

The Polylingual Labeled Topic Model

Lisa Posch^{1,2}, Arnim Bleier¹, Philipp Schaer¹, and Markus Strohmaier^{1,2}

¹ GESIS – Leibniz Institute for the Social Sciences
Cologne, Germany

² Institute for Web Science and Technologies
University of Koblenz-Landau, Germany
{firstname.lastname}@gesis.org

Abstract. In this paper, we present the *Polylingual Labeled Topic Model*, a model which combines the characteristics of the existing *Polylingual Topic Model* and *Labeled LDA*. The model accounts for multiple languages with separate topic distributions for each language while restricting the permitted topics of a document to a set of predefined labels. We explore the properties of the model in a two-language setting on a dataset from the social science domain. Our experiments show that our model outperforms LDA and Labeled LDA in terms of their held-out perplexity and that it produces semantically coherent topics which are well interpretable by human subjects.

1 Introduction

Topic models are a popular and widely used method for the analysis of textual corpora. *Latent Dirichlet Allocation (LDA)* [2], one of the most popular topic models, has been adapted to a multitude of different problem settings, such as modeling labeled documents with *Labeled LDA (L-LDA)* [9] or modeling multilingual documents with *Polylingual Topic Models (PLTM)* [7]. Textual corpora often exhibit both of these characteristics, containing documents in multiple languages which are also annotated with a classification system. However, there is currently no topic model which possesses the ability to process multiple languages while simultaneously incorporating the documents' labels.

To close this gap, this paper introduces the *Polylingual Labeled Topic Model (PLL-TM)*, a model which combines the characteristics of PLTM and L-LDA. PLL-TM models multilingual labeled documents by generating separate distributions over the vocabulary of each language, while restricting the permitted topics of a document to a set of predefined labels. We explore the characteristics of our model in a two-language setting, with German natural language text as the first language and the controlled *SKOS* vocabulary of a thesaurus as the second language. The labels of the documents, in our setting, are classes from the classification system with which our corpus is annotated.

Contributions. The main contribution of this paper is the presentation of the PLL-TM. We present the model's generative storyline as well as an easy-to-implement inference strategy based on Gibbs sampling. For evaluation, we compute the held-out perplexity and conduct a word intrusion task with human

subjects using a dataset from the social science domain. On this dataset, the PLL-TM outperforms LDA and L-LDA in terms of its predictive performance and generates semantically coherent topics. To the best of our knowledge, PLL-TM is the first model which accounts for multiple vocabularies and, at the same time, possesses the ability to restrict the topics of a document to its labels.

2 Related Work

Topic models are generative probabilistic models for discovering latent topics in documents and other discrete data. One of the most popular topic models, LDA, is a generative Bayesian model which was introduced by Blei et al. [2]. In this section, we review LDA, as well as the two other topic models whose characteristics we are going to integrate into PLL-TM.

LDA. Beginning with LDA [2], we follow the common notation of a document d being a vector of N_d words, \mathbf{w}_d , where each word w_{di} is chosen from a vocabulary of V terms. A collection of documents is defined by $\mathcal{D} = \{\mathbf{w}_1, \dots, \mathbf{w}_D\}$. LDA’s generative storyline can be described by the following steps.

1. For each document $d \in \{1, \dots, D\}$, a distribution θ_d over topics is drawn from a symmetric K-dimensional Dirichlet prior parametrized by α :

$$\theta_d \sim Dir(\alpha) . \quad (1)$$

2. Then, for each topic $k = \{1, \dots, K\}$, a distribution ϕ_k over the vocabulary is drawn from a V-dimensional Dirichlet distribution parametrized by β :

$$\phi_k \sim Dir(\beta) . \quad (2)$$

3. In the final step, the i^{th} word in document d is generated by first drawing a topic index z_{di} and subsequently, a word w_{di} from the topic indexed by z_{di} :

$$w_{di} \sim Cat(\phi_{z_{di}}) , \quad z_{di} \sim Cat(\theta_d) . \quad (3)$$

Labeled LDA. Ramage et al. [9] introduced L-LDA, a supervised version of LDA. In L-LDA, a document d ’s topic distribution θ_d is restricted to a subset of all possible topics $\mathbf{A}_d \subseteq \{1, \dots, K\}$. Here, collection of documents is defined by $\mathcal{D} = \{(\mathbf{w}_1, \mathbf{A}_1), \dots, (\mathbf{w}_D, \mathbf{A}_D)\}$. The first step in L-LDA’s generative storyline draws the distribution of topics θ_d for each document $d \in \{1, \dots, D\}$

$$\theta_d \sim Dir(\alpha \boldsymbol{\mu}_d) , \quad (4)$$

where α is a continuous positive valued scalar and $\boldsymbol{\mu}_d$ is a K-dimensional vector

$$\mu_{dk} = \begin{cases} 1 & \text{if } k \in \mathbf{A}_d \\ 0 & \text{otherwise} , \end{cases} \quad (5)$$

indicating which topics are permitted. Once these label-restricted topic distributions are drawn, the process of generating documents continues identically to the generative process of LDA. In the case of $\mathbf{A}_d = \{1, \dots, K\}$ for all documents, no restrictions are active and L-LDA is reduced to LDA.




