Published in final edited form as:

Psychon Bull Rev. 2014 April ; 21(2): 312–331. doi:10.3758/s13423-013-0506-0.

Prototypes, Exemplars, and the Natural History of Categorization

J. David Smith

Department of Psychology, Center for Cognitive Science, University at Buffalo, The State University of New York

Abstract

The article explores—from a utility/adaptation perspective—the role of prototype and exemplar processes in categorization. The author surveys important category tasks within the categorization literature from the perspective of the optimality of applying prototype and exemplar processes. Formal simulations reveal that organisms will often (not always!) receive more useful signals about category belongingness if they average their exemplar experience into a prototype and use this as the comparative standard for categorization. This survey then provides the theoretical context for considering the evolution of cognitive systems for categorization. In the article's final sections, the author reviews recent research on the performance of nonhuman primates and humans in the tasks analyzed in the article. Diverse species share operating principles, default commitments, and processing weaknesses in categorization. From these commonalities, it may be possible to infer some properties of the categorization ecology these species generally experienced during cognitive evolution.

Keywords

categorization; prototypes; exemplars; comparative cognition; cognitive evolution

Introduction

Categorization is a focus of human and animal research because it is an essential cognitive function (e.g., Ashby & Maddox, 2005; Brooks, 1978; Feldman, 2000; Hampton, 2010; Jitsumori, 1994; Kemler Nelson, 1984; Knowlton & Squire, 1993; Lazareva & Wasserman, 2010; Lea & Wills, 2008; Malt, 1995; Medin & Schaffer, 1978; Murphy, 2003; Nosofsky, 1987; Pearce, 1994; Smith, Minda, & Washburn, 2004; Vauclair, 2002). This article considers the utility that organisms derive from category knowledge and how the form of that knowledge affects the utility it confers.

Categorization could be important enough to receive redundant expression in cognition. When organisms must identify mates or caregivers, categorization would best have a particularizing emphasis that identifies individuals. In contrast, when organisms must identify a whole class of objects (e.g., a predator species), categorization would best have a generalizing emphasis that supports adaptive behavior toward all class members. Accordingly, there has been interest in the possibility that learners store category exemplars as individuated memory representations (e.g., Medin & Schaffer, 1978; Nosofsky, 1987) and in the possibility that learners average their exemplar experience to form a general category representation or prototype (Homa, Sterling, & Trepel, 1981; Reed, 1972; Smith & Minda, 2002).

In fact, research has begun to document the tradeoffs among processes in categorization, giving each their due. For example, different processes dominate categorization at early and late stages of learning (Cook & Smith, 2006; Reed, 1978; Smith, Chapman, & Redford, 2010; Smith & Minda, 1998; Wasserman, Keidinger, & Bhatt, 1988), when categories are small or large (Homa et al., 1981; Minda & Smith, 2001), when categories are well or poorly structured (Blair & Homa, 2003; Smith, Murray, & Minda, 1997), and when category rules are easy or difficult to verbalize (Ashby, Alfonso-Reese, Turken, & Waldron, 1998; Ashby & Maddox, 2010). In the spirit of this balanced research literature, the present article shows that a prototype process sometimes fails disastrously, but other times brings performance benefits.

Based on extensive research like that just described, some theorists have adopted a multiple-systems theoretical perspective toward human categorization (Ashby & Maddox, 2010, Erickson & Kruschke, 1998; Homa et al., 1981; Minda & Smith, 2001, Rosseel, 2002; Smith & Minda, 1998), based on the idea that organisms have multiple categorization utilities that learn different distributional aspects of the environment. However, some still favor the goal to explain categorization broadly by assuming a unitary, exemplar-based process (e. g., Nosofsky & Johansen, 2000; Nosofsky, Little, Donkin, & Fific, 2011). In some respects, this article supports the exemplar perspective by describing conditions under which that process would be highly adaptive.

The present article takes a distinctive approach toward the idea of multiple categorization processes. Previously, researchers asked what subjects *do* do as they learn categories. They asked which processing assumption predicted best a target performance, no matter how accurate that target performance was. To this end, they used formal models to fit humans' or animals' categorization data. In contrast, the present article explores what subjects *should* do as they learn categories. That is, I consider how subjects might glean from their category experience the clearest signals of category membership, the clearest bases for discriminating category members and nonmembers, and therefore the most efficient categorization as they navigate their worlds. To this end, I use formal models to evaluate how objectively, absolutely well the prototype or exemplar process can learn different category structures. This approach bears on which process is inherently optimal for performing different tasks. Briscoe and Feldman (2011) adopted a similar optimality approach. It is remarkable how little attention has been given to how well humans and animals perform in category tasks, or to the strategies they should adopt to perform optimally.

The optimality approach complements the traditional fitting approach rather than supplanting it. The two approaches pursue different theoretical goals to potentially different outcomes. For example, the fitting approach might show that subjects use exemplar processing in a task even as the optimality approach showed that this strategy reduced absolute performance. In the present article, therefore, a prototype-process or exemplar-process advantage will not imply a fit advantage as the prototype or exemplar models fit data, but a performance advantage as the prototype and exemplar processes actually perform a category task.

Given this distinctive optimality approach, I will focus here on contrasting two categorization processes—exemplar and prototype processes—in an initial exploration of these issues. The prototype-exemplar debate has been highly influential within the categorization literature and thus it makes for an apt case study. However, I emphasize that other processes—especially rule-based processes—have also been essential parts of the multiple-systems literature (Ashby & Maddox, 2010). Going forward, parallel and equally important theoretical questions could be raised about the ecological optimality of these other processes. Indeed, some of these questions have begun to be addressed (Smith, Ashby, et al.,

2011; Smith, Beran, Crossley, Boomer, & Ashby, 2010; Smith, Crossley, Boomer, Church, Beran, & Ashby, 2012).

I proceed as follows. First, I analyze formally performance on a category task organized by family resemblance and prototypes. I evaluate through simulations how optimally humans or animals would perform in this task if they applied prototype or exemplar processes. Thus, I quantify the performance advantage achievable from matching correctly the categorization process to the category structure.

Second, I test the limits on this advantage from matching, by analyzing a task that presents the additional challenge of linear nonseparability as defined below. I evaluate again how prototype and exemplar processes would fare, now with simulations applied to prototype-exception category structures. Through this comparison, one can ask whether nonlinearity in category tasks affects the relative optimality of different categorization processes. The answer is surprising.

Third, I test the ultimate limits on the advantage from matching strategy to category structure, by analyzing a task that presents the ultimate nonlinearity challenge. I evaluate how prototype and exemplar processes would fare when applied to exclusive-or categories. These simulations represent a “win” for exemplar processing by identifying conditions that would make prototype processing distinctively nonoptimal.

These three category structures have influenced categorization research for decades. They have been essential in the field’s important debates (prototypes-exemplars, linear-separability constraints, multiple systems, etc.). They span the field’s history and theoretical approaches. They cover a range of degrees of linear separability. They resonate with recent cross-species research in categorization. For these reasons, they were good choices for exploring the general optimality of prototype and exemplar processes. This survey of category structures then provides the theoretical context for considering the evolution of cognitive systems for categorization.

The first consideration is that performance optimality could have fitness implications. The possibility of gains in foraging efficiency and predator avoidance could have exerted evolutionary pressure so that organisms developed well-tailored categorization systems for achieving these gains. It would be an important insight about the emergence of categorization systems if they have been adaptively tuned during cognitive evolution to the categorization ecologies within which they developed. Such evolutionary considerations have rarely been raised in the human literature. Indeed, anything regarding the relative fitness of different categorization processes has received scant attention, because in recent decades human categorization was often conceptualized as a unitary system that did not comprise different processes.

The second consideration is that the category structures generally experienced during cognitive evolution could determine the answer to the fitness or tuning question. For example, this article will show that exemplar processes provide a strong optimality advantage in learning some category problems. If the natural ecology were suffused with these problems, there would have been strong evolutionary pressure toward the development of robust and sensitive exemplar processes. Different categorization ecologies might have exerted different evolutionary pressures. Thus, I suggest that the course of categorization’s evolution could have been different, and could be different for different species, depending on the underlying statistical properties of natural kinds.

Because this article may seem to challenge the field’s dominant theoretical narrative, it is important to say that it does not matter which way this evolutionary tuning might go. If

exemplar-privileged categories were dominant in nature, they would be the dominant force in shaping the evolution of cognitive systems for categorization. Or, if prototype-privileged categories were dominant in nature, they would be the shaping force. Either dynamic would be profoundly interesting. Therefore, the theoretical approach and formal methods in this article are inherently neutral regarding either outcome, and I will deliberately demonstrate the advantage of both prototype and exemplar processes.

The third consideration is that the structure of cognitive systems for categorization might illuminate the statistical properties of natural kinds. For example, suppose one found that the categorization capabilities of a species were dominated by prototype-based processes that allow the efficient learning of family-resemblance categories. This could suggest that the species evolved and became adaptively tuned within a family-resemblance categorization ecology. This possible inference has also received scant attention in the human literature for this reason: if one supposes that categorization is a unitary system able to learn different category structures equivalently well, there would be no adaptive tuning of process to structure and so nothing to infer.

Yet research is bringing this inference within reach. Later in the article, I will review recent research on the performance of nonhuman primates and humans in the tasks formally analyzed in this article. This research suggests that diverse species emphasized during cognitive evolution one set of operating principles and default commitments, and one balance point among the processes they can bring to category tasks. From this balance point, one may be able to infer some properties of the categorization ecology these evolutionary lineages generally experienced.

