



Computational Rationality: Linking Mechanism and Behavior Through Bounded Utility Maximization

Richard L. Lewis,^a Andrew Howes,^b Satinder Singh^c

^a*Department of Psychology, University of Michigan*

^b*School of Computer Science, University of Birmingham*

^c*Computer Science and Engineering, University of Michigan*

Received 11 December 2012; received in revised form 8 July 2013; accepted 16 July 2013

Abstract

We propose a framework for including information-processing bounds in rational analyses. It is an application of *bounded optimality* (Russell & Subramanian, 1995) to the challenges of developing theories of mechanism and behavior. The framework is based on the idea that behaviors are generated by cognitive mechanisms that are adapted to the structure of not only the environment but also the mind and brain itself. We call the framework *computational rationality* to emphasize the incorporation of computational mechanism into the definition of rational action. Theories are specified as *optimal program problems*, defined by an adaptation environment, a bounded machine, and a utility function. Such theories yield different classes of explanation, depending on the extent to which they emphasize adaptation to bounds, and adaptation to some ecology that differs from the immediate local environment. We illustrate this variation with examples from three domains: visual attention in a linguistic task, manual response ordering, and reasoning. We explore the relation of this framework to existing “levels” approaches to explanation, and to other optimality-based modeling approaches.

Keywords: Cognitive modeling; Rational analysis; Bounded optimality; Utility maximization; Bounded rationality; Cognitive architecture; Rationality

1. Introduction: Rational analyses and information-processing bounds

Top-down, rational analyses in the cognitive and biological sciences offer the potential for deep “why” explanations of behavior by appealing to assumptions about goals, adaptive pressures, or functions that should be computed by brains to generate effective behavior (e.g., Anderson, 1990; Griffiths, Chater, Norris, & Pouget, 2012; Marr, 1982;

Correspondence should be sent to Richard L. Lewis, Department of Psychology, University of Michigan, Ann Arbor, MI 48109. E-mail: rickl@umich.edu

Stephens & Krebs, 1986). These assumptions can be understood as defining problems of adaptive or rational behavior, and are formulated independently of the mechanisms that might be used to solve them (“how” explanations). This view is reflected in a key property of most meta-theoretical frameworks in cognitive science: a separation of higher rational levels that identify goals from lower mechanism levels that approximately implement those goals (Marr, 1982; Newell, 1982). In such accounts, observed gaps between how the agent *should* behave in some environment and how it *actually* behaves are explained in one of two ways. One way is to appeal to mechanism limitations that preclude the computation of optimal behavior (as in approaches emphasizing heuristics for choice and problem solving; e.g., Gigerenzer & Selten, 2001; Newell & Simon, 1972; Simon, 1956, 1990). Another way is to redefine the problem of rational behavior, by positing goals or environments of adaptation different from the immediate local environment (as in many prominent rational analyses, and evolutionary psychology; Anderson & Schooler, 1991; Cosmides, Barrett, & Tooby, 2010; Oaksford & Chater, 1994). We depict this standard view schematically in Fig. 1.

Newell and Simon’s (1972) seminal work on problem solving provides a clear exposition of the distinction between rational and mechanistic explanations. They state that “... one of our main tasks will be to understand what is demanded by the task

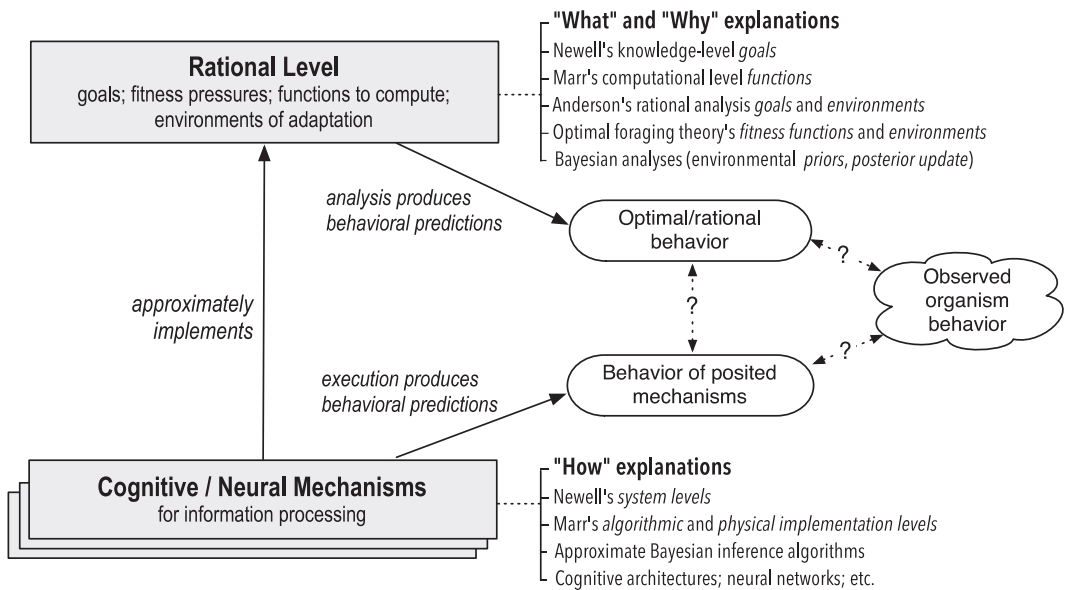


Fig. 1. The standard view in cognitive science of how rational analysis links behavior and mechanism. Rational analyses may be conducted via a variety of techniques, including Bayesian analysis, and can produce predictions of optimal behavior that may be compared to observed behavior. Rational analysis levels are the locus of “what” and “why” accounts and potentially explain behavior in a way that abstracts away from mechanism. Cognitive/neural mechanisms are assumed to provide approximate implementations of the goals/functions posited at the rational level. The execution of these mechanisms produces behaviors that may be compared with both observed behaviors and the behavior calculated by the rational analysis.

environment, so that we can understand—by elimination—what aspects of behavior are determined by the psychology of the problem solver” (p. 79). For example, the behavior of a human making correct moves in an end-game of chess (as derived, perhaps, by an unbounded minimax algorithm) may be explained by what is demanded by the task and the goal of winning, that is, a rational analysis. Departures from correct moves, or systematic biases in selecting among equally good moves, might be explained by appeal to a set of information processing mechanisms that approximate the function of winning—for example, a bounded heuristic search process. These two kinds of explanations—rational and mechanistic—represent the dominant ways of theorizing in cognitive science.

The purpose of this paper is to make explicit and illustrate an alternative to this standard view—one that has significant theoretical advantages, while retaining the benefits of rational analysis. The distinctive feature of this alternative is that it allows for information-processing capacities and bounds to be included as first-class elements of definitions of rational behavior, formalized as problems of utility maximization. Fig. 2 illustrates this alternative schematically. The early precedents of this alternative are signal detection theory (Tanner & Swets, 1954)¹ and its development as ideal observer analysis (Geisler, 2011); one of Simon’s (1955) early candidate definitions of bounded rationality²; Baron and Kleinman’s (1969) optimal control models of visual attention and manual control³; and Anderson’s (1990) original method of rational analysis.⁴

We call this alternative *computational rationality* to emphasize the incorporation of computational mechanism into the definition of rational action. The framework allows a rigorous exploration of the idea that behaviors are generated by cognitive mechanisms that are adapted to the structure of not only the environment but also the mind itself. It places utility maximization at the heart of psychological theory, and we think it provides

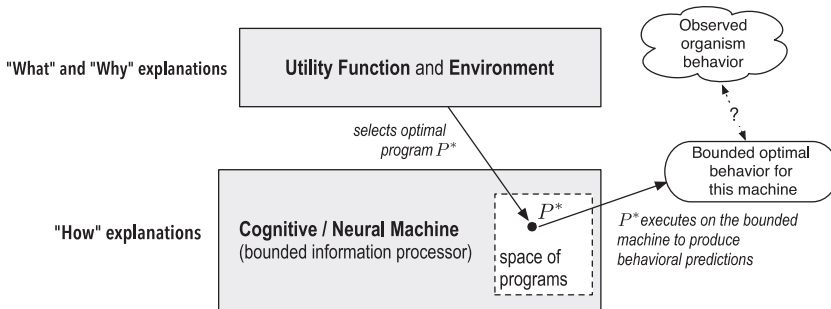


Fig. 2. Computational rationality: an alternative view of how rational analysis links mechanism and behavior, based on *bounded optimality*. Unlike the standard approach in Fig. 1, the rational analysis makes direct contact with information-processing mechanisms by selecting an optimal program for a bounded machine that maximizes utility in some environment. The execution of the optimal program generates behavior that is the optimal behavior for that machine; it cannot do better in the given environment. Although the output of the analysis is an optimum, it is nevertheless directly constrained by the posited information-processing bounds, and the processes of optimization used in the analysis are not ascribed to the agent (rather they are tools of the theorist). The bounded optimal behavior may be compared to observed behavior as well as behavior calculated by unbounded rational analyses.

a useful way to unify psychology's dual aims of generating explanations of behavior and generating explanations of cognitive mechanism. In a nutshell, the framework extends the question posed by standard rational approaches (what should a rational agent do in this environment?) to include processing bounds: What should a rational (utility-maximizing) agent, *with its available information-processing mechanisms*, do in this environment?

The technical innovation offered here is to apply *bounded optimality*—a precisely defined concept in artificial intelligence (Russell & Subramanian, 1995)—to these explanatory aims. A bounded optimality analysis provides a clear answer to the question above: A bounded agent should do whatever the best program running on its information-processing architecture would do. We believe that this framing has value because of its generality and the analytic rigor and clarity that it brings to the problems of stating and testing the implications of different theoretical assumptions. We introduce the term *computational rationality* to refer to the application of bounded optimality to psychology, to avoid confusion with bounded rationality, which has come to be associated with a rejection of optimality analyses.

In the remainder of the paper, we first present the framework of computational rationality, and show how it yields a broad space of useful explanatory theory. Theories and explanations in this space differ in the extent to which they emphasize adaptation to information-processing bounds, and adaptation to some ecological environment that differs from the immediate local environment. Rich mixtures of both kinds of constraint are possible. We introduce a minimal amount of formalism along the way, focusing on defining the key underlying concepts. Next, we illustrate the framework with three models in the domains of response ordering, eye movements in a linguistic task, and logical reasoning, each of which illustrates a different kind of explanation. We conclude with a discussion of the relation of computational rationality to existing approaches in cognitive science.

2. Computational rationality as a framework for explanation and prediction

In this section, we apply computational rationality to the challenge of theorizing about human cognitive mechanisms, behaviors, and ecologies. Computational rationality is strongly influenced by bounded optimality (Russell & Subramanian, 1995). Russell and Subramanian (1995) defined a bounded optimal agent as one that is executing the program (from its space of possible programs) that maximizes some performance measure in some environments of interest. Understanding computational rationality therefore requires a *bounded rational analysis*. It requires defining and solving an *optimal program problem*. This analysis provides an ontology of analytic elements that precisely defines such notions as agent bounds, goals, behaviors, and distinct ecological and evaluation environments. This ontology is summarized in Table 1, with examples of how the different elements map onto concepts from three familiar frameworks: signal detection theory, cognitive architectures, and classic rational analysis.

