Sequential dependencies in human behavior offer insights into cognitive control

Michael C. Mozer
Department of Computer Science and
Institute of Cognitive Science
University of Colorado
Boulder, CO 80309-0430 USA

Sachiko Kinoshita
MACCS and
Department of Psychology
Macquarie University
Sydney, NSW 2109
Australia

Michael Shettel
Department of Computer Science
University of Colorado
Boulder, CO 80309-0430 USA

Abstract

We present a perspective on cognitive control that is motivated by an examination of *sequential dependencies* in human behavior. A sequential dependency is an influence of one incidental experience on subsequent experience. Sequential dependencies arise in psychological experiments when individuals perform a task repeatedly or perform a series of tasks, and one task trial influences behavior on subsequent trials. For example, in a naming task, individuals are faster to name a word after having just named easy (e.g., orthographically regular) words than after having just named difficult words. And in a choice task, individuals are faster to press a response key if the same response was made on recent trials than if a different response had been made. We view sequential dependencies as reflecting the fine tuning of cognitive control to the structure of the environment. We discuss the two sequential phenomena just mentioned, and present accounts of the phenomena in terms of the adaptation of cognitive control. For each phenomenon, we characterize cognitive control in terms of constructing a predictive model of the environment and using this model to optimize future performance. This same perspective offers insight not only into adaptation of control, but how task instructions can be translated into an initial configuration of the cognitive architecture.

INTRODUCTION

In this chapter, we present a particular perspective on cognitive control that is motivated by an examination of *sequential dependencies* in human behavior. At its essence, a sequential dependency is an influence of one incidental experience on subsequent experience. Sequential dependencies arise both in naturalistic settings and in psychological experiments when individuals perform a task repeatedly or perform a series of tasks, and performing one task trial influences behavior on subsequent trials. Measures of behavior are diverse, including response latency, accuracy, type of errors produced, and interpretation of ambiguous stimuli.

To illustrate, consider the three columns of addition problems in Table 1. The first column is a series of easy problems; individuals are quick and accurate in naming the sum. The second column is a series of hard problems; individuals are slower and less accurate in responding. The third column contains a mixture of easy and hard problems. If sequential dependencies arise in repeatedly naming the sums, then the response time or accuracy to an easy problem will depend on the preceding context, i.e., whether it appears in an easy or mixed block; similarly, performance on a hard problem will depend on whether it appears in a hard or mixed block. Exactly this sort of dependency has been observed (Lupker, Kinoshita, Coltheart, & Taylor, 2003): responses to a hard problem are faster but less accurate in a mixed block than in a pure

TABLE 1. Three blocks of addition problems

EASY BLOCK	HARD BLOCK	MIXED BLOCK
3 + 2	9 + 4	3 + 2
1 + 4	7 + 6	7 + 6
10 + 7	8 + 6	10 + 7
5 + 5	6 + 13	6 + 13

block; similarly, responses to an easy problem are slower and more accurate in a mixed block than in a pure block of easy trials. Essentially, the presence of recent easy problems causes response-initiation processes to treat a hard problem as if it were easier, speeding up responses but causing them to be more error prone; the reverse effect occurs for easy problems in the presence of recent hard problems.

Sequential dependencies reflect cortical adaptation operating on the time scale of seconds, not—as one usually imagines when discussing learning—days or weeks. Sequential dependencies are robust and nearly ubiquitous across a wide range of experimental tasks. Table 2 presents a catalog of sequential dependency effects, spanning a variety of components of the cognitive architecture, including perception, attention, language, stimulus-response mapping, and response initiation. Sequential dependencies arise in a variety of experimental paradigms. The aspect of the stimulus that produces the dependency—which we term the *dimension of dependency*—ranges from the concrete, such as color or identity, to the abstract, such as cue validity and item difficulty. Most sequential dependencies are fairly short lived, lasting roughly five intervening trials, but some varieties span hundreds of trials and weeks of passing time (e.g., global display configuration, Chun and Jiang, 1998; syntactic structure, Bock, 2002).

Sequential dependencies may be even more widespread than Table 2 suggests, because they are ignored in the traditional psychological experimental paradigm. In a typical experiment, participants perform dozens of practice trials during which data is not collected, followed by experimental trials that are randomized such that when aggregation is performed over trials in a particular experimental condition, sequential effects are cancelled. When sequential effects are studied, they are often larger than other experimental effects explored in the same paradigm; for example, in visual search, sequential effects can modulate response latency by 100 msec given latencies in the 700 msec range (e.g., Wolfe et al., 2003).

Sequential dependencies are often described as a sort of *priming*, facilitation of performance due to having processed similar stimuli or made similar responses in the past. We prefer not to characterize sequential dependencies using the term priming for two reasons. First, priming is often viewed as an experimental curiosity used to diagnose the nature of cognitive representations, one which has little bearing on naturalistic tasks and experience. Second, many sequential dependencies are not due to repetitions of specific stimulus identities or features, but rather to a more abstract type of similarity. For example, in the arithmetic-problem difficulty manipulation described earlier, problem *difficulty*, not having experience on a specific problem, induces sequential dependencies; and in language, syntactic structure induces sequential dependencies, not particular words or semantic content.

COGNITIVE CONTROL

We view sequential dependencies as a strong constraint on the operation of *cognitive control*. Cognitive control allows individuals to flexibly adapt behavior to current goals and task demands. Aspects of cognitive control include the deployment of visual attention, the selection of responses, forming arbitrary associations between stimuli and responses, and using working memory to subserve ongoing processing. At its essence, cognitive control involves translating a task specification into a configuration of the cognitive architecture appropriate for performing that task. But cognitive control involves a secondary, more subtle, ability—that of fine tuning the operation of the cognitive architecture to the environment. For example, consider searching for a key in a bowl of coins versus searching for a key on a black leather couch. In the

TABLE 2. A catalog of sequential dependency effects

component of architecture	experimental paradigm	dimension of dependency	example citations
perception	figure-ground	stimulus color	Vecera (2005)
	identification	stimulus shape and identity	Bar & Biederman (1998); Ratcliff & McKoon (1997)
	intensity judgement	stimulus magnitude	Lockhead (1984, 1995)
	categorization	stimulus features	Johns & Mewhort (2003); Stewart et al. (2002)
	ambiguous motion	previous judgements	Maloney et al. (2005)
stimulus-response mapping	task switching	task set	Rogers & Monsell (1995)
language	semantic judgement	syntactic structure	Bock & Griffin (2000)
response initiation	word naming	task difficulty	Taylor & Lupker (2001)
	choice		Kiger & Glass (1981); Strayer & Kramer (1994a)
		response	Jentszch & Sommer (2002); Jones et al. (2003)
attention	cued detection and identification	cue validity	Bodner & Masson (2001); Posner (1980)
	visual search	stimulus features	Maljkovic & Nakayama (1996); Wolfe et al. (2003)
		scene configuration and statistics	Chun & Jiang (1998, 1999)

former case, the environment dictates that the most relevant feature is the size of the key, whereas in the latter case, the most relevant feature is the metallic luster of the key.

