Computers, brains and AI

8: ChatGPT

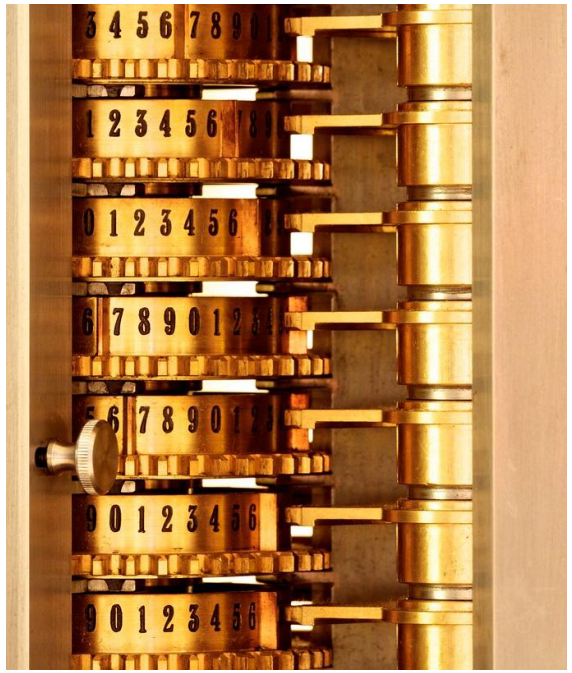



Difference engine by Marcin Wichary / CC BY
Ada Lovelace by Science Museum Group / CC0

1) Intro

2) Foundations
3) Silicon

4) Deep learning
5) Deep learning

6) ChatGPT
7) ChatGPT
8) ChatGPT

9) The future of AI
10) The future of AI

- Training
- Transformers
- Fine tuning
- Reasoning



The Riemann hypothesis 1859



Claude helps find a proof 20 Nov 2026

Reactions to the proof

James Maynard described the proof, if correct, as "a historic milestone".[19] Brian Conrey said of the development, "It's a monumental moment in human history. It's really the biggest thing I think that's ever happened in mathematics."[20]

https://www.youtube.com/watch?v=sD0NjbwqlYw

3Blue1Brown

https://en.wikipedia.org/wiki/Riemann_hypothesis

Training

AlexNet

labelled images



dog

Fei-Fei Li



ImageNet (2009)

14 million ***labelled*** images

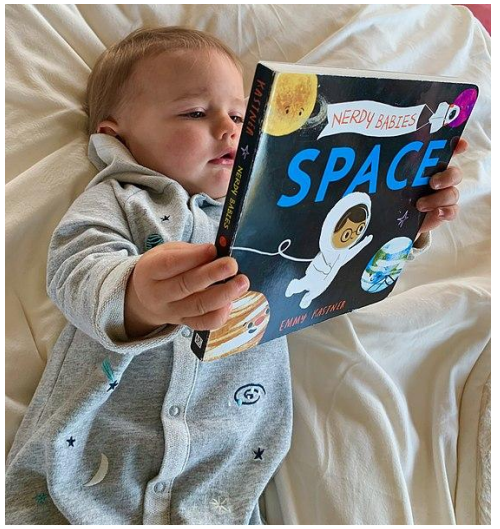

Training

AlexNet

labelled images



dog

Language models

Unlabelled text

The daytime sky appears blue

The daytime sky appears blue because

Training data

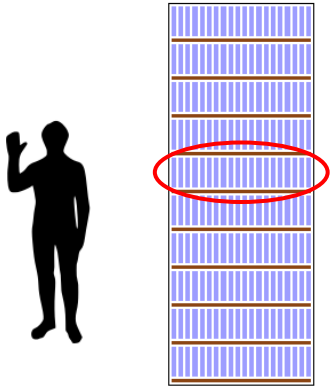

Encyclopædia Britannica

WIKIPEDIA
The Free Encyclopedia

(English edition)

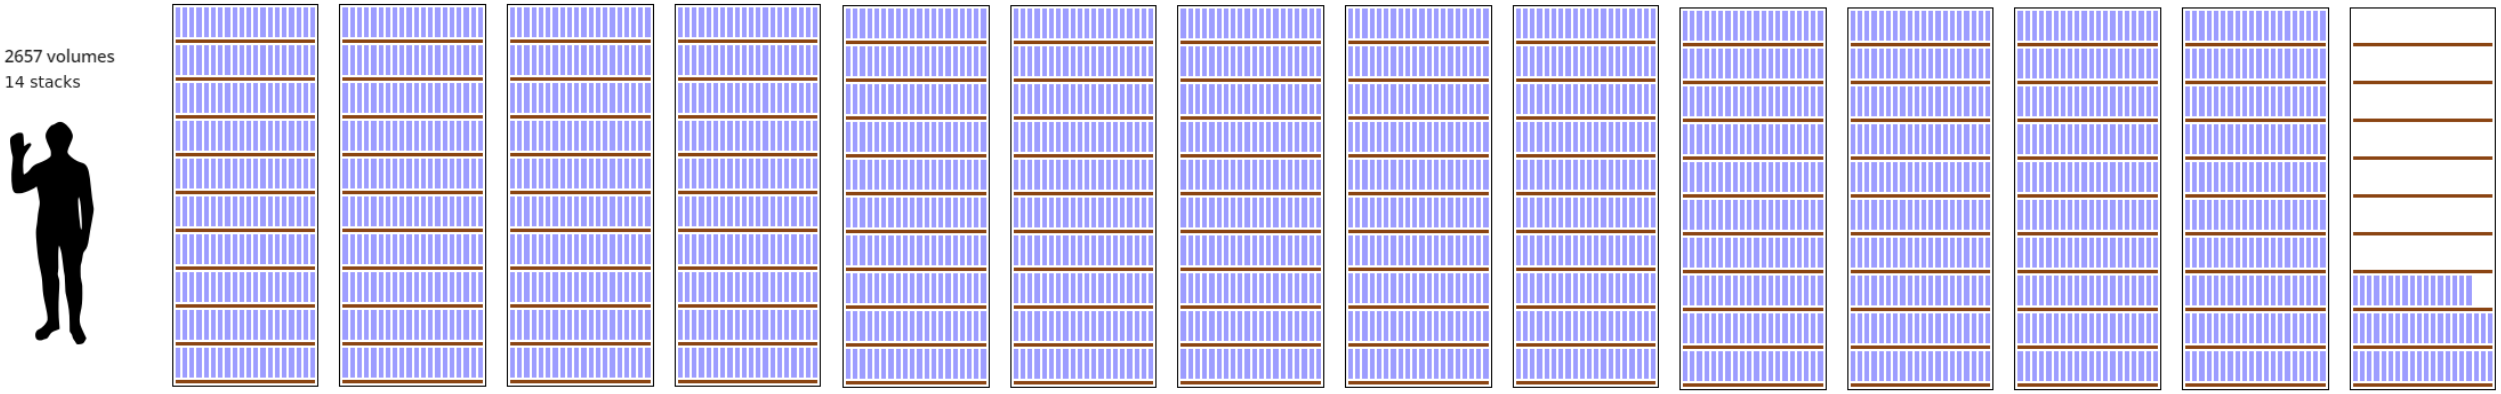

https://en.wikipedia.org/wiki/Wikipedia:Size_in_volumes

Wikipedia size reformatted by The_Anome / CC BY-SA 4.0

2657 volumes
14 stacks



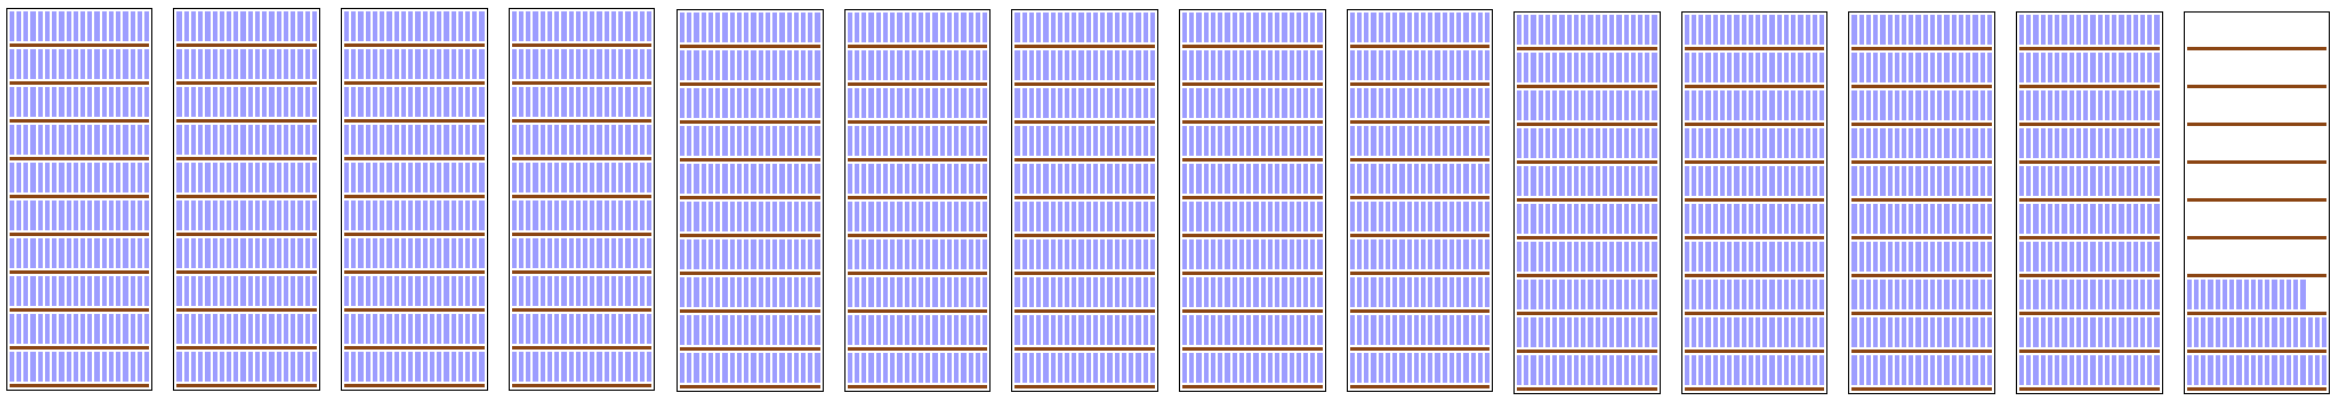



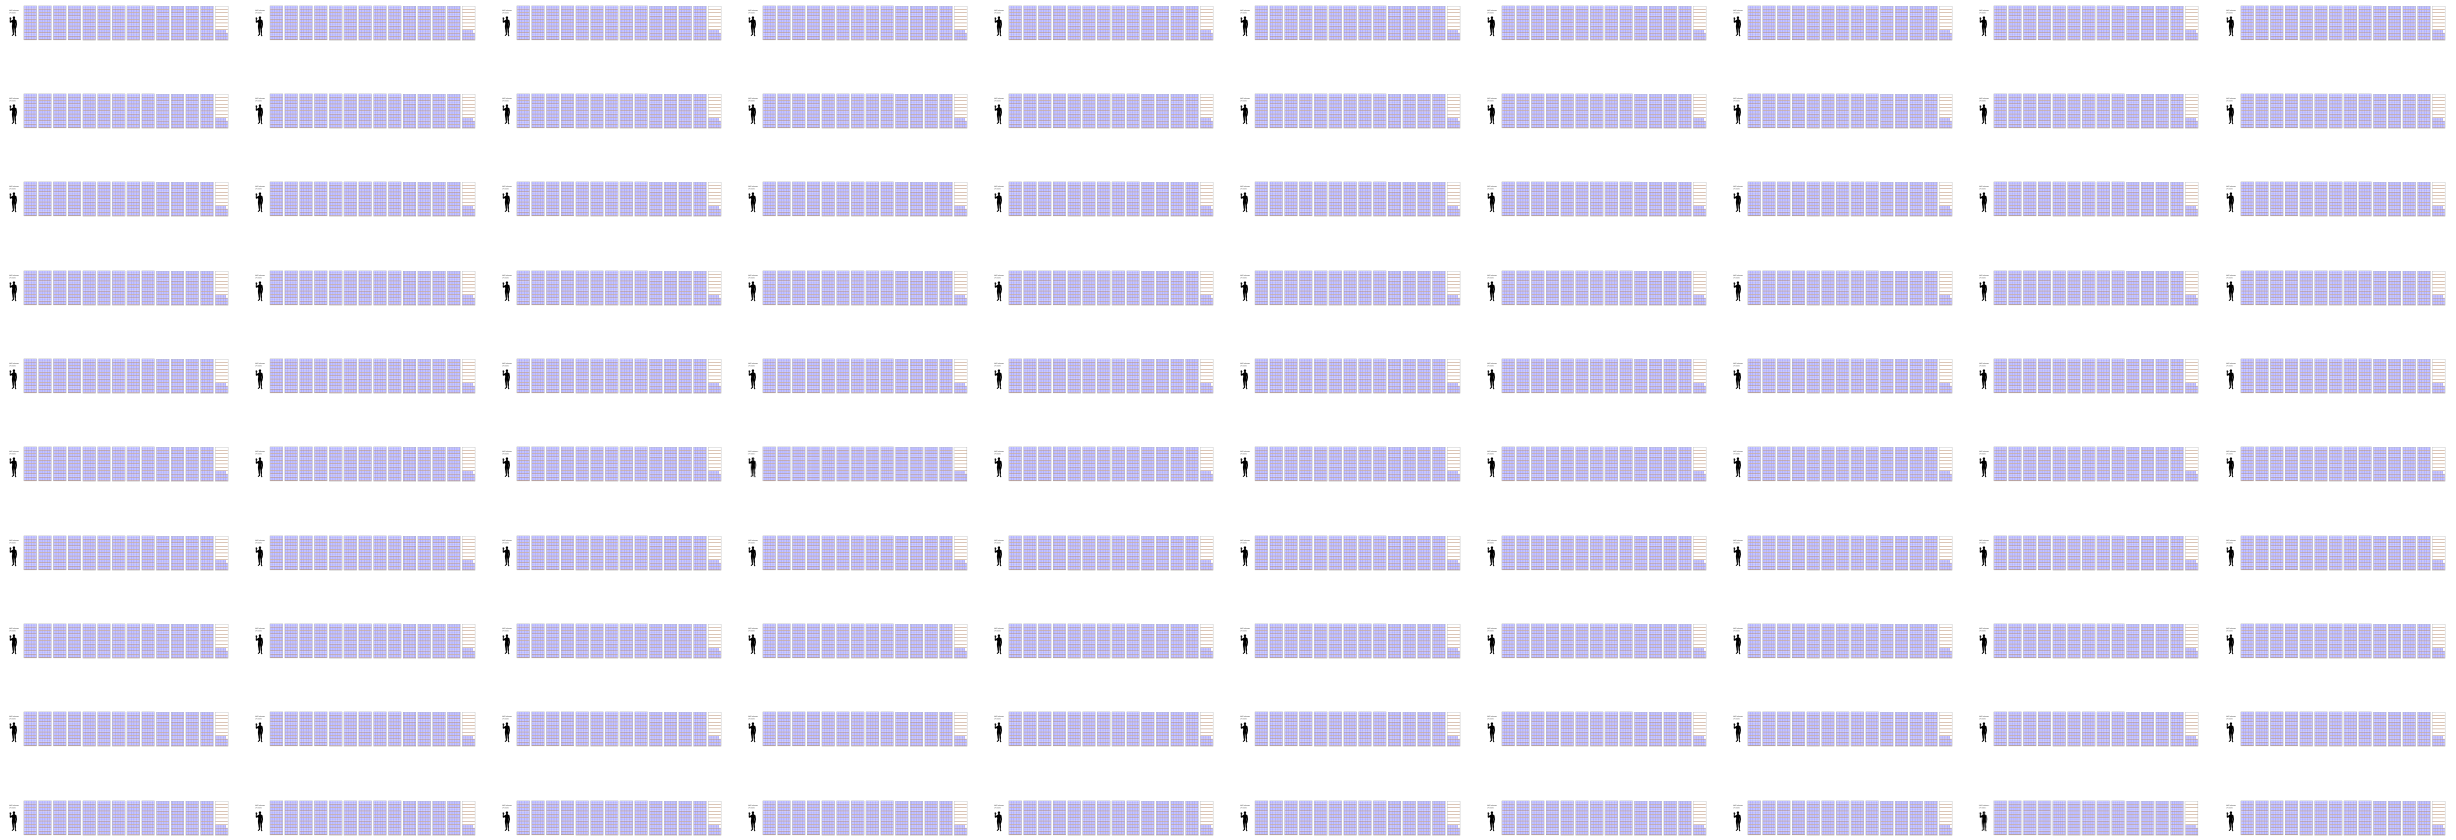

Training data for GPT-3

May 2020

Dataset	Quantity (tokens)
Common Crawl (filtered)	410 billion
WebText2	19 billion
Books1	12 billion
Books2	55 billion
Wikipedia	3 billion

https://arxiv.org/abs/2005.14165

Learning

- Start with a newborn ~~AlexNet~~ LLM
- Improve things a *tiny* step at a time



  - (by adjusting all the weights)



Training

The sky









weights

appears

dog

bus

cat

AlexNet













LLM

- Grammar
- The world
- *Anything!*

Earth from Apollo 8 by NASA / CC0
Shadow by © European Union, 2024 / CC BY 4.0
Mirror Lake By Oregon's Mt. Hood Territory / CC0

ChatGPT *chatbot*

GPT

Generative transformer

large language model

(LLM)



Transformers

280,000 citations

Attention Is All You Need - 2017

https://arxiv.org/pdf/1706.03762

a + b

c + d

parallel ☺

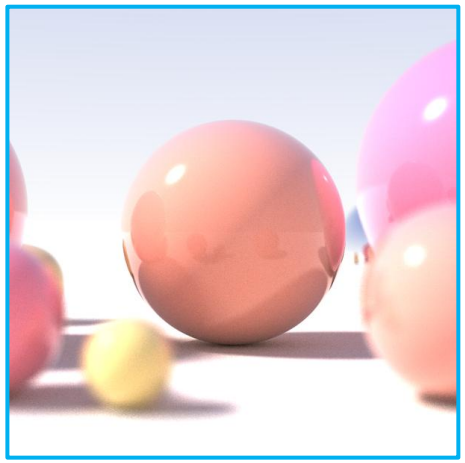

x + y + z ✗

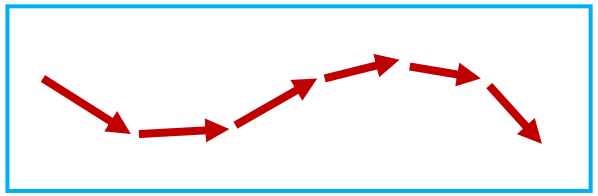

?

