

Deep Learning

8.2 Word Embeddings (CBOW)

Dr. Konda Reddy Mopuri
kmopuri@iittp.ac.in
Dept. of CSE, IIT Tirupati

Word Embeddings

- ① Text processing with NNs require to encoding into vectors

One-hot encoding

① One-hot encoding: N words encoded as binary vectors of length N

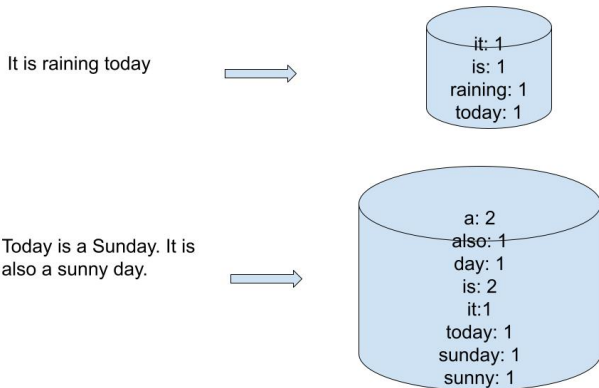
Dictionary

Word Representation

| | | | | | | |
|------|---|---|---|-------|---|---|
| A | 1 | 0 | 0 | | 0 | 0 |
| Bus | 0 | 1 | 0 | | 0 | 0 |
| Cat | 0 | 0 | 1 | | 0 | 0 |
| ⋮ | | | | | | |
| Tide | 0 | 0 | 0 | | 1 | 0 |
| Zone | 0 | 0 | 0 | | 0 | 1 |

Bag of Words (BoW)

- 1 Bag of Words: Collection and frequency of words



Drawbacks

- ① Space inefficient

Drawbacks

- ① Space inefficient
- ② Word order is lost

Drawbacks

- ① Space inefficient
- ② Word order is lost
- ③ Doesn't capture language structure

Word Embeddings: idea

- ① Learn embeddings from the words into vectors: $W(\text{word})$

Word Embeddings: idea

- ① Learn embeddings from the words into vectors: $W(\text{word})$
- ② Expecting that similar words fall nearby in the space

Word Embeddings

- ① What is the dimension of the embedding?

Word Embeddings

- ① What is the dimension of the embedding?
- ② Trade-off: greater capacity vs. efficiency

Word Embeddings

- ① Finding W : as a part of a prediction task that involves neighboring words

Word Embeddings: word2vec


① T Mikolov et al. (2013)

Word Embeddings: word2vec


- ① T Mikolov et al. (2013)
- ② Predict words from the context

Word Embeddings: word2vec

- ① T Mikolov et al. (2013)
- ② Predict words from the context
- ③ Two versions: Continuous Bag of Words (CBow) and Skip-gram



CBOW




Skip-gram

Word Embeddings: CBoW

- ① Considers the embeddings of 'n' words before and 'n' words after the target word

Word Embeddings: CBoW


- ① Considers the embeddings of 'n' words before and 'n' words after the target word
- ② Adds them (order is lost) for predicting the target word



CBoW

Word Embeddings: CBoW

The dog slept on couch



Word Embeddings: CBoW

- ① Size of the vocabulary = V

Word Embeddings: CBoW

- ① Size of the vocabulary = V
- ② Dimension of the embeddings = N

Word Embeddings: CBoW

- ① Size of the vocabulary = V
- ② Dimension of the embeddings = N
- ③ Input layer will have the weight matrix $W_{N \times V}$

Word Embeddings: CBoW

- ① Size of the vocabulary = V
- ② Dimension of the embeddings = N
- ③ Input layer will have the weight matrix $W_{N \times V}$
- ④ Projects the words in to N dimensional space

Word Embeddings: CBoW

- ① Size of the vocabulary = V
- ② Dimension of the embeddings = N
- ③ Input layer will have the weight matrix $W_{N \times V}$
- ④ Projects the words in to N dimensional space
- ⑤ Projections of all the $(2n)$ words in context are summed (result is an N d vector)

Word Embeddings: CBoW

- ① Next layer has a weight matrix $W'_{V \times N}$

Word Embeddings: CBoW

- ① Next layer has a weight matrix $W'_{V \times N}$
- ② Projects the accumulated embeddings onto the vocabulary

Word Embeddings: CBoW

- ① Next layer has a weight matrix $W'_{V \times N}$
- ② Projects the accumulated embeddings onto the vocabulary
- ③ That is, V - way classification \rightarrow (after a softmax) maximizes the probability for the target word

Word Embeddings: CBoW


- ① $W_{N \times V}$ or $W'_{V \times N}$ can be considered as the word embeddings

Word Embeddings: CBoW

- ① $W_{N \times V}$ or $W'_{V \times N}$ can be considered as the word embeddings
- ② Or, take the average of both the representations

Word Embeddings: Skipgram


- ① Predicts surrounding words given current word



Skip-gram

Word Embeddings: Skipgram

- 1 Predicts surrounding words given current word
- 2 Pick a word in the context randomly, and predict that word



Skip-gram

Word Embeddings: interesting results

① $W(\text{Paris}) - W(\text{France}) + W(\text{Italy}) = W(\text{Rome})$

Word Embeddings: interesting results

- ① $W(\text{Paris}) - W(\text{France}) + W(\text{Italy}) = W(\text{Rome})$
- ② $W(\text{Man}) - W(\text{Woman}) + W(\text{King}) = W(\text{Queen})$

Word Embeddings: Applications

- ① Key for the success of many NLP tasks such as PoS tagging, parsing, semantic role labeling, etc.

Word Embeddings: Applications

- ① Key for the success of many NLP tasks such as PoS tagging, parsing, semantic role labeling, etc.
- ② Can serve projecting multi-modal data (e.g. multiple languages, images and text, etc.)

References

- ① Mikolov, Tomas; et al. (2013). "Efficient Estimation of Word Representations in Vector Space". arXiv:1301.3781