Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

Natural Language Processing

Deep Learning — Units 5 & 6

Dr. Jon Krohn
jon@untapt.com

Slides available at jonkrohn.com/talks

August 4th, 2018

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

Outline

1. Review Take-Home Exercise
2. The Power and Elegance of Deep Learning for NLP
   - Introduction to DL for NLP
   - Computational Representations of NL
   - NLP Applications
3. Word Vectors
   - Vector-Space Embedding
   - word2vec
   - Creating Word Vectors with word2vec
4. Modeling Natural Language Data
   - Best Practices for Preprocessing NLP Data
   - The Area Under the ROC Curve
   - Sentiment Classification
5. Recurrent Neural Networks
   - Simple RNNs
   - LSTMs
6. Parallel Network Architectures

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

Outline

1. Review Take-Home Exercise
2. The Power and Elegance of Deep Learning for NLP
   - Introduction to DL for NLP
   - Computational Representations of NL
   - NLP Applications
3. Word Vectors
   - Vector-Space Embedding
   - word2vec
   - Creating Word Vectors with word2vec
4. Modeling Natural Language Data
   - Best Practices for Preprocessing NLP Data
   - The Area Under the ROC Curve
   - Sentiment Classification
5. Recurrent Neural Networks
   - Simple RNNs
   - LSTMs
6. Parallel Network Architectures

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

Outline

1. Review Take-Home Exercise
2. The Power and Elegance of Deep Learning for NLP
   - Introduction to DL for NLP
   - Computational Representations of NL
   - NLP Applications
3. Word Vectors
   - Vector-Space Embedding
   - word2vec
   - Creating Word Vectors with word2vec
4. Modeling Natural Language Data
   - Best Practices for Preprocessing NLP Data
   - The Area Under the ROC Curve
   - Sentiment Classification
5. Recurrent Neural Networks
   - Simple RNNs
   - LSTMs
6. Parallel Network Architectures

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

Outline

1. Review Take-Home Exercise
2. The Power and Elegance of Deep Learning for NLP
   - Introduction to DL for NLP
   - Computational Representations of NL
   - NLP Applications
3. Word Vectors
   - Vector-Space Embedding
   - word2vec
   - Creating Word Vectors with word2vec
4. Modeling Natural Language Data
   - Best Practices for Preprocessing NLP Data
   - The Area Under the ROC Curve
   - Sentiment Classification
5. Recurrent Neural Networks
   - Simple RNNs
   - LSTMs
6. Parallel Network Architectures

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

Outline

1. Review Take-Home Exercise
2. The Power and Elegance of Deep Learning for NLP
   - Introduction to DL for NLP
   - Computational Representations of NL
   - NLP Applications
3. Word Vectors
   - Vector-Space Embedding
   - word2vec
   - Creating Word Vectors with word2vec
4. Modeling Natural Language Data
   - Best Practices for Preprocessing NLP Data
   - The Area Under the ROC Curve
   - Sentiment Classification
5. Recurrent Neural Networks
   - Simple RNNs
   - LSTMs
6. Parallel Network Architectures

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

Outline

1. Review Take-Home Exercise
2. The Power and Elegance of Deep Learning for NLP
   - Introduction to DL for NLP
   - Computational Representations of NL
   - NLP Applications
3. Word Vectors
   - Vector-Space Embedding
   - word2vec
   - Creating Word Vectors with word2vec
4. Modeling Natural Language Data
   - Best Practices for Preprocessing NLP Data
   - The Area Under the ROC Curve
   - Sentiment Classification
5. Recurrent Neural Networks
   - Simple RNNs
   - LSTMs
6. Parallel Network Architectures

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

Outline

1. Review Take-Home Exercise
2. The Power and Elegance of Deep Learning for NLP
   - Introduction to DL for NLP
   - Computational Representations of NL
   - NLP Applications
3. Word Vectors
   - Vector-Space Embedding
   - word2vec
   - Creating Word Vectors with word2vec
4. Modeling Natural Language Data
   - Best Practices for Preprocessing NLP Data
   - The Area Under the ROC Curve
   - Sentiment Classification
5. Recurrent Neural Networks
   - Simple RNNs
   - LSTMs
6. Parallel Network Architectures

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

Take-Home Exercise: VGGNet

Talk through the purpose of every line in the [VGGNet notebook], including all of the following terms:

- ReLU
- cross-entropy
- epoch
- parameters
- hyperparams
- SGD
- learning rate
- batch size
- Adam
- dropout
- batchnorm
- input layer
- dense/FC layer
- convolutional
- max-pooling
- flatten
- softmax layer

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

Take-Home Exercise: VGGNet

Talk through the purpose of every line in the [VGGNet notebook], including all of the following terms:

- ReLU
- cross-entropy
- epoch
- parameters
- hyperparams
- SGD
- learning rate
- batch size
- Adam
- dropout
- batchnorm
- input layer
- dense/FC layer
- convolutional
- max-pooling
- flatten
- softmax layer

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

Take-Home Exercise: VGGNet

Talk through the purpose of every line in the [VGGNet notebook], including all of the following terms:

- ReLU
- cross-entropy
- epoch
- parameters
- hyperparams
- SGD
- learning rate
- batch size
- Adam
- dropout
- batchnorm
- input layer
- dense/FC layer
- convolutional
- max-pooling
- flatten
- softmax layer

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

Take-Home Exercise: VGGNet

Talk through the purpose of every line in the [VGGNet notebook], including all of the following terms:

- ReLU
- cross-entropy
- epoch
- parameters
- hyperparams
- SGD
- learning rate
- batch size
- Adam
- dropout
- batchnorm
- input layer
- dense/FC layer
- convolutional
- max-pooling
- flatten
- softmax layer

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

Take-Home Exercise: VGGNet

Talk through the purpose of every line in the [VGGNet notebook], including all of the following terms:

- ReLU
- cross-entropy
- epoch
- parameters
- hyperparams
- SGD
- learning rate
- batch size
- Adam
- dropout
- batchnorm
- input layer
- dense/FC layer
- convolutional
- max-pooling
- flatten
- softmax layer

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

Take-Home Exercise: VGGNet

Talk through the purpose of every line in the [VGGNet notebook], including all of the following terms:

- ReLU
- cross-entropy
- epoch
- parameters
- hyperparams
- SGD
- learning rate
- batch size
- Adam
- dropout
- batchnorm
- input layer
- dense/FC layer
- convolutional
- max-pooling
- flatten
- softmax layer

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

Take-Home Exercise: VGGNet

Talk through the purpose of every line in the [VGGNet notebook], including all of the following terms:

