In Memoriam

NEWS.MIT.EDU

Professor Patrick Winston, former director of MIT's Artificial Intelligence Laboratory, dies at 76

Cognitive Modeling, Categorization, and Object Recognition

- This course explores computational models of basic human cognitive processes, to include categorization, inference, problem solving, decision making, learning, and creativity.

- The course is intended to whet your appetite for research at the nexus of computation and (natural) cognition

- Cognitive models are computational models (i.e., algorithms described in a formal language) that are intended to explain basic cognitive processes

- Computational models of cognition grew out of mathematical models of natural cognitive phenomena, notably in mathematical psychology

Cognitive Modeling, Categorization, and Object Recognition

Computational and mathematical models have the advantage that

- they force researchers to be precise in their hypotheses;
- they enable quantitative predictions of behavioral variables like response time, as well as qualitative characteristics of behaviors;
- they are therefore falsifiable;
- they are generalizable (we hope) to behaviors other than the ones that they were designed for;
- they enable exploration of variant models, such as through ablation, to hypothesize about what analog changes to the natural system being modeled would imply; and
- they may be jumping off points for powerful technology, quite separated from their fit to natural cognition.
    - A good example of this is decision tree induction, which began with psychologists (Hovland, Hunt, Marin, Stone) who modeled how humans learned disjunctive concepts, then picked up by computer scientists (Quinlan, Michie), then leading to still popular methods of learning decision forests (notably in medical research)

Thomas Standish
taught me
data structures

Tarow Indow
taught me
memory and cognition

And I had an epiphany one day walking across campus that I could model human behaviors with certain data structures and algorithms – I was naïve, but it was still a singular moment

An example of mathematical modeling

Heathcote, A., Brown, S., & Mewhort, D. J. K. (2000). The power law repealed: The case for an exponential law of practice. Psychonomic Bulletin & Review, 7, 185-207.

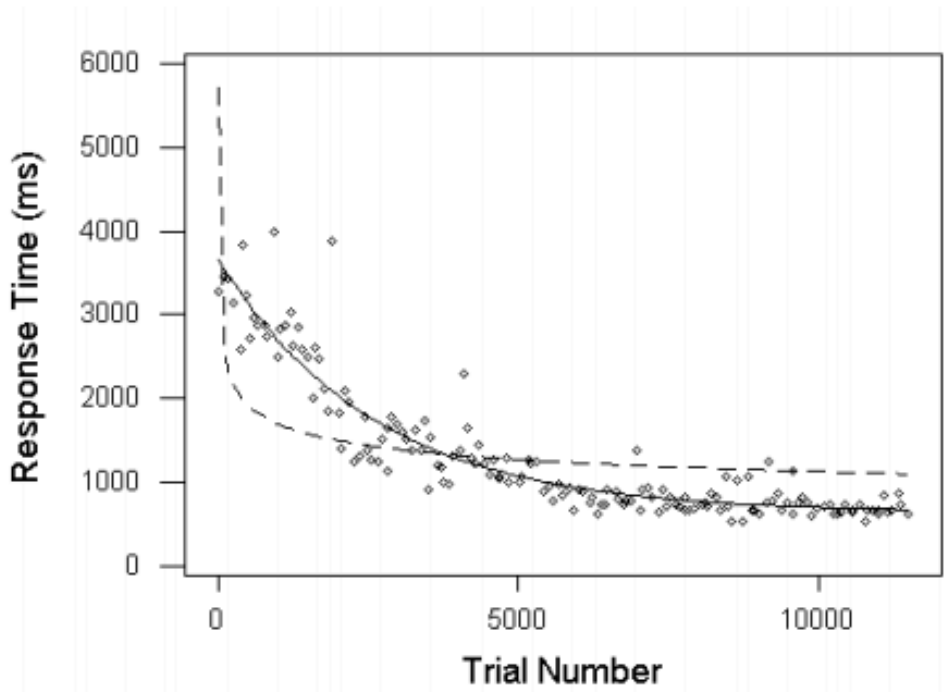

This describes the data.
It does not explain it

- - - - $RT = t^{-\beta}$ (power) ——— $RT = e^{-\alpha t}$ (exponential)

