

# Topic Modeling based on Keywords and Context

Johannes Schneider, University of Liechtenstein  
Michail Vlachos, IBM Research, Zurich

October 12, 2017

## Abstract

Current topic models often suffer from discovering topics not matching human intuition, unnatural switching of topics within documents and high computational demands. We address these shortcomings by proposing a topic model and an inference algorithm based on automatically identifying characteristic keywords for topics. Keywords influence the topic assignments of nearby words. Our algorithm learns (key)word-topic scores and self-regulates the number of topics. The inference is simple and easily parallelizable. A qualitative analysis yields comparable results to those of state-of-the-art models, such as LDA, but with different strengths and weaknesses. Quantitative analysis using nine datasets shows gains regarding classification accuracy, PMI score, computational performance, and consistency of topic assignments within documents, while most often using fewer topics.

## 1 Introduction

Topic modeling deals with the extraction of latent information, i.e., topics, from a collection of text documents. Classical approaches, such as probabilistic latent semantic indexing (PLSA) [11] and more so its generalization, Latent Dirichlet Allocation (LDA) [4] enjoy widespread popularity despite a few shortcomings. They neglect word order within documents, i.e., documents are not treated as a sequence of words but rather as a set or ‘bag’ of words. This might be one reason for unnatural word-topic assignments within documents, where topics change after almost every word (Figure 1). Many attempts have been made to remedy the ‘bag-of-words’ assumption (eg. [8, 9, 23, 25]), but improvement comes usually at a price, a strong increase in computational demands and often model complexity.

The second shortcoming is an unnatural assignment of word-topic probabilities. Better solutions to the LDA model (in terms of its loss function) might result in worse human interpretability of topics [5]. For

PLSA (and LDA) the frequency of a word within a topic heavily influences its probability (within the topic) whereas the frequencies of the word in other topics have a lesser impact. Thus, a very common word that occurs equally frequently across all topics might have a large probability in each topic. This makes sense for models based on document “generation”, because frequent words should be generated more often. But using our rationale that a likely word should also be typical for a topic, high scores (or probabilities) should only be assigned to words that occur in a few topics. Although there are techniques that provide a weighing of words, they either do not fully cover the above rationale and perform static weighing e.g., TF-IDF, or the outcome comes at high computational demands with limited improvements [13].

Third, a topic modeling technique should discover an appropriate number of topics for a corpus and a specific task. Hierarchical models based on LDA [22, 17] limit the number of topics automatically, but increase computational costs. PLSA and LDA do not tend to self-regulate the number of topics. They return as many topics as specified by a parameter. We proclaim that this parameter should be seen as an upper bound. Choosing too large a number of topics leads to overfitting. Furthermore, it is unclear how to choose the number of topics. Manual investigation of topics also benefits from self-regulation, i.e., a reduction of the number of topics to consider.

Other concerns of topic models are the computational performance and the implementation complexity. We provide a novel, holistic model and inference algorithm that addresses all of the above mentioned shortcomings of existing models – at least partially – by introducing a novel topic model based on the following rationale: i) For each word in each topic, we compute a keyword score stating how likely a word belongs to the topic. Roughly speaking, the score depends on how common the word is within the topic and how common it is within other topics. Classical generative models compute a word-topic probability distribution that states how likely a word is generated given a topic. In those models, the probability of a word



**Figure 1:** Words colored by topics for our proposed topic keyword model, TKM (left) and LDA (right). TKM assigns topics more consistently.

given a topic depends less strongly (only implicitly) dependent on the frequency of a word in other topics. ii) We assume that topics of a word in a sequence of words are heavily influenced by words in the same sequence. Therefore, words near a word with high keyword score might be assigned to the topic of that word, even if they themselves are unlikely for the topic. Thus, the order of words has an impact in contrast to bag-of-words models.

iii) The topic-document distribution depends on a simple and highly efficient summation of the keyword scores of the words within a document. Like LDA does, it has a prior for topic-document distributions. iv) An intrinsic property of our model is to limit the number of distinct topics. Redundant topics are removed explicitly.

After modeling the above ideas and describing an algorithm for inference, we evaluate it on several datasets, showing improvements compared with two other methods (LDA and BTM [25]) on the PMI score, classification accuracy, performance and a new metric denoted by topic-change likelihood. This metric gives the probability that two consecutive words have different topics. From a human perspective, it seems natural that multiple consecutive words in a document should very often belong to the same topic. Models such as LDA tend to assign consecutive words to different topics (Figure 1).

## 2 Topic Keyword Model

Our topic keyword model (TKM) uses ideas from the aspect model [12] which defines a joint probability  $D \times W$  (for notation see Table 1). The standard aspect model assumes conditional independence of words and documents given a topic:

$$\begin{aligned}
 p(d, w) &:= p(d) \cdot p(w|d) \\
 p(w|d) &:= \sum_t p(w|t) \cdot p(t|d)
 \end{aligned}
 \tag{1}$$

Our model (Equations 2 – 4) maintains the core ideas of the aspect model, but we account for context by taking the position  $i$  of a word into account and we use a keyword score  $f(w, t)$ . Taking the context of a word into account implies that if a word occurs multiple times in the same document but with different nearby words, it might have different probabilities. To express context, we include the index  $i$ , denoting position of the  $i$ th word in the document (see Equation 2). The distribution  $p(d)$  is proportional to  $|d|$ . We simply use a uniform distribution. We also assume a uniform distribution for  $p(i|d)$ , as we do not consider any position in the document as more likely (or important) than any other. In principle, one could, for example, give higher priority to initially occurring words that might correspond to a summarizing abstract of a text.

