

Word Embeddings for Named Entity Recognition

Fabienne Braune¹

¹LMU Munich

December 13th, 2017

Outline

- 1 Named Entity Recognition
- 2 Feedforward Neural Networks: recap
- 3 Neural Networks for Named Entity Recognition
- 4 Example
- 5 Adding Pre-trained Word Embeddings
- 6 Word2Vec models
- 7 Bilingual Word Embeddings

NAMED ENTITY RECOGNITION

Task

Find segments of **entity mentions** in input text and tag with **labels**.

Example inputs:

- *Trump attacks BMW and Mercedes*
- *U.N. official Ekeus heads for Baghdad*

Example labels (coarse grained):

- persons **PER**
- locations **LOC**
- organizations **ORG**
- names **NAME**
- other **MISC**

Labeled data

Desired outputs:

- *Trump PER attacks BMW ORG and Mercedes ORG*
- *U.N. ORG official Ekeus PER heads for Baghdad LOC*

Example annotations (CoNLL-2003):

Surface	POS	Sh-synt	Tag
U.N.	NNP	I-NP	I-ORG
official	NN	I-NP	O
Ekeus	NNP	I-NP	I-PER
heads	VBZ	I-VP	O
for	IN	I-PP	O
Baghdad	NNP	I-NP	I-LOC
.	.	O	O

Classification-based approaches

Given input segment, train classifier to tell:

- Is this segment a Named Entity ?
- Give the corresponding Tag

Classification task:

Trump attacks BMW and Mercedes

Is **Trump** a named entity ?

Yes, it is a **person (PER)**

Classification-based approaches

- Classifier combination with engineered features (Florian et al. 2003)
 - ▶ Manually engineer features
 - ▶ Use large (external) gazetteer
 - ▶ Combine classifiers (ME, MRR, HMM) trained on annotated data
 - ▶ 88.76 F1


- Semi-supervised learning with linear models (Ando and Zhang 2005)
 - ▶ Train linear model on annotated data
 - ▶ Add non-annotated data
 - ▶ 89.31 F1

Classification-based approaches

- Use feedforward neural networks (Collobert et al. 2011):
 - ▶ With raw words 81.74
 - ▶ With pre-trained word embeddings 88.67
 - ▶ Using a gazetteer 89.59
- Use sequential models:
 - ▶ Linear Chain CRF (linear)
 - ▶ LSTM networks (deep)
 - Achieve best performance but not covered here

FEEDFORWARD NEURAL NETWORKS: RECAP

Motivation



Cannot be solved using a linear model

Motivation

a	b	a XNOR b		
0	0	1	b • 0	• 1
0	1	0		
1	0	0		
1	1	1	• 1	• 0
				a

Features : a, b

Feature values : binary

Cannot be solved using a linear model


Motivation

Linear models not suited to learn non-linear decision boundaries.

Neural networks can do that

- Through composition of **non-linear** functions
- Learn relevant features from (almost) raw text
 - No need for manual feature engineering
 - learned by network

Feedforward Neural Network



Computation of hidden layer \mathbf{H} :

- $A_1 = \sigma(X \cdot \Theta_1)$
- $A_2 = \sigma(X \cdot \Theta_2)$
- $B_0 = 1$ (bias term)

Computation of output unit $h(X)$:

- $h(X) = \sigma(\mathbf{H} \cdot \Theta_3)$


Feedforward Neural Network

Feedforward neural network with:

- 1 input layer X (feature vector)
- 2 weight matrices $U = (\Theta_1, \Theta_2)$ and $V = \Theta_3$
- 1 hidden layer \mathbf{H} composed of:
 - ▶ 2 activations $A_1 = \sigma(Z_1)$ and $A_2 = \sigma(Z_2)$ where:
 - ★ $Z_1 = X \cdot \Theta_1$
 - ★ $Z_2 = X \cdot \Theta_2$
- 1 output unit $h(X) = \sigma(Z_3)$ where:
 - ▶ $Z_3 = \mathbf{H} \cdot \Theta_3$

Non-linear activation function

The **sigmoid function** $\sigma(Z)$ is often used



Feedforward neural network

Trump attacks BMW and Mercedes


Binray NER task: Is the segment from position 1 to 2 a Named Entity?

