

From Task Parameters to Motor Synergies: A Hierarchical Framework for Approximately Optimal Control of Redundant Manipulators

• • • • •

• • • • •

Emanuel Todorov*

*Department of Cognitive Science
University of California San Diego
La Jolla, California 92093-0515
e-mail: todorov@cogsci.ucsd.edu*

Weiwei Li

*Department of Mechanical and Aerospace
Engineering
University of California San Diego
La Jolla, California 92093-0411*

Xiuchuan Pan

*Department of Cognitive Science
University of California San Diego
La Jolla, California 92093-0515*

Received 4 June 2005; accepted 7 June 2005

We present a hierarchical framework for approximately optimal control of redundant manipulators. The plant is augmented with a low-level feedback controller, designed to yield input-output behavior that captures the task-relevant aspects of plant dynamics but has reduced dimensionality. This makes it possible to reformulate the optimal control problem in terms of the augmented dynamics, and optimize a high-level feedback controller without running into the curse of dimensionality. The resulting control hierarchy compares favorably to existing methods in robotics. Furthermore, we demonstrate a number of similarities to (nonhierarchical) optimal feedback control. Besides its engineering applications, the new framework addresses a key unresolved problem in the neural control of movement. It has long been hypothesized that coordination involves selective control of task parameters via muscle synergies, but the link between these parameters and the synergies capable of controlling them has remained elusive. Our framework provides this missing link. © 2005 Wiley Periodicals, Inc.

*To whom all correspondence should be addressed.

1. INTRODUCTION

The control of complex redundant systems is a challenging problem, of interest in both robotics and biological motor control. The nonlinear dynamics and high-dimensional state and control spaces of such systems prevent the use of many traditional methods for controller design. One way to tackle very complex problems is through divide-and-conquer strategies. Indeed, the most advanced control system known to date—the nervous system—appears to rely on such strategies. Sensorimotor control occurs simultaneously on many levels.^{1–3} Lower-level circuits interact with the musculoskeletal system directly: They receive rich sensory input, and generate corresponding motor output before the rest of the brain has had time to react to that input. Higher-level circuits interact with an augmented plant, that consists of the lower levels and the musculoskeletal system. The lower levels perform a (not well understood) transformation, allowing higher levels to operate on increasingly more abstract and more goal-related movement representations.⁴

Here, we propose a hierarchical control scheme inspired by this general organization of the sensorimotor system, as well as by prior work on hierarchical control in robotics.^{5–7} We focus on two-level feedback control hierarchies as illustrated in Figure 1. The low-level controller receives information about the plant state x , and generates an abstract and more compact state representation $y(x)$ that is sent to the high level. The high-level controller monitors task progress, and issues commands $v(y)$ which in general specify how y should change. The job of the low-level controller is to compute energy-efficient controls $u(v, x)$ consistent with v . Thus the low-level controller does not solve a specific subtask (as usually assumed in hierarchical reinforcement learning),^{8,9} but instead performs an instantaneous feedback

transformation. This enables the high level to control y unencumbered by the full details of the plant.

While the proposed scheme is designed to resemble the sensorimotor system on a structural level, achieving functional resemblance is equally important. Functionally, sensorimotor control is best described as being near optimal.¹⁰ It may seem surprising that a hierarchical controller can closely approximate an optimal controller. But as we have shown elsewhere,^{11,12} optimal feedback controllers for redundant systems exhibit hierarchical organization similar to Figure 1, even when such organization is not imposed by design. This finding provides yet another motivation for the present scheme: If full-blown optimization on redundant tasks is known to yield hierarchical structure, it makes sense to restrict the optimization to an (appropriately chosen) family of hierarchical controllers.

The general idea that the brain monitors a small number of task parameters y instead of the full state x , generates abstract commands v , and maps them into muscle activations u using motor synergies, has been around for a long time.^{13,14} A number of concrete models of end-effector control have been formulated in the context of reaching tasks.^{15–20} The high-level state in such models is assumed to be hand position, the abstract command is desired velocity in hand space or in joint space, and the high-level controller is a simple positional servo. While these models are related to our work, in some sense they leave all the hard questions unanswered: It is unclear how the task parameters are actually controlled (i.e., what the corresponding muscle synergies are), and whether this choice of task parameters can yield satisfactory performance. We address these questions here.

Our framework is related in interesting ways to input-output feedback linearization^{21,22} as well as to the operational space formulation⁶—which also fit in the general scheme in Figure 1. These methods yield linear dynamics on the high level, by cancelling the plant nonlinearities at the low level. However, many systems of interest cannot be linearized, and furthermore it is not clear that linearization is desirable in the first place. Suppressing the natural plant dynamics may require large control signals—which are energetically expensive, and also increase error in systems subject to control-multiplicative noise (a universal characteristic of biological movement).^{23–25} In contrast, we summarize the plant dynamics on the high level and thus create opportunities for exploiting them. Recent work in biped locomotion²⁶ underscores the potential of such approaches. In general, our objective is dimensionality reduction rather than




Figure 1. Schematic illustration of the proposed framework.

linearization. This is because we are pursuing optimality, and what makes optimal control hard is the curse of dimensionality. We are also pursuing neurobiological relevance, and it is clear that a number of behaviors (including locomotion²⁷ and arm movements)²⁸ are optimized with respect to the non-linear musculoskeletal dynamics.

The rest of the paper is organized as follows. Section 2 summarizes our previous findings relating optimal feedback control to hierarchical control, introduces relevant concepts from biological motor control, and illustrates them with experimental data. This motivates the formal development of the new framework in Section 3. The relation to existing methodologies is elaborated in Section 4. Numerical comparisons to robotic control methods as well as to optimal control are presented in Section 5. Preliminary results have been reported in conference proceedings.²⁹

2. BIOLOGICAL MOTIVATION, AND RELATION TO OPTIMAL CONTROL

Optimal control models of biological movement have a long history, and provide satisfying computational accounts of many behavioral phenomena.¹⁰ The majority of these models are formulated in open loop, and plan a desired trajectory while ignoring the role of online feedback. Desired trajectory planning 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 unlikely.⁴ Consequently, we and others^{30–33} have focused on optimal feedback control models, that predict not only average behavior but also the task-specific sensorimotor contingencies used to generate intelligent adjustments online. Such adjustments enable biological systems to “solve a control problem repeatedly rather than repeat its solution,”¹³ and thus afford remarkable levels of performance in the presence of noise, delays, internal fluctuations, and unpredictable changes in the environment. The latter models have significantly extended the range of phenomena addressed by open-loop optimization.¹⁰ Some of the key results from this line of work are directly relevant to hierarchical control, as described next.

2.1. Minimal Intervention, Synergies, and Uncontrolled Manifolds

Our work on motor coordination^{11,12} revealed that when the task is redundant, the optimal feedback




Figure 2. Properties of optimal feedback controllers in redundant tasks (reprinted from Ref. 10).

controller exhibits hierarchical structure even though it is not specifically designed to be hierarchical. Figure 2 illustrates this finding in the context of a simple redundant task—where two state variables x_1, x_2 with dynamics $x_i(t+1) = ax_i(t) + u_i(t)(1 + \epsilon_i(t))$, $i = 1, 2$ have to be controlled so that $x_1 + x_2$ equals a given target T . The noise model is control multiplicative; details can be found in ref. 11. The optimal controller pushes the system state toward the nearest point on the manifold of acceptable states (“redundant direction” in Figure 2). This clearly requires less control energy compared to pushing toward a predefined point (say $x_1 = x_2 = T/2$). It also causes less variability in the direction that matters (compare state covariance to gray circle). Note however that variability in the redundant direction is increased—because the optimal controller is not correcting deviations in that direction. Analysis of optimal value functions for stochastic nonlinear control problems has shown that the structure in Figure 2 is a general property of optimal feedback controllers: Deviations from the average behavior are corrected selectively, only when they compromise task performance.¹¹ We have called this the *minimal intervention* principle.

This is an important result because redundancy is ubiquitous in the motor system. The phenomenon of increased variability along redundant directions (illustrated below) has been observed in a wide range of behaviors,¹¹ and has been quantified via the “uncontrolled manifold” method for comparing task-relevant and task-irrelevant variance.³⁴ It implies that the substantial yet structured variability of biological movements is not due to sloppiness, but on the contrary, is a signature of an exceptionally well-designed sensorimotor system.

The hierarchical nature of the optimal controller is evident in Figure 2: u_1, u_2 are coupled, and equal to some function of $x_1 + x_2$ rather than being functions of the individual state variables x_1, x_2 . In the terminology of Figure 1, the high-level state is $y = x_1 + x_2$, the high-level control is $v = f(y)$, and the vector of actual control signals is $\mathbf{u} = [v; v]$. In more complex tasks, our analysis has shown that the optimal control law still has this property, but for a high-level state \mathbf{y} which may not be easy to guess. This is because \mathbf{y} is related to the shape of the optimal value function (defined as the sum of future costs under an optimal control law) rather than the cost function, and the optimal value function is hard to compute. Nevertheless, these results^{11,12} provide strong motivation for hierarchical approximations to optimal control in redundant tasks.

2.2. Empirical Illustration of Structured Motor Variability

Selective control of task-relevant parameters—which is a signature of both optimal feedback control and hierarchical control schemes—implies increased variability in task-irrelevant parameters. This phenomenon is illustrated here with data from the following experiment. Six subjects were asked to move repeatedly between three targets attached to their body (one target on each leg, and one on the left upper arm). Subjects held a 25 cm wooden pointer in their right hand, and always used the same tip (TIP 1) to touch the center of the targets. We measured the position and orientation of the center of the pointer with a Polhemus Liberty sensor (240 Hz sampling rate), which allowed us to calculate the positions of both TIP 1 and TIP 2. The experimental setup is shown in Figure 3(a). The goal of the experiment was to quantify the positional variance of the movement paths.

The data were analyzed as follows. The start and end times of each movement were found using a velocity threshold, applied when TIP 1 was within 10 cm of a target. Outlier movements, that failed to reach within 10 cm of their target, were eliminated. Each movement was resampled in space, at equal intervals with length 2% of the total path length for that movement (this resampling is needed to suppress temporal misalignment). The mean and covariance for each sample point were computed, separately for each subject and movement target. Positional variance was defined as the trace of the covariance matrix. The means and covariances were then averaged over subjects, and plotted in Figure

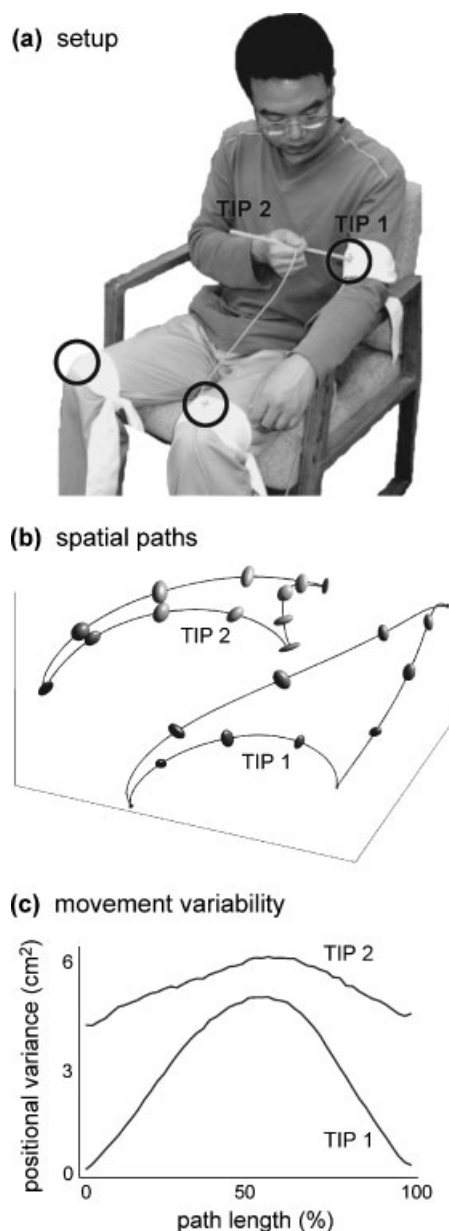


Figure 3. Experimental illustration of increased variability in redundant dimensions.

