Lecture 9:
Introduction to Language

(new Gen Z friendly lecture)

Administrative: Assignment 1

- Grades released yesterday
- Please submit regrade requests by tomorrow night

Administrative: Assignment 3

Due 3/02 11:59pm

- Multi-layer Neural Networks,

- Image Features,

- Optimizers

# "Vanilla" Neural Network

one to one
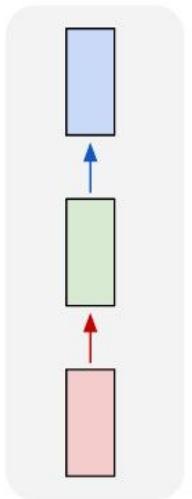
Vanilla Neural Networks

Recurrent Neural Networks: Process Sequences
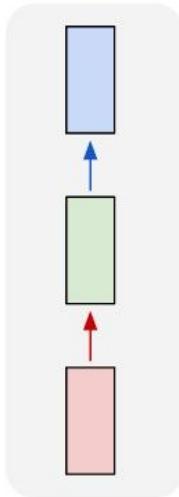
one to one
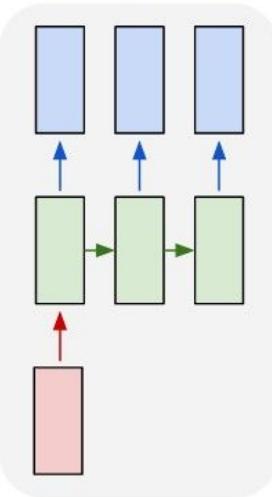
one to many

e.g. **Image Captioning**
image -> sequence of words

Recurrent Neural Networks: Process Sequences
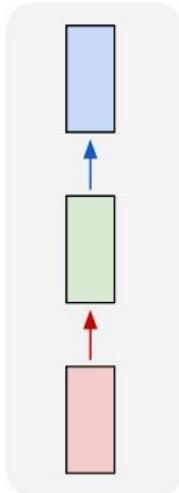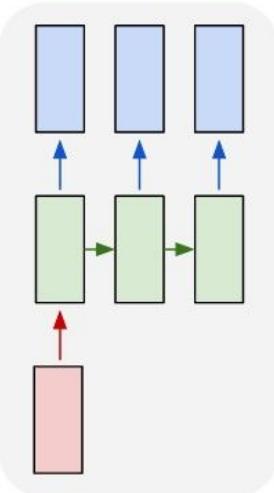
one to one   one to many   many to one
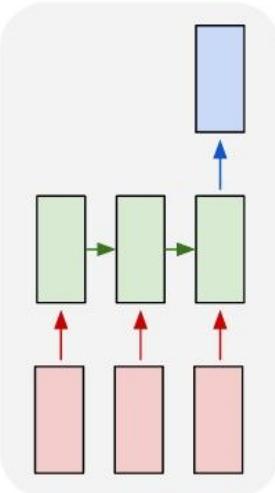
e.g. **action prediction, sentiment classification**
sequence of video frames -> action class

Recurrent Neural Networks: Process Sequences
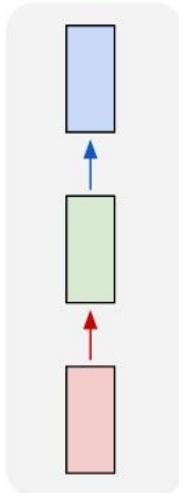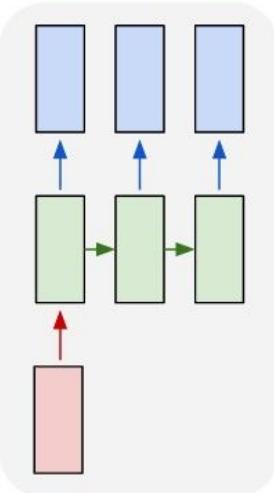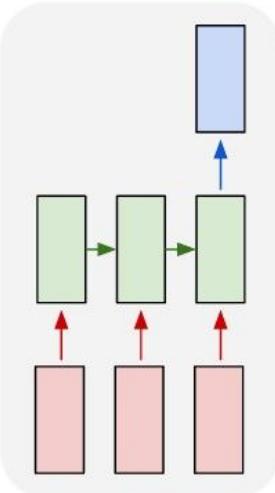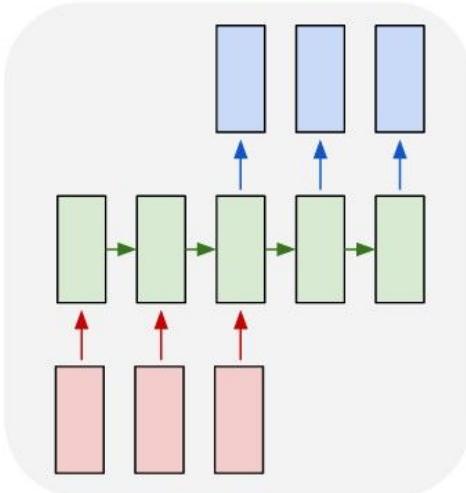
E.g. **Video Captioning**
Sequence of video frames ->
caption

Recurrent Neural Networks: Process Sequences
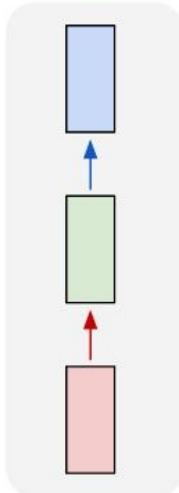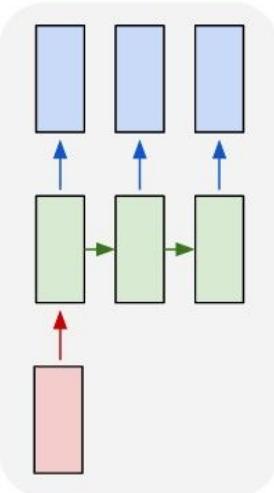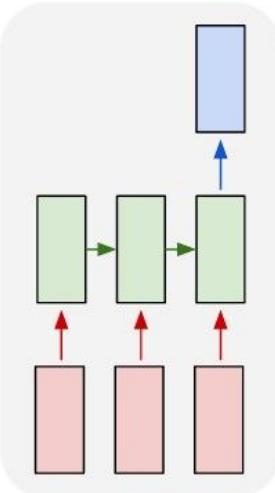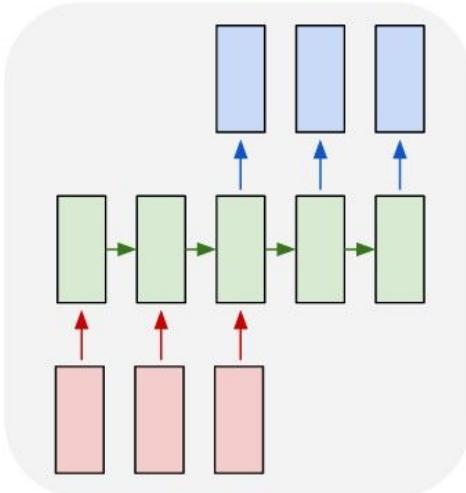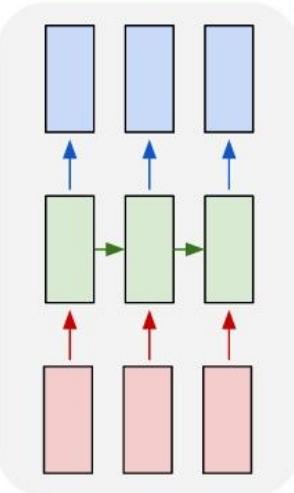
e.g. **Video classification on frame level**
Fill in the blanks with words

Sequential Processing of Non-Sequence Data

Classify images by taking a
series of "glimpses"
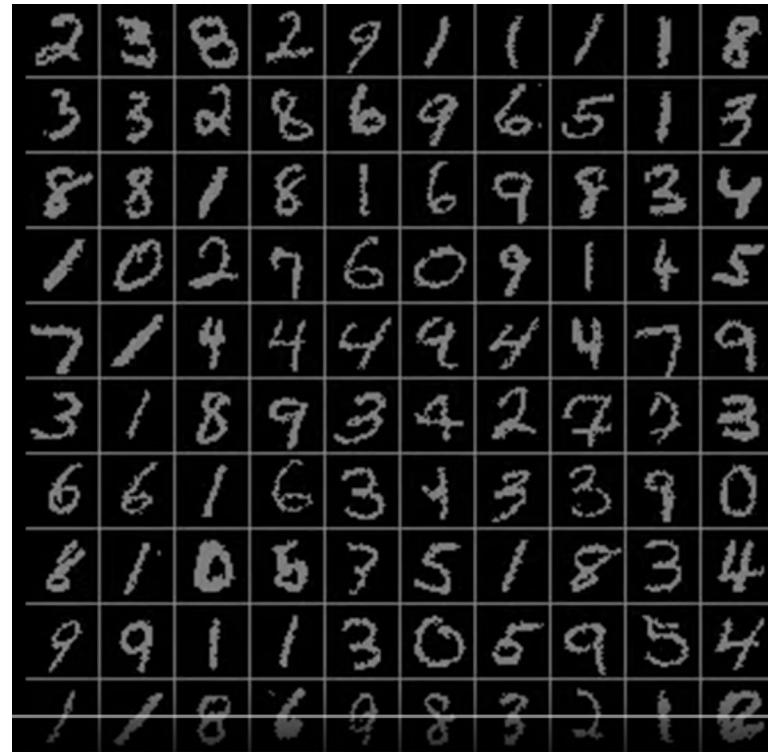


Ba, Mnih, and Kavukcuoglu, "Multiple Object Recognition with Visual Attention", ICLR 2015.
Gregor et al, "DRAW: A Recurrent Neural Network For Image Generation", ICML 2015
Figure copyright Karol Gregor, Ivo Danihelka, Alex Graves, Danilo Jimenez Rezende, and Daan Wierstra,
2015. Reproduced with permission.

Sequential Processing of Non-Sequence Data
Generate images one piece at a time!
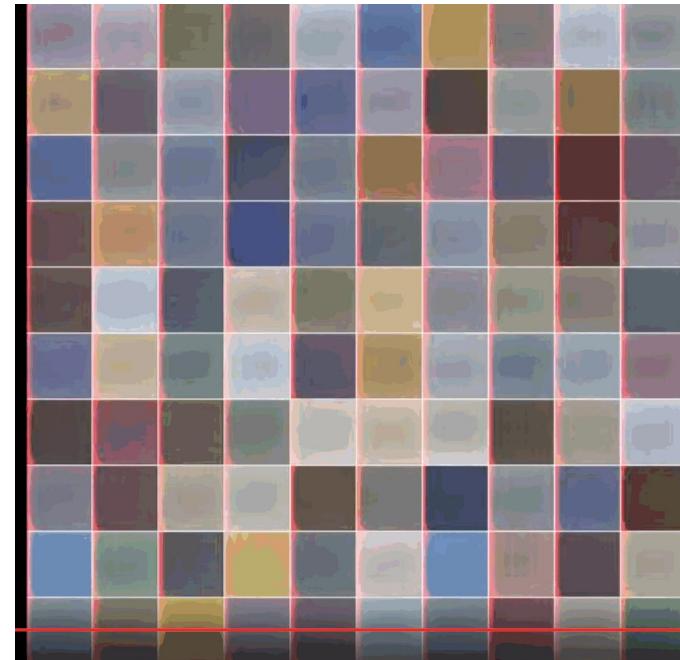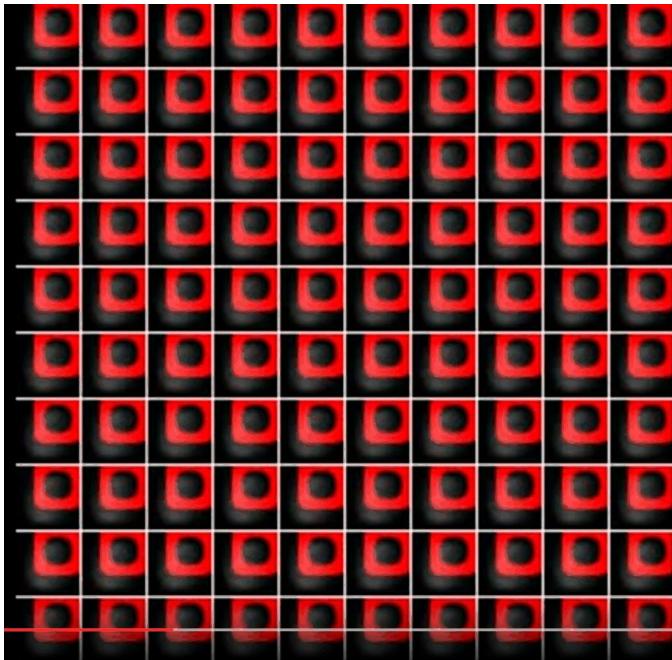
Gregor et al, "DRAW: A Recurrent Neural Network For Image Generation", ICML 2015
Figure copyright Karol Gregor, Ivo Danihelka, Alex Graves, Danilo Jimenez Rezende, and Daan Wierstra, 2015. Reproduced with permission.

