Comp. by: RMervin Date:27/2/08 Time:15:34:44 Stage:First Proof File Path://
pchns1301/WOMAT/Production/PRODENV/0000000001/0000007779/
0000000016/0000772317.3D Proof by: QC by: ProjectAcronym:bs:PLM
Volume:00005

THE CONCRETE SUBSTRATES OF ABSTRACT RULE USE

Bradley C. Love, Marc Tomlinson, and Todd M. Gureckis

Abstract

We live in a world consisting of concrete experiences, yet we appear to form abstractions that transcend the details of our experiences. In this contribution, we argue that the abstract nature of our thought is overstated and that our representations are inherently bound to the examples we experience during learning. We present three lines of related research to support this general point. The first line of research suggests that there are no separate learning systems for acquiring mental rules and storing exceptions to these rules. Instead, both items types share a common representational substrate that is grounded in experienced training examples. The second line of research suggests that representations of abstract concepts, such as *same* and *different* that can range over an unbounded set of stimulus properties, are rooted in experienced examples coupled with analogical processes. Finally, we consider how people perform in dynamic decision tasks in which short- and long-term rewards are in opposition. Rather than invoking explicit reasoning processes and planning, people's performance is best explained by reinforcement learning procedures that update estimates of action values in a reactive, trial-by-trial fashion. All three lines of research implicate mechanisms of thought that are capable of broad generalization, yet inherently local in terms of the procedures

DOI: 10.1016/S0079-7421(08)00005-4

Comp. by: RMervin Date:27/2/08 Time:15:34:44 Stage:First Proof File Path://
pchns1301/WOMAT/Production/PRODENV/0000000001/0000007779/
0000000016/0000772317.3D Proof by: QC by: ProjectAcronym:bs:PLM
Volume:00005

used for updating mental representations and planning future actions. We end by considering the benefits of designing systems that operate according to these principles.

I. Introduction

We live in a world consisting of concrete experiences. Moment to moment, we learn and make decisions based on these experiences. Given this characterization of our environment, it is unclear what would constitute a true abstraction, let alone how we could acquire and retrieve such knowledge based on the sensory cues provided by the environment. After all, we are situated in the same world as pigeons, grasshoppers, and newborn humans, not in a platonic realm of universals that exist independently of the particulars of our concrete experiences.

This view of our environment has strong implications for theories of learning and representation. It predicts that our knowledge structures are updated in a moment-to-moment fashion, instead of being calculated as global statistics or parameters. Observed ordering effects in learning support this view (Clapper & Bower, 1994; Medin & Bettger, 1994). Our view also predicts that it is not possible to acquire true abstractions. Even when people seem to have acquired and successfully applied an abstraction, we predict that the concrete basis of this behavior should become apparent on closer inspection. Following this view, findings in category learning demonstrate that experienced examples influence classification decisions in rule-governed tasks. For example, similarity effects exert themselves even when people apply explicit rules (Allen & Brooks, 1991) and exemplar frequency affects performance in seemingly rule-based tasks (Nosofsky, 1991; Rouder & Ratcliff, 2006).

In light of these findings, one reasonable position is that there is both an abstract and concrete basis to cognition. Dual processing theories are prevalent in cognitive science (Ashby & Maddox, 2005; Sloman, 1996). Sloman (1996) distinguishes between a fast associative system that operates according to similarity to past experiences and an analytical system that is slower and more deductive in nature. Reasoning fallacies, such as the conjunctive fallacy (Tversky & Kahneman, 1983) and susceptibility to framing effects (Tversky & Kahneman, 1981), are attributable to associative processes, whereas the ability to appreciate and overcome these pitfalls is attributable to the analytical system.

One fundamental challenge for dual accounts is establishing the necessity of the abstract system. Adapting to the local statistics of the environment, calculating similarities, and forming analogies to past examples offer substantial reasoning abilities without invoking rule systems (Forbus, Gentner, &

Comp. by: RMervin Date:27/2/08 Time:15:34:45 Stage:First Proof File Path://
pchns1301/WOMAT/Production/PRODENV/0000000001/0000007779/
0000000016/0000772317.3D Proof by: QC by: ProjectAcronym:bs:PLM
Volume:00005

Law, 1994; Ross & Kennedy, 1990). It is not clear that non-rule-based systems are inferior in terms of the range of problems that can be successfully addressed. Accounting for human-level competency might not require positing rule-based systems. For example, language is often viewed as a uniquely human competence, but even in this domain arguments regarding the necessity of rule systems are common. For example, the language processing literature offers ample examples of debates between proponents of single and dual accounts for domains like the past tense (Pinker & Prince, 1988; Rumelhart & McClelland, 1987). Challenges to these dual route accounts bring into question the necessity of positing abstract rule-based processes and representations.

In this chapter, we argue that the prevalence of abstract representations in cognition is overestimated. We do not argue that people cannot create or follow sequences of rules. For example, baking a cake according to a recipe involves executing an algorithm that was constructed by another human. Instead, we argue that seemingly abstract rule representations are built on and are intimately linked to concrete experiences. We predict that our representations never fully transcend these concrete experiences that they are built on and, therefore, are not truly abstract in nature. Upon closer inspection, the messy underpinnings of these seemingly abstract rules are revealed.

We offer three supportive cases drawn from our own research. In each case, seemingly abstract processes and representations appear to govern performance, but on closer examination performance is instead governed by incremental learning processes operating over experienced examples. We offer model-based explanations for all three cases considered. Each model learns on a trial-by-trial basis by adjusting its memory representations based on the current example. These learning mechanisms do not follow from a rational analysis in that they update their category estimates based on local statistics from the immediate context (e.g., the discrepancy between the current memory representation of a category and the current stimulus).

In the first case study, we consider how people learn and apply rules, as well as encode exceptions to rules. We find evidence that people do not posses separate rule and exception systems, but instead rely on a single cluster system that is intimately influenced by the concrete details of the training set. Even when subjects mentally rehearse rules, much like a baker following a recipe step-by-step, recognition performance following learning reveals that the representation of rules is intimately linked to the distribution of experienced examples.

In the second case study, we extend this analysis to abstract rules in which the acquired rules are not bound by concrete stimulus aspects. For example, abstract notions like *same* and *different* do not range over specific stimulus properties, like *has blue eyes* or *has two wheels*. For example, two identical

Comp. by: RMervin Date:27/2/08 Time:15:34:45 Stage:First Proof File Path://
pchns1301/WOMAT/Production/PRODENV/0000000001/0000007779/
0000000016/0000772317.3D Proof by: QC by: ProjectAcronym:bs:PLM
Volume:00005

wheelbarrows are the same as are two identical comets, and this notion of sameness is not based on a shared property spanning wheelbarrows and comets. We find evidence that such seemingly abstract notions are rooted in experienced exemplars. Analogy to these stored exemplars enables seemingly abstract rule performance, but the concrete substrate enabling this performance can be seen by considering stimulus probes that fall in between relational categories. The results from these probes reveal a family resemblance structure to abstract concepts that is rooted in experienced exemplars.

In the final case study, we examine how people learn to reason about rewards in a dynamic decision task that puts short- and long-term rewards in opposition. People's representation of the underlying system, which can be manipulated by introducing perceptual cues, strongly constrains performance. People eventually gain an understanding of the dynamics of the system, but the trajectory of learning implicates, as in the previous two cases, a trial-by-trial learning mechanism that makes local updates to estimates of the value of particular actions. All three cases suggest that our abstract knowledge supervenes on concrete representations and processes.

II. When Rules are not Rules: Rule-Plus-Exception Category Learning

In our first case study, we consider how people acquire rules and master exceptions to these rules in category learning tasks. The results we review suggest that even when people explicitly acquire and apply a rule, closer inspection of the data reveals a strong dependency on the concrete examples experienced during training. Recent studies in human rule-plus-exception learning strongly suggest that people do not fully abstract away the details of the training set when acquiring rules. Rather than propose a dual system model to account for these results, we propose a clustering model. This model does not contain a rule route, but nevertheless explains rule-like behavior, as well as subtle, but diagnostic, deviations from rule-like behavior that will be discussed later in this section.

Proposals for category representation are diverse, ranging from exemplar based (Medin & Schaffer, 1978) to prototype based (Smith & Minda, 1998) and include proposals between these two extremes (Love, Medin, & Gureckis, 2004). Determining the best psychological model can be difficult as one model may perform well in one situation but be bested by a competing model in a different situation. One possibility is that there is not a single "true" model.

In category learning, this line of reasoning has led to the development of models containing multiple learning systems. These more complex models

Comp. by: RMervin Date:27/2/08 Time:15:34:45 Stage:First Proof File Path://
pchns1301/WOMAT/Production/PRODENV/0000000001/0000007779/
0000000016/0000772317.3D Proof by: QC by: ProjectAcronym:bs:PLM
Volume:00005

hold that category learning behavior reflects the contributions of different systems organized around discrepant principles that utilize qualitatively distinct representations. The idea that multiple learning systems support category learning behavior enjoys widespread support in the cognitive neuroscience of category learning (see Ashby & O'Brien, 2005, for a review and Nosofsky & Zaki, 1998, for a dissenting opinion). Multiple systems models typically include a rule system.

Multiple system models of category learning detail the relative contributions of the component learning systems. The relative contributions can depend on the circumstances. For example, ATRIUM (Erickson & Kruschke, 1998) contains a rule and exemplar learning system. Which system is operable is determined by a gating system, allowing different classification procedures to be applied to different parts of the stimulus space. For example, familiar items could be classified by the exemplar system, whereas rules could be applied to unfamiliar items. The power to apply qualitatively different procedures to different stimuli is the hallmark of multiple systems models.

Proposing multiple systems begs the questions of how many systems are present and how do they interact. Are there two, three, or thirty-four systems? Do some systems combine outputs whereas others shunt each other? These questions are not trivial to answer. For example, a two system model may suffice for one data set, but a new manipulation could provide evidence for a third system. As systems propagate, the complexity of the overall system dramatically increases. Building in this degree of complexity complicates model evaluation.

Instead of proposing a complex model of category learning containing multiple systems, we advocate a complex systems approach to category learning modeling in which multiple learning systems emerge from a flexible and adaptive clustering mechanism's interactions with the environment. We evaluate the hypothesis that a relatively small set of learning principles can effectively "grow" knowledge structures that satisfy the needs that multiple systems models are intended to address. Our proposal does not require positing a separate rule route and instead maintains some details about the experienced examples. The forces of trial-by-trial learning grow knowledge structures that approximate abstract rules, but are in fact more consistent with human performance.

A. Past Work and Current Challenges

Previous work with the SUSTAIN model, which is the precursor to the model that we introduce here, has partially delivered on the promise of flexibly building needed knowledge structures. SUSTAIN is a clustering model that starts simple and recruits clusters in response to surprising events,

Comp. by: RMervin Date:27/2/08 Time:15:34:45 Stage:First Proof File Path://
pchns1301/WOMAT/Production/PRODENV/0000000001/0000007779/
0000000016/0000772317.3D Proof by: QC by: ProjectAcronym:bs:PLM
Volume:00005

such as encountering an unfamiliar stimulus in unsupervised learning or making an error in supervised learning (cf. Carpenter & Grossberg, 2003). Surprising events are indicative that the existing clusters do not satisfy the learner's current goals and that the model should grow new knowledge structures (i.e., clusters). These clusters are modified by learning rules that adjust their position to center them amidst their members. Dimension-wide attention is also adjusted to accentuate stimulus properties that are most predictive across clusters.

Although simple, these growth dynamics allow SUSTAIN to address a wide range of human learning data across various paradigms including unsupervised, inference, and classification learning (Love et al., 2004). Depending on the circumstances of the learning situation (i.e., depending on what the task stresses and target categories), SUSTAIN can evolve clusters that resemble prototypes, exemplars, or rules (Love, 2005). Careful behavioral experiments support the conclusion that SUSTAIN is not merely mimicking these other models, but that human learners' and SUSTAIN's representations are in accord (Sakamoto & Love, 2004). Importantly, SUSTAIN's clusters maintain a great deal of information about the experienced exemplars and their distributions of feature values. Although some abstraction occurs when items cluster together, the degree of information preservation appears to be in accord with human learners and is to some extent dictated by the nature of the task and domain. In summary, SUSTAIN accounts for both classical studies of category learning and the more contemporary work that suggests that conceptual organization is determined by the interplay of information structures in the environment and task pressures or goals (Markman & Ross, 2003).

Despite these successes, considerable challenges remain. Two basic challenges are (1) to formalize the notion of a goal or task pressure and specify how such factors direct learning and (2) to endow learning models with the flexibility to develop representations that approach the range, and richness of the representations that human learners build when learning from examples.

