Concept-based embeddings for natural language processing

Ma, Yukun

2018

Ma, Y. (2018). Concept-based embeddings for natural language processing. Doctoral thesis, Nanyang Technological University, Singapore.

<http://hdl.handle.net/10356/75838>

<https://doi.org/10.32657/10356/75838>

CONCEPT-BASED EMBEDDINGS FOR NATURAL LANGUAGE PROCESSING

YUKUN MA

SCHOOL OF COMPUTER SCIENCE AND ENGINEERING

2018

CONCEPT-BASED EMBEDDINGS FOR NATURAL LANGUAGE PROCESSING

YUKUN MA

School of Computer Science and Engineering
A thesis submitted to the Nanyang Technological University
in fulfillment of the requirement for the degree of
Doctor of Philosophy

2018

Acknowledgments

First and foremost, I am deeply grateful to my supervisor Dr. Erik Cambria who has been extremely supportive during the last two years of my Ph.D. Without his guidance and assistance, it would be impossible for me to come to this final stage.

I also would like to extend my thanks to Dr. Jungjae Kim, my former supervisor, for bringing to me this program and helping me to set up the initial research profile.

I would like to thank Rolls-Royce and National Research Foundation Singapore for providing financial support for my research as well as my living in Singapore. I also would like to thank San Linn, the manager of our research project, for his effective coordination and kind support throughout the whole process.

Special thanks go to Dr. Benjamin Bigot, Dr. Tahir Khan, who are my colleagues and friends in the Rolls-Royce Corp Lab, for continuously encouraging and inspiring me with great thinkings.

I also have greatly benefited from the friendship of Dr. Luu Ann Tuan, Dr. Tran Ha Nguyen, Sa Gao, Haiyun Peng, Wen Song, Chai Woonhuei, Guanghao Zhang, and Zehong Hu. Thank Pei-hung Chen and Dr. Chieh Lu who are great friends of us and had stayed by my side during my toughest days in the small and hopeless student hostel. Thanks to all of them, my Ph.D. could end up as a pleasant and rewarding journey.

I would like to give extra credit to my parents, Jie Yang and Jianguo Ma, for that all the achievements in this thesis are results of their wise advice on selecting my undergraduate program. I owe enormous gratitude to my parents-in-law, Dongrong Tong and Xiaoying Xu, who have always supported our decisions and life both financially and emotionally. Finally, Qiying Tong, my wife and the most valuable of my life, definitely deserves my endless gratitude and love for her bearing with my stupidity and dullness since our first day. This thesis is fully dedicated to them.

Abstract

Concepts are critical semantics capturing the high-level knowledge of human language. As a way to go beyond the word-level analysis, representing and leveraging the concept-level information is an important add-on to existing natural language processing (NLP) systems. More specifically, the concepts are critical for understanding opinions of people. For example, people express their opinion towards particular entities such as products or sentiment aspects in online reviews, where these entities are mentions of concepts rather than just words.

As compared with words, the mentions of abstract concepts may be compounded phrases (either consecutive or non-consecutive) that are likely to form a large vocabulary. Furthermore, there might be semantic properties (e.g., relations or attributes) attached to the concepts, which increases the dimensionality of concepts. In short, using concepts is faced with the curse of dimensionality. On the other hand, information from only a single level does not suffice for a thorough understanding of human language, and meaningful representation is required at any point to encode the correlation and dependency between abstract concepts and words. In this thesis, we thus focus on effectively leveraging and integrating information from concept-level as well as word-level via projecting concepts and words into a lower dimensional space while retaining most critical semantics. In a broad context of opinion understanding system, we investigate the use of the fused embedding for several core NLP tasks: named entity detection and classification, automatic speech recognition reranking, and targeted sentiment analysis.

We first propose a novel method to inject the entity-based information into a word embedding space. The word embeddings are learned from a set of named entity features instead of merely contextual words. We demonstrate that the new word embedding is a better feature representation for detecting and classifying named entities from the stream of telephone conversations.

Apart from learning input feature embeddings, we then explore encoding the entity types (i.e., concept categories) in a label embedding space. Our label embeddings

mainly leverage two types of information: label hierarchy and label prototype. Since our label embedding is computed prior to the training process, it has exactly the same computation complexity at run-time. We evaluate the resulting label embeddings on multiple large-scale datasets built for the task of fine-grained named entity typing. As compared with the state-of-the-art methods, our label embedding method can achieve superior performance.

Next, we demonstrate that a binary embedding of the named entities can help reranking the speech-to-text hypothesis. Named entities are encoded using a Restricted Boltzmann Machine (RBM) and used as a prior knowledge in the discriminative reranking model. We also extend the training of RBM to work with speech recognition hypothesis.

Finally, we investigate the problem of using embeddings of commonsense concepts for the task of targeted sentiment analysis. The task is also entity-centered. Namely, given a targeted entity in a sentence, the task is to resolve the correct aspects categories and corresponding sentiment polarity of the target. We propose a new computation structure of Long Short-Term Memory (LSTM) that can more effectively incorporate the embeddings of commonsense knowledge.

In summary, this thesis proposes novel solutions of representing and leveraging concept-level and word-level information in a series of NLP tasks that are key to understanding the opinion of people.

Contents

Acknowledgments	
Abstract	iii
List of Figures	viii
List of Tables	x
1 Introduction	1
1.1 Preliminaries	2
1.1.1 Concepts	2
1.1.2 Concept-based Knowledge Resources	3
1.1.3 Concept-centered NLP for Opinion Understanding	4
1.2 Challenges and Problems Faced by Concept-centered NLP	7
1.3 Contribution	7
1.4 Organization	9
2 Background	1
2.1 Named Entity Recognition	1
2.1.1 Rule-based NER	1
2.1.2 Learning-based NER Model	2
2.1.3 Deep Neural NER Models	6
2.1.4 NER Features	7
2.2 Fine-grained NER	8
2.2.1 Fine-grained Labels	9
2.2.2 Automatic Generation of Training Data	9
2.2.3 Features	11
2.2.4 Segmentation and Multi-label Classification	11
2.2.5 Label Noise Reduction	14
2.2.6 Neural Architecture for FNET	15
2.3 Targeted Sentiment Analysis	16
2.3.1 Aspect-Based Sentiment Analysis	16

2.3.2	Targeted Sentiment Analysis	17
2.3.3	Targeted Aspect-Based Sentiment Analysis	17
2.3.4	Incorporating External Knowledge	18
2.4	ASR Hypotheses Reranking	18
2.5	Embedding Models	20
2.5.1	Count-based Models	21
2.5.2	Prediction-based Models	23
2.6	Summary	24
3	Learning Word Embedding with Multi-task Entity-based Features	25
3.1	Introduction	25
3.2	Methods	27
3.2.1	Baseline NER method	27
3.2.2	Skip-gram model	28
3.2.3	Feature extraction for word embeddings	29
3.2.4	Learning word embeddings from entity-based features	29
3.2.5	Multi-task interpretation of the proposed method	31
3.2.6	Using word embeddings as NER features	31
3.3	Experiments	32
3.3.1	Corpora	32
3.3.2	Evaluations	33
3.3.3	Summary	35
4	Label Embedding for Fine-grained Named Entity Typing	36
4.1	Introduction	36
4.2	Background	37
4.3	Embedding Methods for FNET	38
4.3.1	Joint Embedding Model	38
4.3.2	Zero-shot FNET Extension	39
4.4	Methods	40
4.4.1	Prototype-driven Label Embedding	40
4.4.2	Hierarchical Label Embedding	41
4.4.3	Prototype-driven Hierarchical Label Embedding	42
4.4.4	Type Inference	42
4.5	Experiments	43
4.5.1	Experiment Setup	43

4.5.2	Generating ProtoLE	45
4.5.3	Few-shots Fine-grained Entity Typing	45
4.5.4	Zero-shot Fine-grained Entity Typing	47
4.6	Summary	49
5	Binary Entity Embedding for Reranking ASR Hypotheses	50
5.1	Introduction	50
5.2	Training RBM for ASR Reranking	52
5.2.1	Restricted Boltzmann Machine	52
5.2.2	Maximum Margin Training for RBM-based Reranker	53
5.2.3	Training with Prior Knowledge	54
5.2.4	Scoring ASR Hypotheses	56
5.3	Related Work	56
5.4	Experiment	57
5.4.1	Dataset	57
5.4.2	Baseline	58
5.4.3	RBM Setup	58
5.4.4	Evaluation	59
5.5	Summary	62
6	Embedding Concepts for Target-based Sentiment Analysis	63
6.1	Introduction	63
6.2	Methodology	65
6.2.1	Task Definition	66
6.2.2	Overview	66
6.2.3	Long Short-Term Memory Network	67
6.2.4	Target-Level Attention	68
6.2.5	Sentence-Level Attention Model	68
6.2.6	Commonsense Knowledge	69
6.2.7	Sentic LSTM	70
6.2.8	Prediction and Parameter Learning	72
6.3	Experiments	72
6.3.1	Dataset and Resources	72
6.3.2	Experiment Setting	73
6.3.3	Performance Comparison	74
6.3.4	Results of Attention Model	76

6.3.5	Visualization of Attention	77
6.3.6	Results of Knowledge-Embedded LSTM	78
6.3.7	AffectiveSpace v.s. SentiWordNet	78
6.4	Summary	79
7	Summary and Future Work	81
7.1	Summary of Proposed Methods	81
7.2	Future Work	83
7.2.1	Limitation of Our Work	83
7.2.2	Graph-based Concept Embedding	84
7.2.3	Concept-oriented Graph Model	85
7.2.4	New Sentiment Analysis Application using Concepts	85

List of Figures

1.1	NLP curve for the evolution of NLP techniques by Cambria et al. [1]	4
1.2	A pipeline system [2] illustrating the key components of an opinion understanding system. Tasks selected as the focus of this thesis are highlighted in red boxes.	5
2.1	Structure of linear-chain CRF [3]	5
2.2	Neural architecture of LSTM-CRF used in Guillaume et al. [4]	7
2.3	Hierarchical categories used in Ling et al. [5]	10
2.4	Example of anchor text on a Wikipedia page.	10
2.5	Label in the projected space [6]	14
2.6	Neural Architecture used by Shimaoka et al. [7]	15
2.7	Training process of Skip-gram model.	24
3.1	NER Framework with NER-specific word embeddings	27
3.2	Performance with varying size of NER training data	34
4.1	t-SNE visualization of the prototype-driven label embeddings for BBN dataset [8]	44
4.2	BNN Dataset	46
4.3	OntoNotes Dataset	46
4.4	OntoNotes Dataset	46
4.5	Performance changes on the development set with regard to the sizes of prototype list [8]	46
5.1	Structure of RBM with Entity-related Prior [9]	55
5.2	WER reduction for words versus TF-IDF scores [9]	61
6.1	Example of Sentihood data set	64
6.2	Overview of the attentive neural architecture	66
6.3	A sketch of SenticNet semantic network [10]	75
6.4	Example of sentence-level attention	78

6.5	Example of target-level attention	79
-----	---	----

List of Tables

2.1	Example BIO tagging for a sample of MUC-7 [11]	3
2.2	Features shared by all the four related work	11
2.3	The dependency tuples for sentence <i>I WENT TO THE AN- THE DOC-TER</i>	20
3.1	Templates of features used in the CRF baseline	28
3.2	Features learned from word embeddings for NER	32
3.3	Named entity distribution in Switchboard	32
3.4	Details on baseline Word Embeddings preparation	33
3.5	Performance of NER methods in terms of F-score	33
3.6	Impact of feature subsets on Skip _{NER} [12]	34
3.7	Top similar words returned by Skip-gram & Skip _{NER}	35
4.1	Feature Template	43
4.2	Statistics of datasets	44
4.3	Example prototypes learned by PMI for types in BBN dataset	45
4.4	Performance of FNET in a few-shots learning on 3 benchmark datasets	48
4.5	Performance of zero-shot entity typing	49
5.1	Characteristics of the datasets used in experiments	58
5.2	Most activated word by p-dRBM	59
5.3	Example topics learned by p-dRBM	60
5.4	Most-similar words for queries using p-dRBM word embeddings	60
5.5	Performance of reranking model on TedLium corpus	60
6.1	Example of SenticNet assertions	70
6.2	SentiHood dataset	73
6.3	System performance on SentiHood dataset	74
6.4	System performance on Semeval-2015 dataset	75

6.5	Comparison of systems using AffectiveSpace and SentiWordNet (SH stands for Sentihood, and SE stands for SemEval-15)	79
7.1	Summarization of our proposed work	83

List of Abbreviations

AI	Artificial Intelligence
NLP	Natural Language Processing
ABSA	Aspect-based Sentiment Analysis
ASR	Automatic Speech Recognition
DNN	Deep Neural Network
RNN	Recurrent Neural Network
NER	Named Entity Recognition
FNET	Fine-grained Named Entity Typing
RBM	Restricted Boltzman Machine
LSTM	Long Short-Term Memory
Bi-LSTM	Bidirectional LSTM
TDLSTM	Target Dependent LSTM
TCLSTM	Target Connection LSTM
TSA	Targeted Sentiment Analysis
CRF	Conditional Random Field
PMI	Pointwise Mutual Information

Chapter 1

Introduction

The understanding of human language involves processing information beyond the words. A complete process should be capable of handling information from various levels such as concepts or thoughts. Concepts, as one of the most basic units of language, refer to aspects of the world that are used by people in their language. Concepts are constituents including mostly phrases such as noun phrases, verb phrases, adjective phrases and so on and so forth. In addition, the definition of concepts is broad enough to include mentions of named entities which can be deemed as concepts belonging to certain categories, e.g., location, organization, or person.

Recently, a variety of knowledge resources such as ConceptNet have been built from the huge unstructured information thanks to the rapid growth of online communities and social network. These knowledge resources are based on concepts rather than single words and are in the form of either a directed or undirected graph with concepts being the nodes and their relations being the edges. Additional attributes or properties might also be incorporated in order to better represent the semantics of concepts. These concept-based resources might benefit many NLP systems via providing accesses to external knowledge.

On the other hand, concepts are playing a core role in understanding the opinions of people. For example, people tend to express their sentiment towards a particular targeted entity (e.g., purchased products) in online reviews. By our definition of concepts, these entities are actually mentions of concepts. Therefore, the opinion understanding system would benefit from the use of concept-level information. However,

due to the complexity of natural language, a complete analysis of people’s opinions requires performing a series of subtasks such as concept extraction (or named entity recognition) and sentiment analysis, and it is hard to train a single end-to-end trained model that can perform well on all the subtasks of an opinion understanding system. Instead, it typically needs to build a pipeline composed of a series of natural language processing modules. In this thesis, we choose to extend the use of concepts to three important modules in the whole context: named entity recognition, targeted sentiment analysis and ASR hypotheses reranking. These tasks are core components of a multimodal opinion understanding system which is capable of processing not only text data but also audios. One of the natural obstacle hindering concepts from being employed in large scale is that concepts are generally of high dimensionality. Hence, our solutions focus on generating and leveraging concept-based embeddings for solving concept-centered NLP problems.

1.1 Preliminaries

We first introduce the definition of concepts as well as two knowledge resources that have been widely adopted for NLP tasks. We then give a brief introduction of concept-centered NLP tasks that we are interested in addressing in this thesis: named entity recognition, fine-grained named entity typing, targeted sentiment analysis, and ASR hypotheses reranking.

1.1.1 Concepts

- **Concepts** are used by people to describe the objects or actions perceived. Concepts might have different forms in different contexts. Admittedly, the surface form of concepts are composed of single or multiple words. For example, the concept **small room** is composed of two words **small** and **room**. Inflections of concepts are prevalent in natural language. Namely, the same objects or action might be referred to as different concepts due to many reasons including cultural differences, the existence of synonyms and acronyms, and so on. We could deem a concept as an abstract term having a canonical form as well as several surface

forms or mentions. For example, the concept of which the surface is “New York” might have mentions like “New York city” or “big apple”.

- Named entities [11, 13] can also be deemed as concepts. In the very first challenge of named entity recognition [11], named entities are defined as noun phrases in its surface form and refer to objects belonging to categories such as person, location, and organization. Although different challenges such as CoNLL 2003 shared tasks have extended the definition to involve other concepts, a standard definition of named entities has been constrained to consider only flat categories (i.e., non-hierarchical). Notably, a recent study on fine-grained named entities has further extended the concern with named entities to involve a hierarchical structure where IsA relation between concepts is taken into consideration.

1.1.2 Concept-based Knowledge Resources

Recently, knowledge bases such as ConceptNet [14] have linked concepts to more structural semantic information, e.g., relations and attributes. These knowledge bases are built automatically by employing rule-based and machine learning techniques to extract facts from Web texts. Therefore, the coverage of concepts could keep growing over the time.

1.1.2.1 ConceptNet

ConceptNet is automatically built as a combination and filtering of many existing and ever-growing multilingual knowledge resources such as Open Mind. ConceptNet keeps updating itself according to the source knowledge bases. One notable feature of ConceptNet is that it contains a large number of relational knowledge. Namely, we could be able to know the relations between concepts. For example, from ConceptNet, it can be known that **the white house** (argument 1) is **located** (relation) in **america** (argument 2) and **a type of** (relation) **building** (argument 3).

Figure 1.1: NLP curve for the evolution of NLP techniques by Cambria et al. [1]

1.1.2.2 SenticNet

Same as ConceptNet, SenticNet [15] merges the information in many existing knowledge resources. At the core of SenticNet is AffecNet where each concept is attached to a set of affective properties.

1.1.3 Concept-centered NLP for Opinion Understanding

The evolution of NLP techniques clearly shows a trend moving from word-level to concept-level (as illustrated in Fig. 1.1). Merely relying on word-level analysis of natural language is limited in the power of expressiveness since the natural language is composed of different levels of abstraction. It is, therefore, advantageous to use concepts for NLP tasks for that it allows for access to the rich semantic features of particular concepts. For example, collecting information about battery life and price is helpful when comparing the two products “iPhone X” and “Samsung S8”.

The concept-level analysis has been adopted for NLP tasks such as sentiment analysis and named entity recognition. However, concepts are very sparse due to the

Figure 1.2: A pipeline system [2] illustrating the key components of an opinion understanding system. Tasks selected as the focus of this thesis are highlighted in red boxes.

overwhelming number of unique concepts as well as concept inflections and features. A more compressed and effective representation is required to encode the concepts. On the other hand, a dense and lower dimensional representation also facilitate the exploration of deep learning models that have recently gained great successes in many NLP tasks. With the access to knowledge bases, we are able to augment NLP solutions with concept-level semantics.

On the other hand, understanding the opinions expressed by human beings is one of the ultimate objectives of NLP. As shown in Fig. 1.2, an opinion understanding system requires the composition of multiple NLP modules involving natural language concepts, from which we identify the following tasks that either requires performing concept-level analysis or can benefit from using concept-level information:

- **Named entity recognition (NER).** Named entities are typically belonging to several semantic classes defined according to the particular interest of downstream applications. For example, a flight booking assistant system might require recognizing the location names corresponding to departure and arrival cities is the most critical step for retrieving the best matched flight. A typical setting of NER is composed of both tasks of detection and classification.

In general, an NER model might take as input a sequence of tokens as well as a predefined set of entity categories, and then outputs the boundaries and semantic types of named entities presenting in the input.

- **Fine-grained named entity typing (FNET)**. Typical setting of NER models aims at resolving only the coarse-grained types of named entities. As an extension of the classification task in NER, FNET aims at recognizing the fine-grained types of a named entity, which are typically organized in a tree-shape hierarchy [16]. A natural formulation of FNET is a multi-label multi-class classification problem of which one entity is assigned with multiple categories.
- **Targeted sentiment analysis (TSA)**. The targeted sentiment analysis is another concept-centered NLP task resembling the task of FNET. The input to TSA is identical to FNET in which the region of the targeted entity is presumed known, and the ultimate goal is mapping the targeted entity to its sentiment aspects and polarity labels (e.g., positive and negative). More specifically, the task of TSA is composed of two subtasks: 1) aspect categorization and 2) polarity classification. For example, the desired outputs for a product review *Pixel 2 XL is cheaper than Iphone X* are the sentiment aspect, **price**, and the corresponding sentiment polarity label, **negative**, given the target is **Iphone X**.

Apart from these concept-centered NLP tasks, we also extend the use of concept-based information to automatic speech recognition (ASR) candidate re-ranking, which could augment the opinion understanding system with the capability to perform on data of multi-modality. We assume that concepts such as named entities deserve more attention than functional words for their critical roles for downstream tasks such as spoken language understanding systems or a multimodal sentiment analysis system. Therefore, given a set of ASR decoding candidates, we hope the entity-based information could help us select the candidate the makes fewer mistakes in recognizing important keywords such as person names and location names.

1.2 Challenges and Problems Faced by Concept-centered NLP

In the past decade, great progresses have been made on concept-centered tasks, especially for tasks such as named entity recognition that have been studied for decades. In general, whenever concepts are concerned in these NLP tasks, it is naturally desirable to have an effective representation as well as the corresponding computational structure to leverage the concept-level information.

For knowledge base such as ConceptNet that extracts conceptual facts from the Web automatically, the number of concepts will grow over time. On the other hand, the number of semantic properties (or relations) associated with each concept might also be an obstacle for employing conceptual analysis.

Therefore, a better representation of concepts basically requires tackling two problems: 1) producing a condensed representation of the original feature space, and 2) retain the semantics encoded in the raw data. Ideally, the compressed representation minimizes the dimensionality and information loss at the same time. However, it is hard to fulfill in practice, and a trade-off must be taken in such case.

Apart from an efficient and effective dense representation of concepts, leveraging conceptual information also needs well-designed computation structure in order to be integrated with information flowing from other levels (e.g., word level).

Most importantly, we believe using only information from a single-level does not suffice to achieve a complete analysis of the natural language. Therefore, the fusion of multi-level information is required for advancing the NLP research.

1.3 Contribution

The focus of this thesis is to leverage concept embeddings for concept-centered natural language processing tasks discussed in the last section and investigating on how to fuse the information of multiple levels.

- We investigate several novel solutions on fusing the concept-level and word-level information. We demonstrate the effectiveness of such fusion with multiple concept-centered NLP tasks.
- We propose a novel method to inject concept-based information into the word embedding space which then being used as effective features to recognize concepts/entities from input sentences.

(**Yukun Ma**, Jung-jae Kim, Benjamin Bigot, and Tahir Muhammad Khan. “Feature-enriched word embeddings for named entity recognition in open-domain conversations.” In Acoustics, Speech and Signal Processing (ICASSP), 2016 IEEE International Conference on, pp. 6055-6059. IEEE, 2016.)

