Brain Categorization: Learning, Attention, and Consciousness

Stephen Grossberg, Gail A. Carpenter, Bilgin Ersoy Center for Adaptive Systems and Department of Cognitive and Neural Systems Boston University, 677 Beacon Street, Boston, MA 02215 USA. [steve, gail, bersoy]@cns.bu.edu http://cns.bu.edu/

Abstract – How do humans and animals learn to recognize objects and events? Two classical views are that exemplars or prototypes are learned. A hybrid view is that a mixture, called rule-plus-exceptions, is learned. None of these models learn their categories. A distributed ARTMAP neural network with selfsupervised learning incrementally learns categories that match human learning data on a class of thirty diagnostic experiments called the 5-4 category structure. Key predictions of ART models have received behavioral, neurophysiological, and anatomical support. The ART prediction about what goes wrong during amnesic learning has also been supported: A lesion in its orienting system causes a low vigilance parameter.

I. INTRODUCTION

Some scientists believe that exemplars, or individual experiences, can be learned and remembered, like those of familiar faces. Unfortunately, storing every exemplar can lead to a combinatorial explosion of memory, as well as to unwieldy memory retrieval. Others believe that we learn prototypes that represent more general properties of the environment, such as that everyone has a face. But then how do we learn specific episodic memories? Popular cognitive models of these processes also do not describe how this information is learned. This article briefly summarizes recent results showing that a variant of distributed Adaptive Resonance Theory (ART) can incrementally learn categories in a way that allows quantitative fits of human categorization data, while clarifying how both specific and general information can be incrementally learned in a context-appropriate way. The model also sheds light on amnesic categorization data.

More generally, these results support the hypothesis that brain processes underlying categorization are part of a larger system whereby the brain is designed to learn about a changing world. In particular, the processes whereby our brains continue to learn about a changing world in a stable fashion throughout life are proposed to lead to conscious experiences. These processes include the learning of bottom-up adaptive filters that activate recognition categories, the read-out of top-down expectations by these categories, the matching of these expectations against bottom-up data, the focusing of attention upon the expected clusters of information, and the development of resonant states between bottom-up and top-down processes as they reach a predictive and attentive consensus between what is expected and what is there in the outside world. It is suggested that all conscious states in the brain are resonant states, and that these resonant states trigger learning of sensory and cognitive representations when they amplify and synchronize distributed neural signals that are bound together by the resonance. Thus, processes of learning, categorization, intention, attention, synchronization, and consciousness are intimately linked. ART explicates this predicted link.

Illustrative psychophysical and neurobiological data have been explained and quantitatively simulated using these concepts in the areas of early vision, visual object recognition, auditory streaming, and speech perception, among others [1-5]. These articles summarize recent neurobiological experiments that provide convergent evidence for ART predictions, including the predicted link expectations, attention, resonant between learned synchronization, and learning, with top-down expectations computed by modulatory on-center off-surround networks that can prime the brain to get ready for bottom-up information that may or may not occur, and match or mismatch such information when it does occur, focusing attention upon patterns of critical features that match the modulatory on-center, thereby leading to synchronization and gain amplification of these features, while suppressing mismatched features.

II. UNIFYING EXEMPLARS AND PROTOTYPES USING ATTENTIONALLY MODULATED CRITICAL FEATURE PATTERNS

In the cognitive literature on recognition, and more specifically on object categorization, prototype and exemplar descriptions have lead to prominent models of the human categorization process. In the prototype-based approaches [6-9], a single center of a category is extracted from many exemplars, to-be-categorized items are compared to these category prototypes, and they are assigned to the category of the most similar prototype. The alternative exemplar-based approach [10-13] does not assume a single category center. Instead, a more distributed representation of the category domain is assumed to exist, wherein memorized sets of individual exemplars are the core representational units in memory. A new item is compared to each of the exemplars and similarity measures are obtained in terms of these comparisons.

Both of these approaches have advantages and disadvantages. Because the exemplar approach codes individual events, it is plausible that individual events, like a particular face in a particular pose, can be recognized. On the other hand, this approach raises the problem of how to recognize novel variations of familiar events; that is, where should category boundaries be drawn? Said more generally, how can one determine the proper level of abstraction when only exemplars are stored in memory? In addition, how can one search such a large memory in an efficient way? How can one avoid a combinatorial explosion as more and more exemplars are learned and searched as life proceeds? In particular, why does the reaction time to a familiar recognition event remain rapid even as the total number of exemplars that are stored in memory increases?