3(b) (± 1 standard deviation ellipsoids are shown at 25% intervals). To facilitate comparison, we further averaged over the three targets and plotted the positional variances of the two tips of the pointer in Figure 3(c).

The main result shown in Figure 3(c) nicely illustrates the typical structure of motor variability. The position of the task-relevant TIP 1 is less variable than the task-irrelevant TIP 2. Furthermore,

both variances are smaller in the task-relevant portions of the movement (when the pointer approaches the target) compared to the middle of the movement. The latter effect is much stronger for the task-relevant TIP 1.

These observations make it clear that the nervous system is not using the classic robotics approach—which is to plan a trajectory in joint space, and then execute that trajectory using some combination of computed torque and servo control. This is not surprising. Since a joint-space trajectory plan contains no information regarding task relevance, the execution system would have no choice but to track faithfully all details of the plan—including the irrelevant ones. This would result in increased energy consumption, and increased errors due to control-dependent noise.

3. HIERARCHICAL CONTROL FRAMEWORK

Key to our framework is the low-level controller $\mathbf{u}(\mathbf{v}, \mathbf{x})$ —whose design we address first, assuming that the high-level parameters $\mathbf{y}(\mathbf{x})$ and their desired dynamics have been given. We pay specific attention to the case when the controls \mathbf{u} do not affect the high-level parameters \mathbf{y} instantaneously. We then discuss how the desired \mathbf{y} -dynamics can be adapted to the natural plant dynamics, and how the high-level controller $\mathbf{v}(\mathbf{y})$ can be designed using optimal control techniques.

3.1. Low-Level Controller Design

Consider the dynamical system

$$\dot{\mathbf{x}}(t) = \mathbf{a}(\mathbf{x}(t)) + B(\mathbf{x}(t))\mathbf{u}(t), \quad (1)$$

where $\mathbf{x} \in \mathbb{R}^{n_x}$ is the state vector, $\mathbf{u} \in \mathbb{R}^{n_u}$ is the control vector, $\mathbf{a}(\mathbf{x})$ are the passive dynamics, and $B(\mathbf{x})\mathbf{u}$ are the control-dependent dynamics (assumed linear in the control). We are interested in controllers that achieve low cumulative cost, for a cost-per-step function of the form

$$\ell(t, \mathbf{x}(t), \mathbf{u}(t)) = q(t, \mathbf{x}(t)) + r(\mathbf{u}(t), \mathbf{x}(t)). \quad (2)$$

The high-level state vector $\mathbf{y} \in \mathbb{R}^{n_y}$, $n_y \leq n_x$, is a static function of the plant state:

$$\mathbf{y} = \mathbf{h}(\mathbf{x}). \quad (3)$$

The function \mathbf{h} defining the high-level state \mathbf{y} is supplied by the user. What we provide is a way to “control” \mathbf{y} , that is, to generate controls \mathbf{u} which affect \mathbf{x} in such a way that the corresponding changes in \mathbf{y} are as desired. We assume that \mathbf{y} contains enough information to compute the state-dependent cost $q(t, \mathbf{x})$; in other words, there exists a function \tilde{q} such that

$$\tilde{q}(t, \mathbf{h}(\mathbf{x})) = q(t, \mathbf{x}). \quad (4)$$

Other than that, the function \mathbf{h} should be chosen to yield the lowest-dimensional representation which allows satisfactory control. Different choices of \mathbf{h} will be explored later. The control cost $r(\mathbf{u}, \mathbf{x})$ cannot be represented exactly on the high level, because it typically contains independent contributions from all components of \mathbf{u} . This is the reason why the proposed hierarchical scheme is only an approximation to optimal control.

Differentiating Eq. (3) w.r.t. t and using Eq. (1), the dynamics of \mathbf{y} become

$$\dot{\mathbf{y}} = H(\mathbf{x})(\mathbf{a}(\mathbf{x}) + B(\mathbf{x})\mathbf{u}), \quad (5)$$

where $H(\mathbf{x}) = \partial \mathbf{h}(\mathbf{x}) / \partial \mathbf{x}$ is the Jacobian of the function \mathbf{h} . Note that \mathbf{y} will often contain more information than end-effector position.

Our design method seeks to create an augmented system with prescribed high-level dynamics

$$\dot{\mathbf{y}} = \mathbf{f}(\mathbf{y}) + G(\mathbf{y})\mathbf{v}, \quad (6)$$

where $\mathbf{v} \in \mathbb{R}^{n_v}$ is the high-level control signal generated by some high-level controller. The functions \mathbf{f}, G define the meaning of the high-level control \mathbf{v} , as discussed later. For now, we simply assume that \mathbf{f}, G are some known functions.

The low-level controller design problem is the following: choose $\mathbf{u}(\mathbf{v}, \mathbf{x})$ so that the prescribed [Eq. (6)] and actual [Eq. (5)] \mathbf{y} dynamics are identical; when multiple solutions exist, use the control cost $r(\mathbf{u}, \mathbf{x})$ to resolve redundancy. The control \mathbf{u} is thus defined at each time t as the solution to the following constrained optimization problem:

given \mathbf{v} and \mathbf{x} , find \mathbf{u} that minimizes
 $r(\mathbf{u}, \mathbf{x})$ subject to

$$H(\mathbf{x})B(\mathbf{x})\mathbf{u} = \mathbf{f}(\mathbf{y}) + G(\mathbf{y})\mathbf{v} - H(\mathbf{x})\mathbf{a}(\mathbf{x}). \quad (7)$$

In addition to the above equality constraint, we can incorporate inequality constraints on \mathbf{u} . The latter are particularly important in biological movement control where muscle activations are always non-negative. For a general $r(\mathbf{u}, \mathbf{x})$, problem (7) has to be solved numerically at each time step (using Sequential Quadratic Programming, or another efficient method). Since the solution is likely to be similar from one time step to the next, we can use the current solution to initialize the search for the next solution.

In practice, the control cost will normally have the quadratic form

$$r(\mathbf{u}, \mathbf{x}) = \frac{1}{2}\mathbf{u}^T R(\mathbf{x})\mathbf{u}, \quad (8)$$

where $R(\mathbf{x})$ is a symmetric positive-definite matrix (often proportional to the identity matrix). When $r(\mathbf{u}, \mathbf{x})$ is in the form (8), and control inequality constraints are absent, we can solve Eq. (7) explicitly. The solution relies on the following fact:

$$\begin{aligned} &\text{the minimum of } \frac{1}{2}\mathbf{u}^T W\mathbf{u} \text{ subject to } X\mathbf{u} = \mathbf{c} \text{ is} \\ &\mathbf{u} = X_W^\dagger \mathbf{c}. \end{aligned} \quad (9)$$

The weighted pseudo-inverse is defined as

$$X_W^\dagger = W^{-1}X^T(XW^{-1}X^T)^{-1}, \quad (10)$$

where W is symmetric positive definite, and X has full row rank. Note that X_1^\dagger is the Moore–Penrose pseudo-inverse $X^T(XX^T)^{-1}$; in that case we use the shorthand X^\dagger .

Applying this to problem (7) for a cost in the form (8), and assuming that $H(\mathbf{x})B(\mathbf{x})$ has full row-rank (a condition which we will return to), the unique solution is

$$\mathbf{u}(\mathbf{v}, \mathbf{x}) = (H(\mathbf{x})B(\mathbf{x}))_{R(\mathbf{x})}^\dagger (\mathbf{f}(\mathbf{y}) + G(\mathbf{y})\mathbf{v} - H(\mathbf{x})\mathbf{a}(\mathbf{x})). \quad (11)$$

In case of inequality constraints on \mathbf{u} , we cannot obtain an explicit solution, and have to resort to numerical optimization. But the quadratic form (8) of the control cost allows us to use Quadratic Programming—which is numerically very efficient when the cost is convex.

Another form of cost function, that can be used on the low level to resolve redundancy, is

$$r(\mathbf{u}, \mathbf{x}) = \frac{1}{2}\mathbf{u}^T R(\mathbf{x})\mathbf{u} + (\mathbf{a}(\mathbf{x}) + B(\mathbf{x})\mathbf{u})^T \nabla_{\mathbf{x}} g(\mathbf{x}). \quad (12)$$

The new term encourages movement against the gradient of some potential function $g(\mathbf{x})$. With this $r(\mathbf{u}, \mathbf{x})$ problem (7) can still be solved explicitly, using the following fact:

the minimum of $\frac{1}{2}\mathbf{u}^T W\mathbf{u} + \mathbf{u}^T \mathbf{d}$ subject to $X\mathbf{u} = \mathbf{c}$ is

$$\mathbf{u} = \begin{bmatrix} X \\ N_X^T W \end{bmatrix}^{-1} \begin{bmatrix} \mathbf{c} \\ -N_X^T \mathbf{d} \end{bmatrix} = X_W^\dagger \mathbf{c} - N_X (N_X^T W N_X)^{-1} N_X^T \mathbf{d}, \quad (13)$$

where N_X is a matrix whose columns form an orthonormal basis for the null space of X . Note that solving the system $[X; N_X^T W]\mathbf{u} = [\mathbf{c}; -N_X^T \mathbf{d}]$ with a linear solver is numerically preferable to explicit computation of the pseudo-inverses in Eq. (13). A similar point has been made in Ref. 35 regarding velocity control methods.

Applying Eq. (13) to our problem in the case when control inequality constraints are absent, and suppressing the functional dependencies for clarity, we have

$$\begin{aligned} \mathbf{u} = &(HB)_R^\dagger (\mathbf{f} + G\mathbf{v} - H\mathbf{a}) \\ &- N_{HB} (N_{HB}^T R N_{HB})^{-1} N_{HB}^T \nabla_{\mathbf{x}} g. \end{aligned} \quad (14)$$

In comparison to Eq. (11), the controller in Eq. (14) has an extra term acting in the redundant subspace where variations in \mathbf{x} do not affect \mathbf{y} . Thus $g(\mathbf{x})$ can be used to modify the behavior of the controlled system in that subspace; the choice of $g(\mathbf{x})$ is presently left to the user.

An alternative way of computing $\mathbf{u}(\mathbf{v}, \mathbf{x})$ is to transform the above constrained optimization problem into an unconstrained one, and use a mixed cost that absorbs all constraints. Let $c(\mathbf{u}, \mathbf{x})$ be a potential function replacing any control inequality constraints: $c(\mathbf{u}, \mathbf{x})=0$ when \mathbf{u} satisfies those constraints and $c(\mathbf{u}, \mathbf{x})>0$ when it does not. Then we can define \mathbf{u} at each time t as the solution to the following unconstrained optimization problem:

given \mathbf{v} and \mathbf{x} , find \mathbf{u} that minimizes

$$\lambda_1 \|\mathbf{f}(\mathbf{y}) + G(\mathbf{y})\mathbf{v} - H(\mathbf{x})\mathbf{a}(\mathbf{x}) - H(\mathbf{x})B(\mathbf{x})\mathbf{u}\|^2 + \lambda_2 c(\mathbf{u}, \mathbf{x}) + r(\mathbf{u}, \mathbf{x}). \quad (15)$$

The constants λ_1, λ_2 set the relative importance of satisfying constraints versus minimizing energy. For large λ 's the solution to this problem is numerically indistinguishable from the solution to the constrained problem, in the case when all constraints can be satisfied. However, the advantage of Eq. (15) is that it applies even when the constraints on \mathbf{u} cannot be satisfied—which can happen when $H(\mathbf{x})B(\mathbf{x})$ becomes row-rank deficient, or when the high-level control \mathbf{v} calls for a low-level control \mathbf{u} exceeding its allowed range.