- We explore to use the pre-trained representation of concept categories (entity labels) in the task of fine-grained named entity typing. We demonstrate that the concept category embedding can help resolve the fine-grained categories of entities. (**Yukun Ma**, Erik Cambria, and Sa Gao. “Label Embedding for Zero-shot Fine-grained Named Entity Typing.” In COLING, pp. 171-180. 2016.)
- We propose a binary embedding of named entities to help ranking hypotheses of speech recognition. (**Yukun Ma**, Erik Cambria, and Benjamin Bigot. “ASR Hypothesis Reranking using Prior-informed Restricted Boltzmann Machine.” In International Conference on Computational Linguistics and Intelligent Text Processing (CiCLing), 2017)
- We propose an extension of the Long short-term memory model that could leverage both concept-level and word-level information with more effectiveness. (**Yukun Ma**, Haiyun Peng, and Erik Cambria. “Targeted Aspect-Based Sentiment Analysis via Embedding Commonsense Knowledge into an Attentive LSTM”, In the Thirty-Second AAAI Conference on Artificial Intelligence (AAAI-18), pp. 5876-5883, 2018)

1.4 Organization

The rest of this thesis is organized as follows: Chapter 2 surveys the background and existing work on concept embeddings and concept-centered NLP tasks; Chapter 3 describes our method on injecting entity-based information into word embedding space to help recognize named entities from telephone conversations; Chapter 4 describes the concept-category embedding used for resolve fine-grained categories of named entities; Chapter 5 presents our binary entity embedding for reranking ASR hypotheses; finally, Chapter 6 presents our work on leveraging commonsense concepts for analyzing sentiment with respect to targeted entities.

Chapter 2

Background

In this chapter, we review the related work on NER, ASR reranking, TSA, as well as recent progress on learning concept or word embeddings.

2.1 Named Entity Recognition

Recognizing named entities from texts such as news articles is a concept-level NLP task has been studied by the community. Typically, recognition of named entity requires jointly resolving the location and category of a named entity in its context. In other words, the task could also be split into two tasks: detection and classification (or called Entity Typing). We start with introducing the related work on NER ranging from rule-based approaches of early days to most recent deep neural networks. We then extend our interest to fine-grained entity typing that has recently attracted increasing attention.

2.1.1 Rule-based NER

Early attempts to build NER systems mainly rely on hand-written grammar and rules designed by domain experts for NER challenges, e.g., MUC-7 [17]. The context of named entities is critical for rule-based approaches for it characterizing the mention of an entity. For example, one rule can be that any single word following “Mr” should be tagged as “PER”. Such rules encode information such as graphical information (capitalization, punctuation), syntactic (Part-of-speech tag, syntactic structure) or even semantics (word meanings) related to the current word itself as well as its context.

Synsets (country, person etc.) in WordNet are used as the reference of named entity types [18]. All the words in the same synset are sent as queries to a search engine to retrieve relevant documents. They then collected words appearing in the context of the query words and calculated the frequency. For an input word along with its context, they compared the context words with the contexts of synsets. The word will be assigned with the corresponding named entity type of the synset whose context is most similar to that of the given word.

Evans et al. [19] built a document-specific typology for named entities using manually designed patterns, e.g., “Y such as X” or “Y like X”. Then, they performed clustering on the set of hypernyms in order to identify potential named entity types. Afterwards, a classification was trained with the typology.

Rule-based approaches rely on hand-crafted rules that are expensive to obtain and require additional efforts when adapted to new data or domain. In addition, rule-based approaches are also vulnerable to noises and variations, even though some of the rules can be encoded as Finite-State-Transducer which increases its capability of being generalized.

2.1.2 Learning-based NER Model

Rather than using manual rules encoding the knowledge of experts, machine learning approaches derive knowledge in a data-driven manner and have been widely used for the task of NER. Machine learning approaches typically model the two subtasks of NER as a sequence prediction task using machine learning based approach: given a word sequence $W = w_0 w_1 w_2 \dots w_n$, the objective of NER is to find the optimal tag sequence $T = t_0 t_1 t_2 \dots t_n$, where the tag set is pre-defined. One of the most popular tagging schema used by NER approaches is the BIO encoding. The set of BIO tags includes **B**eginning, **I**nside, and **O**utside tags. For example, the BIO scheme for person names includes B_{Pers} , I_{Pers} and O , which indicate the beginning of a person name, an internal word of a person name and the word belonging to none of the named entity types. Table 2.1 shows a more specific example of MUC7 where the start and end of *Martin Puris* is indicated by “B-PER” and “I-PER” respectively. Given the

tagging schema and word sequence, the sequence labeling model searches for the best tag sequence with the parameters learned from training process. In the following, we review the popular sequence labeling models having been used for NER.

Mr.(O) Dooner(B-PER) met(O) with(O) Martin(B-PER) Puris(I-PER) ,(O) president(O) and(O) chief(O) executive(O) officer(O) of(O) Ammirati(B-ORG) &(I-ORG) Puris(I-ORG) ,(O) about(O) McCann(B-ORG) s(O) acquiring(O) the(O) agency(O) with(O) billings(O) of(O) \$400(B-MONEY) million(I-MONEY) ,(O) but(O) nothing(O) has(O) materialized(O) .(O)

Table 2.1: Example BIO tagging for a sample of MUC-7 [11]

Charniak et al. [20] proposes a Hidden Markov Model(HMM) for Part-of-speech tagging. Given the word sequence $W = w_0 w_1 w_2 \dots w_n$, the best tagging sequence $T = t_0 t_1 t_2 \dots t_n$ is defined as

$$\hat{T}_{1\dots n} = \arg \max_{T_{1\dots n}} P(T_{1\dots n}, W_{1\dots n})$$

HMM model assumes that the current word w_i is generated from a multinomial distribution conditioned on current tag while current tag is only conditioned on previous tags. In a first order HMM, the above joint probability can be expanded to

$$P(T_{1\dots n}, W_{1\dots n}) \approx \prod_{i=1}^n P(t_i | t_{i-1}) P(w_i | t_i)$$

However, the assumption made by HMM is too strong for NER as the current tag depends on not only the previous tags but also the other types of characteristics in context. For example, in Table 2.1, HMM model assumes that the tag **B-ORG** for word **Ammirati** only depends on the previous tag which is **O**. This makes less sense than assuming **B-ORG** depends on the previous word **officer**. Classic HMM does not allow such dependency between tags and words. Many variations of HMM model have been proposed for addressing this problem, most of which relax the Markov assumption to allow the dependency between previous tags and observations.

HMM models are not flexible enough to train with a rich feature set as the dependency between words and tags are strictly constrained. In order to overcome this shortcoming, Maximum Entropy (ME) model is proposed as an alternative to HMM.

It has been widely applied to tagging tasks such as POS tagging for it has the freedom of using arbitrary features.

Similar to HMM-based NER, ME-based model assigns each word in the input with either a tag corresponding to the named entity type and indicating if it is the beginning, or a tag indicating it is not part of an named entity.

Given the word sequence $W_{1...n}$, the probability of tag sequence $T_{1...n}$ is calculated as:

$$P(T_{1...n}|W_{1...n}) = \frac{\exp \sum_j^m \alpha_j \phi_j(T_{1...n}, W_{1...n})}{Z(T_{1...n})}$$

where $Z(T_{1...n})$ is the partition function defined as:

$$Z(T_{1...n}) = \sum_T \exp \sum_j^m \alpha_j \phi_j(T_{1...n}, W_{1...n})$$

The objective of learning process of ME model is to find the optimal setting of parameters α_j , and during the inference phase, a Viterbi decoding algorithm can be used for finding the most likely tagging sequence for the given input.

One problem with ME model is known as the *label bias* problem. When there is no alternative transition from one certain state to another, the model will ignore the observation. In other words, when estimating $P(t_i|t_{i-1}, w_i)$, the observed word w_i will not have any influence on the model if t_i is the only outgoing transition from t_{i-1} . As an alternative to ME model, Conditional Random Field (CRF) [3] is proposed in order to avoid such label bias problem by estimating $P(t_i|t_{i-1}, w_i)$ based on the complete sequence of tags. CRF shares the same exponential probabilistic form as ME model except that CRF model is an undirected graph model. Both HMM and ME are directed models while CRF uses an undirected structure which allows the dependency between words and tags to go in both directions. However, long-range dependency becomes intractable as the size of possible features grows exponentially with the number of words in the sequence. To significantly reduce the possible searching space, a simplified version of CRF models allowing only dependency of tags between consecutive time slots, i.e., t_{i-1} and t_i , which is called linear-chain CRF (as shown in Figure 2.1).

Figure 2.1: Structure of linear-chain CRF [3]

As defined by Lafferty et al. [3], linear-chain CRF defines the conditional probability of T given the word sequence W :

$$P(T|W) = \frac{\exp \sum_i \sum_j \alpha_j \phi_j(t_i, t_{i-1}, W)}{Z(T)}$$

while Z is the partition function

$$Z(T) = \sum_t \exp \sum_i \sum_j \alpha_j \phi_j(t_i, t_{i-1}, W)$$

, where T is the tag sequence and W is the word sequence; α_j is the weight associated with the j^{th} feature; and ϕ_j is the j^{th} feature function. Based on the above definition, one of the advantages of CRF is that it dynamically models the dependency between possible tags. In addition, it allows us to define arbitrary feature functions for T and W .

Like ME model, CRF model provides the freedom for defining arbitrary features on the sequence of words and tags. However, note that this does not mean that there are no constraints on the dependency. Long range dependency will result in intractable Z as the size of possible tagging sequence is exponential to the length of the word sequence. Similar to first order HMM or second order HMM, only previous two tags are included in the feature set for tractable computation.

2.1.3 Deep Neural NER Models

Most recently, sequence modeling has also benefited from the series of successes achieved by deep neural models. In general, deep neural models relieve the requirement for manually designed features and thus can be effortlessly adapted to problems sharing similar settings. For example, Collobert et al. [21] builds a deep multi-layer network that can be applied to both POS tagging and NER without the requirement of any adaptation. From a different point of view, deep neural models could serve as better feature extractors for discriminative models such as CRF and MaxEnt. In practice, as have been discussed before, many discriminative models such as MaxEnt and CRF have difficulties in using global features and usually have to resort to window-based simplification (e.g., using n-gram-based features). In comparison, deep neural nets such as Recurrent Neural Network (RNN) are capable of encoding the information of the whole sequence thanks to its computational structure (e.g., recurrent connection).

Recurrent neural networks are powerful in modeling sequential data. A RNN takes as input a sequence of words $x_1, x_2, \dots, x_n$, and outputs a sequence of hidden states $h_1, h_2, \dots, h_n$, with each h_i representing the subsequence from the beginning to i th time step. Despite its theoretical capability of encoding the long term dependency, it turns out that the vanilla setting of RNN tends to capture only the most recent input. Besides, the vanilla setting of RNN is also faced with a serious problem known as gradient vanishment. Namely, the gradient would either vanish or explode when being propagated through many time steps. To overcome this shortcoming, a variant of RNN, called LSTM [22], is proposed with a memory cell as well as several gates controlling the information flow from time to time. Another issue with a single RNN is that it encodes only the unidirectional sequential information. This is suboptimal in the sense that the hidden output at each timestep provides no access to the information of the incoming word sequence. Therefore, a simple yet effective extension is to simultaneously use the outputs of a forward RNN and a backward RNN, which is referred to as bidirectional RNN.

Figure 2.2: Neural architecture of LSTM-CRF used in Guillaume et al. [4]

RNNs became natural choices for modeling NER problem as soon as they are invented. As compared with previous work such as CRF or HMM, RNNs are shown better at leveraging global information lying with the whole sequence. But, RNNs fail to model the dependency between outputs of different time steps which can be easily achieved by models like CRF or HMM. Most recently, NER models [23, 4] using a neural structure combining the LSTM with CRF components (as shown in Figure 2.2) have also been proposed. In the new architecture, a bidirectional LSTM is used as a feature extractor, while a linear-chain CRF component is used for modeling the dependency between the output tags of neighboring timesteps.

2.1.4 NER Features

All the aforementioned supervised NER models rely on a feature set to represent the inputting sequence of words. Therefore, it is critical to design and use features capturing the most discriminative information marking the presence and categories of a named entity. These features basically fall into two categories: manual designed features and automatically learned features. The following are a subset of features used by Collin et al. [24],

Features

Current word

Previous word

Next word

Previous tag

Tag bi-grams (previous two tags)

If the word is the first word of a sentence

If the word appears as uppercase more frequently than lower case

Type of the initial word, i.e., uppercase, lower case, digit, etc.

Word shape, i.e., each character in the word is mapped to its type

It can be found that some of the manually designed features shares the same intuition of the hand-crafted rules used by a rule-based system. Basically, the feature set consists of grapheme features of the word at the current step as well as previous and next steps as contextual markers of the entities.

The manually designed features require human knowledge and are hard to be generalized to new domains or new datasets. In comparison, the second class of features are data-driven and can be automatically learned without the need for intervention. Most notably, NER models [21, 4, 23] using pre-trained word embeddings (or phrase embedding) as features have successfully achieved state-of-the-art performance. Word embedding is a type of vector representation of words that is typically learned via auxiliary tasks such as language modeling [25]. We will cover more reviews on the technique of learning embeddings of words or concepts in section 2.5.

2.2 Fine-grained NER

Fine-grained named entity recognition refers to the task of recognizing named entities belonging to a fine-grained set of entity categories. As compared with coarse-grained NER, the fine-grained entity types are organized into a hierarchy. In contrast to coarse-grained NER, the fine-grained NER is generally modeled in two separate steps: 1) segmentation, which is to detect the boundaries of named entity mentions; and 2) classification, which is to classify the detected named entity into fine-grained types. It is also likely that the classifier might have accesses to pre-existing entity boundaries rather than reading the outputs of the detection model. In such case, classification of

entity mentions becomes the only focus of Fine-grained NER, which is also known as the task of Fine-grained Entity Typing (FNET).

One of the biggest challenges faced by fine-grained named entity recognition is the shortage of training data. Since the number of fine-grained types is much bigger than coarse-grained of named entities, it is harder to have sufficient training data for each category. Fortunately, structural information in public resources such as Wikipedia and Freebase can be used for automatically building annotation of fine-grained named entity types. In this section, we review the most recent work on FINER or FNET: FIGER [16], HYENA [26], GFT [5], the embedding based method [yogatama:2015](#), as well as methods based on deep neural networks [7, 27].

2.2.1 Fine-grained Labels

The first step to creating a benchmark for fine-grained NER is identifying the set of labels (or classes) of named entities. Different from coarse-grained NER, whose label set is typically limited to a few categories such as Organization, Location, and Person, there’s actually no standards for fine-grained types. However, previous works have used online knowledge base for building the label set as well as their hierarchy. FIGER derives a set of 112 named entity types from Freebase, an online knowledge base built on Wikipedia and WordNet, even though the hierarchy only has two levels. HYENA uses a much larger tagset composed of 505 entity types by selecting top 100 subtypes for each top-level types of YAGO knowledge base [28, 29]. Additionally, the labels used in their approach are organized as a taxonomic hierarchy. More recently, GFT manually refined and expanded the label hierarchy used by FIGER into three levels. As shown in Fig. 2.3, the labels are organized into a three-level hierarchy, where the first level contains only four categories, i.e., PERSON, LOCATION, ORGANIZATION, and OTHER (or MISC), which are typically used in coarse NER.

2.2.2 Automatic Generation of Training Data

Training classifiers for fine-grained entity types requires a considerable number of labeled data, which is expensive and time-consuming. Instead, previous methods

PERSON	LOCATION	ORGANIZATION	OTHER	
artist actor author director music	structure airport government hospital hotel restaurant sports facility theatre	company broadcast news	art broadcast film music stage writing	language programming language
education student teacher	geography body of water island mountain	education government military music political party sports league sports team stock exchange transit	event accident election holiday natural disaster protest sports event violent conflict	living thing animal
athlete business coach doctor legal military political figure religious leader title	transit bridge railway road		product camera car computer mobile phone software weapon	
	celestial city country park		health malady treatment	food heritage internet legal religion scientific sports & leisure supernatural
			award body part currency	

Figure 2.3: Hierarchical categories used in Ling et al. [5]

have been trying to use online resources such as Wikipedia to automatically construct labeled data. Most of the times, anchor text in Wikipedia is in fact marking a mention of named entity. The link of anchor text redirects to the Wiki page corresponding to the classic form of that given named entity mention (as shown in Fig 2.4). Despite that the Wikipedia page has already the categories of that named entity, they are very noisy as many of them are automatically given by the system of Wikipedia. Instead of using the Wikipedia categories, Ling et al. [16] propose to use the freebase categories which contains fewer noises. By creating the mappings from the classic form to a freebase category and from freebase to manually refined labels, a large annotated corpus (approx. 3 million sentences) is built from the anchor texts of Wikipedia pages and Freebase.

altitudes significantly higher than the rest of the mountains. From the north
[Mount St. Helens](#), and [Mount Adams](#). Mount St. Helens is currently the only
active volcanoes. Mount Rainier ([Mount Adams \(Washington\)](#)) is in the state, [6] is 5
visible. The 14,411-foot-tall (4,392 m) Mt. Rainier is considered the most
[metropolitan area](#), and most dangerous in the continental U.S. according to

Figure 2.4: Example of anchor text on a Wikipedia page.

However, the labels in the annotated data built by [16] still contain noises, because the mapping from classic forms to freebase concepts do not consider the impacts of

context during tagging. For example, *ishington* can be a *military_leader* or a *city* according to Freebase, but in FIGER data, no matter what the contexts is, *Washington* is always annotated with both the two labels. To overcome this shortcoming, [5] used the context of named entities to refine the labels with the help of a set of heuristic rules which takes into account the contexts of named entities: **Sibling**, using the parent label, if the mention is tagged with multiple labels having the same parents; **Coarse-type**, using softmax classifiers trained on ACE corpus, which classify named entity mentions into top-level labels based on their local context, and removing fine-grained labels that are inconsistent with the results of coarse-grained classification; **Minimum count**, removing the labels appearing less than a certain number of times within the current documents. These heuristic rules actually result in training instances with higher confidence as compared to FIGER.

2.2.3 Features

Fine-grained NER methods convert the input into a vector of features. Most of the features used in fine-grained NER are shared across existing methods. In Table 2.2, we list the features have been used in all of the four systems. We observed that most of the features are adapted from coarse-grained NER.

Feature	Description	Example
Phrase head	Head word of a named entity mention	“Obama”
Non-head	All the non-head words of the mention	“Barack”, “H”
Head cluster	Word cluster ID of the head word	“59”
Character n-gram	character n-grams of the head word	“:ob”, “oba”, ...
Word Shape	The word shape of the mention	“Aa A Aa”
Role	Dependency label of the head word	“nsubj”
Context	context word of the mention	“B:who”, “A:first”
Parent	Lexical parent in a dependency tree	“picker”

Table 2.2: Features shared by all the four related work

2.2.4 Segmentation and Multi-label Classification

Depending on the settings of experiments, different methods might use different strategies to obtain the segmentation of named entity mentions. For example, human-

annotated segmentation of testing data might be directly used as input [5, 6, 7, 27], even though the assumption that the entity boundaries are pre-existing sometimes does not hold in reality. HYENA uses the segmentation provided by the anchored text of Wikipedia page that is not as reliable as manual annotations. In contrast, the setting of FIGER is the most reasonable and practical one. Their testing set is manually labeled. They train a CRF model with BIO tagging schema for segmentation of text into named entity mentions and non-mention during testing.

The classification of fine-grained NER is typically modeled as a multi-label classification problem. According to Silla et al. [30], most of the existing multi-label classifiers fall into four categories: 1) *flat*, referring to those using a single multi-class classifier; 2) *local*, referring to methods using one-versus-rest binary classifiers for each node in the hierarchy; 3) *local per parent node*, associating a multi-class classifier with each internal node; and 4) *global*, using a single multi-class classifier with objective function encoding the similarity between labels. As an example of the first category (*flat*), [16] adapted a simple linear classifier based on single perceptron [31], which maps an input x to a label y given the feature function $F(x, y) = [f_1(x, y), \dots, f_n(x, y)]$,

$$\hat{y} = \operatorname{argmax}_y \sum_{i=1}^n w_i \cdot f_i(x, y)$$

where $w = [w_1, \dots, w_n]$ are the vector of weights learned from training set. During training, w is updated by

$$w \leftarrow w + \alpha \left(\sum_y F(x, y) - \sum_{\hat{y}} F(x, \hat{y}) \right)$$

where $\{y\}$ are the set of true predictions, and $\{\hat{y}\}$ are the set of mispredictions, which is discouraged by decreasing the weights associated with the corresponding features. In order to get multiple labels for one input x , all the tags y with positive score are selected as the predicted labels.

HYENA trained one binary SVM for each node in the hierarchy using all the siblings of that node as negative examples, and all instances of children node as positive examples. In this sense, HYENA falls into the *local* category. Since the

hierarchy of labels might be incomplete, it creates an additional child node *OTEHR* for each inner node in the hierarchy. HYENA classifies the mention of named entity in a top-down manner, assigning each mention with all labels accepted by the classifier, and stopping at the current level when classifier chooses the OTHER label. GFT experiments with both *local* and *flat* classifiers. They also use SVM as the choice for local classification in a way similar to HYENA, while using a softmax classifier as the local approach.

All the above three methods, i.e., FIGER, HYNNA, and GFT, do not consider the similarity between labels, which would be very useful for handling the data sparseness when a big label set is in use. Instead of using traditional classification methods, WsabeeNER used an embedding-based method, called Wsabee [32], which maps both labels and features to same lower dimensional space using liner transformations

$$F(x) \rightarrow AF(x)$$

$$y \rightarrow By$$

where $F(x)$ is the vector of features, and y is the vector of labels. The classification is done by scoring a pair (x, Y) with the products of the transformed vectors,

$$s(x, y) = AF(x)By$$

To learn the parameters, A and B , from training data, they optimize an objective function called weighted approximate pairwise (WARP) [32], which is designed to rank the true prediction y higher than the misprediction $\hat{y}$. The mathematical description of WARP is

$$\sum_y \sum_{\hat{y}} R(rank(x, y)) max(1 - s(x, y) + s(x, \hat{y}), 0)$$

where $rank(x, y)$ is margin infused counting function defined as

$$rank(x, y) = \sum_{\hat{y}} \mathbb{I}(s(x, y) - 1 < s(x, \hat{y}))$$

, and $R(k) = \sum_{i=1}^k \frac{1}{i}$ is a function turning the rank into a weight. By optimizing this objective function, $F(x)$ and y are mapped to lower dimensional space where labels

cooccurring with similar feature vector $F(x)$ are close to each other. Fig.2.5 shows that in 2-dimension space (after applying PCA, and labels are colored based on their top-level labels), labels sharing parent node or superclass tend to be grouped together. As the prediction is made based on the products of label and feature embeddings, this also suggests that the classifier is less discriminative to labels of a lower level, which tends to be close in new space, while more discriminative to labels of a higher level.

Figure 2.5: Label in the projected space [6]

2.2.5 Label Noise Reduction

As mentioned above, data used for training FNET models are weakly tagged. It is unavoidable that the weak tagging process would introduce label noises. More specifically, the noisy instances are mostly over-labeled because of the heuristic rule of weak tagging. For example, in the dataset created by Ling et al. [16], each instance is assigned with all the possible entity types if the surface form matches with multiple canonical forms. Some previous studies [33] have formulated this problem as a partial label learning problem where the true label set is a subset of the observed label. Namely,

$$\hat{Y} \cap Y = \hat{Y}$$

where $\hat{Y}$ is the true label set and Y refers to the set of observed labels. The learning process can then be based on loss functions considering only a subset of the observed

Figure 2.6: Neural Architecture used by Shimaoka et al. [7]

label set. For example, one of the loss function used in the previous work [33] has the form like

$$\max\{0, 1 - \max_{y \in Y} s(x, y) + \max_{y' \in Y'} s(x, y')\}$$

where Y' refers to the set of negative labels. The intuitive behind this loss function is to learn from only the most confident label of a noisily labeled instance.

