The Grammar of Neural Networks

Dieuwke Hupkes

Institute for Logic, Language and Computation
University of Amsterdam

December 7, 2017

Structures in Language

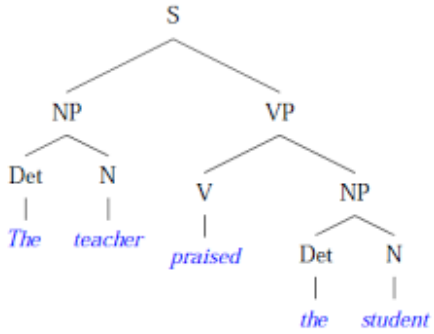

Structures in Language

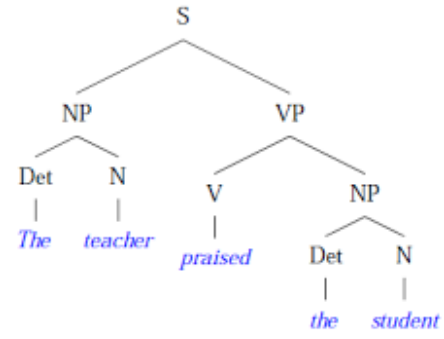

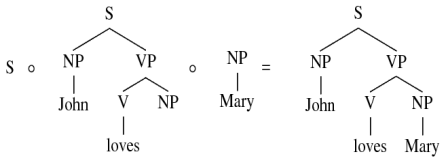

Structures in Language

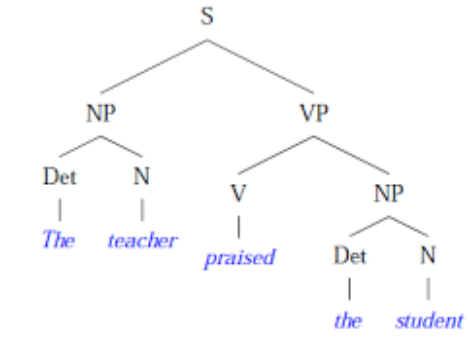

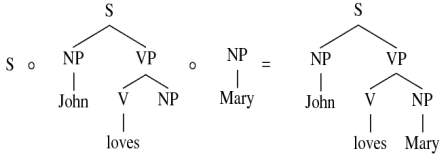

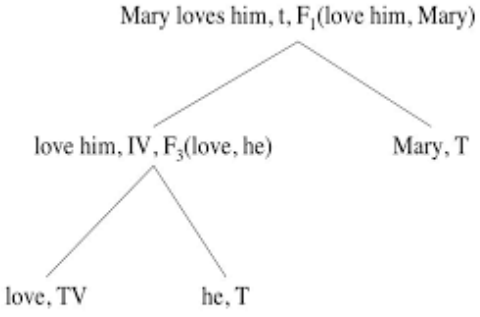

Structures in Language

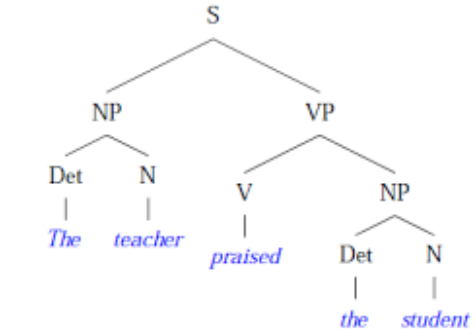

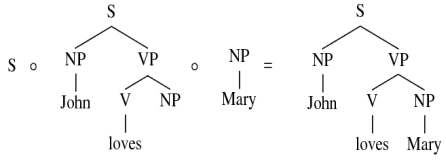

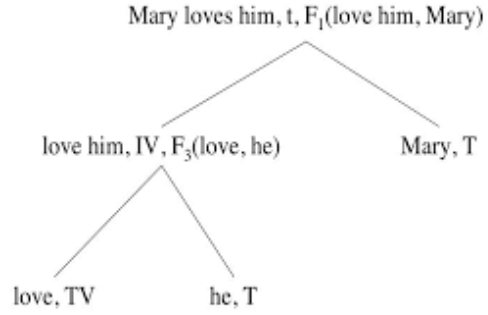

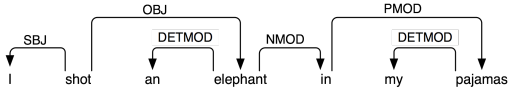

Structures in Language

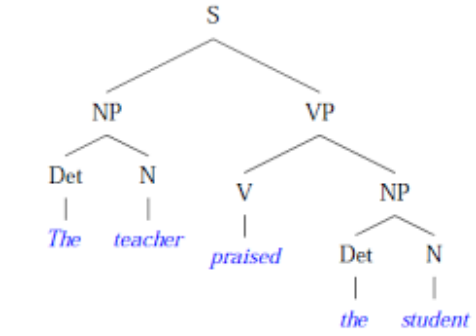

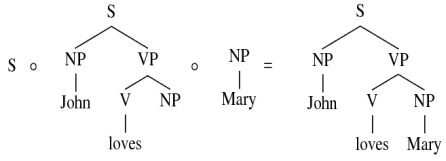

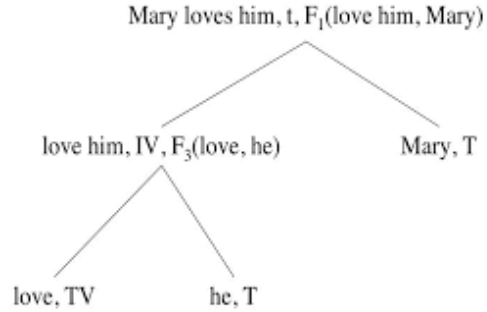

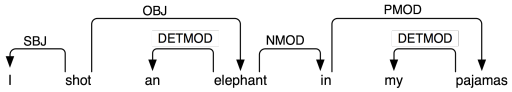

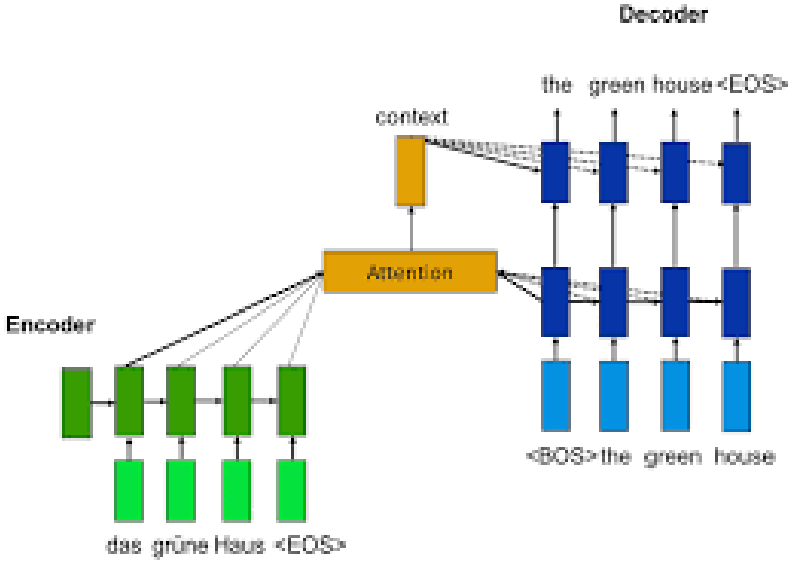

Structures in Language

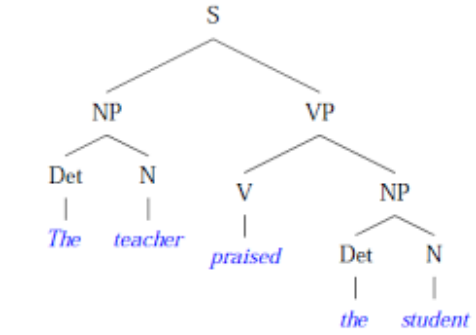

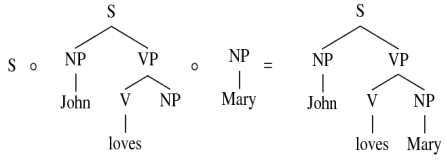

