

MUSE: Modularizing Unsupervised Sense Embeddings

Guang-He Lee^{*†} Yun-Nung Chen[†]

^{*}Massachusetts Institute of Technology, Cambridge, MA

[†]National Taiwan University, Taipei, Taiwan

guanghe@mit.edu y.v.chen@ieee.org

Abstract

This paper proposes to address the word sense ambiguity issue in an unsupervised manner, where word sense representations are learned along a word sense selection mechanism given contexts. Prior work focused on designing a single model to deliver both mechanisms, and thus suffered from either coarse-grained representation learning or inefficient sense selection. The proposed modular approach, *MUSE*, implements flexible modules to optimize distinct mechanisms, achieving the first purely sense-level representation learning system with linear-time sense selection. We leverage reinforcement learning to enable joint training on the proposed modules, and introduce various exploration techniques on sense selection for better robustness. The experiments on benchmark data show that the proposed approach achieves the state-of-the-art performance on synonym selection as well as on contextual word similarities in terms of MaxSimC.

1 Introduction

Recently, deep learning methodologies have dominated several research areas in natural language processing (NLP), such as machine translation, language understanding, and dialogue systems. However, most of applications usually utilize word-level embeddings to obtain semantics. Considering that natural language is highly ambiguous, the standard word embeddings may suffer from polysemy issues. Neelakantan et al. (2014) pointed out that, due to triangle inequality in vector space, if one word has two different senses but is restricted to *one embedding*, the sum of

the distances between the word and its synonym in each sense would upper-bound the distance between the respective synonyms, which may be mutually irrelevant, in embedding space¹. Due to the theoretical inability to account for polysemy using a single embedding representation per word, multi-sense word representations are proposed to address the ambiguity issue using multiple embedding representations for different senses in a word (Reisinger and Mooney, 2010; Huang et al., 2012).

This paper focuses on unsupervised learning from the unannotated corpus. There are two key mechanisms for a multi-sense word representation system in such scenario: 1) a sense selection (decoding) mechanism infers the most probable sense for a word given its context and 2) a sense representation mechanism learns to embed word senses in a continuous space.

Under this framework, prior work focused on designing a single model to deliver both mechanisms (Neelakantan et al., 2014; Li and Jurafsky, 2015; Qiu et al., 2016). However, the previously proposed models introduce side-effects: 1) mixing word-level and sense-level tokens achieves efficient sense selection but introduces ambiguous word-level tokens during the representation learning process (Neelakantan et al., 2014; Li and Jurafsky, 2015), and 2) pure sense-level tokens prevent ambiguity from word-level tokens but require exponential time complexity when decoding a sense sequence (Qiu et al., 2016).

Unlike the prior work, this paper proposes *MUSE*²—a novel modularization framework incorporating sense selection and representation learning models, which implements flexible modules to optimize distinct mechanisms. Specifically,

¹ $d(\text{rock}, \text{stone}) + d(\text{rock}, \text{shake}) \geq d(\text{stone}, \text{shake})$

²The trained models and code are available at <https://github.com/MiuLab/MUSE>.

MUSE enables linear time sense identity decoding with a *sense selection module* and purely sense-level representation learning with a *sense representation module*.

With the modular design, we propose a novel joint learning algorithm on the modules by connecting to a reinforcement learning scenario, which achieves the following advantages. First, the decision making process under reinforcement learning better captures the sense selection mechanism than probabilistic and clustering methods. Second, our reinforcement learning algorithm realizes the first single objective function for modular unsupervised sense representation systems. Finally, we introduce various exploration techniques under reinforcement learning on sense selection to enhance robustness.

In summary, our contributions are five-fold:

- *MUSE* is the first system that maintains purely sense-level representation learning with linear-time sense decoding.
- We are among the first to leverage reinforcement learning to model the sense selection process in sense representations system.
- We are among the first to propose a single objective for modularized unsupervised sense embedding learning.
- We introduce a sense exploration mechanism for the sense selection module to achieve better flexibility and robustness.
- Our experimental results show the state-of-the-art performance for synonym selection and contextual word similarities in terms of MaxSimC.

2 Related Work

There are three dominant types of approaches for learning multi-sense word representations in the literature: 1) clustering methods, 2) probabilistic modeling methods, and 3) lexical ontology based methods. Our reinforcement learning based approach can be loosely connected to clustering methods and probabilistic modeling methods.

Reisinger and Mooney (2010) first proposed multi-sense word representations on the vector space based on clustering techniques. With the power of deep learning, some work exploited neural networks to learn embeddings with sense selection based on clustering (Huang et al., 2012; Neelakantan et al., 2014). Chen et al. (2014) replaced the clustering procedure with a word sense

disambiguation model using WordNet (Miller, 1995). Kågebäck et al. (2015) and Vu and Parker (2016) further leveraged a weighting mechanism and interactive process in the clustering procedure. Moreover, Guo et al. (2014) leveraged bilingual resources for clustering. However, most of the above approaches separated the clustering procedure and the representation learning procedure without a joint objective, which may suffer from the error propagation issue. Instead, the proposed approach, MUSE, enables joint training on sense selection and representation learning.

Instead of clustering, probabilistic modeling methods have been applied for learning multi-sense embeddings in order to make the sense selection more flexible, where Tian et al. (2014) and Jauhar et al. (2015) conducted probabilistic modeling with EM training. Li and Jurafsky (2015) exploited Chinese Restaurant Process to infer the sense identity. Furthermore, Bartunov et al. (2016) developed a non-parametric Bayesian extension on the skip-gram model (Mikolov et al., 2013b). Despite reasonable modeling on sense selection, all above methods mixed word-level and sense-level tokens during representation learning—unable to conduct representation learning in the pure sense level due to the complicated computation in their EM algorithms.

