

KATE: K-Competitive Autoencoder for Text

Yu Chen

Rensselaer Polytechnic Institute
Troy, NY, USA 12180
cheny39@rpi.edu

Mohammed J. Zaki

Rensselaer Polytechnic Institute
Troy, NY, USA 12180
zaki@cs.rpi.edu

ABSTRACT

Autoencoders have been successful in learning meaningful representations from image datasets. However, their performance on text datasets has not been widely studied. Traditional autoencoders tend to learn possibly trivial representations of text documents due to their confounding properties such as high-dimensionality, sparsity and power-law word distributions. In this paper, we propose a novel k-competitive autoencoder, called KATE, for text documents. Due to the competition between the neurons in the hidden layer, each neuron becomes specialized in recognizing specific data patterns, and overall the model can learn meaningful representations of textual data. A comprehensive set of experiments show that KATE can learn better representations than traditional autoencoders including denoising, contractive, variational, and k-sparse autoencoders. Our model also outperforms deep generative models, probabilistic topic models, and even word representation models (e.g., Word2Vec) in terms of several downstream tasks such as document classification, regression, and retrieval.

CCS CONCEPTS

- **Information systems** → **Data mining; Document representation; Retrieval tasks and goals;** • **Computing methodologies** → **Neural networks; Information extraction;**

KEYWORDS

Representation Learning, Autoencoders, Competitive Learning, Text Analytics

ACM Reference format:

Yu Chen and Mohammed J. Zaki. 2017. KATE: K-Competitive Autoencoder for Text. In *Proceedings of KDD '17, Halifax, NS, Canada, August 13-17, 2017*, 10 pages.
<https://doi.org/10.1145/3097983.3098017>

1 INTRODUCTION

An autoencoder is a neural network which can automatically learn data representations by trying to reconstruct its input at the output layer. Many variants of autoencoders have been proposed recently [14, 24–26, 35, 39]. While autoencoders have been successfully applied to learn meaningful representations on image datasets

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

KDD '17, August 13-17, 2017, Halifax, NS, Canada

© 2017 Copyright held by the owner/author(s). Publication rights licensed to Association for Computing Machinery.

ACM ISBN 978-1-4503-4887-4/17/08...\$15.00
<https://doi.org/10.1145/3097983.3098017>

(e.g., MNIST [20], CIFAR-10 [15]), their performance on text datasets has not been widely studied. Traditional autoencoders are susceptible to learning trivial representations for text documents. As noted by Zhai and Zhang [42], the reasons include that fact that textual data is extremely high dimensional and sparse. The vocabulary size can be hundreds of thousands while the average fraction of zero entries in the document vectors can be very high (e.g., 98%). Further, textual data typically follows power-law word distributions. That is, low-frequency words account for most of the word occurrences. Traditional autoencoders always try to reconstruct each dimension of the input vector on an equal footing, which is not quite appropriate for textual data.

Document representation is an interesting and challenging task which is concerned with representing textual documents in a vector space, and it has various applications in text processing, retrieval and mining. There are two major approaches to represent documents: 1) **Distributional Representation** is based on the hypothesis that linguistic terms with similar distributions have similar meanings. These methods usually take advantage of the co-occurrence and context information of words and documents, and each dimension of the document vector usually represents a specific semantic meaning (e.g., a topic). Typical models in this category include Latent Semantic Analysis (LSA) [7], probabilistic LSA (pLSA) [12] and Latent Dirichlet Allocation (LDA) [2]. 2) **Distributed Representations** encode a document as a compact, dense and lower dimensional vector with the semantic meaning of the document distributed along the dimensions of the vector. Many neural network-based distributed representation models [5, 18, 22, 27, 36] have been proposed and shown to be able to learn better representations of documents than distributional representation models.

In this paper, we try to overcome the weaknesses of traditional autoencoders when applied to textual data. We propose a novel autoencoder called KATE (for **K**-competitive **A**utoencoder for **T**ext), which relies on competitive learning among the autoencoding neurons. In the feedforward phase, only the most competitive k neurons in the layer fire and those k “winners” further incorporate the aggregate activation potential of the remaining inactive neurons. As a result, each hidden neuron becomes better at recognizing specific data patterns and the overall model can learn meaningful representations of the input data. After training the model, each hidden neuron is distinct from the others and no competition is needed in the testing/encoding phase. We conduct comprehensive experiments qualitatively and quantitatively to evaluate KATE and to demonstrate the effectiveness of our model. We compare KATE with traditional autoencoders including basic autoencoder, denoising autoencoder [39], contractive autoencoder [35], variational autoencoder [14], and k-sparse autoencoder [24]. We also compare with deep generative models [22], neural autoregressive [18] and variational inference [27] models, probabilistic topic models such as

LDA [2], and word representation models such as Word2Vec [28] and Doc2Vec [19]. KATE achieves state-of-the-art performance across various datasets on several downstream tasks like document classification, regression and retrieval.

2 RELATED WORK

Autoencoders. The basic autoencoder is a shallow neural network which tries to reconstruct its input at the output layer. An autoencoder consists of an encoder which maps the input \mathbf{x} to the hidden layer: $\mathbf{z} = g(\mathbf{W}\mathbf{x} + \mathbf{b})$ and a decoder which reconstructs the input as: $\hat{\mathbf{x}} = o(\mathbf{W}'\mathbf{z} + \mathbf{c})$; here \mathbf{b} and \mathbf{c} are bias terms, \mathbf{W} and \mathbf{W}' are input-to-hidden and hidden-to-output layer weight matrices, and g and o are activation functions. Weight tying (i.e., setting $\mathbf{W}' = \mathbf{W}^T$) is often used as a regularization method to avoid overfitting. While plain autoencoders, even with perfect reconstructions, usually only extract trivial representations of the data, more meaningful representations can be obtained by adding appropriate regularization to the models. Following this line of reasoning, many variants of autoencoders have been proposed recently [14, 24–26, 35, 39]. The denoising autoencoder (DAE) [39] inputs a corrupted version of the data while the output is still compared with the original uncorrupted data, allowing the model to learn patterns useful for denoising. The contractive autoencoder (CAE) [35] introduces the Frobenius norm of the Jacobian matrix of the encoder activations into the regularization term. When the Frobenius norm is 0, the model is extremely invariant to perturbations of input data, which is thought as good. The variational autoencoder (VAE) [14] is a generative model inspired by variational inference whose encoder $q_\phi(\mathbf{z}|\mathbf{x})$ approximates the intractable true posterior $p_\theta(\mathbf{z}|\mathbf{x})$, and the decoder $p_\theta(\mathbf{x}|\mathbf{z})$ is a data generator. The k-sparse autoencoder (KSAE) [24] explicitly enforces sparsity by only keeping the k highest activities in the feedforward phase.