- ReLU
- cross-entropy
- epoch
- parameters
- hyperparams

Take-Home Exercise: VGGNet

Talk through the purpose of every line in the [VGGNet notebook], including all of the following terms:

- ReLU
- cross-entropy
- epoch
- parameters
- hyperparams
- SGD
- learning rate
- batch size
- Adam
- dropout
- batchnorm
- input layer
- dense/FC layer
- convolutional
- max-pooling
- flatten
- softmax layer

Take-Home Exercise: VGGNet

Talk through the purpose of every line in the [VGGNet notebook], including all of the following terms:

- ReLU
- cross-entropy
- epoch
- parameters
- hyperparams
- SGD
- learning rate
- batch size
- Adam
- dropout
- batchnorm
- input layer
- dense/FC layer
- convolutional
- max-pooling
- flatten
- softmax layer

Take-Home Exercise: VGGNet

Talk through the purpose of every line in the [VGGNet notebook], including all of the following terms:

- ReLU
- cross-entropy
- epoch
- parameters
- hyperparams
- SGD
- learning rate
- batch size
- Adam
- dropout
- batchnorm
- input layer
- dense/FC layer
- convolutional
- max-pooling
- flatten
- softmax layer

Take-Home Exercise: VGGNet

Talk through the purpose of every line in the [VGGNet notebook], including all of the following terms:

- ReLU
- cross-entropy
- epoch
- parameters
- hyperparams
- SGD
- learning rate
- batch size
- Adam
- dropout
- batchnorm
- input layer
- dense/FC layer
- convolutional
- max-pooling
- flatten
- softmax layer

Take-Home Exercise: VGGNet

Talk through the purpose of every line in the [VGGNet notebook], including all of the following terms:

- ReLU
- cross-entropy
- epoch
- parameters
- hyperparams
- SGD
- learning rate
- batch size
- Adam
- dropout
- batchnorm
- input layer
- dense/FC layer
- convolutional
- max-pooling
- flatten
- softmax layer

Take-Home Exercise: VGGNet

Talk through the purpose of every line in the [VGGNet notebook], including all of the following terms:

- ReLU
- cross-entropy
- epoch
- parameters
- hyperparams
- SGD
- learning rate
- batch size
- Adam
- dropout
- batchnorm
- input layer
- dense/FC layer
- convolutional
- max-pooling
- flatten
- softmax layer

Take-Home Exercise: VGGNet

Talk through the purpose of every line in the [VGGNet notebook], including all of the following terms:

- ReLU
- cross-entropy
- epoch
- parameters
- hyperparams
- SGD
- learning rate
- batch size
- Adam
- dropout
- batchnorm
- input layer
- dense/FC layer
- convolutional
- max-pooling
- flatten
- softmax layer

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

Take-Home Exercise: VGGNet

Talk through the purpose of every line in the [VGGNet notebook], including all of the following terms:

- ReLU
- cross-entropy
- epoch
- parameters
- hyperparams
- SGD
- learning rate
- batch size
- Adam
- dropout
- batchnorm
- input layer
- dense/FC layer
- convolutional
- max-pooling
- flatten
- softmax layer

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

Take-Home Exercise: VGGNet

Talk through the purpose of every line in the [VGGNet notebook], including all of the following terms:

- ReLU
- cross-entropy
- epoch
- parameters
- hyperparams
- SGD
- learning rate
- batch size
- Adam
- dropout
- batchnorm
- input layer
- dense/FC layer
- convolutional
- max-pooling
- flatten
- softmax layer

Take-Home Exercise: VGGNet

Talk through the purpose of every line in the [VGGNet notebook], including all of the following terms:

- ReLU
- cross-entropy
- epoch
- parameters
- hyperparams
- SGD
- learning rate
- batch size
- Adam
- dropout
- batchnorm
- input layer
- dense/FC layer
- convolutional
- max-pooling
- flatten
- softmax layer

Take-Home Exercise: VGGNet

Talk through the purpose of every line in the [VGGNet notebook], including all of the following terms:

- ReLU
- cross-entropy
- epoch
- parameters
- hyperparams
- SGD
- learning rate
- batch size
- Adam
- dropout
- batchnorm
- input layer
- dense/FC layer
- convolutional
- max-pooling
- flatten
- softmax layer

Take-Home Exercise: VGGNet

Talk through the purpose of every line in the [VGGNet notebook], including all of the following terms:

- ReLU
- cross-entropy
- epoch
- parameters
- hyperparams
- SGD
- learning rate
- batch size
- Adam
- dropout
- batchnorm
- input layer
- dense/FC layer
- convolutional
- max-pooling
- flatten
- softmax layer

Take-Home Exercise: VGGNet

Talk through the purpose of every line in the [VGGNet notebook], including all of the following terms:

- ReLU
- cross-entropy
- epoch
- parameters
- hyperparams
- SGD
- learning rate
- batch size
- Adam
- dropout
- batchnorm
- input layer
- dense/FC layer
- convolutional
- max-pooling
- flatten
- softmax layer

Take-Home Exercise: VGGNet

Talk through the purpose of every line in the [VGGNet notebook], including all of the following terms:

- ReLU
- cross-entropy
- epoch
- parameters
- hyperparams
- SGD
- learning rate
- batch size
- Adam
- dropout
- batchnorm
- input layer
- dense/FC layer
- convolutional
- max-pooling
- flatten
- softmax layer

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

Outline

1. Review Take-Home Exercise
2. The Power and Elegance of Deep Learning for NLP
   - Introduction to DL for NLP
   - Computational Representations of NL
   - NLP Applications
3. Word Vectors
   - Vector-Space Embedding
   - word2vec
   - Creating Word Vectors with word2vec
4. Modeling Natural Language Data
   - Best Practices for Preprocessing NLP Data
   - The Area Under the ROC Curve
   - Sentiment Classification
5. Recurrent Neural Networks
   - Simple RNNs
   - LSTMs
6. Parallel Network Architectures

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

Outline

1. Review Take-Home Exercise
2. The Power and Elegance of Deep Learning for NLP
   - Introduction to DL for NLP
   - Computational Representations of NL
   - NLP Applications
3. Word Vectors
   - Vector-Space Embedding
   - word2vec
   - Creating Word Vectors with word2vec
4. Modeling Natural Language Data
   - Best Practices for Preprocessing NLP Data
   - The Area Under the ROC Curve
   - Sentiment Classification
5. Recurrent Neural Networks
   - Simple RNNs
   - LSTMs
6. Parallel Network Architectures

