
Candidates vs. Noises Estimation for Large Multi-Class Classification Problem

Lei Han¹ Yiheng Huang¹ Tong Zhang¹

Abstract

This paper proposes a method for multi-class classification problems, where the number of classes K is large. The method, referred to as *Candidates vs. Noises Estimation* (CANE), selects a small subset of candidate classes and samples the remaining classes. We show that CANE is always consistent and computationally efficient. Moreover, the resulting estimator has low statistical variance approaching that of the maximum likelihood estimator, when the observed label belongs to the selected candidates with high probability. In practice, we use a tree structure with leaves as classes to promote fast beam search for candidate selection. We further apply the CANE method to estimate word probabilities in learning large neural language models. Extensive experimental results show that CANE achieves better prediction accuracy over the Noise-Contrastive Estimation (NCE), its variants and a number of the state-of-the-art tree classifiers, while it gains significant speedup compared to standard $\mathcal{O}(K)$ methods.

1. Introduction

In practice one often encounters multi-class classification problem with a large number of classes. For example, applications in image classification (Russakovsky et al., 2015) and language modeling (Mikolov et al., 2010) usually have tens to hundreds of thousands of classes. Under such cases, training the standard softmax logistic or one-against-all models becomes impractical.

One promising way to handle the large class size is to use sampling. In language models, a commonly adopted technique is *Noise-Contrastive Estimation* (NCE) (Gutmann & Hyvärinen, 2012). This method is originally proposed

for estimating probability densities and has been applied to various language modeling situations, such as learning word embeddings, context generation and neural machine translation (Mnih & Teh, 2012; Mnih & Kavukcuoglu, 2013; Vaswani et al., 2013; Sordani et al., 2015). NCE reduces the problem of multi-class classification to binary classification problem, which discriminates between a target class distribution and a noise distribution and a few noise classes are sampled as a representation of the entire noise space. In general, the noise distribution is given a priori. For example, a power-raised unigram distribution has been shown to be effective in language models (Mikolov et al., 2013; Ji et al., 2015; Mnih & Teh, 2012). Recently, some variants of NCE have been proposed. The *Negative Sampling* (Mikolov et al., 2013) is a simplified version of NCE that ignores the numerical probabilities in the distributions and discriminates between only the target class and noise samples; the *One vs. Each* (Titsias, 2016) solves a very similar problem motivated by bounding the softmax logistic log-likelihood. Two other variants, *BlackOut* (Ji et al., 2015) and *complementary sum sampling* (Botev et al., 2017), employ parametric forms of the noise distribution and use sampled noises to approximate the normalization factor. In summary, NCE and its variants use (only) the observed class versus the noises; by sampling the noises, these methods avoid the costly computation of the normalization factor to achieve fast training speed. In this paper, we will generalize the idea by using a subset of classes (which can be automatically learned), called candidate classes, against the remaining noise classes. Compared to NCE, this approach can significantly improve the statistical efficiency when the true class belongs to the candidate classes with high probability.

Another type of popular methods for large class space is the tree structured classifier (Beygelzimer et al., 2009; Bengio et al., 2010; Deng et al., 2011; Choromanska & Langford, 2015; Daume III et al., 2017; Jernite et al., 2017). In these methods, a tree structure is defined over the classes which are treated as leaves. Each internal node of the tree is assigned with a local classifier, routing the examples to one of its descendants. Decisions are made from the root until reaching a leaf. Then, the multi-class classification problem is reduced to solving a number of small local models defined by a tree, which typically admits a logarithmic complexity on the total number of classes. Generally, tree classifiers

¹Tencent AI Lab, Shenzhen, China. Correspondence to: Lei Han <lshan@tencent.com>, Yiheng Huang <arnold-huang@tencent.com>, Tong Zhang <tongzhang@tongzhang-ml.org>.

gain training and prediction speed while suffering a loss of accuracy. The performance of tree classifier may rely heavily on the quality of the tree (Mnih & Hinton, 2009). Earlier approaches use fixed tree, such as the Filter Tree (Beygelzimer et al., 2009) and the Hierarchical Softmax (HSM) (Morin & Bengio, 2005). Recent methods are able to adjust the tree and learn the local classifiers simultaneously, such as the LOMTree (Choromanska & Langford, 2015) and Recall Tree (Daume III et al., 2017). Our approach is complementary to these tree classifiers, because we study the orthogonal issue of consistent class sampling, which in principle can be combined with many of these tree methods. In fact, a tree structure will be used in our approach to select a small subset of candidate classes. Since we focus on the class sampling aspect, we do not necessarily employ the best tree construction method in our experiments.

In this paper, we propose a method to efficiently deal with the large class problem by paying attention to a small subset of candidate classes instead of the entire class space. Given a data point \mathbf{x} (without observing y), we select a small number of competitive candidates as a set \mathcal{C}_x . Then, we sample the remaining classes, which are treated as noises, to represent the entire noise space in the large normalization factor. The estimation is referred to as *Candidates vs. Noises Estimation* (CANE). We show that CANE is consistent and its computation using stochastic gradient method is independent of the class size K . Moreover, the statistical variance of the CANE estimator can approach that of the maximum likelihood estimator (MLE) of the softmax logistic regression when \mathcal{C}_x can cover the target class y with high probability. This statistical efficiency is a key advantage of CANE over NCE, and its effect can be observed in practice.

We then describe two concrete algorithms: the first one is a generic stochastic optimization procedure for CANE; the second one employs a tree structure with leaves as classes to enable fast beam search for candidate selection. We also apply CANE to solve the word probability estimation problem in neural language modeling. Experimental results conducted on both classification and neural language modeling problems show that CANE achieves significant speedup compared to the standard softmax logistic regression. Moreover, it achieves superior performance over NCE, its variants, and a number of the state-of-the-art tree classifiers.

2. Candidates vs. Noises Estimation

Consider a K -class classification problem (K is large) with n training examples $(\mathbf{x}_i, y_i)_{i=1}^n$, where \mathbf{x}_i is from an input space \mathcal{X} and $y_i \in \{1, \dots, K\}$. The softmax logistic regression solves

$$\max_{\theta} \frac{1}{n} \sum_{i=1}^n \sum_{k=1}^K \mathbb{I}(y_i = k) \log \frac{e^{s_k(\mathbf{x}_i, \theta)}}{\sum_{k'=1}^K e^{s_{k'}(\mathbf{x}_i, \theta)}}, \quad (1)$$

where $s_k(\mathbf{x}, \theta)$ for $k = 1, \dots, K$ is a model parameterized by θ . Solving Eq. (1) requires computing a score for every class and the summation in the normalization factor, which is very expensive when K is large.

