Hierarchical Compositionality in Recurrent Neural Networks

Dieuwke Hupkes

Institute for Logic, Language and Computation
University of Amsterdam

December 15, 2017

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?

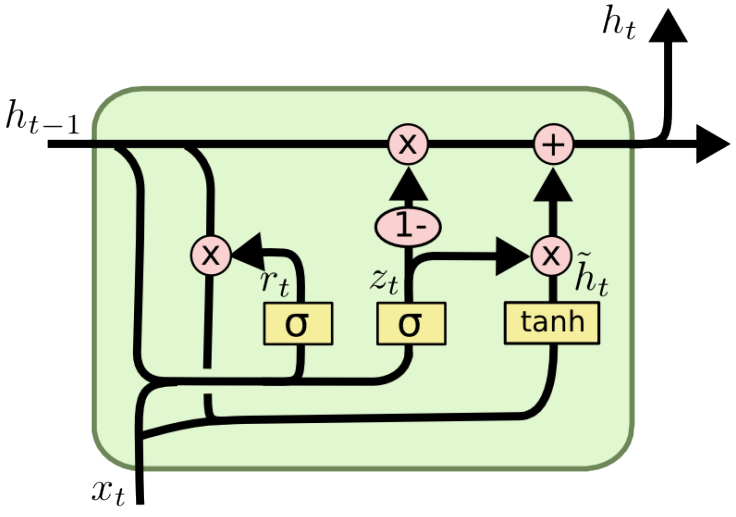

$$z_t = \sigma\left(W_z \cdot [h_{t-1}, x_t]\right)$$

$$r_t = \sigma\left(W_r \cdot [h_{t-1}, x_t]\right)$$

$$\tilde{h}_t = \tanh\left(W \cdot [r_t * h_{t-1}, x_t]\right)$$

$$h_t = (1 - z_t) * h_{t-1} + z_t * \tilde{h}_t$$

Arithmetic Language

Name	*#digits*	*Example*
L1	1	`minus three`
L2	2	`( five plus seven )`
L3	3	`( three - ( one + minus two ) )`
...		
L5R	5	`( ( ( ( nine + six) + seven ) + five ) - seven )`
L5L	5	`( eight + ( six- ( two - ( ten + nine ) ) ) )`

Arithmetic Language

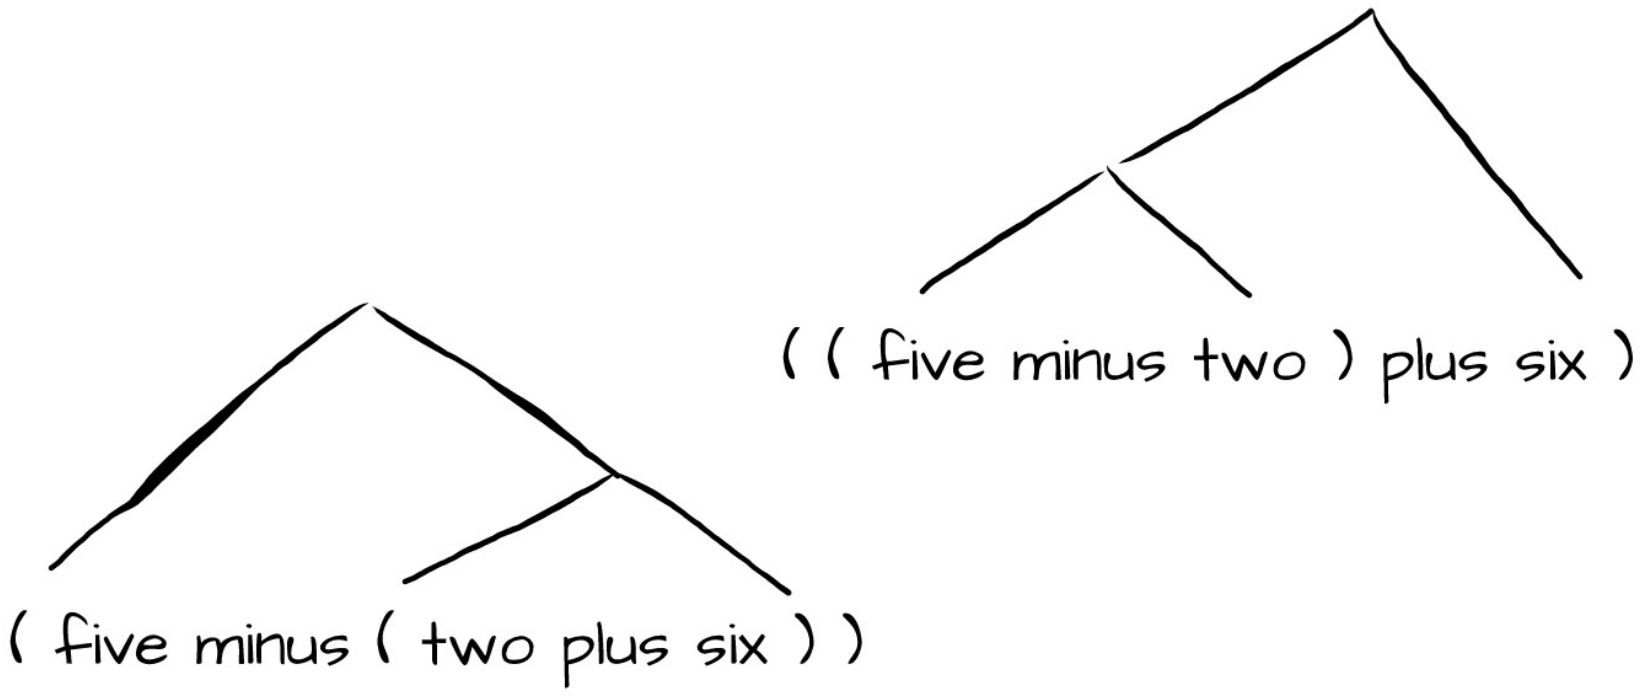

Results

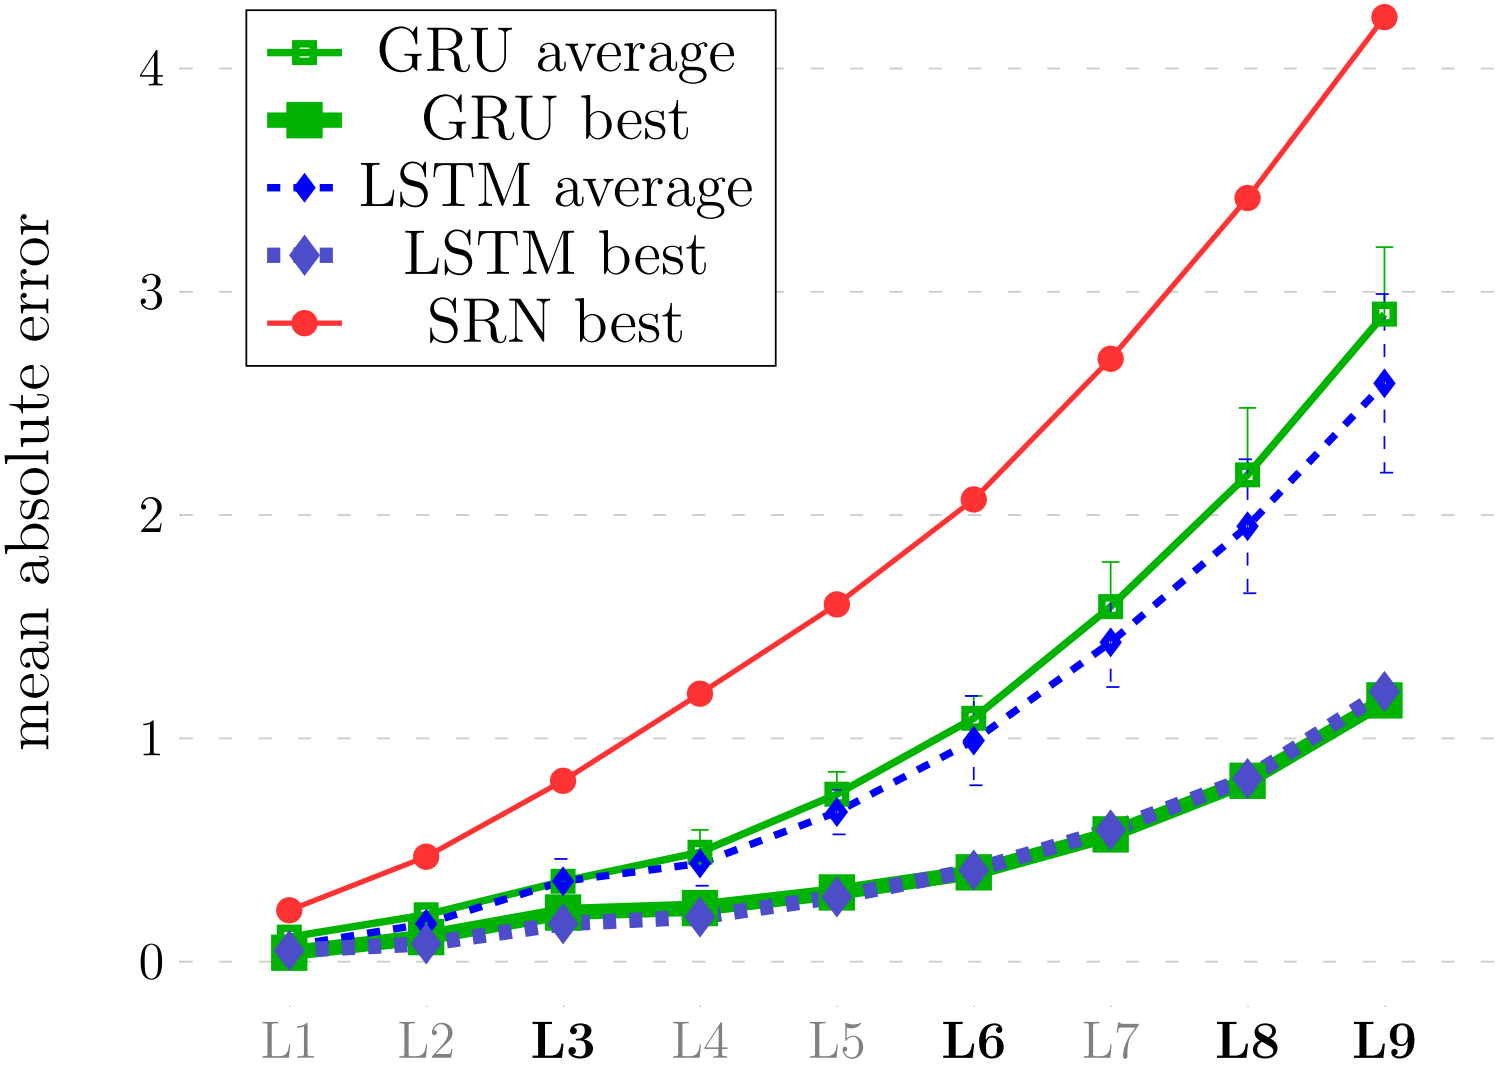
GRU average
GRU best
LSTM average
LSTM best
SRN best
mean absolute error
0
1
2
3
4
L1
L2
L3
L4
L5
L6
L7
L8
L9

