

Deep Memory Networks for Attitude Identification

Cheng Li
School of Information
University of Michigan
Ann Arbor
lichengz@umich.edu

Xiaoxiao Guo
Computer Science and Eng.
University of Michigan
Ann Arbor
guoxiao@umich.edu

Qiaozhu Mei
School of Information
University of Michigan
Ann Arbor
qmei@umich.edu

ABSTRACT

We consider the task of identifying attitudes towards a given set of entities from text. Conventionally, this task is decomposed into two separate subtasks: target detection that identifies whether each entity is mentioned in the text, either explicitly or implicitly, and polarity classification that classifies the exact sentiment towards an identified entity (the target) into positive, negative, or neutral.

Instead, we show that attitude identification can be solved with an end-to-end machine learning architecture, in which the two subtasks are interleaved by a deep memory network. In this way, signals produced in target detection provide clues for polarity classification, and reversely, the predicted polarity provides feedback to the identification of targets. Moreover, the treatments for the set of targets also influence each other – the learned representations may share the same semantics for some targets but vary for others. The proposed deep memory network, the AttNet, outperforms methods that do not consider the interactions between the subtasks or those among the targets, including conventional machine learning methods and the state-of-the-art deep learning models.

1. INTRODUCTION

In many scenarios, it is critical to identify people’s attitudes¹ towards a set of entities. Examples include companies who want to know customers’ opinions about their products, governments who are concerned with public reactions about policy changes, and financial analysts who identify daily news that could potentially influence the prices of securities. In a more general case, attitudes towards all entities in a knowledge base may be tracked over time for various in-depth analyses.

Different from a *sentiment* which might not have a target (e.g., “I feel happy”) or an *opinion* which might not have a polarity (e.g., “we should do more exercise”), an *attitude* can be roughly considered as a sentiment polarity towards a particular entity (e.g., “WSDM is a great conference”). Therefore, the task of attitude identification has been conventionally decomposed into two separate

¹ “The way you think and feel about someone or something,” as defined by Merriam-Webster. <http://www.merriam-webster.com/dictionary/attitude>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

WSDM 2017, February 06-10, 2017, Cambridge, United Kingdom

© 2017 ACM. ISBN 978-1-4503-4675-7/17/02...\$15.00

DOI: <http://dx.doi.org/10.1145/3018661.3018714>

rate subtasks: target detection that identifies whether an entity is mentioned in the text, either explicitly or implicitly, and polarity classification that classifies the exact sentiment towards the identified target, usually into three categories: positive, negative, and neutral.

Solving the two subtasks back-to-back is by no means unreasonable, but it may not be optimal. Specifically, intrinsic interactions between the two subtasks may be neglected in such a modularized pipeline. Indeed, signals identified in the first subtask – both the words that refer to the target and the positions of these words, could provide useful information for the polarity of sentiments. For example, the identified target in the sentence “*this Tiramisu cake is _*” indicates a high probability that a sentimental word would appear in the blank and is highly likely to be related to flavor or price. On the other hand, sentimental expressions identified in the second subtask and their positions could in turn provide feedback to the first task and signal the existence of a target. For example, the positive sentiment in “*the new Keynote is user friendly*” provides good evidence that “*Keynote*” is a software (the target) instead of a speech (not the target). In addition, models learned for certain targets and their sentiments may share some important dimensions with each other while differ on other dimensions. For example, two targets *food* and *service* may share many sentimental expressions, but the sentence “*we have been waiting for food for one hour*” is clearly about the service instead of the food. Failure to utilize these interactions (both between tasks and among targets) may compromise the performance of both subtasks.

Recent developments of deep learning has provided the opportunity of a better alternative to modularized pipelines, in which machine learning and natural language processing tasks can be solved in an *end-to-end* manner. With a carefully designed multi-layer neural network, learning errors backpropagate from upper layers to lower layers, which enables deep interactions between the learning of multi-grained representations of the data or multiple subtasks. Indeed, deep learning has recently been applied to target-specific sentiment analysis (mostly the second subtask of attitude identification) and achieved promising performance, where a given target is assumed to have appeared exactly once in a piece of text and the task is to determine the polarity of this text [38, 45, 36]. A deep network structure learns the dependency between the words in the context and the target word.

In another related topic known as multi-aspect sentiment analysis, where the goal is to learn the fine-grained sentiments on different aspects of a target, some methods have attempted to model aspects and sentiments jointly. Aspects are often assumed to be mentioned explicitly in text, so that the related entities can be extracted through supervised sequence labeling methods [21, 19, 46]; aspects mentioned implicitly can be extracted as fuzzy representa-

tions through unsupervised methods such as topic models [22, 41, 32]. While unsupervised methods suffer from low accuracy, it is usually difficult for supervised methods, like support vector machines (SVMs) [17], to interleave aspect extraction and sentiment classification.

In this paper, we show that the accuracy of attitude identification can be significantly improved through effectively modeling the interactions between subtasks and among targets. The problem can be solved with an end-to-end machine learning architecture, where the two subtasks are interleaved by a deep memory network. The proposed model, called the AttNet, also allows different targets to interact with each other, by sharing a common semantic space and simultaneously keep their own space, making it possible for all targets to be learned in a unified model. The proposed deep memory network outperforms models that do not consider the subtask or target interactions, including conventional supervised learning methods and state-of-the-art deep learning models.

The rest of the paper is organized as follows. Section 2 summarizes the related literature. In Section 3, we describe how the deep neural network is designed to incorporate the interaction both between subtasks and among targets. We present the design and the results of empirical experiments in Section 4 and Section 5, and then conclude the paper in Section 6.

2. RELATED WORK

Sentiment analysis has been a very active area of research [27, 30]. While sentiment in general does not need to have a specific target, the notion of *attitude* is usually concerned with a sentiment towards a target entity (someone or something). As one category of sentiment analysis, there is much existing work related to attitude identification, which generally takes place in three domains: multi-aspect sentiment analysis in product reviews, stance classification in online debates, and target-dependent sentiment classification in social media posts. Below we categorize existing work by the problem settings, e.g., whether the target is required to be explicitly mentioned.

Explicitly tagged targets. There has been a body of work that classifies the sentiment towards a particular target that is explicitly mentioned and tagged in text, mostly applied to social media text such as Tweets. Due to the short length of Tweets, many models assume that targets appear exactly once in every post. Jiang et al. [14] developed seven rule-based target-dependent features, which are fed to an SVM classifier. Dong et al. [6] proposed an adaptive recursive neural network that propagates sentiment signals from sentiment-bearing words to specific targets on a dependence tree. Vo et al. [38] split a Tweet into a left context and a right context according to a given target, and used pre-trained word embeddings and neural pooling functions to extract features. Zhang et al. [45] extended this idea by using gated recursive neural networks. The paper most relevant to ours is Tang et al. [36], which applied Memory Networks [35] to the task of multi-aspect sentiment analysis. Aspects are given as inputs, assuming that they have already been annotated in the text. Their memory network beat all LSTM-based networks but did not outperform SVM with hand-crafted features.

Model structures for target-dependent sentiment classification heavily rely on the assumption that the target appears in the text *explicitly*, and exactly *once*. These models could degenerate when a target is implicitly mentioned or mentioned multiple times. Additionally, they do not consider the interactions between the subtasks (target detection and sentiment classification) or among the targets.

Given target, one per instance. In the problem of stance classification, the target, mentioned explicitly or implicitly, is given but not tagged in a piece of text. The task is only to classify the sen-

timent polarity towards that target. Most methods train a specific classifier for each target and report performance separately per target. Many researchers focus on the domain of online debates. They utilized various features based on n-grams, part of speech, syntactic rules, and dialogic relations between posts [40, 10, 7, 31]. The workshop SemEval-2016 presented a task on detecting stance from tweets [24], where an additional category is added for the given target, indicating the absence of sentiment towards the target. Mohammad et al. [25] beat all teams by building an SVM classifier for each target.