We adopt the perspective that sequential dependencies reflect this fine tuning of cognitive control to the structure of the environment. We discuss two distinct sequential phenomena, and present accounts of the phenomena in terms of the adaptation of cognitive control. For each phenomenon, we assume that cognitive control involves constructing a predictive model of the environment, and using this model to optimize future performance.

SEQUENTIAL EFFECTS INVOLVING RESPONSE REPETITION

In this section, we model a speeded discrimination paradigm in which individuals are asked to classify a sequence of stimuli (Jones et al., 2003). The stimuli are letters of the alphabet, ***A***–***Z***, presented in rapid succession, and individuals are asked to press one response key if the letter is an ***X*** or another response key for any letter other than ***X*** (as a shorthand, we will refer to the alternative responses as R_1 and R_2). Jones et al. (2003) manipulated the relative frequency of R_1 and R_2; the ratio of presentation frequency was either 1:5, 1:1, or 5:1. Response conflict arises when the two stimulus classes are unbalanced in frequency, resulting in more errors and slower reaction times. For example, when R_1's are frequent but R_2 is presented, individuals are predisposed toward producing the R_1 response, and this predisposition must be overcome by the perceptual evidence from the R_2. Cognitive control is presumed to be required in situations involving response conflict. In this task, response repetition is key, rather than stimulus repetition, because effects are symmetric for R_1 and R_2, even though one of the responses corresponds to many distinct stimuli, and those stimuli are *not* repeated.

A probabilistic information transmission model

The heart of our account is an existing model of probabilistic information transmission (PIT) that explains a variety of facilitation effects that arise from long-term repetition priming (Colagrosso, 2004; Colagrosso & Mozer, 2005; Mozer, Colagrosso, & Huber, 2003), and more broadly, that addresses changes in the nature of information transmission in neocortex due to experience. We give a brief overview of the aspects of this model essential for the present work.

The model posits that the cognitive architecture can be characterized by a collection of information-processing *pathways*, and any act of cognition involves coordination among pathways. To model a simple discrimination task, we might suppose a *perceptual pathway* to map the visual input to a semantic representation, and a *response pathway* to map the semantic representation to a response. The model is framed in terms of probability theory: pathway inputs and outputs are random variables and inference in a pathway is carried out by Bayesian belief revision.

To elaborate, consider a pathway whose input at time t is a discrete random variable, denoted X(t), which can assume values 1, 2, 3, ... n_x corresponding to alternative input states. Similarly, the output of the pathway at time t is a discrete random variable, denoted Y(t), which can assume values 1, 2, 3, ... n_y. For example, the input to the perceptual pathway in the discrimination task is one of $n_x = 26$ visual patterns corresponding to the letters of the alphabet, and the output is one of $n_y = 26$ letter identities.[1] To present a particular input alternative, i, to the model for T time steps, we clamp $X(t) = i$ for $t = 1 \ldots T$. The model computes a probability distribution over Y given X , i.e., P(Y(t) | X(1) ... X(t)), the probability over the output states given the input sequence. [2]

A pathway is modeled as a dynamic Bayes network; the minimal version of the model used in the present simulations is simply a hidden Markov model, where the X(t) are observations and the Y(t) are inferred state (see Figure 1, left panel).[3] To understand the diagram, ignore the directionality of the arrows, and notes simply that Y(t) is linked to both Y(t–1) and X(t), meaning that Y(t) is constrained by these other two variables. To compute P(Y(t) | X(1) ... X(t)), it is necessary to specify three probability distributions.

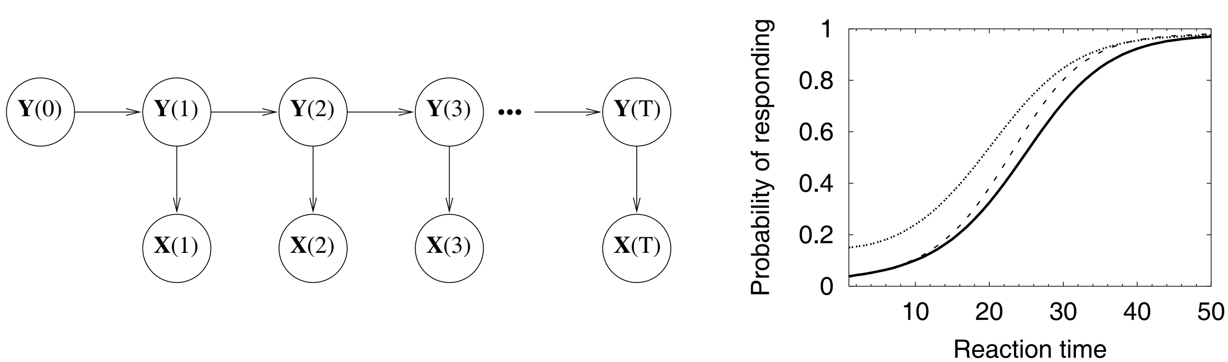

FIGURE 1. (left panel) basic pathway architecture; (right panel) time course of inference in a pathway

1. This model is highly abstract: the visual patterns are enumerated, but the actual pixel patterns are not explicitly represented in the model. Nonetheless, the similarity structure among inputs can be captured, but we skip a discussion of this issue because it is irrelevant for the current work.

2. A brief explanation of probability notation: If V is a random varaible, then P(V) denotes a distribution over values that the variable can take on. In the case of a discrete random variable, P(V) denotes a vector of values. For example, if V can take on the values v_1, v_2, and v_3, then P(V) might represent the probability vector [0.3, 0,6, 0.1], meaning that V has value v_1 with probability 0.3, and so forth. To denote the probability of V taking on a certain value, we use the standard notation P(V=v_1), and in this example, P(V=v_1)=0.3. The notation P(V|W) denotes the probability vector for V given a specific, yet unspecified value of W.