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

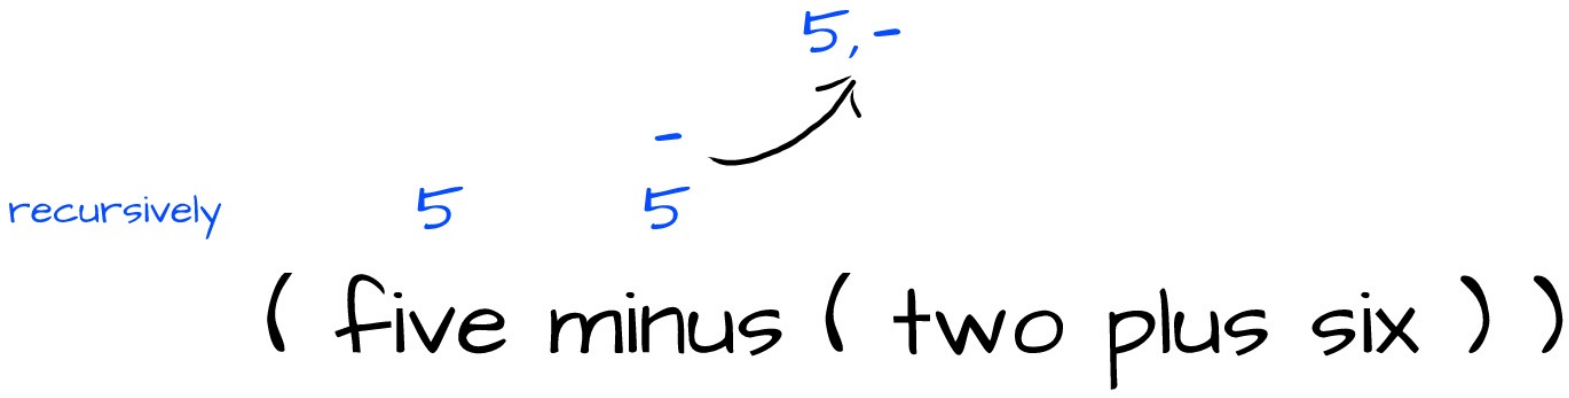

Symbolic solutions

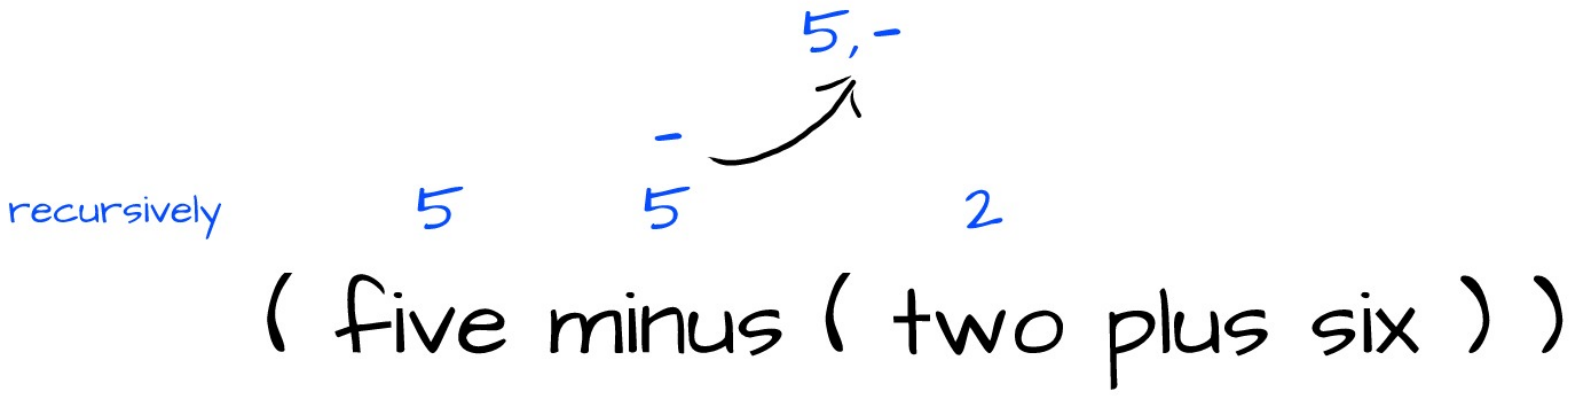

Symbolic solutions

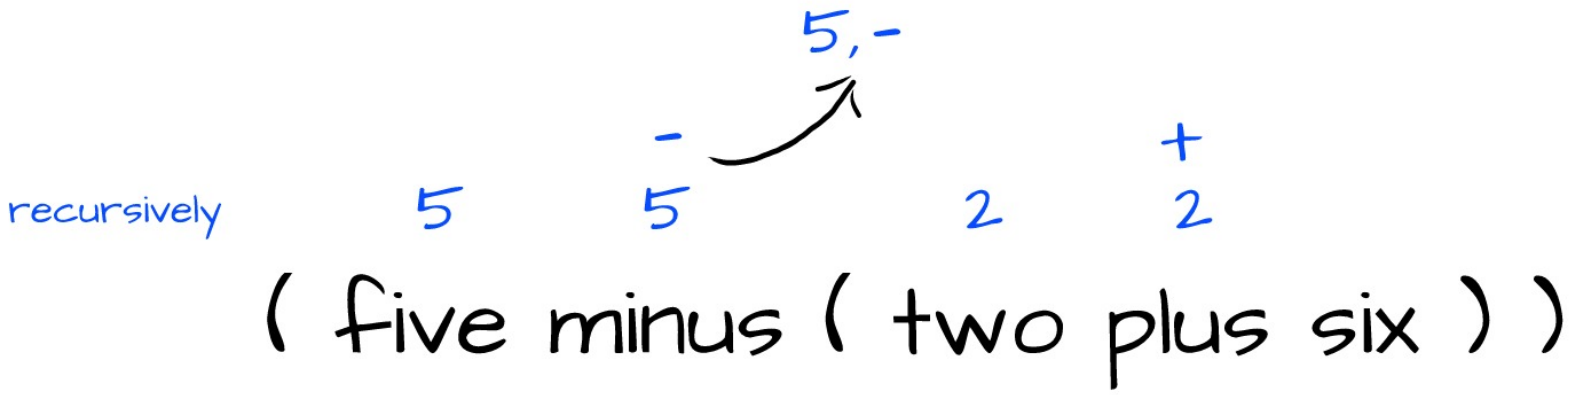

Symbolic solutions

