
Prediction Focused Topic Models via Feature Selection

Jason Ren*

jason_ren@college.harvard.edu

Russell Kunes*

rk3064@columbia.edu

Finale Doshi-Velez

finale@seas.harvard.edu

Abstract

Supervised topic models are often sought to balance prediction quality and interpretability. However, when models are (inevitably) misspecified, standard approaches rarely deliver on both. We introduce a novel approach, the prediction-focused topic model, that uses the supervisory signal to retain only vocabulary terms that improve, or at least do not hinder, prediction performance. By removing terms with irrelevant signal, the topic model is able to learn task-relevant, coherent topics. We demonstrate on several data sets that compared to existing approaches, prediction-focused topic models learn much more coherent topics while maintaining competitive predictions.

1 Introduction

Topic models such as Latent Dirichlet Allocation (LDA) [Blei et al., 2003] learn a small set of topics that capture the co-occurrence of terms in discrete count data. Given the topics, any datum can be represented as a mixture over topics; these mixture weights can then be used as input for some downstream supervised task. By reducing the dimensionality of the data, topic models help make predictions more interpretable, especially if the topics are themselves interpretable. This makes topic models valuable in contexts such as health care (e.g. Hughes et al. [2017a]) or criminal justice (e.g. Kuang et al. [2017]) where interpretability is important for downstream validation.

However, learning topics that are both interpretable and predictive is tricky, and there generally exists a trade-off between the ability of the topics to predict the target well and explain the count data well. In real world settings in which the data do not truly come from a topic modeling generative process, the signal captured by a standard topic model may not be predictive of the target. Supervised topic models attempt to mitigate this issue by including the targets during training, but the effect of including the targets is often minimal due to a cardinality mismatch between the data and the targets ([Zhang and Kjellström, 2014], [Halpern et al., 2012]).

In this work, we focus on one common reason why topic models fail: documents often contain common, correlated terms that are irrelevant to the task. For example, a topic model trained on movie reviews may form a topic that assigns high probability to words like “comedy”, “action”, “character”, and “plot,” which may be nearly orthogonal to a sentiment label. The existence of features irrelevant to the supervised task complicates optimization of the trade-off between prediction quality and explaining the count data, and also renders the topics less coherent.

To address this issue, we introduce a novel supervised topic model, prediction-focused sLDA (pf-sLDA), that explicitly allows features irrelevant to predicting the response variable to be ignored when learning topics. During training, pf-sLDA simultaneously learns which features are irrelevant and fits a supervised topic model on the relevant features. It modulates the trade-off between prediction quality and explaining the count data with an interpretable model parameter. We demonstrate that compared to existing approaches, pf-sLDA is able to learn more coherent topics while maintaining competitive prediction accuracy on several data sets.

2 Related Work

There are two main lines of related work: work to improve prediction quality in supervised topic models, and work to focus topics toward areas of interest.

Improving prediction quality. Since the original supervised LDA (sLDA) work of McAuliffe and Blei [2008], many works have incorporated information about the prediction target into the topic model training process in different ways to improve prediction quality, including power-sLDA [Zhang and Kjellström, 2014], med-LDA [Zhu et al., 2012], and BP-sLDA [Chen et al., 2015].

Hughes et al. [2017b] pointed out a number of shortcomings of these previous methods and introduced a new objective that weighted a combination of the conditional target likelihood and marginal data likelihood: $\lambda \log p(\mathbf{y}|\mathbf{w}) + \log p(\mathbf{w})$. They demonstrated the resulting method, termed prediction-constrained sLDA (pc-sLDA), achieves better empirical results in optimizing the trade-off between prediction quality and explaining the count data and justify why this is the case. However, irrelevant but common terms are not handled well in their formulation; their topics are often polluted by irrelevant terms. Our pf-sLDA objective enjoys analogous theoretical properties but effectively removes irrelevant terms, and thus achieves more coherent topics.

There are some supervised topic modeling approaches that model the existence of some sort of irrelevance in the data. NUF-sLDA [Zhang and Kjellström, 2014] models entire topics as noise. Topic models that encourage sparsity in topic vocabulary distributions, such as dual sparse topic models [Lin et al., 2014], reduce the number of features per topic but do not easily allow an irrelevant feature to be ignored across all topics. To our knowledge, our approach is the first supervised topic modeling approach in which the model is able to learn which features are irrelevant to the supervised task and ignore them when fitting the topics.