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

Two Core Concepts

1. Deep Learning
2. Natural Language Processing (NLP)

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

Two Core Concepts

1. Deep Learning
2. Natural Language Processing (NLP)

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

TML vs Deep Learning

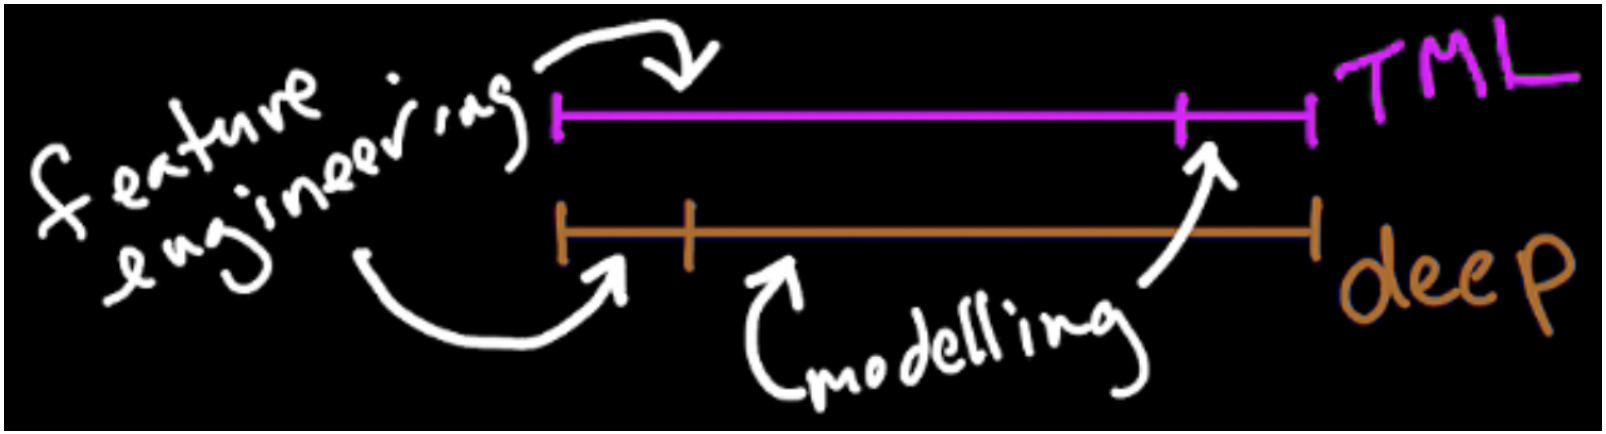

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

Two Core Concepts

1. Deep Learning
2. Natural Language Processing (NLP)

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

Two Core Concepts

1. Deep Learning
2. Natural Language Processing (NLP)

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

Natural Language Processing

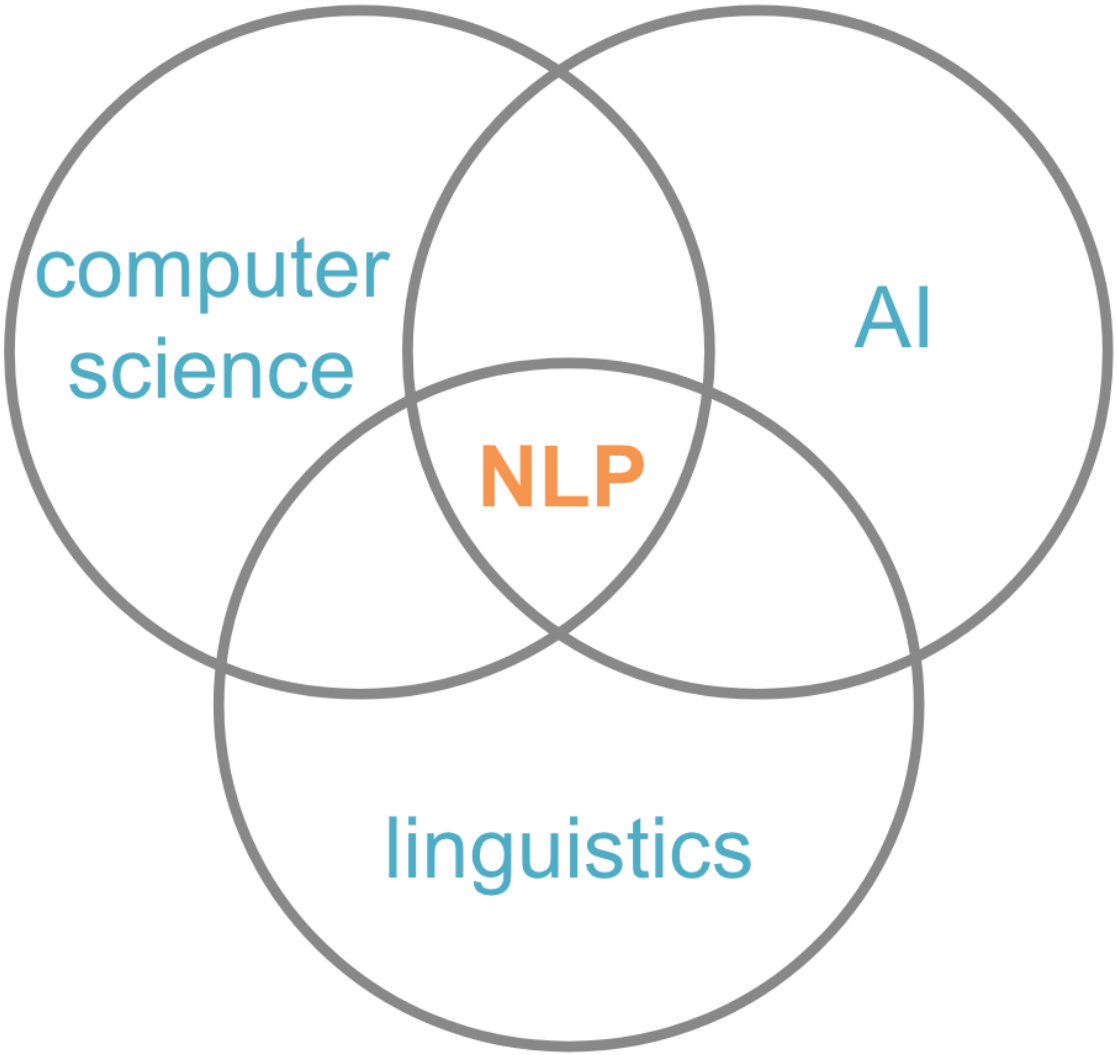

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

Examples of NLP in Industry

- speech recognition (Echo, Siri, Cortana)
- search (typed into omnibox, spoken)
- classifying documents
- language translation
- chatbots

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

Examples of NLP in Industry

- speech recognition (Echo, Siri, Cortana)
- search (typed into omnibox, spoken)
- classifying documents
- language translation
- chatbots