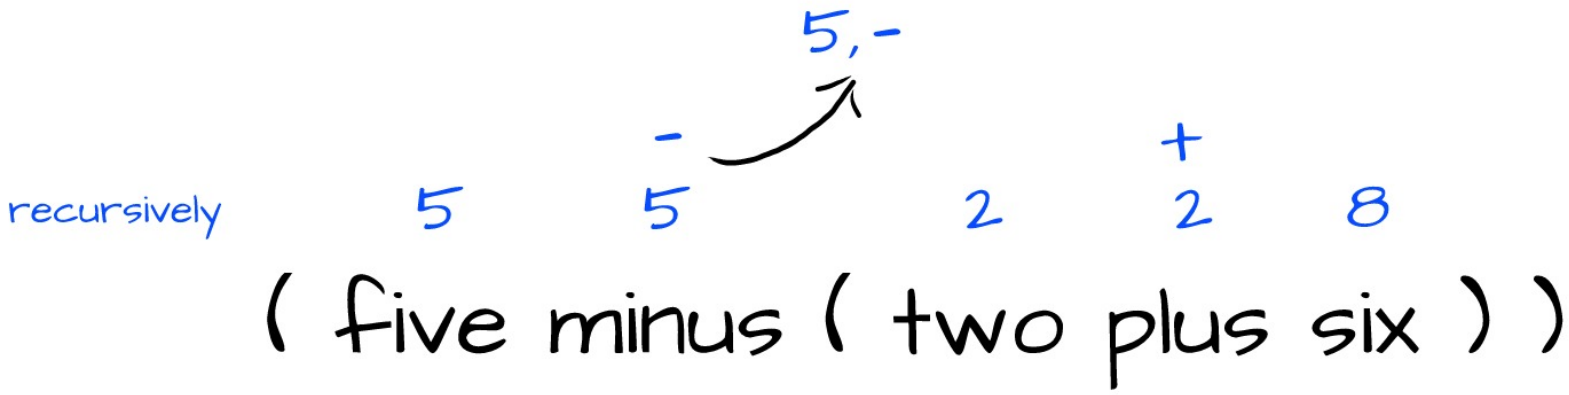

Symbolic solutions

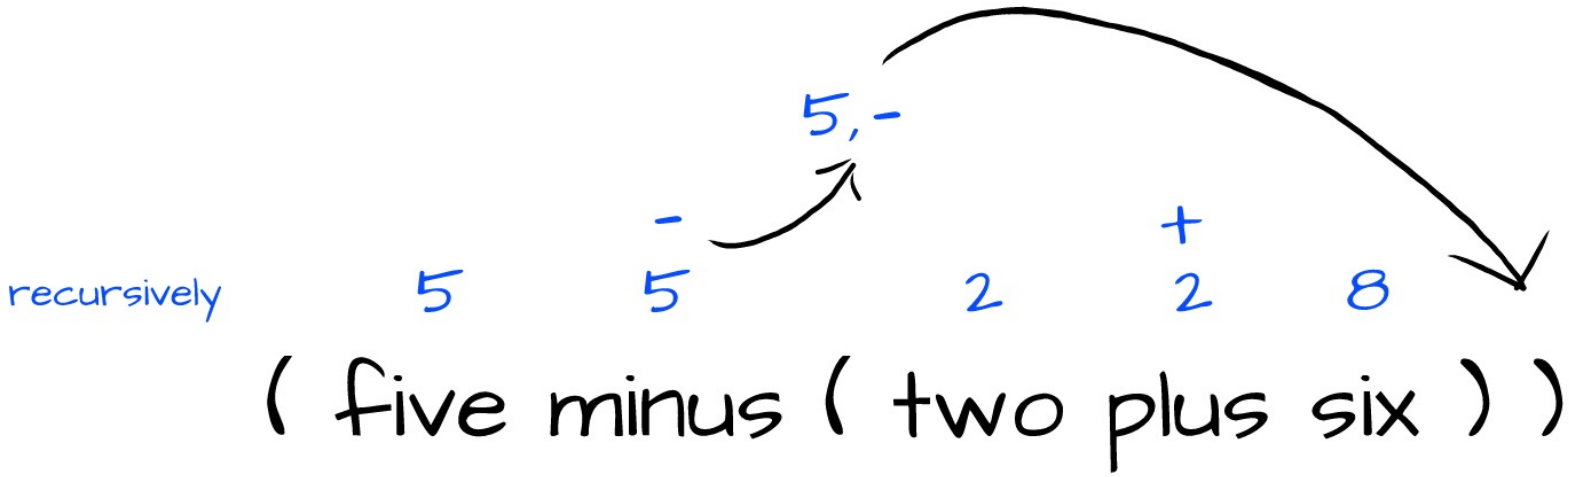

Symbolic solutions

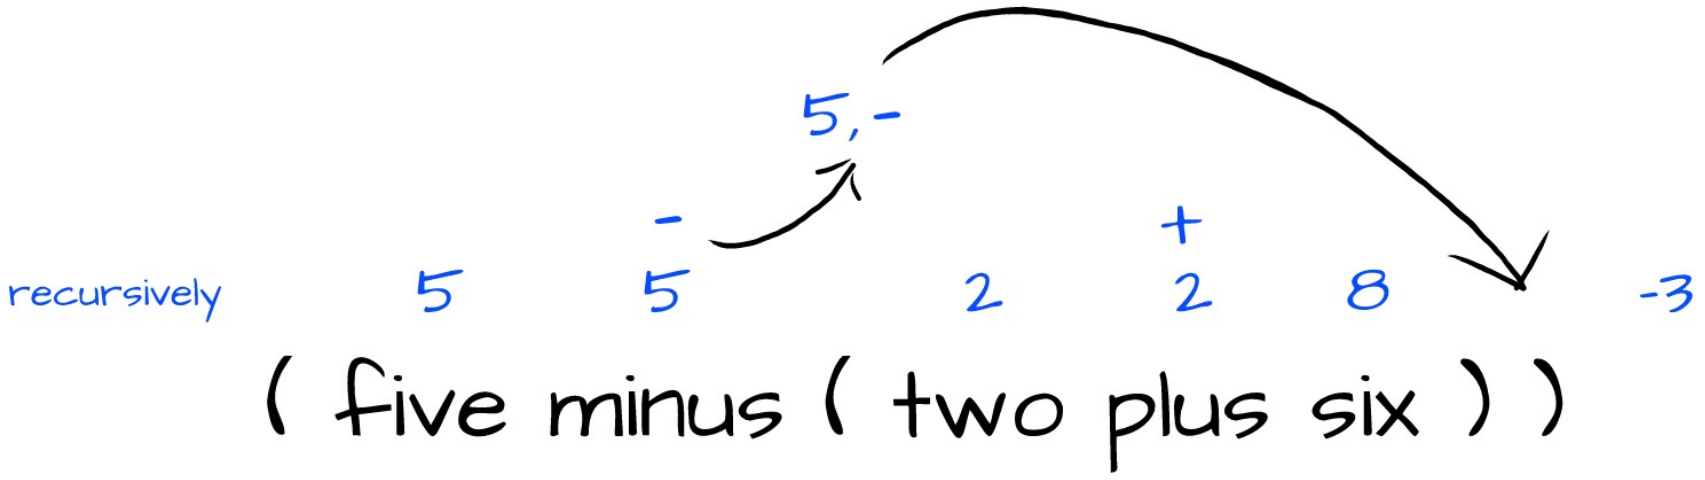

Symbolic solutions

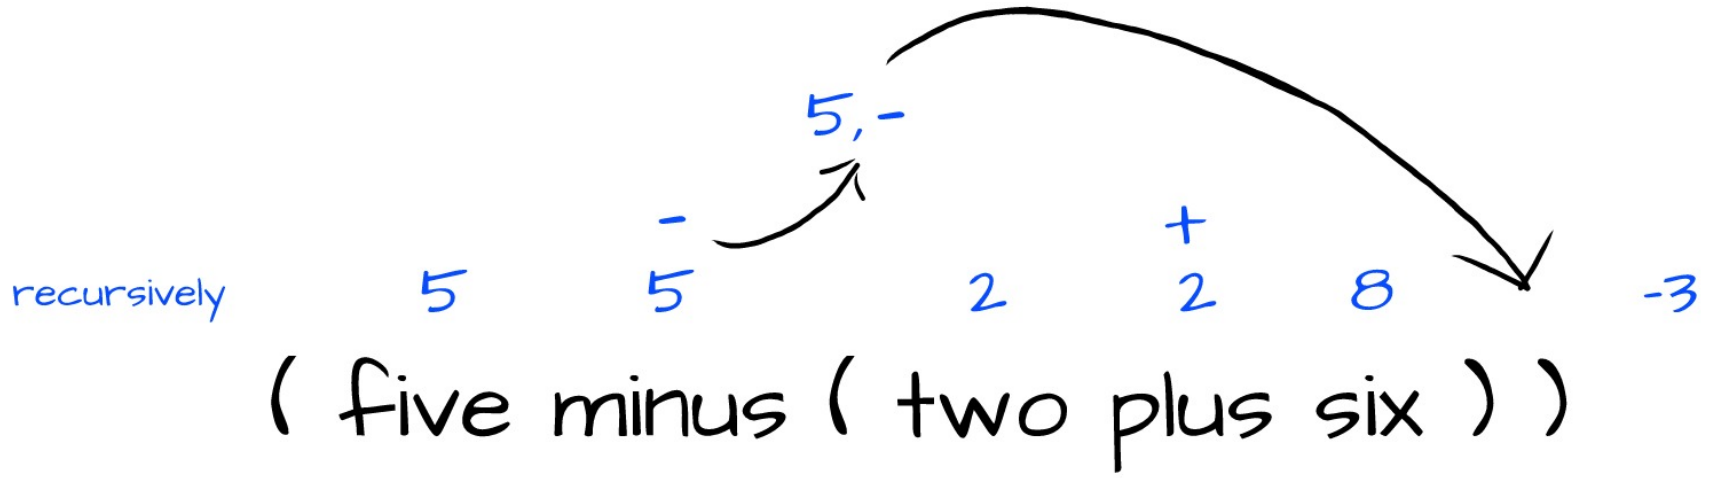