### Model Equations

$$p(d, w, i) := p(d) \cdot p(i|d) \cdot p(w_i = w|d, i) \tag{2}$$

$$p(w|d, i) := \max_{t, j \in R_i} \{f(w, t) + f(w_{i+j}, t)\} \cdot p(t|d) \tag{3}$$

$$\text{with } R_i := [\max(0, i - L), \min(|d| - 1, i + L)]$$

$$p(t|d) := \frac{(\sum_{i \in [0, |d| - 1]} f(w_i, t))^\alpha}{\sum_t (\sum_{i \in [0, |d| - 1]} f(w_i, t))^\alpha} \tag{4}$$

To compute the probability of a word at a certain position in a document (Equation 3) we use latent variables, i.e., topics, as done in the aspect model (Equations 1). Instead of the generative probability  $p(w|t)$  that word  $w$  occurs in topic  $t$ , we use a keyword score  $f(w, t)$ . A keyword for a topic should be somewhat frequent and also characteristic for that topic only. The keyword scoring function  $f$  computes a keyword score for a particular word and topic that might depend on multiple parameters such as  $p(w|t)$ ,  $p(t|w)$  and  $p(t)$ , whereas the generative probability  $p(w|t)$  is only based on the relative number of occurrences of a word in a topic. We shall discuss such functions in

Section 2.1. We use the idea that a word with a high keyword score for a topic might “enforce” the topic onto a nearby word, even if that word is just weakly associated with the topic. For a word  $w_i$  at position  $i$  in the document, all words within  $L$  words to the left and right, i.e., words  $w_{i+j}$  with  $j \in [-L, L]$ , could contain a word (with high keyword score) that determines the topic assignment of  $w_i$ . To account for boundary conditions as the beginning and end of a document  $d$ , we use  $j \in [\max(0, i - L), \min(|d| - 1, i + L)]$ . There are multiple options how a nearby keyword might impact the topic of a word. The addition of the scores  $f(w_i, t)$  and  $f(w_{i+j}, t)$  exhibits a linear behavior that is suitable for modeling the assumption that one word might determine the topic of a nearby word, even if the other word is only weakly associated with the topic. Generative models following the bag-of-words model imply the use of multiplication of probabilities, i.e., keyword scores, which does not capture our modeling assumption: A word that is weakly associated with a topic, i.e., has a score close to zero, would have a low score for the topic even in the presence of strong keywords for the topic. Furthermore, each occurrence of a word in a document is assumed to stem from exactly one topic, which is expressed by taking the maximum in Equation 3.

We compute  $p(t|d)$  dependent on keyword scores  $f(w, t)$  of the words in the document  $d$  (Equation 4). We model the idea of looking for keywords in a document and aggregating their score, i.e.,  $f(w, t)$ . The parameter  $\alpha$  impacts the number of topics per document. The larger  $\alpha$  the more concentrated the topic-document distribution, i.e., the fewer topics per document.

Essentially, these equations allow us to derive an algorithm for inference that is more efficient than LDA, while avoiding overfitting and allowing to model a prior on topic concentration.

## 2.1 Modeling Keywords

Here we state one way how to compute the keyword score  $f(w, t)$  given a word and a topic. Alternative options are discussed in the Related Work Section. A word obtains a high keyword score for a topic if it is assigned often to the topic and the relative number of assignments to the topic is high compared with other topics. The first aspect relates to the frequency of the word  $n(w, t)$  within the topic. The second captures how characteristic a keyword is for the topic relative to others, i.e.,  $p(t|w)$ . If the topic-word distribution  $p(t|w)$  is uniform, the word is not characteristic of any topic. If it is highly concentrated then it is. This can be captured using the inverse entropy  $H(w)$ :

$$H(w) := - \sum_t p(t|w) \cdot \log(p(t|w)) \quad (5)$$

The entropy  $H(w)$  is maximized for a uniform distribution, i.e.,  $p(w|t) = 1/|T|$  giving  $H(w) = \log |T|$ , where  $|T|$  is the number of (current) topics, i.e., initially  $k$ . Thus,  $1/H(w)$  is a measure that increases the more concentrated the assignment of words to topics is. However, if all occurrences of a word are assigned to a single topic, the entropy is zero and the inverse infinity. Therefore, we add one in the denominator, i.e.,  $1/(1 + H(w))$ . Furthermore, if the occurrences of a word  $n(w)$  in the entire corpus are fewer than the number of topics  $|T|$ , then the word’s entropy can be at most  $\log n(w) < \log |T|$ . Ignoring this results in a high keyword score for a rare word even if each occurrence is assigned to a different topic. Thus, we ensure that rare words are not preferred too much by using the factor  $\log \min(|T|, n(w) + 1)$ , where the addition of one is to ensure non-zero weights for words that occur once. An optional weight parameter  $\delta$  allows more or less emphasis to be put on the concentration (relative to the frequency within a topic). Overall, our concentration score is:

$$con(w) := \left( \frac{\log(\min(|T|, n(w) + 1))}{1 + H(w)} \right)^\delta \quad (6)$$

The second aspect of keyword scores relates to the frequency of the word within a topic, i.e.  $p(w, t) \cdot \sum_{d \in D} |d|$ . Damped frequencies, e.g.,  $\log(1 + p(w, t) \cdot \sum_{d \in D} |d|)$ , work better than using raw frequencies for inference and classification, because classification relies more on concentration, i.e., being certain that a word belongs to a topic. For humans, the words with the highest keyword score are often too specific – they might only be familiar to experts on the topic. Therefore, we propose a second keyword distribution targeted for human understanding that puts more emphasis on frequency using raw counts. Combining the word frequency and the concentration score we get a score  $f(w, t)$  that prefers rather specific keywords and a second one that emphasizes more widely used (known) words  $f_{hu}(w, t)$ . Both can be normalized. We add a prior  $\beta$  for  $f(w, t)$ , similar to LDA and other models, stating that a word is assumed to occur for a each topic at least  $\beta$  times.