Focused topics. The notion of focusing topics in relevant directions is also present in the unsupervised topic modeling literature. For example, Wang et al. [2016] focus topics by seeding them with keywords; Kim et al. [2012] introduce variable selection for LDA, which models some of the vocabulary as irrelevant. Fan et al. [2017] similarly develop stop-word exclusion schemes. However, these approaches adjust topics based on some general notions of “focused,” whereas the goal of our pf-sLDA is to remove terms to explicitly manage a trade-off between pre-

Symbol	Meaning
K	Number of topics
M	Number of documents
V	Vocabulary
\mathbf{w}	Documents (bag of words)
\mathbf{y}	Targets
β	Topics
θ	Topic-document distributions
α	Topic prior parameter
η, δ	Target GLM parameters
ξ	pf-sLDA channel switches
p	pf-sLDA word inclusion prior
π	pf-sLDA additional topic

Table 1: Notation

diction quality and explaining the count data.

3 Background and Notation

We briefly review supervised Latent Dirichlet Allocation (sLDA) [McAuliffe and Blei, 2008]. sLDA models count data (words) as coming from a mixture of K topics $\{\beta_k\}_{k=1}^K$, where each topic $\beta_k \in \Delta^{|V|-1}$ is a categorical distribution over a vocabulary V of $|V|$ discrete words. The count data are represented as a collection of M documents, with each document $\mathbf{w}_d \in \mathbb{N}^{|V|}$ being a vector of counts over the vocabulary. Each document \mathbf{w}_d is associated with a target y_d . Additionally, each document has an associated topic distribution $\theta_d \in \Delta^{K-1}$, which generates both the words and the target.

The generative process is:

```

for each document do
  Draw topic distribution  $\theta \sim \text{Dir}(\alpha)$ 
  for each word do
    Draw topic  $z \sim \text{Cat}(\theta)$ 
    Draw word  $w \sim \text{Cat}(\beta_z)$ 
  Draw target  $y \sim \text{GLM}(\theta; \eta, \delta)$ 

```

4 Prediction Focused Topic Models

We now introduce prediction-focused sLDA (pf-sLDA), a novel, dual-channel topic model. The fundamental assumption that pf-sLDA builds on is that the vocabulary V can be divided into two disjoint components, one of which is irrelevant to predicting the target variable. pf-sLDA separates out the words irrelevant to predicting the target, even if they have latent structure, so that the topics can focus on only modeling structure that is relevant to predicting the




Figure 1: *Left:* pf-sLDA graphical model. *Right:* pf-sLDA generative process.

target. The goal is to learn interpretable, focused topics that result in high prediction quality.

Generative Model. The pf-sLDA latent variable model has the following components: one channel of pf-sLDA models the count data as coming from a mixture of K topics $\{\beta_k\}_{k=1}^K$, identical to sLDA. The second channel of pf-sLDA models the data as coming from an additional topic $\pi \in \Delta^{|V|-1}$. (This modeling choice of a simple second channel is explained later in this section.) The target only depends on the first channel, so the second channel acts as an outlet for words irrelevant to predicting the target. We constrain β and π such that $\beta_k^\top \pi = 0$ for all k , such that each word is always either relevant or irrelevant to predicting the target. Which channel a word comes from is determined by its corresponding Bernoulli switch ξ , which has prior p . The generative process of pf-sLDA is given in Figure 1. In Appendix 9.4, we prove that a lower bound to the pf-sLDA likelihood is:

$$\log p(\mathbf{y}, \mathbf{w}) \geq E_{\xi}[\log p_{\beta}(\mathbf{y}|\mathbf{w}, \xi)] + pE_{\theta}[\log p_{\beta}(\mathbf{w}|\theta)] + (1-p)\log p_{\pi}(\mathbf{w}) \quad (1)$$

Here, the notation emphasizes the fact that the conditional pmf p_{β} is free of π and the multinomial pmf p_{π} is free of β .