Because prototypes code abstractions of multiple events, prototype models plausibly explain how the learning of abstract information, such as the fact that all humans have a face, may occur. On the other hand, then one is faced with the problem of how to recognize individual events, such as the particular face of a friend. Here, too, the problem of abstraction is again raised, but from the opposite end of the concreteness-abstractness continuum.

In order to deal with these concerns, a third approach, which often is called by Nosofsky the *rule-plus-exceptions* model [14-17], attempts to incorporate the strengths of both the exemplar and prototype approaches, while overcoming their most obvious weaknesses. Here it is assumed that categories are represented mainly by prototypes but, in addition, a few exemplars are allowed that are located usually at points that are distant from the category centers or in regions where class boundaries based on distance from prototypes would give erroneous results.

Despite the significant progress represented by these three modeling approaches, they all experience several shortcomings. A key difficulty is that all the models take the form of formal equations for response probabilities. None of them actually learns their exemplars or prototypes using the type of real-time incremental learning process that humans typically experience during a new categorization task. Prototype models define prototypes a priori even though these prototypes might not be the ones that are actually used by human subjects. None of these models explains how exemplar or prototype information may be stored or retrieved in real time as part of the brain's information processing dynamics. In particular, the successful exemplar models all use combinations of exemplars, not individual exemplars, to derive formal response probabilities, but the real-time process whereby these combinations are derived from stored individual exemplars is not specified. Finally, none of these models sheds light upon the types of brain categorization processes for which neurophysiological data have been accumulated in cortical areas like inferotemporal cortex, or IT, from awake behaving monkeys as they learn and perform categorization tasks [18-26].

III. THIRTY COGNITIVE EXPERIMENTS USING THE 5-4 CATEGORY STRUCTURE

A substantial body of the debate over the question of what model best describes human cognitive data has been based on a particular data structure, the so called 5-4 category structure (Table 1). Starting in the early 1980's, exemplar-based models gave consistently better fits to experimental data than prototype-based models. Smith and Minda [8] have shown, however, using thirty data sets of this category structure, that when allowed greater flexibility, prototype models produce results that overcome some of the earlier problems, but this claim has been challenged [27, 28].

Table 1. Schematic of the 5-4 category structure: Binary features on four dimensions (D1,...,D4) define the exemplars in the two categories (A and B). There are five exemplars in category A (A1,..., A5) and four in category B (B1,..., B4).

Type and	Dimension (D)			
Stimulus	D1	D2	D3	D4
Category A				
A1	1	1	1	0
A2	1	0	1	0
A3	1	0	1	1
A4	1	1	0	1
A5	0	1	1	1
Category B				
B1	1	1	0	0
B2	0	1	1	0
B3	0	0	0	1
B4	0	0	0	0
Transfer (T)				
T10	1	0	0	1
T11	1	0	0	0
T12	1	1	1	1
T13	0	0	1	0
T14	0	1	0	1
T15	0	0	1	1
T16	0	1	0	0

Experiments with the 5-4 category structure have used geometric shapes, Brunswick faces, yearbook photos, verbal descriptions, and rocket ship drawings. There are four dimensions with binary values. The whole sample space, therefore, has $2^4 = 16$ different samples. Five samples are labeled as *Class A* and four as *Class B*. The other seven Transfer samples are unlabeled. In many studies that use this category structure [12, 13, 16, 17, 29] it is assumed that class prototypes are the two extreme points of the sample

space; namely, [1, 1, 1, 1] for Class A, and [0, 0, 0, 0] for Class B. Prototype [1, 1, 1, 1] is not a Class A exemplar.

In general, items in *Class A* share more features with the [1, 1, 1, 1] prototype, with the exception of the A2 exemplar. Two of the four exemplars in *Class B*, B3 and B4, share more features with the [0, 0, 0, 0] prototype. For the exceptional exemplars in both categories, the two prototypes are equally well represented. No feature is perfectly diagnostic, as it is not possible to correctly separate items into the two classes based on knowledge of only one dimension.

An index of within-category coherence and betweencategory differentiation, used by Smith and Minda [8], is the *structural ratio*. It is defined as the ratio of withincategory similarity to between-category similarity. The two similarity measures for this category structure are 2.4 and 1.6, respectively. The structural ratio is thus 1.5, which is quite low; a structural ratio of 1.0 implies no differentiation, and a structural ratio greater than 3.0 implies easy differentiation.



