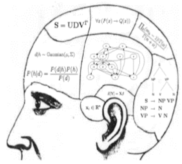

First-order probabilistic models of human cognition

Josh Tenenbaum
MIT

Why logic + probability?

- Two complementary formalisms:
  - Logic: a framework for knowledge representation.
  - Probability: a framework for inductive inference.
- Why do we need to integrate them?
  - To account for our knowledge about the structure of the world
    - its form and content.
    - how it is used and acquired.
  - To capture linguistic meaning and use.

Why logic + probability?

- To capture deep inductive biases

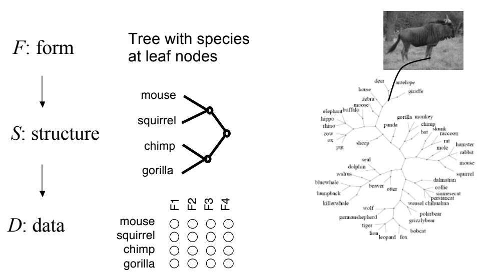

Outline

- The traditional debate in cognitive science: logic *versus* probability
  - The case of connectionism
  - Examples: knowledge about biology, social relations, language, visual objects
- Models of human reasoning that integrate probabilistic inference and logical representations
  - Ecological reasoning
  - Causal reasoning
- Other directions
  - Social reasoning, physical reasoning

Movitations for connectionism

- Build models with more neural plausibility.
- Overcome the problems of symbolic representations.
  - Inference is too rigid and brittle.
  - No general way to learn new representations.

Semantic networks
(Quillian, 1968)

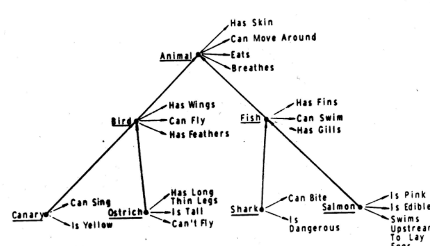

- Useful for compression (memory)
- Useful for predicting properties of new objects.

Reaction time tests of hierarchy

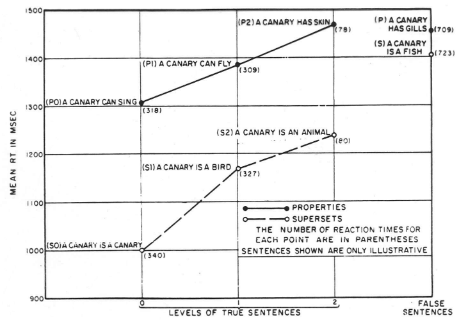

Problems

- Typicality effects.
  - "canary is a bird" faster than "chicken is a bird".
- Violations of hierarchy for atypical items.
  - "chicken is an animal" faster than "chicken is a bird."
- How could this knowledge representation be learned in an unsupervised way?

An alternative architecture

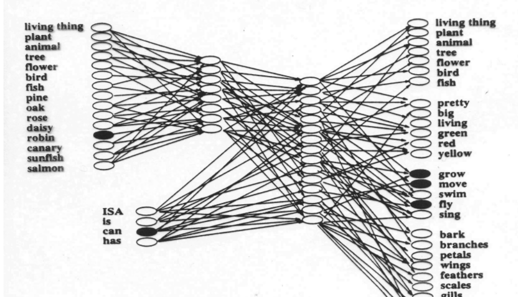

(Rogers and McClelland, 2004)

Training set

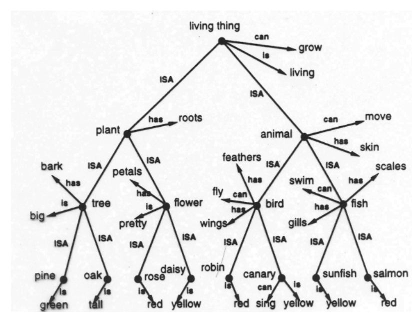

Learned distributed representation

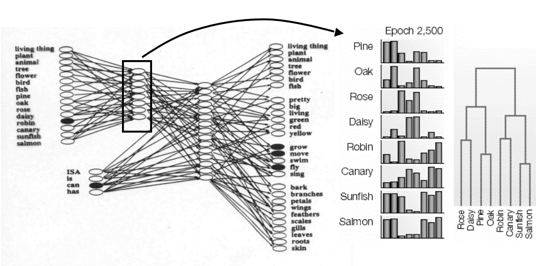

Problems

- Does not actually capture generalization behavior very well.
  - Correlations: $r < 0.7$ on basic property tasks.
  - Inductive bias is too weak
- Missing crucial abstract knowledge about the domain.
  - e.g., ISA(x,y) <= ISA(x,z) & ISA(z,y)
- Can we combine the best of logical representation and statistical learning and inference?

Learning family relationships
(Hinton, 1986)

- A family tree structure:

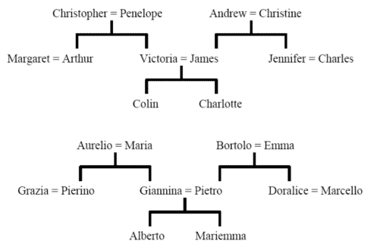

The family relations dataset

father(Christopher, Arthur)
father(Christopher, Victoria)
father(Andrew, James)
father(Andrew, Jennifer)
father(James, Colin)
father(James, Charlotte)

mother(Penelope, Arthur)
mother(Penelope, Victoria)
mother(Christine, James)
mother(Christine, Jennifer)
mother(Victoria, Colin)
mother(Victoria, Charlotte)

husband(Christopher, Penelope)
husband(Andrew, Christine)
husband(Arthur, Margaret)
husband(James, Victoria)
husband(Charles, Jennifer)

wife(Penelope, Christopher)
wife(Christine, Andrew)
wife(Margaret, Arthur)
wife(Victoria, James)
wife(Jennifer, Charles)

son(Arthur, Christopher)
son(Arthur, Penelope)
son(James, Andrew)
son(James, Christine)
son(Colin, Victoria)
son(Colin, James)

daughter(Victoria, Christopher)
daughter(Victoria, Penelope)
daughter(Jennifer, Andrew)
daughter(Jennifer, Christine)
daughter(Charlotte, Victoria)
daughter(Charlotte, James)

brother(Arthur, Victoria)
brother(James, Jennifer)
brother(Colin, Charlotte)

sister(Victoria, Arthur)
sister(Jennifer, James)
sister(Charlotte, Colin)

uncle(Arthur, Colin)
uncle(Charles, Colin)
uncle(Arthur, Charlotte)
uncle(Charles, Charlotte)

aunt(Jennifer, Colin)
aunt(Margaret, Colin)
aunt(Jennifer, Charlotte)
aunt(Margaret, Charlotte)

nephew(Colin, Arthur)
nephew(Colin, Jennifer)
nephew(Colin, Margaret)
nephew(Colin, Charles)

niece(Charlotte, Arthur)
niece(Charlotte, Jennifer)
niece(Charlotte, Margaret)
niece(Charlotte, Charles)

Learning family relationships
(Hinton, 1986)

- Network architecture:

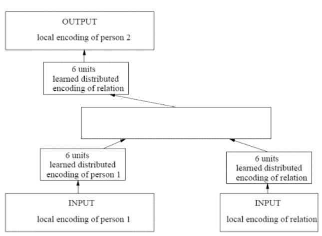

Learning family relationships
(Hinton, 1986)

- 112 possible facts of the form:
  - <person1, relation, person2>
  - <Christopher, father-of, Victoria>,
  - <Colin, son-of, Victoria>,
  - <Jennifer, aunt-of, Colin> . . .
- Trained on 108 examples, network usually generalizes well to the other 4.
  - Doesn't work well with less training.

