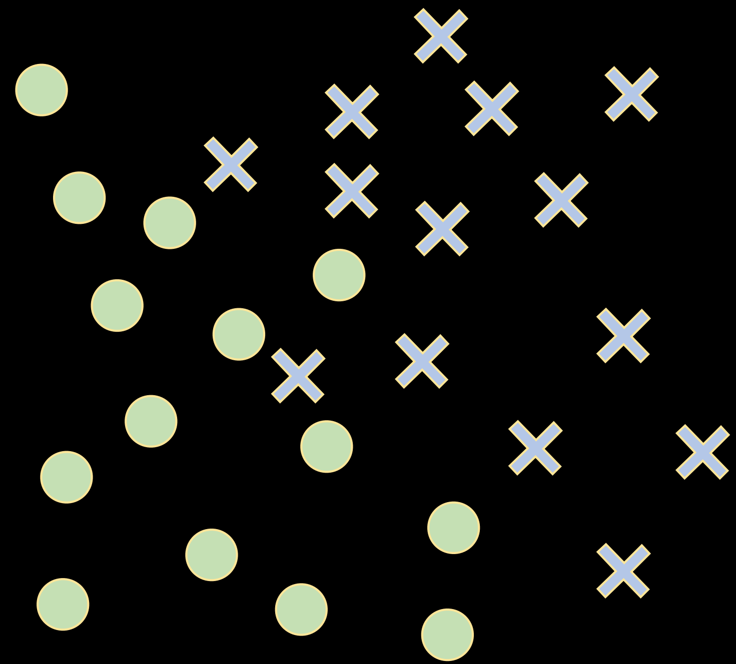

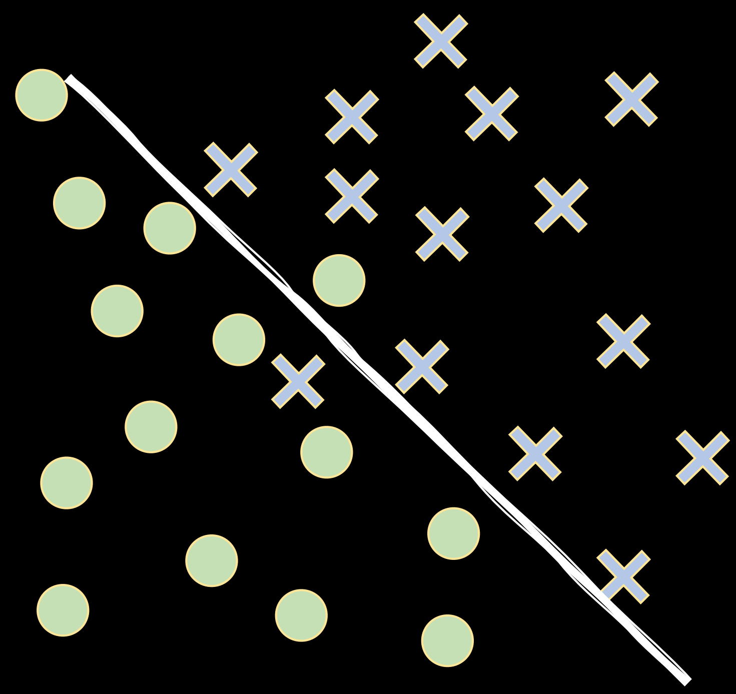

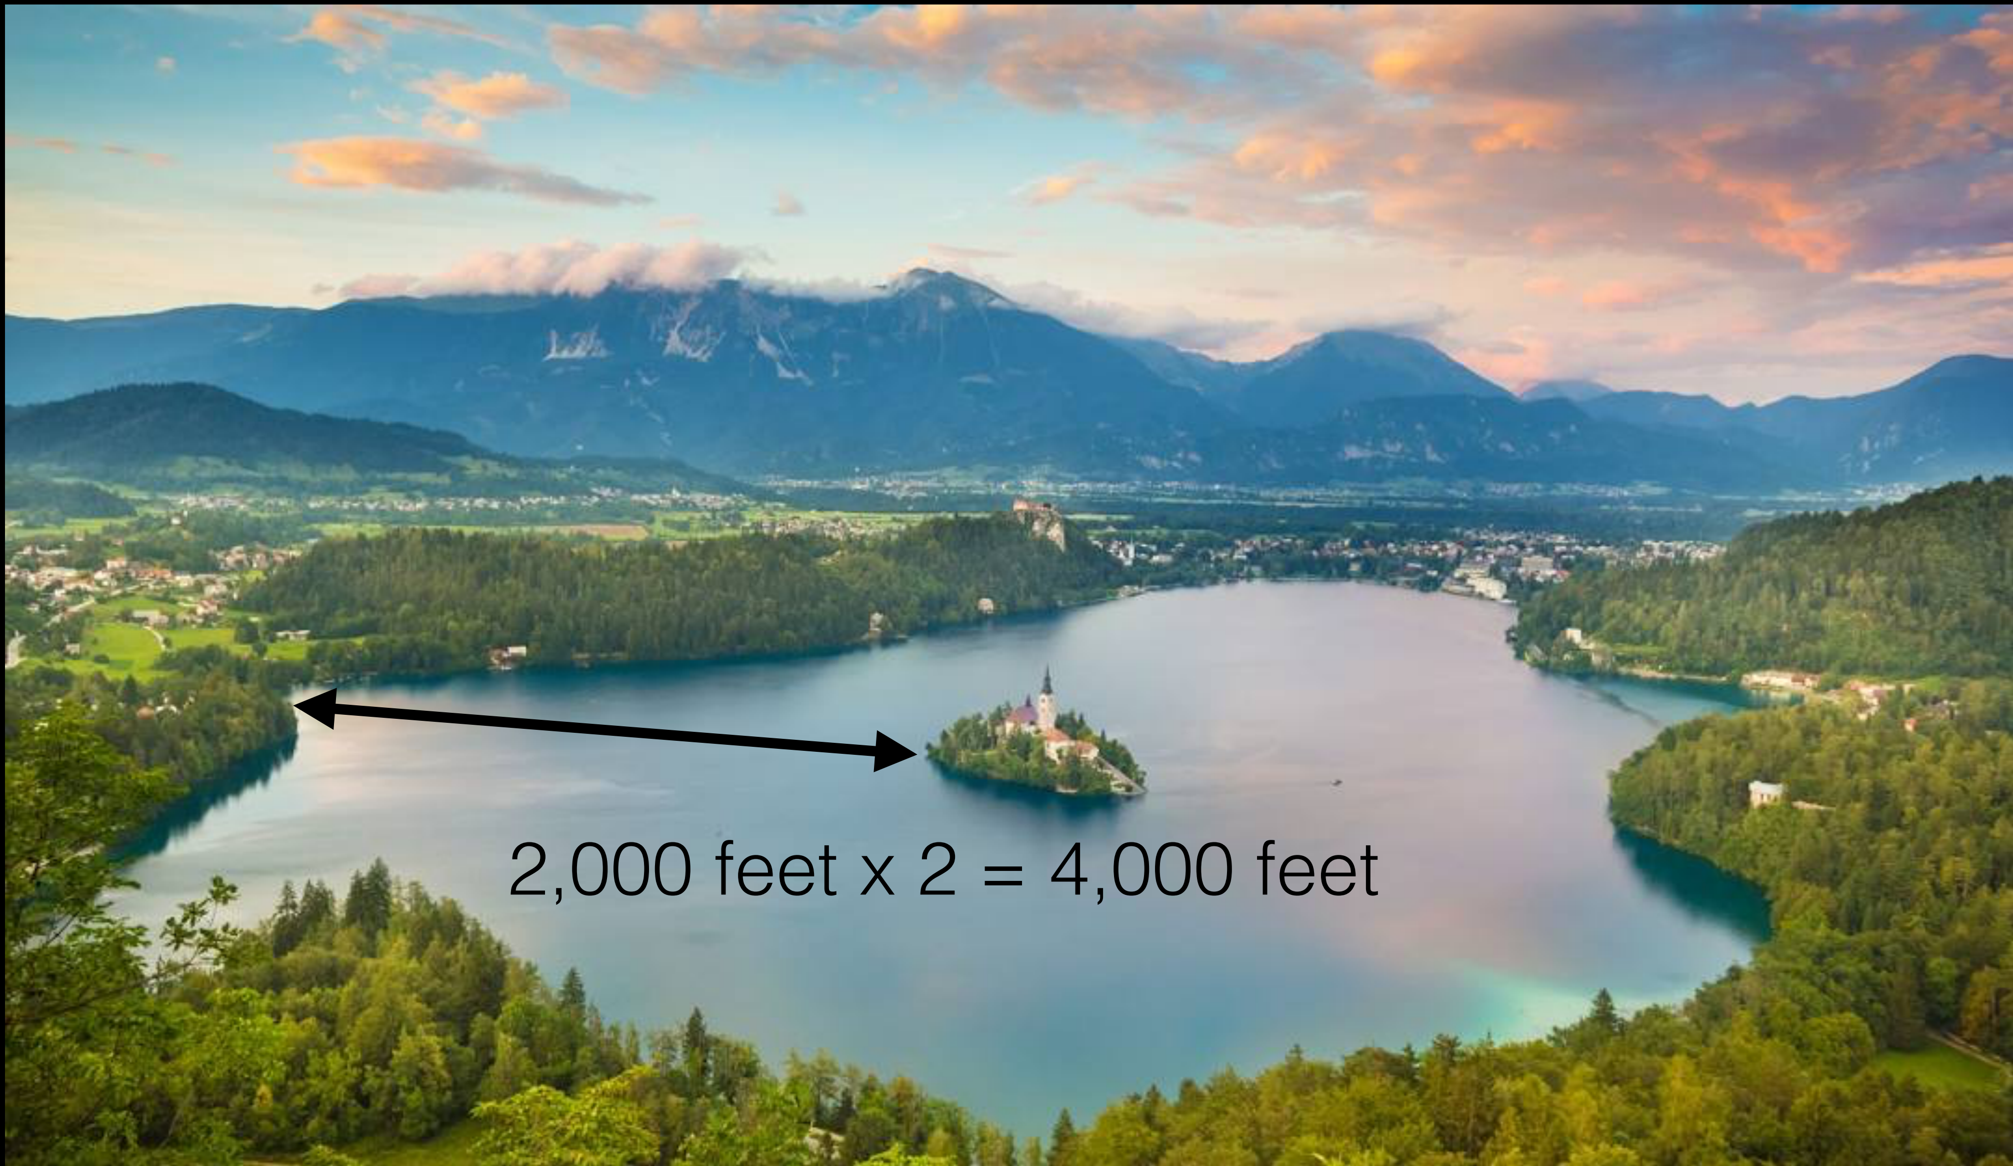
2,000 feet x 2 = 4,000 feet

11:06 PM

Your Assistant for Safe Swimming in Bled

How long did you sleep

7 hours

Hours/week workout?

at least 10 hours

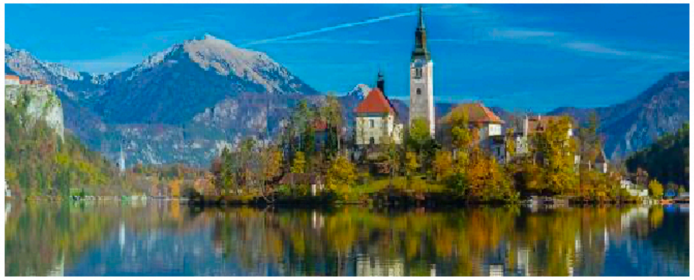

Tip: You can safely swim to the island and back. Enjoy!

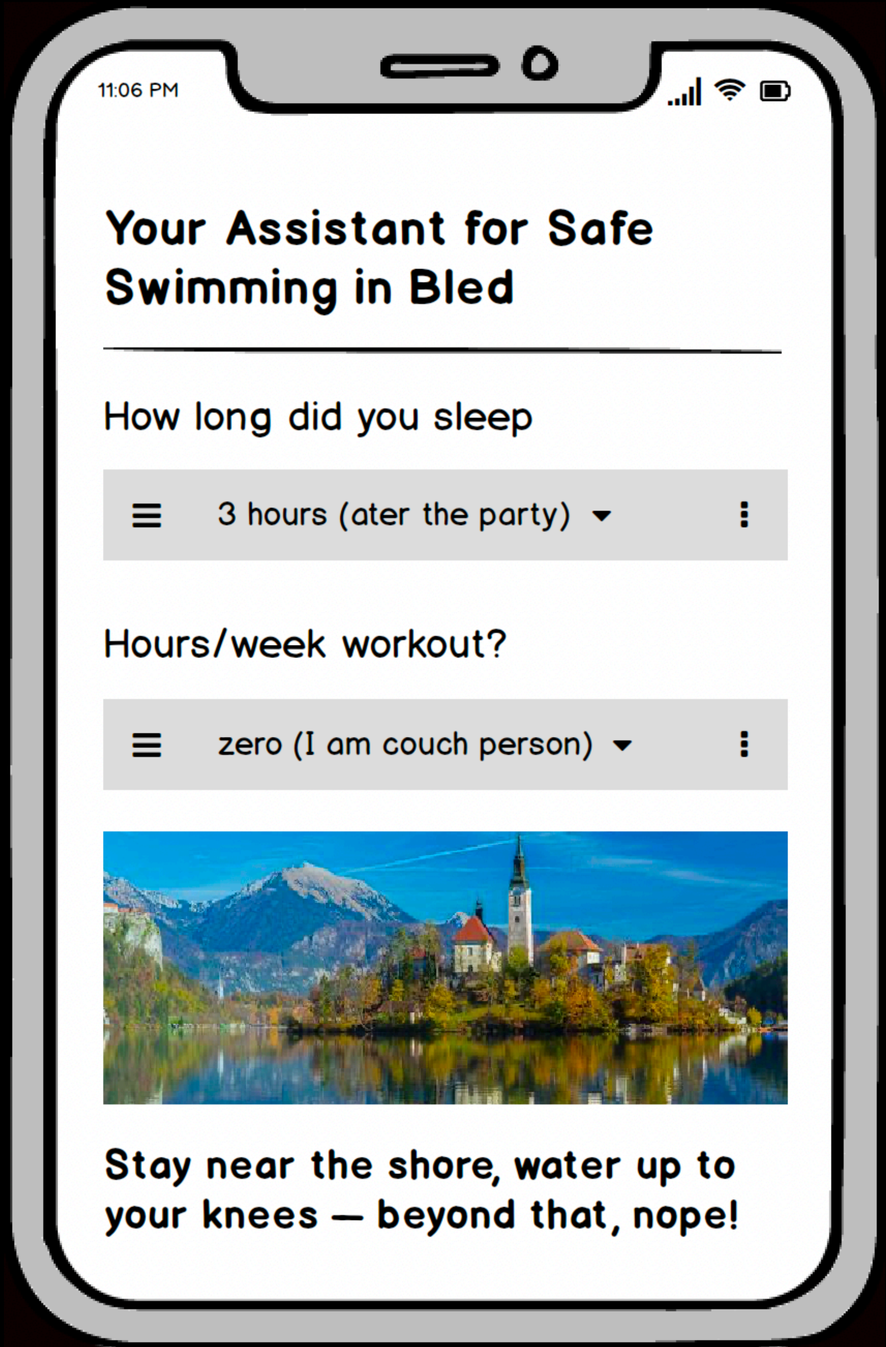
11:06 PM
Your Assistant for Safe Swimming in Bled
How long did you sleep
3 hours (ater the party)
Hours/week workout?
zero (I am couch person)
Stay near the shore, water up to your knees — beyond that, nope!

Data Table

	name	exercise	sleep	activity
1	Ana	2	8	bathing
2	Marko	4	7	bathing
3	Petra	3	6	bathing
4	Janez	5	6	bathing
5	Maja	9	6	bathing
6	Andrej	5	5	bathing
7	Nika	2	5	bathing
8	Luka	5	4	bathing
9	Katarina	4	3	bathing
10	Miha	6	4	bathing
11	Irena	9	3	bathing
12	Jure	8	4	bathing
13	Tjaša	11	3	bathing
14	Boštjan	7	5	bathing
15	Helena	4	9	swimming
16	Matej	4	8	swimming
17	Alenka	7	8	swimming
18	Gregor	10	9	swimming
19	Polona	9	8	swimming
20	Sašo	10	7	swimming

30 30 | 30

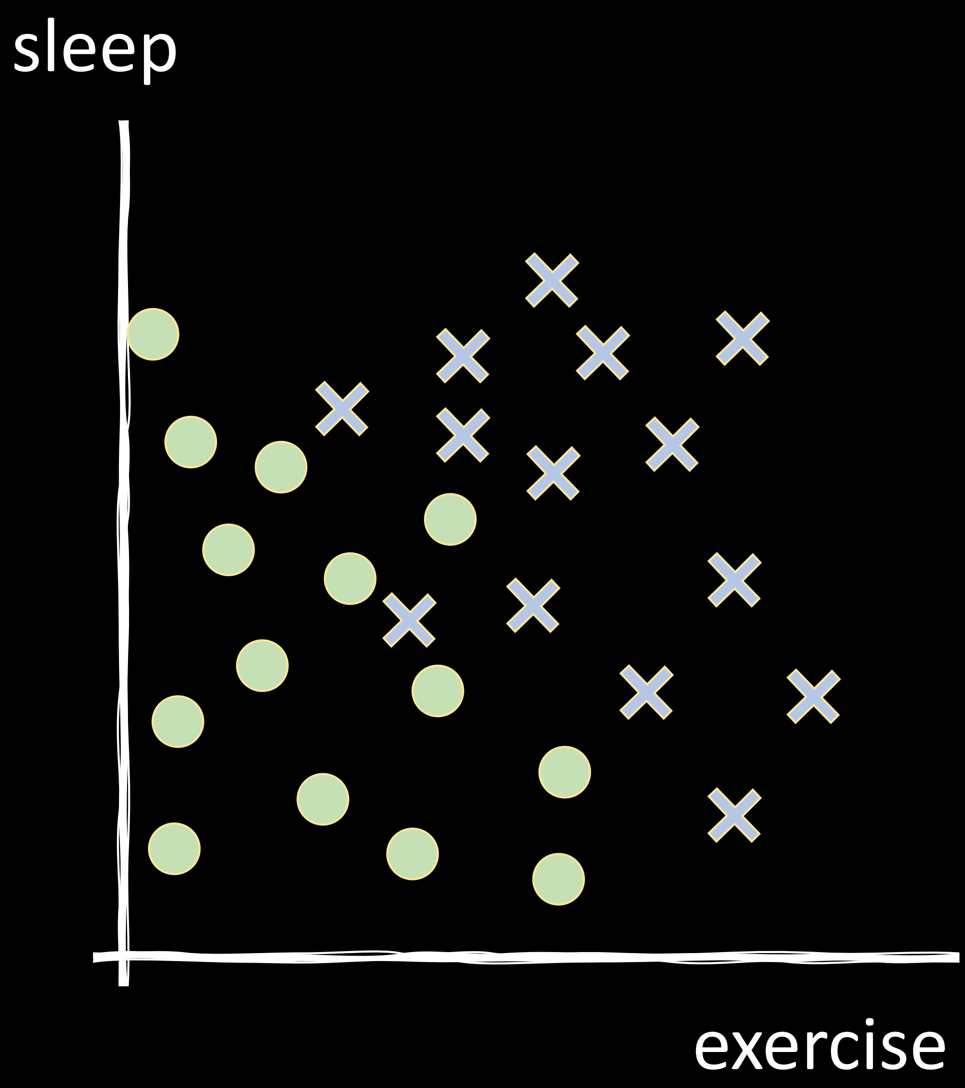

Data Table

	name	exercise	sleep	activity
1	Ana	2	8	bathing
2	Marko	4	7	bathing
3	Petra	3	6	bathing
4	Janez	5	6	bathing
5	Maja	9	6	bathing
6	Andrej	5	5	bathing
7	Nika	2	5	bathing
8	Luka	5	4	bathing
9	Katarina	4	3	bathing
10	Miha	6	4	bathing
11	Irena	9	3	bathing
12	Jure	8	4	bathing
13	Tjaša	11	3	bathing
14	Boštjan	7	5	bathing
15	Helena	4	9	swimming
16	Matej	4	8	swimming
17	Alenka	7	8	swimming
18	Gregor	10	9	swimming
19	Polona	9	8	swimming
20	Sašo	10	7	swimming

