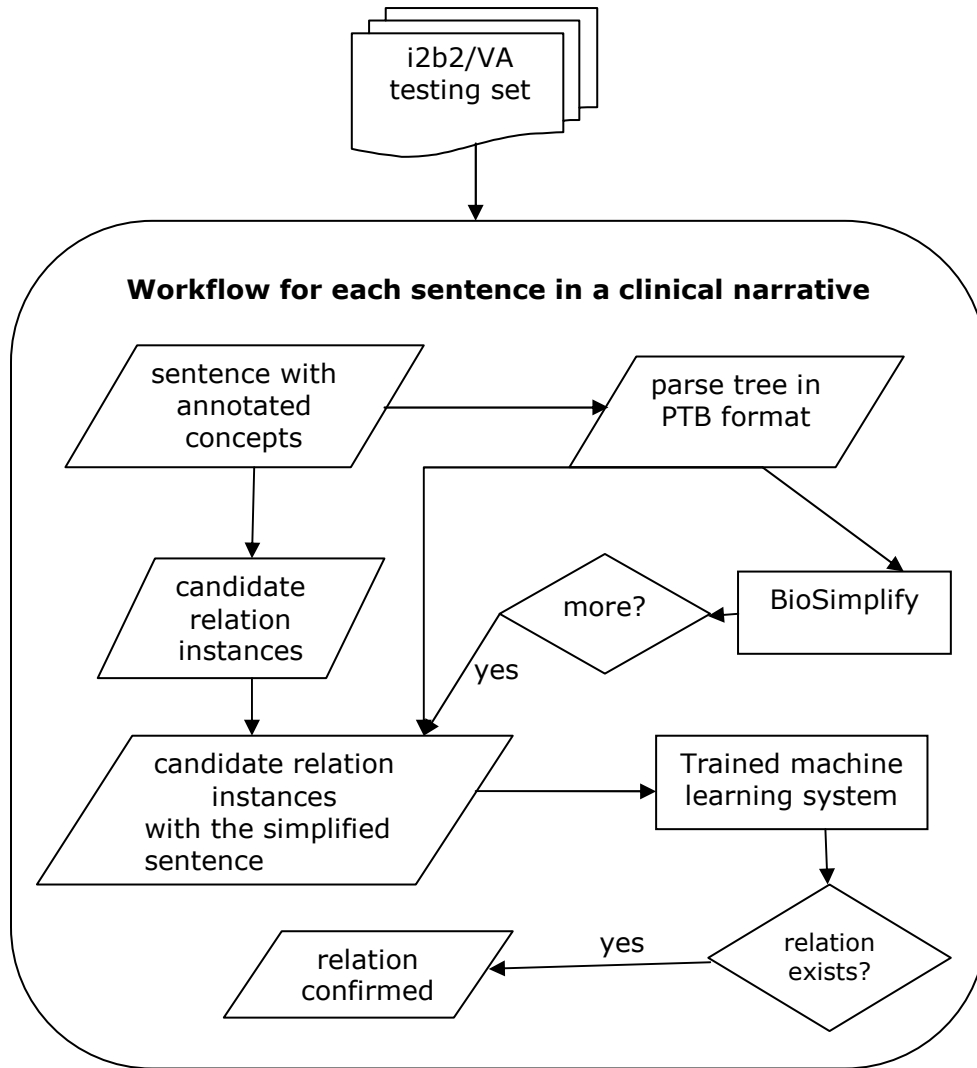
Our system placed 5th according to F-score among the 16 international teams that submitted annotations to the shared task. We rank second as per the levels of significance (shown in color).

8.4 Using sentence simplification as the post-processing step

The algorithm of BioSimplify is described in (Siddhartha Jonnalagadda & Graciela Gonzalez, 2010). Since the parse tree is available from the statistical parser used for creating link path related features, it is reused. Since creating

the parse tree has the highest time complexity (cubic in number of tokens in the sentence), using BioSimplify (linear in the number of clauses in the parse tree or approximately quadratic in the number of tokens in the sentence) does not significantly increase the total time taken.

Figure 8.1: Architecture for testing the impact of sentence simplification



BioSimplify produces a “bag of simplified sentences”, the collection of all possible simplified sentences from the original sentence including the original sentence and a “shotgun approach” can be used to search for relations in the

simplified sentence. Figure 8.1 describes the architecture for the analysis of the impact of sentence simplification. This architecture is similar to that is used in Chapter 7.

8.5 Impact of sentence simplification

Using the sentences resulted from BioSimplify with the trained machine learning system helped in converting 237 false negatives into true positives. However, 825 true negatives were also converted to false positives. The net effect leads to a slight decrease in F-score. The newly created false positives were examined manually. In almost all of the cases, the output of BioSimplify proved accurate. This means the errors were introduced because of errors in the trained machine learning system. A system with a high precision and low recall would have found the use of sentence simplification beneficial; however, the baseline system had high recall (75%) and low precision (65%) on the test set.

Noting F-score disregards the confidence of the system in annotating a relation, comparison of the systems is made using the area of curve, more specifically AUC of the interpolated precision–recall curve (AUC iP/R) that is proposed in the FEBS Letters/BioCreativeII.5 experiment (Leitner et al., 2010). In cases where there are more negative instances than positive instances, Leitner et al. argue that an ROC curve with recall (= sensitivity) on the x-axis and precision (= true positive rate) is more reasonable than using recall on x-axis and false positive rate (1 - specificity). The false positive rate ($FP/(FP+TN)$) is closer to zero for most systems given that the number of true negatives is usually very large. This is the also the case with i2b2/VA relation extraction task where the number of possible relations per sentence is in the order of square of the number of concepts per sentence, whereas the number

of actual relations annotated is usually much lesser than even the number of concepts. For example, in the training set, there were 9K relations and 45K concepts and in the testing set, there were 5K relations and 27K concepts.


Further, an interpolated precision-recall curve, estimates the highest possible precision at each achievable recall, using the formula:

$$A(f_{pr}) = \sum_{j=1}^n (p_{i_j} * (r_j - r_{j-1}))$$
$$p_i(r) = \max_{r' \geq r} p(r')$$

where n is the total number of correct hits and p_i is the highest interpolated precision for the correct hit j at r_j , the recall at that hit. Interpolated precision p_i is calculated for each recall r by taking the highest precision at r or any $r' > r$.

Figure 8.2 below compared the iP/R curves of the annotations with (blue) and without (red) simplification.

Figure 8.2: iP/R curves of annotations with and without use of sentence simplification




The AUC for the annotations which used simplification is 0.684 and that when simplification is not used is 0.673. There is a slight increment in the area under curve.

Further, Figure 8.3 below shows that sentence simplification helps in increasing F-score significantly when the system annotates only the relations it is confident of. This measure of confidence was available as part of Mallet's output along with the label assigned to the instance. Though there is a slight drop in F-score initially, for annotations the system is highly confident about, the F-score is significantly higher. In a hospital situation, it is essential to only consider system's output when there is a high probability that the output is correct, even at the cost of not being able to use the system's guess for some cases.

Wilcoxon signed-rank test (<http://faculty.vassar.edu/lowry/wilcoxon.html>) was used to measure the significance of increase of F-score over the entire


range of minimum confidence. Wilcoxon signed-rank test is used an alternative to the paired Student's t-test, when the samples may not be assumed to be normally distributed. The increase in F-scores over the entire range of confidence was found to be significant ($P < 0.0001$).