Fig. 1. Best fits to the mean of the 30 responses and average of individual best fits shown for one set of data (squares: average of experimental data; circles: best fit to mean experimental data; triangles: average of best fits to individual data).

IV. DISTRIBUTED ARTMAP WITH SELF-SUPERVISION

A. Procedure, Parameters, and Goodness-of-Fit Values

A version of distributed ARTMAP [30] with selfsupervision was used in the simulations. A set of 32 4dimensional input-output pairs was formed from the 16 different stimuli characterizing the 5-4 category structure (each pair included twice). The same 9 pairs (from categories A and B) used in the 30 experiments studied here were used as training data. The stimuli themselves were fed as inputs a to the dARTMAP system and their corresponding category labels were fed as inputs b. In each of the 100 runs, the training and test inputs were randomized and the network was trained until 100% correct categorization was reached in the training phase. In the testing phase, the entire 32-exemplar set (including Transfer samples) was presented sequentially. Categorization scores for each exemplar for each individual run were recorded. When an exemplar is correctly classified in both of its presentations, its score is 2. If it is once misclassified and once correctly classified, its score is 1. Finally, if it is misclassified in both presentations, its score is 0. The mean score for each item is the sum of its score over the 100 runs divided by the total number of presentations of the item (200 = 2x100). These mean scores are the simulation results that are compared with the experimental results.

The only parameters in the model that are tuned by the user to fit the data are the vigilance parameter and the learning rate for unsupervised learning. The first parameter determines how big a mismatch the network can tolerate before searching for a new category. Vigilance thus influences how general a category is and thus the number of memories (adaptive weights) that are needed to learn to categorize all the training inputs. The learning rate determines how much a new exemplar can change on a single learning trial. Simulations for each of the 800 parameter pairs - 20 for vigilance (ranging from 0.05 to 1) and 40 for unsupervised learning rate (ranging from 0.025 to 1) - were run and the pair giving the best fit was picked. The model was tuned to fit both each individual data set and the average of the data (Fig. 1). Statistical analysis of these two parameters indicates that: (1) The best vigilance parameter for the fit to mean data was identical to the best vigilance parameter for 50% of the 30 experiments; (2) The best unsupervised learning rate for the fit to mean data was identical to the best unsupervised learning rate for 67% of the 30 experiments.

B. Prototypes or Exemplars

For each category, we analyzed the distribution of the number of hyper-boxes that were created by learning and the distribution of their sizes. The size of a hyper-box measures how general, or prototype-like, the category is. The number of hyper-boxes measures how distributed the category representation is. The mean number of boxes created for each class over the entire set of simulation runs was 2.2 and 2.1 for Class A and Class B, respectively. This result indicates that there was not a single category center for each region, thus eliminating the possibility of a pure prototype-based representation. On the other hand, this number of hyper-boxes is too small to support a claim for a pure exemplar representation. Instead, the network learns larger hyper-boxes that span most of the category space, but in addition learns 1 or 2 smaller boxes for the items that occupy more marginal parts of the feature space that are close to the category boundary or to the edges of the feature space. Indeed, the histogram of box sizes indicates a bimodal distribution for both Class A boxes and Class B boxes with peaks at the big-size and point-size. This

distribution supports a *rule-plus-exceptions* type of category representation.

C. Predictive Power of Each Dimension

One theoretical measure of the predictive power of each dimension is the ratio of correct category values along that dimension, for all training items to the number of training items (Table 1). For example, along the first dimension, there are four 1s for Class A items and three 0s for Class B items. Thus, the predictive power of this dimension is $(4 + 3)/9 = 7/9 \cdot 0.78$. The further from 0.5 this value is, the more predictive power it has. The values of this measure of predictive power for the other three dimensions are 0.56, 0.78, and 0.67, in order. This index suggests that subjects should use mostly the first and the third dimensions in their categorization decisions and not rely on the second dimension.

