

The Pupil Has Become the Master: Teacher-Student Model-Based Word Embedding Distillation with Ensemble Learning

Bonggun Shin¹, Hao Yang², Jinho D. Choi¹

¹Department of Computer Science, Emory University, Atlanta, GA

²Visa Research, Palo Alto, CA

bonggun.shin@emory.edu, haoyang@visa.com, jinho.choi@emory.edu

Abstract

Recent advances in deep learning have facilitated the demand of neural models for real applications. In practice, these applications often need to be deployed with limited resources while keeping high accuracy. This paper touches the core of neural models in NLP, word embeddings, and presents a new embedding distillation framework that remarkably reduces the dimension of word embeddings without compromising accuracy. A novel distillation ensemble approach is also proposed that trains a high-efficient student model using multiple teacher models. In our approach, the teacher models play roles only during training such that the student model operates on its own without getting supports from the teacher models during decoding, which makes it eighty times faster and lighter than other typical ensemble methods. All models are evaluated on seven document classification datasets and show significant advantage over the teacher models for most cases. Our analysis depicts insightful transformation of word embeddings from distillation and suggests a future direction to ensemble approaches using neural models.

1 Introduction

As deep learning starts dominating the field of machine learning, there have been growing interests in deploying deep neural models for real applications. [Hinton *et al.*, 2014] stated that academic research on model development had mostly focused on accuracy improvement, whereas the deployment of deep neural models would also require the optimization of other practical aspects such as speed, memory, storage, power, etc. To satisfy these requirements, several neural model compression methods have been proposed, which can be categorized into the following four: weight pruning [Denil *et al.*, 2013; Han *et al.*, 2015; Jurgovsky *et al.*, 2016], weight quantization [Han *et al.*, 2016; Jurgovsky *et al.*, 2016; Ling *et al.*, 2016], lossless compression [Van Leeuwen, 1976; Han *et al.*, 2015], and distillation [Mou *et al.*, 2016]. This paper focuses on distillation methods that can remarkably reduce the model size, resulting in much less memory usage and fewer computations.

Distillation aims to extract core elements from a complex network and transfer them to a simpler network so it gives comparable results to the complex network. It has been shown that the core elements can be transferred to various types of networks i.e., deep to shallow networks [Ba and Caruana, 2014], recurrent to dense networks [Chan *et al.*, 2015], and vice versa [Romero *et al.*, 2014; Tang *et al.*, 2016]. Lately, embedding distillation was suggested [Mou *et al.*, 2016], which transferred the output of the projection layer in the source network as input to the target network, although accuracy drop was expected with this approach. Considering the upper bound of a distilled network, that is the accuracy achieved by the original network [Ba and Caruana, 2014], enough room is left for the improvement of embedding distillation. Distilled embeddings can significantly enhance the efficiency of deep neural models in NLP, where the majority of model space is occupied by word embeddings.

In this paper, we first propose a new embedding distillation method based on three teacher-student frameworks, which is a more advanced way of embedding distillation, because the previous one [Mou *et al.*, 2016] is a standalone embedding distillation (Section 2.4) with limited knowledge transfer. Our distilled embeddings not only enable the target network to outperform the previous state of the art [Mou *et al.*, 2016], but also are eight times smaller than the original word embeddings yet allow the target network to achieve compatible (sometimes higher) accuracy to the source network. We then present a novel ensemble approach which extends this distillation framework by allowing multiple teacher models when training a student model. After learning from multiple teachers during training, the student model runs on its own during decoding such that it performs faster and lighter than any of the teacher models yet pushes the accuracy much beyond them with just 1.25% ($5^0/4000$) the size of other typical ensemble models. All models are evaluated on seven document classification datasets; our experiments show the effectiveness of the proposed frameworks, and our analysis illustrates an interesting nature of the distilled word embeddings. To the best of our knowledge, this is the first time that embedding distillation is thoroughly examined for natural language processing and used in ensemble to achieve such promising results.¹

¹https://github.com/bgshin/distill_demo

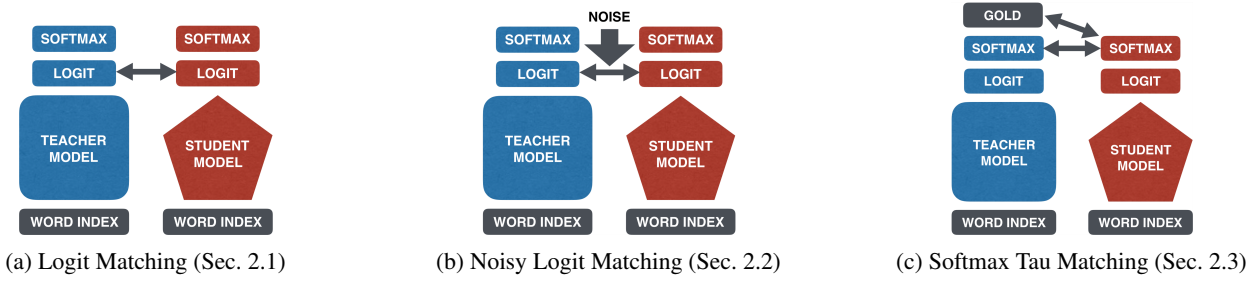


Figure 1: Three teacher-student methods described in Background section, which uses different cost functions to transfer trained knowledge from the teacher model to the student model.