cummulatively

Symbolic solutions

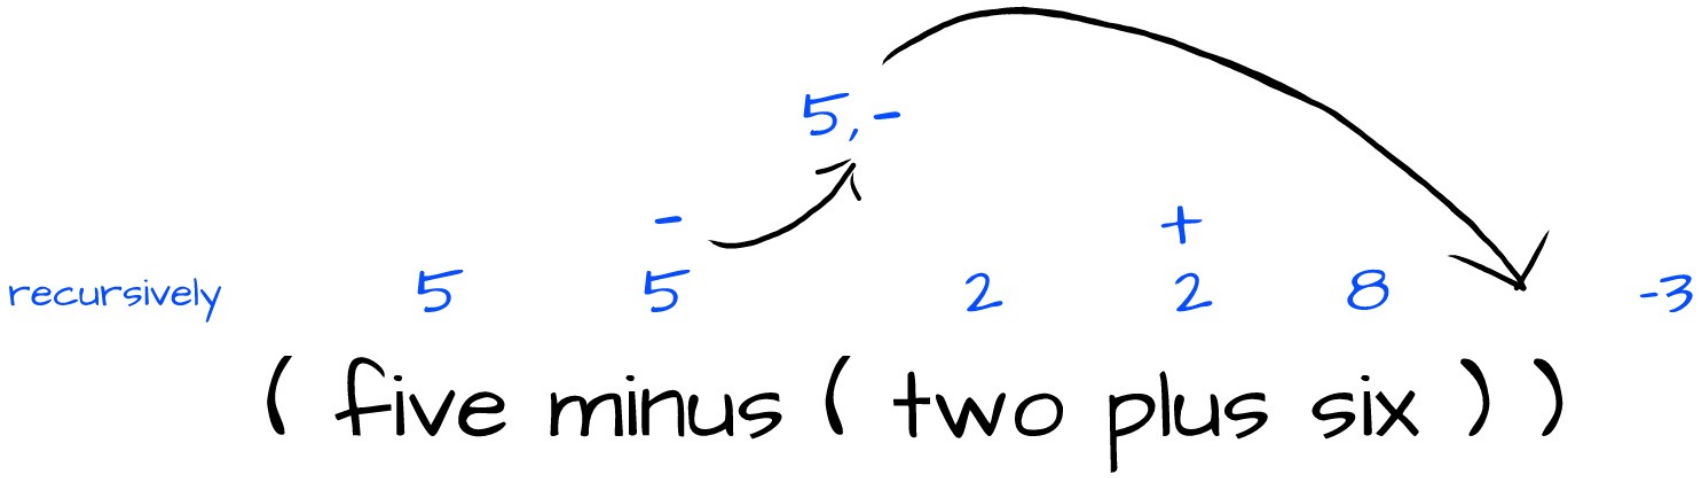

Symbolic solutions

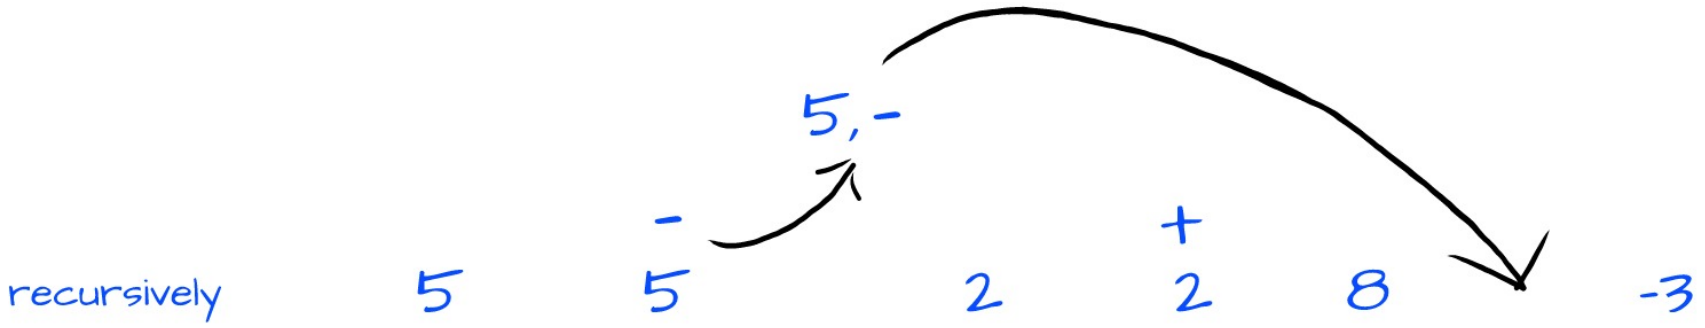

Symbolic solutions

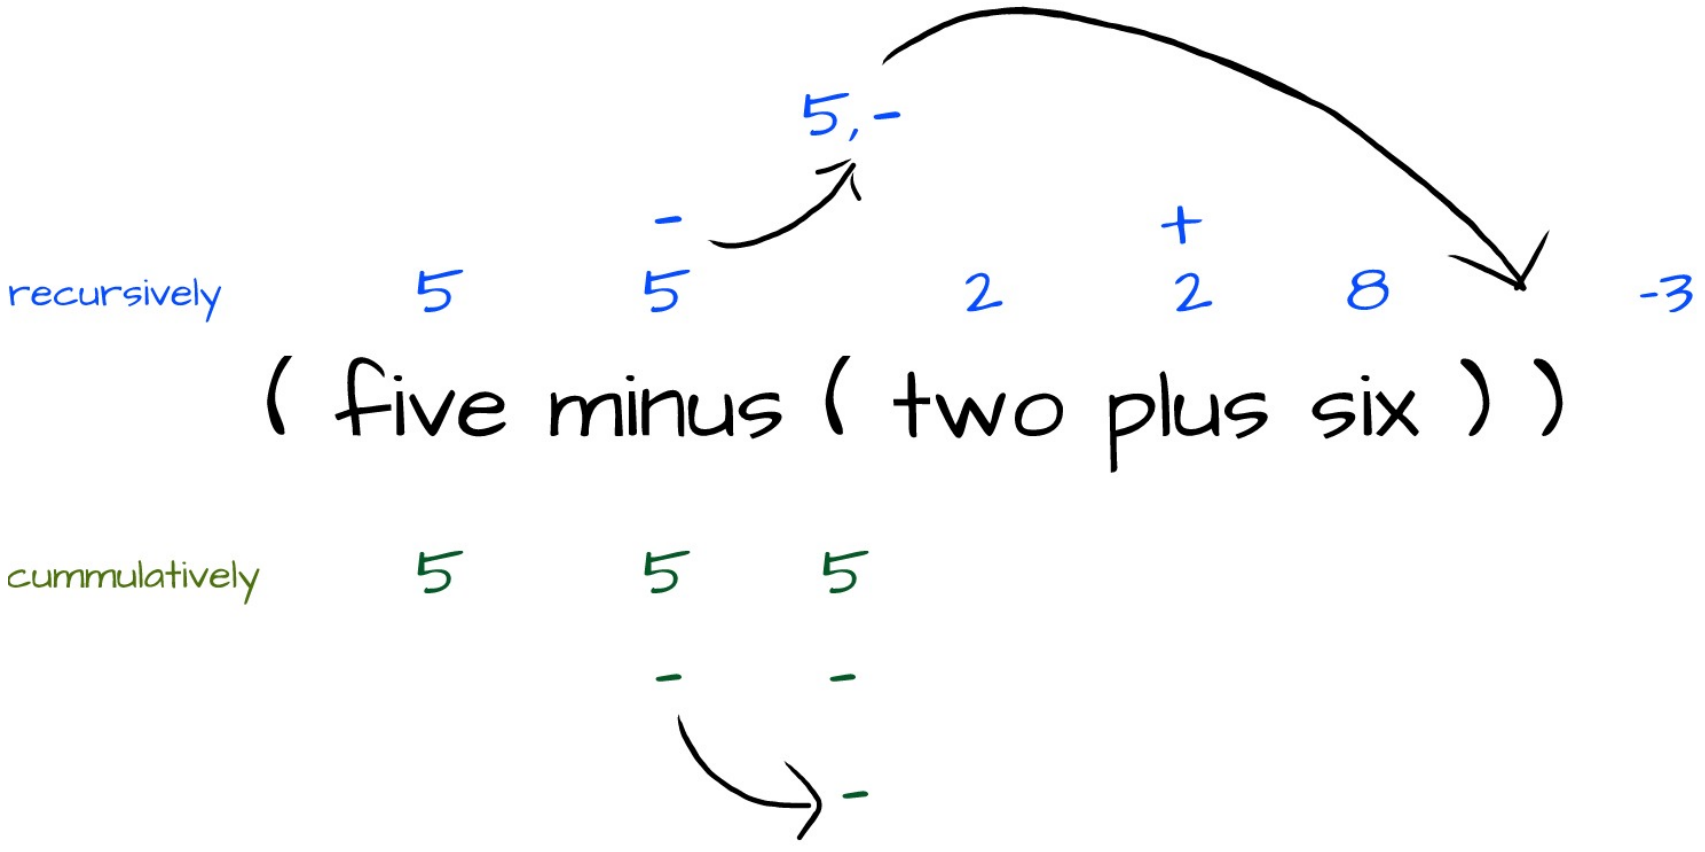

Symbolic solutions