Why existing neural networks are insufficient?

Variable sequence length inputs and outputs!

Example task: video captioning

Input video can have variable number of frames

Output captions can be variable length.

Input Video

Output Captions

T Frames

A lady joins the man and sings along to the music.
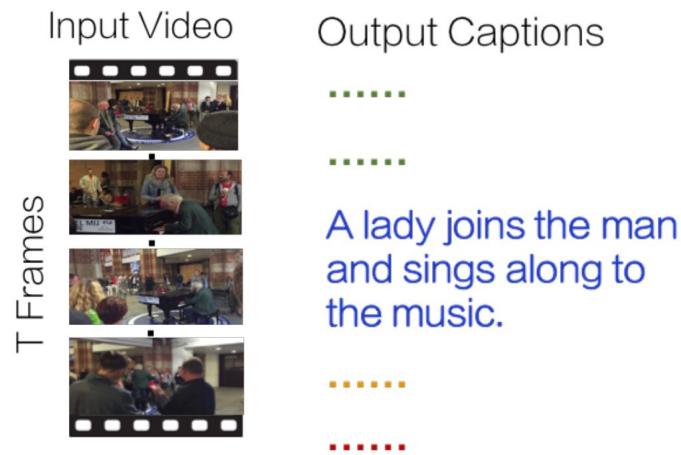
Krishna, et al. Dense captioning Events in Videos. ICCV 2019
Figure copyright Ranjay Krishna, Kenji Hata, Frederic Ren, Li Fei-Fei, and Juan Carlos Niebles, 2019. Reproduced with permission.

Today: Introduction to language

Language models are revolutionizing search

User: when did Kendrick lamar's first album come out?

GPT: July 2, 2011

Search bar: when did Kendrick lamar's first album come out?

wikipedia.org
https://en.wikipedia.org › wiki › Kendrick_Lamar_dis...

Kendrick Lamar discography - Wikipedia

Lamar's debut studio album Section.80, was released on **July 2, 2011**, and issued on independent record label Top Dawg Entertainment (TDE). It peaked at number ...

Albums · Singles · Other charted and certified songs · Other guest appearances

wikipedia.org
https://en.wikipedia.org › wiki › Kendrick_Lamar

Kendrick Lamar - Wikipedia

Following the success of his alternative rap debut album **Section.80 (2011)**, Lamar secured a joint contract with Aftermath Entertainment and Interscope Records.

Discography · Damn (Kendrick Lamar album) · Kendrick Lamar (EP) · Baby Keem

genius.com
https://genius.com › artists › Kendrick-lamar › first-al...

When did Kendrick Lamar's first album release? - Genius

Kendrick Lamar's first album Y.H.N.I.C. (Hub City Threat: Minor of the Year) released in **2004**.

allmusic.com
https://www.allmusic.com › artist › kendrick-lamar-m...

Kendrick Lamar Songs, Albums, Reviews, Bio & More | AllMusic

Lamar's major-label debut, good kid, m.A.A.d city, was released in **October 2012** and entered the Billboard 200 at number two. Three of its singles -- "Swimming ...

Language models can learn to follow examples

S: I broke the window.
Q: What did I break?
S: I gracefully saved the day.
Q: What did I gracefully save?
S: I gave John flowers.

What will GPT generate?

Language models can learn to follow examples

S: I broke the window.
Q: What did I break?
S: I gracefully saved the day.
Q: What did I gracefully save?
S: I gave John flowers.
Q: Who did I give flowers to?

Language models can learn to follow examples

S: I broke the window.
Q: What did I break?
S: I gracefully saved the day.
Q: What did I gracefully save?
S: I gave John flowers.
Q: Who did I give flowers to?
S: I gave her a rose and a guitar.

What will GPT generate?

Language models can learn to follow examples

S: I broke the window.
Q: What did I break?
S: I gracefully saved the day.
Q: What did I gracefully save?
S: I gave John flowers.
Q: Who did I give flowers to?
S: I gave her a rose and a guitar.
Q: Who did I give a rose and a guitar to?

Language models can even write code / sql

How many users have signed up since the start of 2020?
SELECT count(id) FROM users WHERE created_at > '2020-01-01'

What is the average number of influencers each user is subscribed to?
SELECT avg(count) FROM (SELECT user_id, count(*) FROM subscribers GROUP BY user_id) AS avg_subscriptions_per_user

Compared to Vision (ImageNet, MNIST), language benchmarks reach super-human fast!
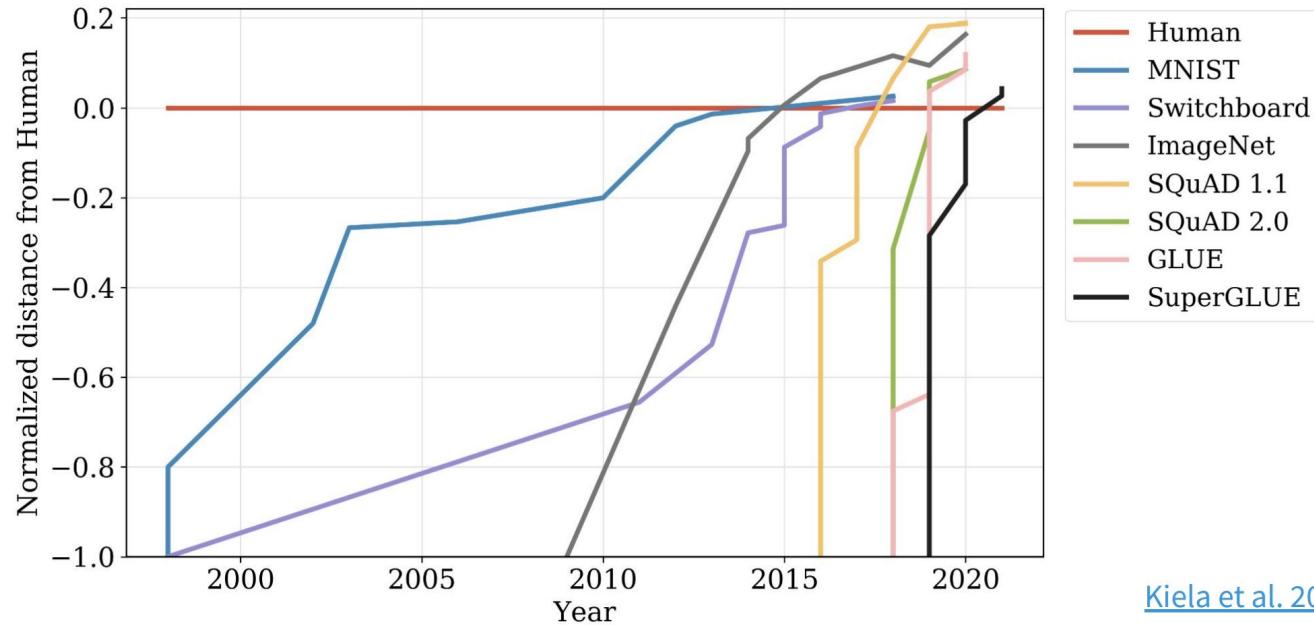
Kiela et al. 2021

Kiela et al. Dynabench: Rethinking Benchmarking in NLP NAACL 2021

An example that you all will resonate with:

GPT capabilities today

You

Please write up a polite letter to my professor, Ranjay Krishna, explaining why I need two more late days to submit my assignment 3. Make up a good excuse. I live in seattle.

ChatGPT

Subject: Request for Extension on Assignment 3 Submission

Dear Professor Ranjay Krishna,

I hope this message finds you well. I am writing to respectfully request an extension for the submission of Assignment 3.

Unfortunately, I have encountered an unforeseen and challenging situation that has significantly hindered my ability to complete the assignment on time. The recent severe weather conditions in Seattle, coupled with unexpected power outages in my area, have greatly impacted my ability to work effectively. The power outages have not only disrupted my internet access, which is crucial for research and accessing course materials, but have also made it challenging to maintain a conducive environment for studying.

I have been making every effort to keep up with the coursework and had planned my schedule to complete the assignment on time. However, these circumstances have been beyond my control and have unexpectedly set back my progress.

In light of this, I kindly request a two-day extension for the submission of Assignment 3. This additional time would be immensely helpful in allowing me to complete the assignment to the best of my ability, despite the challenges posed by the current situation.

Make it more deferential!

ChatGPT

Subject: Respectful Request for Extension on Assignment 3 Submission

Dear Professor Ranjay Krishna,

I hope this message finds you in good health and spirits. I am writing to humbly request your consideration for an extension on the submission of Assignment 3. I deeply respect the high academic standards of your course and it is with great reluctance that I find myself in a position to ask for this concession.

Make it MORE DEFERENTIAL!!

You

Please write up a polite letter to my professor, Ranjay Krishna, explaining why I need two more late days to submit my assignment 3. Make up a good excuse. I live in seattle.

ChatGPT

Subject: Request for Extension on Assignment 3 Submission

ChatGPT

Subje

Dear F

I hope

your c

respec

I find r

ChatGPT

Subject: Humble Request for Extension on Assignment 3 Submission Due to Unforeseen Circumstances

Dear Esteemed Professor Ranjay Krishna,

I hope this message reaches you in the best of health and spirits. I am writing to you with a deep sense of respect and humility, to request an extension for the submission of Assignment 3. It is with a heavy heart that I find myself in a position where I must seek your understanding and compassion due to unforeseen circumstances beyond my control.

pectfully request an extension for

challenging situation that has gnment on time. The recent expected power outages in my ely. The power outages have not research and accessing course in a conducive environment for

ursework and had planned my er, these circumstances have ack my progress.

In light of this, I kindly request a two-day extension for the submission of Assignment 3. This additional time would be immensely helpful in allowing me to complete the assignment to the best of my ability, despite the challenges posed by the current situation.

Ranjay Krishna　　　　　　　　　Lecture 9

Make it MORE DEFERENTIAL!!!!!

ChatGPT

Subject: Request for Extension on Assignment 3 Submission

ChatGPT

Subjec...

Dear F...

I hope...
your c...
respec...
I find n...

ChatGPT

Subject: Hu...
Unforesee...

Dear Estee...

I hope this...
with a deep...
of Assignm...
seek your u...
my control...

ChatGPT

Subject: Most Respectful Request for Extension on Assignment 3 Due to Extenuating Circumstances

Honorable Professor Ranjay Krishna,

I hope this message finds you in the best of health and circumstances. It is with the utmost respect and a deep sense of humility that I find myself compelled to write to you, seeking your compassionate understanding and consideration. I am fully aware of the gravity of my request and approach you with great reverence, hoping for your empathy in a matter that significantly affects my academic responsibilities.

assignment to the best of my ability, despite the challenges posed by the current situation.