$$\begin{aligned} f(w, t) &\propto \log(1 + p(w, t) \cdot \sum_{d \in D} |d| + \beta) \cdot con(w) \\ f_{hu}(w, t) &\propto (p(w, t) \cdot \sum_{d \in D} |d|) \cdot con(w) \end{aligned} \quad (7)$$

## 3 Inference

We want to find parameters that maximize the likelihood of the data, i.e.,  $\prod_d \prod_{i \in [0, |d| - 1]} p(d, w, i)$ . A key challenge for inference is the fairly complex model

formalized in Equations (2–4) and (7). Although methods such as Gibbs sampling might be used, they would be rather inefficient. In particular, their optimizations for faster inference are harder to apply, such as integrating out (collapsing) variables, for a Gibbs Sampler. To derive an efficient inference mechanism, we follow the expectation-maximization(EM) approach combined with standard probabilistic reasoning based on word-topic assignment frequencies. The general idea of EM is to perform two steps. In the E-step latent variables are estimated, i.e., the probability  $p(t|w, i, d)$  of a topic given word  $w$  and position  $i$  in document  $d$ . In the M-step the loss function is maximized with respect to the parameters using the topic distribution  $p(t|w, i, d)$ . In our model we assume that a word at some fixed position in a document can only have one topic as expressed in Equation (3). Therefore, the topic  $t(w, i, d)$  is simply the most likely topic of that word in that context, i.e., adjusting Equation (3) accordingly we get:

$$t(w, i, d) := \operatorname{argmax}_t \{(f(w, t) + f(w_{i+j}, t)) \cdot p(t|d) | j \in R_i\}$$

$$p(t|w, i, d) = \begin{cases} 1 & t_{w,i,d} = t \\ 0 & t_{w,i,d} \neq t \end{cases} \quad (8)$$

This differs from PLSA and LDA, where each word within a document is assigned a distribution typically with non-zero probabilities for all topics. In the M-Step, we want to optimize  $\Theta$ :

$$\Theta^{new} = \operatorname{argmax}_{\Theta} Q(\Theta, \Theta^{old})$$

$$\text{with } Q(\Theta, \Theta^{old}) := \sum_{d,i,t} p(t|D, \Theta^{old}) \log p(D, t|\Theta) \quad (9)$$

The optimization problem in Equation (9) might be tackled using various methods, e.g. using Lagrange multipliers. Unfortunately, simple analytical solutions based on these approaches are intractable given the complexity of the model equations (2-4). However, one might also look at the inference of  $p(w|t)$ ,  $p(t)$  and  $p(w)$  differently. Assume that we are given the assignments of words  $w$  to topics  $t$  for a collection of documents  $D$ , i.e.,  $n(w, t)$ . Our frequentist inference approach uses the empirical distribution: The probability of a word given a topic equals the fraction of words that have been assigned to the topic of all words assigned to the topic. Under mild assumptions the maximum likelihood distribution equals the empirical distribution (see eg. 9.2.2 in [2]):

$$p(w|t) := \frac{n(w, t)}{\sum_w n(w, t)} \quad (10)$$

Note, that  $n(w, t)$  is computable by summing across the assignments from the E-step, i.e.,  $p(t|w, i, d)$ , because each word within a context is assigned to one topic only (Equation 8):

$$n(w, t) := \sum_{i,d} p(t|w_i, i, d) \quad (11)$$

To compute  $p(t|w)$  we use Bayes' Law to obtain  $p(t|w) = p(w|t) \cdot p(t)/p(w)$ . Therefore, the only free parameters, we need to estimate are  $p(t)$  and  $p(w)$ , i.e.,  $k+|W|$  rather than  $k \cdot |W|$ . We have  $p(t) = \frac{\sum_w n(w, t)}{\sum_{w,t} n(w, t)}$  and  $p(w) = \frac{\sum_t n(w, t)}{\sum_{w,t} n(w, t)}$ , thus:

$$p(t|w) = p(w|t) \cdot \frac{p(t)}{p(w)}$$

$$= \frac{n(w, t)}{\sum_{w'} n(w', t)} \cdot \frac{(\sum_{w'} n(w', t)) / \sum_{w', t'} n(w', t')}{(\sum_{t'} n(w, t')) / \sum_{w', t'} n(w', t')}$$

$$= \frac{n(w, t)}{\sum_{t'} n(w, t')}$$

We also need to compute the keyword score  $f(w, t)$ . We use that  $\sum_{d \in D} |d| = \sum_{w,t} n(w, t)$ , since each word in each document is assigned to one topic.

$$f(w, t) \propto \log(1 + p(w, t) \cdot \sum_{d \in D} |d| + \beta) \cdot \operatorname{con}(w)$$

$$\propto \log(1 + p(w|t) \cdot p(t) \cdot \sum_{w,t} n(w, t) + \beta) \cdot \operatorname{con}(w)$$

$$\propto \log(1 + \frac{n(w, t)}{\sum_w n(w, t)} \cdot \frac{\sum_w n(w, t)}{\sum_{w,t} n(w, t)} \cdot \sum_{w,t} n(w, t) + \beta) \cdot \operatorname{con}(w)$$

$$\propto \log(1 + n(w, t) + \beta) \cdot \operatorname{con}(w)$$

## 4 Self-Regulation of Topics

We only keep word-topic distributions that are significantly different from each other. Redundant topics can be removed either during inference (as done in Algorithm 1) or after inference. To measure the difference between two word-topic distributions we use the symmetrized Kullback–Leibler divergence.

$$KL(t_i, t_j) := \sum_w p(w|t_i) \cdot \log\left(\frac{p(w|t_i)}{p(w|t_j)}\right) \quad (12)$$

$$SKL(t_i, t_j) := KL(t_i, t_j) + KL(t_j, t_i)$$

The set of indexes of (significantly) distinct word-topic distributions  $DT$  is such that for any two topics  $i, j \in DT \subseteq T$ , it holds that  $SKL(t_i, t_j) \geq \gamma$ . We used  $\gamma := 0.25$ .

