Attention and Augmented Neural Networks

Korbinian Riedhammer

Based on https://distill.pub/2016/augmented-rnns/

Recurrent Neural Networks

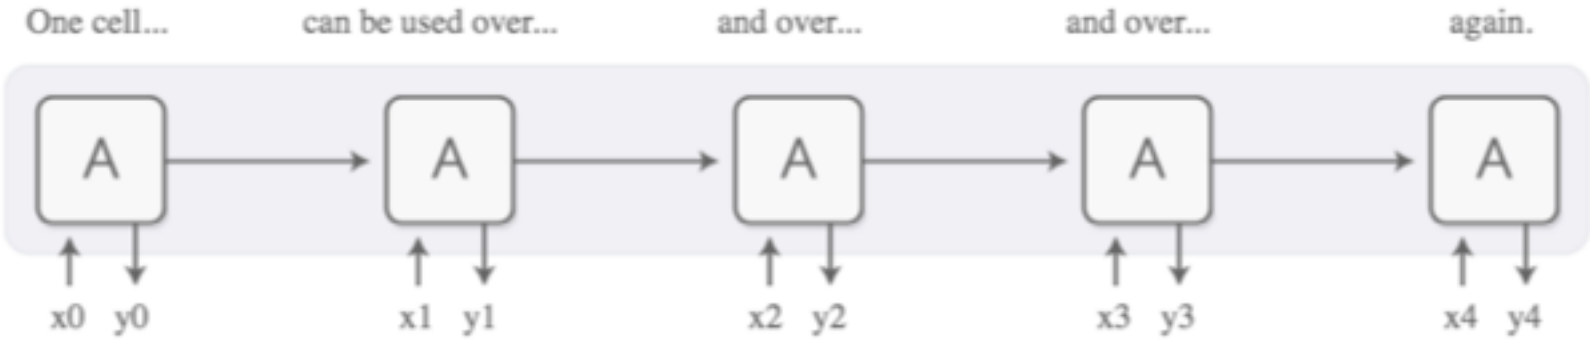

Variants

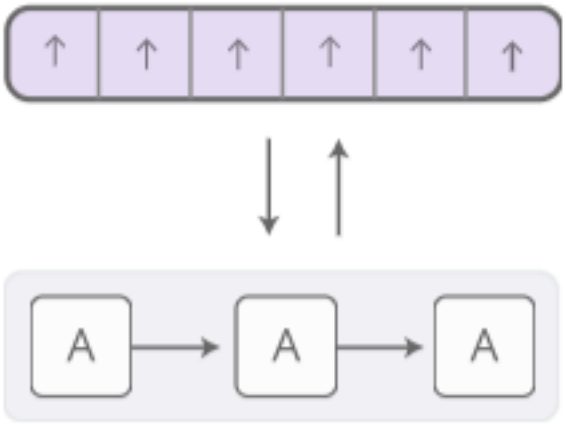

Neural Turing Machines
have external memory that they can read and write to.

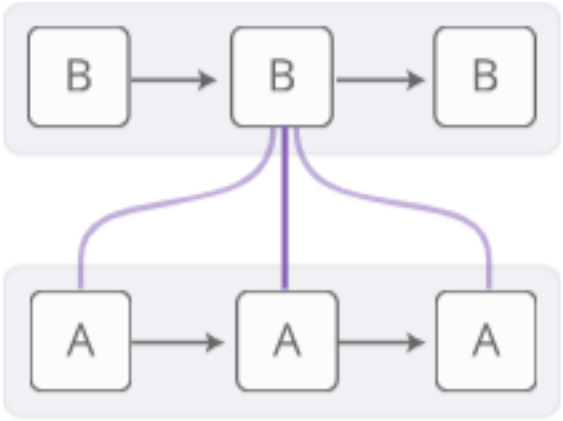

Attentional Interfaces
allow RNNs to focus on parts of their input.

Variants

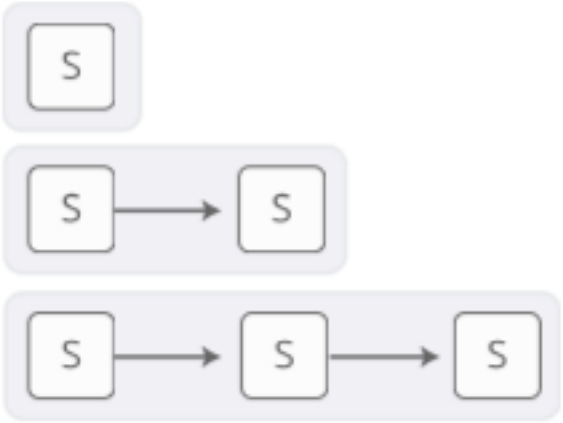

Adaptive Computation Time
allows for varying amounts of computation per step.

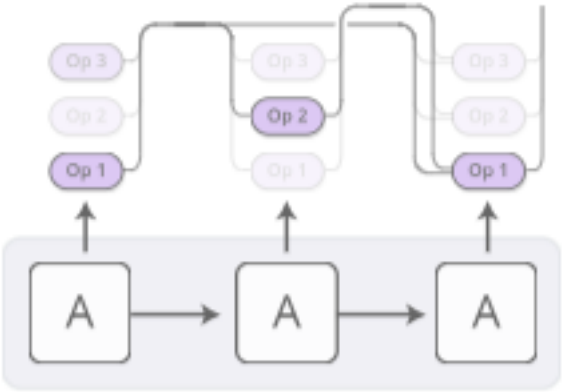

Neural Programmers
can call functions, building programs as they run.