Sketch of a computational model of practice (of passwords)

Constructing a password: start with three delicious coffees

Deadman~~’s~~ Reach   Raven’s Brew   Wicked Wolf

Interleave names to obtain password

DRWeaiavcdekmnea’dnsWRBoerlaefcwh

Using and learning password

- Initially, with visual aid
- Memorization of interleaving rule and individual letters
- Chunking of subsets such as DRW, eaia, …, ‘dns, … WRBoar, …
- Continue chunking into larger substrings, until entire string automatized

Sketch of a computational model of practice (of passwords)

- Just a sketch, but could be formalized at some level of granularity by looking at propensity for certain chunks over others (e.g., WRBoer looks like a person's name)
- But how is "automatization" implemented? How are chunks stored/retrieved in/from memory
- Look for analog discussions in Palmeri and Cottrell on levels of granularity that a model explains
- Does model explain "interruption" phenomena (i.e., when interrupted in typing, must start again)?
- Does model inform password construction?

Examples of non-conscious cognition
Look for some of these in Palmeri and Cotrell, and Anderson)

Examples of non-conscious human categorization phenomena:

Basic Level effects : "Psychological studies have shown that, within hierarchical classification schemes, there appears to be a basic level preferred by human subjects. For example, in a hierarchy containing {animal, vertebrate, mammal, dog, collie}, subject behavior may indicate that "dog" lies at the basic level. Rosch, Mervis, Gray, Johnson, and Boyes-Braem (1976) used a target-recognition task to show that subjects are quicker to confirm that a test item is a member of such a basic category (e.g., dog) than they are for a superordinate (e.g., animal) or subordinate (e.g., collie) category. In a forced naming task (Jolicour, Gluck, & Kosslyn, 1984; Rosch et al., 1976), a subject is shown a picture of a particular item and asked to respond with its identity – most subjects respond with the basic level category name. Third, basic level category names typically have the shortest words (e.g., dog as apposed to animal and collie)."

Typicality effects: "Psychological experiments have repeatedly shown that human subjects do not treat concept instances equally, but regard certain members as more typical than others. For example, in a target-recognition task, subjects must determine if a test instance is a member of a target category (e.g., "Is a robin a bird?"). Several studies (Rips, Shoben, & Smith 1973: Rosch & Mervis, 1975) indicate that subjects consistently respond affirmatively more quickly to certain positive instances than to others. For example, they may more quickly affirm that a robin is a bird than they will affirm that a chicken is a bird. The relative ranking of positive test items corresponds to a typicality ranking of category members, and this conclusion is bolstered by results in a variety of other experimental tasks (Mervis & Rosch 1981; Smith & Medin, 1981)." In a forced naming task, where instances of a concept are named by a subject, subjects tend to list names roughly from more typical to less typical.

Fan effects: "Observations with frequently encountered features may be more difficult to recognize than observations with relatively unique features, given that exposure across observations is relatively constant." Intuitively, its easier to remember that you have observed a rarely encountered thing that it is to recognize a frequently encountered thing. We explained fan effects as a degenerate case of typicality effects.

Much of AI study is largely unconcerned with learning and reasoning as people do, but there is a field of AI known as Cognitive Modeling that is very concerned with modeling human thinking, at least to an approximation. If interested in more, both about cognitive modeling methodology and the categorization phenomena above, see Doug's paper on "*The Structure and Formation of Natural Categories*" (all quotes above from this paper), and "Categorization, Concept Learning, and Problem Solving: A Unifying View".

What is this?

What is this?

What is this?

What is this?

