Outline

- Theory-based Bayesian framework for property induction
- Causal structure induction
  - Constraint-based (bottom-up) learning
  - Theory-based Bayesian learning

The Bayesian approach

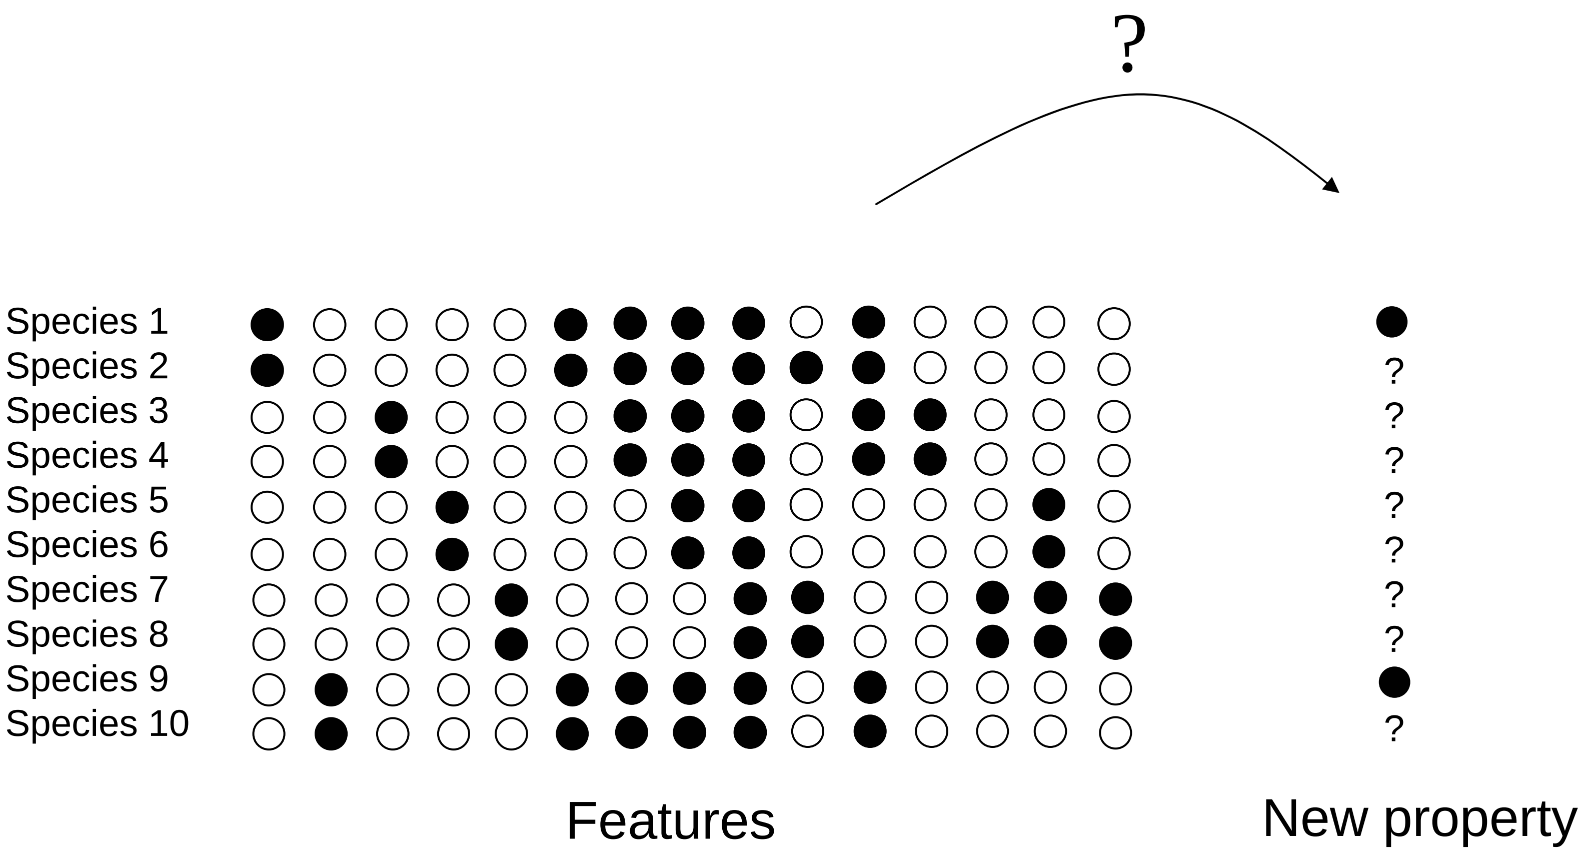

Mutation process generates $p(h|T)$:

- Choose label for root.
- Probability that label mutates along branch b : $\dfrac{1-e^{-2\lambda|b|}}{2}$