Although SUSTAIN made strides in capturing the influence of goals, its notion of goal is underdeveloped. In particular, SUSTAIN is sensitive in an all-or-none fashion to whether a particular stimulus dimension is queried (e.g., an unknown perceptual feature or category label). Ideally, the notion of goal would be more encompassing and continuous to capture all possible cases from pure classification learning in which the only goal is to predict category membership to pure unsupervised learning in which the goal is to predict every feature (i.e., to capture the correlational structure of the environment in an unbiased fashion). Importantly, the formal notion of goal should directly affect the recruitment and modification of clusters in a principled way. Learning rules should update clusters to reflect the goal

Comp. by: RMervin Date:27/2/08 Time:15:34:45 Stage:First Proof File Path://
pchns1301/WOMAT/Production/PRODENV/0000000001/0000007779/
0000000016/0000772317.3D Proof by: QC by: ProjectAcronym:bs:PLM
Volume:00005

measure, and clusters should be recruited in light of how well the current clusters satisfy the current goal measure.

In regards to the second basic challenge, current models like SUSTAIN are too limited in terms of the range of knowledge structures they can construct. For instance, SUSTAIN's attentional mechanism accentuates certain features that are predictive in the current task, but is constrained such that every cluster is focused on the same set of properties. In contrast, people stress different properties in different domains. For example, when shopping for clothing, color is important, but when shopping for a computer the type of processor is important (a feature not even relevant to clothing). To evolve these kinds of knowledge structures and to apply different "procedures" to different parts of the stimulus space as multiple systems models do, each cluster needs to be able to accentuate the features that satisfy the learning goals for the stimulus aspects it represents. A related challenge is storing information and capturing regularities at different scales ranging from very specific (e.g., Jim's dog Fido) to very broad (e.g, living things). To address these issues, clusters need to fine tune their level of specificity to satisfy the goal measure. As in the case of adjusting attention at the individual cluster level, adjusting specificity at the individual cluster level allows for different criteria to be applied to different parts of the stimulus space, as in multiple systems models.

The model that is introduced in the next section, CLUSTer Error Reduction (CLUSTER), meets these stated challenges. CLUSTER incorporates a formal goal measure that directs cluster development. CLUSTER has sufficient flexibility to evolve conceptual structures (i.e., clusters) that reflect key aspects of human knowledge representation. After introducing the model, a supportive simulation will be discussed. The simulation illustrates how CLUSTER can evolve cluster organizations that serve the functions of multiple systems. Importantly, CLUSTER does not posit a separate rule system and maintains key aspects of the training examples that allow it to capture human performance. Finally, we will consider how CLUSTER is consistent with cognitive neuroscience findings advocating multiple memory systems, and briefly discuss work that is being done to further develop and verify the model. The mathematical details of the model are presented in Love and Jones (2006).

B. Overview of CLUSTER

CLUSTER is an autoassociative model of human category learning in which the "hidden" layer consists of clusters (see Fig. 1). A cluster is a bundle of related features. A presented stimulus activates the existing clusters, which pass their activation to the output layer via connection weights. Like other

Comp. by: RMervin Date:27/2/08 Time:15:34:45 Stage:First Proof File Path://
pchns1301/WOMAT/Production/PRODENV/0000000001/0000007779/
0000000016/0000772317.3D Proof by: QC by: ProjectAcronym:bs:PLM
Volume:00005

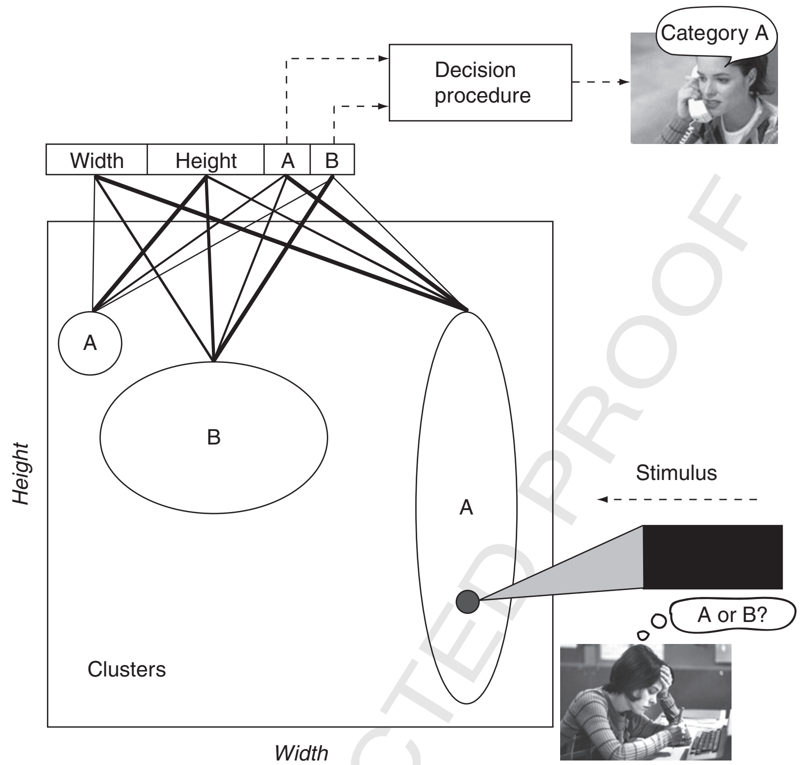

Fig. 1. CLUSTER is an autoassociative learning model in which the hidden layer consists of clusters that adjust their position, attention, and association weights to minimize an error term that reflects the learner's goals. In the illustrated example, three clusters have been recruited and the model is being asked to infer the category label. The different shape of each cluster indicates its unique attention profile. The thickness of the links from each cluster to the output layer reflect the strength of the association weights. The input layer, which mirrors the output layer, is not shown. Instead, each cluster's position is shown geometrically for the width and height dimensions. Category dimensions are indicated by the labels A and B.

autoassociative models (e.g., Kurtz, 2007), CLUSTER attempts to replicate the input layer at the output layer and in the process develops internal representations that seize on key regularities.

CLUSTER differs from other autoassociative models in a critical way. The error term CLUSTER minimizes and does not uniformly weight reconstruction error equally across features. Instead, each feature's error is weighted according to its goal relevance. For example, pure classification learning places all the error term weighting on the category label features and error Au1 associated with reconstructed perceptual features is disregarded (as in most category learning models). At the other extreme, pure unsupervised learning weights the reconstruction error uniformly across features (as in most autoassociative models). CLUSTER can capture every conceivable case in between these extremes, which is critical as the extremes are likely caricatures that do not correspond to human learning (e.g., people incidentally learn about feature correlations in classification learning and place more importance on predicting certain features in unsupervised learning).

The error term (with goal weights on each feature) reflects the discrepancy between what CLUSTER predicts and what is observed. Importantly, discrepancies for features that are goal-relevant are weighted more highly and, therefore, have more influence on learning. To satisfy this goal, clusters adjust their position, attention, and weights to minimize the error term through gradient descent learning. Thus, depending on the goal weights, different cluster organizations will emerge. Unlike most models, each cluster can adjust its own attention to minimize error and attention does not sum to a fixed number (i.e., clusters can vary in overall attention or specificity). These changes allow additional flexibility for clusters to emphasize different features and to vary in specificity (e.g., a specific dog vs. dogs in general). Although in Fig. 1 grouping of features implies dimensional structure, CLUSTER departs from the majority of models that utilize selective attention mechanisms (e.g., Nosofsky, 1986) in that it does not assume a dimensional structure. Not assuming dimensional structure allows for additional flexibility (e.g., the presence or absence of red can be critical to a cluster, whereas the presence or absence of blue can be somewhat irrelevant).[1]

CLUSTER begins with one cluster centered on the first training example, and recruits additional clusters when the existing clusters are not supportive of the current goal. Each newly recruited cluster is centered on the current stimulus. Like CLUSTER's other operations, the algorithm for cluster recruitment is consistent across all induction tasks (there are no special cases). Despite its consistency across situations, CLUSTER retains the flexibility to build representations that capture many of the competencies of human learners without proposing distinct learning systems. CLUSTER is highly principled (all of its operations are tied to the goal-weighted error term), but minimal structure is built into the model. Instead, CLUSTER evolves the knowledge structures needed to solve the current task.

[1] Interestingly, in cases in which contrasts are consistent (e.g., when red is present, blue is absent, and vice versa), CLUSTER attends equally to the contrasting features within each cluster. Thus, CLUSTER may prove to provide some insight into how dimensional structure arises.

Comp. by: RMervin Date:27/2/08 Time:15:34:46 Stage:First Proof File Path://
pchns1301/WOMAT/Production/PRODENV/0000000001/0000007779/
0000000016/0000772317.3D Proof by: QC by: ProjectAcronym:bs:PLM
Volume:00005

C. Illustrative Simulation

Findings from previous studies exploring rule-plus-exception learning have been problematic for exemplar models and have been used to support multiple systems models, like the RULEX model of category learning (Nosofsky, Palmeri, & McKinley, 1994). RULEX proposes that rule-following items are captured by a rule system whereas exception items reside in an exemplar store. Interestingly, both SUSTAIN and CLUSTER can account for such findings (e.g., Palmeri & Nosofsky, 1995). Such fits alone cast doubt on the necessity of rule systems. Here, we go farther and present more challenging findings that establish that human rule-plus-exception learning is inconsistent with rule-plus-exception models, but does follow from CLUSTER. We demonstrate that CLUSTER can accommodate such findings by applying different procedures to different parts of the stimulus space and in fact provides an account superior to RULEX's.

To test between this dual route account (i.e., rules and exceptions) and a clustering account, Sakamoto and Love (2004) revisited the rule-plus-exception design with the twist that one rule was twice as frequent as the other. Subjects sequentially classified stimulus items into categories A and B and received corrective feedback. Each category was defined by a rule (e.g., if large, then A; if small, then B). Additionally, each category contained an exception (e.g., a small member of A; a large member of B). Table I provides the design details of Sakamoto and Love's variation in which one experienced category had twice as many rule-following items as the contrasting category. Because subjects reason from stimulus dimensions to categories in classification learning, the exception in the smaller category violates the more frequent rule in Table I (i.e., if value 1 on the first stimulus dimension, then A). Following learning, recognition memory was assessed. In contrast to RULEX's predictions (across all explored parameter values), the exception violating the more frequent rule was better remembered than the exception violating the less frequent rule (see Fig. 2). This result is surprising given that the ''rules'' during learning were cued on the screen and that subjects reported internally rehearsing these rules. The fact that the item violating the more salient or frequent rule was remembered better than the other exception is evidence that the mental substrate of rule-plus-exception learning does not consist of separate rule and exception routes.

CLUSTER was applied to the data to illustrate its ability to ''evolve'' multiple systems. Each stimulus dimension shown in Table I and category membership were represented by 2 features for a total of 12 features. In contrast to RULEX (which requires eight parameters to CLUSTER's seven for the simulation), multiple sets of parameters replicated the basic pattern of results, indicating that these findings follow from CLUSTER's

TABLE I

THE ABSTRACT STIMULUS STRUCTURE FOR SAKAMOTO AND LOVE'S (2004) EXPERIMENT 1 IS SHOWN

Learning items	Dimension values	Novel items	Dimension values
Category A			
→ A1	21112	N1	11221
A2	12122	N2	12112
A3	11211	N3	12221
A4	12211	N4	12212
A5	11122	N5	12222
A6	12111	N6	21221
A7	11222	N7	22112
A8	11212	N8	22221
A9	12121	N9	22212
Category B		N10	22222
→ B1	11121		
B2	22122		
B3	21211		
B4	22211		
B5	21122		

Items A1 and B1 (indicated by the arrows) violate the imperfect rule of the first stimulus dimension. subjects completed 10 training blocks where each block consisted of each item below presented in a random order. Following learning, items A1–5 and B1–5 were paired with all combinations of novel foils that matched on the first dimension in forced choice recognition. The actual stimuli were simple geometric figures. for example, for some subjects the first dimension was size with a 1 indicating a small figure and 2 indicating a large figure.

basic operation and that additional work is necessary to establish default parameters for CLUSTER. These and other model evaluation issues, such as consideration of nested models within CLUSTER's formalism, are topics currently being intensely pursued, but are set aside here in favor of demonstrating CLUSTER's promise to evolve multiple learning systems. Here, we focus on the results and rationale, rather than the details of the simulations. Readers interested in the details of the simulations are directed to Love and Jones (2006).

Using these parameters, CLUSTER was simulated 10,000 times, adopting methods paralleling the human study (e.g., 10 blocks of training) and the results were averaged. CLUSTER correctly predicts that the exceptions are recognized best with the exception from the small category recognized better (.88 vs. .80) than the exception from the large category. Rule items from the small and large categories are recognized at similar rates (.58 vs. .59, respectively). In summary, CLUSTER's predictions matched the basic pattern observed in human subjects (see Fig. 2).