3.2. Dynamic Compatibility Between Levels of Control

A potential problem with the above method is lack of “dynamic compatibility” between the two levels. This occurs when $\dot{\mathbf{y}}$ does not explicitly depend on \mathbf{u} , and leads to $H(\mathbf{x})B(\mathbf{x})=0$. For example, suppose \mathbf{y} is end-effector position, \mathbf{v} is end-effector velocity, and \mathbf{u} is joint torque. This is problematic because torque cannot affect velocity instantaneously. However torque has a predictable effect when applied over time, suggesting that the “instantaneous” \mathbf{a}, B should be replaced with functions $\tilde{\mathbf{a}}, \tilde{B}$ incorporating temporal prediction. Next we present a method for doing so.

Consider the discrete-time representation of system (1) with time step Δ :

$$\mathbf{x}(t + \Delta) = \mathbf{x}(t) + \Delta(\mathbf{a}(\mathbf{x}(t)) + B(\mathbf{x}(t))\mathbf{u}(t)). \quad (16)$$

We now analyze this system within a single step, from time t to time $t + \Delta$. Define the shortcut notation $\mathbf{x}_\tau = \mathbf{x}(t + \tau)$, $\tau \in [0, \Delta]$, so that $\mathbf{x}_0 = \mathbf{x}(t)$. The control

term $\mathbf{b} = B(\mathbf{x}(t))\mathbf{u}(t)$ will be treated as constant, while the changes in $\mathbf{a}(\mathbf{x})$ will be captured to first order through the linearization

$$\mathbf{a}(\mathbf{x}_\tau) \approx \mathbf{a}(\mathbf{x}_0) + A(\mathbf{x}_\tau - \mathbf{x}_0), \quad A = \left. \frac{\partial \mathbf{a}(\mathbf{x})}{\partial \mathbf{x}} \right|_{\mathbf{x} = \mathbf{x}_0}. \quad (17)$$

Then, we have continuous dynamics

$$\dot{\mathbf{x}}_\tau = A\mathbf{x}_\tau + (\mathbf{a}(\mathbf{x}_0) + \mathbf{b} - A\mathbf{x}_0), \quad (18)$$

where the term in the brackets is constant. In the case when $x, a, b, A = a'(x_0)$ are scalar, we can solve this initial-value ODE analytically, and after some rearrangement obtain

$$x_\Delta = x_0 + \frac{\exp(\Delta a'(x_0)) - 1}{a'(x_0)}(a(x_0) + b). \quad (19)$$

This scalar result indicates that in the vector case we can expect a similar solution involving a matrix exponential and an inverse, and we can also expect complications when A is singular.

We will now solve Eq. (18) in the vector case, by performing Euler integration in n smaller time steps of duration $d = \Delta/n$ each, and then taking the limit $n \rightarrow \infty$. Each step $k < n$ is in the form

$$\mathbf{x}_{d(k+1)} = (\mathbf{I} + dA)\mathbf{x}_{dk} + d(\mathbf{a}(\mathbf{x}_0) - A\mathbf{x}_0 + \mathbf{b}). \quad (20)$$

Applying this recursion n times, and noting that $\mathbf{x}_{dn} = \mathbf{x}_\Delta$, we have

$$\begin{aligned} \mathbf{x}_\Delta &= (\mathbf{I} + dA)^n \mathbf{x}_0 + (\mathbf{I} + (\mathbf{I} + dA) + \dots \\ &\quad + (\mathbf{I} + dA)^{n-1}) d(\mathbf{a}(\mathbf{x}_0) - A\mathbf{x}_0 + \mathbf{b}). \end{aligned} \quad (21)$$

Using the identity $(\mathbf{I} + X + \dots + X^{n-1})(X - \mathbf{I}) = (X^n - \mathbf{I})$ along with $dA\mathbf{x}_0 = ((\mathbf{I} + dA) - \mathbf{I})\mathbf{x}_0$, the above equality simplifies to

$$\mathbf{x}_\Delta = \mathbf{x}_0 + \Delta T(\Delta A, n)(\mathbf{a}(\mathbf{x}_0) + \mathbf{b}), \quad (22)$$

where the matrix function T is defined as

$$T(X,n) = \frac{1}{n} \left(\mathbf{I} + \left(\mathbf{I} + \frac{X}{n} \right) + \cdots + \left(\mathbf{I} + \frac{X}{n} \right)^{n-1} \right). \quad (23)$$

When X is invertible, T further simplifies to

$$T(X,n) = \left(\left(\mathbf{I} + \frac{X}{n} \right)^n - \mathbf{I} \right) X^{-1}. \quad (24)$$

In the limit $n \rightarrow \infty$, the identity $\lim_{n \rightarrow \infty} (\mathbf{I} + X/n)^n = \exp(X)$ can be used to obtain

$$T(X,\infty) = (\exp(X) - \mathbf{I})X^{-1}. \quad (25)$$

When X is singular we can approximate $T(X,\infty)$ by either adding $\epsilon \mathbf{I}$ to X and applying Eq. (25), or by using $T(X,n)$ evaluated with a large n . Numerically, we have found that $T(X,n)$ converges even for singular X , although we do not have a closed-form solution in that case.

Switching back to continuous notation, we have the following modified integration scheme:

$$\mathbf{x}(t + \Delta) = \mathbf{x}(t) + \Delta T(\Delta A, \infty) (\mathbf{a}(\mathbf{x}(t)) + B(\mathbf{x}(t))\mathbf{u}(t)), \quad (26)$$

which differs from Euler integration (16) only by T . This motivates the construction of the “predictive” dynamics $\tilde{\mathbf{a}}, \tilde{B}$ from the following condition: Euler integration (16) with $\tilde{\mathbf{a}}, \tilde{B}$ should be equivalent to the modified scheme (26) with the original \mathbf{a}, B . This condition is satisfied by

$$\begin{aligned} \tilde{\mathbf{a}}(\mathbf{x}) &= T(\Delta A, \infty) \mathbf{a}(\mathbf{x}), \\ \tilde{B}(\mathbf{x}) &= T(\Delta A, \infty) B(\mathbf{x}). \end{aligned} \quad (27)$$

Explicit integration of the functions $\tilde{\mathbf{a}}, \tilde{B}$ resembles implicit integration of the original \mathbf{a}, B . To see the relationship to implicit integration, consider the implicit Euler scheme

$$\mathbf{x}_\Delta = \mathbf{x}_0 + \Delta (\mathbf{a}(\mathbf{x}_\Delta) + \mathbf{b}), \quad (28)$$

where \mathbf{x}_Δ is now the unknown in a nonlinear equation. One way to approximate the solution is to linearize \mathbf{a} around \mathbf{x}_0 as we did before, yielding

$$\mathbf{x}_\Delta = \mathbf{x}_0 + \Delta (\mathbf{a}(\mathbf{x}_0) + A\mathbf{x}_\Delta - A\mathbf{x}_0 + \mathbf{b}). \quad (29)$$

Moving $\Delta A\mathbf{x}_\Delta$ to the left hand side and left multiplying by $(\mathbf{I} - \Delta A)^{-1}$, the solution is

$$\mathbf{x}_\Delta = \mathbf{x}_0 + \Delta \bar{T}(\Delta A) (\mathbf{a}(\mathbf{x}_0) + \mathbf{b}), \quad (30)$$

where the matrix function \bar{T} is defined as

$$\bar{T}(X) = (\mathbf{I} - X)^{-1}. \quad (31)$$

This is very similar to Eq. (22), except that \bar{T} now replaces T . However Eq. (25) is preferable to Eq. (31), because it captures the exact solution of the linearized ODE (18) while the implicit Euler method relies on a finite-difference approximation. Note that implicit integration is in general more stable and accurate, so the functions $\tilde{\mathbf{a}}, \tilde{B}$ can be used instead of \mathbf{a}, B even when dynamic compatibility is not an issue. Their only drawback is the increased computational cost per step—which may be offset by the larger time steps that implicit integrators can safely take.

Here is a simple example of how Eq. (27) can enforce dynamic compatibility. Suppose p and v are the position and velocity of a one-dimensional point with mass 1, and u is the applied force, so that $\dot{v} = u$. The plant state is $\mathbf{x} = [p; v]$. Let the high-level state be $y = p$. Then, the dynamics functions are $\mathbf{a}(\mathbf{x}) = [v; 0]$ and $B = [0; 1]$. Therefore, $H = \partial y / \partial \mathbf{x} = [1, 0]$, and so $HB = 0$. To apply Eq. (27), we first compute $A = \partial \mathbf{a} / \partial \mathbf{x} = [0, 1; 0, 0]$. Numerically $T(\Delta A, \infty)$ converges to the matrix $[1, \Delta/2; 0, 1]$. Thus the modified dynamics are $\tilde{\mathbf{a}}(\mathbf{x}) = \mathbf{a}(\mathbf{x})$ and $\tilde{B} = [\Delta/2; 1]$, yielding $H\tilde{B} = \Delta/2$. The fact that $H\tilde{B} \neq 0$ makes it possible to design the low-level controller using Eq. (11).

3.3. High-Level Dynamics and Cost Models

It is clear from Eq. (11) that as long as $H(\mathbf{x})B(\mathbf{x})$ has full row-rank, we can choose \mathbf{f}, G arbitrarily and thereby instantiate whatever high-level dynamics we desire. But what should we desire? One possibility is to choose $\mathbf{f}(\mathbf{y})$ linear and $G(\mathbf{y})$ constant—which yields linear \mathbf{y} dynamics, and is related to feedback linearization as discussed below. However we believe that the \mathbf{y} dynamics should mimic the \mathbf{x} dynamics to the extent possible, so that the high-level controller can exploit the natural plant dynamics

while operating on a lower-dimensional system. Matching the passive dynamics in Eqs. (5) and (6) yields

$$\mathbf{f}(\mathbf{h}(\mathbf{x})) \approx H(\mathbf{x})\mathbf{a}(\mathbf{x}). \quad (32)$$

Matching the control-dependent terms in Eqs. (5) and (6) yields

$$G(\mathbf{h}(\mathbf{x}))\tilde{\mathbf{v}}(\mathbf{u}, \mathbf{x}) \approx H(\mathbf{x})B(\mathbf{x})\mathbf{u}, \quad (33)$$

where the function $\tilde{\mathbf{v}}(\mathbf{u}, \mathbf{x})$ remains to be defined. In addition, optimization of the high-level controller will be facilitated if the cost function can be captured on the high level. The state-dependent cost term $q(t, \mathbf{x})$ is by definition computable given $\mathbf{y} = \mathbf{h}(\mathbf{x})$, but we would also like (an approximation to) the control energy term $r(\mathbf{u}, \mathbf{x})$. This requires a function \tilde{r} such that

$$\tilde{r}(\tilde{\mathbf{v}}(\mathbf{u}, \mathbf{x}), \mathbf{h}(\mathbf{x})) \approx r(\mathbf{u}, \mathbf{x}). \quad (34)$$

Functions \mathbf{f}, G, \tilde{r} that approximately satisfy Eqs. (32)–(34) can be constructed by choosing a suitable parameterization (basis functions, neural networks, etc.), collecting data, and fitting a model using supervised learning procedures. The dataset consists of observation pairs $(\mathbf{x}^{(k)}, \mathbf{u}^{(k)})$, where $\mathbf{u}^{(k)}$ is the signal generated by some existing controller when the plant is in state $\mathbf{x}^{(k)}$. Given $\mathbf{x}^{(k)}$ and $\mathbf{u}^{(k)}$, we can compute learning targets for \mathbf{f}, G, \tilde{r} by evaluating the right hand sides of Eqs. (32)–(34). Thus the application of supervised learning is straightforward—as long as we define the function $\tilde{\mathbf{v}}(\mathbf{u}, \mathbf{x})$. This function corresponds to the meaning of the high-level control signals, and should be chosen using physical intuition. For example, if \mathbf{u} is joint torque and \mathbf{y} is end-effector position and velocity, $\tilde{\mathbf{v}}(\mathbf{u}, \mathbf{x})$ could be either end-effector acceleration or force. In the former case G will be constant and will not capture anything about plant dynamics. In the latter case, G will be related to inverse inertia in end-effector space. Regardless of how we choose the function $\tilde{\mathbf{v}}$, after learning we will have $\tilde{\mathbf{v}}(\mathbf{u}(\mathbf{v}, \mathbf{x}), \mathbf{x}) \approx \mathbf{v}$. Note that whenever the functions \mathbf{f}, G change (through learning or otherwise) the low-level controller must be redesigned.

