

A Joint Model for Topic-Sentiment Modeling from Text

Mohamed Dermouche
AMI Software R&D, and
Université de Lyon
(ERIC LYON 2), France
mde@amisw.com

Leila Kouas
AMI Software R&D
France
lkh@amisw.com

Julien Velcin, Sabine Loudcher
Université de Lyon
(ERIC LYON 2), France
julien.velcin@univ-lyon2.fr
sabine.loudcher@univ-lyon2.fr

ABSTRACT

Traditional topic models, like LDA and PLSA, have been efficiently extended to capture further aspects of text in addition to the latent topics (e.g., time evolution, sentiment etc.). In this paper, we discuss the issue of joint topic-sentiment modeling. We propose a novel topic model for topic-specific sentiment modeling from text and we derive an inference algorithm based on the Gibbs sampling process. We also propose a method for automatically setting the model parameters. The experiments performed on two review datasets show that our model outperforms other state-of-the-art models, in particular for sentiment prediction at the topic level.

Categories and Subject Descriptors

H.3.3 [Information Systems]: Information Storage and Retrieval; G.3 [Mathematics of Computing]: Probability and Statistics—*Probabilistic Algorithms*

General Terms

Algorithm, opinion mining, sentiment analysis.

Keywords

Joint topic sentiment models, topic models, sentiment analysis, opinion mining.

1. INTRODUCTION

A major research issue in text mining today is the joint extraction of latent topics and sentiments from text. The two text aspects, yet tightly related to each other, have been mainly treated separately. In the recent years, several works mainly based on probabilistic topic models have been proposed to tackle this issue [7, 9, 10, 11, 12]. In this paper, we discuss the issue of joint topic and sentiment modeling, in particular, we focus on extracting topic-relative sentiments, i.e., the sentiment proportions specific to a topic at hand.

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. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SAC'15 April 13-17, 2015, Salamanca, Spain.

Copyright 2015 ACM 978-1-4503-3196-8/15/04

<http://dx.doi.org/10.1145/2695664.2695726> ...\$15.00.

In the literature, the problem is usually treated as “top-down”, i.e., topics are learnt once the sentiment polarity has been stated. The approach can be viewed as a hierarchy where the first level represent the sentiment polarities (e.g., positive and negative) and the second level represent the topics. This strategy has many limitations that we summarize in the following: (i) topics are definitely assigned with sentiment polarities whereas it is more appropriate to assign them with probability distributions over sentiment polarities, (ii) a post-processing is required in order to match topics from different polarities. In addition, the topic-sentiment correlations are document-specific, which prevents from having an overall overview of these correlations that may be useful for many applications (Web intelligence, reputation management etc.).

To address these issues, we propose the Topic-Sentiment (TS) model. TS model is built upon the state-of-art LDA model [1], and adopts a “bottom-up” approach, unlike the existing models. It has two main features that are not jointly addressed by the existing models: first, the topic-specific sentiment is learnt from all the documents at once, which allows to extract the overall topic-sentiment correlation. Second, no post-processing is needed to match the same topics from different polarities. Each topic has a dual distribution over words, one under each sentiment polarity.

We demonstrate the effectiveness of our model for extracting topics and topic-specific sentiments from Web data. As the main goal of TS model is the estimation of sentiments specific to topics, we use an evaluation framework based on the sentiment prediction at the topic level. TS model outperforms the state-of-art models on two product-review datasets. Moreover, and to facilitate the deployment of TS model for real-world applications, we propose a method to automatically set the model parameters. Our method achieves better performance compared to the well-known Maximum-Likelihood-based parameter estimation method.

2. RELATED WORK

Traditional topic models like LDA and PLSA have been efficiently deployed for topic discovery from text. The success of this approach has motivated the creation of numerous other models that extend these models in order to capture other text aspects, such as author community, time evolution, sentiment etc. In this line, several works have been designed to capture the topic-sentiment correlation from textual data [7, 9, 10, 11, 12].

JST [10] is an earlier topic-sentiment model that extends LDA with a new sentiment layer. Thus, the generation of

Table 1: A comparison of the proposed model with existing models.