2 Background

Our embedding distillation framework is based on teacher-student models [Ba and Caruana, 2014; Sau and Balasubramanian, 2016; Hinton *et al.*, 2014], where teacher models are trained on deep neural networks and transfer their knowledge to student models on simpler networks. The following subsections describe three popular teacher-student methods applied to our framework. The main difference between these three methods is in their cost functions (Figure 1). The last subsection discusses embedding encoding that is used to extract distilled embeddings from the projection layer.

Throughout this section, a *logit* refers to a vector representing the layer immediately before the softmax layer in a neural network, where z_i and v_i are the teacher’s and student’s logit values for the class i , respectively. Note that the student models are not necessarily optimized for only the gold labels but also optimized for the logit values from the teacher models in these methods.

2.1 Logit Matching (LM)

Proposed by [Ba and Caruana, 2014], the cost function of this teacher-student method is defined by the logit differences between the teacher and the student models (D : the total number of classes):

$$L_{LM} = \frac{1}{2 \cdot D} \sum_{i=1}^D |z_i - v_i|^2$$

2.2 Noisy Logit Matching (NLM)

Proposed by [Sau and Balasubramanian, 2016], this method is similar to Logit Matching except that Gaussian noise is introduced during the distillation, simulating variations in the teacher models, which gives a similar effect for the student to learn from multiple teachers. The cost function takes random noise η drawn from Gaussian distribution such that the logit of each teacher model is $z'_i = (1 + \eta) \cdot z_i$. Thus, the final cost function becomes $L_{NLM} = \frac{1}{2D} \sum_{\forall i} |z'_i - v_i|^2$.

2.3 Softmax Tau Matching (STM)

Proposed by [Hinton *et al.*, 2014], this method is based on softmax matching where softmax values are compared between the teacher and the student models instead of logits. Later, [Hinton *et al.*, 2014] added two hyperparameters to further generalize this method. The first hyperparameter, λ , is for the weighted average of two sub-cost functions, where the first sub-cost function measures a cross-entropy between the

student’s softmax predictions and the truth values, represented as $L_1 = -\sum_{\forall i} y_i \log p_i$ (i indexes classes, y is the gold label, $p_i \in (0, 1)$ is the prediction for a sample). Another cost function involves the second hyperparameter, τ , that is a temperature variable normalizing the output of the teacher’s logit value:

$$s_i(z, \tau) = \frac{e^{z_i/\tau}}{\sum_{j=1}^D e^{z_j/\tau}}$$

Given s_i , the second cost function can be defined as $L_2 = -\sum_{\forall i} s_i(z, \tau) \log p_i$. Therefore, the final cost function becomes $L_{STM} = \lambda L_1 + (1 - \lambda) L_2$. If λ weights more on L_1 , the student model values more on the gold labels than teacher’s predictions. If τ is greater, the teacher’s output becomes more uniformed, implying that the probability values are spread out more throughout all classes.

2.4 Embedding Encoding (ENC)

Embedding distillation was first proposed by [Mou *et al.*, 2016] for NLP tasks. Unlike our framework, their method does not rely on teacher-student models, but rather directly trains a single model with an encoding layer inserted between the embedding layer and its upper layer in the network (Figure 2a). Each word w_i is entered to an embedding layer ϕ that yields a large embedding vector $\phi(w_i)$. This vector is projected into a smaller embedding space by W_{enc} with an activation function f . As a result, a smaller embedding vector $\phi'(w_i)$ is produced for w_i as follows (b_{enc} : a bias for the projection):

$$\phi'(w_i) = f(W_{enc} \cdot \phi(w_i) + b_{enc})$$

The smaller embedding $\phi'(w_i)$ generated by this projection contains distilled knowledge from the larger embedding $\phi(w_i)$. The cost function of this method simply measures the cross-entropy between gold labels and the softmax output values.

3 Embedding Distillation

Our proposed embedding distillation framework begins by training a teacher model using the original embeddings. After training, the teacher model generates the corresponding logit value for each input in the training data. Then, a student model that comprises a projection layer is optimized for the logit (or softmax) values from the teacher model. After training the student model, small embeddings are distilled from the projection layer (Figure 3a).