It is interesting to ask under what conditions Eqs. (32)–(34) can be satisfied exactly. We focus on the passive dynamics (32), which do not depend on our choice of $\tilde{\mathbf{v}}(\mathbf{u}, \mathbf{x})$ and are also likely to be the most worthwhile fitting. Since $\mathbf{x} \rightarrow \mathbf{y}$ is a many-to-

one mapping, a function \mathbf{f} satisfying Eq. (32) does not always exist. The necessary and sufficient condition for the existence of \mathbf{f} is

$$H(\mathbf{x}_1)\mathbf{a}(\mathbf{x}_1) = H(\mathbf{x}_2)\mathbf{a}(\mathbf{x}_2) \text{ whenever } \mathbf{h}(\mathbf{x}_1) = \mathbf{h}(\mathbf{x}_2). \quad (35)$$

Define the set $\mathcal{M}[\mathbf{x}_0] = \{\mathbf{x} : \mathbf{h}(\mathbf{x}) = \mathbf{h}(\mathbf{x}_0)\}$ and suppose \mathbf{h} is such that \mathcal{M} is a smooth manifold. This is the redundant (or uncontrolled) manifold where the plant state \mathbf{x} can fluctuate without affecting the high-level state \mathbf{y} . If the passive dynamics act within this manifold, i.e., $\mathbf{a}(\mathbf{x}) \in T_{\mathbf{x}}\mathcal{M}[\mathbf{x}]$, then $H(\mathbf{x})\mathbf{a}(\mathbf{x}) = 0$ and Eq. (35) is trivially satisfied. Therefore, Eq. (35) can only be violated by the component of the vector field \mathbf{a} normal to \mathcal{M} . In particular, note that the row vectors of the matrix $H(\mathbf{x})$ span the subspace normal to $\mathcal{M}[\mathbf{x}]$, and so $H(\mathbf{x})\mathbf{a}(\mathbf{x})$ is a form of projection of the vector field in the normal subspace. For nonlinear $\mathbf{h}(\mathbf{x})$ this projection is difficult to understand because $H(\mathbf{x})$ varies with \mathbf{x} , but in the linear case $\mathbf{h}(\mathbf{x}) = H\mathbf{x}$ condition (35) becomes

$$H(\mathbf{a}(\mathbf{x}_1) - \mathbf{a}(\mathbf{x}_2)) = 0 \text{ whenever } H(\mathbf{x}_1 - \mathbf{x}_2) = 0. \quad (36)$$

The latter condition has a simple interpretation: if the difference between two states lies within the redundant subspace, then the difference in the passive dynamics between these states should also lie within the redundant subspace.

3.4. High-Level Controller Optimization

A high-level dynamics model together with an (approximate) high-level cost model lead to an optimal control problem formulated on the high level. This problem can be solved via standard techniques, or via a new implicit method (see below) that is specific to hierarchical control schemes.

While a survey of standard techniques for optimal controller design is beyond the scope of this paper, three points are in order. First, mapping the original control problem to the high level reduces dimensionality, and thereby facilitates the application of numerical methods for optimal control that suffer from the curse of dimensionality. Second, a high-level dynamics model is required in order to design the low-level controller, even if the high-level controller is built via model-free methods (such as proportional-derivative control, or reinforcement

learning). Third, it may be advantageous to use iterative design: Start with some high-level model, design the two-level control scheme and apply it, use the resulting data to fit a high-level model more compatible with the plant dynamics, and iterate. The need to iterate comes from the fact that the approximation [Eqs. (32)–(34)] depends on what regions of the state and control spaces are visited, and that in turn depends on what controller is being used during data collection.

In addition to standard techniques, the hierarchical framework allows a novel implicit method for optimal controller design. In this method, the passive dynamics are no longer modeled explicitly on the high level and the discrepancy between $\mathbf{f}(\mathbf{y})$ and $H(\mathbf{x})\mathbf{a}(\mathbf{x})$ is no longer compensated on the low level. Instead, the high-level controller is given online access to $H(\mathbf{x})\mathbf{a}(\mathbf{x})$ and is responsible for dealing with that term. The constraint in Eq. (7) simplifies to $H(\mathbf{x})B(\mathbf{x})\mathbf{u}=G(\mathbf{y})\mathbf{v}$, and the low-level controller given by Eq. (11) now becomes

$$\mathbf{u}(\mathbf{v}, \mathbf{x}) = (H(\mathbf{x})B(\mathbf{x}))_{R(\mathbf{x})}^{\dagger} G(\mathbf{y})\mathbf{v}. \quad (37)$$

When condition (35) is violated the input-output behavior of the augmented plant resulting from Eq. (37) is no longer autonomous in terms of \mathbf{y} , but instead can depend on the specific value of \mathbf{x} —because $H(\mathbf{x})\mathbf{a}(\mathbf{x})$ can vary while \mathbf{y} remains constant. In that case $\mathbf{f}(\mathbf{y})$ can still be computed as long as the plant state \mathbf{x} is known. But how can we know \mathbf{x} when $\mathbf{x} \rightarrow \mathbf{y}$ is a many-to-one mapping? We cannot if all we are given is an arbitrarily chosen \mathbf{y} . However, suppose that we reached \mathbf{y} by initializing the plant in some state $\mathbf{x}(0)$, applying some high-level control sequence $\mathbf{v}(t)$, computing the low-level controls from Eq. (37), and integrating the plant dynamics. Then we always have a unique underlying state \mathbf{x} . Implicit computation of $\mathbf{f}(\mathbf{y})$ can be used in conjunction with trajectory-based methods for optimal control—such as ODE methods,³⁶ differential dynamic programming,³⁷ or the iterative linear-quadratic-Gaussian method³⁸ we have developed (illustrated later). This modified version of our hierarchical scheme is likely to yield better performance, but comes at the price of not having an explicit high-level controller $\mathbf{v}(\mathbf{y})$ that is independent of the initial \mathbf{x} and its trajectory. When condition (35) is satisfied, and an exact passive dynamics model $\mathbf{f}(\mathbf{y}) = H(\mathbf{x})\mathbf{a}(\mathbf{x})$ is available, the implicit method does not improve performance.

4. RELATION TO EXISTING APPROACHES

Hierarchical control is an appealing idea which has been pursued in a number of fields. Consequently, several existing approaches are related to ours. Some are related only in spirit (and are summarized first), while others allow formal comparisons developed in the subsections below.

In reinforcement learning, researchers have attempted to alleviate the curse of dimensionality via autonomous low-level control policies^{8,9} that solve specific subproblems over time (such as exiting a room in a maze). In contrast our low-level controller performs an instantaneous feedback transformation of plant dynamics, and is continuously driven by high-level commands. We believe this is more appropriate for control of complex redundant manipulators (and is a more plausible model of biological sensorimotor control) while hierarchical reinforcement learning is more appropriate for non-articulated “agents” solving navigation problems.

In control theory, there is increasing interest in hybrid systems^{39–41} with continuous low-level dynamics and discrete event-driven high-level dynamics. In contrast, the high-level dynamics we instantiate are continuous. Consider for example a walking mechanism. Our framework can be applied to yield augmented dynamics in terms of the center-of-mass and feet positions. The resulting high level could be further decomposed via a hybrid systems approach, with ground contacts being the discrete events. Thus, the two approaches are in some cases complementary, and a closer examination of the link between them is needed.

In robotics, examples of hierarchical control include the subsumption architecture⁵ as well as virtual model control.⁴² The latter is similar to our approach in terms of overall strategy—which is to design high-level feedback controllers without considering the full details of the plant. However virtual model control is a kinematic method: It maps forces produced by virtual actuators into real actuator commands, and does not exploit knowledge of plant dynamics. Another difference is that the high-level controller design is intuitive rather than model based, and furthermore a model of high-level dynamics independent of the specific high-level controller is not available. Finally, virtual model control takes dynamic compatibility for granted and does not readily apply to biomechanical systems with higher-order dynamics.

4.1. Feedback Linearization

There are interesting parallels between our method and feedback linearization (FL)—which is one of the few general methods for nonlinear control.^{21,22} For comparison purposes, we first give a brief summary of input-output FL in the single-input-single-output case. Consider the system

$$\begin{aligned}\dot{\mathbf{x}} &= \mathbf{a}(\mathbf{x}) + \mathbf{b}(\mathbf{x})u, \\ y &= h(\mathbf{x}),\end{aligned}\quad (38)$$

where the control u is now called “input”, the high-level state y is “output”, the high-level control v is “external reference input”, and u, y, v are for simplicity scalar. Define the Lie derivative of a function $h(\mathbf{x})$ with respect to a vector field $\mathbf{a}(\mathbf{x})$ as the directional derivative of h along \mathbf{a} :

$$L_{\mathbf{a}}h = \frac{\partial h(\mathbf{x})}{\partial \mathbf{x}} \mathbf{a}(\mathbf{x}). \quad (39)$$

System (38) is said to have relative degree r in a region \mathcal{D} if for all $\mathbf{x} \in \mathcal{D}$, we have

$$\begin{aligned}L_{\mathbf{b}}L_{\mathbf{a}}^k h(\mathbf{x}) &= 0 \text{ for all } k < r - 1, \\ L_{\mathbf{b}}L_{\mathbf{a}}^{r-1} h(\mathbf{x}) &\neq 0.\end{aligned}\quad (40)$$

Then the low-level control law (called “state feedback control”) given by

$$u(v, \mathbf{x}) = \frac{1}{L_{\mathbf{b}}L_{\mathbf{a}}^{r-1}h(\mathbf{x})}(v - L_{\mathbf{a}}^r h(\mathbf{x})) \quad (41)$$

yields linear y dynamics of the form $y^{(r)} = v$, where $y^{(r)}$ is the r th-order time derivative of y .