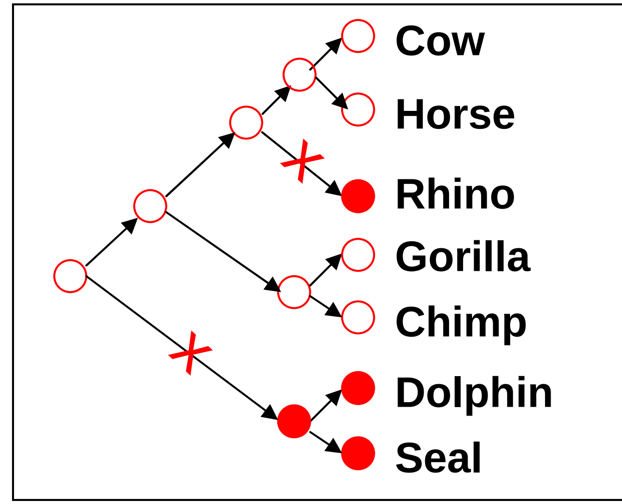

λ = mutation rate

$|b|$ = length of branch b

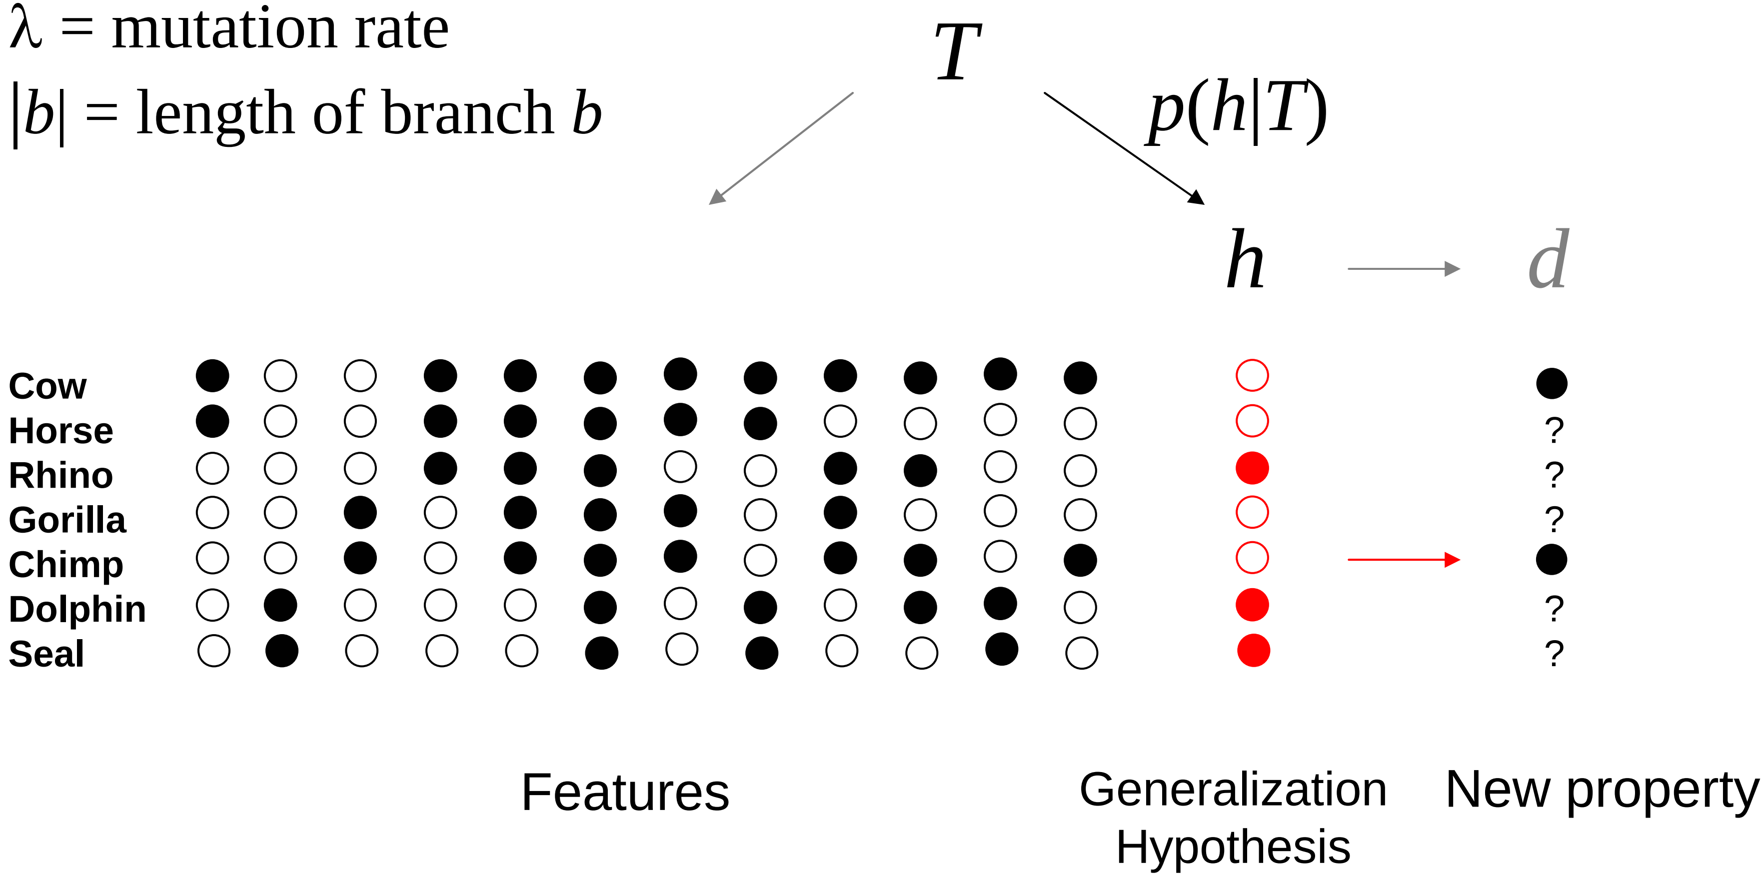

Results

"Theory-based"
Bayes

Images removed due to
copyright considerations.

"Empiricist"
Bayes

Max-sim

A Bayesian dream

Prior based on mutations over tree structure addresses all the challenges to traditional Bayesian concept learning (Mitchell, Tenenbaum, etc.)

- Assign a reasonable prior over all logically possible concepts (labelings) in a potentially unbounded domain, with natural Occam’s razor.
- Efficiently integrate over all logically possible concepts consistent with the training data.
- Robust with respect to label noise.
- PAC-style guarantees of generalization.

Bayes with alternative theories

- Taxonomic Bayes (strictly taxonomic hypotheses, with no mutation process)

Results

Theory-based Bayes

Bias is just right!

Taxonomic Bayes

Images removed due to copyright considerations.

