

Variational Weakly Supervised Sentiment Analysis with Posterior Regularization

Ziqian Zeng, Yangqiu Song

Department of CSE, HKUST, Hong Kong, China

{zzengae, yqsong}@cse.ust.hk

Abstract

Sentiment analysis is an important task in natural language processing (NLP). Most of existing state-of-the-art methods are under the supervised learning paradigm. However, human annotations can be scarce. Thus, we should leverage more weak supervision for sentiment analysis. In this paper, we propose a posterior regularization framework for the variational approach to the weakly supervised sentiment analysis to better control the posterior distribution of the label assignment. The intuition behind the posterior regularization is that if extracted opinion words from two documents are semantically similar, the posterior distributions of two documents should be similar. Our experimental results show that the posterior regularization can improve the original variational approach to the weakly supervised sentiment analysis and the performance is more stable with smaller prediction variance.

1 Introduction

Sentiment analysis is a task of identifying the sentiment polarity expressed in textual data (Liu, 2012). Most state-of-the-art sentiment analysis methods in the literature are supervised methods which require many labeled training data. However, human annotations in the real world are scarce. While we assume there is abundant annotated data to train more and more complex models, there is still a need to consider weakly supervised methods that require less human annotation.

One way to perform weakly supervised sentiment analysis is using a predefined lexicon (Turney, 2002; Taboada et al., 2011). A lexicon consists of many opinion words. For each opinion word, its polarity (positive or negative) and strength (the degree to which the opinion word is positive or negative) are annotated by domain experts. lexicon-based weakly supervised methods perform a dictionary lookup and assign a polarity according to all

opinion words extracted from a document. A good lexicon requires high precision and high coverage, which needs a lot of human effort.

Another way to do weakly supervised sentiment analysis is using limited keywords (Meng et al., 2018; Zeng et al., 2019). Compared with lexicon-based methods, user-provided keywords require less human effort. Among keyword-based methods, there are two directions. First, (Meng et al., 2018) leveraged limited keywords to expand more keywords and generate pseudo-labeled data, and then performed self-training on real unlabeled data for model refinement. Possible improvements of this direction include investigating more advanced keywords expansion techniques to generate better pseudo-labeled samples (Miller et al., 2012) and developing more advanced self-training algorithms (Codon et al., 2014).

Second, the Variational Weakly Supervised (VWS) sentiment analysis (Zeng et al., 2019) used target-opinion word pairs as supervision signal. Its objective function is to predict an opinion word given a target word. For example, in a sentence “the room is big,” “room” is a target word and “big” is an opinion word. By introducing a latent variable (the sentiment polarity), they can learn a well-approximated posterior distribution via optimizing the evidence lower bound. The posterior probability here is the probability of a possible polarity (e.g., positive or negative) given a document, which is a typical sentiment classifier.

A potential issue with VWS is that optimizing the objective function may not guide the role of the latent variable to be sentiment polarity. For example, when half of reviews mention “big room” and half of reviews mention “small room,” the latent variable is possibly related to the size of rooms, but the expected role of the latent variable is the sentiment polarity of rooms. Hence how to control and regularize the posterior distribution is very im-

portant. One indirect way to control the posterior distribution is clever initialization (Ganchev et al., 2010). Originally, VWS aims to predict the sentiment polarity of each aspect for the multi-aspect sentiment analysis. So it uses the overall sentiment polarity to pretrain the model so that the posterior distribution has a good initialization. The overall polarity and polarity of each aspect are highly correlated. Thus, the initialization is highly likely to be similar to the true posterior distribution.

In this paper, we propose to use posterior regularization to regularize the VWS approach for sentiment analysis. There are two types of side information we can leverage to regularize the latent variable. First, calculating the similarity between keywords and extracted opinion words from a document can guide the model to decide which polarity that the document belongs to. Second, calculating the similarity between extracted opinion words from two documents can guide the model to decide whether two documents belong to the same polarity. The first type of side information is very easy to leverage, and it is reflected by our pretraining process. When a document is similar to keywords associated with a specific polarity, then we enforce that the posterior probability of a specific polarity should be large. In this case, in the pretraining process, we assign a pseudo label to the document. The second type of side information does not directly suggest which sentiment polarity that a document should be assigned to. It enforces pairwise constraints to the model. Our proposed posterior regularization leverage the second side information to ensure that when two documents are similar (dissimilar), the regularization enforces the posterior distribution of two documents to be similar (dissimilar).

Our contributions are summarized as follows,

- We develop a posterior regularization framework for the variational weakly supervised sentiment analysis.
- The experimental results show that the proposed regularization can improve the VWS model, make the results more stable, and outperform other weakly supervised baselines.

Our code is available at <https://github.com/HKUST-KnowComp/VWS-PR>.

2 Methodology

In this section, we first review the variational weakly supervised (VWS) sentiment analysis

method in Section 2.1. Then we introduce our posterior regularization in Section 2.2.

