Translating From Perceptual to Cognitive Coding

Tyler Davis (thdavis@mail.utexas.edu)
Bradley C. Love (brad_love@mail.utexas.edu)
W. Todd Maddox (maddox@psy.utexas.edu)
Department of Psychology, The University of Texas at Austin
Austin, TX 78712 USA

Abstract

An important question in cognitive science is how conceptual knowledge interacts with perception. For example, how does knowing that trucks have wheels and windshields impact a person's recognition of a truck? Through a series of categorization tasks, we explore two pathways by which humans can learn to identify objects. One pathway involves the acquisition of perceptual expertise through extensive task-relevant experience, and leads to the rapid identification of individual stimuli. A second pathway involves a translation process in which percepts of objects are encoded into symbolic descriptions. The use of this pathway results in a similar ability to identify individual stimuli, but requires more time to operate due to the demands of the online translation process. These results are discussed in terms of related work in categorization, perceptual expertise, and influence of language on thought.

Keywords: expertise; object identification; symbolic coding; translation; categorization; language and thought

Introduction

One of the most fundamental psychological capacities is the ability to recognize sameness. In categorization and object recognition this capacity is typically studied in terms of the formation and use of equivalence classes, however there are numerous occasions where a person's goals require going beyond category membership to the level of individual identity. Whether it be taking the right car home from work, or the right dog home from the dog park, many times getting things exact makes all of the difference. This ability to recognize that something is the same as or equivalent to something seen on a previous occasion, termed reidentification (cf. Millikan, 1998), is fundamental to our everyday survival. The present study investigates two capacities, perceptual expertise and translation that interact to allow humans to maximize their ability to track objects across time.

Perceptual expertise refers to a variety of representational changes that occur due to experience with a particular object or within a particular domain. These changes alter the way in which people process the objects for which they have expertise such that they become as efficient at recognizing them and as knowledgeable about their unique characteristics at subordinate levels (e.g. as an individual or a species) as they are at naming and discussing them as members of a basic level category (Tanaka & Taylor, 1991).

However, perceptual expertise is acquired rather slowly, with the representations formed being largely configural, and thus specific to the metric relationships between features typical to a certain class of objects or a particular individual (Gauthier & Tarr, 1997; Tanaka & Gauthier, 1997). While these representational changes have been shown to facilitate the learning of new similar classes or new individuals within a class (Gauthier, et al., 1998; Tanaka, Curran, & Sheinberg, 2005), they likely do not generalize beyond a given basic level category. Thus, perceptual expertise is extremely useful for developing an automatic, stable, and permanent reidentification capacity for a class or an individual object, but it is too slow and specific to subserve all of the occasions where people need to be able to engage in subordinate level identification.

Another capacity, which may bootstrap reidentification performance while perceptual expertise is developing, is a process we will refer to as translation. Translation will be defined as the mapping of sensory information onto a compositional alphabet of symbols that refer to sensory qualities or featural characteristics of objects. As with perceptual expertise, we argue that many of these symbols are learned, but instead of being fused configural entities, they refer to independent features, and can be applied across a wider variety of domains. These symbols may be acquired from a variety of sources, most particularly through the learning of a language (e.g. Quine, 1960), explicit instruction (Experiment 2), or the learning of task specific invariants (e.g Schyns, Goldstone, & Thiabut, 1998).

For these symbols to be useful for object reidentification, a person must be able to use them systematically, and productively. That is, if an individual possesses a set of feature labels that they are able to attach to one object, they should at the very least be able to identify those qualities in any of their combinations in any other object. Another further constraint on the success of translation in a given context is that the symbols must stand in a one to one relationship with the sensory qualities to which they refer. For example, the term green may be sufficient for identifying a particular object when the only colors in the environment are green or blue, but when the colors of the objects vary more continuously the term green will quickly lose its utility.

Though we are treating translation and the development of expertise as distinct, in many cases they may be intertwined in a manner similar to the algorithmic and automatic processes that operate in tasks like mental math (e.g. Logan, 1988; see Figure 1 for summary of proposed pathways). To this end, translation may not only be well suited for bootstrapping reidentification during the early stages of perceptual expertise, but it might continue to operate

	Perceptual Expertise	Translation
Representation	**Configural/continuous**	**Featural/symbolic**
Acquisition	**Extensive direct experience**	**Culture, instruction, or feature learning**
Function	**Permanent and rapid reidentification ability**	**Bootstrapping performance in the absence of expertise, etc.**
Characteristics	**Automatic and domain specific**	**Algorithmic and domain general**

Figure 1. Proposed pathways to object identification

depending on the context in which an individual finds themselves. For example, a bird expert may recognize a golden- cheeked warbler automatically due to extensive experience, but might engage in the more algorithmically based and costly translation process when they witness an instance of golden- cheeked warbler exhibit an exceptional trait and they wish to be able to track it over future occasions. Further, even if they are confident in their abilities to track this exceptional warbler, they may nonetheless wish to translate it so that they might be able to convince a skeptical comrade of its existence with a verbal description of its warbler and non-warbler-like features.

