

1

# **Spherical Text Embedding**

Yu Meng<sup>1</sup>, Jiaxin Huang<sup>1</sup>, Guangyuan Wang<sup>1</sup>, Chao Zhang<sup>2</sup>, Honglei Zhuang<sup>3</sup>, Lance Kaplan<sup>4</sup> and Jiawei Han<sup>1</sup>

University of Illinois at Urbana-Champaign<sup>1</sup>, Georgia Institute of Technology<sup>2</sup>, Google Research<sup>3</sup>, U.S. Army Research Laboratory<sup>4</sup>

Code: <u>https://github.com/yumeng5/Spherical-Text-Embedding</u>

### Outline



- Motivations & Introduction
- □ Model: Spherical Text Embedding
- Method: Optimization on the Sphere
- **Experiments**
- Conclusions

## **Preliminaries: Text Embedding**

#### □ A milestone in NLP and ML:

- Unsupervised learning of text representations—No supervision needed
- Embed one-hot vectors into lower-dimensional space—Address "curse of dimensionality"
- Word embedding captures useful properties of word semantics
  - □ Word similarity: Words with similar meanings are embedded closer
  - □ Word analogy: Linear relationships between words (e.g. king queen = man woman)



## **Preliminaries: Text Embedding**

□ How to use text embeddings? Mostly directional similarity (i.e., cosine similarity)

□ Word similarity is derived using cosine similarity





- □ Word clustering (e.g. TaxoGen) is performed on a sphere
- Better document clustering performances when embeddings are normalized and spherical clustering algorithms are used



### Outline

Preliminaries

Motivations & Introduction

Model: Spherical Text Embedding

Method: Optimization on the Sphere

**Experiments** 

Conclusions



### Motivations

□ Issues with previous word embedding frameworks:

- Although directional similarity has shown effective for various applications, previous embeddings (e.g. Word2Vec, GloVe, fastText) are trained in the Euclidean space
- □ A gap between training space and usage space: Trained in Euclidean space but used on sphere



Embedding Training in Euclidean Space

Embedding Usage on the Sphere (Similarity, Clustering, etc.)



### Motivations

- What is the consequence of the inconsistency between word embedding training and usage space?
  - □ The objective we optimize during training is not really the one we use
  - Regardless of the different training objective, Word2Vec, GloVe and fastText all optimize the embedding dot product during training, but cosine similarity is what actually used in applications





### Motivations

- What is the consequence of the inconsistency between word embedding training and usage space?
  - □ The objective we optimize during training is not really the one we use
  - E.g. Consider two pairs of words, A: lover-quarrel; B: rock-jazz. Pair B has higher ground truth similarity than pair A in WordSim353 (a benchmark testset)
  - □ Word2Vec assigns higher dot product to pair B, but its cosine similarity is still smaller than pair A

	Metrics	A: lover-quarrel	B: rock-jazz
Training	Dot Product	5.284	<b>&lt;</b> 6.287
Usage	- Cosine Similarity	0.637	> 0.628
		Inconsistency	



### Motivations

- Apart from the training/usage space inconsistency issue, previous embedding frameworks only leverage local contexts to learn word representations
  - Local contexts can only partly define word semantics in unsupervised word embedding learning

t Local contexts of "harmful"

If I hear someone screwing with my car (ie, setting off the **alarm**) and **taunting** me to come out, you can be very sure that my Colt Delta Elite will also be coming with me. It is not the screwing with the car that would get them **shot**, it is the potential physical **danger**. If they are **taunting** like that, it's very possible that they also intend to **rob** me and or do other physically *harmful* things. Here in Houston last year a woman heard the sound of someone ...

