The child as scientist

Learning as “theory building”, not “data analysis”.
Knowledge grows through hypothesis- and explanation-driven interpretations of sparse data, causal learning, learning theories, learning compositional abstractions, learning to learn, exploratory learning, social learning.
[Carey, Karmiloff-Smith, Gopnik, Schulz, Feigenson…]

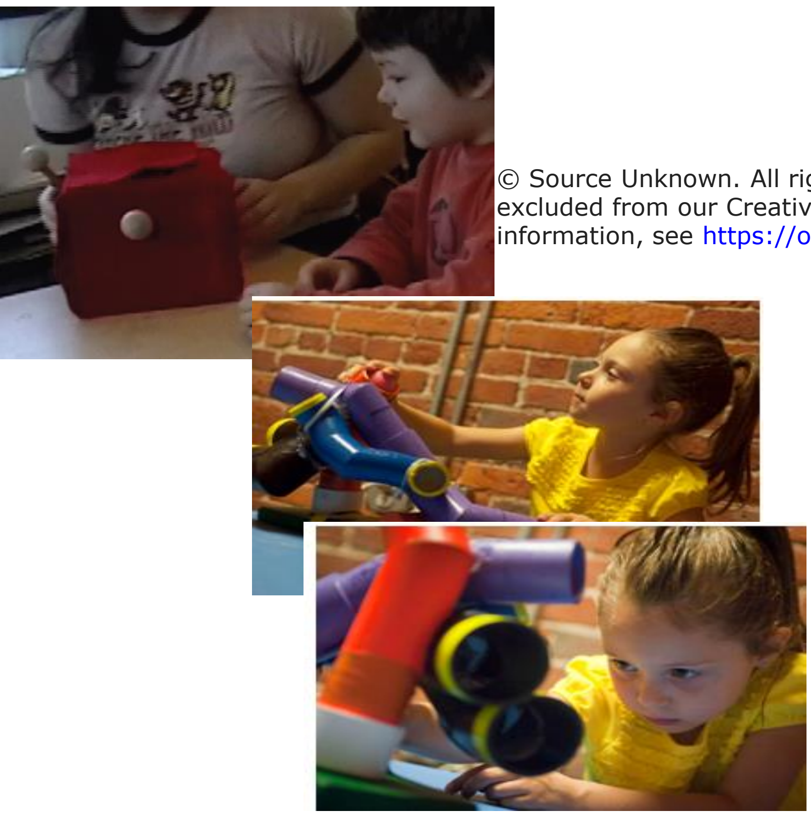

© Source Unknown. All rights reserved. This content is excluded from our Creative Commons license. For more information, see https://ocw.mit.edu/help/faq-fair-use/.

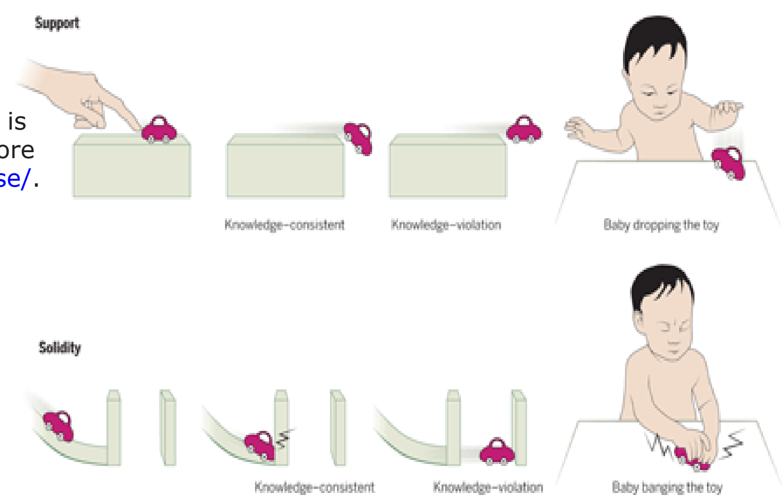

© AAAS. All rights reserved. This content is excluded from our Creative Commons license. For more information, see https://ocw.mit.edu/help/faq-fair-use/.

© Patrick Gillooly / MIT News. All rights reserved.This content is excluded from our Creative Commons license. For more information, see https://ocw.mit.edu/help/faq-fair-use/.

Probabilistic programs for model building (“program-learning” programs)

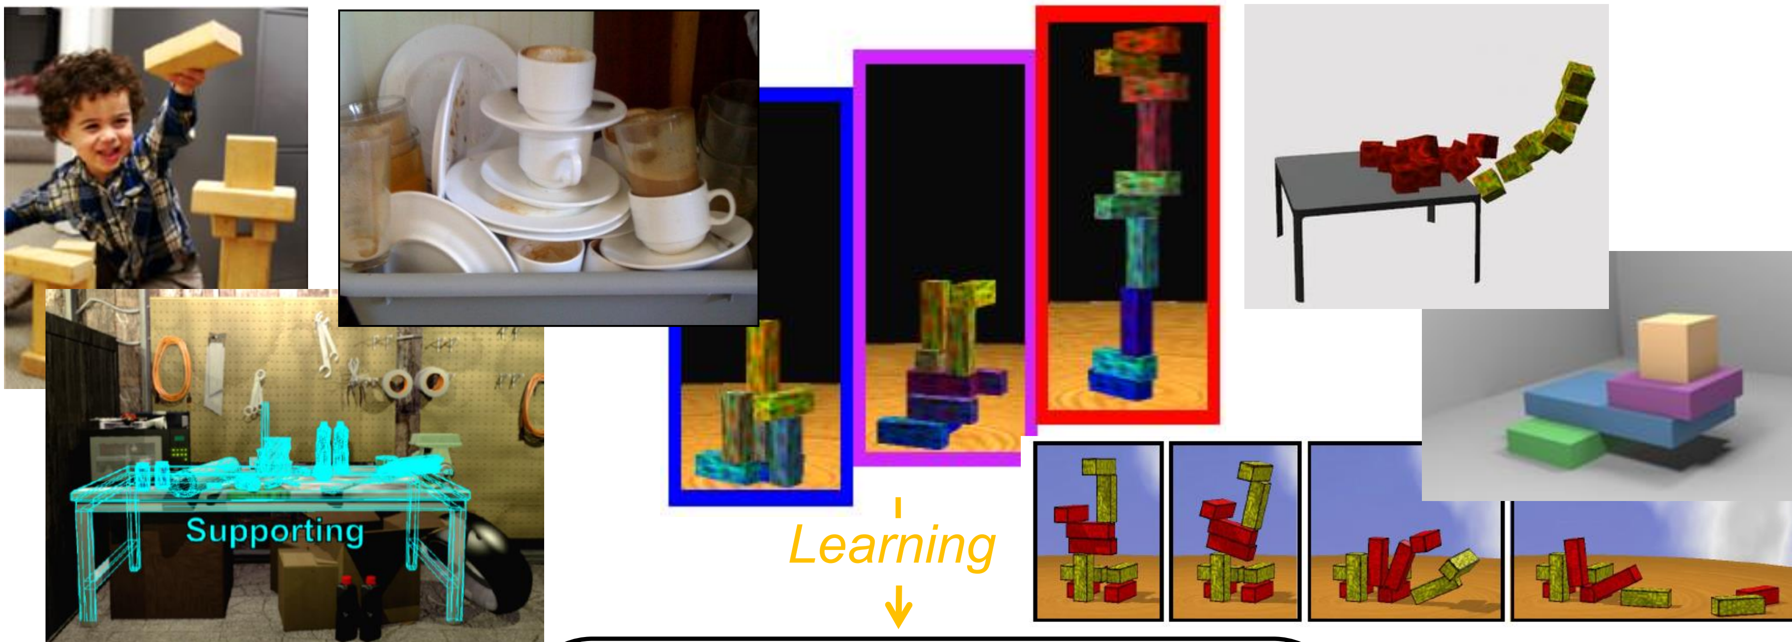

Learning

physics

World state (t) → World state (t+1) . . .

graphics

Image (t)   Image (t+1)

© Proceedings of the National Academy of Sciences. All rights reserved. This content is excluded from our Creative Commons license. For more information, see https://ocw.mit.edu/help/faq-fair-use/.
Source: Battaglia, Peter W., Jessica B. Hamrick, and Joshua B. Tenenbaum. "Simulation as an engine of physical scene understanding." Proceedings of the National Academy of Sciences 110, no. 45 (2013): 18327-18332.

Learning and generalization for object concepts

© AAAS. All rights reserved. This content is excluded from our Creative Commons license. For more information, see https://ocw.mit.edu/help/faq-fair-use/.
Source: Tenenbaum, Joshua B., Charles Kemp, Thomas L. Griffiths, and Noah D. Goodman. "How to grow a mind: Statistics, structure, and abstraction." Science 331, no. 6022 (2011): 1279-1285.

Learning and generalization for object concepts

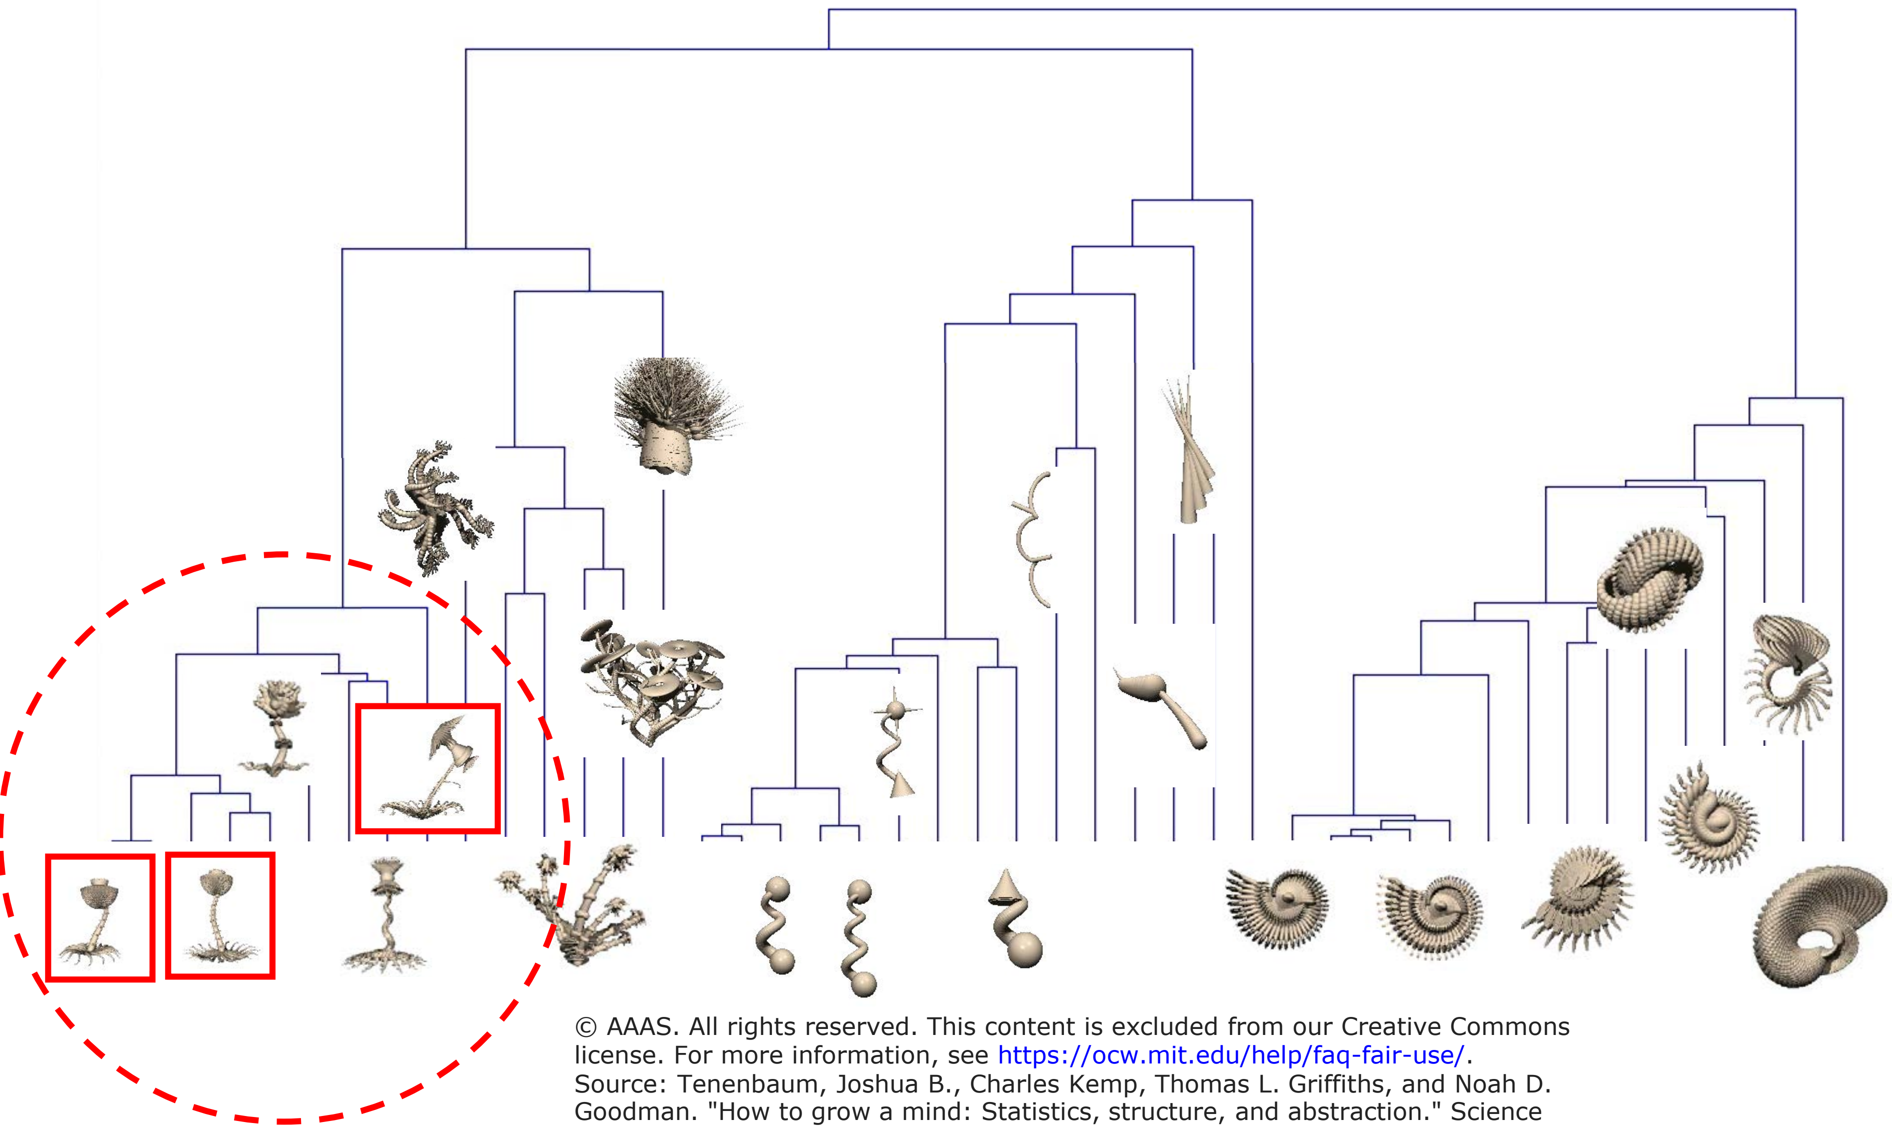

© AAAS. All rights reserved. This content is excluded from our Creative Commons license. For more information, see https://ocw.mit.edu/help/faq-fair-use/.
Source: Tenenbaum, Joshua B., Charles Kemp, Thomas L. Griffiths, and Noah D. Goodman. "How to grow a mind: Statistics, structure, and abstraction." Science 331, no. 6022 (2011): 1279-1285.