In the present series of experiments we provide evidence for the existence of the translation process while shedding light on some of the similarities and differences it has in relation to perceptual expertise. In Experiment 1, we show that the reidentification of individual perceptual figures in the early stages of perceptual expertise depends upon how readily they can be translated. In Experiment 2A & 2B, we provide evidence for the algorithmic nature of translation by showing that learning a vocabulary by which to translate perceptual stimuli can lead to reidentification as successful as perceptual expertise, but sufficient time is necessary at encoding for this to occur. Finally, in Experiment 3, we look further into translation and perceptual expertise by investigating how readily they generalize to new categories. This results in the replication of findings from Experiment 2A & 2B as well as findings from the expertise literature which have suggested that expertise generalizes well to new items when they come from the same general class.

Experiment 1

A critical component of the translation hypothesis is that for reidentification of an object to occur in the early stages of perceptual expertise, a person must possess an adequate vocabulary upon which to map the object. To begin to test the adequacy of the translation account, we will use a perceptual classification task in which participants have to reidentify individual exemplars, and where the conditions differ in terms of how easily the stimuli are mapped onto symbolic descriptions. Rule-plus-exception designs (e.g. Nosofsky et al.,1994) are precisely the type of task where these two requirements can be met. In these tasks, the majority of stimuli can be classified according to an imperfect rule, however there are also several exception stimuli that do not follow this rule. Participants learning a category that includes exceptions must individuate these items, and reidentify them throughout the course of the task. As a result of this individuation, the exception items become more prominent in memory and are recognized to a greater extent in transfer memory tasks (e.g. old/new judgments) in comparison to rule-following items (Palmeri & Nosofsky, 1995; Sakamoto & Love, 2004). Thus, the rule-plus-exception task requires the sort of reidentification of individual objects that we wish to test, and the exception advantage effect can be used as a measure of the ability for participants to successfully individuate a particular stimulus.

Studies of perceptual classification also offer a ready solution to the manipulation of translatability. In perceptual classification tasks, stimuli typically come in one of two possible varieties that are used as if they were virtually interchangeable, discrete (usually binary), and continuous. The translation hypothesis suggests however, that these types of stimuli may be different in terms of participants' abilities to reidentify them on the individual level. That is, discrete stimuli are easily mapped onto values that stand in a one to one correspondence with the sensory qualities of the stimulus (e.g. large, red, triangle), whereas those with continuous values do not.

Experiment 1 will test whether translation is important for reidentification by using a rule-plus-exception design which forces participants to identify exceptions at the individual level, and by manipulating translatability through the use of stimuli that can be made to appear continuous or discrete depending on the availability of additional verbal cues. If translatability affects participants' abilities to reidentify objects in the absence of perceptual expertise participants in the continuous condition should perform worse in learning to classify the exception items, and show little to no advantage for them in a follow-up recognition task.

Method

Stimuli

Each stimulus was a box that varied in terms of its height and the position of a line segment along its lower boundary (see Figure 2). In the discrete version of the stimuli, tic marks and labels marked each value along both of the dimensions. The only difference between the continuous discrete stimuli was the absence of these labels in the former condition (For illustrative purposes and ease of presentation, these groups will be referred to as Tic Marks and No Tic Marks for the discrete and continuous conditions, respectively). The categories were constructed according to a rule-plus-exception design where there was an imperfect rule orthogonal to one of the dimensions, and each category contained an exception item that appeared as if it belonged to the opposite category (see Figure 3). There were 36 stimuli which formed a 6X6 grid, with each category containing 14 regular rule following members, 1 exception,

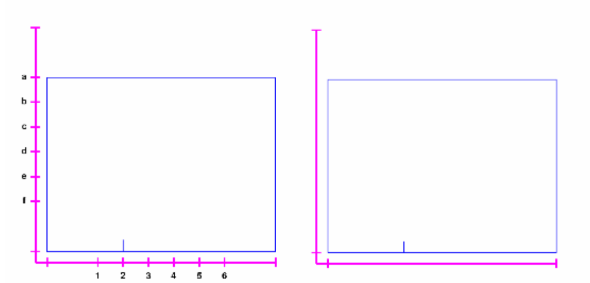

Figure 2. An example of a Tick Mark (discrete) (A) and a No Tick Mark (continuous) (B) stimulus from Experiment 1. The stimuli vary in terms of their height and the position of the line segment along their lower boundary.