Generally speaking, given \mathbf{x} , only a small number of classes in the entire class space might be competitive to the true class. Therefore, we propose to find a small subset of classes as a candidate set $\mathcal{C}_x \subset \{1, \dots, K\}$ and treat the classes outside \mathcal{C}_x as noises, so that we can focus on the small set \mathcal{C}_x instead of the entire K classes. We will discuss one way to choose \mathcal{C}_x in Section 4. Denote the remaining $K - |\mathcal{C}_x|$ noises as a set \mathcal{N}_x , so \mathcal{N}_x is the complementary set of \mathcal{C}_x . We propose to sample some noise class $j \in \mathcal{N}_x$ to represent the entire \mathcal{N}_x . That is, we replace the partial summation $\sum_{j \in \mathcal{N}_x} e^{s_j(\mathbf{x}, \theta)}$ in the denominator of Eq. (1) by $e^{s_j(\mathbf{x}, \theta)}/q_x(j)$ using some sampled class j with an *arbitrary sampling probability* $q_x(j)$, where $q_x(j) \in (0, 1)$ and $\sum_{j \in \mathcal{N}_x} q_x(j) = 1$. Thus, the denominator $\sum_{k'=1}^K e^{s_{k'}(\mathbf{x}, \theta)}$ will be approximated as $\sum_{k' \in \mathcal{C}_x} e^{s_{k'}(\mathbf{x}, \theta)} + e^{s_j(\mathbf{x}, \theta)}/q_x(j)$. Given example (\mathbf{x}, y) and its candidate set \mathcal{C}_x , if $y \in \mathcal{C}_x$, then for some sampled noise class j , we will focus on maximizing the approximated probability

$$\frac{e^{s_y(\mathbf{x}, \theta)}}{\sum_{k' \in \mathcal{C}_x} e^{s_{k'}(\mathbf{x}, \theta)} + e^{s_j(\mathbf{x}, \theta)}/q_x(j)}; \quad (2)$$

otherwise, if $y \notin \mathcal{C}_x$, we maximize

$$\frac{e^{s_y(\mathbf{x}, \theta)}}{\sum_{k' \in \mathcal{C}_x} e^{s_{k'}(\mathbf{x}, \theta)} + e^{s_y(\mathbf{x}, \theta)}/q_x(y)} \quad (3)$$

alternatively, where y is treated as the sampled noise in place. Now, with Eqs. (2) and (3), in expectation, we will need to solve the following objective:

$$\begin{aligned} & \text{maximize } R(\theta) = \\ & \mathbb{E}_{\mathbf{x}} \left[\sum_{k \in \mathcal{C}_x} p(y = k | \mathbf{x}) \sum_{j \in \mathcal{N}_x} q_x(j) \log \frac{e^{s_k(\mathbf{x}, \theta)}}{\sum_{k' \in \mathcal{C}_x} e^{s_{k'}(\mathbf{x}, \theta)} + \frac{e^{s_j(\mathbf{x}, \theta)}}{q_x(j)}} \right. \\ & \left. + \sum_{k \in \mathcal{N}_x} p(y = k | \mathbf{x}) \log \frac{e^{s_k(\mathbf{x}, \theta)}}{\sum_{k' \in \mathcal{C}_x} e^{s_{k'}(\mathbf{x}, \theta)} + \frac{e^{s_k(\mathbf{x}, \theta)}}{q_x(k)}} \right], \quad (4) \end{aligned}$$

and empirically, we will need to solve

$$\begin{aligned} & \text{maximize } \hat{R}_n(\theta) = \\ & \frac{1}{n} \sum_{i=1}^n \left[\mathbb{I}(y_i \in \mathcal{C}_{\mathbf{x}_i}) \sum_{j \in \mathcal{N}_{\mathbf{x}_i}} q_{\mathbf{x}_i}(j) \log \frac{e^{s_{y_i}(\mathbf{x}_i, \theta)}}{\sum_{k' \in \mathcal{C}_{\mathbf{x}_i}} e^{s_{k'}(\mathbf{x}_i, \theta)} + \frac{e^{s_j(\mathbf{x}_i, \theta)}}{q_{\mathbf{x}_i}(j)}} \right. \\ & \left. + \mathbb{I}(y_i \notin \mathcal{C}_{\mathbf{x}_i}) \log \frac{e^{s_{y_i}(\mathbf{x}_i, \theta)}}{\sum_{k' \in \mathcal{C}_{\mathbf{x}_i}} e^{s_{k'}(\mathbf{x}_i, \theta)} + \frac{e^{s_{y_i}(\mathbf{x}_i, \theta)}}{q_{\mathbf{x}_i}(y_i)}} \right]. \quad (5) \end{aligned}$$

Eq. (5) consists of two summations over both the data points and the classes in the noise set \mathcal{N}_x . Therefore, we can employ a ‘doubly’ stochastic gradient optimization method by

sampling both data points $i \in \{1, \dots, n\}$ and noise classes $j \in \mathcal{N}_{\mathbf{x}_i}$. It is not difficult to check that each stochastic gradient is bounded under reasonable conditions, which means that the computational cost for solving (5) using stochastic gradient is independent of the class number K . Since we only choose a small number of candidates in $\mathcal{C}_{\mathbf{x}}$, the computation for each stochastic gradient in Eq. (5) is efficient. The above method is referred to as Candidates vs. Noises Estimation (CANE).

3. Properties

In this section, we investigate the statistical properties of CANE. The parameter space of the softmax logistic model in Eq. (1) has redundancy, observing that adding any function $h(\mathbf{x})$ to $s_k(\mathbf{x}, \boldsymbol{\theta})$ for $k = 1, \dots, K$ will not change the objective. Similar situation happens for Eqs. (4) and (5). To avoid this redundancy, one can add some constraints on the K scores or simply fix one of them as zero, e.g., let $s_K(\mathbf{x}, \boldsymbol{\theta}) = 0$. To facilitate the analysis, we will fix $s_K(\mathbf{x}, \boldsymbol{\theta}) = 0$ and consider $\mathcal{C}_{\mathbf{x}} \cup \mathcal{N}_{\mathbf{x}} = \{1, \dots, K-1\}$ within this section. First, we have the following result.

Theorem 1 (Infinity-Sample Consistency). *By viewing the objective R as a function of $\{s_1, \dots, s_{K-1}\}$, R achieves its maximum if and only if $s_k = \log \frac{p(y=k|\mathbf{x})}{p(y=K|\mathbf{x})}$ for $k = 1, \dots, K-1$.*

In Theorem 1, the global optima is exactly the log-odds function with class K as the reference class. Now, considering the parametric form $s_k(\mathbf{x}, \boldsymbol{\theta})$, there exists a true parameter $\boldsymbol{\theta}^*$ so that $s_k(\mathbf{x}, \boldsymbol{\theta}^*) = \log \frac{p(y=k|\mathbf{x})}{p(y=K|\mathbf{x})}$ if the model $s_k(\mathbf{x}, \boldsymbol{\theta})$ is correctly specified. The following theorem shows that the CANE estimator $\hat{\boldsymbol{\theta}} = \arg \max_{\boldsymbol{\theta}} \hat{R}_n(\boldsymbol{\theta})$ is consistent with the true parameter $\boldsymbol{\theta}^*$.