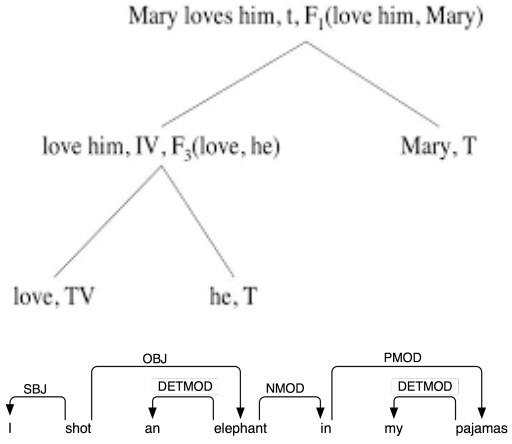

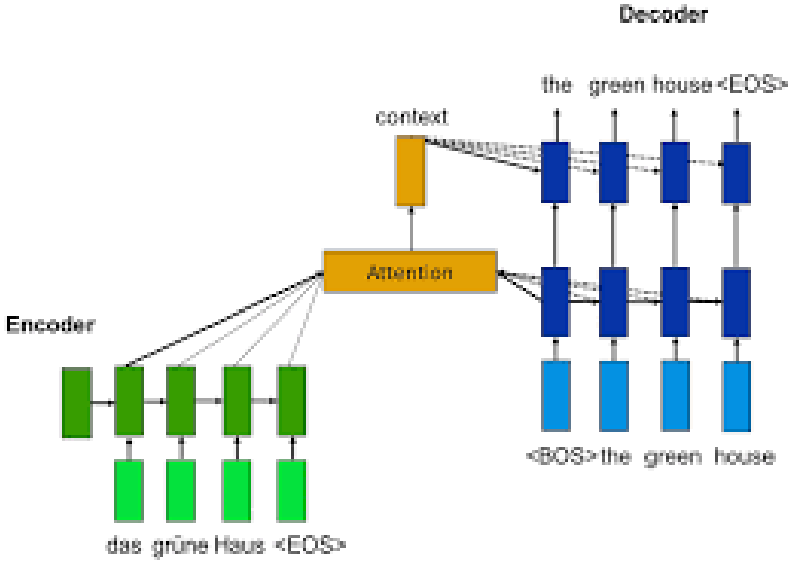

Hierarchical compositionality of language

The scientist who wrote the research paper jumped with joy

Hierarchical compositionality of language

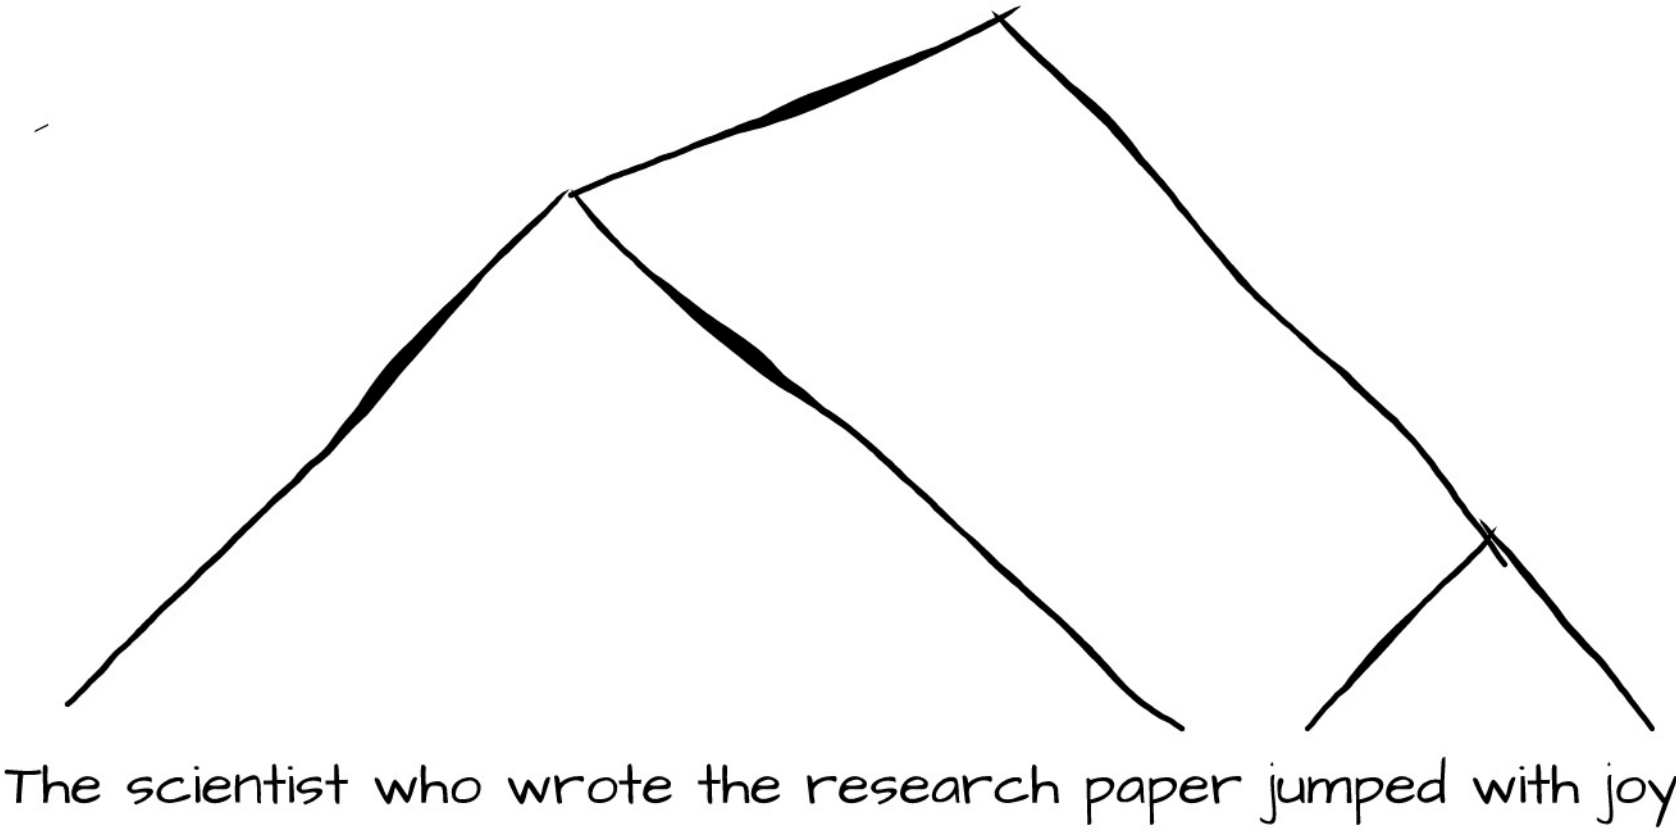

Hierarchical compositionality of language

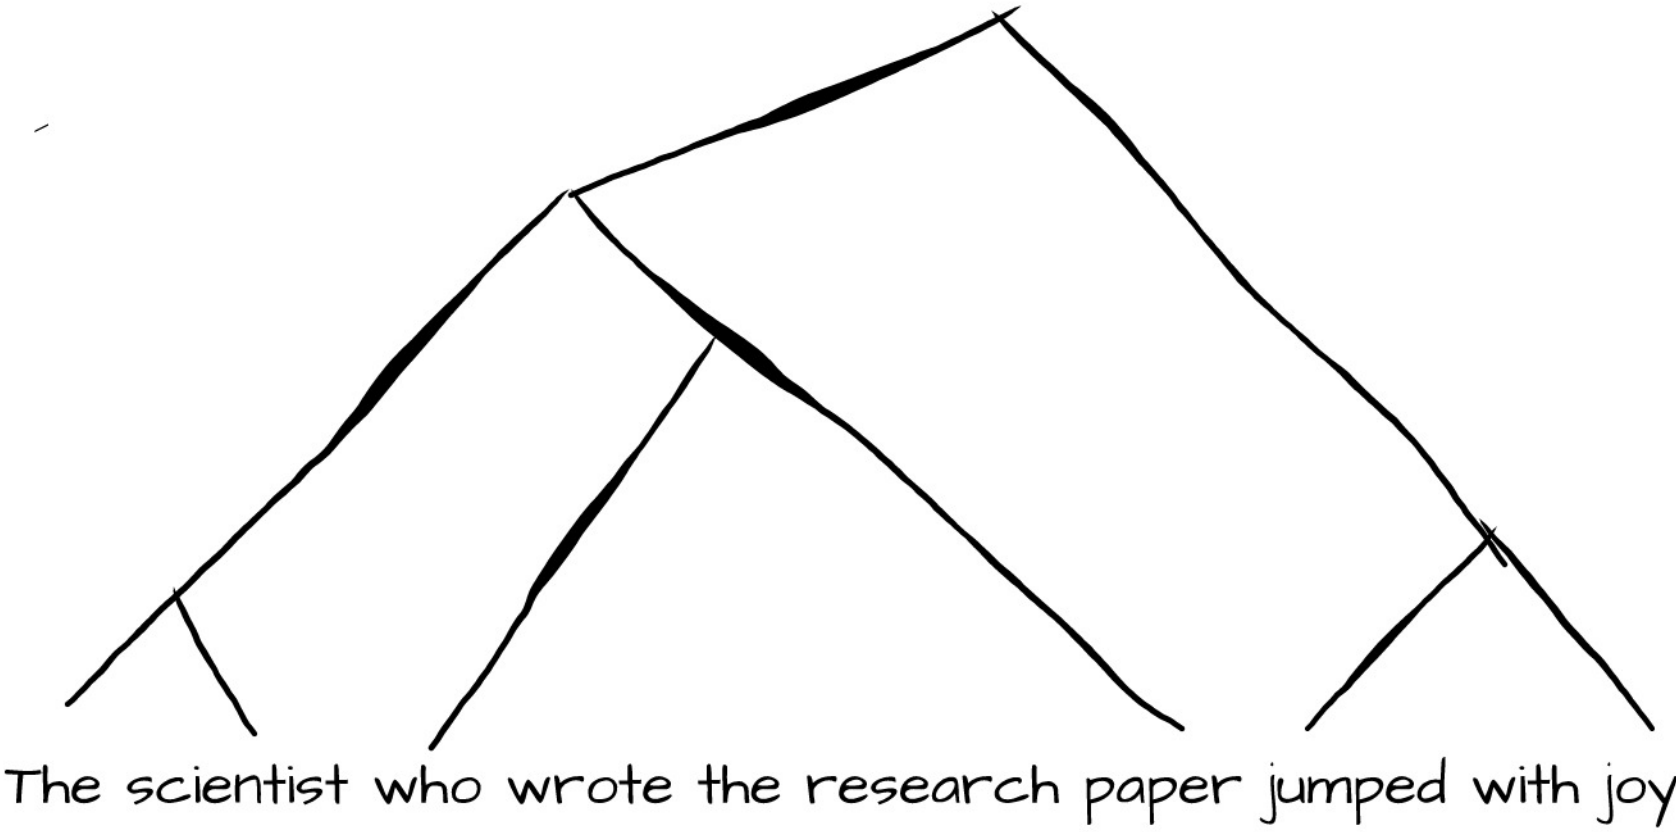

Hierarchical compositionality of language

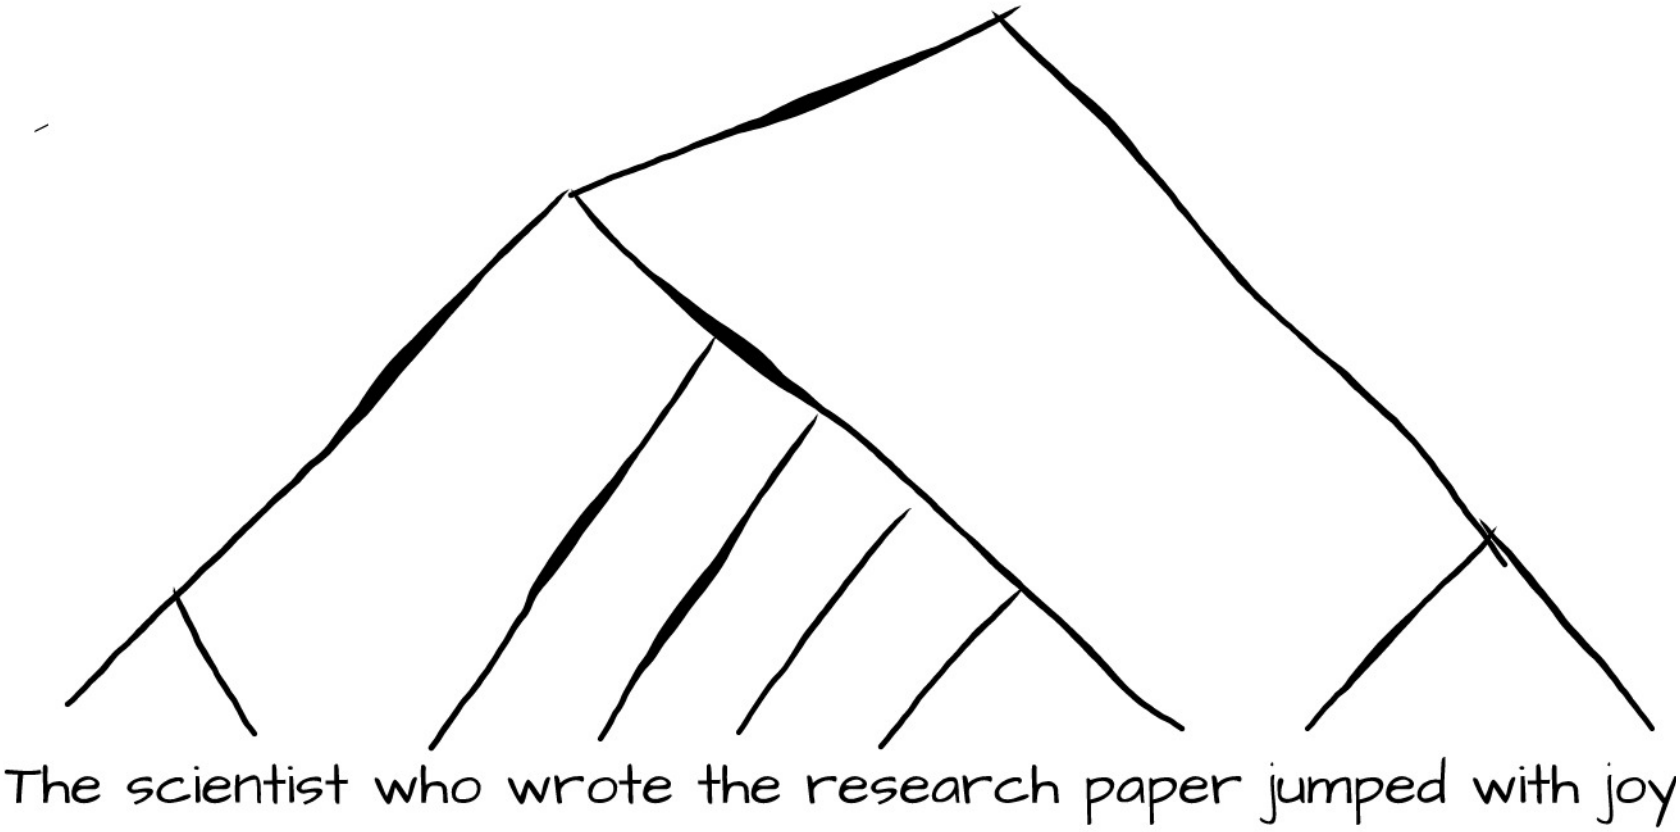

Hierarchical compositionality of language

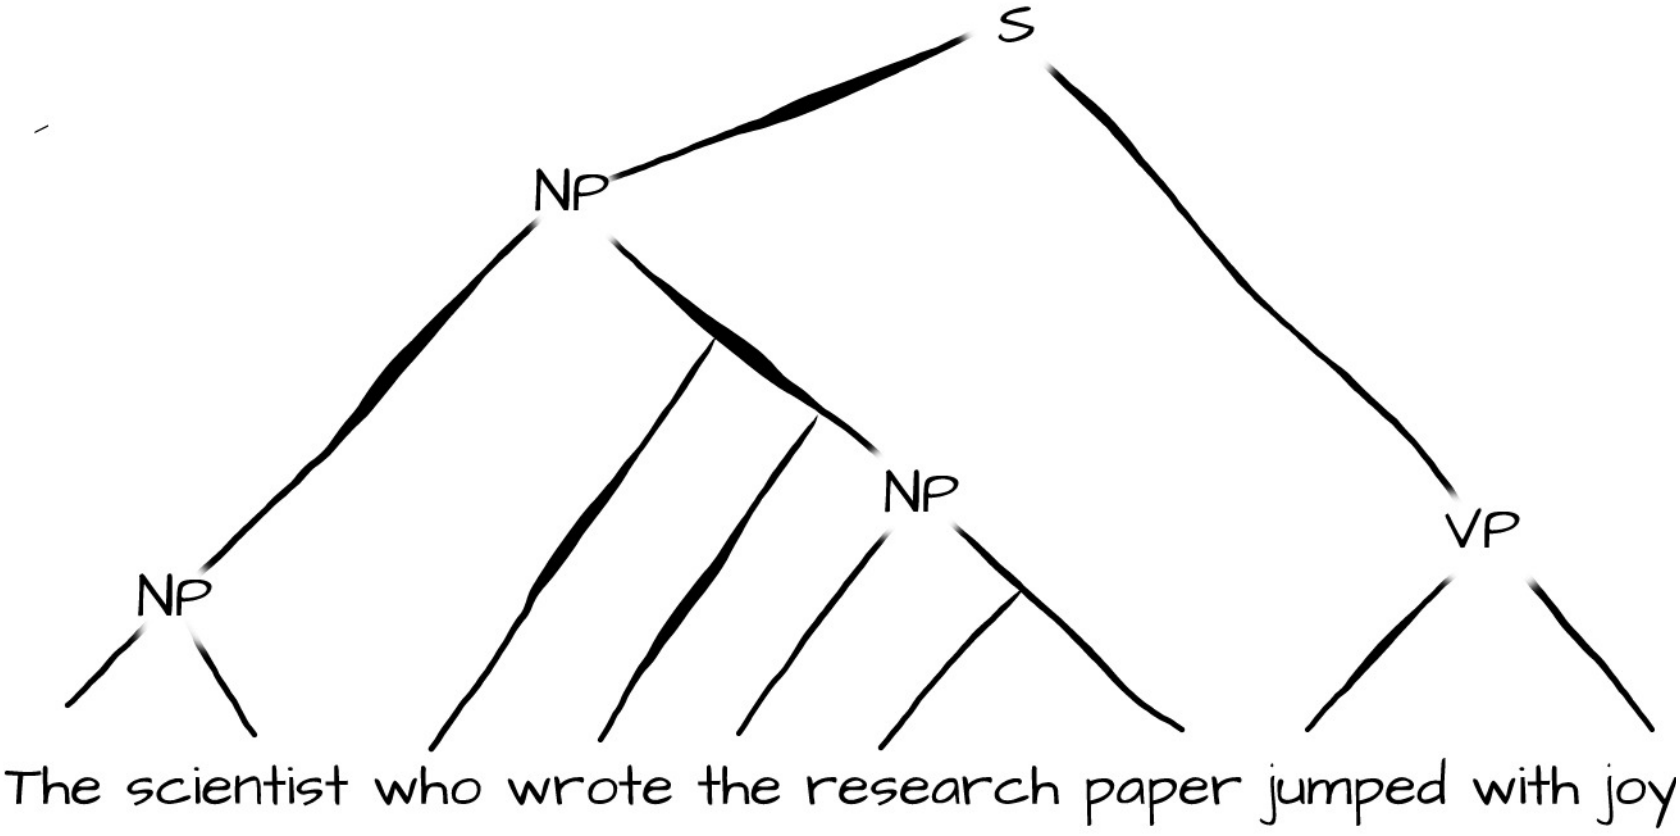

Recurrent Neural Networks

How can hierarchical compositionality be processed incrementally, in linear time, by a recurrent artificial neural network?

Recurrent Neural Networks

How can hierarchical compositionality be processed incrementally, in linear time, by a recurrent artificial neural network?

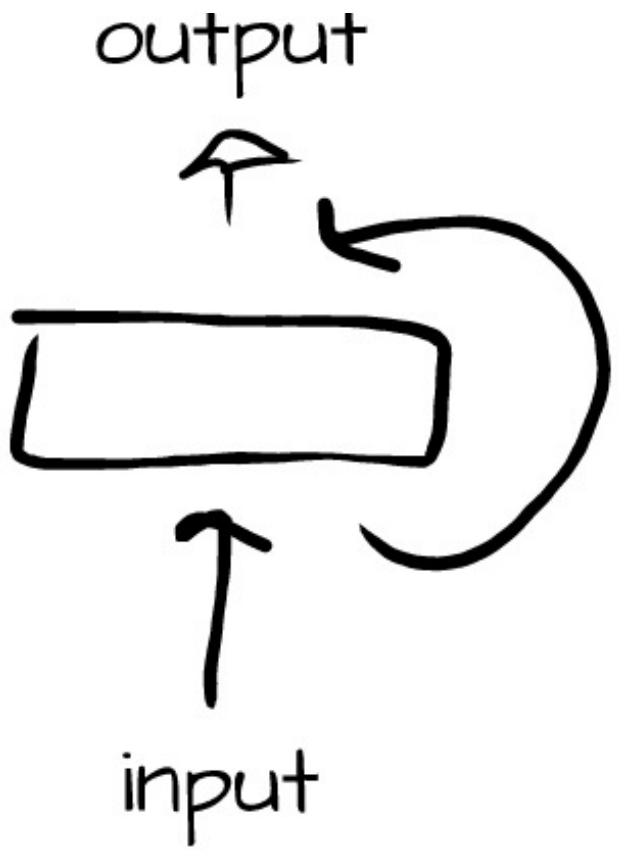

Recurrent Neural Networks

How can hierarchical compositionality be processed incrementally, in linear time, by a recurrent artificial neural network?