Linear Relational Embedding
(Paccanaro and Hinton, 2002)

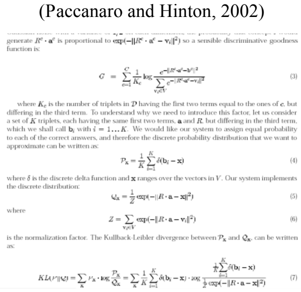

generate $R^c \cdot \mathbf{a}^c$ is proportional to $\exp(-\|R^c \cdot \mathbf{a}^c - \mathbf{v}_i\|^2)$ so a sensible discriminative goodness function is:

$$G = \sum_{c=1}^{C} \frac{1}{K_c} \log \frac{e^{-\|R^c \cdot \mathbf{a}^c - \mathbf{b}^c\|^2}}{\sum_{\mathbf{v}_i \in V} e^{-\|R^c \cdot \mathbf{a}^c - \mathbf{v}_i\|^2}} \quad (3)$$

where K_c is the number of triplets in $\mathcal{D}$ having the first two terms equal to the ones of c, but differing in the third term. To understand why we need to introduce this factor, let us consider a set of K triplets, each having the same first two terms, $\mathbf{a}$ and R, but differing in the third term, which we shall call $\mathbf{b}_i$ with $i = 1 \ldots K$. We would like our system to assign equal probability to each of the correct answers, and therefore the discrete probability distribution that we want to approximate can be written as:

$$\mathcal{P}_{\mathbf{x}} = \frac{1}{K} \sum_{i=1}^{K} \delta(\mathbf{b}_i - \mathbf{x}) \quad (4)$$

where δ is the discrete delta function and $\mathbf{x}$ ranges over the vectors in V. Our system implements the discrete distribution:

$$\mathcal{Q}_{\mathbf{x}} = \frac{1}{Z} \exp(-\|R \cdot \mathbf{a} - \mathbf{x}\|^2) \quad (5)$$

where

$$Z = \sum_{\mathbf{v}_i \in V} \exp(-\|R \cdot \mathbf{a} - \mathbf{v}_i\|^2) \quad (6)$$

is the normalization factor. The Kullback-Leibler divergence between $\mathcal{P}_{\mathbf{x}}$ and $\mathcal{Q}_{\mathbf{x}}$, can be written as:

$$KL(\mathcal{P}\|\mathcal{Q}) = \sum_{\mathbf{x}} \mathcal{P}_{\mathbf{x}} \cdot \log \frac{\mathcal{P}_{\mathbf{x}}}{\mathcal{Q}_{\mathbf{x}}} = \sum_{\mathbf{x}} \frac{1}{K} \sum_{i=1}^{K} \delta(\mathbf{b}_i - \mathbf{x}) \cdot \log \frac{\frac{1}{K}\sum_{i=1}^{K} \delta(\mathbf{b}_i - \mathbf{x})}{\frac{1}{Z}\exp(-\|R \cdot \mathbf{a} - \mathbf{x}\|^2)} \quad (7)$$

- Relatively minor improvement, from 4 to 8 or 12 generalization trials….

A more intuitive representation

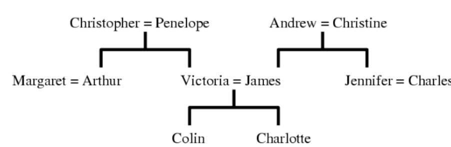

- Relations: spouse ("="), parent (solid line)
- Attribute: male or female (type of name)
- Define other relations in terms of basic relations spouse, parent, and the attributes male, female.

Minimal family description

spouse(Christopher,Penelope)
spouse(Andrew,Christine)
spouse(Arthur,Margaret)
spouse(James,Victoria)
spouse(Charles,Jennifer)

female(Penelope)
female(Christine)
female(Margaret)
female(Victoria)
female(Jennifer)
female(Charlotte)
NOT female(Colin)

parent(Penelope, Arthur)
parent(Penelope, Victoria)
parent(Christine, James)
parent(Christine, Jennifer)
parent(Victoria, Colin)
parent(Victoria, Charlotte)

Abstract theory of kinship

spouse(x,y) <=> spouse(y,x)
NOT female(x) <= spouse(x,y) AND female(y)
parent(x,y) <= spouse(x,z) AND parent(z,y)

father(x,y) <=> parent(x,y) AND NOT female(x)
mother(x,y) <=> parent(x,y) AND female(x)
husband(x,y) <=> spouse(x,y) AND NOT female(x)
wife(x,y) <=> spouse(x,y) AND female(x)
son(x,y) <=> parent(y,x) AND NOT female(x)
daughter(x,y) <=> parent(y,x) AND female(x)
sibling(x,y) <=> parent(z,x) AND parent(z,y) AND ~(x=y)
brother(x,y) <=> sibling(x,y) AND NOT female(x)
sister(x,y) <=>sibling(x,y) AND female(x)
uncle(x,y) <=> (parent(z,y) AND brother(x,z))
OR (aunt(z,y) AND spouse(x,z))
aunt(x,y) <=> (parent(z,y) AND sister(x,z)) OR
OR (uncle(z,y) AND spouse(x,z))
nephew(x,y) <=> (parent(z,x) AND sibling(y,z) AND NOT female(x))
OR (nephew(x,z) AND spouse(y,z))
niece(x,y) <=> (parent(z,x) AND sibling(y,z) AND female(x))
OR (niece(x,z) AND spouse(y,z))

Properties of this representation

- Useful for compression (memory)

Minimal family description

spouse(Christopher,Penelope)
spouse(Andrew,Christine)
spouse(Arthur,Margaret)
spouse(James,Victoria)
spouse(Charles,Jennifer)

female(Penelope)
female(Christine)
female(Margaret)
female(Victoria)
female(Jennifer)
female(Charlotte)
NOT female(Colin)

parent(Penelope, Arthur)
parent(Penelope, Victoria)
parent(Christine, James)
parent(Christine, Jennifer)
parent(Victoria, Colin)
parent(Victoria, Charlotte)

Abstract theory of kinship

spouse(x,y) <=> spouse(y,x)
NOT female(x) <= spouse(x,y) AND female(y)
parent(x,y) <= spouse(x,z) AND parent(z,y)

father(x,y) <=> parent(x,y) AND NOT female(x)
mother(x,y) <=> parent(x,y) AND female(x)
husband(x,y) <=> spouse(x,y) AND NOT female(x)
wife(x,y) <=> spouse(x,y) AND female(x)
son(x,y) <=> parent(y,x) AND NOT female(x)
daughter(x,y) <=> parent(y,x) AND female(x)
sibling(x,y) <=> parent(z,x) AND parent(z,y) AND ~(x=y)
brother(x,y) <=> sibling(x,y) AND NOT female(x)
sister(x,y) <=>sibling(x,y) AND female(x)
uncle(x,y) <=> (parent(z,y) AND brother(x,z))
OR (aunt(z,y) AND spouse(x,z))
aunt(x,y) <=> (parent(z,y) AND sister(x,z)) OR
OR (uncle(z,y) AND spouse(x,z))
nephew(x,y) <=> (parent(z,x) AND sibling(y,z) AND NOT female(x))
OR (nephew(x,z) AND spouse(y,z))
niece(x,y) <=> (parent(z,x) AND sibling(y,z) AND female(x))
OR (niece(x,z) AND spouse(y,z))

Properties of this representation

- Useful for compression (memory)
- Useful for predicting unknown relations
  - Margaret is Arthur's wife. What else do we know about her?

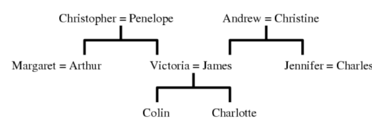

Reasoning about kinship

- Problem: Consider a new person, Boris.
  - Is the mother of Boris's father his grandmother?
  - Is the mother of Boris's sister his mother?
  - Is the daughter of Boris's sister his grandfather?
  - Is the son of Boris's sister his son?

Reasoning about kinship

- Problem: Consider a new person, Boris.
  - Is the mother of Boris's father his grandmother?
  - Is the mother of Boris's sister his mother?
  - Is the daughter of Boris's sister his grandfather?
  - Is the son of Boris's sister his son? (*Note: Boris and his family were stranded on a desert island when he was a young boy.*)
- What this tells us about human knowledge
  - Depends on abstract knowledge about relations.
  - Abstractions must be probabilistic.
  - Knowledge representation and efficient inference on the scale of common-sense reasoning is not going to be easy.

Language

INPUT		OUTPUT	
0000000000000000000000000000010	(woman)	0000000000000000000000000010000	(smash)
0000000000000000000000000010000	(smash)	0000000000000000000001000000000	(plate)
0000000000000000000001000000000	(plate)	0000100000000000000000000000000	(cat)
0000100000000000000000000000000	(cat)	0000000000000000000100000000000	(move)
0000000000000000000100000000000	(move)	0000000000000000100000000000000	(man)
0000000000000000100000000000000	(man)	0001000000000000000000000000000	(break)
0001000000000000000000000000000	(break)	0000100000000000000000000000000	(car)
0000100000000000000000000000000	(car)	0100000000000000000000000000000	(boy)
0100000000000000000000000000000	(boy)	0000000000000000000100000000000	(move)
0000000000000000000100000000000	(move)	0000000000001000000000000000000	(girl)
0000000000001000000000000000000	(girl)	0000000000100000000000000000000	(eat)
0000000000100000000000000000000	(eat)	0010000000000000000000000000000	(bread)
0010000000000000000000000000000	(bread)	0000000010000000000000000000000	(dog)
0000000010000000000000000000000	(dog)	0000000000000000000100000000000	(move)
0000000000000000000100000000000	(move)	0000000000000000001000000000000	(mouse)
0000000000000000001000000000000	(mouse)	0000000000000000001000000000000	(mouse)
0000000000000000001000000000000	(mouse)	0000000000000000000100000000000	(move)
0000000000000000000100000000000	(move)	1000000000000000000000000000000	(book)
1000000000000000000000000000000	(book)	0000000000000001000000000000000	(lion

Elman (1990): Simple recurrent network (SRN) trained to predict the next word in a sentence.