Theorem 2 (Finite-Sample Asymptotic Consistency). *Given \mathbf{x} , denote $\mathcal{C}_{\mathbf{x}}$ as $\{i_1, \dots, i_{|\mathcal{C}_{\mathbf{x}}|}\}$ and $\mathcal{N}_{\mathbf{x}}$ as $\{j_1, \dots, j_{|\mathcal{N}_{\mathbf{x}}|}\}$. Suppose that the parameter space is compact and $\forall \boldsymbol{\theta} \neq \boldsymbol{\theta}^*$ such that $\mathbb{P}_{\mathcal{X}}(s_k(\mathbf{x}, \boldsymbol{\theta}) \neq s_k(\mathbf{x}, \boldsymbol{\theta}^*)) > 0$ for $\mathbf{x} \sim \mathcal{X}$, $k \neq K$. Assume $\|\nabla_{\boldsymbol{\theta}} s_k(\mathbf{x}, \boldsymbol{\theta})\|$, $\|\nabla_{\boldsymbol{\theta}}^2 s_k(\mathbf{x}, \boldsymbol{\theta})\|$ and $\|\nabla_{\boldsymbol{\theta}}^3 s_k(\mathbf{x}, \boldsymbol{\theta})\|$ for $k \neq K$ are bounded under some norm $\|\cdot\|$ defined on the parameter space of $\boldsymbol{\theta}$. Then, as $n \rightarrow \infty$, the estimator $\hat{\boldsymbol{\theta}}$ converges to $\boldsymbol{\theta}^*$.*

The above theorem shows that similar to the maximum likelihood estimator of Eq. (1), the CANE estimator in Eq. (5) is also consistent. Next, we have the asymptotic normality for $\hat{\boldsymbol{\theta}}$ as follows.

Theorem 3 (Asymptotic Normality). *Under the same assumption used in Theorem 2, as $n \rightarrow \infty$, $\sqrt{n}(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*)$ follows the asymptotic normal distribution:*

$$\sqrt{n}(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*) \xrightarrow{d} \mathbb{N}(\mathbf{0}, [\mathbb{E}_{\mathbf{x}} \nabla \mathbf{M} \nabla^{\top}]^{-1}), \quad (6)$$

where

$$\begin{aligned} \mathbf{M} &= \sum_{j \in \mathcal{N}_{\mathbf{x}}} q_{\mathbf{x}}(j) \left[\text{diag}(\mathbf{u}_j) - \frac{1}{p(K, \mathbf{x}) + \sum_{k \in \mathcal{C}_{\mathbf{x}}} p(k, \mathbf{x}) + \frac{p(j, \mathbf{x})}{q_{\mathbf{x}}(j)}} \mathbf{u}_j \mathbf{u}_j^{\top} \right], \\ \mathbf{u}_j &= \left(\underbrace{p(i_1, \mathbf{x}), \dots, p(i_{|\mathcal{C}_{\mathbf{x}}|}, \mathbf{x})}_{\text{The candidate part}}, 0, \dots, \underbrace{p(j, \mathbf{x})/q_{\mathbf{x}}(j)}_{\text{The noise part}}, \dots, 0 \right)^{\top}, \\ &\text{for } j = j_1, \dots, j_{|\mathcal{N}_{\mathbf{x}}|}, \\ \nabla &= \text{diag}([\nabla_{\boldsymbol{\theta}} s_{i_1}(\mathbf{x}, \boldsymbol{\theta}), \dots, \nabla_{\boldsymbol{\theta}} s_{i_{|\mathcal{C}_{\mathbf{x}}|}}(\mathbf{x}, \boldsymbol{\theta}), \nabla_{\boldsymbol{\theta}} s_{j_1}(\mathbf{x}, \boldsymbol{\theta}), \\ &\quad \dots, \nabla_{\boldsymbol{\theta}} s_{j_{|\mathcal{N}_{\mathbf{x}}|}}(\mathbf{x}, \boldsymbol{\theta})]^{\top}). \end{aligned}$$

Theorem 3 shows that the CANE method has a statistical variance of $[\mathbb{E}_{\mathbf{x}} \nabla \mathbf{M} \nabla^{\top}]^{-1}$. As we will see in the next corollary, if one can successfully choose the candidate set $\mathcal{C}_{\mathbf{x}}$ so that it covers the observed label y with high probability, then the difference between the statistical variance of CANE and that of Eq. (1) is small. Therefore, choosing a good candidate set can be important for practical applications. Moreover, under standard conditions, the computation of CANE using stochastic gradient is independent of the class size K because the variance of stochastic gradient is bounded.

Corollary 1 (Low Statistical Variance). *The variance of the maximum likelihood estimator for the softmax logistic regression in Eq. (1) has the form $[\mathbb{E}_{\mathbf{x}} \nabla \mathbf{M}^{mle} \nabla^{\top}]^{-1}$. If $\sum_{k \in \mathcal{C}_{\mathbf{x}} \cup \{K\}} p(k, \mathbf{x}) \rightarrow 1$, i.e., the probability that $\mathcal{C}_{\mathbf{x}} \cup \{K\}$ covers the observed class label y approaches 1, then*

$$[\mathbb{E}_{\mathbf{x}} \nabla \mathbf{M} \nabla^{\top}]^{-1} \rightarrow [\mathbb{E}_{\mathbf{x}} \nabla \mathbf{M}^{mle} \nabla^{\top}]^{-1}.$$

4. Algorithm

In this section, we propose two algorithms. The first one is a general optimization procedure for CANE. The second implementation provides an efficient way to select a competitive set $\mathcal{C}_{\mathbf{x}}$ using a tree structure defined on the classes.

4.1. A General Optimization Algorithm

Eq. (5) suggests an efficient algorithm using a ‘doubly’ stochastic gradient descend (SGD) method by sampling both the data points and classes. That is, by sampling a data point (\mathbf{x}, y) , we find the candidate set $\mathcal{C}_{\mathbf{x}} \subset \{1, \dots, K\}$. If $y \in \mathcal{C}_{\mathbf{x}}$, we sample N_n noises from $\mathcal{N}_{\mathbf{x}}$ according to $q_{\mathbf{x}}$ and denote the selected noises as a set $T_{\mathbf{x}}$ ($|T_{\mathbf{x}}| = N_n$). We then optimize

$$\frac{1}{|T_{\mathbf{x}}|} \sum_{j \in T_{\mathbf{x}}} \log \frac{e^{s_y(\mathbf{x}, \boldsymbol{\theta})}}{\sum_{k' \in \mathcal{C}_{\mathbf{x}}} e^{s_{k'}(\mathbf{x}, \boldsymbol{\theta})} + e^{s_j(\mathbf{x}, \boldsymbol{\theta})/q_{\mathbf{x}}(j)}},$$

with gradient $\nabla_{\boldsymbol{\theta}} \hat{R}$ given by Eq. (7). Otherwise, if $y \notin \mathcal{C}_{\mathbf{x}}$, we optimize

$$\log \frac{e^{s_y(\mathbf{x}, \boldsymbol{\theta})}}{\sum_{k' \in \mathcal{C}_{\mathbf{x}}} e^{s_{k'}(\mathbf{x}, \boldsymbol{\theta})} + e^{s_y(\mathbf{x}, \boldsymbol{\theta})/q_{\mathbf{x}}(y)}},$$

Algorithm 1 A general optimization procedure for CANE.