We notice that most of the successful applications of autoencoders are on image data, while only a few have attempted to apply autoencoders on textual data. Zhai and Zhang [42] have argued that traditional autoencoders, which perform well on image data, are less appropriate for modeling textual data due to the problems of high-dimensionality, sparsity and power-law word distributions. They proposed a semi-supervised autoencoder which applies a weighted loss function where the weights are learned by a linear classifier to overcome some of these problems. Kumar and D’Haro [16] found that all the topics extracted from the autoencoder were dominated by the most frequent words due to the sparsity of the input document vectors. Further, they found that adding sparsity and selectivity penalty terms helped alleviate this issue to some extent.

Deep generative models. Deep Belief Networks (DBNs) are probabilistic graphical models which learn to extract a deep hierarchical representation of the data. The top 2 layers of DBNs form a Restricted Boltzmann Machine (RBM) and other layers form a sigmoid belief network. A relatively fast greedy layer-wise pre-training algorithm [9, 10] is applied to train the model. Maaloe et al. [22] showed that DBNs can be competitive as a topic model. DocNADE [18] is a neural autoregressive topic model that estimates the probability of observing a new word in a given document given the previously observed words. It can be used for extracting meaningful

representations of documents. It has been shown to outperform the Replicated Softmax model [11] which is a variant of RBMs for document modeling. Srivastava et al. [36] introduced a type of Deep Boltzmann Machine (DBM) that is suitable for extracting distributed semantic representations from a corpus of documents; an Over-Replicated Softmax model was proposed to overcome the apparent difficulty of training a DBM. NVDM [27] is a neural variational inference model for document modeling inspired by the variational autoencoder.

Probabilistic topic models. Probabilistic topic models, such as probabilistic Latent Semantic Analysis (pLSA) and Latent Dirichlet Allocation (LDA) have been extensively studied [2, 12]. Especially for LDA, many variants have been proposed for non-parametric learning [1, 37], sparsity [8, 40, 43] and efficient inference [4, 38]. Those models typically build a generative probabilistic model using the bag-of-words representation of the documents.

Word representation models. Distributed representations of words in a vector space can capture semantic meanings of words and help achieve better results in various downstream text analysis tasks. Word2Vec [28] and Glove [33] are state-of-the-art word representation models. Pre-training word embeddings on a large corpus of documents and applying learned word embeddings in downstream tasks has been shown to work well in practice [6, 13, 29]. Doc2Vec [19] was inspired by Word2Vec and can directly learn vector representations of paragraphs and documents. NTM [5], which also uses pre-trained word embeddings, is a neural topic model where the representations of words and documents are combined into a uniform framework.

With this brief overview of existing work, we now turn to our competitive autoencoder approach for text documents.

3 K-COMPETITIVE AUTOENCODER

Although the objective of an autoencoder is to minimize the reconstruction error, our goal is to extract meaningful features from data. Compared with image data, textual data is more challenging for autoencoders since it is typically high-dimensional, sparse and has power-law word distributions. When examining the features extracted by an autoencoder, we observed that they were not distinct from one another. That is, many neurons in the hidden layer shared similar groups of input neurons (which typically correspond to the most frequent words) with whom they had the strongest connections. We hypothesized that the autoencoder greedily learned relatively trivial features in order to reconstruct the input.

To overcome this drawback, our approach guides the autoencoder to focus on important patterns in the data by adding constraints in the training phase via mutual competition. In competitive learning, neurons compete for the right to respond to a subset of the input data and as a result, the specialization of each neuron in the network is increased. Note that the specialization of neurons is exactly what we want for an autoencoder, especially when applied on textual data. By introducing competition into an autoencoder, we expect each neuron in the hidden layer to take responsibility for recognizing different patterns within the input data. Following this line of reasoning, we propose the k-competitive autoencoder, KATE, as described below.

3.1 Training and Testing/Encoding

Algorithm 1 KATE: K-competitive Autoencoder

```

1: procedure TRAINING
2:   Feedforward step:  $z = \tanh(W\mathbf{x} + \mathbf{b})$ 
3:   Apply k-competition:  $\hat{z} = \text{k-competitive\_layer}(z)$ 
4:   Compute output:  $\hat{\mathbf{x}} = \text{sigmoid}(W^T\hat{z} + \mathbf{c})$ 
5:   Backpropagate error (cross-entropy) and iterate

1: procedure ENCODING
2:   Encode input data:  $z = \tanh(W\mathbf{x} + \mathbf{b})$ 

```

The pseudo-code for our k-competitive autoencoder KATE is shown in Algorithm 1. KATE is a shallow autoencoder with a (single) competitive hidden layer, with each neuron competing for the right to respond to a given set of input patterns. Let $\mathbf{x} \in \mathbb{R}^d$ be a d -dimensional input vector, which is also the desired output vector, and let h_1, h_2, \dots, h_m be the m hidden neurons. Let $W \in \mathbb{R}^{d \times m}$ be the weight matrix linking the input layer to the hidden layer neurons, and let $\mathbf{b} \in \mathbb{R}^m$ and $\mathbf{c} \in \mathbb{R}^d$ be the bias terms for the hidden and output neurons, respectively. Let g be an activation function; two typical functions are $\tanh(x) = \frac{e^{2x}-1}{e^{2x}+1}$ and $\text{sigmoid}(x) = \frac{1}{1+e^{-x}}$. In each feed-forward step, the activation potential at the hidden neurons is then given as $z = g(W\mathbf{x} + \mathbf{b})$, whereas the activation potential at the output neurons is given as $\hat{\mathbf{x}} = g(W^Tz + \mathbf{c})$. Thus, the hidden-to-output weight matrix is simply W^T , being an instance of *weight tying*.

In KATE, we represent each input text document as a log-normalized word count vector $\mathbf{x} \in \mathbb{R}^d$ where each dimension is represented as

$$x_i = \frac{\log(1 + n_i)}{\max_{i \in V} \log(1 + n_i)}, \text{ for } i \in V$$

where V is the vocabulary and n_i is the count of word i in that document. Let $\hat{\mathbf{x}}$ be the output of KATE on a given input \mathbf{x} . We use the binary cross-entropy as the loss function, which is defined as

$$l(\mathbf{x}, \hat{\mathbf{x}}) = - \sum_{i \in V} x_i \log(\hat{x}_i) + (1 - x_i) \log(1 - \hat{x}_i)$$

where \hat{x}_i is the reconstructed value for x_i .