Figure 8.3: F-score vs. Confidence with and without use of sentence simplification



On X-axis, the confidence for annotating. On Y-axis, the F-score of respective system if all annotations with at least the corresponding confidence.

Figure 8.4: F-score vs. Confidence with and without use of sentence simplification (for confidence>0.4)



Additionally, when better or equal performance is desired, even at low confidence intervals, labels of candidate instances created from simplified sentences could be used only when the confidence is greater than a set threshold. For example, Figure 8.4 shows the F-score with respect to the confidence of a system that uses the label from simplified sentences only when the confidence is at least 0.4.

8.6 Error analysis

Although the sentence simplification was shown to have a positive impact for clinical relation extraction, it is worth noting that there are high numbers of false positives. In an ideal situation, where the system has 100% precision, there would have been no false positives even after using BioSimplify, but the recall will increase. Errors are introduced because the baseline system is not 100% precise and there is non-zero probability of falsely annotating a

candidate relation instance as positive. The objective of error analysis is to categorize the errors so that solutions to overcome or minimize might be identified. This would not only lead to better use of sentence simplification, but also improve the performance of the system overall.

65 candidate instances (out of 825) were identified among the false positives introduced after sentence simplification was used. These candidate relation instances are manually verified by a medical doctor (Sarada Panchanathan) and by a biomedical informatics doctoral candidate (Siddhartha Jonnalagadda). The categories of errors identified are:

1. Instances that the reviewers identified as true positives (20%):

This could also mean that some of them are "ambiguous", since the local annotations conflict the official annotations. However, the fact that the reviewers identified them as true positives only after thorough inspection might indicate that a lot of "domain" knowledge might be needed for the system (or a reviewer) to annotate them correctly. The medical doctor commented, "I think that obvious relations are detectable by all doctors, but it is the less obvious ones that depend on expertise. Also, someone else who is more of an expert may find a relation that I did not."

2. Instances that occur together frequently (siblings) (25%): This

occurs with respect to the relation type – "Problem Indicating Problem". Concepts that occur together **frequently** do not necessarily indicate each other. Example: In the sentence, "*This is an 80 year old female with a history of Stage I breast cancer , hypertension , hyperlipidemia , who presents with a leukocytosis/bandemia , hypotension in the setting of dehydration and influenza .*" leukocytosis/bandemia and hypotension together might indicate

a different problem such as sepsis; however, they do not indicate one another. This situation also occurs when considering alternate diagnosis.

3. Instances that are recognized, but are incorrectly classified (11%):

These situations occur in some cases because the relationship is expressed in future tense, for example - "will reveal", instead of past tense, for example - "revealed". Example: "*We will be obtaining a 24 - hour Holter monitor on discharge as well as setting Mr. **NAME[AAA] up for a prolonged ambulatory EEG in an attempt to capture these spells .*" The test "a 24 - hour holter monitor" **will be conducted** to diagnose "these spells". However, since the future tense is not recognized and the relation is mistaken as "Test revealed Problem."

In some other cases in this category, the errors can be attributed to the presence of negation (implicit or explicit) that reversed the type of relation. Example: "*Her delirium was felt possibly secondary to withdrawal versus toxic metabolic , but gradually improved slowly , and therefore patient did not receive treatment for barbiturate withdrawal .*" The concept "withdrawal" should have been associated to the concept "treatment" as "Treatment Not Administered for Problem". However, the candidate instance was incorrectly classified as "Treatment Administered for Problem".

4. Inability to parse clinical language, where the concepts are punctuated by clinical language to rule out a relation (8%):

Example: *Insulin 44 units NPH qam ; 8 units qpm ; Vasotec 2.5 qd ; Lasix 40 qd ; Mevacor 40 qd ; Cardizem CD 180 qd ; Ciprofloxacin for urinary tract infection ; aspirin .* The concept "cardizem cd" is not related to "urinary tract

infection". Other ways of separating is through use of medical abbreviations such as s/p, h/o and PMH.

5. Miscellaneous (36%): The reviewers identified these as cases where their occurrence together is coincidental and their classification as a relation is a random chance.

Solutions: In situations where the relation between concepts is ambiguous-based on the information in the sentence, domain knowledge through manually built dictionaries or automatically acquired semantics (T. Cohen & Widdows, 2009) might be useful. Most of the false positives are separated by punctuation marks such as ,./; and also punctuation markers in medical vocabulary such as s/p, h/o and PMH. Detecting negation and time (tense, or temporal order) will also help in correctly classifying instances. All these solutions will not only improve the impact of sentence simplification, but also the overall performance of the system.

In addition, the lexical characteristics of those instances that are annotated as True Positive (error type 1) by two reviewers are compared with the rest. The two characteristics considered are: a) number of words in the sentence, and b) number of clauses in the parse tree of the sentence. The samples of observations are compared against the two sets of instances independently using Mann-Whitney U test (<http://elegans.swmed.edu/~leon/stats/utest.cgi>) with the below results in Tables 8.4 and 8.5. Mann-Whitney U test is used as an alternative to the unpaired Student's t-test when the samples may not be assumed to be normally distributed.

Table 8.4: U Test Results for number of clauses

n₁	n₂	U	P (two-tailed)	P (one-tailed)
51	14	461.5	0.095416*	0.047708*
normal approx z = 1.66758			0.0953998*	0.0476999*

*These values are approximate.

Table 8.5: U Test Results for word length

n₁	n₂	U	P (two-tailed)	P (one-tailed)
51	14	451.5	0.133812*	0.066906*
normal approx z = 1.508			0.1315546*	0.0657773*

*These values are approximate.

n₁ and **n₂** respectively are the number of introduced false positive samples that are annotated as false positives and those that are annotated as true positives. The p-value is less than 0.05 for the one-tailed test using the number of clauses even with such a small sample. **These results suggest that relations may be missed more often by annotators when the sentences are complex.**

A comparison is also made between the lexical characteristics of the true positives introduced with sentence simplification and those of the false positives introduced with sentence simplification. Tables 8.6 and 8.7 fail to indicate a correlation based on word lengths and number of clauses. Table 8.8 however indicates that the sentences in true positives are significantly more complex than the sentences in false positives (based on the average number of clauses per word). **This suggests that sentence simplification might be more useful when relatively complex sentences are used.**

Further, Table 8.9 shows that **the number of concepts in false positive sentences is significantly more than those in true positive sentences.**

This appears to be due to error types 2 and 4 which also correlate with the number of concepts in the sentence.

Table 8.6: U Test Results for word length II

n₁	n₂	U	P (two-tailed)	P (one-tailed)
825	237	102506.5	0.25359*	0.126795*
normal approx z = 1.1399			0.254328*	0.127164*

*These values are approximate.

Table 8.7: U Test Results for number of clauses II

n₁	n₂	U	P (two-tailed)	P (one-tailed)
825	237	99789.5	0.625432*	0.312716*
normal approx z = 0.487053			0.62622*	0.31311*

*These values are approximate.

Table 8.8: U Test Results for average number of clauses per word in a sentence

n₁	n₂	U	P (two-tailed)	P (one-tailed)
825	237	108586.5	0.009056*	0.004528*
normal approx z = 2.60082			0.00930006*	0.00465003*

*These values are approximate.

