

Neural Topic Modeling with Cycle-Consistent Adversarial Training

Xuemeng Hu* Rui Wang* Deyu Zhou† Yuxuan Xiong

School of Computer Science and Engineering, Key Laboratory of Computer Network and Information Integration, Ministry of Education, Southeast University, China
{xuemenghu, rui_wang, d.zhou, yuxuanxiong}@seu.edu.cn

Abstract

Advances on deep generative models have attracted significant research interest in neural topic modeling. The recently proposed Adversarial-neural Topic Model models topics with an adversarially trained generator network and employs Dirichlet prior to capture the semantic patterns in latent topics. It is effective in discovering coherent topics but unable to infer topic distributions for given documents or utilize available document labels. To overcome such limitations, we propose Topic Modeling with Cycle-consistent Adversarial Training (ToMCAT) and its supervised version sToMCAT. ToMCAT employs a generator network to interpret topics and an encoder network to infer document topics. Adversarial training and cycle-consistent constraints are used to encourage the generator and the encoder to produce realistic samples that coordinate with each other. sToMCAT extends ToMCAT by incorporating document labels into the topic modeling process to help discover more coherent topics. The effectiveness of the proposed models is evaluated on unsupervised/supervised topic modeling and text classification. The experimental results show that our models can produce both coherent and informative topics, outperforming a number of competitive baselines.

1 Introduction

Topic models, such as Latent Dirichlet Allocation (LDA) (Blei et al., 2003), aim to discover underlying topics and semantic structures from text collections. Due to its interpretability and effectiveness, LDA has been extended to many Natural Language Processing (NLP) tasks (Lin and He, 2009; McAuley and Leskovec, 2013; Zhou et al., 2017). Most of these models employ mean-field

variational inference or collapsed Gibbs sampling (Griffiths and Steyvers, 2004) for model inference as a result of their intractable posteriors. However, such inference algorithms are model specific and require dedicated derivations.

To address such limitation, neural topic models with black-box inference have been explored, with more flexible training schemes. Inspired by variational autoencoder (VAE) (Kingma and Welling, 2013), Miao et al. (2016) proposed Neural Variational Document Model which interprets the latent code in VAE as topics. Following this way, Srivastava and Sutton (2017) adopted the logistic normal prior rather than Gaussian to mimic the simplex properties of topic distribution. Logistic normal is a Laplace approximation to the Dirichlet distribution (MacKay, 1998). However, logistic normal can not exhibit multiple peaks at the vertices of the simplex as the Dirichlet distribution. Therefore, it is less capable of capturing the multi-modality which is crucial for topic modeling (Wallach et al., 2009).

To overcome such limitation, Wang et al. (2019a) proposed Adversarial-neural Topic Model (ATM), a topic model based on Generative Adversarial Networks (GANs) (Goodfellow et al., 2014) and sampling topics directly from the Dirichlet distribution to impose a Dirichlet prior. ATM employs a generator transforming randomly sampled topic distributions to word distributions, and an adversarially trained discriminator estimating the probability that a word distribution came from the training data rather than the generator. Although ATM was shown to be effective in discovering coherent topics, it can not be used to induce the topic distribution given a document due to the absence of a topic inference module. Such limitation hinders its application to downstream tasks, such as text classification. Moreover, ATM fails to deal with document labels which can help extract more co-

*Equal contribution.

†Corresponding author.

herent topics. For example, a document labeled as ‘*sports*’ more likely belongs to topics such as ‘*basketball*’ or ‘*football*’ rather than ‘*economics*’ or ‘*politics*’.

To address such limitations of ATM, we propose a novel neural topic modeling approach, named Topic Modeling with Cycle-consistent Adversarial Training (ToMCAT). In ToMCAT, topic modeling is cast into the transformation between topic distributions and word distributions. Specifically, the transformation from topic distributions to word distributions is used to interpret topics, and the reverse transformation is used to infer underlying topics for a given document. Under such formulation, ToMCAT employs a generator to transform topic distributions randomly sampled from the Dirichlet prior into the corresponding word distributions, and an encoder to reversely transform documents represented as word distributions into their topic distributions. To encourage the generator/encoder to produce more realistic target samples, discriminators for word/topic distributions are introduced to enable adversarial training. Additional cycle-consistency constraints are utilized to align the learning of the encoder and the generator to prevent them from contradicting each other. Furthermore, for documents with labels, we propose sToMCAT that introduces an extra classifier to regularize the topic modeling process.

The main contributions of the paper are:

- ToMCAT, a novel topic model with cycle-consistent adversarial training is proposed. To the best of our knowledge, it is the first adversarial topic modeling approach with both topic discovery and topic inference.
- sToMCAT, a supervised extension to ToMCAT, is proposed to help discover more coherent topics with available document labels.
- Experimental results on unsupervised/supervised topic modeling and text classification demonstrate the effectiveness of the proposed approaches.

2 Related Work

Our work is related to neural topic modeling and unsupervised style transfer.

2.1 Neural Topic Modeling

Recent advances on deep generative models, such as VAEs (Kingma and Welling, 2013) and GANs

(Goodfellow et al., 2014), attract much research interest in the NLP community.

Based on VAE, Neural Variational Document Model (NVDM) (Miao et al., 2016) encodes documents with variational posteriors in the latent topic space. NVDM employs Gaussian as the prior distribution of latent topics. Instead, Srivastava and Sutton (2017) proposed that Dirichlet distribution is a more appropriate prior for multinomial topic distributions, and constructed a Laplace approximation of Dirichlet to enable reparameterisation (Kingma and Welling, 2013). Furthermore, the word-level mixture is replaced with a weighted product of experts (Srivastava and Sutton, 2017). Later, a non-parametric neural topic model utilizing stick-breaking construction was presented in (Miao et al., 2017). There are some attempts in incorporating supervised information into neural topic modeling. For example, Card et al. (2018) extended the Sparse Additive Generative Model (Eisenstein et al., 2011) in the neural framework and incorporated document metadata such as document labels into the modeling process.