Neural network: $h(X) = \sigma(\mathbf{H} \cdot \Theta_n)$, with:

$$\mathbf{H} = \begin{bmatrix} B_0 = 1 \\ A_1 = \sigma(X \cdot \Theta_1) \\ A_2 = \sigma(X \cdot \Theta_2) \\ \dots \\ A_j = \sigma(X \cdot \Theta_j) \end{bmatrix}$$

Prediction: If $h(X) > 0.5$, yes. Otherwise, no.

Feedforward Neural Network



If **weights** are all **random** output will be random

→ Predictions will be **bad**

→ Get the right weights

Getting the right weights

Training: Find weight matrices $U = (\Theta_1, \Theta_2)$ and $V = \Theta_3$ such that $h(X)$ is the **correct answer** as many times as possible.

- Given a set T of training examples t_1, \dots, t_n with **correct labels** \mathbf{y}_i , find $U = (\Theta_1, \Theta_2)$ and $V = \Theta_3$ such that $h(X) = \mathbf{y}_i$ for as many t_i as possible.
- Computation of $h(X)$ called **forward propagation**
- $U = (\Theta_1, \Theta_2)$ and $V = \Theta_3$ with error **back propagation**

Multi-class classification

- More than two labels
- Instead of “yes” and “no”, predict $c_i \in C = \{c_1, \dots, c_k\}$
- **NER**: Is this segment a **location**, **name**, **person** ...
- **Use k output units**, where k is number of classes
 - ▶ Output layer instead of unit
 - ▶ Use softmax to obtain value between 0 and 1 for each class
 - ▶ Highest value is right class

NEURAL NETWORKS FOR NER

Classification-based NER

Given input segment, train classifier to tell:

- Is this segment a Named Entity ?
- Give the corresponding Tag

Classification task:

Trump attacks BMW and Mercedes

Is **Trump** a named entity ?

Yes, it is a **person (PER)**

Labeled data

Desired outputs:

- *Trump PER attacks BMW ORG and Mercedes ORG*
- *U.N. ORG official Ekeus PER heads for Baghdad LOC*

Annotation:

Surface	POS	Sh-synt	Tag
U.N.	NNP	I-NP	I-ORG
official	NN	I-NP	O
Ekeus	NNP	I-NP	I-PER
heads	VBZ	I-VP	O
for	IN	I-PP	O
Baghdad	NNP	I-NP	I-LOC
.	.	O	O

Feedforward Neural Network for NER

Training example: *Trump attacks **BMW** (ORG) and Mercedes*

Neural network input:

Look at word window around **BMW**

→ **Trump**₋₂ **attacks**₋₁ **BMW** **and**₁ **Mercedes**₂

→ Lookup **feature representation** (LT_i) for window

Give LT_i as input to Feedforward Neural Network


Neural network training:

Predict corresponding label (forward propagation)

→ should be **organization** (ORG)

Train weights by backpropagating error

Feedforward Neural Network for NER



Input: word features LT_i

Output: **predicted label**

Note: Bias terms omitted for simplicity

Feedforward Neural Network

Input layer (X): Word features $LT1, LT2, LT3, LT4$

Weight matrices U, V

Hidden layer (H): $\sigma(X \cdot U + d)$

Output layer (O): $H \cdot V + b$

Prediction: $h(X) = \text{softmax}(O)$

- Predicted class is the one with highest probability (given by softmax)

Weight training

Training: Find weight matrices U and V such that $h(X)$ is the **correct answer** as many times as possible.