To see the similarity to our method, note that definition (39) implies $L_{\mathbf{a}}h(\mathbf{x}) = H(\mathbf{x})\mathbf{a}(\mathbf{x})$ and similarly for $L_{\mathbf{b}}$. Now, in the case $r=1$ we have $L_{\mathbf{a}}^{r-1}h = h$, and the control law (41) becomes

$$u(v, \mathbf{x}) = \frac{1}{H(\mathbf{x})\mathbf{b}(\mathbf{x})}(v - H(\mathbf{x})\mathbf{a}(\mathbf{x})) \quad (42)$$

which is identical to our control law (11) for prescribed high-level dynamics $f(y)=0$, $G(y)=1$ and

control cost $u^2/2$. The parallels between the two methods can be summarized as follows:

- The objective of FL is to make the \mathbf{y} -dynamics linear, in the form $\mathbf{y}^{(r)} = \mathbf{v}$, while our objective is to make it low-dimensional yet similar to the \mathbf{x} dynamics. Our method can instantiate any desired \mathbf{y} dynamics.
- The dynamic compatibility problem discussed above corresponds to a system with relative degree greater than one. FL handles this case by augmenting \mathbf{y} with its first $r-1$ derivatives. We could do the same, but that may defeat the purpose of our method because dimensionality is increased. Instead, we prefer to use the predictive version $\tilde{\mathbf{a}}, \tilde{\mathbf{B}}$ of plant dynamics.
- Most of FL theory is developed for systems with equal numbers of inputs and outputs, while redundant manipulators tend to have many more inputs than task-relevant outputs. Our method does not impose any such restriction. We deal with the mismatch in dimensionality by formulating an optimization problem—which yields a weighted pseudo-inverse solution.
- There is an important step in FL which was omitted in the above summary. It has to do with finding an explicit parameterization of the redundant manifold \mathcal{M} defined above. Such a parameterization is analytically interesting, but does not facilitate the design of a high-level controller (and the low-level controller has already been designed).

4.2. Operational Space Formulation

Our framework also turns out to have a number of interesting similarities and differences with the operational space (OS) formulation⁶—which is another hierarchical scheme aiming to decouple task-level control from details of plant dynamics. For comparison purposes, we summarize the OS scheme here. Let \mathbf{q} be the generalized coordinates of a redundant manipulator. The end-effector coordinates \mathbf{p} are related to \mathbf{q} by the forward kinematics function $\mathbf{p} = \mathbf{k}(\mathbf{q})$, with Jacobian $J(\mathbf{q}) = \partial \mathbf{k}(\mathbf{q}) / \partial \mathbf{q}$. The equations of motion are

$$\boldsymbol{\tau} = M(\mathbf{q})\ddot{\mathbf{q}} + \mathbf{n}(\mathbf{q}, \dot{\mathbf{q}}), \quad (43)$$

where $M(\mathbf{q})$ is the (symmetric positive-definite) generalized inertia matrix, and \mathbf{n} are Coriolis, centrip-

etal, gravitational, and viscoelastic forces. The end-effector inertia matrix is

$$\Lambda(\mathbf{q}) = (J(\mathbf{q})M(\mathbf{q})^{-1}J(\mathbf{q})^T)^{-1}. \quad (44)$$

The OS low-level controller maps desired end-effector accelerations $\ddot{\mathbf{p}}^*$ to generalized forces τ as

$$\tau = J(\mathbf{q})^T \Lambda(\mathbf{q})(\ddot{\mathbf{p}}^* - \dot{J}(\mathbf{q})\dot{\mathbf{q}}) + \mathbf{n}(\mathbf{q}, \dot{\mathbf{q}}). \quad (45)$$

The OS scheme leaves the high-level controller unspecified—as long as that controller can issue desired end-effector accelerations. The high-level dynamics correspond to a point mass ($\ddot{\mathbf{p}} = \ddot{\mathbf{p}}^*$), therefore, the operational space formulation is an instance of feedback linearization.

Let us now apply our method to the same plant. Note that our method is formulated using general first-order dynamics, and so the plant state contains both generalized coordinates and their derivatives. We have $\mathbf{x} = [\mathbf{q}; \dot{\mathbf{q}}]$, $\mathbf{y} = [\mathbf{p}; \dot{\mathbf{p}}]$, and, therefore,

$$\mathbf{h}(\mathbf{x}) = \begin{bmatrix} \mathbf{k}(\mathbf{q}) \\ J(\mathbf{q})\dot{\mathbf{q}} \end{bmatrix}. \quad (46)$$

The second row in the above equation comes from the identity

$$\dot{\mathbf{p}} = J(\mathbf{q})\dot{\mathbf{q}}. \quad (47)$$

The Jacobian of \mathbf{h} differs from the Jacobian of the forward kinematics function, and is given by

$$H(\mathbf{x}) = \begin{bmatrix} J(\mathbf{q}) & 0 \\ \frac{\partial}{\partial \mathbf{q}}(J(\mathbf{q})\dot{\mathbf{q}}) & J(\mathbf{q}) \end{bmatrix}, \quad (48)$$

where the above partial derivative term can be written as $\dot{J}(\mathbf{q})$. This is because

$$\begin{aligned} \frac{\partial}{\partial \mathbf{q}}(J(\mathbf{q})\dot{\mathbf{q}}) &= \frac{\partial}{\partial \mathbf{q}} \left(\frac{\partial \mathbf{k}}{\partial \mathbf{q}} \frac{d\mathbf{q}}{dt} \right) = \frac{\partial}{\partial \mathbf{q}} \left(\frac{d\mathbf{k}}{dt} \right) \\ &= \frac{d}{dt} \left(\frac{\partial \mathbf{k}}{\partial \mathbf{q}} \right) = \frac{d}{dt}(J(\mathbf{q})) = \dot{J}(\mathbf{q}). \end{aligned} \quad (49)$$

From Eq. (43), the passive and control-dependent plant dynamics expressed in first-order form are

$$\mathbf{a}(\mathbf{x}) = \begin{bmatrix} \dot{\mathbf{q}} \\ -M(\mathbf{q})^{-1}\mathbf{n}(\mathbf{q}, \dot{\mathbf{q}}) \end{bmatrix}, \quad B(\mathbf{x}) = \begin{bmatrix} 0 \\ M(\mathbf{q})^{-1} \end{bmatrix}. \quad (50)$$

Suppose the prescribed high-level dynamics correspond to a point mass ($\ddot{\mathbf{p}} = \mathbf{v}$) so that

$$\mathbf{f}(\mathbf{y}) = \begin{bmatrix} \dot{\mathbf{p}} \\ 0 \end{bmatrix}, \quad G(\mathbf{y}) = \begin{bmatrix} 0 \\ \mathbf{I} \end{bmatrix}. \quad (51)$$

Let $R(\mathbf{x}) = \mathbf{I}$, so the control cost is $\mathbf{u}^T \mathbf{u} / 2$. Then the constraint in Eq. (7) becomes

$$\begin{bmatrix} 0 \\ J(\mathbf{q})M(\mathbf{q})^{-1} \end{bmatrix} \mathbf{u} = \begin{bmatrix} \dot{\mathbf{p}} \\ \mathbf{v} \end{bmatrix} - \begin{bmatrix} J(\mathbf{q})\dot{\mathbf{q}} \\ J(\mathbf{q})\dot{\mathbf{q}} - J(\mathbf{q})M(\mathbf{q})^{-1}\mathbf{n}(\mathbf{q}, \dot{\mathbf{q}}) \end{bmatrix}. \quad (52)$$

The first row is always satisfied because of Eq. (47). Thus, the effective constraint on \mathbf{u} is given by the second row of the above equation. Replacing the general \mathbf{u}, \mathbf{v} with their specific to this example values $\tau, \ddot{\mathbf{p}}^*$, the solution produced by our method is

$$\tau = (J(\mathbf{q})M(\mathbf{q})^{-1})^\dagger (\ddot{\mathbf{p}}^* - \dot{J}(\mathbf{q})\dot{\mathbf{q}} + J(\mathbf{q})M(\mathbf{q})^{-1}\mathbf{n}(\mathbf{q}, \dot{\mathbf{q}})). \quad (53)$$

To make the relationship to the operational space formulation more explicit, we omit the dependence on \mathbf{q} for clarity, and observe that the definition of the weighted pseudo-inverse implies $(JM^{-1})^\dagger = MJ_{M_2}^\dagger$ and $J^T \Lambda = MJ_M^\dagger$. The two methods can then be rewritten as

$$\text{operational space: } \tau = MJ_M^\dagger (\ddot{\mathbf{p}}^* - \dot{J}\dot{\mathbf{q}}) + \mathbf{n},$$

$$\text{hierarchical control: } \tau = MJ_{M_2}^\dagger (\ddot{\mathbf{p}}^* - \dot{J}\dot{\mathbf{q}}) + (JM^{-1})^\dagger (JM^{-1})\mathbf{n}. \quad (54)$$

The overall form is very similar, but there are two differences. First, the torques $\mathbf{n}(\mathbf{q}, \dot{\mathbf{q}})$ are fully cancelled in the OS method, while our method only cancels the component acting in end-effector space. Second, the Jacobian pseudo-inverses being used are weighted differently (see next section). Note that when J is in-

vertible, both formulas reduce to the inverse dynamics (43), and map desired accelerations to forces. To see this, differentiate (47) to obtain $J(\mathbf{q})\ddot{\mathbf{q}} = \ddot{\mathbf{p}} - \dot{J}(\mathbf{q})\dot{\mathbf{q}}$.

Despite the similarity, our framework is more general than the operational space formulation—in two ways that are obscured by this example. First, it can handle systems other than second order, without any modification. Indeed, we will later illustrate the control of a third-order system (a model of the human arm). Second, our framework is not constrained to point-mass high-level dynamics in the form (51), but can be used with other high-level models that are better adapted to the natural plant dynamics. Below we will see an example where adapting the high-level dynamics makes a difference in terms of performance.

4.3. Kinematic Redundancy Elimination

A number of kinematic planning methods that map desired end-effector velocities into joint velocities have been studied. Such methods start from the identity (47) and look for ways to invert $J(\mathbf{q})$. The majority of available methods lead to solutions of the form

$$\dot{\mathbf{q}} = J(\mathbf{q})_{W(\mathbf{q})}^{\dagger} \dot{\mathbf{p}}^*, \quad (55)$$

and can be interpreted as minimizing $\dot{\mathbf{q}}^T W(\mathbf{q}) \dot{\mathbf{q}}/2$ subject to the constraint (47). We will explore numerically three such pseudo-inverses:

$$\text{Moore–Penrose: } W(\mathbf{q}) = I,$$

$$\text{Dynamically consistent: } W(\mathbf{q}) = M(\mathbf{q}),$$

$$\text{Impedance control: } W(\mathbf{q}) = I - \frac{\partial}{\partial \mathbf{q}} (J(\mathbf{q})^T \bar{\mathbf{f}}),$$

$$\bar{\mathbf{f}} = (J(\mathbf{q})^T)^{\dagger} (\mathbf{q}_0 - \mathbf{q}). \quad (56)$$

The Moore–Penrose pseudo-inverse minimizes joint velocity, while the “dynamically consistent” one minimizes kinetic energy.⁴³ The pseudo-inverse labeled “impedance control” minimizes a quantity that does not have a name, but has an integrability property which yields repeatable joint motion.⁴⁴ In our

implementation of that method, $\bar{\mathbf{f}}$ is the end-effector force resulting from joint stiffness I and joint displacement $\mathbf{q}_0 - \mathbf{q}$, where \mathbf{q}_0 is the configuration illustrated in Figure 4(a).

Note that the dynamically consistent pseudo-inverse J_M^{\dagger} also appears in the operational space formulation (despite the fact that the OS controller is driven by desired accelerations rather than desired velocities). Instead of J_M^{\dagger} , our method as shown in Eq. (54) involves the pseudo-inverse $J_{M^2}^{\dagger}$. If one were to apply Eq. (55) using our pseudo-inverse, that would correspond to minimizing $\dot{\mathbf{q}}^T M(\mathbf{q})^2 \dot{\mathbf{q}}/2 = \|M(\mathbf{q})\dot{\mathbf{q}}\|^2/2$ —which is the vector norm of the generalized momentum. However, we hesitate to propose this or any other method for kinematic redundancy elimination, because we do not believe that redundancy elimination during planning is a good idea in the first place.