Let H be some subset of hidden neurons; define the **energy** of H as the total activation potential for H , given as: $E(H) = \sum_{h_i \in H} |z_i|$, i.e., sum of the absolute values of the activations for neurons in H . In KATE, in the feedforward phase, after computing the activations z for a given input \mathbf{x} , we select the most competitive k neurons as the “winners” while the remaining “losers” are suppressed (i.e., made inactive). However, in order to compensate for the loss of energy from the loser neurons, and to make the competition among neurons more pronounced, we amplify and reallocate that energy among the winner neurons.

KATE uses \tanh activation function for the k-competitive hidden layer. We divide these neurons into positive and negative neurons based on their activations. The most competitive k neurons are those that have the largest absolute activation values. However, we select the $\lceil k/2 \rceil$ largest positive activations as the positive winners, and reallocate the energy of the remaining positive loser neurons among the winners using an α amplification connection, where

Algorithm 2 K-competitive Layer

```

1: function K-COMPETITIVE-LAYER( $z$ )
2:   sort positive neurons in ascending order  $z_1^+ \dots z_P^+$ 
3:   sort negative neurons in descending order  $z_1^- \dots z_N^-$ 
4:   if  $P - \lceil k/2 \rceil > 0$  then
5:      $E_{pos} = \sum_{i=1}^{P-\lceil k/2 \rceil} z_i^+$ 
6:     for  $i = P - \lceil k/2 \rceil + 1, \dots, P$  do
7:        $z_i^+ := z_i^+ + \alpha \cdot E_{pos}$ 
8:     for  $i = 1, \dots, P - \lceil k/2 \rceil$  do
9:        $z_i^+ := 0$ 
10:    if  $N - \lceil k/2 \rceil > 0$  then
11:       $E_{neg} = \sum_{i=1}^{N-\lceil k/2 \rceil} z_i^-$ 
12:      for  $i = N - \lceil k/2 \rceil + 1, \dots, N$  do
13:         $z_i^- := z_i^- + \alpha \cdot E_{neg}$ 
14:      for  $i = 1, \dots, N - \lceil k/2 \rceil$  do
15:         $z_i^- := 0$ 
16:    return updated  $z_1^+ \dots z_P^+, z_1^- \dots z_N^-$ 

```




Figure 1: Competitions among neurons. Input and hidden neurons, and hidden and output neurons are fully connected, but we omit these to avoid clutter.

α is a hyperparameter. Finally, we set the activations of all losers to zero. Similarly, the $\lfloor k/2 \rfloor$ lowest negative activations are the negative winners, and they incorporate the amplified energy from the negative loser neurons, as detailed in Algorithm 2. We argue that the α amplification connections are a critical component in the k-competitive layer. When $\alpha = 0$, no gradients will flow through loser neurons, resulting in a regular k-sparse autoencoder (regardless of the activation functions and k-selection scheme). When $\alpha > 2/k$, we actually boost the gradient signal flowing through the loser neurons. We empirically show that amplification helps improve the autoencoder model (see Sec. 4.4.1 and 4.5). As an example, consider Figure 1, which shows an example feedforward step for $k = 2$. Here, h_1 and h_6 are the positive and negative winners, respectively, since the absolute activation potential for h_1 is $|z_1| = 0.8$, and for h_6 it is $|z_6| = 0.6$. The positive winner h_1 takes away the energy from the positive losers h_2 and h_3 , which is $E(\{h_2, h_3\}) = 0.2 + 0.1 = 0.3$.

Likewise, the negative winner h_6 takes away the energy from the negative losers h_4 and h_5 , which is $-E(\{h_4, h_5\}) = -(|-0.1| + |-0.3|) = -0.4$. The hyperparameter α governs how the energy from the loser neurons is incorporated into the winner neurons, for both positive and negative cases. That is h_1 's net activation becomes $z_1 = 0.8 + 0.3\alpha$, and h_6 's net activation is $z_6 = -0.6 - 0.4\alpha$. The rest of the neurons are set to zero activation.

Finally, as noted in Algorithm 1 we use weight tying for the hidden to output layer weights, i.e., we use W^T as the weight matrix, with different biases c . Also, since the inputs x are non-negative for document representations (e.g., word counts), we use the *sigmoid* activation function at the output layer to maintain the non-negativity. Note that in the back-propagation procedure, the gradients will first flow through the winner neurons in the hidden layer and then the loser neurons via the α amplification connections. No gradients will flow directly from the output neurons to the loser neurons since they are made inactive in the feedforward step.

Testing/Encoding. Once the k-competitive network has been trained, we simply encode each test input as shown in Algorithm 1. That is, given a test input x , we map it to the feature space to obtain $z = \tanh(Wx + b)$. No competition is required for the encoding step since the hidden neurons are well trained to be distinctive from others. We argue that this is one of the superior features of KATE.

3.2 Relationship to Other Models

KATE vs. K-Sparse Autoencoder. The k-sparse autoencoder [24] is closely related to our model, but there are several differences. The k-sparse autoencoder explicitly enforces sparsity by only keeping the k highest activities at training time. Then, at testing time, in order to enforce sparsity, only the αk highest activities are kept where α is a hyperparameter. Since its hidden layer uses a linear activation function, the only non-linearity in the encoder comes from the selection of the k highest activities. Instead of focusing on sparsity, our model focuses on competition to drive each hidden neuron to be distinct from the others. Thus, at testing time, no competition is needed. The non-linearity in KATE's encoding comes from the *tanh* activation function, and the winner-take-all operation via the top k (+ve/-ve) selection and energy reallocation/amplification.

It is important to note that for the k-sparse autoencoder, too much sparsity (i.e., low k) can cause the so-called “dead” hidden neurons problem, which can prevent gradient back-propagation from adjusting the weights of these “dead” hidden neurons. As mentioned in the original paper, the model is prone to behaving in a manner similar to k-means clustering. That is, in the first few epochs, it will greedily assign individual hidden neurons to groups of training cases and these hidden neurons will be re-enforced but other hidden neurons will not be adjusted in subsequent epochs. In order to address this problem, scheduling the sparsity level over epochs was suggested. However, by design our approach does not suffer from this problem since the gradients will still flow through the loser neurons via the α amplification connections in the k-competitive layer.

KATE vs. K-Max Pooling. Our proposed k-competitive operation is also reminiscent of the k-max pooling operation [3] applied in convolutional neural networks. We can intuitively regard k-max

pooling as a global feature sampler which selects a subset of k maximum neurons in the previous convolutional layer and uses only the selected subset of neurons in the following layer. Unlike our k-competitive approach, the objective of k-max pooling is to reduce dimensionality and introduce feature invariance via this downsampling operation.