Neural Turing Machines

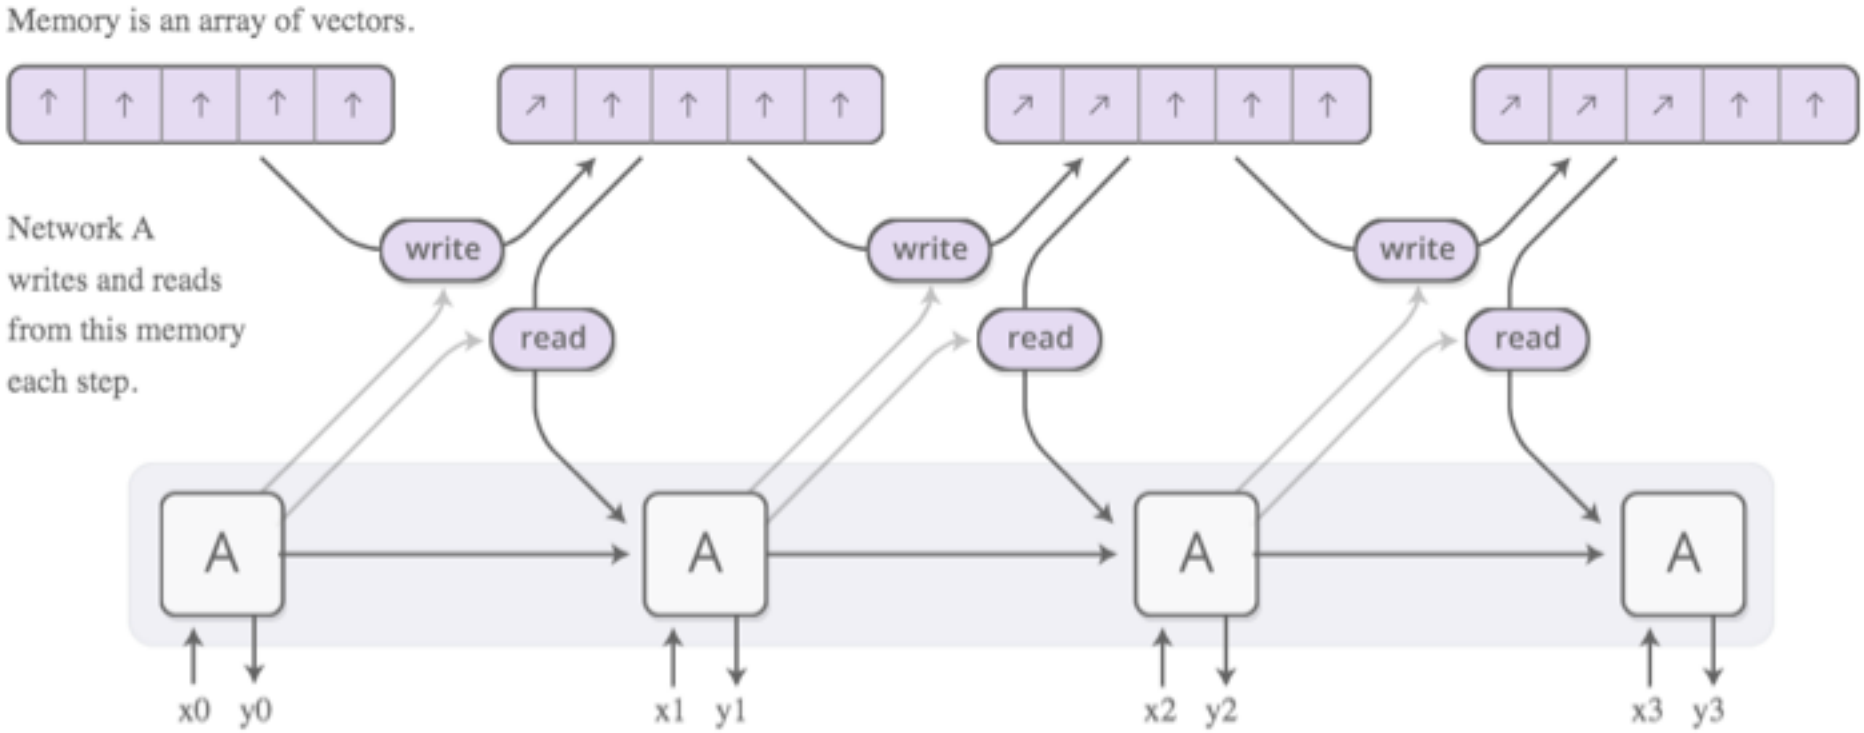

Neural Turing Machines

- Vectors are “natural language” of neural nets
- How to read/write from memory?
- How to differentiate?

In every step, read and write **everywhere!**

“Read” from Memory

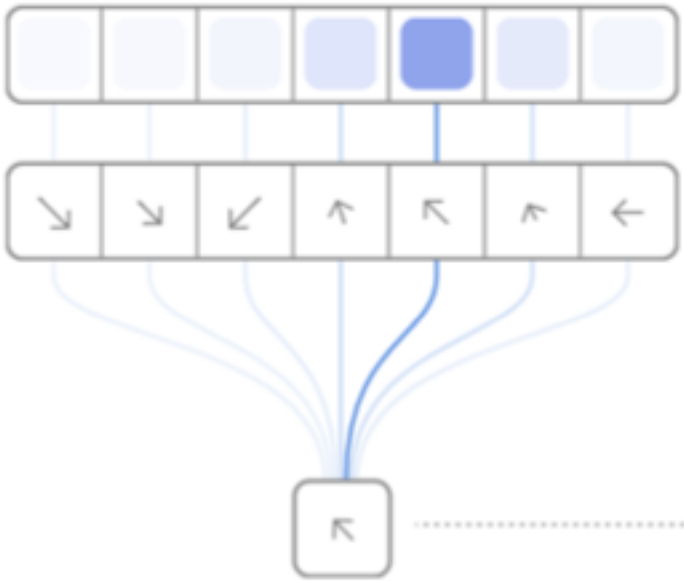

The RNN gives an attention distribution which describe how we spread out the amount we care about different memory positions.

The read result is a weighted sum.

$$r \leftarrow \sum_i a_i M_i$$

“Write” to Memory

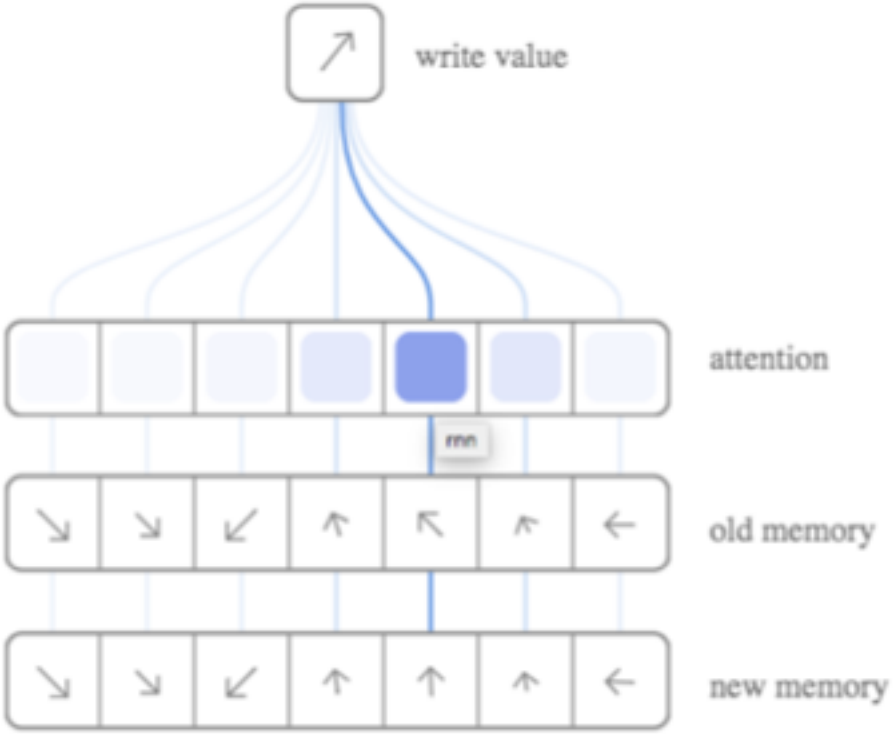

Instead of writing to one location, we write everywhere, just to different extents.

The RNN gives an attention distribution, describing how much we should change each memory position towards the write value.

