a map of meaning

prof. dr. nicolas meseth

1 a map by hand
2 tokens get numbers
3 computing with meaning
4 from tokens to a text
5 the model's map
6 what the map does not know
7 inside the numbers

before a computer can search for meaning, it has to know where each text belongs.

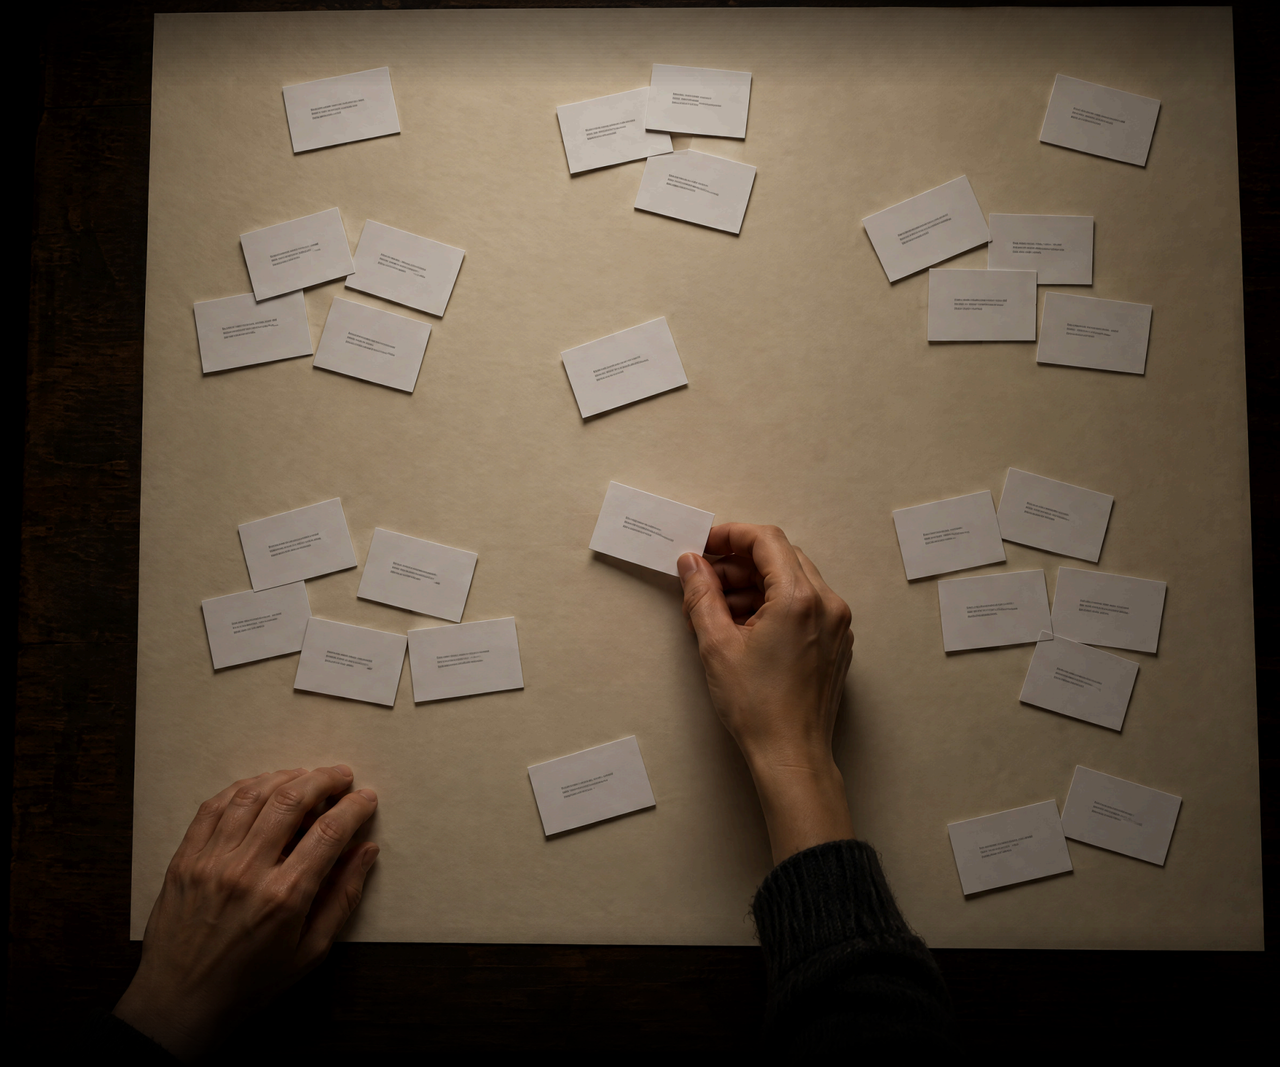
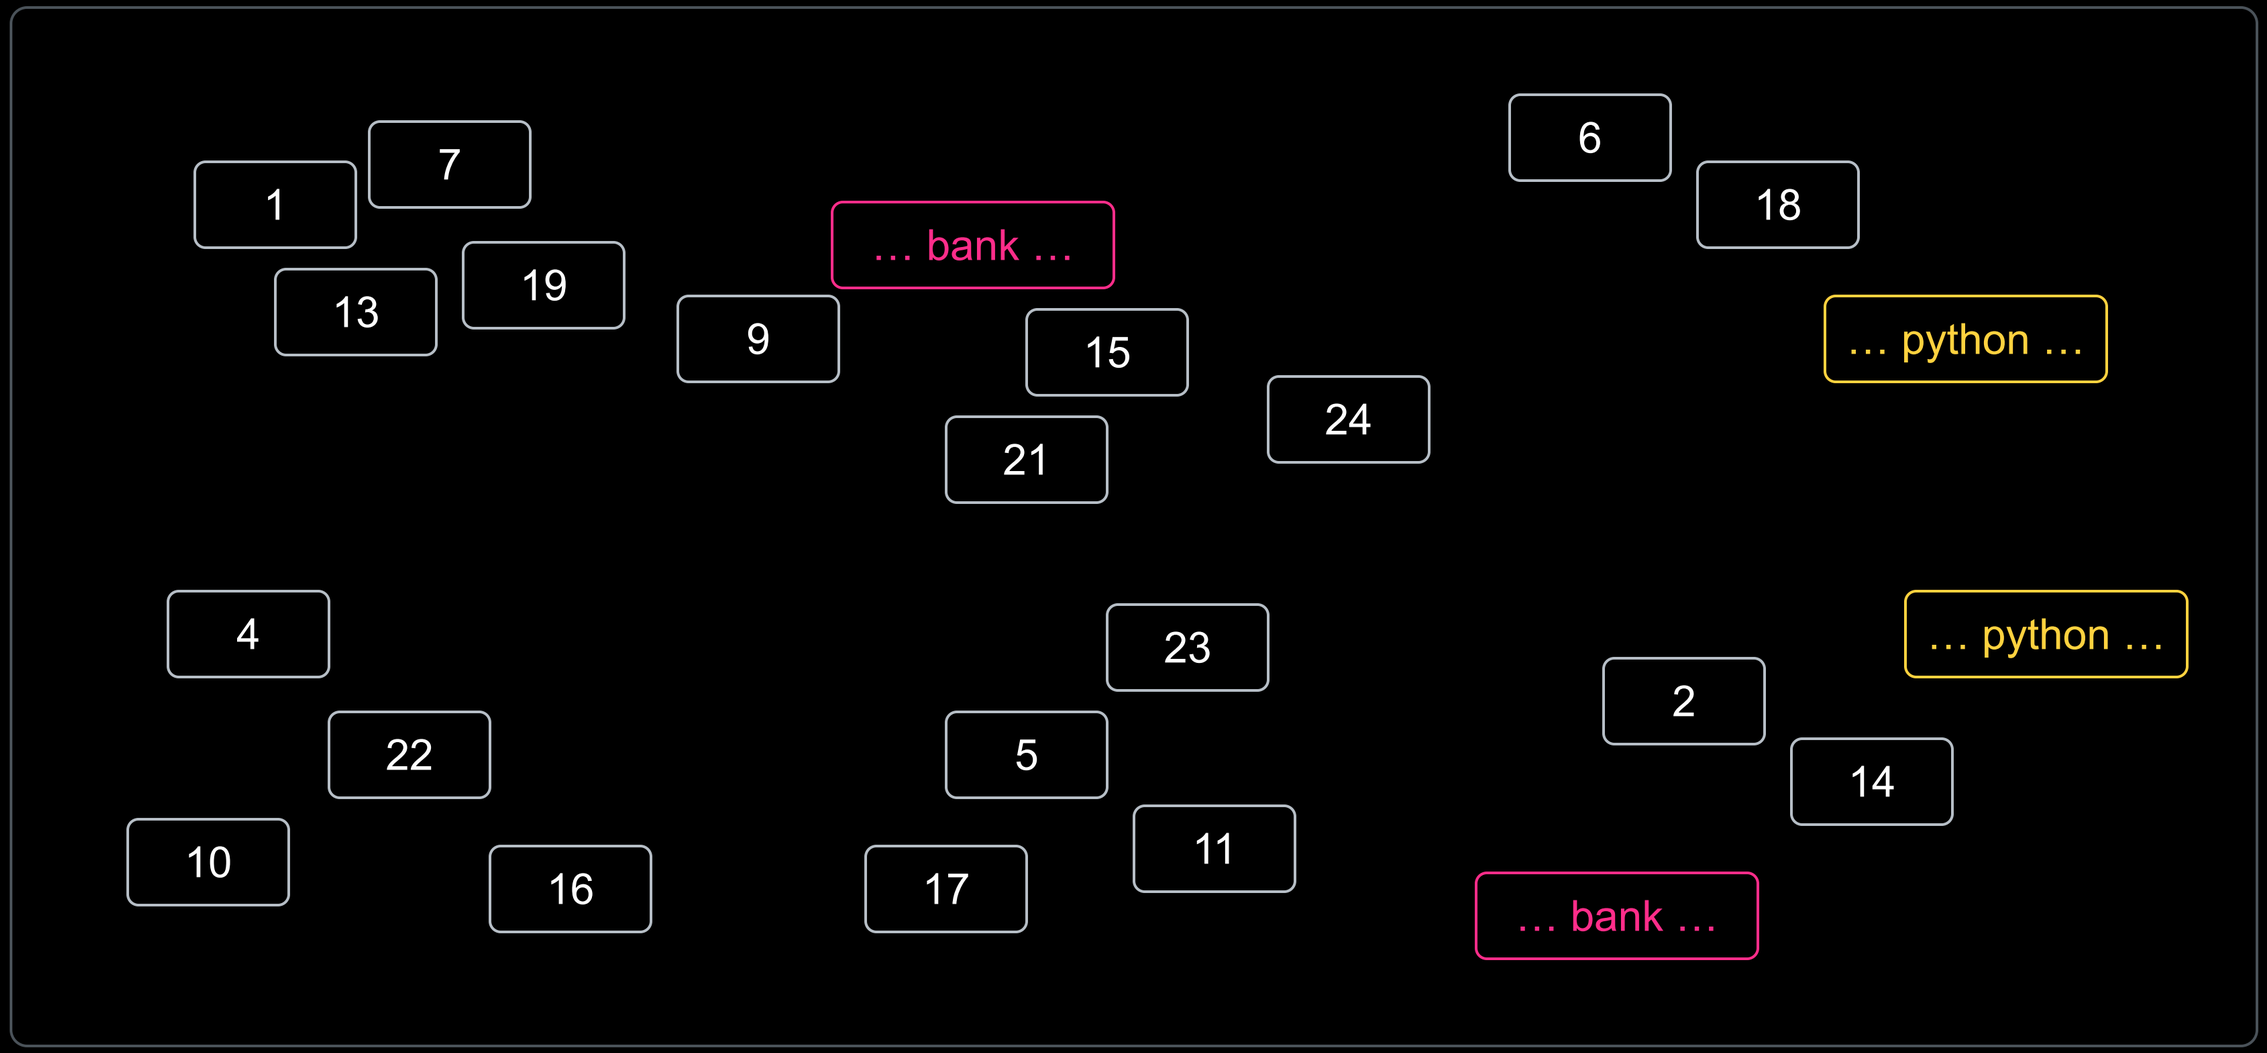


image: ai-generated (gpt-image-2.5-flare)