Ranjay Krishna   Lecture 9

How do we represent the meaning of a word?

Definition: meaning (Webster dictionary)
- the idea that is represented by a word, phrase, etc.
- the idea that a person wants to express by using words, signs, etc.
- the idea that is expressed in a work of writing, art, etc.

Common-est linguistic way of thinking of meaning:

signifier (symbol) ⟺ signified (idea or thing)

tree ⟺ {🌳, 🌲, 🌴, ...}

How do we have usable meaning in a computer?

Previous solution by Linguists: Use a thesaurus (e.g., WordNet) containing lists of synonym sets and hypernyms ("is a" relationships).

e.g., synonym sets containing "good":

```python
from nltk.corpus import wordnet as wn
poses = { 'n':'noun', 'v':'verb', 's':'adj (s)', 'a':'adj', 'r':'adv'}
for synset in wn.synsets("good"):
    print("{}: {}".format(poses[synset.pos()],
            ", ".join([l.name() for l in synset.lemmas()])))
```

```
noun: good
noun: good, goodness
noun: good, goodness
noun: commodity, trade_good, good
adj: good
adj (sat): full, good
adj: good
adj (sat): estimable, good, honorable, respectable
adj (sat): beneficial, good
adj (sat): good
adj (sat): good, just, upright
…
adverb: well, good
adverb: thoroughly, soundly, good
```

How do we have usable meaning in a computer?

Previous solution by Linguists: Use a thesaurus (e.g., WordNet) containing lists of synonym sets and hypernyms ("is a" relationships).

e.g., hypernyms of "panda":

```python
from nltk.corpus import wordnet as wn
panda = wn.synset("panda.n.01")
hyper = lambda s: s.hypernyms()
list(panda.closure(hyper))
```

```
[Synset('procyonid.n.01'),
Synset('carnivore.n.01'),
Synset('placental.n.01'),
Synset('mammal.n.01'),
Synset('vertebrate.n.01'),
Synset('chordate.n.01'),
Synset('animal.n.01'),
Synset('organism.n.01'),
Synset('living_thing.n.01'),
Synset('whole.n.02'),
Synset('object.n.01'),
Synset('physical_entity.n.01'),
Synset('entity.n.01')]
```

IMAGENET

classes were derived from WordNet

22K categories and **14M** images

- Animals
  - Bird
  - Fish
  - Mammal
  - Invertebrate
- Plants
  - Tree
  - Flower
- Food
- Materials
- Structures
- Artifact
  - Tools
  - Appliances
  - Structures
- Person
- Scenes
  - Indoor
  - Geological Formations
- Sport Activities

Deng, Dong, Socher, Li, Li, & Fei-Fei, 2009

Problems with resources like WordNet

A useful resource but missing nuance:
- e.g., "proficient" is listed as a synonym for "good"
  This is only correct in some contexts
- Also, WordNet lists offensive synonyms without any coverage of the connotations or appropriateness of words

Missing new meanings of words:
- e.g., sus, cap, glow up, fam, drip, vibe, G.O.A.T.
- Impossible to keep up-to-date with Gen Z!

Subjective
- Requires human labor to create and adapt
- Can't capture similarity between words

Representing words as discrete symbols

In traditional NLP, we regard words as discrete symbols:
hotel, conference, motel – each has its own symbol.
This is a localist representation

Such symbols for words can be represented by one-hot vectors:
motel = [0 0 0 0 0 0 0 0 0 0 1 0 0 0 0]
hotel = [0 0 0 0 0 0 0 1 0 0 0 0 0 0 0]

Vector dimension = number of words in vocabulary (e.g., 500,000+)

Problem with words as discrete symbols

Example: in web search, if a user searches for "Seattle motel", we would like to match documents containing "Seattle hotel"
But:

motel = [0 0 0 0 0 0 0 0 0 0 1 0 0 0 0]
hotel = [0 0 0 0 0 0 0 1 0 0 0 0 0 0 0]

These two vectors are orthogonal
There is no natural notion of similarity for one-hot vectors!

Can we use WodNet-like taxonomies?

Could try to rely on WordNet's list of synonyms to get similarity?
- But it is well-known to fail badly: incompleteness, etc.

New idea: learn better word representations that encode similarity in the vectors themselves

- Similar to deep learning representations for images!
- Can we learn representations for words?

Just as Neural networks extract linearly separate image features,
can we learn word features?

f

Feature Extraction

10 numbers giving scores for classes

training

Krizhevsky, Sutskever, and Hinton, "Imagenet classification with deep convolutional neural networks", NIPS 2012. Figure copyright Krizhevsky, Sutskever, and Hinton, 2012. Reproduced with permission.

10 numbers giving scores for classes

training
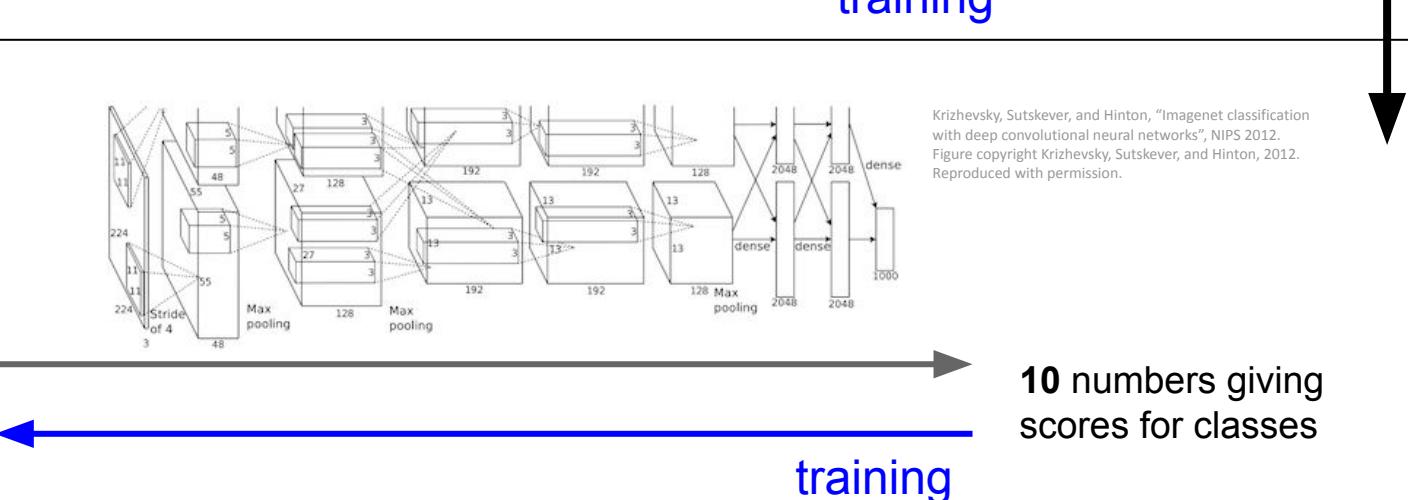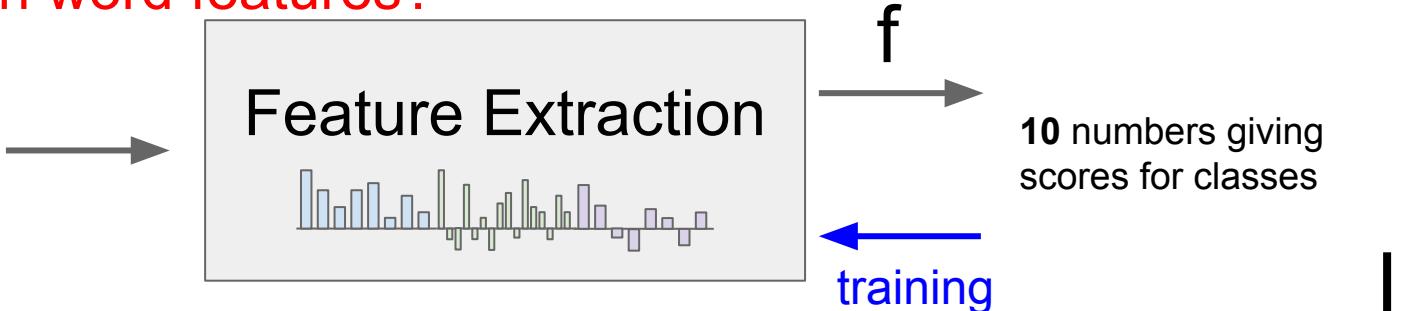
Representing words by their context

Distributional semantics: A word's meaning is given by the words that frequently appear close-by

"You shall know a word by the company it keeps" (J. R. Firth 1957: 11)

One of the most successful ideas of modern statistical language modeling!

High level idea: a word w in text is a function of its contextual words.

Let's try it: Fill in the blanks.

➡️ It's cold today! Don't forget to wear a _____.

The _____ is a popular tourist attraction in Seattle.

I missed ____ bus.

I had 3 pencils and lost one so now I have _____ pencils.

Let's try it: Fill in the blanks.

➡️ It's cold today! Don't forget to wear a <u>jacket / coat / sweater</u>.

The _____ is a popular tourist attraction in Seattle.

I missed ___ bus.

I had 3 pencils and lost one so now I have _____ pencils.

Let's try it: Fill in the blanks.

It's cold today! Don't forget to wear a jacket / coat / sweater.

➡️ The _____ is a popular tourist attraction in Seattle.

I missed ___ bus.

I had 3 pencils and lost one so now I have _____ pencils.

Let's try it: Fill in the blanks.

It's cold today! Don't forget to wear a <u>jacket / coat / sweater</u>.

➡ The <u>Space Needle</u> is a popular tourist attraction in Seattle.

I missed ___ bus.

I had 3 pencils and lost one so now I have _____ pencils.

Let's try it: Fill in the blanks.

It's cold today! Don't forget to wear a <u>jacket / coat / sweater</u>.

The <u>Space Needle</u> is a popular tourist attraction in Seattle.

➡️ I missed ___ bus.

I had 3 pencils and lost one so now I have _____ pencils.

Let's try it: Fill in the blanks.

It's cold today! Don't forget to wear a <u>jacket / coat / sweater</u>.

The <u>Space Needle</u> is a popular tourist attraction in Seattle.

➡️ I missed <u>the</u> bus.

I had 3 pencils and lost one so now I have _____ pencils.

Let's try it: Fill in the blanks.

It's cold today! Don't forget to wear a <u>jacket / coat / sweater</u>.

The <u>Space Needle</u> is a popular tourist attraction in Seattle.

I missed <u>the</u> bus.

➡ I had 3 pencils and lost one so now I have _____ pencils.

Let's try it: Fill in the blanks.

It's cold today! Don't forget to wear a <u>jacket / coat / sweater</u>.

The <u>Space Needle</u> is a popular tourist attraction in Seattle.

I missed <u>the</u> bus.

➡ I had 3 pencils and lost one so now I have <u>2 / two</u> pencils.