Connection to prediction-constrained models. The lower bound above reveals a connection to the pc-sLDA loss function of Hughes et al. [2017b]. The first two terms capture the trade-off between performing the prediction task $E_{\xi}[\log p_{\beta}(\mathbf{y}|\mathbf{w}, \xi)]$ and explaining the words $pE_{\theta}[\log p_{\beta}(\mathbf{w}|\theta)]$, where the word inclusion prior p is used to down-weight explaining the words (or emphasize the prediction task). This is analogous to the prediction-constrained objective, but we manage the trade-off through an interpretable model parameter, the

```

for each document do
  Draw topic distribution  $\theta \sim \text{Dir}(\alpha)$ 
  for each word do
    Draw topic  $z \sim \text{Cat}(\theta)$ 
    Draw switch  $\xi \sim \text{Bern}(p)$ 
    if  $\xi = 1$  then
      Draw word  $w \sim \text{Cat}(\beta_z)$ 
    else if  $\xi = 0$  then
      Draw word  $w \sim \text{Cat}(\pi)$ 
  Draw target  $y \sim \text{GLM}(\theta; \eta, \delta)$ 
  (See Appendix 9.5 for details)

```

word inclusion prior p , rather than a more arbitrary Lagrange multiplier λ . Additionally, unlike pc-sLDA, our formulation is still a valid graphical model, and thus we can perform inference using standard Bayesian techniques. We describe similar connections in the true likelihood, rather than the lower bound, in Appendix 9.6.

Choosing a simple second channel. The second term and third term capture the trade-off between how well a word can be explained by the relevant topics β versus the additional topic π , weighted by p and $1-p$ respectively. This trade-off motivates why it is important to keep the additional topic π simple (e.g. a single topic). If we fix a simple form for π , then π will not be able to explain words as well as β . Thus, there is an inherent cost in considering words irrelevant (i.e. have mass in π and not in β). Words will be considered irrelevant only if including them in the relevant topics β hinders the prediction of the target y in a way that outweighs the cost of considering them irrelevant given p . Decreasing p allows the model to emphasize prediction quality, but it also lowers the cost of considering a word irrelevant due to the weighting terms. A more expressive π would lower the cost of considering words irrelevant even more and causes the model to put even relevant words into π .

All this together informs how pf-sLDA behaves. The word inclusion prior p and the additional topic π have specific interpretations based on generative process: the word inclusion prior p is the prior probability that a word is relevant (has mass in β), while the additional topic π contains the probabilities of each irrelevant word. However, in the context of the loss, they play the role of governing the trade-off between prediction quality and explaining the count data, *while still keeping the model a standard generative model*. As we decrease p , we encourage fewer

words to be considered relevant, with the goal of improving prediction quality by ignoring irrelevant words. This allows pf-sLDA to improve prediction quality while maintaining the ability to model relevant words well. We can then tune p with the goal of encouraging the model to consider the truly irrelevant words to be irrelevant without losing relevant signal. (See Figure 2 in the results for a demonstration of this effect.)

5 Inference

Inference in the pf-sLDA framework corresponds to inference in a graphical model, so advances in Bayesian inference can be applied to solve the inference problem. In this work, we take a variational approach. Our objective is to maximize the evidence lower bound (ELBO), with the constraint that the relevant topics β and additional topics π have disjoint support. The key difficulty is that of optimizing over the non-convex set $\{\beta, \pi : \beta^\top \pi = \mathbf{0}\}$. We resolve this with a strategic choice of variational family, which results in a straightforward training procedure that does not require any tuning parameters. In section 5.1, we discuss model properties that inform our decision of the variational family. In section 5.2, we examine how this choice of variational family and optimization enforce our desired constraint.

5.1 Properties and Variational Family

We first note some properties of pf-sLDA that will help us choose our variational family.

Theorem 1. *If the channel switches ξ_d and the document topic distribution θ_d are conditionally independent in the posterior for all documents, then β and π have disjoint supports over the vocabulary.*

Theorem 2. *$\beta^\top \pi = \mathbf{0}$ if and only if there exists an assignment to all channel switches ξ^* s.t. $p(\xi^* | \mathbf{w}, \mathbf{y}) = 1$*

See Appendix 9.2 for proofs.

