

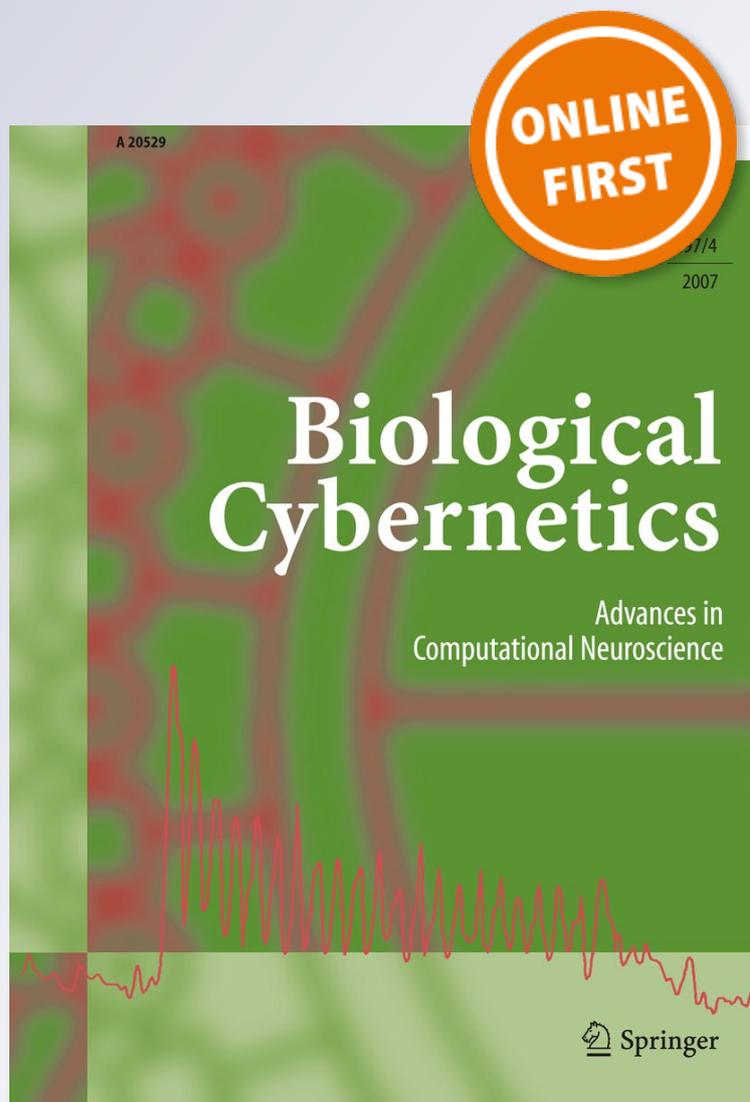
Optimal isn't good enough

Gerald E. Loeb

Biological Cybernetics
Advances in Computational
Neuroscience

ISSN 0340-1200

Biol Cybern
DOI 10.1007/s00422-012-0514-6



 Springer

Your article is protected by copyright and all rights are held exclusively by Springer-Verlag. This e-offprint is for personal use only and shall not be self-archived in electronic repositories. If you wish to self-archive your work, please use the accepted author's version for posting to your own website or your institution's repository. You may further deposit the accepted author's version on a funder's repository at a funder's request, provided it is not made publicly available until 12 months after publication.

Optimal isn't good enough

Gerald E. Loeb

Received: 6 January 2012 / Accepted: 31 July 2012
© Springer-Verlag 2012

Abstract The notion that biological systems come to embody optimal solutions seems consistent with the competitive drive of evolution. It has been used to interpret many examples of sensorimotor behavior. It is attractive from the viewpoint of control engineers because it solves the redundancy problem by identifying the one optimal motor strategy out of many similarly acceptable possibilities. This perspective examines whether there is sufficient basis to apply the formal engineering tools of optimal control to a reductionist understanding of biological systems. For an experimental biologist, this translates into whether the theory of optimal control generates nontrivial and testable hypotheses that accurately predict novel phenomena, ideally at deeper levels of structure than the observable behavior. The methodology of optimal control is applicable when there is (i) a single, known cost function to be optimized, (ii) an invertible model of the plant, and (iii) simple noise interfering with optimal performance. None of these is likely to be true for biological organisms. Furthermore, their motivation is usually good-enough rather than globally optimal behavior. Even then, the performance of a biological organism is often much farther from optimal than the physical limits of its hardware because the brain is continuously testing the acceptable limits of performance as well as just performing the task. This perspective considers an alternative strategy called “good-enough” control, in which the organism uses trial-and-error learning to acquire a repertoire of sensorimotor behaviors that are known to be useful, but not necessarily optimal. This leads to a diversity of solutions that tends to confer robustness on the individual organism and its evolution. It is also more consistent with the capabilities of higher sensorimotor structures,

such as cerebral cortex, which seems to be designed to classify and recall complex sets of information, thereby allowing the organism to learn from experience, rather than to compute new strategies online. Optimal control has been a useful metaphor for understanding some superficial aspects of motor psychophysics. Reductionists who want to understand the underlying neural mechanisms need to move on.

Keywords Sensorimotor · Optimal control · Spinal cord · Internal model · Motor learning · Feedback · Cost · Motor noise

1 Perspective

The notion that biological systems come to embody optimal solutions seems consistent with the competitive drive of evolution. It has been used to interpret many examples of sensorimotor behavior. This perspective examines whether there is sufficient basis to apply the formal engineering tools of optimal control to a reductionist understanding of biological systems. For an experimental biologist, this translates into whether the theory of optimal control generates nontrivial and testable hypotheses that accurately predict novel phenomena, ideally at deeper levels of structure than the observable behavior.

