Psychology of Language

11 Computational models I

Fall 2023
Tues/Thur 5:00-6:15pm

Emma Wing
Drop-in hours:
Wednesdays 3-4pm
& by appointment
[Webex link]

Thanks to Wesley Leong for content and slide inspiration!

Road map

- Review from **10 Discourse processing**
- Unit 2: The Mature System

  11 Computational models I

Review from 10 Discourse processing

- ✓Information flow during language comprehension
- ✓Examples of discourse
- ✓The role of real-world knowledge and inference
- ✓The role of prediction
- ✓Anaphoric reference
- ✓Discourse focus
- ✓Types of discourse representation across time
- ✓Mental models

Learning objectives

- Discuss why we are even looking at computational models of language
- Define semantic space and 1+ ways to create a semantic space
- Describe how Latent Semantic Analysis (LSA) models similarities between words
- Describe how Word2Vec models similarities between words

Why computational models?

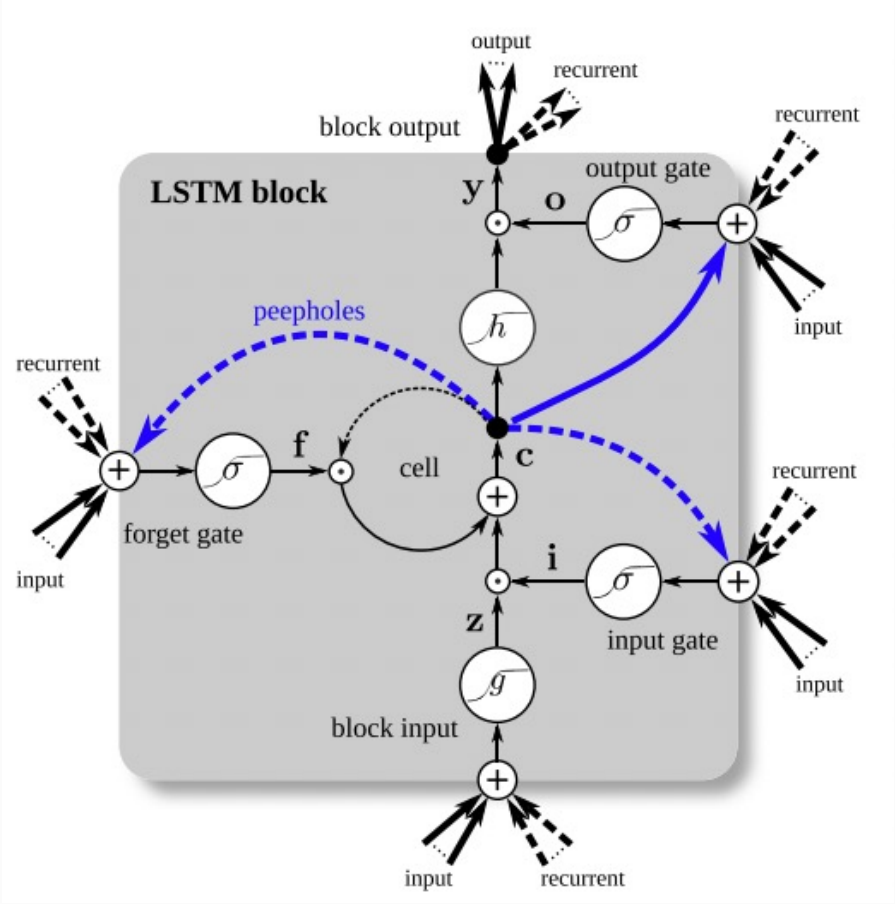

??

Greff et al. (2017)

Why computational models?

"All models are wrong, but some are useful."

-George Box

Why computational models?

- They help establish the principles by which the brain *might* “compute” meaning
  - Some parts of the model might constitute a theory about how, or what, the brain computes
  - Others may just be there to “make it work”
- Contemporary “neural networks” are based on a combination of principles drawn from brain function and models of animal learning

What is meaning?

- What is meaning?
- How might a computer understand what words mean?

- It’s hard for words to mean anything on their own.
    - We might have to place them in comparison with other words

- Rephrase: What behavior might we hope a computer would show?
    - Some representation of **similarity between words**
    - Some understanding of **semantic relationships**
        - Example: dog is to animal as apple is to fruit

Semantic space

- How would you organize these words on a whiteboard?

Cat
Blouse
Window
Treat

Table
Candy
Dog
Apple

- We might want a representation that captures where these words are *in relation* to one another – a semantic space!

