


# 文本表達特徵工程 (Feature Engineering for Text Representation)

Time: 2020/05/29 (Fri) (9:10 -12:00)

Place: 國立臺北護理健康大學 (台北市明德路365號) G210

Host: 祝國忠 院長 (健康科技學院院長)



Min-Yuh Day

戴敏育

Associate Professor

副教授

Dept. of Information Management, Tamkang University

淡江大學 資訊管理學系

<http://mail.tku.edu.tw/myday/>

2020-05-29



# Topics

## 1. 自然語言處理核心技術與文字探勘

(Core Technologies of Natural Language Processing and Text Mining)

## 2. 人工智慧文本分析基礎與應用

(Artificial Intelligence for Text Analytics: Foundations and Applications)

## 3. 文本表達特徵工程

(Feature Engineering for Text Representation)

## 4. 語意分析和命名實體識別

(Semantic Analysis and Named Entity Recognition; NER)

## 5. 深度學習和通用句子嵌入模型

(Deep Learning and Universal Sentence-Embedding Models)

## 6. 問答系統與對話系統

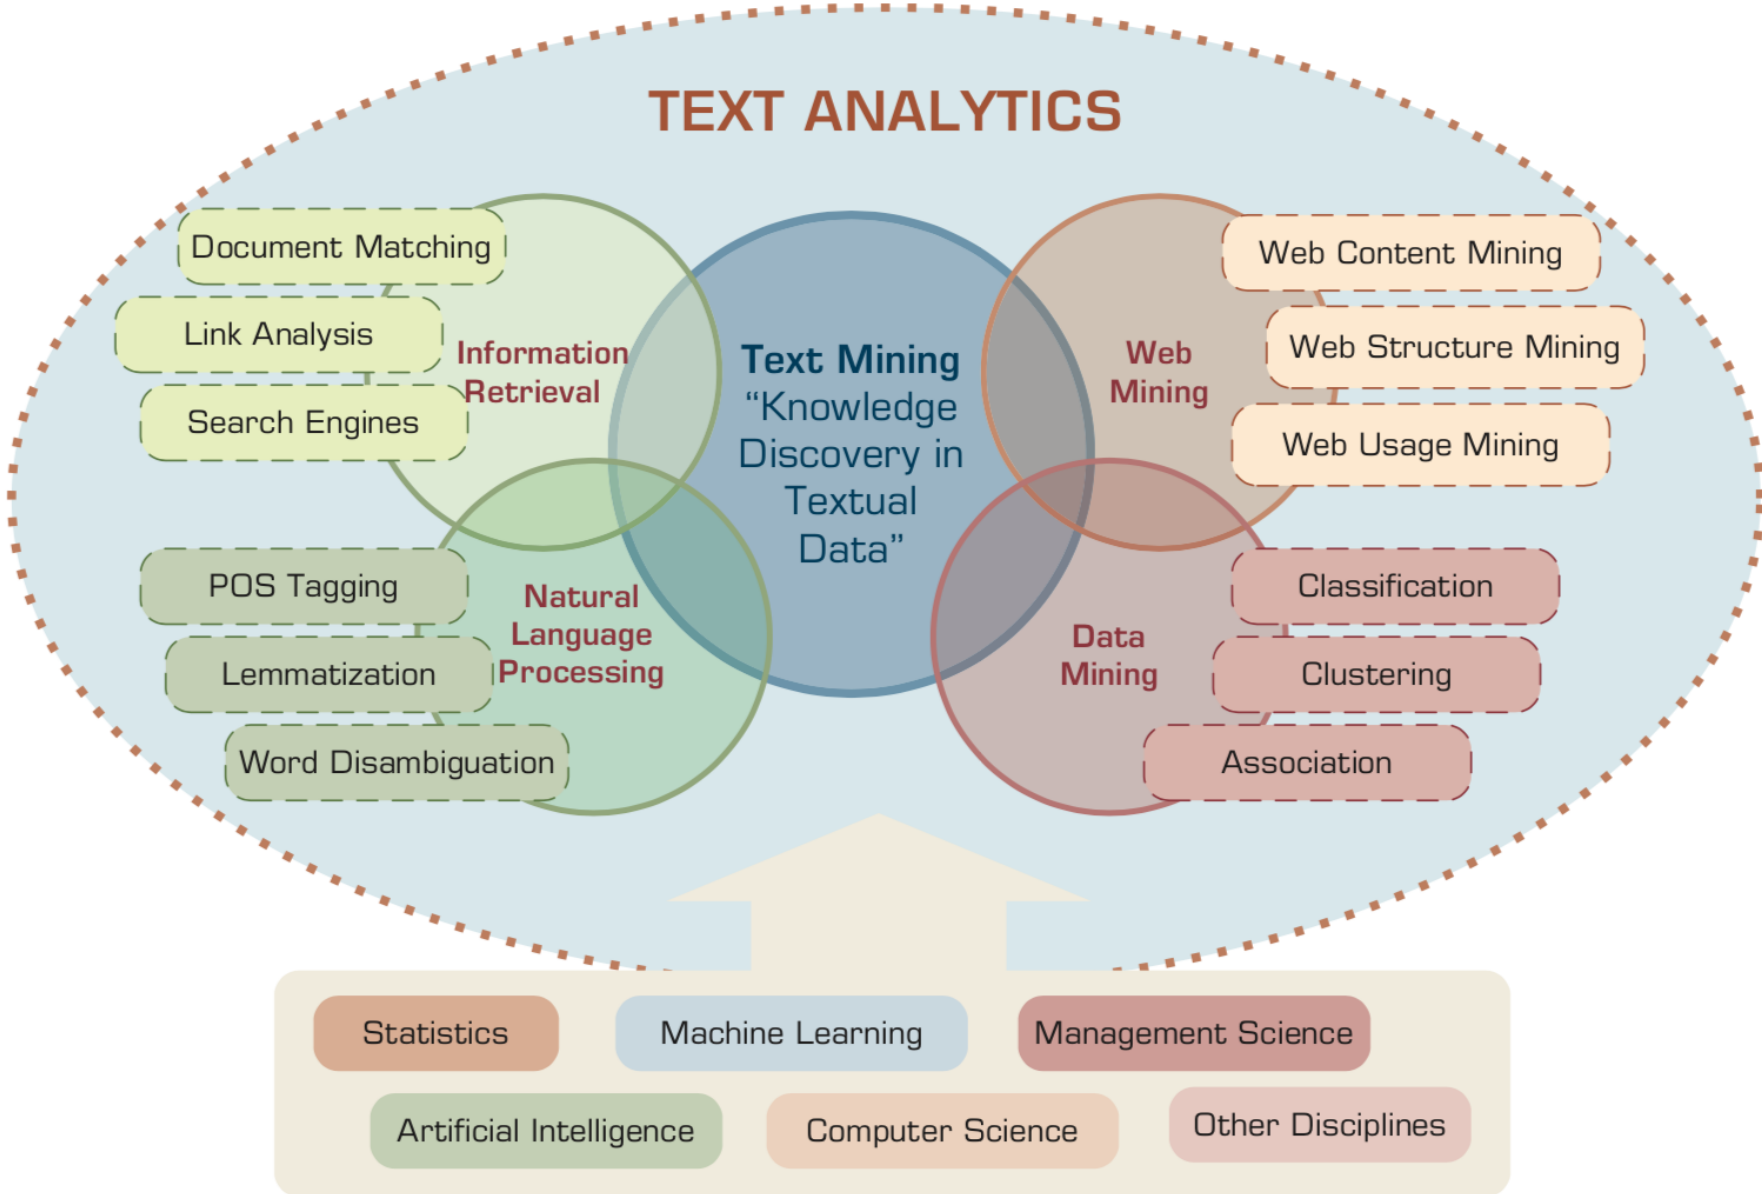
(Question Answering and Dialogue Systems)

# Outline

- Traditional Feature Engineering for Text Data
  - Bag of Words Model
  - Bag of N-Grams Model
  - TF-IDF Model
- Advanced Word Embeddings with Deep Learning
  - Word2Vec Model
  - Robust Word2Vec Models with Gensim
  - GloVe Model
  - FastText Model

# **Feature Engineering for Text Representation**


# Text Analytics and Text Mining




Source: Ramesh Sharda, Dursun Delen, and Efraim Turban (2017), Business Intelligence, Analytics, and Data Science: A Managerial Perspective, 4th Edition, Pearson

# NLP



## Classical NLP




## Deep Learning-based NLP



# Modern NLP Pipeline




# Modern NLP Pipeline






# Deep Learning NLP




# Overview of Text Vectorization Methods

Vectorization Method	Function	Good For	Considerations
Frequency	Counts term frequencies	Bayesian models	Most frequent words not always most informative
One-Hot Encoding	Binarizes term occurrence (0, 1)	Neural networks	All words equidistant, so normalization extra important
TF-IDF	Normalizes term frequencies across documents	General purpose	Moderately frequent terms may not be representative of document topics
Distributed Representations	Context-based, continuous term similarity encoding	Modeling more complex relationships	Performance intensive; difficult to scale without additional tools (e.g., Tensorflow)