## 5 Evaluation

TKM and LDA were compared using quantitative (Experiments 1,2) and qualitative analysis (Experiment 3). We assessed the tendency to self-regulate the number of topics (Experiment 1). We evaluated several metrics, such as the classification accuracy, PMI score

Symbol	Meaning
$D$	corpus, all documents
$d$	document from $D$
$ d $	number of words in $d$
$W$	set of unique words in $D$
$w$	word
$w_i$	$i$ -th word in a document $d$
$k$	(maximal) number of topics
$T$	set (of indexes) of topics $T \subseteq [0, k - 1]$
$DT$	set (of indexes) of distinct topics in $T$
$t$	topic $t$ from $T$
$t(w, i, d)$	topic of word $w$ assuming $i$ -th word $w_i = w$ in $d$
$L$	length of sliding wind. to 1 side
$\alpha, \beta$	topic, word prior
$\delta$	weight for word concentration
$n(w, t)$	number of assignments of word $w$ to topic $t$
$n(w)$	number of occurrences of word $w$ in $D$

**Table 1:** Notation

---

**Algorithm 1** TKM(docs  $D$ , nTopics  $k$ , Priors  $\alpha, \beta$ )

---

```

1:  $\delta := 1.5; L := 7; p(t|d) := 1/k; T := [1, k]$ 
2:  $p(w|t) := 1/|W| + \text{noise}$ 
3: while  $p(w, t)$  “not converged” do
4:    $n(w, t) := 0$ 
5:   for  $d \in D$  do
6:     for  $i = 0$  to  $|d| - 1$  do
7:        $R_i := [\max(0, i - L), \min(|d| - 1, i + L)]$ 
8:        $t(w_i, i, d) := \operatorname{argmax}_t \{ (f(w_i, t) + f(w_j, t)) \cdot p(t|d) | j \in R_i \}$ 
9:        $n(w_i, t(w_i, i, d)) = n(w_i, t(w_i, i, d)) + 1$ 
10:    end for
11:   end for
12:    $p(t|w) := \frac{n(w, t)}{\sum_{t'} n(w, t')}$ 
13:    $H(w) := - \sum_t p(t|w) \cdot \log(p(t|w))$ 
14:    $\text{con}(w) := \left( \frac{\log(\min(|T|, n(w) + 1))}{1 + H(w)} \right)^\delta$ 
15:    $f(w, t) := \frac{\log(1 + n(w, t) + \beta) \cdot \text{con}(w)}{\sum_w \log(1 + n(w, t) + \beta) \cdot \text{con}(w)}$ 
16:    $T := DT$  {Keep only distinct topics (Section 4)}
17:    $p(t|d) := \frac{(\sum_{i \in [0, |d| - 1]} f(w_i, t))^\alpha}{\sum_t (\sum_{i \in [0, |d| - 1]} f(w_i, t))^\alpha}$ 
18: end while
19:  $f_{hu}(w, t) := \frac{n(w, t) \cdot \text{con}(w)}{\sum_{w'} n(w', t) \cdot \text{con}(w')}$ 

```

---

and computation time (Experiment 2), where we also included BTM in the evaluation. Finally, topics were assessed qualitatively for one dataset (Experiment 3).

Dataset	Docs	Unique Words	Avg. Words	Classes
BookReviews	179541	22211	33	8
WikiBig	52024	137155	346	11
Ohsumed	23166	23068	99	23
20Newsgroups	18625	37150	122	20
WikiSmall	10390	60612	388	82
Reuters21578	9091	11098	69	65
CategoryReviews	5809	6596	60	6
WebKB4Uni	4022	7670	136	4
BrownCorpus	500	16514	1006	15

**Table 2:** Datasets

once word-topic distributions no longer changed significantly. For all algorithms, we ran experiments with different parameters for  $\alpha$  and  $\beta$ . We chose the best configuration focusing on classification accuracy, i.e.,  $\alpha = 5/k$  and  $\beta = 0.04$  for LDA,  $\alpha = 50/k$  and  $\beta = 0.02$  for BTM. For TKM,  $\alpha = 2.5$  and  $\beta = 0.05$  were chosen for all experiments.

The datasets in Table 2 are public and most have already been used for text classification. For the distinct review datasets (from Amazon) we predicted either the product, i.e., book, based on a review or the product category. The Wiki benchmarks are based on Wikipedia categories.<sup>3</sup> We performed standard preprocessing, e.g., stemming, stopwords removal and removal of words that occurred only once in the entire corpus. All experiments ran on a server with a 64bit Ubuntu system, 100 GB RAM and 24 AMD cores operating at 2.6 GHz.

## 5.1 Algorithms, Datasets and Setup

We compared an implementation of Algorithm 1 in Python, LDA using a collapsed Gibbs sampler [7] available as Python library<sup>1</sup> and BTM [25] available as C++ library<sup>2</sup>. For all algorithms, we used the same convergence criterion, i.e., computation stopped

<sup>1</sup><https://pypi.python.org/pypi/lda>; The Gensim library implementing variational inference gave worse results. Topic removal (Section 4) for TKM is not done in every iteration.

<sup>2</sup><https://github.com/xiaohuiyan/BTM>

## 5.2 Experiments

**Experiment 1:** We empirically analyzed the convergence of the number of distinct topics  $|DT|$  (See Section 4) depending on the upper bound  $k$  of the number of topics for LDA and TKM.