As stance classification deals with only one given target per instance, it fails to consider the interaction between target detection and sentiment classification. Furthermore, the interplay among targets is ignored by training a separate model per target.

Explicit targets, not tagged. In the domain of product reviews, a specific aspect of a product could be considered as a target of attitudes. When the targets appear in a review but are not explicitly tagged, they need to be extracted first. Most work focuses on extracting explicitly mentioned aspects. Hu et al. [12] extracted product aspects via association mining, and expanded seed opinion terms by using synonyms and antonyms in WordNet. When supervised learning approaches are taken, both tasks of aspect extraction and polarity classification can be cast as a binary classification problem [17], or as a sequence labeling task and solved using sequence learning models such as conditional random fields (CRFs) [21, 19] or hidden Markov models (HMMs) [46].

Implicit targets. There are studies that attempt to address the situation when aspects could be implicitly mentioned. Unsupervised learning approaches like topic modeling treat aspects as topics, so that topics and sentiment polarity can be jointly modeled [22, 41, 32]. The workshop of SemEval-2015 announced a task of aspect based sentiment analysis [28], which separates aspect identification and polarity classification into two subtasks. For aspect identification, top teams cast aspect category extraction as a multi-class classification problem with features based on n-grams, parse trees, and word clusters.

Although aspect identification and polarity classification are modeled jointly here, it is hard to train unsupervised methods in an end-to-end way and directly optimize the task performance.

Deep learning for sentiment analysis. In the general domain of sentiment analysis, there has been an increasing amount of attention on deep learning approaches. In particular, Bespalov et al. [1] used Latent Semantic Analysis to initialize the word embedding, representing each document as the linear combination of n-gram vectors. Glorot et al. [9] applied Denoising Autoencoders for domain adaptation in sentiment classification. A set of models have been proposed to learn the compositionality of phrases based on the representation of children in the syntactic tree [33, 34, 11]. These methods require parse trees as input for each document. However, parsing does not work well on user generated contents, e.g., tweets [8]. Liu et al. [20] used recurrent neural networks to extract explicit aspects in reviews.

Compared to the existing approaches, our work develops a novel deep learning architecture that emphasizes the interplay between target detection and polarity classification, and the interaction among multiple targets. These targets can be explicitly or implicitly mentioned in a piece of text and do not need to be tagged a priori.

3. ATTNET FOR ATTITUDE IDENTIFICATION

We propose an end-to-end neural network model to interleave the target detection task and the polarity classification task. The **target**

detection task is to determine whether a specific target occurs in a given context either explicitly or implicitly. The **polarity classification** task is to decide the attitude of the given context towards the specific target if the target occurs in the context. Formally, a **target detection classifier** is a function mapping pairs of targets and contexts into binary labels, $(\text{context}, \text{target}) \rightarrow \{\text{present}, \text{absent}\}$. A **polarity classifier** is a function mapping pairs of targets and contexts into three attitude labels, $(\text{context}, \text{target}) \rightarrow \{\text{positive}, \text{negative}, \text{neutral}\}$. For example, given a **context**, *if everyone has guns, there would be just mess*, and a **target**, *gun control*, the correct label is *present* for the target detection task and *positive* for polarity classification.

Our model builds on the insight that the target detection task and the polarity classification task are deeply coupled in several ways.

- The polarity classification depends on the target detection because the polarity is meaningful only if the target occurs in the context. Reversely, the polarity classification task provides indirect supervision signals for the target detection task. For example, if the attitude label *positive* is provided for a context-target pair, the target must have occurred in the context following the definition. Such indirect supervision signals are useful especially when the target only occurs in the context implicitly.
- The signal words in the target detection and the polarity classification task are usually position-related: the signal words to determine the polarity are usually the surrounding words of the signal words to detect the target. Moreover, when a context has multiple targets, the signal words usually cluster for different targets [12, 30].
- Different targets interact in both the target detection task and the polarity classification task. Intuitively, some context words could mean the same for many targets, while some context words could mean differently for different targets.

Specifically, our model introduces several techniques building on the interaction between the target detection task and the polarity classification task accordingly.

- The output of the target detection is concatenated as part of the input of the polarity classification task to allow polarity classification to be conditioned on target detection. Polarity classification labels are also used to train the target detection classifier by back-propagating the errors of the polarity classification to the target detection end-to-end.
- The attention of polarity classification over context words are preconditioned by the attention of target detection. The polarity classification task benefits from such precondition especially when there are multiple targets in the context.
- Target-specific projection matrices are introduced to allow some context words to have similar representations among targets and other context words to have distinct representations. These matrices are all learned in an end-to-end fashion.

We propose a deep memory network model, called the AttNet, which implements the above motivation and ideas. In the rest of this section, we give a brief introduction to the memory network model, followed by a description of a single layer version of the model. Then we extend the expressiveness and capability of the model by stacking multiple layers.

3.1 Background: Memory Networks

As one of the recent developments of deep learning, memory networks [35] have been successfully applied to language modeling, question answering, and aspect-level sentiment analysis [36], which generates superior performance over alternative deep learning methods, e.g., LSTM.

Given a context (or document, e.g., “*we have been waiting for food for one hour*”) and a target (e.g., *service*), a memory network layer converts the context into a vector representation by computing a weighted sum of context word vector representations. The weight is a score that measures the relevance between the context word and the target (e.g., a higher score between the words *waiting* and *service*). The vector representation of the context is then passed to a classifier for target detection or polarity classification. An attractive property is that all parameters, including the target embeddings, context word embeddings and scores, are end-to-end trainable without additional supervision signals.

AttNet improves the original memory network models for attitude identification by (1) interleaving the target detection and polarity classification subtasks and (2) introducing target-specific projection matrices in representation learning, without violating the end-to-end trainability.

3.2 Single Layer AttNet

We begin by describing AttNet in the single layer case, shown in Figure 1. Hereafter for simplicity, we refer to the task of target detection as **TD**, and polarity classification as **PC**.

(1) Target Embedding. Each query target is represented as a one-hot vector, $\mathbf{q} \in \mathbb{R}^{N_{\text{target}}}$, where N_{target} is the number of targets. All targets share a target embedding matrix $\mathbf{B} \in \mathbb{R}^{d \times N_{\text{target}}}$, where d is the embedding dimensionality. The matrix \mathbf{B} converts a target into its embedding vector $\mathbf{u} = \mathbf{B}\mathbf{q}$, which is used as the input for the **TD** task.

(2) Input Representation and Attention for TD. We compute match scores between the context (or document) and the target for content-based addressing. The context is first converted into a sequence of one-hot vectors, $\{\mathbf{x}_i \in \mathbb{R}^{N_{\text{voc}}}\}$, where \mathbf{x}_i is the one-hot vector for the i -th word in the context and N_{voc} is the number of words in the dictionary. The entire set of $\{\mathbf{x}_i\}$ are then embedded into a set of input representation vectors $\{\mathbf{m}_i^t\}$ by:

$$\mathbf{m}_i^t = \mathbf{V}_q^t \mathbf{A}^t \mathbf{x}_i,$$

where $\mathbf{A}^t \in \mathbb{R}^{d \times N_{\text{voc}}}$ is the word embedding matrix shared across targets, superscript t stands for the **TD** task, and $\mathbf{V}_q^t \in \mathbb{R}^{d \times d}$ is a target-specific projection matrix for target q , which allows context words \mathbf{x}_i to share some semantic dimensions for some targets while vary for others.