30 30 | 30

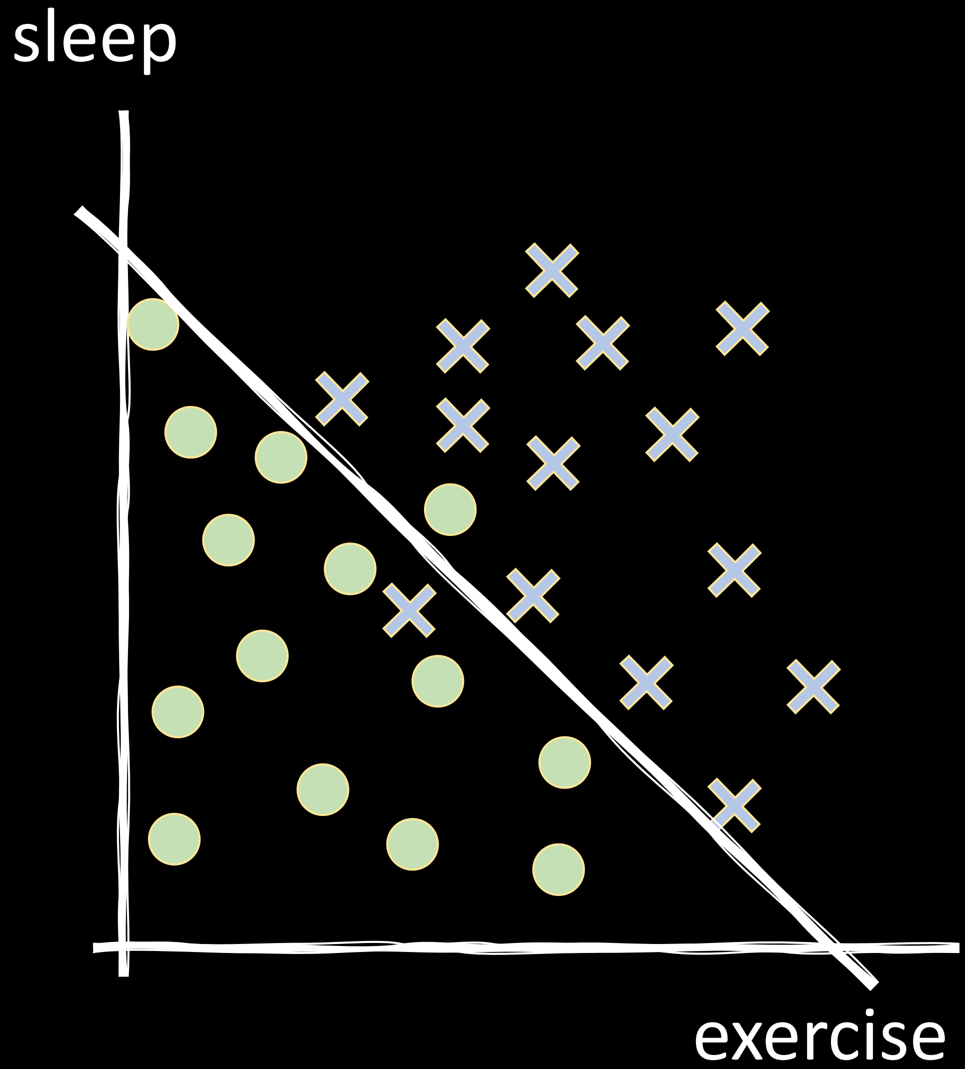
sleep
exercise

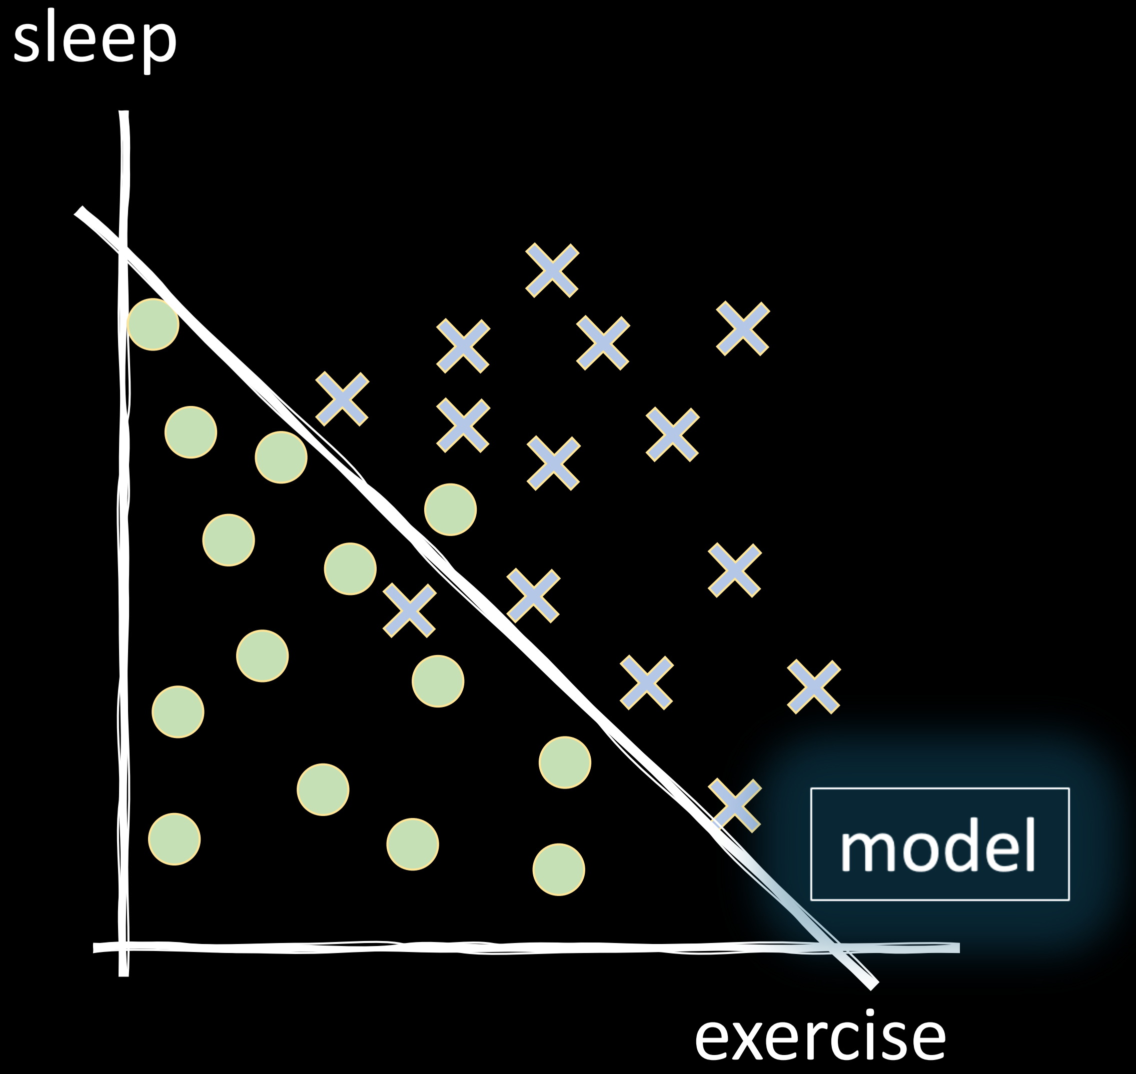
sleep
model
exercise

$$-40 + 2 \cdot \text{exercise} + 4 \cdot \text{sleep} = 0$$

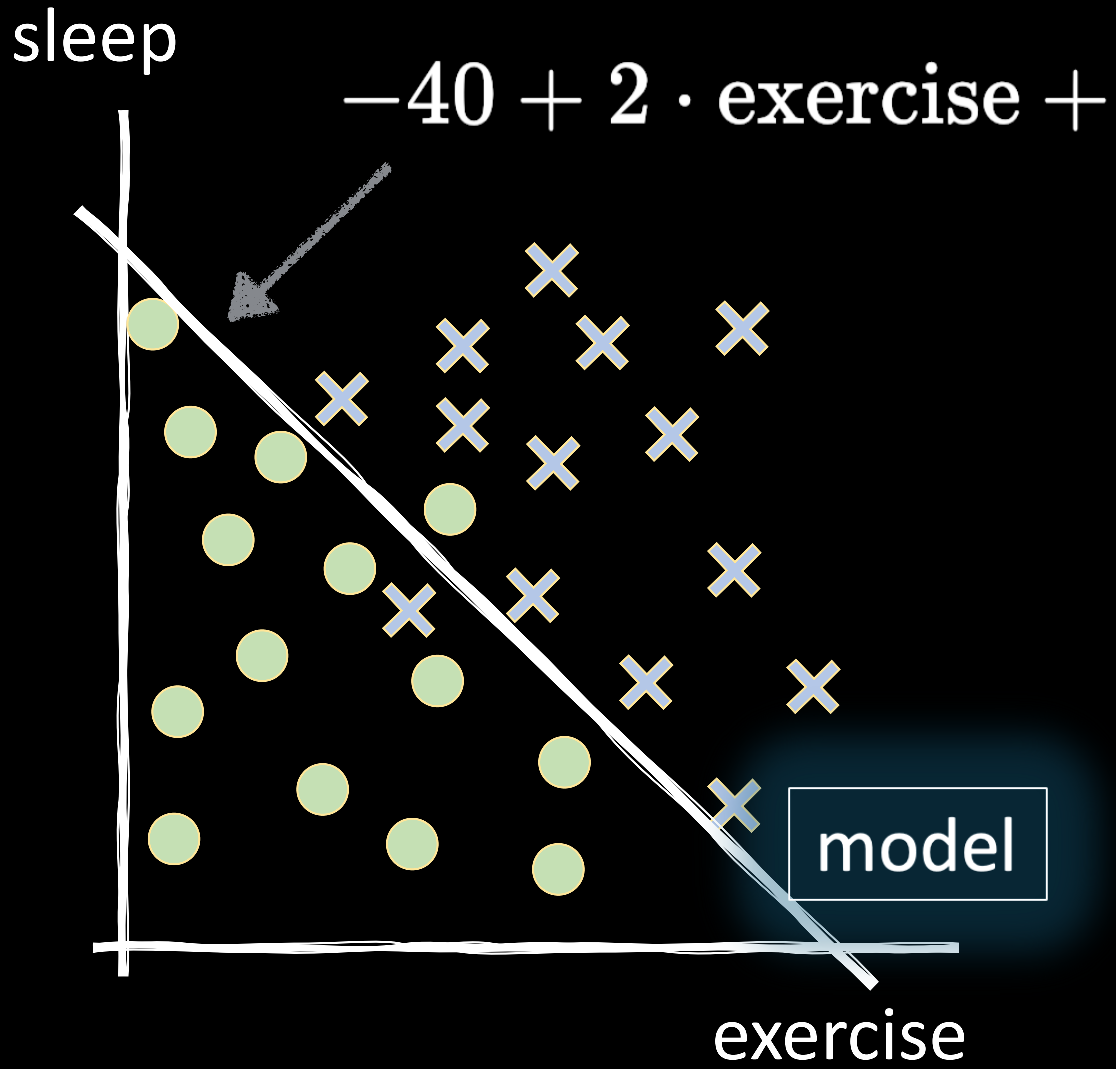

distance to the line

$$-40 + 2 \cdot \text{exercise} + 4 \cdot \text{sleep} = z$$

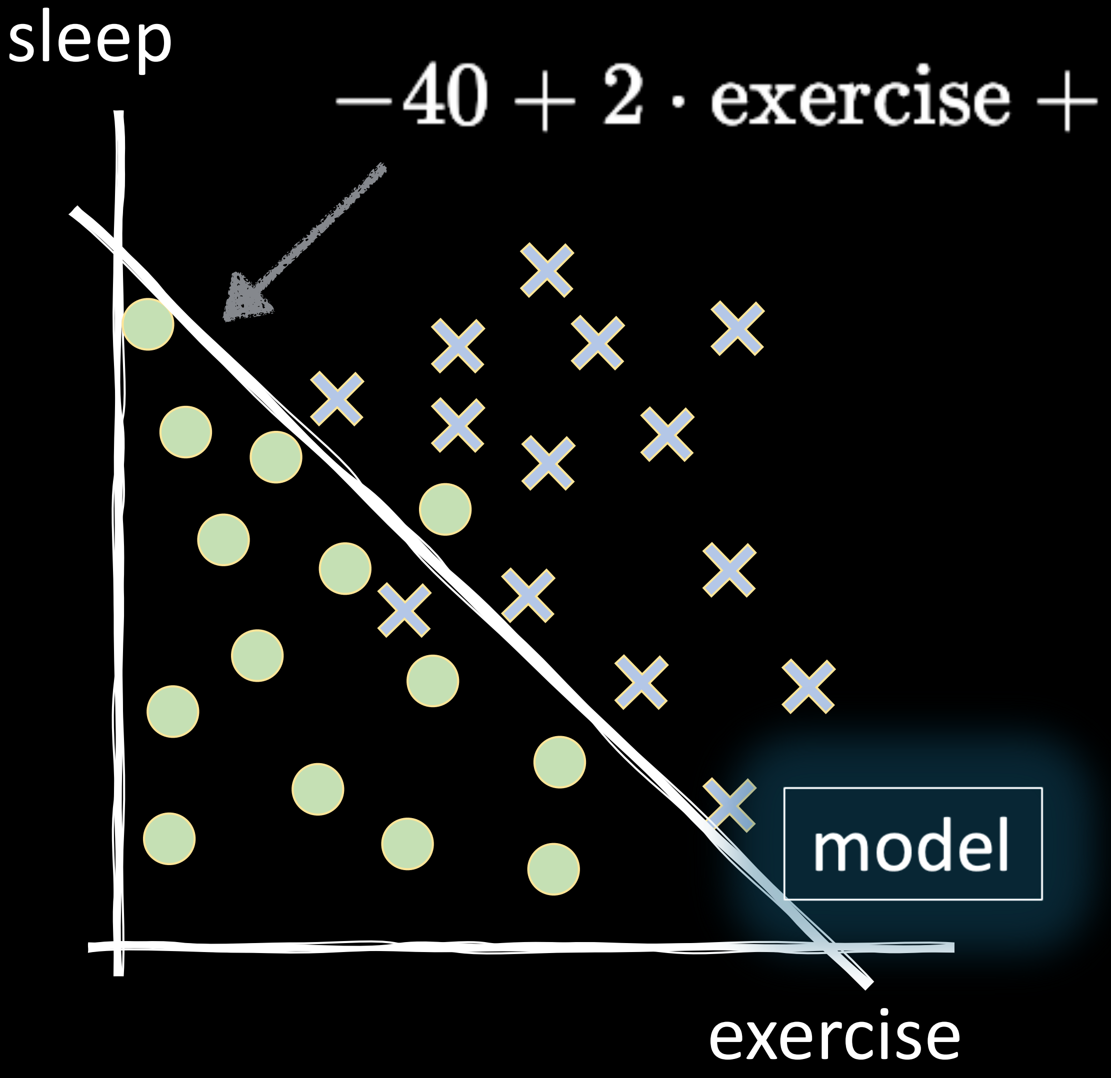

$$-40 + 2 \cdot \text{exercise} + 4 \cdot \text{sleep} = z$$

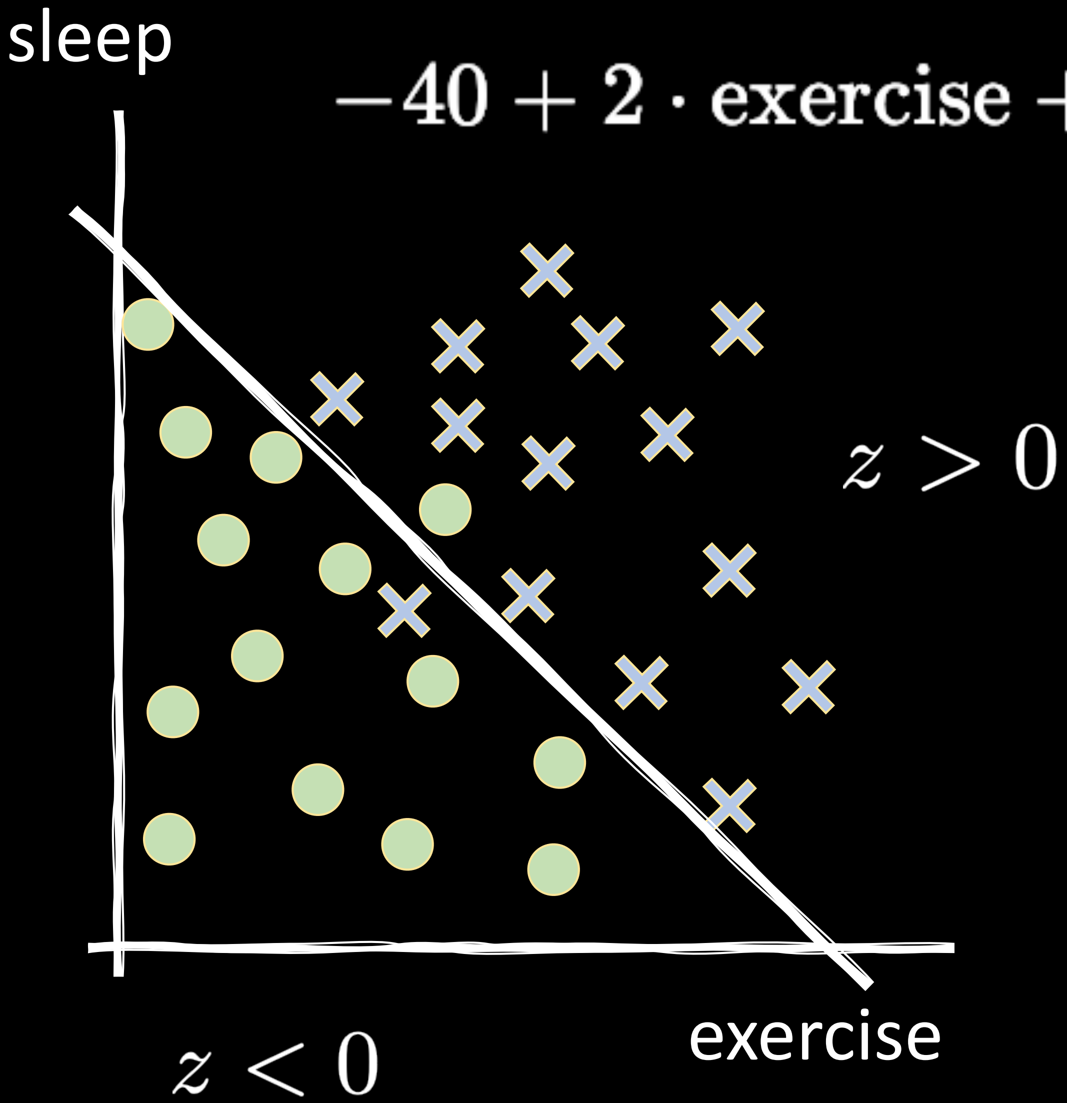

$$-40 + 2 \cdot \text{exercise} + 4 \cdot \text{sleep} = z$$

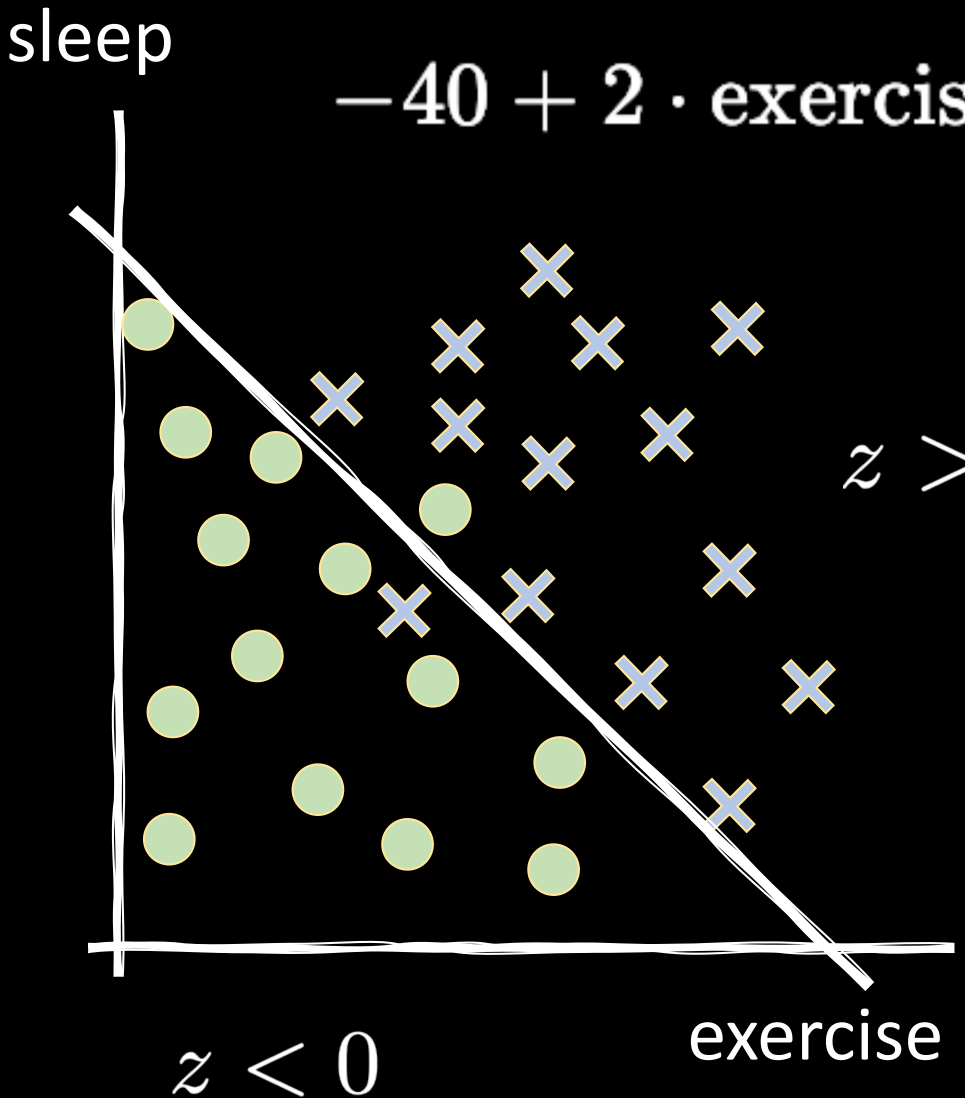

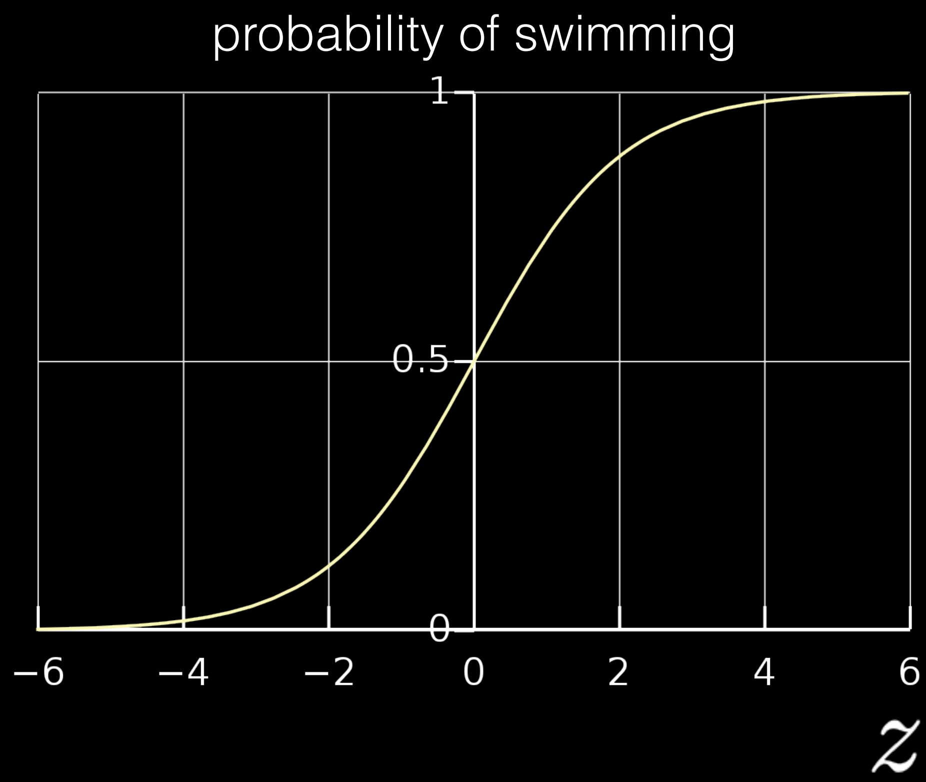

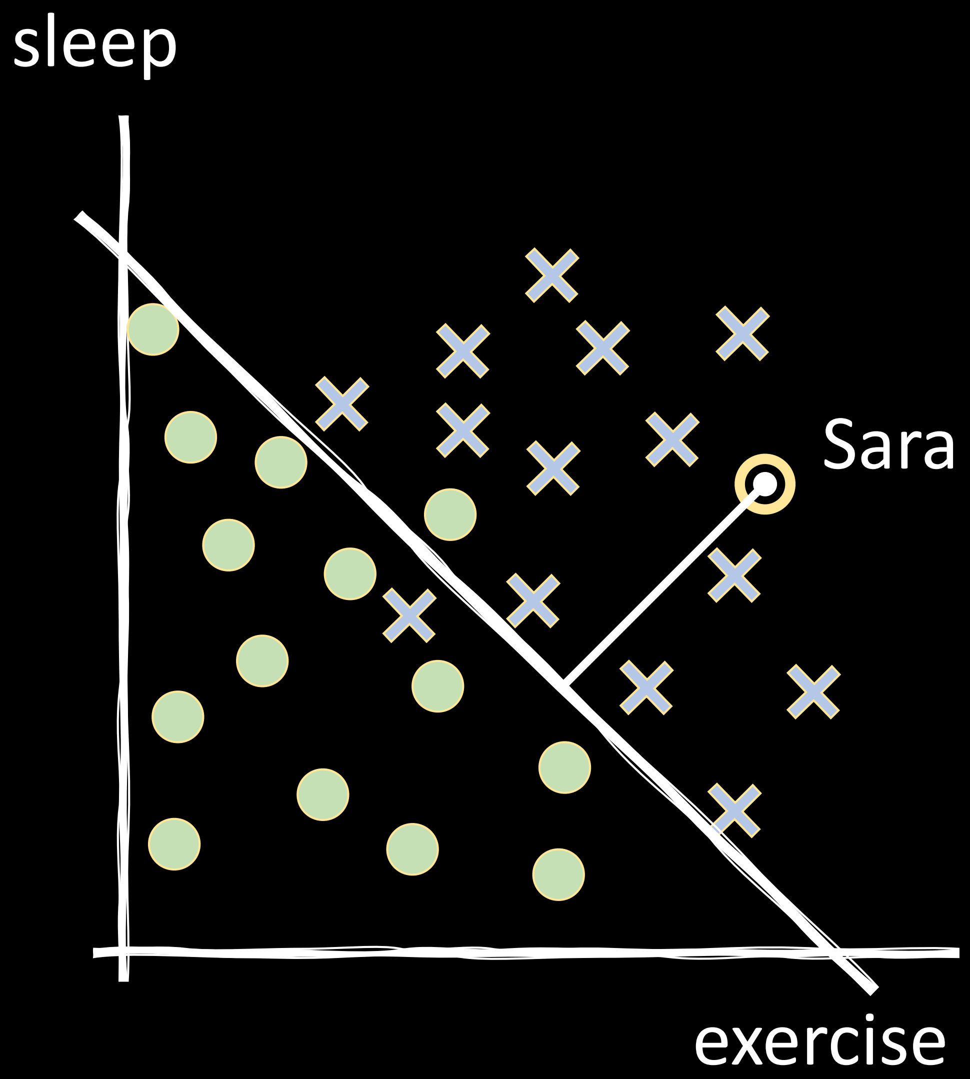
sleep
Sara
exercise

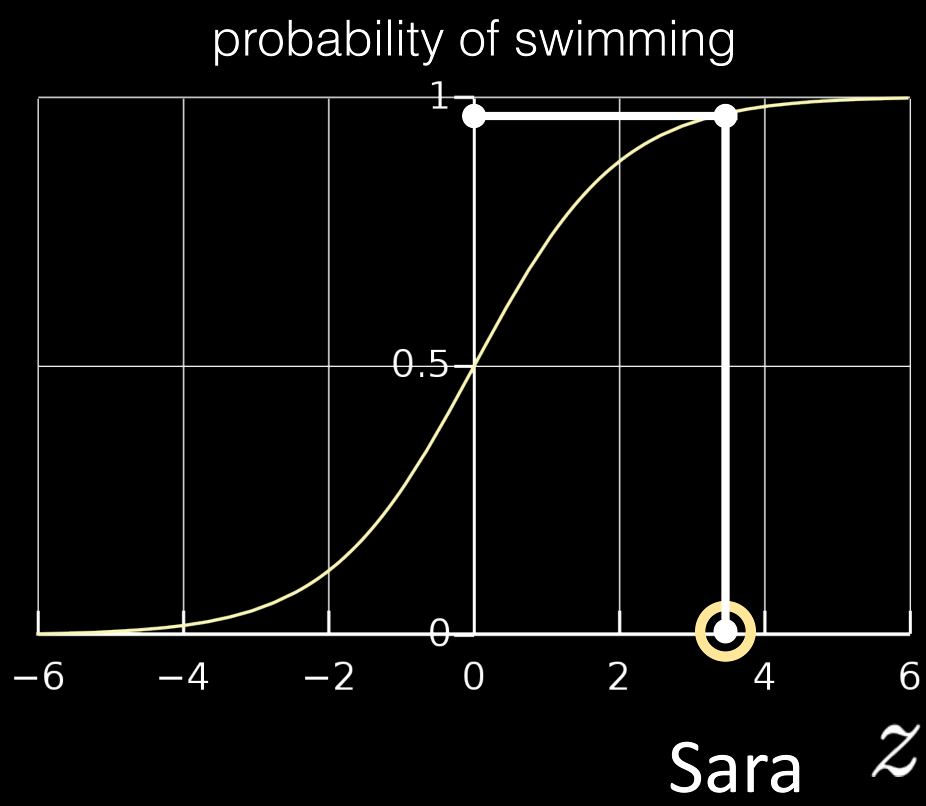
probability of swimming
1
0.5
0
−6
−4
−2
0
2
4
6
Sara
z

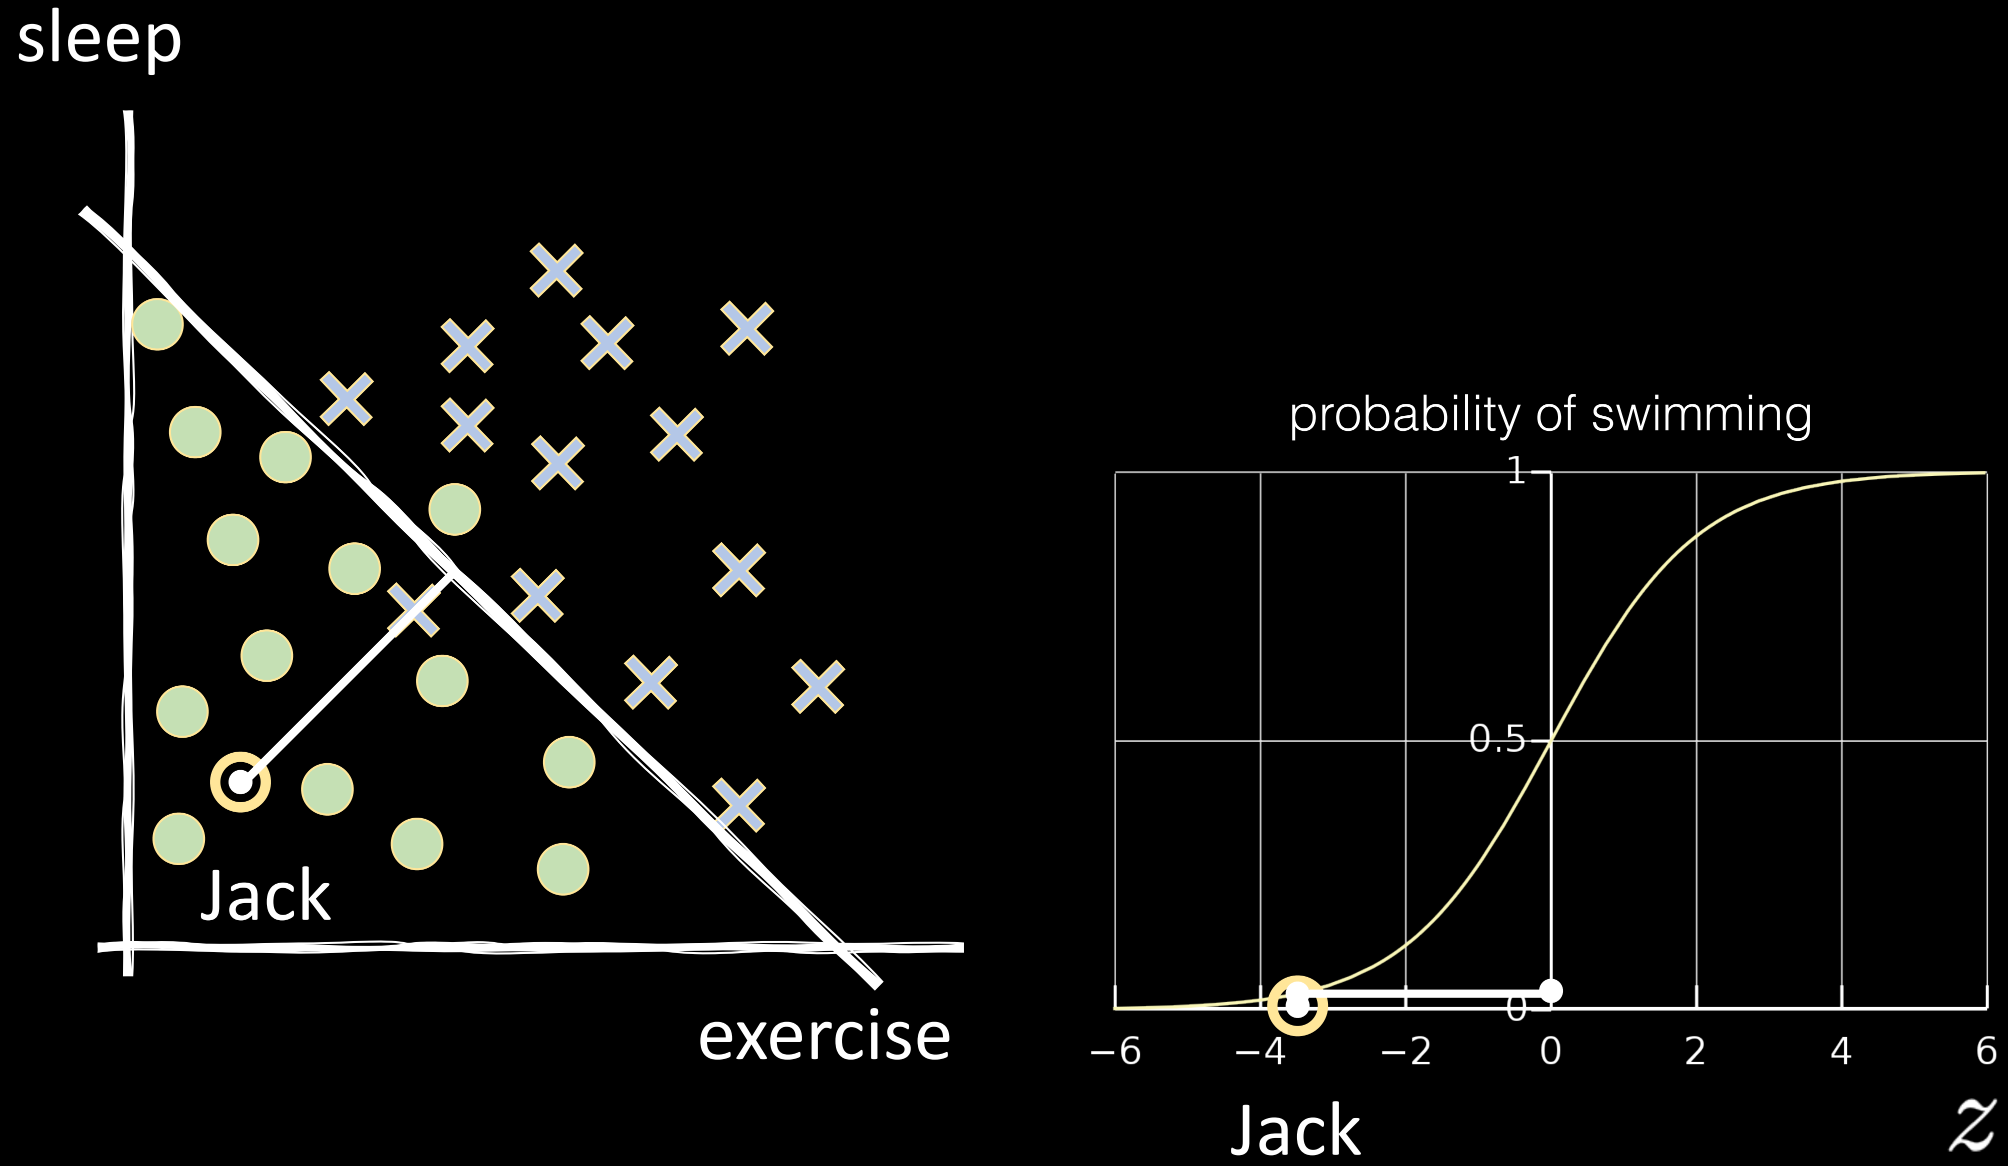
sleep
Jack
exercise
probability of swimming
1
0.5
0
−6
−4
−2
0
2
4
6
Jack
z