1 rule-following item that matched the exception in terms frequency of presentation, and 2 foils taken out for use in a the follow-up recognition phase. The categories were set up to mirror each other with the special items placed in one of four locations such that the exception and frequency-matched rule following items were always diagonal to each other, and with the remaining two positions reserved for foils.

Procedure

Participants were randomly assigned to either the Tic Mark (discrete) or No Tic Mark (continuous) condition with the relevant dimension for the rule following items, and exception placement counter-balanced. Participants completed three blocks of classification training where they classified each regular item one time, and each of the exceptions and frequency matched rule items three times. They were given unlimited time to make each classification decision, followed by 2 seconds of feedback where the stimulus remained on the screen. Following training, participants completed a series of two-alternative forced choice (2AFC) recognition trials, where they were instructed to identify which item they had seen before, with each old item (exception or rule) paired with each foil twice.

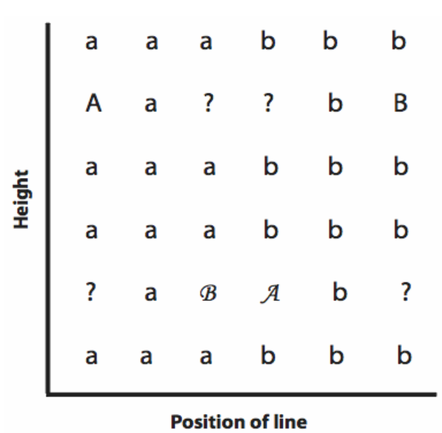

Figure 3. A sample category structure with exceptions in italics and frequency-matched rule items in capital letters. The ?'s represent foils used in the recognition phase.

Results

Classification

Performance in the classification task differed between conditions, where the Tic Mark condition performed better at classifying the exception items with a mean of .55 correct in the final block whereas the No Tick Mark condition had a mean of .33. The groups, however, did not differ to any large extent in how they classified the rule following items with a mean of .87 correct for the Tic Mark condition and .86 for the No Tic Mark condition. The overall interaction was significant, $F(1,101)=17.688$, $p<.001$, and each of the planned comparisons were as expected, with the difference in exception performance between conditions being significant $t(101)=5.895$, $p<.001$, and the difference between rule performance not $t<1$.

Recognition

Performance in the 2AFC recognition task also differed between conditions with the Tic Mark condition showing a sizeable exception advantage with a mean of .78 correct for exceptions and .63 for rule items, whereas the No Tic Mark condition did not exhibit an exception advantage with a mean of .59 for the exceptions and the .57 for the rule items (see Figure 5). This interaction was also significant $F(1,101)=6.043$, $p=.016$, and each of the planned comparisons were as expected where the Tic Mark condition showed a significant exception advantage $t(52)=4.09$, $p<.001$, and the No Tic Mark condition did not, $t<1$.

Discussion

Experiment 1 successfully demonstrated that having a vocabulary upon which to map sensory stimulations is important for tracking individual items in the absence of perceptual expertise, as predicted by the translation hypothesis. Whether an individual was able to track the exception items over the experimental session was shown to depend on whether they were able to map them onto discrete symbols.

Experiment 2A

Although Experiment 1 demonstrated that translatability is important for reidentifying items in the absence of expertise, there is still the question of how people are able to obtain the symbols that they use to translate objects. Unlike the stimuli in Experiment 1, objects in the real world don't often come with labels on them, yet people can still reidentify objects in the absence of expertise. In Experiment 2A, we investigated how the ability to translate is developed, and how it might differ from other types of expertise that do not involve learning a feature vocabulary.

Accordingly, in Experiment 2A, two groups of experts were trained such that one group was taught a vocabulary upon which to map the box stimuli from Experiment 1 prior to learning how to classify them in the rule-plus-exception task. The other group simply attained expertise within the original classification task used in Experiment 1. It was expected that both groups of experts would learn to individuate and reidentify the exception items, but the identification group was expected to be able to do this accurately without having as much experience within the classification task itself. However, since translation is hypothesized to be more algorithmic, it is possible that the ability of participants in the identification condition to reidentify the exceptions will depend on the amount of time available to encode them after feedback.