3. In typical usage, a hidden Markov model (HMM) is presented with a sequence of distinct inputs, whereas we maintain the same input for many successive time steps. Further, in typical usage, an HMM transitions through a sequence of distinct hidden states, whereas we attempt to converge with increasing confidence on a single state. Thus, our model captures the time course of information processing for a single event.)

The particular values hypothesized for these three distributions embody the knowledge of the model and give rise to predictions from the model. The three distributions are:

(1) P(Y(t)|Y(t–1)), which characterizes how the pathway output evolves over time, i.e., how the output at time t, Y(t), depends on the output at time t–1, Y(t–1);

(2) P(X(t) | Y(t)), which characterizes the *strength of association* between inputs and outputs, i.e., how likely it is to observe a given state of the input at some point in time, X(t), if the correct output at that time, Y(t), is a known state; and

(3) P(Y(0)), the *prior* distribution over outputs, i.e., in the absence of any information about the relative likelihood of the various output states.

To give a sense of how PIT operates, the right panel of Figure 1 depicts the time course of inference in a single pathway which has 26 input and output alternatives, with one-to-one associations. The solid line in the Figure shows, as a function of time t, P(Y(t) = 1 | X(1)=1 ... X(t)=1), i.e., the probability that a given input will produce its target output. Due to limited association strengths, perceptual evidence must accumulate over many iterations in order for the target to be produced with high probability. The densely dashed line shows the same target probability when the target prior is increased, and the sparsely dashed line shows the target probability when the association strength to the target is increased. Increasing either the prior or the association strength causes the speed-accuracy curve to shift to the left. In our previous work, we proposed a mechanism by which priors and association strengths are altered following each experience. This mechanism gives rise to sequential effects; we will show that it explains the response-repetition data described earlier.

PIT is a generalization of random walk models and has several advantages. It provides a mathematically principled means of handling multiple alternative responses (necessary for naming) and similarity structure among elements of representation, and characterizes perceptual processing, not just decision making. The counter model (Ratcliff & McKoon, 1997) or connectionist integrator models (e.g., Usher & McClelland, 2001) could also serve us, although PIT has an advantage in that it operates using a currency of probabilities—versus more arbitrary units of *counts* or *activation*—which has two benefits. First, fewer additional assumptions are required to translate model output to predictions of experimental outcomes: If the tendency to make responses is expressed as a probability distribution over alternatives, stochastic sampling can be used to obtain a response, whereas if response tendency is expressed as activation, an arbitrrary tranformation must be invoked to transform activation into a response (e.g., a normalized exponential transform is often used in connectionist models). Second, operating in a currency of probability leads to explicit, interpretable decision criteria and learning mechanisms; for example, Bayes rule can be used to determine an optimal decision criterion or update of beliefs after obtaining evidence, whereas the currency of activation in connectionist models allows for arbitrary threshold and learning rules.

Model details

The simulations we report in this chapter utilize a cascade of two pathways. A perceptual pathway maps visual patterns (26 alternatives) to a letter-identity representation (26 alternatives), and a response pathway maps the letter identity to a response. For the choice task, the response pathway has two outputs, corresponding to the two response keys. The interconnection between the pathways is achieved by copying the output of the perceptual pathway, $Y^p(t)$, to the input of the response pathway, $X^r(t)$, at each time. The free parameters of the model are mostly task and experience related. Nonetheless, in the current simulations we used the same parameter values as Mozer et al. (2004), with one exception: Because the speeded perceptual discrimination task studied here is quite unlike the tasks studied by Mozer et al., we allowed ourselves to vary the association-strength parameter in the response pathway. This parameter has only a quantitative, not qualitative, influence on predictions of the model.

In our simulations, we also use the priming mechanism proposed by Mozer et al. (2004). Essentially, this mechanism constructs a *model of the environment*, which consists of the prior probabilities of the various stimuli and responses. To elaborate, the priors for a pathway are internally represented in a nonnormalized form: the nonnormalized prior for alternative i is p_i, and the normalized prior is

$$P(Y(0) = i) = p_i / \sum_j p_j .$$

The priming mechanism maintains a running average of recent experience. On each trial, the priming mechanism increases the nonnormalized prior of alternative *i* in proportion to its asymptotic activity at final time *T*, and all priors undergo exponential decay:

$$\Delta p_i = \gamma \mathrm{P}(\mathrm{Y}(T) = i | X(1) \ldots X(T)) - \varepsilon p_i ,$$

where γ is the strength of priming, and ε is the decay rate. (The Mozer et al. model also performs priming in the association strengths by a similar rule, which is included in the present simulation although it has a negligible effect on the results here.)

This priming mechanism yields priors on average that match the presentation probabilities in the task, e.g., .17 and .83 for the two responses in the 1:5 condition of the Jones et al. experiment. Consequently, when we report results for overall error rate and reaction time in a condition, we make the assumption of rationality that the model's priors correspond to the true priors of the environment. Although the model yields the same result when the priming mechanism is used on a trial-by-trial basis to adjust the priors, the explicit assumption of rationality avoids any confusion about the factors responsible for the model's performance. We use the priming mechanism on a trial-by-trial basis to account for performance conditional on recent trial history, as explained later.

Control processes and the speed-accuracy trade off

The response pathway of the model produces a speed-accuracy performance function much like that in the right panel of Figure 1. This function characterizes the operation of the pathway, but it does not address the control issue of when in time to initiate a response. A control mechanism might simply choose a threshold in accuracy or in reaction time, but we hypothesize a more general, rational approach in which a *response utility* is computed, and control mechanisms initiate a response at the point in time when a maximum in utility is attained.

When stimulus **S** is presented and the correct response is **R**, we posit a utility of responding at time *T* following stimulus onset:

$$U(T|\mathbf{S}, \mathbf{R}) = \int_{t=0}^{T} [P(Y^r(t) = \mathbf{R}|\mathbf{S}) - \kappa t] dt \tag{1}$$

This utility involves two terms, the accuracy of response and the reaction time. Utility increases with increasing accuracy and decreases with response time. The relative importance of the two terms is determined by κ. This form of utility function leads to an extremely simple stopping rule, which we'll explain shortly.