**Experiment 2:** We compared various metrics for LDA, TKM and BTM using  $k = 100$ . The classification accuracy was measured using a random forest with

<sup>3</sup>We plan to make them publicly accessible. We will be happy to provide detailed references upon request.

100 trees, with 60% of all data for training and 40% for testing. The time to compute the word-topic and topic-document distributions on the training data, the time to compute the topic-document distribution on the test data, the number of distinct topics  $|DT|$  (Equation 12) and the PMI score as proposed in [19] were also compared. PMI measures the co-occurrence of words within topics relative to their co-occurrence within documents (or sequences of words) of a large external corpus, i.e., we used an English Wikipedia dump with about 4M. documents as done in [19]. For each pair of words  $w_i, w_j$  we calculated the fraction of documents in which both occurred.

$$\begin{aligned}
 p(w_i, w_j) &= \frac{|\{d|w_i, w_j \in d, d \in D\}|}{|D|} \\
 p(w_i) &= \frac{|\{d|w_i \in d, d \in D\}|}{|D|} \\
 PMI(w_i, w_j) &= \log \frac{p(w_i, w_j)}{p(w_i)p(w_j)}
 \end{aligned}
 \tag{13}$$

The PMI score for a topic is the median of the PMI score for all pairs of the ten highest-ranked words according to the word-topic distribution, i.e., for TKM the distribution is  $p_{hu}$  as in Equation 7. The PMI score for TKM (and LDA) is the mean of the PMI score of all individual topics. PMI has been reported to correlate better with human judgment than other measures such as perplexity [19].

We introduce a novel measure that captures the consistency of assignments of words to topics within documents as motivated in the end of Section 1. It is given by the probability that two consecutive words stem from different topics. As LDA computes a distribution across all topics, we choose the most likely topic (see Equation 8). The topic change probability  $ToC$  for a corpus  $D$  is defined as

$$ToC := \frac{\sum_{d \in D} \sum_i I_{t(i,d) \neq t(i+1,d)}}{\sum_{d \in D} |d|}
 \tag{14}$$

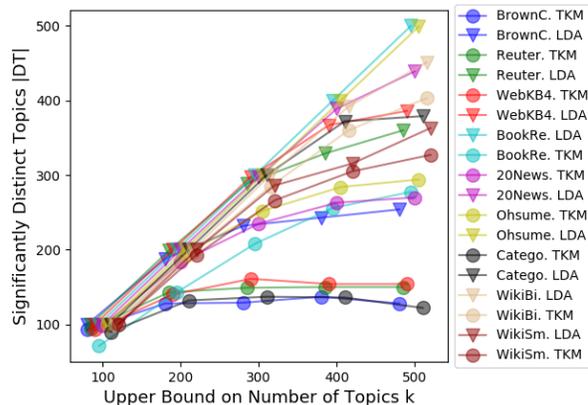
The indicator variable  $I$  is one if two neighboring words have different topics and zero otherwise.

**Experiment 3:** The topics using  $k = 20$  from LDA and TKM for the 20 Newsgroups dataset were assessed qualitatively.

### 5.3 Results

**Experiment 1:** The number of distinct topics  $|DT|$  of TKM converges when the maximal number of topics  $k$  increases (Figure 2).<sup>4</sup> For all datasets LDA showed a significant weaker tendency to limit the number of topics. The fact that topics are typically investigated manually and humans often assign fewer than 100

<sup>4</sup>Figure 2 was created using small offsets on the x-axis for the sake of clarity.



**Figure 2:** Significantly distinct topics given an upper bound.

categories for each dataset (Table 2) suggests that discovering more than a few hundred topics seems less suitable. Furthermore, TKM achieves higher classification accuracy with fewer topics (see results of Experiment 2). LDA’s convergence seems to depend more on the number of documents; that of TKM’s more on the number of unique words. The number of topics  $|DT|$  weakly depends on  $\alpha$  – more so for TKM. LDA showed no significant changes, even when lowering  $\alpha$  to 1/500 of the recommended value of  $50/k$  [7]. Limiting the number of discovered topics, i.e., returning very similar topics, is an intrinsic property of TKM. A word influences the topic of nearby words directly and those of more distant words indirectly. Words with high keyword scores for a topic tend to pull other words to the same topic.

**Experiment 2:** On average, TKM outperforms LDA and TKM on all metrics as shown in Table 3. For classification TKM dominates or matches LDA except for the Ohsumed dataset. The accuracy for Ohsumed is low for all techniques. It seems that there are too many terms that are shared across different categories, which seems to impact TKM and BTM more strongly than LDA. TKM does much better on the Brown corpus and the WikiSmall corpus than LDA. Both are characterized by limited redundancy, i.e., a relatively large diversity of categories and words given the number of documents. BTM seems somewhat better than LDA but worse than TKM for classification except for the BookReviews dataset, which is characterized by relatively short documents that seem to be better handled by BTM, which was designed for short texts. LDA, BTM and TKM can be made to do better on specific datasets, but at the price of classifying worse