Bias is too strong

"Empiricist" Bayes

Bias is too weak

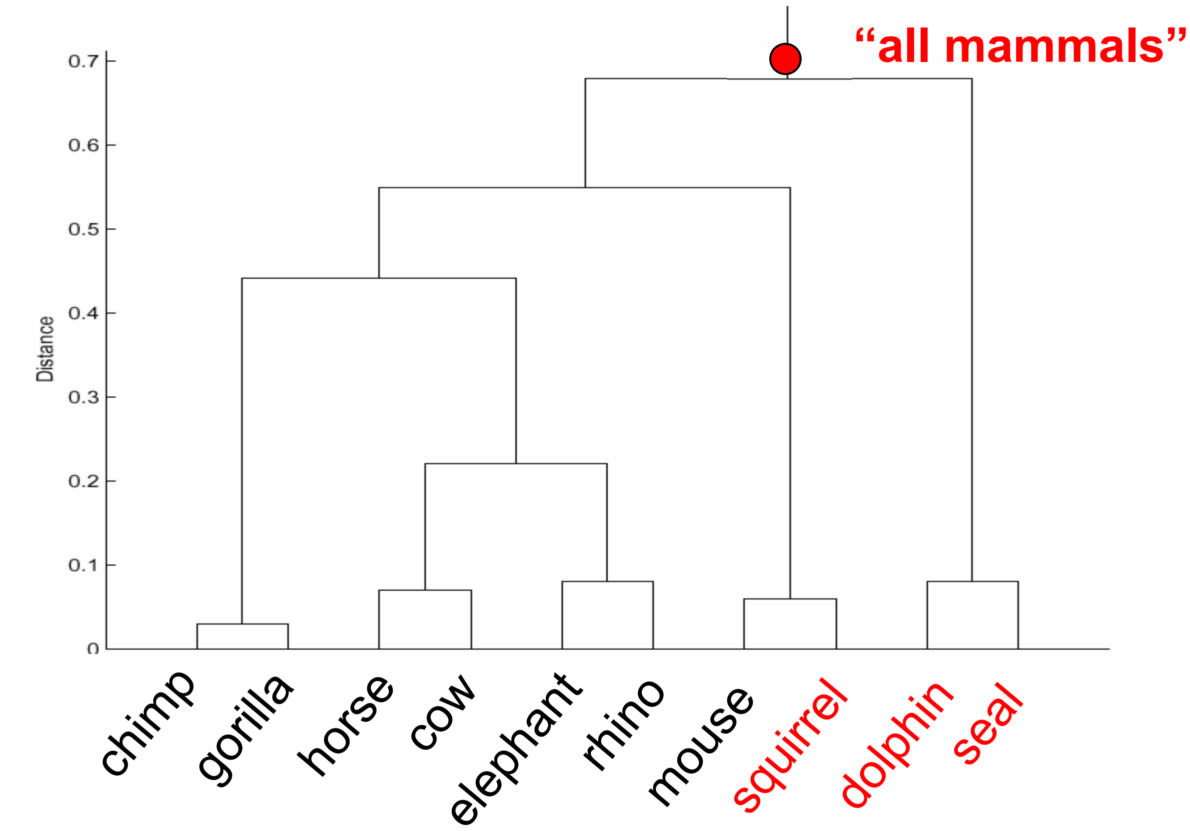

Cows have property P.
Dolphins have property P.
Squirrels have property P.

All mammals have property P.

Strong: 0.76 [max = 0.82]

Seals have property P.
Dolphins have property P.
Squirrels have property P.

All mammals have property P.

Weak: 0.30 [min = 0.14]

Bayes with alternative theories

- Taxonomic Bayes (strictly taxonomic hypotheses, with no mutation process)
- Theory-based Bayes using actual evolutionary tree.

Results

Theory-based Bayes

Bias is just right!

Theory-based Bayes w/ evolutionary tree

Images removed due to copyright considerations.

Bias is wrong

"Empiricist" Bayes

Bias is too weak

Bayes with alternative theories

- Taxonomic Bayes (strictly taxonomic hypotheses, with no mutation process)
- Theory-based Bayes using actual evolutionary tree.
- Replace mutation process with generic "Occam's Razor" prior over branches of tree:
  - e.g., p(feature changes along branch b) = λ, independent of branch length.

Bayes (taxonomy+ mutation)

Bayes (taxonomy+ Occam)

Images removed due to copyright considerations.

Max-sim

Conclusion kind: "all mammals"

Number of examples: 1

Premise typicality effect (Rips, 1975; Osherson et al., 1990):

Strong:

Horses have property P.

All mammals have property P.

Weak:

Seals have property P.

All mammals have property P.

Typicality meets hierarchies

- Collins and Quillian: semantic memory structured hierarchically

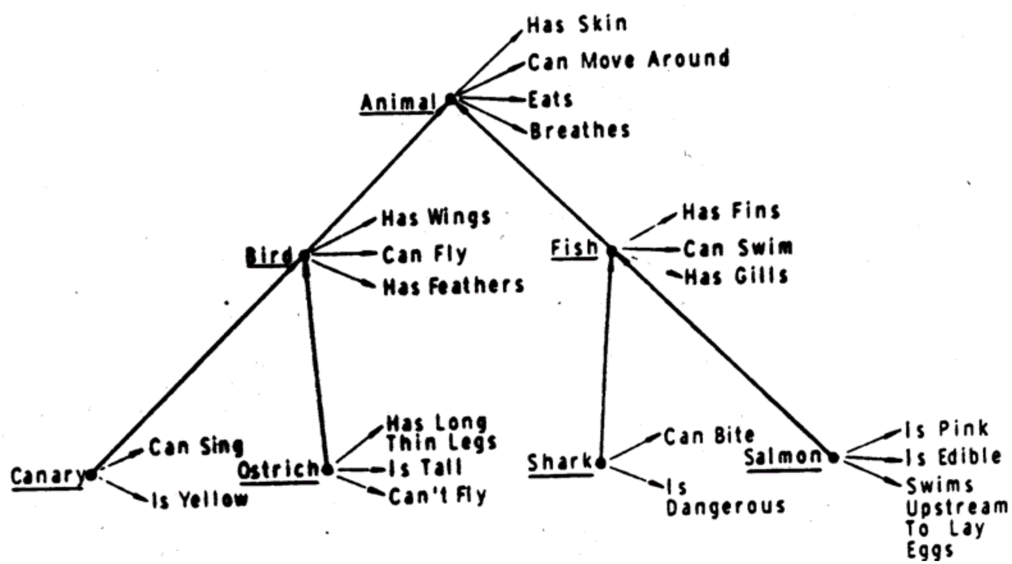

Figure of semantic trees from Quillian (1968). Quillian, M. R. "Semantic Memory." In Semantic *Information Processing*. Edited by M. Minsky. Cambridge, MA: MIT Press,1968, pp. 216-270. Courtesy of the MIT Press. Used with permission.

