

BUILDING AN AI/ML POWERED TEXT SEARCH SYSTEM

NICK BURCH

Berlin Buzzwords 2019

NICK BURCH

CTO, Head of AI Development



@Gagravarr

SOURCE CODE

[HTTPS://NOTEBOOKS.AZURE.COM/G
AGRAVARR/PROJECTS/BBUZZ-2019-
ML-SEARCH](https://notebooks.azure.com/GAGRAVARR/projects/BBUZZ-2019-ML-SEARCH)



OUR TALK TODAY

- Not your typical AI talk...
- The problem to solve
- AI and text - how?
- "Simple" AI for text
- Making it harder
- Neural Networks for text
- Making it better
- Further resources

WHAT IS AI?

AND ML?

AND WHY NOW?

AI - ARTIFICIAL INTELLIGENCE

ML - MACHINE LEARNING

(Larry) Tesler's Theorem - "AI is whatever hasn't been done yet."

Which makes... ML what we can do today!

First big "AI Bubble" was mid-1980s, over \$1 billion in AI startups that were generally touted as "expert systems"

Two problems - not enough training data available, computers much more expensive than the experts they were trying to replace

Moore's Law to the rescue! One million dollars worth of 1985 computing power costs under one dollar today.

Amazon will rent you a machine with 1T of memory for roughly the cost of a latte per hour.

A 4T memory machine is about \$25 / hour, less if reserved!

We have a lot more data available to train our ML models on.


Open Source libraries and frameworks for AI / ML make it easy to get started, mean you can focus on your problem, not the system.

THIS IS YOUR MACHINE LEARNING SYSTEM?

YUP! YOU POUR THE DATA INTO THIS BIG PILE OF LINEAR ALGEBRA, THEN COLLECT THE ANSWERS ON THE OTHER SIDE.

WHAT IF THE ANSWERS ARE WRONG?

JUST STIR THE PILE UNTIL THEY START LOOKING RIGHT.



YOUR TYPICAL AI / ML DEMO

IMAGE CLASSIFICATION!

It's fun, and it's easy!

DOGS OR CATS? DOGS



DOGS OR CATS? KOALA



WHAT ANIMAL? KANGAROO



WHAT ANIMAL? WOMBAT



WHAT ANIMAL? ECHIDNA



IMAGE CLASSIFICATION - COOL!

(Assuming your humans know what everything is, and can correctly give the right labels to your training data)

BUT MY DATA DOESN'T LOOK LIKE THAT....

1 Table of Contents

1	Table of Contents	1
2	Approvals	1
3	Relevance	2
4	Purpose	2
5	Scope	2
6	Policy	2
7	Glossary of Terms	3
7.1	Abbreviations	3
7.2	Definitions	3
8	Procedure	5
8.1	Customer Development Cycle	5

SAS Universal Viewer

File Tools Window Help

Address

c:\users\nickb\desktop\test-columnar.sas7bdat

Library TEST-COLUMNAR

Freeze Hide Show... Format Filter... Font... Find

Table View

	A	B	C	D	E
1	Num=0	0	30DEC99	01JAN00:11:00:00	
2	Num=0.1	0.1	30DEC99	30DEC99:02:24:00	0.1
3	Num=0.25	0.25	30DEC99	30DEC99:06:00:00	0.25
4	Num=0.5	0.5	30DEC99	30DEC99:12:00:00	0.5
5	Num=1	1	01JAN00	01JAN00:00:00:00	
6	Num=1.1	1.1	01JAN00	01JAN00:02:24:00	1.1
7	Num=1.2	1.2	01JAN00	01JAN00:04:48:00	1.2
8	Num=1.5	1.5	01JAN00	01JAN00:12:00:00	1.5
9	Num=2	2	02JAN00	02JAN00:00:00:00	2
10	Num=2.5	2.5	02JAN00	02JAN00:12:00:00	2.5
11	Num=3	3	03JAN00	03JAN00:00:00:00	3
12	Num=4	4	04JAN00	04JAN00:00:00:00	4
13	Num=5	5	05JAN00	05JAN00:00:00:00	5