Handwritten characters

Standard machine learning: MNIST
100s (or more) examples/class

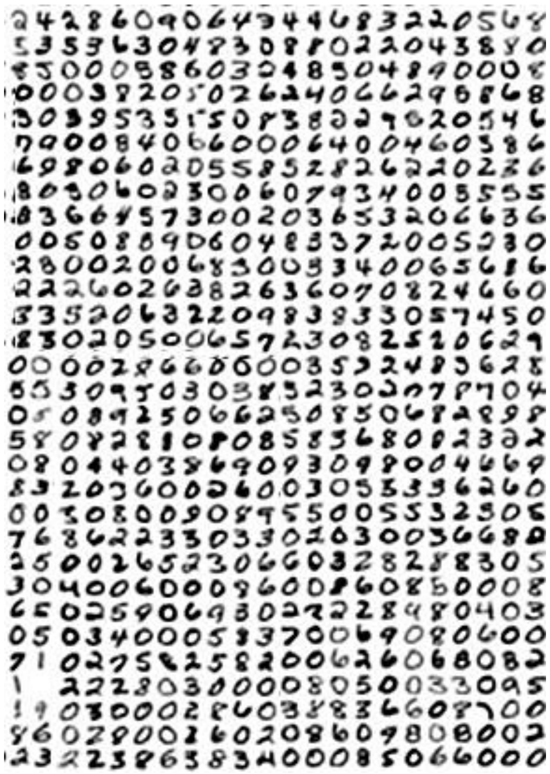

© Mixed National Institute of Standards and Technology. All rights reserved. This content is excluded from our Creative Commons license. For more information, see https://ocw.mit.edu/help/faq-fair-use/.

Our testbed: **Omniglot**
1623 simple visual concepts in 50 alphabets
20 examples/class

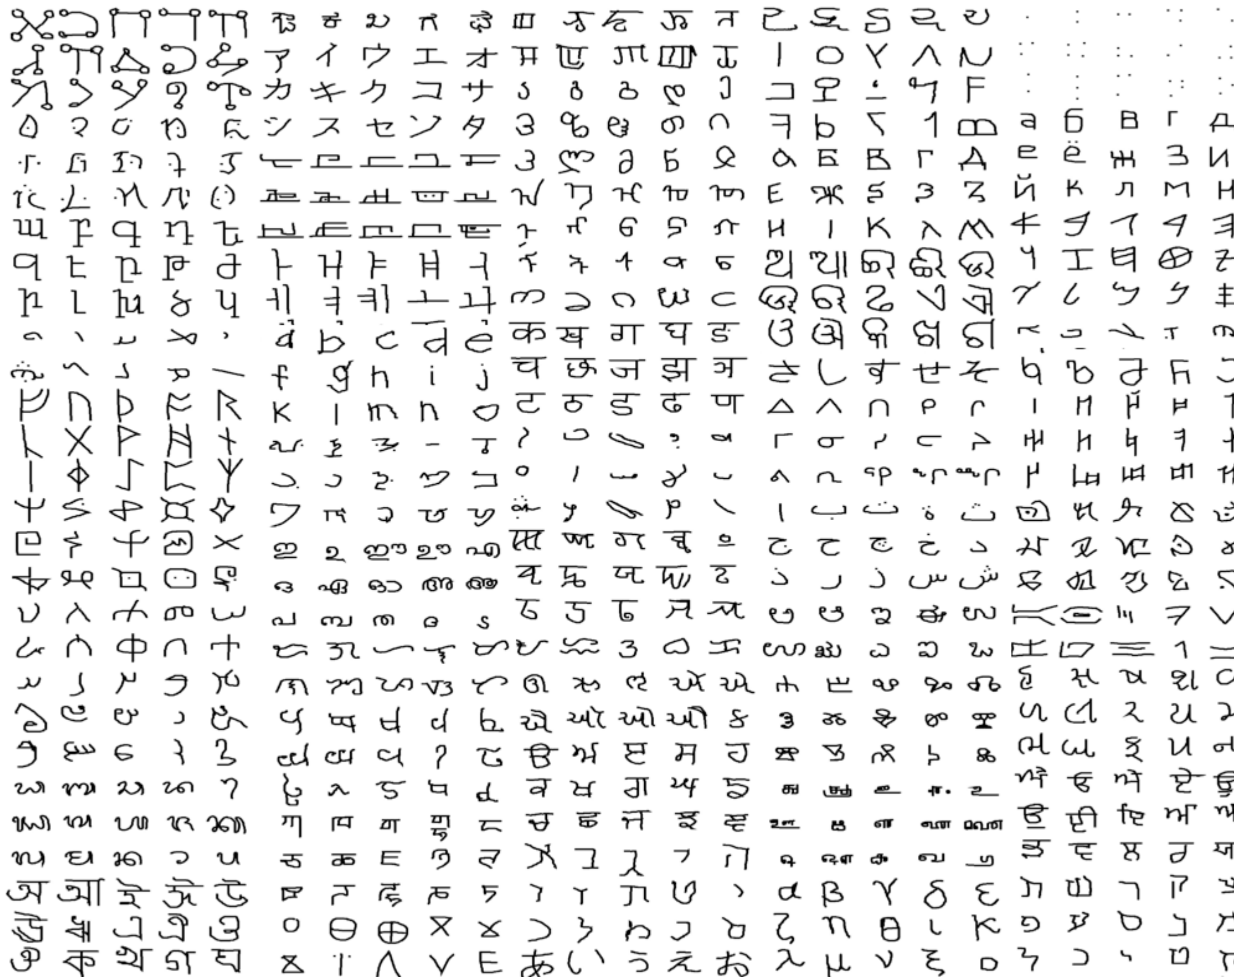

© AAAS. All rights reserved. This content is excluded from our Creative Commons license. For more information, see https://ocw.mit.edu/help/faq-fair-use/.
Source: Lake, Brenden M., Ruslan Salakhutdinov, and Joshua B. Tenenbaum. "Human-level concept learning through probabilistic program induction." Science 350, no. 6266 (2015): 1332-1338.

The omniglot dataset

© AAAS. All rights reserved.This content is excluded from our Creative Commons license. For more information, see https://ocw.mit.edu/help/faq-fair-use/.
Source: Lake, Brenden M., Ruslan Salakhutdinov, and Joshua B. Tenenbaum. "Human-level concept learning through probabilistic program induction." Science 350, no. 6266 (2015): 1332-1338.

The omniglot dataset

Angelic

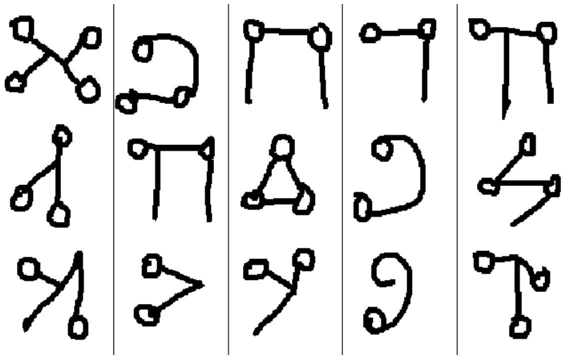

Alphabet of the Magi

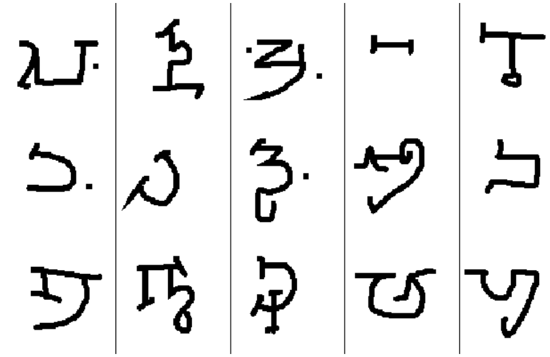

ULOG

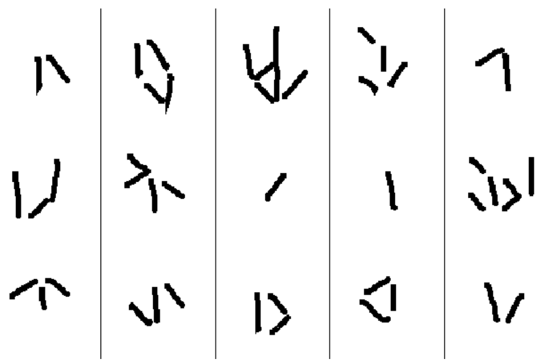

Futurama

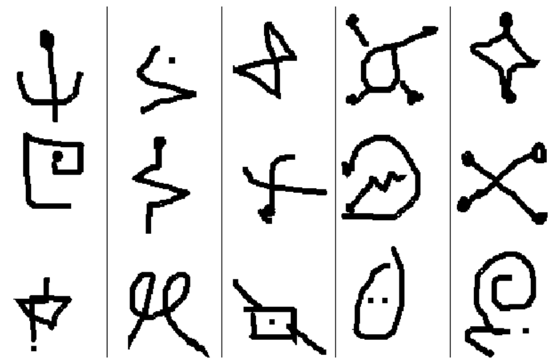

© AAAS. All rights reserved. This content is excluded from our Creative Commons license. For more information, see https://ocw.mit.edu/help/faq-fair-use/.
Source: Lake, Brenden M., Ruslan Salakhutdinov, and Joshua B. Tenenbaum. "Human-level concept learning through probabilistic program induction." Science 350, no. 6266 (2015): 1332-1338.

One-shot learning

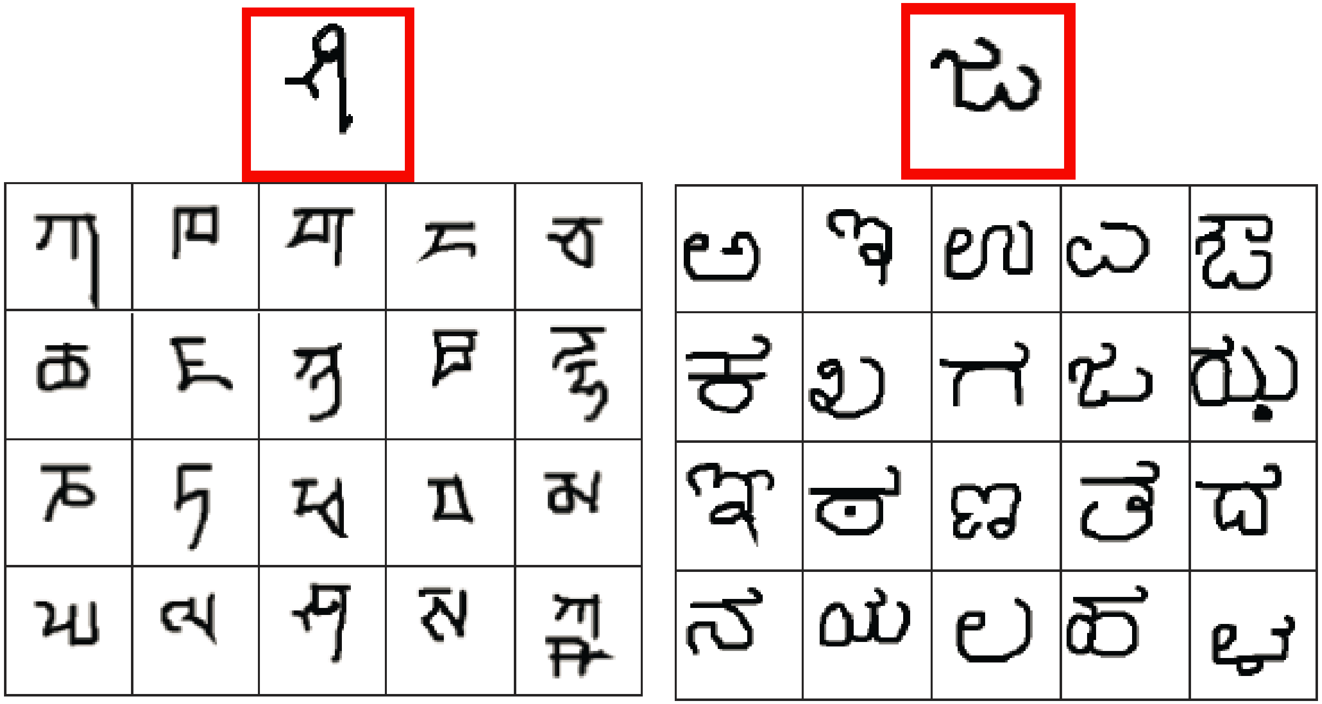

© AAAS. All rights reserved. This content is excluded from our Creative Commons license. For more information, see https://ocw.mit.edu/help/faq-fair-use/.
Source: Lake, Brenden M., Ruslan Salakhutdinov, and Joshua B. Tenenbaum. "Human-level concept learning through probabilistic program induction." Science 350, no. 6266 (2015): 1332-1338.

A multitude of tasks

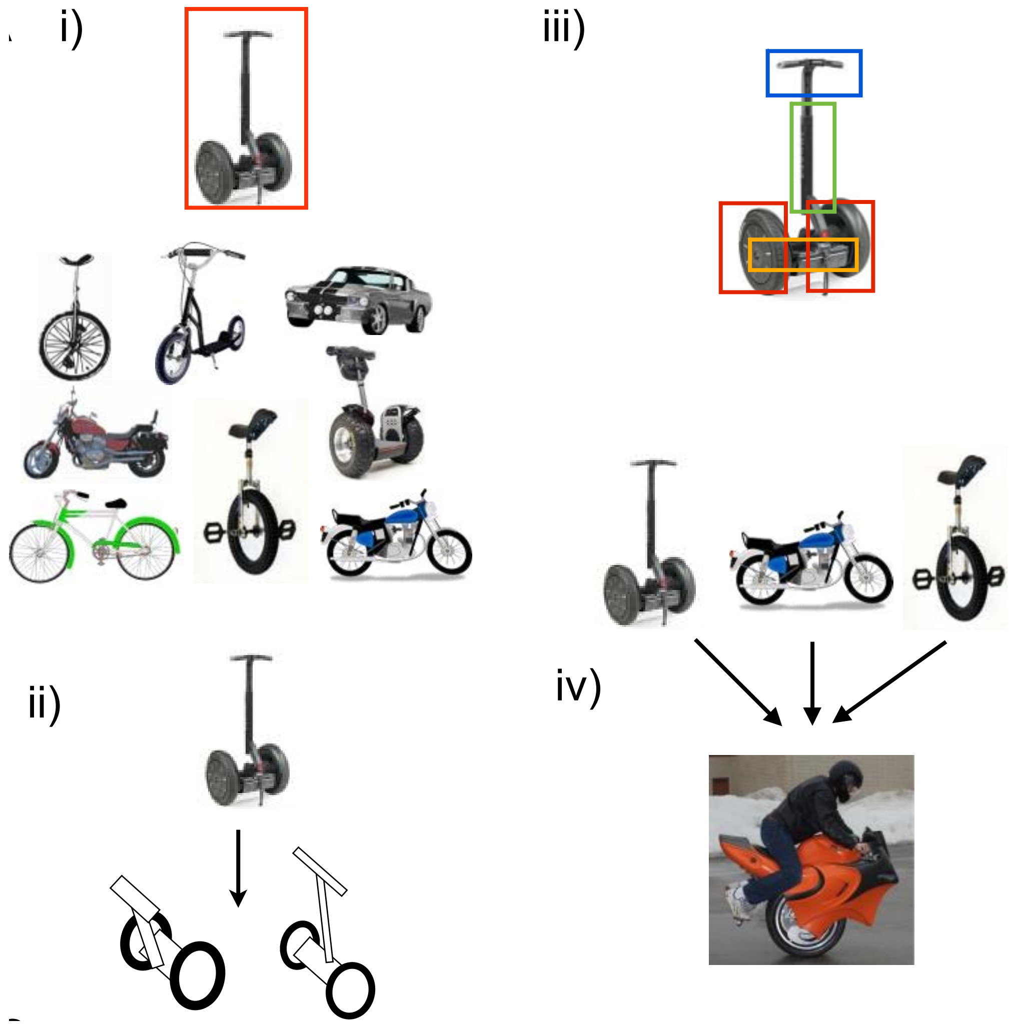

© AAAS. All rights reserved. This content is excluded from our Creative Commons license. For more information, see https://ocw.mit.edu/help/faq-fair-use/.
Source: Lake, Brenden M., Ruslan Salakhutdinov, and Joshua B. Tenenbaum. "Human-level concept learning through probabilistic program induction." Science 350, no. 6266 (2015): 1332-1338.