What we want in the end: word vectors

We will build a dense vector for each word,
- chosen so that it is similar to vectors of words that appear in similar contexts: e.g. jacket / coat / sweater.
- measuring similarity as the vector dot (scalar) product.
- Word vectors are also called (word) embeddings or (neural) word representations

$$\text{Jacket} = \begin{bmatrix} 0.286 \\ 0.792 \\ -0.177 \\ -0.107 \\ 0.109 \\ -0.542 \\ 0.349 \\ 0.271 \end{bmatrix}$$

Remember the **Geometric Viewpoint**
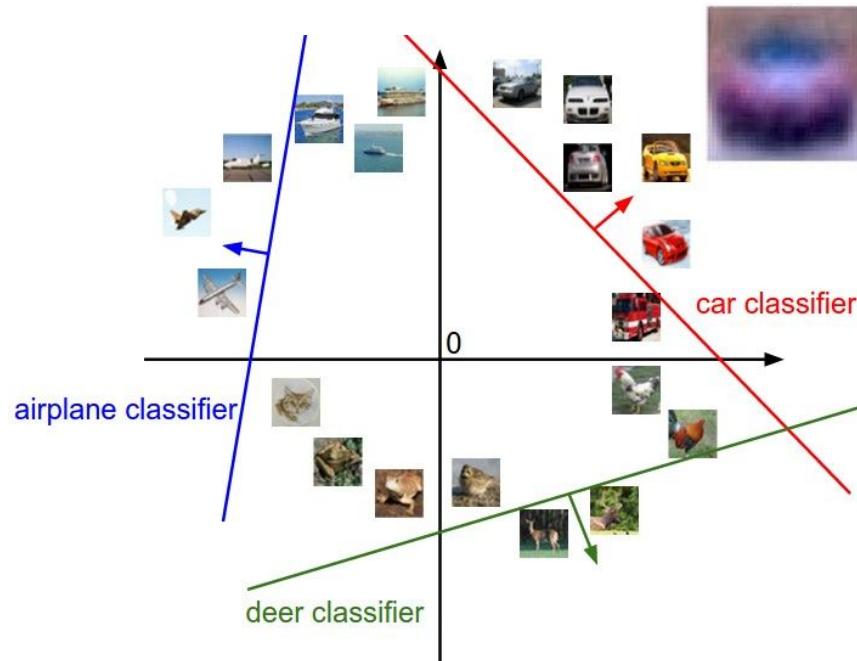
$$f(x,W) = Wx + b$$

Array of **32x32x3** numbers
(3072 numbers total)

Plot created using Wolfram Cloud

Visualizing word vectors in 2D (trained on wine reviews)
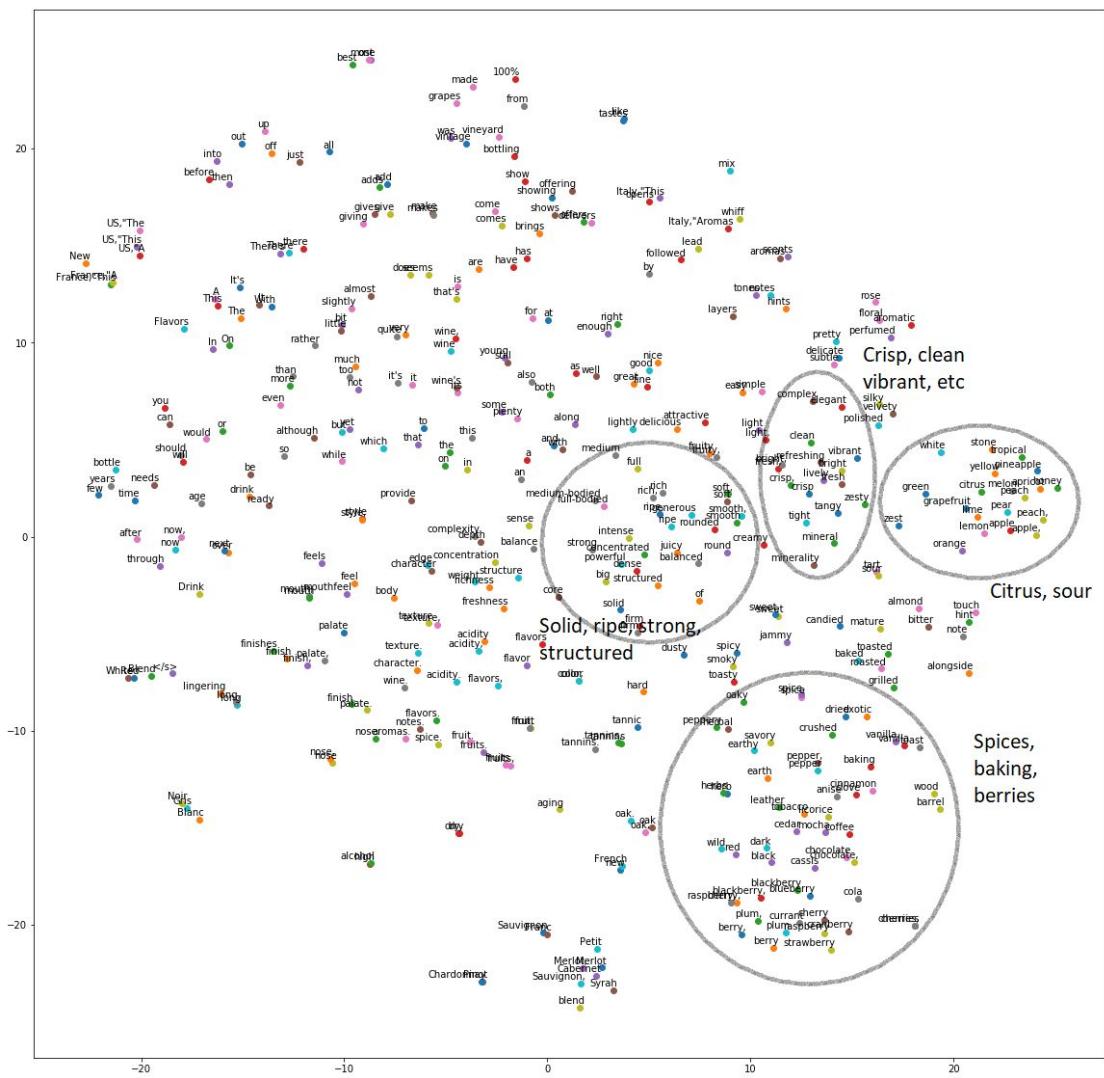
How do we train these word vectors: Word2vec

They won test of time award in December 2023!

Original paper was rejected and never published (has 40K citations!)
Mikolov et al. Efficient Estimation of Word Representations in Vector Space" 2013

Next paper by same authors is what won the award and was published in NeurIPS 2013 (also 40K citations!)
Mikolov, Tomas, et al. "Distributed representations of words and phrases and their compositionality." Advances in neural information processing systems 26 (2013).

Word2vec is a framework for learning word vectors

Idea:
1. Collect a large corpus ("body") of text: a long list of words from the internet
2. Establish a fixed vocabulary of size V.
    a. Usually most frequent V words are used.
    b. All other words are mapped to an <unknown> or <unk> word.
3. Initialize every word by a random vector
4. Go through each position t in the text,
    a. Let the word in that position be c
    b. Let all ("outside") words o be within context
    c. Calculate the probability of o given c (or vice versa)
    d. Maximize this probability

Input   projection   output

w(t)

w(t-2)

w(t-1)

w(t+1)

w(t+2)
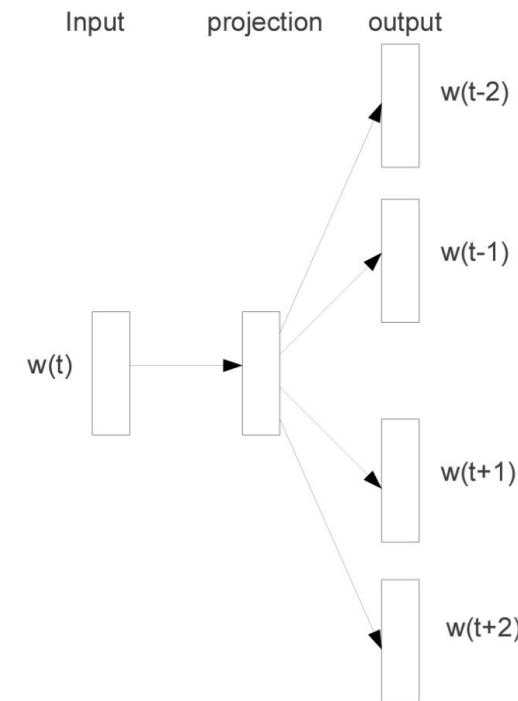
Using a context size of 2!

Example windows and process for computing $P(w_{t+j} \mid w_t)$

$$P(w_{t-2} \mid w_t) \qquad P(w_{t+2} \mid w_t)$$

$$P(w_{t-1} \mid w_t) \qquad P(w_{t+1} \mid w_t)$$

... | *problems* | *turning* | *into* | *banking* | *crises* | *as* | ...

outside context words in window of size 2 | center word at position t | outside context words in window of size 2
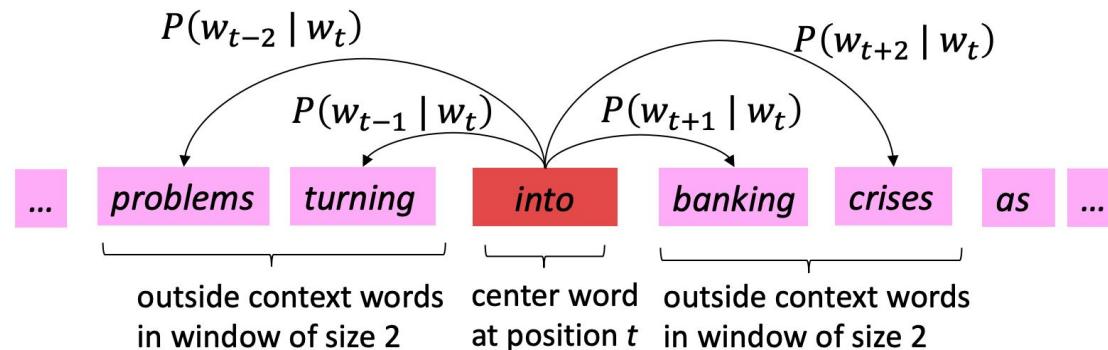
Move to the next word

Example windows and process for computing $P(w_{t+j} \mid w_t)$

$P(w_{t-2} \mid w_t)$ $P(w_{t+2} \mid w_t)$

$P(w_{t-1} \mid w_t)$ $P(w_{t+1} \mid w_t)$

… | problems | turning | into | banking | crises | as | …

outside context words in window of size 2   center word at position t   outside context words in window of size 2
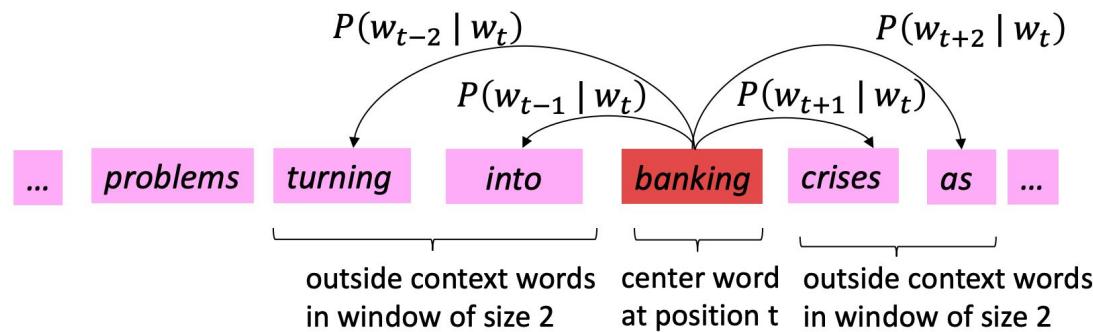
The learning objective is to maximize the probabilities

For each position *t = 1, … , T*, predict context words within a window of fixed size *m*, given center word w_t!. Data likelihood:

Likelihood = $L(\theta) = \prod_{t=1}^{T} \prod_{\substack{-m \le j \le m \\ j \ne 0}} P(w_{t+j} \mid w_t; \theta)$

θ is all variables to be optimized

The loss function

For each position $t = 1, \ldots, T$, predict context words within a window of fixed size m, given center word w_t!. Data likelihood:

Likelihood = $L(\theta) = \displaystyle\prod_{t=1}^{T} \prod_{\substack{-m \leq j \leq m \\ j \neq 0}} P(w_{t+j} \mid w_t; \theta)$

θ is all variables to be optimized

The loss function $L(\theta)$ is the (average) negative log likelihood:

$$L(\theta) = -\frac{1}{T} \log L(\theta) = -\frac{1}{T} \sum_{t=1}^{T} \sum_{\substack{-m \leq j \leq m \\ j \neq 0}} \log P(w_{t+j} \mid w_t; \theta)$$

Minimizing loss function $\Leftrightarrow$ Maximizing predictive accuracy

How do we calculate this probability?

