

Adversarial Contrastive Estimation

Avishek Joey Bose^{1,2,* †} Huan Ling^{1,2,* †} Yanshuai Cao^{1,*}

¹Borealis AI ²University of Toronto

{joey.bose, huan.ling}@mail.utoronto.ca

{yanshuai.cao}@borealisai.com

Abstract

Learning by contrasting positive and negative samples is a general strategy adopted by many methods. Noise contrastive estimation (NCE) for word embeddings and translating embeddings for knowledge graphs are examples in NLP employing this approach. In this work, we view contrastive learning as an abstraction of all such methods and augment the negative sampler into a mixture distribution containing an adversarially learned sampler. The resulting adaptive sampler finds harder negative examples, which forces the main model to learn a better representation of the data. We evaluate our proposal on learning word embeddings, order embeddings and knowledge graph embeddings and observe both faster convergence and improved results on multiple metrics.

1 Introduction

Many models learn by contrasting losses on observed positive examples with those on some fictitious negative examples, trying to decrease some score on positive ones while increasing it on negative ones. There are multiple reasons why such contrastive learning approach is needed. Computational tractability is one. For instance, instead of using softmax to predict a word for learning word embeddings, noise contrastive estimation (NCE) (Dyer, 2014; Mnih and Teh, 2012) can be used in skip-gram or CBOW word embedding models (Gutmann and Hyvärinen, 2012; Mikolov et al., 2013; Mnih and Kavukcuoglu, 2013; Vaswani et al., 2013). Another reason is

modeling need, as certain assumptions are best expressed as some score or energy in margin based or un-normalized probability models (Smith and Eisner, 2005). For example, modeling entity relations as translations or variants thereof in a vector space naturally leads to a distance-based score to be minimized for observed entity-relation-entity triplets (Bordes et al., 2013).

Given a scoring function, the gradient of the model’s parameters on observed positive examples can be readily computed, but the negative phase requires a design decision on how to sample data. In noise contrastive estimation for word embeddings, a negative example is formed by replacing a component of a positive pair by randomly selecting a sampled word from the vocabulary, resulting in a fictitious word-context pair which would be unlikely to actually exist in the dataset. This negative sampling by corruption approach is also used in learning knowledge graph embeddings (Bordes et al., 2013; Lin et al., 2015; Ji et al., 2015; Wang et al., 2014; Trouillon et al., 2016; Yang et al., 2014; Dettmers et al., 2017), order embeddings (Vendrov et al., 2016), caption generation (Dai and Lin, 2017), etc.

Typically the corruption distribution is the same for all inputs like in skip-gram or CBOW NCE, rather than being a conditional distribution that takes into account information about the input sample under consideration. Furthermore, the corruption process usually only encodes a human prior as to what constitutes a hard negative sample, rather than being learned from data. For these two reasons, the simple fixed corruption process often yields only easy negative examples. Easy negatives are sub-optimal for learning discriminative representation as they do not force the model to find critical characteristics of observed positive data, which has been independently discovered in applications outside NLP previously (Shrivastava

* authors contributed equally

† Work done while author was an intern at Borealis AI

et al., 2016). Even if hard negatives are occasionally reached, the infrequency means slow convergence. Designing a more sophisticated corruption process could be fruitful, but requires costly trial-and-error by a human expert.

In this work, we propose to augment the simple corruption noise process in various embedding models with an adversarially learned conditional distribution, forming a mixture negative sampler that adapts to the underlying data and the embedding model training progress. The resulting method is referred to as adversarial contrastive estimation (ACE). The adaptive conditional model engages in a minimax game with the primary embedding model, much like in Generative Adversarial Networks (GANs) (Goodfellow et al., 2014a), where a discriminator net (D), tries to distinguish samples produced by a generator (G) from real data (Goodfellow et al., 2014b). In ACE, the main model learns to distinguish between a real positive example and a negative sample selected by the mixture of a fixed NCE sampler and an adversarial generator. The main model and the generator takes alternating turns to update their parameters. In fact, our method can be viewed as a conditional GAN (Mirza and Osindero, 2014) on discrete inputs, with a mixture generator consisting of a learned and a fixed distribution, with additional techniques introduced to achieve stable and convergent training of embedding models.

In our proposed ACE approach, the conditional sampler finds harder negatives than NCE, while being able to gracefully fall back to NCE whenever the generator cannot find hard negatives. We demonstrate the efficacy and generality of the proposed method on three different learning tasks, word embeddings (Mikolov et al., 2013), order embeddings (Vendrov et al., 2016) and knowledge graph embeddings (Ji et al., 2015).

2 Method

2.1 Background: contrastive learning

In the most general form, our method applies to supervised learning problems with a contrastive objective of the following form:

$$L(\omega) = \mathbb{E}_{p(x^+, y^+, y^-)} l_\omega(x^+, y^+, y^-) \quad (1)$$

where $l_\omega(x^+, y^+, y^-)$ captures both the model with parameters ω and the loss that scores a positive tuple (x^+, y^+) against a negative one (x^+, y^-) . $\mathbb{E}_{p(x^+, y^+, y^-)}(\cdot)$ denotes expectation

with respect to some joint distribution over positive and negative samples. Furthermore, by the law of total expectation, and the fact that given x^+ , the negative sampling is not dependent on the positive label, i.e. $p(y^+, y^- | x^+) = p(y^+ | x^+)p(y^- | x^+)$, Eq. 1 can be re-written as

$$\mathbb{E}_{p(x^+)} [\mathbb{E}_{p(y^+ | x^+)p(y^- | x^+)} l_\omega(x^+, y^+, y^-)] \quad (2)$$

Separable loss

In the case where the loss decomposes into a sum of scores on positive and negative tuples such as $l_\omega(x^+, y^+, y^-) = s_\omega(x^+, y^+) - \tilde{s}_\omega(x^+, y^-)$, then Expression. 2 becomes

$$\mathbb{E}_{p^+(x)} [\mathbb{E}_{p^+(y|x)} s_\omega(x, y) - \mathbb{E}_{p^-(y|x)} \tilde{s}_\omega(x, y)] \quad (3)$$

where we moved the $+$ and $-$ to p for notational brevity. Learning by stochastic gradient descent aims to adjust ω to pushing down $s_\omega(x, y)$ on samples from p^+ while pushing up $\tilde{s}_\omega(x, y)$ on samples from p^- . Note that for generality, the scoring function for negative samples, denoted by \tilde{s}_ω , could be slightly different from s_ω . For instance, \tilde{s} could contain a margin as in the case of Order Embeddings in Sec. 4.2.