Recently, Qiu et al. (2016) proposed an EM algorithm to learn purely sense-level representations, where the computational cost is high when decoding the sense identity sequence, because it takes exponential time to search all sense combination within a context window. Our modular design addresses such drawback, where the sense selection module decodes a sense sequence with linear-time complexity, while the sense representation module remains representation learning in the pure sense level.

Unlike a lot of relevant work that requires additional resources such as the lexical ontology (Pilehvar and Collier, 2016; Rothe and Schütze, 2015; Jauhar et al., 2015; Chen et al., 2015; Iacobacci et al., 2015) or bilingual data (Guo et al., 2014; Ettinger et al., 2016; Šuster et al., 2016), which may be unavailable in some language, our model can be trained using only an unlabeled corpus. Also, some prior work proposed to learn topical embeddings and word embeddings jointly in order to consider the contexts (Liu et al., 2015a,b), whereas this paper focuses on learning multi-sense




Figure 1: The *MUSE* architecture with a 3-step learning algorithm: 1) collocation sampling, 2) sense selection for sense representation learning, and 3) optimizing sense selection with a reward signal from sense representation. Reward signal is only passed to the target word to stabilize model training due to directional architecture in the sense representation module.

word embeddings.

3 Proposed Approach: MUSE

This work proposes a framework to modularize two key mechanisms for multi-sense word representations: a *sense selection module* and a *sense representation module*. The sense selection module decides which sense to use given a text context, whereas the sense representation module learns meaningful representations based on its statistical characteristics. Unlike prior work that must suffer from either inefficient sense selection (Qiu et al., 2016) or coarse-grained representation learning (Neelakantan et al., 2014; Li and Jurafsky, 2015; Bartunov et al., 2016), the proposed modularized framework is capable of performing efficient sense selection and learning representations in pure sense level simultaneously.

To learn sense-level representations, a sense selection model should be first established for sense identity decoding. On the other hand, the sense embeddings should guide the sense selection model when decoding a sense identity sequence. Therefore, these two modules should be tangled. This indicates that a naive two-stage algorithm or two separate learning algorithms proposed by prior work are not optimal.

By connecting the proposed formulation with reinforcement learning literature, we design a novel joint training algorithm. Besides, taking advantage of the form of reinforcement learning, we are among the first to investigate various exploration techniques in sense selection for unsuper-

vised sense embedding learning.

3.1 Model Architecture

Our model architecture is illustrated in Figure 1, where there are two modules in optimization.

3.1.1 Sense Selection Module

Formally speaking, given a corpus C , vocabulary W , and the t -th word $C_t = w_i \in W$, we would like to find the most probable sense $z_{ik} \in Z_i$, where Z_i is the set of senses in word w_i . Assuming that a word sense is determined by the local context, we exploit a local context $\bar{C}_t = \{C_{t-m}, \dots, C_{t+m}\}$ for sense selection according to the Markov assumption, where m is the size of a context window. Then we can either formulate a probabilistic policy $\pi(z_{ik} | \bar{C}_t)$ about sense selection or estimate the *individual* likelihood $q(z_{ik} | \bar{C}_t)$ for each sense identity.

To ensure efficiency, here we exploit a linear neural architecture that takes word-level input tokens and outputs sense-level identities. The architecture is similar to continuous bag-of-words (CBOW) (Mikolov et al., 2013a). Specifically, given a *word* embedding matrix P , the local context can be modeled as the summation of word embeddings from its context \bar{C}_t . The output can be formulated with a 3-mode tensor Q , whose dimensions denote words, senses, and latent variables. Then we can model $\pi(z_{ik} | \bar{C}_t)$ or $q(z_{ik} | \bar{C}_t)$ correspondingly. Here we model $\pi(\cdot)$ as a *categorical*

distribution using a softmax layer:

$$\pi(z_{ik} | \bar{C}_t) = \frac{\exp(Q_{ik}^T \sum_{j \in \bar{C}_t} P_j)}{\sum_{k' \in Z_i} \exp(Q_{ik'}^T \sum_{j \in \bar{C}_t} P_j)}. \quad (1)$$

On the other hand, the likelihood of selecting *distinct* sense identities, $q(z_{ik} | \bar{C}_t)$, is modeled as a *Bernoulli* distribution with a sigmoid function $\sigma(\cdot)$:

$$q(z_{ik} | \bar{C}_t) = \sigma(Q_{ik}^T \sum_{j \in \bar{C}_t} P_j). \quad (2)$$

Different modeling approaches require different learning methods, especially for the unsupervised setting. We leave the corresponding learning algorithms in § 3.2. Finally, with a built sense selection module, we can apply any selection algorithm such as a greedy selection strategy to infer the sense identity z_{ik} given a word w_i with its context C_t .

We note that modularized model enables efficient sense selection by leveraging word-level tokens, while remaining purely sense-level tokens in the representation module. Specifically, if n denotes $\max_k |Z_k|$, decoding L words takes $O(nL)$ senses to be searched due to independent sense selection. The prior work using a single model with purely sense-level tokens (Qiu et al., 2016) requires exponential time to calculate the collocation energy for every possible combination of sense identities within a context window, $O(n^{2m})$, for a *single target sense*. Further, Qiu et al. (2016) took an additional sequence decoding step with quadratic time complexity $O(n^{4m}L)$, based on an exponential number n^{2m} in the base unit. It demonstrates the achievement about sense inference efficiency in our proposed model.

