Fodor & Pylyshyn 1998
Lake & Baroni 2018

Jacob Andreas / MIT 6.884 / Fall 2020

Today

1. F&P: Are there fundamental differences between symbolist / classical accounts of information processing and connectionist / neural ones?

2. How much progress have neural models made towards addressing the concerns raised by F&P?

The research program

F&P:

“The architecture of the cognitive system consists of the set of basic operations, resources, functions, principles, etc (generally the sorts of properties that would be described in a “user’s manual” for that architecture if it were available on a computer), whose domain and range are the representational states of the organism.

It follows that, if you want to make good the Connectionist theory as a theory of cognitive architecture, you have to show that the processes which operate on the representational states of an organism are those which are specified by a Connectionist architecture.”

Historical context

Smolensky 1988:

"Higher-level analyses [of] connectionist models reveal subtle relations to symbolic models. [...] At the lower level, compntation has the character of massively parallel satisfaction of soft numerical constraints; at the higher level, this can lead to competence characterizable by hard rules. Performance will typically deviate from this competence since behavior is achieved not by interpreting hard rules but by satisfying soft constraints."

Historical context

Rumelhart & McClelland 1985:

“Children are typically said to pass through a three-phase acquisition process in which they first learn past tense by rote, then learn the past tense rule and overregularize, and then finally learn the exceptions to the rule. We show that the acquisition data can be accounted for in more detail by dispensing with the assumption that the child [eamns rules and substituting in its place a simple homogeneous learning procedure. We show how ‘rule-like’ behavior can emerge from the interactions among a network of units encoding the root form to past tense mapping.”

The research program

F&P: “Not so fast!

Specific aspects of human *mental representations* and *information processing* seem poorly captured by current connectionist models.”

F&P's argument

1a. Classical representations have combinatorial syntax & semantics; connectionist ones **cannot**.

F&P's argument

1a. Classical representations have combinatorial syntax & semantics; connectionist ones **cannot**.

1b. Classical information processing operations are sensitive to structure; connectionist ones **are not**.

F&P's argument

2a. Human language (& thought?) are productive, which **requires** structure sensitivity and combinatoriality.

F&P's argument

2a. Human language (& thought?) are productive, which **requires** structure sensitivity and combinatoriality.

2b. Ditto for **systematicity** rather than productivity.

F&P's argument

∴ Connectionist models cannot model human language (/ thought).

(But classical models probably can.)

Discussion

Combinatorial structure

Sample task

The cat is on the mat. →

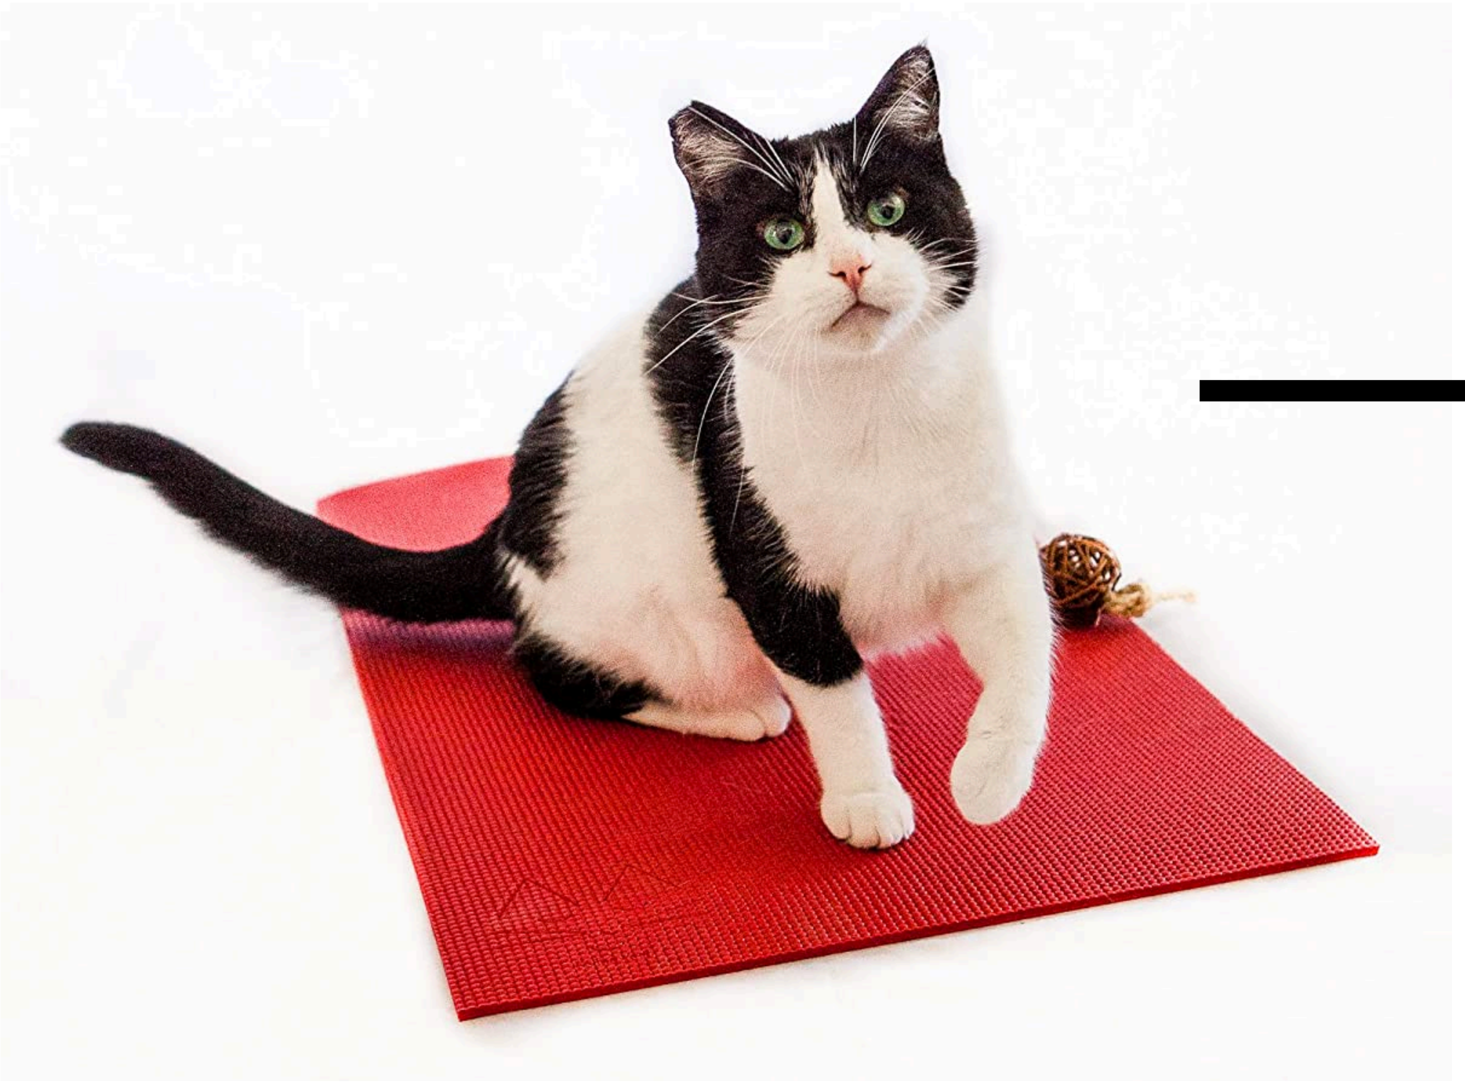 →

→ True

[https://www.amazon.in/Feline-Yogi-Original-Yoga-Cat]

Sample task

The fox is in a box.

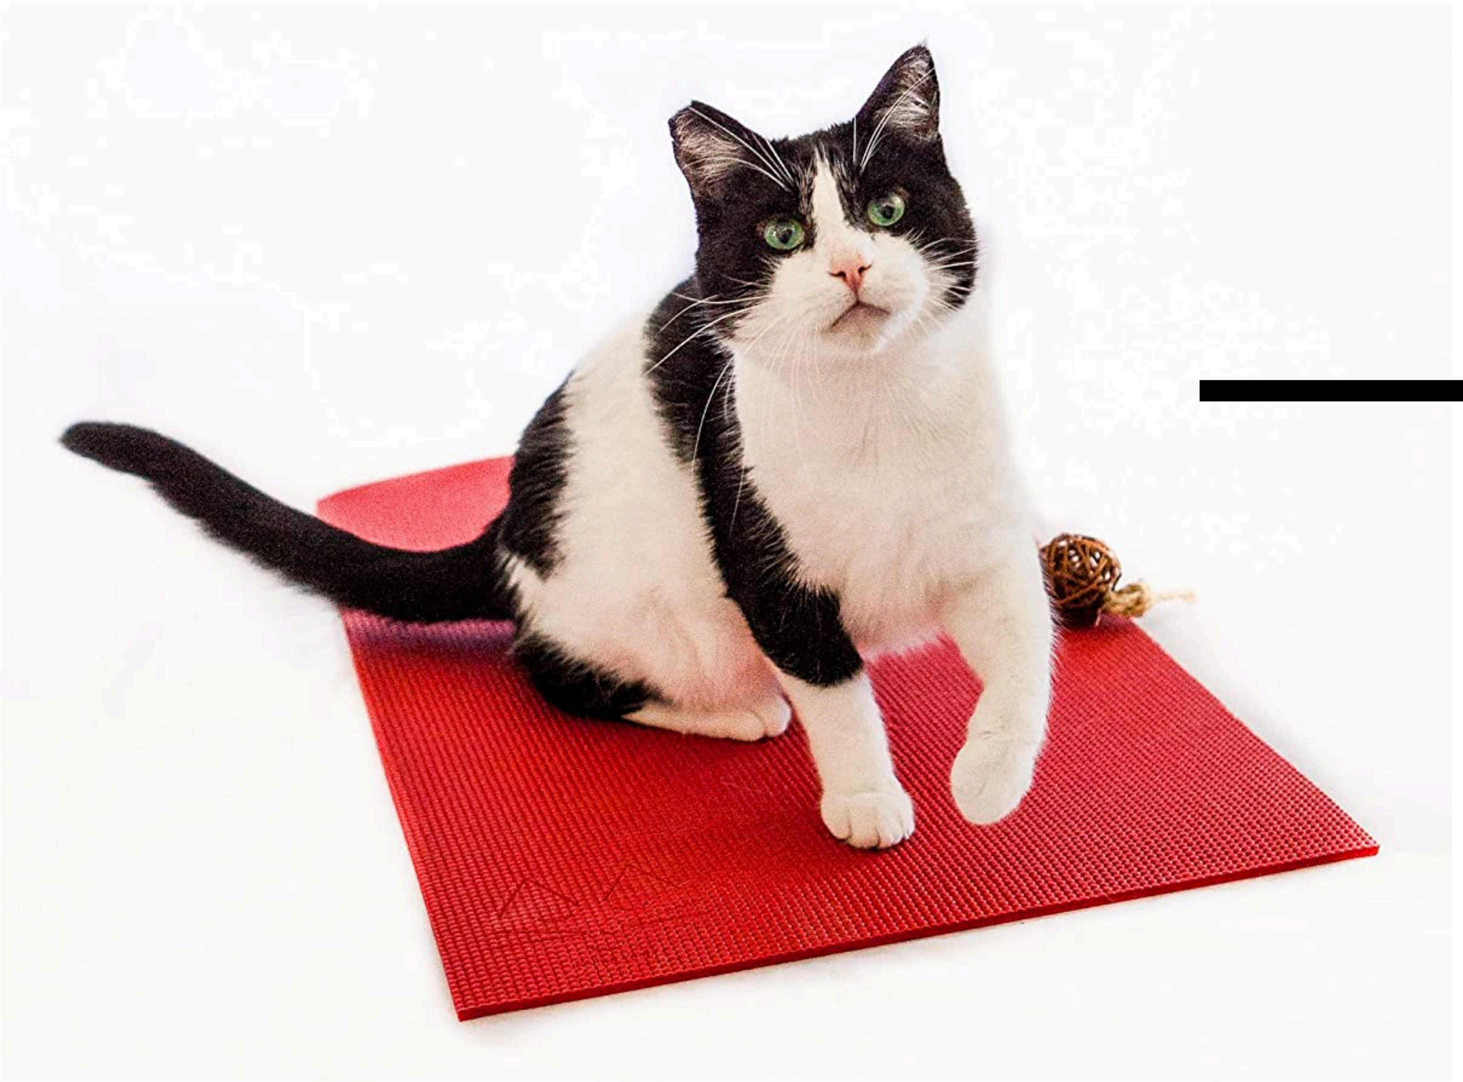

False

[https://www.amazon.in/Feline-Yogi-Original-Yoga-Cat]

A classical implementation

The cat is on the mat. → [[*The cat*] [*is* [*on the mat*]]]

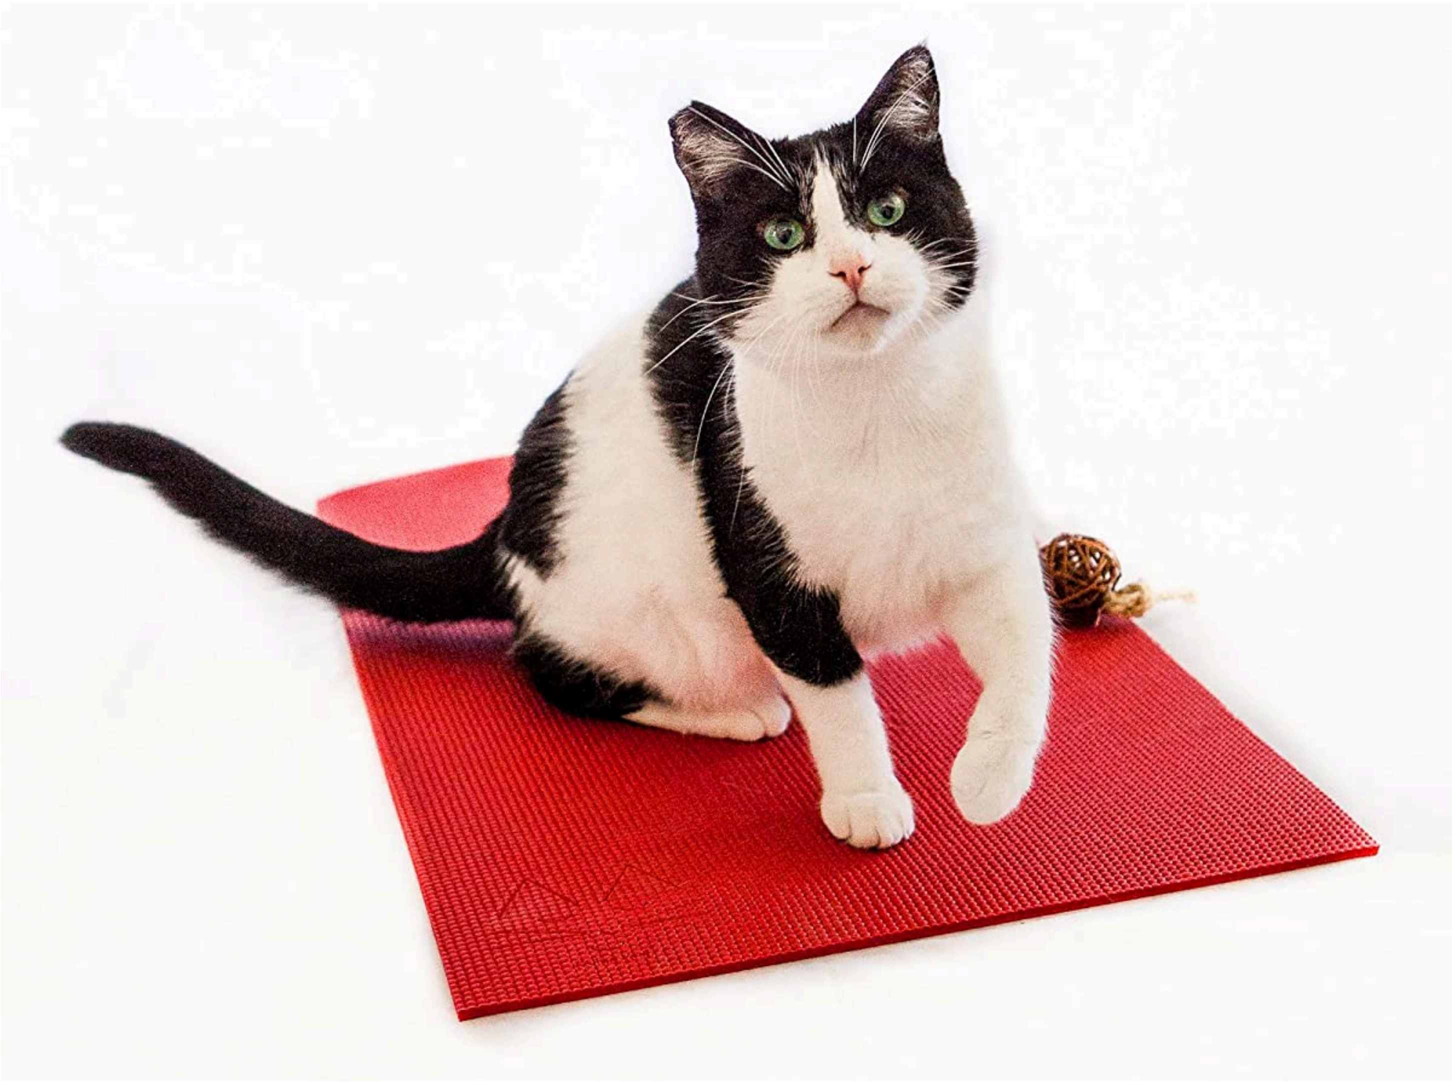

[https://www.amazon.in/Feline-Yogi-Original-Yoga-Cat]

A classical implementation

The cat is on the mat. → [[*The cat*] [*is* [*on the mat*]]]

↓

`cat(x), mat(y), on(x, y)`

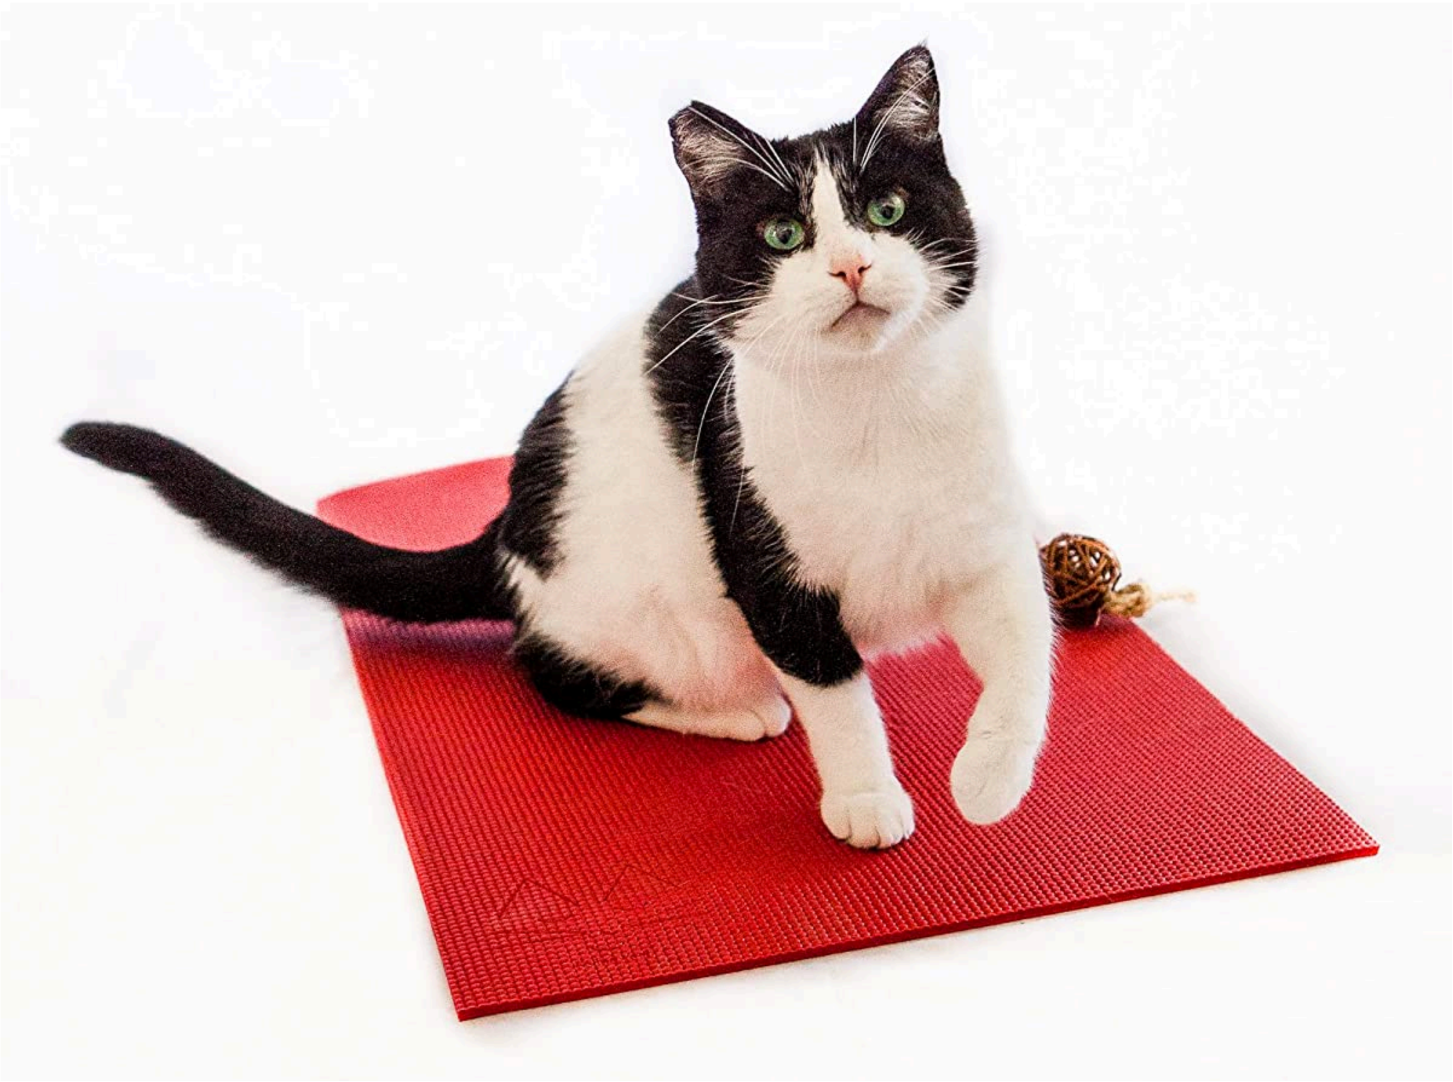

[https://www.amazon.in/Feline-Yogi-Original-Yoga-Cat]

A classical implementation

The cat is on the mat. → [[*The cat*] [*is* [*on the mat*]]]

↓

```
cat(x), mat(y), on(x, y)
```

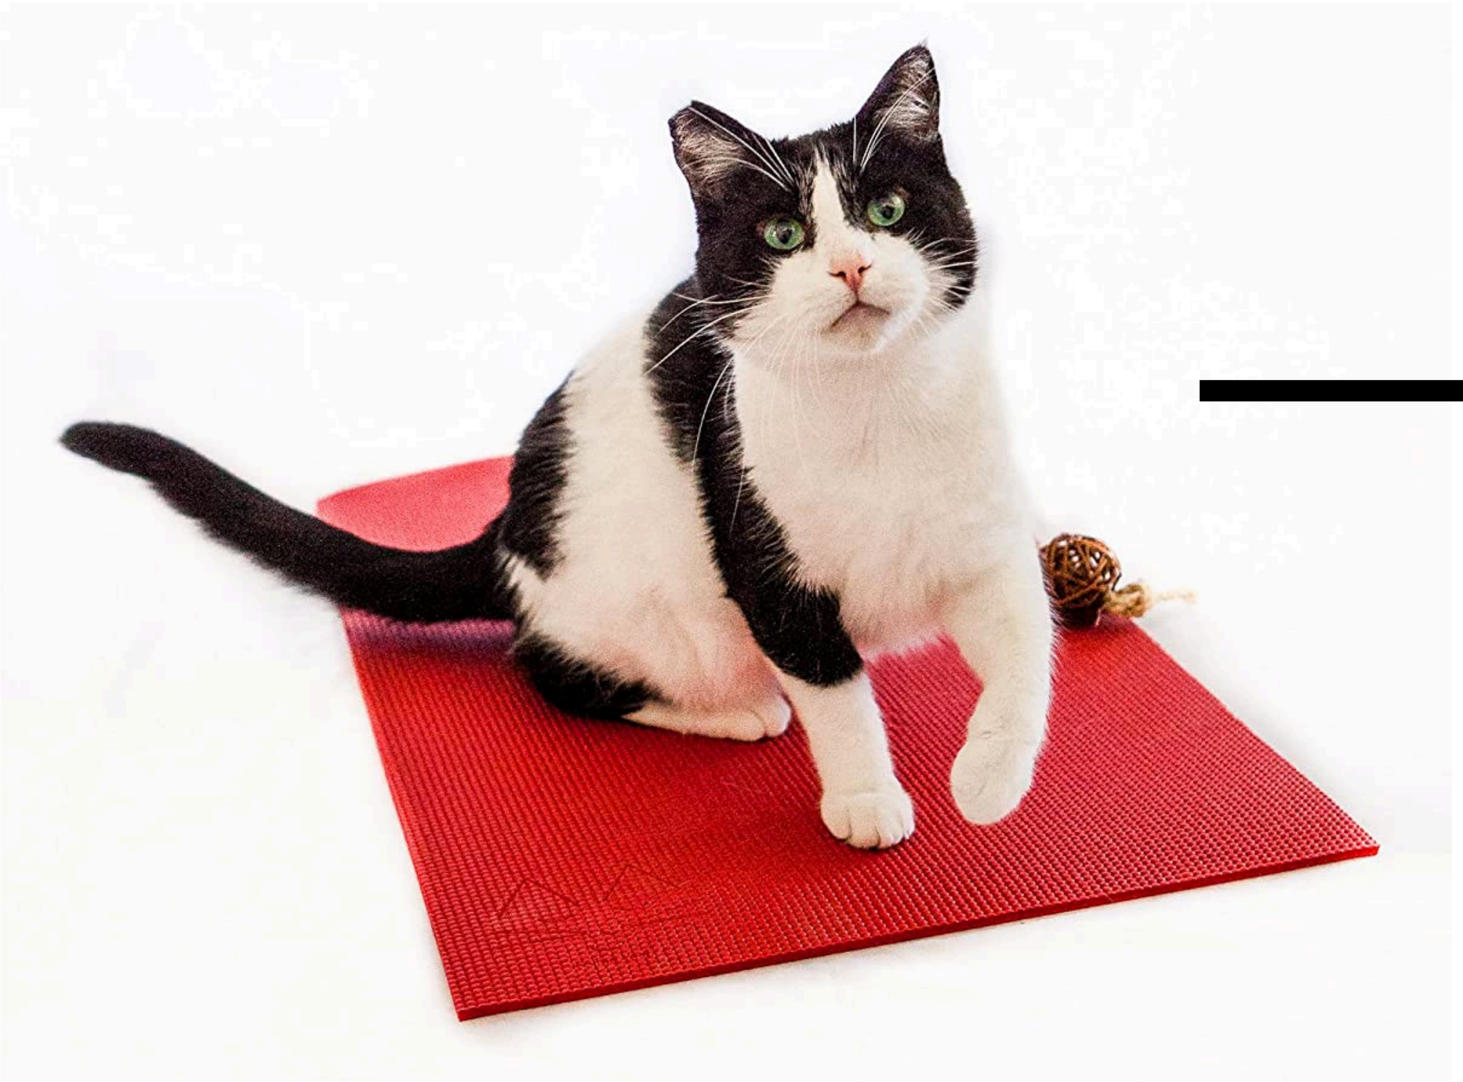

→

```
cat(x)
mat(y)
red(y)
on(x, y)
…
```

[https://www.amazon.in/Feline-Yogi-Original-Yoga-Cat]

A classical implementation

The cat is on the mat. → [[*The cat*] [*is* [*on the mat*]]]

↓

`cat(x), mat(y), on(x, y)`

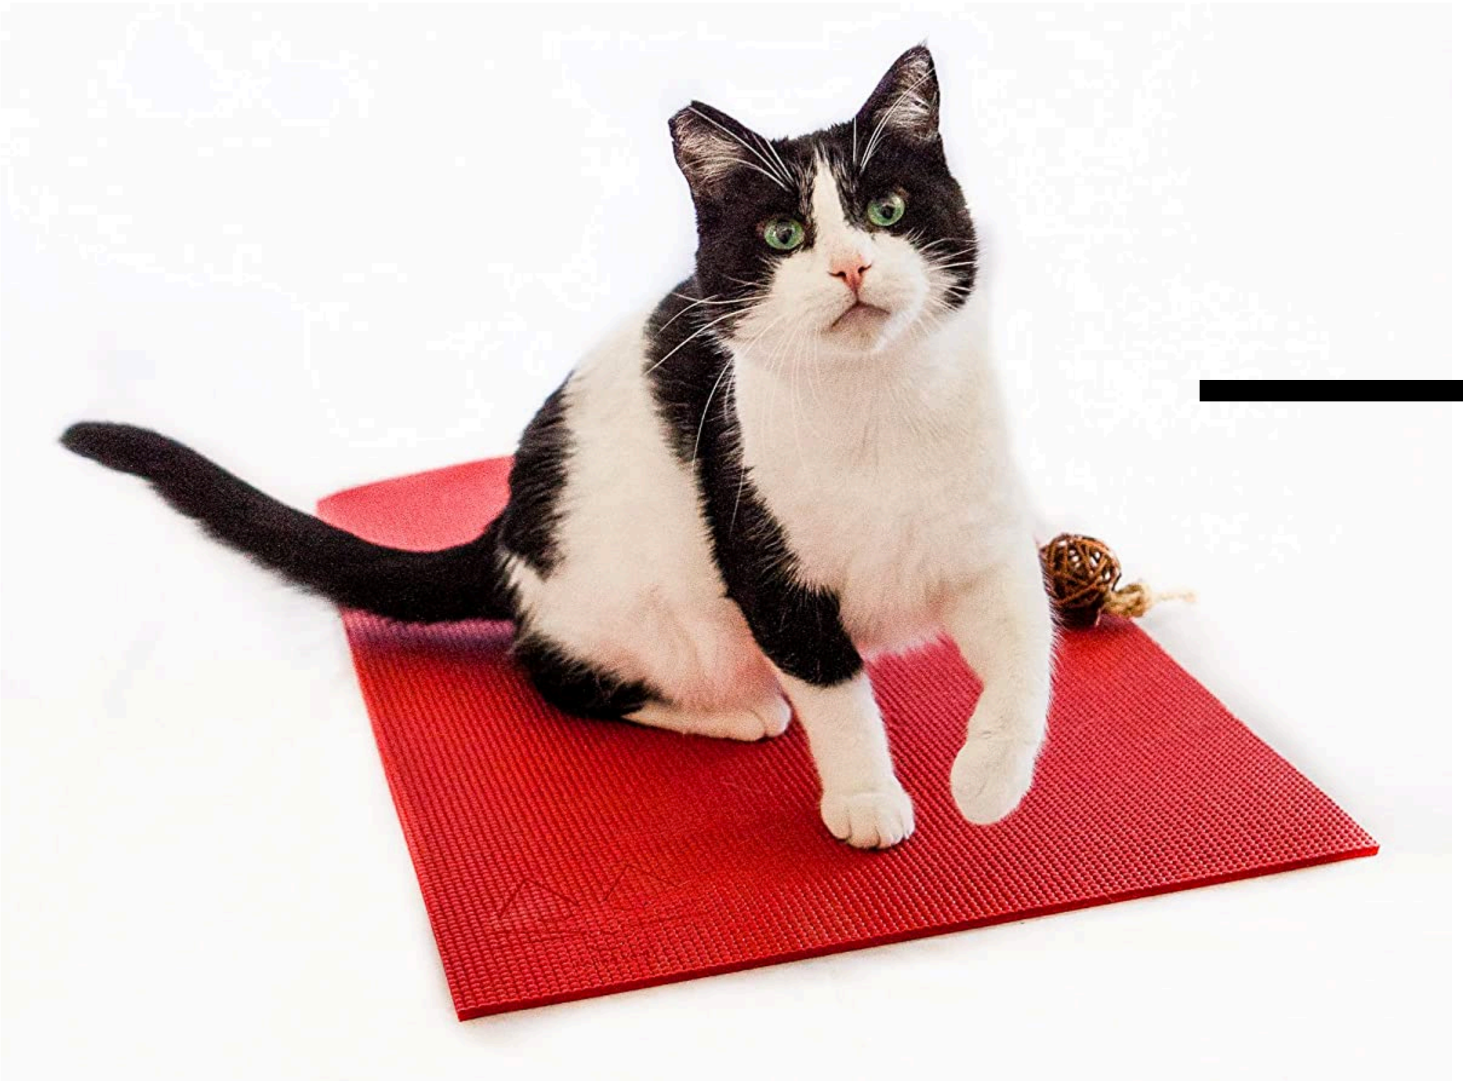

→

```
cat(x)
mat(y)
red(y)
on(x, y)
…
```

→ $\vdash$ → **True**

[https://www.amazon.in/Feline-Yogi-Original-Yoga-Cat]

A connectionist implementation

The cat is on the mat.

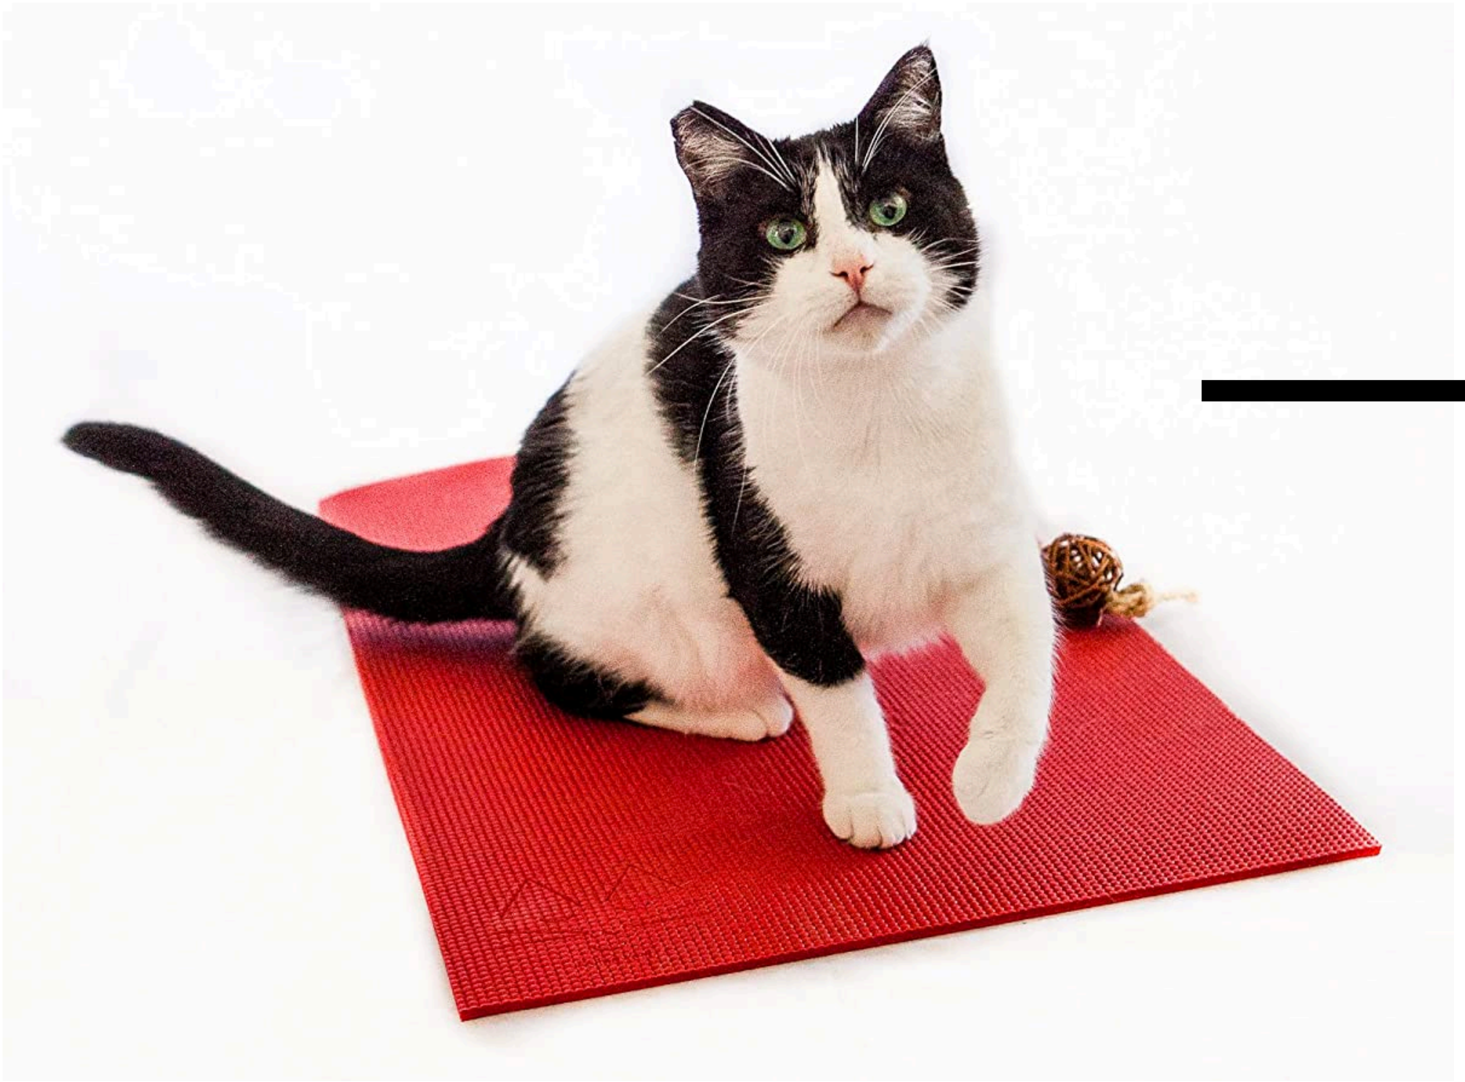

A connectionist implementation

The cat is on the mat.

