

Research Article

A NEURAL NETWORK MODEL OF IMPLICIT MEMORY FOR OBJECT RECOGNITION

Jeffrey N. Rouder, Roger Ratcliff, and Gail McKoon

Northwestern University

Abstract—*People name well-known objects shown in pictures more quickly if they have studied them previously. The most common interpretation of this priming effect is that processing is facilitated by an implicit memory trace in a perceptual representation system. We show that object priming can be explained instead as a bias in information processing, without recourse to an implicit memory system. Assumptions about psychological decision-making processes and bias were added to a neural network model for object identification, and the model accounted for performance both qualitatively and quantitatively in four object identification experiments.*

In a priming experiment with pictured objects, subjects are asked to name the objects and then, as much as a week later, name them again. Repeated objects are named faster than new objects, even by amnesics who cannot consciously recollect them (Cave & Squire, 1992; Mitchell & Brown, 1988; Park & Gabrieli, 1995; Warren & Morton, 1982). These results have been used to support the notion that priming is mediated by an implicit memory system separate from the system that mediates explicit recall (e.g., Schacter & Tulving, 1994; Squire, Knowlton, & Musen, 1993). In this article, we propose instead that the effects of prior encounters come about because of biases in information processing (McKoon & Ratcliff, 1996; Ratcliff & McKoon, 1995, 1996, 1997). Perceptual information is more likely to be interpreted as supporting the identification of objects previously encountered than objects not previously encountered. To model bias, we added psychological mechanisms to an information processing model of object identification (Poggio & Edelman, 1990). In adding bias to such a model, we emphasize that the effect of prior encounters can be understood simply as a by-product of ordinary object perception and identification processes (Morton, 1970); it need not be embodied in a separate memory system.

Repetition “priming” suggests a facilitatory effect, but a prior encounter with a target object actually leads to costs as well as benefits: Although a prior encounter with a target object improves a subject’s ability to name it, a prior encounter with an object that is visually similar to the target but has a different name hurts ability to name the target (Fig. 1), as we showed in a previous experiment (Ratcliff & McKoon, 1996). In the first session of that experiment, subjects were asked to name objects, and then, a week later, they were asked to name objects again. In one condition, the target objects in the second session had been presented in the first session (Session 1: name turtle; Session 2: name turtle). In another condition, objects visually similar to but not the same as the target objects had been presented in the first session (Session 1: name igloo; Session 2: name turtle). In the baseline condition, neither the targets nor the objects similar to them had been presented in the first session. At test, naming times for a target were speeded relative to baseline if the target had been presented in the first

session (a benefit), but they were slowed if the similar object had been presented in the first session (a cost). Our goal for the research described here was to show that an explanation of this bias could be embedded into an information processing, object identification model.

The Poggio and Edelman (1990) model we chose, one of a family of evolving models (Edelman, 1995, 1998; Edelman & Poggio, 1992; Edelman & Weinsall, 1991; O’Toole & Edelman, 1998; Poggio & Gioro, 1989, 1990), is a connectionist network that learns to identify objects from their visual features. For each view that is presented to it of an object it is to learn, the model stores a two-dimensional (2-D), orientation-specific representation of the object. These multiple representations are connected to a single “object” node for the object. By interpolation among the learned orientations, the model can correctly identify the object in orientations that were not taught during learning.

The Poggio and Edelman (1990) model exemplifies one side of a debate about whether multiple 2-D, orientation-specific representations are necessary for object identification or whether a single 3-D, orientation-independent representation alone is sufficient (Biederman & Gerhardstein, 1995; Tarr & Bülhoff, 1995). Our aim was not to contribute to this debate in any direct way, but instead to examine an implicit effect (bias that results from prior encounters) and explain it using an object identification model in such a way as to make testable predictions about naming response times and accuracy rates. The data collected in the experiments described here provided constraints on application of the Poggio and Edelman model, but they are data that could potentially also be well described by a 3-D representation model. The choice of the Poggio and Edelman model as the model to be married with psychological processing mechanisms was determined by the practical requirement of the availability of a full and explicit description of the model. The implementation of the Poggio and Edelman model is designed to serve as an example of what kinds of mechanisms are needed to quantitatively model bias and its associated performance measures.

The four experiments used standard views of line drawings of common objects (Fig. 1). From previous research, we expected that a prior encounter with a target object would lead to facilitation, and that the amount of facilitation would not be significantly affected by whether the orientation of the target object was the same or different on the second presentation as on the first (e.g., Biederman & Gerhardstein, 1993; Srinivas, 1993). However, previous research has not examined whether the costs incurred from prior study of a similar object with a different name are affected by orientation, so Experiments 1 and 2 addressed this question.

