

# Single Training Dimension Selection for Word Embedding with PCA

Yu Wang

Apple

w.y@apple.com

## Abstract

In this paper, we present a fast and reliable method based on PCA to select the number of dimensions for word embeddings. First, we train one embedding with a generous upper bound (e.g. 1,000) of dimensions. Then we transform the embeddings using PCA and incrementally remove the lesser dimensions one at a time while recording the embeddings' performance on language tasks. Lastly, we select the number of dimensions while balancing model size and accuracy. Experiments using various datasets and language tasks demonstrate that we are able to train 10 times fewer sets of embeddings while retaining optimal performance. Researchers interested in training the best-performing embeddings for downstream tasks, such as sentiment analysis, question answering and hypernym extraction, as well as those interested in embedding compression should find the method helpful.

## 1 Introduction

Word embeddings constitute an integral part of the implementation of numerous NLP tasks ranging from sentiment classification (Lin et al., 2015), nationality classification (Ye et al., 2017), classification of behavior on Twitter (Wang et al., 2017), to measuring document similarity (Kusner et al., 2015). The strength of word embeddings stems from their embedding words into low dimensional continuous vector space (Mnih and Kavukcuoglu, 2013; Pennington et al., 2014; Tang et al., 2015; Raunak, 2017).

Various algorithms have been proposed for learning words' vector space representations, including most notably Mikolov et al. (2013b), Pennington et al. (2014), and Nickel and Kiela (2017). However, the exact definition of 'low' dimensionality is rarely explored. As pointed out by Yin and Shen (2018), the most frequently used dimensionality is 300, largely due to the fact that the early

influential papers (Mikolov et al., 2013b; Pennington et al., 2014) used 300. Other often used dimensionalities include 200 (Tang et al., 2015; Ling et al., 2016; Nickel and Kiela, 2017) and 500 (Tang and Liu, 2009; Perozzi et al., 2014).

The impact of vector dimension on embeddings' performance is well known (Lai et al., 2016). With too few dimensions, the model will underfit; with too many dimensions the model will overfit. Both undercut the embeddings' performance (Yin and Shen, 2018). What is also known is that the size of the embeddings will grow linearly with the vector dimension (Ling et al., 2016; Raunak, 2017). What is less known is how to identify the optimal vector dimension given any dataset. The method we propose here helps fill this gap.

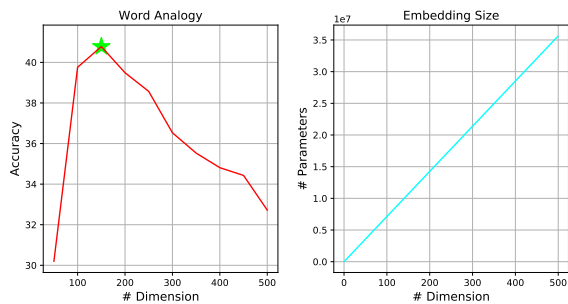


Figure 1: Left: Accuracy for the word analogy task as a function of vector dimension (dimension incremental is 50). Right: The number of embedding parameters as a function of vector dimension. Results are generated using the Text8 dataset with a vocabulary size of 71,291.

We offer a fast and reliable PCA-based method that (1) only needs to train the embeddings once and (2) is able to select vector dimension with competitive performance.<sup>1</sup> First, we train one embedding with a generous upper bound (e.g. 1,000)

<sup>1</sup>As such, the selected dimension is a Pareto-optimal dimension that balances embedding size and performance.

of dimensions. Then we transform the embeddings using PCA and incrementally remove the lesser dimensions one at a time while recording the embeddings’ performance on language tasks. Lastly we calculate the best dimensionality and return the corresponding embeddings.

Experiments using various datasets and language tasks reveal three key observations:

- The optimal dimensionality calculated on the basis of PCA agrees with that by grid search.
- The resulting embedding is competitive against the one selected by grid search.
- Different upper-bound dimensionalities (e.g. 500, 1000) point to the same optimal dimensionality.

Researchers interested in downstream tasks, such as sentiment analysis (Lin et al., 2015), question answering (Devlin et al., 2018) and hypernym extraction (Chen et al., 2018), as well as those interested in embedding compression should find the method helpful.