Language

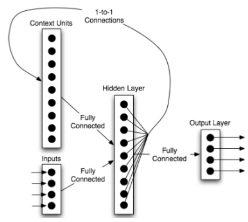

Elman (1990): Simple recurrent network (SRN) trained to predict the next word in a sentence.

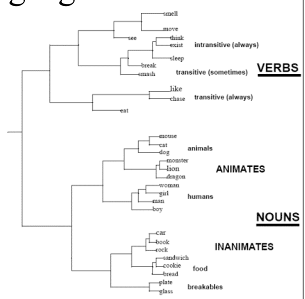

Hierarchical clustering on hidden layer activation vectors.

Semantics with predicate logic

- Bill loves Mary.
  - loves(Bill,Mary)
- Bill thinks that John loves Mary.
  - thinks(Bill,loves(John,Mary)
- Bill thinks that all guys love Mary.
  - thinks(Bill,f.a. *x* guy(*x*) loves(*x*,Mary))
- Mary knows that Bill thinks all guys love her.
  - knows(Mary,thinks(Bill,f.a. *x* guy(*x*) loves(*x*,Mary)))
- Bill is afraid that Mary knows that he thinks all guys love her.
  - afraid(Bill,knows(Mary,thinks(Bill, f.a.*x* guy(*x*) loves(*x*, Mary))))
- Mary wonders if Bill realizes that she knows he thinks all guys love her….

Semantics with predicate logic

If a burkle tumps that one of its gazzers will glip one of its rupples, then the burkle will prin that gazzer.

Outline

- The traditional debate in cognitive science: logic *versus* probability
  - The case of connectionism
  - Examples: knowledge about biology, social relations, language, visual objects
- Models of human reasoning that integrate probabilistic inference and logical representations
  - Ecological reasoning (Shafto, Kemp, et al.)
  - Causal reasoning
- Other directions
  - Social reasoning, physical reasoning

Property induction

Gorillas have T9 hormones.
Seals have T9 hormones.
Squirrels have T9 hormones.

Horses have T9 hormones.

"Similarity", "Typicality", "Diversity"

Gorillas have T9 hormones.
Seals have T9 hormones.
Squirrels have T9 hormones.

Flies have T9 hormones.

Gorillas have T9 hormones.
Chimps have T9 hormones.
Monkeys have T9 hormones.
Baboons have T9 hormones.

Horses have T9 hormones.

Beyond similarity-based induction

- Reasoning based on dimensional thresholds: (Smith et al., 1993)

Poodles can bite through wire.
German shepherds can bite through wire.

Dobermans can bite through wire.
German shepherds can bite through wire.

- Reasoning based on causal relations: (Medin et al., 2004; Coley & Shafto, 2003)

Salmon carry E. Spirus bacteria.
Grizzly bears carry E. Spirus bacteria.

Grizzly bears carry E. Spirus bacteria.
Salmon carry E. Spirus bacteria.

Different sources for priors

Chimps have T9 hormones.
Gorillas have T9 hormones. — Taxonomic similarity

Poodles can bite through wire.
Dobermans can bite through wire. — Strength ordering

Salmon carry E. Spirus bacteria.
Grizzly bears carry E. Spirus bacteria. — Food web relations

Property type
"has T9 hormones" "can bite through wire" "carry E. Spirus bacteria"

Theory Structure
taxonomic tree + diffusion process | directed chain + drift process | directed network + noisy transmission

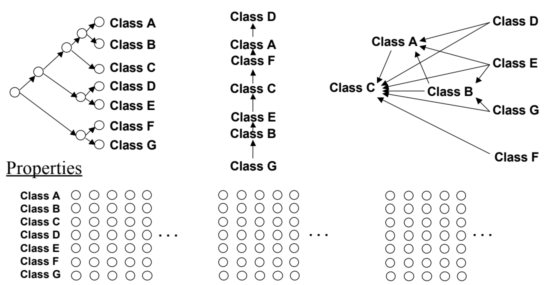

Gazelles carry E. Spirus bacteria.
Lions carry E. Spirus bacteria.

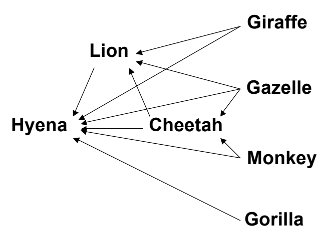

Gazelles carry E. Spirus bacteria.
Lions carry E. Spirus bacteria.

"noisy transmission process"

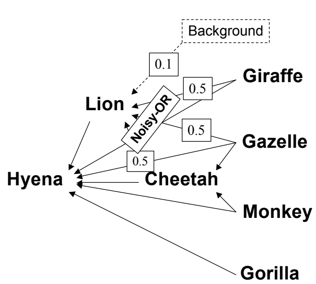

Monkeys carry Gripp's parasite.
Hyenas carry Gripp's parasite.

"noisy transmission process"

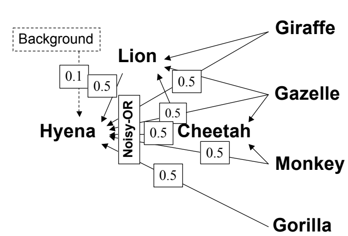

Herring carry E. Spirus bacteria.

Sand sharks carry E. Spirus bacteria.

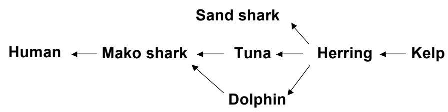

Tuna carry E. Spirus bacteria.

Mako sharks carry E. Spirus bacteria.

"noisy transmission process"

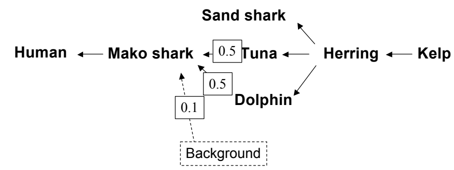

Theory as an RPM

(BLOG syntax)

Types: Species, Disease

Predicates:

Eats(species, species) non-random

Has(species, disease) random

Dependency statements:

Has(s, d) ~ NoisyORAggCPD[0.5 0.1]

({Has(s', d) for species s': Eats(s, s')});

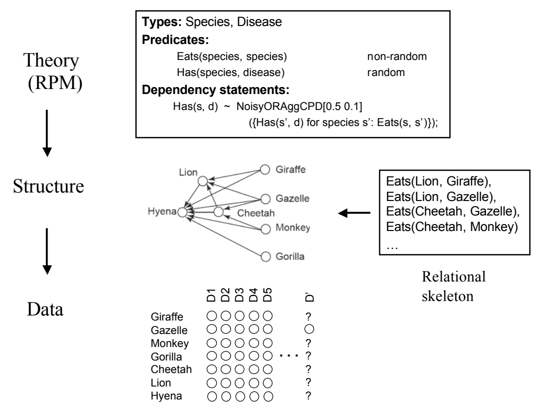

Reasoning with blank food webs

"Given that animal X has disease P, how likely is it that ainimal Y does?"

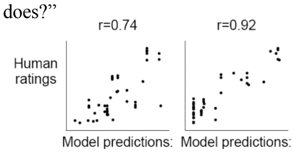

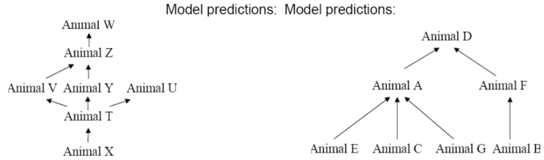

Fitting parameters of the theory

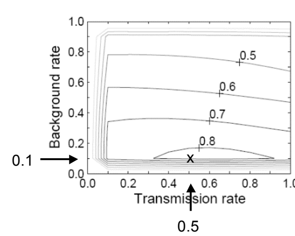

Reasoning with real species and two property types

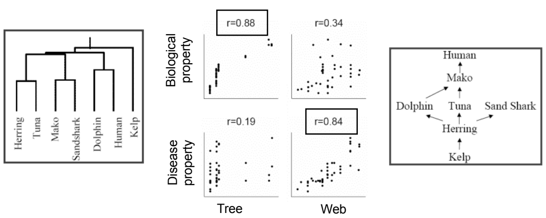

Reasoning with real species and two property types

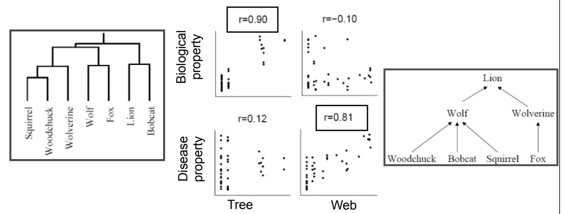

Fitting parameters of the theory

Experiment 1 Experiment 2

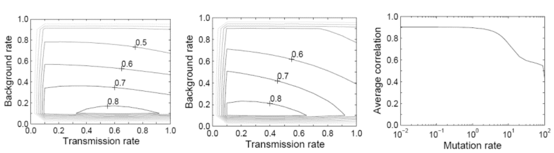

Summary: ecological reasoning

- We can write down simple intuitive theories of ecology as first-order probabilistic models (RPMs).
- Human property induction appears consistent with the use of these theories.
- Theories which fit human inference best are also most appropriate for the structure of the natural environment, both in qualitative structure and quantitative parameter values.
- No natural way of capturing this behavior with traditional with modeling approaches based purely on logic or statistical learning.