- 1: **Input:** K , $(\mathbf{x}_i, y_i)_{i=1}^n$, number of candidates $N_c = |\mathcal{C}_x|$, number of sampled noises $N_n = |T_x|$, sampling strategy \mathbf{q} and learning rate η .
- 2: **Output:** $\hat{\theta}$.
- 3: Initialize θ ;
- 4: **for** every sampled example **do**
- 5: Receive example (\mathbf{x}, y) ;
- 6: Find the candidate set \mathcal{C}_x ;
- 7: **if** $y \in \mathcal{C}_x$ **then**
- 8: Sample N_n noises outside \mathcal{C}_x according to \mathbf{q} and denote the selected noise set as T_x ;
- 9: $\theta \leftarrow \theta + \eta \nabla_{\theta} \hat{R}$ with $\nabla_{\theta} \hat{R}$ given by

$$\nabla_{\theta} s_y(\mathbf{x}, \theta) - \frac{1}{|T_x|} \sum_{j \in T_x} \quad (7)$$

$$\left[\frac{\sum_{k' \in \mathcal{C}_x} e^{s_{k'}(\mathbf{x}, \theta)} \nabla_{\theta} s_{k'}(\mathbf{x}, \theta) + \frac{e^{s_j(\mathbf{x}, \theta)}}{q_{\mathbf{x}}(j)} \nabla_{\theta} s_j(\mathbf{x}, \theta)}{\sum_{k' \in \mathcal{C}_x} e^{s_{k'}(\mathbf{x}, \theta)} + \frac{e^{s_j(\mathbf{x}, \theta)}}{q_{\mathbf{x}}(j)}} \right];$$

10: **else**

11: $\theta \leftarrow \theta + \eta \nabla_{\theta} \hat{R}$ with $\nabla_{\theta} \hat{R}$ given by

$$\nabla_{\theta} s_y(\mathbf{x}, \theta) - \frac{\sum_{k' \in \mathcal{C}_x} e^{s_{k'}(\mathbf{x}, \theta)} \nabla_{\theta} s_{k'}(\mathbf{x}, \theta) + \frac{e^{s_y(\mathbf{x}, \theta)}}{q_{\mathbf{x}}(y)} \nabla_{\theta} s_y(\mathbf{x}, \theta)}{\sum_{k' \in \mathcal{C}_x} e^{s_{k'}(\mathbf{x}, \theta)} + \frac{e^{s_y(\mathbf{x}, \theta)}}{q_{\mathbf{x}}(y)}}; \quad (8)$$

12: **end if**

13: **end for**

with gradient $\nabla_{\theta} \hat{R}$ given by Eq. (8). This general procedure is provided in Algorithm 1. Algorithm 1 has a complexity of $\mathcal{O}(N_c + N_n)$ (where $N_c = |\mathcal{C}_x|$), which is independent of the class size K . In step 6, any method can be used to select \mathcal{C}_x .

4.2. Beam Tree Algorithm

In the second algorithm, we provide an efficient way to find a competitive \mathcal{C}_x . An attractive strategy is to use a tree defined on the classes, because one can perform fast heuristic search algorithms based on a tree structure to prune the uncompetitive classes. Indeed, any structure, e.g., graph or groups, can be used alternatively as long as the structure allows to efficiently prune uncompetitive classes. We will use tree structure for candidate selection in this paper.

Given a tree structure defined on the K classes, the model $s_k(\mathbf{x}, \theta)$ is interpreted as a tree model illustrated in Fig. 1. For simplicity, Fig. 1 uses a binary tree over $K = 8$ labels as example while any tree structure can be used for selecting \mathcal{C}_x . In the example, circles denote internal nodes and squares indicate classes. The parameters are kept in the edges and denoted as $\theta_{(o,c)}$, where o indicates an internal node and c is the index of the c -th child of node o . Therefore, a pair (o, c) represents an edge from node o to its c -th child. The dashed circles indicate that we do not keep any

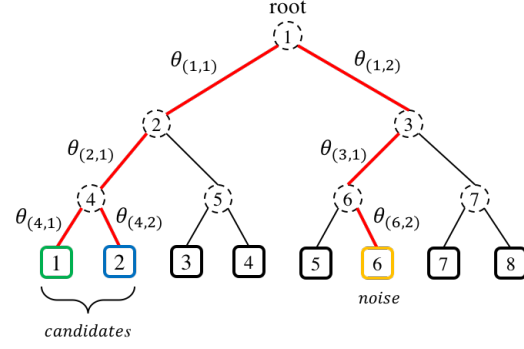


Figure 1. Illustration of the tree model. Suppose an example $(\mathbf{x}, 2)$ is arriving, and two candidate classes 1 and 2 are selected by beam search. The class 6 is sampled as noise.

parameters in the internal nodes. Now, define $s_k(\mathbf{x}, \theta)$ as

$$s_k(\mathbf{x}, \theta) = g_{\psi}(\mathbf{x}) \cdot \sum_{(o,c) \in \mathcal{P}_k} \theta_{(o,c)}, \quad (9)$$

where $g_{\psi}(\mathbf{x})$ is a function parameterized by ψ and it maps the input $\mathbf{x} \sim \mathcal{X}$ to a representation $g_{\psi}(\mathbf{x}) \in \mathbb{R}^{d_r}$ for some d_r . For example, in image classification, a good choice of the representation $g_{\psi}(\mathbf{x})$ of the raw pixels \mathbf{x} is usually a deep neural network. \mathcal{P}_k denotes the path from the root to the class k . Eq. (9) implies that the score of an example belonging to a class is calculated by summing up the scores along the corresponding path. Now, in Fig. 1, suppose that we are given an example (\mathbf{x}, y) with class $y = 2$ (blue). Using beam search, we find two candidates with high scores, i.e., class 1 (green) and class 2. Then, we let $\mathcal{C}_x = \{1, 2\}$. In this case, we have $y \in \mathcal{C}_x$, so we need to sample noises. Suppose we sample one class 6 (orange). According to Eq. (7), the parameters along the corresponding paths (red) will be updated.

Formally, given example (\mathbf{x}, y) , if $y \in \mathcal{C}_x$, we sample noises as a set T_x . Then for $(o, c) \in \mathcal{P}_{\mathcal{C}_x \cup T_x}$, where $\mathcal{P}_{\mathcal{C}_x \cup T_x} = \cup_{k \in \mathcal{C}_x \cup T_x} \mathcal{P}_k$, the gradient with respect to $\theta_{(o,c)}$ is

$$\frac{\partial \hat{R}}{\partial \theta_{(o,c)}} = \frac{1}{|T_x|} \sum_{j \in T_x} \left[\mathbb{I}((o, c) \in \mathcal{P}_y) - \frac{\sum_{k' \in \mathcal{C}_x} \mathbb{I}((o, c) \in \mathcal{P}_{k'}) e^{s_{k'}(\mathbf{x}, \theta)} + \mathbb{I}((o, c) \in \mathcal{P}_j) \frac{e^{s_j(\mathbf{x}, \theta)}}{q_{\mathbf{x}}(j)}}{\sum_{k' \in \mathcal{C}_x} e^{s_{k'}(\mathbf{x}, \theta)} + \frac{e^{s_j(\mathbf{x}, \theta)}}{q_{\mathbf{x}}(j)}} \right] g_{\psi}(\mathbf{x}). \quad (10)$$