Semantic space

window

cat

table

dog

treat

blouse

apple

candy

An example placement

Semantic space

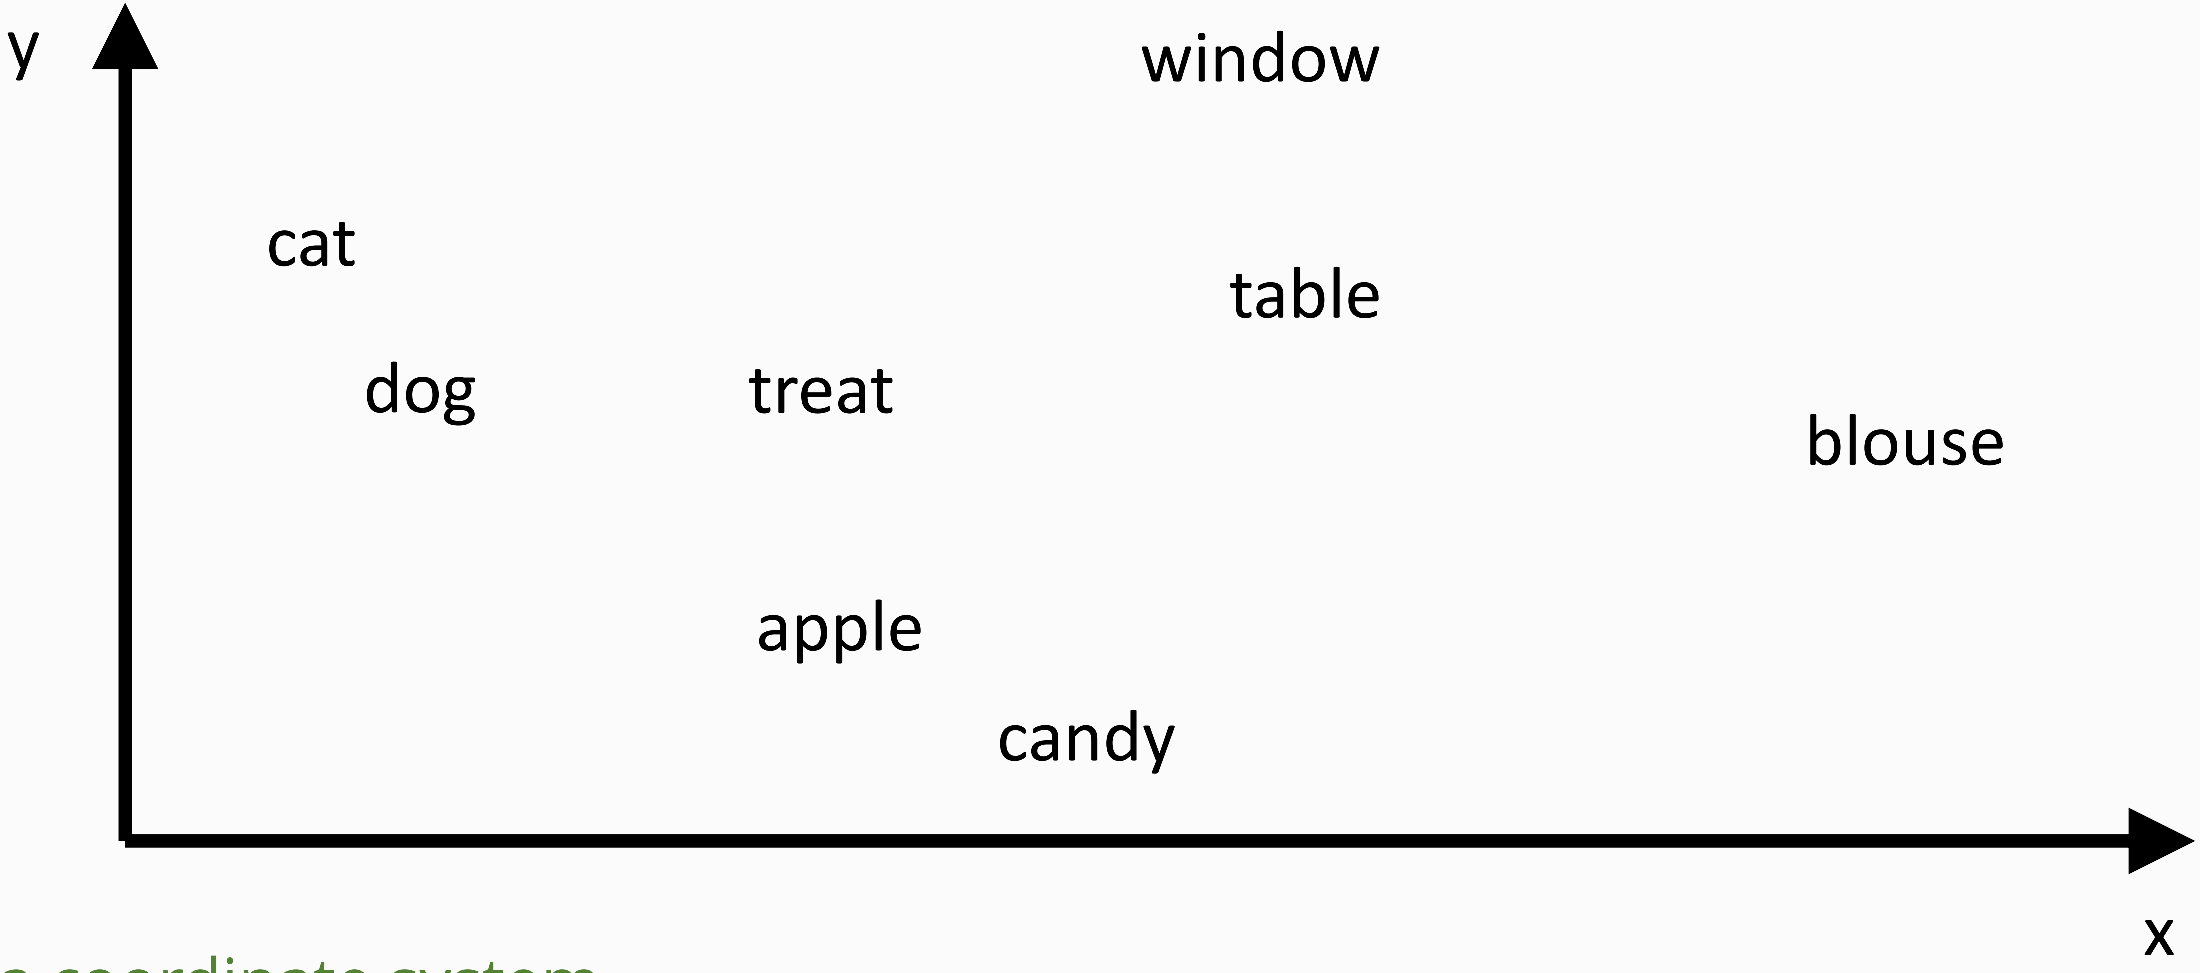

Add a coordinate system

Semantic space

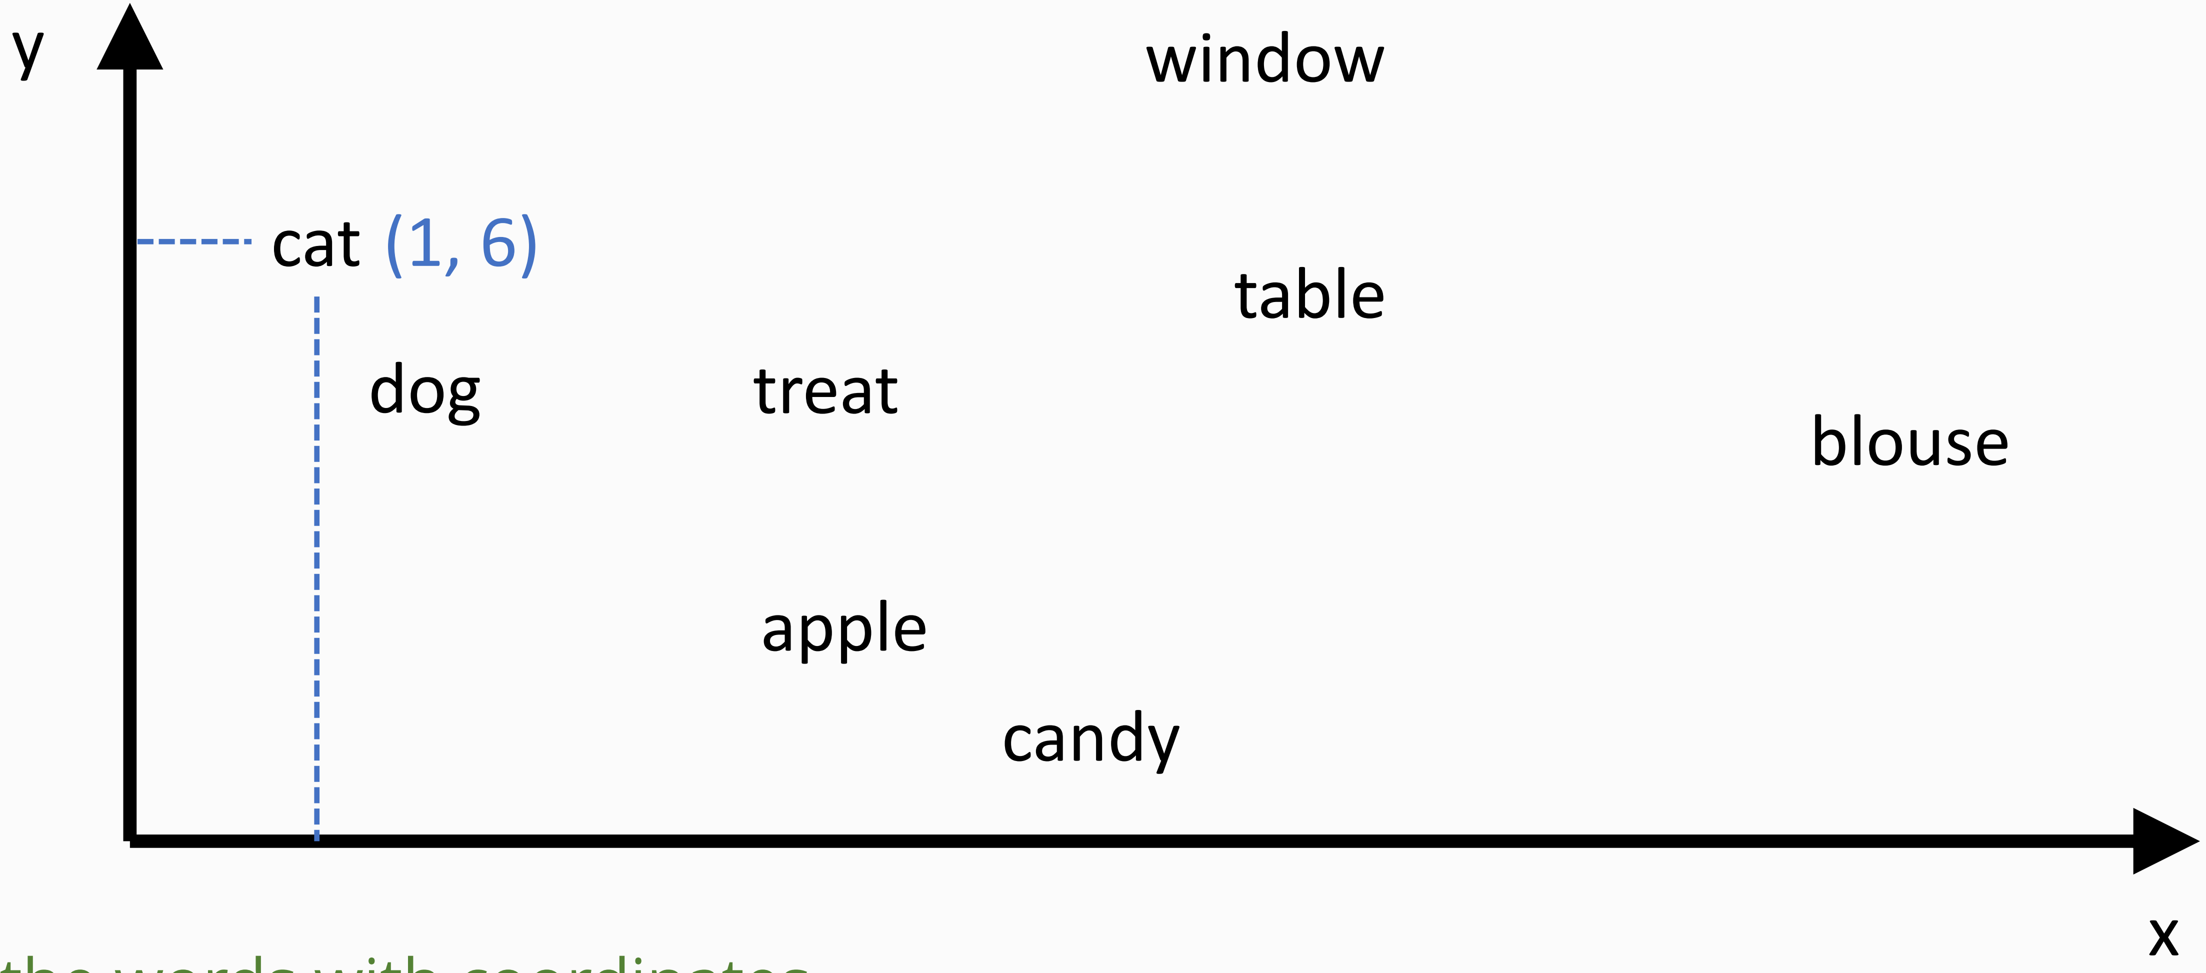

Plot the words with coordinates

Semantic space

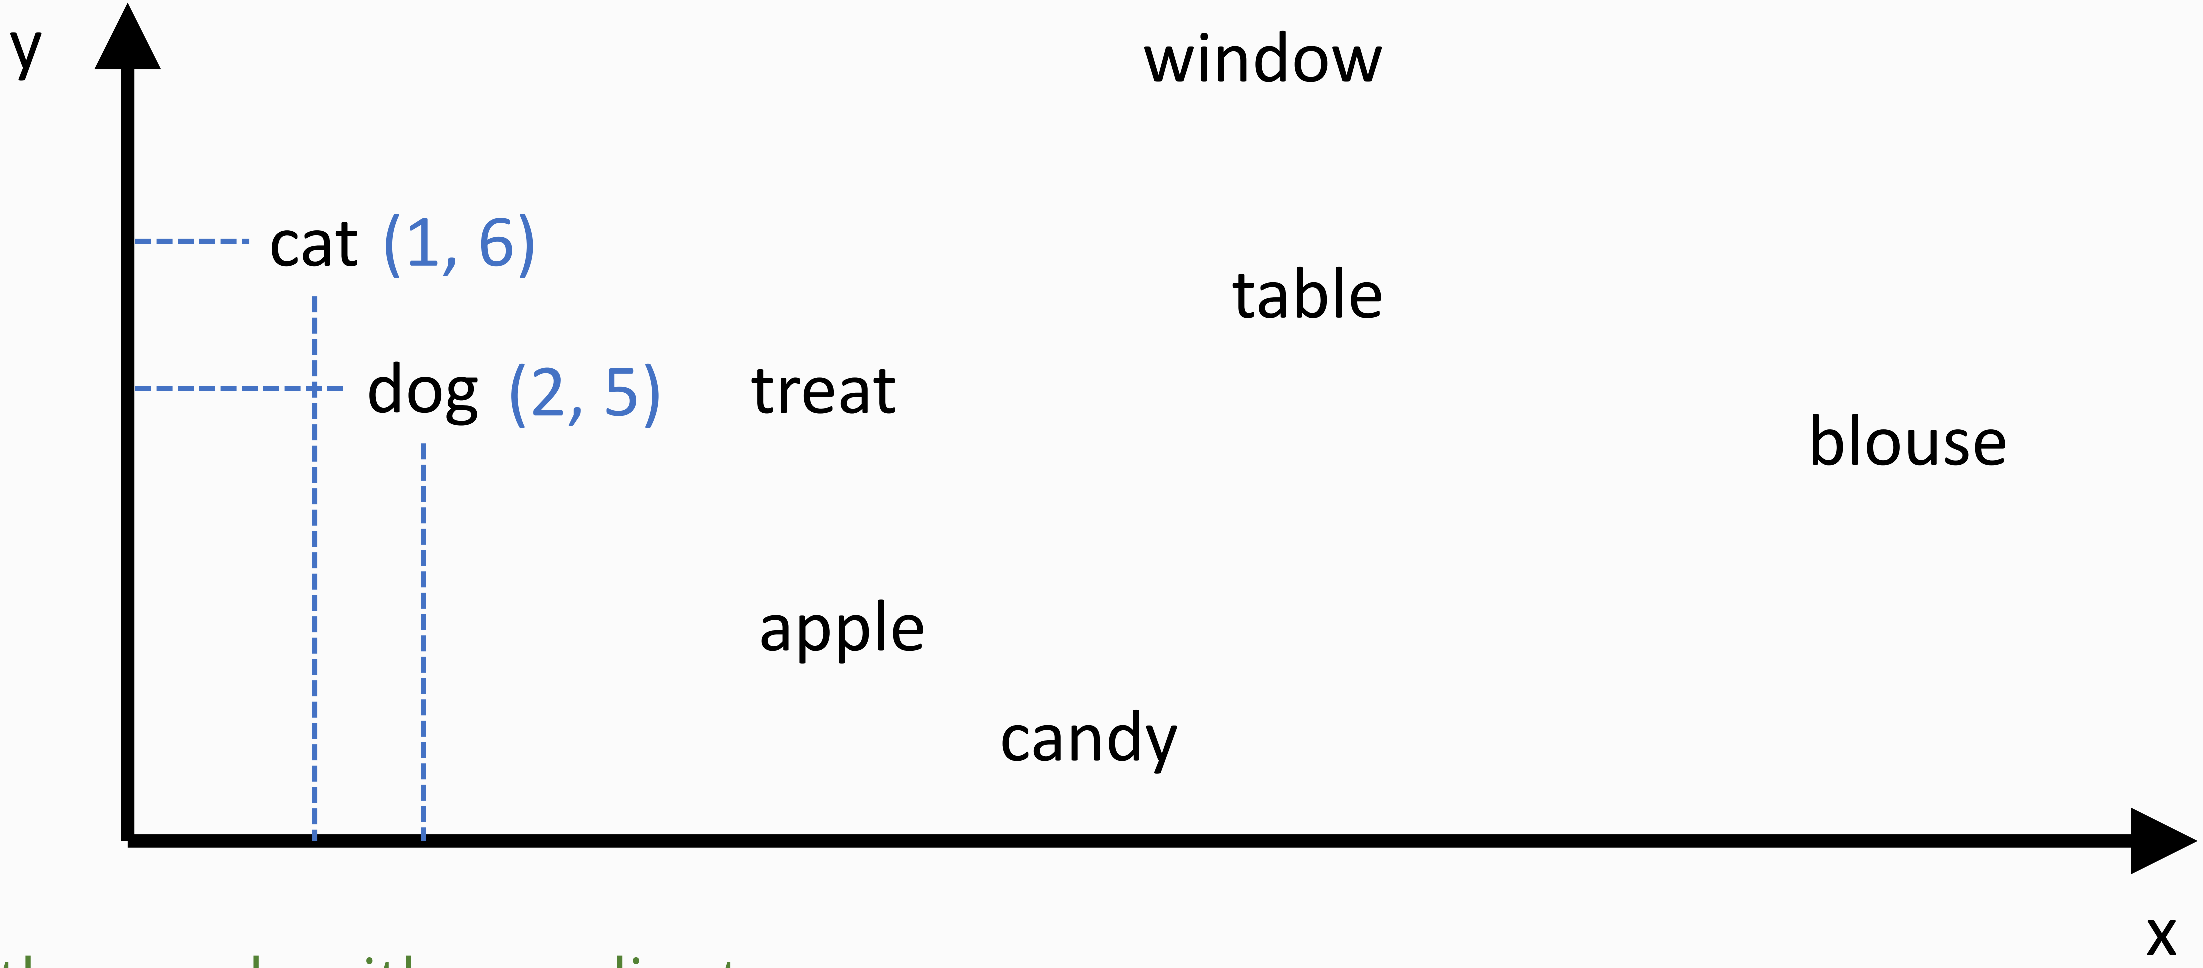

Plot the words with coordinates

Semantic space

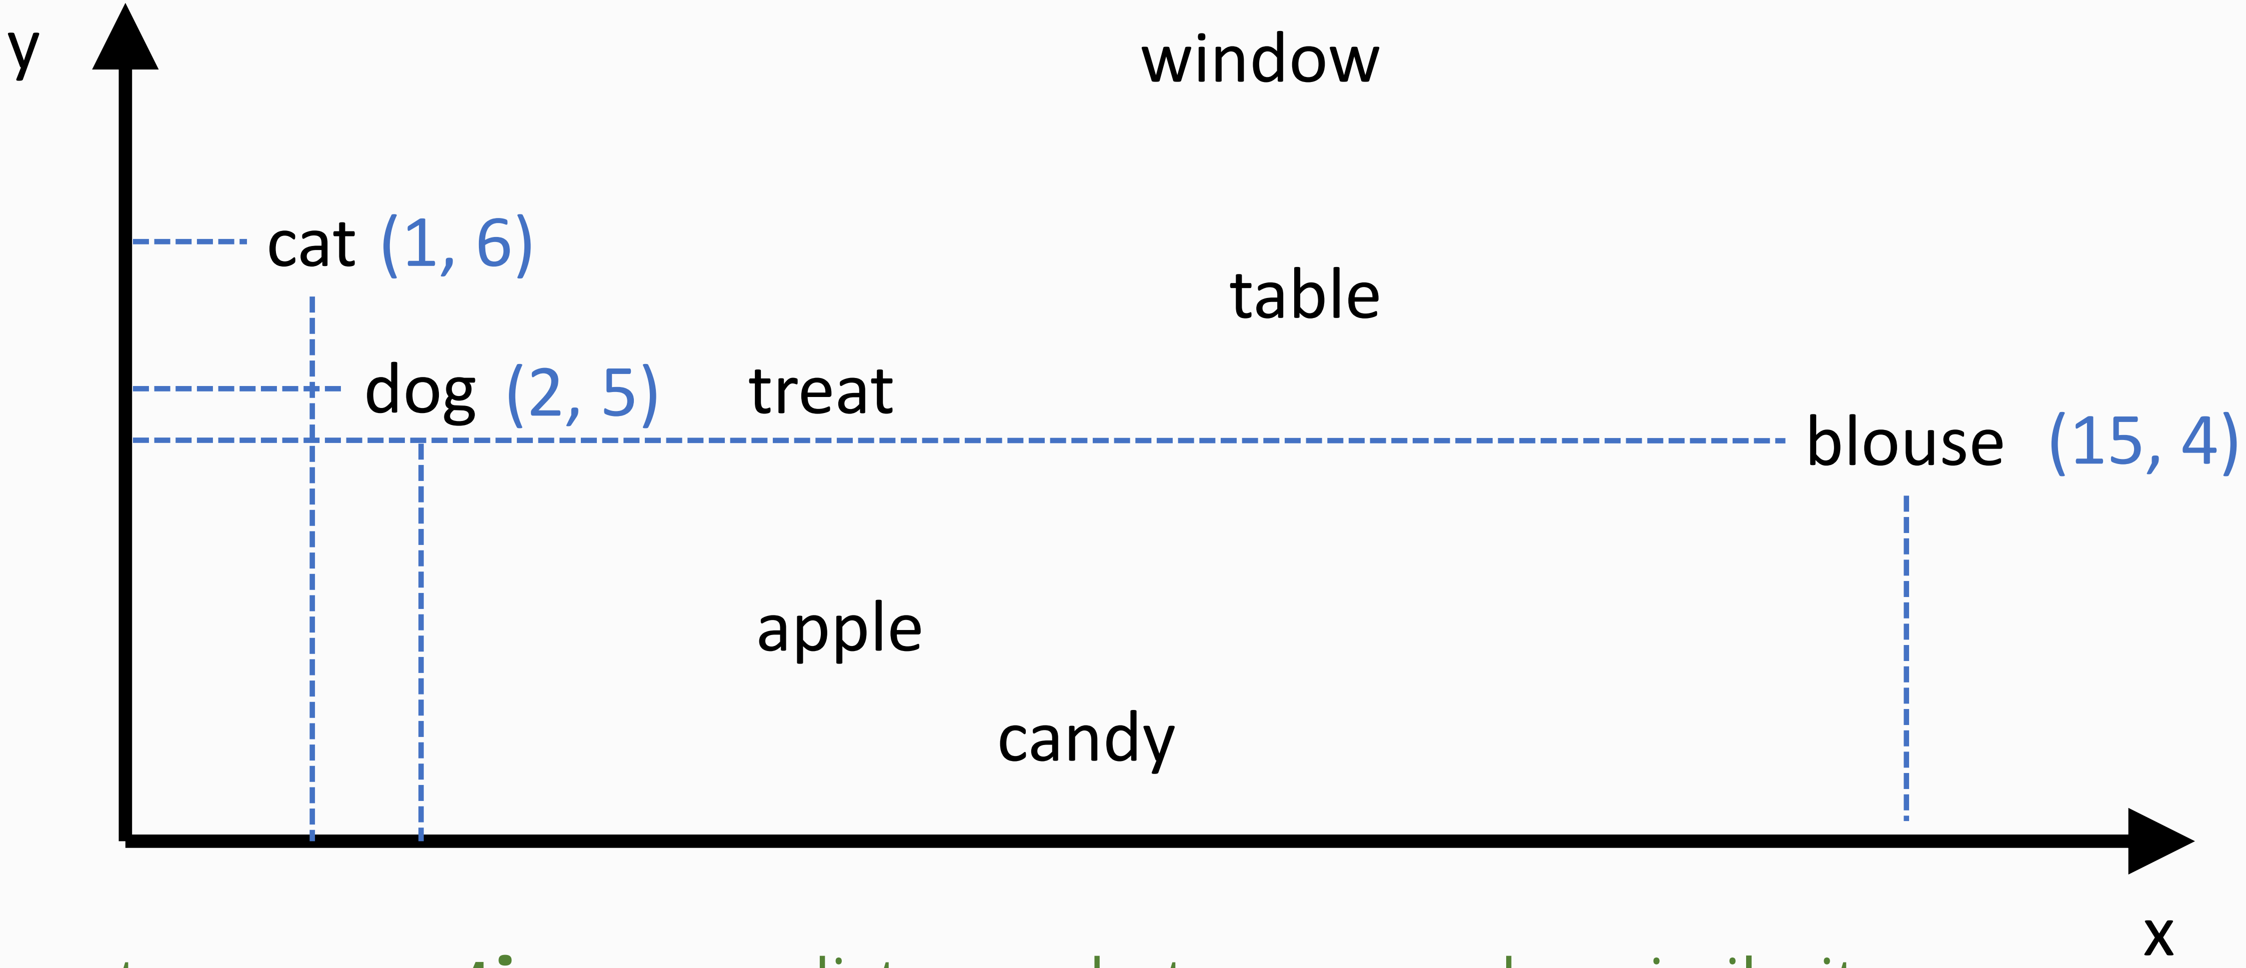

...to create a **semantic space:** distances between words = similarity

Semantic space

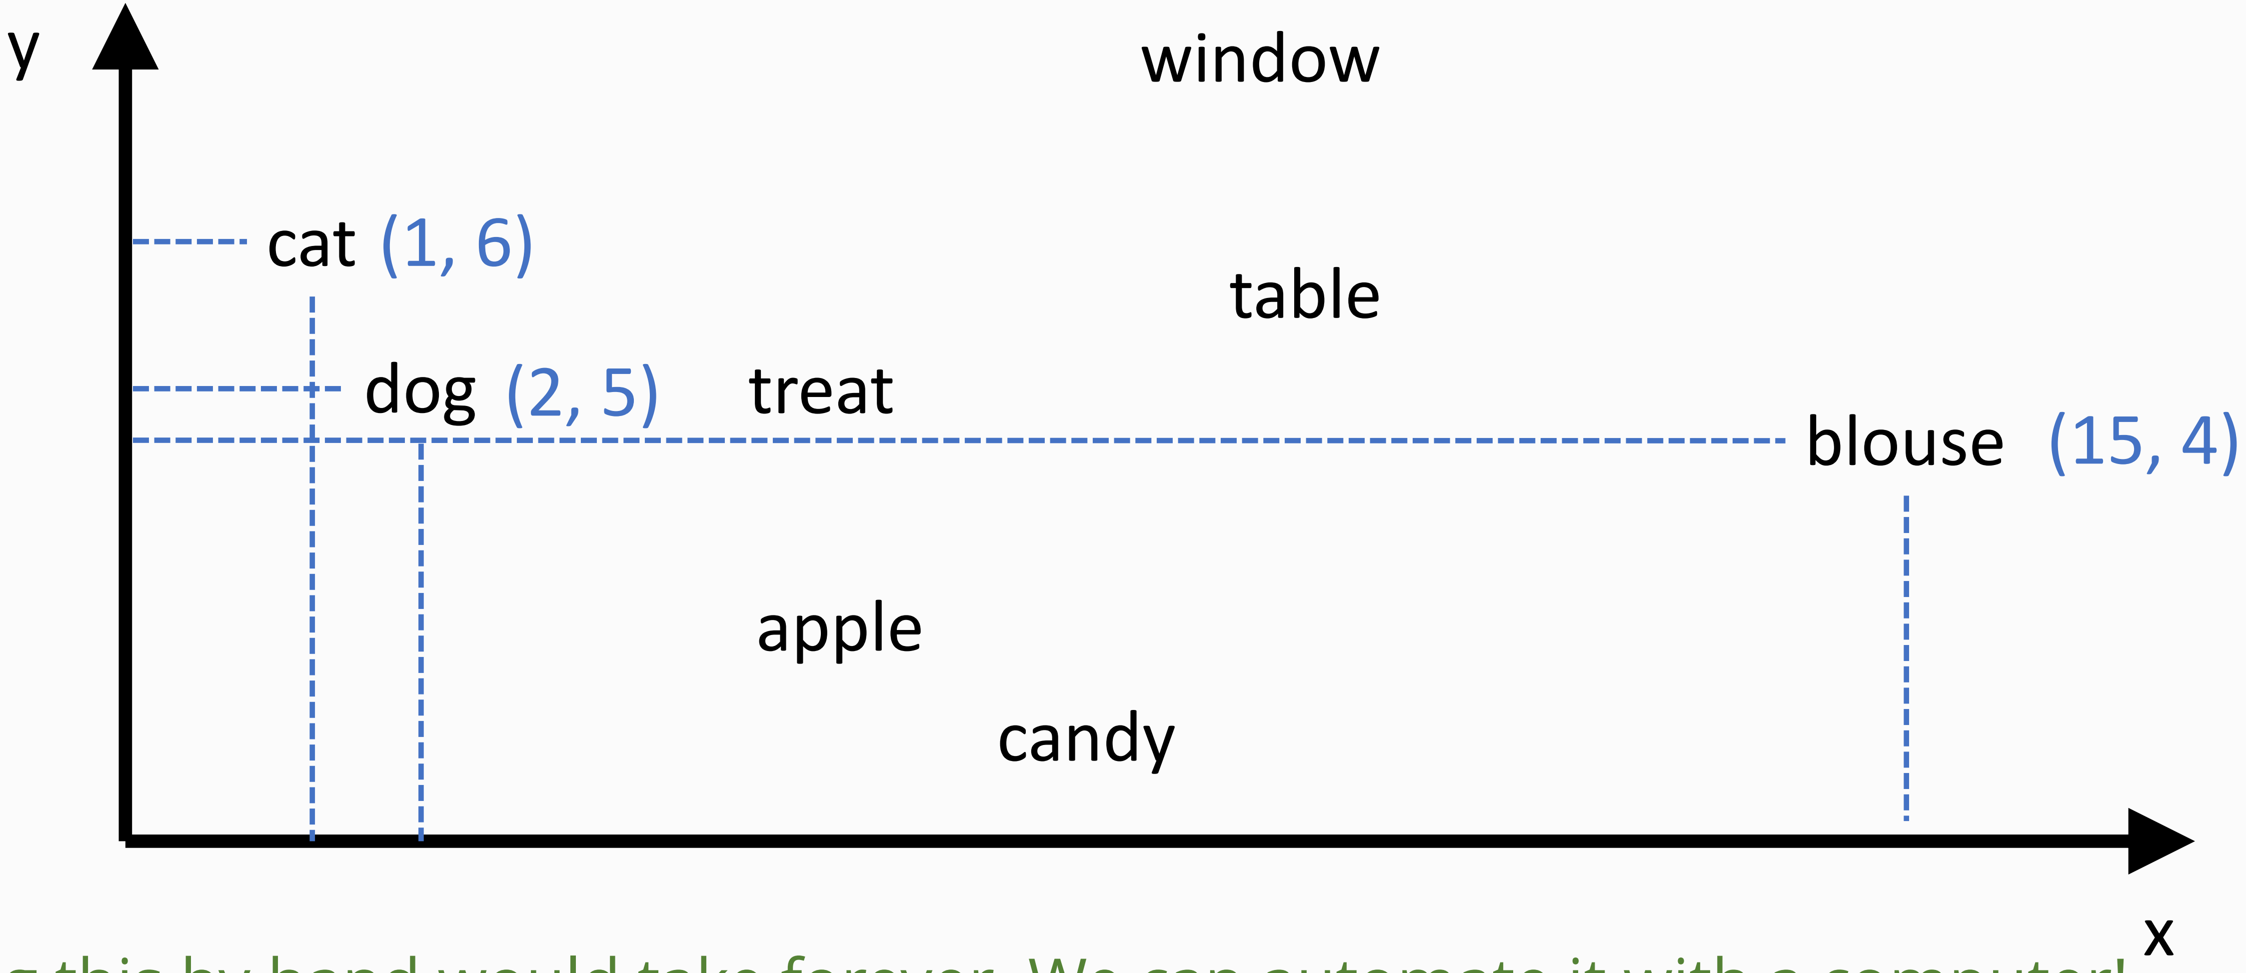

Doing this by hand would take forever. We can automate it with a computer!

Model 1: Latent Semantic Analysis (LSA)

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

- John Rupert Firth

Latent Semantic Analysis (LSA)

- Words that co-occur in bodies of text are often more similar to one another than words that do not.

Text 1: “The brain is an organ that serves as the center of the nervous system in all vertebrate and most invertebrate animals.”

Text 2: “In some systems the drain is for discharge of waste fluids, such as the drain in a sink in which the water is drained when it is no longer needed

Latent Semantic Analysis (LSA)

Example: the words “cat”, “dog”, “dopamine”, “reward”

- Take a bunch of documents (scientific articles, op-eds, etc.)
- Create a table where each <u>word</u> is a <u>row</u> and each <u>column</u> is a <u>document</u>
- For each document, if it contains the corresponding word, put a ✓ in that cell

Context (Document)					
CAT					
DOPAMINE					
DOG					
REWARD					

Landauer & Dumais (1997)

Latent Semantic Analysis (LSA)

Example: the words "cat", "dog", "dopamine", "reward"

- Take a bunch of documents (scientific articles, op-eds, etc.)
- Create a table where each word is a row and each column is a document
- For each document, if it contains the corresponding word, put a ✓ in that cell

Context (Document)	
	#1
CAT	✓
DOPAMINE	
DOG	✓
REWARD	

Landauer & Dumais (1997)

Latent Semantic Analysis (LSA)

Example: the words "cat", "dog", "dopamine", "reward"

- Take a bunch of documents (scientific articles, op-eds, etc.)
- Create a table where each word is a row and each column is a document
- For each document, if it contains the corresponding word, put a ✓ in that cell