Categ.	Model	Output	Topic distrib. over words	Topic polarity		Same topics under different polarities	Document polarity
				Type	Scope		
1	JST, ASUM, STDP	Topics under sentiments	Global	Definitive	Document-specific	No	Yes
2	TSM	Topics	Global	None	None	No	No
	Reverse-JST	Topic-sentiment pairs	For each sentiment	Probability distrib.	Document-specific	Yes	No
	TS	Topic-sentiment pairs	For each sentiment	Probability distrib.	Global	Yes	No

a word depends not only on the topic but also on the sentiment polarity. Reverse-JST [11] is a variant of JST where the topic and the sentiment layers are inverted. ASUM [7] is similar to JST but it works at the sentence level (all the words of a sentence are generated under the same topic). Sentiment-LDA [9] is similar to JST but it only handles a binary sentiment polarity. STDP [8] is another variant of JST where the sentiment polarity of a word depends on its part-of-speech category. TSM [12] is a topic-sentiment model built on the top of PLSA. The topic-sentiment correlation is captured through a post-processing that involves the topic-word and the word-sentiment distributions.

Table 1 summarizes the main features of these models compared to our model, without any post-processing. Existing models can be categorized into two categories based on the order of topic and sentiment layers. In the category #1, the generation of sentiment depends on the generated topic. All the models presented here, except Reverse-JST and TSM, fit into this category. The main limitations of these models can be summarized in the following:

1. A topic is definitely assigned with a sentiment polarity (positive or negative) whereas, in the real world, a topic may occur under different perspectives in the same dataset, even in the same document. For example, in [11], the authors use the Movie Review dataset and show a topic about Titanic movie extracted under the positive polarity. This gives the impression that the Titanic movie is always discussed positively in this dataset, while actually it is not.
2. No correspondence exists between topics from different polarities. For example, the topic about Titanic movie may also appear under the negative polarity, but with a different index. The existing models require a post-processing step in order to match the same topic occurring under different polarities.

In the category #2, the order of topic and sentiment layers is inverted. This strategy allows to overcome the limitation #1 above but in turn it does not capture the overall sentiment relative to a specific topic. The models from this category extract this information for each single document. Consequently, the generated output is unnecessarily detailed and consequently hard to visualize. For many practical applications, the overall sentiment relative to topic is a valuable information (e.g., Web intelligence, reputation management, CRM).

To address these issues, we propose a topic-sentiment model with a different structure. The graphical model, the generative process and the inference algorithm are presented in Section 3.

Table 2: Notation.

D	Number of documents
V	Vocabulary size
T	Number of topics
S	Number of sentiment labels
s	Sentiment labels
z	Topics
Θ	$[\theta_d]$: $D \times T$ matrix of document-specific distributions over topics
Φ	$[\varphi_{z,s}]$: $T \times S \times V$ matrix of topic-sentiment-specific distributions over words
Π	$[\pi_z]$: $T \times S$ matrix of topic-specific distribution over sentiments
n_d	Number of words in document d
$n_{d,j}$	Number of words in document d assigned with topic j
n_j	Number of words assigned with topic j
$n_{j,k}$	Number of words assigned with topic j and sentiment k
$n_{i,j,k}$	Number of times a word i is assigned with topic j and sentiment k
n^{-p}	Count variables excluding word at position p of the current document

3. TOPIC-SENTIMENT (TS) MODEL

As it has been shown in the previous section, number of topic models have been built on the top of the well-known LDA in order to capture topic-sentiment correlation. The idea behind our model is that topical words co-occur with sentiment words, which creates topic-sentiment correlation. Roughly speaking, the topics are characterized by the words that co-occur frequently in the same documents, while the topic-sentiment correlation is characterized by the co-occurrence frequency of topic words with sentiment-bearing words. TS model captures the strength of this correlation and allows to extract topics jointly with sentiments.