KATE as a Regularized Autoencoder. We can also regard our model as a special case of a fully competitive autoencoder where all the neurons in the hidden layer are fully connected with each other and the weights on the connections between them are fully trainable. The difference is that we restrict the architecture of this competitive layer by using a positive adder and a negative adder to constrain the energy, which serves as a regularization method.

4 EXPERIMENTS

In this section, we evaluate our k-competitive autoencoder model on various datasets and downstream text analytics tasks to gauge its effectiveness in learning meaningful representations in different situations. All experiments were performed on a machine with a 1.7GHz AMD Opteron 6272 Processor, with 264G RAM. Our model, KATE, was implemented in Keras (github.com/fchollet/keras) which is a high-level neural networks library, written in Python. The source code for KATE is available at github.com/hugochan/KATE.

dataset	20 news	reuters	wiki10+	mrd
train.size	11,314	554,414	13,972	3,337
test.size	7,532	250,000	6,000	1,669
valid.size	1,000	10,000	1,000	300
vocab.size	2,000	5,000	2,000	2,000
avg.length	93	112	1,299	124
classes/vals	20	103	25	[0, 1]
task	class & DR	MLC	MLC	regression

Table 1: Datasets: Tasks include classification (class), regression, multi-label classification (MLC), and document retrieval (DR).

4.1 Datasets

For evaluation, we use datasets that have been widely used in previous studies [17, 21, 30–32, 44]. Table 1 provides statistics of the different datasets used in our experiments. It lists the training, testing and validation (a subset of training) set sizes, the size of the vocabulary, average document length, the number of classes (or values for regression), and the various downstream tasks we perform on the datasets.

The 20 Newsgroups [17] (www.qwone.com/~jason/20Newsgroups) data consists of 18846 documents, which are partitioned (nearly) evenly across 20 different newsgroups. Each document belongs to exactly one newsgroup. The corpus is divided by date into training (60%) and testing (40%) sets. We follow the preprocessing steps utilized in previous work [18, 27, 36]. That is, after removing stop-words and stemming, we keep the most frequent 2,000 words in the training set as the vocabulary. We use this dataset to show that our model can learn meaningful representations for classification and document retrieval tasks.

The Reuters RCV1-v2 dataset [21] (www.jmlr.org/papers/volume5/lewis04a) contains 804,414 newswire articles, where each document

typically has multiple (hierarchical) topic labels. The total number of topic labels is 103. The dataset already comes preprocessed with stopword removal and stemming. We randomly split the corpus into 554,414 training and 25,000 test cases and keep the most frequent 5,000 words in the training dataset as the vocabulary. We perform multi-label classification on this dataset.

The Wiki10+ dataset [44] ([www.zubiaga.org/datasets/wiki10+/\](http://www.zubiaga.org/datasets/wiki10+/)) comprises English Wikipedia articles with at least 10 annotations on delicious.com. Following the steps of Cao et al. [5], we only keep the 25 most frequent social tags and those documents containing any of these tags. After removing stopwords and stemming, we randomly split the corpus into 13,972 training and 6,000 test cases and keep the most frequent 2,000 words in the training set as the vocabulary for use in multi-label classification.

The Movie review data (MRD) [30–32] ([www.cs.cornell.edu/people/pabo/movie-review-data/\](http://www.cs.cornell.edu/people/pabo/movie-review-data/)) contains a collection of movie-review documents, with a numerical rating score in the interval [0, 1]. After removing stopwords and stemming, we randomly split the corpus into 3,337 training and 1,669 test cases and keep the most frequent 2,000 words in the training set as the vocabulary. We use this dataset for regression, i.e., predicting the movie ratings.

Note that among the above datasets, only the 20 Newsgroups dataset is balanced, whereas both the Reuters and Wiki10+ datasets are highly imbalanced in terms of class labels.

4.2 Comparison with Baseline Methods

We compare our k-competitive autoencoder KATE with a wide range of other models including various types of autoencoders, topic models, belief networks and word representation models, as listed below.

LDA [2]: a directed graphical model which models a document as a mixture of topics and a topic as a mixture of words. Once trained, each document can be represented as a topic proportion vector on the topic simplex. We used the gensim [34] LDA implementation in our experiments.

DocNADE [18]: a neural autoregressive topic model that can be used for extracting meaningful representations of documents. The implementation is available at www.dmi.usherb.ca/~larocheh/code/DocNADE.zip.

DBN [22]: a direct acyclic graph whose top two layers form a restricted Boltzmann machine. We use the implementation available at github.com/larsmaaloe/deep-belief-nets-for-topic-modeling.

NVDM [27]: a neural variational inference model for document modeling. The authors have not released the source code, but we used an open-source implementation at github.com/carpedm20/variational-text-tensor-flow.

Word2Vec [28]: a model in which each document is represented as the average of the word embedding vectors for that document. We use Word2Vec_{pre} to denote the version where we use Google News pre-trained word embeddings which contain 300-dimensional vectors for 3 million words and phrases. Those embeddings were trained by state-of-the-art word2vec skipgram model. On the other hand, we use Word2Vec to denote the version where we train word embeddings separately on each of our datasets, using the gensim [34] implementation.

Doc2Vec [19]: a distributed representation model inspired by Word2Vec which can directly learn vector representations of documents. There are two versions named Doc2Vec-DBOW and Doc2Vec-DM. We use Doc2Vec-DM in our experiments as it was reported to consistently outperform Doc2Vec-DBOW in the original paper. We used the gensim [34] implementation in our experiments.

AE: a plain shallow (i.e., one hidden layer) autoencoder, without any competition, which can automatically learn data representations by trying to reconstruct its input at the output layer.

DAE [39]: a denoising autoencoder that accepts a corrupted version of the input data while the output is still the original uncorrupted data. In our experiments, we found that masking noise consistently outperforms other two types of noise, namely Gaussian noise and salt-and-pepper noise. Thus, we only report the results of using masking noise. Basically, masking noise perturbs the input by setting a fraction v of the elements i in each input vector as 0. To be fair and consistent, we use a shallow denoising autoencoder in our experiments.

CAE [35]: a contractive autoencoder which introduces the Frobenius norm of the Jacobian matrix of the encoder activations into the regularization term.

VAE [14]: a generative autoencoder inspired by variational inference.

KSAE [24]: a competitive autoencoder which explicitly enforces sparsity by only keeping the k highest activities in the feedforward phase.

We implemented the AE, DAE, CAE, VAE and KSAE autoencoders on our own, since their implementations are not publicly available.

