



Published in final edited form as:

Nat Neurosci. 2004 September ; 7(9): 907–915.

Optimality principles in sensorimotor control (review)

Emanuel Todorov

Department of Cognitive Science, University of California San Diego

Abstract

The sensorimotor system is a product of evolution, development, learning, adaptation – processes that work on different time scales to improve behavioral performance. Consequently, many theories of motor function are based on the notion of optimal performance: they quantify the task goals, and apply the sophisticated tools of optimal control theory to obtain detailed behavioral predictions. The resulting models, although not without limitations, has explained a wider range of empirical phenomena than any other class of models. Traditional emphasis has been on optimizing average trajectories while ignoring sensory feedback. Recent work has redefined optimality on the level of feedback control laws, and focused on the mechanisms that generate behavior online. This has made it possible to fit a number of previously unrelated concepts and observations into what may become a unified theoretical framework for interpreting motor function. At the heart of the framework is the relationship between high-level goals, and the realtime sensorimotor control strategies most suitable for accomplishing those goals.

Optimality principles form the basis of many scientific theories. Their appeal lies in the ability to transform a parsimonious performance criterion into elaborate predictions regarding the behavior of a given system. Optimal control models of biological movement^{1–33} have explained behavioral observations on multiple levels of analysis (limb trajectories, joint torques, interaction forces, muscle activations or EMGs) and have arguably been more successful than any other class of models. Their advantages are both theoretical and practical. Theoretically, they are well justified *a priori*. This is because the sensorimotor system is the product of evolution, development, learning, adaptation – all of which are in essence optimization processes that continuously work to improve behavioral performance. Even if skilled performance on a certain task is not exactly optimal, but is just “good enough”, it has been made good enough by processes whose limit is optimality. Thus optimality provides a natural starting point for computational investigations³⁴ of sensorimotor function.

In practice, optimal control modeling affords unsurpassed autonomy and generality. Most alternative methods in engineering control, as well as alternative models of the neural control of movement, require as their input a detailed description of *how* the desired goal should be accomplished. For example, the Equilibrium Point hypothesis^{35,36} explains how a reference trajectory (presumably specified by the CNS) can be used to guide limb movement, but does not tell us how such a trajectory might be computed in tasks more complex than pointing. Similarly, the dynamical systems view³⁷ emphasizes that the composite neuro-musculo-skeletal system is a nonlinear dynamical system that can exhibit interesting phenomena such as bifurcations, but does not predict what nonlinear dynamics we should observe in a new task. In contrast, optimal control methods only require a performance criterion that describes *what* the goal is, and fill in all movement details automatically – by searching for the control strategy (or control law) that achieves the best possible performance. While the search process itself is sometimes cast as a model of behavioral change^{22,23,32,38,39}, most existing models focus on the outcome of that search – which corresponds to skilled performance.

Before the powerful machinery of optimal control can be applied, the modeler has to decide on: (1) a family of control laws under consideration along with a compatible musculo-skeletal model; (2) a quantitative definition of task performance. Performance is measured as the time-

integral of some instantaneously defined quantity called a cost function. The cost is a scalar function which depends on the current set of control signals (e.g. muscle activations) as well as the set of variables describing the current state of the plant and environment (e.g. joint angles and velocities, positions of relevant objects). The choice of cost has attracted a lot of attention in the literature, and in fact optimality models are often named “Minimum X” where X can be jerk, torque change, energy, time, variance. This choice is of course important, and not always transparent (see below). But the more fundamental distinction used here to categorize existing models – a distinction that leads to different views of sensorimotor processing – concerns the type of control law.

The first category of models reviewed below focus on open-loop control: they plan the best sequence of muscle activations (or joint torques, or limb postures), ignore the role of online sensory feedback, and usually assume deterministic dynamics. Such models differ mainly in the cost function they optimize, and often yield detailed and accurate predictions of behavior averaged over multiple repetitions of a task. But open-loop optimization has two serious limitations. First, it implies that the neural processing in the mosaic of brain areas involved in online sensorimotor control does little more than play a prerecorded movement tape – which is highly unlikely⁴⁰. Second, it fails to model trial-to-trial variability³⁰; the stereotypical movement patterns it implies are much more common in constrained laboratory tasks than in the real world⁴¹ – which is why such models have had limited impact in fields emphasizing unconstrained behavior (e.g. Development, Kinesiology, Physical Therapy, Sports Science).

The second category of models focus on closed-loop control: they construct the sensorimotor transformation (or feedback control law) that yields the best possible performance when motor noise as well as sensory uncertainty and delays are taken into account. These models predict not only average behavior, but also the task-specific sensorimotor contingencies that the CNS uses to make intelligent adjustments online. Such adjustments enable biological systems to “solve a control problem repeatedly rather than repeat its solution”⁴¹, and thus afford remarkable levels of performance in the presence of noise, delays, internal fluctuations, and unpredictable changes in the environment. Optimal feedback control has recently made it possible to unify a wide range of concepts and observations (e.g. kinematic regularities, motor synergies and controlled parameters, end-effector control, motor redundancy and structured variability, impedance control, speed-accuracy trade-offs) into a cohesive theoretical framework. It may further allow several important extensions: principles for hierarchical sensorimotor control^{33,42,43}, automated inference of task goals given movement data^{44–46}, and neuronal models of spinal^{24,28,29} as well as motor cortical function⁴⁷ (see also Box 1).

Open-loop optimization: Models of average behavior

The majority of existing optimal control models^{1–23} predict average movement trajectories or EMGs, by optimizing a variety of cost functions. Ideally, the cost assumed in an optimal control model should correspond to what the sensorimotor system is trying to achieve. But how can this be quantified? A rare case where the choice of cost is transparent are behaviors that require maximal effort – e.g. jumping as high as possible¹⁹, or producing maximal isometric fingertip force²⁰. Here one simply optimizes whatever subjects were asked to optimize, under the constraint that each muscle can produce limited force. The model predictions agree with data^{19,20}, especially in well-controlled fingertip force experiments where fine-wire EMGs from all participating muscles have been obtained²⁰.

In many cases however, the cost that is relevant to the sensorimotor system may not directly correspond to our intuitive understanding of “the task”, and so its detailed form should be considered a (relatively) free parameter. A recent attempt to measure that parameter, in a virtual

aiming task⁴⁸, suggests a cost for final position error that is quadratic when the error is small but saturates for larger error. It would be very useful to have a general data analysis procedure that infers the cost function given experimental data and a biomechanical model. Some results along these lines have been obtained in the computational literature^{44–46}, but a method applicable to Motor Control is still lacking. In the absence of empirically derived costs researchers have experimented with a variety of definitions, and found that motor behavior is near-optimal with respect to costs that are flexible and task-dependent. Below, open-loop models are sub-divided according to the cost they optimize.

Energy

Detailed optimal control modeling has its longest history in Biomechanics, and locomotion in particular, where most models minimize energy used by the muscles^{1–7}. While precise models of metabolic energy consumption that reflect the details of muscle physiology are rare, a number of cost functions that increase supra-linearly with muscle activation yield realistic and generally similar predictions⁴. Some models use optimization in a limited sense: they start with the experimentally measured kinematics of the gait cycle, and compute the most efficient muscle activations or joint torques that could cause the observed kinematics^{3,5}. Avoiding this limitation leads to more challenging dynamic optimization problems^{1,6}. In a recent model⁶ incorporating 23 mechanical degrees of freedom and 54 muscle-tendon actuators, only the initial and final posture of the gait cycle were specified empirically. The optimal sequences of muscle activations, joint torques, and body postures were then obtained by minimizing total energy. Considering how many details were predicted simultaneously (after 10000 hours of CPU time!), the agreement with kinematic, kinetic, and EMG measurements is striking⁶.