$$M_i \leftarrow a_i w + (1-a_i) M_i$$

Content based and Location based Attention

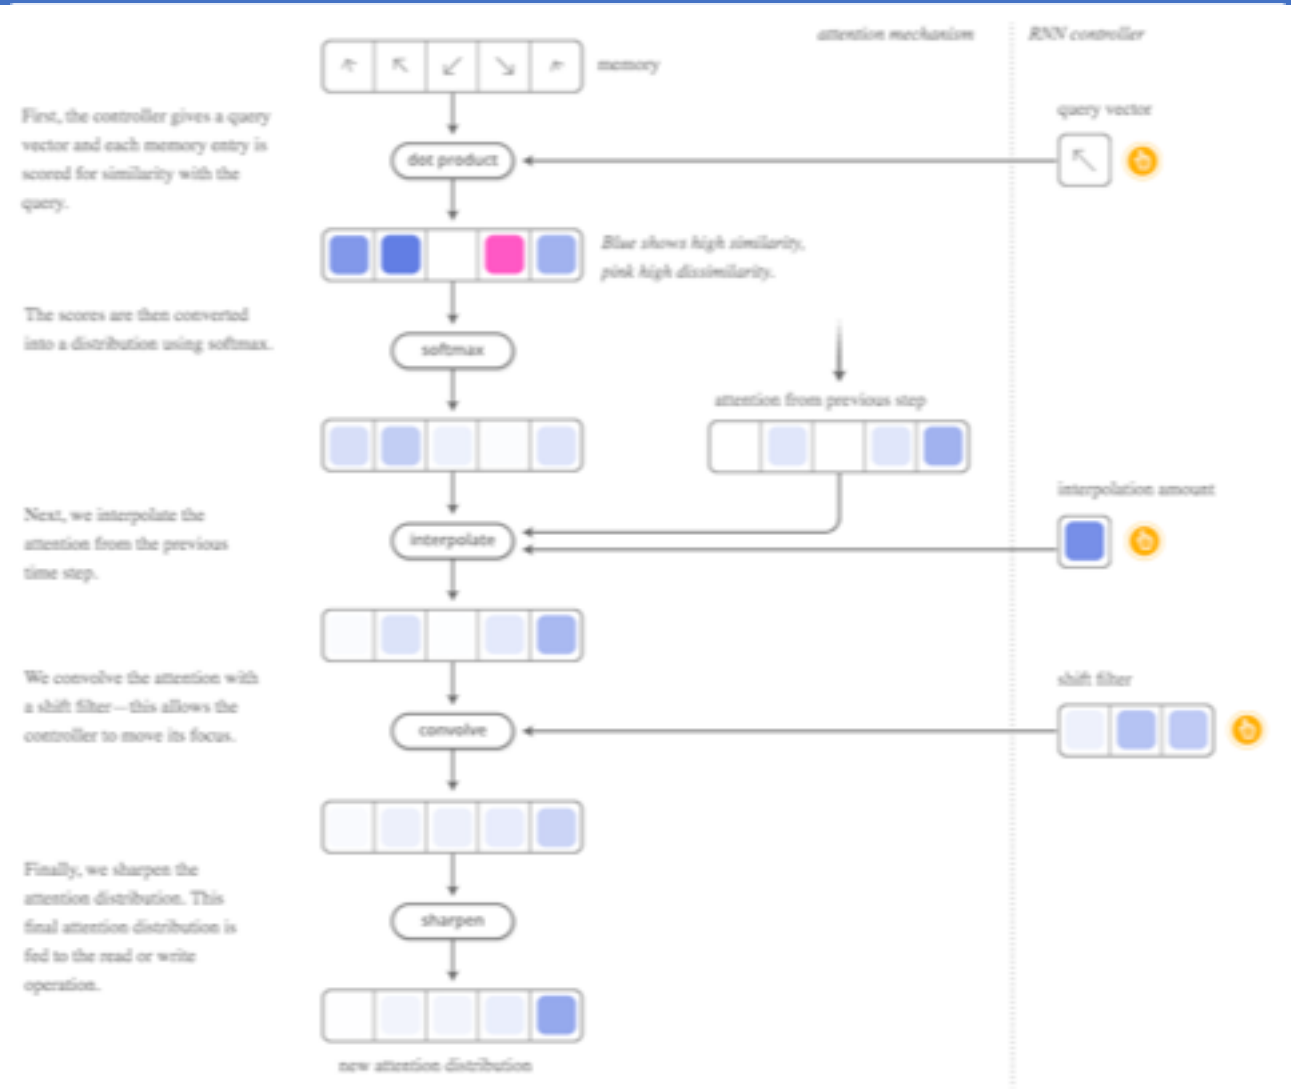

Visualization

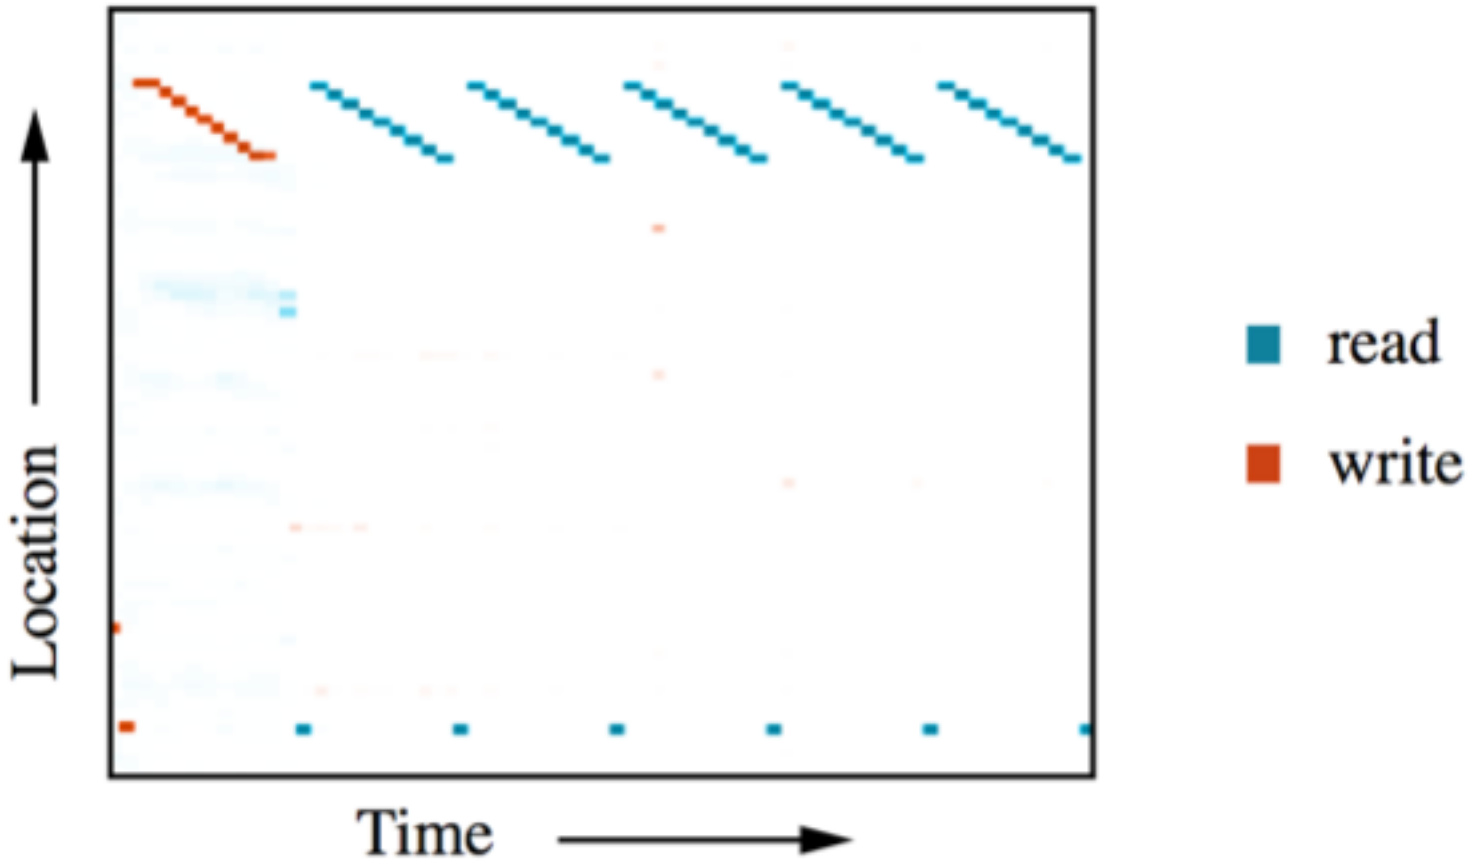

Attentional Interfaces

- Observation: In many complex human tasks (eg. translation, transcription, description, …), you pay attention to different aspects (in time and space)
- Model this “attention” in neural nets?
- How can we incorporate hidden state of previous time steps? …and be differentiable?