Apart from VAE-based approaches, Adversarial-neural Topic Model (ATM) (Wang et al., 2019a)) was proposed to model topics with GANs. The generator of ATM projects randomly sampled topic distributions to word distributions, and is adversarially trained with a discriminator that tries to distinguish real and generated word distributions. Moreover, Wang et al. (2019b) extended ATM for open-domain event extraction by representing an event as a combination of an entity distribution, a location distribution, a keyword distribution and a date distribution. Such joint distributions are adversarially learned in a similar manner as ATM. The proposed ToMCAT is partly inspired by ATM but differs in its capability of inferring document-specific topic distributions and incorporating supervision information. BAT (Wang et al., 2020) is an extension to ATM that employs bidirectional adversarial training (Donahue et al., 2016) for document-specific topic distribution inference. Although BAT similarly utilizes an adversarial training objective to guide the learning of topic distribution, there are some major differences. Apart from different adversarial losses, ToMCAT also incorporates two cycle-consistency constraints which encourage the model to generate informative representations and are shown to be crucial for generating coherent topics as in our experiments.

2.2 Unsupervised Style Transfer

Style transfer, aiming at transforming representations from one style to another, has been found many interesting applications, such as image and text style transfer. However, training data paired between different styles are not available for many tasks. To solve this problem, [Zhu et al. \(2017\)](#) imposed cycle-consistency constraints to align mappings between two styles and proposed CycleGAN for unsupervised image style translation. Similarly, DiscoGAN ([Kim et al., 2017](#)) was proposed to discover the relations between different image styles and transformed images from one style to another without paired data. In the NLP field, [Lee et al. \(2018\)](#) developed a CycleGAN-based approach to transfer the sentiment style (positive, negative) of the text.

Inspired by CycleGAN, Our work views topic modeling as unsupervised distribution transfer and follows the framework of CycleGAN.

3 Methodology

Given a corpus \mathcal{D} consisting of N documents $\{\mathbf{x}_i\}_{i=1}^N$, two main purposes of topic modeling are:

1. Topic discovery. Given a one-hot topic indicating vector $\mathbf{I}_k \in \mathbb{R}^K$ where K is the number of topics and $I_{kk} = 1$, discover the corresponding word distribution $\mathbf{t}_k \in \mathbb{R}^V$ from \mathcal{D} where V is the vocabulary size. More generally, we can consider topic discovery as finding a mapping from topic distribution to word distribution.
2. Topic inference. Infer the topic distribution $\mathbf{z}_j \in \mathbb{R}^K$ of the document $\mathbf{x}_j \in \mathbb{R}^V$. Similarly, the topic inference can be considered as finding a mapping from word distribution to topic distribution.

We now formalize the above observations. Let X be the word distribution set and Z the topic distribution set. Given training samples $\{\mathbf{x}_i\}_{i=1}^N$ where $\mathbf{x}_i \in X$ and document-specific topic distributions $\{\mathbf{z}_j\}_{j=1}^M$ where $\mathbf{z}_j \in Z$, the goal of topic modeling is to learn a mapping function G , called *generator*, to transform samples in Z into X and a reverse function E , called *encoder*, to transform samples in X into Z . However, it should be noted that training samples in X and Z are unpaired since the topic distribution of a document is unknown before topic modeling. Thus, the problem is how

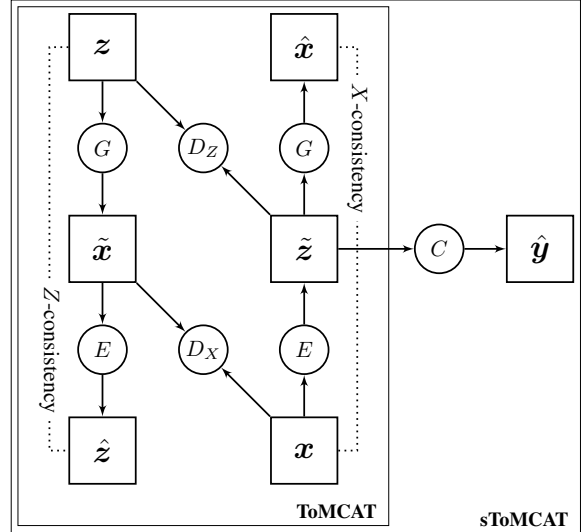


Figure 1: The framework of ToMCAT and sToMCAT. Circles are neural networks, squares are data representations, and arrows indicate the forward pass directions.

to learn G and E to model topics in the absence of paired samples between X and Z .

3.1 ToMCAT

We now introduce the proposed ToMCAT, which is shown in the inner panel of Figure 1.

ToMCAT consists of a generator $G: Z \rightarrow X$, an encoder $E: X \rightarrow Z$, and adversarial discriminators D_X and D_Z of G and E respectively. Following CycleGAN ([Zhu et al., 2017](#)), ToMCAT employs two types of losses, namely adversarial losses and cycle-consistency losses, to guide the training of the encoder E and the generator G . The details of these modules are described below.

3.1.1 Encoder Network E

Encoder E transforms a word distribution $\mathbf{x}_i \in \mathbb{R}^V$ into its corresponding topic distribution $\mathbf{z}_i \in \mathbb{R}^K$. Following ([Wang et al., 2019a](#)), we represent $\mathbf{x}_i \in X$ with the normalized TF-IDF (Term Frequency–Inverse Document Frequency) representation of i -th document:

$$\hat{d}_{ij} = \frac{d_{ij}}{\sum_j d_{ij}} \cdot \log \frac{N}{1 + \sum_{n=1}^N \mathbb{1}(d_{nj} \neq 0)}, \quad (1)$$

$$x_{ij} = \frac{\hat{d}_{ij}}{\sum_v \hat{d}_{iv}}, \quad (2)$$

where d_{ij} is the count of j -th word in i -th document, $\mathbb{1}(\cdot)$ denotes the indicator function. Equation 1 calculates the smoothed TF-IDF of \mathbf{d}_i , which is then normalized to sum to one in Equation 2. We use TF-IDF as the document representation

because TF-IDF generally preserves the relative importance of words in a document and reduce the noise of stop words. As the target distribution of the generator G , such property of TF-IDF will help generate more informative topics.

The implementation of the encoder is a multilayer perception (MLP) with LeakyReLU activation (Maas et al., 2013) and batch normalization (BN) (Ioffe and Szegedy, 2015). The detailed transformations are: [Linear(V, H) \rightarrow LeakyReLU(0.1) \rightarrow BN \rightarrow Linear(H, K) \rightarrow Softmax], where Linear(I, J) denotes a linear transformation from I -dim to J -dim, H is the number of hidden units, and the final Softmax makes sure that the final output is one-normalized to match the input of the generator G . Inputs of E are either sampled from the corpus \mathcal{D} or generated by G .