2.2.6 Neural Architecture for FNET

Following the trend in adopting deep neural networks for NLP tasks, recent work on FNET has adapted the LSTM to modeling the sequence of contextual words as well as of entity mentions. For example, Shimaoka et al. [7](as illustrated in Fig. 2.6) proposes to encode the contextual word sequences using bi-directional LSTM. Their framework also incorporates an attention component to model the salience of each contextual word in resolving the types of the center entity mention. A similar neural network is also adopted by work of Abhishek et al. [27] which replaces the bag-of-words encoder of entity mentions with another uni-directional LSTM encoder. Their work

shows that the neural architecture is able to perform comparably than embedding methods using the manually designed features.

One question is how to encode the correlation of labels. In other words, the labels are not created free of semantics. Each label in the hierarchy, in fact, corresponds to a concept in real world. It is, therefore, useful to model the correlation or dependency of labels. In HYENA [26], although the classification results are passed from parent to its children, each classifier, however, is trained independently. In comparison, the embedding based methods [6, 34] captured the similarity between labels by considering the concurrence of labels and features. But, their methods fail to take into consideration the existing hierarchy that has encoded useful information about the dependency between labels. To address these issues of existing approach, one solution is to consider the Bayesian hierarchical classifier [35], which models the dependency between labels using a hierarchical prior. Their experiments have shown improved performance over *flat* and *local* classifiers on a range of tasks and datasets and is scalable to work with large-scale training data.

2.3 Targeted Sentiment Analysis

In this section, we survey multiple research areas related to the proposed framework, namely: Aspect-based Sentiment Analysis (ABSA), targeted sentiment analysis, targeted ABSA, and finally works on incorporating external knowledge into deep neural models.

2.3.1 Aspect-Based Sentiment Analysis

ABSA is the task of classifying sentiment polarity with respect to a set of aspects. The biggest challenge faced by ABSA is how to effectively represent the aspect-specific sentiment information of the whole sentence. Early works on ABSA have mainly relied on feature-engineering to characterize sentences [36, 37]. Motivated by the success of deep learning in representation learning, many recent works [38, 39, 40, 41] utilize deep neural networks to generate sentence embeddings (dense vector representation of sentences) which are then fed to a classifier as a low-dimensional feature vector.

Moreover, the representation can be enhanced by using the attention mechanism [41], which is typically a multi-layer neural network taking as input the word sequence and aspects. For each word of the sentence, the attention vector quantifies its sentiment salience as well as the relevance to the given aspect. The resulting sentiment representation benefits from the attention mechanism for it overcomes the shortcoming of recurrent neural networks (RNNs), which suffer from information loss when only one single output (e.g., the output at the end of the sequence) is used by the classifier.

2.3.2 Targeted Sentiment Analysis

Targeted sentiment analysis aims to analyze sentiment with respect to targeted entities in the sentence. It is thus critical for targeted sentiment analysis methods, e.g., the target-dependent LSTM (TDLSTM) and target connection LSTM (TCLSTM) [42], to model the interaction between sentiment targets and the whole sentence. In order to obtain the target-dependent sentence representation, TDLSTM directly uses the hidden outputs of a bidirectional-LSTM sentence encoders panning the target mentions, while TCLSTM extends TDLSTM by concatenating each input word vector with a target vector. Similar to ABSA, attention models are also applicable to targeted sentiment analysis. Rather than using a single level of attention, deep memory networks [43] and recurrent attention models [44] have achieved superior performance by learning a deep attention over the single-level attention, as multiple passes (or hops) over the input sequence could refine the attended words again and again to find the most important words. All existing approaches have either ignored the problem of multiple target instances (or words) or simply used an averaging vector over target expressions [43, 45]. Unlike such approaches, our method weights each target word with an attention weight so that a given target is represented by its most informative components.

2.3.3 Targeted Aspect-Based Sentiment Analysis

Two baseline systems [46] are proposed together with SentiHood: a feature-based logistic regression model and an LSTM-based model. The feature-based logistic regression model uses feature templates including n-grams tokens and POS tags extracted

from the context of instances. The LSTM baseline can be seen as an adaptation of TDLSTM that simply uses the hidden outputs at the position of target instances assuming that all target instances are equally important.

2.3.4 Incorporating External Knowledge

External knowledge bases have been typically used as a source of features [47, 48, 49]. Most recently, neural sequential models [50, 51] leverage the lower-dimensional continuous representation of knowledge concepts as additional inputs. However, these approaches have treated the computation of neural sequential models as a black-box without tight integration of knowledge and computational structure. The proposed model, termed Sentic LSTM, is inspired by [52], which adds a knowledge recall gate to the cell state of LSTM. However, our method differs from [52] in the way of using external knowledge to generate the hidden outputs and controlling the information flow.

2.4 ASR Hypotheses Reranking

Another research line to be reviewed is ASR hypotheses reranking. It is a technique to improve the quality of texts automatically transcribed by a speech recognition system. ASR hypotheses reranking is critical for understanding opinions expressed via not only texts but also audios. ASR hypotheses reranking might sometimes be referred to as discriminative language modeling in contrast to generative models such as the n-gram model used by [53].

Typically, ASR systems consist of two modules: 1) acoustic model and 2) language model. The best hypothesis of an ASR system is induced with

$$\arg \max_W P(W|A) = \arg \max_W \log(P_a(A|W)) + \alpha \log(P_l(W))$$

where A is the acoustic representation of an utterance, and W is the hypothesized word sequence, respectively. $P_a(s|a)$ is the acoustic model, and $P_l(s)$ the language model with α controlling the relative importance of the language model.

Speech transcripts contain disfluencies such as repetitions, false starts and fillers, which can interrupt the consecutive n-gram patterns. The conventional n-gram language model, therefore, is unsatisfactory for modelling spoken language, especially for conversational speech which is spontaneously incurred. In order to address the problems caused by the informal nature of speech, some works have modeled the long distance context using syntactic parse trees. Approaches using syntactic information for language modeling either incorporate the parsing process into language modeling or maintain a separate process by including features generated from syntactic parsing. As an example of the first type of approaches, [54] used a shift-reduce parser to maintain a set of possible parses of a given sentence. The probability of a word given its previous context is then defined as the sum of probabilities of candidate parses. As for the second type, [55, 56] used single perceptron model with features extracted from syntactic parse trees and dependency tree. Their experiments show that, when combined with the ASR baseline model, discriminative reranking could further reduce the word error rate, even though by a small margin. A feature function $F(W)$ is defined for a given word sequence W , and used for rescoreing the hypothesis of ASR with a modified log probability

$$\log(P(W, A)) = \log(P_a(A|W)) + \alpha \log(P_l(W)) + \sum_i^N f_i(W) \cdot w_i$$

where f_i is the i^{th} feature in $F(W)$, and w_i is its weight. The syntactic features used by [55] are derived from the parse trees produced by a syntactic parser, including head features such as the syntactic head of a detected noun phrase, and head-to-head dependencies. Instead of using syntactic parsing based features, [56] extracted features from dependency parses, which are mainly composed of a set of dependency tuples. For example, for a sentence

I WENT TO THE AN- THE DOCTER

Stanford parser produces the set of tuples as shown in Table.2.3. Each tuple represents a dependency link in the parsing result, including two nodes and their syntactic relation such as noun subject(*nsubj*) between a verb and subject. All the tuples generated for the given sentence are added to the feature set.

nsubj(went-2, I-1)
root(ROOT-0, went-2)
case(docter-7, to-3)
det:predet(docter-7, the-4)
det(docter-7, an-5)
amod(docter-7, -6)
nmod(went-2, docter-7)

Table 2.3: The dependency tuples for sentence *I WENT TO THE AN- THE DOCTER*

One interesting finding made by Lambert et al. [56] is that a large number of features are actually redundant for reranking. After performing feature selection, they found that the reranker can still achieve comparable results using only 0.5% of the original feature set. In addition, their results also show that approximate 80% features are extracted less than 5 times, which suggests a problem of data sparseness. It is worth a further investigation on whether reranking can benefit from reducing the data sparseness.

2.5 Embedding Models

As discussed in previous sections, many NLP models transform the input into feature vectors of very high dimensionality. Computation on such high-dimensional space is very expensive and sometimes even intractable. To address this issue, dimensionality reduction tricks have been adopted to convert the original feature vector into lower dimensional representations which are also known as embeddings. As compared with one-hot representation whose dimensionality is typically same as the size of feature dictionary, embedding approaches could significantly reduce the dimensionality without suffering a significant loss of information. Depending on the tasks, embeddings can be generated for not only words but also concepts or phrases.

One of the fundamental intuition of learning embeddings for words or concepts is distributional semantics which assumes the distribution of contexts characterizes the semantic meaning of a central word. In other words, two words (or concepts) are close in semantics if having the similar distribution of contexts. Recently, embedding based on distributional semantics have been used as effective features for many NLP tasks,

e.g., named entity recognition [4], dependency parsing [57], semantic relation classification [58], antonym detection [59], sentiment analysis [60], and spoken language understanding [61].

According to the classification by Baroni et al. [62] and Levy et al. [63], existing word embedding methods based on distributional semantics can be grouped into two categories: count-based and prediction-based. The count-based methods can be dated back to methods such as Latent Semantic Analysis (or Latent Semantic Indexing) [64] that are followed by lines of research on decomposing matrices of co-occurrence counts.

2.5.1 Count-based Models

It is a long history that the co-occurrence counts and reweighting tricks have been used for obtaining embeddings. Here, we are only able to review a small subset of them that have recently attracted enough attention.

First of all, the task of learning embeddings from co-occurrence counts can be formulated as a matrix factorization problem. Let M_{ij} be an entry of M . Given the training dataset $D = \{w_1, w_2, \dots, w_n\}$, the value of M_{ij} can be either original or weighted counts of co-occurrence of word w_i and a context c_j . For example, a popular weighting schema is Pointwise Mutual Information (PMI) [65].

We then formulate the process of deriving word embeddings as a mapping from co-occurrence matrix $M \in R^{n \times m}$ to embedding matrix $W^{n \times d}$, with each row of W corresponding to a d -dimensional vector representation of the word w_i .

Next, we will review two of the count-based methods that are recently studied for learning embeddings of words or concepts.

2.5.1.1 Single Value Decomposition

A natural choice to decompose the co-occurrence matrix to lower-dimension vectors is to use traditional dimension reduction methods such as Single Value Decomposition (SVD).

SVD generates an optimal low-rank approximation of the original counts matrix M by factorizing it as

$$M = U\Sigma V^\top$$

where $U \in R^{n \times n}$ and $V \in R^{m \times m}$ are unitary matrices, and $\Sigma \in R^{n \times M}$ being a diagonal matrix. Each diagonal entry of Σ is a singular value of M . To get a d -dimensional vector representation, we just take the dot product of U^* and Σ^* , i.e.,

$$W = U^* \Sigma^*$$

where Σ^* is the same matrix as Σ except that it contains only the top d values, and U^* being the corresponding truncated matrix of U .

2.5.1.2 Random Projection

One problem with SVD is that it is hard to scale up, especially when the original matrix is sparse and too high-dimensional [66].

Cambria et al. [10] propose to a data oblivious method, called random projection, to map the original high-dimensional space to a lower dimensional space. The adopted random projection of concept into the new space is done via a Gaussian matrix. According to Johnson and Lindenstrauss (JL) Lemma [67], points distributed in a m -dimensional space can be safely projected to a d -dimensional space, where d is logarithmic to m , and pairwise distances being preserved with high probability. It is, therefore, possible to use a random matrix to transform each row of original co-occurrence matrix to a lower-dimensional row vector.

Cambria et al. [10] propose to use a fast random projection technique, called Subsampled Randomized Hadamard Transform (SRHT) [68] a random matrix Φ , which can be factorized by

$$\Phi = \sqrt{\frac{m}{d}} R H D$$

where R is a $d \times m$ with rows being sampled from an uniform distribution, and H is a $m \times m$ normalized Walsh-Hadamard matrix, and D is $m \times m$ random matrix.

After computing Φ , the new embedding matrix W can be obtained by linearly transforming M using the transpose of Φ , i.e.,

$$W = M \Phi^\top$$

As shown by Cambria et al. [10], the random projection method could project the very high-dimensional concept matrix to a low-dimensional space with high efficiency.

2.5.2 Prediction-based Models

Apart from working with co-occurrence counts, the past few years have seen a trend of learning embedding by maximizing the probability of contextual terms. The family of embedding learning methods [21, 69, 70] are, therefore, classified as prediction-based methods. Next, we will review one of the most widely-used prediction-based methods, called Skip-gram model, as well as a variant of it, called CBOW model.

2.5.2.1 Skip-gram Models

Skip-gram model [70] is a neural network model that can be efficiently learned from a large-scale corpus with billions of words. Figure.2.7 shows the training process of Skip-gram model. The model first takes a window of the local context of the given word to generate the word pairs. Then, it updates the parameters of the neural network using each word pair. Afterwards, it moves to next position and loop over Step1-4 again until the end of the corpus.

Its objective function is the sum of log probabilities $p(w_{i+j}|w_i)$ over the whole corpus,

$$\sum_{i=1}^N \sum_{j=-k}^k \log p(w_{i+j}|w_i)$$

where w_{i+j} indicates a neighbor word of w_i , k the size of the context window and N is the length of the dataset. And, the basic Skip-gram formulation defines $p(w_{i+j}|w_i)$ using the softmax function as follows:

$$p(w_{i+j}|w_i) = \frac{\exp(v'_{w_{i+j}} v_{w_i})}{\sum_{w \in W} \exp(v'_w v_{w_i})}$$

where W is the vocabulary of words, $v'_{w_{i+j}}$ and v_{w_i} are the context and word embeddings, respectively.

2.5.2.2 Continuous Bag-of-Words (CBOW) Model

The training objective CBOW model [70] is to maximize the sum of probability $p(w_i|w_{i-k}, \dots, w_{i+k})$, which is similar to Skip-gram model.

$$\sum_{i=1}^N p(w_i|w_{i-k}, \dots, w_{i+k})$$

Figure 2.7: Training process of Skip-gram model.

The probability of w_i given contexts words $\{w_i + j\}$ is defined on top of v'_{w_i} , which is the context embedding of w_i and h which is defined as the average vector of words in the context

$$h = \frac{1}{2k} \sum_{j=-k}^k v_{w_i+j}$$

Same as Skip-gram model, the probability $p(w_i | w_{i-k}, \dots, w_{i+k})$ is normalized using softmax function

$$p(w_i | w_{i-k}, \dots, w_{i+k}) = \frac{\exp(v'_{w_i} h)}{\sum_{w' \in W} \exp(v'_{w'} h)}$$

2.6 Summary

In this chapter, we have reviewed recent progress in several tasks that are core to understanding human opinions. We also survey on techniques aiming to learn embedding for words and concepts. All existing solutions have their own technical limits as well as unsolved issues. However, it lacks investigations on leveraging the concepts for specific tasks. In addition, the fusion of information from concept-level and other levels (e.g., word-level) has been missed in existing studies.

Chapter 3

Learning Word Embedding with Multi-task Entity-based Features

3.1 Introduction

Named entity recognition (NER) has been studied and applied successfully to formal [71] and informal texts [72] for many years. Nevertheless the adaptation of NER methods to conversational speech remains challenging due to, for example, case insensitivity, lack of punctuations, un-grammatical structure, repetition, and presence of disfluencies inherent to conversations. In addition, there is not much spoken data annotated with named entities to cover the huge variety of named entity instances likely occurring in speech, and simply increasing the amount of manual annotation is not realistic for reasons of cost, evolution of new spoken terms and diversity.

Several works on NER from spoken contents have already explored the use of external resources like online gazetteers [73] and Wikipedia [74] to overcome the lack of annotations. Gazetteers, for instance, have successfully boosted NER performance for given entities (e.g., Location), but do not convey the information related to the context words surrounding the entity names that are also important for NER. Other lexical resources such as WordNet provide semantic relations like synonymy among common English words, but remain limited for names. A second category of approaches [75, 76, 77, 78, 79] tackling the sparseness of NER training data use unlabeled data to learn low dimensional vector representation of words, called word embeddings. Used either as continuous [75, 76, 79] or discrete features [77, 78], word embeddings have been shown effective in improving the generalization of NER.

For the last two years, an increasing number of studies have suggested that injecting application-specific information into the neural networks used to train word embeddings can further improve the performance of down-stream applications, e.g., dependency parsing [57], semantic relation classification [58], antonym detection [59], spoken language understanding [61]. The task-specific information used by these methods are injected into the training process of word embeddings mainly by expanding or replacing the input or output of the neural network. Passos *et al.* [80], to our knowledge, is the only work trying to learn word embeddings for NER by leveraging lexicons related to named entities. Our proposal is inspired by their approach and can be seen as its generalization in the sense that our word embedding is trained with multiple relevant tasks while their model is trained on only one additional task that predicts whether the current word belonging to a predefined set of semantic classes.

The key to learning word embeddings for the particular problem of NER is to encode entity-related information into the word embedding space. As the first step to investigate a tight fusion of concept-level¹ and word-level information, we propose to augment the learning process of word embedding with entity-related features. More specifically, we introduce an NER framework including a modified word embedding method based on Skip-gram models [70], in which we generalize the training objective to integrate features explicitly designed for NER task and grouped by types. From a different point of view, we extend the training of Skip-gram model to a multi-tasks setting where each feature type corresponds a particular task. The multi-task learning of word embeddings allow the embedding to form a thorough representation of the word by incorporating information from different information sources. We prove the effectiveness of our method in the context of NER. Namely, we find that injecting such NER-specific information as part-of-speech tags, taxonomic relations, and self-training features [74] yields improvements over baseline (i.e., without using word embeddings), when evaluated against conversational speech transcripts (i.e., Switchboard Corpus). Furthermore, the results show that our proposed word embeddings also outperform the task-independent word embeddings.

¹Named entity is a subclass of concepts.

Figure 3.1: NER Framework with NER-specific word embeddings

3.2 Methods

The proposed NER framework illustrated in Figure 3.1 is composed of two parts: 1) a task-specific word embedding learning integrating features specific to NER (Part I in Figure 3.1.), using not only unlabeled data but also resources like knowledge base and baseline NER tagger; and 2) a supervised model of linear-chain CRF trained with the NER features extracted from the labeled corpus and the resultant word embeddings (Part II in Figure 3.1). We denote the proposed word embedding method as Skip_{NER} .

3.2.1 Baseline NER method

In this section, we describe the baseline NER method, a conventional linear-chain CRF with BIO encodings. The features extracted from every word w_i of the training data and used to train the CRF are summarized in Table 3.1. The feature set consists of n-grams, part-of-speech (POS) tags, affixes, and the BIO tag of the i^{th} word (designated as y_i).

Context Features (CF)	
Unigram	$w_{i+k}, \quad -2 \leq k \leq 2$
Bigram	$w_{i+k} \wedge w_{i+k+1}, \quad -2 \leq k \leq 1$
POS	$t_{i+k}, \quad -2 \leq k \leq 2$
POS bigram	$t_{i+k} \wedge t_{i+k+1}, \quad -2 \leq k \leq 1$
Prefix	$Pre(w_{i+k}, l), \quad -2 \leq k \leq 2, 0 \leq l \leq 4$
Suffix	$Suf(w_{i+k}, l), \quad -2 \leq k \leq 2, 0 \leq l \leq 4$
Tag Feature	
Tag&Context	$y_i \wedge c, \quad c \in CF$
Tag Bigram	$y_{i-1}y_i$

Table 3.1: Templates of features used in the CRF baseline

3.2.2 Skip-gram model

Skip-gram model is a neural language model that can be efficiently trained with corpus of billions of words. It assumes that the user would only be concerned about the quality of word embeddings but not the capability of language modeling. The objective of Skip-gram model is to find the representation of word that is useful for predicting the surrounding words instead of just next words. Each word is linked with a d -dimensional vector v_w , and each context c is assigned a vector v'_c which is also of d dimensions, i.e., the objective function is to maximize the log probability

$$\sum_{t=1}^T \sum_{j=-k}^k \log p(c_j = w_{t+j} | w_t)$$

where k is size of context window which are a certain number of words surrounding the central word. The probability $p(c|w)$ is given as

$$p(c|w) = \frac{\exp(v'_c{}^\top v_w)}{\sum_{c'} \exp(v'_c{}^\top v_w)}$$

However, as computing the gradient of $\log p(c|w)$ requires iteration over all pairs of (w, c) , optimizing this objective function is intractable when the size of context vocabulary is large. To approach this problem, [70] uses negative sampling as an alternative training objective. The objective function can be rewritten as

$$\sum_{t=1}^T \sum_{j=-k}^k (\log \sigma(v_c^T v_{w_t}) + \sum_k^N \mathbb{E}_{c' \sim P_w} (\log(\sigma(-v_{c'}^T v_{w_t}))))$$

where $\sigma(x) = 1/(1 + \exp(-x))$. For each (w, c) pair appearing in the corpus, N pairs of (w, c_k) are randomly sampled from a unigram distribution P_w as negative training instances.

3.2.3 Feature extraction for word embeddings

First of all, we extract four types of features (i.e., neighboring words, part of speech (POS), taxonomic, self-trained) from the unlabeled data to train word embeddings. In principal, they are mostly features used for training the baseline NER model, and one of the intuition of our model is to learn a condensed version of these features from unlabelled data to achieve higher capability of generalization.

Neighboring words and POS tags: They are acknowledged to be efficient for NER [73, 74]. In fact, we use these features not only for training our word embeddings, but also for training the baseline model of NER (Section 3.2.1). Their formal definitions can be found in the first 4 lines of Table 3.1. In terms of feature type, we define two separate types for words and POS tags respectively.

Taxonomic features: The generalization of similar entities within a common category (or concept) in a taxonomy may be useful for NER to capture contexts shared by similar entities. We select 280 concepts with at least 100 instances from ConceptNet, which is a large, automatically created semantic graph including taxonomic relations. For each concept g , we add binary taxonomic features, which are defined as $\text{TX}_g(w_{i+k}), 2 \leq k \leq 2$, where i is the index of current word, and $\text{TX}_g(w)$ indicates if the word w belongs to the concept g or not.

Self-trained features: As in [81], these features are generated by automatically labeling the training data using a baseline NER tagger (Section 3.2.1). We use the resultant named entities hypotheses as additional features for word embeddings. The features are defined as T_{i+k} , with T_j the NER label of the j^{th} word.