Case 1: The Family-Resemblance Categorization Task

Some categories are organized by family resemblance (e.g., Rosch & Mervis, 1975, Wittgenstein, 1953). Family resemblance means that category members have variable and probabilistic similarity relationships. A pair of members will share some but not all features the shared features—will change from pair to pair. Family-resemblance categories also have a graded structure—from typical, central category members out to atypical, peripheral members. The family-resemblance organization could be important for some species as they construct the behavioral equivalence classes appropriate for their environments. For example, thinking as monkeys for a moment, some of their important categories—insects, eagles, pythons, seeds, leopards, trees, nest sites, and so forth—do seem to have a family-resemblance structure. The family-resemblance task has been a dominant task in the research literatures on human and animal categorization for decades (Homa et al., 1981; Knowlton & Squire, 1993; Nosofsky & Zaki, 1998; Posner, Goldsmith, & Welton, 1967; Posner & Keele, 1968, 1970; Rosch & Mervis, 1975; Smith, 2002; Smith & Minda, 2001, 2002; Smith, Redford, & Haas, 2008). Figure 1A shows a two-dimensional family-resemblance task. Category A and B members, respectively, lie to the bottom-left and upper-right of the stimulus space. These categories are essentially linearly separable. One could draw a line through the stimulus space that would nearly partition the categories.

For the simulations, I generated an 8-dimensional FR category task as follows. Five hundred Category A training exemplars were generated from the objective category prototype 100 100 100 100 100 100 100 100. Each stimulus was created by subjecting the 8 coordinates to a Gaussian transformation with standard deviation (SD) = 10. This approach reflected the common idea that the expression of traits is approximately Gaussian given multiple determinants. Five hundred Category B training exemplars were generated similarly from the objective category prototype 120 120 120 120 120 120 120 120. The Category A and B exemplars, respectively, would lie in the bottom-left and upper-right of an 8-dimensional

stimulus space. Ashby (1992) and Ashby & Maddox (1990, 1992) endorsed the approach of using multivariate normal variation to build category structures that might be natural and ecological because they show continuous variation along their dimensions, show a unimodal distribution of values along their dimensions, and have indistinct category boundaries. In fact, this approach was central to the development of *General Recognition Theory* (Ashby & Gott, 1988).

Exemplar processors in the simulations were given these 1,000 exemplars, stored whole and veridically. They could compare transfer items to the stored exemplars to evaluate category membership. Given 500 stored exemplars per category, exemplar processors had a rich basis for generalization. Prototype processors in the simulations were given the 8-dimensional central tendencies of the 500 Category A exemplars and 500 Category B exemplars. They could compare transfer items to these two subjective prototypes to evaluate category membership. Given 500-exemplar averages for each category, prototype processors also had a rich basis for generalization.

In turn, I generated 500 Category A transfer items from the objective prototype 100 100 100 100 100 100 100 100, and 500 Category B transfer items from the objective prototype 120 120 120 120 120 120 120 120. I estimated the proportion correct that exemplar and prototype processors would achieve on the transfer items given their different category representations.

To make these comparisons, I incorporated exemplar and prototype models that were very carefully matched. They used the same distance inputs and the same similarity calculations, so that they differed only in their representational assumption. Therefore, they were perfect controls for one another. The two models had the same number of formal parameters, so that they were approximately equally powerful and agile mathematically. This is important because exemplar models are sometimes granted additional parameters to the point that they become mathematical abstractions and cease being psychological models (e.g., Ashby & Alfonso-Reese, 1995)¹. All in all, the exemplar and prototype models used here represented one of the most important minimal pairs in the categorization literature (see Pothos & Wills, 2011, for a description of these and other categorization models).

Within the exemplar model, performance on each transfer item was estimated as follows. First, I calculated the stimulus distance between the transfer item and each Category A exemplar and Category B exemplar. Distance was defined as the 8-dimensional Pythagorean distance between two stimuli. For example, the distance between the stimuli 100 100 100 100 100 100 100 100 and 120 120 120 120 120 120 120 120 was $\sqrt{8 \times 400} = 56.57$.

¹In particular, the gamma parameter should not be added to exemplar models to increase their mathematical power. Gamma (derived by Ashby & Maddox, 1993, to serve a specific mathematical purpose) was incorporated by exemplar theorists into their models as part of the process by which the defining assumptions of exemplar theory were relaxed to keep the exemplar model competitive. By these changes, the exemplar model began to lose its theoretical clarity (Maddox & Ashby, 1998). The problem is that gamma can enormously increase the exemplar model's mathematical complexity (Myung, Pitt, & Navarro, 2007) so that it verges on being the powerful classifier commonly used by professional statisticians (Ashby & Alfonso-Reese, 1995). Thus the model becomes a mathematical abstraction that cannot be used to differentiate psychological theories. Gamma's purpose was to sharpen categorization profiles, producing higher highs and lower lows, but it does so for mathematical reasons that do not relate to human or animal psychology (Smith, Redford, & Haas, 2008). In addition, by this sharpening, gamma contradicts the fundamental character of exemplar-based geometries (Smith, 2002; Smith & Minda, 2001, 2002): exemplar categorization profiles are inherently homogeneous because of the softening tendencies of generalization to multiple exemplars. It has even been shown that gamma can let exemplar processes mimic prototype processes (Smith & Minda, 1998). This representational entanglement is a serious problem when one is trying to differentiate these psychological processes.

This assessment leads me to emphasize that this article is not about criticizing any theoretical approach. Exemplar theory as originally conceived is elegant and fascinating, and this article clearly grants that theory its domain of efficiency and optimality. This article is about considering prototype and exemplar processes from a perspective *above them both*, to analyze when they both may be adaptive. Yet it is important that processing models be instantiated in transparent, intuitive, and psychological ways, and gamma misses these marks.

Then, each distance (d) was transformed into a measure of psychological similarity (sim) by taking $sim = e^{-sens \times d}$. This transformation—using a sensitivity parameter ($sens$)—incorporated the common assumption that psychological similarity is an exponentially decaying function of increasing stimulus distance. The similarities of the transfer item to all 500 Category A exemplars, and separately to all 500 Category B exemplars, were summed. The estimated performance of the simulated observer for each transfer item was finally given by:

$$P(R_{CatA}|S_i) = \frac{sim_A}{sim_A + sim_B}$$

This Category A response proportion directly gave the proportion correct the simulated observer would receive for a Category A transfer item, or the probability of correctly endorsing the Category A transfer item into Category A. That response proportion gave the proportional error rate for a Category B transfer item. For Category B items, the proportion correct was given by $(1 - R_{CatA})$. These procedures were characteristic of exemplar modeling as usually practiced.

The performance estimates made by the prototype model were nearly equivalent, except that initially I calculated the distance of each transfer item from the Category A and Category B subjective prototypes. These distances were transformed to similarities and then to estimates of correct category responses as already described. These procedures were characteristic of prototype modeling as usually practiced.

Results

I analyzed 100 runs of the simulation. In each run, a new set of 1,000 training exemplars and 1,000 transfer items was produced using the techniques already described. Across the 100 runs, the prototype and exemplar processes, respectively, achieved performance levels of .8414 ($SD = 0.0033$) and .8001 ($SD = 0.0036$), respectively, $t(99) = 305.2$, $p = 2.5 \times 10^{-149}$. The prototype-process advantage was present in all 100 runs, with scant variation, and averaged .0413. Figure 2A shows for a representative run the frequency with which transfer items were categorized correctly at different levels. The prototype process produced higher performance more frequently.

I tested the generality of this finding in several ways. First (Figure 2B), I varied the value of the standard-deviation value—governing the within-category spread of exemplars—from 0.8333 to 20.8333 in 25 Steps ($SD = \text{Step}/1.2$). The value used in the representative simulation runs was 10.0. Figure 3A shows the stimulus space of the categories at Step=1, to illustrate (compare Figure 1A) the range of category coherence that was achieved by varying the standard deviation in this way. Second (Figure 2C), I made sensitivity vary from 0.01 up to 0.25 in 25 Steps ($sens = \text{Step}/100$), to explore the prototype-process advantage at different levels of underlying sensitivity for the categorizer. The value used in the representative simulation runs was 0.05. Third (Figure 2D), I made category separation vary in 25 Steps from 2 at Step 1 (eight objective prototype coordinates set at 109 and at 111) up to 50 at Step 25 (eight objective prototype coordinates set at 85 and at 135). The value used in the representative simulation runs was 20. Figure 3B shows the stimulus space of the categories at Step=25, to illustrate (compare Figure 1A) the range of separation between the categories that was achieved by varying the category separation in this way. One sees in Figures 1A, 3A, and 3B that the analyses in the present article were broad-based and inclusive for considering FR categories across a wide range of category separations, or degree/strength of family resemblance.

Across these 75 simulation runs, there was always a performance advantage to the prototype process and never an exemplar-process advantage. The prototype advantage did grow small at the extremes of the ranges tested, where the performances approached floor or ceiling.

Limited exemplar experience

Organisms learning categories initially have limited exemplar experience. This raises the question of whether an organism should abstract the prototype for an FR category based on limited exemplars. Premature prototype estimation might be poor and unwise. Instead, the organism would need to wait to amass exemplar experience, storing exemplars meanwhile. Then, the organism would have to decide when to transition to the prototype. Thus, this transition would be messy, in the sense of being temporally dicey, decisionally complex, and adaptively indeterminate (with the prototype advantage only emerging through time).

The situation would be different, though, if the prototype advantage emerged immediately. Then there would be no waiting period before abstraction was viable and no decisional tipping point. FR categorization could become seamless in a representational and processing sense, because the organism could always be abstracting, always building the best category average possible.

To consider this case, I replicated the basic simulation 500 more times, but this time I made prototype and exemplar processes dependent on Category A and B exemplar sets containing only 1, 2, and up to 500 training exemplars. Both categorizers were given the same amount of category information packaged differently—as an estimated prototype based on N exemplars or as N individuated exemplars.