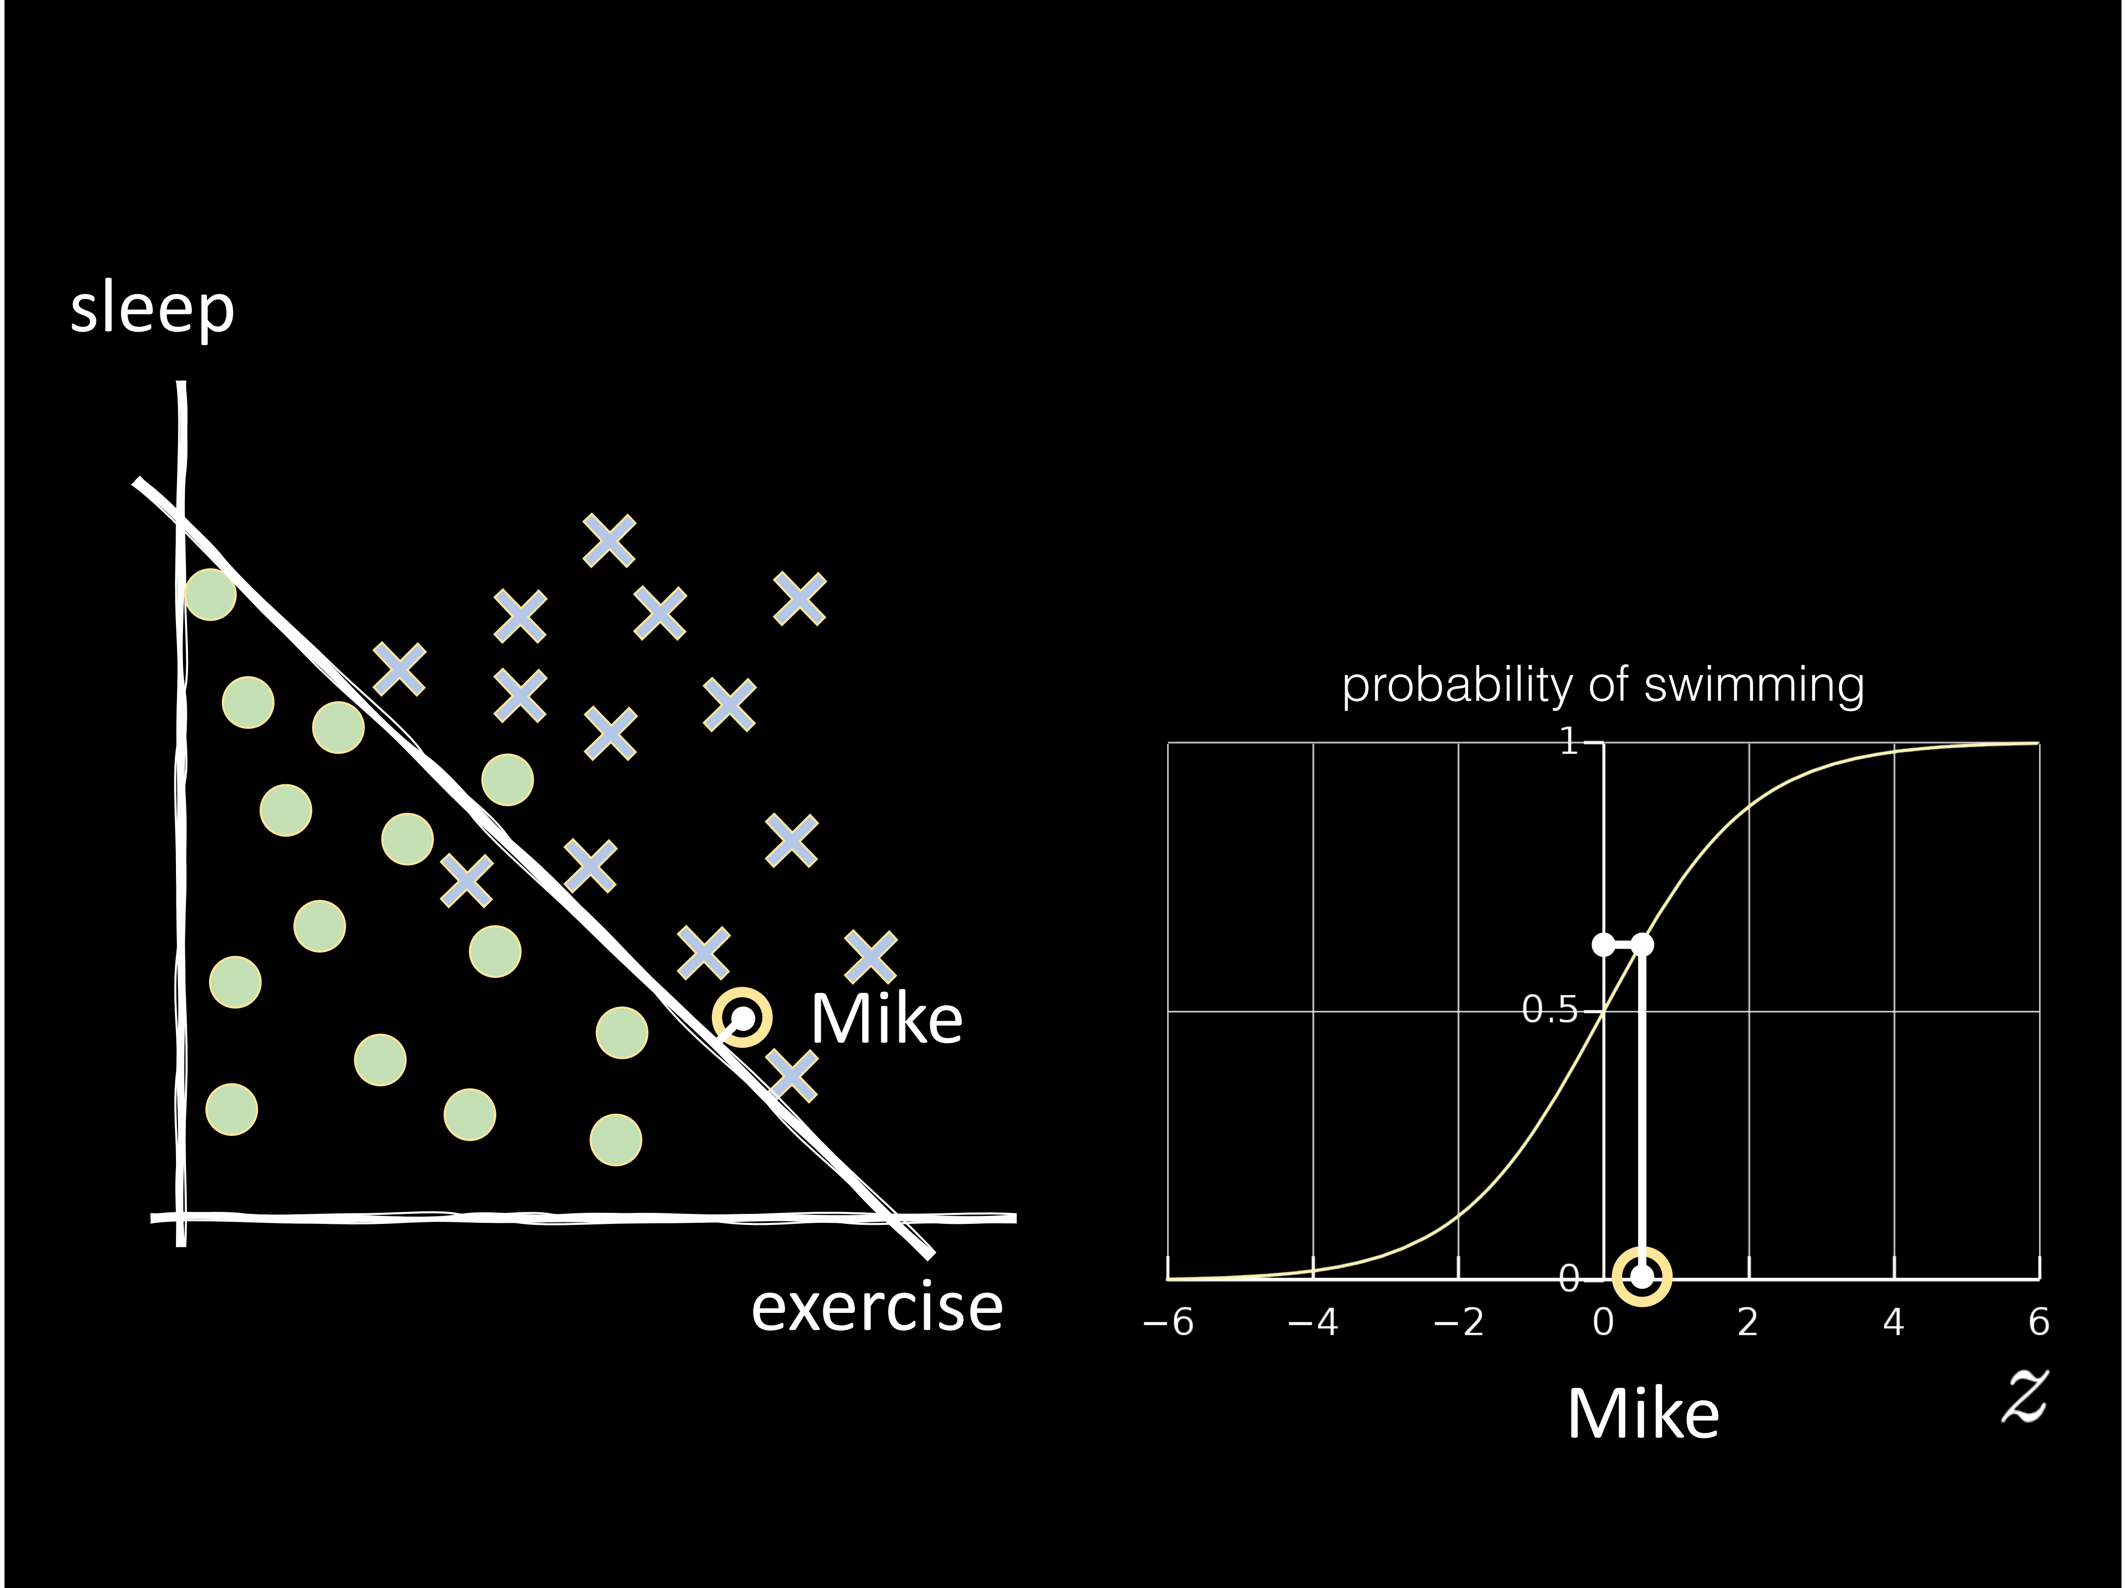

sleep
exercise
Mike
probability of swimming
1
0.5
0
−6
−4
−2
0
2
4
6
Mike
z

$$-40 + 2 \cdot \text{exercise} + 4 \cdot \text{sleep} = z$$

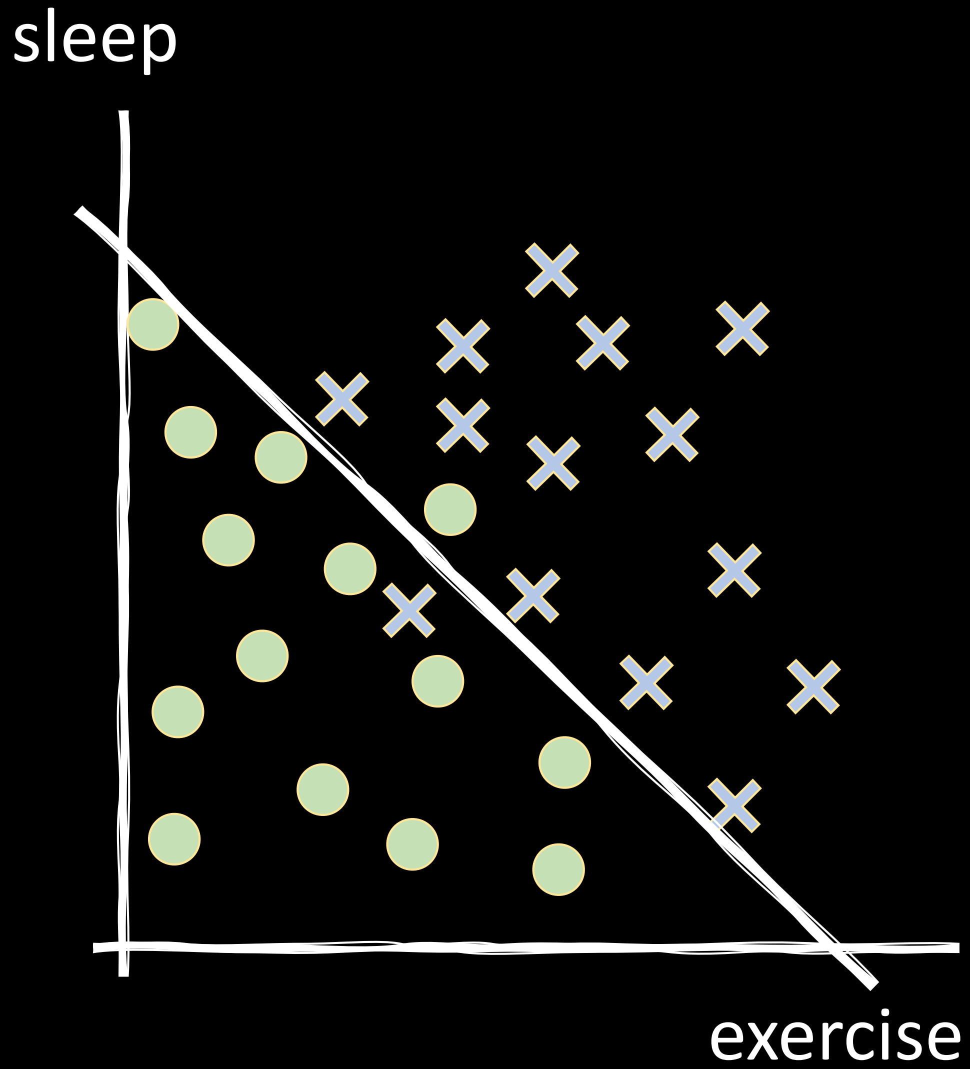

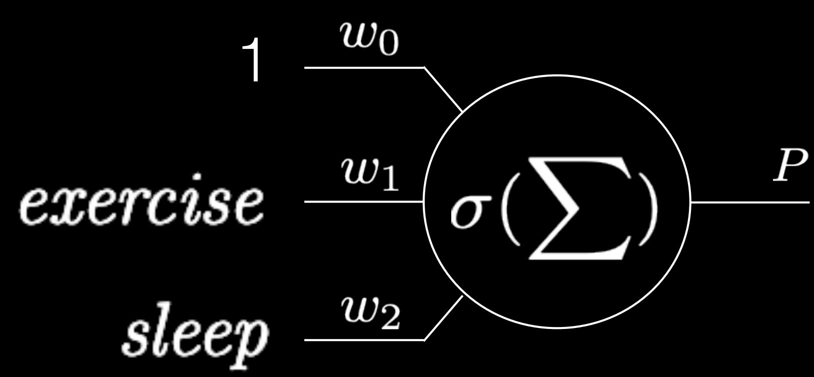

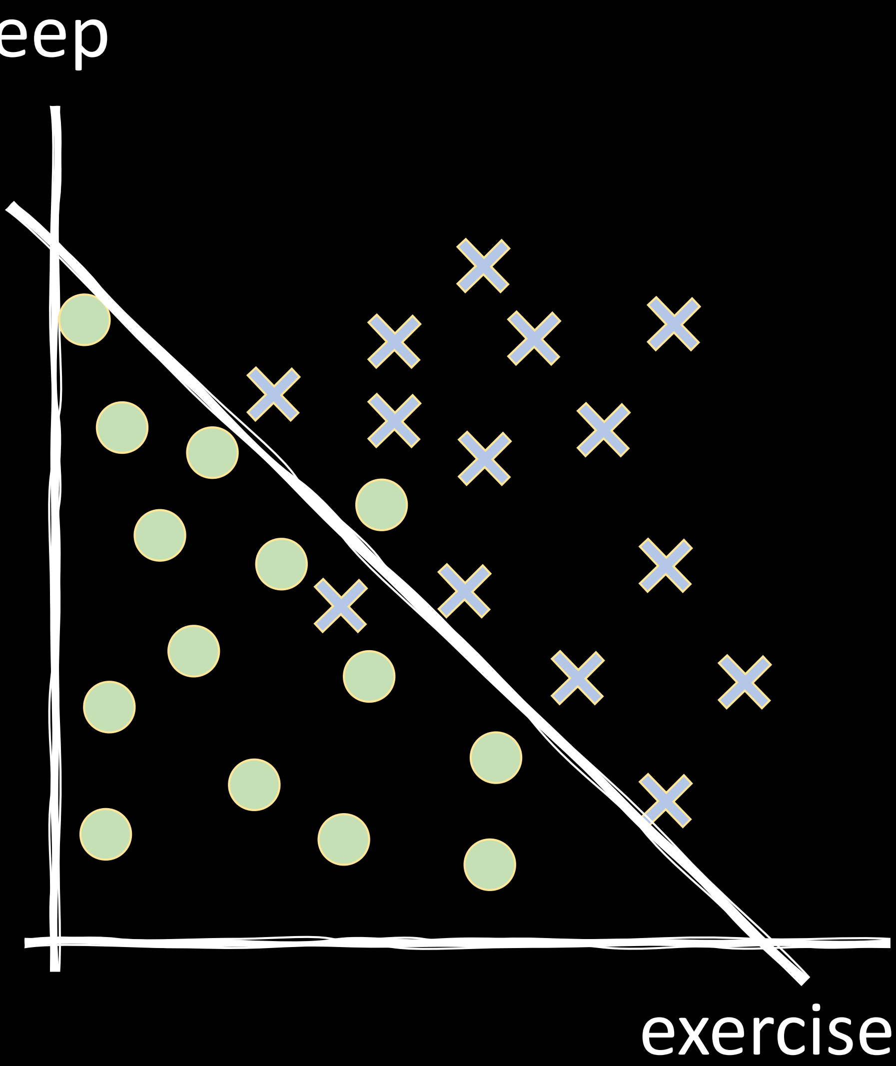

logistic regression

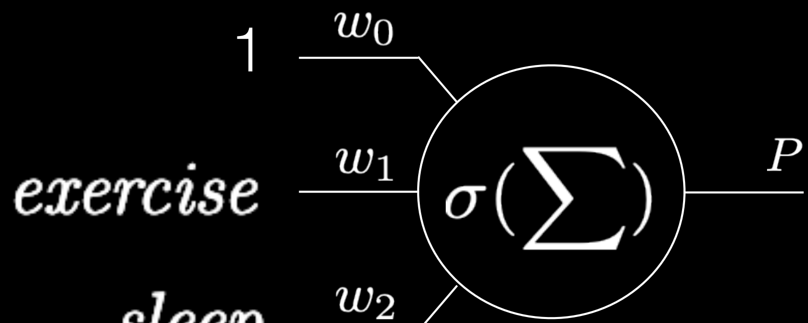

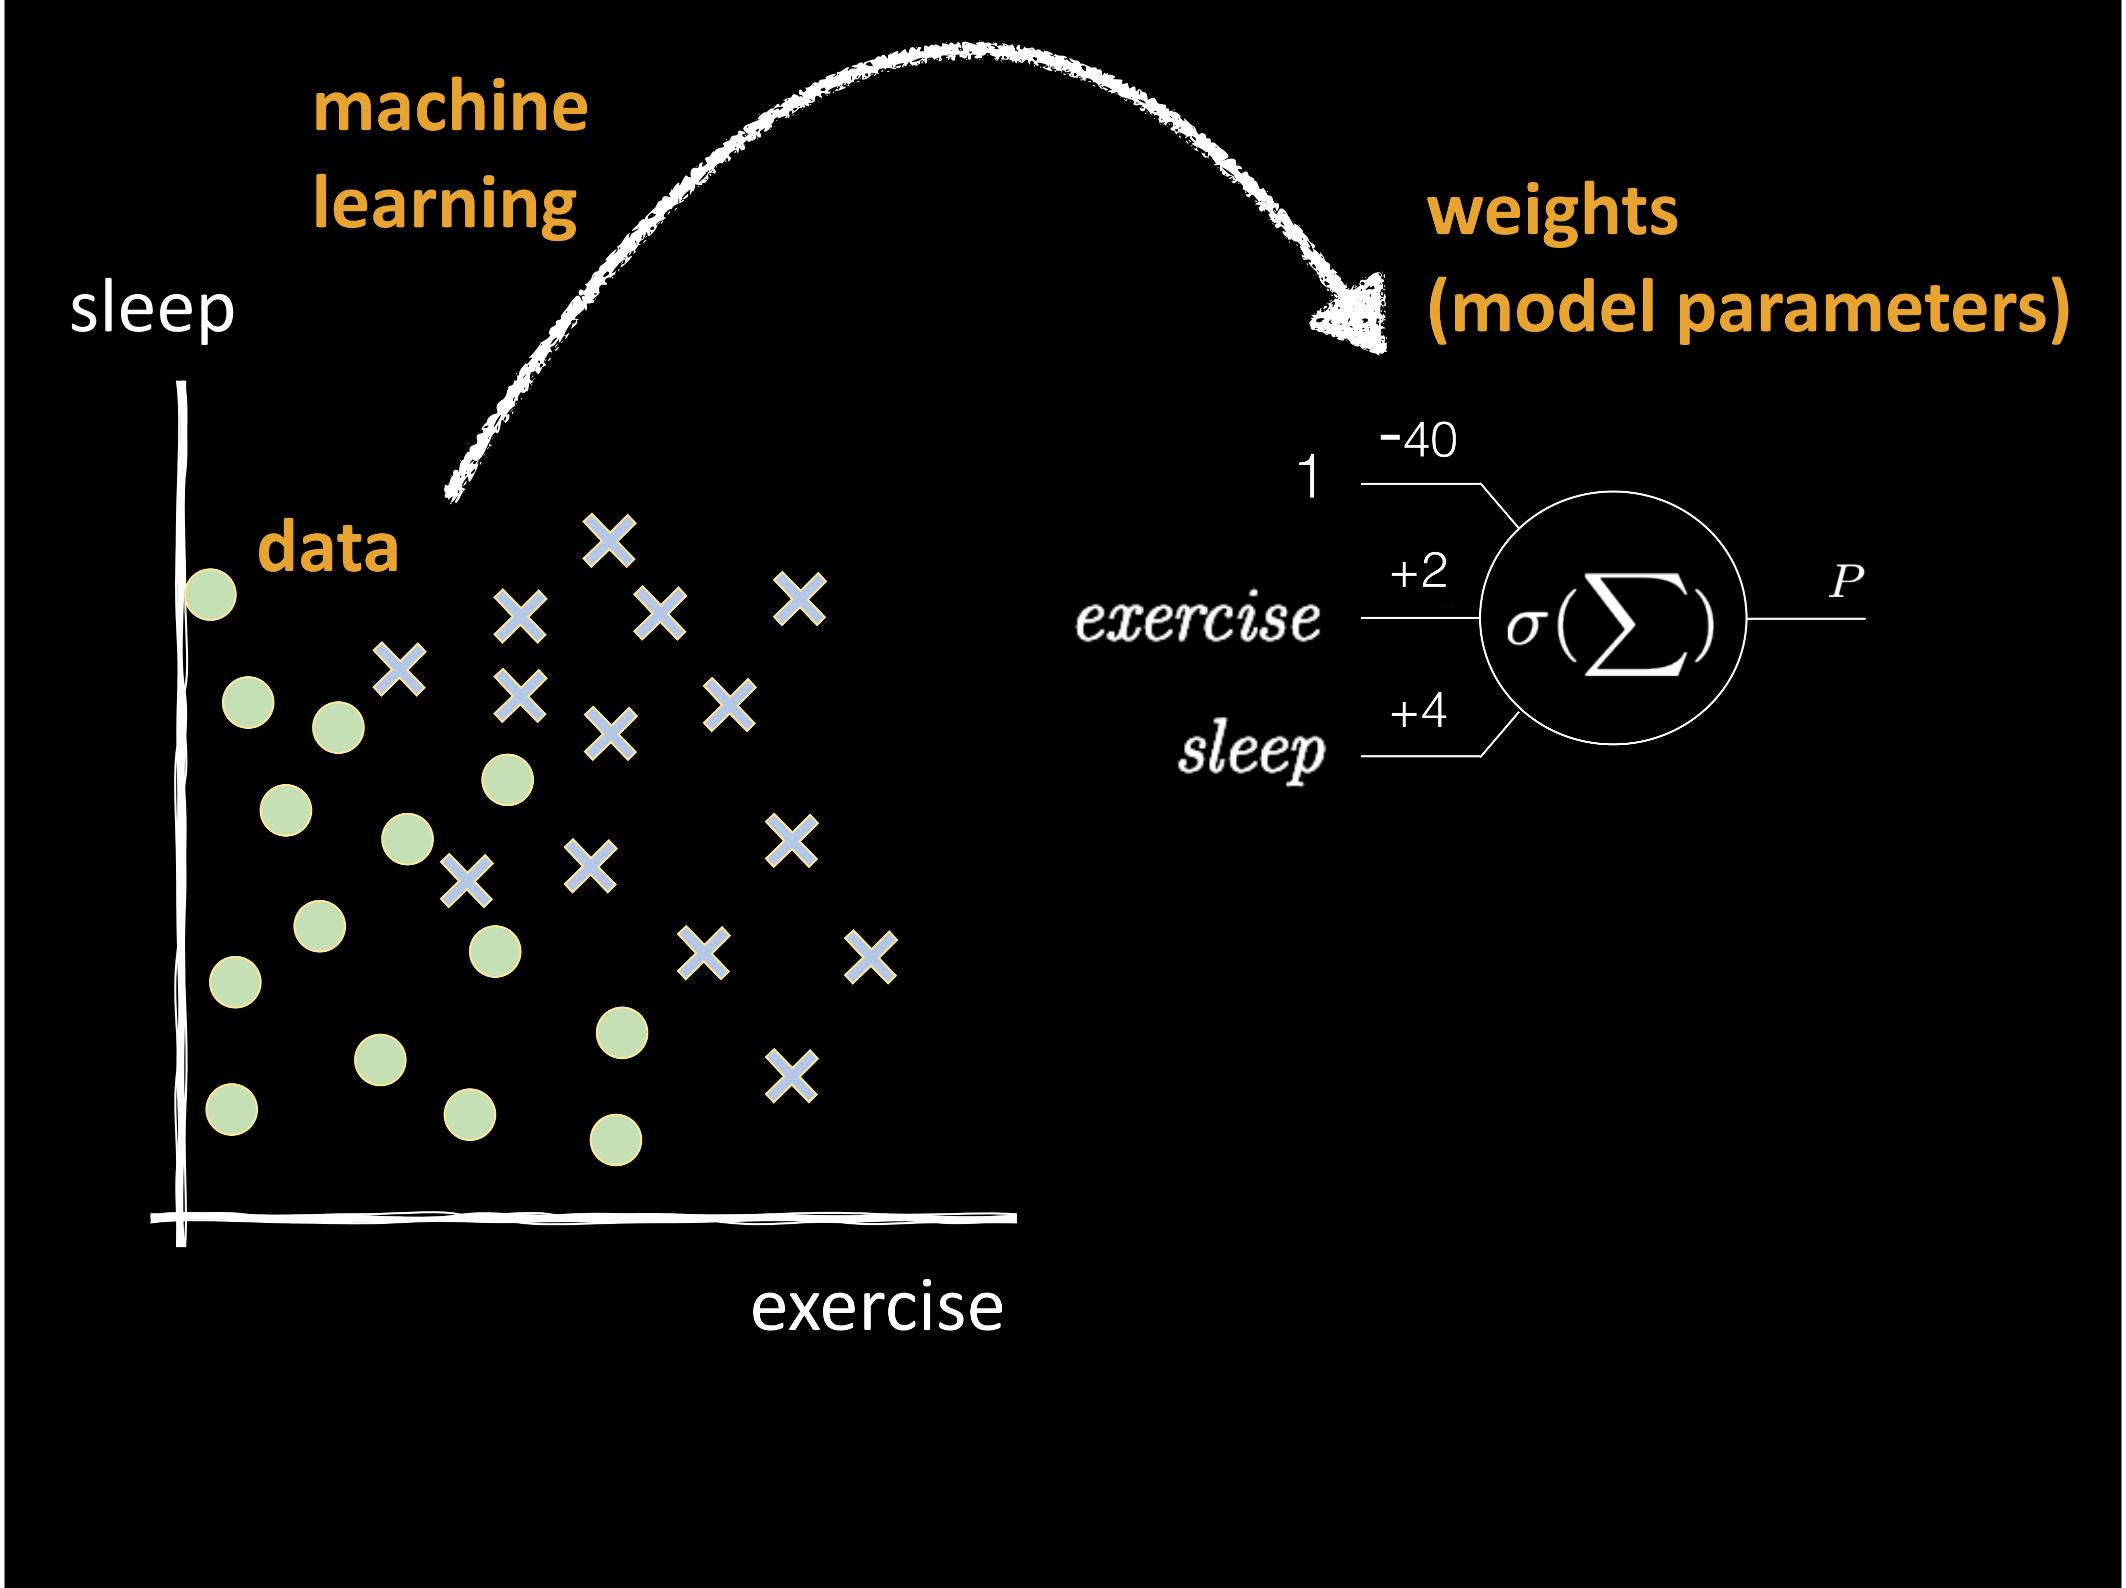
machine
learning
weights
(model parameters)
sleep
data
1
-40
exercise
+2
sleep
+4
$\sigma(\sum)$
P
exercise