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

Examples of NLP in Industry

- speech recognition (Echo, Siri, Cortana)
- search (typed into omnibox, spoken)
- classifying documents
- language translation
- chatbots

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

Examples of NLP in Industry

- speech recognition (Echo, Siri, Cortana)
- search (typed into omnibox, spoken)
- classifying documents
- language translation
- chatbots

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

Examples of NLP in Industry

- speech recognition (Echo, Siri, Cortana)
- search (typed into omnibox, spoken)
- classifying documents
- language translation
- chatbots

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

Outline

1. Review Take-Home Exercise
2. The Power and Elegance of Deep Learning for NLP
   - Introduction to DL for NLP
   - Computational Representations of NL
   - NLP Applications
3. Word Vectors
   - Vector-Space Embedding
   - word2vec
   - Creating Word Vectors with word2vec
4. Modeling Natural Language Data
   - Best Practices for Preprocessing NLP Data
   - The Area Under the ROC Curve
   - Sentiment Classification
5. Recurrent Neural Networks
   - Simple RNNs
   - LSTMs
6. Parallel Network Architectures

One-Hot Word Representations

word	The	cat	sat	on	the	mat.
the	1	0	0	0	1	0
cat	0	1	0	0	0	0
on	0	0	0	1	0	0
⋮						
n_{unique_words}						

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

Outline

1. Review Take-Home Exercise
2. The Power and Elegance of Deep Learning for NLP
   - Introduction to DL for NLP
   - Computational Representations of NL
   - NLP Applications
3. Word Vectors
   - Vector-Space Embedding
   - word2vec
   - Creating Word Vectors with word2vec
4. Modeling Natural Language Data
   - Best Practices for Preprocessing NLP Data
   - The Area Under the ROC Curve
   - Sentiment Classification
5. Recurrent Neural Networks
   - Simple RNNs
   - LSTMs
6. Parallel Network Architectures

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

Easy

- spell checking
- synonym suggestions
- keyword search

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

Easy

- spell checking
- synonym suggestions
- keyword search

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

Easy

- spell checking
- synonym suggestions
- keyword search

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

Intermediate

- reading level
- extracting information
- predicting next words
- classification
- sequence generation
- time-series analysis

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

Intermediate

- reading level
- extracting information
- predicting next words
- classification
- sequence generation
- time-series analysis

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

Intermediate

- reading level
- extracting information
- predicting next words
- classification
- sequence generation
- time-series analysis

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

Intermediate

- reading level
- extracting information
- predicting next words
- classification
- sequence generation
- time-series analysis

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

Intermediate

- reading level
- extracting information
- predicting next words
- classification
- sequence generation
- time-series analysis

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

Intermediate

- reading level
- extracting information
- predicting next words
- classification
- sequence generation
- time-series analysis

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

Complex

- machine translation
- question-answering
- chatbots

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

Complex

- machine translation
- question-answering
- chatbots

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

Complex

- machine translation
- question-answering
- chatbots

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

Outline

1. Review Take-Home Exercise
2. The Power and Elegance of Deep Learning for NLP
   - Introduction to DL for NLP
   - Computational Representations of NL
   - NLP Applications
3. Word Vectors
   - Vector-Space Embedding
   - word2vec
   - Creating Word Vectors with word2vec
4. Modeling Natural Language Data
   - Best Practices for Preprocessing NLP Data
   - The Area Under the ROC Curve
   - Sentiment Classification
5. Recurrent Neural Networks
   - Simple RNNs
   - LSTMs
6. Parallel Network Architectures

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

Outline

1. Review Take-Home Exercise
2. The Power and Elegance of Deep Learning for NLP
   - Introduction to DL for NLP
   - Computational Representations of NL
   - NLP Applications
3. Word Vectors
   - Vector-Space Embedding
   - word2vec
   - Creating Word Vectors with word2vec
4. Modeling Natural Language Data
   - Best Practices for Preprocessing NLP Data
   - The Area Under the ROC Curve
   - Sentiment Classification
5. Recurrent Neural Networks
   - Simple RNNs
   - LSTMs
6. Parallel Network Architectures

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

JR Firth (1957)

"You shall know a word by the company it keeps"