David Marr identified three levels of analysis for an information processing system (Marr 1982). At the top is a theory of computation—the problem to be solved. We must ask what exactly is being optimized and how important it is for the organism to find optimal solutions to sensorimotor control problems. Next is the algorithmic representation of the solution—the computational transformations and procedures that the organism employs to solve that problem. The biological equivalent of the engineering tools for finding optimal

G. E. Loeb (✉)
Department of Biomedical Engineering,
University of Southern California, Los Angeles, CA, USA
e-mail: gloeb@usc.edu

solutions generally require internal models of the plant, existence of which is often assumed in or inferred from the interpretation of psychophysical data. At the lowest level is the physical system that performs the computations required by the algorithms—the hardware and software of computers or brains. In contrast with a theoretical or engineered system, the biological nervous system has constraints on its design, which arise from the phylogenetic evolution of its hardware (neurons) and the ontogenic development of its software (memories). Unless optimal control can succeed on all three levels, it is not a viable theory.

2 History

The interest of neuroscientists in optimal control can be traced to the problem of redundancy posed by the pioneering (and long overlooked) biomechanical studies of Nikolai Bernstein (1967) (English translation of 1934 publication in Russian). The musculoskeletal system exhibits both kinematic and kinetic redundancy when asked to perform a simple task such as moving the hand through a specified trajectory in space. There are more joints in the limb than necessary to follow the trajectory, and so the segmental posture is unspecified (kinematic redundancy). There are more muscles than required to generate the net joint torques required to achieve any given state of limb motion (kinetic redundancy). So how does the brain decide what postures and which muscles to use? Bernstein proposed that the nervous system does not have independent access to all the degrees of freedom in the plant, and instead is limited to certain synergies. This would seem to foreclose valuable solutions to more complex problems such as producing a particular mechanical impedance to resist perturbations while moving along a trajectory (Osu and Gomi 1999). Instead of invoking synergies, engineers solve such problems by adding the constraint that the solution must minimize some cost function. Biologists have developed optimal solutions for the redundancy problem based on minimizing total angular excursion of joints (Pellionisz and Peterson 1998), minimizing total muscle force, joint stress, or other kinetic costs (Pedotti et al. 1978; Crowninshield and Brand 1981; Buchanan and Shreeve 1996; Raikova et al. 2005), or maximizing performance (Pandy et al. 1990, 1992; Anderson and Pandy 1999). If the mechanics of the musculoskeletal plant can be described as an invertible mathematical model, then it is possible to compute such optimal solutions directly using the Moore–Penrose pseudo-inverse or algebraic Riccati equations.

3 What are the goals of sensorimotor control?

Any attempt to improve performance (whether to make it optimal or simply better) begs a quantitative definition of

quality of performance. The optimal solutions to biological problems cited above were all based on a simple cost function that was intuited, based on what the investigator guessed was important to the organism *and* limited by the state variables available in the model of the system to be optimized. For example, energy consumption seems likely to be more important than muscle force, but only recently has there been an accurate energetics model for biological muscles under dynamic conditions (Tsianos et al. 2012). Furthermore, biological systems are likely to have multiple competing goals (e.g., speed, accuracy, and effort), relative importance of which is itself learned and task dependent. Given a sufficiently rich body of performance data, it is possible to solve an inverse problem to identify the cost function that is being optimized. This can be done by comparing predicted to observed behavior for various weightings of terms in the cost function, either informally (He et al. 1991) or formally (Terekhov and Zatsiorsky 2011). The problem rapidly becomes intractable, however, if more than two types of cost must be weighted differentially (e.g., position, impedance, and effort). It may be possible to use higher-order goals of the organism to lump some of these cost terms into a unifying principle such as “free-energy minimization” (Friston 2010), but implementing this cost for real behavior generally requires many implicit assumptions.

3.1 Optimal control: globally minimized cost

The word “optimal” actually originated in evolutionary biology to mean “most favorable” (Harper 2010). It now has a very specific but closely related meaning that is essential to engineering concepts and mathematical tools. It is a superlative, referring exclusively to *the one state* that is better than all others as formally defined by a quantitative cost function. As noted above, these optimal control methods have been applied to biological systems because they eliminate the redundancy problem, providing a unique solution to any problem. As noted below, computing optimal solutions is extremely difficult, currently impossible for most systems containing nonlinear elements and for dynamically changing cost functions.

3.2 Good-enough control: locally defined utility

Any adaptive control system will tend gradually toward a locally stable state if one exists, but there is no guarantee that this state is globally optimal. The actual state that will be reached depends on the starting state and nature of the system. Consider a hyperspace with many independent variables representing the control inputs to the system. The cost of any point in this hyperspace can be computed as a function of the state variables of the system. For a sufficiently complex system, the hypersurface of this cost function will