3.1.2 Sense Representation Module

The semantic representation learning is typically formulated as a maximum likelihood estimation (MLE) problem for collocation likelihood. In this paper, we use the skip-gram formulation (Mikolov et al., 2013b) considering that it requires less training time, where only two sense identities are required for stochastic training. Other popular candidates, like GloVe (Pennington et al., 2014) and CBOW (Mikolov et al., 2013a), require more sense identities to be selected as input and thus not suitable for our scenario. For example, GloVe (Pennington et al., 2014) takes computationally expensive collocation *counting* statistics

for each token in a corpus as input, which requires sense selection for every occurrence of the target word across the whole corpus for a single optimization step.

To learn the representations, we first create input sense representation matrix U and collocation estimation matrix V as the learning targets. Given a target word w_i and collocated word w_j with corresponding local contexts, we map them to their sense identities as z_{ik} and z_{jl} by the sense selection module, and maximize the sense collocation log likelihood $\log \mathcal{L}(\cdot)$. A natural choice of the likelihood function is formulated as a categorical distribution over all possible collocated senses given the target sense z_{ik} :

$$\max_{U,V} \log \mathcal{L}(z_{jl} | z_{ik}) = \log \frac{\exp(U_{z_{ik}}^T V_{z_{jl}})}{\sum_{z_{uv}} \exp(U_{z_{ik}}^T V_{z_{uv}})}. \quad (3)$$

Instead of enumerating all possible collocated senses which is computationally expensive, we use the skip-gram objective (4) (Mikolov et al., 2013b) to approximate (3) as shown in the green block of Figure 1.

$$\max_{U,V} \log \bar{\mathcal{L}}(z_{jl} | z_{ik}) = \log \sigma(U_{z_{ik}}^T V_{z_{jl}}) \quad (4)$$

$$+ \sum_{v=1}^M \mathbb{E}_{z_{uv} \sim p_{neg}(z)} [\log \sigma(-U_{z_{ik}}^T V_{z_{uv}})],$$

where $p_{neg}(z)$ is the distribution over all senses for negative samples. In our experiment with $|Z_i|$ senses for word w_i , we use $(1/|Z_i|)$ word-level unigram as sense-level unigram for efficiency and the 3/4-th power trick in Mikolov et al. (2013b).

We note that our modular framework can easily maintain purely sense-level tokens with an arbitrary representation learning model. In contrast, most related work using probabilistic modeling (Tian et al., 2014; Jauhar et al., 2015; Li and Jurafsky, 2015; Bartunov et al., 2016) binded sense representations with the sense selection mechanism, so efficient sense selection by leveraging word-level tokens can be achieved only at the cost of mixing word-level and sense-level tokens in their representation learning process.

3.2 Learning

Without the supervised signal for the proposed modules, it is desirable to connect two modules in a way where they can improve each other by their own estimations. First, a trivial way is to forward the prediction of the sense selection module

to the representation module. Then we cast the estimated collocation likelihood as a reward signal for the selected sense for effective learning.

To realize the above procedure, we cast the learning problem a one-step Markov decision process (MDP) (Sutton and Barto, 1998), where the state, action, and reward correspond to context \bar{C}_t , sense z_{ik} , and collocation log likelihood $\log \bar{\mathcal{L}}(\cdot)$, respectively. Based on different modeling methods ((1) or (2)) in the sense selection module, we connect the model to respective reinforcement learning algorithms to solve the MDP. Specifically, we refer (1) to policy distribution and refer (2) to Q-value estimation in the reinforcement learning literature.

The proposed MDP framework embodies several nuances of sense selection. First, the decision of a word sense is Markov: taking the whole corpus into consideration is not more helpful than a handful of necessary local contexts. Second, the decision making in MDP exploits a hard decision for selecting sense identity, which captures the sense selection process more naturally than a joint probability distribution among senses (Qiu et al., 2016). Finally, we exploit the reward mechanism in MDP to enable joint training: the estimation of sense representation is treated as a reward signal to guide sense selection. In contrast, the decision making under clustering (Huang et al., 2012; Neelakantan et al., 2014) considers the similarity within clusters instead of the outcome of a decision using a reward signal as MDP.

3.2.1 Policy Gradient Method

Because (1) fits a valid probability distribution, an intuitive optimization target is the expectation of resulting collocation likelihood among each sense. In addition, as the skip-gram formulation in (4) is unidirectional ($\bar{\mathcal{L}}(z_{ik} | z_{jl}) \neq \bar{\mathcal{L}}(z_{jl} | z_{ik})$), we perform one-side optimization for the target sense z_{ik} to stabilize model training³. That is, for the target word w_i and the collocated word w_j given respective contexts \bar{C}_t and $\bar{C}_{t'}$ ($0 < |t - t'| \leq m$), we first draw a sense z_{jl} for w_j from the policy $\pi(\cdot | \bar{C}_{t'})$ and optimize the expected collocation likelihood for the target sense z_{ik} as follows,

$$\max_{P, Q} \mathbb{E}_{z_{ik} \sim \pi(\cdot | \bar{C}_t)} [\log \bar{\mathcal{L}}(z_{jl} | z_{ik})]. \quad (5)$$

Note that (4) can be merged into (5) as a single objective. The objective is differentiable and

³We observe about 4% performance drop by optimizing input selection z_{ik} and output selection z_{jl} simultaneously.

supports stochastic optimization (Lei et al., 2016), which uses a stochastic sample z_{ik} for optimization.

However, there are two possible disadvantages in this formulation. First, because the policy assumes the probability distribution in (1), optimizing the selected sense must affect the estimation of the other senses. Second, if applying stochastic gradient ascent to optimizing (5), it would always lower the probability estimation for the selected sense z_{ik} even if the model accurately selects the right sense. The detailed proof is in Appendix A.