In the embedding space, we compute the match scores between the target input representation \mathbf{u} and each context word representation \mathbf{m}_i^t by taking the inner product followed by a softmax, $a_i^t = \text{SoftMax}(\mathbf{u}^T \mathbf{m}_i^t)$, where $\text{SoftMax}(w_i) = \exp(w_i) / \sum_j \exp(w_j)$. In this way, \mathbf{a}^t is a soft attention (or probability) vector defined over the context words.

(3) Output Representation for TD. A different embedding matrix, $\mathbf{C}^t \in \mathbb{R}^{d \times N_{\text{voc}}}$, is introduced for flexibility in computing the output representation of context words by:

$$\mathbf{c}_i^t = \mathbf{V}_q^t \mathbf{C}^t \mathbf{x}_i$$

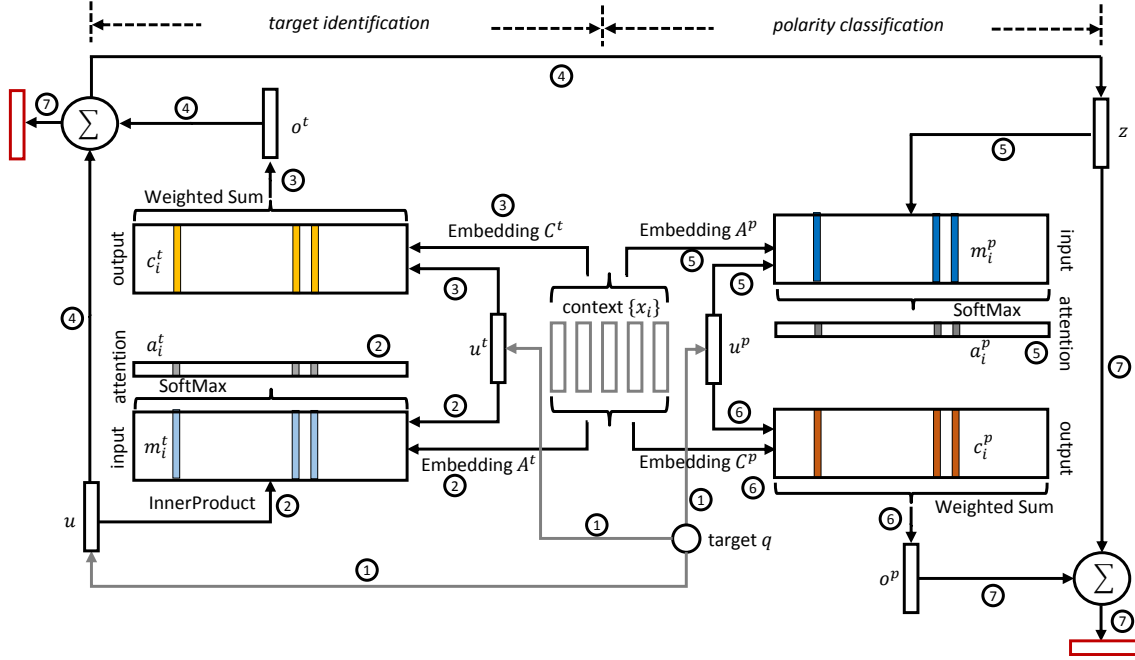


Figure 1: A single layer version of AttNet. Key submodules are numbered and correspondingly detailed in the text.

The response output vector \mathbf{o}^t is then a sum over the outputs \mathbf{c}_i^t , weighted by the attention vector from the input: $\mathbf{o}^t = \sum_i a_i^t \mathbf{c}_i^t$.

(4) Interleaving TD and PC. In the single layer case, the sum of the output vector \mathbf{o}^t and the target query embedding \mathbf{u} is then passed to the PC task, $\mathbf{z} = \mathbf{o}^t + \mathbf{u}$.

(5) Input Representation and Attention for PC. Similar to the TD task, we convert the entire set of $\{x_i\}$ into input representation vectors $\{\mathbf{m}_i^p\}$ by:

$$\mathbf{m}_i^p = \mathbf{V}_q^p \mathbf{A}^p \mathbf{x}_i,$$

where $\mathbf{A}^p \in \mathbb{R}^{d \times N_{\text{voc}}}$ is the input embedding matrix for PC. We use separate embedding matrices \mathbf{A}^t and \mathbf{A}^p for TD and PC, as the words could have different semantics in the two tasks. For similar reasons, we use different projection matrices \mathbf{V}_q^t and \mathbf{V}_q^p for the two tasks.

Given the polarity input representation $\{\mathbf{m}_i^p\}$, we also compute the soft attention over the context words for polarity identification, $a_i^p = \text{SoftMax}(\mathbf{z}^T \mathbf{m}_i^p)$.

(6) Output Representation for PC. There is also one corresponding output vector \mathbf{c}_i^p in PC for each \mathbf{x}_i :

$$\mathbf{c}_i^p = \mathbf{V}_q^p \mathbf{C}^p \mathbf{x}_i,$$

where $\mathbf{C}^p \in \mathbb{R}^{d \times N_{\text{voc}}}$ is the polarity output embedding matrix. It has been observed that sentiment-bearing words are often close to the target [12, 30]. Based on this observation, attentions, or positions of important words that identify the target in the first module, could provide prior knowledge to learn the attention of the second module. Therefore we compute the final attention vector as a function of original attentions of both tasks:

$$\mathbf{b}^p = (1 - \lambda) \mathbf{a}^p + \lambda f(\mathbf{a}^t), \quad (1)$$

where $0 \leq \lambda < 1$ controls the importance of the second term, and f is a moving average function which shifts attentions from words of high values to their surrounding neighbors. The output vector is $\mathbf{o}^p = \sum_i b_i^p \mathbf{c}_i^p$.

(7) Prediction for TD and PC. To predict whether a target presents, the sum of the output vector of target classification \mathbf{o}^t and the target query vector \mathbf{u} is passed through a weight matrix $\mathbf{W}^t \in \mathbb{R}^{2 \times d}$ (2 is the number of classes: *present*, *absent*) and a softmax operator to produce the predicted label, a vector of class probabilities: $\mathbf{y}^t = \text{SoftMax}(\mathbf{W}^t(\mathbf{o}^t + \mathbf{u}))$.

Similarly, the sum of the output vectors \mathbf{o}^p of PC and its input vector \mathbf{z} is then passed through a weight matrix $\mathbf{W}^p \in \mathbb{R}^{3 \times d}$ and a softmax operator to produce the predicted attitude label vector, $\mathbf{y}^p = \text{SoftMax}(\mathbf{W}^p(\mathbf{o}^p + \mathbf{z}))$.

3.3 Multiple Layer AttNet

We now extend our model to stacked multiple layers. Figure 2 shows a three layer version of our model. The layers are stacked in the following way:

Functionality of Each Layer. For TD, the input to the $(k+1)$ -th layer is the sum of the output \mathbf{o}_k^t and the input \mathbf{u}_k from the k -th layer, followed by a sigmoid nonlinearity: $\mathbf{u}_{k+1} = \sigma(\mathbf{H}^t \mathbf{u}_k + \mathbf{o}_k^t)$, where $\sigma(x) = 1/(1 + \exp(x))$ is the sigmoid function and \mathbf{H}^t is a learnable linear mapping matrix shared across layers. For the PC task, the input to the first layer is the transformed sum from the last layer of the TD module, $\mathbf{z}_1 = \sigma(\mathbf{H}^t \mathbf{u}_{K_t} + \mathbf{o}_{K_t}^t)$, where K_t is the number of stacked layers in the TD task. Thus the prediction of polarity would depend on the output of the TD task and reversely the TD task would benefit from indirect supervision from the PC task by backward propagation of errors. Similarly for PC, the input to the $(k+1)$ -th layer is the sum of the output \mathbf{o}_k^p and the input \mathbf{z}_k

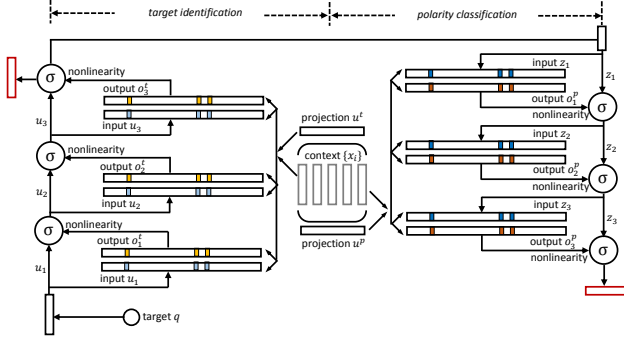


Figure 2: A three layer version of our model. Both the TD and PC modules have three stacked layers.

from the k -th layer, followed by a sigmoid nonlinearity: $\mathbf{z}_{k+1} = \sigma(\mathbf{H}^p \mathbf{z}_k + \mathbf{o}_k^p)$.