Table 1
The elements of a bounded rational analysis, with examples of how they might be instantiated in three familiar theoretical frameworks

Examples From Familiar Frameworks			
Element and Definition	Signal Detection Theory	Cognitive Architectures	Classic Rational Analysis
Bounded agent, M: An information-processing machine with a set of possible observations \mathcal{O}^M and actions \mathcal{A}^M . The machine computes mappings from histories of observations to actions, determined by the program executing on the machine	Thresholded decision machine with $\mathcal{O}^M =$ continuous signal range, subject to some noise, and $\mathcal{A}^M = \{\text{yes, no}\}$ responses	Perceptual, motor, cognitive processes; memories, often including a working memory and a production rule memory	Defines only \mathcal{O}^M and \mathcal{A}^M without defining the machinery that computes mappings from histories to actions
Bounded agent program space, \mathcal{P}^M: The space of possible programs that can run on machine M	Detection thresholds	Set of all possible production rule programs	None
Bounded agent program, P: One of the programs in \mathcal{P}^M . The program P and machine M together define an agent model $\langle M, P \rangle$ that determines the policy computed by M	A specific choice of threshold	A specific production rule program	None
Policy or strategy, $C^{(M,P)}$: A mapping from observation histories to (probabilistic) choice of actions. A bounded machine with its program space \mathcal{P}^M defines a space of possible policies (mappings) that is a strict subset of all the possible policies (mappings)	Function mapping noisy stimulus signals to Yes/No responses	Computed mapping from observations to task responses	Mapping from observations to task responses
Evaluation environment, E_{eval}: The (task) environment in which the agent exhibits the behavior to be explained. An environment must define a distribution over next possible observations given a history of observations and actions	An experimental visual detection experiment specifying the frequencies of targets and distractors.	An experimental task requiring some interaction with a computer via mouse and keyboard (e.g., see Section 3.1)	A local experimental task defined in terms of a set of abstract observations and responses (e.g., see Section 3.3)
Ecological environments \mathcal{E} and distributions, $P(\mathcal{E})$: The environment or distribution of environments used to select aspects of the optimal program for M	Usually identical to the evaluation environment	Usually identical to the evaluation environment	Distributions of typically encountered objects outside the local environment E_{eval}

(continued)

Table 1. (continued)

Examples From Familiar Frameworks			
Element and Definition	Signal Detection Theory	Cognitive Architectures	Classic Rational Analysis
Behaviors or histories, h : Sequences of observations and actions generated by executing a policy in an environment	Yes or no responses to a stimulus	Task-related motor responses in some experiment	Abstract local task responses
Utility (or goals), \mathcal{U} : A function that maps histories (behaviors) to some scalar measure of goodness	The precise payoff matrix for hits, misses, false alarms, correct rejections	Often specified only informally in the task environment	Information gain (e.g. see Section 3.3)
Bounded optimal program, P^* : The program in \mathcal{P}^M that maximizes expected utility in the ecological environments	The optimal threshold given the payoff matrix and noise	Usually not derived, but if \mathcal{U} and \mathcal{E} are made explicit, P^* is well-defined as the optimal production rule program	None
Bounded optimal policy, G^* : The mapping from histories to actions computed by the bounded optimal program P^*	The mapping from noisy signals to responses implied by the optimal threshold	The mapping from observations to motor acts computed by optimal production rule program	None
Bounded optimal behavior : The behavior in E_{eval} generated by M running the bounded optimal program P^*	Correct and error response frequencies for both targets and distractors	Reaction times, error rates, etc. in the task environment	Correct and error responses in the local task environment

2.1. Defining and solving optimal program problems

Optimal program problems demand three inputs (Fig. 2): (a) a *bounded agent*; (b) an *adaptation*, or *ecological, environment*; and (c) a *utility function*. The solution to an OPP is (d) the *optimal program(s)* that, when executed on the bounded agent, will maximize utility in the given adaptation environment. All four elements may carry theoretical content and be put to a variety of uses, as described next.

2.1.1. Bounded information-processing agents

By information-processing bounds or mechanisms, we mean the machinery that computes mappings from perceptual inputs to effector outputs. This definition is broad enough to include constraints that are associated with both cognitive and perceptual/motor systems, such as memory limitations and perceptual or motor noise (Faisal, Selen, & Wolpert, 2008), or the duration of primitive computations. Thus, nothing hinges on a distinction between “cognitive” and other types of bounds. But the definition is narrow enough to exclude machines that do not map between perception and action. One example of a class of machines satisfying the definition are *cognitive architectures* in the sense advocated by Newell (1990), Meyer and Kieras (1997), and Anderson (2007).

2.1.1.1. Definitions: We define a bounded agent as a machine M with a space of possible observations \mathcal{O}^M , a space of possible actions \mathcal{A}^M , and a space of programs \mathcal{P}^M that can run on the machine. The agent is thus a machine that interacts with the environment via a set of sensors providing \mathcal{O}^M and effectors realizing \mathcal{A}^M . Depending on the kind of model being built by the scientist, the space of observations could be low level (e.g., retinal images) or high level (e.g., inputs already abstractly categorized). The space of actions can be external motor actions (e.g., low-level muscle control signals) or high-level abstract actions (e.g., a plan to travel to a distant city). The programs \mathcal{P}^M may include internal cognitive-processes—“mental actions” such as rehearsal, visual imagery and rotation, or memory retrievals.

Choosing a program $P \in \mathcal{P}^M$ specifies an agent-model $\langle M, P \rangle$. When an agent-model interacts with some environment, it generates a history or trajectory of observations and actions. More precisely, it generates a *random history* at time t , denoted h_0^t , which is the sequence of alternating observations and actions from the beginning of time, that is, $h_0^t = o_0, a_0, o_1, a_1, \dots, o_{t-1}, a_{t-1}, o_t$. We call such histories the *behaviors* of the agent-model. (We assume the observation-action cycle happens in discrete time for ease of exposition only.)

An agent-model intensionally defines a *policy* or a mapping $G^{(M, P)}$ from histories to distributions over actions, that is, for all t

$$G^{(M, P)} : (OA)^t \mathcal{O} \times A \rightarrow [0, 1]. \quad (1)$$

We will henceforth abbreviate $G^{(M, P)}$ simply as G . Thus, the probability of action a in history h_0^t is $G(h_0^t, a)$. Each machine-program combination induces a particular G -mapping.

In the context of a standard experimental cognitive psychology task, a policy can be thought of as a *strategy* for performing the task.

We can now precisely state the distinction between bounded and unbounded machines. We denote the space of all computable mappings from histories to distributions over actions as $\bar{\mathcal{G}}$. This represents the *unbounded* space of all agent-models that have observation set \mathcal{O}^M and action set \mathcal{A}^M . $\bar{\mathcal{G}}$ is defined only in terms of \mathcal{O}^M and \mathcal{A}^M , and makes no reference to—and thus is not constrained by—the mechanisms in M . But any *bounded* machine M will implicitly select out a strict subset of $\bar{\mathcal{G}}$, which we denote as the (implicitly bounded) subset \mathcal{G} .

At first glance, the notion that an agent-model is a fixed mapping (Eq. 1) seems rather limiting, but, in fact, this is the most general definition possible. It admits of any agent whatsoever (within the constraints of discrete time interaction that can also be relaxed). In particular, it allows the agent-model's behavior to be arbitrarily influenced by history, thus allowing for all manner of learning.⁵

2.1.1.2. Theoretical uses: The machine M may play a number of theoretical roles. It may define a comprehensive cognitive architecture, in the sense of ACT-R (Anderson, 2007), EPIC (Meyer & Kieras, 1997), or Soar (Laird, 2012; Newell, 1990) and thus entail a huge space \mathcal{P}^M corresponding to the set of possible production rule programs. It may define a comprehensive architecture, but one whose program space is more narrowly limited to a range of policies or strategies of interest to the scientist. It may describe a machine architecture with bounds that represent assumptions about knowledge gained from previous experience. It may describe a particular instantiation or subset of a broader architecture, specifying just those bounds thought to be relevant to the analysis at hand. It may represent processing bounds at an extremely low level—perhaps resource constraints on neurotransmitters—combined with strategic adaptation parameters at a very high level.

2.1.2. Ecological environment

Any assertion of optimality is relative to the definition of some environment to which the agent has adapted, an insight extensively leveraged in work on rational analyses (e.g., Anderson & Schooler, 1991; Oaksford & Chater, 1994). The key insight here is that the optimality, or otherwise, of an adaptive system should be partly determined by the statistics of the ecological environment to which the agent has adapted. For example, Anderson and Schooler argued that forgetting is optimal, given the probability of needing the information in the ecological environment, which decays as a power-law function of time. In such analyses, behavior can be understood as optimal relative to that ecology, but it may appear to be suboptimal when observed in a local evaluation environment with differing statistical properties.

2.1.2.1. Definitions: In the computational rationality framework, distributions over histories or behaviors are determined jointly by environments and agent models. We denote histories h sampled from an agent-model $\langle M, P \rangle$ interacting with an environment E as $h \sim (\langle M, P \rangle, E)$. We denote the environment in which the agent behavior is to be

explained as E_{eval} , and we denote as \mathcal{E} the distribution of environments to which the agent is assumed to have been adapted, that is, which serves to select the set of optimal programs.

2.1.2.2. Theoretical uses: The evaluation environment E_{eval} and adaptation environments \mathcal{E} may play a number of theoretical roles. In the usual setting of an experimental cognitive psychology task, \mathcal{E} may be some mix of the immediate local environment—for example, the “training” or “practice” phase, and prior experience with the stimuli. In a classic rational analysis (e.g., the seminal analysis of conditional reasoning by Oaksford & Chater, 1994, taken up in detail below), \mathcal{E} may specify ecological distributions of objects and events thought to be encountered in the agent’s lifetime. In an evolutionary biology setting, \mathcal{E} may specify assumptions about “the environment of evolutionary adaptation”—a theoretical construct that need not be narrowly identified with a particular place or time (Tooby & Cosmides, 2005).

2.1.3. Utility functions

2.1.3.1. Definitions: We use a utility function \mathcal{U} to determine the goodness of programs that control behavior. \mathcal{U} maps histories h from the set of possible histories \mathcal{H} to some scalar:

$$\mathcal{U} : \mathcal{H} \rightarrow \mathbb{R} \quad (2)$$

Note that \mathcal{U} is predicated on behaviors—interactions with the environment—not mechanisms. The implication for mechanism is through the selection of the optimal program that maximizes expected utility, made precise below in Eq. 3.

The approach thus follows many formal approaches to modeling behavior that require utility and reward functions (e.g., optimal control theory (e.g. Stengel, 1994); dynamic programming (Bellman, 1957); decision theory (Von Neumann & Morgenstern, 1947); reinforcement learning (Sutton & Barto, 1998)).

2.1.3.2. Theoretical uses: The utility function may play a number of theoretical roles. In a standard experimental psychology setting, it may specify the agent’s *task goals*, which might be given an objective grounding in the instructions. It may also specify a theory of *internal subjective utility*, including factors such as novelty traditionally ascribed to “intrinsic motivation” (Singh, Lewis, Barto, & Sorg, 2010), temporal discounting, or sensitivity to risk. It may specify desired *speed–accuracy trade-offs*, which may have subjective or objective grounding (or both). In evolutionary biology settings, the utility function may specify assumptions about what determines organism fitness. In a cultural evolution settings, the utility function may also specify assumptions about what determines organism fitness.

2.1.4. Bounded optimal programs and behaviors

Together, an ecological environment distribution \mathcal{E} , information-processing machine M , and a utility function \mathcal{U} fully specify an optimal program problem.