Mapping desired end-effector velocities to joint velocities has traditionally been used for open-loop planning. In contrast, our method makes it possible to perform closed-loop control with end-effector velocity as the high-level control signal. All we have to do is redefine \mathbf{y} as $\mathbf{y} = \mathbf{p}$, while the plant is still $\mathbf{x} = [\mathbf{q}; \dot{\mathbf{q}}]$. Now $\mathbf{h}(\mathbf{x}) = \mathbf{k}(\mathbf{x})$ and so $H(\mathbf{x}) = [J(\mathbf{q}), 0]$. The prescribed high-level dynamics are $\mathbf{f}(\mathbf{y}) = 0$, $G(\mathbf{y}) = I$, and the plant dynamics are the same as in Eq. (50). In this case, we have a dynamic compatibility problem: $H(\mathbf{x})B(\mathbf{x}) = 0$, and so in order to apply our method, the functions \mathbf{a}, B have to be replaced with their predictive form $\tilde{\mathbf{a}}, \tilde{B}$.

5. NUMERICAL SIMULATIONS

5.1. Comparison to Robotic Control Methods

Our method was compared to existing methods in Robotics on a family of robotic manipulators [Figure 4(a)], which were simulated with the Matlab Robotics Toolbox.⁴⁵ The manipulators had between 2 and 10 hinge joints and moved in the plane. The kinematics were scaled so that all links of a given manipulator were equal in length, and the end effector (filled circle) was 1 m away from the base in the shown configuration. The dynamics were also scaled, so that all links had the same mass and the sum of all link masses was 1 kg. Material density was kept constant, and set so that a 1 m cylindrical link with diameter 0.1 m had mass 1 kg. The link diameter for each manipulator was computed given

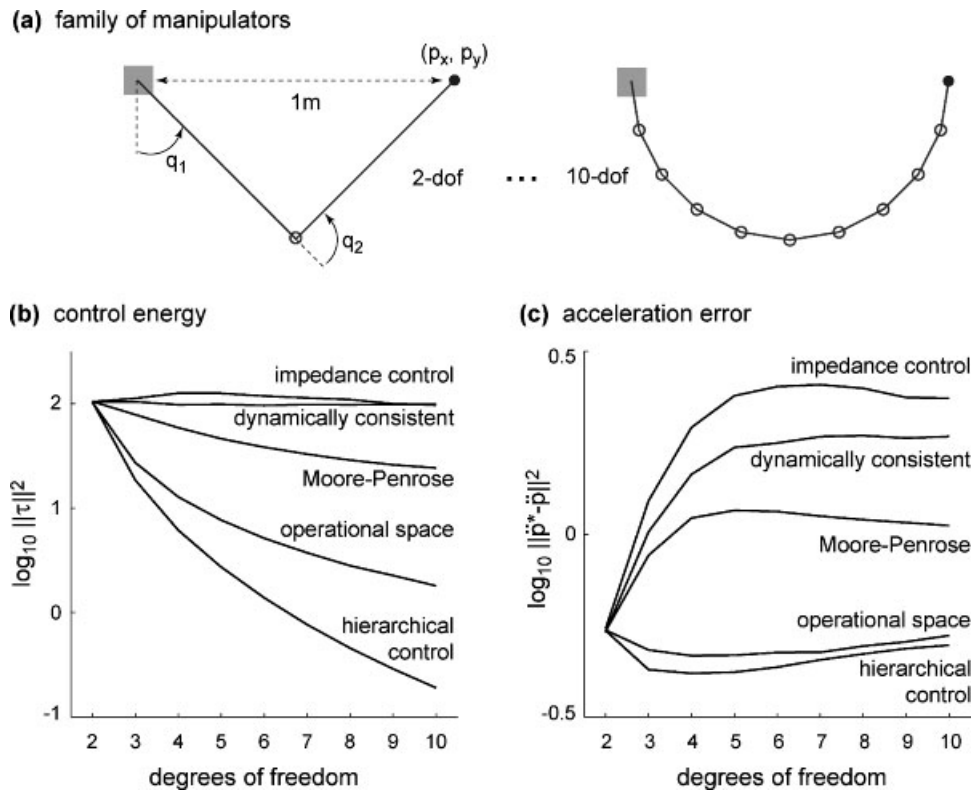


Figure 4. Comparison of control methodologies on a family of robotic manipulators.

the constraint on total mass and the fixed material density. Link inertia was then found assuming uniform density.

We compared the three kinematic methods given in Eq. (56) and the two dynamics methods given in Eq. (54), as follows. A database of joint positions \mathbf{q} and joint velocities $\dot{\mathbf{q}}$, along with desired end-effector accelerations $\ddot{\mathbf{p}}^*$, was generated randomly. For each manipulator, the database had 50,000 entries sampled uniformly from $q_i \in [-\pi, \pi]$, $\dot{q}_i \in [-1, 1]$, $\ddot{p}_i^* \in [-10, 10]$. Each method was used to compute a torque τ that yields end-effector acceleration $\ddot{\mathbf{p}}^*$ when the plant is in state $[\mathbf{q}; \dot{\mathbf{q}}]$. The dynamic methods (54) perform that computation directly. Using the kinematic methods (56) is more complicated, but possible. We first computed the current end-effector velocity $\dot{\mathbf{p}}$ from Eq. (47), and the desired velocity at a time $\Delta = 0.01$ s later as $\dot{\mathbf{p}}^*(\Delta) = \dot{\mathbf{p}} + \Delta \ddot{\mathbf{p}}^*$. Then, $\dot{\mathbf{p}}^*(\Delta)$ was mapped to $\dot{\mathbf{q}}^*(\Delta)$ using each kinematic method, the desired joint acceleration was computed as $\ddot{\mathbf{q}}^* = (\dot{\mathbf{q}}^*(\Delta) - \dot{\mathbf{q}}) / \Delta$, and used in the inverse dynamics Eq. (43) to obtain the torque τ . De-

spite this complication, both kinematic and dynamic methods achieved the desired end-effector acceleration in the absence of noise.

We compared two indices of performance, averaged over the entire database for each manipulator. The first index was control energy [Figure 4(b)]. In the nonredundant case [degree of freedom (DOF) = 2] all methods are identical, but as redundancy increases we see clear differences. The dynamic methods outperform the kinematic ones, and in particular our method (hierarchical control) outperforms all others. The second index is the error in end-effector acceleration, caused by control-multiplicative noise injected in the computed torque: $\tau_i^{\text{actual}} = (1 + 0.1\varepsilon_i)\tau_i^{\text{computed}}$, where $\varepsilon_i \sim \mathcal{N}(0, 1)$. The results in Figure 4(c) are not identical to Figure 4(b), despite the fact that noise is proportional to torque, because the nonlinear mapping from torques to end-effector accelerations makes a difference. But the rank order of the five methods remains unchanged.

The above simulations are encouraging—especially since they do not reflect the full power of our method. This is because the version (54) of the

method was derived for trivial high-level dynamics (51), not adapted to the plant dynamics. The next example illustrates the advantage of such adaptation. We defined an optimal control problem with the following total cost:

$$\|\mathbf{p}(T) - \mathbf{p}^*\|^2 + w\Delta \sum_{t=1}^{T-1} \|\boldsymbol{\tau}(t)\|^2, \quad (57)$$

where \mathbf{p}^* is the target of a reaching movement executed with the 2-DOF manipulator, $T=500$ is the number of time steps, step duration is $\Delta=0.001$ s (resulting in a 0.5 s movement), and $w=0.0001$ is the weight for the control cost. Targets were 0.3 m away from the starting position, in eight equally spaced directions. This simulation also included gravity (downward direction in the figures).

We compared the performance of our method, with $\mathbf{y}=[\mathbf{p};\dot{\mathbf{p}}]$, applied to trivial high-level dynamics (51) versus high-level dynamics adapted to the plant. The latter were obtained by fitting a linear dynamical system model to data from the former set of simulations. The dataset contained entries of the form $(\mathbf{p}, \dot{\mathbf{p}}, \ddot{\mathbf{p}}, \tilde{\mathbf{v}}, r)$, where $\tilde{\mathbf{v}}$ is the end-effector force $\tilde{\mathbf{v}}=(JJ^T)^{-1}J\boldsymbol{\tau}$, and r is the control cost $r=w\Delta\|\boldsymbol{\tau}\|^2$. There were 4,000 such entries, corresponding to the 500 time steps in each of the eight movement directions. The model fitting was based on the equations

$$\begin{aligned} \ddot{\mathbf{p}} &= F[\mathbf{p};\dot{\mathbf{p}}] + G\tilde{\mathbf{v}}, \\ r &= \tilde{\mathbf{v}}^T R \tilde{\mathbf{v}}. \end{aligned} \quad (58)$$

These equations are linear in the unknown matrices F, G, R , and therefore can be solved in a least squares sense via linear regression. The fit accounted for 87% of the variance in $\ddot{\mathbf{p}}$ and 95% of the variance in r , despite the nonlinear dynamics of the manipulator. Once F, G, R were computed, the adapted high-level dynamics could be defined as Eq. (58) by replacing $\tilde{\mathbf{v}}$ with \mathbf{v} . The fit also yielded a control cost model $\mathbf{v}^T R \mathbf{v}$ on the high level. In case of the trivial dynamics, that cost model was in the same form but with $R=w\Delta I$. Thus, in both cases we have linear high-level dynamics, and quadratic cost

$$\|\mathbf{p}(T) - \mathbf{p}^*\|^2 + \sum_{t=1}^{T-1} \mathbf{v}(t)^T R \mathbf{v}(t), \quad (59)$$

which corresponds to a linear-quadratic optimal control problem. The optimal high-level controller in




Figure 5. Effects of adapting the high-level dynamics to the plant dynamics.

each case was found with standard linear-quadratic techniques, and the resulting two-level control scheme was applied to the manipulator.

Results are shown in Figure 5. For the trivial dynamics (before learning) the end-effector paths were straight as expected. For the adapted dynamics (after learning) the paths became systematically curved. As a result of this strategy, the total cost (57) was reduced by 23%. The improvement was largely due to the adapted controller's ability to take advantage of gravity.