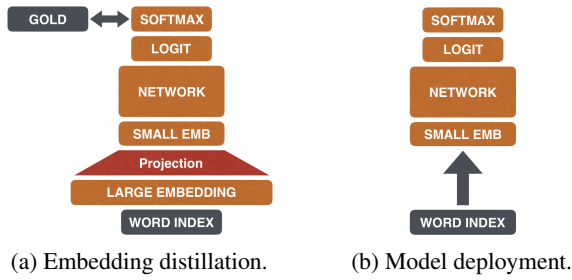


Figure 2: (a) The core elements from large embeddings are distilled during training using gold labels and transferred to smaller embeddings. (b) Only the small embeddings are kept for the deployment, resulting less space and fewer computations.

The original large embeddings as well as weights in the projection layer are discarded for deployment such that the small embeddings can be referenced directly from the word indices in the student model during decoding (Figure 3b). Such distillation significantly reduces the model size and computations in the network, which is welcomed in production.

3.1 Distillation via Teacher-Student Models

Projecting a vector into a lower dimensional space generally entails information loss, although it does not have to be the case under two conditions. First, the source embeddings comprise both relevant and irrelevant information for the target task, therefore, there is a room to discard the irrelevant information. Second, the projection layer in the target network is capable of preserving the relevant information from the source embeddings. The first condition is met for NLP because most vector space models such as Word2Vec [Mikolov *et al.*, 2013], Glove [Pennington *et al.*, 2014], or FastText [Bojanowski *et al.*, 2017] are trained on a vast amount of text, where only a small portion is germane to a specific task.

To meet the second condition, teacher-student models are adapted to our distillation framework, where the projection layer in a student model learns the relevant information for the target task from a teacher model. The output dimension of the projection layer is generally much smaller than the size of the embedding layer in a teacher model. Note that it is possible to integrate multiple projection layers in a student model; Experiment section discusses performance difference by adding different numbers of projection layers in the student model.

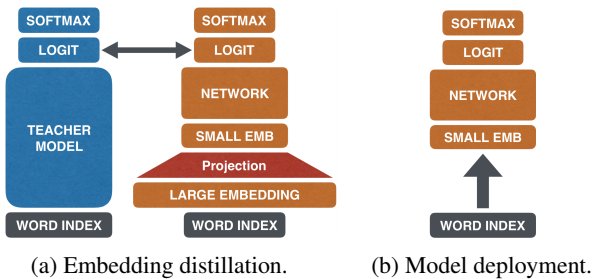


Figure 3: (a) Our proposed embedding distillation framework using teacher-student models. (b) Only the small embeddings are kept for deployment, similarly to Figure 2b.

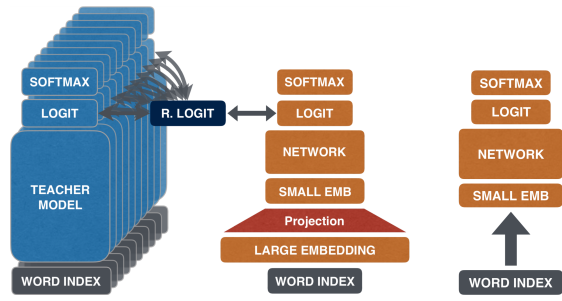


Figure 4: (a) Our proposed embedding distillation ensemble model. One representing logit (R. LOGIT) is calculated from a set of multiple teachers’ logits by the proposed ensemble methods. (b) No need to evaluate teachers at deployment, unlike other ensemble methods.

Figure 4: (a) Our proposed embedding distillation ensemble model. One representing logit (R. LOGIT) is calculated from a set of multiple teachers’ logits by the proposed ensemble methods. (b) No need to evaluate teachers at deployment, unlike other ensemble methods.

3.2 Projection Layer Initialization

Unlike [Mou *et al.*, 2016] who randomly initialized the projection layer, it is initialized with vectors pre-trained by an autoencoder [Hinton and Zemel, 1994] in our framework. This initialization stabilizes and improves optimization of neural networks during training, resulting more robust models.

4 Distillation Ensemble

Ensemble methods generally achieve higher accuracy than a standalone model; however, slow speed is expected due to the runs from multiple models in ensemble. This section presents a novel ensemble approach based on our distillation framework using logit matching that produces a light-weighted student model trained by multiple teachers (Figure 4).

The premise of this approach is that it is possible to have multiple teacher models train a student model by combining their logit values such that the student no longer needs the teachers during decoding because it already learned “enough” from them during training. As a result, our ensemble approach ensures higher efficiency for the student model than for any of the teacher models during decoding. The following sections describe two different ensemble methods applied to our framework.

4.1 Routing by Agreement Ensemble (RAE)

This method gives more weights to the majority, by adopting the dynamic routing algorithm presented by [Sabour *et al.*, 2017]. It first collects the consensus of all teachers, then boosts weights of the teachers who strongly agree with that consensus whereas suppresses the influence of the teachers who do not agree as much. The procedure of calculating the representing logit is described in Algorithm 1.