SECTION 5: Business Impact Analysis (BIA)	
Business Critical?	<u>[Y/N]</u>
Data Stored Within System?	<u>[Y/N]</u>
Probability of Failure:	<u>[1 - 5]</u> (1 is minimum, 5 is maximum)
Impact of Failure:	<u>[1 - 5]</u> (1 is minimum, 5 is maximum)
Overall BIA Rating:	<u>[1 - 25]</u> (Product of Impact and Probability)

IF YOU WANT AN IMAGE CLASSIFICATION TALK

Large Scale Landuse Classification of Satellite Imagery

<https://berlinbuzzwords.de/18/session/large-scale-landuse-classification-satellite-imagery>

IF YOU WANT A NUMERIC MODELLING / FITTING TALK

Real World AirBnB Data Science and Pricing Bot

<https://berlinbuzzwords.de/17/session/weekend-project-real-world-airbnb-data-science-and-pricing-bot>

STILL HERE? LET'S TALK TEXT!

MY PROBLEM - I HAVE LOTS OF DOCUMENTS

Including Policies, Procedures, Training Guides, Help Information
Including Questionnaires, Request For Proposals

(I have data too, but that's out-of-scope for today!)

**I NEED TO DO INEXACT SEARCHING OVER THE
CONTENTS OF ALL THESE DOCUMENTS**

But I'm not allowed to show you most of those documents....

BERLIN BUZZWORDS TALKS TO THE RESCUE!

Let's use past Talk Titles and Abstracts as our test data

Building an AI/ML powered text search system

Search

Nick Burch

06/18/2019 - 12:20 to 13:00

Frannz Salon

long talk (40 min)

Beginner

Session abstract:

There are some great Open Source text search / information retrieval systems, such as Apache SOLR and

Large Scale Landuse Classification of Satellite Imagery

Scale

Suneel Marthi

06/11/2018 - 14:00 to 14:40

Moon Lounge

long talk (40 min)

Intermediate

Session abstract:

With the abundance of Remote Sensing satellite imagery, the possibilities are endless as to the kind of insights

Partly clustering - what talks are similar to what other talks? What words are similar to other words?

Partly recommending - people using these search terms also found these talks relevant

It isn't - exact matching

It doesn't have - classification labels, certain answers

It can cope with - people who gave their talks fun / cool titles!

AI / ML FRAMEWORKS DON'T LIKE WORD DOCUMENTS / SPREADSHEETS / RAW HTML

[Apache Tika](#) lets us turn all of these documents into clean, semantically meaningful HTML

Then split the HTML into chunks (eg document sections, sheets etc), and finally generate something like JSON

For BBuzz talks, used BeautifulSoup to process talk pages into JSON

SAMPLE DATA

```
{
  "level": "Beginner",
  "track": "Scale",
  "abstract": "Whether we're in the position of a team lead.....",
  "title": "Building and Scaling a High Performing Development Team by Finding the Facts and Avoid",
  "url": "https://berlinbuzzwords.de/19/session/building-and-scaling-high-performing-development-t",
  "speaker": "Will Hayes"
},
{
  "level": "Intermediate",
  "track": "Search",
  "abstract": "Search is fundamental feature of mobile.de platform and we as Data Team work hard to",
  "title": "Architecture of relevancy search at mobile.de",
  "url": "https://berlinbuzzwords.de/19/session/architecture-relevancy-search-mobilede",
  "speaker": "Richard Knox"
},
{
  "level": "Beginner",
  "track": "Scale",
  "abstract": "Non-code contributions, like project and community management, are essential to the",
  "title": "Non-Code contributions: The hidden gem in Open Source Projects",
  "url": "https://berlinbuzzwords.de/19/session/non-code-contributions-hidden-gem-open-source-proj",
  "speaker": "Griselda Cuevas"
},
```

WE HAVE OUR TALKS AS JSON

So, we're set right? Just feed the JSON into the AI, and magic happens?

