Proceedings of the 26th Annual International Conference of the IEEE EMBS
San Francisco, CA, USA • September 1-5, 2004

HIERARCHICAL OPTIMAL CONTROL OF REDUNDANT BIOMECHANICAL SYSTEMS

Weiwei Li[1], *Emanuel Todorov*[2] *and Xiuchuan Pan*[2]

[1] Mechanical and Aerospace Engineering, University of California San Diego, La Jolla, CA 92093-0411
[2] Department of Cognitive Science, University of California San Diego, La Jolla, CA 92093-0515

ABSTRACT

Sensorimotor control occurs simultaneously on multiple levels. We present a general approach to designing feedback control hierarchies for redundant biomechanical systems, that approximate the (non-hierarchical) optimal control law but have much lower computational demands. The approach is applied to the task of reaching, using a detailed model of the human arm. Our hierarchy has two levels of feedback control. The high level is designed as an optimal feedback controller operating on a simplified virtual plant. The low level is responsible for transforming the dynamics of the true plant into the desired virtual dynamics. The new method may be useful not only for modelling the neural control of movement, but also for designing Functional Electric Stimulation systems that have to achieve task goals by activating muscles in real time.

Keywords— Hierarchical control, Optimal feedback control, Redundant systems

1. INTRODUCTION

Neural control of movement is accomplished by a complex hierarchy of recurrently connected brain regions. Understanding how the multiple levels of the sensorimotor system cooperate to produce integrated action has been a central theme in Neuroscience since the pioneering work of Sherrington. Perhaps the most thorough investigation of the levels of human motor control was undertaken by Bernstein [1], who concluded that every motor skill involves at least two levels of feedback control: a leading level that monitors progress and steers the biomechanical plant towards the achievement of the task goals, and a background level that provides various automatisms and corrections that help the leading level. More recently, recordings in a variety of motor areas have revealed a progression of increasingly abstract (and increasingly hard to describe) neural representations. But despite this long standing interest and the wealth of relevant data and intuitive notions, quantitative theories of hierarchical sensorimotor control are surprisingly rare [3]. Indeed, the computational ideas that can provide a solid foundation for such theories have emerged only recently. The objective of the present paper is to synthesize these ideas into a general approach to hierarchical feedback control, and illustrate the approach with a specific model of arm movement control.

This work is supported by the National Institutes of Health Grant R01-NS045915.

2. GENERAL APPROACH

A distinguishing feature of biomechanical plants is their massive redundancy. Let x be a state vector (representing joint angles, joint velocities, muscle activation states, and relevant aspects of the environment), and $q(x)$ be a cost function that quantifies the notion of a "task". Redundancy means that the function $q(x)$ depends on the state x only through a reduced set of (more abstract) variables $p(x)$, i.e. $q(x)$ can be written as $q(p(x))$. In the task of reaching, for example, performance depends only on the position of the hand relative to the target.

Our goal is to design controllers for the dynamical system $\dot{x} = \mathrm{f}(x, u_x)$ that minimize the performance criterion

$$J = h(x(T)) + \int_0^T q(x(t)) + r\|u_x(t)\|^2 dt \qquad (1)$$

Our recent theory of coordination [5] has revealed a somewhat surprising fact about the optimal way to control redundant systems. We have shown that optimal feedback controllers obey a "minimal intervention" principle: rather than correcting all deviations from a desired trajectory, such controllers only act in the task-relevant subspace and leave task-irrelevant deviations uncorrected. Loosely speaking, this is done by extracting a small set of task-relevant features (through a set of sensory synergies), performing feedback control in that feature space, and mapping the resulting abstract controls into control signals for the real actuators (through a set of motor synergies). The difficulty with optimal feedback controllers for complex systems is that they are extremely hard to design – perhaps even for the brain.

The minimal intervention principle suggests a natural approximation to optimal feedback control, using a two-level feedback control hierarchy. The basic idea (see Fig

1) is to design optimal feedback controllers that optimize the true cost function, but for a simpler virtual plant with state $y(x)$. This corresponds to Bernsteins leading level of control, and generates an abstract control signal u_y interpreted as desired change in y. The background level then has the job of transforming the dynamics of the actual plant into that of the virtual plant, i.e. computing the real control signal $u_x(x, u_y)$ such that the desired change in $y(x)$ is accomplished, with minimal control effort. We propose the following principles for designing the virtual dynamics $\dot{y} = g(y; w) + u_y$:

1. It should be sufficiently simple, so that optimal control methods become feasible;
2. It should be sufficiently close to the true dynamics, so that the transformer does not have to guess blindly (recall that the transformer does not know what the task is);
3. It should contain $p(x)$ as state variables, so that the true cost $q(p(x))$ can be measured and optimized by the leading level;
4. The internal model of the virtual dynamics (needed by the leading level to design feedback control laws) should be improved through system identification, so as to compensate for unavoidable errors in the translation process. This corresponds to learning the parameters w via supervised learning.