Context (Document)			
	#1	#2	
CAT	✓		
DOPAMINE			
DOG	✓	✓	
REWARD		✓	

Landauer & Dumais (1997)

Latent Semantic Analysis (LSA)

Example: the words “cat”, “dog”, “dopamine”, “reward”

- Take a bunch of documents (scientific articles, op-eds, etc.)
- Create a table where each word is a row and each column is a document
- For each document, if it contains the corresponding word, put a ✓ in that cell

Context (Document)				
	#1	#2	#3	
CAT	✓			
DOPAMINE			✓	
DOG	✓	✓		
REWARD		✓	✓	

Landauer & Dumais (1997)

Latent Semantic Analysis (LSA)

Example: the words “cat”, “dog”, “dopamine”, “reward”

- Take a bunch of documents (scientific articles, op-eds, etc.)
- Create a table where each word is a row and each column is a document
- For each document, if it contains the corresponding word, put a ✓ in that cell

Context (Document)					
	#1	#2	#3	#4	
CAT	✓				
DOPAMINE			✓	✓	
DOG	✓	✓			
REWARD		✓	✓	✓	

Landauer & Dumais (1997)

Latent Semantic Analysis (LSA)

- CAT and DOG are more similar to each other than they are to DOPAMINE or REWARD.

Context (Document)					
	#1	#2	#3	#4	#5
CAT	✓				✓
DOPAMINE			✓	✓	
DOG	✓	✓			✓
REWARD		✓	✓	✓	

Latent Semantic Analysis (LSA)

- DOPAMINE and REWARD are more similar to each other than they are to CAT or DOG.

Context (Document)					
	#1	#2	#3	#4	#5
CAT	✓				✓
DOPAMINE			✓	✓	
DOG	✓	✓			✓
REWARD		✓	✓	✓	

Latent Semantic Analysis (LSA)

DOG and REWARD are similar to each other.

Context (Document)					
	#1	#2	#3	#4	#5
CAT	✓				✓
DOPAMINE			✓	✓	
DOG	✓	✓			✓
REWARD		✓	✓	✓	

Latent Semantic Analysis (LSA)

- Reduce to 3 columns for the sake of the example

Context (Document)				
	#1	#2	#3	
CAT	✓			
DOPAMINE			✓	
DOG	✓	✓		
REWARD		✓	✓	

Latent Semantic Analysis (LSA)

- Change each check to a 1 and nothing to a 0

Context (Document)				
	#1	#2	#3	coordinate
CAT	1	0	0	1, 0, 0
DOPAMINE	0	0	1	0, 0, 1
DOG	1	1	0	1, 1, 0
REWARD	0	1	1	0, 1, 1

Latent Semantic Analysis (LSA)

- …to create a coordinate.
- How many dimensions are there?

Context (Document)				
	#1	#2	#3	coordinate
CAT	1	0	0	1, 0, 0
DOPAMINE	0	0	1	0, 0, 1
DOG	1	1	0	1, 1, 0
REWARD	0	1	1	0, 1, 1

3 dimensions now

Latent Semantic Analysis (LSA)

- …to create a coordinate.
- How many dimensions are there?