Non separable loss

Eq. 1 is the general form that we would like to consider because for certain problems, the loss function cannot be separated into sums of terms containing only positive (x^+, y^+) and terms with negatives (x^+, y^-) . An example of such a non-separable loss is the triplet ranking loss (Schroff et al., 2015): $l_\omega = \max(0, \eta + s_\omega(x^+, y^+) - s_\omega(x^+, y^-))$, which does not decompose due to the rectification.

Noise contrastive estimation

The typical NCE approach in tasks such as word embeddings (Mikolov et al., 2013), order embeddings (Vendrov et al., 2016), and knowledge graph embeddings can be viewed as a special case of Eq. 2 by taking $p(y^- | x^+)$ to be some unconditional $p_{nce}(y)$.

This leads to efficient computation during training, however, $p_{nce}(y)$ sacrifices the sampling efficiency of learning as the negatives produced using a fixed distribution are not tailored toward x^+ , and as a result are not necessarily hard negative examples. Thus, the model is not forced to discover discriminative representation of observed positive

data. As training progresses, more and more negative examples are correctly learned, the probability of drawing a hard negative example diminishes further, causing slow convergence.

2.2 Adversarial mixture noise

To remedy the above mentioned problem of a fixed unconditional negative sampler, we propose to augment it into a mixture one, $\lambda p_{nce}(y) + (1 - \lambda)g_\theta(y|x)$, where g_θ is a conditional distribution with a learnable parameter θ and λ is a hyperparameter. The objective in Expression. 2 can then be written as (conditioned on x for notational brevity):

$$L(\omega, \theta; x) = \lambda \mathbb{E}_{p(y^+|x)p_{nce}(y^-)} l_\omega(x, y^+, y^-) + (1 - \lambda) \mathbb{E}_{p(y^+|x)g_\theta(y^-|x)} l_\omega(x, y^+, y^-) \quad (4)$$

We learn (ω, θ) in a GAN-style minimax game:

$$\min_{\omega} \max_{\theta} V(\omega, \theta) = \min_{\omega} \max_{\theta} \mathbb{E}_{p^+(x)} L(\omega, \theta; x) \quad (5)$$

The embedding model behind $l_\omega(x, y^+, y^-)$ is similar to the discriminator in (conditional) GAN (or critic in Wasserstein (Arjovsky et al., 2017) or Energy-based GAN (Zhao et al., 2016), while $g_\theta(y|x)$ acts as the generator. Henceforth, we will use the term discriminator (D) and embedding model interchangeably, and refer to g_θ as the generator.

2.3 Learning the generator

There is one important distinction to typical GAN: $g_\theta(y|x)$ defines a categorical distribution over possible y values, and samples are drawn accordingly; in contrast to typical GAN over continuous data space such as images, where samples are generated by an implicit generative model that warps noise vectors into data points. Due to the discrete sampling step, g_θ cannot learn by receiving gradient through the discriminator. One possible solution is to use the Gumbel-softmax reparametrization trick (Jang et al., 2016; Maddison et al., 2016), which gives a differentiable approximation. However, this differentiability comes at the cost of drawing N Gumbel samples per each categorical sample, where N is the number of categories. For word embeddings, N is the vocabulary size, and for knowledge graph embeddings, N is the number of entities, both leading to infeasible computational requirements.

Instead, we use the REINFORCE (Williams, 1992) gradient estimator for $\nabla_{\theta} L(\theta, x)$:

$$(1 - \lambda) \mathbb{E} [-l_\omega(x, y^+, y^-) \nabla_{\theta} \log(g_\theta(y^-|x))] \quad (6)$$

where the expectation \mathbb{E} is with respect to $p(y^+, y^-|x) = p(y^+|x)g_\theta(y^-|x)$, and the discriminator loss $l_\omega(x, y^+, y^-)$ acts as the reward.

With a separable loss, the (conditional) value function of the minimax game is:

$$L(\omega, \theta; x) = \mathbb{E}_{p^+(y|x)} s_\omega(x, y) - \mathbb{E}_{p_{nce}(y)} \tilde{s}_\omega(x, y) - \mathbb{E}_{g_\theta(y|x)} \tilde{s}_\omega(x, y) \quad (7)$$

and only the last term depends on the generator parameter ω . Hence, with a separable loss, the reward is $-\tilde{s}(x^+, y^-)$. This reduction does not happen with a non-separable loss, and we have to use $l_\omega(x, y^+, y^-)$.

2.4 Entropy and training stability

GAN training can suffer from instability and degeneracy where the generator probability mass collapses to a few modes or points. Much work has been done to stabilize GAN training in the continuous case (Arjovsky et al., 2017; Gulrajani et al., 2017; Cao et al., 2018). In ACE, if the generator g_θ probability mass collapses to a few candidates, then after the discriminator successfully learns about these negatives, g_θ cannot adapt to select new hard negatives, because the REINFORCE gradient estimator Eq. 6 relies on g_θ being able to explore other candidates during sampling. Therefore, if the g_θ probability mass collapses, instead of leading to oscillation as in typical GAN, the min-max game in ACE reaches an equilibrium where the discriminator wins and g_θ can no longer adapt, then ACE falls back to NCE since the negative sampler has another mixture component from NCE.