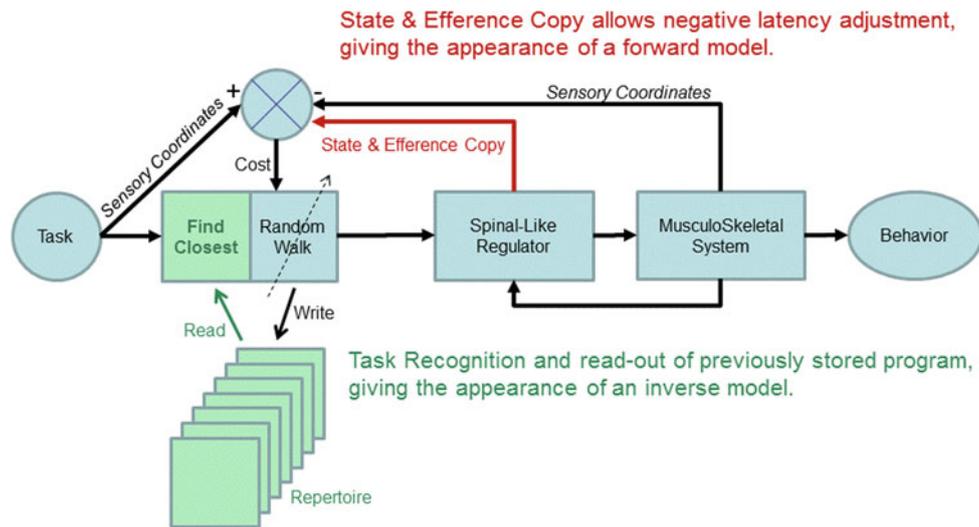


Fig. 1 A hierarchical system that behaves as if it used internal models to compute optimal control but does not. The controller at the left learns by trial-and-error and saves useful programs that it can recall when faced with a similar task, giving the appearance that it possesses and uses an internal inverse model of the musculoskeletal system to compute motor programs online. The programs from the controller operate on a lower-level regulator that mixes them with concurrent somatosensory feedback from the musculoskeletal system. These sensory signals

plus efference copies of inputs and internal states of the regulator are compared with the preferred feedback associated with the task, providing a teacher for the controller. It also results in short-latency and even negative-latency error signals that trigger corrective movements before a motor error appears in the kinematics, giving the appearance that the controller possesses an internal forward model of the musculoskeletal system

have many peaks and valleys. If the state equations are not invertible, then there is no way to determine whether any local valley represents the global minimum cost in the hyperspace, short of exhaustively computing all possible states. Various formal strategies for systematic exploration of such hyperspaces have been developed to maximize the probability of discovering the global minimum (or one sufficiently close to it) in the shortest possible time (Gill et al. 1982; Lagarias et al. 1998; Boyd and Vandenberghe 2004; Venkataraman 2009; Yuan et al. 2010). Such strategies entail their own costs, such as the time required to find an acceptable solution and the number of intermediate results that must be stored and reviewed to guide the next step of the exploration. The exploration strategy that is most effective thus depends on both the shape of the hyperspace as well as the goals, strengths, and weaknesses of the controller exploring the space.

In the remainder of this article, we will consider good-enough control as an alternative to optimal control. The general notion is that most biological organisms (unlike an engineer or an Olympic athlete) do not care about finding globally optimal solutions to problems; instead they are motivated by how quickly they can find any solution that is good enough to obtain a reward or escape a punishment. This utilitarian goal has important implications for the algorithms that might be employed to explore the hyperspace (see below). If the goal really is optimal control, then one would prefer to keep the number of dimensions (free variables to be explored) as low as possible. This sort of thinking led Bernstein to propose that the lower levels of the nervous system introduced

constraints (i.e., fixed synergies or primitives) to simplify the control problem faced by the higher level controllers. If one is looking only for good-enough rather than optimal solutions, however, then the addition of more dimensions may be useful if they increase the density of such good-enough local minima in the hyperspace. If there are lots of good-enough local minima in all parts of the hyperspace, then any learning algorithm will be more likely to find one of them quickly regardless of where it starts in the hyperspace.

4 What is the algorithm for sensorimotor learning?

4.1 Good-enough control: trial-and-error learning of a repertoire

The motivation for good-enough control comes from models of the musculoskeletal system plus spinal cord (see Fig. 1). The spinal cord consists of several types of well-described interneurons that mix various sources of somatosensory signals to provide feedback from the plant to the motoneurons (Pierrot-Deseilligny and Burke 2005). The effective gains of these multi-input–multi-output elements are set by descending control from the brain, suggesting the term “spinal-like regulator” for the computational model. The spinal cord is the phylogenetically oldest part of the vertebrate sensorimotor nervous system. Its connectivity is largely determined genetically during embryological development rather than resulting from behavioral experience and learning. Thus, it is appropriate to consider it as a part of the plant that the

higher level systems in the brain must learn to control. The spinal-like regulator presents a much larger number of non-linear elements to be controlled than the individual muscles, immediately compounding the problems of redundancy and invertibility. For a two-dimensional (2D) model of the wrist with four muscles, there were about 200 elements to be controlled (Raphael et al. 2010); for a 2D model of the coplanar elbow and shoulder with six muscles, there were about 400 such elements (Tsianos et al. 2011). Surprisingly, extremely simple trial-and-error learning tended to find good-enough solutions even for complex sensorimotor problems such as rapid movements in a viscous curl field (Raphael et al. 2010). The hyperspace appeared to have many good-enough solutions scattered throughout and discovered easily. As discussed later, an organism could accumulate a repertoire of such learned movements, thereby reducing most new tasks to the process of looking up the closest, previously learned task(s) and performing some minor interpolations and refinements (Loeb 1983).

4.2 Optimal control: computation from an internal model

All of the computational methods for guaranteeing a globally optimal solution (short of exhaustive exploration of all possible states) require an inverse model of the plant to be controlled [but see (Li et al. 2012) for an approximately optimal nonlinear controller for locally linear conditions]. It is difficult enough to come up with an accurate forward model of a biological neuromusculoskeletal system and impossible to invert such models without making gross simplifications, usually involving linearity. Valero-Cuevas and colleagues reviewed the application of optimal control methods to musculoskeletal systems and pointed out how uncertainties in these models would likely lead to instability (Valero-Cuevas et al. 2009).