3.2.2 Value-Based Method

To address the above issues, we apply the Q-learning algorithm (Mnih et al., 2013). Instead of maintaining a probabilistic policy for sense selection, Q-learning estimates the Q-value (resulting collocation log likelihood) for each sense candidate directly and independently. Thus, the estimation of unselected senses may not be influenced by the selected one. Note that in one-step MDP, the reward is equivalent to the Q-value, so we will use reward and Q-value interchangeably, hereinafter, based on the context.

We further follow the convention of recent neural reinforcement learning by reducing the reward range to aid model training (Mnih et al., 2013). Specifically, we replace the *log likelihood* $\log \bar{\mathcal{L}}(\cdot) \in (-\infty, 0]$ with the *likelihood* $\bar{\mathcal{L}}(\cdot) \in [0, 1]$ as the reward function. Due to the monotonic operation in $\log(\cdot)$, the relative ordering of the reward remains the same.

Furthermore, we exploit the probabilistic nature of likelihood for Q-learning. To elaborate, as Q-learning is used to approximate the Q-value for each action in typical reinforcement learning, most literature adopted square loss to characterize the discrepancy between the target and estimated Q-values (Mnih et al., 2013). In our setting where the Q-value/reward is a likelihood function, our model exploits cross-entropy loss to better capture the characteristics of probability distribution.

Given that the collocation likelihood in (4) is an *approximation* to the original categorical distribution with a softmax function shown in (3) (Mikolov et al., 2013b), we revise the formulation by omitting the negative sampling term. The resulting formulation $\hat{\mathcal{L}}(\cdot)$ is a Bernoulli distribution indicating whether z_{jl} collocates or not given z_{ik} :

$$\hat{\mathcal{L}}(z_{jl} | z_{ik}) = \sigma(U_{z_{ik}}^T V_{z_{jl}}). \quad (6)$$

There are three advantages about using $\hat{\mathcal{L}}(\cdot)$ instead of approximated $\tilde{\mathcal{L}}(\cdot)$ and original $\mathcal{L}(\cdot)$. First, regarding the variance of estimation, $\hat{\mathcal{L}}(\cdot)$ better captures $\mathcal{L}(\cdot)$ than $\tilde{\mathcal{L}}(\cdot)$ because $\tilde{\mathcal{L}}(\cdot)$ involves sampling:

$$\text{Var}(\tilde{\mathcal{L}}(\cdot)) \geq \text{Var}(\hat{\mathcal{L}}(\cdot)) = \text{Var}(\mathcal{L}(\cdot)) = 0. \quad (7)$$

Second, regarding the relative ordering of estimation, for any two collocated senses z_{jl} and $z_{j'l'}$ with a target sense z_{ik} , the following equivalence holds:

$$\begin{aligned} \mathcal{L}(z_{jl} | z_{ik}) &< \mathcal{L}(z_{j'l'} | z_{ik}) & (8) \\ \Leftrightarrow \tilde{\mathcal{L}}(z_{jl} | z_{ik}) &< \tilde{\mathcal{L}}(z_{j'l'} | z_{ik}) \\ \Leftrightarrow \hat{\mathcal{L}}(z_{jl} | z_{ik}) &< \hat{\mathcal{L}}(z_{j'l'} | z_{ik}) \end{aligned}$$

Third, for collocation computation, $\mathcal{L}(\cdot)$ requires all sense identities and $\tilde{\mathcal{L}}(\cdot)$ requires $(M+1)$ sense identities, whereas $\hat{\mathcal{L}}(\cdot)$ only requires 1 sense identity. In sum, the proposed $\hat{\mathcal{L}}(\cdot)$ approximates $\mathcal{L}(\cdot)$ with no variance, no ‘‘bias’’ (in terms of relative ordering), and significantly less computation.

Finally, because both target distribution $\hat{\mathcal{L}}(\cdot)$ and estimated distribution $q(\cdot)$ in (2) are Bernoulli distributions, we follow the last section to conduct one-side optimization by fixing a collocated sense z_{jl} and optimize the selected sense z_{ik} with cross entropy as

$$\min_{P,Q} H(\hat{\mathcal{L}}(z_{ik} | z_{jl}), q(z_{ik} | \bar{C}_t)). \quad (9)$$

3.2.3 Joint Training

To jointly train sense selection and sense representation modules, we first select a pair of the collocated senses, z_{ik} and z_{jl} , based on the sense selection module with any selecting strategy (e.g. greedy), and then optimize the sense representation module and the sense selection module using the above derivations. Algorithm 1 describes the proposed MUSE model training procedure.

As modular frameworks, the major distinction between our modular framework and two-stage clustering-representation learning framework (Neelakantan et al., 2014; Vu and Parker, 2016) is that we establish a reward signal from the sense representation to the sense selection module to enable immediate and joint optimization.

3.3 Sense Selection Strategy

Given a fitness estimation for each sense, exploiting the greedy sense is the most popular strategy for clustering algorithms (Neelakantan et al.,

Algorithm 1: Learning Algorithm

```

for  $w_i = C_t \in C$  do
  sample  $w_j = C_{t'}$  ( $0 < |t' - t| \leq m$ );
   $z_{ik} = \text{select}(C_t, w_i)$ ;
   $z_{jl} = \text{select}(C_{t'}, w_j)$ ;
  optimize  $U, V$  by (4) for the sense
  representation module;
  optimize  $P, Q$  by (5) or (9) for the sense
  selection module;

```

2014; Kågebäck et al., 2015) and hard-EM algorithms (Qiu et al., 2016; Jauhar et al., 2015) in literature. However, there are two incentives to conduct exploration. First, in the early training stage when the fitness is not well estimated, it is desirable to explore underestimated senses. Second, due to high ambiguity in natural language, sometimes multiple senses in a word would fit in the same context. The dilemma between exploring sub-optimal choices and exploiting the optimal choice is called exploration-exploitation trade-off in reinforcement learning (Sutton and Barto, 1998).