3.1.2 Generator Network G

The generator G performs the reverse operation of the encoder by transforming a topic distribution $z_j \in \mathbb{R}^K$ into a word distribution $x_j \in \mathbb{R}^V$, where the input z_j is generated by the encoder or sampled from the prior distribution. To draw the topic distribution z_j , a common practice for topic modeling is to use the Dirichlet distribution, the conjugate prior of the multinomial distribution. We also stick with this choice in our model. Specifically, we draw topic distributions from a symmetric Dirichlet distribution with parameters $\mathbf{A} \in \mathbb{R}^K$ where $A_k = \alpha$ for $1 \leq k \leq K$.

After sampling a topic distribution z_j from the Dirichlet prior, the generator then maps z_j from Z to X , and the transformations is similar to the encoder: [Linear(K, H) \rightarrow LeakyReLU(0.1) \rightarrow BN \rightarrow Linear(H, V) \rightarrow Softmax], where the final output is also normalized by the Softmax to match the input of the encoder.

3.1.3 Training Objective

Following CycleGAN (Zhu et al., 2017), we employ adversarial losses and cycle-consistency losses to guide the training of G and E . The adversarial losses encourage G and E to generate samples matching the data distribution in the target space (X for G and Z for E) while the cycle-consistency losses align G and E in these two distribution spaces to prevent them from contradicting each other.

Adversarial Loss Generator G is adversarially trained with a discriminator D_X , which takes as

input either real samples from training data, i.e., $x \sim p_{\text{data}}(x)$, or fake samples generated by G , i.e., $G(z)$. The goal of D_X is to distinguish real samples from fake ones, while G instead aims to fool D_X by generating samples similar to x . Therefore, the adversarial training encourages G to mimic the pattern of X and produce realistic word distributions. We employ a Wasserstein GAN (WGAN) (Arjovsky et al., 2017) based adversarial loss to G and D_X :

$$\mathcal{L}_{\text{adv}}(G, D_X) = \mathbb{E}_{x \sim p_{\text{data}}(x)}[D_X(x)] - \mathbb{E}_{z \sim p_{\text{data}}(z)}[D_X(G(z))], \quad (3)$$

where D tries to maximize $\mathcal{L}_{\text{adv}}(G, D_X)$ while G tries to minimize it. Similarly, the adversarial loss applied to E and D_Z is:

$$\mathcal{L}_{\text{adv}}(E, D_Z) = \mathbb{E}_{z \sim p_{\text{data}}(z)}[D_Z(z)] - \mathbb{E}_{x \sim p_{\text{data}}(x)}[D_Z(E(x))]. \quad (4)$$

Discriminators D_X and D_Z are implemented with MLPs, and we use the same architecture for them: [Linear(S, H) \rightarrow LeakyReLU(0.1) \rightarrow BN \rightarrow Linear($H, 1$)], where S equals to V for D_X and K for D_Z . Since we are using WGAN rather than the original GAN loss as in CycleGAN, we do not apply a sigmoid transformation to discriminator outputs.

Cycle-Consistency Loss Adversarial training might lead to generating samples identically distributed as corresponding target samples (Goodfellow et al., 2014). However, the relationship between the source distributions and the transformed distributions is unconstrained. Zhu et al. (2017) argued that adversarial losses alone is not able to fulfill this task and that the learned mappings should be cycle-consistent to reduce the search space of possible mapping functions, i.e., $x \rightarrow E(x) \rightarrow G(E(x)) \approx x$ and $z \rightarrow G(z) \rightarrow E(G(z)) \approx z$. To this end, two cycle-consistency losses $\overrightarrow{\mathcal{L}}_{\text{cyc}}(G, E)$ and $\overleftarrow{\mathcal{L}}_{\text{cyc}}(G, E)$ are added to the training objective, as shown in the inner panel (dotted lines) of Figure 1. Specifically,

$$\begin{aligned} \overrightarrow{\mathcal{L}}_{\text{cyc}}(G, E) &= \mathbb{E}_{x \sim p_{\text{data}}(x)}[\|G(E(x)) - x\|_1], \\ \overleftarrow{\mathcal{L}}_{\text{cyc}}(G, E) &= \mathbb{E}_{z \sim p_{\text{data}}(z)}[\|E(G(z)) - z\|_1], \end{aligned} \quad (5)$$

where $\|\cdot\|_1$ denotes L1 norm.

Overall Objective Summing up adversarial losses in Equation 3, 4 and cycle-consistency losses in Equation 5, the overall objective of ToMCAT is:

$$\mathcal{L}(G, E, D_X, D_Z) = \mathcal{L}_{\text{adv}}(G, D_X) + \mathcal{L}_{\text{adv}}(E, D_Z) + \lambda_1 \overrightarrow{\mathcal{L}}_{\text{cyc}}(G, E) + \lambda_2 \overleftarrow{\mathcal{L}}_{\text{cyc}}(G, E), \quad (6)$$

where λ_1 and λ_2 respectively control the relative importance of $\overrightarrow{\mathcal{L}}_{\text{cyc}}(G, E)$ and $\overleftarrow{\mathcal{L}}_{\text{cyc}}(G, E)$ w.r.t. adversarial losses.

3.2 sToMCAT

The encoder E transforms the word distribution \mathbf{x} to corresponding topic distribution \mathbf{z} , which effectively captures the key semantic information of \mathbf{x} and can be directly used to downstream tasks, e.g., text classification. Therefore, for labeled documents we extend ToMCAT with a classifier C to allow the incorporation of label information, as shown in Figure 1. We name the supervised version as sToMCAT.

For a word distribution \mathbf{x} and its one-hot label \mathbf{y} , \mathbf{x} is first encoded by the encoder E into the topic distribution \mathbf{z} , and then \mathbf{z} is fed to the classifier C to predict the probability of \mathbf{y} . The predictive objective is defined as:

$$\mathcal{L}_{\text{cls}}(E, C) = -\mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim p_{\text{data}}(\mathbf{x}, \mathbf{y})} [\mathbf{y} \log C(E(\mathbf{x}))], \quad (7)$$

where L is the dimension of \mathbf{y} . We employ an MLP classifier: [Linear(K, H) \rightarrow LeakyReLU(0.1) \rightarrow BN \rightarrow Linear(H, L) \rightarrow Softmax].