Sadly not... ML frameworks don't generally work on Text-in-JSON

ML NEEDS A MATRIX OF NUMERIC VALUES

Ideally mostly -1.0 to 1.0 or 0.0 to 1.0

One value for each *feature* of the thing to learn / predict on

Can be sparse (only a few non-zero values) or dense (mostly non-zero)

More features requires more memory and more CPU, so a tradeoff!

How can we turn our text into something like that?

AN ASIDE - SOME TERMINOLOGY

Feature - An input variable, a value for one aspect of the thing to predict. eg
height / weight / 1st RGB channel

Label - The value to be predicted / trained for, eg Dog / Cat / price of house the
features describe

Training - Creating / learning a model to map from the features to the label

Inference - Applying the model to something new to get a prediction

<https://developers.google.com/machine-learning/crash-course/framing/ml-terminology>

AN ASIDE - SOME TERMINOLOGY

Regression - A model to predict continuous values. eg "Given the location and number of bedrooms, what's the likely price of a house"

Classification - A model to predict discrete values. eg "Is this an image of a Dog, a Cat or a Wombat?"

Clustering - A model to group similar things together. If those are labelled, it's classification. If those aren't labelled, it's clustering, and relies on unsupervised machine learning

<https://developers.google.com/machine-learning/clustering/overview>

VECTOR SPACE FOR TEXT

First up, we need to break our text down into smaller units. We call this *tokenisation*.

The simplest way to do that is to split on whitespace and punctuation, so we get one token per word.

(We'll cover more advanced things, eg stopwords and stemming, a bit later on)

See any Lucene introduction talk for more!

Lucene in Action or Taming Text have good stuff on this in early chapters too.

VECTOR SPACE FOR TEXT

Next, build up a *term dictionary* for all the different tokens in our text, assigning each a unique index

The mouse ran up the clock

The mouse ran down

```
{ 'the':1, 'mouse':2, 'ran':3, 'up':4, 'clock':5,  
  'down':6 }
```

The mouse ran up the clock = [1, 2, 3, 4, 5]

The mouse ran down = [1, 2, 3, 6]

VECTOR SPACE FOR TEXT

When one word occurs multiple times, what then? Two easy options

One-Hot encoding - 1 if present, otherwise 0

CBOW Count - how many times across the *continuous bag of words* each term occurs. (This loses position information though)

The mouse ran up the clock = [2, 1, 1, 1, 1, 0]

The mouse ran down = [1, 1, 1, 0, 0, 1]

TF-IDF

If a document contains a given term a lot, we want to rate that document higher for that term.

If most documents have a term, it's probably a less interesting thing to look for, and not that specific. If only a few do, that term is probably more interesting.

If a document is very long, its counts will naturally be higher than a very short document's counts, as it'll have more text.

Weigh rare terms higher, common terms lower, across all documents. Within a document, rate the term higher the more it is used. Weight shorter documents higher than longer ones.

TF-IDF = TERM FREQUENCY – INVERSE DOCUMENT FREQUENCY

TF-IDF is the simplest common way to do this, implemented in most ML libraries you'll come across.

For each talk, tokenise then do CBOW count, then apply TF-IDF. (Generally 1-2 lines of python code!). Gives us, for each talk, a value between 0 and 1 for each term. A big ML friendly matrix!

Other relevancy techniques exist, for more details see <https://berlinbuzzwords.de/16/session/bm25-demystified> and <https://2017.berlinbuzzwords.de/17/session/bm25-so-yesterday-modern-techniques-better-search-relevance>

OUR NEXT CHALLENGE - HOW TO TRAIN THE AI?

We don't know what the right answer is for which talk(s) go with which query.

We don't have the training data, the labels or the scoring function.

For many techniques, we would need to gather that data first! Perhaps by watching what users did, or sitting down and manually classifying loads of things.

However, if we just did things on text similarity, would that get us close enough for now?