Note that an edge may be included in multiple selected paths. For example, \mathcal{P}_1 and \mathcal{P}_2 share edges $(1, 1)$ and $(2, 1)$ in Fig. 1. The case of $y \notin \mathcal{C}_x$ can be illustrated similarly. The gradient with respect to $\theta_{(o,c)}$ when $y \notin \mathcal{C}_x$ is

$$\frac{\partial \hat{R}}{\partial \theta_{(o,c)}} = \left[\mathbb{I}((o, c) \in \mathcal{P}_y) - \frac{\sum_{k' \in \mathcal{C}_x} \mathbb{I}((o, c) \in \mathcal{P}_{k'}) e^{s_{k'}(\mathbf{x}, \theta)} + \mathbb{I}((o, c) \in \mathcal{P}_y) \frac{e^{s_y(\mathbf{x}, \theta)}}{q_{\mathbf{x}}(y)}}{\sum_{k' \in \mathcal{C}_x} e^{s_{k'}(\mathbf{x}, \theta)} + \frac{e^{s_y(\mathbf{x}, \theta)}}{q_{\mathbf{x}}(y)}} \right] g_{\psi}(\mathbf{x}). \quad (11)$$

Algorithm 2 The Beam Tree Algorithm.

-
- 1: **Input:** $K, (\mathbf{x}_i, y_i)_{i=1}^n$, representation function $g_\psi(\mathbf{x})$, number of candidates $N_c = |\mathcal{C}_x|$, number of sampled noises $N_n = |T_x|$, sampling strategy \mathbf{q} and learning rate η .
 - 2: **Output:** θ .
 - 3: Construct a tree on the K classes;
 - 4: Initialize θ ;
 - 5: **for** every sampled example **do**
 - 6: Receive example (\mathbf{x}, y) ;
 - 7: Given \mathbf{x} , use beam search to find the N_c classes with high scores to compose \mathcal{C}_x ;
 - 8: **if** $y \in \mathcal{C}_x$ **then**
 - 9: Sample N_n noises outside \mathcal{C}_x according to \mathbf{q} and denote the selected noise set as T_x ;
 - 10: Find the paths with respect to the classes in $\mathcal{C}_x \cup T_x$;
 - 11: **else**
 - 12: Find the paths with respect to the classes in $\mathcal{C}_x \cup \{y\}$;
 - 13: **end if**
 - 14: Sum up the scores along each selected path for the corresponding class;
 - 15: $\theta_{(o,c)} \leftarrow \theta_{(o,c)} + \eta \frac{\partial \hat{R}}{\partial \theta_{(o,c)}}$ for each (o, c) included in the selected paths according to Eqs. (10) and (11);
 - 16: $\psi \leftarrow \psi + \eta \frac{\partial \hat{R}}{\partial g} \frac{\partial g}{\partial \psi}$; // if g is parameterized.
 - 17: **end for**
-

The gradients in Eqs. (10) and (11) enjoy the following property.

Proposition 1. *At each iteration of Algorithm 2, if an edge (o, c) is included in every selected path, then $\theta_{(o,c)}$ does not need to be updated.*

The proof of Proposition 1 is straightforward that if (o, c) belongs to every selected path, then the gradients in Eqs. (10) and (11) are $\mathbf{0}$. The above property allows a fast detection of those parameters which do not need to be updated in SGD and hence can save computations. In practice, the number of shared edges is related to the tree structure.

Since we use beam search to choose the candidates in a tree structure, the proposed algorithm is referred to as *Beam Tree*, which is depicted in Algorithm 2.¹ For the tree construction method in step 3, we can use some hierarchical clustering based methods which will be detailed in the experiments and supplementary material. In the algorithm, the beam search needs $\mathcal{O}(N_c \log_b K)$ operations, where b is a constant related to the tree structure, e.g., binary tree for $b = 2$. The parameter updating needs $\mathcal{O}((N_c + N_n) \log_b K)$ operations. Therefore, Algorithm 2 has a complexity of $\mathcal{O}((2N_c + N_n) \log_b K)$ which is logarithmic with respect to K . The term $\log_b K$ is from the tree structure used in this specific candidate selection method, so it does not conflict with the complexity of the general Algorithm 1, which is independent of K . Another advantage of the Beam Tree

¹The beam search procedure in step 7 is provided in the supplementary material.

algorithm is that it allows fast predictions and can naturally output the top- J predictions using beam search. The prediction time has an order of $\mathcal{O}(J \log_b K)$ for the top- J predictions.

5. Application to Neural Language Modeling

In this section, we apply the CANE method to neural language modeling which solves a probability density estimation problem. In neural language models, the conditional probability distribution of the target word w given context h is defined as

$$P_h(w) = \frac{e^{s_w(h, \theta)}}{\sum_{w' \in \mathcal{V}} e^{s_{w'}(h, \theta)}},$$

where $s_w(h, \theta)$ is the scoring function with parameter θ . A word w in the context h will be represented by an embedding vector $\mathbf{u}_w \in \mathbb{R}^d$ with embedding size d . Given context h , the model computes the score for the target word w as

$$s_w(h, \theta) = g_\varphi(\mathbf{u}_h) \mathbf{v}_w,$$

where $\theta = \{\mathbf{u}, \mathbf{v}, \varphi\}$, $g_\varphi(\cdot)$ is a representation function (parameterized by φ) of the embeddings in the context h , e.g., a LSTM modular (Hochreiter & Schmidhuber, 1997), and \mathbf{v}_w is the weight parameter for the target word w . Both the word embedding \mathbf{u} and weight parameter \mathbf{v} need to be estimated. In language models, the vocabulary size $|\mathcal{V}|$ is usually very large and the computation of the normalization factor is expensive. Therefore, instead of estimating the exact probability distribution $P_h(w)$, sampling methods such as NCE and its variants (Mnih & Kavukcuoglu, 2013; Ji et al., 2015) are typically adopted to approximate $P_h(w)$.

In order to apply the CANE method, we need to select the candidates given any context h . For multi-class classification problem, we have devised a Beam Tree algorithm in Algorithm 2 that uses a tree structure to select candidates, and the tree can be obtained by some hierarchical clustering methods over x before learning. However, different from the classification problem, the word embeddings in the language model are not known before training, and thus obtaining a hierarchical structure based on the word embeddings is not practical. In this paper, we construct a simple tree with only one layer under the root, where the layer contains N subsets formed by splitting the words according to their frequencies. At each iteration of Algorithm 2, we route the example by selecting the subset with the largest score (in place of beam search) and then sample the candidates from the subset according to some distribution. For the noises in CANE, we directly sample words out of the candidate set according to \mathbf{q} . Other methods can be used to select the candidates alternatively, for example, one can choose candidates conditioned on the context h using a lightly pre-trained N-gram model.