The basic level varies with "expertise" and experience

The basic level for a judge of "Best in Show" might include individual dog breeds, rather than "dog" (relates to one possible topic for a blog post on Palmeri and Cottrell)

The basic level (if we could measure) for a brand new infant might include "Mom-entity" plus animate, rather than specializing below that to dogs, cats, etc

The Ethics of AI class also participated in the exercises above intended to illustrate basic level and typicality effects

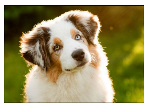

All answered "dog"
Dog is a basic level concept as identified through a convergence of behavioral data

---

Most answered "dog" but about 5 answered "Scottish Terrier" or "Scotty Dog"
A "Scotty" is a mildly atypical dog. The basic level may be "overridden" in the case of a moderately atypical instance (relates to one possible topic for a blog post on Palmeri and Cottrell)

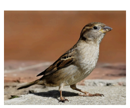

All answered "bird"
Bird is a basic level concept as identified through a convergence of behavioral data

---

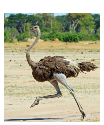

NO ONE ANSWERED "bird". All answered "Ostrich". An Ostrich" is a very atypical bird
The basic level will almost always be "overridden" in the case of a very atypical instance

What strikes you as funny (or interesting) about each of these "No Dogs Allowed" signs?

Either that atypical breeds are used to represent the basic level ‘dog’ (interesting)

or

that its only the atypical breeds that are banned (funny)

More fun with (a)typicality and meaning

???

Study these images for a few seconds.
You will be asked to say whether you have seen an image or not







Have you seen this image?

Yes! Its easier to recognize that an entity of a infrequent entity was observed or not (fan effects)

Have you seen this image?

No? Its harder to recognize whether an entity of a frequent class was observed or not (fan effects)

Have you seen this image?

Even if you answered all these questions, response times would likely show that you were faster to (dis)confirm the "chicken questions" than the hummingbird questions

Complete aside:

Fan effects seem intuitive when viewed in an evolutionary context of survival
similar looking and frequent things (your villagers)
are glossed over as you walk about, but anomalies (possible threats) stand out

But fan effects may seem less than optimal from a “computational” standpoint – shouldn’t it be easier to recognize those things that are normative and more common than it is to recognize those things that are rare?

Taking both the evolutionary perspective and the computational standpoint into account, Fan effects may drive category formation

Modeling basic level, typicality, and fan effects

Category learning over data through unsupervised learning

See references on slide 8, as well as "Knowledge Acquisition Via Incremental Conceptual Clustering" (Fisher, 1987)