Smoothness

Energy minimization alone fails to account for average behavior in arm movements⁸, eye movements¹⁷, and some full-body movements such as standing from a chair¹¹. The usual remedy in such cases is a smoothness cost, that penalizes the time-derivative of hand acceleration^{9,10,12,13} (i.e. jerk), or joint torque^{14,15}, or muscle force¹¹. These models are less “ecological”: while the nervous system has obvious reasons to care about energetic efficiency or accuracy, it is much less clear *a priori* why smoothness might be important. Nevertheless, smoothness optimization has been rather successful in predicting average trajectories – particularly in arm movement tasks. The idea was first introduced in the minimum-jerk model^{9,10}, where it accounts for the straight paths and bell-shaped speed profiles of reaching movements, as well as a number of trajectory features in via-point tasks (where the hand is required to pass through a sequence of “via-points”). A more accurate but also more phenomenological model fits the speed profiles of arbitrary arm trajectories, by computing the minimum-jerk speed profile along the empirical movement path¹². It also captures the inverse relationship between speed and curvature, better than the 2/3 power law⁴⁹ previously used to quantify that phenomenon. The minimum-jerk model has been extended to the task of grasping, by using a cost that includes the smoothness of each fingertip trajectory plus a term that encourages perpendicular approach of the fingertips to the object surface¹³. This model of independent fingertip control explains many observations from grasping experiments (in particular the effects of object size on hand opening and timing), without having to invoke the previously proposed separation between hand transport and hand shaping. Smoothness optimization has also been formulated on the level of dynamics – by minimizing the time-derivative of joint torque^{14,15}. Interestingly, despite the nonlinear dynamics of the arm, the hand reaching trajectories predicted by this model are roughly straight in Cartesian space; in fact, their mild curvature agrees with experimental data. The minimum torque-change model also accounts for the lack of mirror symmetry observed in some via-point tasks^{14,15} – a phenomenon inconsistent with models that ignore the nonlinear arm dynamics.

Accuracy

The above models yield average behavior that achieves the task goals accurately. But one can be perfectly accurate on average and yet make substantial variable errors on individual trials. Within the limits of open-loop optimization, this issue was addressed by the minimum variance model¹⁶ where the sequence of muscle activations is planned so as to minimize the resulting variance of final hand positions (i.e. endpoint variance). Note that motor noise is known to be control-dependent – its magnitude is proportional to muscle activation^{21,50,51} – and so the choice of control signals affects movement variability. This form of optimization is loosely related to smoothness: nonsmooth movements require abrupt changes of muscle force, which require large EMG signals (to overcome the low-pass filtering properties of muscles), which lead to increased control-dependent noise. Indeed, this model predicts reaching trajectories very similar to the (successful) predictions of the minimum torque-change model, and also accounts for the inverse speed-curvature relationship found in elliptic movements⁴⁹. It will be interesting to see if the more general relationship between path and speed, that is captured by the path-constrained minimum-jerk model¹², can also be explained. The minimum variance model predicts, in impressive detail, the magnitude-dependent speed profiles of saccadic eye movements^{16,17}. For eye movements the predictions are also accurate on the level of muscle activations, but it is not yet clear if the same holds for arm movements. An extension to obstacle avoidance tasks¹⁸ accounts for the empirical relationship between the direction-dependent arm inertia and the margin by which the hand clears the obstacle⁵²; this is another phenomenon inconsistent with kinematic models. Note that in addition to average trajectories the minimum variance model also predicts the pattern of variability; but since it ignores feedback, and variability is significantly affected by feedback (especially in movements of longer duration), the latter prediction is less reliable.

Multi-attribute costs

Exploration of simple costs has illuminated the performance criteria relevant in different tasks. But the true performance criterion in most cases is likely to involve a mix of cost terms⁵³. Even if accuracy in the minimum variance sense¹⁶ can completely subsume smoothness optimization, it is clear that energetics is also a factor in many tasks. A cost function combining accuracy (under control-dependent noise) and energy was used to predict muscle directional tuning, i.e. how the activation of individual muscles varies with the direction of desired net muscle force²¹. Under quite general assumptions, it was shown analytically that the optimal tuning curve is either a full cosine or a truncated cosine – as observed empirically⁵⁴. Cosine tuning curves (for wrist muscles) were also predicted by a recent model that only minimizes energy²².

While open-loop models tend to optimize simple costs subject to boundary constraints (e.g. hand position, velocity and acceleration specified at the beginning and end of the movement¹⁰), such constraints are inapplicable to stochastic problems where the final state is affected by noise. Instead, stochastic models have to use final accuracy costs in addition to whatever costs are defined during the movement^{30,31}. The modeler then has the non-trivial task of assigning relative weights to quantities that have different units (e.g. metabolic energy and endpoint error). It may be possible to automate this, using an algorithm⁵⁵ that converts probabilistic constraints (e.g. a threshold on endpoint variance) into multi-attribute costs. Note that such costs can significantly enrich optimality models: different weight settings yield different predictions, which can be tested experimentally by varying the relative importance of different aspects of task performance. The speed-accuracy trade-offs discussed below are an example of that. Also, the weights that define a multi-attribute cost could be used as command signals by higher-level control centers; this may play an important role in future optimality models that have hierarchical structure (see last section).

Closed-loop optimization: Models of sensorimotor integration

Instead of focusing on average behavior, which reflects neural information processing rather indirectly, sensorimotor integration can be modeled much more directly via closed-loop optimization^{24–27,30–33}. Here both sensory and motor noise are incorporated in the biomechanical model, and the optimization is performed over the family of all possible feedback control laws. As explained next, this class of models can address all phenomena that open-loop models address, and much more.

Optimal feedback control – a unifying framework

What is optimal feedback control, and how is it related to optimal open-loop control? Both optimization procedures start with a cost defining the task goals, as well as an initial state (fig 1). Open-loop optimization then yields a “desired” movement. Since open-loop control makes little sense in the presence of noise, the movement plan is usually thought to be executed by a feedback controller – which uses some servo mechanism to cancel the instantaneous deviations between the desired and actual state of the plant. That mechanism however is predefined, and is not taken into consideration in the optimization phase. In contrast closed-loop optimization treats the feedback mechanism as being fully programmable, i.e. it constructs the best possible transformation from states of the plant and environment into control signals. The resulting controller does whatever is needed to accomplish the task: instead of relying on preconceived notions of what control schemes the sensorimotor system might use, optimal feedback control lets the task and biomechanical model dictate the control scheme that best suits them. This may yield a force-control scheme in an isometric task where a target force level is specified, or a position-control scheme in a postural task where a target limb position is specified; but in less trivial tasks the optimal control scheme will generally be one that we do not yet have a name for. Such flexibility, however hard to grasp, matches the flexibility and resourcefulness apparent in motor behavior⁴¹.

The numerical methods used to approximate optimal feedback controllers are rather involved^{56–61}, and computationally expensive. One class of such methods – Temporal Difference Reinforcement Learning⁶⁰ – have been used with remarkable success to model many aspects of reward-related neural activity^{62–65}. Almost all available methods are based on the fundamental concept of long-term performance, quantified by an optimal cost-to-go (or value) function. For every state and point in time, this function tells us how much cost (or reward) is expected to accumulate from now until the end of the movement, assuming we behave optimally. Box 1 clarifies the optimal cost-to-go function, its importance in the computation of optimal controls, and its potential role in future analyses and models of motor cortical activity.

The above discussion implied that feedback controllers map actual states into control signals. But when the state of a stochastic plant is observable only through delayed and noisy sensors, the controller has to rely on an internal estimate of the state (fig 1b). The resulting controls are optimal only when the state estimator is also optimal, i.e. Bayesian. Such an estimator takes into account sensory data, recent control signals, knowledge of plant dynamics, as well as its earlier output, and weights all these sources of information regarding the current state in proportion to their reliability. In modeling practice one typically uses a Kalman filter – which is the optimal estimator when the plant dynamics are linear and the noise is Gaussian, and provides a good approximation in other cases⁵⁸. A number of studies suggest that perception in general⁶⁶, and online state estimation in particular^{67–69}, are based on the principles of Bayesian inference. A key feature of optimal estimators is their ability to anticipate state changes before the corresponding sensory data has arrived. This requires either explicit or implicit knowledge of plant dynamics, i.e. an “internal model”. There is a growing body of both psychophysical^{70–73} and neurophysiological^{74,75} evidence in support of this notion,

although critics point out that some of it is rather indirect⁷⁶. The formation of internal models through adaptation was initially interpreted in the context of movement planning⁷⁰; recent results^{77–81} however paint a much more complex picture, and suggest³⁹ the kind of flexibility that optimal feedback control affords. Note that here I am only referring to what are usually called internal forward models – as distinguished from internal inverse models. The latter are thought to transform task goals into motor commands; but since this is the job of a controller, I believe the “inverse model” terminology should be avoided.