For sToMCAT, the topic model and the classifier are trained jointly, and its overall objective is defined as:

$$\mathcal{L}_{\text{sup}}(G, E, D_X, D_Z, C) = \mathcal{L}(G, E, D_X, D_Z) + \lambda_3 \mathcal{L}_{\text{cls}}(E, C). \quad (8)$$

3.3 Training Details

The proposed ToMCAT and sToMCAT are trained with the Adam optimizer (Kingma and Ba, 2014), whose learning rate and momentum term β_1 are set to 0.0001 and 0.5 respectively for (G, E, D_X) and D_Z , while 0.001 and 0.9 for the classifier C . The hidden unit numbers are set to 100 for all modules. Besides, to enforce the Lipschitz constraints required by WGAN, a weight clipping of 0.01 is adopted (Arjovsky et al., 2017).¹

¹We also experiment with the gradient-penalty WGAN (Gulrajani et al., 2017), but the weight clipping version performs better in general.

During training, the parameters of discriminators D_X, D_Z and mappings G, E are alternately updated. Specifically, at each training iteration, firstly we optimize D_X and D_Z for 5 steps with adversarial losses, and then another training step is taken to optimize G and E with adversarial losses and cycle-consistency losses (Equation 6). When the model is trained in a supervised way, the predictive objective is additionally applied to E and C at the last training step (Equation 8).

We found that relatively good choices of λ_1 and λ_2 fall into different regions for different datasets and topic number settings, which implies a further tuning of these hyperparameters is needed. To ease this kind of burden, we apply a gradient-based mechanism to adversarial losses and cycle-consistency losses. It balances these two types of losses with the L2 norms of their gradients w.r.t. the output of their preceding mapping functions. E.g., for $\mathcal{L}_{\text{adv}}(G, D_X)$ and $\overrightarrow{\mathcal{L}}_{\text{cyc}}(G, E)$, we replace λ_1 in Equation 6 with:

$$\lambda_1 = \hat{\lambda}_1 \frac{\|\partial \mathcal{L}_{\text{adv}}(G, D_X) / \partial G(\mathbf{z})\|_2}{\|\partial \overrightarrow{\mathcal{L}}_{\text{cyc}}(G, E) / \partial G(\mathbf{z})\|_2}, \quad (9)$$

where $\hat{\lambda}_1$ is the new balancing factor and $\|\cdot\|_2$ denotes L2 norm. Similarly, $\mathcal{L}_{\text{adv}}(E, D_Z)$ and $\overleftarrow{\mathcal{L}}_{\text{cyc}}(G, E)$, $\mathcal{L}_{\text{adv}}(E, D_Z)$ and $\mathcal{L}_{\text{cls}}(E, C)$ are also balanced in this way with $\hat{\lambda}_2$ and $\hat{\lambda}_3$. The resulting $\hat{\lambda}_1, \hat{\lambda}_2$ and $\hat{\lambda}_3$ are set to 2, 0.2 and 1 respectively for all datasets and topic number settings in our experiments, thus avoiding the time-consuming hyperparameter tuning process.

4 Experiments

In this section, we first describe datasets and compared baselines. Then we present topic modeling results under both unsupervised and supervised settings. Finally, we report the text classification results.

Dataset	#Train	#Test	Vocab Size	#Class
NYT	99,992	-	12,604	-
GRL	29,762	-	15,276	-
DBP	99,991	69,993	9,005	14
20NG	11,258	7,492	2,000	20

Table 1: Dataset statistics.

Dataset	Metric	Unsupervised							Supervised		
		NVDM	ProdLDA	ATM	BAT	LDA	Scholar	ToMCAT	sLDA	Scholar	sToMCAT
NYT	C_A	0.0770	0.1841	0.2292	0.2356	0.2145	0.1949	0.2444	–	–	–
	C_P	–0.5368	0.1255	0.3330	0.3749	0.3230	0.0451	0.3879	–	–	–
	NPMI	–0.1461	0.0155	0.0806	0.0952	0.0814	–0.0290	0.0956	–	–	–
GRL	C_A	0.0715	0.1483	0.2203	0.2108	0.1960	0.2064	0.2285	–	–	–
	C_P	–0.5188	–0.0651	0.2576	0.2312	0.1974	0.2150	0.2752	–	–	–
	NPMI	–0.1225	–0.0193	0.0655	0.0608	0.0533	0.0592	0.0808	–	–	–
DBP	C_A	0.1385	0.2653	0.2928	0.2355	0.2756	0.3010	0.3410	0.2216	0.2966	0.3568
	C_P	–0.2970	0.2149	0.3397	0.3749	0.3516	0.2369	0.4327	0.2581	0.1834	0.4981
	NPMI	–0.1171	0.0212	0.1100	0.0951	0.1033	0.0661	0.1434	0.0685	0.0526	0.1661
20NG	C_A	0.1115	0.1776	0.1833	0.1991	0.1862	0.1777	0.2082	0.1771	0.1811	0.2248
	C_P	–0.0632	0.0709	0.2572	0.2962	0.2816	0.2120	0.3137	0.2621	0.2443	0.3563
	NPMI	–0.0495	–0.0439	0.0379	0.0555	0.0637	0.0426	0.0656	0.0554	0.0486	0.0709

Table 2: Average topic coherence of 5 topic number settings (20, 30, 50, 75, 100) on 4 datasets. Bold values indicate the best performing models for each dataset/metric/supervision setting.

4.1 Experimental Setup

We evaluate the performance of proposed models on four datasets: NYTimes² (NYT), Grolier³ (GRL), DBpedia ontology classification dataset (DBP) (Zhang et al., 2015) and 20 Newsgroups⁴ (20NG). For NYTimes and Grolier datasets, we use the processed version of (Wang et al., 2019a). For the DBpedia dataset, we first sample 100,000 documents from the whole training set, and then perform preprocessing including tokenization, lemmatization, removal of stopwords, and low-frequency words. The same preprocessing is also applied to the 20 Newsgroups dataset. The statistics of the processed datasets are shown in Table 1.

We choose the following approaches as our baselines:

- LDA (Blei et al., 2003). We use GibbsLDA++, an implementation using Gibbs sampling for parameter estimation and inference⁵.
- sLDA⁶ (Mcauliffe and Blei, 2008), a supervised extension to LDA.
- NVDM⁷ (Miao et al., 2016), a VAE-based model that employs Gaussian prior for the latent topics.
- ProdLDA⁸ (Srivastava and Sutton, 2017), a

²<http://archive.ics.uci.edu/ml/datasets/Bag+of+Words>

³<https://cs.nyu.edu/~roweis/data>

⁴<http://qwone.com/~jason/20Newsgroups>

⁵<http://gibbslda.sourceforge.net>

⁶<https://github.com/blei-lab/class-slda>

⁷<https://github.com/ysmiao/nvdm>

⁸https://github.com/akashgit/autoencoding_vi_for_topic_models

VAE-based model that replaces the mixture model with a product of experts.