Fig. 1: **The PLL-TM in plate notation.** Random variables are represented by nodes. Shaded nodes denote the observed words and labels, bare symbols indicate the fixed priors α and β^l . Directed edges between the nodes then define conditional probabilities, where the child node is conditioned on its parents. The rectangular plates indicate replication over data-points and parameters. Colors indicate the parts which are inherited from L-LDA (**blue**) and PLTM (**green**). Black is used for the LDA base.

Polylingual Topic Model. Ni et al. [8] extended the generative view of LDA to multilingual documents. Mimno et al. [7] elaborated on this concept, introducing the *Polylingual Topic Model* (PLTM). PLTM assumes that the documents are available in L languages. A document d is represented by $[\mathbf{w}_d^1, \dots, \mathbf{w}_d^L]$, where for each language $l \in 1, \dots, L$, the vector \mathbf{w}_d^l consists of N_d^l words which are chosen from a language specific vocabulary with V^l terms. A collection of documents is then defined by $\mathcal{D} = \{[\mathbf{w}_1^1, \dots, \mathbf{w}_1^L], \dots, [\mathbf{w}_D^1, \dots, \mathbf{w}_D^L]\}$. The generative storyline is equivalent to LDA's except that steps 2 and 3 are repeated for each language. Hence, for each topic $k = \{1, \dots, K\}$ in each language $l \in \{1, \dots, L\}$, a language specific topic distribution ϕ_k^l over the vocabulary of length V^l is drawn:

$$\phi_k^l \sim \text{Dir}(\beta^l). \quad (6)$$

Then, the i^{th} word of language l in document d is generated by drawing a topic index z_{di}^l and subsequently, a word w_{di}^l from a language specific topic distribution indexed by z_{di}^l :

$$w_{di}^l \sim \text{Cat}(\phi_{z_{di}^l}^l), \quad z_{di}^l \sim \text{Cat}(\theta_d). \quad (7)$$

Note that in the special case of just one language, i.e. $L = 1$, PLTM is reduced to LDA.

3 The Polylingual Labeled Topic Model

In this section, we introduce the *Polylingual Labeled Topic Model (PLL-TM)*, which integrates the characteristics of the models described in the previous section into a single model. Figure 1 depicts the PLL-TM in plate notation. Here, a collection of documents is defined by $\mathcal{D} = \{[\mathbf{w}_1^1, \dots, \mathbf{w}_1^L, \mathbf{A}_1], \dots, [\mathbf{w}_D^1, \dots, \mathbf{w}_D^L, \mathbf{A}_D]\}$.

The generative process follows three main steps:

1. For each document $d \in \{1, \dots, D\}$, we draw the distribution of topics

$$\theta_d \sim \text{Dir}(\alpha \boldsymbol{\mu}_d), \quad (8)$$

where $\boldsymbol{\mu}_d$ is computed according to Equation 5.

2. For each topic $k \in \{1, \dots, K\}$ in each language $l \in \{1, \dots, L\}$, we draw a distribution over the vocabulary of size V^l :

$$\phi_k^l \sim \text{Dir}(\beta^l), \quad (9)$$

3. Next, for each word in each language l of document d , we draw a topic

$$w_{di}^l \sim \text{Cat}(\phi_{z_{di}^l}^l), \quad z_{di}^l \sim \text{Cat}(\theta_d). \quad (10)$$

Note that PLL-TM contains both PLTM and L-LDA as special cases.

For inference, we use collapsed Gibbs sampling [6] for the indicator variables \mathbf{z} , with all other variables integrated out. The full conditional probability for a topic k is given by

$$P(z_{di}^l = k \mid w_{di}^l = t, \dots) \propto \frac{n_{dk}^{-di} + \alpha}{n_d^{-di} + K\alpha} \times \frac{n_{kt}^{l-di} + \beta^l}{n_k^{l-di} + V^l \beta^l}, \quad (11)$$

where n_{dk} is the number of tokens allocated to topic k in document d , and n_{kt}^l is the number of tokens of word $w_{di}^l = t$ which are assigned to topic k in language l . Furthermore, \cdot is used in place of a variable to indicate that the sum over its values (i.e. $n_d = \sum_k n_{dk}$) is taken and $-di$ to mark the current token as excluded. While the full conditional posterior distribution is reminiscent of the one used in PLTM, the assumptions of the L-LDA model restrict the probability $P(z_{di}^l = k)$ to those $k \in \mathbf{A}_d$ with which document d is labeled.