Theorems 1 and 2 tell us that if the posterior distribution of the channel switches ξ_d is independent of the posterior distribution of the document topic distribution θ_d for all documents, then the true relevant topics β and additional topic π must have disjoint support, and moreover the posterior of the channel switches ξ is a point mass. This suggests that to enforce that our generative model constraint that β and π are disjoint, we should choose the variational family such that ξ and θ are independent.

If ξ and θ are conditionally independent in the posterior, then the posterior can factor as $p(\xi, \theta | \mathbf{y}, \mathbf{w}) = p(\xi | \mathbf{y}, \mathbf{w})p(\theta | \mathbf{y}, \mathbf{w})$. In this case, the posterior for the channel switch of the n th word in document d , ξ_{dn} , has no dependence d , which can be seen directly from the graphical model. Thus, choosing $q(\xi | \varphi)$ to have no dependence on document naturally pushes our assumption into the variational posterior. We specify the full form of our variational family below:

$$q(\theta, \mathbf{z}, \xi | \phi, \varphi, \gamma) = \prod_d q(\theta_d | \gamma_d) \prod_n q(\xi_{dn} | \varphi) q(z_{dn} | \phi_{dn})$$

$$\theta_d | \gamma_d \sim \text{Dir}(\gamma_d)$$

$$z_{dn} | \phi_{dn} \sim \text{Cat}(\phi_{dn})$$

$$\xi_{dn} | \varphi \sim \text{Bern}(\varphi_{w_{dn}})$$

where d indexes over the documents and n indexes over the words in each document. Our choices for the variational distributions for θ and \mathbf{z} match those of McAuliffe and Blei [2008]. We choose $q(\xi_{dn} | \varphi_{w_{dn}})$ to be a Bernoulli probability mass function with parameter $\varphi_{w_{dn}}$ indexed only by the word w_{dn} . This distribution acts as a relaxation of a true point mass posterior, allowing us to use gradient information to optimize over $[0, 1]$ rather than directly over $\{0, 1\}$. Moreover, this parameterization allows us to naturally use the variational parameter φ as a feature selector; low estimated values of φ indicate irrelevant words, while high values of φ indicate relevant words. We refer to φ as our variational feature selector.

5.2 ELBO and Optimization

Inference in the pf-sLDA corresponds to optimizing the following lower bound:

$$\text{ELBO}(\Theta, \Lambda) = \log p_\Theta(\mathbf{y}, \mathbf{w}) - \text{KL}(q_\Lambda(\zeta) || p_\Theta(\zeta | \mathbf{w}, \mathbf{y}))$$

where we denote the full set of model parameters as Θ , the full set of variational parameters as Λ , and our latent variables as $\zeta := \{\xi, \theta, \mathbf{z}\}$.

We show how optimization of the ELBO, combined with our choice of variational family, will push β and π to be disjoint. First, we note that maximizing the ELBO over Θ, Λ is equivalent to:

$$\max_{\Theta} \left(\log p_\Theta(\mathbf{y}, \mathbf{w}) - \min_{\Lambda} \text{KL}(q_\Lambda(\zeta) || p_\Theta(\zeta | \mathbf{w}, \mathbf{y})) \right)$$

Therefore, estimating parameters based on the ELBO can be viewed as *penalized* maximum likelihood estimation with penalty:

$$\text{pen}(\Theta) = \min_{\Lambda} \text{KL}(q_\Lambda(\zeta) || p_\Theta(\zeta | \mathbf{w}, \mathbf{y}))$$

In the pf-sLDA model, the KL penalty for the channel switches ξ is minimized for the subset of the model parameter space $\{\beta, \pi : \beta^\top \pi = \mathbf{0}\}$. On this set, $\min_{\Lambda} \text{KL}(q_{\varphi}(\xi) \| p_{\Theta}(\xi | \mathbf{w}, \mathbf{y})) = 0$ so long as $q_{\varphi}(\xi)$ includes the set of point masses, as shown by Theorems 1 and 2. The analogous penalty everywhere else in the parameter space is greater than zero, based on the contrapositive of Theorem 1. We note that the size of this penalty and whether it is strong enough to enforce a disjoint solution for $\{\beta, \pi\}$ has to do with how restrictive the choice of variational family $q_{\Lambda}(\zeta)$ is. In our case, we found that our parameterization of $q_{\Lambda}(\zeta)$ resulted in estimates with β and π disjoint as desired.