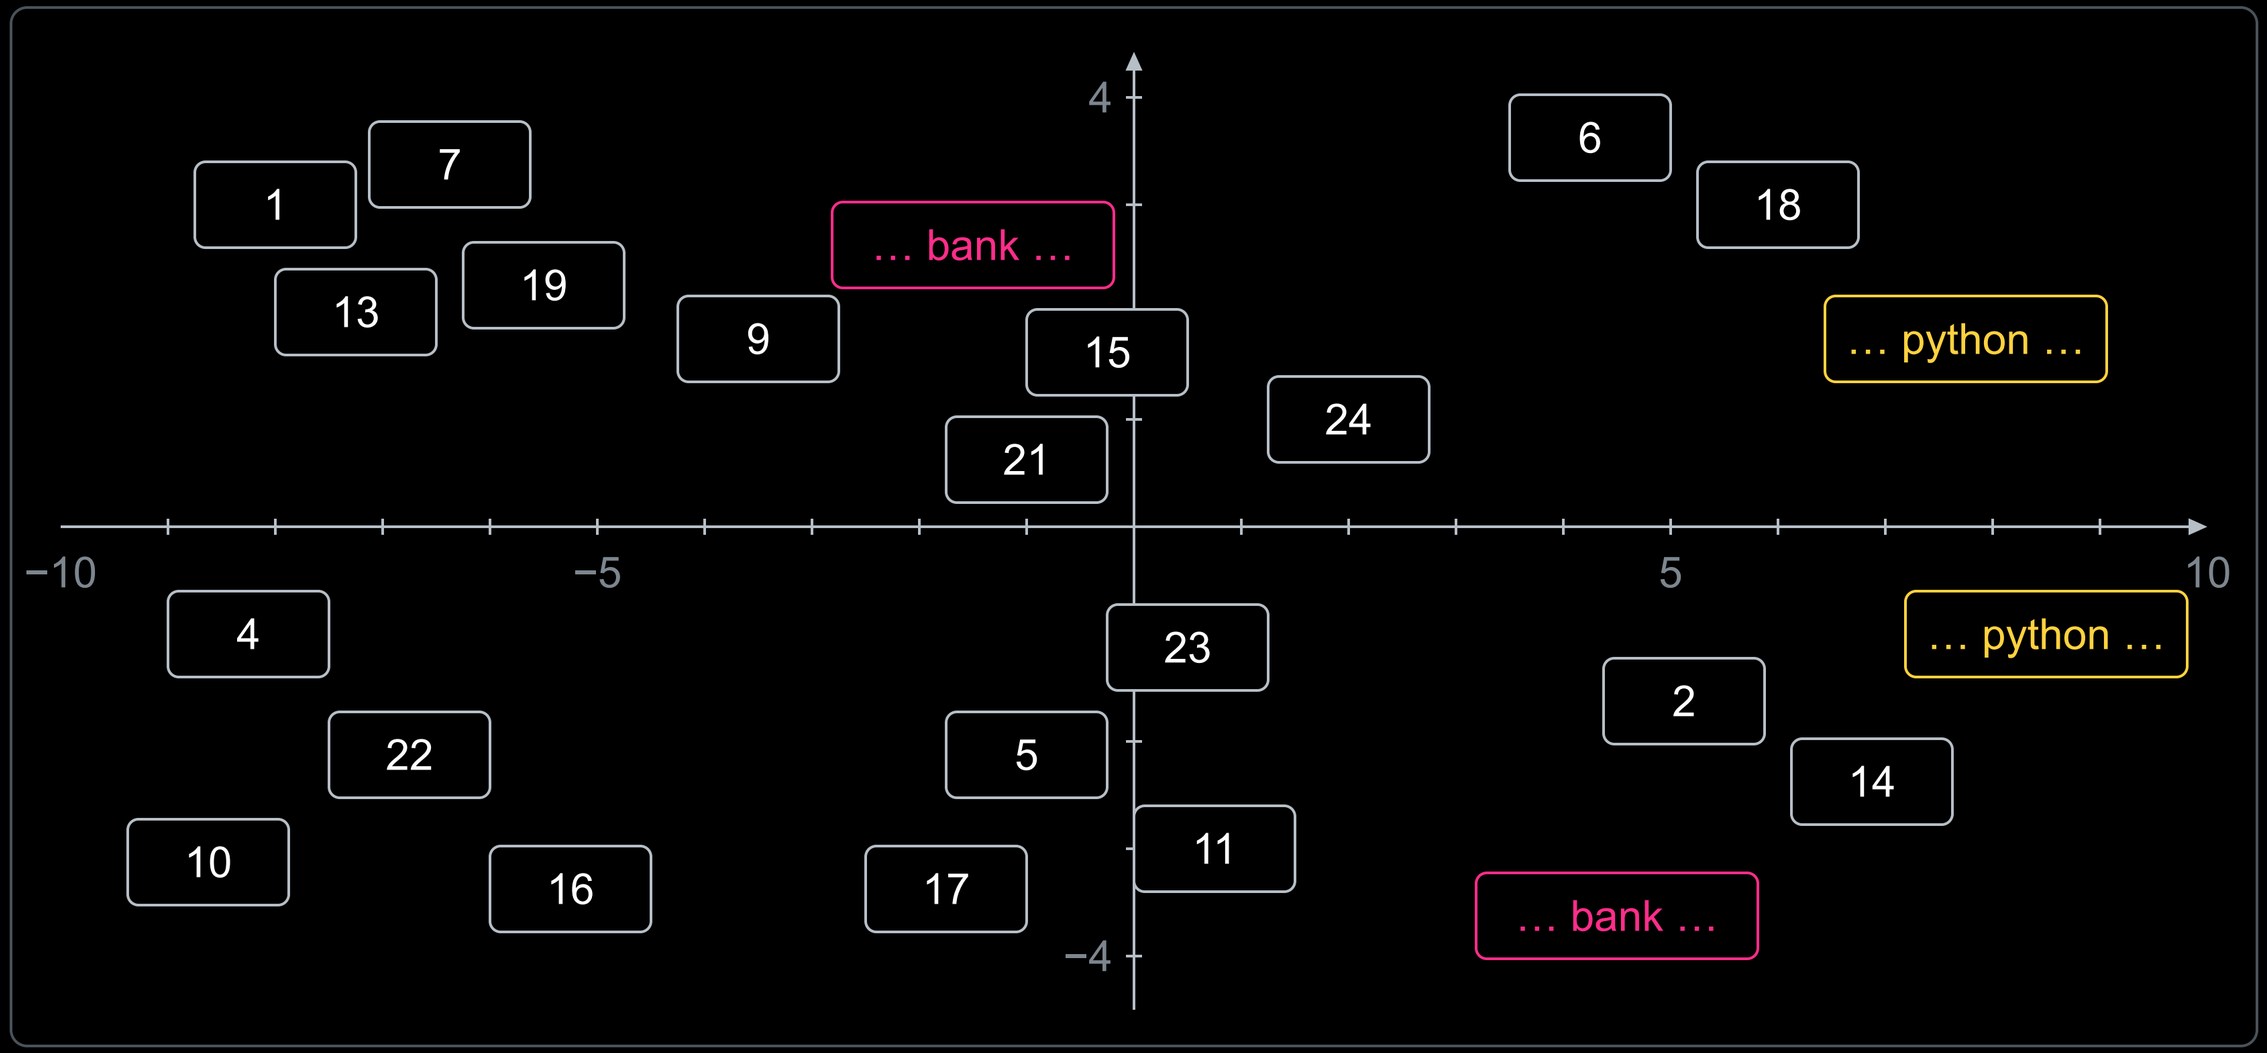
part 1

a map by hand

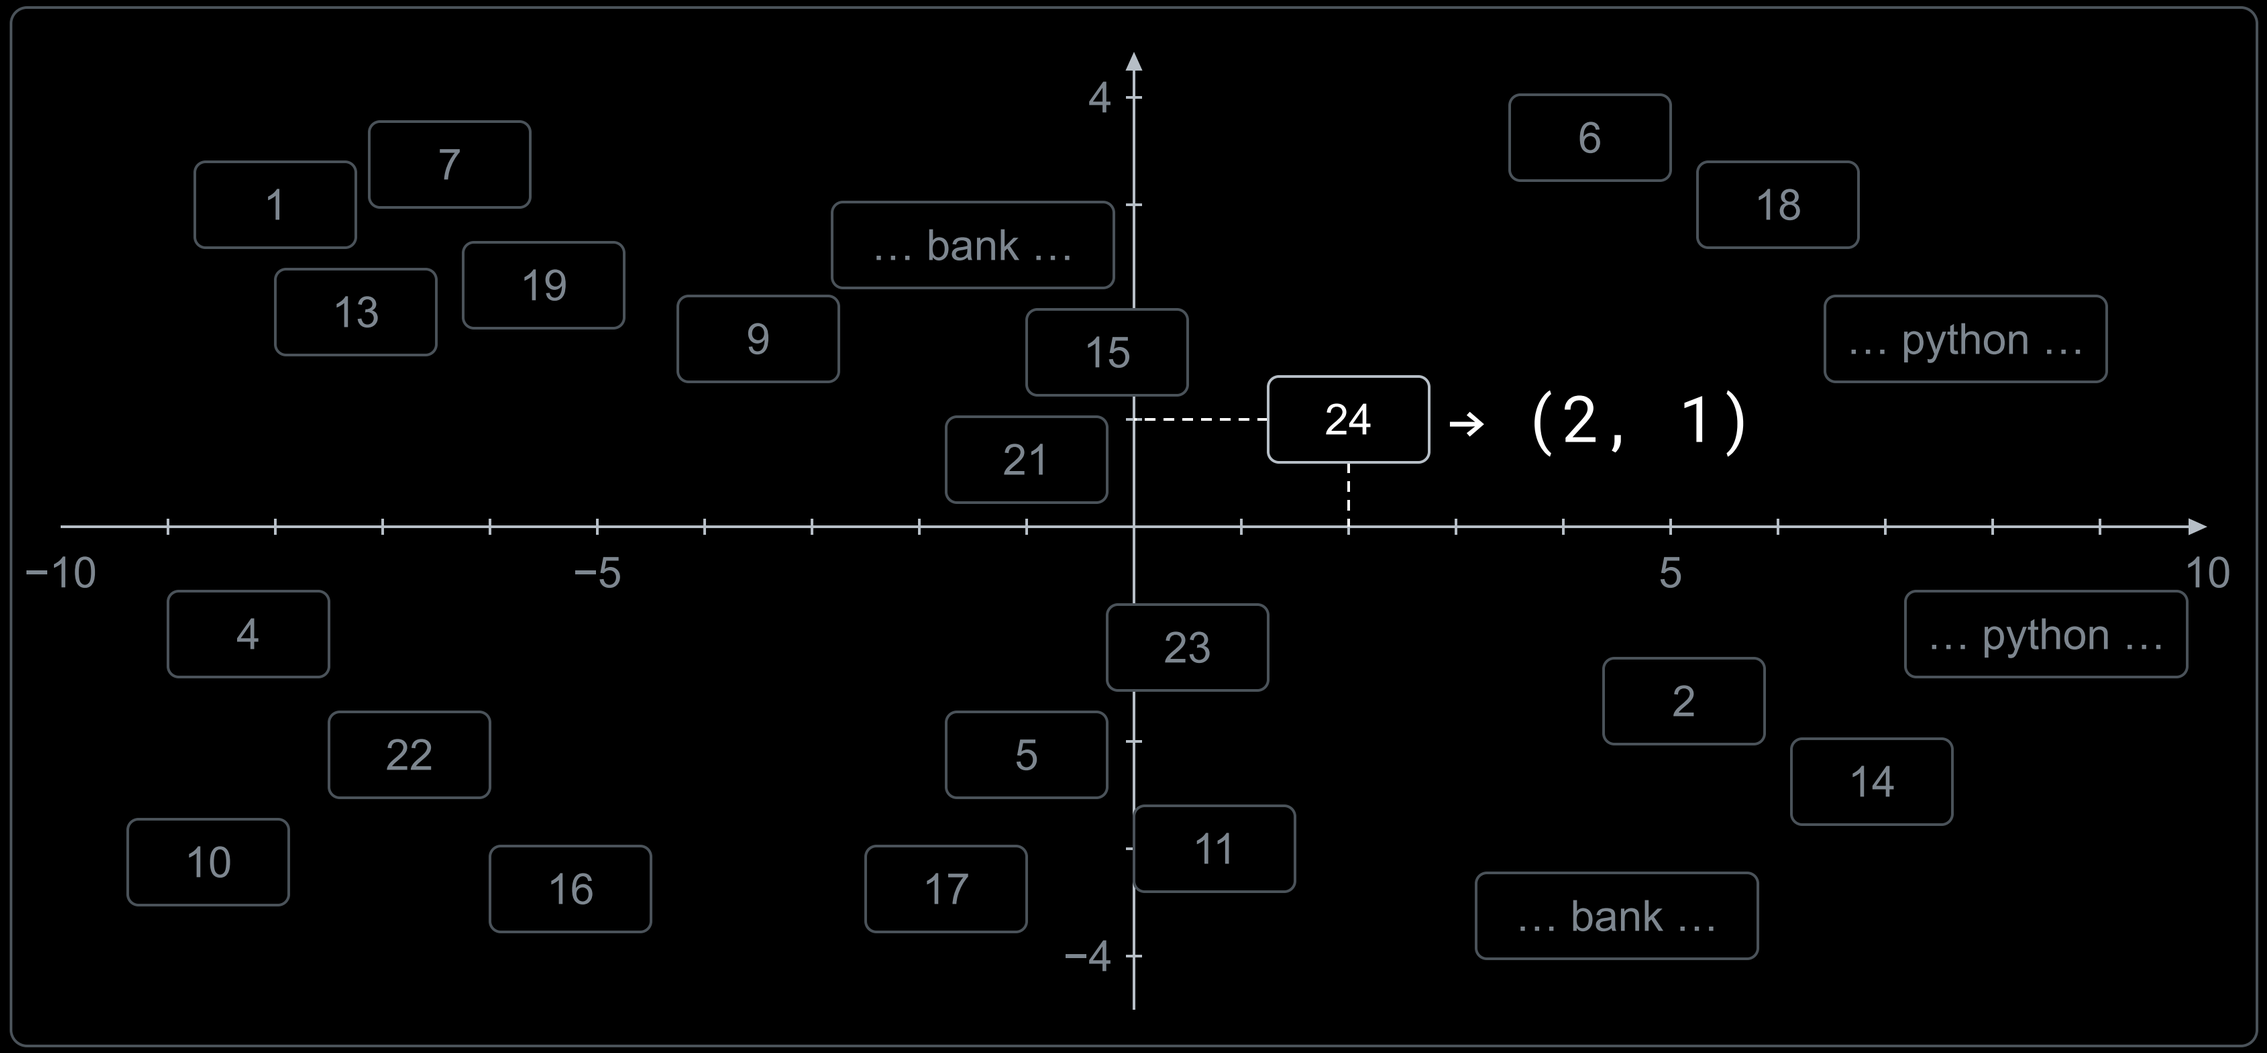
close in meaning, close on the map



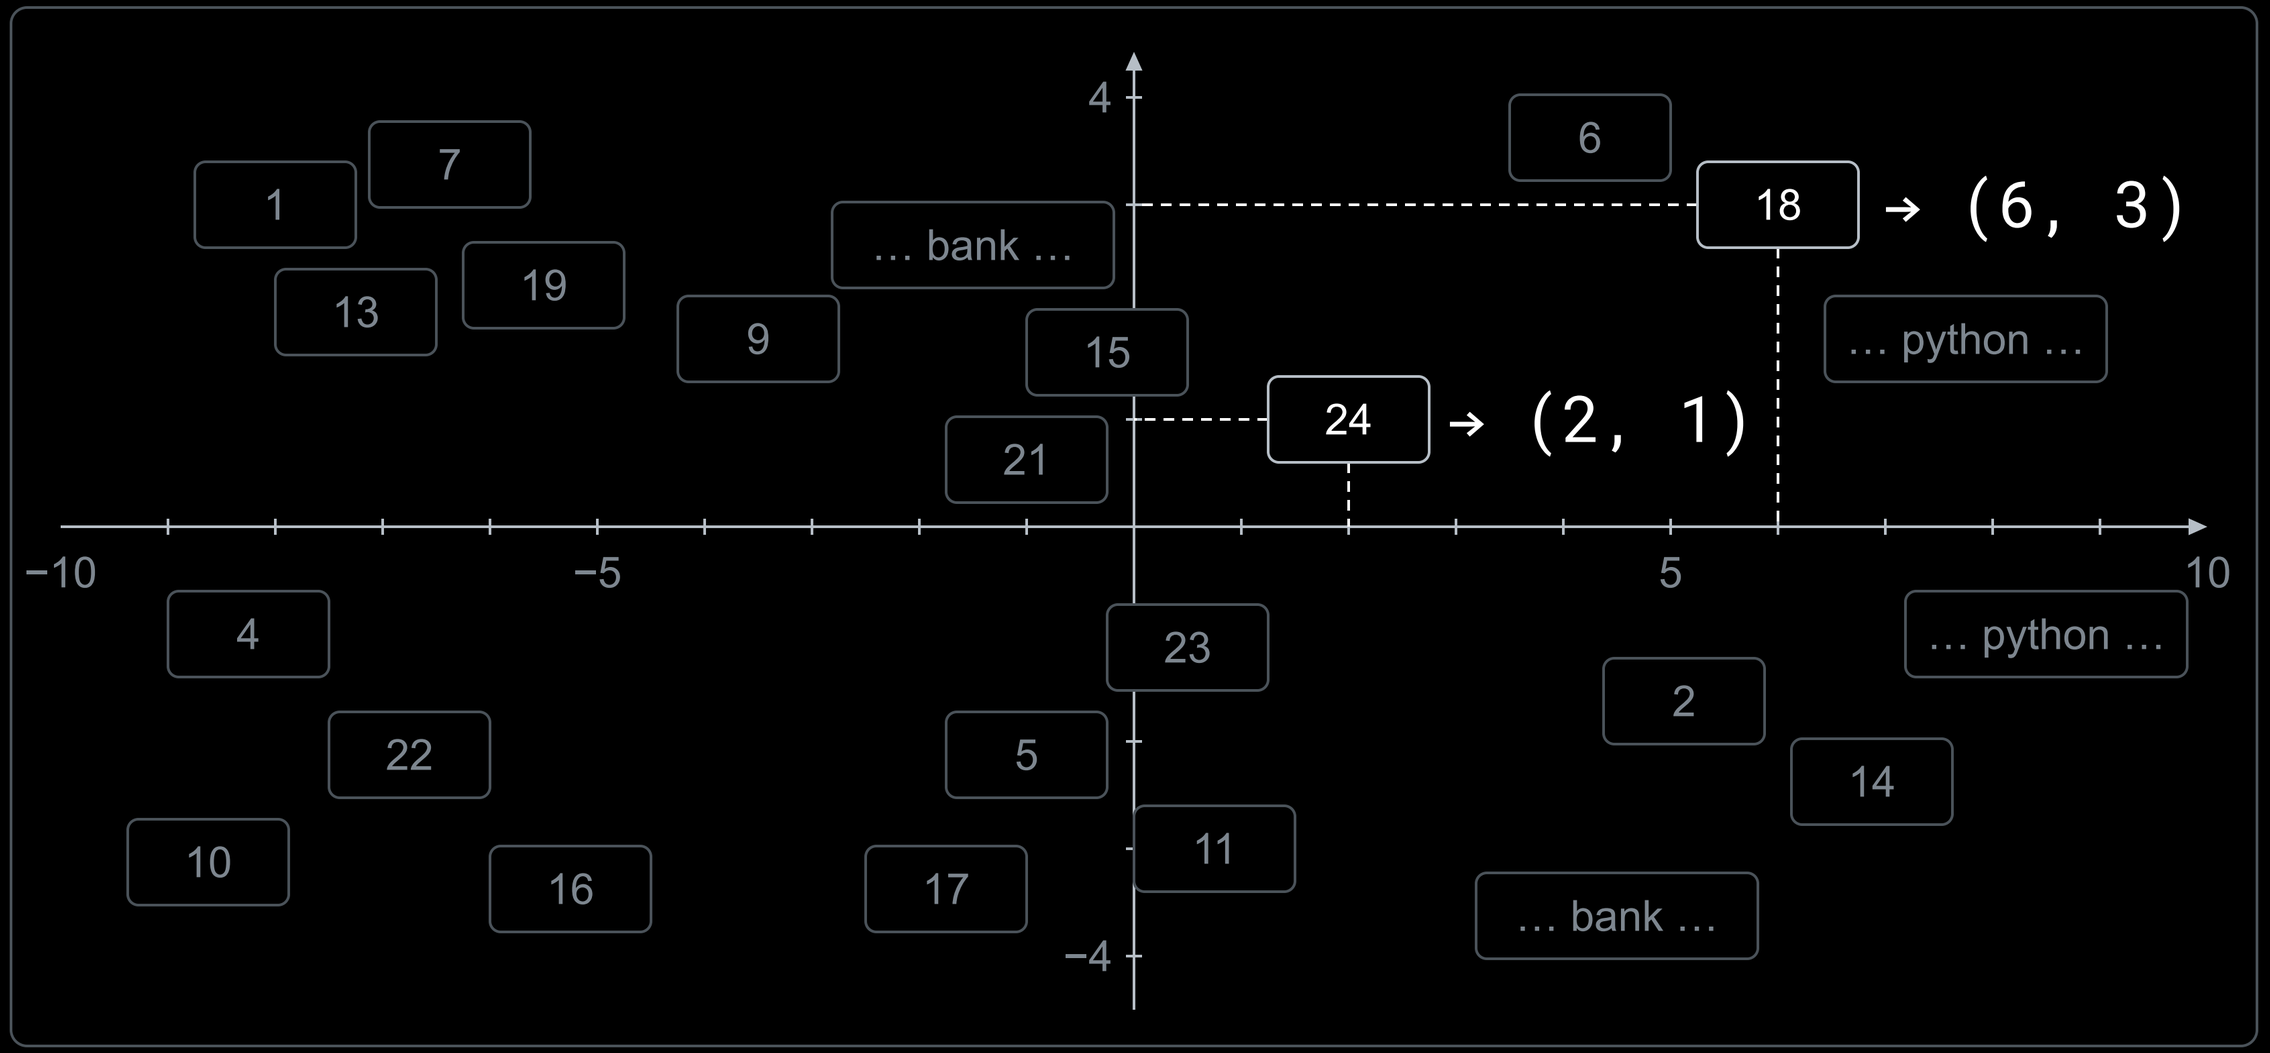
a place is a list of numbers