Outline

- The traditional debate in cognitive science: logic *versus* probability
  - The case of connectionism
  - Examples: knowledge about biology, social relations, language, visual objects
- Models of human reasoning that integrate probabilistic inference and logical representations
  - Ecological reasoning
  - Causal reasoning (Griffiths, Kemp, Goodman et al.)
- Other directions
  - Social reasoning, physical reasoning

"Backwards blocking"

(Sobel, Tenenbaum & Gopnik, 2004)

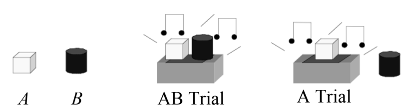

- Two objects: A and B
- Trial 1: A B on detector – detector active
- Trial 2: A on detector – detector active
- 4-year-olds judge whether each object is a blicket
  - A: a blicket (100% say yes)
  - B: probably not a blicket (34% say yes)

All logically hypotheses

A = 1 if Contact(block A, detector), else 0
B = 1 if Contact(block B, detector), else 0
E = 1 if Active(detector), else 0

$$P(\boxed{A \rightarrow E} \mid d) = \sum_{h_j \in H} P(A \rightarrow E \mid h_j) P(h_j \mid d)$$

"A is a blicket"

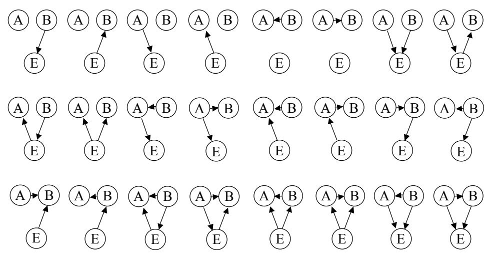

Theory-based hypothesis space

A = 1 if Contact(block A, detector), else 0
B = 1 if Contact(block B, detector), else 0
E = 1 if Active(detector), else 0

[redacted]

"A is a blicket"

	$P(h_{00}) = (1-q)^2$	$P(h_{01}) = (1-q)\,q$	$P(h_{10}) = q(1-q)$	$P(h_{11}) = q^2$

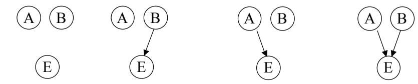

$P(E{=}1 \mid A{=}0, B{=}0)$:	0	0	0	0
$P(E{=}1 \mid A{=}1, B{=}0)$:	0	0	1	1
$P(E{=}1 \mid A{=}0, B{=}1)$:	0	1	0	1
$P(E{=}1 \mid A{=}1, B{=}1)$:	0	1	1	1

Theory as an RPM
(BLOG syntax)

Types: Block, Detector, Trial
Predicates:

Contact(block, detector, trial)	non-random
Activates(block, detector)	random
Active(detector, trial)	random

Dependency statements:

Activates(b, d) ~ TabularCPD[[(1-q) q]];
Active(d, t) ~ NoisyORAggCPD[1 0]
({Contact(b, d, t) for block b: Activates(b, d)});
OR:
Active(d, t) ~ Exists Block b: Contact(b, d, t) and Activates(b, d);

Theory
(RPM)

Types: Block, Detector, Trial
Predicates:

Contact(block, detector, trial)	non-random
Activates(block, detector)	random
Active(detector, trial)	random

Dependency statements:

Activates(b, d) ~ TabularCPD[[q (1-q)]];
Active(d, t) ~ NoisyORAggCPD[1 0]
({Contact(b, d, t) for block b: Activates(b, d)});

Structure

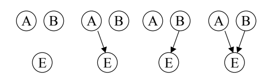

Activates(A, detector)?
Activates(B, detector)?

Relational
Skeleton
(unknown)

Data

Contact(A, detector, t1)
Contact(B, detector, t1)
Active(detector, t1)

Contact(A, detector, t2)
~Contact(B, detector, t2)
Active(detector, t2)

Manipulating plausibility
(n = 12 per condition)

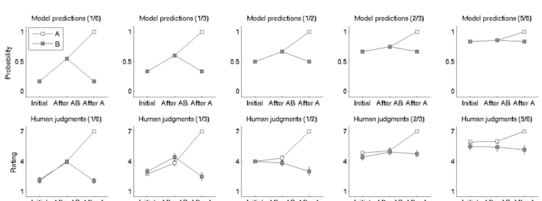

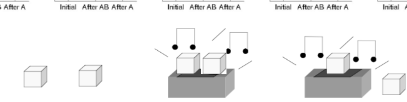

Initial AB Trial A Trial

Inferences from ambiguous data

I. Pre-training phase: Blickets are rare . . . .

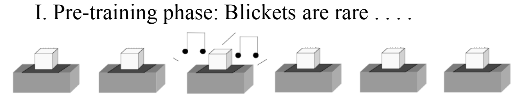

II. Two trials: A B → detector, B C → detector

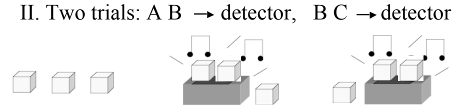

A *B* *C* Trial 1 Trial 2

After each trial, adults judge the probability that each object is a blicket.

Same domain theory generates hypothesis space for 3 objects:

- Hypotheses:

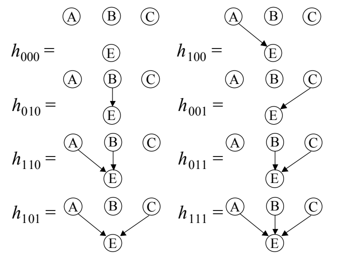

- Likelihoods: $P(E=1 \mid A, B, C; h) = 1$ if $A = 1$ and $A \rightarrow E$ exists, or $B = 1$ and $B \rightarrow E$ exists, or $C = 1$ and $C \rightarrow E$ exists, else 0.

- "Rare" condition: First observe 12 objects on detector, of which 2 set it off.

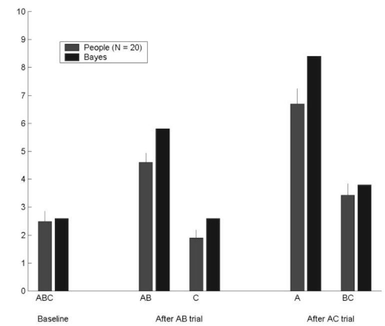

Manipulating the priors of 4-year-olds

(Sobel, Tenenbaum & Gopnik, 2004)

I. Pre-training phase: Blickets are rare.

II. Backwards blocking phase:

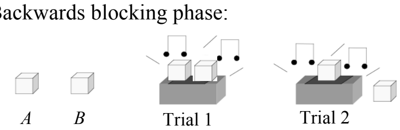

Rare condition:

A: 100% say "a blicket"

B: 25% say "a blicket"

Common condition:

A: 100% say "a blicket"

B: 81% say "a blicket"

Ambiguous data with 4-year-olds

I. Pre-training phase: Blickets are rare.

II. Two trials: A B → detector, B C → detector

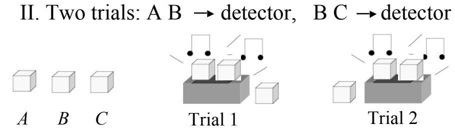

Final judgments:

A: 87% say "a blicket"

B or *C*: 56% say "a blicket"

Outline

- The traditional debate in cognitive science: logic *versus* probability
  - The case of connectionism
  - Examples: knowledge about biology, social relations, language, visual objects
- Models of human reasoning that integrate probabilistic inference and logical representations
  - Ecological reasoning
  - Causal reasoning (Griffiths, Kemp, Goodman et al.)
- Other directions
  - Social reasoning, physical reasoning