In word identification, one of the key findings that motivated a bias account of the effects of prior encounters (Ratcliff & McKoon, 1997) was that bias is observed in a forced-choice paradigm only when the forced-choice alternatives are similar to each other, not when they are dissimilar (Ratcliff, McKoon, & Verwoerd, 1989). For example, if *died* is a target word (flashed briefly, then masked), and *died* and *lied* are the choices, then there is a benefit from prior study of *died* and a cost from prior study of *lied*. But if *died* is the target word, and *died*

Address correspondence to Jeffrey Rouder, Department of Psychology, 210 McAlester Hall, University of Missouri, Columbia, MO 65211.




Fig. 1. Examples of stimuli used in the experiments. The turtle and igloo are similar to each other but dissimilar from the crutch.

and *sofa* are the choices, then there is no benefit from prior study of *died* and no cost from prior study of *sofa*.

The finding of bias with similar but not dissimilar alternatives has two important consequences for modeling. One is obvious—that similarity must play a role in processing. The second is that a prior encounter with a stimulus cannot affect some property of the representation in memory of the stimulus itself. A prior encounter cannot, for example, increase the resting level of activation for a word or lead to the creation of a new representation of a word in memory that facilitates later processing of the word. If this were the case, then bias would be observed in all contexts; it would be observed with dissimilar as well as similar alternatives in forced choice. Instead, a prior encounter with a stimulus must affect processing without affecting the representation of the stimulus itself (Ratcliff & McKoon, 1997).

Experiments 3 and 4 were designed to examine bias in forced choice with objects with similar versus dissimilar alternatives in order to determine how bias should be handled in an object identification model—as an effect on some property of the representation of an object or as a bias of processing mechanisms.

We begin with the experiments, then describe our implementation of the Poggio and Edelman (1990) model, and then show how added assumptions that account for bias and the speed and variability of subjects' performance give a good account of the data.

EXPERIMENTS 1 AND 2

These experiments looked for costs and benefits due to a prior encounter. In Experiment 1, there were two blocks of trials, the second block immediately following the first. A target object in the second block (always presented in upright orientation) was itself presented in

the first block, in either upright or rotated orientation, or was similar to an object that had been presented in the first block, in either upright or rotated orientation. The task for both blocks was to name the objects as quickly as possible, and the dependent measure was response time. In Experiment 2, there were also two blocks of trials, but they were 1 week apart. The same variables were manipulated as in Experiment 1: Either a target object itself or a similar object (in either upright or rotated orientation) was presented in the first block. In the second block, the objects were masked, and the dependent measure was accuracy.

Method

Thirty Northwestern University undergraduates participated in Experiment 1, and 18 in Experiment 2. There were 40 pairs of objects, with the members of a pair similar to each other, plus 30 objects used for fillers in the study and test lists. The objects were drawn black on white and varied in width and height from 2 cm to 6 cm. They were displayed on a PC-VGA gray-scale monitor with the objects either upright or rotated 135° from upright.

In the second block, the target objects from the 40 pairs were always tested in their upright orientation. Either the target itself or its similar pair-mate, in upright or rotated orientation, or neither had appeared in the first block (see Table 1). Targets were counterbalanced across the five conditions. In the first block, 15 fillers were presented in upright orientation and 15 rotated, and in the second block, all fillers were presented in rotated orientation.

In both blocks of Experiment 1 and the first block of Experiment 2, each object was preceded by a fixation point displayed for 500 ms, and then the object was displayed until it was named aloud. Subjects were instructed to name the objects as quickly as possible. In the second block of Experiment 2, the sequence for each trial was as follows: the fixation point for 500 ms, the test object for 67 ms, and then a mask (random shapes and lines) for 200 ms. Subjects were instructed to do their best to name the flashed objects. For both experiments, subjects were instructed to say "no" whenever they could not name a displayed object. Order of presentation of all items was random.

Results

For both experiments (Table 1), there was a bias effect—facilitation relative to baseline when the target object had been pre-

Table 1. Results from Experiments 1 and 2

Condition for first block	Experiment 1		Experiment 2		
	Naming latency (ms)	Response probability		Response probability	
		Target	Pair-mate	Target	Pair-mate
Study target	843 (828)	.93	.01	.77 (.76)	.06 (.03)
Study rotated target	863 (850)	.92	.02	.75 (.75)	.06 (.03)
Study similar pair-mate	1,074 (1,085)	.90	.03	.53 (.50)	.17 (.19)
Study rotated similar pair-mate	1,055 (1,066)	.89	.07	.54 (.52)	.16 (.17)
No study	998 (1,015)	.94	.05	.56 (.57)	.07 (.06)

Note. The numbers in parentheses are the best fit of the neural network model.

sented in the first session and inhibition when the target's similar pair-mate had been presented, with little effect of orientation (as in Biederman & Gerhardstein, 1993; Srinivas, 1993). Analysis of variance on the four nonbaseline conditions showed a significant effect of which was presented in the first block, the target or its pair-mate: Experiment 1, $F(1, 29) = 50.0$; Experiment 2, $F(1, 17) = 41.9$ ($p < .05$ throughout this article). There was no significant effect of orientation or significant interaction ($F_s < 1.0$). The standard errors of the means were 28 ms and .04, respectively.