While the distinction between open- and closed-loop control was traditionally seen as a dichotomy worth debating, it is increasingly realized that one is simply a special case of the other⁸². Because the optimal feedback controller is driven by an optimal state estimate rather than raw sensory input, it responds appropriately to any information supplied by the estimator – regardless of whether that information reflects immediate sensory data, or past experiences, or predictions about the future. The anticipative capabilities of the estimator allow the controller to counteract disturbances before they cause errors – by generating net muscle force when the direction and time course of the disturbance are predictable^{70,83}, or by adjusting muscle coactivation (and thereby limb stiffness and damping) when only the magnitude of the disturbance is known^{78,84}. It is technically straightforward to incorporate an open-loop control signal in the feedback control law, by treating time as another state variable; however, adaptation experiments⁸⁵ reveal a strong preference for associating time-varying forces with limb positions and velocities rather than time. Optimal feedback control yields a servo controller when the task explicitly specifies a limb trajectory to be traced, and approaches open-loop control when sensory noise or delays become very large. Thus the differences between the two classes of models are most salient when: (i) the movement duration allows time for sensory-guided adjustments; (ii) the task goal can be achieved by a large variety of movements. Most everyday behaviors have these properties.

Scaling laws and online corrections

Stochastic optimal control explains one of the most thoroughly investigated properties of discrete movements: the scaling of duration with amplitude and desired accuracy^{86,87}. This scaling is quantified by Fitts’ law, which states that movement duration is a linear function of the logarithm of movement amplitude divided by target width. Fitts’ law is predicted by both intermittent⁸⁸ and continuous²⁵ optimal feedback control models of reaching. The essential ingredient of these models is the control-dependent nature of motor noise^{21,50,51} – which makes faster movements less accurate. The recent minimum-variance model¹⁶ also predicts Fitts’ law, via open-loop optimization. But feedback control explains²⁵ an important additional observation: the increased duration of more accurate movements is due to a prolonged deceleration phase, making the speed profiles significantly skewed⁸⁹. Such skewing is optimal because the largest motor commands (and consequently most of the control-dependent noise) are generated early in the movement, and so the feedback mechanism has more time to detect and correct the resulting errors.

Optimal feedback control provides a natural framework for studying the responses to experimentally-induced perturbations. In human standing, various features of ankle-hip trajectories observed during postural adjustments have been explained²⁷. The smooth correction of the hand movement towards a displaced target⁸² – which occurs even when subjects are unaware of the displacement – is also well explained²⁵. An interesting phenomenon, that appears to contradict optimality models (as well as other models⁹⁰), is the systematic under-compensation for target displacements introduced late in the movement⁹¹. But this turns out to have an explanation³⁹, as follows. The feedback controller is optimized for a world where targets do not jump. So it takes advantage of stationarity – by lowering positional feedback gains towards the end of a reaching movement, and using negative velocity

and force feedback to stop without oscillation. Consequently, large positional errors introduced in the stopping phase are not fully compensated.

Redundancy, synergies, uncontrolled manifolds, and minimal intervention

A recent theory of motor coordination^{30,92}, based on optimal feedback control, brings together a number of key ideas which have stimulated Motor Control research since the pioneering work of Bernstein⁴¹. Here they are illustrated intuitively in context of the simplest redundant task (fig 2); for a mathematically precise description see⁹².

Redundancy means that the same goal can be achieved in many different ways – e.g. many limb trajectories can bring the fingertip to the same target. This is characteristic of most everyday tasks⁴¹ and raises the problem of choosing one out of all possible solutions, i.e. the “redundancy problem”. Such abundance of solutions is actually beneficial to the sensorimotor system (since it makes the search for a solution more likely to succeed), but is a problem for the researcher trying to understand how this choice is made. An appealing aspect of optimal control is the principled manner in which it resolves redundancy – by choosing the best possible control law. The latter is usually well defined: in the presence of noise, different feedback controllers yield different performance even when the resulting average behavior is the same⁴². Mathematical analysis reveals that optimal feedback controllers resolve redundancy online. They obey a “minimal intervention” principle: make no effort to correct deviations away from the average behavior unless those deviations interfere with task performance^{30,92}. This is because acting (and making corrections in particular) is expensive – due to control-dependent noise and energy costs. It follows that solving a redundant task according to a detailed movement plan (e.g. tracking a prespecified reach trajectory) is a suboptimal strategy, regardless of how the plan is chosen. It also follows that experimentally-induced perturbations should be resisted in a goal-directed manner, so that task performance is recovered although the corrected movement may differ from baseline; this has been observed repeatedly^{41,93–95}.

Optimal feedback control confirms Bernstein’s intuition⁴¹ that the substantial yet structured variability of biological movements is not due to sloppiness, but on the contrary, is an indication that the sensorimotor system is exceptionally well designed. If we think of body configurations as being vectors in a multi-dimensional state space (fig 2), a redundant task is one where the state vector can vary in certain directions without interfering with the goal. A control law that obeys the minimal intervention principle has the effect of “pushing” the state vector orthogonally to the redundant directions. The state is also being perturbed by motor noise, and so the probability distribution of observed states reflects the balance of these two influences. When motor noise is isotropic, the resulting covariance ellipse is elongated in redundant directions. Such effects have been observed in a wide range of behaviors^{30,41,96–98}, and were recently quantified by the “uncontrolled manifold” method for comparing task-relevant vs. redundant variances⁹⁶. Note that optimal feedback control literally creates an uncontrolled manifold: there are directions in which the control law does not act. A different control law that acts in all directions, e.g. by pushing the state towards the center of the plot in fig 2, can further reduce variance in the redundant direction but only by increasing variance in the task-relevant direction. Thus, allowing large variability in the redundant direction is necessary for achieving optimal performance. Note however that variability structure does not necessarily arise from redundancy; instead it may reflect structure in motor noise. An example is the distribution of reach endpoints – which is known to be elongated in the (non-redundant) movement direction^{99,100}. This is likely due to the fact that muscles pulling along the movement are more active, and therefore more affected by control-dependent noise. In movements of longer duration the anisotropy of endpoint distributions is reduced^{99,100}, probably because the feedback controller has more time to make corrections. Optimal feedback

control reproduces³⁹ these findings for the reasons just outlined. Another example is template-drawing, where the variance of hand position is modulated similarly to hand speed (both in experimental data and optimal control simulations³⁰) even though the drawing task suppresses positional redundancy.

Lack of control action in certain directions implies correlations among control signals (fig 2 is an extreme case where the two controls are always equal). Principal Components Analysis (PCA) on a correlated dataset always yields a small number of principal components (PCs) that account for a large percentage of the variance. Thus, optimal feedback control predicts that PCA-related methods applied to EMG data will find evidence for reduced dimensionality – which is indeed the case^{101,102}. Such PCs correspond to the idea of motor synergies^{41, 103}, or high-level “control knobs” thought to affect a few important features of the state while leaving many others uncontrolled. This is precisely what the control law in fig 2 does: it extracts from the two-dimensional state vector a single task-relevant feature, and generates controls which selectively affect that feature. Fig 3 illustrates the emergence of synergies and structured variability in a more complex task (see legend). By representing the feedback controller as a neural network, we can see more clearly how it compresses the 12-dimensional state vector into just two task-relevant features, and then expands them into a vector of 5 control signals (whose variance can therefore be accounted for by only two PCs). We also see the emergence of movement regularities that reflect biomechanical structure rather than control objectives.