2.1 VWS Sentiment Analysis

Before formally introducing the VWS framework, we give a concrete example to illustrate how VWS works. Let \mathbf{x} be the representation of a document x , e.g., bag of words or feature outputs of neural networks. Let C be a random variable, indicating the sentiment polarity of a document. The possible value assignment of C can be *positive* or *negative*, or rating from 1 to 10. Suppose there is a document x where we extract an opinion word “terrific.” The objective function is to maximize the probability of opinion word “terrific.” By introducing a latent variable C , the objective function is split into two probabilities, corresponding to two classifiers, namely, sentiment polarity classifier and opinion word classifier. The input of sentiment classifier is the document representation \mathbf{x} , and it produces a probability distribution of sentiment polarity, i.e., $p(C = \textit{positive}|\mathbf{x})$ and $p(C = \textit{negative}|\mathbf{x})$. The input of opinion word classifier is extracted opinion words and the estimated sentiment polarity distribution, and it produces a probability distribution of opinion word given estimated sentiment polarity distribution, i.e., $p(\textit{“terrific”}|C = \textit{positive})$ and $p(\textit{“terrific”}|C = \textit{negative})$.

2.1.1 Sentiment Polarity Classifier

The sentiment polarity classifier aims to estimate a distribution $q(C|\mathbf{x})$, where C is a discrete random variable representing the sentiment polarity of a document. Let c denote a possible value of the random variable C , representing a possible value of sentiment polarity, e.g., *positive* or *negative*. The sentiment classifier estimates the probability as

$$q(C = c|\mathbf{x}) = \frac{\exp(\mathbf{w}_c^T \mathbf{x})}{\sum_{c'} \exp(\mathbf{w}_{c'}^T \mathbf{x})}, \quad (1)$$

$$\mathbf{x} = \text{CNN}(x), \quad (2)$$

where \mathbf{w}_c is a trainable vector associated with a sentiment polarity c , x is a document, and \mathbf{x} is the document representation. The representation of a document \mathbf{x} can be various. We use Convolutional Neural Network (CNN) in the experiment.

2.1.2 Opinion Word Classifier

The opinion word classifier aims to estimate the probability of an opinion word w_o given a possible

value of sentiment polarity c :

$$p(w_o|c) = \frac{\exp(\varphi(w_o, c))}{\sum_{w'_o} \exp(\varphi(w'_o, c))}, \quad (3)$$

where $\varphi(\cdot)$ is a scoring function taking opinion word w_o and a possible value of sentiment polarity c as inputs. The nature of the scoring function is about the frequency of occurrence. If an opinion word and a possible value of sentiment polarity co-occur frequently, the score will be high, otherwise, it will be low. Specifically, we define:

$$\varphi(w_o, c) = \mathbf{a}_c^T \mathbf{w}_o, \quad (4)$$

where \mathbf{w}_o is the trainable word embedding of opinion word w_o , \mathbf{a}_c is a trainable vector associated with c . The scoring function can be various, e.g., multilayer perceptron (MLP). Here we only introduce the simplest case.

Given a possible value of sentiment polarity c , VWS aims to maximize the probability of opinion words that frequently occurred with c . For example, the opinion word “good” is usually occurred with sentiment polarity *positive*, and the opinion word “terrible” is usually occurred with sentiment polarity *negative*.

2.1.3 Training Objective

The objective function of VWS is to maximize the log-likelihood of an opinion word w_o . After introducing a latent variable (i.e., the sentiment polarity of a document) to the objective function, we can derive a variational lower bound of the log-likelihood which can incorporate two classifiers. The first one corresponds to the sentiment classifier. The second one corresponds to the opinion word classifier. The variational lower bound of log-likelihood is shown as follows:

$$\begin{aligned} \mathcal{L}_1 &= \sum_{x \in X} \sum_{w_o \in \mathcal{P}_x} \log p(w_o) \\ &= \sum_{x \in X} \sum_{w_o \in \mathcal{P}_x} \log \sum_c p(w_o, c) \\ &= \sum_{x \in X} \sum_{w_o \in \mathcal{P}_x} \log \sum_c q(c|\mathbf{x}) \left[\frac{p(w_o, c)}{q(c|\mathbf{x})} \right] \\ &\geq \sum_{x \in X} \sum_{w_o \in \mathcal{P}_x} \sum_c q(c|\mathbf{x}) \left[\log \frac{p(w_o, c)}{q(c|\mathbf{x})} \right] \\ &= \sum_{x \in X} \sum_{w_o \in \mathcal{P}_x} \mathbb{E}_{q(C|\mathbf{x})} [\log p(w_o|c)p(c)] \\ &+ \sum_{x \in X} \sum_{w_o \in \mathcal{P}_x} H(q(C|\mathbf{x})), \end{aligned} \quad (5)$$