## 2 Related Work

Our work draws inspiration from Yin and Shen (2018). The authors build on the Latent Semantics Analysis (LSA) approach and slide  $k$  from a lower bound (e.g. 10) to a generous upper bound (e.g. 1,000) in  $E = U_{1:k} D_{1:k,1:k}^\alpha$ , where  $U$  and  $D$  come from the singular-value decomposition of the signal matrix and  $\alpha$  is a hyperparameter to be tuned. For each  $k$ , the authors generate one corresponding embedding and compare it with the simulated oracle embedding. The  $k$  that yields the smallest loss is selected. In a similar vein, our work bypasses the problem of training multiple embeddings, often necessitated by grid search, by sliding over all the  $k$  values of PCA. Compared with Yin and Shen (2018), our method is easier to implement, as we do not rely on, e.g. Monte Carlo simulations of the oracle embeddings.

At a deeper level, our work is also connected to Yin and Shen (2018) in terms of the trade-off between bias and variance. Yin and Shen (2018) propose pairwise inner product (PIP) loss to measure the quality of an embedding. They decompose the PIP loss into a bias term and a variance term, where reducing the dimension increases the bias term but reduces the variance. They show

that the bias-variance trade-off reflects the signal-to-noise ratio in dimension selection. While there is no exact 1-1 mapping from their theorem to our work, we do have analogous discussion in Section 3. The PCA step in our algorithm enables us to identify and drop dimensions that (1) contribute less to the explained variance in the embedding and yet (2) contribute equally to cosine similarity. In essence, our PCA step is removing dimensions with low signal-to-noise ratios.

Our work also draws strength from the literature on post-processing embeddings. Mu and Viswanath (2018) demonstrate that removing the top dominating directions in the PCA-transformed embeddings helps improve the embeddings’ performance in word analogy and similarity tasks. Building on that, Raunak (2017) shows that by performing another iteration of PCA and dropping the bottom directions, one can further improve a model’s performance as well as reduce its size. Both works focus on improving pre-trained embeddings’ performance in terms of accuracy and size. By contrast, our algorithm selects the optimal dimensionality before the actual training.

In addition, our work is related to a few recent studies on model compression (Faruqui et al., 2015; Ling et al., 2016; Shu and Nakayama, 2018). In particular, Ling et al. (2016) seek to drop the least significant digits to reduce the embeddings’ size. By comparison, our method removes the least significant dimensions (in terms of explained variance (Bishop, 2006)). It should be noted that the two methods complement each other, as one focuses on dimension selection whereas the other on limited precision representation.

## 3 Algorithm

In this section, we formally describe how to select a competitive dimensionality by training one embedding. We state the proposed algorithm in Algorithm 1.

First, we note that the PCA transformation (when retaining all the  $N$  dimensions) does not affect embeddings’ performance on word similarity tasks. Any potential performance gain should come from dropping the lesser dimensions. By “lesser dimensions,” we mean the dimensions that contribute little to the explanation of variance (Bishop, 2006).

In Figure 2, we first transform an embedding

---

**Algorithm 1** Select the Optimal Dimensionality using PCA

---

Set dimension upper bound  $N$ ; select language task from {word analogy, similarity}  
Train embedding  $E$  with  $N$  dimensions  
Transform  $E$  using PCA:  $(u_1, u_2, \dots, u_N; \tilde{E}) \leftarrow \text{PCA}(E)$ , where  $u_1, u_2, \dots, u_N$  are the new basis vectors,  $\tilde{E}$  represents the transformed coefficients  
**for**  $i = N$  to 2 **do**  
     $E = E - \tilde{E}_{:,i} \cdot u_i$ , where  $\tilde{E}_{:,i}$  represents the  $i$ th column of  $\tilde{E}$  and each scalar in  $\tilde{E}_{:,i}$  scales vector  $u_i$   
    Evaluate  $E$  on the selected language task, record ( $i$ , metric)  
**end for**  
Return the selected dimension:  $[\underset{i}{\operatorname{argmax}} f(i, \text{metric})] - 1$ , where  $f$  is a score function that balances performance and model size and  $i$  is between 2 and  $N$