The conclusion to be drawn is clear: Just as previous research has shown that the benefit to object naming given by a prior encounter with a target object is not significantly affected by changes in orientation, Experiments 1 and 2 have shown that the cost of a prior encounter with a similar object is likewise not significantly affected by orientation.¹ Therefore, we implemented bias in the Poggio and Edelman (1990) model at a stage of processing at which an object is represented as an object, not as a particular orientation of the object or as particular individual features of the object.

EXPERIMENTS 3 AND 4

Experiment 3 investigated whether bias appears in a forced-choice paradigm with both similar and dissimilar alternatives. Target objects were flashed briefly, then masked, and subjects were asked to decide which of two choices matched the target object. The similar-versus-dissimilar variable was manipulated between subjects (to maximize power), and, so that accuracy rates would be in about the same range for the two groups of subjects, longer flash times were used for the subjects tested with similar objects. Experiment 4 compared these two conditions directly with flash time held constant in a within-subjects design.

Method

Ninety Northwestern undergraduates participated in Experiment 3, and 23 in Experiment 4, each for one session. Thirty triples of objects were used in Experiment 3 and 36 in Experiment 4. Each triple consisted of two similar objects and one dissimilar object, always presented in upright orientation.

In the study phase, objects were displayed one at a time for 1.3 s each. Subjects were instructed to learn them for a later (unspecified) memory test. The test phase immediately followed the study phase. For Experiment 3, with similar forced-choice alternatives, either the target itself, its similar triple-mate, or neither was studied. With dissimilar forced-choice alternatives, either the target itself, its dissimilar triple-mate, or neither was studied (Table 2). For 45 subjects, the

1. In the first session of Experiment 1, subjects named objects either upright or rotated 135° from upright. We divided the data into objects that are typically seen in only the upright orientation (e.g., igloo) and objects that are seen in a variety of orientations (e.g., crutch). For objects of the former kind, responses to upright presentations (response time = 1,062 ms) were 135 ms faster than responses to rotated presentations. For objects of the latter kind, this difference was only 22 ms (upright response time = 1,042 ms). Models with 2-D representations could account for this result (by assuming that the object normally viewed in only one orientation has fewer stored views), but models with only 3-D representations would require additional assumptions (e.g., viewpoint-dependent processes as opposed to representations; Biederman & Gerhardstein, 1995).

Table 2. Results from Experiment 3

Study condition	Probability correct	
	Similar alternatives	Dissimilar alternatives
Study target	.73 (.72)	.72 (.72)
Study alternative	.60 (.60)	.71 (.72)
Study neither	.66 (.66)	.73 (.72)

Note. The numbers in parentheses are the best fit of the neural network model.

alternatives were always similar, and for the other 45, they were always dissimilar. Triples of objects were counterbalanced across the three study conditions. The sequence of events for test items was as follows: a fixation point for 500 ms, the target flashed briefly, a mask for 200 ms, and the two alternatives displayed side by side until the subject's response. The flash time was set individually for each subject, in a calibration phase of the experiment that preceded the study list, so that performance would be neither at chance nor at ceiling. The median flash time was 38 ms for subjects in the dissimilar condition and 52 ms for subjects in the similar condition. Subjects pressed the "z" key on the computer keyboard to indicate that the left-hand alternative matched the flashed target and the "/" key to indicate that the right-hand alternative matched. Order of presentation of all items was random.

To maximize power in Experiment 4, we included no study phase, only a test phase. The 36 objects were counterbalanced across the four test conditions (within subjects, Table 3). The procedure for the test phase was the same as for Experiment 3 except that flash time for a target was constant at either 42 ms or 54 ms.

Results

For Experiment 3, the data (Table 2) show significant costs and benefits with similar forced-choice alternatives but not with dissimilar alternatives (a significant interaction, $F[2, 180] = 3.6$; other effects, $F_s < 1.0$). The standard error of the means was .02.