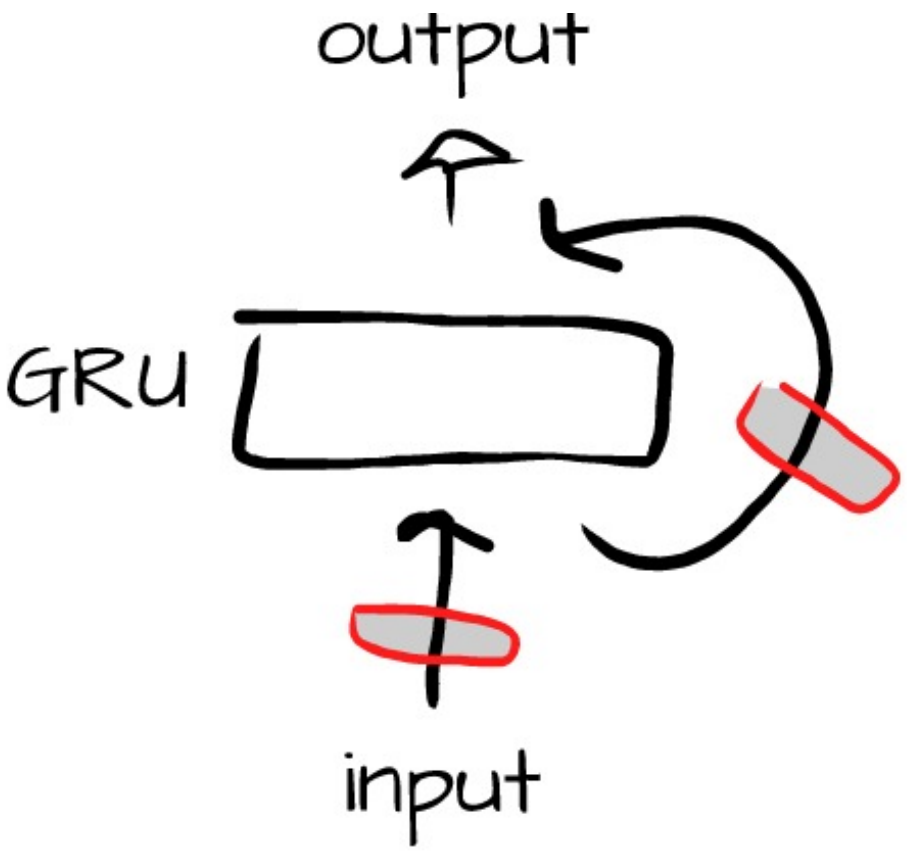

Arithmetic Language

((five minus two) plus six)

(five minus (two plus six))