Figure 4A shows the performance of both processes given 1 to 500 exemplars. The prototype advantage rose quickly to the level already reported. Figure 4B zooms in on the performance of both processes given 1 to 20 experienced exemplars. Prototype formation is immediately viable as an approach to category representation in FR category tasks. Even using the average of just 2 exemplars as a standard in categorization is better than using 2 exemplars separately as standards. The organism need not defer until a tipping point of exemplar experience is reached. It needs no decisional transition between strategies. The best approach in FR categorization tasks—from the beginning—is to form the best running average of exemplar experience one can. That strategy is applicable throughout the course of experience and learning.

This result also addresses the potential misconception that exemplar and prototype processes should converge as the exemplar set increases toward infinity. This is definitely not the case. Comparing to-be-categorized items to exemplars scattered in psychological space—even infinite numbers of exemplars—is inherently different (for mathematical and geometrical reasons discussed shortly) from comparing those items to the unitary, central prototype of the category's psychological space.

Additional Modeling Perspectives

To increase the generalizeability of the claim about FR categories and the relative advantage of prototype and exemplar processes applied to them, I tested variants on the models that made different assumptions about attention and about the metric of psychological distance. First, I considered the unequal allocation of attention. So, I let the attention to the 8 stimulus dimensions vary over 4 orders of magnitude from 0.04 up to 0.32—respectively, stimulus dimensions barely attended and barely affecting categorization and stimulus dimension sharply attended and dominantly affecting categorization. The attentional weights across the 8 stimulus dimensions still summed to 1.0, presumed to be the limit of attention's capacity.

Despite attention's allocation, there was still a robust prototype-process advantage. Across 100 runs of this selective-attention simulation, prototype and exemplar processes, respectively, achieved proportions correct of .8486 and .8066 ($t[99]=256.1$, $p = 8.2 \times 10^{-142}$).

Second, I restricted the breadth of attention until the models verged on using unidimensional rules. So, I modeled prototype- and exemplar-based performances when attention was restricted to 8, 4, 2, or only 1 stimulus dimension of the eight. Unattended dimensions made no contribution to the categorization decision. Despite attention's narrowing, there was still a robust prototype-process advantage, though narrowed attention reduced the amount of category-relevant information available and reduced overall performance. Across 400 runs of this attention-narrowing simulation (100 runs with 8-, 4-, 2, and 1-dimensional attention), prototype categorizers achieved performance levels of .8416, .7641, .6979, and .6415, respectively. Exemplar categorizers achieved performance levels of .8001, .7236, .6602, and .6070.

Third, the simulations already reported had used the Euclidean psychological-distance metric, appropriate if the stimulus dimensions were integral (i.e., selectively attended with difficulty). Now, I extended the simulation to include the City-Block metric, appropriate for separable stimulus dimensions (i.e., selectively attended with ease). So, now I calculated interstimulus distance simply by summing up the eight dimensional distances between stimuli. Across 100 runs of this modeling system, prototype and exemplar processes, respectively, achieved performance levels of .8597 and .8084.

Fourth, the simulations already reported had used a multiplicative similarity calculation, in which the distance between a to-be-categorized item and a category representation (either prototype or exemplar) is converted into a psychological similarity scaled from 0 to 1.0 using an exponential-decay function governed by a sensitivity parameter. Another possibility was to incorporate the additive calculation, in which the exemplar- or prototype-based distances lead directly to an estimate of performance level with no intervening similarity calculation. Accordingly, I instantiated an additive modeling system. Across 100 runs of this modeling system, the prototype and exemplar processes, respectively, achieved performance levels of .7162 and .6489, respectively.

These diverse simulations show that the advantage accruing to the prototype process in FR categorization persists across different assumptions about attention and psychological distance. Now I will explain why—statistically and geometrically—the present results are fundamental given FR categories organized around prototypes.

Discussion

It has been debated in cognitive science whether categorizers do sometimes use prototypes in FR category tasks (e.g., Nosofsky & Zaki, 1998; Smith & Minda, 2001). Here I addressed the qualitatively different question of whether FR categorizers might be advantaged by using prototypes. Simulations showed that it was advantageous for organisms to blend their exemplar experience into the category prototype and use this representation as the standard for classifying transfer items. Remember that advantage here means improved levels of performance, not improved model-fitting closeness.

There are converging explanations for this advantage. First, FR category members are variations on the prototype theme. Therefore, a test item will differ from the prototype in one set of ways (i.e., in its non-prototypical idiosyncrasies). However, crucially, it will differ from stored exemplars in two sets of ways (its own idiosyncrasies and the exemplars' different idiosyncrasies). Thus, a test item will commonly be farther from the category's

exemplars than from the prototype, making the exemplars a weaker basis for generalization and transfer.

In addition, exemplars contain atypical features that are irrelevant to or misleading about the category decision. Then, the evidence base for a category endorsement is undercut because, when the exemplars are retrieved, the irrelevant features are retrieved, too, potentially misdirecting the category decision. In a sense, an item-to-exemplar comparison compares the item to a signal-plus-noise bundle. The item, even if a category member, will not match the noise and it will therefore relate less strongly to the exemplar. In contrast, the prototype contains no extraneous features, and so the signal supporting correct categorization can be stronger and clearer as reflected in Figure 2A's histogram. In this case, the item is compared to pure signal.

Now one might suppose that the comparisons to the signal-noise bundles would work out because one is comparing to many stored exemplars. Wouldn't the noise components cancel out and even out in the end? This is a crucial potential misunderstanding to fend off. The exemplar model includes the combined signal-noise for each stored exemplar every time. Each of these comparisons thus produces a weaker similarity signal. Then the model averages over these weaker similarity signals, but at this point it is too late to remove the noise components, and the average similarity signal will be weaker, too. In contrast, the prototype model drops away the noise components as the prototype is stored. Then, to-be-categorized items are compared to a signalful, noiseless category representation. Now, the noise has been removed from the comparison (except the idiosyncratic information in the to-be-categorized item), and the similarity signal can be stronger.

Briscoe and Feldman (2011) gave a related perspective on the prototype-process advantage. They discussed how the exemplar model may sometimes overcommit to exemplar experience. It may store information in too much detail and at the wrong grain for generalization to new tokens. Their perspective is similar to the signal-detection perspective given earlier on exemplar strategies. Note that the need for generalization to novel tokens is frequently required in nature, as when, for example, consumed prey items have no way to appear again as exemplars in the environment. If a predator's prey-selection system seized too narrowly on one phenotype (e.g., the seasonal phenotype of white hares in winter), this could present serious transfer problems.

Geometrically, one sees in Figure 1A that the exemplars of an FR category are spread out in psychological space around the prototype. A to-be-categorized item cannot be close to all of a category's stored exemplars at once. It will always be close to some but far from others. Its exemplar-based category-belongingness signal will arise from a mix of near and far comparisons (or strong and weak similarities). This mix will produce a middling category signal that cannot get very strong no matter how close the item is to the center of the category. In contrast, a test item can get indefinitely close to the prototype because geometrically the prototype lies at an approachable, singular place in the category space (the center). Thus, the test item can be close to all the category information, and the categorization signal can grow indefinitely strong. Systematic psychophysical-scaling research spanning 35 years has confirmed these descriptions of the geometrical properties of exemplar-based and prototype-based stimulus comparisons (Posner et al., 1967; Smith, 2002; Smith & Minda, 2001, 2002).

Accordingly, one sees that the prototype advantage for FR categories is deeply grounded in the geometry of stimulus spaces and the mathematics of prototype and exemplar comparisons. The models only let the geometry flow through up to the level of predicted performance, using their transformations of stimulus differences into distances, distances

into similarities, and similarities into predicted performance levels. The models do not alter that geometry or change its consequences

The prototype advantage for FR categories would apply not only in simulations, but also as organisms solve FR category problems in their natural ecologies. The prototype advantage would be inherently, latently there—as an invariant or affordance in Gibson's (1979) terminology—for organisms prepared to take advantage of it. They would only need a cognitive mechanism for averaging over exemplar experience to grant themselves a stronger signal for FR categorization. Below I will consider whether organisms do have this mechanism. For now, the simulations suggest why they might, *depending on how prevalent the family-resemblance principle is in the natural world.*

Case 2: The Prototype-Exception Categorization Task

However, family resemblance is just one organizing principle for categories. Categories can also include exception items, logically disjoint category sub-clusters, and other elements that introduce nonlinearity into categorization. Accordingly, I consider now the nonlinearity created by giving categories perceptually misleading exception items. The prototype-exception category structure has also been a dominant task in the recent research literature (Blair & Homa, 2001; Cook & Smith, 2006; Nosofsky & Johansen, 2000; Smith & Minda, 1998; Smith et al., 2010; Smith et al., 1997). It has been important for asking how particularly participants can remember individual exemplars and make correct but non-intuitive categorization responses toward them. It has figured prominently in the prototype-exemplar debate and in the debate over whether categorizers operate with a linear-separability constraint. Figure 1B depicts a 2-dimensional prototype-exception category structure. Many Category A members lie to the bottom-left of the stimulus space. But some lie confuseably in the upper right of the stimulus space and deceptively seem to be Category B members. Category B members are positioned analogously. These categories are not linearly separable. One could not draw a line through the stimulus space that would adequately partition the members of the two categories.

In the prototype-exception simulations, the 500 Category A training exemplars and transfer items comprised 333 stimuli generated from the category sub-prototype 100 100 100 100 100 100 100 100 and 167 stimuli generated from the category sub-prototype 120 120 120 120 120 120 120 120. The Category B stimuli were distributed oppositely. The transfer items were generated identically. Exemplar processors used the 1,000 training exemplars as the comparative standards for categorization. Prototype processors used the 8-dimensional averages of the 500 Category A exemplars and 500 Category B exemplars. These averages were displaced by the exception items, so that the 8 coordinates for the subjective Category A and Category B prototypes were around 107 and 113, respectively. Now one can evaluate the performance effect of making comparisons back to exceptional exemplars or back to prototypes that incorporate the information from exceptional exemplars. The performance levels achieved by both processes were found using the equivalent exemplar and prototype models that were already described.