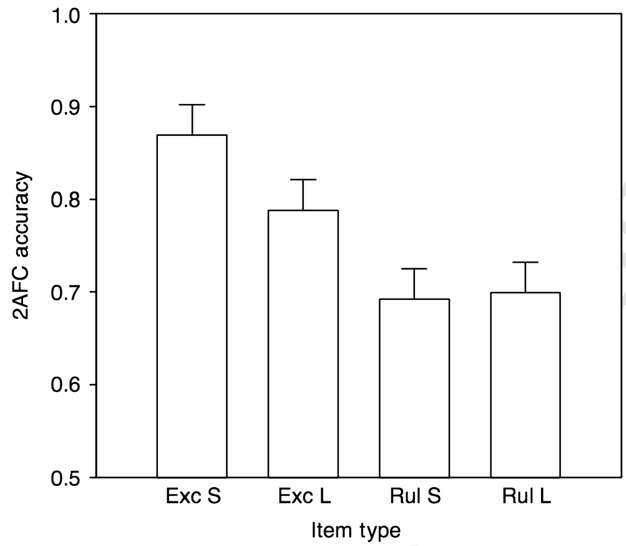

Fig. 2. Mean accuracies in the recognition phase of Sakamoto and Love's (2004) Experiment 1 are shown along with 95% within-subjects confidence intervals (see Loftus & Masson, 1994). Exc S are the exception of the small category, Exc L are the exception of the large category, Rul S are the rule-following items of the small category, and Rul L are the rule-following items of the large category. Au16

CLUSTER recruited 11.4 clusters on average (the median was 11) to represent the 14 training items. The number of clusters recruited followed a normal distribution with solutions ranging from 4 to 23 clusters with a standard deviation of 2.3. Every solution examined involved devoting at least one cluster to encoding each exception with many simulations devoting multiple clusters to each exception. Because CLUSTER is a distributed model and its predictions for an item depend on the responses of all clusters, an analysis of the four item types was conducted that factored in all clusters.

One explanation for CLUSTER's ability to accommodate the results is that it increased attention for clusters playing prominent roles in coding the exceptions, particularly for nonrule stimulus features. Encoding these items at a different specificity than rule-following items would help reduce confusions between these items and rule-following items, resulting in both reduced error during training and in enhanced recognition for exceptions. The pressure to enhance attention should be greatest for the exception violating the more frequent rule as every rule-following item from the contrasting category provides an impetus to enhance attention. This process is illustrated and explained in greater detail in Fig. 3.

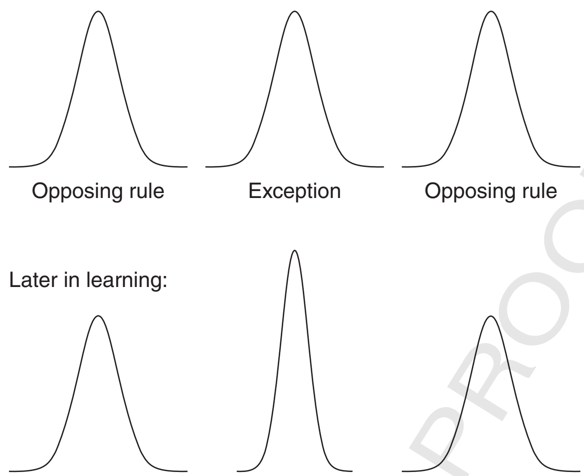

Fig. 3. Better recognition of exceptions in CLUSTER is driven by two factors. The first factor (shown in the top panel) is that exception items elicit surprise and are subsequently encoded by their own recruited cluster. In contrast, clusters responding to rule-following items tend to encode for multiple items, which leads to the loss of some item individuating information that in turn reduces later item recognition. The second factor is that the tuning or specificity of exception-encoding clusters tends to become peaked in order to reduce confusion with rule-following items from the opposing category. This process is illustrated in the bottom panel. Every time a cluster encoding a rule-following item from the contrasting category is activated, the cluster encoding the exception from the opposing category is inappropriately activated as this cluster tends to be similar to rule-following items from the contrasting category but votes for the opposite response. This inappropriate activation of the exception cluster leads to error and the tuning of the exception cluster increases to reduce this error by limited generalization or coactivation. Increasing the exception cluster's tuning boosts recognition of that exception item. This process is more pronounced for exception clusters that are more confusable with rule-following items from the contrasting category because of either the numerosity of these rule-following items (e.g., Sakamoto & Love, 2004) or because of their high similarity to the exception cluster (e.g., Sakamoto & Love, 2006). Both of these findings are inconsistent with rule-based accounts of rule-plus-exception learning.

To evaluate this explanation, following training, study items were presented to CLUSTER and a weighted sum of attention to nonrule features was calculated by multiplying each cluster's sum of attention for nonrule features by its activation. Then, these products were summed and normalized by dividing by the sum of all cluster activations. The results for items of the same type were averaged. The mean results for the four item types (averaged over 10,000 simulations) are 1.36, 1.32, 1.28, and 1.29 for exceptions from the

Comp. by: RMervin Date:27/2/08 Time:15:34:46 Stage:First Proof File Path://
pchns1301/WOMAT/Production/PRODENV/0000000001/0000007779/
0000000016/0000772317.3D Proof by: QC by: ProjectAcronym:bs:PLM
Volume:00005

small category, exceptions from the large category, rule items from the small category, and rule items from the large category, respectively. As predicted, these sums perfectly track item recognition. Exceptions (particularly the exception violating the more frequent rule) were stored as ‘‘hot spots’’ of focused activity, whereas clusters coding for rule items were more broadly tuned and were less apt to code item specific differences. Distinct representations emerge for the item types.

CLUSTER provides a similar account of related data sets in which exception memory was manipulated by varying the similarity between exception types and contrasting rule items instead of manipulating rule token frequency (Sakamoto & Love, 2006). SUSTAIN cannot account for these data. The Sakamoto and Love (2006) studies demonstrate that enhanced recognition of exceptions violating frequent rules is not due to frequency per se. Instead, increased confusions in memory enhance exception encoding (see Fig. 3). Sakamoto and Love (2006) equated the frequency of each rule and instead varied how similar rule-following items were to the exception from the opposing category. The more similar the exception, the more the specificity of the exception increased to offset confusions with rule-following items from the contrasting category. This enhancement boosted recognition. These results provide further support that one learning system underlies rule-plus-exception learning, and that there is not a separate rule system governing performance.

D. Discussion

Human learners display flexibility in how they represent category information that outstrips the capacities of traditional single system models. In response, the field has developed multiple system models that are themselves not without problems. These multiple system models often contain a separate rule system. One such model, RULEX, could not account for the results presented here despite its complexity. Here, we pursue a theoretical approach that diverges from straightforward single system models and multiple system models—knowledge structures evolve as needed to satisfy the learner’s goals.

CLUSTER embodies this third position. CLUSTER has a formally defined notion of goal that spans induction tasks, recruits clusters when existing clusters fail to support the learner’s goals, and adjusts clusters’ positions, attention, and association weights to reduce goal mismatch. These operations are sufficient to apply different procedures to different parts of the stimulus space, as multiple systems models do.

How do we reconcile our position with impressive evidence from cognitive neuroscience that multiple systems underlie human category learning

Comp. by: RMervin Date:27/2/08 Time:15:34:46 Stage:First Proof File Path://
pchns1301/WOMAT/Production/PRODENV/0000000001/0000007779/
0000000016/0000772317.3D Proof by: QC by: ProjectAcronym:bs:PLM
Volume:00005

performance? We do not deny that multiple learning systems underlie human category learning. A nonexhaustive list of systems includes a dopaminergic procedural learning system, a working memory system engaging cortical–thalamic loops, and a PFC-hippocampal-perirhinal learning system. The last system is marked by its flexibility and is adept at creating new conjunctive representations that link features (i.e., clusters). SUSTAIN (the precursor to CLUSTER) is readily put in correspondence with this learning circuit and has successfully simulated populations with hippocampal deficits by reducing the model's ability to form new clusters (Love & Gureckis, 2007). CLUSTER likely corresponds to the hippocampal system as well. We believe that a fast learning hippocampal system is shadowing the other learning systems. For instance, the literature is replete (including Sakamoto & Love, 2004) with cases in which learners are clearly applying a rule stored in working memory, but are nevertheless storing additional information about rule-following examples. Critically, we believe our empirical studies and simulations demonstrate that positing multiple systems is not enough in itself to explain human behavior. We believe that at least one of the systems posited must have the ability to build new representations flexibly in response to a learner's goals, like CLUSTER does.

The proposal that CLUSTER relates to a learning circuit involving the hippocampus is consistent with our position that experienced examples exert an influence on performance following learning, as opposed to experienced examples being completely abstracted away by a rule representation. The hippocampal memory system is closely tied to episodic memory, which requires memory for experienced examples. Love and Gureckis's (2007) account of CLUSTER and the hippocampal system posits that every semantic memory begins as an episodic memory (surprising events are stored as separate clusters). Although recruited clusters can be activated and modified by subsequent events, in many cases clusters will retain some information about the original episode.

Of course, much work remains to be done. Efforts are underway to apply CLUSTER to all the studies to which SUSTAIN has been applied. The results so far are promising. Additionally, we are applying CLUSTER to studies exploring how people partition knowledge and appear to apply different procedures depending on context (e.g., Yang & Lewandowsky, 2004). Finally, CLUSTER has been successfully applied to Kruschke's (1993) filtration and condensation tasks that were intended to demonstrate the necessity of dimensional attention (CLUSTER has cluster and feature-specific attention). Although CLUSTER does not have a built in notion of dimensional attention, dimensional attention emerges (i.e., there is advantage for aligning all clusters along the same contrasting features) much like how what looks like multiple learning systems emerges out of the Sakamoto and

Comp. by: RMervin Date:27/2/08 Time:15:34:46 Stage:First Proof File Path://
pchns1301/WOMAT/Production/PRODENV/0000000001/0000007779/
0000000016/0000772317.3D Proof by: QC by: ProjectAcronym:bs:PLM
Volume:00005

Love (2004) simulations. While CLUSTER itself is still evolving, it appears it has the necessary constraints built in to account for human learning and no more. CLUSTER demonstrates that learning can be both intimately tied to the specifics of the training examples (in terms of presentation order, numerosity, similarity relations, etc.) and display the kinds of flexibility that human learners show. This flexibility outstrips multiple system models consisting of separate rule systems.

III. Learning Abstract Rules from Examples

The previous section suggested that people do not use mental rules to capture rule-like regularities defined over stimulus features. This section makes a bolder claim. Here, we claim that similar mechanisms rooted in concrete examples can explain our ability to acquire and apply seemingly abstract rules or concepts. Abstract rules are not defined over any particular set of stimulus properties. In the introduction, the concepts same and different were provided as examples of abstract rules. To provide another example, Chomsky (1957) described the rules of language as being abstract in nature. According to Chomsky, transformative and generative rules govern whether a sentence is grammatical. On this view, the set of grammatical sentences in a language is infinite and no specific set of features (e.g., ''The'' appears before ''boy'') can be used to classify sentences as grammatical or ungrammatical. The difficulty in imagining how such a complex and abstract concept could be acquired is one of the motivating pillars of the nativist position.

In this section, we argue that humans can learn seemingly abstract concepts by making analogy to concrete exemplars. Our view predicts that abstract concepts should display a family resemblance structure (cf. Rosch & Mervis, 1975) determined by similarity to stored exemplars. Evidence of such similarity effects would suggest a concrete basis to abstract concepts. To support our conjecture, we present an exemplar model of category learning that can learn seemingly abstract concepts by analogy to stored exemplars. To foreshadow our simulation results, the model, Building Relations through Instance Driven Gradient Error Shifting (BRIDGES), accounts for the acquisition of seemingly abstract concepts and correctly predicts that these concepts have a family resemblance structure rooted in experienced examples.

Exemplar models of category learning hold that all abstraction or generalization occurs through similarity-based activations of concrete examples. In contrast, CLUSTER formed abstractions directly in memory in the form of clusters. We will consider the relationship between CLUSTER and BRIDGES in the General Discussion. In exemplar models like BRIDGES,

Comp. by: RMervin Date:27/2/08 Time:15:34:46 Stage:First Proof File Path://
pchns1301/WOMAT/Production/PRODENV/0000000001/0000007779/
0000000016/0000772317.3D Proof by: QC by: ProjectAcronym:bs:PLM
Volume:00005

each experienced instance is stored in memory. When a new item is encountered, the similarity between the item and each exemplar in memory is calculated. The stimulus is predicted to belong to the category with the greatest sum of pairwise similarity (Medin & Schaffer, 1978). Thus, exemplar models clearly link experienced events to later generalization.