Arithmetic Language

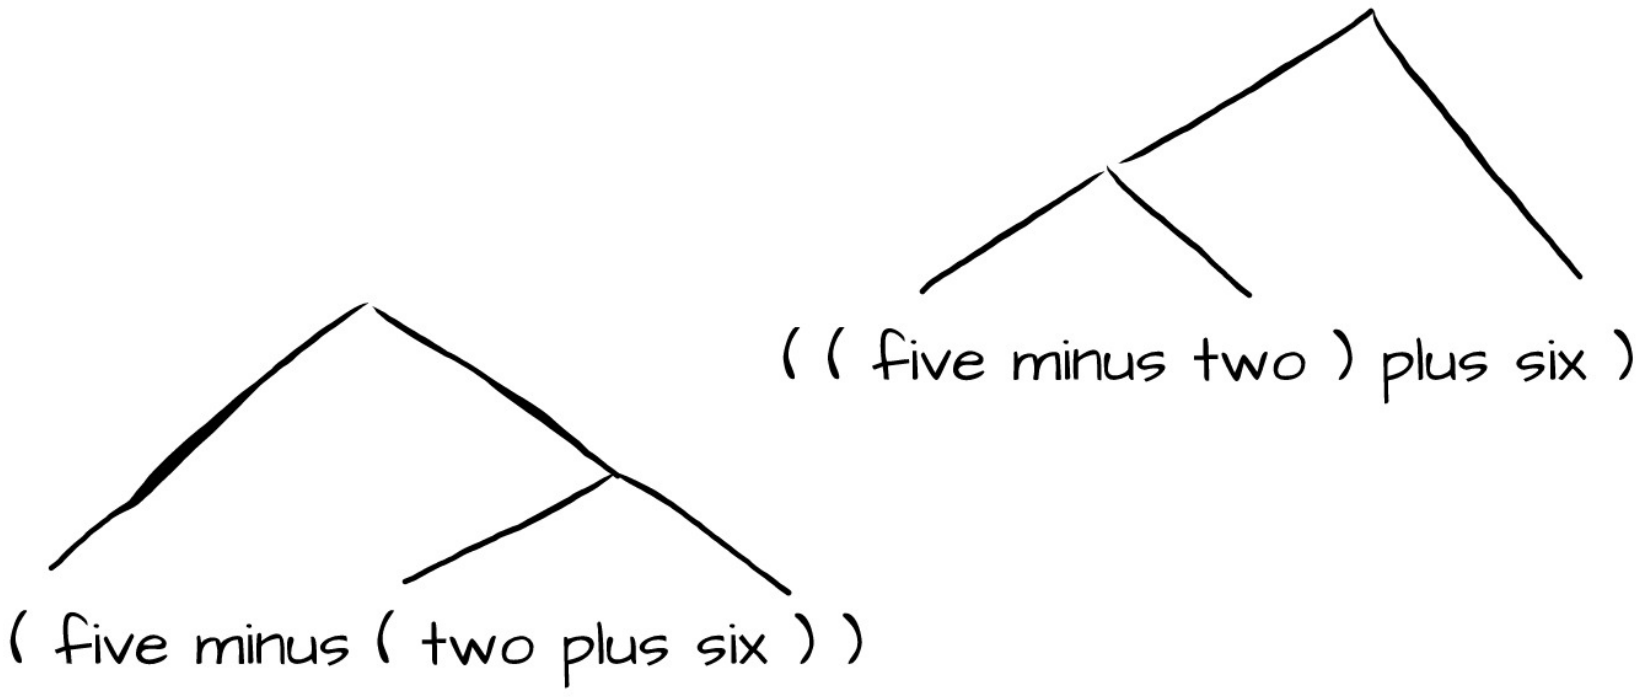

Results

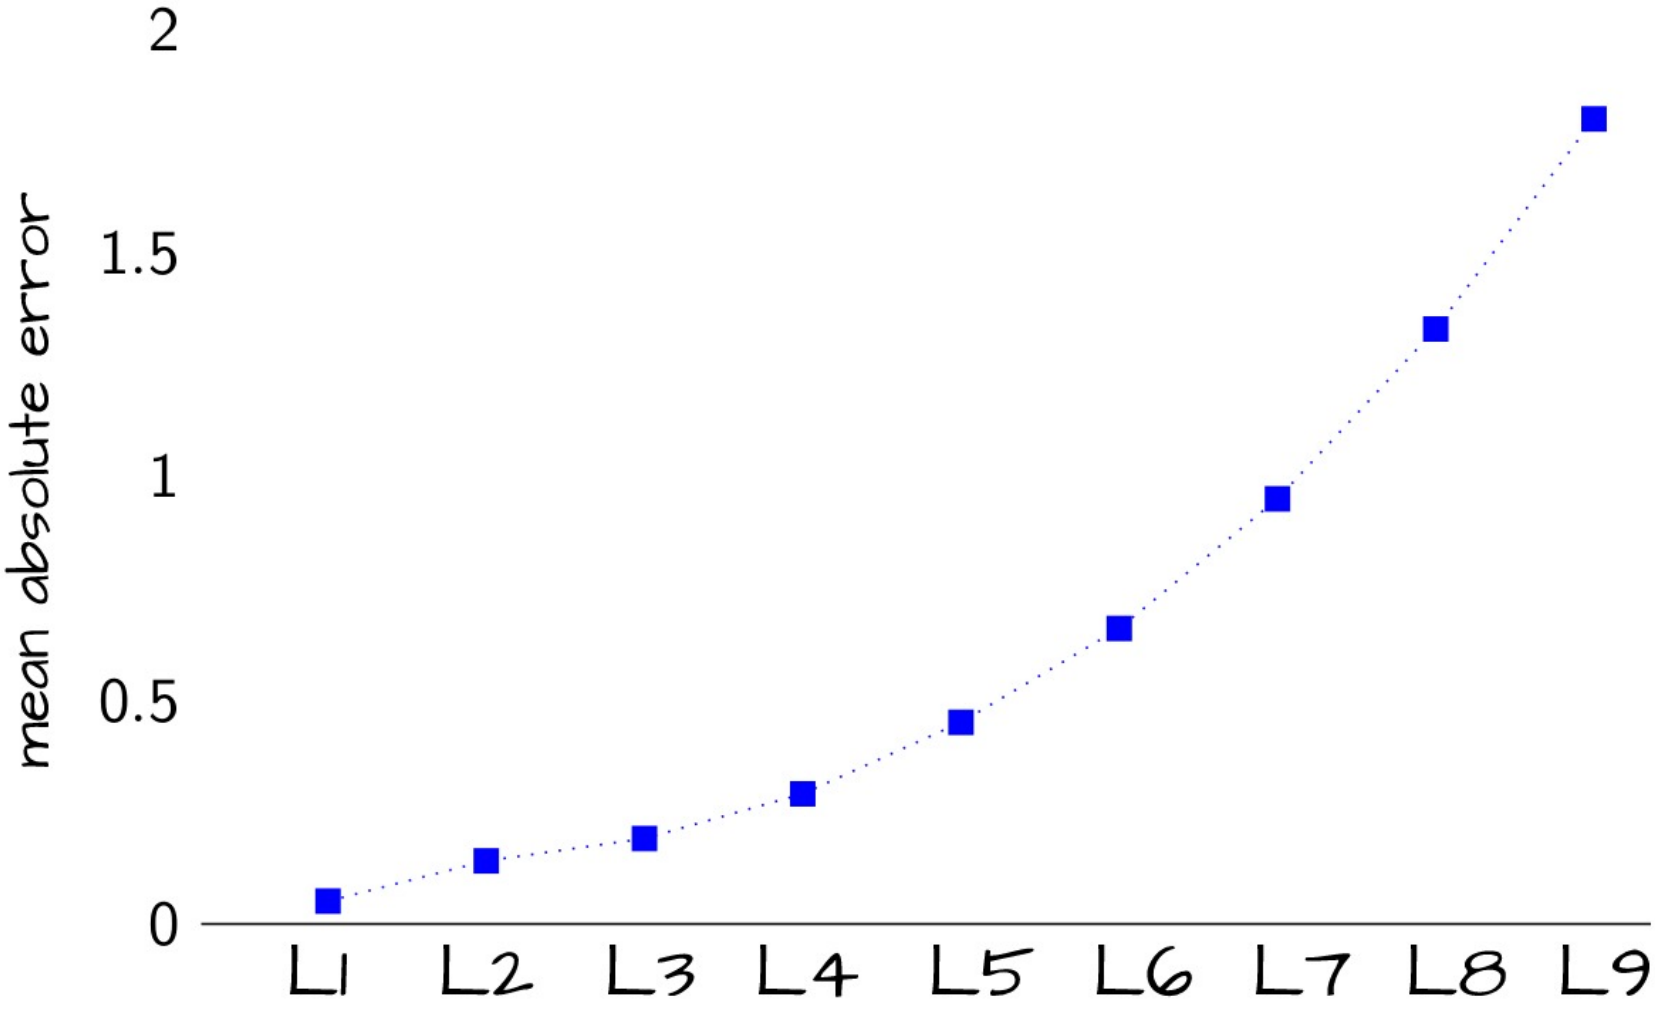

Results

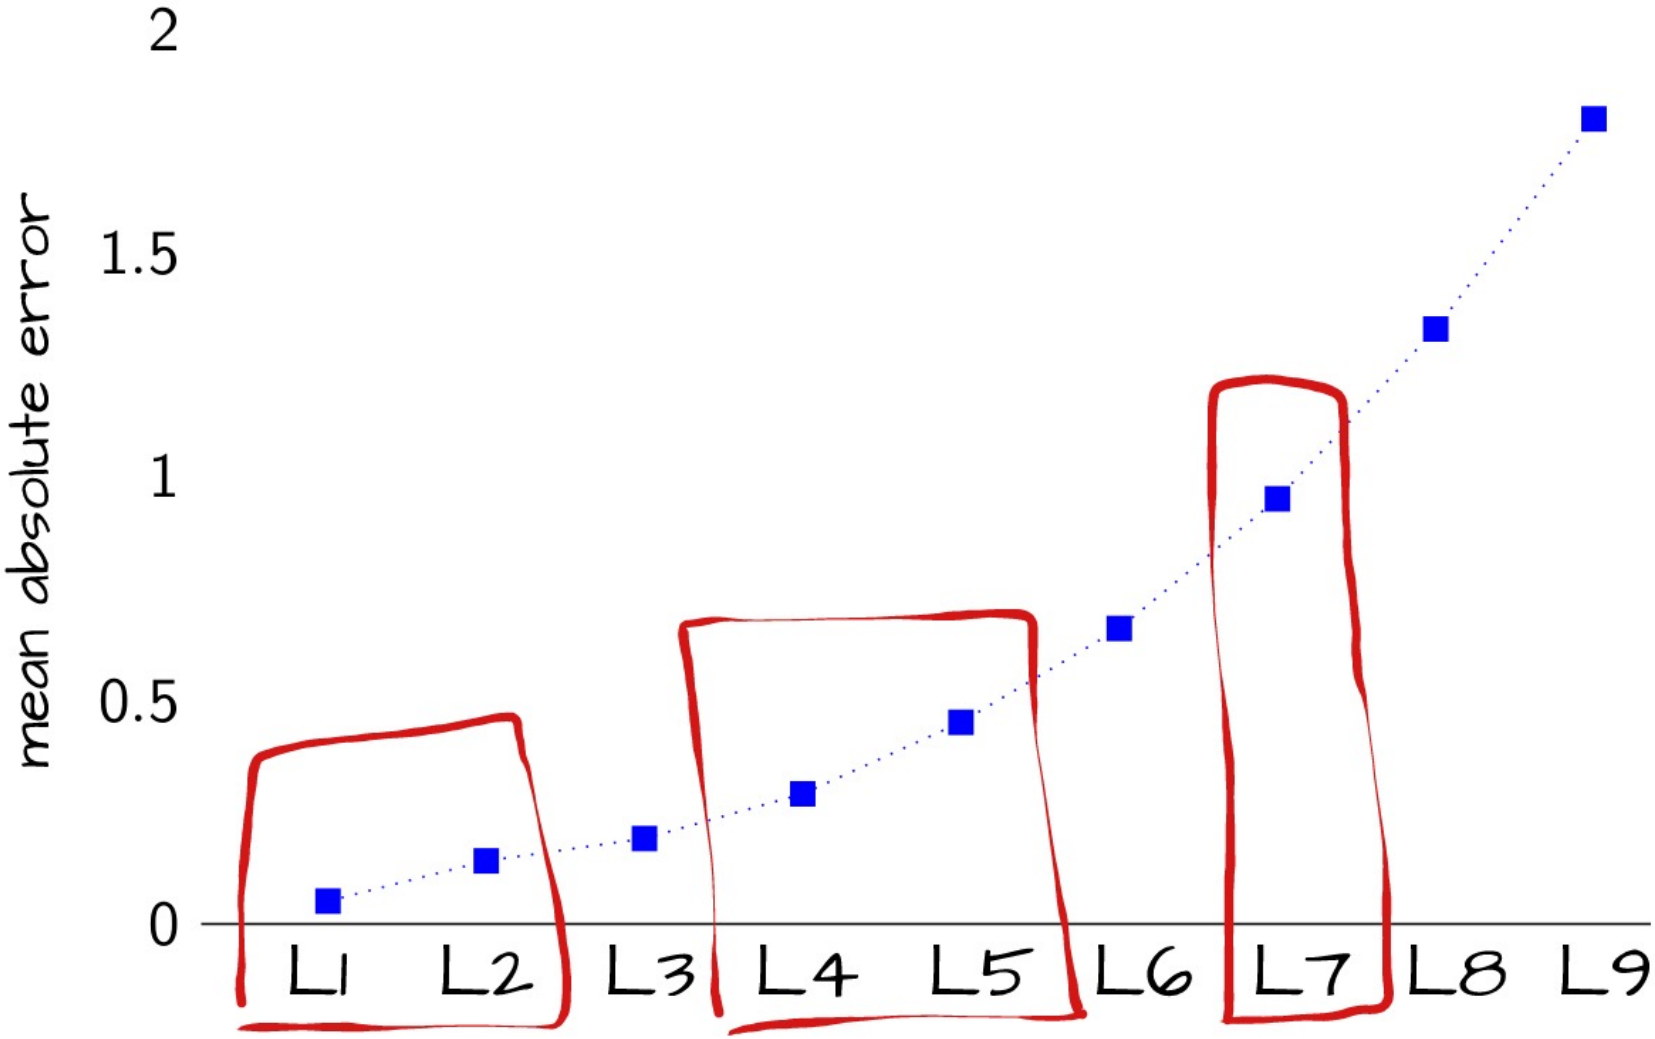

Symbolic solutions

(five minus (two plus six))

Symbolic solutions

recursively

(five minus (two plus six))

Symbolic solutions

recursively 5

(five minus (two plus six))

Symbolic solutions

recursively

5 - 5

(five minus (two plus six))