We assume that κ depends on task instructions: if individuals are told to make no errors, κ should be small to emphasize the error rate; if individuals are told to respond quickly and not concern themselves with occasional errors, κ should be large to emphasize the reaction time. We picked a value of κ to obtain the best fit to the human data.

The utility cannot be computed without knowing the correct response **R**. Nonetheless, the control mechanism could still compute an *expected* cost over the n_y^r alternative responses based on the model's current estimate of the likelihood of each:

$$\bar{U}(T|\mathbf{S}) = \sum_r P(Y^r(T) = r|\mathbf{S}) U(T|\mathbf{S}, r) \tag{2}$$

The optimal point in time at which to respond is the value of T that yields the maximum utility. This point in time can be characterized in a simple, intuitive manner by rearranging Equations 1 and 2. Based on the response probability distribution $P(Y^r(T) = r|\mathbf{S})$, an *estimate of response accuracy* for the current stimulus **S** at time *T* can be computed, even without knowing the correct response:

$$\bar{A}(T|\mathbf{S}) = \sum_r P(Y^r(T) = r|\mathbf{S})^2 \tag{3}$$

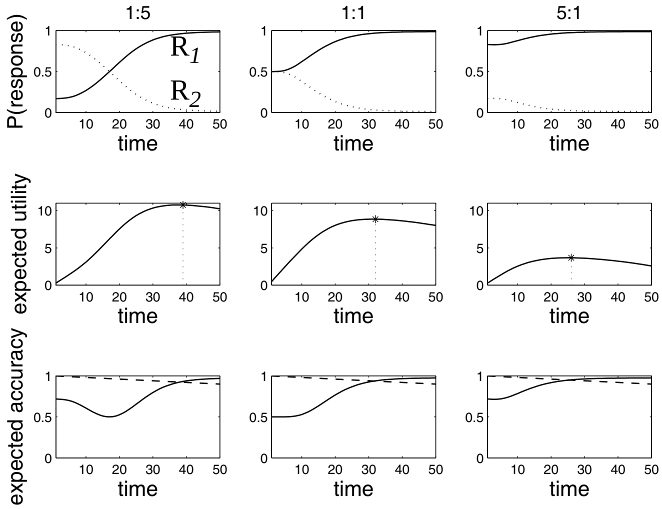

FIGURE 2. (top row) Output of PIT response pathway as a function of time when stimulus S, associated with response R_1, is presented, and relative frequency of R_1 (solid line) and the alternative response, R_2 (dotted line), is 1:5, 1:1, and 5:1. (middle row) Expected cost of responding; the asterisk shows the optimal point in time. (bottom row) PIT's internal estimate of accuracy over time (solid line) and time-decreasing criterion, $1 - \kappa T$ (dashed line)

This equation is the expectation, under the current response distribution, of a correct response assuming that the actual probability of a response being correct matches the model's internal estimate. In terms of $\bar{A}$, the optimal stopping time according to Equation 2 occurs at the earliest time T when

$$\bar{A}(T|\mathbf{S}) = 1 - \kappa T. \tag{4}$$

The optimal stopping time can be identified by examination of $\bar{A}(T|\mathbf{S})$ and T at two consecutive time steps, satisfying the essential requirement for real-time performance.

Results

Figure 2 illustrates the model's performance on the choice task when presented with a stimulus, **S**, associated with a response, R_1, and the relative frequency of R_1 and the alternative response, R_2, is 1:5, 1:1, or 5:1 (left, center, and right columns, respectively). The top row plots the probability of R_1 and R_2 against time. Although R_1 wins out asymptotically in all three conditions, it must overcome the effect of its low prior in the 1:5 condition. The middle row plots the expected utility over time. Early on, the high error rate leads to low utility; later on, reaction time leads to decreasing utility. Our rational analysis suggests that a response should be initiated at the global maximum—indicated by asterisks in the figure—implying that both the reaction time and error rate will decrease as the response prior is increased. The bottom row plots the model's estimate of its accuracy, $\bar{A}(T|\mathbf{S})$, as a function of time. Also shown is the $1 - \kappa T$ line (dashed), and it can be seen that the utility maximum is obtained when Equation 4 is satisfied.

Figure 3 presents human and simulation data for the choice task. The data consist of mean reaction time and accuracy for the two target responses, R_1 and R_2, for the three conditions corresponding to different R_1:R_2 presentation ratios. The qualities of the model giving rise to the fit can be inferred by inspection of Figure 2—namely, accuracy is higher and reaction times are faster when a response is expected.

The model provides an extremely good fit not only to the overall pattern of results, but also sequential effects. Figure reveals how the recent history of experimental trials influences reaction time and error rate. The trial *context* along the x-axis is coded as $v_4 v_3 v_2 v_1$, where v_i specifies that trial n–i required the same (``S'') or different (``D'') response as trial n–i+1. For example, if the five trials leading up to and including the current trial are—in forward temporal order—R_2, R_2, R_2, R_1, and R_1, the current trial's context would be coded as "SSDS." The correlation coefficient between human and simulation data is .960 for reaction time and .953 for error rate.

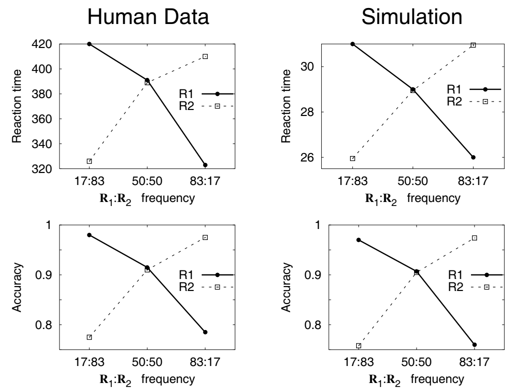

FIGURE 3. Human data (left column) and simulation results (right column) for the choice task. Human data from Jones et al. (2003). The upper and lower rows show mean reaction time and accuracy, respectively, for the two responses R_1 and R_2 in the three conditions corresponding to different R_1:R_2 frequencies.