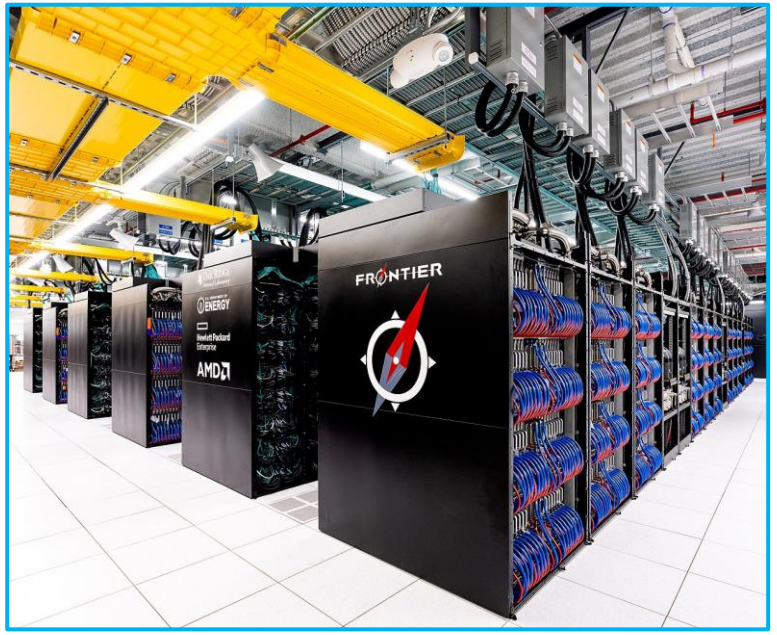



Frontier by OLCF / CC BY 2.0

Transformers

Attention Is All You Need - 2017

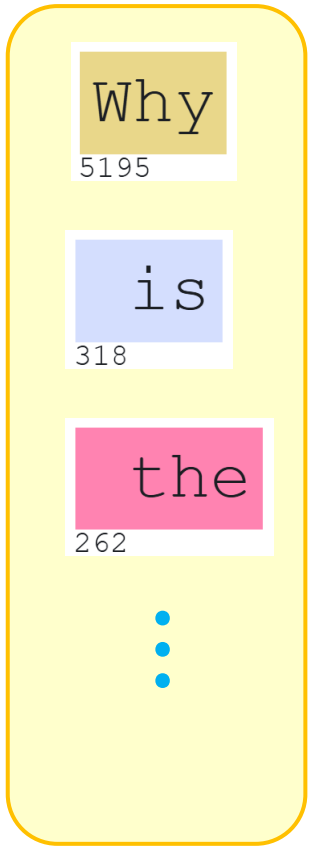

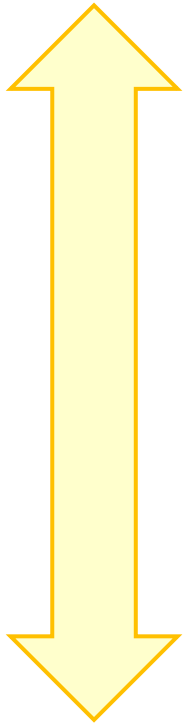

- *Parallel*
- *Large context window*
- *Expressive*

Discovered by Google…

…picked up by OpenAI



Not just text

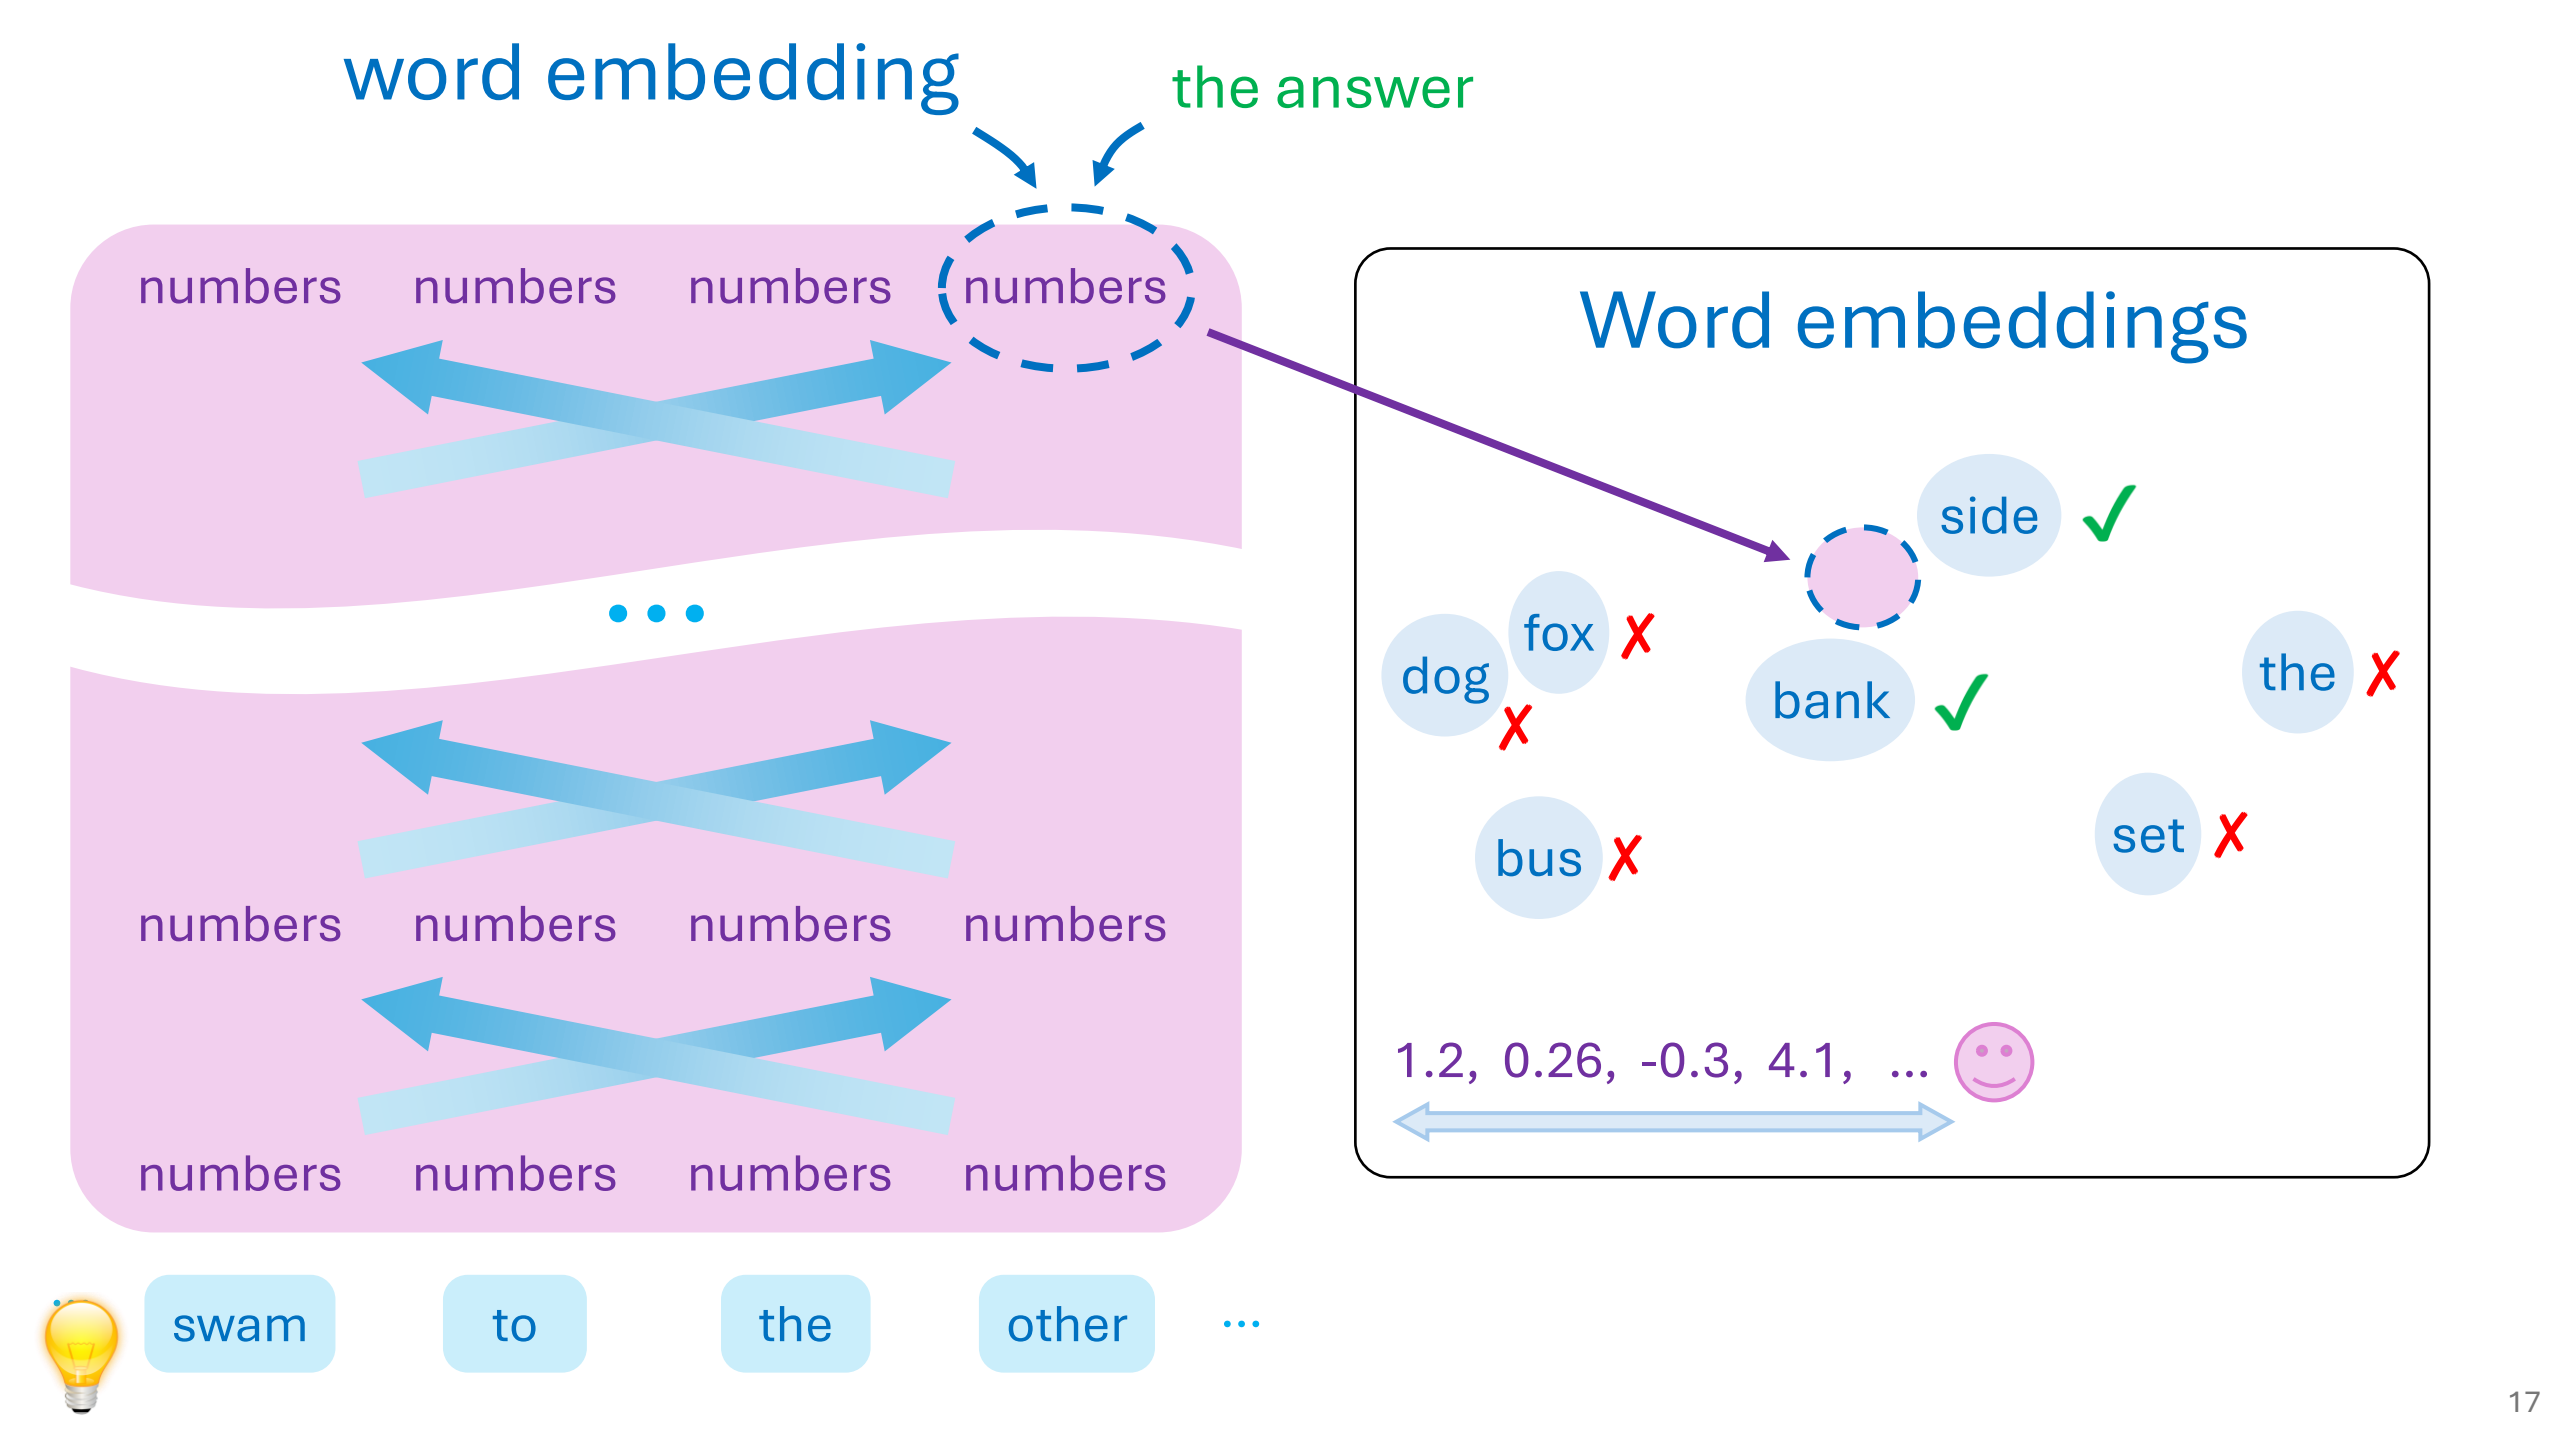
word embedding
the answer
numbers numbers numbers numbers
numbers numbers numbers numbers
numbers numbers numbers numbers
Word embeddings
side
fox
dog
bank
the
bus
set
1.2, 0.26, -0.3, 4.1, ...
swam
to
the
other
...

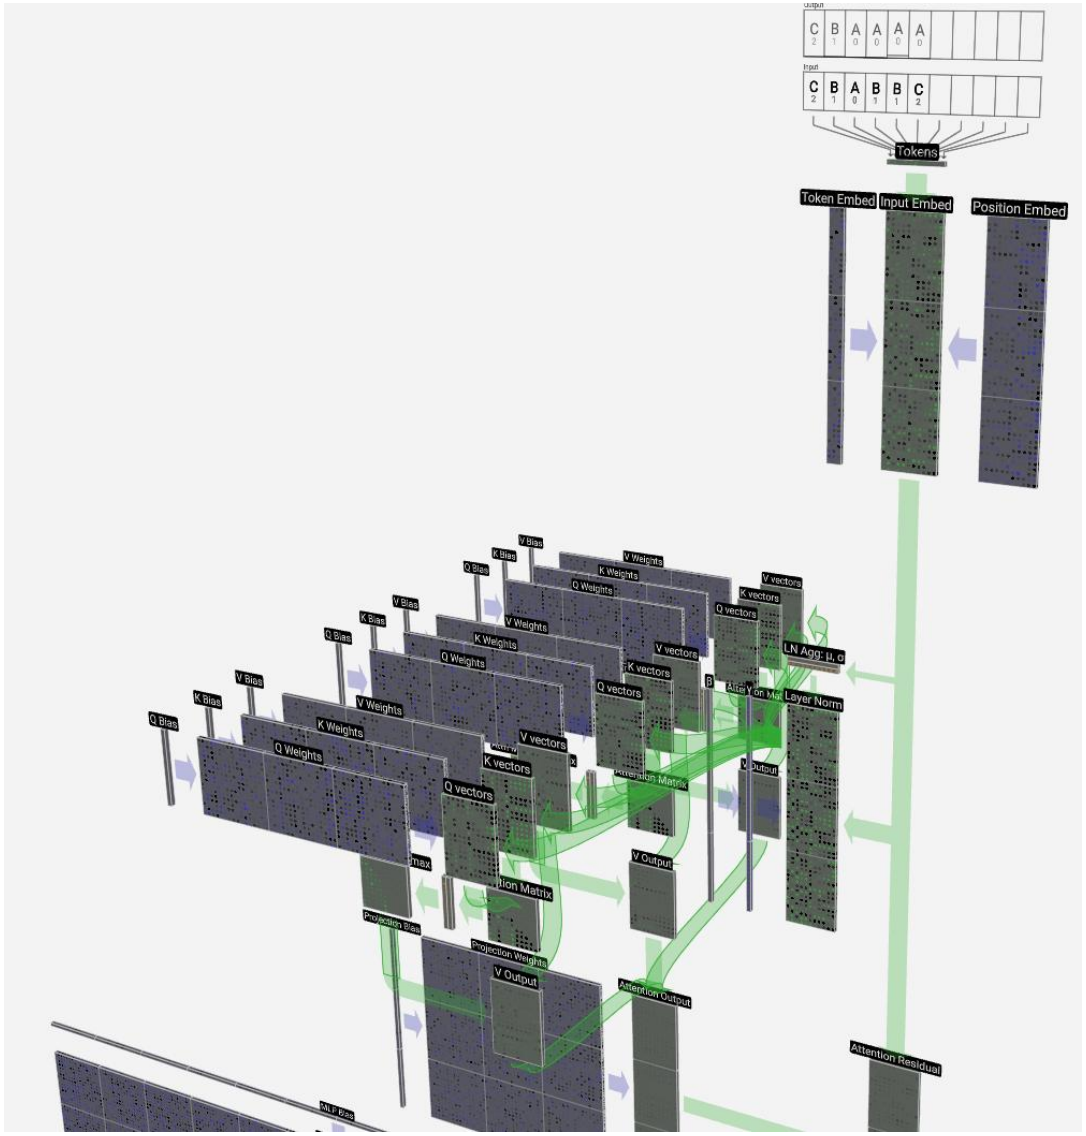

https://bbycroft.net/llm

See also

https://jalammar.github.io/illustrated-transformer/

https://jalammar.github.io/illustrated-gpt2/

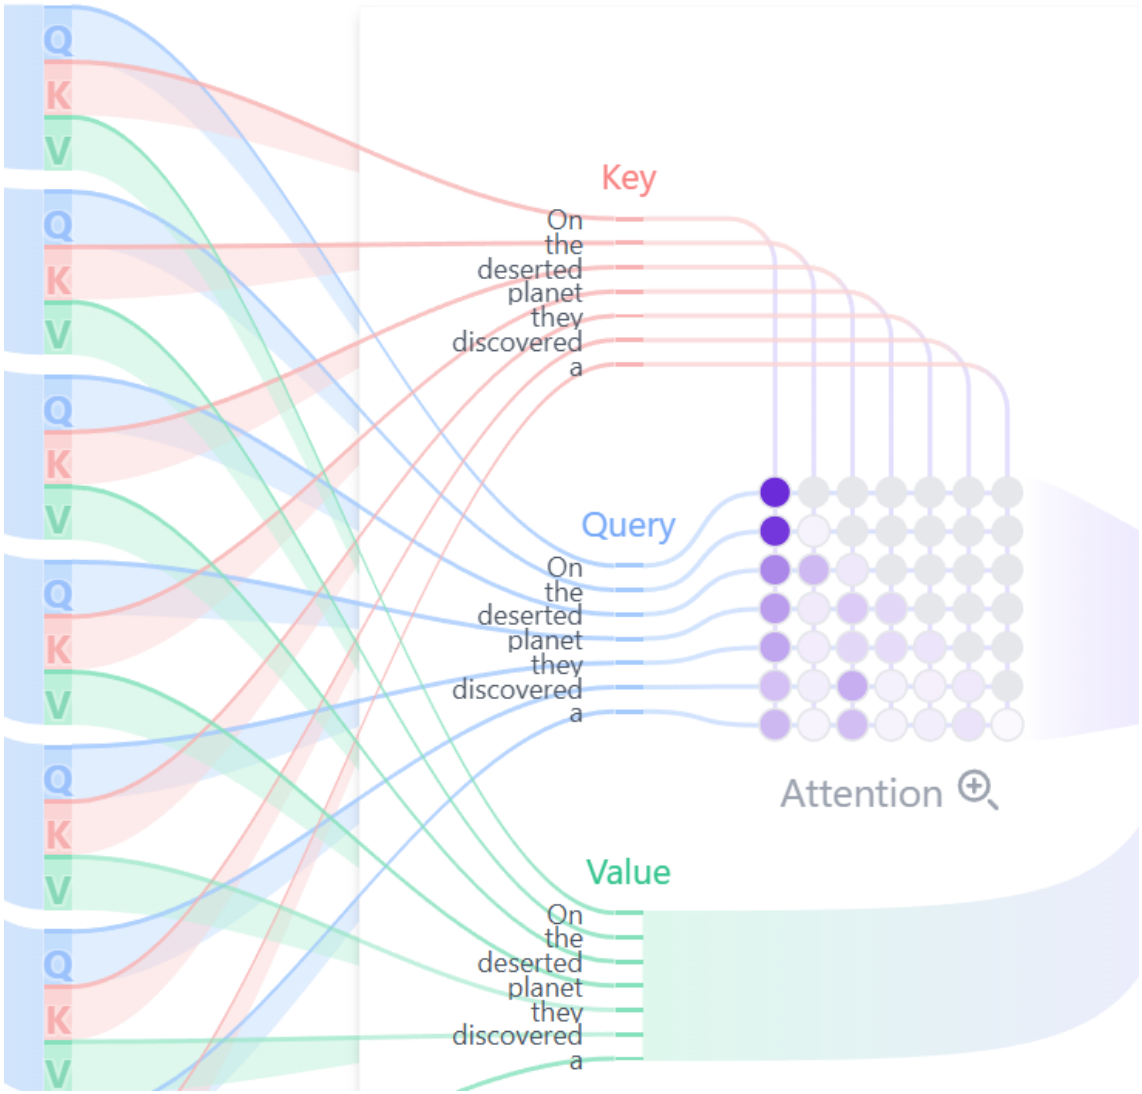

https://poloclub.github.io/transformer-explainer/

The GPT-3 Architecture, on a Napkin

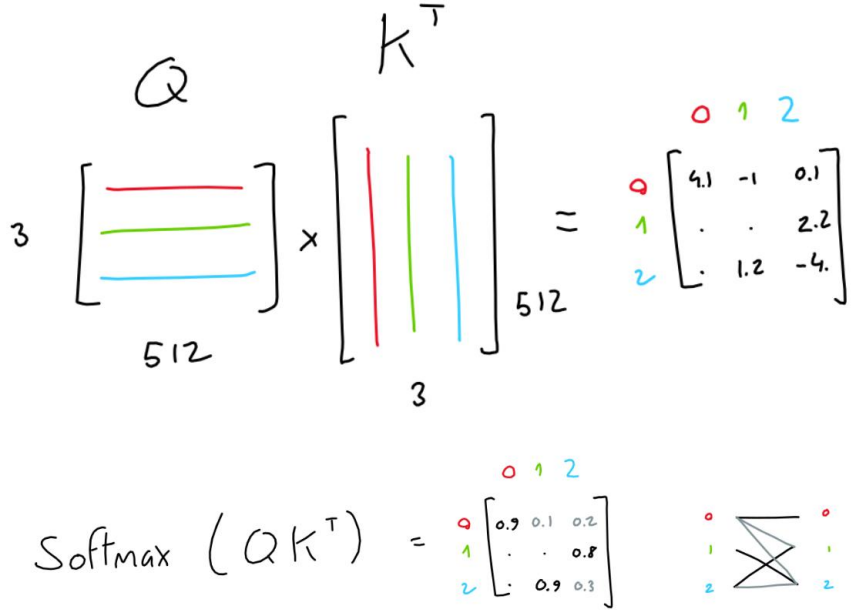

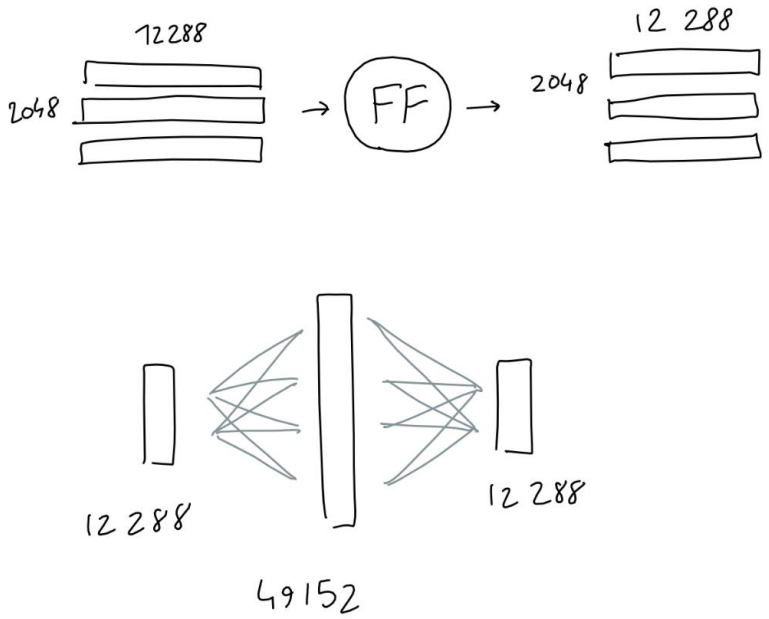

https://dugas.ch/artificial_curiosity/GPT_architecture.html

The old man the boats

adjective *noun*

Where's the verb?!

people

The old man the boats

noun *verb*



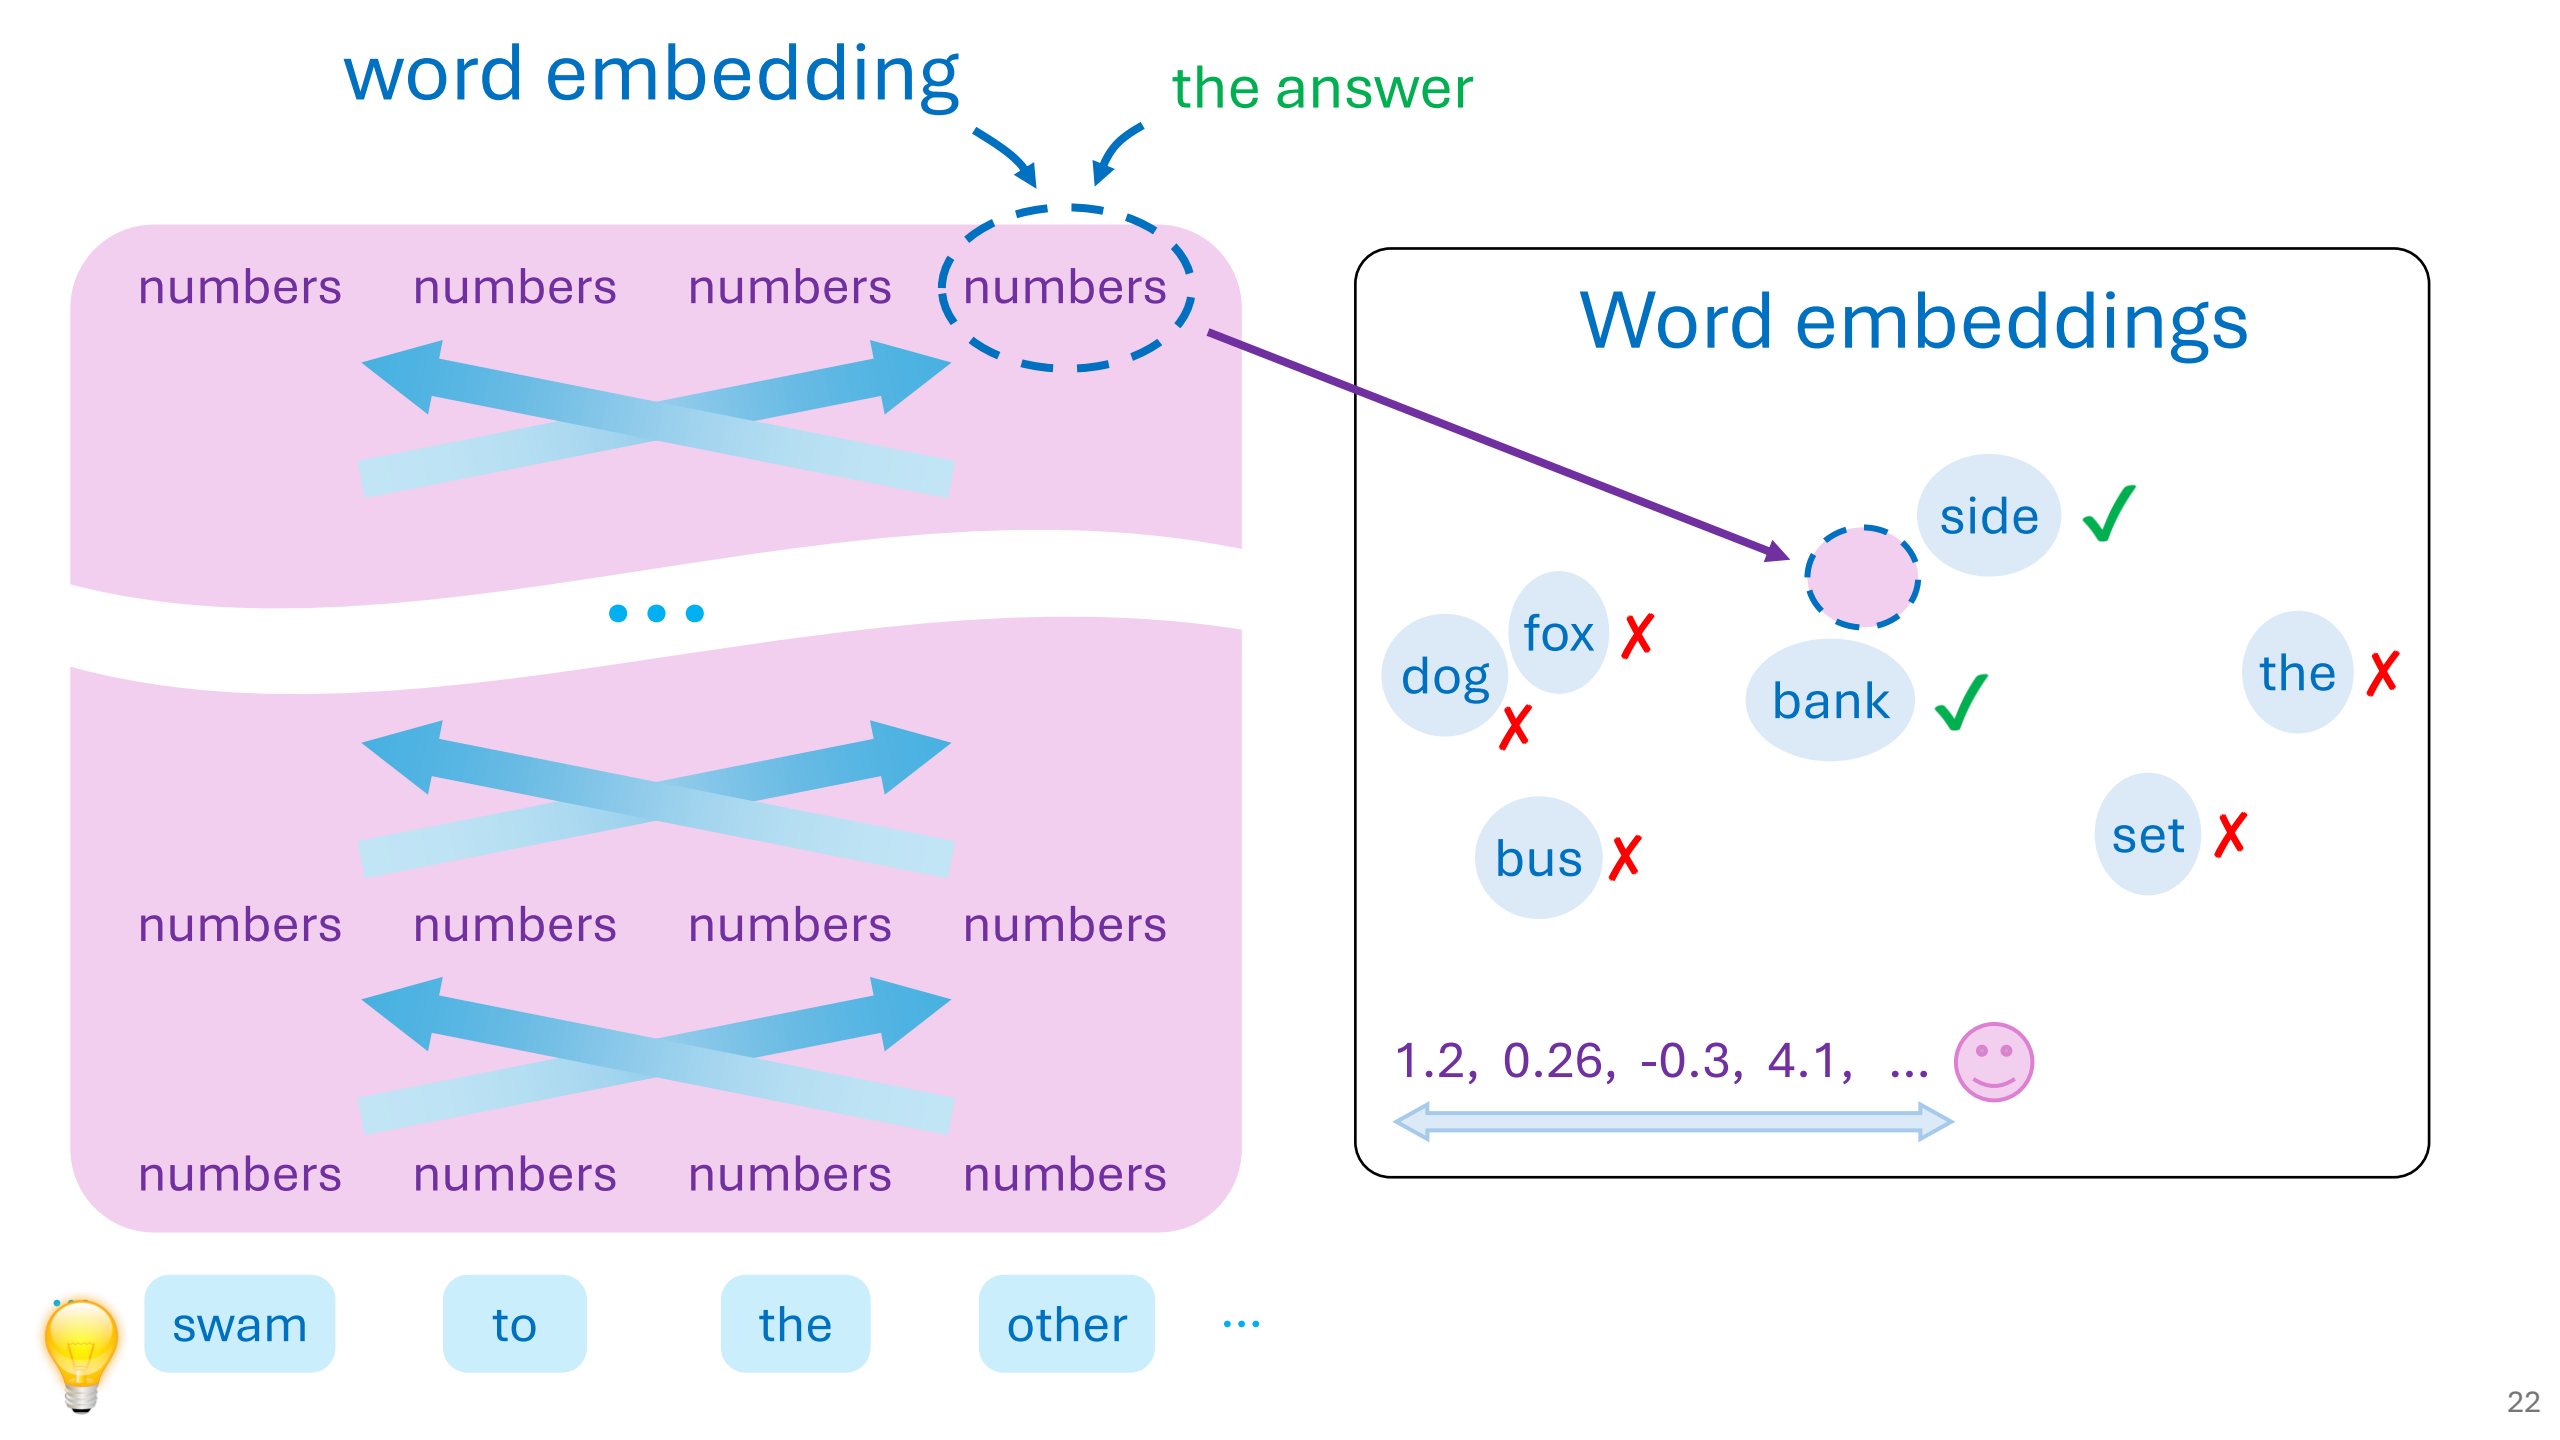
word embedding
the answer
numbers numbers numbers numbers
numbers numbers numbers numbers
numbers numbers numbers numbers
Word embeddings
side
fox
dog
bank
the
bus
set
1.2, 0.26, -0.3, 4.1, ...
swam
to
the
other
...