We introduce exploration mechanisms for sense selection for both policy gradient and Q-learning. For policy gradient, we sample the policy distribution to approximate the expectation in (5). Because of the flexible formulation of Q-learning, the following classic exploration mechanisms are applied to sense selection:

- *Greedy*: selects the sense with the largest Q-value (no exploration).
- ϵ -*Greedy*: selects a random sense with ϵ probability, and adopts the greedy strategy otherwise (Mnih et al., 2013).
- *Boltzmann*: samples the sense based on the Boltzmann distribution modeled by Q-value. We directly use (1) as the Boltzmann distribution for simplicity.

We note that Q-learning with Boltzmann sampling yields the same sampling process as policy gradient but different optimization objectives. To our best knowledge, we are among the first to explore several exploration strategies for unsupervised sense embedding learning.

In the following sections, MUSE-Policy denotes the proposed MUSE model with policy learning and MUSE-Greedy denotes the model using corresponding sense selection strategy for Q-learning.

4 Experiments

We evaluate our proposed MUSE model in both quantitative and qualitative experiments.

4.1 Experimental Setup

Our model is trained on the April 2010 Wikipedia dump (Shaoul and Westbury, 2010), which contains approximately 1 billion tokens. For fair comparison, we adopt the same vocabulary set as Huang et al. (2012) and Neelakantan et al. (2014). For preprocessing, we convert all words to their lower cases, apply the Stanford tokenizer and the Stanford sentence tokenizer (Manning et al., 2014), and remove all sentences with less than 10 tokens. The number of senses per word in Q is set to 3 as the prior work (Neelakantan et al., 2014).

In the experiments, the context window size is set to 5 ($|\bar{C}_t| = 11$). Subsampling technique introduced by word2vec (Mikolov et al., 2013b) is applied to accelerate the training process. The learning rate is set to 0.025. The embedding dimension is 300. We initialize Q and V as zeros, and P and U from uniform distribution $[-\sqrt{1/100}, \sqrt{1/100}]$ such that each embedding has unit length in expectation (Lei et al., 2015). Our model uses 25 negative senses for negative sampling in (4). We use $\epsilon = 5\%$ for ϵ -Greedy sense selection strategy

In optimization, we conduct mini-batch training with 2048 batch size using the following procedure: 1) select senses in the batch; 2) optimize U, V using stochastic training within the batch for efficiency; 3) optimize P, Q using mini-batch training for robustness.

4.2 Experiment 1: Contextual Word Similarity

To evaluate the quality of the learned sense embeddings, we compute the similarity score between each word pair given their respective local contexts and compare with the human-judged score using Stanford’s Contextual Word Similarities (SCWS) dataset (Huang et al., 2012). Specifically, given a list of word pairs with corresponding contexts, $S = \{(w_i, \bar{C}_t, w_j, \bar{C}_{t'})\}$, we calculate the Spearman’s rank correlation ρ between human-judged similarity and model similarity estimations⁴. Two major contextual similarity esti-

⁴For example, human-judged similarity between “... east bank of the Des Moines River ...” and “... basis of all money laundering ...” is 2.5 out of 10.0 in SCWS dataset (Huang

Method	MaxSimC	AvgSimC
Huang et al. (2012)	26.1	65.7
Neelakantan et al. (2014)	60.1	69.3
Tian et al. (2014)	63.6	65.4
Li and Jurafsky (2015)	66.6	66.8
Bartunov et al. (2016)	53.8	61.2
Qiu et al. (2016)	64.9	66.1
MUSE-Policy	66.1	67.4
MUSE-Greedy	66.3	68.3
MUSE- ϵ -Greedy	67.4 [†]	68.6
MUSE-Boltzmann	67.9[†]	68.7

Table 1: Spearman’s rank correlation $\rho \times 100$ on the SCWS dataset. [†] denotes superior performance to all unsupervised competitors.

mations are introduced by Reisinger and Mooney (2010): AvgSimC and MaxSimC. AvgSimC is a *soft* measurement that addresses the contextual information with a probability estimation:

$$\text{AvgSimC}(w_i, \bar{C}_t, w_j, \bar{C}_{t'}) = \frac{|Z_i| |Z_j|}{\sum_{k=1}^{|Z_i|} \sum_{l=1}^{|Z_j|} \pi(z_{ik}|\bar{C}_t)\pi(z_{jl}|\bar{C}_{t'})} d(z_{ik}, z_{jl}),$$

where $d(z_{ik}, z_{jl})$ refers to the cosine similarity between $U_{z_{ik}}$ and $U_{z_{jl}}$. AvgSimC weights the similarity measurement of each sense pair z_{ik} and z_{jl} by their probability estimations. On the other hand, MaxSimC is a *hard* measurement that only considers the most probable senses:

$$\begin{aligned} \text{MaxSimC}(w_i, \bar{C}_t, w_j, \bar{C}_{t'}) &= d(z_{ik}, z_{jl}), \\ z_{ik} &= \arg \max_{z_{ik'}} \pi(z_{ik'}|\bar{C}_t), \\ z_{jl} &= \arg \max_{z_{jl'}} \pi(z_{jl'}|\bar{C}_{t'}). \end{aligned}$$

The baselines for comparison include classic clustering methods (Huang et al., 2012; Neelakantan et al., 2014), EM algorithms (Tian et al., 2014; Qiu et al., 2016; Bartunov et al., 2016), and Chinese Restaurant Process (Li and Jurafsky, 2015)⁵, where all approaches are trained on the same corpus except Qiu et al. (2016) used more recent Wikipedia dumps. The embedding sizes of all baselines are 300, except 50 in Huang et al. (2012). For every competitor with multiple settings, we report the best performance in each similarity measurement setting and show in Table 1.

et al., 2012).