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

Vector Representations of Words

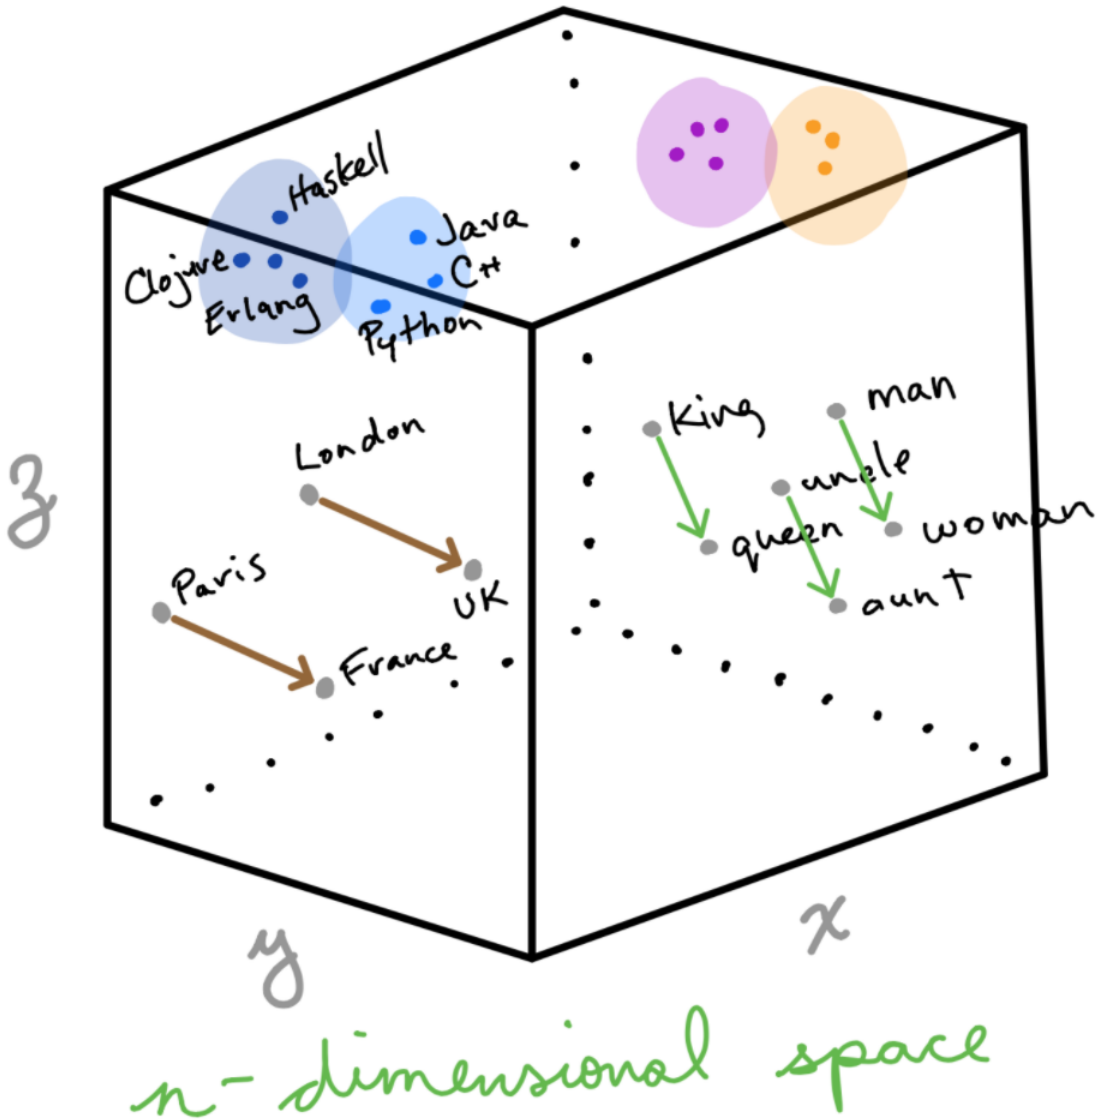

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

Word Vector Arithmetic

$$v_{king} - v_{man} + v_{woman} = v_{?}$$

$$v_{jeff\_bezos} - v_{amazon} + v_{facebook} = v_{?}$$

$$v_{windows} - v_{microsoft} + v_{google} = v_{?}$$

$$v_{cu} - v_{copper} + v_{gold} = v_{?}$$

Units 5 and 6 — NLP

Review

DL for NLP
Intro
Representations
NLP Applications

Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors

Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification

RNNs
Simple RNNs
LSTMs

Parallel Nets

[word2viz demo]

Word Representations

One-Hot	Vector-Based
lack nuance	extremely **nuanced**
handle new words poorly	seamlessly incorporate **new words**
subjective	**driven by** natural language **data**
laborious, manual taxonomies	fully-**automatic**
word similarity ignored	**word similarity** = closeness in space
unwieldy with large vocabulary	accommodate **large vocabularies**

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

Outline

1. Review Take-Home Exercise
2. The Power and Elegance of Deep Learning for NLP
   - Introduction to DL for NLP
   - Computational Representations of NL
   - NLP Applications
3. Word Vectors
   - Vector-Space Embedding
   - word2vec
   - Creating Word Vectors with word2vec
4. Modeling Natural Language Data
   - Best Practices for Preprocessing NLP Data
   - The Area Under the ROC Curve
   - Sentiment Classification
5. Recurrent Neural Networks
   - Simple RNNs
   - LSTMs
6. Parallel Network Architectures

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

JR Firth (1957)

"You shall know a word by the company it keeps"

Word Representations

	predicts	**relative strengths**
Skip-Gram (SG)	context given target	• small data set • rare words
CBOW	target given context	• many times faster • slightly better for frequent words

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

Evaluating Word Vectors

1. intrinsic
2. extrinsic

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

Evaluating Word Vectors

1. intrinsic
2. extrinsic

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

word2vec Hyperparameters

1. *n* dimensions
2. window size (SG ~10, CBOW ~5)
3. *n* iterations
4. data set size

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

word2vec Hyperparameters

1. *n* dimensions
2. window size (SG ~10, CBOW ~5)
3. *n* iterations
4. data set size

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

word2vec Hyperparameters

1. *n* dimensions
2. window size (SG ~10, CBOW ~5)
3. *n* iterations
4. data set size

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

word2vec Hyperparameters

1. *n* dimensions
2. window size (SG ~10, CBOW ~5)
3. *n* iterations
4. data set size

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

Outline

1. Review Take-Home Exercise
2. The Power and Elegance of Deep Learning for NLP
   - Introduction to DL for NLP
   - Computational Representations of NL
   - NLP Applications
3. Word Vectors
   - Vector-Space Embedding
   - word2vec
   - Creating Word Vectors with word2vec
4. Modeling Natural Language Data
   - Best Practices for Preprocessing NLP Data
   - The Area Under the ROC Curve
   - Sentiment Classification
5. Recurrent Neural Networks
   - Simple RNNs
   - LSTMs
6. Parallel Network Architectures

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

[*creating word vectors* notebook]

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

Outline

1. Review Take-Home Exercise
2. The Power and Elegance of Deep Learning for NLP
   - Introduction to DL for NLP
   - Computational Representations of NL
   - NLP Applications
3. Word Vectors
   - Vector-Space Embedding
   - word2vec
   - Creating Word Vectors with word2vec
4. Modeling Natural Language Data
   - Best Practices for Preprocessing NLP Data
   - The Area Under the ROC Curve
   - Sentiment Classification
5. Recurrent Neural Networks
   - Simple RNNs
   - LSTMs
6. Parallel Network Architectures

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

Outline

1. Review Take-Home Exercise
2. The Power and Elegance of Deep Learning for NLP
   - Introduction to DL for NLP
   - Computational Representations of NL
   - NLP Applications
3. Word Vectors
   - Vector-Space Embedding
   - word2vec
   - Creating Word Vectors with word2vec
4. Modeling Natural Language Data
   - Best Practices for Preprocessing NLP Data
   - The Area Under the ROC Curve
   - Sentiment Classification
5. Recurrent Neural Networks
   - Simple RNNs
   - LSTMs
6. Parallel Network Architectures

Best Practices for Preprocessing NLP Data

[*NL preprocessing best practices* notebook]

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

Outline

1. Review Take-Home Exercise
2. The Power and Elegance of Deep Learning for NLP
   - Introduction to DL for NLP
   - Computational Representations of NL
   - NLP Applications
3. Word Vectors
   - Vector-Space Embedding
   - word2vec
   - Creating Word Vectors with word2vec
4. Modeling Natural Language Data
   - Best Practices for Preprocessing NLP Data
   - The Area Under the ROC Curve
   - Sentiment Classification
5. Recurrent Neural Networks
   - Simple RNNs
   - LSTMs
6. Parallel Network Architectures