Context (Document)				
	#1	#2	#3	coordinate
CAT	1	0	0	1, 0, 0
DOPAMINE	0	0	1	0, 0, 1
DOG	1	1	0	1, 1, 0
REWARD	0	1	1	0, 1, 1

The coordinates (or strings of numbers) are called **vectors**

Latent Semantic Analysis (LSA)

- Now, if we treat these words as coordinates in semantic space…

Context (Document)				
	#1	#2	#3	coordinate
CAT	1	0	0	1, 0, 0
DOPAMINE	0	0	1	0, 0, 1
DOG	1	1	0	1, 1, 0
REWARD	0	1	1	0, 1, 1

Latent Semantic Analysis (LSA)

- Now, if we treat these words as coordinates in semantic space…

Context (Document)				
	#1	#2	#3	coordinate
CAT	1	0	0	1, 0, 0
DOPAMINE	0	0	1	0, 0, 1
DOG	1	1	0	1, 1, 0
REWARD	0	1	1	0, 1, 1

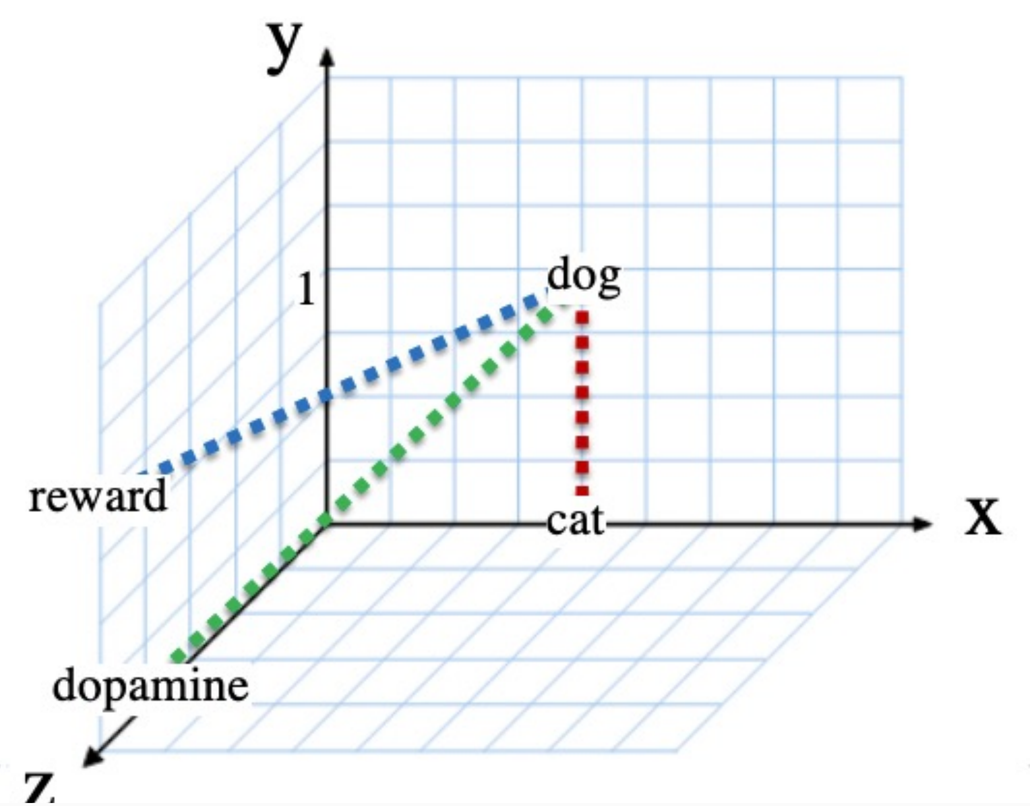

Latent Semantic Analysis (LSA)

- Now, if we treat these words as coordinates in semantic space…

Context (Document)				
	#1	#2	#3	coordinate
CAT	1	0	0	1, 0, 0
DOPAMINE	0	0	1	0, 0, 1
DOG	1	1	0	1, 1, 0
REWARD	0	1	1	0, 1, 1

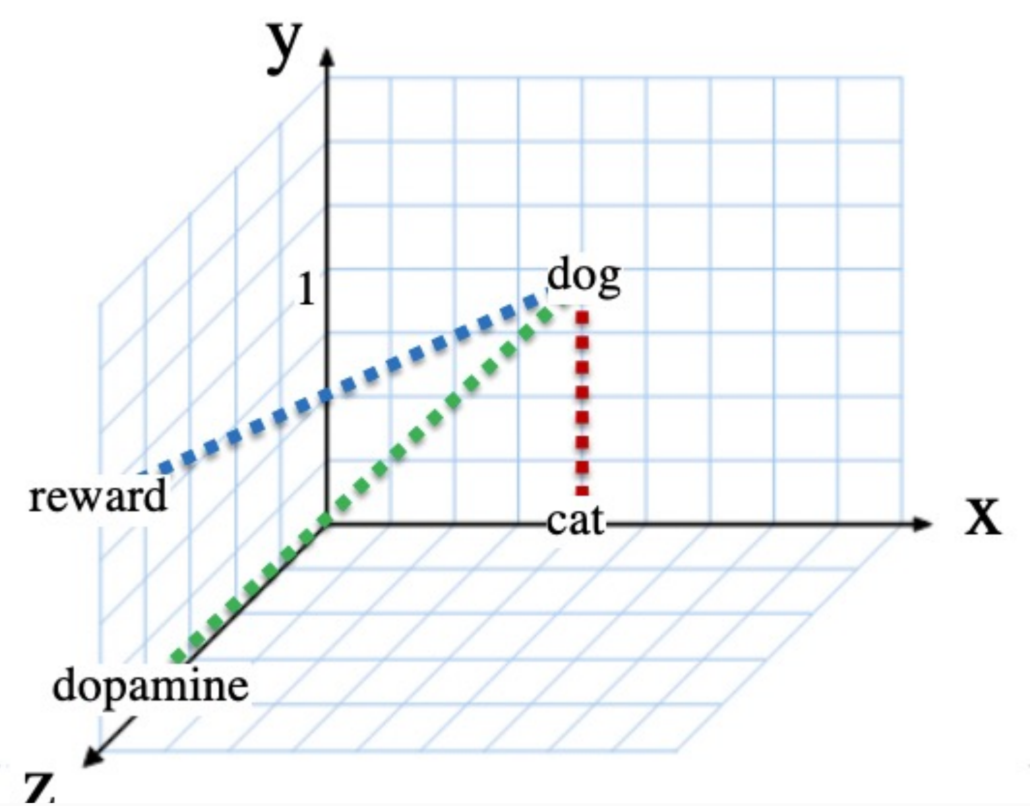

A way to capture relationships between words.

Latent Semantic Analysis (LSA)

- How does LSA think of meaning?

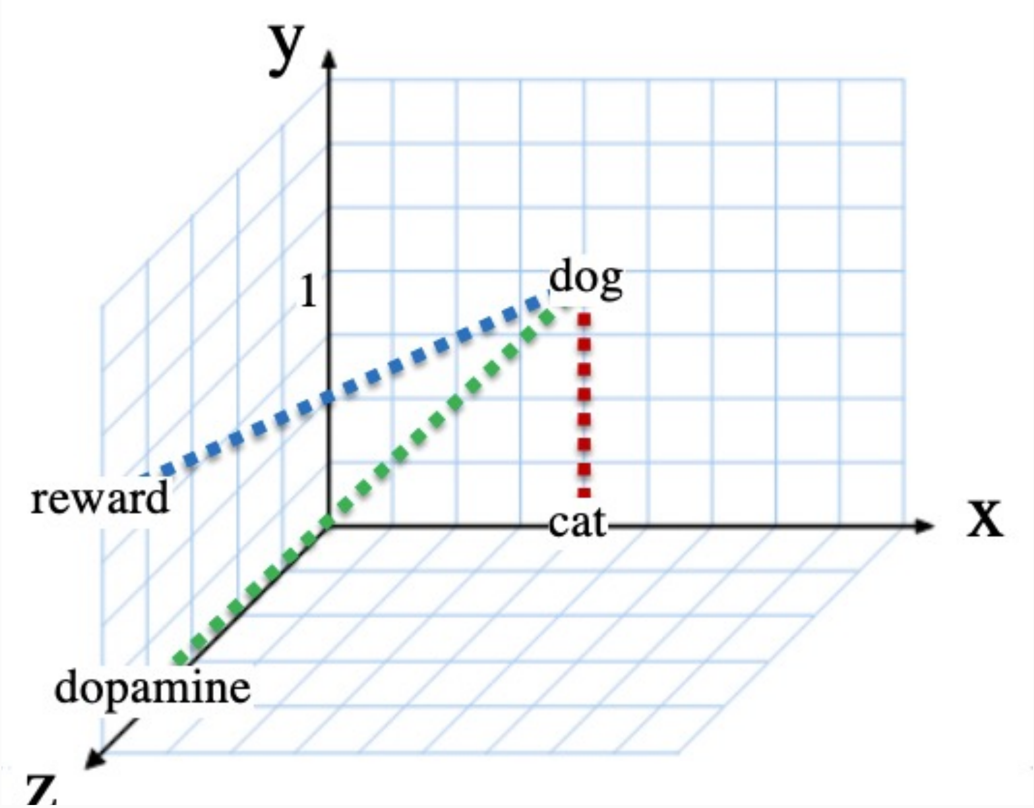

A way to capture relationships between words.

Latent Semantic Analysis (LSA)

- How does LSA think of meaning?

- 'Distance' in semantic space = distance between points on a graph

- Distance correlates with priming and other measures of semantic overlap

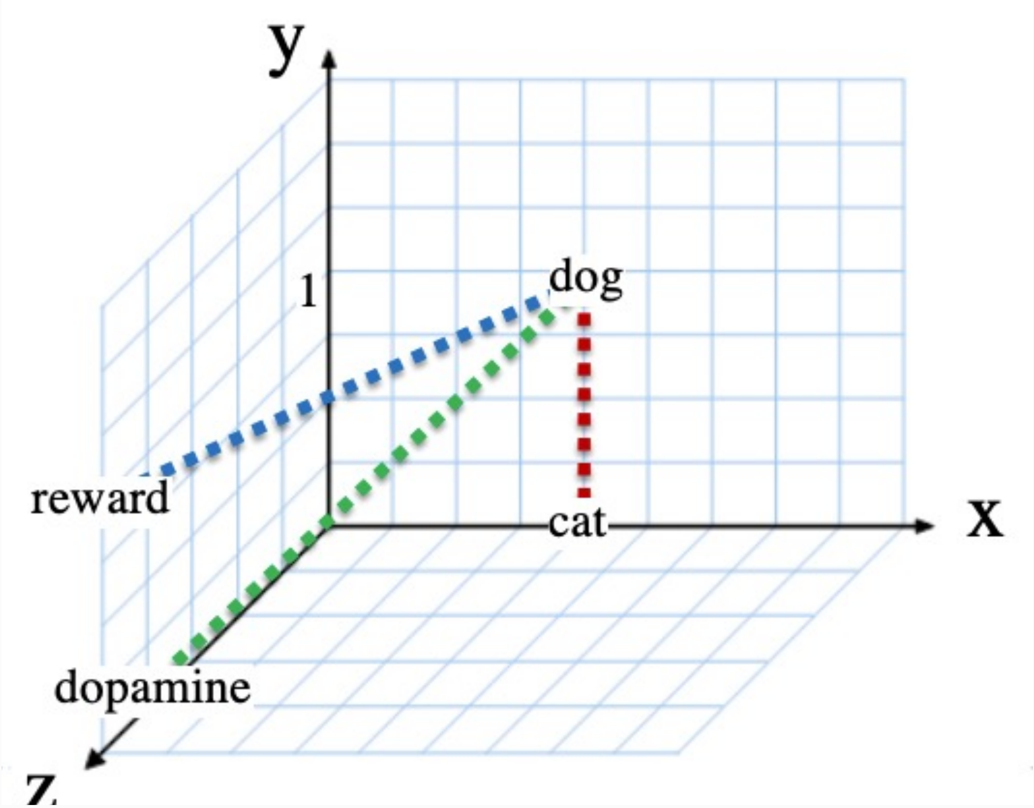

A way to capture relationships between words.

Model 2: Word2Vec

"Never ask for the meaning of a word in isolation, but only in the context of a sentence."

- Gottlob Frege

Word2Vec

- A slightly different approach to LSA, but same assumption
  - Words with similar meanings can be used in similar local contexts
  - What words could go into these contexts?

1. I almost got run over by a ________ yesterday.
2. Mary just had _________ for dinner.

Word2Vec

- With the model Word2Vec, you train it to predict…
  - a word given its context

    Saw a stray ____ wandering around the…

  - a neighboring word given a word

    ___ _____ ___ cat ____ _____ _____

Activity: make a sentence as a class.

Word2Vec

- Preparation before the model
  - Take all the words you want (could be all the words in the English language)
  - Line them up in a list
  - Give them each a different (random) string of 0s and 1s
    - Example:
      - Cat 1 0 0 0 0 …
      - Dog 0 1 0 0 0 …
      - Dopamine 0 0 1 0 0 …
      - Treat 0 0 0 1 0 …
      - Apple 0 0 0 0 1 …
      - …

Background on models

No need to know the details! Just know that although the model isn't *explicitly* trained to do this, it can develop representations of words.

During the training phase, we train the model

Note: we will discuss more about how the model works in the next lecture

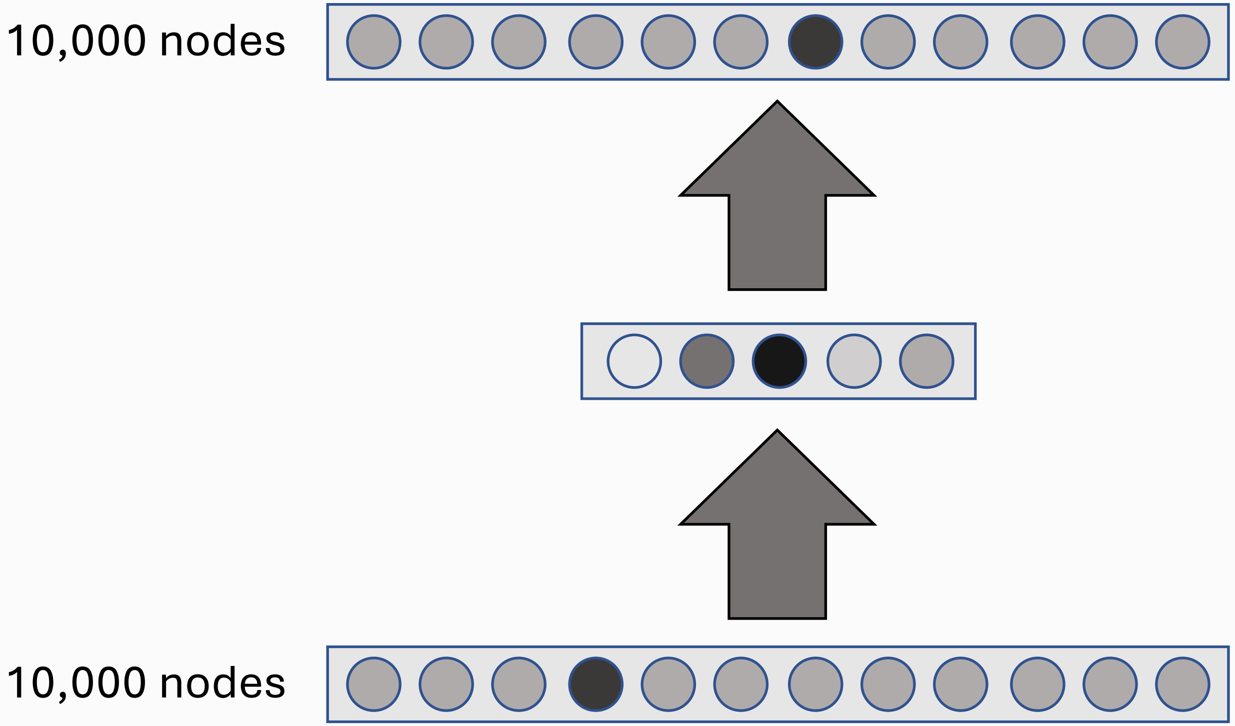

Background on models

No need to know the details! Just know that although the model isn't *explicitly* trained to do this, it can develop representations of words.