Similar to NTM, focus **everywhere**
but with different amount

Attentional Interfaces

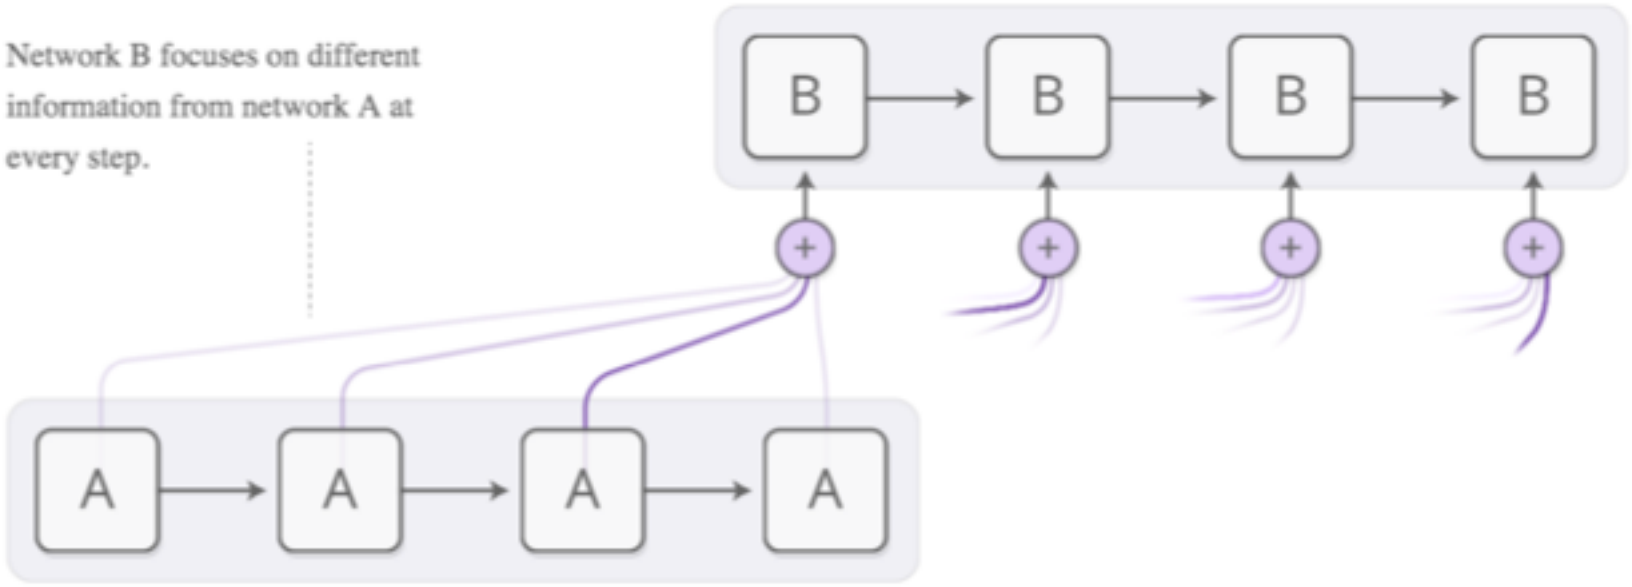

Content-Based Attention

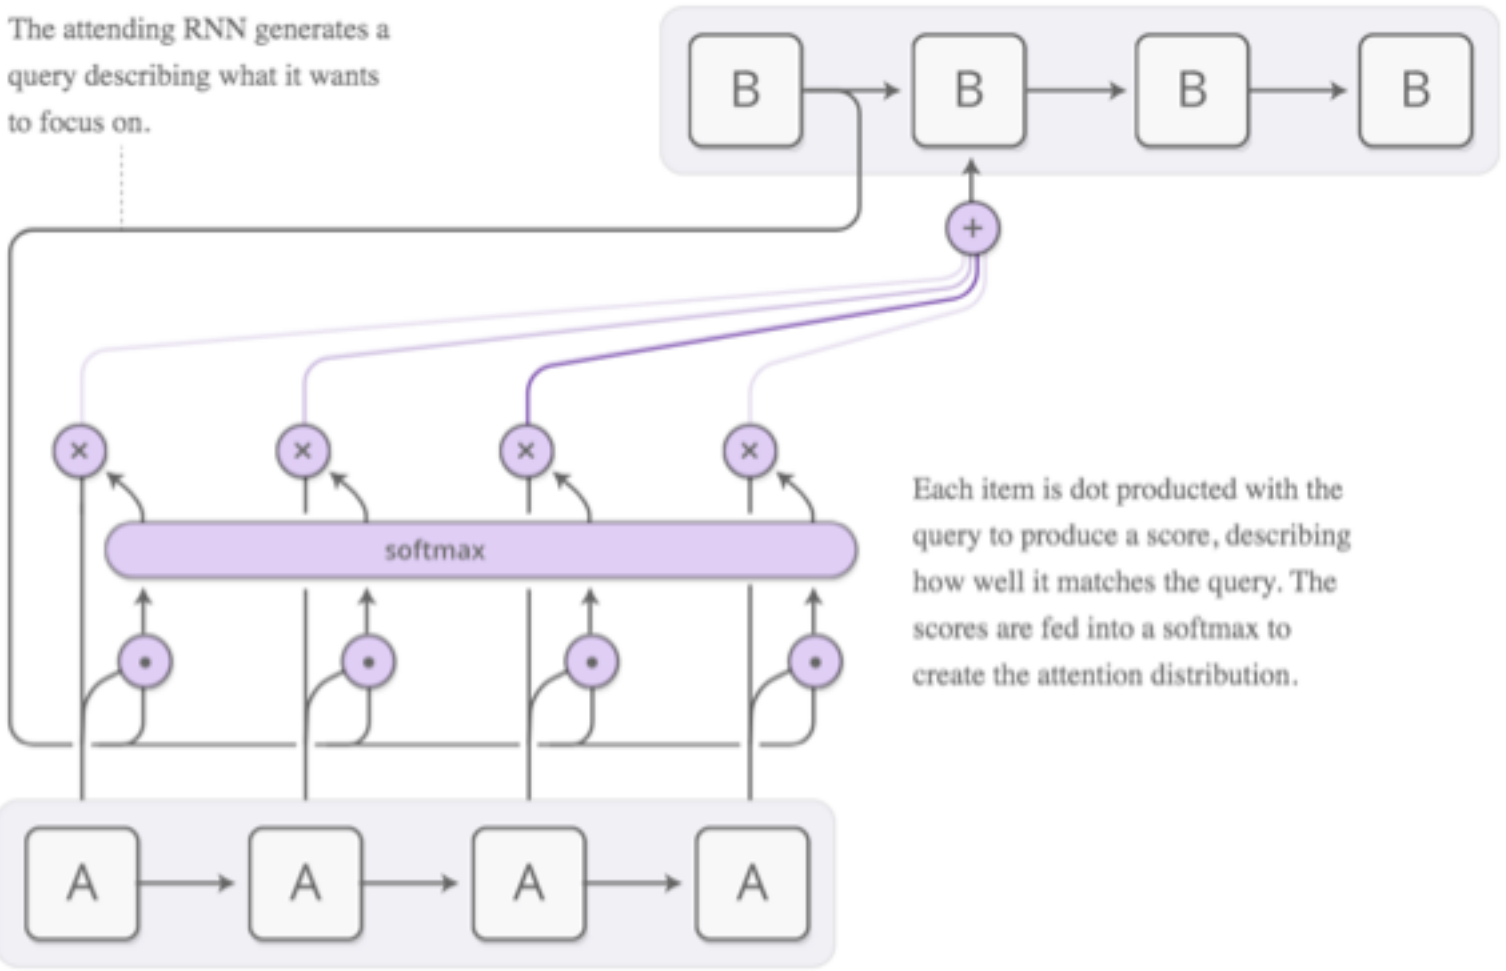

Example: direct dependency

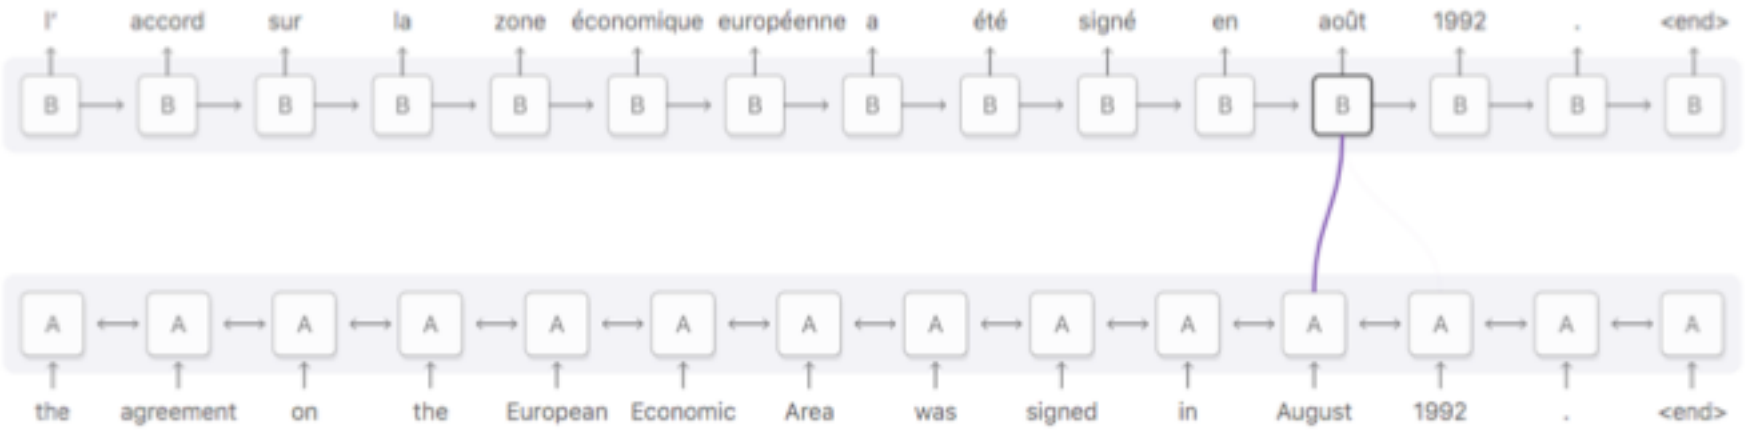

Diagram derived from Fig. 3 of Bahdanau, *et al.* 2014

Example: multi-dependency

Diagram derived from Fig. 3 of Bahdanau, *et al.* 2014

Attention in Speech

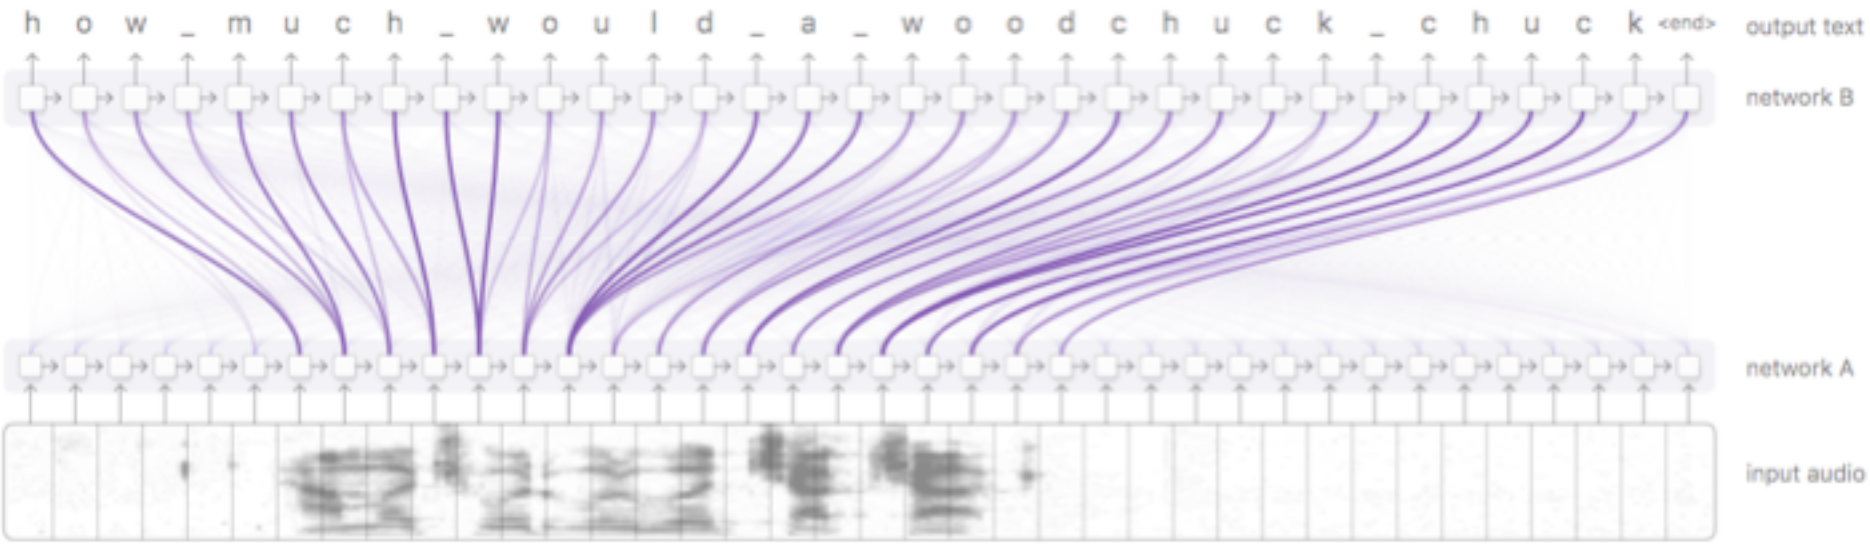

Figure derived from Chan, *et al.* 2015