2.1.4.1. *Definitions:* Solving this problem requires identifying the set of optimal programs $\mathcal{P}^{M^*} = \{P^*\}$, defined by the following bounded optimality equation, adapted from Russell and Subramanian (1995):

$$\langle M, \mathcal{P}^{M^*} \rangle = \{ \langle M, P^* \rangle \} = \arg \max_{P \in \mathcal{P}^M} \mathbb{E}_{E \sim P(\mathcal{E})} \mathbb{E}_{h \sim (\langle M, P \rangle, E)} \{ \mathcal{U}(h) \}. \quad (3)$$

where the inner expectation is over histories (behaviors) in some environment and the outer expectation is over a distribution of environments. Each $\langle M, P^* \rangle$ defines an optimal policy G^* and so the set of optimal policies $\mathcal{G}^* = \{G^*\} = \{ \langle M, P^* \rangle \}$.

Again, the identification of optimal program problems may play a number of theoretical roles that depend on the nature of the theories of the bounds, environments, and utility function. The optimization step may be taken as an *abstraction over a variety of possible natural adaptation mechanisms*, including biological and cultural evolution, and individual agent learning. In particular, it identifies the limits of those adaptive processes, given the posited bounds. For example, in the typical setting of a cognitive psychology experiment, where the utility function \mathcal{U} encodes assumptions about immediate local task goals, and the adaptation environment E may include training trials or indeed the entire evaluation environment, solving Eq. 3 is a way of abstracting over any number of adaptive mechanisms, including procedural reinforcement learning, instruction taking, local look-ahead planning, and cultural transmission. An optimal program problem is thus a way to specify a problem that has some stability over a region of time and space such that some adaptive mechanism(s) of interest may operate.

The output of the optimization ($\{ \langle M, P^* \rangle \}$) itself may be put to a variety of uses. If the cognitive mechanisms themselves are of interest, then what is important is that ($\{ \langle M, P^* \rangle \}$) constitutes a *derivation of mechanism*, and its structure may be analyzed. If strategies (policies) are of interest, then what is important is that $\{ \langle M, P^* \rangle \}$ computes the set of *optimal strategies* (policies) \mathcal{G}^* given the bounds, and the structure of these strategies may be analyzed. If behavior (histories) are of interest, then what is important is that $\{ \langle M, P^* \rangle \}$ *derives a distribution of behaviors* $h \sim (\langle M, P \rangle, E)$ in environments, and properties of this behavior can be analyzed, including correspondence to observed biological behaviors in some evaluation environment.

2.1.5. *Important properties of bounded rational analysis*

There are three crucial points that must be understood to avoid misinterpreting what a bounded rational analysis does:

1. **It is important to distinguish two kinds of computational costs: the cost of finding the optimal program, and the cost of executing the optimal program.** In general, the individual agent need not be assumed to have arrived at the optimal program through an internal process of “optimization” over the posited program space \mathcal{P}^M , and the method of solving Eq. 3 itself carries no theoretical content. Put

another way, the computational cost (or time- and space-complexity) of the algorithm used by the scientist to solve the OPP is *not* borne by the agent. The computational complexity of the selected optimal program *is* ascribed to the agent—because the optimal program executing on the agent’s processing architecture constitutes the cognitive mechanisms asserted to produce the agent’s behavior. The former computational cost is by definition associated with “optimization”; the latter need not be.

2. **It is important to distinguish between the *optimality of programs* and the *optimality of choices or behaviors*.** What is selected by Eq. 3 is a program (or set of programs), and not behavior—although of course the program has implications for behavior. Optimal programs can be given a precise definition that takes into account the constraints of the information-processing machine; optimal behavior cannot (without implicitly adopting the optimal program perspective). This distinction is the foundational insight of Russell and Subramanian (1995): An agent is bounded optimal if its *program* is a solution to the constrained optimization problem presented by its architecture and the task environment.
3. **The data to be explained play no role in the optimal program derivation.** Thus, it is not an analytic error to simply set $\mathcal{E} = \{E_{\text{eval}}\}$, because only \mathcal{U} , and not fit to data, is used to select P^* . All of the explanatory power of the bounded optimality analyses rests on this distinction. Nevertheless, bounded optimality analyses may be used in concert with methods for estimating parameters of the bounded machine from data. We illustrate two such methods in the examples that follow.

2.2. Varieties of explanations

Different choices of information-processing machine and environments of adaptation will lead to different kinds of optimality-based explanation. We focus here on four natural subtypes that arise by varying the optimal program problem inputs in discrete ways. In practice models and explanations may be some mix of these types, but it is useful to lay out the simplest cases. Note that classes III and IV, but not classes I and II, define classes of theory in which it is assumed that an organism is computationally rational.

2.2.1. Optimality explanations

Consider first the case where the machine M is unbounded (that is, the program space \mathcal{P}^M does not restrict the set of possible computable mappings, and so $\mathcal{G} = \bar{\mathcal{G}}$) and the adaptation environment is the evaluation environment ($\mathcal{E} = \{E_{\text{eval}}\}$). Then solving Eq. 3 derives the best possible policy that an agent with observations \mathcal{O}^M and actions \mathcal{A}^M could in principle execute in E_{eval} . To the extent that the predicted behaviors or policy structure resulting from this analysis corresponds to observed behavior, then the behavior has been explained. Indeed, it has been given the most powerful explanation possible because no appeal to machine bounds or prior environments of adaptation is required; equivalently, the behavior of the observed agent provides no evidence for such bounds (Newell & Simon, 1972).

Type I Optimality explanations can be thought of as rational analyses in which only the local evaluation environment plays a role in determining behavior.

2.2.2. *Ecological-optimality explanations*

The second class of explanation appeals to an ecology of adaptation represented by a distribution of environments $P(\mathcal{E})$ —but not to information-processing bounds. In the formal notation, $\mathcal{E} \neq \{E_{\text{eval}}\}$ but $\mathcal{G} = \bar{\mathcal{G}}$. To the extent that the predicted behaviors or policy structure correspond to what is observed, then the observed behavior has been explained. A natural explanatory requirement to impose simultaneously is that Type I Optimality explanations *fail* to produce correspondence—otherwise it can be argued that no evidence has been gained for the assumptions about \mathcal{E} . Of course, it is possible that a different distribution of environments or even a different bounded machine might have also yielded correspondence.

Type II Ecological Optimality explanations can be thought of as rational analyses in which ecological environments of adaptation have shaped behavior (Anderson, 1990; Oaksford & Chater, 1994, 2007), but processing bounds play no role.

2.2.3. *Bounded-optimality explanations*

The third class of explanation appeals to information-processing bounds: The machine M with its program space \mathcal{P}^M implies a set of policies $\mathcal{G} \subsetneq \bar{\mathcal{G}}$, that is, a strict subset of the set of all computable mappings. But the adaptation environment is the evaluation environment ($\mathcal{E} = \{E_{\text{eval}}\}$). Then solving Eq. 3 derives the best possible policy (and thus behavior) that the bounded agent M could in principle execute in E_{eval} . To the extent that the predicted behaviors, policy structure, or mechanisms correspond to what is observed, then the observations have been explained. Again, a natural explanatory requirement to impose simultaneously is that Type I Optimality explanations *fail* to produce correspondence—otherwise it can be argued that no evidence has been gained for the assumptions about the bounds of M . Of course, it is possible that a different bounded machine or even a different distribution of environments might have also yielded correspondence.

Examples of Type III explanations in the literature include classic signal detection and ideal observer analyses (Geisler, 2011; Tanner & Swets, 1954), and the perceptual-motor control models of Trommershaeuser, Maloney, and Landy (2008) and Wolpert (2007). In Section 3.1 we provide an example of a Type III analysis with a fairly rich program space over both cognitive and motor processes.

2.2.4. *Ecological-bounded-optimality explanations*

The fourth class of explanation appeals to both information-processing bounds and an ecology of adaptation distinct from the evaluation environment. In notation, $\mathcal{G} \subsetneq \bar{\mathcal{G}}$, and $\mathcal{E} \neq \{E_{\text{eval}}\}$. To the extent that the predicted behaviors, policy structure, or mechanisms correspond to what is observed, then the observations have been explained. A strong explanatory requirement to impose simultaneously is that changing the OPP by dropping either the bounds (yielding Type II), the ecology (yielding Type III), or both (yielding Type I) *fails* to produce correspondence—otherwise it can be argued that no evidence has

been gained for the assumptions about combined effects of the bounds and the environment M . Again, it is possible that a different bounded machine or even a different distribution of environments might have also yielded correspondence.

Ecological-bounded-optimality explanations offer perhaps the richest classes of psychological theory, although there are fewer clear examples in the literature in this type (however, see Geisler, 2011, for examples of optimal visual feature selection for performing object identification tasks, given the statistical properties of natural scenes). In Section 3.2 we illustrate this type of explanation with a model of eye movements in a linguistic task, where program (strategic) parameters of the model are adapted to the local task structure, an ecology of experience prior to the task, and the dynamics of the bounded processing architecture.

3. Three exemplar models

Table 2 gives an overview of the three exemplar models, summarizing the different elements of the optimal program problem specifications. The first example (Section 3.1), a model of response ordering in a dual task (Howes, Lewis, & Vera, 2009), illustrates a bounded optimal (Type III) explanation in which the implications of bounds are explored by deriving optimal strategies given different architectures for the cognitive component of response selection. The program space is a combinatoric space of possible production rule programs that may run on the different machines. Utility is fixed by the quantitative payoff used in the experiment. The second example (Section 3.2), a model of sequential eye movements in a linguistic task (Lewis, Shvartsman, & Singh, 2013), illustrates ecological bounded optimality (Type IV) explanation, in which the implications of varying both bounds and utility are explored, and in which utility is experimentally manipulated in the human experiments. The ecological adaptation is in the form of priors adapted to previous experience with English words (approximated by corpus frequencies). The third example (Section 3.3), a model of the Wason Selection task (Wason, 1966), a simple conditional reasoning task, illustrates an ecological optimality (Type II) explanation, in which the implications of both different utility functions and ecologies are explored. This is a redescription and extension of the rational analysis of this task first advanced by Oaksford and Chater (1994).

3.1. Bounded optimal response ordering (bounded optimality explanation)

In this section, we illustrate the bounded optimality explanation type by redescribing Howes et al.'s (2009) analysis of response ordering. The analysis shows that response ordering in elementary dual tasks can be predicted by defining an OPP that includes internal information-processing assumptions in the bounded machine. These processes require computational resources (that are limited), they require time to compute, and they must be scheduled. The nature of the available resources, whether serial or parallel, has been a matter of scientific debate, and we describe how a comparison of optimality-based

analyses each with a different set of resource bounds as inputs to the OPP can inform this debate.

3.1.1. *The psychological refractory period task*

Consider a situation in which a verbal response is given to the pitch of a tone and a manual response is given to a visual pattern. An experimental points regime is imposed that favors trials on which the two responses are ordered such that the verbal response is issued before the manual response, and it also favors faster responses over slower responses. Thus, smaller positive inter-response intervals have the highest reward, but because of intrinsic noise in the cognitive machine, these small intervals carry a higher risk of response reversal. Monetary awards are assigned by the experimenter in accordance with the points regime.

