CSE 40171:
Artificial Intelligence

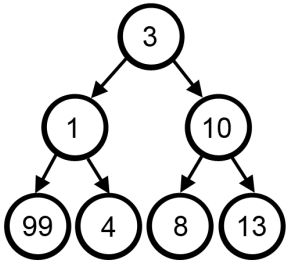


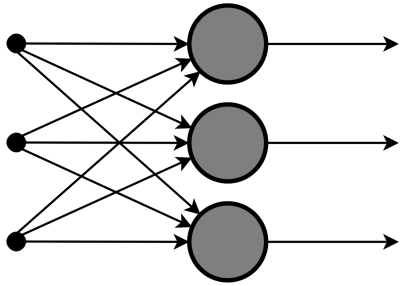

The Brain: Cognition

Conscious vs. Unconscious Thinking

University of Notre Dame's God Quad

close up of a white king on a chess board

What's the difference between viewing a scene and playing a game?

Bottom-Up Strategies

- Piecing together smaller systems to give rise to a larger system
- Information may be linked in a hierarchical fashion
- Physical characteristics of a stimulus drive perception

Top-Down Strategies

- Knowledge, expectations and reflection influence decisions
- Constructivist view: “Perception is not determined simply by stimulus patterns; rather it is a dynamic searching for the best interpretation of the available data.” (Gregory, 1966)

How does the mind get so much from so little?

“These are goats”

Few-shot Learning

Tenenbaum et al. Science 2011

Problem of Induction

Why do we believe the conclusions of arguments?

- Weak: "all swans I have seen are white, and, therefore, all swans are white”
- Strong: “the laws of physics will hold as they have always been observed to hold”

Three central questions posed by Tenenbaum et al.

1. How does abstract knowledge guide learning and inference from sparse data?

2. What forms does abstract knowledge take, across different domains and tasks?

3. How is abstract knowledge itself acquired?

Nativism

Chomsky:

“Intrinsic (psychological) structure is rich . . . and diverse.”

“We may usefully think of the language faculty, the number faculty, and others as 'mental organs,' analogous to the heart or the visual system or the system of motor coordination and planning. There appears to be no clear demarcation line between physical organs, perceptual and motor systems and cognitive faculties in the respects in question.”

N. Chomsky, Rules and Representations, Oxford 1980

Nativism

Arguments against nativism:

Significant of evidence for learning, especially during development.

Information capacity of the genome is, relatively speaking, small.

Connectionism

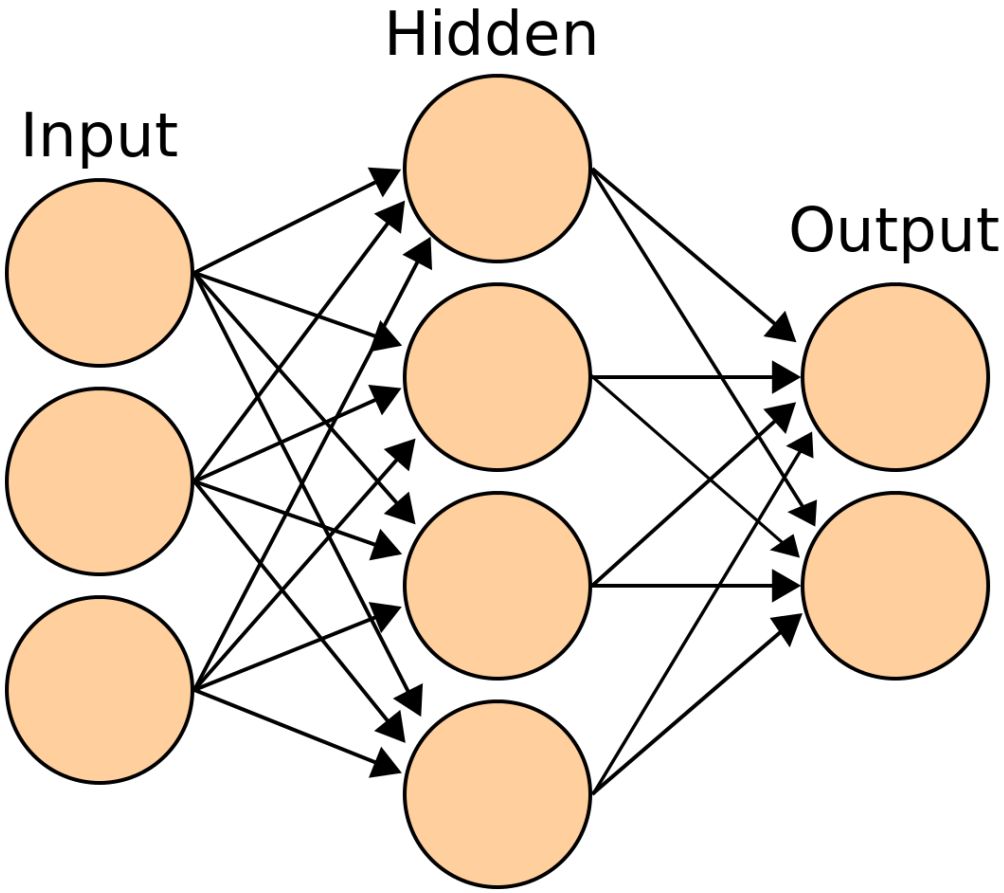

J. McClelland, D. Rumelhart, Eds., Parallel Distributed Processing: Explorations in the Microstructure of Cognition, MIT Press, 1986