BRIDGES is derived from ALCOVE (Kruschke, 1992), a connectionist-based exemplar model of classification. ALCOVE learns to weight different aspects of the exemplars differently when calculating similarity. Through a process of trial-by-trial backpropagation of error, when computing the similarity between stimuli and exemplars, ALCOVE learns to ignore those aspects of the stimuli that are not predictive of correct classification. ALCOVE learns to generalize over the irrelevant features. When first learning about mammals, fur, or giving birth to live young might be considered important, but with enough experience possession of a mammary gland will become the sole predictor of category membership. Thus, ALCOVE can exhibit abstract behavior. However, ALCOVE, like many models of learning, is limited to only learning regularities over features.

BRIDGES generalizes ALCOVE by extending the notion of attention shifting and similarity to include relational match. The model supports the notion that analogy to stored experiences and attention shifting are the only attributes required of a model to appreciate abstract relationships. Furthermore, by incorporating ALCOVE's attentional shifting mechanisms into BRIDGES, we forward an explanation of how perceived similarity can change over the course of learning as more predictive stimulus properties are accentuated. Such attentional shifts will prove useful in demonstrating how seemingly abstract understandings can arise from analogies to concrete experiences.

Numerous accounts of how people detect and grade these analogical similarities exist (CAB, Larkey, & Love, 2003; LISA, Hummel, & Holyoak, 1997; SME, Falkenhainer, Forbus, & Gentner, 1989). Most approaches assume that an analogy is a mapping from items in one domain to items in another domain (Gentner, 1983). The differences between the various models occur in how the domains are represented and exactly how the mapping process progresses. Although we illustrate BRIDGES using a variation of the theory behind the Structure-Mapping Engine (SME), other methods for detecting analogical fit could have instead been implemented. In fact, BRIDGES is readily implemented using radically different match procedures, such as transformation approaches. The transformation account holds that one analog is transformed into the other over a series of steps until they match (Hahn, Chater, & Richardson, 2003). The fewer the steps and the smaller their cost, the higher the resulting similarity is. Learning would progress in much the same way with attention shifting to more predictive transformations.

Au2

Comp. by: RMervin Date:27/2/08 Time:15:34:46 Stage:First Proof File Path://
pchns1301/WOMAT/Production/PRODENV/0000000001/0000007779/
0000000016/0000772317.3D Proof by: QC by: ProjectAcronym:bs:PLM
Volume:00005

Structure mapping holds that people encode stimuli (e.g., objects, scenes, events) in terms of predicate representations that capture relations among entities (e.g., Revolves(planets,sun)). Relations can serve as arguments to other relations (e.g., Causes(GreaterMass(sun,planets), Revolves(planets,sun))). Structure mapping posits that people align structured representations to find the most satisfying correspondences. Satisfying correspondences are those that map elements playing identical roles in corresponding relations. Higher-order relations (i.e., relations serving as arguments to other relations) can serve to disambiguate and strengthen candidate mappings between analogs. In the solar system/atom example, an analogical sounds alignment places the sun in correspondence with the nucleus and the planets in correspondence with the electrons. Sounder mappings lead to increased perceived similarity (Gentner & Markman, 1997).

A. Abstraction Through Attention Shifting

By bridging the work from analogy to category learning, the BRIDGES model is able to demonstrate relational abstraction by extending the traditional definition of similarity used in category learning to make use of analogical alignment. For any given comparison between a stimulus and a stored exemplar, BRIDGES considers all of the possible one-to-one mappings between the stimulus and each exemplar.[2] For each exemplar, similarity is determined according to a difference measure that incorporates notions of featural and relational mismatch (see Fig. 4). A relational mismatch of 1 occurs when a relation does not exhibit parallel connectivity[3] (i.e. the mapped entities play different roles in their respective relations, see the right panel of Fig. 4). A featural mismatch of 1 occurs when nonidentical entities or entities containing mismatching features are mapped to one another (see the right panel of Fig. 4). As Fig. 4 illustrates how these two measures can be at odds.

Both types of mismatch are weighted by attention weights and the sum of these attention weighted mismatches yields an overall differences measure that is inversely proportional to similarity. The mapping that maximizes similarity (i.e., minimizes attention weighted difference) is chosen. These exemplar similarity values serve as exemplar unit activations and are passed across association weights to category units (e.g., predator and prey).

[2] Models of analogical alignment avoid this exhaustive search by using heuristics to guide the mapping process. BRIDGES could be extended to incorporate these shortcuts, but instead we focus on the basic ideas underlying BRIDGES.

[3] Systematicity has been left out as a constraint because it falls out as a natural consequence of parallel connectivity (Larkey & Love, 2003).

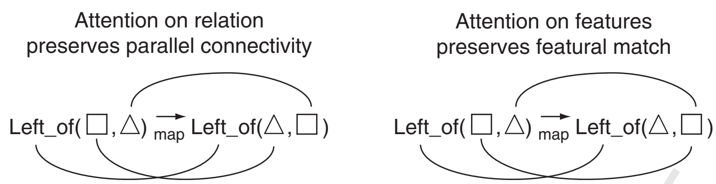

Fig. 4. There are two possible ways to map the elements in these corresponding relations. The example on the left preserves parallel connectivity by mapping elements that play the same role in each relation to one another. This solution is high in relational match, but low in featural match because the corresponding elements differ in shape features. The situation is reversed in the mapping shown in the right example. Attention weighting of mismatches determines which of these two possible mappings will be preferred by BRIDGES. BRIDGES chooses the mapping that minimizes attention-weighted mismatch.

The stimulus tends to be classified into the category whose unit has the greatest activation.

After feedback is provided, attention weights and association weights between exemplars and category units are adjusted to reduce error. Changes in attentional weights can lead to different future mappings. When attention shifts away from features and toward relations, parallel connectivity (i.e., analogical match) is stressed over featural similarity and BRIDGES demonstrates abstract understanding of a domain. Conversely, when featural matches lead to successful predictions, attention shifts toward features and BRIDGES is governed by featural similarity. When relational information is not discriminative or present, BRIDGES reduces to the standard ALCOVE model. A complete formalism for BRIDGES can be viewed in Tomlinson and Love (2006).

B. Relations Without Rules

Marcus, Vijayan, Bandi Rao, & Vishton (1999) found that infants could discriminate between abstract patterns or grammars of speech sounds. Importantly, this discrimination could not be accomplished by any weighting of phonetic features. Because featural regularities could not be leveraged to discriminate between grammars, Marcus et al. proposed that infants utilized variable binding in conjunction with algebraic rules to master such learning tasks.

We posit, based on evidence provided by BRIDGES, that infants do not use algebraic rules. Instead, the rule-like behavior arises from a comparison process between stimuli and previous exemplars coupled with a learned attention shift away from the concrete features of the stimuli to the relations within the stimuli.

Comp. by: RMervin Date:27/2/08 Time:15:34:46 Stage:First Proof File Path://
pchns1301/WOMAT/Production/PRODENV/0000000001/0000007779/
0000000016/0000772317.3D Proof by: QC by: ProjectAcronym:bs:PLM
Volume:00005

In Marcus et al. (1999), 7-month-old infants were exposed to sentences that followed either an AAB pattern or an ABB pattern. The sentences were made up of simple monosyllable sounds (words) such as "GA TI TI." After training on one of the patterns, the infants were tested with novel phonemes following the same patter or the opposite pattern (e.g., "BA BA GU"). The researchers found that 15 of the 16 infants were able to distinguish between the two grammars using novel phonemes.

BRIDGES was fit to the study. Each sentence (e.g., "GA TI TI") was represented as an exemplar. BRIDGES's exemplar representation for "GA TI TI" is shown in Table II. Each syllable is represented as an entity involved in a type-token relation with an abstract representation of the type. Each syllable's position in the speech stream is encoded by a positional feature. These syllables have a number of phonetic features that are not represented in these simulations. Not including such features follows Marcus et al.'s presumption that no significant regularities exist across these features. Importantly, including uncorrelated features does not alter the pattern of our simulation results.

Critically, relational information was included in BRIDGES's representations. BRIDGES makes a distinction between tokens and types. In effect, we assume that infants have developed categories of speech sounds (Eimas, Siqueland, Jusczyk, & Vigorito, 1971). These type relations allow for abstract patterns to be uncovered through analogy to stored exemplars as one type of sound can be mapped to another.

Following Love et al. (2004), unsupervised learning was modeled by a network consisting of a single category output unit with a target value of 1 for all stimuli. In effect, this category unit is a familiarity detector. Association and attention weights in the model were learned which uncovered the underlying regularities across the sentences to yield consistently high familiarity (i.e., high output values for the category unit).

During habituation, the 16 unique sentences were presented three times each to BRIDGES and stored as exemplars. On each presentation, association and attention weights were updated. Though not critical, we assumed

TABLE II

BRIDGES'S REPRESENTATION OF "GA TI TI"

Entities	Features	Relations
GA_1	Position $(GA_1) = 1$	Type of (GA_1, GA)
TI_1	Position $(TI_1) = 2$	Type of (TI_1, TI)
TI_2	Position $(TI_2) = 3$	Type of (TI_2, TI)

Comp. by: RMervin Date:27/2/08 Time:15:34:46 Stage:First Proof File Path://
pchns1301/WOMAT/Production/PRODENV/0000000001/0000007779/
0000000016/0000772317.3D Proof by: QC by: ProjectAcronym:bs:PLM
Volume:00005

that the saliency of positional features is sufficiently great to constrain the mapping process (i.e., words in sentences align temporally). Besides position (which does not discriminate between grammars), no regularities across features or entities existed. However, parallel connectivity was perfect for members of the same grammar. For instance, "GA TI TI" is isomorphic to "LI NA NA" in that all token and types in the type relation (see Table II) can be mapped to one another and preserve parallel connectivity. This degree of perfect match caused BRIDGES to shift attention to the type relation. This shift makes BRIDGES sensitive to the underlying grammar and immune to changes of the phonemes, rendering novel sentences following the original grammar somewhat familiar. Sentences not following the learned grammar can be mapped to the stored exemplars, but parallel connectivity is violated making these items less familiar and resulting in greater looking time as infants dishabituate.

BRIDGES was able to learn to discriminate between abstract patterns on the basis of analogical similarity. Storing concrete exemplars, shifting attention, and analogical matching are sufficient to show generalization to novel items. BRIDGES's success calls into question Marcus et al.'s (1999) claim that algebraic rules underlie infant performance. However, BRIDGES's success is attributable to its ability to bind arguments to relations, which is consistent Marcus et al.'s claim that infants bind variables.

Marcus (1999) has criticized other accounts (Seidenberg & Elman, 1999) of these results for including a same/different detector. The BRIDGES simulations do not explicitly label speech sounds as identical, rather the model assumes that infants can categorize speech sounds, as embodied by the type/token distinction. BRIDGES's solution does not hinge on a same detector. In fact, the patterns that can be discriminated by analogical mapping (even in simple domains in which only the type relation is present) are more encompassing than the concepts *same* and *different*. The analogical mapping process in these simulations aligned the current stimulus to stored exemplars—BRIDGES did not label words within sentences as same or different nor did it shift attention to a same feature. Abstract responding arose through analogy to stored exemplars and attention shifting from concrete features to relations.

C. Graded Responses as Similarity

One of the key assumptions of BRIDGES is that the basis for understanding abstract relations is similarity based and therefore inherently graded. The design of the Marcus et al. study did not allow for assessment of this possibility because stimuli were either grammatical or ungrammatical. According to BRIDGES, learners can both respond abstractly (i.e., generalize to featurally

Comp. by: RMervin Date:27/2/08 Time:15:34:46 Stage:First Proof File Path://
pchns1301/WOMAT/Production/PRODENV/0000000001/0000007779/
0000000016/0000772317.3D Proof by: QC by: ProjectAcronym:bs:PLM
Volume:00005

novel stimuli) and show evidence for the influence of past examples. If BRIDGES is correct, category membership in relationally defined categories is graded as it is in natural categories (Rosch & Mervis, 1975).

To evaluate BRIDGES's predictions, we will consider results from a series of studies exploring how pigeons and humans learn notions of same and different. To illustrate how BRIDGES learns the concepts same and different, we applied BRIDGES to Young and Wasserman's (1997) study demonstrating that pigeons can learn to discriminate between arrays of same and different icons. To foreshadow, Young and Wasserman's results indicate that pigeons can master a notion of same and different that cannot be explained by featural similarity. At the same time, the pigeons are sensitive to the particular examples they experienced during training and display a graded notion of same and different. Although fascinating, it would be easy to dismiss these results as relevant to pigeon cognition, but not human cognition. However, later work found the same pattern of performance with human subjects (Castro, Young, & Wasserman, 2006; Young & Wasserman, 2001). Humans as a group are slightly more deterministic than pigeons, but this group difference is within the range of individual differences. The bottom and top 20% of humans clearly bracket the mean performance of pigeons.