The simple priming mechanism proposed previously by Mozer et al. (2004), which aims to adapt the model's priors rapidly to the statistics of the environment, is responsible for the model's performance: On a coarse time scale, the mechanism produces priors in the model that match priors in the environment. On a fine time scale, changes to and decay of the priors results in a strong effect of recent trial history, consistent with the human data: The graphs in Figure 4 show that the fastest and most accurate trials are clearly those in which the previous two trials required the same response as the current trial (the leftmost four contexts in each graph). The fit to the data is all the more impressive given that Mozer et al. priming mechanism was used to model perceptual priming, and here the same mechanism is used to model response priming.

Discussion

We introduced a model that accounts for sequential effects of response repetition in a simple choice task. The model was based on the principle that control processes incrementally estimate response prior probabilities. The PIT model, which performs Bayesian inference, utilizes these response priors to determine the optimal point in time at which to initiate a response. The probabilistic framework imposes strong constraints on the model and removes arbitrary choices and degrees of freedom that are often present in psychological models.

Jones et al. (2003) proposed a neural network model to address response conflict in a speeded discrimination task. Their model produces an excellent fit to the data too, but involves significantly more machinery, free parameters, and ad hoc assumptions. In brief, their model is an associative net mapping activity from stimulus units to response units. When response units R_1 and R_2 both receive significant activation,

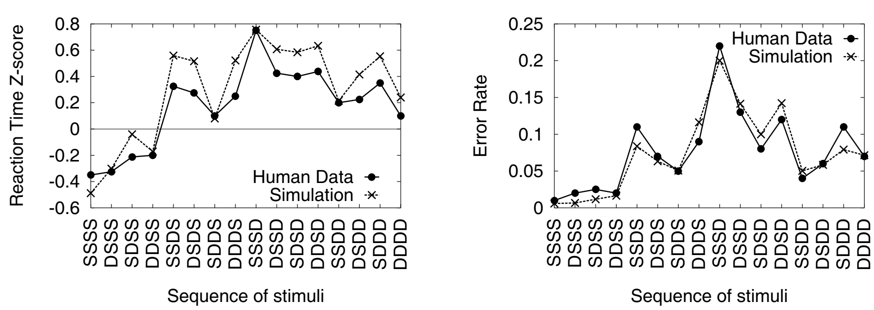

FIGURE 4. Reaction time (left curve) and accuracy (right curve) data for humans (solid line) and model (dashed line), contingent on the recent history of experimental trials.

noise in the system can push the inappropriate response unit over threshold. When this conflict situation is detected, a control mechanism acts to lower the baseline activity of response units, requiring them to build up more evidence before responding and thereby reducing the likelihood of noise determining the response. Their model includes a priming mechanism to facilitate repetition of responses, much as we have in our model. However, their model also includes a secondary priming mechanism to facilitate *alternation* of responses, which our model does not require. Both models address additional data; for example, a variant of their model predicts a neurophysiological marker of conflict called error-related negativity (Yeung, Botvinick, & Cohen, 2000).

Jones et al. (2003) also performed an fMRI study of this task and found that anterior cingulate cortex (ACC) becomes activated in situations involving response conflict. Specifically, when one stimulus occurs infrequently relative to the other, event-related fMRI response in the ACC is greater for the low frequency stimulus. According to the Jones et al. model, the role of ACC is to conflict detection. Our model allows for an alternative interpretation of the fMRI data: ACC activity may reflect the expected utility of decision making on a fine time grain. Specifically, the ACC may provide the information needed to determine the optimal point in time at which to initiate a response, computing curves such as those in the bottom row of Figure 2. If ACC activity is related to the height of the utility curves, then fMRI activation—which reflects a time integral of the instantaneous response—should be greater when the response prior is lower, i.e., when conflict is present. Recent neuropsychological data has shown a deficit in performance with a simple RT task following ACC damage (Fellows & Farah, 2005). These data are consistent with our interpretation of the role of ACC, but not with the conflict-detection interpretation.

SEQUENTIAL EFFECTS INVOLVING TASK DIFFICULTY

In this section, we return to the sequential dependency on item difficulty described in the introduction to the chapter. To remind the reader, Table 1 shows three columns of addition problems. Some problems are intrinsically easier than others, e.g., 10+3 is easier than 5+8, whether due to practice or the number of cognitive operations required to determine the sum. By definition, individuals have faster RTs *and* lower error rates to easy problems. However, when items are presented in a sequence or *block*, reaction time (*RT*) and error rate to an item depend on the composition of the block. When presented in a mixed block (column 3 of Table 1), easy items slow down relative to a pure block (column 1 of Table 1) and hard items speed up relative to a pure block (column 2 of Table 1). However, the convergence of RTs for easy and hard items in a mixed block is not complete. Thus, RT depends both on the stimulus type and the composition of the block.

This phenomenon, sometimes called a *blocking effect*, occurs across diverse paradigms, including naming, arithmetic verification and calculation, target search, and lexical decision (e.g., Lupker, Brown & Columbo, 1997; Lupker, Kinoshita, Coltheart, & Taylor, 2003; Taylor & Lupker, 2001). It is obtained when stimulus or response characteristics alternate from trial to trial (Lupker et al., 2003). Thus, the blocking effect is not associated with a specific stimulus or response pathway. Because blocking effects influence the speed-accuracy trade off, they appear to reflect the operation of a fundamental form of cognitive control—the mechanism that governs the initiation of a behavioral response. The blocking effect shows that control of response initiation depends not only on information from the current stimulus, but also on recent stimuli in the trial history.

Explaining the Blocking Effect

Any explanation of the blocking effect must specify how response-initiation processes are sensitive to the composition of a block. Various mechanisms of control adaptation have been proposed, including: domain specific mechanisms (Rastle & Coltheart, 1999; Meyer, Roelofs, & Levelt, 2003), adjustment of the rate of processing (Kello & Plaut, 2003), adjustment of an evidence criterion in a random walk model (e.g., Strayer & Kramer, 1994b). In Mozer and Kinoshita (in preparation), we present a detailed critique of these accounts.

We propose an alternative account. By this account, response-initiation mechanisms are sensitive to the statistical structure of the environment for the following reason. An accurate response can be produced only when when the evidence reaching the response stage from earlier stages of processing is reliable.

Because the point in time at which this occurs will be earlier or later depending on item difficulty, some estimate of the difficulty is required. This estimate can be explicit or implicit; an implicit estimate might indicate the likelihood of a correct response at any point in time given the available evidence. If only noisy information is available to response systems concerning the difficulty of the current trial, a rational strategy is to increase reliability by incorporating estimates of difficulty from recent—and presumably similar—trials.