### Introduction

□ In this work, we propose "Spherical Text Embedding"

- □ "Spherical": Embeddings are trained on the unit sphere, where vector norms are ignored and directional similarity is directly optimized
- "Text Embedding": Instead of training word embeddings only, we jointly train paragraph (document) embeddings with word embeddings to capture the local and global contexts in text embedding
- **Contributions**:
  - A spherical generative model that jointly exploits word-word (local) and word-paragraph (global) cooccurrence statistics
  - □ An efficient optimization algorithm in the spherical space with convergence guarantee
  - State-of-the-art performances on various text embedding applications



### Outline

 $\bigtriangledown$ 

- Preliminaries
- Motivations & Introduction
- Model: Spherical Text Embedding



- **Experiments**
- Conclusions

## **Model: Spherical Text Embedding**

U We design a generative model on the sphere that follows how humans write articles:

- We first have a general idea of the paragraph/document, and then start to write down each word in consistent with not only the paragraph/document, but also the surrounding words
- Assume a two-step generation process:



## **Model: Spherical Text Embedding**

How to define the generative model in the spherical space?



• We prove a theorem connecting the above generative model with a spherical probability distribution:

**Theorem 1.** When the corpus has infinite vocabulary, *i.e.*,  $|V| \to \infty$ , the analytic forms of  $p(u \mid d) \propto \exp(\cos(u, d))$  and  $p(v \mid u) \propto \exp(\cos(v, u))$  are given by the von Mises-Fisher (vMF) distribution with the prior embedding as the mean direction and constant 1 as the concentration parameter, *i.e.*,

$$\lim_{|V|\to\infty} p(v \mid u) = v \mathsf{MF}_p(v; u, 1), \quad \lim_{|V|\to\infty} p(u \mid d) = v \mathsf{MF}_p(u; d, 1).$$

## **Model: Spherical Text Embedding**

### Understanding the spherical generative model



## **Model: Spherical Text Embedding**

#### □ Training objective:

□ The final generation probability:

 $p(v, u \mid d) = p(v \mid u) \cdot p(u \mid d) = vMF_p(v; u, 1) \cdot vMF_p(u; d, 1)$ 

Maximize the log-probability of a real co-occurred tuple (v, u, d), while minimize that of a negative sample (v, u', d), with a max-margin loss:

$$\mathcal{L}_{\text{joint}}(\boldsymbol{u}, \boldsymbol{v}, \boldsymbol{d}) = \max \left( 0, m - \log \left( c_p(1) \exp(\cos(\boldsymbol{v}, \boldsymbol{u})) \cdot c_p(1) \exp(\cos(\boldsymbol{u}, \boldsymbol{d})) \right) \right) \text{ Positive Sample} \\ + \log \left( c_p(1) \exp(\cos(\boldsymbol{v}, \boldsymbol{u}')) \cdot c_p(1) \exp(\cos(\boldsymbol{u}', \boldsymbol{d})) \right) \right) \text{ Negative Sample} \\ = \max \left( 0, m - \cos(\boldsymbol{v}, \boldsymbol{u}) - \cos(\boldsymbol{u}, \boldsymbol{d}) + \cos(\boldsymbol{v}, \boldsymbol{u}') + \cos(\boldsymbol{u}', \boldsymbol{d}) \right),$$



### Outline

 $\checkmark$ 

- Preliminaries
- Motivations & Introduction
- Model: Spherical Text Embedding
- Method: Optimization on the Sphere
- **Experiments**
- Conclusions

## **Method: Optimization on the Sphere**

□ The constrained optimization problem:

$$\min_{\boldsymbol{\Theta}} \mathcal{L}_{\text{joint}}(\boldsymbol{\Theta}) \quad \text{s.t.} \quad \forall \boldsymbol{\theta} \in \boldsymbol{\Theta} : \|\boldsymbol{\theta}\| = 1 \qquad \boldsymbol{\Theta} = \{\boldsymbol{u}_i\} \Big|_{i=1}^{|V|} \cup \{\boldsymbol{v}_i\} \Big|_{i=1}^{|V|} \cup \{\boldsymbol{d}_i\} \Big|_{i=1}^{|\mathcal{D}|}$$