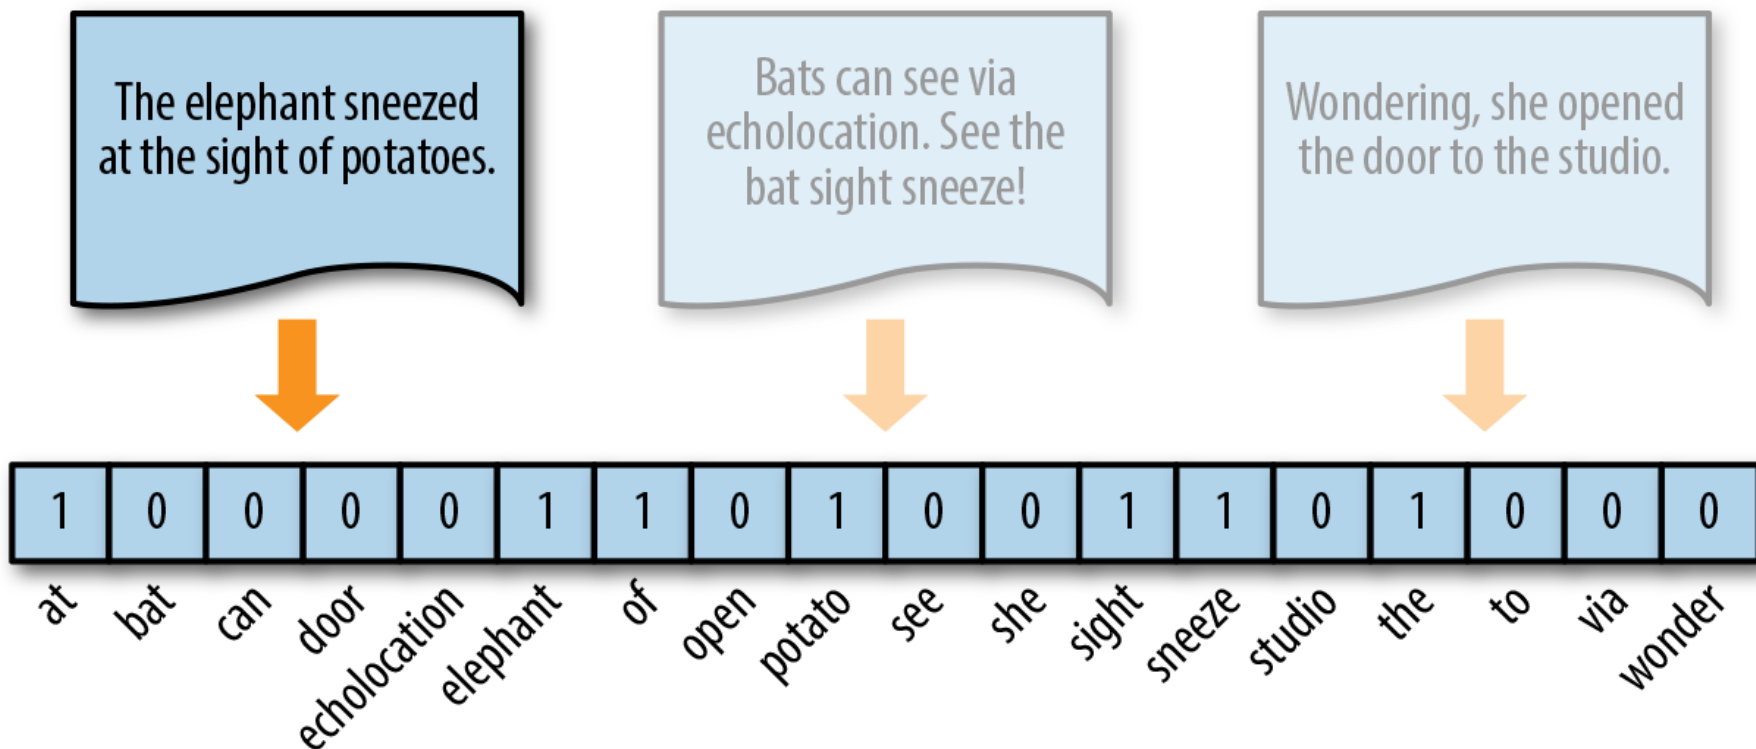
# Encoding Documents as Vectors



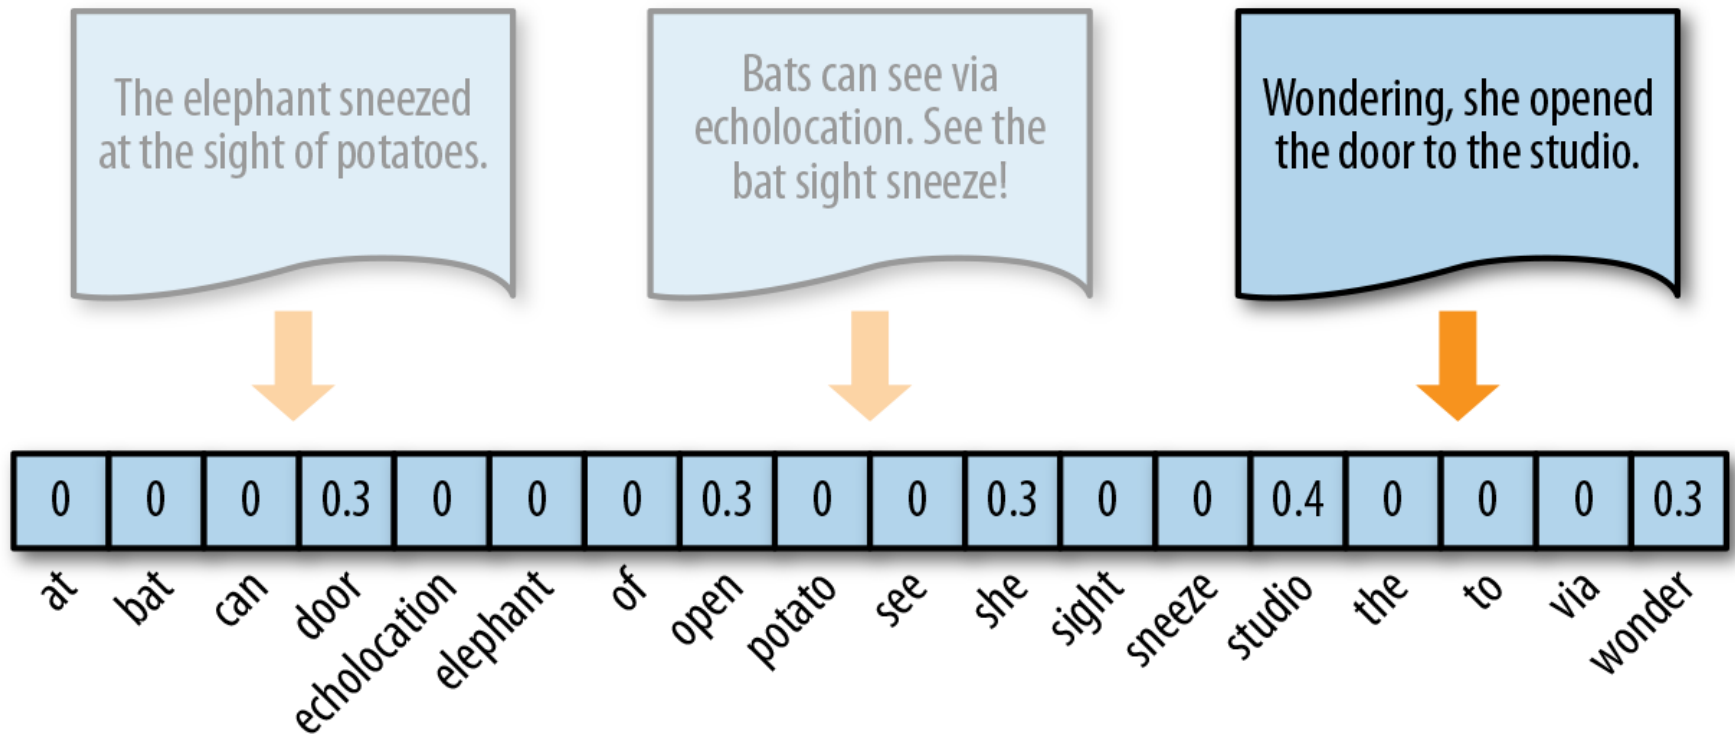
# Token Frequency as Vector Encoding



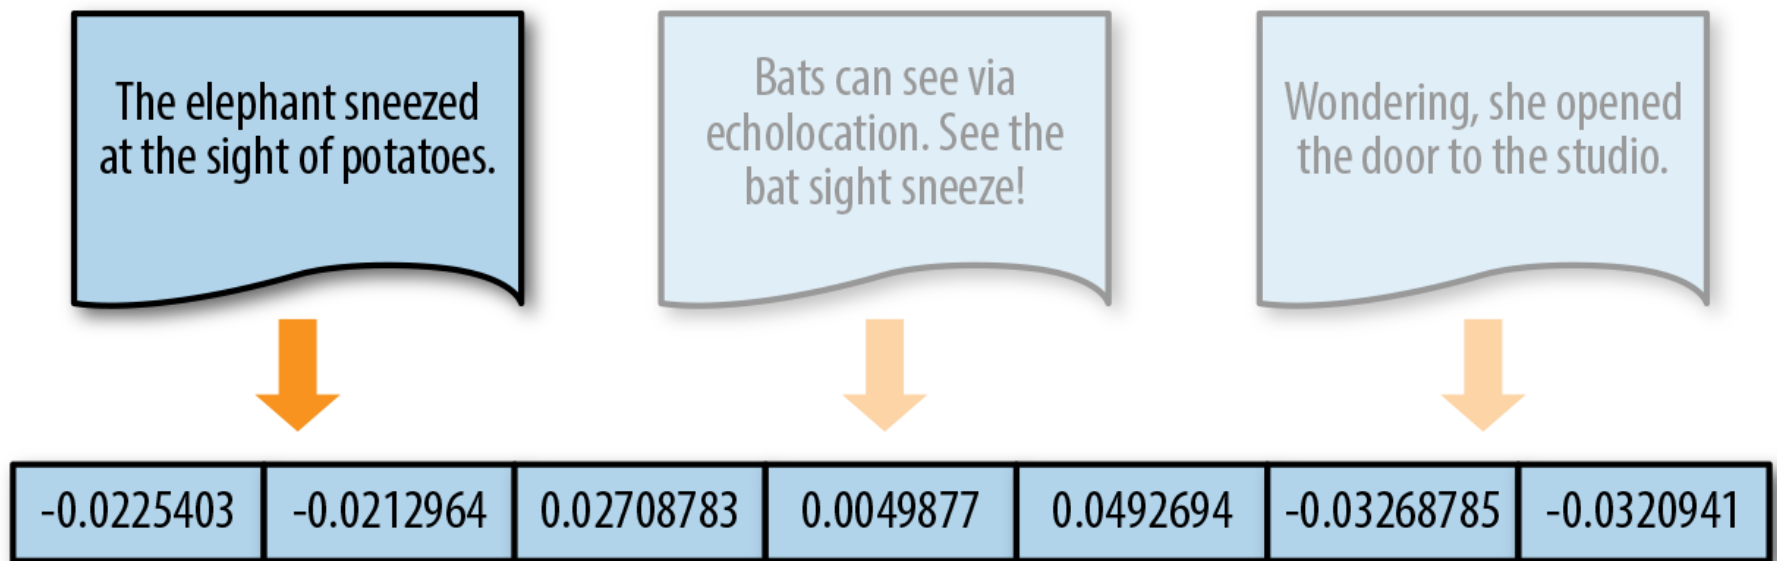
# One-hot Encoding



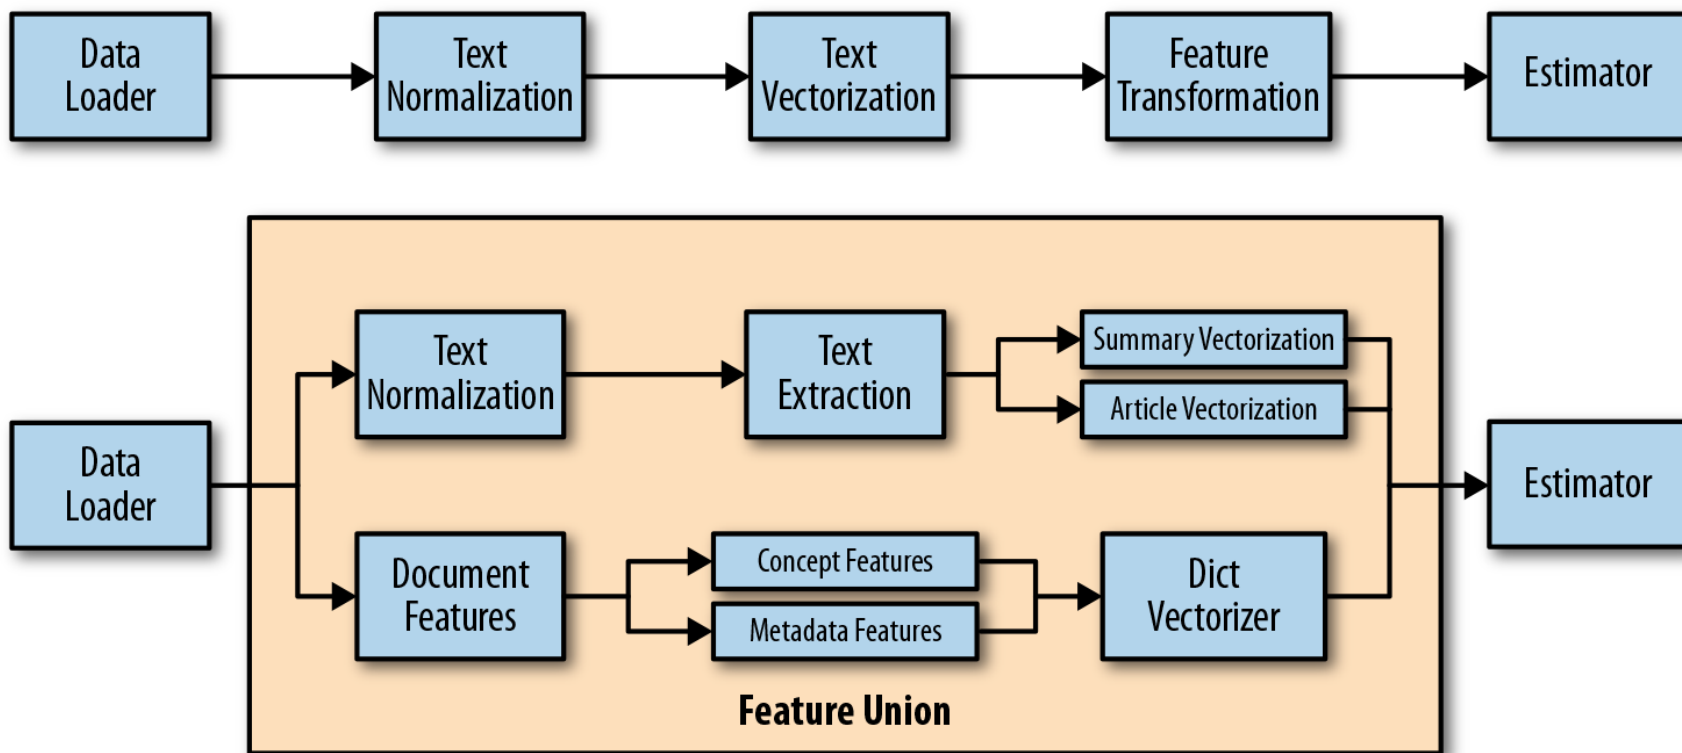
# TF-IDF Encoding



# Distributed Representation

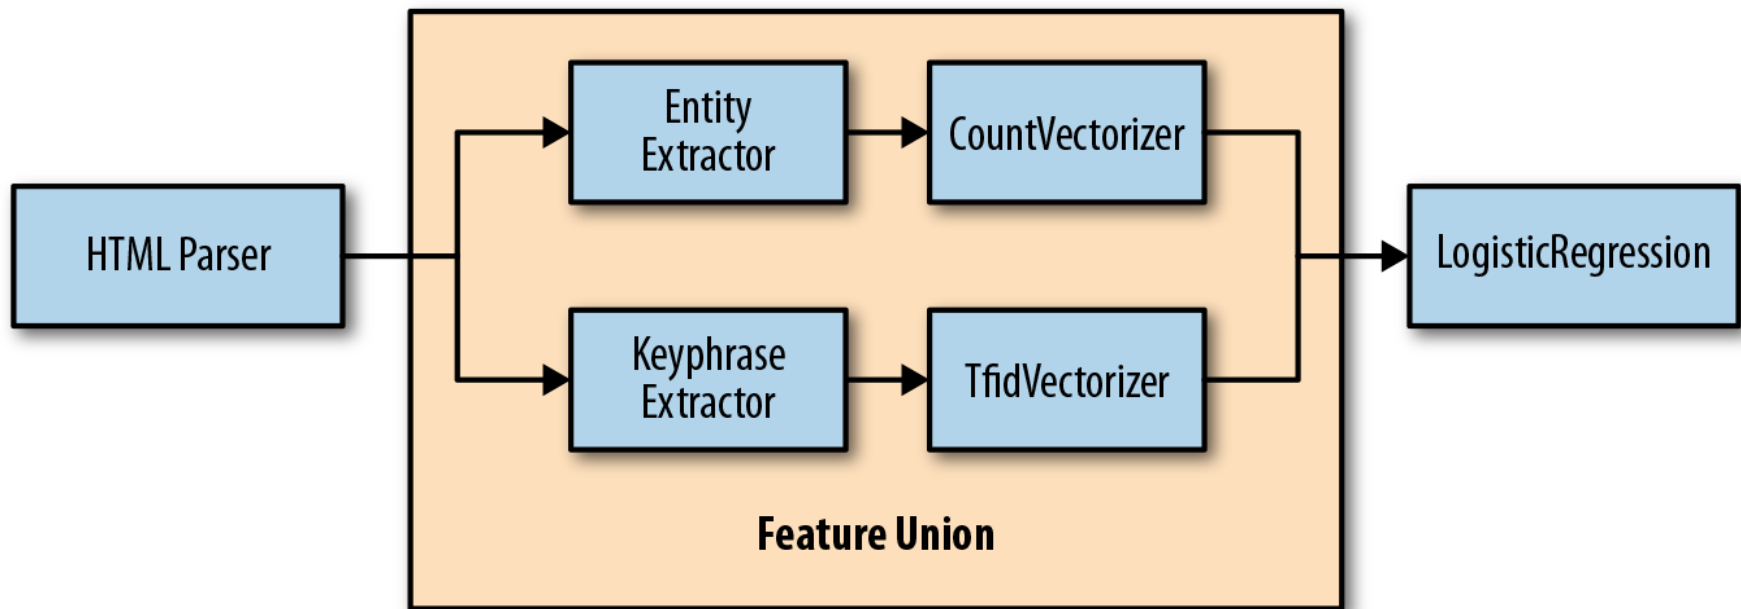


# Pipelines for Text Vectorization and Feature Extraction






# Feature Unions for Branching Vectorization

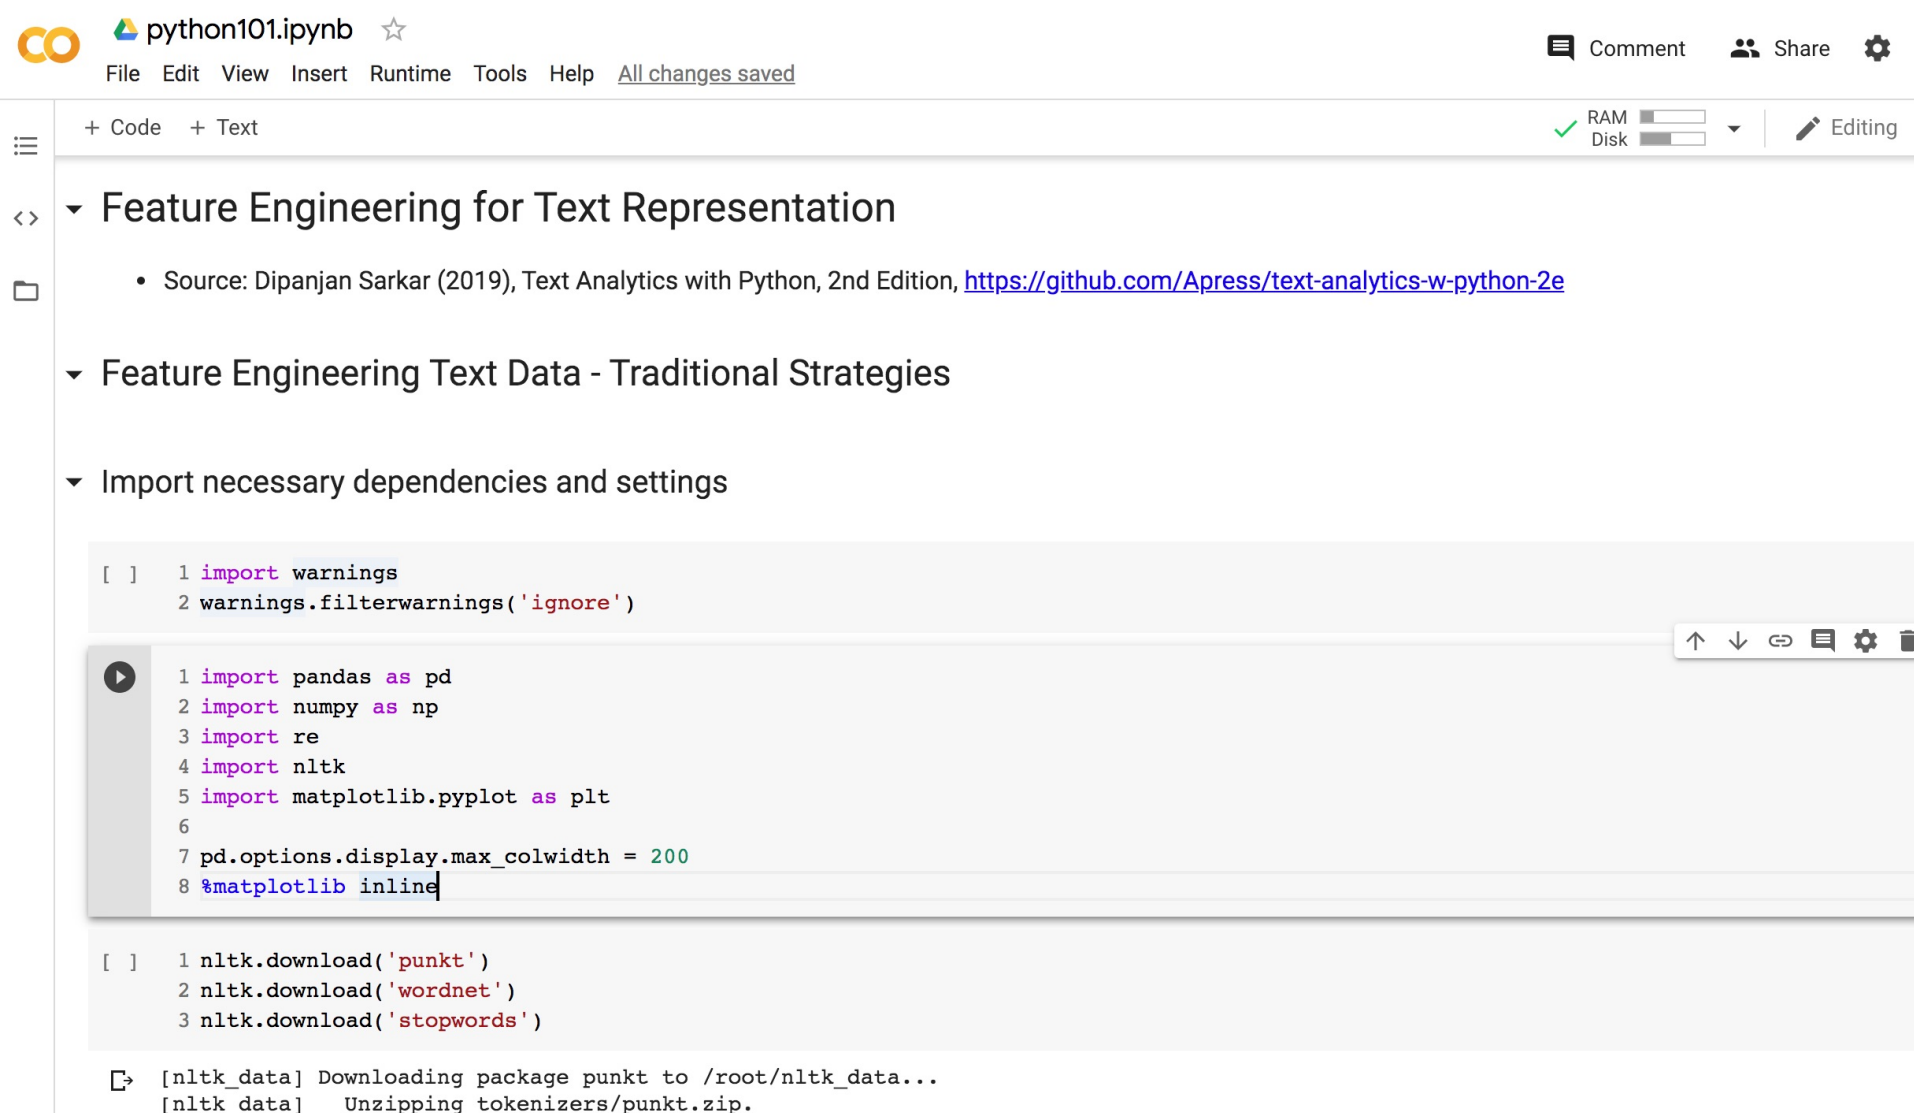


# Feature Extraction and Union



# Python in Google Colab (Python101)

<https://colab.research.google.com/drive/1FEG6DnGvwfUbeo4zJ1zTunjMqf2RkCrT>



The screenshot shows a Google Colab notebook interface. At the top, the notebook is titled 'python101.ipynb' and has a star icon. The menu bar includes 'File', 'Edit', 'View', 'Insert', 'Runtime', 'Tools', and 'Help', with a status indicator 'All changes saved'. On the right, there are icons for 'Comment', 'Share', and 'Settings'. Below the menu, there are controls for '+ Code' and '+ Text', and a status bar showing 'RAM' and 'Disk' usage with a green checkmark, and an 'Editing' mode indicator.

The notebook content is organized into a table of contents with expandable sections:

- Feature Engineering for Text Representation
  - Source: Dipanjan Sarkar (2019), Text Analytics with Python, 2nd Edition, <https://github.com/Apress/text-analytics-w-python-2e>
- Feature Engineering Text Data - Traditional Strategies
- Import necessary dependencies and settings

Below the table of contents, there are two code cells. The first cell contains the following code:

```
[ ] 1 import warnings
    2 warnings.filterwarnings('ignore')
```

The second cell contains the following code:

```
[ ] 1 import pandas as pd
    2 import numpy as np
    3 import re
    4 import nltk
    5 import matplotlib.pyplot as plt
    6
    7 pd.options.display.max_colwidth = 200
    8 %matplotlib inline
```

Below the code cells, there is a terminal output showing the execution of the second cell:

```
[ ] 1 nltk.download('punkt')
    2 nltk.download('wordnet')
    3 nltk.download('stopwords')
```

The terminal output shows the following messages:

```
[nltk_data] Downloading package punkt to /root/nltk_data...
[nltk_data] Unzipping tokenizers/punkt.zip.
```

<https://tinyurl.com/imtkupython101>

```
corpus = ['The sky is blue and beautiful.',
'Love this blue and beautiful sky!',
'The quick brown fox jumps over the lazy dog.',
"A king's breakfast has sausages, ham, bacon, eggs, toast and
beans",
'I love green eggs, ham, sausages and bacon!',
'The brown fox is quick and the blue dog is lazy!',
'The sky is very blue and the sky is very beautiful today',
'The dog is lazy but the brown fox is quick!'
]
labels = ['weather', 'weather', 'animals', 'food', 'food',
'animals', 'weather', 'animals']

corpus = np.array(corpus)
corpus_df = pd.DataFrame({'Document': corpus,
'Category': labels})
corpus_df = corpus_df[['Document', 'Category']]
corpus_df
```

```
corpus = np.array(corpus)
corpus_df = pd.DataFrame({'Document': corpus,
                          'Category': labels})
corpus_df = corpus_df[['Document', 'Category']]
corpus_df
```

	Document	Category
0	The sky is blue and beautiful.	weather
1	Love this blue and beautiful sky!	weather
2	The quick brown fox jumps over the lazy dog.	animals
3	A king's breakfast has sausages, ham, bacon, eggs, toast and beans	food
4	I love green eggs, ham, sausages and bacon!	food
5	The brown fox is quick and the blue dog is lazy!	animals
6	The sky is very blue and the sky is very beautiful today	weather
7	The dog is lazy but the brown fox is quick!	animals

```
wpt = nltk.WordPunctTokenizer()
stop_words = nltk.corpus.stopwords.words('english')

def normalize_document(doc):
    # lower case and remove special characters\whitespaces
    doc = re.sub(r'[^a-zA-Z\s]', '', doc, re.I|re.A)
    doc = doc.lower()
    doc = doc.strip()
    # tokenize document
    tokens = wpt.tokenize(doc)
    # filter stopwords out of document
    filtered_tokens = [token for token in tokens if token not in
stop_words]
    # re-create document from filtered tokens
    doc = ' '.join(filtered_tokens)
    return doc

normalize_corpus = np.vectorize(normalize_document)
norm_corpus = normalize_corpus(corpus)
norm_corpus
```

```
from sklearn.feature_extraction.text import CountVectorizer
# get bag of words features in sparse format
cv = CountVectorizer(min_df=0., max_df=1.)
cv_matrix = cv.fit_transform(norm_corpus)
cv_matrix
```

```
# view non-zero feature positions in the sparse matrix
print(cv_matrix)
```

```
# view dense representation
# warning might give a memory error if data is too big
cv_matrix = cv_matrix.toarray()
cv_matrix
```

```
array([[
0, 0, 1, 1, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 1, 0, 0],
[0, 0, 1, 1, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 1, 0, 0, 1, 0, 0],
[0, 0, 0, 0, 0, 1, 1, 0, 1, 0, 0, 1, 0, 1, 0, 1, 0, 0, 0, 0],
[1, 1, 0, 0, 1, 0, 0, 1, 0, 0, 1, 0, 1, 0, 0, 0, 1, 0, 1, 0],
[1, 0, 0, 0, 0, 0, 0, 1, 0, 1, 1, 0, 0, 0, 1, 0, 1, 0, 0, 0],
[0, 0, 0, 1, 0, 1, 1, 0, 1, 0, 0, 0, 0, 1, 0, 1, 0, 0, 0, 0],
[0, 0, 1, 1, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 2, 1],
[0, 0, 0, 0, 0, 1, 1, 0, 1, 0, 0, 0, 0, 1, 0, 1, 0, 0, 0, 0]])
```

```
# get all unique words in the corpus
vocab = cv.get_feature_names()
# show document feature vectors
pd.DataFrame(cv_matrix, columns=vocab)
```

```
1 # get all unique words in the corpus
2 vocab = cv.get_feature_names()
3 # show document feature vectors
4 pd.DataFrame(cv_matrix, columns=vocab)
```

	bacon	beans	beautiful	blue	breakfast	brown	dog	eggs	fox	green	ham	jumps	kings	lazy	love	quick	sausages	sky	toast	today
0	0	0	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	1	0	0
1	0	0	1	1	0	0	0	0	0	0	0	0	0	0	1	0	0	1	0	0
2	0	0	0	0	0	1	1	0	1	0	0	1	0	1	0	1	0	0	0	0
3	1	1	0	0	1	0	0	1	0	0	1	0	1	0	0	0	1	0	1	0
4	1	0	0	0	0	0	0	1	0	1	1	0	0	0	1	0	1	0	0	0
5	0	0	0	1	0	1	1	0	1	0	0	0	0	1	0	1	0	0	0	0
6	0	0	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	2	0	1
7	0	0	0	0	0	1	1	0	1	0	0	0	0	1	0	1	0	0	0	0



```

1 # you can set the n-gram range to 1,2 to get unigrams as well as bigrams
2 bv = CountVectorizer(ngram_range=(2,2))
3 bv_matrix = bv.fit_transform(norm_corpus)
4
5 bv_matrix = bv_matrix.toarray()
6 vocab = bv.get_feature_names()
7 pd.DataFrame(bv_matrix, columns=vocab)

```

	bacon eggs	beautiful sky	beautiful today	blue beautiful	blue dog	blue sky	breakfast sausages	brown fox	dog lazy	eggs ham	eggs toast	fox jumps	fox quick	green eggs	ham bacon	ham sausages	jumps lazy	kings breakfast
0	0	0	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0
1	0	1	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0
2	0	0	0	0	0	0	0	1	0	0	0	1	0	0	0	0	1	0
3	1	0	0	0	0	0	1	0	0	0	1	0	0	0	1	0	0	1
4	0	0	0	0	0	0	0	0	0	1	0	0	0	1	0	1	0	0
5	0	0	0	0	1	0	0	1	1	0	0	0	1	0	0	0	0	0
6	0	0	1	0	0	1	0	0	0	0	0	0	0	0	0	0	0	0
7	0	0	0	0	0	0	0	1	1	0	0	0	1	0	0	0	0	0

```

1 from sklearn.feature_extraction.text import TfidfTransformer
2
3 tt = TfidfTransformer(norm='l2', use_idf=True, smooth_idf=True)
4 tt_matrix = tt.fit_transform(cv_matrix)
5
6 tt_matrix = tt_matrix.toarray()
7 vocab = cv.get_feature_names()
8 pd.DataFrame(np.round(tt_matrix, 2), columns=vocab)

```

	bacon	beans	beautiful	blue	breakfast	brown	dog	eggs	fox	green	ham	jumps	kings	lazy	love	quick	sausages	sky	toast	today
0	0.00	0.00	0.60	0.53	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.60	0.00	0.0
1	0.00	0.00	0.49	0.43	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.57	0.00	0.00	0.49	0.00	0.0
2	0.00	0.00	0.00	0.00	0.00	0.38	0.38	0.00	0.38	0.00	0.00	0.53	0.00	0.38	0.00	0.38	0.00	0.00	0.00	0.0
3	0.32	0.38	0.00	0.00	0.38	0.00	0.00	0.32	0.00	0.00	0.32	0.00	0.38	0.00	0.00	0.00	0.32	0.00	0.38	0.0
4	0.39	0.00	0.00	0.00	0.00	0.00	0.00	0.39	0.00	0.47	0.39	0.00	0.00	0.00	0.39	0.00	0.39	0.00	0.00	0.0
5	0.00	0.00	0.00	0.37	0.00	0.42	0.42	0.00	0.42	0.00	0.00	0.00	0.00	0.42	0.00	0.42	0.00	0.00	0.00	0.0
6	0.00	0.00	0.36	0.32	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.72	0.00	0.5
7	0.00	0.00	0.00	0.00	0.00	0.45	0.45	0.00	0.45	0.00	0.00	0.00	0.00	0.45	0.00	0.45	0.00	0.00	0.00	0.0

```

1 from sklearn.feature_extraction.text import TfidfVectorizer
2
3 tv = TfidfVectorizer(min_df=0., max_df=1., norm='l2',
4                       use_idf=True, smooth_idf=True)
5 tv_matrix = tv.fit_transform(norm_corpus)
6 tv_matrix = tv_matrix.toarray()
7
8 vocab = tv.get_feature_names()
9 pd.DataFrame(np.round(tv_matrix, 2), columns=vocab)