We elaborate this idea in a mathematical model of response initiation. The model utilizes the PIT framework described previously to characterize the temporal dynamics of information processing, and the optimal decision criterion used for response initiation. As described earlier, PIT proposes that the transmission of stimulus information to response systems is gradual and accumulates over time, and that control mechanisms respond at the point in time that maximizes a utility measure that depends on both expected accuracy and time. In the previous model we described, we assumed that the response distribution is available for control processes to estimate the expected accuracy, $\bar{A}$ (Equation 3, and depicted in the bottom row of Figure 2, solid lines). However, if the response distribution obtained is noisy, $\bar{A}$ will be a high variance estimate of accuracy. Rather than relying solely on $\bar{A}$, the variance can be lowered by making the ecological assumption that the environment is relatively constant from one trial to the next, and therefore, the estimates over successive trials can be averaged.[4]

We use the phrase *current accuracy trace (CAT)* to denote the complete time-varying trace of $\bar{A}$, i.e.,

$$\mathrm{CAT} \equiv \{\bar{A}(t|S),\, t = 1 \ldots T\}\,.$$

To implement averaging over trials, the model maintains a *historical accuracy trace (HAT)*, and the trace used for estimating utility—the *mean accuracy trace (MAT)*—is a weighted average of CAT and HAT, i.e.,

$$\mathrm{HAT}(n) \;=\; \lambda\mathrm{CAT}(n-1) + (1-\lambda)\mathrm{HAT}(n-1)\,,$$

where n is an index over trials, and

$$\mathrm{MAT}(n) \;=\; \theta\mathrm{CAT}(n) + (1-\theta)\mathrm{HAT}(n)\,;$$

λ and θ are averaging weights. Figure 5a depicts the CAT, HAT, and MAT. The thin and thick solid curves represent CATs for easy and hard trials, respectively; these same curves also represent the MATs for pure blocks. The dotted curve represents the expected HAT in a mixed block—an average of easy and hard CATs. The thin and thick dashed curves represent the MATs for easy and hard trials in a mixed block, respectively, formed by averaging the HAT and corresponding CAT. Because the CAT and HAT are time-varying functions, the notion of averaging is ambiguous; possibilities include averaging the accuracy of

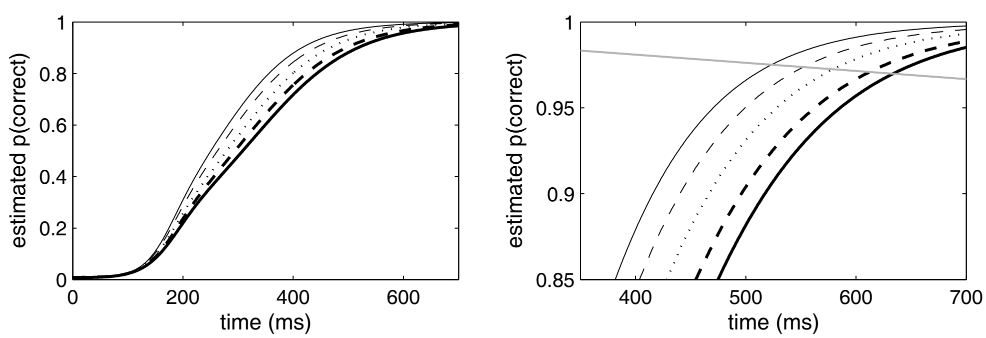

FIGURE 5. (a) easy CAT (thin solid line, which is also the MAT in a pure block of easy items) and hard CAT (thick solid, which is also the MAT in a pure block of hard items), mixed block HAT (dotted), and easy and hard MAT in a mixed block (thin and thick dashed) (b) close up of the traces, along with the time threshold (grey solid)

4. The claim of noise in the response system's estimation of evidence favoring a decision is also made in what is perhaps the most successful model of decision processes, Ratcliff's (1979) diffusion model. This assumption is reflected both in the diffusion process itself, and the assumption of trial to trial variability in drift rates.

points with the same time value and times of points with the same accuracy value. It turns out that the choice has no qualitative impact on the simulation results we present.

Results

Figure 5b provides an intuition concerning the model's ability to replicate the basic blocking effect. The mean RT for easy and hard items in a pure block is indicated by the point of intersection of the CAT with the time threshold (Equation 4). The mean RT for easy and hard items in a mixed block is indicated by the point of intersection of the MAT with the time threshold. The easy item slows down, the hard item speeds up. Because the rate of processing is not affected by the blocking manipulation, the error rate will necessarily drop for easy items and rise for hard items. Although the RTs for easy and hard items come together, the convergence is not complete as long as $\theta > 0$.

A signature of the blocking effect concerns the relative magnitudes of easy-item slow down and hard-item speed up. Significantly more speed up than slow down is never observed in experimental studies. The trend is that speed up is less than slow down—indeed, some studies show no reliable speed up—although equal magnitude effects are observed. Empirically, the model we propose never yields more speed up than slow down. The slow down is represented by the shift of the easy MAT in mixed versus pure blocks (the thin dashed and thin solid lines in Figure 5b, respectively), and the speed up is represented by the shift of the hard MAT in mixed versus pure blocks (the thick dashed and thick solid lines). Comparing these two sets of curves, one observes that the hard MATs hug one another more closely than the easy MATs, at the point in time of response initiation. The asymmetry is due to the fact that the easy CAT reaches asymptote before the hard CAT. Blocking effects are more symmetric in the model when responses are initiated at a point when both easy and hard CATs are ascending at the same rate. The invalid pattern of more speed up than slow down will be obtained only if the hard CAT is more negatively accelerated than the easy CAT at the point of response initiation; but by the definition of the easy and hard items, the hard CAT should reach asymptote after the easy CAT, and therefore should never be more negatively accelerated than the easy CAT.

The theory thus explains the key phenomena of the blocking effect. The theory is also consistent with three additional observations: (1) Blocking effects occur across a wide range of tasks, and even when tasks are switched trial to trial; and (2) blocking effects occur even in the absence of overt errors; and (3) blocking effects occur only if overt responses are produced; if responses are not produced, the response-accuracy curves need not be generated, and the averaging process that underlies the effect cannot occur.