- Traditional story: Simple hierarchical structure uncomfortable with typicality effects & exceptions.
- New story: Typicality & exceptions compatible with rational statistical inference over hierarchy.

Bayes with alternative theories

- Taxonomic Bayes (strictly taxonomic hypotheses, with no mutation process)
- Theory-based Bayes using actual evolutionary tree.
- Replace mutation process with generic "Occam's Razor" prior over branches of tree.
- Infinite flat mixture model (essentially, Anderson's model of categorization)

Best Cluster Structure

Beaver
Otter
Rat
Weasel
Raccoon
Chihuahua
Persian Cat
Siamese Cat
Dalmatian
Collie
German Shepherd
Lion
Tiger
Leopard
Wolf
Bobcat
Fox
Polar Bear
Grizzly Bear

Cow
Pig
Ox
Sheep
Buffalo
Moose
Horse
Zebra
Antelope
Deer
Giraffe
Rhinoceros
Elephant
Hippopotamus
Giant Panda

Rabbit
Mouse
Hamster
Mole
Skunk
Squirrel

Gorilla
Chimp
Monkey
Bat

Dolphin
Seal
Humpback Whale
Blue Whale
Walrus
Killer Whale

Results with flat mixture model

Specific

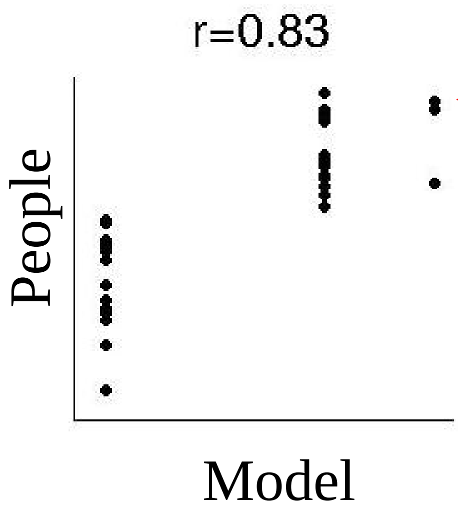

Cows can catch Disease X
Rhinos can catch Disease X

Horses can catch Disease X

General

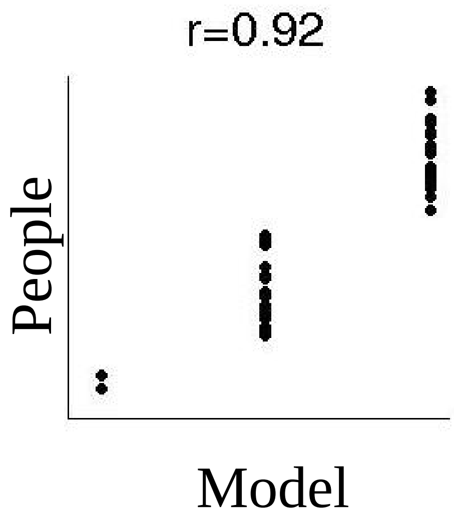

Gorillas can catch Disease X
Mice can catch Disease X
Seals can catch Disease X

All mammals can catch Disease X

Results with flat mixture model

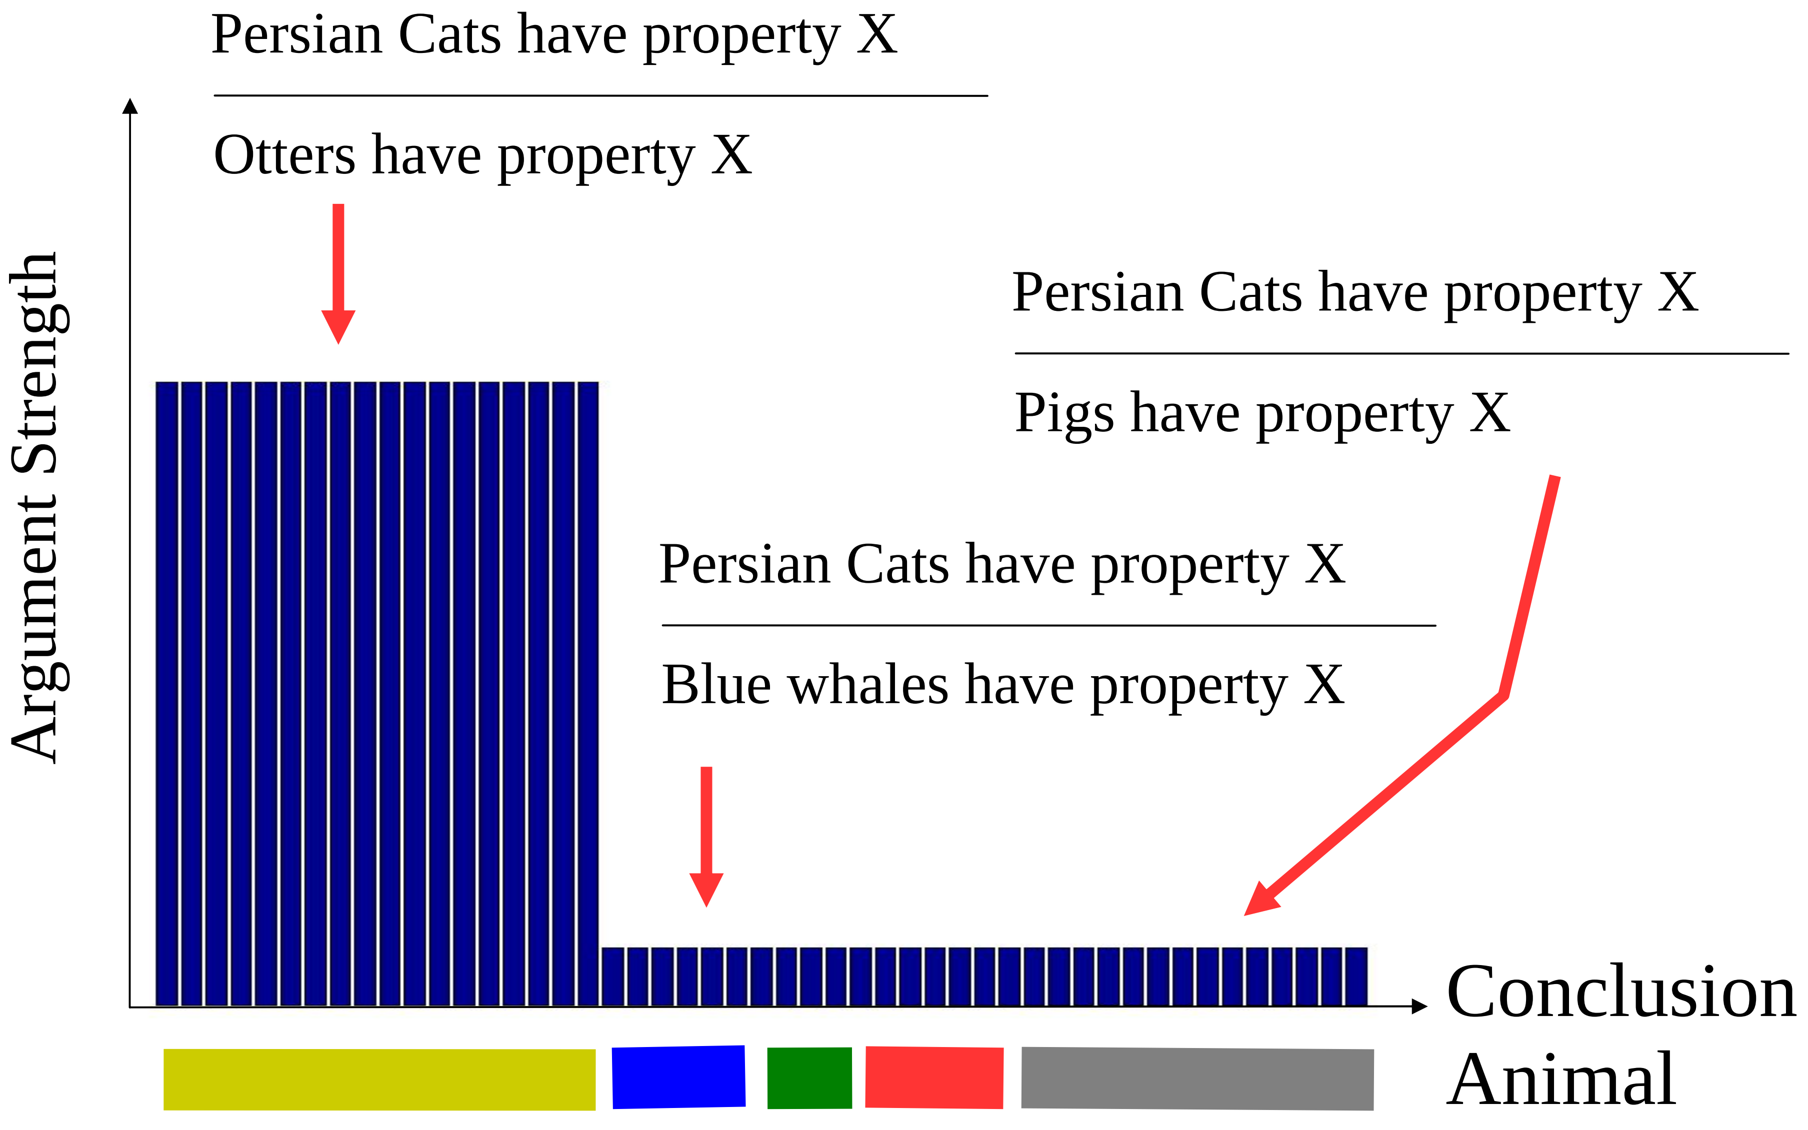

Beyond similarity-based induction

- Reasoning based on known dimensions: (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.

Beyond similarity-based induction

- Reasoning based on known dimensions: (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.

Property type

"has T4 neurons" "can bite through wire" "carry E. Spirus bacteria"