Hierarchical sensorimotor control in the context of optimality

Sensorimotor function results from multiple feedback loops that operate simultaneously: tunable muscle stiffness and damping provide instantaneous feedback; the spinal cord generates the fastest neural feedback; slower but more adaptable loops are implemented in somatosensory and motor cortices, visuo-motor loops involve parietal cortex, etc. The different latencies may be arranged so that slower but more intelligent loops respond to perturbations just when the faster loops are about to run out of steam¹⁰⁴. An important step towards understanding how this complex mosaic produces integrated action was Bernstein’s analysis¹⁰⁵, translated in part only recently¹⁰⁶. It suggested a four-level functional hierarchy for human motor control: posture and muscle tone, muscle synergies, dealing with three-dimensional space, and organizing complex actions that pursue more abstract goals. It also suggested that any one behavior involves at least two levels of neural feedback control: a leading level that monitors progress and exploits the many different ways of achieving the goal, and a background level that provides automatism and corrections without which the leading level could not function. Note that although the task-relevant features illustrated in fig 2 and 3 are reminiscent of a higher level of control, all feedback models discussed so far involved a single sensorimotor transformation.

Computational modeling that aims to capture the essence of feedback control hierarchies – via optimization^{29,33,42,43} or otherwise¹⁰⁷ – is still in its infancy. Anatomically specific models^{29,42} emphasize the distinction between spinal and supra-spinal processing and take into account our partial understanding of spinal circuitry; such models may prove very useful in elucidating spinal cord function, especially in lower species. Models^{33,43} that aim to explain complex behavior emphasize a functional hierarchy: the low level of neural feedback augments or transforms the dynamics of the musculo-skeletal system, so that the high level “sees” a composite dynamical system that is easier to (learn how to) control optimally. One way to construct an appropriate low level is through unsupervised learning – which captures the statistical regularities present in the flow of motor commands and corresponding sensory data⁴³. Note that unsupervised learning has a long history in the sensory domain¹⁰⁸ where it has been used to model neural coding in primary visual cortex¹⁰⁹ as well as the auditory nerve¹¹⁰. Another approach is inspired by the minimal intervention principle: if we guess the

task-relevant features that the optimal feedback controller will use in the context of a specific task, then we can design a low level feedback controller that extracts those features, sends them to the high level, and maps the descending commands (which signal desired changes in the task features) into appropriate muscle activations³³. When coupled with optimal feedback control on the high level, both of these approaches yield hierarchical controllers that are approximately optimal – at a fraction of the computational effort required to optimize a non-hierarchical feedback controller^{33,43}. Related ideas have also been pursued in robotics^{111, 112}. Apart from modeling sensorimotor function, such hierarchical methods hold promise for real-time control of complex robotic prostheses as well as electrical stimulators¹¹³ implanted in multiple muscles.

A number of “end-effector” control models^{26,88,114–117}, formulated in the context of reaching tasks, are related to the present discussion. They postulate kinematic feedback control mechanisms that monitor progress of the hand towards the target, and issue desired changes in hand position or joint angles. These models do not specify how muscles are activated by a low-level controller. In the stochastic iterative correction model⁸⁸, the hand and target positions are compared and corrective submovements towards the target are (intermittently) generated; submovement magnitudes are chosen so as to minimize control-dependent noise. The vector-integration-to-endpoint model¹¹⁶ is similar, except that here the comparison is continuous, and the hand-target vector is multiplied by a time-varying GO signal which maps distance into speed. The minimum-jerk¹⁰ model of trajectory planning has also been transformed into a feedback control model: at each point in time, a new minimum-jerk trajectory appropriate for the remainder of the movement is computed (starting at the current state), and its initial portion is used for instantaneous control²⁶. The latter is an example of model-predictive control – an approach gaining popularity in engineering¹¹⁸. Two related models^{114,117} take us a step further – to the level of joint kinematics. An earlier model¹¹⁴ proposed that each joint is moved autonomously, in proportion to how much moving that joint alone affects the hand-target distance. In retrospect, the joint increments computed in this way correspond to the gradient (i.e. list of partial derivatives) of the hand-target distance with respect to the joint angles. However, the idea of gradient-following was formalized and compared to data only recently¹¹⁷. A related (although rather abstract) approach to coordination in the presence of redundancy is afforded by Tensor Network Theory¹¹⁵. Note that the latter models are in effect optimization models: the gradient is the minimal change in joint angles that results in unit displacement of the hand towards the target.

It is not yet clear how the predictions of hierarchical models of optimal sensorimotor control will differ from non-hierarchical ones. But as the following example illustrates, consideration of low-level feedback loops may be essential for reaching correct conclusions. One of the most prominent arguments³⁶ in favor of trajectory planning is the observation¹¹⁹ that a limb perturbed in the direction of a target “fights back”, as if to return to some moving virtual attractor. This is indeed inconsistent with kinematic feedback control that always pushes the hand towards the target. But in the context of hierarchical optimal control, the phenomenon has a simple explanation that does not involve a hypothetical trajectory plan: (i) neural feedback is delayed while tunable muscle stiffness and damping act instantaneously; (ii) the neural controller knows that, and coactivates muscles preemptively so as to ensure an immediate response when unexpected perturbations arise; (iii) muscles are dumb feedback controllers which, once coactivated, resist perturbations both away from the target and towards it. Thus, a response which appears task-inappropriate from a kinematic perspective may actually be optimal when muscle properties, noise, sensorimotor delays, and the need for stability are taken into account⁴².

Acknowledgements

We thank Gerald Loeb and Jochen Triesch for their comments on the manuscript. This work was supported by NIH grant R01-NS045915.