Specifically, the goal of the second experiment was to test whether learning to track the exception items was possible without any tic marks, and if it was, whether it required actual experience in classifying them, or if it could be achieved through translation.

Method

Stimuli

Stimulus design and category structure for the classification, identification, and recognition tasks were identical to Experiment 1, except that no stimuli included the tic marks or labels.

Procedure

Participants were randomly assigned to either identification (ID) or classification training. Both groups were trained in their respective condition for 14 blocks over 3 sessions on consecutive days, and an additional 3 blocks in the final session. During these training sessions, the classification training condition simply practiced the identical task that they would be tested on, with the same design as Experiment 1. The ID condition practiced naming the different stimulus dimensions according to a grid (see Figure 4), where they made an independent response for each dimension and then received feedback for the entire stimulus for 2 seconds. After the 3 additional training blocks in the final session, the testing procedure was identical to experiment 1, where participants ran the original classification task, followed by 2AFC recognition for exceptions and frequency matched rule items.

Results

Both training conditions gained a high deal of proficiency in their respective training regime by the time that they began the final day assessment. Performance was at or near ceiling for all participants and all items.

Classification

There were differences in performance between groups where the classification condition performed better at classifying exceptions than the ID condition with .93 and .53 correct respectively, but both performed similarly in terms of their rule item performance (.97 and .97). The interaction between condition and item type was marginally significant, $F(1,8)=5.236$, $p=.051$.

Recognition

There were also differences between conditions for 2AFC recognition performance, where the classification condition showed a robust exception advantage with mean .96 and .58 correct for exceptions and rule following items respectfully, but the ID group failed to show a significant exception advantage with .51 correct for exceptions and .40 for rule following items. This resulted in a substantial main effect where-by the classification outperformed the ID group in recognition overall, $F(1,8)=34.524$, $p<.001$ (see Figure 5). Planned comparisons for item and condition showed that these results were due to a significant exception advantage in the classification condition, $t(4)=3.457$, $p=.026$, but a non-significant difference between exception and rule recognition in the ID condition, $t<1$. Post-hoc tests showed that recognition of neither the exception nor the rule was significantly different from chance in the ID condition, $t<1$, and $t(4)=1.020$, $p=.365$, for the exception and rule items respectfully (see figure 5 for Experiment 2A & 2B recognition results).

Discussion

In Experiment 2A there were large differences between groups in their performance with the exception items. Whereas the classification condition performed well, participants in the ID condition showed a deficit in their ability classify and recognize exceptions correctly. These results are understandable in light of the translation hypothesis as well as predictions derived from Experiment 1. In Experiment 1, it was shown that participants do not typically encode the values of the features fully before making a classification decision, as evidenced by the fact that participants in the No Tick Mark condition were not impaired at rule-following items. Thus, where a subject is likely to resort to translation is in the event of an error, meaning they will have to sufficiently encode the values of the stimulus within the feedback time. It is a possibility then, that the feedback time used in Experiment 2A was not long enough for subjects to translate the exception stimuli.