Training phase

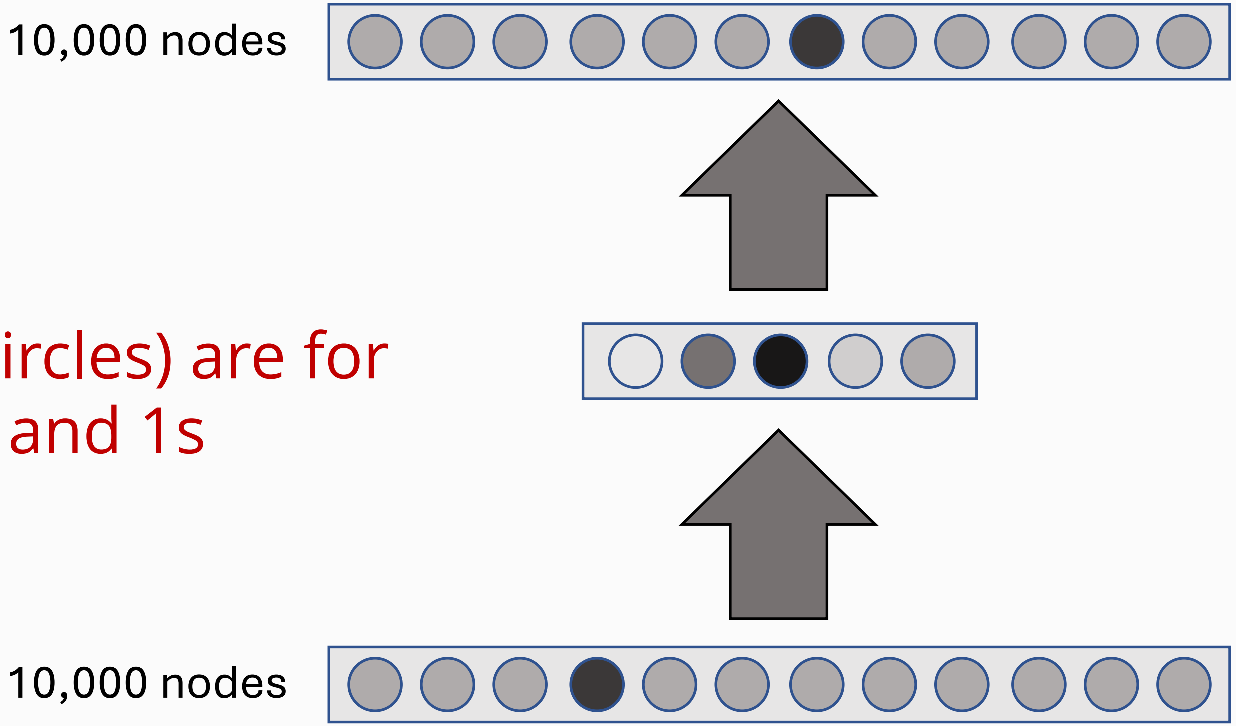

Nodes (circles) are for
0s and 1s

Background on models

No need to know the details! Just know that although the model isn't *explicitly* trained to do this, it can develop representations of words.

Training phase

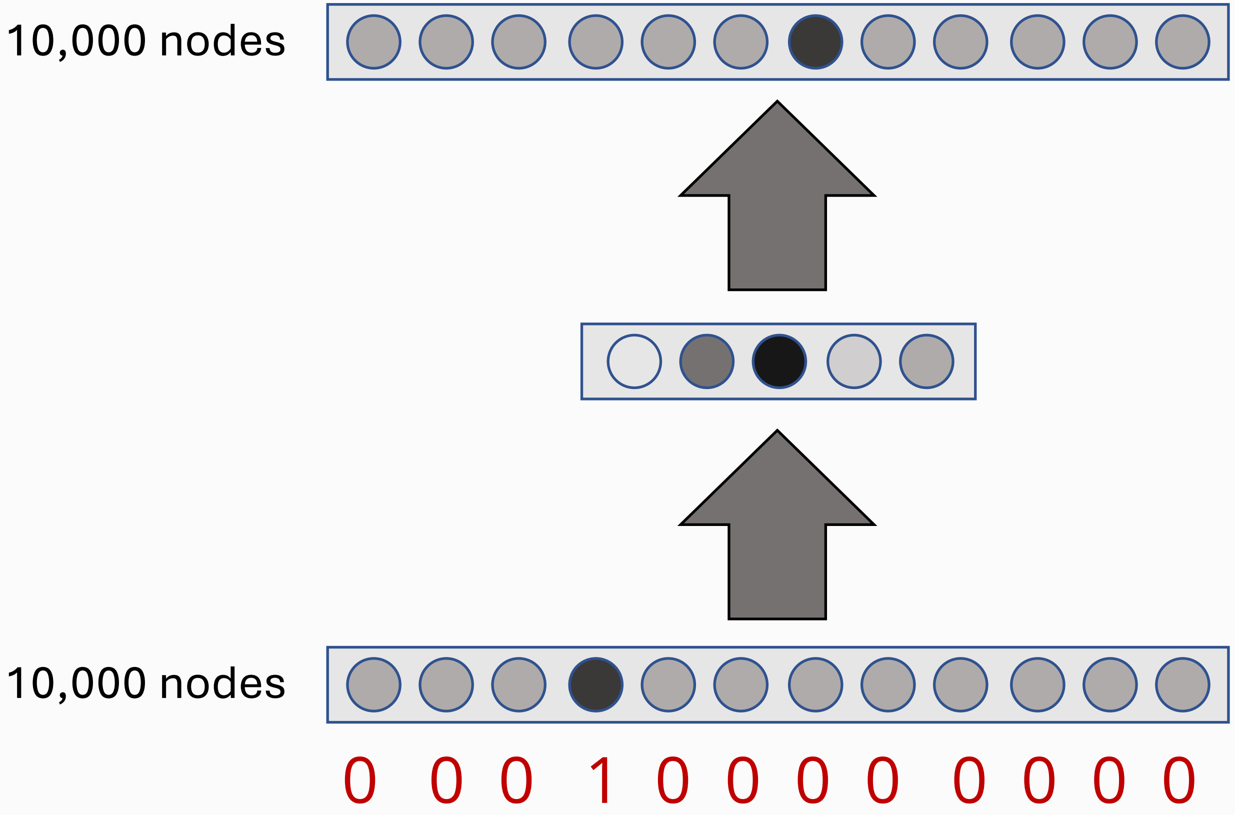

Input the string for "treat"

Background on models

Training phase

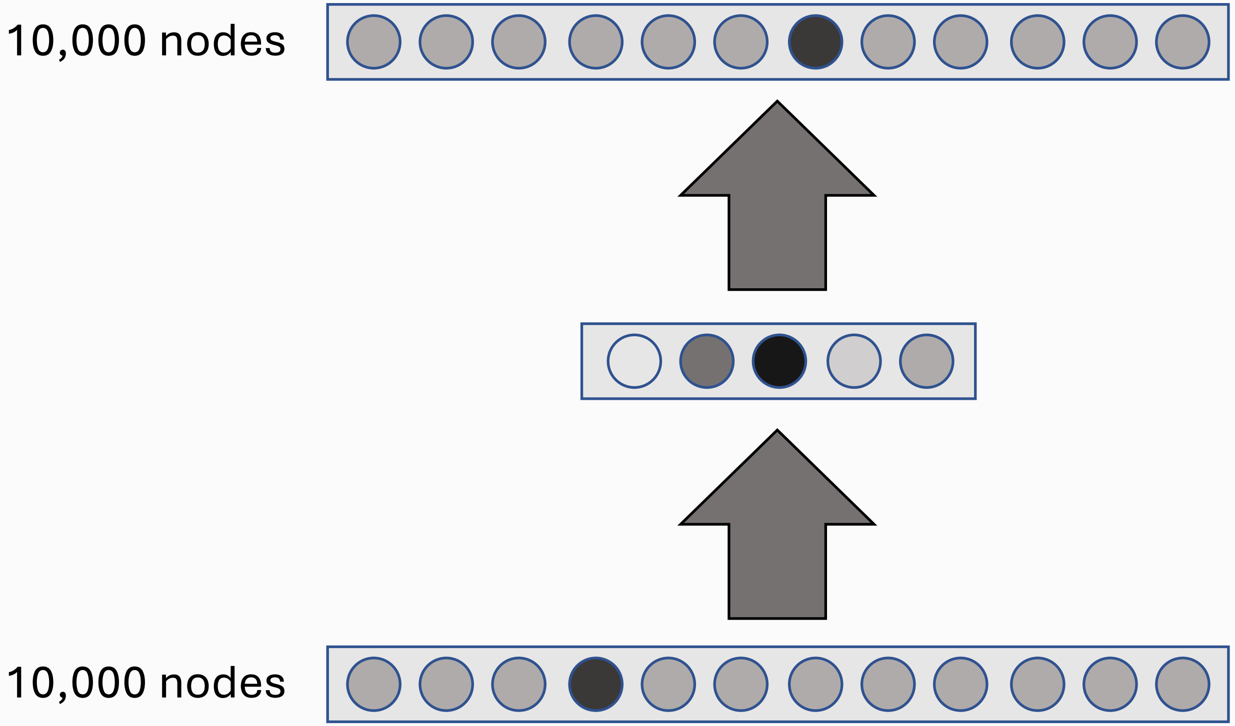

Hidden layer of nodes
It’s a black box
(meaning we don’t know
what it does!)

Input word goes here (e.g. TREAT)