Results

Across 100 comparisons of the models, the prototype and exemplar processes, respectively, achieved performance levels of .6497 ($SD = 0.0024$) and .5539 ($SD = 0.0040$), respectively, $t(99) = 216.6$, $p = 1.3 \times 10^{-134}$. The prototype-process advantage was always present and averaged .0958. Figure 5A shows in one representative run the frequency with which transfer items were classified correctly at different levels. The prototype process produces highly accurate categorization for the typical Category A transfer items, but fails spectacularly for the less frequent Category A exception items. The exemplar process

produces less-good performance and less-bad performance. Again the prototype process produced a stronger category-membership signal, here for the better (typical items) or worse (exception items).

I tested the breadth of this prototype advantage by varying over 25 steps the standard-deviation value that controlled the within-category coherence (Figure 5B), the sensitivity of the simulated categorizers (Figure 5C), and the separation between categories in the stimulus space (Figure 5D). Across these 75 simulation runs, there was always a prototype-process-advantage and never an exemplar-process advantage.

Discussion

Prototype-exception categories have been an important realization of nonlinearity in category tasks. One would have supposed that the prototype process would be disadvantaged in a task that presents exceptions that the prototype process is bound to miscategorize. One would have supposed that the exemplar process would be advantaged, because it would have stored exception items that it could refer back to in categorizing transfer exception items. Yet, not only was there still a pervasive prototype advantage in this case, the advantage was larger than for well-behaved FR categories.

To explain this phenomenon I conducted additional analyses to “x-ray” the exemplar process. Its similarity comparison is exponential. Similarity decays rapidly with increasing stimulus distance. The close proximity between a transfer item and a stored representation can produce a similarity number so large that that comparison alone can capture and determine the category decision, outweighing other exemplars that make far smaller contributions to overall similarity. Therefore, the categorization decision can be derailed—for example—when a typical Category A transfer item has close proximity to a stored Category B exception. The prototype process is not subject to these decisional captures, because there are no comparisons to single stored items that could capture the decision. To the contrary, the prototypes, though sullied by the exceptions, are stable, similarity-comparison anchors. The simulations in Case 2 widen the adaptive sphere of the prototype-processing advantage.

Case 3: The Exclusive-Or Categorization Task

Toward finding the limit on the prototype-process advantage, Case 3 considers exclusive-or categories that comprise logically disjoint category sub-groups that no prototype can encompass. In Figure 1C, Category A and B members lie to the bottom-left and upper-right. Category B members lie to the bottom-right and upper-left. Neither category is perceptually coherent. No straight line through the stimulus space could partition these categories. They are profoundly linearly nonseparable. XOR categories have been crucial to the debates about prototypes and exemplars, linear-separability constraints in categorization, and rule-based systems of categorization that step in when categories lack perceptual coherence (Feldman, 2000; Nosofsky, Gluck, Palmeri, & McKinley, 1994; Shepard, Hovland, & Jenkins, 1961; Smith, Coutinho, & Couchman, 2011; Smith et al., 2004). XOR categories may provide the strongest evidence for exemplar-based categorization processes (Medin, Altom, Edelson, & Freko, 1982).

In the XOR simulations, the 500 Category A training exemplars and transfer items comprised 250 stimuli each generated from the category sub-prototypes 100 100 100 100 100 100 100 100 and 120 120 120 120 120 120 120 120. The Category B stimuli were generated from the category sub-prototypes 100 100 100 100 120 120 120 120 and 120 120 120 100 100 100 100. The transfer items were generated identically. Exemplar processors were given the 1,000 training exemplars as the comparative standards for

categorization. Prototype processors were given the 8-dimensional average of the 500 Category A exemplars and 500 Category B exemplars. The performance levels for both processes were found using the equivalent exemplar and prototype models already described. Now, one can evaluate the effect of making comparisons back to exemplars that lie in disparate sub-clusters within the stimulus space, or back to prototypes that have averaged across exemplars lying in disparate subclusters.

Results

Across 100 comparisons of the models, the prototype and exemplar processes, respectively, achieved performance levels of .5000 ($SD = 0.0011$) and .8620 ($SD = 0.0055$), respectively, $t(99) = -638.6$, $p = 4.6 \times 10^{-181}$. The exemplar advantage was present in all 100 runs and averaged .3620. Figure 6A shows in one representative run the frequency with which transfer items were classified correctly at different levels. The prototype process produced performance that hovered near chance levels.

I tested the breadth of this exemplar advantage by varying over 25 steps the standard-deviation value that controlled the within-category coherence (Figure 6B), the sensitivity of the simulated categorizers (Figure 6C), and the separation between categories in the stimulus space (Figure 6D). Across these 75 simulation runs, there was always an exemplar-process-advantage and never a prototype-process advantage.

Discussion

The exclusive-or simulations were important for several reasons. The exclusive-or task, so influential in the literature, was naturally included in this survey of important category structures. It was also important to explore the limits of the prototype advantage. It was important to show that the processing advantage can shift radically depending on the category problem presented. It was valuable, as a matter of theoretical balance, to present a complementary case (favoring the exemplar process) to the family-resemblance case (favoring the prototype process).

Nonetheless, the explanation for the results of the exclusive-or simulations is transparent. Averaging over the two subclusters in each category, the Category A and Category B subjective prototypes were almost identical and coincident in perceptual space. The prototype process had no representational basis for categorizing differentially because its two category representations did not differ. The predicted categorization proportion— $\text{Sim A} / (\text{Sim A} + \text{Sim B})$ —was always about 0.5.

Indeed, one might think this simulation was circular, because it took exemplar-privileged categories and confirmed the privilege. If this were just about a modeler choosing a category structure and a matching model by artifice, it would be circular. But the issue is deeper than that. If the natural world were suffused with exclusive-or categories, and if species had responded to this fact with exemplar-emphasizing paths of cognitive evolution that let them learn adaptively, then that tailoring of mind to natural structure would not be artifice and it would not be circular. Thus, the exclusive-or simulations, like the family-resemblance simulations earlier, could provide an insight about the emergence of cognitive systems for categorization, *depending on how prevalent the exclusive-or principle is within the natural world*.

Once again it is important to see that this article is not just another round in the prototype-exemplar debate. Its theoretical goals transcend that debate. Here I am discussing the tuning of systems for categorization to the structure of the natural world—whatever the directionality of that tuning. Tuning is an equally important and interesting theoretical

possibility whether exemplar processes or prototype processes are favored by that tuning, and the exclusive-or simulation makes this point clearly.

Interim Summary

This article surveys influential category structures from the literature, by evaluating the relative optimality of prototype and exemplar processes applied to them. The family-resemblance structure—across a range of within-category similarity, observer sensitivity, and between-category separation—afforded about a 4% prototype performance advantage. There was never an exemplar performance advantage.

The 4% performance advantage is not small. This improved foraging efficiency could mean 120 additional mosquitoes nightly for an insectivorous bat or 34 additional nectar flowers daily for a honeybee. Moreover, it would need only a small increase in overall fitness for prototype abstraction to have become widespread in extant species, given the long and patient time-frame of cognitive evolution. Fifty percent of mortality among vervet monkeys is caused by predation (e.g., eagles, leopards, pythons, baboons—all family-resemblance categories). Increasing predator avoidance and decreasing the false alarms that disrupt foraging would make life safer and more efficient. Indeed, predator categories are so crucial that vervets have developed calls to communicate about them (Cheney & Seyfarth, 1990). It makes sense that they also would have evolved to represent these categories optimally.

This possibility gains in plausibility given the article's simulations with the prototype-exception category structure. Now there was a larger prototype advantage. There was never an exemplar advantage, though this task presented some nonlinearity that could have favored the exemplar process.

The article's simulations with the exclusive-or category structure finally eliminated the prototype advantage and produced instead a profoundly large exemplar advantage.

Given these results, one can frame important questions about the structure of natural categories and the evolution of cognitive systems for categorization. The principal question is about the relative frequencies of family-resemblance, prototype-exception, exclusive-or, and related kinds of category structures among biological kinds and other natural-kind categories.

One possible answer is that something like exclusive-or nonlinearity dominates the natural categorization ecology. Then, exemplar processing would be strongly advantageous and there would have been an urgent adaptive pressure toward the inclusion of that process within animals' categorization systems. It might take only the occurrence of occasional exclusive-or categories in nature to justify a significant cognitive investment in a non-linear categorization capability, because the large performance advantage each time would provide a significant fitness boost.

Another possible answer is that family-resemblance dominates the natural categorization ecology. Then, prototype processes would be inherently advantageous in foraging and predator avoidance. This might have created an adaptive pressure toward the inclusion of prototype processes within animals' categorization systems. This adaptive pressure could operate even if the categorization ecology included some nonlinearity. For example, prototype-exception categories still afford a prototype-process advantage.

Categorization by Nonhuman Primates

To what extent have animals developed exemplar processes as part of their categorization system, or invested in prototype processing? Research is beginning to answer these questions. In fact, there is an empirical confluence between studies of human and animal categorization. We know now how members of one nonhuman species process the three tasks that have influenced the human literature and for this reason were featured in this article's simulations. Animals' data patterns reveal their operating principles and default commitments in categorization, and perhaps also the structure of the natural ecology within which their categorization capacity evolved.

Performance in family-resemblance tasks

Smith, Redford, and Haas (2008) gave rhesus macaques (*Macaca mulatta*) the FR category structure that was the focus of this article's first simulations. They used the dot-distortion task that is commonly used with humans (Homa et al., 1981; Knowlton & Squire, 1993; Smith & Minda, 2001, 2002; Nosofsky & Zaki, 1998; Posner & Keele, 1968, 1970).