FIRST APPROACH - PREDICT ONE TALK, SUGGEST SIMILAR

1. Build a classification model based on talks
2. Train model on text from talk (title, abstract), using TF-IDF to make feature vector for each talk
3. Ask model to classify our query
4. Result is talk "most like" our query
5. Return other "similar" talks too, based on text similarity

IN CODE - BUILD THE MODEL

```
# How to do the TF-IDF conversion
tf_settings = dict( analyzer="word", ngram_range=(1,2),
                    sublinear_tf=True, min_df=0, stop_words='english' )

# Build the TF-IDF over all the talks
# Use title + category + abstract for our text
tfidf = TfidfVectorizer(**tfs)
tfidf_matrix = tfidf.fit_transform(talks["learn"])
print(tfidf_matrix.shape)    # eg 294 x 30076

# Build the similarities of each talk against every other talk
# We'll use this for scoring
tfidf_similarities = linear_kernel(tfidf_matrix, tfidf_matrix)

# Build a model, using Multinomial Naive Bayes
# Model the text of the talk, to predict the talk's index
classifier = MultinomialNB()
model = make_pipeline(tfidf, classifier)
learn_text = talks["learn"]
model.fit( list(learn_text), list(learn_text.index) )
```

IN CODE - USE THE MODEL

```
query = "apache tika"

# Ask the model to compare our query against every talk,
# then pick the talk it thinks is the most similar
pred_idx = model.predict([query])

# The prediction should be the index of that talk
print("Best match - talk %d" % pred_idx)

# Get the pairwise similarity scores of all other talks with that one
# Filter for ones high enough, and sort so highest scores come first
similarities = tfidf_similarities.[pred_idx]
sim_scores = list( [idx,s] for idx,s in enumerate(similarities[0]) if s > 0.01 )
sim_scores = sorted(sim_scores, key=lambda x: x[1], reverse=True)

# Get the scores of the x most similar talks
sim_scores = sim_scores[0:max_hits]

# Grab those talks
indexes = [i[0] for i in sim_scores]
return talks.iloc[indexes]
```

HOW WELL DOES IT WORK?

Moderately well?

(But we don't have a set of known-good examples, so we can't calculate an exact score)

LET'S TRY A LIVE DEMO!

<https://bbuzz2019mlsearch-gagravarr.notebooks.azure.com/j/notebooks/BuildAndRecommend.ipynb>

NEXT APPROACH - CLUSTER TALKS FIRST

Clustering - Grouping similar, but unlabelled things together

Rather than matching then finding similar things, as we did with the Naive Bayes approach, what if we clustered first?

K-Means, originally from signal processing, will let us group n things into k groupings, based on minimising the sum of the squared errors.

But how many clusters should we make?

CLUSTER SIZING

If we have k terms in our TF-IDF, then k clusters with fit with zero error. 1 cluster will be maximum error. Neither helps much...

We need something in-between, where most things are grouped together, and not too many things are a long way from their neighbours.

Various measures can be used to test the effectiveness of the clustering, *Average Silhouette* and *Gap Statistic* methods seem quite popular.

Need to run a few times at each k value (with different init seeds to avoid local-minima), then compute effectiveness measure, then pick best.