The squash function in lines 4 and 9 is a non-linear activation function that ensures the norm of the output vector to be in $[0, 1]$ [Sabour *et al.*, 2017]. This vectorized activation is important for a routing algorithm because both magnitude and direction play an important role in the agreement calculation by the dot product, which enforces the consensus into the direction with strong confidence.

Algorithm 1: Get R. LOGIT for RAE and RDE

Input: Teachers’ logits $Z \in \mathbb{R}^{T \times C}$, and an algorithm selector $b \in \{\text{RAE}, \text{RDE}\}$

- 1 . **Output:** The representing logit, $z^{rep} \in \mathbb{R}^C$.
- 2 $k \leftarrow 1$ **if** b is RAE **else** -1
- 3 **for** $t \in \{1, \dots, T\}$ **do**
- 4 $x^t \leftarrow \text{squash}(z^t)$, $w_t \leftarrow 0$
- 5 **while** n iterations **do**
- 6 $c \leftarrow \text{softmax}(w)$
- 7 $z^{rep} = \sum_{t=1}^T c_t \cdot z^t$
- 8 $s \leftarrow \text{squash}(z^{rep})$
- 9 **if** not last iteration **then**
- 10 **for** $t \in \{1, \dots, T\}$ **do**
- 11 $w_t \leftarrow w_t + kx^t \cdot s$
- 12 **return** z^{rep}

Dataset	C	TRN	DEV	TST
MR [Pang and Lee, 2005]	2	7,684	990	1,988
SST-1 [Socher <i>et al.</i> , 2013]	5	8,544	1,101	2,210
SST-2 [Socher <i>et al.</i> , 2013]	2	6,920	872	1,821
Subj [Pang and Lee, 2004]	2	7,199	907	1,894
TREC [Li and Roth, 2002]	6	4,952	500	500
CR [Hu and Liu, 2004]	2	2,718	340	717
MPQA [Wiebe <i>et al.</i> , 2005]	2	7,636	955	2,015

Table 1: Seven datasets used for our experiments. C: number of classes, TRN/DEV/TST: number of instances in training/development/evaluation set.

4.2 Routing by Disagreement Ensemble (RDE)

This method focuses on the minority vote instead, because minority opinions may cast important information for the task. The algorithm is the same as RAE, except for the sign of the weight update (line 2 in Algorithm 1).

5 Experiments

5.1 Datasets

All models are evaluated on seven document classification datasets in Table 1. MR, SST-*, and CR are targeted at the task of sentiment analysis while Subj, TREC, and MPQA are targeted at the classifications of subjectivity, question types, and opinion polarity, respectively. About 10% of the training sets are split into development sets for SST-* and TREC, and about 10/20% of the provided resources are divided into development/evaluation sets for the other datasets, respectively.

5.2 Word Embeddings

For sentiment analysis, raw text from the Amazon Review dataset² is used to train word embeddings, resulting 2.67M word vectors. For the other tasks, combined text from Wikipedia and the New York Times Annotated corpus³ are used, resulting 1.96M word vectors. For each group, two sets

of embeddings are trained with dimensions of 50 and 400 by Word2Vec [Mikolov *et al.*, 2013]. While training, default hyper-parameters are used without an explicit hyper-parameter tuning.

5.3 Network Configuration

Two types of teacher models are developed using Convolutional Neural Networks (CNN) and Long Short-Term Memory Networks (LSTM); comparing different types of teacher models provides more generalized insights for our distillation framework. All teacher models use 400 dimensional word embeddings, and all student models are based on CNN. The CNN-based teacher and student models share the followings: filter sizes = [2, 3, 4, 5], # of filters = 32, dimension of the hidden layer right below the softmax layer = 50. Teacher models add a dropout of 0.8 to the hidden layer, whereas student models add dropouts of 0.1 to both the hidden layer and the projection layer. On the other hand, the LSTM-based teacher models use two bidirectional LSTM layers. Only the last output vector is fed into the 50 dimensional hidden layer, which becomes the input to the softmax layer. A dropout of 0.2 is applied to all hidden layers, both in and out of the LSTM. For both CNN and LSTM ensembles, all 10 teachers share the same model structures with different initializations. Although this limited teacher diversity, our ensemble method produces remarkably good results (Section 5.7).

Each student model may consist of one or two projection layers. The one-layered projection adds one 50 dimensional hidden layer above the embedding layer that transfers core knowledge from the original embeddings to the distilled embeddings. The two-layered projection comprises two layers; the size of the lower layer is the same as the size of teacher’s embedding layer, 400, and the size of the upper layer is 50, which is the dimension of the distilled embeddings. This two-layered projection is empirically found to be more robust than various combinations of network architectures including wider and deeper layers from our experiments.⁴

5.4 Pre-trained Weights

An autoencoder comprising a 50-dimensional encoder and a 400-dimensional decoder is used to pre-train weights for the two-layered projection in student models, where the encoder and the decoder have the same and inversed shapes as the upper and lower layers of the projection. Note that results by using pre-trained weights for the one-layered projection are not reported in Table 2 due to the limited space, but we consistently see robust improvement using pre-trained weights for the projection layers.