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

The Area Under the ROC Curve

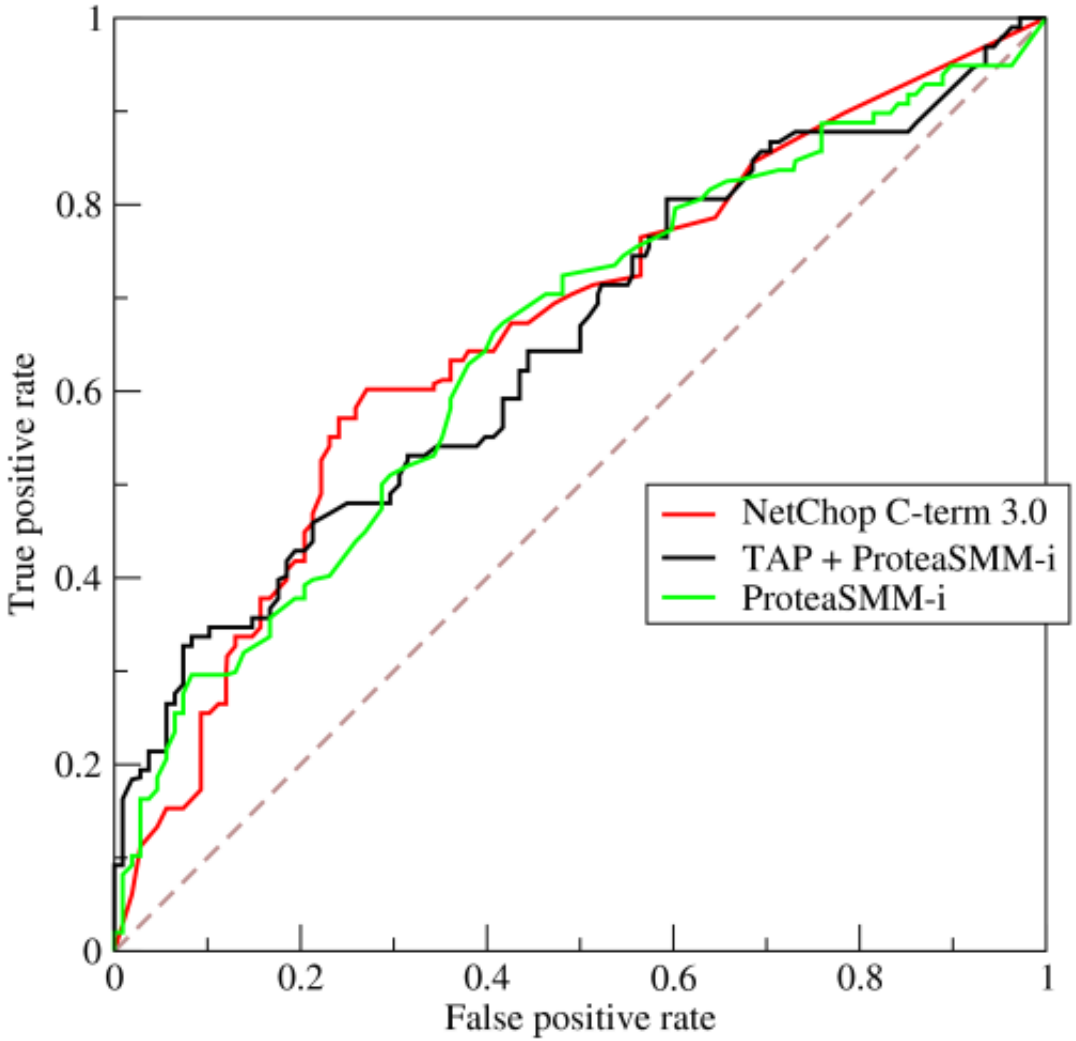

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

Outline

1 Review Take-Home Exercise
2 The Power and Elegance of Deep Learning for NLP
- Introduction to DL for NLP
- Computational Representations of NL
- NLP Applications

3 Word Vectors
- Vector-Space Embedding
- word2vec
- Creating Word Vectors with word2vec

4 Modeling Natural Language Data
- Best Practices for Preprocessing NLP Data
- The Area Under the ROC Curve
- Sentiment Classification

5 Recurrent Neural Networks
- Simple RNNs
- LSTMs

6 Parallel Network Architectures

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

Dense Net Classification

[*dense sentiment classifier* notebook]

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

ConvNet Classification

[*convolutional sentiment classifier* notebook]

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

Outline

1. Review Take-Home Exercise
2. The Power and Elegance of Deep Learning for NLP
   - Introduction to DL for NLP
   - Computational Representations of NL
   - NLP Applications
3. Word Vectors
   - Vector-Space Embedding
   - word2vec
   - Creating Word Vectors with word2vec
4. Modeling Natural Language Data
   - Best Practices for Preprocessing NLP Data
   - The Area Under the ROC Curve
   - Sentiment Classification
5. Recurrent Neural Networks
   - Simple RNNs
   - LSTMs
6. Parallel Network Architectures

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

Outline

1. Review Take-Home Exercise
2. The Power and Elegance of Deep Learning for NLP
   - Introduction to DL for NLP
   - Computational Representations of NL
   - NLP Applications
3. Word Vectors
   - Vector-Space Embedding
   - word2vec
   - Creating Word Vectors with word2vec
4. Modeling Natural Language Data
   - Best Practices for Preprocessing NLP Data
   - The Area Under the ROC Curve
   - Sentiment Classification
5. Recurrent Neural Networks
   - Simple RNNs
   - LSTMs
6. Parallel Network Architectures