where X is the training set containing all documents, and \mathcal{P}_x is the set of all opinion words extracted from a document x , $H(\cdot)$ refers to the Shannon entropy, and $q(c|\mathbf{x})$ is short for $q(C = c|\mathbf{x})$. By applying Jensen’s inequality, the log-likelihood is lower-bounded by Eq. (5). The equality holds if and only if the KL-divergence of two distributions, $q(C|\mathbf{x})$ and $p(C|w_o)$, equals to zero. Maximizing the evidence lower bound is equivalent to minimizing the KL-divergence. Hence, VWS can learn a sentiment classifier that can produce a similar distribution to the true posterior $p(C|w_o)$. We assume that the training set is perfectly balanced, which means the prior distribution of sentiment polarity, i.e., $p(C)$, is a uniform distribution. Hence, $p(c)$ is a constant, which can be ignored.

2.1.4 Approximation

The partition function in Eq. (3) requires the summation over all opinion words in the vocabulary. Since the size of the opinion word vocabulary is large, VWS uses the negative sampling technique (Mikolov et al., 2013) to approximate Eq. (3). Specifically, VWS approximates $p(w_o|c)$ in the objective (3) with the following objective function:

$$\log \sigma(\varphi(w_o, c)) + \sum_{w'_o \in \mathcal{N}} \log (1 - \sigma(\varphi(w'_o, c))), \quad (6)$$

where w'_o is a negative sample in opinion words vocabulary, \mathcal{N} is the set of negative samples and $\sigma(\cdot)$ is the sigmoid function. In order to ensure that the approximation part and the entropy term are on the same scale (Marcheggiani and Titov, 2016), a hyper-parameter α is added to the entropy term. The objective function becomes:

$$\begin{aligned} \mathcal{L}_2 &= \sum_{x \in X} \sum_{w_o \in \mathcal{P}_x} \mathbb{E}_{q(C|\mathbf{x})} [\log \sigma(\varphi(w_o, c))] \\ &+ \sum_{w'_o \in \mathcal{N}} \log (1 - \sigma(\varphi(w'_o, c))) + \log p(c)] \\ &+ \sum_{x \in X} \sum_{w_o \in \mathcal{P}_x} \alpha H(q(C|\mathbf{x})). \end{aligned} \quad (7)$$

2.2 Posterior Regularization

As pointed out by (Ganchev et al., 2010), controlling the posterior distribution is crucial for models that estimate posterior distribution by maximizing the likelihood of the observed data via marginalizing over the latent variables. We need side information to regularize the posterior distribution. The

side information we leveraged is that if the opinion words extracted from two documents are similar semantically, then these two documents probably are in the same class, and if the opinion words are opposite semantically, then these two documents are probably not in the same class. For example, if one document x^i contains opinion words “great” and “awesome,” another document x^j contains opinion words “great” and “excellent,” and another document x^k contains opinion words “awful” and “terrible,” it is highly possible that x^i and x^j belong to the same class because their extracted opinion words are similar semantically, and x^i and x^k do not belong to the same class because their extracted opinion words are opposite semantically.

We formulate our posterior regularization as:

$$\mathcal{R}(x^i, x^j) = -d(x^i, x^j) \cdot \mathcal{S}(\mathcal{O}(x^i), \mathcal{O}(x^j)), \quad (8)$$

where $d(x^i, x^j)$ is short for $d(q(C|\mathbf{x}^i), q(C|\mathbf{x}^j))$, meaning the distance of two posterior distributions, and we use Euclidean distance metric; $\mathcal{S}(\cdot, \cdot)$ is a score function which measures the similarity or dissimilarity between two sets of opinion words; $\mathcal{O}(x^i)$ represents all opinion words extracted from a document x^i . We finally maximize Eq. (8) in the objective function. When $\mathcal{S}(\mathcal{O}(x^i), \mathcal{O}(x^j))$ is positive (suggesting similar), this regularization enforces the distance to be small, and when $\mathcal{S}(\mathcal{O}(x^i), \mathcal{O}(x^j))$ is negative (suggesting dissimilar), this regularization enforces the distance to be large. When $\mathcal{S}(\mathcal{O}(x^i), \mathcal{O}(x^j))$ is zero, it suggests comparison between opinion words cannot decide whether two documents are similar or not.