A multitude of tasks

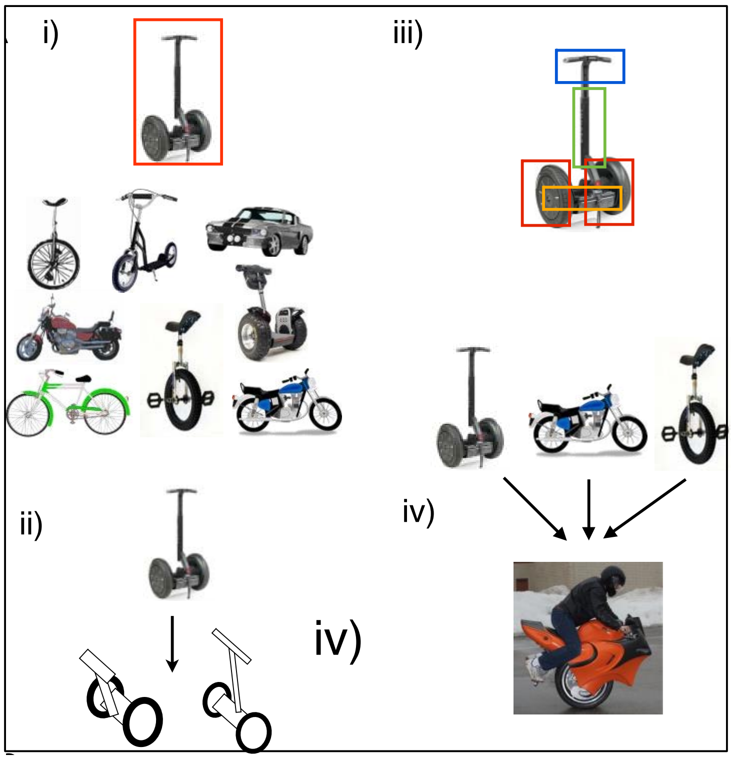

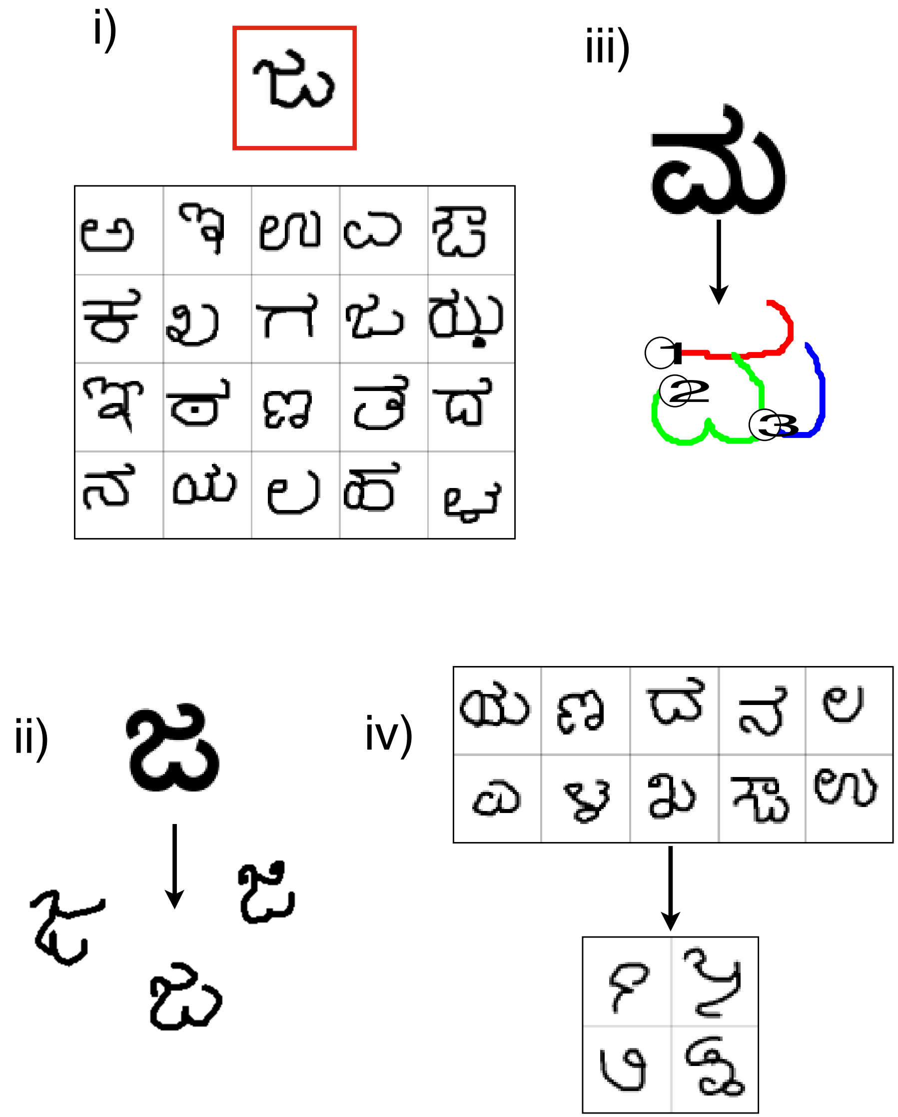

© AAAS. All rights reserved. This content is excluded from our Creative Commons license. For more information, see https://ocw.mit.edu/help/faq-fair-use/.
Source: Lake, Brenden M., Ruslan Salakhutdinov, and Joshua B. Tenenbaum. "Human-level concept learning through probabilistic program induction." Science 350, no. 6266 (2015): 1332-1338.

Bayesian Program Learning
(Lake, Salakhutdinov, Tenenbaum, NIPS 2013; in prep)

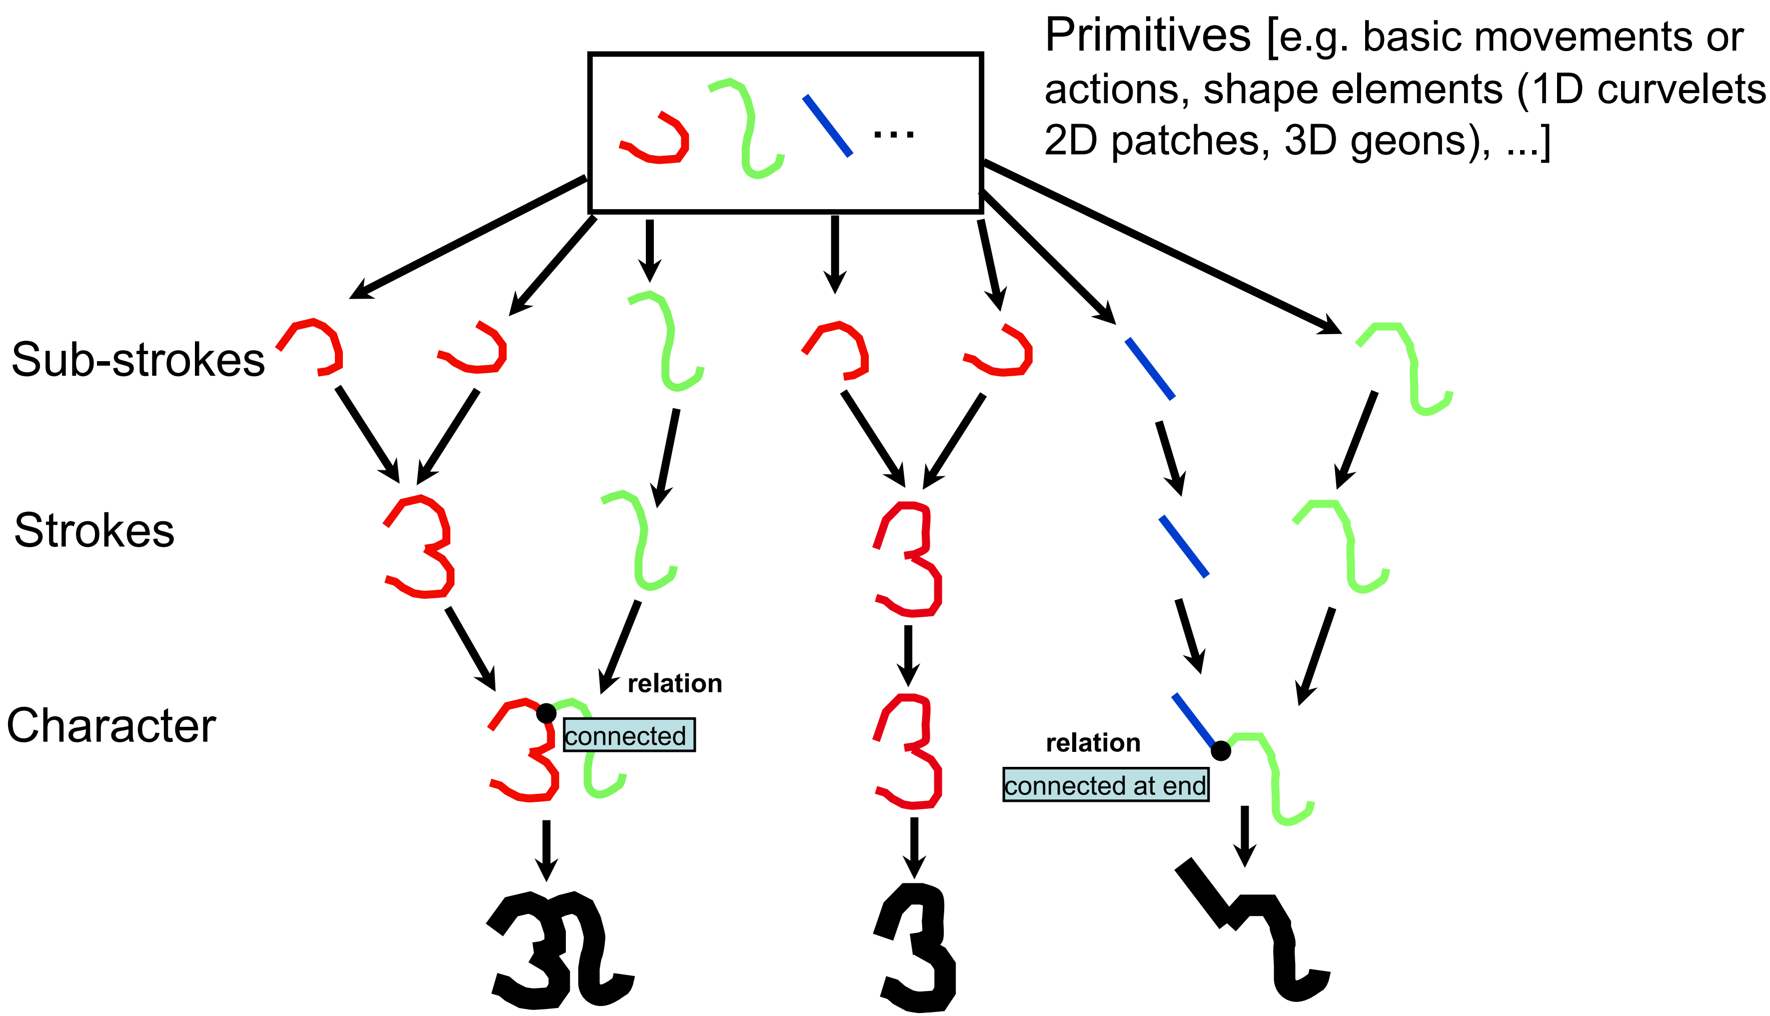

© AAAS. All rights reserved. This content is excluded from our Creative Commons license. For more information, see https://ocw.mit.edu/help/faq-fair-use/.
Source: Lake, Brenden M., Ruslan Salakhutdinov, and Joshua B. Tenenbaum. "Human-level concept learning through probabilistic program induction." Science 350, no. 6266 (2015): 1332-1338.

Bayesian Progr
(Lake, Salakhutdinov, T

procedure GENERATETYPE
- $\kappa \leftarrow P(\kappa)$ — Sample number of parts
- **for** $i = 1 \ldots \kappa$ **do**
  - $z_i \leftarrow P(z_i)$ — Sample sub-parts
  - **for** $j = 1 \ldots n_i$ **do**
    - $x_{ij} \leftarrow P(x_{ij}|z_{ij})$ — Transform sub-parts
  - **end for**
  - $R_i \leftarrow P(R_i|z_1, \ldots, z_{i-1})$
- **end for** — Sample part relations
- $\psi \leftarrow \{\kappa, R, z, x\}$
- **return** @GENERATETOKEN(ψ)

end procedure — Handle to stochastic program

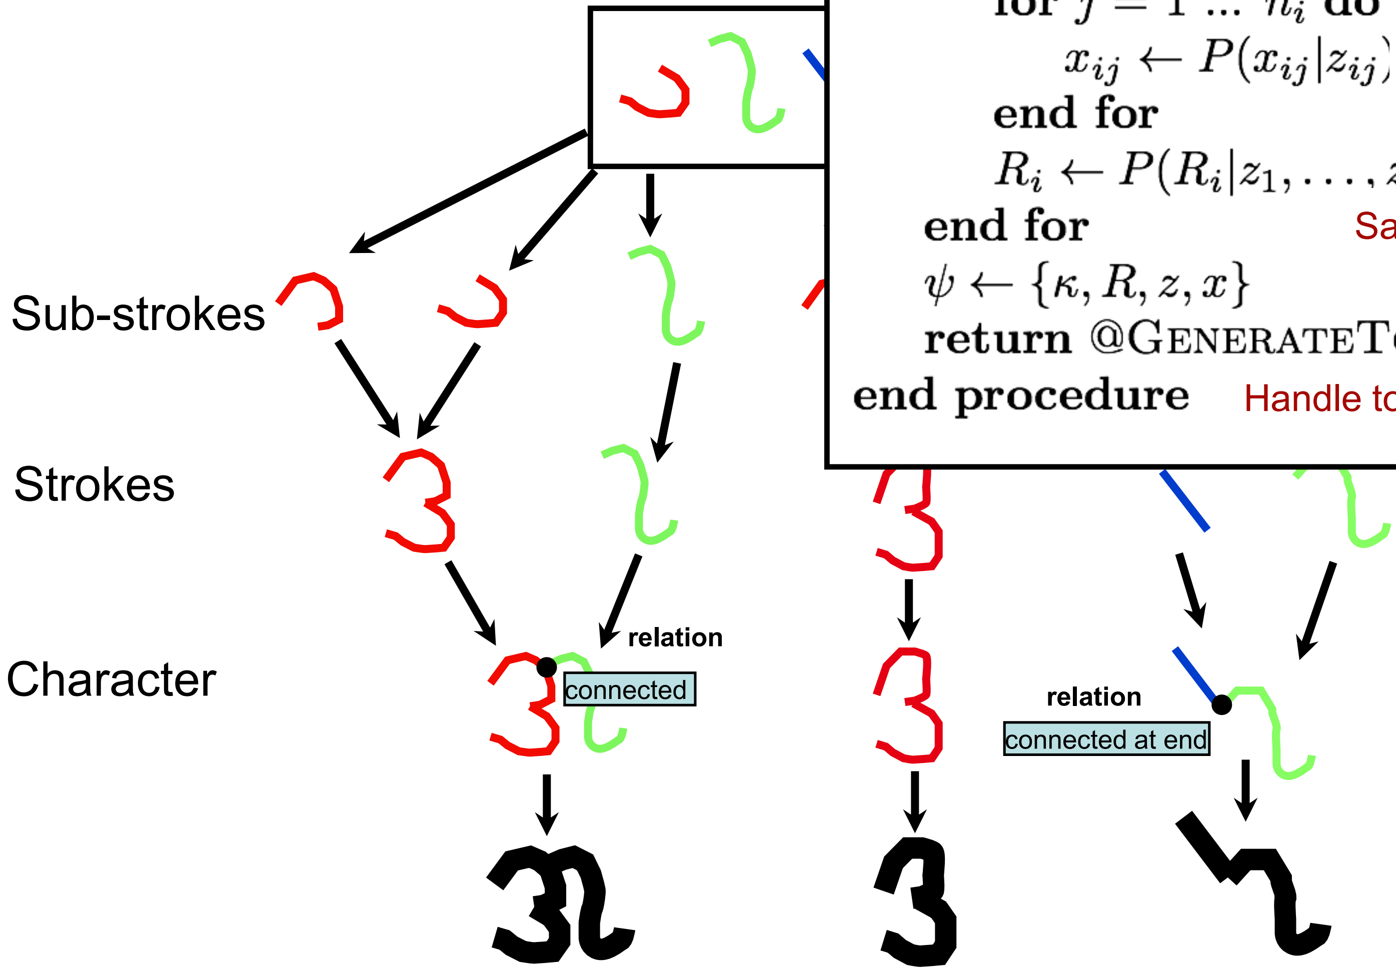

© AAAS. All rights reserved. This content is excluded from our Creative Commons license. For more information, see https://ocw.mit.edu/help/faq-fair-use/.

Source: Lake, Brenden M., Ruslan Salakhutdinov, and Joshua B. Tenenbaum. "Human-level concept learning through probabilistic program induction." Science 350, no. 6266 (2015): 1332-1338.