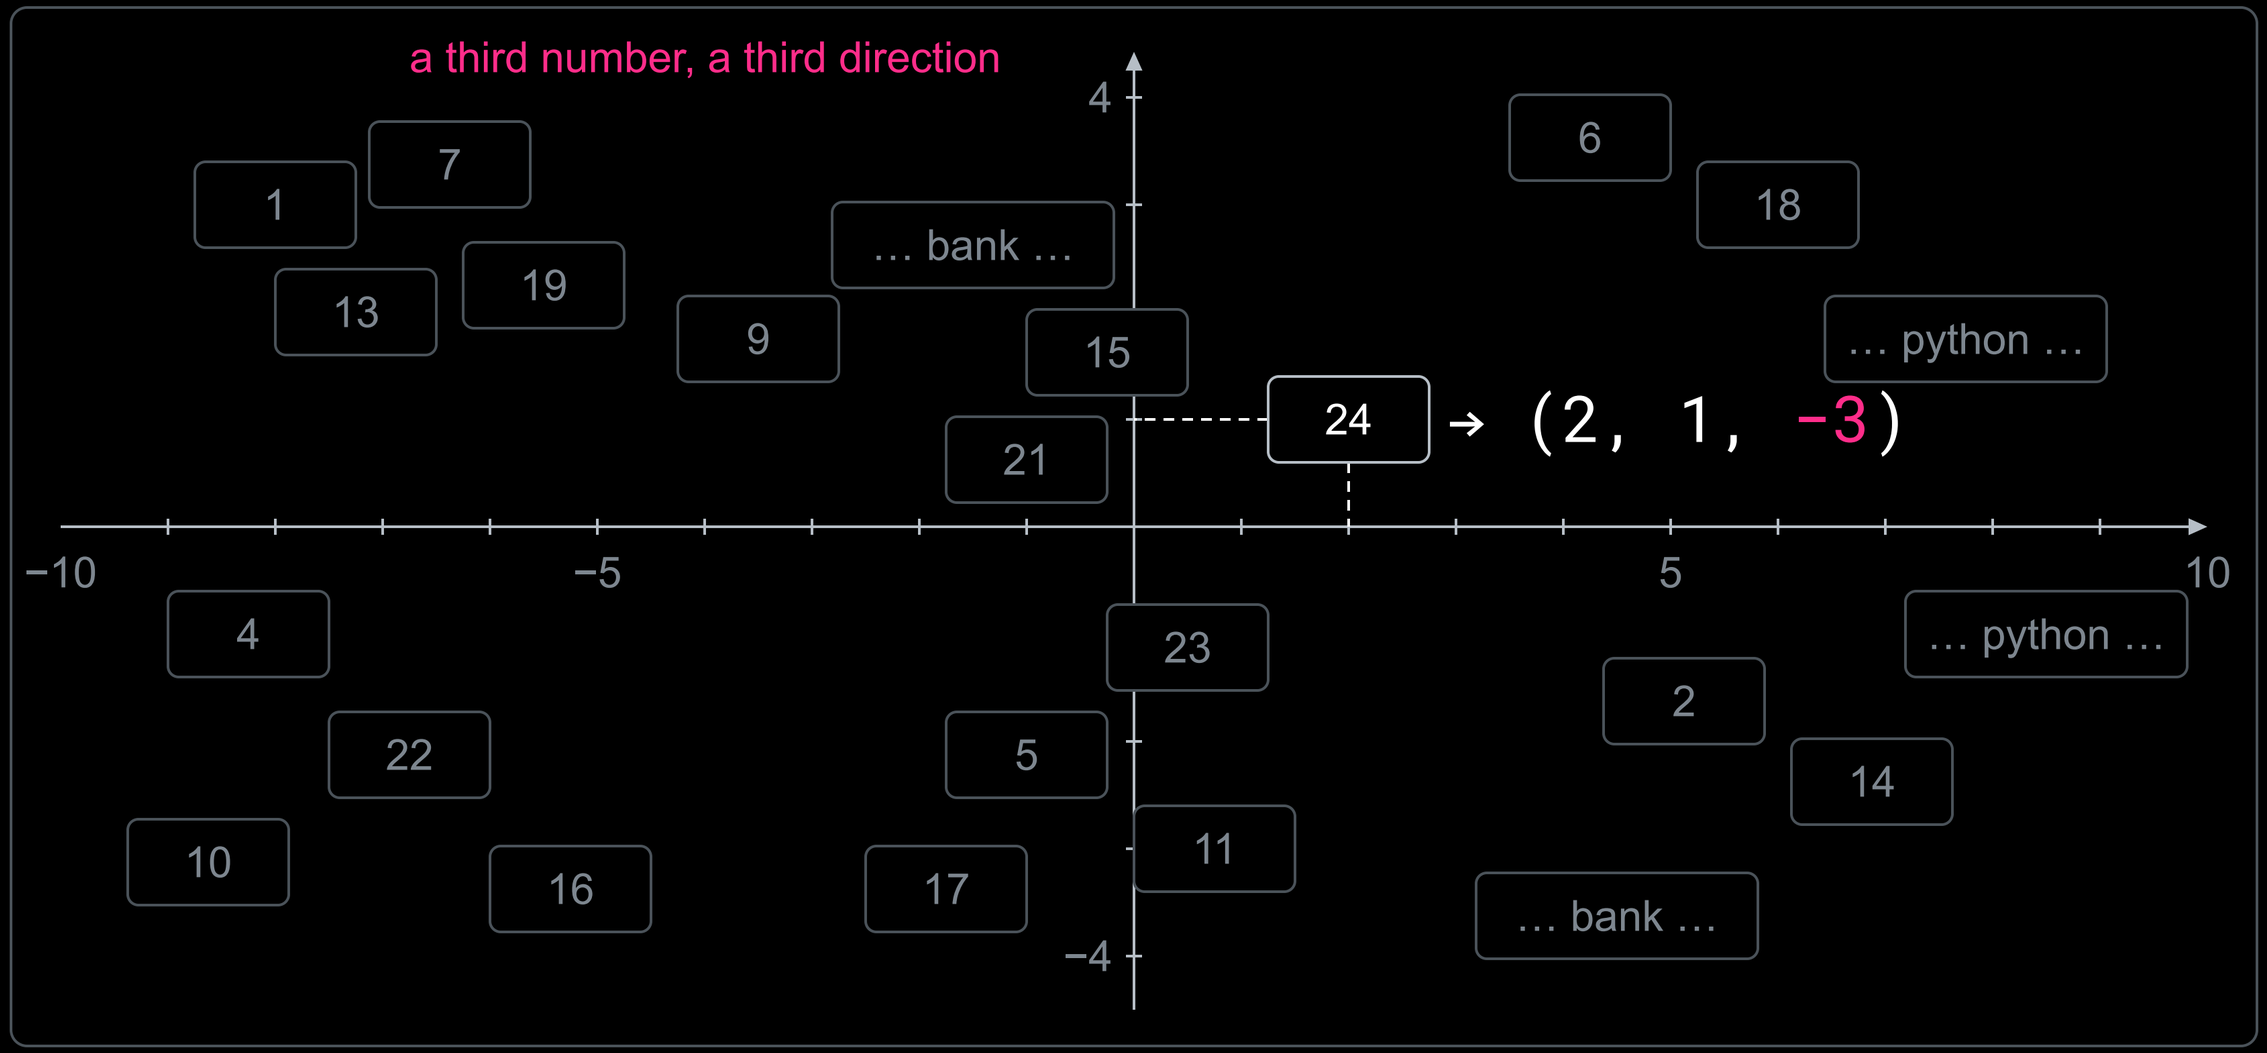
a place is a list of numbers



a place is a list of numbers



a place is a list of numbers



an embedding is a list of numbers that says where something lies on a map of meaning.