In Young and Wasserman (1997), pigeons learned to respond differentially to displays containing 16 identical and 16 different icons. On each trial, the 16 icons were randomly placed within a 5×5 grid. The pigeons were reinforced for pushing a green button when presented with a same stimulus and a red button when presented with a different stimulus. Training consisted of blocks of 16 same stimuli and 16 different stimuli in a random order. An identical set of icons was used to form stimuli for both the same and different items, making it impossible to correctly associate an icon or icon feature with a response. The pigeons were trained until 80% accuracy and then tested.

The test phase consisted of intermediate stimuli that were somewhat similar to both the same and different stimuli experienced in the training phase. Some examples of the intermediate stimuli are shown in Fig. 5. These stimuli can be viewed as forming a continuum between the pure same stimuli (all 16 icons identical) and the pure different stimuli (all 16 icons different) used during the training phase. The pigeon's performance in these intermediate conditions, as well as BRIDGES's predictions, is shown in Fig. 6.

Like the Marcus et al. (1999) simulations, we adopted a minimal approach to stimulus representation in fitting the BRIDGES model. Each stimulus's icon was represented as an entity. Each of the 16 entities participated in a type relation as in the Marcus et al. simulations (see Table II).

Through training BRIDGES discovered analogical mappings among presented stimuli and exemplars stored in memory that correctly predicted the label for the training stimuli. For example, consider aligning a stimulus

Comp. by: RMervin Date:27/2/08 Time:15:34:46 Stage:First Proof File Path://
pchns1301/WOMAT/Production/PRODENV/0000000001/0000007779/
0000000016/0000772317.3D Proof by: QC by: ProjectAcronym:bs:PLM
Volume:00005

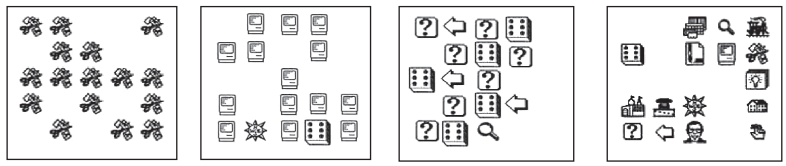

Fig. 5. Pigeons were trained on pure same or different stimuli like those shown on the far left and far right. At test, intermediate cases between pure same and different stimuli were shown like the middle two stimuli.

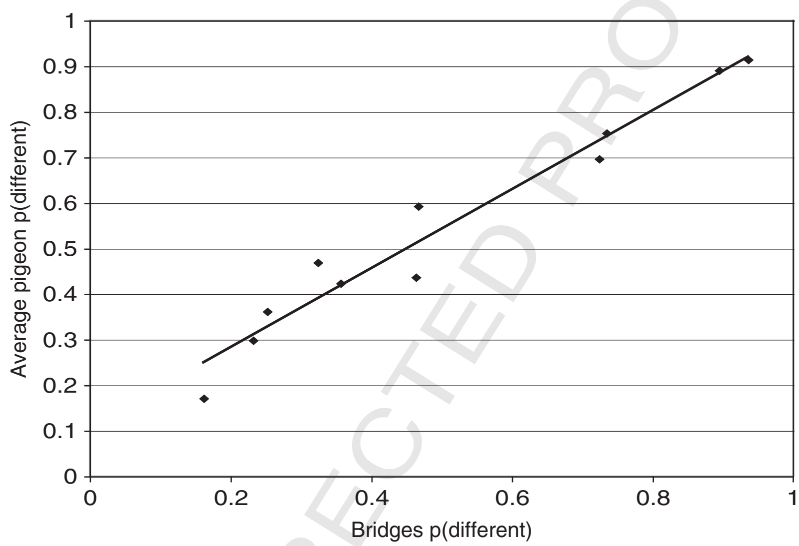

Fig. 6. The results from Young and Wasserman's (1997) studies and BRIDGES's predictions are shown. The 11 intermediate conditions form a continuum between pure same and pure different stimuli. Figure 5 provides two examples of such intermediate cases. Detailed descriptions of these intermediate can be found in Young and Wasserman (1997).

containing 16 squares to another stimulus containing 16 triangles. Each triangle entity is put into correspondence with a square entity. This results in a perfect feature mismatch, but parallel connectivity is preserved. Within each type relation, the type triangle maps to the type square. This alignment leads to attention shifting toward the type relation and away from the entities. In contrast, only 1 out of 16 type relations will exhibit parallel connectivity when aligning a different stimulus with a same stimulus. Thus, it is straightforward for BRIDGES to discriminate between same and different stimuli in the absence of featural support, by learning to pay attention to the number of type relations that match between the stimuli.

BRIDGES used the same similarity measure with the intermediate stimuli during testing. For example, the second stimulus from the right in Fig. 5 would result in 7 of 16 (seven identical icons) relational matches to each same exemplar and 4 of 16 (four distinct icon types) relational matches to each different exemplar. Similarity-based activations are not all or none and these intermediate cases activate stored examplars to varying degrees, leading to gradation in response. This gradation of response almost exactly matches the responding of the pigeons to the intermediate stimuli ($R^2 = .95$).

This simulation demonstrates how abstract concepts can be acquired through storage and analogy to concrete examples. BRIDGES's excellent fit of the intermediate conditions is a natural consequence of similarity-based processing. Like natural categories, BRIDGES predicts that relational categories have a graded structure.

D. Discussion

By combining insights from the category learning and analogy literatures, BRIDGES provides an account of how people and animals can gain abstract understandings of domains based solely on experience with concrete instances. BRIDGES's power arises from using a notion of similarity informed by work in both analogy and category learning. Structural alignment processes allow BRIDGES to appreciate analogical similarities, while attention shifting modifies BRIDGES notion of similarity over the course of learning. Integrating these mechanisms allows BRIDGES to grasp abstract patterns by shifting attention to relations which drive the alignment process.

In the supportive simulations, BRIDGES offered an explanation of how infants become sensitive to abstract grammars and how people and pigeons develop the concepts of same and different irrespective of a stimulus's features. BRIDGES's exemplar-based representation of experienced examples and trial-by-trial error-driven learning were sufficient to capture seemingly abstract concepts. Consistent with the stance that abstract concepts are similarity based, the relational concepts same and different display graded membership like natural categories.

BRIDGES is not the first model to use analogical alignment to support category learning. SEQL can acquire category structures through a process of repeated abstraction of a structured category representation and has been successfully applied to the infant grammar learning studies considered here (Kuehne, Gentner, & Forbus, 2003). While SEQL stresses building abstract representations, abstraction in BRIDGES arises from shifting attention. Some relative strengths of BRIDGES are that it extends an existing model

Comp. by: RMervin Date:27/2/08 Time:15:34:46 Stage:First Proof File Path://
pchns1301/WOMAT/Production/PRODENV/0000000001/0000007779/
0000000016/0000772317.3D Proof by: QC by: ProjectAcronym:bs:PLM
Volume:00005

of category learning (ALCOVE is a special case of BRIDGES) and incorporates attentional mechanisms.

One challenge for BRIDGES is incorporating new relational information into its exemplar representations. Although BRIDGES can learn relational concepts, BRIDGES is not yet able to incorporate acquired relations directly into its exemplar representations (see Doumas & Hummel, 2005, for an example of a predicate discovery system).

IV. Learning to Reason About Rewards in Dynamic Environments

The previous examples demonstrate how abstract rule-like behavior may be rooted in concrete experience. As a final example, we examine how people learn to control a dynamic system that continually evolves in response to their ongoing interaction. Like the categorization studies reviewed above, these experiments and simulations emphasize the way people learn (what appear to be) abstract rules that help them achieve their current goal. However, like our previous simulations, we will show how this rule-like behavior may emerge, not from abstract reasoning processes, but through simple trial-by-trial updates of concrete task representations.

A. Learning in Dynamic Tasks

Flexible adaptation to our environment requires us to learn in a variety of ways. Fortunately, a teacher is often times available to provide us with corrective feedback or instruction (i.e., supervised learning). However, in other situations we must discover for ourselves the relevant contingencies in the world through trial and error. For example, many video games require players to discover sequences of actions that allow them to advance to a new level through trial-and-error learning. In a set of recent studies (Gureckis & Love, xxxxa,b), we have examined how people discover response strategies that maximize their long-term benefit in an interactive task environment where the structure of rewards continually change in response to the actions of the individual (see Berry & Broadbent, 1988, or Stanley, Mathew, Russ, & Kotler-Cope, 1989, for a similar approach). Au3

In our task (called the "Farming on Mars" task), subjects are given a cover story about controlling two fictitious farming robots on a distant planet in order to generate oxygen for future human inhabitants. On each trial, subjects are asked to make repeated choices between these two alternatives (i.e., robots) with the goal of maximizing the rewards (i.e., oxygen) they receive over the entire experiment. On any given trial, one robot always generated more oxygen than the other. However, each time the subject selects

Comp. by: RMervin Date:27/2/08 Time:15:34:46 Stage:First Proof File Path://
pchns1301/WOMAT/Production/PRODENV/0000000001/0000007779/
0000000016/0000772317.3D Proof by: QC by: ProjectAcronym:bs:PLM
Volume:00005

this more attractive alternative, the long-range expected utility of both robots is lowered on the following trial. Thus, the strategy which provides the most reward over the entire experiment is to choose what appears to be the immediately inferior robot on every trial. Critically, the nature of this contingency is only revealed to subjects through their ongoing interaction with the system (i.e., responses and associated rewards) and subjects must learn for themselves how to maximize their earnings.

The conflict between short- and long-term rewards in our task is relevant to many realworld decision making situations. For example, drug addicts are often unable to overcome the desire for an immediate high despite the long-term consequences to their health. Interestingly, a number of studies that have examined choice behavior in tasks where short- and long-term rewards conflict have found that both humans and other animals often fail to inhibit the tendency to select an initially attractive option even when doing so leads to lower rates of reinforcement overall, a phenomena referred to as melioration (Herrnstein, 1991; Herrnstein & Prelec, 1991; Neth, Sims, & Gray, 2006; Tunney & Shanks, 2002). This pattern of results appears at odds with rational accounts, which dictate that decision makers follow a strategy that maximizes their long-term expected utility (see Tunney & Shanks, 2002, for a similar discussion). However, the rational account fails to specify how this optimal strategy should be determined in an unknown environment.

Indeed, one interesting question concerns the process by which learners arrive at effective strategies in these kinds of highly interactive and dynamic tasks. By one account, subjects reason in an abstract way about the contingencies in the task and effortfully evaluate the costs and benefits of various choice strategies. Such a strategy might involve explicit planning and sequencing of future choices, as well as evaluation of candidate plans through hypothesis testing. Instead, we propose that people engage in a trial-by-trial value estimation procedure using reinforcement learning. Success under this scheme is closely linked to people's representation of the underlying state of the system, which can be inferred from perceptual cues in the environment. Here, as in our previous examples, abstract behavior appears to emerge naturally from a learning process which is more directly grounded in concrete experience.

B. Bootstrapping Abstract Strategies from Concrete Cues

In one set of experiments (Gureckis & Love, xxxxb), we evaluated how subjects use cues about their current state to support their learning and decision making abilities. The term state simply refers to a discrete situation that a situated agent may be in with respect to their environment. For example, a state might correspond to an agent being in a particular location Au4

Comp. by: RMervin Date:27/2/08 Time:15:34:47 Stage:First Proof File Path://
pchns1301/WOMAT/Production/PRODENV/0000000001/0000007779/
0000000016/0000772317.3D Proof by: QC by: ProjectAcronym:bs:PLM
Volume:00005

in a maze. For agents interacting in an unknown environment, successful identification of the current state can provide important source of information that can help structure their behavior. For example, when an agent determines they are in a particular location in a maze, they can use that knowledge to help determine which direction to head next in order to increase the chances of successful escape. While representations of the current state can help determine which actions to take in particular situations, identifying task-relevant states in an unknown environment can be challenging. For example, in a maze it is often possible to arrive at physically different locations which are perceptually identical (for example, two hallways which have the same configuration of doorways and junctions). In this case, the agent must deal with the problem of perceptual aliasing (McCallum, 1993; Whitehead & Ballard, 1991), where multiple percepts may map onto a single state or situation in the world.

