What will it take to solve human cognitive development?

Jay McClelland
Stanford University

Visiting at

DeepMind

Questions

- Where does our knowledge come from?
- How does it evolve over time?
- How do we go from very primitive initial abilities to all of the amazing achievements of human thought?

My Basic Tenets

- Gradient not discrete
- Connection- not Proposition-based
- Learned not built in
- General not specific

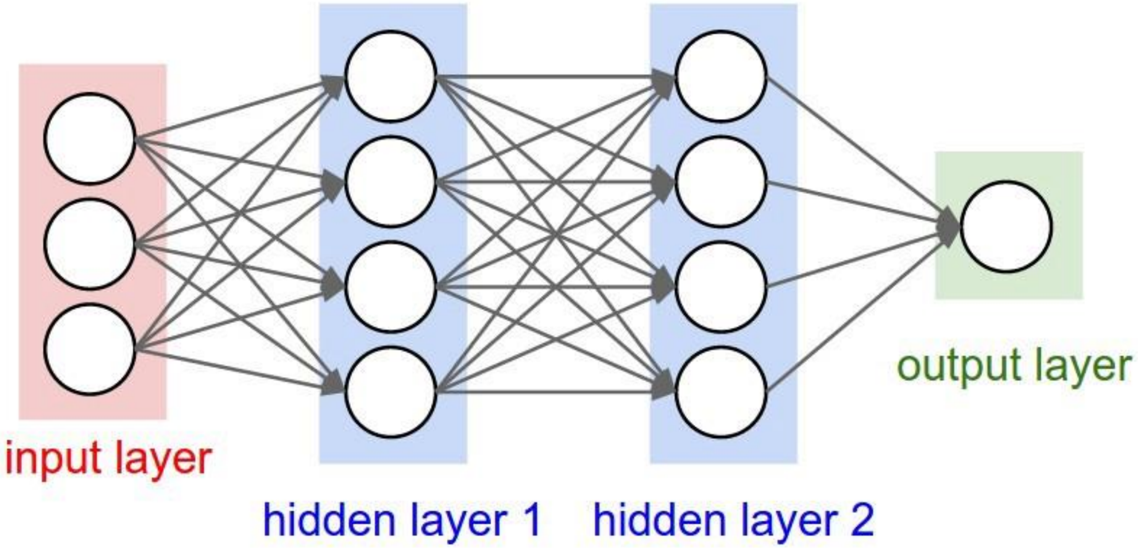

Philosophy behind our Approach

1. Learning vs. Handcrafted
2. General vs. Specific
3. Grounded vs. Logic-Based
4. Active vs. Passive

DeepMind

Prediction-error driven learning

Adjust each parameter of the mind to reduce the discrepancy between predicted and observed events

- Could this be the engine that drives cognitive development and the discovery of new knowledge?

McClelland, J. L. (1989). Parallel distributed processing: Implications for cognition and development. In Morris, R. (Ed)., *Parallel distributed processing: Implications for psychology and neurobiology.* (pp. 8-45). New York: Oxford University Press.

McClelland, J. L. (1994). The interaction of nature and nurture in development: A parallel distributed processing perspective. In P. Bertelson *et al.* (Eds.), *International perspectives on psychological science, Volume 1: Leading themes.* United Kingdom: Erlbaum.

Aspects of Human Developmental Change: A Case Study (McClelland, 1989)

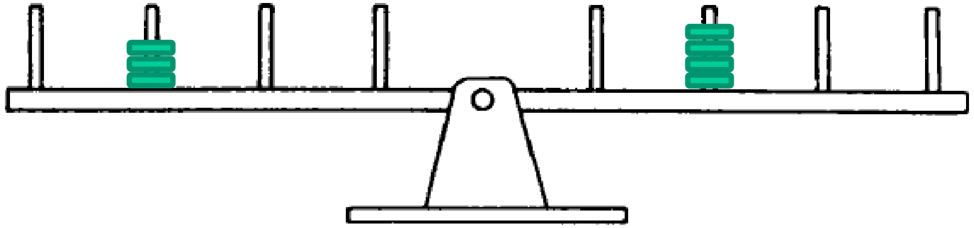

- Stage-like progression:
  - Early incomprehension
  - Systematic errors
  - Gradual progression to approximate intuitive mastery

- Readiness to learn:
  - Gradual developmental change creates conditions for fast learning
  - Consistent new knowledge can be added to existing knowledge easily
  - Inconsistent new knowledge is much harder to learn