3.2.4 Learning word embeddings from entity-based features

Our proposal is based-on the Skip-gram model [70], a neural language model that can be efficiently trained on a large corpus of billions of words. Its objective function is

the sum of log probabilities $p(w_{i+j}|w_i)$ over the whole corpus,

$$\sum_{i=1}^N \sum_{j=-k}^k \log p(w_{i+j}|w_i)$$

where w_{i+j} indicates a neighbor word of w_i , k the size of the context window and N is the length of the dataset. And, the basic Skip-gram formulation defines $p(w_{i+j}|w_i)$ using the softmax function as follows:

$$p(w_{i+j}|w_i) = \frac{\exp(v'_{w_{i+j}}{}^\top v_{w_i})}{\sum_{w \in W} \exp(v'_w{}^\top v_{w_i})}$$

where W is the vocabulary of words, $v'_{w_{i+j}}$ and v_{w_i} are the embeddings of context word and current word, respectively. We modify the Skip-gram model to predict the set of features $F(w_i)$ extracted for the given word w_i at the i^{th} position of a corpus. The objective function can be rewritten:

$$\sum_{i=1}^N \sum_{f \in F(w_i)} \log p(f|w_i)$$

To estimate $p(w_{i+j}|w_i)$, the basic Skip-gram model assumes a single distribution over all the words. However, for $p(f|w_i)$, since the features we use for training application-specific word embeddings are heterogeneous, they may have different distributions. We thus split the whole set of features (S) into subsets (S_X), where X indicates one of the four feature types aforementioned and the relative position to the center word. For example, the subset $S_{pos:-1}$ includes all the features that tell the POS tags of the previous word. We define $C_S(f)$ as the function that returns the subset of S that contains the feature f . We define the probability of extracting f for a given word w from the training data as follows:

$$p(f|w) = \frac{\exp(v_f^\top v_w)}{\sum_{f' \in C_S(f)} \exp(v_{f'}^\top v_w)}$$

where v_w and v_f are the vectors associated with center word w and feature f respectively, and both are the parameters to be learned. Note that, with the new definition, the objective function becomes a linear combination of the objective functions of multiple classifiers with equal weights. We optimize this objective function

using stochastic gradient descent and negative sampling method previously proposed for Skip-gram [70], which rewrites the objective function as:

$$\sum_{i=1}^N \sum_{f \in F(w_i)} \left(\log \sigma(v_f^T v_{w_i}) + \sum_{f' \in Z(f)} \log (\sigma(-v_{f'}^T v_{w_i})) \right)$$

In this expression $\sigma(x) = 1/(1 + \exp(-x))$, and the set of features $Z(f)$ is created by randomly selecting n negative samples from a unigram distribution over features in $C_S(f)$

3.2.5 Multi-task interpretation of the proposed method

Since each type of features are heterogeneous to one another, we could also view the prediction of each feature type as performing a separate task. Therefore, the iteration over the set of feature type (i.e., F) can be seen as performing multi-task learning over a set of tasks. For example, predicting neighboring words and POS tags corresponds to the tasks of language modeling and POS tagging respectively.

3.2.6 Using word embeddings as NER features

We convert word embeddings to additional features of the CRF model as in [77, 78] by binarizing vector elements of word embeddings, and clustering of words based on similarity of word embeddings. The vectors are binarized using the following rules,

$$D_{mn} = \begin{cases} 1 & \text{if } W_{mn} \geq \overline{W^+}_m. \\ -1 & \text{if } W_{mn} \leq \overline{W^-}_m. \\ 0 & \text{else} \end{cases}$$

where W is the original word embedding matrix, and D the binarized matrix, $\overline{W^+}_m$ the mean of all positive values of the m^{th} dimension, and $\overline{W^-}_m$ the mean of all negative values. Only non-zero values are added to the feature set. $\vec{V}_{D_w}$ denotes the column of D corresponding to the word w . Then we cluster all words based on Euclidean distance using K-Means [78, 77]. We use different numbers of clusters, K , to let the clustering reflect different levels of granularity. $C_K(w)$ is the cluster of the word w , where K indicates the number of clusters in the K-Means ($500 \leq K \leq 3000$). The features learned from word embeddings are summarized in Table 3.2.

Binarized vector	$\vec{V}_{Dw_{i+k}}$	$-2 \leq k \leq 2$
Cluster (Unigram)	$C^K w_{i+k}$	$-2 \leq k \leq 2$
Cluster (Bigram)	$C^K w_{i+k} \wedge C^K w_{i+k+1}$	$-2 \leq k \leq 1$
Cluster (Disjunct)	$C^K w_{i-1} \wedge C^K w_{i+1}$	

Table 3.2: Features learned from word embeddings for NER

3.3 Experiments

3.3.1 Corpora

We have used the **ukWaC corpus** [82] as unlabeled text dataset to train word embedding models. This is one of the largest English text corpus (about 1.8 billion tokens), crawled from the Web on the *co.uk* domains, and has been tagged with POS².

The labeled open-conversation dataset is the Switchboard corpus [83]. It is a large collection of 2-speaker conversational telephonic speech recordings. We use the NER annotations of the corpus including four classes – location (LOC), person (PER), organization (ORG) and miscellaneous (MISC), and the train/test partitions provided by Surdeanu et al. [73]. Detailed numbers are presented in Table 3.3.

	LOC	PER	ORG	MISC
Train	14,397	4,257	5,311	8,484
Test	631	342	296	687

Table 3.3: Named entity distribution in Switchboard

We compare the proposed Skip-gram model (Skip_{NER}) with non-task specific word embeddings models previously used for NER. We retrained the word embeddings for Skip-gram, CBOW [70], and Glove [79] with the ukWaC corpus, using two-word context windows for each side of the current word. We used the pre-trained word embeddings models³ of HUANG [69] and SENNA [75], because their training is slow. Table 3.4 shows some statistics of the data for the models. Following [76, 75, 78, 79], the number of vector dimensions is set to 50 for all experiments.

²<http://www.cis.uni-muenchen.de/~schmid/tools/TreeTagger/>

³<http://ai.stanford.edu/~ehhuang/>, <http://ronan.collobert.com/senna/>

Word embeddings	Tokens	Vocab.	Dataset
Skip _{NER} , Skip-gram	1.8B	167K	ukWaC
Glove, CBOW	1.8B	167K	ukWaC
HUANG	1.8B	100K	Wikipedia
SENNA	1.8B	130K	Wikipedia

Table 3.4: Details on baseline Word Embeddings preparation

3.3.2 Evaluations

We report in Table 3.5 the performance of the proposed NER method (Skip_{NER}) and the baseline method with the other word embeddings when tested against the manual transcripts of open-domain conversations. Our method outperforms all the other methods and achieves 2% absolute improvement of F-score in comparison to the original Skip-gram model.

System	F-score	System	F-score
Baseline	66.51		
Baseline + Skip	67.89	Baseline + Huang	67.88
Baseline + CBOW	68.45	Baseline + Senna	68.29
Baseline + Glove	67.44	Baseline + Skip _{NER}	70.19

Table 3.5: Performance of NER methods in terms of F-score

In the second experiment, we investigate the impact of the amount of labeled training data used to train the NER model. As illustrated in Figure 3.2, the Skip_{NER}-based NER system always outperforms the other systems, independent of the amount of training data. Also, the NER systems with the generic (or non task-specific) word embeddings also perform better than the baseline. It suggests that the generalization brought by word embeddings consistently helps NER models in dealing with data sparsity. As more training data are used (80%-100% of training data), the gap of performance between all the baseline systems reduces. However, the gap between the baseline and Skip_{NER} further increases even using more training data, which may mean that our method robustly handles data sparsity and extracts more discriminative information when more data are available.

We also study the impact of each subset of features used to train the Skip_{NER} model. As reported in Table 3.6, we find that using the whole set of features outper-

Figure 3.2: Performance with varying size of NER training data

forms each of the feature subsets used alone. This result may indicate that the feature subsets are complementary to each other, thus supporting the proposed method of integrating NER-specific features for training word embeddings. Also, note that the Words only subset corresponds to the skip-gram model. Location, and Person names are particularly well recognized by using the POS subset, while self-trained features and taxonomic relations seem more effective to discriminate the class ORG.

Skip _{NER} Feature	PER	LOC	ORG	MISC	All
Words	80.6	79.3	51.7	57.0	68.5
POS	82.3	80.6	51.1	57.8	69.4
self-trained	80.8	79.8	53.6	57.2	68.9
taxonomic	81.0	80.0	54.0	57.4	69.3
All	81.7	80.9	55.7	57.9	70.2

Table 3.6: Impact of feature subsets on Skip_{NER} [12]

Table 3.7 lists the five words most similar to each of three example words (i.e., cowboys, texas, batman), which are computed using the Skip-gram and Skip_{NER}. For instance of the keyword *texas*, both Skip_{NER} and Skip-gram models return locations in the United States (US), but Skip_{NER} seems to produce more ‘relevant’ results

Word	texas	cowboys	batman
Skip-gram	lousiana	texans	superman
	spokane	cowboy	superhero
	kansas	bandits	remake
	sacramento	cowgirls	spiderman
	biloxi	impersonators	catwoman
Skip _{NER}	kentucky	cheerleaders	superman
	kansas	texans	shrek
	lousiana	yankees	starsky
	florida	redskins	scooby-doo
	minnesota	broncos	spiderman

Table 3.7: Top similar words returned by Skip-gram & Skip_{NER}

than Skip-gram since *texas* and all its five most similar words (e.g., *kentucky*) are the names of states in US, while the results of Skip-gram include city names (e.g., *spokane*, *sacramento*), thus of different granularity. For ambiguous words, Skip_{NER} seems to focus on the semantics of the words related to the NER task. For instance of the keyword *cowboys*, it may be the plural form of the noun ‘cowboy’ or may indicate the American football team “Dallas Cowboys”. While its similar words from Skip-gram reflect the ambiguity, all the top results of Skip_{NER} are related to the latter meaning of the keyword. We observe a similar difference in the results for the keyword *batman*. These observations suggest Skip_{NER} is able to capture more discriminative information for NER compared to generic word embeddings.

3.3.3 Summary

In this work, we propose a novel method to train task-specific word embeddings for NER by incorporating NER related multi-task learning. Through several experiments, we have shown that, on the manual transcripts of open-domain conversations, how our proposed feature-enriched word embeddings can out- perform baseline NER method and systems using task inde- pendent word embeddings.

Chapter 4

Label Embedding for Fine-grained Named Entity Typing

4.1 Introduction

Named entity typing (NET) is the task of inferring types of named entity mentions in text. NET is a useful pre-processing step for many NLP tasks, e.g., auto-categorization and sentiment analysis. Named entity linking, for instance, can use NET to refine entity candidates of a given mention [16]. Besides, NET is capable of supporting applications based on a deeper understanding of natural language, e.g., knowledge completion [84] and question answering [85, 86]. Standard NET approaches or sometime known as named entity recognition [11, 13, 87] are concerned with coarse-grained types (e.g., person, location, organization) that are flat in structure. In comparison, fine-grained named entity typing (FNET) [16], which has been studied as an extension of standard NET task, uses a tree-structured taxonomy including not only coarse-grained types but also fine-grained types of named entities. For instance, given “[*Intel*] said that over the past decade”, standard NET only classifies *Intel* as *organization*, whereas FNET further classifies it as *organization/corporation*.

FNET is faced with two major challenges: growing type set and label noises. Since the type hierarchy of entities is typically built from knowledge bases such as DBpedia, which is regularly updated with new types (especially fine-grained types) and entities, it is natural to assume that the type hierarchy is growing rather than fixed over time. However, current FNET systems are impeded from handling a growing type set for that information learned from training set cannot be transferred to unseen types.

Another problem with FNET is that the weakly supervised tagging process used for automatically generating labeled data inevitably introduces label noises. Current solutions rely on heuristic rules [88] or embedding method [33] to remove noises prior to training the multi-label classifier. In order to address these two problems at the same time, we propose a simple yet effective method for learning prototype-driven label embeddings that works for both seen and unseen types and is robust to the label noises. The label embeddings, which might also be deemed as concept embeddings, are inferred from a pre-trained word embedding space via a set of prototypes. From the perspective of fusing information from multiple levels, we made an attempt to fuse word-level information (i.e., word embedding of prototypes) to represent the canonical form of concepts (i.e., concept types). Another contribution of this work is that we combine prototypical and hierarchical information for learning label embeddings.

The remainder of this chapter is organized as follows: Section 4.2 proposes a survey of prior works related to FNET; Section 4.3 introduces the embedding-based FNET method and its zero-shot extension; Section 4.4 describes our label embedding method; Section 4.5 illustrates experiments and analysis for both few-shot and zero-shot settings; finally, Section 4.6 concludes the paper and discusses future work.

4.2 Background

There is little related work specifically on zero-shot FNET but several research lines are considered related to this work: fine-grained named entity recognition, prototype-driven learning, and multi-label classification models based on embeddings. As FNET works with a much larger type set as compared with standard NET, it becomes difficult to have sufficient training data for every entity type if relying on only manual annotation. Fortunately, semi-structural data such as Wikipedia pages [16] turns out to be a good source for automatically generating training data. This auto-annotating practice has been followed by later work on FNET [26, 6, 33]. However, since the automated tagging process is not perfect all the time, a number of label noises are propagated to supervised training and have consequently a negative impact on the classification performance.

The starting point of this work is the embedding method, WSABIE [32], which have been adapted by [6] to FNET. WSABIE maps input features and labels to a joint space, where information is shared among correlated labels. However, the joint embedding method stills suffers from label noises. In addition, WSABIE fails to handle unseen types for it uses labeled training set as the only source learning label embeddings.

In addition to WSABIE, this work is also inspired by recent progress on image annotation based on embedding methods. For example, DeVISE [89] is proposed for annotating image with words or phrases, where embedding of single words, e.g., fruit, that are pre-learned from textual data using Skip-gram word embeddings (see Section 2.5.2.1 for more details) [70] are directly used for representing labels. In addition to label itself, existing work has also attempted to derive label embeddings from side information such as attributes [90], manually-written descriptions [91], taxonomy of types [92, 90, 93], and so on and so forth.

Another related line of research is prototype-driven learning. In [94], a sequence labeling model uses prototypes as features and is applied to NLP tasks such as part-of-speech (POS) tagging. Prototype-based features [78] are then adapted for coarse-grained named entity recognition task. Even though we select prototypes in the same way as [78], we use prototypes in a very different manner: we consider prototypes as the basis for representing labels, whereas prototypes are mainly used as additional features in prior works [94, 78]. In other words, prototypes are previously used on the input side, while we use them on the label side.

4.3 Embedding Methods for FNET

In this section, we introduce the embedding method for FNET proposed by [6] and its extension to zero-shot entity typing.

4.3.1 Joint Embedding Model

Each entity mention m is represented as a feature vector $x \in \mathbb{R}^V$; and each label $y \in Y$ is a one-hot vector, where Y is the set of true labels associated with x . $\bar{Y}$ denotes

the set of false labels of the given entity mention. The bi-linear scoring function for a given pair of x and y is defined as follows:

$$f(x, y, W) = x'Wy$$

where $W \in \mathbb{R}^{M \times N}$ matrix with M the dimension of feature vector and N the number of types.

Instead of using a single compatibility matrix, WSABIE [32, 6] considers an alternate low-rank decomposition of W , i.e., $W = A^\top B$, in order to reduce the number of parameters. WSABIE rewrites the scoring function as

$$f(x, y, A, B) = \phi(x, A) \cdot \theta(y, B) = x'A^\top By$$

, which maps feature vector x and label vector y to a joint space. Note that it actually defines feature embeddings and label embeddings as

$$\phi(x, A) : x \rightarrow Ax,$$

$$\theta(y, B) : y \rightarrow By,$$

where $A \in \mathbb{R}^{D \times M}$ and $B \in \mathbb{R}^{D \times N}$ are matrices corresponding to lookup tables of feature embeddings and label embeddings, respectively. The embedding matrices A and B are the only parameters to be learned from supervised training process. In [32], the learning is formulated as a learning-to-rank problem using weighted approximate-rank pairwise (WARP) loss,

$$\sum_{y \in Y} \sum_{y' \in \bar{Y}} L(\text{rank}(x, y)) \max(1 - f(x, y, A, B) + f(x, y', A, B), 0),$$

where the ranking function $\text{rank}(x, y) = \sum_{y' \in \bar{Y}} \mathbb{I}(1 + f(x, y', A, B) > f(x, y, A, B))$, and $L(k) = \sum_{i=1}^k \frac{1}{i}$ which maps the ranking to a floating-point weight.

4.3.2 Zero-shot FNET Extension

A zero-shot extension of above WSABIE method can be done by introducing pre-trained label embeddings into the framework. The pre-trained label embeddings are

learned from additional resources, e.g., text corpora, to encode semantic relation and dependency between labels. Similar to [90], we use two different methods for incorporating pre-trained label embeddings. The first one is to fully trust pre-trained label embeddings. Namely, we fix B as the pre-trained $\tilde{B}$ and only learn A in an iterative process. The second method is to use pre-trained label embedding as prior knowledge while adjusting both A and B according to the labeled data, i.e., adding a regularizer to the WARP loss function,

$$\sum_{y \in Y} \sum_{y' \in \tilde{Y}} L(\text{rank}(x, y)) \max(1 - f(x, y, A, B) + f(x, y', A, B), 0) + \lambda \|B - \tilde{B}\|_F^2,$$

where $\|\cdot\|_F$ is the Frobenius norm, and λ is the trade-off parameter.

4.4 Methods

4.4.1 Prototype-driven Label Embedding

Joint embedding methods such as WSABIE learn label embeddings from the whole training set including noisy labeled instances resulting from weak supervision. It is inevitable that the resulting label embeddings are affected by noisy labels and fail to accurately capture the semantic correlation between types. Another issue is that zero-shot frameworks such as DeVISE are not directly applicable to FNET as conceptually complex types, e.g., *GPE* (Geo-political Entity) cannot be simply mapped to a single natural word or phrase.

To address this issue, we propose a simple yet effective solution which is referred to as prototype-driven label embedding (ProtoLE), and henceforth we use $\tilde{B}^P$ to denote the label embedding matrix learned by ProtoLE. The first step is to learn a set of prototypes for each type in the type set. ProtoLE does not fully rely on training data to generate label embeddings. Instead, it selects a subset of entity mentions as the prototypes of each type. These prototypes are less ambiguous and noisy compared to the rest of the full set.

Even though it is already far less labor-intensive to manually select prototypes than annotating entity mentions one by one, we consider an alternative automated

process using Normalized Point-wise Mutual Information (NPMI) as the particular criterion for prototype selection. The NPMI between a label and an entity mention is computed as:

$$\text{NPMI}(y, m) = \frac{\text{PMI}(y, m)}{-\ln p(y, m)},$$

where $\text{NPMI}(\cdot, \cdot)$ is the point-wise mutual information computed as follows:

$$\text{PMI}(y, m) = \log \frac{p(y, m)}{p(y)p(m)},$$

where $p(y)$, $p(m)$ and $p(y, m)$ are the probability of entity mention m , label y and their joint probability. For each label, NPMI is computed for all the entity mentions and only a list of top k mentions are selected as prototypes. Note that NPMI is not applicable to unseen labels. In such case, it is necessary to combine manual selection and NPMI.

Word embeddings methods such as Skip-gram model [70] are shown capable of learning distributional semantics of words from unlabeled text corpora. To further avoid affected by label noises, we use pre-trained word embeddings as the source to compute prototype-driven label embeddings. For each label y_i , we compute its label embedding as the average of pre-trained word embeddings of the head words of prototypes, i.e.,

$$\tilde{B}_i^P = \frac{1}{k} \sum_{j=1}^k v_{m_{ik}},$$

where $v_{m_{ik}}$ denotes the word embedding of k th word in the prototype list of label y_i . In the case of using phrase embeddings, the full strings of multi-word prototypes could be used directly.

4.4.2 Hierarchical Label Embedding

Another side information that is available for generating label embeddings is the label hierarchy. We adapt the Hierarchical Label Embeddings (HLE) [90] to FNET task. Unlike [90], which uses the WordNet hierarchy, FNET systems typically have direct access to predefined tree hierarchy of type set. We denote the label embedding matrix resulting from label hierarchy as $\tilde{B}^H$. Each row in $\tilde{B}^H$ corresponds to a binary label

embedding and has a dimension equal to the size of label set. For each label, the sets $\tilde{B}_{ij}^H$ to 1 when y_j is the parent of y_i or $i = j$, and 0 to the remainder,

$$\tilde{B}_{ij}^H = \begin{cases} 1 & \text{if } i = j \text{ or } y_j \in \text{Parent}(y_i) \\ 0 & \text{otherwise} \end{cases}.$$

HLE explicitly encodes the hierarchical dependency between labels by scoring a type y_i given m using not only y_i but also its parent type $\text{Parent}(y_i)$. The underlying intuition is that recognition of a child type should be also based on the recognition of its parent.

4.4.3 Prototype-driven Hierarchical Label Embedding

One shortcoming of HLE is that it is too sparse. A natural solution is combining HLE with ProtoLE, which is denoted as Proto-HLE. Since $\tilde{B}^H \in \mathbb{R}^{N \times N}$ and $\tilde{B}^P \in \mathbb{R}^{D \times N}$, the combined embedding matrix $\tilde{B}^{HP}$ can be obtained by simply multiplying $\tilde{B}^H$ by $\tilde{B}^P$, i.e.,

$$\tilde{B}^{HP} = \tilde{B}^P \tilde{B}^{H\top}.$$

Note that $\tilde{B}^{HP}$ has the same shape as $\tilde{B}^P$, and it is actually representing the child label as a linear combination of the ProtoLE vectors of its parent and itself.

4.4.4 Type Inference

Having computed the scoring function for each label given a feature vector of the mention, we conduct type inference to refine the top k type candidates. In the setting of few-shots FNET, k is typically set to the maximum depth of type hierarchy, while different values for k may be used for a better prediction of unseen labels in zero-shot typing. For top k type candidates, we greedily remove the labels that conflict with others. However, unlike [6], we use a relative threshold t to decide whether the selected type should remain in the final results, which is more consistent with the margin-infused objective function than a global threshold. Namely, a type candidate will be passed to type inference only if the difference of score from the 1-best is less than a threshold.

4.5 Experiments

4.5.1 Experiment Setup

Our method uses feature templates similar to what have been used by state-of-the-art FNET methods [16, 88, 6, 34]. Table 4.1 illustrates the full set of feature templates used in this work. We evaluate the performance of our methods on three benchmark datasets that have been used for the FNET task: BBN dataset [95], OntoNotes dataset [96] and Wikipedia dataset [16]. Xiang Ren et al. [34] has pre-processed the training sets of BBN and OntoNotes using DBpedia Spotlight¹github.com/dbpedia-spotlight/dbpedia-spotlight. Entity mentions in the training set are automatically linked to a named entity in Freebase and assigned with the Freebase types of induced named entity. As shown in Table 4.2, BBN dataset contains 2.3K news articles of Wall Street Journal, which includes 109K entity mentions belonging to 47 types. OntoNotes contains 13.1K news articles and 223.3K entity mentions belonging to 89 entity types. The size of Wikipedia dataset is much larger than the other two with 2.69M entity mentions of 113 types extracted from 780.5K Wikipedia articles. Each dataset has a test set that is manually annotated for purpose of evaluation. To tune parameters such as the type inference threshold t and trade-off parameter λ , we randomly sample 10% instances from each testing set as the development sets and use the rest as evaluation sets.