Next, we will introduce the scoring function $\mathcal{S}(\cdot, \cdot)$. Suppose a document x^i contains a set of opinion words $\mathcal{O}(x^i) = \{w_{o_1}^i, w_{o_2}^i, \dots, w_{o_k}^i\}$ and a document x^j contains a set of opinion words $\mathcal{O}(x^j) = \{w_{o_1}^j, w_{o_2}^j, \dots, w_{o_k}^j\}$. We define an operation $\text{cos}(\mathcal{O}(x^i), \mathcal{O}(x^j))$ over two sets of opinion words. It will return all cosine similarity values of all valid opinion word pairs where one word must come from $\mathcal{O}(x^i)$ and the other must come from $\mathcal{O}(x^j)$. We represent opinion words using embeddings, i.e., the embeddings in the opinion word classifier in Eq. (4). If there are k opinion words in each set, $\text{cos}(\mathcal{O}(x^i), \mathcal{O}(x^j))$ will return $\frac{k*(k-1)}{2}$ cosine similarity values. When we want to know whether two documents are similar in opinion words, we pay attention to the maximum value, i.e., $\text{max_cos} = \text{max}(\text{cos}(\mathcal{O}(x^i), \mathcal{O}(x^j)))$. When we want to know whether two documents

are dissimilar in opinion words, we pay attention to the minimum value, i.e., $\text{min_cos} = \text{min}(\text{cos}(\mathcal{O}(x^i), \mathcal{O}(x^j)))$. So we define:

$$\mathcal{S}(\cdot, \cdot) = \begin{cases} \text{max_cos}, & \text{max_cos} > \gamma_1 \quad \& \\ & \text{min_cos} \geq \gamma_2, \\ \text{min_cos}, & \text{max_cos} \leq \gamma_1 \quad \& \\ & \text{min_cos} < \gamma_2, \\ \delta, & \text{max_cos} > \gamma_1 \quad \& \\ & \text{min_cos} < \gamma_2, \\ 0, & \text{otherwise}, \end{cases} \quad (9)$$

where $\mathcal{S}(\cdot, \cdot)$ is short for $\mathcal{S}(\mathcal{O}(x^i), \mathcal{O}(x^j))$ due to space limit. The first condition means two documents have some semantically similar opinion words ($\text{max_cos} > \gamma_1$) and have no semantically dissimilar opinion words ($\text{min_cos} \geq \gamma_2$). The value returned by the function score is max_cos . It should be a positive value. The second condition means two documents have no semantically similar opinion words ($\text{max_cos} \leq \gamma_1$) and have some semantically dissimilar opinion words ($\text{min_cos} < \gamma_2$). The value returned by function score is min_cos . It should be a negative number. The third condition means two documents have some semantically similar opinion words ($\text{max_cos} > \gamma_1$) and also have some semantically dissimilar opinion words ($\text{min_cos} < \gamma_2$). This condition corresponds to a real-world situation that when some customers want to express some negative sentiment, they usually point out some positive aspects first, and then start with a “but”, and emphasize some negative aspects. The opinion words sets extracted from these type of documents have both negative and positive opinion words. When we compare two of them, they will have some similar opinion words and dissimilar opinion words. In this case, we tend to assume they are in the same class. If the third condition is satisfied, it will return a non-negative value $\delta \in [0, 1]$; The final condition means two documents have no semantically similar opinion words ($\text{max_cos} \leq \gamma_1$) and have no semantically dissimilar opinion words ($\text{min_cos} \geq \gamma_2$). It will return 0.

The mechanism of the regularization is that if the posterior distributions $q(C|\mathbf{x}^i)$ and $q(C|\mathbf{x}^j)$ are different from each other, i.e., $d(q(C|\mathbf{x}^i), q(C|\mathbf{x}^j))$ is large, but opinion words suggest that these two documents should be in the same cluster i.e., $s(w_d^i, w_d^j)$

Datasets	Yelp	IMDB	Amazon
Size of training set	38,000	20,000	20,000
Size of development set	3,800	2,000	2,000
Average doc length	155.12	252.90	88.28
Opinion vocabulary size	1,097	688	394

Table 1: Statistics of Yelp, IMDB, and Amazon dataset.

is large, then $d(q(C|\mathbf{x}^i), q(C|\mathbf{x}^j))$ will be encouraged to be small by applying the regularization. Oppositely, if the posterior distributions $q(C|\mathbf{x}^i)$ and $q(C|\mathbf{x}^j)$ are similar, i.e., $d(q(C|\mathbf{x}^i), q(C|\mathbf{x}^j))$ is small, but opinion words suggest that these two documents should be in the different cluster i.e., $s(w_d^i, w_d^j)$ is small, then $d(q(C|\mathbf{x}^i), q(C|\mathbf{x}^j))$ will be encouraged to be large by applying the regularization.

The final objective function with posterior regularization is as follows,

$$\mathcal{J} = \mathcal{L}_2 + \beta \sum_{x^i \in X} \sum_{x^j \in X} R(x^i, x^j), \quad (10)$$

where x^i and x^j are documents in the training set X . The constraints are defined in a $|X| \times |X|$ space. In practice, we train our model batch by batch. So we only apply the constraints within a mini-batch. There are at most $|X_b| \times |X_b|$ constraints in a mini-batch, where $|X_b|$ is the number of samples in a mini-batch.