- Given a set T of training examples t_1, \dots, t_n with **correct labels** y_i , find U and V such that $h(X) = y_i$ for as many t_i as possible.
- Computation of $h(X)$ with **forward propagation**
- C , U and V with error **back propagation**

Training data

Training example	POS	Sh-synt	Tag
U.N.	NNP	I-NP	I-ORG
official	NN	I-NP	O
Ekeus	NNP	I-NP	I-PER
heads	VBZ	I-VP	O
for	IN	I-PP	O
Baghdad	NNP	I-NP	I-LOC
.	.	O	O

Weight training


Training: Find weight matrices U and V such that $h(X)$ is the **correct answer** as many times as possible.

- Computation of $h(X)$ with **forward propagation**
- U and V with error **back propagation**

For each batch of training examples

- 1 Forward propagation to get predictions
- 2 Backpropagation of error
 - ▶ Gives gradient of E wrt. **weights**
- 3 Modify weights (gradient descent)
- 4 Goto 1 until convergence

Lookup Layer



Lookup Layer

- Each word encoded into index vector $w_i = \begin{bmatrix} 0 \\ 1 \\ 0 \\ 0 \\ 0 \end{bmatrix}$
- LT_i is dot product of **weight matrix** C with index of w_i
→ C is **shared** among all words

Dot product with (trained) weight vector

$W = \{\text{the,cat,on,table,chair}\}$

$$w_{table} = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 1 \\ 0 \end{bmatrix} \quad C = \begin{bmatrix} 0.02 & 0.1 & 0.05 & 0.03 & 0.01 \\ 0.15 & 0.2 & 0.01 & 0.02 & 0.11 \\ 0.03 & 0.1 & 0.04 & 0.04 & 0.12 \end{bmatrix}$$

$$LT_{table} = w_{table} \cdot C^T = \begin{bmatrix} 0.03 \\ 0.02 \\ 0.04 \end{bmatrix}$$

Words get mapped to lower dimension

→ Hyperparameter to be set

Dot product with (initial) weight vector


$$W = \{\text{the,cat,on,table,chair}\}$$

$$w_{table} = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 1 \\ 0 \end{bmatrix} \quad C = \begin{bmatrix} 0.01 & 0.01 & 0.01 & 0.01 & 0.01 \\ 0.01 & 0.01 & 0.01 & 0.01 & 0.01 \\ 0.01 & 0.01 & 0.01 & 0.01 & 0.01 \end{bmatrix}$$

$$LT_{table} = w_{table} \cdot C^T = \begin{bmatrix} 0.01 \\ 0.01 \\ 0.01 \end{bmatrix}$$

Feature vectors same for all words.

Feedforward Neural Network with Lookup Table



Note: Bias terms omitted for simplicity

Dot product with (trained) weight vector

$W = \{\text{the, cat, on, table, chair}\}$

$$w_{table} = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 1 \\ 0 \end{bmatrix} \quad C = \begin{bmatrix} 0.02 & 0.1 & 0.05 & 0.03 & 0.01 \\ 0.15 & 0.2 & 0.01 & 0.02 & 0.11 \\ 0.03 & 0.1 & 0.04 & 0.04 & 0.12 \end{bmatrix}$$

$$LT_{table} = w_{table} \cdot C^T = \begin{bmatrix} 0.03 \\ 0.02 \\ 0.04 \end{bmatrix}$$

Each word gets a **specific** feature vector

Training data

Training example	POS	Sh-synt	Tag
U.N.	NNP	I-NP	I-ORG
official	NN	I-NP	O
Ekeus	NNP	I-NP	I-PER
heads	VBZ	I-VP	O
for	IN	I-PP	O
Baghdad	NNP	I-NP	I-LOC
.	.	O	O

- Lookup vector C trained with NER training data
- Word feature vectors are trained towards NER