Feature	Description
Tokens	Unigram words in the mentions
Head	Head word of the mention
Cluster	Brown Cluster IDs of the head word
POS Tag	POS tag of the mention
Character	Lower-cased character trigrams in the head word
Word Shape	The word shape of words in the mention
Context	Unigram/bigram words in context of the mention
Dependency	Dependency relations involving the head word

Table 4.1: Feature Template

Following prior works [16], we evaluate our methods and baseline systems using both loose and strict metrics, i.e., Macro-F1, Micro-F1, and strict Accuracy

Dataset		Types	Documents	Sentences	Mentions
BBN	train	47	2.3K	48.8K	109K
	test		459	6.4K	13.8K
OntoNotes	train	89	13.1K	147.7K	223.3K
	test		76	1.3K	9.6K
Wikipedia	train	113	780.5K	1.15M	2.69M
	test		-	434	563

Table 4.2: Statistics of datasets

Figure 4.1: t-SNE visualization of the prototype-driven label embeddings for BBN dataset [8]

(Acc.). Given the evaluation set D , we denote Y_m as the ground truth types for entity mention $m \in D$ and $\hat{Y}_m$ as the predicted labels. Strict accuracy (Acc) can be computed as: $\text{Acc} = \frac{1}{|D|} \sum_{m \in D} \sigma(Y_m = \hat{Y}_m)$, where $\sigma(\cdot)$ is an indicator function. Macro-F1 is based on Macro-Precision (Ma-P) and Micro-Recall (Ma-R), where $\text{Ma-P} = \frac{1}{|D|} \sum_{m \in D} \frac{|Y_m \cap \hat{Y}_m|}{Y_m}$, and $\text{Ma-R} = \frac{1}{|D|} \sum_{m \in D} \frac{|Y_m \cap \hat{Y}_m|}{\hat{Y}_m}$. And Micro-F1 is based on Micro-Precision (Mi-P) and Micro-Recall (Mi-R), where $\text{Mi-P} = \frac{\sum_{m \in D} |Y_m \cap \hat{Y}_m|}{\sum_{m \in D} \hat{Y}_m}$, and $\text{Mi-R} = \frac{\sum_{m \in D} |Y_m \cap \hat{Y}_m|}{\sum_{m \in D} Y_m}$.

4.5.2 Generating ProtoLE

Our ProtoLE embeddings use Continuous-Bag-of-Words (CBOW) word embedding model [70] trained on Wikipedia dump using a window of 2 words to both directions. We use 300 dimensions for all embedding methods except HLE. Table 4.3 illustrates examples of prototypes learned for types in BBN dataset. It can be observed that most of the top-ranked mentions are correctly linked to types, even though there are still some noises, e.g., north.american for /LOCATION/CONTINENT. It also shows that prototypes of related types such as /LOCATION and /GPE are also semantically related. Figure 4.1 visualizes the prototype-driven label embeddings for BBN dataset using Distributed Stochastic Neighbor Embedding (t-SNE) [97]. It can be easily observed that semantic related types are close to each other in the new space, which proves that prototype-driven label embeddings can capture the semantic correlation between labels.

Type	Prototypes
/LOCATION	areas connaught earth lane brooklyn
/LOCATION/CONTINENT	north.america europe africa north.american asia
/LOCATION/LAKE_SEA_OCEAN	big_bear lake_erie champ lake_geneva fujisawa
/LOCATION/RIVER	hudson thompson mississippi_river james_river tana
/GPE	soviet edisto canada china france
/GPE/STATE_PROVINCE	california texas ohio arizona jersey

Table 4.3: Example prototypes learned by PMI for types in BBN dataset

Figure 4.5 shows the Micro-F1 score of FNET with regard to the number of PMI prototypes used by ProtoLE. It shows that the Micro-F1 score does not change significantly on BBN and Wikipedia dataset, whereas using fewer prototypes per type (≤ 40) results in a drop of Micro-F1. Since the definitions of several types, especially the coarse-grained types, are actually very general, it may introduce bias into the label embeddings if using too few prototypes. We use $K = 60$ for all our experiments for that it achieves decent performance on all three datasets.

4.5.3 Few-shots Fine-grained Entity Typing

In this section, we compare performances of FNET methods in the setting of few-shots FNET where the training set covers all types. Methods compared in this section are

Figure 4.2: BBN Dataset

Figure 4.3: OntoNotes Dataset

Figure 4.4: OntoNotes Dataset

trained using the entire type set. We use evaluation metrics for our experiments: macro-F1, micro-F1 and accuracy. As in section 4.3.2, we train our label embeddings in two different ways: 1) non-adaptive training where label embeddings are fixed during training; and 2) adaptive training where label embeddings are also updated. Table 4.4 shows the comparison with the state-of-the-art FNET methods: FIGER [16], HYENA [26] and WSABIE [6]. We make several findings from the results.

Firstly, embedding methods with WARP loss function consistently outperform non-embedding methods (i.e., FIGER and HYENA) on all three datasets. The performance gaps are huge for BBN and OntoNotes, where the best embedding method achieves 10%-20% absolute improvement over the best non-embedding method (FIGER). However, the gap is much smaller on Wikipedia dataset whose size is significantly larger than the other two.

Secondly, non-adaptive embedding methods always outperform their adaptive versions except HLE on Wikipedia dataset. Performance of adaptive label embeddings are all close to WSABIE, which suggests that adaptive label embeddings might suffer from same label noise problem as WSABIE does.

Thirdly, our ProtoLE and its combination with HLE consistently outperform both non-embedding and embedding baselines. Using the prototype information and non-adaptive framework results in absolute 3%-5% improvement with both loose and strict evaluation metrics. Non-adaptive HLE performs poorer than other embedding methods, which is most likely due to its sparsity in representing labels. However, ProtoHLE performs very close to ProtoLE on BBN and Wiki, while it improves all three measures by another absolute $\approx 2.5\%$ on OntoNotes.

4.5.4 Zero-shot Fine-grained Entity Typing

In this section, we evaluate our method’s capability recognizing mentions of unseen fine-grained types. We assume that the training set contains only coarse-grained types (i.e., Level-1), and Level-2 types are unseen types to be removed from the training set. Table 4.5 shows the Micro-Precision for Level-1 and Level-2 types using top k type

Method	Adapt	BBN		
		Ma-F1	Mi-F1	Acc.
FIGER	NA	67.28	60.70	46.92
HYENA	NA	51.38	52.85	45.01
WSABIE	NA	71.28	72.08	66.22
HLE	Y	70.84	71.61	65.74
	N	68.86	70.00	63.32
ProtoLE	Y	72.67	73.54	67.58
	N	75.78	76.50	70.43
Proto-HLE	Y	71.97	72.89	67.05
	N	74.54	74.38	69.46
OntoNotes				
FIGER	NA	58.77	52.37	38.01
HYENA	NA	47.65	43.97	26.56
WSABIE	NA	62.03	55.83	43.61
HLE	Y	61.54	49.16	43.25
	N	59.52	54.01	41.60
ProtoLE	Y	60.90	54.68	42.82
	N	65.91	59.08	46.94
Proto-HLE	Y	62.71	56.64	44.81
	N	68.23	61.27	49.30
Wiki				
FIGER	NA	68.28	64.71	47.37
HYENA	NA	45.51	43.80	30.67
WSABIE	NA	67.97	64.49	48.28
HLE	Y	67.09	65.65	47.01
	N	65.29	62.53	45.19
ProtoLE	Y	66.96	65.78	49.18
	N	68.06	66.53	53.54
Proto-HLE	Y	67.85	65.74	50.27
	N	66.61	65.29	50.45

Table 4.4: Performance of FNET in a few-shots learning on 3 benchmark datasets

candidates for type inference. NPMI is computed for Level-1 types. We manually build prototype lists for unseen types by choosing from a randomly sampled list of entity mentions. Level-3 types are ignored for OntoNotes as Level-3 types never show in top-10 list produced by all methods. As the prediction for coarse-grained types are the same with regard to k , we only list the results using $k = 3$.

One interesting finding on all three datasets is that combining hierarchical and prototypical information results in better classification of coarse-grained types. It suggests that embeddings of unseen fine-grained types contains information complementary to the embeddings of coarse-grained types. Since HLE actually produces random prediction on Level-2 types due to its sparse representation, HLE perform

poorly on Level-2 types.

Dataset	Method	Micro-Precision @k Level 1	Micro-Precision @k Level 2		
		3	3	5	10
BBN	ProtoLE	76.71	42.95	36.61	42.34
	HLE	70.44	13.08	13.16	12.82
	Proto-HLE	76.89	42.35	35.18	30.16
OntoNotes	ProtoLE	73.26	21.01	13.72	12.22
	HLE	66.96	7.13	6.14	6.23
	Proto-HLE	76.33	7.09	11.43	9.91
Wiki	ProtoLE	65.52	12.50	21.28	17.91
	HLE	65.13	0.00	8.82	8.99
	Proto-HLE	67.41	20.01	31.25	24.24

Table 4.5: Performance of zero-shot entity typing

ProtoLE outperforms HLE by 100%-300% in terms of Micro-Precision. However, again the combination of prototypes and hierarchy achieves similar or better results than ProtoLE on BBN and Wikipedia dataset. The drop of precision of Proto-HLE on OntoNotes is likely due to a different nature of annotation. It is more prevalent in the test set of OntoNotes that one entity mention is annotated with multiple Level-1 types, and the presence of fine-grained types are less constrained by the label hierarchy. In such case, hierarchical constraints enforced by Proto-HLE might have negative impacts on type inference.

4.6 Summary

In this paper, we present a prototype-driven label embedding method for fine-grained named entity typing (FNET). It shows that our method outperforms state-of-the-art embedding-based FNET methods in both few-shots and zero-shots settings. It also shows that combining prototype-driven label embeddings and type hierarchy can improve the prediction of coarse-grained types.

Chapter 5

Binary Entity Embedding for Reranking ASR Hypotheses

In this chapter, we address the problem of reranking ASR hypotheses via binary embeddings of entities and words.

5.1 Introduction

Reranking models have been shown effective for reducing errors in a variety of Natural Language Processing tasks such as Named Entity Recognition [24, 8], Syntactic Parsing [98, 99] and Statistical Machine Translation [100]. A reranking model typically treats the baseline system as a black box and is trained to rank the competing hypotheses based on more complex or global information.

In an ASR system, discriminative language model (DLM) is first introduced by Roark et al. [101] for reranking ASR hypotheses. They adopt a single perceptron to modify the confidence scores of hypotheses generated by a baseline ASR system. By using only n-gram features, their reranking model is shown capable of reducing the Word Error Rate (WER) of an ASR system. His work is followed by several variants with a variety of feature choices such as syntactic features [55, 56], which try to capture correlation between simple features on the feature level. However, existing DLMs still suffer from poor generalization power and are vulnerable to a shortage of training data, because most of them rely on linear or log-linear models that fail to take into consideration the correlation of input features on the model level. Apart

from feature engineering, using hidden variables encoding semantic information helps to improve the generalization power.

Koo et al. [99] proposes a hidden-variable model to rerank syntactic parsing trees. By linking input features to hidden states corresponding to word senses or classes, they achieve improved accuracy over a linear baseline. Inspired by the success of Koo et al., we propose to use the computational structure of Restricted Boltzmann Machine (RBM) [102] for the task of ASR hypotheses reranking. RBM is a neural network composed of one hidden layer and one input layer. The hidden layer of RBM can be viewed as embeddings of the input and has been shown capable of capturing high-order correlation and semantic information in the context of language modeling [103, 104]. These approaches model the probability of a fixed length of word sequences, i.e., N -grams, using only local information, and are trained with a generative objective function. However, RBM cannot be directly used for ASR reranking due to its generative training manner.

In this chapter, we propose two modifications to train RBM in a discriminative manner. We modify the energy function of RBM to incorporate the ASR confidence score, which has been proved critical for reranking by previous DLMs [101, 55, 56]. We then propose a novel discriminative objective function for training RBM with N -best lists of ASR hypotheses. Our method differs from existing RBM-based language models [103, 104] in two major aspects. Firstly, the proposed RBM reranker is trained discriminatively. Secondly, RBM in our method represents sentences of variable length as global feature vectors. The most attractive property of RBM is that the computational structure is flexible enough to incorporate prior knowledge [105]. As function words have little meanings and are less important for language understanding [106], we decide to focus more on content words. In order to allow the reranking model to have a better knowledge of content words, we propose to fuse the hidden layer with information about a particular class of natural language concepts, i.e., named entities. For technical simplicity, we constrain the hidden layer of RBM to take only binary values. Therefore, the hidden layer can be seen as a binary embedding of the input.

In this regard, the resulting binary embedding encodes the fusion of both word-level and concept-level information.

To our knowledge, this work is the first to consider using hidden layer and fusion of knowledge from both concept and word levels in the context of ASR hypotheses reranking. The remainder of this chapter is structured as follows: Section 2 describes in detail the proposed work; Section 3 shows the empirical results as well as analyses; finally, Section 4 concludes this work and provide more discussion on future work.

5.2 Training RBM for ASR Reranking

5.2.1 Restricted Boltzmann Machine

A Restricted Boltzmann Machine [102] (see Figure 5.1) is a neural network composed of : one n -dimension input feature layer $\phi(\mathbf{t}) = [\phi_1(t), \phi_2(t), \dots, \phi_n(t)]$, which is a global feature vector extracted for a raw input t , and one d -dimension binary hidden layer $\mathbf{h} = [h_1, h_2, \dots, h_d]$. The joint probability $P_{\text{RBM}}(t, h)$ of hidden variables and raw input is defined as

$$P_{\text{RBM}}(t, h) = \frac{e^{-E_{\text{RBM}}(t, h)}}{\sum_{t, h} e^{-E_{\text{RBM}}(t, h)}}$$

$$E_{\text{RBM}}(t, h) = -\phi(t)^T W h - b^T \phi(t) - c^T h,$$

where $W \in R^{n \times d}$ is the matrix specifying the weights of connections between hidden and input layer, and $b \in R^n$ and $c \in R^d$ are the bias vectors of the two layers. $E_{\text{RBM}}(t, h)$ is called the energy function of RBM. The probability of a raw input t is then defined as the marginal probability of t

$$P_{\text{RBM}}(t) = \sum_h P_{\text{RBM}}(t, h)$$

, and the training objective is to maximize the log likelihood of training data D

$$\sum_{t \in D} \ln P_{\text{RBM}}(t)$$

5.2.2 Maximum Margin Training for RBM-based Reranker

The goal of generative training of RBM is to learn a probability distribution, which is not necessary for choosing correct ASR hypotheses. Instead, the discriminative training allows the model to explicitly select ASR hypotheses containing fewer errors. In this section, we describe our discriminatively trained RBM-based reranking model, denoted as **dRBM**.

Before introducing the training objective function, we first introduce the energy function of RBM model. ASR posterior probabilities produced by the baseline ASR system have been shown useful for reranking in previous works on DLM [101, 55, 56]. We hence add ASR posterior to the energy function of RBM. The modified energy function is expressed as

$$\begin{aligned} E_{\text{dRBM}}(t, h) &= E_{\text{RBM}}(t, h) + E_{\text{asr}}(t) \\ E_{\text{asr}}(t) &= -w_0 \ln(P(t|a)), \end{aligned}$$

where $P(t|a)$ is the posterior probability of a given ASR hypothesis t given the acoustic input a , and w_0 is the weight of ASR confidence score fixed during training. We represent each hypothesis as a global feature vector $\phi(t)$ using a predefined set of feature functions. In this chapter, we mainly consider using unigram features, yet using more complicated features does not need to change the model.

Inspired by the maximum margin training for Bayesian Networks [107], we adopt a discriminative objective function L using likelihood ratio,

$$L = \frac{1}{|D|} \sum_{a \in D} \sum_{t' \in \text{GEN}(a)} \max(1 - \ln \frac{P_{\text{dRBM}}(\hat{t})}{P_{\text{dRBM}}(t')}, 0),$$

where D is the training set for the discriminative training of RBM and $|D|$ denotes the number of utterances in training set. $\text{GEN}(a)$ refers to the list of N -best hypotheses generated by the baseline ASR system for the acoustic input a , while $\hat{t}$ is the oracle-best in the N -best list of t . Intuitively, the learning process finds the parameter setting maximizing the margin between the oracle-best hypotheses and other hypotheses in the N -best list. The subgradient of the objective function is

$$\frac{\partial L}{\partial \theta} = \sum_{a \in D} \sum_{t' \in \text{GEN}(a)} \mathbb{I}(\mathcal{F}(\hat{t}) - \mathcal{F}(t') < 1) \left(\frac{\partial \mathcal{F}(t')}{\partial \theta} - \frac{\partial \mathcal{F}(\hat{t})}{\partial \theta} \right),$$

where $\mathcal{F}(\cdot)$ is the free energy of RBM defined as

$$\mathcal{F}(t) = -\ln \sum_h e^{-E_{\text{dRBM}}(t, h)}$$

Algorithm 1: Discriminative training for RBM

Input:

D : the training dataset

$\text{GEN}(a)$: N -best list for an utterance t in the reference

λ : learning rate

```

1 for  $k=1:K$  do
2   for  $a \in D$  do
3     Positive:
4      $\hat{t} = \text{argmin}_{t' \in \text{GEN}(a)} \text{WER}(t')$ 
5     Negative:
6      $T^- = \{t' | 1 + \text{Score}(t') > \text{Score}(\hat{t}), t' \in \text{GEN}(t)\}$ 
7     for  $t' \in T^-$  do
8        $\theta \leftarrow \theta + \lambda \frac{\partial -\mathcal{F}(\hat{t})}{\partial \theta}$ 
9        $\theta \leftarrow \theta - \lambda \frac{\partial -\mathcal{F}(t')}{\partial \theta}$ 
    
```

The training algorithm is described in Algorithm 1. For each acoustic input in training set, we select a set of hypotheses T^- , which are ranked higher than the oracle best hypothesis. Based on our analysis of the loss function, we boost the score of oracle best with its derivate of negative free energy and penalize hypothesis in T^- with their derivates of negative free energy. Note that, as compared with standard RBM training [108], which iterates over input space or samples of input space, our discriminative training needs only to iterate over the N -best list which grows linearly with the size of training data and N -best list.

5.2.3 Training with Prior Knowledge

The binary hidden layer of RBM allows for easily incorporating prior knowledge into the reranking model. We consider using two types of prior knowledge: named entity

Figure 5.1: Structure of RBM with Entity-related Prior [9]

labels and pre-trained latent layer from texts. Firstly, to improve the capability of recognizing content words, we capture prior of a special class of content words – named entities. As entity related prior also encodes information about word classes, it helps improving the generalization power of language models [109] as well.

Specifically, we extract pairs of named entity words and their classes from texts using a named entity tagger, which annotates the text with 3 widely-adopted named entity classes, i.e., LOCATION, ORGANIZATION and PERSON. As show in Fig. 5.1, 3 variables in the hidden layer of RBM are used to represent named entity classes. For purpose of reducing ambiguity, we remove words belonging to multiple entity classes. We denote the list of entity-class pairs as $G = \{w, e\}$, where w is an index of the unigram feature in the input layer, and e an index of entity-class variable in the hidden layer. The objective function is then augmented with an entity-related regularizer,

$$L - \lambda \ln \prod_{w, e \in G} (P(h_e = 1 | \phi_w) - 1)^2,$$

$$P(h_e | \phi_w) = \sigma(c_e + W_{e,w} \phi_w).$$

As introduced in Wang et al. [105], $P(h_e | \phi_w)$ denotes the probability of a hidden variable h_e being activated by a given input feature ϕ_w .

To handle the data sparsity, we initiate connection matrix W of RBM with values pre-trained using a large text corpus and the generative training. The pretraining captures the distributional semantics of input features [110].

5.2.4 Scoring ASR Hypotheses

To score a given hypothesis, we propose two scoring functions using our RBM-based reranker and its combination with SLP. First of all, the RBM-based reranking score S_{RBM} is defined as the logarithm of the unnormalized probability $\tilde{P}_{\text{dRBM}}(t)$ assigned by the RBM-based reranker solely,

$$\begin{aligned} S_{\text{RBM}}(t) &= \ln \tilde{P}_{\text{dRBM}}(t) \\ &= \underbrace{w_0 \ln(P(t|a))}_{\text{ASR posterior}} + \underbrace{\sum_i^n b_i \phi_i(t)}_{\text{linear part}} + \underbrace{\sum_j^d \ln(1 + e^{(c_j + W_j \phi(t))})}_{\text{hidden variable part}}. \end{aligned}$$

As shown above, the re-scoring function is composed of the original ASR posterior, a linear bias, and the hidden variable component. In addition, SLP and RBM are likely to have encoded information complementary to each other due to their different structures and training methods. Therefore, we propose a late fusion of the two methods, which combines their confidence scores in the testing phase. The combined reranking score is

$$S(t) = S_{\text{RBM}}(t) + \alpha S_{\text{SLP}}(t),$$

where $S_{\text{SLP}}(t)$ is the single perceptron based confidence score weighted by α .

5.3 Related Work

DLM has been first introduced by Roark et al. in [101], where simple features like N -gram is shown able to effectively reduce WER. This previous work is using a Single Layer Perceptron (SLP) to modify the original posterior probabilities of the outputs of a baseline ASR system using a linear function,

$$\log P(t|a) + \sum_i w_i f_i(t),$$

where $\log P(t|a)$ is the log probability of a word sequence t given the acoustic signal a , and $\{f_i(\cdot)\}$ are the set of feature functions of an utterance weighted by $\{w_i\}$. Different types of features extracted from syntactic trees [55] and dependency trees [56] have also been used to enrich the feature set.

Apart from feature engineering and using a linear combination of feature functions, inferring hidden variables from the observed input captures semantic information related to word classes and word senses. Our work is closely related to Koo et al. [99] who proposed a hidden-variable model to rerank syntactic parsing trees. For the tractability of their model, they put constraints on the connections between latent variables and visible variables (i.e., input layer) by splitting features into two sets. However, the way they divide features is specific to syntactic parsing and thus is not applicable to our task. Our model differs from Koo et al. [99] in the sense that the connection is not constrained by their feature type, but instead relying on the structure of RBM to build connections between input and hidden layer.

RBM-based models [103, 104] have been explored for language modeling. Both approaches model the probability of a fixed length of word sequences, i.e., N -grams, and trained with a generative objective function. Our method differs from these methods in two major aspects: the training of proposed RBM reranker is discriminative, and it represents sentences of variable length as global feature vectors.

5.4 Experiment

5.4.1 Dataset

We evaluate our work on the latest release of TedLium Corpus [111] which is a set of audio and manually transcribed texts of Ted talks. As shown in Table 5.1. We split the training set of TedLium Version 2 into two parts: former TedLium Training set Version 1 and the rest. The Version 1 part is a set of 774 Ted talks consisting of 56,800 utterances and more than 1.7 million words, while the remaining of the TedLium training set contains another 718 talks. The evaluation set of our experiment is the testing set of TedLium corpus, which is composed of 11 talks. Our text corpus is the ukWaC corpus [82], which is a collection of texts containing about 1.8 billion words.

	Utterances	Talks	Words
ASR AM Train	56.8K	774	1.7M
Reranking Train(speech)	36.2K	718	0.9M
Reranking Train(text)	24M	-	1.8B
Reranking Test	1.15K	11	29K

Table 5.1: Characteristics of the datasets used in experiments

5.4.2 Baseline

The baseline ASR system is based on KALDI¹ toolkit [112] including a DNN-based acoustic model. It uses a pre-trained language model², which is released as part of Sphinx project [113] and has achieved a perplexity of 158.3 on a corpus of Ted Talks. The acoustic model is trained using the training set of TedLium Version 1. The rest of training set of TedLium Version 2 is used for training reranking models.

The baseline SLP reranker is trained by following the work of Lambert et al. [56] that randomly selects K pairs of hypotheses from the N -best lists. Specifically, we randomly select 100 pairs from the 100-best list. We ran the training process for 10 iterations, as we cannot observe further WER reduction with more iterations.