Attention for PC. In the single layer case, the attention for PC is based on that of the TD module. When layers are stacked, all layers of the first module collectively identify important attention words to detect the target. Therefore we compute the averaged attention vector across all layers in the TD module $\bar{\mathbf{a}}^t = \frac{1}{K_t} \sum_{k=1}^{K_t} \mathbf{a}_k^t$. Accordingly for k -th layer of the PC module, the final attention vector is $\mathbf{b}_k^p = (1 - \lambda) \mathbf{a}_k^p + \lambda f(\bar{\mathbf{a}}^t)$, and the output vector is $\mathbf{o}_k^p = \sum_i b_{k,i}^p \mathbf{c}_{k,i}^p$.

Tying Embedding and Projection Matrices. The embedding matrices and projection matrices are constrained to ease training and reduce the number of parameters [35]. The embedding matrices and the projection matrices are shared for different layers. Specifically, using subscription (k) denote the parameters in the k -th layer, for any layer k , we have $\mathbf{A}^{t(1)} \equiv \mathbf{A}^{t(k)}$, $\mathbf{C}^{t(1)} \equiv \mathbf{C}^{t(k)}$, $\mathbf{A}^{p(1)} \equiv \mathbf{A}^{p(k)}$, $\mathbf{C}^{p(1)} \equiv \mathbf{C}^{p(k)}$, $\mathbf{V}_q^{t(1)} \equiv \mathbf{V}_q^{t(k)}$ and $\mathbf{V}_q^{p(1)} \equiv \mathbf{V}_q^{p(k)}$.

Predictions for TD and PC. The prediction stage is similar to the single-layer case, with the prediction based on the output of the last layer K_t (for TD) and K_p (for PC). For the TD task, $\mathbf{y}^t = \text{SoftMax}(\mathbf{W}^t \sigma(\mathbf{H}^t \mathbf{u}_{K_t} + \mathbf{o}_{K_t}^t))$, while for PC, $\mathbf{y}^p = \text{SoftMax}(\mathbf{W}^p \sigma(\mathbf{H}^p \mathbf{z}_{K_p} + \mathbf{o}_{K_p}^p))$.

3.4 End-to-End Multi-Task Training

We use cross entropy loss to train our model end-to-end given a set of training data $\{ct_i, q_j, g_{ij}^t, g_{ij}^p\}$, where ct_i is the i -th context (or document), q_j is the j -th target, g_{ij}^t and g_{ij}^p are the ground-truth labels for the TD and the PC tasks respectively. The training is to minimize the objective function:

$$\mathcal{L} = - \sum_i \sum_j \left(\log(\mathbf{y}_{ij}^t(g_{ij}^t)) + \mathbb{1}_{g_{ij}^t} \log(\mathbf{y}_{ij}^p(g_{ij}^p)) \right),$$

where \mathbf{y}_{ij}^t is a vector of predicted probability for each class of TD, $\mathbf{y}_{ij}^t(s)$ selects the s -th element of \mathbf{y}_{ij}^t , $\mathbb{1}_{g_{ij}^t}$ equals to 1 if g_{ij}^t equals to class *present* and 0 otherwise. Note that when a target is not mentioned in a given context, the polarity term plays no role in the objective because the value of $\mathbb{1}_{g_{ij}^t}$ is zero.

4. EXPERIMENT SETUP

In the experiments, we compare AttNet to conventional approaches and alternative deep learning approaches on three real world data

sets, and we show the superior performance of our model. We also experiment with variants of AttNet as credit assignments for the key components in our model.

4.1 Data Sets

We examine AttNet on three domains that are related to attitude classification: online debates (Debates), multi-aspect sentiment analysis on product review (Reviews), and stance in tweets (Tweets).

Debates. This data set is from the Internet Argument Corpus version 2². The data set consists of political debates on three Internet forums³. In these forums, a person can initiate a debate by posting a topic and taking positions such as *favor* vs. *against*. Examples of topics are *gun control*, *death penalty* and *abortion*. Other users participate in these debates by posting their arguments for one of the sides.

Tweets. This data set comes from a task of the workshop SemEval-2016 on detecting stance on tweets [24]. Targets are mostly related to ideology, e.g., *atheism* and *feminist movement*⁴.

Review. This data set includes reviews of restaurants and laptops from SemEval 2014 [29] and 2015 [28], where subtasks of identifying aspects and classifying sentiments are provided. We merge two years' data to enlarge the data set, and only include aspects that are annotated in both years.

To guarantee enough training and test instances, for all the data sets we filter out targets mentioned in less than 100 documents. The original train-test split is used if provided, otherwise we randomly sample 10% data into test set. We further randomly sample 10% training data for validation. Text pre-processing includes stopword removal and tokenization by the CMU Twitter NLP tool [8]. The details of the data sets are shown in Table 1.

Table 1: Statistics of each data set.

Data set	Set	#docs	#pos	#neg	#neutral	#absent
Debates	train	24352	13891	10711	0	0
	val	2706	1530	1203	0	0
	test	3064	1740	1371	0	0
Tweets	train	2614	682	1253	0	679
	val	291	71	142	0	78
	test	1249	304	715	0	230
Review	train	5485	2184	1222	210	2336
	val	610	260	121	17	277
	test	1446	496	455	60	634

#pos means the number of documents with positive sentiment for each target. If one document contains positive sentiment towards two targets, it will be counted twice. *#absent* counts the number of documents without any attitude towards any target.

²<https://nlds.soe.ucsc.edu/iac2>.

³4forums(<http://www.4forums.com/political/>), ConvinceMe(<http://www.convinceme.net/>) and CreateDebate(<http://www.createdebate.com/>)

⁴Since there are less than 10 tweets with neutral stance, we only consider *positive* and *negative* attitude by discarding these neutral tweets.

4.2 Metrics

For our problem, each data set has multiple targets, and each target can be classified into one of the outcomes: *absent* (do not exist), *neutral*, *positive*, and *negative*. If we treat each outcome of each target as one *category*, we can adopt common metrics for multi-class classification. Since most targets do not appear in most instances, we have a highly skewed class distribution, where measures like accuracy are not good choices [3].

Apart from *precision*, *recall* and *AUC*, we also use the *macro-average F-measure* [44]. Let ρ_i and π_i be recall and precision respectively for a particular category i , $\rho_i = \frac{TP_i}{TP_i + FN_i}$, $\pi_i = \frac{TP_i}{TP_i + FP_i}$, where TP_i , FP_i , FN_i are the number of true positive, false positive, and false negative for category i . Given ρ_i and π_i , F-score of category i is computed as $F_i = \frac{2\pi_i\rho_i}{\pi_i + \rho_i}$. The macro-average F-score is obtained by taking the average over all categories. The final precision and recall are also averaged over individual categories. There is another micro-averaged F-measure, which is equivalent to accuracy. Therefore, we do not include it.