EXAMPLE

Example

Trump PER attacks BMW ORG and Mercedes ORG

$W = \{\text{Trump, BMW, Mercedes, attacks, and}\}$

$$w_{\text{Trump}} = \begin{bmatrix} 1 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}$$

$$w_{\text{attacks}} = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 1 \\ 0 \end{bmatrix}$$

$$w_{\text{BMW}} = \begin{bmatrix} 0 \\ 1 \\ 0 \\ 0 \\ 0 \end{bmatrix}$$

$$w_{\text{and}} = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 1 \end{bmatrix}$$

$$w_{\text{Mercedes}} = \begin{bmatrix} 0 \\ 0 \\ 1 \\ 0 \\ 0 \end{bmatrix}$$

Example

Window: Trump attacks **BMW** and Mercedes


$$W_{window} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \end{bmatrix}$$

$$C = \begin{bmatrix} 0.01 & 0.8 & 0.05 & 0.02 & 0.01 \\ 0.03 & 0.2 & 0.08 & 0.01 & 0.02 \\ 0.04 & 0.1 & 0.04 & 0.02 & 0.04 \end{bmatrix}$$

C is randomly initialized

$$LT = W_{window} \cdot C^T$$

Example



Output of lookup table given as input
Note: Bias terms omitted for simplicity

Example


$$LT = \begin{bmatrix} 0.01 & 0.03 & 0.04 \\ 0.05 & 0.08 & 0.04 \\ 0.01 & 0.02 & 0.04 \\ 0.8 & 0.2 & 0.1 \\ 0.02 & 0.01 & 0.4 \end{bmatrix} \quad U = \begin{bmatrix} 0.04 & 0.6 & 0.01 & 0.02 & 0.06 & 0.03 \\ 0.01 & 0.9 & 0.02 & 0.05 & 0.03 & 0.05 \\ 0.02 & 0.3 & 0.05 & 0.07 & 0.09 & 0.01 \\ 0.02 & 0.4 & 0.02 & 0.03 & 0.04 & 0.02 \\ 0.01 & 0.8 & 0.01 & 0.01 & 0.03 & 0.07 \end{bmatrix}$$

U is randomly initialized

$$\mathbf{Z} = LT^T \cdot U^T$$

$$\mathbf{A} = \sigma(\mathbf{Z})$$

Example



Output of lookup table given as input
Note: Bias terms omitted for simplicity

Example

- Repeat same procedure for each hidden layer
- Apply softmax on output (last) layer
- **Predict** label


Example

- Compute error between **prediction** (e.g. LOCATION) and **true label**
 - Given in training data (BMW is ORG)
- Backpropagate error through network and adjust weights
- Redo same procedure with adjusted weights
- Stop at convergence
 - Or early stopping on held-out dataset

ADDING PRE-TRAINED WORD EMBEDDINGS


Word Embeddings

- Representation of words in vector space



Word Embeddings

- Similar words are close to each other
→ Similarity is the cosine of the angle between two word vectors



Learning word embeddings

Count-based methods:

- Compute cooccurrence statistics
- Learn high-dimensional representation
- Map sparse high-dimensional vectors to small dense representation

Neural networks:


- Predict a word from its neighbors
- Learn (small) embedding vectors

Word vectors with Neural Networks

- LM Task: Given k previous words, predict the current word
 - For each word w in V , model $P(w_t | w_{t-1}, w_{t-2}, \dots, w_{t-n})$
 - **Learn embeddings C of words**
 - Input for task

- Task: Given k context words, predict the current word
 - **Learn embeddings C of words**


Network architecture



Given words w_{t-2} , w_{t-1} , w_{t+1} and w_{t+2} , predict w_t

Note: Bias terms omitted for simplicity


Network architecture



We want the **context vectors** \rightarrow embed words in shared space

Note: Bias terms omitted for simplicity

Getting the Word Embeddings



Same as **lookup table** but trained on a **language model task** (predict w_t)
NER lookup table was trained on **NER task** (predict NE label)

Word Embeddings for NER

- Train word embeddings using language model task:
 - Labels are words w_t
 - No need for NER training data
 - Use large amounts of **non-annotated** data
- Replace lookup table **C** (randomly initialized) with **C** (pre-trained)

Example

Window: Trump attacks **BMW** and Mercedes

$$W_{window} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \end{bmatrix}$$

$$C = \begin{bmatrix} 0.01 & 0.8 & 0.05 & 0.02 & 0.01 \\ 0.03 & 0.2 & 0.08 & 0.01 & 0.02 \\ 0.04 & 0.1 & 0.04 & 0.02 & 0.04 \end{bmatrix}$$

C is **randomly initialized**

Before NER training, word embeddings are **very bad**.