6. Related Algorithms

We provide a discussion comparing CANE with the existing techniques for solving the large class space problem. Given (x, y) , NCE and its variants (Gutmann & Hyvärinen, 2012; Mnih & Kavukcuoglu, 2013; Mikolov et al., 2013; Ji et al., 2015; Titsias, 2016; Botev et al., 2017) use the observed class y as the only ‘candidate’, while CANE chooses a subset of candidates \mathcal{C}_x according to x . NCE assumes the entire noise distribution $P_{noise}(y)$ is known (e.g., a power-raised unigram distribution). However, in general multi-class classification problems, when the knowledge of the noise distribution is absent, NCE may have unstable estimations using an inaccurate noise distribution. CANE is developed for general multi-class classification problems and does not rely on a known noise distribution. Instead, CANE focuses on a small candidate set \mathcal{C}_x . Once the true class label is contained in \mathcal{C}_x with high probability, CANE will have low statistical variance. The variants of NCE (Mikolov et al., 2013; Ji et al., 2015; Titsias, 2016; Botev et al., 2017) also sample one or multiple noises to replace the normalization factor while according theoretical guarantees on the consistency and variance are rarely discussed. NCE and its variants can not speed up prediction while the Beam Tree algorithm can reduce the prediction complexity to $O(\log K)$.

The Beam Tree algorithm is related to some tree classifiers, while CANE is a general procedure and we only use tree structure to select candidates. The Beam Tree method itself is also different from existing tree classifiers. Most of the state-of-the-art tree classifiers, e.g., LOMTree (Choromanska & Langford, 2015) and Recall Tree (Daume III et al., 2017), store local classifiers in their internal nodes, and route examples through the root until reaching the leaf. Differently, the Beam Tree algorithm shown in Fig. 1 does not maintain local classifiers, and it only uses the tree structure to perform global heuristic search for candidate selection. We will compare our approach to some state-of-the-art tree classifiers in the experiments.

7. Experiments

We evaluate the CANE method in various applications in this section, including both multi-class classification problems and neural language modeling. We compare CANE with NCE, its variants and some state-of-the-art tree classifiers that have been used for large class space problems. The competitors include the standard softmax, the NCE (Mnih & Kavukcuoglu, 2013; Mnih & Teh, 2012), the BlackOut (Ji et al., 2015), the hierarchical softmax (HSM) (Morin & Bengio, 2005), the Filter Tree (Beygelzimer et al., 2009) implemented in Vowpal-Wabbit (VW, a learning platform)²,

²https://github.com/JohnLangford/vowpal_wabbit/wiki

the LOMTree (Choromanska & Langford, 2015) in VW and the Recall Tree (Daume III et al., 2017) in VW.

7.1. Classification Problems

In this section, we consider four multi-class classification problems, including the Sector³ dataset with 105 classes (Chang & Lin, 2011), the ALOI⁴ dataset with 1000 classes (Geusebroek et al., 2005), the ImageNet-2010⁵ dataset with 1000 classes, and the ImageNet-10K⁵ dataset with 10K classes (ImageNet Fall 2009 release). The data from Sector and ALOI is split into 90% training and 10% testing. In ImageNet-2010, the training set contains 1.3M images and we use the validation set containing 50K images as the test set. The ImageNet-10K data contains 9M images and we randomly split the data into two halves for training and testing by following the protocols in (Deng et al., 2010; Sánchez & Perronnin, 2011; Le, 2013). For ImageNet-2010 and ImageNet-10K datasets, similar to (Oquab et al., 2014), we transfer the mid-level representations from the pre-trained VGG-16 net (Simonyan & Zisserman, 2014) on ImageNet 2012 data (Russakovsky et al., 2015) to our case. Then, we concatenate CANE or other compared methods above the partial VGG-16 net as the top layer. The parameters of the partial VGG-16 net are pre-trained⁶ and kept fixed. Only the parameters in the top layer are trained on the target datasets, i.e., ImageNet-2010 and ImageNet-10K.

We use b -nary tree for CANE and set $b = 10$ for all classification problems. We trade off $|\mathcal{C}_x|$ and $|T_x|$ to see how these parameters affect the learning performance. Different configurations will be referred to as ‘CANE- $(|\mathcal{C}_x|$ vs. $|T_x|)$ ’. We always let $|\mathcal{C}_x| + |T_x|$ equal the number of noises used by NCE and BlackOut, so that these methods will have the same number of considered classes. We use ‘NCE- k ’ and ‘BlackOut- k ’ to denote the corresponding method with k noises. Generally, a large $|\mathcal{C}_x| + |T_x|$ and k will lower the variance of CANE, NCE and BlackOut and improve their performance, but this also increases the computation. We set $k = 10$ for Sector and ALOI and $k = 20$ for ImageNet-2010 and ImageNet-10K. We uniformly sample noises in CANE. For NCE and BlackOut, by following (Mnih & Teh, 2012; Mnih & Kavukcuoglu, 2013; Ji et al., 2015; Botev et al., 2017), we use the power-raised unigram distribution with the power factor selected from $\{0, 0.1, 0.3, 0.5, 0.75, 1\}$ to sample the noises. However, when the classes are balanced as in many cases of the classification datasets, this distribution reduces to the uniform distribution. For the com-

³<http://www.cs.utexas.edu/~xrhuang/PDSparse/>

⁴<http://www.csie.ntu.edu.tw/~cjlin/libsvmtools/datasets/multiclass.html>

⁵<http://image-net.org>

⁶http://www.robots.ox.ac.uk/~vgg/research/very_deep/

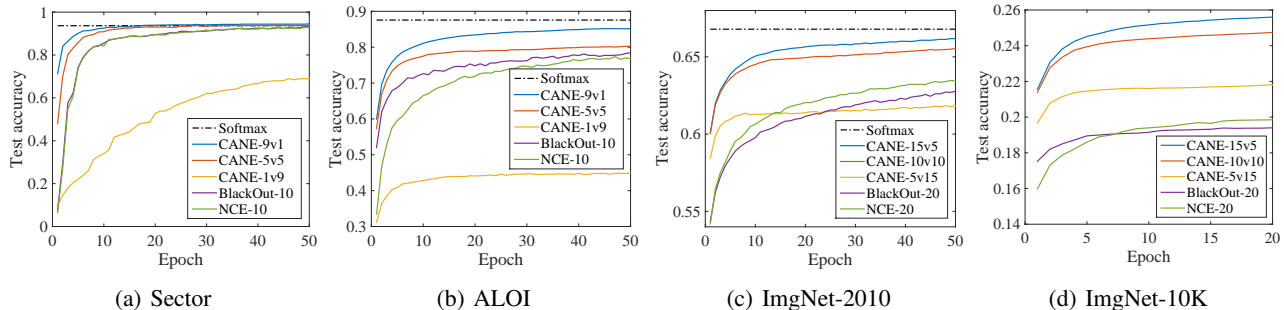


Figure 2. Results of test accuracy vs. epoch on different classification datasets.

Table 1. Training / testing time of the sampling methods. Running Softmax and testing NCE and BlackOut on large datasets are time consuming. We use multi-thread implementation for these methods and estimate the running time. ‘~’ indicates the estimated time.