4.3 Baselines

We compare baseline methods from two large categories: conventional methods and alternative deep learning methods.

Each baseline method has various configurations, based on whether: (1) it trains a *single model* or *two separate models* for the **TD** and **PC** subtasks, and (2) it trains *one universal model* for all targets or *separated models* for *different* targets. To distinguish different configurations, we append *-sgl* when using a single model for the two subtasks, and *-sep* when using separate models for each subtask. Taking SVM as an example, SVM-*sgl* directly classify targets into four classes: *absent*, *neutral*, *positive*, and *negative*. In contrast, SVM-*sep* first classifies each target into two classes: *absent*, *present*, and use a second model to classify polarity: *neutral*, *positive*, and *negative*. Moreover, we append *-ind* when individual targets are trained on separate models, or *-all* when one model is trained for all targets.

4.3.1 Conventional baselines

SVM+features. SVM using a set of hand-crafted features has achieved the state-of-the-art performance in stance classification of SemEval 2016 task [25], online debates [10], and aspect-based sentiment analysis [36]. SVM has also demonstrated superior performance in document-level sentiment analysis compared with conditional random field methods [42]. Therefore we include all features from these methods that are general across domains, and use a linear kernel SVM implemented by LIBSVM [2] for classification. We list the set of features:

Document info: basic counting features of a document, including the number of characters, the number of words, the average words per document and the average word length.

N-grams: word unigrams, bigrams, and trigrams. We insert symbols that represent the start and end of a document to capture cue words [40].

Sentiment: the number of positive and negative words counted from the NRC Emotion Lexicon [26], Hu and Liu Lexicon [12], and the MPQA Subjectivity Lexicon [43].

Target: presence of the target phrase in the text. Furthermore, if the target is present, we generate a set of target dependent features according to [14]. To get a sense of these features, for the target *iPhone* in text *I love iPhone*, a feature *love_arg* will be generated.

POS: the number of occurrences of each part-of-speech tag (POS).

Syntactic dependency: a set of triples obtained by Stanford dependency parser [5]. More specifically, the triple is of the form

(rel, w_i, w_j) , where *rel* represents the grammatical relation between word w_i and w_j , e.g., *is subject of*.

Generalized dependency: the first word of the dependency triple is “backed off” to its part-of-speech tag [39]. Additionally, words that appear in sentiment lexicons are replaced by positive or negative polarity equivalents [39].

Embedding: the element-wise averages of the word vectors for all the words in a document. We use three types of word embeddings. Two of them are from studies on target-dependent sentiment classification [38, 45], which are the skip-gram embeddings of Mikilov et al. [23] and the sentiment-driven embeddings of Tang et al. [37]. The first type of embedding is trained on 5 million unlabeled tweets that contain emoticons, which guarantees that more sentiment related tweets are included. The second type of embedding is of 50 dimensions, which is publicly available⁵. The third type of embedding is also 50-dimensional, released by Collobert et al. [4] and trained on English Wikipedia⁶.

Word cluster: the number of occurrences of word clusters for all words in text. We perform K-means clustering on the word vectors.

Apart from two standard SVM model configurations, *SVM-sep-ind* and *SVM-sgl-ind*, we also compare with a hybrid model *SVM-cmb-ind*, whose prediction is *absent* if *SVM-sep-ind* says so, and otherwise it follows the decisions of *SVM-sgl-ind*.⁷

4.3.2 Deep Learning Baselines

BiLSTM, MultiBiLSTM and Memnet. We also compare to the bidirectional LSTM (**BiLSTM**) model, the state-of-the-art in target-dependent sentiment classification [45]. Their variant of BiLSTM model assumes that the given target always appears exactly once, and can be tagged in text by starting and ending offsets. When such assumption fails, their model is equivalent to standard BiLSTM. We include the standard multi-layered bidirectional LSTM (**MultiBiLSTM**) [13] as an extension. Recently, Tang et al. [36] applied memory networks (**Memnet**) to multi-aspect sentiment analysis. Their results show memory network performs comparably with feature based SVM and outperforms all LSTM-related methods in their tasks.

CNN and ParaVec. We include related deep learning techniques beyond the sentiment analysis domain, such as the convolutional neural networks (**CNN**) [15] and **ParaVec** [18]. **ParaVec** requires a huge amount of training data to reach decent performance. We enhance the performance of the ParaVec model by training over the merged training set of all data sets, plus the 5 million unlabeled tweets mentioned above.

Parser-dependent deep learning methods have also been applied to sentiment analysis [33, 34, 11]. These models are limited in our attitude identification problem for two reasons. First, they often work well with phrase-level sentiment labels, but only document-level sentiment labels are provided in our problems. Second, their parsers do not extend to user generated content, such as Tweets and Debates [8]. Our preliminary results show these methods work poorly on our problems and we do not include their results for simplicity.

For all deep learning methods, we report their *-sep-all* and *-sgl-all* version. Unlike SVM, deep methods perform quite well when using a single model for all targets, by casting the problem as a multi-task multi-class classification. Though not scalable, for the strongest baselines (BiLSTM and MultiBiLSTM), we in addition

⁵<http://ir.hit.edu.cn/~dyltang/>

⁶<http://ronan.collobert.com/senna/>

⁷*SVM-sgl-all* and *SVM-sep-all* have performance degeneration due to the interference of different targets. We do not include their results for simplicity.

train a separate model for each target. Since *-sep-ind* works better than *-sgl-ind*, we only report the former one. The variants of memory networks are detailed below.

4.4 Variants of AttNet

To assign the credit of key components in our model, we construct a competing model *AttNet-*. Unlike our proposed model, for *AttNet-* the target-specific projection matrices \mathbf{V}_q^p and \mathbf{V}_q^t are replaced by the identity matrix and are fixed during training. Thus the *AttNet-* model interleaves the target detection and polarity classification subtasks, but do not consider the interactions among targets. We refer our proposed model as *AttNet*, which allows the projection matrices to be learned during training, and thus word semantics could vary for targets.

For *AttNet-*, we report two settings in our experiments: *AttNet-ind* and *AttNet-all*. The former makes all targets share the same embedding, while the latter separates the embedding space completely for each target, i.e., targets are trained on separate models.

Table 2: Hyper-parameters for our method AttNet.

Hyper-parameters	Tweets	Review	Debates
L1 coeff	1e-6	1e-4	1e-6
L2 coeff	1e-4	1e-8	1e-8
init learning rate	0.05	0.01	0.005
#layers(target)	4	4	3
#layers(sentiment)	4	8	6
prior attention λ	0.5	0.1	0.5

The embedding size is set to 100 for all data sets. The sliding window size of the moving average function in Equation 1 is set to 3. $\#layers(target)$ is the number of memory layers for target detection, and $\#layers(sentiment)$ is the number for sentiment classification. *prior attention* λ is the weight for prior attention in Equation 1.

4.5 Training Details

All hyper-parameters are tuned to obtain the best performance of F-score on validation set. The candidate embedding size set is $\{50, 100, 200, 300\}$ for LSTM-related methods, SVM and CNN. The candidate number of clusters for K-means is $\{50, 100, 150\}$. The candidate relaxing parameter C for SVM model is $\{2^7, 2^6, \dots, 2^{-3}\}$. The CNN model has three convolutional filter sizes and their filter size candidates are $\{\{1, 2, 3\}, \{2, 3, 4\}, \{3, 4, 5\}, \{2, 4, 6\}\}$, and the candidate number of filters is $\{50, 100, 200, 300\}$. For ParaVec, we experiment with both skip-gram model or bag-of-words model, and select the hidden layer size from $\{2^6, 2^7, \dots, 2^{10}\}$.