Dataset	Classification Accuracy			PMI			Distinct Topics			Topic Change Probability			Training Time[Min]			Inference Time[Min]		
	TKM	BTM	LDA	TKM	BTM	LDA	TKM	BTM	LDA	TKM	BTM	LDA	TKM	BTM	LDA	TKM	BTM	LDA
20Newsg.	<b>0.79</b>	0.78	0.76	<b>1.58</b>	0.85	1.33	<b>99</b>	99	100	<b>0.14</b>	0.65	0.62	<b>3.7</b>	66	10	<b>0.6</b>	2.9	2.0
Reuters.	<b>0.9</b>	0.87	0.83	<b>1.36</b>	1.26	1.28	<b>95</b>	100	100	<b>0.14</b>	0.63	0.59	<b>1.1</b>	17	2.3	<b>0.3</b>	0.9	0.6
WebKB4U.	0.81	0.81	0.81	1.34	1.33	1.33	<b>93</b>	100	100	<b>0.25</b>	0.75	0.72	<b>0.8</b>	16	1.9	<b>0.1</b>	0.6	0.3
WikiBig.	<b>0.95</b>	0.93	0.93	<b>2.17</b>	1.33	1.53	100	100	100	<b>0.13</b>	0.62	0.53	<b>20</b>	516	71	<b>2.0</b>	19	14
BrownCo.	<b>0.45</b>	0.42	0.3	1.34	1.21	1.37	<b>92</b>	100	100	<b>0.38</b>	0.88	0.79	<b>1.0</b>	15	2.1	<b>0.0</b>	0.6	0.2
Ohsumed.	0.24	0.23	<b>0.36</b>	2.2	2.07	<b>2.43</b>	<b>99</b>	100	100	<b>0.12</b>	0.66	0.66	<b>2.9</b>	66	9.1	<b>0.7</b>	3.0	2.0
Categor.	<b>0.9</b>	<b>0.88</b>	0.87	<b>1.5</b>	1.41	1.44	<b>90</b>	100	100	<b>0.24</b>	0.78	0.78	<b>0.6</b>	9.6	1.3	<b>0.1</b>	0.4	0.3
BookRev.	0.77	<b>0.79</b>	0.75	<b>1.42</b>	1.3	1.37	<b>73</b>	100	100	<b>0.24</b>	0.7	0.67	<b>8.1</b>	146	25	<b>5.0</b>	11	14
WikiSma.	<b>0.85</b>	0.81	0.75	<b>1.81</b>	1.31	1.5	100	100	100	<b>0.14</b>	0.62	0.62	<b>5.7</b>	127	17	<b>1.0</b>	5.1	3.0

**Table 3:** Comparison of TKM, BTM and LDA. ‘Bold’ is better at a 99.5% significance level.

on others.

PMI assesses co-occurrence and TKM reports the best results. TKM’s word-topic distribution  $p_{hu}$  (Equation 7) is based on weighing words by frequency and concentration. Thus, words that are frequently assigned to one topic but also to others might still get a low score. In turn, LDA and BTM focus on frequency only. It is intuitive that words that are frequent within a few topics only co-occur with higher likelihood and, thus, have a larger PMI score than words that are frequent in many topics.

The fact that TKM tends to learn less complex models, i.e., discovers fewer topics, has been investigated in depth in Experiment 1. Less complex models are preferable when the models perform well. The topic change probability (Equation 14) is significantly lower for TKM, indicating more consistent assignments of words to topics as visualized in Figure 1 for the initial words of 30 documents of the 20Newsgroups dataset. TKM tends to assign the same topic to a sequence of words, if one word in the sequence has a high keyword score for that topic. This is often the case, i.e., TKM switches topics only every 5 to 8 words. LDA and BTM switch topics after almost every other word. LDA and BTM have a parameter that influences the number of topics per document. However, as expected for a bag-of-words model, even if a document covers few topics, the change probability tends to be high. Tuning parameters have little impact, i.e., we used  $\alpha = 5/k$  to reduce the number of topics (compared to the recommended value of  $50/k$  [7]). We found that making fewer topics per document (lower  $\alpha$ ) results in worse classification results, but still does not reduce the number of topic changes significantly.

The computational performance is clearly best for TKM. During training in our baseline implementation, the complexity per iteration is  $O(L \cdot k \cdot (\sum_{d \in D} |d|))$ . For LDA [7] it is only  $O(k \cdot (\sum_{d \in D} |d|))$ . The stopping criterion is the same for both algorithms. TKM requires fewer iterations and computations are simple,

e.g., TKM does not sample from distributions as LDA does. Inference time of document-topic distributions is faster for TKM, as we only process a document once, whereas LDA potentially requires multiple iterations. For BTM a matrix containing all biterns is needed. This alone makes computation fairly expensive.

Experiment 3: Table 5 shows the highest ranked words per topic of LDA and TKM, i.e., using  $p_{hu}(w|t)$ . Using cosine similarity, topics of both methods were matched in Table 5. The categories for the 20 Newsgroups dataset are shown in Table 4. We noted considerable variation of topics for LDA and TKM for each execution of the algorithms but overall found qualitatively comparable results. Both methods tend to find some expected topics as suggested in Table 4 but missing or mixing others. Overall, LDA tends to give high probabilities to general words more often, e.g., ‘need’, ‘problem’, ‘go’, ‘com’, whereas TKM prefers rather specific words, e.g., ‘gif’, ‘duke’, ‘Laurentian’, ‘BMW’, ‘HIV’. TKM tends to mix topics using indicative words of different topics. LDA tends to find topics that are uninformative owing to the generality of the words. The first 14 topics in Table 5 are rather similar for both methods except Topic 12 for TKM, which makes limited sense. Topic 3 of LDA mixes autos and motorcycles. For LDA, topics 15 to 20 seem to make limited sense, as most words are common and non-characteristic of any topic. For TKM topics 15 to 20 seem to be slightly better, i.e., they tend to be more of a mix of topics with one topic often dominating. For example, topic 16 has elements of religion and sexuality, topic 18 of hockey, topic 19 of atheism and windows.x and topic 20 of autos.