This behavior of gracefully falling back to NCE is more desirable than the alternative of stalled training if $p^-(y|x)$ does not have a simple p_{nce} mixture component. However, we would still like to avoid such collapse, as the adversarial samples provide greater learning signals than NCE samples. To this end, we propose to use a regularizer to encourage the categorical distribution $g_\theta(y|x)$ to have high entropy. In order to make the the regularizer interpretable and its hyperparameters easy to tune, we design the following form:

$$R_{ent}(x) = \min(0, c - H(g_\theta(y|x))) \quad (8)$$

where $H(g_\theta(y|x))$ is the entropy of the categorical distribution $g_\theta(y|x)$, and $c = \log(k)$ is the entropy of a uniform distribution over k choices, and k is a hyper-parameter. Intuitively, R_{ent} expresses the prior that the generator should spread its mass over more than k choices for each x .

2.5 Handling false negatives

During negative sampling, $p^-(y|x)$ could actually produce y that forms a positive pair that exists in the training set, i.e., a false negative. This possibility exists in NCE already, but since p_{nce} is not adaptive, the probability of sampling a false negative is low. Hence in NCE, the score on this false negative (true observation) pair is pushed up less in the negative term than in the positive term.

However, with the adaptive sampler, $g_\omega(y|x)$, false negatives become a much more severe issue. $g_\omega(y|x)$ can learn to concentrate its mass on a few false negatives, significantly canceling the learning of those observations in the positive phase. The entropy regularization reduces this problem as it forces the generator to spread its mass, hence reducing the chance of a false negative.

To further alleviate this problem, whenever computationally feasible, we apply an additional two-step technique. First, we maintain a hash map of the training data in memory, and use it to efficiently detect if a negative sample (x^+, y^-) is an actual observation. If so, its contribution to the loss is given a zero weight in ω learning step. Second, to update θ in the generator learning step, the reward for false negative samples are replaced by a large penalty, so that the REINFORCE gradient update would steer g_θ away from those samples. The second step is needed to prevent null computation where g_θ learns to sample false negatives which are subsequently ignored by the discriminator update for ω .

2.6 Variance Reduction

The basic REINFORCE gradient estimator is poised with high variance, so in practice one often needs to apply variance reduction techniques. The most basic form of variance reduction is to subtract a baseline from the reward. As long as the baseline is not a function of actions (i.e., samples y^- being drawn), the REINFORCE gradient estimator remains unbiased. More advanced gradient estimators exist that also reduce variance (Grathwohl et al., 2017; Tucker et al., 2017; Liu et al., 2018), but for simplicity we use the

self-critical baseline method (Rennie et al., 2016), where the baseline is $b(x) = l_\omega(y^+, y^*, x)$, or $b(x) = -\tilde{s}_\omega(y^*, x)$ in the separable loss case, and $y^* = \operatorname{argmax}_i g_\theta(y_i|x)$. In other words, the baseline is the reward of the most likely sample according to the generator.

2.7 Improving exploration in g_θ by leveraging NCE samples

In Sec. 2.4 we touched on the need for sufficient exploration in g_θ . It is possible to also leverage negative samples from NCE to help the generator learn. This is essentially off-policy exploration in reinforcement learning since NCE samples are not drawn according to $g_\theta(y|x)$. The generator learning can use importance re-weighting to leverage those samples. The resulting REINFORCE gradient estimator is basically the same as Eq. 6 except that the rewards are reweighted by $g_\theta(y^-|x)/p_{nce}(y^-)$, and the expectation is with respect to $p(y^+|x)p_{nce}(y^-)$. This additional off-policy learning term provides gradient information for generator learning if $g_\theta(y^-|x)$ is not zero, meaning that for it to be effective in helping exploration, the generator cannot be collapsed at the first place. Hence, in practice, this term is only used to further help on top of the entropy regularization, but it does not replace it.

3 Related Work

Smith and Eisner (2005) proposed contrastive estimation as a way for unsupervised learning of log-linear models by taking implicit evidence from user-defined neighborhoods around observed datapoints. Gutmann and Hyvärinen (2010) introduced NCE as an alternative to the hierarchical softmax. In the works of Mnih and Teh (2012) and Mnih and Kavukcuoglu (2013), NCE is applied to log-bilinear models and Vaswani et al. (2013) applied NCE to neural probabilistic language models (Yoshua et al., 2003). Compared to these previous NCE methods that rely on simple fixed sampling heuristics, ACE uses an adaptive sampler that produces harder negatives.

In the domain of max-margin estimation for structured prediction (Taskar et al., 2005), loss augmented MAP inference plays the role of finding hard negatives (the hardest). However, this inference is only tractable in a limited class of models such structured SVM (Tsochantaridis et al., 2005). Compared to those models that use exact

maximization to find the hardest negative configuration each time, the generator in ACE can be viewed as learning an approximate amortized inference network. Concurrently to this work, [Tu and Gimpel \(2018\)](#) proposes a very similar framework, using a learned inference network for Structured prediction energy networks (SPEN) ([Be-langer and McCallum, 2016](#)).

Concurrent with our work, there have been other interests in applying the GAN to NLP problems ([Fedus et al., 2018](#); [Wang et al., 2018](#); [Cai and Wang, 2017](#)). Knowledge graph models naturally lend to a GAN setup, and has been the subject of study in [Wang et al. \(2018\)](#) and [Cai and Wang \(2017\)](#). These two concurrent works are most closely related to one of the three tasks on which we study ACE in this work. Besides a more general formulation that applies to problems beyond those considered in [Wang et al. \(2018\)](#) and [Cai and Wang \(2017\)](#), the techniques introduced in our work on handling false negatives and entropy regularization lead to improved experimental results as shown in [Sec. 5.4](#).