Although in Experiment 3 the flash times were different for the similar and dissimilar conditions, flash times were constant in Experiment 4, and in that experiment, performance was better with dissimilar than similar alternatives, as expected.

Table 3. Results from Experiment 4

Target flash time	Probability correct	
	Similar alternatives	Dissimilar alternatives
42 ms	.55 (.57)	.71 (.68)
54 ms	.57 (.60)	.78 (.75)

Note. The numbers in parentheses are the best fit of the neural network model.

NEURAL NETWORK MODEL

To successfully model the data of Experiments 1 through 4, a model should be able to (a) correctly identify objects and in addition correctly identify them no matter what their orientation; (b) exhibit bias as a result of prior encounters, with the amount of bias independent of orientation differences between first and second encounter; (c) exhibit bias in forced choice with similar but not dissimilar objects; and (d) show human performance characteristics as reflected in accuracy rates, response latencies, and variability across trials. Poggio and Edelman's (1990) model provides the first requirement—the model can identify the objects it has learned, and it can identify them in orientations that were not presented to it during learning. We added to the model a mechanism to give appropriate bias effects. To account for response latencies, we added assumptions about how processing progresses over time, and to account for variability in response latencies and accuracy rates, we added noise into the system.

Objects in the Poggio and Edelman (1990) model are represented as “feature points,” one point for each of an object's vertices, line crossings, and highly curved line segments. For our simulations, we used five feature points per object. The feature coordinates of an object are input to the system (the x, y coordinates of its feature points are scaled to the object's center of mass and always in the same order; Bennett, Hoffman, & Prakash, 1993; Poggio & Edelman, 1990). We added within-trial variability by having the coordinates of the feature points fluctuate moment to moment by small random amounts. We also added across-trial variability by having the coordinates of the feature points vary across trials (Ratcliff, 1978; Ratcliff & Rouder, 1998) in order to model differences in perception of a stimulus across trials and to model inaccurate extraction of stimulus information for short stimulus durations (Ratcliff & Rouder, in press).

Activation flows from the input nodes to object “modules,” one module for each object the model knows (Fig. 2). At the view layer, each node corresponds to one view (one orientation) of an object,

“storing” one particular orientation of the object. Activation (r_i) at a view node i is a function of the difference between the vector of coordinate values at the input layer (\mathbf{P}) and the vector of coordinate values to which view node i maximally responds (\mathbf{Q}_i):

$$r_i = \exp(-k_0 \|\mathbf{P} - \mathbf{Q}_i\|^2).$$

The maximum activation of a view node occurs when the difference is smallest, that is, when the orientation of the input object matches the view node's designated orientation. k_0 (set to 0.0002 in our simulations) is a tuning coefficient that determines how much activation decreases as a function of the difference.

Activation flows from the view nodes to the object layer nodes through weighted connections, v_{ij} . The nodes at the object layer represent the x, y coordinates of the object in its canonical (for our purposes, upright) orientation. The activation value for each object layer node is the sum of the weighted activations from the view nodes: $s_j = \sum_i v_{ij} r_i$. The weights are set so that if the input object matches any one of the stored views of the object, then the activation values at the object layer nodes will be the x, y coordinates of the object in its upright orientation. In other words, a module gives approximate orientation invariance; it serves to translate an object from its input orientation to a standard (upright) orientation. If the orientation of the input object does not exactly match one of the stored orientations, then the coordinate values at the object layer nodes are somewhat distorted. (Of course, the coordinate values are also distorted because of the variability introduced at the input layer percolating through the system.)

To map the model's output onto human performance, we added an output node and a decision node to each module. The amount of activation at the output node at time t represents the similarity between the coordinate values at the object layer nodes, the vector $\mathbf{S} =$