## 6 Related Work

Hofmann [11] introduced probabilistic latent semantic analysis (PLSA) as an improvement over latent semantic analysis. Its generalization LDA [4] adds priors with hyperparameters to sample from distributions of

comp.graphics, comp.os.ms-windows.misc, comp.sys.ibm.pc.hardware, comp.sys.mac.hardware, comp.windows.x, rec.autos, rec.motorcycles, rec.sport.baseball, rec.sport.hockey, sci.crypt, sci.electronics, sci.med, sci.space, misc.forsale,talk.politics.misc,talk.politics.guns, talk.politics.mideast, talk.religion.misc, alt.atheism, soc.religion.christian

**Table 4:** The 20 Newsgroups

	TKM Topics	LDA Topics	Sim
1	ax giz tax myer presid think chz go pl ms	ax giz chz gk pl fij uy fyn ah ei	1.0
2	armenian turkish armenia turk turkei azerbaijan azeri	armenian turkish muslim turkei turk armenia peopl war	0.78
3	com bike motorcycl dod bmw ride ca write biker articl	car com bike write articl edu engin ride drive new	0.65
4	chip clipper encryption wire privaci nsa kei escrow	kei us com chip clipper secur encryption would system	0.64
5	game stephanopoulo playoff pitch espn score pitcher pt	game team plai player year edu ca win hockei season	0.64
6	monitor simm mhz mac card centri duo us edu modem	card drive us edu mac driver system work disk problem	0.64
7	gun firearm atf weapon stratu fire handgun bd amend edu	gun com edu would fire write articl peopl koresh fbi	0.64
8	church christ scriptur bibl koresh faith sin god cathol	god christian jesu bibl church us christ sai sin love	0.64
9	medic cancer drug patient doctor hiv health newslett	us medic studi diseas patient effect drug doctor food	0.63
10	edu israel isra arab palestinian articl write adl uci	right israel state isra peopl arab edu war write jew	0.59
11	space drive scsi orbit shuttl disk nasa id mission hst	space nasa gov launch orbit earth would mission moon edu	0.55
12	peopl us would batteri like jew right know make year	presid would state us mr think go year monei tax	0.49
13	imag jpeg gif format file graphic xv color pixel viewer	imag us window graphic avail system server softwar data	0.47
14	window mous font server microsoft driver client xterm us	file us window program imag entri format jpeg need	0.44
15	oil kuwait ac write rushdi edu islam uk entri contest	edu write articl com know uiuc cc would anyone cs	0.33
16	homosexu rutger cramer christian sexual gai msg food edu	edu peopl sai write would think com articl moral us	0.33
17	insur kei detector radar duke de phone system edu ripem	book post list mail new edu inform send address email	0.24
18	team player jesu plai water roger edu laurentian hockei	us would like write edu articl good com look get	0.23
19	moral atheist god widget atheism openwindow belief exist	go like get would time know us sai peopl think	0.14
20	car brake candida yeast engin militia vitamin steer	us com edu need power work help ca mous anyone	0.12

**Table 5:** Topics by TKM and LDA for 20Newsgroups dataset for  $k = 20$

topics and words. LDA has been extended and varied in many ways, e.g., [14, 18, 6, 3, 25, 26, 22]. Whereas our model has little in common with PLSA and LDA except its rooting in the aspect model, extensions and modifications of LDA and PLSA mostly did not touch upon key generative aspects such as how document-topic distributions are determined. Most extensions with the exception of [8, 1, 20, 9, 23, 24, 25] rely on the bag-of-words assumption. In contrast, we assume that a word influences topics of nearby words. In [8] each sentence is assigned to one topic using a Markov chain to model topic-to-topic transitions after sentences. Multiple works have used bigrams for latent topic models, e.g., [1, 20, 9, 23]. [1, 20, 9] use a bigram model to improve PLSA for speech recognition. For a bigram  $(w_i, w_j)$  [1, 20] both multiply probabilities containing conditionals  $w_i|w_j$ . [9] models  $p(t|w_i, w_j) \propto p(w_i|t) \cdot p(w_j|t)$ . They use distanced n-grams, i.e., for a fixed distance  $d$  between two words they estimate a probability distribution  $p(w_i|t)$  using all word pairs at distance  $d$  from the corpus. N-gram statistics and latent topic variables have been combined in [23] and later work, e.g., [24, 25]. A key underlying modeling assumption of [23] is inferring the probability of one word given its predecessor using smoothed bigram estimators. The sparsity of short

texts was the motivation in [25] to use biterns yielding the BTM model. In the BTM model the probability of a bitern equals  $p(w_i, w_j) = \sum_t p(t) \cdot p(w_i|t) \cdot p(w_j|t)$ . In [25] an increases of the time complexity by about a factor of 3 is reported together with improvements otherwise. We also consider all biterns within a window and therefore also compare against [25]. Aside from that, there are few similarities.

Relatively little work has been conducted on limiting the number of topics. The hierarchical topic models [22, 17] do not require the specification of the number of topics but come with increased complexity. Keyword extraction often relies on using co-occurrence data, POS tagging [21] and external sources or TF-IDF [10]. Typically, these methods first extract key-phrases and then rank them. For example, [21] first splits words into phrases using sentence delimiters and stop words. They compute keywords using the ratio of the frequency of a word as well as its degree, i.e., all words that are within a specific distance for any occurrence within a phrase. We do not use any of the typical preprocessing, e.g., POS tagging or phrase extraction, though this might be beneficial. We also tested the metric of [21] and found that it gave overall slightly worse result. Topic modeling and keyword extraction are related. For example [15]

extracts keywords using topic-word distributions obtained from LDA. Key word extraction and clustering are also related. [16] uses clustering as a preprocessing step to obtain keyword candidates stemming from medoids of these clusters. [16, 15] stick to the concept that keywords are extracted based on a preprocessing phase. In contrast, we perform a dynamic iterative approach, in which words are assigned to topics and the distribution of word assignments across topics determines the keyword score. [13] uses term-weighting to enhance topic modeling, i.e., LDA. Their keyword score corresponds to the variance of the word-topic distribution. They do not state a thorough derivation of their Gibbs sampler, i.e., an explicit integration of Equation (4) in [13] to get (6) and (7). Their classification performance reported on the 20Newsgroup dataset for their best algorithm is significantly lower than the performance we observed for LDA (as well as that of our algorithms).