- Scholar⁹ (Card et al., 2018), a ProdLDA-based model that enables optional incorporation of metadata.
- ATM (Wang et al., 2019a), a neural topic model utilizing adversarial training.
- BAT (Wang et al., 2020), a neural topic model utilizing bidirectional adversarial training.

4.2 Topic Modeling

We evaluate the performance of the proposed models and baselines using topic coherence measures. Topic coherence measures are metrics for quantifying the understandability of the extracted topics, which are shown highly correlated with human subjects (Newman et al., 2010; Aletras and Stevenson, 2013). Since a topic is typically represented as a word distribution over the vocabulary or n top-weighted words (i.e., topic words) in this distribution, we calculate the coherence of a topic by measuring the relatedness between its topic words. The word relatedness scores are estimated based on some kind of word co-occurrence statistics on Wikipedia, for example, by applying a sliding window over the Wikipedia corpus and collecting word co-occurrences to calculate NPMI (Normalized Pointwise Mutual Information) (Bouma, 2009) for word pairs. We refer readers to (Röder et al., 2015) for detailed calculation and comparison of different topic coherence measures. In our experiments, we use top-10 topic words of each topic to calculate topic coherence and report the results of 3 topic

⁹<https://github.com/dallascard/scholar>

coherence measures: C_A (Aletras and Stevenson, 2013), C_P (Röder et al., 2015), and NPMI (Aletras and Stevenson, 2013). The topic coherence scores are calculated using Palmetto¹⁰.

4.2.1 Unsupervised Topic Modeling

To make a more comprehensive comparison of our model with baselines for topic modeling, we experiment on each dataset with five topic number settings: 20, 30, 50, 75, 100. The average topic coherence scores of 5 settings are presented in Table 2. We can see from the left part of Table 2 that, among all unsupervised topic models, our model achieves the highest scores on all datasets and topic coherence measures.

With an improper Gaussian prior, NVDM shows the worst performance among all neural topic models with no exception. The logistic-normal based ProLDA and Scholar achieve higher topic coherence scores compared to NVDM, but are still largely underperformed compared to our model. BAT achieves the second-best place most of the time in unsupervised topic modeling experiments. Compared to ToMCAT, BAT has a similar adversarial objective but lacks the cycle-consistency constraints. Therefore, the generator and encoder of BAT only aim to fool the discriminator by mimicking the pattern of the joint distribution of real documents and topics. With the incorporation of two cycle-consistency losses, ToMCAT is explicitly encouraged to generate not only realistic but also informative representations in order to reduce the cycle-consistency losses.

To give an insight into the generated topics, 8 out of 50 topics discovered by ToMCAT on NYTimes are presented in Table 3, where a topic is represented by the ten words with the highest probability in the topic. We can observe that the extracted topics are highly coherent and interpretable. The corresponding full list of topics can be found in the appendix.

4.2.2 Supervised Topic Modeling

Supervised topic modeling aims to leverage available document labels to benefit topic modeling. Therefore we only conduct experiments on labeled datasets, i.e., DBpedia and 20 Newsgroups. The experimental results are shown on the right part of Table 2. We expect the topic extraction results would be improved with the incorporation of topic labels. However, this is not always the case as shown in

Vehicle	Election	Court	Fashion
car	voter	court	fashion
tires	poll	lawsuit	designer
fuel	campaign	case	leather
driver	percent	ruling	wear
truck	primary	antitrust	dress
vehicle	republican	suit	clothes
vehicles	vote	plaintiff	skirt
gas	democratic	judge	white
gasoline	states	settlement	shirt
engine	democrat	federal	pant
Cooking	Baseball	Disease	Art
cup	inning	patient	artist
tablespoon	run	cancer	painting
pepper	hit	doctor	art
teaspoon	homer	hospital	collection
garlic	game	drug	exhibition
sauce	yankees	disease	photograph
onion	pitcher	medical	museum
chopped	season	therapy	images
add	hitter	surgery	gallery
butter	pitch	treatment	exhibit

Table 3: 8 topics discovered by our model on NYT.

Table 2. The supervised Scholar outperforms its unsupervised version on 20 Newsgroups but the unsupervised one achieves higher coherence scores on DBpedia. While sLDA fails to surpass its unsupervised counterpart LDA on both DBpedia and 20 Newsgroups. On the contrary, improvements of sToMCAT over the unsupervised ToMCAT can be observed under all settings. The results show that the incorporation of the supervised information seems to be more effective in our proposed model, probably contributing to the gradient-based loss balancing mechanism. Overall, our model consistently outperforms sLDA and Scholar on all datasets and all topic coherence measures.

4.2.3 Impact of Topic Numbers

To investigate how topic coherence scores vary with respect to different topic number settings, we show in Figure 2 the topic coherence measures on four datasets for all models. Although there are exceptions that some baselines achieve higher scores on specific experimental settings, the general conclusion is that our models perform the best in both unsupervised and supervised topic modeling tasks. On DBpedia and 20 Newsgroups datasets, sToMCAT consistently outperforms ToMCAT, indicating the additional supervision helps generate more co-