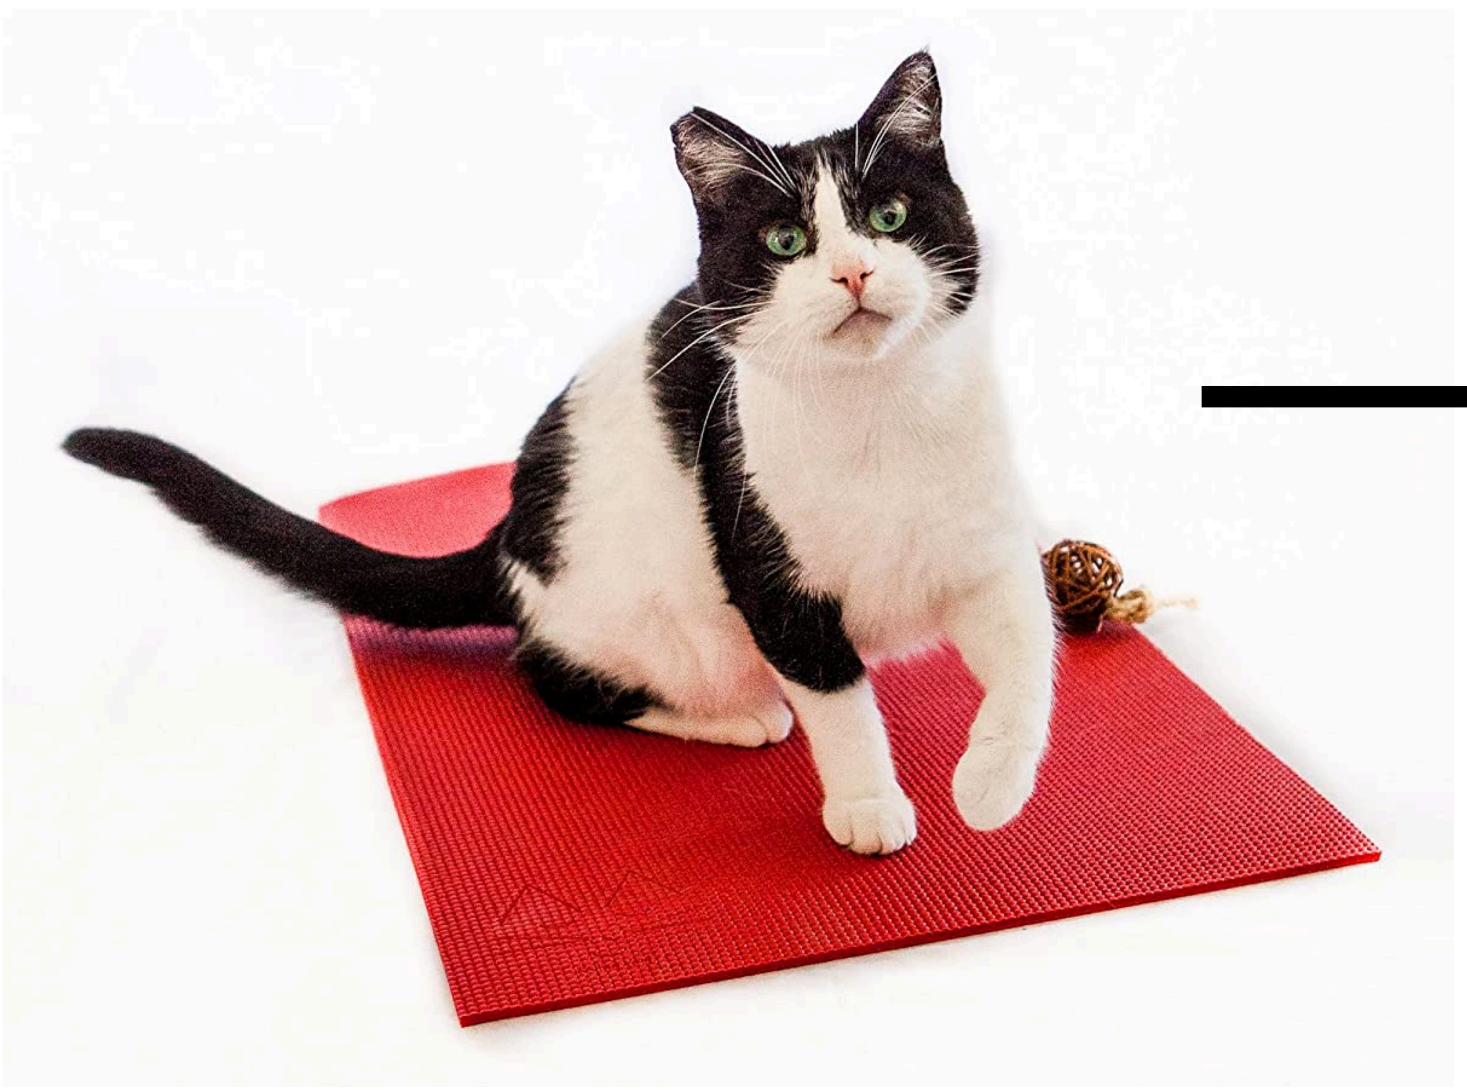

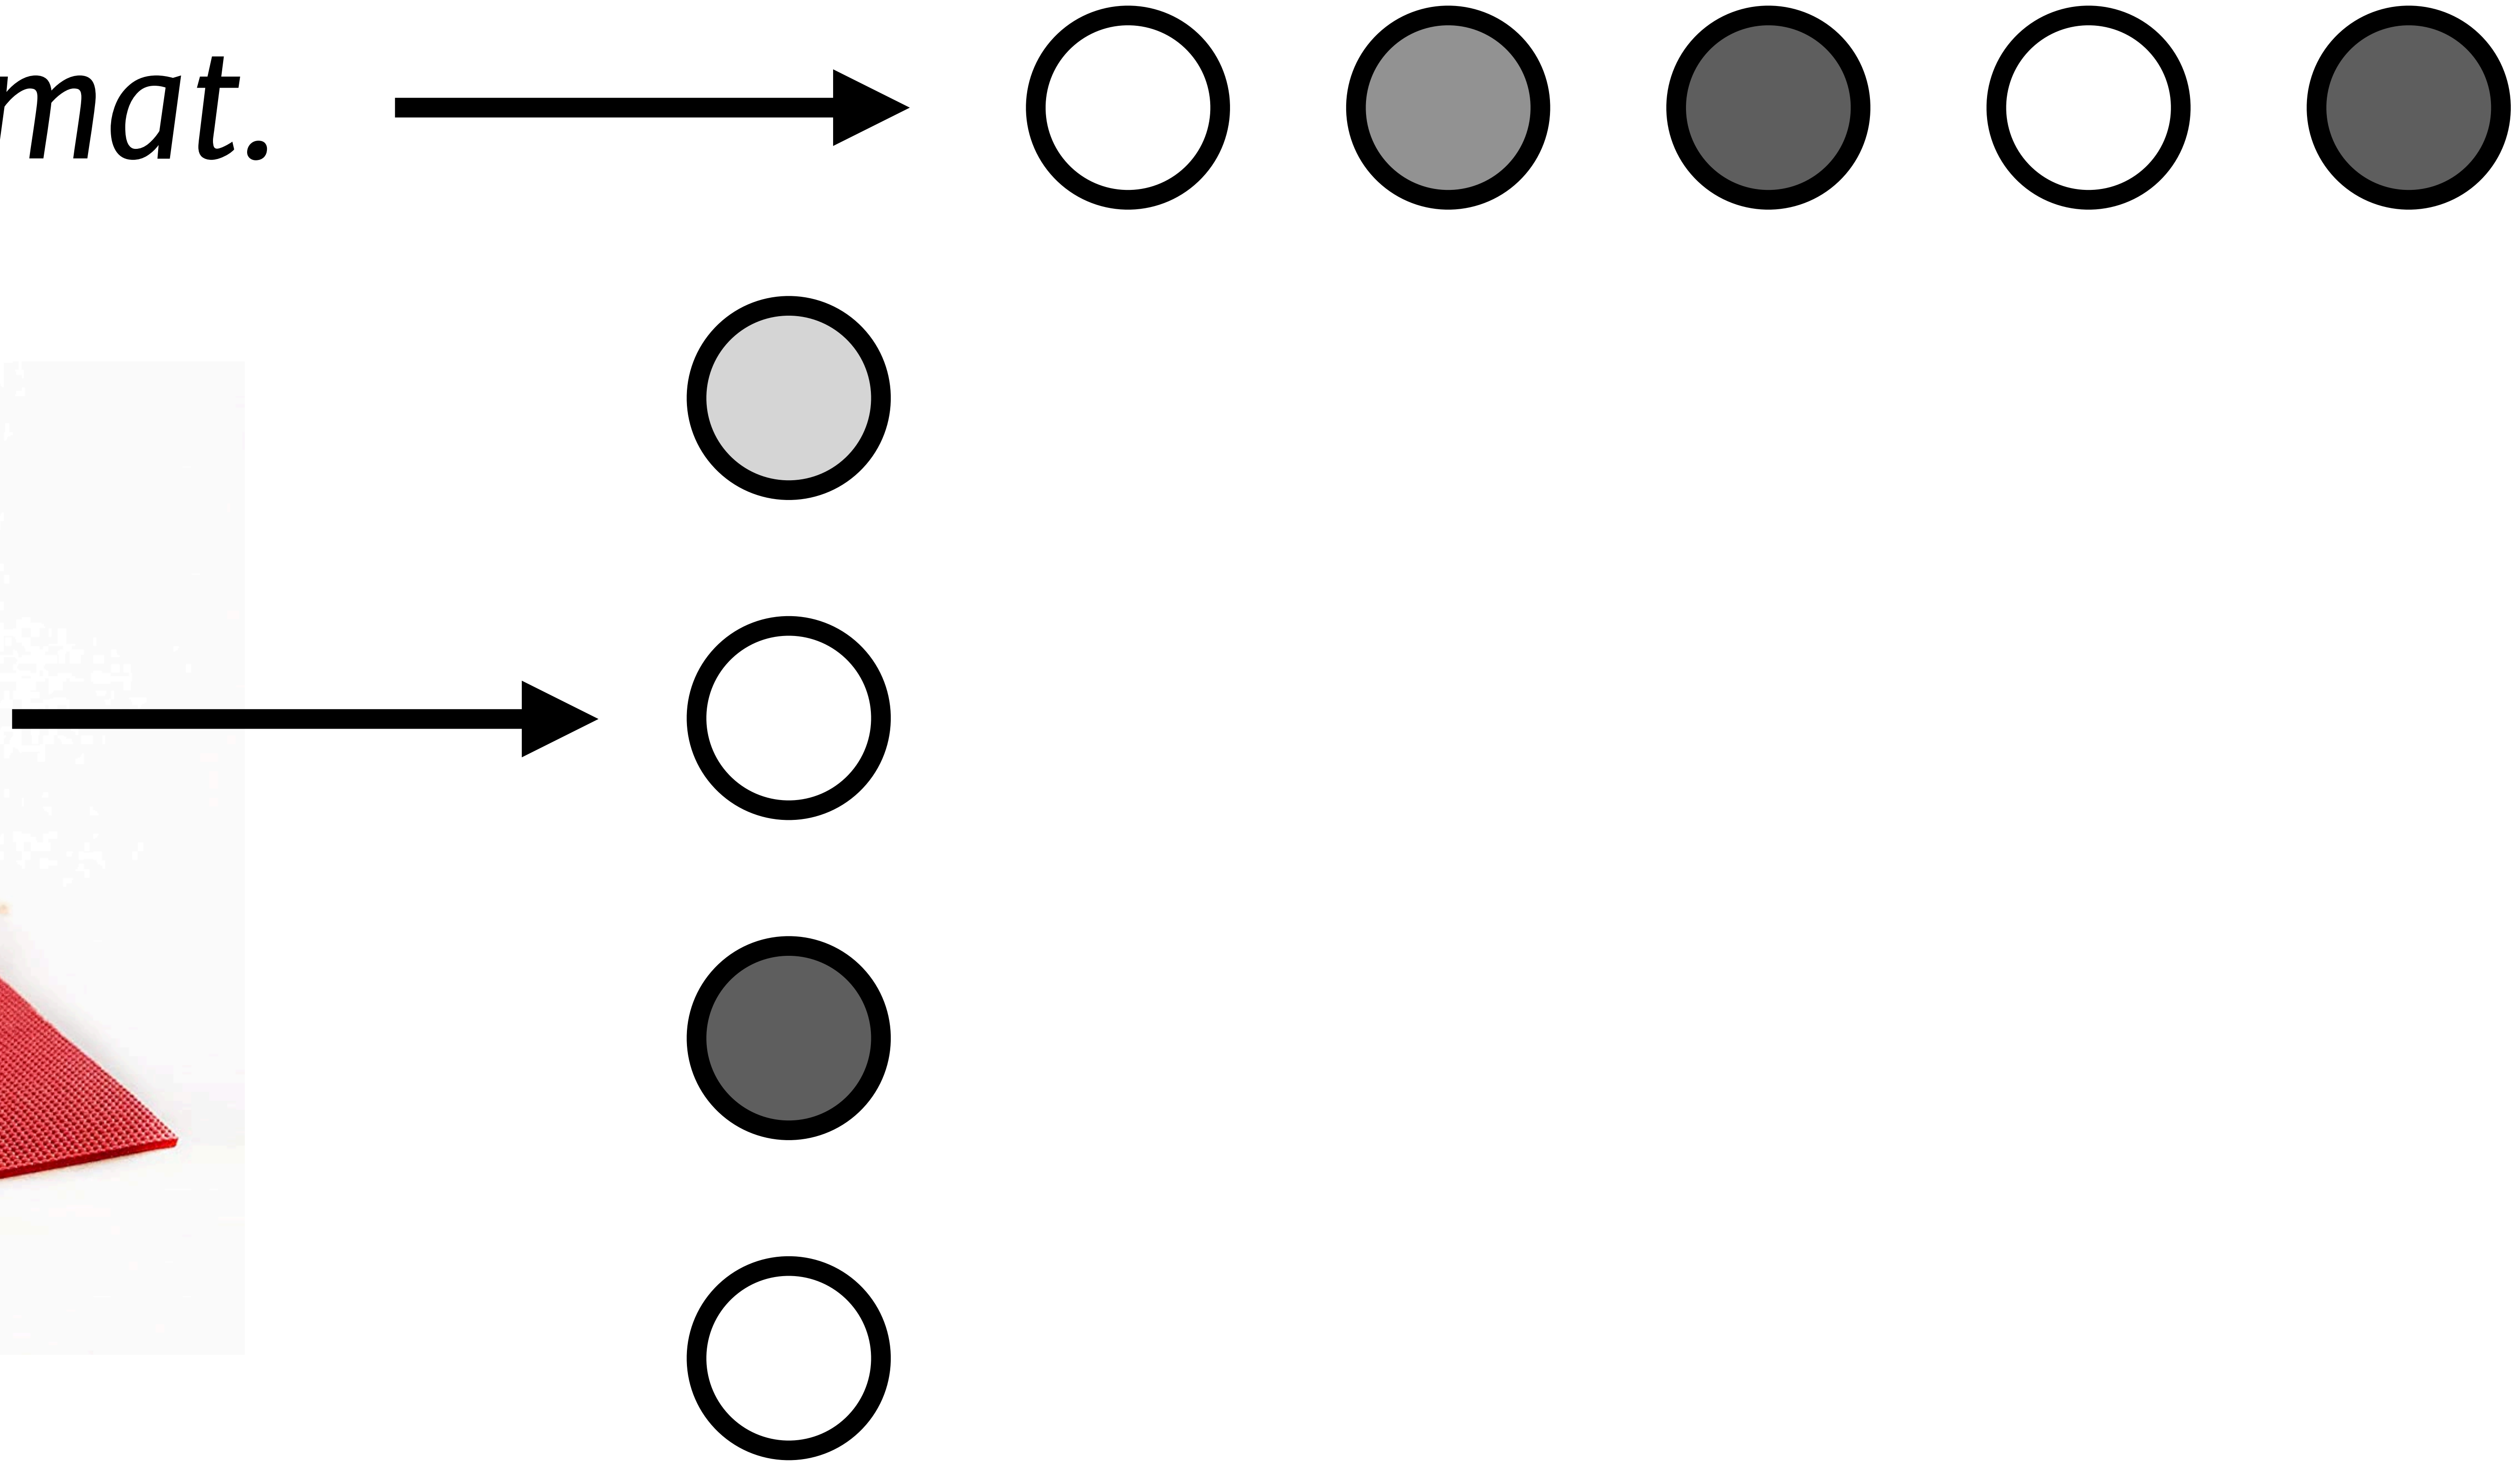

A connectionist implementation

The cat is on the mat.

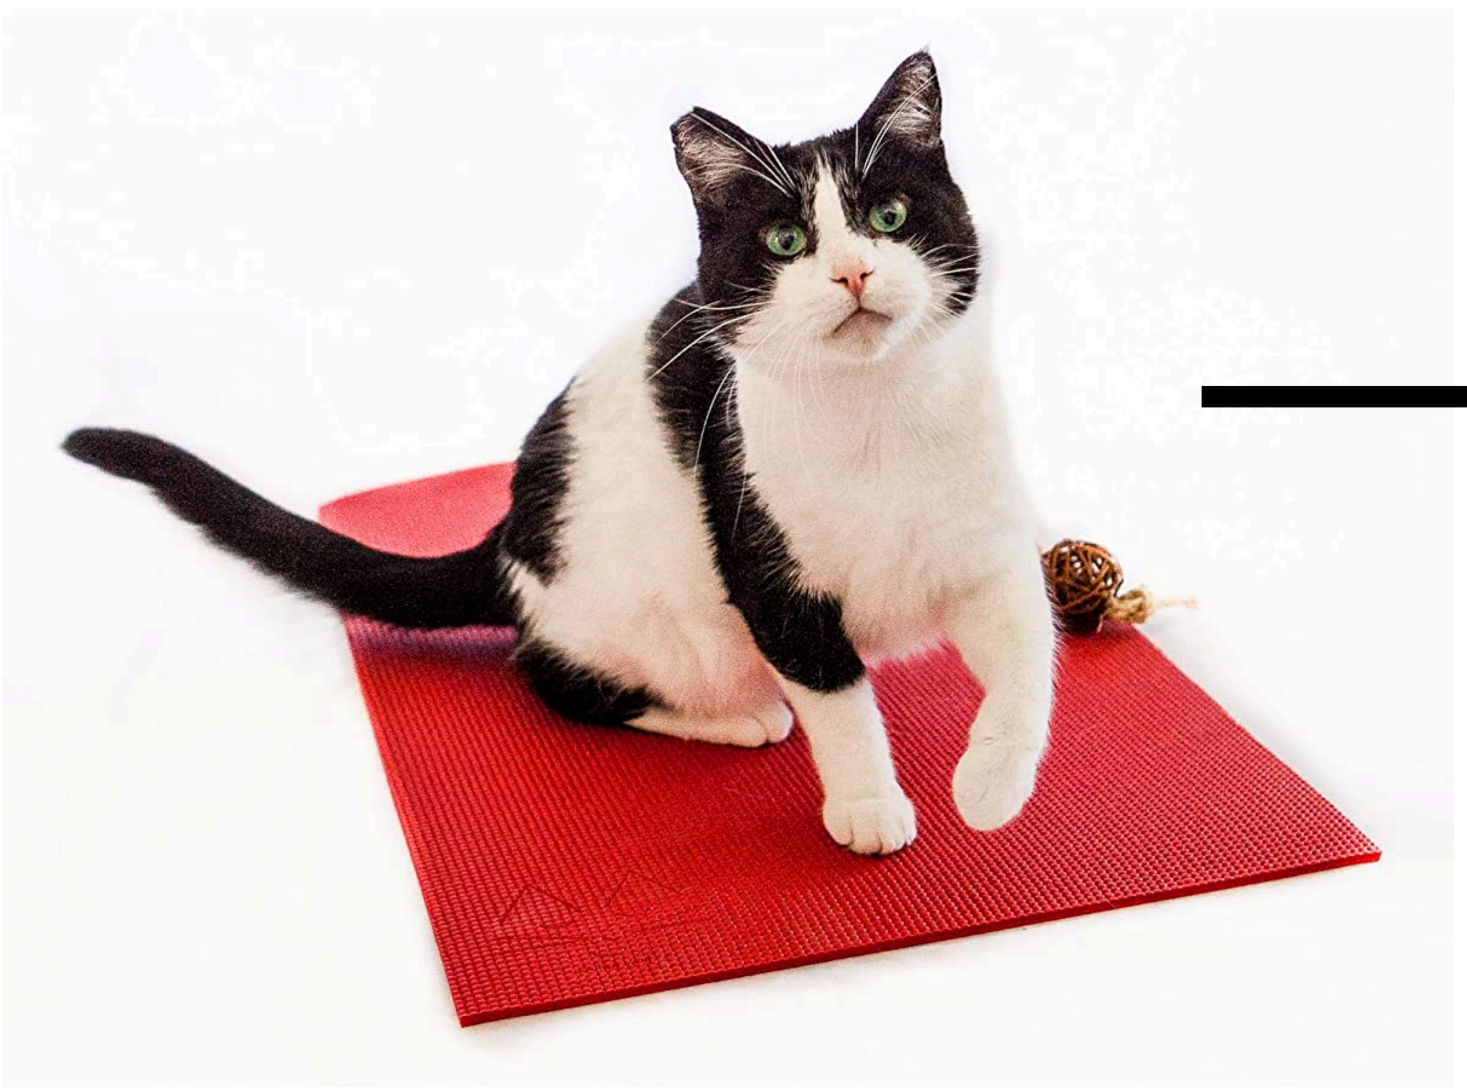

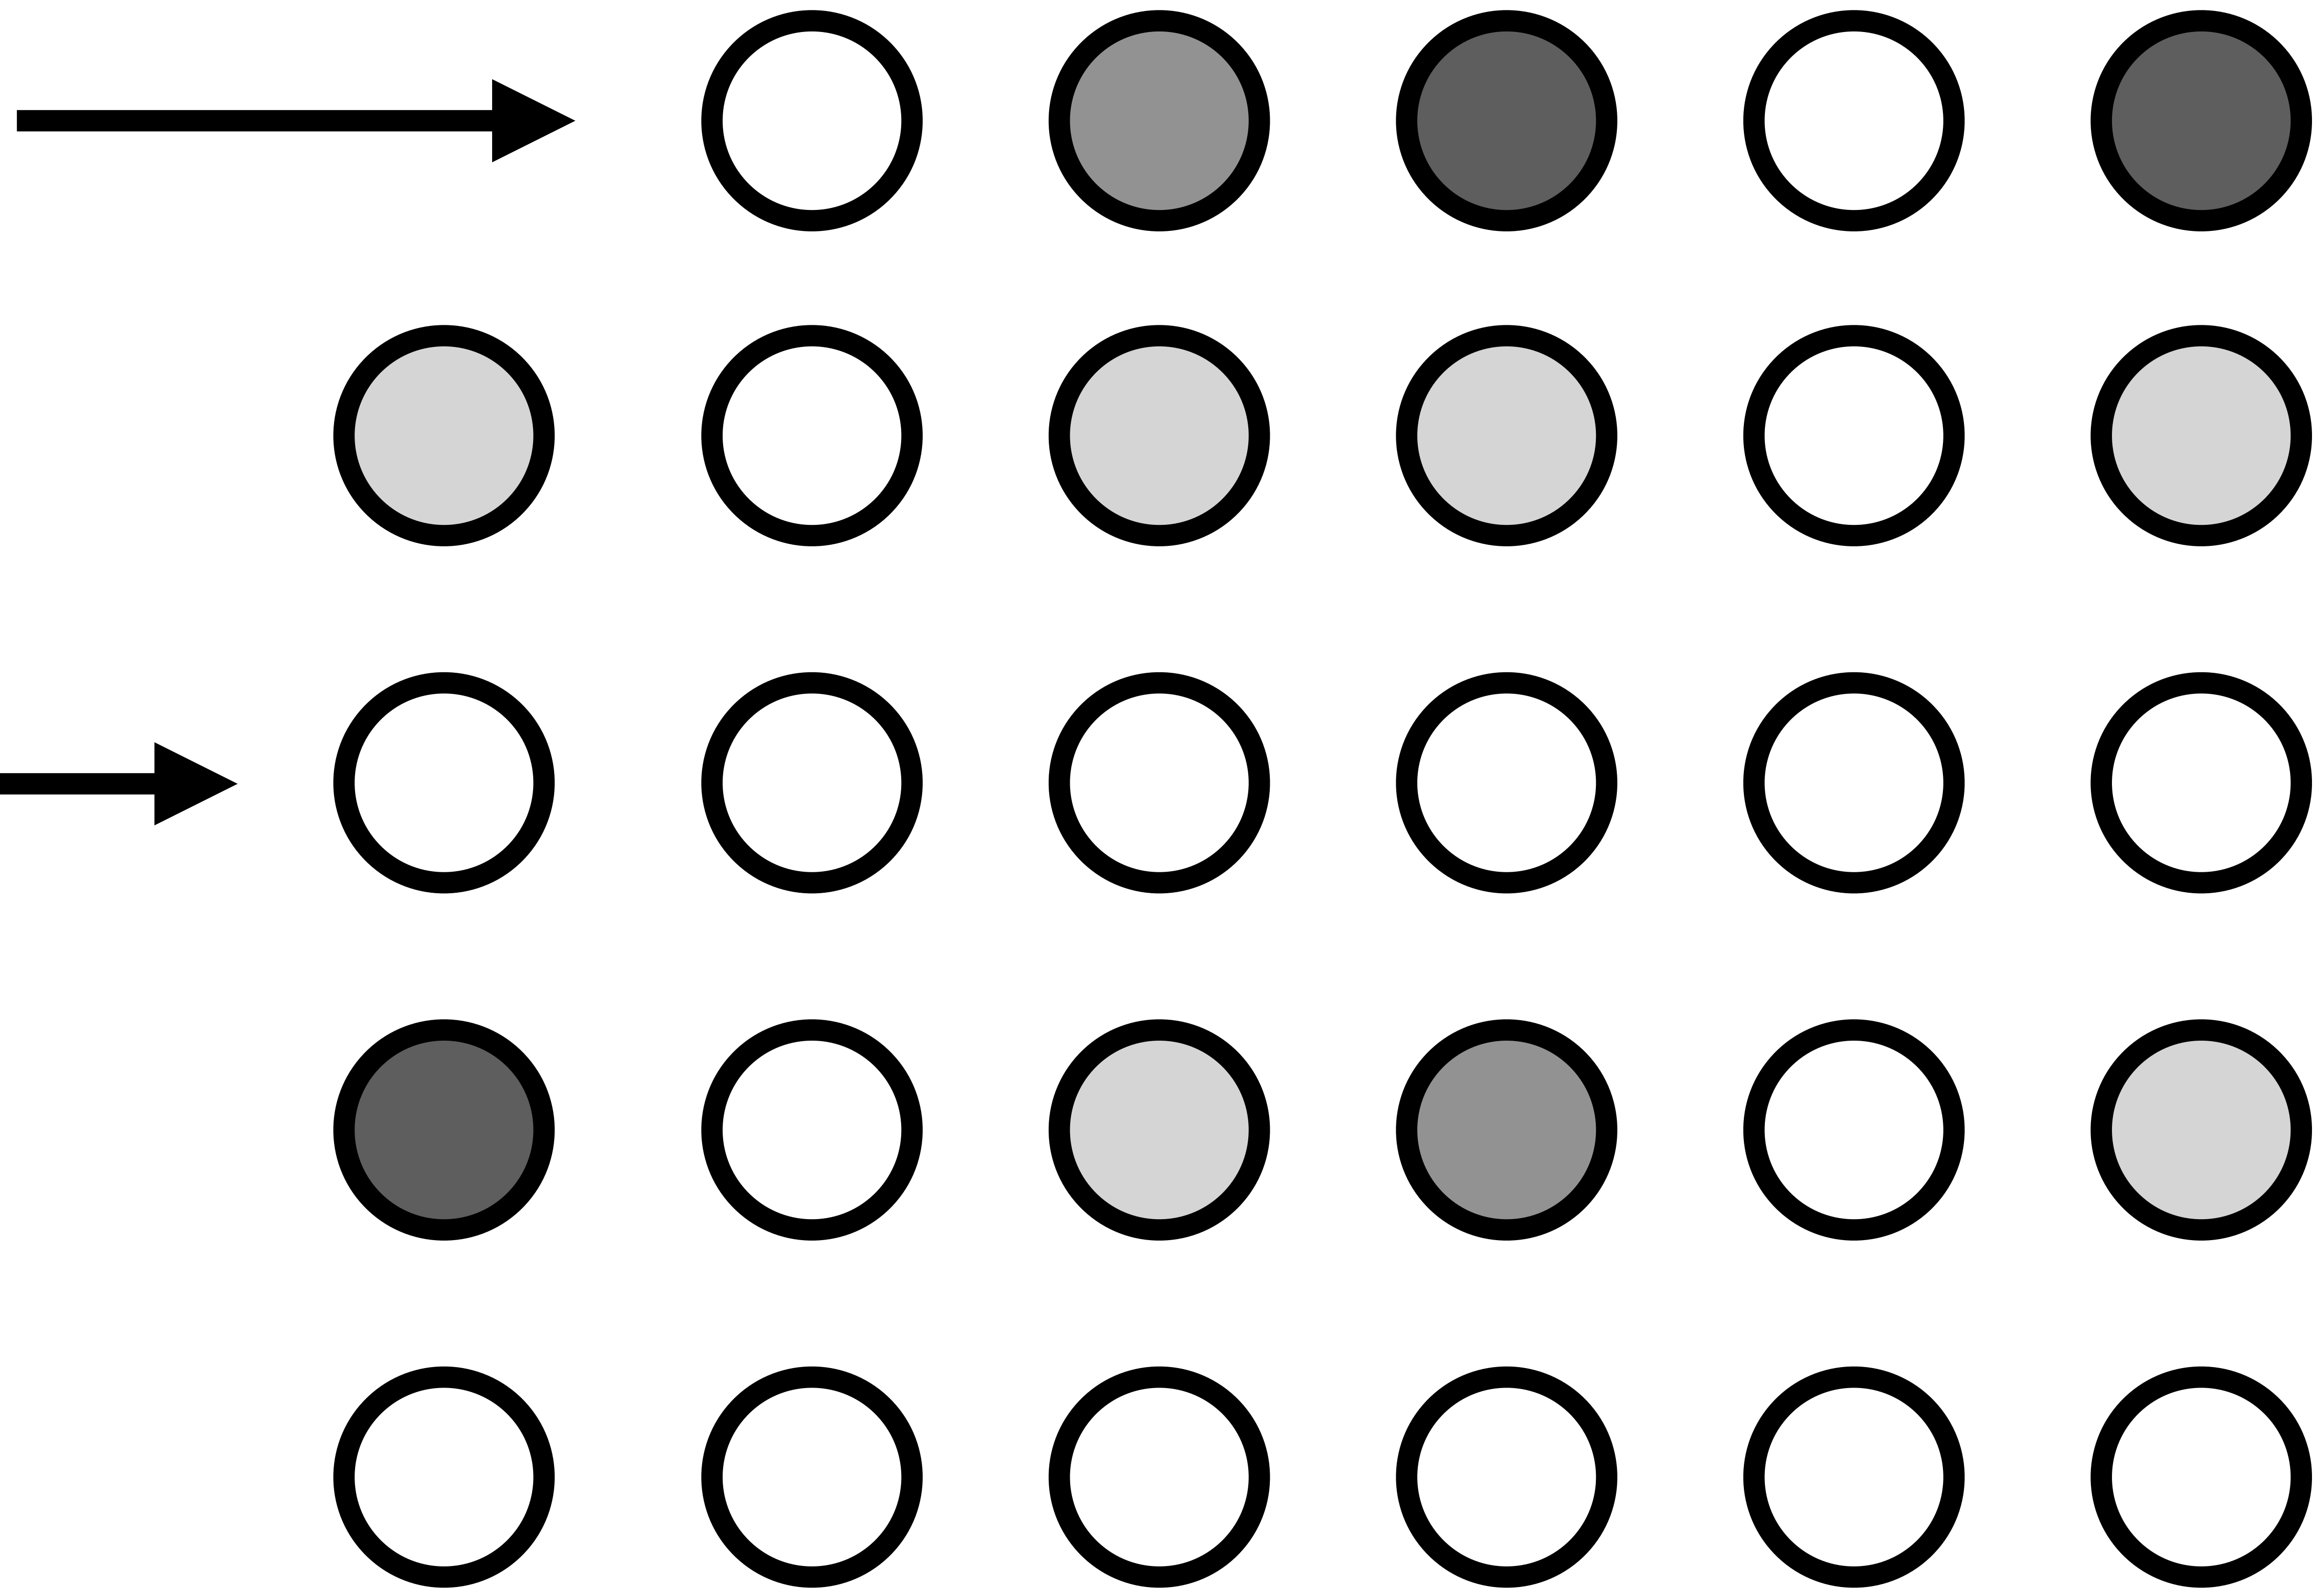

A connectionist implementation

The cat is on the mat.