Propositional Theory (Deterministic)

- Scenario with students, courses, profs

  Dr. Pavlov teaches CS1 and CS120
  Matt takes CS1
  Judy takes CS1 and CS120

- Propositional theory

PavlovDemanding → CS1Hard
PavlovDemanding → CS120Hard

CS1Hard → MattTired
CS1Hard → JudyTired
CS120Hard → JudyTired

¬CS1Hard → MattGetsAInCS1
¬CS1Hard → JudyGetsAInCS1
¬CS120Hard → JudyGetsAInCS120

MattSmart ∧ CS1Hard → MattGetsAInCS1
JudySmart ∧ CS1Hard → JudyGetsAInCS1
JudySmart ∧ CS120Hard → JudyGetsAInCS120

Social Reasoning

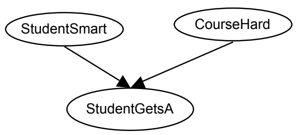

Social Reasoning

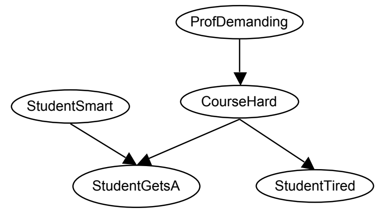

Probabilistic Theory (Propositional)

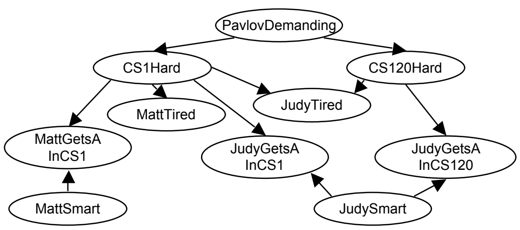

- Specific to particular scenario (who takes what, etc.)
- No generalization of knowledge across objects

First-Order Theory

- General theory:
  - $\forall$ p $\forall$ c [Teaches(p, c) $\wedge$ Demanding(p) $\rightarrow$ Hard(c)]
  - $\forall$ s $\forall$ c [Takes(s, c) $\wedge$ Hard(c) $\rightarrow$ Tired(s)]
  - $\forall$ s $\forall$ c [Takes(s, c) $\wedge$ Easy(c) $\rightarrow$ GetsA(s, c)]
  - $\forall$ s $\forall$ c [Takes(s, c) $\wedge$ Hard(c) $\wedge$ Smart(s) $\rightarrow$ GetsA(s, c)]
- Relational skeleton:

Teaches(Pavlov, CS1)	Teaches(Pavlov, CS120)
Takes(Matt, CS1)	
Takes(Judy, CS1)	Takes(Judy, CS120)

- Compact, generalizes across scenarios and objects
- But deterministic

First-Order Probabilistic Model

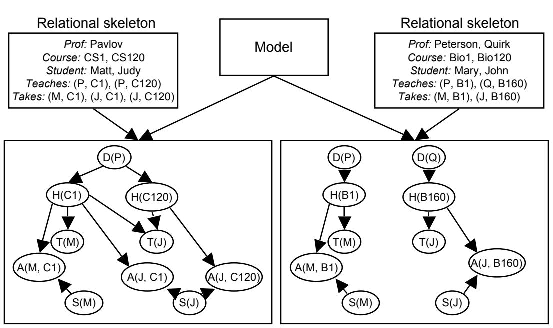

RPM (BLOG syntax)

Types: Professor, Course, Student

Predicates:

TaughtBy(course) $\rightarrow$ professor	non-random
Takes(student, course)	non-random
Demanding(professor)	random
Smart(student)	random
Hard(course)	random
Tired(student)	random
GetsA(student, course)	random

RPM (BLOG syntax)

Dependency statements:

```
Demanding(p) ~ TabularCPD[[0.8 0.2]];
Smart(s) ~ TabularCPD[[0.3 0.7]];
Hard(c) ~ TabularCPD[[0.6 0.4],
                     [0.1 0.9]]
                     (Demanding(TaughtBy(c)));
Tired(s) ~ CumGeomCPD[0.5]
           (#TRUE({Hard(c) for course c: Takes(s, c)}));
GetsA(s, c) ~ TabularCPD[[0.5 0.5],
                         [0.1 0.9],
                         [0.9 0.1],
                         [0.7 0.3]]
                        (Hard(c), Smart(s));
```

Generative Process for Aircraft Tracking

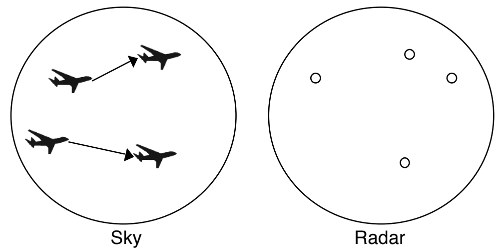

BLOG Model for Aircraft Tracking

World ...

```
#Aircraft ~ NumAircraftDistrib();
State(a, t)
  if t = 0 then ~ InitState()
  else ~ StateTransition(State(a, Pred(t)));
```

Observations

```
#Blip: (Source, Time) -> (a, t)
  ~ NumDetectionsDistrib(State(a, t));
#Blip: (Time) -> (t)
  ~ NumFalseAlarmsDistrib();
ApparentPos(r)
  if (Source(r) = null) then ~ FalseAlarmDistrib()
  else ~ ObsDistrib(State(Source(r), Time(r)));
```

Apparent motion

- Visual system parses ambiguous experience into objects under several assumptions:
  - Objects typically do not disappear and appear spontaneously.
  - Objects typically follow "simple" space-time trajectories.

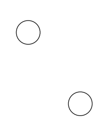

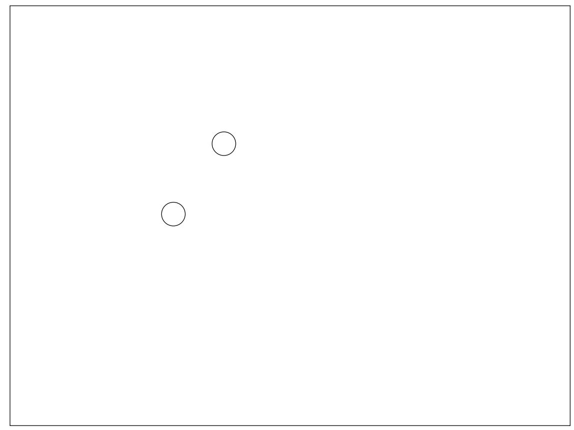

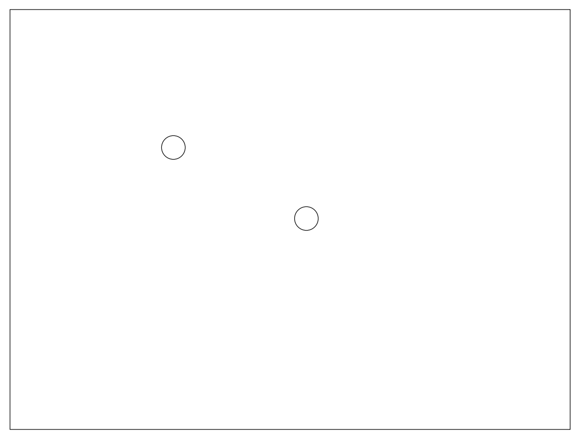

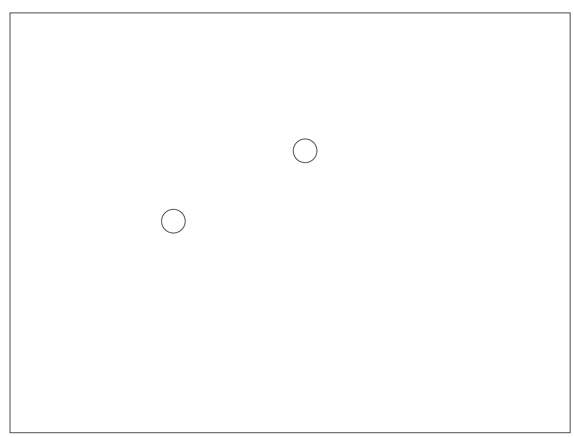

Not so simple

- Parsing depends on
  - inferences about occlusion and visibility.
  - dynamic interactions among (potentially invisible) objects.
  - inferences about object shape, color, and other static properties.