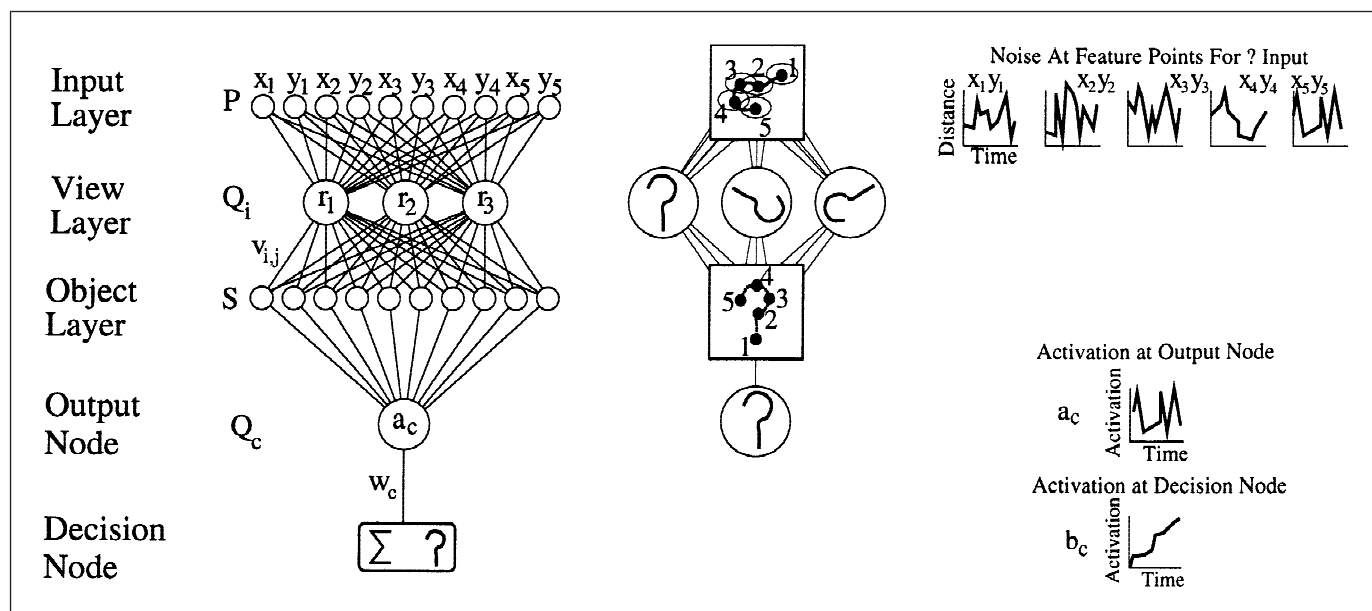


Fig. 2. A Poggio-Edelman module that recognizes question marks. The left side depicts the nodes of the module, the center gives a schematic of what the nodes represent, and the right side shows input and activation values over time.

(s_1, s_2, \dots), and the coordinate values for the canonical orientation, the vector \mathbf{Q}_c , calculated by Shepard's (1987) similarity metric, with the scaling parameter k (set to 0.1 in our simulations):

$$a_c = \exp(-k\|\mathbf{S} - \mathbf{Q}_c\|);$$

this gives activation values greater than 0 for object c .

Because of variability in the x, y coordinates at the input level, the value of a_c fluctuates moment to moment. The decision node for an object module sums, over time, the object's a_c values, weighted by the amount of bias w_c toward the object, relative to the a_c values for all the other objects in the system (Luce's, 1959, choice rule); the amount of activation at the decision node at time t is

$$b_c(t) = \frac{\sum_{\tau} [w_c a_c(\tau)]}{\sum_d w_d a_d(\tau)}.$$

The stopping rule is that the model makes an identification decision when the amount of activation in the decision node with the largest amount of activation exceeds the amount of activation in the decision node with the next largest amount of activation by some criterial amount (e.g., Audley & Pike, 1965; Ratcliff & McKoon, 1997).

The amount of bias w_c for an object reflects some learned likelihood of the system seeing object c relative to all other objects. We assumed that the effect of a prior encounter with an object in an experiment was to increase the bias toward that object. The effect of increasing w_c is to increase activation at the decision node for object c (b_c) and decrease activation at the decision nodes for all other objects.

The model predicts greater bias effects for forced choice with similar choices than for forced choice with dissimilar choices. This is because changes in weights (w_c) have a large effect when the amount of activation (b_c) is in the middle of the activation range (e.g., .5) but a very small effect when the amount of activation is very small or very large. For example, a change of 10% is .05 when activation is .5, but only .001 when activation is .01 or .99. When the choices are similar, if activation (b_c) is in the middle of the range for one of the similar choices, it is likely to be in the middle of the range for the other. But when the alternatives are dissimilar, if one is in the middle of the range, the other is usually small or large.

Figure 3 illustrates the decision process. Fluctuation of the input values over time makes the amounts of activation at the output nodes fluctuate, which makes decision time variable. The shape of the distribution of decision times is in part determined by the use of the relative stopping rule (Ratcliff & McKoon, 1997), and the distribution is right skewed, corresponding to empirically observed response time distributions. The fluctuations in activation values also lead to errors because the amount of activation in some other object's decision node can sometimes be greater than the amount in the input object's decision node.