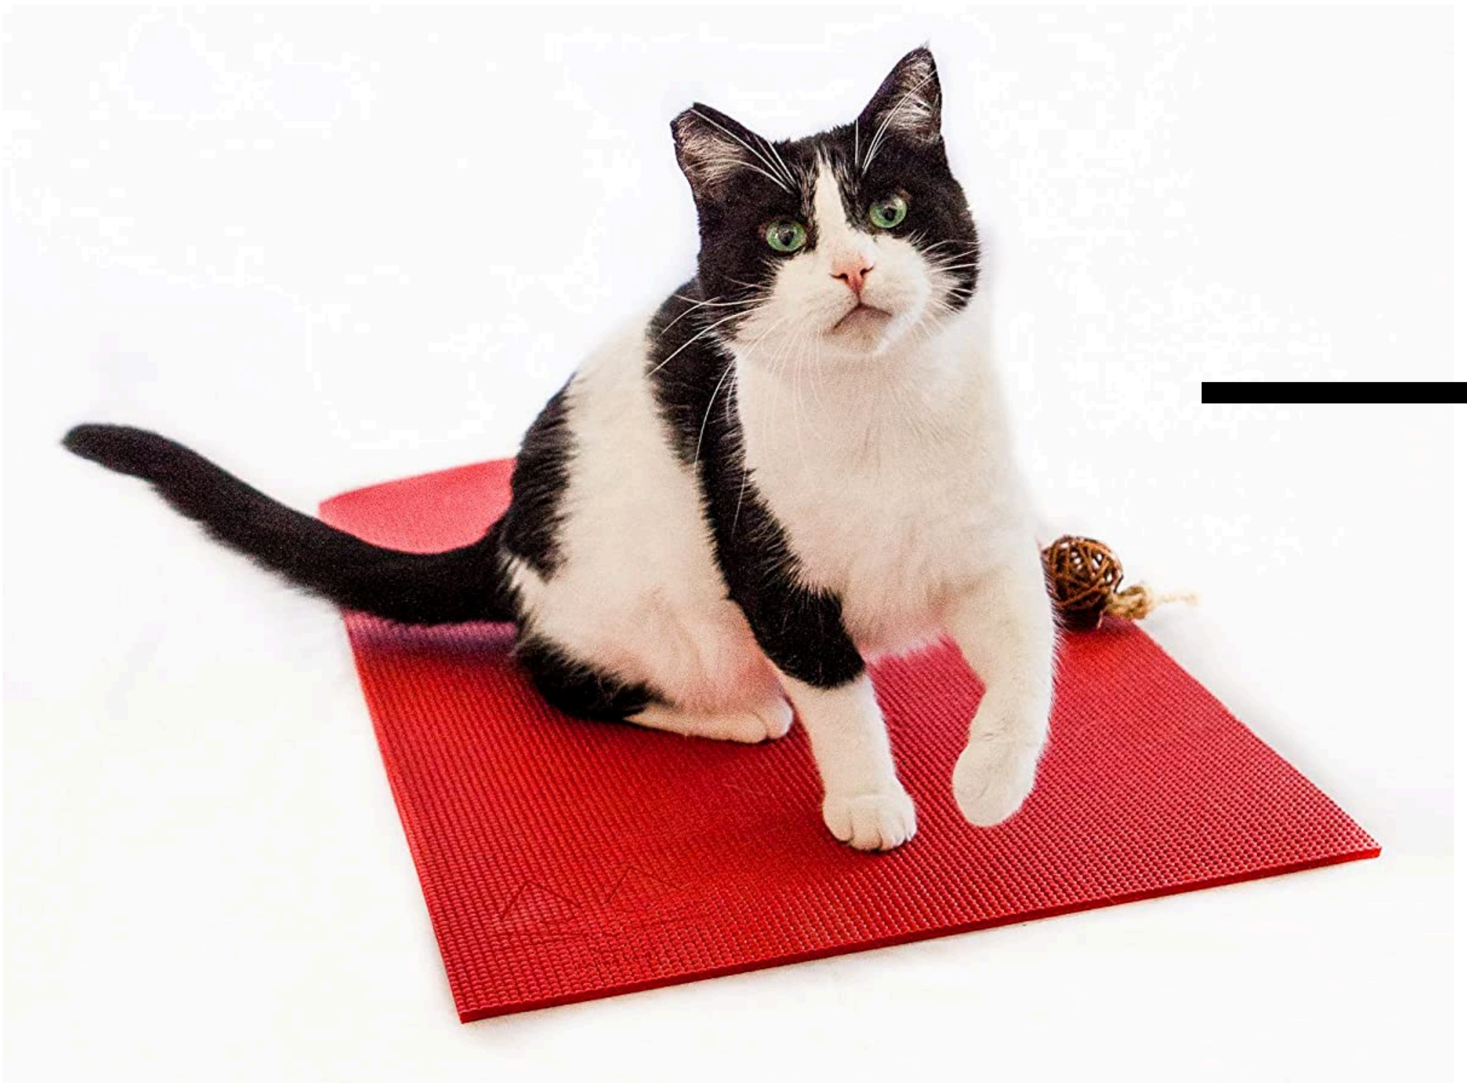

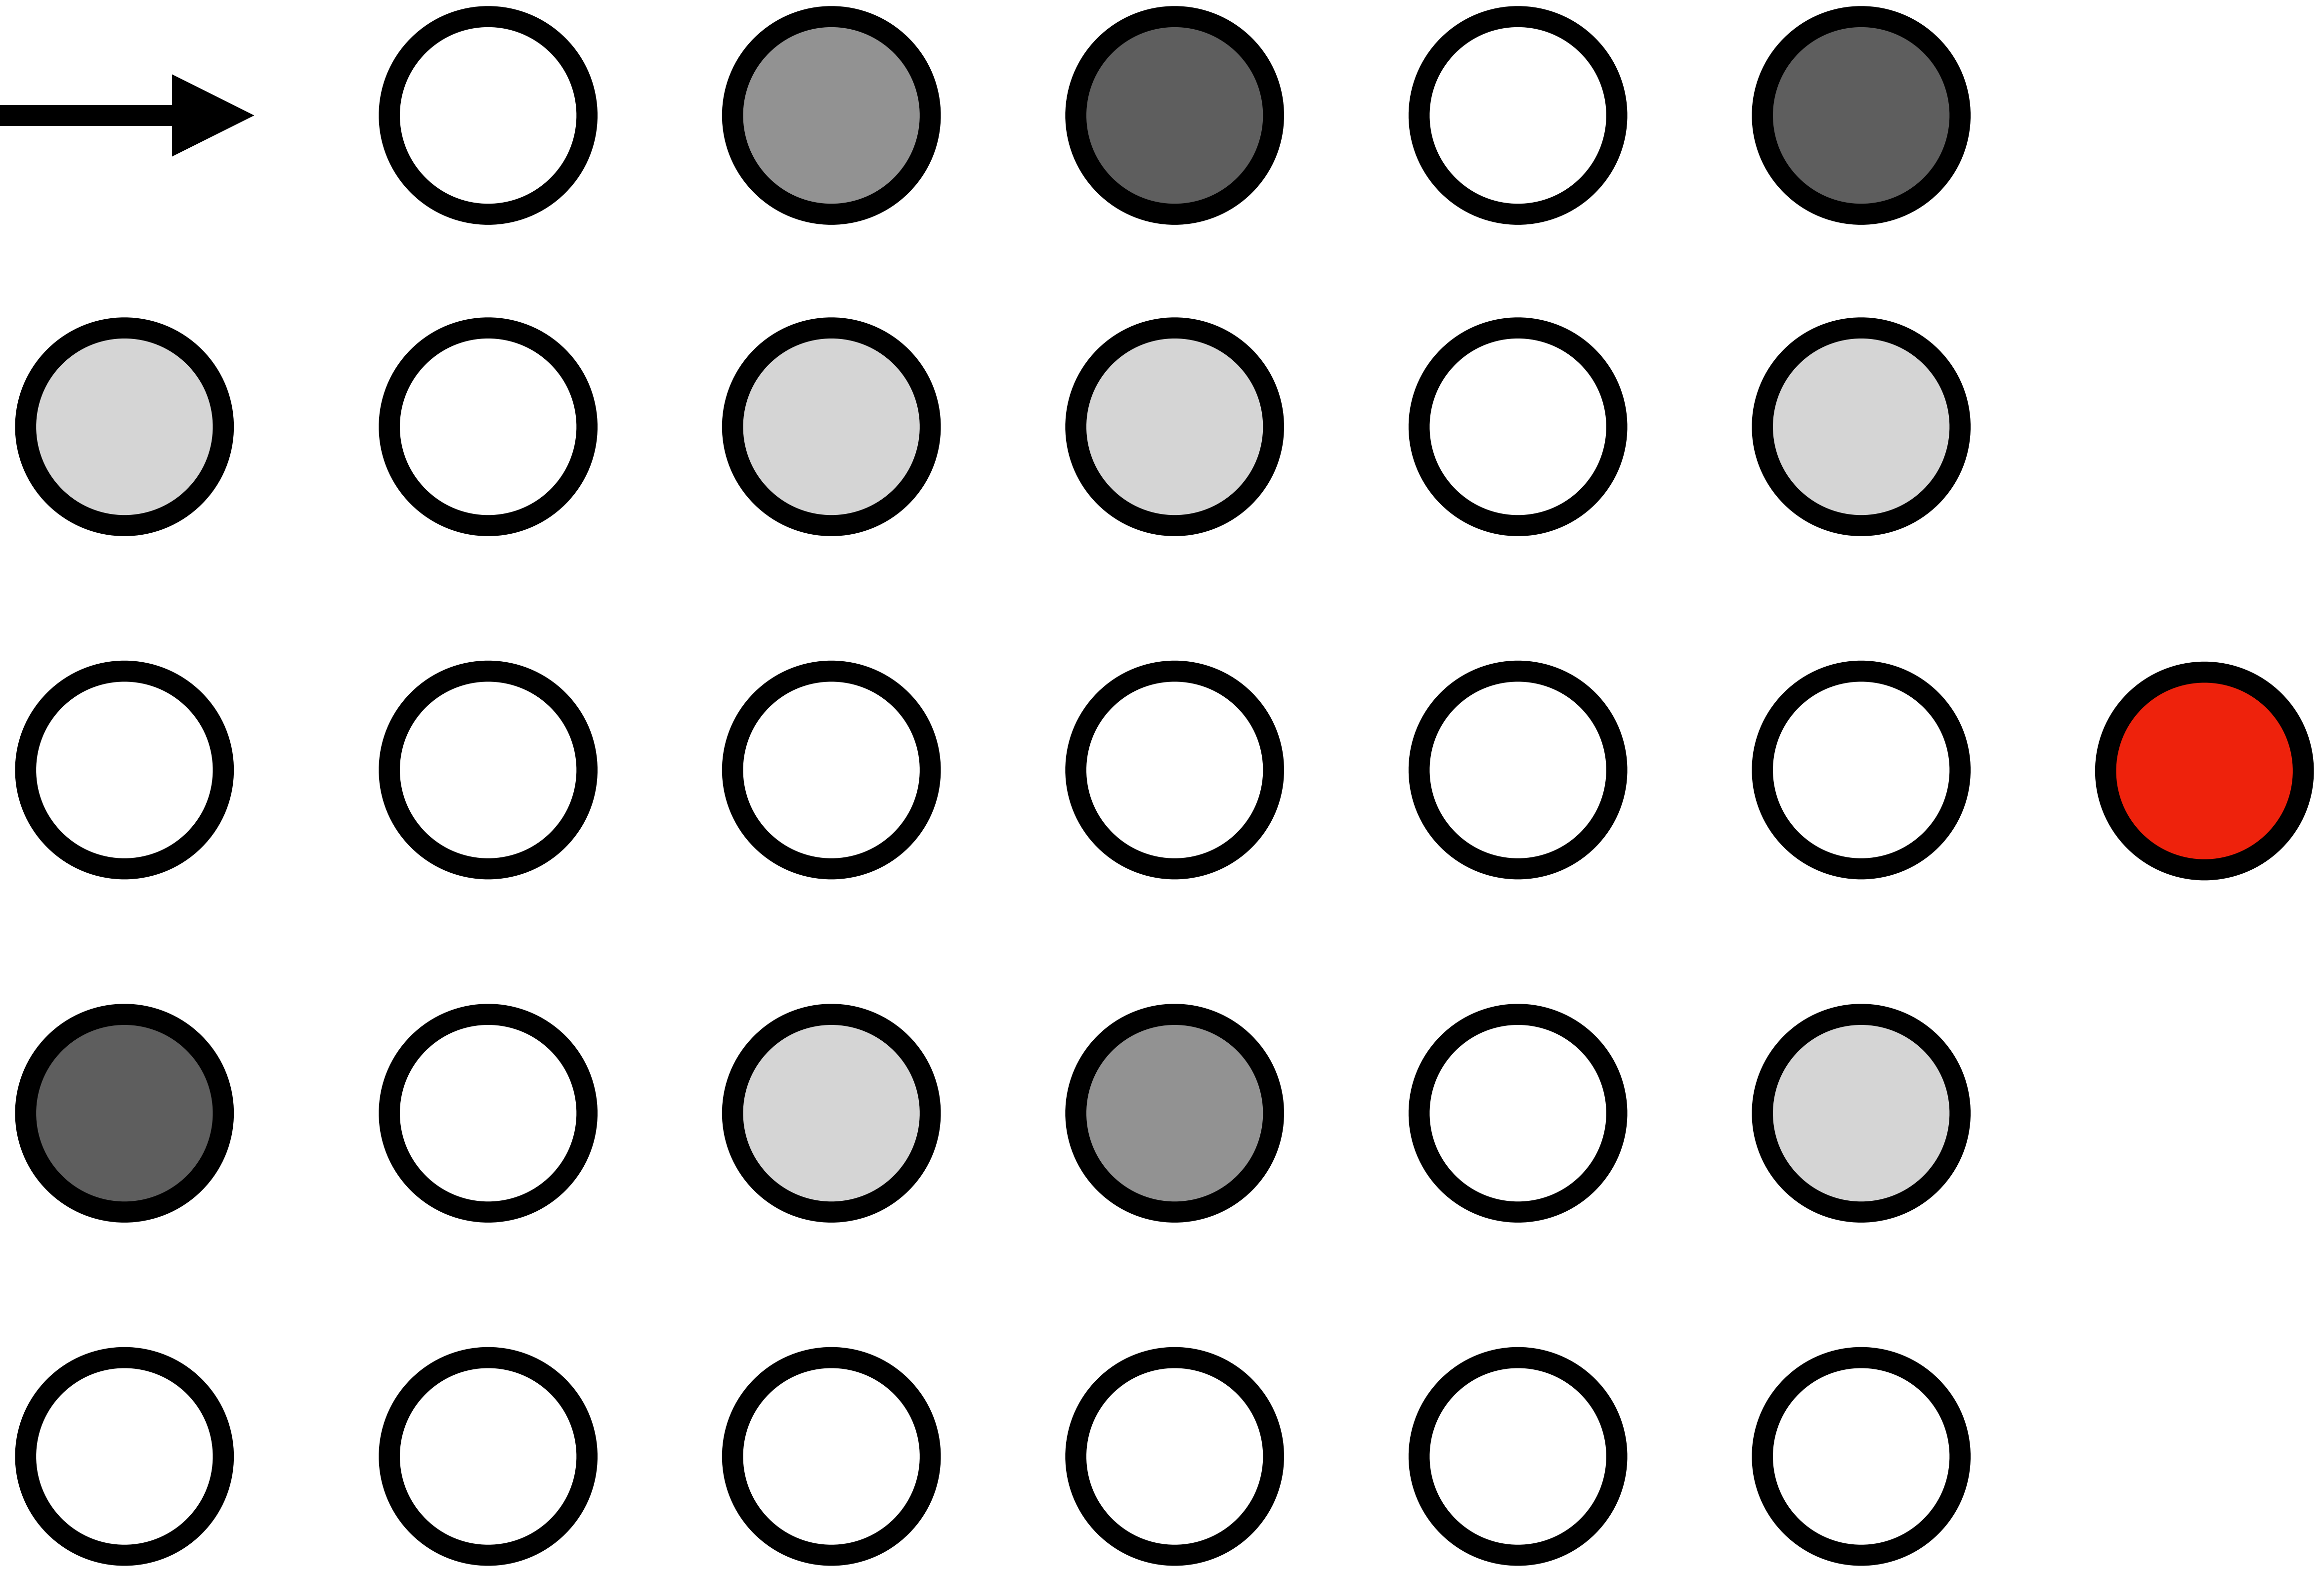

A connectionist implementation

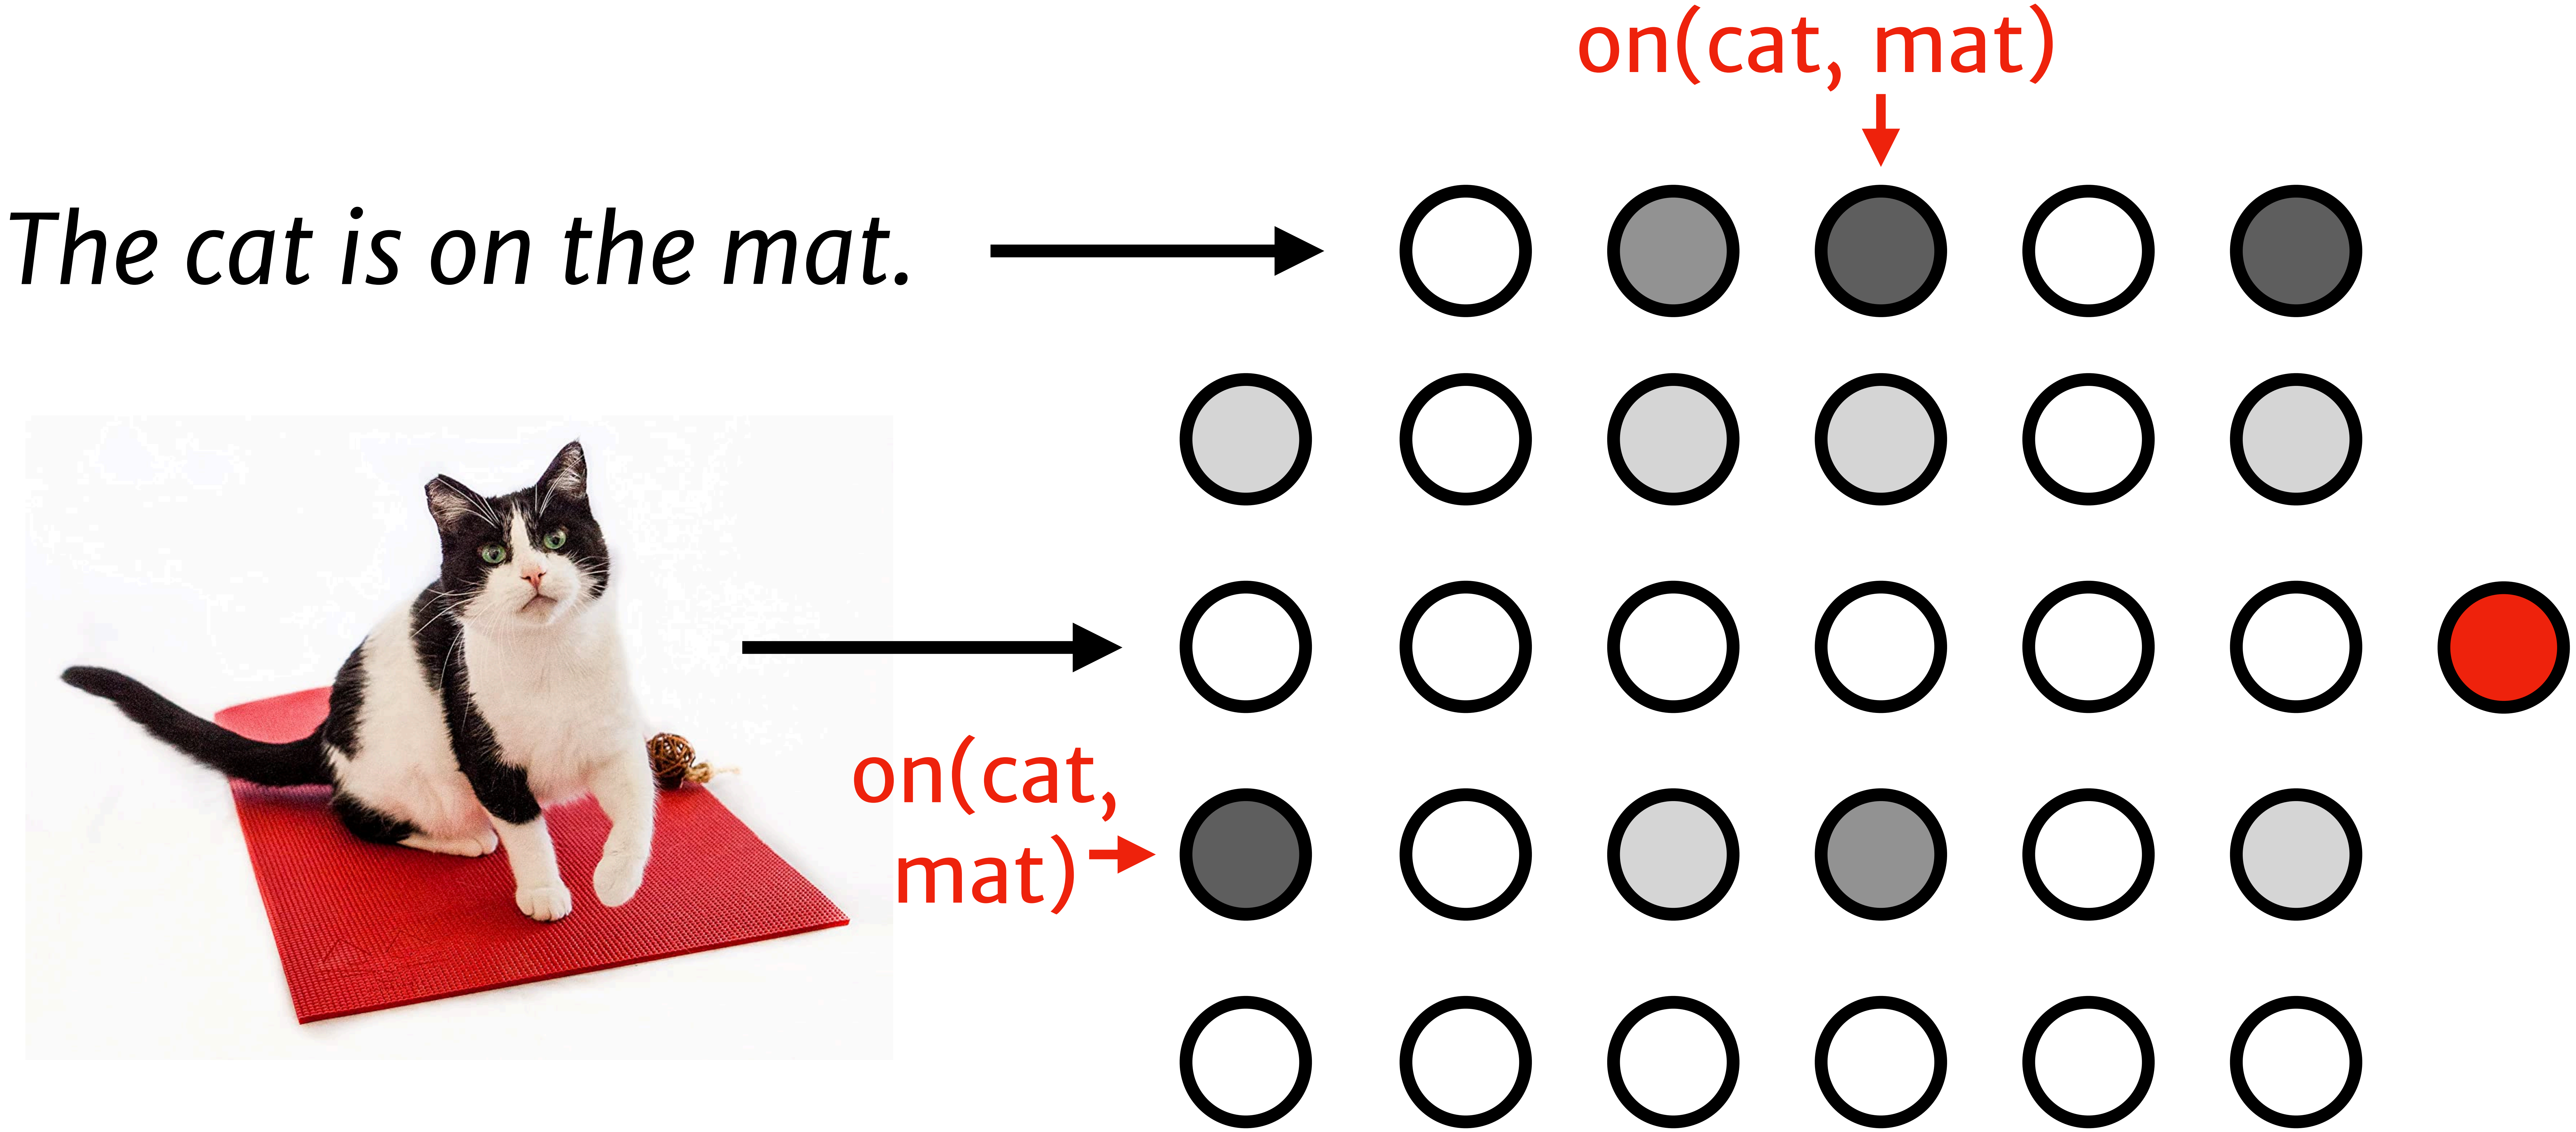

A modern neural implementation

The cat is on the mat.

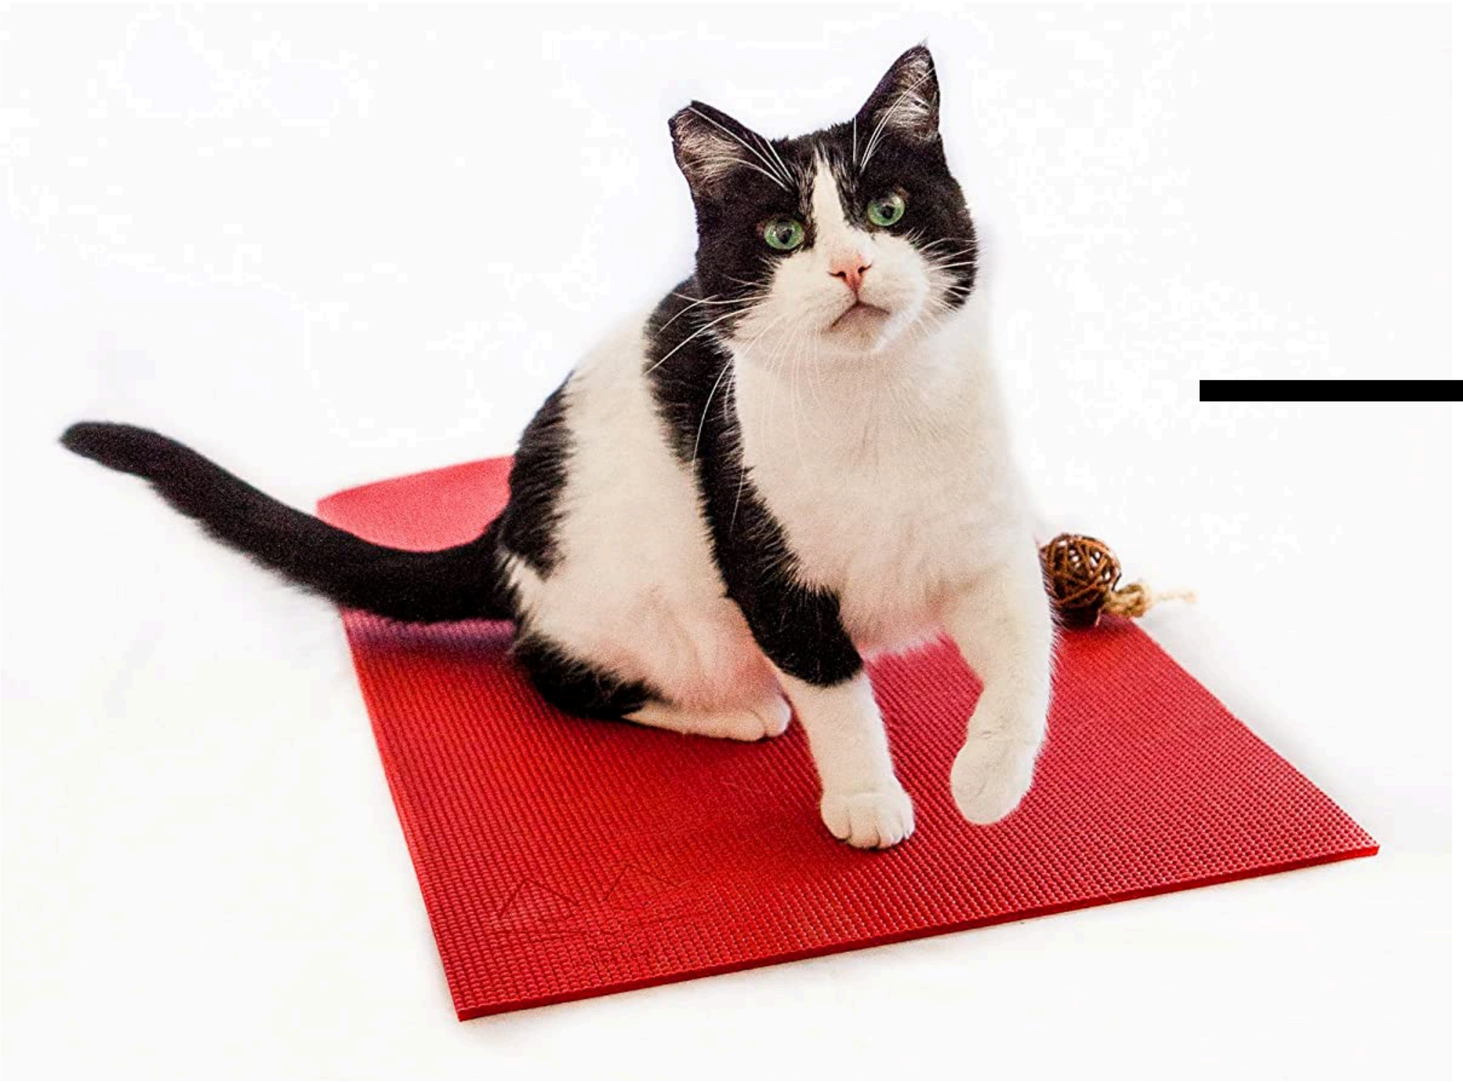

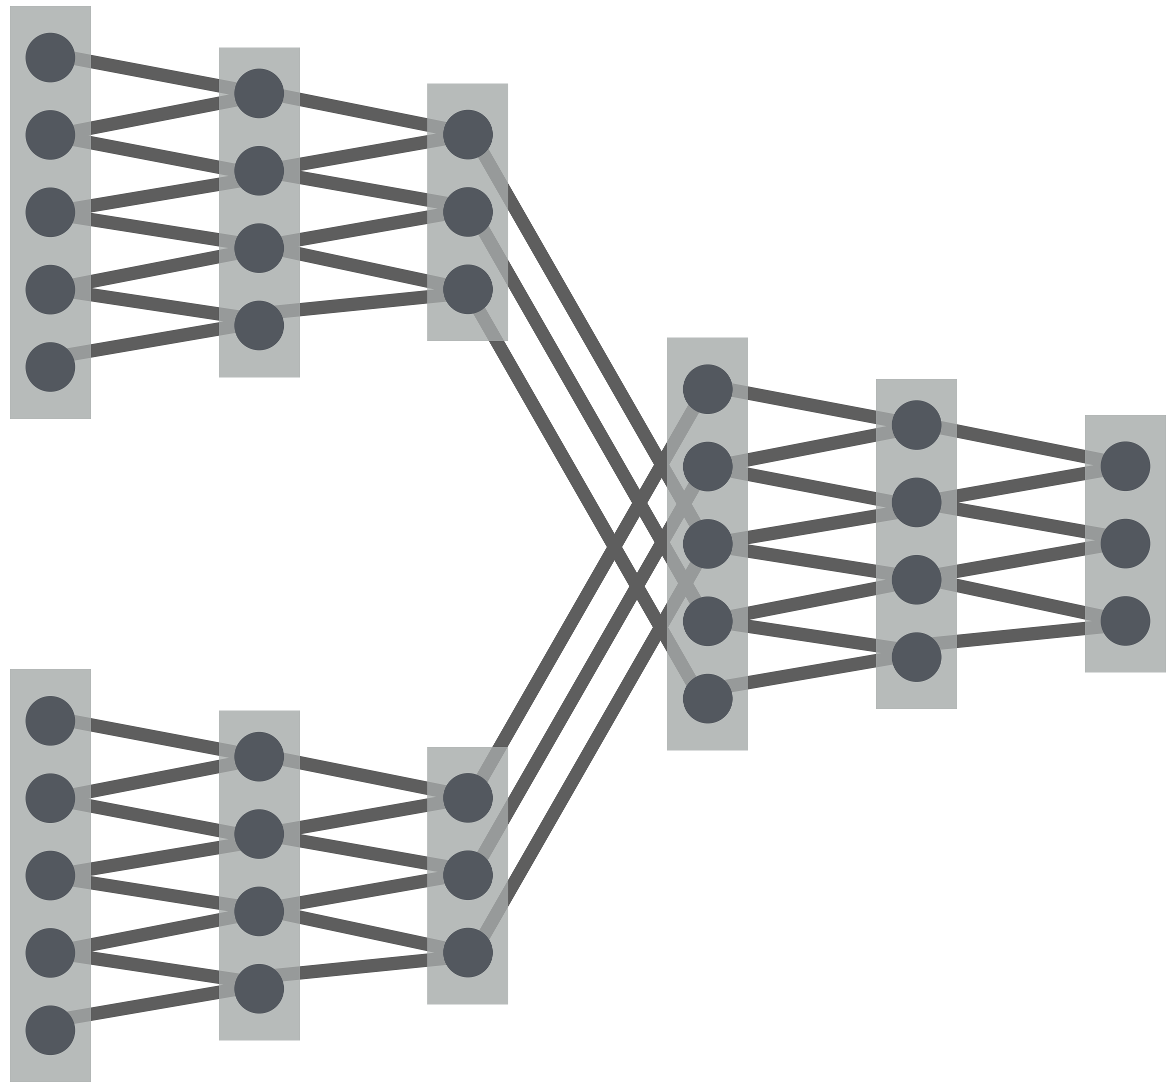

True

[https://www.amazon.in/Feline-Yogi-Original-Yoga-Cat]

A classical implementation

The cat is on the mat and the fox is in a box. →

[[[*The cat*] [*is* [*on the mat*]]] [*and* [*the fox* [*is* [*in a box*]]]]]

↓

`cat(x), mat(y), box(z), on(x, y), …`

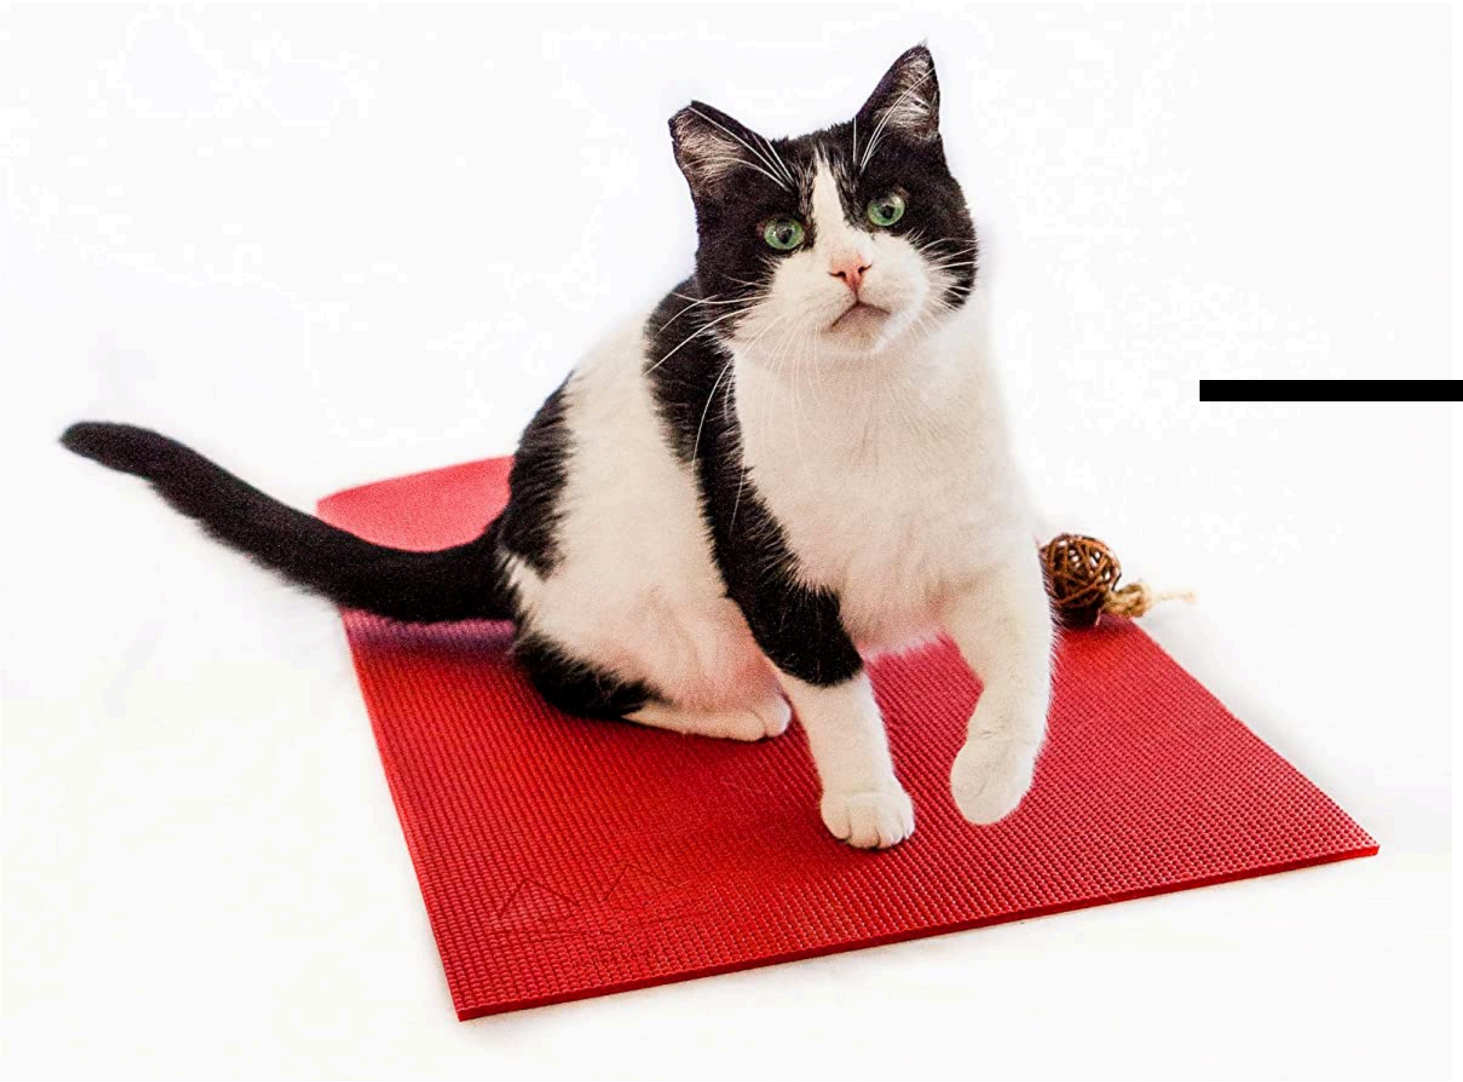

→

```
cat(x)
mat(y)
red(y)
on(x, y)
…
```

→ $\vdash$ → **True**

[https://www.amazon.in/Feline-Yogi-Original-Yoga-Cat]

A connectionist implementation

The cat is on the mat and the fox is in a box.

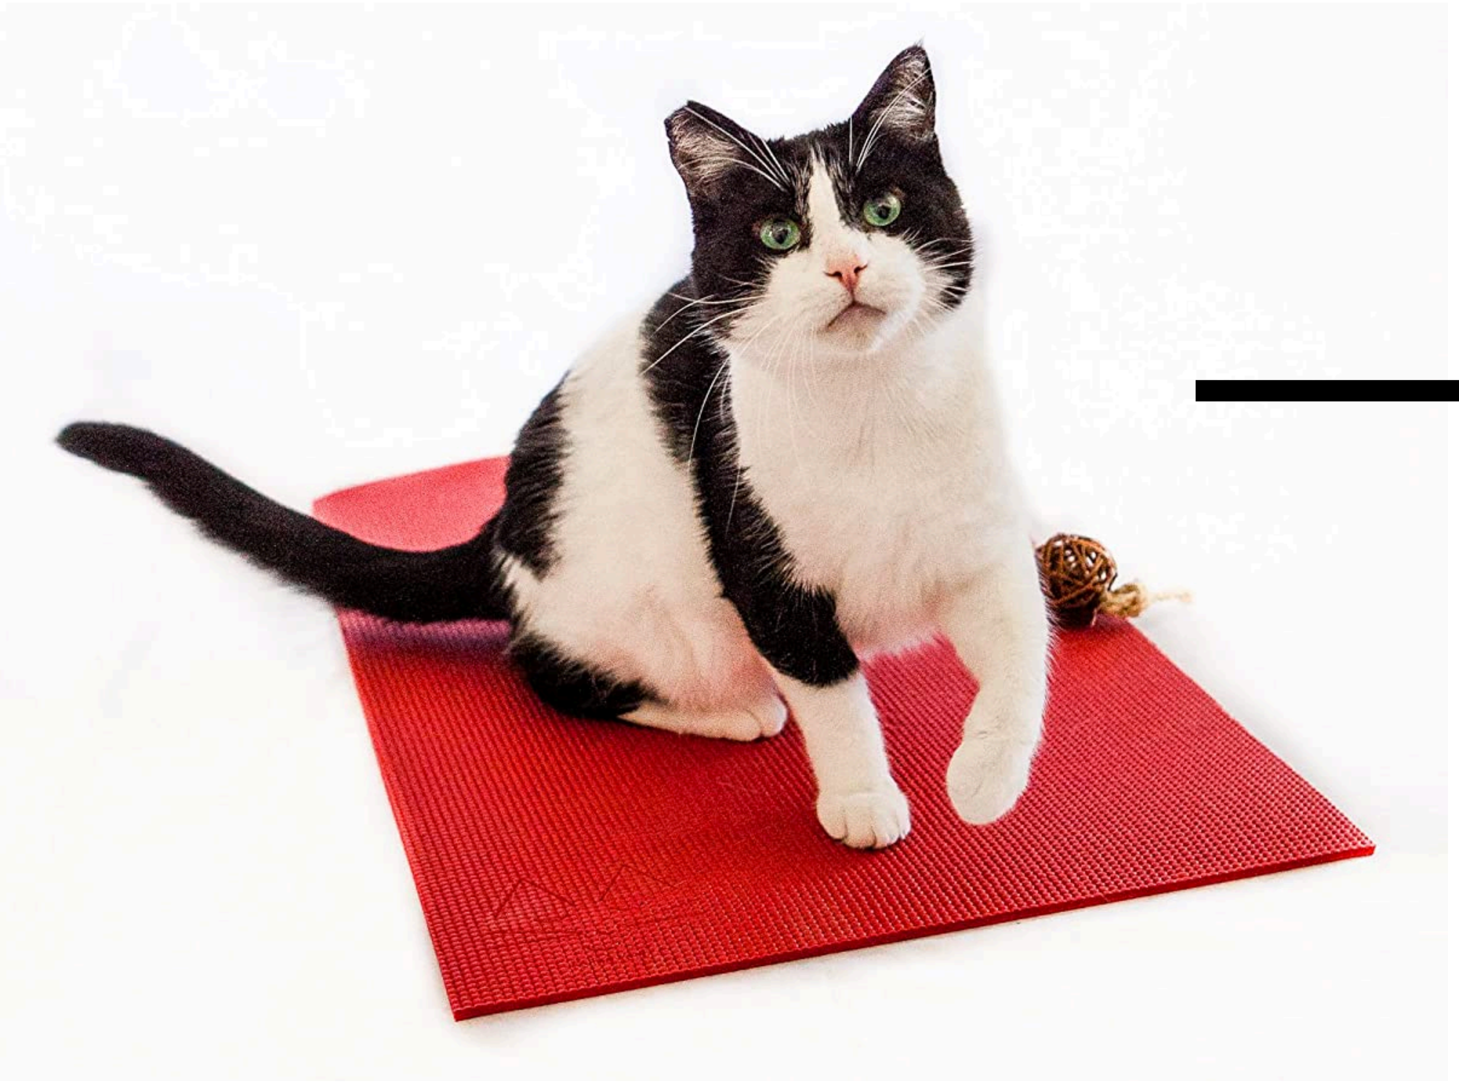

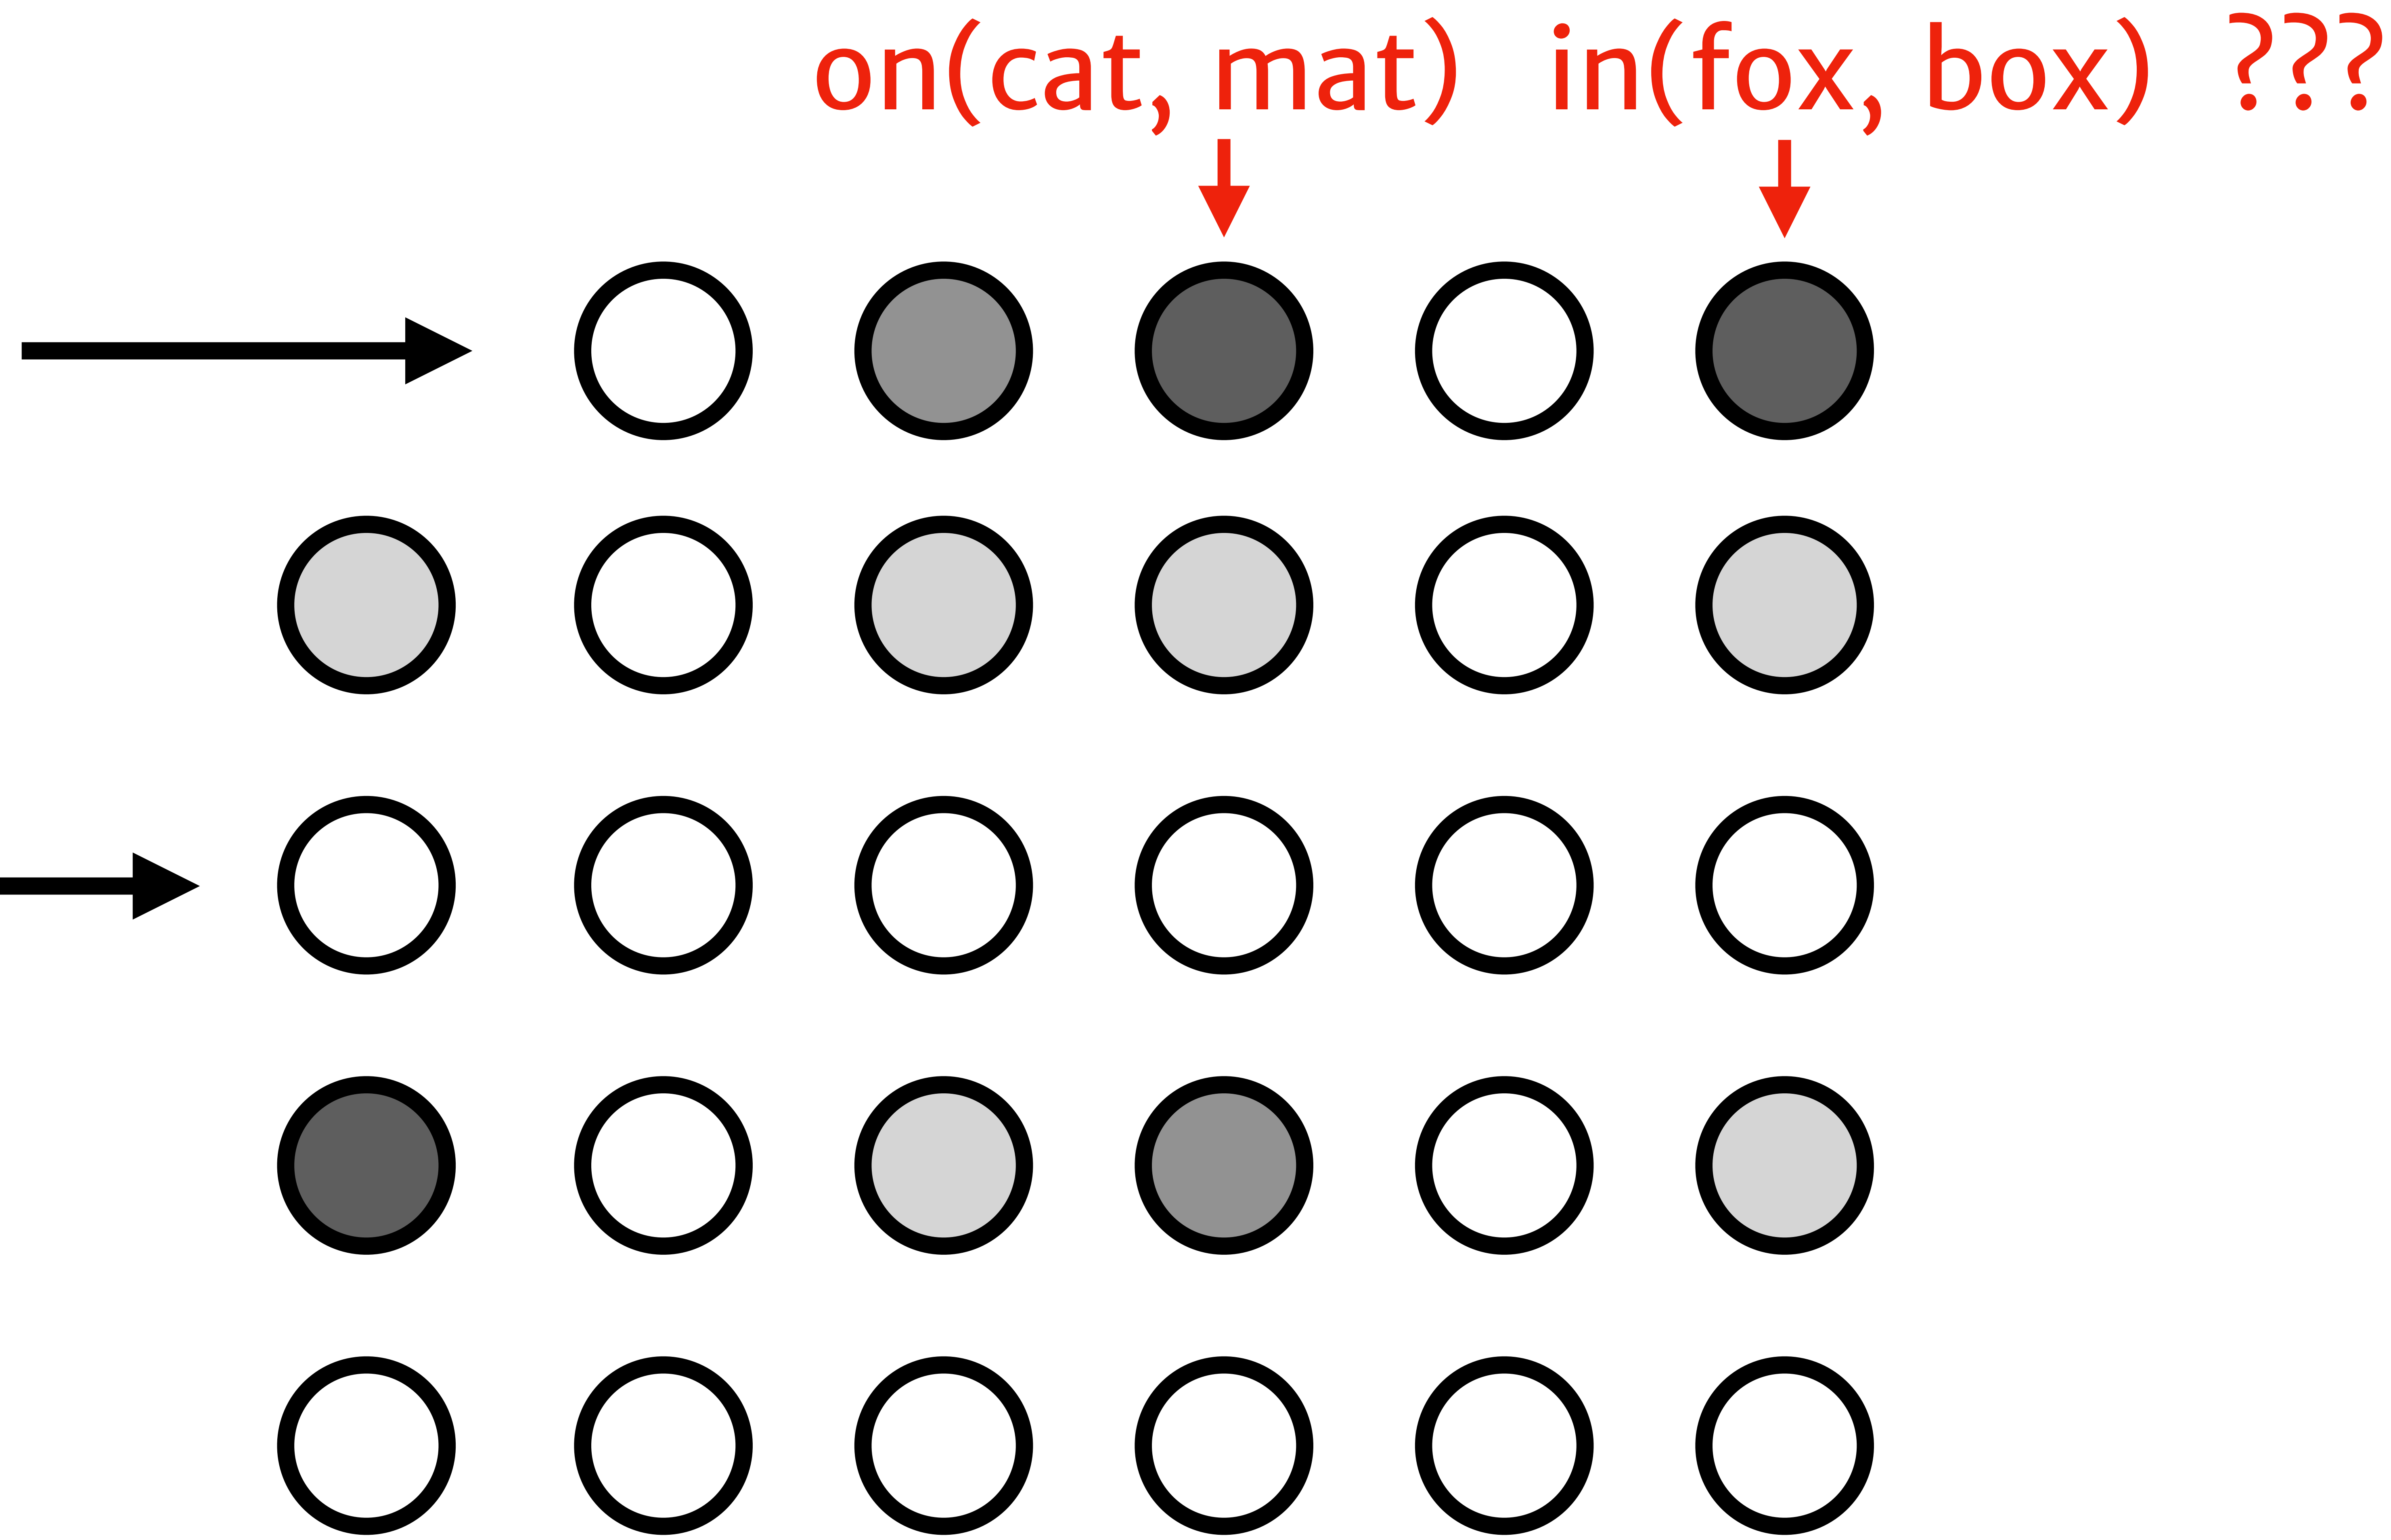

A connectionist implementation

on1(., cat)   on2(., mat)   ???