3 Experiments

In this section, we evaluate the empirical performance of our method on binary sentiment classification tasks.

3.1 Datasets

We use three corpora to evaluate the performance of our proposed method. All corpora have two classes and perfectly balanced. For all methods, we use a development set for hyper-parameter tuning. For all methods, we use the training set as the test set since all methods do not use the ground truth in the training set.

(1) **Yelp Review**: We use the Yelp reviews polarity dataset from (Zhang et al., 2015) and take its test set containing 38,000 documents as the corpus for evaluation. For hyper-parameter tuning, we also extract 3,800 documents from the original training set of (Zhang et al., 2015) to serve as a development set.

(2) **IMDB Review**: We use the IMDB reviews polarity dataset from (Maas et al., 2011) and randomly extract 20,000 reviews from its original test set as the corpus for evaluation. For hyper-parameter tuning, we also extract 2,000 documents from the original training set of (Maas et al., 2011) to serve as a development set.

(3) **Amazon Review**: We use the Amazon reviews polarity dataset from (Zhang et al., 2015) and randomly extracted 20,000 reviews from its original test set as the corpus for evaluation. For hyper-parameter tuning, we also extract 2,000 documents from the original training set of (Zhang et al., 2015) to serve as a development set.

Table 1 provides the details of these datasets.

3.2 Compared Methods

Lexicon uses an opinion lexicon to assign sentiment polarity to a document (Read and Carroll, 2009; Pablos et al., 2015). We combine two popular opinion lexicons used by (Hu and Liu, 2004) and (Wilson et al., 2005) to get a larger lexicon. If an extracted opinion is in the positive (negative) lexicon, it votes for positive (negative). When the opinion word is with a negation word such as “no” and “not”, its polarity will be the opposite. Then, the polarity of a document is determined by using majority voting among all extracted opinion words. When the number of positive and negative words is equal, the document will be randomly assigned a polarity.

WeSTClass (Meng et al., 2018) first generates pseudo labels for documents which contain user-provided keywords. Keywords are expanded to generate more pseudo samples. It pretrains a CNN/LSTM model using pseudo samples as the training set and then performs a self-training process. Here, we use CNN because it empirically outperforms LSTM. The CNN architecture we used here is the same as the one described in (Meng et al., 2018).

Keyword Pretrain generates pseudo labels for

Datasets	Positive			Negative		
Yelp	terrific	amazing	awesome	horrible	worst	bad
IMDB	great	fantastic	awesome	awful	worst	bad
Amazon	great	fantastic	awesome	poor	worst	bad

Table 2: Keywords of Yelp, IMDB, and Amazon datasets.

documents in training set which contain user-provided keywords. We pretrain a CNN model using pseudo samples as the training set. In the CNN model, four different filter sizes $\{2, 3, 4, 5\}$ are applied, and a max-pooling layer is applied to each convolutional layer, and each convolutional layer has 100 filters.

VWS (Zeng et al., 2019) uses target opinion word pairs as supervision signal. It trains a sentiment polarity classifier and opinion word classifier simultaneously via optimizing the variational lower bound. We use a CNN model as the sentiment polarity classifier. And we pretrain it by generating pseudo labels for documents which contain user-provided keywords in the training set. The CNN architecture is the same as the one in **Keyword Pretrain**.

VWS-PR is VWS method with proposed posterior regularization.

3.3 Keywords and Opinion Word Extraction

We manually select three keywords for each class. The details of keywords of three datasets are shown in table 2.

For opinion word extraction, we adopt four rules proposed by VWS (Zeng et al., 2019) in the implementation. All rules rely on dependency parser (Chen and Manning, 2014). When a target word and an opinion word satisfy a dependency relation, we will extract the opinion word. The details of dependency relation and examples are provided in Table 3. When a pair of words satisfy one rule, there are still some restrictions on head and tails words to be satisfied. There is no restriction for Rule 1. For Rule 2, the head word should be an adjective and the tail word should a noun. For Rule 3, the head word should be one of the following four words: “like,” “dislike,” “love,” and “hate.” For Rule 4, the head word should be one of the following word: “seem,” “look,” “feel,” “smell,” and “taste.”

3.4 Result Analysis

Table 4 shows that our method VWS-PR outperforms VWS by 4%, 2%, and 1% on Yelp, IMDB, and Amazon datasets respectively. Compared with WeSTClass and VWS, our method is more stable, i.e., smaller standard deviation, which shows that the regularization confine the posterior distribution to a smaller space. The performance of lexicon method is bad across three datasets. The main reason is that it does not involve any learning process. Keyword pretraining method can outperform lexicon method. But pseudo labels are not ground truths, hence the pseudo training set contains noises. Also, user provided keywords are limited, so the training samples with pseudo labels are restricted to some samples which contain certain keywords. For example, reviews with an extreme polarity (only expressing positive polarity or only express negative polarity) are likely to be pseudo samples. But most of reviews express mixed polarities. This will hinder the generalization ability. WeSTClass outperforms the keyword pretraining method on Yelp and Amazon dataset due to keyword expansion and self-training process. But in IMDB dataset, WeSTClass is slightly worse than the keyword pretraining method. Possible reason would be keyword expansion involve some harmful keywords and self-training procedure amplifies errors. VWS outperforms WeSTClass on IMDB and Amazon datasets and is comparable to Yelp dataset.