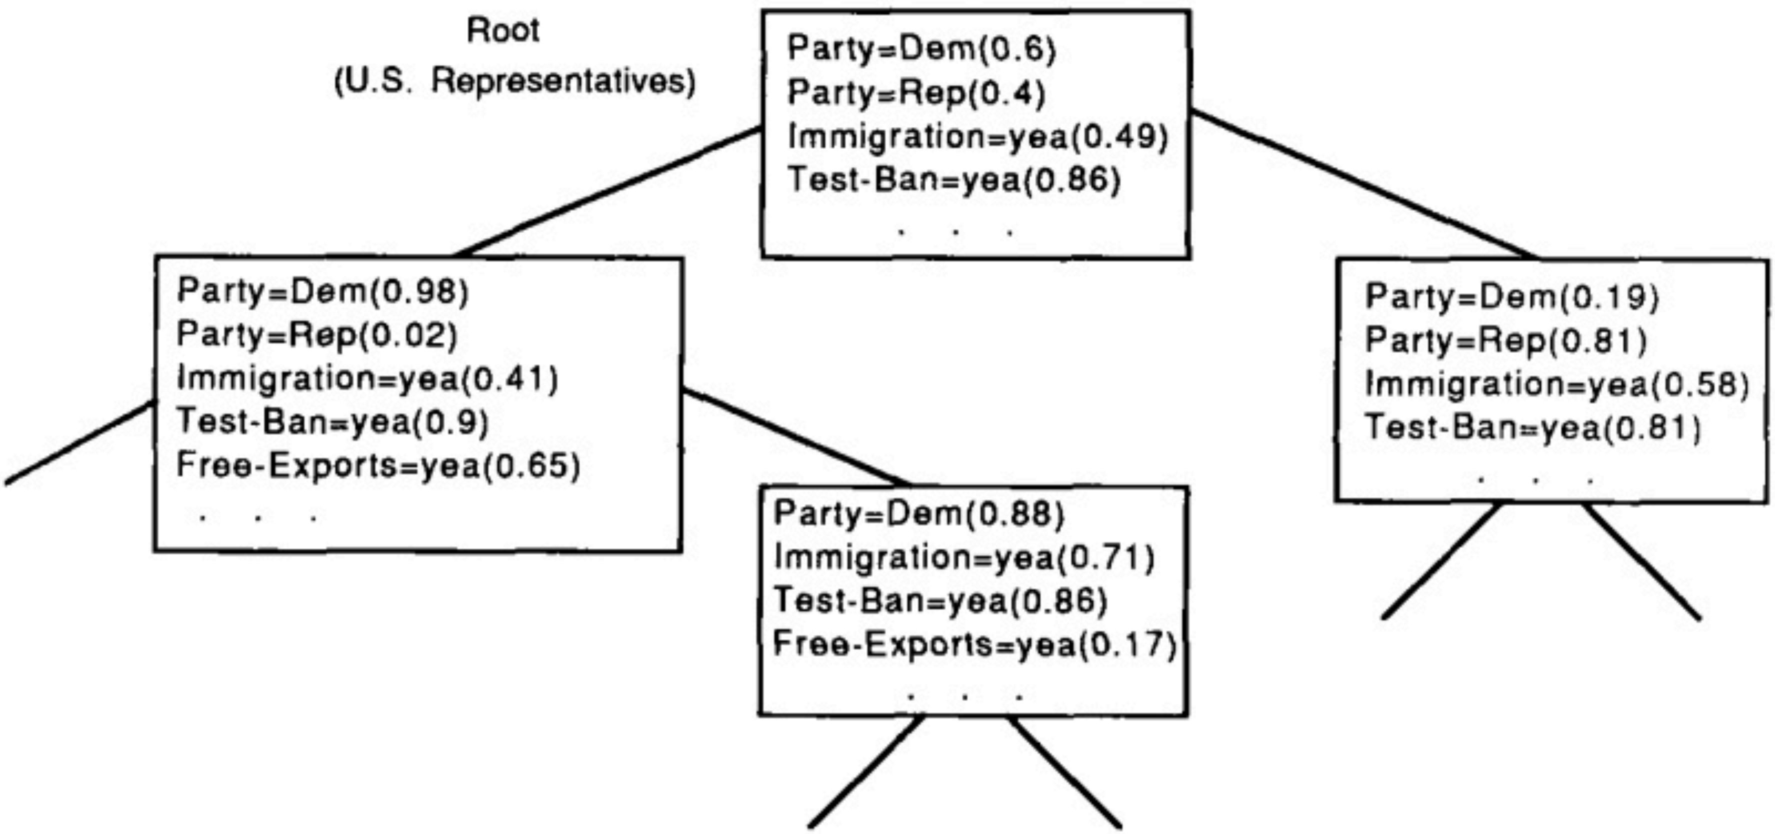

Fig. 2. A sample probabilistic concept tree over congressional voting records.