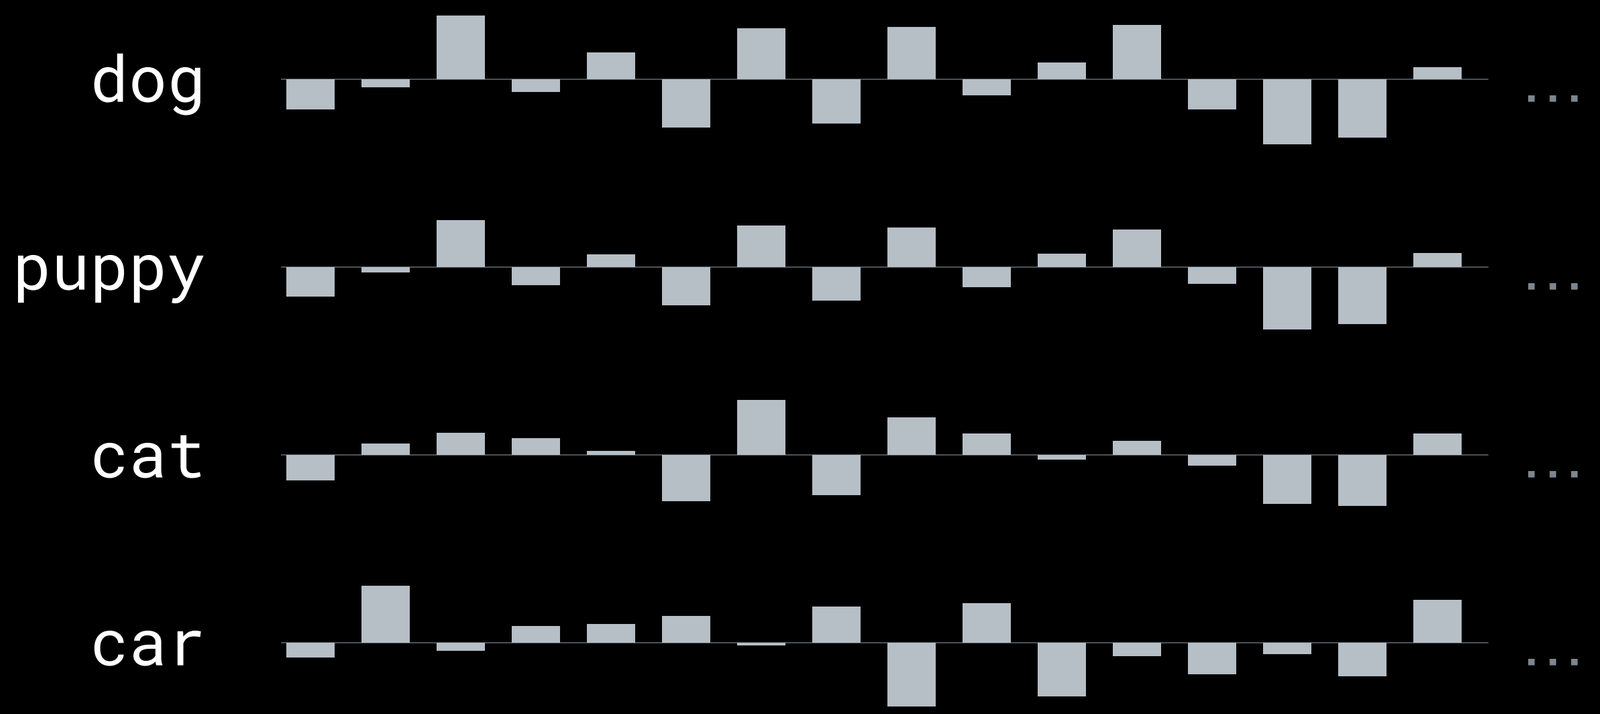
your sheet: 2 numbers per card · the model*: 384 numbers per token and per text