---

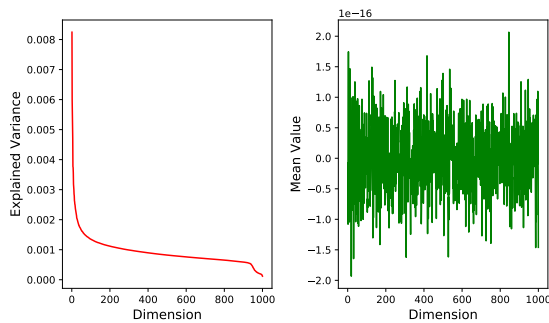


Figure 2: Left: explained variance drops sharply after the top 100 dimensions. Right: all the 1,000 dimensions have roughly the same mean value.

using PCA, so that each new dimension represents a principal component. We show that the explained variance goes drastically down for dimensions beyond the 100th and stays relatively stable for 100th-1,000th dimensions. In terms of magnitude, the first dimension explains 69.2 times more variance than the last dimension.

While different dimensions contribute differently to the explained variance, they nonetheless contribute equally to the calculation of inner product (Figure 2, right). Therefore, the lesser dimensions, with less variance but equal weighting, effectively decreases the discriminative power of the model. Removing these lesser dimensions enables us to focus on the more discriminative dimensions. To identify optimal dimensionality, beyond which all dimensions are considered lesser, we turn to experiments in Section 4.

## 4 Experiments

Given the popularity of the word2vec model (Mikolov et al., 2013b), we use Skip-gram as the embedding algorithm. Following Yin and Shen (2018) and Grave et al. (2017), we use the widely

used benchmark datasets, Text8 (Mahoney, 2011) and WikiText-103 (Merity et al., 2017), as the training datasets. For ground truth, we train 20 embeddings, with dimensions ranging from 50 to 1,000 at an interval of 50 as well as two embeddings with dimensions of 5 and 25.<sup>2</sup> Here we have made the implicit assumption that 1,000 is an upper bound for the embedding space’s dimensionality. For each embedding, we train 200 epochs (Pennington et al., 2014; Shazeer et al., 2016) and keep only the checkpoint that performs best on the word analogy task (Mikolov et al., 2013a). Our experiments focus on (1) comparing our method with grid-search based ground truth and (2) examining consistency between different upper bounds.

### 4.1 Performance Compared with Grid Search

In this subsection, we compare the optimal dimensionality that our method calculates with the ground truth (Figure 3). We perform the comparison across three testing datasets: Wordsim 353 (Finkelstein et al., 2002), RW Stanford (Luo et al., 2013) and MTurk 771 (Halawi et al., 2012). Figure 3 demonstrates that our PCA based method (with one training only) is able to uncover the optimal dimensionality. Table 1 reports the distance in selected dimensionalities.

We also observe that the optimal embedding that results from our method (with retraining) is competitive against the optimal embedding found using grid search. In Figure 3, we mark out the respective optimal performances of the two approaches in similarity tasks. In Table 2, we further

---

<sup>2</sup>A larger interval will save researchers more time, but may result in a sub-optimal dimensionality. While a smaller interval may give more accurate results, it requires training more sets of embeddings proportionally. Our proposed method, by contrast, can scan through all the dimensions at the finest granularity, 1, at virtually no extra cost.