Training Details: For all the autoencoder models (including AE, DAE, CAE, VAE, KSAE, and KATE), we represent each input document as a log-normalized word count vector, using binary cross-entropy as the loss function and Adadelta [41] as the optimizer. Weight tying is also applied. For CAE and VAE, additional regularization terms are added to the loss function as mentioned in the original papers. As for VAE, we use \tanh as the nonlinear activation function while as for AE, DAE and CAE, sigmoid is applied. As for KSAE, we found that omitting sparsity in the testing phase gave us better results in all experiments.

When training models, we randomly extract a subset of documents from the training set as a validation set, as noted in Table 1, which is used for tuning hyperparameters and early stopping. Early stopping is a type of regularization used to avoid overfitting when training an iterative algorithm. We stop training after 5 successive epochs with no improvement on the validation set. All baseline models were optimized as recommended in original sources. For KATE, we set α as 6.26, learning rate as 2, batch size as 100 (for the Reuters dataset) or 50 (for other datasets) and k as 6 (for the 20 topics case), 32 (for the 128 topics case) or 102 (for the 512 topics case), as determined from the validation set.

4.3 Qualitative Analysis

In this set of qualitative experiments, we compare the topics generated by KATE to other representative models including AE, KSAE, and LDA. Even though KATE is not explicitly designed for the purpose of word embeddings, we compared word representations

learned by KATE with the Word2Vec model to demonstrate that our model can learn semantically meaningful representations from text. We evaluate the above models on the 20 Newsgroups data. Matching the number of classes, the number of topics is set to exactly 20 for all models. For both KSAE and KATE, k (the sparsity level/number of winning neurons) is set as 6.

soc.religion.christian	AE	line subject organ articl peopl post time make write good
	KSAE	peopl origin bottom applic mind subject pad europ role christian
	KATE	god christian jesu moral rutger bibl exist religion apr christ
	LDA	god christian jesu church bibl peopl christ man time life
sci.crypt	KATE	govern articl law key encrypt clipper chip secur case distribut
	LDA	key encrypt chip clipper secur govern system law escrow privaci
comp.os.ms-windows.misc	KATE	ca system univers window problem card file driver drive scsi
	LDA	drive card gun control disk scsi system driver hard bu

Table 2: Topics learned by various models.

4.3.1 Topics Generated by Different Models. Table 2 shows some topics learned by various models. As for autoencoders, each topic is represented by the 10 words (i.e., input neurons) with the strongest connection to that topic (i.e., hidden neuron). As for LDA, each topic is represented by the 10 most probable words in that topic. The basic AE is not very good at learning distinctive topics from textual data. In our experiment, all the topics learned by AE are dominated by frequent common words like *line*, *subject* and *organ*, which were always the top 3 words in all the 20 topics. KSAE learns some meaningful words but only alleviates this problem to some extent, for example, *line*, *subject*, *organ* and *white* still appears as top 4 words in 6 topics. For this reason, the output of AE and KSAE is shown for only one of the newsgroups (*soc.religion.christian*). On the other hand, we find that KATE generates 20 topics that are distinct from each other, and which capture the underlying semantics very well. For example, it associates words such as *god*, *christian*, *jesu*, *moral*, *bibl*, *exist*, *religion*, *christ* under the topic *soc.religion.christian*. It is worth emphasizing that KATE belongs to the class of distributed representation models, where each topic is “distributed” among a group of hidden neurons (the topics are therefore better interpreted as “virtual” topics). However, we find that KATE can generate competitive topics compared with LDA, which explicitly infers topics as mixture of words.

4.3.2 Word Embeddings Learned by Different Models. In AE, KSAE and KATE, each input neuron (i.e., a word in the vocabulary set) is connected to each hidden neuron (i.e., a virtual topic) with different strengths. Thus, each row i of the input to hidden layer weight matrix $W \in \mathbb{R}^{d \times m}$ is taken as an m -dimensional word embedding for word i . In order to evaluate whether KATE can capture semantically meaningful word representations, we check if similar or related words are close to each other in the vector space. Table 3 shows the five nearest neighbors for some query words in the word representation space learned by AE, KSAE, KATE and

Word2Vec. KATE performs much better than AE and KSAE. For example KATE lists words like *arm*, *crime*, *gun*, *firearm*, *handgun* among the nearest neighbors of query word *weapon* while neither AE nor KSAE is able to find relevant words. One can observe that KATE can learn competitive word representations compared to Word2Vec in terms of this word similarity task.

4.3.3 Visualization of Document Representations. A good document representation method is expected to group related documents, and to separate the different groups. Figure 2 shows the PCA projections of the document representations taken from the six main groups in the 20 Newsgroups data. As we can observe, neither AE nor the KSAE methods can learn good document representations. On the other hand, KATE successfully extracts meaningful representations from the documents; it automatically clusters related documents in the same group, and it can easily distinguish the six different groups. In fact, KATE is very competitive with LDA (arguably even better on this dataset, since LDA confuses some categories), even though the latter explicitly learns documents representations as mixture of topics, which in turn are mixture of words. Figure 3 shows the T-SNE based visualization [23] of the above document representations and we can draw a similar conclusion.

4.4 Quantitative Experiments

We now turn to quantitative experiments to measure the effectiveness of KATE compared to other models on tasks such as classification, multi-label classification (MLC), regression, and document retrieval (DR). For classification, MLC and regression tasks, we train a simple neural network that uses the encoded test inputs as feature vectors, and directly maps them to the output classes or values. A simple softmax classifier with cross-entropy loss was applied for the classification task, and multi-label logistic regression classifier with cross-entropy loss was applied for the MLC task. For the regression task we used a two-layer neural regression model (where the output layer is a sigmoid neuron) with squared error loss. The same architecture is used for all methods to ensure fairness. Note that when comparing various methods, the same number of features were learned for all of them except for Word2Vec_{pre} which uses 300-dimensional pre-trained word embeddings and thus its number of features was fixed as 300 in all experiments.

4.4.1 Mean Squared Cosine Deviation among Topics. We first quantify how distinct are the topics learned via different methods. Let \mathbf{v}_i denote the vector representation of topic i , and let there be m topics. The cosine of the angle between \mathbf{v}_i and \mathbf{v}_j , given as $\cos(\mathbf{v}_i, \mathbf{v}_j) = \frac{\mathbf{v}_i^T \mathbf{v}_j}{\|\mathbf{v}_i\| \|\mathbf{v}_j\|}$, is a measure of how similar/correlated the two topic vectors are; it takes values in the range $[-1, 1]$. The topics are most dissimilar when the vectors are orthogonal to each other, i.e., with the angle between them is $\pi/2$, with the cosine of the angle being zero. Define the pair-wise mean squared cosine deviation among m topics as follows

$$MSCD = \sqrt{\frac{2}{m(m-1)} \sum_{i,j>i} \cos^2(\mathbf{v}_i, \mathbf{v}_j)}$$

Thus, $MSCD \in [0, 1]$, and smaller values of $MSCD$ (closer to zero) imply more distinctive, i.e., orthogonal topics.