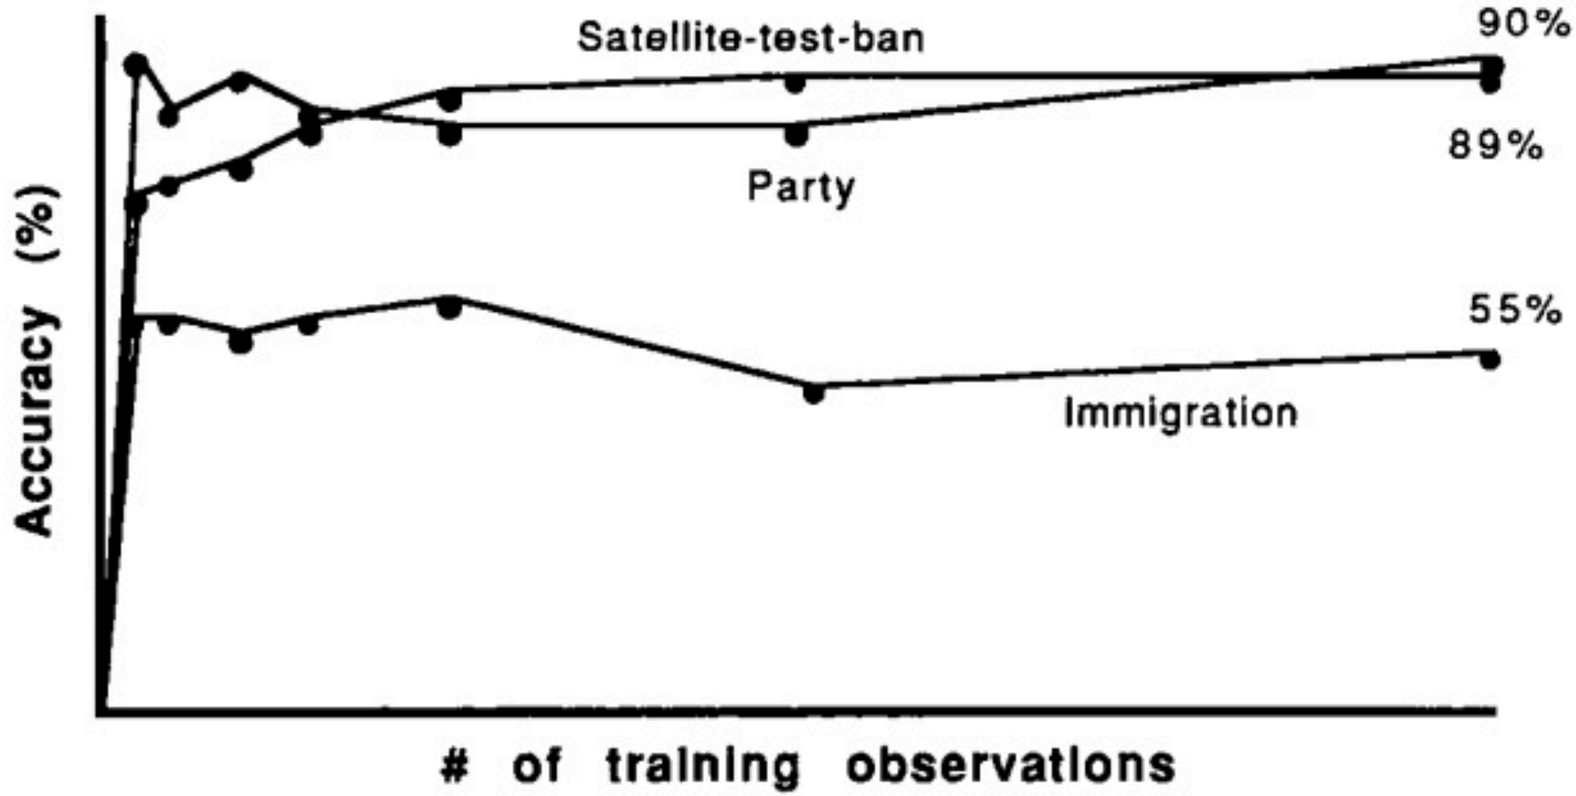

Fig. 3. Learning curves for three attributes in the congressional domain.

Category Match

$$P(C_k)\sum_j [P(V_j|C_k)^2 - P(V_j)^2]$$

- C_k is category k in a set of categories (e.g., a level in a categorization tree)
- V_j is an observation's value along the jth attribute
- The category match score is highly, positively correlated with behavior variables like response time across a large number of studies