Table 8.9: U Test Results for number of concepts in a sentence

n₁	n₂	U	P (two-tailed)	P (one-tailed)
818	237	116163.5	8e-06*	4e-06*
normal approx z = 4.65584			3.22656e-06*	1.61328e-06*

*These values are approximate

8.7 Conclusion

The impact of sentence simplification in clinical relationship extraction is studied using state of the art baseline – a MaxEnt-based system using several NLP features. In sum, sentence simplification significantly aids clinical relationship extraction, especially in situations where relatively high confidence of annotation is expected from the system.

In error analysis, the errors were categorized to help improve the overall system as well as the impact of sentence simplification. Further, the statistical tests suggest that human annotators face more difficulties with finding relations in complex sentences. The true positives found with sentence simplification are significantly more complex than the corresponding false positives. The number of concepts in these true positives is also significantly higher than those in the corresponding false positives.

SECTION D: OVERALL DISCUSSION AND CONCLUSIONS

This dissertation examined the possibility of reducing the limitations when automatically extracting information from biomedical texts caused by the scarcity of sufficiently large annotated corpora. The chapters in Section B presented an unsupervised approach that addressed the challenge of limited vocabulary in the training data. This approach included the use of distributional semantics to automatically generate lexical resources. These resources were shown to be more useful for concept extraction than the traditionally used manually curated lexicons. An overall improvement was also observed in the accuracy of extracting named entities from biomedical literature and clinical notes when using this approach. Section C addressed the challenge of small corpora not having enough syntactic patterns for accurate recognition of relationships between concepts. Experiments in that section showed that sentence simplification can help association extraction.

This section discusses future directions in concept and relationship extraction research suggested by the observed results and presents the overall conclusions one can draw from them.

9. DISCUSSION

A summary of the overall accomplishments with references to the corresponding chapters is first presented. Implications of these accomplishments on the field of informatics are then discussed. Limitations of the work are also described. Finally, potential applications and extensions of the research are elaborated upon.

9.1 Summary

In Chapter 3, the method to adapt distributional semantics for use in concept extraction or named entity recognition was discussed. An evaluation of the method was performed using the i2b2/VA concept extraction task corpus thus demonstrating that the distributional semantic features significantly improve the performance. In Chapter 4, the domain adaptability of the above method was demonstrated using the BioCreative II Gene Mention extraction task corpus. Further, a comparison of different types of distributional semantic measures was made. In Chapter 5, the applicability of distributional semantic measures for the more complex multi-label classification of named entities was demonstrated.

In Chapter 6, a sentence simplification method using Link Grammar parser was presented. In Chapter 7, the sentence simplification method was improved to be able to use the output of any statistical parser as input and the utility of simplification was demonstrated for protein-protein interaction extraction. In Chapter 8, the same system was used to simplify sentences in clinical notes. Sentence simplification was shown to improve clinical relationship extraction, especially in situations where high confidence of annotation is expected from the system.

9.2 Implications for the field

Through the work on concept extraction, two diverse research areas - distributional semantics and information extraction are amalgamated. A novel finding of this work is that semantic vector models could be used in extracting named entities. Previously, Gibb's sampling is used to incorporate features from words that are not in immediate neighborhood (Jenny Rose Finkel, Grenager, & C. Manning, 2005) to address the label inconsistency problem. For example, from the sentence "... the news agency *Tanjung* reported ...", they were able to correctly classify the named entity *Tanjung* used at the end of another sentence ("... airport, *Tanjung* said.") as an organization rather than as a person. Finkel et al. evaluated the impact of the sampling features by measuring the increment in performance after adding them to a state of the art baseline system. The increment in performance was 1.3% in the task of extracting concepts from general English. Although 1.3% increment might seem slight, Finkel et al.'s work has been cited more than 250 times, and is generally considered significant. In comparison to Finkel et al.'s work, the features proposed in this work not only use information from sentences in the same paragraph or a document, but also from those sentences present elsewhere in the unlabeled documents. The method proposed here using distributional semantics achieved a performance improvement of 2.0% in the task of extracting medical problems, tests, and treatments from clinical notes and 1.9% in the task of extracting gene mentions from biomedical literature. These improvements were measured against state of the art systems. It is hoped that this research, when published, compels similar interest in the NLP scientific community.

It is shown previously (J. R. Finkel & C. D. Manning, 2009a; Turian et al., 2010) that statistically created word clusters (Brown et al., 1992; Clark, 2000) could be used to improve named entity recognition. However, only a single feature can be derived from the clusters. Using distributional representations for concept extraction or similar NLP tasks has been considered an unsolved problem (Turian et al., 2010). This work not only demonstrated how distributional semantics could be used for concept extraction, but also that they are better than the feature based on clustering.

Although there are other semi-supervised machine learning approaches such as ASO (Ando, 2007), most of these systems restrict themselves to a specific framework and algorithm. This is disadvantageous considering most state of the art concept extraction systems use supervised machine learning algorithms, such as CRF, and it would not interest researchers to replace these systems with a completely different framework and algorithm. Our approach was to use NLP features generated through unsupervised means within a supervised machine learning system. While this approach renders all the advantages offered in other semi-supervised machine learning systems, it also permits maximum flexibility in the choice of the basic framework of the set-up. In fact, a rule-based system that uses these unsupervised features or measures was also demonstrated in Chapter 5.

The features created using the semantic vectors proved to be much more valuable than manually created lexical resources in the biomedical and clinical concept extraction tasks. This shows that these features help in domain adaptation for concept extraction which has immediate practical applications. In the medical domain, concept extraction is human intensive since different

systems focus on extremely specific problems such as extracting concepts related to congestive heart failure from chest radiology reports (Jeff Friedlin & McDonald, 2006) and identifying pancreatic cancer patients (J. Friedlin et al., 2010). Systems that could be easily adapted to different domains help in decreasing human effort by minimizing the need to customize systems manually to hundreds of such tasks.

Sentence simplification is shown to aid association extraction from biomedical literature and clinical notes. Sentences in biomedical literature and clinical notes are significantly more complex than articles in typical English text. After simplifying sentences using link grammar dependencies, the accuracy of Link Grammar parser and Charniak-McClosky parser in parsing biomedical literature increased. In addition, protein-protein interactions were extracted more accurately after simplifying sentences. A shotgun approach to using simplified sentences for association extraction is proposed in chapter 7. The same approach was applied to extract relations from clinical notes also with positive results. The methods proposed for sentence simplification and for use of simplified sentences in association extraction are unique and novel.

The hypothesis that sentence simplification could be used to improve association extraction was subsequently verified by other researchers (Miwa, Sætre, Miyao, & Jun'ichi Tsujii, 2010). They also simplified sentences using deep parser and tested the impact using several protein-protein interaction corpora. Sentence simplification improved the accuracy of extracting protein-protein interactions in the case of all the corpora.

9.3 Limitations

The unlabeled data currently used are clinical trials from Medline. The improvement to concept extraction is expected to be higher if one uses clinical notes from hospital records since that would also capture semantics of non standard words used in the clinical notes. The semantic vectors of the terms could be extracted from the clinical notes in clinical data warehouses. However, it is not known whether the vectors themselves could reveal protected health information of patients. Research needs to be done to determine such a possibility and if necessary redact information that compromise protected health information.

Future work should include comparison of the features with other models, such as Collobert and Weston's word embeddings (Collobert & Weston, 2008), which utilize neural network models to represent the syntactic and semantic information about the words in a dense, low dimensional feature space.