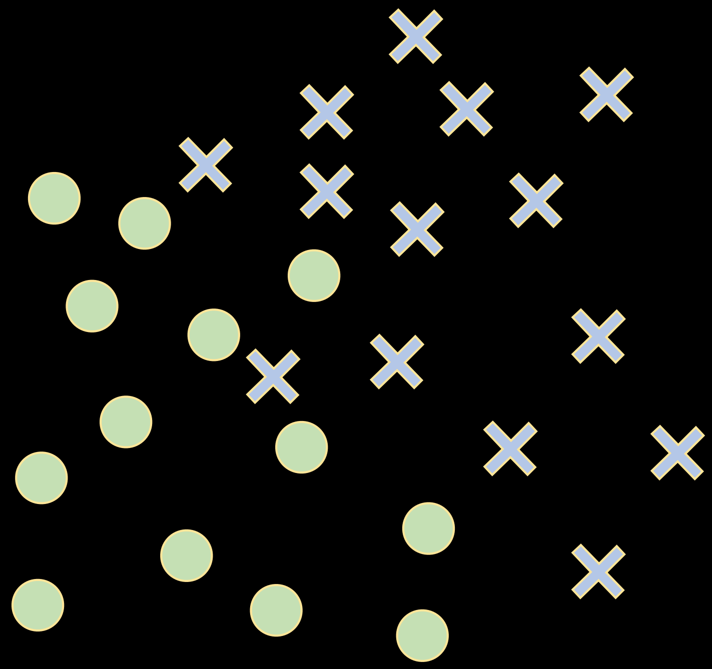

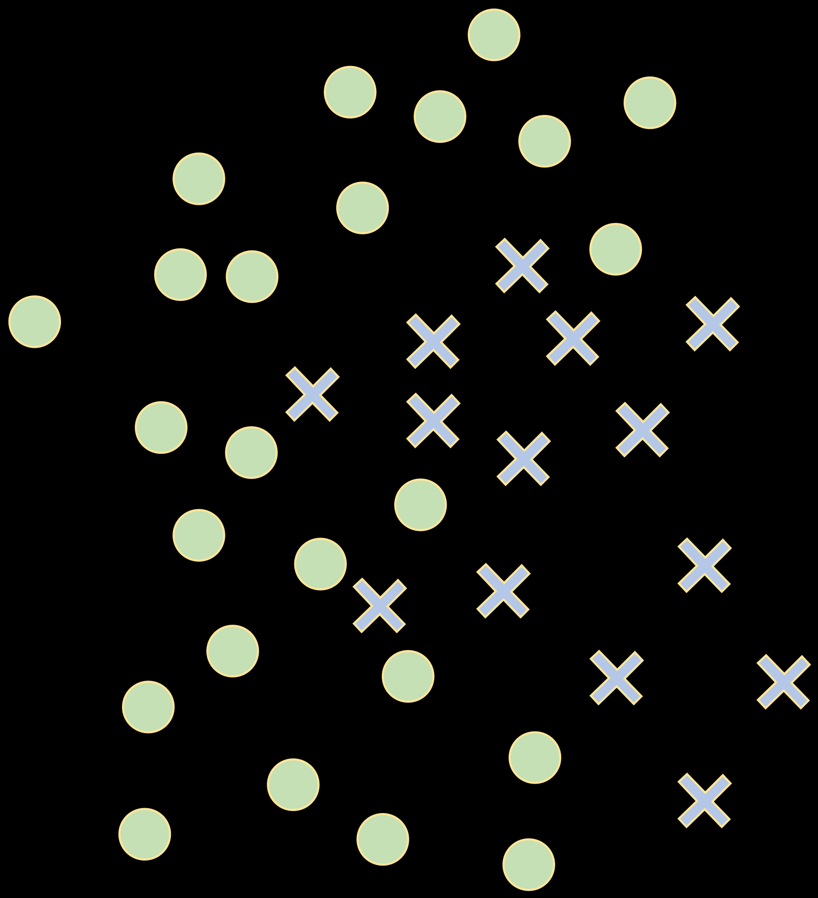

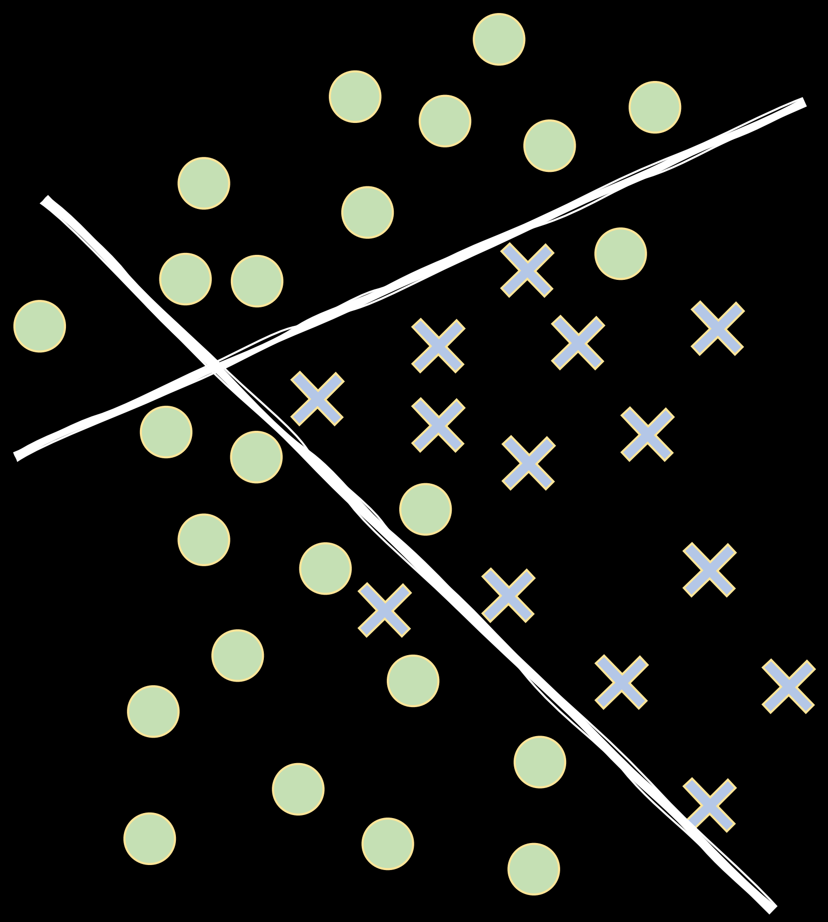

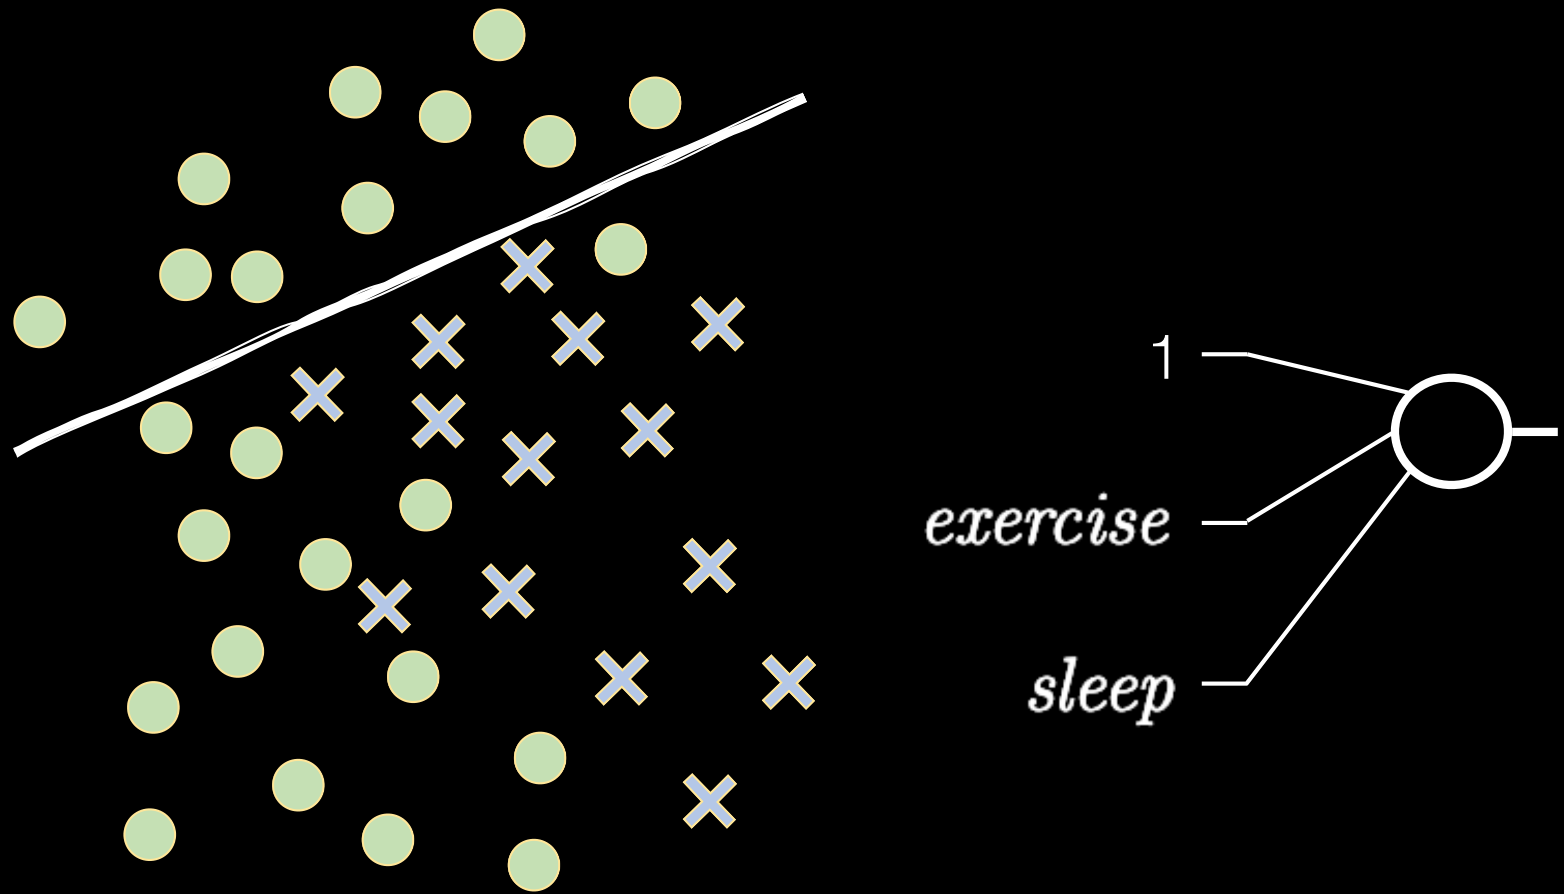
1
exercise
sleep

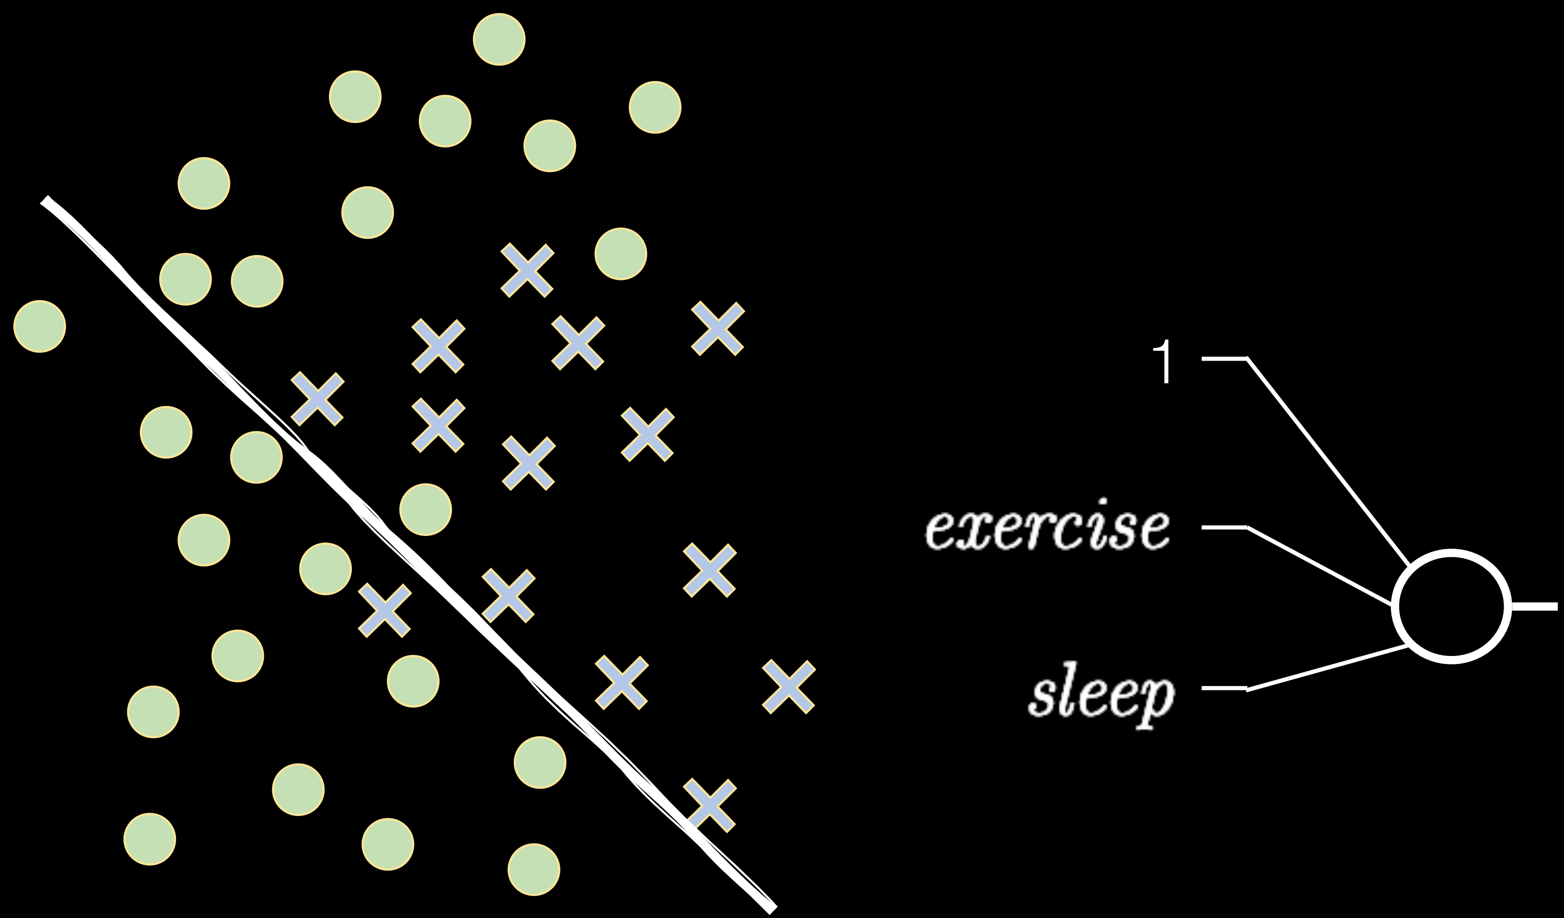

1
exercise
sleep

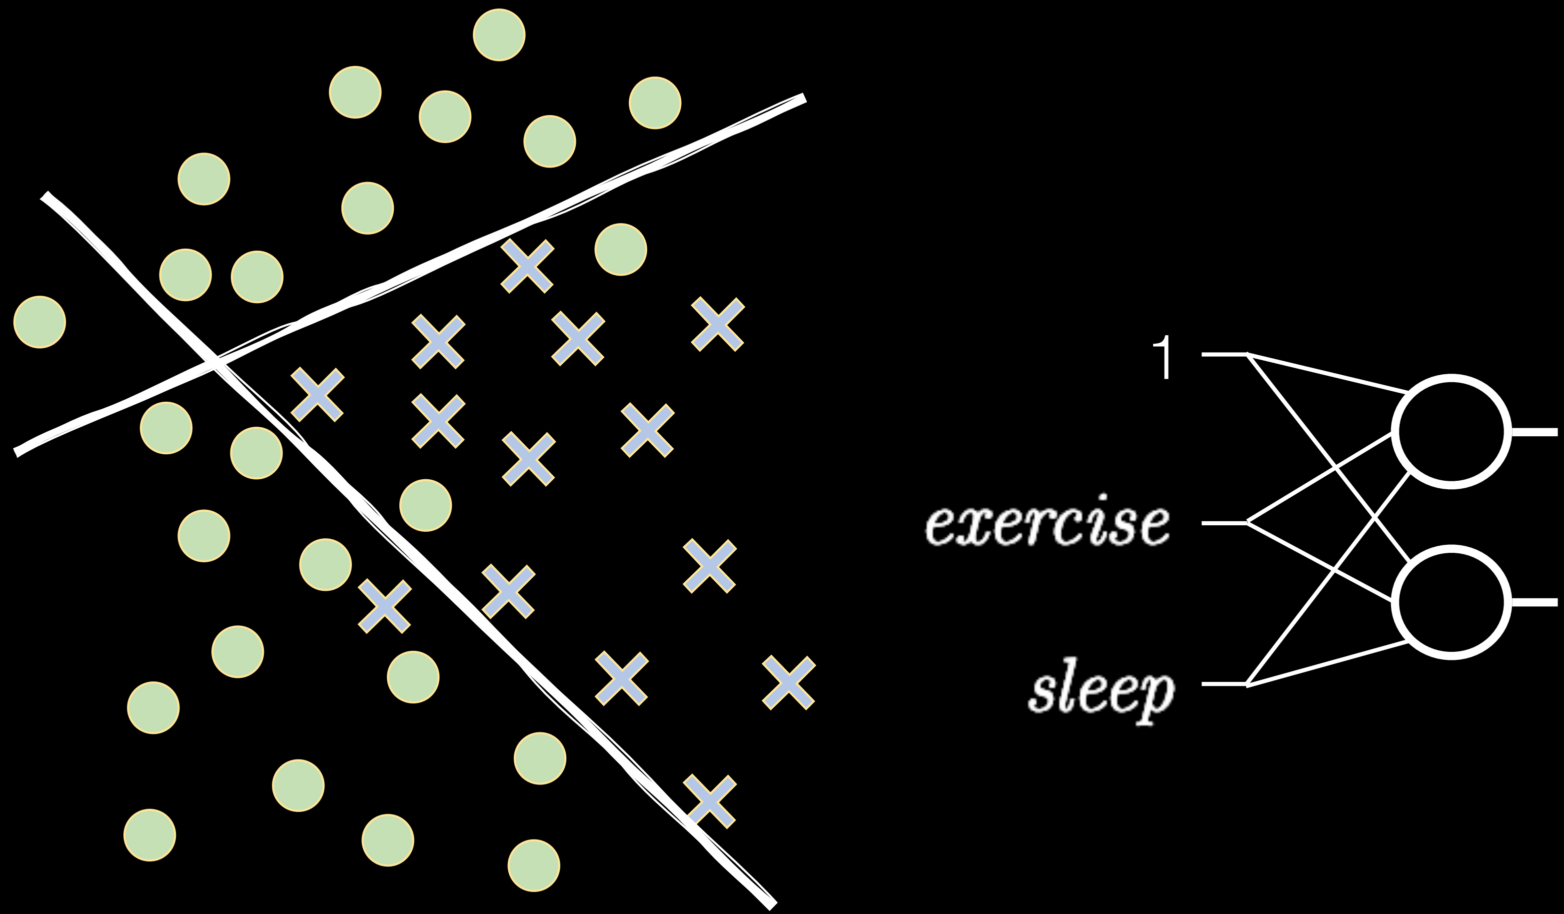
1
exercise
sleep

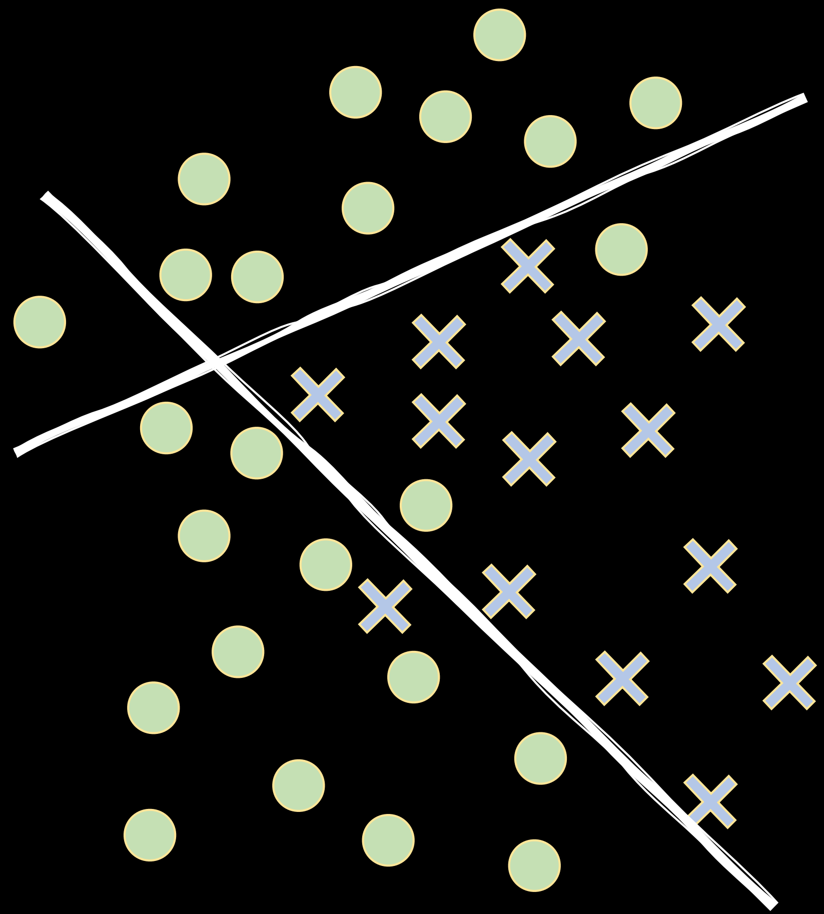

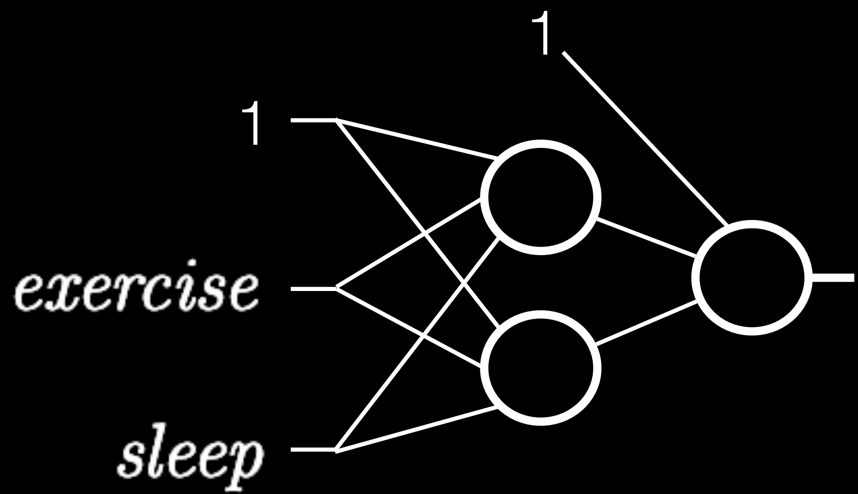
1
1
exercise
sleep

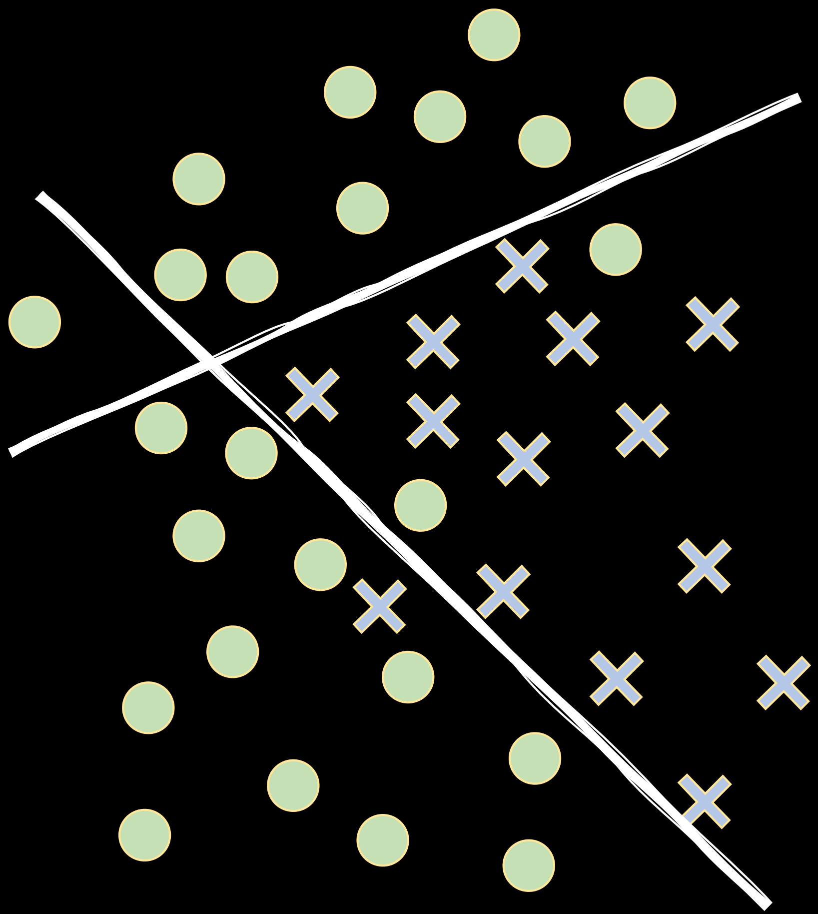

neural network

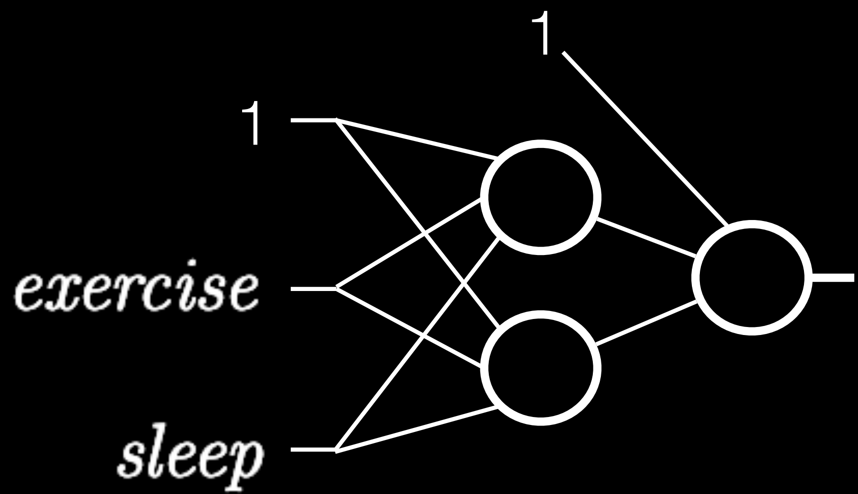

6 + 3 = 9 weights
(model parameters)

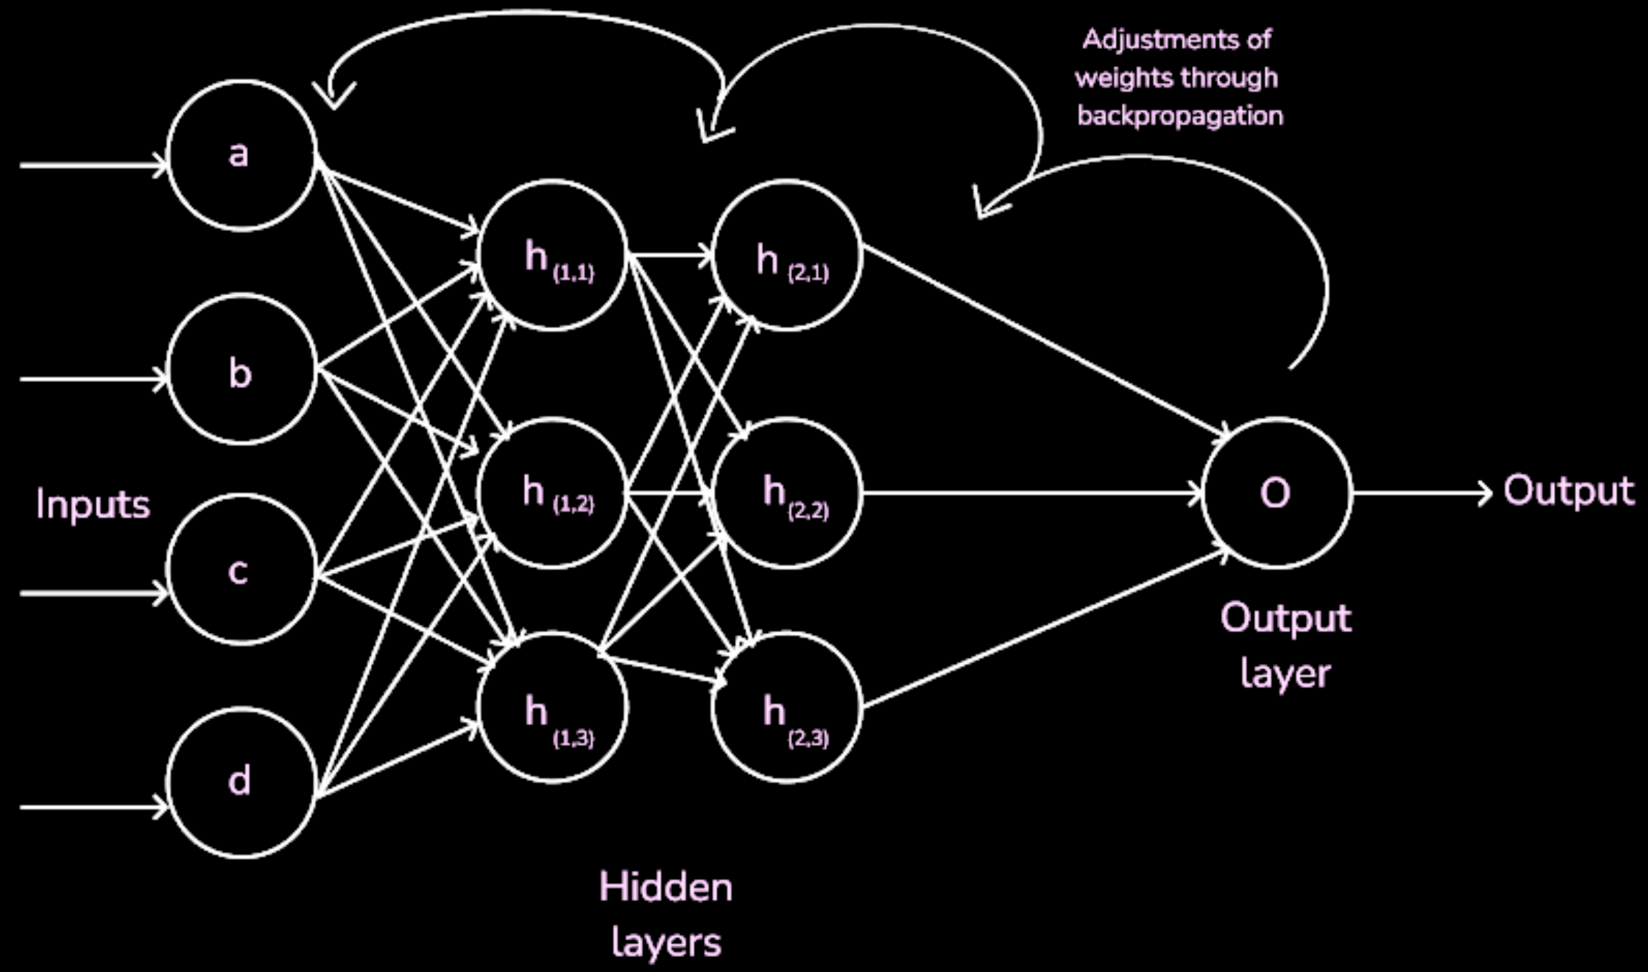
Adjustments of
weights through
backpropagation
a
b
c
d
Inputs
h (1,1)
h (1,2)
h (1,3)
h (2,1)
h (2,2)
h (2,3)
O
Output
Output
layer
Hidden
layers

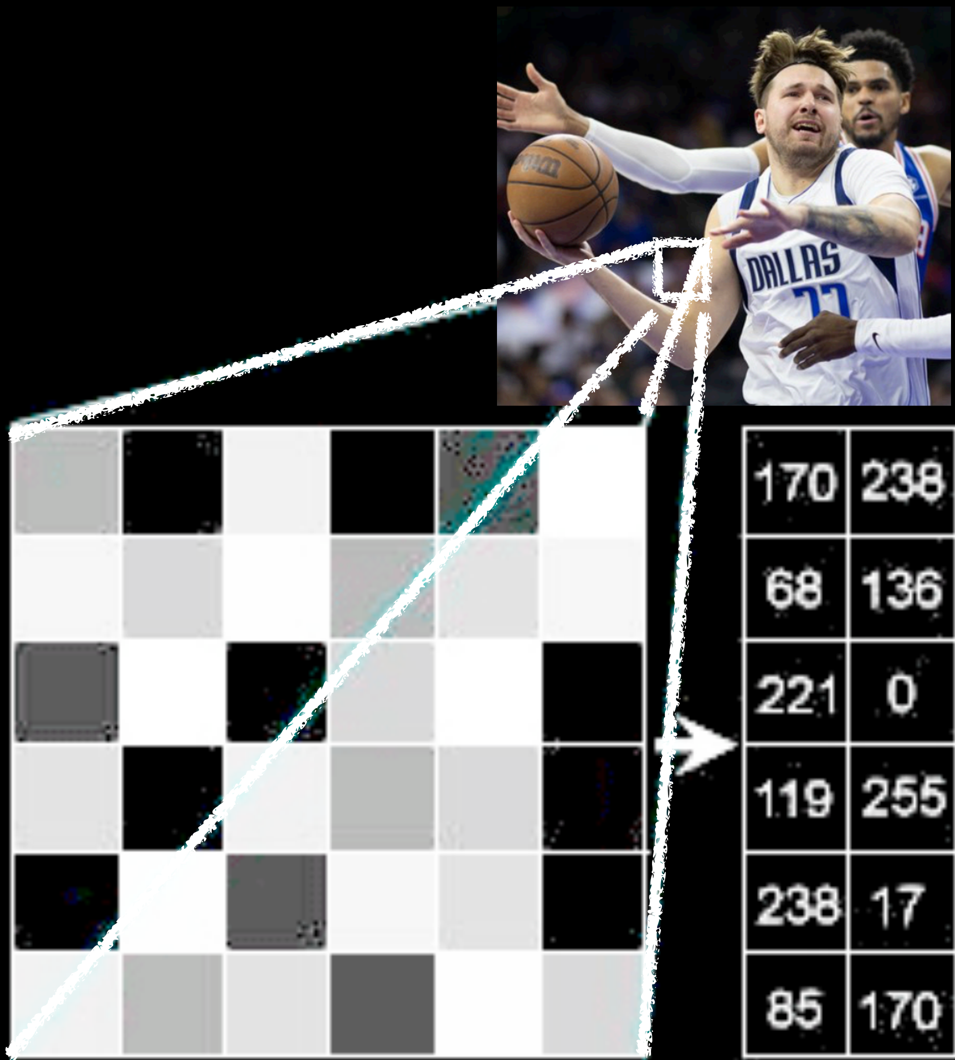

170	238	85	255	221	0
68	136	17	170	119	68
221	0	238	136	0	255
119	255	85	170	136	238
238	17	221	68	119	255
85	170	119	221	17	136

image pixels

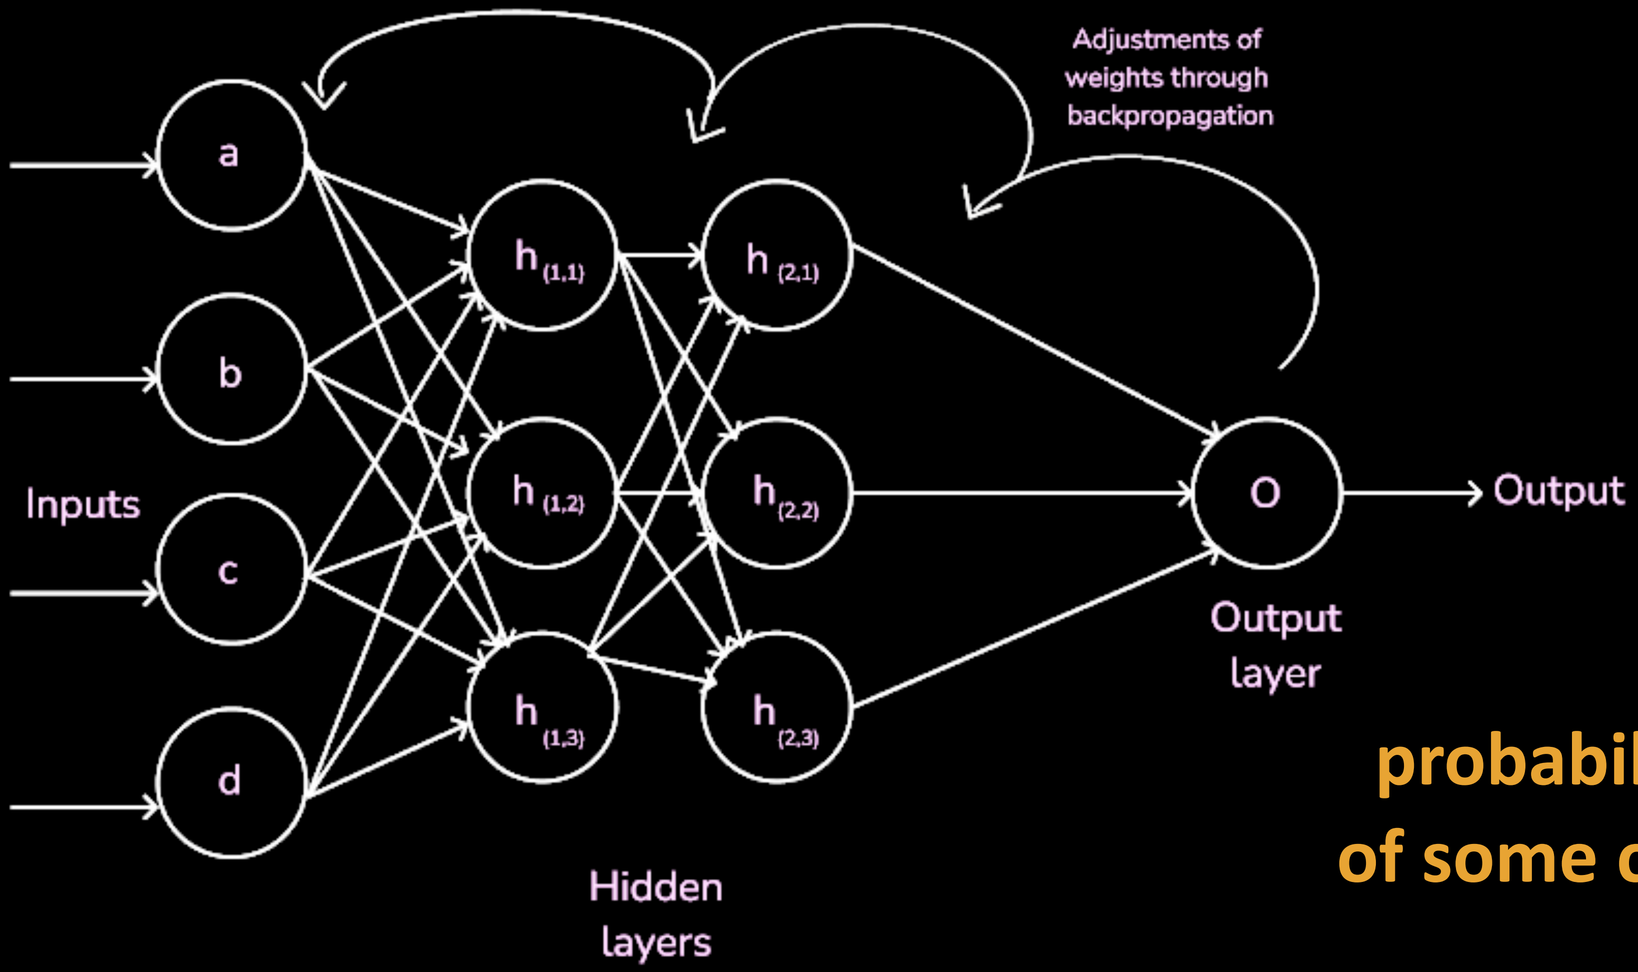

probability of some class

Jason Salavon
American, born 1970

Little Infinity (v.MFAH), 2020

Site-specific wallpaper

Museum commission funded by the Caroline Wiess Law Accessions Endowment Fund, 2020.141

Using software he designed, Jason Salavon has pulled nearly a quarter million images from the ImageNet dataset, a picture collection used by researchers and corporations to train artificial intelligence systems. By sampling each of the dataset's 20,000 categories and arranging those images from darkest to lightest within each category, Salavon presents a cascade of photographs suggesting the flood of imagery that inundates our daily lives. Exposing ImageNet's odd inclusions and notable absences in each category, *Little Infinity (v.MFAH)* poses essential questions about the nature of images and the ways in which they sway any understanding of the world.