- Challenge: Parameters must be always updated on the sphere, but Euclidean optimization methods (e.g. SGD) are not constrained on a curvature space
- □ Need to consider the nature of the spherical space

## **Method: Optimization on the Sphere**

#### **Riemannian manifold:**

**D** The sphere is a Riemannian manifold with constant positive curvature



## **Method: Optimization on the Sphere**

**Riemannian optimization with Riemannian SGD:** 

Riemannian gradient:

grad 
$$f(\boldsymbol{x}) \coloneqq \left(I - \boldsymbol{x} \boldsymbol{x}^{\top}\right) \nabla f(\boldsymbol{x})$$

**Exponential mapping (maps from the tangent plane to the sphere):** 

$$\exp_{\boldsymbol{x}}(\boldsymbol{z}) \coloneqq \begin{cases} \cos(\|\boldsymbol{z}\|)\boldsymbol{x} + \sin(\|\boldsymbol{z}\|)\frac{\boldsymbol{z}}{\|\boldsymbol{z}\|}, & \boldsymbol{z} \in T_{\boldsymbol{x}} \mathbb{S}^{p-1} \setminus \{\boldsymbol{0}\}, \\ \boldsymbol{x}, & \boldsymbol{z} = \boldsymbol{0}. \end{cases}$$

**Riemannian SGD:** 

$$\boldsymbol{x}_{t+1} = \exp_{\boldsymbol{x}_t} \left( -\eta_t \operatorname{grad} f(\boldsymbol{x}_t) \right)$$

□ Retraction (first-order approximation of the exponential mapping):

$$R_{\boldsymbol{x}}\left(\boldsymbol{z}
ight)\coloneqqrac{\boldsymbol{x}+\boldsymbol{z}}{\|\boldsymbol{x}+\boldsymbol{z}\|}$$



## **Method: Optimization on the Sphere**

#### **Training details:**

Incorporate angular distances into Riemannian optimization



Multiply the Euclidean gradient with its angular distance from the current point

$$\boldsymbol{x}_{t+1} = R_{\boldsymbol{x}_t} \left( -\eta_t \left( 1 + \frac{\boldsymbol{x}_t^\top \nabla f(\boldsymbol{x}_t)}{\|\nabla f(\boldsymbol{x}_t)\|} \right) \left( I - \boldsymbol{x}_t \boldsymbol{x}_t^\top \right) \nabla f(\boldsymbol{x}_t) \right)$$

## **Method: Optimization on the Sphere**

□ Convergence guarantee of the optimization procedure:

**Theorem 2.** When the update rule given by Equation (7) is applied to  $\mathcal{L}(\boldsymbol{x})$ , and the learning rate satisfies the usual condition in stochastic approximation, *i.e.*,  $\sum_t \eta_t^2 < \infty$  and  $\sum_t \eta_t = \infty$ ,  $\boldsymbol{x}$  converges almost surely to a critical point  $\boldsymbol{x}^*$  and grad  $\mathcal{L}(\boldsymbol{x})$  converges almost surely to 0, *i.e.*,

$$\Pr\left(\lim_{t\to\infty}\mathcal{L}(\boldsymbol{x}_t)=\mathcal{L}(\boldsymbol{x}^*)\right)=1, \quad \Pr\left(\lim_{t\to\infty}\operatorname{grad}\mathcal{L}(\boldsymbol{x}_t)=0\right)=1.$$



### Outline

- Preliminaries
- Motivations & Introduction
- Model: Spherical Text Embedding

Method: Optimization on the Sphere

Experiments

Conclusions

### Experiments

- □ Word similarity:
  - □ Train 100-d word embedding on the latest Wikipedia dump (~13G)
  - Compute embedding **cosine similarity** between word pairs to obtain a ranking of similarity
  - Benchmark datasets contain human rated similarity scores
  - □ The more similar the two rankings are, the better embedding reflects human thoughts
  - **Spearman's rank correlation** is used to measure the ranking similarity