An extended exploration of these ideas

- Human semantic cognition
  - Progressive differentiation
  - U-shaped developmental change
  - Reorganization of conceptual knowledge
  - Acquired domain-specific inductive biases through domain general error-correcting learning

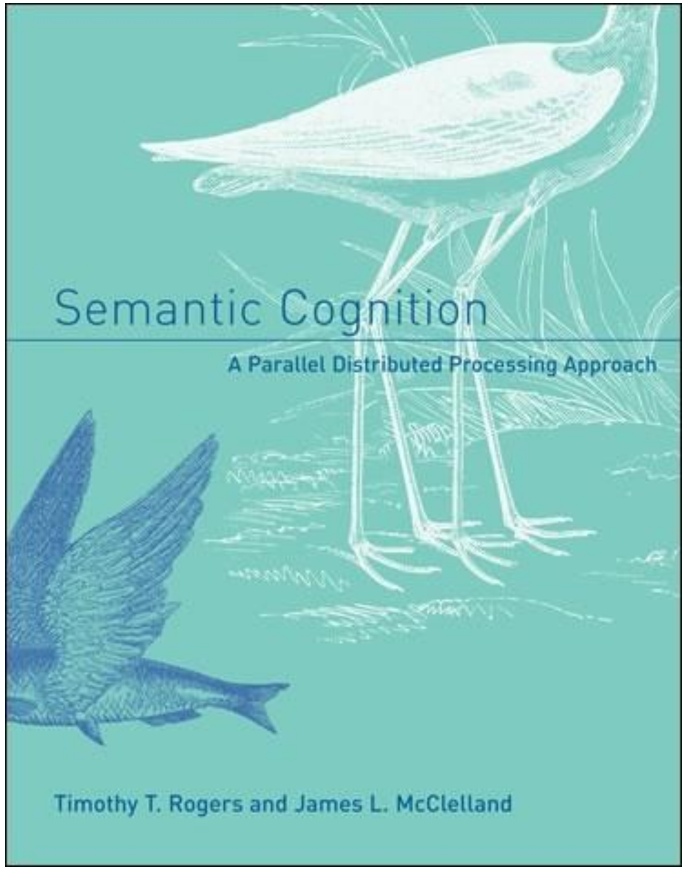

Shortcomings of these models

- Toy vs real problems
- Abstract vs grounded inputs and outputs
- Early-acquired intuitions vs advanced cognitive abilities
- Do not exhibit exploration and discovery

- Do not generalize beyond a specific task space
- Do not exhibit explicit understanding
- Do not benefit from explicit instruction

- Wherein lie the solutions?

DeepMind provides key aspects of the solution

Real Progress - But Challenges Remain

- ✓ Real problems
- ✓ Grounded inputs and outputs
- ✓ Advanced cognitive abilities
- ✓ Exhibit exploration and discovery

- Do not generalize beyond a specific task space
- Do not exhibit explicit understanding
- Do not benefit from explicit instruction

- Wherein lie the solutions?

My Plan

- Focus on human mathematical cognition
  - From number, to arithmetic, algebra, geometry, and beyond
- Apply the principles and resources of DeepMind
- Extend them to address the challenges

Why Mathematical Cognition?

A Human Solves a Geometry Proof Problem

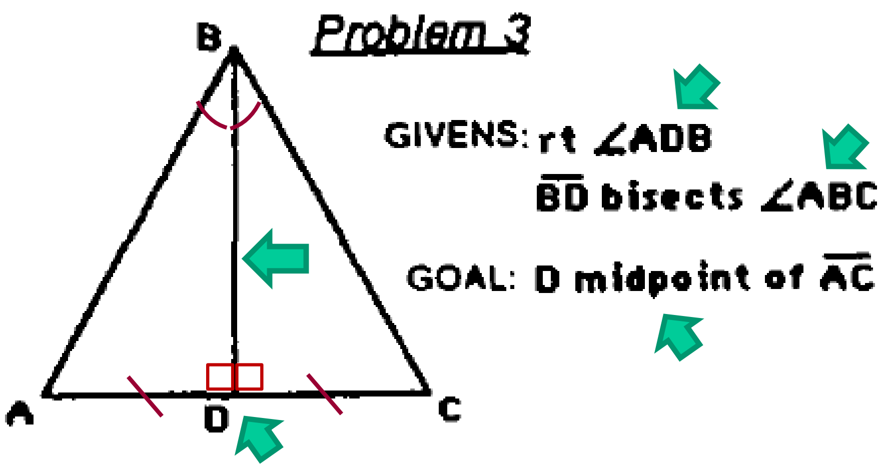

We're given a right angle.
This is a right angle,
perpendicular on both sides
BD bisects angle ABC
And we're done.