The cat is on the mat and the fox is in a box.

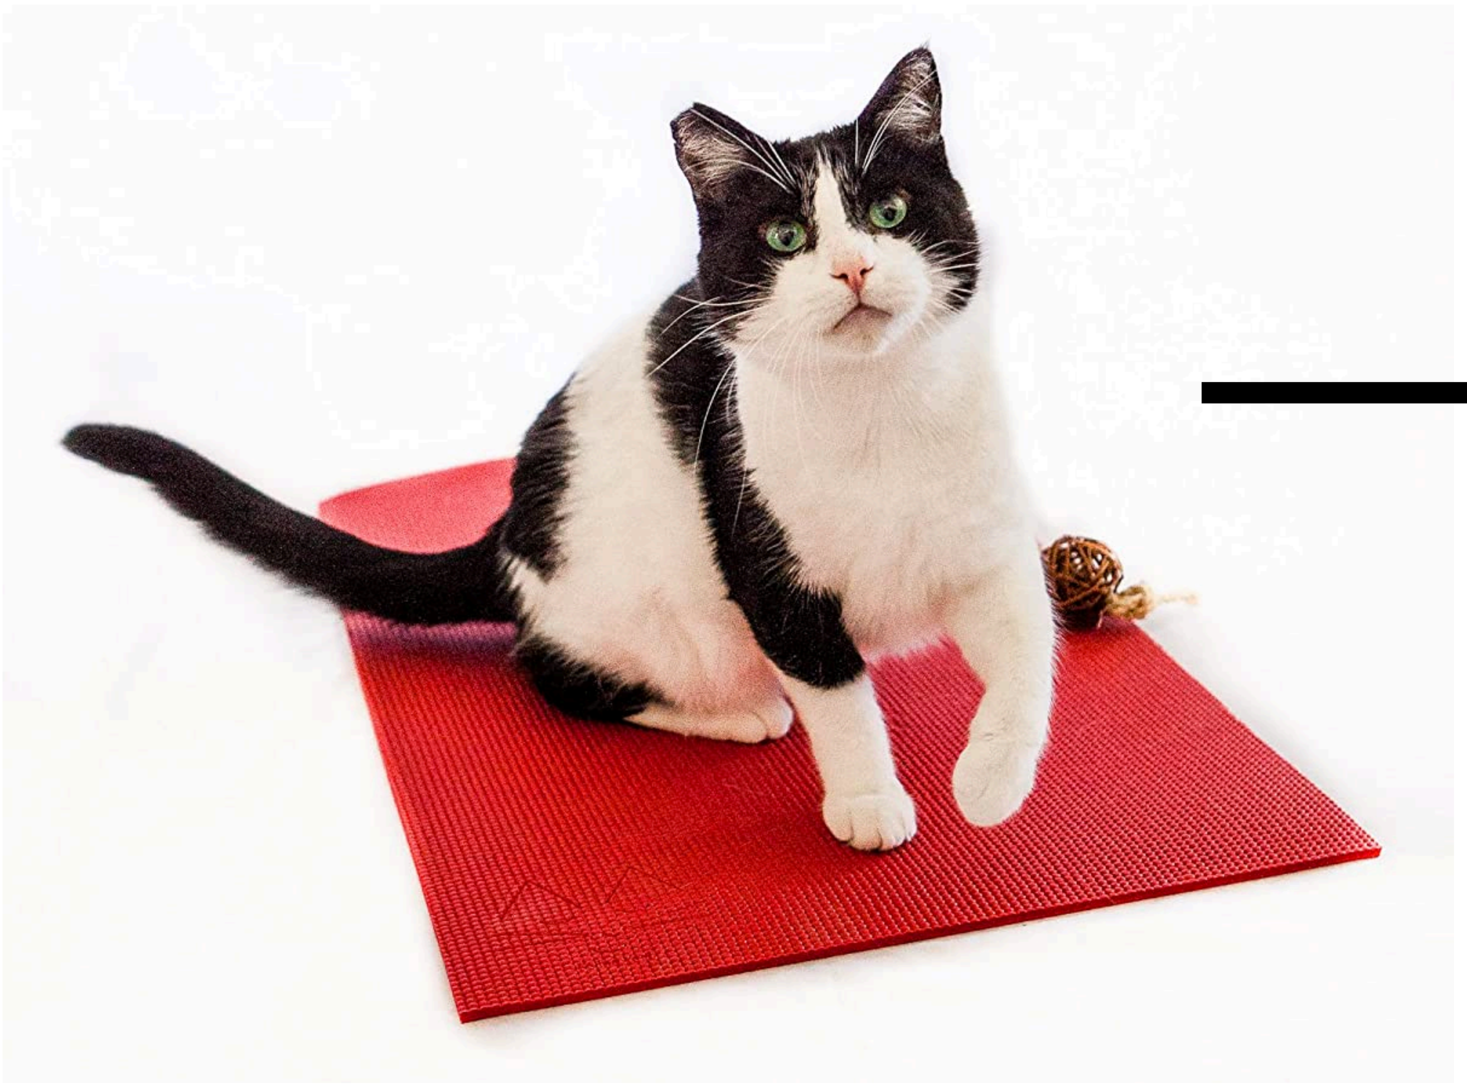

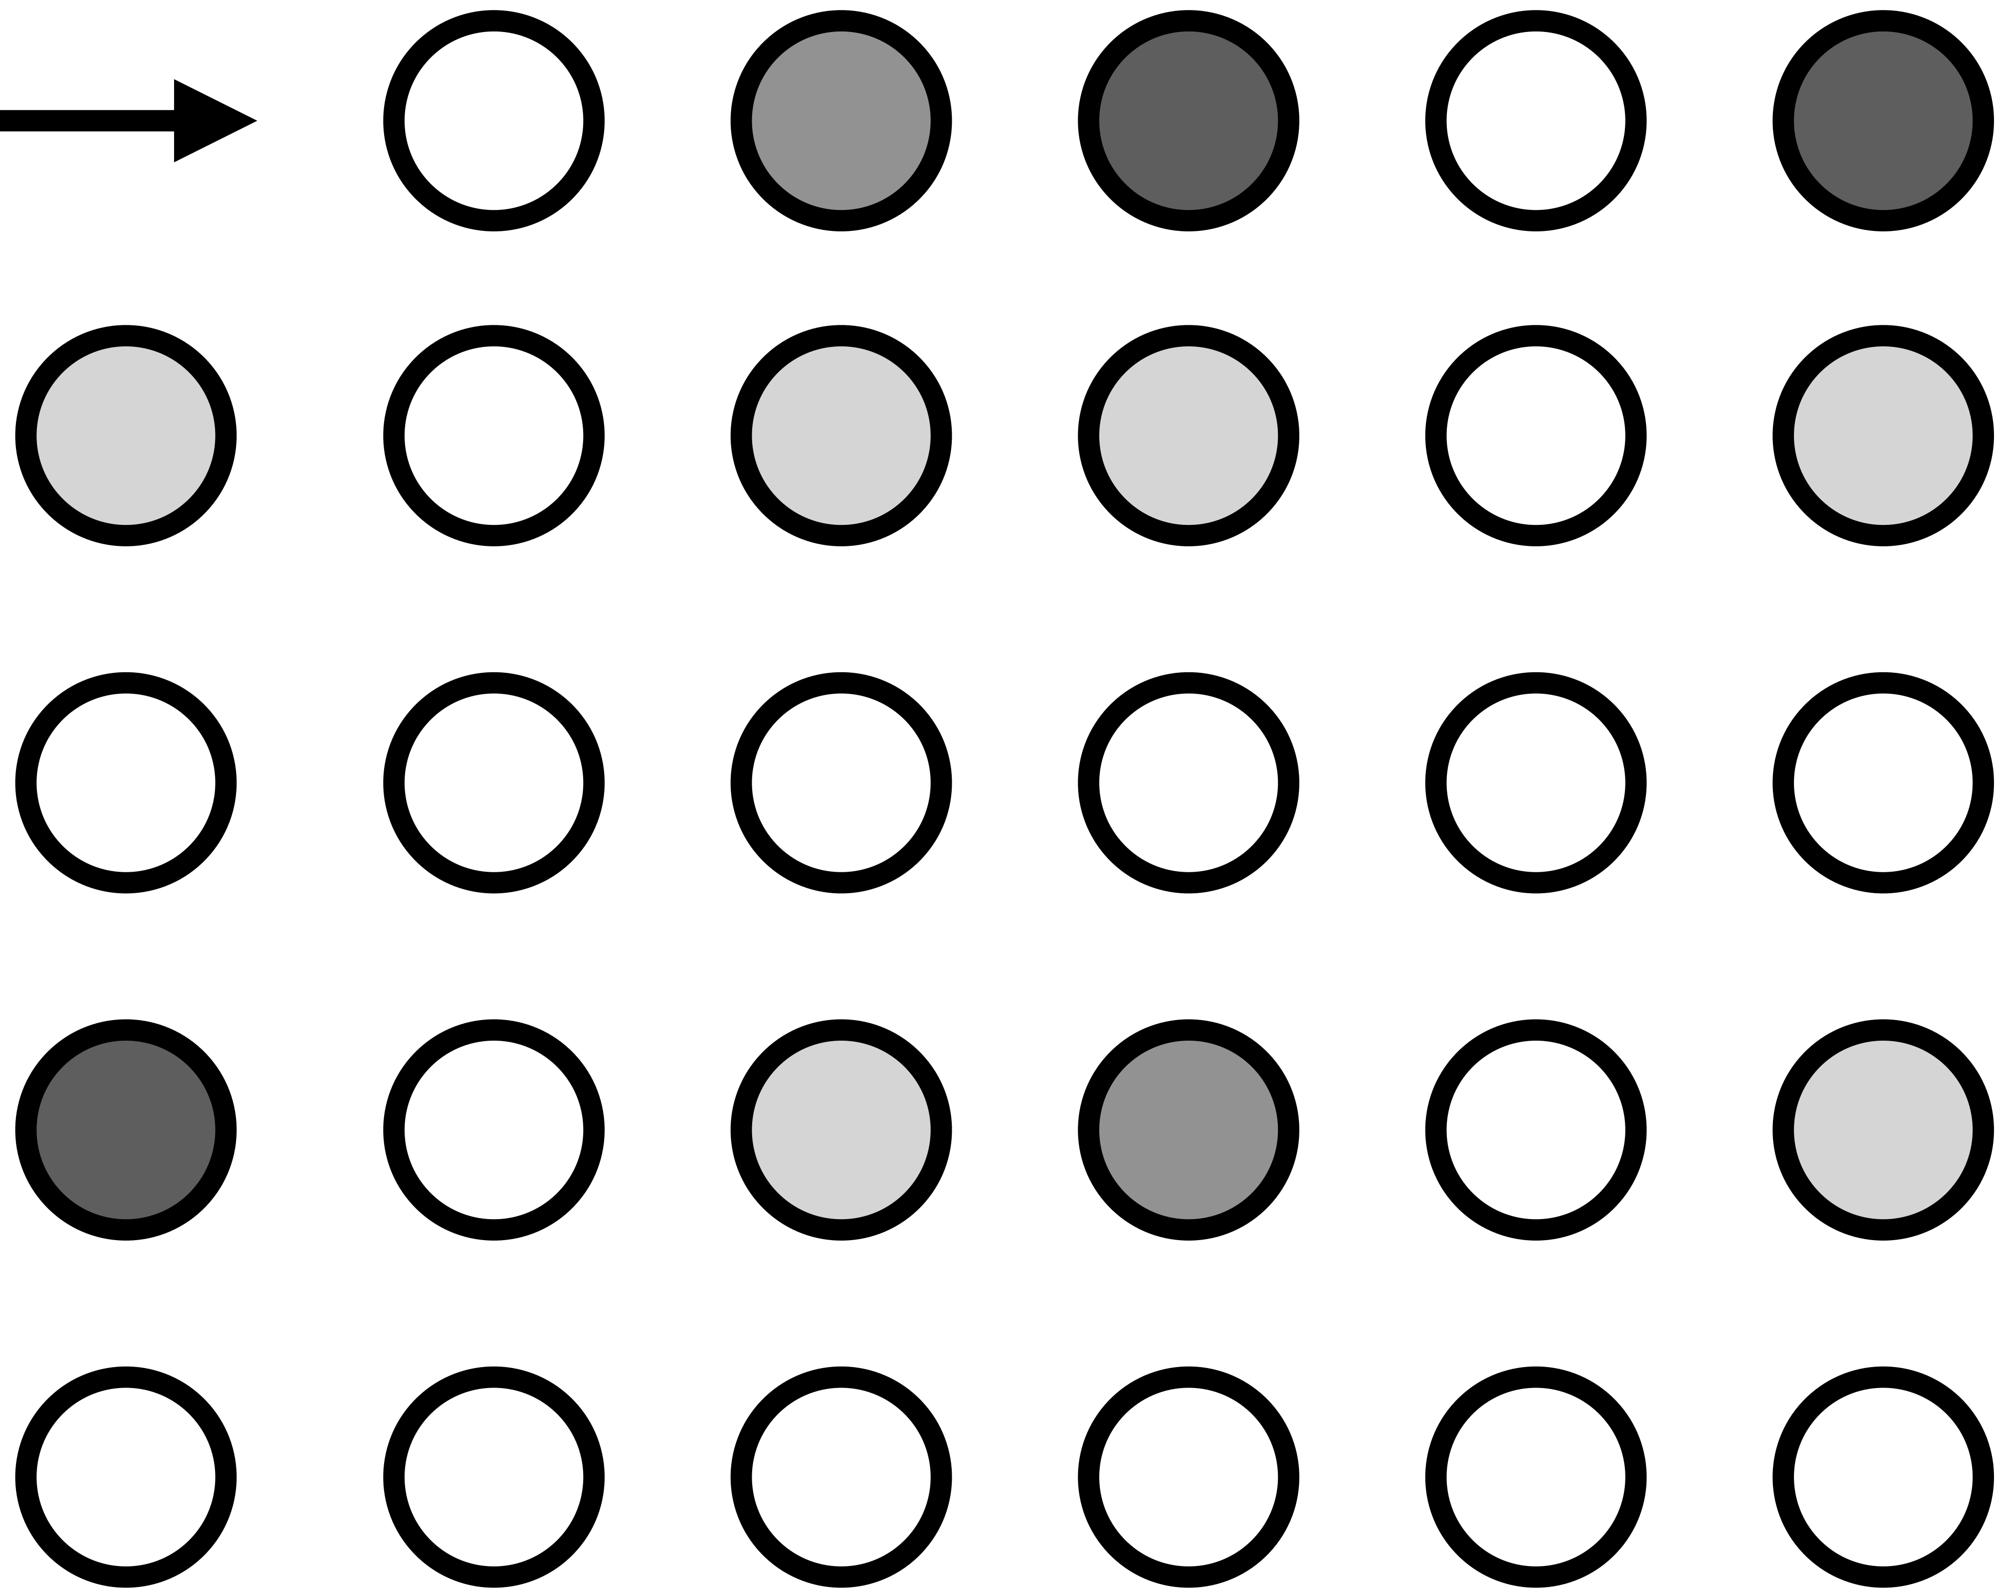

A modern neural implementation

The cat is on the mat and the fox is in a box.

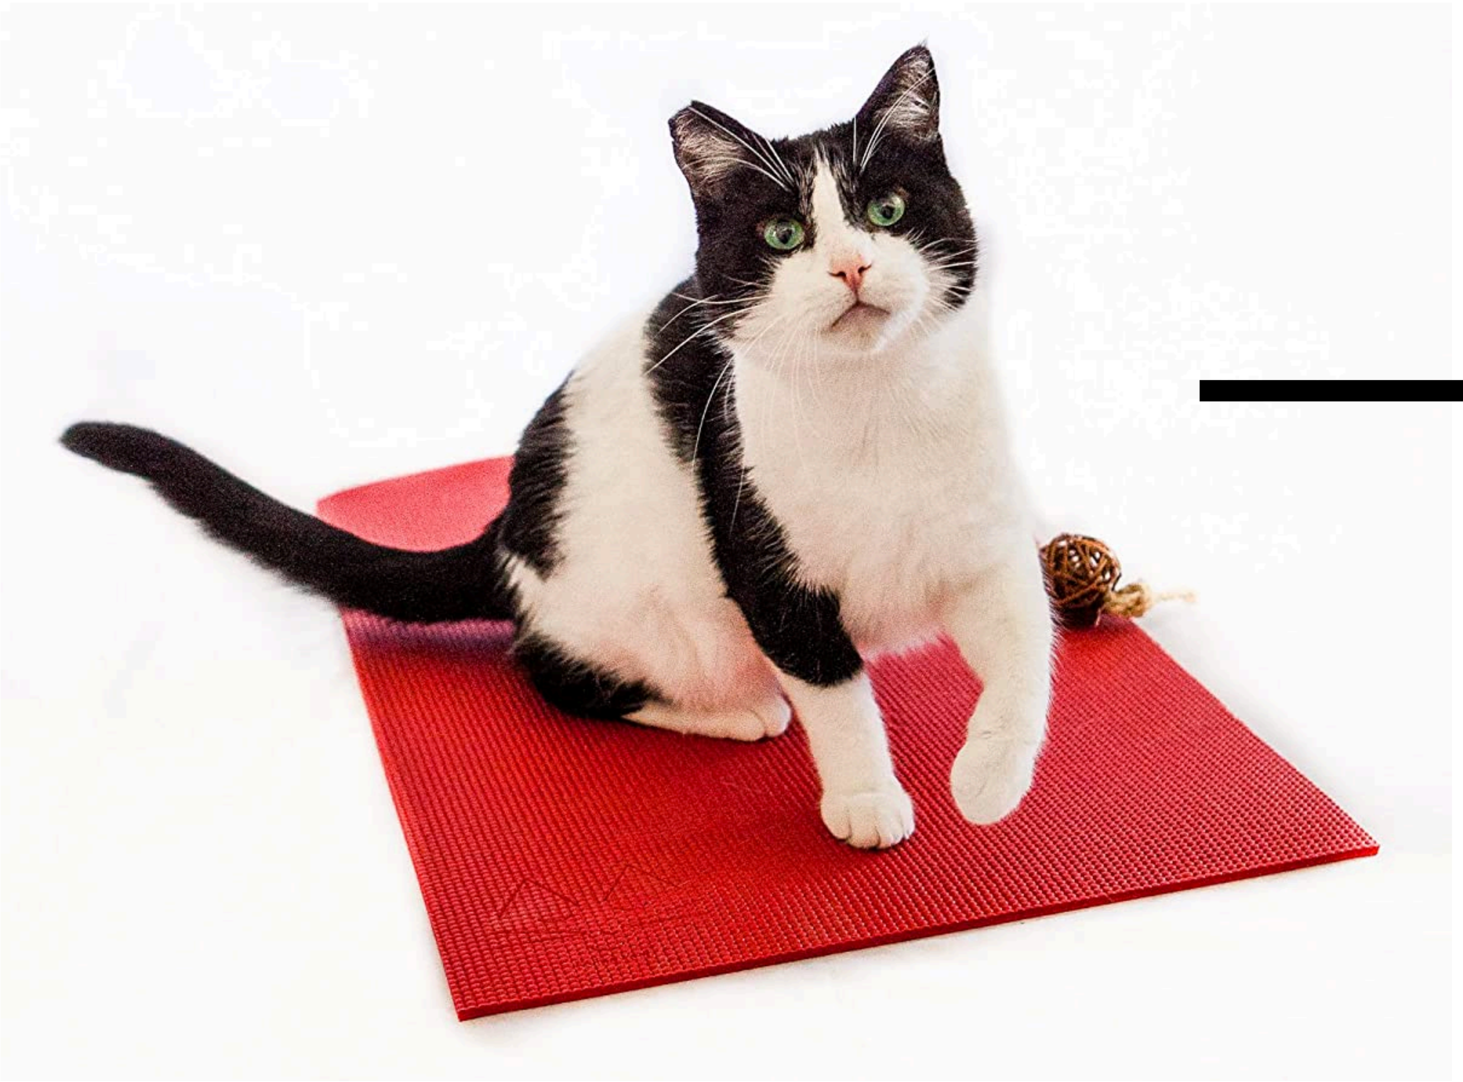

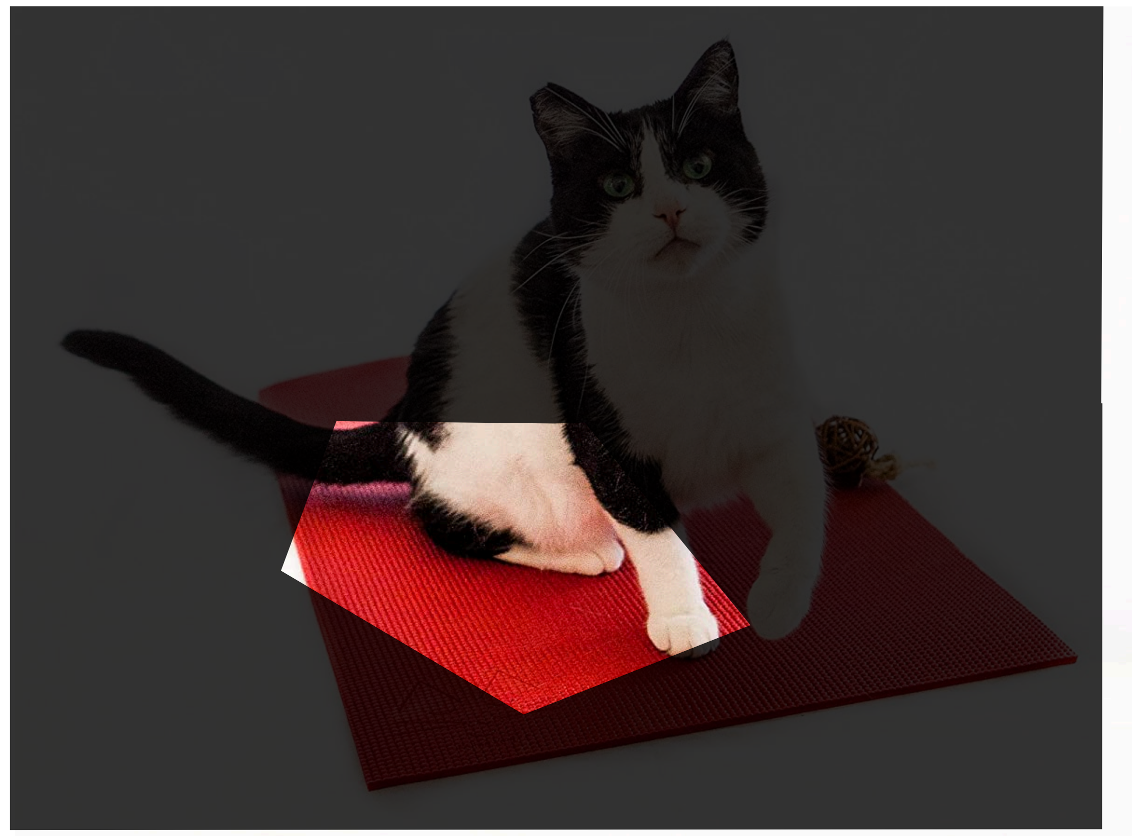

True

[https://www.amazon.in/Feline-Yogi-Original-Yoga-Cat]

Classical representations contain their constituents

[[[*The cat*] [*is* [*on the mat*]]]
[*and* [*the fox* [*is* [*in a box*]]]]]

[[*The cat*] [*is* [*on the mat*]]]

Classical representations contain their constituents

[[[*The cat*] [*is* [*on the mat*]]]
[*and* [*the fox* [*is* [*in a box*]]]]]

[[*The cat*] [*is* [*on the mat*]]]

Constituents of connectionist representations?

The cat is on the mat and the fox is in the box.

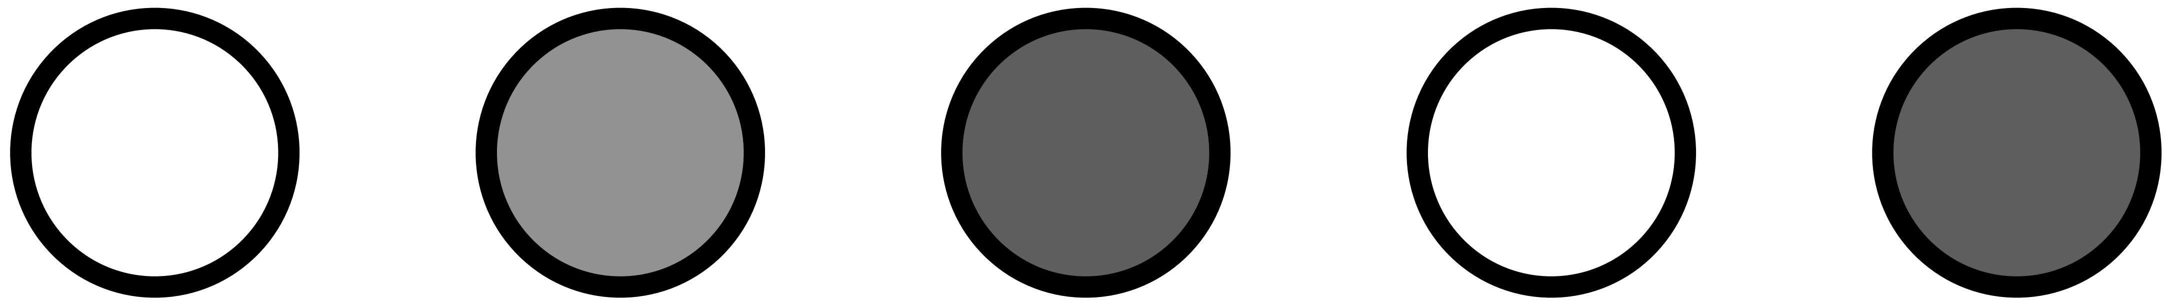

The cat is on the mat.

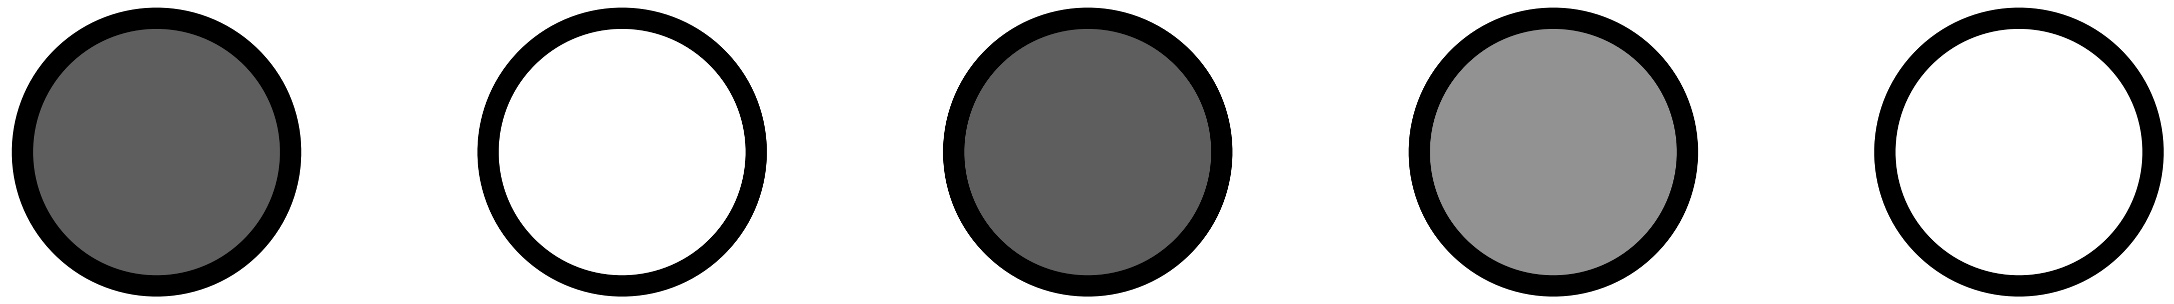

Algebraic structure

[*the fox* [*is* [*in a box*]]]

$*$

[[*The cat*] [*is* [*on the mat*]]]

=

[[[*The cat*] [*is* [*on the mat*]]]

[*and* [*the fox* [*is* [*in a box*]]]]]

Combinatorial structure

in(fox, box)

*

on(cat, mat)

=

and(in(fox, box), on(cat, mat))