After NER training, word embeddings are **good for NER**.

Example

Window: Trump attacks **BMW** and Mercedes

$$W_{window} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \end{bmatrix}$$

$$C = \begin{bmatrix} 0.01 & 0.8 & 0.05 & 0.02 & 0.01 \\ 0.03 & 0.2 & 0.08 & 0.01 & 0.02 \\ 0.04 & 0.1 & 0.04 & 0.02 & 0.04 \end{bmatrix}$$

C is **pre-trained** on LM task

Before NER training, word embeddings are **good word embeddings**.

NER trained word embeddings

Word embeddings trained on **NER task**

- (Collobert et al. 2011)

→ **Small** amount of **annotated** data.

- Closest words to **France**
 - ▶ Persuade
 - ▶ Faw
 - ▶ Blackstock
- Closest words to **XBOX**
 - ▶ Decadent
 - ▶ Divo
 - ▶ Versus

NER trained word embeddings

Word embeddings trained on **LM task**

→ **Large** amount of **non-annotated** data.

- Closest words to **France**
 - ▶ Austria
 - ▶ Belgium
 - ▶ Germany
- Closest words to **XBOX**
 - ▶ Amiga
 - ▶ Playstation
 - ▶ MSX

Results

Feedforward Neural Networks for NER (Collobert et al. 2011):

- With raw words 81.74
- With pre-trained word embeddings 88.67
- Using a gazetteer 89.59

Classifier combination with engineered features (Florian et al. 2003)

- 88.76 F1

Semi-supervised learning with linear models (Ando and Zhang 2005)

- 89.31 F1


Results

- **Pre-trained** word embeddings yield comparable results to state of the art NER systems
- To beat the best system, **additional features** are needed
 - ▶ Indicate if word is in Gazetteer or not

WORD EMBEDDINGS

Word Embeddings

- Similar words are close to each other
→ Similarity is the cosine of the angle between two word vectors



Learning word embeddings

Count-based methods:

- Compute cooccurrence statistics
- Learn high-dimensional representation
- Map sparse high-dimensional vectors to small dense representation

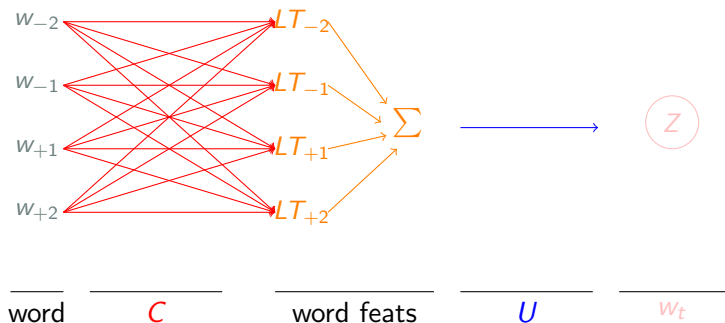
Neural networks:

- Predict a word from its neighbors
- Learn (small) embedding vectors

Word2Vec


- Software train word embeddings (Mikolov. 2013)
 - very fast
- Two models:
 - ▶ **BOW model:**
 - ★ Input is w_{t+2} , w_{t+1} , w_{t-1} and w_{t-2}
 - ★ Prediction is w_t
 - ▶ **Skip-gram model:**
 - ★ Input is w_t
 - ★ Prediction is w_{t+2} , w_{t+1} , w_{t-1} and w_{t-2}

Learning word embeddings with CBOW



Note: Bias terms omitted for simplicity

Learning word embeddings with skip-gram




Note: Bias terms omitted for simplicity

BILINGUAL WORD EMBEDDINGS

Bilingual Word Spaces

Representation of words in two languages in same semantic space:


- Each word is one dimension
- Each word represented respective to all others



Bilingual Word Spaces

Representation of words in two languages in same semantic space:

- Similar words are close to each other
- Given by cosine



Exercise


How is this related to translation?