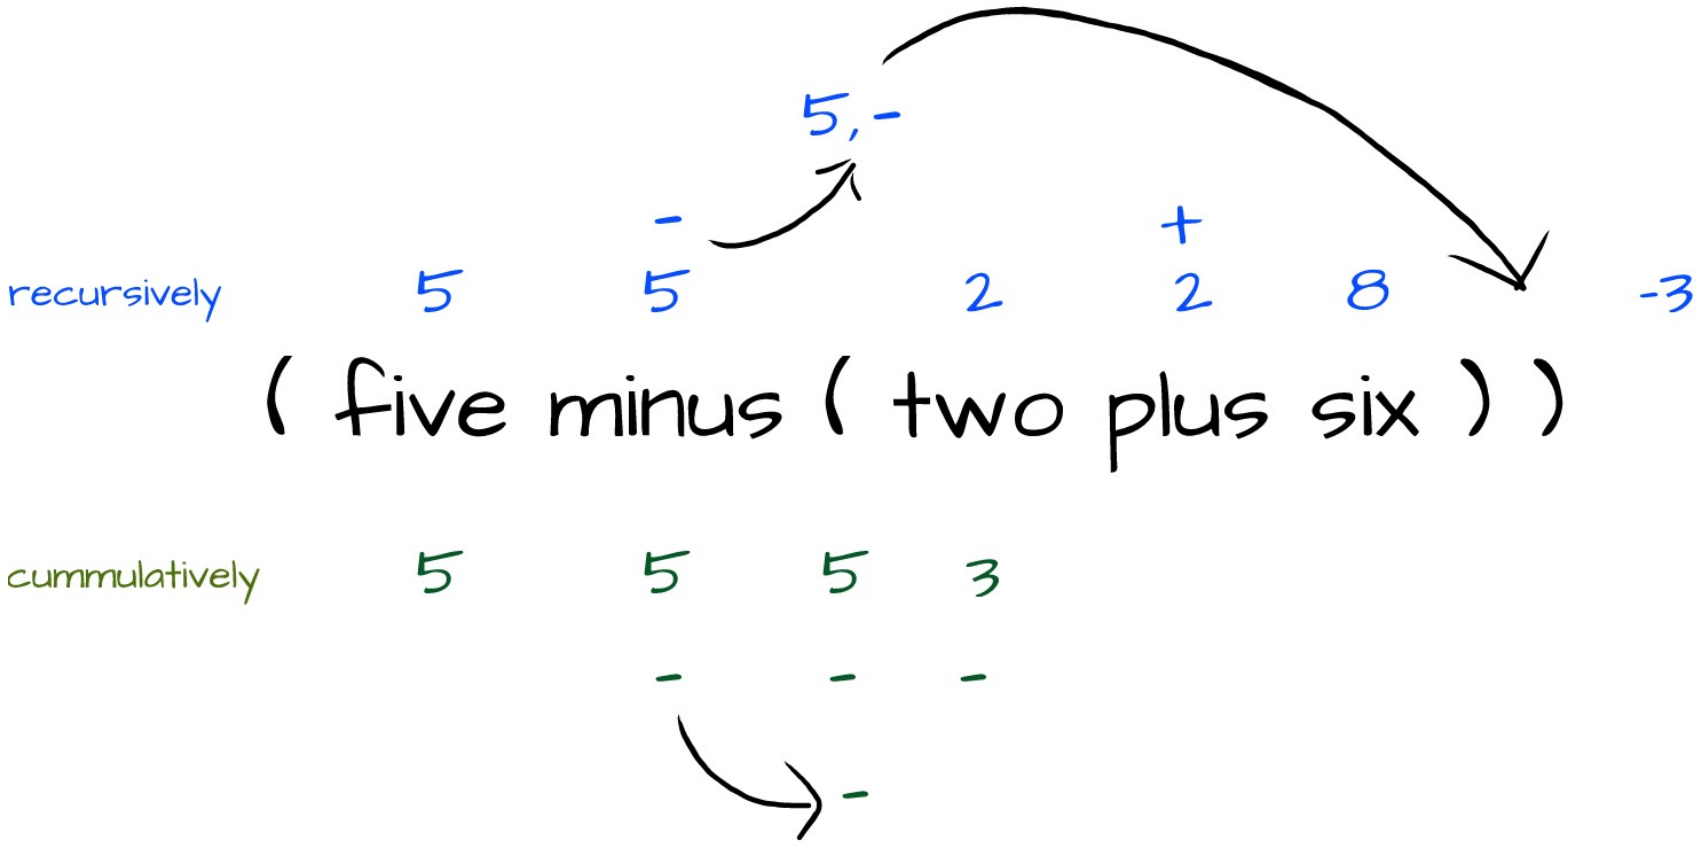

Symbolic solutions

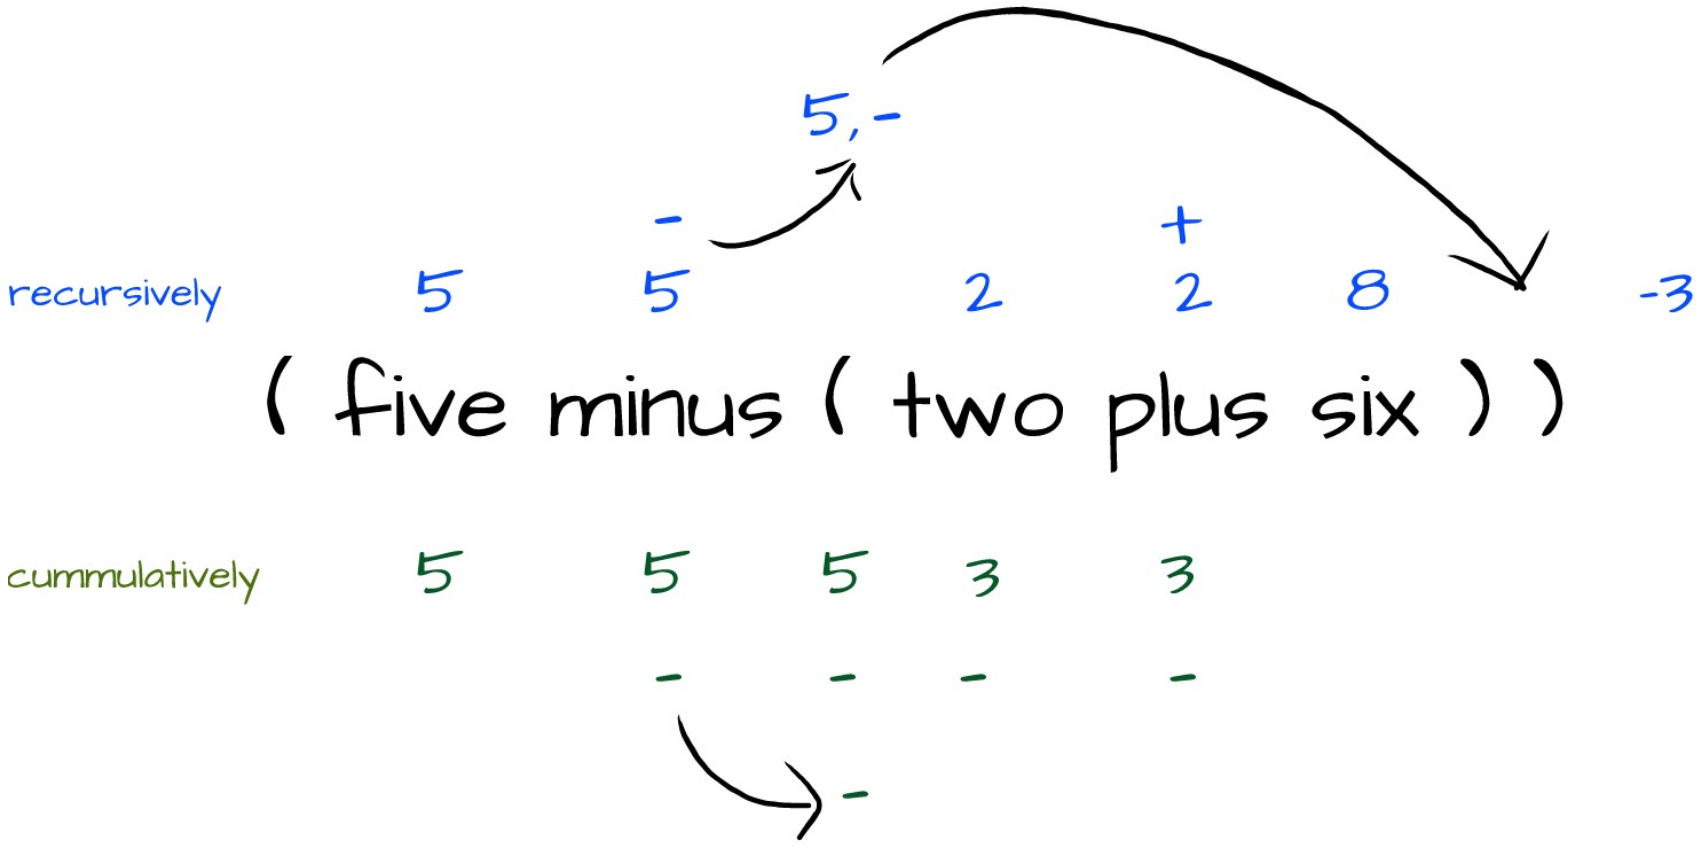

Symbolic solutions

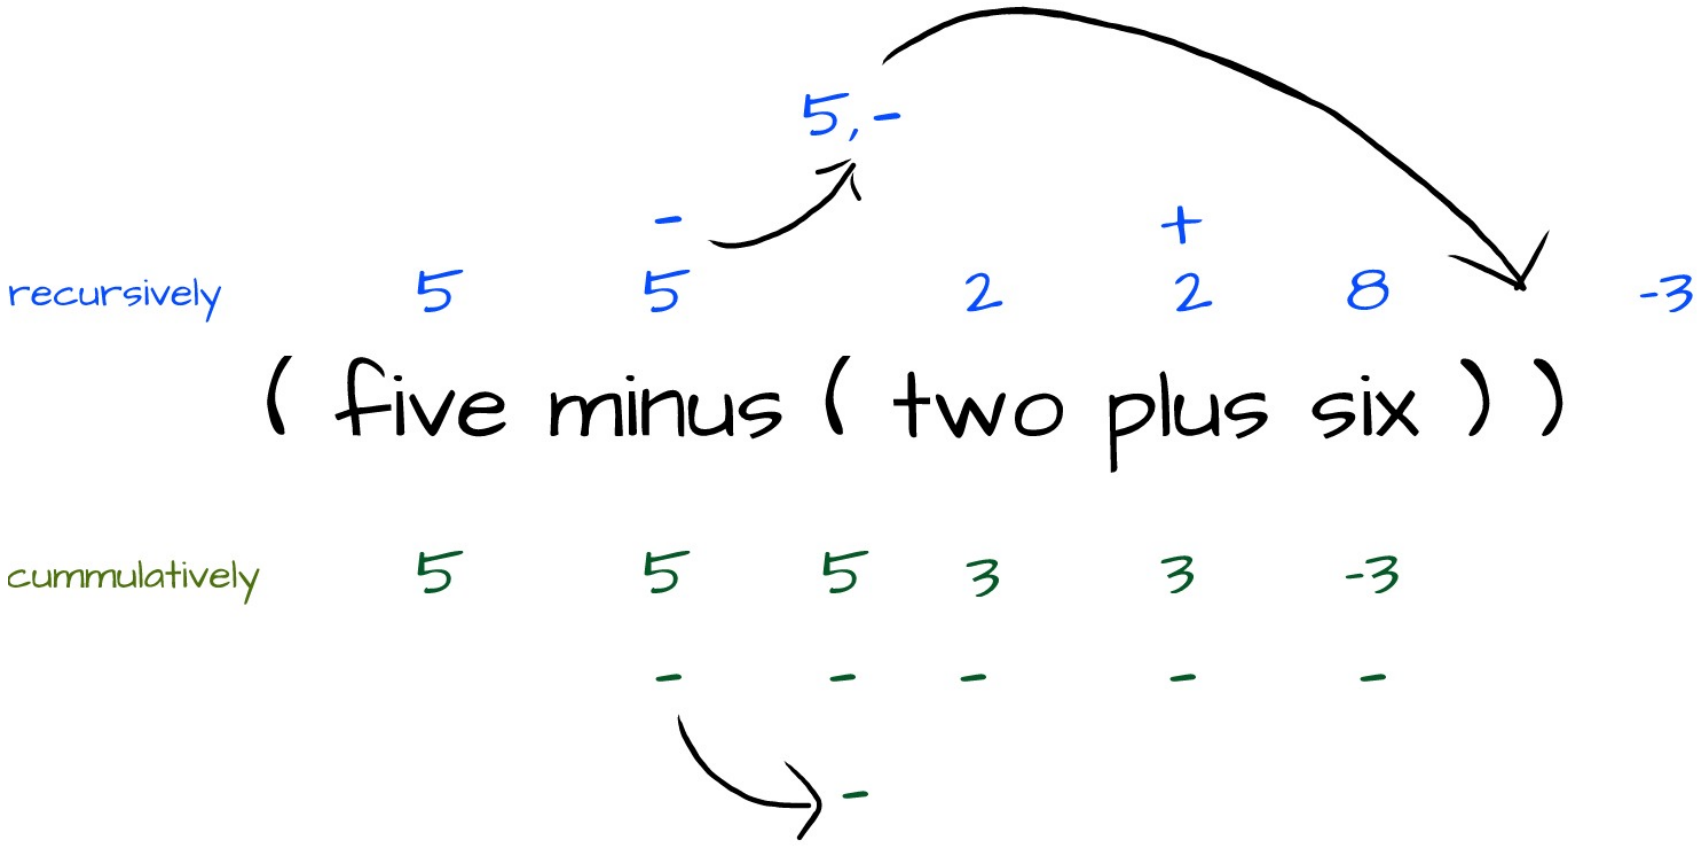

Symbolic solutions

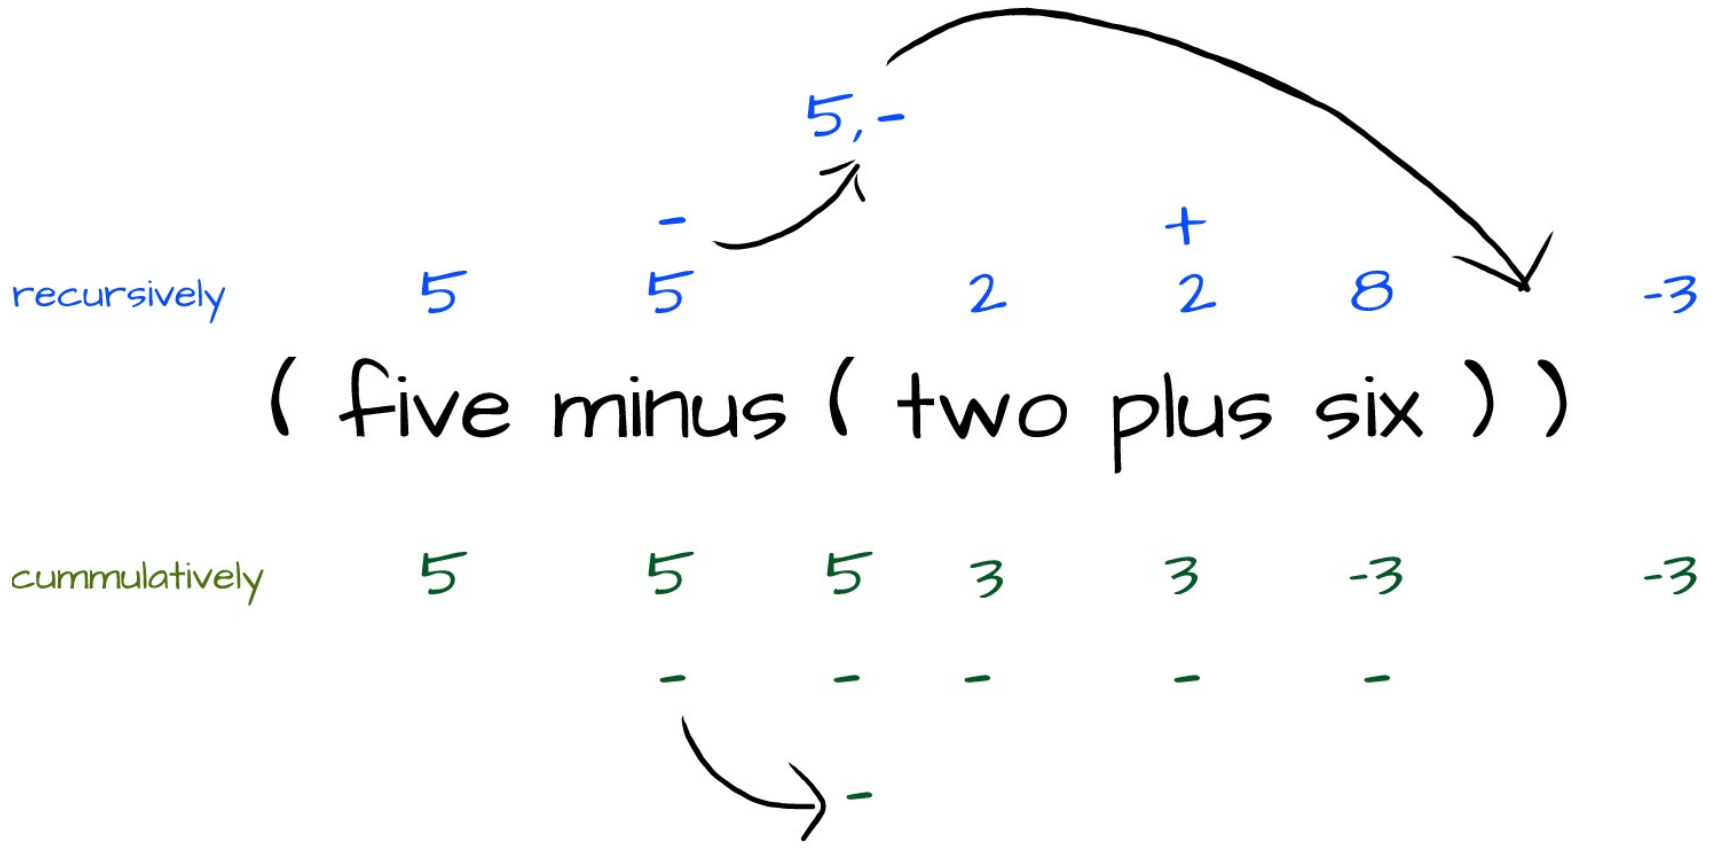

How do we study the network?