If one assumes that the only escape from the redundancy problem is optimal control, then one is forced to postulate a *deus ex machina*—the nervous system somehow learns such an inverse model as a result of many years of practice (Kawato 1999), the equivalent of system identification. The presumption that there is a sufficiently accurate internal model somewhere in the nervous system and that it is used to compute optimal solutions online has led to some useful summaries of observable motor psychophysics (Flanders 2011). The question is whether the predictions of this algorithm are different from those for a stored repertoire of good-enough solutions. If one uses the term “internal inverse model” to describe the performance rather than the structure of the control system, then it becomes a tautological description of any system that behaves successfully without requiring extensive iteration (see Fig. 1). If the control scheme needs to account for more complex emergent phenomena such as the shape of the learning curve and the persistence of motor habits

(see below), then good-enough control looks more plausible. If experimenters actually go looking for an internal model, then where should they look and what would they find?

5 What is the machinery for sensorimotor learning?

5.1 Optimal control: cerebellum?

Everyone acknowledges that the brain does not solve algebraic Riccati equations, but it is not clear how the nervous system might implement any other inverse internal model, much less a rule to learn and update such a model behaviorally. Kawato and Gomi proposed a mechanism for the cerebellum to learn such a model, but it required other parts of the brain to compute an error signal in motor coordinates rather than the sensory coordinates that seem more plausible for cerebellar inputs (Kawato and Gomi 1992). Their example was drawn from the well-known role of the cerebellum in fine-tuning the vestibulo-ocular reflex. That is a very special plant, however, for which the sensorimotor dynamics are simple and largely orthogonal rather than redundant and for which there is little or no proprioceptive regulation between command signals and muscle activation.

5.2 Good-enough control: neuronal hardware + experiential software

Evolution is not the same as engineering design (Partridge 1982). Optimality got its start in Darwinian evolution, but when applied to the nervous system, it conflates two subsystems for which optimal means very different things: the “hardware” of neuromuscular systems and the “software” of learned sensorimotor behaviors. The hardware includes the complete phylogenetic set of articulated segments, neuromuscular actuators, sensors, and central neurons available to the species, thereby delimiting its ultimate capabilities, which evolution presumably seeks to optimize. The software includes whatever ontogenic modifications and memories a given individual has managed to accumulate, thereby delimiting its current performance. The process of genetic mutation has no mechanism to insure that the cumulative effects of these mutations results in an organism that is optimal according to physical capabilities; rather it must merely be better at occupying some ecological niche than any other organism that has yet occurred. Nevertheless, given sufficient time and consistent selection pressure, many capabilities do appear to asymptote toward optima within the constraints of the basic schema. By contrast, the individual confronted with a new task has no motivation to find a solution that is optimal according to physical performance criteria; rather, the motivation is to find quickly a solution that is good enough to get rewarded without expending more time or effort than the

reward is perceived to be worth. Developing a new sensorimotor program from scratch is slow and difficult; hacking a previously saved program is usually the best strategy. Actual performance is likely to be noisy and idiosyncratic, forcing psychophysicists to average a lot of data (see Sect. 6.1) before proclaiming universal principles of optimal control. Those principles tend to provide more insight into the banality and ambiguity of most experimental tasks than into adaptive control by the brain. Rather than focusing on the observable solutions to tasks, neuroscientists need to focus on the process of finding those solutions.

Optimal control is a concept from engineering that has been used with some success in very well-defined industrial problems, but it is not actually common practice even in engineering. The following explores a broader and perhaps more realistic analogy between engineered and biological systems that better reflects the hardware–software distinction introduced above.

- *A modern engineer* faced with designing, building, and programming a complex system starts first with a set of tasks that it must perform, then figures out the simplest strategies for performing those tasks, then selects the most cost-effective general purpose components that can embody those strategies, and finally assembles the hardware and writes the software as a foregone conclusion. The wise engineer knows that “mission creep” is inevitable, and so he/she designs in as much extra capacity as possible based on currently available “off-the-shelf” technologies, eschewing expensive and risky custom development unless absolutely unavoidable. These stock components themselves represent decades of intensive and expensive development that simultaneously empowers the engineer to achieve a high level of performance at low cost while confining that performance to some arbitrary limits, an acceptable bargain if that performance is “good enough” to meet the original requirements. Inevitably, the product eventually becomes obsolete as newer off-the-shelf technologies emerge (i.e., hardware evolution) and are incorporated into competing products, enabling their programmers to write software that serves new demands that thereby become requirements for competing systems.
- *A biological organism* evolves as an assembly of “off-the-shelf” cellular and molecular components that themselves represent hundreds of millions of years of evolution. These components and the specific organism are optimal only in the sense that their survival indicates that no better designs have yet occurred in the particular ecological niches that they occupy. For most of biological history, performance depended almost entirely on such hardware. The “learning” of primitive organisms is akin to fixed rules for adaptive control that are themselves embodied

in the genetic firmware of modulation and messengers, thereby limiting the behavioral repertoire of the organism. At some relatively recent point in evolution, nervous systems became capable of something akin to programming a general purpose computer. It should be noted that the precipitating factor is probably not to be found in the neural hardware itself. The underlying physics of a digital computer are no different from those of a transistor radio, but the emergent properties of those products force us to treat them differently. In particular, the actual performance of the computer depends on both the slow and expensive evolution of its hardware and firmware (now requiring billion dollar silicon foundries) and the rapid and cheap experiments with writing code (suitable for “dotcom” start-up companies). As intelligent biological organisms, humans inherit the best current neural components and system architectures that Mother Nature has yet devised and then spend our countless individual lives hacking together software.