AlexNet (Alex Krizhevsky, Ilya Sutskever, and Geoffrey Hinton, 2012)
60 million parameters

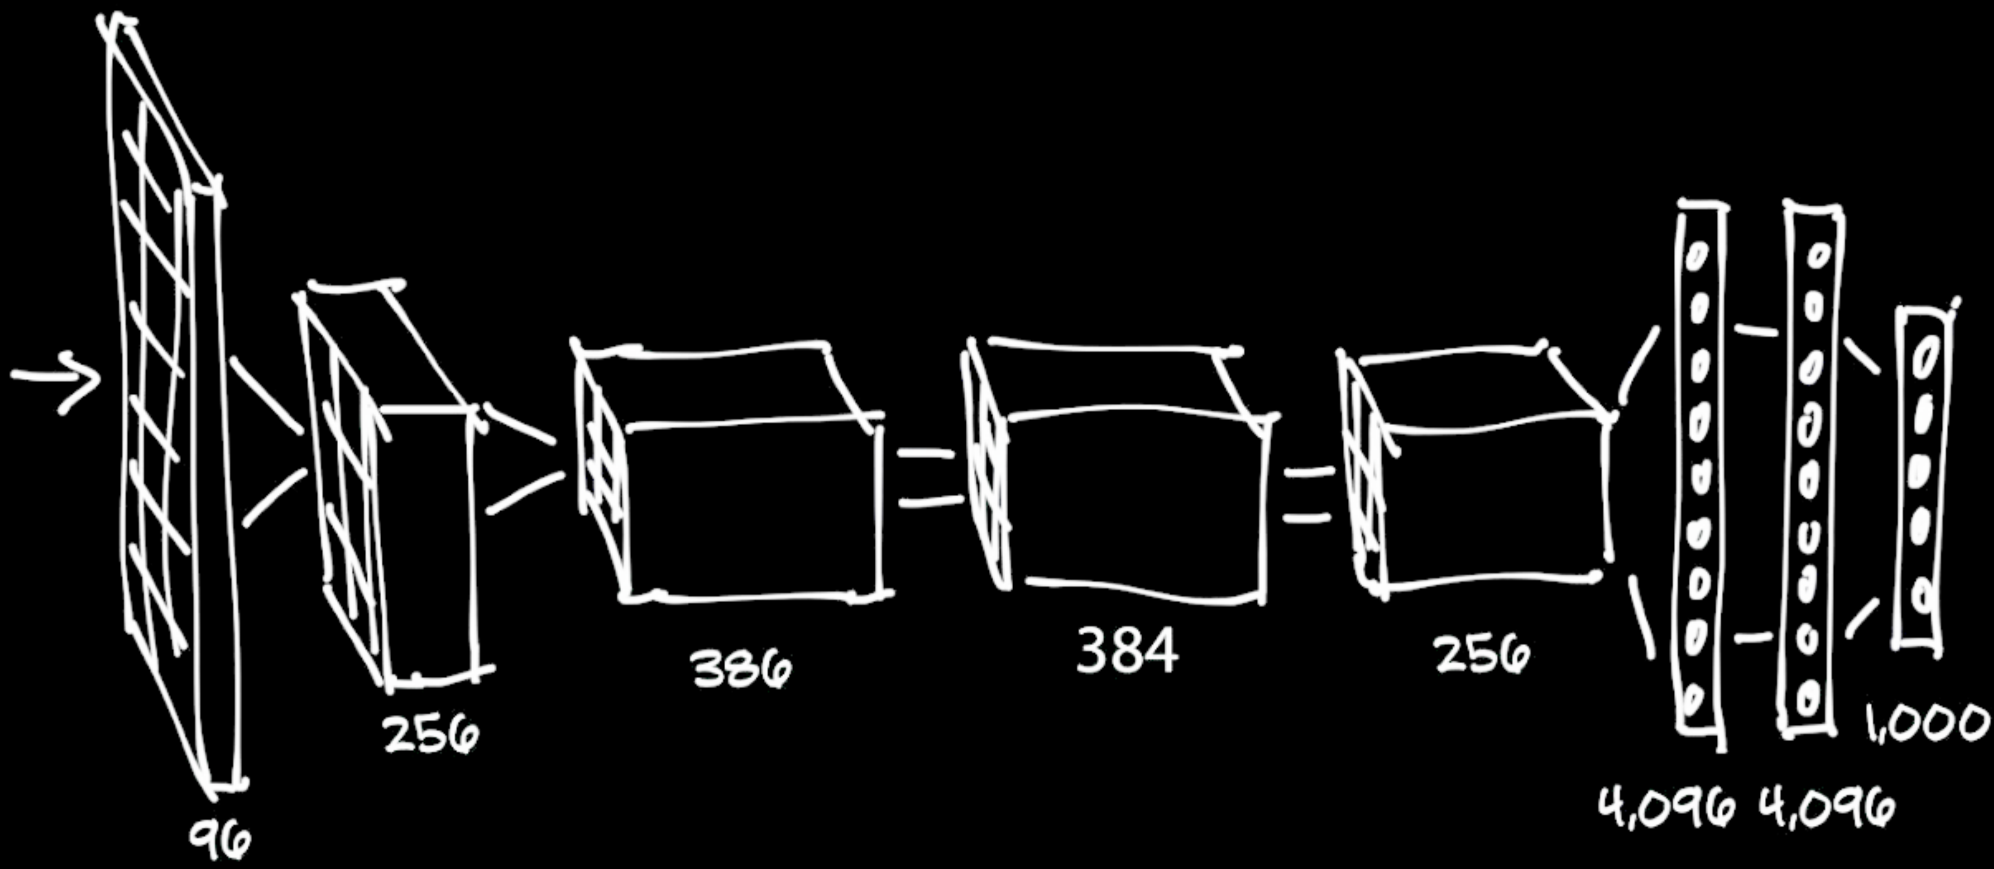

AlexNet (Alex Krizhevsky, Ilya Sutskever, and Geoffrey Hinton, 2012)
60 million parameters

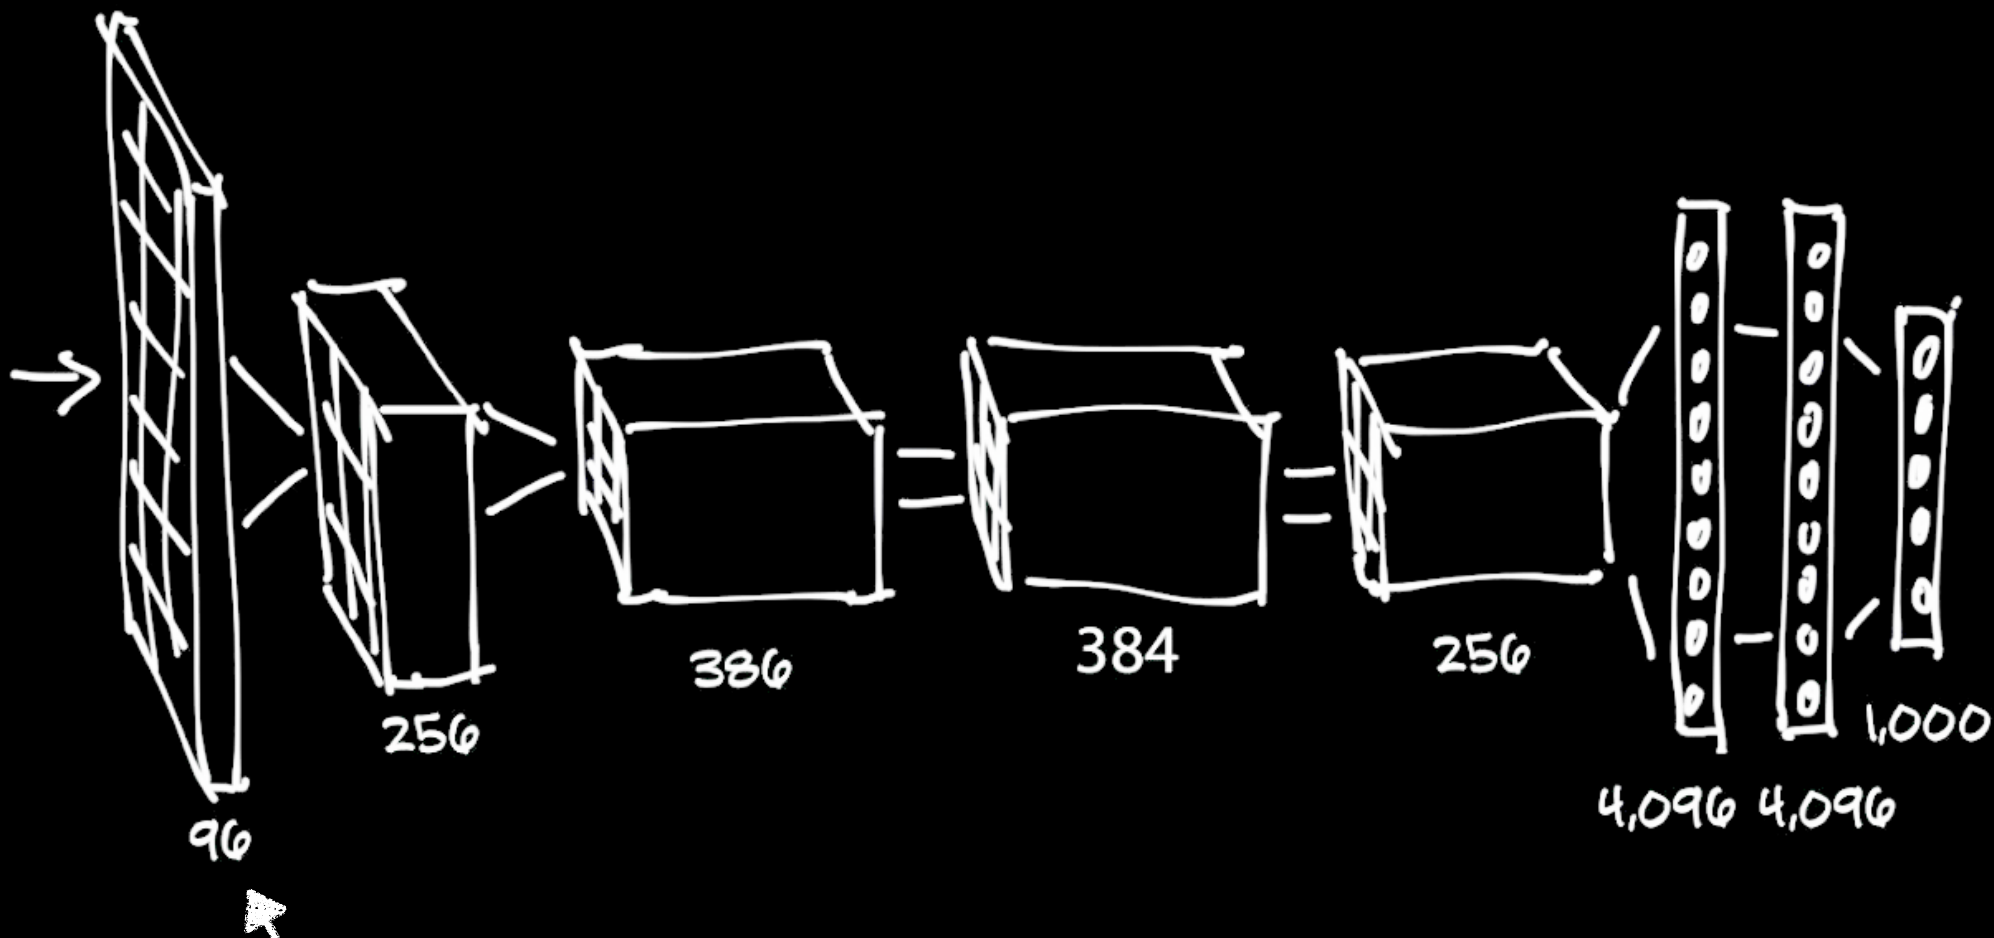

input image

AlexNet (Alex Krizhevsky, Ilya Sutskever, and Geoffrey Hinton, 2012)
60 million parameters

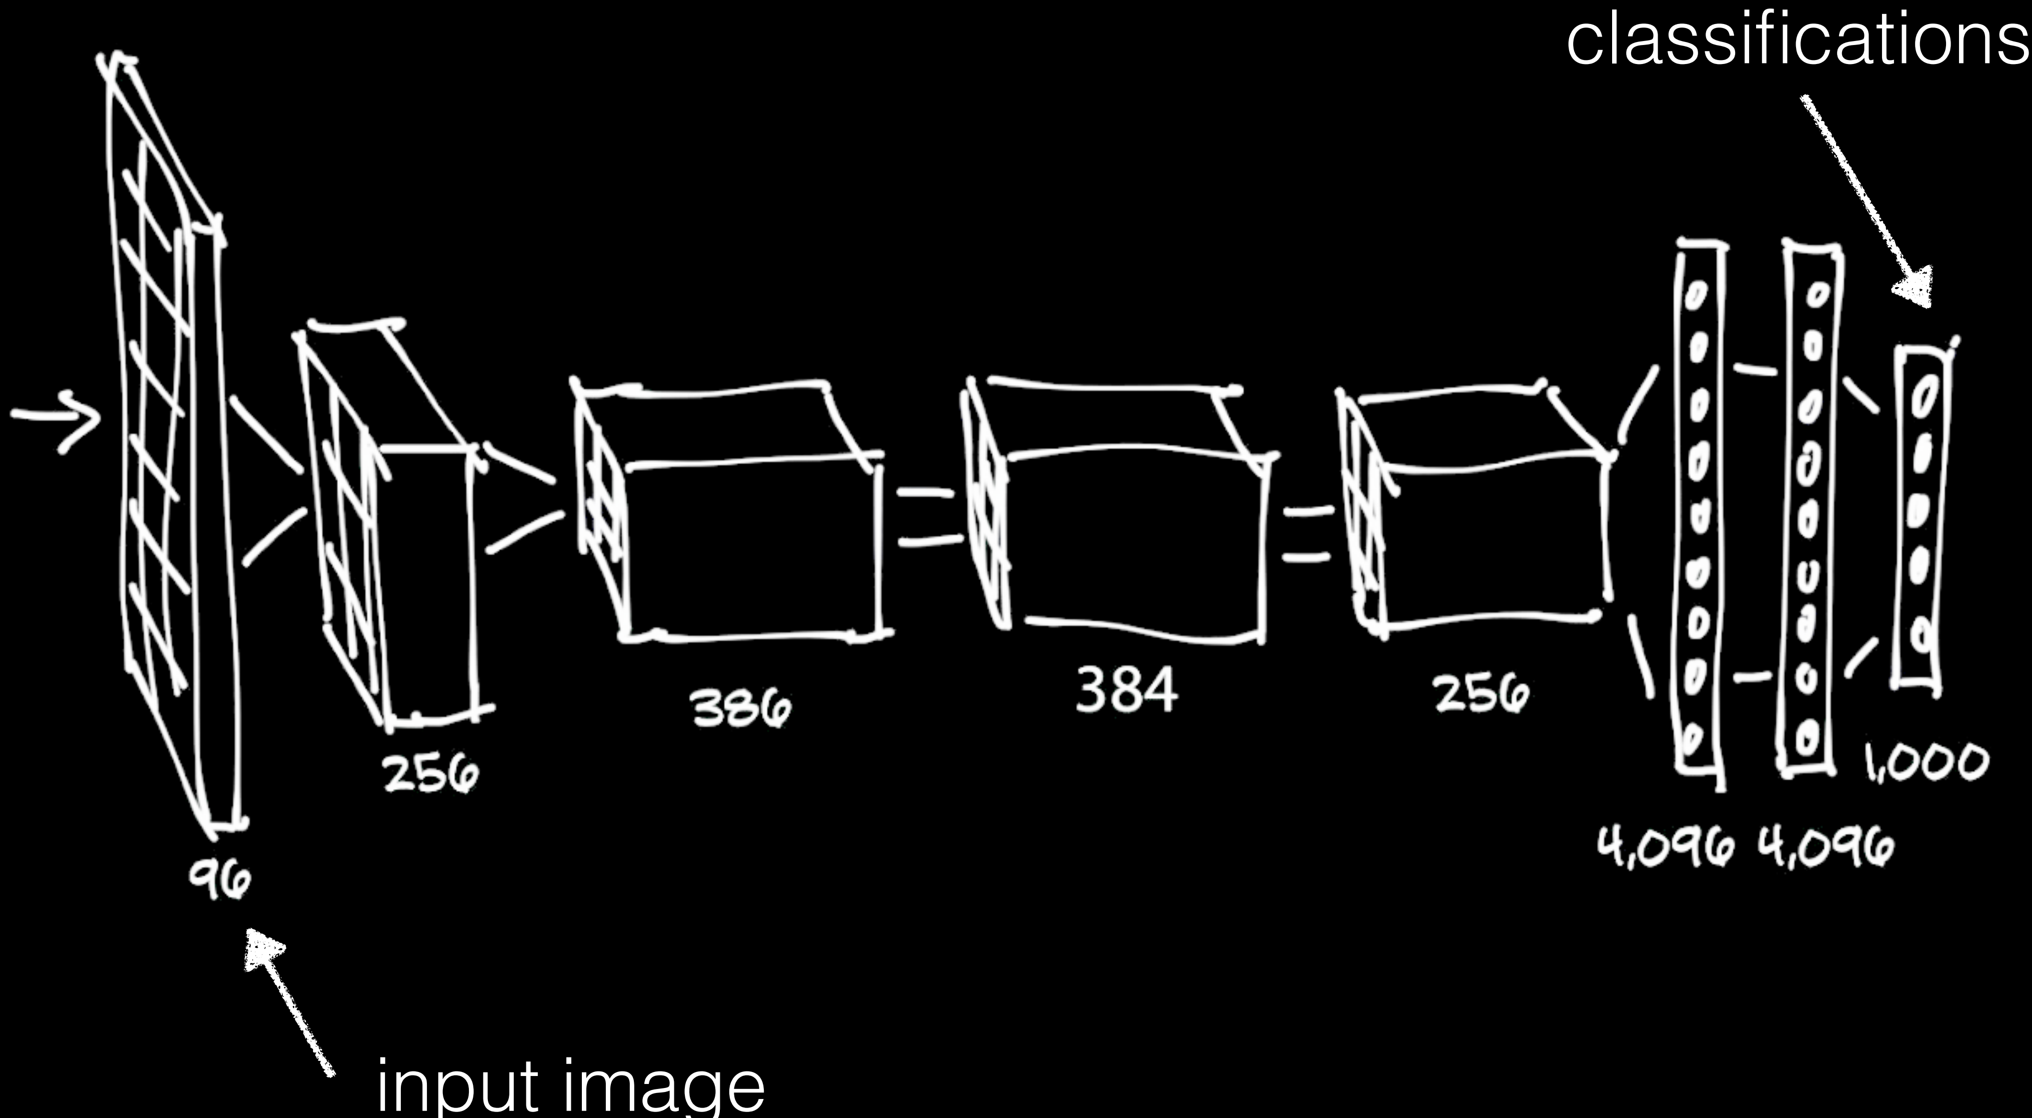

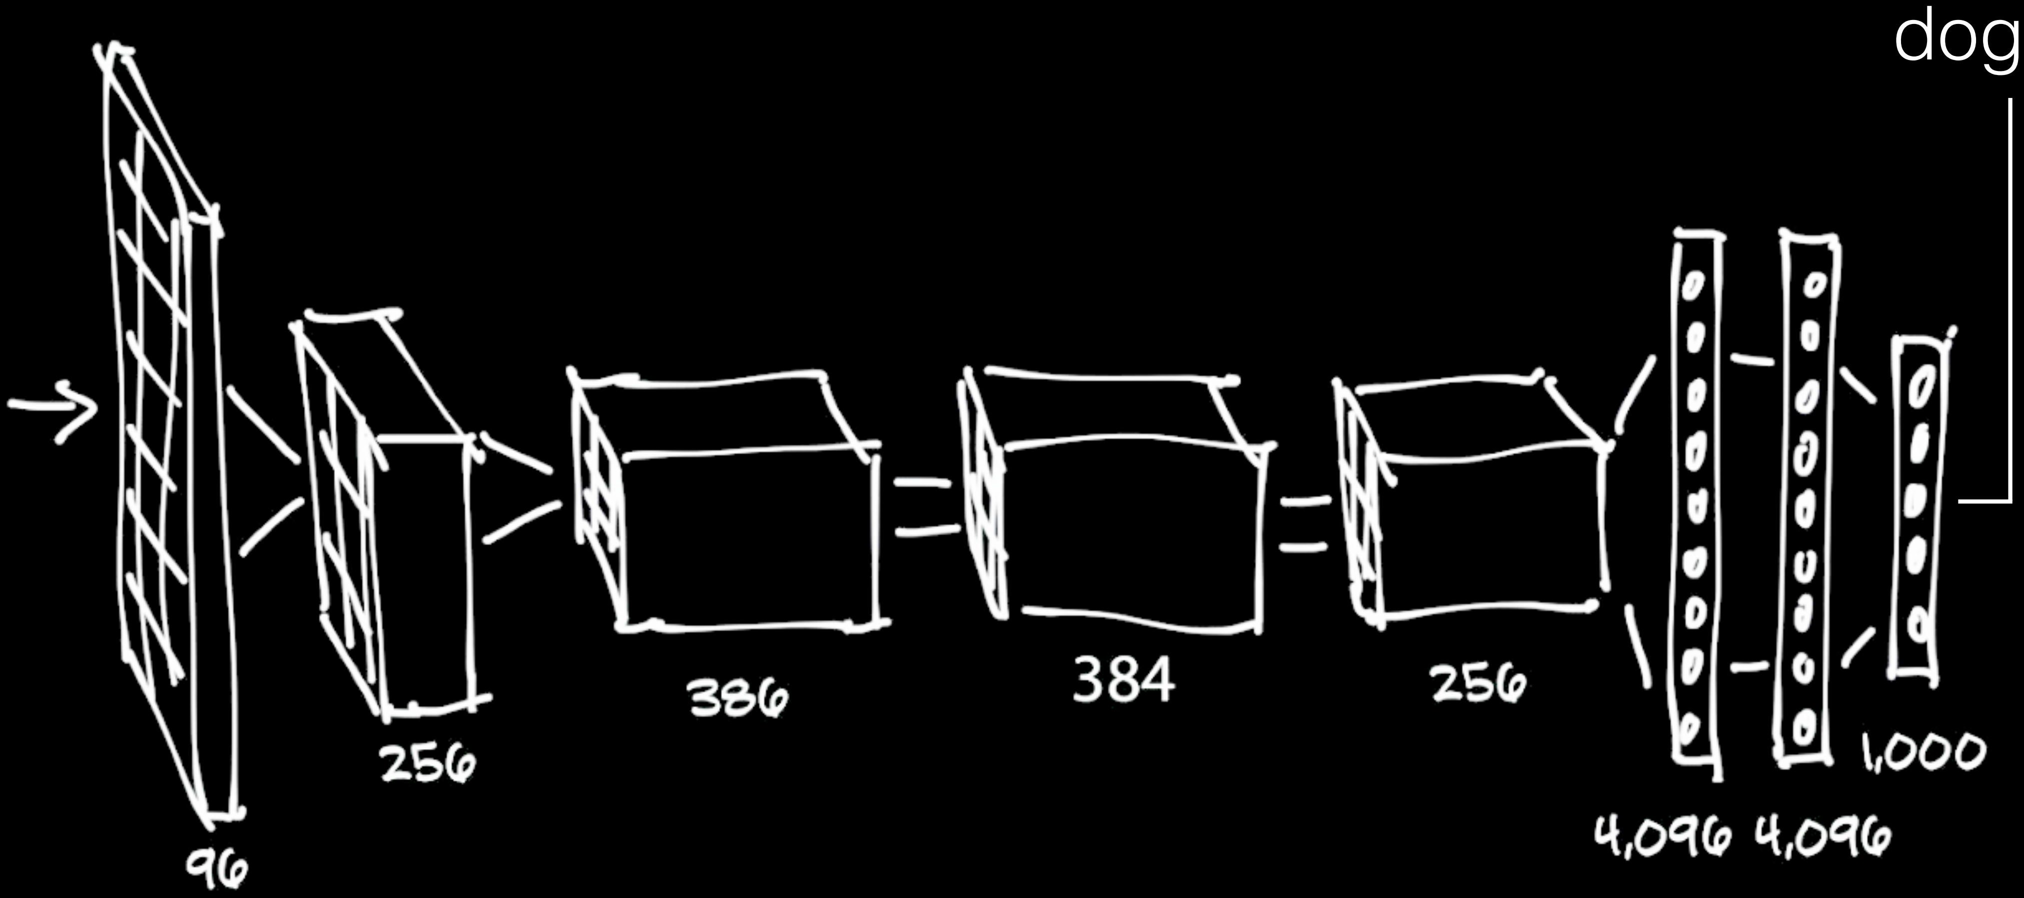
dog
96
256
386
384
256
4,096
4,096
1,000

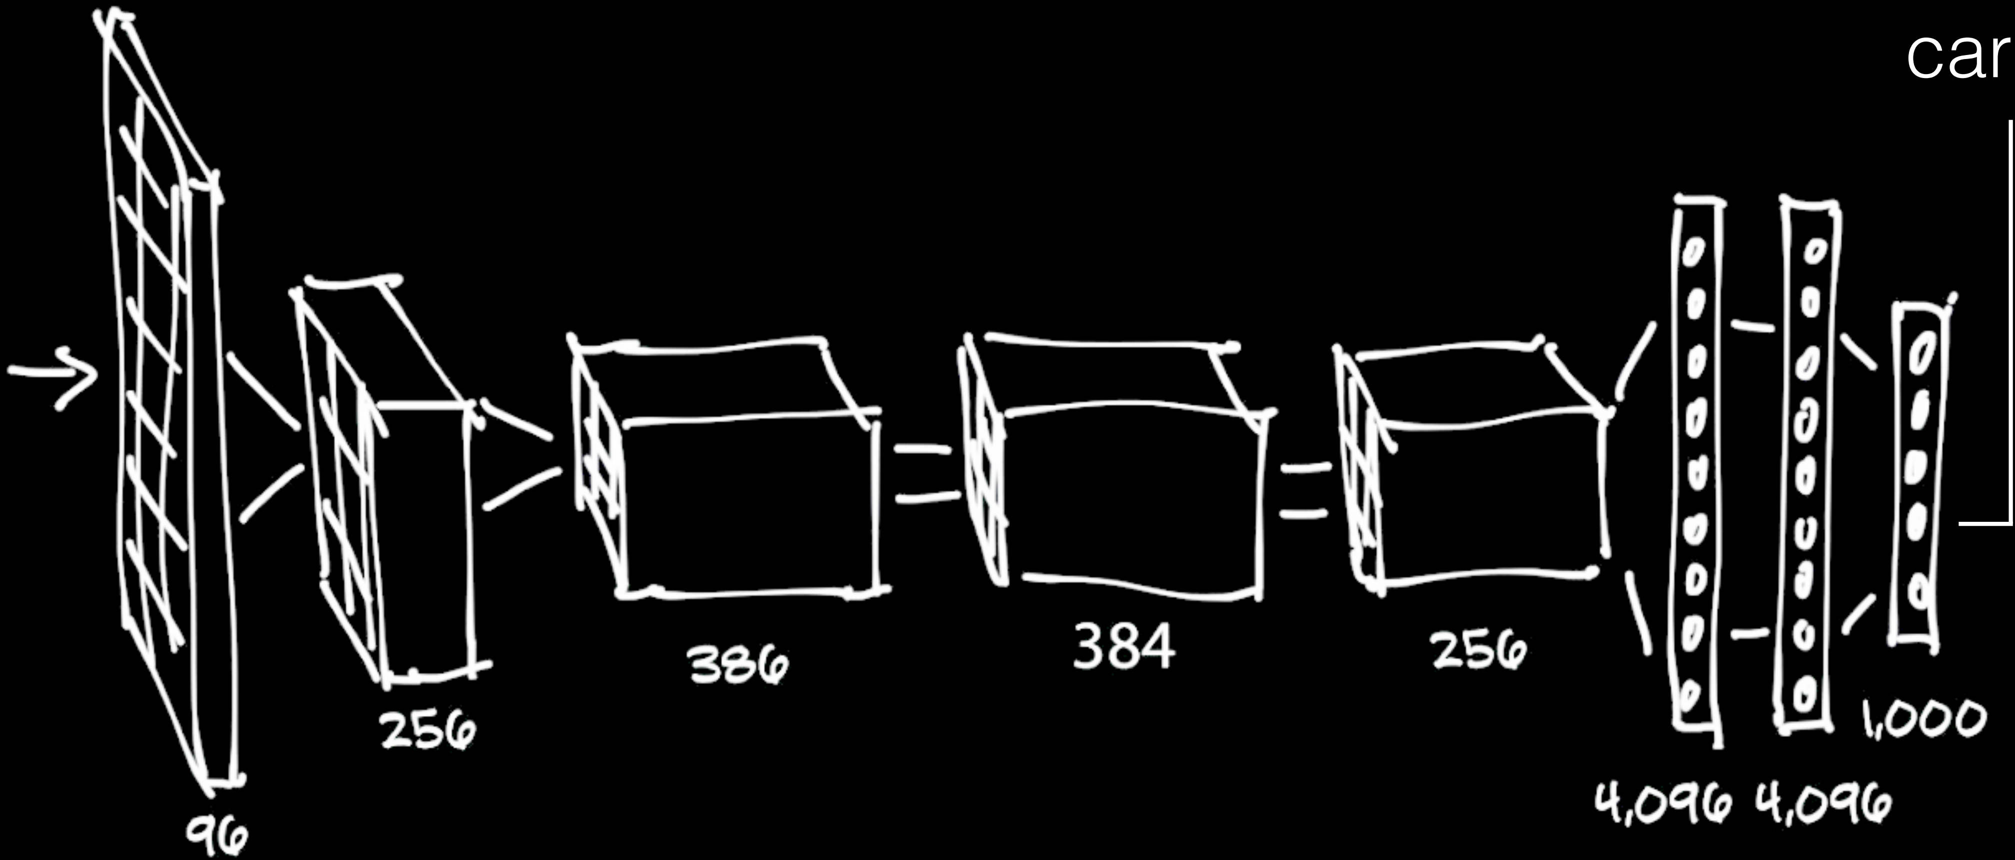
car
96
256
386
384
256
4,096
4,096
1,000