A limitation in using distributional semantics is that the meanings of words often differ with the passage of time and for specialized disciplines. This phenomenon is known as semantic change or semantic shift. For example, the word "guy" which is originally used to mean "grotesque person" (as named after Guy Fawkes hanged an assassination attempt in 1605), now means "fellow". The meaning of the word "case" changes with specialization (Traugott, 2002). In general domain, it means "circumstances in which one is". In medical domain, it means "a patient". In legal domain, it means "a law suit". The extent of reliability of the semantic vectors created automatically, given the semantic shift possibility, need to be analyzed.

It is also curious that while pattern-based systems were efficient for extracting protein-protein interactions (Leitner et al., 2010), they were outperformed by machine learning systems while extracting relationships in clinical notes (Chapter 8). Would prior knowledge about the concepts help in some cases to determine whether there is a relationship between them? It would be interesting to analyze if distributional semantics is useful to automatically gain this automatic knowledge.

The clinical relation extraction system could potentially perform more accurately if machine learning features to detect negation and temporality are employed. Given the flexibility of the clinical relationship extraction system's architecture, existing systems for detecting negations and temporality (W. Chapman, 2001; L. Zhou, Parsons, & Hripcsak, 2008) could be integrated in. Post-processing using rules based on abbreviations such as s/p, h/o and PMH that impact the relationship between concepts would also improve performance.

9.4 Potential applications and Future work

Automatic coding of concepts and relations in clinical notes: Since systems for extracting concepts and relationships are still being explored, automatic coding is an emerging technology. In text mining pipeline, there is a further step after extracting concepts known as normalization and disambiguation. This step involves disambiguating polysemy and synonymy. Polysemy is the event of the same entity having different semantic meanings in two different contexts. For example, PPI could mean Inorganic Pyro Phosphate when the subject of discussion is about enzymes; it could mean Peptidyl prolyl cis-trans-isomerase when the discussion is about proteins; it

could also mean protein-protein interaction in other situations. The first step in normalization will be to assign different identifiers to polysemous entities. Synonymy is the event of two different entities having the same semantic meanings in all contexts. For example, AD and SDAT; poliomyelitis and polio are synonyms. The task of normalization for entities like Proteins involves using a dictionary, such as UniProt, where each entity is mapped to a unique id. Automatic coding systems not only need concept extraction, but also require normalization using different schemas, such as the International Classification of Diseases (ICD) or Current Procedural Terminology (CPT), and use-case-specific, non-standardized schemas, such as the presence or absence of a given condition. The entities in the schemas could be pre-coordinated (coding as a single field) or post-coordinated (reassembled after storing individual components separately). Two problems associated with pre-coordination (Lussier & Bodenreider, 2007) are: a) updates to the cross index could be slow because of the knowledge resources needed and b) computation ambiguity of the reuse of the concept. Given the ability of distributional semantics to detect words appearing in similar context and to predict the appropriateness of a concept to a context, both these problems could be addressed automatically. Distributional semantics also helps in automatically mapping "pre-coordinated" terms with "post-coordinated" expressions and vice-versa.

Secondary use of clinical notes: The clinical notes in the hospital systems could be used to automatically create a conceptual thesaurus for the words in the notes. This thesaurus would be useful for annotating concepts in a machine learning system or a rule-based system. This thesaurus would be more useful for annotating the clinical notes within the hospital system, since

it is customized for that genre of text. However, chapter 3 used the clinical notes from the University of Texas, Houston to create a thesaurus and extracted relations in the i2b2/VA task (coming from a different hospital system). The F-score that is obtained is significantly better than the best combination thus far. In terms of comparison with other systems, the F-score of 91.3% with inexact match and 82.3% with exact match, ranks 2nd in terms of inexact match and 3rd in terms of exact match out of 22 international teams.

Domain independent concept extraction: To adapt a machine learning system of a particular domain to another domain, domain specific information is needed. Firstly, domain specific resources, such as thesauri and lexicons, need to be created. In chapter 4, it is demonstrated that such resources can be automatically created using distributional information in text. In fact, the automatically created resources were more useful than the manually created, curated and compiled dictionaries. Secondly, and more importantly, one needs corpora annotated with labels indicating concepts. Is it possible to construct labeled corpora automatically? Can one use semantic information retrieval to find suitable text belonging to the domain? Can publicly available ontologies be employed to simultaneously annotate concepts in the retrieved text?

Summarization and question answering: Summarization of documents is useful for clinicians and policymakers to obtain a summary of guidelines or related literature. It is also useful to aphasia patients who would like to read summaries made of shorter sentences. Our work with sentence simplification

can be extended in the future for the task of summarization (as demonstrated through SumBasic (Vanderwende et al., 2007)).

In addition, the simplified sentences could be translated to logical facts. Compilations of logical facts over large amount of text will make it possible to construct reasoning systems and potentially be useful for creating question answering systems such as Watson.

10. CONCLUSION

How does one overcome the inherent ambiguity present in human written text – be it in literature or in narratives, be it in semantics or syntax, be it in identifying concepts or expressing relations between two concepts in a sentence? This is the challenge that inspired this research. Ideally, all the ambiguity should be resolved by annotating all possible examples necessary to identify concepts and relationships. However, this is not possible due to the amount of domain expertise and effort needed. Section B proposed an effective way to minimize this limitation using distributional semantics and sentence simplification.

Humans learn semantics through experience. Attempts to impart semantic information to concept extraction systems is validated by the success of distributional semantic features in outperforming manual dictionary features. It is demonstrated that a computer can learn semantics through distributional statistics or geometry. Using concept extraction from the clinical notes as the first example, it is shown that distributional semantics is more useful than traditionally used dictionaries. This was measured by the increase in F-score accuracy on a baseline system that extracts medical problems, tests and treatments from clinical notes. Using distributional semantic measures calculated automatically saves the time spent in building and compiling dictionaries automatically. When the same methods are used for extracting protein names from biomedical literature, significant increase in F-measure is again achieved. This demonstrates that the distributional semantic features are significantly useful in automatically acquiring knowledge of any domain.

Humans deal with complex syntax by first breaking the sentence into simple, yet meaningful parts. Hence, an attempt is made to help a relation extracting system by paraphrasing or simplifying the sentence in different ways so that a relation might be identified in one of them. This experiment also yielded positive results in both biomedical literature and clinical notes. On the task of PPI extraction, use of sentence simplification improved the F-score by around 7%, with an improvement in recall of around 20%. Even for clinical notes, the F-score of the clinical relation extraction system increased significantly in situations where high confidence on the annotations were expected from the system.