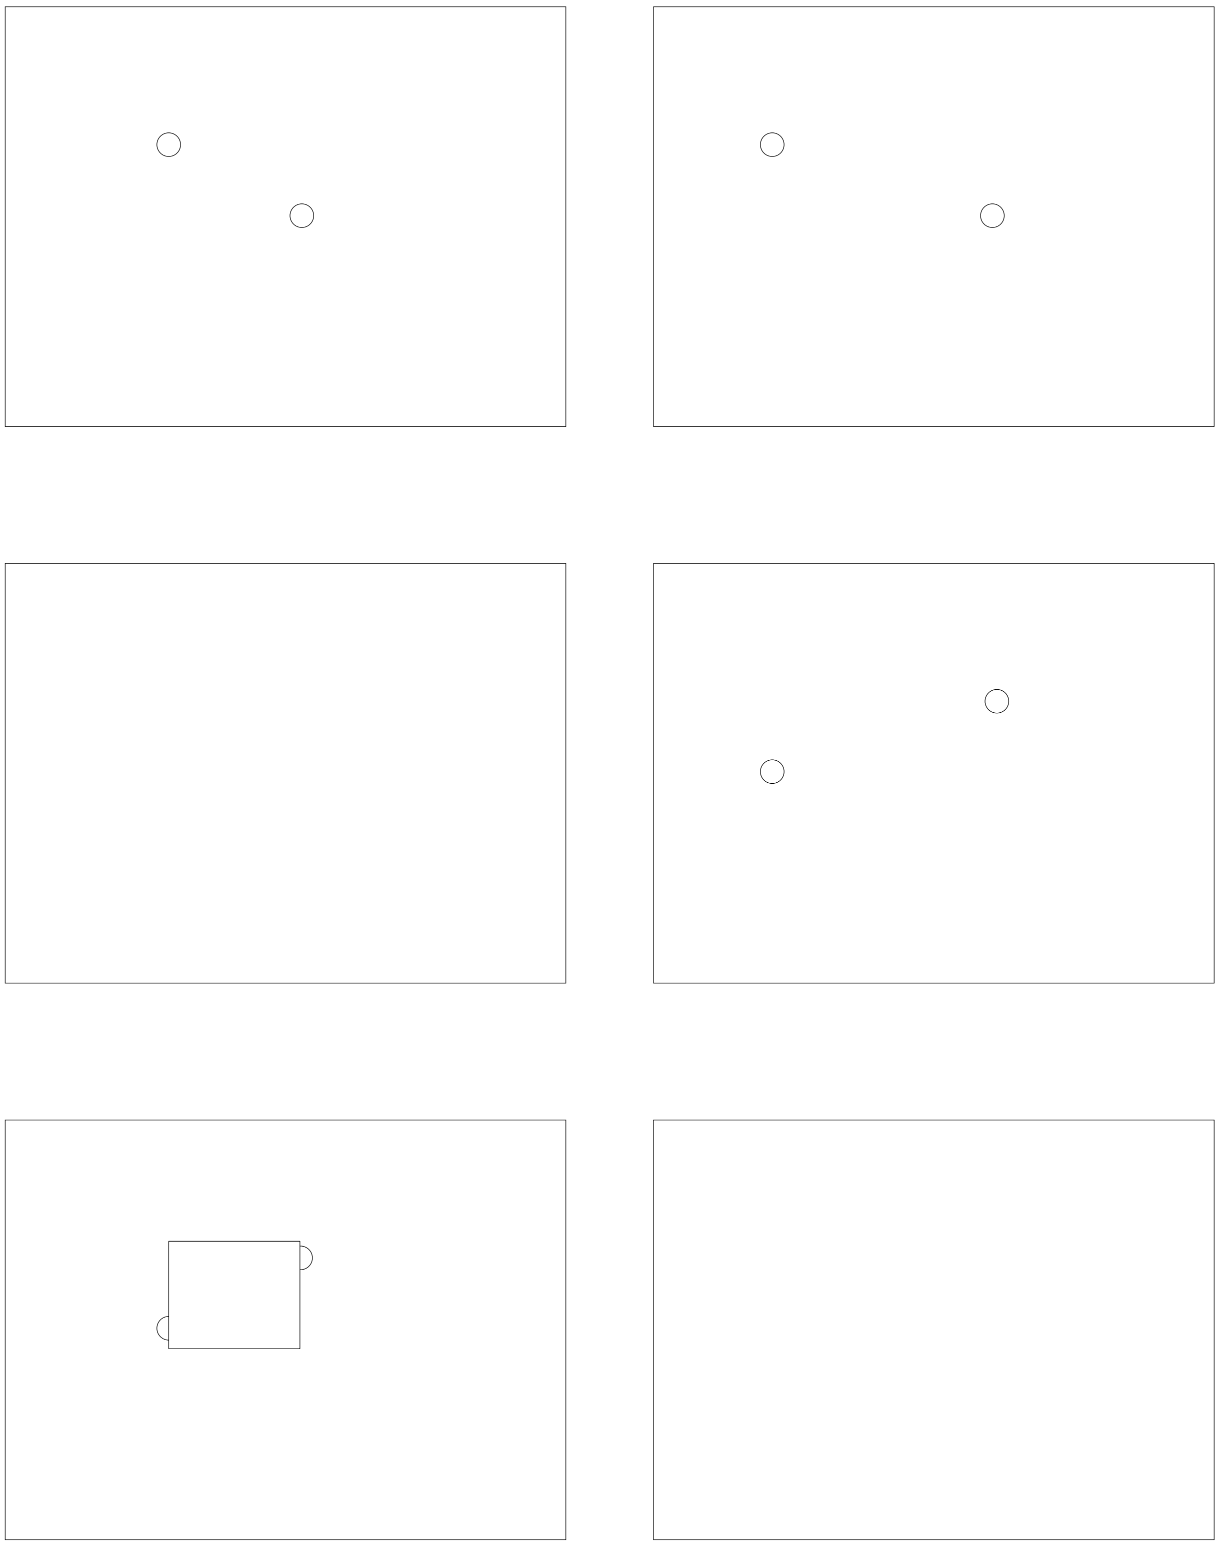

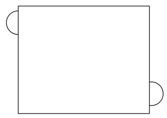

Not so simple

- Parsing depends on
  - inferences about occlusion and visibility.
  - dynamic interactions among (potentially invisible) objects.
  - inferences about object shape, color, and other static properties.
- Theory of objects must allow uncertainty in how many objects are in the scene.