Learning to generate types

("generative model for generative models")

HBPL (and other models) were trained on 30 "background alphabets" that weren't seen again.

number of strokes

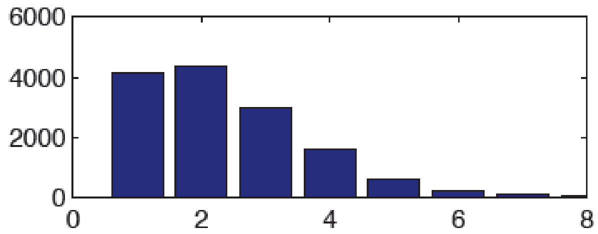

learned (motor) primitives

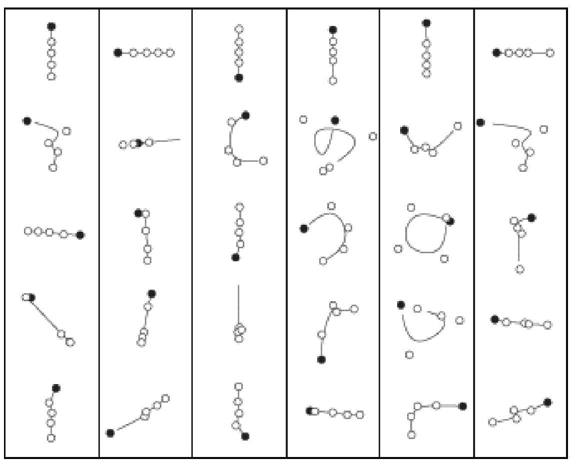

1000 primitives and their bigrams.
transformations: control point variability and scale

relations (stroke attachment)

independent (70%)

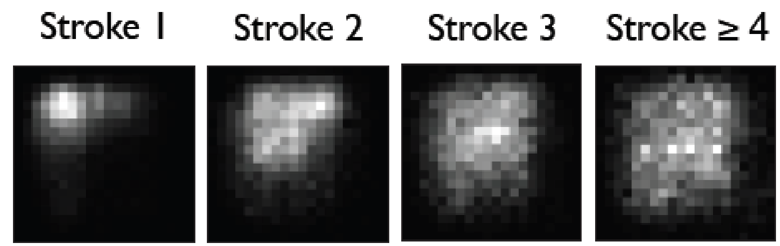

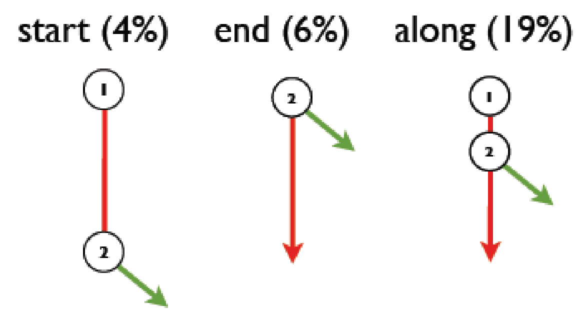

token-level transformations

$$\theta^{(m)} \leftarrow \textsc{GenerateToken}(\psi)$$

- Gaussian noise on continuos variables
- global object scale/translation
- adaptive image blur
- adaptive pixel noise

© AAAS. All rights reserved. This content is excluded from our Creative Commons license. For more information, see https://ocw.mit.edu/help/faq-fair-use/.

Source: Lake, Brenden M., Ruslan Salakhutdinov, and Joshua B. Tenenbaum. "Human-level concept learning through probabilistic program induction." Science 350, no. 6266 (2015): 1332-1338.

Inferring a program from a single example

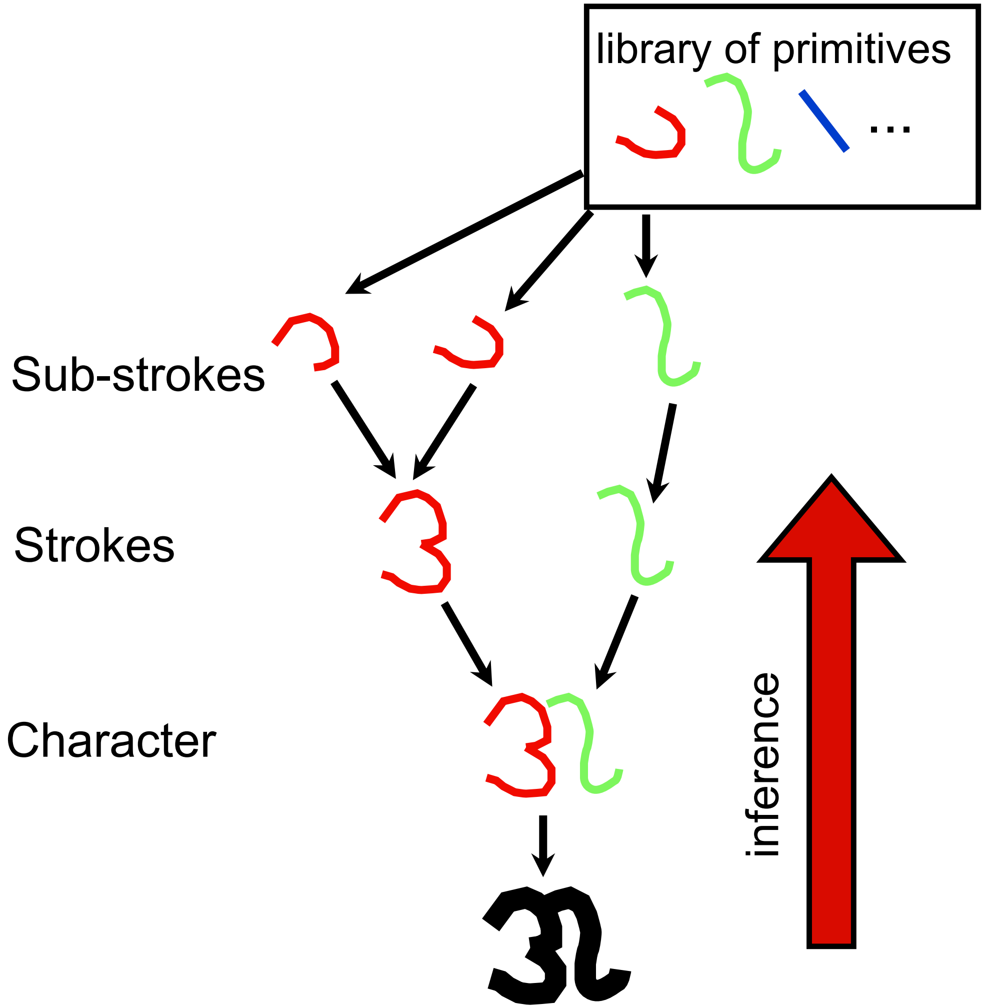

θ latent variables

I image

Discrete approximation to posterior in the form of several parses θ_i.

$$P(\theta|I) \approx \frac{\sum_i w_i 1\{\theta = \theta_i\}}{\sum_i w_i}$$

such that

$$w_i = P(\theta_i|I)$$

Intuition:
Fit strokes to the observed pixels as closely as possible, while also:

- minimizing the number of strokes
- choosing high-probability sub-strokes and maintaining their shape
- choosing stroke start positions that match dataset statistics and abide by stroke relations

© AAAS. All rights reserved. This content is excluded from our Creative Commons license. For more information, see https://ocw.mit.edu/help/faq-fair-use/.
Source: Lake, Brenden M., Ruslan Salakhutdinov, and Joshua B. Tenenbaum. "Human-level concept learning through probabilistic program induction." Science 350, no. 6266 (2015): 1332-1338.

Classifying with probabilistic programs

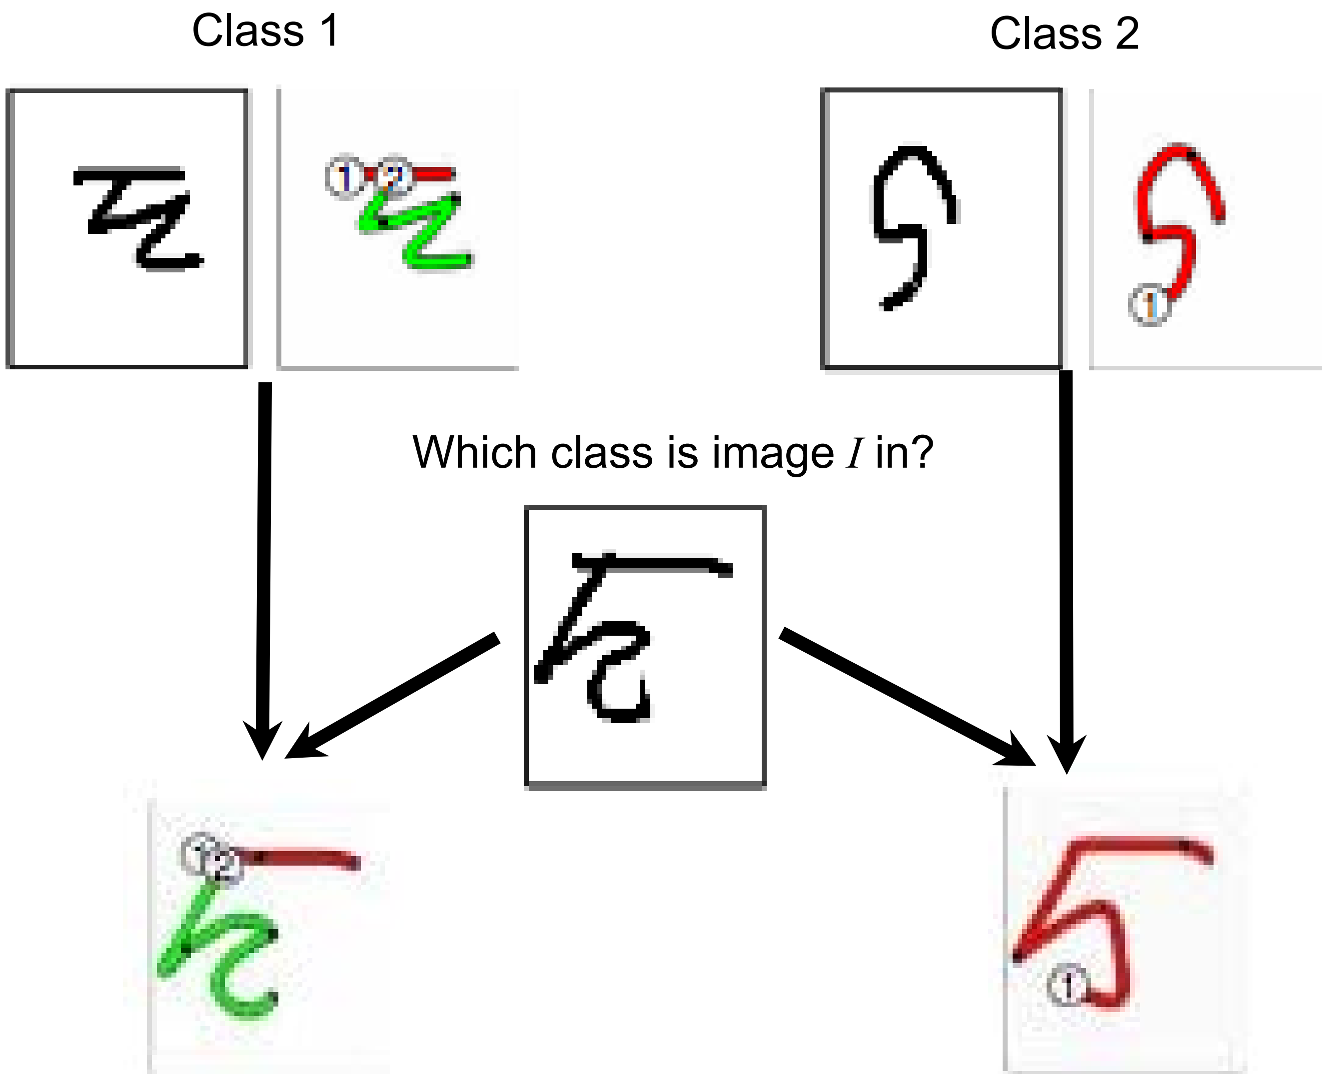

$$\log P(I|\text{class 1}) \approx -758 \qquad \log P(I|\text{class 2}) \approx -1880$$

© AAAS. All rights reserved. This content is excluded from our Creative Commons license. For more information, see https://ocw.mit.edu/help/faq-fair-use/.
Source: Lake, Brenden M., Ruslan Salakhutdinov, and Joshua B. Tenenbaum. "Human-level concept learning through probabilistic program induction." Science 350, no. 6266 (2015): 1332-1338.

One-shot classification

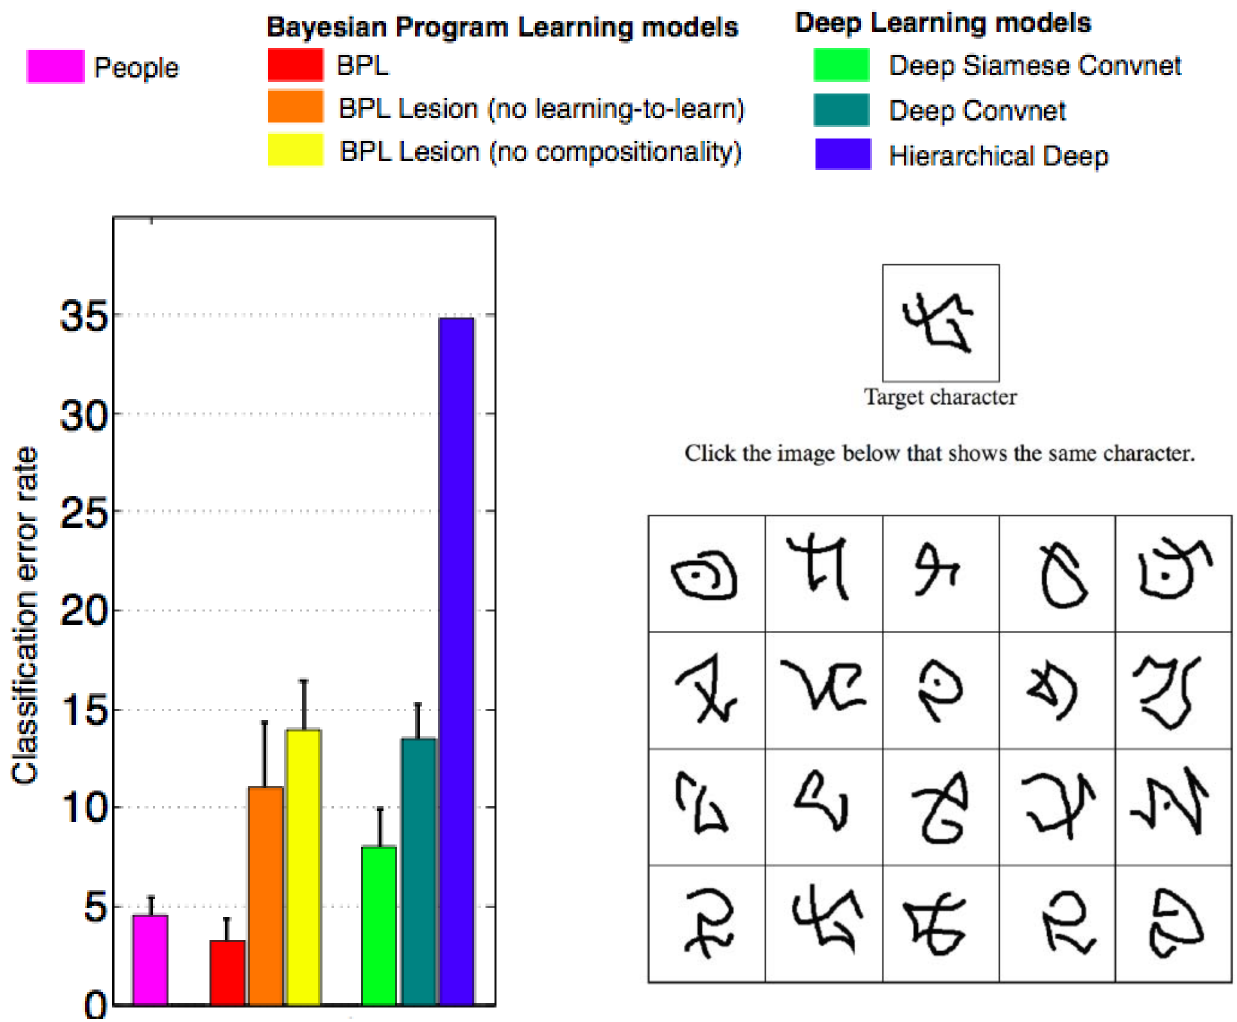

© AAAS. All rights reserved. This content is excluded from our Creative Commons license. For more information, see https://ocw.mit.edu/help/faq-fair-use/.
Source: Lake, Brenden M., Ruslan Salakhutdinov, and Joshua B. Tenenbaum. "Human-level concept learning through probabilistic program induction." Science 350, no. 6266 (2015): 1332-1338.