Symbolic solutions

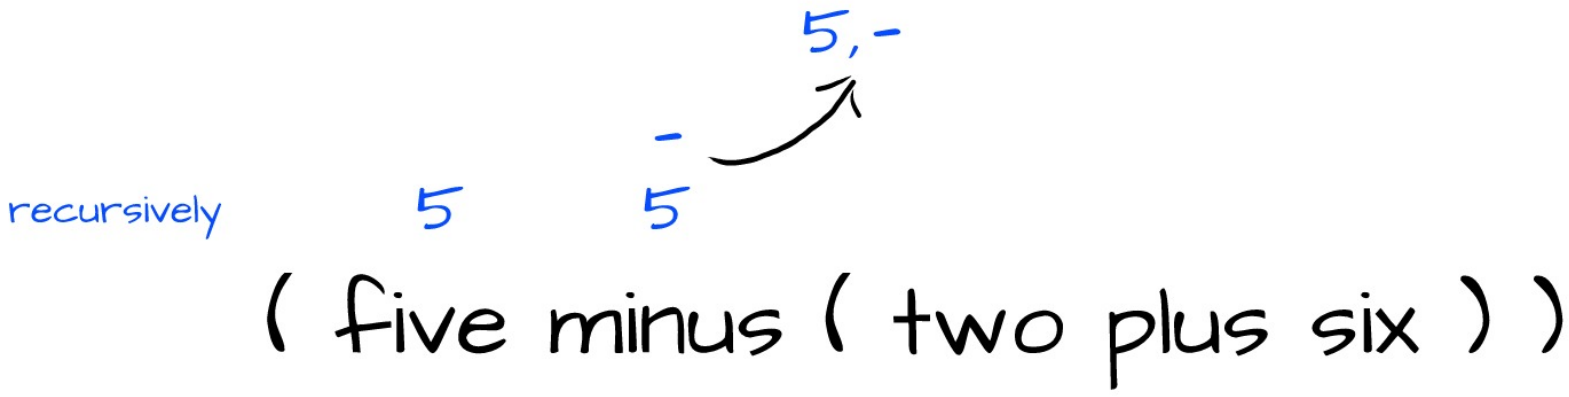

Symbolic solutions

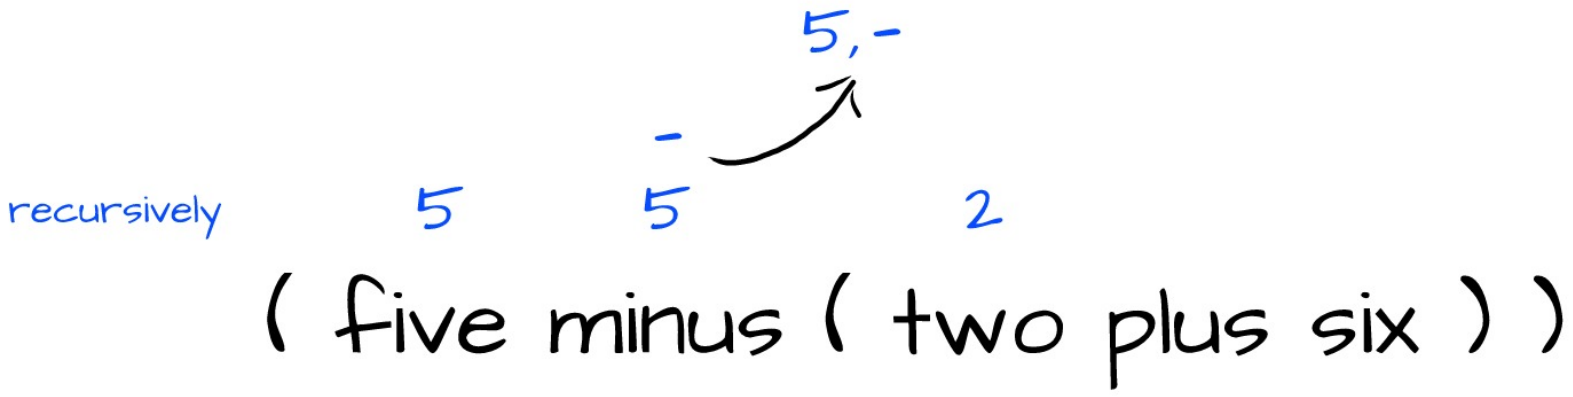

Symbolic solutions

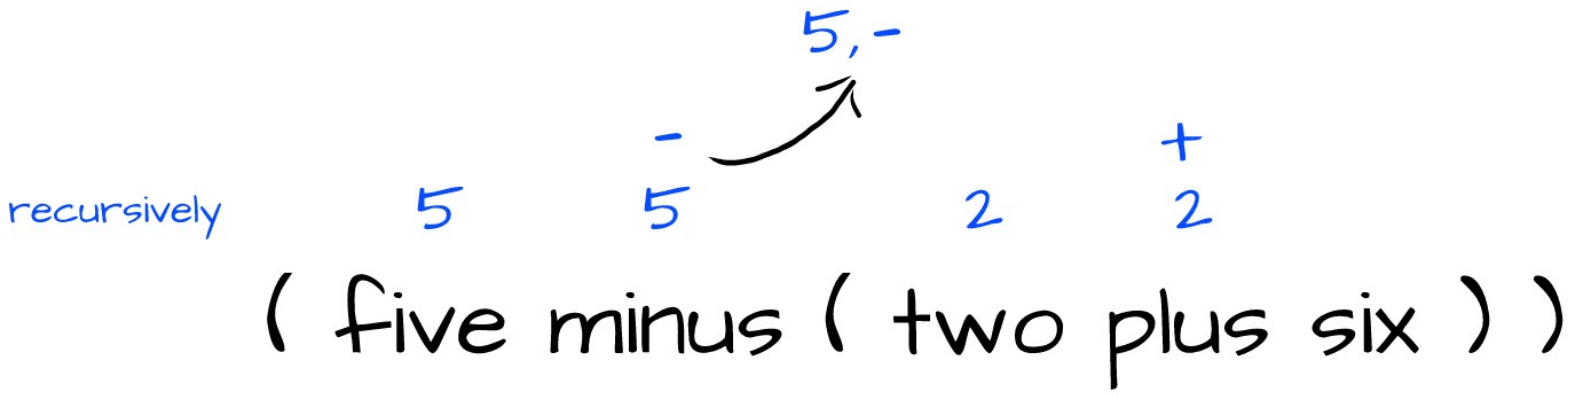

Symbolic solutions

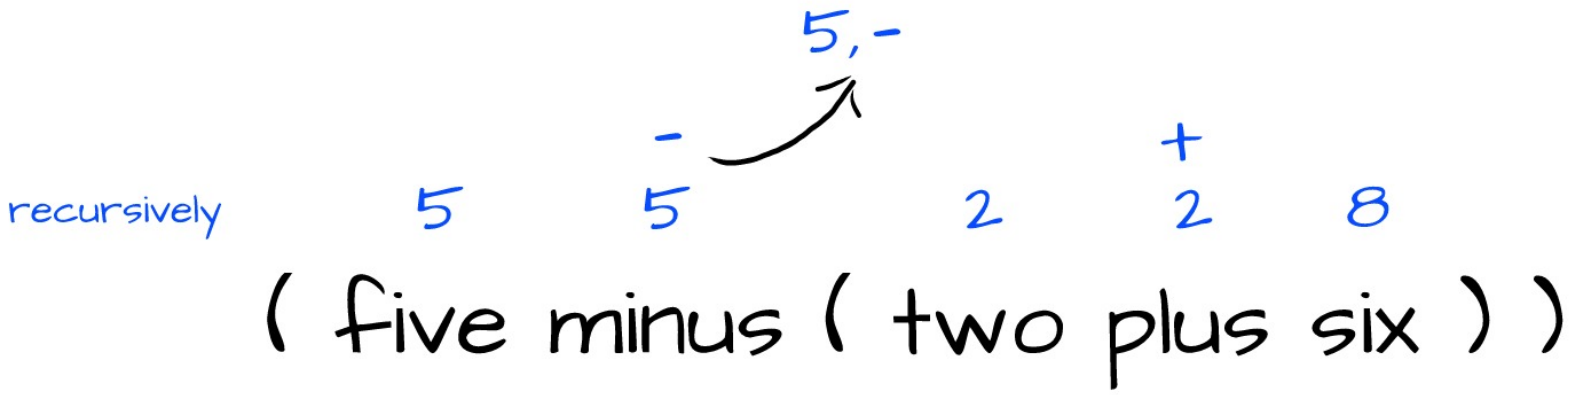

Symbolic solutions

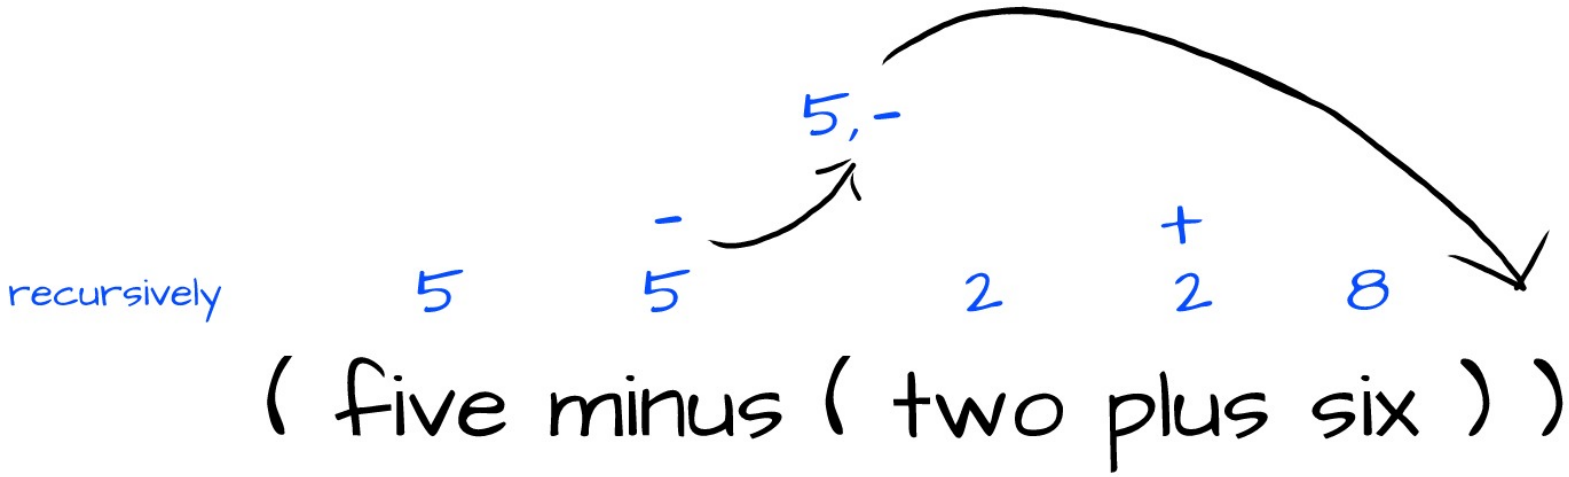

Symbolic solutions

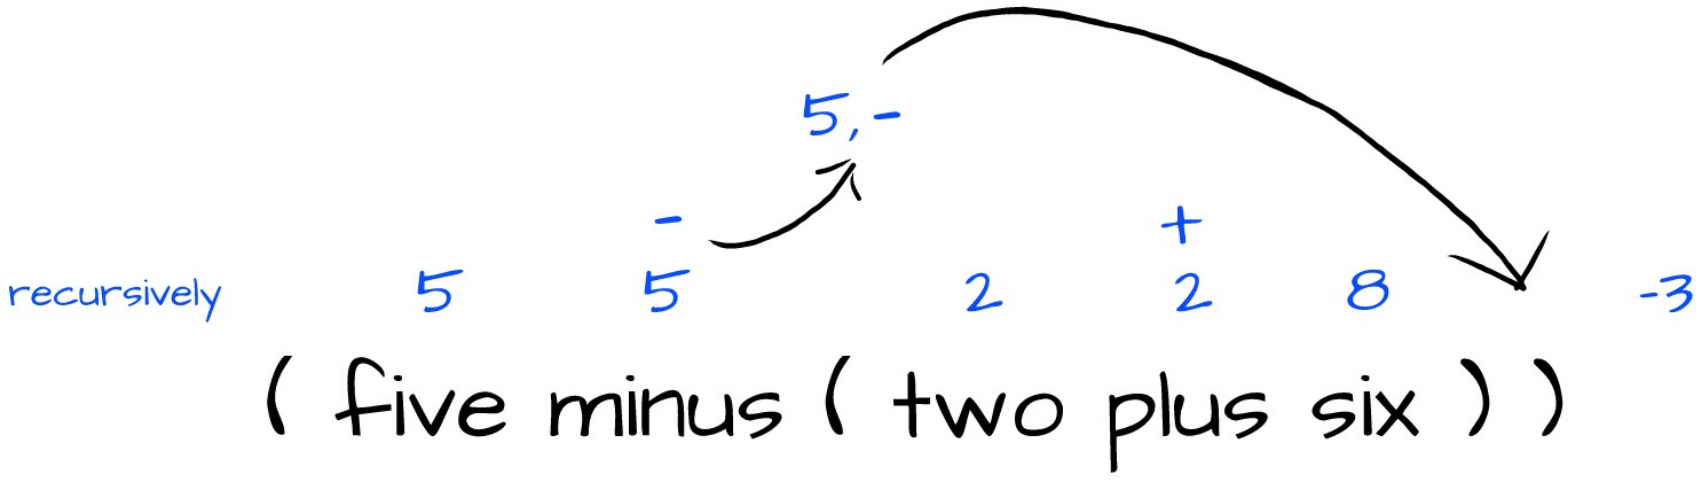

Symbolic solutions