To capture the sentiment relative to topics, we extend LDA with a new sentiment layer, denoted s (see Table 2 for notation). The node corresponding to s is inserted after the node z because the sentiment depends on the topic at hand (cf. Figure 1). The generative process of TS model operates as follows:

1. Draw $T \times S$ multinomials $\varphi_{z,s} \sim \text{Dir}(\beta)$
2. Draw T multinomials $\pi_z \sim \text{Dir}(\gamma)$
3. For each document d , draw a multinomial $\theta_d \sim \text{Dir}(\alpha)$, then for each word w_i in d :
 - (a) Draw a topic $z_i \sim \theta_d$
 - (b) Draw a sentiment label $s_i \sim \pi_{z_i}$
 - (c) Draw a word $w_i \sim \varphi_{z_i,s_i}$

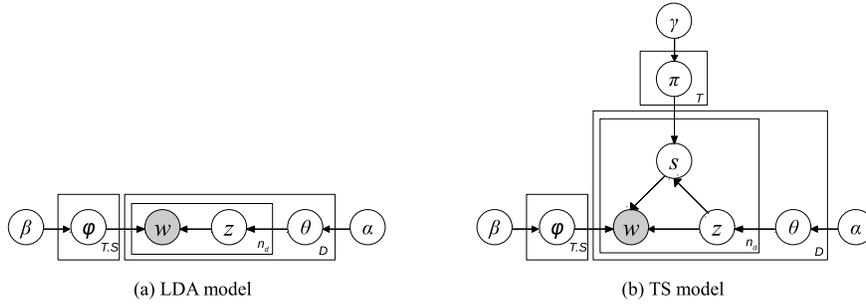


Figure 1: (a) LDA and (b) TS graphical models.

3.1 Inference

Gibbs sampling is a popular approach to parameter estimation (inference) in topic models [8, 11, 14]. We adopt this approach because it often yields relatively simple algorithms. Due to space limitation, we only give the final formulas (for details on Gibbs Sampling for topic models, see [5]).

Joint distribution.

Using Bayes rule, the joint probability of words, topics, and sentiments can be factored as follows:

$$p(\mathbf{w}, \mathbf{s}, \mathbf{z} | \alpha, \beta, \gamma) = p(\mathbf{w} | \mathbf{s}, \mathbf{z}, \beta) \cdot p(\mathbf{s} | \mathbf{z}, \gamma) \cdot p(\mathbf{z} | \alpha). \quad (1)$$

The first term is obtained by integrating over φ .

$$p(\mathbf{w} | \mathbf{s}, \mathbf{z}, \beta) = \left(\frac{\Gamma(V\beta)}{\Gamma(\beta)^V} \right)^{T \cdot S} \prod_j \prod_k \frac{\Gamma(n_{i,j,k} + \beta)}{\Gamma(n_{j,k} + V\beta)}, \quad (2)$$

Where Γ denotes Gamma function. Subscripts i, j, k are used to loop over words, topics, and sentiments respectively.

The remaining terms of Equation 1 are obtained in the same way by integrating over π and θ respectively.

$$p(\mathbf{s} | \mathbf{z}, \gamma) = \left(\frac{\Gamma(\sum_k \gamma_k)}{\prod_k \Gamma(\gamma_k)} \right)^T \prod_j \prod_k \frac{\Gamma(n_{j,k} + \gamma_k)}{\Gamma(n_j + \sum_k \gamma_k)}, \quad (3)$$

$$p(\mathbf{z} | \alpha) = \left(\frac{\Gamma(\sum_j \alpha_j)}{\prod_j \Gamma(\alpha_j)} \right)^D \prod_d \prod_j \frac{\Gamma(n_{d,j} + \alpha_j)}{\Gamma(n_d + \sum_j \alpha_j)}, \quad (4)$$

Posterior distribution.