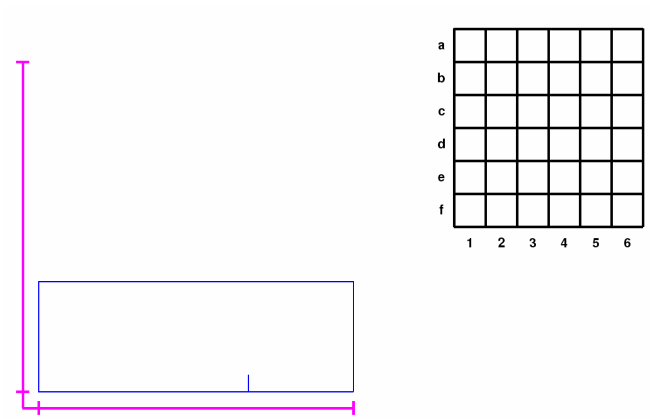

Figure 4. An illustration of the identification learning task

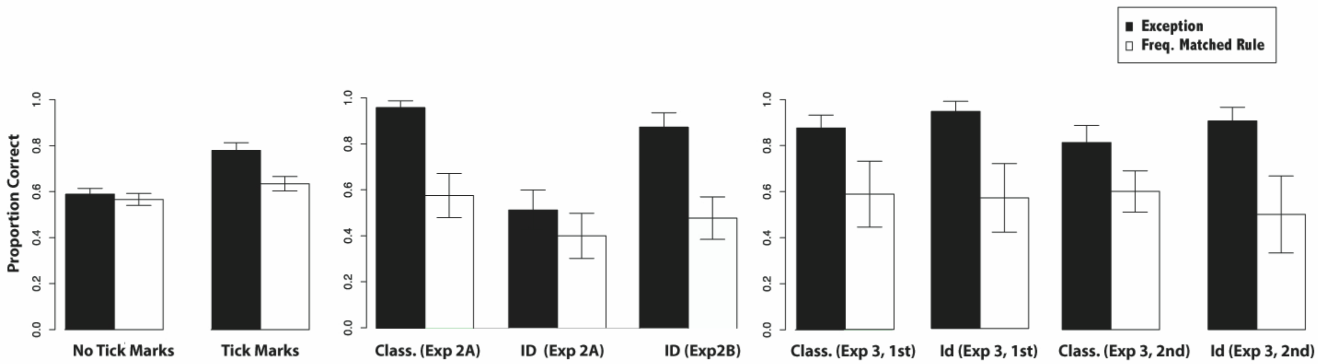

Figure 5. Recognition result Experiment 1 (left) Experiment 2&2B (center), and Experiment 3 (right).

Experiment 2B

To test whether the feedback time provided in Experiment 2A was responsible for the suboptimal performance witnessed in the ID condition, in Experiment 2B, we had the ID experts from Experiment 2A return for an identical procedure to the previous session, except that they were trained on the opposite rule dimension, and were given 4 seconds of feedback rather than 2. The expectation being that longer feedback should increase recognition performance for the exception items, but not the rule items leading to a more robust exception advantage for the identification trained participants.

Accordingly, the classification results improved to where there was no significant difference between performance with exceptions and rule following items, $t<1$. Further, exception recognition increased in 2AFC to where there was a significant exception advantage $t(4)=3.45$, $p=.026$, and in comparison to the previous session there was a significant main effect showing an increase in recognition performance between sessions $F(1,4)=9.954$ (see Figure 5). Planned comparisons revealed that this significant change in performance was likely due to an increase in exception recognition between sessions, $t(4)=3.076$, $p=.037$, and not in rule item recognition $t<1$.

Thus Experiment 2B taken with Experiment 2A demonstrates that perceptual expertise and translation are unique pathways to object identification. Whereas perceptual expertise results in the automatic recognition of objects for which an individual has developed expertise, as evidenced by the lack of impairment in the classification condition in Experiment 2A, translation's algorithmic characteristics make it such that its ability to lead to successful reidentification depends upon the amount of time available at encoding. When encoding time is not sufficient translation does not provide a benefit over novice performance.

Experiment 3

Since Experiment 2B was potentially confounded by knowledge of the experimental hypothesis, and participants' awareness of the conditions under which they would have to reidentify the exceptions, Experiment 3 was conducted to test whether the results of Experiment 2B were replicable while extending the practice of having participants transfer to a novel classification task to the classification condition.

Therefore, to overcome the possible confounds in Experiment 2B, and to provide more information on the abilities of the classification experts, in Experiment 3 we trained two groups of experts in a manner identical to Experiment 2. On the final test day, all participants received two classification tasks with a feedback time of 4 seconds followed by a 2AFC recognition phase. It was predicted that the feedback time would be sufficient for the ID participants to show an exception advantage in both tasks. In the second task it was uncertain how the classification condition would perform. Results from the expertise literature suggest that they should be able to assimilate some of the new category structure to their prior expertise, but at the same time they may be expected to exhibit a good deal of interference from their prior learning.

In correspondence with these predictions and the results of Experiments 2A and 2B, there was a significant main effect for exceptions in the first 2AFC recognition phase across groups $F(1,6)=7.229$, $p=.036$ (see Figure 5). In correspondence to the hypothesis that there would be assimilation of the new category structure to prior expertise in the classification condition, this exception advantage carried over into the second 2AFC recognition phase $F(1,6)=16.215$, $p=.007$.