⁵We run Li and Jurafsky (2015)’s released code on our corpus for fair comparison.

Method	ESL-50	RD-300	TOEFL-80
<i>1) Conventional Word Embedding</i>			
Global Context	47.73	45.07	60.87
Skip-Gram	52.08	55.66	66.67
<i>2) Word Sense Disambiguation</i>			
IMS+SG	41.67	53.77	66.67
<i>3) Unsupervised Sense Embeddings</i>			
EM	27.08	33.96	40.00
MSSG	57.14	58.93	78.26
CRP	50.00	55.36	82.61
MUSE-Policy	52.38	51.79	79.71
MUSE-Greedy	57.14	58.93	79.71
MUSE- ϵ -Greedy	61.90 [†]	62.50 [†]	84.06 [†]
MUSE-Boltzmann	64.29[†]	66.07[†]	88.41[†]
<i>4) Supervised Sense Embeddings</i>			
Retro-GC	63.64	66.20	71.01
Retro-SG	56.25	65.09	73.33

Table 2: Accuracy on synonym selection. [†] denotes superior performance to all unsupervised competitors.

Our MUSE model achieves the state-of-the-art performance on MaxSimC, demonstrating superior quality on independent sense embeddings. On the other hand, MUSE achieves comparable performance with the best competitor in terms of AvgSimC (68.7 vs. 69.3), while MUSE outperforms the same competitor significantly in terms of MaxSimC (67.9 vs. 60.1). The results demonstrate not only the high quality of sense representations but also accurate sense selection.

From the application perspective, MaxSimC refers to a typical scenario using single embedding per word, while AvgSimC employs multiple sense vectors simultaneously per word, which not only brings computational overhead but changes existing neural architecture for NLP. Hence, we argue that MaxSimC better characterize practical usage of a sense representation system than AvgSimC.

Among various learning methods for MUSE, policy gradient performs worst, echoing our argument in § 3.2.1. On the other hand, the superior performance of Boltzmann sampling and ϵ -Greedy over Greedy selection demonstrates the effectiveness of exploration.

Finally, replacing $\bar{\mathcal{L}}(\cdot)$ with $\hat{\mathcal{L}}(\cdot)$ as the reward signal yields 2.3 times speedup for MUSE- ϵ -Greedy and 1.3 times speedup for MUSE-Boltzmann to reach 67.0 in MaxSimC, which demonstrates the efficacy of proposed approxima-

tion $\hat{\mathcal{L}}(\cdot)$ over typical $\bar{\mathcal{L}}(\cdot)$ in terms of convergence.

4.3 Experiment 2: Synonym Selection

We further evaluate our model on synonym selection using multi-sense word representations (Jauhar et al., 2015). Three standard synonym selection datasets, ESL-50 (Turney, 2001), RD-300 (Jarmasz and Szpakowicz, 2004), and TOEFL-80 (Landauer and Dumais, 1997), are performed. In the datasets, each question consists of a question word w_Q and four answer candidates $\{w_A, w_B, w_C, w_D\}$, and the goal is to select the most semantically synonymous choice among the four candidates. For example, in the TOEFL-80 dataset, a question shows $\{(Q) \text{ enormously, (A) appropriately, (B) uniquely, (C) tremendously, (D) decidedly}\}$, and the answer is (C). For multi-sense representations system, it selects the synonym of the question word w_Q using the maximum sense-level cosine similarity as a proxy of the semantic similarity (Jauhar et al., 2015).

Our model is compared with the following baselines: 1) conventional word embeddings: global context vectors (Huang et al., 2012) and skip-gram (Mikolov et al., 2013b); 2) applying supervised word sense disambiguation using the IMS system and then applying skip-gram on disambiguated corpus (IMS+SG) (Zhong and Ng, 2010); 3) unsupervised sense embeddings: EM algorithm (Jauhar et al., 2015), multi-sense skip-gram (MSSG) (Neelakantan et al., 2014), Chinese restaurant process (CRP) (Li and Jurafsky, 2015), and the MUSE models; 4) supervised sense embeddings with WordNet (Miller, 1995): retrofitting global context vectors (Retro-GC) and retrofitting skip-gram (Retro-SG) (Jauhar et al., 2015).

Among unsupervised sense embedding approaches, CRP and MSSG refer to the baselines with highest MaxSimC and AvgSimC in Table 1 respectively. Here we report the setting for baselines based on the best average performance in this task. We also show the performance of supervised sense embeddings as an upperbound of unsupervised methods due to the usage of additional supervised information from WordNet.

