Learning to generalize by skill composition

Facebook AI Research

Marco Baroni
joint work with:
Tomas Mikolov, Rahma Chaabouni, Allan Jabri,
Germán Kruszewski, Armand Joulin, Klemen Simonic

Outline

- The fast learning challenge
- Compositional architectures
- Preliminary experiments
- The bigger picture: CommAI

Progress in machine learning

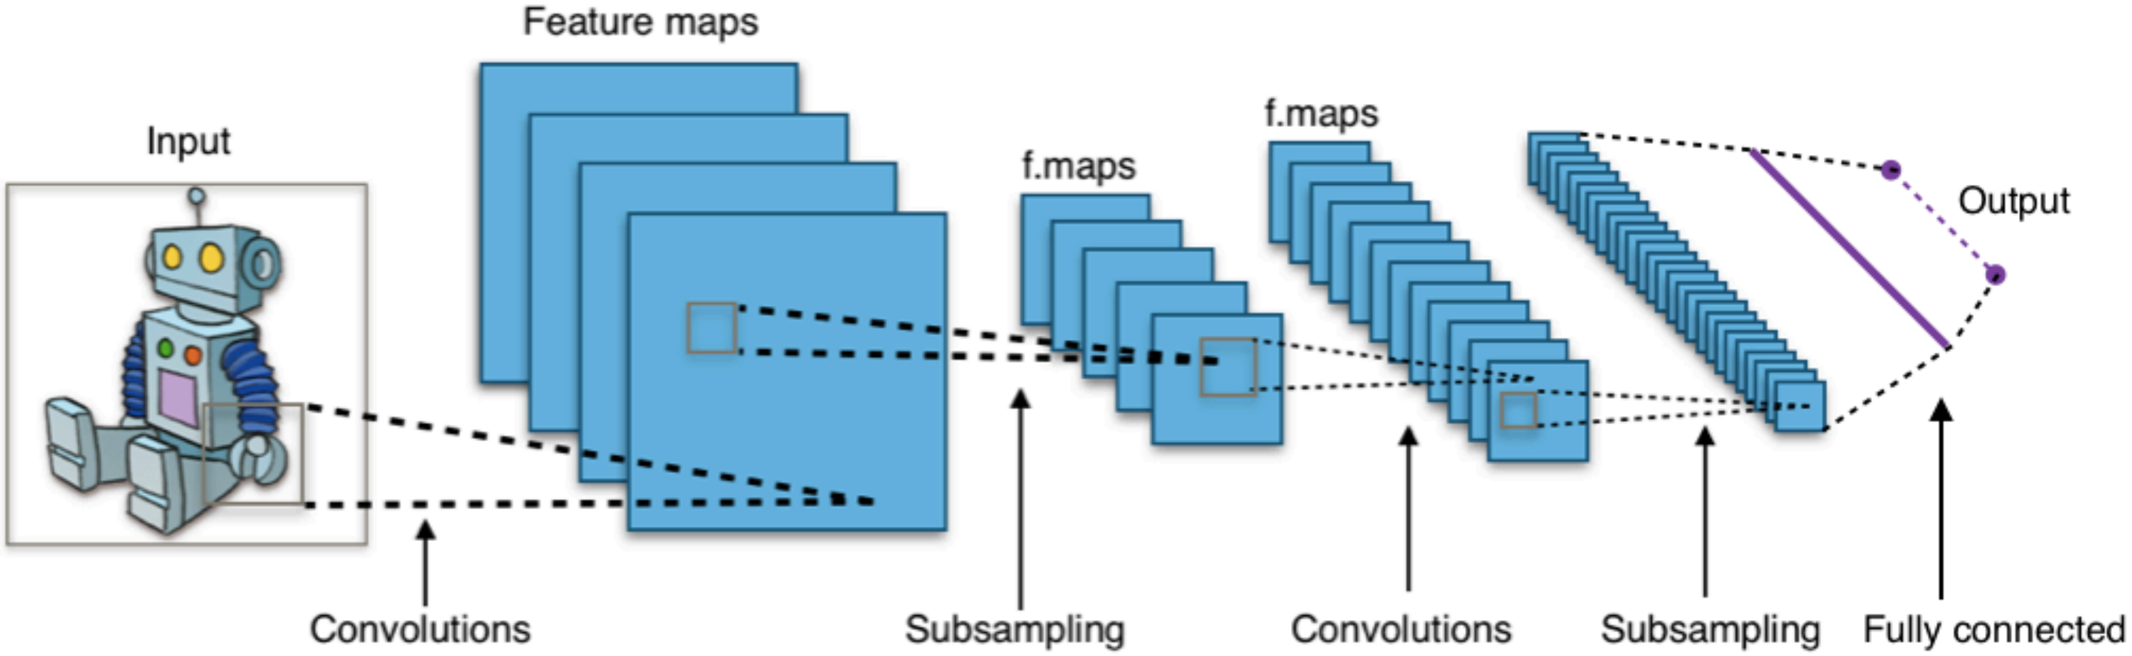

Image credit: Aphex34

Progress in machine learning

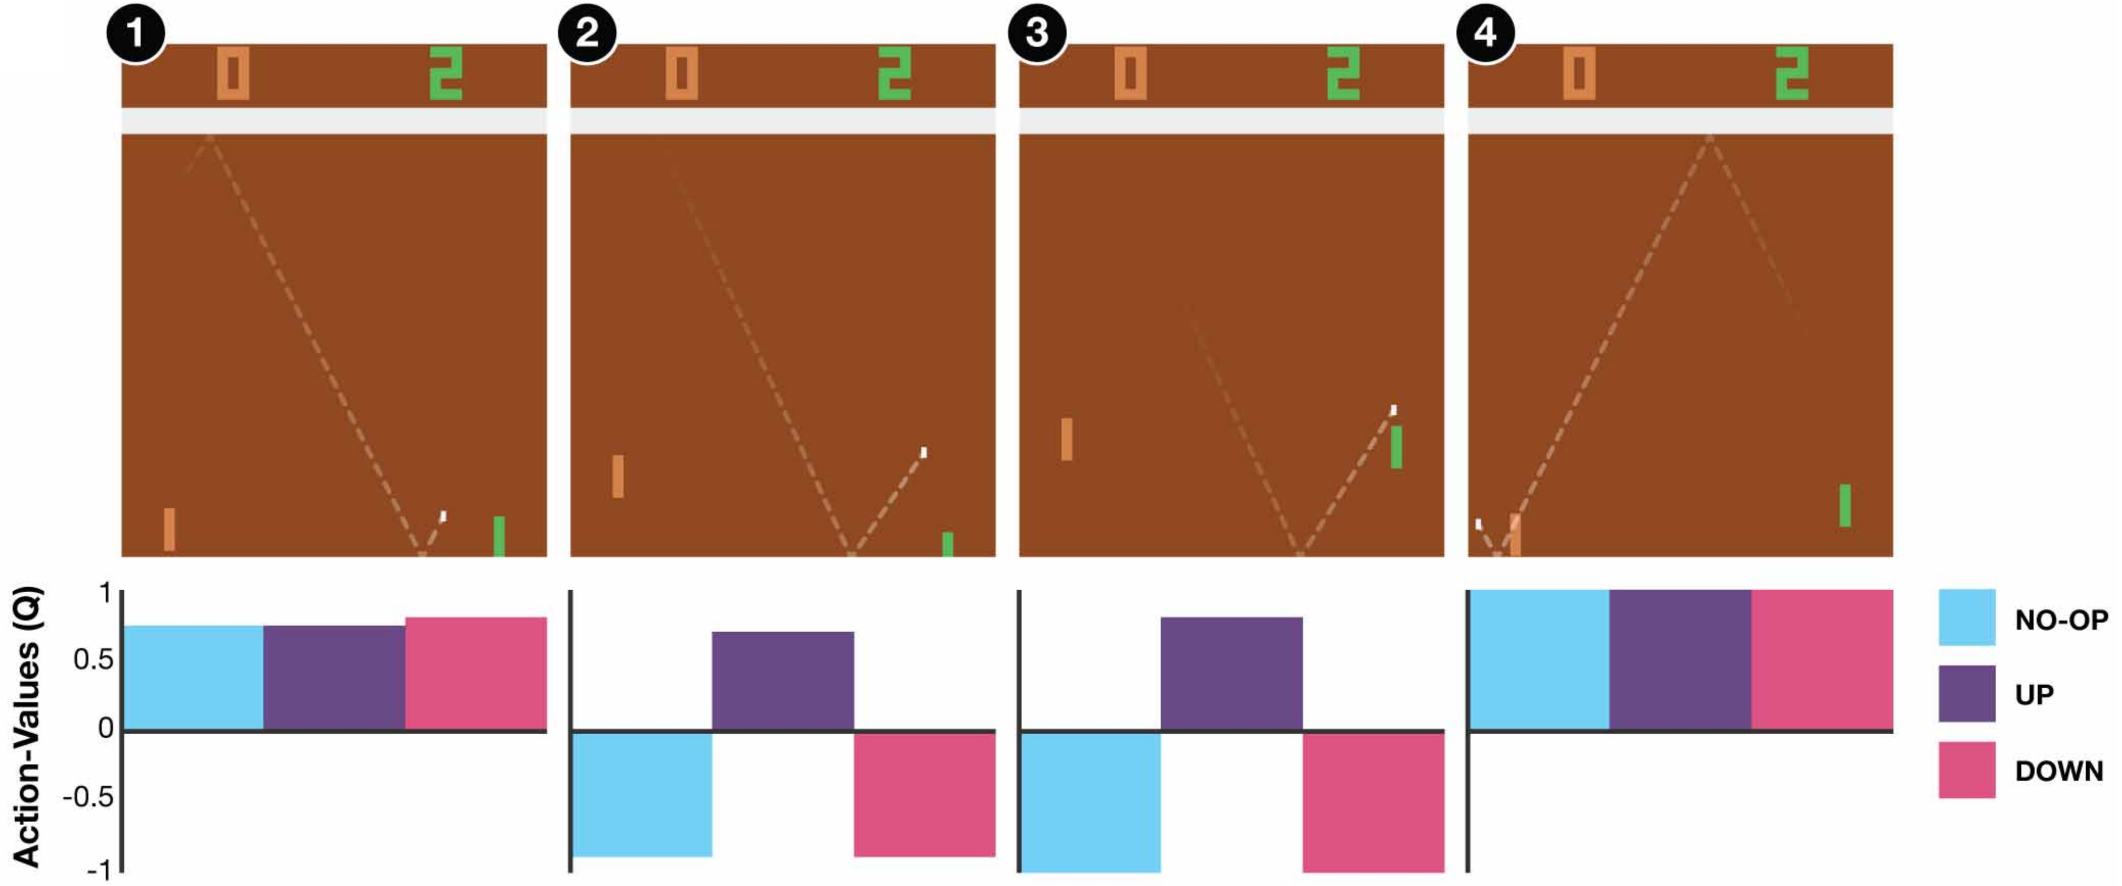

Mnih et al. 2015

The fast learning challenge

- Although the final performance of these agents is impressive, these techniques usually require several orders of magnitude more interactions with their environment than a human in order to reach an equivalent level of expected performance. [Pritzel et al. 2017]
- People learning new concepts can often generalize successfully from just a single example, yet machine learning algorithms typically require tens or hundreds of examples to perform with similar accuracy. [Lake et al. 2015]
- Notably, humans and large primates learn new knowledge even from limited experience. [Shin et al. 2017]
- Learning quickly is a hallmark of human intelligence, whether it involves recognizing objects from a few examples or quickly learning new skills after just minutes of experience. [Finn et al. 2017]
- Notably, performance in such tasks is typically evaluated after extensive, incremental training on large data sets. In contrast, many problems of interest require rapid inference from small quantities of data. [Santoro et al. 2016]

NB: fast learning is a prerequisite of general intelligence!!!