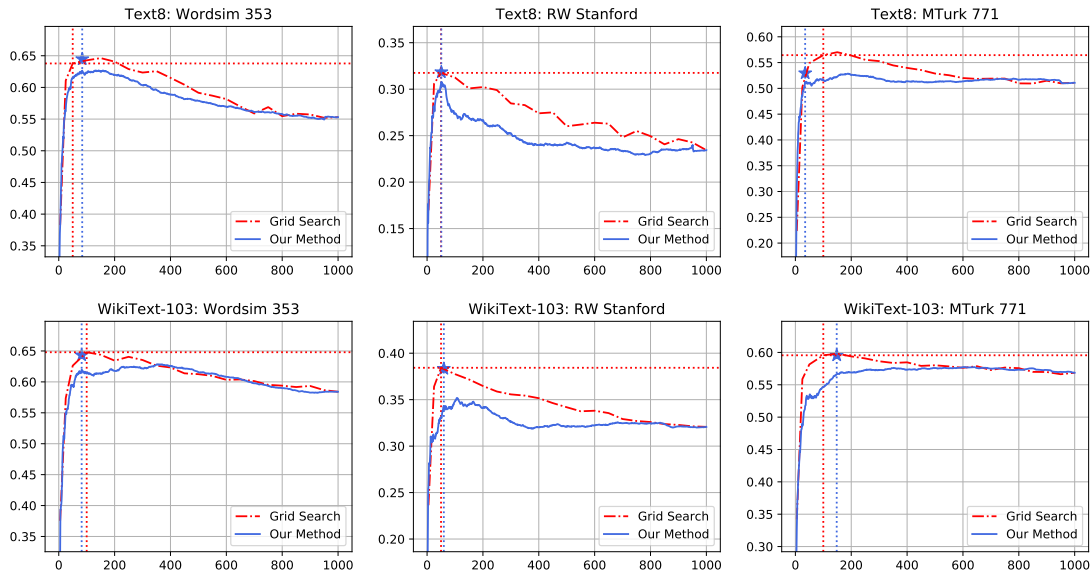


Figure 3: Across different benchmarks and different training datasets, the optimal dimensionality that our method (blue curve) identifies closely matches grid search (red curve). The top row is based on Text8. The bottom row is based on WikiText-103. The vertical dotted lines represent the optimal dimensionality for the respective curves. The horizontal dotted line represents the performance of grid search and the blue star marks the performance of our method. The score function is  $f(i, \text{metric}) = \text{metric} - 50 \times i$ . All the curves are averaged over 5 random runs.

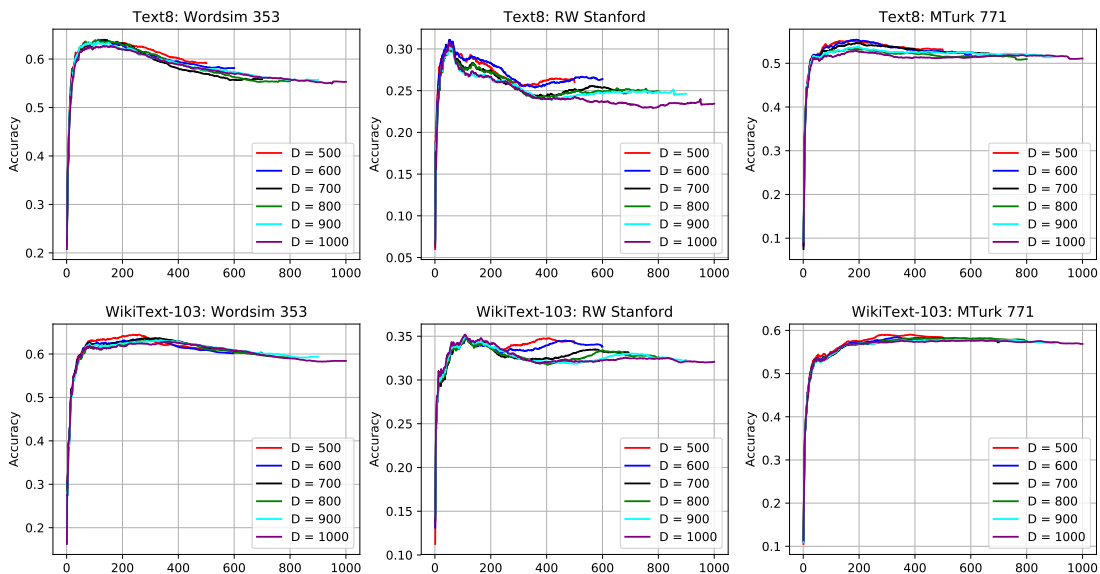


Figure 4: Across different benchmarks and different training datasets, consistency is observed when the upper bound is set to 500, 600, 700, 800, 900 and 1,000. The curves trace closely each other, in particular for WordSim 353 and MTurk 771. The top row is based on the Text8 dataset. The bottom row is based on the WikiText-103 dataset. All the curves are averaged over 5 random runs. Best viewed in color.

report the optimal performance achieved by grid search and our method as well as their relative performance. Even though we have only trained one embedding (and one retraining), our method, on average, is able to achieve 100.2% (WordSim 353) to 96.9% (MTurk 771) of the optimal performance

by grid-searching through 22 sets of embeddings.