The results are shown in Table 2, where our MUSE- ϵ -Greedy and MUSE-Boltzmann significantly outperform all unsupervised sense embeddings methods, echoing the superior quality of our

Context	k-NN Senses
... braves finish the season in tie with the los angeles dodgers ...	scoreless otl shootout 6-6 hingis 3-3 7-7 0-0
... his later years proudly wore tie with the chinese characters for ...	pants trousers shirt juventus blazer socks anfield
... of the mulberry or the blackberry and minos sent him to ...	cranberries maple vaccinium apricot apple
... of the large number of blackberry users in the us federal ...	smartphones sap microsoft ipv6 smartphone
... shells and/or high explosive squash head hesh and/or anti-tank ...	venter thorax neck spear millimeters fusiform
... head was shaven to prevent head lice serious threat back then ...	shaved thatcher loki thorax mao luthor chest
... appoint john pope republican as head of the new army of ...	multi-party appoints unicameral beria appointed

Table 3: Different word senses are selected by MUSE according to different contexts. The respective k-NN (sorted by collocation likelihood) senses are shown to indicate respective semantic meanings.

sense vectors in last section. MUSE-Boltzmann also outperforms the supervised sense embeddings except 1 setting without any supervised signal during training. Finally, the MUSE methods with proper exploration outperform all unsupervised baselines consistently, demonstrating the importance of exploration.

4.4 Qualitative Analysis

We further conduct qualitative analysis to check the semantic meanings of different senses learned by MUSE with k-nearest neighbors (k-NN) using sense representations. In addition, we provide contexts in the training corpus where the sense will be selected to validate the sense selection module. Table 3 shows the results. The learned sense embeddings of the words “tie”, “blackberry”, and “head” clearly correspond to correct senses under different contexts.

Since we address an unsupervised setting that learns sense embeddings from unannotated corpus, the discovered senses highly depend on the training corpus. From our manual inspection, it is common for our model to discover only two senses in a word, like “tie” and “blackberry”. However, we maintain our effort in developing unsupervised sense embeddings learning methods in this work, and the number of discovered sense is not a focus.

5 Conclusion

This paper proposes a novel modularized framework for unsupervised sense representation learning, which supports not only the flexible design of modular tasks but also joint optimization among modules. The proposed model is the first work that implements purely sense-level representation learning with linear-time sense selection, and achieves the state-of-the-art performance on benchmark contextual word similarity and syn-

onym selection tasks. In the future, we plan to investigate reinforcement learning methods to incorporate multi-sense word representations for downstream NLP tasks.

Acknowledgements

We would like to thank reviewers for their insightful comments on the paper. The authors are supported by the Ministry of Science and Technology of Taiwan under the contract number 105-2218-E-002-033, Institute for Information Industry, and MediaTek Inc..

References

- Sergey Bartunov, Dmitry Kondrashkin, Anton Osokin, and Dmitry Vetrov. 2016. Breaking sticks and ambiguities with adaptive skip-gram. *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics*, page 130–138.
- Tao Chen, Ruifeng Xu, Yulan He, and Xuan Wang. 2015. Improving distributed representation of word sense via wordnet gloss composition and context clustering. *Association for Computational Linguistics*.
- Xinxiong Chen, Zhiyuan Liu, and Maosong Sun. 2014. A unified model for word sense representation and disambiguation. In *EMNLP*, pages 1025–1035. Citeseer.
- Allyson Ettinger, Philip Resnik, and Marine Carpuat. 2016. Retrofitting sense-specific word vectors using parallel text. In *Proceedings of NAACL-HLT*, pages 1378–1383.
- Jiang Guo, Wanxiang Che, Haifeng Wang, and Ting Liu. 2014. Learning sense-specific word embeddings by exploiting bilingual resources. In *COLING*, pages 497–507.
- Eric H. Huang, Richard Socher, Christopher D. Manning, and Andrew Y. Ng. 2012. Improving Word Representations via Global Context and Multiple

- Word Prototypes. In *Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Ignacio Iacobacci, Mohammad Taher Pilehvar, and Roberto Navigli. 2015. Senseembed: Learning sense embeddings for word and relational similarity. In *ACL (1)*, pages 95–105.
- Mario Jarmasz and Stan Szpakowicz. 2004. Roget’s thesaurus and semantic similarity. *Recent Advances in Natural Language Processing III: Selected Papers from RANLP*, 2003:111.
- Sujay Kumar Jauhar, Chris Dyer, and Eduard H Hovy. 2015. Ontologically grounded multi-sense representation learning for semantic vector space models. In *HLT-NAACL*, pages 683–693.
- Mikael Kågebäck, Fredrik Johansson, Richard Johansson, and Devdatt Dubhashi. 2015. Neural context embeddings for automatic discovery of word senses. In *Proceedings of NAACL-HLT*, pages 25–32.
- Thomas K Landauer and Susan T Dumais. 1997. A solution to plato’s problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological review*, 104(2):211.
- Tao Lei, Regina Barzilay, and Tommi Jaakkola. 2015. Molding cnns for text: non-linear, non-consecutive convolutions. *EMNLP*.
- Tao Lei, Regina Barzilay, and Tommi Jaakkola. 2016. Rationalizing neural predictions. *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*.
- Jiwei Li and Dan Jurafsky. 2015. Do multi-sense embeddings improve natural language understanding? *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1722–1732.
- Pengfei Liu, Xipeng Qiu, and Xuanjing Huang. 2015a. Learning context-sensitive word embeddings with neural tensor skip-gram model. In *IJCAI*, pages 1284–1290.
- Yang Liu, Zhiyuan Liu, Tat-Seng Chua, and Maosong Sun. 2015b. Topical word embeddings. In *AAAI*, pages 2418–2424.
- Christopher D. Manning, Mihai Surdeanu, John Bauer, Jenny Finkel, Steven J. Bethard, and David McClosky. 2014. [The Stanford CoreNLP natural language processing toolkit](#). In *Association for Computational Linguistics (ACL) System Demonstrations*, pages 55–60.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013a. Efficient estimation of word representations in vector space. *Proceedings of Workshop at ICLR*.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013b. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pages 3111–3119.
- George A Miller. 1995. Wordnet: a lexical database for english. *Communications of the ACM*, 38(11):39–41.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. 2013. Playing atari with deep reinforcement learning. *NIPS Deep Learning Workshop*.
- Arvind Neelakantan, Jeevan Shankar, Alexandre Passos, and Andrew McCallum. 2014. Efficient non-parametric estimation of multiple embeddings per word in vector space. *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing*.
- Jeffrey Pennington, Richard Socher, and Christopher D Manning. 2014. Glove: Global vectors for word representation. volume 14, pages 1532–1543.
- Mohammad Taher Pilehvar and Nigel Collier. 2016. De-conflated semantic representations. *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*.
- Lin Qiu, Kewei Tu, and Yong Yu. 2016. Context-dependent sense embedding. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*.
- Joseph Reisinger and Raymond J Mooney. 2010. Multi-prototype vector-space models of word meaning. In *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, pages 109–117. Association for Computational Linguistics.
- Sascha Rothe and Hinrich Schütze. 2015. Autoextend: Extending word embeddings to embeddings for synsets and lexemes. In *Proceedings of the ACL*.
- Cyrus Shaoul and Chris Westbury. 2010. The westbury lab wikipedia corpus.
- Simon Šuster, Ivan Titov, and Gertjan van Noord. 2016. Bilingual learning of multi-sense embeddings with discrete autoencoders. *NAACL-HLT 2016*.
- Richard S Sutton and Andrew G Barto. 1998. *Reinforcement learning: An introduction*, volume 1. MIT press Cambridge.
- Fei Tian, Hanjun Dai, Jiang Bian, Bin Gao, Rui Zhang, Enhong Chen, and Tie-Yan Liu. 2014. A probabilistic model for learning multi-prototype word embeddings. In *COLING*, pages 151–160.