If ART captures the dynamics of the categorization decision process, it should be able to reproduce these observations. One way to extract this information from the system parameters relies on the fact that the stimuli in the 5-4 category structure are linearly separable. Consequently, the boxes created by the system should not have substantial overlapping regions. This, in turn, implies that almost all of the weights created in the 200 runs and labeled either Class A or Class B should be linearly separable by a hyper-plane that divides the categories. Then, finding the hyper-plane that optimally separates all weights should give a good estimate of the category boundary. The angle at which this hyper-plane intersects each of the axes (dimensions) of the feature space is a direct measure of the predictive power of the corresponding dimension. The closer this angle is to 90° (or to 270°), the more predictive this dimension is; and conversely, the closer this angle is to 0° (or to 180°) the less predictive it is. The plane could be parameterized in such a way that a bigger parameter for one dimension causes a steeper intersection angle with that dimension; namely,

$$0 = a_0 + a_1 x_1 + a_2 x_2 + a_3 x_3 + a_4 x_4, \tag{1}$$

where x_i is a vector with the values of all weights created in the 200 runs along the *j*th dimension and a_j are the parameters of the plane. The first parameter, a_{m} the bias term, does not effect the inclination of the plane but just its distance from the origin. It is therefore not important for our analysis and will be ignored. Moreover, only the amplitude of the parameters a_i matters, not their sign. We find that $[a_i]$ $a_2 a_3 a_4 = [0.48, -0.33, 0.48, 0.38]$. In order to compare them with the predictive power index introduced at the beginning of the section, we normalize both the theoretical values and the experimental plane parameters such that the sum of their absolute values adds up to one. Then, the relative predictive power indices are [0.28, 0.20, 0.28, 0.24] and the relative hyper-plane parameters are [0.29, 0.20]0.29, 0.23]. In summary, the incrementally learned categories are sensitive to the relative predictive power of the experimental features.

V. DISCUSSION: NORMAL AND AMNESIC CATEGORIZATION

The classical prototype and the exemplar models are based on conflicting assumptions about the nature of category representation in humans, yet they both can provide statistical fits of category data. In order to better characterize the dynamics of category learning and information processing, this article adopted a substantially different approach. Instead of trying to come up with an analytical expression that would map successfully the sixteen four-dimensional input data to observations obtained from 5-4 human categorization experiments, we developed an ART model to carry out the incremental learning and decision making process of each individual used in the experiments and then showed how this model could reproduce the experimental results.

Previous studies have shown that ART-based models can fit other data about brain categorization [1-5, 31-35]. In particular, ART posits that both bottom-up and top-down processes contribute to category learning, shows how a subject can learn which critical feature combinations to attend and which features to ignore, and how sufficiently large mismatches between bottom-up data and learned topdown expectations can drive a memory search for a new or better-fitting category. ART also predicts that matched bottom-up and top-down processes can lead to a resonance that can enable fast learning and also give rise to a conscious brain state. ART learning enables the autonomous creation of new categories and the refinement of previously learned critical feature patterns in response to new exemplars. A dynamically controlled vigilance process helps to determine how general a category will become based on its ability to predict the correct classification. Experimental evidence consistent with predicted properties of vigilance control have been reported in macaque inferotemporal cortex during a categorization task [26].

An ART model has also been used to explain data about the type of abnormal learning and memory that occur during medial temporal amnesia [33, 34]. A lesion of the ART orienting system, which is interpreted to model aspects of hippocampal dynamics, eliminates vigilance control; that is, the lesioned model behaves as if it has a very low vigilance.

Knowlton and Squire [36] reported dissociations between categorization and recognition in amnesic individuals and used these data to argue for multiple memory systems to mediate these tasks. However, Nosofsky and Zaki [37] and Zaki et al. [38] have shown that they can quantitatively fit the Knowlton and Squire and their own data using an exemplar model in which they choose a low value of their sensitivity parameter. Their low sensitivity parameter plays a role like the low vigilance parameter in ART. It should be noted that, when an exemplar model is interpreted as a realtime dynamical processing model, its hypotheses look very much like those of an ART model. These parallel approaches may thus become even more closely linked through future research. In this regard, Nosofsky and Johansen [39, p. 375] have argued that many multiplesystem accounts can be replaced by a single system model when "similarity relations among exemplars change systematically because of selective attention to dimensions and because of changes in the level of sensitivity relating judged similarity to distance in psychological space. Adaptive learning principles may help explain the systematic influence of the selective attention process and of modulation in sensitivity settings on judged similarity." ART provides a dynamical account of how subjects can incrementally learn to selectively pay attention to predictive stimulus dimensions and of how they may alter their vigilance, or sensitivity, in a context-sensitive way.