Macaques showed strong performance from the beginning of training—near 90% correct on typical category members. They showed steep typicality gradients—that is, they endorsed items into the trained category much less strongly as those items had less strong similarity relationships within the category. These results suggest that FR tasks are natural and approachable for macaques.

Figure 7 shows these typicality gradients, and the ability of prototype and exemplar processes to reproduce them. The prototype model's predictions were so precise that its P symbols are difficult to discern. In contrast, the exemplar model mispredicted performance on all item types and its indices of misfit were large. Smith et al. tested many variants of the exemplar model but none of them let exemplar processes adequately reproduce the performance data. These results imply that macaques were averaging over their exemplar experience, to create for themselves a centroid within the exemplar space—a prototype of the dot-distortion category. They were then referring to-be-categorized items to this unitary, central prototype, instead of referring to-be-categorized items back to previously experienced whole, stored exemplars.

It is theoretically important that macaques brought to the FR task the prototype process that the present simulations confirm is advantageous, and not exemplar processing. Although the dot-distortion task used by Smith et al. is uniquely diagnostic of prototype processing, this conclusion dovetails with other findings of prototype-enhancement effects in animal categorization (Aydin & Pearce, 1994; Huber & Lenz, 1993; Jitsumori, 1996; von Fersen & Lea, 1990; White, Alsop, & Williams, 1993).

Performance in prototype-exception tasks

Smith et al. (2010) gave macaques the prototype-exception task featured in this article's simulations and common in human studies (e.g., Blair & Homa, 2001; Cook & Smith, 2006; Smith & Minda, 1998; Smith et al., 1997). Each category had a majority membership of 6 typical items with a strong FR structure and 2 exception items derived from the opposing category. The macaques completed more than 160,000 trials trying to master prototype-exception categories.

Figure 8 shows their proportion correct on the typical (T) and exception (E) items by 64-trial block. Again, macaques categorized items related by family resemblance easily and naturally. They reached high performance levels on typical items. However, they exhibited remarkably weak and insensitive exemplar processing as reflected in their learning of

exceptions. One macaque never learned the exceptions, even after 12,000 trials in each of five prototype-exception tasks. Across these tasks, he received about 15,500 exception trials, and 8,800 additional correction trials for errors made on those trials, to no avail. It would be hard for an animal to exhibit a weaker capacity for exemplar processing, for this animal was dealing with four specific exception items at a time that repeated hundreds of times. The other macaques remained below chance on exceptions for thousands of trials. One animal produced the important result that he did systematically worse on exceptions the better he did on typical items—he was clearly responding to the FR gradients in the task. Nonetheless, these two macaques did both finally reach 70–80% correct performance on exceptions, confirming that they had some capacity to finally treat individual category exceptions particularly.

These macaques revealed basic aspects of their category-learning systems. They made a strong assumption about family-resemblance structure entering these tasks, systematically misclassifying exceptions. They showed a strong linear-separability constraint in categorization that made it difficult to overcome this assumption and learn exceptions. This difficulty betrayed a weak and insensitive exemplar process, not the robust exemplar process one might see if nonlinearity characterized the categorization ecology. In short, here, too, macaques' categorization performance was shifted toward the family-resemblance, prototype-processing pole of functioning.

Performance in exclusive-or tasks

Smith et al. (2011) gave 3 macaques the exclusive-or task that was the focus of this article's third set of simulations. Each macaque learned 2-, 3-, and 4-dimensional exclusive-or categories for 5,760 trials. In each task, two dimensions instantiated the exclusive-or relation, and the remaining dimension(s) varied randomly, carrying no useful category information.

Macaques found the exclusive-or problems very difficult to learn. One must attribute this to the structure of the XOR categories, not to the unnaturalness of the stimulus domain. Previous research had confirmed that the stimulus materials used in these tasks were very approachable for macaques—they sometimes learned FR tasks constructed from these materials in less than 100 trials. However, for the exclusive-or categories, they had 22%, 32%, and 38% error rates in the 2-, 3-, and 4-dimensional tasks, respectively. They received 3,731, 5,452, and 6,608 20-s penalties. Even their terminal performance levels—.830, .764, and .656—were modest, though there were only 4, 8, or 16 stimuli to learn to classify, and though those stimuli repeated hundreds or thousands of times.

Here, too, macaques revealed a severe linear-separability constraint in the category structures they find easily learnable. The common assumption would be that exclusive-or performance is supported by exemplar processes. Adopting this assumption, one would conclude that macaques only bring insensitive exemplar processes to exclusive-or category learning. Clearly, they show nothing of the sophisticated exemplar processing one might expect to see if nonlinearity characterized the categorization ecology.

To be clear, the macaque results do evaluate negatively the power of macaques' exemplar processing in category tasks. However, this is neither a general rejection of exemplar processes in category learning (that might occur robustly in other species or in other situations), nor a critique of exemplar theory (that has its place within the multiple-systems theoretical perspective on categorization). Nonetheless, the weakness of macaques' exemplar process, their linearity-separability constraint, and their family-resemblance default assumption will be dominant factors in understanding the overall categorization competence of nonhuman primates.

Categorization by Humans

Parallel research has shown how humans process the category structures featured in this article's simulations. Their data patterns may also reveal their operating principles and their default assumptions.

Family-resemblance category tasks

Humans, like macaques, learn family-resemblance tasks easily with few training trials (e.g., Knowlton & Squire, 1993; Reber, Stark, & Squire, 1998a,b; Smith & Minda, 2001, 2002). Moreover, they show the steep typicality gradients that macaques do and that are diagnostic of prototype processing (Smith, 2002). Figure 9 illustrates these typicality gradients. It also shows the success and failure, respectively, of the prototype (P) and exemplar (E) process in reproducing these gradients.

Though humans could apply a robust exemplar process to these tasks, they do not. The prototype process has precedence. This does not imply that the prototype process always has precedence for humans. Given different category structures that lack a strong FR structure, other processes might take precedence.

Prototype-exception category tasks

Humans also reveal the precedence of prototypes and a linear-separability constraint in prototype-exception tasks. Figure 10 shows humans' correct classification on typical (T) and exception (E) items (Smith et al., 2010). Humans initially misclassified the exception items systematically. Humans did not improve their exception-item performance at all until completing 128 trials. They were unable to respond above chance on the exception items until after 304 trials. In all, the 40 humans made 3,488 (44%) errors in category assignment on exception items, despite the deliberate correction procedure that followed each error. Thus, humans showed continuities with macaques in their prototype-exception performance. In another well-known example (Smith & Minda, 1998), participants in prototype-exception tasks demonstrated large prototype-enhancement effects of about 15%. They also performed far below chance on the exception items in the category, making errors on these items about 75% of the time. The prototype model captured this data pattern; the exemplar model could not. Blair and Homa (2001) contributed a strongly convergent finding. Pothos, Perlman, et al. (2011) also showed that humans strongly prefer to spontaneously form linearly separable groupings or classifications.

The precedence of prototype processing, and the linear-separability constraint in categorization, may extend back hundreds of millions of years in vertebrate evolution. Cook and Smith (2006) placed humans and pigeons in a prototype-exception task. Figure 11 shows the performance of both species during the early and later phases of category training. Early on, humans and pigeons (black circles) showed strong prototype-enhancement effects for prototype items and below-chance exception-item performance. A prototype model (squares) captured this data pattern; an exemplar model (triangles) could not—it underpredicted prototype performance and overpredicted exception-item performance. Later on, one sees that the data pattern for both species became friendlier to the exemplar perspective, a hint that humans and pigeons have processes waiting in the wings that let them manage nonlinear category structures.

Nonetheless, the early learning by both species is consistent only with a default commitment to family resemblance and linear separability as an organizing principle for categories. Remarkably, this default commitment evidently spans 150 million years of vertebrate evolution, and it could suggest that cognitive systems for categorization have become adaptively tuned to the general properties of categories in the natural world.

Exclusive-or category tasks

Finally, Smith et al. (2011) gave humans the same 2-, 3-, and 4-dimensional exclusive-or tasks they gave macaques. In these three tasks, respectively, 2, 8, and 38 participants failed ever to learn the categories to criterion. Human learners needed 74, 178, and 236 trials to reach criterion, and they made 21, 59, and 89 errors before reaching criterion. Humans, like macaques, lack a robust exemplar process that can quickly resolve the category sub-clusters of items.

More evidence supporting humans' weak exemplar processing comes from humans' performance in other poorly structured category tasks. It is seldom appreciated how badly humans perform in tasks that afford exemplar processing but not prototype processing. As examples, 30%, 36%, 40%, 60%, 66%, and 72% of participants failed to reach criterion in various experiments in Medin and Schwanenflugel (1981), Medin, Dewey, and Murphy (1983), Medin and Smith (1981), and Medin and Schaffer (1978). Moreover, even learners had an asymptotic performance of about 80%. This pervasive difficulty is probably only explained through the hypothesis that these tasks make humans operate outside their natural, family-resemblance, linearly separable element, requiring them to bring to bear other processes (e.g., exemplar processes) that they can only recruit weakly and insensitively.

Perhaps the weakness of humans' exemplar system, and their lack of tuning to the exemplar-process pole of categorization, are not surprising because the fundamental roots of human categorization lie in primate categorization. However, in the present theoretical climate, this weakness is an important theoretical conclusion with many implications.

Conclusion

Over the past three decades, the categorization literature contained a dominant theoretical narrative. It emphasized humans' flexibility as categorizers: their ability to provide category structure from their own minds, their independence from category structure dictated perceptually, the flexible, knowledge-based cores of categories, the ad hoc, occasional nature of categories, the importance of correlational information, and humans' comfort with nonlinear categories. This narrative included the suppositions that FR categories are not especially natural, not psychologically privileged, that organisms do not especially value coherent pools of perceptually similar items, and that animals' and humans' systems for category learning do not make an FR assumption or suffer a linear-separability constraint (e.g., Medin & Schaffer, 1978; Medin & Schwanenflugel, 1981; Medin, Wattenmaker, & Hampson, 1987; Nosofsky, 1986, 1992; Nosofsky, Little, Donkin, & Fific, 2011; Nosofsky & Johansen, 2000; Stanton & Nosofsky, 2007).