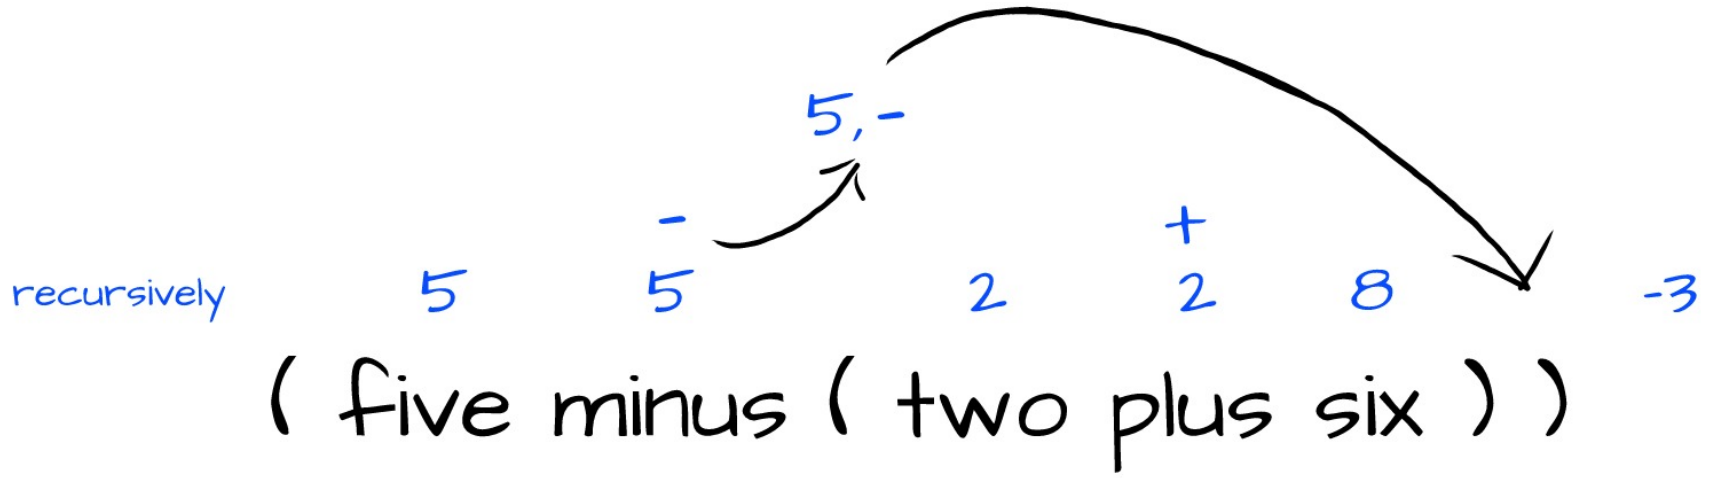

cummulatively

Symbolic solutions

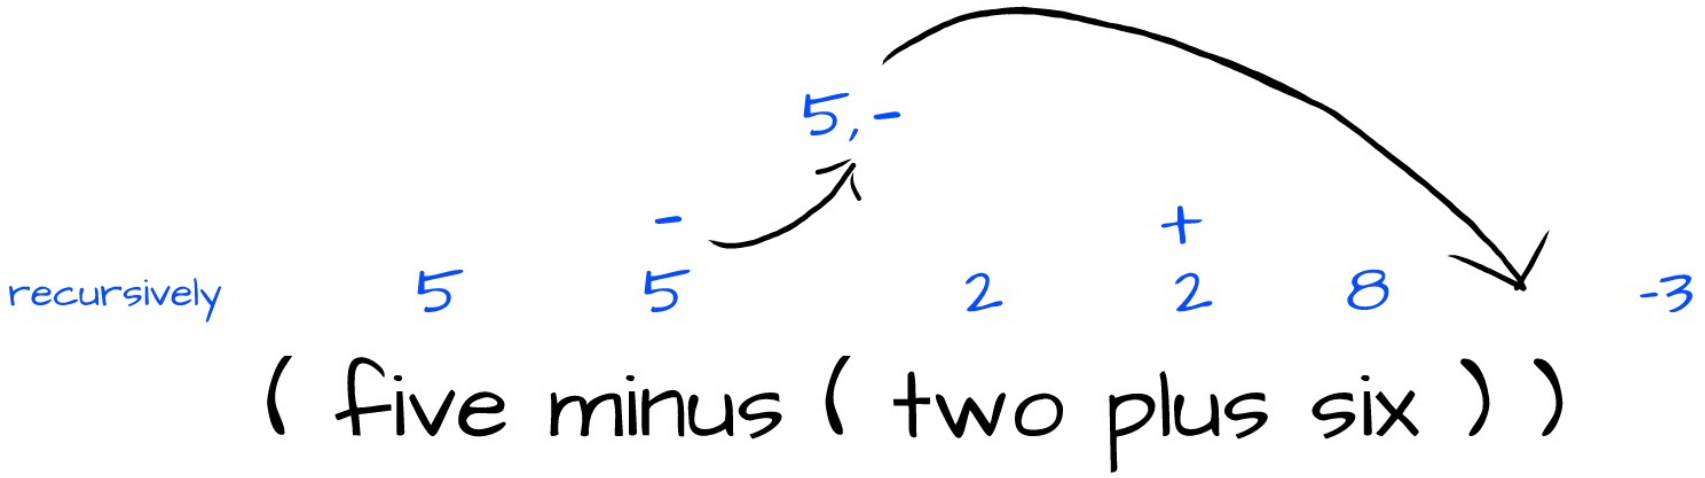

Symbolic solutions

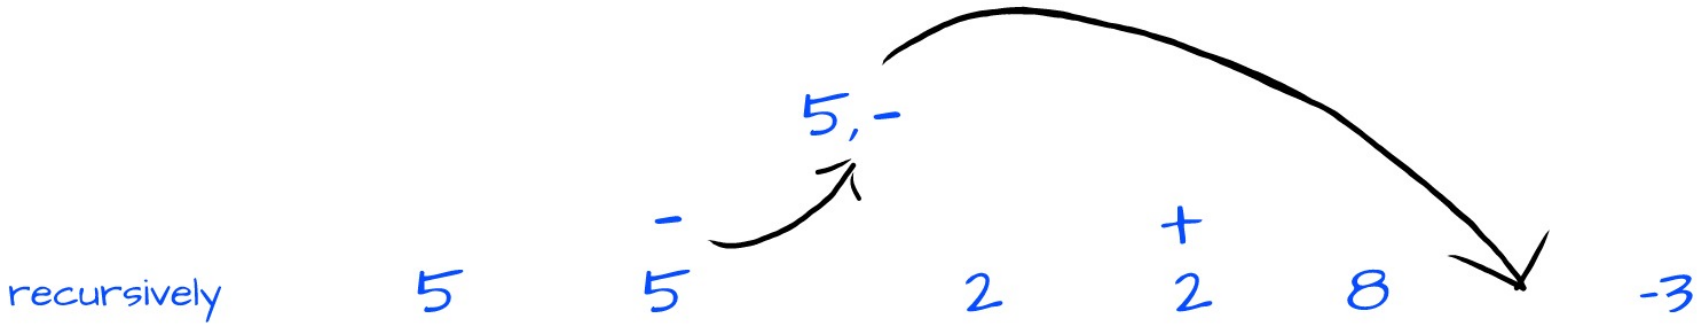

Symbolic solutions

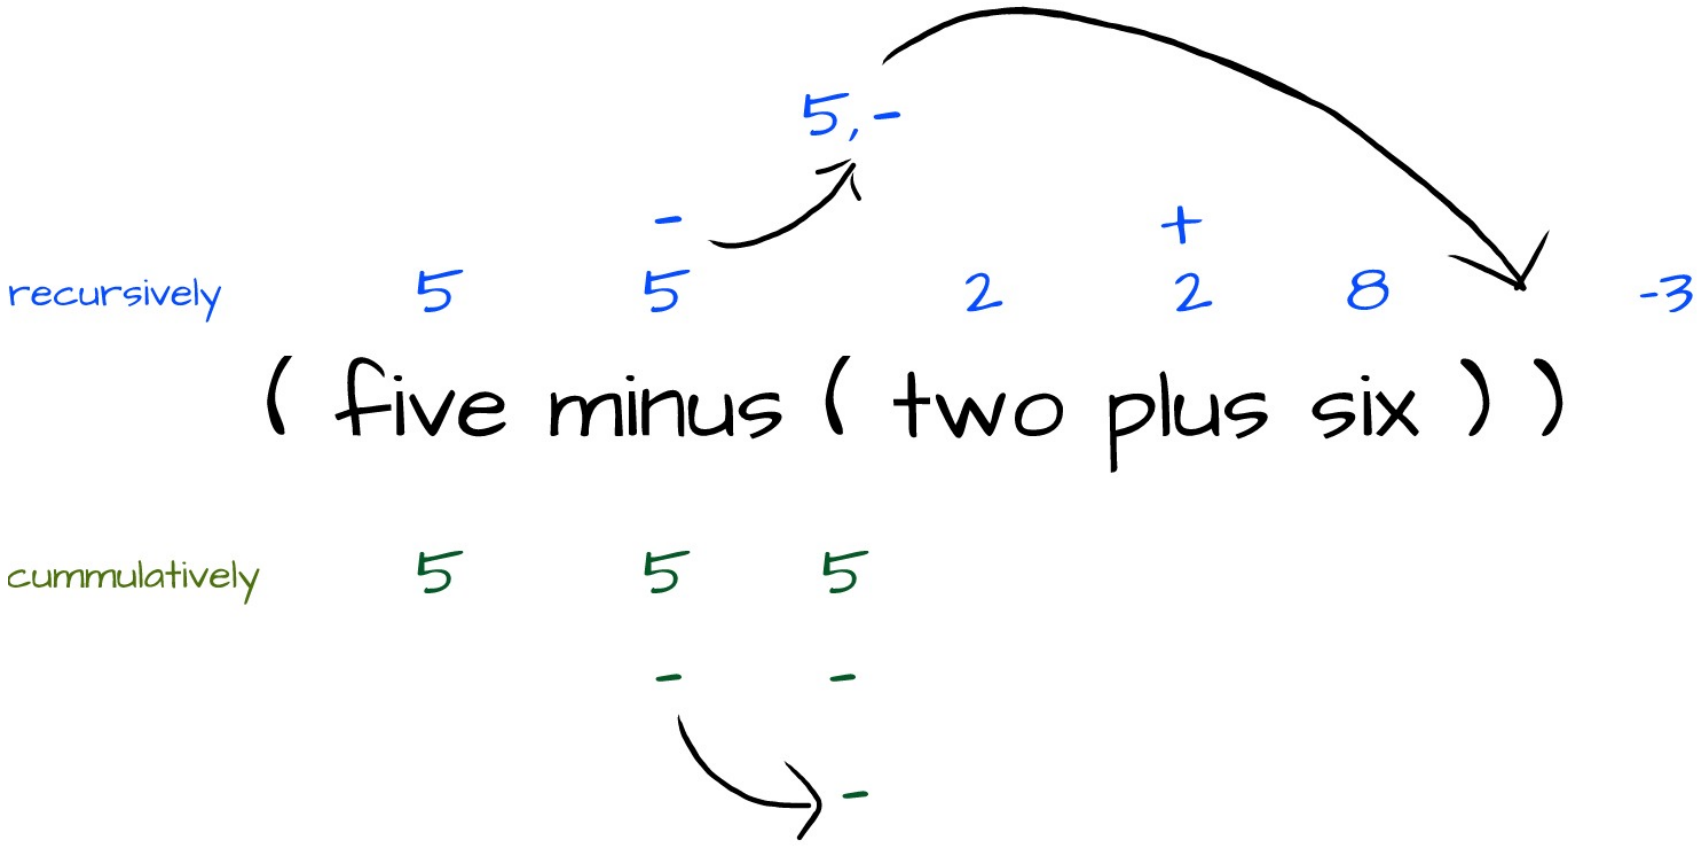

Symbolic solutions

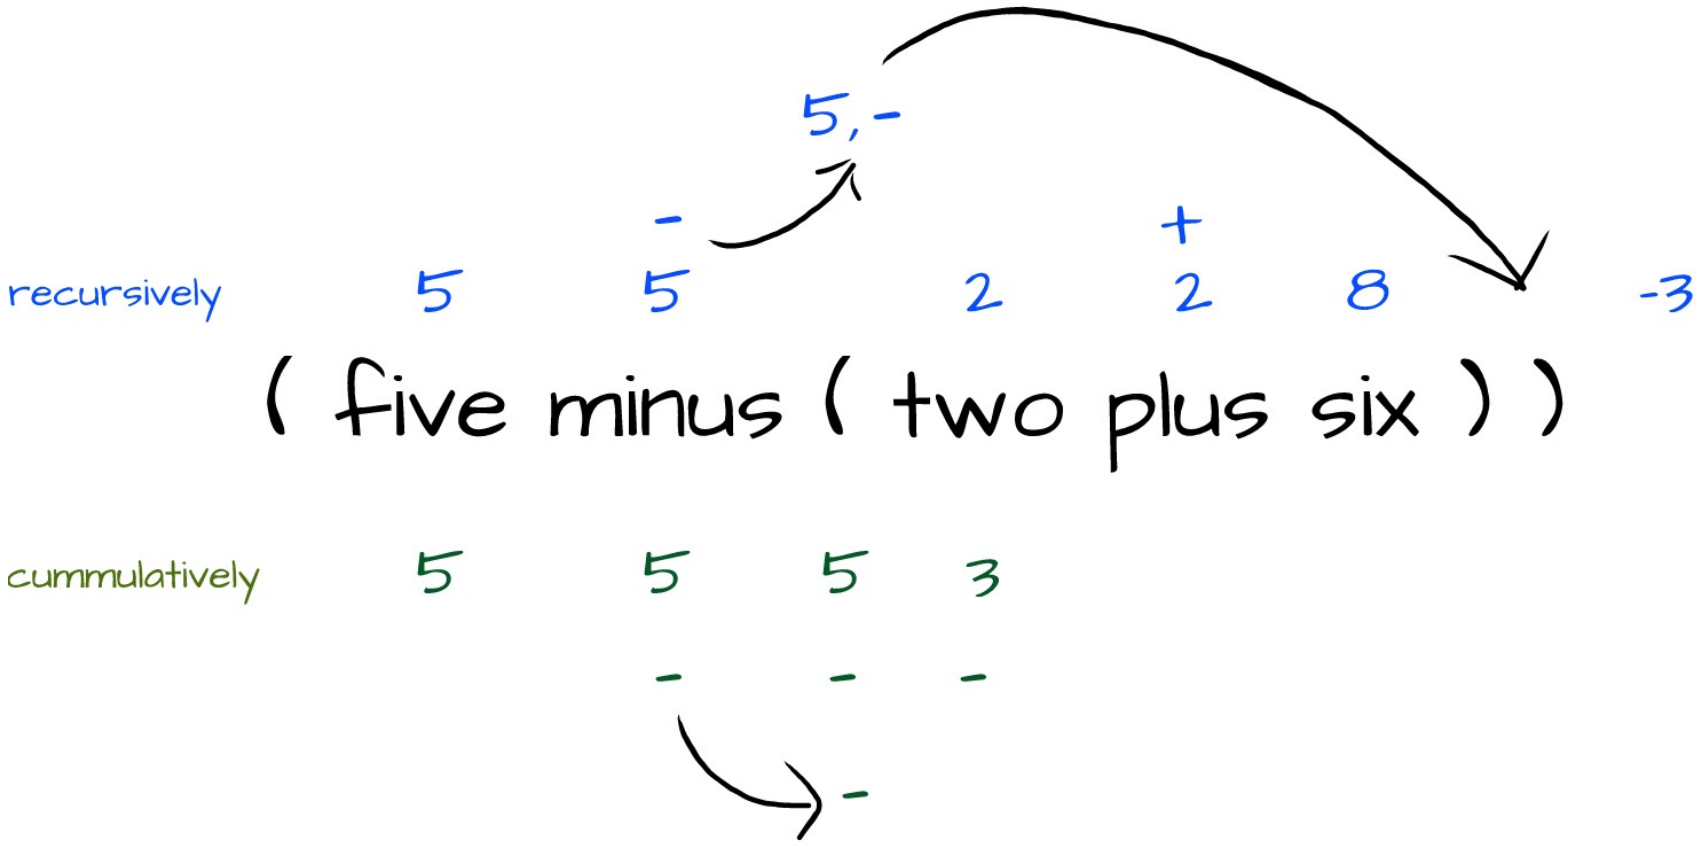

Symbolic solutions

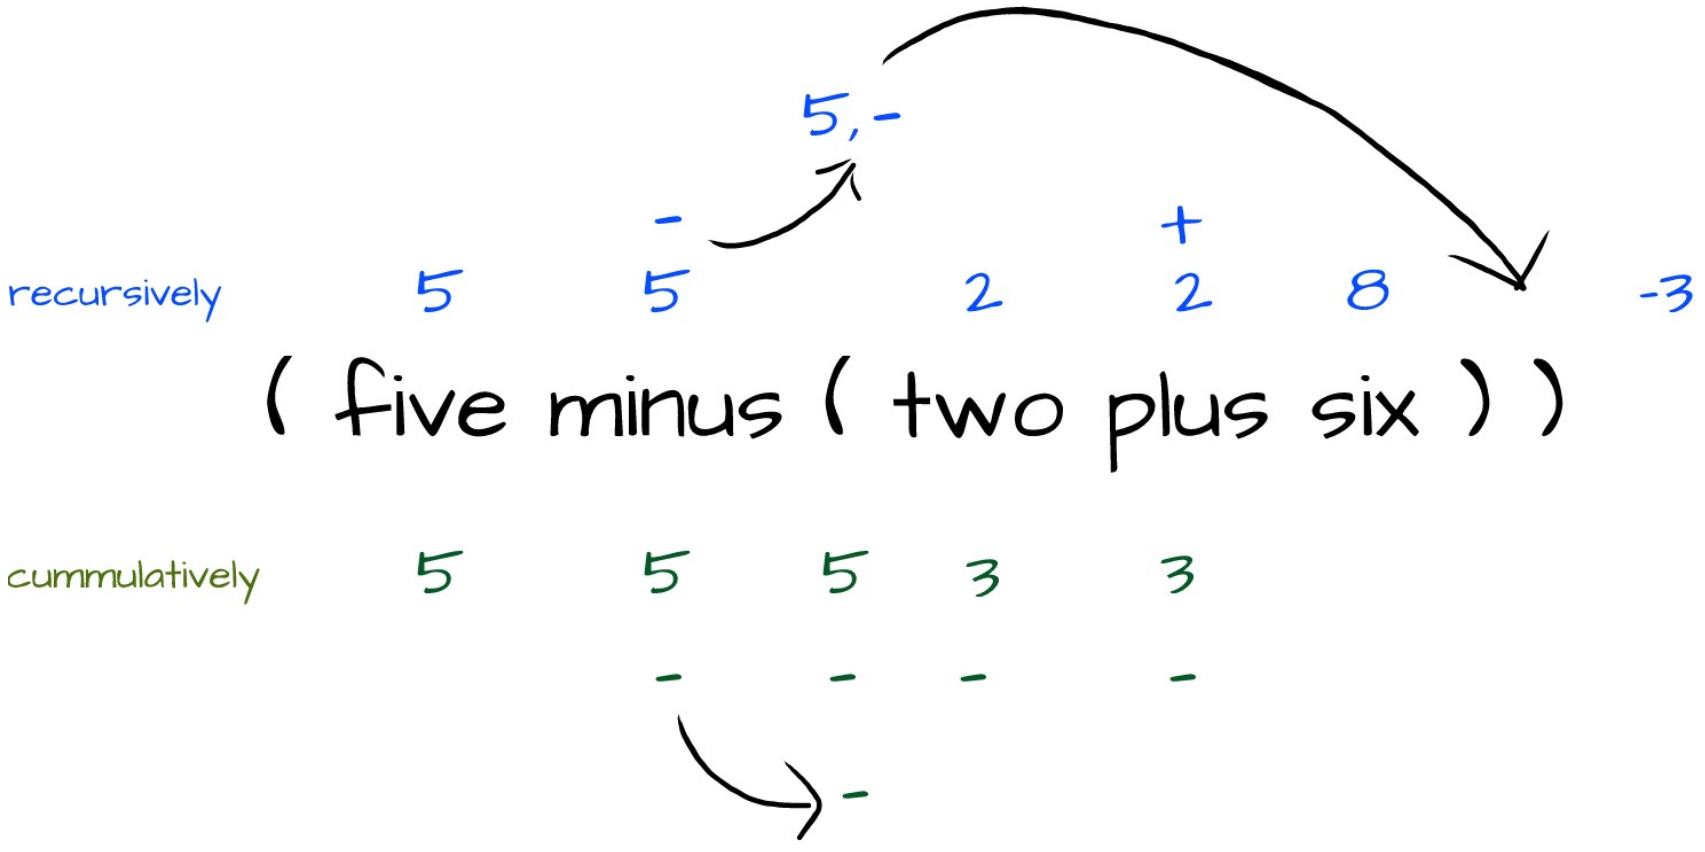

Symbolic solutions

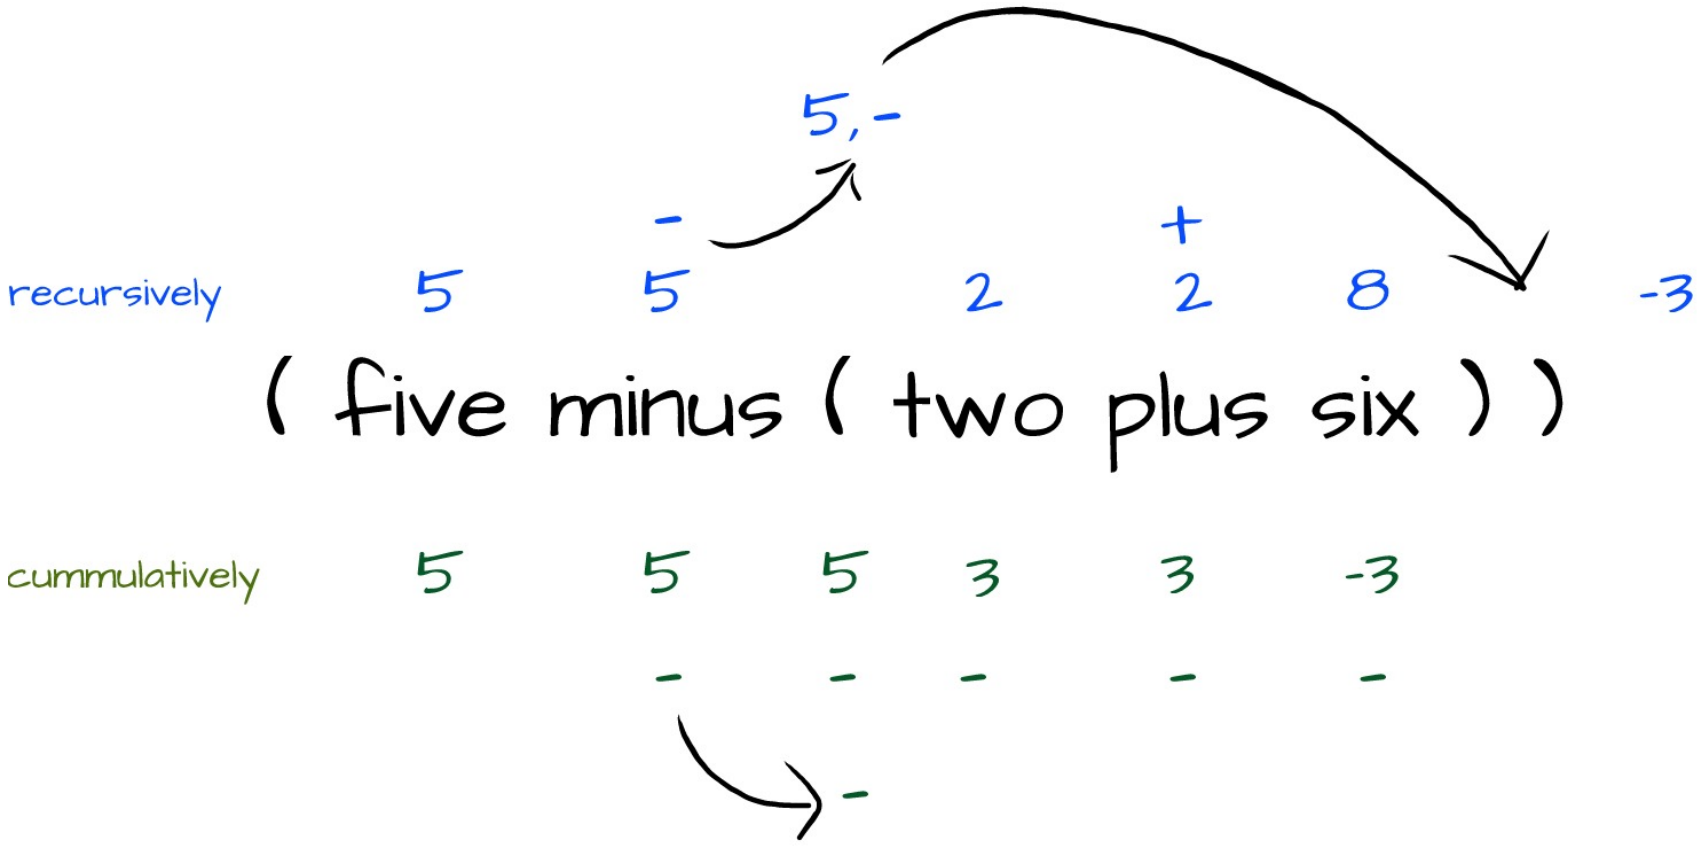