* all-MiniLM-L6-v2 · huggingface.co/sentence-transformers/all-MiniLM-L6-v2

part 2

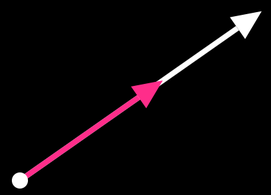
tokens get numbers

a table with one row per token



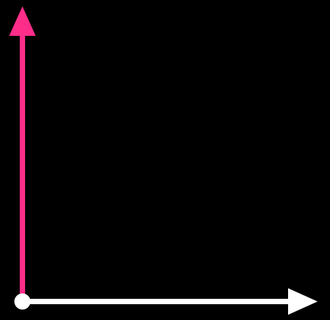
30,522 tokens × 384 numbers

similarity is the angle between two arrows



1

same direction



0

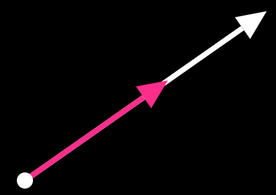
at a right angle



−1

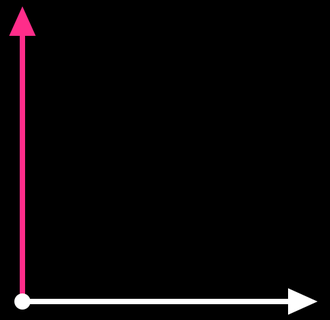
opposite

similarity is the angle between two arrows



1

same direction



0

at a right angle



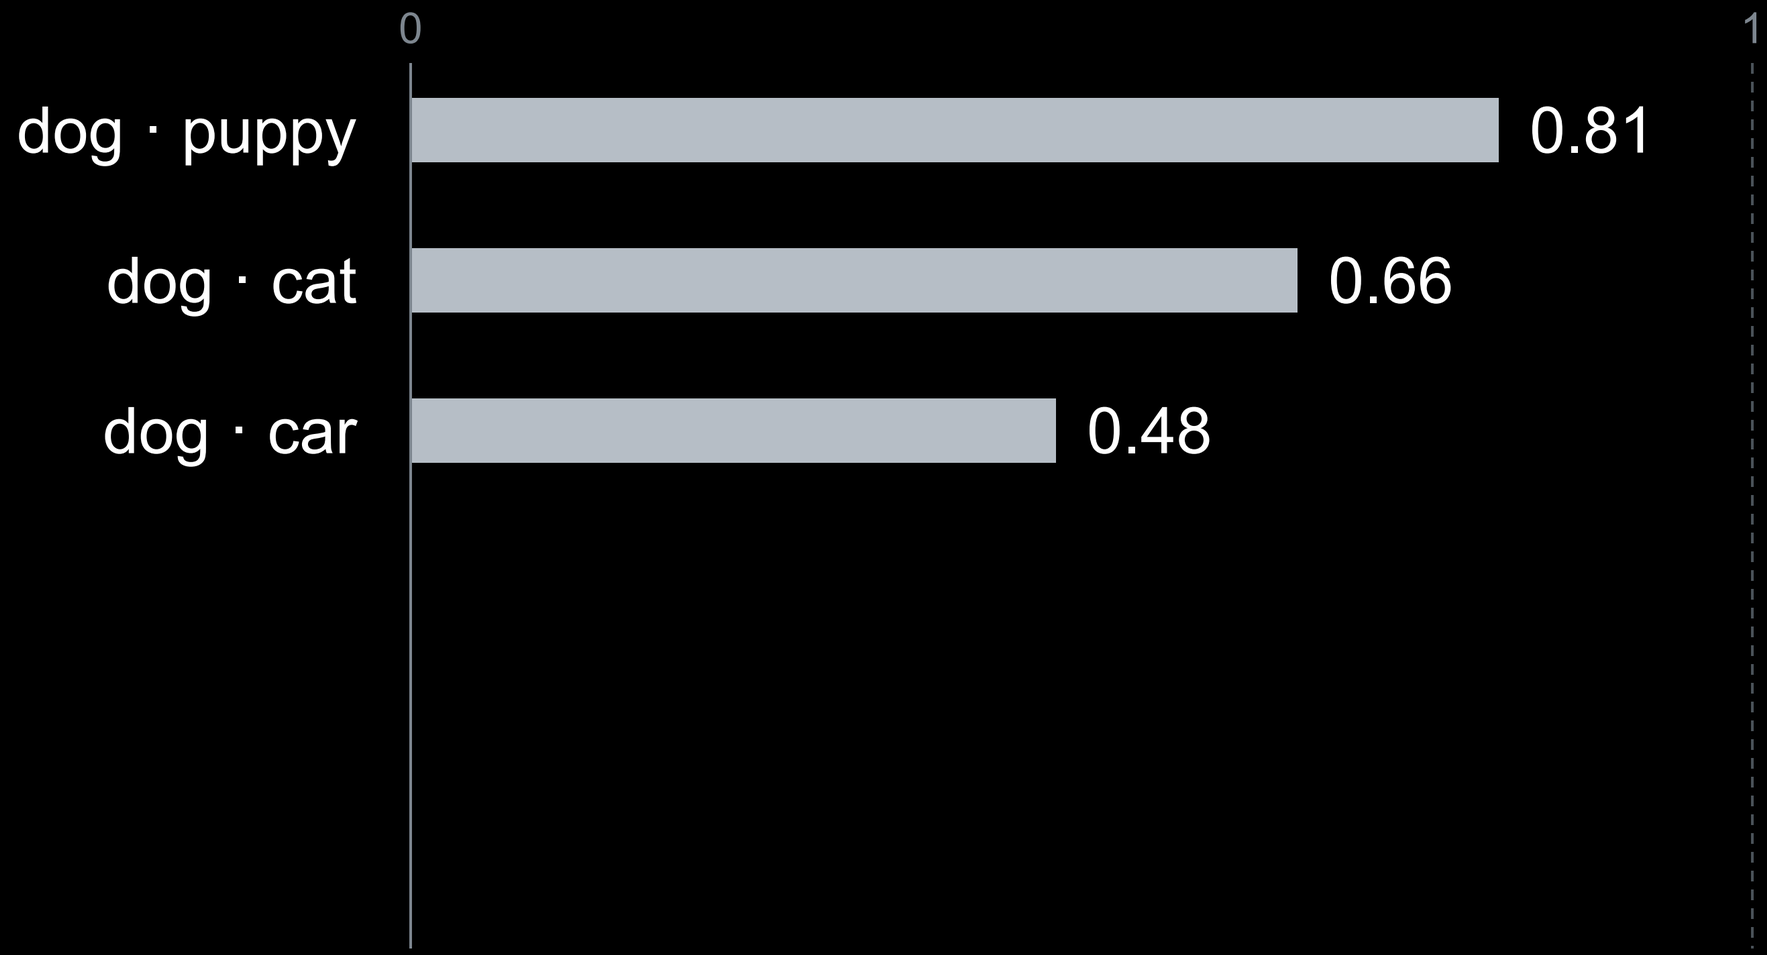
−1

opposite

almost never with the model*:
the lowest of 1.3 million
word pairs is −0.09

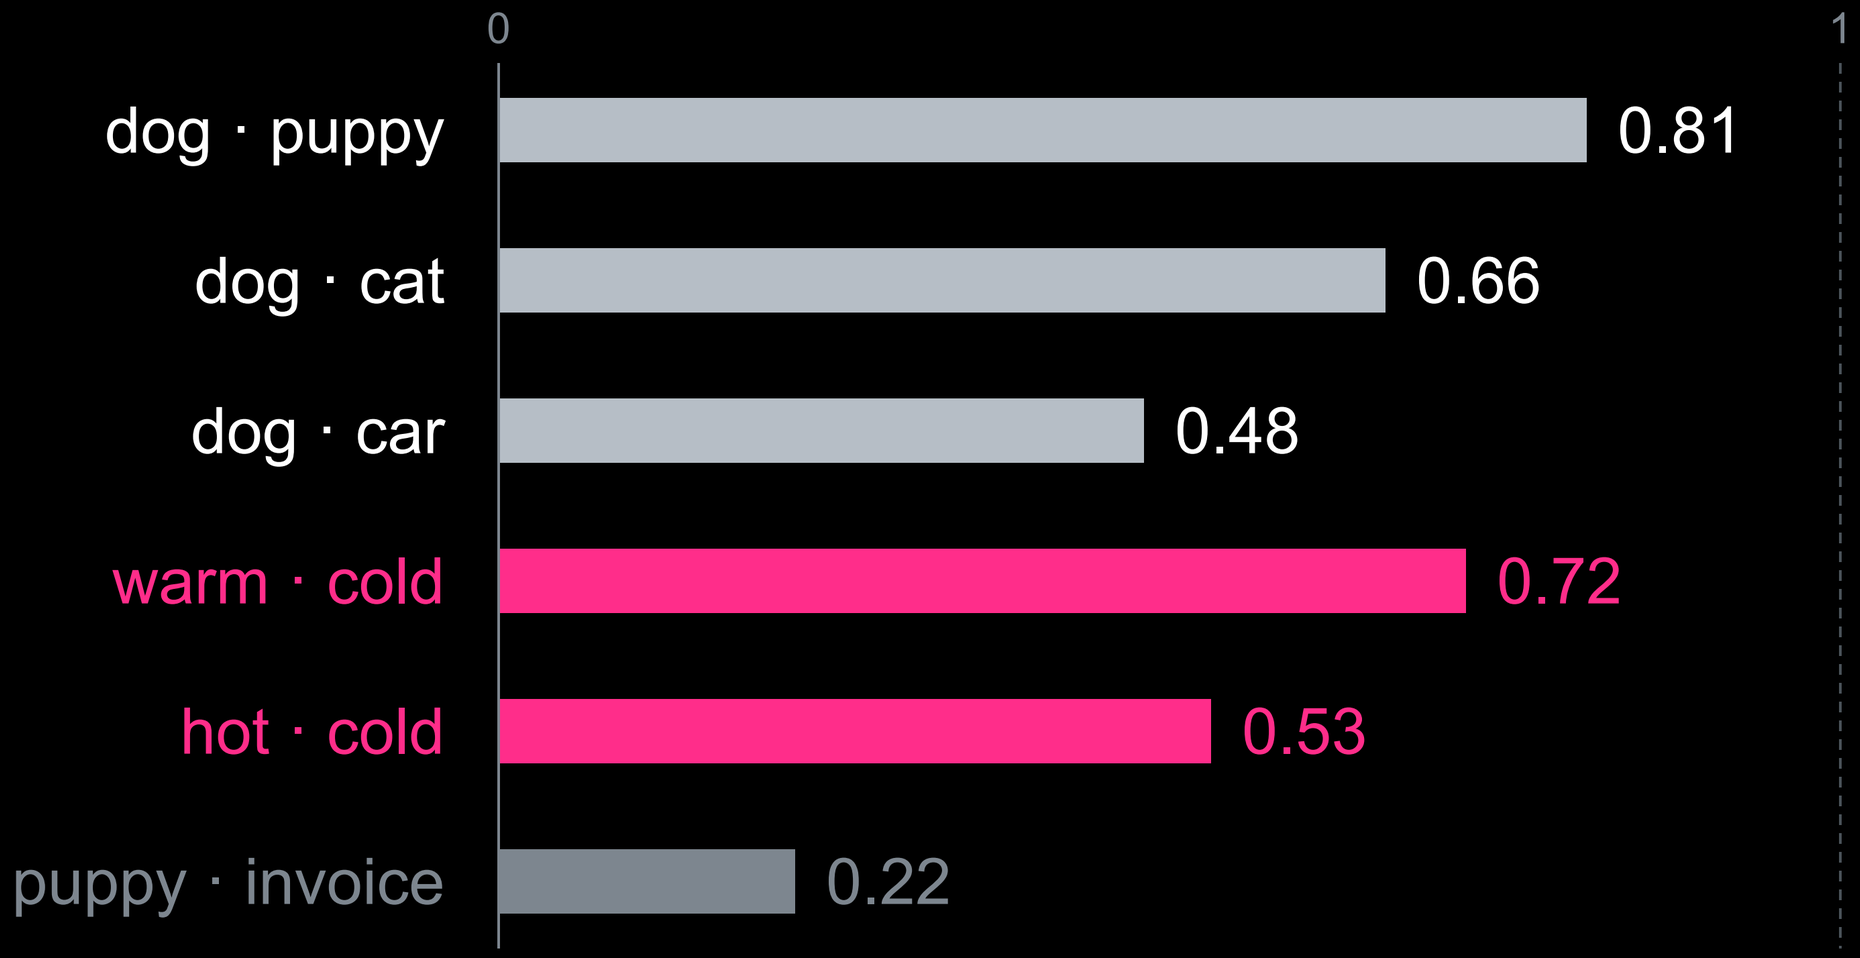
* all-MiniLM-L6-v2 · huggingface.co/sentence-transformers/all-MiniLM-L6-v2

measured with the model*



* all-MiniLM-L6-v2 · huggingface.co/sentence-transformers/all-MiniLM-L6-v2

measured with the model*



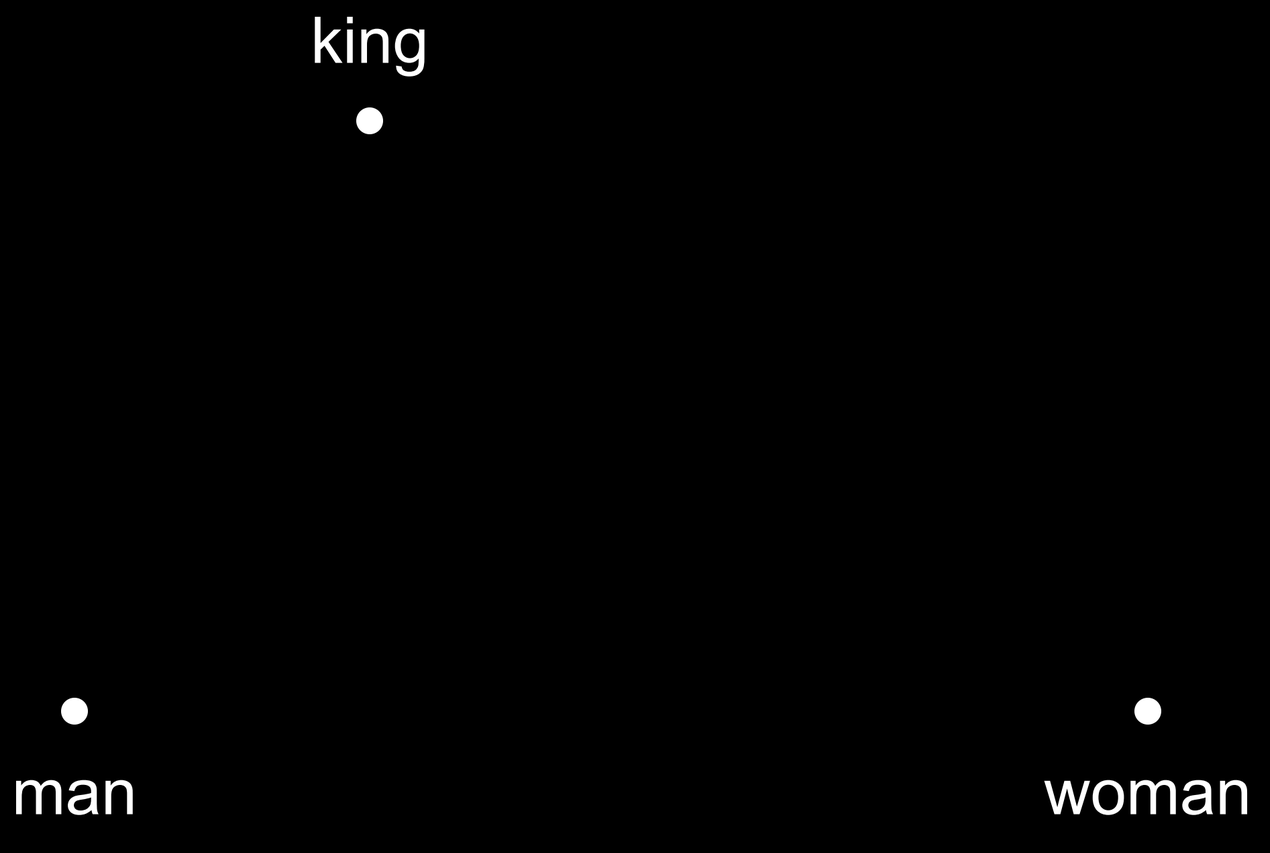
* all-MiniLM-L6-v2 · huggingface.co/sentence-transformers/all-MiniLM-L6-v2