Finally, to optimize the ELBO, we simply run mini-batch gradient descent on the ELBO (full form specified in Appendix 9.3). We also explored using more traditional closed form coordinate ascent updates (see Appendix 9.10), but we found the gradient descent approach was computationally much faster and had very similar quality results.

5.3 Prediction for new documents

At test time, we make predictions given the learned GLM parameters and the MAP estimate of the topic distribution θ for the new document (estimated through gradient descent on the posterior for θ).

6 Experimental Results

In this section, we demonstrate the ability of pf-sLDA to find more coherent topics while maintaining competitive prediction quality when compared to pc-sLDA. We also demonstrate in the well-specified case, pf-sLDA is able to reliably and accurately recover the relevant features. Finally, we show the feature filtering of pf-sLDA outperforms naive feature filtering based on correlation to target.

6.1 Experimental Set-Up

Metrics. We wish to assess prediction quality and interpretability of learned topics. To measure prediction quality, we use RMSE for real targets and AUC for binary targets. To measure interpretability of topics, we use normalized pointwise mutual information coherence, as it was found by Newman et al. [2010] to be the that metric most consistently and closely matches human judgement in evaluating interpretability of topics. See Appendix 9.9 for an explicit procedure on how we calculate coherence.

Baselines. The recent work in Hughes et al. [2017b]

demonstrates that pc-sLDA outperforms other supervised topic modeling approaches, so we use pc-sLDA as our main baseline. We also include standard sLDA [McAuliffe and Blei, 2008] for reference.

Data Sets. We run our model and baselines on three real data sets:

- Pang and Lee’s movie review data set [Pang and Lee, 2005]: 5006 movie reviews, with integer ratings from 1 (worst) to 10 (best) as targets.
- Yelp business reviews [Yelp, 2019]: 10,000 business reviews, with integer stars from 1 (worst) to 5 (best) as targets.
- Electronic health records (EHR) data set of patients with Autism Spectrum Disorder (ASD), introduced in Masood and Doshi-Velez [2018]: 3804 EHRs, with binary indicator of epilepsy as target.

For more details about data sets, see Appendix 9.7.

Implementation details. Refer to Appendix 9.5

6.2 Results

pf-sLDA learns the most coherent topics. Across data sets, pf-sLDA learns the most coherent topics by far (see Figure 2). pc-sLDA improves on topic coherence compared to sLDA, but cannot match the performance of pf-sLDA. Qualitative examination of the topics in Table 2 supports the claim that the pf-sLDA topics are more coherent, more interpretable, and more focused on the supervised task.

Furthermore, if we filter the learned topics of either sLDA or pc-sLDA post hoc based on the variational feature selectors of a trained pf-sLDA model, the topics (and coherence scores) resemble the original pf-sLDA’s. This supports the notion that pf-sLDA is focusing on relevant signal.

Prediction quality of pf-sLDA remains competitive. pf-sLDA produces similar prediction quality compared to pc-sLDA across data sets (see Figure 2). Both pc-sLDA and pf-sLDA outperform sLDA in prediction quality. In the best performing models of pf-sLDA for all 3 data sets, generally between 5% and 20% of the words were considered relevant. Considering both more words or less words relevant hurt performance, as seen in the plots in Figure 2.

In the well-specified case, pf-sLDA can recover relevant features with high precision and

Pang and Lee Movie Reviews		
Model	Coherence	RMSE
sLDA	0.362 (0.101)	1.682 (0.021)
pc-sLDA	0.492 (0.130)	1.298 (0.015)
pf-sLDA	1.987 (0.092)	1.305 (0.024)

Yelp Reviews		
Model	Coherence	RMSE
sLDA	0.848 (0.086)	1.162 (0.017)
pc-sLDA	1.080 (0.213)	0.953 (0.004)
pf-sLDA	3.258 (0.102)	0.952 (0.011)

ASD Dataset		
Model	Coherence	AUC
sLDA	1.412 (0.113)	0.590 (0.013)
pc-sLDA	2.178 (0.141)	0.701 (0.015)
pf-sLDA	2.639 (0.091)	0.748 (0.013)