4 Evaluation

For our evaluation, we use documents from the *Social Science Literature Information System (SOLIS)*. The documents are manually indexed with the SKOS

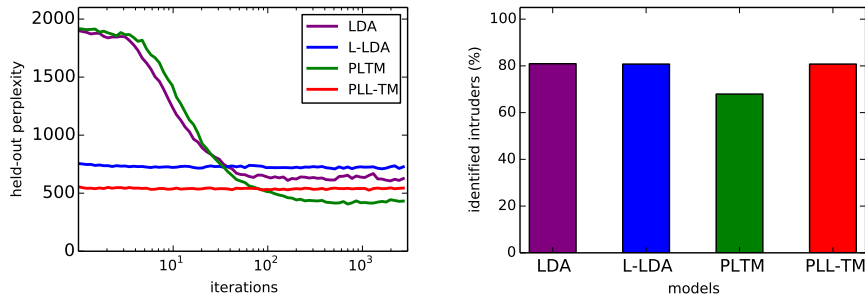
Table 1: This table shows the five most probable terms for two classes in the CSS, generated by PLL-TM, in two languages: *TheSoz* (TS) and German natural language words with their translation (AB).

Population Studies, Sociology of Population:

TS: population development, demographic aging, population, demographic factors, demography
 AB: wandel, demografischen, bevölkerung, deutschland, entwicklung
 (change, demographic, population, germany, development)

Developmental Psychology:

TS: child, developmental psychology, adolescent, personality development, socialization research
 AB: entwicklung, sozialisation, kinder, kindern, identität
 (development, socialization, children, children, identity)



(a) Comparison of the held-out perplexity (lower values are better) as a function of iterations. (b) Comparison of the semantic coherence (word intrusion) of the generated topics.

Fig. 2: **Evaluation of the PLL-TM.** These figures show that on the SOLIS dataset, PLL-TM outperforms LDA and L-LDA in terms of its predictive performance and produces topics with a higher semantic coherence than PLTM.

Thesaurus for the Social Sciences (TheSoz) [10] and manually classified with the *Classification for the Social Sciences (CSS)* by human domain experts. For our experiments, we used all SOLIS documents which were published in the years 2008 to 2013, resulting in a corpus of about 60.000 documents.

We explore the characteristics of our model in a two-language setting, with German natural language text as the first language (*AbstractWords*) and the controlled *SKOS* vocabulary of a thesaurus as the second language (*TheSoz*). The labels of the documents, in our setting, are classes from the CSS. After applying standard preprocessing to remove rare words and stopwords, *TheSoz* consisted of 802.764 tokens over a vocabulary of 7.406 distinct terms, and *AbstractWords* consisted of 5.417.779 tokens over a vocabulary of about 43.000 distinct terms. In our corpus, each document is labeled with an average of 2.14 classes.

We compare four different topic models: LDA, L-LDA, PLTM and PLL-TM. The unilingual models (i.e. LDA and L-LDA) were trained on language *TheSoz*; the polylingual models (i.e. PLTM and PLL-TM) were trained on *TheSoz* and *AbstractWords*. The documents in our corpus were labeled with a total of 131 different classes from the CSS and we trained the unlabeled models with an equal number of topics. α and β^l were specified with 0.1 and 0.01, respectively. Table 1 shows the topics generated by PLL-TM for two classes of the CSS, reporting the five most probable terms for the languages *TheSoz* and *AbstractWords*.

Language Model Evaluation. For an evaluation of the predictive performance, we computed the held-out perplexity for all models. We held out 1.000 documents as test set \mathcal{D}_{test} and, with the remaining data \mathcal{D}_{train} , we trained the four models. We split each test document in the following way:

- \mathbf{x}_{d1} : All words of language *AbstractWords* and a randomly selected 50% of the words in language *TheSoz* which occur in document d .

- \mathbf{x}_{d2} : The remaining 50% of the words in language *TheSoz* which occur in document d .

The test documents for the unilingual models were split analogously, with \mathbf{x}_{d1} consisting of 50% percent of the words in language *TheSoz* which occur in document d . For each document d , we computed the perplexity of \mathbf{x}_{d2} .