4 Application of ACE on three tasks

4.1 Word Embeddings

Word embeddings learn a vector representation of words from co-occurrences in a text corpus. NCE casts this learning problem as a binary classification where the model tries to distinguish positive word and context pairs, from negative noise samples composed of word and false context pairs. The NCE objective in Skip-gram ([Mikolov et al., 2013](#)) for word embeddings is a separable loss of the form:

$$L = - \sum_{w_t \in V} [\log p(y = 1 | w_t, w_c^+)] + \sum_{c=1}^K \log p(y = 0 | w_t, w_c^-) \quad (9)$$

Here, w_c^+ is sampled from the set of true contexts and $w_c^- \sim Q$ is sampled k times from a fixed noise distribution. [Mikolov et al. \(2013\)](#) introduced a further simplification of NCE, called “Negative Sampling” ([Dyer, 2014](#)). With respect to our ACE framework, the difference between NCE and Negative Sampling is inconsequential, so we continue the discussion using NCE. A drawback of this sampling scheme is that it favors more common words as context. Another issue

is that the negative context words are sampled in the same way, rather than tailored toward the actual target word. To apply ACE to this problem we first define the value function for the minimax game, $V(D, G)$, as follows:

$$V(D, G) = \mathbb{E}_{p^+(w_c)} [\log D(w_c, w_t)] - \mathbb{E}_{p_{nce}(w_c)} [-\log(1 - D(w_c, w_t))] - \mathbb{E}_{g_\theta(w_c|w_t)} [-\log(1 - D(w_c, w_t))] \quad (10)$$

with $D = p(y = 1 | w_t, w_c)$ and $G = g_\theta(w_c | w_t)$.

Implementation details

For our experiments, we train all our models on a single pass of the May 2017 dump of the English Wikipedia with lowercased unigrams. The vocabulary size is restricted to the top 150k most frequent words when training from scratch while for finetuning we use the same vocabulary as [Pennington et al. \(2014\)](#), which is 400k of the most frequent words. We use 5 NCE samples for each positive sample and 1 adversarial sample in a window size of 10 and the same positive subsampling scheme proposed by [Mikolov et al. \(2013\)](#). Learning for both G and D uses Adam ([Kingma and Ba, 2014](#)) optimizer with its default parameters. Our conditional discriminator is modeled using the Skip-Gram architecture, which is a two layer neural network with a linear mapping between the layers. The generator network consists of an embedding layer followed by two small hidden layers, followed by an output softmax layer. The first layer of the generator shares its weights with the second embedding layer in the discriminator network, which we find really speeds up convergence as the generator does not have to relearn its own set of embeddings. The difference between the discriminator and generator is that a sigmoid nonlinearity is used after the second layer in the discriminator, while in the generator, a softmax layer is used to define a categorical distribution over negative word candidates. We find that controlling the generator entropy is critical for finetuning experiments as otherwise the generator collapses to its favorite negative sample. The word embeddings are taken to be the first dense matrix in the discriminator.

4.2 Order Embeddings Hypernym Prediction

As introduced in [Vendrov et al. \(2016\)](#), ordered representations over hierarchy can be learned by

order embeddings. An example task for such ordered representation is hypernym prediction. A hypernym pair is a pair of concepts where the first concept is a specialization or an instance of the second.

For completeness, we briefly describe order embeddings, then analyze ACE on the hypernym prediction task. In order embeddings, each entity is represented by a vector in \mathbb{R}^N , the score for a positive ordered pair of entities (x, y) is defined by $s_\omega(x, y) = \|\max(0, y - x)\|^2$ and, score for a negative ordered pair (x^+, y^-) is defined by $\tilde{s}_\omega(x^+, y^-) = \max\{0, \eta - s(x^+, y^-)\}$, where η is the margin. Let $f(u)$ be the embedding function which takes an entity as input and outputs an embedding vector. We define P as a set of positive pairs and N as negative pairs, the separable loss function for order embedding task is defined by:

$$L = \sum_{(u,v) \in P} s_\omega(f(u), f(v)) + \sum_{(u,v) \in N} \tilde{s}(f(u), f(v)) \quad (11)$$

Implementation details

Our generator for this task is just a linear fully connected softmax layer, taking an embedding vector from discriminator as input and outputting a categorical distribution over the entity set. For the discriminator, we inherit all model setting from [Vendrov et al. \(2016\)](#): we use 50 dimensions hidden state and batch size 1000, a learning rate of 0.01 and the Adam optimizer. For the generator, we use a batch size of 1000, a learning rate 0.01 and the Adam optimizer. We apply weight decay with rate 0.1 and entropy loss regularization as described in [Sec. 2.4](#). We handle false negative as described in [Sec. 2.5](#). After cross validation, variance reduction and leveraging NCE samples does not greatly affect the order embedding task.

4.3 Knowledge Graph Embeddings

Knowledge graphs contain entity and relation data of the form *(head entity, relation, tail entity)*, and the goal is to learn from observed positive entity relations and predict missing links (a.k.a. link prediction). There have been many works on knowledge graph embeddings, e.g. TransE ([Bordes et al., 2013](#)), TransR ([Lin et al., 2015](#)), TransH ([Wang et al., 2014](#)), TransD ([Ji et al., 2015](#)), Complex ([Trouillon et al., 2016](#)), DistMult ([Yang et al., 2014](#)) and ConvE ([Dettmers et al., 2017](#)). Many of them use a contrastive learning objective. Here we

take TransD as an example, and modify its noise contrastive learning to ACE, and demonstrate significant improvement in sample efficiency and link prediction results.