General Discussion

The account of object identification described herein focuses on two ways in which people can learn to identify individual objects. The first of these is perceptual expertise, which results in rapid and automatic identification of objects for which a given individual has extensive experience. The second, translation, is an algorithmic process by which sensory stimulations can be mapped onto discrete symbols. In Experiment 1, we provided support for this process by demonstrating that translatability is crucial for performance when a stimulus must be reidentified from amongst other perceptually similar stimuli in the absence of perceptual expertise. In Experiment 2A and 2B, we showed that a vocabulary by which to translate perceptual stimuli is learnable, and can aid in individuating these particular stimuli. However, unlike the more automatic perceptual expertise developed from extensive experience within a particular domain, it requires time in order to operate, thus supporting the characterization of translation as an algorithmic process. Finally, Experiment 3 provided information on how these two learning mechanisms allow people to generalize to new training sets, replicating results from Experiment 2A & 2B, as well as the expertise literature.

There are several implications of the present research on studies of perceptual expertise and classification. First, it suggests that large differences can arise in perceptual classification tasks depending on whether the stimuli have discrete or continuous valued dimensions. As mentioned previously, these types of stimuli are often assumed to be completely interchangeable. However, since their main difference is in translatability, and this factor has presently been demonstrated to have large effects on performance, this practice should now be considered suspect.

Similarly for the domain of expertise, the present research suggests that experts may differ in important ways depending on how they interact with the stimuli for which they have expertise. Bird experts, for instance, possess not only a great deal of experience with birds, but also a large vocabulary by which to translate birds if the need arises (Tanaka & Taylor, 1991). Thus, whether an individual relies upon their perceptual expertise or translation in a particular act of identification will vary widely with the context. A more complete account of expert performance will need to develop a better understanding of how these two mechanisms interact.

Finally, the present series of experiments also has implications for the language and thought literature. For example, the present experiment suggests that where language is going to have the largest impact on recognition performance is in situations that call for translation, and less so in cases of perceptual expertise where items are represented more holistically. Correspondingly, a great deal of research in the language and thought literature has shown enhanced discrimination around linguistically relevant color boundaries (see Kay & Regier, 2006 for review). Since color terms are predominately used to represent object features, this is likely related to the normative constraints put on the use of these symbols by a culture for purposes of translation and communication about objects. Thus, the way that we represent objects in language for other people is probably strongly related to how we represent them to ourselves for purposes of reidentification.

In conclusion, the perceptual expertise/translation dichotomy is a useful way to approach object identification that may help to open a broader dialog between a wide variety of research programs in the cognitive sciences. The present research not only takes a first step in characterizing how these processes interact, but also provides a coherent methodology by which future questions regarding the interplay of language, categorization, and object recognition can be investigated.

Acknowledgments

This work was supported by AFOSR Grant FA9550-04-1-0226 and NSF Grant 0349101 to Bradley C. Love and NIH Grant R01 MH59196 to W. Todd Maddox.

References

Gauthier, I., & Tarr, M.J. (1997). Becoming a “Greeble” expert: Exploring mechanisms for face recognition. *Vision Research, 37,* 1673-1682.

Gauthier, I., Williams, P., Tarr, M.J., & Tanaka, J. (1998). Training ‘greeble’ experts: a framework for studying expert object recognition processes. *Vision Research, 38,* 2401-2428.

Kay, P. & Regier, T. (2006). Language, thought and color: recent developments. *Trends in Cognitive Science, 10,* 51-54.

Logan, G.D. (1988). Toward an instance theory of automatization. *Psychological Review, 95,* 492-527.

Millikan, R.G. (1998). A common structure for concepts of individuals, stuffs, and real kinds: More mama, more milk, and more mouse. *Behavioral and Brain Sciences, 21,* 55-100.

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

Palmeri, T.J., & Nosofsky, R.M. (1995). Recognition memory for exceptions to the category rule. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 21*, 548-569.

Quine, W.V. (1960). Word and Object. Cambridge, MA: MIT Press.

Sakamoto, Y., & Love, B.C. (2004). Schematic influences on category learning and recognition memory. *Journal of Experimental Psychology: General, 133*, 534-553.

Schyns, P.G., Goldstone, R.L., & Thiabut, J.P. (1998). The development of features in object concepts. *Behavioral and Brain Sciences, 21,* 1-54.

Tanaka, J.W., & Taylor, M. (1991). Object categories and expertise: Is the basic level in the eye of the beholder? *Cognitive Psychology, 23,* 457-482.

Tanaka, J.W., Curran, T., & Sheinberg, D.L. (2005). The training and transfer of real-world perceptual expertise. *Psychological Science, 16,* 145-151.

Tanaka, J., & Gauthier, I. (1997). Expertise in object and face recognition. *The Psychology of Learning and Motivation, 36,* 83-125.