Figure 2: *Left*: Mean and (SD) across 5 runs for topic coherence (higher is better) and RMSE (lower is better) or AUC (higher is better) on held-out test sets. Final models were chosen based on a combination of validation coherence and RMSE/AUC. pf-sLDA produces topics with much higher coherence across all three data sets, while maintaining similar prediction performance. *Right*: Validation coherence and RMSE on Pang and Lee Movie Review data set as $p = \frac{1}{\lambda}$ varies. This demonstrates the effect of the channel word inclusion prior p controlling the trade-off between prediction quality and explaining the count data.

Pang and Lee Movie Reviews		
Filtering Method	Coherence	RMSE
$p = 10$, pf-sLDA	1.916 (0.105)	1.418 (0.024)
$p = 10$, correlation	1.115 (0.092)	1.729 (0.031)
$p = 15$, pf-sLDA	1.509 (0.077)	1.313 (0.008)
$p = 15$, correlation	0.875 (0.088)	1.694 (0.017)




Figure 3: *Left*: Mean and (SD) across 5 runs for topic coherence (higher is better) and RMSE (lower is better) for sLDA on a filtered vocabulary on a held-out test set. Filtering by pf-sLDA as compared to correlation results better coherence and RMSE. *Right*: Each dot represents one word in the vocabulary. We plot the variational feature selector φ after training vs. the correlation to the target for each word

recall. Our real world examples demonstrate that pf-sLDA achieves better topic coherency with similar prediction performance to state-of-the-art ap-

proaches in supervised topic modeling. To test how well pf-sLDA removes irrelevant features, we now turn to a simulated setting in which words are gen-

Pang and Lee Movie Reviews			
	sLDA	pc-sLDA	pf-sLDA
High	motion, way, love, performance, best , picture, films, character, characters, life	best , little, time, good , don, picture, year, rated, films just	brilliant , rare, perfectly , true, oscar , documentary, wonderful , fascinating, perfect , best
Low	plot, time, bad , funny , good , humor , little, isn, action	script, year, little, good , don, look, rated, picture, just, films	awful , stupid , gags, dumb , dull , sequel, flat , worse , ridiculous , bad
Yelp Reviews			
	sLDA	pc-sLDA	pf-sLDA
High	fries, fresh , burger, try, cheese, really, pizza, place, like , good	best , just, amazing , love , good , food, service, place, time, great	fantastic , loved, highly, fun , excellent , awesome , atmosphere, amazing , delicious , great
Low	store, time, want, going, place, know, people, don, like , just	didn, don, said, told, like, place, time, just, service, food	awful , management, dirty , poor , horrible , worst , rude , terrible , money, bad
ASD			
	sLDA	pc-sLDA	pf-sLDA
High	Intellect disability Esophageal reflux Hearing loss Development delay Downs syndrome	Infantile cerebral palsy Congenital quadriplegia Esophageal reflux fascia Muscle/ligament dis Feeding problem	Other convulsions Aphasia Convulsions Central hearing loss Grand mal status
Low	Otitis media Asthma Downs syndrome Scoliosis Constipation	Accommodative esotropia Joint pain-ankle Congenital factor VIII Fragile X syndrome Pain in limb	Autistic disorder Diabetes Type 1 c0375114 Other symbolic dysfunc Diabetes Type 1 c0375116 Diabetes Type 2

Table 2: We list the most probable words in the topics with the highest and lowest regression coefficient for each model and data set. In the context of each data set, words expected to be in a high-regression-coefficient topic are listed in green, and words expected to be in a low-regression-coefficient topic are listed in red. These colors are done by hand and are only meant for ease of evaluation. It is clear that the topics learned by pf-sLDA are the most coherent and contain the most words with task relevance.

Synthetic Data Set, true $p = 0.25$		
p	Precision	Recall
0.15	1.000 (0.000)	0.430 (0.013)
0.25	0.981 (0.006)	0.962 (0.006)
0.35	0.885 (0.008)	0.982 (0.006)

Table 3: Mean and (SD) of precision and recall of relevant features being considered relevant by pf-sLDA with word inclusion prior p on a synthetic data set. In the well specified case, pf-sLDA is able to recover the relevant features with high precision and recall.

erated from the pf-sLDA generative process with $p = 0.25$. (See details on simulation setting in Appendix 9.7.)