Combinatorial structure

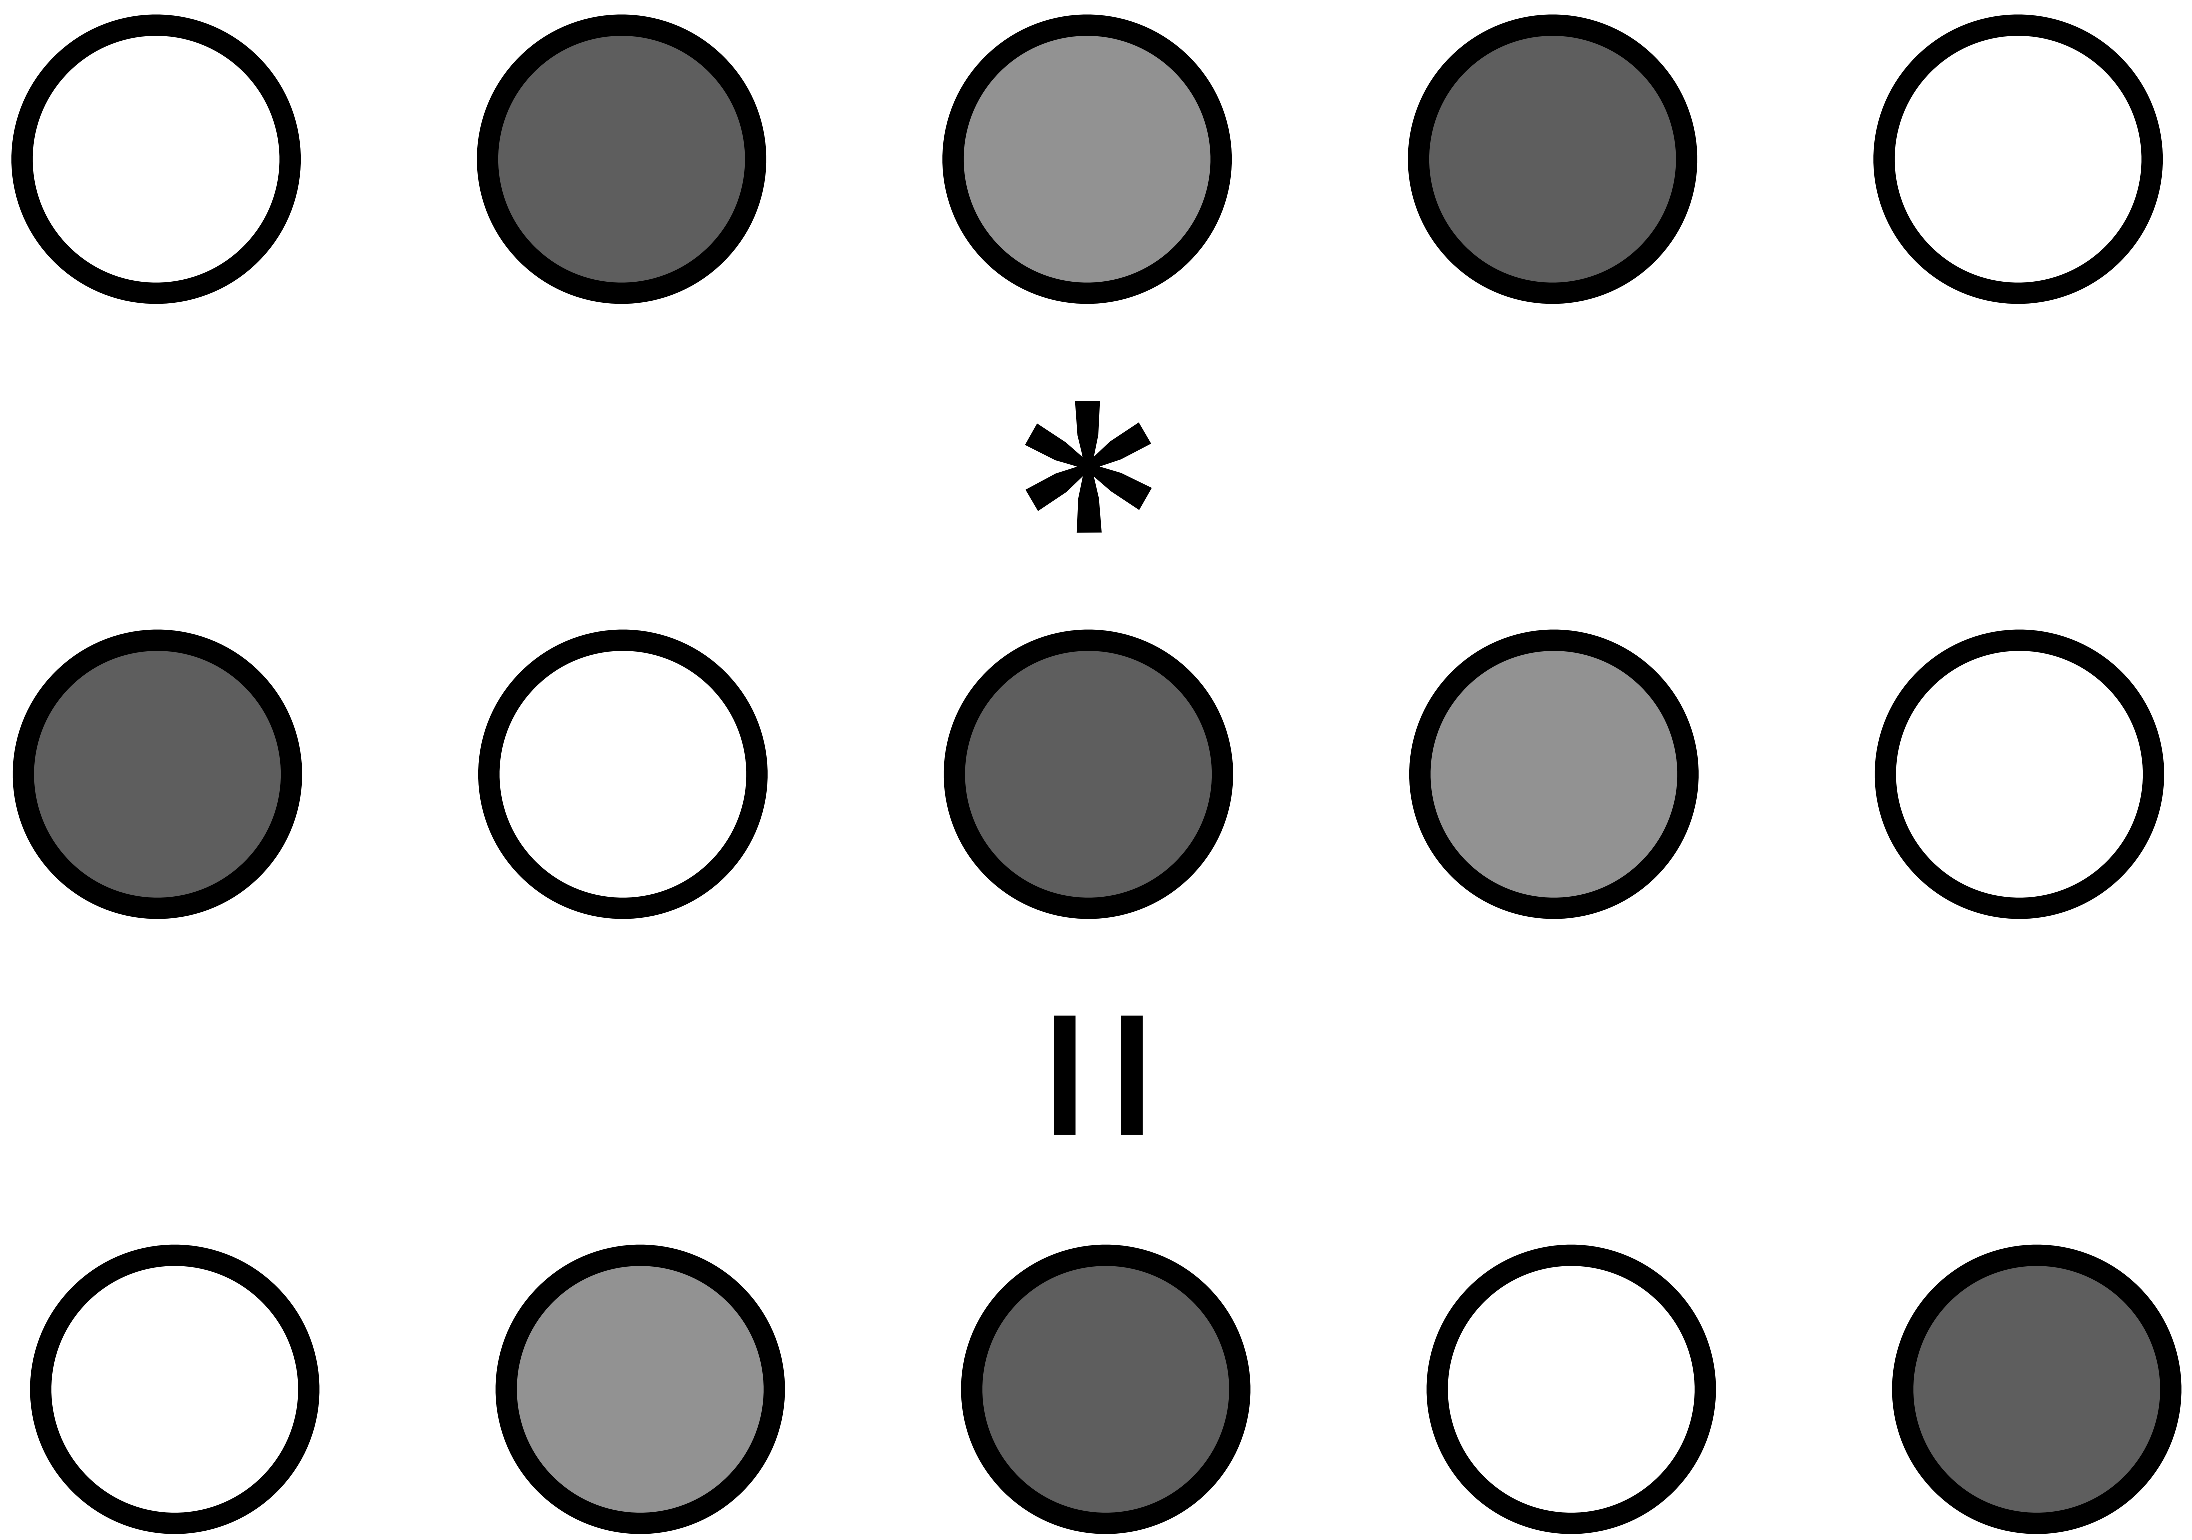

Combinatorial structure

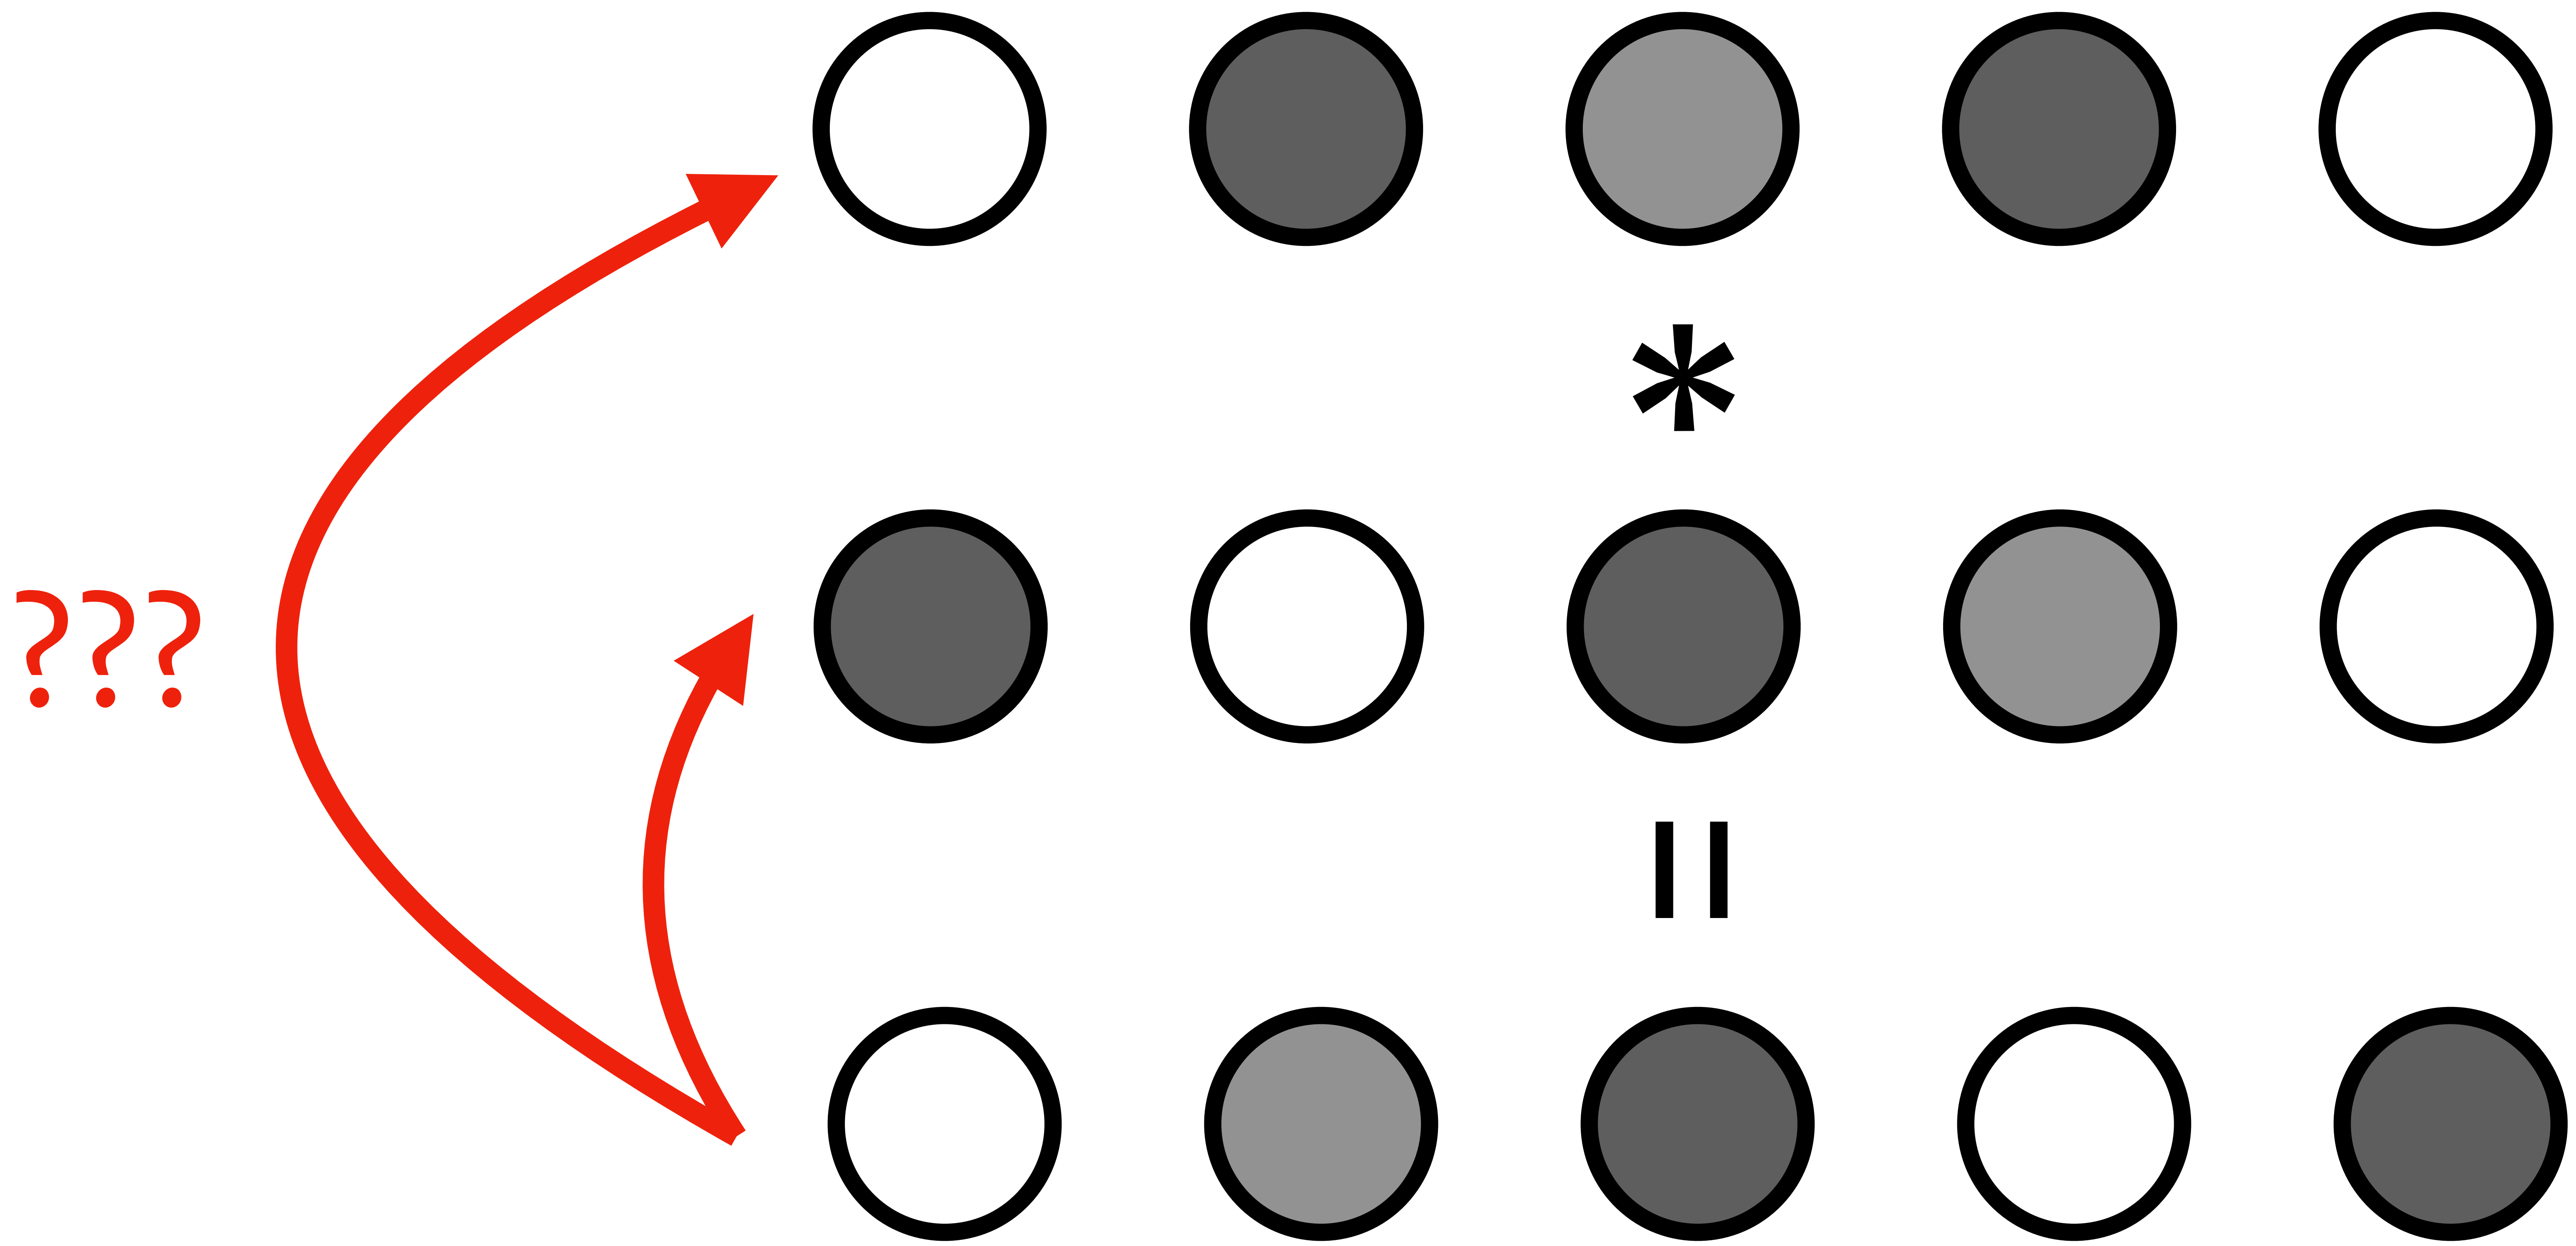

Algebraic structure

[*the fox* [*is* [*in a box*]]]

*

[[*The cat*] [*is* [*on the mat*]]]

=

[[[*The cat*] [*is* [*on the mat*]]]
[*and* [*the fox* [*is* [*in a box*]]]]]

Algebraic structure

$$\alpha$$

$$*$$

$$\beta$$

$$=$$

$$(\alpha * \beta)$$

Discussion

Structure-sensitive processing

The cat is on the mat. → [[*The cat*] [*is* [*on the mat*]]]

↓

`cat(x), mat(y), on(x, y)`

↓

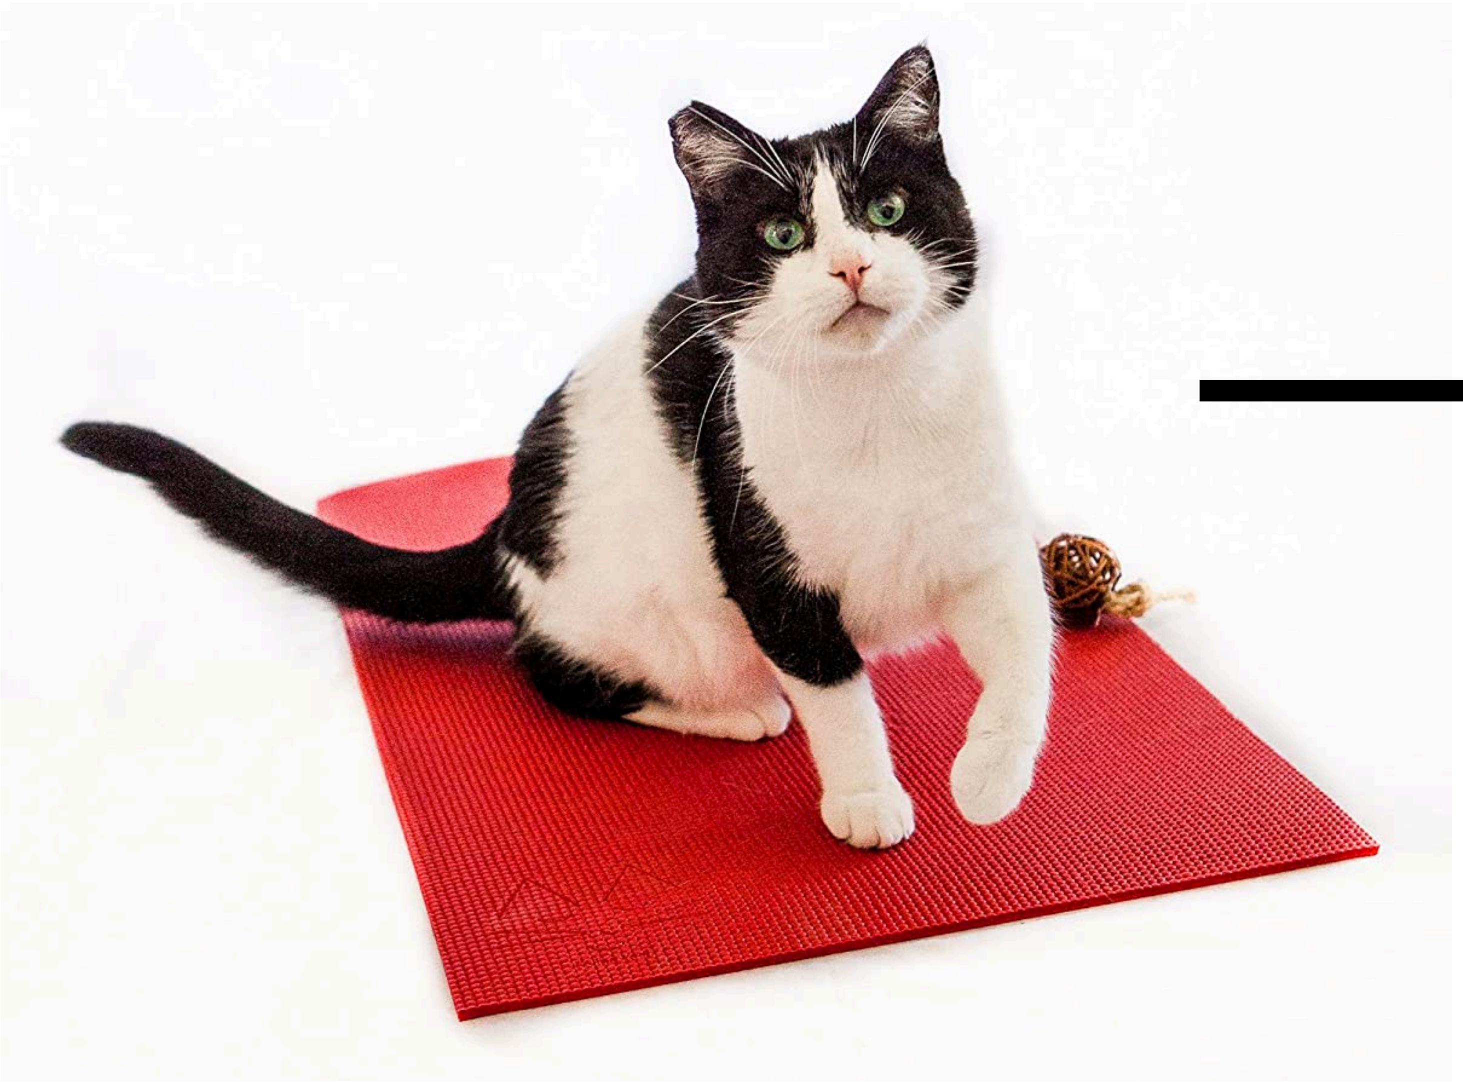

→

```
cat(x)
mat(y)
red(y)
on(x, y)
…
```

→ $\vdash$ → **True**

[https://www.amazon.in/Feline-Yogi-Original-Yoga-Cat]

Structure-sensitive processing

$$\alpha \wedge \beta \rightarrow \beta$$

and(in(fox, box), on(cat, mat))

$\downarrow$

on(cat, mat)

$\vdash \longrightarrow$ True

Structure–sensitive processing

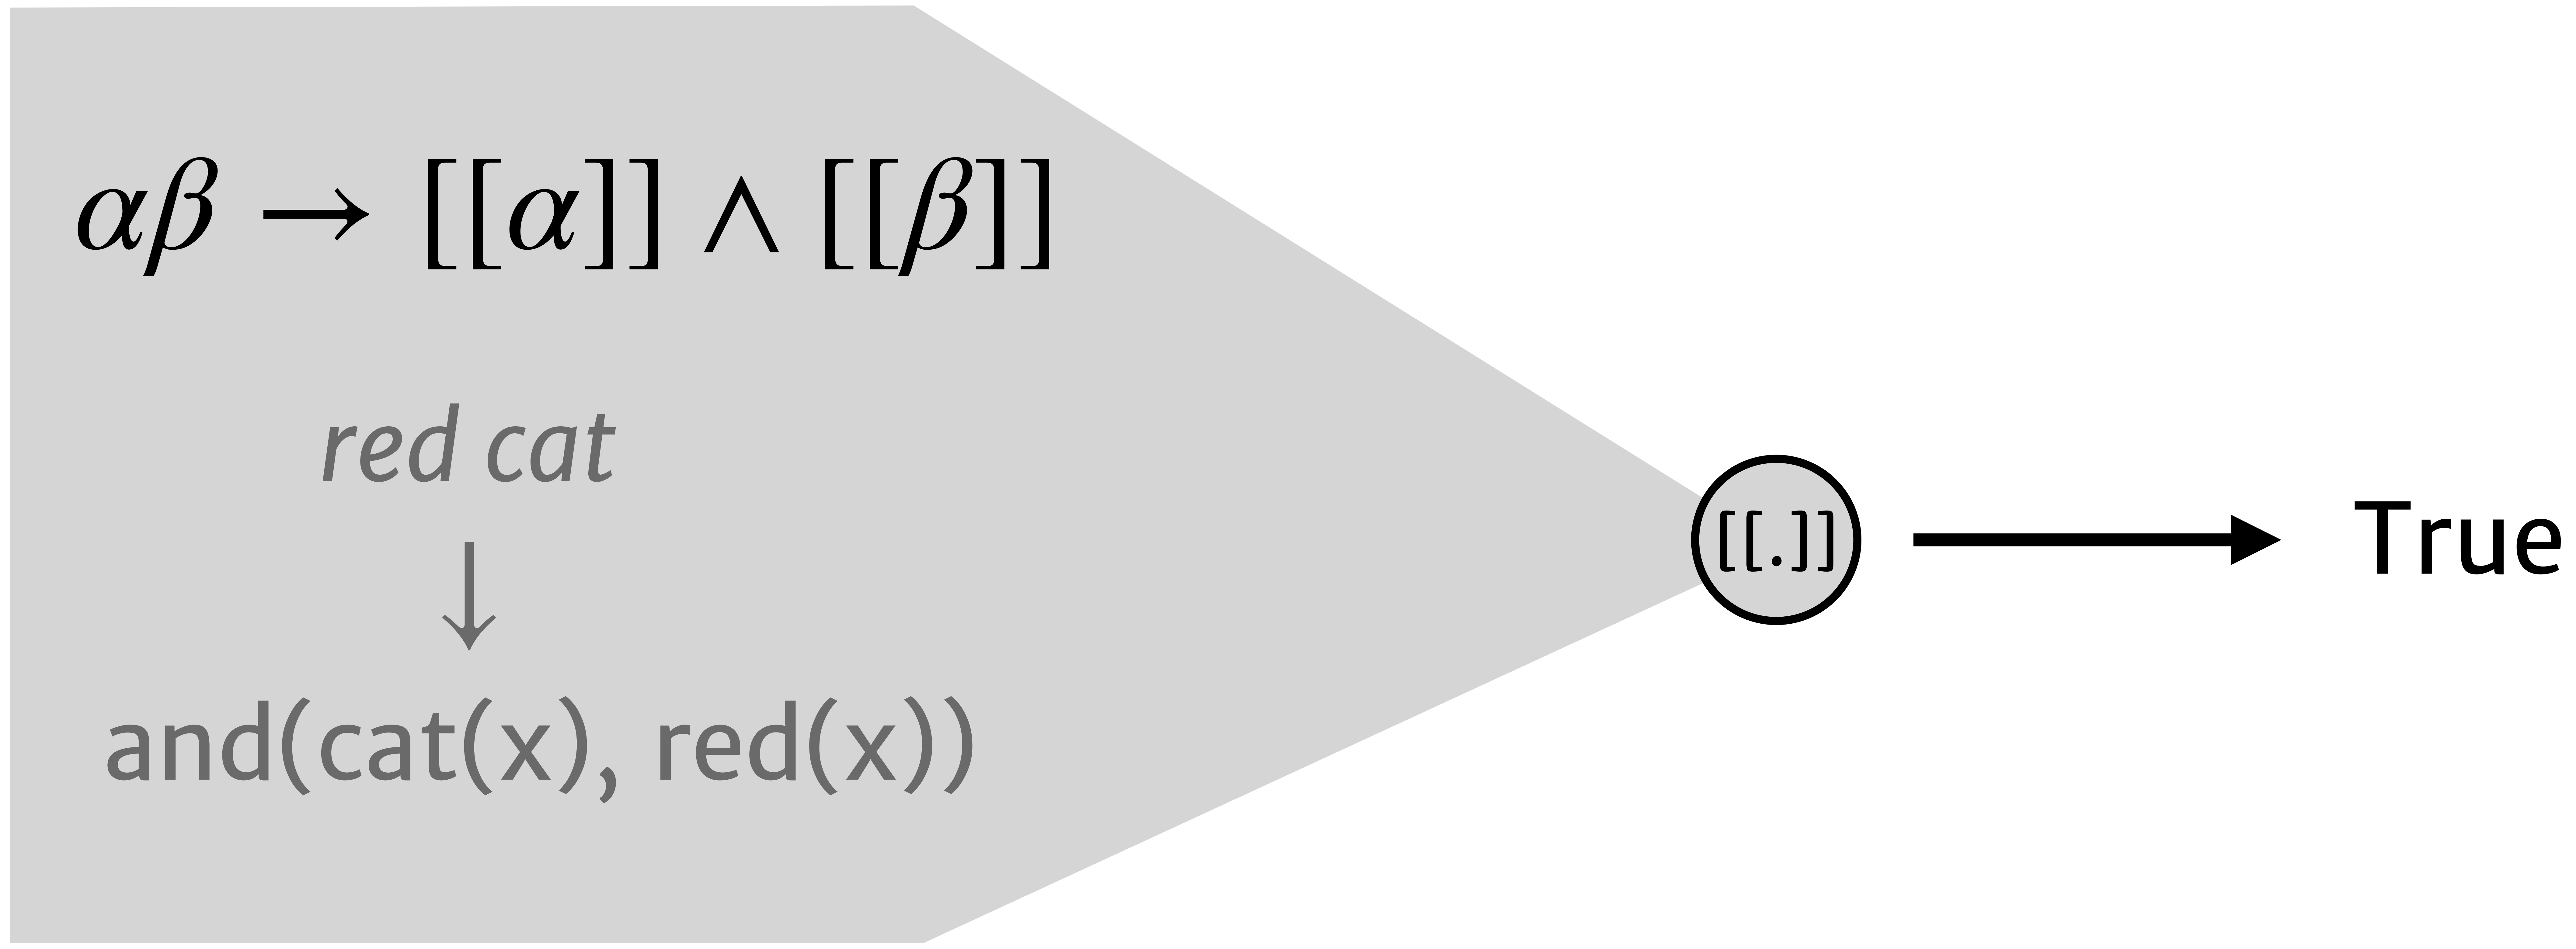

Structure-sensitive processing

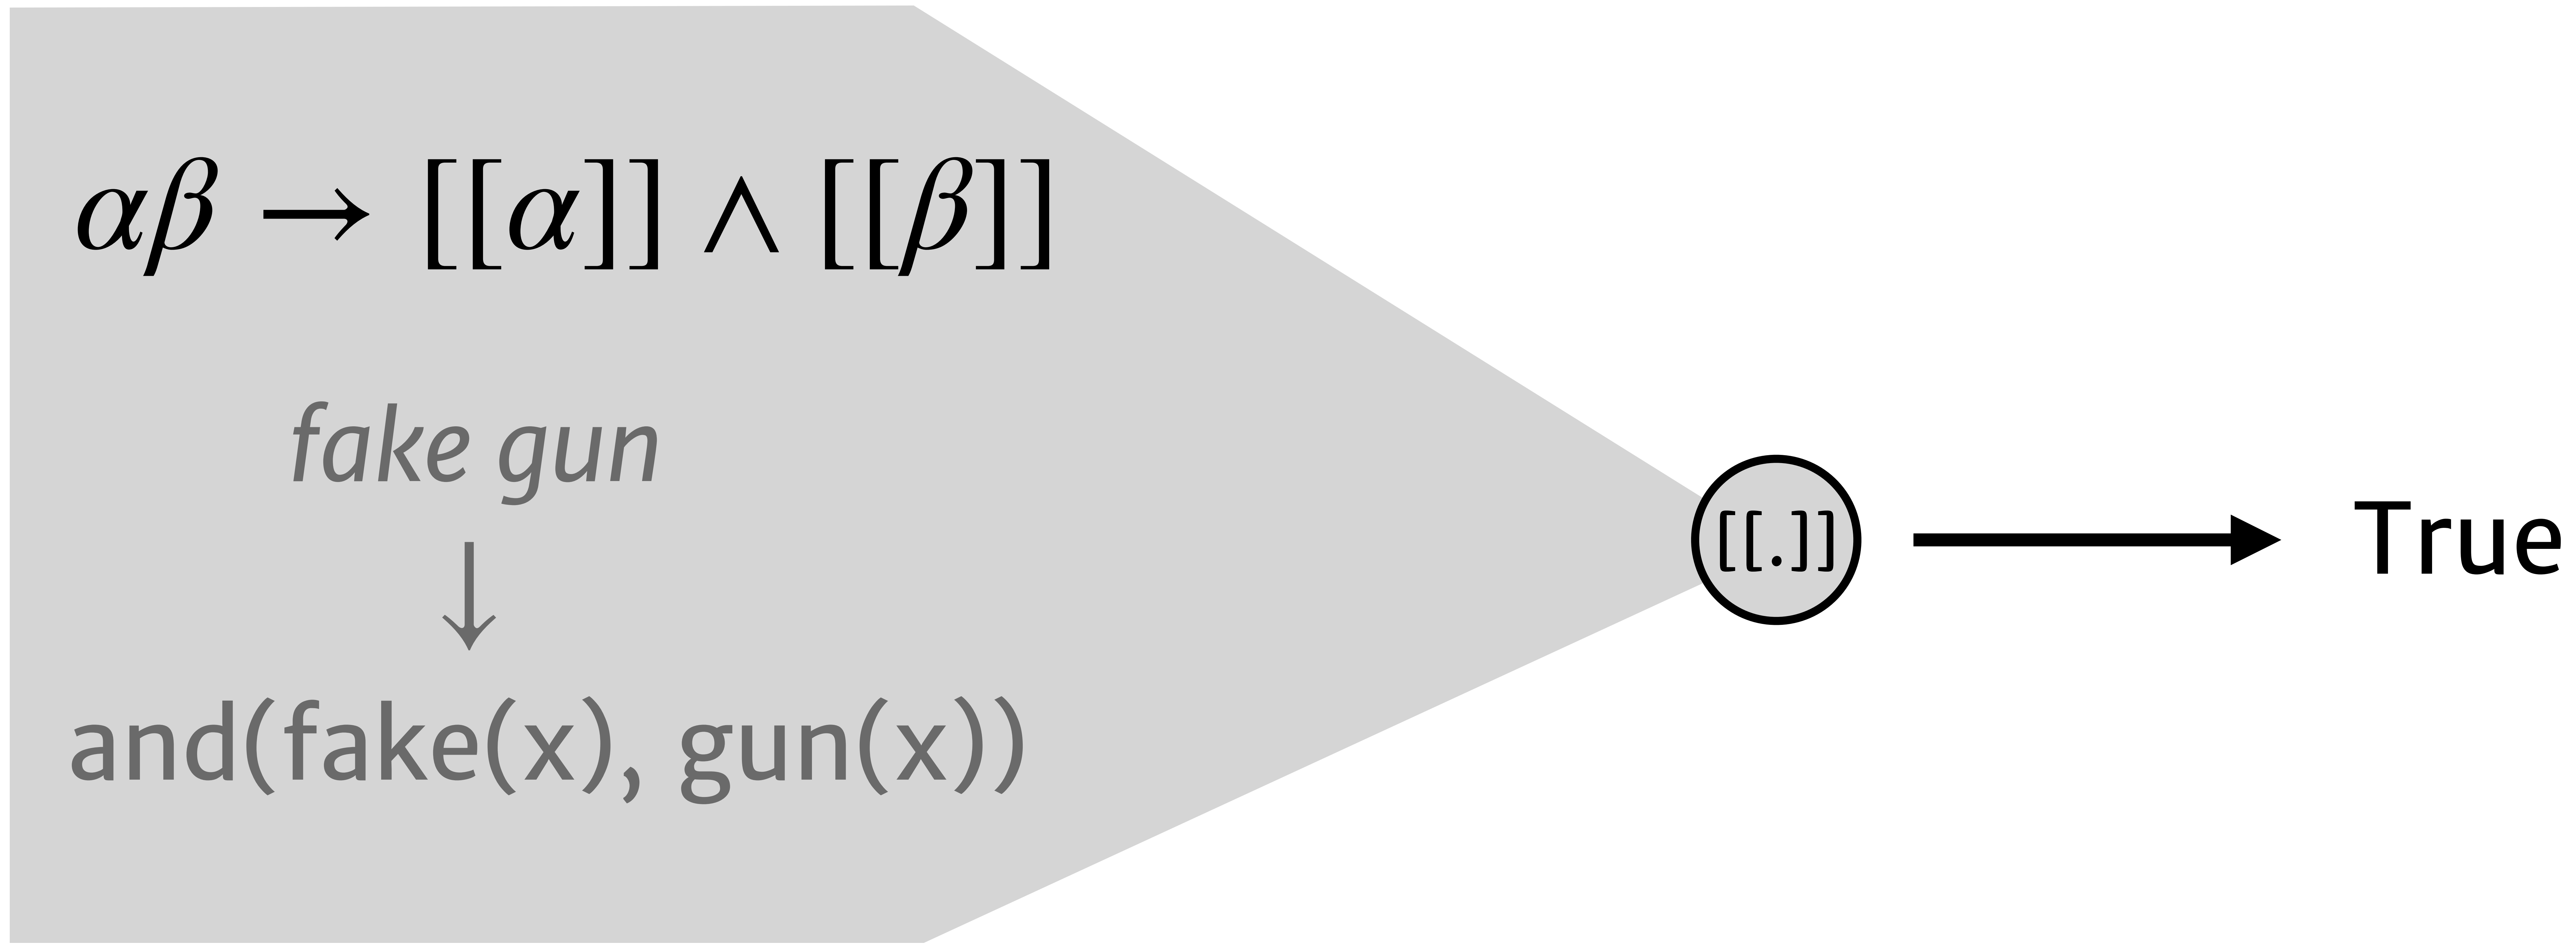

Structure-sensitive processing

The cat is on the mat.

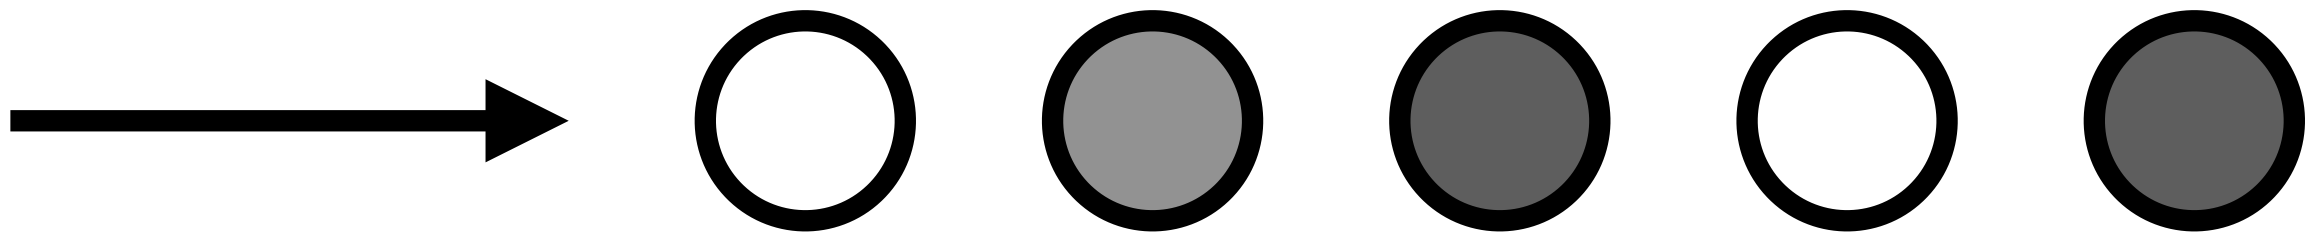

Structure-sensitive processing

The cat is on the mat.