5.3 The role of cerebral cortex

The biological software hacker seems to reside in the cerebral cortex, all phylogenetically proliferating regions of which appear to share a similar computational algorithm (Diamond 1979). This algorithm seems to be particularly suited to classifying and recalling complex sets of information, thereby allowing the organism to learn from experience. Like the modern software hacker, the cortex is particularly good at finding and taking advantage of bits of code that already exist rather than constructing new code from scratch.

Decorticate mammals actually exhibit a fairly rich repertoire of complex sensorimotor behaviors, including conditioned sensorimotor responses (Culler and Mettler 1934; Bjursten et al. 1976). They rely on the reflexes and pattern generators hardwired into subcortical structures like the tectum (Perfiliev et al. 2010), brainstem, and spinal cord (McCrea and Rybak 2008), which subserve the whole behavioral repertoire of intact lower vertebrates. The strategy for learning the repertoire of good-enough motor skills may consist largely of trial-and-error exploration of a high dimensional space that evolution has endowed with an unusually favorable distribution of attractors to desirable behaviors (Raphael et al. 2010). This advantageous hyperspace is the inevitable consequence of these phylogenetically retained subcortical circuits. As the later evolving cerebral cortex gradually asserted dominance over this circuitry, it would have been advantageous, perhaps even necessary, to exert its influence by modulating the gains of this circuitry rather than sending its commands directly to the actuators. Direct control of the actuators is largely absent (Rathelot and Strick 2009) and would be inconsistent with the surprisingly modest consequences of decorticating an adult mammal.

6 Reinterpreting emergent phenomena

6.1 The learning curve

If the central nervous system possessed a reasonably accurate inverse model of the sensorimotor plant, then it would be expected to be able to start with reasonably good performance on a novel motor task and then improve gradually but steadily as the internal model was updated to account for any discrepancies. As (Gallistel et al. 2004) elegantly pointed out, the notion that intelligent organisms gradually improve their performance along an asymptotic curve is a statistical construct that bears little resemblance to actual trial-by-trial performance. Instead, performance on successive trials looks more like a random walk, *running average* of which improves gradually. Even when error rates approach zero, the variability in quantitative performance from trial to trial remains high enough to trigger occasional errors. This is not a performance limitation arising from inherent noise in the motoneurons as postulated by many researchers (Todorov and Jordan 2002; Hamilton and Jones 2004). The variability can usually be decreased if the performance criteria for reward are further tightened. Much of the observed variability appears to arise during motor preparation (Churchland et al. 2006). It is as if the nervous system never ceases its gratuitous exploration of its control space. What purpose might that serve?

6.2 Noise versus intentional exploration

When psychophysicists study motor behavior under controlled conditions in the laboratory, they must make assumptions about the subjects' state of mind. The notion that subjects have an internal inverse model suggests that the subjects can be instructed regarding the cost function imposed by the investigator and can then compute an optimal strategy. The mean performance then represents this optimal strategy and the standard deviation should reflect computational noise in the nervous system and muscles. A cost function is typically divided into a term for performance criteria and a term for effort. Psychophysicists usually specify their performance criteria and make a guess about the subject's motivation to minimize effort, but what really matters is what is going on in the operator's brain. An individual trying to perform a task in the real world is unlikely to be given or to make use of formal performance criteria. Instead, as part of learning to perform the task, the individual simultaneously learns what constitutes "good enough" performance and how to achieve it most efficiently. In fact, investigators already know this, which is why subjects are usually offered the opportunity to "familiarize themselves with the apparatus." The assumption is that after that opportunity is over, everyone can get on with the experiment in a fully informed steady state.

6.2.1 Identifying performance criteria

If the reward is continuously scaled in magnitude according to objective performance criteria such as speed or accuracy, then the reward alone can be used by the subject to define optimality, but not until the subject has explored the space sufficiently to internalize the shape of the reward function. Even then, prior experience with the real world would suggest to the brain that the reward function could be continuously changing, requiring continuous exploration to keep minimizing the complete cost function. More typically, however, the reward is binary and success is determined by threshold performance criteria, resulting in many redundant solutions. The rewarded criteria are unknown to the subject and could be continuously changing. The only information available to the subject is sensory feedback from the task and the presence or the absence of the reward itself. In addition to performing the task to obtain the reward, the subject will be continuously solving a problem in Bayesian estimation, using the results of each trial as data to update a probability distribution of the reward criteria. This exploration usually requires the subject to introduce more variability into the trial-by-trial performance than would otherwise arise from noise. The subject gets the most information about the reward criteria by pushing the behavior to the currently estimated edge of acceptable. An observer has no way to know how much of the observed variability reflects this purposeful exploration versus computational noise [particularly if the measure of variability ignores time-series structure; (Slifkin and Newell 1998)]. The subjects themselves may not know.

6.2.2 Minimizing effort

The effort term in the cost function is even more difficult to pin down than the performance criteria. We can guess at three general types of effort that probably need to be represented: (i) *energetic cost* of the physical task itself; (ii) *attentional cost* in terms of distraction from other activities or events that might be important to pursue in parallel with the task; and (iii) *storage cost* of the current strategy, as defined above in terms of the numbers of neurons and synapses, state of which must be recalled to repeat the performance. In fact, experimentalists have no useful metrics for any of these, although they sometimes explore monotonic terms such as by changing loads to affect energetic cost or adding distracters for attentional cost. Unfortunately, storage cost depends on the entire life history of the subject and is unlikely to be monotonic for specific changes in the task, which might make the task more or less familiar to a given subject. Furthermore, storage cost depends on the number of neurons used by a strategy, which also affects neural noise but inversely. Even if the brain does have its own secret metrics, none of these components of effort is likely to be a function that the brain

can invert and minimize. Similar to the performance criteria, the effort terms must also be explored by deliberately introducing variability into the higher level command signals and considering the consequences according to sensory and efference copy signals from the lower levels. Again, the observer has no way to segregate observable variability into noise vs. purposeful exploration.