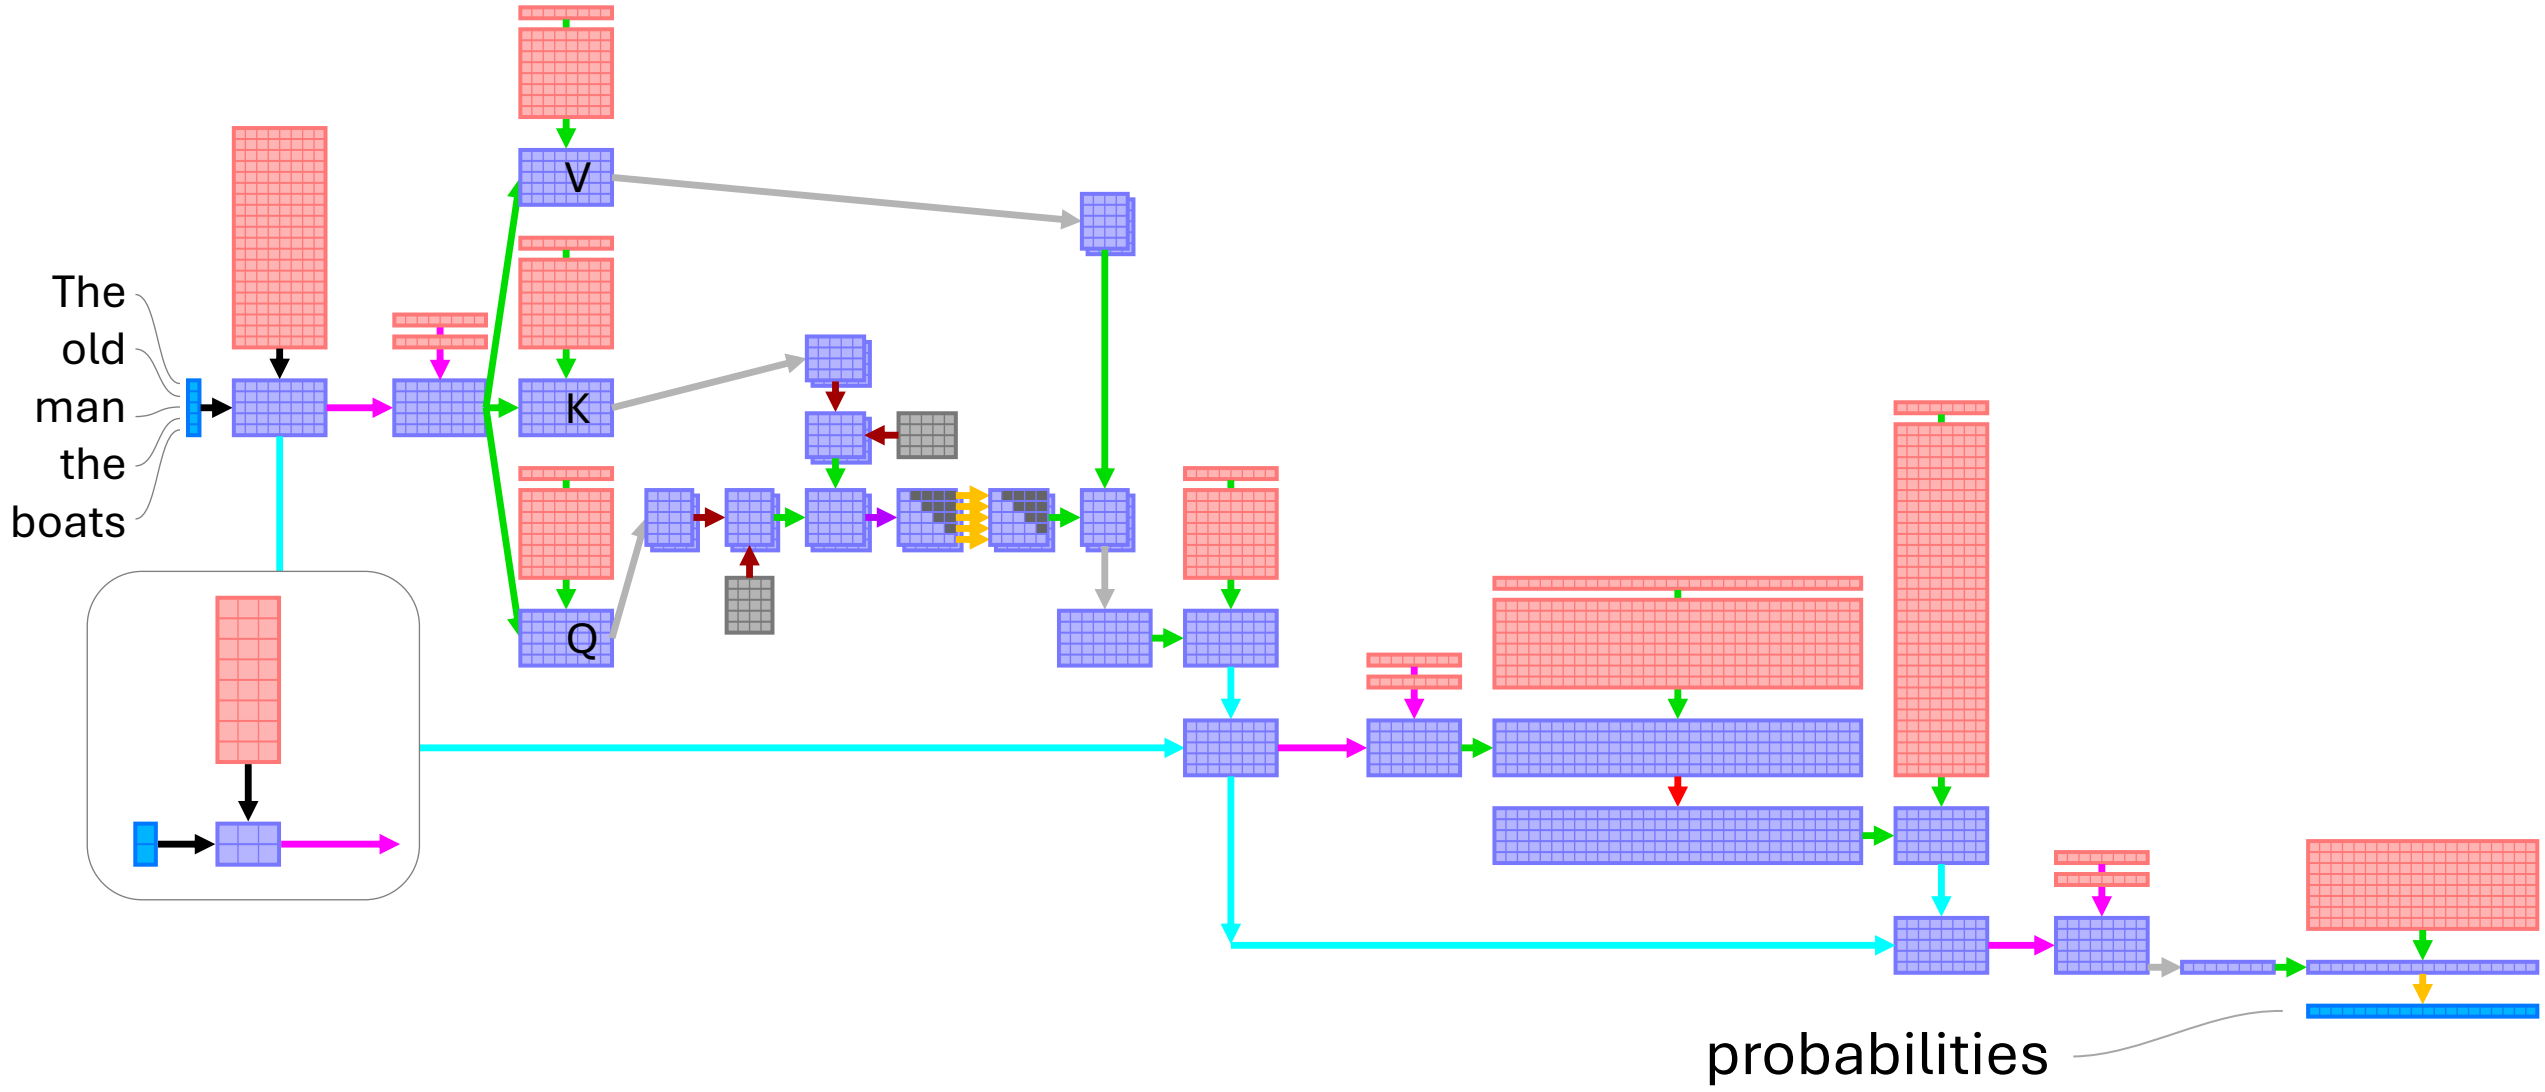
The
old
man
the
boats
V
K
Q
probabilities

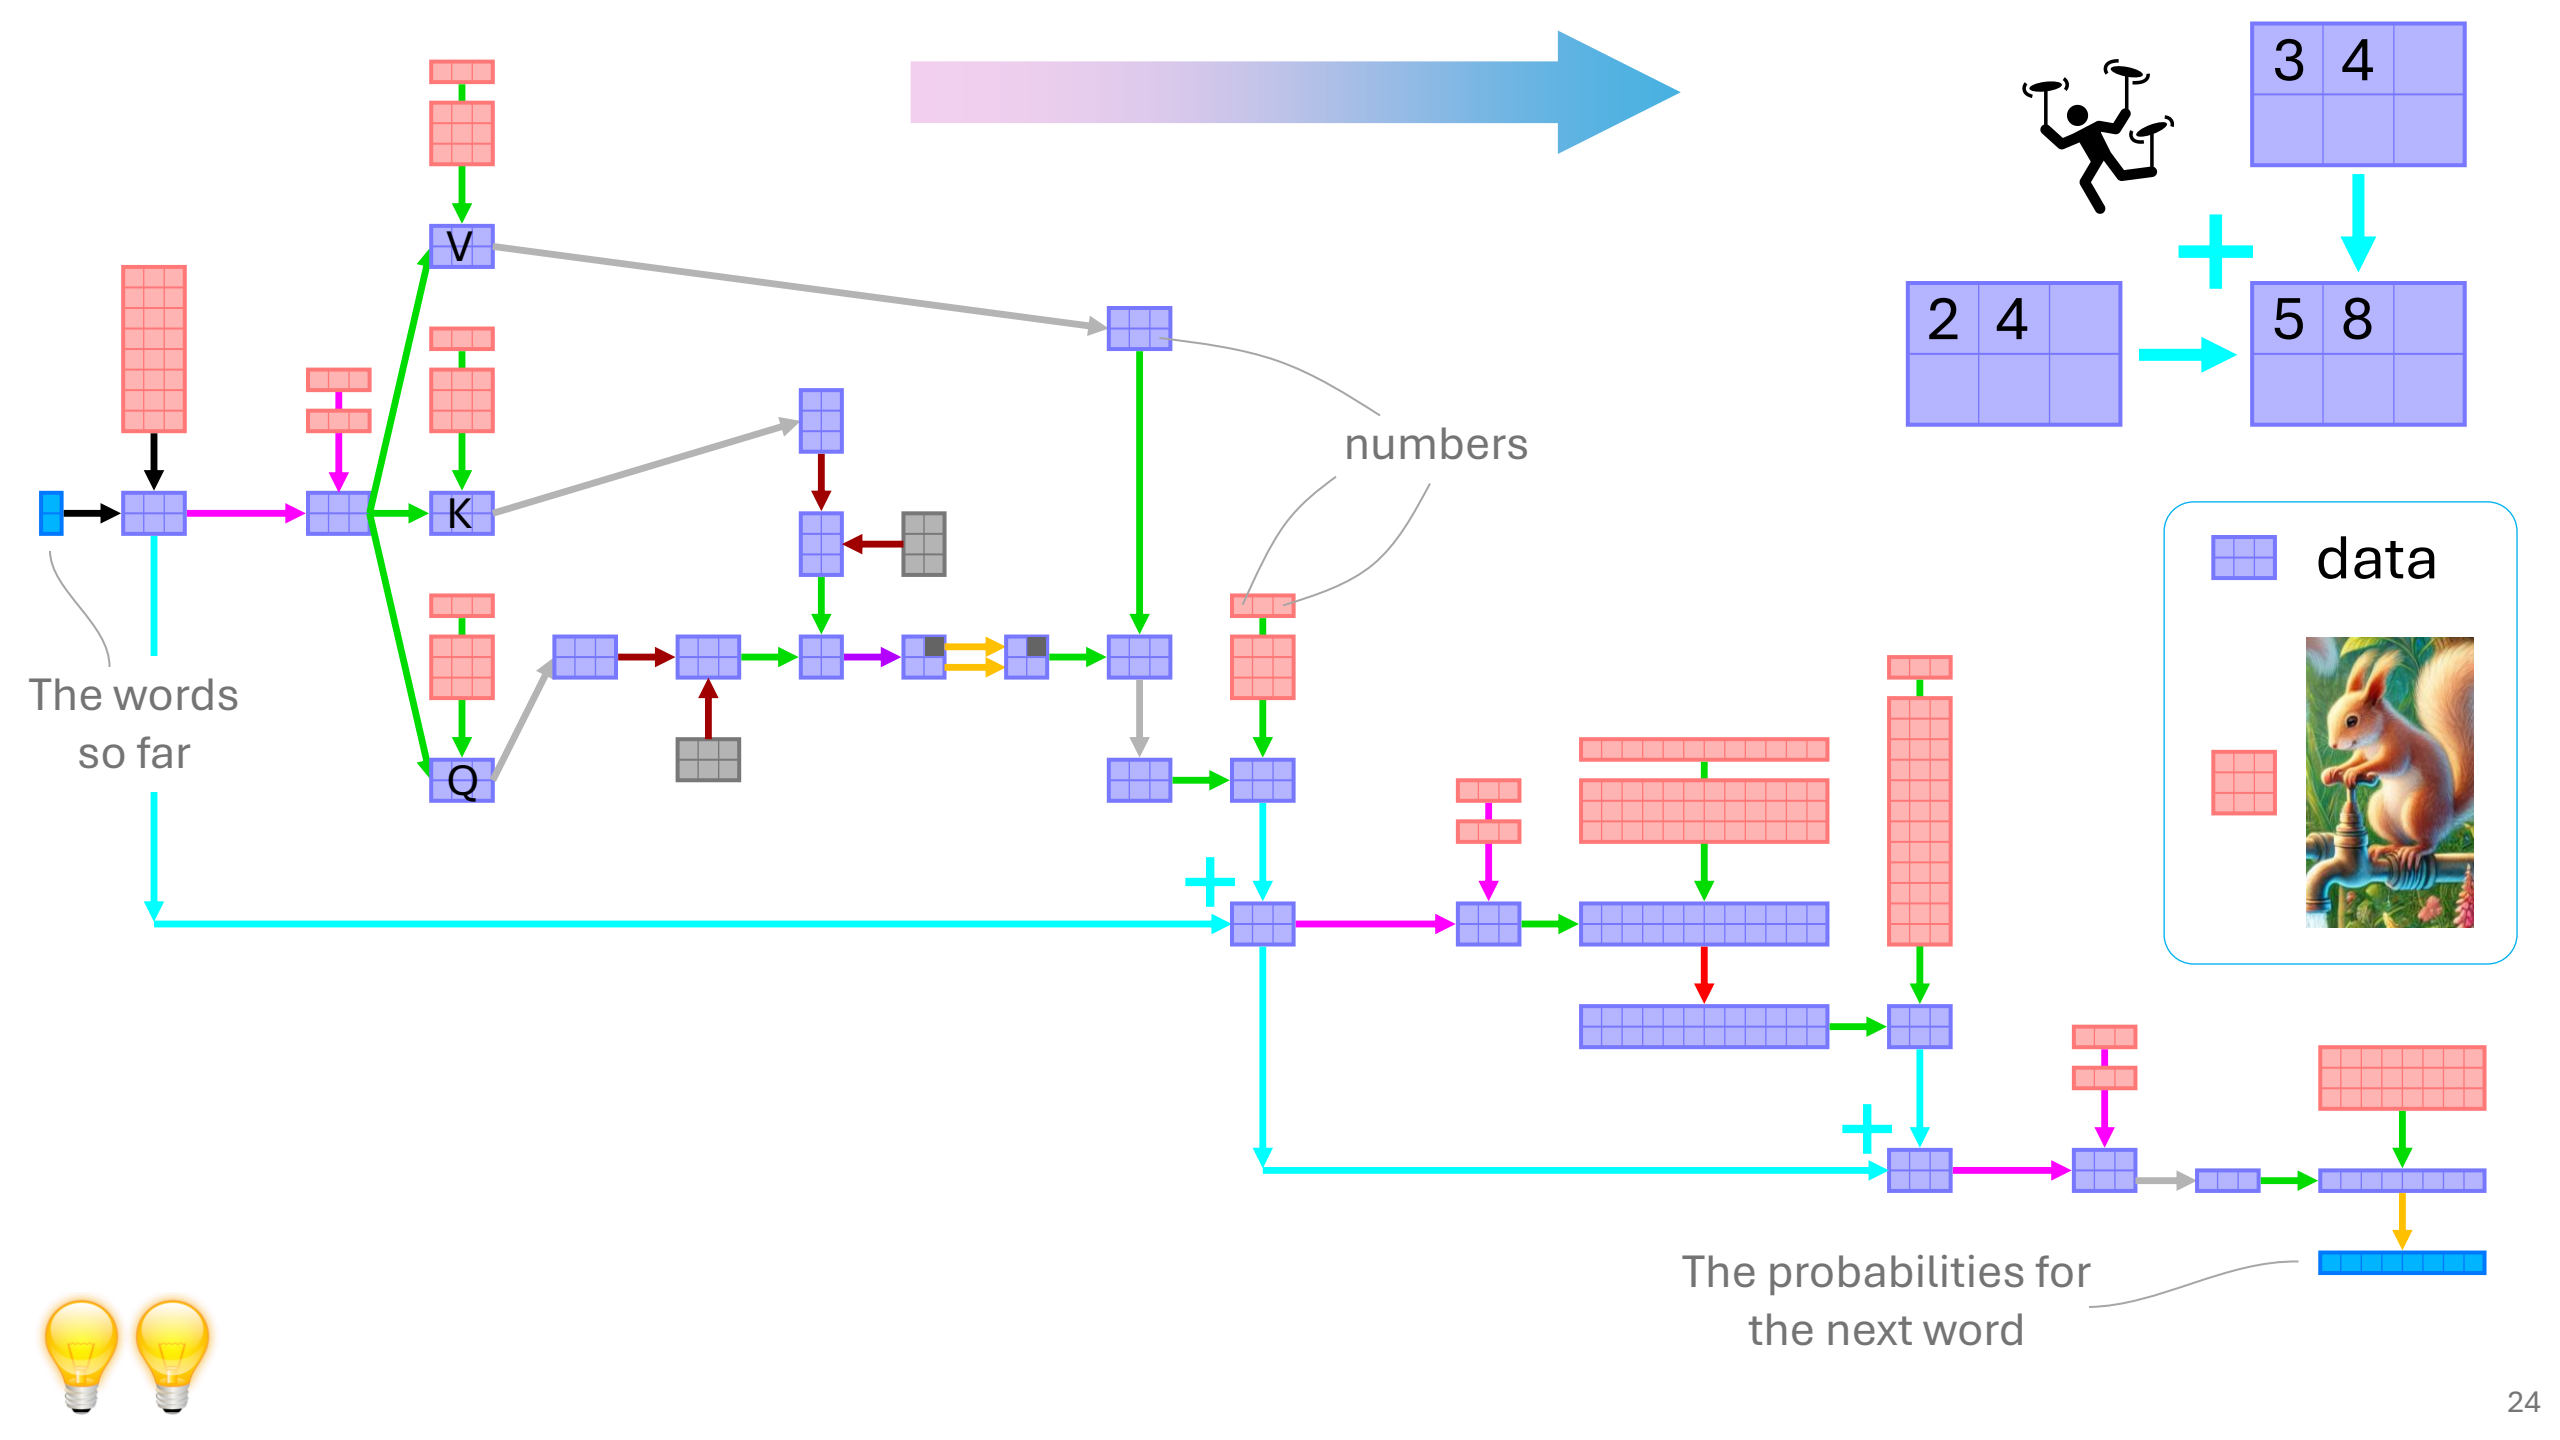
V
K
Q
numbers
The words
so far
3 4
2 4
5 8
data
The probabilities for
the next word

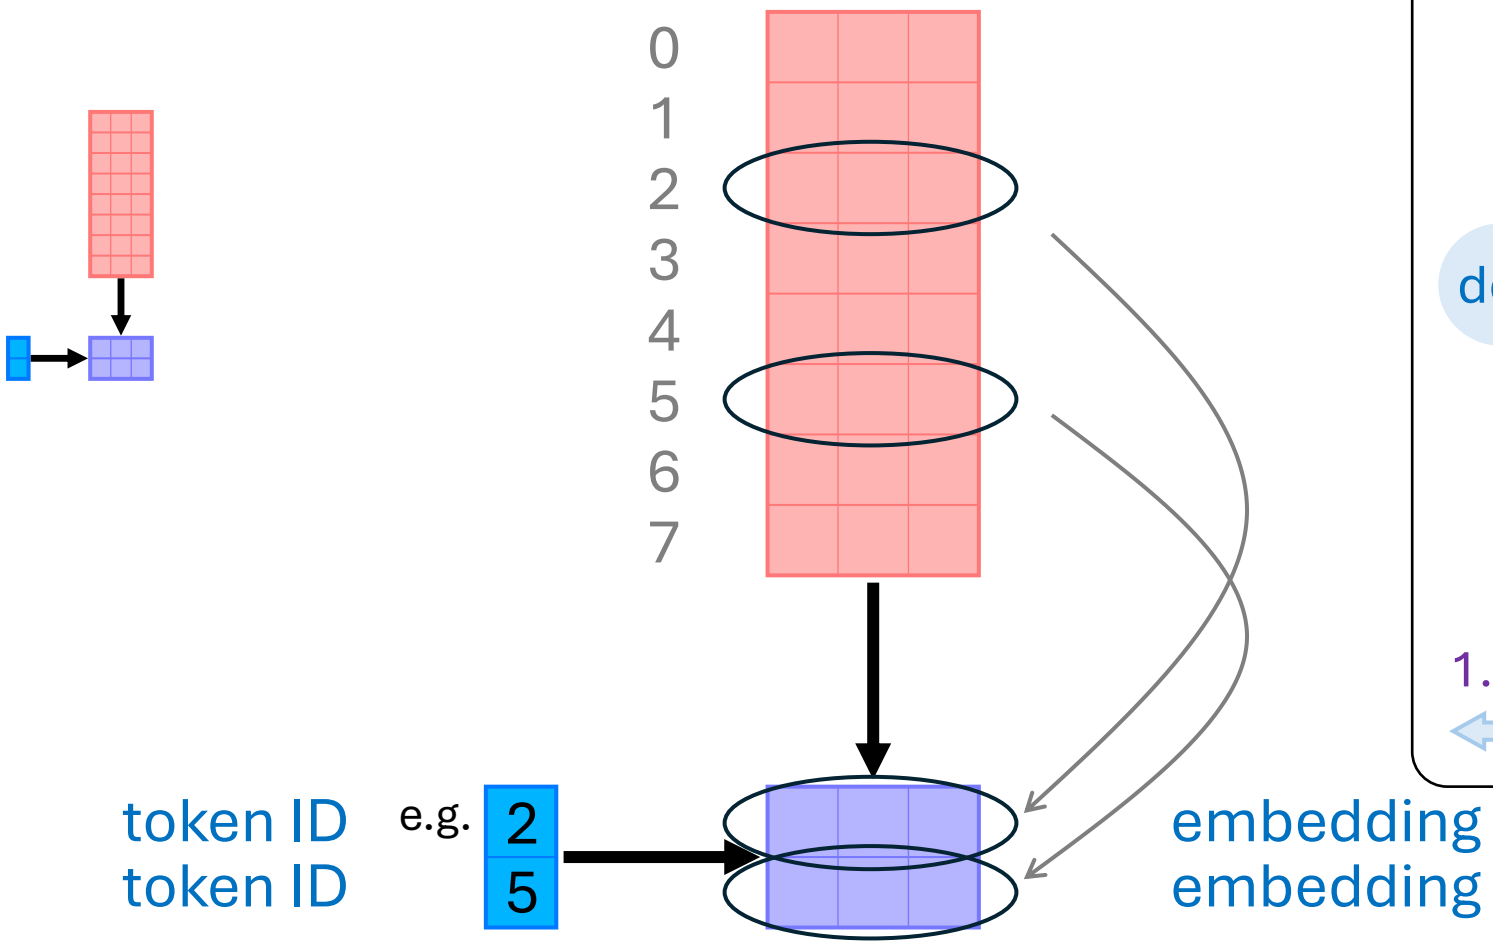

Word embeddings

dog
fox
bank
the
bus
set

1.2, 0.26, -0.3, 4.1, ...