AlexNet (Alex Krizhevsky, Ilya Sutskever, in Geoffrey Hinton, 2012)
60 million parameters

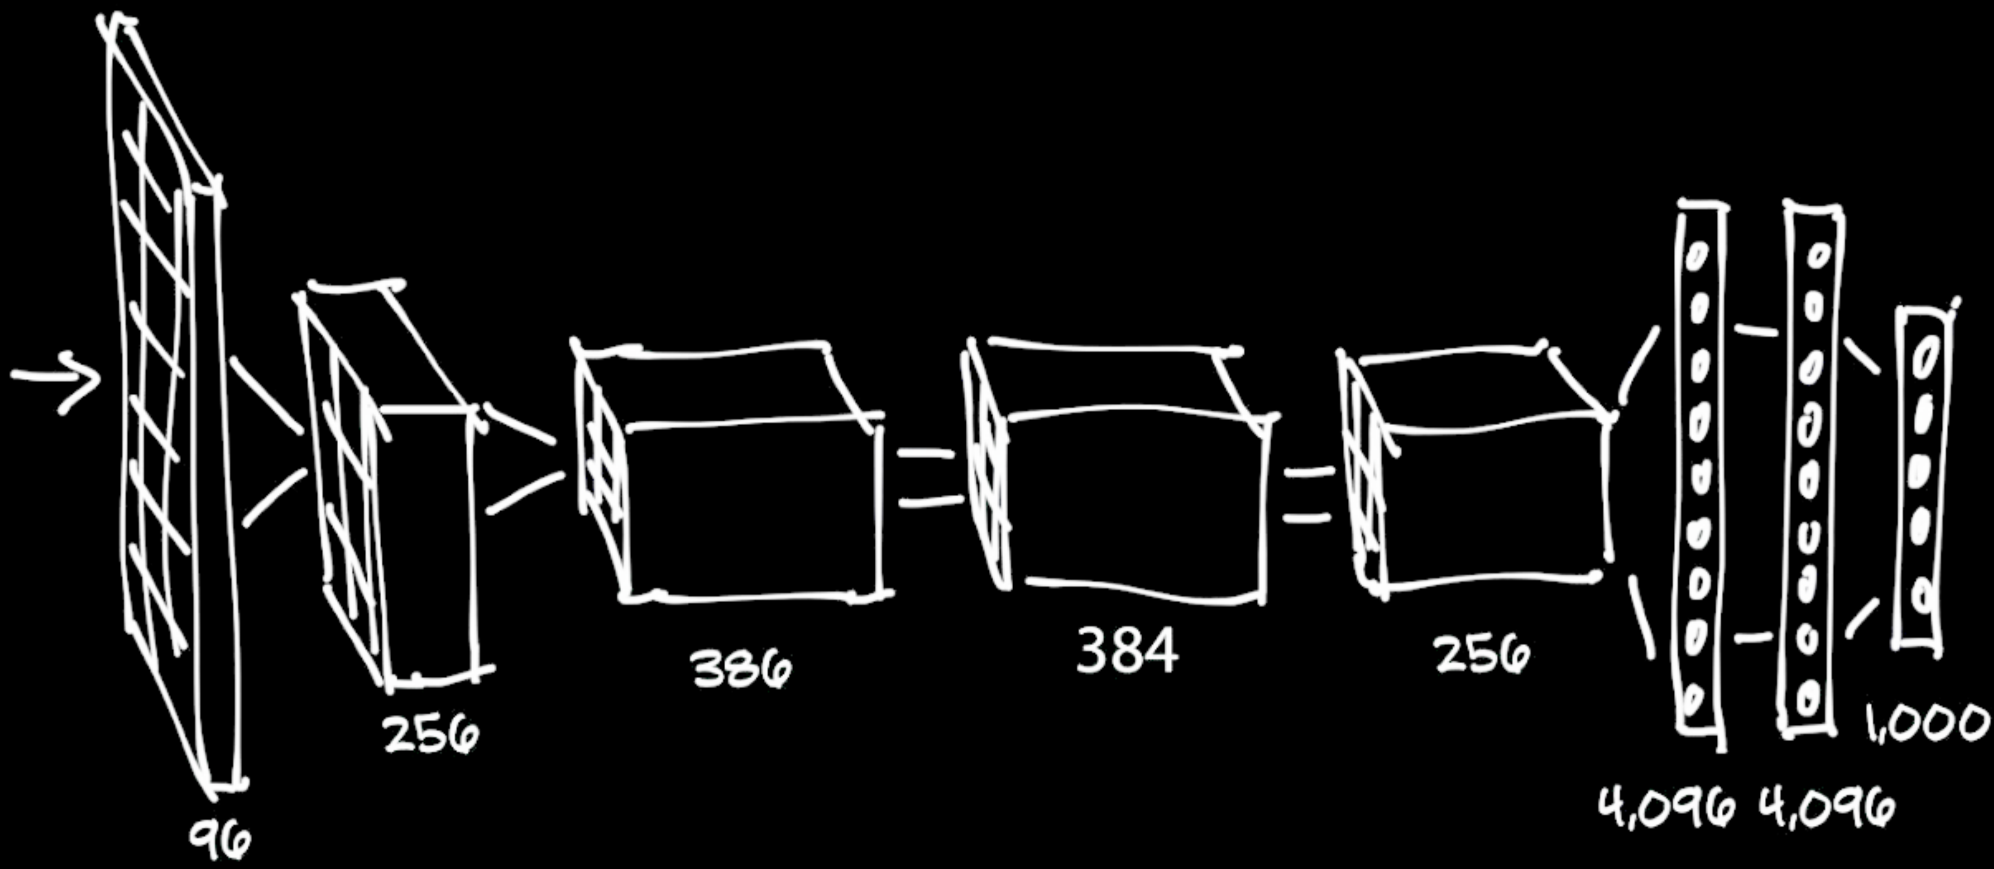

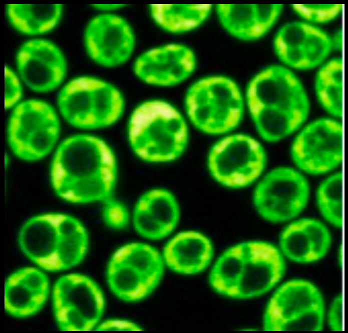

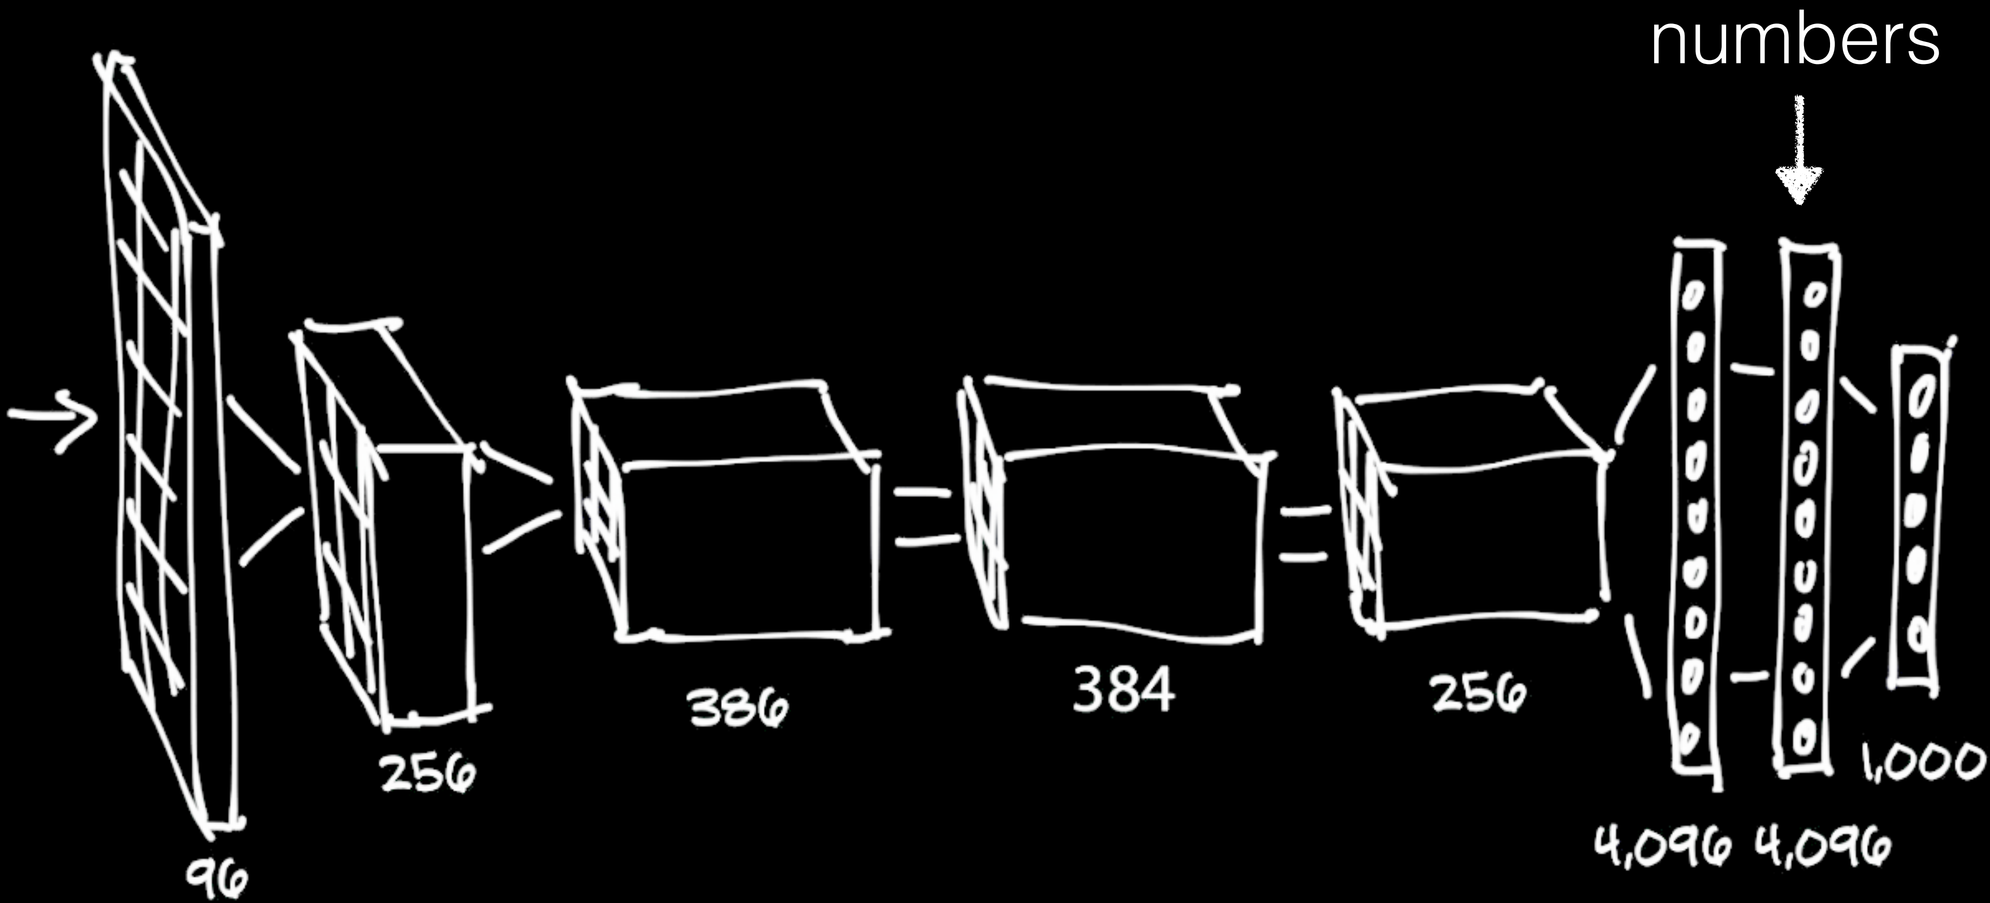
numbers
96
256
386
384
256
4,096
4,096
1,000

Godec et al. (2019) Democratized image analytics by visual programming through integration of deep models and small-scale machine learning, *Nature Communications* 10:4551.

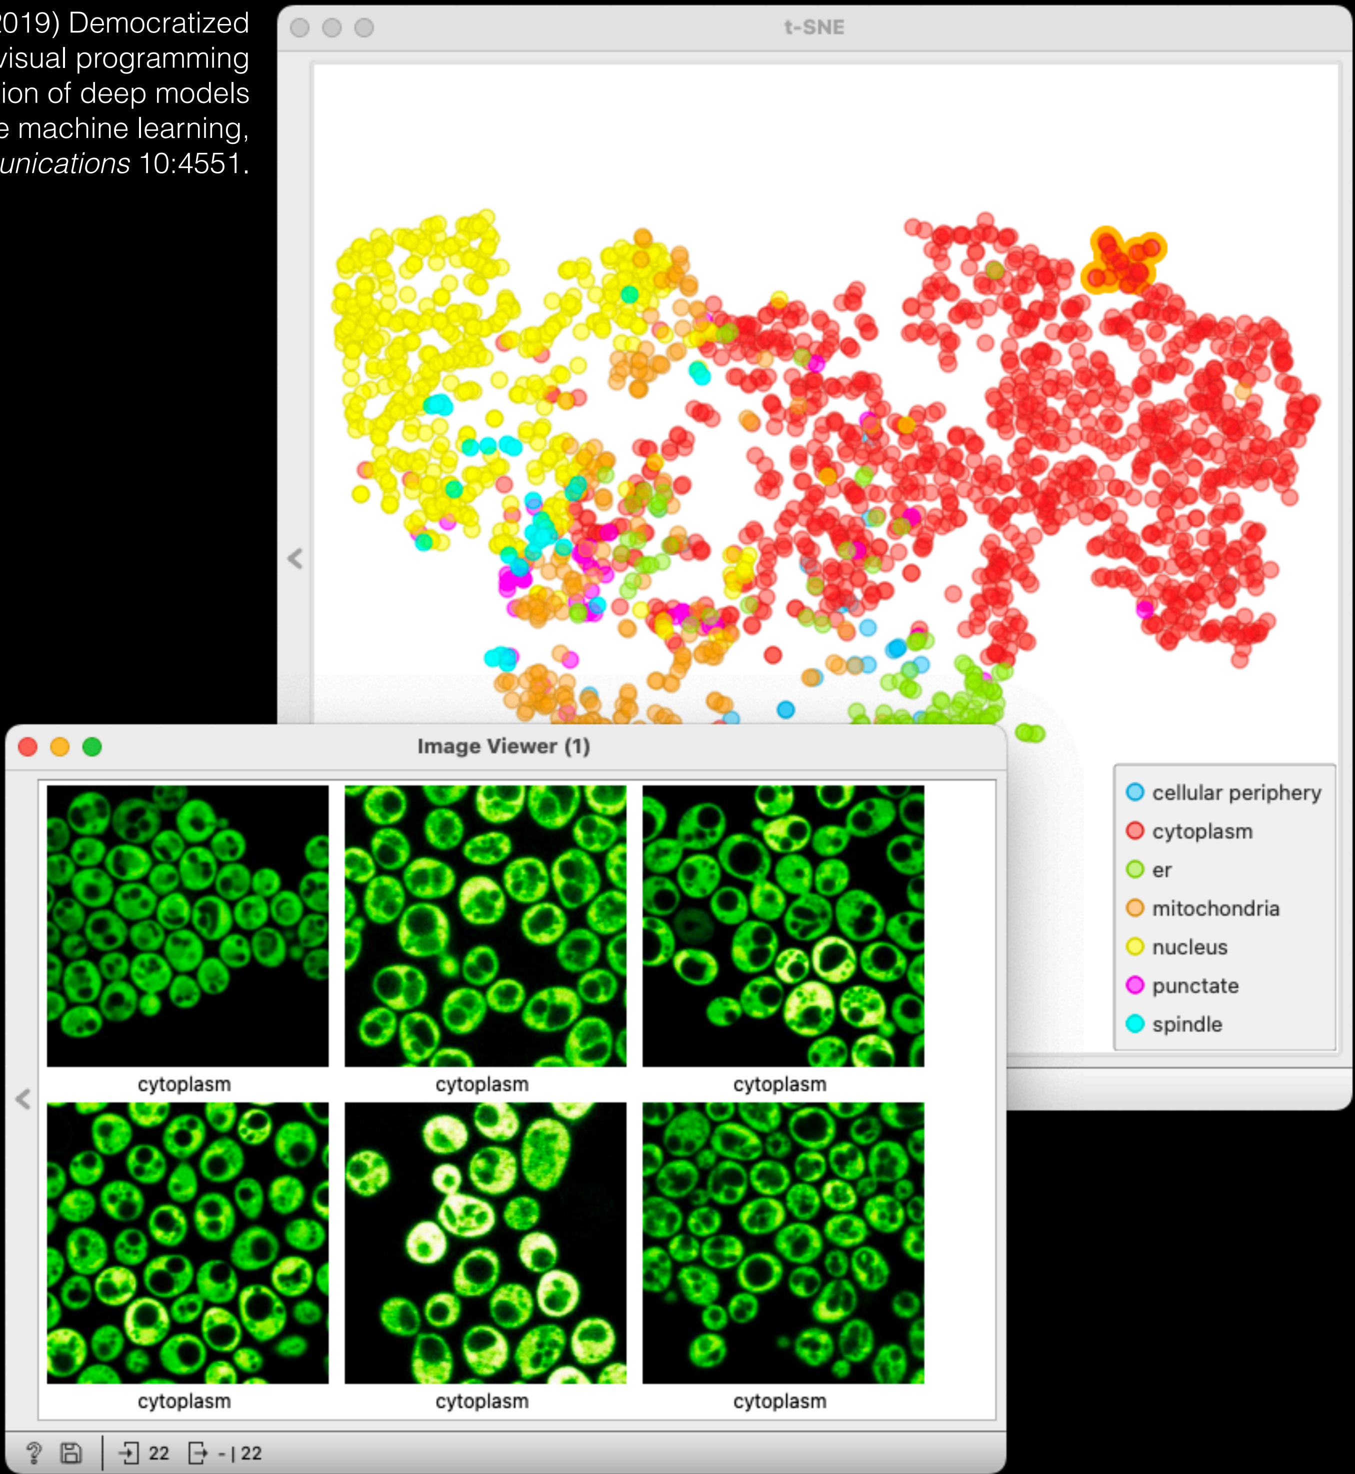

Images are simple.
How about words?