6.3 The robustness of multiple solutions

The *raison d'être* of optimal control is to find and employ the one best strategy for performing a task. If one assumes that a biological system has successfully employed optimal control, then observed variability in the observed strategies can be ascribed to noise. Evidence is accumulating, however, that this variability is not simply noise, even in systems that have no “intelligent behavior” and are incapable of intent. Eve Marder’s lab has measured the many parameters that define membrane channels and synaptic coupling among the neurons that generate oscillatory stomach contractions in the crab (Grashow et al. 2010; Goaillard et al. 2009). They found wide ranges of values among specimens, but only certain combinations of these values result in functional behavior. Only those combinations are found in nature, and so pooling the parametric data across specimens results in the erroneous interpretation that the parameters are independent and that their observed individual variability reflects biological noise. From a control perspective, the crabs seem to know that the high dimensional state space contains a large number of good-enough local minima in terms of the usual performance criteria. Why might it be useful for individuals to adopt various good-enough but suboptimal strategies rather than for all to use the optimal one? If the conditions under which the system must operate change, then some of their strategies may become dysfunctional, and so will be rejected either through ontogenic or phylogenic adaptation (Tang et al. 2010); the other, lucky ones survive. Apparently, Mother Nature has already discovered the lack of robustness inherent in optimal control (Valero-Cuevas et al. 2009). Similarly, the central pattern generator for locomotion in the cat produces many apparently stable variants in its broadly similar behaviors (Loeb 1993); the appearance of those variants can be influenced by introducing mechanical alterations into the musculoskeletal apparatus at an early age (Loeb 1999).

The existence of a state space with many good-enough solutions may actually be a precondition for the evolution of species. First, it guarantees population diversity even in the absence of genetic diversity, thereby increasing the likelihood that some individuals will survive a sudden change in environmental conditions. Second, it improves the likelihood that a given individual will survive the sort of mutation that is required for the evolution of new species. For example, it is highly unlikely that a given mutation in musculoskeletal

form of an individual will be accompanied by a simultaneous mutation in the control system that might be required to make optimal use of the new body form. The mechanically mutated organism must have a high likelihood of being able to achieve a functional state in order for the mutation it carries to persist long enough for subsequent evolutionary refinement to occur. All of the extant species that we study represent the product of countless evolutionary selections, each built upon a substrate that has necessarily been chosen for robust as well as successful performance as populations rather than individuals.

6.4 Optimal feedback control

It is important to distinguish useful predictions of a theory from testable hypotheses about underlying mechanisms. Both optimal control and good-enough control predict that organisms will usually take advantage of strategies that offer substantial improvements in performance. In contrast to good-enough control (at least so far), optimal control offers formal mathematical tools to predict features of useful strategies that experimentalists might observe. Many of these concern short-latency responses to brief perturbations during behavioral tasks, which are attractive because fewer neural pathways can contribute to such responses. If the entire plant (i.e., muscles, sensors, linkages, and their dynamics) can be described mathematically and the performance criteria stated formally as a cost function (e.g., combination of kinematic accuracy and level of effort), then it is possible to use an engineering tool called linear quadratic regulator (LQR) design to compute the optimal set of gains from all sensors to all actuators that will minimize costs over the set of all small perturbations for which the system remains locally linear (Loeb et al. 1990; He et al. 1991). The results bear a striking resemblance to many of the classically identified interneuronal circuits of the mammalian spinal cord. Unfortunately, this particular tool does not afford attribution of the computed net gains to any particular combination of such known circuits nor does it identify how background activity in such circuitry could contribute to the nominal patterns of muscle activation seen in the absence of perturbations. Furthermore, analytic computation of the LQR is not possible if the state equations governing the system cannot be inverted. These shortcomings of this “top-down” approach to understanding the spinal cord motivated the more recent “bottom-up” models in which the known circuitry of the spinal cord is modeled explicitly (Raphael et al. 2010; Tsianos et al. 2011). The resulting solutions predict both the nominal behavior and the response to perturbations. It remains to be seen whether the diversity of good-enough solutions found for the SLR model can be found in the internal states of biological organisms performing such behaviors.

Optimal feedback control has been employed more recently to explain the seemingly arbitrary optimization criteria that have been postulated to deal with redundancy of inverse models (Todorov and Jordan 2002). The inverse models assume perfect performance, whereas real components (especially biological ones) are noisy and sloppy. By assuming that the actuators generate motor noise with constant variance (Jones et al. 2002), it is possible to show that arbitrary constraints such as minimizing jerk (Flash and Hogan 1985) are actually emergent consequences of optimizing performance in the face of such noise. Optimal feedback control also exposes a fundamental flaw in the stepwise processes of inverting the task as outlined above (see History). For example, if the controller's strategy starts with computing the complete trajectory of the intended movement to a desired goal, then presumably the controller should use feedback during the execution of the movement to correct for any deviations in the trajectory that might arise due to external perturbations or internal noise. In fact, biological organisms tolerate a lot of deviations during the trajectory, and instead make only those adjustments that actually improve the accuracy at the goal (Todorov and Jordan 2002). In some situations, the organism actually tailors its nominal trajectory to minimize the different consequences of noise occurring at different times in the movement (Sabes and Jordan 1997). Optimal feedback control predicts that these behaviors are useful and should happen, but without explaining how they might be achieved by the nervous system. Good-enough control provides a biologically plausible mechanism (trial-and-error learning) that should converge on similar solutions when driven by an appropriate cost-function.