References

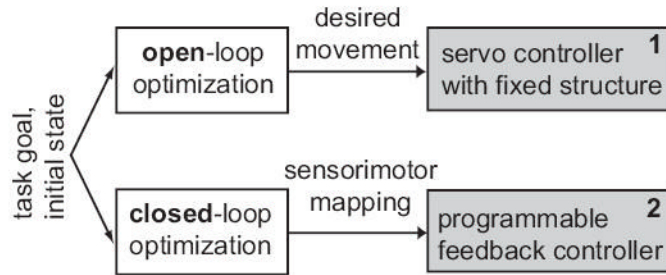
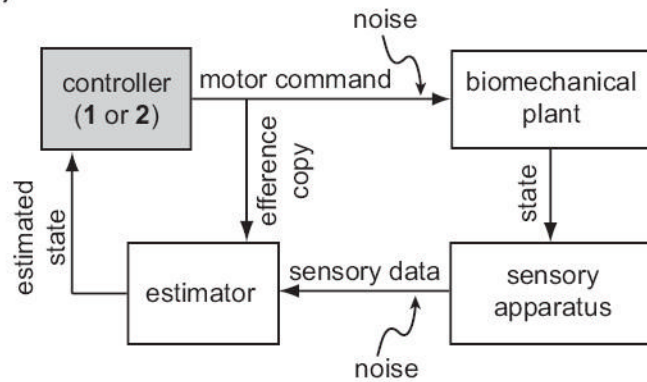
1. Chow CK, Jacobson DH. Studies of human locomotion via optimal programming. *Math Biosciences* 1971;10:239–306.
2. Hatze H, Buys JD. Energy-optimal controls in the mammalian neuromuscular system. *Biol Cybern* 1977;27:9–20. [PubMed: 889959]
3. Davy DT, Audu ML. A dynamic optimization technique for predicting muscle forces in the swing phase of gait. *J Biomech* 1987;20:187–201. [PubMed: 3571299]
4. Collins JJ. The redundant nature of locomotor optimization laws. *J Biomech* 1995;28:251–267. [PubMed: 7730385]
5. Popovic D, Stein RB, Oguztoreli N, Lebedowska M, Jonic S. Optimal control of walking with functional electrical stimulation: a computer simulation study. *IEEE Trans Rehabil Eng* 1999;7:69–79. [PubMed: 10188609]
6. Anderson FC, Pandy MG. Dynamic optimization of human walking. *J Biomech Eng* 2001;123:381–390. [PubMed: 11601721]
7. Rasmussen J, Damsgaard M, Voigt M. Muscle recruitment by the min/max criterion - a comparative numerical study. *J Biomech* 2001;34:409–415. [PubMed: 11182135]
8. Nelson WL. Physical principles for economies of skilled movements. *Biological Cybernetics* 1983;46:135–147. [PubMed: 6838914]
9. Hogan N. An Organizing Principle for a Class of Voluntary Movements. *Journal of Neuroscience* 1984;4:2745–2754. [PubMed: 6502203]
10. Flash T, Hogan N. The coordination of arm movements: an experimentally confirmed mathematical model. *The Journal of Neuroscience* 1985;5:1688–1703. [PubMed: 4020415]
11. Pandy MG, Garner BA, Anderson FC. Optimal control of non-ballistic muscular movements: a constraint-based performance criterion for rising from a chair. *J Biomech Eng* 1995;117:15–26. [PubMed: 7609479]
12. Todorov E, Jordan MI. Smoothness maximization along a predefined path accurately predicts the speed profiles of complex arm movements. *Journal of Neurophysiology* 1998;80:696–714. [PubMed: 9705462]
13. Smeets JB, Brenner E. A new view on grasping. *Motor Control* 1999;3:237–271. [PubMed: 10409797]
14. Uno Y, Kawato M, Suzuki R. Formation control of optimal trajectory in human multijoint arm movement: Minimum torque-change model. *Biological Cybernetics* 1989;61:89–101. [PubMed: 2742921]
15. Nakano E, et al. Quantitative Examinations of Internal Representations for Arm Trajectory Planning: Minimum Commanded Torque Change Model. *Journal of Neurophysiology* 1999;81:2140–2155. [PubMed: 10322055]
16. Harris CM, Wolpert DM. Signal-dependent noise determines motor planning. *Nature* 1998;394:780–784. [PubMed: 9723616]
17. Harwood MR, Mezey LE, Harris CM. The spectral main sequence of human saccades. *J Neurosci* 1999;19:9098–9106. [PubMed: 10516327]
18. Hamilton AF, Wolpert DM. Controlling the statistics of action: obstacle avoidance. *J Neurophysiol* 2002;87:2434–2440. [PubMed: 11976380]
19. Pandy MG, Zajac FE, Sim E, Levine WS. An optimal control model for maximum-height human jumping. *J Biomech* 1990;23:1185–1198. [PubMed: 2292598]
20. Valero-Cuevas FJ, Zajac FE, Burgar CG. Large index-fingertip forces are produced by subject-independent patterns of muscle excitation. *Journal of Biomechanics* 1998;31:693–703. [PubMed: 9796669]

21. Todorov E. Cosine tuning minimizes motor errors. *Neural Computation* 2002;14:1233–1260. [PubMed: 12020444]
22. Fagg AH, Shah A, Barto AG. A computational model of muscle recruitment for wrist movements. *J Neurophysiol* 2002;88:3348–3358. [PubMed: 12466451]
23. Ivanchenko V, Jacobs RA. A developmental approach aids motor learning. *Neural Computation* 2003;15:2051–2065. [PubMed: 12959665]
24. Loeb GE, Levine WS, He J. Understanding sensorimotor feedback through optimal control. *Cold Spring Harb Symp Quant Biol* 1990;55:791–803. [PubMed: 2132855]
25. Hoff, B. A computational description of the organization of human reaching and prehension. Ph.D. Thesis, University of Southern California, (1992).
26. Hoff B, Arbib MA. Models of trajectory formation temporal interaction of reach grasp. *J Mot Behav* 1993;25:175–192. [PubMed: 12581988]
27. Kuo AD. An optimal control model for analyzing human postural balance. *IEEE Transactions on Biomedical Engineering* 1995;42:87–101. [PubMed: 7851935]
28. Shimansky YP. Spinal motor control system incorporates an internal model of limb dynamics. *Biological Cybernetics* 2000;83:379–389. [PubMed: 11039702]
29. Ijspeert AJ. A connectionist central pattern generator for the aquatic terrestrial gaits of a simulated salamander. *Biological Cybernetics* 2001;84:331–348. [PubMed: 11357547]
30. Todorov E, Jordan M. Optimal feedback control as a theory of motor coordination. *Nature Neuroscience* 2002;5:1226–1235.
31. Guigon, E., Baraduc, P. & Desmurget, M. Constant effort computation as a determinant of motor behavior. In *Advances in Computational Motor Control 2* (2003). www.acmc-conference.org
32. Shimansky YP, Kang T, He JP. A novel model of motor learning capable of developing an optimal movement control law online from scratch. *Biological Cybernetics* 2004;90:133–145. [PubMed: 14999480]
33. Li, W., Todorov, E. & Pan, X. Hierarchical optimal control of redundant biomechanical systems. *Proceedings of the 26th Annual International Conferences of the IEEE Engineering in Medicine and Biology Society*, in press (2004).
34. Marr, D. *Vision*. Freeman and Company, (1982).
35. Feldman AG, Levin MF. The origin use of positional frames of reference in motor control. *Behavioral and Brain Sciences* 1995;18:723–744.
36. Bizzi E, Hogan N, Mussa-Ivaldi FA, Giszter S. Does the nervous system use equilibrium-point control to guide single multiple joint movements? *Behavioral and Brain Sciences* 1992;15:603–613.
37. Kelso, J.A.S. *Dynamic patterns: The self-organization of brain and behavior*. MIT Press, Cambridge, MA (1995).
38. Sporns O, Edelman GM. Solving Bernstein's problem: a proposal for the development of coordinated movement by selection. *Child Dev* 1993;64:960–81. [PubMed: 8404271]
39. Todorov, E. Interpreting motor adaptation results within the framework of optimal feedback control. In *Advances in Computational Motor Control 1* (2002). www.acmc-conference.org
40. Kalaska, J., Sergio, L.E. & Cisek, P. Sensory guidance of movement: Novartis Foundation Symposium. Glickstein, M. (ed.) (John Wiley & Sons, Chichester, UK, 1998).
41. Bernstein, N.I. *The Coordination and Regulation of Movements*. Pergamon Press, (1967).
42. Loeb GE, Brown IE, Cheng EJ. A hierarchical foundation for models of sensorimotor control. *Exp Brain Res* 1999;126:1–18. [PubMed: 10333003]
43. Todorov, E. & Ghahramani, Z. Unsupervised learning of sensory-motor primitives. In *Proceedings of the 25th Annual International Conference of the IEEE Engineering in Biology and Medicine Society*, pp 1750–1753 (2003).
44. Kalman R. When is a linear control system optimal? *Trans AMSE J Basic Eng, Ser D* 1964;86:51–60.
45. Moylan P, Anderson B. Nonlinear regulator theory and an inverse optimal control problem. *IEEE Trans Automatic Control* 1973;AC-18:460–465.
46. Ng, A. & Russell, S. Algorithms for inverse reinforcement learning. In *Proceedings of the 17th International Conference on Machine Learning* (2000).