The fast learning challenge

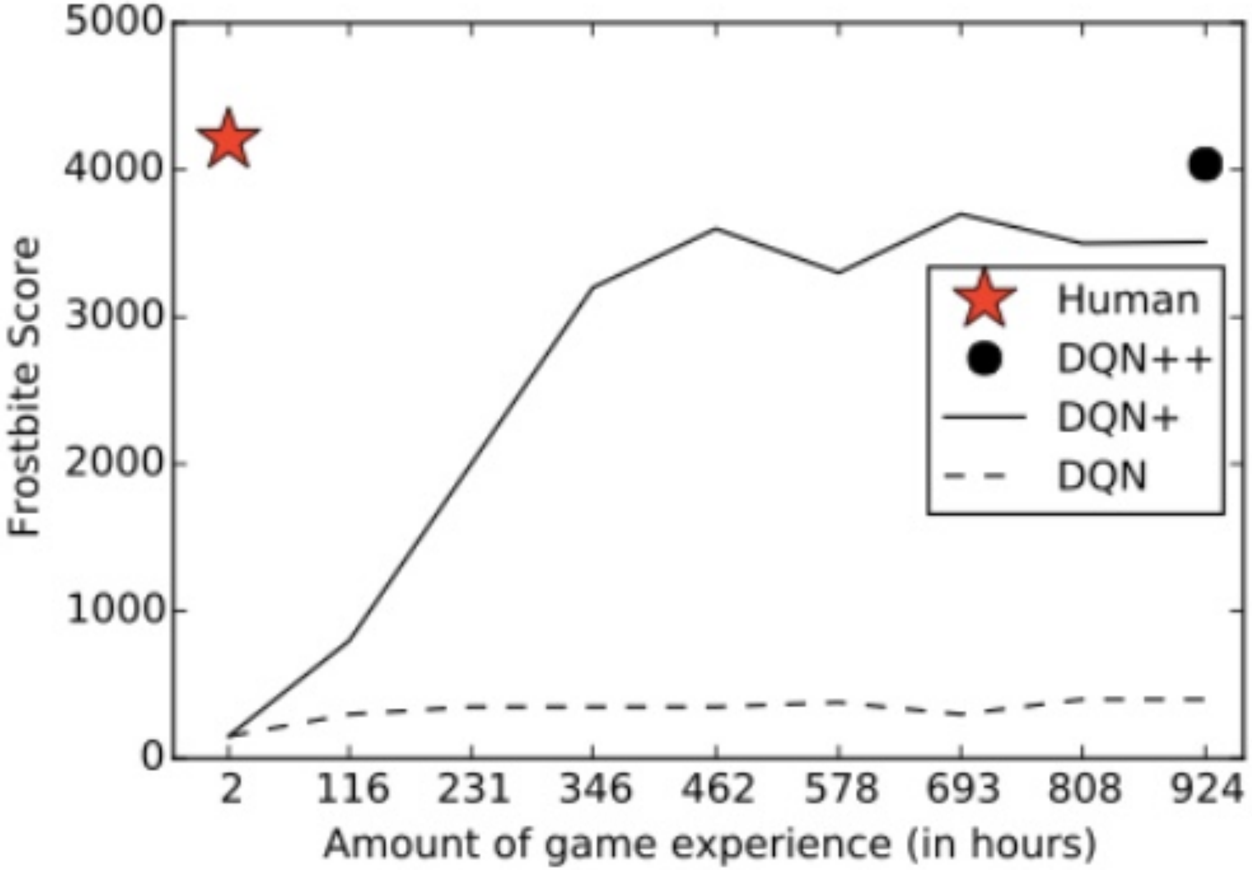

Lake et al. 2015

Not so fast!

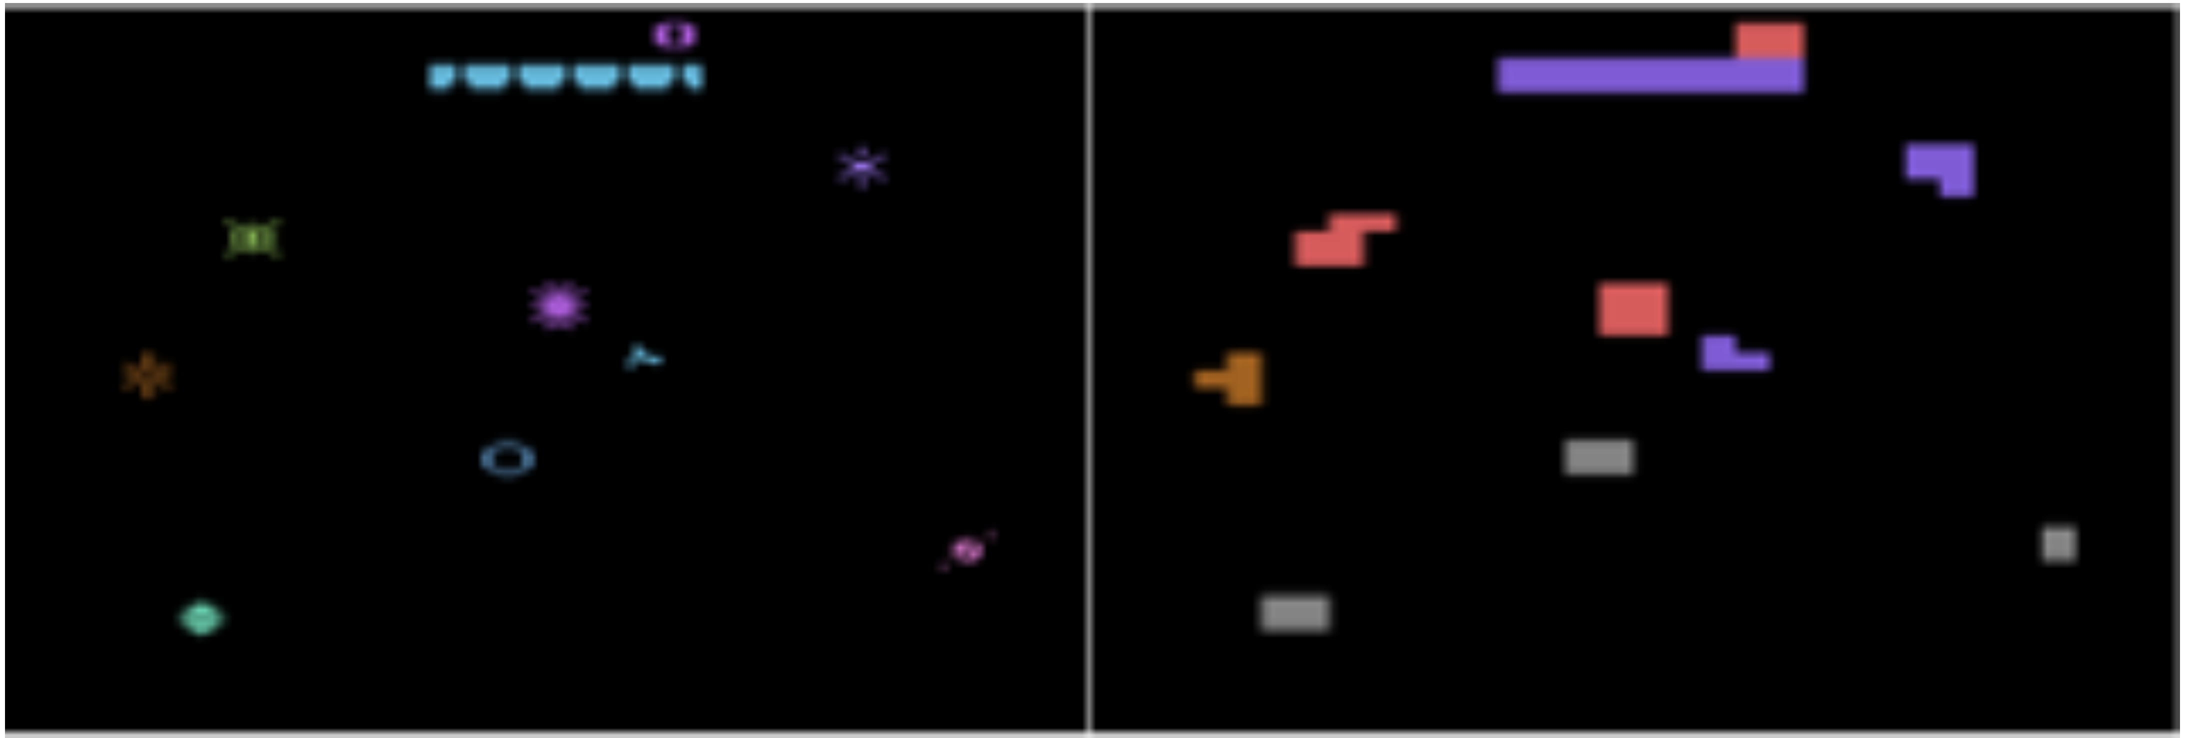

Kelly and Heywood 2017

When are we fast?