Attention in Images

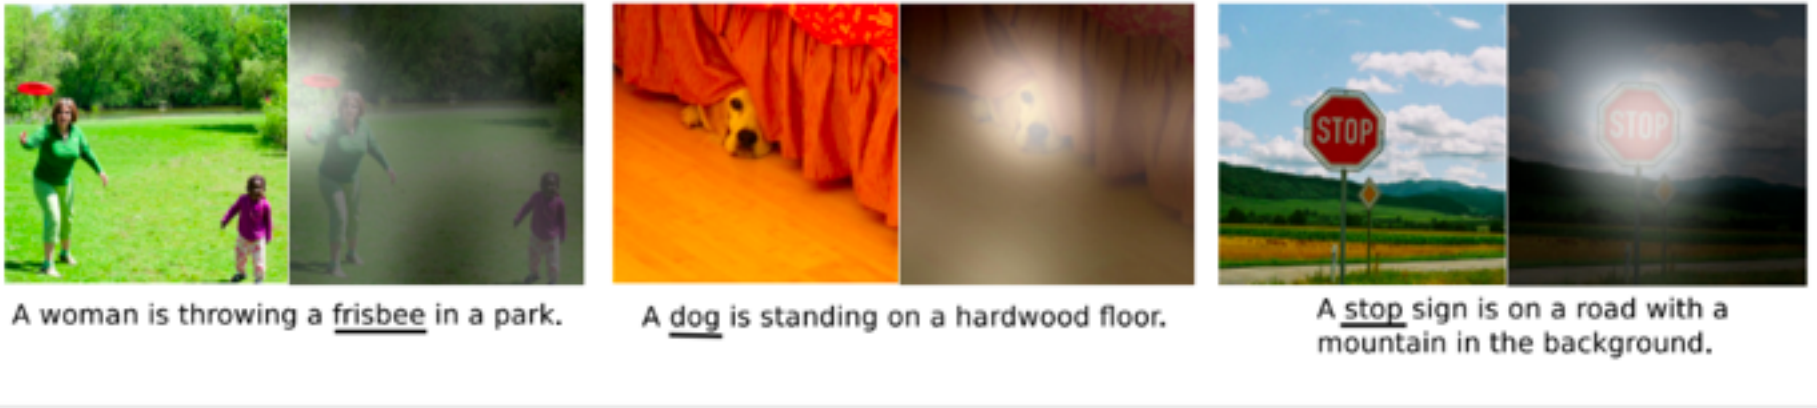

Figure from [3]

General Principle

- Everything in neural net needs to be differentiable (→ learning with backprop!)
- Model discrete selections (single outputs) as continuous selections (select all with different weight)
- Neural networks become “computational graphs”

Adaptive Computation Time

- Allow RNN to execute variable amounts of computation for each timestep?
- How many timesteps? …attention!