Generating new examples

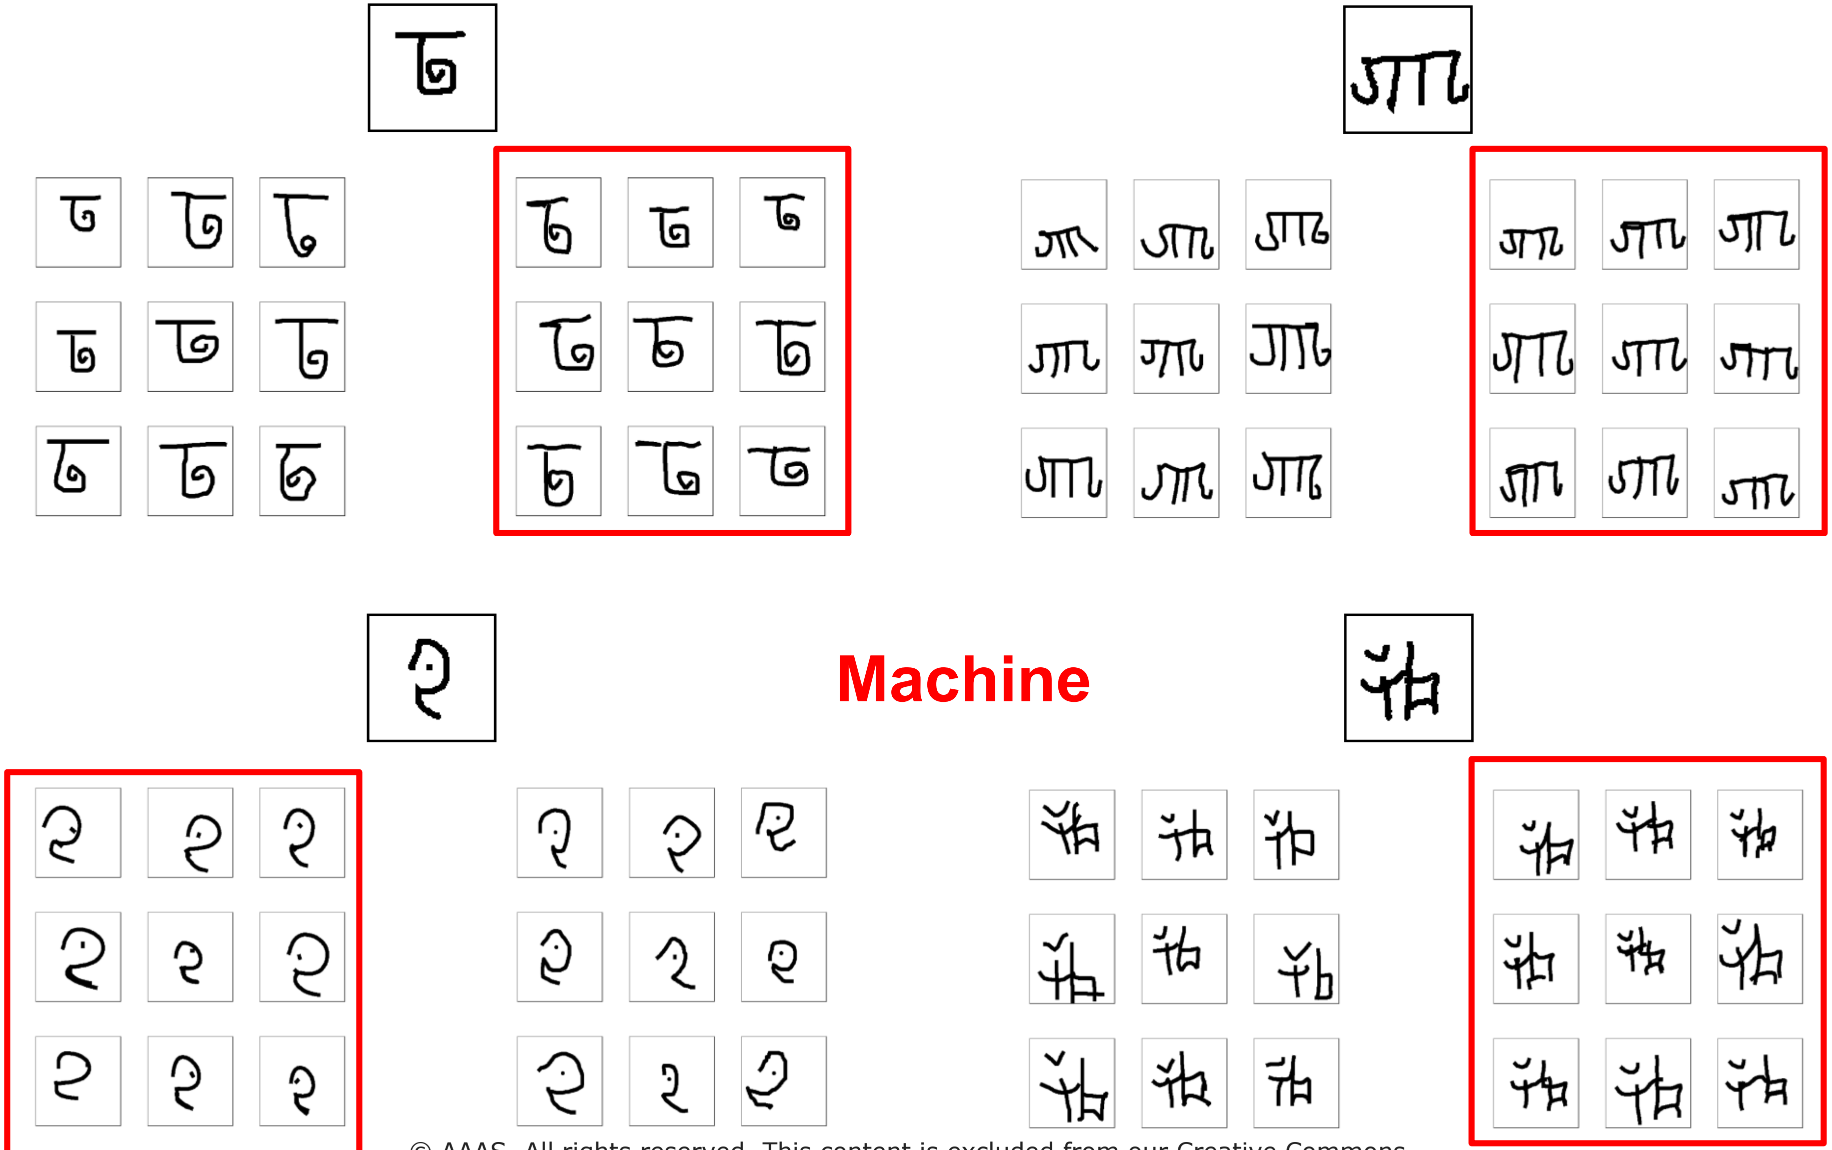

© AAAS. All rights reserved. This content is excluded from our Creative Commons license. For more information, see https://ocw.mit.edu/help/faq-fair-use/.
Source: Lake, Brenden M., Ruslan Salakhutdinov, and Joshua B. Tenenbaum. "Human-level concept learning through probabilistic program induction." Science 350, no. 6266 (2015): 1332-1338.

Generating new examples

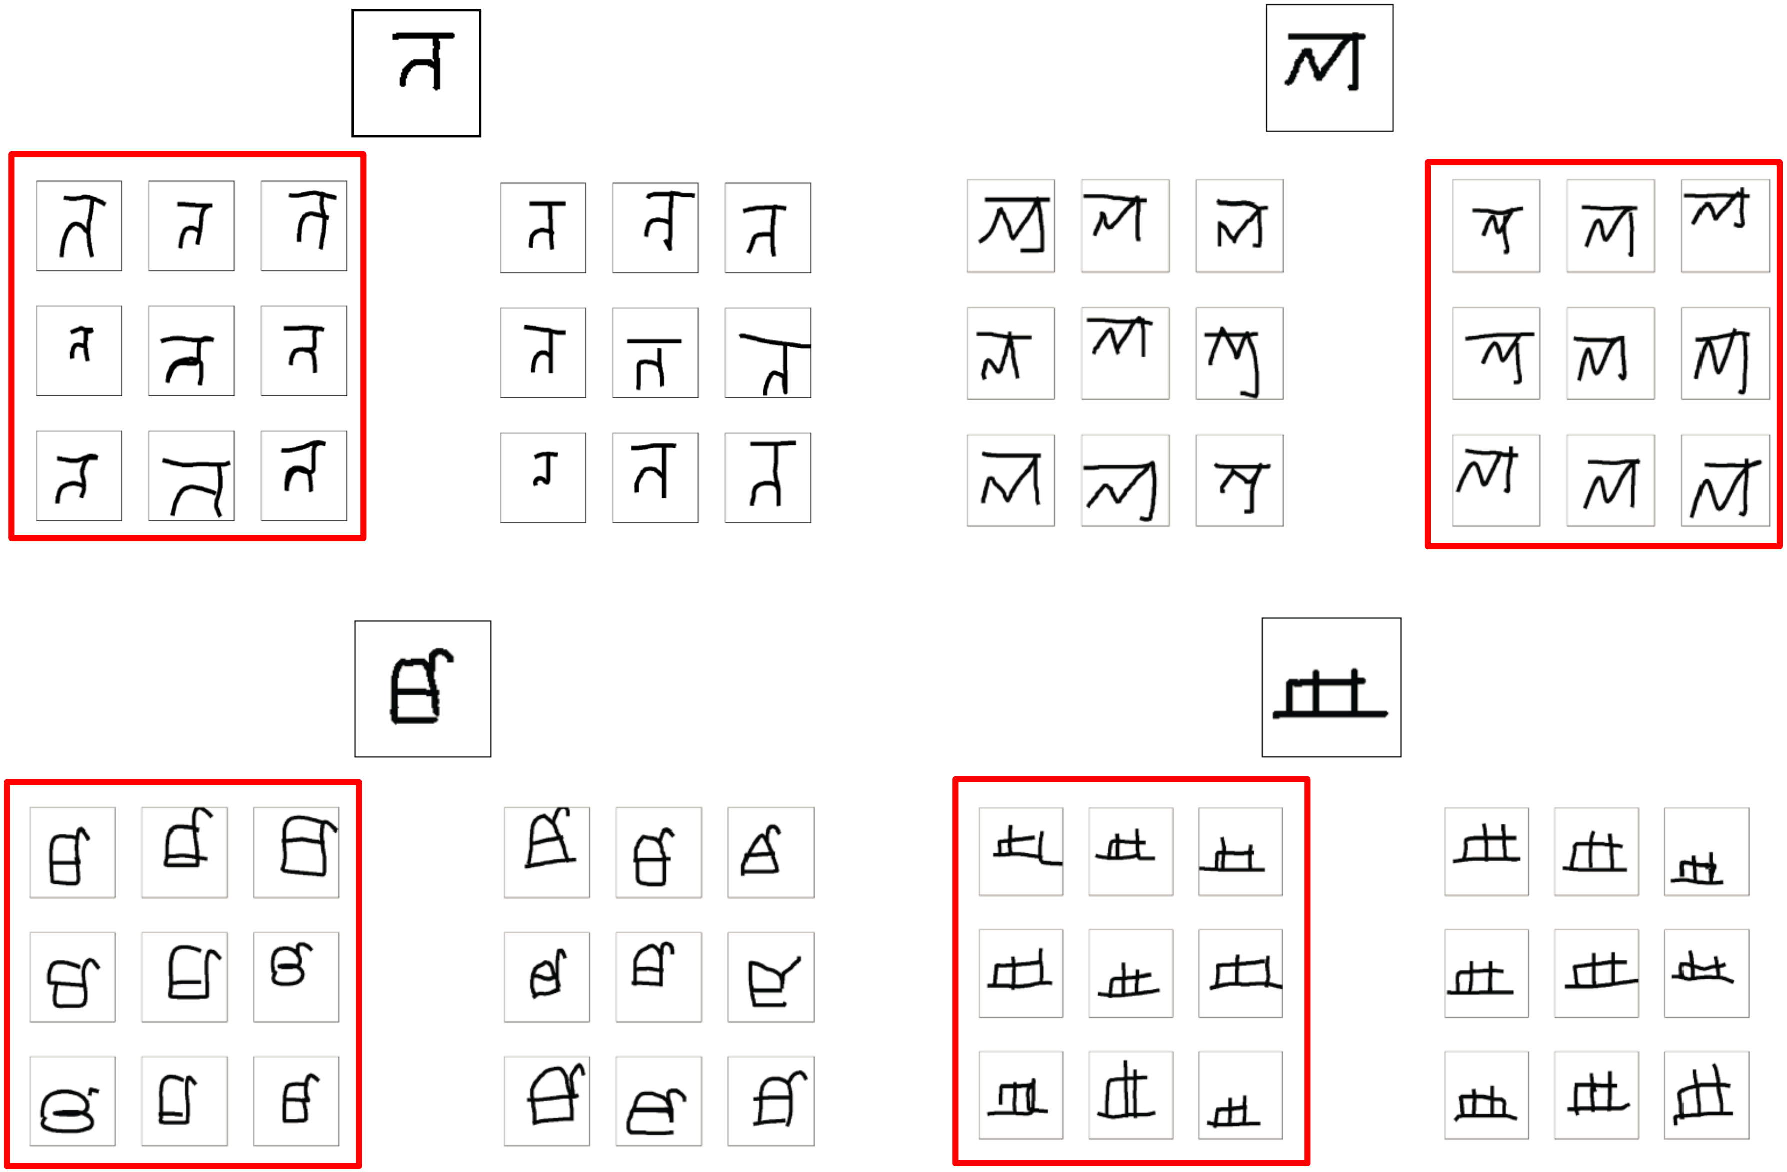

© AAAS. All rights reserved. This content is excluded from our Creative Commons license. For more information, see https://ocw.mit.edu/help/faq-fair-use/.
Source: Lake, Brenden M., Ruslan Salakhutdinov, and Joshua B. Tenenbaum. "Human-level concept learning through probabilistic program induction." Science 350, no. 6266 (2015): 1332-1338.

Turing test: Can people tell the humans from the machine?

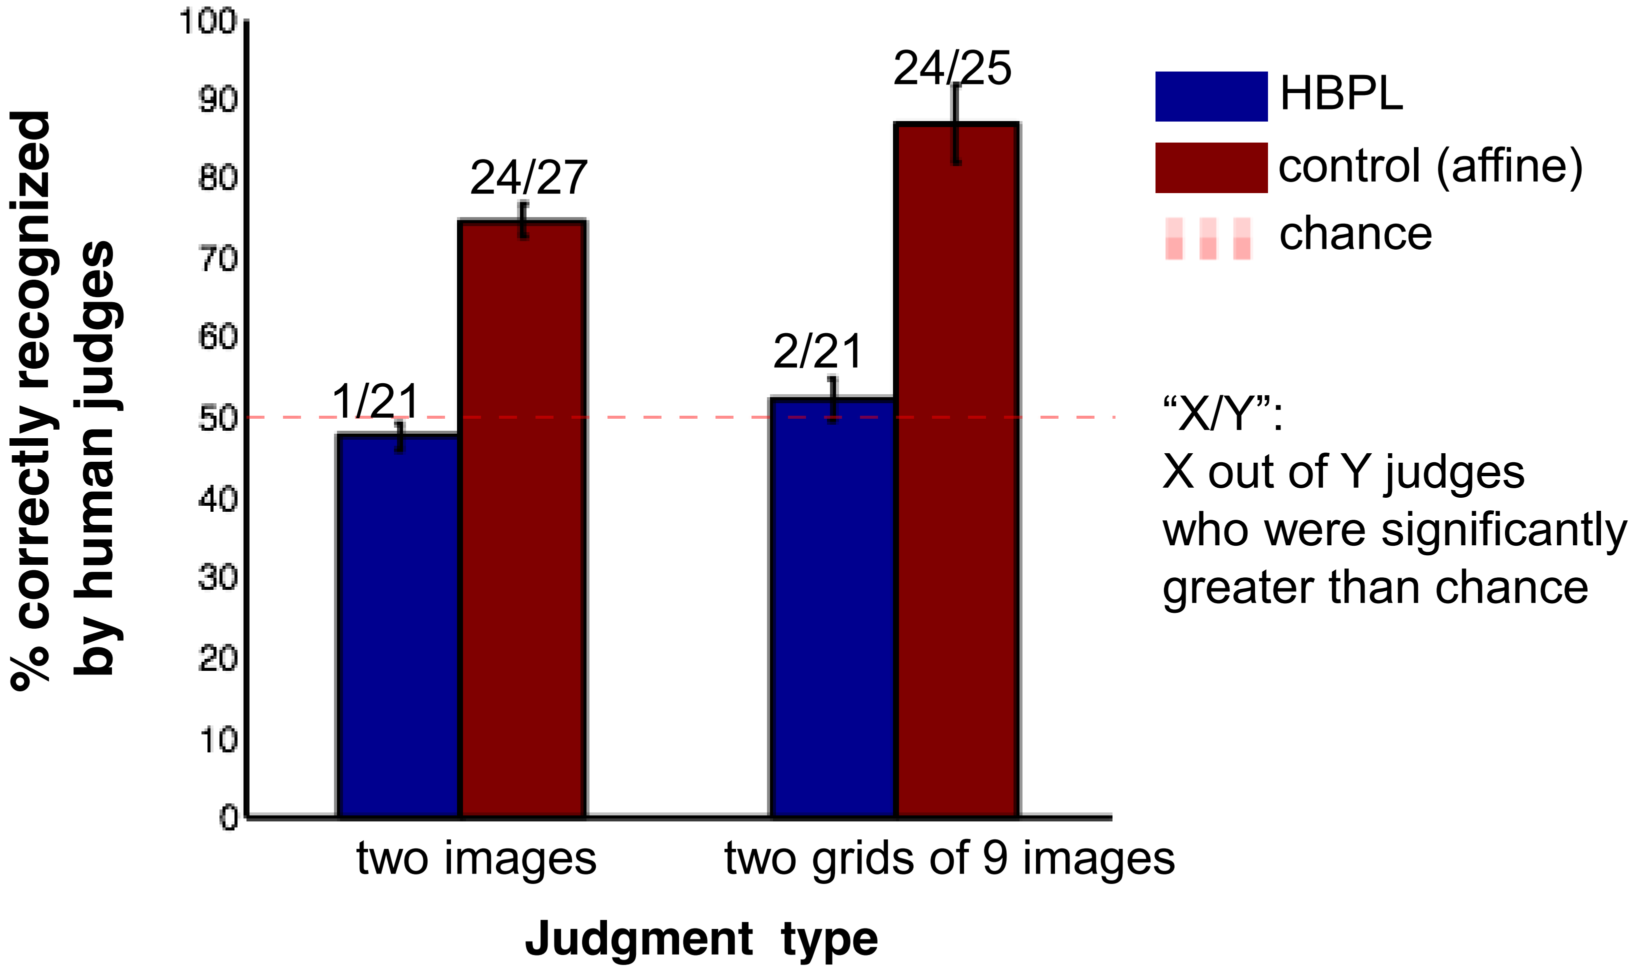

© AAAS. All rights reserved. This content is excluded from our Creative Commons license. For more information, see https://ocw.mit.edu/help/faq-fair-use/.
Source: Lake, Brenden M., Ruslan Salakhutdinov, and Joshua B. Tenenbaum. "Human-level concept learning through probabilistic program induction." Science 350, no. 6266 (2015): 1332-1338.