Beyond providing qualitative explanations for key phenomena, the model fits specific experimental data. Taylor and Lupker (2001, Expt. 1) instructed participants to name high frequency words (easy items) and nonwords (hard items). Table 3 compares mean RTs and error rates for human participants and the simulation. One should not be concerned with the error-rate fit, because measuring errors in a naming task is difficult and subjective. (Over many experiments, error rates show a speed-accuracy trade off.) Taylor and Lupker further analyzed RTs in the mixed block conditional on the context—the 0, 1, and 2 preceding items. Figure 6 shows the RTs conditional on context. The model's fit is excellent. Trial *n* is most influenced by trial *n–1*, but trial *n–2* modulates behavior as well; this is well modeled by the exponentially decaying HAT.

Simulation details

Parameters of the PIT model were chosen to obtain pure-block mean RTs comparable to those obtained in the experiment and asymptotic accuracy of 100% for both easy and hard items. We added noise to the transmission rates to model item-to-item and trial-to-trial variability, but found that this did not affect the expected RTs and error rates. We fixed the HAT and MAT averaging terms, λ and θ, at 0.5, and picked κ to

TABLE 3. Expt. 1 of Taylor & Lupker (2001): Human data and simulation

	Human Data			Simulation		
	Pure	**Mixed**	**Difference**	**Pure**	**Mixed**	**Difference**
Easy	519 ms (0.6%)	548 ms (0.7%)	29 ms (0.1%)	524 ms (2.4%)	555 ms (1.7%)	31 ms (–0.7%)
Hard	631 ms (2.9%)	610 ms (2.9%)	–21 ms (0.0%)	634 ms (3.0%)	613 ms (3.7%)	–21 ms (0.7%)

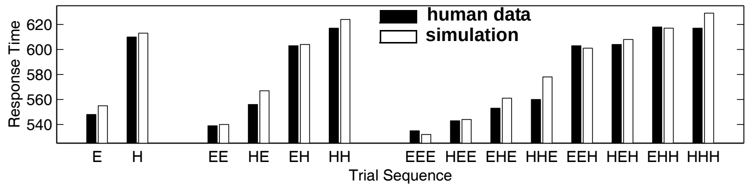

FIGURE 6. RTs from human subjects (black) and simulation (white) for easy and hard items in mixed block, conditional on 0, 1, and 2 previous item types. Last letter in a string indicates the current trial and first letters indicate context. Thus, "EHH" means a hard item preceded by another hard item preceded by an easy item.

TABLE 4. Context experiment: Human data and simulation

	Human Data			Simulation		
	Same Context	Diff. Context	Switch Effect	Same Context	Diff. Context	Switch Effect
Easy	432 ms	488 ms	56 ms	437 ms	493 ms	56 ms
Hard	514 ms	467 ms	–47 ms	514 ms	470 ms	–44 ms

obtain error rates in the pure block of the right order. Thus, the degrees of freedom at our disposal were used for fitting pure block performance; the mixed block performance (Figure 6) emerged from the model.

Testing model predictions

In the standard blocking paradigm, the target item is preceded by a context in which roughly half the items are of a different difficulty level. We conducted a behavioral study in which the context was maximally different from the target. Each target was preceded by a context of ten items of homogeneous difficulty, either the *same* or *different* difficulty as the target. This study allows us to examine the asymptotic effect of context switching. We performed this study for two reasons. First, Taylor and Lupker (2001) obtained results suggesting that a trial was influenced by only the previous two trials; our model predicts a cumulative effect of all context, but diminishing exponentially with lag. Second, several candidate models we explored predict that with a strong context, speed up of hard is significantly larger than slow down of easy; the model we've described does not.

The results are presented in Table 4. The model provides an excellent fit to the data. Significantly larger context effects are obtained than in the previous simulation (~50 ms in contrast to ~25 ms), and—given the strong context—the easy items become slower than the hard (although this effect is not statistically reliable in the experimental data). Further, both data and model show more slow down than speed up, a result that allowed us to eliminate several competing models.[5]

We have conducted a variety of other behavioral experiments testing predictions of the model. For example, in Kinoshita and Mozer (in press), we explore the conditions giving rise to symmetric versus asymmetric blocking effects. We have also shown that various other phenomena involving blocked performance comparisons can be interpreted as blocking effects (Mozer & Kinoshita, in preparation).

CONCLUSIONS

Theories in cognitive science often hand the problem of cognitive control to an unspecified homunculus. Other theories consider cognitive control in terms of a central, unitary component of the cognitive architecture. In contrast, we view cognitive control as a collection of simple, specialized mechanisms. We described two such mechanisms in this chapter, one that determines the predisposition to produce specific

5. For this simulation, we fit parameters of the PIT model to the same-context results. We also treated the MAT averaging constant, θ, as a free parameter on the rational argument that this parameter can be tuned to optimize performance: if there is not much variability among items in a block, there should be more benefit to suppressing noise in the CAT using the HAT, and hence θ should be smaller. We used 0.35 for this simulation, in contrast to 0.5 for the first simulation.

responses, and another that determines how long to wait following stimulus onset before initiating a response. We characterized the nature and adaptation of these bottom-up control mechanisms by accounting for two types of sequential dependencies. The central claim of our accounts is that bottom-up cognitive control constructs a predictive model of the environment—response priors in one case, item difficulty in the other case—and then uses this model to optimize performance on subsequent trials. Although we focused on mechanisms of response initiation, predictive models of the environment can be useful for determining where in the visual field to look, what features to focus attention on, and how to interpret and categorize objects the visual field (e.g., Mozer, Shettel, & Vecera, 2005).

Our accounts are based on the premise that the goal of cognition is optimal and flexible performance across a variety of tasks and environments. In service of this goal, cognition must be sensitive to the statistical structure of the environment, and must be responsive to changes in the structure of the environment. We view sequential dependencies as reflecting continual adaptation to the ongoing stream of experience, wherein each sensory and motor experience can affect subsequent behavior. Sequential dependencies suggest that learning should be understood not only in terms of changes that occur on the time scale of hours or days, but also in terms of changes that occur from individual incidental experiences that occur on the scale of seconds.

Acknowledgements

This research was supported by NSF IBN Award 9873492, NSF BCS Award 0339103, and NIH/IFOPAL R01 MH61549–01A1. This chapter greatly benefited from the the thorough reviews and critiques by Hansjörg Neth and Wayne Gray. We also thank Andrew Jones for generously providing the raw data from the Jones et al. (2003) study.