tokens get numbers

part 3

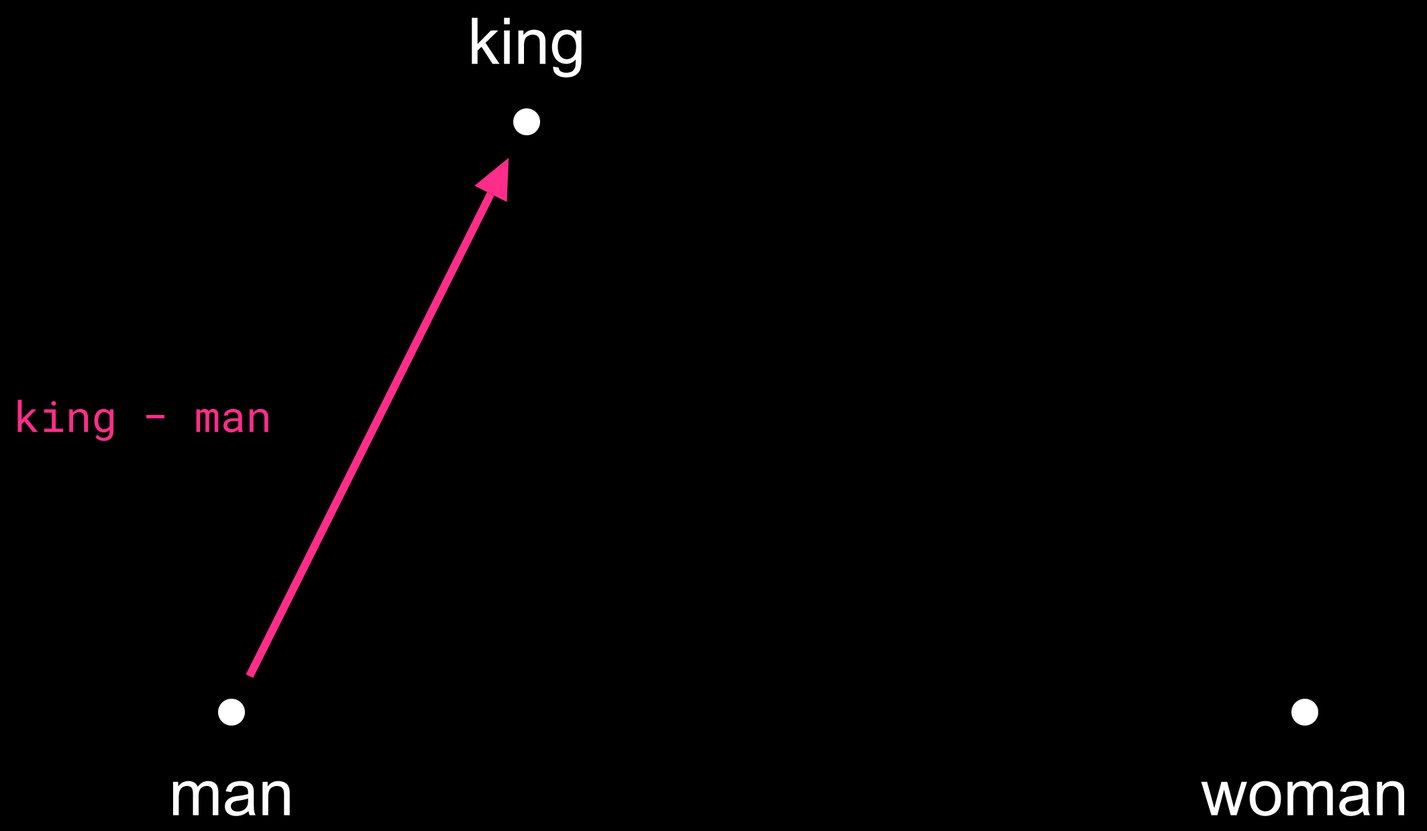
computing with meaning

king − man + woman

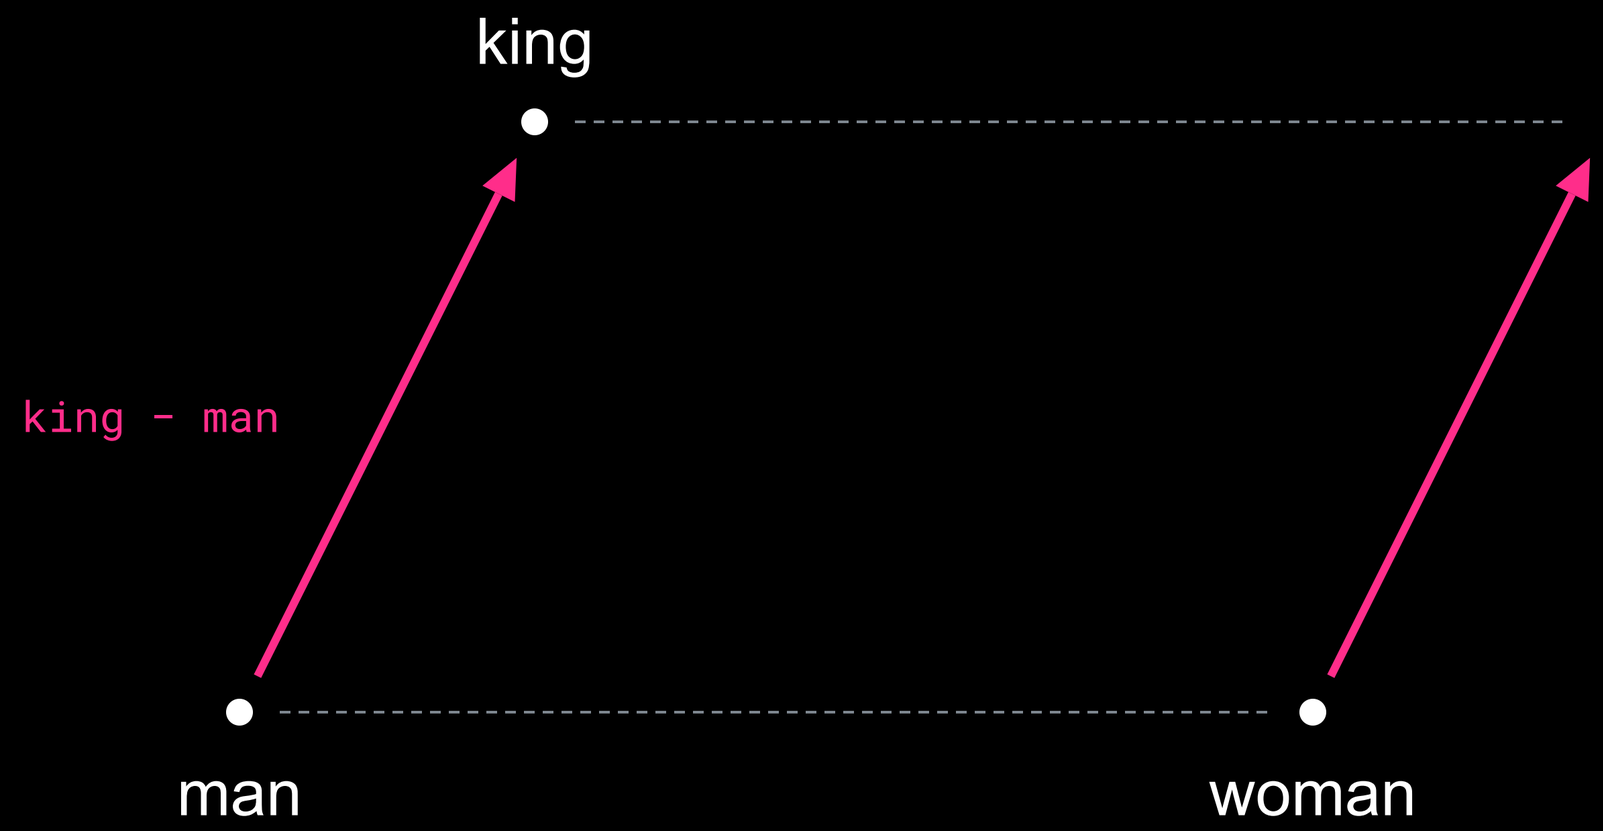


a sketch, not to scale

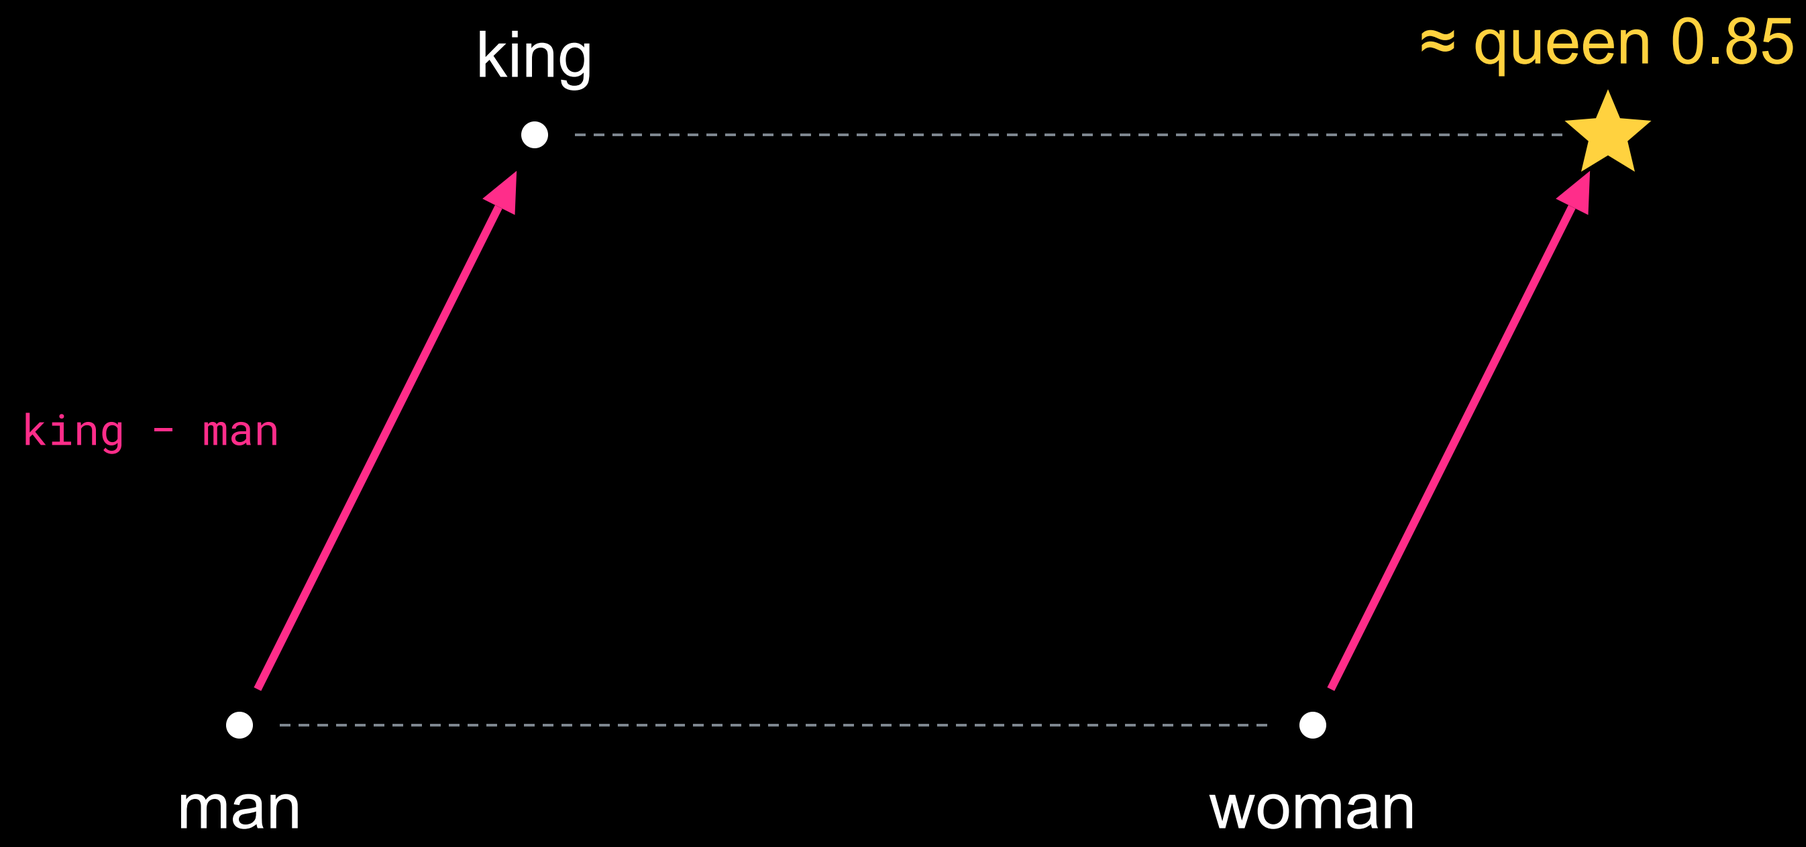
king − man + woman



a sketch, not to scale



a sketch, not to scale



a sketch, not to scale

```
paris - france + germany

puppy - dog + cat

doctor - hospital + school

walked - walk + swim
```

```
paris - france + germany    →  berlin 0.79

puppy - dog + cat

doctor - hospital + school

walked - walk + swim
```

```
paris - france + germany    → berlin 0.79

puppy - dog + cat           → kitten 0.87

doctor - hospital + school

walked - walk + swim
```

```
paris - france + germany    → berlin 0.79

puppy - dog + cat           → kitten 0.87

doctor - hospital + school  → teacher 0.82

walked - walk + swim
```

```
paris - france + germany    → berlin 0.79
puppy - dog + cat           → kitten 0.87
doctor - hospital + school  → teacher 0.82
walked - walk + swim        → swimming 0.91  swam 0.88
```

```
queen - woman + man
```



part 4

from tokens to a text

what the model* reads

I can't get home tonight	[CLS] i can ' t get home tonight [SEP]
unbelievably	[CLS] un ##bel ##ie ##va ##bly [SEP]
ChatGPT	[CLS] chat ##gp ##t [SEP]
Ich bin krank	[CLS] ich bin k ##rank [SEP]