47. Scott S. Optimal feedback control the neural basis of volitional motor control. *Nature Neuroscience Reviews* 2004;5:534–546.
48. Kording K, Wolpert D. The loss function of sensorimotor learning. *PNAS* 2004;101:9839–9842. [PubMed: 15210973]
49. Lacquaniti F, Terzuolo C, Viviani P. The law relating the kinematic figural aspects of drawing movements. *Acta Psychol* 1983;54:115–130.
50. Sutton GG, Sykes K. The variation of hand tremor with force in healthy subjects. *Journal of Physiology* 1967;191(3):699–711. [PubMed: 6051798]
51. Schmidt RA, Zelaznik H, Hawkins B, Frank JS, Quinn JT Jr. Motor-output variability: a theory for the accuracy of rapid motor acts. *Psychol Rev* 1979;86:415–451. [PubMed: 504536]
52. Sabes PN, Jordan MI, Wolpert DM. The role of inertial sensitivity in motor planning. *J Neurosci* 1998;18:5948–5957. [PubMed: 9671681]
53. Rosenbaum, D.A. *Human motor control*. Academic press, San Diego (1991).
54. Hoffman DS, Strick PL. Step-Tracking Movements of the Wrist. {IV}. Muscle Activity Associated with Movements in Different Directions. *Journal of Neurophysiology* 1999;81:319–333. [PubMed: 9914292]
55. Zhu G, Rotea MA, Skelton R. A convergent algorithm for the output covariance constraint control problem. *Siam Journal on Control and Optimization* 1997;35:341–361.
56. Bryson, A. & Ho, Y. *Applied optimal control*. Blaisdell Publishing Company, Massachusetts (1969).
57. Kirk, D. *Optimal control theory: An introduction*. Prentice Hall, New Jersey (1970).
58. Davis, M.H.A. & R.B., V. *Stochastic Modelling and Control*. Chapman and Hall, (1985).
59. Bertsekas, D. & Tsitsiklis, J. *Neuro-dynamic programming*. Athena Scientific, Belmont, MA (1997).
60. Sutton, R.S. & Barto, A.G. *Reinforcement Learning: An Introduction*. MIT Press, (1998).
61. Li, W. & Todorov, E. Iterative linear quadratic regulator design for nonlinear biological movement systems. *Proceedings of the 1st International Conference on Informatics in Control, Automation, and Robotics*, in press (2004).
62. Montague PR, Dayan P, Person C, Sejnowski TJ. Bee Foraging in Uncertain Environments Using Predictive Hebbian Learning. *Nature* 1995;377:725–728. [PubMed: 7477260]
63. Schultz W, Dayan P, Montague PR. A neural substrate of prediction reward. *Science* 1997;275:1593–1599. [PubMed: 9054347]
64. O'Doherty J, et al. Dissociable roles of ventral and dorsal striatum in instrumental conditioning. *Science* 2004;304:452–454. [PubMed: 15087550]
65. Seymour B, et al. Temporal difference models describe higher-order learning in humans. *Nature* 2004;429:664–667. [PubMed: 15190354]
66. Knill, D. & Richards, W. *Perception as Bayesian inference*. Cambridge University Press, (1996).
67. Wolpert D, Gharahmani Z, Jordan M. An Internal Model for Sensorimotor Integration. *Science* 1995;269:1880–1882. [PubMed: 7569931]
68. Kording KP, Wolpert DM. Bayesian integration in sensorimotor learning. *Nature* 2004;427:244–247. [PubMed: 14724638]
69. Saunders JA, Knill DC. Visual feedback control of hand movements. *Journal of Neuroscience* 2004;24:3223–3234. [PubMed: 15056701]
70. Shadmehr R, Mussa-Ivaldi FA. Adaptive representation of dynamics during learning of a motor task. *Journal of Neuroscience* 1994;14:3208–3224. [PubMed: 8182467]
71. Flanagan JR, Wing AM. The role of internal models in motion planning control: Evidence from grip force adjustments during movements of hand-held loads. *Journal of Neuroscience* 1997;17:1519–1528. [PubMed: 9006993]
72. Kawato M. Internal models for motor control trajectory planning. *Curr Opin Neurobiol* 1999;9:718–27. [PubMed: 10607637]
73. Flanagan JR, Lolley S. The inertial anisotropy of the arm is accurately predicted during movement planning. *J Neurosci* 2001;21:1361–1369. [PubMed: 11160407]

74. Li CS, Padoa-Schioppa C, Bizzi E. Neuronal correlates of motor performance motor learning in the primary motor cortex of monkeys adapting to an external force field. *Neuron* 2001;30:593–607. [PubMed: 11395017]
75. Gribble PL, Scott SH. Overlap of internal models in motor cortex for mechanical loads during reaching. *Nature* 2002;417:938–941. [PubMed: 12087402]
76. Ostry DJ, Feldman AG. A critical evaluation of the force control hypothesis in motor control. *Experimental Brain Research* 2003;153:275–288.
77. Thoroughman KA, Shadmehr R. Learning of action through adaptive combination of motor primitives. *Nature* 2000;407:742–747. [PubMed: 11048720]
78. Burdet E, Osu R, Franklin DW, Milner TE, Kawato M. The central nervous system stabilizes unstable dynamics by learning optimal impedance. *Nature* 2001;414:446–449. [PubMed: 11719805]
79. Wang T, Dordevic GS, Shadmehr R. Learning the dynamics of reaching movements results in the modification of arm impedance long-latency perturbation responses. *Biol Cybern* 2001;85:437–448. [PubMed: 11762234]
80. Flanagan JR, Vetter P, Johansson RS, Wolpert DM. Prediction precedes control in motor learning. *Curr Biol* 2003;13:146–150. [PubMed: 12546789]
81. Korenberg, A. Computational and psychophysical studies of motor learning. PhD Thesis, University College London, (2003).
82. Desmurget M, Grafton S. Forward modeling allows feedback control for fast reaching movements. *Trends Cogn Sci* 2000;4:423–431. [PubMed: 11058820]
83. Gribble PL, Ostry DJ. Compensation for interaction torques during single-multijoint limb movement. *J Neurophysiol* 1999;82:2310–2326. [PubMed: 10561408]
84. Humphrey, D.R. & Reed, D.J. *Advances in Neurology: Motor Control Mechanisms in Health and Disease*. Desmedt, J.E. (ed.), pp. 347–372 (Raven Press, New York, 1983).
85. Condit MA, Mussa-Ivaldi FA. Central representation of time during motor learning. *PNAS* 1999;96:11625–11630. [PubMed: 10500227]
86. Woodworth RS. The accuracy of voluntary movement. *Psychol Rev Monogr* 1899;3
87. Fitts P. The information capacity of the human motor system in controlling the amplitude of movements. *J Exp Psychol* 1954;47:381–391.
88. Meyer DE, Abrams RA, Kornblum S, Wright CE, Smith JEK. Optimality in human motor performance: Ideal control of rapid aimed movements. *Psychological Review* 1988;95:340–370. [PubMed: 3406245]
89. Milner TE, Ijaz MM. The effect of accuracy constraints on three-dimensional movement kinematics. *Neuroscience* 35 1990;365:374.
90. Flash T, Henis E. Arm Trajectory Modifications During Reaching Towards Visual Targets. *Journal of Cognitive Neuroscience* 1991;3:220–230.
91. Komilis E, Pelisson D, Prablanc C. Error processing in pointing at randomly feedback-induced double-step stimuli. *Journal of Motor Behavior* 1993;25:299–308. [PubMed: 15064196]
92. Todorov, E. & Jordan, M. *Advances in neural information processing systems 15*. Becker, S., Thrun, S. & Obermayer, K. (eds.), pp. 27–34 (MIT Press, Cambridge, MA, 2002).
93. Cole KJ, Abbs JH. Kinematic and electromyographic responses to perturbation of a rapid grasp. *J Neurophysiol* 1987;57:1498–510. [PubMed: 3585477]
94. Gracco VL, Abbs JH. Dynamic control of the perioral system during speech: kinematic analyses of autogenic nonautogenic sensorimotor processes. *J Neurophysiol* 1985;54:418–32. [PubMed: 4031995]
95. Robertson EM, Miall RC. Multi-joint limbs permit a flexible response to unpredictable events. *Exp Brain Res* 1997;117:148–52. [PubMed: 9386013]
96. Scholz JP, Schoner G. The uncontrolled manifold concept: identifying control variables for a functional task. *Exp Brain Res* 1999;126:289–306. [PubMed: 10382616]
97. Scholz JP, Schoner G, Latash ML. Identifying the control structure of multijoint coordination during pistol shooting. *Exp Brain Res* 2000;135:382–404. [PubMed: 11146817]
98. Li ZM, Latash ML, Zatsiorsky VM. Force sharing among fingers as a model of the redundancy problem. *Exp Brain Res* 1998;119:276–86. [PubMed: 9551828]