5.4.3 RBM Setup

We refer to the system integrating prior knowledge with dRBM as p-dRBM. Both dRBM and p-dRBM use 200 hidden units and are trained using the same dataset and 100-best hypotheses as SLP reranker. Since our focus is not on feature engineering, and for simplicity of interpreting our experiment results, we use only unigram features (i.e., single words) in our experiments, which have also been shown as the most effective features by previous works [56]. For training p-dRBM, we first crawled down a set of text summaries of ted talks from Ted website. We then create a list of 20 words for each entity category by tagging the collected text summaries with Stanford named entity recognition tool [114] following description in section 5.2.3. A basic RBM is trained using ukWaC corpus and used for initiating the connection matrix W of p-dRBM. The late fusion of SLP and our proposed methods are denoted as SLP +

¹<http://kaldi.sourceforge.net/>

²<http://cmusphinx.sourceforge.net/2013/01/a-new-english-language-model-release/>

dRBM and SLP + p-dRBM. We use $\lambda = 0.01$ and $\alpha = 1.0$ as weights of entity-related regularizer and SLP scores in late fusion.

5.4.4 Evaluation

First of all, we analyze the behavior of p-dRBM by computing the most-activated words by p-dRBM as shown in Table 5.2. As shown in the previous section, the scoring function used by our method is a combination of ASR confidence score, a linear component (denoted as p-dRBM-L) and a hidden-variable component (denoted as p-dRBM-H). p-dRBM then takes as input one-hot vectors of words to compute their reranking scores. It shows that the linear component is mainly accounting for the function words, while the hidden component favors content words that are mostly nouns and adjectives. The final scoring function is a trade-off between function words and content words through a combination of the two components.

p-dRBM-H	p-dRBM-L	p-dRBM
integrated	of	integrated
demeanor	and	demeanor
disgust	the	disgust
tattoo	to	tattoo
formula	a	formula

Table 5.2: Most activated word by p-dRBM

To investigate what is captured by RBM and potentially effective for improving the Word Error Rate (WER), we represent each hidden variable as a vector of words. These vectors represent how much a word is activated by a given hidden variable. Table 5.3 shows a selected set of hidden variables that can be seen as a set of topics. We found that RBM can capture meaningful topics by using only sentence-level co-occurrence. We then represent each word as a vector of hidden variables by taking the rows of the matrix $W \in R^{|V| \times d}$ of p-dRBM.

We rank words based on their cosine similarity with the queries and select the top 5 words for four query words. As shown in Table 5.4, the top-ranked words all seem very relevant to the query words. Since our RBM is trained with sentence-level concurrence, which is different from the window-based methods, the word ‘similarity’

working	media	higher education	entertainment
security	news	cambridge	scene
services	forum	mary	story
office	business	professor	tv
home	press	william	songs
for	new	royal	moving

Table 5.3: Example topics learned by p-dRBM

seems more like a topical relatedness rather than syntactical similarity. In general, we can conclude that the resultant RBM-based reranking model to some extent captures the distributional semantics related to the topics of words.

japan	film	bible	computer
india	story	greatest	software
italy	music	holy	database
asia	beautiful	truth	digital
germany	famous	gospel	user
china	classic	spirit	server

Table 5.4: Most-similar words for queries using p-dRBM word embeddings

As shown in Table 5.5, we evaluate WER of reranking systems. It shows that the proposed discriminatively trained RBM produces greater WER reduction than baseline SLP rerankers. Effectiveness of using prior knowledge is validated by further improving WER over dRBM. The greatest absolute WER reduction (1.3%) is achieved by the late fusion of SLP and p-dRBM, which confirms that our reranker captures information complementary to SLP.

	WER	WER (TF-IDF ≥ 3)
ASR 1-best	18.23	46.9
Oracle 1-best	11.42	36.1
SLP	17.76	46.3
dRBM	17.51	44.6
p-dRBM	17.36	43.8
SLP + dRBM	17.11	45.2
SLP + p-dRBM	16.91	44.2

Table 5.5: Performance of reranking model on TedLium corpus

Since the latent layer of p-dRBM incorporates prior knowledge related to content words (e.g., named entities), it is desirable that the proposed method can better

recognize content words, which are more critical for downstream applications such as spoken language understanding. To evaluate the performance of our proposed methods on recognizing content words in a more general way, we words that have higher TF-IDF scores are more likely to be content words. We hence assign more weight to errors involving a set of keywords with high TF-IDF instead of treating all words equally. Specifically, the list of keywords are chosen based on TF-IDF scores (≥ 3.0) computed from whole TedLium corpus. We use the weighted-word-scoring implementation in NIST SCLITE tool ³ by aggressively assigning weight 1.0 to words on the list and 0.0 to the rest.

Figure 5.2: WER reduction for words versus TF-IDF scores [9]

Table 5.5 clearly shows that baseline reranking systems (SLP) fail to reduce much WER for selected keywords. In comparison, proposed RBM rerankers, especially p-dRBM, have reduced more errors on chosen keywords without sacrificing overall performance. We further break down the TF-IDF scores into 3 bins. Figure 5.2 shows the WER reduction by all three approaches. Thanks to hidden variables, our methods are capable of better capturing the discriminative information for most content words. In particular, p-dRBM is shown working significantly better than other methods on words with medium TF-IDF scores (<5), which is a result of injecting named entity words, e.g., washington (TF-IDF= 3.87), that mostly have TF-IDF between 3.0 and 5.0.

³<http://www1.icsi.berkeley.edu/Speech/docs/sctk-1.2/sclite.htm>

5.5 Summary

In this chapter, we proposed an RBM-based reranking model that is discriminatively trained for reranking ASR hypotheses. The proposed reranking learns a binary embedding of words by fusing information of concept-level and word-level in its hidden layer. In comparison with single perceptron based reranker, our proposed approach achieves higher WER reduction. The success of combining single perceptron and RBM-based reranker suggests that two models actually capture complementary information useful for selecting less erroneous ASR hypotheses. In addition, we found that introducing concept-level knowledge to RBM-based reranker results in a better recognition of content words. In the future, we would like to explore the use of lexical knowledge obtained from different resources, e.g., WordNet [115] or SenticNet [15], as additional prior knowledge for the proposed model.

Chapter 6

Embedding Concepts for Target-based Sentiment Analysis

In this chapter, we focus on leveraging commonsense concepts for target-based sentiment analysis.

6.1 Introduction

In the past decade, sentiment analysis [116] has become increasingly popular for processing social media data on online communities, blogs, wikis, microblogging platforms, and other online collaborative media. Sentiment analysis is a branch of affective computing research [117] that aims to classify text into either positive or negative, but sometimes also neutral [118]. Most of the literature is on English language but recently an increasing number of publications is tackling the multilinguality issue [119].

While most works approach it as a simple categorization problem, sentiment analysis is actually a suitcase research problem [2] that requires tackling many NLP tasks, including named entity recognition [8], word polarity disambiguation [120], personality recognition [121], sarcasm detection [122], and aspect extraction. The last one, in particular, is an extremely important subtask that, if ignored, can consistently reduce the accuracy of sentiment classification in the presence of multiple opinion targets.

Hence, aspect-based sentiment analysis (ABSA) [123, 124, 125] extends the typical setting of sentiment analysis with a more realistic assumption that polarity is associated with specific aspects (or product features) rather than the whole text unit.

For example, in the sentence “The design of the space is good but the service is horrible”, the sentiment expressed towards the two aspects (“space” and “service”) is completely opposite. Through aggregating sentiment analysis with aspects, ABSA allows the model to produce a fine-grained understanding of people’s opinion towards a particular product.

Targeted (or target-dependent) sentiment classification [42, 38, 45], instead, resolves the sentiment polarity of a given target in its context, assuming that a sentence might express different opinions towards different targeted entities. For instance, in the sentence “I just log on my [facebook]. [Transformers] is boring”, the sentiment expressed towards [Transformers] is negative, while there is no clear sentiment for [facebook]. Recently, targeted ABSA [46] has attempted to tackle the challenges of both ABSA and targeted sentiment analysis. The task is to jointly detect the aspect category and resolve the polarity of aspects with respect to a given target.

Deep learning methods [40, 41, 42, 43, 45] have achieved great accuracy when applied to ABSA and targeted sentiment analysis. Especially, neural sequential models, such as LSTM networks [22], are of growing interest for their capacity of representing sequential information. Moreover, most of these sequence-based methods incorporate the attention mechanism, which has its root in the alignment model of machine translation [126]. Such mechanism takes an external memory and representations of a sequence as input and produces a probability distribution quantifying the concerns in each position of the sequence.

Figure 6.1: Example of Sentihood data set

Despite these advances in sentiment analysis, we identify three problems remaining unsolved in current state-of-the-art methods. Firstly, a given target might consist of multiple instances (mentions of the same target) or multiple words in a sentence, existing research assumes all instances are of equal importance and simply computes an average vector over such instances. This oversimplification conflicts with the fact that one or more instances of the target are often more tightly tied to sentiment than others. Secondly, hierarchical attention exploited by existing methods only implicitly models the process of inferring the sentiment-bearing words related to the given target and aspect as black-box. Last but not least, existing research falls short in effectively incorporating into the deep neural network external knowledge, e.g., affective or commonsense knowledge, that could directly contribute to the identification of aspects and sentiment polarity. Without any constraints, moreover, the global attention model might tend to encode task-irrelevant information. To address these problems, our method simultaneously learns a target-specific instance attention as well as a global attention. In particular, our contribution is three-fold:

- (i) We propose a hierarchical attention model that explicitly attends to first the targets and then the whole sentence;
- (ii) We extend the classic LSTM cell with components accounting for integration with external knowledge;
- (iii) We fuse affective commonsense concepts with word-level information in a deep neural network.

6.2 Methodology

In this section, we describe the proposed attention-based neural architecture in detail: we first proposed the task definition of targeted ABSA, followed by an overview of the whole neural architecture; afterwards, we describe instance attention and global attention model; lastly, we describe the proposed knowledge-embedded extension of LSTM cell.

6.2.1 Task Definition

A sentence s consists of a sequence of words. Similar to [45], we consider all mentions of the same target as a single target. A target t composed of m words in sentence s , denoted as $T = \{t_1, t_2, \dots, t_i, \dots, t_m\}$ with t_i referring to the position of i th word in the target expression, the task of targeted ABSA can be divided into two subtasks. Firstly, it resolves the aspect categories of t belonging to a predefined set. Secondly, it classifies the sentiment polarity with respect to each aspect category associated with t .

For example, the sentence “I live in [West London] for years. I like it and it is safe to live in much of [west London]. Except [Brent] maybe. ” contains two targets, $[WestLondon]$ and $[Brent]$. Our objective is to detect the aspects and classify the sentiment polarity. The desired output for $[WestLondon]$ is $[general:\textbf{positive}; safety:\textbf{positive}]$, while output for $[Brent]$ should be $[general:\textbf{negative}; safety:\textbf{negative}]$.

6.2.2 Overview

In this section, we provide an overview of the proposed method. Our neural architecture consists of two components: the sequence encoder and a hierarchical attention component.

Fig. 6.2 illustrates how the neural architecture works. Given a sentence $s = \{w_1, w_2, \dots, w_L\}$, a look-up operation is first performed to convert input words into

Figure 6.2: Overview of the attentive neural architecture

word embeddings $\{v_{w_1}, v_{w_2}, \dots, v_{w_L}\}$. The sequence encoder, which is based on a bidirectional LSTM, transforms the word embeddings into a sequence of hidden outputs. The attention component is built on top of the hidden outputs. The target-level attention takes as input the hidden outputs at the positions of target expression (highlighted in brown) and computes a self-attention vector over these words.

The output of target-level attention component is a representation of the target. Afterwards, the target representation together with the aspect embeddings is used for computing a sentence-level attention transforming the whole sentence into a vector. The sentence-level attention component returns one sentence vector for each aspect and target pair. The aspect-based sentence vector is then fed into the corresponding multi-class (e.g., None, Neural, Negative, and Positive for a 4-class setting; or None, Negative, and Positive for a 3-class setting) classifier to resolve the sentiment polarity.

6.2.3 Long Short-Term Memory Network

The sentence is encoded using an extension of RNN [127], termed LSTM [22], which is firstly introduced by [22] to solve the vanishing and exploding gradient problem faced by the vanilla RNN. A typical LSTM cell contains three gates: forget gate, input gate and output gate. These gates determine the information to flow in and flow out at the current time step. The mathematical representations of the cell are as follows:

$$\begin{aligned}
 f_i &= \sigma(W_f[x_i, h_{i-1}] + b_f) \\
 I_i &= \sigma(W_I[x_i, h_{i-1}] + b_I) \\
 \tilde{C}_i &= \tanh(W_C[x_i, h_{i-1}] + b_C) \\
 C_i &= f_i * C_{i-1} + I_i * \tilde{C}_i \\
 o_i &= \sigma(W_o[x_i, h_{i-1}] + b_o) \\
 h_i &= o_i * \tanh(C_i)
 \end{aligned} \tag{6.1}$$

where f_i , I_i and o_i are the forget gate, input gate and output gate, respectively. W_f , W_I , W_o , b_f , b_I and b_o are the weight matrix and bias scalar for each gate. C_i is the cell state and h_i is the hidden output. A single LSTM typically encodes the sequence from only one direction. However, two LSTMs can also be stacked to be

used as a bidirectional encoder, referred to as bidirectional LSTM. For a sentence $s = \{w_1, w_2, \dots, w_L\}$, bidirectional LSTM produces a sequence of hidden outputs,

$$H = [h_1, h_2 \dots h_L] = \begin{bmatrix} \overrightarrow{h_1} & \overrightarrow{h_2} & \dots & \overrightarrow{h_L} \\ \overleftarrow{h_1} & \overleftarrow{h_2} & \dots & \overleftarrow{h_L} \end{bmatrix}$$

where each element of H is a concatenation of the corresponding hidden outputs of both forward and backward LSTM cells.

6.2.4 Target-Level Attention

Based on the attention mechanism, we calculate an attention vector for a target expression. A target might consist of a consecutive or non-consecutive sequence of words, denoted as $T = \{t_1, t_2, \dots, t_m\}$, where t_i is the location of an individual word in a target expression. The hidden outputs corresponding to T is denoted as $H' = \{h_{t_1}, h_{t_2}, \dots, h_{t_m}\}$. We compute the vector representation of a target t as

$$v_t = H' \alpha = \sum_j \alpha_j h_{t_j} \quad (6.2)$$

where the target attention vector $\alpha = \{\alpha_1, \alpha_2, \dots, \alpha_m\}$ is distributed over target word sequence T . The attention vector α is a self-attention vector that takes nothing but the hidden output itself as input. The attention vector α of target expression is computed by feeding the hidden output into a bi-layer perceptron, as shown in Equation 6.3.

$$\alpha = \text{softmax}(W_a^{(2)} \tanh(W_a^{(1)} H')) \quad (6.3)$$

where $W_a^{(1)} \in R^{d_m \times d_h}$ and $W_a^{(2)} \in R^{1 \times d_m}$ are parameters of the attention component.

6.2.5 Sentence-Level Attention Model

Following the target-level attention, our model learns a target-and-aspect-specific sentence attention over all the words of a sentence. Given a sentence s of length L , the hidden outputs are denoted as $H = [h_1, h_2, \dots, h_L]$. An attention model computes a linear combination of the hidden vectors into a single vector, i.e.,

$$v_{s,t}^a = H \beta = \sum_i \beta_i h_i \quad (6.4)$$

where the vector $\beta = [\beta_1, \beta_2, \dots, \beta_L]$ is called the sentence-level attention vector. Each element β_i encodes the salience of the word w_i in the sentence s with respect to the aspect a and target T . Existing research on targeted sentiment analysis or ABSA mostly uses targets or aspect terms as queries.

At first, each h_i is transformed to a d_m dimensional vector by a multi-layer neural network with a *tanh* activation function, followed by a dense softmax layer to generate a probability distribution over the words in sentence s , i.e.,

$$\beta_a = \text{softmax}(v_a^T \tanh(W_m(H' \odot v_t))) \quad (6.5)$$

where v_a is the aspect embedding of aspect a , $H \odot v_t$ is the operation concatenating v_t to each h_i ; $W_m^{(1)} \in R^{d_m \times d_h}$ is the matrix mapping row vectors of H to a d_m dimensional space, and $W_m^{(2)} \in R^{1 \times d_m}$ maps each new row vector to a unnormalized attention weight.

6.2.6 Commonsense Knowledge

In order to improve the accuracy of sentiment classification, we use commonsense knowledge as our knowledge source to be embedded into the sequence encoder. In particular, we use SenticNet [15], a commonsense knowledge base that contains 100,000 concepts associated with a rich set of affective properties (Table 6.1). These affective properties provide not only concept-level representation but also semantic links to the aspects and their sentiment. For example, the concept ‘rotten fish’ has property “KindOf-food” that directly relates with aspects such as ‘restaurant’ or ‘food quality’, but also emotions, e.g., ‘joy’, which can support polarity detection (Fig. 6.3).

However, the high dimensionality of SenticNet hinders it from being used in deep neural models. AffectiveSpace [10] has been built to map the concepts of SenticNet to continuous low-dimensional embeddings without losing the semantic and affective relatedness of the original space. We have included a brief introduction of efficiently generating low-dimensional concept embeddings using random projection in Section 2.5.1.2. Based on this new space of concepts, we embed concept-level information into deep neural sequential models to better classify both aspects and sentiment in natural language text.

6.2.7 Sentic LSTM

In order to leverage SenticNet’s affective commonsense knowledge efficiently, we propose an affective extension of LSTM, termed Sentic LSTM. It is reasonable to assume that SenticNet concepts contain information complementary to the textual word sequence as, by definition, commonsense knowledge is about concepts that are usually taken for granted and, hence, absent from the text. Sentic LSTM aims to entitle the concepts with two important roles: 1) assisting with the filtering of information flowing from one time step to the next and 2) providing complementary information to the memory cell. At each time step i , we assume that a set of knowledge concept candidates can be triggered and mapped to a d_c dimensional space. We denote the set of K concepts as $\{\mu_{i,1}, \mu_{i,2}, \dots, \mu_{i,K}\}$. First, we combine the candidate embeddings into a single vector as follows:

$$\mu_i = \frac{1}{K} \sum_j \mu_{i,j} \quad (6.6)$$

As we realized that there are only up to 4 extracted concepts for each time step, we simply use the average vector (although a more sophisticated attention model can also be easily employed to replace the averaging function).

Table 6.1: Example of SenticNet assertions

SenticNet	IsA-pet	KindOf-food	Arises-joy	...
dog	0.981	0	0.789	...
cupcake	0	0.922	0.910	...
songbird	0.672	0	0.862	...
gift	0	0	0.899	...
sandwich	0	0.853	0.768	...
rotten fish	0	0.459	0	...
win lottery	0	0	0.991	...
bunny	0.611	0.892	0.594	...
police man	0	0	0	...
cat	0.913	0	0.699	...
rattlesnake	0.432	0.235	0	...

$$\begin{aligned}
 f_i &= \sigma(W_f[x_i, h_{i-1}, \mu_i] + b_f) \\
 I_i &= \sigma(W_I[x_i, h_{i-1}, \mu_i] + b_I) \\
 \tilde{C}_i &= \tanh(W_C[x_i, h_{i-1}] + b_C) \\
 C_i &= f_i * C_{i-1} + I_i * \tilde{C}_i \\
 o_i &= \sigma(W_o[x_i, h_{i-1}, \mu_i] + b_o) \\
 o_i^c &= \sigma(W_{co}[x_i, h_{i-1}, \mu_i] + b_{co}) \\
 h_i &= o_i * \tanh(C_i) + o_i^c * \tanh(W_c \mu_i)
 \end{aligned} \tag{6.7}$$

Our affective extension of LSTM is illustrated in Equation 6.7. At first, we assume that affective concepts are meaningful cues to control the information of token-level information. For example, a multi-word concept ‘rotten fish’ might indicate that the word ‘rotten’ is a sentiment-related modifier of its next word ‘fish’ and, hence, less information should be filtered out at next time step. We thus add knowledge concepts to the forget, input, and output gate of standard LSTM to help to filter the information. The presence of affective concepts in the input gate is expected to prevent the memory cell from being affected by input tokens conflicting with pre-existing knowledge. Similarly, the output gate uses such knowledge to filter out irrelevant information stored in the memory.

Another important feature of Sentic LSTM is based on the assumption that the information from the concept-level output is complementary to the token level. Therefore, we extended the regular LSTM with an additional knowledge output gate o_i^c to output concept-level knowledge complementary to the token-level memory. Since AffectiveSpace is learned independently, we leverage a transformation matrix $W_c \in R^{d_h \times d_\mu}$ to map it to the same space as the memory outputs. In other words, o_i^c models the relative contributions of token level and concept level.

Moreover, we notice that $o_i^c * \tanh(W_c \mu_i)$ actually resembles the functionality of the sentinel vector used by [51], which allows the model to choose whether to use affective knowledge or not.

6.2.8 Prediction and Parameter Learning

The objective to train our classifier is defined as minimizing the sum of the cross-entropy losses of prediction on each target-aspect pair, i.e.,

$$\mathcal{L}_s = \frac{1}{|D|} \sum_{s \in D} \sum_{t \in s} \sum_{a \in A} \log p_{c,t}^a$$

where A is the set of predefined aspects, and $p_{c,t}^a$ is the probability of the gold-standard polarity class c given target t with respect to a sentiment category a , which is defined by a softmax function,

$$p_{c,t}^a = \text{softmax}(W^p v_{s,t}^a + b_s^a)$$

where W^p and b_s^a are the parameters to map the vector representation of target t to the polarity label of aspect a . To avoid overfitting, we add a dropout layer with dropout probability of 0.5 after the embedding layer. We stop the training process of our model after 10 epochs and select the model that achieves the best performance on the development set.

6.3 Experiments

6.3.1 Dataset and Resources

We evaluate our method on two datasets: SentiHood [46] and a subset of Semeval 2015 [128]. SentiHood is built by querying Yahoo! Answers with location names of London city. Table 6.2 shows statistics of SentiHood. The whole dataset is split into train, test, and development set by the authors. Overall, the entire dataset contains 5,215 sentences, with 3,862 sentences containing a single target and 1,353 sentences containing multiple targets. It also shows that there are approximately two third of targets annotated with aspect-based sentiment polarity (train set: 2476 out of 2977; test set: 1241 out of 1898; development set: 619 out of 955). On average, each sentiment-bearing target has been annotated with 1.37 aspects. To show the generalizability of our methods, we build a subset of the dataset used by Semeval-2015. We remove sentences containing no targets as well as *NULL* targets. To be comparable with SentiHood, we combine targets with the same surface form within

the same sentence as mentions of the same target. In total, we have 1,197 targets left in the training set and 542 targets left in the testing set. On average, each target has 1.06 aspects.

To inject the commonsense knowledge, we use a syntax-based concept parser [github.com/senticnet](https://github.com/senticnet/senticnet) to extract a set of concept candidates at each time step, and use AffectiveSpaces sentic.net/downloads as the concept embeddings. In case no concepts are extracted, a zero vector is used as the concept input.