Good fits to data with the 5-4 category structure were achieved by an ART model with the following selfsupervision refinement: Each test exemplar can perturb those memories that had already been learned in the training phase. This memory change represents a kind of selfsupervised learning. It clarifies why in the testing phase less than 100% classification is observed for exemplars that subjects had previously been trained to perfect performance. This learning scheme fits the data and provides new insights into the prototype-exemplar debate. The simulation results suggest that, for this data structure, subject learning leads to what Nosofsky calls a rule-plus-exceptions approach for categorization: the model created, on average, 2 prototypes per category (as opposed to 1, if it were a purely prototypebased classification) of which one covered a large region of the feature space and the other covered a very small region. These results also clarify why small populations of cells in inferotemporal cortex can be used to categorize many objects in the world. Finally, The ART concept of a vigilance-controlled, attentive, critical feature pattern allows this model to overcome problems of classical prototype and exemplar models while explicating in a real-time neural processing framework the intuitive concepts that made these models so appealing.

ACKNOWLEDGMENTS

This work was supported by research grants from the Air Force Office of Scientific Research (AFOSR F49620-01-1-0397 and F49620-01-1-0423), the Defense Advanced Research Projects Agency and the Office of Naval Research (ONR N00014-01-1-0624), the National Geospatial-Intelligence Agency (NMA 201-01-1-2016), the National Science Foundation (NSF EIA-01-30851, NSF IIS-97-20333, and NSF SBE-0354378), and the Office of Naval Research (ONR N00014-95-1-0657 and ONR N00014-01-1-0624).

REFERENCES

- [1] Grossberg, S. (1999), The link between brain learning, attention, and consciousness. *Consciousness and Cognition*, 8, 1-44.
- [2] Grossberg, S. (2003), Resonant neural dynamics of speech perception. *Journal of Phonetics*, 31, 423-445.
- [3] Grossberg, S., Govindarajan, K.K., Wyse, L.L, & Cohen, M.A. (2004), ARTSTREAM: A neural network model of auditory scene analysis and source segregation. *Neural Networks*, 17, 511-536.
- [4] Grossberg, S. & Myers, C. W. (2000), The resonant dynamics of conscious speech: Interword integration and duration-dependent backward effects. *Psychological Review*, 107, 735-767.
- [5] Raizada, R.D.S. & Grossberg, S. (2003), Towards a theory of the laminar architecture of cerebral cortex: Computational clues from the visual system. *Cerebral Cortex*, 13, 100-113.
- [6] Posner, M.I. & Keele, S.W. (1970). Retention of abstract ideas. *Journal of Experimental Psychology*, 83, 304-308.
- [7] Smith, J. D., & Minda, J. P. (1998). Prototypes in the mist: the early epochs of category learning. J. Experimental Psychology: Learning, Memory, and Cognition, 24, 1411-1430.
- [8] Smith, J. D., & Minda, J. P. (2000). Thirty categorization results in search of a model. J. Experimental Psychology: Learning, Memory, and Cognition, 26, 3-27.
- [9] Smith, J. D., Murray, M. J., & Minda, J. P. (1997). Straight talk about linear separability. J. Experimental Psychology: Learning, Memory, and Cognition, 23, 659-680.
- [10] Estes, W.K. (1994), Classification and cognition. Oxford: Oxford University Press.
- [11] Medin, D. L., Dewey, G. I., & Murphy, T. D. (1983), Relationships between item and category learning: Evidence that abstraction is not automatic. J. Experimental Psychology: Learning and Memory, 9, 607-625.
- [12] Medin, D. L., & Schaffer, M. M. (1978), Context theory of classification learning. *Psychological Review*, 85, 207-238.
- [13] Medin, D. L., & Smith, E. E. (1981), Strategies and classification learning. J. Experimental Psychology: Human Learning and Memory, 7, 241-253.
- [14] Nosofsky, R. M. (1984). Choice, similarity, and the identification-categorization relationship. J. Experimental Psychology: Learning, Memory, and Cognition, 10, 104-114.
- [15] Nosofsky, R. M. (1987). Attention and learning processes in the identification-categorization of integral stimuli. J. Experimental Psychology: Learning, Memory, and Cognition, 13, 87-108.
- [16] Nosofsky, R. M., Kruschke, J.K., &McKinley, S. C. (1992) Combining exemplar-based category representation and connectionist learning rules. J. Experimental Psychology: Learning, Memory, and Cognition, 18, 211-233.