Theory type

taxonomic tree + mutation

directed chain + random threshold

directed network + noisy transmission

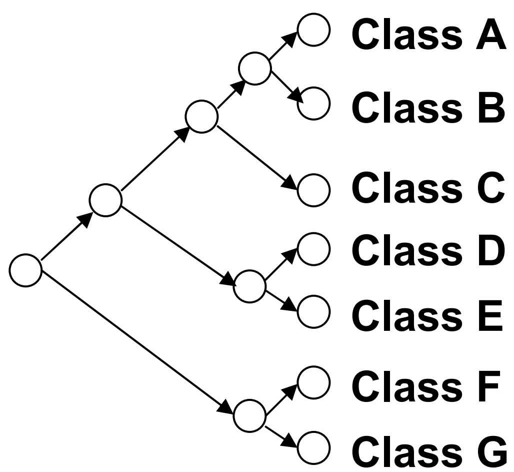

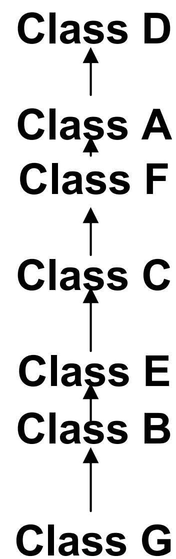

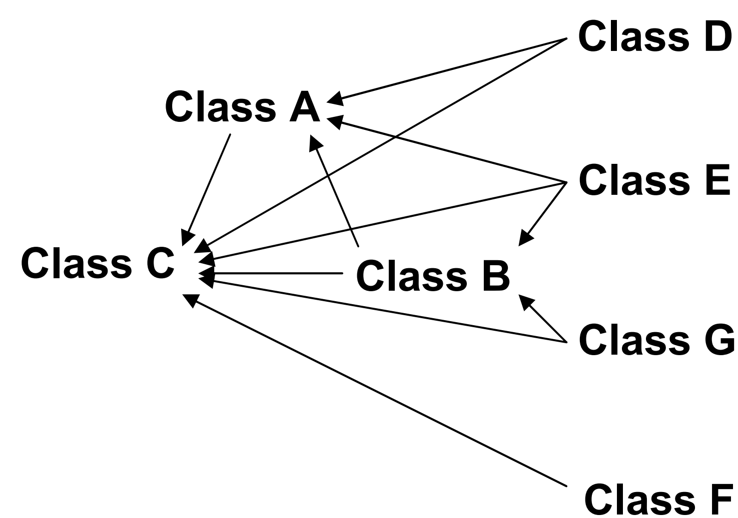

Hypotheses

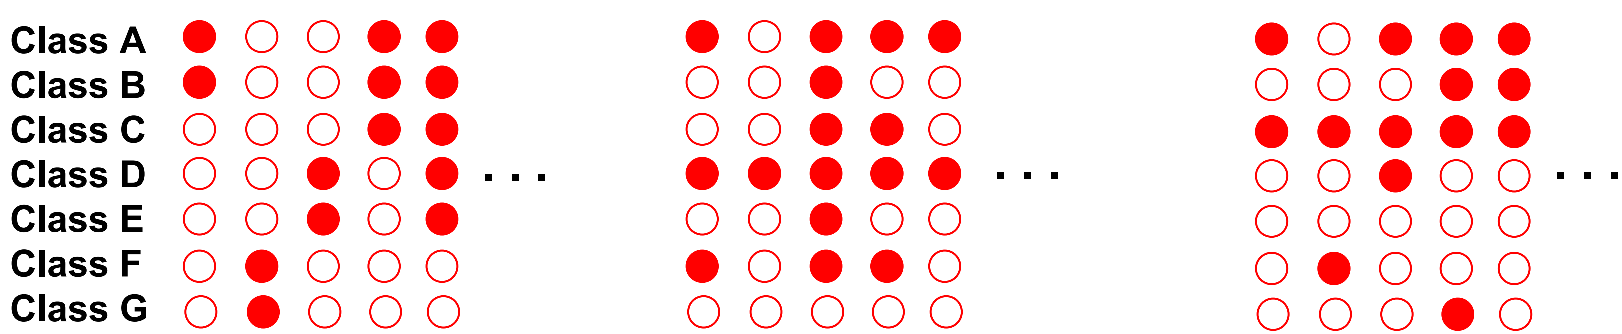

Property type

"can bite through wire"

Theory type

directed chain

+ random threshold

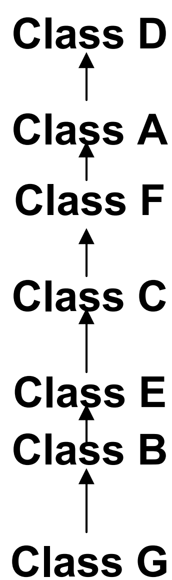

Reasoning based on known dimensions (Smith et al., 1993):

Poodles can bite through wire.

German shepherds can bite through wire.

Dobermans can bite through wire.

German shepherds can bite through wire.

Models / Datasets	Bayes (chain)	Bayes (tree)	Sim.-Cover.
Smith et al. (1993):			
- night vision			
- thick skin			
Blok et al. (2002):			
- 1 premise			
- 2 premises			
- 1 premise (pos. and neg.)			
- 2 premises (pos. and neg.)			

Datasets \ Models	Bayes (chain)	Bayes (tree)	Sim.-Cover.
Smith et al. (1993):			
- night vision	$r =$ 0.84		
- thick skin	0.94		
Blok et al. (2002):			
- 1 premise	0.97		
- 2 premises	0.98		
- 1 premise (pos. and neg.)	0.91		
- 2 premises (pos. and neg.)	0.90		

Datasets \ Models		Bayes (chain)	Bayes (tree)	Sim.-Cover.
Smith et al. (1993):				
- night vision	r =	0.84	0.49	0.51
- thick skin		0.94	0.32	0.27
Blok et al. (2002):				
- 1 premise		0.97	0.07	0.32
- 2 premises		0.98	0.47	0.47
- 1 premise (pos. and neg.)		0.91	0.46	N/A
- 2 premises (pos. and neg.)		0.90	0.67	N/A

Property type
"carry E. Spirus bacteria"

Theory type
directed network
+ noisy transmission

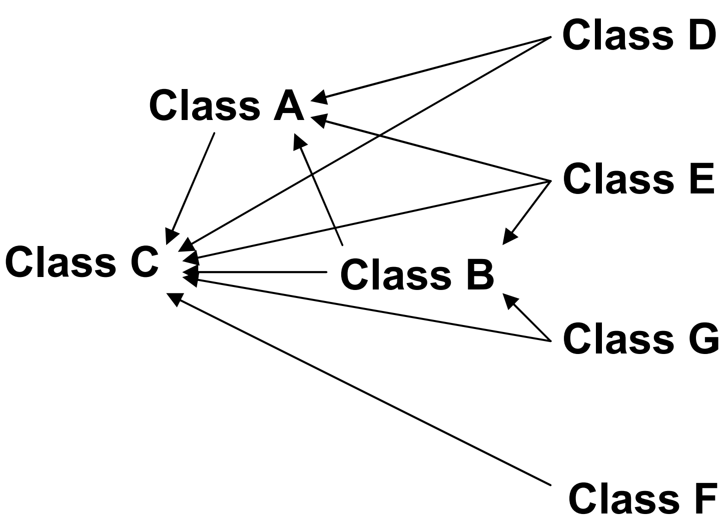

Reasoning based on causal relations (Medin et al., 2004; Coley & Shafto, 2003):

Salmon carry E. Spirus bacteria.

Grizzly bears carry E. Spirus bacteria.

Grizzly bears carry E. Spirus bacteria.

Salmon carry E. Spirus bacteria.