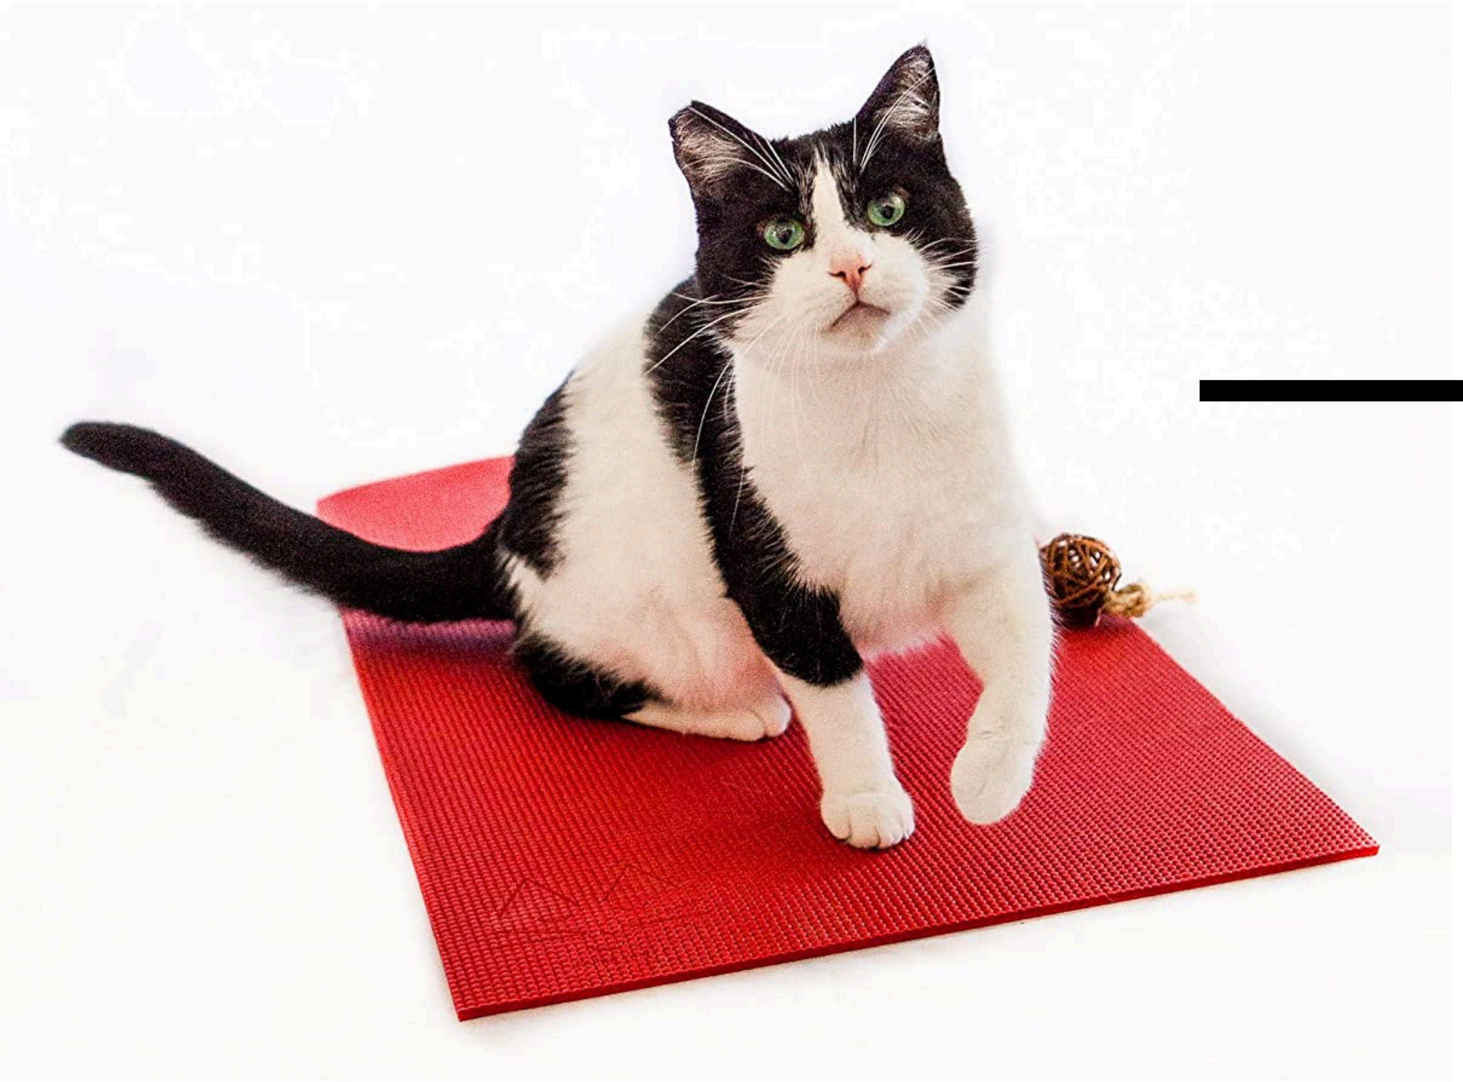

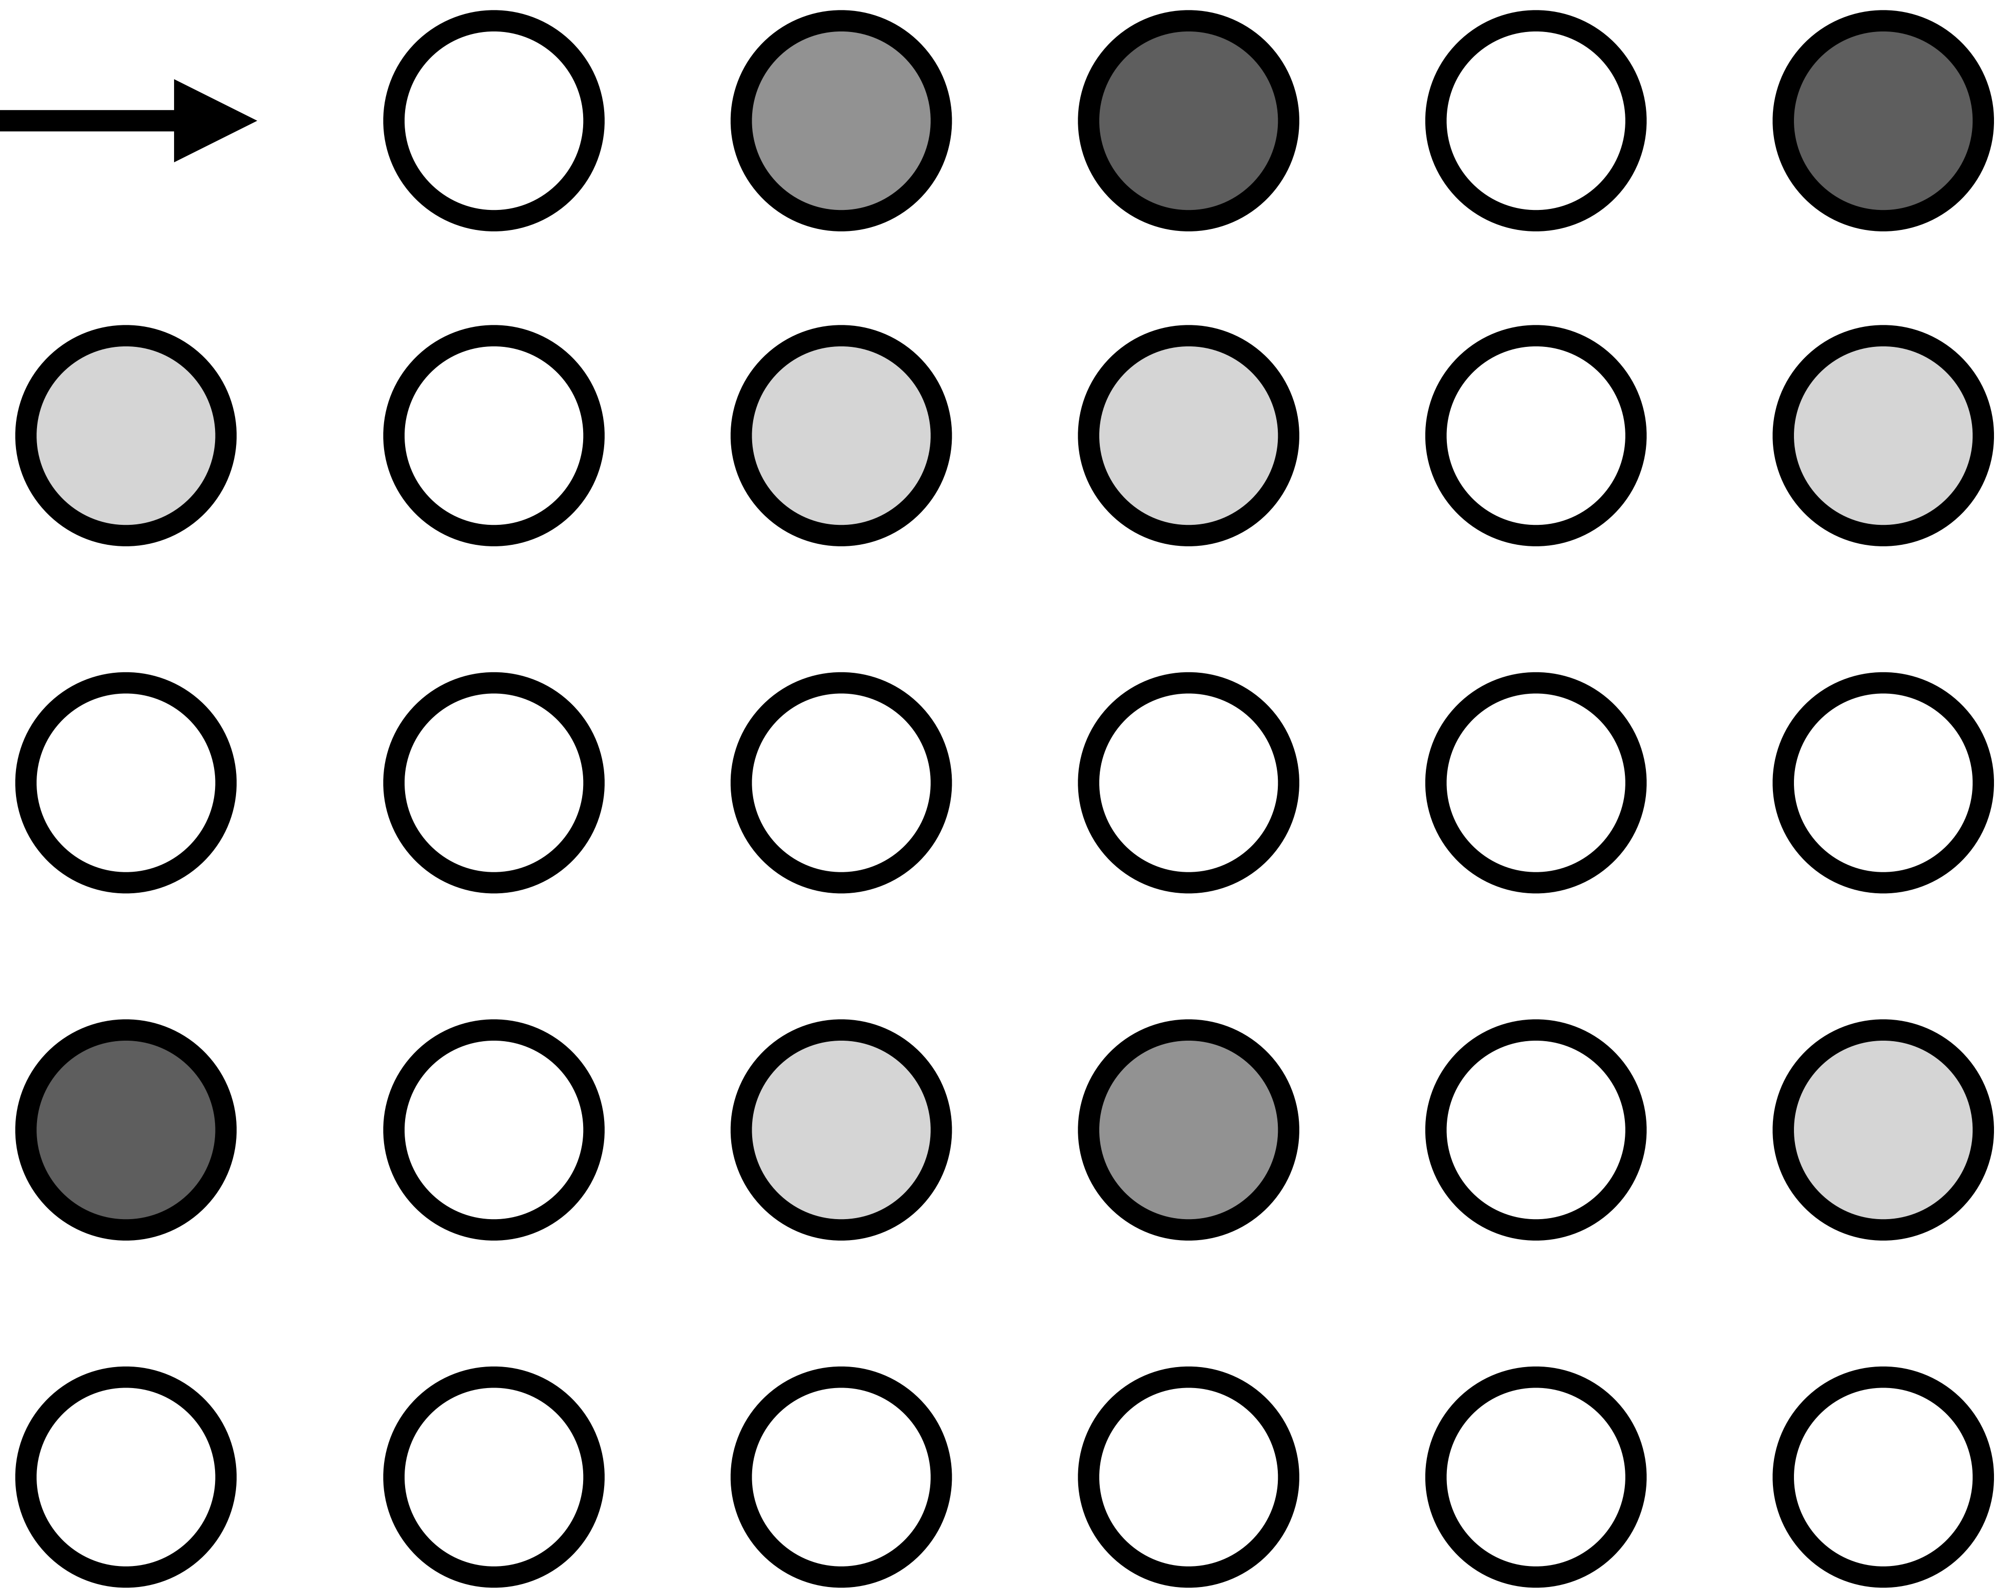

Structure–sensitive processing

The cat is on the mat.

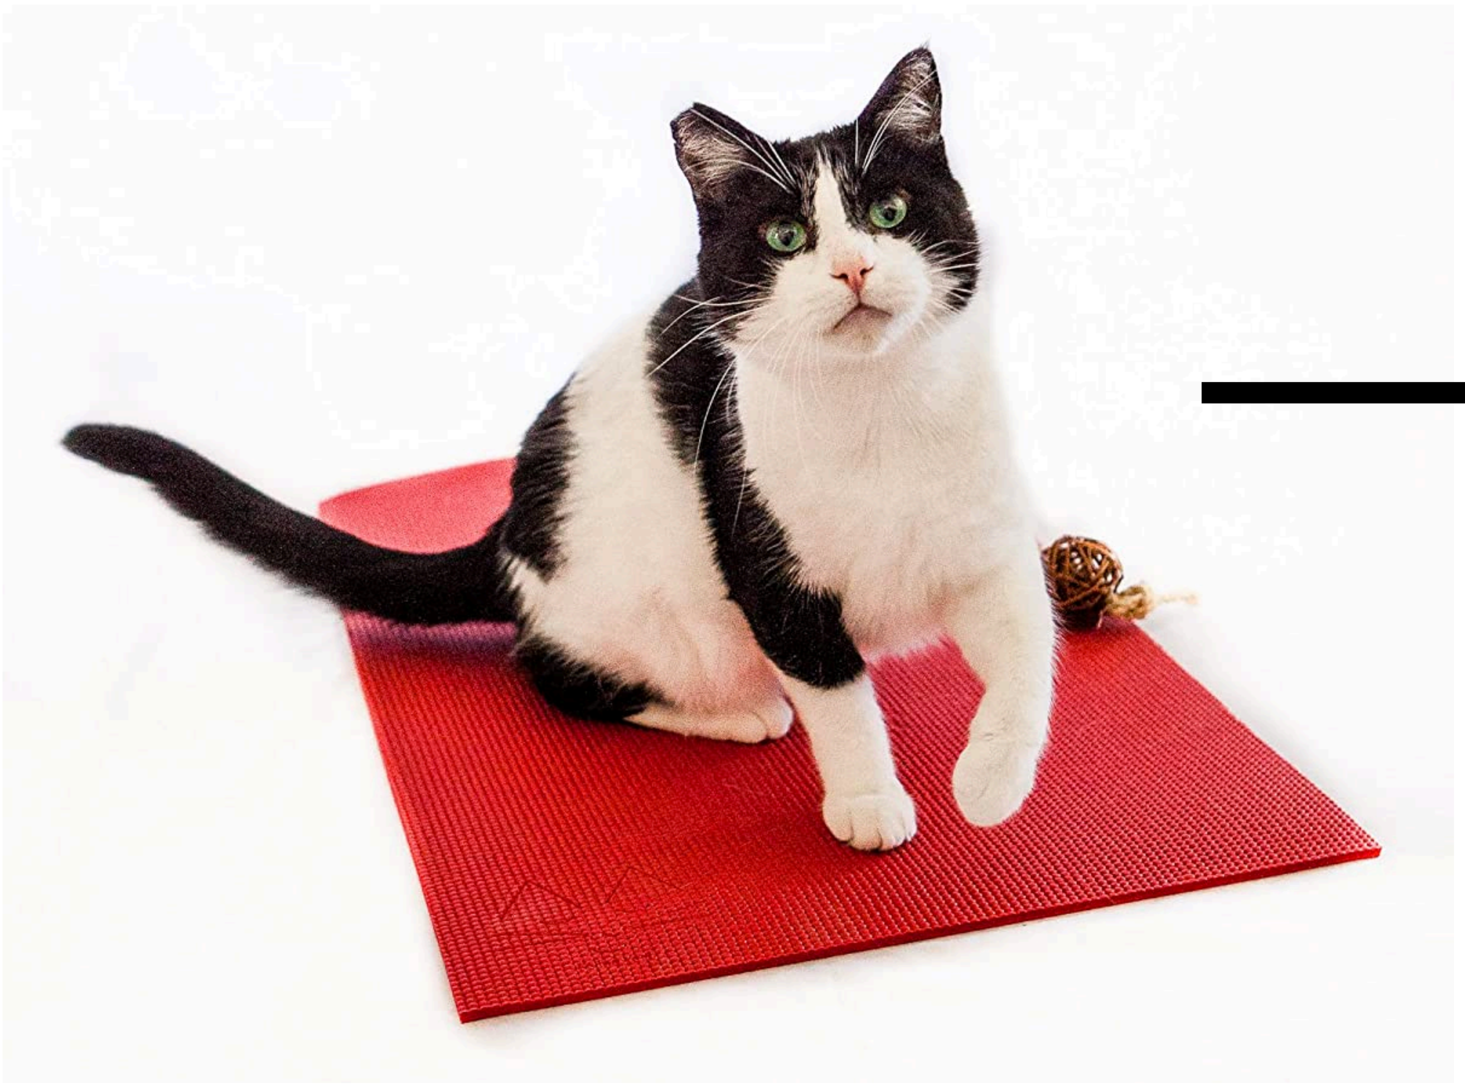

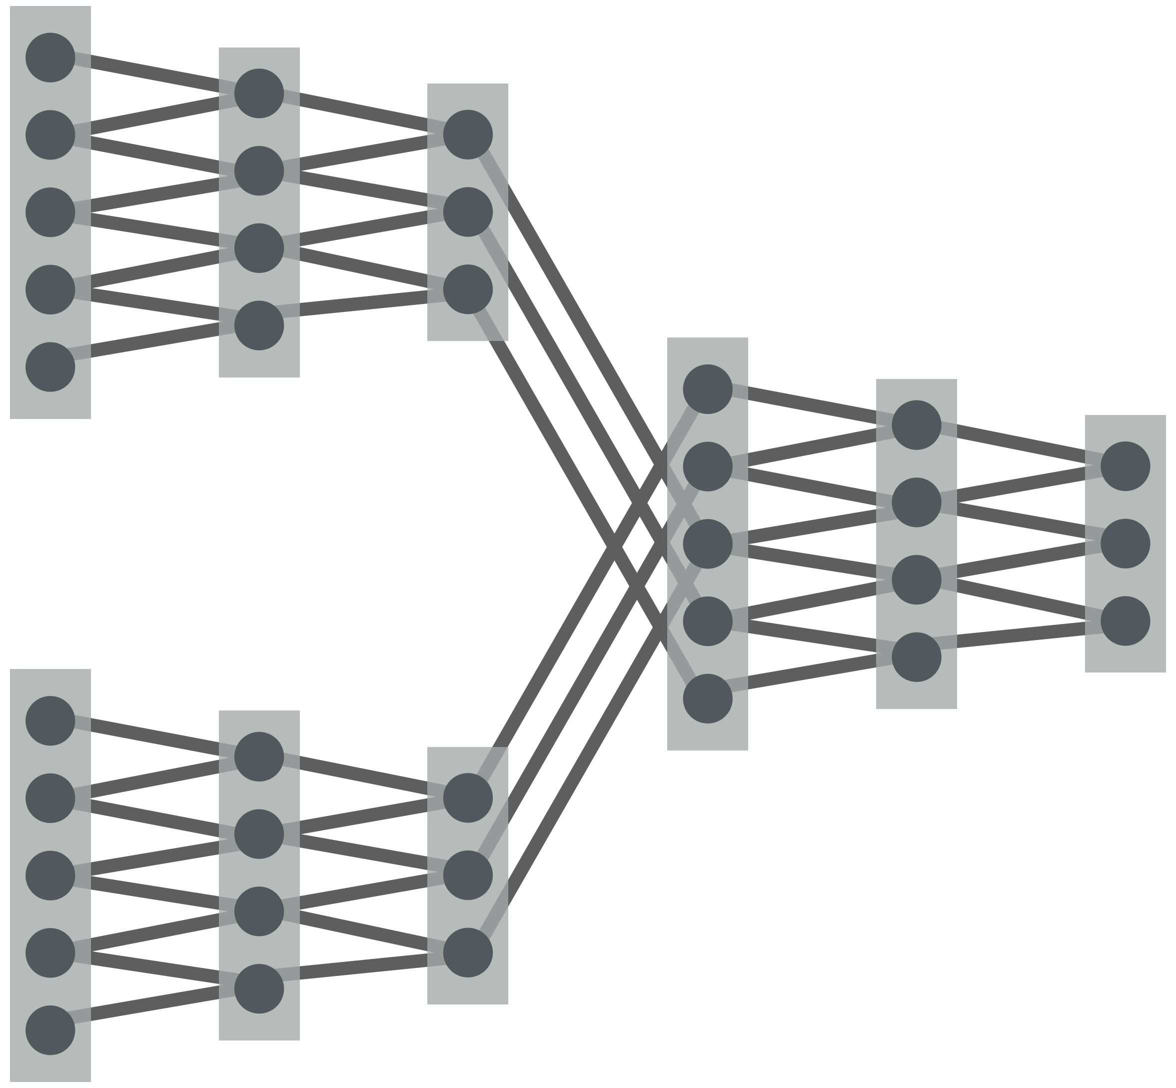

True

Structure–sensitive processing

The cat is on the mat. →

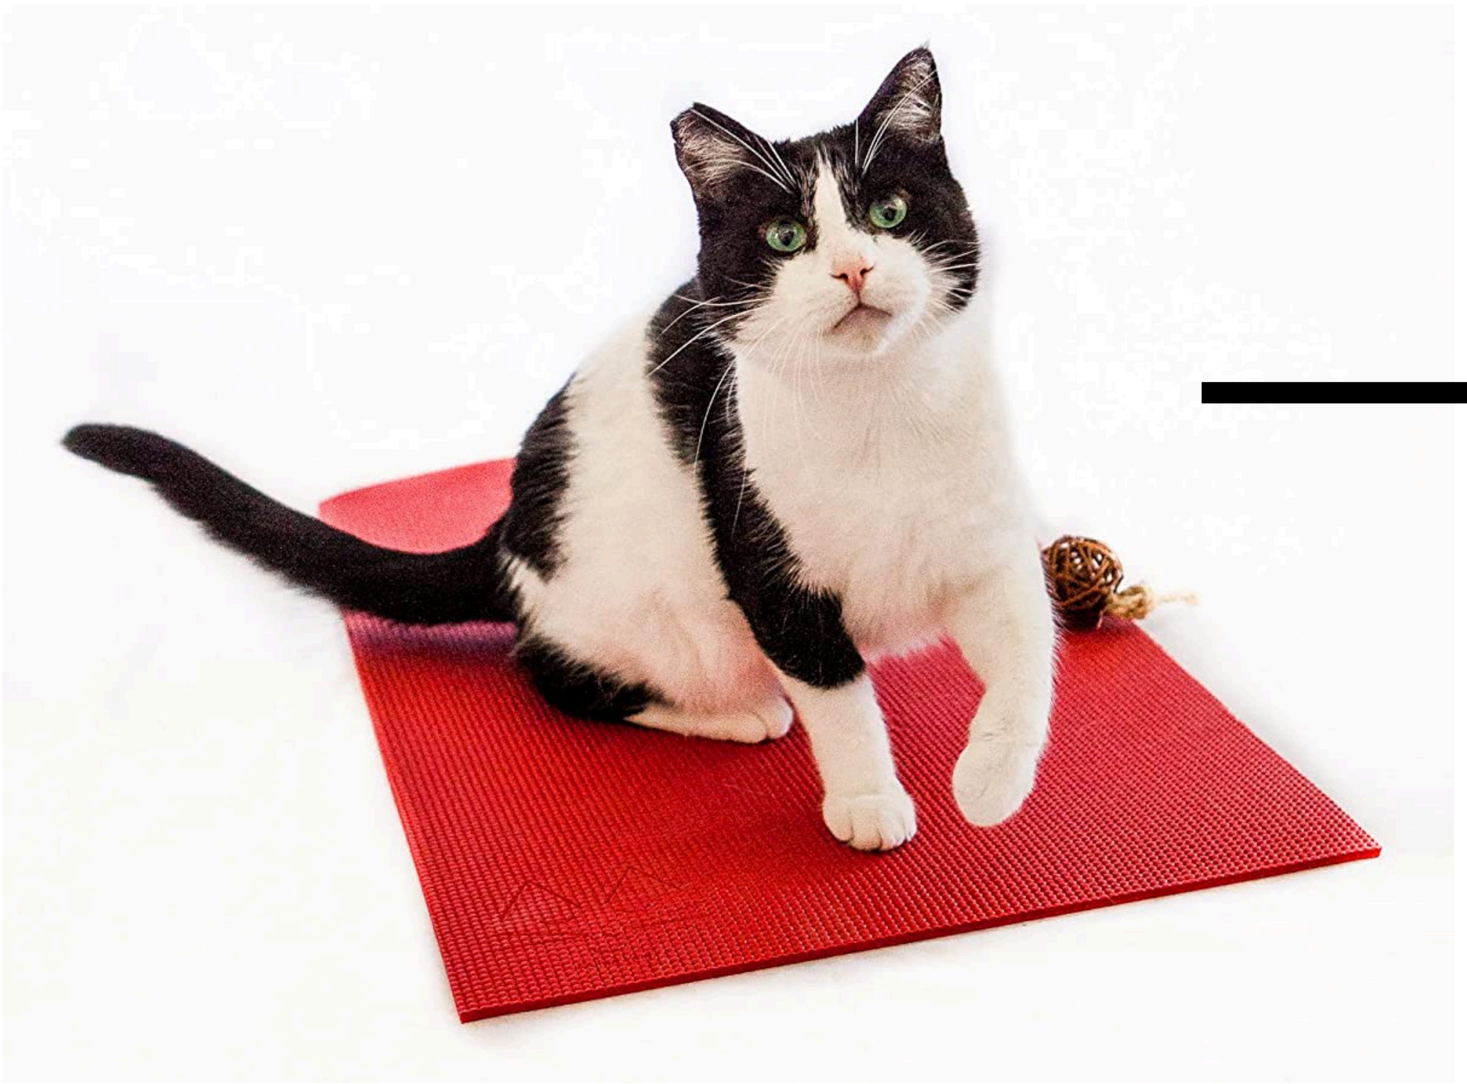

→

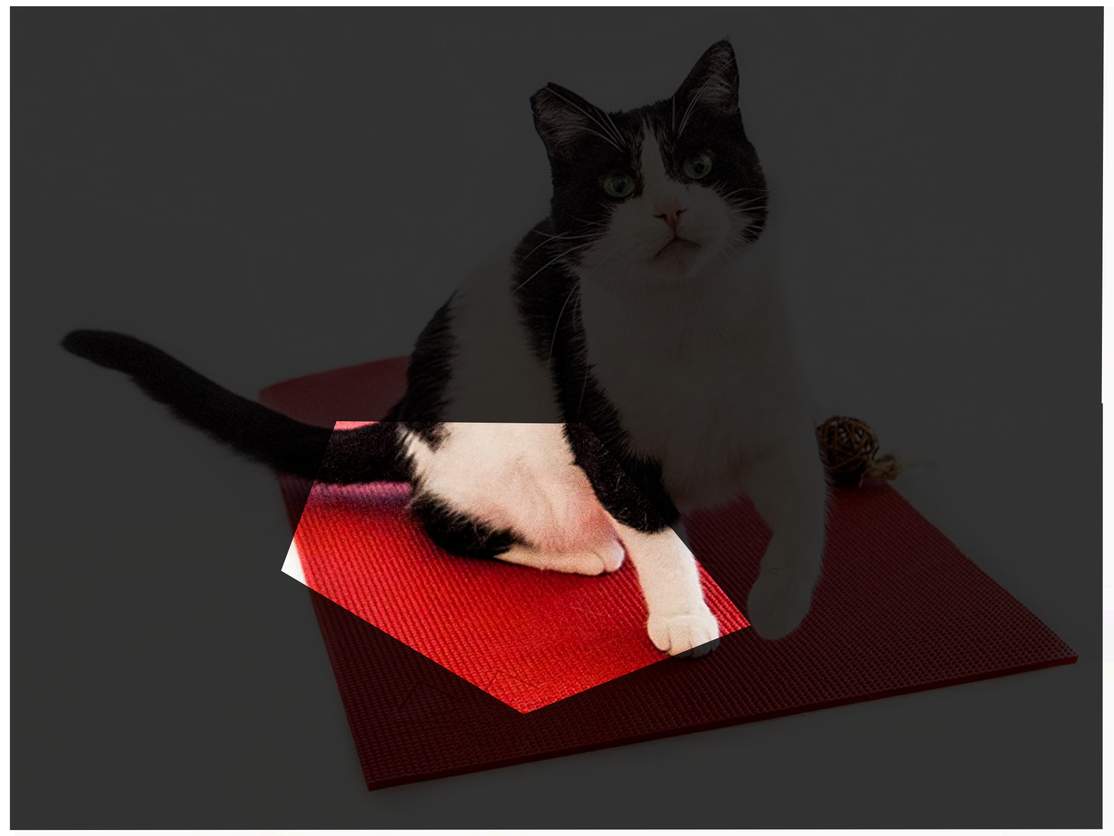

→ True

Discussion

Break

Linguistic productivity

“Infinite use of finite means”

W. von Humboldt

this is the dog that chased the cat that ate the rat that lived in the house that Jack built…

The competence/performance distinction

Chomsky 1965: *Linguistic theory is concerned primarily with an ideal speaker-listener, in a completely homogeneous speech-community, who knows its (the speech community's) language perfectly and is unaffected by such grammatically irrelevant conditions as memory limitations, distractions, shifts of attention and interest, and errors (random or characteristic) in applying his knowledge of this language in actual performance.*

Linguistic competence (including claims about productivity of language) concerns this idealized speaker.

The competence/performance distinction?

Labov 1971: *It is now evident to many linguists that the primary purpose of the [performance/competence] distinction has been to help the linguist exclude data which he finds inconvenient to handle.*

Productivity in classical models

Claim: like humans, the classical model can interpret arbitrarily complex sentences:

The cat is on the mat. → [[*The cat*] [*is* [*on the mat*]]]

↓

`cat(x), mat(y), on(x, y)`

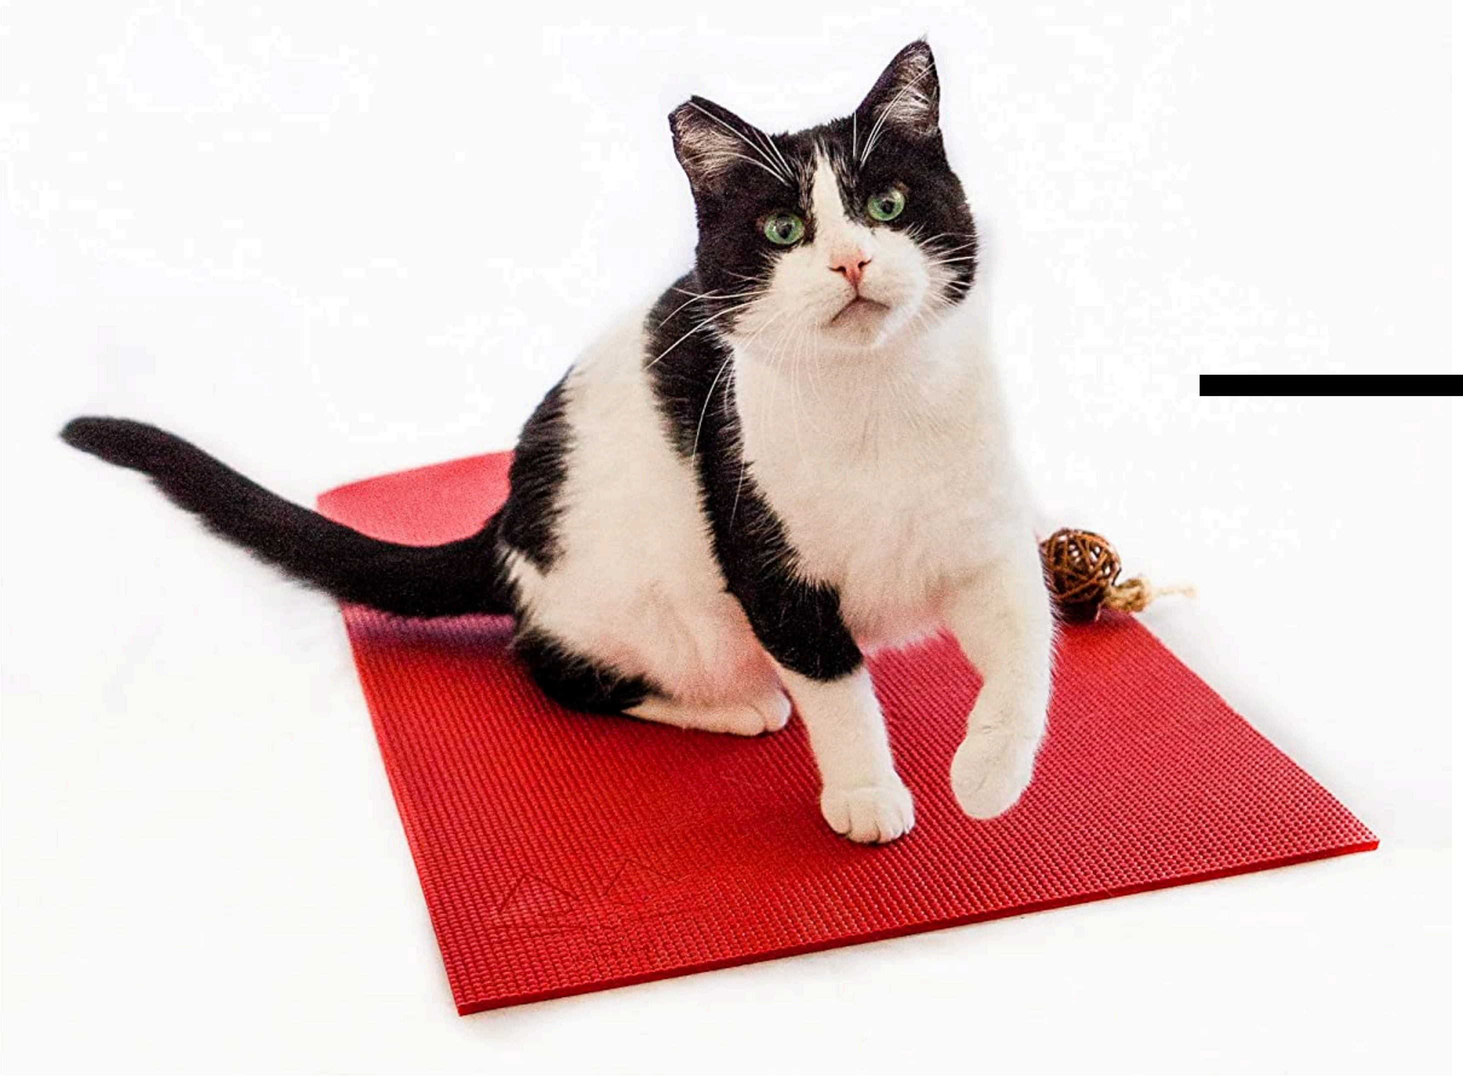

→

```
cat(x)
mat(y)
red(y)
on(x, y)
…
```

→ $\vdash$ → True

Productivity in classical models

Claim: like humans, the classical model can interpret arbitrarily complex sentences:

The cat is on the mat. → [[*The cat*] [*is* [*on the mat*]]]

↓

`cat(x), mat(y), on(x, y)`

↓

→ `cat(x)`

Need more processing power? Just add RAM!

…

Productivity in connectionist models

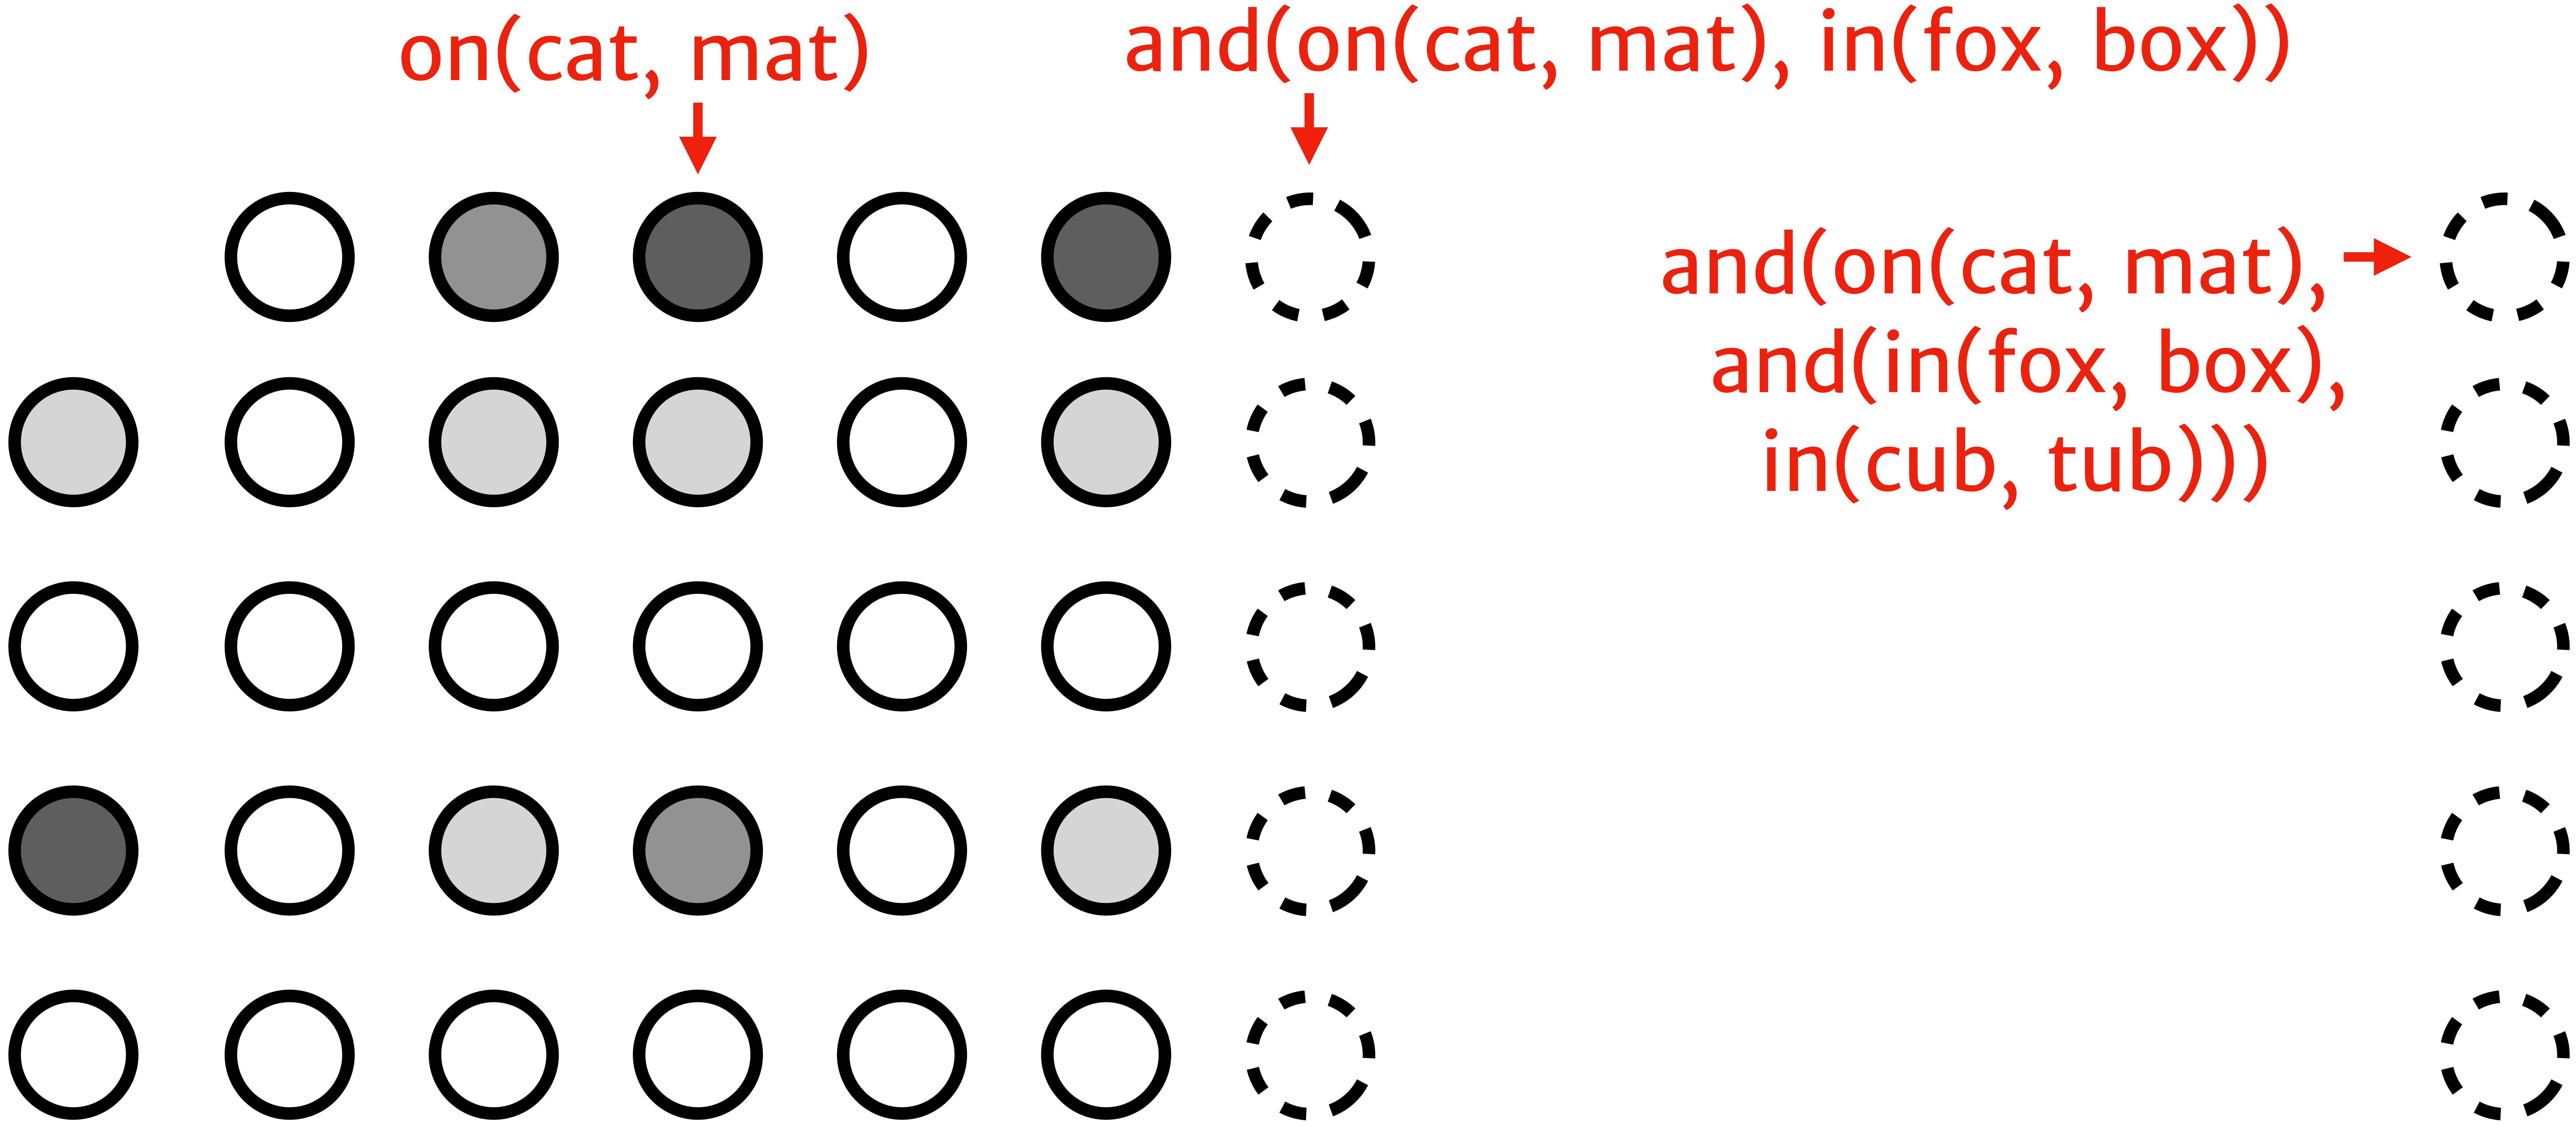

[Ray Mooney, ca. 2014]