Learning Bilingual Word Embeddings

- Learn monolingual word embeddings and map using seed lexicon
Mikolov et al. (2013); Faruqui and Dyer (2014); Lazaridou et al. (2015)
Need seed lexicon
- Learn bilingual embeddings or lexicon from document-aligned data
Vulic and Moens (2015); Vulic and Korhonen (2016)
Need document-aligned data
- Learn bilingual embeddings from parallel data
Hermann and Blunsom (2014), Gouws et al. (2015), Gouws and Sjøgaard (2015), Duong et al. (2016)
Need for parallel data


Post-hoc mapping (with seed lexicon)

- Learn monolingual word embeddings
- Learn a linear mapping W



Post-hoc mapping

- Project source words into target space



Post-hoc Mapping with seed lexicon

- ① Train **monolingual** word embeddings (Word2vec) in **English**
 - ▶ Need **English** monolingual data
- ② Train **monolingual** word embeddings (Word2vec) in **German**
 - ▶ Need **German** monolingual data
- ③ Learn mapping **W** using a seed lexicon
 - ▶ Need a list of **5000 English words and their translation**

Learning W with Ridge Regression

Ridge regression (Mikolov et al. (2013))

$$\mathbf{W}^* = \arg \min_{\mathbf{W}} \sum_i^n \| \mathbf{x}_i \mathbf{W} - \mathbf{y}_i \|^2$$

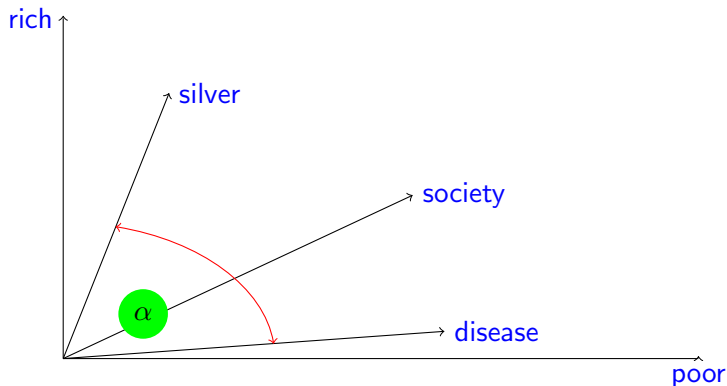
\mathbf{x}_i : **embedding** of i-th **source** (English) word in the seed lexicon.

\mathbf{y}_i : **embedding** of i-th **target** (German) word in the seed lexicon.

Learning W with Ridge Regression

x_i : **embedding** of i -th **source** (English) word in the seed lexicon.

→ **vector** representing **silver** in **monolingual** word embedding



Learning W with Ridge Regression

Ridge regression (Mikolov et al. (2013))

$$\mathbf{W}^* = \arg \min_{\mathbf{W}} \sum_i^n \| \mathbf{x}_i \mathbf{W} - \mathbf{y}_i \|^2$$

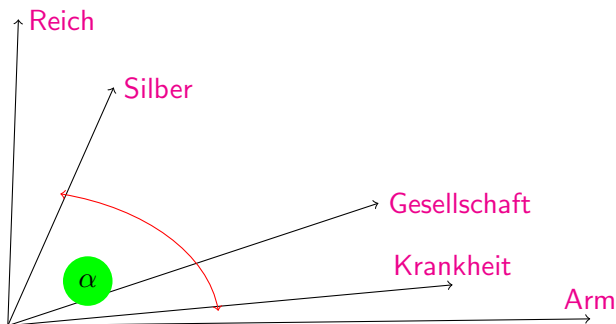
\mathbf{x}_i : **embedding** of i-th **source** (English) word in the seed lexicon.