FITTING THE MODEL TO DATA

The simulations used 21 objects. A target object was constructed by randomly sampling five points in the x - y plane, with the average

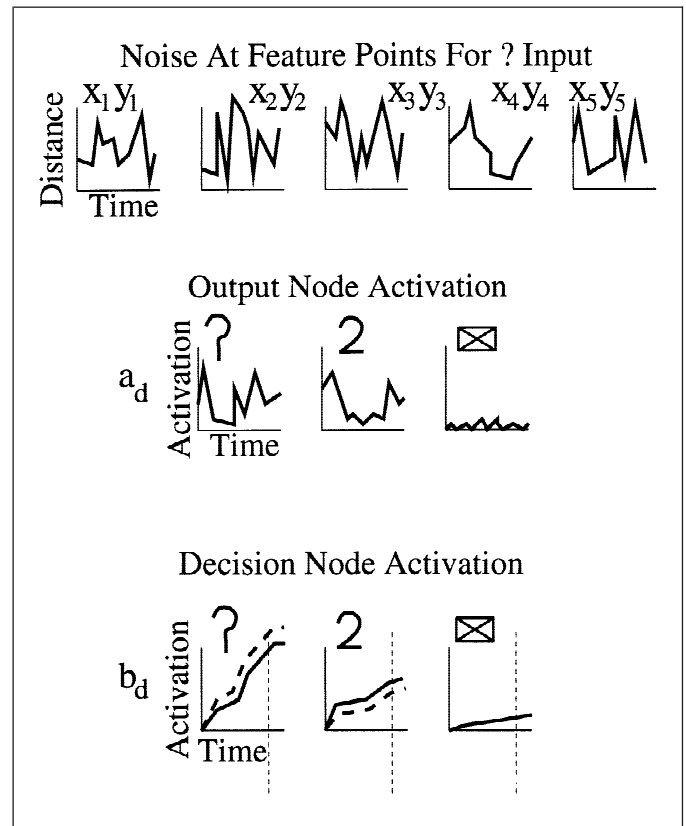


Fig. 3. Input and activation values over time for a question mark as an input object and three modules, a question mark, the number 2, and an envelope. The dashed lines show the effects of prior study of the question mark.

radial distance from a point to the center of the object being 28 pixels. An object "most similar" to the target was constructed by moving each of the points 4.0 pixels in a random direction. Six "highly similar" objects were constructed by moving the points 4.2 pixels, and 12 "similar" objects were constructed by moving the points 12 pixels. One dissimilar object (to represent all dissimilar objects) was constructed by sampling five new points in the x - y plane. For each object, there were six view nodes (six different orientations). We implemented the model as already having learned the weights between the view nodes and the object layer (v_{ij}) that allowed it to correctly identify all 21 objects. Without any variability in the input feature points, the output-node activation for the target object was .9999 if it was tested in an orientation that had been learned and .990 if it was tested in an unlearned orientation. The output-node activation for the most similar, highly similar, and similar objects ranged from .84 to .05, and the output-node activation for the dissimilar object was less than .001.

Activation in the model was updated in discrete cycles with within-trial variability added to the pixel locations at each cycle. For forced choice, activation was accumulated only by the decision nodes for the two alternative choices (Ratcliff & McKoon, 1997). The number of cycles to reach criterion-determined response time (RT) is

$$RT = mN + T_{er},$$

Object Recognition and Implicit Memory

where m is time per cycle (set at 13 ms) and T_{er} is the time taken by nondecisional processes like stimulus encoding and response execution (set at 380 ms). The final parameter was the effect on bias, w_d , of a prior encounter: If an object was upright for both encounters, w_d increased from 1.00 to 1.040; if the orientations were different, w_d increased from 1.00 to 1.035.

Once the number of objects, their similarities, and the scaling parameters (k_0 , k) are set, the model has six free parameters. The model was fit to the data from all conditions and experiments simultaneously with T_{er} , time per cycle m , within-trial variability, and bias w_d fixed, and between-trial variability in pixel location and the response criterion allowed to vary across conditions and experiments (see Table 4). The best fits of the model to the data are shown in Tables 1, 2, and 3. For Experiment 1, the model also does a reasonable job of predicting the shapes of the naming latency distributions (Fig. 4).

It should be emphasized that the parameters of the model are not free parameters as, say, in a multiple regression model; rather, the parameters are all part of an integrated framework, so that changes in any one parameter result in changes to all the predictions of the model. Thus, there is considerably less model freedom than might at first appear.