We run pf-sLDA with random initializations 10 times on each of the simulated data sets. We report mean and standard deviation of precision and recall of relevant features being labeled as relevant by pf-sLDA. We consider pf-sLDA labelling a feature as relevant if the corresponding φ is greater than 0.99.

Under the well specified case when we have the correct word inclusion prior p , we see that pf-sLDA is able to correctly identify which features are relevant with high precision and recall (see Table 3). We get results as expected if we misspecify p . Precision increases and recall decreases with lower p , since this encourages the model to believe fewer of the words are actually relevant. Precision decreases and recall increases with higher p , since this encourages the model to believe more of the words are actually relevant.

evant. Furthermore, the model is able to reproduce high precision and recall consistently across random initialization, as seen by the low standard deviations.

The variable selection of pf-sLDA outperforms naive variable selection based on correlation to target. Another question of interest was whether using pf-sLDA as a feature selector is any better than simply retaining words that correlate to the target. To test this, we consider two different filtering schemes. The first scheme retains words that under pf-sLDA have their variational feature selector $\varphi_v > 0.99$. The second scheme retains the top- N highest correlated words to the target, where N is chosen to be the same number as the words retained by pf-sLDA. We then run a standard sLDA on the filtered vocabulary. We run this experiment on Pang and Lee’s movie review data set.

From Figure 3, we see that running sLDA on a filtered vocab based on pf-sLDA produces better results in terms of both topic coherence and RMSE than on a filtered vocab based on correlation. Additionally, it is clear pf-sLDA throws out words that have close to zero correlation with the target. For words that have some correlation, pf-sLDA identifies and retains those that are more useful for prediction.

7 Discussion

The introduction of pf-sLDA was motivated by the observation that the presence of irrelevant features in real world data sets hampers the ability of topic models to recover topics that are simultaneously coherent and relevant for a supervised task.

Learning and Inference. From a learning perspective, one of the advantages of our approach compared to prediction constrained training is that pf-sLDA can be specified via a graphical model. Thus, we can enjoy the benefits of a principled loss and graphical model inference. For example, while we used SGD to optimize the ELBO, we could have easily incorporated classic updates for some of the parameters as in McAuliffe and Blei [2008]. The most challenging aspect was maintaining orthogonality between the relevant topics β and the additional topic π . In all the settings we tested, our parameterization of $q(\xi|\varphi)$ that removes dependence on document enforced $\beta^\top \pi = \mathbf{0}$. Proving properties of this choice and the approach of incorporating model constraints into the variational family and inference is an area for future work.

Interpretable Predictions. Our experiments show that the topics learned by pf-sLDA are more

coherent than other methods, using a metric that aligns closely with human judgement of interpretability (Newman et al. [2010]). Moreover, qualitative examination of the topics and selected words for a variety of supervised tasks supports this conclusion: topics learned by pf-sLDA are highly coherent while maintaining task relevance. For example, in a deeper qualitative analysis of the Pang and Lee movie review data set, where targets are integers from 1 to 10 (10 being most positive and 1 being most negative), pf-sLDA produces topics with corresponding regression coefficients close to 10 that are highly positive in sentiment, and topics with coefficients close to 0 that are highly negative in sentiment, while topics with coefficients close to 5 and 6 are more mild in sentiment (see Table 4 in Appendix 9.8). This contrasts with the pc-sLDA and sLDA topics, in which topic sentiments are not particularly clear.

Feature Selection and Applications. The pf-sLDA approach is particularly well suited for tasks where it is suspected that there are a large number of features unrelated to a target variable of interest. For example, beyond sentiment and epilepsy detection, one could imagine pf-sLDA being useful for tasks such as genome wide association studies where a majority of genetic data is irrelevant to a certain disease risk but there exists a latent structure describing a small set of relevant genes.

In Figure 3, we demonstrated the effectiveness of the feature selection of pf-sLDA by showing retaining features based on the variational feature selectors (φ) of pf-sLDA outperforms simply retaining features with strong bivariate correlation with the target. Intuitively, pf-sLDA preserves co-occurrence relationships between the features that may be useful for a wide range of possible downstream tasks, which selection based solely on the bivariate correlations may ignore.