- When evolution has done the slow learning work for us
  - Naïve physics and psychology, motor skills, language, reasoning...
- When new problems can be solved by combining old tricks (compositionality)

Fast learning

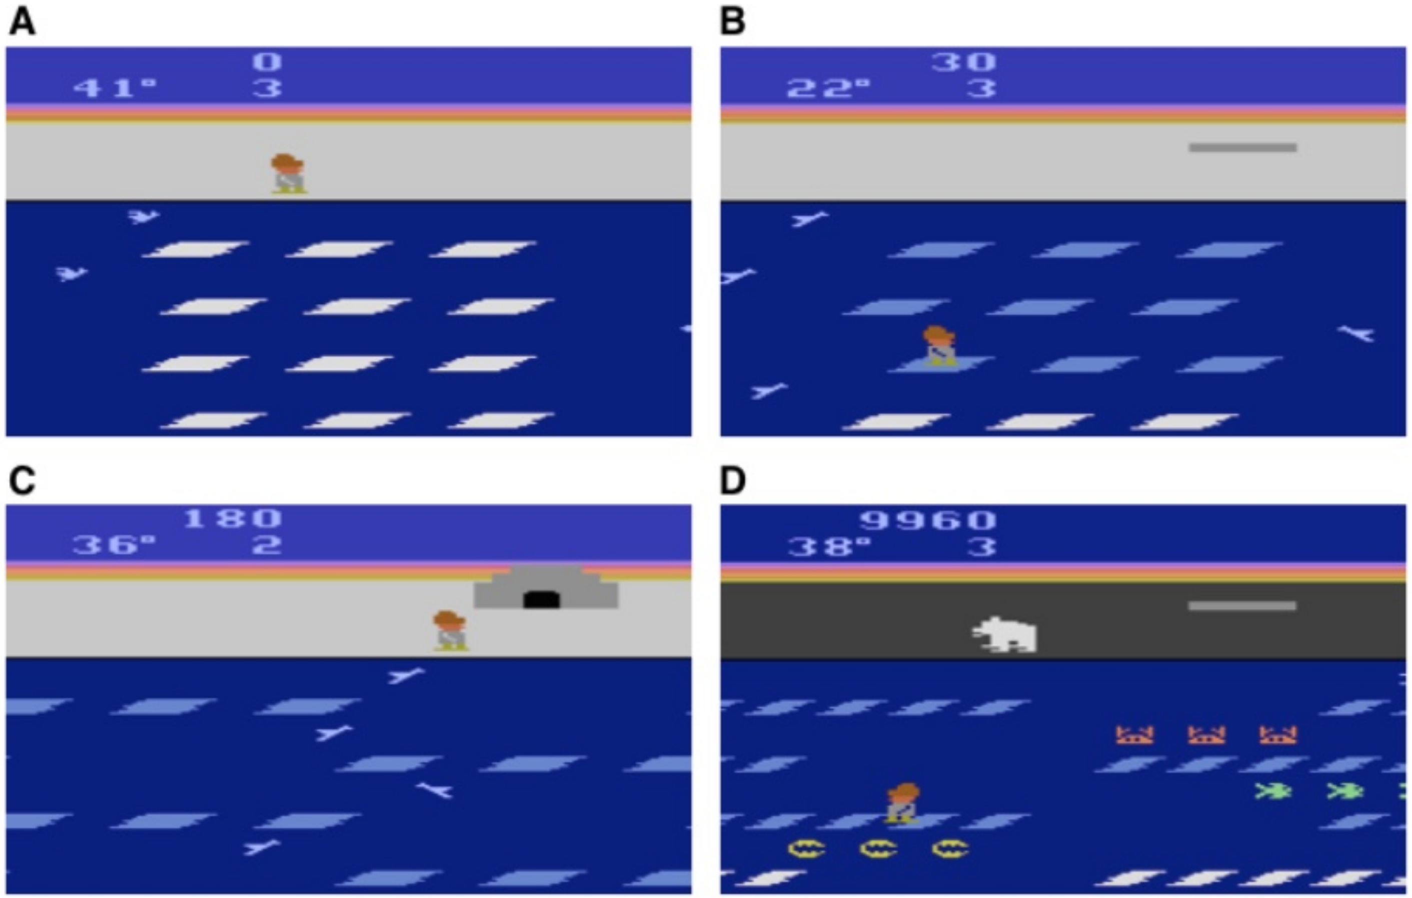

Lake et al. 2015

Fast learning

- Can you pick Angelina today at 5pm?
- Sure, where?
- She's at the Lluïsos in Plaça del Nord.
- And where is that?
- You go to the supermarket where we always shop, turn left after you passed it, and continue for about 100 meters... the theater is the big red building on the right

How to build faster learners

- Speed up evolution by programming the right innate skills
  - Very challenging: we are better at designing learning machines than at hand-coding the right kind of knowledge
  - We'll have to tackle this eventually, though!

- Develop architectures capable of compositional learning
  - Also challenging, but we have clearer ideas of what is required
  - Potentially very useful, as we can leverage skills acquired with effective slow learning techniques to accomplish a combinatorial explosion of new tasks

Desiderata for a compositional learner

- Learner must still be able to acquire new skills directly from data
- Learner must discover when/what/how to compose
- New compositions should be fast, a lot faster than acquiring a new skill
- Frequently needed composed skills can be memorized and accessed just like basic data-induced skills
- As few constraints as possible on modes of composition (e.g., not excluding recursion)

Compositional architectures: a (nice) example

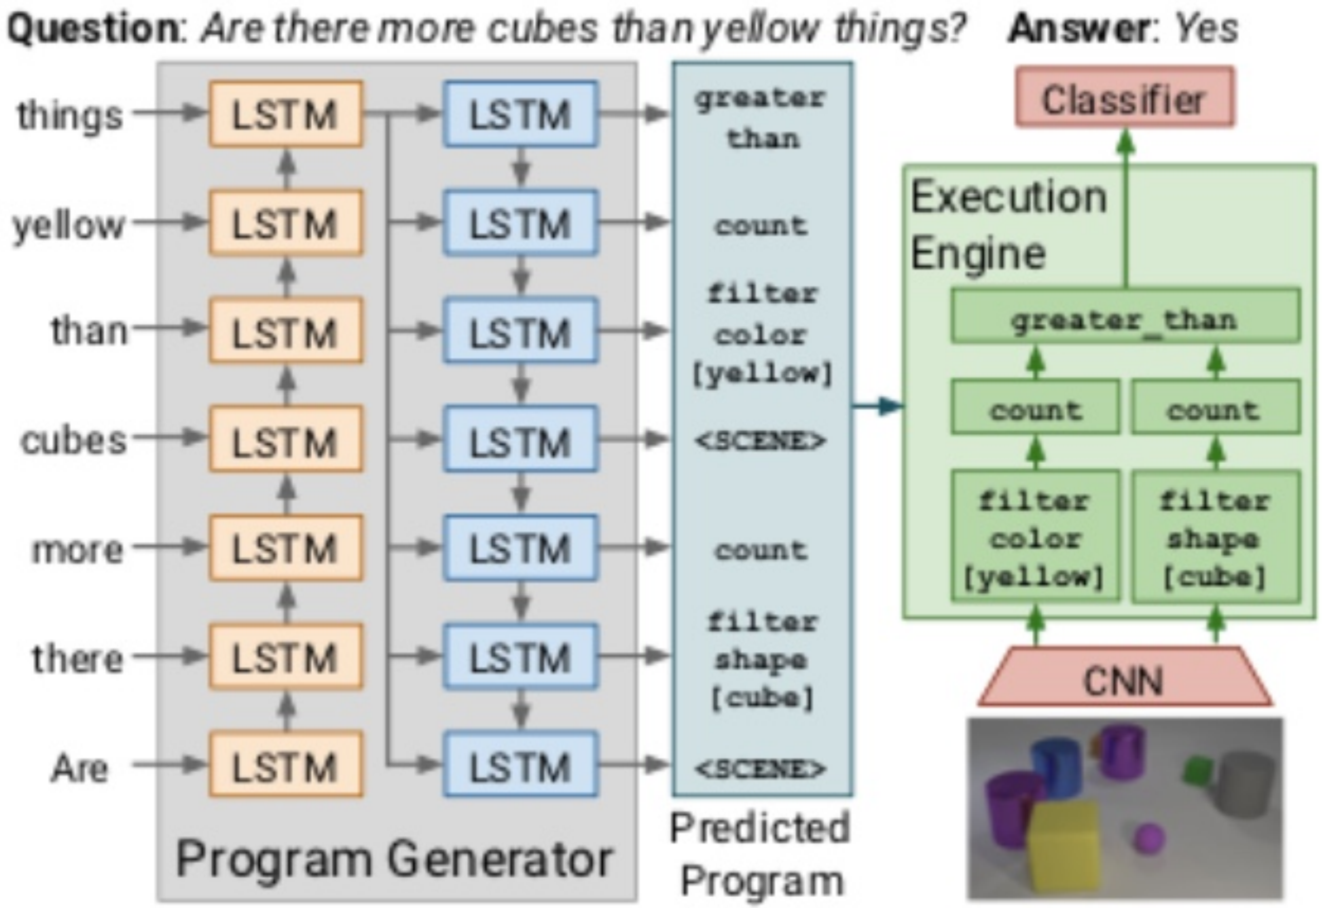

Johnson et al. 2017

Compositional architectures: a (nice) example

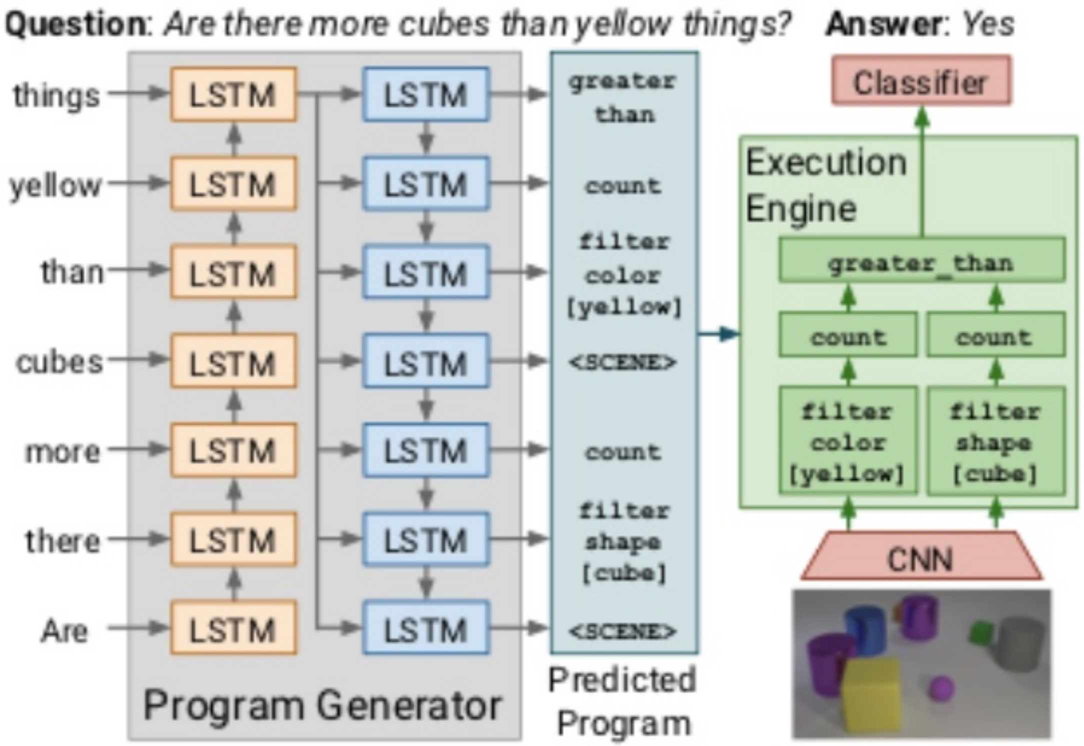

Johnson et al. 2017

- Acquiring new skills 🙁
- Discovering when/what/how to compose 🙃
- Composing is fast 🙂
- Composed skill memorization 🙁
- Few constraints on modes of composition 🙃

WARNING

In the next slides, I will sketch a model that currently only exists in our rosiest dreams. Do not use it to operate heavy vehicles.