Figure 2a shows the results of this evaluation. One can see that the labeled models both start out with a lower perplexity and need less iterations to achieve a good performance, which is due to the fact that the labels provide additional information to the model. In contrast, the unlabeled models need almost 100 iterations to achieve a comparable performance. On our corpus, PLL-TM outperformed LDA and L-LDA, and even though PLL-TM had a higher perplexity than PLTM, it is important to keep in mind that PLTM does not possess the ability to produce topics which correspond to the classes of the CSS.

Human Evaluation of the Topics. Chang et al. [4] proposed a formal setting in which humans evaluate the latent space of a topic model. For evaluating the topics’ semantic coherence, they proposed a *word intrusion* task: Crowdworkers were shown six terms, five of which were highly probable terms in a topic and one was an “intruder” – an improbable term for this topic which had a high probability in some other topic.

We conducted the word intrusion task for the four topic models on CrowdFlower [1], with ten distinct workers for each topic in each model. Figure 2b shows the results of this evaluation for the different models. For each model, the figure depicts the percentage of topics for which the ten workers collectively detected the correct intruder. The collective decision was based on CrowdFlower’s *confidence score*, i.e. the level of agreement between workers weighted by each worker’s percentage of correctly answered test questions. The results show that PLL-TM produces topics which are equally coherent as unilingual models, and more coherent than the topics produced by PLTM.

5 Discussion and Conclusions

In this paper, we presented PLL-TM, a joint model for multilingual labeled documents. The results of our evaluation showed that PLL-TM was the only model which produced both highly interpretable topics and achieved a good predictive performance. Compared to L-LDA, the only other model capable of incorporating label information, our model produced equally well interpretable topics while achieving a better predictive performance. Compared to PLTM, the only other model capable of dealing with multiple languages, PLL-TM had a lower predictive performance, but produced topics with a higher semantic coherence. For future work, we plan an evaluation of the model in a label prediction task and an application of the model in a setting with more than two natural languages. Furthermore, we plan an evaluation on a larger dataset using a more memory-friendly inference strategy such as *Stochastic Collapsed Variational Bayesian Inference* [5], which has been shown to be applicable outside of its original LDA application [3].

References

1. L. Biewald. Massive multiplayer human computation for fun, money, and survival. In *Current Trends in Web Engineering - Workshops, Doctoral Symposium, and Tutorials, Held at ICWE 2011, Paphos, Cyprus, June 20-21, 2011. Revised Selected Papers*, pages 171–176, 2011.
2. D. M. Blei, A. Y. Ng, and M. I. Jordan. Latent dirichlet allocation. *Journal of Machine Learning Research*, 3:993–1022, 2003.
3. A. Bleier. Practical collapsed stochastic variational inference for the hdp. In *NIPS Workshop on Topic Models: Computation, Application, and Evaluation*, 2013.
4. J. Chang, J. L. Boyd-Graber, S. Gerrish, C. Wang, and D. M. Blei. Reading tea leaves: How humans interpret topic models. In *Advances in Neural Information Processing Systems 22: 23rd Annual Conference on Neural Information Processing Systems 2009. Proceedings of a meeting held 7-10 December 2009, Vancouver, British Columbia, Canada.*, pages 288–296, 2009.
5. J. R. Foulds, L. Boyles, C. DuBois, P. Smyth, and M. Welling. Stochastic collapsed variational bayesian inference for latent dirichlet allocation. In *The 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD 2013, Chicago, IL, USA, August 11-14, 2013*, pages 446–454, 2013.
6. T. L. Griffiths and M. Steyvers. Finding scientific topics. *Proceedings of the National Academy of Sciences*, 2004.
7. D. M. Mimno, H. M. Wallach, J. Naradowsky, D. A. Smith, and A. McCallum. Polylingual topic models. In *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing, EMNLP 2009, 6-7 August 2009, Singapore, A meeting of SIGDAT, a Special Interest Group of the ACL*, pages 880–889, 2009.
8. X. Ni, J. Sun, J. Hu, and Z. Chen. Mining multilingual topics from wikipedia. In *Proceedings of the 18th International Conference on World Wide Web, WWW 2009, Madrid, Spain, April 20-24, 2009*, pages 1155–1156, 2009.
9. D. Ramage, D. L. W. Hall, R. Nallapati, and C. D. Manning. Labeled LDA: A supervised topic model for credit attribution in multi-labeled corpora. In *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing, EMNLP 2009, 6-7 August 2009, Singapore, A meeting of SIGDAT, a Special Interest Group of the ACL*, pages 248–256, 2009.
10. B. Zapolko, J. Schaible, P. Mayr, and B. Mathiak. Thesoz: A SKOS representation of the thesaurus for the social sciences. *Semantic Web*, 4(3):257–263, 2013.