```

	bacon	beans	beautiful	blue	breakfast	brown	dog	eggs	fox	green	ham	jumps	kings	lazy	love	quick	sausages	sky	toast	today
0	0.00	0.00	0.60	0.53	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.60	0.00	0.0
1	0.00	0.00	0.49	0.43	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.57	0.00	0.00	0.49	0.00	0.0
2	0.00	0.00	0.00	0.00	0.00	0.38	0.38	0.00	0.38	0.00	0.00	0.53	0.00	0.38	0.00	0.38	0.00	0.00	0.00	0.0
3	0.32	0.38	0.00	0.00	0.38	0.00	0.00	0.32	0.00	0.00	0.32	0.00	0.38	0.00	0.00	0.00	0.32	0.00	0.38	0.0
4	0.39	0.00	0.00	0.00	0.00	0.00	0.00	0.39	0.00	0.47	0.39	0.00	0.00	0.00	0.39	0.00	0.39	0.00	0.00	0.0
5	0.00	0.00	0.00	0.37	0.00	0.42	0.42	0.00	0.42	0.00	0.00	0.00	0.00	0.42	0.00	0.42	0.00	0.00	0.00	0.0
6	0.00	0.00	0.36	0.32	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.72	0.00	0.5
7	0.00	0.00	0.00	0.00	0.00	0.45	0.45	0.00	0.45	0.00	0.00	0.00	0.00	0.45	0.00	0.45	0.00	0.00	0.00	0.0

```

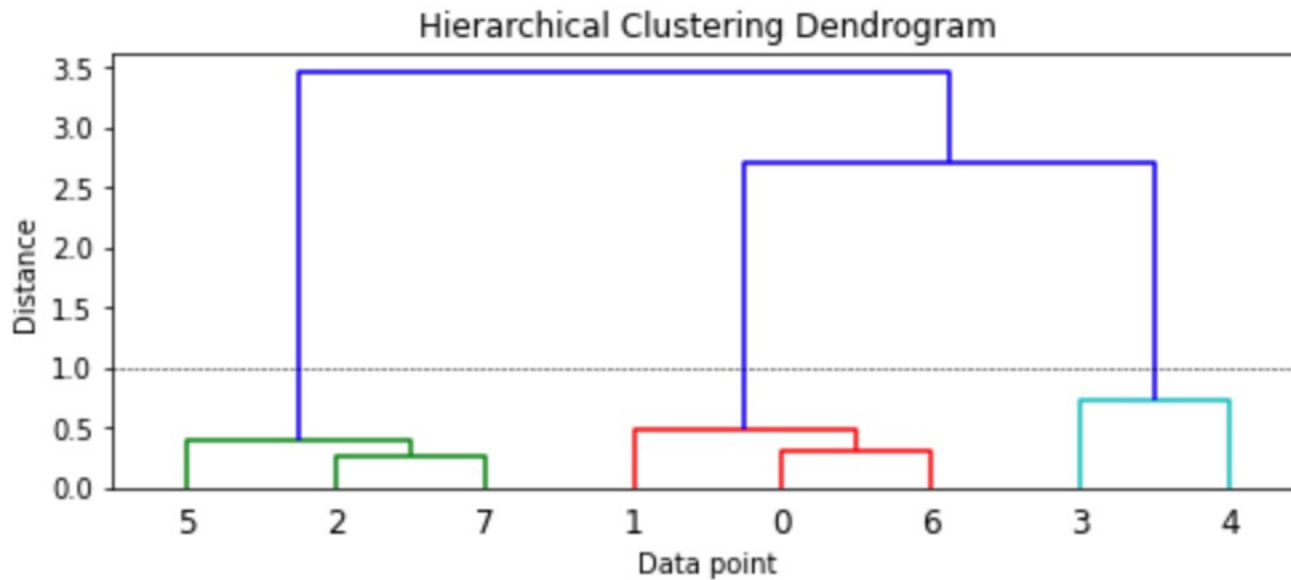
1 from scipy.cluster.hierarchy import dendrogram, linkage
2
3 Z = linkage(similarity_matrix, 'ward')
4 pd.DataFrame(Z, columns=['Document Cluster 1', 'Document Cluster 2',
5                          'Distance', 'Cluster Size'], dtype='object')

```

	Document Cluster 1	Document Cluster 2	Distance	Cluster Size
0	2	7	0.253098	2
1	0	6	0.308539	2
2	5	8	0.386952	3
3	1	9	0.489845	3
4	3	4	0.732945	2
5	11	12	2.69565	5
6	10	13	3.45108	8

```
1 plt.figure(figsize=(8, 3))
2 plt.title('Hierarchical Clustering Dendrogram')
3 plt.xlabel('Data point')
4 plt.ylabel('Distance')
5 dendrogram(Z)
6 plt.axhline(y=1.0, c='k', ls='--', lw=0.5)
```

<matplotlib.lines.Line2D at 0x7ff7b5d793c8>



```

1 from scipy.cluster.hierarchy import fcluster
2 max_dist = 1.0
3
4 cluster_labels = fcluster(Z, max_dist, criterion='distance')
5 cluster_labels = pd.DataFrame(cluster_labels, columns=['ClusterLabel'])
6 pd.concat([corpus_df, cluster_labels], axis=1)

```

	Document	Category	ClusterLabel
0	The sky is blue and beautiful.	weather	2
1	Love this blue and beautiful sky!	weather	2
2	The quick brown fox jumps over the lazy dog.	animals	1
3	A king's breakfast has sausages, ham, bacon, eggs, toast and beans	food	3
4	I love green eggs, ham, sausages and bacon!	food	3
5	The brown fox is quick and the blue dog is lazy!	animals	1
6	The sky is very blue and the sky is very beautiful today	weather	2
7	The dog is lazy but the brown fox is quick!	animals	1

```
1 from sklearn.decomposition import LatentDirichletAllocation
2 lda = LatentDirichletAllocation(n_components=3, max_iter=10000, random_state=0)
3 #lda = LatentDirichletAllocation(n_topics=3, max_iter=10000, random_state=0)
4 dt_matrix = lda.fit_transform(cv_matrix)
5 features = pd.DataFrame(dt_matrix, columns=['T1', 'T2', 'T3'])
6 features
```

	<b>T1</b>	<b>T2</b>	<b>T3</b>
<b>0</b>	0.832191	0.083480	0.084329
<b>1</b>	0.863554	0.069100	0.067346
<b>2</b>	0.047794	0.047776	0.904430
<b>3</b>	0.037243	0.925559	0.037198
<b>4</b>	0.049121	0.903076	0.047802
<b>5</b>	0.054902	0.047778	0.897321
<b>6</b>	0.888287	0.055697	0.056016
<b>7</b>	0.055704	0.055689	0.888607

```
1 tt_matrix = lda.components_  
2 for topic_weights in tt_matrix:  
3     topic = [(token, weight) for token, weight in zip(vocab, topic_weights)]  
4     topic = sorted(topic, key=lambda x: -x[1])  
5     topic = [item for item in topic if item[1] > 0.6]  
6     print(topic)  
7     print()  
  
[('sky', 4.332439442470133), ('blue', 3.373774254787669), ('beautiful', 3.3323650509884386), ('today', 1.3325579855138987), ('love', 1.3323473548404405), ('bacon', 2.33269586574902), ('eggs', 2.33269586574902), ('ham', 2.33269586574902), ('sausages', 2.33269586574902), ('love', 1.3323473548404405), ('brown', 3.3323473548404405), ('dog', 3.3323473548404405), ('fox', 3.3323473548404405), ('lazy', 3.3323473548404405), ('quick', 1.3323473548404405)]
```



```

1 from sklearn.cluster import KMeans
2
3 km = KMeans(n_clusters=3, random_state=0)
4 km.fit_transform(features)
5 cluster_labels = km.labels_
6 cluster_labels = pd.DataFrame(cluster_labels, columns=['ClusterLabel'])
7 pd.concat([corpus_df, cluster_labels], axis=1)

```

	Document	Category	ClusterLabel
0	The sky is blue and beautiful.	weather	1
1	Love this blue and beautiful sky!	weather	1
2	The quick brown fox jumps over the lazy dog.	animals	2
3	A king's breakfast has sausages, ham, bacon, eggs, toast and beans	food	0
4	I love green eggs, ham, sausages and bacon!	food	0
5	The brown fox is quick and the blue dog is lazy!	animals	2
6	The sky is very blue and the sky is very beautiful today	weather	1
7	The dog is lazy but the brown fox is quick!	animals	2


```
from gensim.models import word2vec


# tokenize sentences in corpus
wpt = nltk.WordPunctTokenizer()
tokenized_corpus = [wpt.tokenize(document) for document in norm_bible]

# Set values for various parameters
feature_size = 100 # Word vector dimensionality
window_context = 30 # Context window size
min_word_count = 1 # Minimum word count
sample = 1e-3 # Downsample setting for frequent words

w2v_model = word2vec.Word2Vec(tokenized_corpus, size=feature_size,
window=window_context, min_count=min_word_count,
sample=sample, iter=50)

# view similar words based on gensim's model
similar_words = {search_term: [item[0] for item in
w2v_model.wv.most_similar([search_term], topn=5)]
for search_term in ['god', 'jesus', 'noah', 'egypt', 'john', 'gospel',
'moses', 'famine']}
similar_words
```





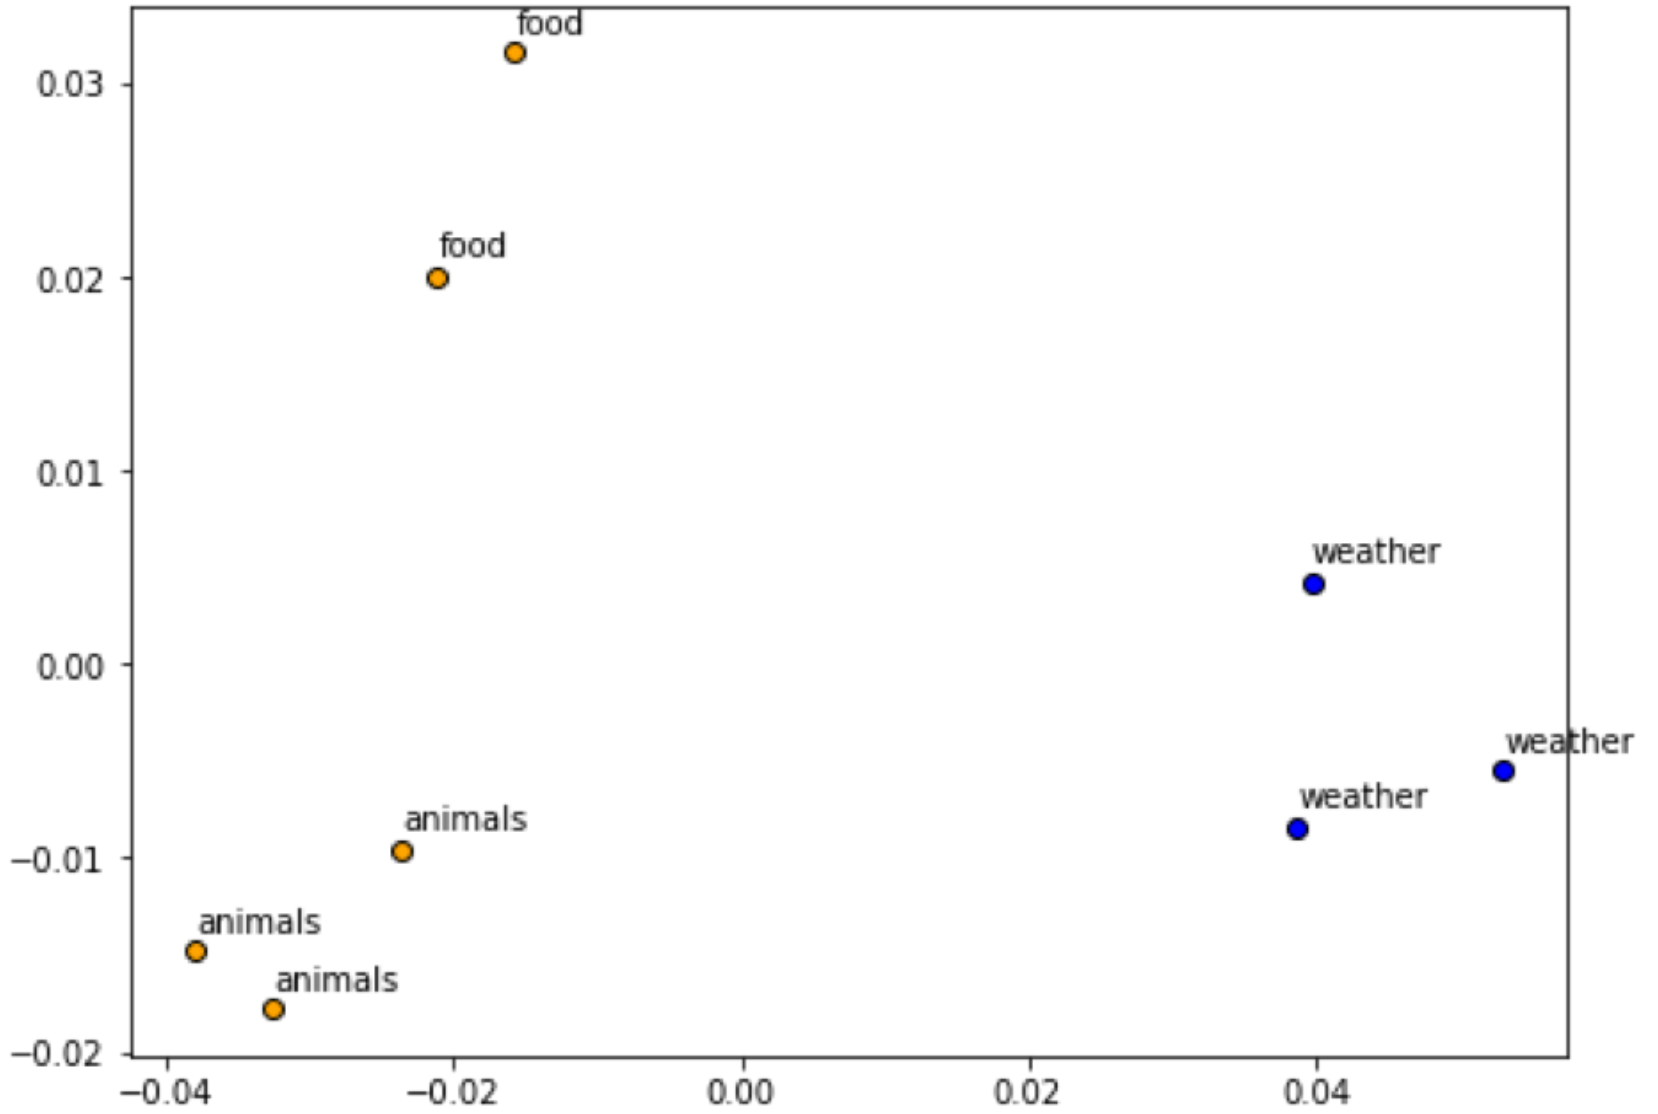
```
w2v_model.wv.most_similar([search_term], topn=5)]
```

```
{'egypt': ['egyptians', 'pharaoh', 'bondage', 'flowing',  
'rod'], 'famine': ['pestilence', 'peril', 'deaths',  
'morever', 'sword'], 'god': ['lord', 'worldly', 'soberly',  
'reasonable', 'unto'], 'gospel': ['christ', 'faith',  
'repentance', 'sufferings', 'afflictions'], 'jesus':  
['peter', 'messias', 'immediately', 'apostles',  
'synagogue'], 'john': ['james', 'baptist', 'devine',  
'peter', 'simon'], 'moses': ['congregation', 'elisheba',  
'naashon', 'joshua', 'children'], 'noah': ['shem',  
'japheth', 'ham', 'noe', 'hoglah']}
```

```
from sklearn.decomposition import PCA

pca = PCA(n_components=2, random_state=0)
pcs = pca.fit_transform(w2v_feature_array)
labels = ap.labels_
categories = list(corpus_df['Category'])
plt.figure(figsize=(8, 6))

for i in range(len(labels)):
    label = labels[i]
    color = 'orange' if label == 0 else 'blue' if label == 1
    else 'green'
    annotation_label = categories[i]
    x, y = pcs[i]
    plt.scatter(x, y, c=color, edgecolors='k')
    plt.annotate(annotation_label, xy=(x+1e-4, y+1e-3),
xytext=(0, 0), textcoords='offset points')
```



# BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding

**BERT: Pre-training of Deep Bidirectional Transformers for  
Language Understanding**


**Jacob Devlin   Ming-Wei Chang   Kenton Lee   Kristina Toutanova**  
Google AI Language

`{jacobdevlin, mingweichang, kentonl, kristout}@google.com`



# BERT

## Bidirectional Encoder Representations from Transformers



## Pre-training model architectures


**BERT** uses a bidirectional Transformer.

**OpenAI GPT** uses a left-to-right Transformer.

**ELMo** uses the concatenation of independently trained left-to-right and right-to-left LSTM to generate features for downstream tasks.


Among three, only BERT representations are jointly conditioned on both left and right context in all layers.

# BERT input representation




The input embeddings is the sum of the token embeddings, the segmentation embeddings and the position embeddings.

# BERT Sequence-level tasks




(a) Sentence Pair Classification Tasks:  
MNLI, QQP, QNLI, STS-B, MRPC,  
RTE, SWAG




(b) Single Sentence Classification Tasks:  
SST-2, CoLA

# BERT Token-level tasks



(c) Question Answering Tasks:  
SQuAD v1.1



(d) Single Sentence Tagging Tasks:  
CoNLL-2003 NER

# General Language Understanding Evaluation (GLUE) benchmark

## GLUE Test results

System	MNLI-(m/mm) 392k	QQP 363k	QNLI 108k	SST-2 67k	CoLA 8.5k	STS-B 5.7k	MRPC 3.5k	RTE 2.5k	Average -
Pre-OpenAI SOTA	80.6/80.1	66.1	82.3	93.2	35.0	81.0	86.0	61.7	74.0
BiLSTM+ELMo+Attn	76.4/76.1	64.8	79.9	90.4	36.0	73.3	84.9	56.8	71.0
OpenAI GPT	82.1/81.4	70.3	88.1	91.3	45.4	80.0	82.3	56.0	75.2
BERT <sub>BASE</sub>	84.6/83.4	71.2	90.1	93.5	52.1	85.8	88.9	66.4	79.6
BERT <sub>LARGE</sub>	<b>86.7/85.9</b>	<b>72.1</b>	<b>91.1</b>	<b>94.9</b>	<b>60.5</b>	<b>86.5</b>	<b>89.3</b>	<b>70.1</b>	<b>81.9</b>

**MNLI:** Multi-Genre Natural Language Inference

**QQP:** Quora Question Pairs

**QNLI:** Question Natural Language Inference

**SST-2:** The Stanford Sentiment Treebank

**CoLA:** The Corpus of Linguistic Acceptability

**STS-B:** The Semantic Textual Similarity Benchmark

**MRPC:** Microsoft Research Paraphrase Corpus

**RTE:** Recognizing Textual Entailment

Source: Devlin, Jacob, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova (2018).

"BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding." arXiv preprint arXiv:1810.04805

# Facebook Research FastText

Pre-trained word vectors

Word2Vec

wiki.zh.vec (861MB)

332647 word

300 vec

Pre-trained word vectors for 90 languages,  
trained on Wikipedia using fastText.

These vectors in dimension 300 were obtained using  
the skip-gram model with default parameters.

<https://github.com/facebookresearch/fastText/blob/master/pretrained-vectors.md>

# Facebook Research FastText

## Word2Vec: wiki.zh.vec

(861MB) (332647 word 300 vec)