We want to minimize the objective function:

$$L(\theta) = -\frac{1}{T}\sum_{t=1}^{T}\sum_{\substack{-m \le j \le m \\ j \ne 0}} \log P\left(w_{t+j} \mid w_t; \theta\right)$$

Question: How to calculate $P\left(w_{t+j} \mid w_t; \theta\right)$

How do we calculate this probability?

We want to minimize the objective function:

$$L(\theta) = -\frac{1}{T} \sum_{t=1}^{T} \sum_{\substack{-m \leq j \leq m \\ j \neq 0}} \log P\left(w_{t+j} \mid w_t; \theta\right)$$

Question: How to calculate $P\left(w_{t+j} \mid w_t; \theta\right)$

The best idea #1: Use a neural network whenever you need to model a function

How did Word2Vec design $f_j()$?

We will use two vectors per word w:
-   v_w when w is a center word
-   u_w when w is a context word

Let $f_1(w_t, w_{t+1}) = w_t^T w_{t+1}$ be a simple dot product
-   so now learnable parameters except the word vectors.

Then for a center word c and a context word o:

$$P(o|c) = \frac{\exp(u_o^T v_c)}{\sum_{w \in V} \exp(u_w^T v_c)}$$

Remember the **Softmax (cross-entropy) Classifier**

Want to interpret raw classifier scores as **probabilities**

$$s = f(x_i; W)$$

$$P(Y = k | X = x_i) = \frac{e^{s_k}}{\sum_j e^{s_j}}$$
Softmax Function

Probabilities must be >= 0

Probabilities must sum to 1

$$L_i = -\log P(Y = y_i | X = x_i)$$

cat	**3.2**	**24.5**	**0.13**	**1.00**
car	5.1	164.0	0.87	0.00
frog	-1.7	0.18	0.00	0.00

exp → normalize → compare ←

Unnormalized log-probabilities / logits

unnormalized probabilities

probabilities

Correct probs

Cross Entropy

$$H(P, Q) = H(p) + D_{KL}(P \| Q)$$

Understanding the calculation

② Exponentiation makes anything positive

① Dot product compares similarity of o and c.
$$u^T v = u.v = \sum_{i=1}^{n} u_i v_i$$
Larger dot product = larger probability

$$P(o|c) = \frac{\exp(u_o^T v_c)}{\sum_{w \in V} \exp(u_w^T v_c)}$$

③ Normalize over entire vocabulary to give probability distribution

To train the model: Optimize word vectors of all words to minimize loss using backprop

Recall: θ represents all the model parameters, in one long vector
- d-dimensional vectors
- V - many words,
- every word has two vectors

$$\theta = \begin{bmatrix} v_{aardvark} \\ v_a \\ \vdots \\ v_{zebra} \\ u_{aardvark} \\ u_a \\ \vdots \\ u_{zebra} \end{bmatrix} \in \mathbb{R}^{2dV}$$

Use stochastic gradient descent with batches of N center words

Why did this algorithm use two vectors?

We will use two vectors per word *w*:
- v_w when *w* is a center word
- u_w when *w* is a context word

$$P(o|c) = \frac{\exp(u_o^T v_c)}{\sum_{w \in V} \exp(u_w^T v_c)}$$

Hint: this paper came out in 2013, 1 year after AlexNet.

Why did this algorithm use two vectors?

We will use two vectors per word *w*:
- v_w when *w* is a center word
- u_w when *w* is a context word

$$P(o|c) = \frac{\exp(u_o^T v_c)}{\sum_{w \in V} \exp(u_w^T v_c)}$$

Hint: this paper came out in 2013, 1 year after AlexNet.

Do you think the results would be better if they used 1 instead of 2 vectors?

What are the final word vectors?

It is the average of the two vectors:

$\frac{1}{2}(v_w + u_w)$

Variants of algorithm: The skip-gram model versus CBOW

Skip-gram predicts context words given center word.

CBOW (continuous bag of words) predicts center word given context words.

What's another issue with this calculation?

$$P(o|c) = \frac{\exp(u_o^T v_c)}{\sum_{w \in V} \exp(u_w^T v_c)}$$

Solution: negative sampling

$$\log P(o|c) = -\log \sigma(\boldsymbol{u}_o^T \boldsymbol{v}_c) - \sum_{k \in \{K\ sampled\ indices\}} \log \sigma(-\boldsymbol{u}_k^T \boldsymbol{v}_c)$$

Replace denominator with randomly sampled k vocabulary instances.
New hyperparameter: k

How should we sample?
- Rare words (aardvark) are unlikely to be helpful.
- You can sample a word *w* based on its probability of occurrence:
  - p(w) = U(w) / Z
  - where Z is total number of words
  - U(w) is the unigram frequency of word w (i.e., number of times it appears in the dataset).

Solution: negative sampling

$$\log P(o|c) = -\log \sigma(\boldsymbol{u}_o^T \boldsymbol{v}_c) - \sum_{k \in \{K \text{ sampled indices}\}} \log \sigma(-\boldsymbol{u}_k^T \boldsymbol{v}_c)$$

Replace denominator with randomly sampled k vocabulary instances.
New hyperparameter: k

How should we sample?
- Rare words (aardvark) are unlikely to be helpful.
- You can sample a word *w* based on its probability of occurrence:
  - p(w) = U(w) / Z
  - where Z is total number of words   In practice, p(w) = U(w)$^{3/4}$ / Z
  - U(w) is the unigram frequency of word w (i.e., number of times it appears in the dataset).

Problem of sparse gradients

We iteratively take gradients at each *m* window for SGD
- In each window, we only have at most *2m + 1* words
- plus *2km* negative words with negative sampling
- so the gradient for each update is sparse for V of size 500K

$$\nabla_\theta J_t(\theta) = \begin{bmatrix} 0 \\ \vdots \\ \nabla_{v_{like}} \\ \vdots \\ 0 \\ \nabla_{u_I} \\ \vdots \\ \nabla_{u_{learning}} \\ \vdots \end{bmatrix} \in \mathbb{R}^{2dV}$$

Implementation detail

Most DL packages represent word vectors using a special embedding layer

Rows represent words
- even though we usually talk about words as column vectors
- In implementations, they are row vectors.
- It is a hash table with look up and write functions to avoid writing to the entire V x d matrix.

Indirectly, skim-gram is trying to calculate co-occurance of words

It does this using backprop and by iterating through the entire corpus of text data multiple times.

Can we do better?

Can we build a co-occurrence matrix directly?
- Calculate co-occurrence of words within a window.
- captures some syntactic and semantic information ("word space")
- If window is too large (size of entire articles or documents):
  - Co-occurrence matrix will represent general topics
  - Example, all sports words will have similar entries

Example co-occurrence matrix

- Let's try an example with window length 1 (it is more common to use 5–10)
- Symmetric (irrelevant whether left or right context)

Example corpus:
- I like deep learning
- I like UW
- I enjoy class

counts	I	like	enjoy	deep	learning	UW	class	.
I	0	2	1	0	0	0	0	0
like	2	0	0	1	0	1	0	0
enjoy	1	0	0	0	0	0	1	0
deep	0	1	0	0	1	0	0	0
learning	0	0	0	1	0	0	0	1
UW	0	1	0	0	0	0	0	1
class	0	0	1	0	0	0	0	1
.	0	0	0	0	1	1	1	0

Co-occurrence vectors

Problem with simple count co-occurrence vectors
- Vectors increase in size with vocabulary
- Very high dimensional: require a lot of storage (though sparse)

Idea: Low-dimensional vectors
- Idea: store "most" of the important information in a fixed, small number of dimensions: a dense vector
- Usually 25–1000 dimensions, similar to word2vec
- **But how should we reduce the dimensionality from 500K to <1000?**

Classic Method: Dimensionality Reduction

From linear algebra: Singular Value Decomposition of co-occurrence matrix X

U and V are orthogonal matrices
Σ is a diagonal matrix of singular values.

Classic Method: Dimensionality Reduction

From linear algebra: Singular Value Decomposition of co-occurrence matrix X

We can discard all except the largest d singular values and their corresponding multiplicative values in U and V
New d-dimensional word vector representations are: top-d(U) * top-d(Σ)

Making co-occurrence counts work

Running an SVD on raw counts doesn't work well.

Problem: function words (the, he, has) are too frequent à syntax has too much impact.
Some fixes:
- Use log the frequencies instead
- Limit the maximum values: min(X, t), with t ≈ 100
- Ignore the function words
- Ramped windows that count closer words more than further away words

Interesting semantic patterns emerge
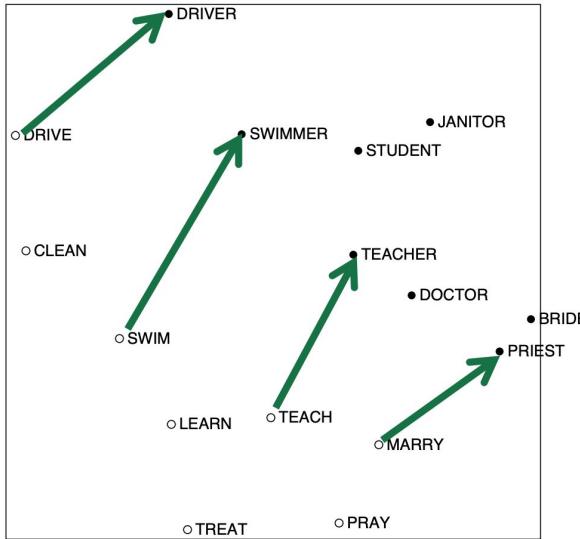
COALS model from Rohde et al. ms., 2005. An Improved Model of Semantic Similarity Based on Lexical Co-Occurrence

Comparing the two methods:

Co-occurrence vectors
- Fast training
  - Single iteration over data
  - SVD is fast as long as vocabulary is reasonable.
- Good for capturing word similarities
- Needs hacks to work
- Not good for anything beyond similarities
- SVD is very slow for large vocabularies

Skip-gram algorithm
- Scales well with increasing vocabulary size
- Good for many other tasks
- Not as good for word similarities
- Needs multiple iterations across the dataset as backprop is slow

Can we combine the strengths of both methods?

Log bilinear model:
- Let every word be a d-dimensional vector
- Remember from skip-gram that dot product is the probability of one word given its context

$$w_i \cdot w_j = \log P(i|j)$$

GloVe [Pennington, Socher, and Manning, EMNLP 2014]: Encoding meaning components in vector differences

Can we combine the strengths of both methods?

Log bilinear model:
- Let every word be a d-dimensional vector
- Remember from skip-gram that dot product is the probability of one word given its context

$$w_i \cdot w_j = \log P(i|j)$$

Main idea: Similarity between two words should be proportional to their co-occurrence count.
- Log of count used as a hack

$$L(\theta) \;=\; \sum_{i,j=1}^{V} f\left(X_{ij}\right)\left(w_i^T \tilde{w}_j + b_i + \tilde{b}_j - \log X_{ij}\right)^2$$

GloVe [Pennington, Socher, and Manning, EMNLP 2014]: Encoding meaning components in vector differences

Can we combine the strengths of both methods?