Predicting basic level effects

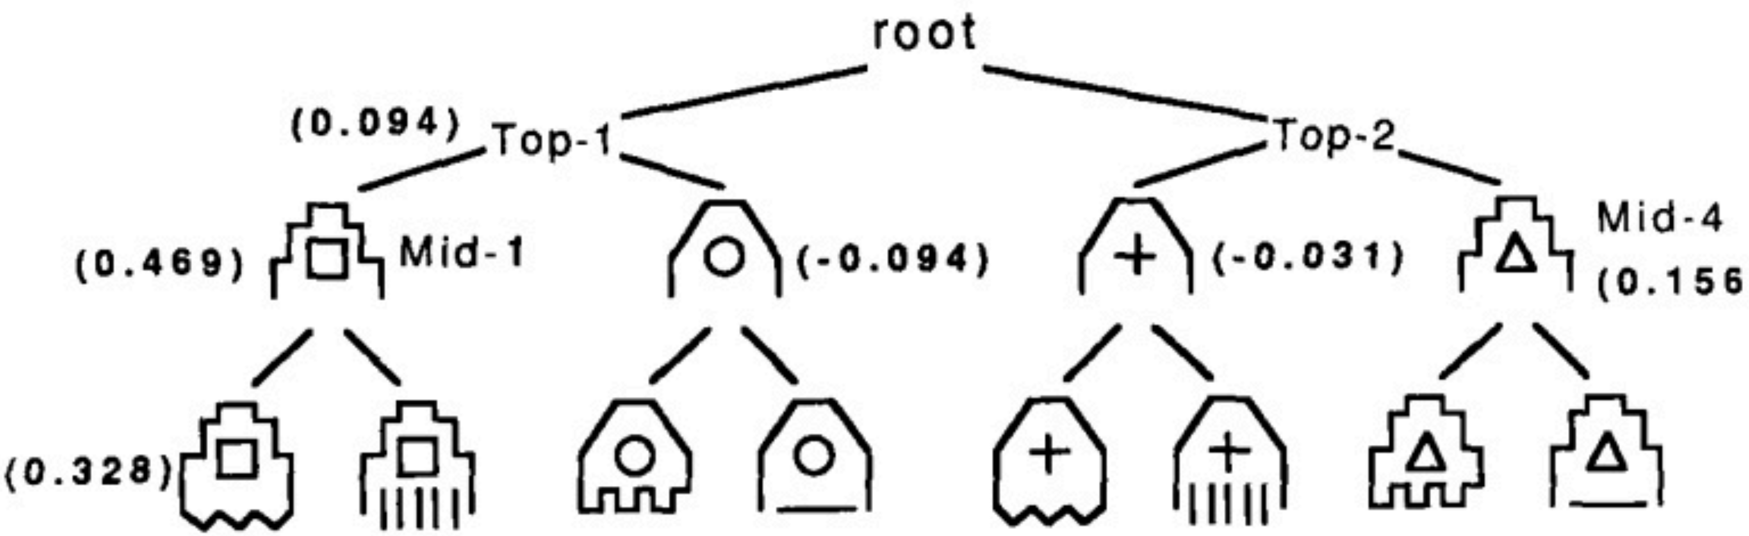

Fig. 4. Approximation of a tree from Hoffman and Ziessler basic-level studies.

Predicting basic level effects

TABLE II

ENCODED HOFFMAN AND ZIESSLER (1983) TREE

Classification tree			Attribute values		
Superordinate	Basic	Subordinate	Outer	Inside	Bottom
Top-1	Middle-1	Leaf-1	0	0	0
		Leaf-2	0	0	1
	Middle-2	Leaf-3	1	2	2
		Leaf-4	1	2	3
Top-2	Middle-3	Leaf-5	1	3	0
		Leaf-6	1	3	1
	Middle-4	Leaf-7	0	4	2
		Leaf-8	0	4	3

Note that the "perception" problem is grossly oversimplified (recall this as you read Palmeri and Cottrell)