Productivity in neural models

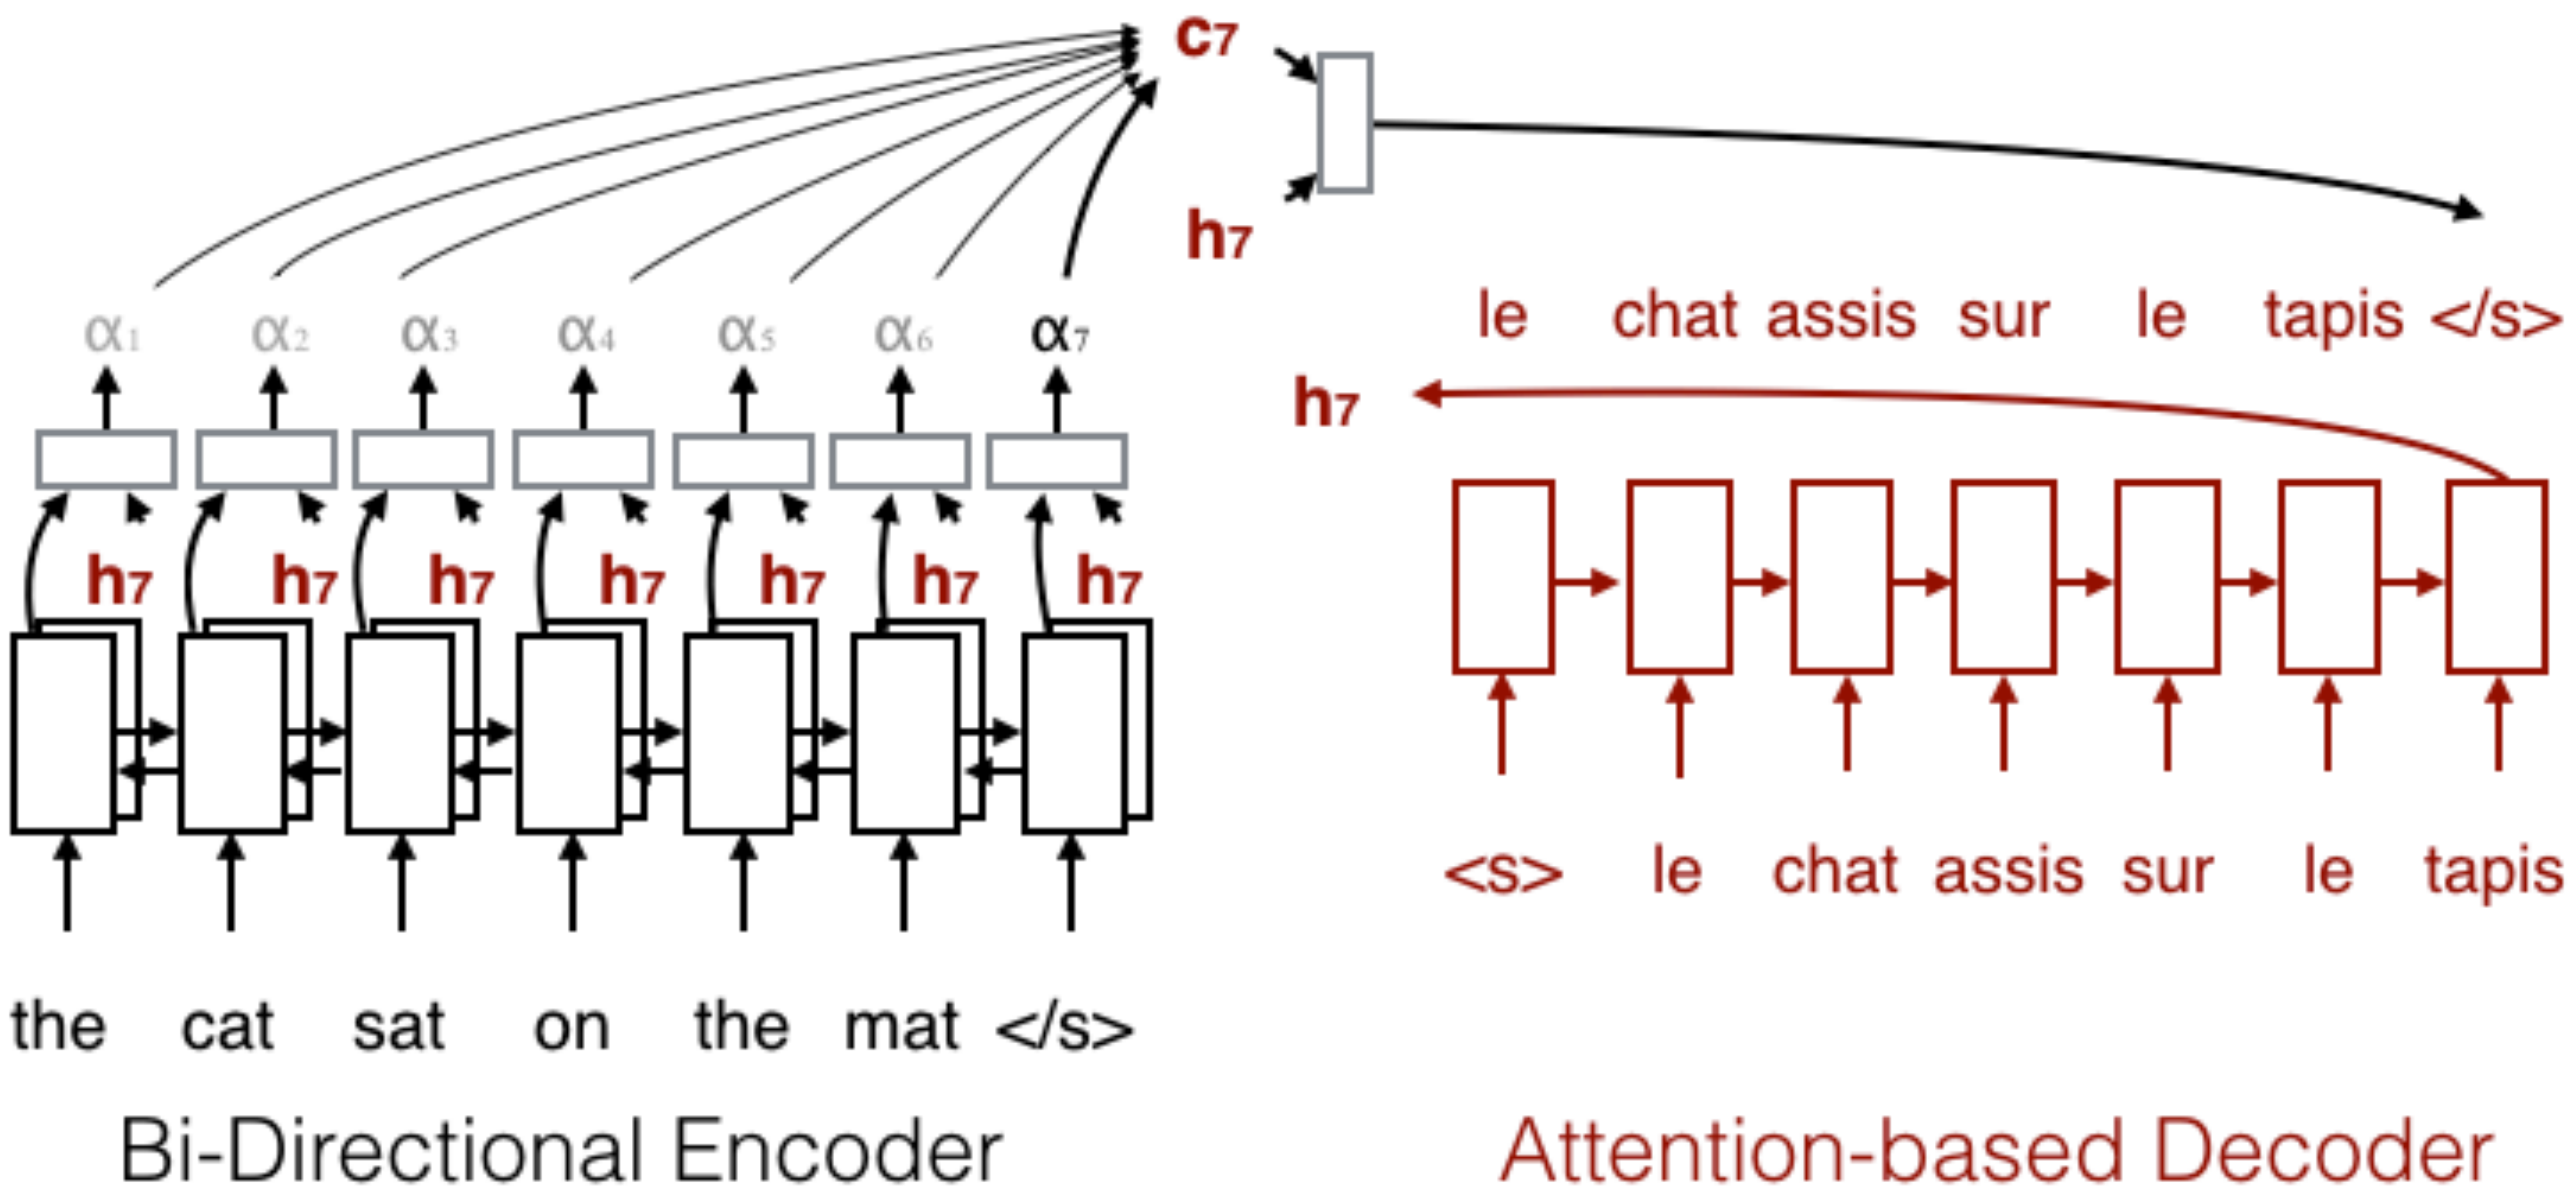

[Bahdanau 2015]

Productivity in neural models

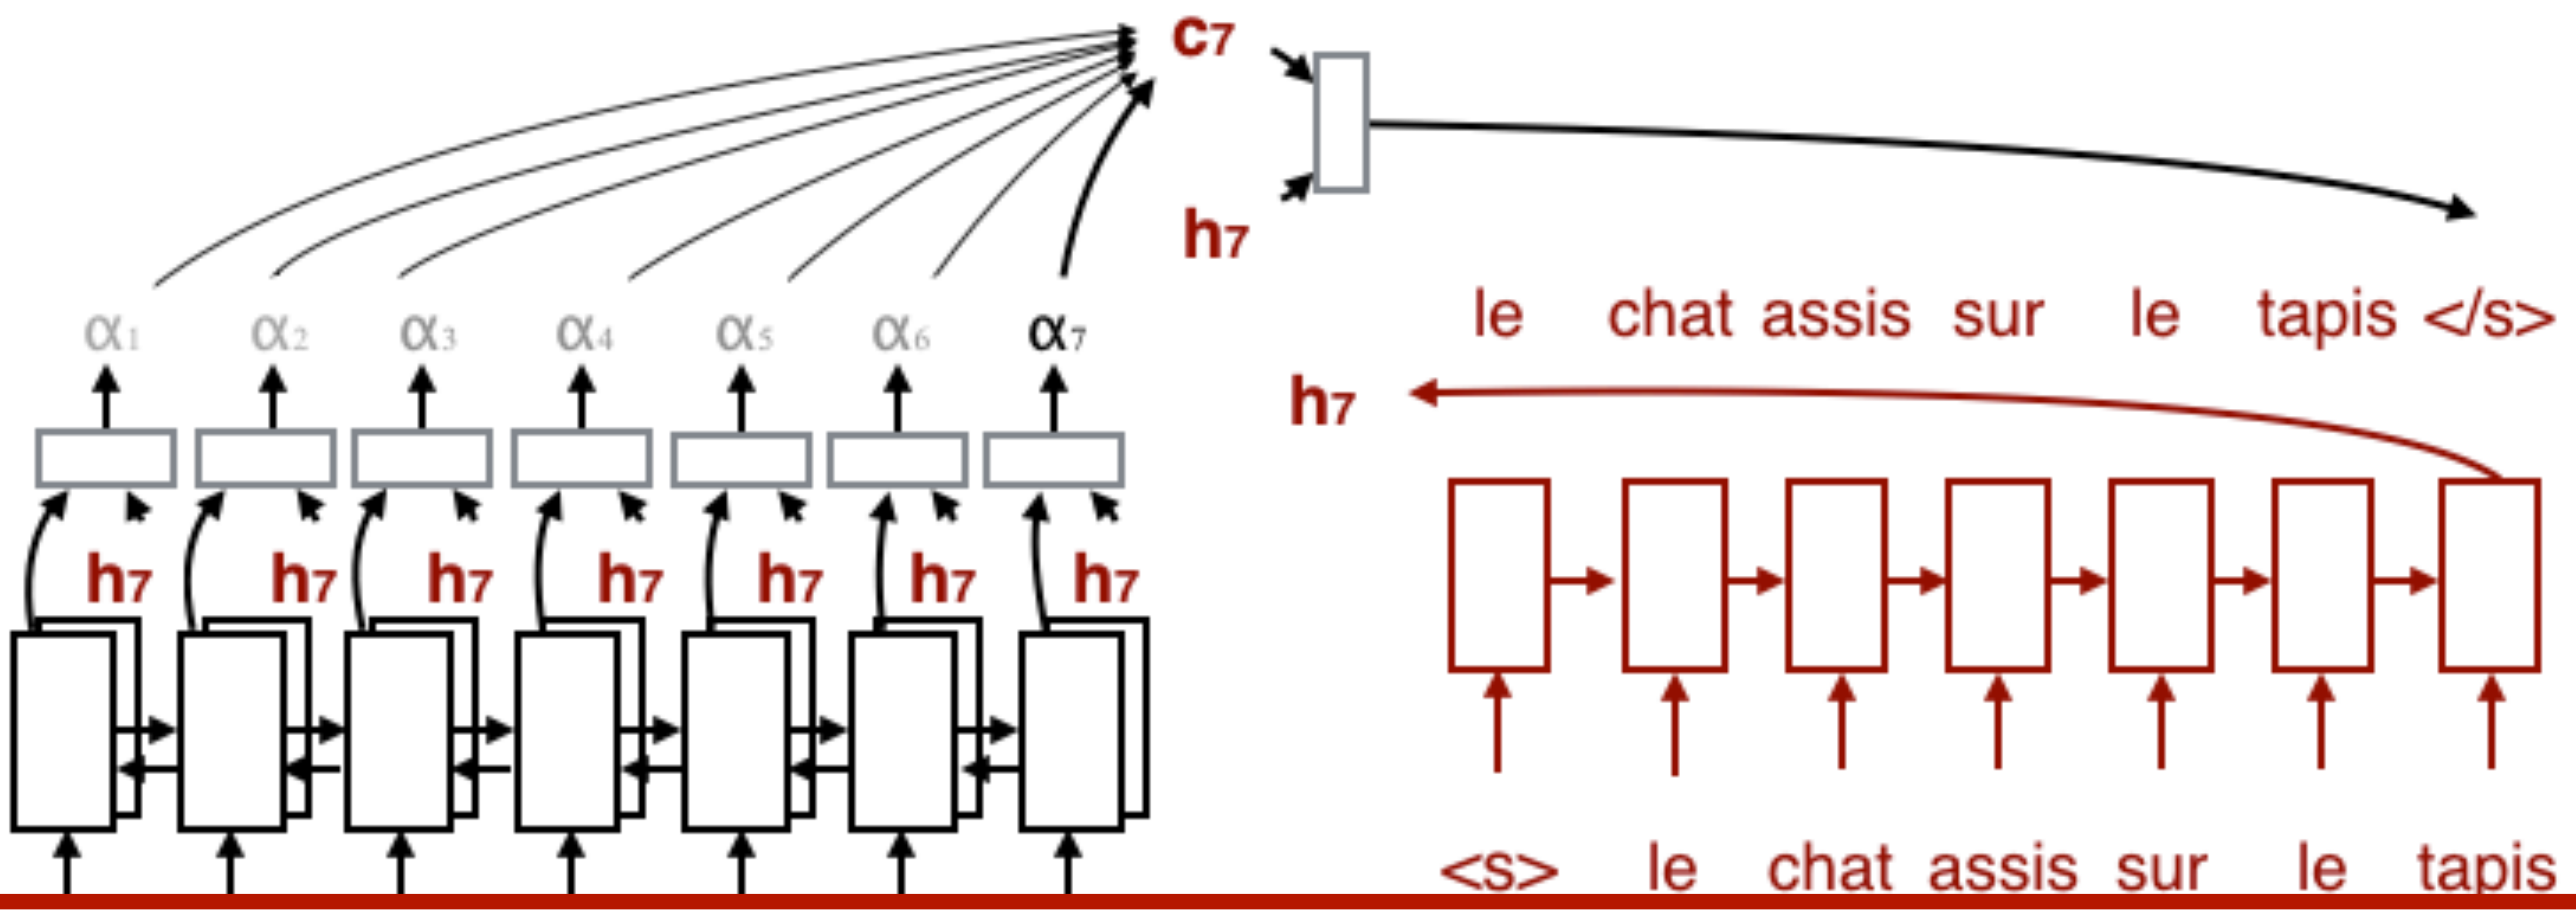

Need more processing power? Just add steps/layers/precision!

[Bahdanau 2015]

Logical labels for neurons

Unit 314 `operating room OR castle OR bathroom`
IoU 0.05

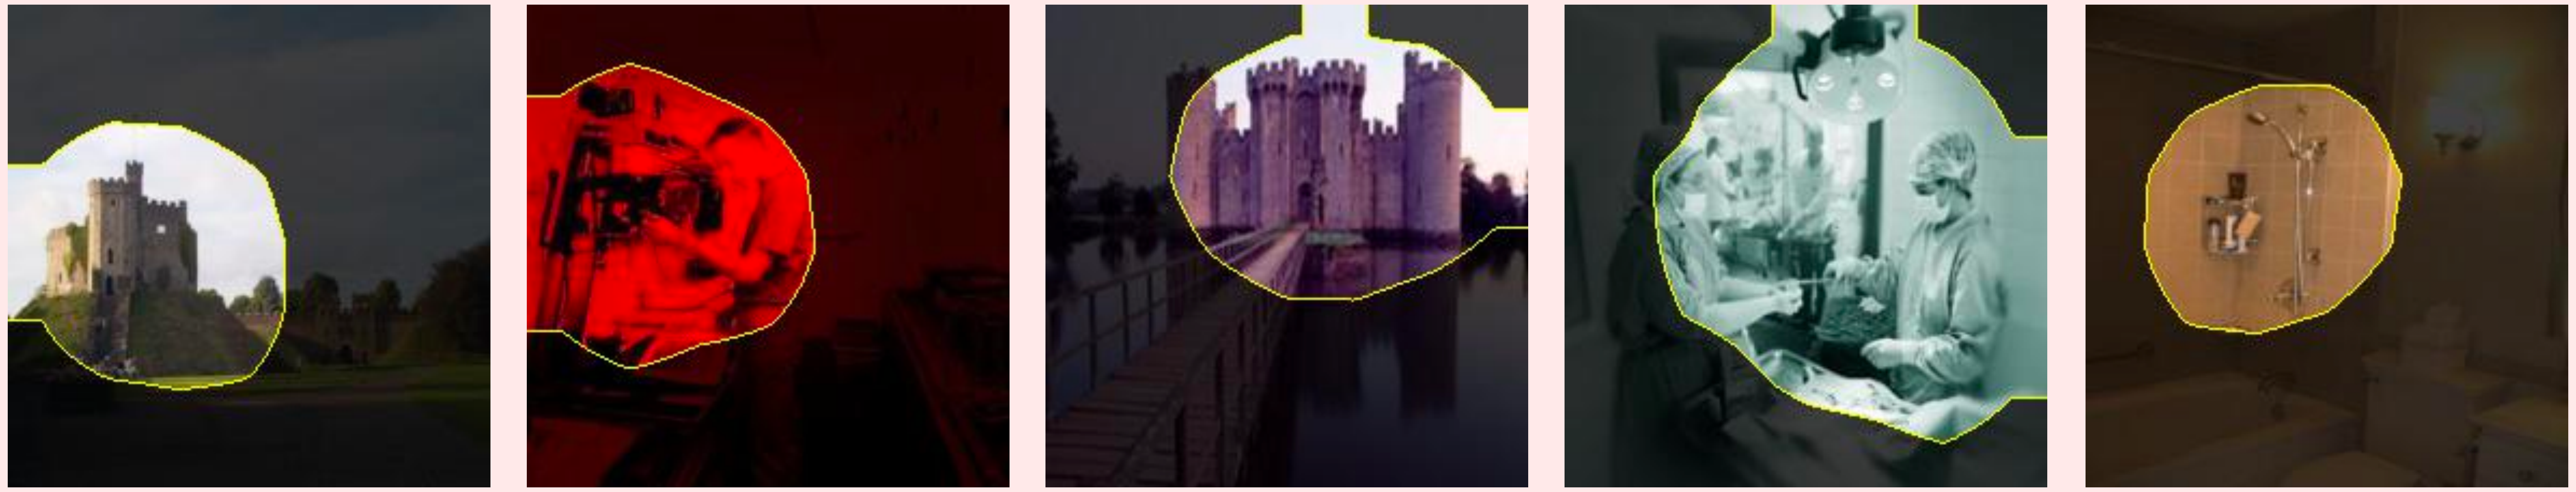

Unit 439 `bakery OR bank vault OR shopfront`
IoU 0.08

[Mu and Andreas 2020; c.f. Bau et al. 2017, Dalvi et al. 2018]

Logical labels for neurons

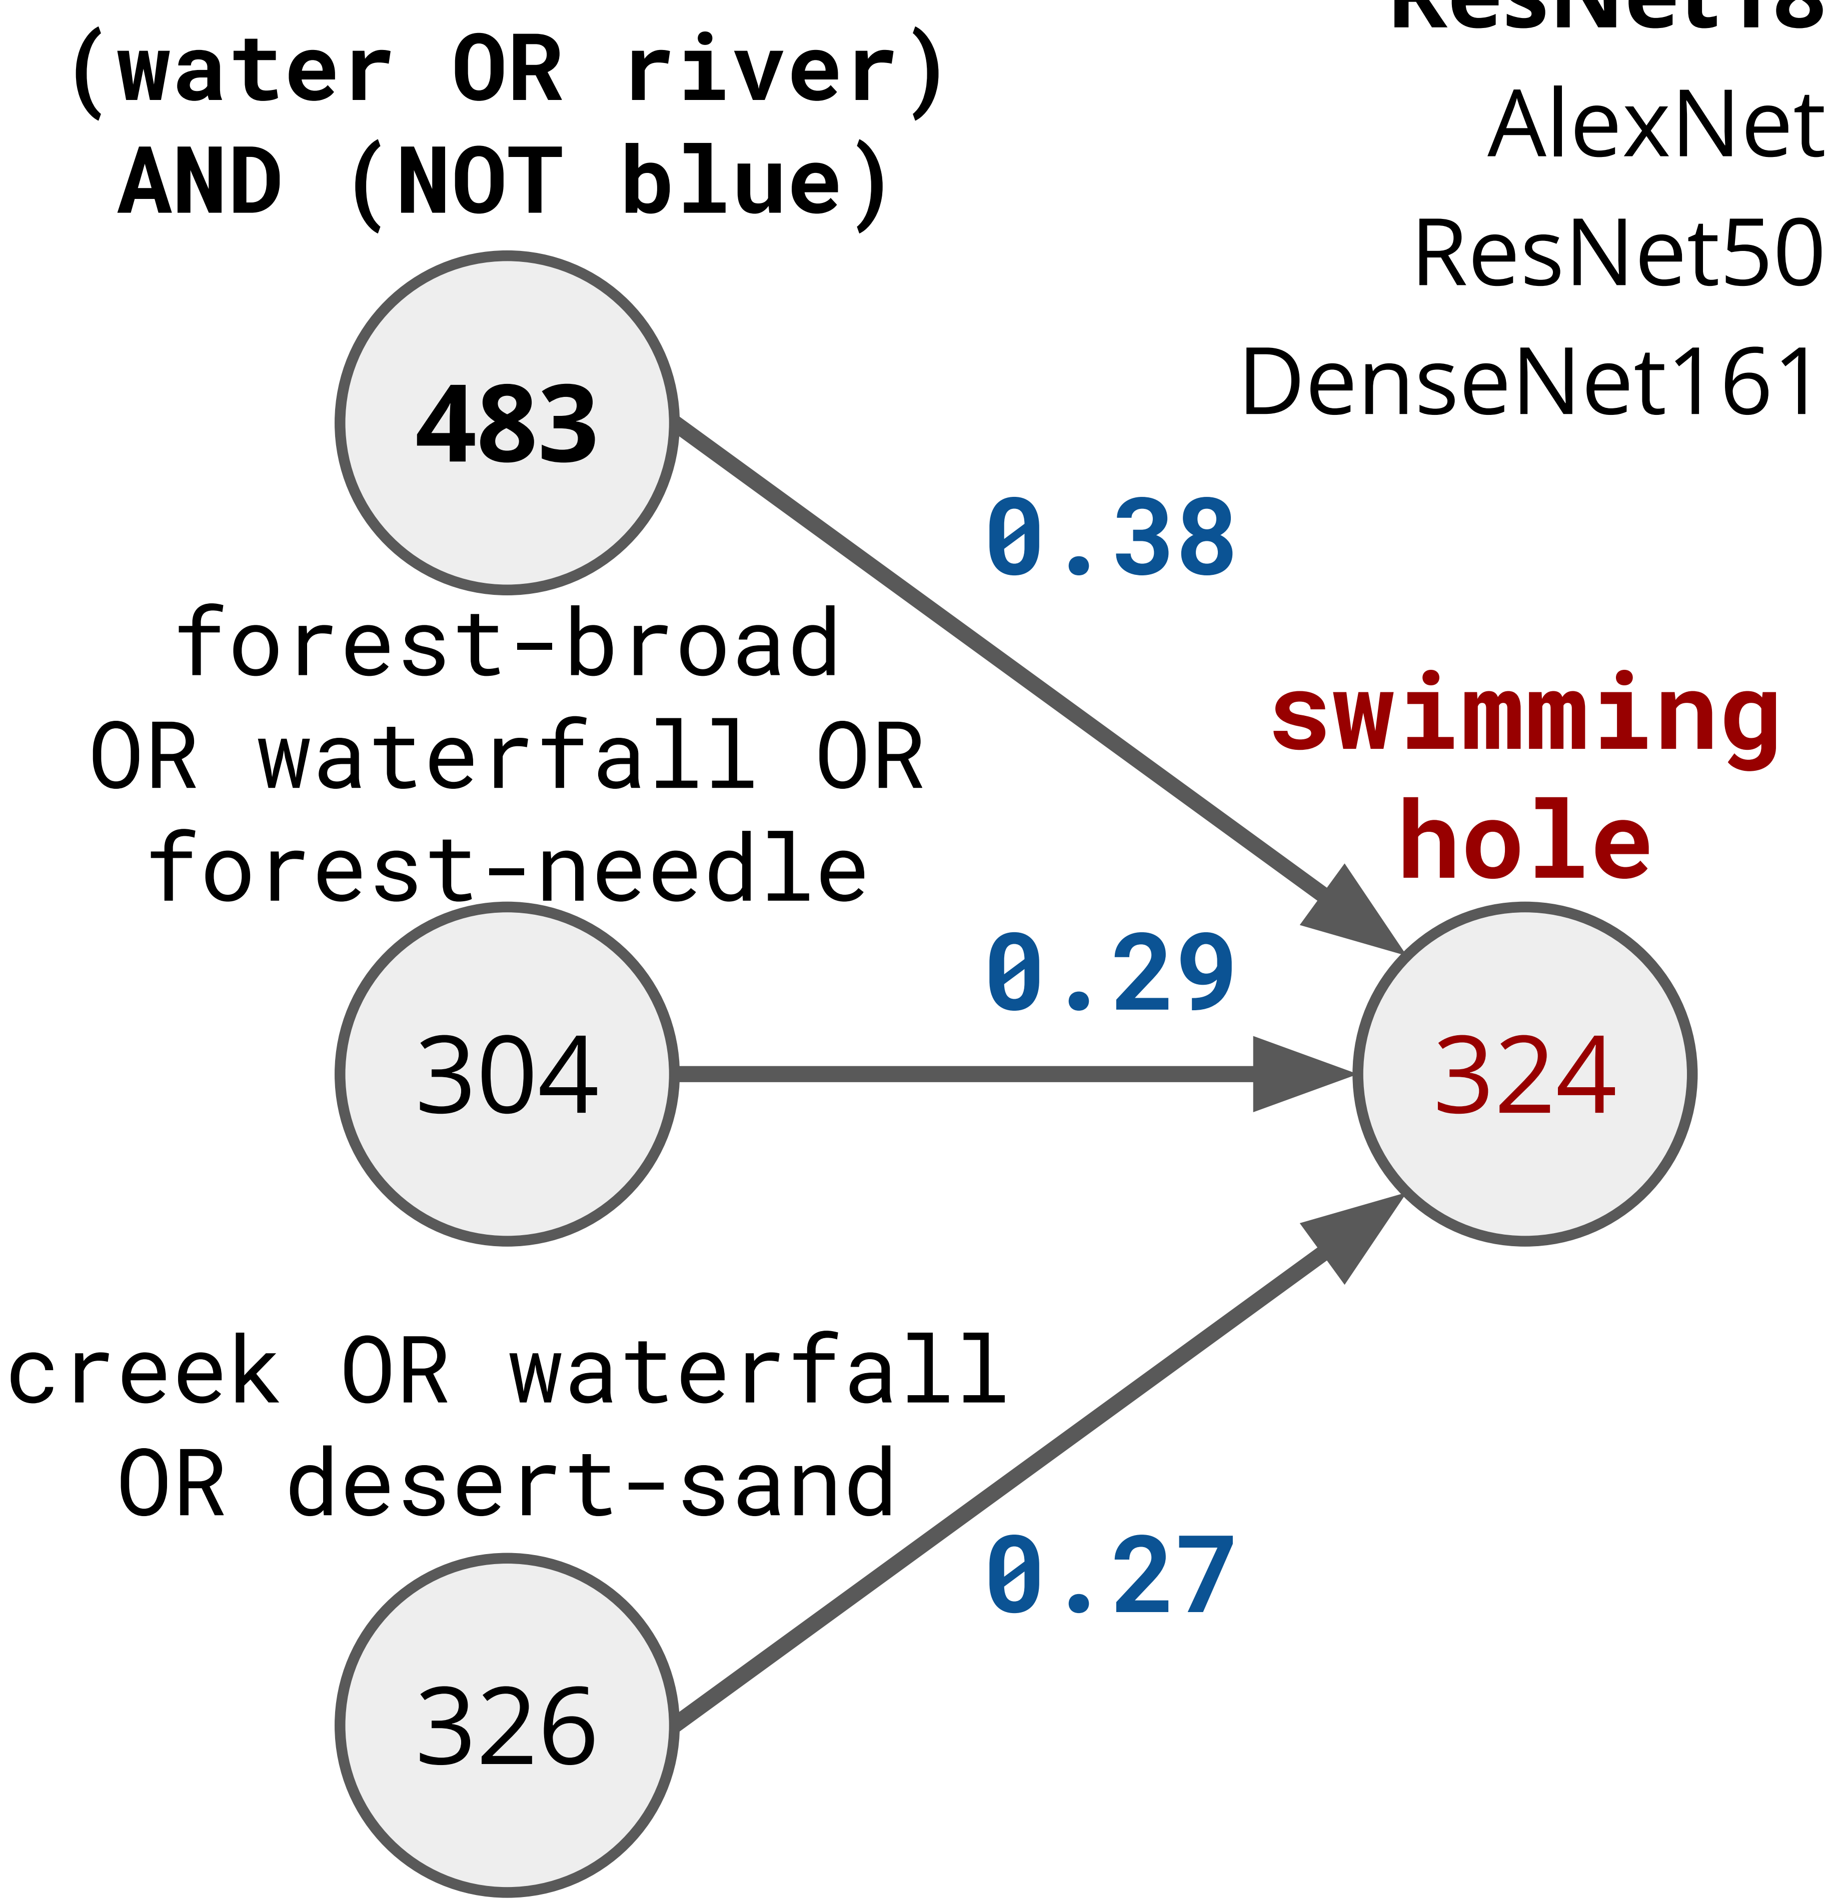

ResNet18	swimming hole	**grotto**
AlexNet	swimming hole	**grotto**
ResNet50	swimming hole	**grotto**
DenseNet161	swimming hole	hot spring

[Mu and Andreas 2020; c.f. Bau et al. 2017, Dalvi et al. 2018]

Discussion

Systematicity

F&P: *What we mean when we say that linguistic capacities are* systematic *is that the ability to produce / understand some sentences is* intrinsically *connected to the ability to produce / understand certain others.*

the cat is on the mat ⟷ *the mat is on the cat*

Systematicity

$$\text{NP} \rightarrow \text{NP V PP}$$

$$\text{PP} \rightarrow \text{on NP}$$

$$\text{NP} \rightsquigarrow \textit{the cat sat on the mat}$$

Systematicity

$\mathrm{NP} \rightarrow \mathrm{NP\ V\ PP}$

$\mathrm{PP} \rightarrow \mathrm{on\ NP}$

$\mathrm{NP} \rightsquigarrow$ *the cat sat on the mat*

$\Rightarrow \mathrm{NP} \rightsquigarrow$ *the mat sat on the cat*

Connectionist models *permit* non-systematicity

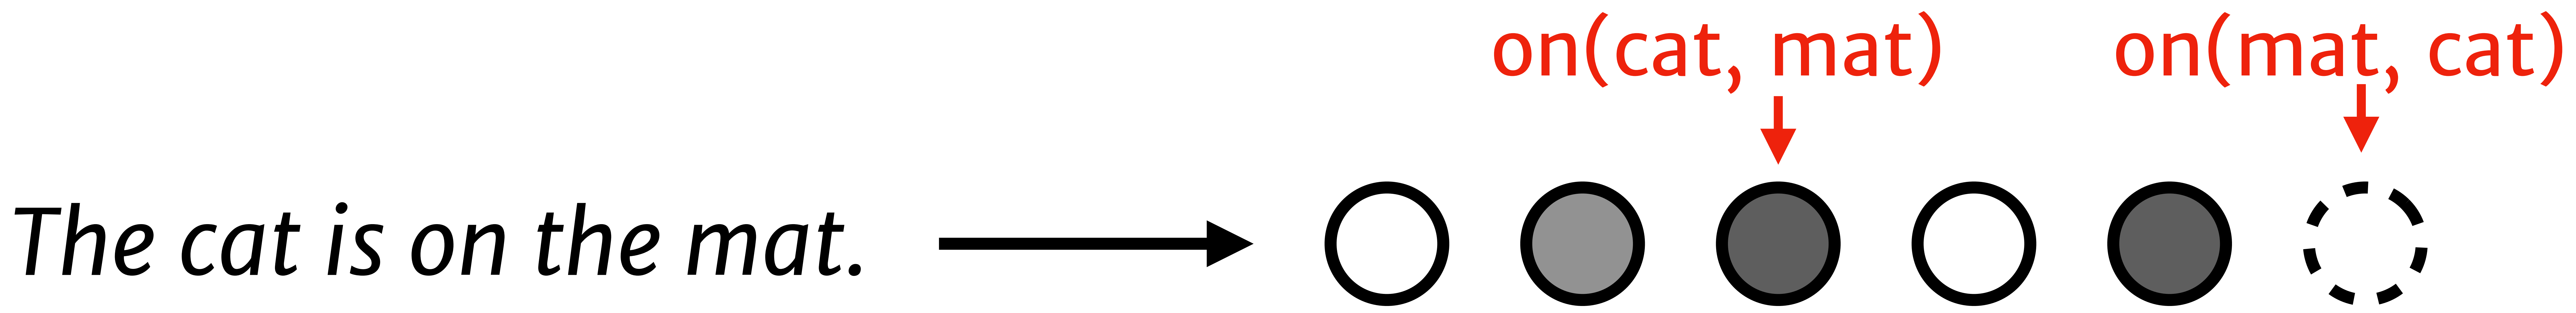

(but so do classical ones)

NP → NP1 V PP

PP → on NP2

NP ⇝ *the cat sat on the mat*

⇏ NP ⇝ *the mat sat on the cat*

(but so do classical ones)

$\mathrm{NP} \to \mathrm{NP1}\ \mathrm{V}\ \mathrm{PP}$

$\mathrm{PP} \to \text{on}\ \mathrm{NP2}$

$\mathrm{NP} \rightsquigarrow$ *the cat sat on itself*

$\not\Rightarrow \mathrm{NP} \rightsquigarrow$ **itself sat on the cat*

Takeaway

Systematicity is a property of a **parameterization**, not just a model class!

Discussion

F&P's conclusions

F&P: By contrast, since the Connectionist architecture recognizes no combinatorial structure in mental representations, gaps in cognitive competence should proliferate arbitrarily. It's not just that you'd expect to get them from time to time; it's that, on the 'no-structure' story, *gaps are the unmarked case*. It's the *systematic* competence that the theory is required to treat as an embarrassment. But, as a matter of fact, inferential competences are *blatantly* systematic. So there must be something deeply wrong with Connectionist architecture.

F&P's conclusions

F&P: By contrast, since the Connectionist architecture recognizes no combinatorial structure in mental representations, gaps in cognitive competence should proliferate arbitrarily. It’s not just that you’d expect to get them from time to time; it’s that, on the ‘no-structure’ story, *gaps are the unmarked case*. It's the *systematic* competence that the theory is required to treat as an embarrassment. But, as a matter of fact, inferential competences are *blatantly* systematic. So there must be something deeply wrong with Connectionist architecture.

[...but] we have no objection at all to networks as potential implementation models, nor do we suppose that any of the arguments we’ve given are incompatible with this proposal.