We explored three weight initialization methods of word embeddings for LSTM-related and CNN baselines: (1) sampling weights from a zero-mean Gaussian with 0.1 standard deviation; (2) initializing from the pre-trained embedding matrix, and (3) using a fixed pre-trained embedding matrix.

Memory network models, including our model, are initialized by sampling weights from a zero-mean Gaussian with unit standard deviation. The candidate number of memory layers is $\{2, 3, \dots, 9\}$. The prior attention parameter λ of our model is selected from $\{0, 0.1, 0.5, 0.9, 1\}$. The capacity of memory, which has limited impact on the performance, is restricted to 100 words without further tuning. A null symbol was used to pad all documents to this fixed size. To reduce the model complexity, the projection matrices are initialized in the way that each column is a one-hot vector.

Deep learning models are optimized by Adam [16]. The initial learning rate is selected from $\{0.1, 0.05, 0.01, 0.005, 0.001\}$, and L1-coefficient and L2-coefficient of regularizers are selected from

$\{10^{-2}, 10^{-4}, \dots, 10^{-10}\}$. The hyper-parameters of our model *AttNet+* for different data sets are listed in Table 2.

5. EXPERIMENT RESULTS

5.1 Overall Performance

The overall performance of all competing methods over data sets are shown in Table 3⁸. Evaluating with F-score and AUC, we make the following observations. Our method *AttNet* outperforms all competing methods significantly. This empirically confirms that interleaving target detection and polarity classification subtasks combined with target-specific representations could benefit attitude identification.

The variants of our model, *AttNet-all* and *AttNet-ind*, have already gained significant improvement over the strongest baselines on all data sets. More importantly, the two methods significantly outperform the *Memnet-sep-all* and *Memnet-sep-all* baselines, which do not interleave the subtasks. Such empirical findings cast light on that interleaving the subtasks indeed improves the attitude identification performance. In contrast, separating the two subtasks of attitude identification could lead to performance degeneration.

Our model *AttNet* also outperforms its variants, *AttNet-all* and *AttNet-ind*, on all data sets. The performance advantage of *AttNet* comes from the adoption of target-specific projection matrices in representation learning, since these matrices are the only differences between *AttNet* and *AttNet-*. Even though the improvement from adopting target-specific projection matrices is not as marked as from the techniques of interleaving the subtasks, the improvement is still significant. This result confirms that attitude identification could benefit from the learned representations that share the same semantics for many targets but vary for some targets.

By examining the *precision* and *recall* results, we find that the superior performance of our model is mainly from the significant improvement of *recall*, though both *precision* and *recall* are improved significantly on the Debates data set.

5.2 Performance on Subtasks

We have established that our models outperform competing methods on all data sets. In order to further assign the credits of the improvement of our methods, we evaluate our models on the two subtasks: target detection and polarity classification, with results given in Table 4 and 5 respectively. Since different configurations of the same method work similarly, we only present the results where separate models are trained for each task. It can be seen from Table 4 that the target detection task is relatively easy, as all methods can achieve quite high scores. This also means that it is hard to improve any further on this task. In terms of precision and recall, SVM performs quite well on the precision metric, especially for the Review data set. While most deep learning methods focus more on enhancing recall. When considering both precision and recall, most deep learning methods are still better, as the F-score shows.

The second task is only evaluated on documents with ground-truth sentiments towards particular targets, with F-scores averaged over all targets and three sentiment classes: *positive*, *negative*, and *neutral*. We make several notes for this evaluation. (1) In order to achieve a high score in the second task, it is still important to classify correctly the presence of a target. (2) In general the scores for all methods in the second task are low, due to that the classifier might predict a target as *absent*, even though the ground-truth

⁸The performance of all methods on the Review data set is lower than the other two because Review data set handles three polarities while the others only need to handle two polarities as shown in Table 1.

Table 3: Performance of competing methods: AttNet achieves top performance.

Method	Tweets				Review				Debates			
	F-score	AUC	Precision	Recall	F-score	AUC	Precision	Recall	F-score	AUC	Precision	Recall
SVM-sep-ind	59.93	69.20	68.70	55.69	38.43	57.99	51.22	36.83	58.30	72.10	64.48	57.81
SVM-sgl-ind	57.44***	66.64***	69.87	52.45***	36.06**	56.84**	50.79	34.07**	59.75	72.25	66.39	57.67
SVM-cmb-ind	57.09***	66.46***	69.84	52.25***	35.71***	56.61***	50.73	33.78***	59.86	71.68	66.28	56.48
ParaVec-sep-all	53.17***	62.88***	56.75***	48.29***	34.02***	55.26***	38.04***	30.47***	56.32**	68.12***	59.09***	49.41***
ParaVec-sgl-all	54.15***	63.41***	57.52***	48.76***	34.26***	55.31***	38.26***	30.89***	55.35***	67.48***	59.46***	49.82***
CNN-sep-all	58.05*	70.10	62.43***	56.19	37.15*	57.55	43.73***	33.24**	57.38	70.70**	61.81**	52.94***
CNN-sgl-all	58.69	70.71*	61.83***	56.64	35.45***	56.29*	44.65***	32.83**	56.23**	69.92**	60.75**	52.29***
BiLSTM-sep-all	61.16	71.26**	63.45*	59.87***	40.78*	61.01***	42.54***	39.01**	59.83*	71.91	65.94*	57.65
BiLSTM-sgl-all	60.86	71.02**	62.58*	59.61***	39.68	60.84**	41.88***	38.81*	58.66	72.01	64.87	57.89
BiLSTM-sep-ind	59.49	71.92**	61.44***	57.86	40.42	62.25***	42.68***	39.78**	58.75	72.83	64.73	57.95
MultiBiLSTM-sep-all	60.53	71.51**	64.81*	57.76	40.47	60.71**	44.89***	37.67*	59.24*	72.24	64.75	58.43
MultiBiLSTM-sgl-all	60.59	71.32**	64.27*	57.97	39.38	59.68*	43.22***	37.92*	58.98	71.18	63.46	57.26
MultiBiLSTM-sep-ind	59.71	71.16**	63.72*	57.94	40.81	61.27**	44.76***	38.02*	58.36	72.93	64.15	57.14
Memnet-sep-all	59.44	71.68**	63.22*	59.80***	41.75**	61.82***	45.61***	39.25**	60.42**	73.84*	65.37	58.92*
Memnet-sgl-all	60.69	71.80**	63.48*	59.97***	41.65**	61.53***	45.23***	39.13**	59.67*	73.31**	64.27	58.83
Proposed methods												
AttNet-all	63.42**	72.94***	68.78	60.57***	43.91***	63.18***	47.89**	42.77**	64.23***	76.13***	67.19**	62.17**
AttNet-ind	63.09**	72.73***	68.33	59.68**	44.15***	63.53***	50.02	40.58**	65.01***	76.35***	70.08***	60.70**
AttNet	64.62**	74.76***	68.40	62.09**	45.93**	65.58***	50.34	44.95**	67.68***	78.48***	74.55***	66.31**

*(**, ***) indicate that one method is statistically significantly better or worse than *SVM-sep-ind* (which is in general the best configuration among all SVM models) according to t-test [44] at the significance level of 0.05(0.01,0.001). $\diamond\diamond\diamond$ indicate *AttNet* outperforms the better one between *Memnet-sep-all* and *Memnet-sgl-all* at the significance level of 0.01(0.001). $\nabla\nabla(\nabla\nabla)$ indicate *AttNet* outperforms the better one between *AttNet-all* and *AttNet-ind* at the significance level of 0.01(0.001).

Table 4: Performance on target detection for -sep models.