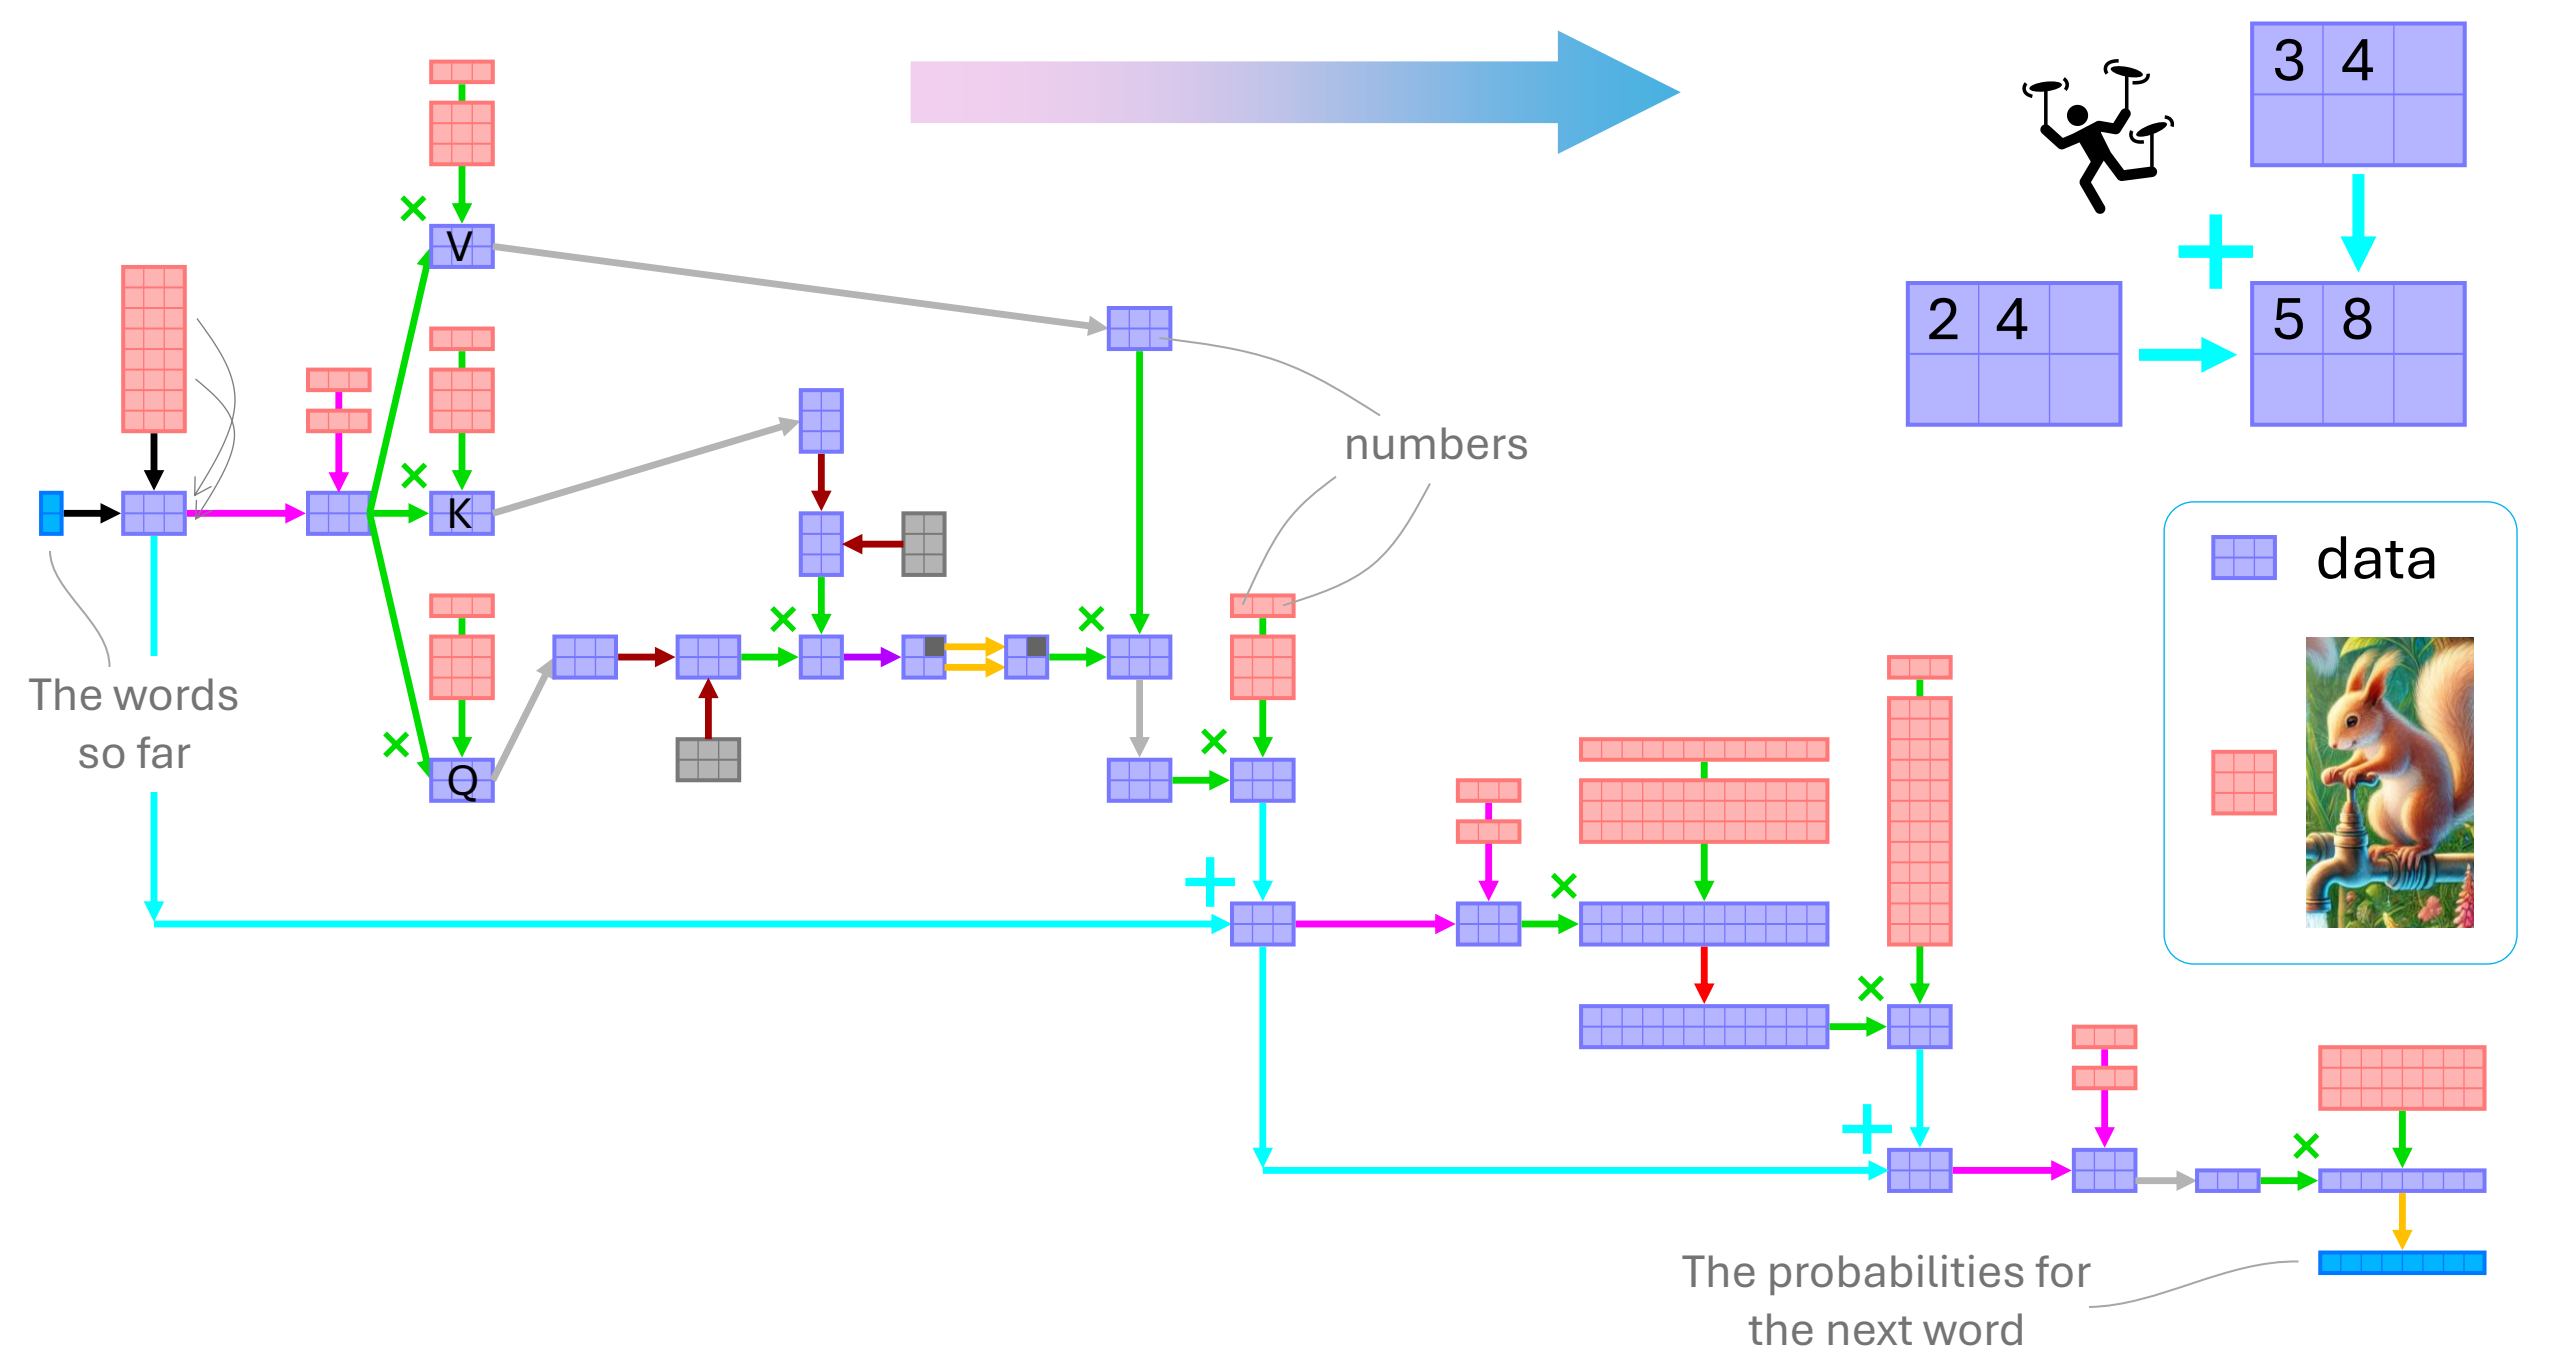
V
K
Q
numbers
The words
so far
3 4
2 4
5 8
data
The probabilities for
the next word

$$x_1 + x_2 + \dots + x_n$$

$$w_1 \times x_1 + w_2 \times x_2 + \dots + w_n \times x_n$$

Synapse

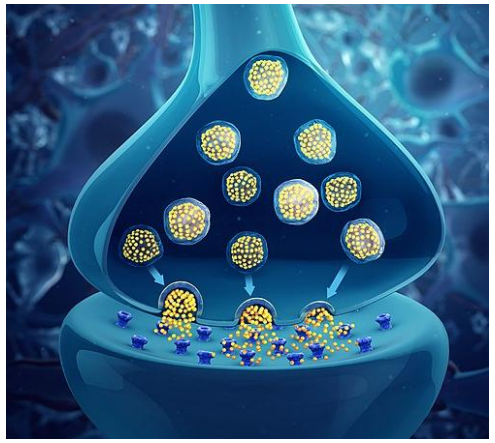

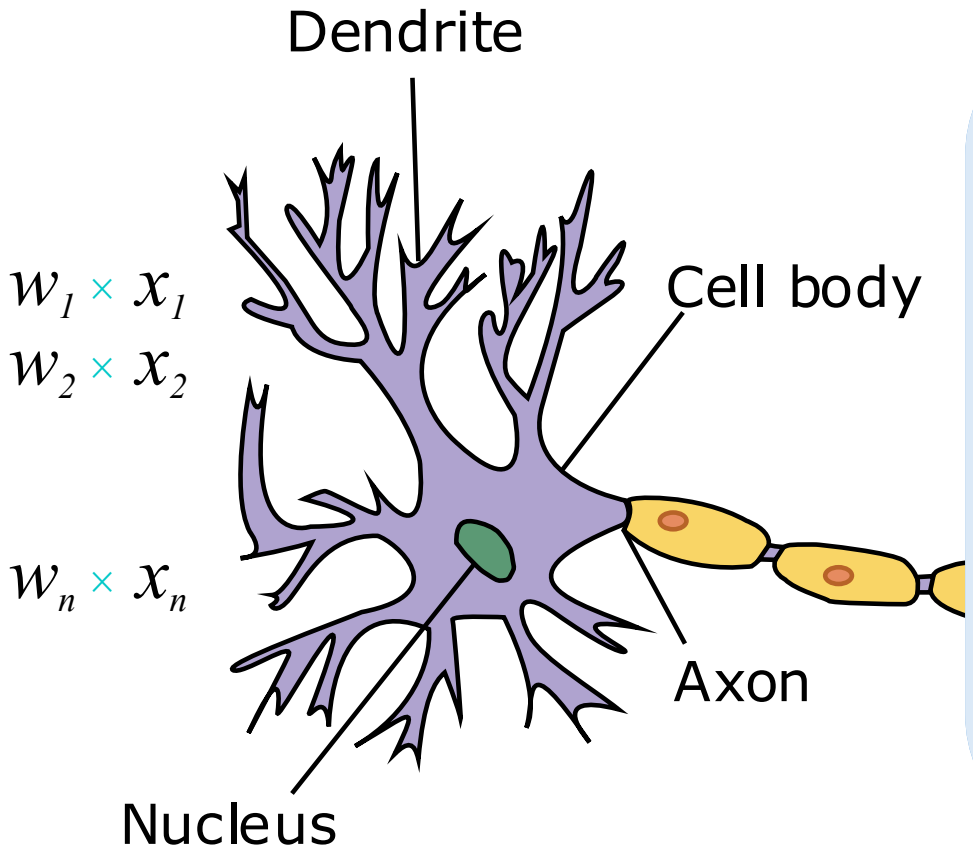

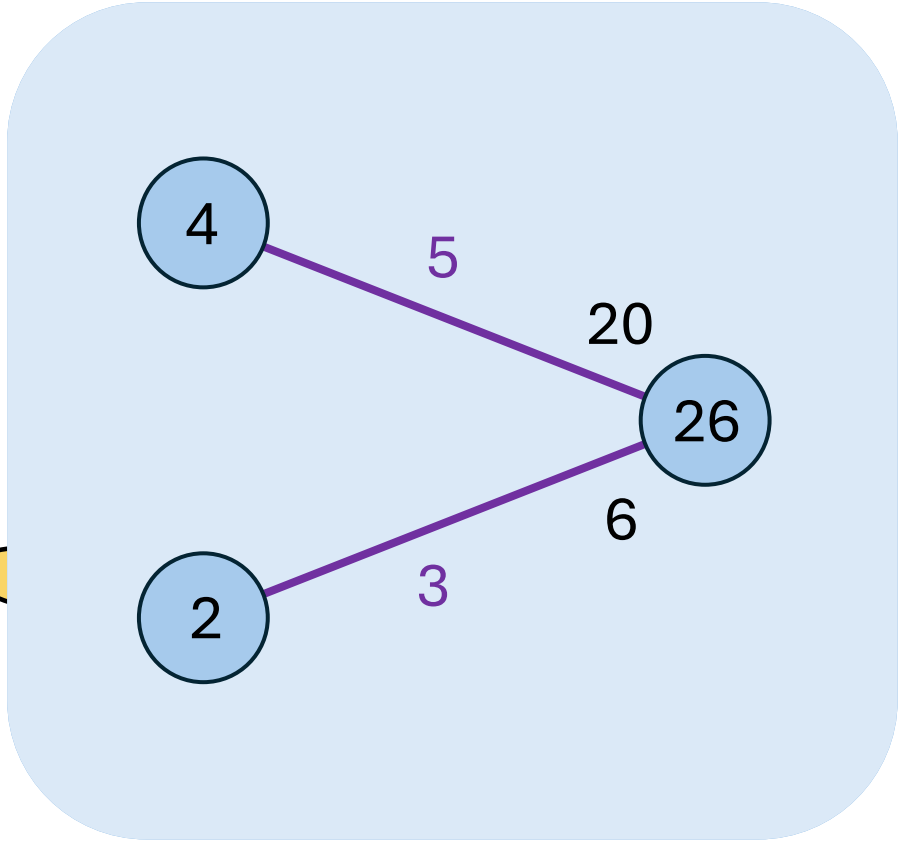

Synapse by www.scientificanimations.com / CC BY-SA 4.0
Neuron by Dhp1080, the US National Cancer Institute's (SEER) Program / CC BY-SA 3.0

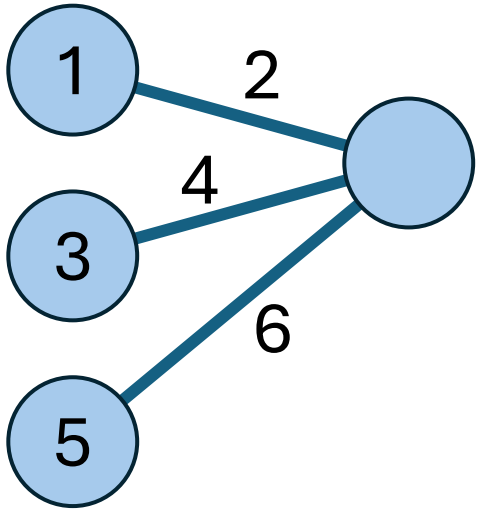

$$1 \times 2 + 3 \times 4 + 5 \times 6$$

$$\begin{bmatrix} 1 & 3 & 5 \end{bmatrix} \begin{bmatrix} 2 \\ 4 \\ 6 \end{bmatrix} \qquad \begin{bmatrix} \bullet & \bullet & \bullet \end{bmatrix} \begin{bmatrix} \bullet \\ \bullet \\ \bullet \end{bmatrix} = \begin{bmatrix} \bullet \end{bmatrix}$$



$1 \times 2 + 3 \times 4 + 5 \times 6$

$$\begin{bmatrix} 1 & 3 & 5 \end{bmatrix} \begin{bmatrix} 2 \\ 4 \\ 6 \end{bmatrix}$$