To keep things in the right perspective, I will call this model MARCONE, for **MA**gical **R**ecurrent **CO**mpositional **NE**twork

Starting with a very general architecture

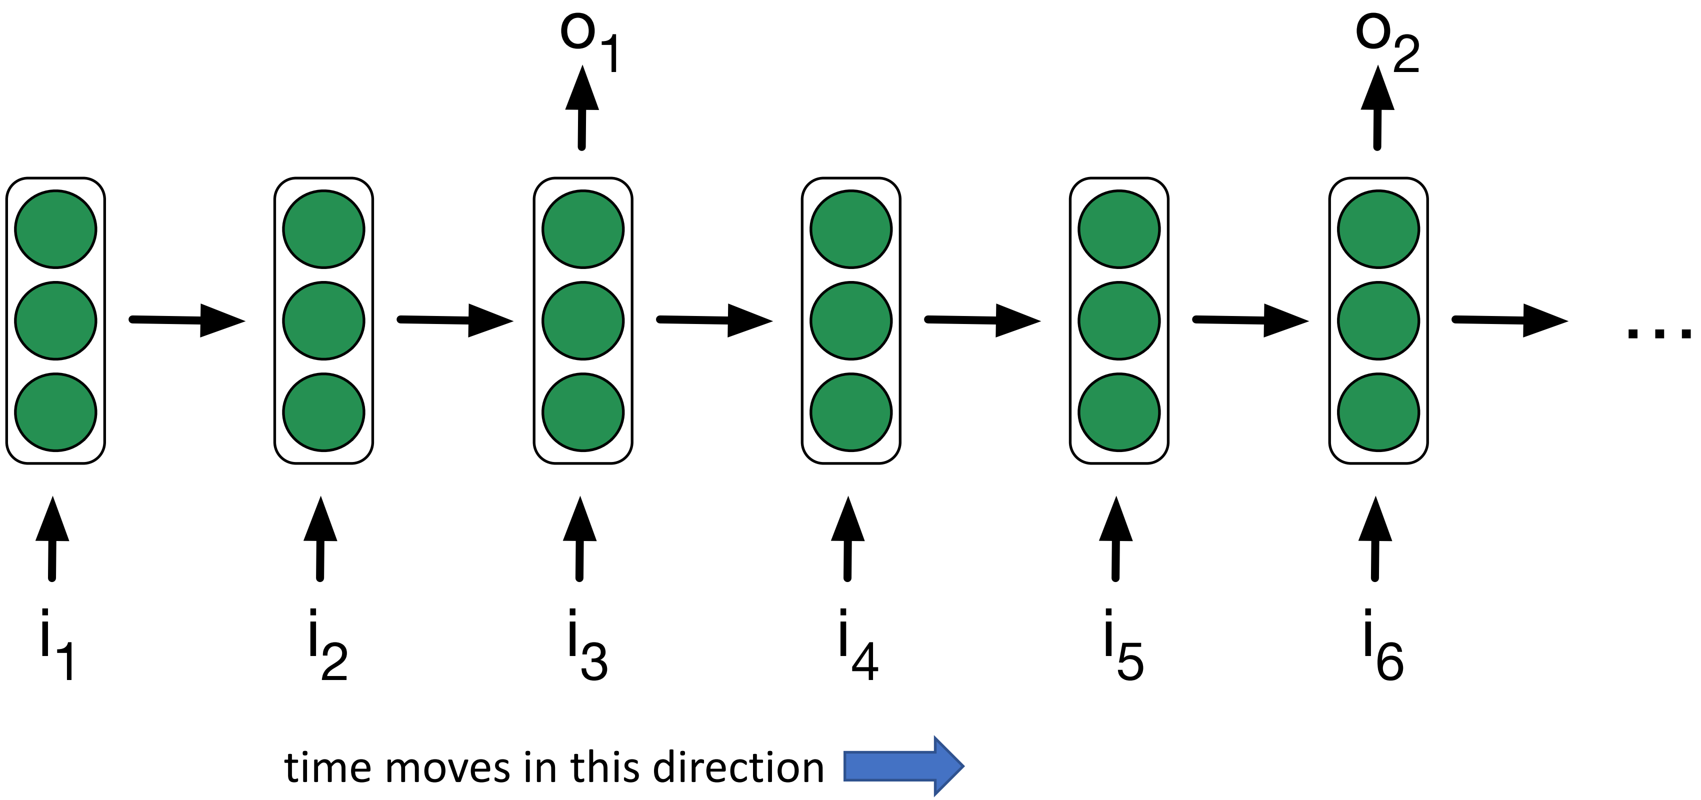

Elman 1990, Jordan 1986, Hochreiter and Schmidhuber 1997, ...,
Mikolov et al. 2010, Graves et al. 2013, ...

The RNN as a compositional architecture?

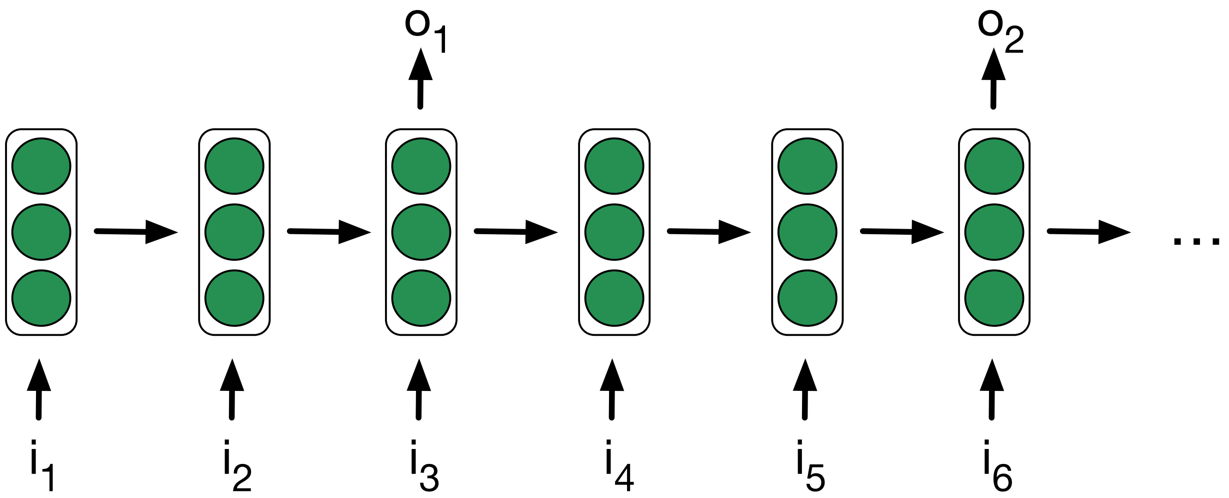

- Powerful learning capabilities 🙂
- Designed for sequential processes, natural choice for continual learning setups 🙂
- Agnostic, imposes few constraints on how tasks must be solved 🙂

- To compose skills, it should first avoid forgetting them 🙁
- Not clear how composition would be performed 🙁

Challenge #1: Long-term skill memory

Catastrophic forgetting (McCloskey and Cohen 1989):