3.5 Hyper-parameters Sensitivity Analysis

We first show F1 scores on three datasets with varied β in Figure 1(a). It shows that optimal β values of our model on three datasets are different. When they achieve optimal β value, the standard deviation is much smaller than others. The regularization makes models more stable. Our method on IMDB and Amazon is not very sensitive. The changes are within 2%. Our method on Yelp is more sensitive. But we could still find a range, e.g., 0.1 to 0.5, where the changes are within 2%. When β keeps growing, the performance in Yelp deteriorates a lot.

Rule	Dependency Relation	Example	Extracted Word
1	adjectival modifier	they have delicious food	delicious
2	nominal subject	the room is big	big
3	direct object	i like it	like
4	open clausal complement	i feel comfortable	comfortable

Table 3: Opinion words extraction rules.

Methods	Yelp		IMDB		Amazon	
	Mean	Std	Mean	Std	Mean	Std
Lexicon	0.5982	0.0006	0.5998	0.0013	0.5754	0.0014
Keyword Pretraining	0.7441	0.0060	0.7496	0.0050	0.6375	0.0171
WeSTClass (Meng et al., 2018)	0.8061	0.0105	0.7354	0.0096	0.7374	0.0082
VWS (Zeng et al., 2019)	0.8014	0.0179	0.7825	0.0045	0.7530	0.0025
VWS-PR	0.8431	0.0013	0.8025	0.0041	0.7644	0.0020

Table 4: F1 scores of weakly supervised sentiment analysis methods on Yelp, IMDB, and Amazon.

The main reason is probably that the opinion word vocabulary size in Yelp is much larger than other datasets, and hence it is likely that the vocabulary in Yelp contains more noisy opinion words. When β is large, the noises may harm the performance.

We then show F1 scores on three datasets with varied γ_1 in Figure 1(b). The optimal γ_1 values of three datasets are the same, i.e., 0.7. The trends are consistent on three datasets. F1 score first increases and then decreases. When γ_1 is small, the constraints are easier to satisfy and the performance is bad because it may involve more noisy constraints. It probably enforces two samples to be similar but in reality, they are not similar. When γ_1 is large, the performance is also bad because it has fewer constraints that enforce two samples to be similar. It could have made use of more constraints.

We show F1 scores on three datasets with varied γ_2 in Figure 1(c). The optimal γ_2 values of three datasets are the same, i.e., -0.1 . The trends are consistent on three datasets. F1 score first increases and then decreases. When γ_2 is small, the performance is bad because it has fewer constraints that enforce two samples to be dissimilar. It could have made use of more constraints. When γ_2 is large, the performance is bad because it has more noisy constraints.


3.6 Error Analysis

We show some incorrectly predicted documents by VWS-PR on three datasets in Table 5. For the

first document, the customer emphasizes price a lot. Our method cannot extract opinion words on snippets such as “not at that price” and “for half the price.” For the second document, the reviewer loves this movie because he/she loves the basketball player. The reviewer thinks that the movie itself does not deserve a high score. Our method detects both positive and negative polarities on this document, so it tends to predict as negative. Because most mixed polarities are likely to be negative polarity. The regularization enforces this pattern. This document obviously is different from other documents with mixed opinion words. For the last document, our method cannot extract opinion words on snippets such as “no wrist strap” and “without a place to attach a wrist strap.” Our method is good at extracting words on subjective expression such as “nice light,” but not on descriptive expression such as “no wrist strap.” Our method fails because no other knowledge source indicates that “no wrist strap” is negative.

3.7 Implementation Details

For WeSTClass and VWS, we used code released by (Meng et al., 2018) and (Zeng et al., 2019) respectively, and followed their preprocessing steps and optimal settings. For VWS and VWS-PR, we pretrain a CNN model using pseudo-labeled samples. After that, the embeddings are untrainable. The rest of parameters are trainable. For our method, the hyperparameter settings of VWS part



(a) Varying β with fixed γ_1 and γ_2 . (b) Varying γ_1 with fixed β and γ_2 . (c) Varying γ_2 with fixed β and γ_1 .

Figure 1: Hyper-parameters sensitivity analysis.