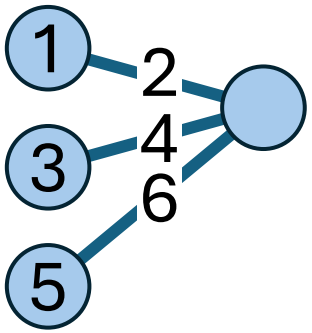

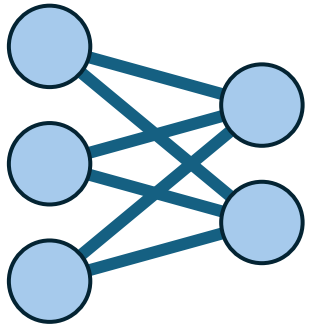

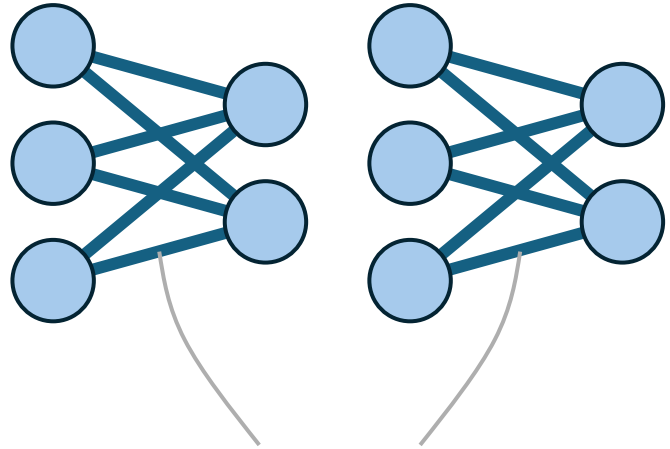

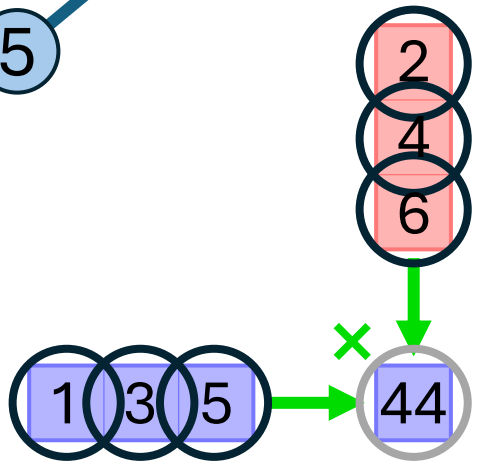

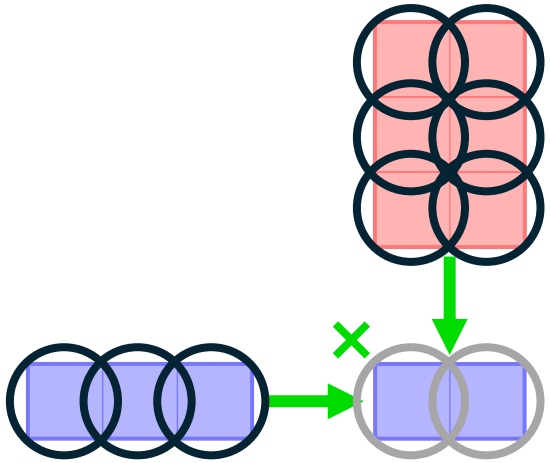

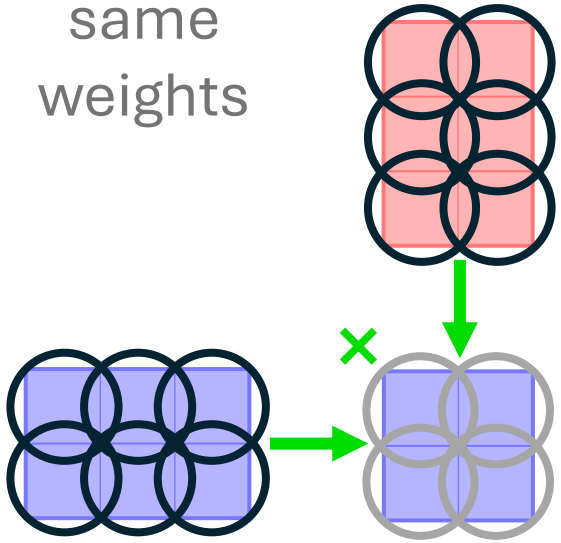

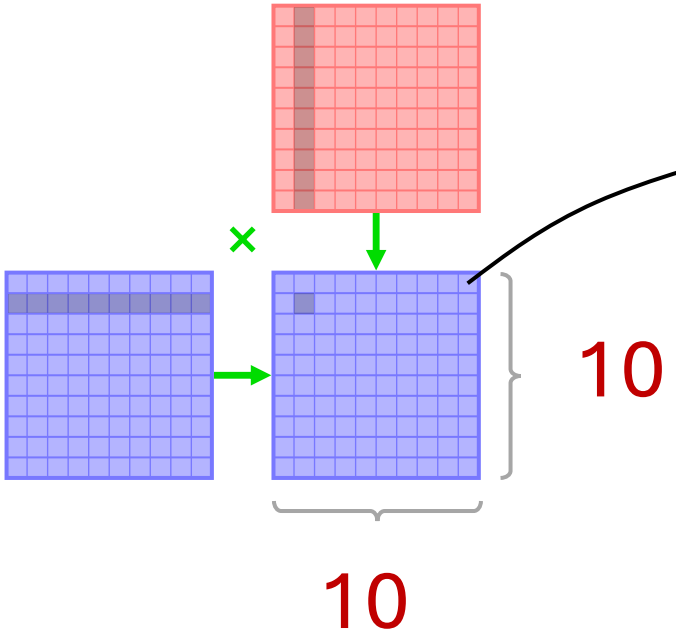

10 by 10

100 numbers

1000 multiplies

10,000 by 10,000

100,000,000

1,000,000,000,000

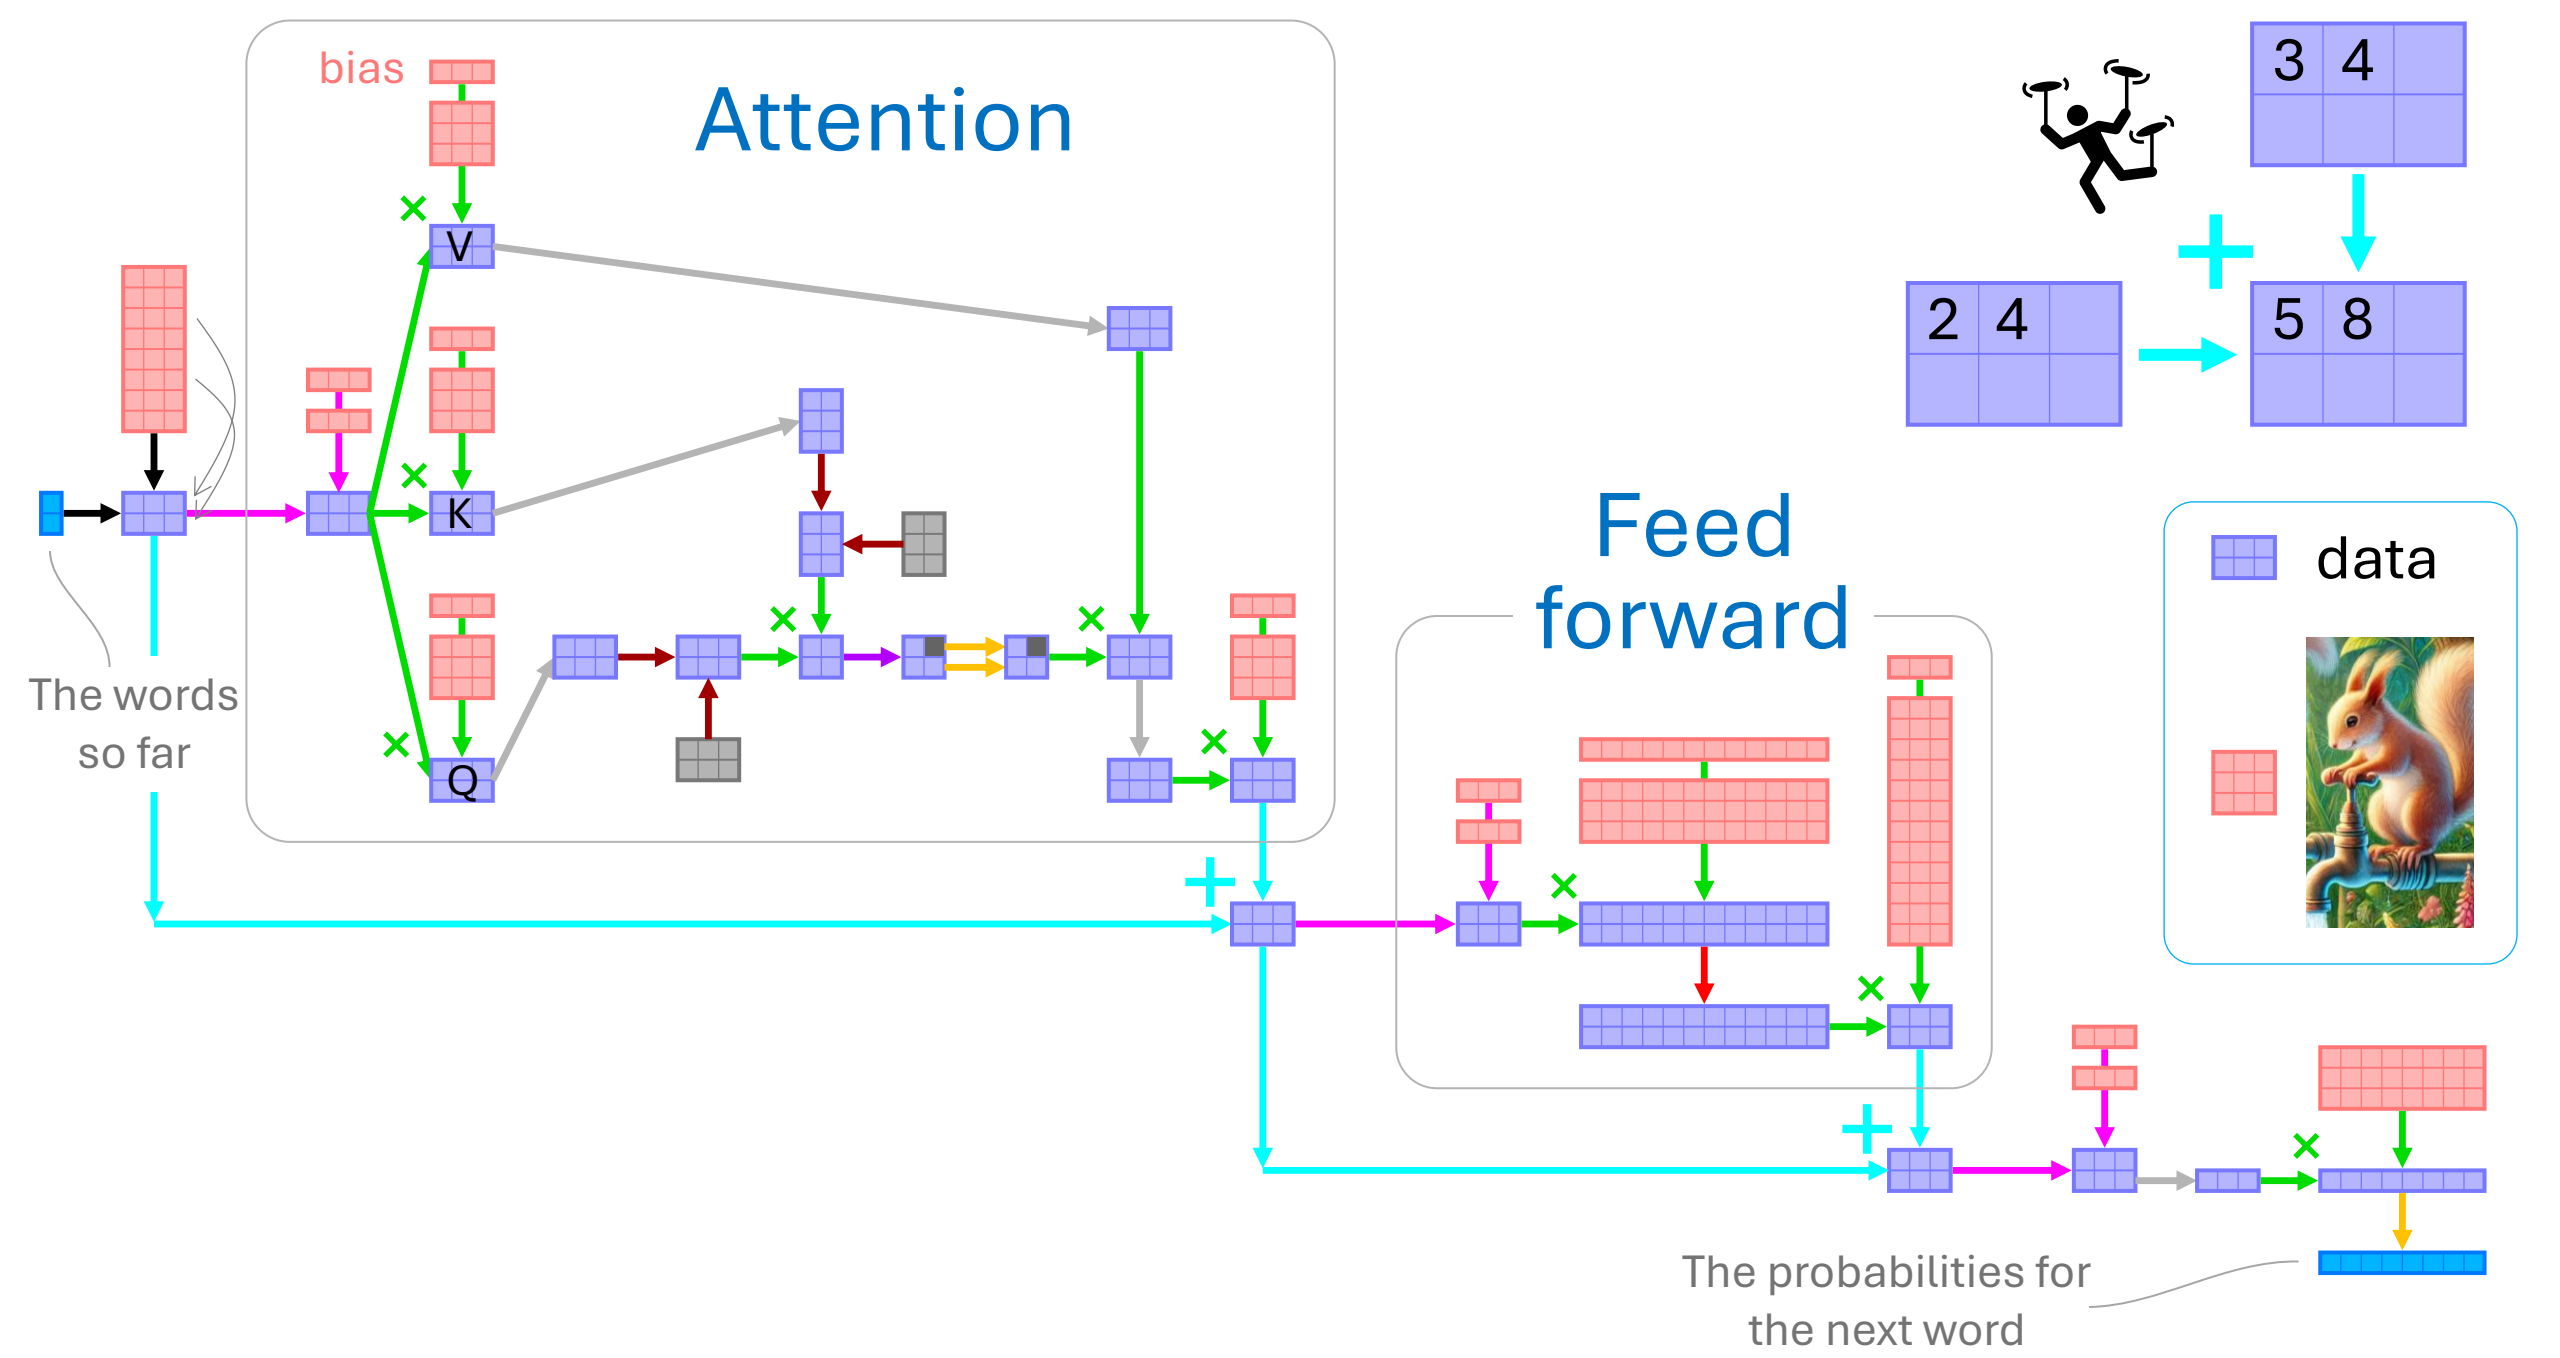
bias
Attention
V
K
Q
The words so far
Feed forward
3 4
2 4
5 8
data
The probabilities for the next word

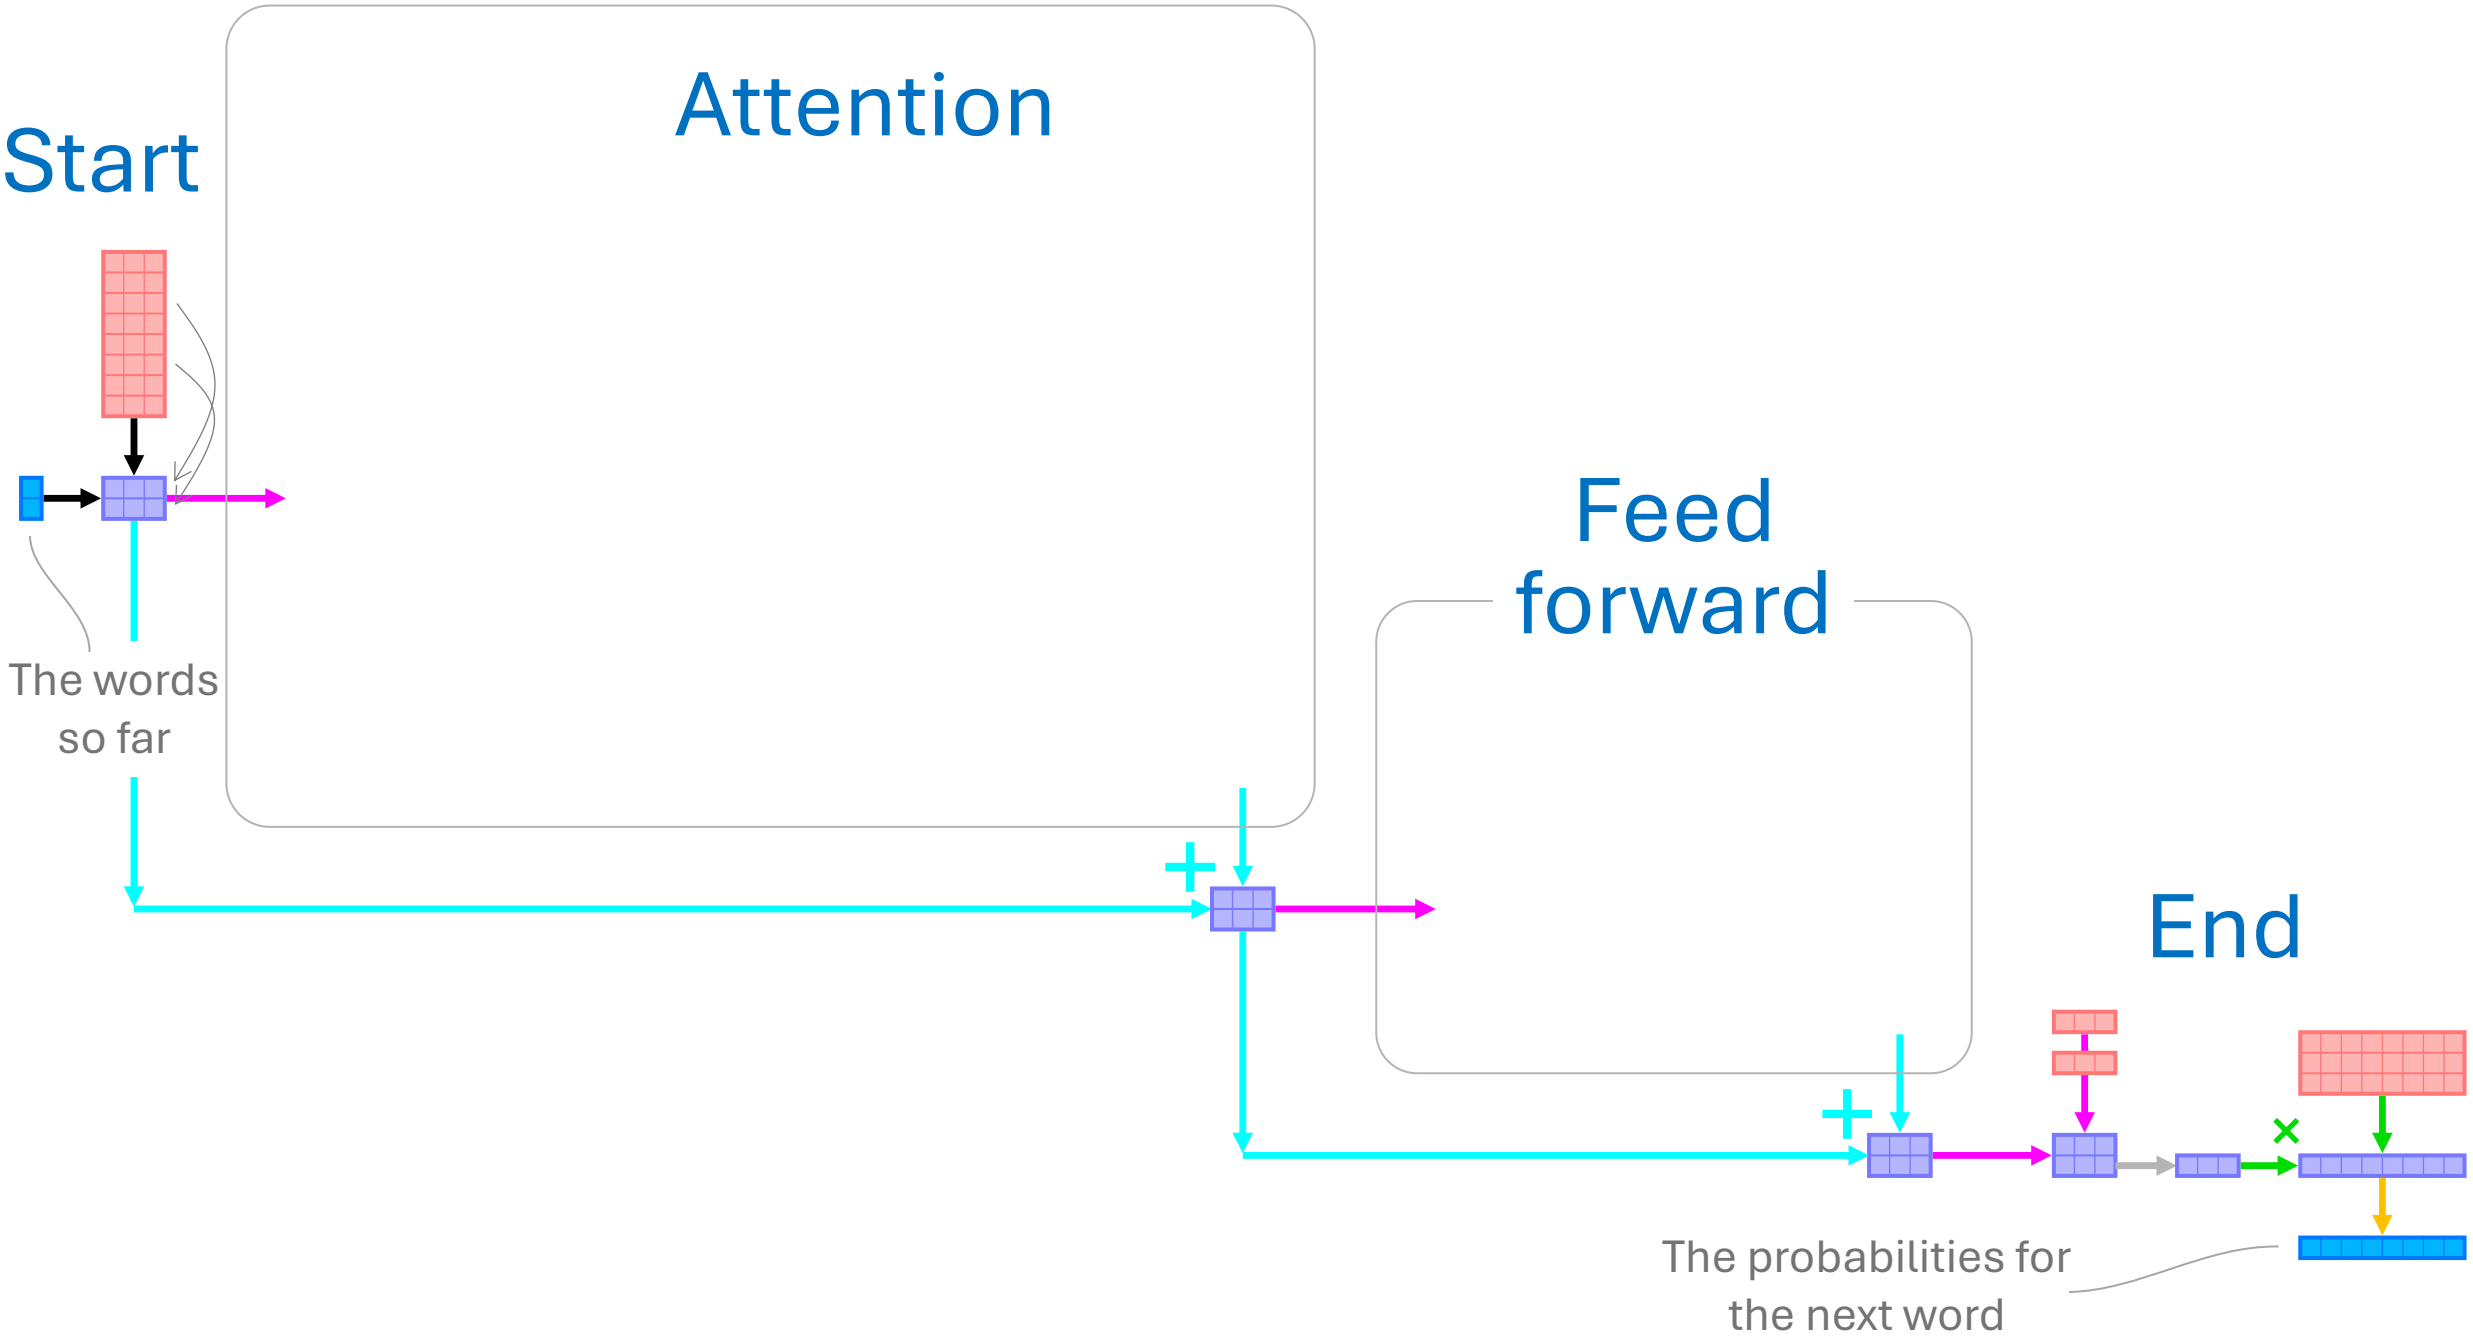
Start
Attention
Feed forward
End
The words so far
The probabilities for the next word