IN CODE - BUILD DIFFERENT CLUSTERS

```
# Try a range of k-sizes
krange = range(25, talks.shape[0]-25)

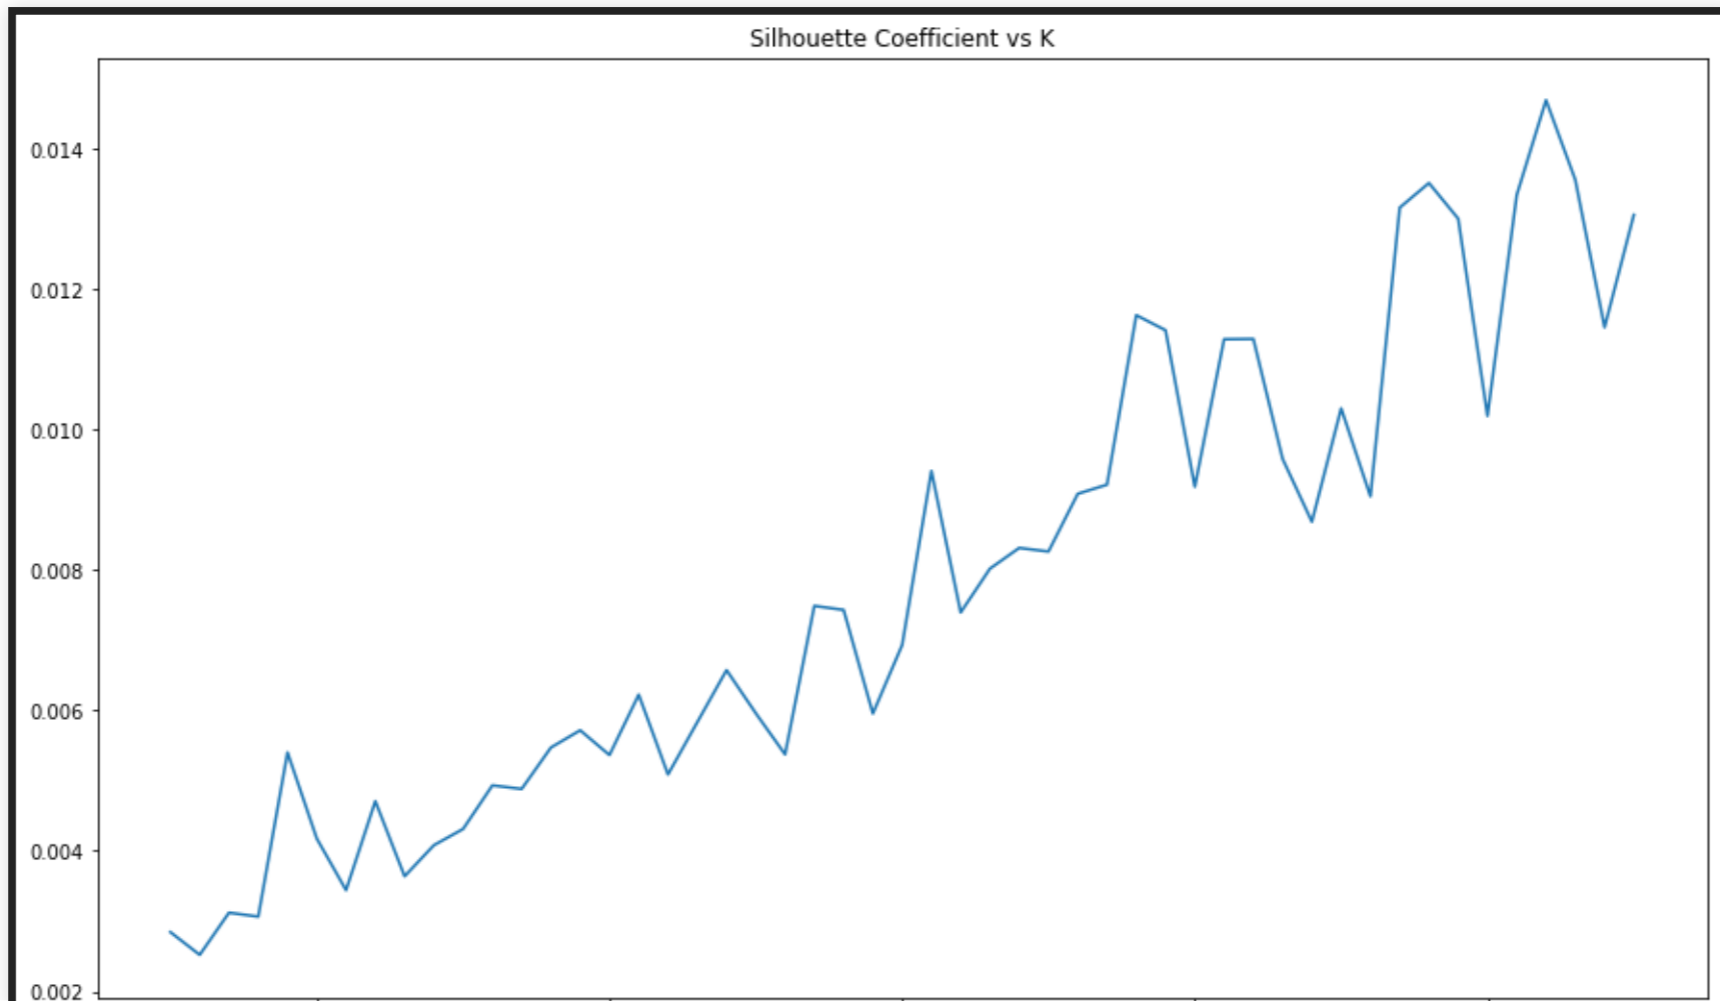
# Try each one with a few different inits, to avoid getting stuck in local-minima
kms = []
for k in krange:
    tfidf = tfidf_word

    print("Building k-means cluster of size %d" % k)
    km = KMeans(n_clusters=k, init='k-means++', max_iter=100, n_init=15, verbose=False)
    km.fit(tfidf.matrix)

    scoff = metrics.silhouette_score(tfidf.matrix, km.labels_, sample_size=1000)
    print(" - Silhouette Coefficient: %0.4f" % scoff)

    kms.append( [k,scoff] )

# Which was the best?
kms.sort(key=lambda x: x[1], reverse=True)
print("Best K-Means found with a cluster-size (k) of %d" % kms[0][0])
print("That had a Silhouette Coefficient: %0.4f" % kms[0][1])
```