This task, known as a Psychological Refractory Period (PRP) task, has been used extensively in an effort to understand the bounds imposed by cognition on response selection. Performance is often thought to reveal the presence of a serial bottleneck in response selection (Pashler, 1998). However, some researchers have argued that the results reflect strategic choices that participants make in response to the task demands (Meyer & Kieras, 1997). The debate can be characterized as one in which some believe that the evidence suggests bounds and others believe that it suggests choices of program that people make in response to the utility function that people adopt, given the experimental instructions. Here, we define an OPP that captures a specific set of PRP experiments reported in Schumacher, Lauber, Glass, Zurbriggen, Gmeindl, and Kieras (1999).

3.1.2. *Step 1: Defining the elements of the optimal program problem*

3.1.2.1. *Information-processing machine and its bounds:* Here, the bounded machine is a theory of the invariant cognitive-neural architecture in which there are multiple processors, each of which is capable of executing one or more processes at any one time (Fig. 3). The duration of each process is a random variable with a Gamma distribution. Fig. 3(a) represents the constraints imposed by one theory of the bounds, ACT-R (Anderson et al., 2004), and Fig. 3(b) represents those imposed by another, EPIC (Meyer & Kieras, 1997). Each bounded machine in the figure consists of an identical set of processors and processes, but it differs in whether the cognitive processor has the capacity to select only one response at any one time (ACT-R) or multiple responses (EPIC). We denote the machines as M_{ACTR} and M_{EPIC} . The set of actions \mathcal{A}^M for both machines consists of the manual (button presses) and vocal task responses, and the set of observations \mathcal{O}^M for both machines consists of high-level categorized visual and auditory percepts (see Howes et al., 2009 for details). Thus, the only difference between the two machines is in the computational architecture that computes the mappings between observations and actions—so $\mathcal{P}^{M_{\text{ACTR}}} \neq \mathcal{P}^{M_{\text{EPIC}}}$. Consistent with the production-rule architecture of ACT-R and EPIC, Howes et al. (2009) formalize both spaces as sets of condition-action pairs, where conditions may include both the percepts \mathcal{O}^M , internal memory codes and the internal states of motor processors, and actions include both the external actions in \mathcal{A}^M and cognitive actions to set internal memory codes and internal clocks.

Table 2
 Summary of the three exemplar models in this paper, in terms of the variation in the structure of the optimal program problems

Model	Bounded Machine M	Program Space \mathcal{P}^M	Utility \mathcal{U}	Evaluation Env. E_{eval}	Ecological Environments \mathcal{E}	Behaviors to Be Explained
Response ordering (Section 3.1)	Simplifications of ACT-R and EPIC	Subset of possible production rules	Payoff given in human experiment	Local dual-task experiment	Local dual-task experiment	Reaction times
Task-driven eye movements (Section 3.2)	Variations of oculomotor control architecture and Bayesian decision architecture	Thresholds for saccades, motor responses; lexicon priors	Three payoffs given in human experiments	Local task experiment	Local task experiment, prior lexical experience	Fixation durations; reaction times; errors
Wason selection task (Section 3.3)	None	None	Local task utility; information gain	Local task experiment	Objects w/common or rare properties	Errors

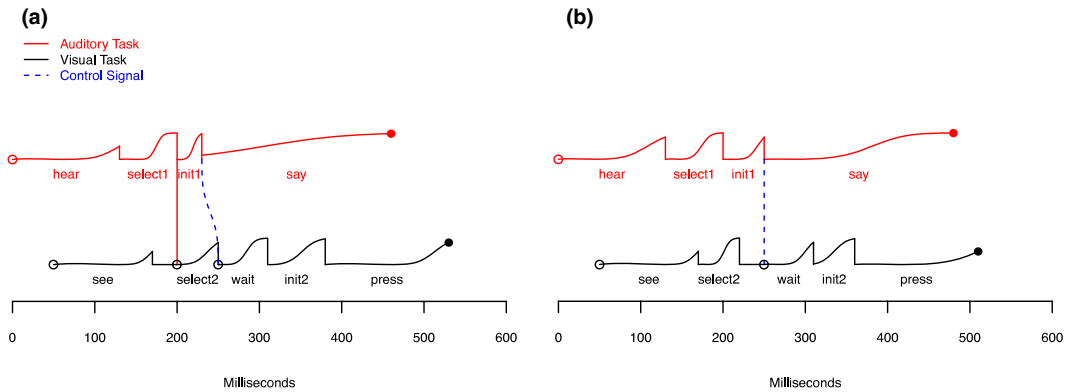


Fig. 3. The multiple dimensions of choice (duration of wait and control signals) for two trials using two bounded machines. In each panel, there is an auditory task and a visual task. Each requires multiple stages of processing, including perceptual, cognitive, and motor processing. Each stage duration is stochastic, which is represented in the figure with cumulative probability of completion in the y-axis. In both panels, the visual task finishes after the auditory task, indicating that, on this trial, a response reversal was avoided.

The bounds imposed by the machines include the following: (a) Response to a signal requires, minimally, a sequence of one perceptual, one cognitive, and one motor process, each of which has a duration sampled from a particular distribution. (b) In M_{ACTR} , but not M_{EPIC} , the process that selects the manual response cannot start until after the completion of the process that selects the verbal response. (c) The availability of control signals between processes imposes bounds on how tightly the manual and verbal responses can be synchronized. (d) A wait process can be used to slow the response to the visual-manual task, but its intended duration must be selected prior to the completion of the auditory task. Other bounds are also possible but not pursued here. Furthermore, a theory of the bounds for each individual participant can be defined by calibrating process duration distributions to individual participant task data (Howes et al., 2009).

3.1.2.2. Adaptation environment as ecological environment: Howes et al. (2009) assumed that, through thousands of trials, participants adapted to the evaluation environment. Crucially, the evaluation environment provided participants with an explicit feedback signal with which it was possible to ascertain the contribution of programs to utility. To emphasize the equivalence of the adaptation environment to the specific local evaluation environment in Schumacher et al. (1999), we say $E_{eval} = \mathcal{E} = \mathcal{E}_{Schu}$.

3.1.2.3. Utility function: The utility function $\mathcal{U} = \mathcal{U}_{Schu}$ was set to the monetary points regime given to the participants (Schumacher et al., 1999). The assumption was that the subjective utility function adopted by participants corresponded precisely to the objective utility set by the experimenter. More points were awarded for going fast, but points were deducted if the two responses were given in the wrong order. Money was awarded to each participant in proportion to his or her points total at the end of the experiment. Only

this utility function was considered for the analyses presented here, and it is used for both bounded machines.

3.1.3. Step 2: Select the optimal programs

The optimal programs can be derived for each machine according to Eq. 3, instantiated here for each machine:

$$\langle M_{\text{EPIC}}, \mathcal{P}^{*M_{\text{EPIC}}} \rangle = \{ \langle M_{\text{EPIC}}, P^* \rangle \} = \arg \max_{P \in \mathcal{P}^{M_{\text{EPIC}}}} \mathbb{E}_{h \sim ((M_{\text{EPIC}}, P), \mathcal{E}_{\text{Schu}})} \{ \mathcal{U}_{\text{Schu}}(h) \} \quad (4)$$

$$\langle M_{\text{ACTR}}, \mathcal{P}^{*M_{\text{ACTR}}} \rangle = \{ \langle M_{\text{ACTR}}, P^* \rangle \} = \arg \max_{P \in \mathcal{P}^{M_{\text{ACTR}}}} \mathbb{E}_{h \sim ((M_{\text{ACTR}}, P), \mathcal{E}_{\text{Schu}})} \{ \mathcal{U}_{\text{Schu}}(h) \} \quad (5)$$

Given the OPP, the set of optimal programs must order and select processes that result in an optimal inter-response interval (IRI). Optimal programs are those that select scheduling signals and a wait process duration that maximizes the utility by balancing the weighted cost of response reversals with the weighted cost of delayed responses—both of which are stochastic. A key emergent property of a given program is the *inter-response interval* (IRI). Fig. 4 plots utility against IRIs of a space of programs, exposing the fundamental trade-off of speed and accuracy for these machines and this task.

3.1.4. Step 3: Derive the predictions of the optimal programs

Predictions of the optimal programs were derived by conducting Monte-Carlo simulation of the optimal programs (i.e., sampling histories) on $\mathcal{E}_{\text{Schu}}$.

3.1.5. Step 4: Compare predictions with observed behavior

We examined the correspondence between the predictions derived from each theory of the bounds and observed human data reported by Schumacher et al. (1999). An analysis of correspondence between models and data from four experiments was reported by Howes et al. (2009). The results of one of these analyses are illustrated in Fig. 4.

The results showed that the average prediction of the response duration for the second task across four experiments had mean R^2 values of 89%, 84%, 77%, and 85% (Howes et al., 2009). The results also show that the serial bounded machine generated significantly lower differences between observed and predicted response times than the parallel bounded machine. However, optimal programs for both machines qualitatively reproduced the classic PRP dual-task interference effect.

3.1.6. Discussion

Two theories of the resource limits on how people process responses to simple stimuli, one with serial response selection and one parallel, were expressed as OPPs. The optimal programs derived as solutions to these OPPs predict inter-response-intervals and their correspondence to human data were measured. Differences in the level of correspondence achieved by the two theories suggested that the serial model offers better predictions than

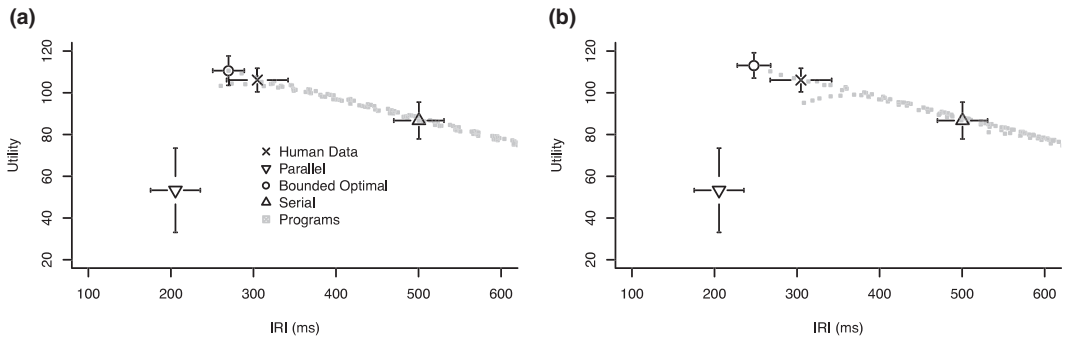


Fig. 4. The utility of programs plotted against the inter-response interval (IRI). (a) is for a bounded machine with a serial response selection bottleneck (ACT-R); (b) is for a bounded machine without such a bottleneck (EPIC). The utility of programs in each program space is plotted. The human data in both panels are from Schumacher et al. (1999) Experiment 2, with 95% confidence intervals. The shortest IRI is achieved by a model that executes the two tasks entirely in parallel with no interference, but neither architecture can achieve this and, regardless, utility is low because of reversals. The bounded optimal programs for each machine are the programs that achieve the highest utility and their mean utility is plotted with 95% C.I.s. The bounded optimal programs for the serial bounded machine are a better predictor of the human data than the bounded optimal program for the parallel bounded machine—although optimal programs for *both* machines reproduce the classic PRP dual-task interference effect.

the parallel model (Howes et al., 2009). But because both machines produce the qualitative dual-task slowing, the analysis suggests that this slowing itself cannot be taken as the empirical signature of a serial response bottleneck.