Model	weapon	christian	compani	israel	law	hockey	comput	space
AE	effort muslim sort america escap	close test larg result answer	hold simpl serv commit societi	cost isra arab fear occupi	made live give power reason	plane tie sex english intel	inform run program base author	studi data answer origin unit
KSAE	qualiti challeng tire 7u learn	god power lie logic simpl	commit student age hold consist	occupi enhanc azeri rpi sleep	back govern reason call answer	int cco monash rsa pasadena	run inform part case start	data process answer heard version
KATE	arm crime gun firearm handgun	belief god believ faith bibl	market dealer manufactur expens cost	arab isra palestinian occupi jew	citizen constitut court feder govern	playoff nhl team wing coach	scienc dept cs math univ	launch orbit mission shuttl flight
Word2Vec	assault militia possess automat gun	understand belief believ god truth	insur feder manufactur industri pay	lebanon isra lebanes arab palestinian	court prohibit ban sentenc legitim	sport nhl playoff winner cup	engin colleg umich subject perform	launch jpl nasa moon gov

Table 3: Five nearest neighbors in the word representation space on 20 Newsgroups dataset.




Figure 2: PCA on the 20-dimensional document vectors from 20 Newsgroups dataset.




Figure 3: T-SNE on the 20-dimensional document vectors from 20 Newsgroups dataset.

We evaluate MSCD for topics generated by AE, KSAE, LDA, and KATE. We also evaluate KATE without amplification. In LDA, a topic is represented as its probabilistic distribution over the vocabulary set, whereas for autoencoders, it is defined as the weights on the connections between the corresponding hidden neuron and all the input neurons. We conduct experiments on the 20 Newsgroups

dataset and vary the number of topics from 20 to 128 and 512. Table 4 shows these results. We find that KATE has the lowest MSCD values, which means that it can learn more distinctive (i.e., orthogonal) topics than other methods. Our results are much better than LDA, since the latter does not prevent topics from being similar. On the other hand, the competition in KATE drives topics (i.e., the hidden neurons) to become distinct from each other. Interestingly,

Model	20 News		
	20	128	512
AE	0.976	0.722	0.319
KSAE	0.268	0.198	0.056
LDA	0.249	0.059	0.028
KATE <i>no_amp.</i>	0.154	0.069	0.037
KATE	0.097	0.024	0.014

Table 4: Mean squared cosine deviation among topics; smaller means more distinctive topics.

KATE with amplification (i.e., here we have $\alpha = 6.26$) consistently achieves lower MSCD values than KATE without amplification, which verifies the effectiveness of the α amplification connections in terms of learning distinctive topics.

Model	20 News	
	128	512
LDA	0.657	0.685
DBN	0.677	0.705
DocNADE	0.714	0.735
NVDM	0.052	0.053
Word2Vec _{pre}	0.687	0.687
Word2Vec	0.564	0.586
Doc2Vec	0.347	0.399
AE	0.084	0.516
DAE	0.125	0.291
CAE	0.083	0.512
VAE	0.724	0.747
KSAE	0.486	0.675
KATE	0.744	0.761

Table 5: Classification accuracy on 20 Newsgroups dataset.

4.4.2 Document Classification Task. In this set of experiments, we evaluate the quality of learned document representations from various models for the purpose of document classification. Table 5 shows the classification accuracy results on the 20 Newsgroups dataset (using 128 topics). Traditional autoencoders (including AE, DAE, CAE) do not perform well on this task. We observed that the validation set error was oscillating when training these classifiers (also observed in the regression task below), which indicates that the extracted features are not representative and consistent. KSAE consistently achieves higher accuracies than other autoencoders and does not exhibit the oscillating phenomenon, which means that adding sparsity does help learn better representations. VAE even performs better than KSAE on this dataset, which shows the advantages of VAE over other traditional autoencoders. However, as we will see later, VAE fails to consistently perform well across different datasets and tasks. Word2Vec_{pre} performs on par with DBN and LDA even though it just averages all the word embeddings in a document, which suggests the effectiveness of pre-training word embeddings on a large external corpus to learn general knowledge. Not surprisingly, DocNADE works very well on this task as also reported in previous work [5, 18, 36]. Our KATE model significantly outperforms all other models. For example, KATE obtains 74.4% accuracy which is significantly higher than the 72.4% accuracy achieved by VAE.

Table 6 shows multi-label classification results on Reuters and Wiki10+ datasets. Here we show both the Macro-F1 and Micro-F1

scores (reflecting a balance of precision and recall) for different number of features. Micro-F1 score biases the metric towards the most populated labels, while Macro-F1 biases the metric towards the least populated labels. Both Reuters and Wiki10+ are highly imbalanced. For example in Wiki10+, the documents belonging to ‘wikipedia’ or ‘wiki’ account for 90% of the corpus while only around 6% of the documents are relevant to ‘religion’. Similarly, in Reuters, the documents belonging to ‘CCAT’ account for 47% of the corpus while there are only 5 documents relevant to ‘GMIL’. DocNADE works the very well on this task, but the sparse and competitive autoencoders also perform well. KATE outperforms KSAE on Reuters and remains competitive on Wiki10+. We don’t report the results of DBN on Reuters since the training did not end even after a long time.

4.4.3 Regression Task. In this set of experiments, we evaluate the quality of learned document representations from various models for predicting the movie ratings in the MRD dataset, as shown in Table 7 (using 128 features). The coefficient of determination, denoted r^2 , from the regression model was used to evaluate the methods. The best possible r^2 statistic value is 1.0; negative values are also possible, indicating a poor fit of the model to the data. In general, other autoencoder models perform poorly on this task, for example, AE even gets a negative r^2 score. Interestingly, Word2Vec_{pre} performs on par with DocNADE, indicating that word embeddings learned from a large external corpus can capture some semantics of emotive words (e.g., good, bad, wonderful). We observe that KATE significantly outperforms all other models, including Word2Vec_{pre}, which means it can learn meaningful representations which are helpful for sentiment analysis.




Figure 4: Document retrieval on 20 Newsgroups dataset (128 features).

4.4.4 Document Retrieval Task. We also evaluate the various models for document retrieval. Each document in the test set is used as an individual query and we fetch the relevant documents from the training set based on the cosine similarity between the document representations. The average fraction of retrieved documents which share the same label as the query document, i.e., precision, was used as the evaluation metric. As shown in Figure 4, VAE performs the best on this task followed by DocNADE and KATE. Among the other models, DBN and LDA also have decent performance, but the other autoencoders are not that effective.