wiki.zh.vec

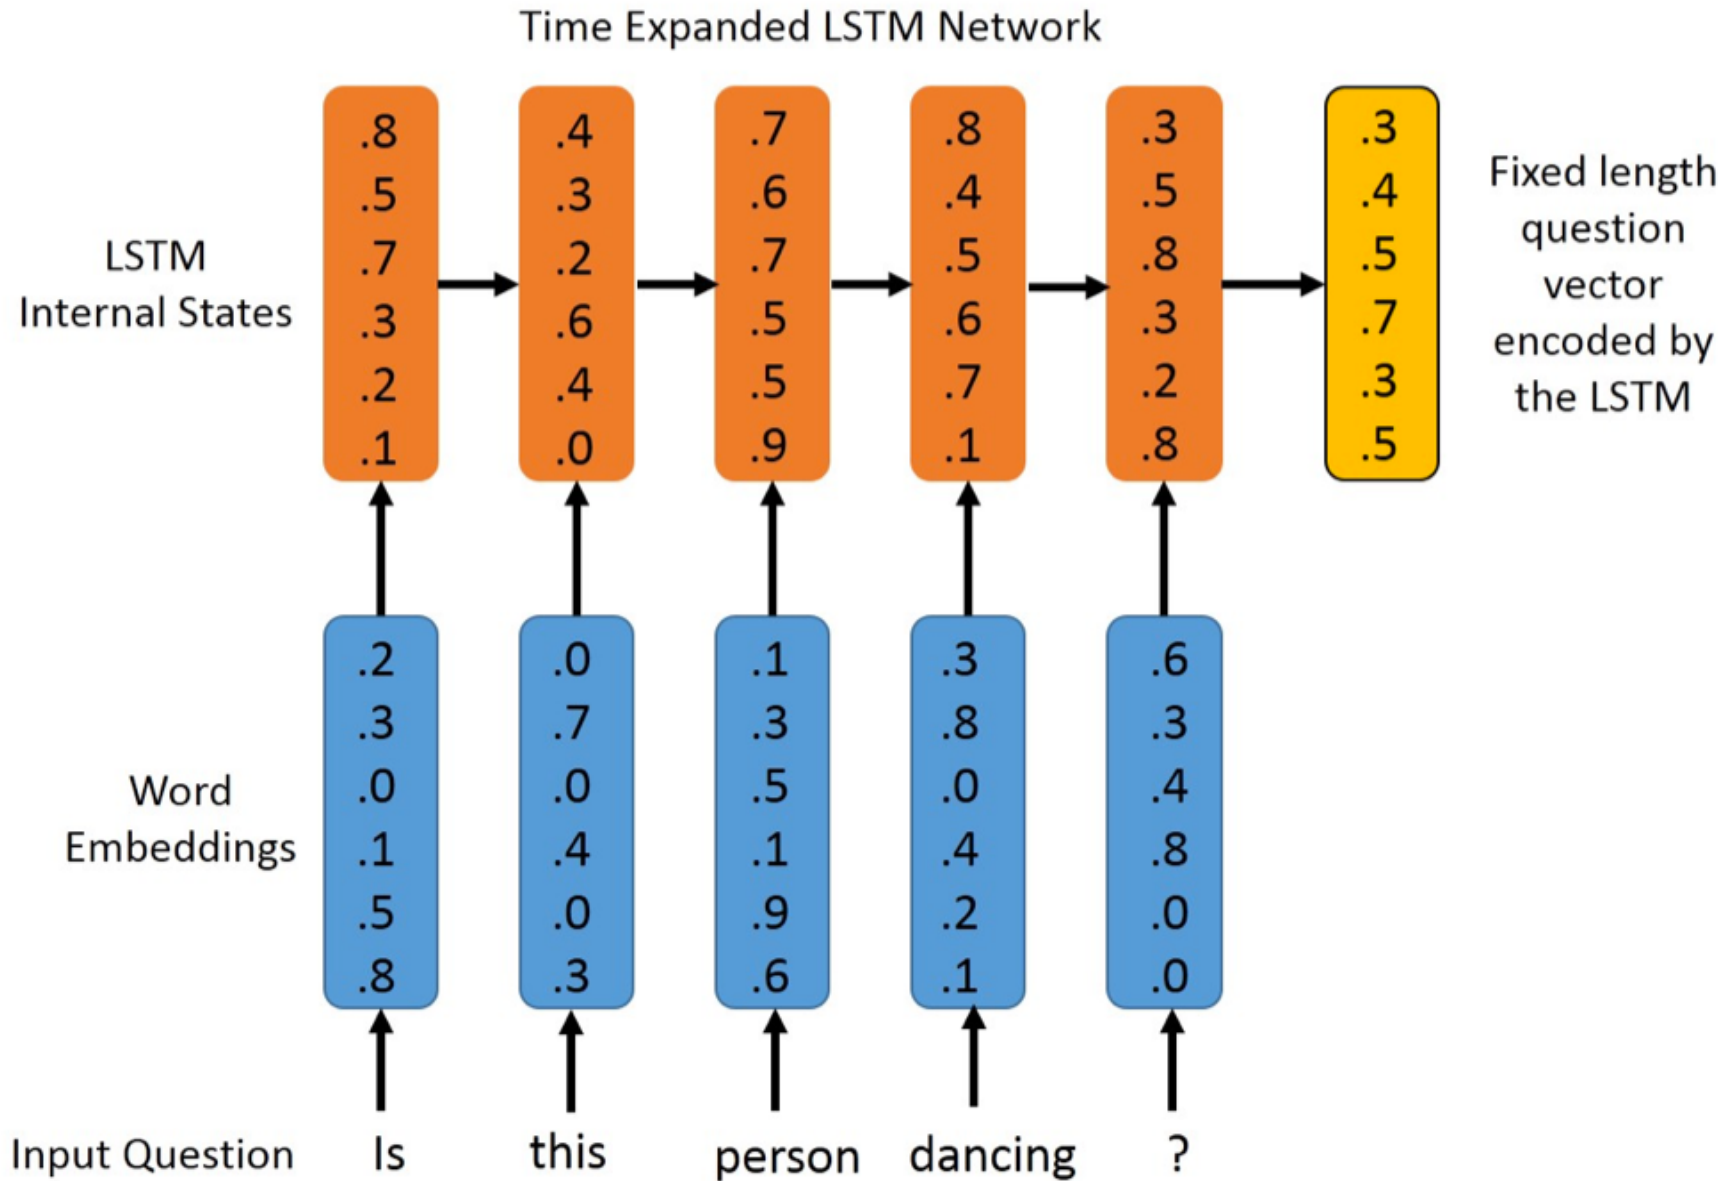
31845 yg -0.3978 0.49084 -0.54621 0.078991 0.8584 -0.26163 -0.45787 0.060828 0.36513 -0.03771 0.80791 0.16613 1.4828 -0.89862 0.085965  
31846 迴圈 -0.034834 0.71651 -0.4377 0.48344 0.31117 -0.51783 -0.40156 -0.057097 0.31535 -0.088301 0.23436 0.30884 1.2932 -0.6704 0.218  
31847 ぶっ -0.23267 0.39349 -0.90806 -0.53805 0.59308 -0.31819 -0.64229 0.16871 0.10086 0.09342 1.0914 -0.16019 1.6954 -0.70604 -0.218  
31848 三公 0.54129 0.55641 -0.4348 0.25094 0.1631 -0.10326 -0.54099 0.064742 0.13175 0.10217 0.84938 -0.10287 1.312 -0.74969 0.24025 -0  
31849 水貨 -0.14451 0.80455 -0.6145 0.55905 0.58307 -0.02559 -0.41088 -0.19056 -0.09178 0.33935 1.1927  
31850 剛才 0.19347 0.553 -0.64736 0.26358 0.83816 -0.24098 -0.83997 -0.16232 -0.024786 -0.2483 0.69732  
31851 無知 -0.0089777 0.90866 -0.25306 0.72983 0.67791 -0.3285 -0.63835 0.075295 0.4774 -0.04134 0.7210  
31852 好轉 -0.026068 0.92676 -0.47469 0.50129 0.67343 -0.32509 -0.32917 0.066499 0.3875 0.0011722 0.66  
31853 紀事 0.40541 0.67654 -0.5351 0.30329 0.43042 -0.24675 -0.19287 0.34207 0.35516 -0.076331 0.85916  
31854 變回 -0.089933 0.88136 -0.43524 0.59963 0.6403 -0.70981 -0.56788 -0.074018 0.16905 -0.086594 0.6  
31855 牟尼 -0.26578 0.6434 0.028982 -0.044001 0.88297 -0.17646 -0.64672 0.040483 0.43653 0.084908 0.74  
31856 埋藏 -0.0985 0.85082 -0.33363 0.24784 0.71518 -0.59054 -0.73731 0.050949 0.36726 -0.076886 0.817  
31857 正大 0.21069 0.27605 -0.83862 -0.099698 0.47894 -0.32196 -0.38288 -0.01892 0.40548 -0.029619 0.7  
31858 kis -0.30595 0.18482 -0.71287 -0.314 0.44776 -0.44245 -0.36447 -0.23723 0.00098801 -0.2528 0.60  
31859 合奏 0.1841 0.60874 -0.51376 -0.48002 0.21506 -0.55515 -0.71746 0.030735 0.39508 -0.40856 0.6226  
31860 精兵 0.25619 0.77186 -0.48847 0.23118 0.27254 0.21305 -0.3517 0.47305 0.24882 -0.34756 1.025 0.1  
31861 疲勞 -0.072521 1.0381 -0.51933 0.19421 0.67573 -0.45204 -0.20126 0.22704 0.44196 0.018401 0.3473  
31862 襪 -0.11771 1.4272 -1.0849 0.77532 0.87026 -0.6892 -0.3521 0.036517 0.42727 -0.1871 0.82789 -0.0  
31863 小貓 -0.21554 0.73988 -0.39628 0.044656 1.0602 -0.67047 -0.54102 0.11888 0.1693 0.19343 1.0841 0.  
31864 lai -0.25451 0.31596 -0.29228 -0.19144 0.99059 -0.24459 -0.66342 0.063093 -0.061142 -0.22749 0.6  
31865 偏東 -0.50835 1.0943 0.043918 0.29173 1.0161 -0.32493 -0.27305 0.026946 0.46811 -0.3874 1.4049 0.  
31866 大约是 -0.35726 -0.03476 -0.28672 0.075447 0.18175 -0.39421 -0.32088 0.025225 0.34808 0.074744 0.  
31867 franch -0.6046 -0.3235 0.024041 -0.2756 0.74761 -0.14654 0.0082566 -0.10071 0.53593 -0.17374 0.2  
31868 brazilian -0.54029 -0.63905 -0.094006 -0.68768 0.33263 -0.1583 -0.060424 0.20644 0.46234 -0.0764  
31869 夹竹桃 -0.4361 0.011429 -0.078896 -0.078186 0.37747 -0.052101 -0.096683 0.10769 0.62661 -0.37252  
31870 continent -0.37761 -0.72151 -0.42248 -0.81768 0.5016 -0.48569 0.13464 0.12644 0.32292 0.18099 0.  
31871 我还是 0.097443 0.28929 -0.14202 0.034027 0.50621 -0.1647 -0.45849 -0.16198 0.13965 -0.33451 0.61  
31872 vienna -0.25827 -0.050966 0.050502 -0.63466 0.4949 -0.17448 -0.59978 0.20269 0.37532 0.059419 0.  
31873 固态 -0.12678 0.4556 -0.27108 0.12506 0.52106 -0.058477 -0.69296 0.12162 0.26508 -0.089028 0.752  
31874 吉普 -0.33693 0.48335 -0.58455 0.13722 0.74856 -0.24529 -0.41125 -0.13832 0.33871 -0.12051 0.864  
31875 實物 0.030096 0.65756 -0.67982 0.2203 0.38492 -0.19001 -0.53136 -0.10322 0.24523 0.15287 0.92591  
31876 教职 0.11559 0.67087 -0.5111 0.14955 0.61417 -0.51571 -0.47901 0.29445 0.37629 -0.24232 0.4608 -0  
31877 惕 0.50469 1.5357 -0.64393 0.48668 0.69479 -0.23443 -0.47863 0.16288 0.3347 -0.51673 0.86777 0.0  
31878 岸上 0.088323 0.85815 -0.485 0.30383 0.75965 -0.25031 -0.76678 0.12805 0.37641 -0.088752 0.65012  
31879 议和 0.26835 0.94854 -0.27972 0.097623 0.43305 -0.031361 -0.57406 0.21608 0.3324 -0.36823 0.6987  
31880 aka -0.21332 0.11216 -0.48872 -0.18531 0.79093 -0.34221 -0.51122 0.10067 0.29963 -0.075253 0.642  
31881 滑鐵盧 -0.28726 0.88014 -0.39751 -0.056992 0.37408 -0.16967 -0.20673 -0.048533 -0.1978 -0.13107 0

### Models

The models can be downloaded from:

- Afrikaans: [bin+text](#), [text](#)
- Albanian: [bin+text](#), [text](#)
- Arabic: [bin+text](#), [text](#)
- Armenian: [bin+text](#), [text](#)
- Asturian: [bin+text](#), [text](#)
- Azerbaijani: [bin+text](#), [text](#)
- Bashkir: [bin+text](#), [text](#)
- Basque: [bin+text](#), [text](#)
- Belarusian: [bin+text](#), [text](#)
- Bengali: [bin+text](#), [text](#)
- Bosnian: [bin+text](#), [text](#)
- Breton: [bin+text](#), [text](#)
- Bulgarian: [bin+text](#), [text](#)
- Burmese: [bin+text](#), [text](#)
- Catalan: [bin+text](#), [text](#)
- Cebuano: [bin+text](#), [text](#)
- Chechen: [bin+text](#), [text](#)
- Chinese: [bin+text](#), [text](#)
- Chuvash: [bin+text](#), [text](#)
- Croatian: [bin+text](#), [text](#)
- Czech: [bin+text](#), [text](#)


# Word Embeddings in LSTM RNN





# Transformer (Attention is All You Need)

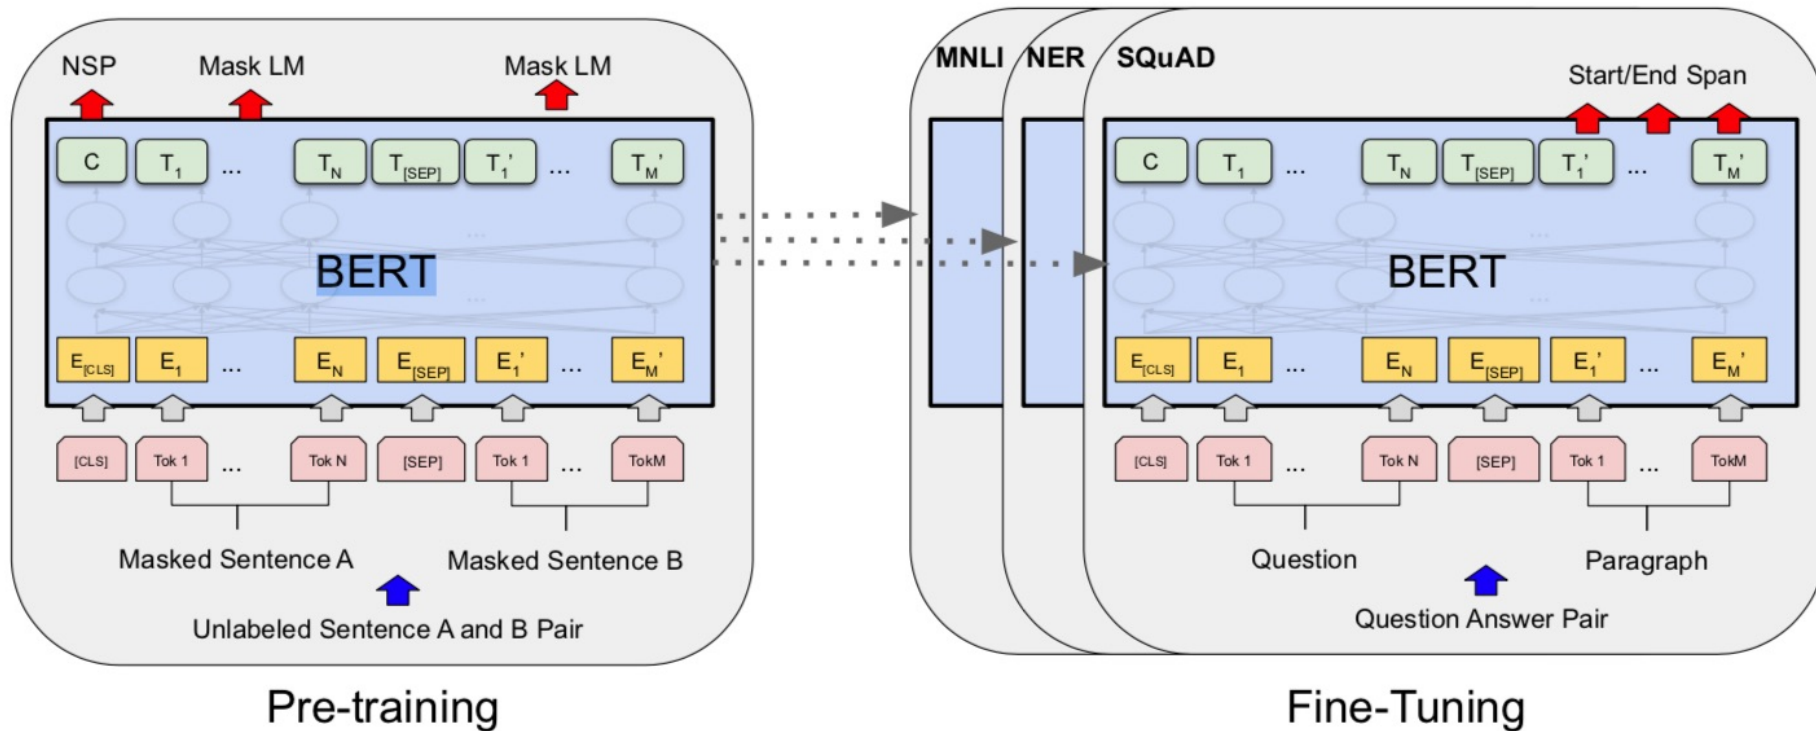
(Vaswani et al., 2017)



# BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding

BERT (Bidirectional Encoder Representations from Transformers)


Overall pre-training and fine-tuning procedures for BERT






# BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding

BERT (Bidirectional Encoder Representations from Transformers)

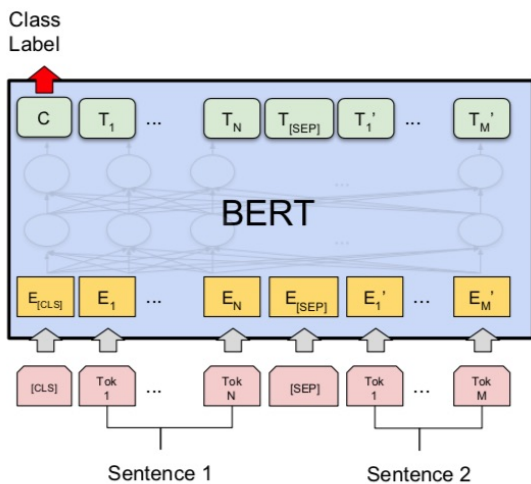
## BERT input representation




# BERT, OpenAI GPT, ELMo



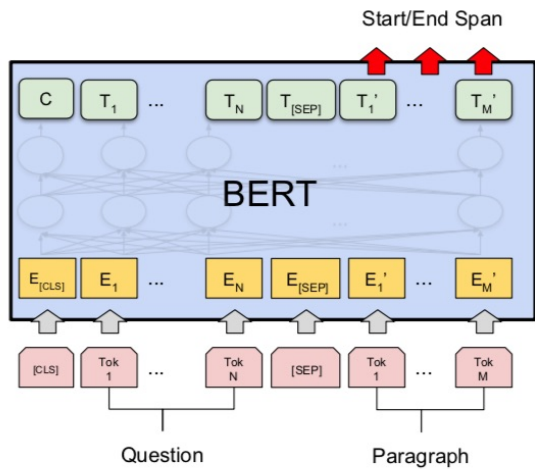
# Fine-tuning BERT on Different Tasks




(a) Sentence Pair Classification Tasks:  
MNLI, QQP, QNLI, STS-B, MRPC,  
RTE, SWAG



(b) Single Sentence Classification Tasks:  
SST-2, CoLA



(c) Question Answering Tasks:  
SQuAD v1.1




(d) Single Sentence Tagging Tasks:  
CoNLL-2003 NER

Source: Devlin, Jacob, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova (2018).


"Bert: Pre-training of deep bidirectional transformers for language understanding." arXiv preprint arXiv:1810.04805.

# Pre-trained Language Model (PLM)



By Xiaozhi Wang & Zhengyan Zhang @THUNLP

# Turing Natural Language Generation (T-NLG)



# Transformers Transformers

## State-of-the-art Natural Language Processing for TensorFlow 2.0 and PyTorch

- Transformers
  - pytorch-transformers
  - pytorch-pretrained-bert
- provides state-of-the-art general-purpose architectures
  - (BERT, GPT-2, RoBERTa, XLM, DistilBert, XLNet, CTRL...)
  - for Natural Language Understanding (NLU) and Natural Language Generation (NLG) with over 32+ pretrained models in 100+ languages and deep interoperability between TensorFlow 2.0 and PyTorch.




# Transfer Learning in Natural Language Processing

Source: Sebastian Ruder, Matthew E. Peters, Swabha Swayamdipta, and Thomas Wolf (2019), "Transfer learning in natural language processing." In Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Tutorials, pp. 15-18.

# NLP Benchmark Datasets

Task	Dataset	Link
Machine Translation	WMT 2014 EN-DE WMT 2014 EN-FR	<a href="http://www-lium.univ-lemans.fr/~schwenk/csmlm_joint_paper/">http://www-lium.univ-lemans.fr/~schwenk/csmlm_joint_paper/</a>
Text Summarization	CNN/DM Newsroom DUC Gigaword	<a href="https://cs.nyu.edu/~kcho/DMQA/">https://cs.nyu.edu/~kcho/DMQA/</a> <a href="https://summari.es/">https://summari.es/</a> <a href="https://www-nlpir.nist.gov/projects/duc/data.html">https://www-nlpir.nist.gov/projects/duc/data.html</a> <a href="https://catalog.ldc.upenn.edu/LDC2012T21">https://catalog.ldc.upenn.edu/LDC2012T21</a>
Reading Comprehension Question Answering Question Generation	ARC CliCR CNN/DM NewsQA RACE SQuAD Story Cloze Test NarrativeQA Quasar SearchQA	<a href="http://data.allenai.org/arc/">http://data.allenai.org/arc/</a> <a href="http://aclweb.org/anthology/N18-1140">http://aclweb.org/anthology/N18-1140</a> <a href="https://cs.nyu.edu/~kcho/DMQA/">https://cs.nyu.edu/~kcho/DMQA/</a> <a href="https://datasets.maluuba.com/NewsQA">https://datasets.maluuba.com/NewsQA</a> <a href="http://www.qizhexie.com/data/RACE_leaderboard">http://www.qizhexie.com/data/RACE_leaderboard</a> <a href="https://rajpurkar.github.io/SQuAD-explorer/">https://rajpurkar.github.io/SQuAD-explorer/</a> <a href="http://aclweb.org/anthology/W17-0906.pdf">http://aclweb.org/anthology/W17-0906.pdf</a> <a href="https://github.com/deepmind/narrativeqa">https://github.com/deepmind/narrativeqa</a> <a href="https://github.com/bdhingra/quasar">https://github.com/bdhingra/quasar</a> <a href="https://github.com/nyu-dl/SearchQA">https://github.com/nyu-dl/SearchQA</a>
Semantic Parsing	AMR parsing ATIS (SQL Parsing) WikiSQL (SQL Parsing)	<a href="https://amr.isi.edu/index.html">https://amr.isi.edu/index.html</a> <a href="https://github.com/jkkummerfeld/text2sql-data/tree/master/data">https://github.com/jkkummerfeld/text2sql-data/tree/master/data</a> <a href="https://github.com/salesforce/WikiSQL">https://github.com/salesforce/WikiSQL</a>
Sentiment Analysis	IMDB Reviews SST Yelp Reviews Subjectivity Dataset	<a href="http://ai.stanford.edu/~amaas/data/sentiment/">http://ai.stanford.edu/~amaas/data/sentiment/</a> <a href="https://nlp.stanford.edu/sentiment/index.html">https://nlp.stanford.edu/sentiment/index.html</a> <a href="https://www.yelp.com/dataset/challenge">https://www.yelp.com/dataset/challenge</a> <a href="http://www.cs.cornell.edu/people/pabo/movie-review-data/">http://www.cs.cornell.edu/people/pabo/movie-review-data/</a>
Text Classification	AG News DBpedia TREC 20 NewsGroup	<a href="http://www.di.unipi.it/~gulli/AG_corpus_of_news_articles.html">http://www.di.unipi.it/~gulli/AG_corpus_of_news_articles.html</a> <a href="https://wiki.dbpedia.org/Datasets">https://wiki.dbpedia.org/Datasets</a> <a href="https://trec.nist.gov/data.html">https://trec.nist.gov/data.html</a> <a href="http://qwone.com/~jason/20Newsgroups/">http://qwone.com/~jason/20Newsgroups/</a>
Natural Language Inference	SNLI Corpus MultiNLI SciTail	<a href="https://nlp.stanford.edu/projects/snli/">https://nlp.stanford.edu/projects/snli/</a> <a href="https://www.nyu.edu/projects/bowman/multinli/">https://www.nyu.edu/projects/bowman/multinli/</a> <a href="http://data.allenai.org/scitail/">http://data.allenai.org/scitail/</a>
Semantic Role Labeling	Proposition Bank OneNotes	<a href="http://propbank.github.io/">http://propbank.github.io/</a> <a href="https://catalog.ldc.upenn.edu/LDC2013T19">https://catalog.ldc.upenn.edu/LDC2013T19</a>

Aurélien Géron (2019),  
**Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow:  
Concepts, Tools, and Techniques to Build Intelligent Systems, 2nd Edition**  
O'Reilly Media, 2019




<https://github.com/ageron/handson-ml2>

# Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow

## Notebooks

- [1. The Machine Learning landscape](#)
- [2. End-to-end Machine Learning project](#)
- [3. Classification](#)
- [4. Training Models](#)
- [5. Support Vector Machines](#)
- [6. Decision Trees](#)
- [7. Ensemble Learning and Random Forests](#)
- [8. Dimensionality Reduction](#)
- [9. Unsupervised Learning Techniques](#)
- [10. Artificial Neural Nets with Keras](#)
- [11. Training Deep Neural Networks](#)
- [12. Custom Models and Training with TensorFlow](#)
- [13. Loading and Preprocessing Data](#)
- [14. Deep Computer Vision Using Convolutional Neural Networks](#)
- [15. Processing Sequences Using RNNs and CNNs](#)
- [16. Natural Language Processing with RNNs and Attention](#)
- [17. Representation Learning Using Autoencoders](#)
- [18. Reinforcement Learning](#)
- [19. Training and Deploying TensorFlow Models at Scale](#)




# Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow

github.com/ageron/handson-ml2

ageron loss = metric \* mean of sample weights, fixes #63

datasets	Fix vertical bars	
docker	Remove pyvirtualdisplay from environment.yml and add it to the Docker...	
images	Add breakout.gif	
work_in_progress	Remove from __future__ imports as we move away from Python 2	
.gitignore	Add jsb_chorales dataset to .gitignore	
01_the_machine_learning_landsc...	Fix typo on import urllib	
02_end_to_end_machine_learning_...	Make notebooks 1 to 9 runnable in Colab without changes	
03_classification.ipynb	Make notebooks 1 to 9 runnable in Colab without changes	
04_training_linear_models.ipynb	Make notebooks 1 to 9 runnable in Colab without changes	
05_support_vector_machines.ipynb	Make notebooks 1 to 9 runnable in Colab without changes	
06_decision_trees.ipynb	Make notebooks 1 to 9 runnable in Colab without changes	
07_ensemble_learning_and_rando...	Make notebooks 1 to 9 runnable in Colab without changes	13 days ago
08_dimensionality_reduction.ipynb	Make notebooks 1 to 9 runnable in Colab without changes	13 days ago
09_unsupervised_learning.ipynb	Make notebooks 1 to 9 runnable in Colab without changes	13 days ago
10_neural_nets_with_keras.ipynb	Make notebooks 10 and 11 runnable in Colab without changes	13 days ago
11_training_deep_neural_networks....	Make notebooks 10 and 11 runnable in Colab without changes	13 days ago
12_custom_models_and_training_...	loss = metric * mean of sample weights, fixes #63	6 days ago
13_loading_and_preprocessing_da...	Make notebook 13 runnable in Colab without changes	13 days ago
14_deep_computer_vision_with_cn...	Make notebooks 14 to 19 runnable in Colab without changes	13 days ago
15_processing_sequences_using_r...	Make notebooks 14 to 19 runnable in Colab without changes	13 days ago



Aurélien Géron

# Sequences using RNNs and CNNs




15\_processing\_sequences\_using\_rnn\_and\_cnns.ipynb

File Edit View Insert Runtime Tools Help Last edited on November 6 by ageron

Share


+ Code + Text Copy to Drive

Connect Editing



```
[ ] 1 np.random.seed(43)
2
3 series = generate_time_series(1, 50 + 10)
4 X_new, Y_new = series[:, :50, :], series[:, 50:, :]
5 Y_pred = model.predict(X_new[:, -1][..., np.newaxis])
```

```
[ ] 1 plot_multiple_forecasts(X_new, Y_new, Y_pred)
2 plt.show()
```





TensorFlow

# TensorFlow



Install

Learn ▾

API ▾

Resources ▾

More ▾

🔍 Search

Language ▾

GitHub

Sign in

Missed TensorFlow World? Check out the recap.