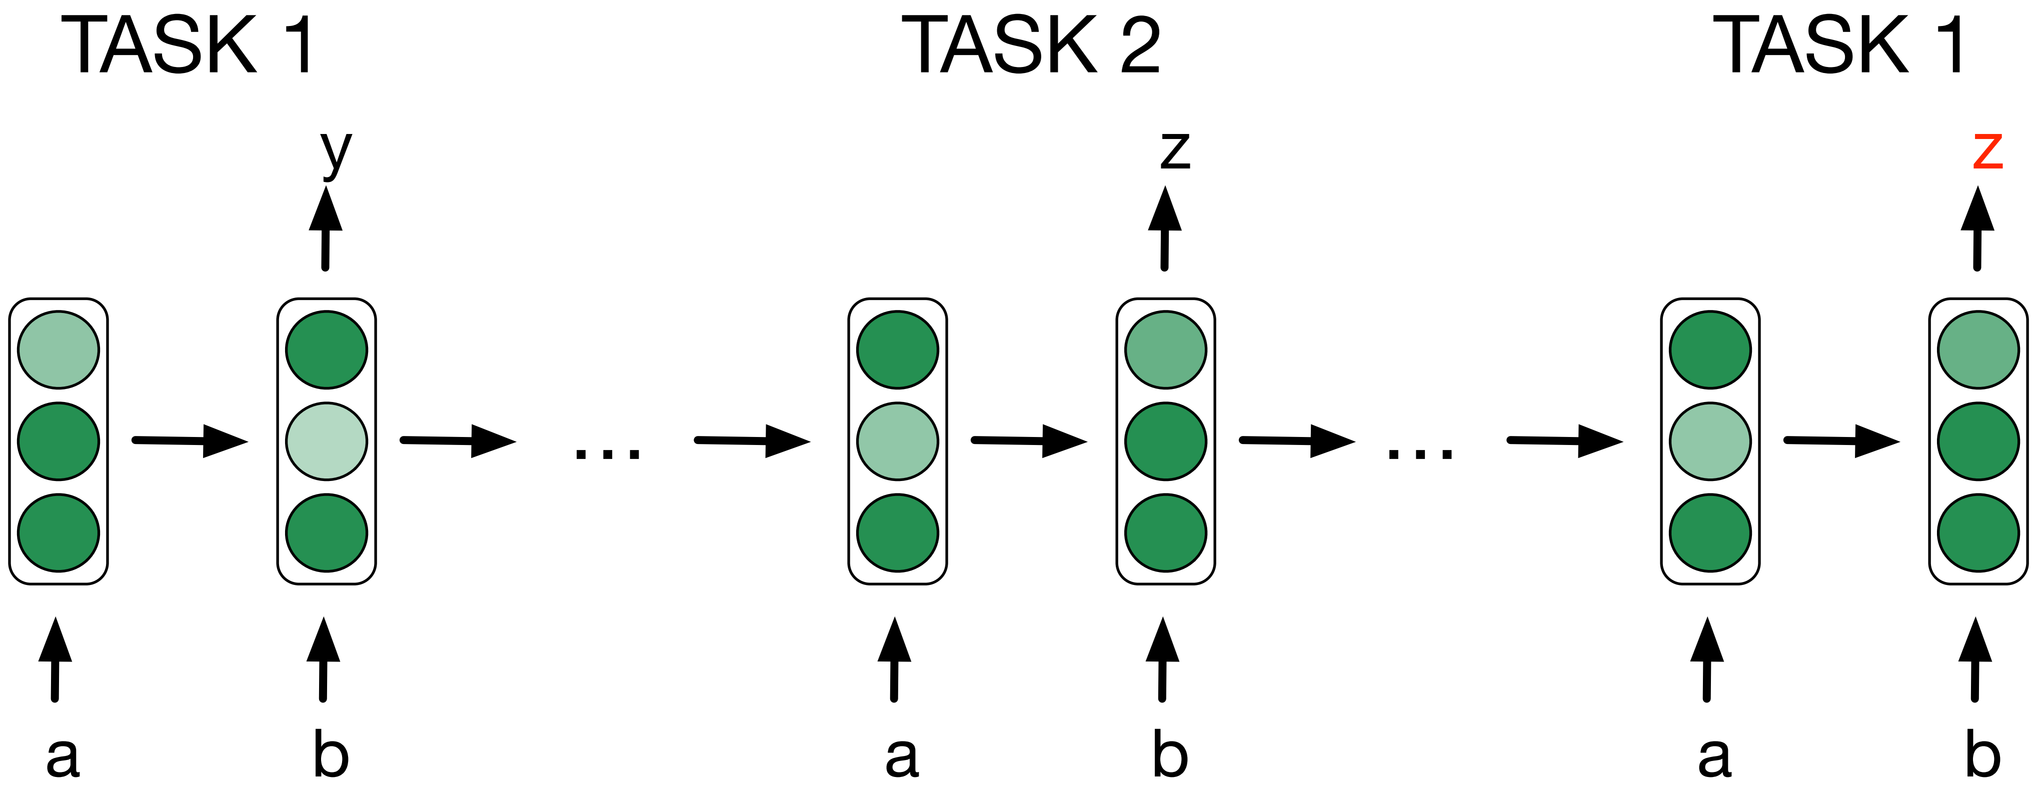

Catastrophic forgetting

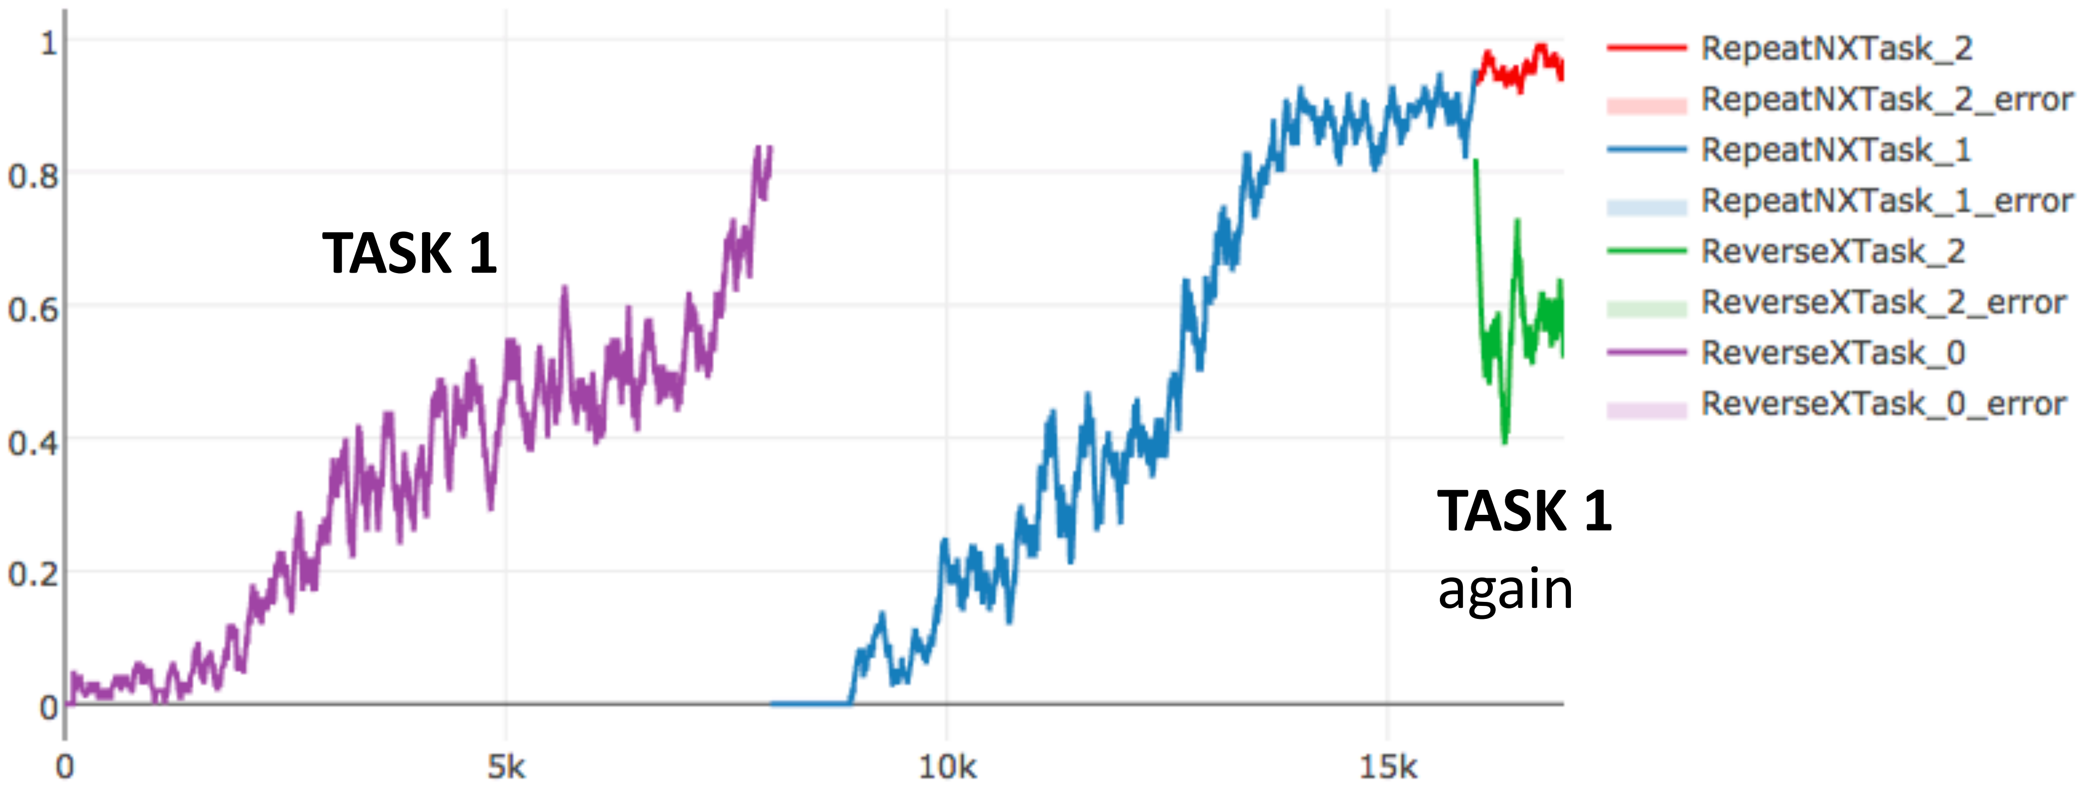

The MARCONE solution

Related: Lu et al. 2016, Kickpatrick et al. 2017, Lopez-Paz and Ranzato 2017, ...