Symbolic solutions

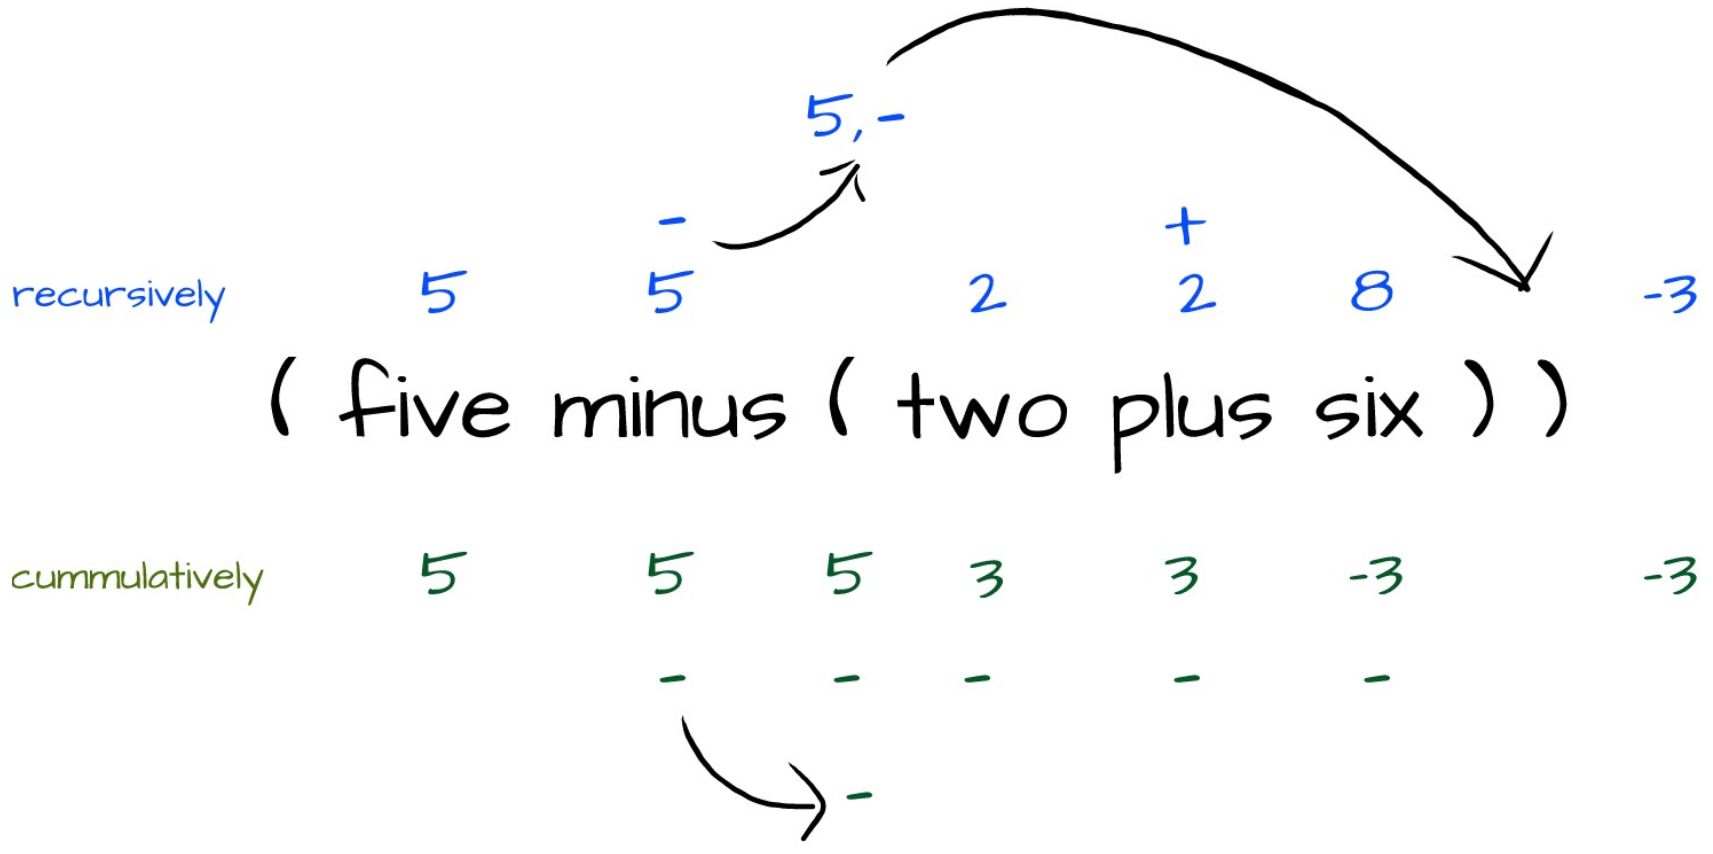

How do we study the network?

Plotting activation values

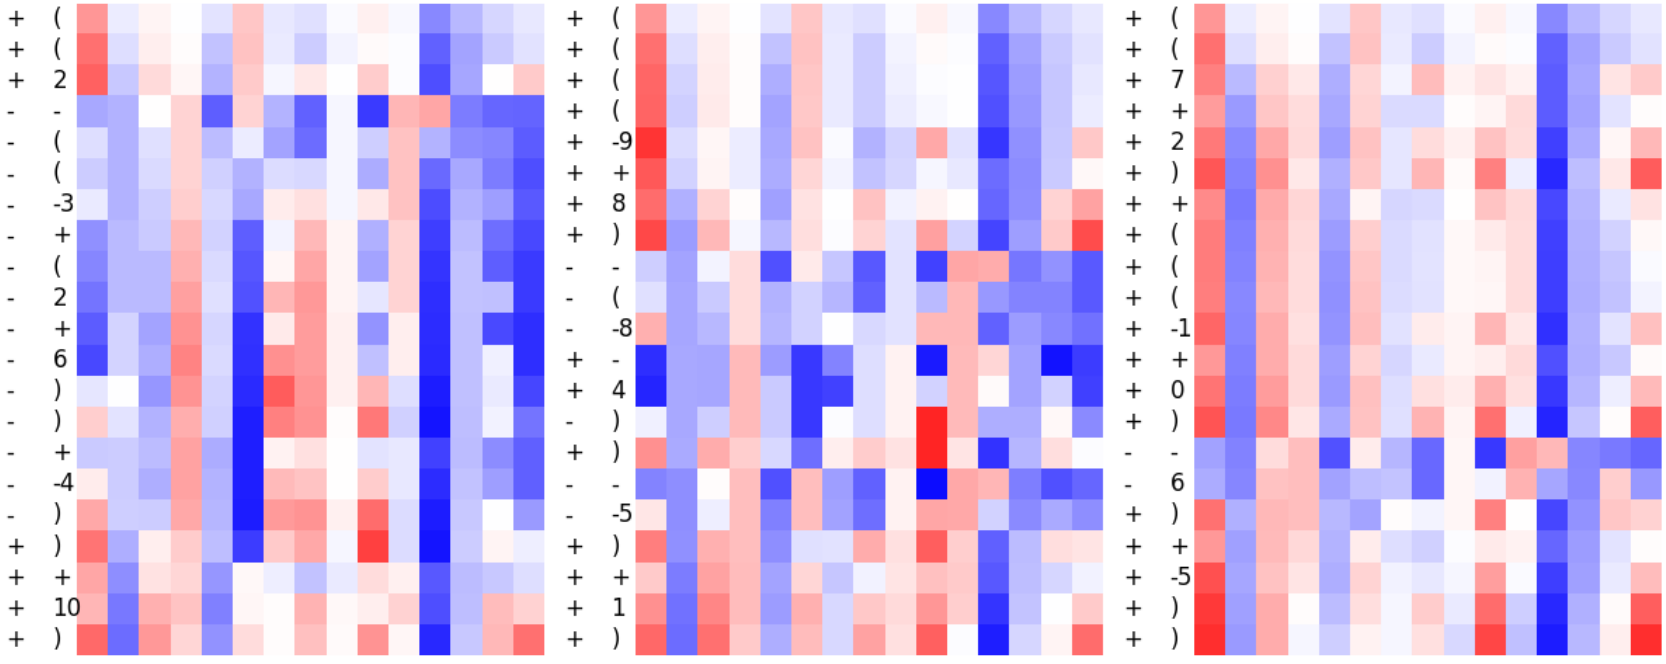

How do we study the network?

Diagnostic Classification

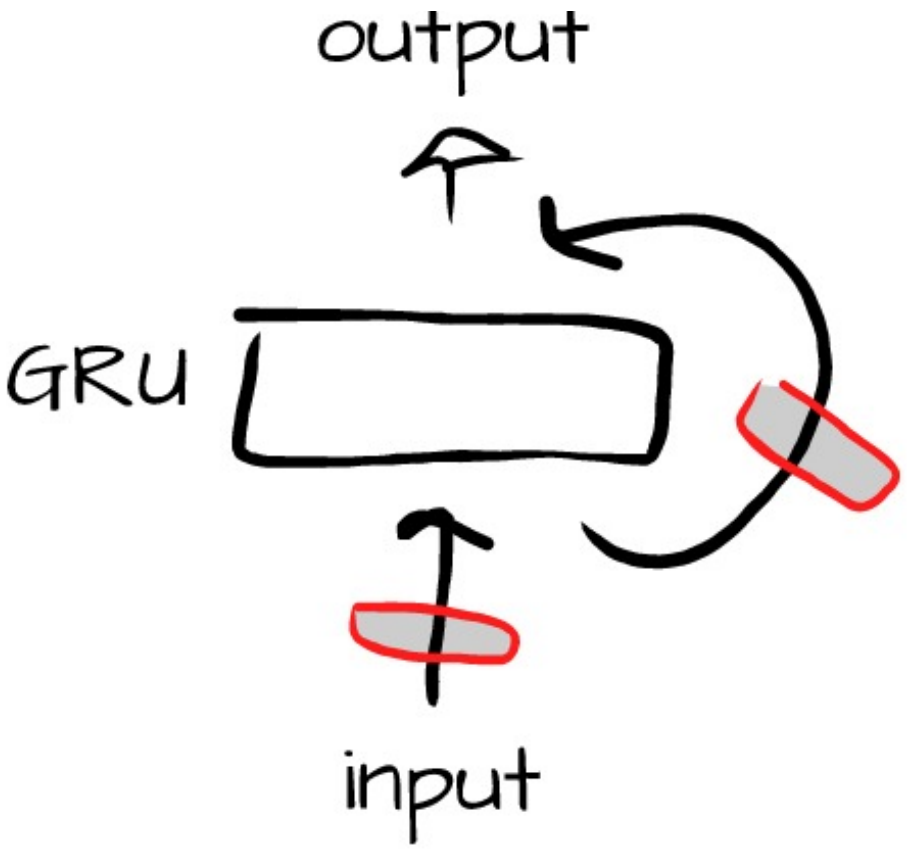

How do we study the network?

Diagnostic Classification

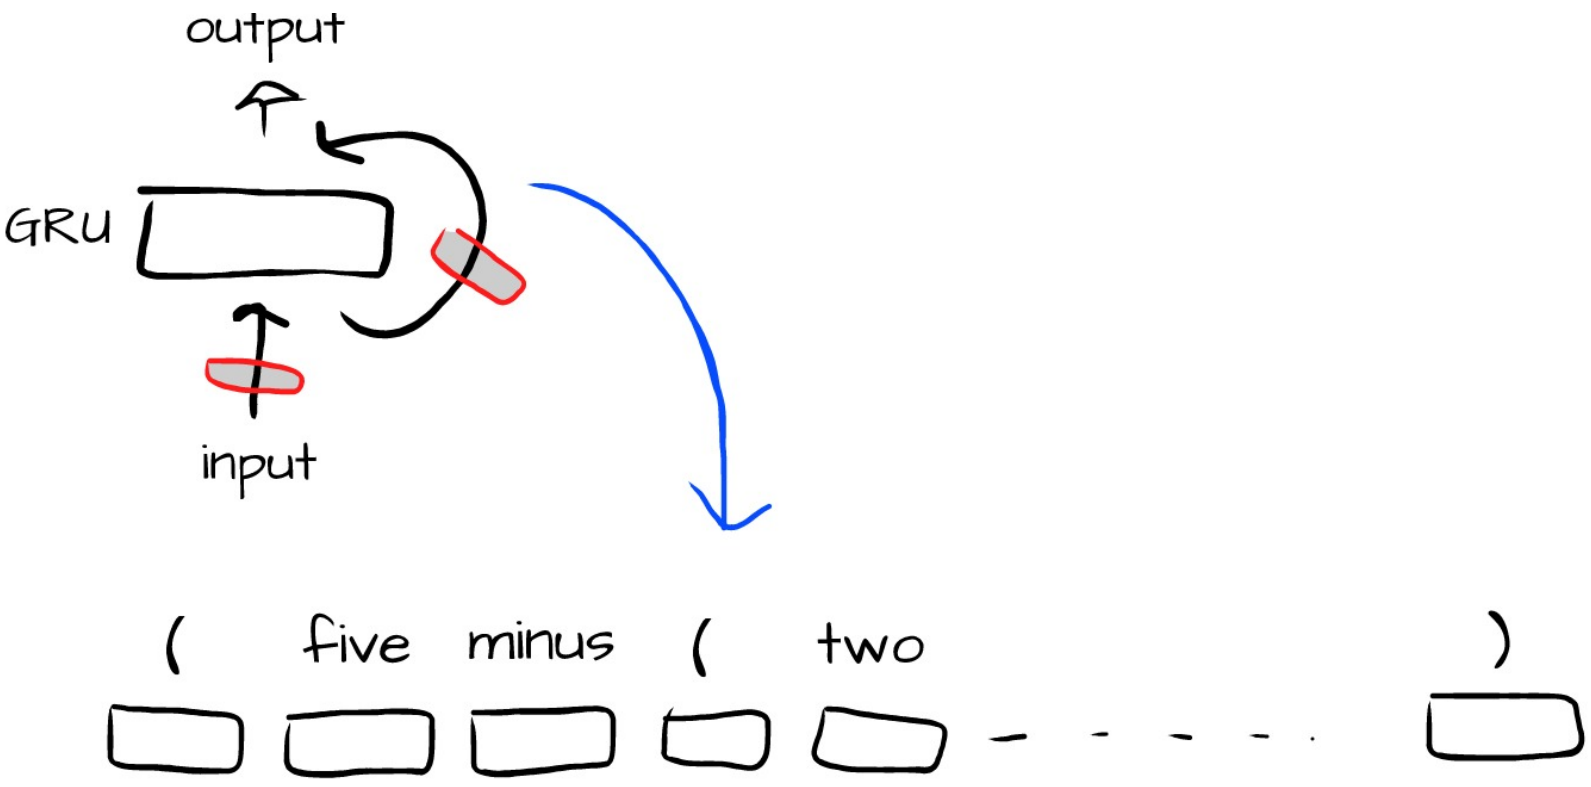

How do we study the network?