Constructivism

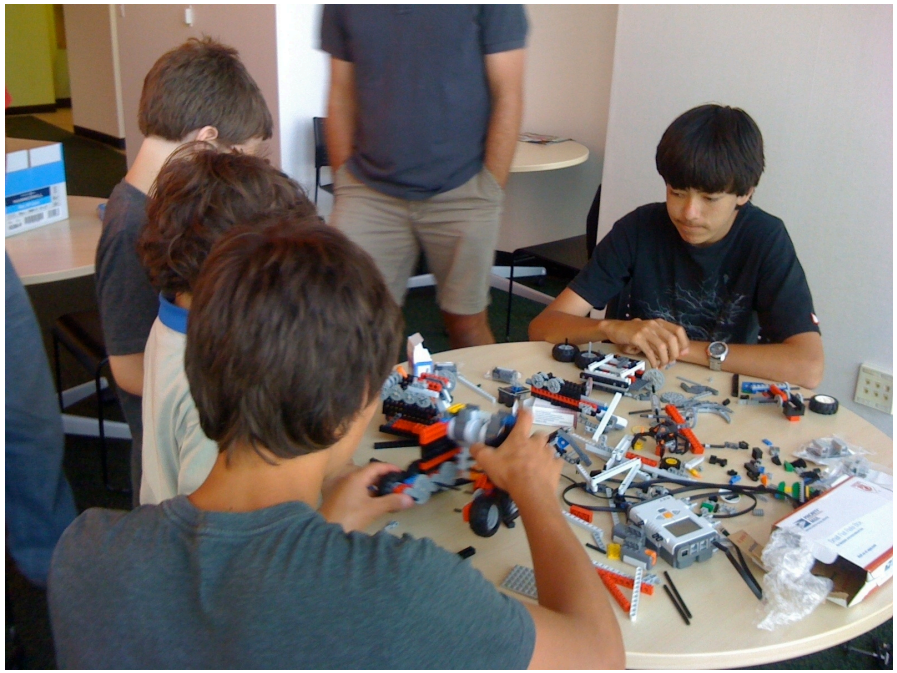

hands on learning

- Less formal than connectionism and other theories
- Learners construct knowledge through experience

Bayesian Inference: An explanatory model of the brain **and** a model for machine learning

Aspects of higher-level cognition that can be modeled this way:

- ✓ sense of similarity
- ✓ representativeness
- ✓ randomness
- ✓ coincidences as a cue to hidden causes
- ✓ judgments of causal strengths
- ✓ judgements of evidential support
- ✓ diagnostic and conditional reasoning
- ✓ predictions about the future of everyday events

An example

Why is John coughing?

h_1: John has a cold

h_2: John has lung cancer

h_3: John has heartburn

An example

The likelihood favors h_1 and h_2

The prior, in contrast, favors h_1 and h_3 over h_2

Bayes's rule thus favors h_1 as an explanation, because it scores highly on both terms

Learning from Sparse Data

h_1: "horse"

h_2: "all horses except Clydesdales"

h_3: "all animals"

Bayesian inference over a tree-structured domain representation

Tenenbaum et al. Science 2011

Form of Abstract Knowledge

- Trees are only one option
  - Two-dimensional spaces or grids
  - One-dimensional orders
  - Directed networks for causally transmitted properties