5.2. Application to Arm Movements, and Comparison to Optimal Control

Here, we apply our method to a model of the human arm [Figure 6(a)]. The details of this model can be found elsewhere,⁴⁶ so we will keep the description brief. The skeletal kinematics and dynamics are identical to the above 2-DOF robotic manipulator (apart from parameter settings), but the muscle actuators make this model more complex: The dynamics become third order, and the controls become constrained. The muscles acting on the human arm in the horizontal plane can be organized into six actuator groups: Shoulder flexors (SF) and extensors (SX), elbow flexors (EF), and extensors (EX), biarticular flexors (BF), and extensors (BX). The tension produced by muscle i depends on its physiological cross-sectional area (PCSA) and activation state a_i , as well as the muscle length and velocity. The substantial length-and-velocity dependence is illustrated in

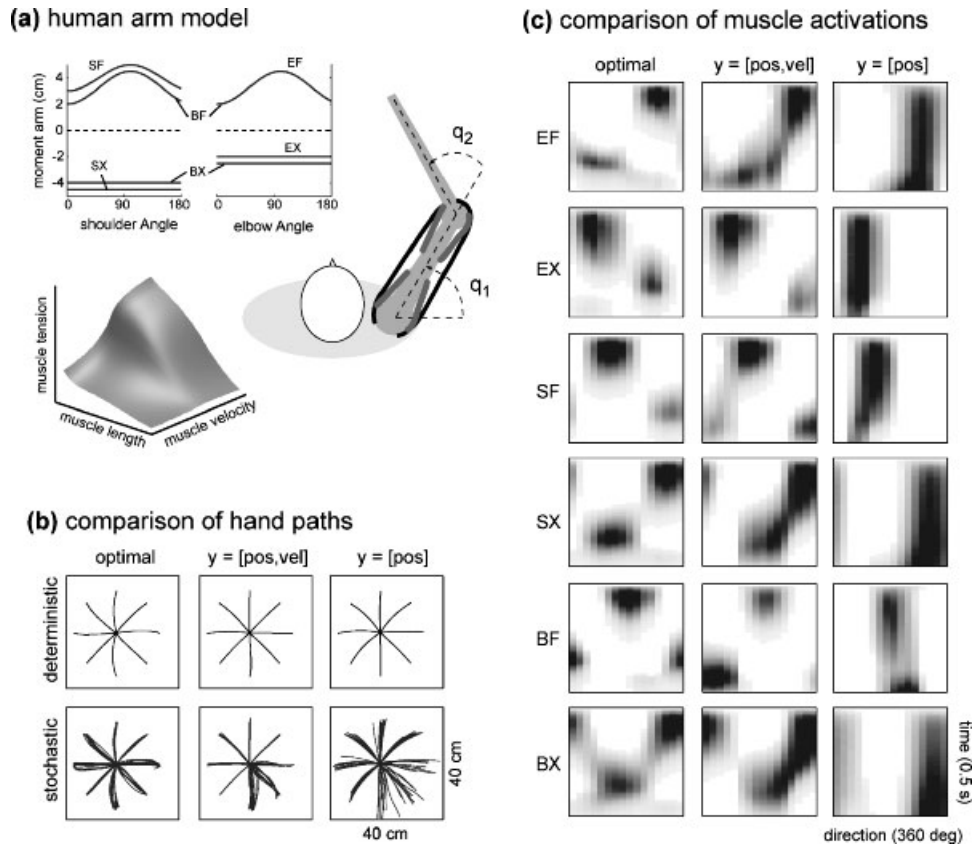


Figure 6. Comparison to the optimal controller in human reaching.

Figure 6(a), for maximal activation $a=1$. This surface is based on the Virtual Muscle model,⁴⁷ which provides a state-of-the-art fit to a range of physiological data. Tensions are multiplied by moment arms to yield joint torques. Moment arms are posture dependent [Figure 6(a)], in a way consistent with experimental data. They also determine the mapping from joint positions \mathbf{q} and velocities $\dot{\mathbf{q}}$ to muscle lengths and velocities. Muscle activation states have first-order low-pass filter dynamics $\dot{a}_i = (u_i - a_i) / \tau$. The controls u_i are constrained to the range $[0,1]$. The complete system has a 10D state vector $\mathbf{x} = [q_1; q_2; \dot{q}_1; \dot{q}_2; a_1; \dots; a_6]$ and a 6D control vector $\mathbf{u} = [u_1; \dots; u_6]$.

The task is reaching: the arm starts from rest at $(q_1 = \pi/4, q_2 = \pi/2)$, and has to reach a specified target in 0.5 s, with minimal control energy. The cost per step is

$$s(t) \|\mathbf{p}(t) - \mathbf{p}^*\|^2 + \mathbf{u}(t)^T \mathbf{R} \mathbf{u}(t), \quad (60)$$

where $s(t)=1$ when $t > 0.4$ s and 0 otherwise (we want the hand to remain stationary at the target for

the last 0.1 s). The diagonal matrix \mathbf{R} has entries proportional to PCSA_i —because larger muscles generate more force per unit activation, and therefore consume more metabolic energy. Targets are arranged in a circle with 15 cm radius around the start position.

We designed two hierarchical controllers, for $\mathbf{y} = \mathbf{p}$ and $\mathbf{y} = [\mathbf{p}; \dot{\mathbf{p}}]$, respectively. Because of the third-order dynamics, our method relied on the predictive functions $\tilde{\mathbf{a}}, \tilde{\mathbf{B}}$ in both cases. The control inequality constraints made the explicit solutions Eqs. (11) and (37) inapplicable, and therefore, the low-level controller used Quadratic Programming at each time step to compute $\mathbf{u}(\mathbf{v}, \mathbf{x})$. Instead of modeling the high-level dynamics, here we used the implicit method described above (which yields a high-level feedback controller). In addition to the two hierarchical controllers, we also computed the optimal feedback controller in each case, using our iterative linear-quadratic-Gaussian (iLQG) method.³⁸

Figure 6(b) shows deterministic and stochastic hand paths for each controller. Stochastic trajectories were simulated by corrupting each control signal

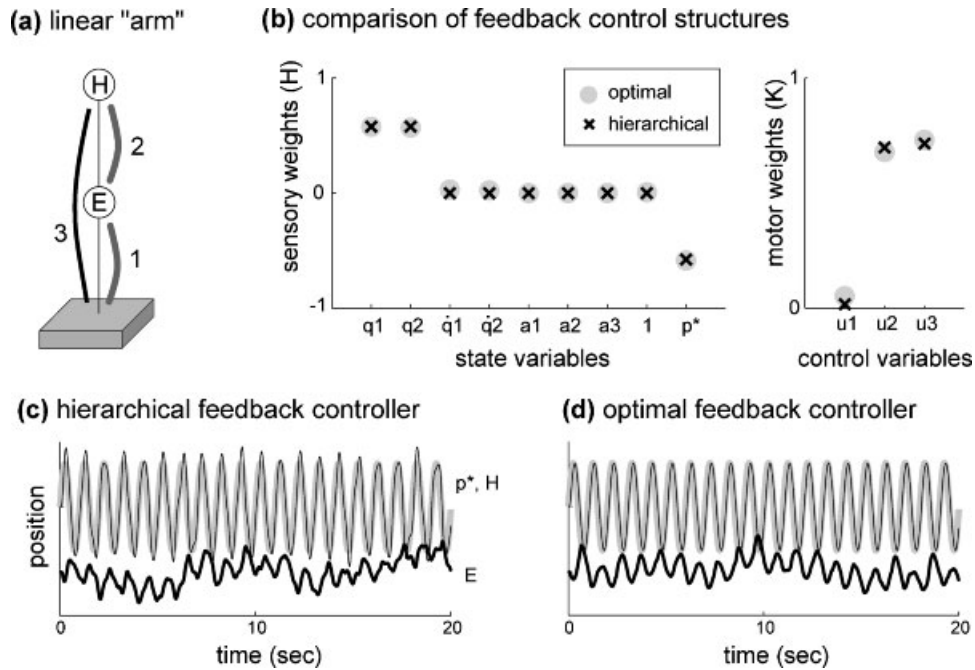


Figure 7. Internal dynamics and sensorimotor synergies.

with 100% control-multiplicative noise. The differences on the level of kinematics are small, although they increase (especially for the $\mathbf{y}=\mathbf{p}$ controller) when noise is added. Note that although the $\mathbf{y}=[\mathbf{p};\dot{\mathbf{p}}]$ controller appears to be more accurate than the optimal controller, that is not the case—the optimal controller allows more variability during the movement but reaches the target more accurately. Note also that the controllers can sustain such a large amount of motor noise because we are not modeling sensory noise and delays.

Figure 6(c) shows the (open-loop) muscle activations. Each subplot is one muscle-controller combination. The horizontal axis corresponds to movement direction, while the vertical axis is time during the movement (increasing downward). Dark means higher activation. We now see a much more clear distinction between the two hierarchical controllers. The muscle activations found by the $\mathbf{y}=[\mathbf{p};\dot{\mathbf{p}}]$ controller are quite similar to the optimal muscle activations, and furthermore resemble many features of experimentally observed muscle activations (but that is beyond the scope of the paper). On the other hand, the $\mathbf{y}=\mathbf{p}$ controller misses the elaborate temporal pattern of muscle activation, although it still activates the appropriate muscles.

Apart from demonstrating the power and gener-

ality of the proposed method, these results are important with regard to prior models of end-effector arm movement control.^{15–20} Our method completes these models by going all the way to muscle activations. It also reveals a problem. The relatively poor performance of the $\mathbf{y}=\mathbf{p}$ controller suggests that controlling hand position through instantaneous velocity commands is not a good idea—because such commands are too far removed from the muscles that have to carry them out. In contrast, high-level commands related to hand force rather than velocity afford hierarchical control that is much closer to optimal.

5.3. Internal Dynamics and Sensorimotor Synergies

Here, we study a tracking task executed with a redundant linear plant. Linearity allows us to look more closely at the structure of the hierarchical controller, and compare it to the structure of the corresponding optimal controller. We also explore the behavior within the redundant subspace where the high-level controller cannot act.

The plant [Figure 7(a)] consists of two points (H,E) with mass 1, driven by three linear muscle-like actuators that can both push and pull. Actuator i has

activation state a_i and generates force $a_i - b\dot{q}_i - k(q_i - \bar{q}_i)$, where $b=10$, $k=10$ are intrinsic damping and stiffness, q_i is the length of the actuator, and $\bar{q}_i=0.3$ is the resting length of the spring. We can think of point H as the “hand,” and point E as the “elbow” of a telescopic linear “arm.” At each time step, we are given a target position p^* . The task is to track the target with the hand. The state vector is $\mathbf{x} = [q_1; q_2; \dot{q}_1; \dot{q}_2; a_1; a_2; a_3; 1; p^*]$. The constant 1 is needed to implement the spring model, and p^* is included in the state because it varies over time. Note that q_3 is not included in the state because $q_3 = q_1 + q_2$. The activation dynamics are $\dot{a}_i = (u_i - a_i) / \tau$ with time constant 40 ms. The control vector is $\mathbf{u} = [u_1; u_2; u_3]$. Tracking is formulated as an optimal control problem with cost-per-step $\mathbf{x}^T Q \mathbf{x} + \mathbf{u}^T R \mathbf{u}$. We set $R=10^{-5}$, and $Q=H^T H$, where $H = [1, 1, 0, 0, 0, 0, 0, 0, -1]$. This encodes the tracking error $(q_1 + q_2 - p^*)^2$, where $q_1 + q_2$ is hand position. Representing $p^*(t)$ as Brownian motion, the problem becomes LQG and can be solved with standard methods. The optimal control law is $\mathbf{u} = L^\circ \mathbf{x}$, where L° is a 3×9 matrix of optimal feedback gains. L° will be compared to the hierarchical controller below.

The high-level state is defined as the instantaneous tracking error:

$$y = q_1 + q_2 - p^*. \quad (61)$$

Note that $y = H\mathbf{x}$, with H as defined above. The high-level dynamics and cost are $\dot{y} = v$ and y^2 , respectively. Because of the linearity of the plant, the mapping from v to \mathbf{u} is now linear: $\mathbf{u} = K v$, where K is computed from Eq. (11). The optimal high-level control law is also linear: $v = g y$, where the gain g is found via LQG techniques. Then the hierarchical controller is

$$\mathbf{u} = L\mathbf{x}, \quad \text{where } L = KgH. \quad (62)$$

To compare the structure of the hierarchical and optimal controllers, we applied singular value decomposition to the optimal L° and found that it has singular values (447.6, 1.9, 0)—which is essentially rank 1. So, we can write the optimal controller as

$$\mathbf{u} = L^\circ \mathbf{x}, \quad \text{where } L^\circ \approx K^\circ g^\circ H^\circ \quad (63)$$

and K° and H° are the left and right singular vectors of L° corresponding to the largest singular value. The elements of vectors H and H° (normalized to unit length) are compared in Figure 7(b) (left), and K

and K° are compared in Figure 7(b) (right). The similarity in both cases is striking. Note that of these four vectors, the only one we set manually was H . Then K was computed automatically by our method, while K° and H° fell out of (nonhierarchical) optimal feedback control. In summary, the structure of the hierarchical and optimal feedback controllers can be very similar.