Implementation details

Let a positive entity-relation-entity triplet be denoted by $\xi^+ = (h^+, r^+, t^+)$, and a negative triplet could either have its head or tail be a negative sample, i.e. $\xi^- = (h^-, r^+, t^+)$ or $\xi^- = (h^+, r^+, t^-)$. In either case, the general formulation in [Sec. 2.1](#) still applies. The non-separable loss function takes on the form:

$$l = \max(0, \eta + s_\omega(\xi^+) - s_\omega(\xi^-)) \quad (12)$$

The scoring rule is:

$$s = \|\mathbf{h}_\perp + \mathbf{r} - \mathbf{t}_\perp\| \quad (13)$$

where \mathbf{r} is the embedding vector for r , and \mathbf{h}_\perp is projection of the embedding of h onto the space of \mathbf{r} by $\mathbf{h}_\perp = \mathbf{h} + \mathbf{r}_p \mathbf{h}_p^\top \mathbf{h}$, where \mathbf{r}_p and \mathbf{h}_p are projection parameters of the model. \mathbf{t}_\perp is defined in a similar way through parameters \mathbf{t} , \mathbf{t}_p and \mathbf{r}_p .

The form of the generator $g_\theta(t^- | r^+, h^+)$ is chosen to be $f_\theta(\mathbf{h}_\perp, \mathbf{h}_\perp + \mathbf{r})$, where f_θ is a feedforward neural net that concatenates its two input arguments, then propagates through two hidden layers, followed by a final softmax output layer. As a function of (r^+, h^+) , g_θ shares parameter with the discriminator, as the inputs to f_θ are the embedding vectors. During generator learning, only θ is updated and the TransD model embedding parameters are frozen.

5 Experiments

We evaluate ACE with experiments on word embeddings, order embeddings, and knowledge graph embeddings tasks. In short, whenever the original learning objective is contrastive (all tasks except Glove fine-tuning) our results consistently show that ACE improves over NCE. In some cases, we include additional comparisons to the state-of-art results on the task to put the significance of such improvements in context: the generic ACE can often make a reasonable baseline competitive with SOTA methods that are optimized for the task.

For word embeddings, we evaluate models trained from scratch as well as fine-tuned Glove models ([Pennington et al., 2014](#)) on word similarity tasks that consist of computing the similarity

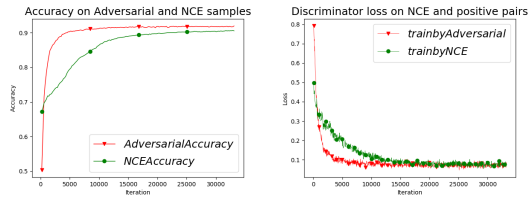


Figure 1: **Left:** Order embedding Accuracy plot. **Right:** Order embedding discriminator Loss plot on NCE sampled negative pairs and positive pairs.

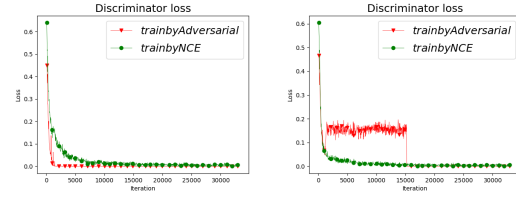


Figure 2: loss curve on NCE negative pairs and ACE negative pairs. **Left:** without entropy and weight decay. **Right:** with entropy and weight decay

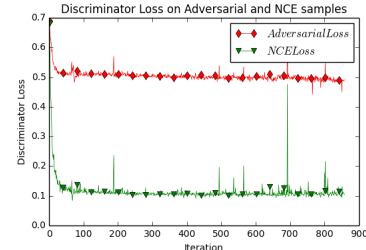
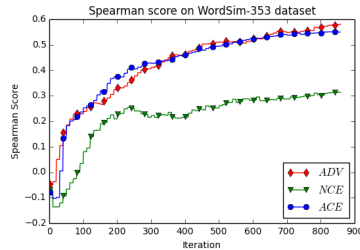
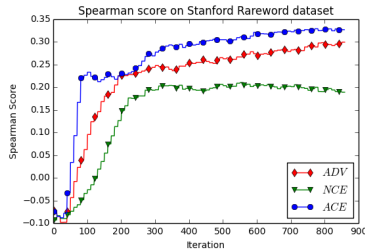


Figure 3: **Left:** Rare Word, **Right:** WS353 similarity scores during the first epoch of training.

Figure 4: Training from scratch losses on the Discriminator

between word pairs where the ground truth is an average of human scores. We choose the Rare word dataset (Luong et al., 2013) and WordSim-353 (Finkelstein et al., 2001) by virtue of our hypothesis that ACE learns better representations for both rare and frequent words. We also qualitatively evaluate ACE word embeddings by inspecting the nearest neighbors of selected words.

For the hypernym prediction task, following Vendrov et al. (2016), hypernym pairs are created from the WordNet hierarchy’s transitive closure. We use the released random development split and test split from Vendrov et al. (2016), which both contain 4000 edges.

For knowledge graph embeddings, we use TransD (Ji et al., 2015) as our base model, and perform ablation study to analyze the behavior of ACE with various add-on features, and confirm that entropy regularization is crucial for good performance in ACE. We also obtain link prediction results that are competitive or superior to the state-of-the-art on the WN18 dataset (Bordes et al., 2014).

5.1 Training Word Embeddings from scratch

In this experiment, we empirically observe that training word embeddings using ACE converges significantly faster than NCE after one epoch. As shown in Fig. 3 both ACE (a mixture of p_{nce} and g_{θ}) and just g_{θ} (denoted by ADV) significantly outperforms the NCE baseline, with an absolute improvement of 73.1% and 58.5% respectively on RW score. We note similar results on WordSim-353 dataset where ACE and ADV outperforms

NCE by 40.4% and 45.7%. We also evaluate our model qualitatively by inspecting the nearest neighbors of selected words in Table. 1. We first present the five nearest neighbors to each word to show that both NCE and ACE models learn sensible embeddings. We then show that ACE embeddings have much better semantic relevance in a larger neighborhood (nearest neighbor 45-50).