This article's simulations, and the lines of research just summarized, suggest that it is time to consider a complementary narrative as follows.

The natural ecology may be suffused with FR categories organized around coherent similarity relationships. Ashby (1992, p. 454) endorsed the idea that "the structure of many natural categories can be effectively modeled by the multivariate normal distribution." That is, many natural categories show continuous variation along their dimensions, have a symmetrical and unimodal distribution of values along those dimensions, are describable as being centered around the prototypical centroid of the category, have potential overlap with neighboring categories, and thus have indistinct category boundaries (also Ashby & Maddox, 1990, 1992).

Fried and Holyoak (1986, p. 235; also Flannagan, Fried, & Holyoak, 1984) suggested that normal distributions have particular ecological importance, especially given that many

“basic-level categories seem to consist of a dense central region of typical instances, surrounded by sparser regions of atypical instances.” “People may therefore expect new categories to be unimodal and to have roughly symmetrical density functions, which may be well approximated by multidimensional normal distributions.” This expectation—by humans, macaques and pigeons—has been confirmed (Blair & Homa, 2001; Cook & Smith, 2006; Smith & Minda, 1998; Smith et al., 1997). Palmeri & Nosofsky (2001, p. 198) also pointed out that FR categories may mimic the structure of many natural categories. Black (1954, p. 28) argued that “if we examine instances of the application of any biological term, we shall find ranges, not classes—specimens (i.e., individuals or species) arranged according to the degree of their variation from certain typical or ‘clear’ cases.”

Rosch is well known to have suggested that many natural-kind categories, especially those at the basic level, are overlapping bundles of probabilistic features, organized by family resemblance, with typicality gradients running from central, prototypical items out to peripheral, boundary items. Her extensive research also demonstrated that humans are responsive to these aspects of FR categories: in their naming, in their verification of category membership, in their typicality ratings, and in the developmental course of category learning (Mervis & Rosch, 1981; Rosch, 1973, 1975; Rosch & Mervis, 1975; Rosch, Mervis, Gray, Johnson, & Boyes-Braem, 1976; Rosch, Simpson, & Miller, 1975).

Ethnobiologists also emphasize that many biological kinds—especially taxa near the specificity level of the genus—represent natural groupings of organisms with salient clusters of correlated, probabilistic attributes that are separated in feature space from other groupings, just as instantiated in the present simulations. These biological categories are highly salient, highly privileged, and highly consensual among humans, across cultural groups and even across folk and formal taxonomic systems (e.g., Berlin, Breedlove, & Raven, 1973; Boster, 1987; Bulmer, 1970; Hunn, 1975; Waddy, 1988; see Malt, 1995, for an extraordinary review of this extensive literature).

Therefore, contrary to the dominant theoretical narrative in the categorization literature, there may well be a privileged, FR category structure underlying natural kinds.

In that case, FR categories may have played a significant role in categorization’s evolution. The potential advantage from prototypes, demonstrated in this article’s first and second group of simulations, has been present in FR category problems through the hundreds of millions of years that animals have confronted these problems. There has likely been a phylogeny-long, gentle pressure exerted toward the evolution in many species of a prototype-abstraction capability, so that prototype abstraction has been a common solution to the problem of FR categorization. This could be an example of an evolutionary tailoring phenomenon by which animals become more cognitively adept within their natural world (see also Shepard, 1987, 1994, 2001).

Ashby and Alfonso-Reese (1995) provided an additional reason why cognitive evolution might have progressed toward prototype processing. They analyzed categorization as a problem of estimating the statistical distribution of the category’s exemplars. The hardest aspect of discovering category structure may lie in discovering the character of its underlying exemplar distribution. Given FR structure that kept repeating in nature, organisms would be advantaged if they became parametric categorizers that assume the structure of the next category. In this context, it is interesting that pigeons, macaques, and humans do make this kind of family-resemblance, linear-separability presumption entering category tasks. That assumption would leave the observer only needing to learn the central tendency and dimensional variances organizing the new category—an easier learning goal than discovering the next qualitatively unknown statistical distribution. Ashby and Alfonso-

Reese (1995) further suggested that organisms might be tuned in this way because natural categories do have a repeating FR structure. Thus organisms would reap the performance benefits described in this article.

Briscoe and Feldman (2011, also Feldman, 2003) took a similar approach. They described exemplar processes as high-variance--low-bias processes. They can embody very complex hypotheses, allowing exemplar models to fit disparate data patterns. However, exemplar processes may generalize poorly, because they overcommit to specifics based on the input patterns given them, and thus they may be poorly equipped to categorize novel tokens. In contrast, Briscoe and Feldman described prototype models as low-variance--high-bias processes. They entertain simple, inflexible hypotheses about a common conceptual form (the prototype). But these models may nonetheless generalize well if the structure they commit to is actually present within the category. Then, once again, they would reap the performance benefits described here.

Briscoe and Feldman (2011) share a good deal with Ashby and Alfonso Reese (1995) in their description of the statistical commitments and performance affordances of the two processes. They share a good deal with this article in evaluating empirically how humans are tuned for categorization, and how humans should be tuned. Between the lines in all three articles lie important questions about the evolutionary emergence of cognitive systems for categorization in the vertebrate line over half a billion years—its natural context, the pressures exerted by that context, the adaptive tuning caused by those pressures. These questions point to a broad area that deserves far more theoretical attention within categorization science.

In contrast to FR categories, the natural ecology appears to present few category problems that embody the exclusive-or relation or related kinds of nonlinearity. In considering the daily rounds of the primate species that I know the most about, I have not found any such category problems in their foraging or predator-avoidance lives (e.g., Cheney & Seyfarth, 1990). That is, it seemingly never occurs that a preferred food source is dimorphic (producing palatable orange, three-lobed berries and palatable blue, four-lobed berries) whereas a to-be-avoided food source is exclusively-or dimorphic (producing toxic orange, four-lobed berries and toxic blue, three-lobed berries). It is a graceful state of nature that the genetics of biological kinds—fruits, mates, and predators—creates well-behaved family-resemblance categories, not exclusive-or double-crosses. Even in cases of species mimicry, the mimicry clearly lets one species benefit from another's linear separability constraint. It is adaptive to perceptually resemble the bark of a tree, because you will be less likely to be seen as a member of a prey-species category. It is adaptive to perceptually resemble a toxic moth or poisonous snake because it allows inclusion within an FR category that predators will avoid.

Consequently, there has been scant pressure toward the development of categorization processes that robustly address nonlinearities in categorization, and in particular no large investment in the development of sensitive exemplar processes that could naturally serve that function. These aspects of animals' and humans' categorization systems remain weak and slow to learn, even on the extreme repetition of few exceptional items. Categorization in many vertebrate species is probably not tuned toward the non-linear, exemplar-processing pole of functioning in category-learning tasks. I would suggest this is because the natural ecology has not presented these kinds of category structure, or pressed for systems that manage these category structures well, or tuned organisms to manage those structures well.

An insightful reviewer's perspective strengthens this conclusion. He noted that logically, there are far more nonlinearly separable category structures that would favor exemplar

processes than there are linearly separable category structures that would favor prototype processing. But the more this is true, the more telling it is that humans and animals are so poor with the exemplar strategy. For this suggests that the reviewer's important logical point does not extend to the structure of natural kinds. It suggests that human and animal forebears did not have deep experience with the broad range of nonlinear structures—despite their logical dominance—and that instead they had deep experience with the narrow range of prototype-privileged category structures—perhaps because these are naturally prevalent though not logically dominant. In contrast, recent categorization research has focused selectively on nonlinear category structures that would favor exemplar processing (discussion in Feldman, 2003, 2004; Smith & Minda, 2000; Smith, Murray, & Minda, 1997). This divergence of the research laboratory from the natural world emphasizes once again the importance of the present article's evolutionary and fitness perspective.

In the end, the recent dominant theoretical narrative in human categorization research is undermined by an intuitive look at the natural world, by the macaque evidence, by the human evidence, even by the pigeon evidence, and by the evidence of formal simulations that show that prototypes would have a valuable role to fill in many cognitive systems. I recommend that the field give equal theoretical attention to the complementary theoretical narrative just outlined that has many strengths. It is grounded in the categorization ecologies that animals may have experienced. It has phylogenetic breadth and evolutionary depth. It explains many aspects of categorization performance by humans, nonhuman primates, and probably vertebrates more broadly. It reveals important continuities in the categorization systems of humans and nonhumans. It accomplishes the constructive goal of integrating synergistically the human and animal research literatures on categorization. Finally, it points theory toward the important possibility that cognitive systems for categorization have been shaped and molded by the structure of the natural kinds in the natural world.

Acknowledgments

The preparation of this article was supported by Grant HD-060563 from the *National Institute of Child Health and Human Development*.