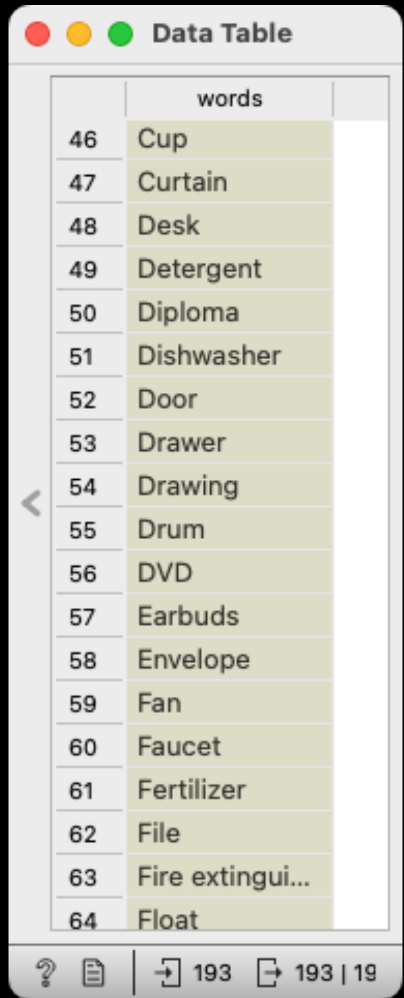

words - n-grams - neural network - classification

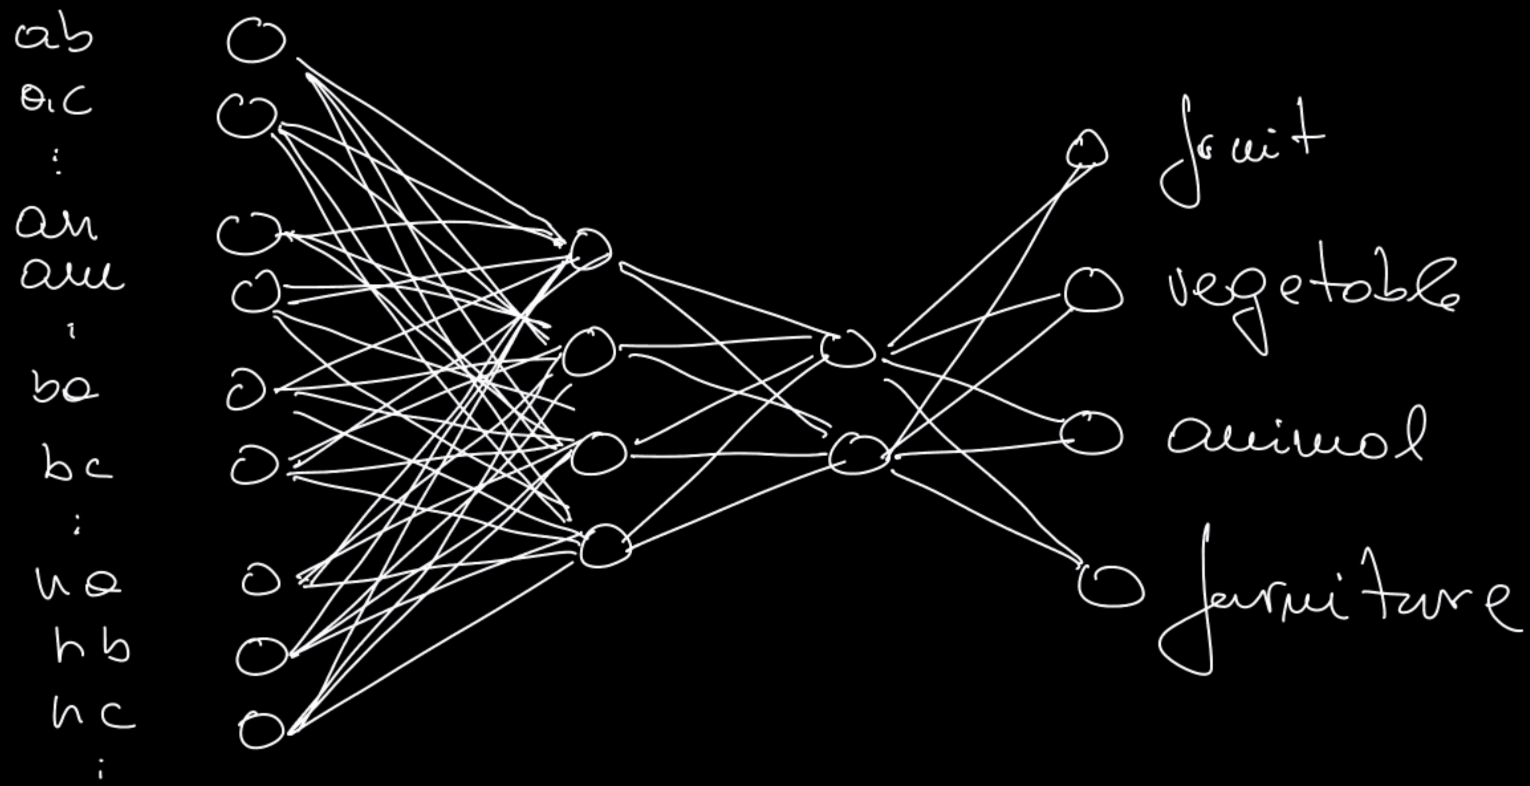

words - n-grams - neural network - classification

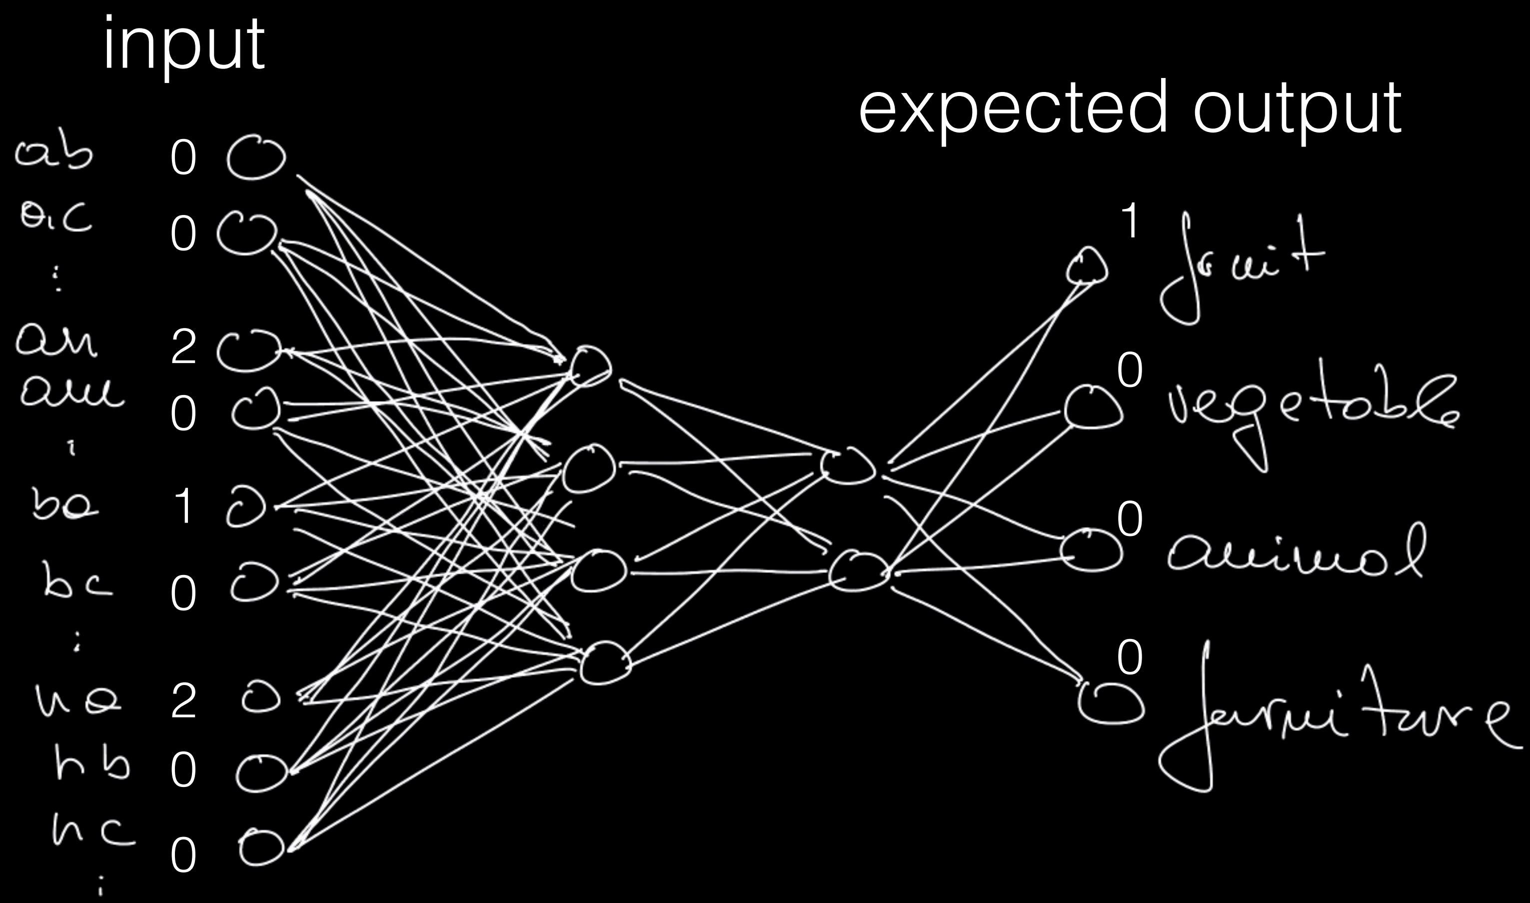

banana – ba, an, na, an, na

encoding of
input word

prediction
of next word

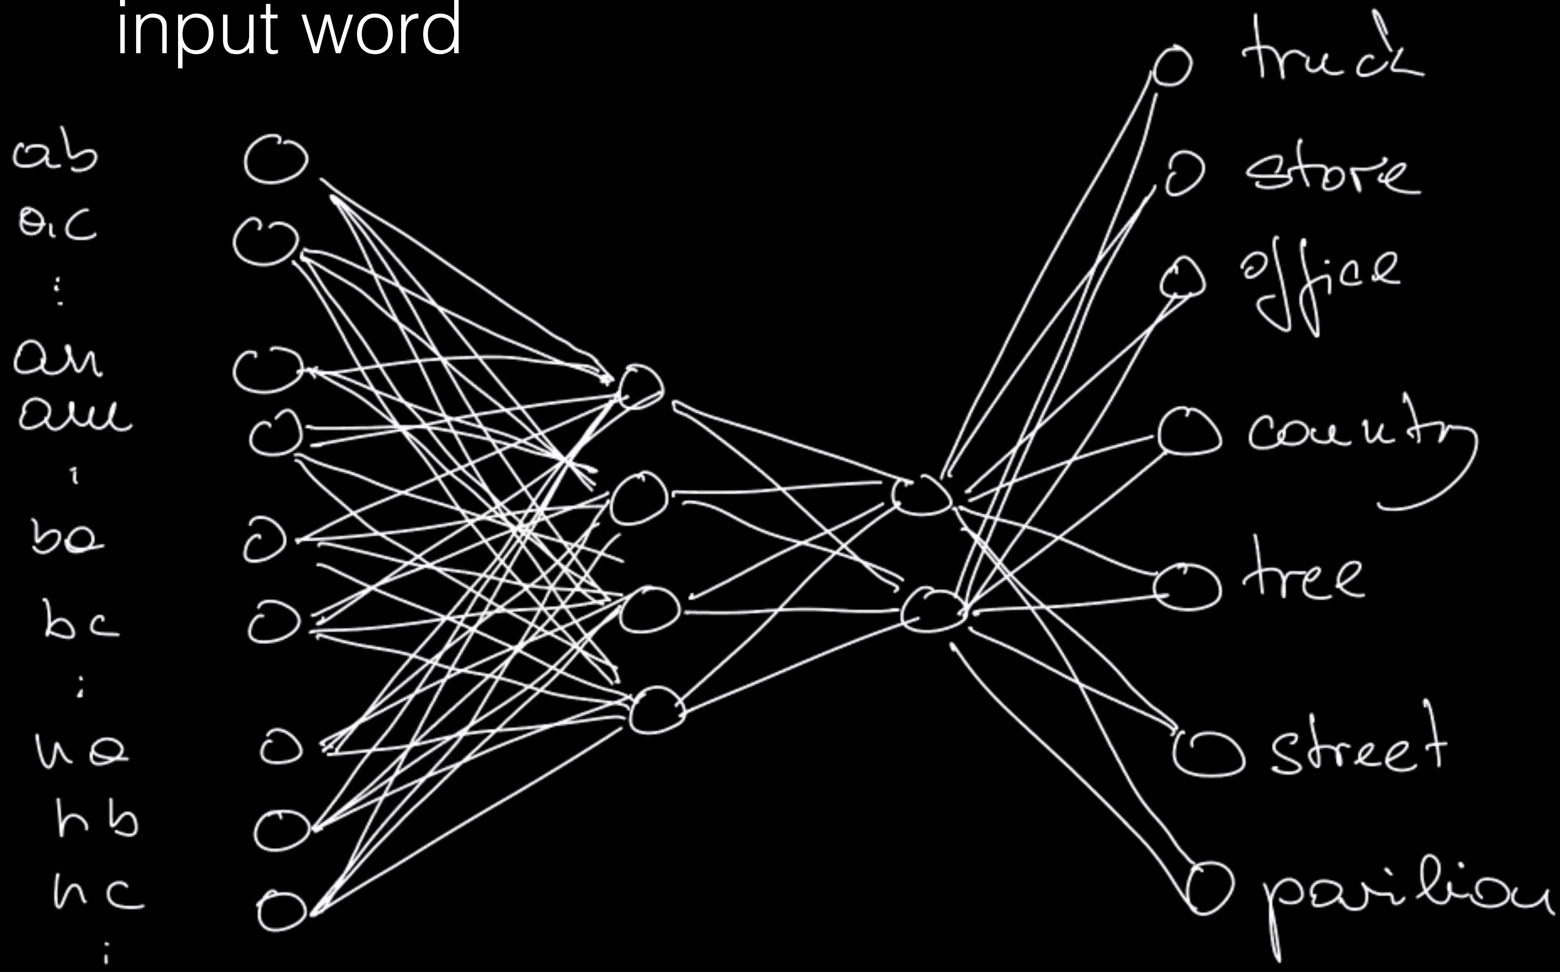

encoding of
input string

prediction
of next token

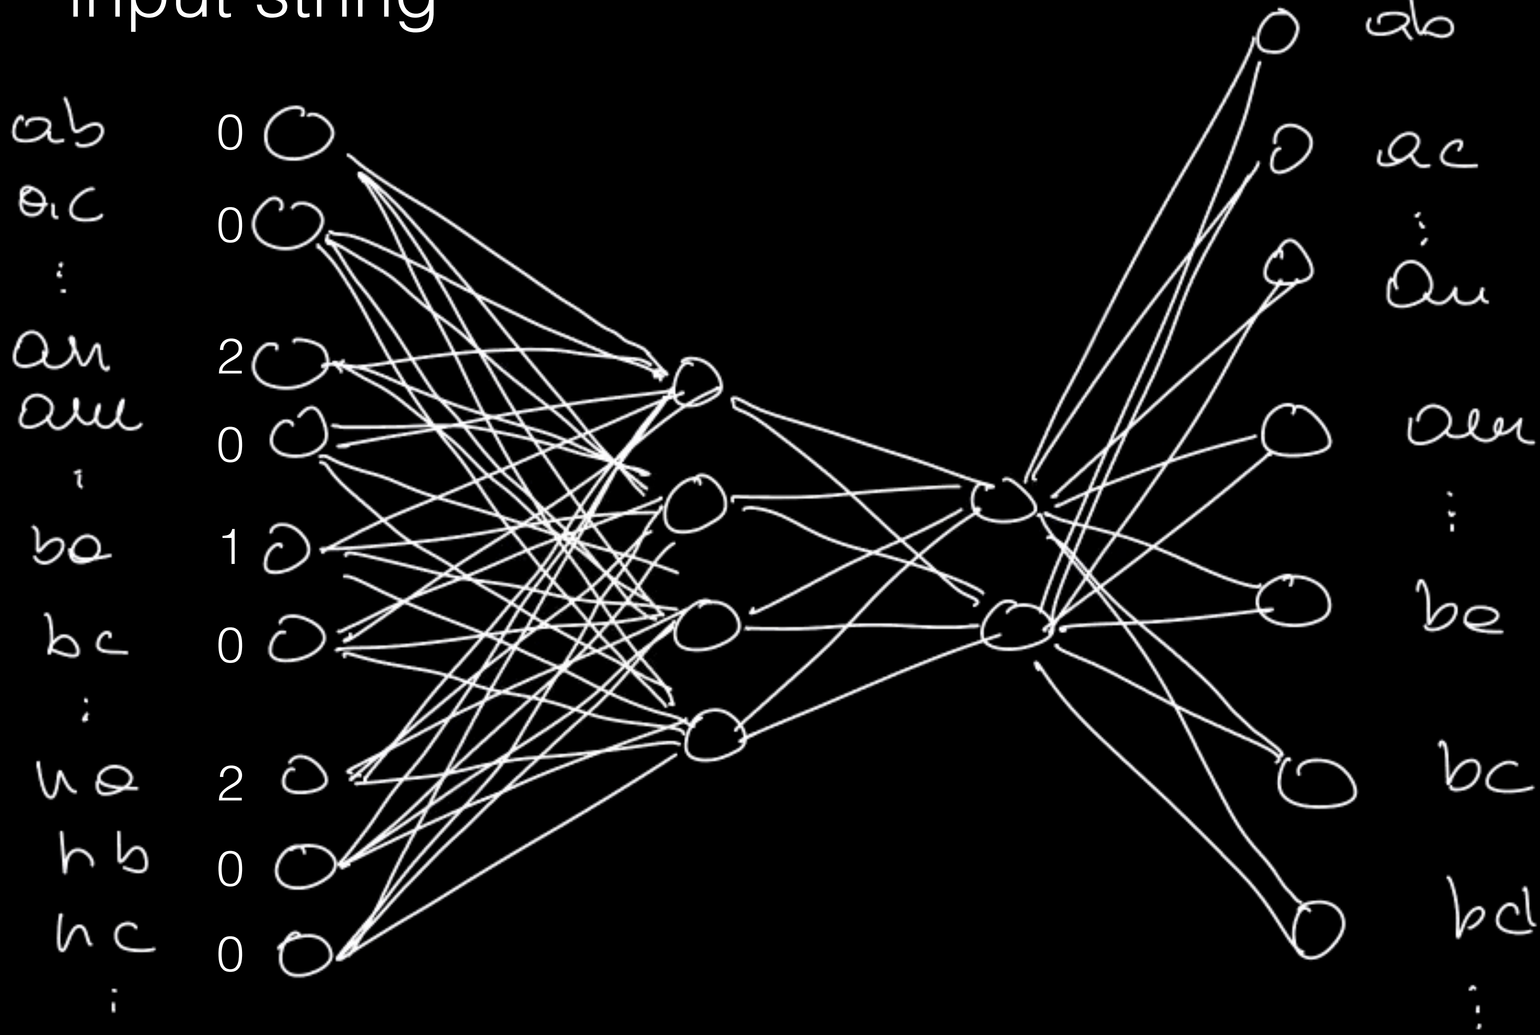

banana – ba, an, na, an, na

she peeled the **banana** quickly

monkey grabbed **banana** from the visitor

this **banana** is too green

he dropped his **banana** slice to the floor

fresh **banana** smoothie with berries

encoding of
input string

prediction
of next token

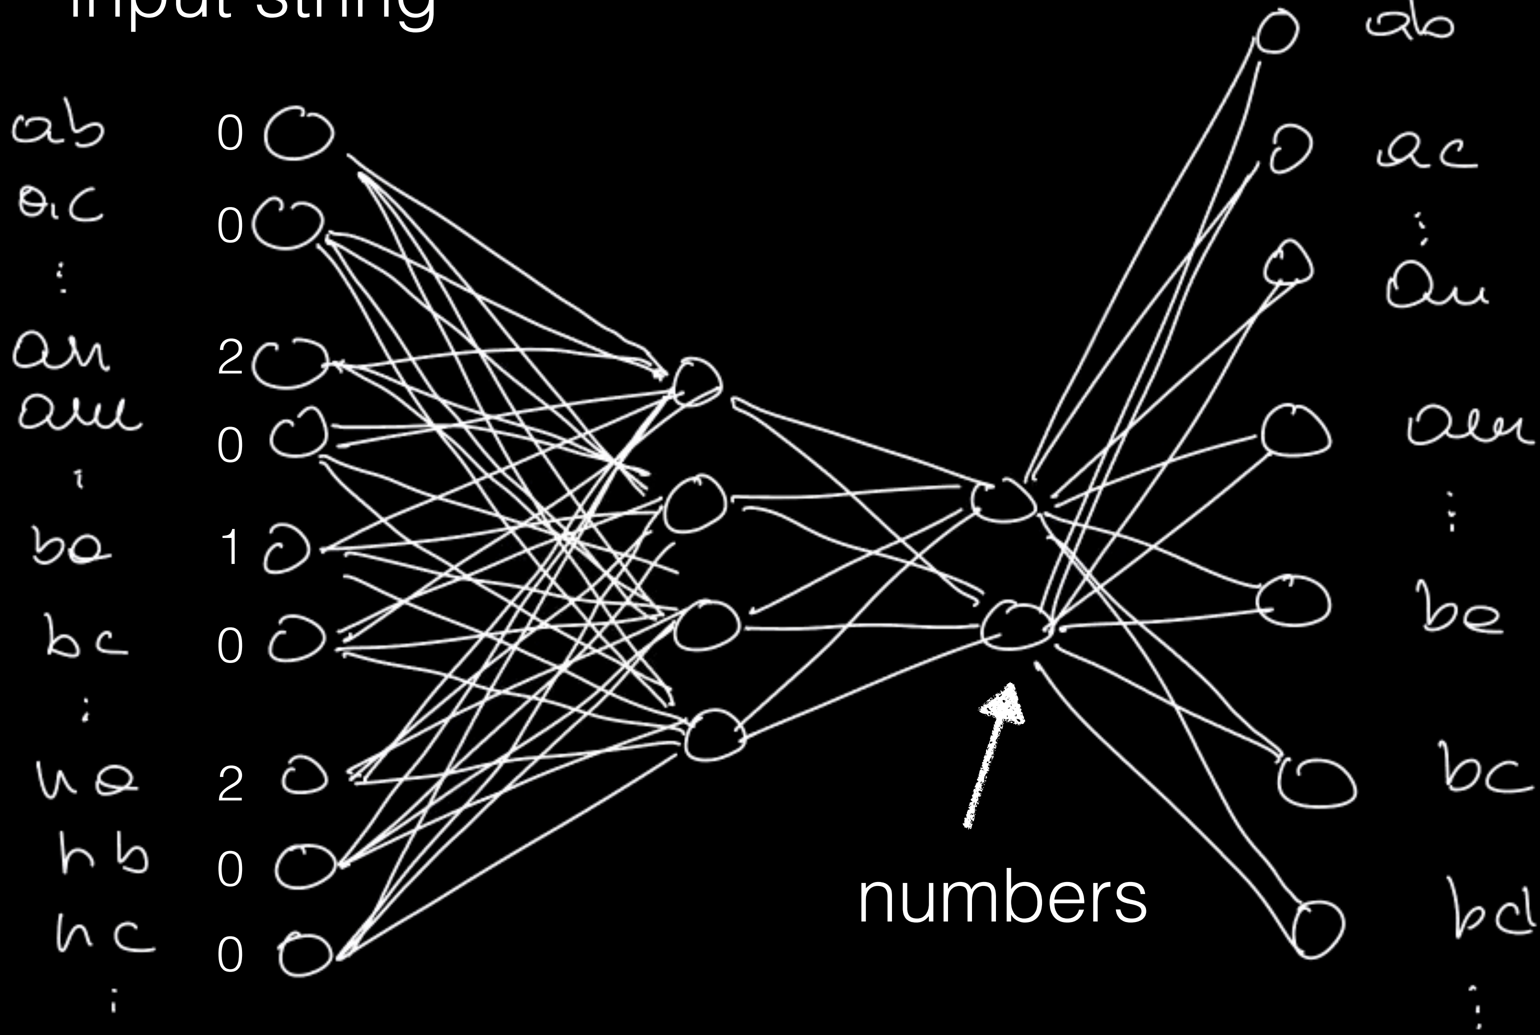

banana – ba, an, na, an, na

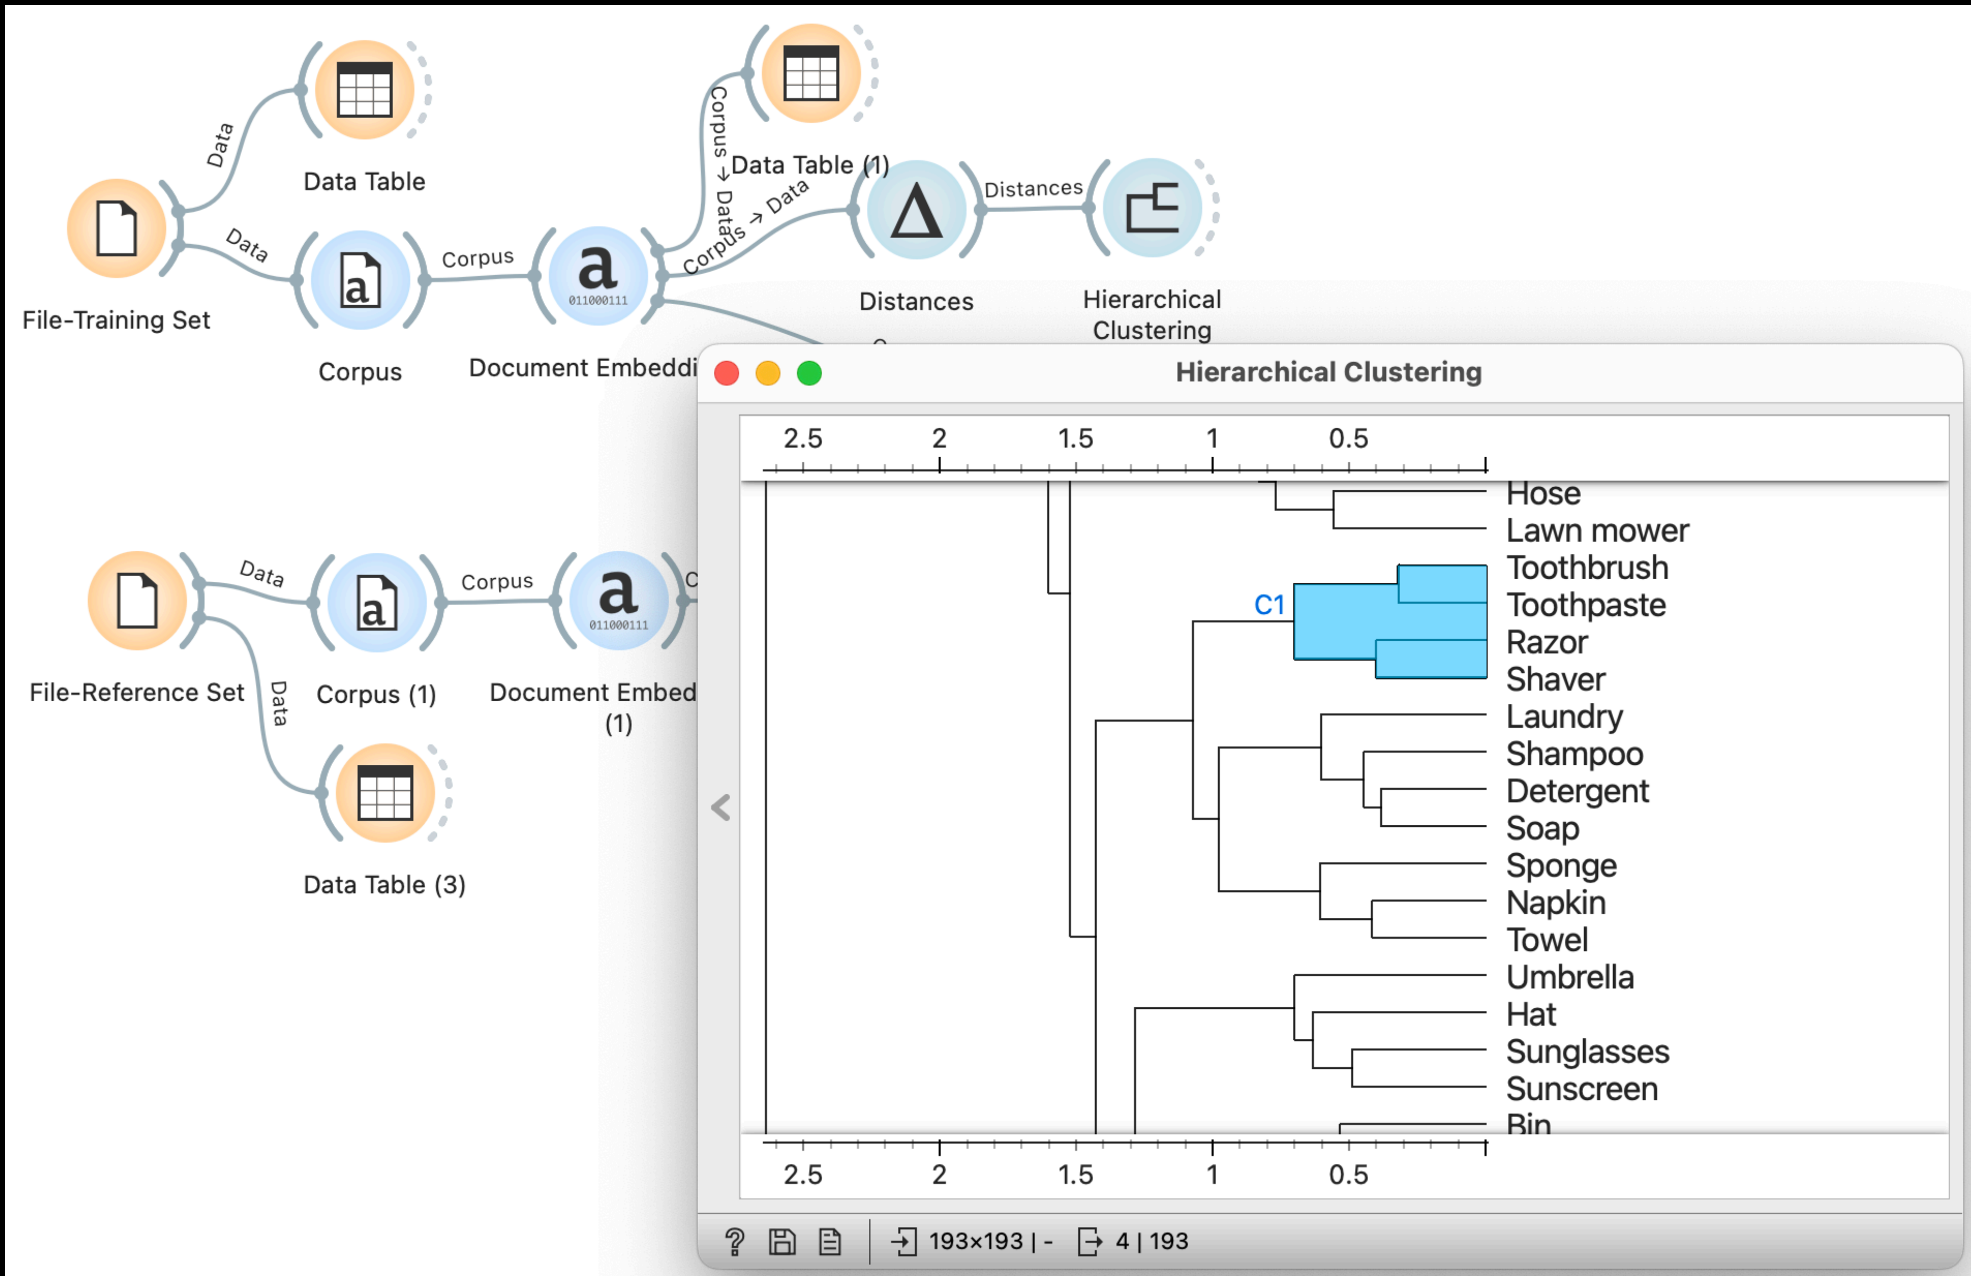
Data Table
Data
Data
File-Training Set
Corpus
Corpus
Document Embeddi
Data Table (1)
Corpus → Data
Corpus → Data
Distances
Distances
Hierarchical Clustering
Data
File-Reference Set
Data
Corpus (1)
Corpus
Document Embed (1)
Data
Data Table (3)
Hierarchical Clustering
2.5
2
1.5
1
0.5
Hose
Lawn mower
Toothbrush
Toothpaste
Razor
Shaver
C1
Laundry
Shampoo
Detergent
Soap
Sponge
Napkin
Towel
Umbrella
Hat
Sunglasses
Sunscreen
Bin
193×193 | -
4 | 193

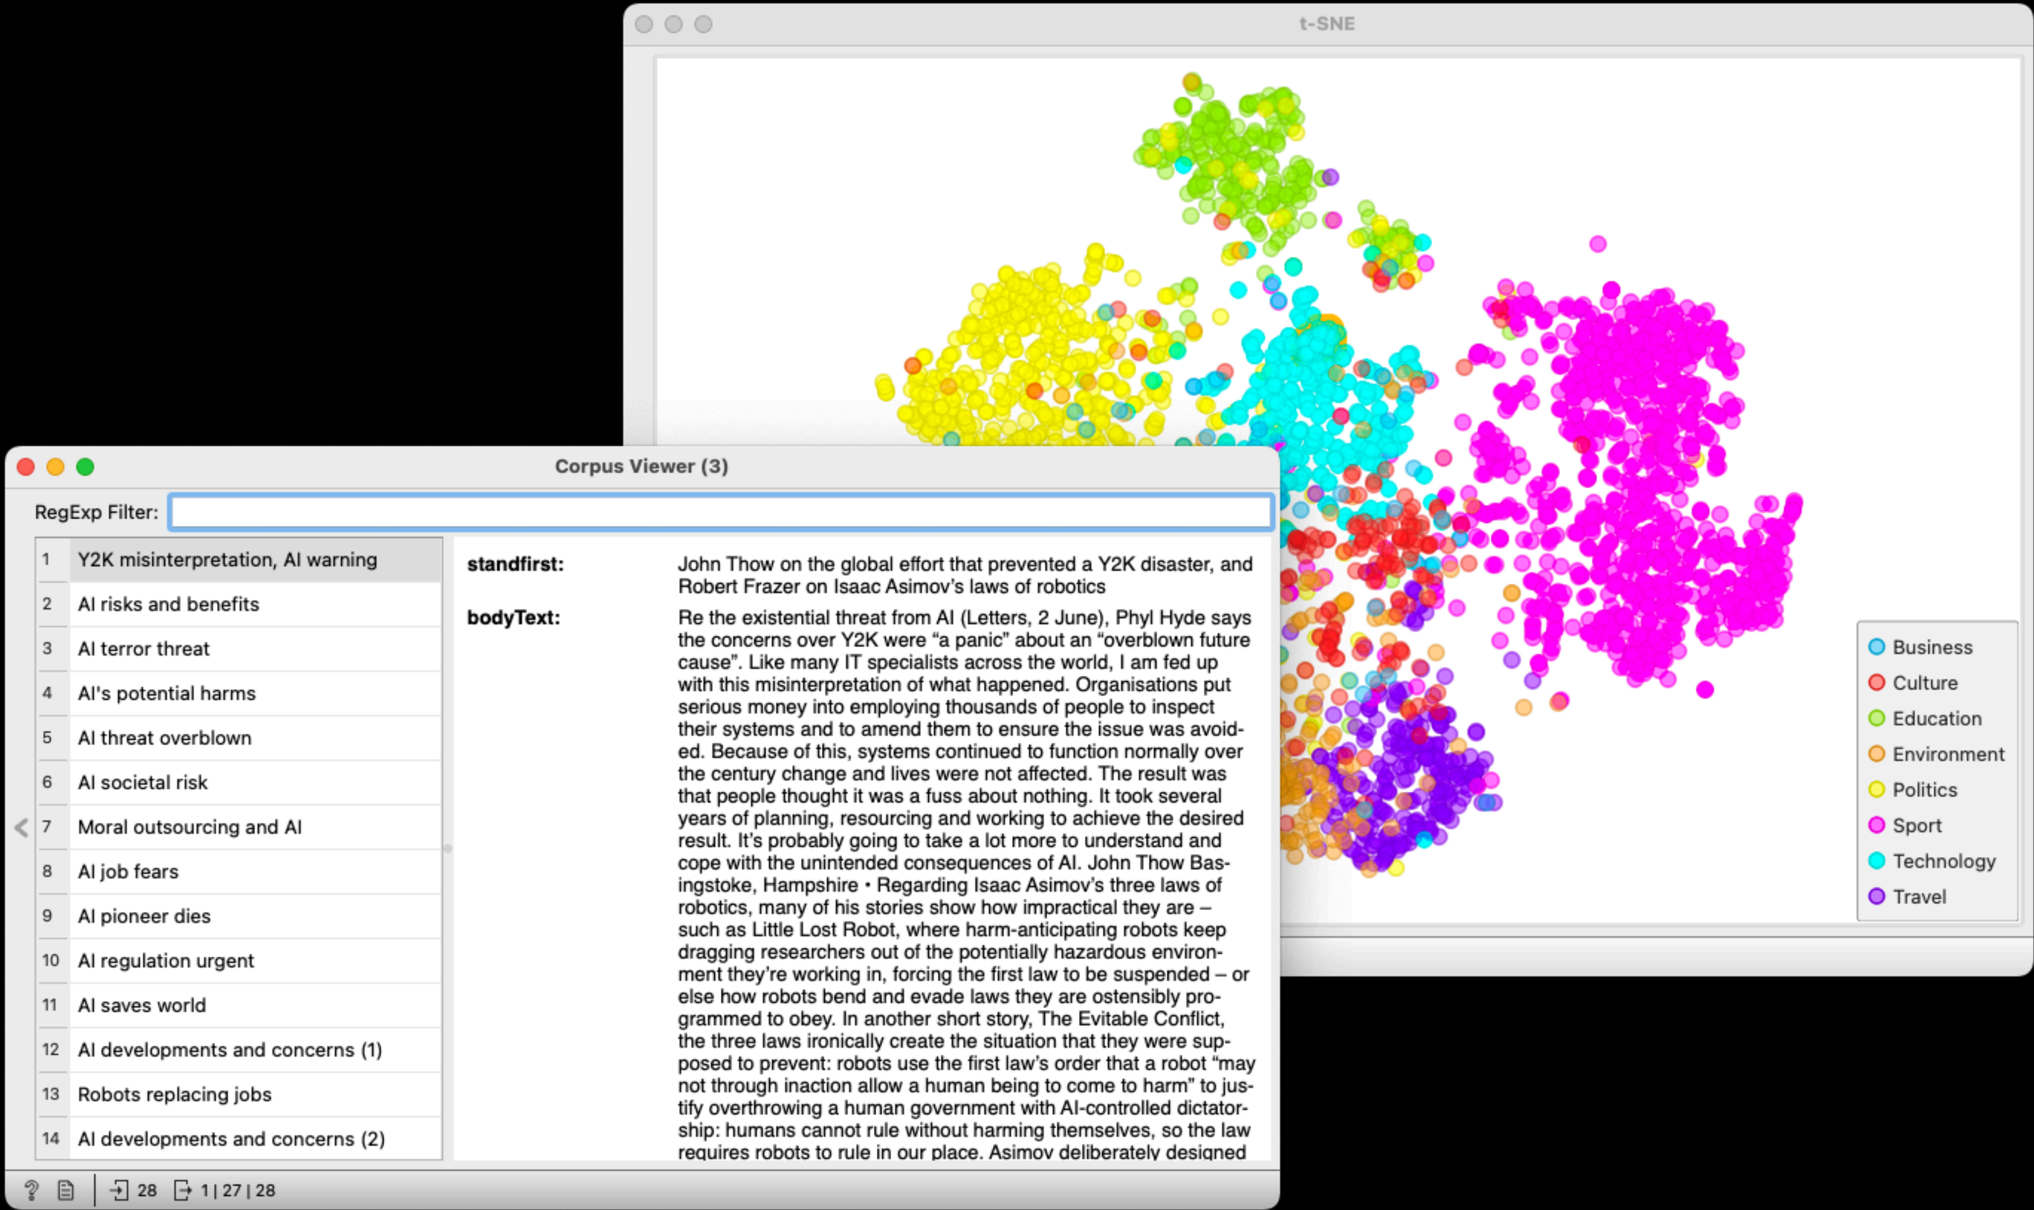
t-SNE
Business
Culture
Education
Environment
Politics
Sport
Technology
Travel
Corpus Viewer (3)
RegExp Filter:
1 Y2K misinterpretation, AI warning
2 AI risks and benefits
3 AI terror threat
4 AI's potential harms
5 AI threat overblown
6 AI societal risk
7 Moral outsourcing and AI
8 AI job fears
9 AI pioneer dies
10 AI regulation urgent
11 AI saves world
12 AI developments and concerns (1)
13 Robots replacing jobs
14 AI developments and concerns (2)
standfirst:
John Thow on the global effort that prevented a Y2K disaster, and Robert Frazer on Isaac Asimov's laws of robotics
bodyText:
Re the existential threat from AI (Letters, 2 June), Phyl Hyde says the concerns over Y2K were "a panic" about an "overblown future cause". Like many IT specialists across the world, I am fed up with this misinterpretation of what happened. Organisations put serious money into employing thousands of people to inspect their systems and to amend them to ensure the issue was avoided. Because of this, systems continued to function normally over the century change and lives were not affected. The result was that people thought it was a fuss about nothing. It took several years of planning, resourcing and working to achieve the desired result. It's probably going to take a lot more to understand and cope with the unintended consequences of AI. John Thow Basingstoke, Hampshire • Regarding Isaac Asimov's three laws of robotics, many of his stories show how impractical they are – such as Little Lost Robot, where harm-anticipating robots keep dragging researchers out of the potentially hazardous environment they're working in, forcing the first law to be suspended – or else how robots bend and evade laws they are ostensibly programmed to obey. In another short story, The Evitable Conflict, the three laws ironically create the situation that they were supposed to prevent: robots use the first law's order that a robot "may not through inaction allow a human being to come to harm" to justify overthrowing a human government with AI-controlled dictatorship: humans cannot rule without harming themselves, so the law requires robots to rule in our place. Asimov deliberately designed
28
1 | 27 | 28

Generative AI

Type of artificial intelligence that can create new content, like text, images, or music, by learning patterns from existing data.