Learning a tree of animals

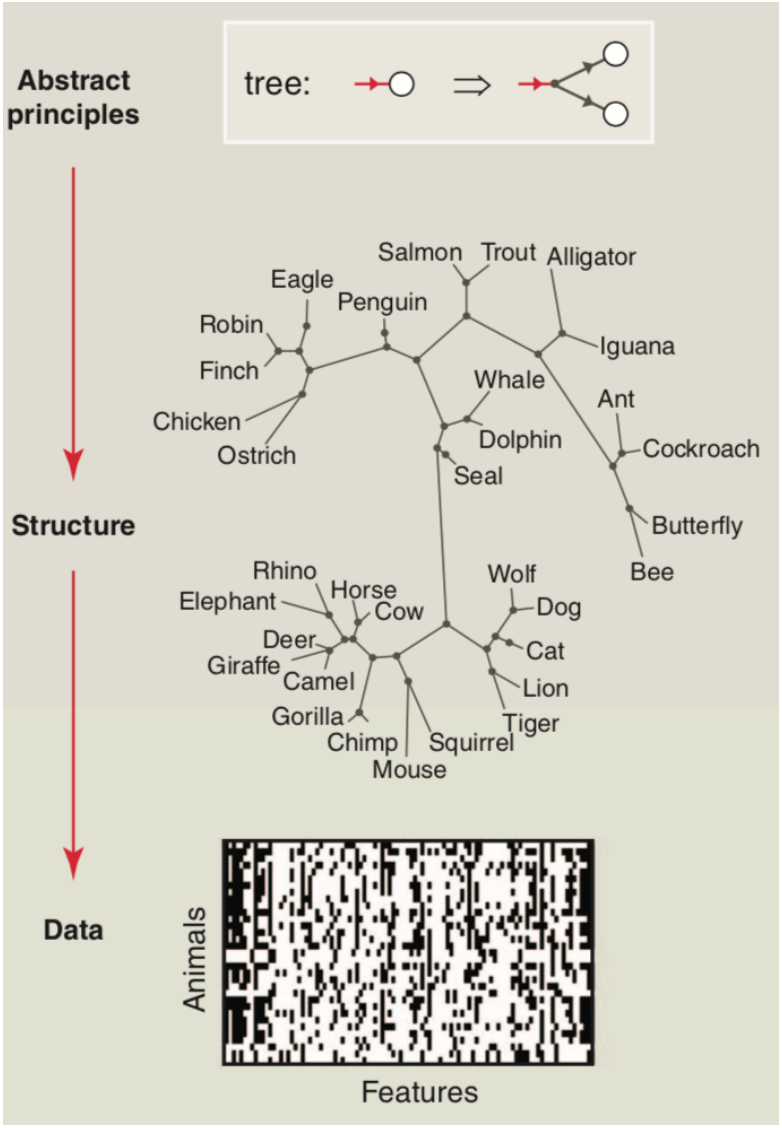

Tenenbaum et al. Science 2011

Origins of Abstract Knowledge

How do learners learn what they need to know to make learning possible?

How does a child know which tree structure is the right way to organize hypotheses for word learning?

How can a learner know that a given domain of entities and concepts should be represented by using a tree at all?

How do people come to correct framework theories with causal links?