References

- Ashby, FG. Multidimensional models of categorization. In: Ashby, FG., editor. *Multidimensional models of perception and cognition*. Hillsdale, NJ: Erlbaum; 1992. p. 449-483.
- Ashby FG, Alfonso-Reese LA. Categorization as probability density estimation. *Journal of Mathematical Psychology*. 1995; 39:216-233.
- Ashby FG, Alfonso-Reese LA, Turken AU, Waldron EM. A neuropsychological theory of multiple systems in category learning. *Psychological Review*. 1998; 105:442-481. [PubMed: 9697427]
- Ashby FG, Gott RE. Decision rules in the perception and categorization of multidimensional stimuli. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. 1988; 14:33-53.
- Ashby FG, Maddox WT. Integrating information from separable psychological dimensions. *Journal of Experimental Psychology: Human Perception and Performance*. 1990; 16:598-612. [PubMed: 2144574]
- Ashby FG, Maddox WT. Complex decision rules in categorization: Contrasting novice and experienced performance. *Journal of Experimental Psychology: Human Perception and Performance*. 1992; 18:50-71.
- Ashby FG, Maddox WT. Human category learning. *Annual Review of Psychology*. 2005; 56:149-178.
- Ashby FG, Maddox WT. Human category learning 2.0. *Annals of the New York Academy of Sciences*. 2010; 1224:147-161. [PubMed: 21182535]
- Aydin A, Pearce JM. Prototype effects in categorization by pigeons. *Journal of Experimental Psychology: Animal Behavior Processes*. 1994; 20:264-277.

- Berlin B, Breedlove DE, Raven PH. General principles of classification and nomenclature in folk biology. *American Anthropology*. 1973; 75:214–242.
- Black, M. Problems of analysis. Ithaca, NY: Cornell University Press; 1954. (Collected Essays)
- Blair M, Homa D. Expanding the search for a linear separability constraint on category learning. *Memory & Cognition*. 2001; 29:1153–1164. [PubMed: 11913752]
- Blair M, Homa D. As easy to memorize as they are to classify: The 5-4 categories and the category advantage. *Memory & Cognition*. 2003; 31:1293–1301. [PubMed: 15058690]
- Boster J. Agreement between biological classification systems is not dependent on cultural transmission. *American Anthropologist*. 1987; 89:914–919.
- Briscoe E, Feldman J. Conceptual complexity and the bias/variance tradeoff. *Cognition*. 2011; 118:2–16. [PubMed: 21112048]
- Brooks, LR. Nonanalytic concept formation and memory for instances. In: Rosch, E.; Lloyd, BB., editors. *Cognition and categorization*. Hillsdale, NJ: Erlbaum; 1978. p. 169–211.
- Bulmer, RNH. Which came first, the chicken or the egghead?. In: Pouillon, J.; Maranda, P., editors. *Échanges et communications: Mélange offerts à Claude Lévi-Strauss à l'occasion de son 60ème anniversaire*. Vol. 2. The Hague: Mouton; 1970. p. 1069–1091.
- Cheney, DL.; Seyfarth, RM. How monkeys see the world. Chicago: University of Chicago Press; 1990.
- Cook R, Smith JD. Stages of abstraction and exemplar memorization in pigeons' category learning. *Psychological Science*. 2006; 17:1059–1067. [PubMed: 17201788]
- Couchman JJ, Coutinho MVC, Smith JD. Rules and resemblance: Their changing balance in the category learning of humans (*Homo sapiens*) and monkeys (*Macaca mulatta*). *Journal of Experimental Psychology: Animal Behavior Processes*. 2010; 36:172–183. [PubMed: 20384398]
- Erickson MA, Kruschke JK. Rules and exemplars in category learning. *Journal of Experimental Psychology: General*. 1998; 127:107–140. [PubMed: 9622910]
- Feldman J. Minimization of Boolean complexity in human concept learning. *Nature*. 2000; 407:630–633. [PubMed: 11034211]
- Feldman J. The simplicity principle in human category learning. *Current Directions in Psychological Science*. 2003; 12:227–232.
- Feldman J. How surprising is a simple pattern? Quantifying “Eureka! *Cognition*. 2004; 93:199–224.
- Flannagan MJ, Fried LS, Holyoak KJ. Distributional expectations and the induction of category structure. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. 1986; 12:241–256.
- Fried LS, Holyoak KJ. Induction of category distributions: A framework for classification learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. 1984; 10:234–257.
- Gibson, JJ. The ecological approach to visual perception. Boston, MA: Houghton Mifflin; 1979.
- Hampton, JA. Concepts in human adults. In: Mareschal, D.; Quinn, P.; Lea, SEG., editors. *The making of human concepts*. Oxford: Oxford University Press; 2010. p. 293–311.
- Homa D, Sterling S, Trepel L. Limitations of exemplar-based generalization and the abstraction of categorical information. *Journal of Experimental Psychology: Human Learning and Memory*. 1981; 7:418–439.
- Huber L, Lenz R. A test of the linear feature model of polymorphous concept discrimination with pigeons. *Quarterly Journal of Experimental Psychology*. 1993; 46B:1–18.
- Hunn E. A measure of the degree of correspondence of folk to scientific biological classification. *American Ethnologist*. 1975; 3:309–327.
- Jitsumori M. A prototype effect and categorization of artificial polymorphous stimuli in pigeons. *Journal of Experimental Psychology: Animal Behavior Processes*. 1996; 22:405–419. [PubMed: 8865609]
- Kemler Nelson DG. The effect of intention on what concepts are acquired. *Journal of Verbal Learning and Verbal Behavior*. 1984; 23:734–759.
- Knowlton BJ, Squire LR. The learning of categories: Parallel brain systems for item memory and category knowledge. *Science*. 1993; 262:1747–1749. [PubMed: 8259522]

- Lazareva, OF.; Wasserman, EA. Category learning and concept learning in birds. In: Mareschal, D.; Quinn, PC.; Lea, SEG., editors. *The making of human concepts*. Oxford: Oxford University Press; 2010.
- Lea SEG, Wills AJ. Use of multiple dimensions in learned discriminations. *Comparative Cognition and Behavior Reviews*. 2008; 3:115–133.
- Malt BC. Category coherence in cross-cultural perspective. *Cognitive Psychology*. 1995; 29:85–148.
- Medin DL, Altom MW, Edelson SM, Freko D. Correlated symptoms and simulated medical classification. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. 1982; 8:37–50.
- Medin DL, Dewey GI, Murphy TD. Relationships between item and category learning: Evidence that categorization is not automatic. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. 1983; 9:607–625.
- Medin DL, Schaffer MM. Context theory of classification learning. *Psychological Review*. 1978; 85:207–238.
- Medin DL, Schwanenflugel PJ. Linear separability in human classification learning. *Journal of Experimental Psychology: Human Learning & Memory*. 1981; 7:355–368.
- Medin DL, Smith EE. Strategies and classification learning. *Journal of Experimental Psychology: Human Learning and Memory*. 1981; 7:241–253.
- Medin DL, Wattenmaker WD, Hampson SE. Family resemblance, conceptual cohesiveness, and category construction. *Cognitive Psychology*. 1987; 19:242–279. [PubMed: 3581758]
- Mervis CB, Rosch E. Categorization of natural objects. *Annual review of psychology*. 1981; 82:89–115.
- Minda JP, Smith JD. Prototypes in category learning: The effects of category size, category structure, and stimulus complexity. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. 2001; 27:775–799.
- Murphy, GL. *The big book of concepts*. Cambridge, MA: MIT Press; 2003.
- Myung JI, Pitt MA, Navarro DJ. Does response scaling cause the generalized context model to mimic a prototype model. *Psychonomic Bulletin & Review*. 2007; 14:1043–1050. [PubMed: 18229473]
- Nosofsky RM. Attention, similarity, and the identification-categorization relationship. *Journal of Experimental Psychology: General*. 1986; 115:39–57. [PubMed: 2937873]
- Nosofsky RM. Attention and learning processes in the identification and categorization of integral stimuli. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. 1987; 13:87–108.
- Nosofsky, RM. Exemplar-based approach to relating categorization, identification, and recognition. In: Ashby, FG., editor. *Multidimensional models of perception and cognition*. Hillsdale, NJ, England: Lawrence Erlbaum Associates, Inc; 1992. p. 363–393. *Scientific psychology series*
- Nosofsky RM, Gluck MA, Palmeri TJ, McKinley SC. Comparing models of rule-based classification learning: A replication and extension of Shepard, Hovland, and Jenkins (1961). *Memory & Cognition*. 1994; 22:352–369. [PubMed: 8007837]
- Nosofsky RM, Johansen MK. Exemplar-based accounts of “multiple-system” phenomena in perceptual categorization. *Psychonomic Bulletin and Review*. 2000; 7:375–402. [PubMed: 11082849]
- Nosofsky RM, Little DR, Donkin C, Fific M. Short-term memory scanning viewed as exemplar-based categorization. *Psychological Review*. 2011; 118:280–315. [PubMed: 21355662]
- Nosofsky RM, Zaki SR. Dissociations between categorization and recognition memory in amnesic and normal individuals: An exemplar-based interpretation. *Psychological Science*. 1998; 9:247–255.
- Palmeri TJ, Nosofsky RM. Central tendencies, extreme points, and prototype enhancement effects in ill-defined perceptual categorization. *Quarterly Journal of Experimental Psychology*. 2001; 54A: 197–235. [PubMed: 11216316]
- Pearce, JM. Discrimination and categorization. In: Mackintosh, NJ., editor. *Animal learning and cognition*. New York: Academic Press; 1994. p. 109–134.
- Posner MI, Goldsmith R, Welton KE. Perceived distance and the classification of distorted patterns. *Journal of Experimental Psychology*. 1967; 73:28–38. [PubMed: 6047706]