Start ⇨ Attention ⇨ Ffwd ⇨ Attention ⇨ Ffwd ⇨ Attention ⇨ Ffwd ⇨ End

Block Block Block

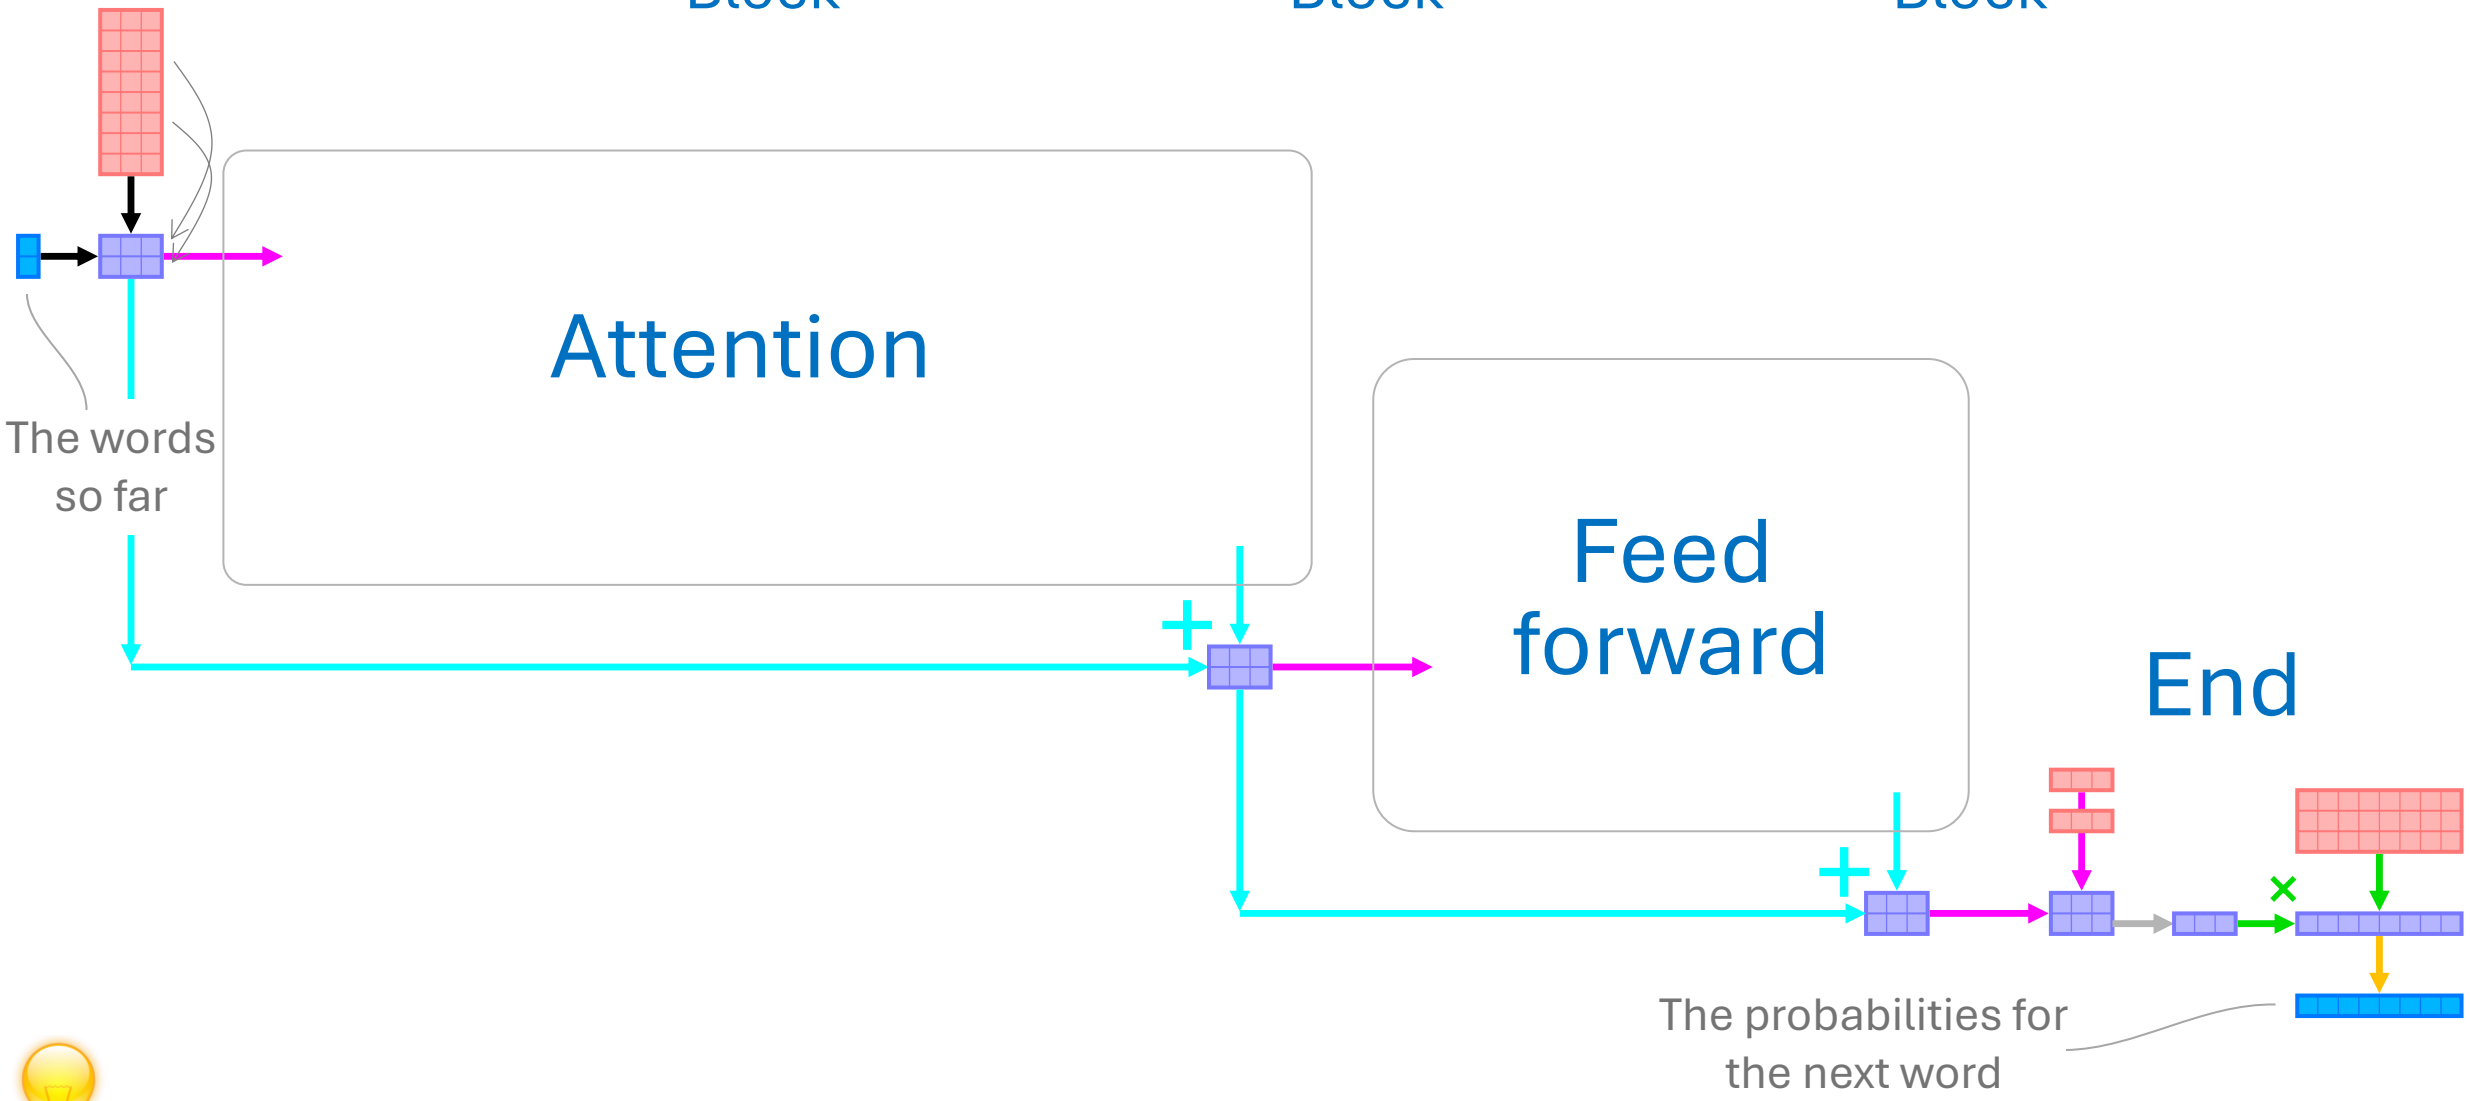

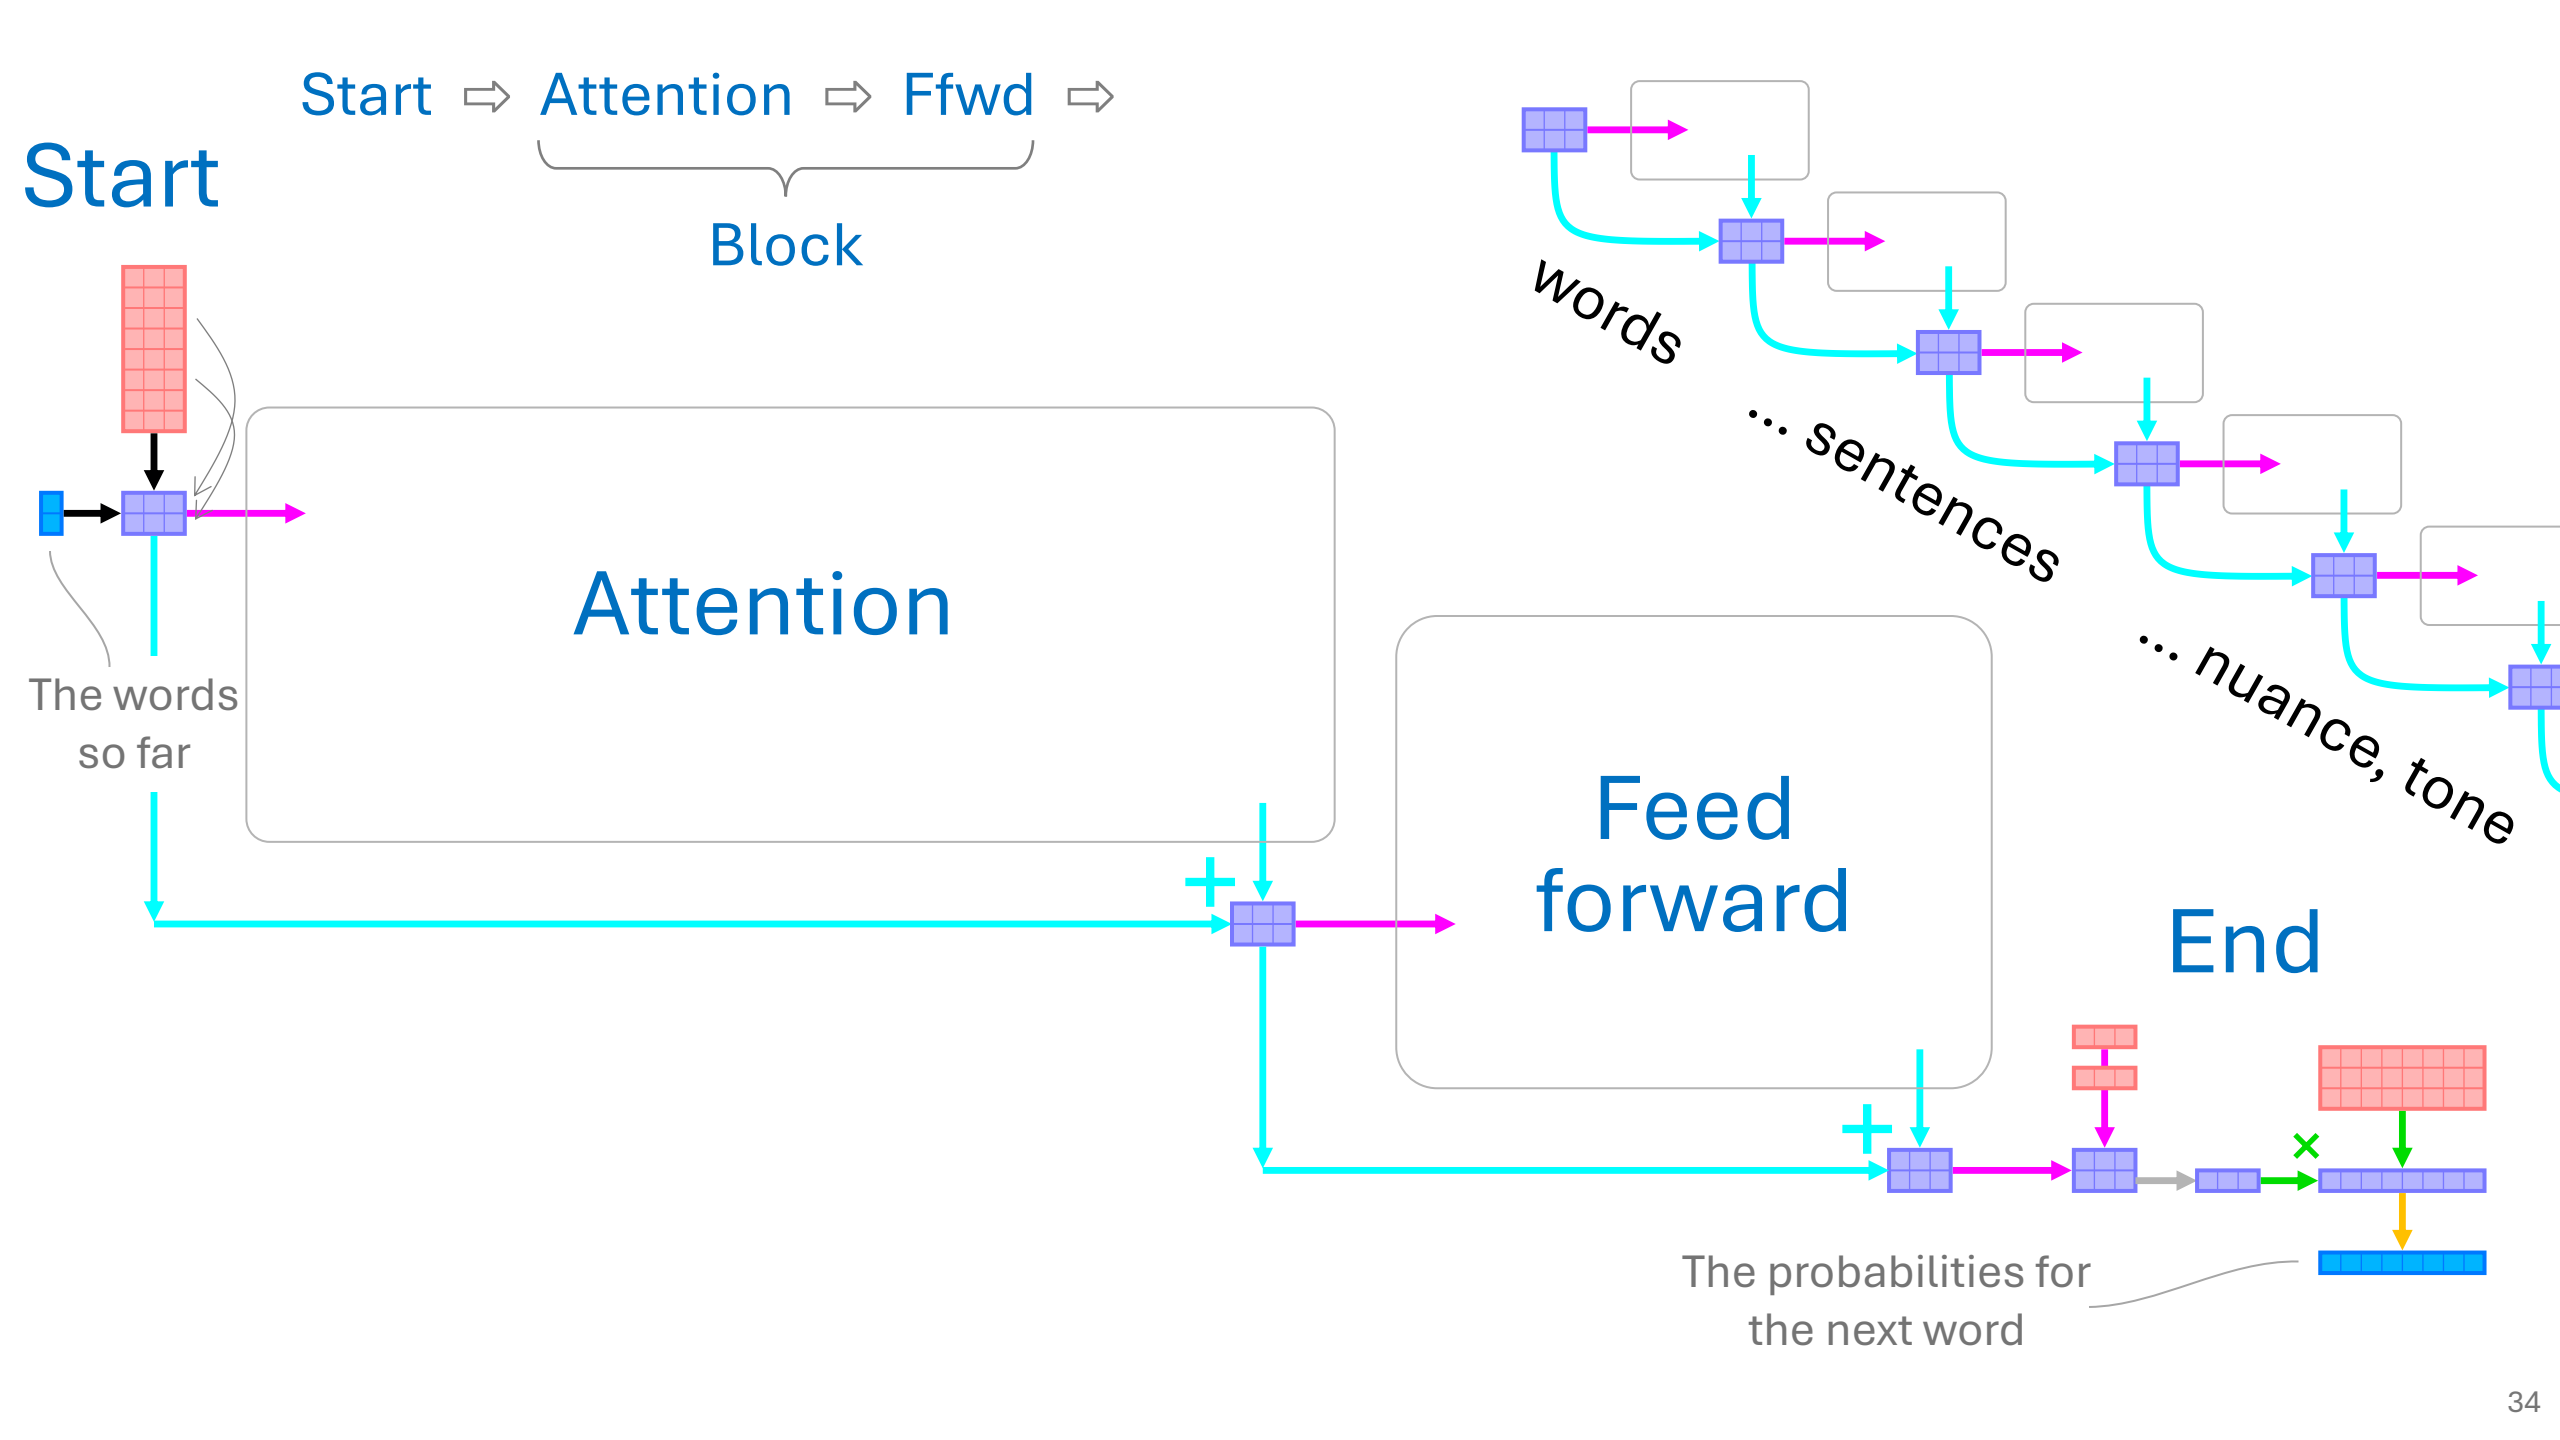
Start ⇨ Attention ⇨ Ffwd ⇨
Start
Block
Attention
The words so far
words
… sentences
… nuance, tone
Feed forward
End
The probabilities for the next word

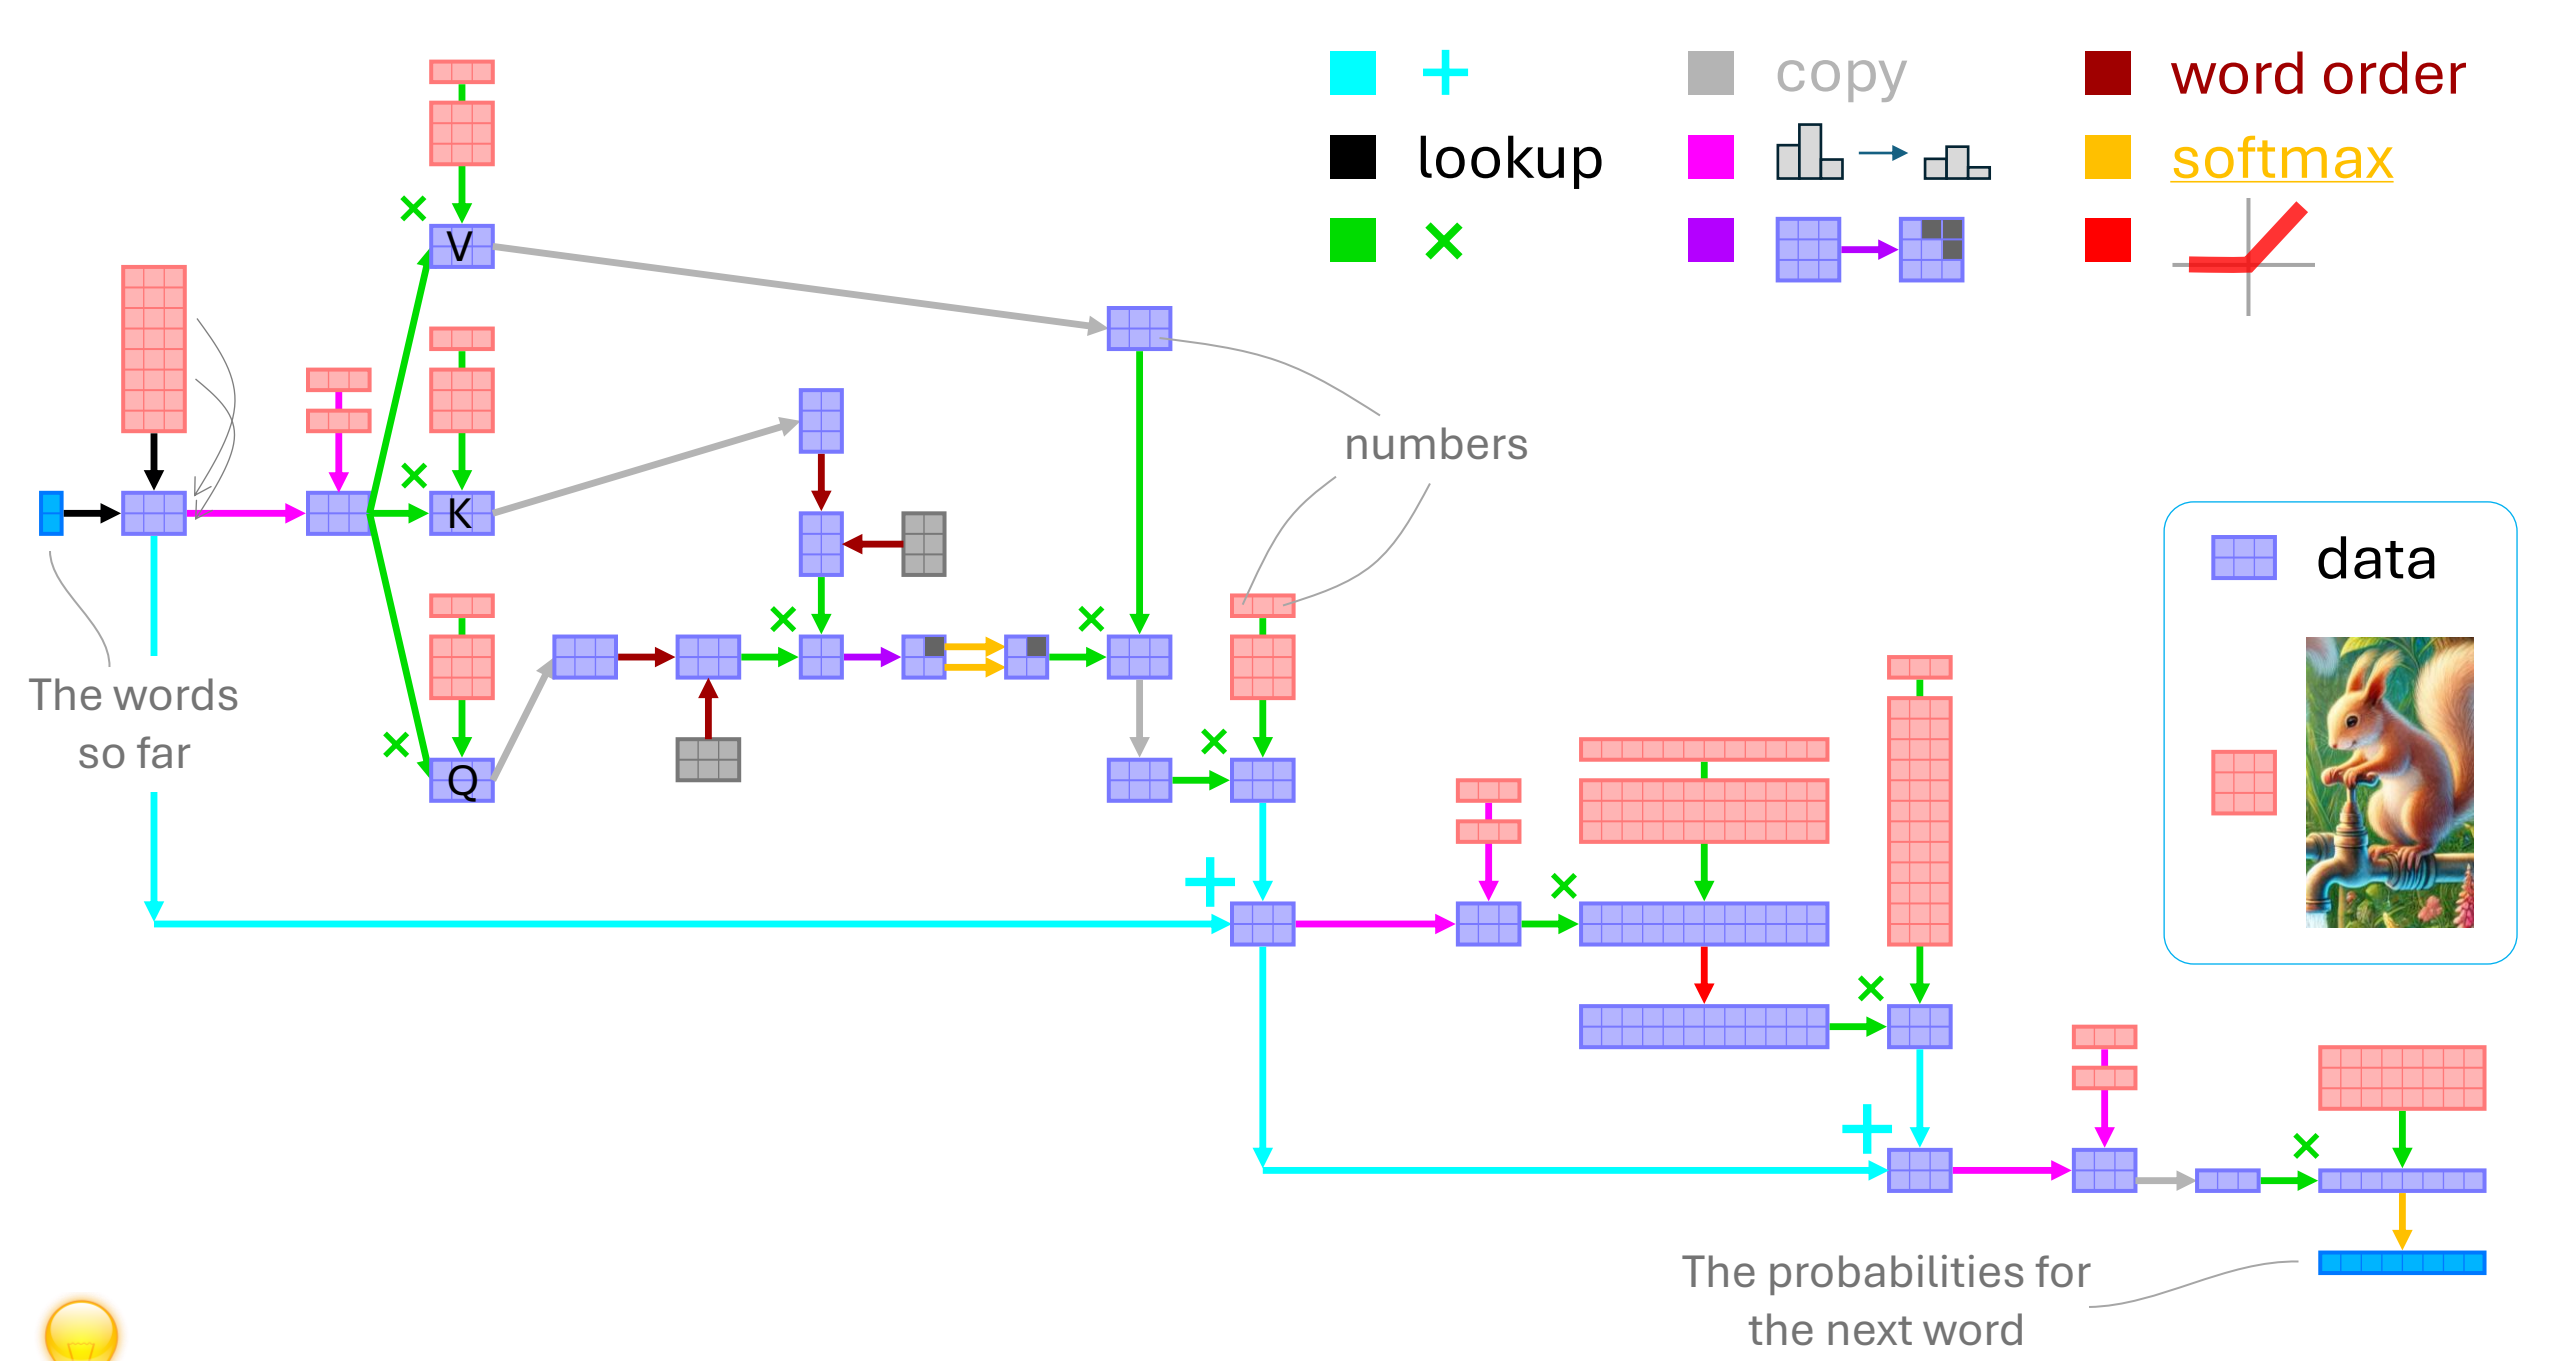
+
copy
word order
lookup
softmax
×
V
K
Q
numbers
The words so far
data
The probabilities for the next word