Posterior distribution is estimated by sampling the variables z, s given all other variables. We use the superscript $-p$ to denote the quantity of data that excludes the word at position p of the current document d . Posterior probability can be derived from joint probability as follows:

$$p(s_p = k, z_p = j | \mathbf{w}, \mathbf{s}^{-p}, \mathbf{z}^{-p}, \alpha, \beta, \gamma) \propto \frac{n_{d,j}^{-p} + \alpha_j}{n_d^{-p} + \sum_j \alpha_j} \cdot \frac{n_{w_p,j,k}^{-p} + \beta}{n_{j,k}^{-p} + V\beta} \cdot \frac{n_{j,k}^{-p} + \gamma_k}{n_j^{-p} + \sum_k \gamma_k}. \quad (5)$$

Samples obtained from the Markov chain are then used to estimate the distributions φ, θ , and π as follows:

$$\varphi_{j,k,i} = \frac{n_{i,j,k} + \beta}{n_{j,k} + V\beta}, \theta_{d,j} = \frac{n_{d,j} + \alpha_j}{n_d + \sum_j \alpha_j}, \pi_{j,k} = \frac{n_{j,k} + \gamma_k}{n_j + \sum_k \gamma_k} \quad (6)$$

Gibbs sampling algorithm.

The complete Gibbs sampling procedure is given in Algorithm 1. Let W be the number of words of all documents in the learning data ($W = \sum_{d \in D} n_d$). The computational complexity of sampling a topic and a sentiment for a word (line 6 of Algorithm 1) is $O(S \cdot T)$. Consequently, the computational complexity of each Gibbs sampling iteration is $O(W \cdot S \cdot T)$. For the sake of clarity, the time complexity of drawing a multinomial has been assumed to be $O(1)$.

Algorithm 1 Inference on TS

Require: α, β, γ, T

- 1: Initialize matrices Φ, Θ, Π .
 - 2: **for** iteration $c = 1$ **to** $nbGibbsIterations$ **do**
 - 3: **for** document $d = 1$ **to** D **do**
 - 4: **for** $p = 1$ **to** n_d **do**
 - 5: Exclude word w_p from d and update count variables
 - 6: Sample a topic and a sentiment label for word w_p using Equation 5
 - 7: Update count variables with the new samples
 - 8: **end for**
 - 9: **end for**
 - 10: **end for**
 - 11: Update matrices Φ, Θ, Π using Equations 6
-

3.2 Incorporating Prior Knowledge

Following the work in [7, 8, 9, 10], we use prior knowledge, represented in the form of a sentiment lexicon (a list of words annotated with prior sentiment labels), in order to guide the sentiment discovery. Prior knowledge is incorporated when sampling a sentiment for a word. Thus, if the word is in the lexicon, it is assigned with the corresponding sentiment label taken from the lexicon. Otherwise, the sentiment label is generated using Equation 5.

In all experiments, we use a subset of MPQA subjectivity lexicon [15] (words tagged with “strong subjectivity”), to which we add a number of manually-tagged words. The final stemmed lexicon comprises 1502 positive and 2541 negative stems.

4. EXPERIMENT

4.1 Dataset

We use two different datasets for experiments: MDS and MDSfr. MDS dataset (Multi-Domain Sentiment dataset [2]) consists of reviews for different types of products from Amazon¹. Documents are annotated with 24 different topics

¹<http://www.amazon.com/>

(books, apparel, software, kitchen etc.) and sentiment labels (positive or negative). In the original dataset, the positive polarity prevails on the negative one (23 from the 24 topics are annotated positive). However, in machine learning, the models are usually tested on balanced datasets in order to assess their ability to handle the different sentiment polarities. For our experiments, we rebalance the MDS dataset so as to have an equal number of positive and negative topics.

The second dataset MDSfr is a collection of French-written reviews for different types of products. We have collected this dataset on Amazon France². Each review is annotated with the topic (product category from a set of 17 categories) and polarity (positive for 5-star reviews, and negative for 1-star reviews).

As preprocessing, we perform stopword removal and word stemming using Porter stemmer. Basic statistics on the pre-processed datasets are given in Table 3.

Table 3: Dataset statistics. The last row gives the number of most likely positive/negative topics.

Dataset	MDS	MDSfr
Type	Reviews	Reviews
Language	English	French
D	27065	10668
V	42010	12773
# pos./neg. topics	12/12	9/8

4.2 Evaluation Framework

In this section, we present our approach to evaluate TS model as well as two baseline models: JST and ASUM. Other models may be relevant for this comparison, such as [12, 16] but, for ease of reproducibility we only choose models with available source code.