Learn more

An end-to-end open  
source machine  
learning platform

TensorFlow

For JavaScript

For Mobile & IoT

For Production

The core open source library to help you develop and train ML models. Get started quickly by running Colab notebooks directly in your browser.

Get started with TensorFlow







**TensorFlow**  
is an  
**Open Source**  
**Software Library**  
for  
**Machine Intelligence**

# TensorFlow 2.0

```
import tensorflow as tf
mnist = tf.keras.datasets.mnist

(x_train, y_train), (x_test, y_test) = mnist.load_data()
x_train, x_test = x_train / 255.0, x_test / 255.0

model = tf.keras.models.Sequential([
    tf.keras.layers.Flatten(input_shape=(28, 28)),
    tf.keras.layers.Dense(128, activation='relu'),
    tf.keras.layers.Dropout(0.2),
    tf.keras.layers.Dense(10, activation='softmax')
])

model.compile(optimizer='adam',
              loss='sparse_categorical_crossentropy',
              metrics=['accuracy'])

model.fit(x_train, y_train, epochs=5)
model.evaluate(x_test, y_test)
```

# TensorFlow Playground

Tinker With a **Neural Network** Right Here in Your Browser.  
Don't Worry, You Can't Break It. We Promise.



Iterations  
000,582

Learning rate  
0.03

Activation  
Tanh

Regularization  
None

Regularization rate  
0

Problem type  
Classification

## DATA

Which dataset do you want to use?



Ratio of training to test data: 50%



Noise: 0



Batch size: 10



## INPUT

Which properties do you want to feed in?

$X_1$



$X_2$



$X_1^2$



$X_2^2$



$X_1 X_2$



## 3 HIDDEN LAYERS

4 neurons



2 neurons



2 neurons




The outputs are mixed with varying **weights**, shown by the thickness of the lines.

This is the output from one **neuron**. Hover to see it larger.

## OUTPUT

Test loss 0.000  
Training loss 0.000



# Tensor

- **3**
  - # a rank 0 tensor; this is a **scalar** with shape []
- **[1., 2., 3.]**
  - # a rank 1 tensor; this is a **vector** with shape [3]
- **[[1., 2., 3.], [4., 5., 6.]]**
  - # a rank 2 tensor; a **matrix** with shape [2, 3]
- **[[[1., 2., 3.]], [[7., 8., 9.]]]**
  - # a rank 3 **tensor** with shape [2, 1, 3]

**Scalar**

80

**Vector**

[50 60 70]

**Matrix**

$$\begin{bmatrix} 50 & 60 & 70 \\ 55 & 65 & 75 \end{bmatrix}$$

**Tensor**

$$\begin{bmatrix} [50 & 60 & 70] & [70 & 80 & 90] \\ [55 & 65 & 75] & [75 & 85 & 95] \end{bmatrix}$$


# Deep Learning and Neural Networks

# Deep Learning and Neural Networks

**Input Layer  
(X)**

**Hidden Layer  
(H)**

**Output Layer  
(Y)**




# Deep Learning and Neural Networks

Input Layer  
(X)

Hidden Layer  
(H)

Output Layer  
(Y)






# Deep Learning and Neural Networks

Input Layer  
(X)

Hidden Layers  
(H)


Output Layer  
(Y)

Deep Neural Networks  
Deep Learning




# Convolutional Neural Networks


(CNN or Deep Convolutional Neural Networks, DCNN)



# Recurrent Neural Networks (RNN)




# Long / Short Term Memory (LSTM)




Hochreiter, Sepp, and Jürgen Schmidhuber. "Long short-term memory." *Neural computation* 9.8 (1997): 1735-1780.

Source: <http://www.asimovinstitute.org/neural-network-zoo/>

# Gated Recurrent Units (GRU)




# Generative Adversarial Networks (GAN)



Goodfellow, Ian, et al. "Generative adversarial nets." Advances in Neural Information Processing Systems. 2014.

Source: <http://www.asimovinstitute.org/neural-network-zoo/>

# Support Vector Machines (SVM)



Cortes, Corinna, and Vladimir Vapnik. "Support-vector networks." *Machine learning* 20.3 (1995): 273-297.


Source: <http://www.asimovinstitute.org/neural-network-zoo/>

# Neural Networks

**Input Layer**  
(X)

**Hidden Layer**  
(H)

**Output Layer**  
(Y)






# Neural Networks


**Input Layer**  
(X)

**Hidden Layer**  
(H)


**Output Layer**  
(Y)



# The Neuron




# Neuron and Synapse



# The Neuron


$$y = F\left(\sum_i w_i x_i\right)$$




$$F(x) = \max(0, x)$$

$$y = \max ( 0, -0.21 * x_1 + 0.3 * x_2 + 0.7 * x_3 )$$

Weights



# Neural Networks




# Neural Networks

**Input Layer**  
(X)

**Hidden Layer**  
(H)

**Output Layer**  
(Y)




# Neural Networks

Input Layer  
(X)

Hidden Layers  
(H)

Output Layer  
(Y)

Deep Neural Networks  
Deep Learning






# Neural Networks

Input Layer  
(X)

Hidden Layer  
(H)

Output Layer  
(Y)




# Neural Networks

**Input Layer**  
(X)

**Hidden Layer**  
(H)

**Output Layer**  
(Y)




# Neural Networks

Input Layer  
(X)

Hidden Layer  
(H)

Output Layer  
(Y)




# Neural Networks

**Input Layer**  
(X)

**Hidden Layer**  
(H)


**Output Layer**  
(Y)



<b>X</b>		<b>Y</b>
<b>Hours Sleep</b>	<b>Hours Study</b>	<b>Score</b>
<b>3</b>	<b>5</b>	<b>75</b>
<b>5</b>	<b>1</b>	<b>82</b>
<b>10</b>	<b>2</b>	<b>93</b>
<b>8</b>	<b>3</b>	<b>?</b>


	<b>X</b>		<b>Y</b>
	<b>Hours Sleep</b>	<b>Hours Study</b>	<b>Score</b>
<b>Training</b>	<b>3</b>	<b>5</b>	<b>75</b>
	<b>5</b>	<b>1</b>	<b>82</b>
	<b>10</b>	<b>2</b>	<b>93</b>
<b>Testing</b>	<b>8</b>	<b>3</b>	<b>?</b>

$$Y = WX + b$$






$$W X + b = Y$$



Scores  $\longrightarrow$  Probabilities

# SoftMAX

$$W X + b = Y$$



Logits


Scores

Probabilities

$$S(y_i) = \frac{e^{y_i}}{\sum_j e^{y_j}} = \frac{e^{2.0}}{e^{2.0} + e^{1.0} + e^{0.1}} = \frac{2.7182^{2.0}}{2.7182^{2.0} + 2.7182^{1.0} + 2.7182^{0.1}} = 0.7$$

$$S(y_i) = \frac{e^{y_i}}{\sum_j e^{y_j}} = \frac{e^{1.0}}{e^{2.0} + e^{1.0} + e^{0.1}} = \frac{2.7182^{1.0}}{2.7182^{2.0} + 2.7182^{1.0} + 2.7182^{0.1}} = 0.2$$

$$S(y_i) = \frac{e^{y_i}}{\sum_j e^{y_j}} = \frac{e^{0.1}}{e^{2.0} + e^{1.0} + e^{0.1}} = \frac{2.7182^{0.1}}{2.7182^{2.0} + 2.7182^{1.0} + 2.7182^{0.1}} = 0.1$$



**Logits**

**Scores**

**Probabilities**

**Training a Network**  
**=**  
**Minimize the Cost Function**


# Training a Network

=


Minimize the **Cost** Function

Minimize the **Loss** Function


**Error = Predict Y - Actual Y**  
**Error : Cost : Loss**



**Error = Predict Y - Actual Y**  
**Error : Cost : Loss**



**Error = Predict Y - Actual Y**  
**Error : Cost : Loss**





# Activation Functions




# Activation Functions

**Sigmoid**

**TanH**

**ReLU**

(Rectified Linear Unit)




**[0, 1]**

**[-1, 1]**


**$f(x) = \max(0, x)$**

# Activation Functions


Sigmoid



TanH



ReLU



# Loss Function

# Binary Classification: 2 Class

**Activation Function:  
Sigmoid**

**Loss Function:  
Binary Cross-Entropy**


# Multiple Classification: 10 Class

**Activation Function:  
SoftMAX**


**Loss Function:  
Categorical Cross-Entropy**

# Dropout

Dropout: a simple way to prevent neural networks from overfitting



(a) Standard Neural Net



(b) After applying dropout.

Source: Srivastava, Nitish, Geoffrey E. Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov.

"Dropout: a simple way to prevent neural networks from overfitting." *Journal of machine learning research* 15, no. 1 (2014): 1929-1958.

# Learning Algorithm

While not done:

Pick a random training example “(input, label)”


Run neural network on “input”

Adjust weights on edges to make output closer to “label”



$$y = \max ( 0, -0.21 * x_1 + 0.3 * x_2 + 0.7 * x_3 )$$

Weights




Next time:


$$y = \max(0, -0.23 * x_1 + 0.31 * x_2 + 0.65 * x_3)$$


~~$$y = \max(0, -0.21 * x_1 + 0.3 * x_2 + 0.7 * x_3)$$~~

Weights




# Optimizer: Stochastic Gradient Descent (SGD)





*This shows a function of 2 variables: real neural nets are functions of hundreds of millions of variables!*

# Neural Network and Deep Learning



Source: 3Blue1Brown (2017), But what \*is\* a Neural Network? | Chapter 1, deep learning, <https://www.youtube.com/watch?v=aircAruvnKk>

# Gradient Descent

## how neural networks learn


Average cost of  
all training data...

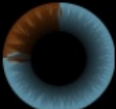
Cost of




$$\left\{ \begin{array}{l} (0.18 - 0.00)^2 + \\ (0.29 - 0.00)^2 + \\ (0.58 - 0.00)^2 + \\ (0.77 - 0.00)^2 + \\ (0.20 - 0.00)^2 + \\ (0.36 - 0.00)^2 + \\ (0.93 - 0.00)^2 + \\ (1.00 - 0.00)^2 + \\ (0.95 - 1.00)^2 + \\ (0.35 - 0.00)^2 \end{array} \right.$$

What's the "cost" <sup>i</sup>  
of this difference?



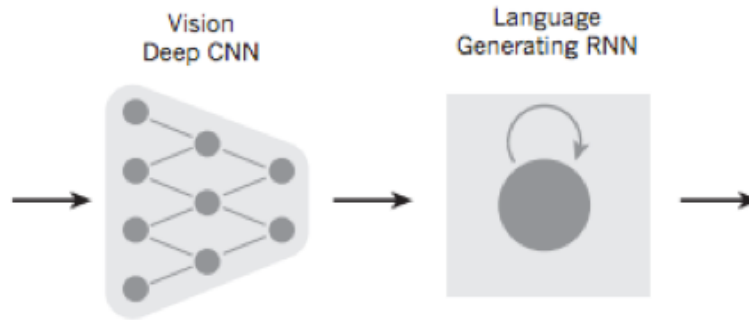
Utter trash 

# Backpropagation



Source: 3Blue1Brown (2017), What is backpropagation really doing? | Chapter 3, deep learning, <https://www.youtube.com/watch?v=llg3gGewQ5U>

# From image to text



A group of people shopping at an outdoor market.

There are many vegetables at the fruit stand.




A woman is throwing a **frisbee** in a park.




A **dog** is standing on a hardwood floor.



A **stop** sign is on a road with a mountain in the background



A little **girl** sitting on a bed with a teddy bear.



A group of **people** sitting on a boat in the water.



A giraffe standing in a forest with **trees** in the background.



# From image to text

Image: deep convolution neural network (CNN)

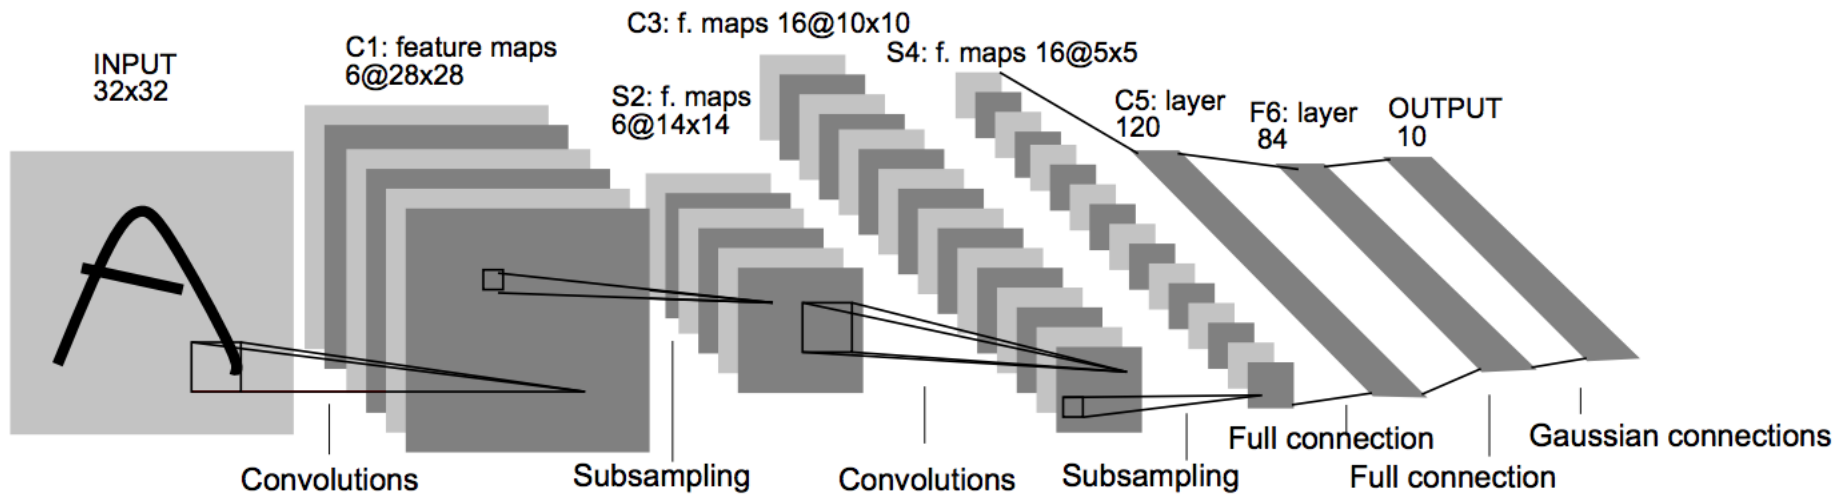
Text: recurrent neural network (RNN)



A group of **people** sitting on a boat in the water.

# Convolutional Neural Networks (CNN)

# Convolutional Neural Networks (CNN)



## Architecture of LeNet-5 (7 Layers) (LeCun et al., 1998)

Source: <http://yann.lecun.com/exdb/publis/pdf/lecun-01a.pdf>


Source: LeCun, Yann, Léon Bottou, Yoshua Bengio, and Patrick Haffner.

"Gradient-based learning applied to document recognition." *Proceedings of the IEEE* 86, no. 11 (1998): 2278-2324.

# Convolutional Neural Networks (CNN)

- Convolution
- Pooling
- Fully Connection (FC) (Flattening)

# CNN Architecture



# CNN Convolution Layer

Convolution is a mathematical operation to merge two sets of information

3x3 convolution

1	1	1	0	0
0	1	1	1	0
0	0	1	1	1
0	0	1	1	0
0	1	1	0	0

Input

1	0	1
0	1	0
1	0	1

Filter / Kernel

# CNN Convolution Layer

## Input x Filter --> Feature Map

receptive field: 3x3

1x1	1x0	1x1	0	0
0x0	1x1	1x0	1	0
0x1	0x0	1x1	1	1
0	0	1	1	0
0	1	1	0	0

Input x Filter

4		

Feature Map

# CNN Convolution Layer

## Input x Filter --> Feature Map

receptive field: 3x3

1	1x1	1x0	0x1	0
0	1x0	1x1	1x0	0
0	0x1	1x0	1x1	1
0	0	1	1	0
0	1	1	0	0

Input x Filter

4	3	

Feature Map



# CNN Convolution Layer

1	1	1	0	0
0	1	1	1	0
0	0	1	1	1
0	0	1	1	0
0	1	1	0	0

Input

1	0	1
0	1	0
1	0	1

Filter / Kernel


1x1	1x0	1x1	0	0
0x0	1x1	1x0	1	0
0x1	0x0	1x1	1	1
0	0	1	1	0
0	1	1	0	0

4		

Example convolution operation shown in 2D using a 3x3 filter

# CNN Convolution Layer


10 different filters 10 feature maps of size 32x32x1



final output of the convolution layer:  
a volume of size 32x32x10


# CNN Convolution Layer

## Sliding operation at 4 locations




# CNN Convolution Layer

two feature maps




# CNN Convolution Layer

**Stride** specifies how much we move the convolution filter at each step




Stride 1




Feature Map

# CNN Convolution Layer

**Stride** specifies how much we move the convolution filter at each step



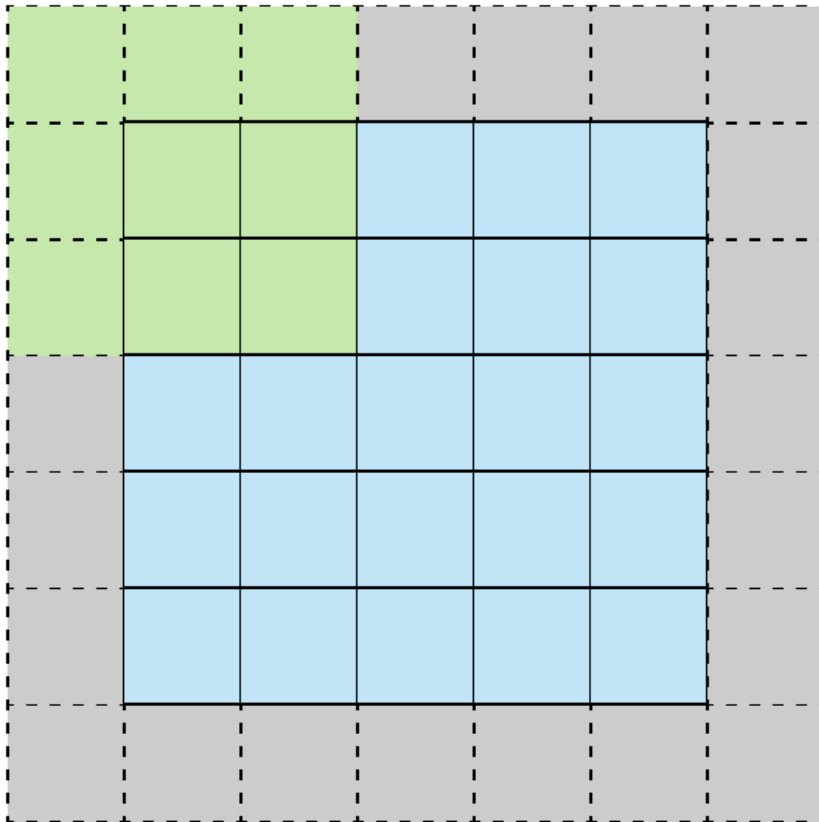
Stride 2




Feature Map

# CNN Convolution Layer

## Stride 1 with Padding



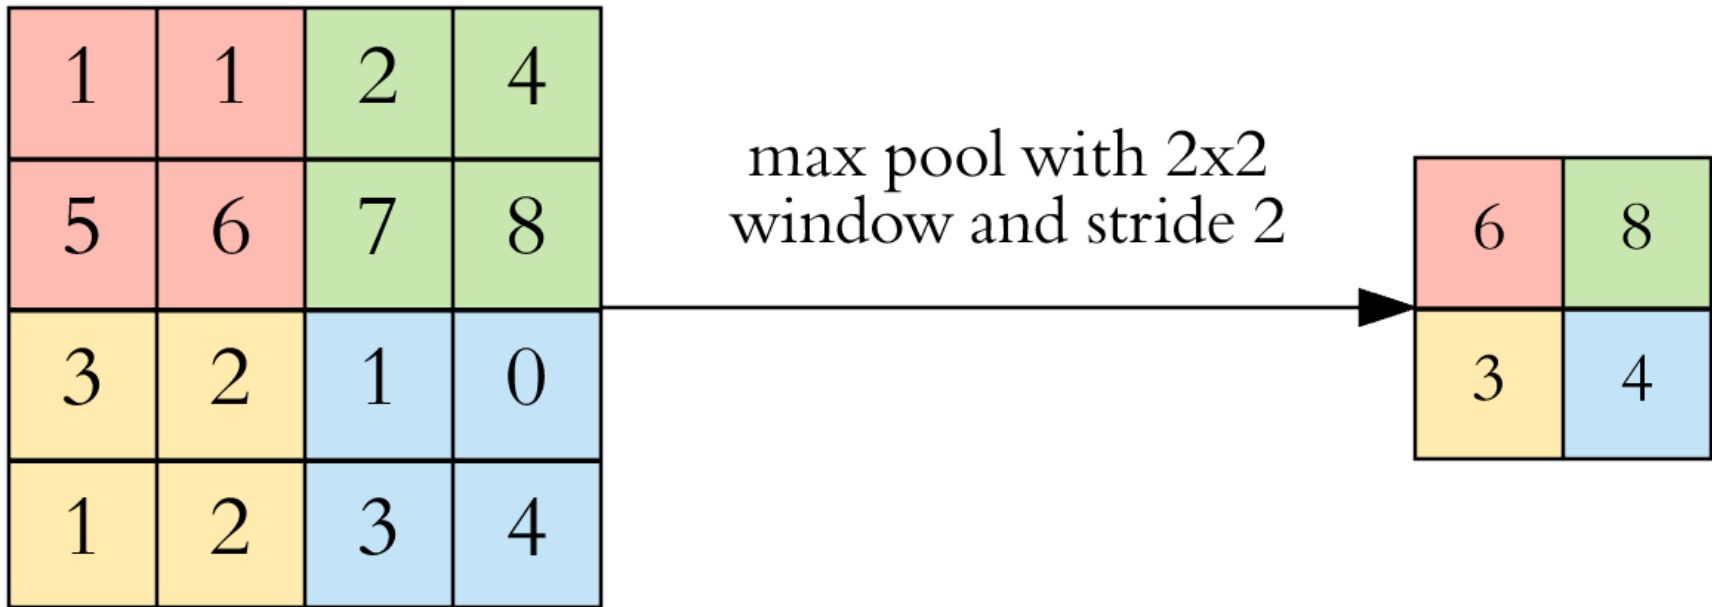
Stride 1 with Padding



Feature Map


# CNN Pooling Layer

## Max Pooling






# CNN Pooling Layer



# CNN Architecture

## 4 convolution + pooling layers, followed by 2 fully connected layers




# CNN Architecture

## 4 convolution + pooling layers, followed by 2 fully connected layers


<https://gist.github.com/ardendertat/0fc5515057c47e7386fe04e9334504e3>

```
model = Sequential()
model.add(Conv2D(32, (3, 3), activation='relu', padding='same', name='conv_1',
                input_shape=(150, 150, 3)))
model.add(MaxPooling2D((2, 2), name='maxpool_1'))
model.add(Conv2D(64, (3, 3), activation='relu', padding='same', name='conv_2'))
model.add(MaxPooling2D((2, 2), name='maxpool_2'))
model.add(Conv2D(128, (3, 3), activation='relu', padding='same', name='conv_3'))
model.add(MaxPooling2D((2, 2), name='maxpool_3'))
model.add(Conv2D(128, (3, 3), activation='relu', padding='same', name='conv_4'))
model.add(MaxPooling2D((2, 2), name='maxpool_4'))
model.add(Flatten())
model.add(Dropout(0.5))
model.add(Dense(512, activation='relu', name='dense_1'))
model.add(Dense(128, activation='relu', name='dense_2'))
model.add(Dense(1, activation='sigmoid', name='output'))
```