Furthermore, in the one experiment with sufficient individual variation, defining OPPs at the level of varying individual participant architectures (by calibrating process duration distributions on data other than that to be explained) yielded individual models that accounted for a significant proportion of between-participant variation. This provides further evidence that the human behavior represents an adaptation to the bounds of the processing architecture.

3.2. Adaptive eye-movement control (ecological bounded optimality explanation)

We turn to a model of visual attention, which provides direct evidence for adaptation to variation in utility, and also illustrates an explanation grounded in *ecological bounded optimality*. This is a re-description of a model of sequential eye movements in an elementary word reading task, and an associated set of experiments with human participants performing the task while their eye movements were monitored (Lewis et al., 2013).

Because the model integrates control of saccades and control of task-level responses, it provides a way to explore how low-level eye-movement decisions are influenced by higher level task goals. Optimization of task utility provides the analytic means to translate variation in goals into millisecond-level variation in eye-movements. More specifically, the theoretical claim explored here is that eye-movement strategies in reading are

precisely adapted to the joint constraints of (a) local task structure, including differential pressures on speed and accuracy; (b) the natural ecology of words, approximated by lexical frequencies in linguistic corpora; and (c) oculo-motor processing architecture, including temporal dynamics and perceptual/representational noise.

3.2.1. The List Lexical Decision Task

The *List Lexical Decision Task* (LLDT) is a simple extension of a paradigm introduced by Meyer and Schvaneveldt (1971). On each trial, participants are presented with a list of alphabetic character strings and must make a single decision: Does the list contain only words, or is there a nonword in the list? In the human and modeling experiments summarized below, there are six strings in a horizontal array; each string is four letters long. There is at most one nonword per list and no words are repeated in the same list. Fig. 5 shows the stimulus for a typical “all-words” trial.

3.2.2. Step 1: Defining the elements of the optimal control problem

3.2.2.1. *Information-processing machine and its bounds*: The machine bounds consist of a small set of independently motivated assumptions about the architecture of saccadic control, decision-making, and motor control: (a) saccadic control is a “rise-to-threshold” system (Brodersen et al., 2008) conditioned on task-specific decision variables that

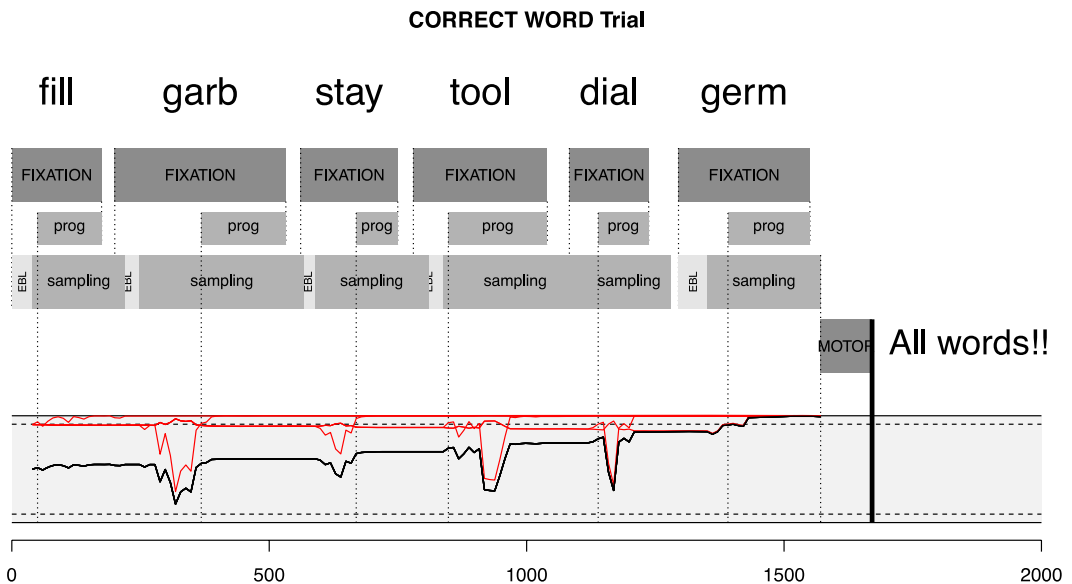


Fig. 5. Simulated model trace for a correct word trial (adapted from Lewis et al., 2013). The filled rectangles show the timing and duration of fixation durations, saccade programming (prog), eye-brain-lag (EBL), perceptual sampling, and motor response preparation and execution. At the bottom is the random walk of the belief probabilities, with the bottom representing 0 and the top 1. The black line is the trial-level belief (and so starts at 0.5), and the red lines are the string-level beliefs (and so start at 0.82, the prior probability that a given string is a word).

reflect the integration of noisy evidence over time; (b) there are separate decision variables for the control of saccades and task decisions indicated by motor presses; and (c) the perceptual, oculomotor, and manual motor system has specific dynamics that place limits on information-processing rates. These properties include saccade programming duration, eye-brain-lag, saccade execution duration, manual motor programming duration, and representational noise. Lewis et al. (2013) describe how they are motivated by prior empirical and theoretical work on eye movements in reading (Engbert, Nuthmann, Richter, & Kliegl, 2005; Reichle, Warren, & McConnell, 2009) and immediate response tasks (Meyer & Kieras, 1997). Only the representational noise parameter is not fixed in advance; we describe below how this parameter is set.

The machine properties can be understood by tracing the dynamics of processing in a single trial, as shown in Fig. 5. The first fixation starts on the leftmost string. During each fixation, noisy information about the fixated string is acquired at every timestep, with some delay (the eye-brain-lag, VanRullen & Thorpe, 2001). This information is used for a Bayesian update of beliefs about the status of the current string and the whole trial. The sampling continues until either the string-level or the trial-level belief reaches some threshold, at which point either a saccade is initiated (if the string-level *saccade threshold* is reached), or a manual response is initiated (if the trial-level *decision threshold* is reached). Information acquisition continues while the saccade or manual response is being programmed and until the saccade begins execution. Once saccade programming and execution is complete, the model fixates on the following string (if there are strings remaining) or initiates a response otherwise.

How the machine is bounded may be appreciated by considering the nature of the unbounded program space $\bar{\mathcal{G}}$ and the bounded subset \mathcal{G} . The action set \mathcal{A}^M consists of three actions: $\{\text{move-eyes-to-next-string}, \text{push-button-yes}, \text{push-button-no}\}$. The observations \mathcal{O}^M may be defined in one of two ways: as noiseless, thereby ascribing noise to the agent as a bound, or noisy, thereby ascribing noise to the environment. We adopt the latter formulation for the analytic purpose of understanding the effects of the machine *dynamics* as bounds. Thus, the unbounded program space consists of all possible mappings from history of observations and actions to probabilities of selecting the three actions: $(OA)^t O \times A \rightarrow [0,1]$. The program space for the bounded machine defined above consists of all possible *threshold pairs* $\langle [0,1], [0,1] \rangle$, which intensionally defines a *subset* \mathcal{G} of all possible mappings $\bar{\mathcal{G}}$. Crucially, this subset is determined by the temporal dynamics of the perceptual and motor architecture described above. It must be a strict subset ($\mathcal{G} \subsetneq \bar{\mathcal{G}}$) because the optimal program for the unbounded machine may condition its saccade and motor actions on *all* information obtained up to the time step immediately before the action, while the bounded machine imposes noisy delays that do not permit this optimal conditioning on history. Rather than attempt to analytically derive properties of this subset, we simulate performance of the optimal programs for the bounded and unbounded machines.

3.2.2.2. Adaptation environment as ecological environment: Lists of six unrelated four-letter strings do not have a natural ecology, but individual words do. The model assumes

that the adaptation environment has the structure of the local task (lists of six string strings; half of lists contain a nonword in a position chosen uniformly randomly), but that the distribution of the four letter words in the lists mirrors their frequency in the broader linguistic experience of the participants. We approximate this frequency via corpus frequency counts and this information is encoded in the priors in the model's lexicon; in this way, the optimal programs are adapted to this ecological word distribution.

3.2.2.3. Utility functions: As in the Response Ordering model, the utility function is derived directly from the payoff scheme used in the experiments. In the experiments and models reported in Lewis et al. (2013), there were three distinct payoffs that imposed different speed–accuracy trade-offs. We refer to the utility functions based on these three schemes as $\mathcal{U}_{\text{speed}}$ (weighted toward speed), \mathcal{U}_{acc} (weighted toward accuracy), and an intermediate function $\mathcal{U}_{\text{medium}}$. All three utility functions depend only on the manual button press, not the history of eye movements.

3.2.3. Step 2: Select the optimal programs

We are interested in deriving optimal programs for each of the three utilities $\mathcal{U}_{\text{speed}}$ (weighted toward speed), \mathcal{U}_{acc} (weighted toward accuracy), and the intermediate function $\mathcal{U}_{\text{medium}}$. We derive optimal *ecological bounded* programs, varying the bounds concerning dynamics constraints to explore their implications (this variation is described below).

The optimal programs were found through an exhaustive grid search over the two thresholds values, using Monte-Carlo simulation to determine expected utility. Optimal saccade and button-press thresholds were both higher under utility $\mathcal{U}_{\text{accuracy}}$, because raising thresholds trades speed for accuracy. Fig. 6, upper left panel, shows a portion of the payoff surface over the two dimensional threshold space, for $\mathcal{U}_{\text{accuracy}}$, with optimal points identified. The horizontal axis is saccade threshold, and each separate line corresponds to a different manual response threshold.

3.2.4. Step 3: Derive the predictions of the optimal programs

Once the optimal programs (thresholds) are determined, the behavioral predictions follow. Because the model both performs the task and controls the eyes, it yields a wide range of predictions, including overall RTs for correct and incorrect trials, single-fixation durations, frequency effects, accuracy effects, lexical-status effects, list position effects, and their modulation by task payoff. One parameter is not fixed in advance: the variance of the mean-zero Gaussian noise added to the perceptual samples. We explored a range of such noise settings and computed optimal policies across this range, and then compared the resulting behavioral predictions to the human data as described next.

3.2.5. Step 4: Compare predictions with observed behavior

We focus here on those comparisons that illustrate how evidence was obtained bearing on the three key theoretical assumptions (a)–(c) above.