6.3.2 Experiment Setting

We evaluate our method on two sub-tasks of targeted ABSA: 1) aspect categorization and 2) aspect-based sentiment classification. Following Saeidi et al. [46], we treat the outputs of aspect-based classification as hierarchical classes. For aspect categorization, we output the label (e.g., in the 3-class setting, it outputs ‘Positive’, ‘Negative’, or ‘None’) with the highest probability for each aspect. For aspect-based sentiment classification, we ignore the scores of ‘None’. For evaluating the aspect-based sentiment classification, we simply calculate the accuracy averaged over aspects. We evaluate aspect categorization as a multi-label classification problem so that results are averaged over targets instead of aspects.

We evaluate our methods and baseline systems using both loose and strict metrics. We report scores of three widely used evaluation metrics of multi-label classifier: Macro-F1, Micro-F1, and strict Accuracy. Given the dataset D , the ground-truth aspect categories of the target $t \in D$ is denoted as Y_t , while the predicted aspect categories denoted as $\hat{Y}_t$. The three metrics can be computed as

- Strict accuracy (Strict Acc.): $\frac{1}{D} \sum_{t \in D} \sigma(Y_t = \hat{Y}_t)$, where $\sigma(\cdot)$ is an indicator function.

Table 6.2: SentiHood dataset

	Train	Dev	Test
Targets	3,806	955	1,898
Targets w/ Sentiment	2,476	619	1,241
Aspects per Target(w/ Sentiment)	1.37	1.37	1.35

Table 6.3: System performance on SentiHood dataset

	Sentiment	
	Sentiment Acc. (%)	
	dev	test
TDLSTM	82.60	81.82
LSTM + TA	83.80	84.29
LSTM + TA + SA	86.00	86.75
LSTM + TA + DMN SA	84.80	83.36
LSTM + TA + SA + KB Feat	87.00	88.70
LSTM + TA + SA + KBA	87.40	87.98
Recall LSTM + TA + SA	86.80	86.85
Sentic LSTM + TA + SA	88.80	89.32

	Aspect Categorization					
	Strict Acc. (%)		Macro F1 (%)		Micro F1 (%)	
	dev	test	dev	test	dev	test
TDLSTM	50.27	50.83	59.03	58.17	55.72	55.78
LSTM + TA	54.17	52.02	62.90	61.07	60.56	59.02
LSTM + TA + SA	68.83	66.42	79.36	76.69	79.14	76.64
LSTM + TA + DMN SA	60.66	60.14	68.89	70.19	67.28	68.37
LSTM + TA + SA + KB Feat	69.38	64.76	80.00	76.33	79.79	76.08
LSTM + TA + SA + KBA	68.08	65.12	78.68	76.40	78.73	76.46
Recall LSTM + TA + SA	68.64	64.66	78.44	75.61	78.53	75.91
Sentic LSTM + TA + SA	69.20	67.43	78.84	78.18	79.09	77.66

- Macro-F1 = $2 \frac{\text{Ma-P} \times \text{Ma-R}}{\text{Ma-P} + \text{Ma-R}}$, which is based on Macro-Precision (Ma-P) and Micro-Recall (Ma-R) with $\text{Ma-P} = \frac{1}{|D|} \sum_{t \in D} \frac{|Y_t \cap \hat{Y}_t|}{\hat{Y}_t}$, and $\text{Ma-R} = \frac{1}{|D|} \sum_{t \in D} \frac{|Y_t \cap \hat{Y}_t|}{Y_t}$.
- Micro-F1 = $2 \frac{\text{Mi-P} \times \text{Mi-R}}{\text{Mi-P} + \text{Mi-R}}$, which is based on Micro-Precision (Mi-P) and Micro-Recall (Mi-R), where $\text{Mi-P} = \frac{\sum_{t \in D} |Y_t \cap \hat{Y}_t|}{\sum_{t \in D} \hat{Y}_t}$, and $\text{Mi-R} = \frac{\sum_{t \in D} |Y_t \cap \hat{Y}_t|}{\sum_{t \in D} Y_t}$.

6.3.3 Performance Comparison

We compare our proposed method with the methods that have been proposed for targeted ABSA as well as methods proposed for ABSA or targeted sentiment analysis but applicable to targeted ABSA.

Furthermore, we also compare the performances of several variants of our proposed method in order to highlight our technical contribution. We run the model for multiple times and report the results that perform best in the development set. For

Figure 6.3: A sketch of SenticNet semantic network [10]

Table 6.4: System performance on Semeval-2015 dataset

	Aspect Categorization			Sentiment
	Strict Acc.	Macro F1	Micro F1	Sentiment Acc.
TDLSTM	65.49	70.56	69.00	68.57
LSTM+TA	66.42	71.71	70.06	69.24
LSTM+TA+SA	63.46	70.73	66.18	74.28
LSTM+TA+DMN SA	48.33	52.73	51.39	69.07
LSTM+TA+SA+KB Feat	65.68	74.46	70.71	76.13
LSTM+TA+SA+KBA	67.34	74.36	71.78	73.10
Recall LSTM + TA + SA	66.05	72.90	69.66	74.11
Sentic LSTM + TA + SA	67.34	76.44	73.82	76.47

Semeval-2015 dataset, we report the results of the final epoch.

- **TDLSTM**: TDLSTM [42] adopts Bi-LSTM to encode the sequential structure of a sentence and represents a given target using a vector averaged on the hidden outputs of target instances.
- **LSTM + TA**: Our method learns an instance attention on top of the outputs of LSTM to model the contribution of each instance.
- **LSTM + TA + SA**: In addition to target instance attention, we add a

sentence-level attention to the model.

- **LSTM + TA + DMN SA:** The sentence-level attention is replaced by a dynamic memory network with multiple hops [43]. We run the memory network with different numbers of hops and report the result with 4 hops (best performance on development set of SentiHood). We exclude the case of zero hops as it corresponds to Bi-LSTM + TA + SA.
- **LSTM + TA + SA + KB Feat:** Concepts are fed into the input layer as additional features.
- **LSTM + TA + SA + KBA:** This is an integration of the method proposed by [51], which learns an attention over the concept embeddings (the embeddings are combined with the hidden output before being fed into the classifier).
- **Recall LSTM + TA + SA:** LSTM is extended with a recall knowledge gate as in [52].
- **Sentic LSTM + TA + SA:** The encoder is replaced with the proposed knowledge-embedded LSTM.

The word embedding of the input layer is initialized by a pre-trained skip-gram model [25] with 150 hidden units on a combination of Yelp¹ and Amazon review dataset [129] and 50 hidden units for the bi-directional LSTM.

6.3.4 Results of Attention Model

Table 6.3 and Table 6.4 show the performance on SentiHood and Semeval-15 dataset, respectively. In comparison with the non-attention baseline (Bi-LSTM+Avg.), we can find that our best attention-based model significantly improves aspect categorization (by more than 20%) and sentiment classification (approximately 10%) on SentiHood.

¹yelp.com.sg/dataset/challenge

However, it is notable that, on the Semeval-2015 dataset, the improvement is relatively smaller. We conjecture the reason is that SentiHood has masked the target as a special word “LOCATION”, which resulted less informative than the full name of aspect targets that are used by Semeval-2015.

Hence, using only the hidden outputs regarding the target does not suffice to represent the sentiment of the whole sentence in SentiHood. Compared with target averaging model, the target-level attention achieves some improvement (even though not significant), as the target attention is capable of identifying the part of target expressions with higher sentiment salience. On the other hand, it is notable that the two-step attention achieves significant improvement on both aspect categorization and sentiment classification, indicating that the target- and aspect-dependent sentence attention could retrieve information relevant to both tasks.

To our surprise, using multiple hops in the sentence-level attention fails to bring in any improvement. The performance even falls down significantly on Semeval-2015 with a much smaller number of training instances but larger aspect set than SentiHood. We conjecture the reason is that using multi-hops increases the number of parameter to learn, which makes it less applicable to small and sparse datasets.

6.3.5 Visualization of Attention

We visualize the attention vectors of sentence-level attention in Figure 6.4 with respect to “Transition-location” and “Price” aspects. The two attention vectors have encoded quite different concerns in the word sequence.

In the first example, the ‘Transition-location’ attention attends to the word “long”, which is expressing a negative sentiment towards the target. In comparison, the ‘Price’ attention attends more to the word ‘cheap’, which is related to the aspect. That is to say, the two attention vectors are capable of distinguishing information related to different aspects. As visualized in Figure 6.5, the target-level attention is capable of selecting the part of target expression of which the aspect or sentiment is easier to be resolved.

6.3.6 Results of Knowledge-Embedded LSTM

It can be seen from Table 6.3 and 6.4 that injecting knowledge into the model improves the performance in general. Since AffectiveSpace encodes the information about affective properties that are semantically related to the aspects, it is reasonable to find out that it can improve performance on both tasks. The results show that our proposed Sentic LSTM outperforms baseline methods, even if not significantly.

An important outcome of the experiments is that Sentic LSTM significantly outperforms a baseline (LSTM + TA + SA + KB feat) feeding the knowledge features to the input layer, which confirms the efficacy of using a knowledge output gate to control the flow of background knowledge. Furthermore, the superior performance of Sentic LSTM over Recall LSTM and KBA indicates that the activated knowledge concepts can also help filtering the information that conflicts with the background knowledge.

6.3.7 AffectiveSpace v.s. SentiWordNet

At last, we compare systems using different sentiment knowledge bases. SentiWordNet is a lexicon-based sentiment knowledge base consisting of word sense synsets annotated with sentiment polarity. However, it is notable that SentiWordNet contains

Figure 6.4: Example of sentence-level attention

Figure 6.5: Example of target-level attention

Table 6.5: Comparison of systems using AffectiveSpace and SentiWordNet (SH stands for Sentihood, and SE stands for SemEval-15)

	Aspect Categorization						Sentiment	
	Strict Acc. (%)		Macro F1 (%)		Micro F1 (%)		Sentiment Acc. (%)	
	SH	SE	SH	SE	SH	SE	SH	SE
AffectiveSpace	67.71	70.47	78.05	77.90	78.19	75.12	89.63	78.31
SentiWordNet	60.23	66.42	70.03	73.60	69.48	71.18	85.01	75.79

neither commonsense concepts nor affective properties, which are the key features of AffectiveSpace. Consequently, we have to use randomly initialized embeddings for representing the SentiWordnet synsets. Word synsets are mapped to the same 100 dimension embeddings as AffectiveSpace does. Each word in the sentence is mapped to its word sense with the help of a Word Sense Disambiguation tool. We deliberately remove the neutral synsets (i.e., those having zero values for both Positive and Negative) to emphasize on the sentiment-bearing words. Table 6.5 shows the comparison of the two knowledge base. We report the results using our best-performed system (Sentic LSTM + TA + SA). It shows that using the AffectiveSpace achieves superior performance than using SentiWordNet. We conjecture the reason is that the word sense synsets is not as informative as affective properties. Moreover, probably because the link between word senses and aspects are not straightforward, we find the gap in sentiment classification is smaller than aspect categorization.

6.4 Summary

In this chapter, we propose a neural architecture for the task of targeted ABSA. We explicitly modeled the attention as a two-step model which encodes targets and full

sentence. The target-level attention learns to attend to the sentiment-salient part of a target expression and generates a more accurate representation of the target, while the sentence-level attention searches for the target- and aspect-dependent evidence over the full sentence. Moreover, we propose an extension of the LSTM cell so that it could more effectively incorporate affective commonsense knowledge when encoding the sequence into a vector. In the future, we would like to collectively analyze the sentiment of multiple targets co-occurring in the same sentence and investigate the role of commonsense knowledge in modeling the relation between targets.

Chapter 7

Summary and Future Work

In this final chapter, we summarize our contribution to the problem of leveraging concept-level information for natural language processing, followed by possible directions of research that we would like to explore in the future.

7.1 Summary of Proposed Methods

Concepts are playing a critical role in natural languages, especially in the context of opinion understandings. In this thesis, we address the problem of leveraging concept-level information for concept-centered natural language processing tasks that are core modules in a pipeline opinion understanding system with a particular focus on fusing information from multiple levels. We identify our main contribution as follows:

- In Chapter 3, we demonstrate that incorporating entity-based features into the learning process of word embedding could better encode entity-related information in the word embedding space. We present a novel method to learn word embeddings in a multi-task setting, where entity-level information is fused with the word-level by setting the training objective to predicting the entity-related features. Each entity-related feature type is modeled as a separate subtask of prediction, while the word embedding layer is shared by all subtasks. We then demonstrate the effectiveness of the resulting word embedding by feeding word embeddings as features into a CRF-based NER tagger.
- In Chapter 4, we capture the semantic correlation between fine-grained entity types. We propose a process to represent the entity labels with their prototypes.

The process can be implemented via either manual selection or an automated algorithm based on PMI. At the core of our proposed method are word embeddings learned from a large corpus. The embedding of each label is then defined as the average word embedding of its prototypes. In doing so, we are able to project the entity labels into the word embedding space that has encoded distributional semantics. We validate the proposed label embedding with two different settings of fine-grained named entity typing: 1) few-shots learning, where all labels are assumed having presented in the training data; and 2) zero-shot learning, where the second level entity types are assumed absent from the training set. Our experiment shows that the resulting label embedding outperforms existing methods in both few-shots and zero-shots learning setting.

- In Chapter 5, we study the problem of using entity-type information to help to rank ASR hypotheses. We propose an RBM based ASR reranking model that can be trained using discriminative objective function. The hidden layer of RBM are in fact binary embeddings of the input. We design a regularization term to softly constrain the binary embeddings such that it can encode a better knowledge of entity types. The resulting binary embedding could also be deemed as a fusion of entity-related prior information and word-level input. Our experiment shows that the proposed ASR reranking model can outperform the popularly used single-perceptron baseline thanks to a better knowledge of natural language concepts.
- In Chapter 6, we address the problem of targeted sentiment analysis which is an important subtask of opinion understanding. Our solution leveraged common-sense concepts via an extended LSTM network. Our contribution also includes a two-layer attention neural architecture that jointly learns the representation of the targeted entity as well as a target-aware representation of the whole sentence. More specifically, our framework explicitly fuses the concept-level and word-level information via a separate output gate of LSTM cell. The experiment shows that the knowledge of AffectiveSpace provides complementary semantics

to assist resolving both aspect categories as well as the corresponding sentiment polarity.

We further summarize the results of our work into Table 7.1.

Table 7.1: Summarization of our proposed work

Proposed Methods	Major Contribution	Results
Feature Word Embedding	Encode information from multiple levels	Better representation of entity-related words; improve NER
Label Embedding of Entities	Encode the concept-level dependency using prototypes; combine the force of word-level and concept-level information.	Improve FNET at both few-shots and zero-shot setting.
RBM-based reranking model	Incorporate the concept-level prior knowledge into the binary embedding layer of RBM.	Improve ASR reranking algorithm using concept-level prior.
Sentic-LSTM	Utilize commonsense knowledge for targeted sentiment analysis; propose an effective extension of LSTM to incorporate the concept-level and word-level inputs.	Improve both aspect categorization and polarity detection of targeted sentiment analysis

At last, the results presented in this thesis suggests that the concept-level knowledge is generally helpful as long as concepts (including named entities) are playing a key roles in the task. On the other hand, as concepts are typically of high dimensionality, the embedding layer together with efficient computational structure could make the concept-level information more accessible.

7.2 Future Work

We would discuss possible directions we like to explore in the future.

7.2.1 Limitation of Our Work

One limit of our solution to leveraging concepts as well as fusing concept-level and word-level information is that our usage of concepts is constrained to the original form

of input data. In other words, the use of concepts does not change the form of inputs that are mostly fed into the model as a sequence of words. However, when considering the interaction between concepts, keeping to the original sequential order of words might become redundant or sometimes might even cause unnecessary confusion. For example, for a sentence 'this place is safe, but areas close to it are very dangerous', if the target is to analyze the sentiment towards the first place, the second clause might be cause confusion if being processed in its original order. In contrast, parsing the sentence into concepts and relations is able to shorten the sentence into a simple tuple (place,be,safe) which could not only help remove the redundant information but also create the dependency between a target and a correct sentiment-related modifier. Inspired by previous work reforming the sequential input into structures such as trees [130] or graphs [131], we would like to investigate exploiting non-sequential structure to better perform concept-level analysis of natural language.

7.2.2 Graph-based Concept Embedding

Concepts are naturally organized in graphs. For example, ConceptNet defines a set of relations such as Isa relation between concepts. A graph can be built with concept being the nodes and relations being the edges. Similarly, AffectiveSpace can also be converted into a directed graph, even though it uses affective properties (relation + argument). So far, the concept embeddings [10] are built by merely considering decomposition of the concept-property matrix. In other words, the result embedding of concepts are only affected by immediate neighbors.

Recent work on embedding graphs [132] has shed light upon propagating information over a graph so that the embedding of a node is able to also reflect the properties of its indirect neighbors. We could try to infer concept embeddings via the graph structure. Different from existing work, our method will also consider the relation types which are also important for modeling information propagation in AffectiveSpace.

7.2.3 Concept-oriented Graph Model

Since concepts are very informative and their interaction is key to properly understanding the meaning of natural language, our first proposal for future work is to extend the usage of concepts to reforming the input sequence. One way to do so is to transform the input to local graphs of concepts by considering their explicit or implicit relations. Explicit relations can be the semantic relation given by knowledge bases such as ConceptNet or SenticNet. Implicit relations, for example, can be the presence of concepts in each other's contexts. By converting the word sequence into concept graphs, the NLP model can explicitly avoid affected by redundant and misleading structures and take into account the interaction between concepts.

In addition, another extension is to consider the global interaction of concepts. This can be fairly critical for tasks such as sentiment analysis of geographical entities [46]¹. For example, people might share the same safety concerns with neighboring locations. A graph built to represent the relation of concepts allows us to perform collective inference via approaches such as label propagation [133].

7.2.4 New Sentiment Analysis Application using Concepts

In future, we also would like to explore new sentiment-related applications using concepts. It is known that sentiment is playing an important role in expressing opinions. In addition to decoding (or interpreting) people's opinion, it is also of our interest to generate sentiment-aware responses. Especially, in the context of e-commerce, automatically generating sentiment-aware responses to customers' reviews would largely improve the customer experience. Moreover, rather than generating a sequence of words, we are more interested in generating a mix of words and commonsense concepts. We believe that, by leveraging commonsense concepts as well as concept relations, it could help improve the consistency and comprehensiveness of generated responses.

¹Similar to named entities, we could also consider geographical entities as a type of concepts.

Appendix A

Authors Publications

- **Yukun Ma**, Jung-jae Kim, Benjamin Bigot, and Tahir Muhammad Khan. “Feature-enriched word embeddings for named entity recognition in open-domain conversations.” In Acoustics, Speech and Signal Processing (ICASSP), 2016 IEEE International Conference on, pp. 6055-6059. IEEE, 2016.
- Khan, Muhammad Tahir, **Yukun Ma**, and Jung-jae Kim. “Term Ranker: A Graph-Based Re-Ranking Approach.” In FLAIRS Conference, pp. 310-315. 2016.
- **Yukun Ma**, Erik Cambria, and Sa Gao. “Label Embedding for Zero-shot Fine-grained Named Entity Typing.” In COLING, pp. 171-180. 2016..
- **Yukun Ma**, Erik Cambria, and Benjamin Bigot. “ASR Hypothesis Reranking using Prior-informed Restricted Boltzmann Machine.” In International Conference on Computational Linguistics and Intelligent Text Processing (CiCLing), 2017
- Sa Gao, Zhenchang Xing, **Yukun Ma**, Deheng Ye and Shang-Wei Lin, “Enhancing Knowledge Sharing in Stack Overflow via Automatic External Web Resources Linking”, International Conference on Engineering of Complex Computer Systems (ICECCS), Fukuoka, Japan, 2017.
- **Yukun Ma**, Haiyun Peng, and Erik Cambria. “Targeted Aspect-Based Sentiment Analysis via Embedding Commonsense Knowledge into an Attentive LSTM”, In the Thirty-Second AAAI Conference on Artificial Intelligence (AAAI-18), pp. 5876-5883, 2018

- Haiyun Peng, **Yukun Ma**, Erik Cambria, Yang Li, “Learning Multi-grained Aspect Target Sequence for Chinese Sentiment Analysis”, Vol. 1, pp.167-176, Knowledge-based Systems.

References

- [1] E. Cambria and B. White, “Jumping nlp curves: A review of natural language processing research,” *IEEE Computational intelligence magazine*, vol. 9, no. 2, pp. 48–57, 2014.
- [2] E. Cambria, S. Poria, A. Gelbukh, and M. Thelwall, “Sentiment analysis is a big suitcase,” *IEEE Intelligent Systems*, vol. 32, no. 6, pp. 74–80, 2017.
- [3] J. Lafferty, “Conditional random fields: Probabilistic models for segmenting and labeling sequence data,” in *ICML*, 2001, pp. 282–289.
- [4] G. Lample, M. Ballesteros, S. Subramanian, K. Kawakami, and C. Dyer, “Neural architectures for named entity recognition,” in *NAACL-HLT*, San , California, 2016, pp. 260–270.
- [5] D. Gillick, N. Lazic, K. Ganchev, J. Kirchner, and D. Huynh, “Context-dependent fine-grained entity type tagging,” *CoRR*, vol. abs/1412.1820, 2014.
- [6] D. Yogatama, D. Gillick, and N. Lazic, “Embedding methods for fine grained entity type classification,” in *ACL*, 2015, pp. 291–296.
- [7] S. Shimaoka, P. Stenetorp, K. Inui, and S. Riedel, “Neural architectures for fine-grained entity type classification,” in *EACL*, 2017, pp. 1271–1280.
- [8] Y. Ma, E. Cambria, and S. Gao, “Label embedding for zero-shot fine-grained named entity typing,” in *COLING*, Osaka, 2016, pp. 171–180.
- [9] Y. Ma, E. Cambria, and B. Bigot, “Asr hypothesis reranking using prior-informed restricted boltzmann machine,” in *CiCLing*, 2017.

- [10] E. Cambria, J. Fu, F. Bisio, and S. Poria, “AffectiveSpace 2: Enabling affective intuition for concept-level sentiment analysis,” in *AAAI*, Austin, 2015, pp. 508–514.
- [11] N. Chinchor, “Message understanding conference (muc) 7 ldc2001t02.” in *Philadelphia: Linguistic Data Consortium*, 2001.
- [12] Y. Ma, J.-j. Kim, B. Bigot, and T. M. Khan, “Feature-enriched word embeddings for named entity recognition in open-domain conversations,” in *ICASSP*. IEEE, 2016, pp. 6055–6059.
- [13] E. F. Tjong Kim Sang and F. De Meulder, “Introduction to the conll-2003 shared task: Language-independent named entity recognition,” in *HLT-NAACL*, 2003, pp. 142–147.
- [14] R. Speer and C. Havasi, “Conceptnet 5: A large semantic network for relational knowledge.” Springer Berlin Heidelberg, 2013, pp. 161–176.
- [15] E. Cambria, S. Poria, D. Hazarika, and K. Kwok, “SenticNet 5: Discovering conceptual primitives for sentiment analysis by means of context embeddings,” in *AAAI*, 2018, pp. 1795–1802.
- [16] X. Ling and D. S. Weld, “Fine-grained entity recognition,” in *AAAI*, 2012.
- [17] D. M. Bikel, R. Schwartz, and R. M. Weischedel, “An algorithm that learns what’s in a name,” *Mach. Learn.*, pp. 211–231, 1999.
- [18] E. Alfonseca and S. Manandhar, “An unsupervised method for general named entity recognition and automated concept discovery,” in *In: the 1 st International Conference on General WordNet*, 2002, pp. 34–43.
- [19] R. Evans, “A framework for named entity recognition in the open domain,” in *RANLP*, 2003, pp. 137–144.