Data	NCE-10	BlackOut-10	CANE-1v9	CANE-5v5	CANE-9v1	Softmax
Sector	0.4m / 0.8s	0.4m / 0.8s	1m / 0.1s	1.8m / 0.1s	2.3m / 0.2s	6.1m / 0.9s
ALOI	3m / 6s	3m / 6s	4m / 0.1s	7m / 0.3s	8m / 0.5s	28m / 7s
Data	NCE-20	BlackOut-20	CANE-5v15	CANE-10v10	CANE-15v5	Softmax
ImgNet-2010	3.5h / 8m	3.5h / 8m	4h / 0.4m	5.8h / 0.7m	6.4h / 0.9m	96h / 8.7m
ImgNet-10K	13h / ~5d	12h / ~5d	20h / 1h	33h / 1.5h	39h / 2h	~140d / ~5d

Table 2. Accuracy and training / testing time of the tree classifiers.

Data	HSM	Filter Tree	LOMTree	Recall Tree
Sector	91.36%	84.67%	84.91%	86.89%
	0.5m / 0.1s	0.4m / 0.4s	0.5m / 0.2s	0.7m / 0.2s
ALOI	65.69%	20.07%	82.70%	83.03%
	1m / 0.4s	1m / 0.2s	3.3m / 1s	2.5m / 0.2s
ImgNet	47.68%	48.29%	49.87%	61.28%
2010	4.7h / 0.5m	6.8h / 0.1m	17.8h / 0.3m	32h / 0.5m
ImgNet	17.31%	4.49%	9.72%	22.74%
10K	14h / 1h	22h / 0.3h	23h / 0.3h	68h / 1.2h

pared tree classifiers, the HSM adopts the same tree used by CANE, the Filter Tree generates a fixed tree itself in VW, the LOMTree and Recall Tree use binary trees and they are able to adjust the tree structure automatically.

All the methods use SGD with learning rate selected from $\{0.0001, 0.001, 0.01, 0.05, 0.1, 0.5, 1.0\}$. The Beam Tree algorithm requires a tree structure and we use some tree generated by a simple hierarchical clustering method on the centers of the individual classes.⁷ We run all the methods 50 epochs on Sector, ALOI and ImageNet-2010 datasets and 20 epochs on ImageNet-10K to report the accuracy vs. epoch curves. All the methods are implemented using a standard CPU machine with quad-core Intel Core i5 processor.

Fig. 2 and Table 1 show the accuracy vs. epoch plots and

⁷The method is provided in the supplementary material.

the training / testing time for NCE, BlackOut, CANE and Softmax. The tree classifiers in the VW platform require the number of training epochs as input and do not take evaluation directly after each epoch, so we report the final results of the tree classifiers in Table 2. For ImageNet-10K data, the Softmax method is very time consuming (even with multi-thread implementation) and we do not report this result. As we can observe, by fixing $|\mathcal{C}_x| + |T_x|$, using more candidates than noises in CANE will achieve better performance, because a larger \mathcal{C}_x will increase the chance to cover the target class y . The probability that the target class is included in the selected candidate set on the test data is reported in Table 3. On all the datasets, CANE with larger candidate set achieves considerable improvement compared to other methods in terms of accuracy. The speed of processing each example of CANE is slightly slower than that of NCE and BlackOut because of beam search, however, CANE shows faster convergence to reach higher accuracy. Moreover, the prediction time of CANE is much faster than those of NCE and BlackOut. It is worth mentioning that CANE exceeds some state-of-the-art results on the ImageNet-10K data, e.g., 19.2% top-1 accuracy reported in (Le, 2013) and 21.9% top-1 accuracy reported in (Mensink et al., 2013) which are conducted from $\mathcal{O}(K)$ methods; but it underperforms the recent result 28.4% in (Huang et al., 2016). This is probably because the VGG-16 net works better than the neural network structure used in (Le, 2013) and the distance-based method in (Mensink et al., 2013), while the method in (Huang et al., 2016) adopts a better feature embedding, which leads to superior prediction performance on this dataset.

Table 3. The probability that the true label is included in the selected candidate set on the test set, i.e., the top- $|\mathcal{C}_x|$ accuracy.

$ \mathcal{C}_x $	Sector	ALOI	$ \mathcal{C}_x $	ImgNet-2010	ImgNet-10K
1	68.89%	44.84%	5	76.59%	39.59%
5	96.57%	86.47%	10	87.29%	53.28%
9	97.92%	93.59%	15	91.17%	60.22%

7.2. Neural Language Modeling

In this experiment, we apply the CANE method to neural language modeling. We test the methods on two benchmark corpora: the Penn TreeBank (PTB) (Mikolov et al., 2010) and Gutenberg⁸ corpora. The Penn TreeBank dataset contains 1M tokens and we choose the most frequent 12K words appearing at least 5 times as the vocabulary. The Gutenberg dataset contains 50M tokens and the most frequent 116K words appearing at least 10 times are chosen as the vocabulary. We set the embedding size as 256 and use a LSTM model with 512 hidden states and 256 projection size. The sequence length is fixed as 20 and the learning rate is selected from $\{0.025, 0.05, 0.1, 0.2\}$.

The tree classifiers evaluated in multi-class classification problems can not be directly applied to solve the language modeling problem, so we omit their comparison and focus on the evaluation of the sampling methods. We sample 40, 60 and 80 noises for NCE and Blackout respectively and use power-raised unigram distribution with the power factor selected from $\{0, 0.25, 0.5, 0.75, 1\}$. For CANE, we adopt the one-layer tree structure discussed in Section 5 with $N = 6$ subsets, split by averaging over the word frequencies. We uniformly sample the candidates when reaching any subset. For efficiency consideration, we respectively sample 40, 60 and 80 candidates plus one more uniform noise for CANE. The experiments in this section are implemented on a machine with NVIDIA Tesla M40 GPUs.

The test perplexities are shown in Fig. 3. As we can observe, the CANE method always achieves faster convergence and lower perplexities (approaching that of Softmax) compared to NCE and Blackout under various settings. Generally, when the number of selected candidates / noises decrease, the test perplexities of all the methods increase on both datasets, while the performance degradation of CANE is not obvious. By using GPUs, all the methods can finish training within a few minutes on the PTB dataset; for the Gutenberg corpus, CANE and BlackOut have similar training time that is around 5 hours on all the three settings, while NCE spends around 6-8 hours on these tasks and Softmax uses 35 hours to finish the training.