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

RNN Theory

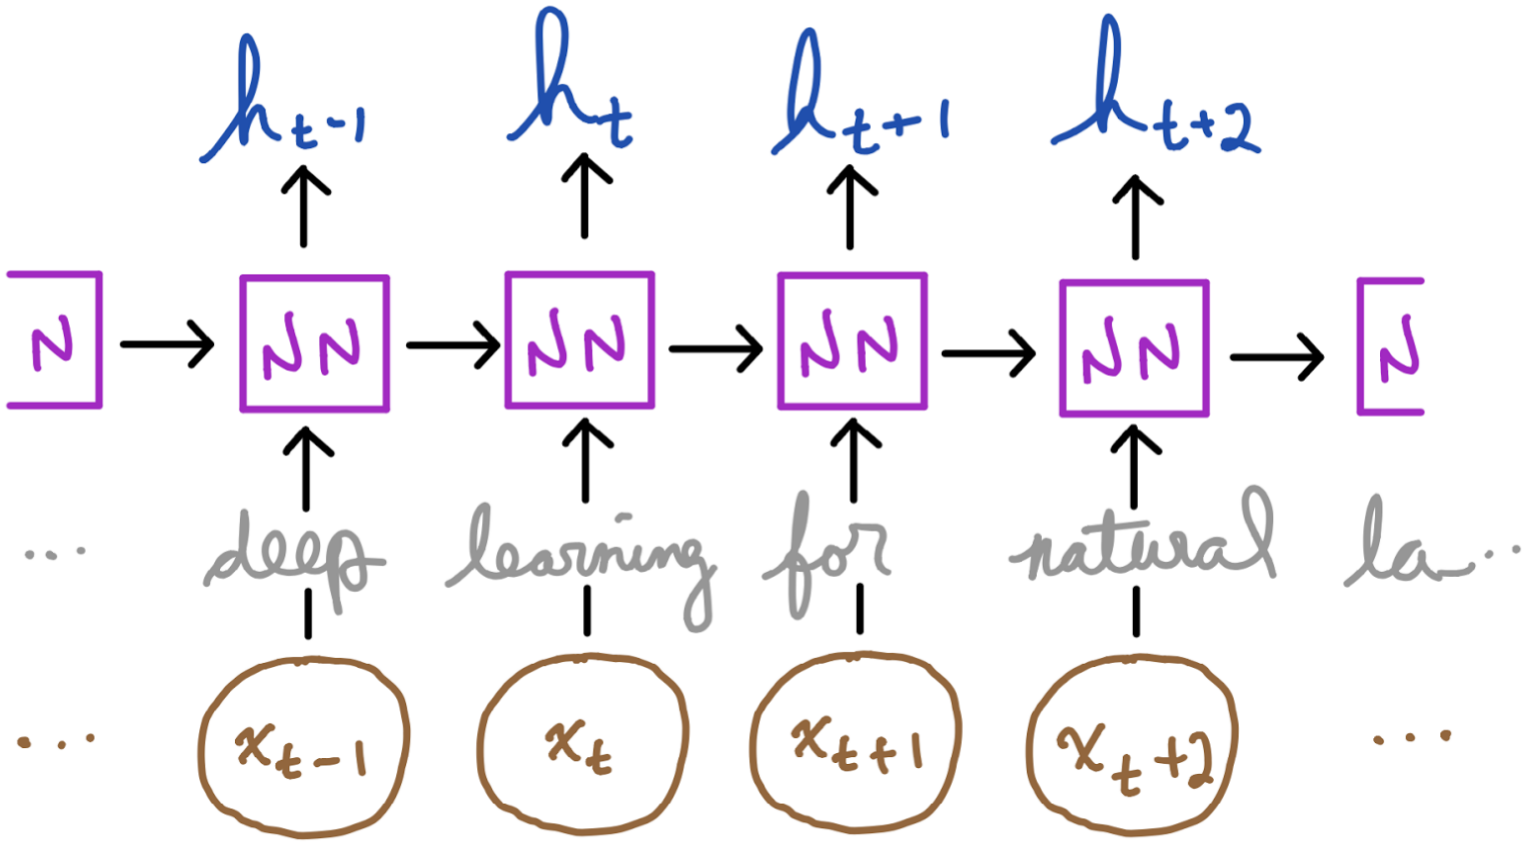

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

RNNs in Practice

[*rnn* notebook]

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

Outline

1. Review Take-Home Exercise
2. The Power and Elegance of Deep Learning for NLP
   - Introduction to DL for NLP
   - Computational Representations of NL
   - NLP Applications
3. Word Vectors
   - Vector-Space Embedding
   - word2vec
   - Creating Word Vectors with word2vec
4. Modeling Natural Language Data
   - Best Practices for Preprocessing NLP Data
   - The Area Under the ROC Curve
   - Sentiment Classification
5. Recurrent Neural Networks
   - Simple RNNs
   - LSTMs
6. Parallel Network Architectures