Log bilinear model:
- Let every word be a d-dimensional vector
- Remember from skip-gram that dot product is the probability of one word given its context

$$w_i \cdot w_j = \log P(i|j)$$

Main idea: Similarity between two words should be proportional to their co-occurrence count.
- Log of count used as a hack

$$L(\theta) = \sum_{i,j=1}^{V} f\left(X_{ij}\right)\left(w_i^T \tilde{w}_j + b_i + \tilde{b}_j - \log X_{ij}\right)^2$$

GloVe [Pennington, Socher, and Manning, EMNLP 2014]: Encoding meaning components in vector differences

Can we combine the strengths of both methods?

Log bilinear model:
- Let every word be a d-dimensional vector
- Remember from skip-gram that dot product is the probability of one word given its context

$$w_i \cdot w_j = \log P(i|j)$$

Main idea: Similarity between two words should be proportional to their co-occurrence count.
- Log of count used as a hack

$$L(\theta) \;=\; \sum_{i,j=1}^{V} f\left(X_{ij}\right)\left(w_i^T \tilde{w}_j + b_i + \tilde{b}_j - \log X_{ij}\right)^2$$

GloVe [Pennington, Socher, and Manning, EMNLP 2014]: Encoding meaning components in vector differences

Can we combine the strengths of both methods?

Log bilinear model:
- Let every word be a d-dimensional vector
- Remember from skip-gram that dot product is the probability of one word given its context

$$w_i \cdot w_j = \log P(i|j)$$

Main idea: Similarity between two words should be proportional to their co-occurrence count.
- Log of count used as a hack
- f() is a threshold for large values

$$L(\theta) = \sum_{i,j=1}^{V} \boxed{f\left(X_{ij}\right)} \left(w_i^T \tilde{w}_j + b_i + \tilde{b}_j - \log X_{ij}\right)^2$$

GloVe [Pennington, Socher, and Manning, EMNLP 2014]: Encoding meaning components in vector differences

Word vectors are very good at analogies

a:b :: c:?

man:woman :: king:?

Word vectors are very good at analogies

a:b :: c:?

man:woman :: king:?

$$d = \arg\max_i \frac{(x_b - x_a + x_c)^T x_i}{||x_b - x_a + x_c||}$$

(man - woman + king)* w_i

Word vectors are very good at analogies

a:b :: c:?

→

$$d = \arg\max_i \frac{(x_b - x_a + x_c)^T x_i}{||x_b - x_a + x_c||}$$

man:woman :: king:?

(man - woman + king)* w_i
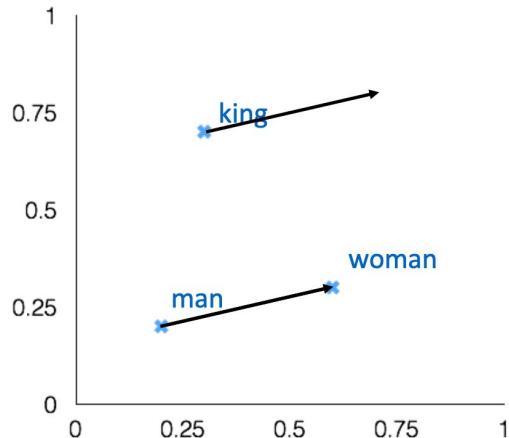
More visualizations
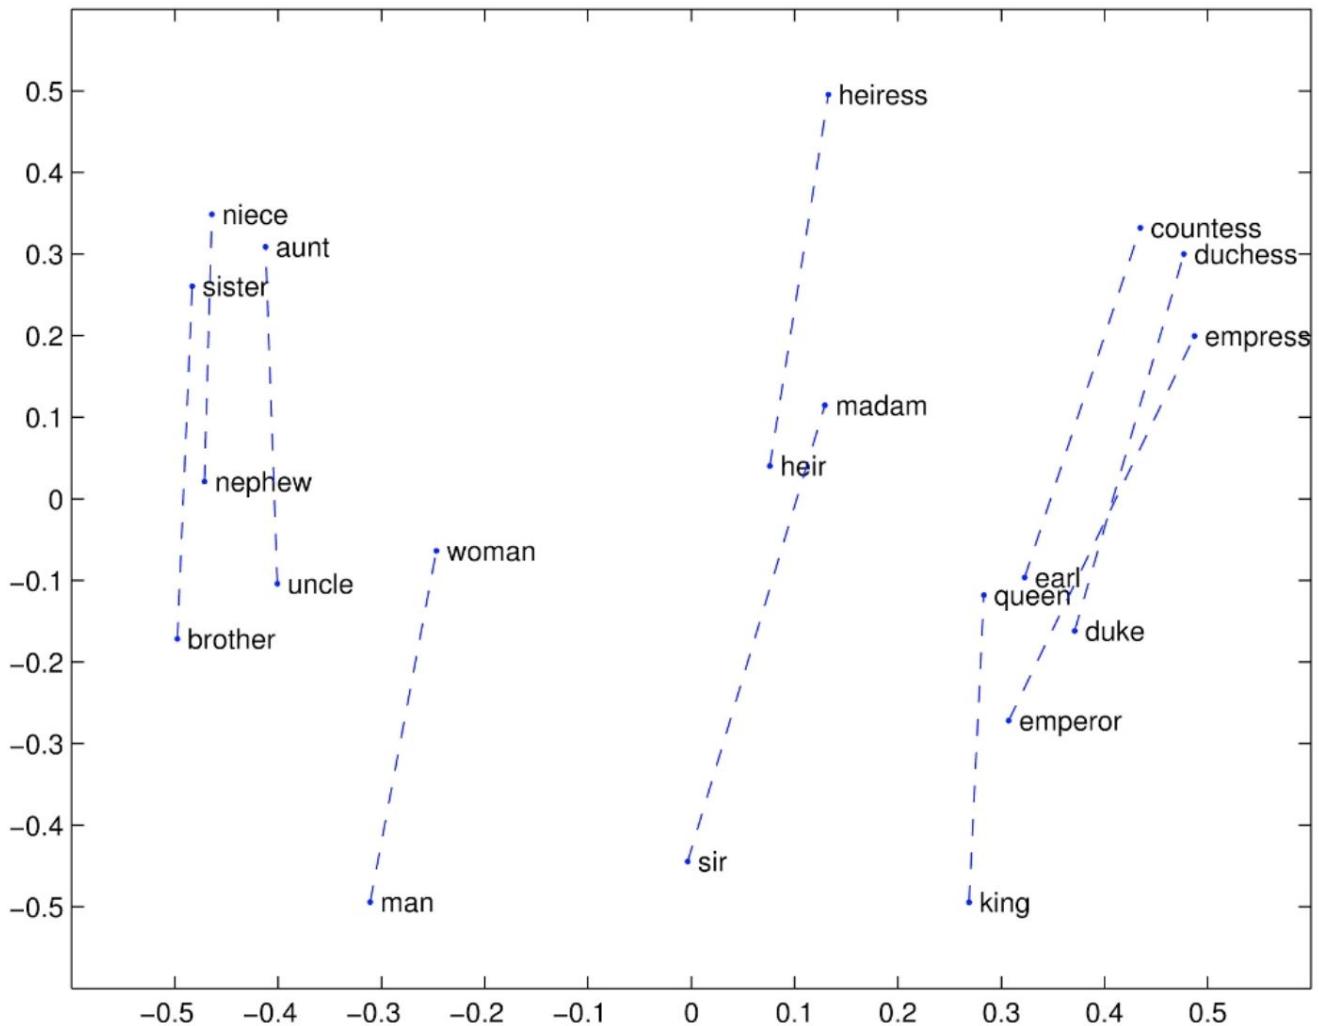
Word vectors correlate with human judgement

Linguists have created a dataset of word similarity judgements
- Word vector distances and their correlation with human judgments
- Example dataset: WordSim353

Word 1	Word 2	Human (mean)
tiger	cat	7.35
tiger	tiger	10
book	paper	7.46
computer	internet	7.58
plane	car	5.77
professor	doctor	6.62
stock	phone	1.62
stock	CD	1.31
stock	jaguar	0.92

Problem of **polysemy**

Word senses and word sense ambiguity

Cap
- Used to only be similar to *Hat*
- But now, thanks to gen z, should also be closer to *Lying*

Can one vector capture all these meanings? Probably not!

Researchers have tried to segregate words into multiple vectors, each with its own meaning
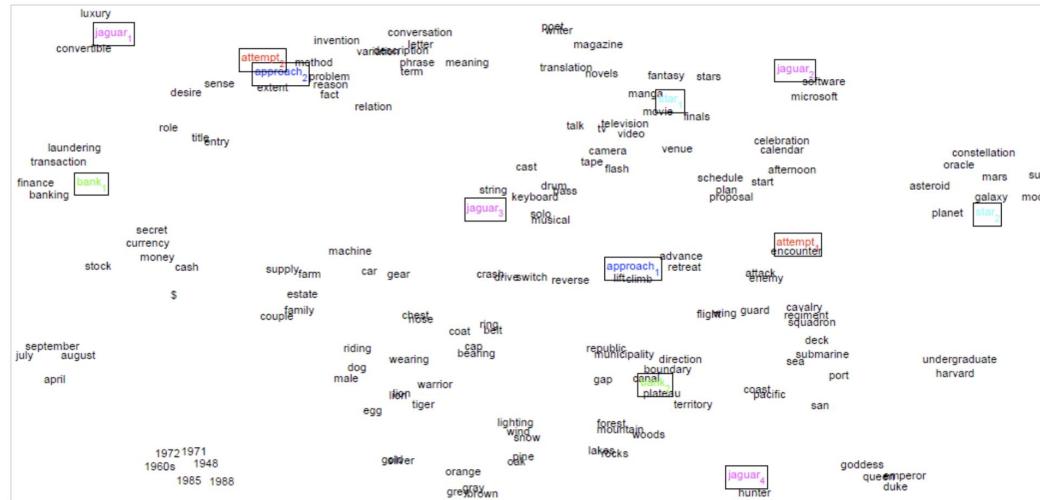
But it doesn't work well. A word's usage in a sentence defines its meaning.

Words should be a function of not just its context but its position in the sentence

-> **Next lecture**

So far: vectors are associated with words

Our vocabulary was comprised of all of the **words in a language**
Problems:

- 500,000 words Webster's English Dictionary (3rd edition)
- Language is changing all of the time
    - 690 words were added to Merriam Webster's in September 2023 ("rizz", "goated", "mid")
- Long tail of infrequent words.
    - Zipf's law: word frequency is inversely proportional to word rank
- Some words may not appear in a training set of documents
- No modeled relationship between words - e.g., "run", "ran", "runs", "runner" are all separate entries despite being linked in meaning
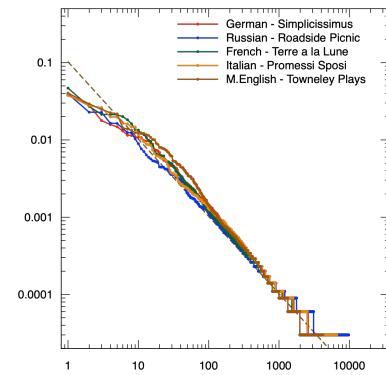
Zipf's Law: Word Rank vs. Word Frequency for Several Languages

Character level vectors instead?

What about assigning a vector to every character instead?

(Maybe add capital letters, punctuation, spaces, …)
Pros:
- Small vocabulary size (for English)
- Complete coverage (unseen words are represented by letters)

Cons:
- Encoding a single sentence becomes very long!
    - # chars instead of # words
- Characters mean very different things in different words!
    - Even worse for representing multiple meanings

Subword tokenization!

How can we combine
1. the **high coverage** of character-level representations
2. with the efficiency of word-level representation?