Hierarchical Bayesian models

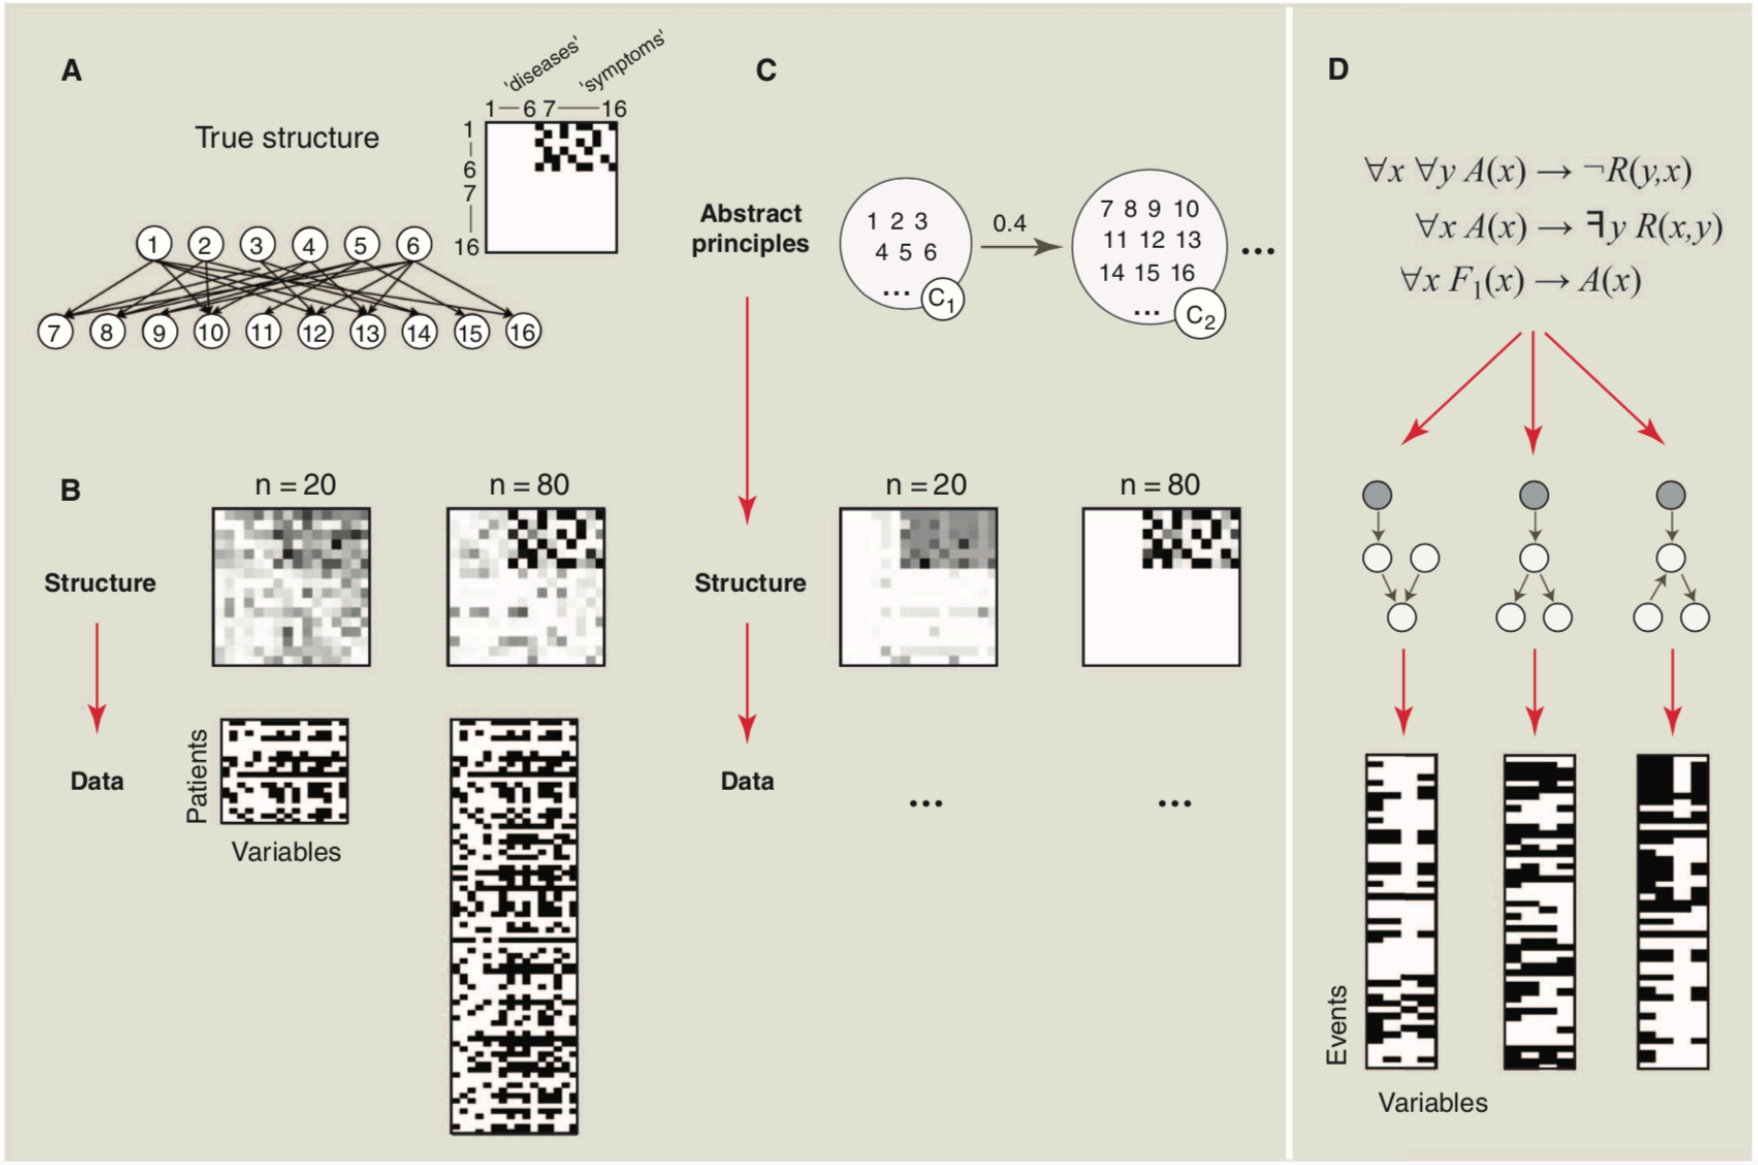

Tenenbaum et al. Science 2011

Direct application to Machine Learning

HBMs have been used for:

Transfer learning (Kemp et al. Dev. Sci. 2007)

Discovery of the form of structure governing similarity (Kemp et al. Psychol. Rev. 2009)

Learning of abstract causal knowledge (Kemp et al. Cognition 2010)

Open Questions

- How does it all start?
- What about hard tasks in cognitive development?
- The process likely requires multiple levels of analysis — what are the rest of the pieces?
- How are these principles implemented in neural circuits?