- [20] E. Charniak, C. Hendrickson, N. Jacobson, and M. Perkowitz, “Equations for part-of-speech tagging,” in *the Eleventh National Conference on Artificial Intelligence*, 1993, pp. 784–789.
- [21] R. Collobert, J. Weston, L. Bottou, M. Karlen, K. Kavukcuoglu, and P. Kuksa, “Natural language processing (almost) from scratch,” *JMLR*, vol. 12, pp. 2493–2537, 2011.
- [22] S. Hochreiter and J. Schmidhuber, “Long short-term memory,” *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [23] X. Ma and E. Hovy, “End-to-end sequence labeling via bi-directional lstm-cnns-crf,” in *ACL*, Berlin, Germany, 2016, pp. 1064–1074.
- [24] M. Collins, “Ranking algorithms for named-entity extraction: Boosting and the voted perceptron,” in *ACL*, 2002, pp. 489–496.
- [25] T. Mikolov, I. Sutskever, K. Chen, G. S. Corrado, and J. Dean, “Distributed representations of words and phrases and their compositionality,” in *NIPS 2013*, 2013, pp. 3111–3119.
- [26] M. A. Yosef, S. Bauer, J. Hoffart, M. Spaniol, and G. Weikum, “HYENA: Hierarchical type classification for entity names,” in *COLING 2012: Posters*, 2012, pp. 1361–1370.
- [27] A. Abhishek, A. Anand, and A. Awekar, “Fine-grained entity type classification by jointly learning representations and label embeddings,” in *EACL*, Valencia, Spain, 2017, pp. 797–807.
- [28] F. M. Suchanek, G. Kasneci, and G. Weikum, “Yago: A core of semantic knowledge,” in *WWW*, 2007, pp. 697–706.
- [29] J. Hoffart, F. M. Suchanek, K. Berberich, and G. Weikum, “Yago2: A spatially and temporally enhanced knowledge base from wikipedia,” *Artificial Intelligence*, vol. 194, pp. 28–61, 2013.

- [30] C. N. Silla Jr and A. A. Freitas, “A survey of hierarchical classification across different application domains,” *Data Mining and Knowledge Discovery*, vol. 22, pp. 31–72, 2011.
- [31] F. Rosenblatt, “The perceptron: A probabilistic model for information storage and organization in the brain,” *Psychological Review*, vol. 65, pp. 386–408, 1958.
- [32] J. Weston, S. Bengio, and N. Usunier, “Wsabie: Scaling up to large vocabulary image annotation,” in *IJCAI*, 2011, pp. 2764–2770.
- [33] X. Ren, W. He, M. Qu, C. R. Voss, H. Ji, and J. Han, “Label noise reduction in entity typing by heterogeneous partial-label embedding,” *arXiv preprint arXiv:1602.05307*, 2016.
- [34] C. W. Xiang Ren, Ahmed El-Kishky, “Clustype: Effective entity recognition and typing by relation phrase-based clustering,” in *SIGKDD’15*, 2015, pp. 995–1004.
- [35] S. Gopal, Y. Yang, B. Bai, and A. Niculescu-mizil, “Bayesian models for large-scale hierarchical classification,” in *NIPS*. Curran Associates, Inc., 2012, pp. 2411–2419.
- [36] J. Wagner, P. Arora, S. Cortes, U. Barman, D. Bogdanova, J. Foster, and L. Tounsi, “Dcu: Aspect-based polarity classification for semeval task 4,” in *SemEval*, Dublin, Ireland, 2014, pp. 223–229.
- [37] S. Kiritchenko, X. Zhu, C. Cherry, and S. Mohammad, “Nrc-canada-2014: Detecting aspects and sentiment in customer reviews,” in *SemEval*, Dublin, Ireland, 2014, pp. 437–442.
- [38] L. Dong, F. Wei, C. Tan, D. Tang, M. Zhou, and K. Xu, “Adaptive recursive neural network for target-dependent twitter sentiment classification,” in *the 52nd ACL*, 2014, pp. 49–54.

- [39] H. Lakkaraju, R. Socher, and C. Manning, “Aspect specific sentiment analysis using hierarchical deep learning,” in *NIPS Workshop on Deep Learning and Representation Learning*, 2014.
- [40] T. H. Nguyen and K. Shirai, “Phrasernn: Phrase recursive neural network for aspect-based sentiment analysis,” in *EMNLP 2015*, Lisbon, Portugal, 2015, pp. 2509–2514.
- [41] Y. Wang, M. Huang, x. zhu, and L. Zhao, “Attention-based lstm for aspect-level sentiment classification,” in *the EMNLP 2016*, Austin, Texas, 2016, pp. 606–615.
- [42] D. Tang, B. Qin, X. Feng, and T. Liu, “Effective lstms for target-dependent sentiment classification,” in *COLING 2016*, 2016, pp. 3298–3307.
- [43] D. Tang, B. Qin, and T. Liu, “Aspect level sentiment classification with deep memory network,” in *the EMNLP 2016*, 2016, pp. 214–224.
- [44] P. Chen, Z. Sun, L. Bing, and W. Yang, “Recurrent attention network on memory for aspect sentiment analysis,” in *EMNLP 2017*, 2017, pp. 463–472.
- [45] B. Wang, M. Liakata, A. Zubiaga, and R. Procter, “Tdpars: Multi-target-specific sentiment recognition on twitter,” in *the 15th EACL*, Valencia, Spain, 2017, pp. 483–493.
- [46] M. Saeidi, G. Bouchard, M. Liakata, and S. Riedel, “Sentihood: Targeted aspect based sentiment analysis dataset for urban neighbourhoods,” in *COLING 2016*, 2016, pp. 1546–1556.
- [47] L. Ratnov and D. Roth, “Design challenges and misconceptions in named entity recognition,” in *the 13th CoNLL*, 2009, pp. 147–155.
- [48] A. Rahman and V. Ng, “Coreference resolution with world knowledge,” in *the 49th ACL*, 2011, pp. 814–824.

- [49] N. Nakashole and T. M. Mitchell, “A knowledge-intensive model for prepositional phrase attachment.” in *ACL (1)*, 2015, pp. 365–375.
- [50] S. Ahn, H. Choi, T. Pärnamaa, and Y. Bengio, “A neural knowledge language model,” *arXiv preprint arXiv:1608.00318*, 2016.
- [51] B. Yang and T. Mitchell, “Leveraging knowledge bases in lstms for improving machine reading,” in *the 55th ACL*, 2017, pp. 1436–1446.
- [52] Z. Xu, B. Liu, B. Wang, C. Sun, and X. Wang, “Incorporating loose-structured knowledge into lstm with recall gate for conversation modeling,” *arXiv preprint*, 2016.
- [53] M. Hatmi, C. Jacquin, E. Morin, and S. Meignier, “Incorporating named entity recognition into the speech transcription process.” in *INTERSPEECH*, 2013, pp. 3732–3736.
- [54] C. Chelba and F. Jelinek, “Exploiting syntactic structure for language modeling,” in *ACL*, 1998, pp. 225–231.
- [55] M. Collins, B. Roark, and M. Saraclar, “Discriminative syntactic language modeling for speech recognition,” in *ACL*, 2005, pp. 507–514.
- [56] B. Lambert, B. Raj, and R. Singh, “Discriminatively trained dependency language modeling for conversational speech recognition.” in *INTERSPEECH*, 2013, pp. 3414–3418.
- [57] M. Bansal, K. Gimpel, and K. Livescu, “Tailoring continuous word representations for dependency parsing,” in *ACL*, 2014, pp. 809–815.
- [58] K. Hashimoto, P. Stenetorp, M. Miwa, and Y. Tsuruoka, “Task-oriented learning of word embeddings for semantic relation classification,” *CoRR*, 2015.
- [59] M. Ono, M. Miwa, and Y. Sasaki, “Word embedding-based antonym detection using thesauri and distributional information,” in *HLT-NAACL*, 2015, pp. 984–989.

- [60] Y. Li, Q. Pan, T. Yang, S. Wang, J. Tang, and E. Cambria, “Learning word representations for sentiment analysis,” *Cognitive Computation*, vol. 9, no. 6, pp. 843–851, 2017.
- [61] T. Anastasakos and A. Deoras, “Task specific continuous word representations for mono and multi-lingual spoken language understanding,” in *ICASSP*, 2014, pp. 3246–3250.
- [62] M. Baroni, G. Dinu, and G. Kruszewski, “Don’t count, predict! a systematic comparison of context-counting vs. context-predicting semantic vectors.” in *ACL*, 2014, pp. 238–247.
- [63] O. Levy, Y. Goldberg, and I. Dagan, “Improving distributional similarity with lessons learned from word embeddings,” *Transactions of the Association for Computational Linguistics*, vol. 3, pp. 211–225, 2015.
- [64] S. Deerwester, S. T. Dumais, G. W. Furnas, T. K. Landauer, and R. Harshman, “Indexing by latent semantic analysis,” *Journal of the American society for information science*, vol. 41, no. 6, p. 391, 1990.
- [65] K. W. Church and P. Hanks, “Word association norms, mutual information, and lexicography,” *Computational linguistics*, vol. 16, no. 1, pp. 22–29, 1990.
- [66] D. Balduzzi, “Randomized co-training: from cortical neurons to machine learning and back again,” *arXiv preprint arXiv:1310.6536*, 2013.
- [67] W. B. Johnson, J. Lindenstrauss, and G. Schechtman, “Extensions of lipschitz maps into banach spaces,” *Israel Journal of Mathematics*, vol. 54, no. 2, pp. 129–138, 1986.
- [68] Y. Lu, P. Dhillon, D. P. Foster, and L. Ungar, “Faster ridge regression via the subsampled randomized hadamard transform,” in *NIPS*, 2013, pp. 369–377.
- [69] E. H. Huang, R. Socher, C. D. Manning, and A. Y. Ng, “Improving word representations via global context and multiple word prototypes,” in *ACL*, 2012, pp. 873–882.

- [70] T. Mikolov, I. Sutskever, K. Chen, G. S. Corrado, and J. Dean, “Distributed representations of words and phrases and their compositionality,” in *NIPS’13*, 2013, pp. 3111–3119.
- [71] D. Nadeau and S. Sekine, “A survey of named entity recognition and classification,” *Linguisticae Investigationes*, pp. 3–26, 2007.
- [72] C. Li, A. Sun, J. Weng, and Q. He, “Tweet segmentation and its application to named entity recognition,” *IEEE Transactions on Knowledge and Data Engineering*, vol. 27, no. 2, pp. 558–570, 2015.
- [73] M. Surdeanu, J. Turmo, and E. Comelles, “Named entity recognition from spontaneous Open-Domain speech,” in *Interspeech*, 2005, pp. 3433–3436.
- [74] F. Bechet and E. Charton, “Unsupervised knowledge acquisition for extracting named entities from speech.” in *ICASSP*, 2010, pp. 5338–5341.
- [75] R. Collobert and J. Weston, “A unified architecture for natural language processing: Deep neural networks with multitask learning,” in *ICML’08*, 2008, pp. 160–167.
- [76] J. Turian, L. Ratinov, and Y. Bengio, “Word representations: A simple and general method for semi-supervised learning,” in *ACL*, 2010, pp. 384–394.
- [77] M. Yu, T. Zhao, D. Dong, H. Tian, and D. Yu, “Compound embedding features for semi-supervised learning,” in *HLT-NAACL*, 2013, pp. 563–568.
- [78] J. Guo, W. Che, H. Wang, and T. Liu, “Revisiting embedding features for simple semi-supervised learning,” in *Proc. of EMNLP*, 2014, pp. 110–120.
- [79] J. Pennington, R. Socher, and C. Manning, “Glove: Global vectors for word representation,” in *EMNLP*, Doha, Qatar, 2014, pp. 1532–1543.
- [80] A. Passos, V. Kumar, and A. McCallum, “Lexicon infused phrase embeddings for named entity resolution,” in *CoNLL*, 2014, pp. 78–86.

- [81] Y. Qi, P. P. Kuksa, R. Collobert, K. Sadamasa, K. Kavukcuoglu, and J. Weston, “Semi-supervised sequence labeling with self-learned features.” in *ICDM*, 2009, pp. 428–437.
- [82] A. Ferraresi, E. Zanchetta, M. Baroni, and S. Bernardini, “Introducing and evaluating ukwac, a very large web-derived corpus of english,” in *WAC-4*, 2008.
- [83] J. Godfrey and H. Edward, “Switchboard-1 release 2 LDC97S62,” 1993.
- [84] X. Dong, E. Gabrilovich, G. Heitz, W. Horn, N. Lao, K. Murphy, T. Strohmman, S. Sun, and W. Zhang, “Knowledge vault: A web-scale approach to probabilistic knowledge fusion,” in *SIGKDD*, 2014, pp. 601–610.
- [85] T. Lin, O. Etzioni *et al.*, “No noun phrase left behind: detecting and typing unlinkable entities,” in *EMNLP*, 2012, pp. 893–903.
- [86] A. Fader, L. Zettlemoyer, and O. Etzioni, “Open question answering over curated and extracted knowledge bases,” in *SIGKDD*, 2014, pp. 1156–1165.
- [87] G. R. Doddington, A. Mitchell, M. A. Przybocki, L. A. Ramshaw, S. Strassel, and R. M. Weischedel, “The automatic content extraction (ace) program-tasks, data, and evaluation.” in *LREC*, 2004, p. 1.
- [88] D. Gillick, N. Lazic, K. Ganchev, J. Kirchner, and D. Huynh, “Context-dependent fine-grained entity type tagging,” *arXiv preprint arXiv:1412.1820*, 2014.
- [89] A. Frome, G. S. Corrado, J. Shlens, S. Bengio, J. Dean, T. Mikolov *et al.*, “Devise: A deep visual-semantic embedding model,” in *NIPS*, 2013, pp. 2121–2129.
- [90] Z. Akata, F. Perronnin, Z. Harchaoui, and C. Schmid, “Label-embedding for attribute-based classification.” in *CVPR*, 2013, pp. 819–826.
- [91] H. Larochelle, D. Erhan, and Y. Bengio, “Zero-data learning of new tasks.” in *AAAI*, 2008, p. 3.

- [92] K. Q. Weinberger and O. Chapelle, “Large margin taxonomy embedding for document categorization,” in *NIPS*, 2009, pp. 1737–1744.
- [93] Z. Akata, S. Reed, D. Walter, H. Lee, and B. Schiele, “Evaluation of output embeddings for fine-grained image classification,” in *CVPR*, 2015, pp. 2927–2936.
- [94] A. Haghighi and D. Klein, “Prototype-driven learning for sequence models,” in *HLT-NAACL*, 2006, pp. 320–327.
- [95] R. Weischedel and A. Brunstein, “BBN pronoun coreference and entity type corpus,” Linguistic Data Consortium, Philadelphia, Tech. Rep., 2005.
- [96] R. Weischedel, M. Palmer, M. Marcus, E. Hovy, S. Pradhan, L. Ramshaw, N. Xue, A. Taylor, J. Kaufman, M. Franchini *et al.*, “Ontonotes release 4.0,” *LDC2011T03, Philadelphia, Penn.: Linguistic Data Consortium*, 2011.
- [97] L. v. d. Maaten and G. Hinton, “Visualizing data using t-sne,” *Journal of Machine Learning Research*, vol. 9, no. Nov, pp. 2579–2605, 2008.
- [98] M. Collins and T. Koo, “Discriminative reranking for natural language parsing,” *Computational Linguistics*, vol. 31, no. 1, pp. 25–70, 2005.
- [99] T. Koo and M. Collins, “Hidden-variable models for discriminative reranking,” in *HLT-EMNLP*, Vancouver, British Columbia, Canada, 2005, pp. 507–514.
- [100] Z. Li and S. Khudanpur, “Large-scale discriminative n-gram language models for statistical machine translation,” in *AMTA*, 2009, pp. 133–142.
- [101] B. Roark, M. Saraclar, M. Collins, and M. Johnson, “Discriminative language modeling with conditional random fields and the perceptron algorithm,” in *ACL*, Barcelona, Spain, 2004, pp. 47–54.
- [102] P. Smolensky, “Parallel distributed processing: Explorations in the microstructure of cognition, vol. 1,” *Information Processing in Dynamical Systems: Foundations of Harmony Theory*, pp. 194–281, 1986.

- [103] J. Niehues and A. Waibel, “Continuous space language models using restricted boltzmann machines.” in *IWSLT*, 2012, pp. 164–170.
- [104] G. E. Dahl, R. P. Adams, and H. Larochelle, “Training restricted boltzmann machines on word observations,” *arXiv preprint arXiv:1202.5695*, 2012.
- [105] L. Wang, K. Liu, Z. Cao, J. Zhao, and G. de Melo, “Sentiment-aspect extraction based on restricted boltzmann machines,” in *ACL*, 2015, pp. 616–625.
- [106] C. C. Fries, in *The Structure of English*. New York: Harcourt Brace, 1952.
- [107] F. Pernkopf, M. Wohlmayr, and S. Tschitschek, “Maximum margin bayesian network classifiers,” *IEEE Transaction Pattern Analysis and Machine Intelligence*, vol. 34, no. 3, pp. 521–532, 2012.
- [108] G. E. Hinton, “Training products of experts by minimizing contrastive divergence,” *Neural Computation*, vol. 14, no. 8, pp. 1771–1800, 2002.
- [109] S. C. Michael Levit, Sarangarajan Parthasarathy, “Word-phrase-entity language models: Getting more mileage out of n-grams,” in *nterspeech*, 2014, pp. 666–670.
- [110] R. Salakhutdinov and G. E. Hinton, “Replicated softmax: an undirected topic model.” in *NIPS*, vol. 22, 2009, pp. 1607–1614.
- [111] A. Rousseau, P. Deléglise, and Y. Estève, “Enhancing the ted-lium corpus with selected data for language modeling and more ted talks,” in *LREC*, 2014, pp. 3935–3939.
- [112] D. Povey, A. Ghoshal, G. Boulianne, L. Burget, O. Glembek, N. Goel, M. Hannemann, P. Motlicek, Y. Qian, P. Schwarz, J. Silovsky, G. Stemmer, and K. Vesely, “The kaldi speech recognition toolkit,” in *ASRU*, 2011.
- [113] W. Walker, P. Lamere, P. Kwok, B. Raj, R. Singh, E. Gouvea, P. Wolf, and J. Woelfel, “Sphinx-4: A flexible open source framework for speech recognition,” Mountain View, CA, USA, Tech. Rep., 2004.

- [114] J. R. Finkel, T. Grenager, and C. Manning, “Incorporating non-local information into information extraction systems by gibbs sampling,” in *ACL*, 2005, pp. 363–370.
- [115] C. Fellbaum, *WordNet: An Electronic Lexical Database (Language, Speech, and Communication)*. The MIT Press, 1998.
- [116] E. Cambria, D. Das, S. Bandyopadhyay, and A. Feraco, *A Practical Guide to Sentiment Analysis*. Cham, Switzerland: Springer, 2017.
- [117] S. Poria, E. Cambria, R. Bajpai, and A. Hussain, “A review of affective computing: From unimodal analysis to multimodal fusion,” *Information Fusion*, vol. 37, pp. 98–125, 2017.
- [118] I. Chaturvedi, E. Ragusa, P. Gastaldo, R. Zunino, and E. Cambria, “Bayesian network based extreme learning machine for subjectivity detection,” *Journal of The Franklin Institute*, pp. 1780 – 1797, 2017.
- [119] S. L. Lo, E. Cambria, R. Chiong, and D. Cornforth, “Multilingual sentiment analysis: From formal to informal and scarce resource languages,” *Artificial Intelligence Review*, vol. 48, no. 4, pp. 499–527, 2017.
- [120] Y. Xia, E. Cambria, A. Hussain, and H. Zhao, “Word polarity disambiguation using bayesian model and opinion-level features,” *Cognitive Computation*, vol. 7, no. 3, pp. 369–380, 2015.
- [121] N. Majumder, S. Poria, A. Gelbukh, and E. Cambria, “Deep learning-based document modeling for personality detection from text,” *IEEE Intelligent Systems*, vol. 32, no. 2, pp. 74–79, 2017.
- [122] S. Poria, E. Cambria, D. Hazarika, and P. Vij, “A deeper look into sarcastic tweets using deep convolutional neural networks,” in *COLING*, 2016, pp. 1601–1612.

- [123] M. Pontiki, D. Galanis, J. Pavlopoulos, H. Papageorgiou, I. Androutsopoulos, and S. Manandhar, “Semeval-2014 task 4: Aspect based sentiment analysis,” in *SemEval*, 2014, pp. 27–35.
- [124] M. Pontiki, D. Galanis, H. Papageorgiou, I. Androutsopoulos, S. Manandhar, M. AL-Smadi, M. Al-Ayyoub, Y. Zhao, B. Qin, O. De Clercq, V. Hoste, M. Apidianaki, X. Tannier, N. Loukachevitch, E. Kotelnikov, N. Bel, S. M. Jiménez-Zafra, and G. Eryigit, “Semeval-2016 task 5: Aspect based sentiment analysis,” in *SemEval*, 2016, pp. 19–30.
- [125] S. Poria, E. Cambria, and A. Gelbukh, “Aspect extraction for opinion mining with a deep convolutional neural network,” *Knowledge-Based Systems*, vol. 108, pp. 42–49, 2016.
- [126] D. Bahdanau, K. Cho, and Y. Bengio, “Neural machine translation by jointly learning to align and translate,” *arXiv preprint arXiv:1409.0473*, vol. abs/1409.0473, 2014.
- [127] M. Schuster and K. K. Paliwal, “Bidirectional recurrent neural networks,” *IEEE Transactions on Signal Processing*, vol. 45, no. 11, pp. 2673–2681, 1997.
- [128] M. Pontiki, D. Galanis, H. Papageorgiou, S. Manandhar, and I. Androutsopoulos, “Semeval-2015 task 12: Aspect based sentiment analysis,” in *SemEval*, Denver, Colorado, 2015, pp. 486–495.
- [129] R. He and J. McAuley, “Ups and downs: Modeling the visual evolution of fashion trends with one-class collaborative filtering,” in *the 25th WWW*, 2016, pp. 507–517.
- [130] K. S. Tai, R. Socher, and C. D. Manning, “Improved semantic representations from tree-structured long short-term memory networks,” in *ACL-IJCNN*, Beijing, China, 2015, pp. 1556–1566.

- [131] N. Peng, H. Poon, C. Quirk, K. Toutanova, and W.-t. Yih, “Cross-sentence n-ary relation extraction with graph lstms,” *Transactions of the Association for Computational Linguistics*, vol. 5, 2017.
- [132] W. L. Hamilton, R. Ying, and J. Leskovec, “Representation learning on graphs: Methods and applications,” *arXiv preprint arXiv:1709.05584*, 2017.
- [133] X. Zhu and Z. Ghahramani, “Learning from labeled and unlabeled data with label propagation,” Tech. Rep., 2002.