Determining the mapping from observations in the world to relevant task states about which the agent can learn is often a nontrivial problem. For example, in the standard version of the Farming on Mars task, successive states of the task environment are not clearly distinguished from one another. Figure 7 and the associated caption describe the payoff structure of the basic version of the task. Note that each time a player makes a choice, the underlying state of the task system can change so that reward possibilities on the next trial are different than they were on the previous trial. However, no discriminative cues are available about how the current state of the system has changed as a result of the agents actions. As a result, functionally distinct task states which lead to different rewards tend to alias together, making it difficult for the learner to detect the subtle contingency between their past actions and future rewards. Indeed, one hypothesis tested in our studies is that the aliasing of functionally distinct task states might partially explain why subjects seem to prefer short-term, melioration strategies.

We consider how to limit this aliasing and encourage learners to adopt useful representations of the state structure of the task. Subjects in our experiment were assigned to one of three conditions which provided different kinds of perceptual cues about the current state of the Mars farming system. In the no-cue condition, subjects were tested in the simple two-choice task described above and were given no additional information. Subjects in this condition were susceptible to perceptual aliasing of successive task states as there was no discriminative information available to indicate how the task environment was changing in response to the learners' actions. In two other conditions (the shuffled-cue and consistent-cue conditions) subjects display was augmented to include a row of indicator lights. Only one of these lights was active at any point in time, the position of which correlated with the current state of the system.

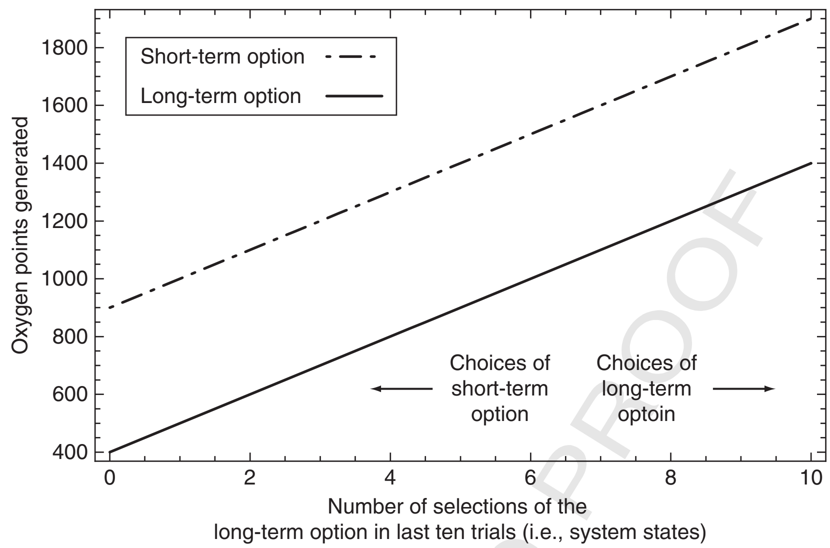

Fig. 7. The payoff structure for the Farming on Mars task used in Gureckis & Love (xxxxb, Exp. 1). Unknown to subjects, one option (i.e. robot) always generates more oxygen than the other on any given trial. For example, at the midpoint along the horizontal axis, selecting the more immediately productive option (referred to as the short-term option) would generate 1300 oxygen units, whereas selecting the other option (referred to as the long-term option) would generate only 800 oxygen units. However, each time the short-term option is selected, the expected output of both options is lowered on the following trial (i.e., the state of the system shifts to the left). Selections of the long-term option behave in the opposite fashion. The current state of the system (i.e. position along the horizontal axis) depends on the number of long-term selections subjects make over the preceding ten trials. Note that the reward received from repeatedly selecting the long-term option exceeds that from always selecting the short-term option (i.e., the highest point of the long-term option curve is above the lowest point for the short-term option curve). As a result, the optimal strategy is to select the long-term option on every trial, even though selecting the short-term option would earn more on any single trial. Au17

In the consistent-cue condition, the indicator lights were organized in a regular fashion such that the active light moved one position either to the left or to the right as the state was updated. In other words, the light veridically represented the current position of the task environment along the horizontal axis in Fig. 7. In this condition, neighboring states in the task corresponded to neighboring positions of the indicator light. The shuffled-cue condition also featured indicator lights, but the relationship between neighboring states was obscured by randomly assigning each state to a particular light. In the shuffled-cue condition, the position of the light was still perfectly predictive of the underlying state of the farming system, but the relationship between successive states and the magnitude of the reward signal was less obvious

Comp. by: RMervin Date:27/2/08 Time:15:34:47 Stage:First Proof File Path://
pchns1301/WOMAT/Production/PRODENV/0000000001/0000007779/
0000000016/0000772317.3D Proof by: QC by: ProjectAcronym:bs:PLM
Volume:00005

because neighboring states are unlikely to have neighboring indicator lights (e.g., the lit indicator could "jump" as one moves from state to state). We predicted that systematic cues which afford generalization between successive states (such as those in the consistent-cue condition) would be most effective for learning as they allow experience in one state to be easily generalized to related states. Note that the addition of the perceptual cue in the consistent-cue and shuffled-cue conditions does not render the task trivial—subjects still need to learn that the short-term option yields less reward over the course of the experiment than the long-term option to excel at the task.

Figure 8 shows the basic pattern of results from Experiment 1 of Gureckis and Love (xxxxb). In the no-cue condition, subjects had no strong preference Au5 for either option (the proportion of trials in which the long-term maximizing response was selected was .52). In contrast, subjects in the consistent-cue condition made significantly more selections of the long-term, reward-maximizing option (.76) compared to either the shuffled-cue (.57) or no-cue conditions. This pattern of results suggests that an important factor limiting performance in the standard version of the task is in identifying the current task state. When given a simple cue which reflected the underlying state of the system (the purpose of which was never explicitly explained), subjects performance drastically improved. Note that the advantage for state information also extends to the shuffled-cue condition. Subjects in this condition started out with a tendency to select the short-term, impulsive option, but later switched to moderately favor the long-term, maximizing-response option. However, subjects' performance appears to benefit most in the consistent-cue condition where the dynamics of the state indicator cue were consistent with the true underlying structure of successive task states.

In summary, subjects in our experiment generally failed to find a reward maximizing strategy unless they were given cues about the current state of the task environment. These results are particularly interesting, given previous work showing that across a variety of manipulations designed to encourage long-term responding, subjects seem to prefer suboptimal strategies (Neth et al., 2006; Tunney & Shanks, 2002). Across the three experimental conditions, subjects' task was held constant. The only difference concerned the presence and dynamic structure of cues that reflected of the underlying state of the task environment. Subjects's learning abilities appear to bootstrap off these concrete perceptual features. Without cues, learners had difficulty constructing an abstract representation of the task dynamics which impaired their ability to detect the relationship between their actions and future rewards. The ability to effectively discover and use an abstract rule may first depend on identifying the correct mental representation of the task environment, which itself maybe tied to concrete, perceptual aspects of the learning environment.

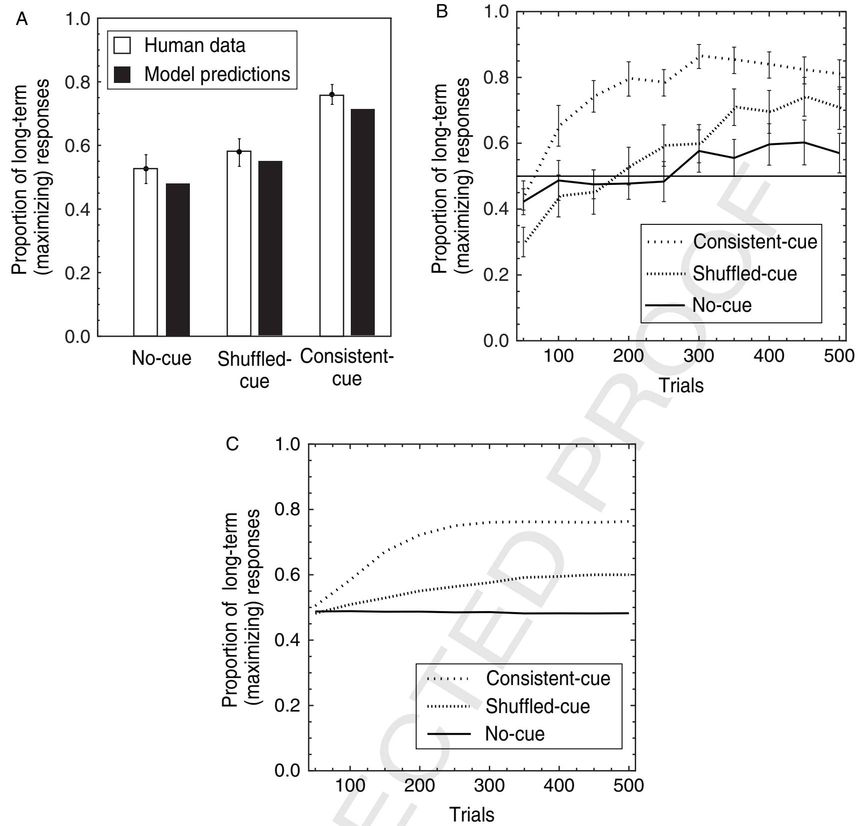

Fig. 8. Overall results of Experiment 1 and accompanying simulations from Gureckis and Love (xxxxb). Panel A shows the average proportion of maximizing response made throughout the experiment as a function of condition for both human subjects and the model. Panel B presents performance curves for human subjects calculated as the average proportion of maximizing responses considered in blocks of 50 trials at a time for all three condition. Error bars are standard errors of the mean. Panel C shows the same performance curves for the model. Au18

The time course of learning in the task reveals interesting details about the way that subjects learned about the reward contingencies in the task. In all three conditions, subjects started out with a bias toward the short-term impulsive option, but gradually increased the proportion of trials on which they selected the long-term option (Fig. 8B). The specific shape of this pattern (early melioration followed by a gradual shift toward maximizing) suggests that subjects incrementally explored the task environment. Early in learning

Comp. by: RMervin Date:27/2/08 Time:15:34:47 Stage:First Proof File Path://
pchns1301/WOMAT/Production/PRODENV/0000000001/0000007779/
0000000016/0000772317.3D Proof by: QC by: ProjectAcronym:bs:PLM
Volume:00005

and prior to much exploration, subjects settled on a suboptimal short-term strategy. However, later in the task, as random, exploratory actions accumulated performance gradually improved. Closer analyses of individual subjects response trajectory showed that the shift from short- to long-term responding was in most cases gradual, taking place over a period of 100 or more trials.

C. Using Reinforcement Learning to Model Task Performance

In order to understand the cognitive processes that underlie performance in the task we developed a simple computational model within the Reinforcement Learning (RL) framework (Sutton & Barto, 1998). RL is an agent-based approach to learning through interaction with the environment. In recent years, RL has attracted considerable attention based on its success in both practical applications (such as flying helicopters, Bagnell & Schneider, 2001, and teaching computers to play backgammon, Tesauro, 1994), and in Au6, 7 the modeling of biological systems (Daw & Touretzky, 2002; Montague, Dayan, & Sejnowski, 1996; Montague, Dayan, Person, & Sejnowski, 1995; Schultz, Dayan, & Montague, 1997; Suri, Bargas, & Arbib, 2001). What makes RL an interesting tool for understanding higher-level cognition is the fact that it details how situated learners can bootstrap effective behavioral strategies through self-guided interactions with an unknown environment.

A key feature of the simple RL models considered here is their ability to develop strategies that take into account delayed outcomes. However, unlike classical planning methods from artificial intelligence, our model can learn what appear to be far-sighted strategies in an incremental, trial-by-trial fashion. Our account is fundamentally shortsighted and reactive: The agent's only goal is to greedily choose whichever action is estimated to be best on each trial. However, through a trial-by-trial update procedure the model is able to learn estimates which reflect the long-term value of actions. Through a simple bootstrapping procedure (related to the well-known temporal-difference algorithm, Sutton & Barto, 1998), immediate online experience is passed backwards to adjust the estimates of previous executed actions. Through a combination of learning, exploration, and bootstrapping, globally optimal rule-like behavior can emerge in the model through direct error-driven learning. In the interest of brevity, we refer the reader to Gureckis and Love (xxxxb) for the mathematical details of the model, and instead Au8 elaborate three psychological principles that motivate the model.