8 Conclusion

In this paper, we introduced prediction-focused supervised LDA, whose feature selection procedure improves both predictive accuracy and topic coherence of supervised topic models. The model enjoys good theoretical properties, inferential properties, and performed well on real and synthetic examples. Future work could include establishing additional theoretical properties of the pf-sLDA variable selection procedure and applying our trick of managing trade-offs within a graphical model for variable selection in other generative models.

Acknowledgements

FDV and JR acknowledge support from NSF CAREER 1750358 and Oracle Labs. JR also acknowledges support from Harvard’s PRISE program over Summer 2019.

References

- David M. Blei, Andrew Y. Ng, and Michael I. Jordan. Latent dirichlet allocation. *J. Mach. Learn. Res.*, 3:993–1022, 2003.
- Jianshu Chen, Ji He, Yelong Shen, Lin Xiao, Xiaodong He, Jianfeng Gao, Xinying Song, and Li Deng. End-to-end learning of lda by mirror-descent back propagation over a deep architecture. In *Advances in Neural Information Processing Systems*, pages 1765–1773, 2015.
- Angela Fan, Finale Doshi-Velez, and Luke Miratrix. Prior matters: simple and general methods for evaluating and improving topic quality in topic modeling. 2017.
- Yoni Halpern, Steven Horng, Larry A Nathanson, Nathan I Shapiro, and David Sontag. A comparison of dimensionality reduction techniques for unstructured clinical text. In *Icml 2012 workshop on clinical data analysis*, volume 6, 2012.
- Michael C Hughes, Gabriel Hope, Leah Weiner, Thomas H McCoy, Roy H Perlis, Erik B Sudderth, and Finale Doshi-Velez. Prediction-constrained topic models for antidepressant recommendation. *arXiv preprint arXiv:1712.00499*, 2017a.
- Michael C Hughes, Leah Weiner, Gabriel Hope, Thomas H McCoy Jr, Roy H Perlis, Erik B Sudderth, and Finale Doshi-Velez. Prediction-constrained training for semi-supervised mixture and topic models. *arXiv preprint arXiv:1707.07341*, 2017b.
- Dongwoo Kim, Yeonseung Chung, and Alice H. Oh. Variable selection for latent dirichlet allocation. *ArXiv*, abs/1205.1053, 2012.
- Da Kuang, P. Jeffrey Brantingham, and Andrea L. Bertozzi. Crime topic modeling. *Crime Science*, 6:1–20, 2017.
- Tianyi Lin, Wentao Tian, Qiaozhu Mei, and Hong Cheng. The dual-sparse topic model: mining focused topics and focused terms in short text. In *Proceedings of the 23rd international conference on World wide web*, pages 539–550. ACM, 2014.
- M. Arjumand Masood and Finale Doshi-Velez. A particle-based variational approach to bayesian non-negative matrix factorization. *J. Mach. Learn. Res.*, 20:90:1–90:56, 2018.
- Jon D McAuliffe and David M Blei. Supervised topic models. In *Advances in neural information processing systems*, pages 121–128, 2008.
- David Newman, Jey Han Lau, Karl Grieser, and Timothy Baldwin. Automatic evaluation of topic coherence. In *HLT-NAACL*, 2010.
- Bo Pang and Lillian Lee. Seeing stars: Exploiting class relationships for sentiment categorization with respect to rating scales. In *ACL*, 2005.
- Shuai Wang, Zhiyuan Chen, Geli Fei, Bing Liu, and Sherry Emery. Targeted topic modeling for focused analysis. In *KDD*, 2016.
- Yelp. Yelp dataset challenge, 2019. data retrieved from <https://www.yelp.com/dataset/challenge>.
- Cheng Zhang and Hedvig Kjellström. How to supervise topic models. In *European Conference on Computer Vision*, pages 500–515. Springer, 2014.
- Jun Zhu, Amr Ahmed, and Eric P Xing. Medlda: maximum margin supervised topic models. *Journal of Machine Learning Research*, 13(Aug):2237–2278, 2012.