REFERENCES

- Ade, A., Wright, Z., & Jagadish, H. (2009). *The NLP web service [Internet]*. Ann Arbor (MI): National Center for Integrative Biomedical Informatics. Retrieved from <http://nlp.ncibi.org>
- Ahmed, S. T., Chidambaram, D., Davulcu, H., & Baral, Chitta. (2005). IntEx: A Syntactic Role Driven Protein-Protein Interaction Extractor for Bio-Medical Text. *BioLink 2005*.
- Alias-i. (2006). *LingPipe 2.3.0*.
- Ando, R. K. (2007). BioCreative II Gene Mention Tagging System at IBM Watson. *Proceedings of the Second BioCreative Challenge*.
- Aronson, A. R. (2001). Effective mapping of biomedical text to the UMLS Metathesaurus: the MetaMap program. *AMIA*.
- Bach, N., & Badaskar, S. (2007). *A survey on relation extraction*. Language Technologies Institute, Carnegie Mellon University. Retrieved from <http://www.ark.cs.cmu.edu/LS2/images/9/97/BachBadaskar.2007.pdf>
- Baral, Chitta, Davulcu, H., Gonzalez, Graciela, Joshi-Topee, G., Nakamura, M., Singh, P., Tari, Luis, et al. (2005). CBioC: Web-based Collaborative Curation of Molecular Interaction Data from Biomedical Literature. *The Genetics Society of America 1st International Biocurator Meeting*.
- Baumgartner, W., Lu, Z., Johnson, H., Caporaso, J. G., Paquette, J., Lindemann, A., White, E., et al. (2008). Concept recognition for extracting protein interaction relations from biomedical text. *Genome Biology*, 9(Suppl 2), S9.
- Blaschke, C. (1999). Automatic extraction of biological information from scientific text: protein-protein interactions. *International Conference on Intelligent System Molecular Biology*, 60-67.
- Brown, P. F., Desouza, P. V., Mercer, R. L., Pietra, V. J. D., & Lai, J. C. (1992). Class-based n-gram models of natural language. *Computational linguistics*, 18(4), 467-479.
- Bunescu, R., Ge, R., Kate, R. J., Marcotte, E. M., Mooney, R. J., Ramani, A. K., & Wong, Y. W. (2005). Comparative experiments on learning information extractors for proteins and their interactions. *Artificial Intelligence in Medicine*, 33(2), 139-155.
- Byrne, K. (2007). Nested named entity recognition in historical archive text. *Proceedings of International Conference on Semantic Computing*.
- Carroll, J., Minnen, G., Canning, Y., Devlin, S., & Tait, J. (1998). Practical simplification of English newspaper text to assist aphasic readers.

Workshop on Integrating Artificial Intelligence and Assistive Technology.

- Chandrasekar, R., & Srinivas, B. (1997). Automatic induction of rules for text simplification. *Knowledge-Based Systems, 10*, 183-190.
- Chapman, W. (2001). A Simple Algorithm for Identifying Negated Findings and Diseases in Discharge Summaries. *Journal of Biomedical Informatics, 34*(5), 301-310. doi:10.1006/jbin.2001.1029
- Chen, E. S., Hripcsak, G., Xu, H., Markatou, M., & Friedman, C. (2008). Automated acquisition of disease–drug knowledge from biomedical and clinical documents: an initial study. *Journal of the American Medical Informatics Association, 15*(1), 87–98.
- Christensen, L. M., Harkema, H., Haug, P. J, Irwin, J. Y., & Chapman, W. W. (2009). ONYX: A system for the semantic analysis of clinical text. *Proceedings of the Workshop on BioNLP* (p. 19–27).
- Chuang, J.-H., Friedman, Carol, & Hripcsak, George. (2002). A comparison of the Charlson comorbidities derived from medical language processing and administrative data. *AMIA Symposium* (pp. 160-164). Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/12463807>
- Clark, A. (2000). Inducing syntactic categories by context distribution clustering. *Proceedings of the 2nd workshop on Learning language in logic and the 4th conference on Computational natural language learning-Volume 7* (pp. 91-94). Association for Computational Linguistics Morristown, NJ, USA.
- Cohen, T., & Widdows, D. (2009). Empirical distributional semantics: Methods and biomedical applications. *Journal of Biomedical Informatics, 42*(2), 390–405.
- Collobert, R., & Weston, J. (2008). A unified architecture for natural language processing: Deep neural networks with multitask learning. *Proceedings of the 25th international conference on Machine learning* (p. 160–167).
- Cortes, C., & Vapnik, V. (1995). Support-Vector Networks. *Machine Learning, 20*, 273-297.
- De Vine, L., & Bruza, P. (2010). Semantic Oscillations: Encoding Context and Structure in Complex Valued Holographic Vectors. *2010 AAAI Fall Symposium Series*.
- Dietterich, T. G. (1998). Approximate statistical tests for comparing supervised classification learning algorithms. *Neural computation, 10*(7), 1895-1923.
- Dijkstra, E. W. (1959). A note on two problems in connexion with graphs. *Numerische Mathematik, 1*(1), 269-271. doi:10.1007/BF01386390

- Ding, J., Berleant, D., Xu, J., & Fulmer, A. W. (2003). Extracting Biochemical Interactions from MEDLINE Using a Link Grammar Parser. *ICTAI '03: Proceedings of the 15th IEEE International Conference on Tools with Artificial Intelligence*. IEEE Computer Society.
- Elhadad, N. (2006). *User-sensitive text summarization: application to the medical domain* (PhD). Columbia University.
- Fan, R. E., Chen, P. H., & Lin, C. J. (2005). Working set selection using second order information for training support vector machines. *The Journal of Machine Learning Research*, 6, 1889–1918.
- Finkel, J. R., & Manning, C. D. (2009a). Joint parsing and named entity recognition. *Proceedings of Human Language Technologies: The 2009 Annual Conference of the North American Chapter of the Association for Computational Linguistics on ZZZ* (pp. 326-334). Association for Computational Linguistics.
- Finkel, J. R., & Manning, C. D. (2009b). Nested Named Entity Recognition. *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing*.
- Finkel, Jenny Rose, Grenager, T., & Manning, C. (2005). Incorporating non-local information into information extraction systems by Gibbs sampling. *Proceedings of the 43rd Annual Meeting on Association for Computational Linguistics* (pp. 363-370). Ann Arbor, Michigan.
- Fizman, M., Chapman, W. W., Evans, S. R., & Haug, P J. (1999). Automatic identification of pneumonia related concepts on chest x-ray reports. *AMIA* (pp. 67-71).
- Fox, C. (1989). A stop list for general text (Vol. 24, p. 21). ACM.
- Friedlin, J., Overhage, M., Al-Haddad, M. A., Waters, J. A., Aguilar-Saavedra, J. J. R., Kesterson, J., & Schmidt, M. (2010). Comparing Methods for Identifying Pancreatic Cancer Patients Using Electronic Data Sources. *AMIA Annual Symposium Proceedings* (Vol. 2010, p. 237).
- Friedlin, Jeff, & McDonald, C. J. (2006). A natural language processing system to extract and code concepts relating to congestive heart failure from chest radiology reports. *AMIA Annual Symposium Proceedings* (pp. 269-273). Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/17238345>
- Friedman, C., Kra, P., & Rzhetsky, A. (2002). Two biomedical sublanguages: a description based on the theories of Zellig Harris. *Journal of Biomedical Informatics*, 35(4), 222–235.
- Friedman, Carol. (1997). Towards a comprehensive medical language processing system: methods and issues. *AMIA*.