Property type
"carry E. Spirus bacteria"

Theory type
directed network
+ noisy transmission

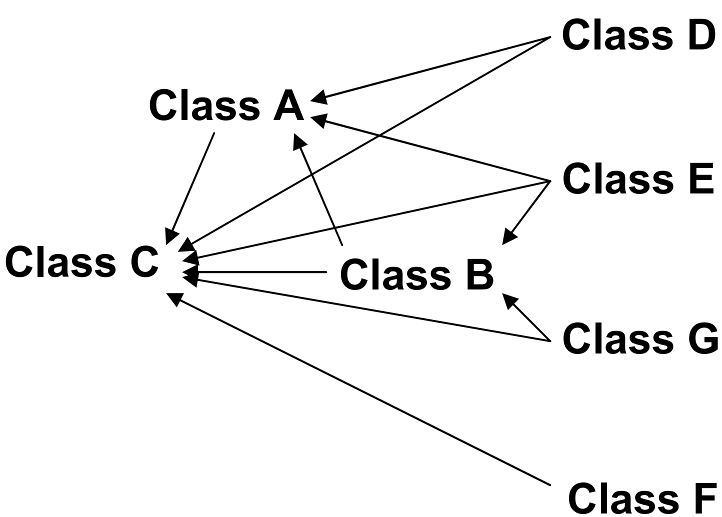

Experiment w/ Pat Shafto, Liz Baraff & John Coley:

- Participants taught two systems of relations:
  - Food web
  - Taxonomic tree
- Asked to reason about two kinds of properties:
  - Diseases
  - Genetic properties
- Two different ecosystems:
  - Mammals, Island

Island ecosystem

Taxonomy

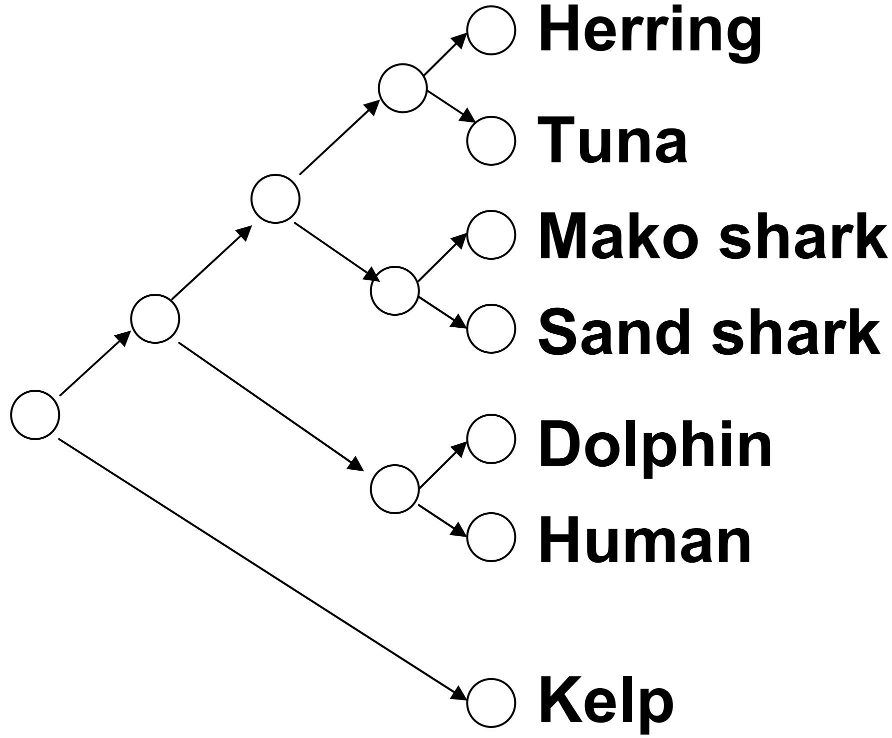

Food web

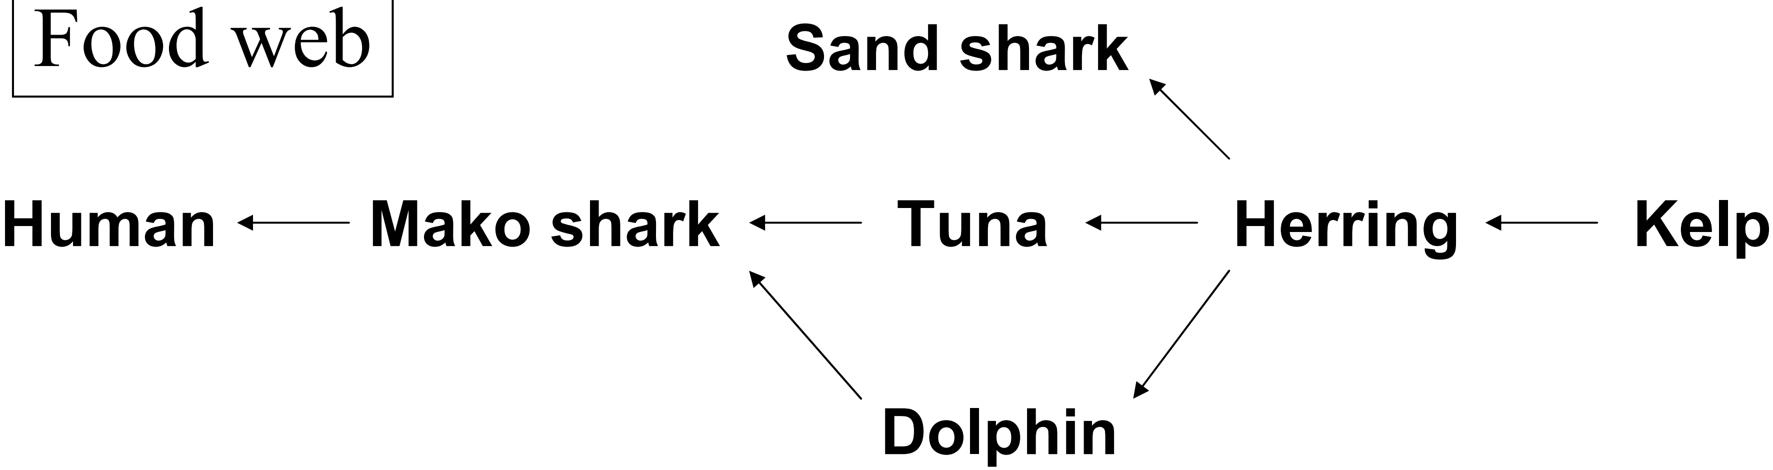

Models / Datasets	Bayes (food web)	Bayes (tree)	Sim.-Cover.
Mammal ecosystem:			
- disease			
- genetic property			
Island ecosystem:			
- disease			
- genetic property			