6.5 Habitual control and coaching

The ability to learn entirely new patterns of neuromuscular coordination such as required to make and use tools arose only very recently. It may include mutations of musculoskeletal structures (e.g., opposable thumb) the associated regulatory circuitry of which is not appropriate for these newly enabled tasks. Presumably these are the tasks for which the cortex requires much longer periods of trial-and-error experimentation to achieve even modest dexterity; consider a child learning to drink from a cup. Once learned, such fine motor habits tend to persist even when they are clearly suboptimal (de Rugy et al. 2012). Subjects learned to overcome real and virtual changes in their musculoskeletal systems but did so by rescaling their prior motor habits rather than recomputing optimal controls as would be expected if they could modify an internal inverse model.

As suggested above, various individuals performing tasks that look similar to an observer may have learned idiosyncratic controls for the highly redundant spinal-like regulator. This has important implications for learning a new task

such as required to play most sports. If the closest motor program in the individual's repertoire happens to be close to a good-enough program for the new task, then learning should be relatively easy. If it happens to be distant, then progress will be unsteady and slow. It may cease entirely unless and until the individual is forced into a distant part of the hyperspace. This would explain why sport coaches confronted with would-be athletes stuck in suboptimal motor habits often prescribe unusual exercises to force their pupils out of their comfort zone.

7 Conclusion

The methodology of optimal control is applicable when there is (i) a single, known cost function to be optimized, (ii) an invertible model of the plant, and (iii) simple noise interfering with optimal performance. None of these is likely to be true for biological organisms. Furthermore, their motivation is usually good-enough rather than globally optimal behavior. Even then, the performance of a biological organism is often much further from optimal than the physical limits of its hardware because the controller carries the software baggage of its entire ontogeny. Much of what has been called motor noise probably is not noise at all but rather a reflection of systematic exploratory behavior, probing that baggage and identifying the performance criteria. The behavior of the organism is thus required to be suboptimal to satisfy the requirements of system identification. The evolution of the organism itself is shaped at least as much by robustness as by optimality. The processes of identifying the requirements of tasks and building a useful repertoire of solutions are more instructive than the solutions themselves, which tend to be consistent with both computation of solutions online (i.e., optimal control) and look-up of previously stored solutions (i.e., good-enough control). Optimal control has been a useful metaphor for understanding some superficial aspects of motor psychophysics. Reductionists who want to understand the underlying neural mechanisms need to move on.

Acknowledgements Based in part on research funded by the DARPA REPAIR program. The opinions expressed are those of the author alone.

References

- Anderson FC, Pandy MG (1999) A dynamic optimization solution for vertical jumping in three dimensions. *Comput Methods Biomech Biomed Eng* 2:201–231
- Bernstein NA (1967) The co-ordination and regulation of movement. In: Whiting HTA (ed) (Translation) *Human motor actions: Bernstein reassessed*. Elsevier, Amsterdam
- Bjursten LM, Norrsell K, Norrsell U (1976) Behavioural repertory of cats without cerebral cortex from infancy. *Exp Brain Res* 25: 115–130