CAN WE SEE HOW IT WORKED?

Our input data had 20-30k dimensions, our cluster has ~50


The human brain struggles with much more than about 4...

A 3D + colour plot is about the limit for most people

Techniques like t-SNE and PCA let us throw away a load of the dimensionality,
whilst still keeping some (but not all) of the info



t-SNE visualisation of raw TF-IDF



NEXT APPROACH - CLUSTER TALKS FIRST

First, identify our "optimal" cluster size

Next, build a k-means clustering of our talks

Match the text query to a cluster, and find the cluster centre

Now match talks based on similarity to that cluster centre (rather than one specific talk), possibly boosting slightly talks from our cluster

IN CODE - USING CLUSTERS

```
# Identify which cluster each talk belongs to
talk_clusters = km.predict(tfidf.matrix)
cluster_talks = defaultdict(list)
for talk_id, cluster_id in enumerate(talk_clusters):
    cluster_talks[cluster_id].append(talk_id)

# Find best cluster for our query, cluster centre and talks
query_tf = tfidf.transform([query])
cluster_id = int(km.predict(query_tf))

c_centre_tfidf = km.cluster_centers_[cluster_id]
c_talk_ids = cluster_talks[cluster_id]

# Compare all talks with this cluster centre, then order
similarities = linear_kernel(c_centre_tfidf, tfidf.matrix)

tscores = [[i,s] for i,s in enumerate(similarities[0])]
for idx, score in tscores:
    if idx in c_talk_ids:
        tscores[idx][1] = score+0.1
tscores = sorted(tscores, key=lambda x: x[1], reverse=True)
```

HOW DOES K-MEANS CLUSTERING DO?

Fairly similar? Maybe a tiny bit worse?

(Still don't have a set of known-good examples, so we still can't calculate an exact score!)

Clusters look to have sensible terms in them, results seem mostly what we'd expect

LET'S TRY ANOTHER LIVE DEMO!

<https://bbuzz2019mlsearch-gagravarr.notebooks.azure.com/j/notebooks/BuildAndRecommend.ipynb>

ADDING TALK YEAR TO SCORING?

Can we prefer newer talks to older ones?

Can't easily add it as a feature, since we don't have a year at query time, and we can't teach the model directly what to prefer

We can add it at scoring time, by multiplying scores by year factor. However, that only applies once we have matched to a talk / cluster.

Ideally need to feed in before we build the model, not after.

If we reduce TF-IDF weights slightly for older talks, that will de-boost them before model is built, but may affect how we cluster.

OUR DATA ISN'T LONG-TERM STATIC

Next year, there'll be another Berlin Buzzwords! That means more talks (Likewise we're also adding new training data to our onboarding chatbot, answering more RFPs, building more complex SDTM mapping rules etc)

Let's say we took the time to manually identify the best talks for some query terms, either explicitly, or via user feedback in a webapp

We boost / prioritise / train based on these "known correct" answers, and it all looks good!