⁸www.gutenberg.org

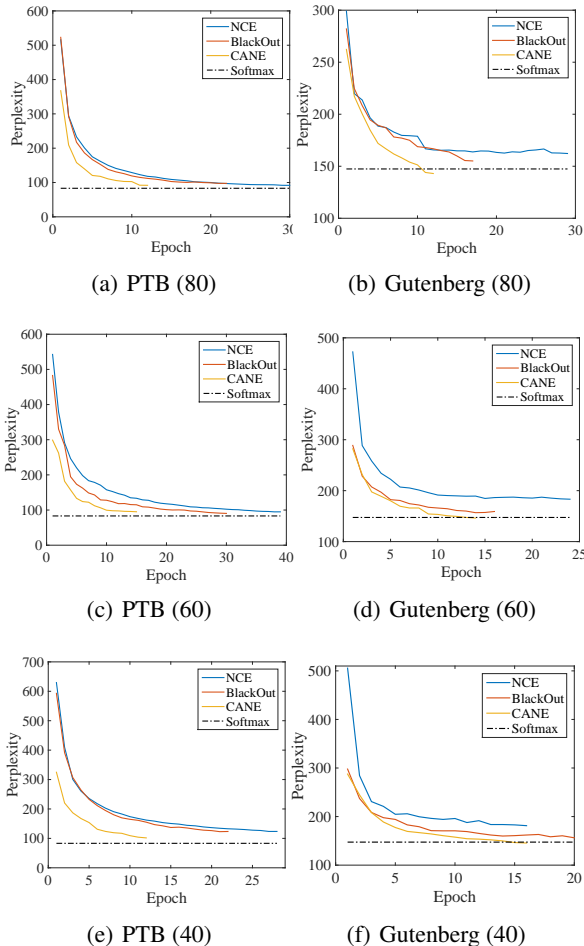


Figure 3. Test perplexity vs. training epoch on PTB and Gutenberg datasets. Numbers in the brackets indicate the number of selected candidates / noises.

8. Conclusion

We proposed Candidates vs. Noises Estimation (CANE) for fast learning in multi-class classification problems with many labels and applied this method to the word probability estimation problem in neural language models. We showed that CANE is consistent and the computation using SGD is always efficient (that is, independent of the class size K). Moreover, the new estimator has low statistical variance approaching that of the softmax logistic regression, if the observed class label belongs to the candidate set with high probability. Empirical results demonstrated that CANE is effective for speeding up both training and prediction in multi-class classification problems and CANE is effective in neural language modeling. We note that this work employs a fixed distribution (i.e., the uniform distribution) to sample noises in CANE. However it can be very useful in practice to estimate the noise distribution, i.e., q , during training, and select noise classes according to this distribution.

References

- Bengio, S., Weston, J., and Grangier, D. Label embedding trees for large multi-class tasks. In *Advances in Neural Information Processing Systems (NIPS)*, pp. 163–171, 2010.
- Beygelzimer, A., Langford, J., and Ravikumar, P. Error-correcting tournaments. In *International Conference on Algorithmic Learning Theory (ALT)*, pp. 247–262, 2009.
- Botev, A., Zheng, B., and Barber, D. Complementary sum sampling for likelihood approximation in large scale classification. In *Artificial Intelligence and Statistics (AISTATS)*, pp. 1030–1038, 2017.
- Chang, C.-C. and Lin, C.-J. Libsvm: a library for support vector machines. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 2(3):27, 2011.
- Choromanska, A. E. and Langford, J. Logarithmic time online multiclass prediction. In *Advances in Neural Information Processing Systems (NIPS)*, pp. 55–63, 2015.
- Daume III, H., Karampatziakis, N., Langford, J., and Mineiro, P. Logarithmic time one-against-some. In *International Conference on Machine Learning (ICML)*, pp. 923–932, 2017.
- Deng, J., Berg, A. C., Li, K., and Fei-Fei, L. What does classifying more than 10,000 image categories tell us? In *European Conference on Computer Vision (ECCV)*, pp. 71–84, 2010.
- Deng, J., Satheesh, S., Berg, A. C., and Li, F. Fast and balanced: Efficient label tree learning for large scale object recognition. In *Advances in Neural Information Processing Systems (NIPS)*, pp. 567–575, 2011.
- Geusebroek, J.-M., Burghouts, G. J., and Smeulders, A. W. The amsterdam library of object images. *International Journal of Computer Vision (IJCV)*, 61(1):103–112, 2005.
- Gutmann, M. U. and Hyvärinen, A. Noise-contrastive estimation of unnormalized statistical models, with applications to natural image statistics. *Journal of Machine Learning Research (JMLR)*, 13(Feb):307–361, 2012.
- Hochreiter, S. and Schmidhuber, J. Long short-term memory. *Neural Computation*, 9(8):1735–1780, 1997.
- Huang, C., Loy, C. C., and Tang, X. Local similarity-aware deep feature embedding. In *Advances in Neural Information Processing Systems*, pp. 1262–1270, 2016.
- Jernite, Y., Choromanska, A., and Sontag, D. Simultaneous learning of trees and representations for extreme classification and density estimation. In *International Conference on Machine Learning (ICML)*, pp. 1665–1674, 2017.
- Ji, S., Vishwanathan, S., Satish, N., Anderson, M. J., and Dubey, P. Blackout: Speeding up recurrent neural network language models with very large vocabularies. *arXiv preprint arXiv:1511.06909*, 2015.
- Le, Q. V. Building high-level features using large scale unsupervised learning. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 8595–8598, 2013.
- Mensink, T., Verbeek, J., Perronnin, F., and Csurka, G. Distance-based image classification: Generalizing to new classes at near-zero cost. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 35(11):2624–2637, 2013.
- Mikolov, T., Karafiát, M., Burget, L., Černocký, J., and Khudanpur, S. Recurrent neural network based language model. In *Eleventh Annual Conference of the International Speech Communication Association*, 2010.
- Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., and Dean, J. Distributed representations of words and phrases and their compositionality. In *Advances in Neural Information Processing Systems (NIPS)*, pp. 3111–3119, 2013.
- Mnih, A. and Hinton, G. E. A scalable hierarchical distributed language model. In *Advances in Neural Information Processing Systems (NIPS)*, pp. 1081–1088, 2009.
- Mnih, A. and Kavukcuoglu, K. Learning word embeddings efficiently with noise-contrastive estimation. In *Advances in Neural Information Processing Systems (NIPS)*, pp. 2265–2273, 2013.
- Mnih, A. and Teh, Y. W. A fast and simple algorithm for training neural probabilistic language models. In *Proceedings of the 29th International Conference on Machine Learning (ICML)*, pp. 1751–1758, 2012.
- Morin, F. and Bengio, Y. Hierarchical probabilistic neural network language model. In *The International Conference on Artificial Intelligence and Statistics (AISTATS)*, volume 5, pp. 246–252, 2005.
- Oquab, M., Bottou, L., Laptev, I., and Sivic, J. Learning and transferring mid-level image representations using convolutional neural networks. In *Proceedings of The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 1717–1724, 2014.
- Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., C. Berg, A., and Fei-Fei, L. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision (IJCV)*, 115(3):211–252, 2015.

- Sánchez, J. and Perronnin, F. High-dimensional signature compression for large-scale image classification. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 1665–1672, 2011.
- Simonyan, K. and Zisserman, A. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- Sordoni, A., Galley, M., Auli, M., Brockett, C., Ji, Y., Mitchell, M., Nie, J.-Y., Gao, J., and Dolan, B. A neural network approach to context-sensitive generation of conversational responses. *arXiv preprint arXiv:1506.06714*, 2015.
- Titsias, M. K. One-vs-each approximation to softmax for scalable estimation of probabilities. In *Advances in Neural Information Processing Systems (NIPS)*, pp. 4161–4169, 2016.
- Vaswani, A., Zhao, Y., Fossum, V., and Chiang, D. Decoding with large-scale neural language models improves translation. In *The Conference on Empirical Methods on Natural Language Processing (EMNLP)*, pp. 1387–1392, 2013.