¹⁰<https://github.com/AKSW/Palmetto>

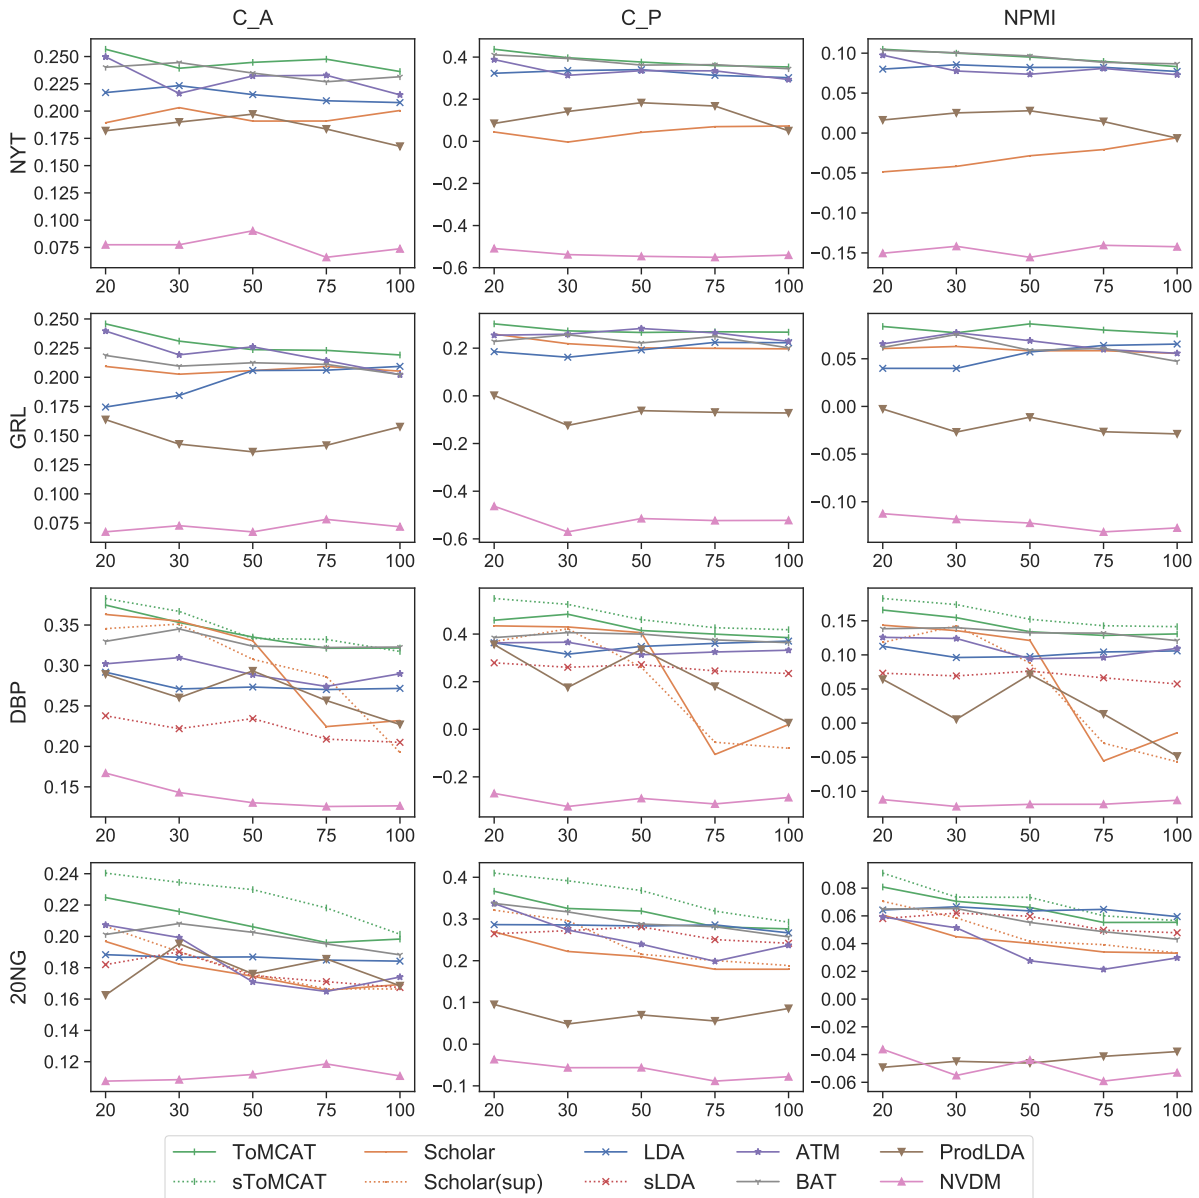


Figure 2: Topic coherence (C_A, C_P, NPMI) w.r.t. topic numbers on 4 datasets. Dotted lines denote supervised topic models.

herent topics. We also notice that although the topic coherence measures of our models remain relatively stable across topic numbers, there are slight drops on the DBpedia and 20 Newsgroups datasets when the topic number becomes bigger. This phenomenon may result from the fact that DBpedia and 20 Newsgroups datasets are less diverse than others. There are only 14 and 20 categories in DBpedia and 20 Newsgroups datasets, respectively. When the topic number is much larger than the ground-truth category number, discriminating different topics would be more challenging. Nevertheless, the overall superiorities of our models are significant as in Figure 2.

4.3 Text Classification

We now report text classification results of supervised topic models: sLDA, Scholar, and sToMCAT. To show that our model can learn both coherent and informative topics concurrently, we use the same models as in the topic modeling experiments to classify test set documents, and do not perform any further fine-tuning. In our experiments, we found that the text classification performance is influenced by topic numbers. Therefore we conduct experiments with five topic number settings: 20, 30, 50, 75, and 100.

Classification results are presented in Table 4. We can see that our model not only achieves the

Dataset	Model	20 (↑)	30 (↑)	50 (↑)	75 (↑)	100 (↑)	Min (↑)	Avg (↑)	Max (↑)	Δ (↓)
DBP	sLDA	0.871	0.906	0.909	0.918	0.922	0.871	0.905	0.922	0.051
	Scholar	0.949	0.951	0.948	0.920	0.900	0.900	0.934	0.951	0.051
	sToMCAT	0.951	0.951	0.953	0.936	0.928	0.928	0.944	0.953	0.025
20NG	sLDA	0.529	0.572	0.563	0.608	0.613	0.529	0.577	0.613	0.084
	Scholar	0.523	0.576	0.598	0.617	0.610	0.523	0.585	0.617	0.094
	sToMCAT	0.642	0.628	0.616	0.617	0.616	0.616	0.624	0.642	0.026

Table 4: Classification accuracy of supervised topic models with different topic numbers (20, 30, 50, 75, 100). ‘Min/Avg/Max’ shows the minimum/average/maximum accuracy among different topic numbers. ‘Δ’ shows the variance of the classification accuracy across different topic numbers.

best overall performance (the Max and Avg column), but also has the highest accuracies on all dataset and topic number settings. Compared to Scholar, our model achieves a slightly higher accuracy on DBpedia and an accuracy improvement of 2.5% on 20 Newsgroups. The performance gain of our model over sLDA is more significant. In addition to better classification results, our model is also more robust to the change of topic numbers (the Δ column). With the topic number increasing from 20 to 100, the variance of the classification accuracy of our model is only 0.025 and 0.026 on DBpedia and 20 Newsgroups respectively, which is much lower than that of sLDA and Scholar.