Comp. by: RMervin Date:27/2/08 Time:15:34:47 Stage:First Proof File Path://
pchns1301/WOMAT/Production/PRODENV/0000000001/0000007779/
0000000016/0000772317.3D Proof by: QC by: ProjectAcronym:bs:PLM
Volume:00005

1. *Principle 1: Learning Depends on Subjects' Mental Representation of the Task*

On each trial, the model assumes that people attempt to choose the actions (or robot) that they estimate will earn the most reward over the long run. To do this, the model maintains estimates the long-term value of selecting a particular action *a* in state *s*, as *a* value referred to as *Q*(*s*, *a*) (or more generally a *Q*-value). Note, however, that because the estimate of any action depends on the current task state, *s*, subjects' identification of the true state of the environment plays a fundamental role in learning. Indeed, the model's estimate of the *Q*-value of a particular action on any trial is a simple linear function of the current input cues (which correspond to the indicator lights presented to subjects in the task). The model attempts to learn the mapping between input cues and the *Q*-values associated with each state using a simple single-layer network (Widrow & Hoff, 1960). Changing the types of input cues the model is given modulates the ability of the model to learn the appropriate representation of the state structure of the task and ultimately influences its ability to uncover an optimal response strategy.

2. *Principle 2: Learning Involves Taking into Account Future Rewards*

Each time an action is selected, the model updates its current estimate of the corresponding *Q*-value according to the temporal-difference error between the reward received as a result of that action and a current estimate of the future reward available from the following state-action pair. The model's asymptotic estimate of the value of each action depends on the relative weight given to immediate versus delayed rewards. In our model, the degree to which learners value short- or long-term rewards is determined by a simple discounting parameter (γ). When $\gamma = 0$, the learning procedure in the model reduces to the standard delta rule (Rescorla & Wagner, 1972; Wagner & Rescorla, 1972; Widrow & Hoff, 1960). Under these conditions, the model strongly favors immediate rewards (and thus predicts melioration behavior in the Farming on Mars task). As the value of γ increases, the model gives more weight to future rewards, eventually allowing it to favor selections of the long-term robot.

3. *Principle 3: Learning Involves Balancing Between Exploration and Exploitation*

Successful behavior in an unknown environment requires learners to balance the exploration of unknown alternatives versus the exploitation of resources or strategies that are known to be productive. In order to capture this

Comp. by: RMervin Date:27/2/08 Time:15:34:47 Stage:First Proof File Path://
pchns1301/WOMAT/Production/PRODENV/0000000001/0000007779/
0000000016/0000772317.3D Proof by: QC by: ProjectAcronym:bs:PLM
Volume:00005

tradeoff in the model, the estimated values of each state-action pair are input to a probabilistic choice mechanism that generates a final choice (Luce, 1959). Early in the task, this choice function favors exploration of actions which may appear to be less profitable. However, overtime the operation of this choice mechanism becomes more deterministic capturing the intuition that early in the experiment subjects are more willing to explore different alternatives, but later their behavior shifts to exploit the known resources they have uncovered.

In our simulations, the model was treated as an active participant in the Farming on Mars task and was given the same number of trials as human subject in which to explore the operation of the system and to uncover the optimal strategy. On each simulated trial the model selected either the short- or long-term option, the appropriate rewards were delivered, and the current state of the task system was updated. Through its ongoing experience the task, the model learns estimates of the value of choosing either the short- or long-term option via an adaptive learning procedure described above (Sutton, 1996; Sutton & Barto, 1998). Au9

Changing the types of cues the model is given about the current state of the system modulates the representation of the task that the model utilizes and ultimately influences its ability to uncover the optimal strategy. In simulations of the no-cue condition, the model (like human subjects) was given no discriminative information about the current task state. As a result, the model was forced to treat the task as consisting of just one, highly aliased state. In the consistent-cue condition, the model was presented with a set of input units that distinctly encoded the current position of the indicator light on the display. The input representation given to the model roughly matches the input provided to subjects, and allowed prediction and generalization between nearby states. Finally, in the shuffled-cue simulations, the input to the model was erratic. In some cases, there were source of predictability that the model could take advantage of, but in others the input cues could be misleading or force the model to become trapped in local, suboptimal solutions.

Figure 8A and C show the basic pattern of results. Using a single set of parameters across all three conditions the model was able to provide an excellent fit to the data. While the best fit parameters used to generate the learning curves in Fig. 8 tended to emphasize considerable exploration early on (thus the near-chance responding in the first few trials of the task), many other parameter combinations allowed the model to capture the trial-by-trial dynamics of human performance curves. For example, after a few early selections and prior to much exploration of the system, the model generally favored the short-term strategy of selecting the impulsive option on each trial. However, after a short period of time following this impulsive strategy, the continued accumulation of exploratory decisions eventually helps to shift the model toward a reward maximizing strategy.

Critically, learning in the model takes the form of trial-by-trial updates to estimates of the long-term value of each action in each possible state. In each state, the model is biased to choose the action whose estimated earnings are highest. However, there is no global representation in the model that represents a general rule about how to behave. Instead, performance is heavily tied to moment-to-moment estimates of estimated values of particular actions. In addition, our simulations show that learning critically depends on the structure of cues in the environment, which, in turn, structure the model's mental representation of the task. In the no-cue condition, the model fails to show any learning because all states collapse together. Building from this representation, the model has trouble developing a strategy which favors either option. In contrast, in the shuffled-cue and consistent-cue conditions, the structure of concrete, perceptual information in the environment helps to provide a framework for interpreting and integrating feedback.

D. Discussion

The experimental and simulation results reviewed in this section suggest that reasoning about complex systems can be accomplished through concrete means. All three conditions in our experiment involved the same reward structure, but subjects' performance across conditions varied as a function of the perceptual cues provided. We had success in helping subjects enrich their representation of the task environment by providing perceptual cues which reflected the underlying task state.

Overall, our results suggest that debates about rationality or optimality may be ill conceived. Learning in our task appears optimal with respect to the representation the learner adopts of the task environment which may or may not be congruent to the actual environment dynamics. The ability to approach the true "optimal" depended on the structure of concrete cues about system state in the environment. In this sense, abstract performance was shown to depend on concrete features of the environment.

Finally, our results show that subjects' ability to learn an effective control strategy is well accounted for by a simple RL learning based on trial-by-trial learning updates and simple task representations. Unlike other methods, the RL agent simulations did not explicitly plan, but instead reasoned in a reactive manner. Rather than engage in hypothesis testing strategies, the RL agent estimated the value of possible options by trial-by-trial updates. In addition, we offer a novel take on previous findings showing that humans and other animals prefer impulsive, short-term gains over higher utility future outcomes by suggesting that some of these failures may actually derive

Comp. by: RMervin Date:27/2/08 Time:15:34:47 Stage:First Proof File Path://
pchns1301/WOMAT/Production/PRODENV/0000000001/0000007779/
0000000016/0000772317.3D Proof by: QC by: ProjectAcronym:bs:PLM
Volume:00005

from how learners represent the relevant states and dynamics underlying their environment (as opposed to a failure in the ability to properly discount future outcomes).

V. General Discussion

When observing a complex system, it is tempting to ascribe sophisticated intentions and abilities. For instance, people sometimes believe machines, like personal computers and cars, are conspiring against them. The job of the cognitive psychologist is to observe the most complex of machines and entertain theories of how it works. Perhaps it is not surprising that many researchers believe humans possess transcendent and abstract thinking abilities. After all, humans appear to acquire abstract understandings of domains and reason in ways that indicate a deep understanding. In this chapter, we argued through three cases that, while humans can perform impressively and demonstrate broad generalization, our performance is intimately tied and colored by the concrete examples we experience during learning. In effect, we never fully transcend our training set.

In the first case study, we examined how people appear to acquire and apply rules. One popular account proposes multiple systems in which one system applies abstract rules and another system stores items that violate these rules (e.g., Nosofsky et al., 1994).

Instead, we found evidence for a single system in which both rule-following and rule-violating items are stored in a common substrate. Recognition performance for stimuli that violated rules indicated that the rules themselves were rooted in experienced examples. In particular, items that violated more frequent rules were remembered better. This effect was also found in studies that manipulated the similarity of rule-following items to exceptions from the opposing category. The results from both manipulations were explained by a clustering model that sharpened its memory representations for exception items in order to reduce confusions with rule-following items from the opposing category. The more extreme the confusions (because of either numerosity or similarity relations), the greater the enhanced memory for exception items. These results are both supportive of a clustering model that updates its representation by trial-by-trial learning, and inconsistent with dual route models that posit an abstract rule system. In summary, even when people mentally rehearse and consciously apply rules, the concrete details of the training examples affect performance.

The second case study makes a much stronger claim—abstract concepts and rules (i.e., concepts not reducible to a finite set of features) are in fact grounded in experienced examples. A model that stored each example in

Comp. by: RMervin Date:27/2/08 Time:15:34:47 Stage:First Proof File Path://
pchns1301/WOMAT/Production/PRODENV/0000000001/0000007779/
0000000016/0000772317.3D Proof by: QC by: ProjectAcronym:bs:PLM
Volume:00005

memory as an exemplar and made analogies to these stored exemplars was able to learn seemingly abstract concepts. In the first simulation, the model, BRIDGES, successfully simulated infant learning of abstract rules consisting of sound patterns. The model learned to extend the pattern to novel cases by making analogies to previous cases. This simulation alone provides a powerful existence proof that abstract rules may not be necessary. The second simulation tested a key prediction of BRIDGES that membership in seemingly abstract concepts is graded and determined by the similarity to stored examples. Data exploring how pigeons (and humans) learn the concepts of same and different were supportive of this prediction. After initial training on pure cases of the two concepts, intermediate cases partially activated exemplars from both concepts leading to graded responding. The exact nature of this gradient was predicted by BRIDGES notion of exemplar similarity, which incorporates a notion of analogical match.

In the final case study, we considered how people reason about rewards in dynamic environments in which learner's actions affect the underlying state of the system (e.g., much like how humans affect the climate). This study brings our analysis closer to the kinds of learning problems that people face in their daily lives. In these situations, do people reason explicitly about system dynamics, form a plan for future actions, and test hypotheses about the nature of the underlying system? While climatologists might undertake such analyses, our results indicate that nonexperts learn about such systems by exploring myopically and adjusting estimates of future rewards in a trial-by-trial fashion.

People's performance was consistent with simple reinforcement learning models that incrementally learn about rewards and use these tentative estimates of reward to bootstrap learning the long-term value of particular actions. People's performance was consistent with simple reinforcement learning models. These models incrementally learn about reward. Current estimates of future reward are used to bootstrap improved estimates of the long-term value of particular actions. While reinforcement learning models may appear to follow farsighted strategies, they are essentially reactive and do not explicitly plan ahead. Such models are exquisitely sensitive to how the state of the system is represented. Human subjects also proved to be sensitive to state representation. When human subjects were given a consistent cue to the underlying state of the system, they were able overcome an option that was more attractive in the short term and instead learned to prefer the option that was optimal in the long term.

These three cases demonstrate the power and richness of learning systems that learn by exposure to concrete examples. All three case studies involve fairly sophisticated and adaptive behaviors. The three models reach human proficiency in each case and correctly predict the idiosyncrasies of human behavior. These idiosyncrasies arise from our concepts being rooted in experienced examples and updated by incremental learning processes. One general lesson from these studies is that one should be cautious in making claims about the representational machinery needed to capture human behavior.

The three case studies we considered here each invoke a different model-based explanation. One key question is whether the models from the three case studies are the same in some sense. Certainly, the models differ in their details, but all three models learn incrementally, are myopic, are bound by experienced examples, and do not engage in hypothesis testing. Numerous relations exist among the three models. The BRIDGES model from the second case study generalizes ALCOVE, which itself is conceptually related to the CLUSTER model from the first case study. Work is currently underway on a version of BRIDGES that, instead of exemplars, represents concepts by clusters that are recruited in response to surprising events as in SUSTAIN and CLUSTER. Matt Jones and the first author have developed a reinforcement learning version of the CLUSTER model. This model offers a connection to the reinforcement learning model showcased in the third case study. One exciting possibility is incorporating analogical notions of similarity into reinforcement learning models. Doing so would enable generalization across analogical states, which could greatly speed learning. All three models considered here incorporate error terms that are minimized. Although all three models are distinct, common themes emerge and avenues for unifying these three models are within reach.

One common theme that emerges across models is that details of the training set are retained. Humans appear to operate in a similar fashion. One question is why would evolutionary processes or within individual learning processes settle on such a solution. One answer is robustness. Grounding performance and generalization in concrete cues present in the environment is a conservative strategy that increases system robustness and stability. Models grounded in these details are less likely to become decoupled from environmental feedback and run amok. Also, details that appear irrelevant at the current moment may prove critical in the future. The complexity and dynamic nature of our environment might dictate such learning mechanisms.

Our discussion has been titled in favor of single system models over multiple system models, particularly in regards to multiple system models containing a rule route. To clarify our stance, we are not opposed to multiple system proposals in general, but we do believe that successful multiple system proposals will be bounded by experienced examples. Of course, proposing multiple systems does not in itself guarantee success. The second case study makes this point by demonstrating that existing dual system accounts do not account for human rule-plus-exception learning behavior. We take our results to indicate that models must be both adaptable and grounded in experienced examples, like the CLUSTER model. We are agnostic about whether such proposals are best described as single or multiple learning system models.