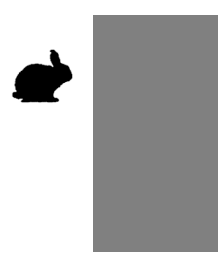

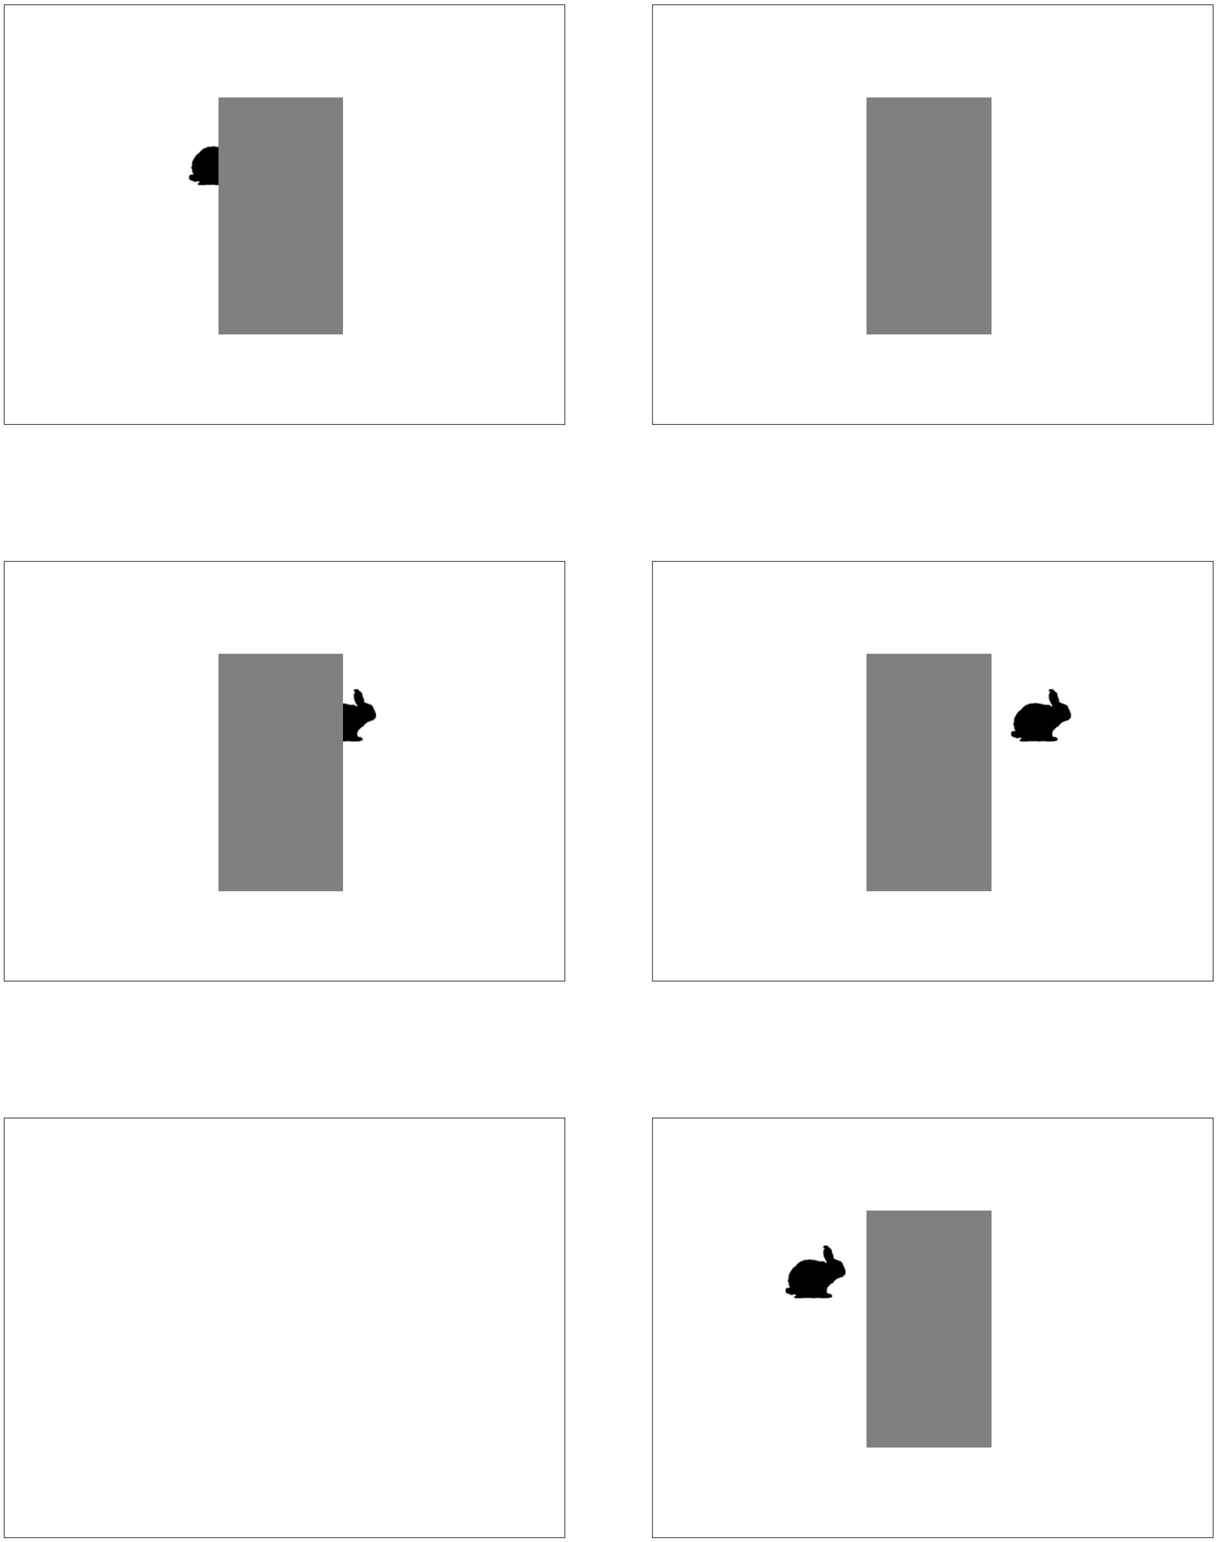

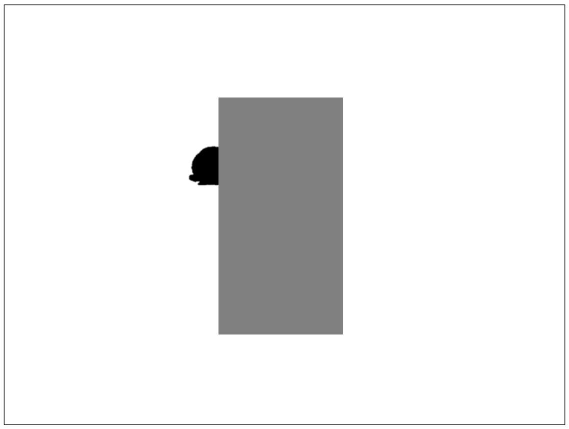

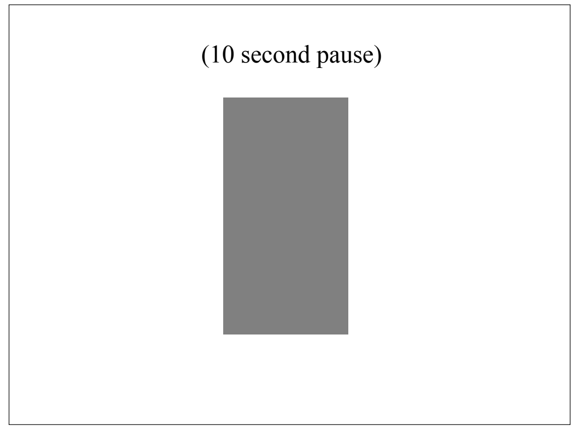
(10 second pause)