- Fundel, Katrin, Kuffner, R., & Zimmer, R. (2007). RelEx–Relation extraction using dependency parse trees. *Bioinformatics*, 23(3), 365-371.
- Gee, J. P. (1999). *An introduction to discourse analysis: Theory and method*. London: Routledge.
- Gemoets, D., Roseblat, G., Tse, T., & Logan, R. (2004). Assessing readability of consumer health information: an exploratory study. *Proceedings of the 11th World Congress on Medical Informatics* (p. 869).
- Gu, B. (2006). Recognizing nested named entities in GENIA corpus. *Proceedings of the Workshop on Linking Natural Language Processing and Biology: Towards Deeper Biological Literature Analysis* (pp. 112-113). Association for Computational Linguistics.
- Gurulingappa, H., Klinger, R., Hofmann-Apitius, M., & Fluck, J. (2010). An Empirical Evaluation of Resources for the Identification of Diseases and Adverse Effects in Biomedical Literature. *2nd Workshop on Building and Evaluating Resources for Biomedical Text Mining* (p. 15).
- Hakenberg, J., Leaman, R., Nguyen, V., Jonnalagadda, S, Sullivan, R., Miller, C., Tari, L, et al. (2010). Efficient extraction of protein-protein interactions from full-text articles. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*.
- Halliday, M. A. K. (1985). *An introduction to functional grammar*. London: Arnold.
- Harris, Z. S. (2002). The structure of science information. *Journal of biomedical informatics*, 35(4), 215–221.
- Harris, Z. (1954). Distributional Structure. *Word*, 10(2/3), 146–62.
- Harris, Z. (1991). *A theory of language and information : a mathematical approach*. Oxford [England]: Clarendon Press.
- Haug, P J, Koehler, S., Lau, L. M., Wang, P., Rocha, R., & Huff, S. M. (1995). Experience with a mixed semantic/syntactic parser. *Annual Symposium on Computer Applications in Medical Care*, 284-288.
- Haug, P J, Ranum, D. L., & Frederick, P. R. (1990). Computerized extraction of coded findings from free-text radiologic reports. Work in progress. *Radiology*, 174(2), 543-548.
- Hearst, M. A. (1999). Untangling text data mining. *Proceedings of the 37th annual meeting of the Association for Computational Linguistics* (pp. 3-10). College Park, Maryland. doi:10.3115/1034678.1034679
- Heinze, D. T., Morsch, M. L., & Holbrook, J. (2001). Mining free-text medical records. *Proceedings of the AMIA Symposium*, 254-258.

- Hirschman, L., Grishman, R., & Sager, N. (1976). From text to structured information: automatic processing of medical reports. *Proceedings of the National Computer Conference and Exposition* (p. 267). New York, NY, USA. doi:10.1145/1499799.1499842
- Hsu, C.-N., Chang, Y.-M., Kuo, C.-J., Lin, Y.-S., Huang, H.-S., & Chung, I.-F. (2008). Integrating high dimensional bi-directional parsing models for gene mention tagging. *Bioinformatics*, 24(13), i286-i294. doi:10.1093/bioinformatics/btn183
- Hu, M., & Liu, B. (2004). Mining and summarizing customer reviews. *Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining* (pp. 168-177). Seattle, WA, USA: ACM.
- Huang, M., Zhu, X., & Li, M. (2006). A hybrid method for relation extraction from biomedical literature. *International Journal of Medical Informatics*.
- Hunter, Lawrence, Lu, Zhiyong, Firby, J., Baumgartner, William, Johnson, Helen, Ogren, P., & Cohen, K. Bretonnel. (2008). OpenDMap: An open source, ontology-driven concept analysis engine, with applications to capturing knowledge regarding protein transport, protein interactions and cell-type-specific gene expression. *BMC Bioinformatics*, 9(1), 78.
- Jang, H., Lim, J., Lim, J.-H., Park, S.-J., Park, S.-H., & Lee, K.-C. (2006). Extracting Protein-Protein Interactions in Biomedical Literature Using an Existing Syntactic Parser. *Knowledge Discovery in Life Science Literature* (pp. 78-90).
- Johnson, W. B., & Lindenstrauss, J. (1984). Extensions of Lipschitz mappings into a Hilbert space. *Contemporary mathematics*, 26(189-206), 1-1.1.
- Jones, M. N., & Mewhort, D. J. K. (2007). Representing word meaning and order information in a composite holographic lexicon. *Psychological review*, 114(1), 1-37.
- Jonnalagadda, S., & Gonzalez, G. (2009). Sentence Simplification Aids Protein-Protein Interaction Extraction. *Languages in Biology and Medicine*.
- Jonnalagadda, Siddhartha, & Gonzalez, Graciela. (2010). BioSimplify: an open source sentence simplification engine to improve recall in automatic biomedical information extraction. *AMIA Annual Symposium Proceedings*.
- Kabiljo, R., Clegg, A. B., & Shepherd, A. J. (2009). A realistic assessment of methods for extracting gene/protein interactions from free text. *BMC Bioinformatics*, 10, 233.

- Kanerva, P., Kristofersson, J., & Holst, A. (2000). Random indexing of text samples for latent semantic analysis. *Proceedings of the 22nd Annual Conference of the Cognitive Science Society* (Vol. 1036). Citeseer.
- Kim, J. D., Ohta, T., Pyysalo, S., Kano, Y., & Tsujii, J. (2009). Overview of BioNLP'09 shared task on event extraction. *Proceedings of the Workshop on BioNLP: Shared Task* (p. 1-9).
- Kim, J. D., Ohta, T., Tateisi, Y., & Tsujii, J. (2003). GENIA corpus-a semantically annotated corpus for bio-textmining. *Bioinformatics-Oxford*, 19(1), 180-182.
- Kim, S., Shin, S. Y., Lee, I. H., Kim, S. J., Sriram, R., & Zhang, B. T. (2008). PIE: an online prediction system for protein-protein interactions from text. *Nucleic acids research*.
- Klein, D., & Manning, C. D. (2003). Accurate unlexicalized parsing. *Proceedings of the 41st Annual Meeting on Association for Computational Linguistics-Volume 1* (pp. 423-430). Association for Computational Linguistics Morristown, NJ, USA.
- Klinger, R., & Friedrich, C. M. (2009). Feature Subset Selection in Conditional Random Fields for Named Entity Recognition. *International Conference on Recent Advances in NLP*.
- Krallinger, M., Leitner, F., & Valencia, A. (2007). Assessment of the Second BioCreative PPI task: Automatic Extraction of Protein-Protein Interactions. *Proceedings of the Second BioCreative Challenge Workshop*, 41-54.
- Lafferty, J., McCallum, A., & Pereira, F. (2001). Conditional random fields: Probabilistic models for segmenting and labeling sequence data. *Proceedings of the Eighteenth International Conference on Machine Learning* (pp. 282-289).
- Leaman, R., & Gonzalez, G. (2008). BANNER: an executable survey of advances in biomedical named entity recognition. *Pacific Symposium in Bioinformatics*.
- Lease, M., & Charniak, E. (2005). Parsing Biomedical Literature. *International Joint Conference on Natural Language Processing*.
- Leitner, F., Chatr-aryamontri, A., Mardis, S. A., Ceol, A., Krallinger, M., Licata, L., Hirschman, L., et al. (2010). The FEBS Letters/BioCreative II.5 experiment: making biological information accessible. *Nature Biotechnology*, 28(9), 897-899. doi:10.1038/nbt0910-897
- Li, Y., Lin, H., & Yang, Z. (2009). Incorporating rich background knowledge for gene named entity classification and recognition. *BMC Bioinformatics*, 10(1), 223. doi:10.1186/1471-2105-10-223