	Tweets	Review	Debates
SVM	79.74	67.84	84.59
ParaVec	89.12 ,75.00	84.13 ,63.52	91.47,81.00
CNN	80.62	65.34**	80.21**
BiLSTM	81.50**	70.41**	85.05
MultiBiLSTM	81.53**	69.26**	85.33
Memnet	81.29**	70.71**	86.29*
AttNet-all	82.58**	71.84***	89.02***
AttNet-ind	82.74**	71.95***	88.89**
AttNet	84.89**	72.59***	89.35***

The first row of each method shows F-score, followed by precision and recall on the second row.

class can only be drawn from three sentiment classes. (3) It is possible for a method to outperform SVM on both tasks, while still obtain close results when two tasks are evaluated together. This results from our method of evaluation on the second task, where a document is included only when it expresses sentiment towards a particular target.

Based on the results from Table 5, we can see that the percentage of improvement over SVM is much higher than that of the first task. Intuitively, the sentiment task requires better modeling of the non-linear interaction between the target and the context, while for the target detection task, presence of certain signal words might be enough.

5.3 Training Time Analysis

In order to measure the training speed of each model, we train all deep learning methods on a server with a single TITAN X GPU. For SVM, it is trained on the same server with a 2.40 GHz CPU and 120 G RAM. All methods are trained sequentially without parallelization.

Table 5: Performance on polarity classification for -sep models.

	Tweets	Review	Debates
SVM	44.45	21.37	42.52
ParaVec	64.50,34.28	53.25,14.66	57.37,38.29
CNN	39.20***	17.69***	30.59***
BiLSTM	56.91**	26.46***	31.05***
MultiBiLSTM	42.95	19.42*	35.15***
Memnet	58.20**	35.71	49.82**
AttNet-all	61.35*	41.79***	54.58,35.59
AttNet-ind	46.26**	24.06**	41.87
AttNet	60.56**	39.36**	49.69*
AttNet-all	47.81**	25.47**	44.61
AttNet-ind	61.20**	40.52**	54.40*
AttNet	50.91***	32.43***	50.46***
AttNet-all	66.39	42.00	60.72
AttNet-ind	49.16*	32.79**	51.88**
AttNet	64.01	41.74***	62.70**
AttNet-all	52.23	35.34	55.93
AttNet-ind	65.54,44.57	59.99	71.53

The first row of each method shows F-score, followed by precision and recall on the second row.

SVM can finish training in less than one hour, but its required training time increase linearly as the number of targets increases.

For all deep learning methods, the number of epochs required for training is in general very close, which is around 20 epochs averaged over all data sets.

Comparing the training time per epoch, ParaVec and CNN are much faster than other methods (less than 5 seconds / epoch). Despite the training efficiency, their effectiveness is a problem. When all targets share a single model, LSTM has a speed of 200 seconds/epoch, while standard memory networks have a speed of 150 seconds/epoch. Memory networks in many tasks, e.g., language modeling, are much faster than LSTM, due to the expensive recursive operation of LSTM. However in our problem setting, each target has to be forwarded one by one for every document, lowering the efficiency of memory networks. When individual targets are trained on separate LSTMs, LSTMs require far more training time (1000 seconds/epoch).

AttNet consumes 200 seconds per epoch. Comparing to stan-

dard memory networks, *AttNet* produces some additional overhead by introducing the interaction between subtasks, and by adding a projection matrix. But this overhead is small.

The efficiency of deep learning methods could be improved by parallelization. Since there are already many work on this topic, which could increase the speed without sacrificing effectiveness, we do not go further into this direction.

Summary: empirical experiments demonstrated that the proposed deep memory networks, *AttNet* and its variants outperforms conventional supervised learning methods. This is promising but perhaps not surprising given the success of deep learning in general. It is encouraging to notice that *AttNet* also improves the state-of-the-art deep learning architectures. This improvement is statistically significant, and can be observed for both subtasks and for attitude identification as a whole. The improvement of effectiveness does not compromise learning efficiency.

5.4 Visualization of attention

In order to better understand the behavior of our models, we compare the attention weights given by our model *AttNets* and the competing method *Memnet*.

1. It is admittedly to have them for policy . if everyone have guns there would be just mess . (Truth: gun control+. Predict + given gun control.)
2. Highly impressed from the decor to the food to the great night ! (Truth: service+, ambience+, food+. Predict + given ambience.)
3. When we inquired about ports - the waitress listed off several but did not know taste variations or cost . (Truth: service-. Predict absent given drink.)

(a) Attention given by *AttNets*.

1. It is admittedly to have them for policy . if everyone have guns there would be just mess . (Predict - given gun control.)
2. Highly impressed from the decor to the food to the great night ! (Predict absent given ambience.)
3. When we inquired about ports - the waitress listed off several but did not know taste variations or cost . (Predict - given drink.)

(b) Attention given by *Memnet*.

Figure 3: Visualization of learned attention. Red patches highlighting the top half of the text indicate model’s attention weight in the target detection task, while green ones highlighting the bottom half show the polarity classification task. Darker colors indicate higher attentions. *Truth: service+* means that the ground-truth sentiment towards *service* is positive, while *Predict + given ambience* gives the predicted positive sentiment given the query target *ambience*.

Figure 3 (a) and (b) list some examples of word attentions generated by different models for the same set of sentences in the test set. In the first sentence, both *them* and *guns* are found as targets by *AttNets*, while words like *mess* and *policy* are found as sentiment words. Though *Memnet* correctly identifies the existence of the attitude towards gun control, it fails to find important words to classify the polarity of sentiment. This suggests the importance of

interleaving the two tasks – successfully identifying mentioned targets could offer clues about the finding of sentiment words for the second task.

The second sentence is from a review of a restaurant, when *ambience* is used as the query target. We can see that the target detection module of *AttNets* captures the word *decor*, which signals the presence of the target *ambience*. The polarity classification module then focuses on extracting sentiment words associated with the target. However for the baseline *Memnet*, it captures both *decor* and *food* in the first task, mistakenly considering all sentiments are only describing food other than the ambience. Consequently, it judges that there is no attitude towards ambience. This example shows us the benefit of using the projection matrices to consider the interaction and distinction between targets. Otherwise the model might easily be confused by to which entity the sentiments are expressed.

From the third sentence, we can see how our model *AttNets* determines that the query target *drink* does not exist. The first module highlights words like ports (wine name), waitress, and the second module extracts negative sentiments *but not know*, which is usually used to describe people, rather than drink. *Memnet* almost has the same attention distribution as *AttNets*, but still fails to produce the correct prediction. Similar to the second case, projection matrices are important for models to figure out the common phrases used to describe different set of entities.

6. CONCLUSION

Attitude identification, a key problem of modern natural language processing, is concerned with detecting one or more target entities from text and then classifying the sentiment polarity towards them. This problem is conventionally approached by separately solving the two subtasks and usually separately treating each target, which fails to leverage the interplay between the two subtasks and the interaction among the target entities. Our study demonstrates that modeling these interactions in a carefully designed, end-to-end deep memory network significantly improves the accuracy of the two subtasks, target detection and polarity classification, and attitude identification as a whole. Empirical experiments proves that this novel model outperforms models that do not consider the interactions between the two subtasks or among the targets, including conventional methods and the state-of-the-art deep learning models.

This work opens the exploration of interactions among subtasks and among contexts (in our case, targets) for sentiment analysis using an end-to-end deep learning architecture. Such an approach can be easily extended to handle other related problems in this domain, such as opinion summarization, multi-aspect sentiment analysis, and emotion classification. Designing specific network architecture to model deeper dependencies among targets is another intriguing future direction.

Acknowledgment

This work is partially supported by the National Science Foundation under grant numbers IIS-1054199 and SES-1131500.