Diagnostic Classification

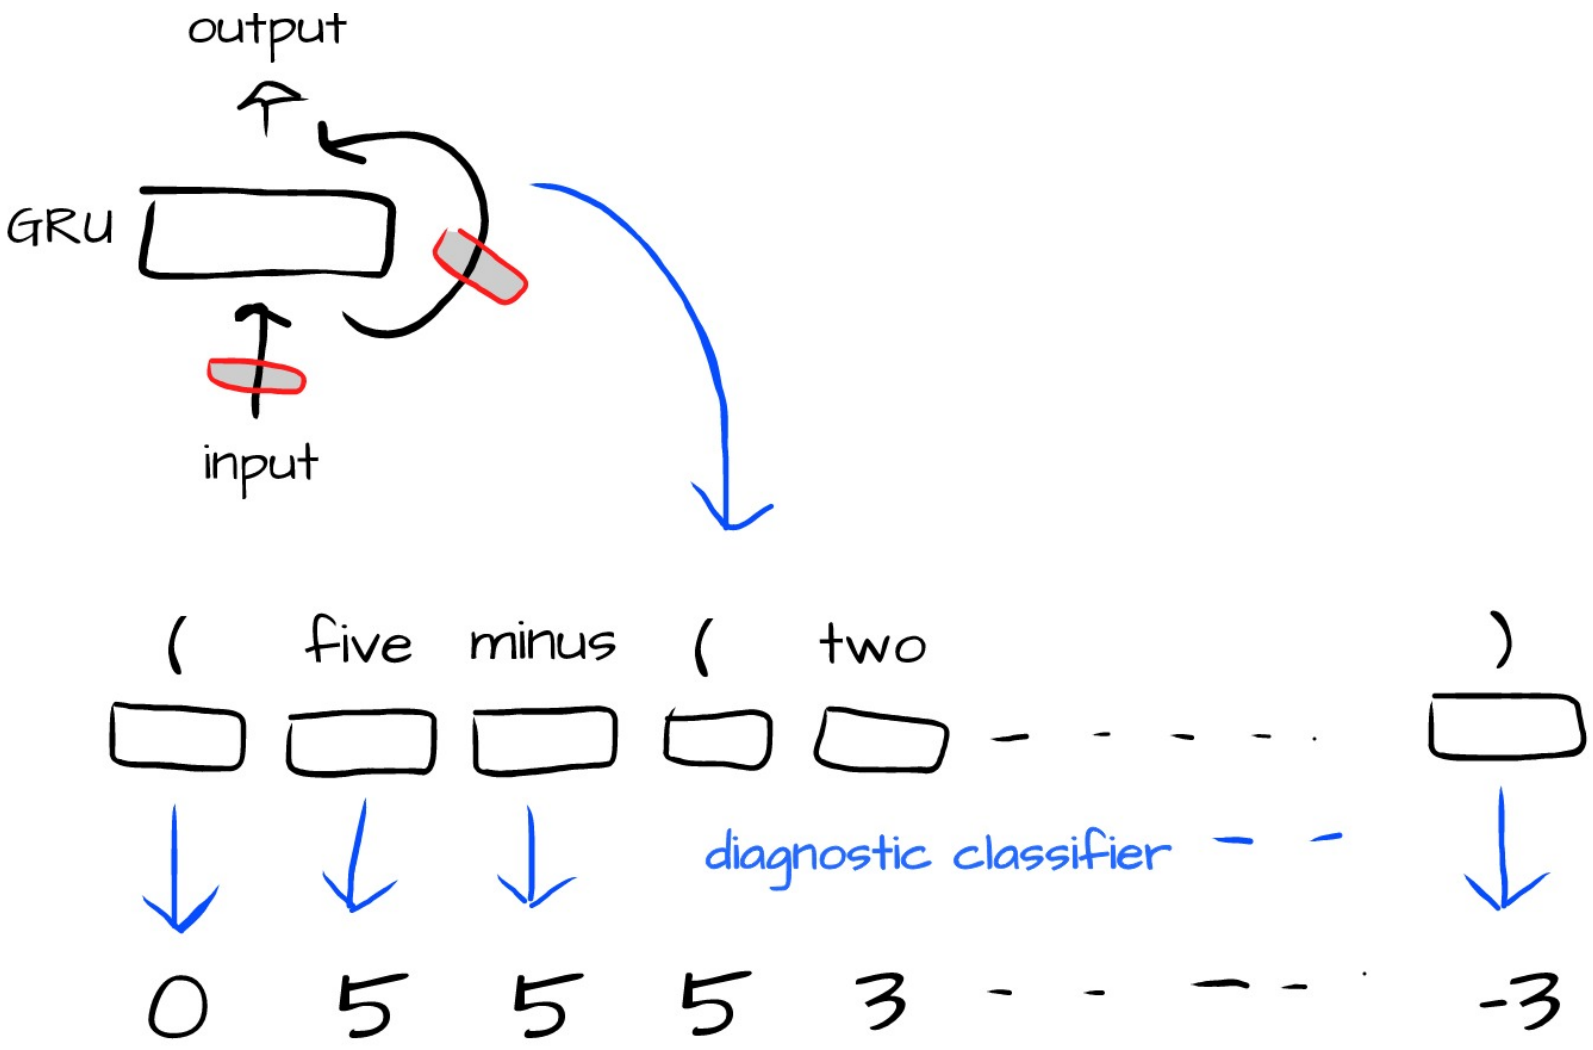

Intermediate results

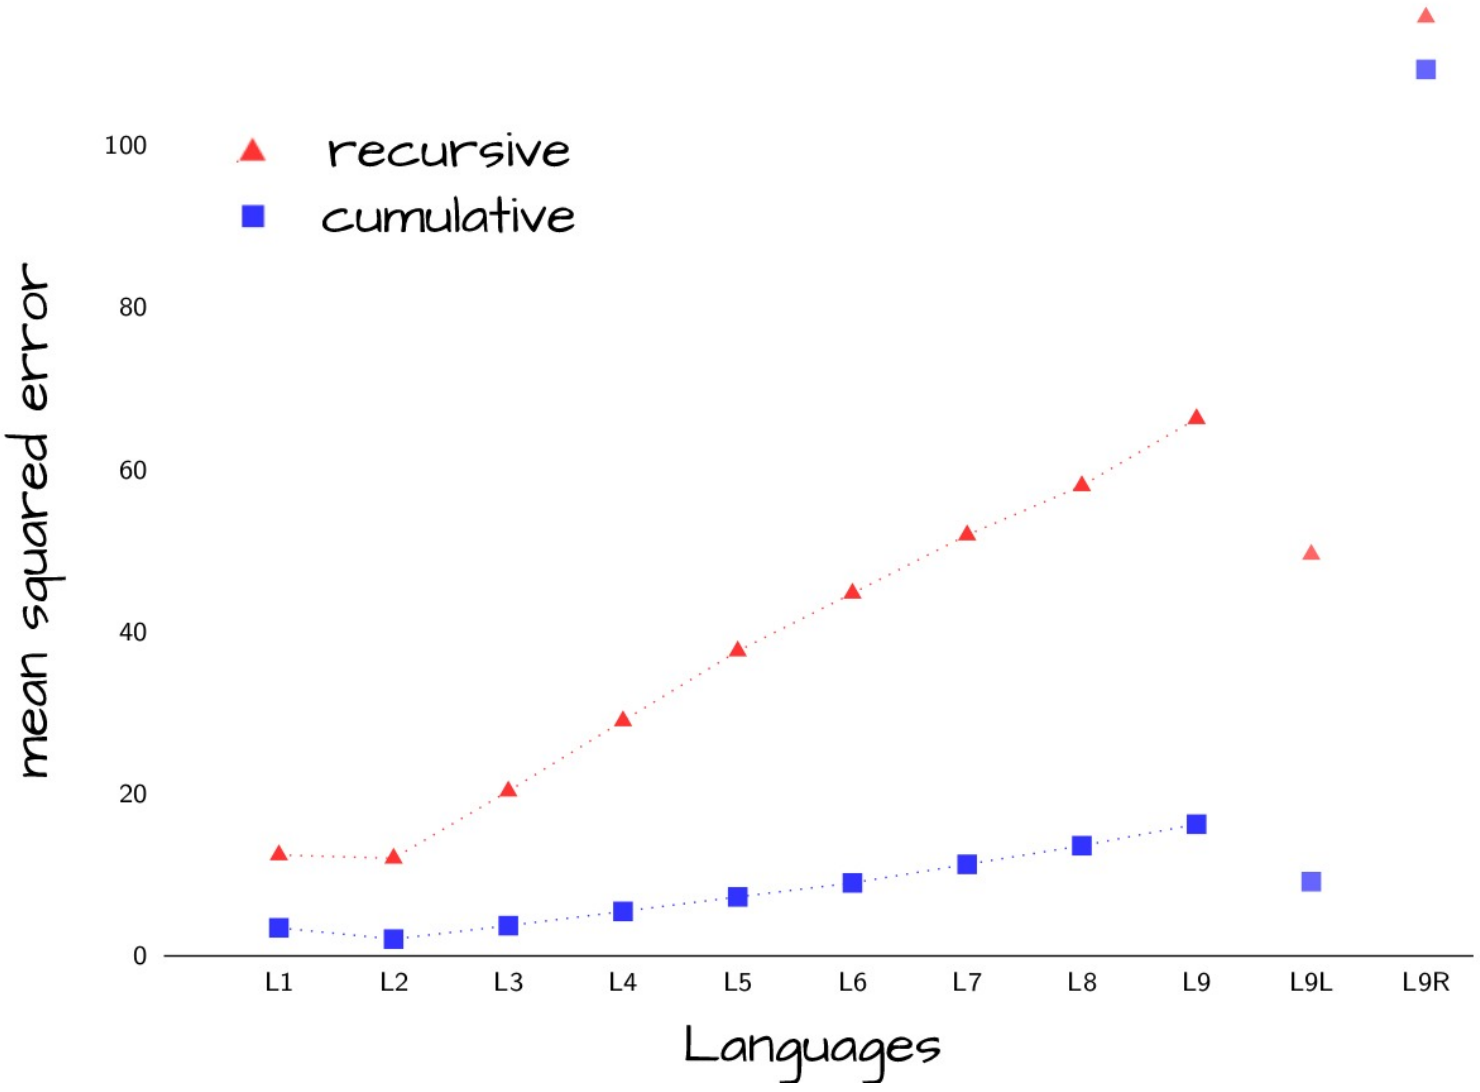

Cumulative strategy, operation mode

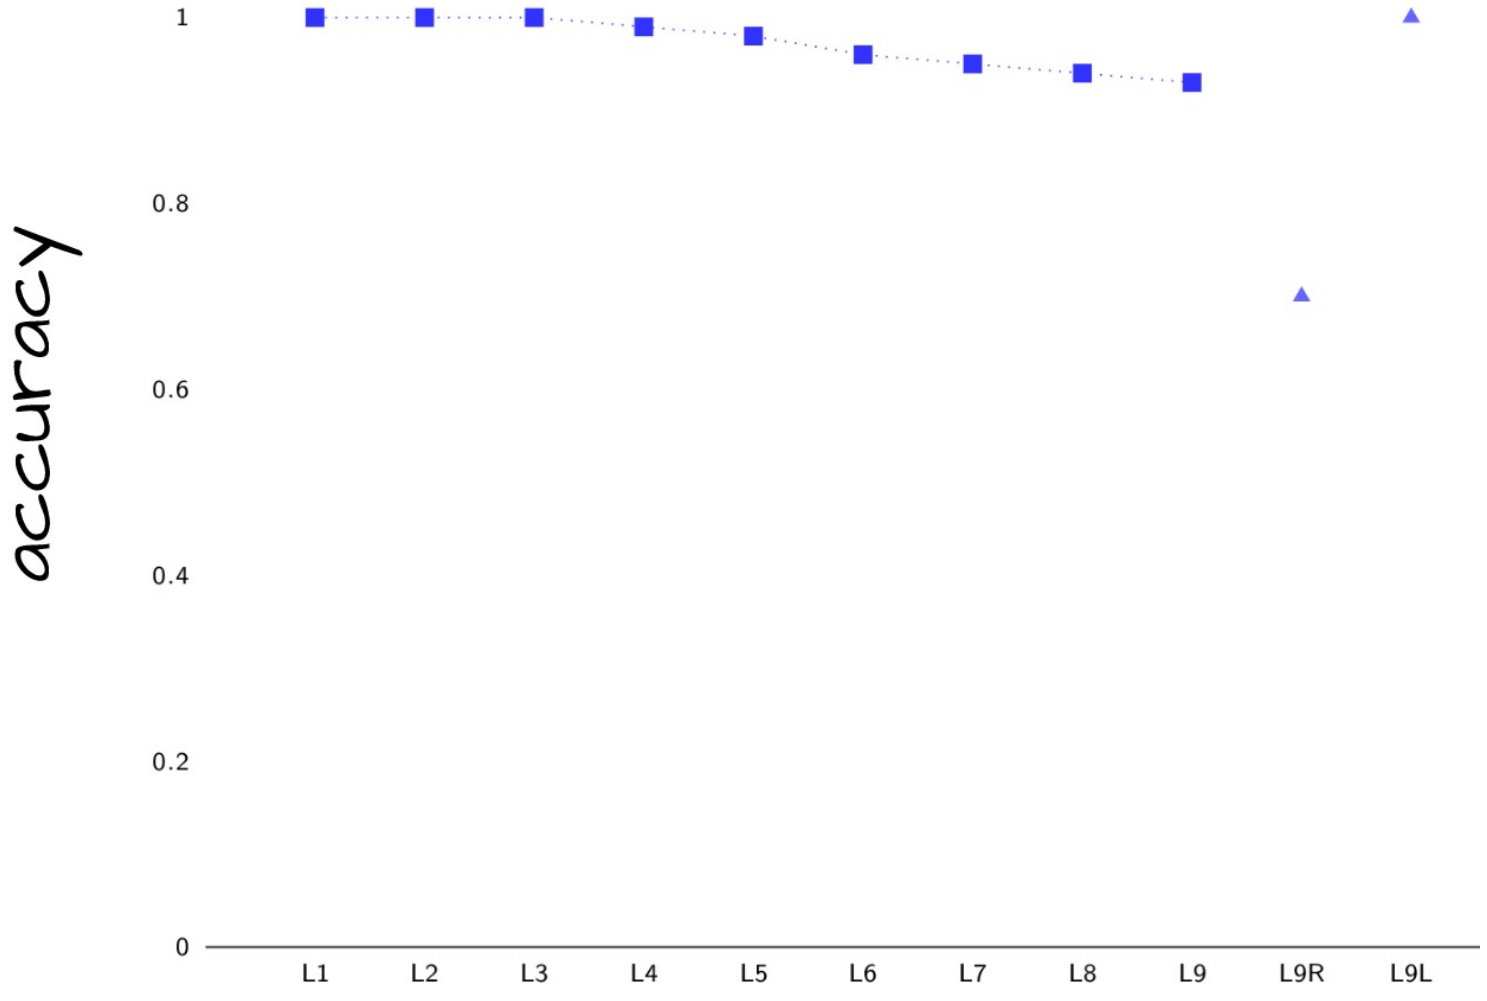

Learning about language by opening the neural network black box

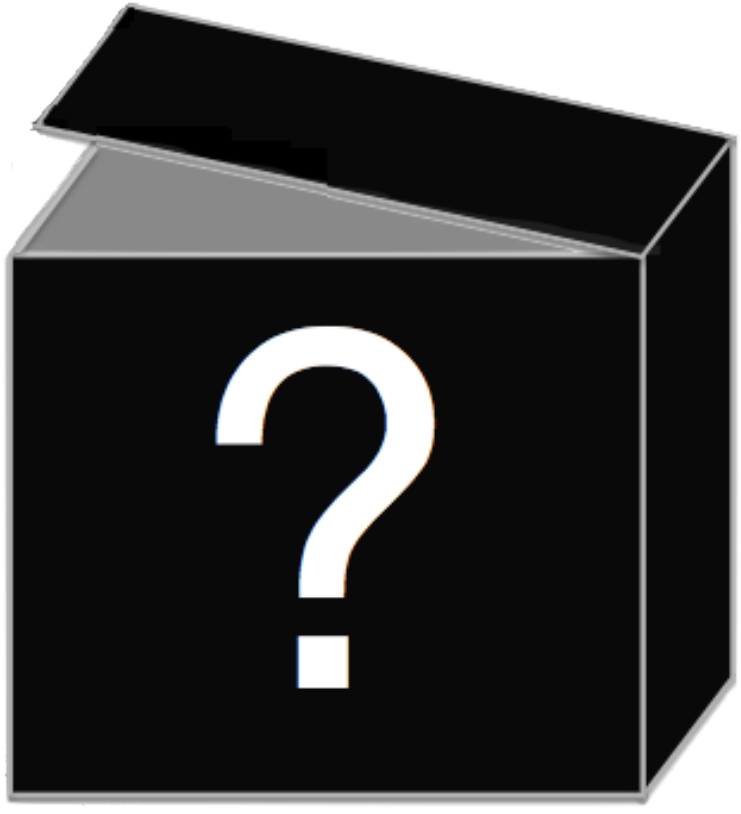

Learning about language by opening the neural network black box

What would Remko have thought about it?