GENERAL DISCUSSION

The model successfully explains performance in three object identification tasks, quantitatively accounting for naming latencies and their distributions, accuracy rates, and probability correct in forced choice. The model predicts both costs and benefits from prior encounters, and it predicts that they will be observed in forced choice with similar but not dissimilar alternatives. Although object identification has been labeled an “implicit memory task” (Schacter, 1994; Schacter & Tulving, 1994; Squire, 1994), the model accounts for performance without recourse to a separate memory system to hold representations of prior encounters with stimuli. Instead, bias from prior encounters is explained in terms of the information processing mechanisms that are used for object identification. Bias comes about because of extra weight given to perceptual evidence for an object previously encountered in the experiment. Because an object’s activation is calculated relative to the activations of all other objects, giving extra weight to




Fig. 4. Cumulative response time distributions from Experiment 1 (constructed by averaging over subjects; Ratcliff, 1979). The letters are the data points, and the lines are the best fit from the model. T = target; S = study object similar to target; N = study neither target nor object similar to target.

evidence for a previously encountered object adds to its activation and takes away from the activations of similar objects that were not previously encountered. This is essentially the same explanation of bias that is given by the counter model for word identification (Ratcliff & McKoon, 1997).

The model we used (based on Poggio & Edelman, 1990) is a connectionist network. Connectionist models are often designed to handle perceptual learning, and so it might have been expected that this particular model would do well, but a successful outcome was not guaranteed. For one thing, no previous model has been applied to make quantitative predictions about bias, response time, and accuracy in object identification. For another, implementing mechanisms to predict response time and accuracy in connectionist models is possible, but not straightforward (Ratcliff, Van Zandt, & McKoon, 1999).

It might, of course, be the case that a model that allows only a 3-D representation for an object (e.g., Hummel & Biederman, 1992), not

Table 4. Target duration and parameter values used in fitting the model

Target duration and parameter values	Experiment 1: Naming	Experiment 2: Naming	Experiment 3: Forced choice, similar alternatives	Experiment 3: Forced choice, dissimilar alternatives	Experiment 4: Forced choice, similar and dissimilar alternatives	Experiment 4: Forced choice, similar and dissimilar alternatives
Target duration (ms)	Until response	67	52	38	42	54
Between-trial variability ^a	1.4	2.9	9.8	33.5	32.3	27.0
Decision criterion	0.5	0.5	5.0	5.0	5.0	5.0

^aStandard deviation in a two-dimensional uniform distribution of the number of pixels away from the feature point.

multiple 2-D representations, could be developed to fit the data from the experiments presented here. There is no aspect of the data that requires 2-D representations. The model we implemented simply provides a standard: Other successful models would have to provide quantitatively accurate predictions for all the same effects and dependent measures.

The model is consistent with dissociations that have been observed between performance in object identification experiments and performance in experiments that require recognition of whether an object was or was not encountered before in the experiment (e.g., Mitchell & Brown, 1988). The object identification processes described in the model likely precede processes that are involved with explicit memories, so they are not likely to be affected by the variables that affect recognition memory or recall. Likewise, because damage to the parts of the brain that are involved with explicit memories will likely not affect parts of the brain involved in object identification, spared identification processing is expected for amnesic patients (Ratcliff & McKoon, 1996).

The model we describe has characteristics of psychological processing embedded well inside of it. The perceptual variability that starts at the input layer percolates all the way through the model to the decision mechanism. The psychological mechanisms integrated into the computational model provide an account of bias, response times, accuracy rates, and variability in performance, and thus give a benchmark for other models that might attempt to model psychological data from object identification tasks.

Acknowledgments—This research is supported by National Institute of Mental Health Grant HD-MH44640 to R. Ratcliff and National Institute for Deafness and Other Communication Disorders Grant R01-DC01240 to G. McKoon.