### Experiments

#### □ Word similarity baselines:

- **Euclidian Space:** 
  - □ Word2Vec: Distributed representations of words and phrases and their compositionality. In NIPS, 2013
  - GloVe: Glove: Global vectors for word representation. In EMNLP, 2014
  - □ fastText: Enriching word vectors with subword information. TACL, 2017
  - BERT: Pre-training of deep bidirectional transformers for language understanding. In NAACL, 2019
- Poincare Space:
  - Deincaré glove: Hyperbolic word embeddings. In ICLR, 2019
- **Spherical Space:** 
  - □ JoSE (our full model)

### Experiments

#### □ Word similarity results:

Embedding Space	Model	WordSim353	MEN	SimLex999	
Euclidean	Word2Vec GloVe fastText BERT	0.711 0.598 0.697 0.477	0.726 0.690 0.722 0.594	0.311 0.321 0.303 0.287	
Poincaré	Poincaré GloVe	0.623	0.652	0.321	
Spherical	JoSE	0.739	0.748	0.339	

Table 1: Spearman rank correlation on word similarity evaluation.

- ❑ Why does BERT fall behind on this task?
- BERT learns contextualized representations, but word similarity is conducted in a context-free manner
- BERT is optimized on specific downstream tasks like predicting masked words and sentence relationships, which have no direct relation to word similarity



### Experiments

#### Document Clustering:

- □ Train document embedding on 20News dataset (20 classes)
- Perform K-means and Spherical K-means (SK-means)
- Metrics: Mutual Information (MI), Normalized Mutual Information (NMI), Adjusted Rand Index (ARI), and Purity
- Run clustering 10 times and report the above metrics with mean and standard deviation



### Experiments

- Document clustering baselines:
  - **Euclidian Space:** 
    - Avg. Word2Vec: Use the averaged word embedding of Word2Vec as document embedding
    - □ SIF: Simple but tough-to-beat baseline for sentence embeddings. In ICLR, 2017
    - BERT: Pre-training of deep bidirectional transformers for language understanding. In NAACL, 2019
    - Doc2Vec: Distributed representations of sentences and documents. In ICML, 2014
  - **Spherical Space:** 
    - □ JoSE (our full model)

### Experiments

#### Document clustering results:

Table 2. Document clustering evaluation on the 20 Newsgroup dataset.					
Embedding	Clus. Alg.	MI	NMI	ARI	Purity
Avg. W2V	K-Means SK-Means	$\begin{array}{c} 1.299 \pm 0.031 \\ 1.328 \pm 0.024 \end{array}$	$\begin{array}{c} 0.445 \pm 0.009 \\ 0.453 \pm 0.009 \end{array}$	$\begin{array}{c} 0.247 \pm 0.008 \\ 0.250 \pm 0.008 \end{array}$	$\begin{array}{c} 0.408 \pm 0.014 \\ 0.419 \pm 0.012 \end{array}$
SIF	K-Means SK-Means	$\begin{array}{c} 0.893 \pm 0.028 \\ 0.958 \pm 0.012 \end{array}$	$\begin{array}{c} 0.308 \pm 0.009 \\ 0.322 \pm 0.004 \end{array}$	$\begin{array}{c} 0.137 \pm 0.006 \\ 0.164 \pm 0.004 \end{array}$	$\begin{array}{c} 0.285 \pm 0.011 \\ 0.331 \pm 0.005 \end{array}$
BERT	K-Means SK-Means	$\begin{array}{c} 0.719 \pm 0.013 \\ 0.854 \pm 0.022 \end{array}$	$\begin{array}{c} 0.248 \pm 0.004 \\ 0.289 \pm 0.008 \end{array}$	$\begin{array}{c} 0.100 \pm 0.003 \\ 0.127 \pm 0.003 \end{array}$	$\begin{array}{c} 0.233 \pm 0.005 \\ 0.281 \pm 0.010 \end{array}$
Doc2Vec	K-Means SK-Means	$\begin{array}{c} 1.856 \pm 0.020 \\ 1.876 \pm 0.020 \end{array}$	$\begin{array}{c} 0.626 \pm 0.006 \\ 0.630 \pm 0.007 \end{array}$	$\begin{array}{c} 0.469 \pm 0.015 \\ 0.494 \pm 0.012 \end{array}$	$\begin{array}{c} 0.640 \pm 0.016 \\ 0.648 \pm 0.017 \end{array}$
JoSE	K-Means SK-Means	$\begin{array}{c} 1.975 \pm 0.026 \\ \textbf{1.982} \pm 0.034 \end{array}$	$0.663 \pm 0.008$ <b>0.664</b> $\pm$ 0.010	$\begin{array}{c} 0.556 \pm 0.018 \\ \textbf{0.568} \pm 0.020 \end{array}$	$\begin{array}{c} 0.711 \pm 0.020 \\ \textbf{0.721} \pm 0.029 \end{array}$