Subword tokenization! (e.g., **Byte-Pair Encoding**)
- Start with character-level representations
- Build up representations from there

Original BPE Paper (Sennrich et al., 2016)

Example of how Byte-pair encoding works

Let's say our entire dataset contains only these 3 sentences:

$$\mathcal{D} = \{\text{``i hug pugs''},\ \text{``hugging pugs is fun''},\ \text{``i make puns''}\}$$

Example of how Byte-pair encoding works

Let's say our entire dataset contains only these 3 sentences:

$$\mathcal{D} = \{\text{"i hug pugs"}, \text{"hugging pugs is fun"}, \text{"i make puns"}\}$$

Initialize the vocabulary as all the individual characters. Current Vocab:

$$\mathcal{V} = \{\text{' '}, \text{'a'}, \text{'e'}, \text{'f'}, \text{'g'}, \text{'h'}, \text{'i'}, \text{'k'}, \text{'m'},$$
$$\text{'n'}, \text{'p'}, \text{'s'}, \text{'u'}\}, |\mathcal{V}| = 13$$

Example of how Byte-pair encoding works

Let's say our entire dataset contains only these 3 sentences:

$$\mathcal{D} = \{\text{``i hug pugs''}, \text{``hugging pugs is fun''}, \text{``i make puns''}\}$$

Let's split it up into words by splitting right **before the whitespace:**

$$\mathcal{D} = \{\text{``i''}, \text{`` hug''}, \text{`` pugs''}, \text{``hugging''}, \text{`` pugs''},$$
$$\text{`` is''}, \text{`` fun''}, \text{``i''}, \text{`` make''}, \text{`` puns''}\}$$

Example of how Byte-pair encoding works

The vocabulary for reference:

$$\mathcal{V} = \{` \text{'}, `a\text{'}, `e\text{'}, `f\text{'}, `g\text{'}, `h\text{'}, `i\text{'}, `k\text{'}, `m\text{'},$$
$$`n\text{'}, `p\text{'}, `s\text{'}, `u\text{'}\}, |\mathcal{V}| = 13$$

Let's split it up into words by splitting right **before the whitespace:**

$$\mathcal{D} = \{\text{``i''}, \text{`` hug''}, \text{`` pugs''}, \text{``hugging''}, \text{`` pugs''},$$
$$\text{`` is''}, \text{`` fun''}, \text{``i''}, \text{`` make''}, \text{`` puns''}\}$$

Example of how Byte-pair encoding works

The vocabulary for reference:

$$\mathcal{V} = \{\ '\ ',\ 'a',\ 'e',\ 'f',\ 'g',\ 'h',\ 'i',\ 'k',\ 'm',$$
$$'n',\ 'p',\ 's',\ 'u'\},\ |\mathcal{V}| = 13$$

Let's represent the dataset with only vocabulary elements:

$$\mathcal{D} = \{\ ['i'],\ [\ '\ ',\ 'h',\ 'u',\ 'g'],\ [\ '\ ',\ 'p',\ 'u',\ 'g',\ 's'],$$
$$['h',\ 'u',\ 'g',\ 'g',\ 'i',\ 'n',\ 'g'],\ [\ '\ ',\ 'p',\ 'u',\ 'g',\ 's'],$$
$$[\ '\ ',\ 'i',\ 's'],\ [\ '\ ',\ 'f',\ 'u',\ 'n'],\ ['i'],$$
$$[\ '\ ',\ 'm',\ 'a',\ 'k',\ 'e'],\ [\ '\ ',\ 'p',\ 'u',\ 'n',\ 's']\}$$

Example of how Byte-pair encoding works

The vocabulary for reference:

$$\mathcal{V} = \{\text{`\,'}, \text{`a'}, \text{`e'}, \text{`f'}, \text{`g'}, \text{`h'}, \text{`i'}, \text{`k'}, \text{`m'},$$
$$\text{`n'}, \text{`p'}, \text{`s'}, \text{`u'}\}, |\mathcal{V}| = 13$$

Now, let's find the most common bi-gram

$$\mathcal{D} = \{\,[\text{`i'}]\,, [\text{`\,'}, \text{`h'}, \text{`u'}, \text{`g'}]\,, [\text{`\,'}, \text{`p'}, \text{`u'}, \text{`g'}, \text{`s'}]\,,$$
$$[\text{`h'}, \text{`u'}, \text{`g'}, \text{`g'}, \text{`i'}, \text{`n'}, \text{`g'}]\,, [\text{`\,'}, \text{`p'}, \text{`u'}, \text{`g'}, \text{`s'}]\,,$$
$$[\text{`\,'}, \text{`i'}, \text{`s'}]\,, [\text{`\,'}, \text{`f'}, \text{`u'}, \text{`n'}]\,, [\text{`i'}]\,,$$
$$[\text{`\,'}, \text{`m'}, \text{`a'}, \text{`k'}, \text{`e'}]\,, [\text{`\,'}, \text{`p'}, \text{`u'}, \text{`n'}, \text{`s'}]\}$$

Example of how Byte-pair encoding works

The vocabulary for reference:

$$\mathcal{V} = \{\text{' '}, \text{'a'}, \text{'e'}, \text{'f'}, \text{'g'}, \text{'h'}, \text{'i'}, \text{'k'}, \text{'m'},$$
$$\text{'n'}, \text{'p'}, \text{'s'}, \text{'u'}\}, |\mathcal{V}| = 13$$

Now, let's find the most common bi-gram

$$\mathcal{D} = \{\,[\text{'i'}], [\text{' '}, \text{'h'}, \text{'u'}, \text{'g'}], [\text{' '}, \text{'p'}, \text{'u'}, \text{'g'}, \text{'s'}],$$
$$[\text{'h'}, \text{'u'}, \text{'g'}, \text{'g'}, \text{'i'}, \text{'n'}, \text{'g'}], [\text{' '}, \text{'p'}, \text{'u'}, \text{'g'}, \text{'s'}],$$
$$[\text{' '}, \text{'i'}, \text{'s'}], [\text{' '}, \text{'f'}, \text{'u'}, \text{'n'}], [\text{'i'}],$$

Create new vocab:

$$v_{14} := \text{concat}(\text{'u'}, \text{'g'}) = \text{'ug'}$$
$$[\text{' '}, \text{'m'}, \text{'a'}, \text{'k'}, \text{'e'}], [\text{' '}, \text{'p'}, \text{'u'}, \text{'n'}, \text{'s'}]\}$$

Example of how Byte-pair encoding works

Update vocabulary with new vocab v_{14}:

$$\mathcal{V} = \{\text{` '}, \text{`a'}, \text{`e'}, \text{`f'}, \text{`g'}, \text{`h'}, \text{`i'}, \text{`k'}, \text{`m'},$$
$$\text{`n'}, \text{`p'}, \text{`s'}, \text{`u'}, \text{`ug'}\}, |\mathcal{V}| = 14$$

Update dataset by replace bigram with new vocab v_{14}:

$$\mathcal{D} = \{\,[\text{`i'}], [\text{` '}, \text{`h'}, \text{`ug'}], [\text{` '}, \text{`p'}, \text{`ug'}, \text{`s'}],$$

Create new vocab:

$$[\text{`h'}, \text{`ug'}, \text{`g'}, \text{`i'}, \text{`n'}, \text{`g'}], [\text{` '}, \text{`p'}, \text{`ug'}, \text{`s'}],$$
$$[\text{` '}, \text{`i'}, \text{`s'}], [\text{` '}, \text{`f'}, \text{`u'}, \text{`n'}], [\text{`i'}],$$

$$v_{14} := \text{concat}(\text{`u'}, \text{`g'}) = \text{`ug'}$$

$$[\text{` '}, \text{`m'}, \text{`a'}, \text{`k'}, \text{`e'}], [\text{` '}, \text{`p'}, \text{`u'}, \text{`n'}, \text{`s'}]\}$$

Example of how Byte-pair encoding works

Current vocabulary:

$$\mathcal{V} = \{\text{' ', 'a', 'e', 'f', 'g', 'h', 'i', 'k', 'm',}$$
$$\text{'n', 'p', 's', 'u', 'ug'}\}, |\mathcal{V}| = 14$$

Find the next common bigram:

$$\mathcal{D} = \{\,[\text{'i'}]\,, [\text{' ', 'h', 'ug'}]\,, [\text{' ', 'p', 'ug', 's'}]\,,$$
$$[\text{'h', 'ug', 'g', 'i', 'n', 'g'}]\,, [\text{' ', 'p', 'ug', 's'}]\,,$$
$$[\text{' ', 'i', 's'}]\,, [\text{' ', 'f', 'u', 'n'}]\,, [\text{'i'}]\,,$$
$$[\text{' ', 'm', 'a', 'k', 'e'}]\,, [\text{' ', 'p', 'u', 'n', 's'}]\}$$

Example of how Byte-pair encoding works

Current vocabulary:

$$\mathcal{V} = \{\text{' ', 'a', 'e', 'f', 'g', 'h', 'i', 'k', 'm',}$$
$$\text{'n', 'p', 's', 'u', 'ug'}\}, |\mathcal{V}| = 14$$

Find the next common bigram:

$$\mathcal{D} = \{\, ['i'], [\text{' '}, 'h', \text{'ug'}], [\text{' '}, 'p', \text{'ug'}, 's'],$$

Create new vocab:

$$['h', \text{'ug'}, 'g', 'i', 'n', 'g'], [\text{' '}, 'p', \text{'ug'}, 's'],$$

$$[\text{' '}, 'i', 's'], [\text{' '}, 'f', 'u', 'n'], ['i'],$$

$$v_{15} := \mathrm{concat}(\text{' '}, 'p') = \text{' p'}$$

$$[\text{' '}, 'm', 'a', 'k', 'e'], [\text{' '}, 'p', 'u', 'n', 's']\}$$

Example of how Byte-pair encoding works

Update vocabulary with new vocab v_{15}:

$$\mathcal{V} = \{\text{' '}, \text{'a'}, \text{'e'}, \text{'f'}, \text{'g'}, \text{'h'}, \text{'i'}, \text{'k'}, \text{'m'},$$
$$\text{'n'}, \text{'p'}, \text{'s'}, \text{'u'}, \text{'ug'}, \text{' p}\}, |\mathcal{V}| = 15$$

Update dataset by replace bigram with new vocab v_{15}:

$$\mathcal{D} = \{\,[\text{'i'}]\,,[\text{' '}, \text{'h'}, \text{'ug'}]\,,[\text{' p'}, \text{'ug'}, \text{'s'}]\,,$$

Create new vocab:

$$[\text{'h'}, \text{'ug'}, \text{'g'}, \text{'i'}, \text{'n'}, \text{'g'}]\,,[\text{' p'}, \text{'ug'}, \text{'s'}]\,,$$
$$[\text{' '}, \text{'i'}, \text{'s'}]\,,[\text{' '}, \text{'f'}, \text{'u'}, \text{'n'}]\,,[\text{'i'}]\,,$$

$$v_{15} := \text{concat}(\text{' '}, \text{'p'}) = \text{' p'}$$

$$[\text{' '}, \text{'m'}, \text{'a'}, \text{'k'}, \text{'e'}]\,,[\text{' p'}, \text{'u'}, \text{'n'}, \text{'s'}]\}$$

Repeat until vocab size reaches the amount you want (20 for example)

Final vocabulary:

$$\mathcal{V} = \{\text{' '}, \text{'a'}, \text{'e'}, \text{'f'}, \text{'g'}, \text{'h'}, \text{'i'}, \text{'k'}, \text{'m'}, \text{'n'}, \text{'p'}, \text{'s'}, \text{'u'},$$
$$\text{'ug'}, \text{' p'}, \text{'hug'}, \text{' pug'}, \text{' pugs'}, \text{'un'}, \text{' hug'}\},$$