Plotting activation values

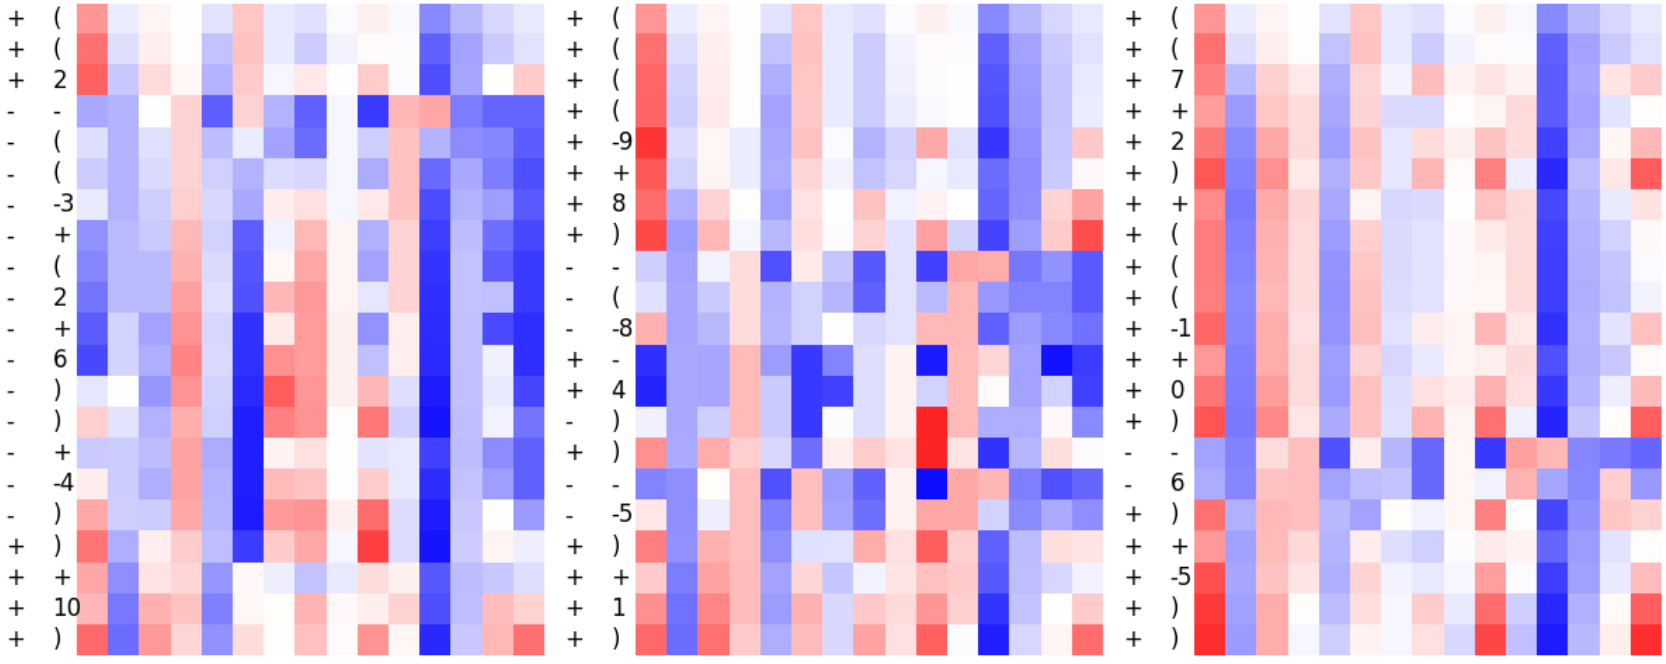

How do we study the network?

Diagnostic Classification

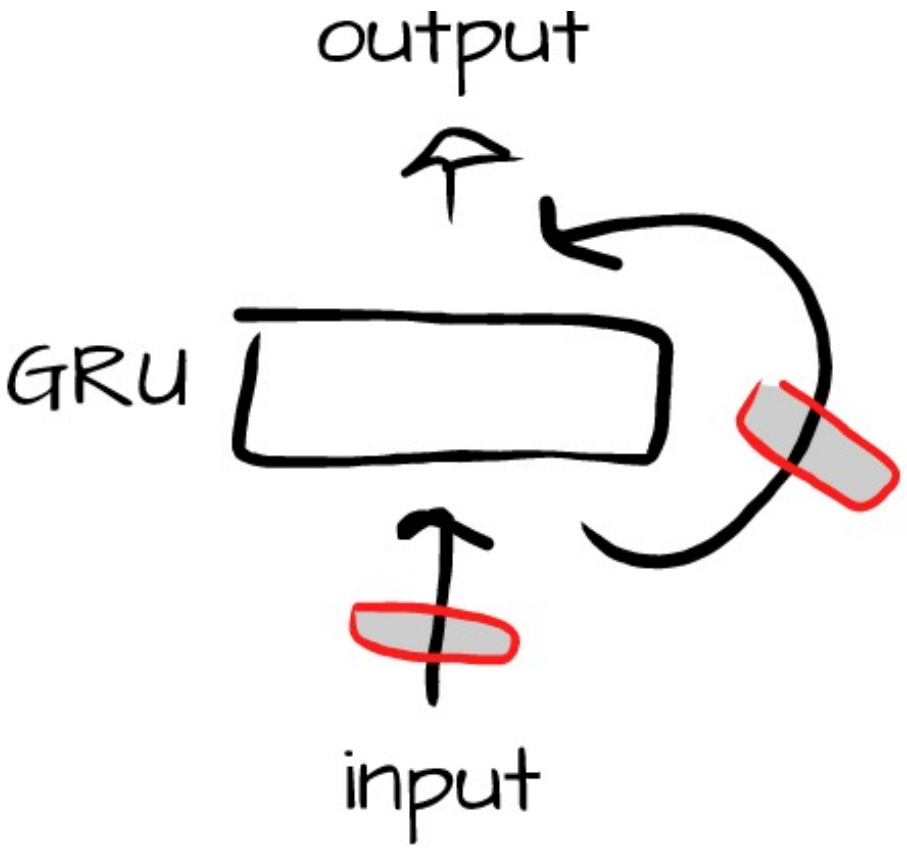

How do we study the network?

Diagnostic Classification

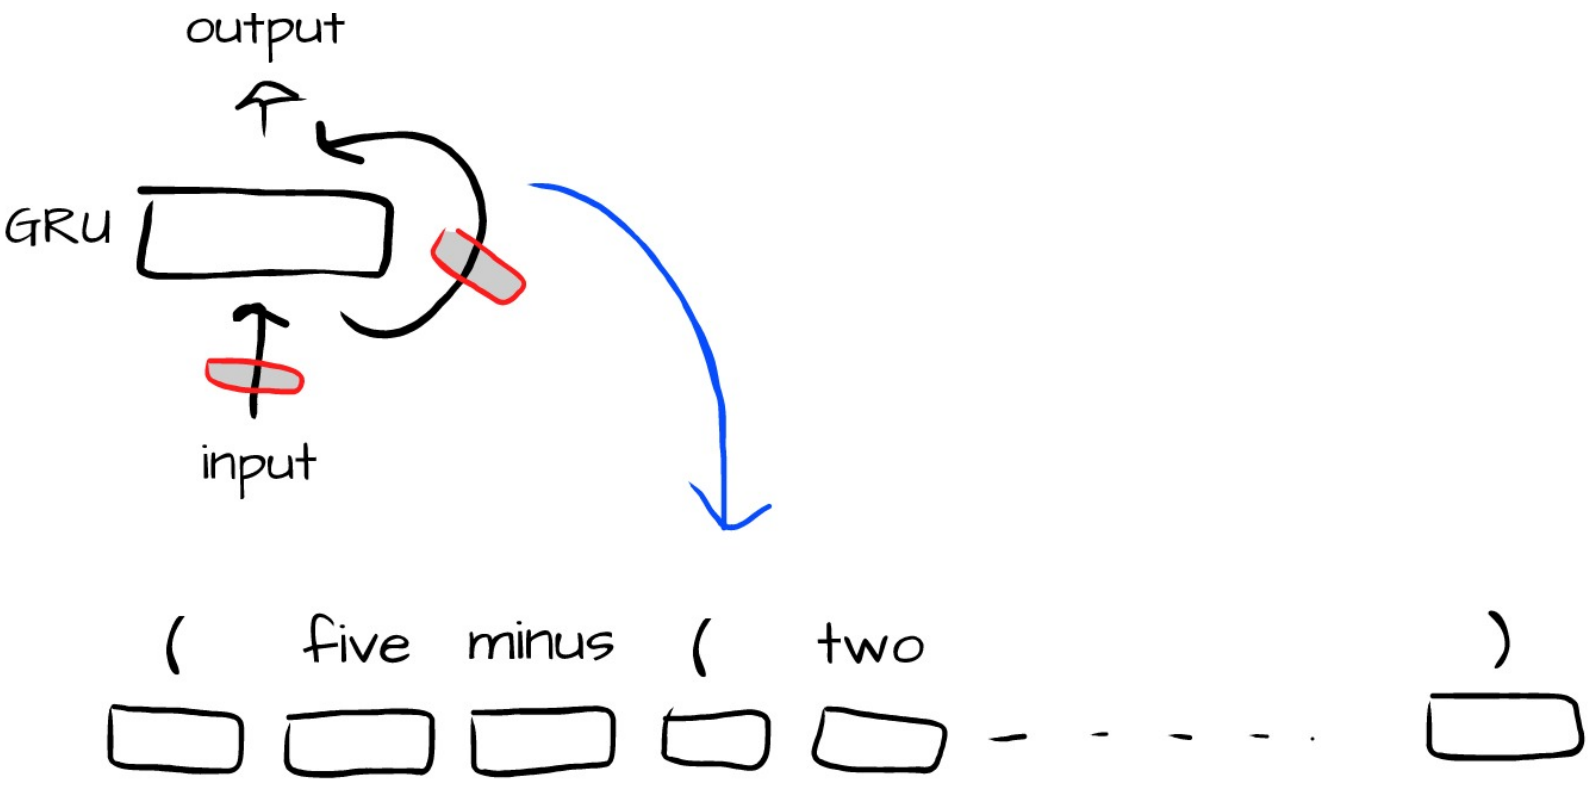

How do we study the network?