Dataset	Document	Prediction	Ground Truth	Extracted Words
Yelp	It was good, but not at that price. There are so many other good italian places in the area for half the price.	good	bad	good
IMDB	I love gheorghe muresan, so i automatically loved this movie. Everything else about it was so so. Billy crystal is a good actor, even if he is annoying.	bad	good	love good annoying
Amazon	No wrist strap. Nice light, well made. But why would anyone design a tactical style flashlight without a place to attach a wrist strap.	good	bad	nice

Table 5: Documents that are predicted incorrectly by VWS-PR.

is the same as described in (Zeng et al., 2019). We implemented our models using TensorFlow (Abadi et al., 2016). When tuning hyper-parameters of regularization term, we perform grid search on $\gamma_1 \in [0.5, 0.6, \dots, 1.0]$ and $\gamma_2 \in [-0.5, -0.4, \dots, 0]$. After than, we fix γ_1 and γ_2 , then tune $\beta \in [0.1, 0.2, \dots, 1.0]$. δ is fixed to 1.

4 Related Work

In this section, we review the related work on weakly supervised sentiment analysis.

Using a lexicon is a typical way to perform weakly supervised sentiment analysis. One line of works perform simple assignment, i.e., majority voting, based on sentiment orientation scores of extracted opinion words. Some methods (Missen and Boughanem, 2009; Tsytarau et al., 2010) used sentiment orientation scores in existing lexicons directly, and aggregated them within a document to

determine polarity. Some methods developed their own semantic orientation estimation algorithm. For example, (Turney, 2002) first identified phrases in the review and then estimated the semantic orientation of each extracted phrases. The semantic orientation of a given phrase is calculated by comparing its similarity to a positive reference word (“excellent”) with its similarity to a negative reference word (“poor”). This method determined the sentiment polarity based on the average semantic orientation of the phrases extracted from the review. (Kamps et al., 2004) used the minimum path distance between a phrase and pivot words (“good” and “bad”) in WordNet to estimate the semantic orientation of extracted phrases.

Another line of works involve learning process when using a lexicon. (Li et al., 2009; Zhou et al., 2014) proposed a constrained non-negative matrix tri-factorization approach to sentiment anal-

ysis, and used a sentiment lexicon as prior knowledge. In these models, a term-document matrix is approximated by three factors that specify soft membership of terms and documents in one of k classes. All three factors are non-negative matrices. The first factor is a matrix representing knowledge in the word space, i.e., each row represents the posterior probability of a word belonging to the k classes. The second factor is a matrix providing a condensed view of the term-document matrix. The third factor is a matrix representing knowledge in document space, i.e., each row represents the posterior probability of a document belonging to the k classes. (Li et al., 2009) applied a regularization to encourage that the first factor is close to prior knowledge. This regularization is different from ours because it requires prior knowledge such as a predefined lexicon. A predefined lexicon needs a lot of human effort. (Zhou et al., 2014) applied a regularization based on an intuition that if two documents are sufficiently close to each other, they tend to share the same sentiment polarity. This intuition of this regularization is similar to ours. But when they compare document similarity, they use textual similarity (e.g., cosine similarity of bag of words) rather than similarity on opinion words. The regularization is applied under matrix factorization framework, it is not straightforward to fit in neural network based models.

Using keywords is another way to perform weakly supervised sentiment analysis. (Meng et al., 2018) leveraged keywords to generate pseudo-labeled samples for model pretraining, and then performed self-training on unlabeled data for model refinement. The possible improvement of this direction would be investigating more advanced keywords expansion techniques to generate better pseudo-labeled samples and developing a more advanced self-training algorithm. LOTClass (Meng et al., 2020), a parallel work to ours, fine-tuned a masked language model to generate relevant words that can replace label name such as “good” and “bad,” and performed self-training on unlabeled data for model refinement. Fine-tuning in LOTClass can be viewed as an advanced keyword expansion process using language models. VWS (Zeng et al., 2019) used target-opinion word pairs as supervision signal. Its objective function is to predict an opinion word given a target word. By introducing a latent variable (the sentiment polarity), they can learn a well-approximated posterior dis-

tribution via optimizing the evidence lower bound. The posterior probability here is the probability of a possible polarity (e.g., positive or negative) given text representation.

5 Conclusion

We propose a posterior regularization framework for the VWS sentiment analysis to better control the posterior distribution. The intuition behind the posterior regularization is that if extracted opinion words from two documents are semantically similar (dissimilar), the posterior distribution of two documents should be similar (dissimilar). Our experiments show that our posterior regularization can improve VWS and the performance is more stable.

Acknowledgments

We thank anonymous reviewers for their valuable comments. This paper was supported by the NSFC Grant U20B2053 from China, the Early Career Scheme (ECS, No. 26206717), the General Research Fund (GRF, No. 16211520), and the Research Impact Fund (RIF, No. R6020-19) from the Research Grants Council (RGC) of Hong Kong.