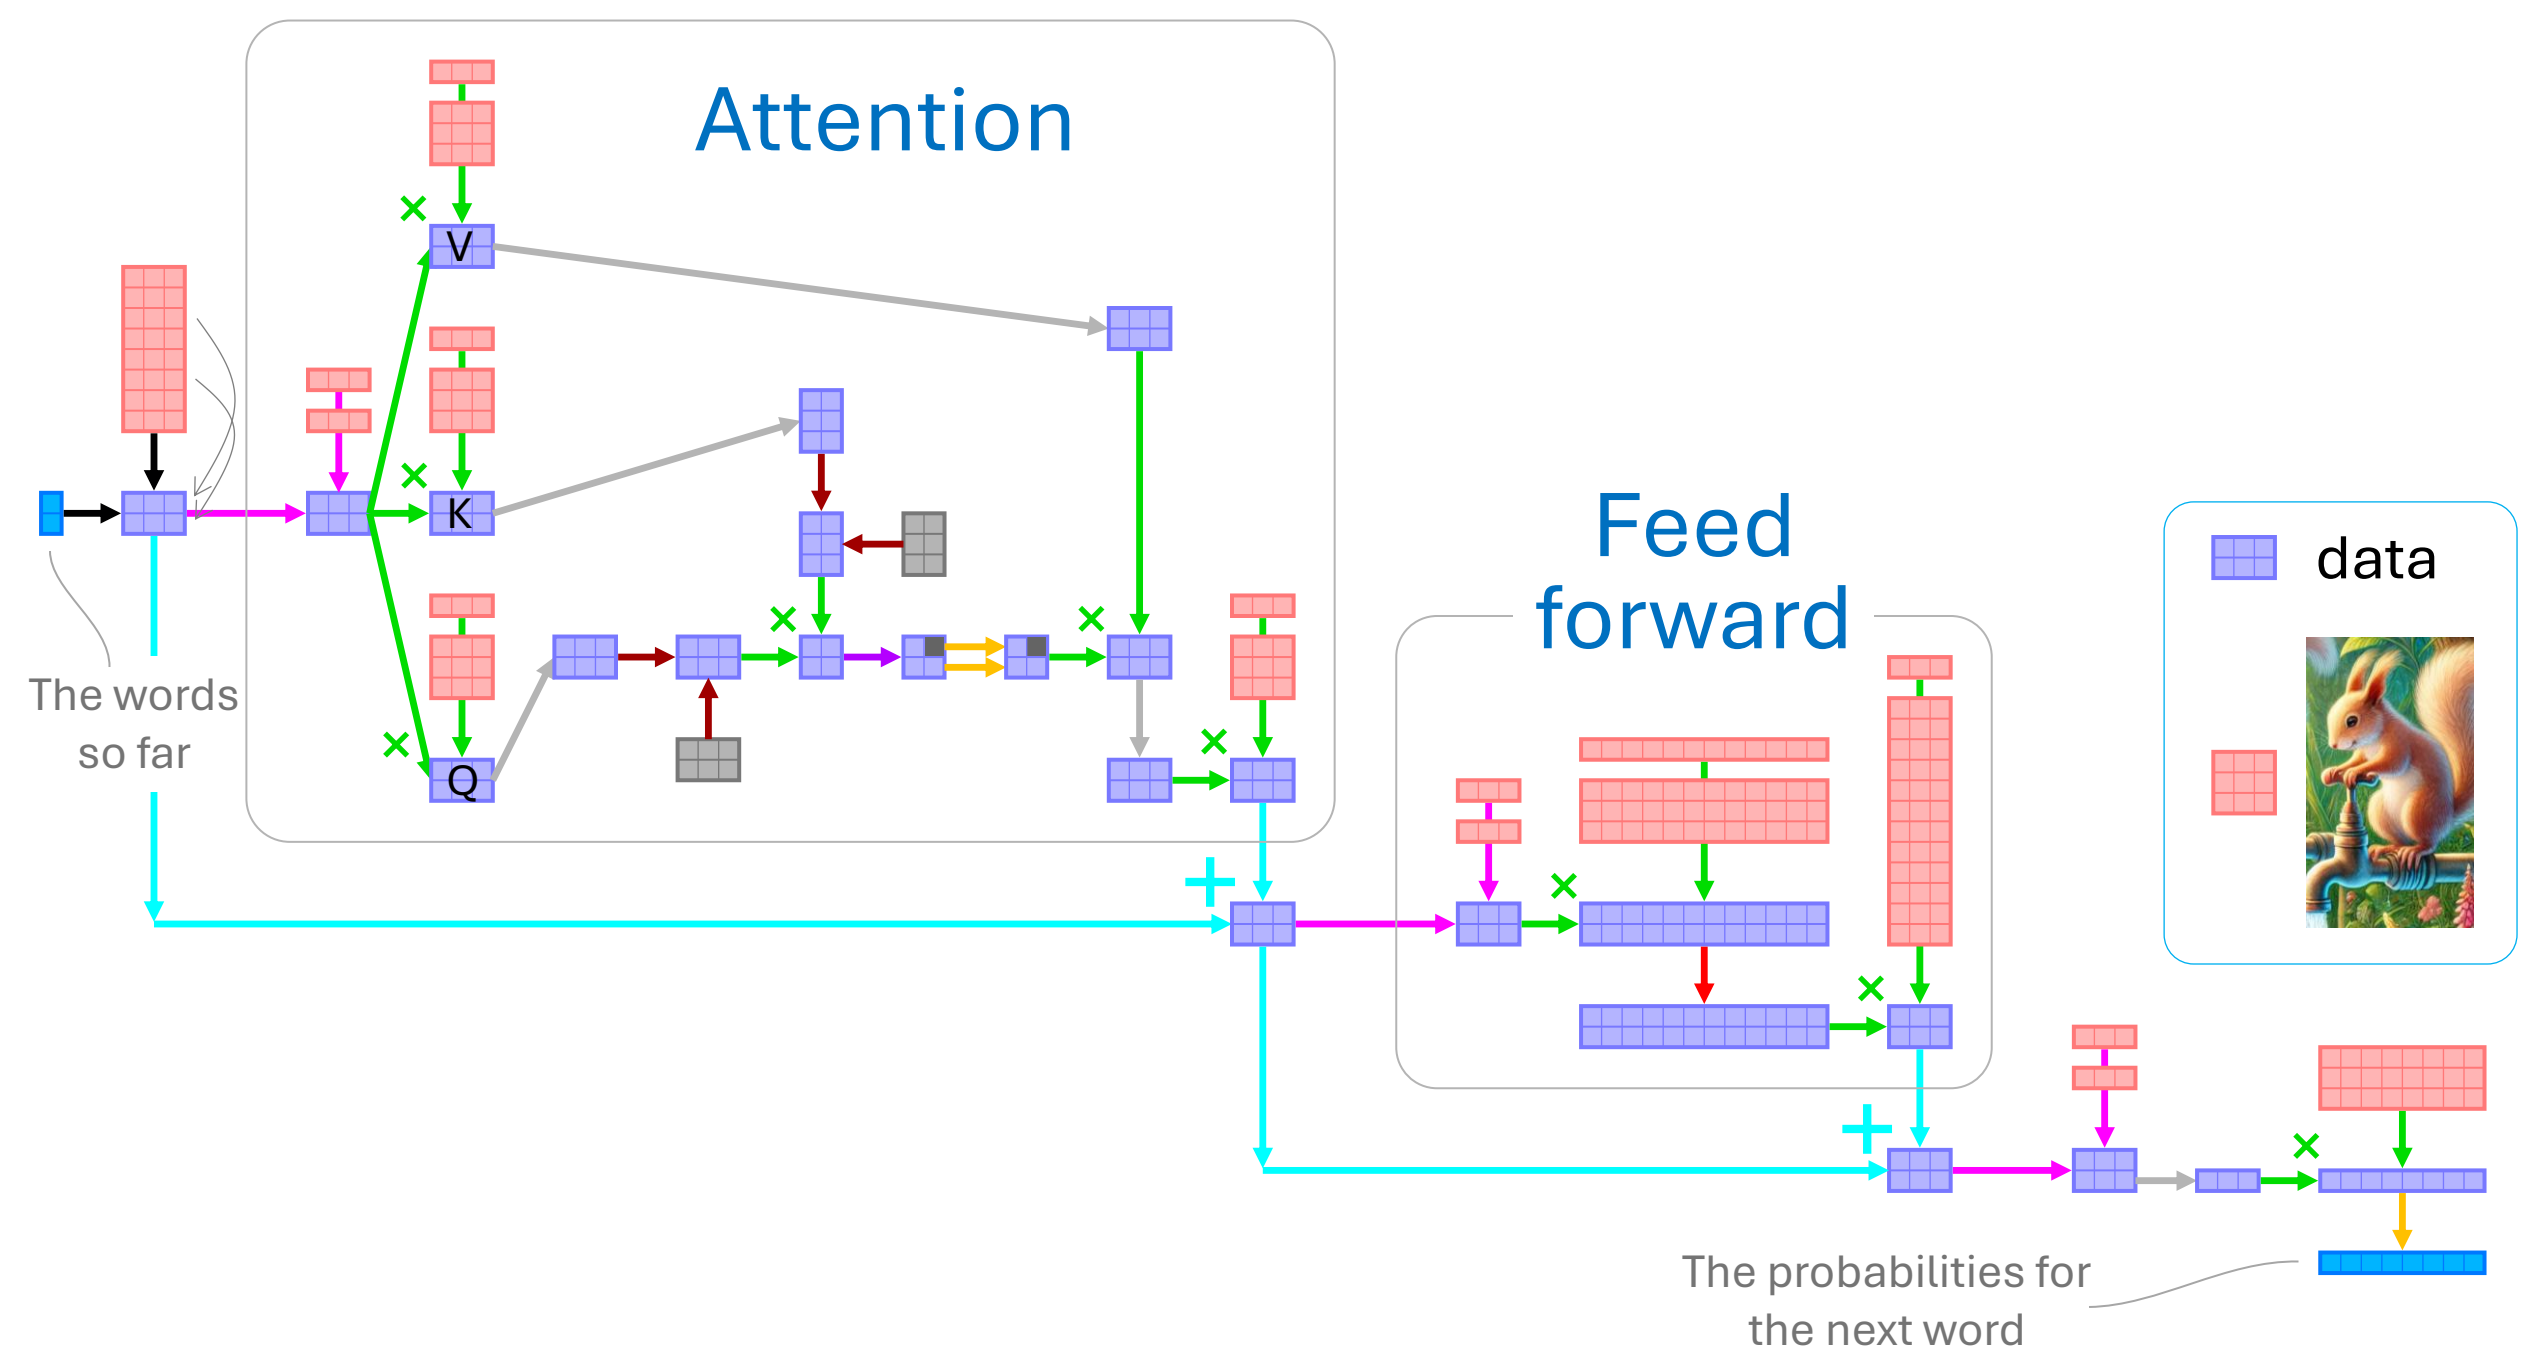
Attention
V
K
Q
The words so far
Feed forward
data
The probabilities for the next word