References

Bar, M., & Biederman, I. (1998). Subliminal visual priming. *Psychological Science*, *9*, 464-469.

Bock, K., & Griffin, Z. M. (2000). The persistence of structural priming: Transient activation or implicit learning? *Journal of Experimental Psychology: General*, *129*, 177–192.

Bodner, G.E., & Masson, M.E. (2001). Prime validity affects masked repetition priming: Evidence for an episodic resource account of priming. *Journal of Memory & Language*, *45*, 616–647.

Chun, M., & Jiang, Y. (1998). Contextual cueing: Implicit learning and memory of visual context guides spatial attention. *Cognitive Psychology*, *36*, 28–71.

Chun, M., & Jiang, Y. (1999). Top-down attentional guidance based on implicit learning of visual covariation. *Psychological Science*, *10*, 360–365.

Colagrosso, M. (2004). A Bayesian cognitive architecture for analyzing information transmission in neocortex. Unpublished Ph.D. Dissertation. Boulder, CO: Department of Computer Science, University of Colorado.

Colagrosso, M. D., & Mozer, M. C. (2005). Theories of access consciousness. In L. K. Saul, Y. Weiss, & L. Bottou (Eds.), *Advances in Neural Information Processing Systems 17* (pp. 289-296). Cambridge, MA: MIT Press.

Fellows, L., & Farah, M. (2005). Is anterior cingulate cortex necessary for cognitive control? *Brain*, *128*, 788–796.

Jentzsch, I., & Sommer, W. (2002). Functional localization and mechanisms of sequential effects in serial reaction time tasks. *Perception & Psychophysics*, *64*, 1169–1188.

Johns, E. E. & Mewhort, D. J. K. (2003). The effect of feature frequency on short-term recognition memory. *Memory & Cognition*, *31*, 285–296.

Kello, C.T. & Plaut, D.C. (2003). Strategic control over rate of processing in word reading: A computational investigation. *Journal of Memory and Language*, *48*, 207-232.

Kiger, J.I., & Glass, A.L. (1981). Context effects in sentence verification. *Journal of Experimental Psychology: Human Perception and Performance, 7,* 688-700.

Kinoshita, S., & Mozer, M. C. (in press). The effects of recent experience: Symmetric versus asymmetric frequency blocking effects. *Memory and Cognition.*

Lockhead, G. R. (1984) Sequential predictors of choice in psychophysical tasks. In S. Kornblum and J. Requin (Eds.), *Preparatory States and Processes,* Hillsdale, N.J., Erlbaum.

Lockhead, G. R. (1995) Context Determines Perception. In F. Kessel (Ed.) *Psychology, Science, and Human Affairs: Essays in Honor of William Bevan* (pp. 125-137). New York, NY: Westview Press.

Lupker, S.J., Brown, P., & Colombo, L. (1997). Strategic control in a naming task: Changing routes or changing deadlines? *Journal of Experimental Psychology: Learning, Memory, and Cognition, 23*, 570–590.

Lupker, S.J., Kinoshita, S., Coltheart, M., & Taylor, T. E. (2003). Mixing costs and mixing benefits in naming words, pictures, and sums. *Journal of Memory and Language*, *49*, 556–575.

Maloney, L. T., Dal Martello, M. F., Sahm, C., & Spillmann, L. (2005). Past trials influence perception of ambiguous motion quartets through pattern completion. *Proceedings of the National Academy of Sciences*, *102*, 3164–3169.

Maljkovic, V. & Nakayama, K. (1996). Priming of popout: II. Role of position. *Perception & Psychophysics, 58*, 977–991.

Meyer, A.S., Roelofs, A., & Levelt W.J.M. (2003). Word length effects in object naming: The role of a response criterion. *Journal of Memory and Language*, *48*, 131–147.

Mozer, M. C., & Kinoshita, S. (in preparation).Control of the speed-accuracy trade off in sequential speeded-response tasks: Mechanisms of adaptation to the stimulus environment.

Mozer, M.C., Colagrosso, M.D., & Huber, D.E. (2003). Mechanisms of long-term repetition priming and skill refinement: A probabilistic pathway model. In *Proceedings of the Twenty Fifth Annual Conference of the Cognitive Science Society.* Hillsdale, NJ: Erlbaum Assoc.

Mozer, M. C., Shettel, M., & Vecera, S. P. (2005). Control of visual attention: A rational account. *Submitted for publication.*

Posner, M. I. (1980). Orienting of attention. *Quarterly Journal of Experimental Psychology, 32*, 3-25.

Rastle, K., & Coltheart, M. (2000). Lexical and nonlexical print-to-sound translation of disyllabic words and nonwords. *Journal of Memory and Language*, *42*, 342-364.

Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review,* 85, 59-108.

Ratcliff, R., & McKoon, G. (1997). A counter model for implicit priming in perceptual word identification. *Psychological Review, 104,* 319–343.

Rogers R. D., & Monsell S. (1995). Costs of a predictible switch between simple cognitive tasks. *Journal of Experimental Psychology: General*, *124*, 207-231

Stewart, N. Brown, G. D. A., & Chater, N. (2002). Sequence effects in categorization of simple perceptual stimuli. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 28*, 3-11.

Strayer, D. L., & Kramer, A. F. (1994a). Strategies and automaticity I: Basic findings and conceptual framework. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 20,* 318-341.

Strayer, D. L., & Kramer, A. F. (1994b). Strategies and automaticity II: Dynamic aspects of strategy adjustment. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *20*, 342-365.

Taylor, T.E., & Lupker, S.J. (2001). Sequential effects in naming: A time-criterion account. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *27*, 117–138.

Usher, M., & McClelland, J.L. (2001). On the time course of perceptual choice: The leaky competing accumulator model. *Psychological Review, 108*, 550-592.

Vecera, S. P. (2005) Sequential effects in figure-ground assignment. Manuscript in preparation.

Wolfe, J. M., Butcher, S. J., Lee, C., & Hyle, M. (2003). Changing your mind: On the contributions of top-down and bottom-up guidance in visual search for feature singletons. *Journal of Experimental Psychology: Human Perception and Performance*, *29*, 483–502.

Yeung, N., Botvinick, M. M., & Cohen, J. D. (2000). The neural basis of error detection: Conflict monitoring and the error-related negativity. Submitted for publication.