Table 2: Document clustering evaluation on the 20 Newsgroup dataset

Embedding quality is generally more important than clustering algorithms: 

- Using spherical K-Means only gives marginal performance boost over K-Means
- JoSE embedding remains optimal regardless of clustering algorithms



### Experiments

- Document Classification:
  - Train document embedding on 20News (20 classes) and Movie review (2 classes) dataset
  - Perform k-NN classification (k=3)
  - Metrics: Macro-F1 & Micro-F1
  - Baselines: Same with document clustering

### Experiments

Document classification results:

Table 3: Document classification evaluation using $k$ -NN ( $k = 3$ ).						
Embedding	20 Newsgroup		Movie Review			
Linocdunig	Macro-F1	Micro-F1	Macro-F1	Micro-F1		
Avg. W2V	0.630	0.631	0.712	0.713		
SIF	0.552	0.549	0.650	0.656		
BERT	0.380	0.371	0.664	0.665		
Doc2Vec	0.648	0.645	0.674	0.678		
JoSE	0.703	0.707	0.764	0.765		

- □ We observe similar comparison results with k from [1, 10]
- k-NN is non-parametric and directly reflect how well the topology of the embedding space captures document-level semantics

### Experiments

□ Training efficiency:

Table 4: Training time	(per iteration) o	on the latest	Wikipedia d	lump.
2	<b>1</b>		1	1

Word2Vec	GloVe	fastText	BERT	Poincaré GloVe	JoSE
0.81 hrs	0.85 hrs	2.11 hrs	> 5 days	1.25 hrs	0.73 hrs

#### □ Why is JoSE training efficient?

Other models' objectives contain many non-linear operations (Word2Vec & fastText's objectives involve exponential functions; GloVe's objective involves logarithm functions), while JoSE only has linear terms in the objective



### Outline

- Preliminaries
- Motivations & Introduction
- Model: Spherical Text Embedding
- Method: Optimization on the Sphere
- **Experiments**



## Conclusions

□ In this work, we introduce a spherical text embedding framework

- Address the discrepancy between the training procedure and the practical usage of text embedding
- Introduce a spherical generative model to jointly learn word and paragraph embedding
- Develop an efficient Riemannian optimization method to train text embedding on the unit hypersphere
- Achieves state-of-the-art performances on common text embedding applications
- **G** Future work:
  - **Exploit spherical embedding space for other tasks like lexical entailment**
  - Incorporate other signals such as subword information into spherical text embedding
  - Benefit other supervised tasks: Word embedding is commonly used as the first layer in DNN
    - □ Add norm constraints to word embedding layer
    - □ Apply Riemannian optimization when fine-tuning the word embedding layer



## Thanks!