In the original papers, JST and ASUM have been evaluated using techniques from supervised learning (sentiment prediction at the document level) but this is not the initial purpose of these models, neither of TS. In addition, such an evaluation does not take into account the sentiment prediction at the topic level. In this paper, we adopt an evaluation based on sentiment prediction at the topic level.

Documents from MDS and MDSfr datasets are annotated with topics (product categories) and sentiment labels. We rely on this annotation in order to build a ground truth about topic’s sentiments in the following way:

1. The dataset of annotated documents is split into T subsets, where each subset D_j is made up of the documents annotated with topic j .
2. Each topic is then annotated positive or negative based on the number of positive and negative documents it contains (annotated with the class of the majority).

The ground truth about topic’s sentiments can then be used to calculate a classical accuracy measure at the topic level. Here, the accuracy is the proportion of correctly classified topics. Finally, the evaluation procedure is as follows:

1. Estimated topics are first matched to real topics. As the number of topics is small, this step is performed manually by asking a human to label each of the extracted topics with a label from the set of topic labels.

2. Once the extracted topics labeled, the sentiment classes of the topics with the same label are compared: on the one hand the real sentiment class of the topic, and on the other hand the estimated sentiment class calculated by maximizing π distribution (the sentiment class m is assigned to a topic z if $m = \arg \max_s \pi_{z,s}$).

It is noted that the annotation of documents and topics is required only for the evaluation purpose. The topic-sentiment modeling with TS remains fully unsupervised.

For the evaluation of JST and ASUM models, a post-processing is required in order to convert their results to the same form as TS. We adopt a similar approach as the creation of the ground-truth: each document d is re-annotated with the sentiment and the topic maximizing probability θ_d .

4.3 Parameter Setting

For topic-sentiment models, the choice of hyperparameter values is important. To set the symmetric hyperparameter α , we follow the technique commonly used in the literature: $\alpha = \frac{50}{T}$ [10]. The symmetric hyperparameter β is set to $\frac{1}{T}$ empirically. Finally, the asymmetric hyperparameter γ have been set for each model/dataset separately by seeking the values that maximize the overall accuracy. Thus, each model is given the chance to produce its best result (cf. Table 4).

Table 4: The best values of hyperparameter (γ_+ , γ_-).

Dataset	T	Model		
		TS	JST	ASUM
MDS	24	(0.1, 200)	(0.1, 2)	(0.1, 50)
MDSfr	17	(0.1, 20)	(0.1, 2)	(0.1, 10)

4.4 Results and Analysis

All the results are obtained at the 300th iteration of Gibbs sampler (empirically determined). In this Section, the results are presented and discussed under the two perspectives: topic discovery and topic-sentiment prediction.

4.4.1 Topic discovery

The first goal of TS model is to extract coherent topics. In Table 5, we show some examples of the extracted topics. From this result, we can draw some conclusions for the two main features of TS model: topic’s coherence and topic’s sentiments.

On the one hand, as can be seen from the table, the topics are quite coherent and meaningful on both datasets. For example, from MDS dataset, z_1 captures the underlying topic of videos and movies, z_2 corresponds to the topic of software and computer industry etc. From MDSfr dataset, the result is the same; most of the discovered topics are coherent and informative. For example, z_6 relates to the topic *chausures* (shoes), z_8 corresponds to the topic *cuisines* (kitchen equipment) etc.

On the other hand, most extracted topics are opinion-bearing because they are clearly described with sentiment-bearing words in addition to the topical words (cf. Table 5). This is an important feature of TS model because one can immediately retrieve the two aspects of the same topic with no additional post-processing. Moreover, the estimated π distribution (topic distribution over sentiment polarities) are often close to the real distributions, that is the topic-specific sentiments are efficiently captured by TS model.

²<http://www.amazon.fr/>

Table 5: Top words of selected topics from MDS (top) and MDSfr (bottom) datasets. Positive and negative words from the lexicon are represented in green (underlined) and red (italic) respectively.