5.5 Embedding Distillation

Six models are evaluated on the seven datasets in Table 1 to show the effectiveness of our embedding distillation framework: logit matching (LM), noisy logit matching (NLM), softmax tau matching (STM) models with teacher-student

²<http://snap.stanford.edu/data/web-Amazon.html>

³<https://catalog.ldc.upenn.edu/LDC2008T19>

⁴Among convolutional, relational, and dense-networks with different configurations, the dense-network with the reported configuration produces the best results.

Model	Emb.Size	MR	SST-1	SST-2	Subj	TREC	CR	MPQA
ENC	50	77.11± 0.97	44.94± 1.26	83.71± 1.41	90.64± 0.49	90.60± 1.10	80.88± 1.22	88.65± 0.60
CNN-400	400	79.07	49.86	86.22	92.34	93.60	83.82	88.78
CNN-50	50	78.07	45.07	84.51	90.81	91.00	80.89	86.40
LM+TSED	50	77.63± 0.37	48.71± 0.73	85.04± 0.64	91.91± 0.29	92.48± 0.73	81.84± 0.57	89.14± 0.25
NLM+TSED	50	78.10± 0.40	48.66± 0.83	85.17± 0.16	92.03± 0.36	92.36± 0.91	83.04± 0.74	88.90± 0.34
STM+TSED	50	77.81± 0.33	49.10± 0.34	84.72± 0.70	92.25± 0.38	92.76± 0.65	80.31± 0.42	89.13± 0.28
LM+TSED+PT+2L	50	79.06± 0.59	49.82± 0.54	85.75± 0.42	92.63± 0.23	92.58± 0.86	83.40± 0.76	89.44± 0.20
NLM+TSED+PT+2L	50	78.60± 0.70	49.90± 0.59	85.31± 0.75	92.26± 0.20	92.80± 0.62	83.82± 0.49	89.61± 0.14
STM+TSED+PT+2L	50	78.77± 0.70	49.19± 0.71	85.83± 0.59	92.38± 0.53	93.48± 0.30	83.57± 0.85	89.95± 0.28
LSTM-400	400	79.28	49.23	86.22	92.71	92.00	82.98	89.73
LSTM-50	50	77.16	43.76	83.36	90.02	86.00	80.06	85.66
LM+TSED	50	78.61± 0.80	48.79± 0.27	85.81± 0.77	91.74± 0.36	92.56± 0.89	82.76± 0.36	89.59± 0.28
NLM+TSED	50	78.89± 0.73	48.81± 0.44	85.55± 0.74	91.87± 0.32	91.80± 1.29	82.96± 0.50	89.63± 0.10
STM+TSED	50	78.85± 0.60	48.77± 0.83	86.11± 0.36	91.99± 0.16	92.36± 0.43	82.96± 0.72	89.60± 0.16
LM+TSED+PT+2L	50	80.33± 0.40	49.37± 0.39	86.12± 0.47	92.53± 0.29	91.91± 0.63	83.54± 0.78	90.15± 0.13
NLM+TSED+PT+2L	50	79.33± 0.66	48.87± 0.53	85.89± 0.43	92.35± 0.18	92.08± 0.46	83.32± 0.61	89.92± 0.35
STM+TSED+PT+2L	50	80.09± 0.49	49.14± 0.62	86.95± 0.44	92.34± 0.49	92.96± 0.62	82.73± 0.25	89.83± 0.30

Table 2: Results from our embedding distillation models on the evaluation sets in Table 1 using the CNN-based (the rows 5-10) and the LSTM-based teacher models (rows 13-18) along with the teacher models (*-400), baseline model (*-50) and previous distillation model (ENC). All models are tuned on the development sets and the best performing models are tested on the evaluation sets. Since neural models produce different results at any training due to the random initialization, five models are developed for each approach to avoid (un)lucky peaks, except for *-400 and *-50 where the results are achieved by selecting the best models among ten trials on the development sets. Each score is based on these five trials and represented as a pair of [Average ± Standard Deviation].