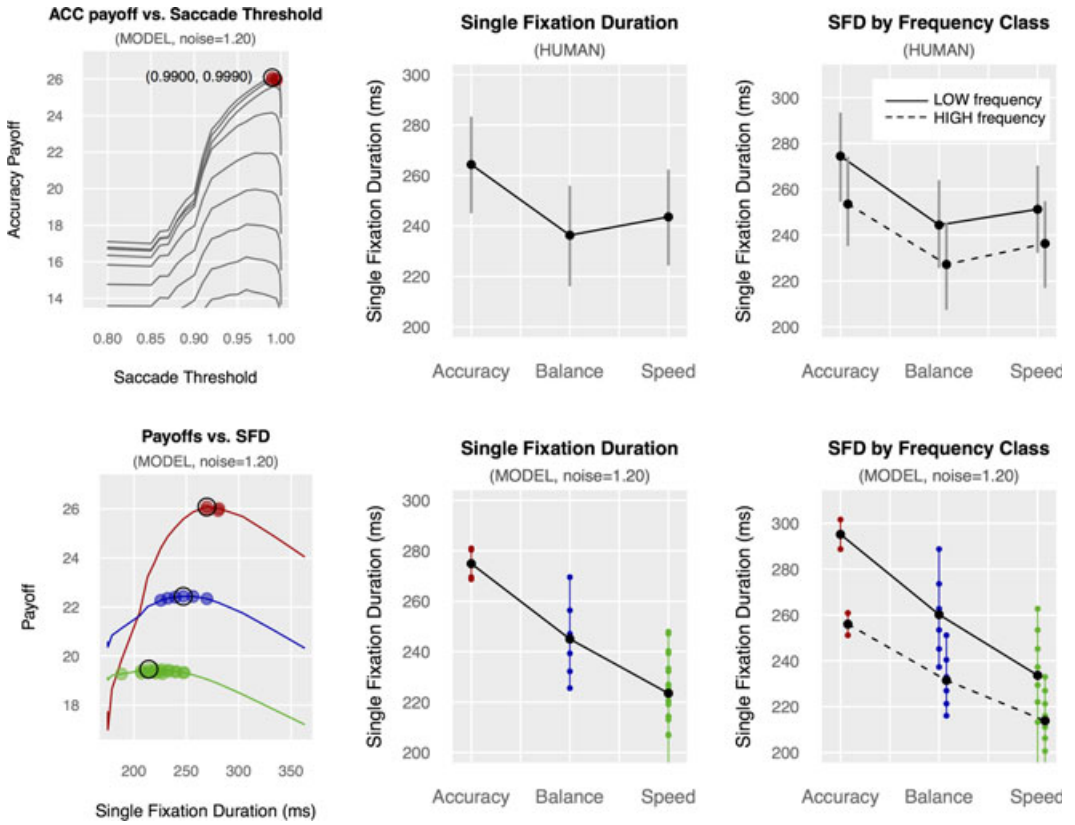


Fig. 6. A subset of modeling outcomes and human participant performance on the List Lexical Decision Task; see text under Steps 2, 3, and 4 for details.

3.2.5.1. *Adaptation to local task payoff*: Fig. 6 (bottom row, left panel) provides a view of the payoff surface for all three utilities, expressed over one of the behavioral outcomes: single fixation durations (SFDs). This view exposes a key prediction of the model: adaptation to changes at the highest level of task definition—task-level utility—should express itself as changes at the lowest levels of saccadic control, manifest in SFDs. Fig. 6 (bottom row, middle panel) replots the predicted SFDs for the optimal and near-optimal programs in the three payoffs; the top row, middle panel, shows the empirical SFDs recorded in the human eye-tracking experiments; the reliable difference between the speed and accuracy conditions confirms the key prediction.

3.2.5.2. *Adaptation to the natural ecology of words*: Evidence for ecological adaptation is obtained by assessing the correspondence of lexical frequency effects predicted by the model to those observed in the human participants. Fig. 6 (bottom row, right panel) shows the predicted SFDs for high- and low-frequency words: Because higher frequency words have a higher prior belief, they need fewer samples to reach the saccade decision

threshold. This effect is borne out in the human data (Fig. 6, top row, right panel). The success of this prediction and the nature of the explanation are inherited from Bayesian Reader (Norris, 2009), an unbounded sequential-sampling model. However, the precise quantitative predictions are sensitive to the bounds of the oculomotor machine, as we describe below. Indeed, with sufficient pressure on speed, the model predicts no frequency effect because a saccade program is initiated as soon as possible—before any samples are obtained.

3.2.5.3. Adaptation to oculo-motor processing architecture: Evidence for adaptation to the oculo-motor constraints is obtained by varying these processing constraints in the model and then re-deriving the optimal programs. Most crucial is the comparison of predictions of the optimal ecological program without the dynamics constraints. Recall the assertion that $\mathcal{G} \subsetneq \bar{\mathcal{G}}$. We now wish to determine whether this strict subset relation manifests in differences in the predictions of the optimal programs, and which of the predictions corresponds best to the human data. Fig. 7 plots deviation from human data (SFDs for the three payoffs; however, other measures and aggregations of measures produce similar results) for the machine without dynamics bounds and three kinds of machine incorporating successively more constraints on dynamics. This deviation is plotted against variation in the free noise parameter. The key result is that the predictions of the fully bounded machine corresponds best to the human data, and the machine without dynamics constraints is worst, even allowing for the noise parameter to be chosen for each machine to provide the best fit.

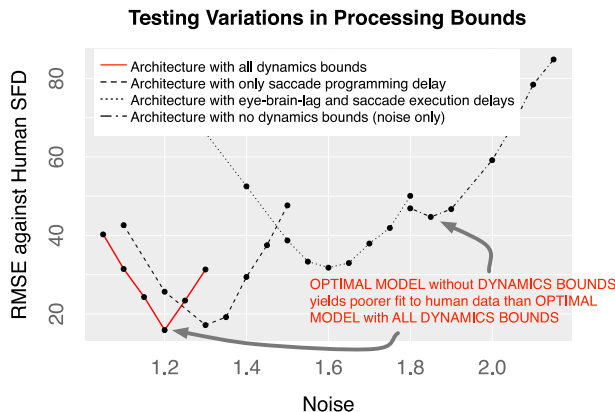


Fig. 7. Root mean squared error of model predictions (against mean single fixation duration [SFD] for the three payoff conditions) for four architectural variants. Optimal control policies are derived separately for each architecture. In red is the bounded machine architecture and includes saccade programming, eye-brain-lag, and saccade execution. The machine without dynamics bounds has none of these delays but does retain noise in perceptual samples. The other two models explore the effect of including and excluding the saccade programming delay but retaining delays imposed by saccade execution and eye-brain-lag.

3.2.6. Discussion

The List Lexical Decision model is an example of an ecologically bounded optimal explanation. It provided a rigorous way to derive and test the implications of the theoretical assumption that eye-movement control is jointly shaped by local task payoff, the bounds imposed by the oculomotor processor, and the ecology of words in linguistic experience. The necessity of the bounds was demonstrated by removing them and analyzing the consequences for correspondence with human data. The adaptation to task payoff was demonstrated empirically by showing that human participants changed their saccadic control at the level of fixation durations. The adaptation to the ecology of prior experience with words was demonstrated by the qualitative and quantitative correspondence of human data with the model's predicted frequency effects.

The model brings together three threads of research: (a) mathematical models of eye movement control in reading, which typically define an architecture, which is parameterized and then fitted to account for data (Engbert et al., 2005; Reichle, Rayner, & Pollatsek, 2003); (b) work on how higher level task goals shape eye movement strategies (Ballard & Hayhoe, 2009; Rothkopf, Ballard, Hayhoe, & Regan, 2007; Salverda, Brown, & Tanenhaus, 2011); and (c) Bayesian sequential sampling models of lexical processing and perception (Norris, 2006 2009; Wagenmakers, Ratcliff, Gomez, & McKoon, 2008). Finally, although this was a machine designed for the LLDT task, note that it is simply an instantiation of a more general architecture for the control of active perception and motor output in service of task goals, one that decomposes the problem into *optimal state estimation* and *optimal control*.

3.3. Reasoning with conditionals (ecological optimality explanation)

Here, we illustrate an explanation grounded in *ecological optimality* with a new description of Oaksford and Chater's (1994) seminal account of how people reason with conditional rules in the Wason selection task (Wason, 1966). The poor performance on this task exhibited by most people was thought to reveal systematic shortcomings in deductive reasoning capacities. Oaksford and Chater's contribution was to provide an alternative explanation based on assumptions about the ecology and utilities that would shape such reasoning capacities. Under these assumptions, Oaksford and Chater argued, the responses most people give to the task can be seen as rational. We now redescribe their analysis by developing it as an OPP, making explicit in a compact form the key assertions and nature of the prediction.

3.3.1. The Wason selection task

Participants were presented with four cards, each with a digit on one side and a letter on the other. They were then asked to verify that a rule holds, e.g., *if there is a letter A on one side, then there is a 7 on the other*. They were instructed that they should pick only the cards that they *must* turn over to verify the rule. Participants were asked to report which cards they would turn over, but not to actually turn cards. There was only one trial and participants did not see the results of turning cards; there was therefore no opportunity for sequential

decision making or learning from feedback. Given that the four cards show an A (p card), a K ($\neg p$ card), a 7 (q card), and a 2 ($\neg q$ card), the correct response, to verify the rule, is to select only the p and the $\neg q$ card. However, only a very small number of the participants made this response; many more selected the p and q cards.

3.3.2. Step 1: Defining the elements of the OPP

3.3.2.1. *Unbounded machine*: We define the available actions \mathcal{A}^M as seven of the possible responses (card selection sets): {none, p , q , p and q , p and $\neg q$, $\neg p$ and $\neg q$, $\neg p$ and q }. The initial observation is always four cards with visible properties p , q , $\neg p$ and $\neg q$. There are no further observations and so the unbounded program space $\bar{\mathcal{G}}$ is simply the set of all policies $G : \mathcal{A}^M \rightarrow [0, 1]$.

3.3.2.2. *Adaptation environment as ecological environment*: The assumption that Oaksford and Chater (1994) made about the ecological environment was that the properties that figure in causal relationships are rare in the environment (for example, when reasoning about the relationship of diseases and symptoms, both the diseases and symptoms are rare). We explore the consequences of this assumption by creating two adaptation environments: one in which the properties p and q are relatively rare (setting the unconditional probability that an object has either property as $Pr(p) = Pr(q) = 0.1$), and one in which the properties are relatively more common ($Pr(p) = Pr(q) = 0.6$). We denote the two adaptation environments as $\mathcal{E}_{\text{rare}}$ and $\mathcal{E}_{\text{common}}$. More precisely, $\mathcal{E}_{\text{rare}}$ and $\mathcal{E}_{\text{common}}$ consist of a distribution of four-tuples of objects o with visible properties $\langle p, q, \neg p, \neg q \rangle$, and “hidden” properties determined by $Pr(q)$, $Pr(p)$, $Pr(\neg q)$ and $Pr(\neg p)$, respectively.

3.3.2.3. *Utility functions*: We consider two utility functions. $\mathcal{U}_{\text{info}}$ awards utility in proportion to the information gained by turning each card (using the measure of Oaksford and Chater (1994); see Supporting Information). $\mathcal{U}_{\text{verify}}$ reflects the standard construal of the Wason task as rule verification: A person should select only those cards that allow the rule to be falsified for the current set of cards. $\mathcal{U}_{\text{verify}}$ awards 100 points for a correct declaration (i.e. p and $\neg q$) and 0 otherwise. $\mathcal{U}_{\text{info}}$ is meant to capture what is useful (and therefore what drives adaptation) in the ecological environment, and $\mathcal{U}_{\text{verify}}$ is meant to capture what is useful in the particular setting of the Wason task. By evaluating the four combinations of utility and environment, we can determine their sufficiency (and, in part, necessity) for accounting for behavior.

3.3.3. Step 2: Select the optimal programs

Because there is only one action and an invariant initial observation, we determine the optimal program by calculating the expected utility of each of the seven actions; the optimal program chooses with probability 1 the action a^* with the highest expected utility:

$$a^* = \arg \max_{a \in \mathcal{A}^M} \mathbb{E}_{o \sim \mathcal{E}} \{ \mathcal{U}(oa) \}. \quad (6)$$

$$G^*(a^*) = 1; G^*(a_i) = 0, \forall a_i \neq a^* \tag{7}$$

where $U(oa)$ is the utility obtained when taking action a upon observing the four-tuple o .

We do this separately for each combination of the two utility functions ($U_{\text{info}}, U_{\text{verify}}$) and the two ecological environments ($\mathcal{E}_{\text{rare}}, \mathcal{E}_{\text{common}}$). Table 3 provides the expected utility (calculated analytically; see Supporting Information) of the two actions corresponding to choosing p and q , and p and $\neg q$ for all four combinations (the other five actions, not shown, always yielded lower utility).