\mathbf{y}_i : **embedding** of i-th **target** (German) word in the seed lexicon.

Learning W with Ridge Regression

y_i : **embedding** of i -th **target** (German) word in the seed lexicon.

→ **vector** representing **Silber** in **monolingual** word embedding



Learning W with Ridge Regression

Ridge regression (Mikolov et al. (2013))

$$\mathbf{W}^* = \arg \min_{\mathbf{W}} \sum_i^n \| \mathbf{x}_i \cdot \mathbf{W} - \mathbf{y}_i \|^2$$

- Predict projection \mathbf{y}^* by computing $\mathbf{x}_i \cdot \mathbf{W}$
- Compute **squared error** between \mathbf{y}^* and \mathbf{y}_i
 - ▶ Correct translation t_i given in seed lexicon
 - ▶ Vector representation \mathbf{y}_i is given by embedding of t_i
- Find \mathbf{W} such that squared error over training set is minimal

Adding Regularization

If \mathbf{W} is too complex the model overfits the data


→ Add regularization term that keeps \mathbf{W} small

→ Add weighted norm of \mathbf{W} to cost function

$$\mathbf{W}^* = \arg \min_{\mathbf{W}} \sum_i^n \| \mathbf{x}_i \cdot \mathbf{W} - \mathbf{y}_i \|^2 + \lambda \| \mathbf{W} \|^2$$

Bilingual lexicon induction

- Task to evaluate bilingual word embeddings extrinsically
- Given a set of source words, find the corresponding translations:
 - ▶ Given **silver**, find its vector in the BWE
 - ▶ Retrieve the **German** word whose vector is closest (cosine distance)



Bilingual lexicon induction with ridge regression

Data: WMT 2011 training data for English, Spanish, Czech

Seed: 5000 most frequent words translated with Google Translate

Test: 1000 next frequent words translated with Google Translate

→ Removed digits, punctuation and transliterations

Languages	top-1	top-5
En-Es	33 %	51 %
Es-En	35 %	50 %
En-Cz	27 %	47 %
Cz-En	23 %	42 %
+ Es-En	53 %	80 %

→ with spanish google news

Recap

- Using neural networks for NER yields good results using (almost) raw representations of words
- Example feedforward neural network for NER
- Word embeddings can be learned automatically on large amounts of non-annotated data
- Giving pre-trained word embeddings as input to neural networks improve end-to-end task

References I

- Duong, L., Kanayama, H., Ma, T., Bird, S., and Cohn, T. (2016). Learning crosslingual word embeddings without bilingual corpora. In *Proc. EMNLP*.
- Faruqui, M. and Dyer, C. (2014). Improving vector space word representations using multilingual correlation. In *Proc. EACL*.
- Gouws, S., Bengio, Y., and Corrado, G. (2015). Bilbowa: Fast bilingual distributed representations without word alignments. In *Proc. ICML*.
- Gouws, S. and Søgaard, A. (2015). Simple task-specific bilingual word embeddings. In *Proc. NAACL*.
- Hermann, K. M. and Blunsom, P. (2014). Multilingual models for compositional distributed semantics. In *Proc. ACL*, pages 58–68, Baltimore, Maryland. Association for Computational Linguistics.
- Lazaridou, A., Dinu, G., and Baroni, M. (2015). Hubness and pollution: Delving into cross-space mapping for zero-shot learning. In *Proc. ACL*.

References II

- Mikolov, T., Le, Q. V., and Sutskever, I. (2013). Exploiting similarities among languages for machine translation. *CoRR*, abs/1309.4168.
- Vulic, I. and Korhonen, A. (2016). On the Role of Seed Lexicons in Learning Bilingual Word Embeddings. In *Proc. ACL*, pages 247–257.
- Vulic, I. and Moens, M. (2015). Bilingual word embeddings from non-parallel document-aligned data applied to bilingual lexicon induction. In *Proc. ACL*.