Predicting basic level effects

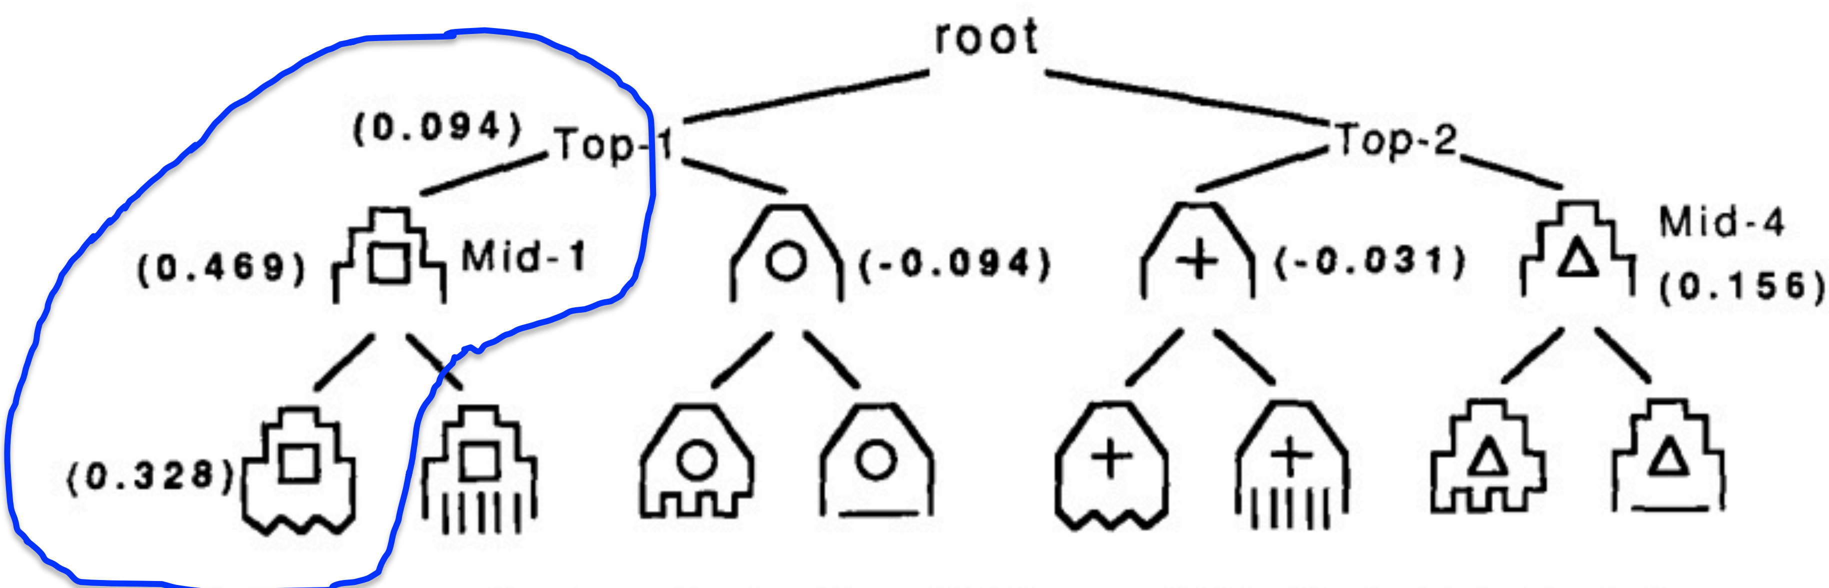

Fig. 4. Approximation of a tree from Hoffman and Ziessler basic-level studies.

TABLE III

OUR ENCODING OF THE MURPHY AND SMITH (1982) TREE

Classification tree			Attribute values			
Superordinate	Basic	Subordinate	Handle	Shaft	Head	Size
Top-1	Middle-1	Sub-1	2	2	0	0,1
		Sub-2	2	2	1	0,1
	Middle-2	Sub-3	0	3	3	0,1
		Sub-4	1	3	3	0,1
Top-2	Middle-3	Sub-5	3	4	4	0,1
		Sub-6	3	4	5	0,1
	Middle-4	Sub-7	4	0	6	0,1
		Sub-8	4	1	6	0,1

TABLE IV

AVERAGE RESPONSE TIMES AND MATCH RANKINGS

	True cases		False cases	
	Response time[a] (msec)	Category match	Response time (msec)	Category match
Superordinate	879	0.21	882	−0.070
Basic	678	0.53	714	−0.035
Subordinate	723	0.36	691	−0.018

[a]Response times from Murphy and Smith (1982).

Predicting basic level effects

Do linear regression between category match (x axis) and (actual human) response time (y axis) to obtain model-predicted response times

Table 1: Human and predicted response times for the Murphy and Smith (1982) data. Adapted from Fisher and Langley (1990).