unbelievably · incredibly 0.35

continues the piece before it, without a space

* all-MiniLM-L6-v2 · huggingface.co/sentence-transformers/all-MiniLM-L6-v2

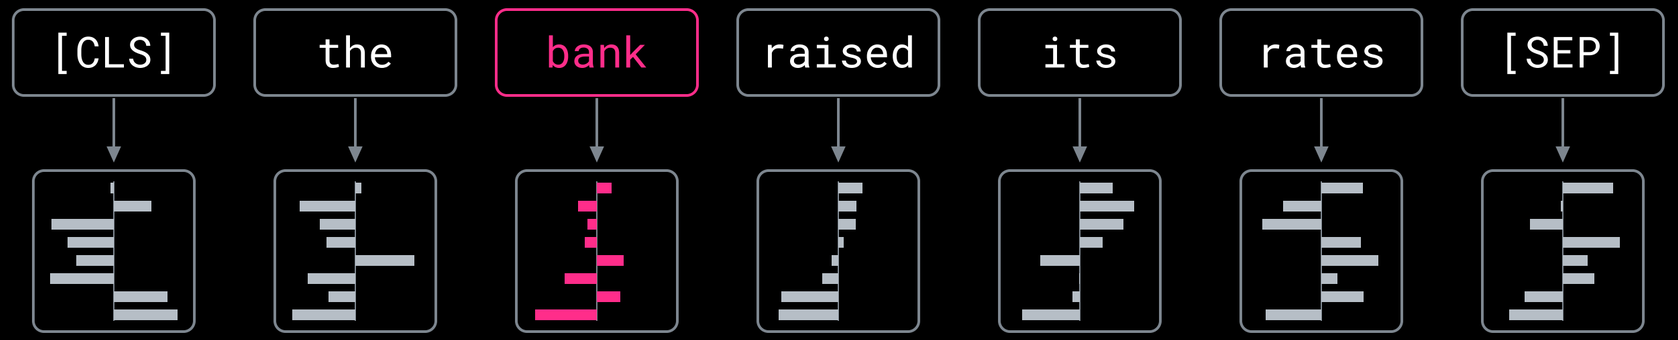
cutting



four steps from tokens to one list

cutting

looking up
384 numbers each



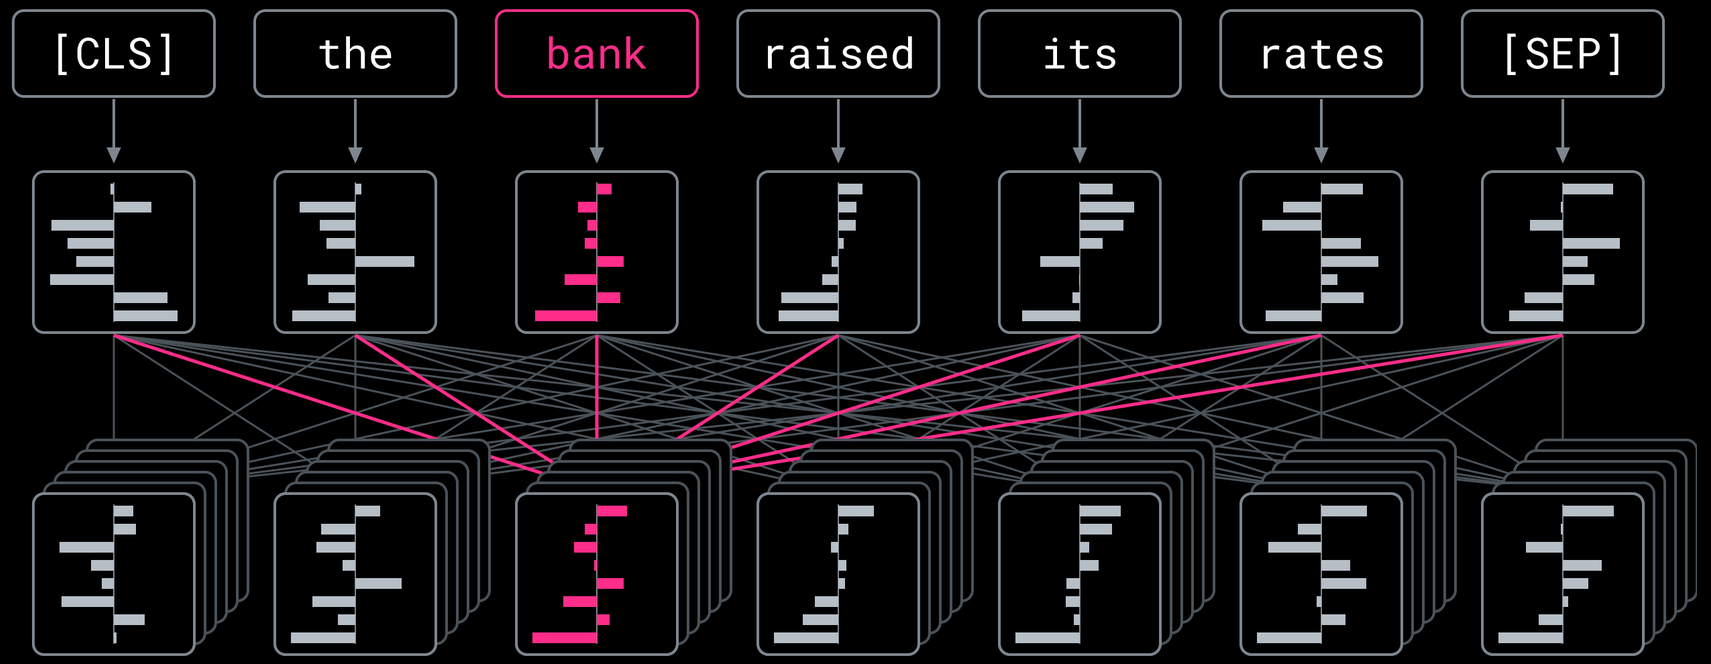
four steps from tokens to one list

cutting

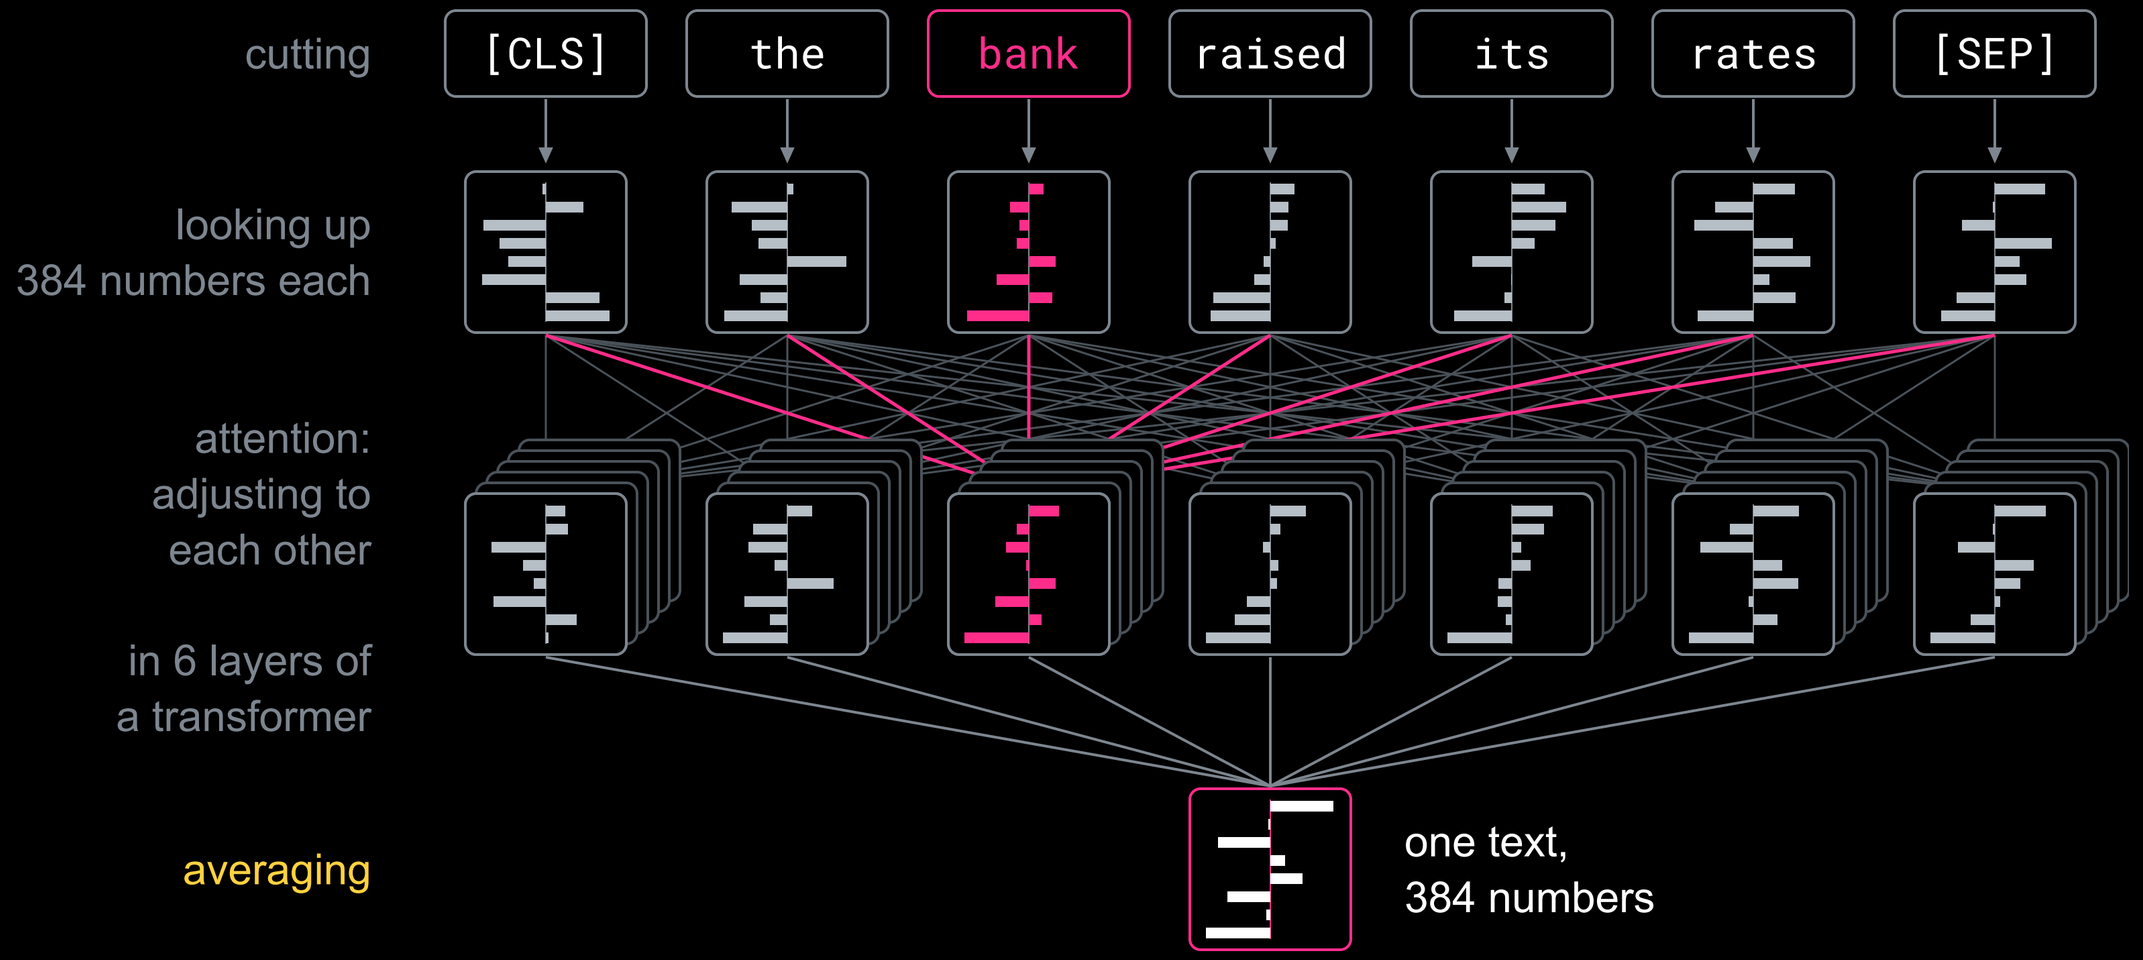
looking up
384 numbers each

attention:
adjusting to
each other

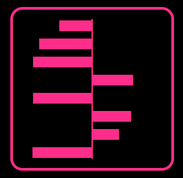
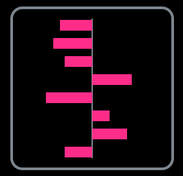
in 6 layers of
a transformer



four steps from tokens to one list



the same token, two meanings

	384 numbers of this “bank”	similarity to the first “bank”
the bank raised its interest rates.		
we watched the ducks from the bank of the river.		0.76
my salary goes straight to the bank.		0.88

all three start from the same row of the table, but after the six layers of the model* every “bank” has 384 numbers of its own

* all-MiniLM-L6-v2 · huggingface.co/sentence-transformers/all-MiniLM-L6-v2

where the numbers come from

stage 1: guess the hidden word

the ____ barked
at the postman.

which word fits?

dog

puppy

words that fit the same gaps
get similar numbers

stage 2: pairs that belong together

how do i reset
my password?

closer

open settings
and click reset.

this soup needs
more salt.

further away

more than 1 billion pairs of texts

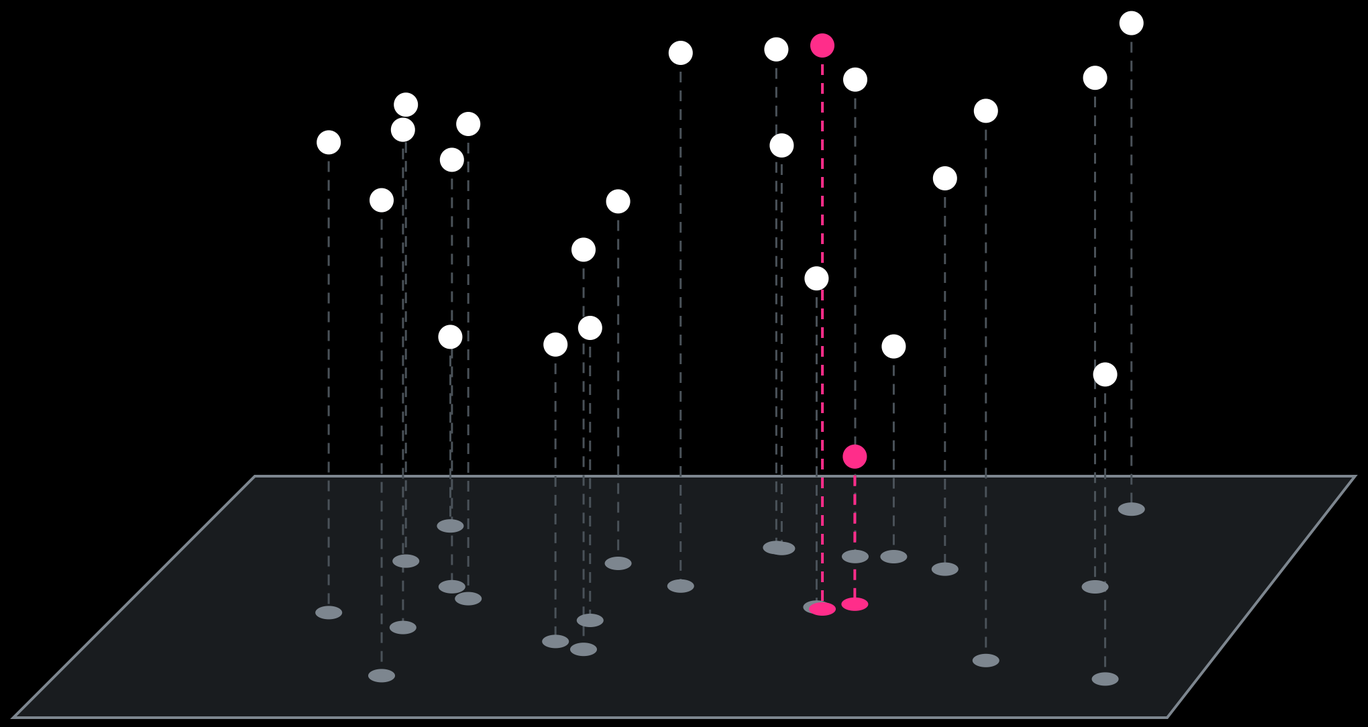
“you shall know a word by the company it keeps.”

John Rupert Firth, 1957

part 5

the model's map

384 directions, drawn on two



principal component analysis:
the two directions in which
the cards differ most

close in the shadow,
far apart in the numbers

only for the picture:
every similarity is still
computed with all 384

★ I can't get home tonight

by meaning

1	My train was cancelled and I missed the last connection.	0.31
2	We landed in Lisbon three hours late.	0.26

---

by keywords

nothing: no card shares a word with the question

one word as the question

river	20	We watched the ducks from the bank of the river.	0.49
	3	The bank raised its interest rates.	0.20
snake	8	A python swallowed a whole rabbit.	0.46
	12	Python is a popular language for data analysis.	0.24

in their sentences, “bank” and “python” have adjusted their numbers

a search by meaning finds texts that share no word with the question.

part 6

what the map does not know

I love this film. · I do not love this film.



0.74

I love this film. · What a wonderful movie!



0.59

I love this film. · I do not love this film.



0.74

I love this film. · What a wonderful movie!



0.59

The dog bit the man. · The man bit the dog.



0.97

Vaccines cause autism. · Vaccines do not cause autism.