#### 4.2 Consistency across Upper Bounds

One hyperparameter involved in our method is the upper bound. Intuitively, we expect the upper bound should be higher for larger datasets. In this

Table 1: Distance in Dimensionalities between Grid Search and Our Proposed Method (Numbers within the bucket size 50 are in bold)

Datasets	WordSim	RW Stanford	MTurk
Text8	<b>34</b>	<b>1</b>	65
WikiText	<b>18</b>	<b>10</b>	<b>48</b>
<i>Average</i>	<b>26</b>	<b>5.5</b>	66.5

Table 2: Performance (*Correlation*) Comparison between Grid Search and Our Proposed Method

Datasets	Method	WordSim	RW Stanford	MTurk
Text8	G.S.	63.8	31.8	56.4
	Ours	64.5 (101.1%)	31.8(100.3%)	52.9 (93.8%)
WikiText	G.S.	64.8	38.4	59.6
	Ours	64.3 (99.3%)	38.3 (99.7%)	59.6 (100%)
<i>Average</i>	-	100.2%	100%	96.9%

subsection, we demonstrate our method is robust against different upper bounds.

In Figure 4, we vary the dimension from 500 to 1,000 at an increment of 100. We observe that the dimensionality our method selects is consistent across different upper bounds. Based on the demonstrated consistency, different upper bounds can be selected and still the optimal dimensionality can be uncovered as long as the chosen upper bound is larger than the optimal dimensionality.

### 4.3 Efficiency Compared with Grid Search

In this subsection, we report the running time of our algorithm and compare it with that of grid search. We have recorded the running time of our experiments and that of the PCA transformations. We average them over 5 random runs (Figure 5).

For Text8, grid search takes 22,801 minutes cumulatively. Training a 1,000-dimension embedding takes 1,724 minutes. The PCA step takes 22 minutes (note that for each embedding we only need one PCA operation). This represents a 13.1x speedup for our method. For WikiText-103, grid search takes 132,652 minutes. Training a 1,000-dimension embedding takes 10,448 minutes. The PCA step takes 34 minutes. This represents a 12.7x speedup.

We note that the comparison results are dependent on grid granularity. A coarser grid search could save researchers more time, at the cost of performance loss.




Figure 5: Left: training time for Text8 for embeddings of different dimensions for 200 epochs using 6 CPUs. Right: training time for WikiText-103 for embeddings of different dimensions for 200 epochs using 10 CPUs. All results are averaged over 5 random runs.

## 5 Conclusion

In this paper, we provided a fast and reliable method based on PCA to select the number of dimensions for training word embeddings. First, we train one embedding with a generous upper bound (e.g. 1,000) of dimensions. Then we transform the embedding using PCA and incrementally remove the lesser dimensions while recording the embeddings’ performance on language tasks. Experiments demonstrate that (1) our method is able to identify the optimal dimensionality, (2) the resulting embedding has competitive performance against grid search, and (3) our method is robust against the selection of the upper bound.

## Acknowledgement

The authors would like to thank Hang Zhao and Srinivasan Venkatachary for supporting this project, would like to thank Russ Webb, Arnab Ghoshal, and Jaewook Chung for their insightful comments on early versions of this paper, and would like to thank the anonymous EMNLP-IJCNLP reviewers for their reviews and suggestions.

## References

- Christopher M. Bishop. 2006. *Pattern Recognition and Machine Learning*. Springer.
- Hong-You Chen, Cheng-Syuan Lee, Keng-Te Liao, and Shou-De Lin. 2018. Word relation autoencoder for unseen hypernym extraction using word embeddings. *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4834–4839.

- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint*.
- Manaal Faruqui, Yulia Tsvetkov, Dani Yogatama, Chris Dyer, and Noah A. Smith. 2015. Sparse over-complete word vector representations. *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics*, pages 1491–1500.
- Lev Finkelstein, Evgeniy Gabrilovich, Yossi Matias, Ehud Rivlin, Gadi Wolfman Zach Solan, and Eytan Ruppin. 2002. Placing search in context: the concept revisited. *ACM Transactions on Information Systems*, 20(1):116–131.
- Edouard Grave, Armand Joulin, and Nicolas Usunier. 2017. Improving neural language models with a continuous cache. *ICLR 2017*.
- Guy Halawi, Gideon Dror, Evgeniy Gabrilovich, and Yehuda Koren. 2012. Large-scale learning of word relatedness with constraints. *KDD*.
- Matt J. Kusner, Yu Sun, Nicholas I. Kolkin, and Kilian Q. Weinberger. 2015. **From Word Embeddings To Document Distances**. *Proceedings of the 32nd International Conference on Machine Learning*.
- Siwei Lai, Kang Liu, Shizhu He, and Jun Zhao. 2016. How to generate a good word embedding. *IEEE Intelligent Systems*, 31(6).
- Yiou Lin, Hang Lei, Jia Wu, and Xiaoyu Li. 2015. **An Empirical Study on Sentiment Classification of Chinese Review using Word Embedding**. *29th Pacific Asia Conference on Language, Information and Computation*, pages 258 – 266.
- Shaoshi Ling, Yangqiu Song, and Dan Roth. 2016. Word embeddings with limited memory. *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*, pages 387–392.
- Thang Luong, Richard Socher, and Christopher Manning. 2013. Better word representations with recursive neural networks for morphology. *Proceedings of the Seventeenth Conference on Computational Natural Language Learning*.
- Matt Mahoney. 2011. Large text compression benchmark.
- Stephen Merity, Caiming Xiong, James Bradbury, and Richard Socher. 2017. Pointer sentinel mixture models. *ICLR 2017*.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013a. Efficient Estimation of Word Representations in Vector Space. *ICLR Workshop*.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013b. **Distributed Representations of Words and Phrases and their Compositionality**. *NIPS’13 Proceedings of the 26th International Conference on Neural Information Processing Systems*, 2:3111–3119.
- Andriy Mnih and Koray Kavukcuoglu. 2013. Learning word embeddings efficiently with noise-contrastive estimation. *NIPS’13 Proceedings of the 26th International Conference on Neural Information Processing Systems*, 2.
- Jiaqi Mu and Pramod Viswanath. 2018. All-but-the-top: Simple and effective postprocessing for word representations. *Sixth International Conference on Learning Representations (ICLR 2018)*.
- Maximillian Nickel and Douwe Kiela. 2017. Poincaré embeddings for learning hierarchical representations. *Advances in Neural Information Processing Systems 30 (NIPS 2017)*.
- Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. GloVe: Global Vectors for Word Representation. In *Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543.
- Bryan Perozzi, Rami Al-Rfou, and Steven Skiena. 2014. **DeepWalk: Online Learning of Social Representations**. *KDD ’14 Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 701–210.
- Vikas Raunak. 2017. Simple and effective dimensionality reduction for word embeddings. *Advances in Neural Information Processing Systems 30 (NIPS 2017) LLD Workshop*.
- Noam Shazeer, Ryan Doherty, Colin Evans, and Chris Waterson. 2016. Swivel: Improving Embeddings by Noticing What’s Missing. <https://arxiv.org/abs/1602.02215>.
- Raphael Shu and Hideki Nakayama. 2018. Compressing word embeddings via deep compositional code learning. *ICLR 2018 Conference*.
- Jian Tang, Meng Qu, Mingzhe Wang, Ming Zhang, Jun Yan, and Qiaozhu Mei. 2015. **LINE: Large-scale Information Network Embedding**. *International World Wide Web Conference*.
- Lei Tang and Huan Liu. 2009. Relational learning via latent social dimensions. *KDD ’09: Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*.
- Yu Wang, Yang Feng, Zhe Hong Ryan Berger, and Jiebo Luo. 2017. How Polarized Have We Become? A Multimodal Classification of Trump Followers and Clinton Followers. *SocInfo 2017: Social Informatics*, pages 440–456.
- Junting Ye, Shuchu Han, Yifan Hu, Baris Coskun, Meizhu Liu, Hong Qin, and Steven Skiena. 2017. **Nationality Classification Using Name Embeddings**. *CIKM’17*, pages 1897–1906.
- Zi Yin and Yuanyuan Shen. 2018. On the dimensionality of word embedding. *Advances in Neural Information Processing Systems 31 (NIPS 2018)*.