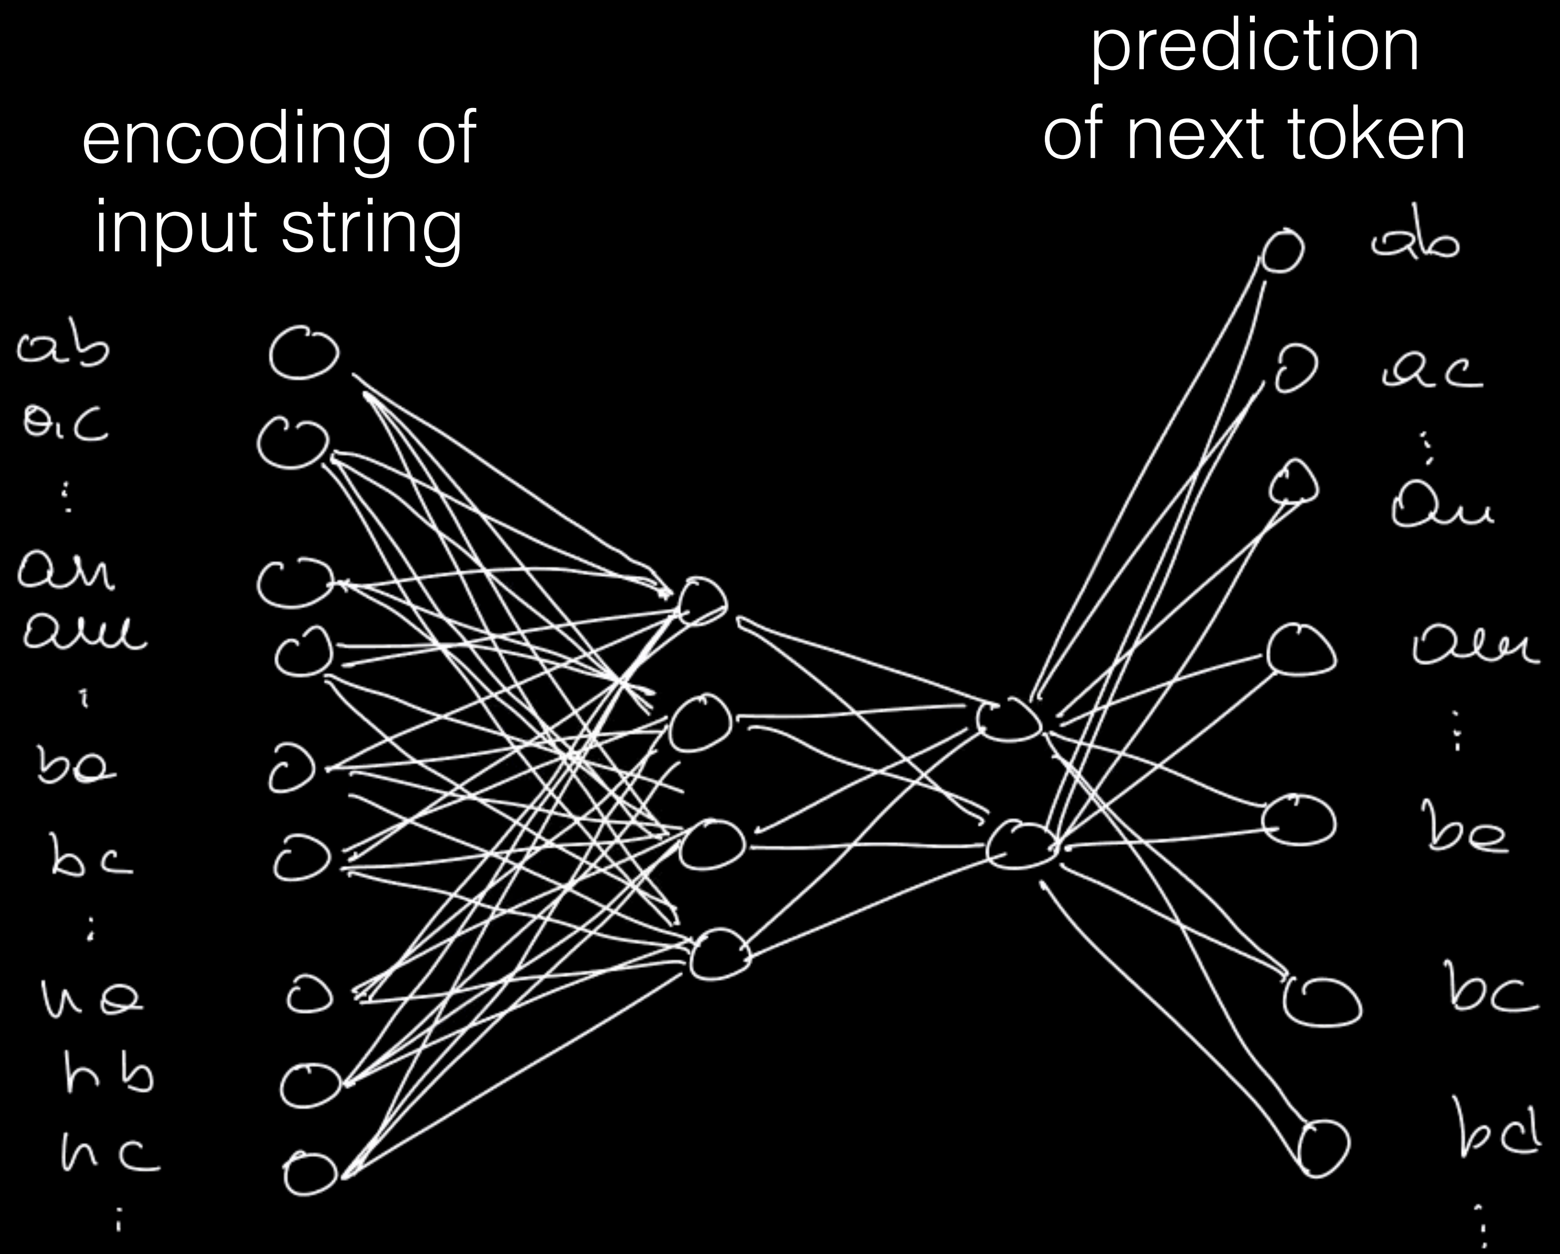
encoding of
input string
prediction
of next token
ab
ac
an
au
ba
bc
na
nb
nc
ab
ac
au
an
be
bc
bd

encoding of
input string

prediction
of next token

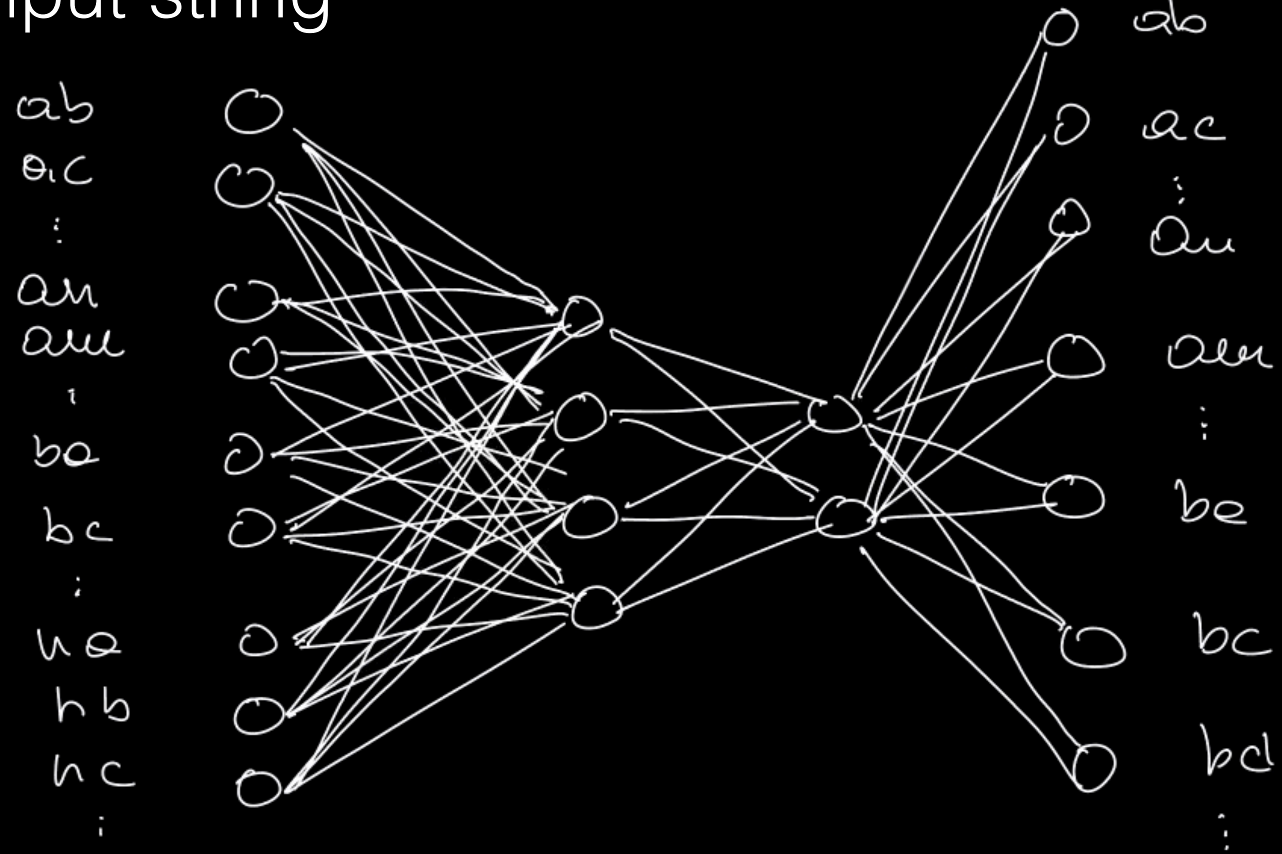

She added a banana to her …

… mo

encoding of
input string

prediction
of next token

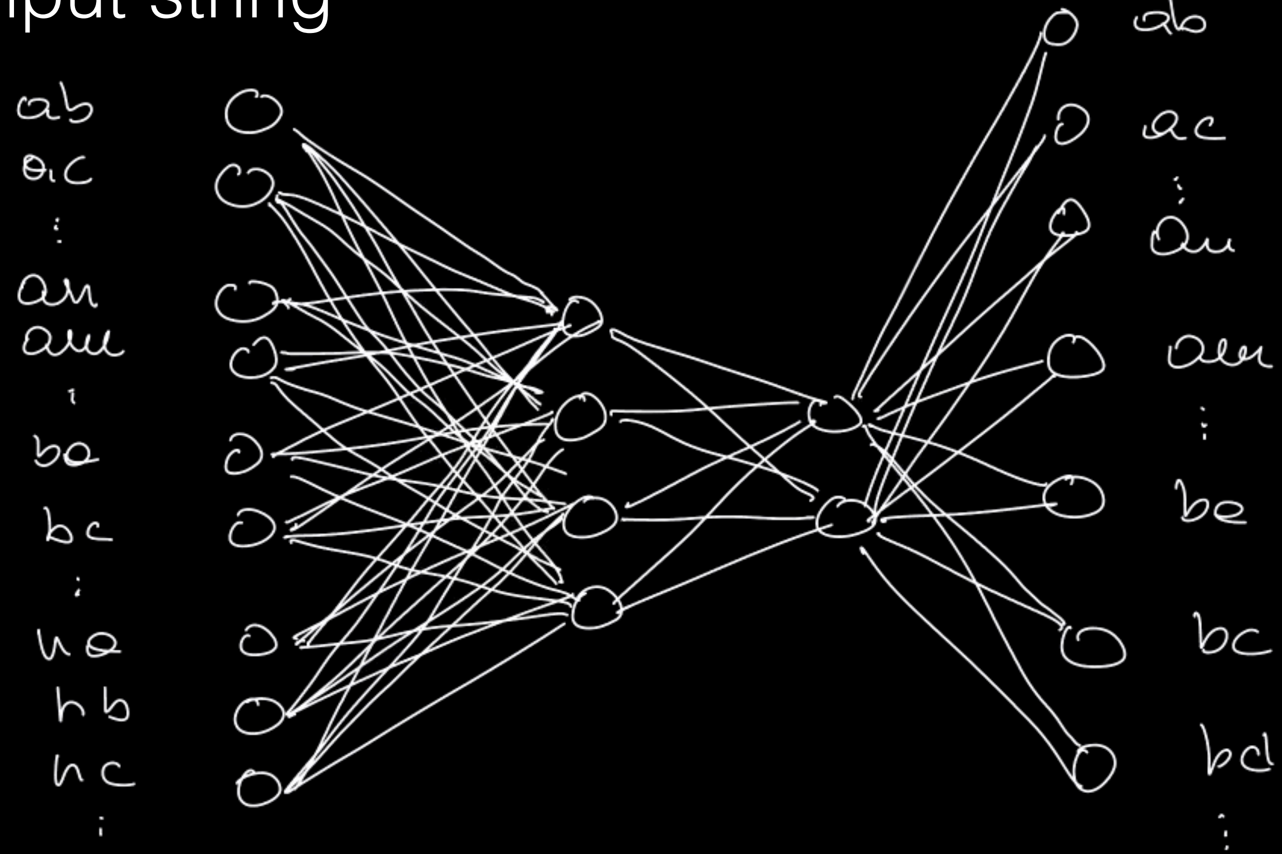

She added a banana to her mo…

… rn

encoding of
input string

prediction
of next token

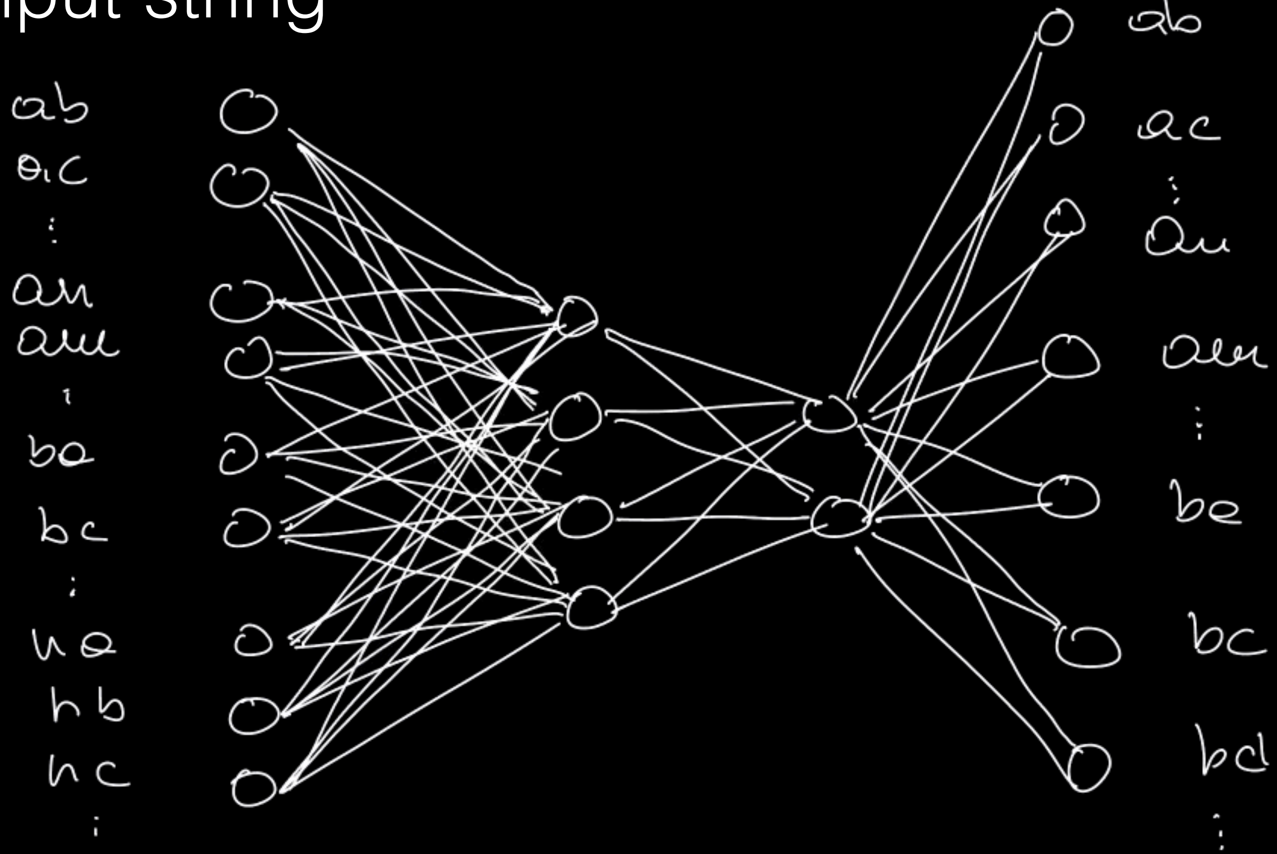

She added a banana to her morn…

… in

encoding of
input string

prediction
of next token

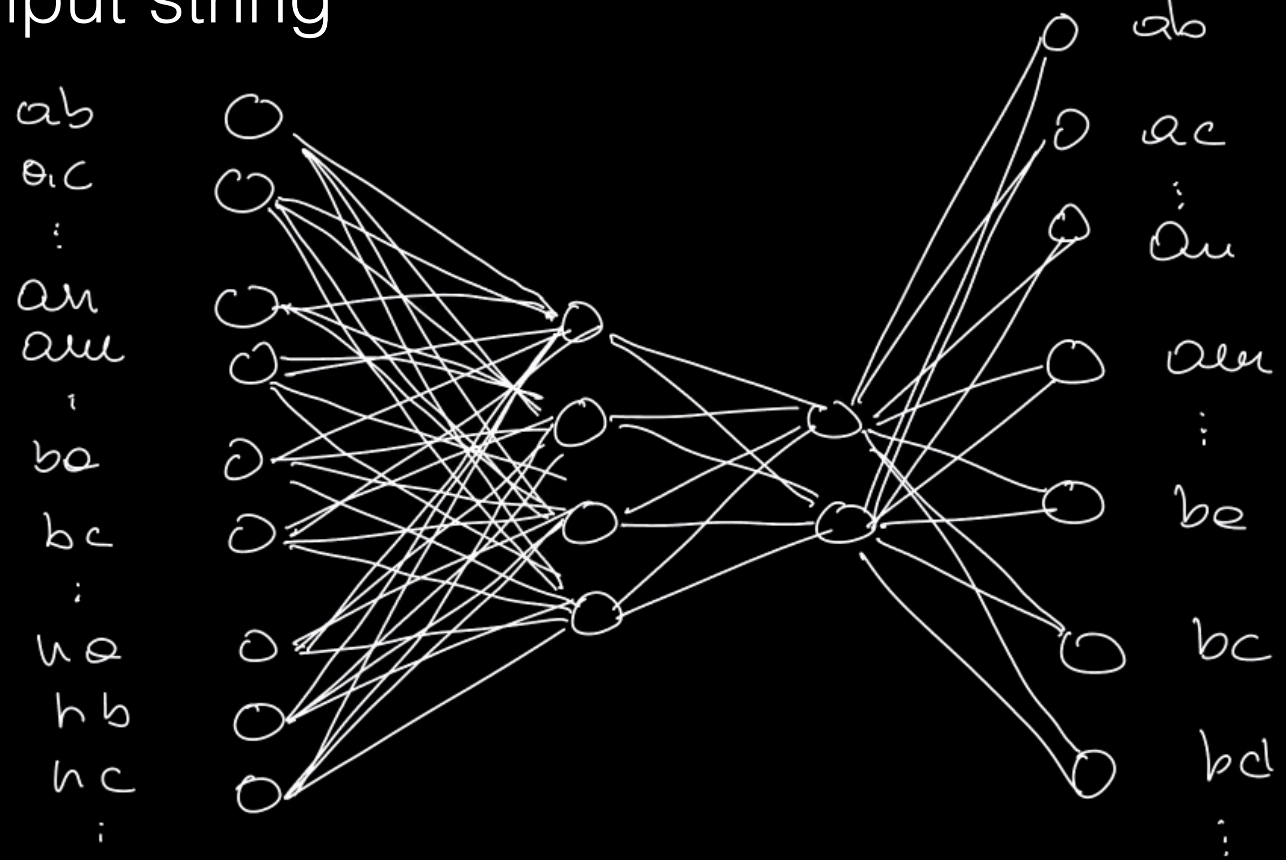

She added a banana to her mornin…

… g

encoding of
input string

prediction
of next token

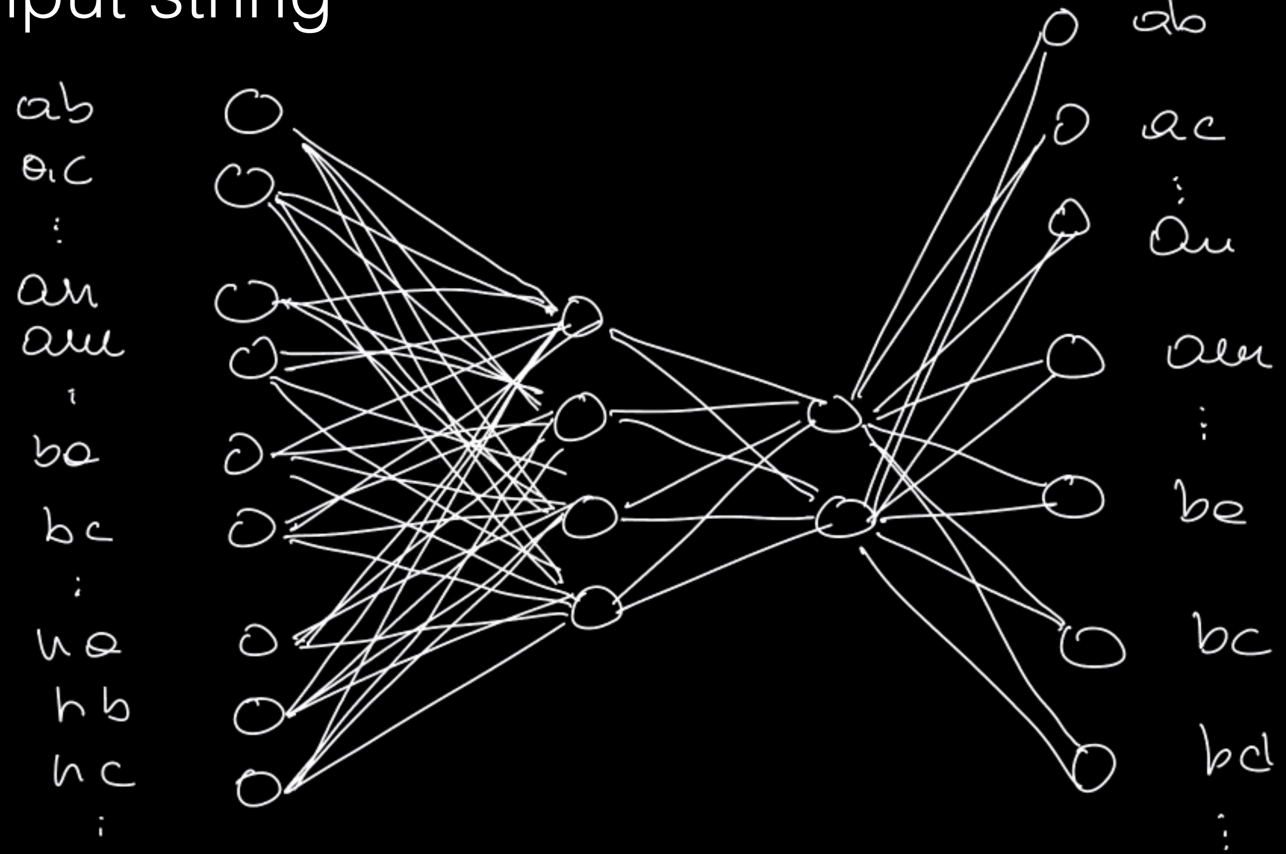

She added a banana to her morning smoothie.

Machine Learning Models Can Be Large

Swimming in Bled … 3 parameters

AlexNet … 60 million parameters

GPT-3.5 … 175 billion parameters

DeepSeek-V3 … 671 billion parameters

Machine Learning Models Can Be Large

Swimming in Bled … 3 parameters

AlexNet … 60 million parameters

GPT-3.5 … 17

DeepSeek-V3 …

Training dataset: 15 trillion tokens

~150 million books

2,000 Nvidia H800 graphics cards

Price per card: €15,000

Development time: two months

Development cost: $5.6 million

Machine Learning Models Can Be Large

Swimming in Bled … 3 parameters
AlexNet … 60 million parameters
GPT-3.5 … 175 billion parameters
DeepSeek-V3 … 671 billion parameters

learning requires a lot of data
many models are freely available
texts, images, sound, sequences
medicine, engineering, chemistry, law, archaeology, telecommunications, anything...

Generative AI and LLM

- Context
- Encoding of context (sequence of tokens)
- (Iterative) output prediction (sequence of tokens)

Training LLMs: Pre-Training

Goal: Teach the model general language understanding.

Data: Massive, diverse text corpora (web data, books, articles…).

- Learns grammar, world knowledge, reasoning patterns.
- Self-supervised learning (predict next token or fill in missing text).
- Produces a foundation model.

Training LLMs: Supervised Fine-Tuning

Goal: Align the model with specific tasks and human-determined behaviors.

Data: Human-curated instruction–response pairs (several hundred individuals, several 100k instructions-responses).

- Trains the model to follow instructions.
- Teaches desired formats and context handling.
- Reduces hallucinations compared to the raw pre-trained model.

Training LLMs: Reinforcement Learning

Goal: Align outputs with human preferences.

Data: Human ratings or preference comparisons (several hundred annotators, several thousand comparisons)

- Model generates multiple responses → human (or AI) ranks them.
- Reward model is trained from rankings.
- Final model learns to optimize for preferred responses.

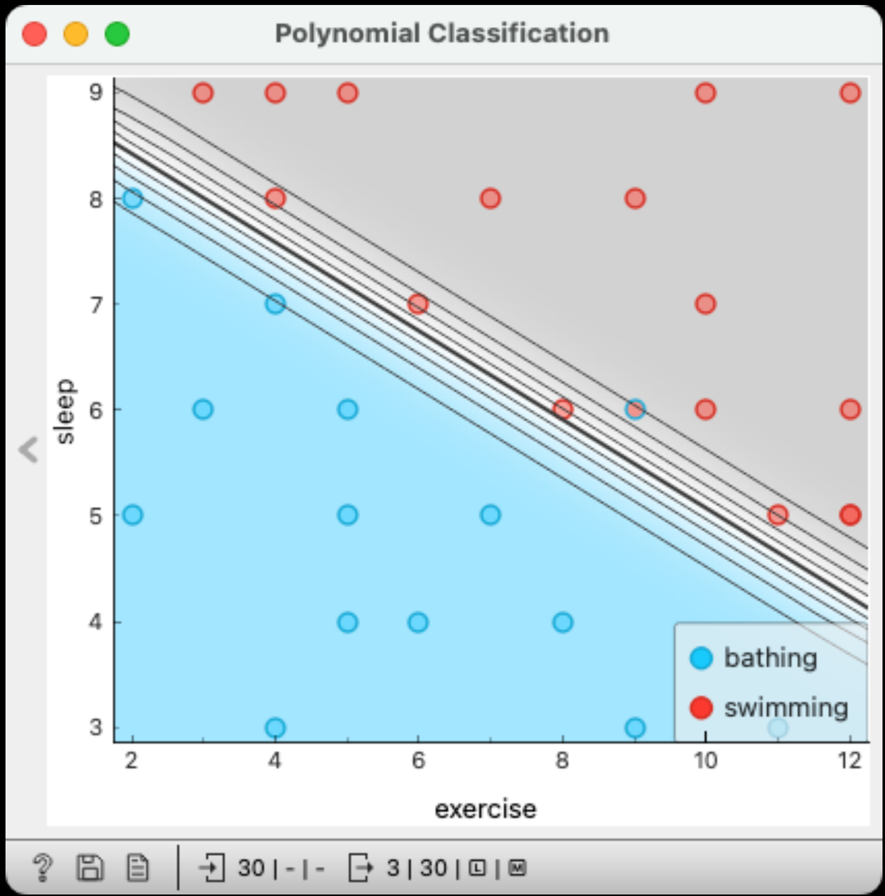

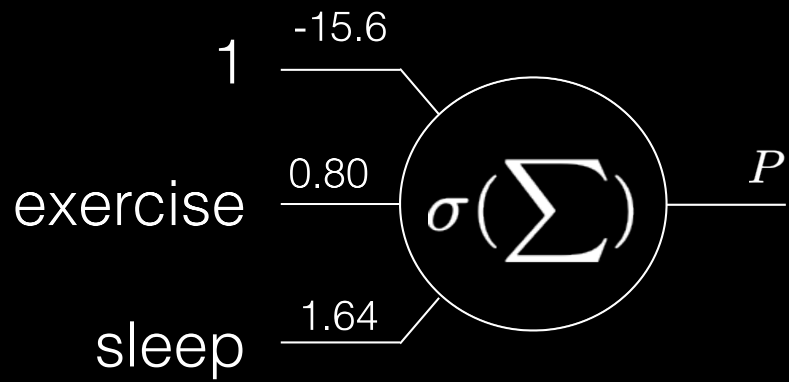

Our swimming model had 3 parameters!

GPT-3.5 has 175 billion parameters!

GPT-4 has 1.76 trillion parameters!

(that's 1,760,000,000,000!)

Chatbots

- A chatbot is a software application that uses an LLM to interact with users through natural language
- Not an AI!
- It’s web-based interface is useful but rather very limited, the more interesting things to come are underneath

Agents

- LLM is only the text engine; the real power (today) comes from agents
- Agents: chatbot built on top of an LLM, extended with tools, memory, planning, and action-taking abilities.
- Agents can:
  - use tools & APIs
  - search the web or databases
  - coordinate multi-step workflows
  - interact with external systems
- Transition from chatbots to action-taking automation

What are then ChatGPT, Copilot, Claude, Le Chat?

- At the interface level: these are chatbots
- At the capability level: they can be agents (depends how you use it)
- The margin between chatbots and agents is becoming very thin

AI, what's Next?

Multi-Modal Models
Understanding and generating across text, images, audio, and video.

Autonomous Agents
LLMs using tools, APIs, and reasoning to complete tasks.

Smarter, Smaller Models
Efficient LLMs with strong performance and lower compute needs.

Personalized AI
Models adapted to individuals, teams, or organizations.

Human-AI Collaboration, Real-Time Learning
AI as a partner in creativity, research, and decision-making.