- Liu, C., & Ng, H. T. (2007). Learning Predictive Structures for Semantic Role Labeling of NomBank. *Proceedings of the 45th Annual Meeting of the Association of Computational Linguistics* (p. 208–215). Prague, Czech Republic: Association for Computational Linguistics. Retrieved from <http://www.aclweb.org/anthology/P07-1027>
- Lund, K., & Burgess, C. (1996). Hyperspace analog to language (HAL): A general model of semantic representation. *Language and Cognitive Processes*.
- Lussier, Y. A., & Bodenreider, O. (2007). Clinical Ontologies for Discovery Applications. In C. J. O. Baker & K.-H. Cheung (Eds.), *Semantic Web* (pp. 101-119). Springer US.
- Marcotte, Edward M., Xenarios, I., & Eisenberg, D. (2001). Mining literature for protein-protein interactions. *Bioinformatics*, 17(4), 359-363.
- Marcus, M. P., Santorini, B., & Marcinkiewicz, M. A. (1993). Building a Large Annotated Corpus of English: The Penn Treebank. *Computational Linguistics*, 19(2), 313-330.
- Marneffe, M.-C. de, MacCartney, B., & Manning, Christopher D. (2006). Generating Typed Dependency Parses from Phrase Structure Parses. *LREC*.
- Marquez, L., Villarejo, L., Marti, M. A., & Taule, M. (2007). Semeval-2007 task 09: Multilevel semantic annotation of Catalan and Spanish. *Proceedings of the 4th International Workshop on Semantic Evaluations* (pp. 42-47). Association for Computational Linguistics.
- McCallum, A., & Li, W. (2003). Early results for named entity recognition with conditional random fields, feature induction and web-enhanced lexicons. *Seventh Conference on Natural Language Learning (CoNLL)*.
- McCallum, AK. (2002). MALLETT: A Machine Learning for Language Toolkit. Retrieved May 9, 2010, from <http://mallet.cs.umass.edu>
- McClosky, D., & Charniak, E. (2008). Self-Training for Biomedical Parsing. *ACL*.
- Meystre, S. M., Savova, G K, Kipper-Schuler, K. C., & Hurdle, J. F. (2008). Extracting information from textual documents in the electronic health record: a review of recent research. *Yearbook of Medical Informatics*, 128-144.
- Meystre, S., & Haug, P. J. (2005). Evaluation of medical problem extraction from electronic clinical documents using MetaMap Transfer (MMTx). *Studies in health technology and informatics*, 116, 823.
- Meystre, Stephane, & Haug, Peter J. (2005). Automation of a problem list using natural language processing. *BMC Medical Informatics and Decision Making*, 5, 30. doi:10.1186/1472-6947-5-30

- Meystre, Stéphane, & Haug, Peter J. (2006). Natural language processing to extract medical problems from electronic clinical documents: performance evaluation. *Journal of Biomedical Informatics*, 39(6), 589-599. doi:10.1016/j.jbi.2005.11.004
- Minka, T. (2005). *Discriminative models, not discriminative training* (Technical Report No. MSR-TR-2005-144). Microsoft Research. Retrieved from ftp://ftp.research.microsoft.com/pub/tr/TR-2005-144.pdf
- Miwa, M., Sætne, R., Miyao, Y., & Tsujii, Jun'ichi. (2010). Entity-focused sentence simplification for relation extraction. *Proceedings of the 23rd International Conference on Computational Linguistics, COLING '10* (p. 788-796). Stroudsburg, PA, USA: Association for Computational Linguistics.
- Morgan, A., Lu, Z., Wang, X., Cohen, A., Fluck, J., Ruch, P., Divoli, A., et al. (2008). Overview of BioCreative II gene normalization. *Genome biology*, 9(Suppl 2), S3.
- Noreen, E. W. (1989). *Computer-intensive methods for testing hypotheses: an introduction*. New York: John Wiley & Sons, Inc.
- Ono, T., Hishigaki, H., Tanigami, A., & Takagi, T. (2001). Automated extraction of information on protein-protein interactions from the biological literature. *Bioinformatics*, 17(2), 155-161.
- Phuong, T. M., Lee, D., & LeeKwang, H. (2003). Learning Rules to Extract Protein Interactions from Biomedical Text. *PAKDD* (pp. 148-158).
- Plate, T. A. (2002). Holographic reduced representations. *Neural Networks, IEEE Transactions on*, 6(3), 623-641.
- Platt, J. (1998). *Sequential minimal optimization: A fast algorithm for training support vector machines* (Technical Report No. MSR-TR-98-14). Microsoft Research.
- Porter, M. F. (1980). An algorithm for suffix stripping. *Program*, 14(3), 313-316.
- Pyysalo, S., Salakoski, T., Aubin, S., & Nazarenko, A. (2007). Lexical Adaptation of Link Grammar to the Biomedical Sublanguage: a Comparative Evaluation of Three Approaches. *BMC Bioinformatics*.
- Quirk, R., Greenbaum, S., Leech, G., & Svartvik, J. (1985). *A comprehensive grammar of the English language*. New York: Oxford Univ Press.
- Rapp, R. (2002). The computation of word associations: Comparing syntagmatic and paradigmatic approaches. *Proceedings of the 19th international conference on Computational linguistics-Volume 1* (p. 7). Association for Computational Linguistics.

- Rau, L. F., Res, G. E., Center, D., & Schenectady, N. Y. (1991). Extracting company names from text. *Seventh IEEE Conference on Artificial Intelligence Applications* (Vol. 1).
- Rebholz-Schuhmann, D., Arregui, M., Gaudan, S., Kirsch, H., & Jimeno, A. (2007). Text processing through Web services: Calling Whatizit. *Bioinformatics*.
- Roberts, A., Gaizauskas, R., Hepple, M., & Guo, Y. (2008). Mining clinical relationships from patient narratives. *BMC bioinformatics*, 9(Suppl 11), S3.
- Roberts, Angus, Gaizauskas, Robert, Hepple, Mark, & Guo, Yikun. (2008). Mining clinical relationships from patient narratives. *BMC Bioinformatics*, 9(Suppl 11), S3. doi:10.1186/1471-2105-9-S11-S3
- Sager, N. (1975). Sublanguage Grammers in Science Information Processing. *Journal of the American Society for Information Science*, 26(1), 10-16.
- Sahlgren, M. (2006). *The Word-space model* (PhD in Computational Linguistics). Stockholm University.
- Sahlgren, M., & Cöster, R. (2004). Using bag-of-concepts to improve the performance of support vector machines in text categorization. *Proceedings of the 20th international conference on Computational Linguistics* (p. 487).
- Sahlgren, M., Holst, A., & Kanerva, P. (2008). Permutations as a means to encode order in word space. *Proceedings of the 30th Annual Meeting of the Cognitive Science Society* (pp. 23-26).
- Saussure, F., Bally, C., Sechehaye, A., Riedlinger, A., Calvet, L. J., & Mauro, T. D. (1922). *Cours de linguistique générale*. Payot, Paris.
- Savova, Guergana K, Kipper-Schuler, K. C., Buntrock, J. D., & Chute, C. G. (2008). UIMA-based Clinical Information Extraction System. *LREC*.
- Schadow, G., & McDonald, C. J. (2003). Extracting structured information from free text pathology reports. *AMIA Symposium*, 584-588.
- Schneider, G., Kaljurand, K., & Rinaldi, F. (2009). Detecting protein-protein interactions in biomedical texts using a parser and linguistic resources. *Proceedings of the 10th International Conference on Computational Linguistics and Intelligent Text Processing, March* (p. 01-07).
- Settles, B. (2004). *Active Learning Literature Survey* (Technical Report No. 1648). Computer Sciences Technical Report (p. 201-221). University of Wisconsin-Madison.
- Settles, Burr. (2005). ABNER: an open source tool for automatically tagging genes, proteins and other entity names in text. *Bioinformatics*, 21(14), 3191-3192.