Diagnostic Classification

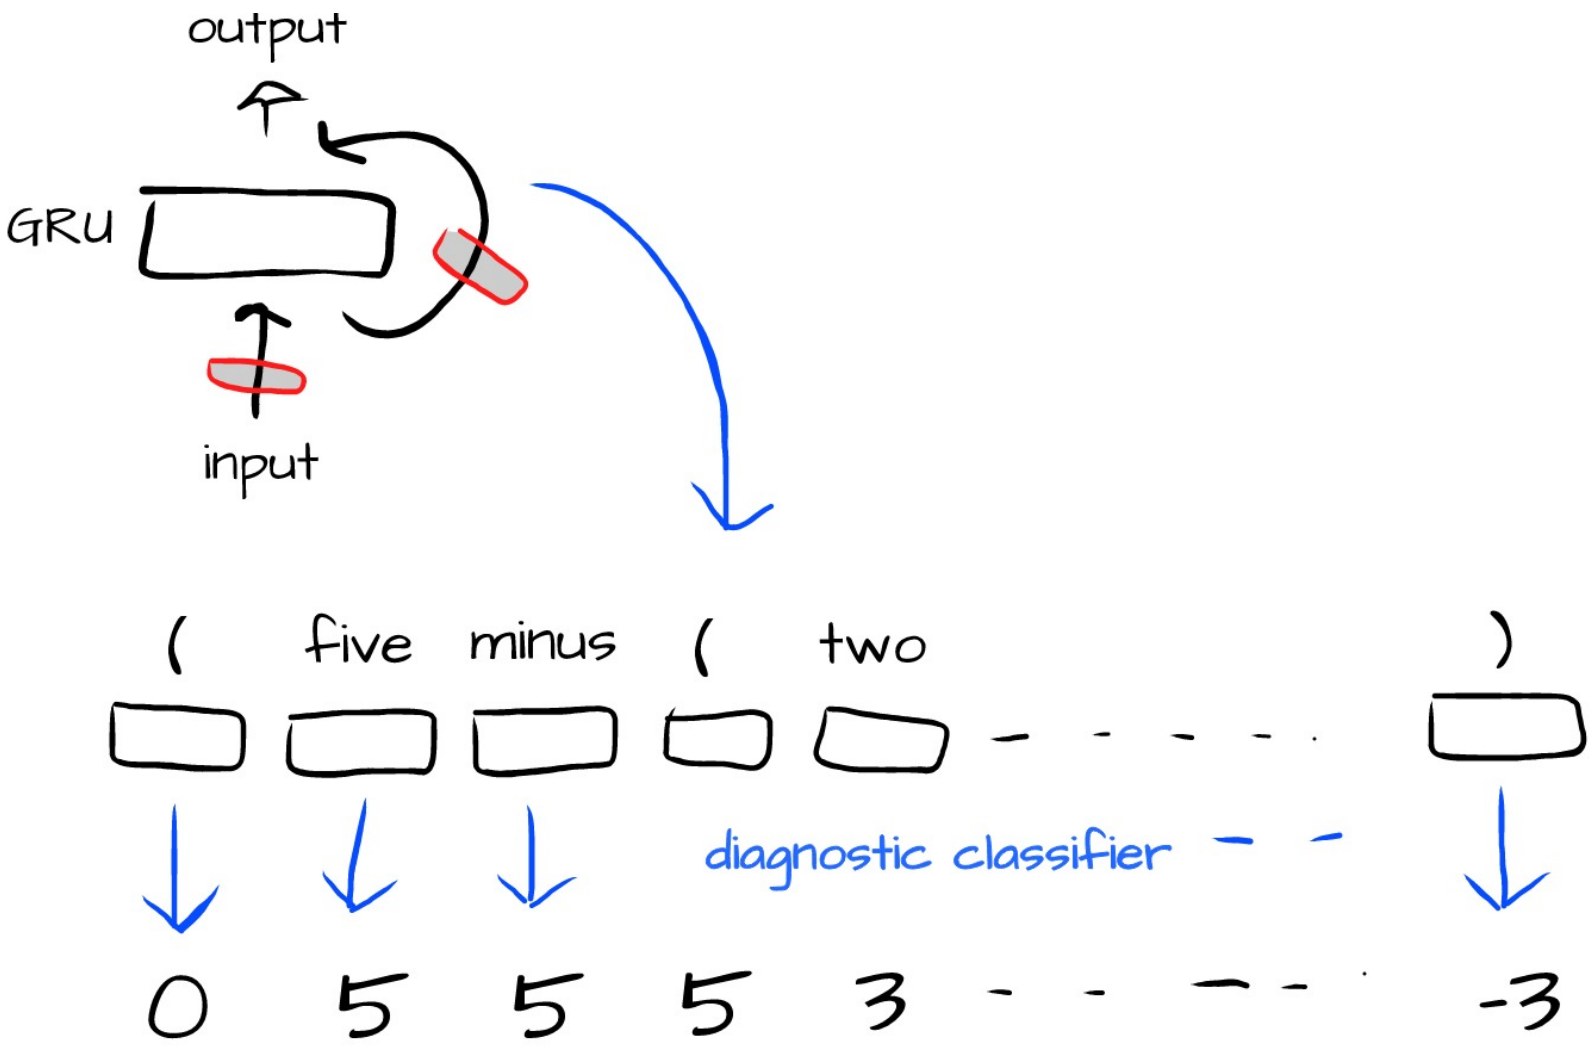

How do we study the network?

Diagnostic Classification

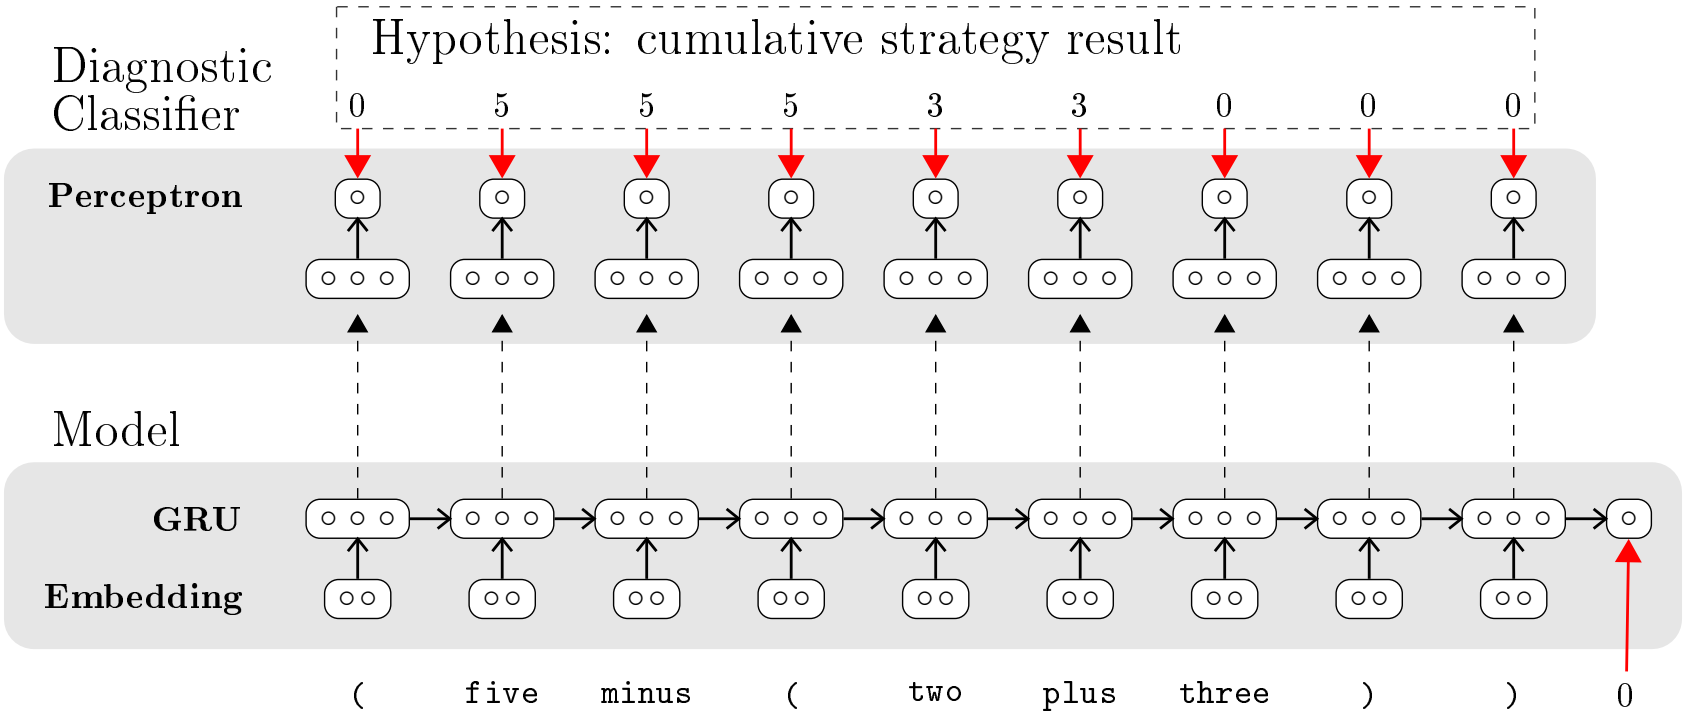

Hypotheses

minus_scope3+														1	1				1	1	1	1											
minus_scope2+								1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1						1	1		
minus_scope1+					1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1			1	1	1	1	1	
close_minus_scope1+	0	0	0	0	1	1	1	2	3	3	3	4	4	4	4	3	2	2	3	3	3	3	2	1	0	0	0	1	1	1	1	0	0
	(	**(**	**-2**	**-**	**(**	**6**	**-**	**(**	**(**	**8**	**+**	**(**	**-3**	**-**	**10**	**)**	**)**	**-**	**(**	**-2**	**-**	**10**	**)**	**)**	**)**	**)**	**-**	**(**	**1**	**-**	**-8**	**)**	**)**
mode	+	+	+	−	−	−	+	+	+	+	+	+	+	−	−	+	+	−	−	−	+	+	−	+	−	+	−	−	−	+	+	−	+
switch_mode				1			1							1		1		1			1		1	1	1	1	1			1		1	1

Hypotheses

minus_scope3+													1	1				1	1	1	1												
minus_scope2+								1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1						1	1		
minus_scope1+					1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1			1	1	1	1	1	
close_minus_scope1+	0	0	0	0	1	1	1	2	3	3	3	4	4	4	4	3	2	2	3	3	3	3	2	1	0	0	0	1	1	1	1	0	0
	(	**(**	**-2**	**-**	**(**	**6**	**-**	**(**	**(**	**8**	**+**	**(**	**-3**	**-**	**10**	**)**	**)**	**-**	**(**	**-2**	**-**	**10**	**)**	**)**	**)**	**)**	**-**	**(**	**1**	**-**	**-8**	**)**	**)**
mode	+	+	+	−	−	−	+	+	+	+	+	+	+	−	−	+	+	−	−	−	+	+	−	+	−	+	−	−	−	+	+	−	+
switch_mode				1			1							1		1		1			1		1	1	1	1	1			1		1	1

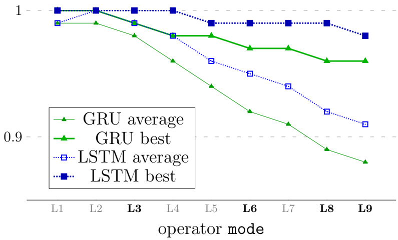

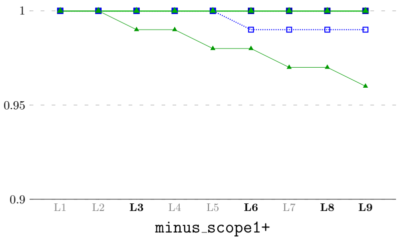

How about the gates?