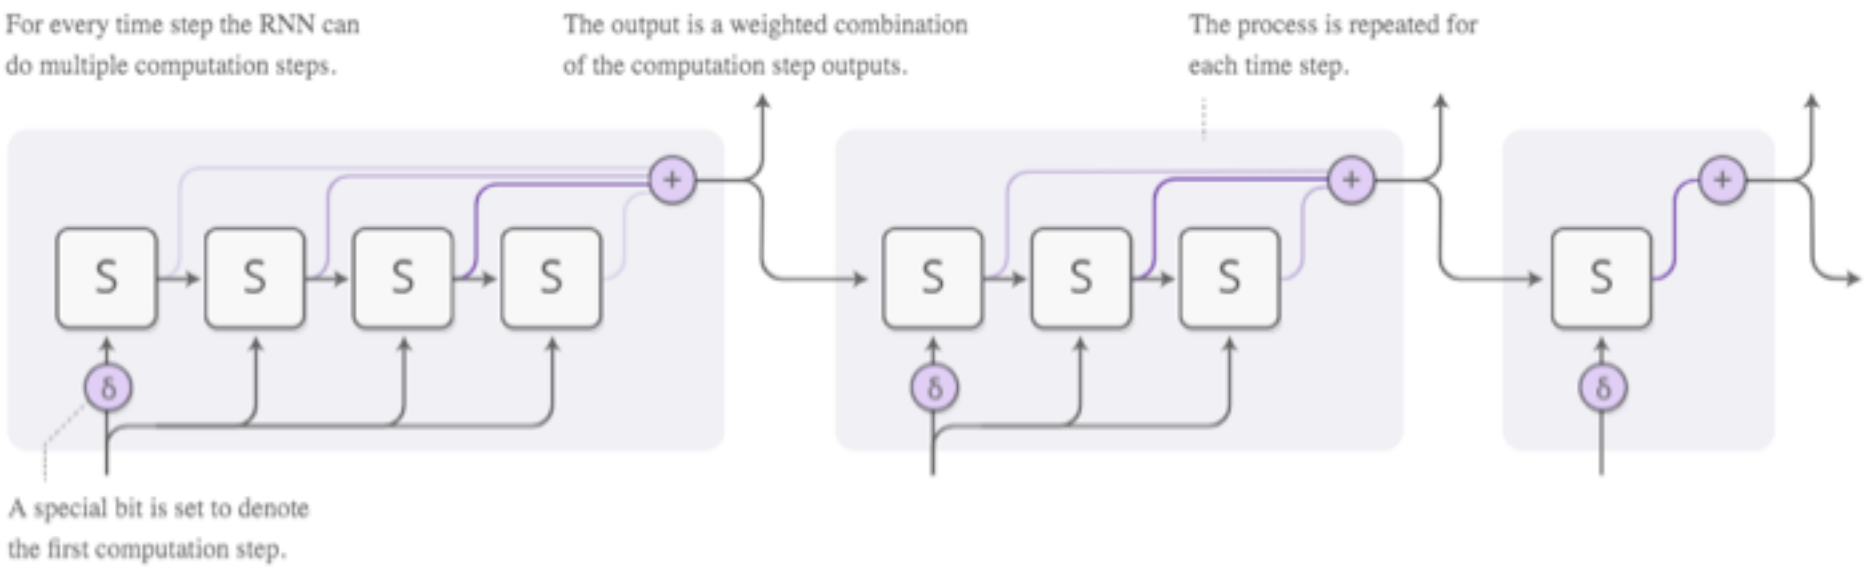

Adaptive Compute Time

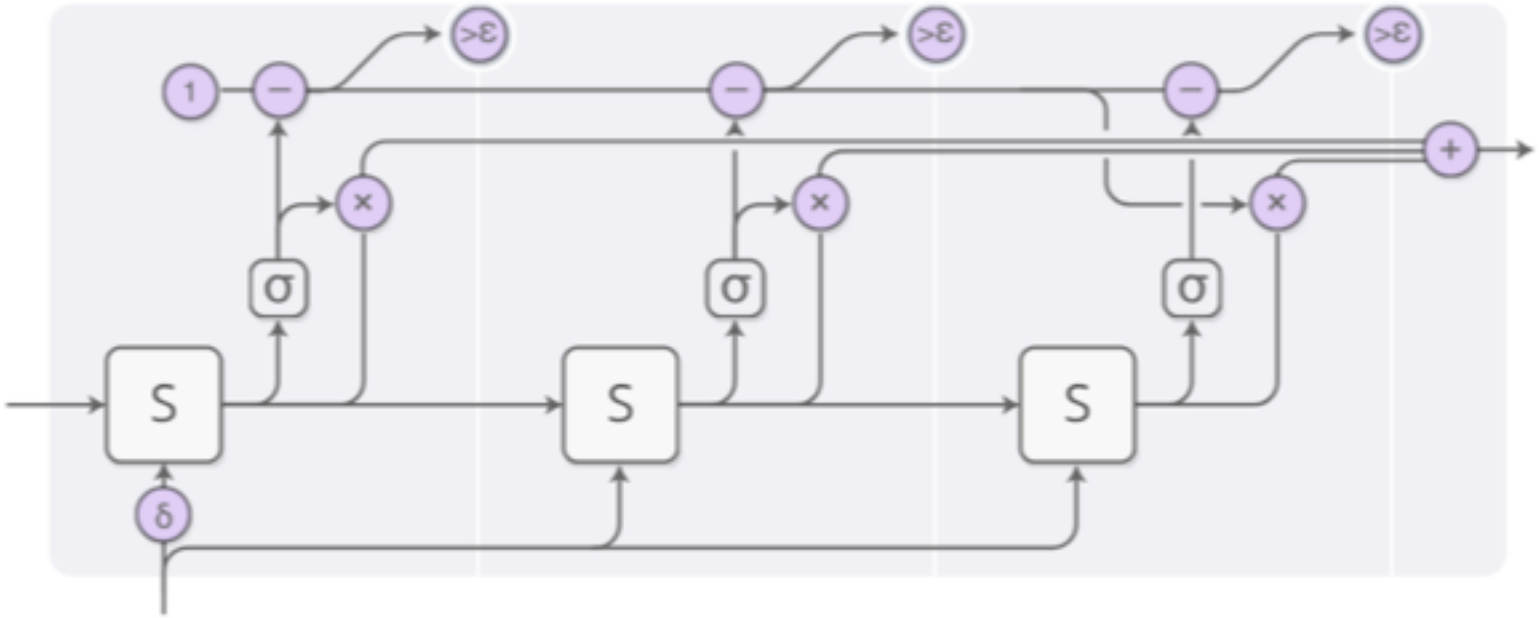

ACT in detail

Adaptive Compute Time

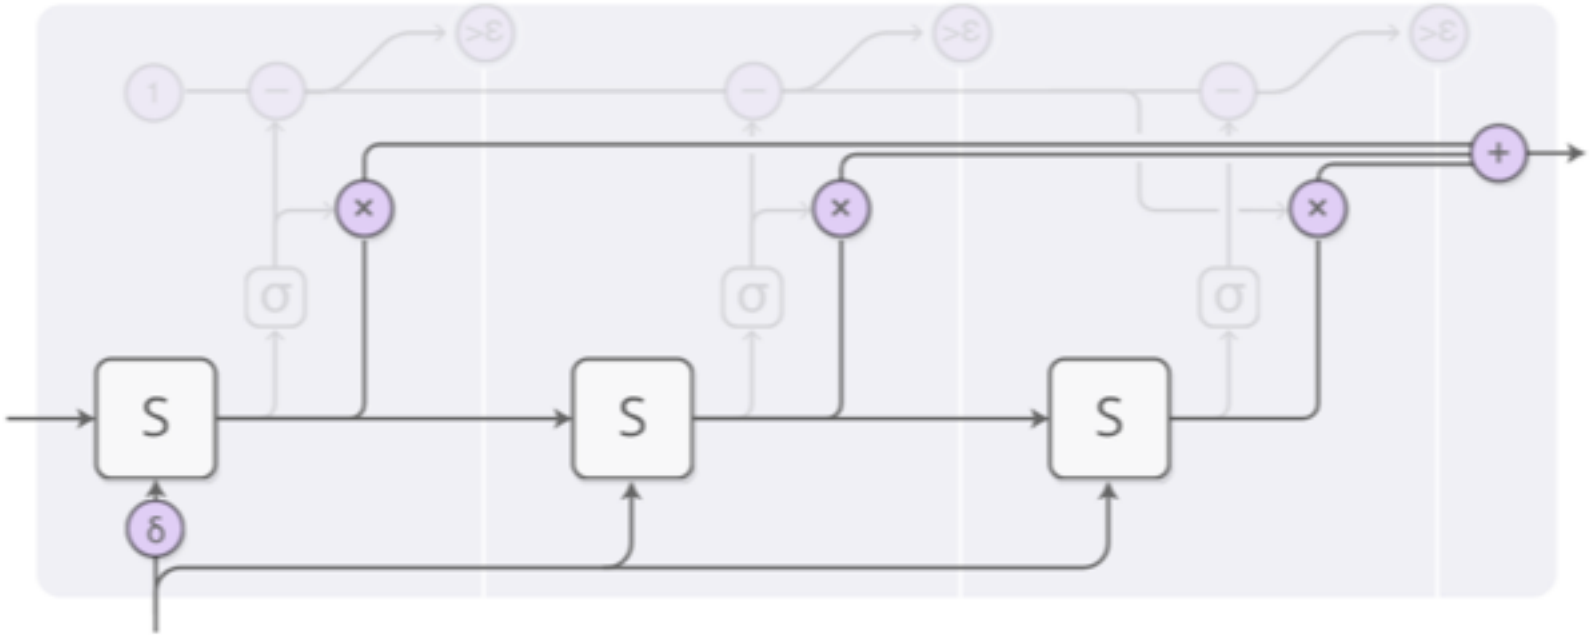

Output: weighted combination of states

Adaptive Compute Time

Individual weights determined by “halting neuron”
(sigmoid activation, read “likelihood to stop here”)