Topic	z ₁ : video		z ₂ : software		z ₃ : toys and games		z ₄ : music		z ₅ : office products	
	<i>negative</i>	<i>positive</i>	<i>negative</i>	<i>positive</i>	<i>negative</i>	<i>positive</i>	<i>negative</i>	<i>positive</i>	<i>negative</i>	<i>positive</i>
Top words	story	movie	version	use	old	kid	guitar	song	use	pen
	<i>bad</i>	film	program	product	toy	<u>love</u>	sing	album	paper	plastic
	performance	<u>good</u>	software	<u>work</u>	year	<u>play</u>	<i>boring</i>	music	one	pocket
	<i>boring</i>	scene	upgrad	<u>easy</u>	daughter	<u>fun</u>	blue	<u>good</u>	ink	<u>nice</u>
	<i>horror</i>	character	try	feature	<i>disappointed</i>	christmas	<i>miss</i>	<u>rock</u>	price	cartridge
	made	actor	<i>problem</i>	new	baby	<u>gift</u>	<i>bad</i>	<u>love</u>	buy	printer
	<i>worst</i>	<u>great</u>	buy	<u>great</u>	age	<u>enjoy</u>	remix	<u>fan</u>	small	design
	<i>stupid</i>	<u>play</u>	install	<u>better</u>	<i>frustrated</i>	learn	<i>lack</i>	<u>original</u>	<i>bad</i>	<u>durable</u>
Real π	0.60	0.40	0.40	0.60	0.80	0.20	0.77	0.23	0.15	0.85
Est. π	0.62	0.38	0.32	0.68	0.47	0.53	0.56	0.44	0.02	0.98

Topic	z ₆ : chaussures (shoes)		z ₇ : téléphones (phones)		z ₈ : cuisine (kitchen)		z ₉ : beauté et parfum (beauty)		z ₁₀ : puériculture (childcare)	
	<i>negative</i>	<i>positive</i>	<i>negative</i>	<i>positive</i>	<i>negative</i>	<i>positive</i>	<i>negative</i>	<i>positive</i>	<i>negative</i>	<i>positive</i>
Top words	taille	couleur	<u>simple</u>	écran	eau	prothèse	savon	huile	<i>encombrant</i>	<u>pratique</u>
	<i>trop</i>	pinceau	<i>bug</i>	téléphone	bébé	électricien	<i>sèche</i>	peau	<i>trop</i>	<u>facile</u>
	coup	chaussures	htc	samsung	chauffer	siemens	odeur	utiliser	<i>mal</i>	ranger
	noir	<u>joli</u>	fenêtre	<u>vrai</u>	bouilloire	thermostat	<i>détruire</i>	produit	fermeture	enfant
	<i>déçu</i>	marque	connecter	<u>autonomie</u>	température	beep	<i>poubelle</i>	douche	<i>compliqué</i>	poussette
	tissu	<u>qualité</u>	paramétrage	<u>bon</u>	plaque	<u>qualité</u>	noir	hydrater	pliage	intérieur
	lavage	paire	<i>loin</i>	prise	<i>mauvais</i>	<u>four</u>	texture	<u>efficace</u>	démonter	biberon
	<i>dommage</i>	ped	ios	internet	onde	<u>bon</u>	marseille	gel	bébé	transport
Real π	0.39	0.61	0.60	0.40	0.20	0.80	0.38	0.62	0.57	0.43
Est. π	0.18	0.82	0.64	0.36	0.41	0.59	0.54	0.46	0.89	0.11

4.4.2 Topic-sentiment prediction

The second goal of TS is to extract topic’s sentiments. In order to examine the accuracy of sentiment prediction at the topic level, we compare our model to JST and ASUM models, using the method presented in Section 4.2.

Table 6: Topic-sentiment prediction results.

Data	Model	Avg. accuracy	Std. dev.
MDS	TS	0.750	0.029
	JST	0.600	0.037
	ASUM	0.667	0.029
MDSfr	TS	0.765	0.042
	JST	0.541	0.049
	ASUM	0.718	0.049

The results achieved on MDS dataset are represented in Table 6 (average and standard deviation based on 5 random initializations). As can be seen, TS model significantly outperforms the two other models JST and ASUM with 15 points and 8 points respectively on MDS dataset. TS model also outperforms JST and ASUM on MDSfr dataset with 23 and 5 points respectively. The good performance of TS compared to other models can be explained by its ability to capture the overall sentiments relative to topics without post-processing. Thus, this information is directly captured with a dedicated variable intrinsic to the model unlike the other models that use post-processing.