Then we add more talks, quite possibly even more relevant ones for our queries, but we our ML won't prioritise them as not on the "good list"

TOKENISATION REVISITED

We need to turn our text into a stream of chunks to feed into the TF-IDF

Stopwords - Very common words that don't help much with the query, and can bloat the TF-IDF.

Stemming - Bringing different forms of a word to a common root, eg `talk`
`talks` `talked` `talking` so they can be matched interchangeably

n-grams - word-based means treating several words as one token, eg word bigrams means word pairs. char-based means splitting word into overlapping chunks, eg trigrams of `hello` are `hel` `ell` `llo`

Need to use the same for learning *and* for querying!

<https://developers.google.com/machine-learning/guides/text-classification/step-3>

EMBEDDINGS AND FEATURE EXTRACTION

Even from just a few hundred talks, our TF-IDF had a lot different terms in it. However, most talks only use a subset of those terms, so lots of TF-IDF matrix values are 0. Our TF-IDF is *sparse*

Bayse and K-Means, amongst others, are fine with sparse matrices. Other techniques, eg Neural Networks, need *dense* matrices, where most values are non-zero

Using stop-words and stemming helps a little bit here, but we need a few more orders of magnitude change!

https://scikit-learn.org/stable/modules/feature_extraction.html#text-feature-extraction

<https://developers.google.com/machine-learning/crash-course/embeddings/obtaining-embeddings>

EMBEDDINGS AND FEATURE EXTRACTION

We need to extract some new features from our existing data, which have much lower dimensionality.

Instead of 20k-30k terms, we want more like 50-250

Our per-talk matrix of features will then be smaller and denser (fewer features/dimensions, more non-zero values)

This new dense set of features is known as an *embedding*

As well as being needed before we can use Neural Networks on our kind of input, they also help with visualisations of you data

Our k-means clusters are a form of term-embedding


WORD2VEC

Word2Vec was originally developed by Google. It's a series of models that produce *word embeddings*.


Based on shallow, 2-layer neural networks, trained to reconstruct linguistic contexts of words

Produces a vector space, typically reduced down to a few hundred dimensions, where words with similar contexts are located nearby.


With enough data, model can be used to make guesses on a word's meaning and association, e.g. “man” is to “boy” what “woman” is to “girl”



Male-Female



Verb Tense



Country-Capital

NEURAL NETWORKS

System built up with a series of layers, with data flowing one way

Input layer accepts your features / embeddings


Output layer constrained to give answer you expect, get "cat or dog"

Potentially many hidden layers between, which do the processing

Training done to build the nodes, the weights of the links between them, if you should add or remove nodes in a layer etc

A simple neural network

input layer hidden layer output layer



TESTING, CROSS VALIDATION

When we know what our model should be predicting for a set of input data (eg what the human-provided labels are for a bunch of input features), we need to measure how well the model does!

Train on one set of known-data, then predict against another set, and see how close the predictions came to what we know are the answers

Our model might do *too well*, if it ends up learning some specific bits of our training data, this is *over-fitting*

Generally we want a train / test split, and maybe validate too, which should all be representative of the data we'll predict with

If we don't have much data available, we can train the model several times with different segments as train/test, and ensure always similar accuracy

HYPER-PARAMETERS AND TUNING

A hyper-parameter is anything we tune / select / change in our ML, that's independent of the data and of the features selected

eg k-means, that includes the range of clusters to try fitting to

Often relate to seeds, steps, number of clusters / layers, how much to change between iterations, and how much to leave alone

Pick the wrong parameters, and your model might get stuck in a local minima a long way off the best answer, or might take ages to converge, or may never even complete!

Along with identifying appropriate features, and cleaning your data, selecting appropriate hyper-parameters is a tough bit of data science

ERRORS

For a binary classification, there are 4 possible states: True Positive, True Negative, False Positive and False Negative

What to aim for depends on your problem, and the distribution of your data

eg if detecting cancer, is it better to give someone the all-clear when they actually have cancer (false negative)? Or to send them for treatment that they don't need (false positive)

Precision is ratio of correct values in our results, **recall** is ratio correct-found to correct-all. **Confidence** is how well the model thinks it did.