REFERENCES

- Audley, R.J., & Pike, A.R. (1965). Some alternative stochastic models of choice. *British Journal of Mathematical and Statistical Psychology*, *18*, 207–225.
- Bennett, B.M., Hoffman, D.B., & Prakash, C. (1993). Recognition polynomials. *Journal of the Optical Society of America*, *10*, 759–764.
- Biederman, I., & Gerhardstein, P.C. (1993). Recognizing depth-rotated objects: Evidence for three-dimensional viewpoint invariance. *Journal of Experimental Psychology: Human Perception and Performance*, *19*, 1162–1182.
- Biederman, I., & Gerhardstein, P. (1995). Viewpoint-dependent mechanisms in visual object recognition: Reply to Tarr and Bülthoff. *Journal of Experimental Psychology: Human Perception and Performance*, *21*, 1506–1514.
- Cave, C., & Squire, L. (1992). Intact and long-lasting repetition priming in amnesia. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *18*, 509–520.
- Edelman, S. (1995). Representation, similarity, and the chorus of prototypes. *Minds and Machines*, *5*, 45–68.
- Edelman, S. (1998). Representation is representation of similarities. *Behavioral and Brain Sciences*, *21*, 449–498.
- Edelman, S., & Poggio, T. (1992). Bringing the grandmother back into the picture: A memory-based view of object recognition. *International Journal of Pattern Recognition and Artificial Intelligence*, *6*, 37–61.
- Edelman, S., & Weinhall, D. (1991). A self-organizing multiple-view representation of 3d objects. *Biological Cybernetics*, *64*, 209–219.
- Hummel, J., & Biederman, I. (1992). Dynamic binding in a neural network for shape recognition. *Psychological Review*, *99*, 480–517.
- Luce, R.D. (1959). *Individual choice behavior*. New York: Wiley.
- McKoon, G., & Ratcliff, R. (1996). Separating implicit from explicit retrieval processes in perceptual identification. *Consciousness and Cognition*, *5*, 500–511.
- Mitchell, D., & Brown, A. (1988). Persistent repetition priming in picture naming and its dissociation from recognition memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *14*, 213–222.
- Morton, J. (1970). A functional model for memory. In D.A. Norman (Ed.), *Models of human memory* (pp. 203–254). New York: Academic Press.
- O'Toole, A.J., & Edelman, S. (1998). Stimulus-specificity effects in face recognition over changes in viewpoint. *Vision Research*, *38*, 2351–2363.
- Park, S.M., & Gabrieli, J.D.E. (1995). Perceptual and non-perceptual components of implicit memory for pictures. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *21*, 1583–1594.
- Poggio, T., & Edelman, S. (1990). A network that learns to recognize three-dimensional objects. *Nature*, *343*, 263–266.
- Poggio, T., & Girosi, F. (1989). *A theory of networks for approximation and learning* (Artificial Intelligence Memo No. 1140). Cambridge, MA: Massachusetts Institute of Technology.
- Poggio, T., & Girosi, F. (1990). Regularization algorithms for learning that are equivalent to multilayer networks. *Science*, *247*, 978–982.
- Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, *85*, 59–108.
- Ratcliff, R. (1979). Group reaction time distributions and an analysis of distribution statistics. *Psychological Bulletin*, *86*, 446–461.
- Ratcliff, R., & McKoon, G. (1995). Bias in the priming of object decisions. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *21*, 754–767.
- Ratcliff, R., & McKoon, G. (1996). Biases in implicit memory tasks. *Journal of Experimental Psychology: General*, *125*, 403–421.
- Ratcliff, R., & McKoon, G. (1997). A counter model for implicit priming in perceptual word identification. *Psychological Review*, *104*, 319–343.
- Ratcliff, R., McKoon, G., & Verwoerd, M. (1989). A bias interpretation of facilitation in perceptual identification. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *15*, 378–387.
- Ratcliff, R., & Rouder, J.N. (1998). Modeling response times for decisions between two choices. *Psychological Science*, *9*, 347–356.
- Ratcliff, R., & Rouder, J.N. (in press). A diffusion model analysis of letter masking. *Journal of Experimental Psychology: Human Perception and Performance*.
- Ratcliff, R., Van Zandt, T., & McKoon, G. (1999). Connectionist and diffusion models of reaction time. *Psychological Review*, *106*, 261–300.
- Schacter, D. (1994). Priming and multiple memory systems: Perceptual mechanisms of implicit memory. In D.L. Schacter & E. Tulving (Eds.), *Memory systems 1994* (pp. 233–268). Cambridge, MA: MIT Press.
- Schacter, D.L., & Tulving, E. (1994). What are the memory systems of 1994? In D.L. Schacter & E. Tulving (Eds.), *Memory systems 1994* (pp. 1–38). Cambridge, MA: MIT Press.
- Shepard, R.N. (1987). Toward a universal law of generalization for psychological science. *Science*, *237*, 1317–1323.
- Squire, L.R. (1994). Declarative and nondeclarative memory: Multiple brain systems supporting learning and memory. In D.L. Schacter & E. Tulving (Eds.), *Memory systems 1994* (pp. 203–231). Cambridge, MA: MIT Press.
- Squire, L.R., Knowlton, B., & Musen, G. (1993). The structure and organization of memory. *Annual Review of Psychology*, *44*, 453–495.
- Srinivas, K. (1993). Perceptual specificity in nonverbal priming. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *19*, 582–602.
- Tarr, M., & Bülthoff, H. (1995). Is human object recognition better described by geostructural-descriptions or by multiple-views? Comment on Biederman and Gerhardstein (1993). *Journal of Experimental Psychology: Human Perception and Performance*, *21*, 1494–1505.
- Warren, C., & Morton, J. (1982). The effects of priming on picture presentation. *British Journal of Psychology*, *73*, 117–129.

(RECEIVED 1/4/99; ACCEPTED 3/9/99)