Background on models

Training phase

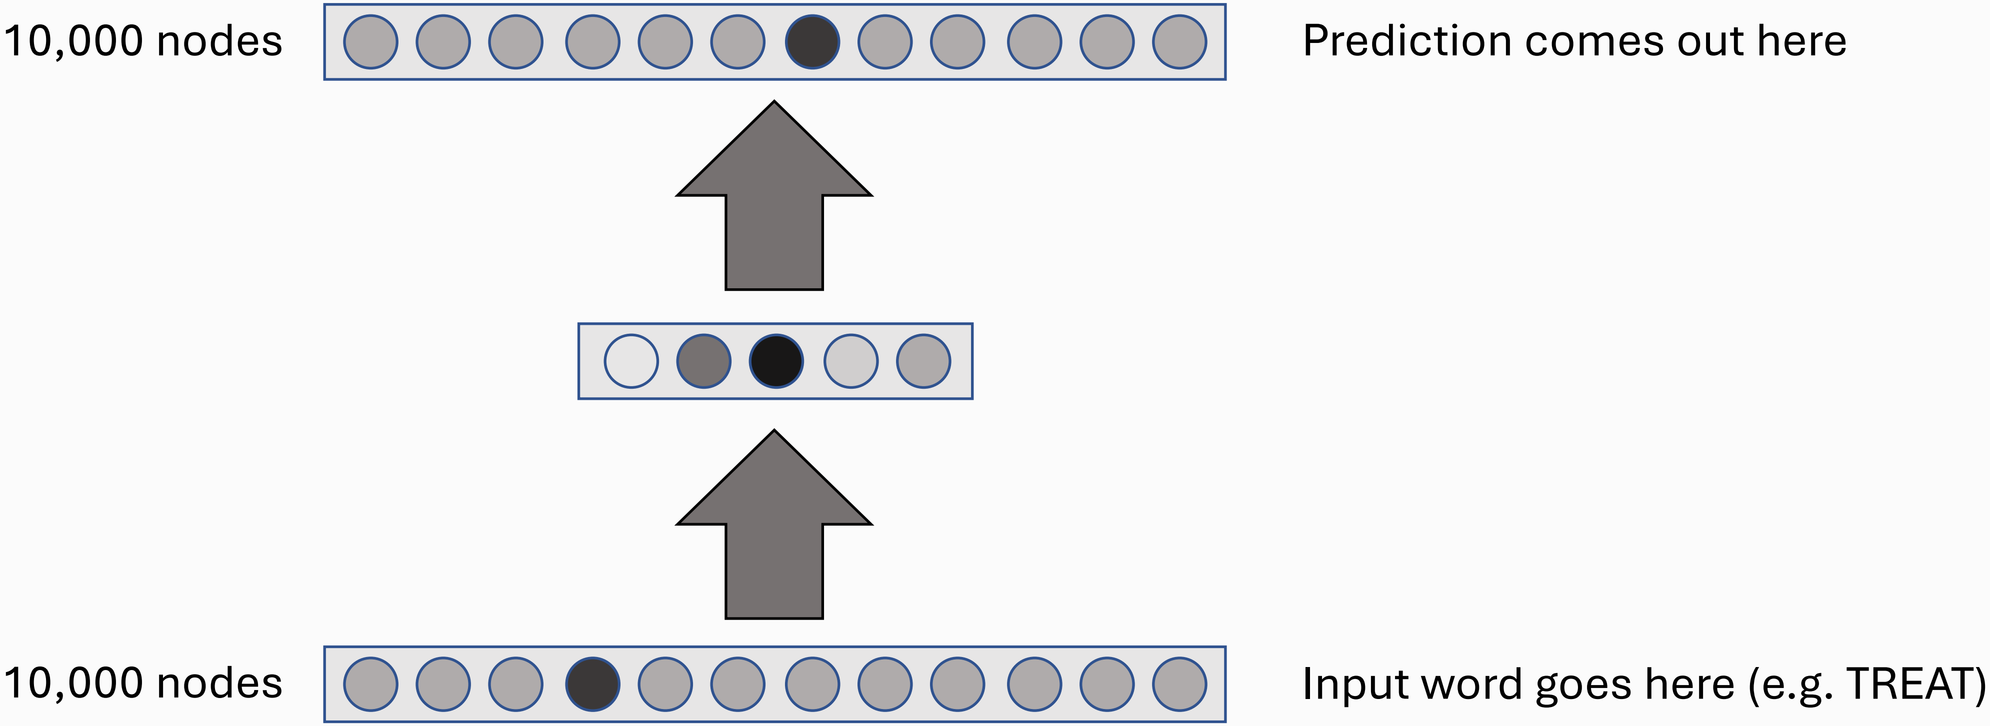

Background on models

Training phase

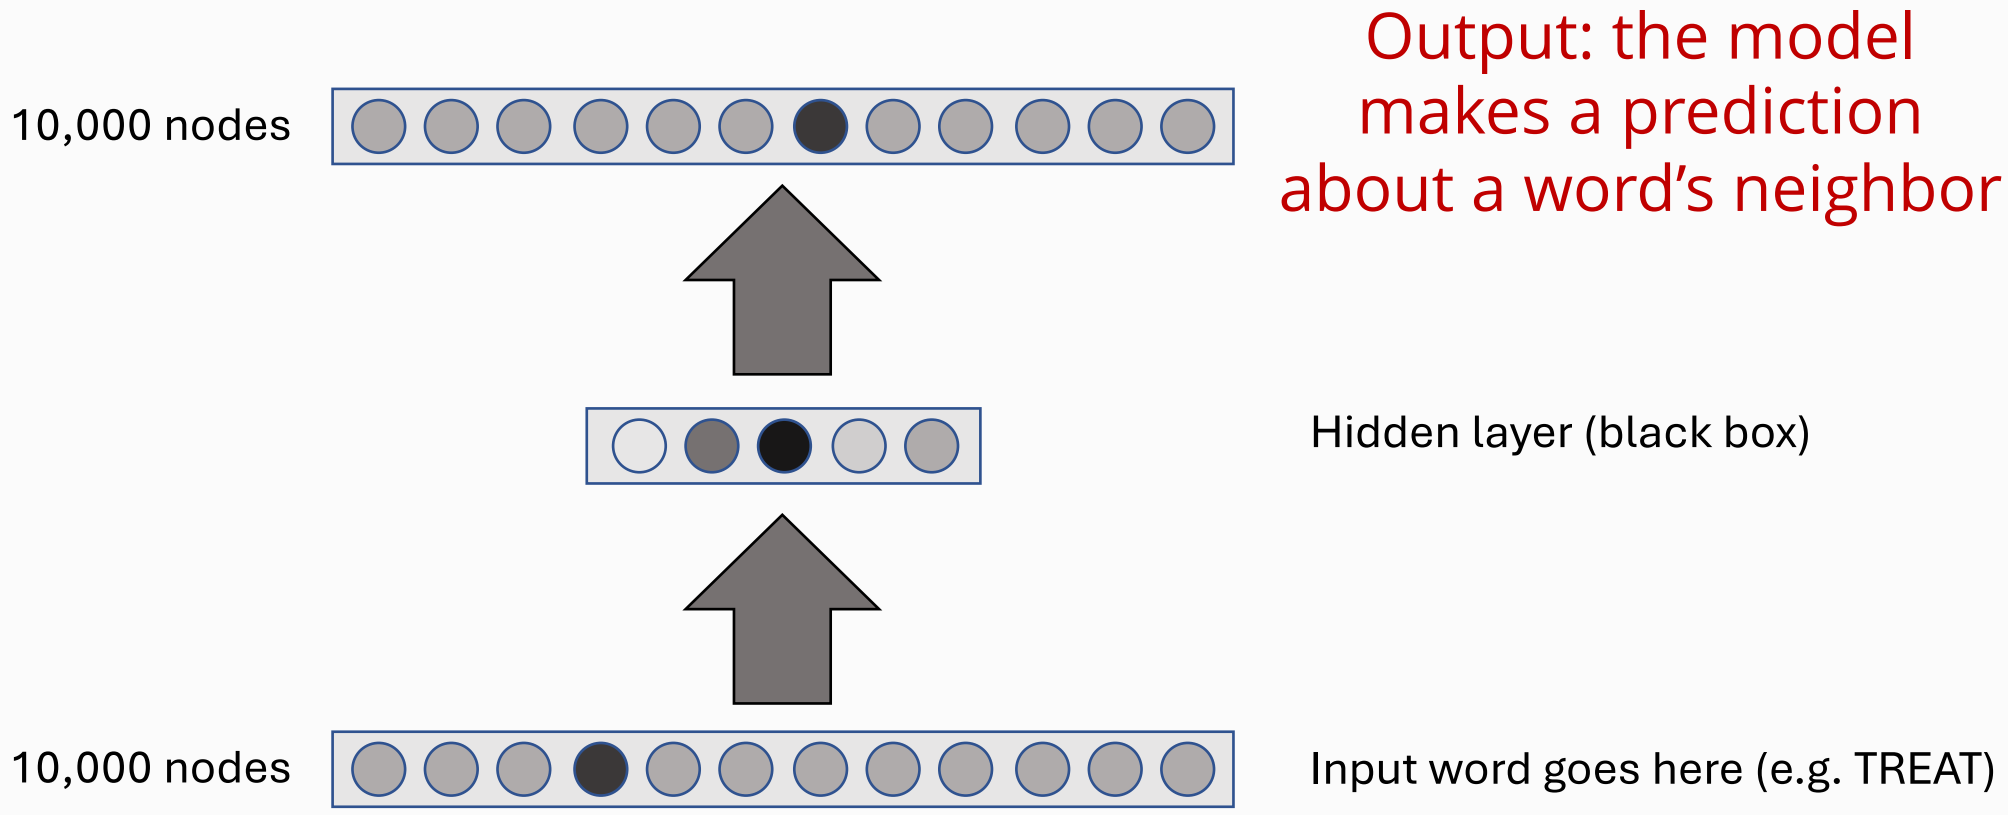

Background on models

Training phase

10,000 nodes

Prediction comes out here

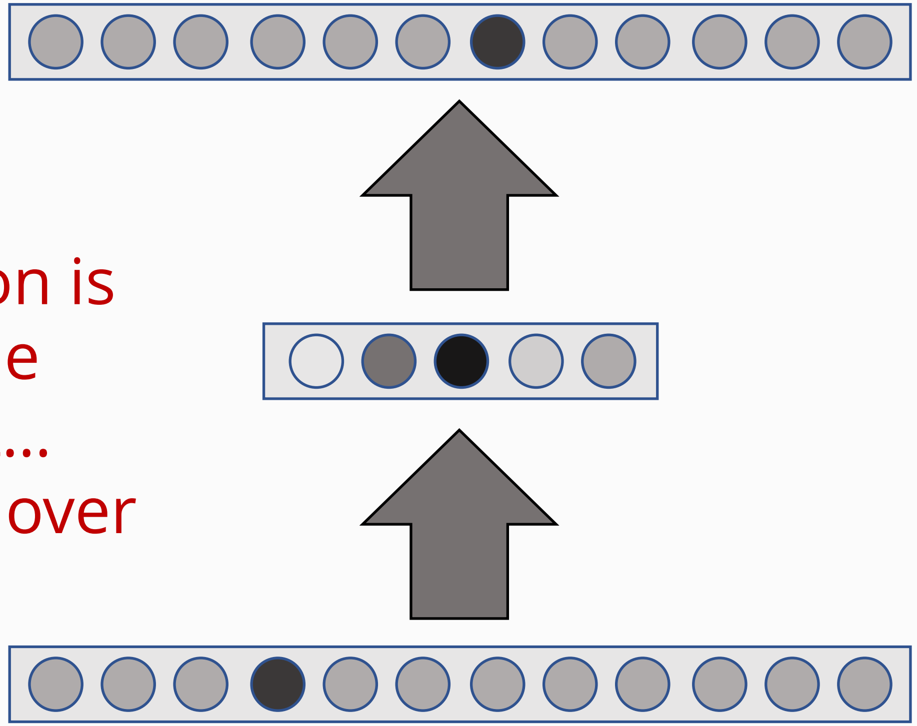

Then the prediction is compared to the "correct" output… over and over and over

Hidden layer (black box)

10,000 nodes

Input word goes here (e.g. TREAT)

Background on models

Test phase: after training, you test by seeing what it predicts without telling it what the output is supposed to be.

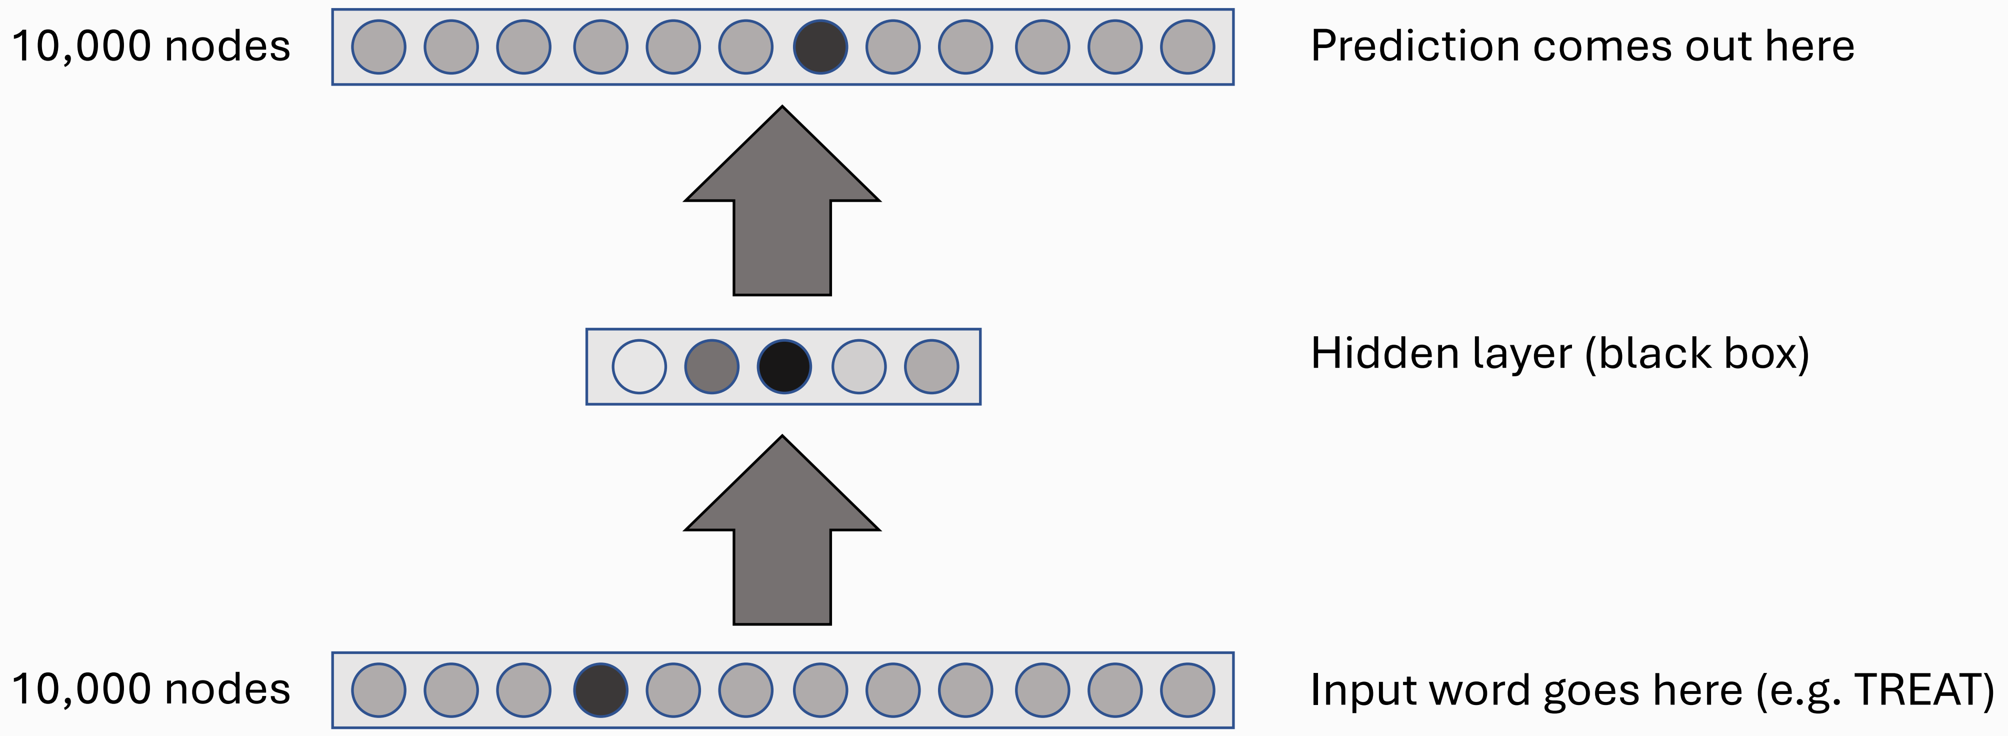

Background on models

Test phase

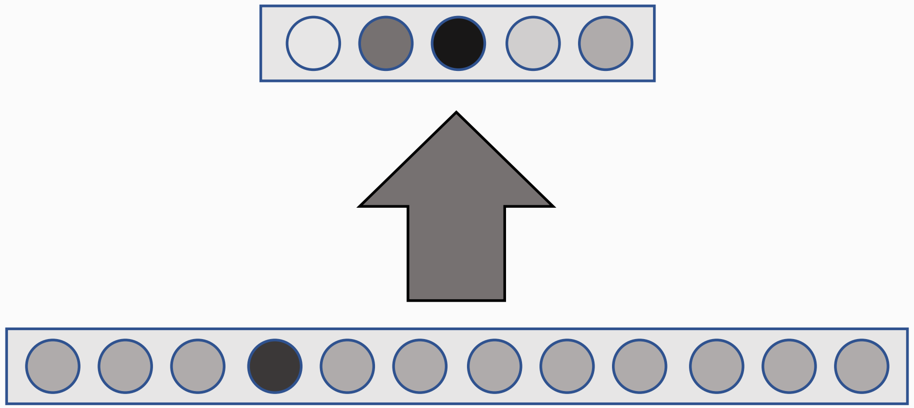

Background on models

Test phase

300-dimension representation of a word is created (a word vector, just like the small ones we made in the LSA example)

Hidden layer (black box)

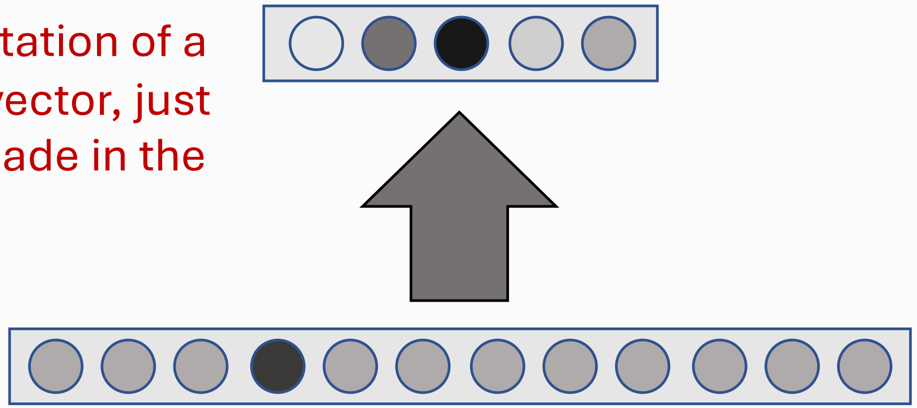

10,000 nodes

Input word goes here (e.g. CAT)

CAT =

[0.26 0.84 0.29 0.54 0.06 0.78 0.32 ...]