Generating entirely new characters

10 Examples Given

25 New Instances Generated

Machine

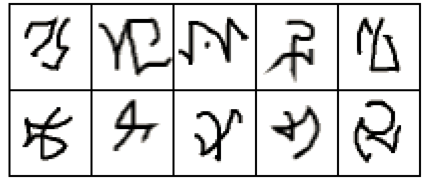

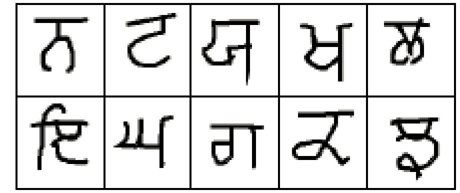

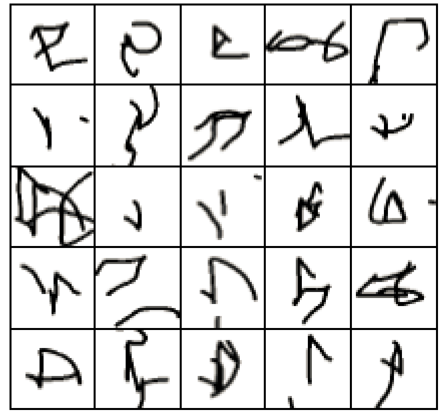

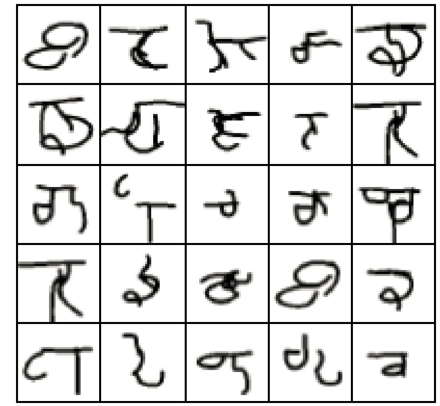

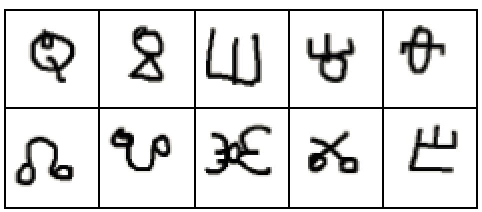

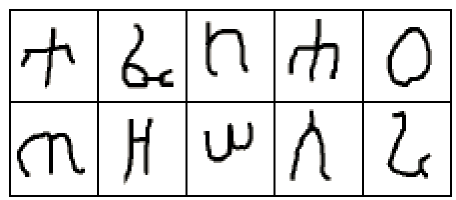

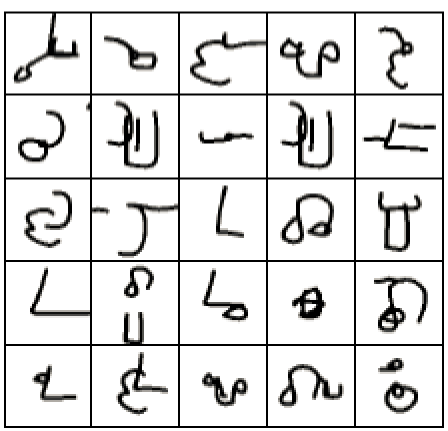

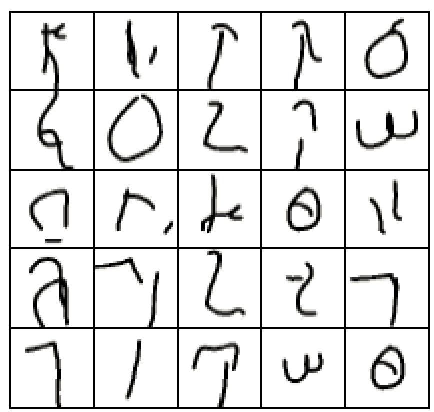

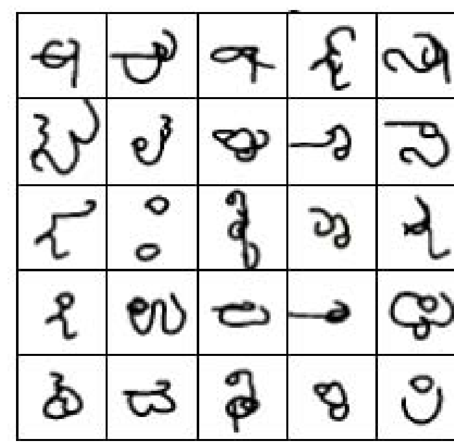

Explaining the dynamics of development? (w/ T. Ullman, Spelke, others)

9 months

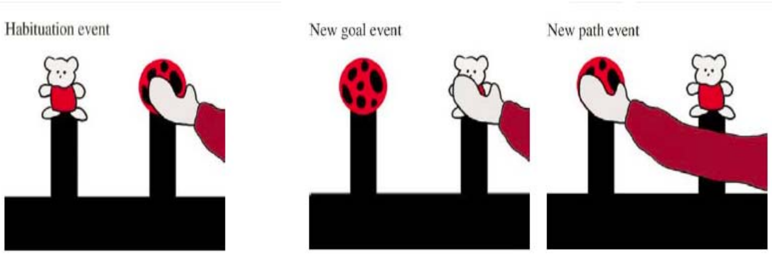

Courtesy of Elsevier, Inc., http://www.sciencedirect.com. Used with permission.
Source: Sommerville, Jessica A., Amanda L. Woodward, and Amy Needham. "Action experience alters 3-month-old infants' perception of others' actions." Cognition 96, no. 1 (2005): B1-B11.

12 months

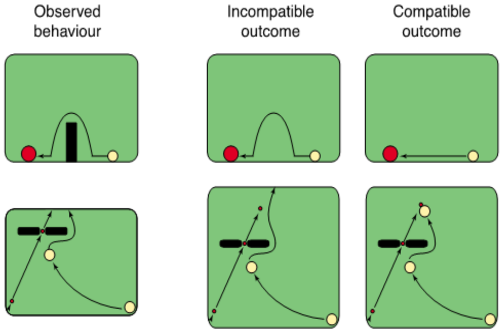

© Elsevier. All rights reserved. This content is excluded from our Creative Commons license. For more information, see https://ocw.mit.edu/help/faq-fair-use/.
Source: Skerry, Amy E., and Elizabeth S. Spelke. "Preverbal infants identify emotional reactions that are incongruent with goal outcomes." Cognition 130, no. 2 (2014): 204-216.

15 months

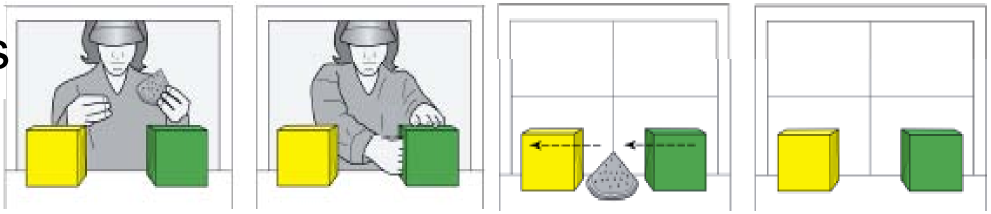

© AAAS. All rights reserved. This content is excluded from our Creative Commons license. For more information, see https://ocw.mit.edu/help/faq-fair-use/.
Source: Onishi, Kristine H., and Renée Baillargeon. "Do 15-month-old infants understand false beliefs?" science 308, no. 5719 (2005): 255-258.

Explaining the dynamics of development? (w/ T. Ullman, Spelke, others)

9 months

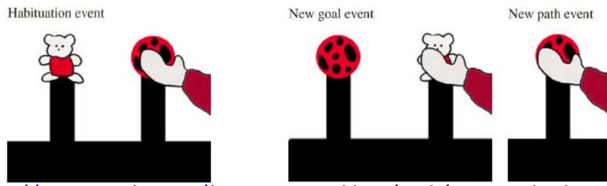

Courtesy of Elsevier, Inc., http://www.sciencedirect.com. Used with permission.
Source: Sommerville, Jessica A., Amanda L. Woodward, and Amy Needham. "Action experience alters 3-month-old infants' perception of others' actions." Cognition 96, no. 1 (2005): B1-B11.

Capture different knowledge stages with a sequence of probabilistic programs?

Explain the trajectory of stages as rational statistical inference in the space of programs?

15 months

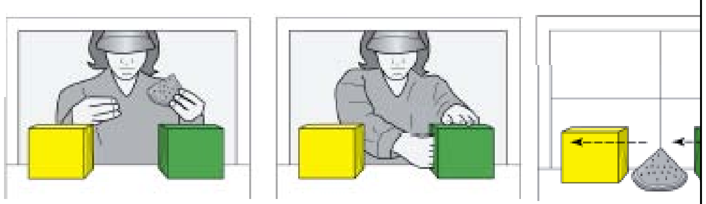

3 months

5 months

6.5 months

12.5 months

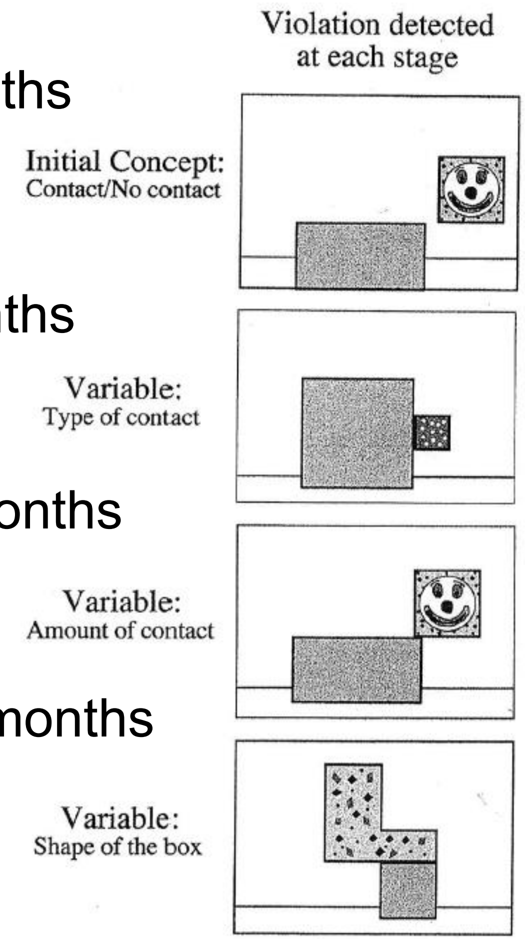

Courtesy of Elsevier, Inc., http://www.sciencedirect.com. Used with permission.
Source: Baillargeon, Renée."Infants' understanding of the physical world." Journal of the Neurological Sciences 143, no. 1-2 (1996): 199.

Learning physics from dynamic scenes

(Ullman, Stuhlmuller, Goodman, Tenenbaum, 2014; under review)

Unobserved properties:
(c.f. parameter learning)

e.g., mass, charge, friction, elasticity, resistance…

See the lecture video to view these video clips

New laws:
(c.f. structure learning)

e.g., presence of forces and their shape, existence of hidden objects, kinds of substances …

Metatheory

Objects


```
(define (construct-particle size position velocity mass)
  (list size position velocity mass))

(define (construct-barrier size elasticity position)
  (list size elasticity position))
```

Inertial dynamics

$$F = m \times a$$

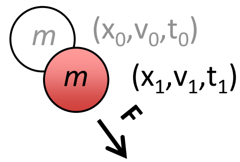


```
(define (next-position objects forces dt)
  (let ((masses       (get-mass objects))
        (position0    (get-position objects))
        (velocity0    (get-velocity objects))
        (a (/ forces masses)))
   (numerical-integration position0 velocity0 a dt)))
```

Theories

Different forces

Coupling

Global

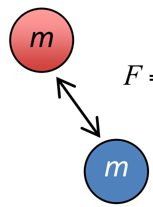

$$F = \frac{C \times m_1 \times m_2}{r^2}$$

```
(define (attraction object1 object2)
      (let ((r  (euclidian-dist objects))
            (m1 (get-mass object1))
            (m2 (get-mass object2)))
        (/ (* C m1 m2) (power r 2))))
```

Different masses

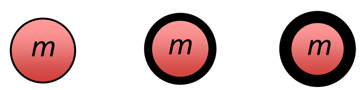


```
(define (heavy-mass 9.0))
```

Different frictions

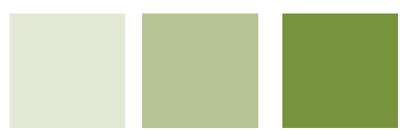

```
(define (smooth-surface 0.0))
```

Events

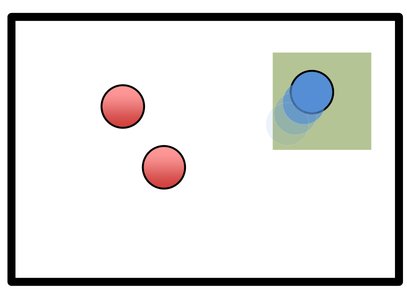

```
(define world-5 (create-world
      (create-forces (collision rr-attract global-left))
      (create-particles (draw-random-particles 3))
      (create-friction (draw-friction-surfaces 1))))
```

© Proceeding of the National Academy of Sciences. All rights reserved.This content is excluded from our Creative Commons license. For more information, see https://ocw.mit.edu/help/faq-fair-use/.

Source: Ullman, Tomer, Andreas Stuhlmüller, Noah Goodman, and Joshua B. Tenenbaum. "Learning physics from dynamical scenes." In Proceedings of the thirty-sixth annual conference of the cognitive science society. 2014.