Adaptive Compute Time

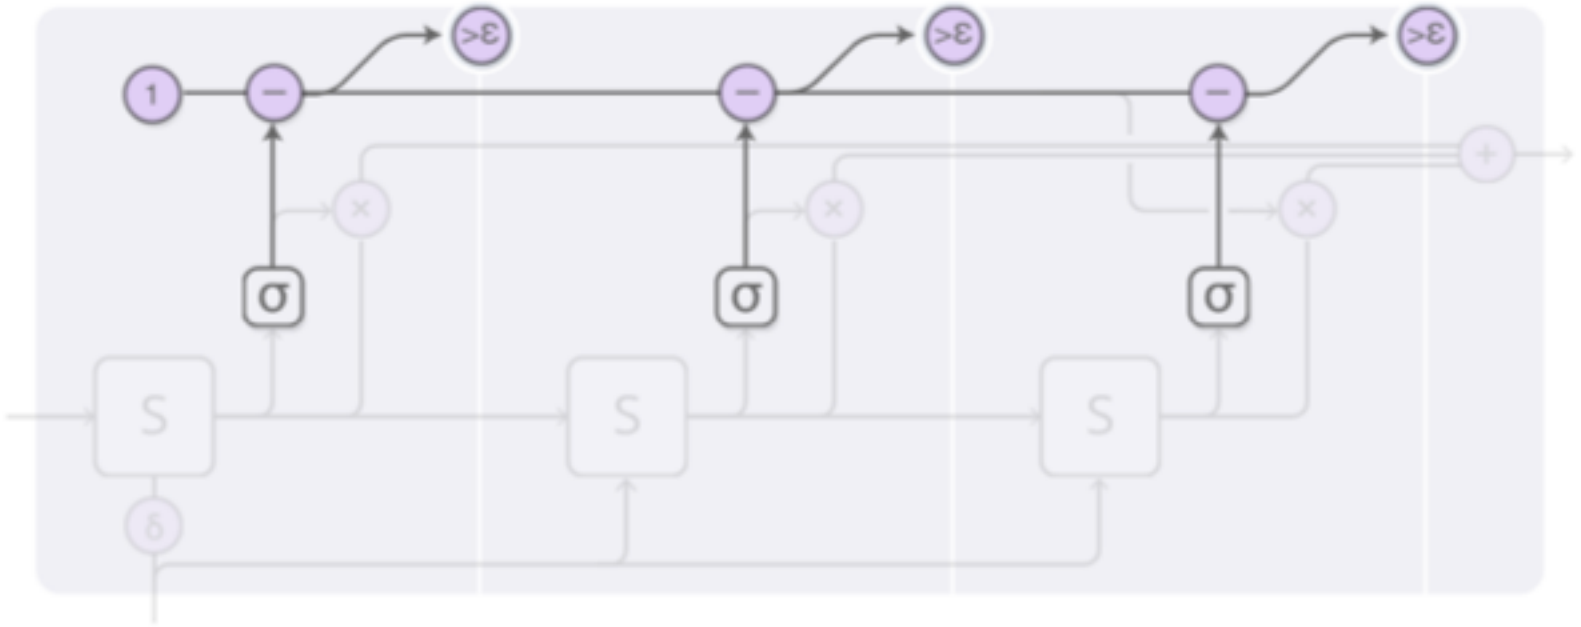

Make sure that weights sum up to 1!
Stop when no weight is left

Adaptive Compute Time

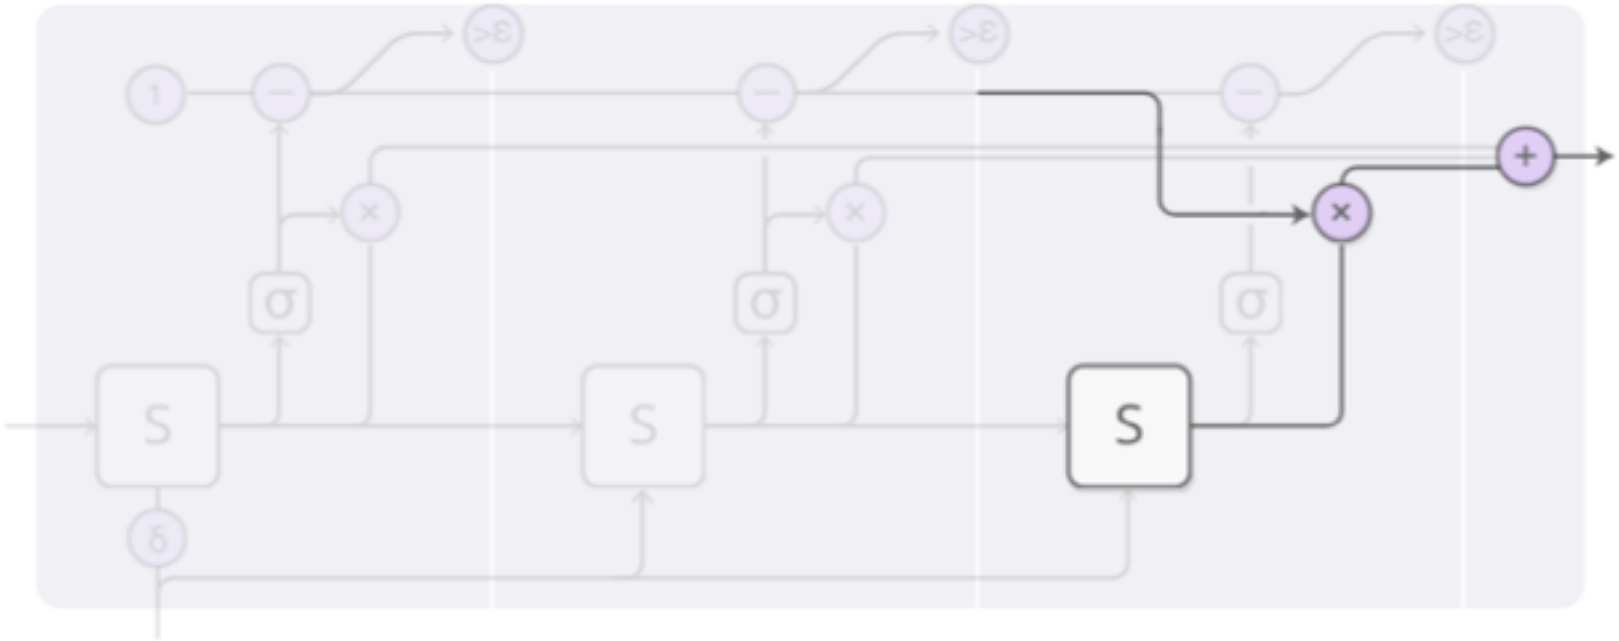

Add residual weight to output by forcing last state

Still not creepy enough?

There is even more!