	True Cases		False Cases	
	Response time (msec)	Predicted time	Response time (msec)	Predicted time
Superordinate	879	869	882	879
Basic	678	646	714	738
Subordinate	723	764	691	669

Predicting typicality effects

	Letter String	Intra-Category Overlap	*Typicality*		Letter String	Inter-Category Overlap	*Typicality*
	JXPHM	low	*low*		HPNWD	low	*high*
	QBLFS	"	"		HPC6B	"	"
A	XPHMQ	medium	*medium*	A	HPNSJ	medium	*medium*
	MQBLF	"	"		4KC6D	"	"
	PHMQB	high	*high*		GKNTJ	high	*low*
	HMQBL	"	"		4KCTG	"	"
	CTRVG				8SJKT		
	TRVGZ				8SJ3G		
B	RVGZK			B	9UJCG		
	VGZKD				4UZC9		
	GZKDW				4UZRT		
	ZKDWN				MSZR5		

Fig. 5. Nonsense strings used to test typicality differences.

Data from Rosch and Mervis, 1975

Predicting typicality effects

Table 2: Human and predicted response times for Rosch and Mervis (1975) data. Adapted from Fisher and Langley (1990).

	Response Time (msec)	Predicted Time
Intraoverlap		
high	560	535
med	617	615
low	692	713
Interoverlap		
low	909	968
med	986	995
high	1125	1062

Rosch and Mervis hypothesized that more typical members of category A were those that

- shared lots of characteristics with other category A members, and that
- shared fewer characteristics with contrast category (e.g., category B) members

Their experiments controlled for intra- and inter- category similarity. Our model fit their data.

Predicting fan effects

Train on lots of two-attribute data like

The doctor is in the park

The teacher is in the bank

Lets say that doctor is seen in only one sentence and park in only one sentence.
Lets say that teacher is seen in lots of different sentences, and so is bank

Table 3: Human (msec) and predicted (in parentheses) response times for Anderson's (1974) fan effect data. Adapted from Fisher and Langley (1990).

		True cases			False cases		
		subject overlap			subject overlap		
		1	2	3	1	2	3
location overlap	1	1111 (1120)	1174 (1157)	1222 (1184)	1197 (1168)	1221 (1240)	1264 (1306)
	2	1167 (1157)	1198 (1195)	1222 (1259)	1250 (1240)	1356 (1312)	1291 (1379)
	3	1153 (1184)	1233 (1259)	1357 (1321)	1262 (1306)	1471 (1379)	1465 (1444)

Predicting fan effects

- Silber and Fisher (and Fisher and Langley, and Fisher and Yoo) suggested that fan effects result as a special/degenerate case of typicality effects, where single-object categories are not differentiated by intra-category similarity (only one object per category), so all differentiation results from inter-category similarity

Also Silber, J., & Fisher, D. (1989). "A Model of Natural Category Structure and its Behavioral Implications," Proceedings of the Eleventh Annual Conference of the Cognitive Science Society, Ann Arbor, MI: Lawrence Erlbaum, 884–891.

Forum Post on Palmeri and Cottrell

A good strategy for your forum posts for this article (and others) is to identify a broad theme or two in the article, and point to examples (with page numbers) of the theme(s). Some broad questions that might be relevant to this article for a Sunday night post include the following.

- How is object processing related to and different from category processing?
- At what depth does the model explain phenomena (e.g., is a fine-grained mechanistic explanation offered, or is a variable like response time simply assumed to be proportional to some model variable like connection strength)?
- Can you think of a way to unify (a) the phenomenon that the entry level into a categorization hierarchy is deeper (i.e., subordinate) for experts than for novices, with (b) the phenomenon that the entry level for atypical basic-level category members (e.g., an ostrich) is deeper (i.e., subordinate), even for many novices?
- How do purely computational models that were developed with no interest in natural cognition inform psychological models, and vice versa?
- How encompassing of various cognitive phenomena should a model be?

There are other strategies that have merit too. For example, you could start with a referenced model and drill down on it, but in most weeks this strategy might be more appropriate to a (Tuesday night) post on a paper that you identify.

When citing references in the article just use page numbers. When citing other material, either readings or lectures or other forum posts, give a more complete pointer.