Comparing models and humans

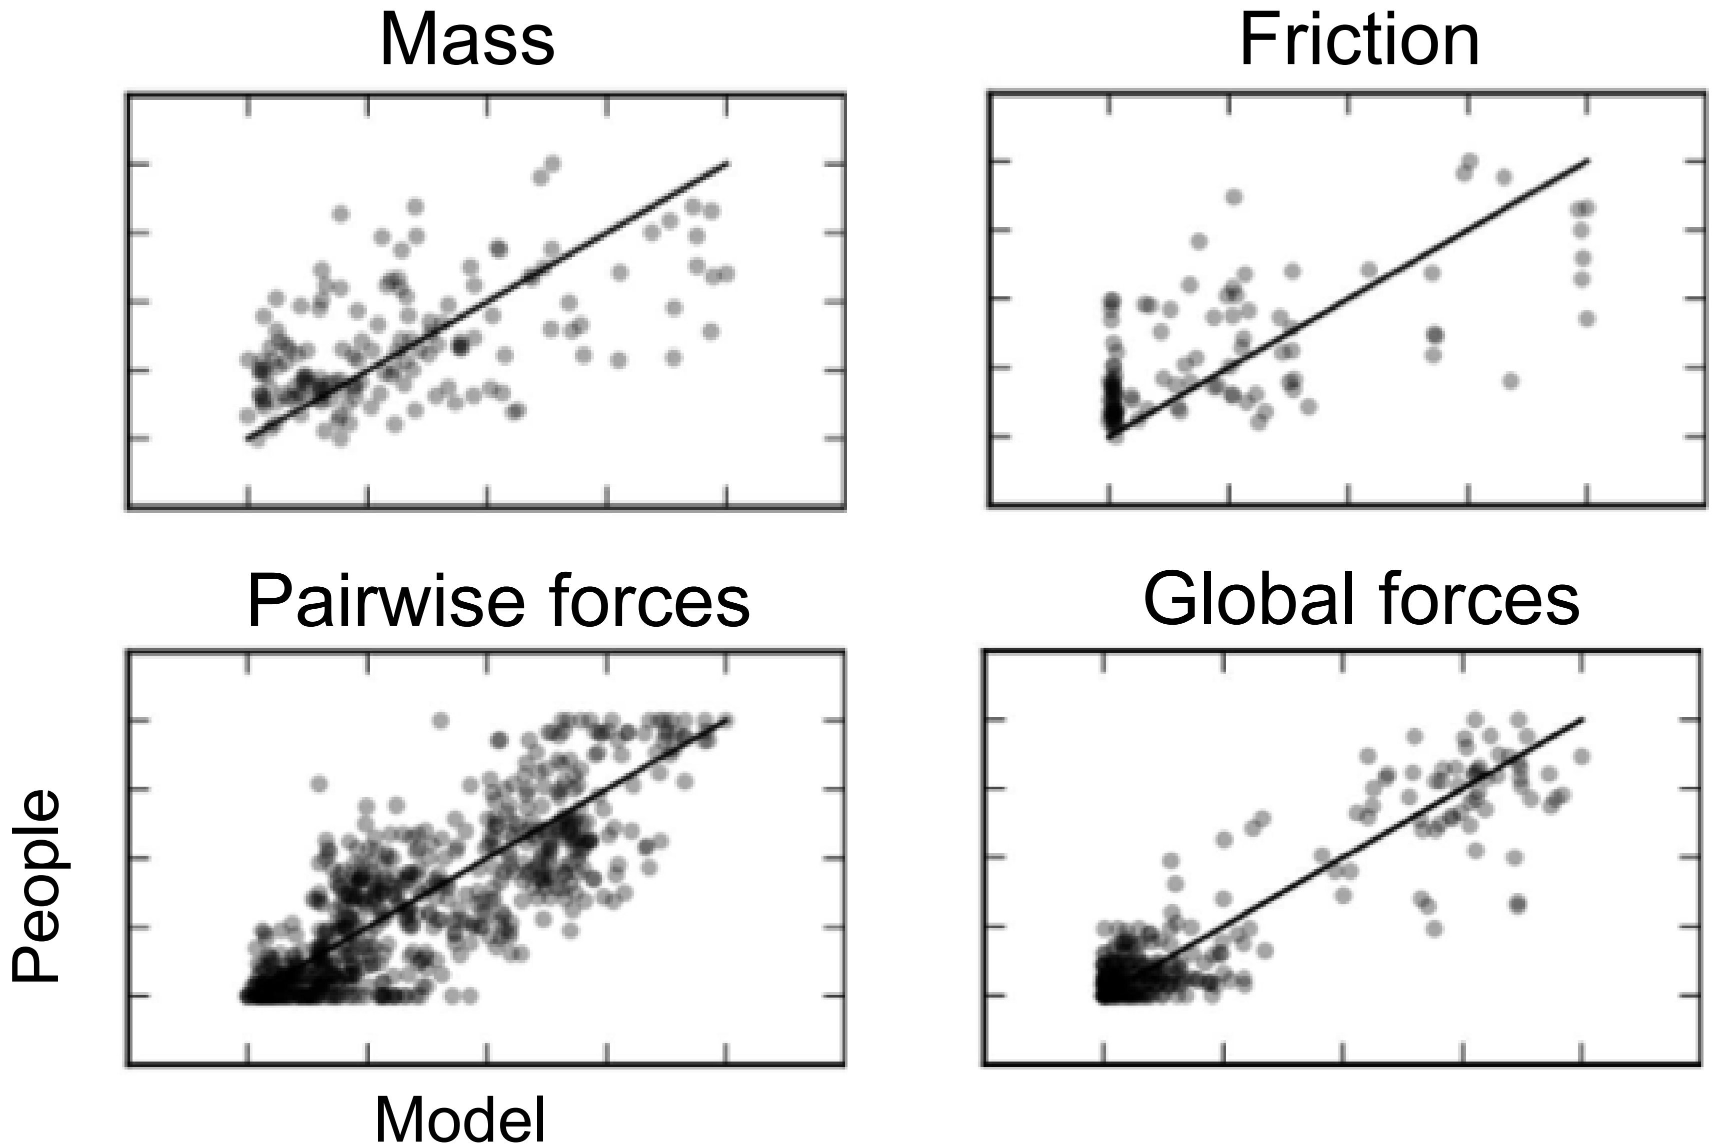

© Proceeding of the National Academy of Sciences. All rights reserved.Thiscontent is excluded from our Creative Commons license. For moreinformation, see https://ocw.mit.edu/help/faq-fair-use/.
Source: Ullman, Tomer, Andreas Stuhlmüller, Noah Goodman, and JoshuaB. Tenenbaum. "Learning physics from dynamical scenes." In Proceedings of the thirty-sixth annual conference of the cognitive science society. 2014.

Learning the form of domain theories?

A really hard problem…

- What’s the right hypothesis space?
- What’s an effective algorithm for searching the space of theories, as fast and as reliably and as flexibly as we see in children’s learning?

Learning the form of domain theories?

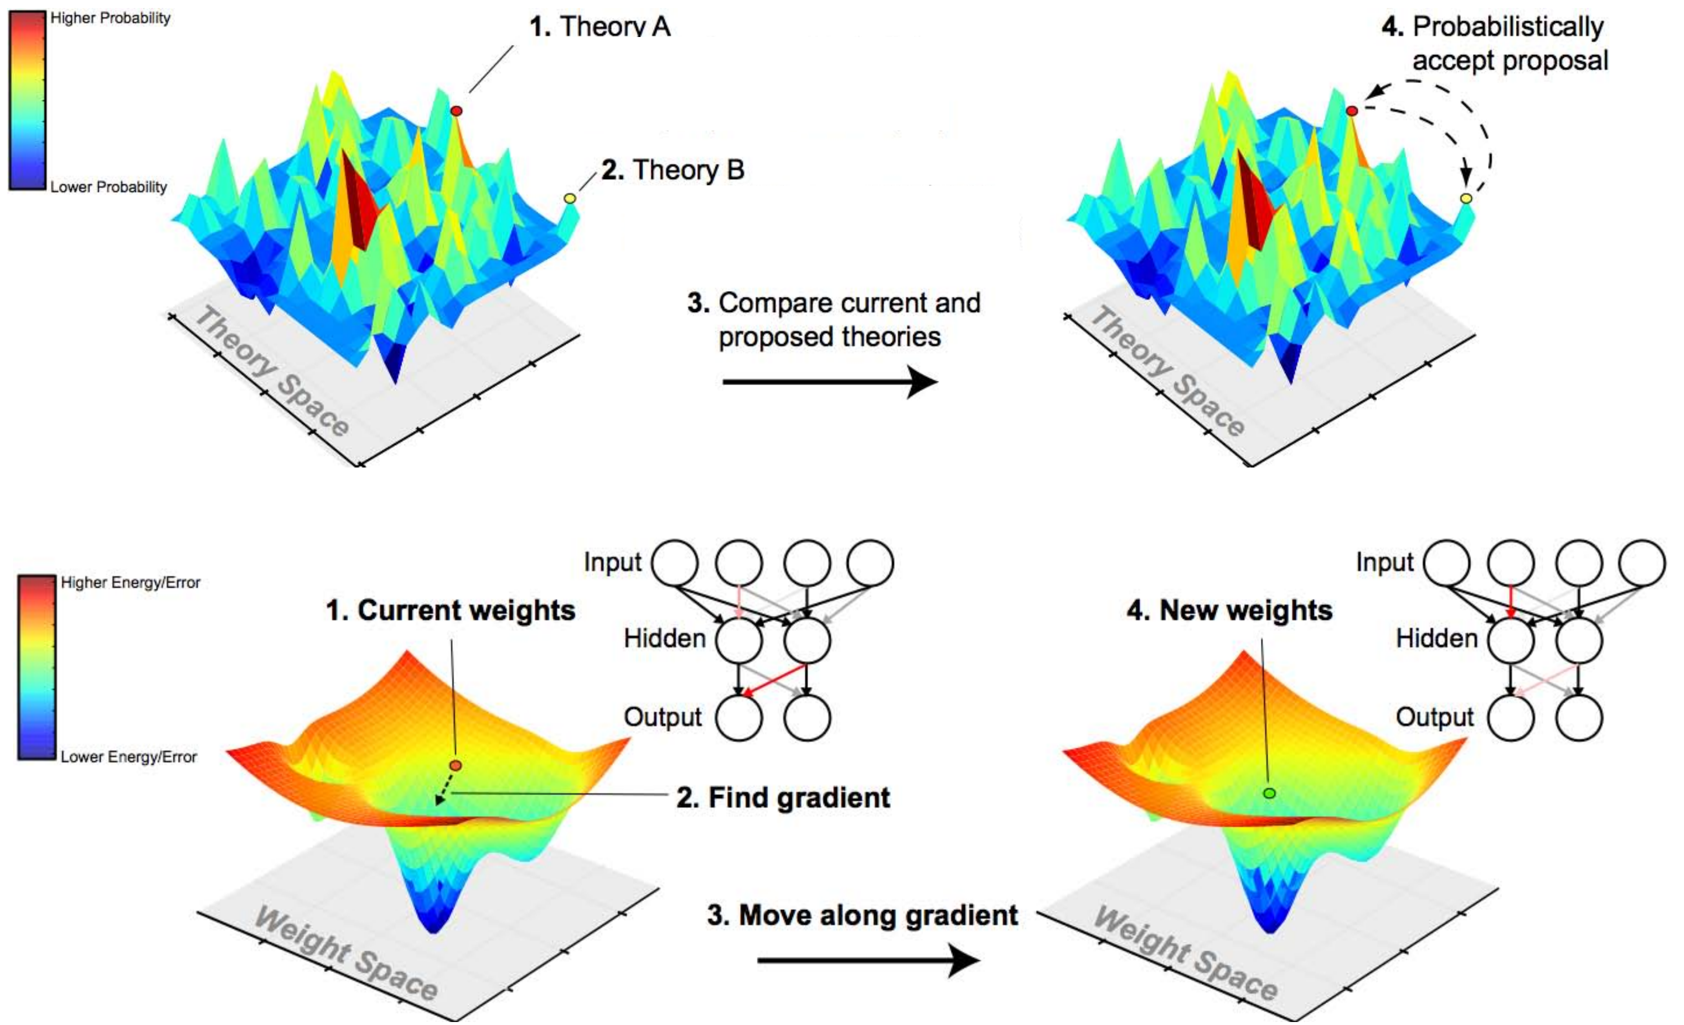

© Elsevier. All rights reserved. This content is excluded from our Creative Commons license. For more information, see https://ocw.mit.edu/help/faq-fair-use/.
Source: Ullman, Tomer D., Noah D. Goodman, and Joshua B. Tenenbaum. "Theory learning as stochastic search in the language of thought." Cognitive Development 27, no. 4 (2012): 455-480.

Hierarchical Bayesian Framework

(Kemp & Tenenbaum, *Psych Review*, 2009)

F: form

S: structure

D: data

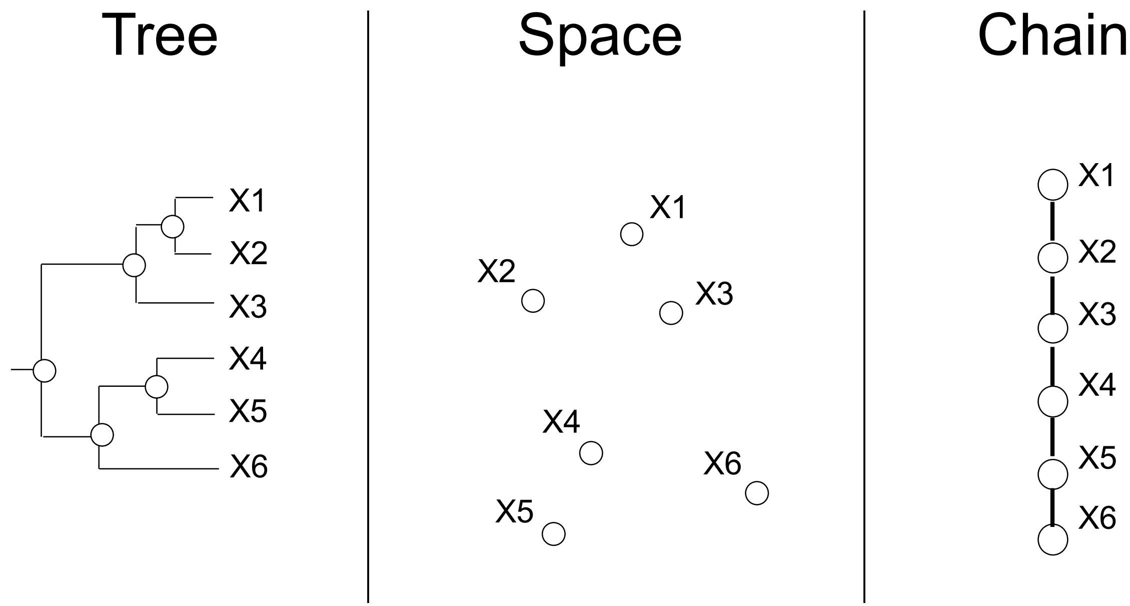

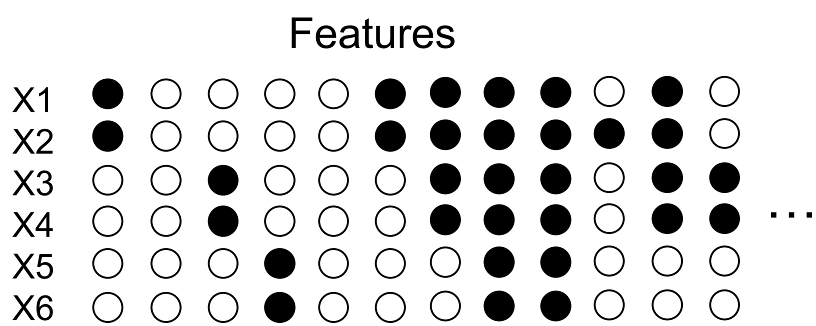

© Psychological Reivew. All rights reserved. This content is excluded from our Creative Commons license. For more information, see https://ocw.mit.edu/help/faq-fair-use/.
Source: Kemp, Charles, and Joshua B. Tenenbaum. "Structured statistical models of inductive reasoning." Psychological review 116, no. 1 (2009): 20.