The worst RNN in the world

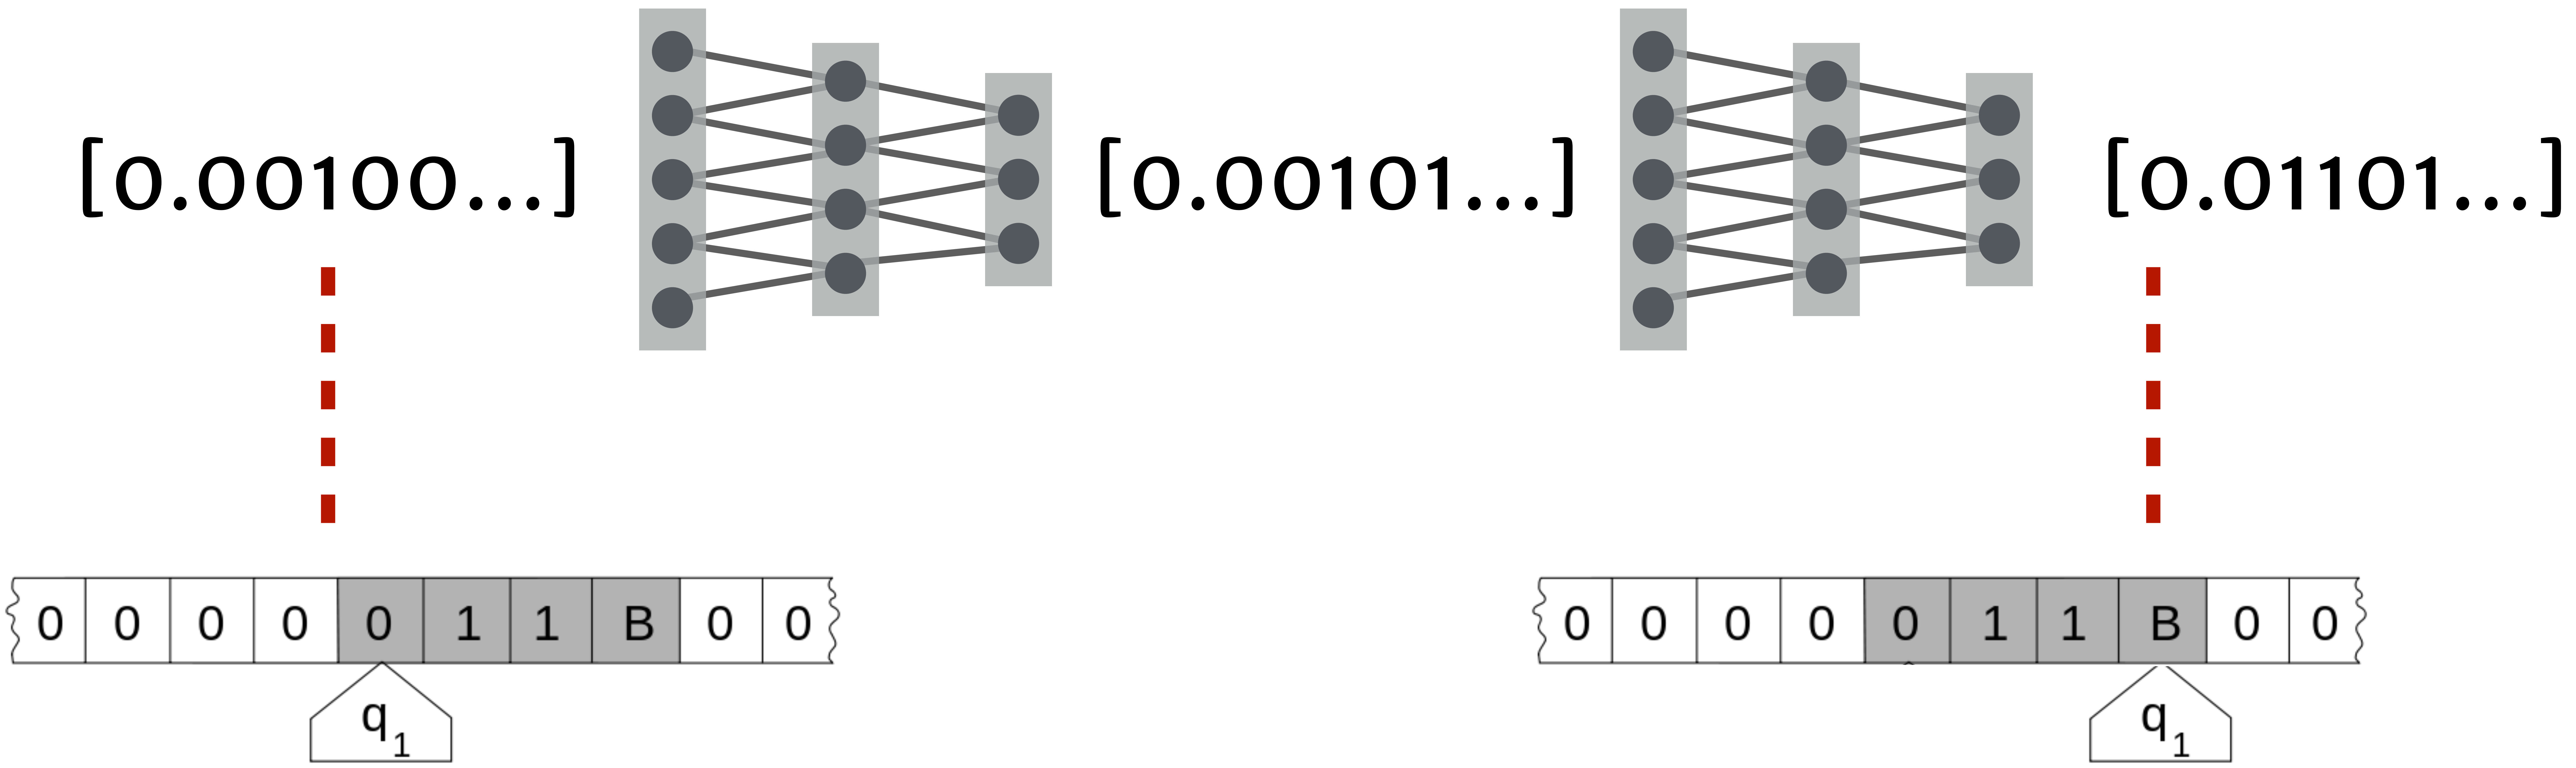

More realistic connectionist symbol processing?

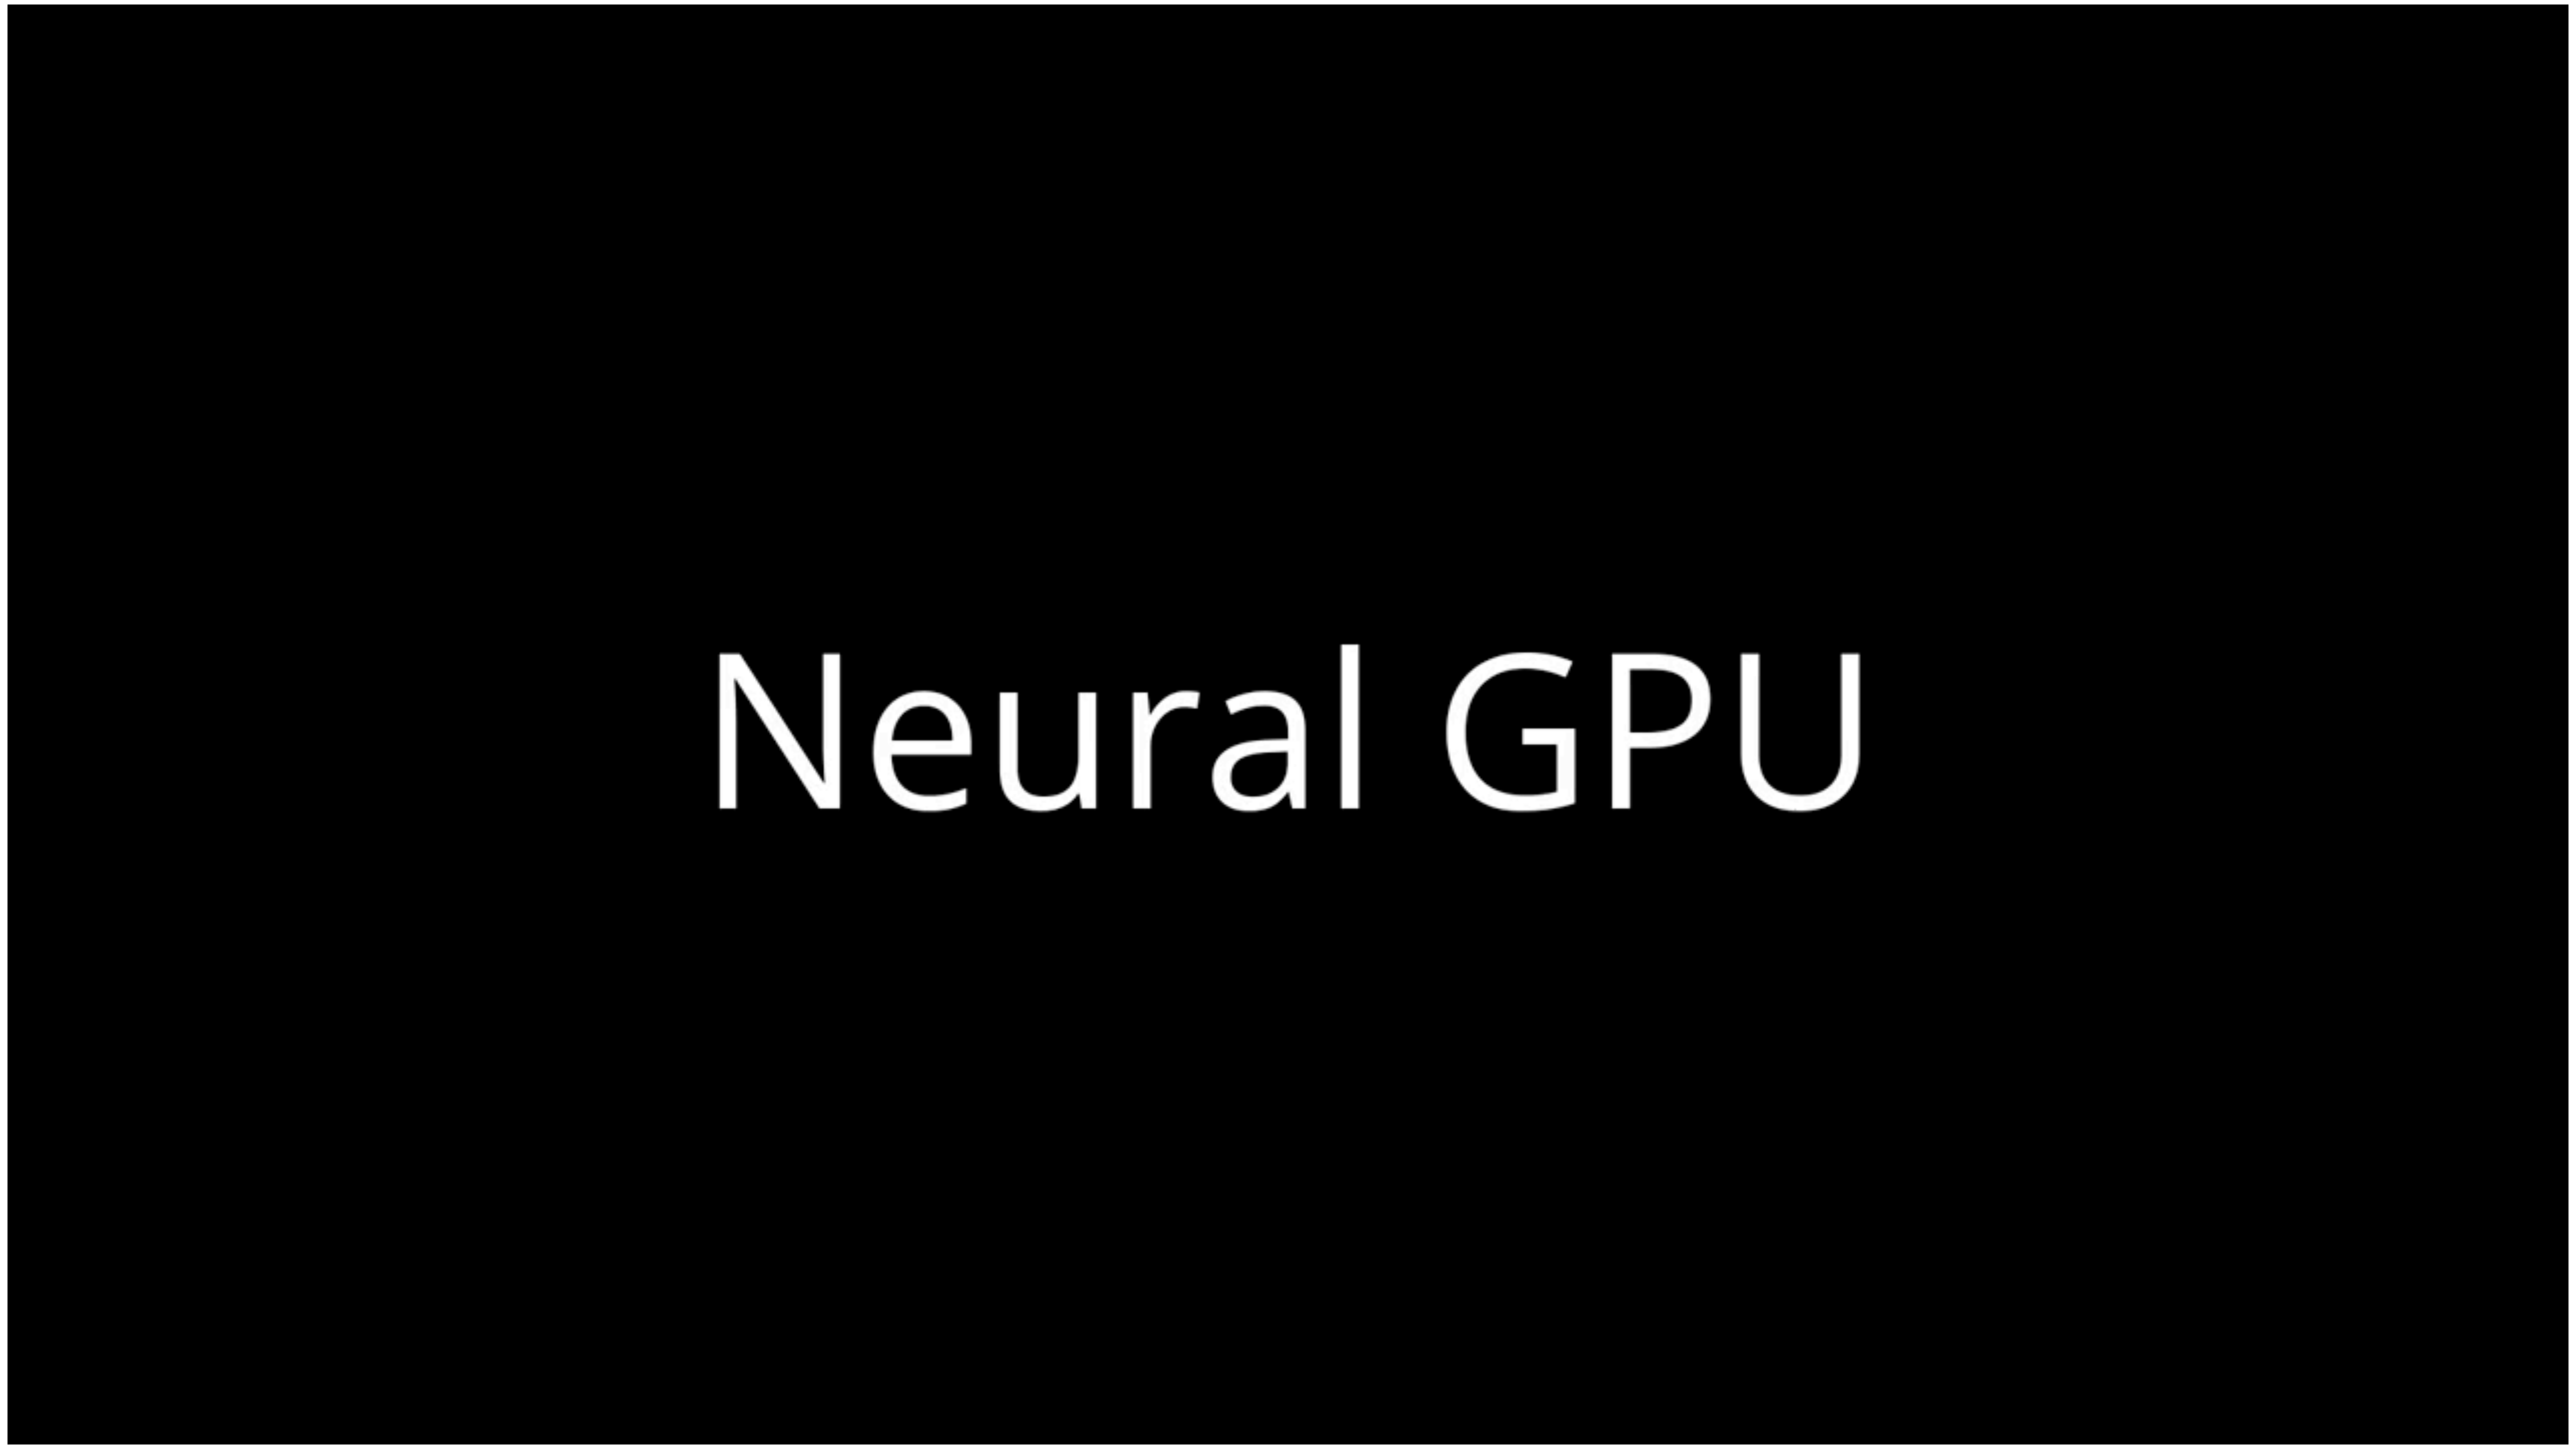

[Kaiser & Sutskever 2015]

Discussion

Empirical results

L&B: connectionist models can be made systematic in principle, but are they systematic in practice?

Operationalizing systematicity

jump	$\Rightarrow$	JUMP
jump left	$\Rightarrow$	LTURN JUMP
jump around right	$\Rightarrow$	RTURN JUMP RTURN JUMP RTURN JUMP RTURN JUMP
turn left twice	$\Rightarrow$	LTURN LTURN
jump thrice	$\Rightarrow$	JUMP JUMP JUMP
jump opposite left and walk thrice	$\Rightarrow$	LTURN LTURN JUMP WALK WALK WALK
jump opposite left after walk around left	$\Rightarrow$	LTURN WALK LTURN WALK LTURN WALK LTURN WALK LTURN LTURN JUMP

Operationalizing systematicity

~~jump~~	~~⇒~~	~~JUMP~~
~~jump left~~	~~⇒~~	~~LTURN JUMP~~
~~jump around right~~	~~⇒~~	~~RTURN JUMP RTURN JUMP RTURN JUMP RTURN JUMP~~
turn left twice	⇒	LTURN LTURN
~~jump thrice~~	~~⇒~~	~~JUMP JUMP JUMP~~
~~jump opposite left and walk thrice~~	~~⇒~~	~~LTURN LTURN JUMP WALK WALK WALK~~
~~jump opposite left after walk around left~~	~~⇒~~	~~LTURN WALK LTURN WALK LTURN WALK LTURN WALK LTURN LTURN JUMP~~

Empirical results

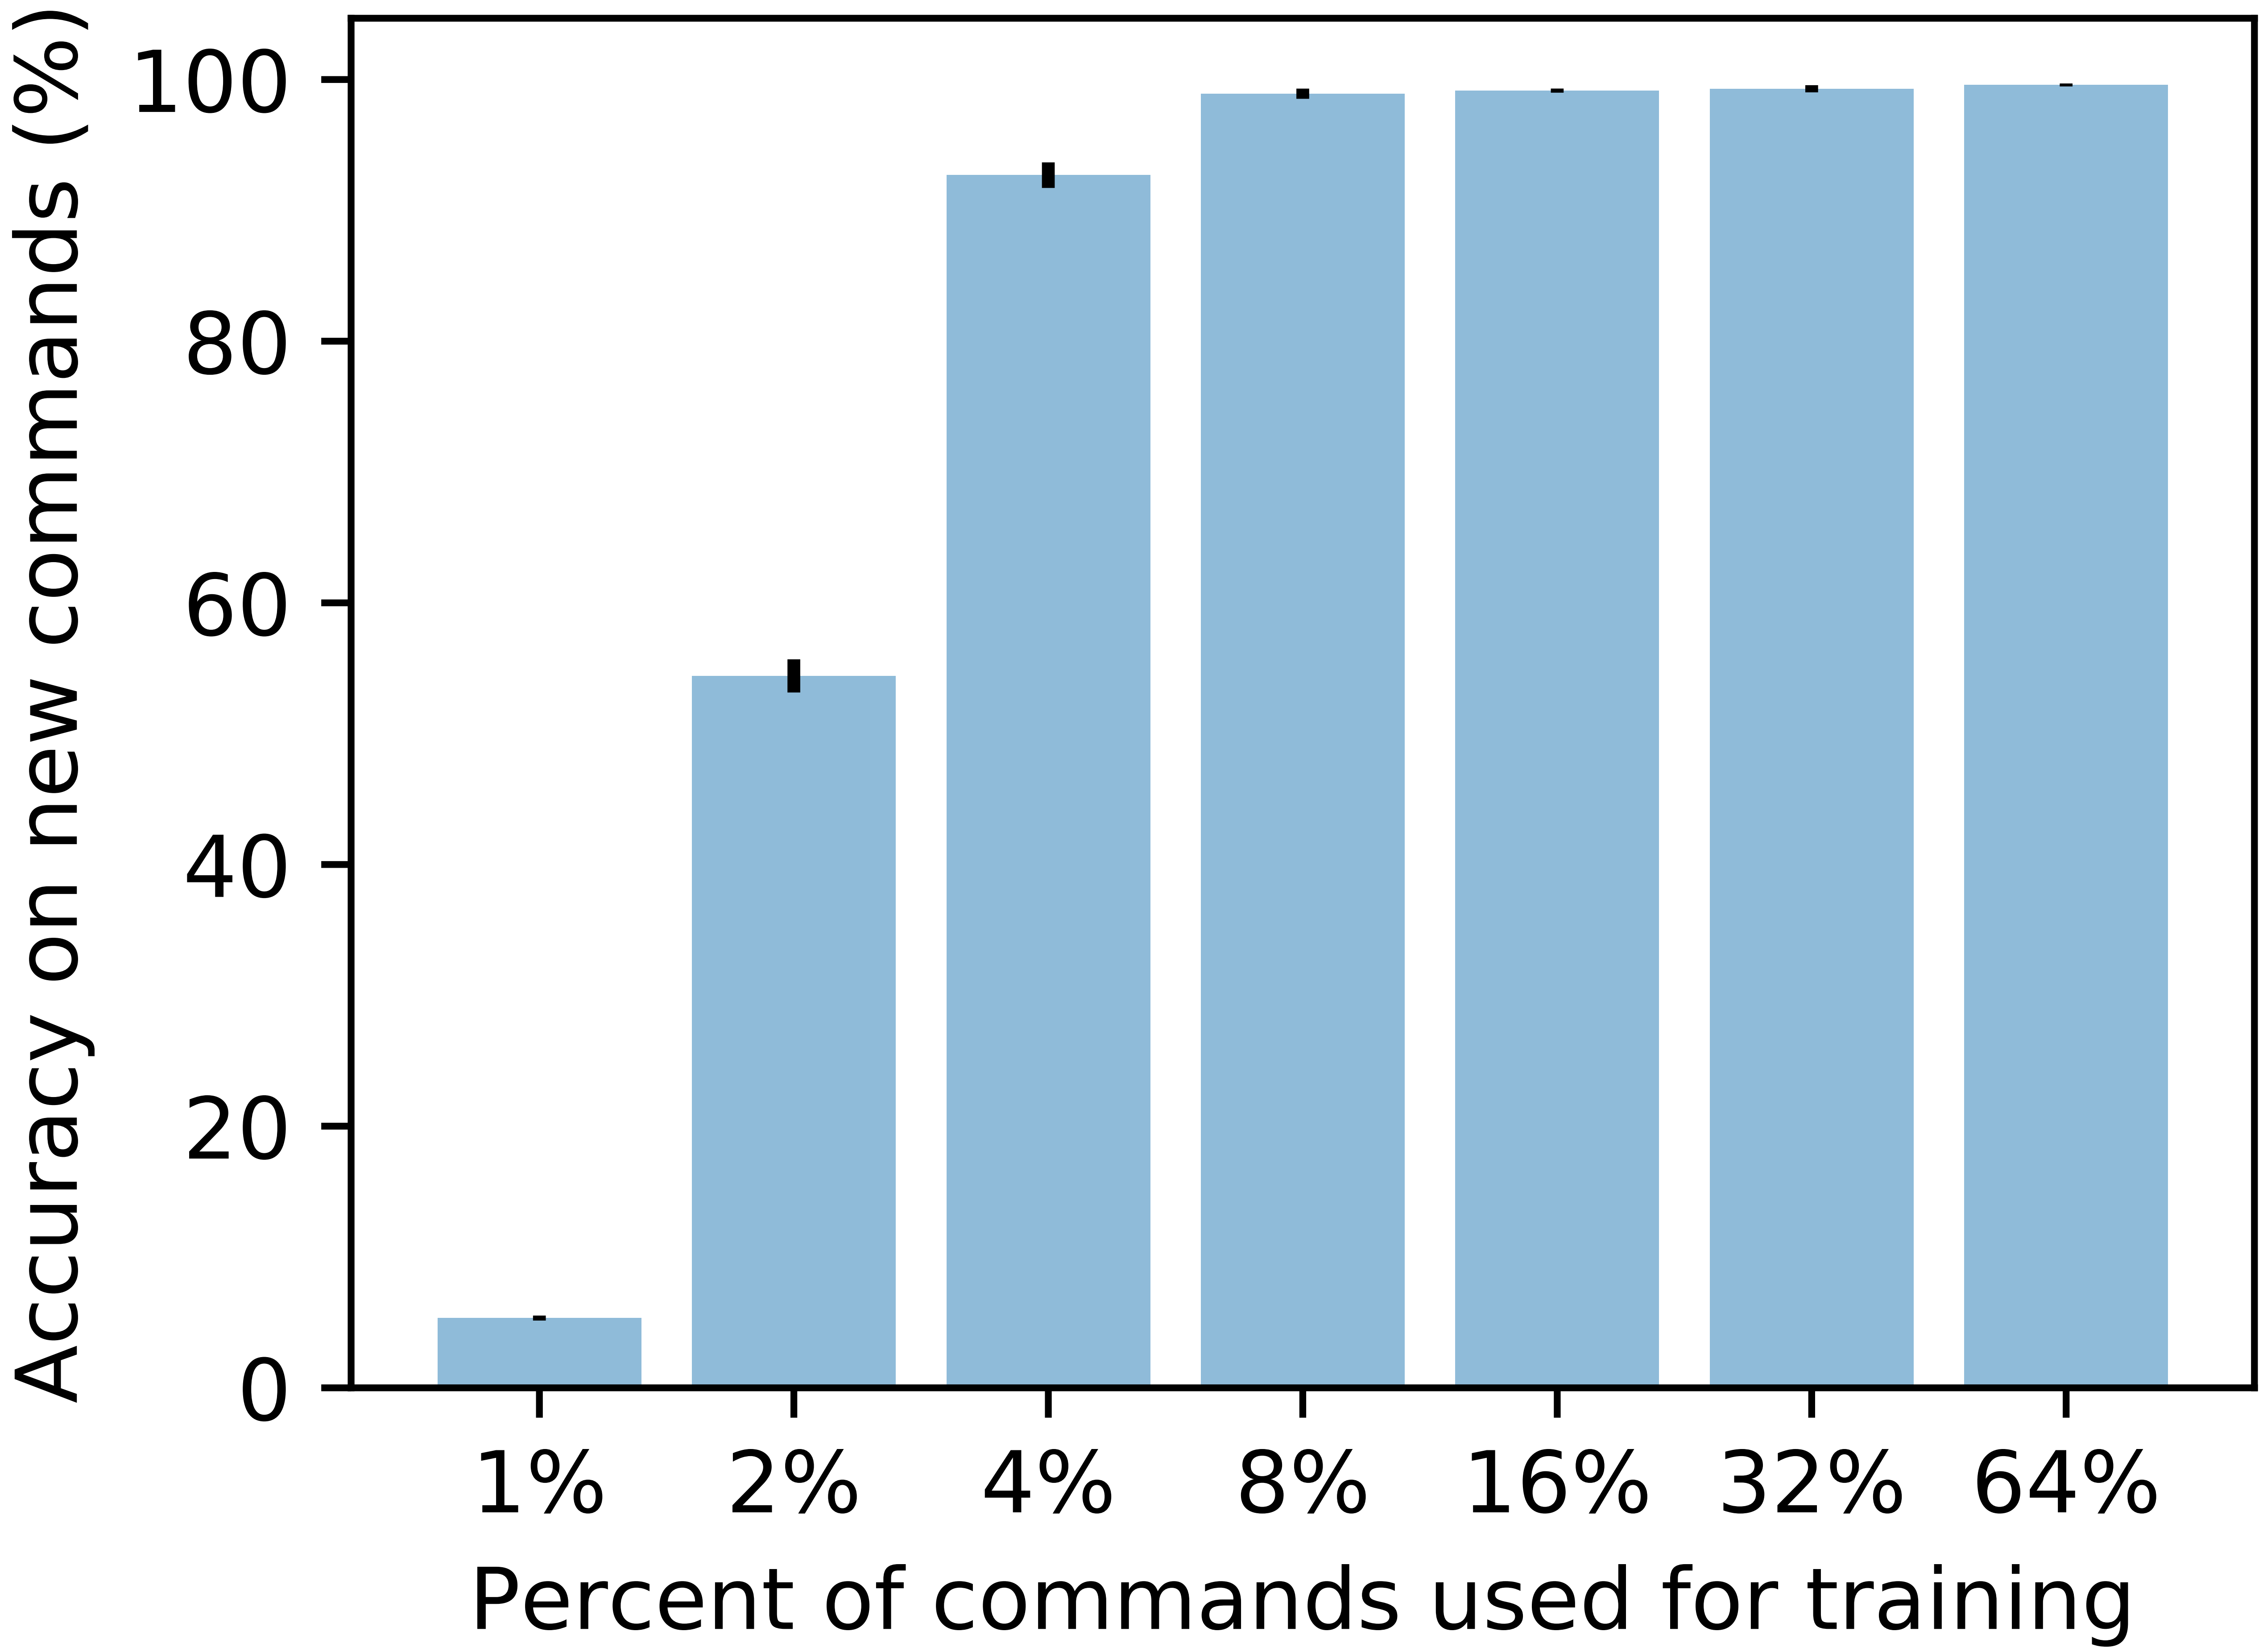

Empirical results

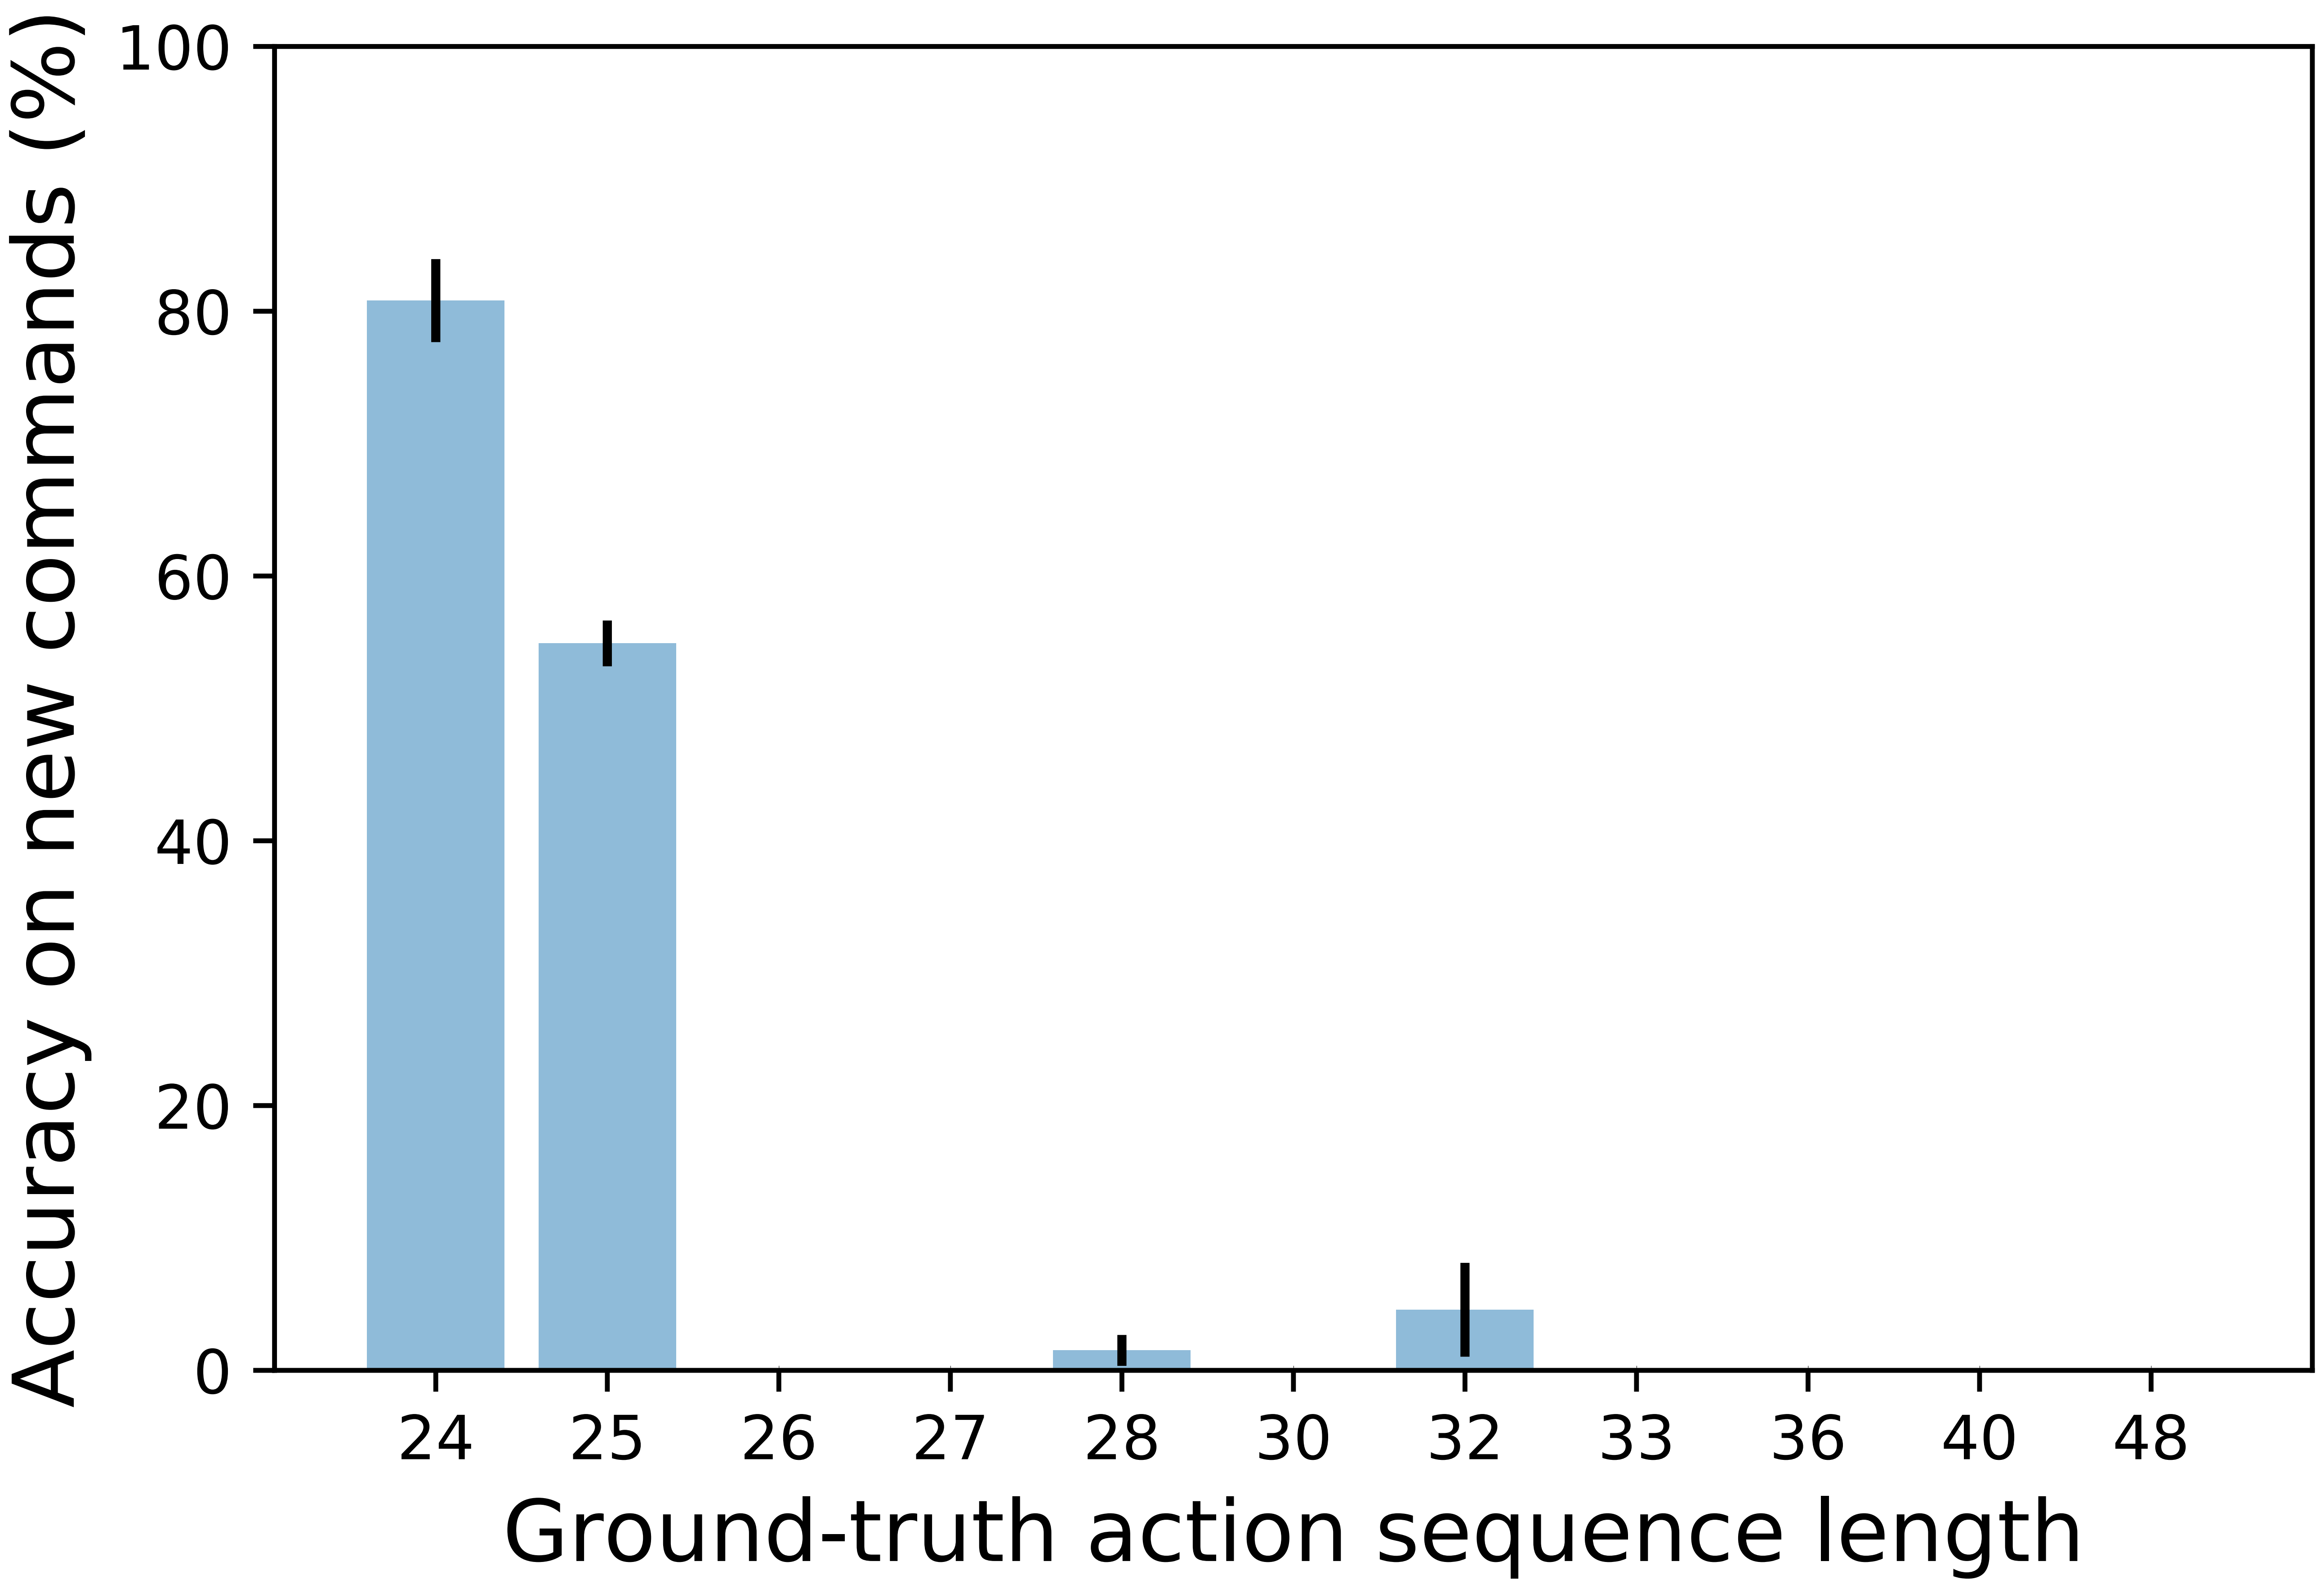

Empirical results

90.3%	1.2%
“turn left”	“jump”

Conclusions

L&B: *Given the astounding successes of seq2seq models in challenging tasks such as machine translation, one might argue that failure to generalize by systematic composition indicates that neural networks are poor models of some aspects of human cognition, but it is of little practical import. However, systematicity is an extremely efficient way to generalize [...] this ability is still beyond the grasp of state-of-the-art neural networks, likely contributing to their striking need for very large training sets. These results give us hope that neural networks capable of systematic compositionality could greatly benefit machine translation, language modeling, and other applications.*

Discussion

See you next week!