References

- Martin Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, Manjunath Kudlur, Josh Levenberg, Rajat Monga, Sherry Moore, Derek G. Murray, Benoit Steiner, Paul Tucker, Vijay Vasudevan, Pete Warden, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. 2016. Tensorflow: A system for large-scale machine learning. In *Proceedings of OSDI*, pages 265–283.
- Danqi Chen and Christopher Manning. 2014. A fast and accurate dependency parser using neural networks. In *Proceedings of EMNLP*, pages 740–750.
- Anni Coden, Dan Gruhl, Neal Lewis, Pablo N Mendes, Meena Nagarajan, Cartic Ramakrishnan, and Steve Welch. 2014. Semantic lexicon expansion for concept-based aspect-aware sentiment analysis. In *Semantic Web Evaluation Challenge*, pages 34–40. Springer.
- Kuzman Ganchev, Joao Gracca, Jennifer Gillenwater, and Ben Taskar. 2010. Posterior regularization for structured latent variable models. *The Journal of Machine Learning Research*, 11:2001–2049.
- Minqing Hu and Bing Liu. 2004. Mining and summarizing customer reviews. In *Proceedings of SIGKDD*, pages 168–177.

- Jaap Kamps, Maarten Marx, Robert J Mokken, Maarten De Rijke, et al. 2004. Using wordnet to measure semantic orientations of adjectives. In *Proceedings of LREC*, pages 1115–1118.
- Tao Li, Yi Zhang, and Vikas Sindhwani. 2009. A non-negative matrix tri-factorization approach to sentiment classification with lexical prior knowledge. In *Proceedings of ACL*, pages 244–252.
- Bing Liu. 2012. Sentiment analysis and opinion mining. *Synthesis lectures on human language technologies*, 5(1):1–167.
- Andrew L. Maas, Raymond E. Daly, Peter T. Pham, Dan Huang, Andrew Y. Ng, and Christopher Potts. 2011. Learning word vectors for sentiment analysis. In *ACL-HLT*, pages 142–150.
- Diego Marcheggiani and Ivan Titov. 2016. Discrete-state variational autoencoders for joint discovery and factorization of relations. *Transactions of the Association for Computational Linguistics*, 4:231–244.
- Yu Meng, Jiaming Shen, Chao Zhang, and Jiawei Han. 2018. Weakly-supervised neural text classification. In *Proceedings of CIKM*, pages 983–992.
- Yu Meng, Yunyi Zhang, Jiaxin Huang, Chenyan Xiong, Heng Ji, Chao Zhang, and Jiawei Han. 2020. Text classification using label names only: A language model self-training approach. In *Proceedings of EMNLP*, pages 9006–9017.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Proceedings of NeurIPS*, pages 3111–3119.
- Tristan Miller, Chris Biemann, Torsten Zesch, and Iryna Gurevych. 2012. Using distributional similarity for lexical expansion in knowledge-based word sense disambiguation. In *Proceedings of COLING*, pages 1781–1796.
- Malik Muhammad Saad Missen and Mohand Boughanem. 2009. Using wordnet’s semantic relations for opinion detection in blogs. In *Proceedings of ECIR*, pages 729–733.
- Aitor García Pablos, Montse Cuadros, and German Rigau. 2015. V3: Unsupervised aspect based sentiment analysis for semeval2015 task 12. In *Proceedings of SemEval*, pages 714–718.
- Jonathon Read and John Carroll. 2009. Weakly supervised techniques for domain-independent sentiment classification. In *Proceedings of CIKM workshop on Topic-sentiment Analysis for Mass Opinion*, pages 45–52.
- Maite Taboada, Julian Brooke, Milan Tofiloski, Kimberly D. Voll, and Manfred Stede. 2011. Lexicon-based methods for sentiment analysis. *Computational Linguistics*, 37(2):267–307.
- Mikalai Tsytsarau, Themis Palpanas, and Kerstin De-necke. 2010. Scalable discovery of contradictions on the web. In *Proceedings of WWW*, pages 1195–1196.
- Peter D Turney. 2002. Thumbs up or thumbs down?: semantic orientation applied to unsupervised classification of reviews. In *Proceedings of ACL*, pages 417–424.
- Theresa Wilson, Janyce Wiebe, and Paul Hoffmann. 2005. Recognizing contextual polarity in phrase-level sentiment analysis. In *Proceedings of EMNLP-HLT*, pages 347–354.
- Ziqian Zeng, Wenxuan Zhou, Xin Liu, and Yangqiu Song. 2019. A variational approach to weakly supervised document-level multi-aspect sentiment classification. In *Proceedings of NAACL-HLT*, pages 386–396.
- Xiang Zhang, Junbo Zhao, and Yann LeCun. 2015. Character-level convolutional networks for text classification. In *Proceedings of NeurIP*, pages 649–657.
- Guangyou Zhou, Jun Zhao, and Daojian Zeng. 2014. Sentiment classification with graph co-regularization. In *Proceedings of COLING*, pages 1331–1340.