- [17] Palmeri, T. J., Nosofsky, R. M. & McKinley, S. K. (1994) Recognition memory for exceptions to the category rule. J. Experimental Psychology: Learning, Memory, and Cognition, 21, 548-568.
- [18] Desimone, R. (1998). Visual attention mediated by biased competition in extrastriate visual cortex. *Philosophical Transactions of the Royal Society of* London, 353, 1245–1255.
- [19] Desimone, R. & Ungerleider, L.G. (1989), Neural mechanisms of visual processing in monkeys. In Handbook of neuropsychology, Volume 2 (F. Boller & J. Grafman, Eds). Amsterdam: Elsevier, pp. 267-299.
- [20] Gochin, P.M., Miller, E.K., Gross, C.G., & Gerstein, G.L. (1991). Functional interactions among neurons in inferior temporal cortex of the awake macaque. *Experimental Brain Research*, 84, 505-516.
- [21] Harries, M.H. & Perrett, D.I. (1991), Visual processing of faces in temporal cortex: Physiological evidence for a modular organization and possible anatomical correlates. *Journal of Cognitive Neuroscience*, 3, 9-24.
- [22] Mishkin, M. (1982), A memory system in the monkey. Philosophical Transactions of the Royal Society of London B, 298, 85-95.
- [23] Mishkin, M. & Appenzeller, T. (1987). The anatomy of memory. Scientific American, 256, 80-89.
- [24] Perrett, D.I., Mistlin, A.J., & Chitty, A.J. (1987). Visual cells responsive to faces. *Trends in Neurosciences*, 10, 358-364.
- [25] Schwartz, E.L., Desimone, R., Albright, T., & Gross, C.G. (1983). Shape recognition and inferior temporal neurons. *Proceedings of the National Academy of Sciences*, 80, 5776-5778.
- [26] Spitzer, H., Desimone, R., & Moran, J. (1988). Increased attention enhances both behavioral and neuronal performance. *Science*, 240, 338-340.
- [27] Nosofsky, R. M. (2000). Exemplar representation without generalization? Comment on Smith and Minda's (2000) "Thirty categorization results in search of a model". J. Experimental Psychology: Learning, Memory, and Cognition, 26, 1735-1743.
- [28] Nosofsky, R. M. & Zaki, S.R. (2002). Exemplar and prototype models revisited: Response strategies, selective attention, and stimulus generalization. J. Experimental Psychology: Learning, Memory, and Cognition, 28, 924-940.
- [29] Medin D. L., Altom, M. W., & Murphy, T. D. (1984), Given versus induced category representation: use of prototype and exemplar information in classification. J. Experimental Psychology: Learning, Memory, and Cognition, 10, 333-352.
- [30] Carpenter, G. A., Milenova, B. L., Noeske B. W. (1998), Distributed ARTMAP: a neural network for distributed supervised learning. *Neural Networks*, 11, 794-813.
- [31] Boardman, I., Grossberg, S., & Cohen, M.A. (1997), Neural dynamics of variable-rate speech categorization. Journal of Experimental Psychology: Human Perception and Performance, 23, 481-503.

- [32] Carpenter, G.A. & Grossberg, S. (1991), Pattern recognition by self-organizing neural networks. Cambridge, MA: MIT Press.
- [33] Carpenter, G.A. & Grossberg, S. (1993), Normal and amnesic learning, recognition, and memory by a neural model of cortico-hippocampal interactions. *Trends in Neurosciences*, 16, 131-137.
- [34] Grossberg, S. & Merrill, J.W.L. (1996), The hippocampus and cerebellum in adaptively timed learning, recognition, and movement. *Journal of Cognitive Neuroscience*, 8, 257-277.
- [35] Grossberg, S. & Stone, G. O. (1986), Neural dynamics of word recognition and recall: Attentional priming, learning, and resonance. *Psychological Review*, 93, 46-74.
- [36] Knowlton, B.J., & Squire, L.R. (1993). The learning of categories: parallel brain systems for item memory and category knowledge. *Science*, 262(5140), 1747-1749.
- [37] Nosofsky, R. M. & Zaki, S.R. (1998). Dissociations between categorization and recognition in amnesic and normal individuals: An exemplar-based interpretation. *Psychological Science*, 9, 247-255.
- [38] Zaki, S.R., Nososfsky, R.M., Jessup, N.M., & Unverzagt, F.W. (2003). Categorization and recognition performance of a memory-impaired group: Evidence for single-system models. *Journal of the International Neuropsychological Society*, 9, 394-406.
- [39] Nosofsky, R. M., & Johansen, M. K. (2000). Exemplarbased accounts of "multiple-system" phenomena in perceptual categorization. *Psychonomic Bulletin & Review*, 7 (3), 375-402.