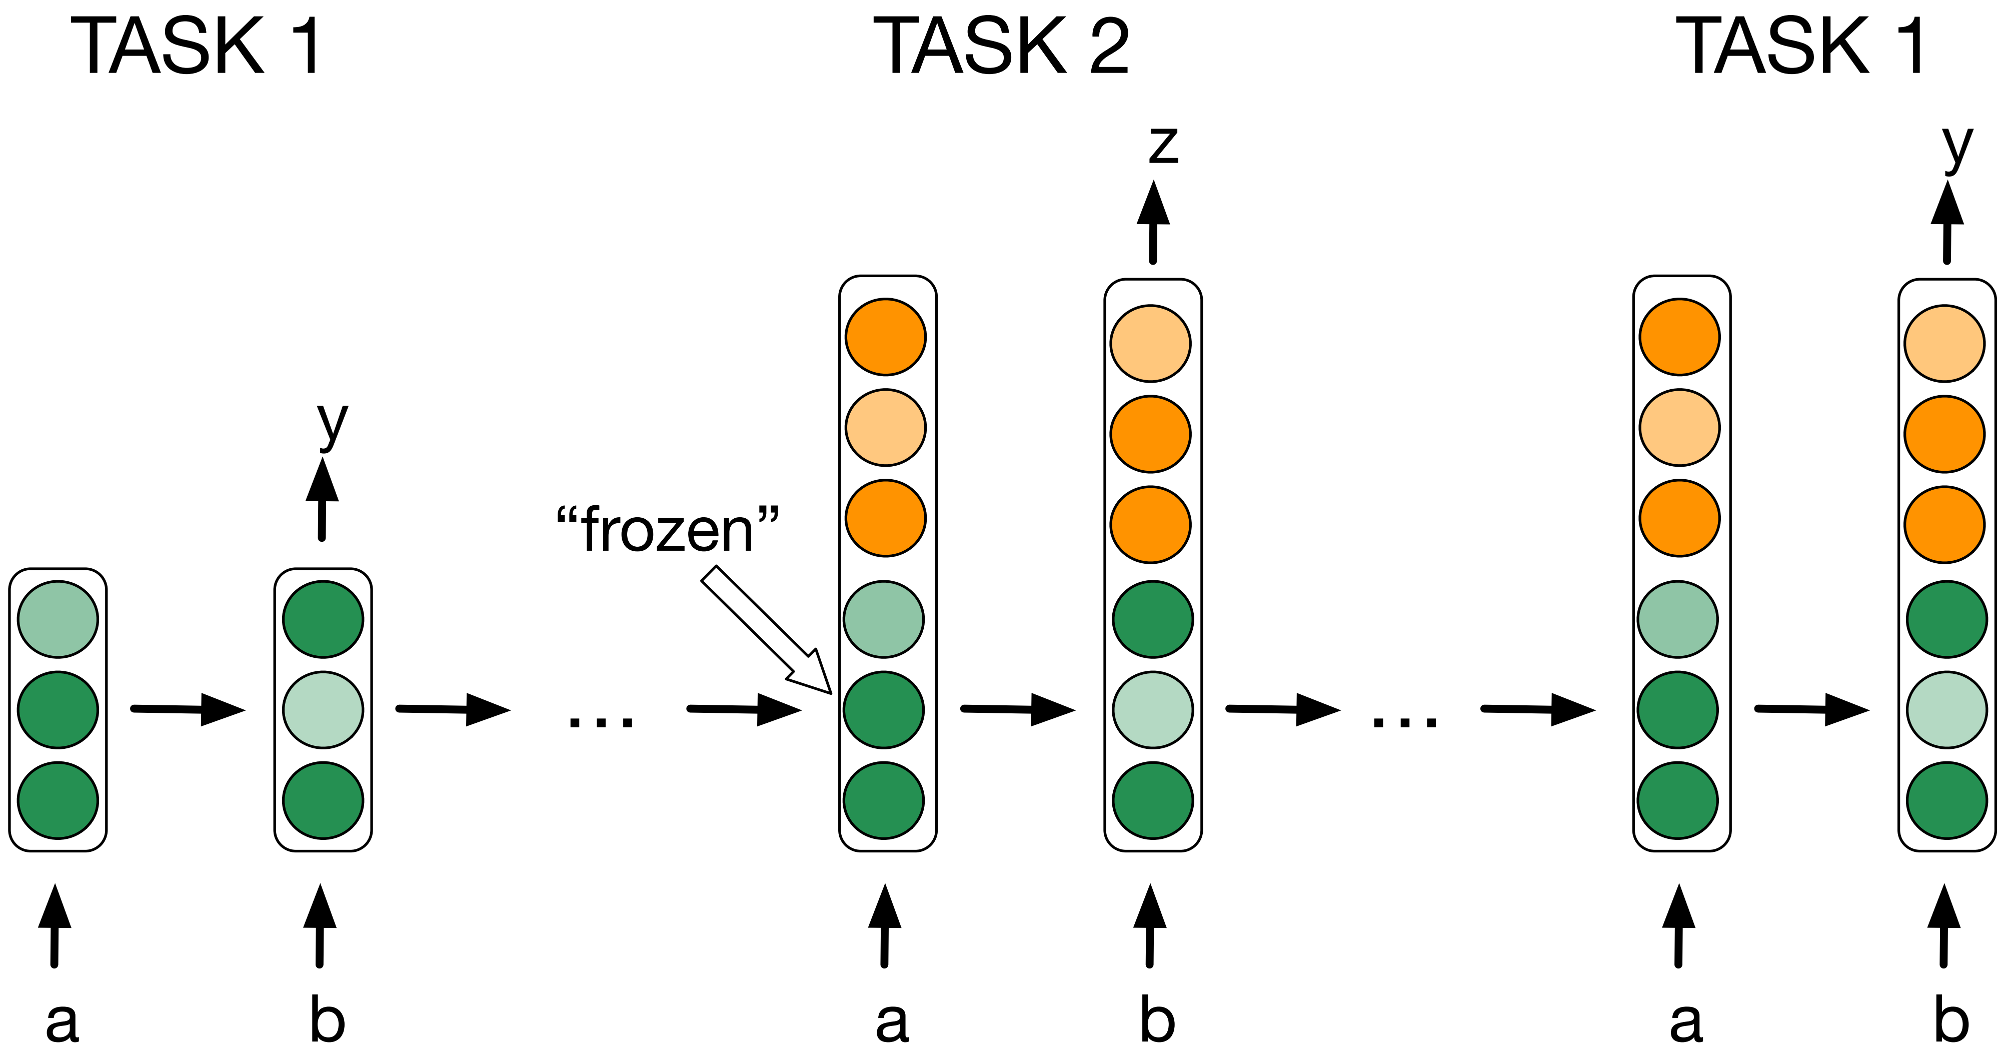

Challenge #2: Composing with a RNN

- Limited form of composition could be learned by growing RNN
- New units could specialize into "controller" role
- If compositional solution simply requires applying two skills in sequence, this might work

e.g.

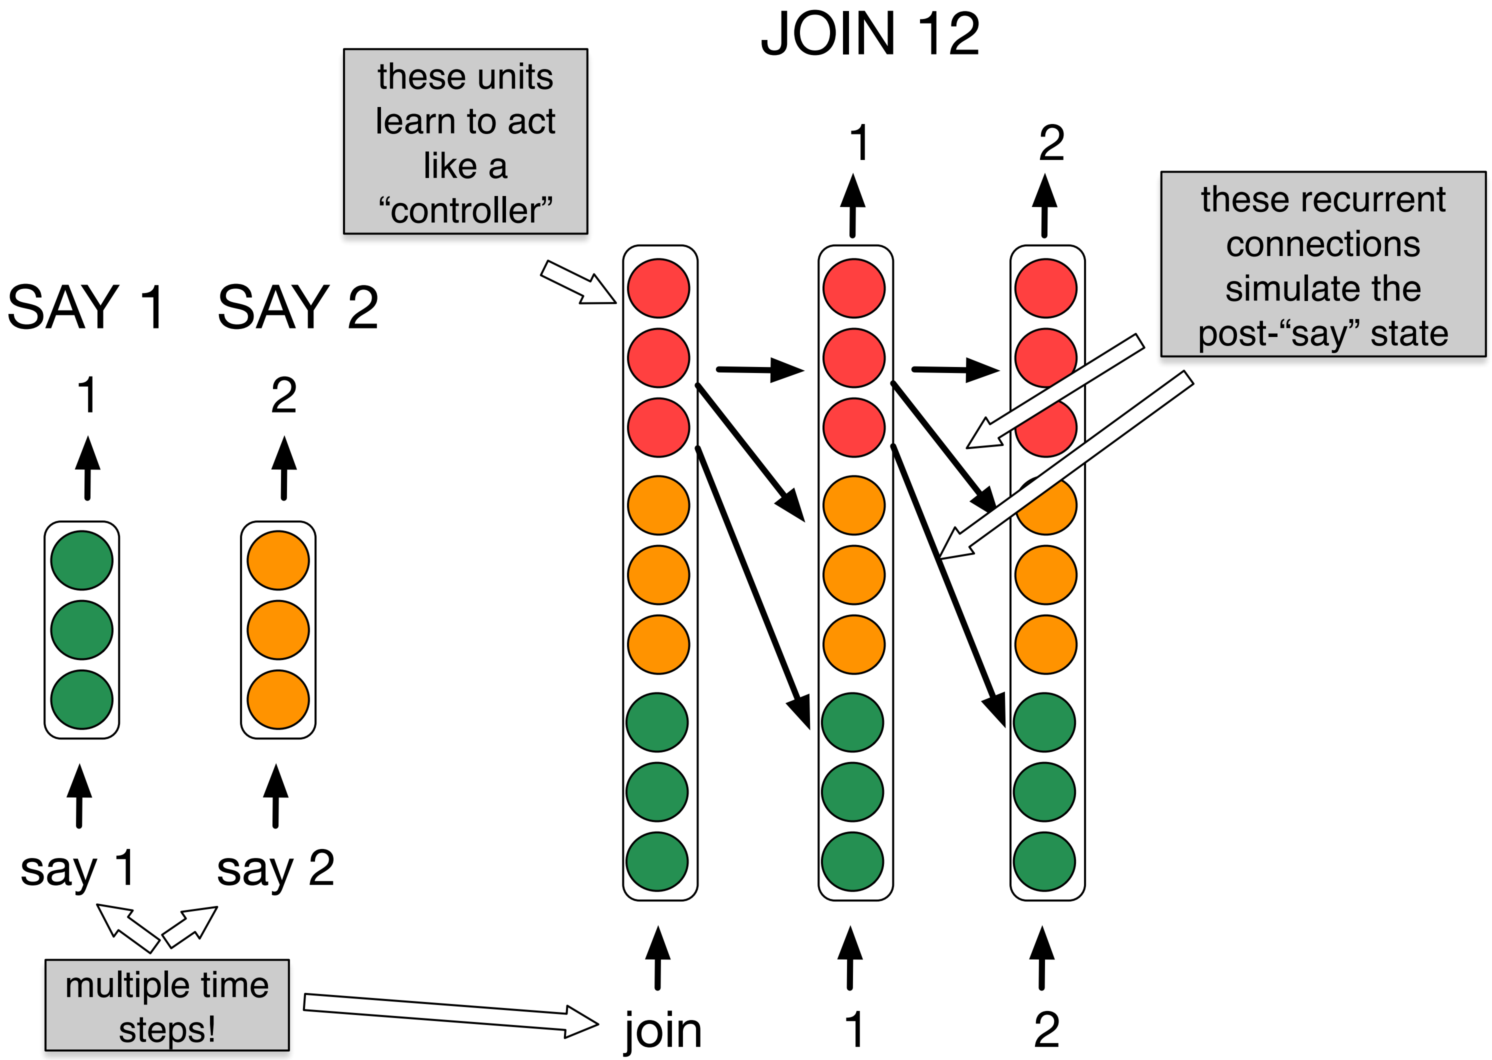

- But what if composition requires multiple reasoning steps?

f(g(x))

- Need dynamic approach (e.g., adding fixed number of layers would, at least, be very inefficient)

f(g(f(x)))

The MARCONE solution

Dissociate network processing (time) steps from input/output processing
Network learns to allocate "pondering time" (Graves 2017, Mozer and Das 1993)