- Boyd SP, Vandenberghe L (2004) Convex optimization. Cambridge University Press, New York
- Buchanan TS, Shreeve DA (1996) An evaluation of optimization techniques for the prediction of muscle activation patterns during isometric tasks. *J Biomech Eng* 118:565–574
- Churchland MM, Afshar A, Shenoy KV (2006) A central source of movement variability. *Neuron* 52:1085–1096
- Crowninshield RD, Brand RA (1981) A physiologically based criterion of muscle force prediction in locomotion. *J Biomech* 14:793–801
- Culler E, Mettler FA (1934) Conditioned behavior in a decorticate dog. *J Comp Psychol* 18:291–303
- de Rugy A, Loeb GE, Carroll TJ (2012) Muscle coordination is habitual rather than optimal. *J Neurosci* 32:7384–7391
- Diamond IT (1979) The subdivisions of neocortex: a proposal to revise the traditional view of sensory, motor, and association areas. *Prog Psychobiol Physiol Psychol* 8:1–43
- Flanders M (2011) What is the biological basis of sensorimotor integration?. *Biol Cybern* 104:1–8
- Flash T, Hogan N (1985) The coordination of arm movements: an experimentally confirmed mathematical model. *J Neurosci* 5:1688–1703
- Friston K (2010) The free-energy principle: a unified brain theory. *Nat Rev Neurosci* 11:127–138
- Gallistel CR, Fairhurst S, Balsam P (2004) The learning curve: implications of a quantitative analysis. *Proc Natl Acad Sci USA* 101:13124–13131
- Gill PE, Murray W, Wright MH (1982) Practical optimization. Emerald Group Publishing Limited, West Yorkshire
- Goaillard J-M, Taylor AL, Schulz DJ, Marder E (2009) Functional consequences of animal-to-animal variation in circuit parameters. *Nat Neurosci* 12:1424–1430
- Grashow R, Brookings T, Marder E (2010) Compensation for variable intrinsic neuronal excitability by circuit-synaptic interactions. *J Neurosci* 30:9145–9156
- Hamilton AFd, Jones KE (2004) The scaling of motor noise with muscle strength and motor unit number in humans. *Exp Brain Res* 157:417–430
- Harper D (2010) optimal. In: Online Etymology Dictionary. http://www.etymonline.com/index.php?allowed_in_frame=0&search=optimal&searchmode=none
- He J, Levine WS, Loeb GE (1991) Feedback gains for correcting small perturbations to standing posture. *IEEE Trans Autom Control* 36:322–332
- Jones KE, Hamilton AFd, Wolpert DM (2002) Sources of signal-dependent noise during isometric force production. *J Neurophysiol* 88:1533–1544
- Kawato M (1999) Internal models for motor control and trajectory planning. *Curr Opin Neurobiol* 9:718–727
- Kawato M, Gomi H (1992) A computational model of four regions of the cerebellum based on feedback-error learning. *Biol Cybern* 68:95–103
- Lagarias JC, Reeds JA, Wright MH, Wright PE (1998) Convergence properties of the Nelder-Mead simplex method in low dimensions. *SIAM J OPTIM* 9:112–147
- Li Y, Levine WS, Loeb GE (2012) A two-joint human posture control model with realistic neural delays. *IEEE Trans Neural Syst Rehabil Eng*. doi:10.1109/TNSRE.2012.2199333
- Loeb GE (1983) Finding common ground between robotics and physiology. *Trends Neurosci* 6:203–204
- Loeb GE (1993) The distal hindlimb musculature of the cat: interanimal variability of locomotor activity and cutaneous reflexes. *Exp Brain Res* 96:125–140
- Loeb GE (1999) Asymmetry of hindlimb muscle activity and cutaneous reflexes after tendon transfers in kittens. *J Neurophysiol* 82:3392–3405
- Loeb GE, Levine WS, He J (1990) Understanding sensorimotor feedback through optimal control. *Cold Spring Harbor Symp Quant Biol* 55:791–803
- Marr D (1982) Vision. W.H. Freeman & Co, New York
- McCrea DA, Rybak IA (2008) Organization of mammalian locomotor rhythm and pattern generation. *Brain Res Rev* 57:134–146
- Osu R, Gomi H (1999) Multijoint muscle regulation mechanisms examined by measured human arm stiffness and EMG signals. *J Neurophysiol* 81:1458–1468
- Pandy MG, Anderson FC, Hull DG (1992) A parameter optimization approach for the optimal control of large-scale musculoskeletal systems. *J Biomech Eng* 114:450–460
- Pandy MG, Zajac FE, Sim E, Levine WS (1990) An optimal control model for maximum-height human jumping. *J Biomech* 23:1185–1198
- Partridge LD (1982) The good enough calculi of evolving control systems: evolution is not engineering. *Am J Physiol* 242:R173–R177
- Pedotti A, Krishnan VV, Stark L (1978) Optimization of muscle-force sequencing in human locomotion. *Math Biosci* 38:57–76
- Pellionisz AJ, Peterson BW (1988) A tensorial model of neck motor activation. In: Peterson BW, Richmond FJR (eds) Control of head movement. Oxford University Press, New York, pp 178–186
- Perfiliev S, Isa T, Johnels B, Steg G, Wessberg J (2010) Reflexive limb selection and control of reach direction to moving targets in cats, monkeys, and humans. *J Neurophysiol* 104:2423–2432
- Pierrot-Deseilligny E, Burke DC (2005) The circuitry of the human spinal cord: its role in motor control and movement disorders. Cambridge University Press, Cambridge
- Raikova RT, Gabriel DA, Aladjov H (2005) Experimental and modeling investigation of learning a fast elbow flexion in the horizontal plane. *J Biomech* 38:2070–2077
- Raphael E, Tsianos GA, Loeb GE (2010) Spinal-like regulator facilitates control of a two-degree-of-freedom wrist. *J Neurosci* 30:9431–9444
- Rathelot JA, Strick PL (2009) Subdivisions of primary motor cortex based on cortico-motoneuronal cells. *PNAS* 106:918–923
- Sabes PN, Jordan MI (1997) Obstacle avoidance and a perturbation sensitivity model for motor planning. *J Neurosci (New York)* 17:7119–7128
- Slifkin AB, Newell KM (1998) Is variability in human performance a reflection of system noise. *Curr Dir Psychol Sci* 7:170–177
- Tang LS, Goeritz ML, Caplan JS, Taylor AL, Fisek M, Marder E (2010) Precise temperature compensation of phase in a rhythmic motor pattern. *PLoS Biol* 8:e1000469
- Terekhov A, Zatsiorsky V (2011) Analytical and numerical analysis of inverse optimization problems: conditions of uniqueness and computational methods. *Biol Cybern* 104:75–93
- Todorov E, Jordan MI (2002) Optimal feedback control as a theory of motor coordination. *Nat Neurosci* 5:1226–1235
- Tsianos GA, Raphael G, Loeb GE (2011) Modeling the potentiality of spinal-like circuitry for stabilization of a planar arm system. *Prog Brain Res* 194:203–213
- Tsianos GA, Rustin C, Loeb GE (2012) Mammalian muscle model for predicting force and energetics during physiological behaviors. *IEEE Trans Neural Syst Rehabil Eng* 20:117–133
- Valero-Cuevas FJ, Hoffmann H, Kurse MU, Kutch JJ, Theodorou EA (2009) Computational models for neuromuscular function. *IEEE Rev Biomed Eng* 2:110–135
- Venkataraman P (2009) Applied optimization with MATLAB programming, Chapt 6. Wiley, Somerset
- Yuan G, Chang K, Hsieh C, Lin C (2010) A comparison of optimization methods for large-scale L1-regularized linear classification. *Linear Classif J Mach Learn Res* 2:3183–3234