- Posner MI, Keele SW. On the genesis of abstract ideas. *Journal of Experimental Psychology*. 1968; 77:353–363. [PubMed: 5665566]
- Posner MI, Keele SW. Retention of abstract ideas. *Journal of Experimental Psychology*. 1970; 83:304–308.
- Pothos, EM.; Wills, AJ. Formal approaches to categorization. Cambridge, UK: Cambridge University Press; 2011.
- Pothos EM, Perlman A, Bailey TM, Kurtz K, Edwards DJ, Hines P, McDonnell JV. Measuring category intuitiveness in unconstrained categorization tasks. *Cognition*. 2011; 121:83–100. [PubMed: 21733497]
- Reber PJ, Stark CEL, Squire LR. Cortical areas supporting category learning identified using functional MRI. *Proceedings of the National Academy of Science*. 1998a; 95:747–750.
- Reber PJ, Stark CEL, Squire LR. Contrasting cortical activity associated with category memory and recognition memory. *Learning and Memory*. 1998b; 5:420–428. [PubMed: 10489259]
- Reed SK. Pattern recognition and categorization. *Cognitive Psychology*. 1972; 3:382–407.
- Reed SK. Category vs. item learning: implications for categorization models. *Memory & Cognition*. 1978; 6:612–621. [PubMed: 723619]
- Rosch, E. On the internal structure of perceptual and semantic categories. In: Moore, TE., editor. *Cognitive development and the acquisition of language*. New York: Academic Press; 1973. p. 111–144.
- Rosch E. Cognitive representations of semantic categories. *Journal of Experimental Psychology: General*. 1975; 104:192–233.
- Rosch E, Mervis CB. Family resemblances: Studies in the internal structure of categories. *Cognitive Psychology*. 1975; 7:573–605.
- Rosch E, Mervis CB, Gray WD, Johnson DM, Boyes-Braem P. Basic objects in natural categories. *Cognitive Psychology*. 1976; 8:382–439.
- Rosch E, Simpson C, Miller RS. Structural bases of typicality effects. *Journal of Experimental Psychology: Human Perception and Performance*. 1976; 2:491–502.
- Rosseel Y. Mixture models of categorization. *Journal of Mathematical Psychology*. 2002; 46:178–210.
- Shepard RN. Toward a universal law of generalization for psychological science. *Science*. Sep; 1987 237(4820):1317–1323. [PubMed: 3629243]
- Shepard RN. Perceptual-cognitive universals as reflections of the world. *Psychonomic Bulletin & Review*. 1994; 1:2–28. [PubMed: 24203412]
- Shepard RN. Perceptual-cognitive universals as reflections of the world. *Behavioral and Brain Sciences*. 2001; 24:581–601. [PubMed: 12048942]
- Shepard RN, Hovland CI, Jenkins HM. Learning and memorization of classifications. *Journal of Experimental Psychology*. 1961; 65:94–102. [PubMed: 13988437]
- Smith JD. Exemplar theory's predicted typicality gradient can be tested and disconfirmed. *Psychological Science*. 2002; 13:437–442. [PubMed: 12219810]
- Smith JD, Ashby FG, Berg ME, Murphy MS, Spiering B, Cook RG, Grace RC. Pigeons' categorization may be exclusively nonanalytic. *Psychonomic Bulletin and Review*. 2011; 18:422–428. [PubMed: 21327381]
- Smith JD, Beran MJ, Crossley MJ, Boomer J, Ashby FG. Implicit and explicit category learning by macaques (*Macaca mulatta*) and humans (*Homo sapiens*). *Journal of Experimental Psychology: Animal Behavior Processes*. 2010; 36:54–65. [PubMed: 20141317]
- Smith JD, Chapman WP, Redford JS. Stages of category learning in monkeys (*Macaca mulatta*) and humans (*Homo sapiens*). *Journal of Experimental Psychology: Animal Behavior Processes*. 2010; 36:39–53. [PubMed: 20141316]
- Smith JD, Coutinho MVC, Couchman JJ. The learning of exclusive-or categories by monkeys (*Macaca mulatta*) and humans (*Homo sapiens*). *Journal of Experimental Psychology: Animal Behavior Processes*. 2011; 37:20–29. [PubMed: 20718556]
- Smith JD, Crossley MJ, Boomer J, Church BA, Beran MJ, Ashby FG. Implicit and explicit category learning by capuchin monkeys (*Cebus apella*). *Journal of Comparative Psychology*. 126:294–304. In press. [PubMed: 22023264]

- Smith JD, Minda JP. Prototypes in the mist: the early epochs of category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. 1998; 24:1411–1436.
- Smith JD, Minda JP. Thirty categorization results in search of a model. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. 2000; 26:3–27.
- Smith JD, Minda JP. Journey to the center of the category: The dissociation in amnesia between categorization and recognition. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. 2001; 27:984–1002.
- Smith JD, Minda JP. Distinguishing prototype-based and exemplar-based processes in category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. 2002; 28:800–811.
- Smith JD, Minda JP, Washburn DA. Category learning in rhesus monkeys: A study of the Shepard, Hovland, and Jenkins tasks. *Journal of Experimental Psychology: General*. 2004; 133:398–414. [PubMed: 15355146]
- Smith JD, Murray MJ, Minda JP. Straight talk about linear separability. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. 1997; 23:659–680.
- Smith JD, Redford JS, Haas SM. Prototype abstraction by monkeys (*Macaca mulatta*). *Journal of Experimental Psychology: General*. 2008; 137:390–401. [PubMed: 18473665]
- Stanton RD, Nosofsky RM. Feedback interference and dissociations of classification: Evidence against the multiple-learning-systems hypothesis. *Memory & Cognition*. 2007; 35:1747–1758. [PubMed: 18062551]
- Vauclair, J. Categorization and conceptional behavior in nonhuman primates. In: Bekoff, M.; Allen, C., editors. *The cognitive animal: Empirical and theoretical perspectives on animal cognition*. Cambridge, MA: MIT Press; 2002. p. 239–245.
- von Fersen L, Lea SEG. Category discrimination by pigeons using five polymorphous features. *Journal of the Experimental Analysis of Behavior*. 1990; 54:69–84. [PubMed: 16812617]
- Waddy, J. Classification of plants and animals from a Groote Eylandt aboriginal point of view. Vol. 1. Darwin, Australia: Australian National University North Australia Research Unit; 1988.
- Wasserman EA, Kiedinger RE, Bhatt RS. Conceptual behavior in pigeons: categories, subcategories, and pseudocategories. *Journal of Experimental Psychology: Animal Behavior Processes*. 1988; 14:235–246.
- White KG, Alsop B, Williams L. Prototype identification and categorization of incomplete figures by pigeons. *Behavioural Processes*. 1993; 30:253–258.
- Wittgenstein, L. *Philosophical investigations*. Oxford, UK: Blackwell; 1953.

A. Family-Resemblance Category Structure **B. Prototype-Exception Category Structure** **C. Exclusive-Or Category Structure**

Figure 1.
A–C. Three category structures.

Figure 2.

A. The frequency with which the prototype and exemplar process (filled and open circles, respectively) categorized transfer items at different accuracy levels in a family-resemblance task. B–D. The overall accuracy of the prototype and exemplar process for different levels of category structure (B), observer sensitivity (C), and category separation (D).

Figure 3.

A. A family-resemblance category structure with very low within-category inter-stimulus distance and thus very high within-category interstimulus similarity. *B.* A family-resemblance category structure with very large between-category inter-stimulus distance and thus very low between-category interstimulus similarity.

Figure 4.

A–B. The overall accuracy of the prototype and exemplar process for different numbers of experienced exemplars (*A*: 1 to 500; *B*: 1 to 20).

Figure 5.

A. The frequency with which the prototype and exemplar process (filled and open circles, respectively) categorized transfer items at different accuracy levels in a prototype-exception category task. B–D. The overall accuracy of the prototype and exemplar process for different levels of category structure (B), observer sensitivity (C), and category separation (D).

Figure 6.

A. The frequency with which the prototype and exemplar process (filled and open circles, respectively) categorized transfer items at different accuracy levels in an exclusive-or category task. B–D. The overall accuracy of the prototype and exemplar process for different levels of sub-category structure (B), observer sensitivity (C), and sub-category separation (D).

Figure 7.

A–B. Two macaques' performance (black circles) in 10 family-resemblance category-learning cycles (Experiment 1, Smith et al., 2008). The measure Endorsement Level expresses how often different item types were accepted (endorsed) into the training category. Also shown is the average of the 10 best-fitting predicted profiles when the exemplar model (E) and prototype model (P) fit each macaque's 10 data sets individually. From "Prototype abstraction by monkeys (*Macaca mulatta*)," by J. D. Smith, J. S. Redford, and S. M. Haas, 2008, *Journal of Experimental Psychology: General*, 127, p. 395. Copyright 2008 by the American Psychological Association. Reprinted with permission.

Figure 8.

A–C. Three macaques' performance in the prototype-exception category task of Smith et al. (2010). Curves T and E, respectively, show the proportion of correct responses made to the six typical items and two exception items in each training category.

Figure 9.

Humans' performance (black circles) in four family-resemblance category-learning studies (controls in Knowlton & Squire, 1993, people with amnesia in Knowlton & Squire, 1993, and the participants in Reber et al., 1998a,b). The measure *Endorsement Level* expresses how often different item types were accepted (endorsed) into the training category. Also shown is the average of the four best-fitting predicted profiles when the exemplar model (E) and prototype model (P) fit individually these four data sets. From "Prototype abstraction by monkeys (*Macaca mulatta*)," by J. D. Smith, J. S. Redford, and S. M. Haas, 2008, *Journal of Experimental Psychology: General*, 127, p. 392. Copyright 2008 by the American Psychological Association. Reprinted with permission.

Figure 10. Humans' performance in the prototype-exception category task of Smith et al. (2010, Experiment 1B). Curves T and E, respectively, show the proportion of correct responses made to the six typical items and two exception items in each training category.

Figure 11.

Humans' and pigeons' observed and predicted accuracy for prototypes, typical items, and exceptions during the early (top panels) and late (bottom panels) stages of learning within the prototype-exception task of Cook and Smith (2006). Observed performances are shown as unconnected, filled, black circles. The best-fitting predictions of prototype and exemplar models are shown, respectively, as connected open squares and open triangles. From "Stages of abstraction and exemplar memorization in pigeon category learning," by R. G. Cook and J. D. Smith, 2006, *Psychological Science*, 17, p. 1063. Copyright 2006 by the Association for Psychological Science. Reprinted with permission.