This linear “arm” model is interesting because the augmented plant has internal dynamics: Point E is free to move without affecting the task outcome. Of course it cannot move arbitrarily, since the actuators have stiffness and damping that couple Points E and H. But the high-level controller has no means of monitoring and correcting the trajectory of E, because it sees only the trajectory of H. To explore the internal dynamics, we applied both the hierarchical and optimal control laws, with 100% control-multiplicative noise. Figures 7(c) and 7(d) show the trajectories of Points H and E. While the “hand” H tracks the target reliably, the “elbow” E exhibits substantial variability for both controllers. The lack of control over E may seem undesirable, but in fact there is nothing undesirable about it—or else the optimal controller would not allow E to fluctuate. If we really want E to follow a certain trajectory, we should be honest about that fact and encode it in the performance criterion. Since the performance criterion for this task is not affected by E, we cannot blame our controller for failing to accomplish something we did not ask for. On the contrary, we should be pleased that it closely resembles the optimal controller.

6. DISCUSSION

We presented a general framework for hierarchical feedback control of redundant systems. The design of the proposed control hierarchy involves: (i) Specifying a set of high-level parameters and their desired dynamics; (ii) designing a low-level feedback controller which yields an augmented plant with the specified input-output behavior; (iii) designing a high-level feedback controller that solves the original control problem but operates on a simplified system. Our focus was on automating the design of the two feedback controllers, as well as the construction of high-level dynamics that mimic the plant dynamics. This provided a way of controlling the specified high-level parameters.

The choice of appropriate high-level parameters is presently an open question, although a few rel-

evant comments can be made. One extreme case would be to use a scalar performance index as the only high-level parameter. Then the high-level commands will essentially say “perform better,” but will not contain information as to how to perform better. Clearly, such commands are not useful to the low-level controller. The other extreme would be to include all the state information on the high level (perhaps transformed nonlinearly), but that would not in general lead to a simplification of the controller design problem. So we seek a middle ground, where the high-level state contains task-related parameters that enter in the computation of the performance index, and perhaps other closely related parameters. The application to reaching movements, for example, revealed a benefit of including velocity as a high-level parameter, even though the state-dependent cost was only a function of position.

Another issue deserving closer examination is the internal dynamics. In the above numerical examples, we did not attempt to shape that dynamics, although our framework provides a mechanism for doing so by specifying a potential function over the plant state. Yet, we saw that both the structure and the behavior of the hierarchical controller can be very similar to the optimal controller. This issue is relevant not only in our framework, but in all methods for hierarchical control. Velocity control methods that yield integrable solutions^{38,48} eliminate the internal dynamics altogether, by enforcing a one-to-one mapping from plant states to high-level (end-effector) states. However, such strategies are suboptimal and unlikely to be used by the nervous system. An alternative approach is to somehow choose high-level parameters that result in asymptotically stable internal dynamics (corresponding to the notion of “minimum phase” in feedback linearization). Asymptotic stability is of course appealing, but its relationship to optimal performance is unclear; how to achieve such stability without unnecessary waste of control energy is also unclear. Since we already know that optimal controllers for redundant tasks are in essence hierarchical, further analysis of their structure can illuminate the choice of both the high-level state and the potential function used to shape the internal dynamics.

ACKNOWLEDGMENTS

This work was supported by US National Institutes of Health grant R01-NS045915. E. Todorov is a Sloan Research Fellow.

REFERENCES

1. C. Sherrington, *The integrative action of the nervous system*, Yale University Press, New Haven, 1906.
2. N. Bernstein, *Dexterity and its development*, in *Dexterity and its development*, M. Latash and M. Turvey (Editors), Lawrence Erlbaum, Mahwah, NJ, 1996.
3. G. Loeb, I. Brown, and E. Cheng, A hierarchical foundation for models of sensorimotor control, *Exp. Brain Res.* 126:(1) (1999), 1–18.
4. J. Kalaska, L. Sergio, and P. Cisek, Cortical control of whole-arm motor tasks, in *Sensory guidance of movement: Novartis Foundation Symposium*, M. Glickstein (Editor), Wiley, Chichester, UK, 1998.
5. R. Brooks, A robust layered control-system for a mobile robot, *IEEE J. Rob. Autom.* 2:(1) (1986), 14–23.
6. O. Khatib, A unified approach to motion and force control of robotic manipulators: The operational space formulation, *IEEE J. Rob. Autom.* 3:(1) (1987), 43–53.
7. J. Pratt, C. Chew, A. Torres, P. Dilworth, and G. Pratt, Virtual model control: An intuitive approach for bipedal locomotion, *Int. J. Robot. Res.* 20:(2) (2001), 129–143.
8. D. Precup and R. Sutton, Multitime models for temporally abstract planning, *Advances in Neural Information Processing Systems*, Vol. 10, M. Jordan, M. Kearns, S. Solla (Editors), MIT Press, Cambridge, MA, 1998, pp. 1050–1056.
9. T. Dietterich, Hierarchical reinforcement learning with the maxq value function decomposition, *J. Artif. Intell. Res.* 13 (2000), 227–303.
10. E. Todorov, Optimality principles in sensorimotor control, *Nat. Neurosci.* 7:(9) (2004), 907–915.
11. E. Todorov and M. Jordan, Optimal feedback control as a theory of motor coordination, *Nat. Neurosci.* 5:(11) (2002), 1226–1235.
12. E. Todorov and M. Jordan, A minimal intervention principle for coordinated movement, *Advances in neural information processing systems*, Vol. 15, S. Becker, S. Thrun, and K. Obermayer (Editors), MIT Press, Cambridge, MA, 2002, pp. 27–34.
13. N. Bernstein, *The coordination and regulation of movements*, Pergamon, New York, 1967.
14. M. Latash, On the evolution of the notion of synergy, *Motor control, today and tomorrow*, G. Gantchev, S. Mori, and J. Massion (Editors), Academic, Sofia, 1999, pp. 181–196.
15. G. Hinton, Parallel computations for controlling an arm, *J. Motor Behav.* 16:(2) (1984), 171–194.
16. A. Pellionisz, Coordination: A vector-matrix description of transformations of overcomplete cns coordinates and a tensorial solution using the Moore-Penrose generalized inverse, *J. Theor. Biol.* 110 (1984), 353–375.
17. D. Meyer, R. Abrams, S. Kornblum, C. Wright, and J. Smith, Optimality in human motor performance: Ideal control of rapid aimed movements, *Psychol. Rev.* 95 (1988), 340–370.
18. D. Bullock and S. Grossberg, Neural dynamics of planned arm movements: emergent invariants and speed-accuracy properties during trajectory formation, *Psychol. Rev.* 95 (1988), 49–90.

19. B. Hoff and M. Arbib, Models of trajectory formation and temporal interaction of reach and grasp, *J. Motor Behav.* 25:(3) (1993), 175–192.
20. E. Torres and D. Zipser, Reaching to grasp with a multi-jointed arm. I: Computational model, *J. Neurophysiol.* 88:(5) (2002), 2355–2367.
21. A. Isidori, *Nonlinear control systems*, Springer, London, 1995.
22. H. Khalil, *Nonlinear systems*, Prentice–Hall, Englewood Cliffs, NJ, 2002.
23. G. Sutton and K. Sykes, The variation of hand tremor with force in healthy subjects, *J. Physiol. (London)* 191:(3) (1967), 699–711.
24. R. Schmidt, H. Zelaznik, B. Hawkins, J. Frank, and J. Quinn, Motor-output variability: A theory for the accuracy of rapid motor acts, *Psychol. Rev.* 86:(5) (1979), 415–451.
25. E. Todorov, Cosine tuning minimizes motor errors, *Neural Comput.* 14:(6) (2002), 1233–1260.
26. A. Collins, A. Ruina, R. Tedrake, and M. Wisse, Efficient bipedal robots based on passive-dynamic walkers, *Science* 307 (2005), 1082–1085.
27. F. Anderson and M. Pandy, Dynamic optimization of human walking, *J. Biomech. Eng.* 123:(5) (2001), 381–390.
28. P. Sabes, M. Jordan, and D. Wolpert, The role of inertial sensitivity in motor planning, *J. Neurosci.* 18:(15) (1998), 5948–5957.
29. W. Li, E. Todorov, and X. Pan, Hierarchical optimal control of redundant biomechanical systems, 26th Annual IEEE Conference on Engineering in Biology and Medicine, 2004.
30. G. Loeb, W. Levine, and J. He, Understanding sensorimotor feedback through optimal control, *Cold Spring Harb Symp. Quant Biol.* 55 (1990), 791–803.
31. B. Hoff, A computational description of the organization of human reaching and prehension, Ph.D. thesis, University of Southern California, 1992.
32. A. Kuo, An optimal control model for analyzing human postural balance, *IEEE Trans. Biomed. Eng.* 42 (1995), 87–101.
33. Y. Shimansky, T. Kang, and J. He, A novel model of motor learning capable of developing an optimal movement control law online from scratch, *Biol. Cybern.* 90 (2004), 133–145.
34. J. Scholz and G. Schoner, The uncontrolled manifold concept: identifying control variables for a functional task, *Exp. Brain Res.* 126:(3) (1999), 289–306.
35. J. English and A. Maciejewski, On the implementation of velocity control for kinematically redundant manipulators, *IEEE Trans on Systems Man Cybern. Part A* 30 (2000), 233–237.
36. A. Bryson and Y. Ho, *Applied optimal control*, Blaisdell, MA, 1969.
37. D. Jacobson and D. Mayne, *Differential dynamic programming*, Elsevier, New York, 1970.
38. E. Todorov and W. Li, A generalized iterative LQG method for locally-optimal feedback control of constrained nonlinear stochastic systems, *American Control Conference*, 2005.
39. M. Branicky, V. Brokar, and S. Mitter, A unified framework for hybrid control: Model and optimal control theory, *IEEE Trans. Autom. Control* 43 (1998), 31–41.
40. S. Sastry, Hybrid control in air traffic management systems, *IEEE Conference on Decision and Control*, vol. 34, 1995.
41. S. Hedlund and A. Rantzer, Optimal control of hybrid systems, *IEEE Conference on Decision and Control*, vol. 43, 1999.
42. J. Pratt and G. Pratt, Intuitive control of a planar bipedal walking robot, *IEEE International Conference on Robotics and Automation (ICRA '98)*, 1998.
43. H. Bruyninckx and O. Khatib, Gauss' principle and the dynamics of redundant and constrained manipulators, *IEEE Int Conference Robotics Automation*, 2000, pp. 2563–2568.
44. A. Mussa-Ivaldi and N. Hogan, Integrable solutions of kinematic redundancy via impedance control, *Int. J. Robot. Res.* 10 (1991), 481–491.
45. P. Corke, A robotics toolbox for Matlab, *IEEE Rob. Autom. Mag.* 3 (1996), 24–32.
46. W. Li and E. Todorov, Iterative linear-quadratic regulator design for nonlinear biological movement systems, *Proc 1st International Conference on Informatics in Control, Automation and Robotics*, 2004.
47. I. Brown, E. Cheng, and G. Loeb, Measured and modeled properties of mammalian skeletal muscle. II. The effects of stimulus frequency on force-length and force-velocity relationships, *J. Muscle Res. Cell Motil.* 20:(7) (1999), 627–643.
48. J. Baillieul, Avoiding obstacles and resolving kinematic redundancy, *Proc IEEE Int Conference Robotics and Automation*, 1986, pp. 1698–1704.