5. OPTIMAL PARAMETER SETTING

It has been shown by Hatzivassiloglou and McKeown that positive words are more frequently used than negative words in English language [3]. This hypothesis is also valid for the French-written dataset MDSfr. By examining the Gibbs

Sampling process, we notice that this phenomenon has an important effect on the algorithm progress. In particular, it guides the assignment of words to sentiment labels (cf. Equation 5). Due to the non-balanced distribution of sentiment-bearing words, the step of assigning sentiment labels with words results more frequently in positive labels, because the negative class is overwhelmed by the positive one. After a sufficient number of iterations, most of words will be assigned with the positive sentiment class. Consequently, most of topics will be characterized with high probability on the positive class too (cf. Equation 6).

The most common way to correct this bias is to set an asymmetric topic-sentiment prior γ with a greater value on the negative sentiment label [7, 10]. This strategy results in a less biased assignment of words with sentiment labels. However, it has two main limitations: (i) The best value of the asymmetric γ is not easy to find because it highly depends on the dataset. Generally, it is empirically set after a number of experiments. (ii) Once γ is fixed for a dataset, it is only valid for a certain number of iterations. For example, if γ has been set based on an experiment with 100-iteration sample, it is not necessary good for a 200-iteration sample because the effect of the negative prior will be too much important, which will make the model more likely to assign the words with the negative sentiment.

5.1 Our Proposal

To address the above issues, we propose to set the γ prior dynamically at each iteration. We denote W_+ , respectively W_- , the number of words assigned with the positive, respectively negative, sentiment polarity by the end of the iteration c . Three situations are possible: either $W_+ > W_-$, or $W_+ < W_-$, or $W_+ = W_-$. The idea is to make the distri-

bution of words balanced over the two sentiment polarities ($W_+ = W_-$). To this end, we need to decrease the prior for the prevailing polarity and increase it for the other one. In the next iteration $c + 1$, we make the new value of the prior equal to the average number of words that should be artificially added to each topic in order to reach the equilibrium ($W_+ = W_-$). Thus, if $W_+ > W_-$, we set $\gamma_+ = \epsilon$ and $\gamma_- = \frac{W_+ - W_-}{T}$. If $W_+ < W_-$, we set $\gamma_+ = \frac{W_- - W_+}{T}$ and $\gamma_- = \epsilon$. Otherwise, we set $\gamma_+ = \gamma_- = \epsilon$. For all the experiments in this paper, $\epsilon = 0.01$.

5.2 Experiment and Results

To evaluate the efficiency of our method, we compare it to the widely-used Maximum-Likelihood (ML) estimation method. We have implemented this method using the fixed-point iteration described in [6]. The results are represented in Table 7 (average and standard deviation based on 5 random initializations).

Table 7: Results obtained with the proposed method for optimal setting of γ parameter.

Data	Method	Avg. accuracy	Std. dev.
MDS	Our method	0.750	0.029
	ML-based method	0.542	0.029
MDSfr	Our method	0.765	0.042
	ML-based method	0.588	0.049

As can be seen from the results, our method outperforms the ML-based method on both MDS and MDSfr datasets. Actually, the ML-based method performs poorly on this specific task, which makes it not adapted for the estimation of γ parameter. However, the ML-based method has been effectively used to estimate hyperparameters for topic models. For example, in [13], it has shown good performance for estimating α and β hyperparameters. In our experiments, the hyperparameter γ does not follow the same logic. The equilibrium constraint that we have incorporated into the learning algorithm significantly contributes to increase the model accuracy in predicting topic’s sentiments.

6. CONCLUSION

In this paper, we have discussed the issue of topic-sentiment modeling. We have underlined the main limitations of the existing models and we have proposed a novel approach for the joint modeling of topics and topic-relative sentiments using probabilistic topic models. Our proposal stands out from the existing models with two key features: (i) different descriptions of the same topic corresponding to the different sentiment polarities, (ii) overall topic-specific distribution over sentiment polarities. The experiments performed on two Web datasets have confirmed the superiority of our model when compared to the state-of-art models, in particular for predicting the topic-specific sentiment polarity.