Model	Reuters				Wiki10+			
	128		512		128		512	
	Macro-F1	Micro-F1	Macro-F1	Micro-F1	Macro-F1	Micro-F1	Macro-F1	Micro-F1
LDA	0.408	0.703	0.576	0.766	0.442	0.584	0.305	0.441
DBN	-	-	-	-	0.330	0.513	0.339	0.536
DocNADE	0.564	0.768	0.667	0.831	0.451	0.585	0.423	0.561
NVDM	0.215	0.441	0.195	0.452	0.187	0.461	0.036	0.375
Word2Vec _{pre}	0.549	0.712	0.549	0.712	0.312	0.454	0.312	0.454
Word2Vec	0.458	0.648	0.595	0.761	0.205	0.318	0.234	0.325
Doc2Vec	0.004	0.082	0.000	0.000	0.289	0.486	0.344	0.524
AE	0.025	0.047	0.459	0.651	0.016	0.040	0.382	0.569
DAE	0.275	0.576	0.489	0.685	0.359	0.560	0.375	0.534
CAE	0.024	0.045	0.549	0.726	0.091	0.168	0.404	0.547
VAE	0.325	0.458	0.490	0.594	0.342	0.497	0.373	0.511
KSAE	0.457	0.660	0.605	0.766	0.449	0.594	0.471	0.614
KATE	0.539	0.716	0.615	0.767	0.445	0.580	0.446	0.580

Table 6: Comparison of MLC F1 score on Reuters RCV1-v2 and Wiki10+ datasets.

Model	MRD	
	128	512
LDA	0.287	0.226
DBN	0.277	0.369
DocNADE	0.404	0.424
NVDM	0.199	0.191
Word2Vec _{pre}	0.409	0.409
Word2Vec	0.143	0.136
Doc2Vec	0.052	0.032
AE	-0.001	0.203
DAE	0.067	0.100
CAE	0.018	0.118
VAE	0.111	0.355
KSAE	0.152	0.365
KATE	0.463	0.516

Table 7: Comparison of regression r^2 score on MRD dataset.

4.4.5 *Timing.* Finally, we compare the training time of various models. Results are shown in Table 8 for the 20 Newsgroups dataset, with 20 topics. Our model is much faster than deep generative models like DBN and DocNADE. It is typically slower than other autoencoders since it usually takes more epochs to converge. Nevertheless, as demonstrated above, it significantly outperforms other models in various text analytics tasks.

4.5 KATE: Effects of Parameter Tuning

Having demonstrated the effectiveness of KATE compared to other methods, we study the effects of various hyperparameter choices in KATE, such as the number of topics (i.e., hidden neurons), the number of winners k and the energy amplification parameter α . The default values for the number of topics is 128, with $k = 32$ and $\alpha = 6.26$. Note when exploring the effect of the number of topics, we also vary k to find its best match to the given number of topics. Figure 5 shows the classification accuracy on the 20 Newsgroups dataset, as we vary these parameters. We observe that as we increase the number of topics or hidden neurons (in Figure 5a), the accuracy continues to rise, but eventually drops off. We use 128 as the default value since it offers the best trade-off in complexity and performance; only relatively minor gains are achieved in

increasing the number of topics beyond 128. Considering the number of winning neurons (see Figure 5b), the main trend is that the performance degrades when we make k larger, which is expected since larger k implies lesser competition. In practice, when tuning k , we find that starting by a value close to around a quarter of the number of topics is a good strategy. Finally, as we mentioned, the α amplification connection is crucial as verified in Figure 5c. When $\alpha = 2/k = 0.0625$, which means there is no amplification for the energy, the classification accuracy is 71.1%. However, we are able to significantly boost the model performance up to 74.6% accuracy by increasing the value of α . We use a default value of $\alpha = 6.26$, which once again reflects a good trade-off across different datasets. It is also important to note that across all the experiments, we found that using the *tanh* activation function (instead of *sigmoid* function) in the k-competitive layer of KATE gave the best performance. For example, on the 20 Newsgroups data, using 128 topics, KATE with *tanh* yields 74.4% accuracy, while with *sigmoid* it was only 56.8%.

5 CONCLUSIONS

We described a novel k-competitive autoencoder, KATE, that explicitly enforces competition among the neurons in the hidden layer by selecting the k highest activation neurons as winners, and reallocates the amplified energy (aggregate activation potential) from the losers. Interestingly, even though we use a shallow model, i.e., with one hidden layer, it outperforms a variety of methods on many different text analytics tasks. More specifically, we perform a comprehensive evaluation of KATE against techniques spanning graphical models (e.g., LDA), belief networks (e.g., DBN), word embedding models (e.g., Word2Vec), and several other autoencoders including the k-sparse autoencoder (KSAE). We find that across tasks such as document classification, multi-label classification, regression and document retrieval, KATE clearly outperforms competing methods or obtains close to the best results. It is very encouraging to note that KATE is also able to learn semantically meaningful representations of words, documents and topics, which we evaluated via both quantitative and qualitative studies. As part of future work, we plan to evaluate KATE on more domain specific datasets, such as bibliographic networks, for example for topic induction and scientific publication retrieval. We also plan to improve the scalability

Model	LDA	DBN	DocNADE	NVDM	Word2Vec	Doc2Vec	AE	DAE	CAE	VAE	KSAE	KATE
Time (s)	399	15,281	4,787	645	977	992	566	361	729	660	489	1,214

Table 8: Training time of various models (in seconds).




Figure 5: Effects of hyper-parameters.

and effectiveness of our approach on much larger text collections by developing parallel and distributed implementations.

ACKNOWLEDGMENTS

This work was supported in part by NSF awards IIS-1302231 and IIS-1738895.