- Shen, D., Zhang, J., Zhou, G., Su, J., & Tan, C. L. (2003). Effective adaptation of a hidden markov model-based named entity recognizer for biomedical domain. *Proceedings of the ACL 2003 workshop on Natural language processing in biomedicine-Volume 13* (p. 56). Association for Computational Linguistics.
- Siddharthan, A. (2006). Syntactic Simplification and Text Cohesion. *Research on Language and Computation*, 4(1), 77-109.
- Sleator, D. (1998). Link Grammar Documentation.
- Smith, T. F., & Waterman, M. S. (1981). Identification of common molecular subsequences. *Journal of Molecular Biology*, 147, 195-197.
- Song, Y., Kim, E., Lee, G. G., & Yi, B. (2004). POSBIOTM-NER in the shared task of BioNLP/NLPBA 2004. *Proceedings of the International Joint Workshop on Natural Language Processing in Biomedicine and its Applications* (pp. 100-103). Association for Computational Linguistics.
- Stanfill, M. H., Williams, M., Fenton, S. H., Jenders, R. A., & Hersh, W. R. (2010). A systematic literature review of automated clinical coding and classification systems. *Journal of the American Medical Informatics Association: JAMIA*, 17(6), 646-651. doi:10.1136/jamia.2009.001024
- Sutton, C., & McCallum, A. (2007). An Introduction to Conditional Random Fields for Relational Learning. *Introduction to statistical relational learning*. Cambridge, Massachusetts, USA: MIT Press.
- Taira, R. K., & Soderland, S. G. (1999). A statistical natural language processor for medical reports. *AMIA Symposium*, 970-974.
- Tari, L., Tu, P. H., Hakenberg, J., Chen, Y., Son, T. C., Gonzalez, G., & Baral, C. (2009). Integrating querying and retrieval for biomedical information extraction. *ICDE*.
- Tateisi, Yuka, & Tsujii, Jun'ichi. (2004). Part-of-Speech Annotation of Biology Research Abstracts. *In the Proceedings of the 4th International Conference on Language Resource and Evaluation*, 1267-1270.
- Temkin, J. M., & Gilder, M. R. (2003). Extraction of protein interaction information from unstructured text using a context-free grammar. *Bioinformatics*, 19(16), 2046-2053.
- Thomas, J., Milward, D., Ouzounis, C., Pulman, S., & Carroll, M. (2000). Automatic extraction of protein interactions from scientific abstracts. *Pacific Symposium on Biocomputing*, 541-552.
- Tomanek, K., Wermter, J., & Hahn, U. (2007). A reappraisal of sentence and token splitting for life sciences documents. *Studies in health technology and informatics*, 129(Pt 1), 524-528.

- Traugott, E. (2002). *Regularity in semantic change*. Cambridge: Cambridge University Press.
- Trefethen, L. N., & Bau, D. (1997). *Numerical linear algebra*. Society for Industrial Mathematics.
- Tsai, R. T., Wu, S. H., Chou, W. C., Lin, Y. C., He, D., Hsiang, J., Sung, T. Y., et al. (2006). Various criteria in the evaluation of biomedical named entity recognition. *BMC bioinformatics*, 7, 92.
- Turian, J., Opérationnelle, R., Ratinov, L., & Bengio, Y. (2010). Word representations: A simple and general method for semi-supervised learning. *ACL* (Vol. 51, p. 61801).
- Uzuner, O., Mailoa, J., Ryan, R., & Sibanda, T. (2010). Semantic relations for problem-oriented medical records. *Artificial Intelligence in Medicine*. doi:10.1016/j.artmed.2010.05.006
- Vanderwende, L., Suzuki, H., Brockett, C., & Nenkova, A. (2007). Beyond SumBasic: Task-focused summarization with sentence simplification and lexical expansion. *Information processing and management*, 43(6), 1606–1618.
- Wang, X., Chused, A., Elhadad, N., Friedman, C., & Markatou, M. (2008). Automated knowledge acquisition from clinical narrative reports. *AMIA Annual Symposium Proceedings* (Vol. 2008, p. 783).
- Wang, X., Hripcsak, G., & Friedman, C. (2009). Characterizing environmental and phenotypic associations using information theory and electronic health records. *BMC bioinformatics*, 10(Suppl 9), S13.
- Wang, Y. (2009). Annotating and recognizing named entities in clinical notes. *Proceedings of the ACL-IJCNLP 2009 Student Research Workshop* (p. 18–26). Stroudsburg, PA, USA: Association for Computational Linguistics.
- Wicentowski, R., & Sydes, M. R. (2008). Using implicit information to identify smoking status in smoke-blind medical discharge summaries. *Journal of the American Medical Informatics Association: JAMIA*, 15(1), 29-31. doi:10.1197/jamia.M2440
- Widdows, D., & Cohen, T. (2010). The Semantic Vectors Package: New Algorithms and Public Tools for Distributional Semantics. *Fourth IEEE International Conference on Semantic Computing* (Vol. 1, p. 43).
- Wilbur, J., Smith, L., & Tanabe, T. (2007). BioCreative 2 Gene Mention Task. *Proceedings of the Second BioCreative Challenge Workshop*, 7-16.
- Witten, I. H., Frank, E., & Hall, M. (2011). *Weka: Practical machine learning tools and techniques with Java implementations*. Morgan Kaufmann.

- Wong, L. (2001). PIES, a protein interaction extraction system. *Pacific Symposium on Biocomputing*, 520-531.
- Xu, Hua, & Friedman, Carol. (2003). Facilitating research in pathology using natural language processing. *AMIA Symposium*, 1057.
- Zeng, Q. T., Goryachev, S., Weiss, S., Sordo, M., Murphy, S. N., & Lazarus, R. (2006). Extracting principal diagnosis, co-morbidity and smoking status for asthma research: evaluation of a natural language processing system. *BMC Medical Informatics and Decision Making*, 6, 30. doi:10.1186/1472-6947-6-30
- Zhou, G. D. (2006). Recognizing names in biomedical texts using mutual information independence model and SVM plus sigmoid. *International journal of medical informatics*, 75(6), 456-467.
- Zhou, G., Zhang, J., Su, J., Shen, D., & Tan, C. (2004). Recognizing names in biomedical texts: a machine learning approach. *Bioinformatics*, 20(7), 1178-1190.
- Zhou, Guodong, Shen, Dan, Zhang, Jie, Su, Jian, & Tan, S. (2005). Recognition of protein/gene names from text using an ensemble of classifiers. *BMC Bioinformatics*, 6(Suppl 1), S7. doi:10.1186/1471-2105-6-S1-S7
- Zhou, L., Parsons, S., & Hripcsak, George. (2008). The Evaluation of a Temporal Reasoning System in Processing Clinical Discharge Summaries. *Journal of the American Medical Informatics Association*, 15(1), 99-106. doi:10.1197/jamia.M2467
- Zhou, X., Han, H., Chankai, I., Prestrud, A., & Brooks, A. (2006). Approaches to text mining for clinical medical records. *Proceedings of the 2006 ACM symposium on Applied computing* (p. 235-239).