Final dataset:

$$\mathcal{D} = \{\; [\text{'i'}], [\text{' hug'}], [\text{' pugs'}],$$
$$[\text{'hug'}, \text{'g'}, \text{'i'}, \text{'n'}, \text{'g'}], [\text{' pugs'}],$$
$$[\text{' '}, \text{'i'}, \text{'s'}], [\text{' '}, \text{'f'}, \text{'un'}], [\text{'i'}],$$
$$[\text{' '}, \text{'m'}, \text{'a'}, \text{'k'}, \text{'e'}], [\text{' p'}, \text{'un'}, \text{'s'}]\}$$

With this vocabulary, can you represent (or, tokenize/encode):

Q: Can you encode "apple"?

$$\mathcal{V} = \{1 : \text{' '}, 2 : \text{'a'}, 3 : \text{'e'}, 4 : \text{'f'}, 5 : \text{'g'}, 6 : \text{'h'}, 7 : \text{'i'},$$
$$8 : \text{'k'}, 9 : \text{'m'}, 10 : \text{'n'}, 11 : \text{'p'}, 12 : \text{'s'}, 13 : \text{'u'},$$
$$14 : \text{'ug'}, 15 : \text{' p'}, 16 : \text{'hug'}, 17 : \text{' pug'}, 18 : \text{' pugs'},$$
$$19 : \text{'un'}, 20 : \text{' hug'}\}$$

With this vocabulary, can you represent (or, tokenize/encode):

Q: Can you encode "apple"?
- No, there is no 'l' in the vocabulary

$$\mathcal{V} = \{1 : \text{' '}, 2 : \text{'a'}, 3 : \text{'e'}, 4 : \text{'f'}, 5 : \text{'g'}, 6 : \text{'h'}, 7 : \text{'i'},$$
$$8 : \text{'k'}, 9 : \text{'m'}, 10 : \text{'n'}, 11 : \text{'p'}, 12 : \text{'s'}, 13 : \text{'u'},$$
$$14 : \text{'ug'}, 15 : \text{' p'}, 16 : \text{'hug'}, 17 : \text{' pug'}, 18 : \text{' pugs'},$$
$$19 : \text{'un'}, 20 : \text{' hug'}\}$$

With this vocabulary, can you represent (or, tokenize/encode):

Q: Can you encode "apple"?
- No, there is no 'l' in the vocabulary

Q: "map"?

$$\mathcal{V} = \{1 : `\ ', 2 : `a', 3 : `e', 4 : `f', 5 : `g', 6 : `h', 7 : `i',$$
$$8 : `k', 9 : `m', 10 : `n', 11 : `p', 12 : `s', 13 : `u',$$
$$14 : `ug', 15 : `\ p', 16 : `hug', 17 : `\ pug', 18 : `\ pugs',$$
$$19 : `un', 20 : `\ hug'\}$$

With this vocabulary, can you represent (or, tokenize/encode):

Q: Can you encode "apple"?
- No, there is no 'l' in the vocabulary

Q: "map"?
Yes - [9,2,11]

$\mathcal{V} = \{1 : `\ ', 2 : `a', 3 : `e', 4 : `f', 5 : `g', 6 : `h', 7 : `i',$
$8 : `k', 9 : `m', 10 : `n', 11 : `p', 12 : `s', 13 : `u',$
$14 : `ug', 15 : `\ p', 16 : `hug', 17 : `\ pug', 18 : `\ pugs',$
$19 : `un', 20 : `\ hug'\}$

With this vocabulary, can you represent (or, tokenize/encode):

Q: Can you encode "apple"?
- No, there is no 'l' in the vocabulary

Q: "map"?
Yes - [9,2,11]

Q: "huge"?

$$\mathcal{V} = \{1 : \text{' '}, 2 : \text{'a'}, 3 : \text{'e'}, 4 : \text{'f'}, 5 : \text{'g'}, 6 : \text{'h'}, 7 : \text{'i'},$$
$$8 : \text{'k'}, 9 : \text{'m'}, 10 : \text{'n'}, 11 : \text{'p'}, 12 : \text{'s'}, 13 : \text{'u'},$$
$$14 : \text{'ug'}, 15 : \text{' p'}, 16 : \text{'hug'}, 17 : \text{' pug'}, 18 : \text{' pugs'},$$
$$19 : \text{'un'}, 20 : \text{' hug'}\}$$

With this vocabulary, can you represent (or, tokenize/encode):

Q: Can you encode "apple"?
- No, there is no 'l' in the vocabulary

Q: "map"?
Yes - [9,2,11]

Q: "huge"?
Yes - [16, 4] or [7,14,4] or [7,13,6,4]

$\mathcal{V} = \{1 : ` `, 2 : `a`, 3 : `e`, 4 : `f`, 5 : `g`, 6 : `h`, 7 : `i`,$
$8 : `k`, 9 : `m`, 10 : `n`, 11 : `p`, 12 : `s`, 13 : `u`,$
$14 : `ug`, 15 : ` p`, 16 : `hug`, 17 : ` pug`, 18 : ` pugs`,$
$19 : `un`, 20 : ` hug`\}$

With this vocabulary, can you represent (or, tokenize/encode):

Q: Can you encode "apple"?
- No, there is no 'l' in the vocabulary

$$\mathcal{V} = \{1 : `\ ', 2 : `a', 3 : `e', 4 : `f', 5 : `g', 6 : `h', 7 : `i',$$
$$8 : `k', 9 : `m', 10 : `n', 11 : `p', 12 : `s', 13 : `u',$$

Q: "map"?
Yes - [9,2,11]

$$14 : `ug', 15 : `\ p', 16 : `hug', 17 : `\ pug', 18 : `\ pugs',$$
$$19 : `un', 20 : `\ hug'\}$$

Q: "huge"?
Yes - [16, 4] or [7,14,4] or [7,13,6,4]

Q: " huge" with a space in the front?

With this vocabulary, can you represent (or, tokenize/encode):

Q: Can you encode "apple"?
- No, there is no 'l' in the vocabulary

$$\mathcal{V} = \{1 : ` \text{'}, 2 : `a\text{'}, 3 : `e\text{'}, 4 : `f\text{'}, 5 : `g\text{'}, 6 : `h\text{'}, 7 : `i\text{'},$$
$$8 : `k\text{'}, 9 : `m\text{'}, 10 : `n\text{'}, 11 : `p\text{'}, 12 : `s\text{'}, 13 : `u\text{'},$$
$$14 : `ug\text{'}, 15 : ` p\text{'}, 16 : `hug\text{'}, 17 : ` pug\text{'}, 18 : ` pugs\text{'},$$
$$19 : `un\text{'}, 20 : ` hug\text{'}\}$$

Q: "map"?
Yes - [9,2,11]

Q: "huge"?
Yes - [16, 4] or [7,14,4] or [7,13,6,4]

Q: " huge" with a space in the front?
Yes - [18, 4]

Benefits of Byte-pair encoding

1. Efficient to run (greedy vs. global optimization)
2. Lossless compression
3. Potentially some shared representations
    a. e.g., the token "hug" could be used both in "hug" and "hugging"

Byte-pair encoding - ChatGPT Example

Call me Ishmael. Some years ago—never mind how long precisely—having little or no money in my purse, and nothing particular to interest me on shore, I thought I would sail about a little and see the watery part of the world. It is a way I have of driving off the spleen and regulating the circulation. Whenever I find myself growing grim about the mouth; whenever it is a damp, drizzly November in my soul; whenever I find myself involuntarily pausing before coffin warehouses, and bringing up the rear of every funeral I meet; and especially whenever my hypos get such an upper hand of me, that it requires a strong moral principle to prevent me from deliberately stepping into the street, and methodically knocking people's hats off—then, I account it high time tozz get to sea as soon as I can. This is my substitute for pistol and ball. With a philosophical flourish Cato throws himself upon his sword; I quietly take to the ship. There is nothing surprising in this. If they but knew it, almost all men in their degree, some time or other, cherish very nearly the same feelings towards the ocean with me.

TEXT   TOKEN IDS

Tokens   Characters
239   1109

[7368, 757, 57704, 1764, 301, 13, 4427, 1667, 4227, 2345, 37593, 4059, 1268, 1317, 24559, 2345, 69666, 2697, 477, 912, 3300, 304, 856, 53101, 11, 323, 4400, 4040, 311, 2802, 757, 389, 31284, 11, 358, 3463, 358, 1053, 30503, 922, 264, 2697, 323, 1518, 279, 30125, 727, 961, 315, 279, 1917, 13, 1102, 374, 264, 1648, 358, 617, 315, 10043, 1022, 279, 87450, 268, 323, 58499, 279, 35855, 13, 43633, 358, 1505, 7182, 7982, 44517, 922, 279, 11013, 26, 15716, 433, 374, 264, 41369, 11, 1377, 73825, 6841, 304, 856, 13836, 26, 15716, 358, 1505, 7182, 4457, 3935, 6751, 7251, 985, 1603, 78766, 83273, 11, 323, 12967, 709, 279, 14981, 315, 1475, 32079, 358, 3449, 26, 323, 5423, 15716, 856, 6409, 981, 636, 1778, 459, 8582, 1450, 315, 757, 11, 430, 433, 7612, 264, 3831, 16033, 17966, 311, 5471, 757, 505, 36192, 36567, 1139, 279, 8761, 11, 323, 1749, 2740, 50244, 1274, 753, 45526, 1022, 2345, 3473, 11, 358, 2759, 433, 1579, 892, 311, 10616, 636, 311, 9581, 439, 5246, 439, 358, 649, 13, 1115, 374, 856, 28779, 369, 40536, 323, 5041, 13, 3161, 264, 41903, 67784, 356, 4428, 3872, 5678, 5304, 813, 20827, 26, 358, 30666, 1935, 311, 279, 8448, 13, 2684, 374, 4400, 15206, 304, 420, 13, 1442, 814, 719, 7020, 433, 11, 4661, 682, 3026, 304, 872, 8547, 11, 1063, 892, 477, 1023, 11, 87785, 1888, 7154, 279, 1890, 16024, 7119, 279, 18435, 449, 757, 13]

TEXT   TOKEN IDS

Weird properties of tokenizers

Token != word
- Spaces are part of token
- "run" is a different token than " run"
- Not invariant to case changes
- "Run" is a different token than "run"

`run run RunRun`

`[6236, 1629, 6588, 6869]`

TEXT   TOKEN IDS

TEXT   TOKEN IDS
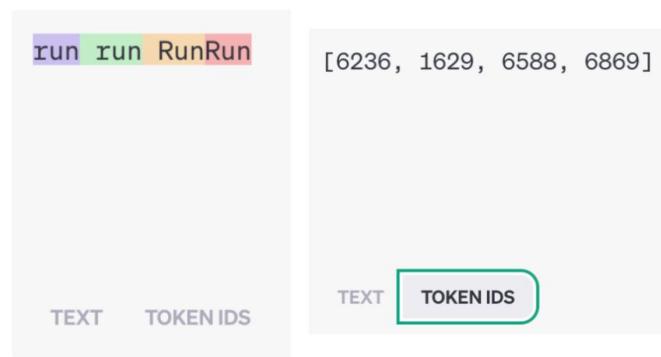
Weird properties of tokenizers

Token != word
- Spaces are part of token
- "run" is a different token than " run"
- Not invariant to case changes
- "Run" is a different token than "run"
- Tokenization fits statistics of your data
  - e.g., while these words are multiple tokens…

Weird properties of tokenizers

Token != word
- Spaces are part of token
- "run" is a different token than " run"
- Not invariant to case changes
- "Run" is a different token than "run"
- Tokenization fits statistics of your data
  - e.g., while these words are multiple tokens…

These words are all 1 token in GPT-3's tokenizer!
Does anyone know why?

attRot
EStreamFrame
 SolidGoldMagikarp
PsyNetMessage
embedreportprint
 Adinida
oreAndOnline
StreamerBot
GoldMagikarp
 externalToEVA
 TheNitrome
 TheNitromeFan
RandomRedditorWithNo
InstoreAndOnline

TEXT   TOKEN IDS

Next time:
RNNs & LSTMs