# Dropout





No Dropout



With Dropout

# Model Performance







# ConvNets

32x32x3 CIFAR-10 image

first Convolutional layer



# ConvNets





# Convolution Demo

Input Volume (+pad 1) (7x7x3)

$x[:, :, 0]$

0	0	0	0	0	0	0
0	1	2	0	2	1	0
0	2	2	2	1	1	0
0	2	2	2	0	1	0
0	2	2	1	2	1	0
0	2	1	2	0	1	0
0	0	0	0	0	0	0

$x[:, :, 1]$

0	0	0	0	0	0	0
0	0	2	2	1	2	0
0	1	2	0	0	2	0
0	0	1	2	1	0	0
0	2	2	2	2	0	0
0	2	2	2	0	2	0
0	0	0	0	0	0	0

$x[:, :, 2]$

0	0	0	0	0	0	0
0	1	0	0	1	0	0
0	0	2	0	0	0	0
0	0	0	1	1	1	0
0	2	2	2	1	2	0
0	1	2	0	0	2	0
0	0	0	0	0	0	0

Filter W0 (3x3x3)

$w0[:, :, 0]$

-1	-1	0
1	1	1
-1	0	1

$w0[:, :, 1]$

0	0	1
0	1	0
0	0	1

$w0[:, :, 2]$

-1	-1	0
1	0	-1
-1	0	-1

Bias  $b0$  (1x1x1)

$b0[:, :, 0]$

1
---

Filter W1 (3x3x3)

$w1[:, :, 0]$

1	-1	0
0	1	1
0	-1	1

$w1[:, :, 1]$

-1	1	0
-1	-1	1
0	0	0

$w1[:, :, 2]$

1	0	-1
0	0	-1
1	0	1

Bias  $b1$  (1x1x1)

$b1[:, :, 0]$

0
---

Output Volume (3x3x2)

$o[:, :, 0]$

6	3	6
7	-1	-2
2	3	-2


$o[:, :, 1]$

7	-1	-3
4	3	2
-1	0	-1

toggle movement

# ConvNets

input volume of size  $[224 \times 224 \times 64]$   
is pooled with **filter** size 2, **stride** 2  
into output volume of size  $[112 \times 112 \times 64]$



# ConvNets

## max pooling

Single depth slice

x

1	1	2	4
5	6	7	8
3	2	1	0
1	2	3	4


y

max pool with 2x2 filters  
and stride 2




6	8
3	4

# Convolutional Neural Networks (CNN) (LeNet)




# Recurrent Neural Networks (RNN)

# Recurrent Neural Networks (RNN)

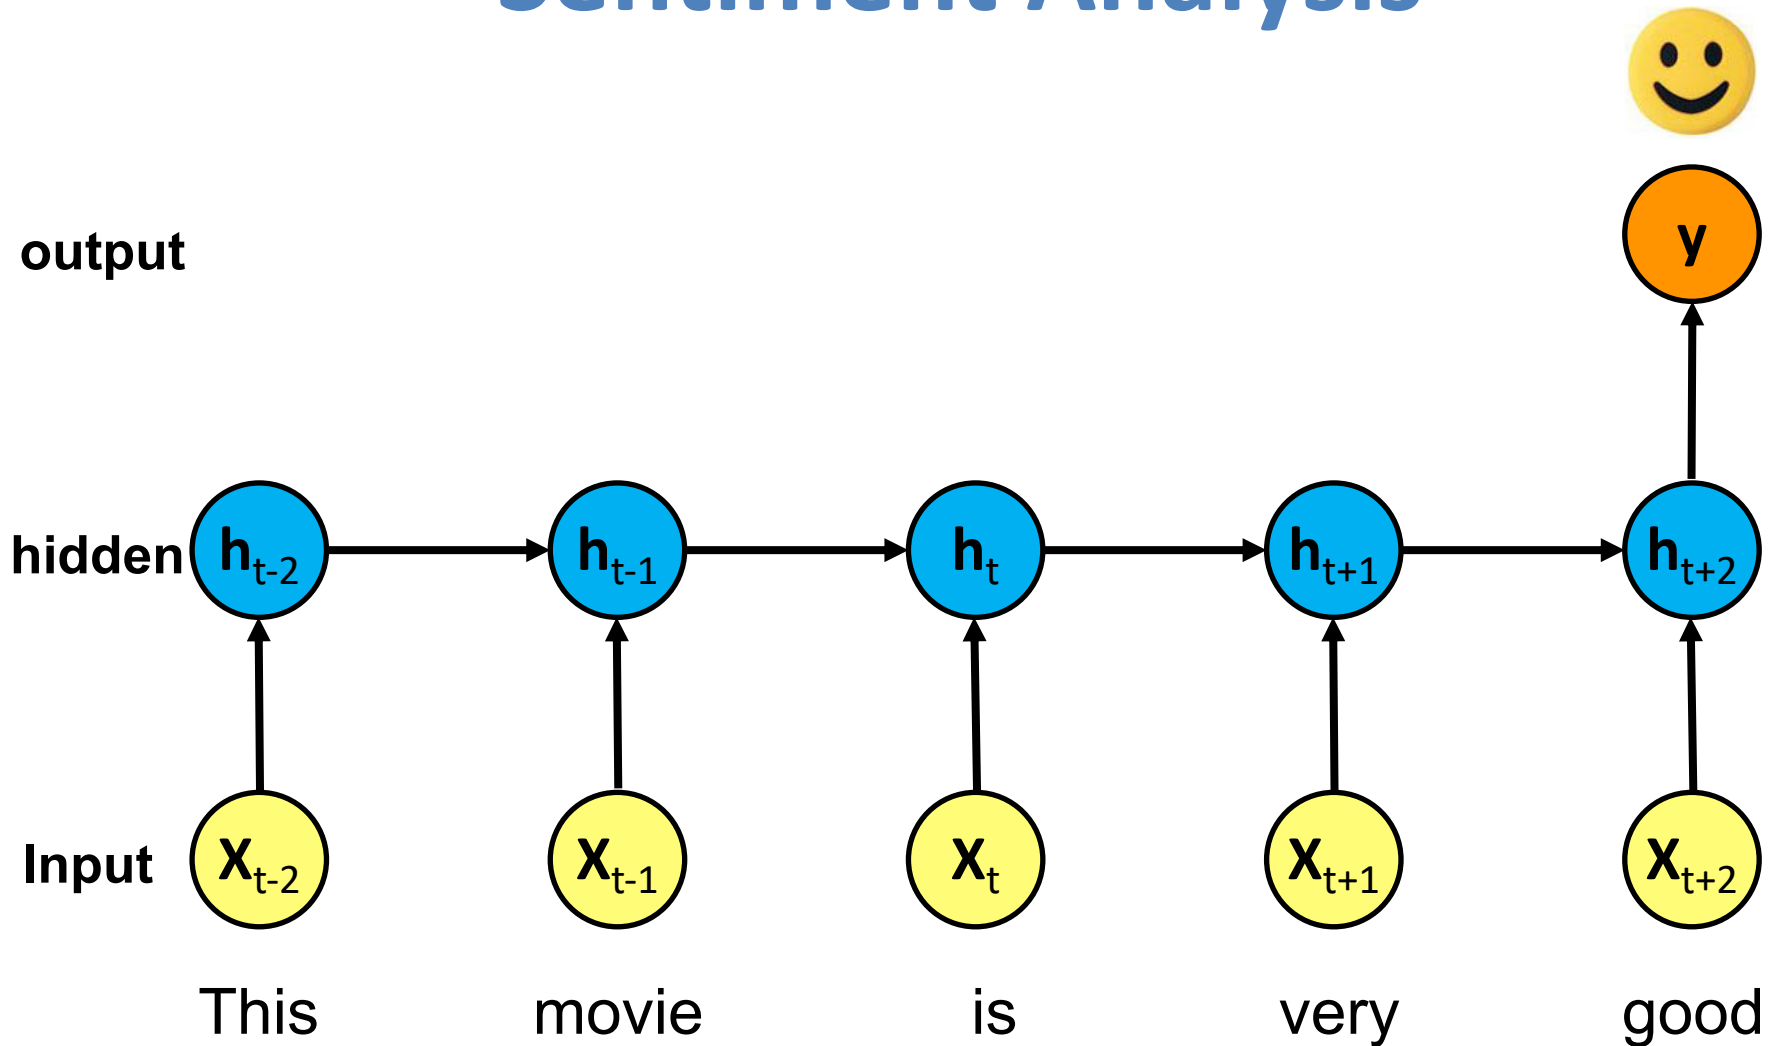


# Recurrent Neural Networks (RNN)



# Recurrent Neural Networks (RNN)

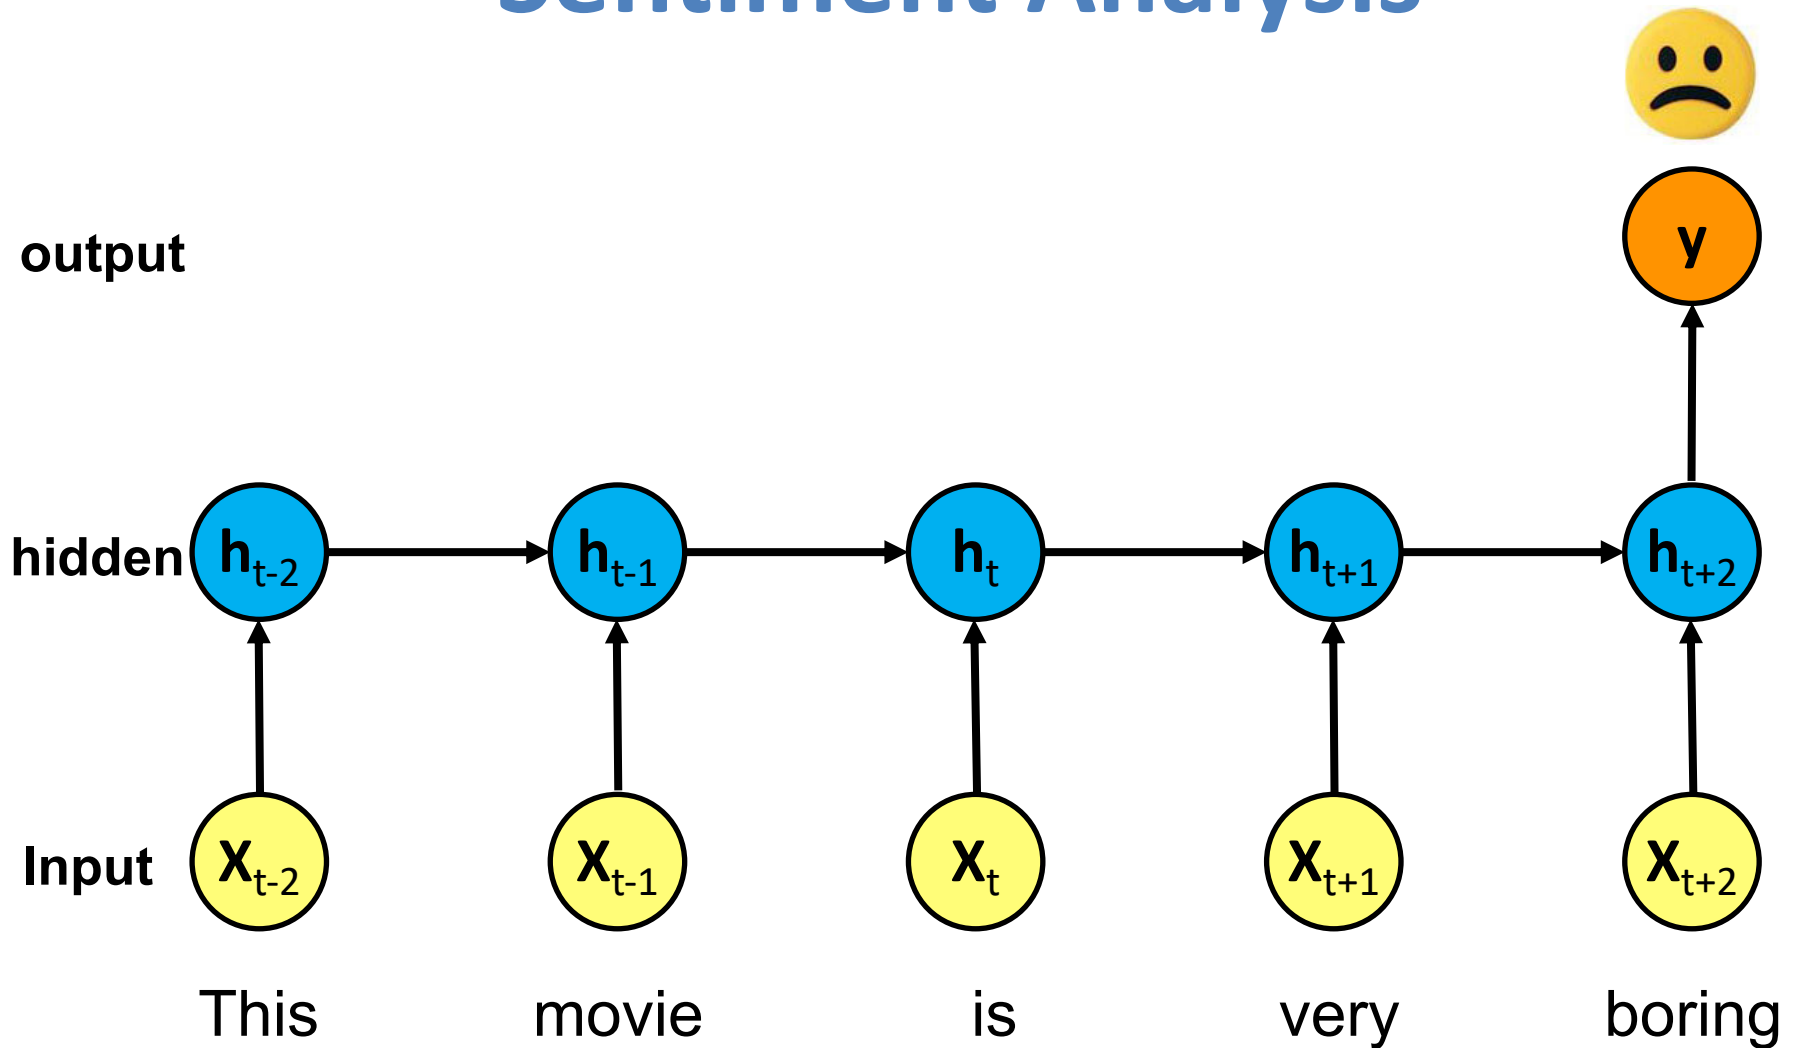
## Sentiment Analysis






# Recurrent Neural Networks (RNN)


## Sentiment Analysis





# Recurrent Neural Network (RNN)



# RNN




# RNN long-term dependencies



I grew up in France... I speak fluent French.