0.86

a map of meaning knows what a text is about, but not whether it is true.

part 7

inside the numbers

similarity is one sum

card 1	-0.0189	-0.0445	0.0417	0.0777	0.0544	0.0546	…
card 7	0.1159	0.0211	0.0192	0.0385	0.0392	-0.0357	…

similarity is one sum

card 1	−0.0189	−0.0445	0.0417	0.0777	0.0544	0.0546	…
	×	×	×	×	×	×	
card 7	0.1159	0.0211	0.0192	0.0385	0.0392	−0.0357	…
products	−0.0022	−0.0009	0.0008	0.0030	0.0021	−0.0019	…

similarity is one sum

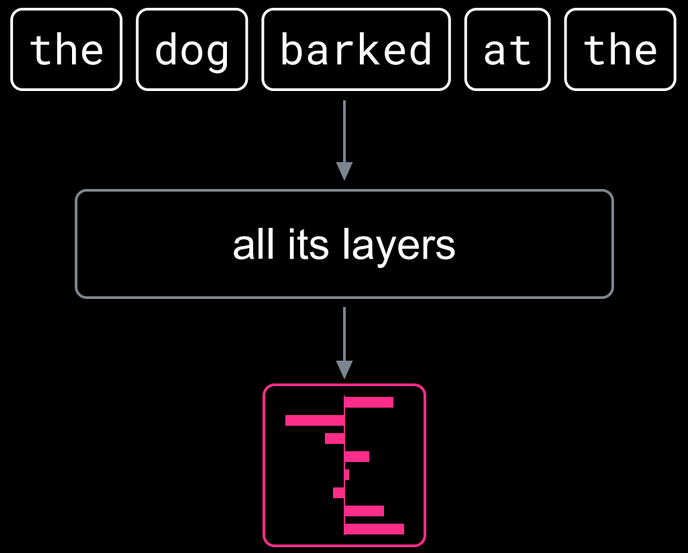
card 1	-0.0189	-0.0445	0.0417	0.0777	0.0544	0.0546	…
	×	×	×	×	×	×	
card 7	0.1159	0.0211	0.0192	0.0385	0.0392	-0.0357	…
products	-0.0022	-0.0009	0.0008	0.0030	0.0021	-0.0019	…

sum of all 384 = 0.32

$$\text{cosine similarity} = a_1b_1 + a_2b_2 + \ldots + a_{384}b_{384}$$

the cosine of the angle between the two arrows, which for arrows of length 1 is just this sum

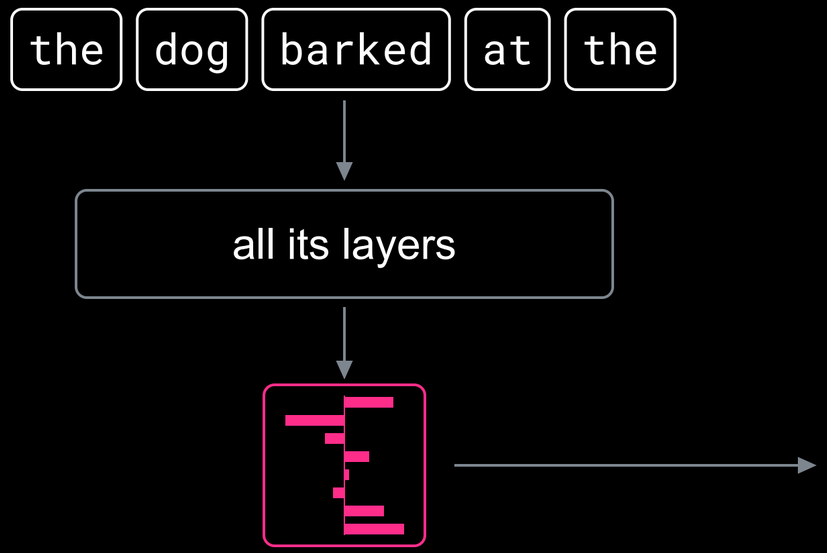
how a language model picks the next token



where the next
token should lie

a sketch, not measured

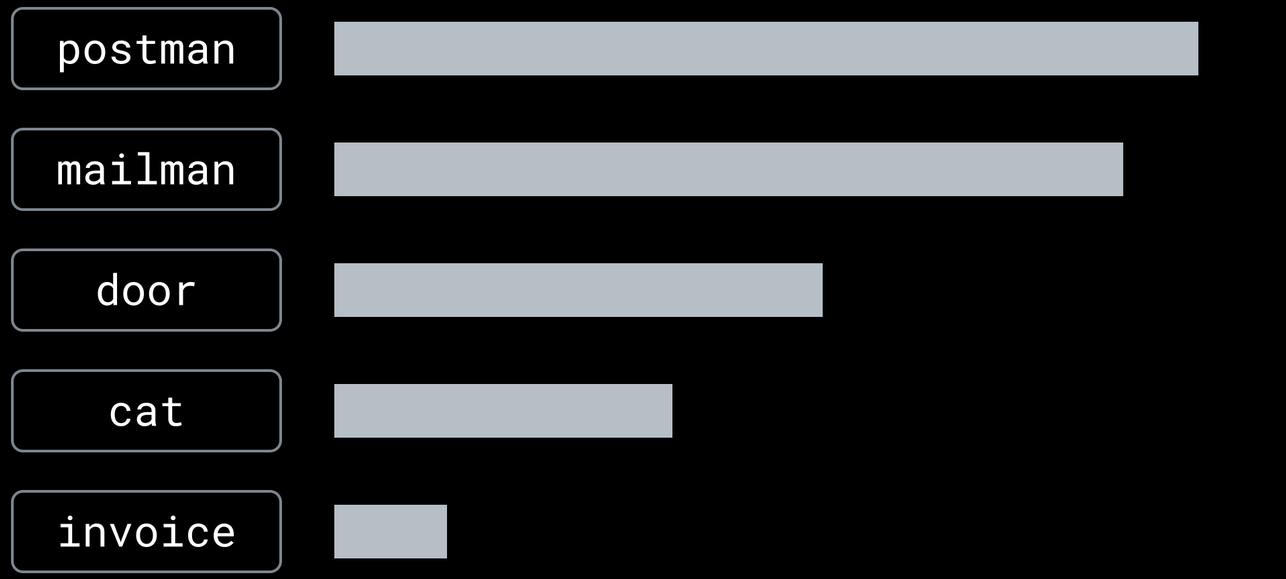
how a language model picks the next token



where the next token should lie

every row of the token table

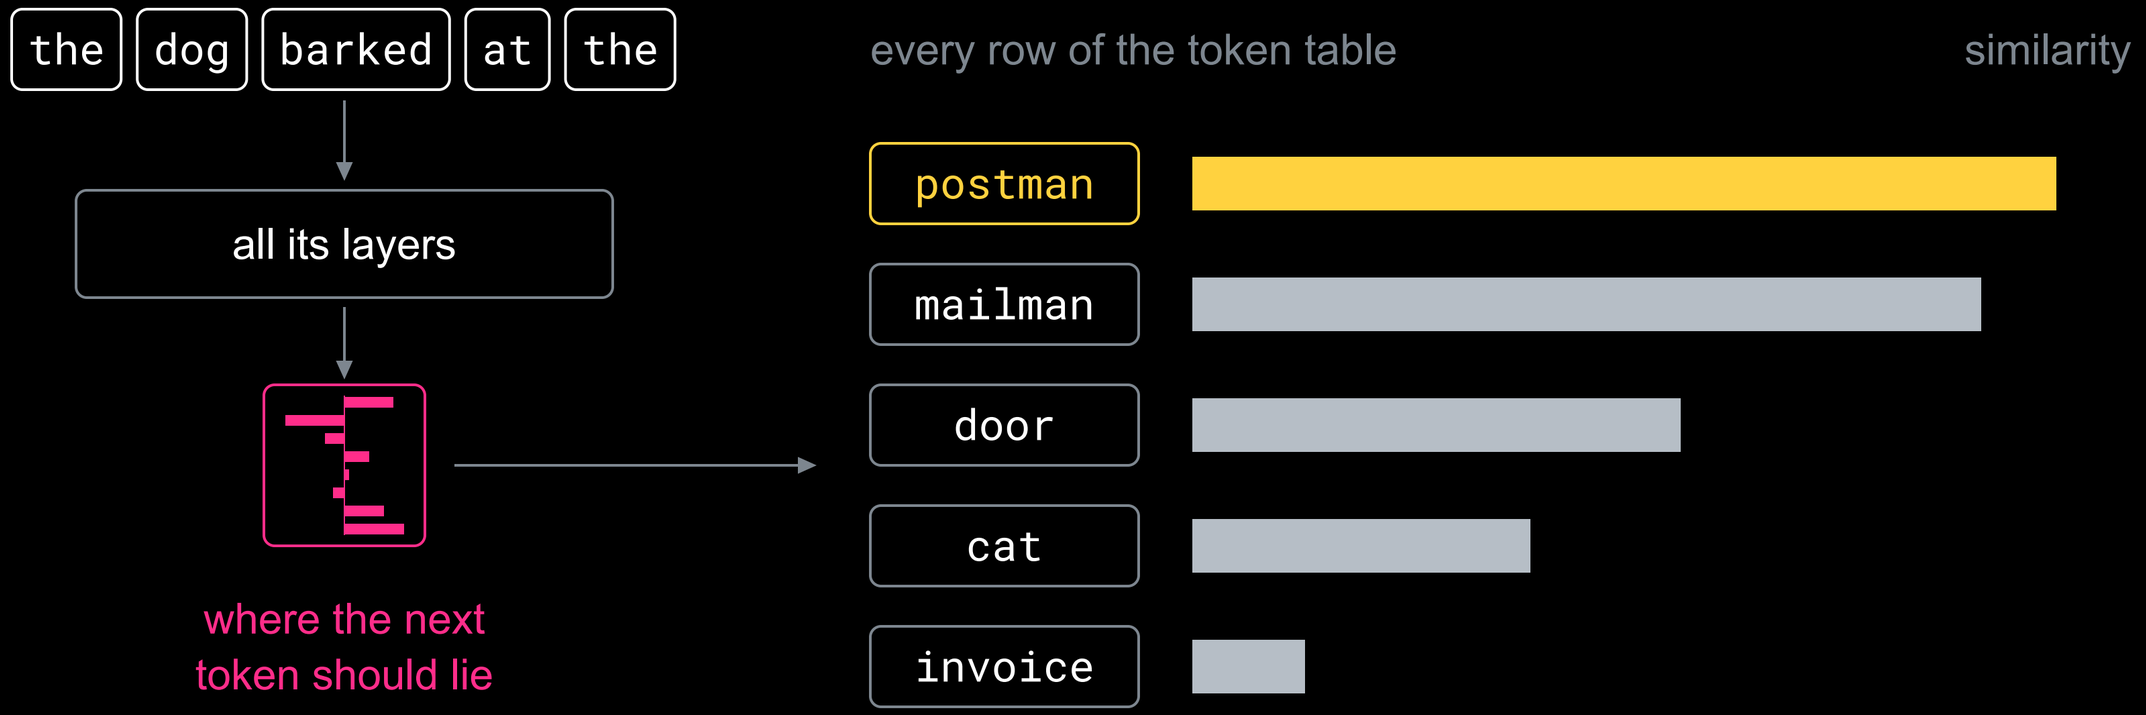
similarity



… one row for every token it knows

a sketch, not measured

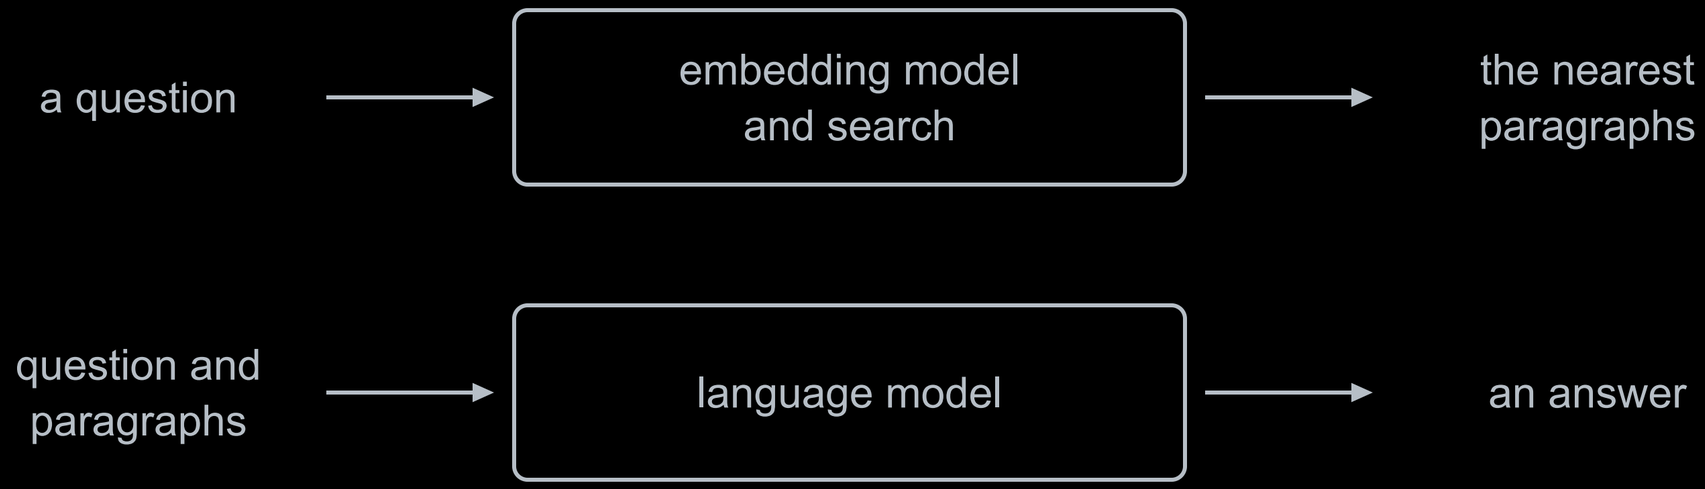
how a language model picks the next token



the dog barked at the postman

a sketch, not measured

a language model predicts where the next token should lie on its map and takes a token that lies close to that place.



retrieval-augmented generation

search · recommendations · duplicate questions