REFERENCES

- [1] David M Blei, Thomas L Griffiths, and Michael I Jordan. 2010. The nested chinese restaurant process and bayesian nonparametric inference of topic hierarchies. *Journal of the ACM (JACM)* 57, 2 (2010), 7.
- [2] David M Blei, Andrew Y Ng, and Michael I Jordan. 2003. Latent dirichlet allocation. *Journal of Machine Learning Research* 3, Jan (2003), 993–1022.
- [3] Phil Blunsom, Edward Grefenstette, and Nal Kalchbrenner. 2014. A convolutional neural network for modelling sentences. In *ACL Meeting*.
- [4] Kevin Robert Canini, Lei Shi, and Thomas L Griffiths. 2009. Online Inference of Topics with Latent Dirichlet Allocation. In *AISTATS*, Vol. 9. 65–72.
- [5] Ziqiang Cao, Sujian Li, Yang Liu, Wenjie Li, and Heng Ji. 2015. A Novel Neural Topic Model and Its Supervised Extension. In *AAAI*. 2210–2216.
- [6] Rajarshi Das, Manzil Zaheer, and Chris Dyer. 2015. Gaussian lda for topic models with word embeddings. In *ACL Meeting*. 795–804.
- [7] Scott Deerwester, Susan T Dumais, George W Furnas, Thomas K Landauer, and Richard Harshman. 1990. Indexing by latent semantic analysis. *Journal of the American Society for Information Science* 41, 6 (1990), 391.
- [8] Jacob Eisenstein, Amr Ahmed, and Eric P Xing. 2011. Sparse additive generative models of text. In *ICML*.
- [9] Geoffrey E Hinton, Simon Osindero, and Yee-Whye Teh. 2006. A fast learning algorithm for deep belief nets. *Neural Computation* 18, 7 (2006), 1527–1554.
- [10] Geoffrey E Hinton and Ruslan R Salakhutdinov. 2006. Reducing the dimensionality of data with neural networks. *Science* 313, 5786 (2006), 504–507.
- [11] Geoffrey E Hinton and Ruslan R Salakhutdinov. 2009. Replicated softmax: an undirected topic model. In *NIPS*. 1607–1614.
- [12] Thomas Hofmann. 1999. Probabilistic latent semantic analysis. In *UAI Conf*. 289–296.
- [13] Yoon Kim. 2014. Convolutional neural networks for sentence classification. In *EMNLP*.
- [14] Diederik P Kingma and Max Welling. 2014. Auto-encoding variational bayes. In *ICLR*.
- [15] Alex Krizhevsky. 2009. The CIFAR-10 dataset. www.cs.toronto.edu/~kriz/cifar.html. (2009).
- [16] Girish Kumar and Luis F. D'Haro. 2015. Deep Autoencoder Topic Model for Short Texts. In *International Workshop on Embeddings and Semantics*.
- [17] Ken Lang. 1995. Newsweeder: Learning to filter netnews. In *ICML*. 331–339.
- [18] Hugo Larochelle and Stanislas Lauly. 2012. A neural autoregressive topic model. In *NIPS*. 2708–2716.
- [19] Quoc V Le and Tomas Mikolov. 2014. Distributed Representations of Sentences and Documents. In *ICML*, Vol. 14. 1188–1196.
- [20] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. 1998. Gradient-based learning applied to document recognition. *Proc. IEEE* 86, 11 (1998), 2278–2324.
- [21] David D Lewis, Yiming Yang, Tony G Rose, and Fan Li. 2004. Rcv1: A new benchmark collection for text categorization research. *Journal of Machine Learning Research* 5, Apr (2004), 361–397.
- [22] Lars Maaløe, Morten Arngren, and Ole Winther. 2014. Deep belief nets for topic modeling. In *Workshop on Knowledge-Powered Deep Learning for Text Mining*.
- [23] Laurens van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-SNE. *Journal of Machine Learning Research* 9, Nov (2008), 2579–2605.
- [24] Alireza Makhzani and Brendan Frey. 2014. k-Sparse Autoencoders. In *ICLR*.
- [25] Alireza Makhzani and Brendan J Frey. 2015. Winner-take-all autoencoders. In *NIPS*. 2791–2799.
- [26] Alireza Makhzani, Jonathon Shlens, Navdeep Jaitly, and Ian Goodfellow. 2016. Adversarial autoencoders. In *ICLR*.
- [27] Yishu Miao, Lei Yu, and Phil Blunsom. 2016. Neural Variational Inference for Text Processing. In *ICML*.
- [28] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *NIPS*. 3111–3119.
- [29] Dat Quoc Nguyen, Richard Billingsley, Lan Du, and Mark Johnson. 2015. Improving Topic Models with Latent Feature Word Representations. *Transactions of the Association for Computational Linguistics* 3 (2015), 299–313.
- [30] Bo Pang and Lillian Lee. 2004. A sentimental education: Sentiment analysis using subjectivity summarization based on minimum cuts. In *ACL Meeting*. 271.
- [31] Bo Pang and Lillian Lee. 2005. Seeing stars: Exploiting class relationships for sentiment categorization with respect to rating scales. In *ACL Meeting*. 115–124.
- [32] Bo Pang, Lillian Lee, and Shivakumar Vaithyanathan. 2002. Thumbs up?: sentiment classification using machine learning techniques. In *EMNLP*. 79–86.
- [33] Jeffrey Pennington, Richard Socher, and Christopher D Manning. 2014. Glove: Global Vectors for Word Representation. In *EMNLP*, Vol. 14. 1532–1543.
- [34] Radim Rehurek and Petr Sojka. 2010. Software Framework for Topic Modelling with Large Corpora. In *LREC 2010 Workshop on New Challenges for NLP Frameworks*. 45–50.
- [35] Salah Rifai, Pascal Vincent, Xavier Muller, Xavier Glorot, and Yoshua Bengio. 2011. Contractive auto-encoders: Explicit invariance during feature extraction. In *ICML*. 833–840.
- [36] Nitish Srivastava, Ruslan R Salakhutdinov, and Geoffrey E Hinton. 2013. Modeling documents with deep boltzmann machines. In *UAI Conf*.
- [37] Yee Whye Teh, Michael I Jordan, Matthew J Beal, and David M Blei. 2006. Hierarchical dirichlet processes. *J. Amer. Statist. Assoc.* 101, 476 (2006), 1566–1581.
- [38] Yee W Teh, Kenichi Kurihara, and Max Welling. 2007. Collapsed variational inference for HDP. In *NIPS*. 1481–1488.
- [39] Pascal Vincent, Hugo Larochelle, Isabelle Lajoie, Yoshua Bengio, and Pierre-Antoine Manzagol. 2010. Stacked denoising autoencoders: Learning useful representations in a deep network with a local denoising criterion. *Journal of Machine Learning Research* 11, Dec (2010), 3371–3408.
- [40] Chong Wang and David M Blei. 2009. Decoupling sparsity and smoothness in the discrete hierarchical dirichlet process. In *NIPS*. 1982–1989.
- [41] Matthew D Zeiler. 2012. ADADELTA: an adaptive learning rate method. *arXiv* 1212.5701 (2012).
- [42] Shuangfei Zhai and Zhongfei Zhang. 2016. Semisupervised Autoencoder for Sentiment Analysis. In *AAAI*.
- [43] Jun Zhu and Eric P Xing. 2011. Sparse topical coding. In *UAI Conf*.
- [44] Arkaitz Zubiaga. 2009. Enhancing Navigation on Wikipedia with Social Tags. In *5th International Conference of the Wikimedia Community*.