We know that *this* is reflexive
We know we have corresponding triangles;
we can determine anything from there
in terms of corresponding parts
And that's what *this* is going to mean…
that *these* are congruent

Koedinger & Anderson, *Cognitive Science*, 1992

Why mathematical cognition?

- Mathematical discoveries are among the highest achievements of human thought
- Mathematical is often hard to learn, yet mastery leads to powerful capabilities
- Some view math as strictly formal, but grounding, intuition and insight play central roles
  - Mathematics can be concretely grounded while still obeying formal rules
  - Grounding facilitates understanding and transfer
- Mathematics includes justification and explanation as well as formal procedures
- Stage-like transitions, and readiness to learn, arise at every step
- The main thread of a mathematics curriculum, from counting to geometry to calculus, appears tractable to explore, since its core contents and grounding structures are circumscribed
- Should complement the existing *Theorem Proving and Maths* effort at DeepMind

What if we could…

- Create a simulated agent that would learn mathematics cumulatively, in a virtual environment, up through the basics of calculus, such that the agent could:
  - Solve novel problems
  - Learn new extensions of its skills quickly through demonstration, explanation and discovery
  - Explain and justify its solutions

Two Conjectures

1. Solving advanced cognitive abilities will depend on building systems that rely on the DeepMind philosophy

   – Learned not Hard-coded
   – General not Specific
   – Grounded not Logic-Based
   – Active not Passive

2. But that alone is not enough

   – Advanced cognitive abilities depend on culturally-constructed tools that leverage human thought

Some technological and conceptual tools for thought

- Writing systems and durable media
  - Chalk and tablet, pencil and paper
  - Electronic documents and editors
- Diagrams and number systems
  - Straight edge and compass,
  - ruler and protractor
  - Place value systems and the abacus
- Inference systems
  - Number and arithmetic,
    algebra and geometry,
    logic and systems of mathematical proof,
    computer programming languages

Human External Memory

External Memory We Compute On

Formal Systems that Structure Human Thought

My bottom line: the basis of insight and discovery

- Intuition and acquired systematic thinking abilities are required for deep understanding and discovery
- Powerful computer-based mathematical processing systems exist but
  - They were programmed by humans
  - They lack insight and intuition
- Modeling the human case will give us insight into how to build artificial systems that incorporate these human-like qualities

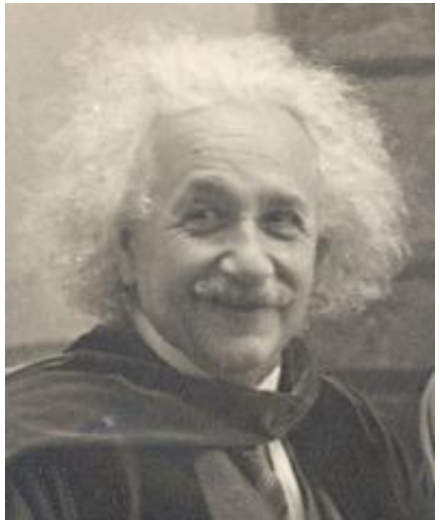

Einstein in Academic Regalia

Understanding the Counting Numbers: A Sudden Discovery or a Gradually Emerging Cognitive Skill?

- Reciting the count list
  - Often learned early, but without apparent meaning
- Answering “how many?”
  - first for sets of 1, then up to 2, then up to 3 or 4 items
- A sudden discovery?
  - If a child tends to succeed for 5, they’ll tend to succeed with larger numbers (though errors continue to occur)
- However the child may fail many other tasks
  - Give-a-number task
  - Which is more task
  - The ‘remove an object and replace it with another’ task

The Long Term Goal

- Understand how expert mathematicians acquire the ability to combine flashes of intuition with the ability to engage in abstract mathematical reasoning

- It remains to be understood how humans acquire the ability to think at such a level

  - They may have to learn to do it over and over as they master more and more advanced domains

  - But I think they can at least achieve a readiness to progress to the next level quickly

- I look forward to thinking and working on these questions with you.