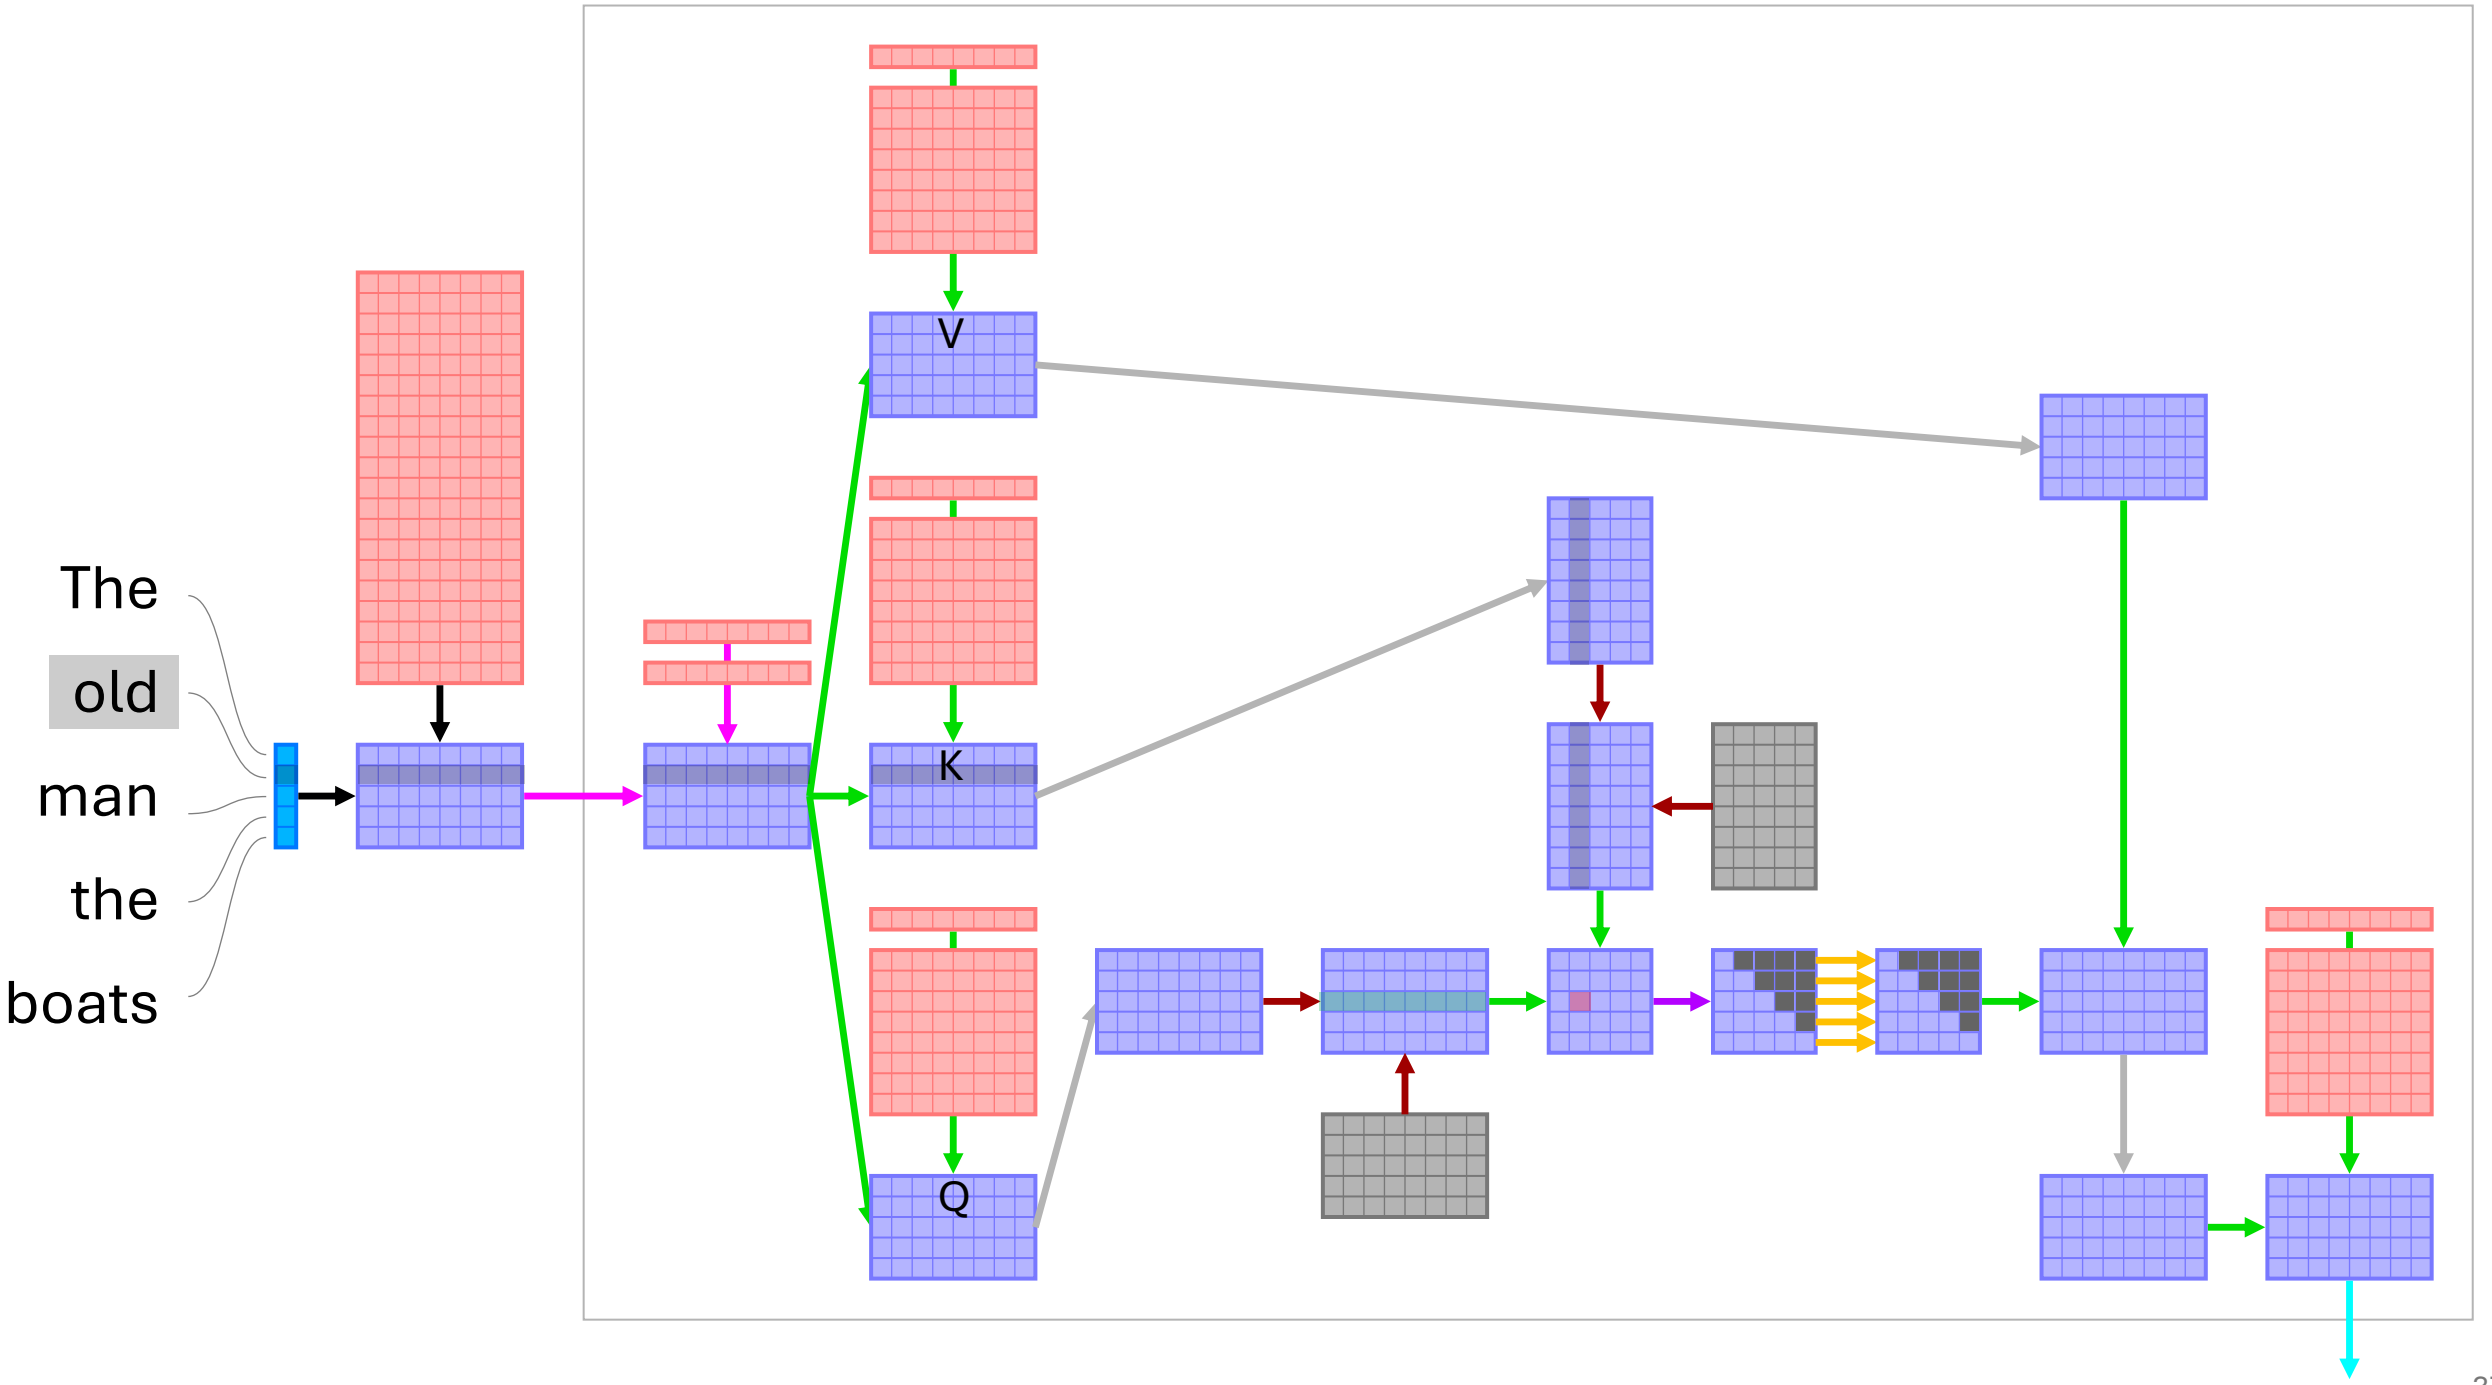
The
old
man
the
boats
V
K
Q

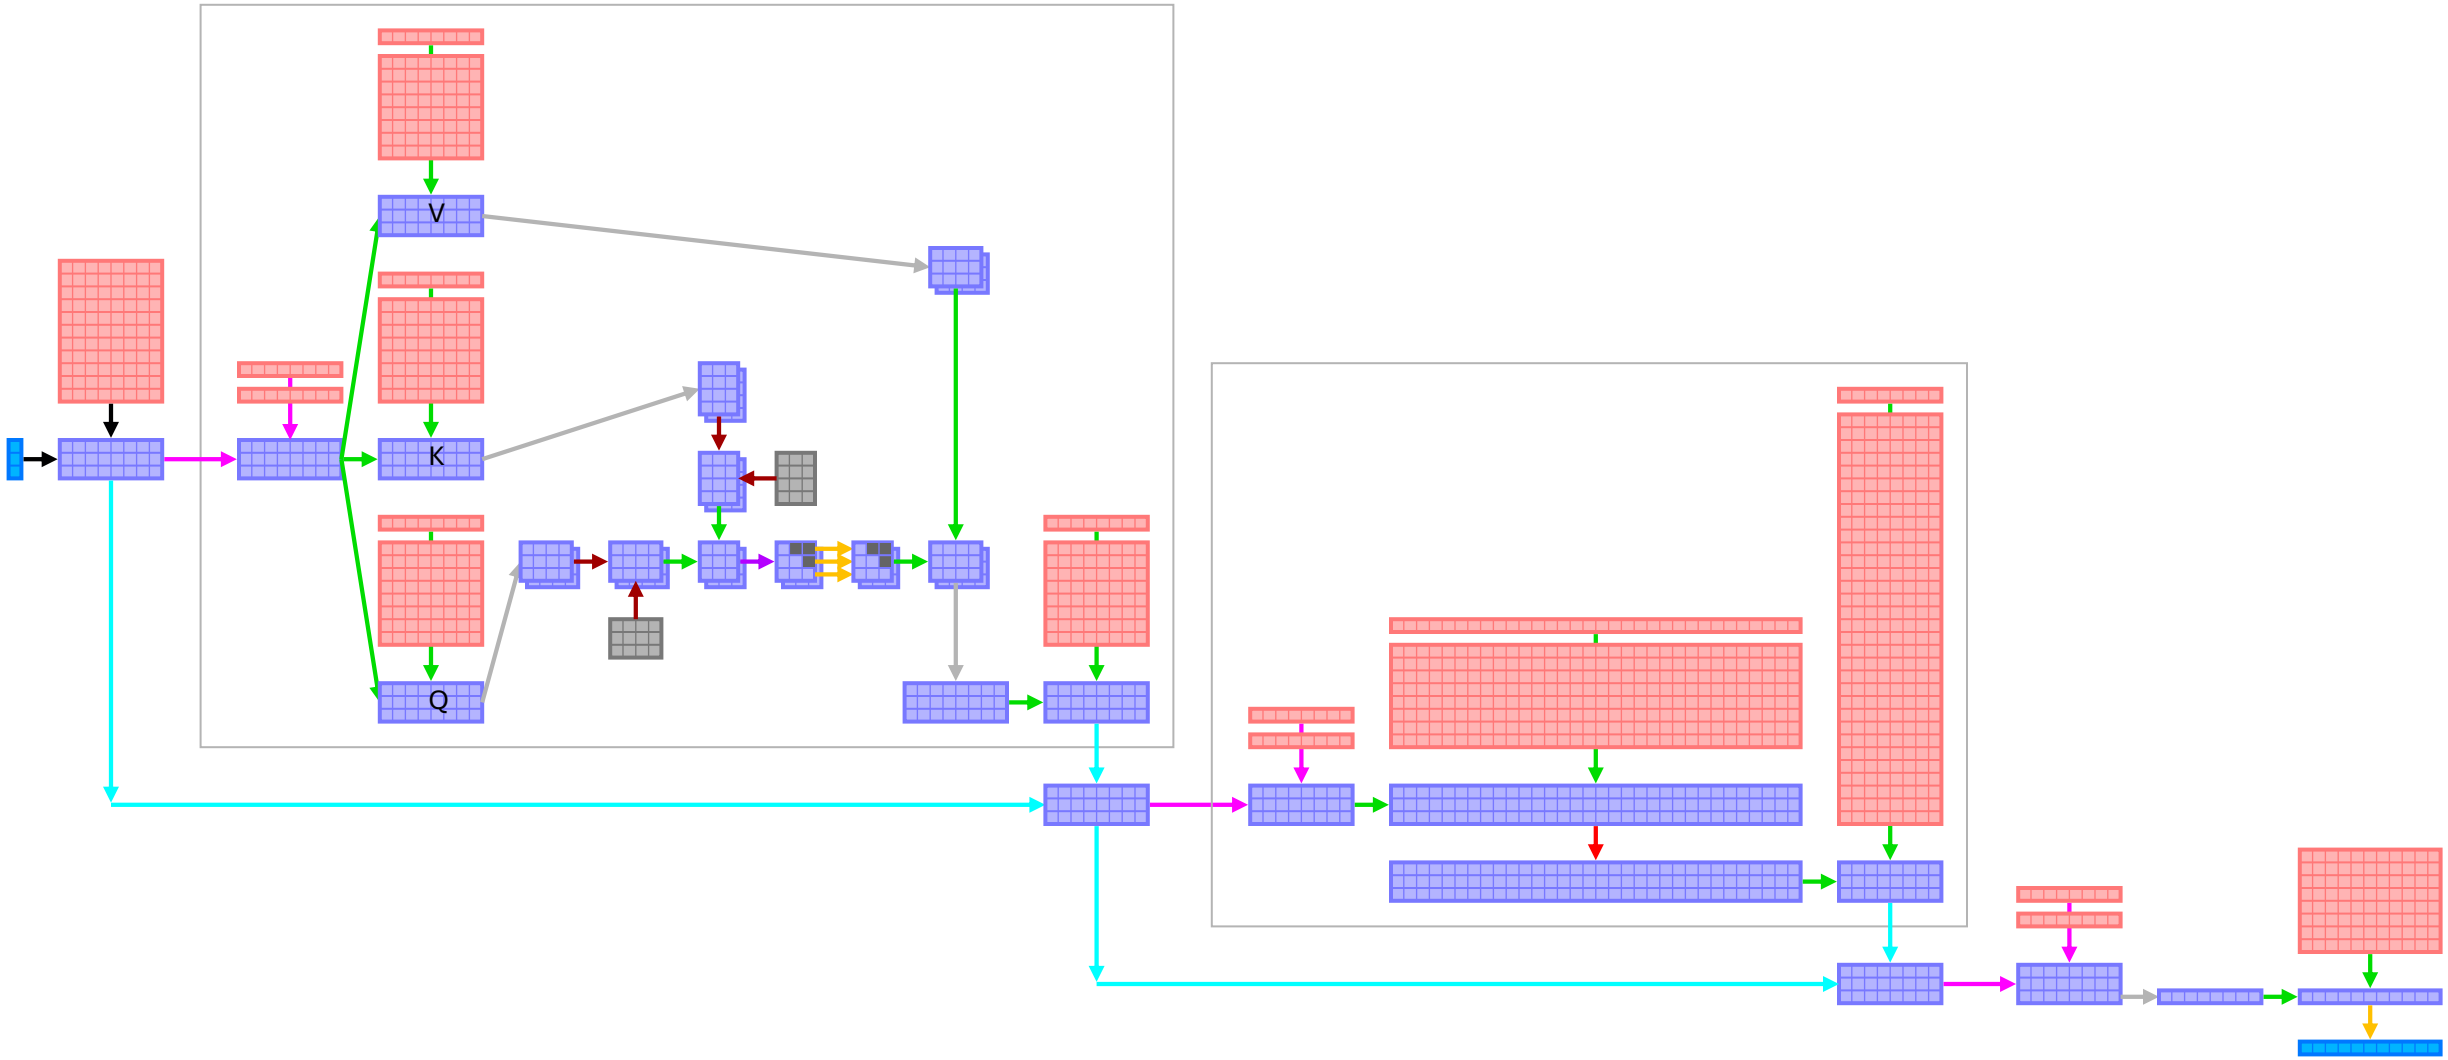
V
K
Q

Not needed for LLMs

Encoder

Decoder

https://arxiv.org/pdf/1706.03762

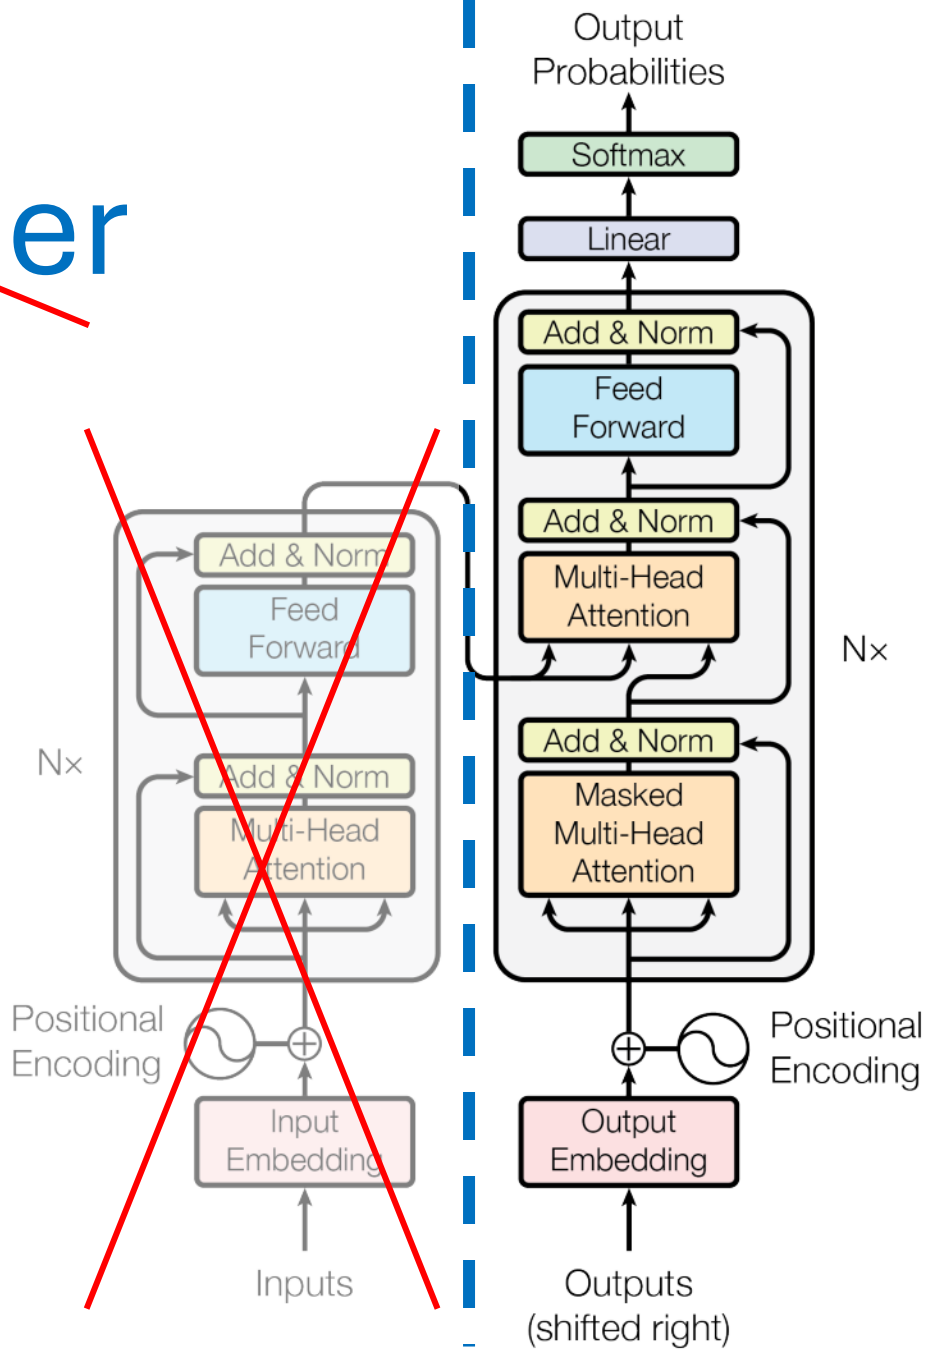

Figure 1: The Transformer - model architecture.

GPT

Generative pre-trained transformer

web



pre
training

Foundation
model

(base model)

fine
tuning

Fine-tuned
model

Foundation model

Why is the sky blue?



What makes a cloud? How is a hurricane born? Why do we have seasons? Answers to these …



https://huggingface.co/spaces/HuggingFaceFW/blogpost-fineweb-v1

🍷 **FineWeb**, a new, large-scale (**15-trillion tokens, 44TB disk space**) dataset for LLM pretraining

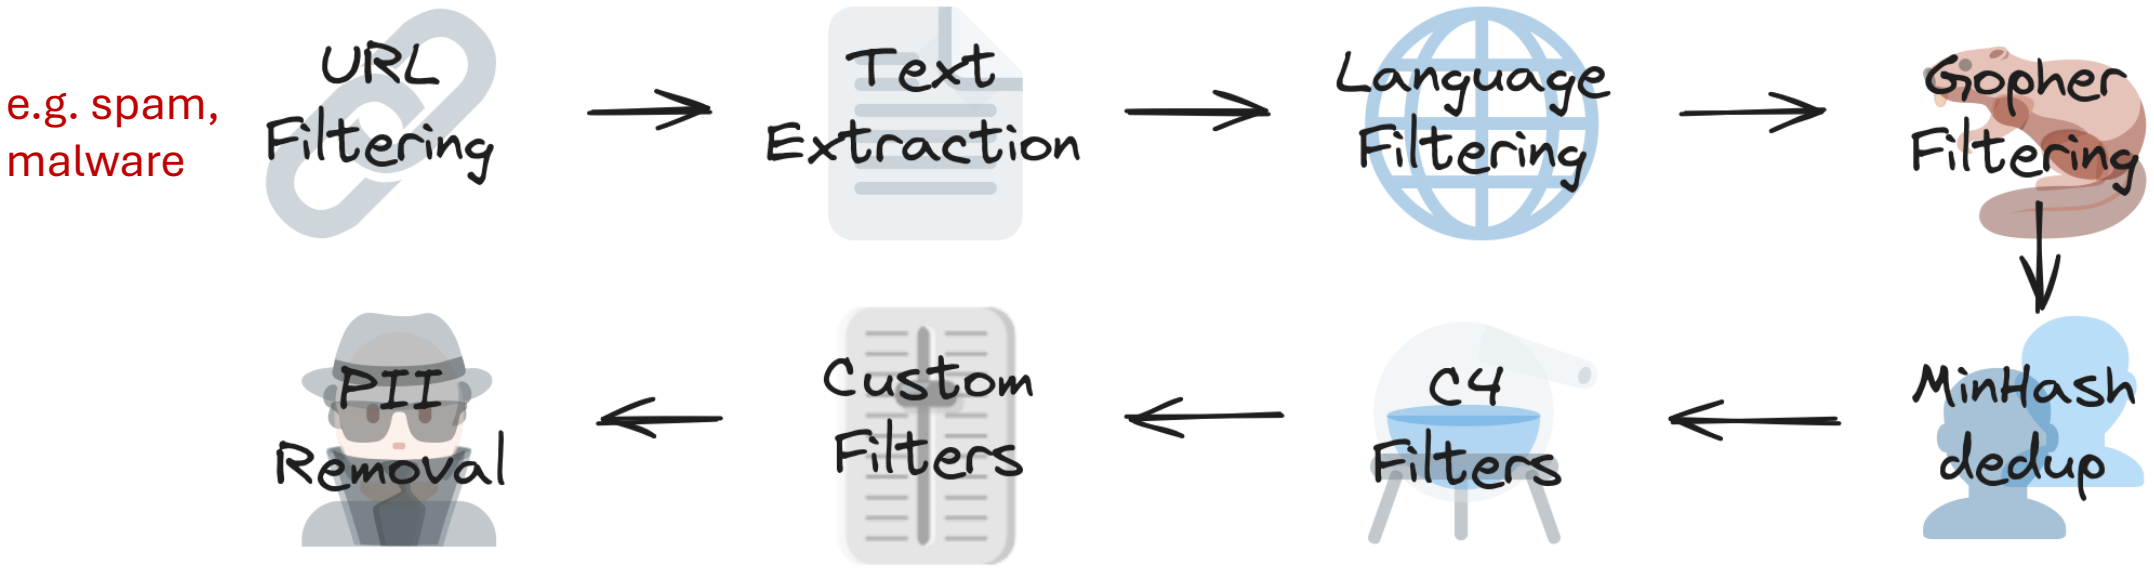

💡 https://en.wikipedia.org/wiki/Wolfgang_Amadeus_Mozart

web



3×10^{23} flops

pre training

Foundation model

(base model)



fine tuning

humans

Fine-tuned model

*A **lossy** compression*

web



3×10^{23} flops

pre training

Foundation model

(base model)



fine tuning

humans

Fine-tuned model

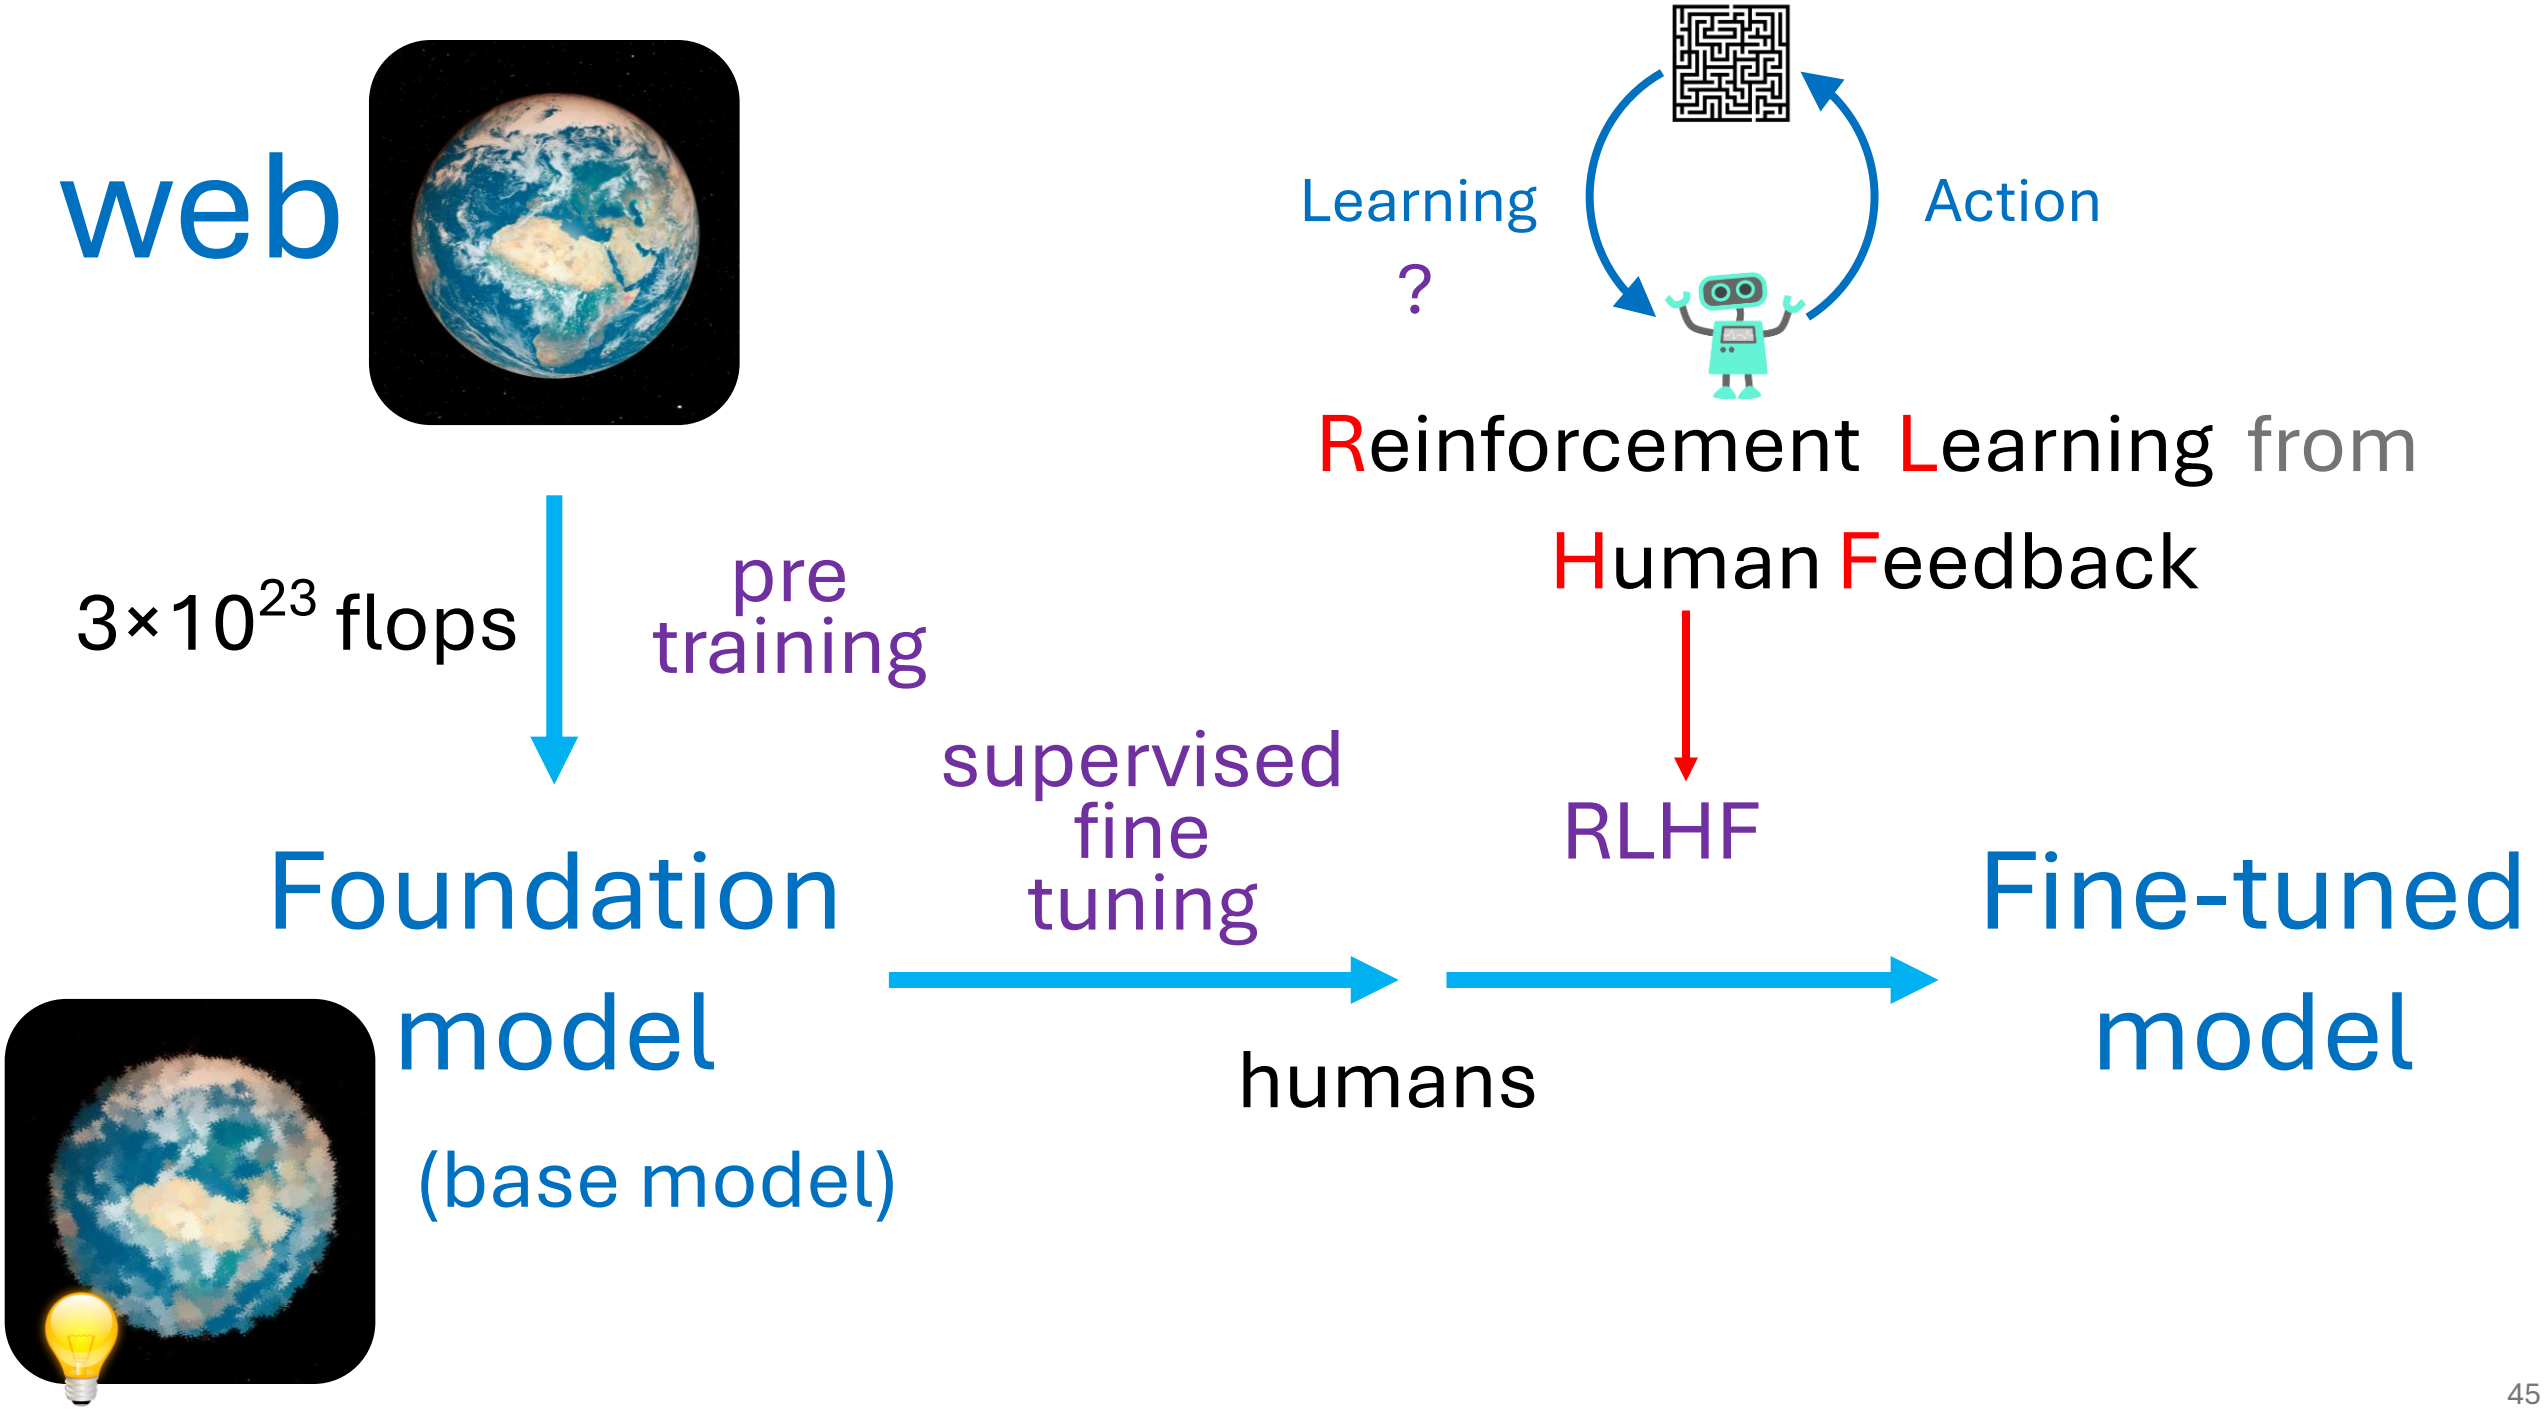
web
Learning
?
Action
Reinforcement Learning from
Human Feedback
3×10^23 flops
pre
training
supervised
fine
tuning
RLHF
Foundation
model
(base model)
humans
Fine-tuned
model

Supervised fine-tuning

Supervised learning



dog

labelled image

Why is the sky blue?

The sky appears blue due to a phenomenon called **Rayleigh scattering** …

How can I hack into a computer?

I'm sorry, I can't help with that.

ChatGPT

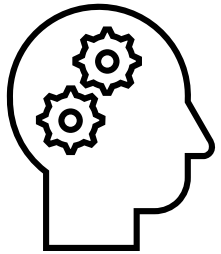

Get the user input
Verify
Tokenize it

Large language model

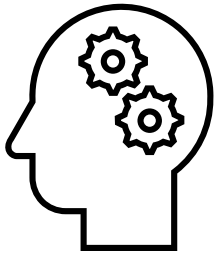

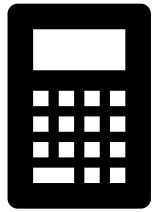

Get new tokens off the LLM
Verify
Display them



https://arxiv.org/pdf/2203.02155

harmless

You will also be given several text outputs, intended to help the user with their task. Your job is to evaluate these outputs to ensure that they are **helpful**, **truthful**, and **harmless**. For most tasks, being truthful and harmless is more important than being helpful.

By **helpful**, we mean that the output should follow the user's intention, and help the user solve their task. Some examples of helpful behavior:

- Writing in clear language.
- Answering the question they meant to ask, even if they mis-asked it.