Tasks:
reverse 100 = 001
chomp 001 = 00
chomp reverse 100 = 00

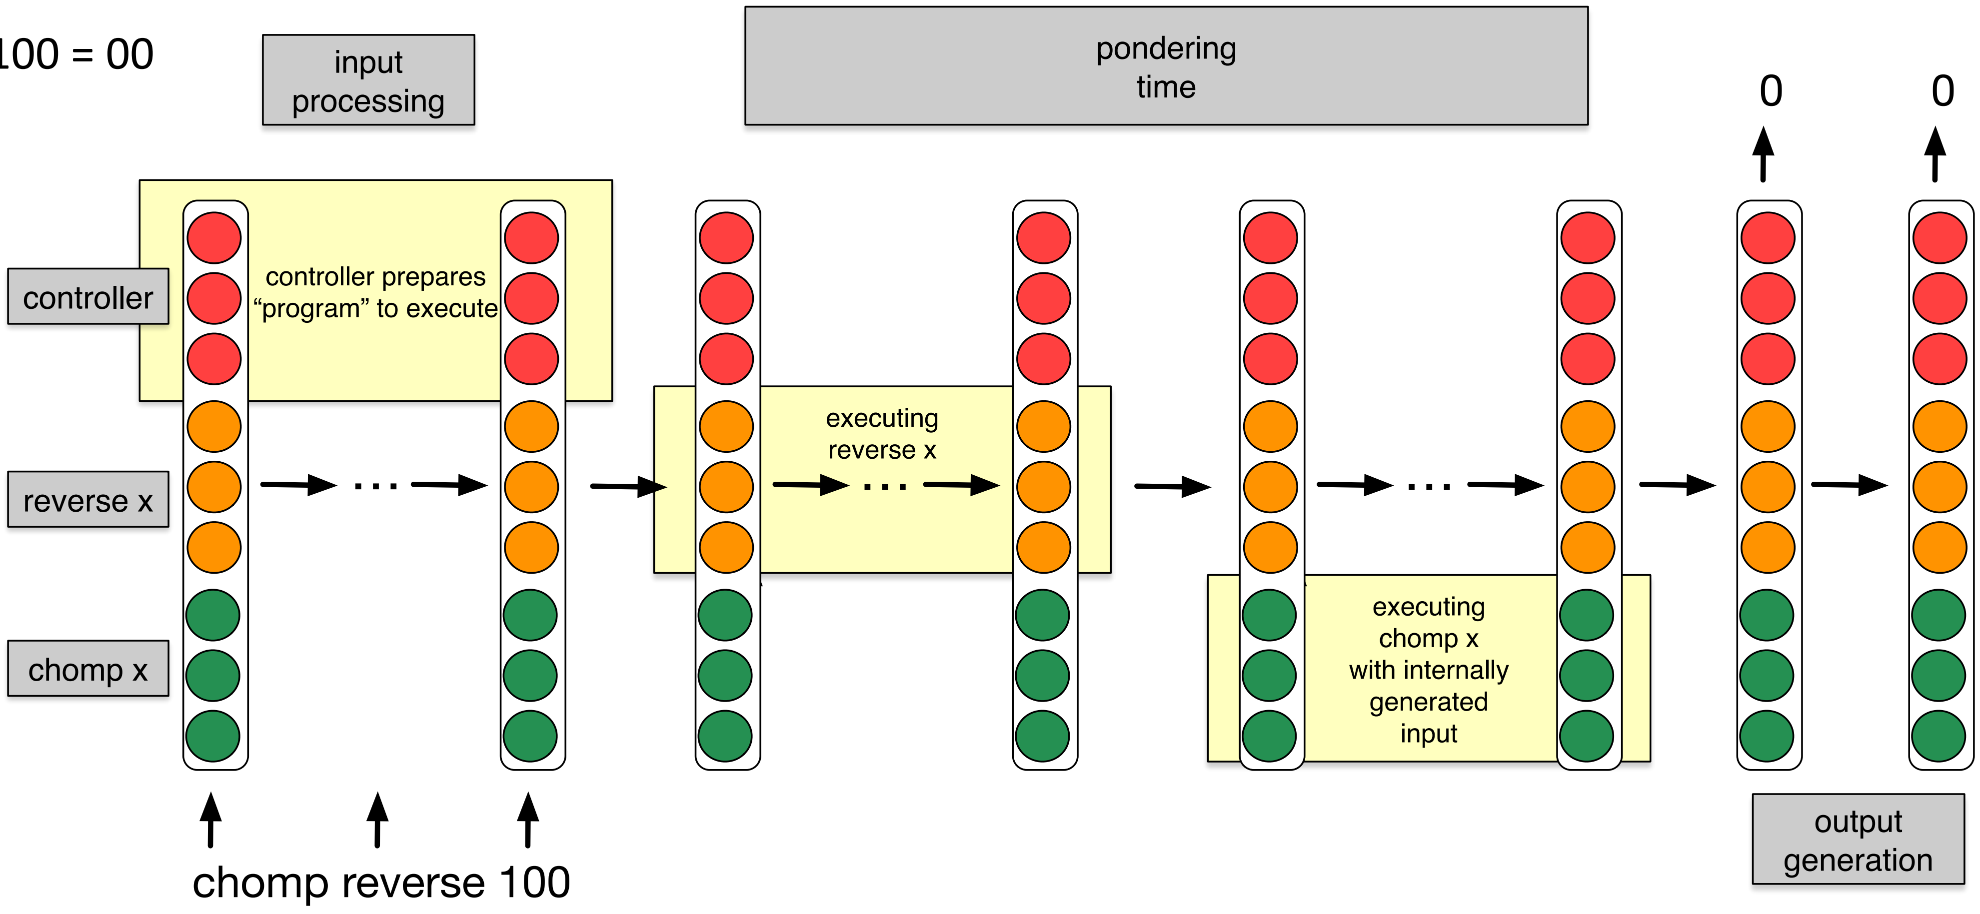

Composition in MARCONE

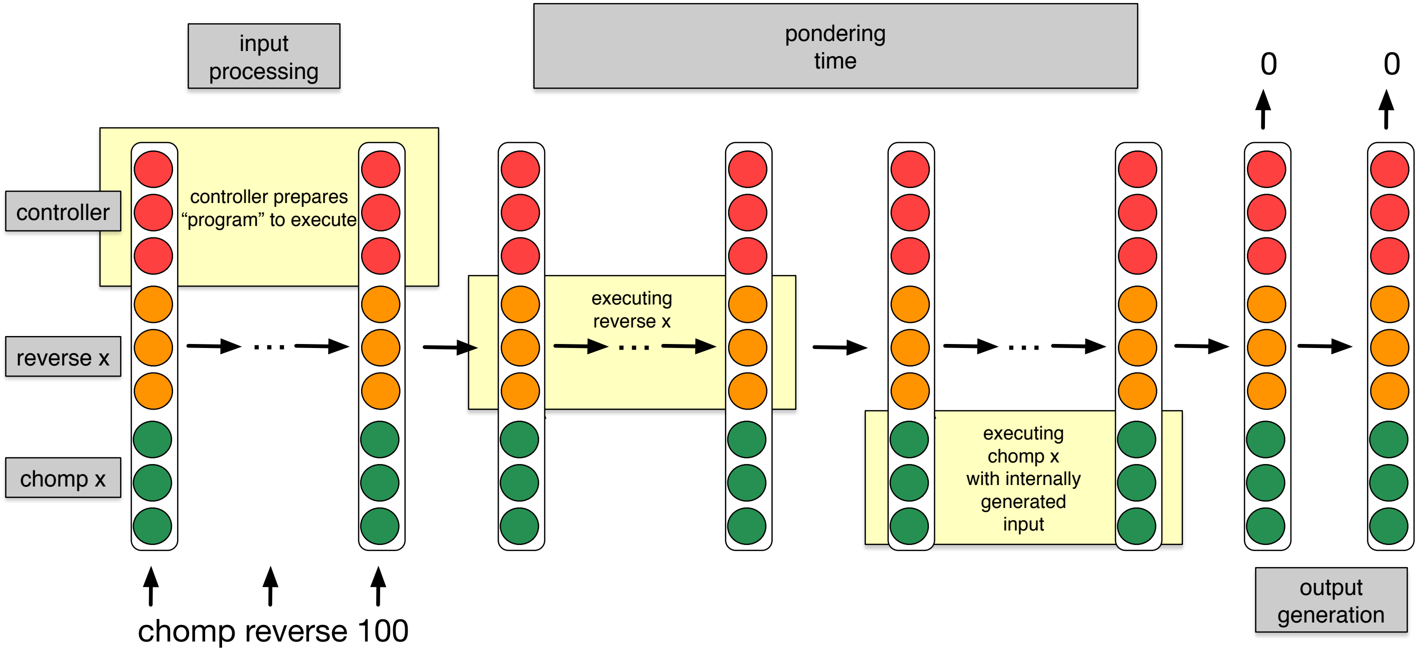

- Acquiring new skills 🙂
- Discovering when/what/how to compose 🙂
- Composing is fast 🙂
- Composed skill memorization 🙂
  - Also, skills can share components
- Few constraints on modes of composition 🙂

MARCONE: research challenges

- Weight consolidation
- Growing new units (avoiding full connectedness)
- Pondering

- Appropriate training regime to encourage compositional learning
- Are gradient descent methods the right learning tools for the discrete, combinatorial side of composition?

INTERMEZZO:
a possibly useful
take-home message

It takes hundreds of thousands of training samples and lots of tricks to teach neural networks to play ATARI games

...

and it takes hundreds of thousands of training samples and lots of tricks to teach neural networks to memorize short bit sequences

The way in which currently popular gradient-based algorithms are learning is far from intuitive expectations about learning

The table lookup tasks

	t1	t2	t3	t4	t5
00	10	11	00	11	10
01	11	00	01	10	00
10	01	01	10	00	01
11	00	10	11	01	11

t1: 00 ⇨ 10, 01 ⇨ 11, 10 ⇨ 01, 11 ⇨ 00
t2: 00 ⇨ 11, 01 ⇨ 00, 10 ⇨ 01, 11 ⇨ 10
t3: 00 ⇨ 00, 01 ⇨ 01, 10 ⇨ 10, 11 ⇨ 11
t4: 00 ⇨ 11, 01 ⇨ 10, 10 ⇨ 00, 11 ⇨ 01
t5: 00 ⇨ 10, 01 ⇨ 00, 10 ⇨ 01, 11 ⇨ 11

…

t1(00)=10
t3(00)=00
t4(t5(01))=11
t5(t4(01))=01
t2(t2(10))=00
t1(t4(t5(11)))=11
t1(t5(t1(10)))=10

thanks to Angeliki Lazaridou
and José Hernandez-Orallo
for the lookup task idea!