Datasets \ Models		Bayes (food web)	Bayes (tree)	Sim.-Cover.
Mammal ecosystem:				
- disease	$r =$	0.75	-0.15	0.07
- genetic property		0.25	0.92	0.87
Island ecosystem:				
- disease		0.79	0.01	0.17
- genetic property		0.31	0.89	0.86

Conclusions

- Beyond classic dichotomies of "domain-specific vs. domain-general", or "structured theories vs. statistical learning".
  - Bayes provides a powerful domain-general statistical engine for generalizing reliably from limited data.
  - Theories generate structured domain-specific priors that provide crucial constraints for Bayesian induction.
- Advantages of Theory-based Bayesian models:
  - Strong quantitative models of generalization behavior, with minimal free parameters or arbitrary assumptions.
  - Flexibility to model different patterns of reasoning that arise with different kinds of properties, using differently structured theories (but the same general-purpose Bayesian engine).
  - Framework for explaining *why* inductive generalization works.

Theory-based Bayesian framework

- The big picture.
  - What do we mean by “theory”?

T1 theory (c.f. theory type, structure grammar, “framework theory”)

taxonomic tree	directed chain	directed network
+ mutation	+ random threshold	+ noisy transmission

T0 theory (c.f. structure, “specific theory”)

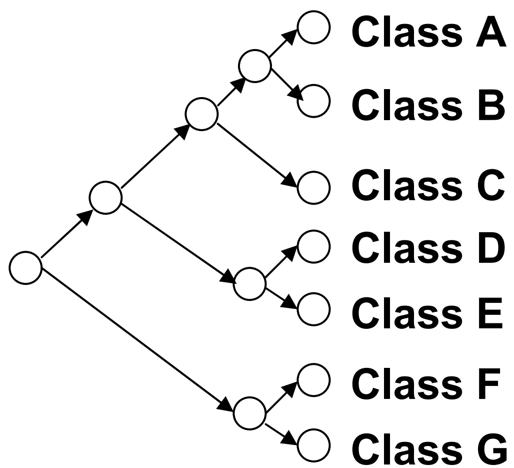

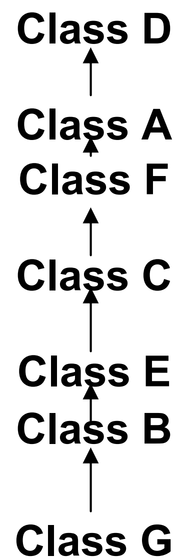

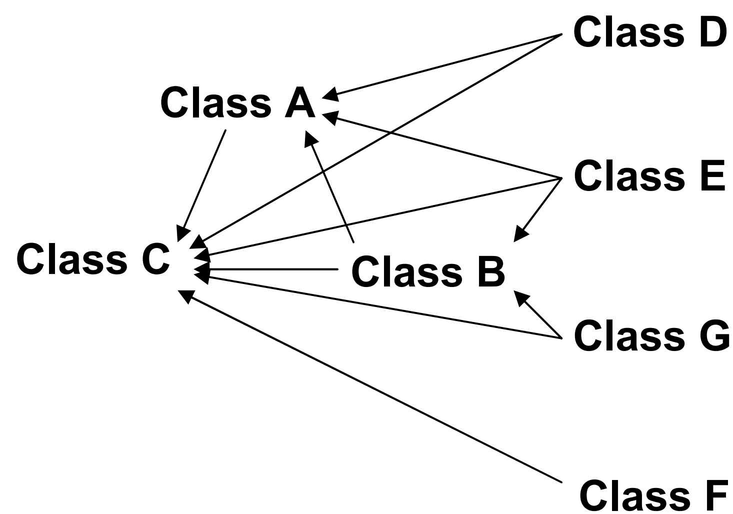

Properties

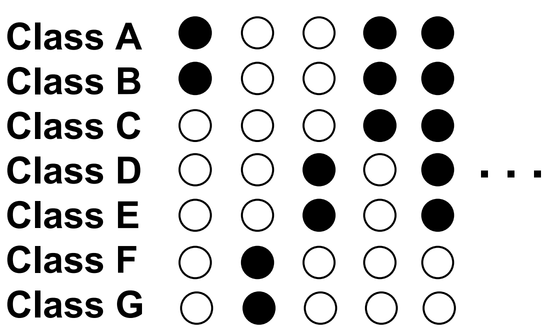

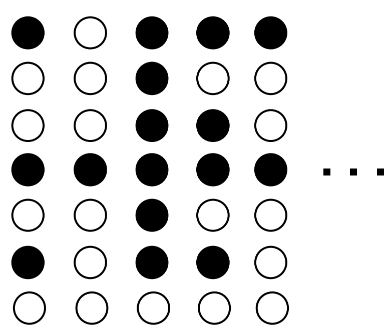

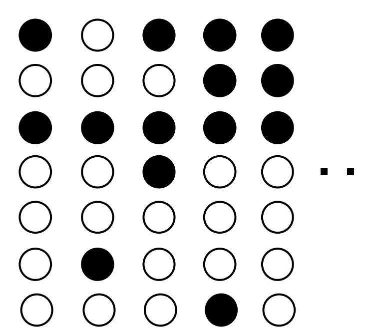

Theory-based Bayesian framework

- The big questions:
  - How are new properties learned, guided by a T0 theory?
  - How is a T0 theory learned, guided by a T1 level theory?
  - How are T1 theories learned?

Theory-based Bayesian framework

- The big questions:
  - How does a T0 theory generate a hypothesis space of properties?
  - How does a T1 theory generate a hypothesis space of T0 theories?
  - What does the hypothesis space of T1 theories look like? (i.e., what are the T2 and higher-level theories?)

Theory-based Bayesian framework

- The big questions:
  - How do we figure out which theory to use for which properties?
  - What structures and relations exist between properties?
    - Clusters, hierarchies
    - Ordered dimensions
    - Causal networks
  - How do structures over properties relate to structures over classes?