Peter D Turney. 2001. Mining the web for synonyms: Pmi-ir versus lsa on toefl. In *European Conference on Machine Learning*, pages 491–502. Springer.

Thuy Vu and D Stott Parker. 2016. K-embeddings: Learning conceptual embeddings for words using context. In *Proceedings of NAACL-HLT*, pages 1262–1267.

Zhi Zhong and Hwee Tou Ng. 2010. It makes sense: A wide-coverage word sense disambiguation system for free text. In *Proceedings of the ACL 2010 System Demonstrations*, pages 78–83. Association for Computational Linguistics.

A Doubly Stochastic Gradient

To derive doubly stochastic gradient for equation (5), we first denote (5) as $J(\Theta)$ with $\Theta = \{P, Q\}$ and resolve the expectation form as:

$$\begin{aligned} J(\theta) &= \mathbb{E}_{z_{ik} \sim \pi(\cdot | \bar{C}_t)} [\log \bar{\mathcal{L}}(z_{jl} | z_{ik})] \\ &= \sum_k \pi(z_{ik} | \bar{C}_t) \log \bar{\mathcal{L}}(z_{jl} | z_{ik}). \end{aligned}$$

Denote $\Theta = \{P, Q\}$ as the parameter set for policy π . The gradient with respect to Θ should be:

$$\begin{aligned} \frac{\partial J(\theta)}{\partial \Theta} &= \frac{\partial}{\partial \Theta} \sum_k \pi(z_{ik} | \bar{C}_t) \log \bar{\mathcal{L}}(z_{jl} | z_{ik}) \\ &= \sum_k \log \bar{\mathcal{L}}(z_{jl} | z_{ik}) \frac{\partial \pi(z_{ik} | \bar{C}_t)}{\partial \Theta} \\ &= \sum_k \log \bar{\mathcal{L}}(z_{jl} | z_{ik}) \left(\frac{\partial \log \pi(z_{ik} | \bar{C}_t)}{\partial \Theta} \right) (\pi(z_{ik} | \bar{C}_t)) \\ &= \mathbb{E}_{z_{ik} \sim \pi(\cdot | \bar{C}_t)} \left[\log \bar{\mathcal{L}}(z_{jl} | z_{ik}) \frac{\partial \log \pi(z_{ik} | \bar{C}_t)}{\partial \Theta} \right] \end{aligned}$$

Accordingly, if we conduct typical stochastic gradient ascent training on $J(\Theta)$ with respect to Θ from samples z_{ik} with a learning rate η , the update formula will be:

$$\Theta = \Theta + \eta \log \bar{\mathcal{L}}(z_{jl} | z_{ik}) \frac{\partial \log \pi(z_{ik} | \bar{C}_t)}{\partial \Theta}.$$

However, the collocation log likelihood should always be non-positive: $\log \bar{\mathcal{L}}(z_{jl} | z_{ik}) \leq 0$. Therefore, as long as the collocation log likelihood $\log \bar{\mathcal{L}}(z_{jl} | z_{ik})$ is negative, the update formula is to minimize the likelihood of choosing z_{ik} , despite the fact that z_{ik} may be good choices. On the other hand, if the log likelihood reaches 0, according to (4), it indicates:

$$\begin{aligned} \log \bar{\mathcal{L}}(z_{jl} | z_{ik}) = 0 &\Rightarrow \bar{\mathcal{L}}(z_{jl} | z_{ik}) = 1 \\ \Rightarrow U_{z_{ik}}^T V_{z_{jl}} &\rightarrow \infty, U_{z_{ik}}^T V_{z_{uv}} \rightarrow \infty, \forall z_{uv}, \end{aligned}$$

which leads to computational overflow from an infinity value.