Background on models

Test phase

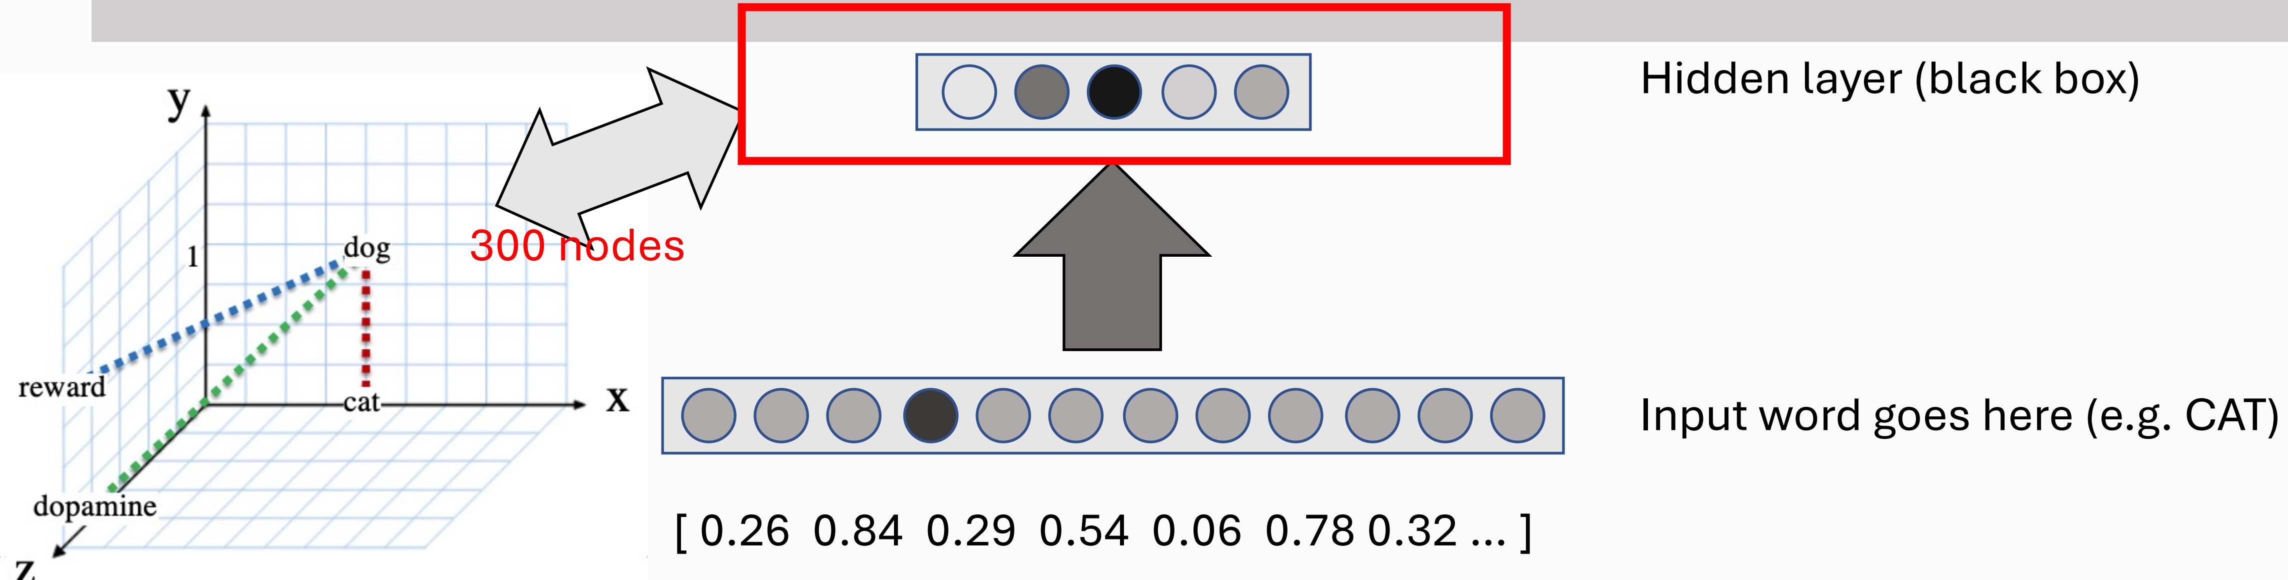

Background on models

We will talk more about this next lecture

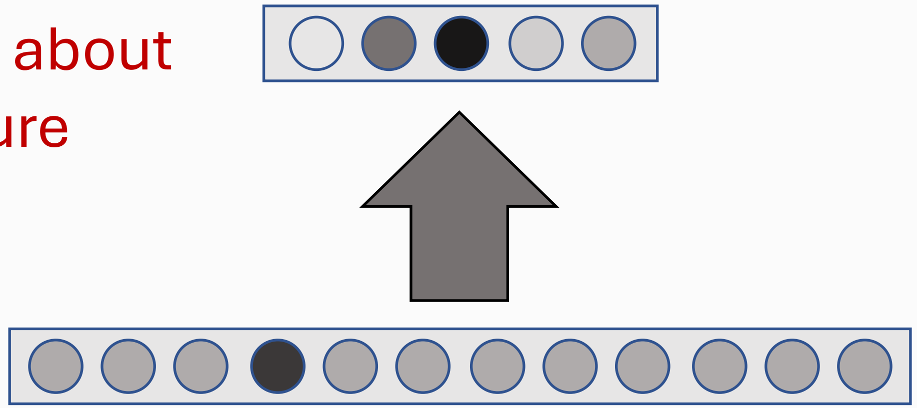

Hidden layer (black box)

10,000 nodes

Input word goes here (e.g. CAT)

CAT =

[0.26 0.84 0.29 0.54 0.06 0.78 0.32 ...]

Are these vectors useful?

- How is Word2Vec defining meaning?

Are these vectors useful?

- What does this kind of modeling have to do with the brain?

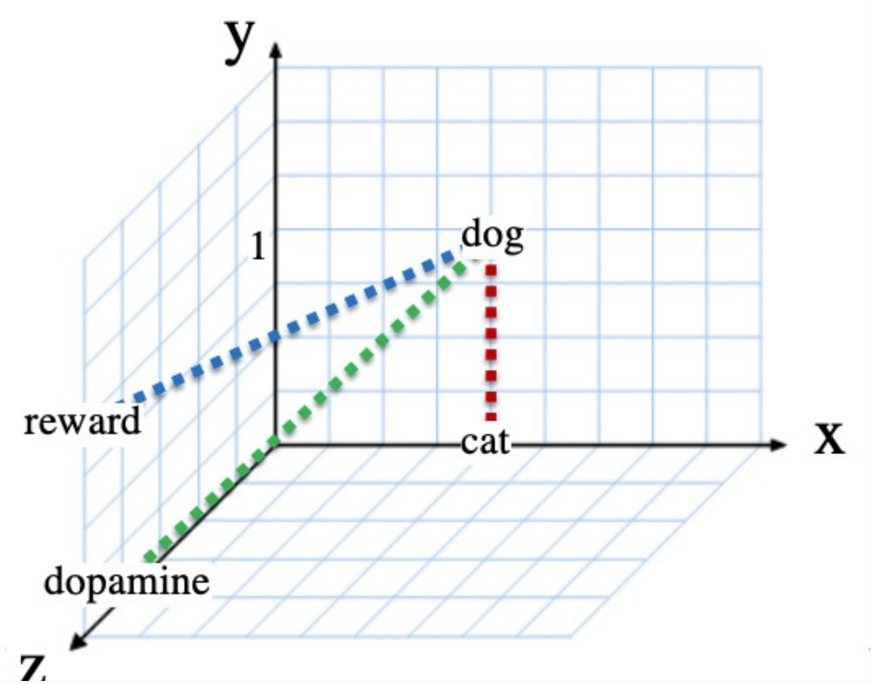

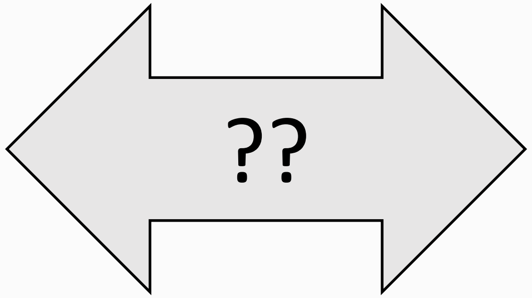

Are these vectors useful?

- Its behavior might match some human behavior when we test meaning.

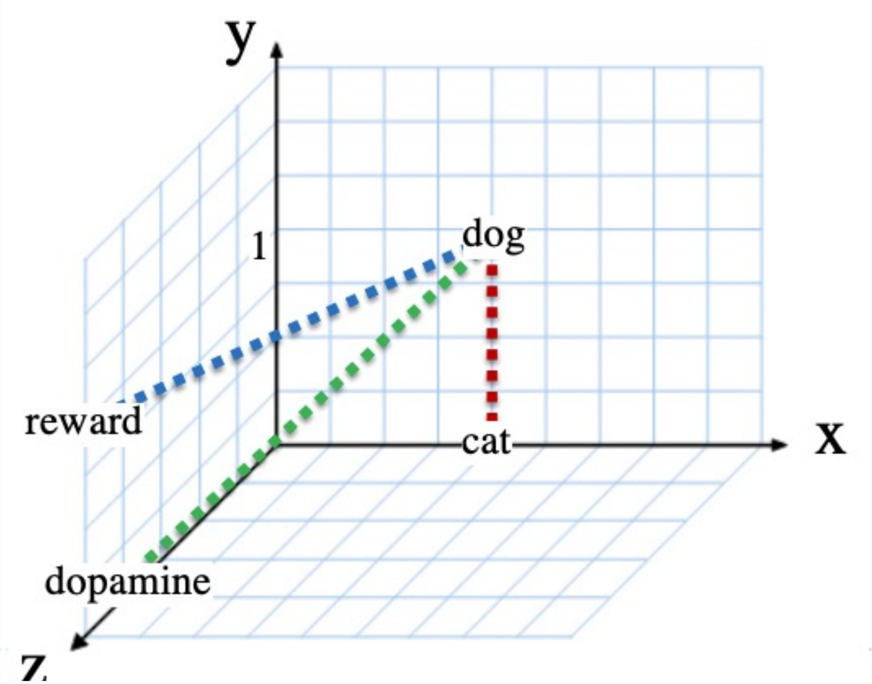

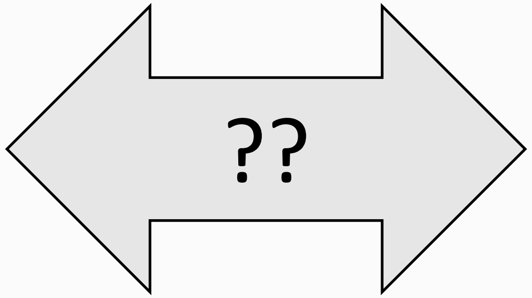

Model behavior versus human behavior

Example: vector manipulation (adding, subtracting, etc.)

Think of each arrow as a vector.

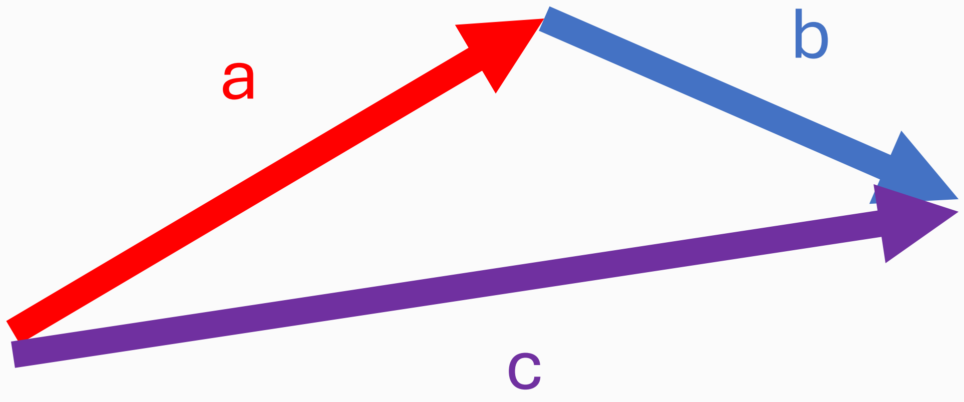

a + b = c

c – b = a

Model behavior versus human behavior

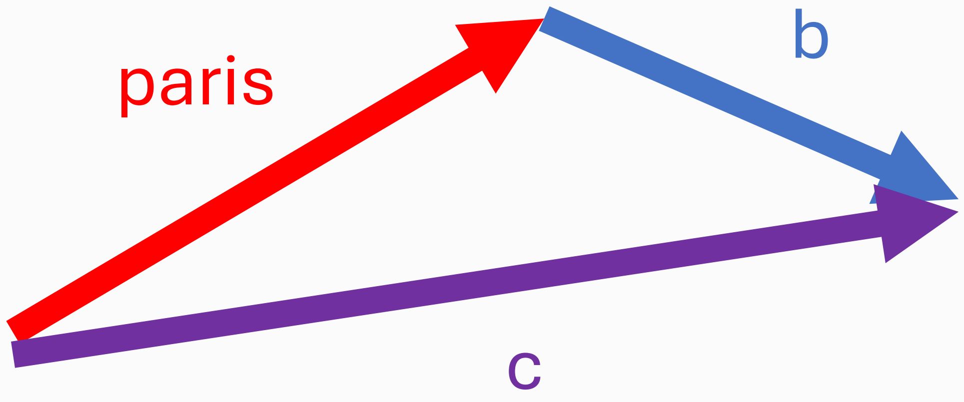

Mikolov et al. (2013)

Model behavior versus human behavior

What would a Human say?

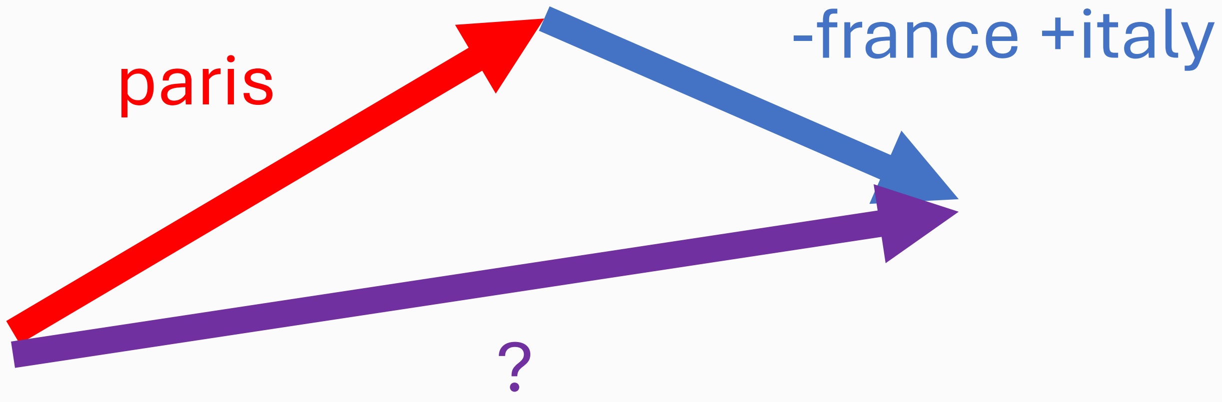

paris - france + italy = ?

Mikolov et al. (2013)

Model behavior versus human behavior

What would a Human say?

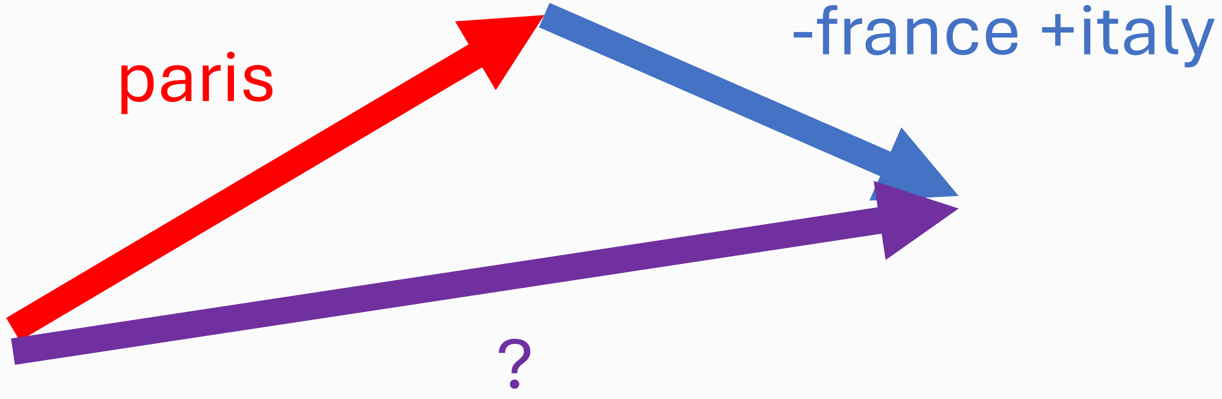

paris - france + italy = rome

Mikolov et al. (2013)

Model behavior versus human behavior

What does a model say?