Remembering lookup tables

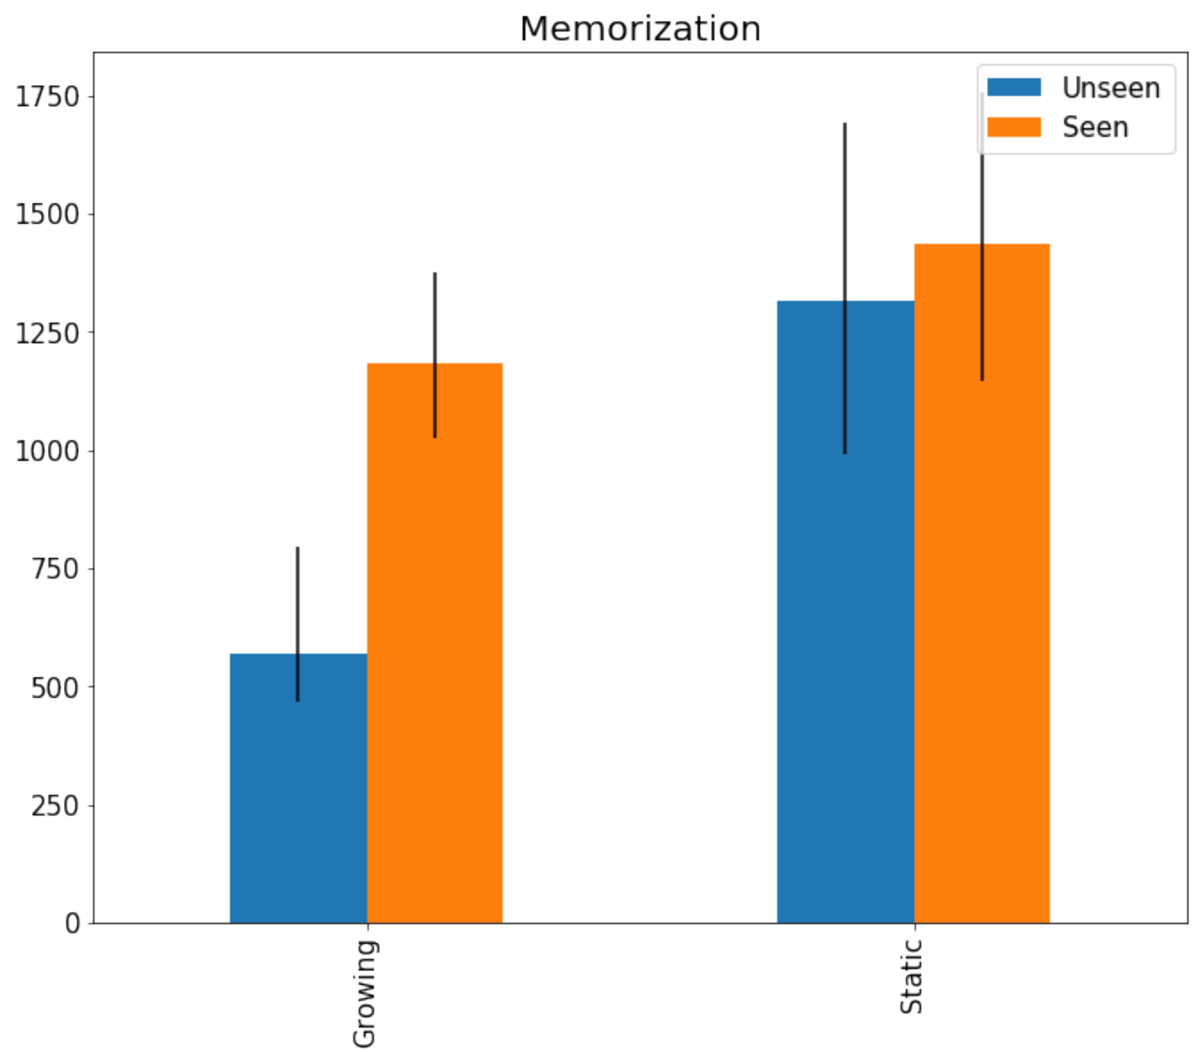

medians of 100 experiments
with 5-bit lookup tables

Concatenating lookup tables

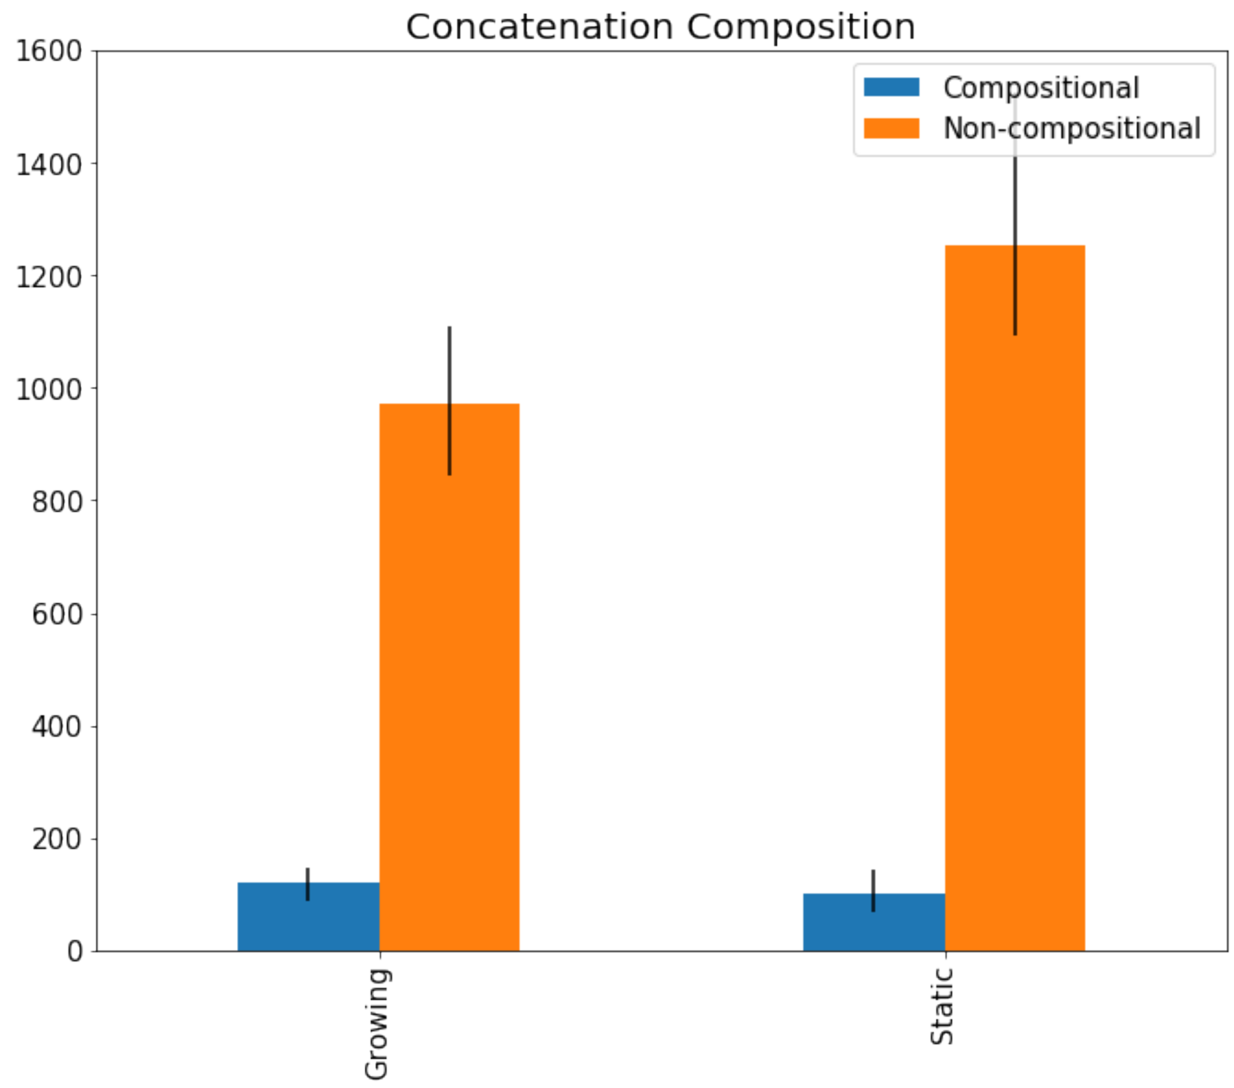

these are concatenations of 4-bit lookup tables

Properly composing lookup tables

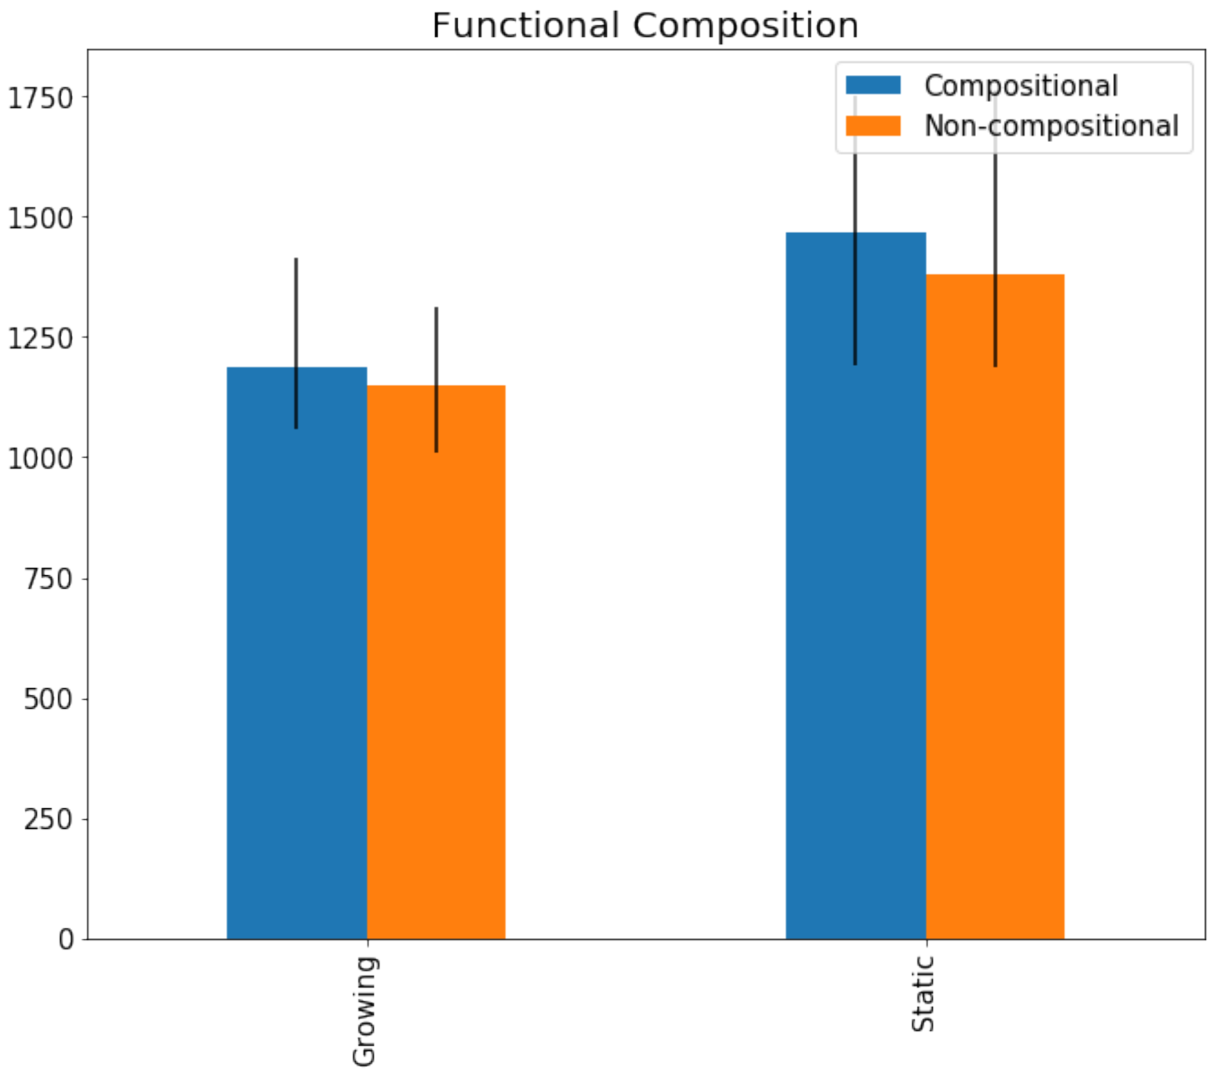

A quick look at the bigger picture:

The intelligent machine we'd want

- Hi Machine, today my blood test results should be available at the clinic, can you help me picking them up?
- Sure, how can I do that?
- Search the number of St. George Clinic, call them, and ask them for their hours. Then, call me a cab for the earliest time at which they're open.

Desideratum #1:
Flexibility, fast adaptation to new tasks

- Hi Machine, today my blood test results should be available at the clinic, can you help me picking them up?
- Sure, how can I do that?
- Search the number of St. George Clinic, call them, and ask them for their hours. Then, call me a cab for the earliest time at which they're open.

compositionality plays a big role here!

Desideratum #2:
Ability to communicate and learn through natural language

- Hi Machine, today my blood test results should be available at the clinic, can you help me picking them up?
- Sure, how can I do that?
- Search the number of St. George Clinic, call them, and ask them for their hours. Then, call me a cab for the earliest time at which they're open.

Desideratum #3:
Learn from light error signals

- Hi Alice, I have booked the cab for 2.30pm.
- Great, thank you!

The CommAI approach: Simple tasks, big challenges

E: concatenate A and K.
L: djksjdkjf.
E: wrong, you should have said AK.

E: reverse KRM.
L: MRK.
E: right. [+1 reward]

E: reverse concatenate K and XYK.
L: KYXK.
E: right. [+1 reward]

E: reverse BRGJ.
L: JGRB.
E: right.

The CommAI initiative

https://research.fb.com/projects/commai/

- Interactive environment to test general learners
  https://github.com/facebookresearch/CommAI-env
- CommAI-env user group
  https://www.facebook.com/groups/1329249007088140/
- GoodAI general AI challenge based on CommAI
  https://www.general-ai-challenge.org
- The CommAI visiting researcher call (call for PhD fellowships coming soon!)
  https://research.fb.com/programs/post-docs-and-sabbaticals/#CommAI_Visiting_Researcher_Program