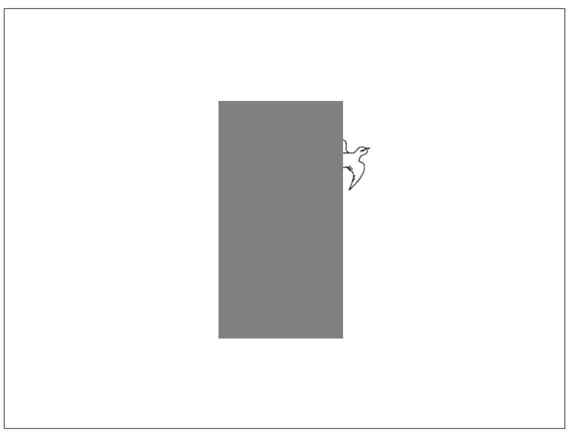

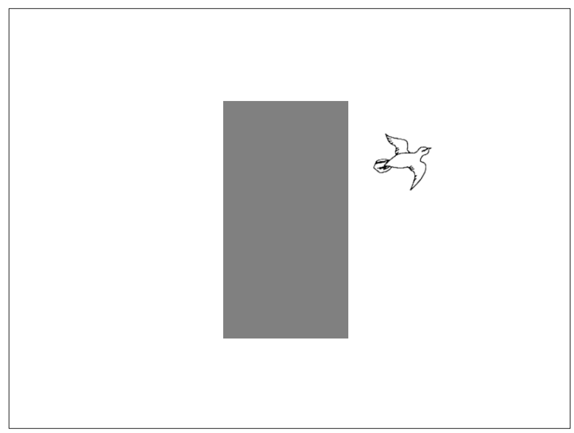

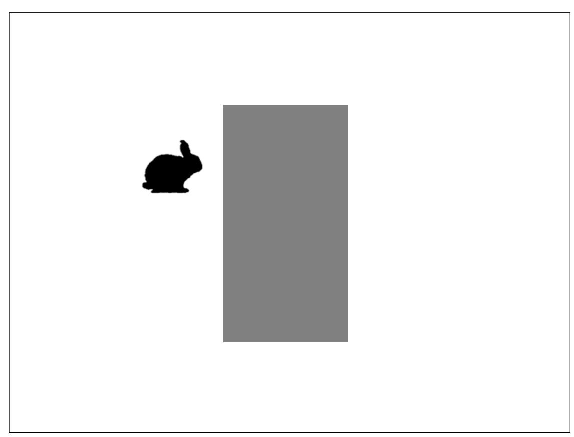

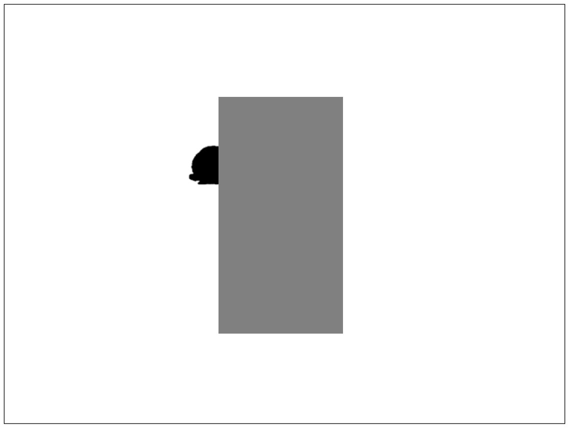

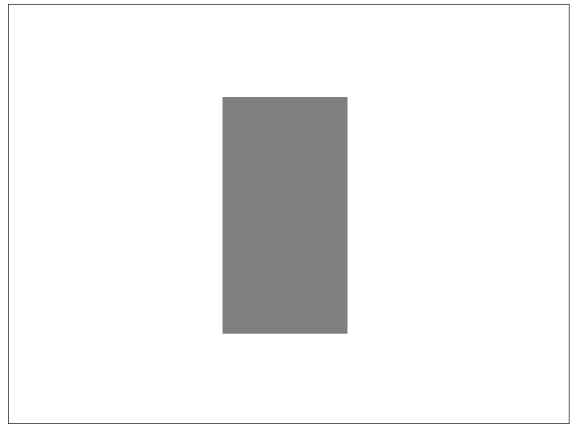

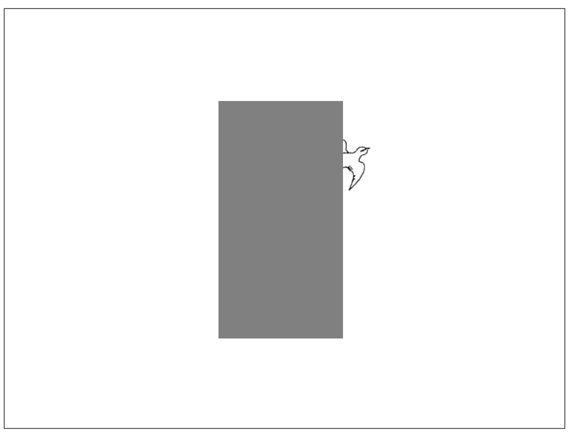

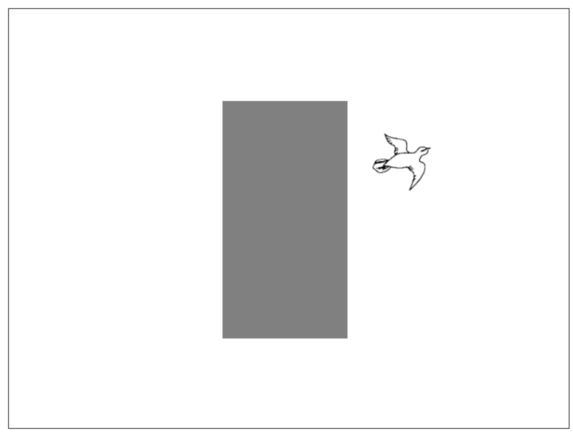

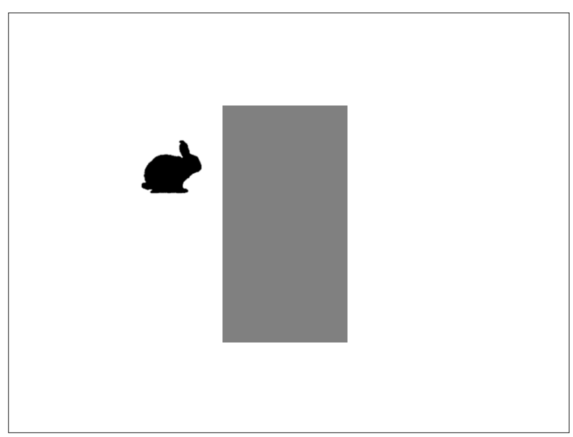

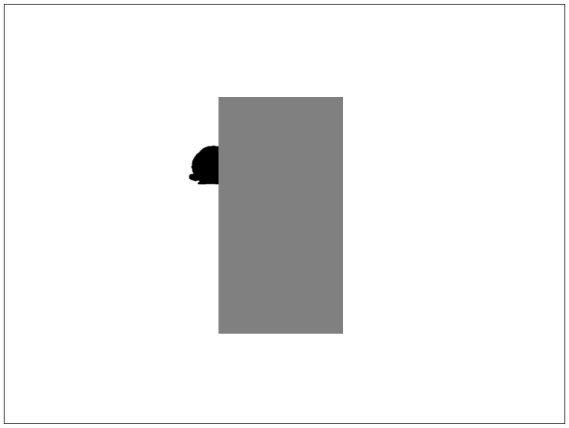

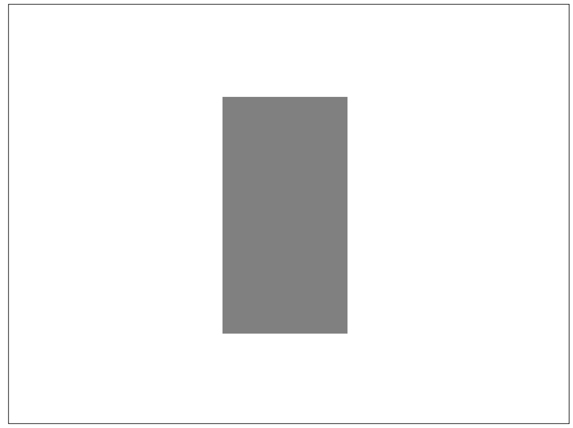

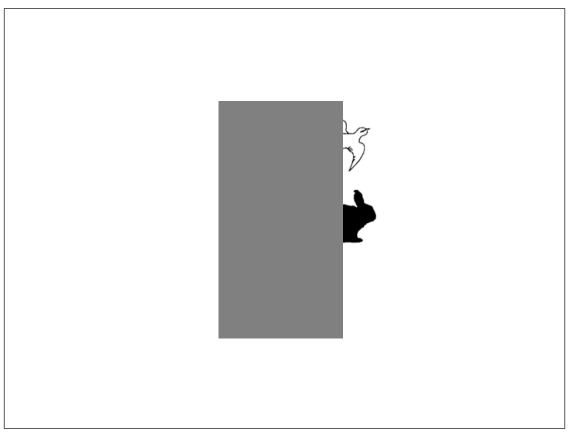

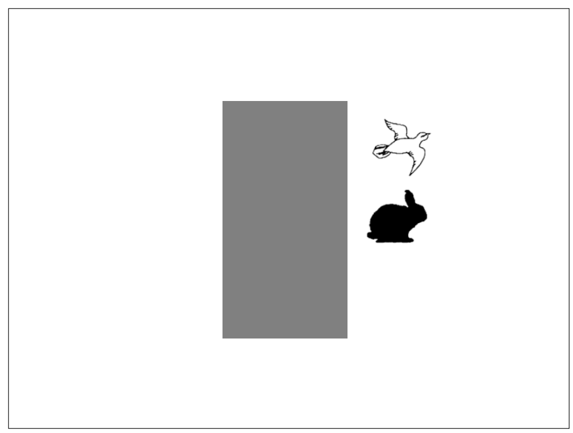

Looking time studies of infants' theory of objects

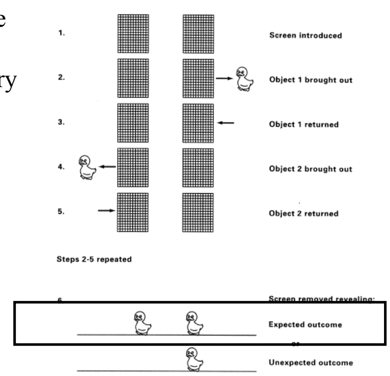

Looking time studies of infants' theory of objects

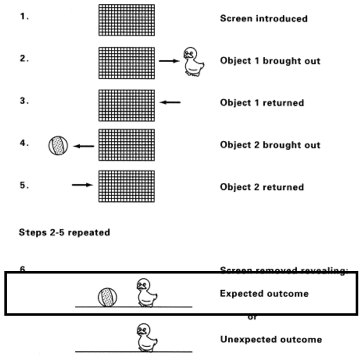

12 months

Looking time studies of infants' theory of objects

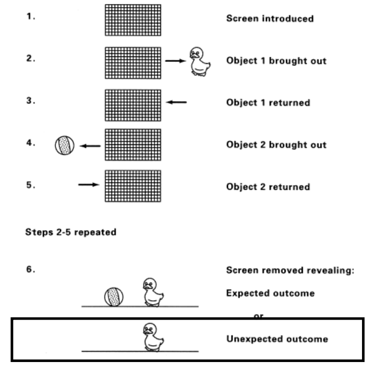

10 months

Big open problems

- Use RPMs, BLOG, etc. to…
  - Formalize intuitive social or physical reasoning.
  - Explain human inferences in social or physical domains.
  - Explain the differences between children's theories at different ages.
  - Explain how children learn these theories.
  - Elucidate core representational capacities of human cognition.