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

LSTM Theory

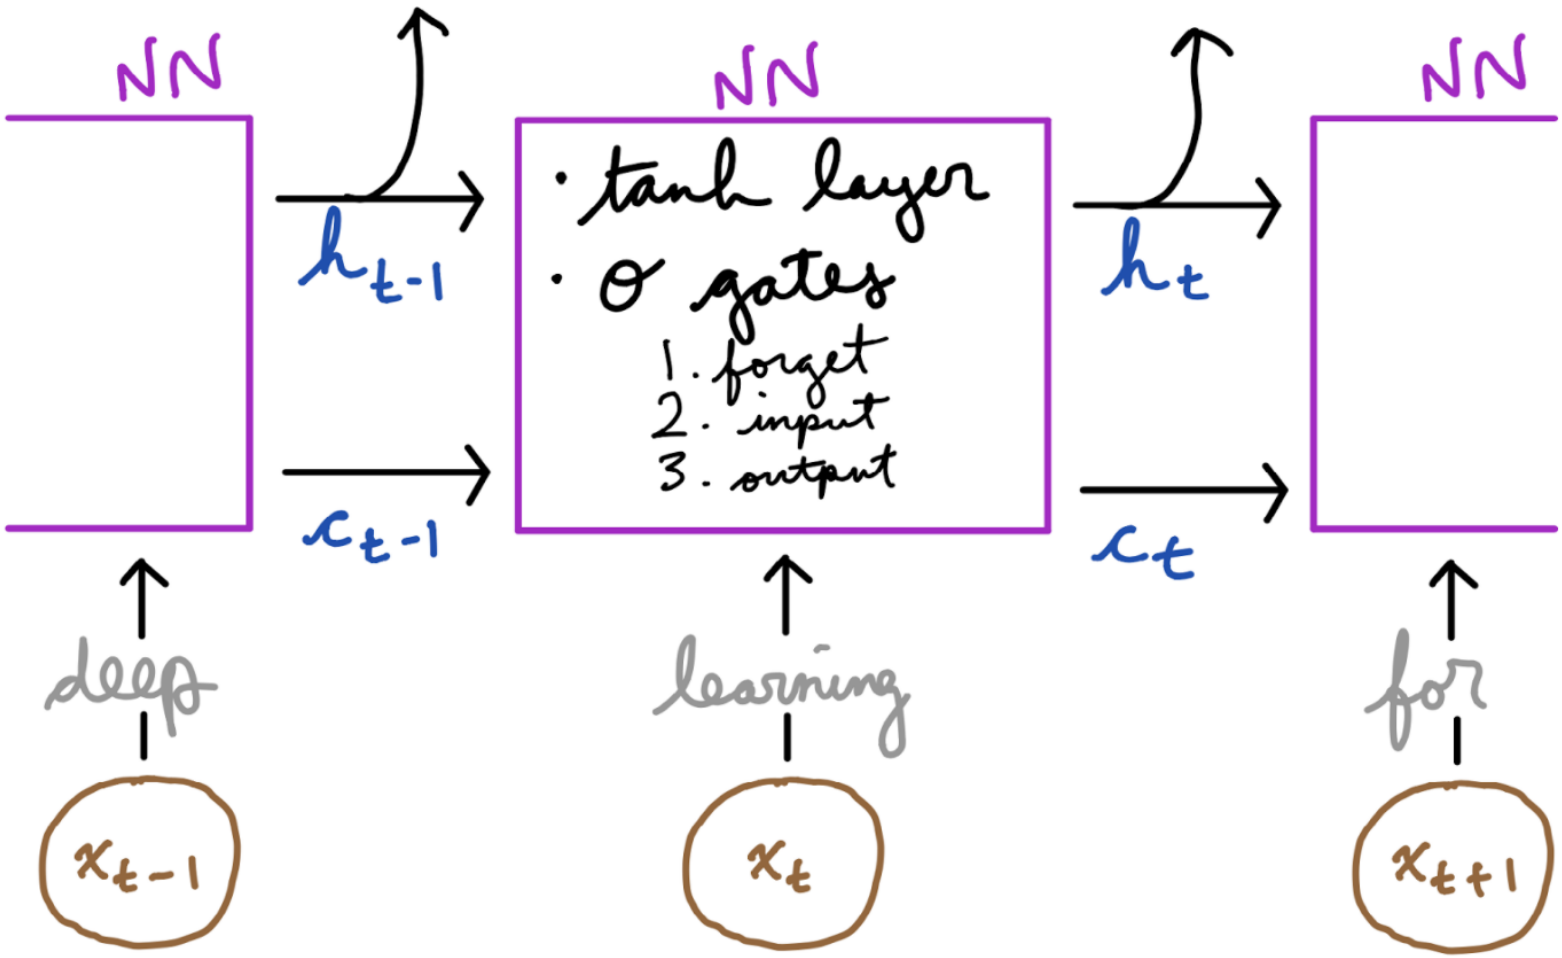

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

LSTMs in Practice

[*vanilla LSTM* and *GRU* notebooks]

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

Bi-Directional LSTMs

[*Bi-LSTM* notebook]

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

Stacked LSTMs

[*stacked LSTM* and *ye olde stackeroo* notebooks]

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

Outline

1. Review Take-Home Exercise
2. The Power and Elegance of Deep Learning for NLP
   - Introduction to DL for NLP
   - Computational Representations of NL
   - NLP Applications
3. Word Vectors
   - Vector-Space Embedding
   - word2vec
   - Creating Word Vectors with word2vec
4. Modeling Natural Language Data
   - Best Practices for Preprocessing NLP Data
   - The Area Under the ROC Curve
   - Sentiment Classification
5. Recurrent Neural Networks
   - Simple RNNs
   - LSTMs
6. Parallel Network Architectures

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

Parallel Network Architectures

[*multi-ConvNet* notebook]

Assessing

Your Deep Learning Project III

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

Assessing

Your Deep Learning Project III

1. split your data
   - training set (80% — for optimizing parameters)
   - validation set (10% — for hyperparameters)
   - test set (10% — don't touch yet!)
2. build and assess architecture
   - get above chance (simplifying problem, if necessary)
   - do existing performance benchmarks exist?
   - if not, use a simple architecture as benchmark
3. "teamwork makes the dream work" (?)

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

Assessing

Your Deep Learning Project III

1. split your data
   - training set (80% — for optimizing parameters)
   - validation set (10% — for hyperparameters)
   - test set (10% — don't touch yet!)
2. build and assess architecture
   - get above chance (simplifying problem, if necessary)
   - do existing performance benchmarks exist?
   - if not, use a simple architecture as benchmark
3. "teamwork makes the dream work" (?)

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

Assessing

Your Deep Learning Project III

1. split your data
   - training set (80% — for optimizing parameters)
   - validation set (10% — for hyperparameters)
   - test set (10% — don't touch yet!)
2. build and assess architecture
   - get above chance (simplifying problem, if necessary)
   - do existing performance benchmarks exist?
   - if not, use a simple architecture as benchmark
3. "teamwork makes the dream work" (?)

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

Assessing

Your Deep Learning Project III

1. split your data
   - training set (80% — for optimizing parameters)
   - validation set (10% — for hyperparameters)
   - test set (10% — don't touch yet!)
2. build and assess architecture
   - get above chance (simplifying problem, if necessary)
   - do existing performance benchmarks exist?
   - if not, use a simple architecture as benchmark
3. "teamwork makes the dream work" (?)

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

Assessing

Your Deep Learning Project III

1. split your data
   - training set (80% — for optimizing parameters)
   - validation set (10% — for hyperparameters)
   - test set (10% — don't touch yet!)
2. build and assess architecture
   - get above chance (simplifying problem, if necessary)
   - do existing performance benchmarks exist?
   - if not, use a simple architecture as benchmark
3. "teamwork makes the dream work" (?)

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

Assessing

Your Deep Learning Project III

1. split your data
   - training set (80% — for optimizing parameters)
   - validation set (10% — for hyperparameters)
   - test set (10% — don't touch yet!)
2. build and assess architecture
   - get above chance (simplifying problem, if necessary)
   - do existing performance benchmarks exist?
   - if not, use a simple architecture as benchmark
3. "teamwork makes the dream work" (?)

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

Assessing

Your Deep Learning Project III

1. split your data
   - training set (80% — for optimizing parameters)
   - validation set (10% — for hyperparameters)
   - test set (10% — don't touch yet!)
2. build and assess architecture
   - get above chance (simplifying problem, if necessary)
   - do existing performance benchmarks exist?
   - if not, use a simple architecture as benchmark
3. "teamwork makes the dream work" (?)

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

Assessing

Your Deep Learning Project III

1. split your data
   - training set (80% — for optimizing parameters)
   - validation set (10% — for hyperparameters)
   - test set (10% — don't touch yet!)
2. build and assess architecture
   - get above chance (simplifying problem, if necessary)
   - do existing performance benchmarks exist?
   - if not, use a simple architecture as benchmark
3. "teamwork makes the dream work" (?)

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

Assessing

Your Deep Learning Project III

1. split your data
   - training set (80% — for optimizing parameters)
   - validation set (10% — for hyperparameters)
   - test set (10% — don't touch yet!)
2. build and assess architecture
   - get above chance (simplifying problem, if necessary)
   - do existing performance benchmarks exist?
   - if not, use a simple architecture as benchmark
3. "teamwork makes the dream work" (?)

Review
DL for NLP
Intro
Representations
NLP Applications
Word Vectors
Vector-Space Embedding
word2vec
Creating Word Vectors
Modeling NL Data
Preprocessing
ROC Curve
Sentiment Classification
RNNs
Simple RNNs
LSTMs
Parallel Nets

Up Next: TensorFlow