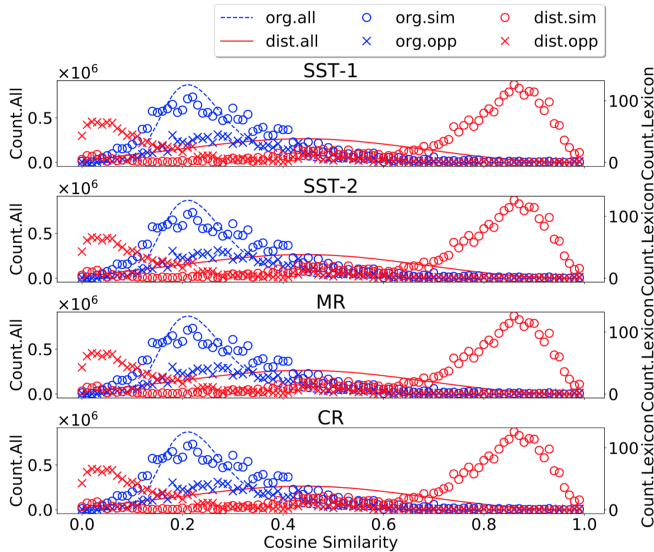


Figure 5: Similarity distributions among sentiment word pairs. Blue and Red colors distinguish histograms from the original and distilled embeddings, respectively. The solid and dashed lines show the distributions from all word pairs regardless of their sentiments. Circles are for sentimentally similar word pairs, that are similarities between positive word pairs and negative word pairs such that $(w_i, w_j) \in (P_\ell \times P_\ell) \cup (N_\ell \times N_\ell)$. Crosses are for sentimentally opposite word pairs, that are similarities across positive and negative word pairs such that $(w_i, w_j) \in P_\ell \times N_\ell$.

based embedding distillation (TSED), and another three models with the autoencoder pre-trained weights (*+PT) and the two layered projection network (*+2L). Teacher models using 400-dim embeddings (*-400) are also presented along with the baseline model using 50-dim word embeddings (*-50) and the previous distillation model (ENC). The comparison to these

two models highlights the strength of our distilled models, significantly outperforming them with the same dimensional word embeddings (50-dim).

Table 2 shows the results achieved by all models. While the two existing models, *-50 and ENC, show marginal differences, our proposed models, *+TSED, outperform the previous embedding distillation SOTA (ENC). Our final models, *+PT+2L, outperform all the other models, reaching similar (or even higher) accuracy to the teacher models (*-400), which confirms that the proposed embedding distillation framework can successfully transfer the core knowledge from the original embeddings with respect to the target tasks, independent from the network structures of the teacher models.

The fact that the best model for each task comes from a different teacher-student strategy appears to be random. However, considering that it is the nature of neural network models whose accuracy deviates at every training, this is not surprising although it signifies the need of ensemble approaches to take advantage of multiple teachers and produce a more robust student model.

5.6 Lexical Analysis

Embedding distillation for a task such as sentiment analysis can be viewed as a vector transformation that adjusts similarities between word embeddings with respect to their sentiments. Thus, we hypothesize that distilled word embeddings from sentiment analysis should bring similar sentiment words together while disperse opposite ones in vector space.

To verify this, lexical analysis on both the original and the distilled word embeddings is conducted using the four sentiment datasets: SST-1, SST-2, MR, and CR. First, positive and negative words are collected from two publicly available lexicons, the MaxDiff Twitter Sentiment Lexicon [Kiritchenko *et al.*, 2014] and the Bing Liu Opinion Lexicon [Hu and Liu, 2004]. Then, two groups of sentiment word sets are con-

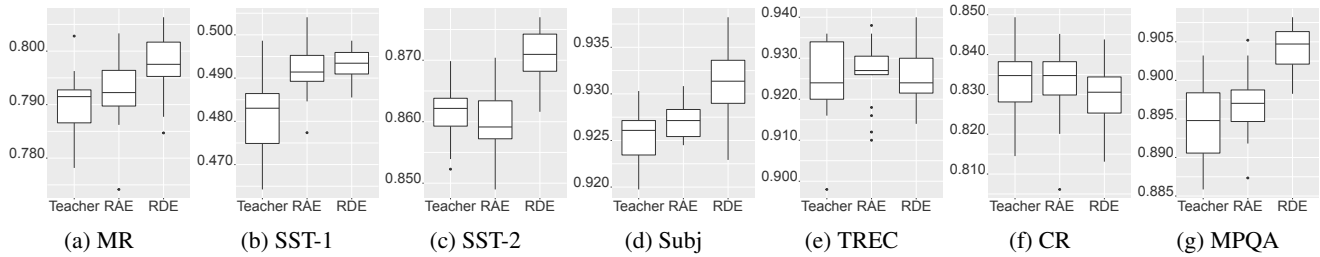


Figure 6: Accuracy comparisons between the ensemble and the teacher models. To avoid (un)lucky peaks, each method is evaluated 20 times where each trial produces a different result. These evaluation results are shown as boxplots in this figure.