5 Conclusion

We have presented ToMCAT, a neural topic model with adversarial and cycle-consistent objectives, and its supervised extension, sToMCAT. ToMCAT employs a generator to capture semantic patterns in topics and an encoder to encode documents into their corresponding topics. sToMCAT further incorporates document labels into topic modeling. The effectiveness of ToMCAT and sToMCAT is verified by experiments on topic modeling and text classification. In the future, we plan to extend our model to cope with external word or document semantics. It would also be interesting to explore alternative architectures other than CycleGAN under our formulation of topic modeling.

Acknowledgments

The authors would like to thank the anonymous reviewers for insightful comments and helpful suggestions. This work was funded in part by the National Key Research and Development Program of China (2016YFC1306704) and the National Natural Science Foundation of China (61772132).

References

- Nikolaos Aletras and Mark Stevenson. 2013. [Evaluating topic coherence using distributional semantics](#). In *Proceedings of the 10th International Conference on Computational Semantics (IWCS 2013) – Long Papers*, pages 13–22, Potsdam, Germany. Association for Computational Linguistics.
- Martin Arjovsky, Soumith Chintala, and Léon Bottou. 2017. [Wasserstein generative adversarial networks](#). In *Proceedings of the 34th International Conference on Machine Learning*, volume 70, pages 214–223, Sydney, Australia. PMLR.
- David M. Blei, Andrew Y. Ng, and Michael I. Jordan. 2003. [Latent dirichlet allocation](#). *J. Mach. Learn. Res.*, 3:993–1022.
- G. Bouma. 2009. Normalized (pointwise) mutual information in collocation extraction. In *From Form to Meaning: Processing Texts Automatically, Proceedings of the Biennial GSCL Conference 2009*, volume Normalized, pages 31–40, Tübingen.
- Dallas Card, Chenhao Tan, and Noah A. Smith. 2018. [Neural models for documents with metadata](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2031–2040, Melbourne, Australia. Association for Computational Linguistics.
- Jeff Donahue, Philipp Krähenbühl, and Trevor Darrell. 2016. [Adversarial feature learning](#). *arXiv preprint arXiv:1605.09782*.
- Jacob Eisenstein, Amr Ahmed, and Eric P. Xing. 2011. [Sparse additive generative models of text](#). In *Proceedings of the 28th International Conference on International Conference on Machine Learning, ICML’11*, page 1041–1048, Madison, WI, USA. Omnipress.
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. [Generative adversarial nets](#). In *Advances in Neural Information Processing Systems 27*, pages 2672–2680. Curran Associates, Inc.

- Thomas L. Griffiths and Mark Steyvers. 2004. [Finding scientific topics](#). *Proceedings of the National Academy of Sciences*, 101(suppl 1):5228–5235.
- Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron C Courville. 2017. [Improved training of wasserstein gans](#). In *Advances in Neural Information Processing Systems 30*, pages 5767–5777. Curran Associates, Inc.
- Sergey Ioffe and Christian Szegedy. 2015. [Batch normalization: Accelerating deep network training by reducing internal covariate shift](#). In *Proceedings of the 32nd International Conference on Machine Learning*, volume 37, pages 448–456, Lille, France. PMLR.
- Taeksoo Kim, Moon-su Cha, Hyunsoo Kim, Jung Kwon Lee, and Jiwon Kim. 2017. [Learning to discover cross-domain relations with generative adversarial networks](#). In *Proceedings of the 34th International Conference on Machine Learning*, volume 70, pages 1857–1865, International Convention Centre, Sydney, Australia. PMLR.
- Diederik P Kingma and Jimmy Ba. 2014. [Adam: A method for stochastic optimization](#). *arXiv preprint arXiv:1412.6980*.
- Diederik P Kingma and Max Welling. 2013. [Auto-encoding variational bayes](#). *arXiv preprint arXiv:1312.6114*.
- C. Lee, Y. Wang, T. Hsu, K. Chen, H. Lee, and L. Lee. 2018. [Scalable sentiment for sequence-to-sequence chatbot response with performance analysis](#). In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6164–6168.
- Chenghua Lin and Yulan He. 2009. [Joint sentiment/topic model for sentiment analysis](#). In *Proceedings of the 18th ACM Conference on Information and Knowledge Management, CIKM '09*, pages 375–384, New York, NY, USA. ACM.
- Andrew L Maas, Awni Y Hannun, and Andrew Y Ng. 2013. [Rectifier nonlinearities improve neural network acoustic models](#). In *ICML Workshop on Deep Learning for Audio, Speech and Language Processing*. Citeseer.
- David J. C. MacKay. 1998. [Choice of basis for laplace approximation](#). *Mach. Learn.*, 33(1):77–86.
- Julian McAuley and Jure Leskovec. 2013. [Hidden factors and hidden topics](#). In *Proceedings of the 7th ACM conference on Recommender systems*. ACM Press.
- Jon D. McAuliffe and David M. Blei. 2008. [Supervised topic models](#). In *Advances in Neural Information Processing Systems 20*, pages 121–128. Curran Associates, Inc.
- Yishu Miao, Edward Grefenstette, and Phil Blunsom. 2017. [Discovering discrete latent topics with neural variational inference](#). In *Proceedings of the 34th International Conference on Machine Learning*, volume 70, pages 2410–2419, International Convention Centre, Sydney, Australia. PMLR.
- Yishu Miao, Lei Yu, and Phil Blunsom. 2016. [Neural variational inference for text processing](#). In *Proceedings of The 33rd International Conference on Machine Learning*, volume 48, pages 1727–1736, New York, New York, USA. PMLR.
- David Newman, Jey Han Lau, Karl Grieser, and Timothy Baldwin. 2010. [Automatic evaluation of topic coherence](#). In *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, pages 100–108, Los Angeles, California. Association for Computational Linguistics.
- Michael Röder, Andreas Both, and Alexander Hinneburg. 2015. [Exploring the space of topic coherence measures](#). In *Proceedings of the Eighth ACM International Conference on Web Search and Data Mining, WSDM '15*, pages 399–408, New York, NY, USA. ACM.
- Akash Srivastava and Charles Sutton. 2017. [Auto-encoding variational inference for topic models](#). *arXiv preprint arXiv:1703.01488*.
- Hanna M. Wallach, David M. Mimno, and Andrew McCallum. 2009. [Rethinking lda: Why priors matter](#). In Y. Bengio, D. Schuurmans, J. D. Lafferty, C. K. I. Williams, and A. Culotta, editors, *Advances in Neural Information Processing Systems 22*, pages 1973–1981. Curran Associates, Inc.
- Rui Wang, Xuemeng Hu, Deyu Zhou, Yulan He, Yuxuan Xiong, Chenchen Ye, and Haiyang Xu. 2020. [Neural topic modeling with bidirectional adversarial training](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 340–350, Online. Association for Computational Linguistics.
- Rui Wang, Deyu Zhou, and Yulan He. 2019a. [ATM: Adversarial-neural topic model](#). *Information Processing & Management*, 56(6):102098.
- Rui Wang, Deyu Zhou, and Yulan He. 2019b. [Open event extraction from online text using a generative adversarial network](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 282–291, Hong Kong, China. Association for Computational Linguistics.
- Xiang Zhang, Junbo Zhao, and Yann LeCun. 2015. [Character-level convolutional networks for text classification](#). In *Advances in Neural Information Processing Systems 28*, pages 649–657. Curran Associates, Inc.