5.2 Finetuning Word Embeddings

We take off-the-shelf pre-trained Glove embeddings which were trained using 6 billion tokens (Pennington et al., 2014) and fine-tune them using our algorithm. It is interesting to note that the original Glove objective does not fit into the contrastive learning framework, but nonetheless we find that they benefit from ACE. In fact, we observe that training such that 75% of the words appear as positive contexts is sufficient to beat the largest dimensionality pre-trained Glove model on word similarity tasks. We evaluate our performance on the Rare Word and WordSim353 data. As can be seen from our results in Table 2, ACE on RW is not always better and for the 100d and 300d Glove embeddings is marginally worse. However, on WordSim353 ACE does considerably better across the board to the point where 50d Glove embeddings outperform the 300d baseline Glove model.

5.3 Hypernym Prediction

As shown in Table 3, with ACE training, our method achieves a 1.5% improvement on accu-

	Queen	King	Computer	Man	Woman
Skip-Gram NCE Top 5	princess king empress pxqueen monarch	prince queen kings emperor monarch	computers computing software microcomputer mainframe	woman boy girl stranger person	girl man prostitute person divorcee
Skip-Gram NCE Top 45-50	sambiria phongsri safrit mcelvoy tsarina	ereric mumbere empress saxonvm pretender	hypercard neurotechnology lgp pcs keystroke	angiomata someone bespectacled hero clown	suitor nymphomaniac barmaid redheaded jew
Skip-Gram ACE Top 5	princess prince elizabeth duke consort	prince vi kings duke iii	software computers applications computing hardware	woman girl tells dead boy	girl herself man lover tells
Skip-Gram ACE Top 45-50	baron abbey throne marie victoria	earl holy cardinal aragon princes	files information device design compatible	kid told revenge magic angry	aunt maid wife lady bride

Table 1: Top 5 Nearest Neighbors of Words followed by Neighbors 45-50 for different Models.

	RW	WS353
Skipgram Only NCE baseline	18.90	31.35
Skipgram + Only ADV	29.96	58.05
Skipgram + ACE	32.71	55.00
Glove-50 (Recomputed based on(Pennington et al., 2014))	34.02	49.51
Glove-100 (Recomputed based on(Pennington et al., 2014))	36.64	52.76
Glove-300 (Recomputed based on(Pennington et al., 2014))	41.18	60.12
Glove-50 + ACE	35.60	60.46
Glove-100 + ACE	36.51	63.29
Glove-300 + ACE	40.57	66.50

Table 2: Spearman score ($\rho * 100$) on RW and WS353 Datasets. We trained a skipgram model from scratch under various settings for only 1 epoch on wikipedia. For finetuned models we re-computed the scores based on the publicly available 6B tokens Glove models and we finetuned until roughly 75% of the vocabulary was seen.

racy over [Vendrov et al. \(2016\)](#) without tuning any of the discriminator’s hyperparameters. We further report training curve in Fig. 1, we report loss curve on randomly sampled pairs. We stress that in the ACE model, we train random pairs and generator generated pairs jointly, as shown in Fig. 2, hard negatives help the order embedding model converges faster.

5.4 Ablation Study and Improving TransD

To analyze different aspects of ACE, we perform an ablation study on the knowledge graph embedding task. As described in Sec. 4.3, the base

Method	Accuracy (%)
order-embeddings	90.6
order-embeddings + Our ACE	92.0

Table 3: Order Embedding Performance

model (discriminator) we apply ACE to is TransD ([Ji et al., 2015](#)). Fig. 5 shows validation performance as training progresses. All variants of ACE converges to better results than base NCE. Among ACE variants, all methods that include entropy regularization significantly outperform without entropy regularization. Without the self critical baseline variance reduction, learning could progress faster at the beginning but the final performance suffers slightly. The best performance is obtained without the additional off-policy learning of the generator.

Table. 4 shows the final test results on WN18 link prediction task. It is interesting to note that ACE improves MRR score more significantly than hit@10. As MRR is a lot more sensitive to the top rankings, i.e., how the correct configuration ranks among the competitive alternatives, this is consistent with the fact that ACE samples hard negatives and forces the base model to learn a more discriminative representation of the positive examples.

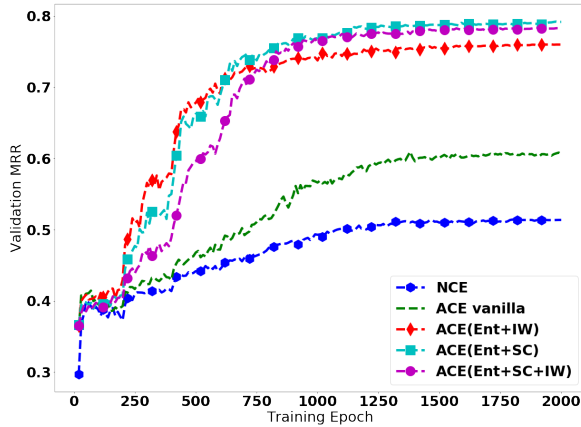


Figure 5: Ablation study: measuring validation Mean Reciprocal Rank (MRR) on WN18 dataset as training progresses.

	MRR	hit@10
ACE(Ent+SC)	<u>0.792</u>	0.945
ACE(Ent+SC+IW)	0.768	0.949
NCE TransD (ours)	0.527	0.947
NCE TransD ((Ji et al., 2015))	-	0.925
KBGAN(DISTMULT) ((Cai and Wang, 2017))	0.772	0.948
KBGAN(COMPLEX) ((Cai and Wang, 2017))	0.779	0.948
Wang et al. ((Wang et al., 2018))	-	0.93
COMPLEX ((Trouillon et al., 2016))	0.941	0.947

Table 4: WN18 experiments: the first portion of the table contains results where the base model is TransD, the last separated line is the COMPLEX embedding model (Trouillon et al., 2016), which achieves the SOTA on this dataset. Among all TransD based models (the best results in this group is underlined), ACE improves over basic NCE and another GAN based approach KBGAN. The gap on MRR is likely due to the difference between TransD and COMPLEX models.