One strength of our theoretical outlook is that it does not demand leaps of imagination to accept. In contrast, we do not see in principle how the human brain could acquire truly abstract representations given that our perceptions and actions are embedded in a noisy world of concrete examples. Evoking a nativist position does not solve the conundrum, as abstract representations would need to be learned over evolutionary time and must be retrieved within the individual's life via perceptual cues in the environment. Our account also has the advantage of not being at odds with sensible intuitions about ontogenetic and phylogenetic continuity. We predict that other species (e.g., pigeons) are cleverer than we expect and that the abstract nature of our own thought is overstated.

Acknowledgments

This work was supported by AFOSR grant FA9550-04-1-0226, ARL grant W911NF-07-2-0023, and NSF CAREER grant 0349101 to B.C. Love.

References

Allen, S. W., & Brooks, L. R. (1991). Specializing the operation of an explicit rule. *Journal of Experimental Psychology: General, 120*, 3–19.

Ashby, F., & O'Brien, J. B. (2005). Category learning and multiple memory systems. *Trends in Cognitive Sciences, 9*, 83–89.

Ashby, F., & Maddox, W. (2005). Human category learning. *Annual Review of Psychology, 56*, 149–178.

Berry, D. C., & Broadbent, D. E. (1988). Interactive tasks and the implicit-explicit distinction. *British Journal of Psychology, 79*, 251–272.

Carpenter, G., & Grossberg, S. (2003). Adaptive resonance theory. In M. Arbib (Ed.), *The handbook of brain theory and neural networks* (pp. 87–90). Cambridge, MA: MIT Press.

Castro, L., Young, M. E., & Wasserman, E. A. (2006). Effects of number of items and visual display variability on same-different discrimination behavior. *Memory & Cognition, 34*(8), 1689–1703.

Chomsky, N. (1957). *Syntactic structures*. The Hague/Paris: Mouton.

Clapper, J. P., & Bower, G. H. (1994). Category invention in unsupervised learning. *Journal of Experimental Psychology: Learning, Memory, & Cognition, 20*, 443–460.

Daw, N., & Touretzky, D. (2002). Long-term reward prediction in td models of the dopamine system. *Neural Computation, 14*, 603–616.

Doumas, L. A. A., & Hummel, J. E. (2005). A symbolic-connectionist model of relational discovery. In *Proceedings of the twenty-seventh annual conference of the cognitive science society*. Hillsdale, NJ: Erlbaum.

Eimas, P., Siqueland, E., Jusczyk, P., & Vigorito, J. (1971). Infant speech perception. *Science, 171*(3968), 303–306.

Erickson, M. A., & Kruschke, J. K. (1998). Rules and exemplars in category learning. *Journal of Experimental Psychology: General, 127*, 107–140.

Falkenhainer, B., Forbus, K. D., & Gentner, D. (1989). The structure mapping engine: Algorithm and examples. *Artificial Intelligence, 41*, 1–63.

Forbus, K. D., Gentner, D., & Law, K. (1994). MAC/FAC: A model of similarity-based retrieval. *Cognitive Science, 19*, 141–205.

Gentner, D. (1983). Structure-mapping: A theoretical framework for analogy. *Cognitive Science, 7*, 155–170.

Gentner, D., & Markman, A. B. (1997). Structure mapping in analogy and similarity. *American Psychologist, 52*, 45–56.

Gureckis, T., & Love, B. (xxxx,a). Solving partially observable markov decision processes via reinforcement learning. Au10, 11

Gureckis, T., & Love, B. (xxxx,b). Short term gains, long term pains: Reinforcement learning in dynamic environments. Hahn, U., Chater, N., & Richardson, L. B. (2003). Similarity as transformation. *Cognition, 87,* 1–32. Au12

Herrnstein, R. (1991). Experiments on stable suboptimality in individual behavior. *The American Economic Review, 81*(2), 360–364.

Herrnstein, R., & Prelec, D. (1991). Melioration: A theory of distributed choice. *The Journal of Economic Perspectives, 5*(3), 137–156.

Hummel, J. E., & Holyoak, K. J. (1997). Distributed representations of structure: A theory of analogical access and mapping. *Psychological Review, 104*, 427–466.

Kruschke, J. K. (1992). ALCOVE: An exemplar-based connectionist model of category learning. *Psychological Review, 99*, 22–44.

Kruschke, J. K. (1993). Human category learning: Implications for backpropagation models. *Connection Science, 5*, 3–36. Au13

Kuehne, S., Gentner, D., & Forbus, K. (2003). Modeling infant learning via symbolic structural alignment. In *Proceedings of the twenty-second annual conference of the cognitive science society.* Hillsdale, NJ*:* Erlbaum.

Kurtz, K. J. (2007). The divergent autoencoder (DIVA) model of category learning. *Psychonomic Bulletin & Review, 14*(4), 560–576. Au14

Larkey, B. L., & Love, B. C. (2003). CAB: Connectionist Analogy Builder. *Cognitive Science, 27*, 781–794.

Love, B. C. (2005). Environment and goals jointly direct category acquisition. *Current Directions in Psychological Science, 14*, 195–199.

Love, B. C., & Gureckis, T. M. (2007). Models in search of a brain. *Cognitive, Affective, & Behavioral Neuroscience, 7*, 90–108.

Love, B. C., & Jones, M. (2006). The emergence of multiple learning systems. *Proceedings of the Cognitive Science Society*. Mahwah, NJ: Erlbaum. Au15

Love, B. C., Medin, D. L., & Gureckis, T. (2004). SUSTAIN: A network model of human category learning. *Psychological Review, 111*, 309–332.

Luce, R. D. (1959). *Individual choice behavior: A theoretical analysis*. Westport, CT: Greenwood Press.

Marcus, G. F. (1999). Do infants learn grammar with algebra or statistics? response to seidenberg & elman, eimas, and negishi *Science, 284*, 436–437.

Marcus, G. F., Vijayan, S., Bandi Rao, S., & Vishton, V. (1999). Rule learning by seven-month-old infants. *Science, 283*(5298), 77–80.

Markman, A. B., & Ross, B. H. (2003). Category use and category learning. *Psychological Bulletin, 129*, 592–613.

McCallum, R. (1993). *Overcoming incomplete perception with utile distinction memory.* In *the proceedings of the Tenth International Machine Learning Conference (ML'93)*. Amherst, MA.

Medin, D. L., & Bettger, J. G. (1994). Presentation order and recognition of categorically related examples. *Psychonomic Bulletin & Review, 1*, 250–254.

Medin, D. L., & Schaffer, M. M. (1978). Context theory of classification learning. *Psychological Review, 85*, 207–238.

Montague, P., Dayan, P., Person, C., & Sejnowski, T. (1995). Bee foraging in uncertain environments using predictive hebbian learning. *Nature, 377*(6551), 725–728.

Montague, P., Dayan, P., & Sejnowski, T. (1996). A framework for mesencephalic dopmaine system based on predictive hebbian learning. *Journal of Neuroscience, 16*(5), 1936–1947.

Neth, H., Sims, C., & Gray, W. (2006). Melioration dominates maximization: Stable suboptimal performance despite global feedback. In R. Sun and N. Miyake (Eds.), *Proceedings of the 28th annual meeting of the cognitive science society*. Hillsdale, NJ: Erlbaum.

Nosofsky, R. M. (1986). Attention, similairty, and the identification-categorization relationship. *Journal of Experimental Psychology: General, 115*, 39–57.

Nosofsky, R. M. (1991). Typicality in logically defined categories: Exemplar similarity versus rule instantiation. *Memory & Cognition, 19*, 131–150.

Nosofsky, R. M., Palmeri, T. J., & McKinley, S. C. (1994). Rule-plus-exception model of classification learning. *Psychological Review, 101*(1), 53–79.

Nosofsky, R. M., & Zaki, S. F. (1998). Dissociations between categorization and recognition in amnesic and normal individuals. *Psychological Science, 9*, 247–255.

Palmeri, T. J., & Nosofsky, R. M. (1995). Recognition memory for exceptions to the category rule. *Journal of Experimental Psychology: Learning, Memory, & Cognition, 21*, 548–568.

Pinker, S., & Prince, A. (1988). On language and connectionism: Analysis of a parallel distributed processing model of language acquisition. *Cognition, 28*, 73–193.

Rescorla, R., & Wagner, A. (1972). A theory of pavolvian conditioning: Variations in the effectiveness of reinforcement and non-reinforcement. In A. Black and W. Prokasy (Eds.), *Classical conditioning ii: Current research and theory* (pp. 64–99). New York: Appleton-Century-Crofts.

Rosch, E., & Mervis, C. B. (1975). Family resemblences: Studies in the internal structure of categories. *Cognitive Psychology, 7*, 573–605.

Ross, B. H., & Kennedy, P. T. (1990). Generalizing from the use of earlier examples in problem solving. *Journal of Experimental Psychology: Learning, Memory, & Cognition, 16*, 42–55.

Rouder, J. N., & Ratcliff, R. (2006). Comparing exemplar- and rule-based models of categorization. *Current Directions in Psychological Science, 15*, 9–13.

Rumelhart, D. E., & McClelland, J. L. (1987). Learning the past tense of English verbs: Implicit rules or parallel distributed processing? In B. MacWhinney (Ed.), *Mechanisms of language aquisition* (pp. 195–248). Hillsdale, NJ: Erlbaum.

Sakamoto, Y., & Love, B. C. (2004). Schematic influences on category learning and recognition memory. *Journal of Experimental Psychology: General, 33*, 534–553.

Sakamoto, Y., & Love, B. C. (2006). Vancouver, Toronto, Montreal, Austin: Enhanced oddball memory through differentiation, not isolation. *Psychonomic Bulletin & Review, 13*, 474–479.

Schultz, W., Dayan, P., & Montague, P. R. (1997). A neural substrate of prediction and reward. *Science, 275*, 1593–1598.

Seidenberg, M., & Elman, J. (1999). Do infants learn grammars with algebra or statistics? *Science, 284*(5413), 433.

Sloman, S. A. (1996). The empirical case for two systems of reasoning. *Psychological Bulletin, 119*, 3–22.

Smith, J. D., & Minda, J. P. (1998). Prototypes in the mist: The early epochs of category learning. *Journal of Experimental Psychology: Learning, Memory, & Cognition, 24*, 1411–1430.

Stanley, W., Mathew, R., Russ, R., & Kotler-Cope, S. (1989). Insight without awareness: On the interaction of verbalization, instruction, and practice in a simulated process control task. *Quartley Journal of Experimental Psychology, 41A*(3), 553–577.

Suri, R., Bargas, J., & Arbib, M. (2001). Modeling functions of striatal dopamine modulation in learning and planning. *Neuroscience, 103*, 65–85.

Sutton, R., & Barto, A. (1998). *Reinforcement learning: An introduction*. Cambridge, MA: MIT Press.

Tomlinson, M., & Love, B. C. (2006). *From pigeons to humans: Grounding relational learning in concrete examples. Proceedings of the Twenty-First National Conference on Artificial Intelligence* (AAAI-2006).

Tunney, R. J., & Shanks, D. R. (2002). A re-examination of melioration and rational choice. *Journal of Behavioral Decision Making, 15*, 291–311.

Tversky, A., & Kahneman, D. (1981). The framing of decisions and the psychology of choice. *Science, 211*, 453–458.

Tversky, A., & Kahneman, D. (1983). Extension versus intuitive reasoning: The conjunction fallacy in probability judgment. *Psychological Review, 90*, 293–315.

Wagner, A., & Rescorla, R. (1972). Inhibition in pavlovian conditioning: Application of a theory. In R. Boake and M. Halliday (Eds.), *Inhibition and learning* (pp. 301–336). London: Academic Press.

Whitehead, S., & Ballard, D. (1991). Learning to perceive and act by trial and error. *Machine Learning, 7*(1), 45–83.

Widrow, B., & Hoff, M. (1960). Adaptive switching circuits. *Institute of Radio Engineers, Western Electronic Show and Convention Record, 4*, 96–104.

Yang, L. X., & Lewandowsky, S. (2004). Context-gated knowledge partitioning in categorization. *Journal of Experimental Psychology: Learning, Memory, & Cognition, 30*, 1045–1064.

Young, M. E., & Wasserman, E. A. (1997). Entropy detection by pigeons: Response to mixed visual displays after same-different discrimination training. *Journal of Experimental Psychology: Animal Behavior Processes, 23*, 157–170.

Young, M. E., & Wasserman, E. A. (2001). Entropy and variability discrimination. *Journal of Experimental Psychology: Learning, Memory and Cognition, 27*(1), 278–293.