It is worth noting that TS model, like the other topic-sentiment models tested in this work, have been very sensitive to the topic-sentiment prior γ . The performance of these models may decrease dramatically when changing the value of this parameter. To overcome this limitation, we have proposed a method to automatically and dynamically set the topic-sentiment prior γ . Here again, our method significantly outperforms the widely-used ML-based method in terms of predicting topic-specific sentiments.

We are now working on the extension of TS model with a time dimension. Some previous works have shown the efficiency of topic models for modeling time jointly with topics and/or sentiments (e.g., [4, 16]). TS model can easily be extended in a way similar to [16] to capture the topic-sentiment dynamics.

7. REFERENCES

- [1] David M. Blei, Andrew Y. Ng, and Michael I. Jordan. Latent Dirichlet Allocation. *JMLR*, 3:993–1022, 2003.
- [2] John Blitzer, Mark Dredze, and Fernando Pereira. Biographies, Bollywood, Boom-boxes and Blenders: Domain Adaptation for Sentiment Classification. In *ACL’07*, pages 440–447, Prague, Czech Republic, 2007.
- [3] Vasileios Hatzivassiloglou and Kathleen R Mckeown. Predicting the Semantic Orientation of Adjectives. In *EACL’97*, pages 174–181, Madrid, Spain, 1997.
- [4] Yulan He, Chenghua Lin, Wei Gao, and Kam Fai Wong. Dynamic Joint Sentiment-Topic model. *TIST*, 9(4), 2014.
- [5] Gregor Heinrich. Parameter estimation for text analysis. Technical report, 2005.
- [6] Jonathan Huang. Maximum likelihood estimation of Dirichlet distribution parameters. Technical report, 2005.
- [7] Yohan Jo and Alice H. Oh. Aspect and sentiment unification model for online review analysis. In *WSDM’11*, pages 815–824, Hong Kong, China, 2011.
- [8] Chengtao Li, Jianwen Zhang, Jian-tao Sun, and Zheng Chen. Sentiment Topic Model with Decomposed Prior. In *SDM’13*, pages 767–776, Austin, TX, USA, 2013.
- [9] Fangtao Li, Minlie Huang, and Xiaoyan Zhu. Sentiment analysis with global topics and local dependency. *AAAI’10*, pages 1371–1376, 2010.
- [10] Chenghua Lin and Yulan He. Joint sentiment/topic model for sentiment analysis. In *CIKM’09*, pages 375–384, Hong Kong, China, 2009.
- [11] Chenghua Lin, Yulan He, Richard Everson, and Stefan Ruger. Weakly Supervised Joint Sentiment-Topic Detection from Text. *TKDE*, 24(6):1134–1145, June 2012.
- [12] Qiaozhu Mei, Xu Ling, Matthew Wondra, Hang Su, and ChengXiang Zhai. Topic sentiment mixture: modeling facets and opinions in weblogs. In *WWW’07*, pages 171–180, Banff, Canada, 2007.
- [13] Hanna M. Wallach, David Mimno, and Andrew McCallum. Rethinking LDA: Why Priors Matter. In *NIPS’2009*, pages 1973–1981, Vancouver, Canada, 2009.
- [14] Xuerui Wang and Andrew McCallum. Topics over time: a non-Markov continuous-time model of topical trends. In *KDD’06*, pages 424–433, Philadelphia, PA, USA, 2006.
- [15] Theresa Wilson, Janyce Wiebe, and Paul Hoffmann. Recognizing contextual polarity in phrase-level sentiment analysis. In *HLT/EMNLP’05*, pages 347–354, Vancouver, Canada, 2005.
- [16] Minjie Zheng, Chaorong Wu, Yue Liu, Xiangwen Liao, and Guolong Chen. Topic Sentiment Trend Model: Modeling facets and sentiment dynamics. In *CSAE’12*, number 1, pages 651–657, Zhangjiajie, China, 2012.