# Vanishing Gradient

# Exploding Gradient

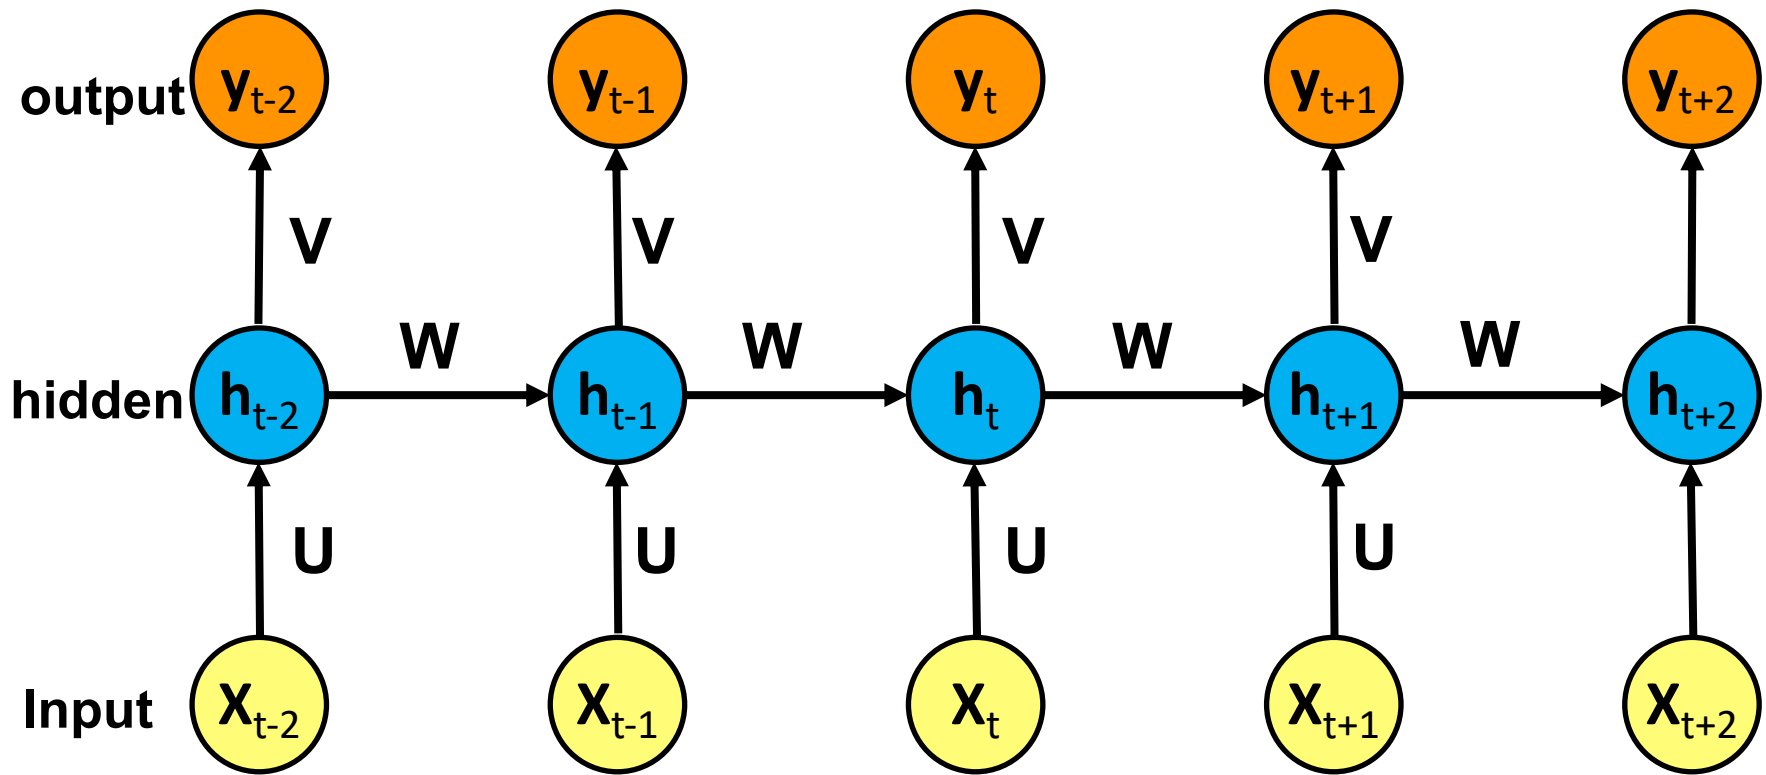


Vanishing Gradient



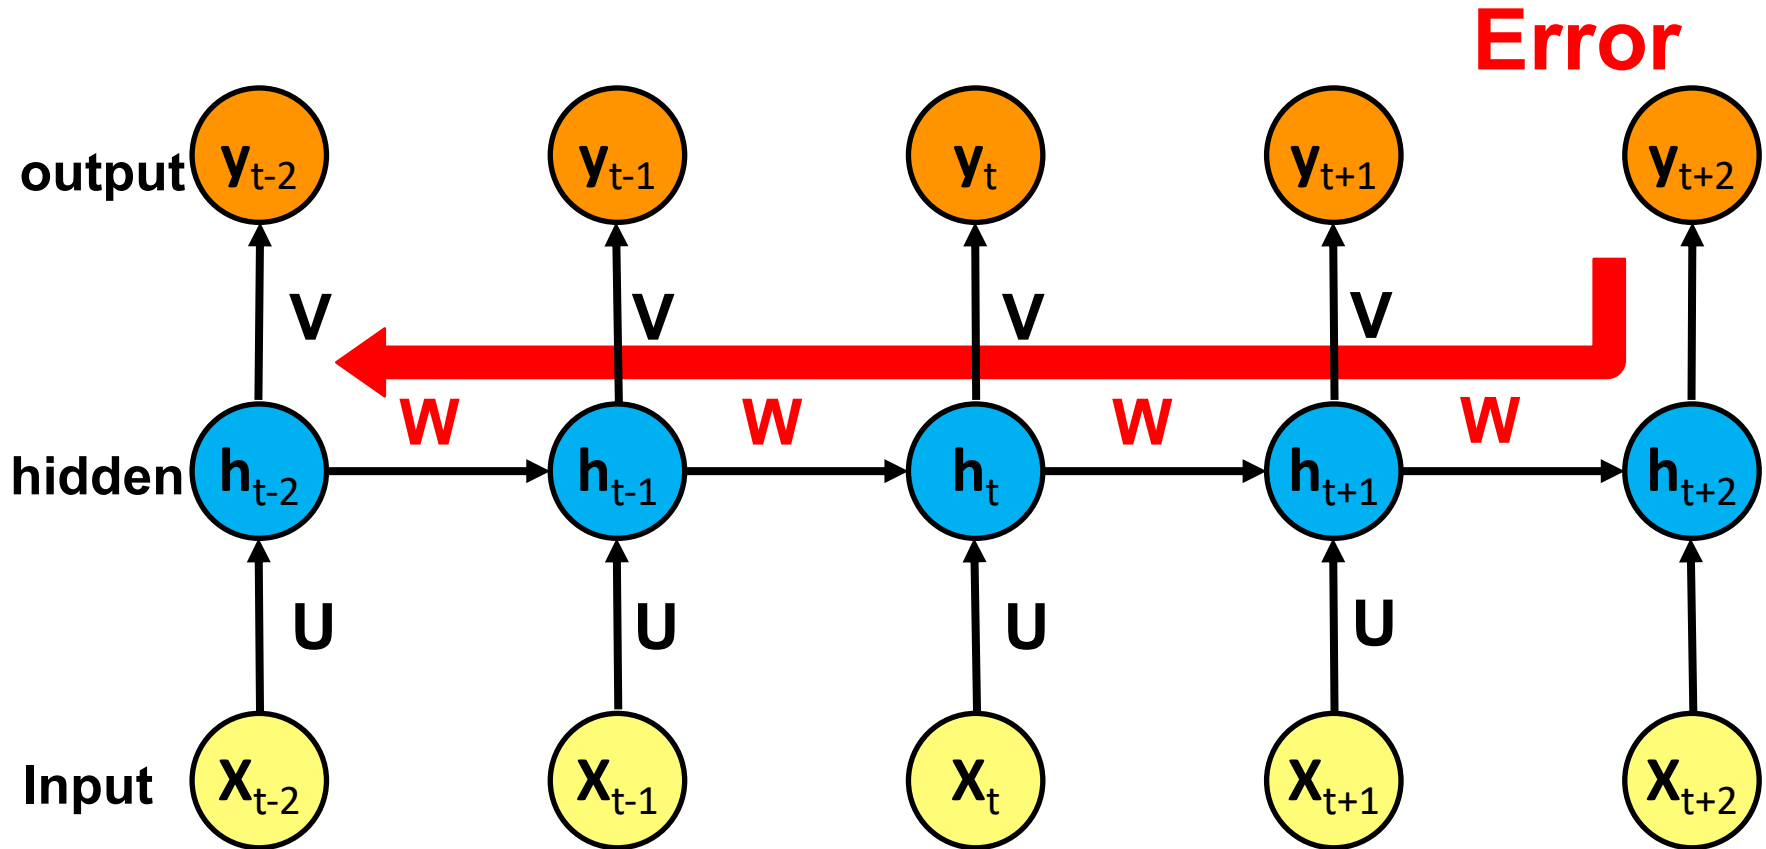
Exploding Gradient

# Recurrent Neural Networks (RNN)



# RNN


## Vanishing Gradient problem Exploding Gradient problem



if  $|W| < 1$  (Vanishing)  
if  $|W| > 1$  (Exploding)

# RNN


## Vanishing Gradient problem






# RNN

## Exploding Gradient problem




# RNN LSTM


RNN




LSTM




# Long Short Term Memory (LSTM)




# Long Short Term Memory (LSTM)




# Gated Recurrent Unit (GRU)



# Gated Recurrent Unit (GRU)



# LSTM



## Inputs:

- $X_t$  Input vector
- $c_{t-1}$  Memory from previous block
- $h_{t-1}$  Output of previous block

## Outputs:

- $C_t$  Memory from current block
- $h_t$  Output of current block

## Nonlinearities:


- $\sigma$  Sigmoid
- $\tanh$  Hyperbolic tangent

Bias: 0

## Vector operations:


- $\times$  Element-wise multiplication
- $+$  Element-wise Summation / Concatenation

# LSTM vs GRU



## LSTM

$i$ ,  $f$  and  $o$  are the **input**, **forget** and **output** gates, respectively.  
 $c$  and  $\tilde{c}$  denote the memory cell and the new memory cell content.




## GRU


$r$  and  $z$  are the **reset** and **update** gates, and  $h$  and  $\tilde{h}$  are the activation and the candidate activation.



# Long Short Term Memory (LSTM)




# Long Short Term Memory (LSTM)




# LSTM

## Memory state (C)



# LSTM


## *forget gate (f)*



$$f_t = \sigma (W_f \cdot [h_{t-1}, x_t] + b_f)$$

# LSTM


## input gate (i)



$$i_t = \sigma(W_i \cdot [h_{t-1}, x_t] + b_i)$$
$$\tilde{C}_t = \tanh(W_C \cdot [h_{t-1}, x_t] + b_C)$$

# LSTM


## Memory state (C)



$$C_t = f_t * C_{t-1} + i_t * \tilde{C}_t$$

# LSTM

## output gate (o)




$$o_t = \sigma(W_o [h_{t-1}, x_t] + b_o)$$

$$h_t = o_t * \tanh(C_t)$$

# LSTM

*forget (f), input (i), output (o) gates*



$$f_t = \sigma (W_f \cdot [C_{t-1}, h_{t-1}, x_t] + b_f)$$


$$i_t = \sigma (W_i \cdot [C_{t-1}, h_{t-1}, x_t] + b_i)$$

$$o_t = \sigma (W_o \cdot [C_t, h_{t-1}, x_t] + b_o)$$



# Gated Recurrent Unit (GRU)

*update (z), reset (r) gates*



$$z_t = \sigma (W_z \cdot [h_{t-1}, x_t])$$


$$r_t = \sigma (W_r \cdot [h_{t-1}, x_t])$$

$$\tilde{h}_t = \tanh (W \cdot [r_t * h_{t-1}, x_t])$$

$$h_t = (1 - z_t) * h_{t-1} + z_t * \tilde{h}_t$$


# LSTM Recurrent Neural Network

one to one




**Traditional  
Neural  
Network**

one to many




**Music  
Generation**

many to one




**Sentiment  
Classification**

many to many



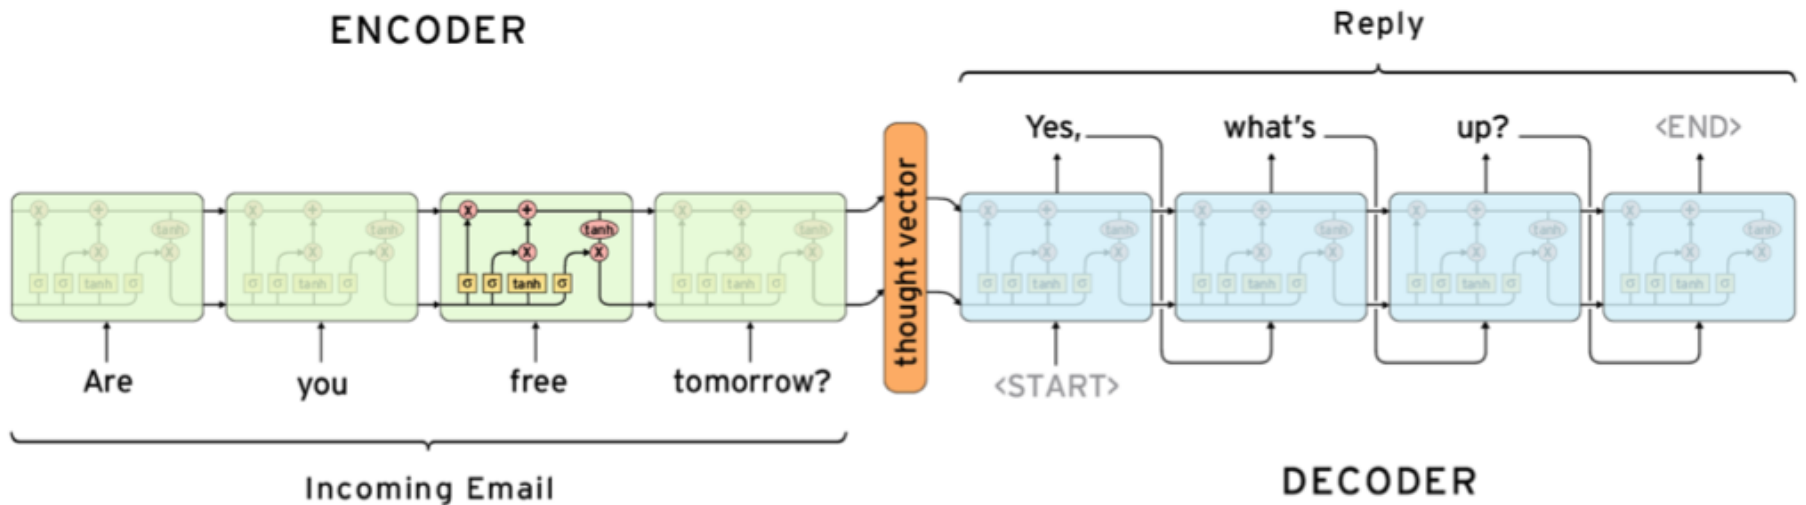
**Name  
Entity  
Recognition**

many to many

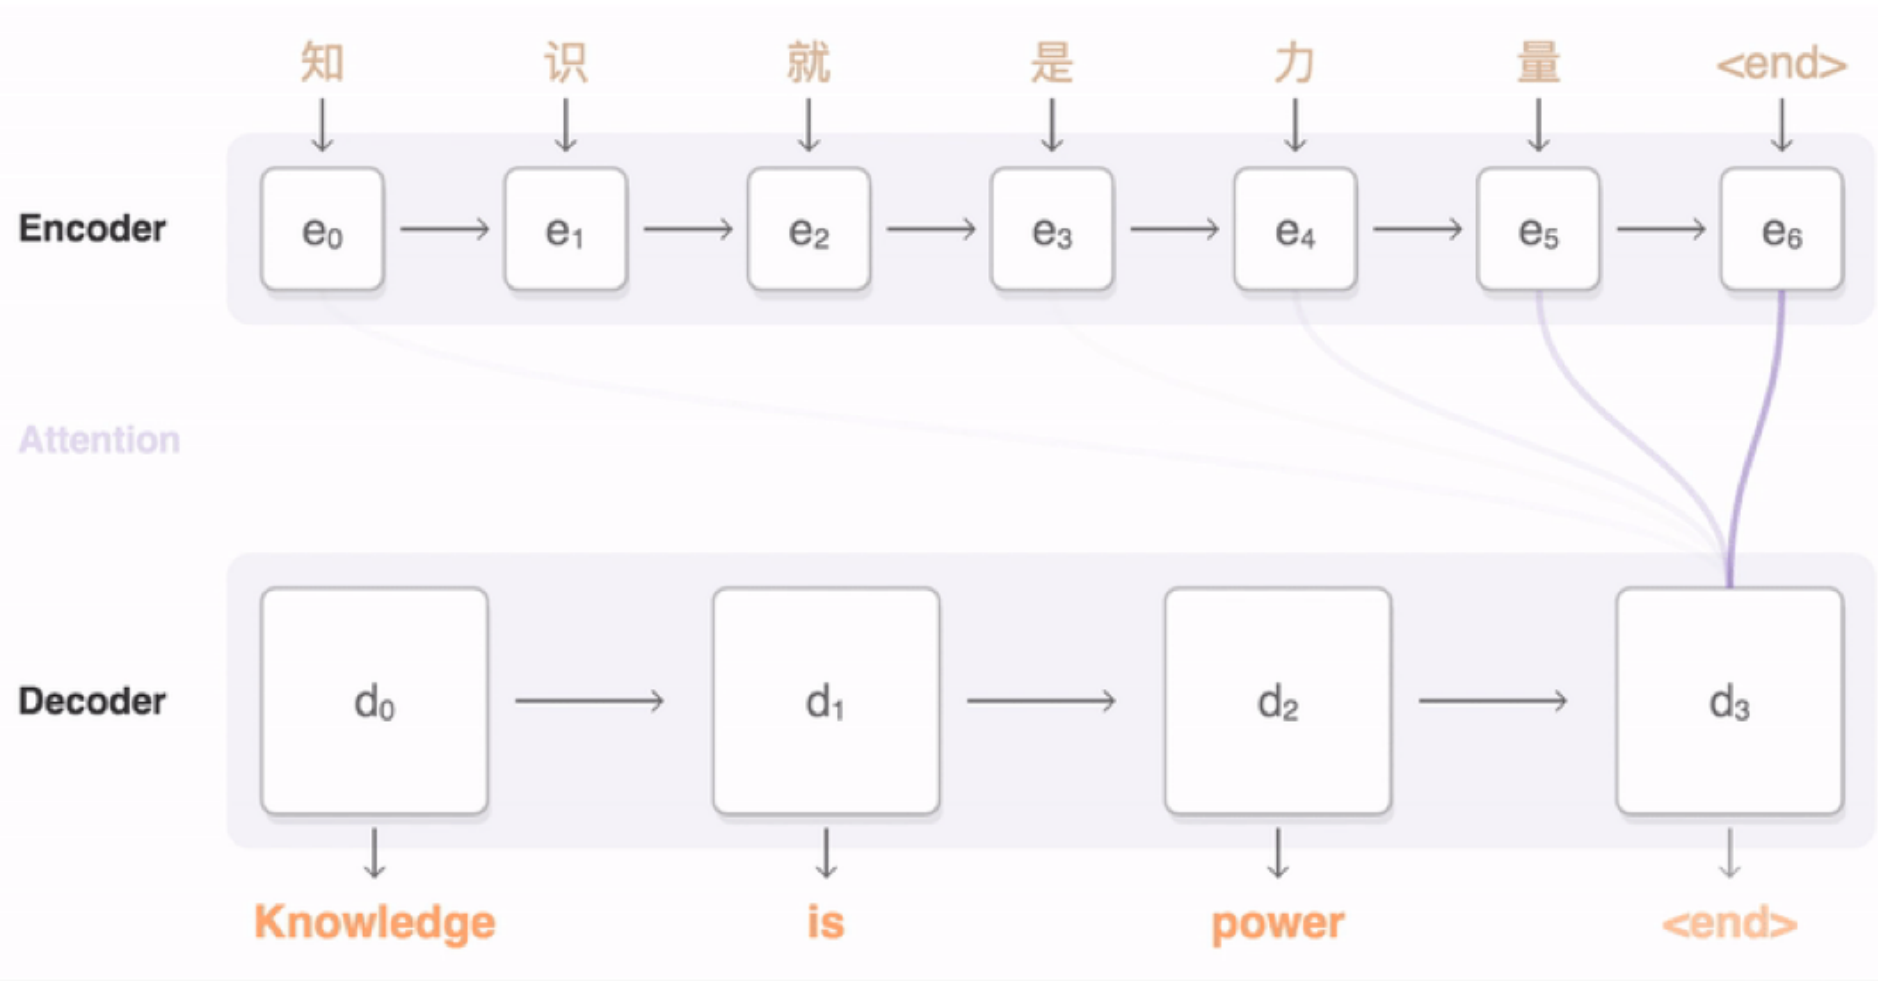


**Machine  
Translation**

# The Sequence to Sequence model (seq2seq)




# Sequence to Sequence (Seq2Seq)



# Transformer (Attention is All You Need)


(Vaswani et al., 2017)