5.5 Hard Negative Analysis

To better understand the effect of the adversarial samples proposed by the generator we plot the discriminator loss on both p_{nce} and g_{θ} samples. In this context, a harder sample means a higher loss assigned by the discriminator. Fig. 4 shows that discriminator loss for the word embedding task on g_{θ} samples are always higher than on p_{nce} samples, confirming that the generator is indeed sampling harder negatives.

For Hypernym Prediction task, Fig.2 shows discriminator loss on negative pairs sampled from NCE and ACE respectively. The higher the loss the harder the negative pair is. As indicated in the left plot, loss on the ACE negative terms collapses

faster than on the NCE negatives. After adding entropy regularization and weight decay, the generator works as expected.

6 Limitations

When the generator softmax is large, the current implementation of ACE training is computationally expensive. Although ACE converges faster per iteration, it may converge more slowly on wall-clock time depending on the cost of the softmax. However, embeddings are typically used as pre-trained building blocks for subsequent tasks. Thus, their learning is usually the pre-computation step for the more complex downstream models and spending more time is justified, especially with GPU acceleration. We believe that the computational cost could potentially be reduced via some existing techniques such as the “augment and reduce” variational inference of (Ruiz et al., 2018), adaptive softmax (Grave et al., 2016), or the “sparsely-gated” softmax of Shazeer et al. (2017), but leave that to future work.

Another limitation is on the theoretical front. As noted in Goodfellow (2014), GAN learning does not implement maximum likelihood estimation (MLE), while NCE has MLE as an asymptotic limit. To the best of our knowledge, more distant connections between GAN and MLE training are not known, and tools for analyzing the equilibrium of a min-max game where players are parametrized by deep neural nets are currently not available to the best of our knowledge.

7 Conclusion

In this paper, we propose Adversarial Contrastive Estimation as a general technique for improving supervised learning problems that learn by contrasting observed and fictitious samples. Specifically, we use a generator network in a conditional GAN like setting to propose hard negative examples for our discriminator model. We find that a mixture distribution of randomly sampling negative examples along with an adaptive negative sampler leads to improved performances on a variety of embedding tasks. We validate our hypothesis that hard negative examples are critical to optimal learning and can be proposed via our ACE framework. Finally, we find that controlling the entropy of the generator through a regularization term and properly handling false negatives is crucial for successful training.

References

- Martin Arjovsky, Soumith Chintala, and Léon Bottou. 2017. Wasserstein GAN. *arXiv preprint arXiv:1701.07875*.
- David Belanger and Andrew McCallum. 2016. Structured prediction energy networks. In *International Conference on Machine Learning*, pages 983–992.
- Antoine Bordes, Xavier Glorot, Jason Weston, and Yoshua Bengio. 2014. A semantic matching energy function for learning with multi-relational data. *Machine Learning*, 94(2):233–259.
- Antoine Bordes, Nicolas Usunier, Alberto Garcia-Duran, Jason Weston, and Oksana Yakhnenko. 2013. Translating embeddings for modeling multi-relational data. In *Advances in neural information processing systems*, pages 2787–2795.
- Liwei Cai and William Yang Wang. 2017. Kbgan: Adversarial learning for knowledge graph embeddings. *arXiv preprint arXiv:1711.04071*.
- Yanshuai Cao, Gavin Weiguang Ding, Kry Yik-Chau Lui, and Ruitong Huang. 2018. [Improving GAN training via binarized representation entropy \(BRE\) regularization](#). In *International Conference on Learning Representations*.
- Bo Dai and Dahua Lin. 2017. Contrastive learning for image captioning. In *Advances in Neural Information Processing Systems*, pages 898–907.
- Tim Dettmers, Pasquale Minervini, Pontus Stenertorp, and Sebastian Riedel. 2017. Convolutional 2d knowledge graph embeddings. *arXiv preprint arXiv:1707.01476*.
- Chris Dyer. 2014. Notes on noise contrastive estimation and negative sampling. *arXiv preprint arXiv:1410.8251*.
- William Fedus, Ian Goodfellow, and Andrew M Dai. 2018. MaskGAN: Better text generation via filling in the .. *arXiv preprint arXiv:1801.07736*.
- Lev Finkelstein, Evgeniy Gabrilovich, Yossi Matias, Ehud Rivlin, Zach Solan, Gadi Wolfman, and Eytan Ruppín. 2001. Placing search in context: The concept revisited. In *Proceedings of the 10th international conference on World Wide Web*, pages 406–414. ACM.
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014a. [Generative adversarial nets](#). In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 27*, pages 2672–2680.
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014b. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680.
- Ian J Goodfellow. 2014. On distinguishability criteria for estimating generative models. *arXiv preprint arXiv:1412.6515*.
- Will Grathwohl, Dami Choi, Yuhuai Wu, Geoff Roeder, and David Duvenaud. 2017. Backpropagation through the void: Optimizing control variates for black-box gradient estimation. *arXiv preprint arXiv:1711.00123*.
- Edouard Grave, Armand Joulin, Moustapha Cissé, David Grangier, and Hervé Jégou. 2016. Efficient softmax approximation for GPUs. *arXiv preprint arXiv:1609.04309*.
- Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron C Courville. 2017. Improved training of wasserstein gans. In *Advances in Neural Information Processing Systems*, pages 5769–5779.
- Michael Gutmann and Aapo Hyvärinen. 2010. Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, pages 297–304.
- Michael U Gutmann and Aapo Hyvärinen. 2012. Noise-contrastive estimation of unnormalized statistical models, with applications to natural image statistics. *Journal of Machine Learning Research*, 13(Feb):307–361.
- Eric Jang, Shixiang Gu, and Ben Poole. 2016. Categorical reparameterization with gumbel-softmax. *arXiv preprint arXiv:1611.01144*.
- Guoliang Ji, Shizhu He, Liheng Xu, Kang Liu, and Jun Zhao. 2015. Knowledge graph embedding via dynamic mapping matrix. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, volume 1, pages 687–696.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Yankai Lin, Zhiyuan Liu, Maosong Sun, Yang Liu, and Xuan Zhu. 2015. Learning entity and relation embeddings for knowledge graph completion. In *AAAI*, volume 15, pages 2181–2187.
- Hao Liu, Yihao Feng, Yi Mao, Dengyong Zhou, Jian Peng, and Qiang Liu. 2018. [Action-dependent control variates for policy optimization via stein identity](#). In *International Conference on Learning Representations*.
- Thang Luong, Richard Socher, and Christopher D Manning. 2013. Better word representations with recursive neural networks for morphology. In *CoNLL*, pages 104–113.