Neural Programmer

- How about modeling actions/operations?
- Like arithmetic, loops, etc.?

Model as distribution of operations

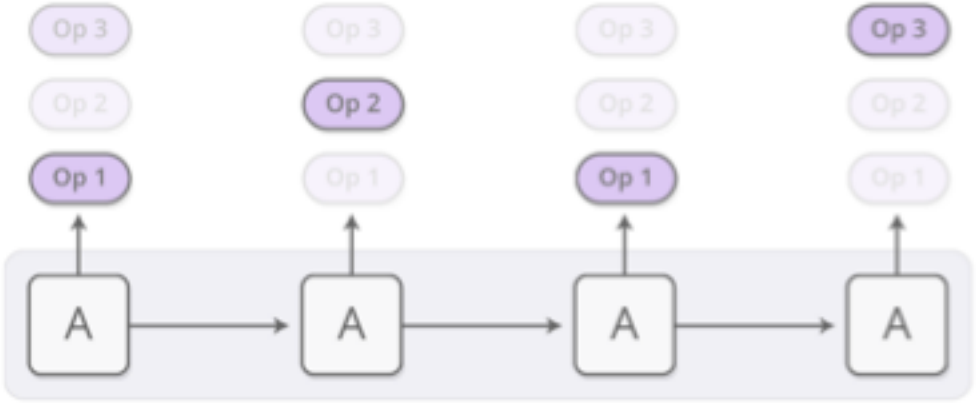

Neural Programmer

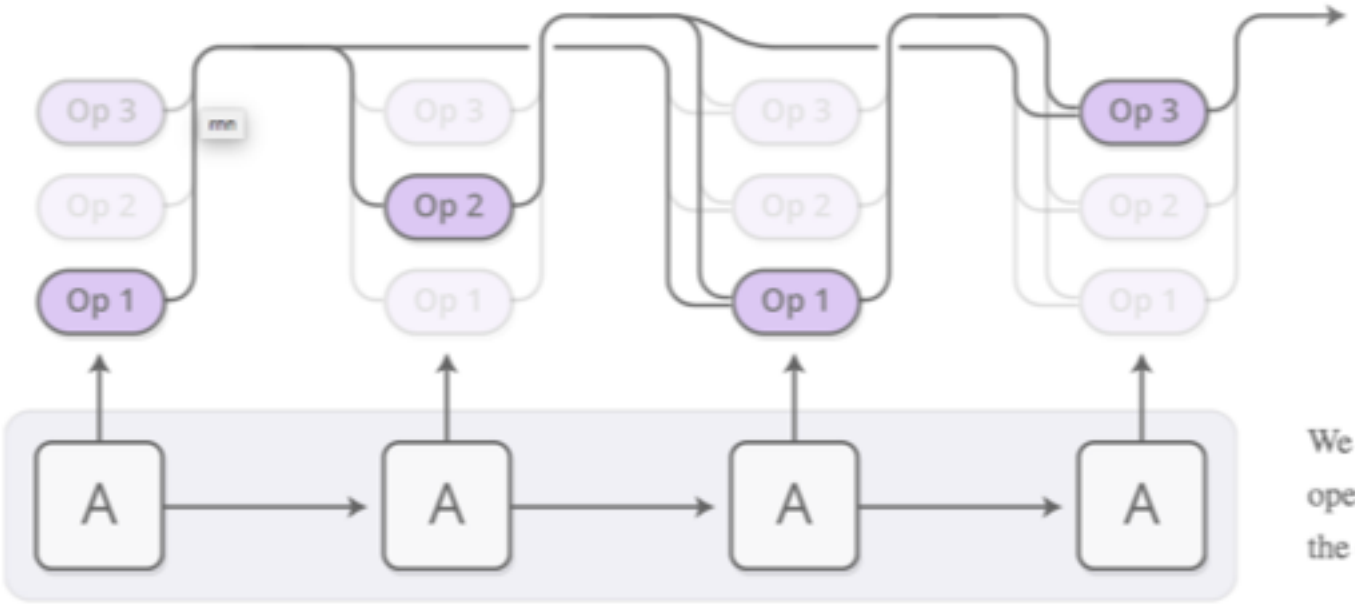

…and use attention to make it differentiable!