## 7 Conclusions

We presented a novel topic model based on keywords. Both the model and its inference are simple and performant. Our experiments report improvements to existing techniques across a number of metrics, thus suggesting that the proposed technique aligns closer with human intuition about topics and keywords.

## References

- [1] M. Bahrani and H. Sameti. A new bigram-plsa language model for speech recognition. *EURASIP Journal on Applied Signal Processing*, 2010, 2010.
- [2] D. Barber. *Bayesian reasoning and machine learning*. Cambridge University Press, 2012.
- [3] D. M. Blei, T. L. Griffiths, and M. I. Jordan. The nested chinese restaurant process and bayesian nonparametric inference of topic hierarchies. *Journal of the ACM (JACM)*, 57(2):7, 2010.
- [4] D. M. Blei, A. Y. Ng, and M. I. Jordan. Latent dirichlet allocation. *Journal of machine Learning research*, 3(Jan):993–1022, 2003.
- [5] J. Chang, J. L. Boyd-Graber, S. Gerrish, C. Wang, and D. M. Blei. Reading tea leaves: How humans interpret topic models. In *Conf. on Neural Information Processing Systems(NIPS)*, pages 1–9, 2009.
- [6] R. Das, M. Zaheer, and C. Dyer. Gaussian lda for topic models with word embeddings. In *Proc. of the Annual Meeting of the Asso. for Computational Linguistics*, 2015.
- [7] T. L. Griffiths and M. Steyvers. Finding scientific topics. *Proceedings of the National academy of Sciences*, 101(suppl 1):5228–5235, 2004.
- [8] A. Gruber, Y. Weiss, and M. Rosen-Zvi. Hidden topic markov models. In *AISTATS*, pages 163–170, 2007.
- [9] M. A. Haidar and D. O’Shaughnessy. PLSA enhanced with a long-distance bigram language model for speech recognition. In *European Signal Processing Conference (EUSIPCO)*, pages 1–5, 2013.
- [10] K. S. Hasan and V. Ng. Automatic keyphrase extraction: A survey of the state of the art. In *Proc. of the Annual Meeting of the Association for Computational Linguistics*, pages 1262–1273, 2014.
- [11] T. Hofmann. Probabilistic latent semantic indexing. In *Proceedings of Research and development in information retrieval*, pages 50–57, 1999.
- [12] T. Hofmann. Unsupervised learning by probabilistic latent semantic analysis. *Machine learning*, 42(1):177–196, 2001.
- [13] S. Lee, J. Kim, and S.-H. Myaeng. An extension of topic models for text classification: A term weighting approach. In *Conf. on Big Data and Smart Computing (BigComp)*, pages 217–224, 2015.
- [14] W. Li and A. McCallum. Pachinko allocation: Dag-structured mixture models of topic correlations. In *Proc. of the int. conference on Machine learning*, pages 577–584, 2006.
- [15] Z. Liu, W. Huang, Y. Zheng, and M. Sun. Automatic keyphrase extraction via topic decomposition. In *Conference on empirical methods in natural language processing*, pages 366–376, 2010.
- [16] Z. Liu, P. Li, Y. Zheng, and M. Sun. Clustering to find exemplar terms for keyphrase extraction. In *Conference on Empirical Methods in Natural Language Processing*, pages 257–266, 2009.
- [17] D. Mimno, W. Li, and A. McCallum. Mixtures of hierarchical topics with pachinko allocation. In *Proc. of the int. conference on Machine learning*, pages 633–640, 2007.
- [18] D. Newman, E. V. Bonilla, and W. Buntine. Improving topic coherence with regularized topic models. In *Adv. in neural information processing systems*, pages 496–504, 2011.
- [19] D. Newman, J. H. Lau, K. Grieser, and T. Baldwin. Automatic evaluation of topic coherence. In *Conf. of the North American Chapter of the Asso. for Computational Linguistics*, pages 100–108, 2010.
- [20] J. Nie, R. Li, D. Luo, and X. Wu. Refine bigram plsa model by assigning latent topics unevenly. In *Workshop on Automatic Speech Recognition & Understanding*, pages 141–146. IEEE, 2007.
- [21] S. Rose, D. Engel, N. Cramer, and W. Cowley. Automatic keyword extraction from individual documents. *Text Mining*, pages 1–20, 2010.
- [22] Y. W. Teh, M. I. Jordan, M. J. Beal, and D. M. Blei. Sharing clusters among related groups: Hierarchical dirichlet processes. In *Advances in neural information processing systems*, pages 1385–1392, 2005.
- [23] H. M. Wallach. Topic modeling: beyond bag-of-words. In *Proc. of the int. conf. on Machine learning*, pages 977–984, 2006.
- [24] X. Wang, A. McCallum, and X. Wei. Topical n-grams: Phrase and topic discovery, with an application to information retrieval. In *International Conference on Data Mining (ICDM)*, pages 697–702, 2007.
- [25] X. Yan, J. Guo, Y. Lan, and X. Cheng. A biterm topic model for short texts. In *Proc. of the international conf. on World Wide Web*, pages 1445–1456, 2013.
- [26] G. Yang, D. Wen, N.-S. Chen, E. Sutinen, et al. A novel contextual topic model for multi-document summarization. *Expert Systems with Applications*, 42(3):1340–1352, 2015.