Deyu Zhou, Xuan Zhang, and Yulan He. 2017. [Event extraction from twitter using non-parametric Bayesian mixture model with word embeddings](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, pages 808–817, Valencia, Spain. Association for Computational Linguistics.

Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A. Efros. 2017. [Unpaired image-to-image translation using cycle-consistent adversarial networks](#). In *2017 IEEE International Conference on Computer Vision (ICCV)*. IEEE.

A Discovered Topics on NYTimes

To gain an insight into the extracted topics, we present the full list of 50 topics on NYTimes discovered by ToMCAT in Table 5. As a comparison, topics discovered by LDA are shown in Table 6.

stock fund firm companies online investment investor broker company customer
car tires fuel driver truck vehicle vehicles gas gasoline engine
voter poll campaign percent primary republican vote democratic states democrat
building project apartment house houses homes resident housing estate square
flight passenger airline plane airport customer carrier pilot airlines tires
yard game touchdown play team quarterback season goal offense pass
show actor producer character series film network award television comedy
computer www user site web window file com files mail
court lawsuit case ruling antitrust suit plaintiff judge settlement federal
officer police investigation mayor prosecutor department charges complaint criminal official
film movie character actor movies director comedy script minutes scenes
company merger companies billion deal cable stock acquisition network market
union worker employees company job contract pay employer manager benefit
school student teacher test program district children education percent parent
friend course article black guy thought movie husband film house
fashion designer leather wear dress clothes skirt white shirt pant
black white protester flag town crowd street protest community school
music album song band jazz artist rock guitar musical singer
campaign political money fund president governor presidential republican election candidates
boy father cuban family relatives mother child son custody grandmother
computer privacy software companies information web user sites internet company
jet coach patriot season player team draft coaching defensive football
cell genome scientist genes human genetic researcher gene disease study
died survived degree film graduated served wife student born article
rebel military war soldier troop attack terrorist civilian forces bombing
abortion religious conservative support european conservatives government republican thunderstorm vote
cup tablespoon pepper teaspoon garlic sauce onion chopped add butter
com commentary daily tduncan information toder holiday eta staffed sport
gun gay women firearm law violence sexual percent bill shooting
war church government country african nation communist black priest leader
drug missile nuclear weapon defense official sanction administration missiles countries
forest bird fire species water land fish fires animal acres
book memoir author bookstores fiction ages nonfiction writer reader witchcraft
palestinian israeli peace israelis jewish syrian violence arab summit lebanese
ballot recount votes election vote counties county count board manual
race medal racing meter gold team track driver races lap
tournament fight round par match tour game champion fighter golf
percent survey population economy immigrant economic million companies worker wage
inning run hit homer game yankees pitcher season hitter pitch
fax syndicate www tour com hotel trip ticket room telex
penalty death execution prosecutor murder trial jury inmates prison lawyer
convention speech party campaign republican democratic delegates president presidential democrat
patient cancer doctor hospital drug disease medical therapy surgery treatment
point game team shot rebound pointer foul guard minutes play
tax taxes bill cut surplus income plan proposal spending billion
stock percent quarter earning market index company analyst cent investor
election party government opposition political minister power president country leader
campaign debate debates candidates presidential aides president vice reporter adviser
artist painting art collection exhibition photograph museum images gallery exhibit
yankees team fan player baseball game games football league stadium

Table 5: Full list of 50 topics on NYTimes discovered by ToMCAT.

official agency investigation letter statement comment public office document interview
percent women number according study survey group found likely million
political government president power leader country party election opposition minister
building project home local town resident area center million house
scientist human research cell science researcher found called brain light
country foreign trade countries government nation european economic american international
need problem feel look right hard happen help trying change
film movie character play actor movies director minutes cast role
election vote ballot votes voter count recount result hand campaign
car driver seat truck drive driving road vehicle model wheel
company companies million business firm deal industry billion executive market
job worker employees union manager president working contract member pay
customer sales sell product buy consumer business price market store
water bird fish weather rain animal wind plant land trees
school student program teacher college high education class children public
guy tell look kid bad big dog right real word
court case law decision lawyer federal legal judge right lawsuit
point game play team goal shot games lead left half
look show art collection fashion designer artist style wear painting
round won sport fight shot player final tournament gold event
drug patient doctor medical health cancer hospital disease treatment care
book author writer writing wrote read published magazine find write
campaign republican president presidential democratic voter political candidates candidate convention
meeting official talk agreement deal leader decision conference president negotiation
article special fax information syndicate contact visit buy separate purchased
computer system software technology user program digital window internet access
com question daily newspaper american today information business sport statesman
police death officer case crime prison criminal prosecutor trial victim
military system security defense nuclear weapon official administration attack arm
web site com www sites mail online information internet telegram
word fact sense question perhap course point matter mean view
black group white religious right gay church jewish member flag
money million tax plan pay billion cut cost fund program
flight plane ship crew pilot air passenger boat airport hour
war palestinian peace soldier israeli military troop violence attack killed
history century french known german today american ago died modern
room house wall door floor hand water window light inside
oil prices plant million gas production energy industry power cost
family father children son mother boy home child parent daughter
percent stock market fund quarter growth economy investor earning analyst
room hotel trip restaurant tour travel night visit visitor dinner
night crowd hour morning reporter hand street told moment left
race won win run track winner running racing place winning
music song band sound record album musical show pop rock
cup food minutes add oil tablespoon fat chicken large pepper
bill group law gun support legislation issue member right federal
run hit game season inning yankees home baseball right games
wanted thought told friend asked knew took felt saw ago
show television network media station commercial series radio viewer rating
team player season game play coach yard games football league

Table 6: Full list of 50 topics on NYTimes discovered by LDA.