structured, (P_t, N_t) and (P_o, N_o) , where P_* and N_* compose positive and negative words, and $*_t$ and $*_o$ are collected from the Twitter and the Opinion lexicons, respectively. Next, the intersection between each type of sentiment word sets is found, that are $P_{to} = P_t \cap P_o$ and $N_{to} = N_t \cap N_o$, where $|P_{to}| = 72$ and $|N_{to}| = 89$. Finally, the intersections between $*_{to}$ and the vocabulary set from each of the four sentiment datasets are found (e.g., $P_\ell = A_\ell \cap P_{to}$, where ℓ is one of the four datasets and A_ℓ is the set of all words in ℓ):

- $\ell = \text{SST}_1 \Rightarrow |P_\ell| = 66$ and $|N_\ell| = 83$.
- $\ell = \text{SST}_2 \Rightarrow |P_\ell| = 67$ and $|N_\ell| = 83$.
- $\ell = \text{MR} \Rightarrow |P_\ell| = 67$ and $|N_\ell| = 82$.
- $\ell = \text{CR} \Rightarrow |P_\ell| = 19$ and $|N_\ell| = 57$.

For each ℓ , cosine similarities are measured for all possible word pairs $(w_i, w_j) \in PN_\ell \times PN_\ell$ where $PN_\ell = P_\ell \cup N_\ell$, using the original and distilled embeddings. Figure 5 illustrates the similarity distributions. It is clear that similar sentiment words generally give high similarity scores with the distilled embeddings (the plots drawn by red circles), whereas opposite sentiment words give low similarity scores (the plots drawn by red crosses). On the other hand, low similarity scores are shown for any case with the original embeddings (the plots drawn by blue circles and crosses). The normal distributions derived from the distilled embeddings (the red lines) are more symmetric and spread than the ones from the original embeddings (the blue lines), implying that distilled embeddings are more stable for the target task.

5.7 Distillation Ensemble

All ensemble models are based on our distillation framework using logit matching (LM) where the teacher models compose of 10 CNN-based or 10 LSTM-based models. Two ensemble methods are evaluated: Routing by Agreement (RAE) and Routing by Disagreement (RDE), and Figure 6 shows comparisons between these ensemble models against the teacher models.

The most notable finding is that RDE significantly outperforms the teacher, if the dataset is big. For example, RDE outperforms the teacher models on average, except for CR and MPQA, whose training set is relatively small (Table 1). The insight behind this trend might be that if there are many data samples, then the probability of exploring different knowledge from minor opinions could be increased, which would positively affect to the ensemble.

	Previous	Proposed	Reduction
Train	$O(400M * 10)$	$O(400M * 11)$	$\times 0.91$
Deploy	$O(400M * 10)$	$O(50M)$	$\times 80$

Table 3: The number of neurons in previous ensemble methods and the proposed distillation ensemble method for training and deploying. M represents the basic unit of model size for the embedding dimension 1. This table assumes ensemble with 10 teachers.

5.8 Model Reduction

The deployment models from either distillation or ensemble are notably smaller than the teacher models. Since word embeddings occupy a large portion of neural models in NLP, reducing the size of word embeddings through distillation decreases the size of the entire model roughly by the same ratio as its reduction; in our case, eight times ($400/50$). Furthermore, if the proposed distillation ensemble method is compared to other typical ensemble methods, this reduction ratio becomes even larger. Table 3 shows the number of neurons required for previous ensemble methods and the proposed one. When training, the proposed one comprises 10% more parameters due to the distillation process. However, when deploying, the reduction of neuron is about eighty times ($4000/50$). This is because the proposed framework doesn't require repetitive evaluation of teacher models when testing, while other ensemble methods require evaluation of all sub models (teachers).

It is worth mentioning that the deployment models produced by distillation ensemble not only outperform the teacher models in accuracy but also are significantly smaller such that they operate much faster and lighter than the teacher models upon deployment. This is very welcoming for those who want to embed these models into low-resource platforms such as mobile environments.

6 Conclusion

This paper proposes a new embedding distillation framework based on several teacher-student methods. Our experiments show that the proposed distillation models outperform the previous distillation model and give compatible accuracy to the teacher models, yet they are significantly smaller. Lexical analysis on sentiments reveals the comprehensiveness of the distilled embeddings. Moreover, a novel distillation ensemble approach is proposed, which shows huge advantage in both speed and accuracy over any teacher model. Our distillation ensemble approach consistently shows more robust results when the size of training data is sufficiently large.