7. REFERENCES

- [1] D. Bespalov, B. Bai, Y. Qi, and A. Shokoufandeh. Sentiment classification based on supervised latent n-gram analysis. In *Proc. of CIKM*, 2011.
- [2] C.-C. Chang and C.-J. Lin. Libsvm: a library for support vector machines. *TIST*, 2011.
- [3] N. V. Chawla. Data mining for imbalanced datasets: An overview. In *Data mining and knowledge discovery handbook*. 2005.

- [4] R. Collobert, J. Weston, L. Bottou, M. Karlen, K. Kavukcuoglu, and P. Kuksa. Natural language processing (almost) from scratch. *JMLR*, 2011.
- [5] M.-C. De Marneffe, B. MacCartney, C. D. Manning, et al. Generating typed dependency parses from phrase structure parses. In *Proc. of LREC*, 2006.
- [6] L. Dong, F. Wei, C. Tan, D. Tang, M. Zhou, and K. Xu. Adaptive recursive neural network for target-dependent twitter sentiment classification. In *ACL*, 2014.
- [7] A. Faulkner. Automated classification of stance in student essays: An approach using stance target information and the wikipedia link-based measure. In *FLAIRS*, 2014.
- [8] K. Gimpel, N. Schneider, B. O'Connor, D. Das, D. Mills, J. Eisenstein, M. Heilman, D. Yogatama, J. Flanigan, and N. A. Smith. Part-of-speech tagging for twitter: Annotation, features, and experiments. In *Proc. of ACL*, 2011.
- [9] X. Glorot, A. Bordes, and Y. Bengio. Domain adaptation for large-scale sentiment classification: A deep learning approach. In *Proc. of ICML*, 2011.
- [10] K. S. Hasan and V. Ng. Stance classification of ideological debates: Data, models, features, and constraints. In *IJCNLP*, 2013.
- [11] K. M. Hermann and P. Blunsom. The role of syntax in vector space models of compositional semantics. In *ACL*, 2013.
- [12] M. Hu and B. Liu. Mining and summarizing customer reviews. In *Proc. of SIGKDD*, 2004.
- [13] O. Irsoy and C. Cardie. Deep recursive neural networks for compositionality in language. In *NIPS*, 2014.
- [14] L. Jiang, M. Yu, M. Zhou, X. Liu, and T. Zhao. Target-dependent twitter sentiment classification. In *Proc. of ACL*, 2011.
- [15] Y. Kim. Convolutional neural networks for sentence classification. 2014.
- [16] D. Kingma and J. Ba. Adam: A method for stochastic optimization. *ICLR*, 2015.
- [17] N. Kobayashi, R. Iida, K. Inui, and Y. Matsumoto. Opinion mining on the web by extracting subject-aspect-evaluation relations. In *AAAI Spring Symposium: Computational Approaches to Analyzing Weblogs*, 2006.
- [18] Q. V. Le and T. Mikolov. Distributed representations of sentences and documents. In *ICML*, 2014.
- [19] F. Li, C. Han, M. Huang, X. Zhu, Y.-J. Xia, S. Zhang, and H. Yu. Structure-aware review mining and summarization. In *Proc. of ACL*, 2010.
- [20] P. Liu, S. Joty, and H. Meng. Fine-grained opinion mining with recurrent neural networks and word embeddings. In *EMNLP*, 2015.
- [21] D. Marcheggiani, O. Täckström, A. Esuli, and F. Sebastiani. Hierarchical multi-label conditional random fields for aspect-oriented opinion mining. In *ECIR*, 2014.
- [22] Q. Mei, X. Ling, M. Wondra, H. Su, and C. Zhai. Topic sentiment mixture: modeling facets and opinions in weblogs. In *Proc. of WWW*, 2007.
- [23] T. Mikolov, I. Sutskever, K. Chen, G. S. Corrado, and J. Dean. Distributed representations of words and phrases and their compositionality. In *NIPS*, 2013.
- [24] S. M. Mohammad, S. Kiritchenko, P. Sobhani, X. Zhu, and C. Cherry. Semeval-2016 task 6: Detecting stance in tweets. In *Proc. of SemEval*, 2016.
- [25] S. M. Mohammad, P. Sobhani, and S. Kiritchenko. Stance and sentiment in tweets. *arXiv preprint arXiv:1605.01655*, 2016.
- [26] S. M. Mohammad and P. D. Turney. Emotions evoked by common words and phrases: Using mechanical turk to create an emotion lexicon. In *Proc. of NAACL*, 2010.
- [27] B. Pang and L. Lee. Opinion mining and sentiment analysis. *Foundations and trends in information retrieval*, 2008.
- [28] M. Pontiki, D. Galanis, H. Papageorgiou, S. Manandhar, and I. Androutsopoulos. Semeval-2015 task 12: Aspect based sentiment analysis. In *Proc. of SemEval*, 2015.
- [29] M. Pontiki, D. Galanis, J. Pavlopoulos, H. Papageorgiou, I. Androutsopoulos, and S. Manandhar. Semeval-2014 task 4: Aspect based sentiment analysis. In *Proc. of SemEval*, 2014.
- [30] A. Popescu and M. Pennacchiotti. Dancing with the stars, nba games, politics: An exploration of twitter users' response to events. In *Proc. of ICWSM*, 2011.
- [31] A. Rajadesingan and H. Liu. Identifying users with opposing opinions in twitter debates. In *Social Computing, Behavioral-Cultural Modeling and Prediction*. 2014.
- [32] C. Sauper and R. Barzilay. Automatic aggregation by joint modeling of aspects and values. *JAIR*, 2013.
- [33] R. Socher, B. Huval, C. D. Manning, and A. Y. Ng. Semantic compositionality through recursive matrix-vector spaces. In *Proc. of EMNLP-CoNLL*, 2012.
- [34] R. Socher, A. Perelygin, J. Y. Wu, J. Chuang, C. D. Manning, A. Y. Ng, and C. Potts. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proc. of EMNLP*, 2013.
- [35] S. Sukhbaatar, J. Weston, R. Fergus, et al. End-to-end memory networks. In *NIPS*, 2015.
- [36] D. Tang, B. Qin, and T. Liu. Aspect level sentiment classification with deep memory network. In *EMNLP*, 2016.
- [37] D. Tang, F. Wei, N. Yang, M. Zhou, T. Liu, and B. Qin. Learning sentiment-specific word embedding for twitter sentiment classification. In *ACL*, 2014.
- [38] D.-T. Vo and Y. Zhang. Target-dependent twitter sentiment classification with rich automatic features. In *IJCAI*, 2015.
- [39] M. A. Walker, P. Anand, R. Abbott, and R. Grant. Stance classification using dialogic properties of persuasion. In *Proc. of NAACL*, 2012.
- [40] M. A. Walker, P. Anand, R. Abbott, J. E. F. Tree, C. Martell, and J. King. That is your evidence?: Classifying stance in online political debate. *Decision Support Systems*, 2012.
- [41] H. Wang, Y. Lu, and C. Zhai. Latent aspect rating analysis without aspect keyword supervision. In *Proc. of SIGKDD*, 2011.
- [42] S. Wang and C. D. Manning. Baselines and bigrams: Simple, good sentiment and topic classification. In *Proc. of ACL*, 2012.
- [43] T. Wilson, J. Wiebe, and P. Hoffmann. Recognizing contextual polarity in phrase-level sentiment analysis. In *Proc. of HLT/EMNLP*, 2005.
- [44] Y. Yang and X. Liu. A re-examination of text categorization methods. In *Proc. of SIGIR*, 1999.
- [45] M. Zhang, Y. Zhang, and D.-T. Vo. Gated neural networks for targeted sentiment analysis. In *Proc. of AAAI*, 2016.
- [46] C. Zirn, M. Niepert, H. Stuckenschmidt, and M. Strube. Fine-grained sentiment analysis with structural features. In *IJCNLP*, 2011.