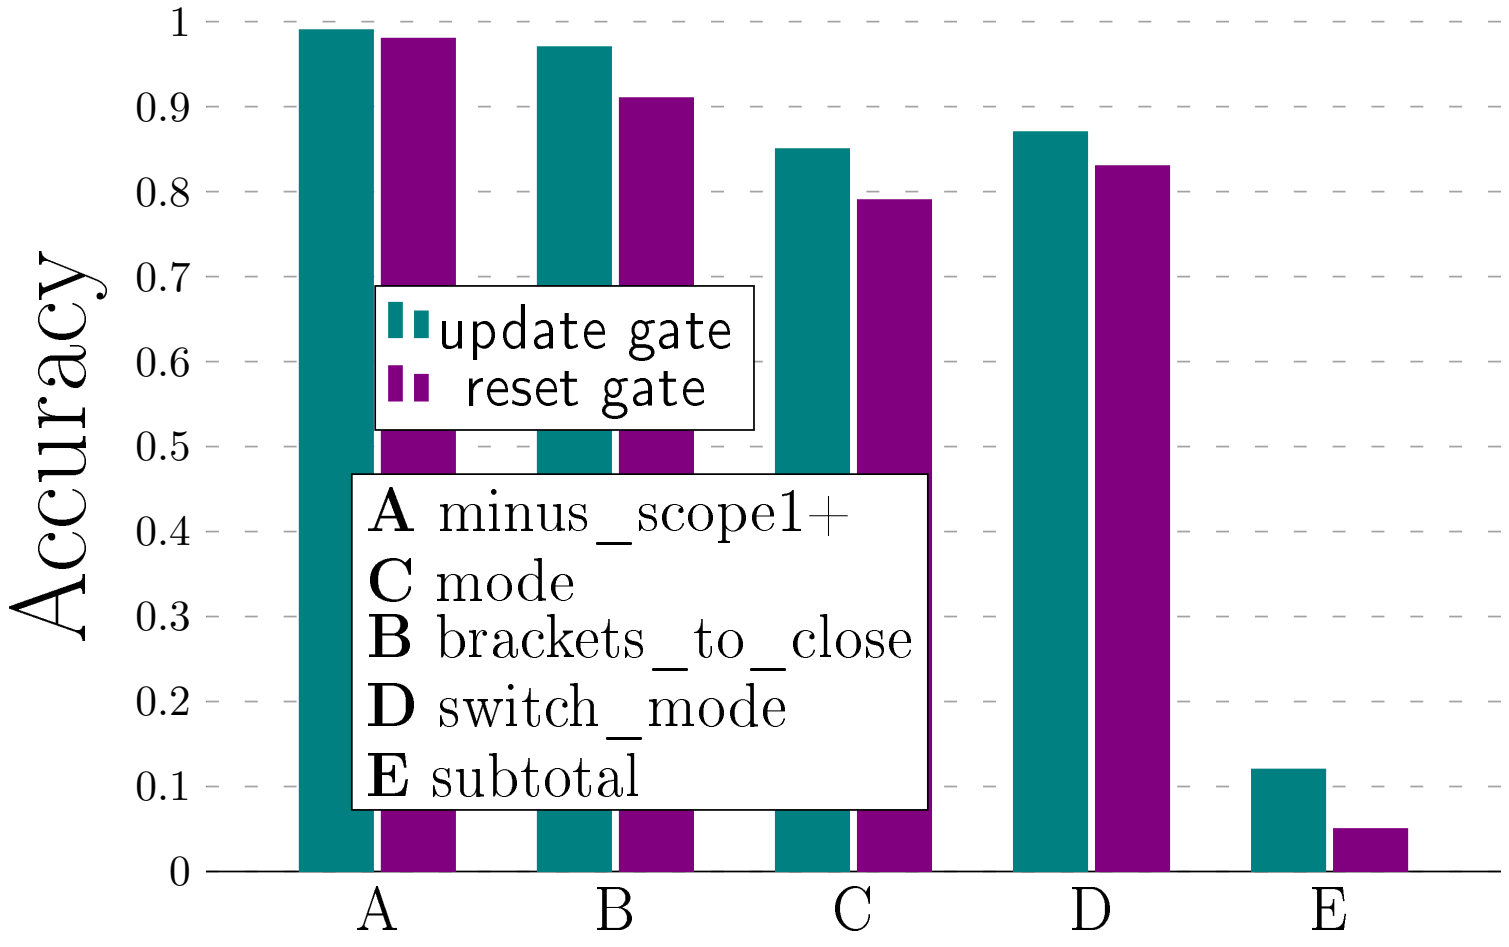

Using diagnostic classifier weights

What happens where?

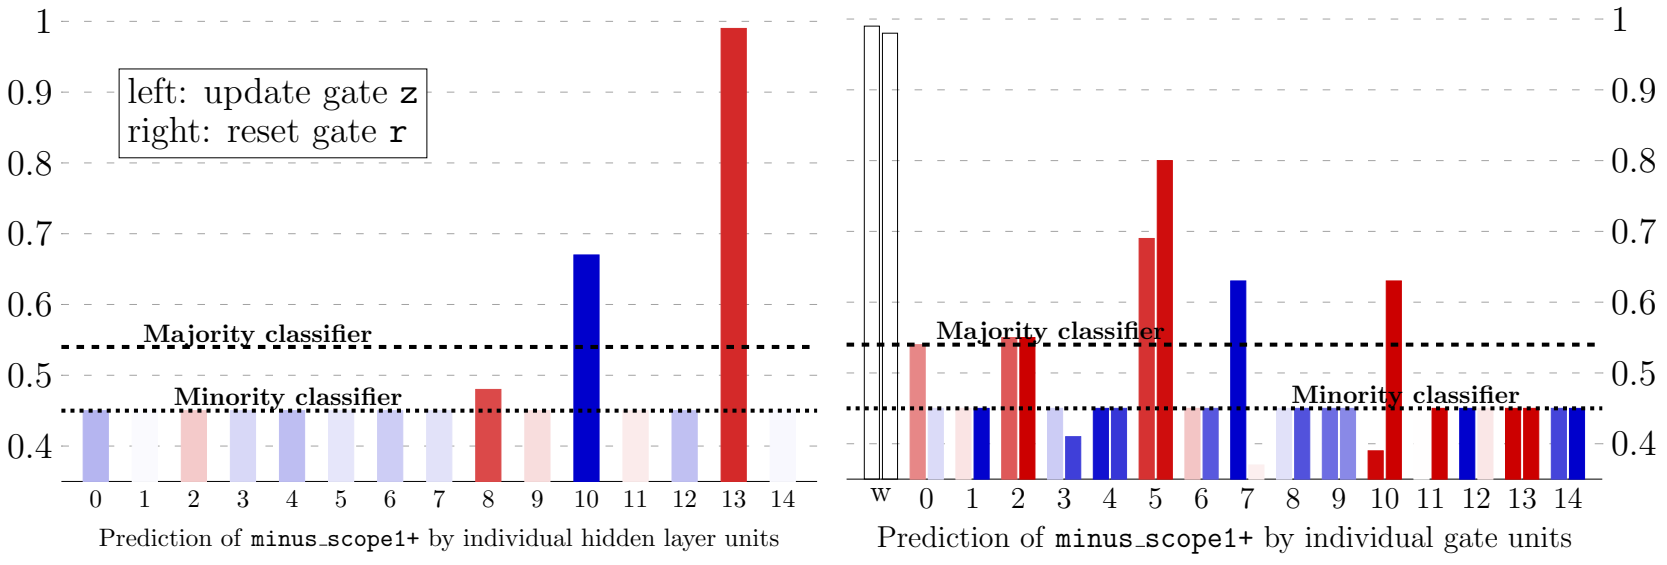

Symbolic Guidance

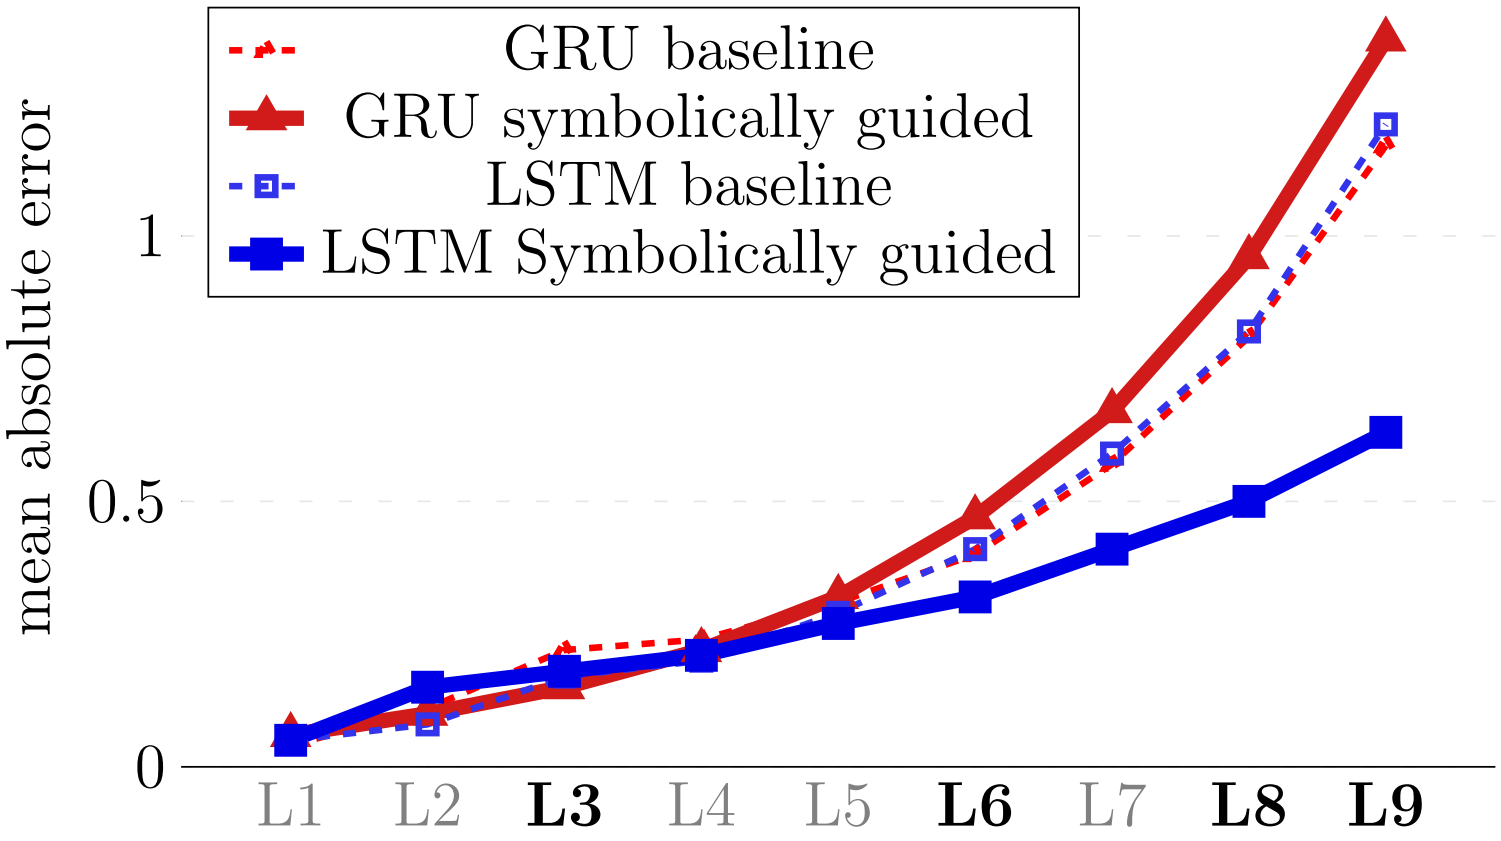

Now what?

1. Understanding what learning biases we need?
2. Injecting symbolic knowledge in neural networks?
3. Understanding if neural networks have linguistic knowledge?
4. ...