99. Gordon J, Ghilardi MF, Cooper S, Ghez C. Accuracy of planar reaching movements. II. Systematic extent errors resulting from inertial anisotropy. *Exp Brain Res* 1994;99:112–130. [PubMed: 7925785]
100. Messier J, Kalaska JF. Comparison of variability of initial kinematics endpoints of reaching movements. *Exp Brain Res* 1999;125:139–52. [PubMed: 10204767]
101. D'Avella A, Saltiel P, Bizzi E. Combinations of muscle synergies in the construction of a natural motor behavior. *Nat Neurosci* 2003;6:300–308. [PubMed: 12563264]
102. Ivanenko YP, et al. Temporal components of the motor patterns expressed by the human spinal cord reflect foot kinematics. *J Neurophysiol* 2003;90:3555–3565. [PubMed: 12853436]
103. Latash, M.L. *Motor Control, Today and Tomorrow*. Gantchev, G., Mori, S. & Massion, J. (eds.), pp. 181–196 (Academic Publishing House "Prof. M. Drinov", Sofia, 1999).
104. Nicols TR, Houk JC. Improvement in Linearity Regulations of Stiffness That Result from Actions of Stretch Reflex. *Journal of Neurophysiology* 1976;39:119–142. [PubMed: 1249597]
105. Bernstein, N.I. *On the construction of movements*. Medgiz, Moscow (1947).
106. Bernstein, N.I. *Dexterity and its development*. Latash, M.L. & Turvey, M. (eds.) (Lawrence Erlbaum, 1996).
107. Cisek P, Grossberg S, Bullock D. A cortico-spinal model of reaching proprioception under multiple task constraints. *Journal of Cognitive Neuroscience* 1998;10:425–444. [PubMed: 9712674]
108. Barlow, H. *Sensory Communication*. Rosenblith, W. (ed.) (MIT Press, Cambridge, MA, 1961).
109. Olshausen BA, Field DJ. Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature* 1996;381:607–609. [PubMed: 8637596]
110. Lewicki MS. Efficient coding of natural sounds. *Nature Neuroscience* 2002;5:356–363.
111. Khatib O. A unified approach to motion and force control of robotic manipulators: The operational space formulation. *IEEE Journal of Robotics and Automation* 1987;RA-3:43–53.
112. Pratt J, Chew CM, Torres A, Dilworth P, Pratt G. Virtual model control: An intuitive approach for bipedal locomotion. *International Journal of Robotics Research* 2001;20:129–143.
113. Loeb GE, Peck RA, Moore WH, Hook K. BIONTM system for distributed neural prosthetic interfaces. *Medical engineering & physics* 2001;23:9–18. [PubMed: 11344003]
114. Hinton GE. Parallel computations for controlling an arm. *Journal of Motor Behavior* 1984;16:171–194. [PubMed: 14713664]
115. Pellionisz A. Coordination: a vector-matrix description of transformations of overcomplete CNS coordinates a tensorial solution using the Moore-Penrose generalized inverse. *Journal of Theoretical Biology* 1984;110:353–375. [PubMed: 6503306]
116. Bullock D, Grossberg S. Neural dynamics of planned arm movements: emergent invariants speed-accuracy properties during trajectory formation. *Psychological Review* 1988;95:49–90. [PubMed: 3281179]
117. Torres EB, Zipser D. Reaching to grasp with a multi-jointed arm. I. Computational model. *Journal of Neurophysiology* 2002;88:2355–2367. [PubMed: 12424277]
118. Camacho, E.F. & Bordons, C. *Model Predictive Control*. Springer-Verlag, London (1999).
119. Bizzi E, Accornero N, Chapple W, Hogan N. Posture control trajectory formation during arm movement. *J Neurosci* 1984;4:2738–44. [PubMed: 6502202]
120. Todorov E. Direct Cortical Control of Muscle Activation in Voluntary Arm Movements: a Model. *Nature Neuroscience* 2000;3:391–398.

(a) planning**(b) execution****Fig 1. Schematic illustration of open- and closed-loop optimization**

(a) The optimization phase, which corresponds to planning or learning, starts with a specification of the task goal and the initial state. Both approaches yield a feedback control law, but in the case of open-loop optimization the feedback portion of the control law is predefined and not adapted to the task. **(b)** Either feedback controller can be used online to execute movements, although controller 2 will generally yield better performance. The estimator needs an efference copy of recent motor commands in order to compensate for sensory delays. Note that the estimator and controller are in a loop; thus they can continue to generate time-varying commands even if sensory feedback becomes unavailable. Noise is typically modeled as a property of the sensorimotor periphery, although a significant portion of it may originate in the nervous system.

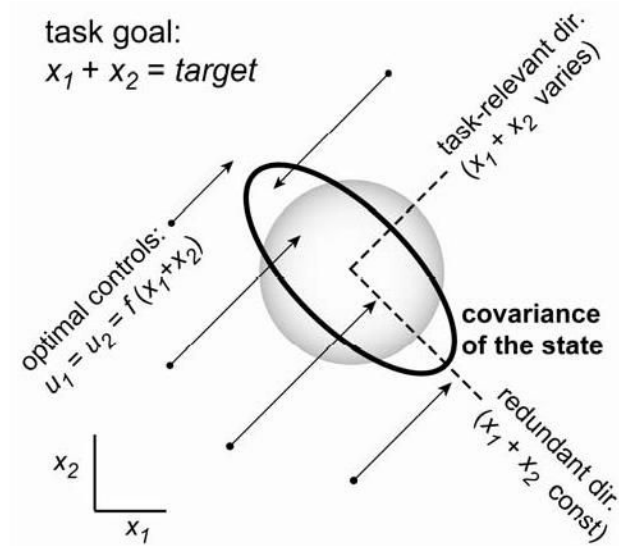
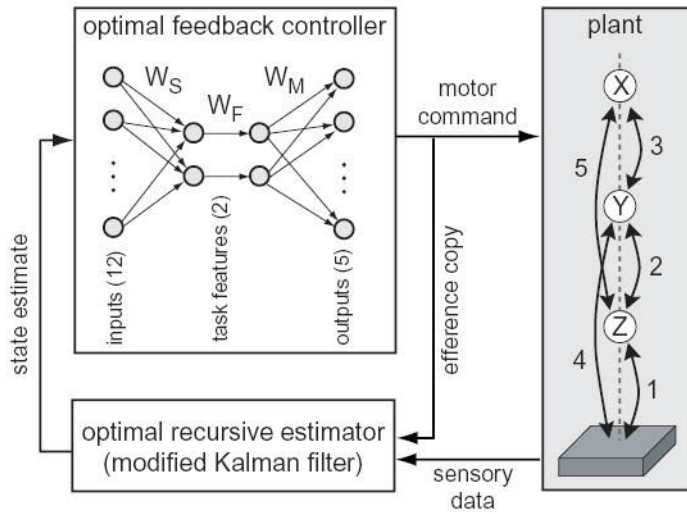


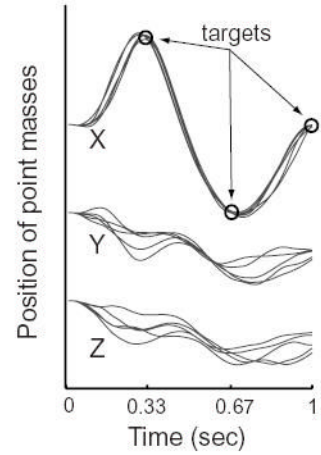
Fig 2. Minimal intervention principle

Illustration of the simplest redundant task, adapted from³⁰. x_1, x_2 are two uncoupled state variables, each driven by a corresponding control signal u_1, u_2 in the presence of control-dependent noise. The task is to maintain $x_1 + x_2 = target$ and use small controls. The optimal u_1 and u_2 are equal – to a function that depends on the task-relevant feature $x_1 + x_2$ but not on the individual values of x_1 and x_2 . Thus u_1 and u_2 form a motor synergy. Arrows show that the optimal controls push the state vector orthogonally to the redundant direction (along which $x_1 + x_2$ is constant). This direction is then an uncontrolled manifold. The black ellipse is the distribution of final states, obtained by sampling the initial state from a circular Gaussian, and applying the optimal control law for one step. The gray circle is the distribution under a different control law, that tries to maintain $x_1 = x_2 = target/2$ by pushing the state towards the center of the plot. Such a control law can reduce variance in the redundant direction as compared to the optimal control law, but only at the expense of increased variance in the task relevant direction, as well as increased control signals (not shown). See³⁰ for technical details.