- Chris J Maddison, Andriy Mnih, and Yee Whye Teh. 2016. The concrete distribution: A continuous relaxation of discrete random variables. *arXiv preprint arXiv:1611.00712*.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pages 3111–3119.
- M. Mirza and S. Osindero. 2014. **Conditional Generative Adversarial Nets**. *ArXiv e-prints*.
- Andriy Mnih and Koray Kavukcuoglu. 2013. Learning word embeddings efficiently with noise-contrastive estimation. In *Advances in neural information processing systems*, pages 2265–2273.
- Andriy Mnih and Yee Whye Teh. 2012. A fast and simple algorithm for training neural probabilistic language models. *arXiv preprint arXiv:1206.6426*.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 1532–1543.
- Steven J Rennie, Etienne Marcheret, Youssef Mroueh, Jarret Ross, and Vaibhava Goel. 2016. Self-critical sequence training for image captioning. *arXiv preprint arXiv:1612.00563*.
- Francisco JR Ruiz, Michalis K Titsias, Adji B Dieng, and David M Blei. 2018. Augment and reduce: Stochastic inference for large categorical distributions. *arXiv preprint arXiv:1802.04220*.
- Florian Schroff, Dmitry Kalenichenko, and James Philbin. 2015. Facenet: A unified embedding for face recognition and clustering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 815–823.
- Noam Shazeer, Azalia Mirhoseini, Krzysztof Maziarz, Andy Davis, Quoc Le, Geoffrey Hinton, and Jeff Dean. 2017. Outrageously large neural networks: The sparsely-gated mixture-of-experts layer. *arXiv preprint arXiv:1701.06538*.
- Abhinav Shrivastava, Abhinav Gupta, and Ross Girshick. 2016. Training region-based object detectors with online hard example mining. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 761–769.
- Noah A Smith and Jason Eisner. 2005. Contrastive estimation: Training log-linear models on unlabeled data. In *Proceedings of the 43rd Annual Meeting on Association for Computational Linguistics*, pages 354–362. Association for Computational Linguistics.
- Ben Taskar, Vassil Chatalbashev, Daphne Koller, and Carlos Guestrin. 2005. Learning structured prediction models: A large margin approach. In *Proceedings of the 22nd international conference on Machine learning*, pages 896–903. ACM.
- Théo Trouillon, Johannes Welbl, Sebastian Riedel, Éric Gaussier, and Guillaume Bouchard. 2016. Complex embeddings for simple link prediction. In *International Conference on Machine Learning*, pages 2071–2080.
- Ioannis Tsochantaridis, Thorsten Joachims, Thomas Hofmann, and Yasemin Altun. 2005. Large margin methods for structured and interdependent output variables. *Journal of machine learning research*, 6(Sep):1453–1484.
- Lifu Tu and Kevin Gimpel. 2018. Learning approximate inference networks for structured prediction. In *International Conference on Learning Representations*.
- George Tucker, Andriy Mnih, Chris J Maddison, John Lawson, and Jascha Sohl-Dickstein. 2017. **Rebar: Low-variance, unbiased gradient estimates for discrete latent variable models**. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 2627–2636. Curran Associates, Inc.
- Ashish Vaswani, Yingdong Zhao, Victoria Fossum, and David Chiang. 2013. Decoding with large-scale neural language models improves translation. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1387–1392.
- Ivan Vendrov, Ryan Kiros, Sanja Fidler, and Raquel Urtasun. 2016. Order-embeddings of images and language. In *International Conference on Learning Representations*.
- Peifeng Wang, Shuangyin Li, and Rong Pan. 2018. Incorporating GAN for negative sampling in knowledge representation learning. In *The Thirty-Second AAAI Conference on Artificial Intelligence (AAAI-18)*.
- Zhen Wang, Jianwen Zhang, Jianlin Feng, and Zheng Chen. 2014. Knowledge graph embedding by translating on hyperplanes. In *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence*, pages 1112–1119. AAAI Press.
- Ronald J Williams. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3-4):229–256.
- Bishan Yang, Wen-tau Yih, Xiaodong He, Jianfeng Gao, and Li Deng. 2014. Embedding entities and relations for learning and inference in knowledge bases. *arXiv preprint arXiv:1412.6575*.

Bengio Yoshua, Ducharme Rejean, Vincent Pascal, and Jauvin Christian. 2003. A neural probabilistic language model. *Journal of Machine Learning Research*.

Junbo Zhao, Michael Mathieu, and Yann LeCun. 2016. Energy-based generative adversarial network. *arXiv preprint arXiv:1609.03126*.