# Transformer




# Transformer Encoder Decoder



# Transformer

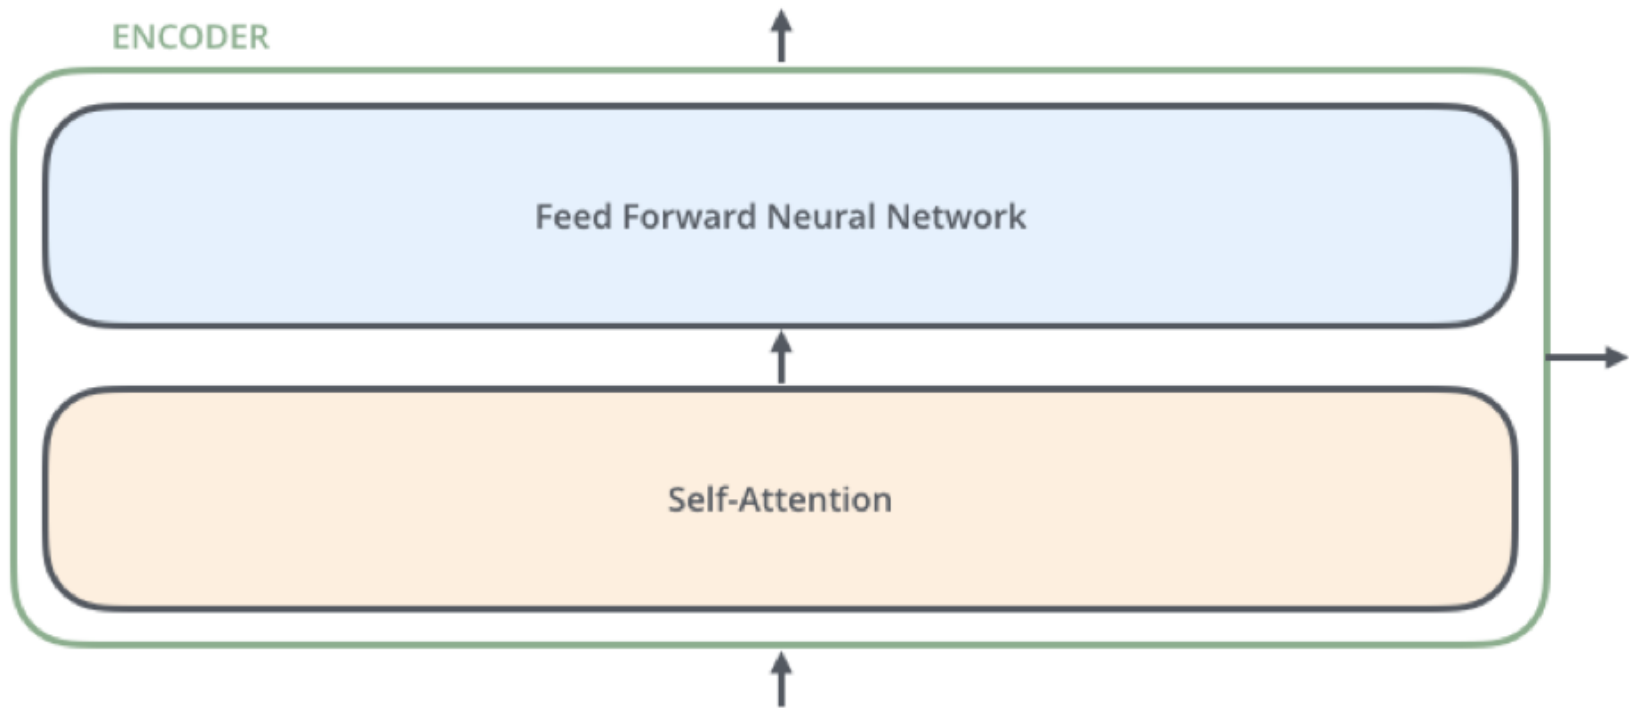
## Encoder Decoder Stack






# Transformer

## Encoder Self-Attention




# Transformer Decoder



# Transformer


## Encoder with Tensors

### Word Embeddings



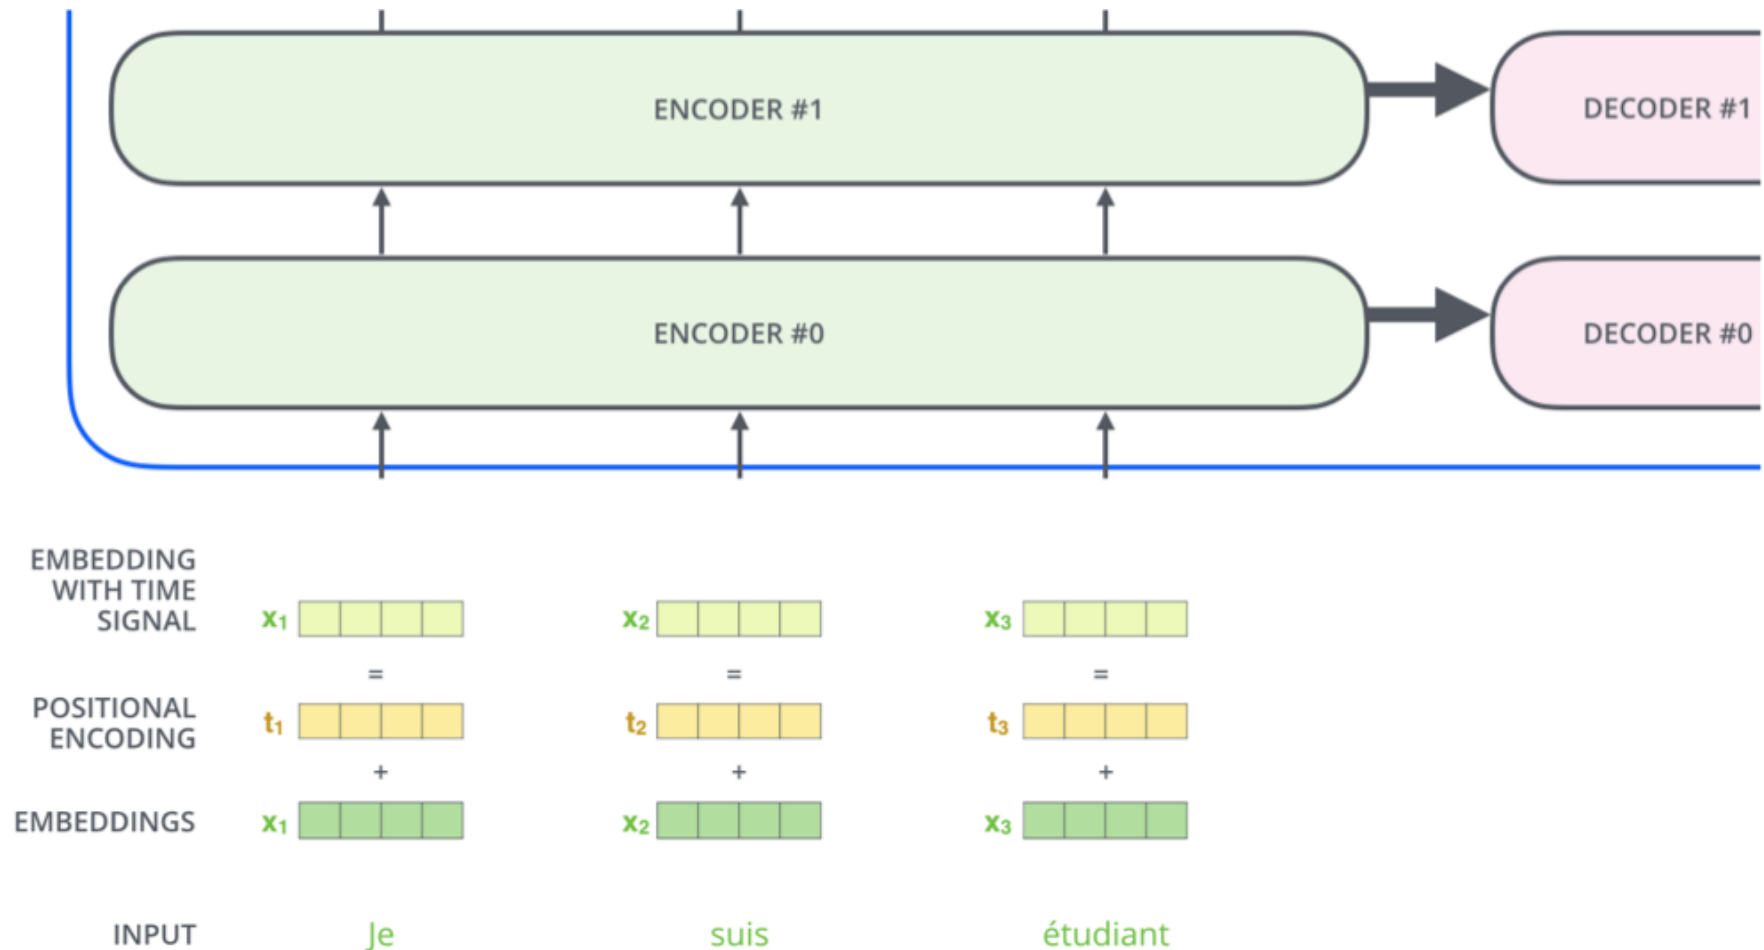
# Transformer

## Self-Attention Visualization



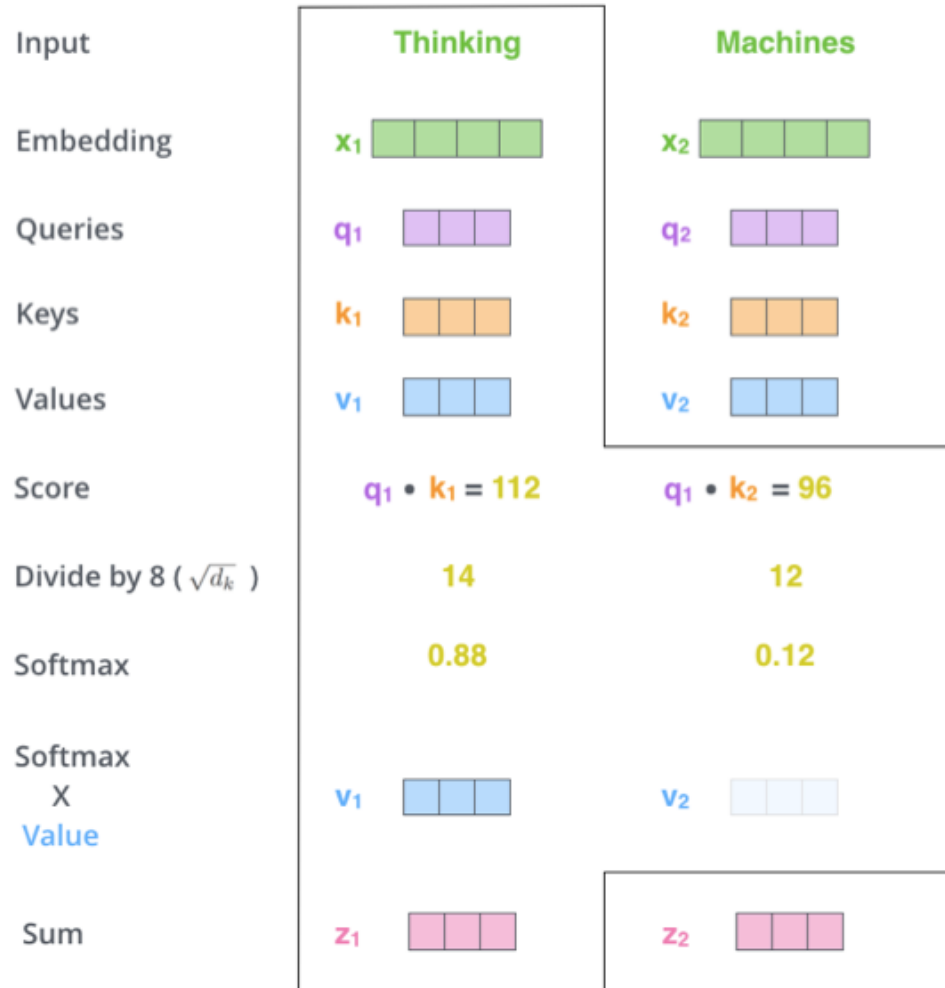
# Transformer

## Positional Encoding Vectors



# Transformer

## Self-Attention Softmax Output




# BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding

BERT

(Bidirectional Encoder Representations from Transformers)

Overall pre-training and fine-tuning procedures for BERT



Source: Devlin, Jacob, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova (2018).

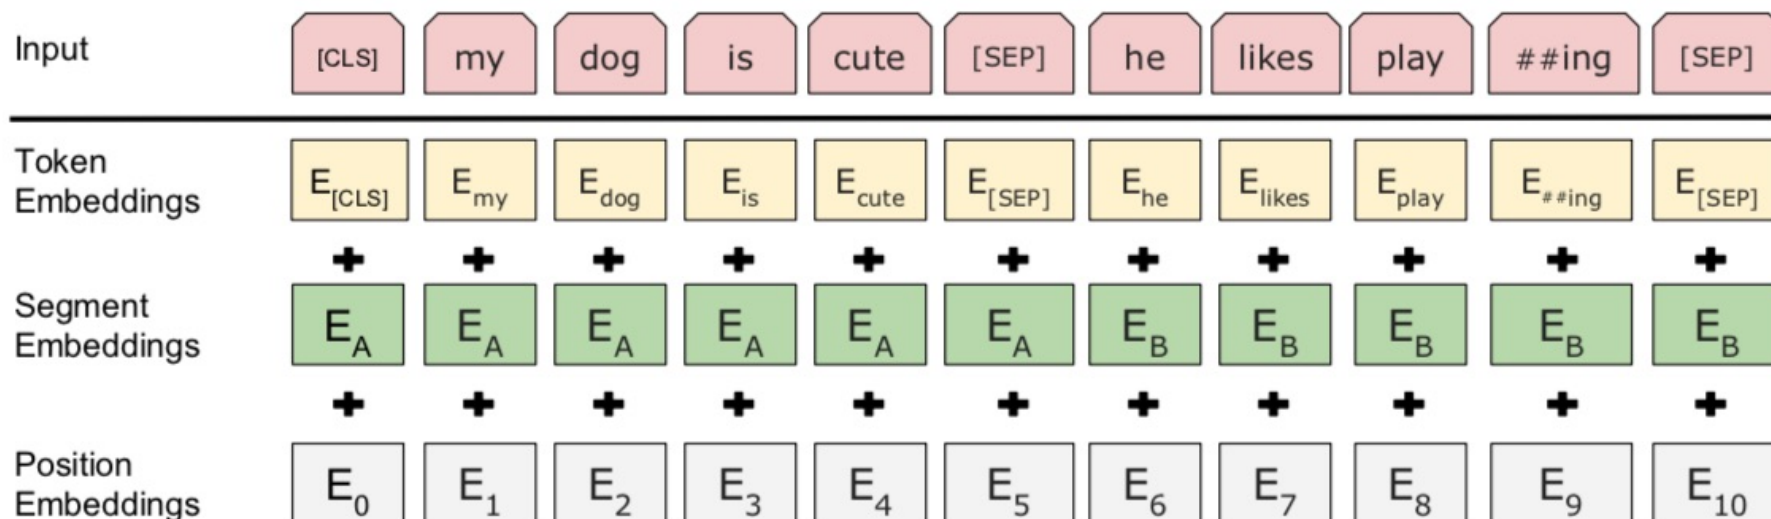
"Bert: Pre-training of deep bidirectional transformers for language understanding." arXiv preprint arXiv:1810.04805.

# BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding

## BERT

(Bidirectional Encoder Representations from Transformers)

### BERT input representation




Source: Devlin, Jacob, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova (2018).


"Bert: Pre-training of deep bidirectional transformers for language understanding." arXiv preprint arXiv:1810.04805.




# Fine-tuning BERT on Different Tasks



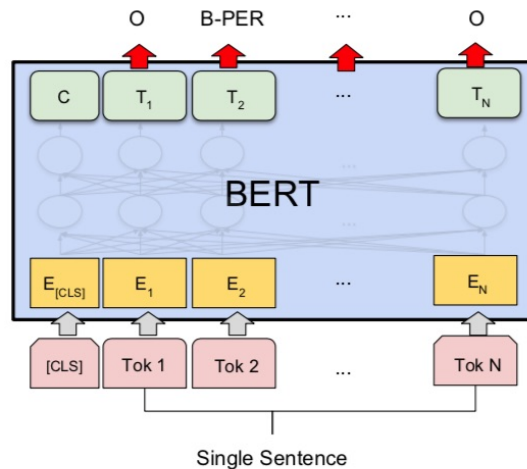
(a) Sentence Pair Classification Tasks:  
MNLI, QQP, QNLI, STS-B, MRPC,  
RTE, SWAG



(b) Single Sentence Classification Tasks:  
SST-2, CoLA



(c) Question Answering Tasks:  
SQuAD v1.1



(d) Single Sentence Tagging Tasks:  
CoNLL-2003 NER

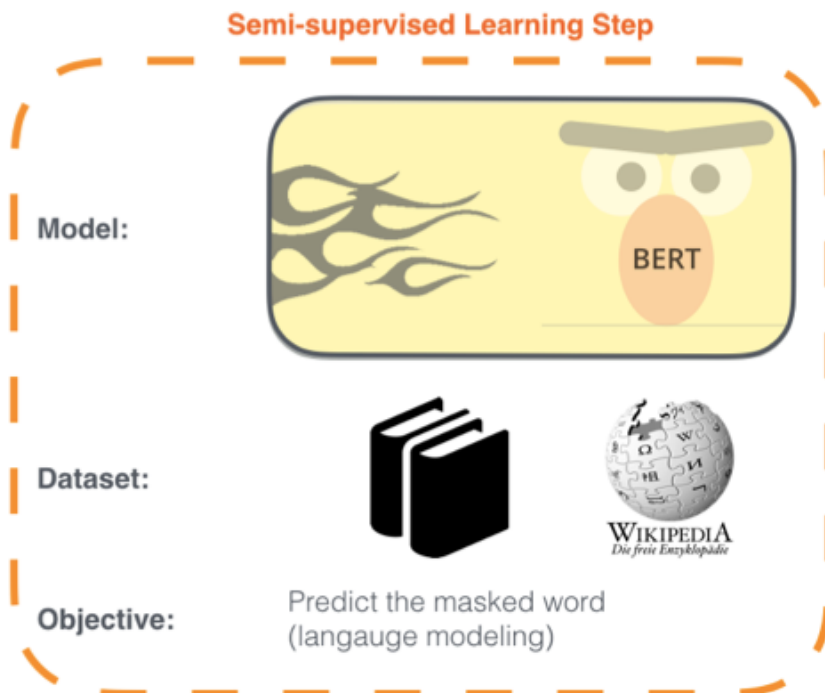
Source: Devlin, Jacob, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova (2018).

"Bert: Pre-training of deep bidirectional transformers for language understanding." arXiv preprint arXiv:1810.04805.


# Illustrated BERT

1 - **Semi-supervised** training on large amounts of text (books, wikipedia..etc).

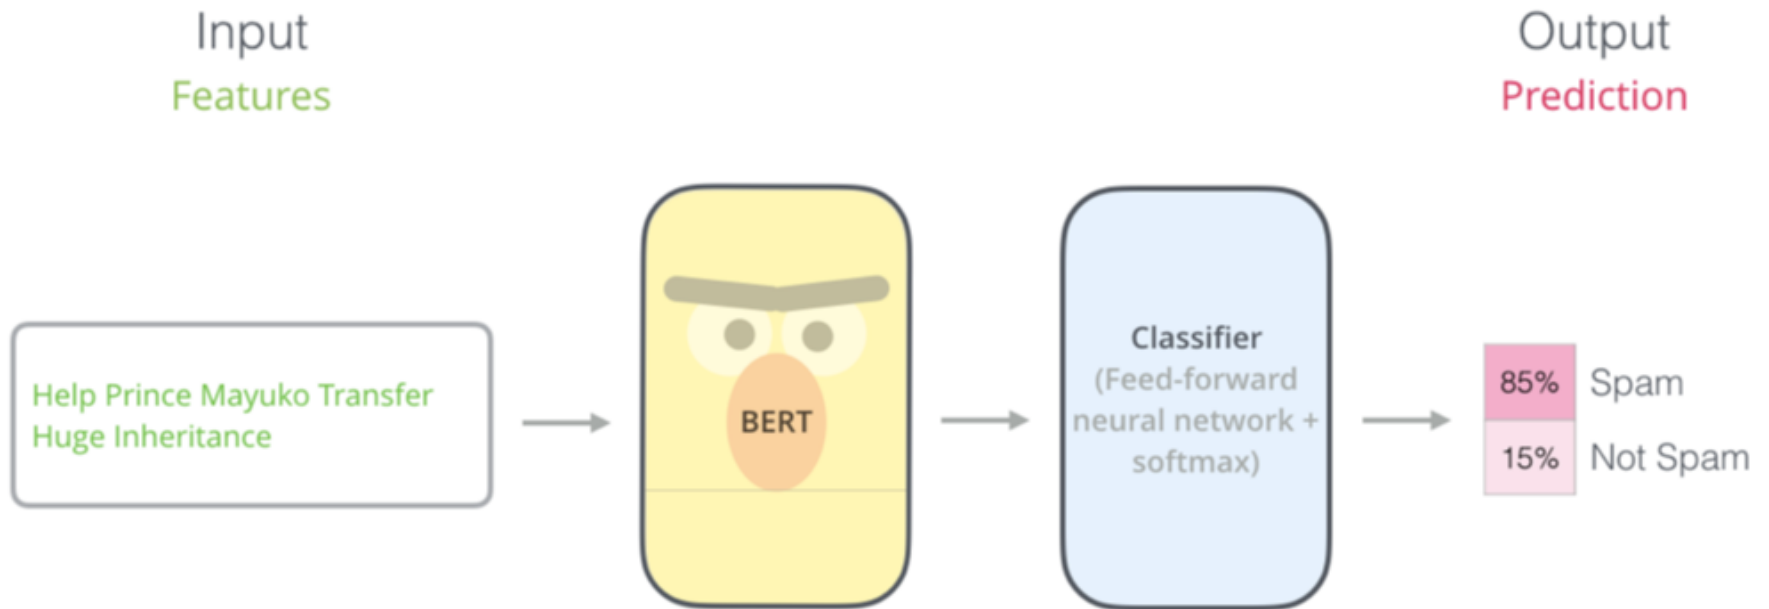
The model is trained on a certain task that enables it to grasp patterns in language. By the end of the training process, BERT has language-processing abilities capable of empowering many models we later need to build and train in a supervised way.




2 - **Supervised** training on a specific task with a labeled dataset.



# BERT Classification Input Output



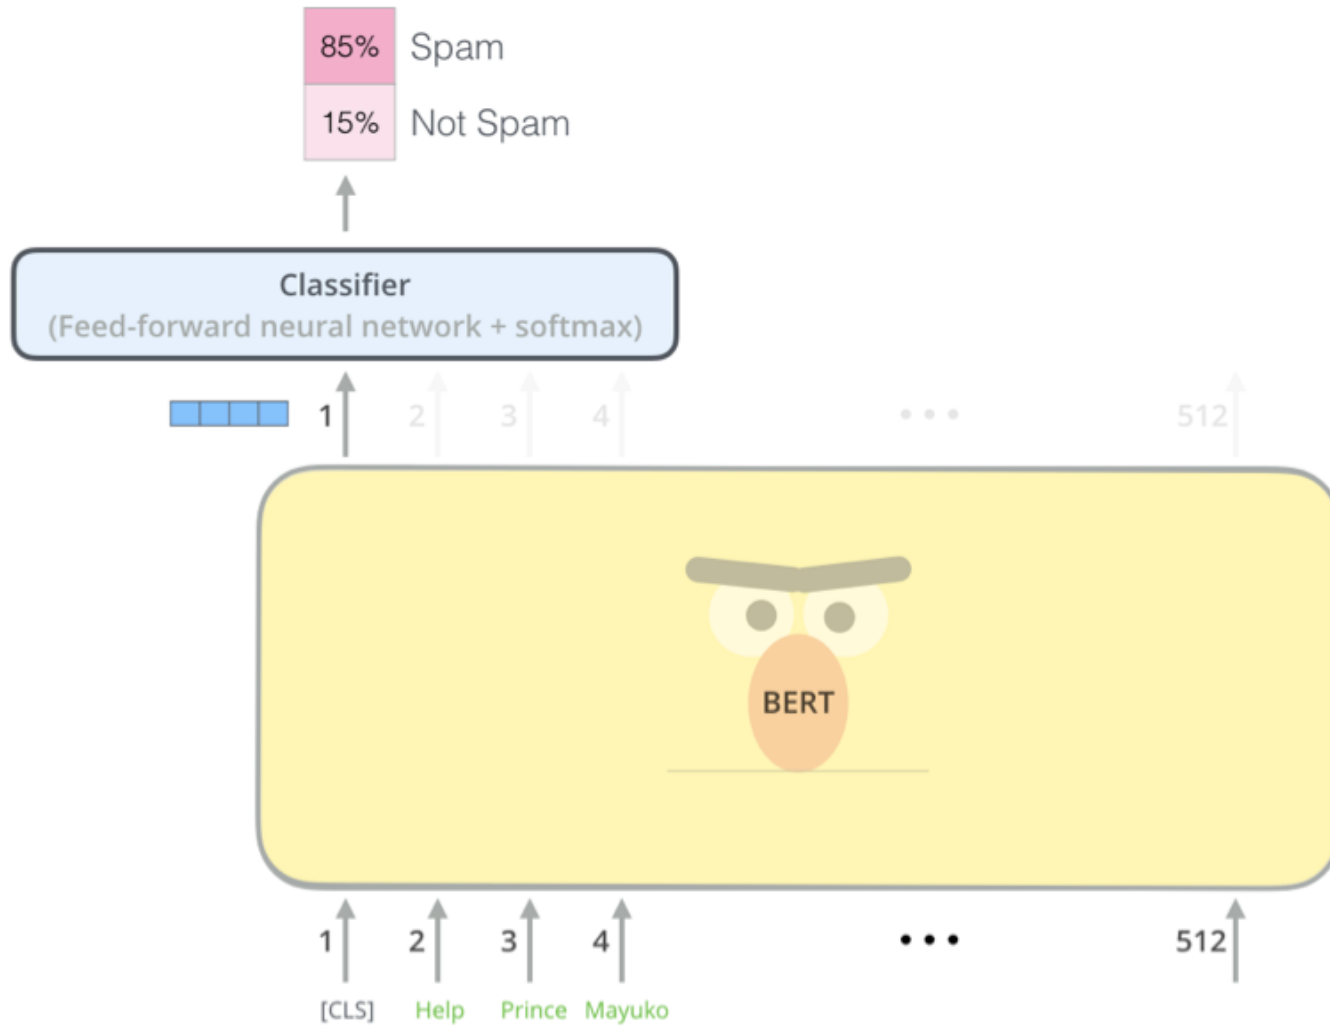
# BERT Encoder Input



BERT

Source: Jay Alamar (2019), The Illustrated BERT, ELMo, and co. (How NLP Cracked Transfer Learning), <http://jalamar.github.io/illustrated-bert/>

# BERT Classifier



Source: Jay Alammar (2019), The Illustrated BERT, ELMo, and co. (How NLP Cracked Transfer Learning), <http://jalammar.github.io/illustrated-bert/>

# Summary

- Traditional Feature Engineering for Text Data
  - Bag of Words Model
  - Bag of N-Grams Model
  - TF-IDF Model
- Advanced Word Embeddings with Deep Learning
  - Word2Vec Model
  - Robust Word2Vec Models with Gensim
  - GloVe Model
  - FastText Model

# References

- Dipanjan Sarkar (2019), Text Analytics with Python: A Practitioner's Guide to Natural Language Processing, Second Edition. APress. <https://github.com/Apress/text-analytics-w-python-2e>
- Benjamin Bengfort, Rebecca Bilbro, and Tony Ojeda (2018), Applied Text Analysis with Python, O'Reilly Media. <https://www.oreilly.com/library/view/applied-text-analysis/9781491963036/>
- Aurélien Géron (2019), Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow: Concepts, Tools, and Techniques to Build Intelligent Systems, 2nd Edition, O'Reilly Media, 2019, <https://github.com/ageron/handson-ml2>
- Devlin, Jacob, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova (2018). "Bert: Pre-training of deep bidirectional transformers for language understanding." arXiv preprint arXiv:1810.04805.
- Deep Learning Basics: Neural Networks Demystified, <https://www.youtube.com/playlist?list=PLiaHhY2iBX9hdHaRr6b7XevZtgZRa1PoU>
- Deep Learning SIMPLIFIED, <https://www.youtube.com/playlist?list=PLjJh1vISEYgvGod9wWiydumYl8hOXixNu>
- 3Blue1Brown (2017), But what \*is\* a Neural Network? | Chapter 1, deep learning, <https://www.youtube.com/watch?v=aircAruvnKk>
- 3Blue1Brown (2017), Gradient descent, how neural networks learn | Chapter 2, deep learning, <https://www.youtube.com/watch?v=IHZwWFHWa-w>
- 3Blue1Brown (2017), What is backpropagation really doing? | Chapter 3, deep learning, <https://www.youtube.com/watch?v=Ilg3gGewQ5U>
- Jay Alamar (2019), The Illustrated Transformer, <http://jalamar.github.io/illustrated-transformer/>
- Jay Alamar (2019), The Illustrated BERT, ELMo, and co. (How NLP Cracked Transfer Learning), <http://jalamar.github.io/illustrated-bert/>
- Min-Yuh Day (2020), Python 101, <https://tinyurl.com/imtkupython101>


# Q & A

# 文本表達特徵工程 (Feature Engineering for Text Representation)

Time: 2020/05/29 (Fri) (9:10 -12:00)

Place: 國立臺北護理健康大學 (台北市明德路365號) G210

Host: 祝國忠 院長 (健康科技學院院長)



Min-Yuh Day

戴敏育

Associate Professor

副教授

Dept. of Information Management, Tamkang University

淡江大學 資訊管理學系

<http://mail.tku.edu.tw/myday/>

2020-05-29