BIASES

If you input data is biased, the model will be biased

If you try to hide some biases from your model, it might still find them from other features.

eg hide Gender, but leave in Name. eg hide Race, but leave in postal code / zipcode, or first / elementary school

Be aware of your data biases, be aware of how people will use your model, try to re-weight your models to counteract biases

Be aware of implicit biases in your data, and impact if model is fed data from somewhere else / something else

Good intro yesterday - <https://berlinbuzzwords.de/19/session/bias-nlp-101>

NLP - NATURAL LANGUAGE PROCESSING

Family of AI/ML related (and often powered) techniques for understanding text

NER Named Entity Recognition - identify proper names in text, eg people / places / organisations

POS Part-of-speech tagging - figure out which words / word sequences are nouns / verbs / adjectives

Sentiment analysis - is the text positive / negative / neutral / etc on the topic it relates to

Plus stemming, word and sentence break identification etc

DETECTING ANSWERS IN TEXT

So far, we've focused on matching a question to some relevant material, rather than a single specific answer

However, if the text has a simple answer embedded in it, potentially NLP could help us pull that out to display / highlight

DrQA - originally developed by Facebook is an ML + IR system for answering factoid questions. Demo is trained on Wikipedia pages, but can be retrained for your documents too!

<https://github.com/facebookresearch/DrQA>

<https://berlinbuzzwords.de/19/session/building-enterprise-natural-language-search-engine-elasticsearch-and-facebooks-drqa>

WOULD LUCENE / ELASTIC HAVE DONE BETTER?

Out of the box, and with our data sizes - no

If we'd spent some time configuring the similarity and clustering features, could probably come pretty close

If we had a lot more data, Lucene / SOLR / Elastic would win - most of the AI techniques need the whole model in memory (Lucene can efficiently pages bits in from disk), and Lucene has lower CPU needs at index/predict

But as an AI / ML learning exercise, a chance to put that into action on a few production projects, and to aid planning of our next data-driven ML projects, it was invaluable!

TRYING OUT AI / ML YOURSELF

Basically everything is open source! Installation and setup isn't always the easiest though...

Jupyter - an open-source web application that allows you to create and share documents that contain live code, equations, visualizations and narrative text. (There's also *Apache Zeppelin*!)

There are several *free* hosted Jupyter instances, which come "batteries included" with all the libraries you need to get started!

Azure Notebooks - <https://notebooks.azure.com/>

Google Colaboratory - <https://colab.research.google.com/>

FURTHER LEARNING RESOURCES - COURSES

Stanford University

- <https://www.coursera.org/learn/machine-learning>

Google

- <https://developers.google.com/machine-learning/crash-course/>
- <https://developers.google.com/machine-learning/guides/rules-of-ml/>
- <https://developers.google.com/machine-learning/guides/text-classification/>

Amazon

- <https://aws.amazon.com/training/learning-paths/machine-learning/>

FURTHER LEARNING RESOURCES - COURSES

Microsoft

- <https://docs.microsoft.com/en-us/azure/machine-learning/>

Jeremy Howard / Fast.AI

- <http://course18.fast.ai/ml>
- <http://course18.fast.ai/part2.html>

O'Reilly Safari has loads!

Others

- <https://scikit-learn.org/stable/tutorial/basic/tutorial.html>
- <https://scipy-lectures.org/>

FURTHER LEARNING RESOURCES - OTHERS

- <http://colah.github.io/posts/2014-03-NN-Manifolds-Topology/>
- <https://www.ben-evans.com/benedictevans/2019/4/15/notes-on-ai-bias>
- <https://www.graphcore.ai/posts/removing-bias-from-machine-learning>
- **DrQA**
 - <https://arxiv.org/abs/1704.00051>
 - <https://github.com/facebookresearch/DrQA>

FURTHER LEARNING RESOURCES - BOOKS

Taming Text - <https://www.manning.com/books/taming-text>

Introduction to Information Retrieval - <https://nlp.stanford.edu/IR-book/>

Loads available from Manning

And from O'Reilly

ANY QUESTIONS?