Hypothesis space of structural forms
(Kemp & Tenenbaum, PNAS 2008)

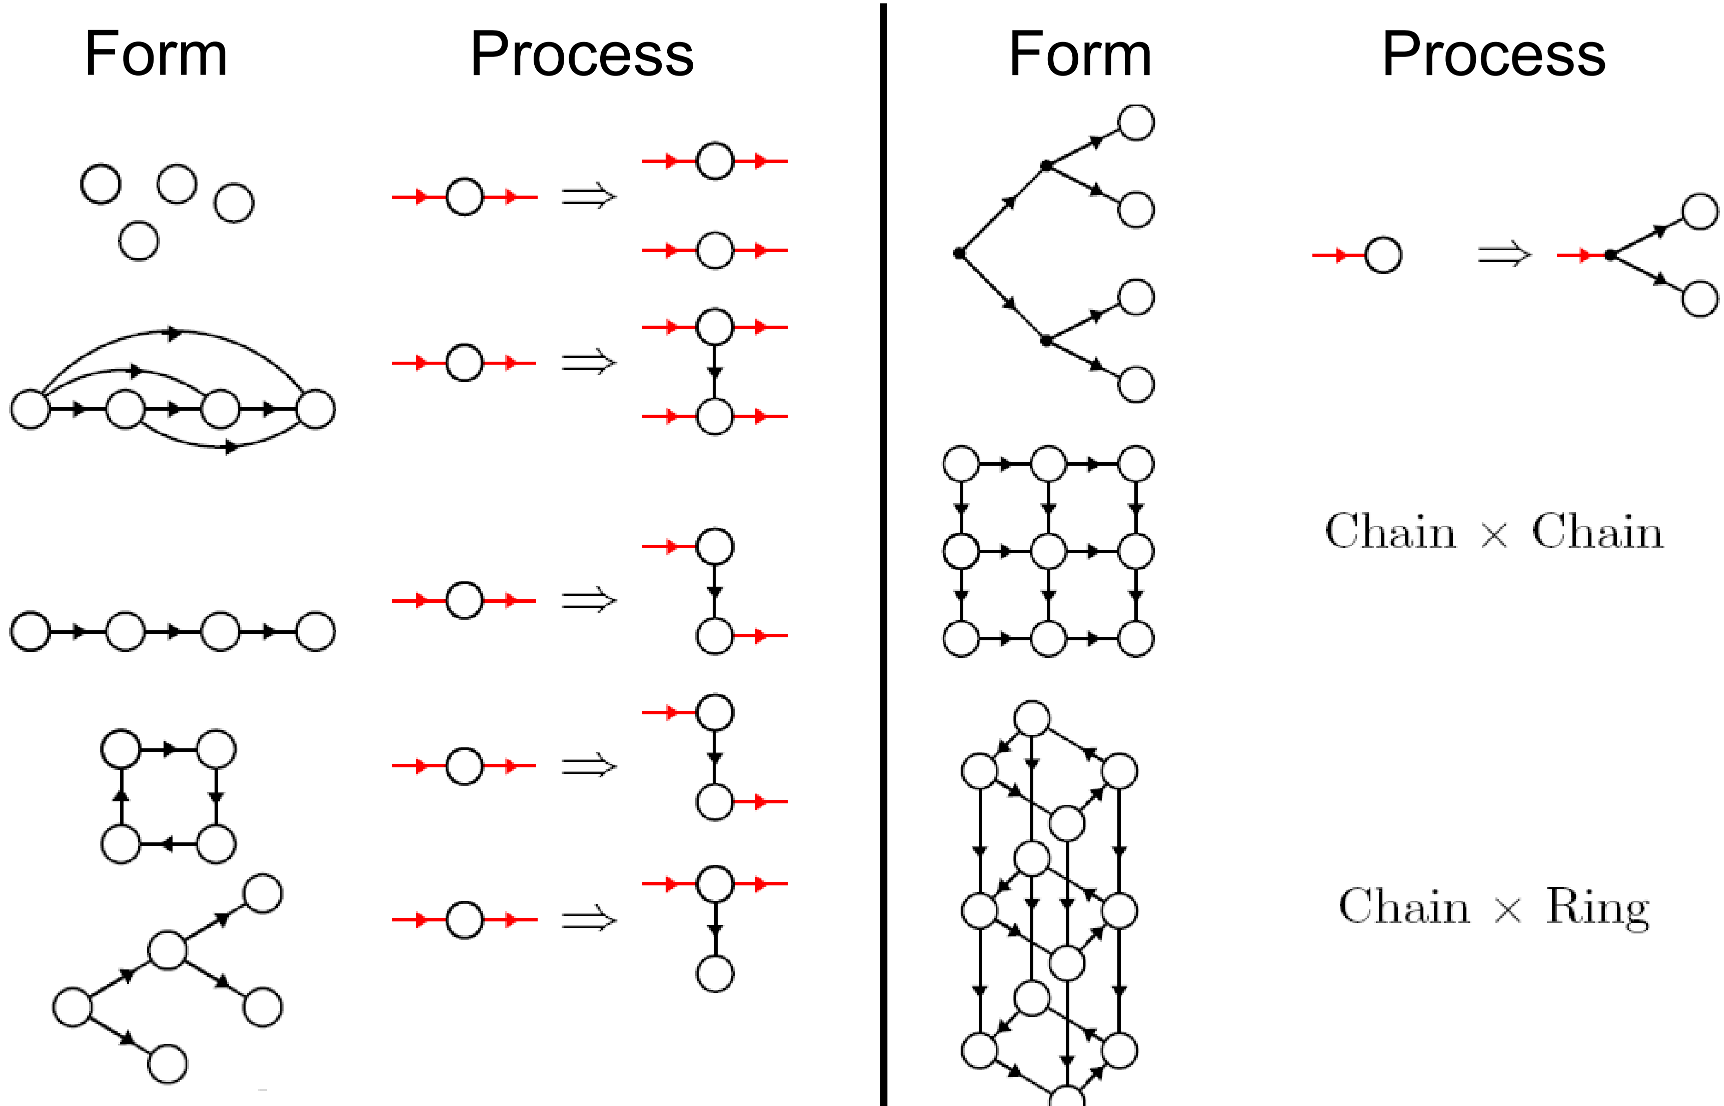

Courtesy of National Academy of Sciences, U. S. A. Used with permission.
Source: Kemp, Charles, and Joshua B. Tenenbaum. "The discovery of structural form." Proceedings of the National Academy of Sciences 105, no. 31 (2008): 10687-10692. Copyright © 2008 National Academy of Sciences, U.S.A.

Discovering the structural form of a domain
(Kemp & Tenenbaum, *PNAS* 2008; *Psych Review*, 2009)

Abstract principles

tree: 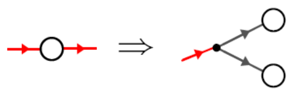

chain: 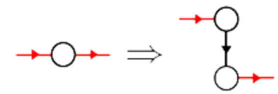

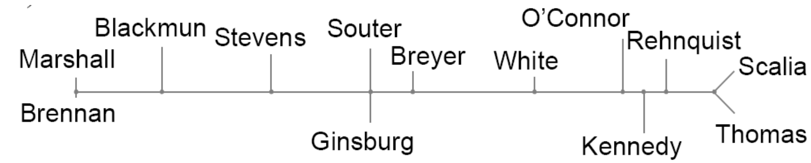

ring: 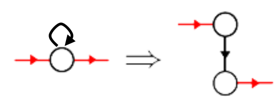

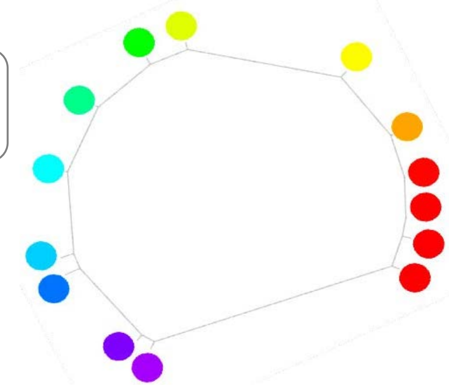

Model

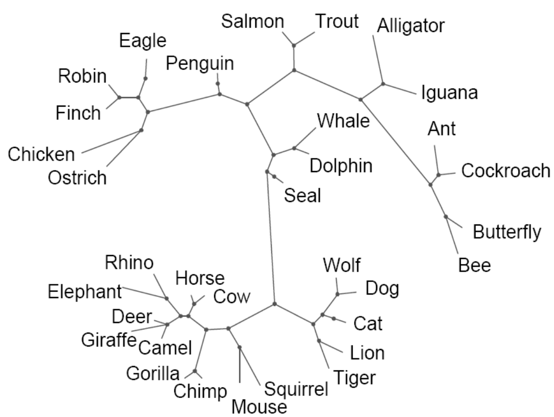

Courtesy of National Academy of Sciences, U. S. A. Used with permission. Source: Kemp, Charles, and Joshua B. Tenenbaum. "The discovery of structural form." Proceedings of the National Academy of Sciences 105, no. 31 (2008): 10687-10692. Copyright © 2008 National Academy of Sciences, U.S.A.

chain x ring

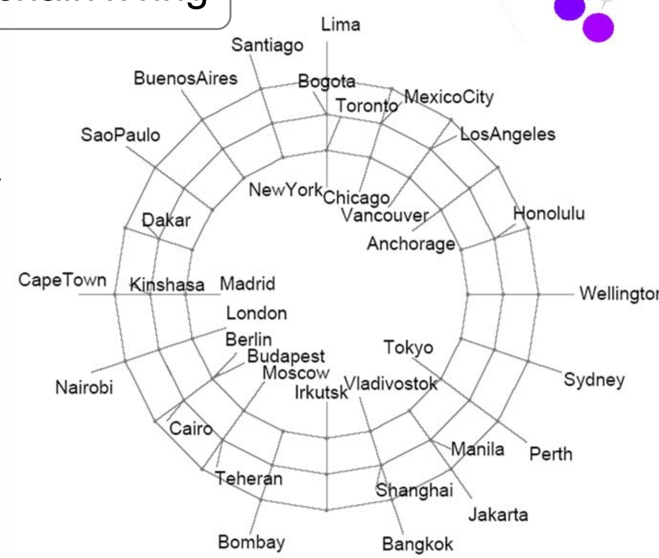

chain x chain

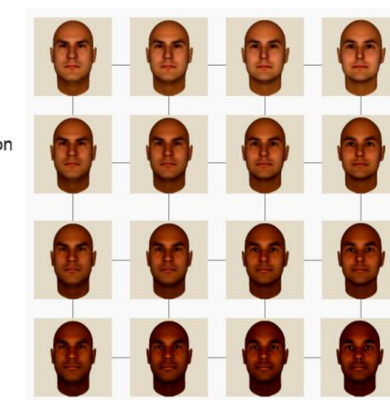

Data

features

animals

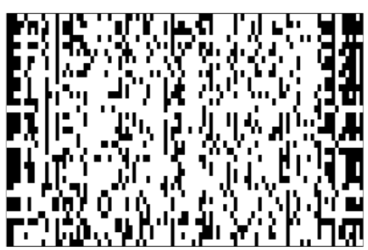

Development of structural forms as more data are observed

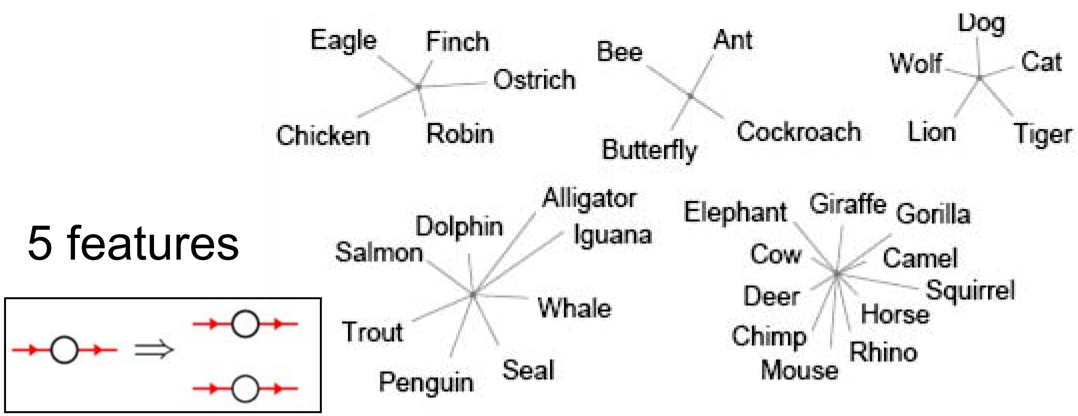

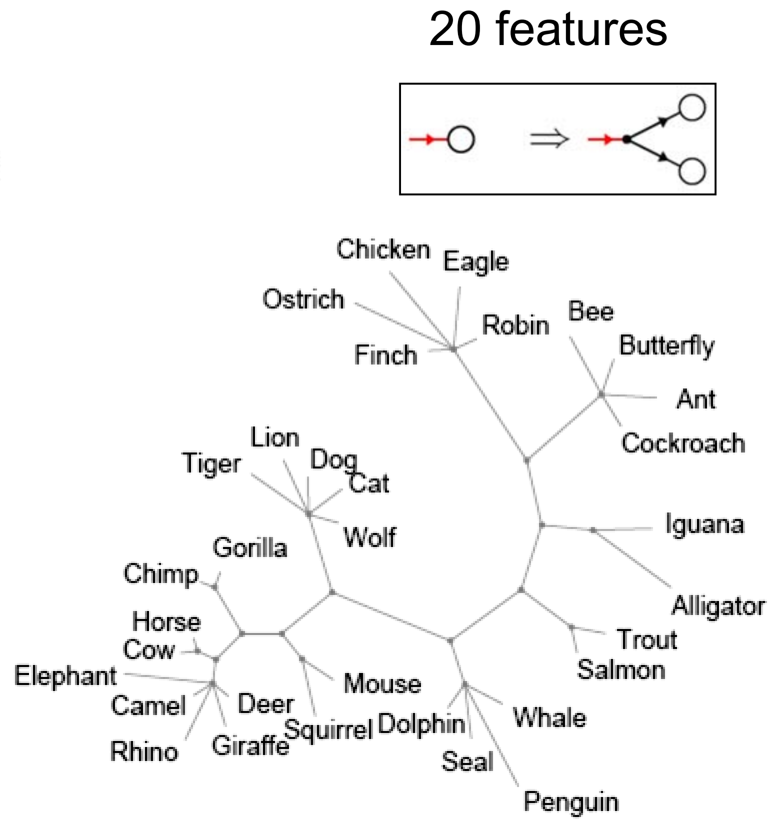

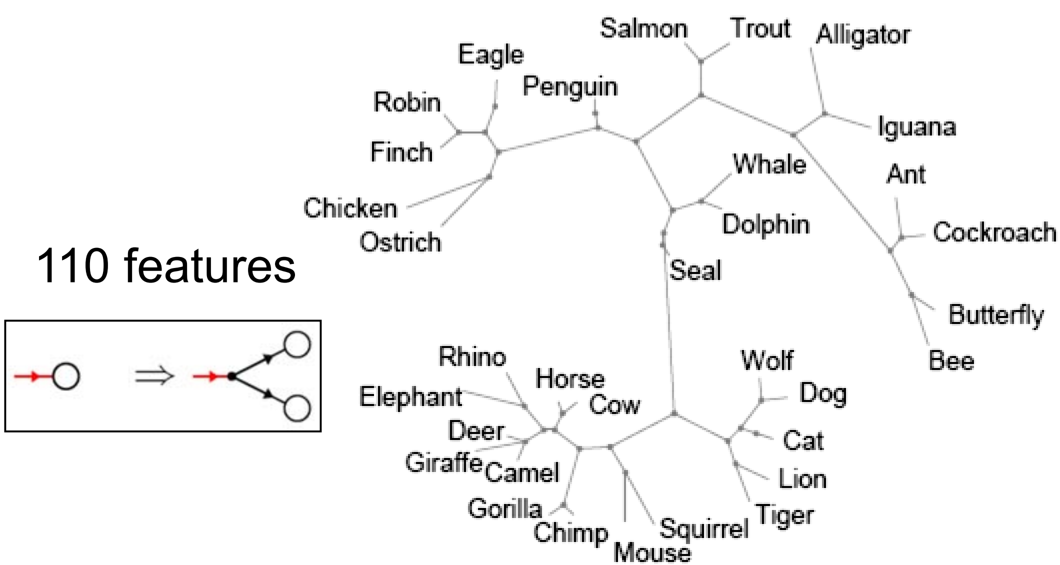

“blessing of abstraction”

Courtesy of National Academy of Sciences, U. S. A. Used with permission.
Source: Kemp, Charles, and Joshua B. Tenenbaum. "The discovery of structural form." Proceedings of the National Academy of Sciences 105, no. 31 (2008): 10687-10692. Copyright © 2008 National Academy of Sciences, U.S.A.

Conclusion

What makes us so smart?

1. **How we start:** Common-sense core theories of intuitive physics and intuitive psychology.
2. **How we grow:** Learning as theory construction, revision and refinement.

The tools of probabilistic programs and program induction are beginning to let us reverse-engineer these capacities, with languages that are:

- Probabilistic.
- Generative.
- Causally structured
- Compositionally structured: flexible, fine-grained dependencies, hierarchical, recursive, unbounded

We have to view the brain not simply as a pattern-recognition device, but as a *modeling engine*, an *explanation engine* – and we have to understand how these views work together.

Much promise but huge engineering and scientific challenges remain… full of opportunities for bidirectional interactions between cognitive science, neuroscience, developmental psychology, AI and machine learning.

MIT OpenCourseWare
https://ocw.mit.edu

Resource: Brains, Minds and Machines Summer Course
Tomaso Poggio and Gabriel Kreiman

The following may not correspond to a particular course on MIT OpenCourseWare, but has been provided by the author as an individual learning resource.

For information about citing these materials or our Terms of Use, visit: https://ocw.mit.edu/terms.