(a) schematic illustration of control system and plant



(c) simulations



(b) weight matrices in neural network

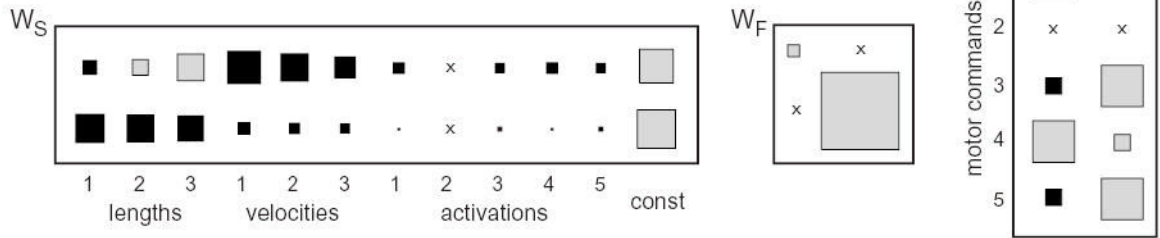
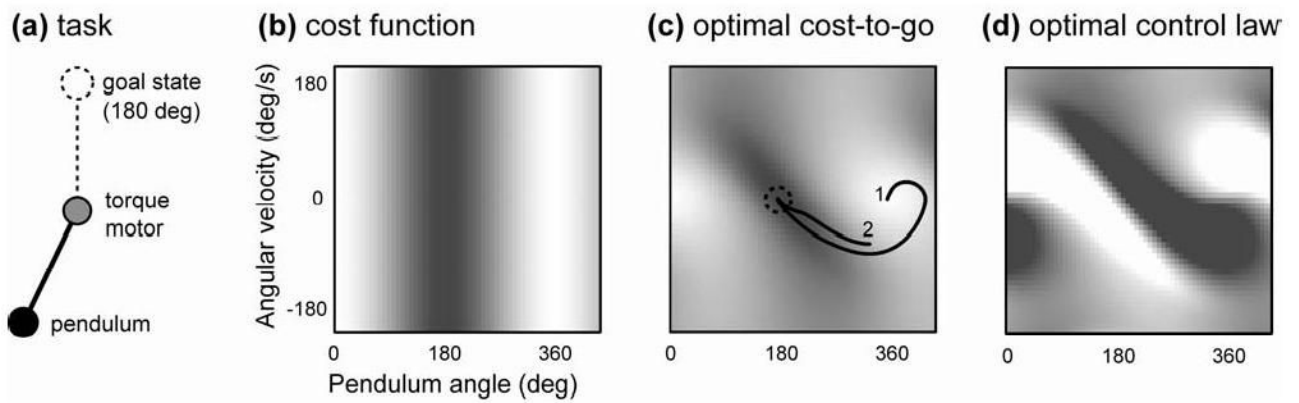


Fig 3. Application of optimal feedback control to a redundant stochastic system

(a) The plant is composed of 3 point masses (X,Y,Z) and 5 actuated visco-elastic links, moving up and down in the presence of gravity³⁰. The task requires point mass X (i.e. the “end-effector”) to pass through specified targets at specified points in time. The state vector includes the lengths and velocities of links 1–3, the activation states of all actuators (modeled as low-pass filters), and the constant 1 (needed for technical reasons). The optimal feedback controller in this case is a 5x12 time-varying matrix. To understand how this matrix transforms estimated states into control signals, it was averaged over time and represented as a linear neural network (using Singular Value Decomposition). (b) Weight matrices in the neural network (color denotes sign, area denotes absolute value, ‘x’ denotes zero weight). The rows of W_S correspond to the task-relevant features being extracted; W_F are feedback gains; the columns of W_M are motor synergies. The bottom feature (with much bigger gain) extracts something closely related to end-effector position, by summing the lengths of links 1–3. The structure of the motor synergies reflects the symmetries of plant: links 3 and 5 (which act on the end-effector) are treated as a unit; links 1 and 4 (which transmit to the ground the forces generated by 3 and 5) are treated as another unit; link 2 is not actuated at all. (c) Trajectories of the point masses from 5 simulation runs. The trajectories of the end-effector are overall more repeatable than the other two point masses; also, the end-effector trajectories themselves show less variability when passing through the targets – as observed in via-point tasks³⁰. These are both examples of variability structure arising from the minimal intervention principle. Note that the distance between the two intermediate point masses Y and Z is kept constant on average; this is an

interesting emergent property due to the structure of the optimal motor synergies (which in turn reflect the structure of the plant).



Box 1. Properties of the optimal cost-to-go function

Consider the task (a) of making a pendulum swing up as quickly as possible. The pendulum is driven by a torque motor with limited output, and has to overcome gravity. Since this is a second-order system, the state vector includes the pendulum angle and angular velocity. The cost function penalizes the vertical distance away from the upright position (b) as well as the squared torque output. If we attempt to minimize this cost greedily, by always pushing up, the pendulum will never rise above some intermediate position where gravity balances the maximal torque the motor can generate. The only way to overcome gravity is to swing in one direction, and then accelerate in the opposite direction. This is similar to hitting and throwing tasks, where we have to move our arm back before accelerating it forward. The important point here is that the cost function itself does not directly suggest such a strategy. Indeed, the relationship between costs and optimal controls is rather subtle, and is mediated by another function: the *optimal cost-to-go*. For each state, this function tells us how much cost we will accumulate from now until the end of the movement, assuming we choose controls optimally. The optimal cost-to-go obeys a self-consistency condition known as Bellman's optimality principle: the optimal cost-to-go at each state (c) is found by considering every possible control at that state, adding the control cost to the optimal cost-to-go for the resulting next state, and taking the minimum of these sums. The latter minimization also yields the optimal control; in (d) the color corresponds to the optimal torque as a function of the pendulum state (black: max negative; white: max positive). Plot (c) shows two optimal trajectories starting at different states. One uses the strategy of swinging back and then forward; the other goes straight to the goal because the initial velocity is sufficient to overcome gravity.

Note that both trajectories in (c) are moving roughly downhill along the optimal cost-to-go surface (i.e. from light to dark). This is because, for a large class of problems, the vector of optimal control signals can be computed by taking the negative gradient of the optimal cost-to-go function, and multiplying it by a matrix that reflects plant dynamics and energy costs^{30,92}. This gradient, known in control theory as the *costate* vector, is a vector with the same dimensionality as the state; it tells us how to change the state so as to increase the cost-to-go most rapidly. Now imagine that the costate vector is encoded by some population of neurons – which would not be surprising given its fundamental role in the computation of optimal controls. Since optimal controls are obtained from the costate via simple matrix multiplication, the activities of these neurons can directly drive muscle activation. This is reminiscent of a model¹²⁰ of direct cortical control of muscle activation, and suggests that the costate vector is something that might be encoded in the output of primary motor cortex. What does the costate look like? As explained above, it is related to how the state varies under the action of the optimal controller; so if the state includes position and velocity, the costate might resemble a mix of velocity and acceleration. But this relationship is loose; the only general way to find the true costate is to solve the optimal control problem.