3.3.4. *Step 3: Derive the predictions of the optimal programs*

What is immediately evident in Table 3 is that choosing p and q is optimal when the utility function is information gain and the ecological environment has rare properties; this is Oaksford and Chater’s main result. This model predicts that all participants will select p and q because it is the optimal solution to the specified control problem. In every other combination of utility and environment, p and $\neg q$ is the optimal, and thus predicted, choice.

3.3.5. *Step 4: Compare predictions with observed behavior*

How well do the models predict the proportion of each subset of cards selected by participants? In a study reported by Johnson-Laird and Wason (1970), the responses were as follows: p and $\neg q$, accounted for 4% of responses; p and q , 46%; p only, 33%; the individual cards p , q , and $\neg q$, collectively 7%. The primary concern of Oaksford and Chater was to explain why the largest proportion of participants selected the p and q cards. The model assuming rarity of properties ($\mathcal{E}_{\text{rare}}$) and information gain utility (U_{info}) provides better correspondence than either of the three alternative models; it is the only model that yields p and q as the optimal.

But not all participants choose p and q ; there is considerable individual variation. While we do not present details here, this analysis can be extended to capture such variation. Some individual variation may be due to adoption of different utility functions.

Table 3

Exploring the implications of variation in utility and ecology in the Wason Selection Task. Table entries are the expected utility of two different card selections for two utility functions (information gain and verification) and two ecological environments (rare properties and common properties; see text for definitions). The other five actions (possible card choices) all have lower utilities in each of the four cases. Assuming information gain utility and rare properties, the highest utility action is to turn over cards p and q , which corresponds to 46% of the human data (Oaksford & Chater, 1994)

Action	Information Gain Utility (U_{info})		Verification Utility (U_{verify})	
	$\mathcal{E}_{\text{rare}}$	$\mathcal{E}_{\text{common}}$	$\mathcal{E}_{\text{rare}}$	$\mathcal{E}_{\text{common}}$
p and q	36.43	46.34	99	76
p and $\neg q$	30.29	86.31	100	100

While the modal function appears to be information gain, it is also possible that some participants, perhaps the 4% observed by Johnson-Laird and Wason (1970), do adopt the verification function. Second, some individuals will be more sensitive to the time costs of thinking and responding than others. Differences in the utility of time costs can explain why some participants select only the p card and no other card.

3.3.6. Discussion

Oaksford and Chater's claim (and ours) is not that people perform optimization to work out the best possible program. This can be seen in the form of Eq. 6: the "arg max" implicates work done by the analyst, not the human or organism being modeled. Rather, the claim is simply that the program that determines the response to the evaluation environment is optimal for the ecological environment. The computation may have been accomplished by some combination of biological evolution, cultural transmission and evolution, and individual learning.

Note that the model is not an ideal observer/actor in a local environment. Rather, the contention is that data selection (choice of cards) can be predicted by defining the optimal program problem faced in ecological environments and that this manifests behavior that is apparently suboptimal in particular evaluation environments.

4. Discussion and conclusion

In summary, we have presented *computational rationality* as a framework for formulating and working with optimality-based theories of human behavior that yields a type of explanation in which top-down rational approaches and bottom-up mechanism approaches are unified. It is an application of *bounded optimality* (Russell & Subramanian, 1995) to the challenges of developing psychological theory. The framework provides a way of pursuing the idea that behaviors are generated by cognitive mechanisms that are rationally adapted to the structure of both the mind and the environment. Theories are specified as *Optimal Program Problems* (OPPs). Such a specification demands three inputs: (a) an adaptation environment, (b) a bounded machine, and (c) a utility function. Given these inputs, a set of optimal programs are derived and provide the solution to the OPP, which determines a set of behavioral predictions. The solutions to OPPs provide a rational answer to the question: What should an agent *with some specific information-processing mechanisms* do in some particular environment? Correspondence with human data can be ascertained and comparisons made with alternative OPP-specified theories that vary the three components of the OPP.

We illustrated the framework with three examples, each of which illustrated theory comparison. The first example was an analysis of response ordering (Howes et al., 2009), which made explicit a comparison between alternative theories of mechanisms underlying response selection (serial vs. parallel). The second example, an analysis of eye movements in a simple reading task, made explicit a comparison between alternative theories of mechanisms of saccadic control, and provided evidence for adaptation of this control to variation

in utility (Lewis et al., 2013). The third example was a redescription of Oaksford and Charter's (1994) analysis of the Wason selection task that made explicit the comparison between alternative theories of the adaptation environment (rare vs. common) and, in addition, alternative theories of the utility function (information gain vs. verification).

The first and second examples illustrate how theories of the internal structure of the organism, that is, theories of the information-processing mechanism, can be informed by a bounded rational analysis. This was achieved in two ways: first, through the explicit derivation of the implications of different theories of the cognitive machine, thereby informing the choice of machine; and second, through the derivation of the optimal programs themselves, thereby informing the choice of the task-specific mechanisms underlying behavior. They also illustrate how bounded optimal and ecologically bounded optimal explanations differ from other rational explanation types. Unlike ecologically optimal explanations (e.g., Anderson, 1990), these explanations do not emphasize the role of the environment over the role of mechanisms. Neither do they involve determining rational functions as a guide to thinking about what mechanisms might be good approximations to these functions (Griffiths et al., 2012). In addition, the examples also illustrate the differences between bounded optimal and ecologically bounded optimal explanations and many mechanism approaches to explaining behavior. The fact that implied behavior is determined through optimization rather than through informal inference, or fitting, increases the rigor with which theories of the mechanisms can be tested. Utility maximization reduces the likelihood that behaviors which are the consequence of discretionary strategies are taken as evidence revealing of invariant mechanisms. This was made very clear in the analysis of the PRP task, where both the bounded optimal serial and parallel models produced qualitative patterns thought to indicate the presence of a serial bottleneck.

The unification of rational and mechanism approaches to psychological theory that we have illustrated turns on showing that the two can be joined with *optimal programs*. Here, we briefly highlight two key theoretical insights associated with optimal programs:

1. Behaviors are generated by organism-internal constructs: *programs*. Programs running on machines compute behavioral policies. Behaviors are not directly shaped by adaptive pressures; programs (mechanisms) are shaped. This last insight is also the basis of the emerging field of "evo-mecho" (McNamara & Houston, 2009), which focuses on understanding how evolution shapes mechanisms for generating behavior, rather than behavior itself.
2. Programs (and so only indirectly behaviors) are shaped by *both* information-processing mechanisms and environment structure. Deriving the consequences of these joint constraints is accomplished with the analytic tool of optimization. When policies are intensionally represented by programs running on specific information-processing machines, they are then assertions about "how" those policies are computed, hence "how" behavior is generated.

The framework is thus not theory agnostic; it carries these assertions. Like any framework (vs. a specific model), it cannot be falsified; it can only be more or less

useful. In fact, the framework can be seen as a sequential control generalization of signal detection theory (Tanner & Swets, 1954) (see Table 1), which has in fact been useful for 50 years. Where signal detection theory specified an optimal program as a scalar parameter, the proposed approach admits optimal programs that are any arbitrary algorithm executable, given the specified information-processing mechanism. This is a view that places the (utility-maximizing) control of behavior as the central problem faced by the organism.

We conclude with a brief comparison of computational rationality to existing theoretical “levels” frameworks in cognitive science and with Bayesian approaches to modeling. Chomsky (1965), Marr (1982), and Newell (1982) all advocate for an approach to cognitive science theory that distinguish useful levels of abstraction, but they differ in important ways. Marr (1982) advocates that the cognitive scientist start with a characterization at the *computational level* of *what function* is to be computed (the function from input to output, e.g. from 2 D percepts to 2 and 1/2 D sketch) and only once this is stated start to think about *how* these functions are implemented. Starting with a specification of functions is a standard top-down approach that is used in cognitive science as discussed earlier (e.g. Anderson, 1990; Griffiths et al., 2012; Oaksford & Chater, 2007). The framework advanced here is also a top-down approach, but rather than starting with a specification of *function*, it starts with a specification of *utility*. In this approach, the function—here the mapping from observations to actions—is not fixed at the start of the analysis; rather, it is *derived* (see the derivation of functions \mathcal{G}^* in Eqs. 3 and others in Section 2.2). The “what” is the derived mapping from observations to actions; the “how” is the (derived) program running on the machine.

Chomsky (1965) advocates that the scientist respect a distinction between *competence* and *performance*, where the former represents an abstract characterization of the organism’s internal capacities to generate behavior (but does not represent the behavior itself), while the latter is concerned with the expression of those capacities as behavior in particular settings. In the computational rationality framework, this distinction is respected: Derived programs are precisely abstract characterizations of competence, and performance is what those programs do in particular environments, given the constraints of the machine. For example, errors or slips might arise during execution of the program because of process noise in the machine or stochasticity in the environment. But these errors and slips need not be intrinsic properties of the programs themselves. Furthermore, it is possible to inquire about how such programs might express their capacities on machines with different bounds—just as Chomsky (1965) famously argued that our internal grammatical competence can generate unbounded self-recursive structures, even though the expression of this capacity behaviorally is prevented by our bounded memories. But unlike earlier competence theories in linguistics, in the computational rationality approach the competence-as-program is derived—not posited—as an adaptation to the joint constraints of machine and environment.⁶ Again, the starting point is utility.

Newell (1982) advances a systems-level analysis of organisms in which each level approximately implements the one above, with the highest level (for humans) being the *knowledge level*, where the principle of action selection is rationality (take the best action

to achieve current goals given the current knowledge). It differs from Marr's and Chomsky's approach in emphasizing rational choice at the highest level, and in this respect, it is more closely aligned with the utility-maximization view of computational rationality. But in Newell's approach, what is rational at the knowledge level does not take into account the mechanism bounds at lower levels (and so in this respect like classic rational analysis; Anderson, 1990; Oaksford & Chater, 2007). The framework advanced here can be understood as a way to allow rationality to do work at any arbitrary level of mechanism abstraction.

Finally, we note that nothing in the computational rationality framework is at odds with Bayesian inference as a way of computing the solutions to unbounded (Type I and Type II in Section 2.2 above) problems that receive a natural Bayesian formulation, or as a way of abstractly specifying subcomponents of an information-processing machine. Indeed, we adopted Bayesian inference as one component of the eye-movement model presented in Section 3.2. But computational rationality does differ from the top-down rational analysis approach associated with the Bayesian approach to cognition, for the reasons outlined above: The starting point is utility, not function; both mechanisms and functions are derived; and rationality does not abstract away from mechanism, but rather is defined in terms of it.

In summary, we have described and illustrated a framework, computational rationality, for formulating theories that explore the idea that behaviors are generated by cognitive mechanisms that are rationally adapted to the structure of both the mind and the environment. In this framework, utility maximization is used to derive optimal programs that execute on bounded information-processing architectures, which then determine the behavioral predictions of those architectures, and which are themselves derivations of cognitive mechanisms. The approach builds especially on the insights of signal detection theory (Tanner & Swets, 1954) and bounded optimality as advanced in AI (Russell & Subramanian, 1995), offering utility maximization as an analytic link between mechanism and behavior. It yields a class of explanation—Ecologically Bounded Optimality—that maintains the rigor and explanatory power of rational analysis in understanding behavior, but that reformulates the problem of rationality by taking into account information-processing bounds.