It picks out a point in the semantic space that is close to "Rome"

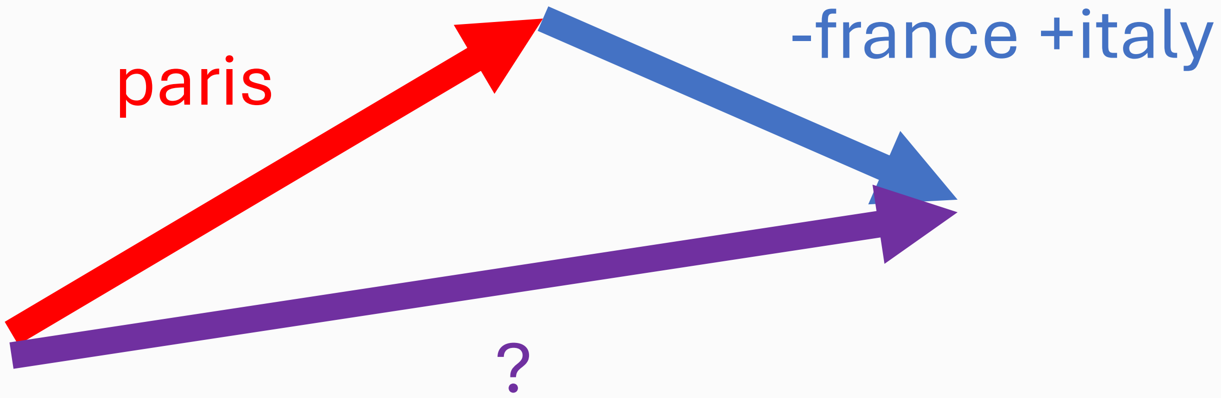

paris - france + italy ≈ rome

Mikolov et al. (2013)

Model behavior versus human behavior

What does a model say?

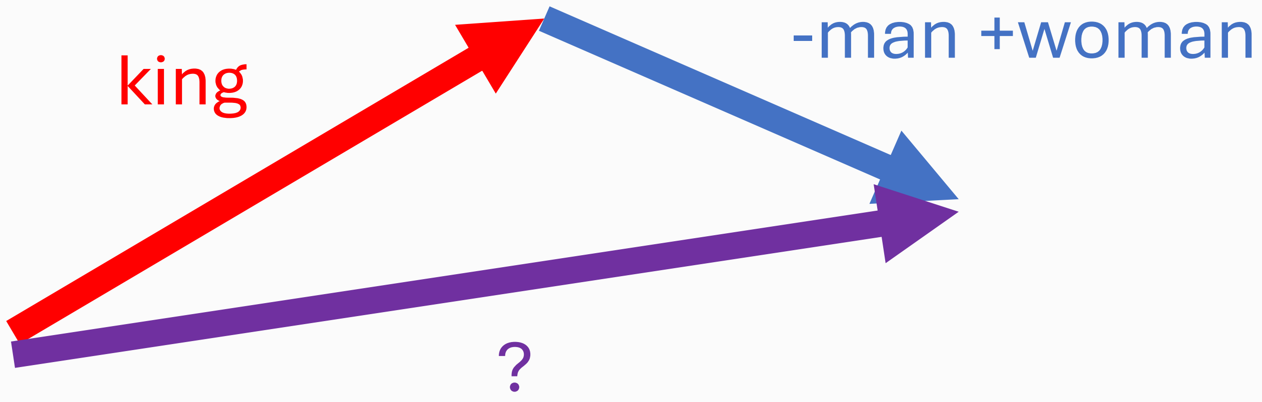

king - man + woman = ?

Mikolov et al. (2013)

Model behavior versus human behavior

What does a model say?

It picks out a point in the semantic space that is close to ”queen”

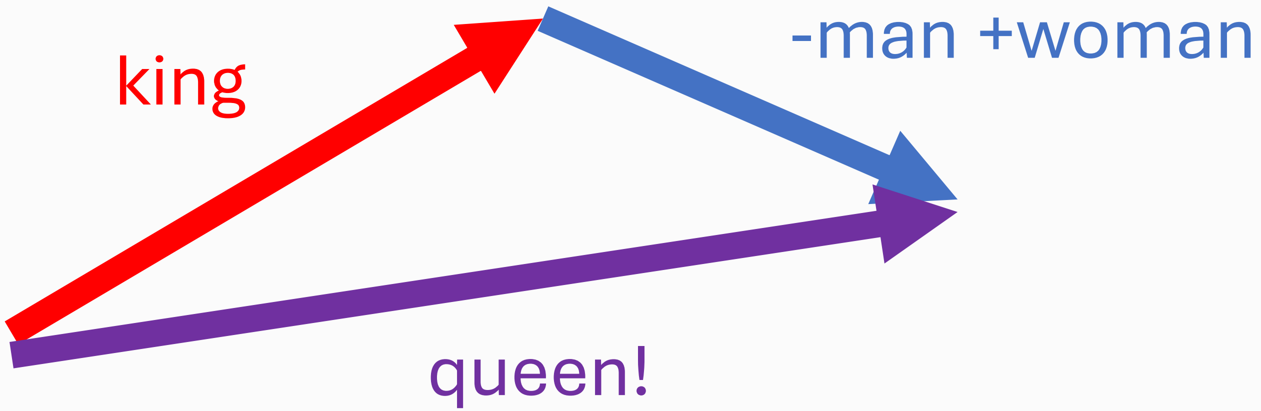

king - man + woman ≈ queen

Mikolov et al. (2013)

Model behavior versus human behavior

- They even appear to categorize words based on dimensions they do not have access to, like size

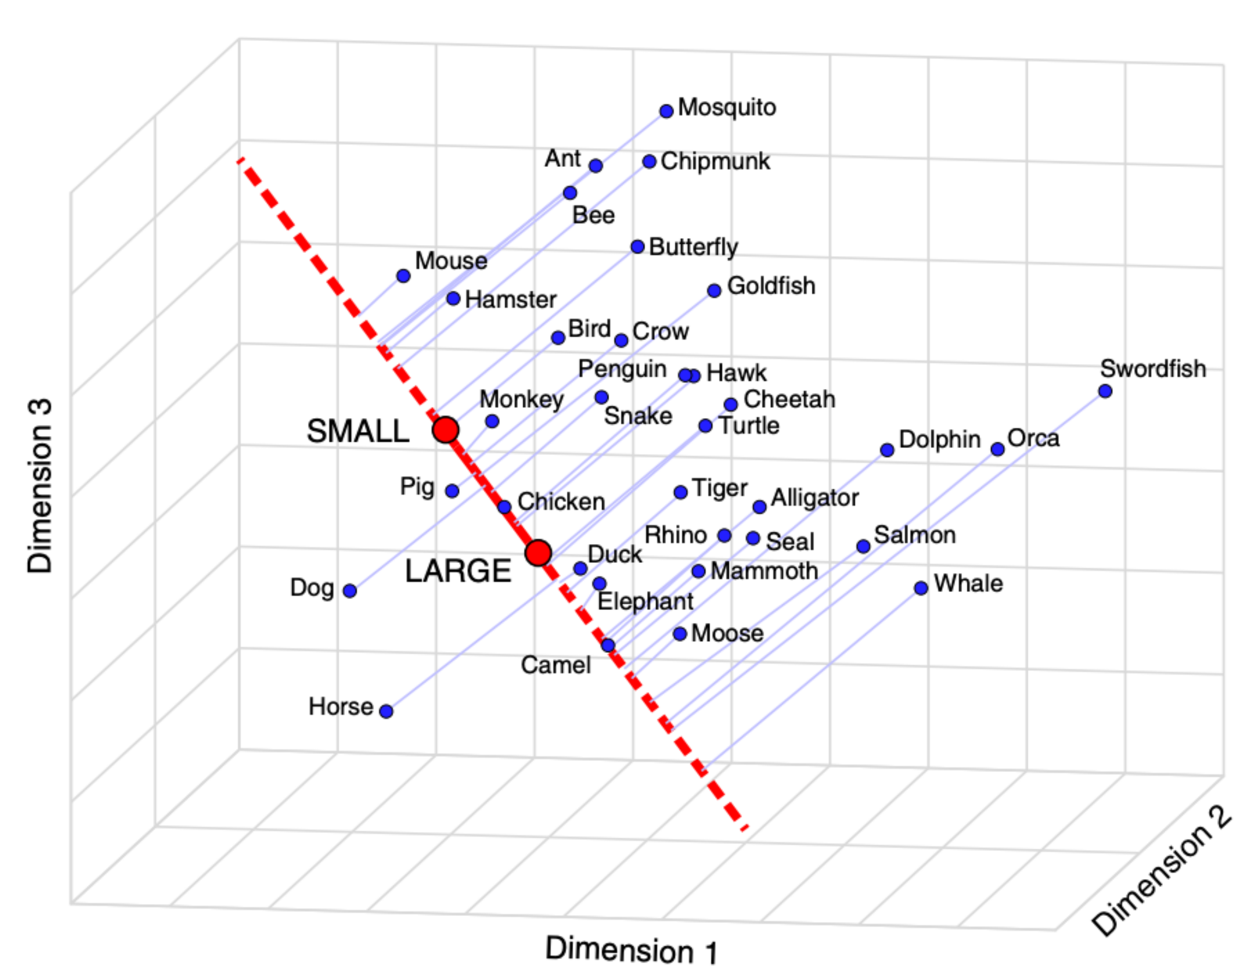

Grand et al. (2022)

Model behavior versus human behavior

- They even appear to categorize words based on dimensions they do not have access to, like size

- …but at a fine-grained level it's not perfect.

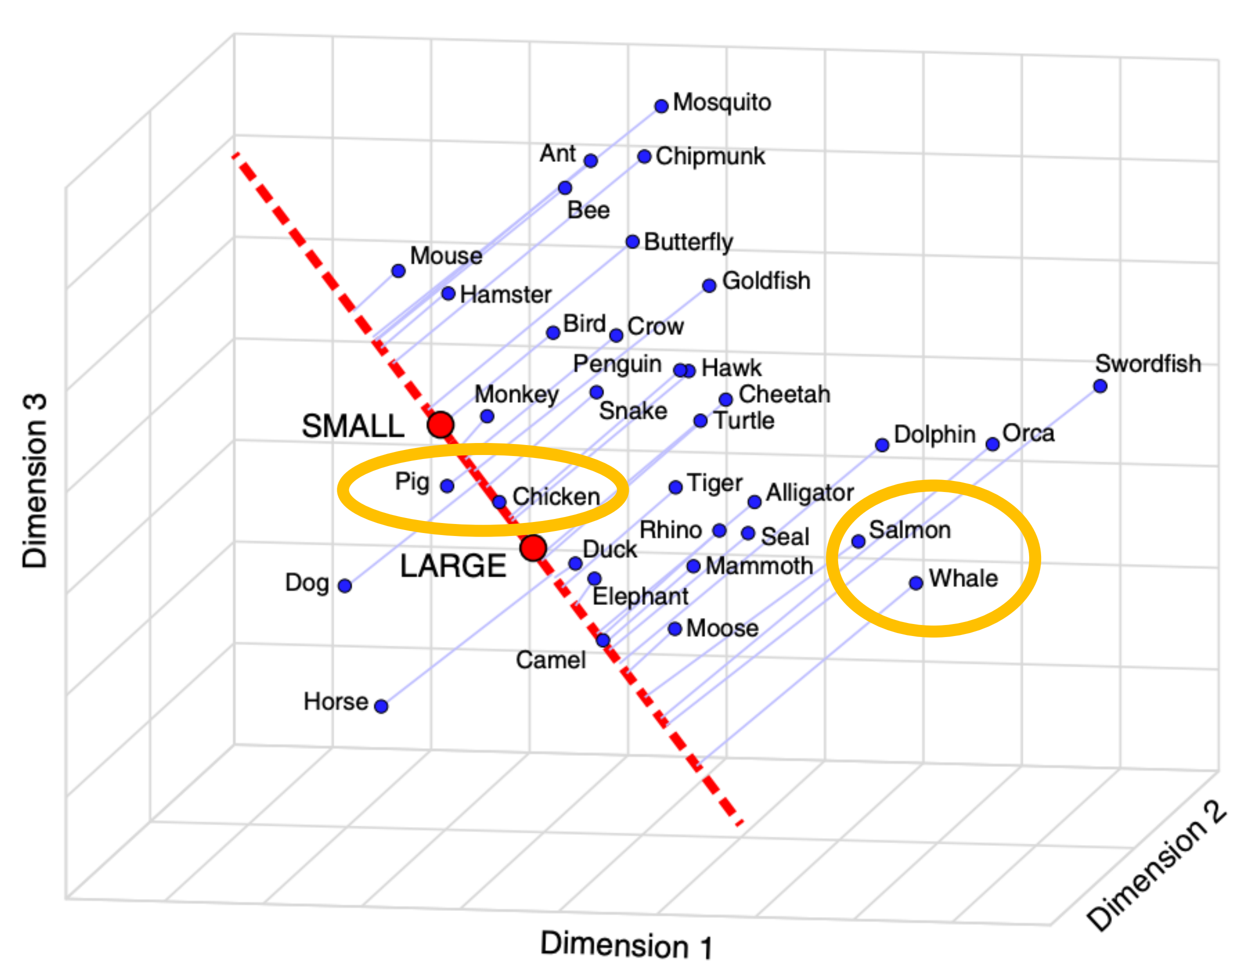

Grand et al. (2022)

Are word vectors good models?

- Are word vectors good models of how we learn language?
  - No.
    - babies get a lot more limited linguistic input
    - they aren’t trained over and over

- But word vectors might be good models of how we organize semantic features.

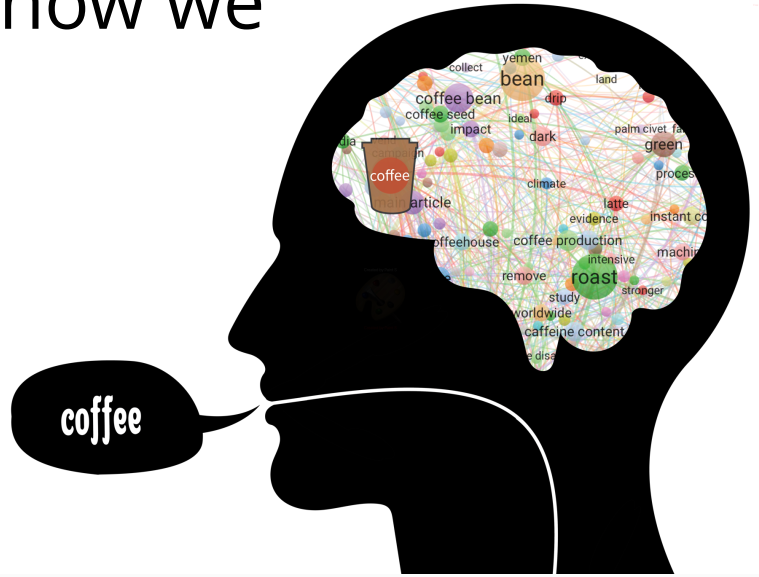

Key concepts

- ✓ Models are wrong but useful
- ✓ Meaning can be represented as a semantic space
  - ✓ Semantic spaces can be created using word vectors
- ✓ Context and co-occurrence matters for these models
- ✓ LSA uses global context
- ✓ Word2Vec uses local context
- ✓ Similarities and differences between computational models for word meanings and what humans do
  - ✓ Acquisition
  - ✓ Organization
  - ✓ Use for prediction