References

- [Ba and Caruana, 2014] Jimmy Ba and Rich Caruana. Do deep nets really need to be deep? In *Advances in neural information processing systems*, pages 2654–2662, 2014.
- [Bojanowski *et al.*, 2017] Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. Enriching word vectors with subword information. *Transactions of the Association of Computational Linguistics*, 5:135–146, 2017.
- [Chan *et al.*, 2015] William Chan, Nan Rosemary Ke, and Ian Lane. Transferring knowledge from a rnn to a dnn. *arXiv preprint arXiv:1504.01483*, 2015.
- [Denil *et al.*, 2013] Misha Denil, Babak Shakibi, Laurent Dinh, Nando de Freitas, et al. Predicting parameters in deep learning. In *Advances in Neural Information Processing Systems*, pages 2148–2156, 2013.
- [Han *et al.*, 2015] Song Han, Jeff Pool, John Tran, and William Dally. Learning both weights and connections for efficient neural network. In *Advances in Neural Information Processing Systems*, pages 1135–1143, 2015.
- [Han *et al.*, 2016] Song Han, Huizi Mao, and William J Dally. Deep compression: Compressing deep neural networks with pruning, trained quantization and Huffman coding. *ICLR*, 2016.
- [Hinton and Zemel, 1994] Geoffrey E Hinton and Richard S Zemel. Autoencoders, minimum description length and helmholtz free energy. In *Advances in neural information processing systems*, pages 3–10, 1994.
- [Hinton *et al.*, 2014] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *NIPS Deep Learning Workshop*, 2014.
- [Hu and Liu, 2004] Minqing Hu and Bing Liu. Mining and summarizing customer reviews. In *Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 168–177. ACM, 2004.
- [Jurgovsky *et al.*, 2016] Johannes Jurgovsky, Michael Granitzer, and Christin Seifert. Evaluating memory efficiency and robustness of word embeddings. In *European conference on information retrieval*, pages 200–211. Springer, 2016.
- [Kiritchenko *et al.*, 2014] Svetlana Kiritchenko, Xiaodan Zhu, and Saif M Mohammad. Sentiment analysis of short informal texts. *Journal of Artificial Intelligence Research*, 50:723–762, 2014.
- [Li and Roth, 2002] Xin Li and Dan Roth. Learning question classifiers. In *Proceedings of the 19th international conference on Computational linguistics-Volume 1*, pages 1–7. Association for Computational Linguistics, 2002.
- [Ling *et al.*, 2016] Shaoshi Ling, Yangqiu Song, and Dan Roth. Word embeddings with limited memory. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, volume 2, pages 387–392, 2016.
- [Mikolov *et al.*, 2013] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pages 3111–3119, 2013.
- [Mou *et al.*, 2016] Lili Mou, Ran Jia, Yan Xu, Ge Li, Lu Zhang, and Zhi Jin. Distilling word embeddings: An encoding approach. In *Proceedings of the 25th ACM International on Conference on Information and Knowledge Management*, pages 1977–1980. ACM, 2016.
- [Pang and Lee, 2004] Bo Pang and Lillian Lee. A sentimental education: Sentiment analysis using subjectivity summarization based on minimum cuts. In *Proceedings of the 42nd annual meeting on Association for Computational Linguistics*, page 271, 2004.
- [Pang and Lee, 2005] Bo Pang and Lillian Lee. Seeing stars: Exploiting class relationships for sentiment categorization with respect to rating scales. In *Proceedings of the 43rd annual meeting on association for computational linguistics*, pages 115–124, 2005.
- [Pennington *et al.*, 2014] Jeffrey Pennington, Richard Socher, and Christopher Manning. Glove: Global vectors for word representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, Doha, Qatar, October 2014. Association for Computational Linguistics.
- [Romero *et al.*, 2014] Adriana Romero, Nicolas Ballas, Samira Ebrahimi Kahou, Antoine Chassang, Carlo Gatta, and Yoshua Bengio. Fitnets: Hints for thin deep nets. *arXiv preprint arXiv:1412.6550*, 2014.
- [Sabour *et al.*, 2017] Sara Sabour, Nicholas Frosst, and Geoffrey E Hinton. Dynamic routing between capsules. In *Advances in Neural Information Processing Systems*, pages 3859–3869, 2017.
- [Sau and Balasubramanian, 2016] Bharat Bhusan Sau and Vineeth N Balasubramanian. Deep model compression: Distilling knowledge from noisy teachers. *arXiv preprint arXiv:1610.09650*, 2016.
- [Socher *et al.*, 2013] Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D Manning, Andrew Ng, and Christopher Potts. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the 2013 conference on empirical methods in natural language processing*, pages 1631–1642, 2013.
- [Tang *et al.*, 2016] Zhiyuan Tang, Dong Wang, and Zhiyong Zhang. Recurrent neural network training with dark knowledge transfer. In *Acoustics, Speech and Signal Processing (ICASSP), 2016 IEEE International Conference on*, pages 5900–5904. IEEE, 2016.
- [Van Leeuwen, 1976] Jan Van Leeuwen. On the construction of Huffman trees. In *ICALP*, pages 382–410, 1976.
- [Wiebe *et al.*, 2005] Janyce Wiebe, Theresa Wilson, and Claire Cardie. Annotating expressions of opinions and emotions in language. *Language resources and evaluation*, 39(2-3):165–210, 2005.