Acknowledgments

We thank Kevin Gluck, Wayne Gray, Glen Gunzelmann, Chris Myers, Michael Shvartsman, and Alonso Vera for useful discussion about this work. This material is based upon work supported by the National Science Foundation under Grant Numbers 1152819 and IIS 0905146; NASA Ames Research Center under Award NNX12AB08A; and Air Force Research Laboratory under Award FA8650-12-2-6353. Any opinions, findings, conclusions, or recommendations expressed here are those of the authors and do not necessarily reflect the views of the sponsors.

Notes

1. As we make clear in Table 1, signal detection theory had forms of many of the key elements of the analyses we advocate here.
2. Bounded rationality is usually associated with a rejection of optimality-based analyses, but Simon (1955) lays out a characterization of bounded rationality very close to the one we present here: “. . . we must be prepared to accept the possibility that what we call ‘the environment’ may lie, in part, within the skin of the biological organism. That is, some of the constraints that must be taken as given in an optimization problem may be physiological and psychological limitations of the organism (biologically defined) itself. . . . Limits on computational capacity may be important constraints entering into the definition of rational choice under particular circumstances.”
3. Baron and Kleinman (1969) derived joint optimal attention strategies and manual control strategies for a tracking task, given constraints on parafoveal vision and noisy motor control.
4. Although Anderson’s rational analysis as it has been usefully pursued (and summarized in Anderson’s own later work; Anderson, 2007) abstracts away from cognitive mechanism, the original informal method of rational analysis includes a step for including the computational constraints of the agent.
5. Whether any predicted behavior is described as learning, or whether the machine M has mechanisms that may be described as learning (or even optimization) mechanisms, are not distinctions explicitly recognized by the framework, although they may be of interest to the scientist. The definition of G is general enough to admit of arbitrary knowledge, Bayesian or otherwise, that determines the mapping. For example, the mapping could be the result of a specific Bayesian decision-making algorithm with some specific prior knowledge about the world. There are no constraints on the algorithm used and the prior knowledge available in determining G .
6. Chomsky’s own recent theoretical syntax work (e.g., Chomsky, 2005) explores the possibility that grammatical competence is shaped by properties of extra-linguistic cognitive architecture. In fact, he advances the conjecture that syntax may be a near perfect solution to these constraints. The approach advocated here may provide a rigorous analytic tool to explore this conjecture; see Bratman, Shvartsman, Lewis, and Singh (2010) for related work.

References

- Anderson, J. R. (1990). *The adaptive character of thought*. Hillsdale, NJ: Lawrence Erlbaum.
- Anderson, J. R. (2007). *How can the human mind occur in the physical universe?* New York: Oxford University Press
- Anderson, J. R., & Schooler, L. J. (1991). Reflections of the environment in memory. *Psychological Science*, 2, 396–408.
- Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S., Lebiere, C., & Qin, Y. (2004). An integrated theory of the mind. *Psychological Review*, 111(4), 1036–1060.

- Ballard, D. H., & Hayhoe, M. M. (2009). Modelling the role of task in the control of gaze. *Visual Cognition*, 17(6–7), 1185–1204.
- Baron, S., & Kleinman, D. L. (1969). The human as an optimal controller and information processor. *IEEE Transactions on Man-Machine Systems*, 10(1), 9–17.
- Bellman, R. E. (1957). *Dynamic programming*. Princeton, NJ: Princeton University Press.
- Bratman, J., Shvartsman, M., Lewis, R. L., & Singh, S. (2010). A new approach to exploring language emergence as boundedly optimal control in the face of environmental and cognitive constraints. In D. Salvucci & G. Gunzelmann (Eds.), *Proceedings of the 10th International Conference on Cognitive Modeling* (pp. 7–12). Philadelphia, PA: Drexel University.
- Brodersen, K. H., Penny, W. D., Harrison, L. M., Daunizeau, J., Ruff, C. C., Duzel, E., Friston, K. J. and Stephan, K. E. (2008). Integrated bayesian models of learning and decision making for saccadic eye movements. *Neural Networks*, 21(9), 1247–1260.
- Chomsky, N. (1965). *Aspects of the theory of syntax*. Cambridge, MA: MIT Press.
- Chomsky, N. (2005). Three factors in language design. *Linguistic Inquiry*, 36(1), 1–22.
- Cosmides, L., Barrett, H. C., & Tooby, J. (2010). Adaptive specializations, social exchange, and the evolution of human intelligence. *Proceedings of the National Academy of Sciences*, 107(Suppl. 2), 9007–9014.
- Engbert, R., Nuthmann, A., Richter, E., & Kliegl, R. (2005). Swift: A dynamical model of saccade generation during reading. *Psychological Review*, 112(4), 777–813.
- Faisal, A. A., Selen, L. P. J., & Wolpert, D. M. (2008). Noise in the nervous system. *Nature Reviews Neuroscience*, 9(4), 292–303.
- Geisler, W. S. (2011). Contributions of ideal observer theory to vision research. *Vision Research*, 51(7), 771–781.
- Gigerenzer, G., & Selten, R. (2001). Rethinking rationality. In G. Gigerenzer & R. Selten (Eds.), *Bounded rationality: The adaptive toolbox* (pp. 1–12). Cambridge, MA: MIT Press.
- Griffiths, T. L., Chater, N., Norris, D., & Pouget, A. (2012). How the bayesians got their beliefs (and what those beliefs actually are): Comment on bowers and davis (2012). *Psychological Bulletin*, 138(3), 415–422.
- Howes, A., Lewis, R. L., & Vera, A. (2009). Rational adaptation under task and processing constraints: implications for testing theories of cognition and action. *Psychological Review*, 116(4), 717–751.
- Johnson-Laird, P., & Wason, P. (1970). A theoretical analysis of insight into a reasoning task. *Cognitive Psychology*, 1(2), 134–148.
- Laird, J. E. (2012). *The soar cognitive architecture*. Cambridge, MA: MIT Press.
- Lewis, R. L., Shvartsman, M., & Singh, S. (2013). The adaptive nature of eye-movements in linguistic tasks: How payoff and architecture shape speed-accuracy tradeoffs. *Topics in Cognitive Science*, 5(3), 1–30.
- Marr, D. C. (1982). *Vision*. New York: Freeman.
- McNamara, J. M., & Houston, A. I. (2009). Integrating function and mechanism. *Trends in Ecology & Evolution*, 24(12), 670–675.
- Meyer, D., & Kieras, D. (1997). A computational theory of executive cognitive processes and multiple-task performance: Part 1. Basic mechanisms. *Psychological Review*, 104, 3–65.
- Meyer, D., & Schvaneveldt, R. (1971). Facilitation in recognizing pairs of words: Evidence of a dependence between retrieval operations. *Journal of Experimental Psychology*, 90, 22–34.
- Newell, A. (1982). The knowledge level. *Artificial Intelligence*, 18, 87–127. (Also in *AI Magazine*, Vol 2, Summer 1981, pp. 1.20).
- Newell, A. (1990). *Unified theories of cognition*. Cambridge, MA: Harvard University Press.
- Newell, A., & Simon, H. A. (1972). *Human problem solving*. Englewood Cliffs, NJ: Prentice-Hall.
- Norris, D. (2006). The Bayesian reader: Explaining word recognition as an optimal Bayesian decision process. *Psychological Review*, 113(2), 327–357.
- Norris, D. (2009). Putting it all together: A unified account of word recognition and reaction-time distributions. *Psychological Review*, 116(1), 207–219.
- Oaksford, M., & Chater, N. (1994). A rational analysis of the selection task as optimal data selection. *Psychological Review*, 101(4), 608.

- Oaksford, M., & Chater, N. (2007). *Bayesian rationality*. Oxford, UK: Oxford University Press.
- Pashler, H. E. (1998). *The psychology of attention*. Cambridge, MA: MIT Press.
- Reichle, E. D., Rayner, K., & Pollatsek, A. (2003). The E-Z Reader model of eye-movement control in reading: Comparisons to other models. *Behavioral and Brain Sciences*, 26(4), 445-+.
- Reichle, E. D., Warren, T., & McConnell, K. (2009). Using E-Z reader to model the effects of higher level language processing on eye movements during reading. *Psychonomic Bulletin Review*, 16, 1–21.
- Rothkopf, C. A., Ballard, D. H., Hayhoe, M. M., & Regan, O. (2007). Task and context determine where you look. *Journal of Vision*, 7, 1–20.
- Russell, S., & Subramanian, D. (1995). Provably bounded-optimal agents. *Journal of Artificial Intelligence Research*, 2, 575–609.
- Salverda, A. P., Brown, M., & Tanenhaus, M. K. (2011). A goal-based perspective on eye movements in visual-world studies. *Acta Psychologica*, 137(2), 172–180.
- Schumacher, E., Lauber, E., Glass, J., Zurbruggen, E., Gmeindl, L., Kieras, D., et al. (1999). Concurrent response-selection processes in dual-task performance: Evidence for adaptive executive control of task scheduling. *Journal of Experimental Psychology*, 25(3), 791–814.
- Simon, H. A. (1955). A behavioral model of rational choice. *The Quarterly Journal of Economics*, 69(1), 99–118.
- Simon, H. A. (1956). Rational choice and the structure of the environment. *Psychological Review*, 63(2), 129–138.
- Simon, H. A. (1990). Invariants of human behavior. *Annual Review of Psychology*, 41, 1–19.
- Singh, S., Lewis, R. L., Barto, A. G., & Sorg, J. (2010). Intrinsically motivated reinforcement learning: An evolutionary perspective. *IEEE Transactions on Autonomous Mental Development*, 2(2), 70–82.
- Stengel, R. F. (1994). *Optimal control and estimation*. New York: Dover Publications.
- Stephens, D. W., & Krebs, J. R. (1986). *Foraging theory*. Princeton, NJ: Princeton University Press.
- Sutton, R., & Barto, A. (1998). *Reinforcement learning: An introduction*. Cambridge, MA: MIT Press.
- Tanner, W. P., & Swets, J. A. (1954). A decision-making theory of visual detection. *Psychological Review*, 61(6), 401.
- Tooby, J., & Cosmides, L. (2005). Conceptual foundations of evolutionary psychology. In Buss, D. M. (Ed.), *The handbook of evolutionary psychology* (pp. 5–67). Hoboken, NJ: Wiley.
- Trommershaeuser, J., Maloney, L. T., & Landy, M. S. (2008). Decision making, movement planning and statistical decision theory. *Trends in Cognitive Sciences*, 12(8), 291–297.
- VanRullen, R., & Thorpe, S. J. (2001). The time course of visual processing: From early perception to decision-making. *Journal of Cognitive Neuroscience*, 13, 454–461.
- Von Neumann, J., & Morgenstern, O. (1947). *Theory of games and economic behavior*. Princeton, NJ: Princeton University Press.
- Wagenmakers, E., Ratcliff, R., Gomez, P., & McKoon, G. (2008). A diffusion model account of criterion shifts in the lexical decision task. *Journal of Memory and Language*, 58(1), 140–159.
- Wason, P. (1966). Reasoning. *New Horizons in Psychology*, 1, 135–151.
- Wolpert, D. M. (2007). Probabilistic models in human sensorimotor control. *Human Movement Science*, 26(4), 511–524.

Supporting Information

Additional Supporting Information may be found in the online version of this article:

Data S1. Analysis of the Wason Task for Table 3.