We emphasize that a virtual plant whose state is only $p(x)$ will most likely violate principle 2. Instead, we include in the virtual dynamical state $y(x)$ both $p(x)$ and its derivatives, so as to match the order of the differential equations governing the true dynamics. In the task of reaching, the virtual state will include hand position, velocity, and net muscle force (expressed in hand space); the virtual control signal will specify desired change in force.

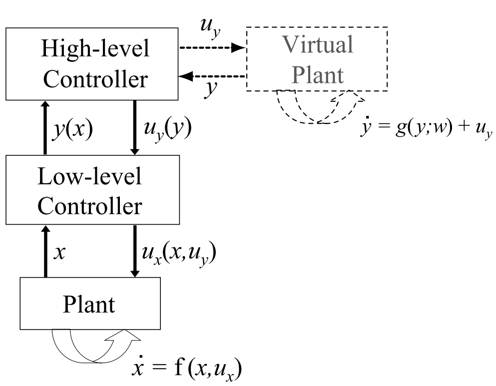

Fig. 1. Illustration of the hierarchical control structure

The optimal feedback controller on the high level will be designed using an iterative Linear-Quadratic-Gaussian (iLQG) method that we have developed recently. This method extends the efficient and well-developed LQG framework to the domain of nonlinear control. Briefly, the iteration starts with some control law that is applied to the nonlinear system – obtaining an average trajectory and control sequence. We then linearize the dynamics and quadratize the cost in the vicinity of that state-control trajectory, apply dynamic programming locally, and use the result to improve the initial control law. The algorithm has quadratic convergence, similar to a Newton's method. For more details see [7].

The low-level has to activate muscles so as to accomplish desired effects in the task-relevant space. Since there are more muscles than task variables, the low-level faces a redundancy problem which will be solved using static optimization. At each point in time, we will obtain a linear relationship between muscle activations and task variables, and then find the minimal feasible muscle activations via Quadratic Programming.

3. APPLICATION TO REACHING MOVEMENT

The approach outlined above will now be applied to the problem of reaching with a realistic arm. The arm model contains 2 links, 6 muscles with varying moment arms, muscle length-velocity-tension curves based on the Virtual Muscle model, and activation dynamics modelled as a nonlinear low-pass filter. See [7] for details.

3.1. The two levels

The high-level dynamics is defined in end-effector space (hand space), and modelled using end-effector equations of motion. The construction of the end-effector dynamic model is achieved by expressing the relationships of hand position p, velocities v, as well as the net muscle force f acting on the hand.

$$\dot{p} = v, \tag{2}$$
$$\dot{v} = f/m + H_v(p, v, f; \mathrm{w}_v), \tag{3}$$
$$\dot{f} = -\alpha(f - mg) + u_y + H_f(p, v, f; \mathrm{w}_f), \tag{4}$$

where m is the average hand mass, u_y is the control signal, H_v and H_f are the function of p, v and f with unknown weighting coefficients w_v and w_f respectively. Since the model (2)-(4) is only an approximation, we need to use some learning processes in the hierarchical control architecture, which can estimate those unknown parameters iteratively online before a successful implementation can be achieved.

Furthermore, we need to establish the relationship between the high-level state (containing hand position p, velocity v, as well as the net muscle force f in hand space), and the state of the low-level dynamic model

$$p = \mathcal{K}(\theta), \tag{5}$$

$$v = J(\theta)\dot{\theta}, \tag{6}$$

$$J^T(\theta) f = \tau_{mus} - \mathcal{G}(\theta), \tag{7}$$

where $\mathcal{K}$ is the transformation of position θ from joint space to the end-effector space (hand space); the Jacobian matrix $J(\theta) = \partial\mathcal{K}/\partial\theta$ transforms velocities in joint space, $\dot{\theta}$, to Cartesian velocity v of the end-effector expressed in hand space.

The control signal u_y is the output of the high-level loop which specifies a desired change of force in the end-effector space. The objective is to find the control u_y that minimizes the performance criterion

$$V = \frac{1}{2}\int_0^T \left[r_1 \|u_y(t)\|^2 + s(t)\, \|p(t) - p^*\|^2 \right] dt, \tag{8}$$

where p^* is the target, $r_1 = 1e - 6$, and $s(t)$ is 0 when the hand is allowed to move, and 1 when the hand is required to be at the target. The optimal feedback controller is designed using the iLQG method [7].

The forward dynamics of the arm can be expressed as

$$\ddot{\theta} = \mathcal{M}(\theta)^{-1}(\tau_{mus} - \mathcal{C}(\theta, \dot{\theta}) - \mathcal{G}(\theta)), \tag{9}$$

where $\theta \in R^2$ is the joint angle vector (shoulder: θ_1, elbow: θ_2), $\mathcal{M}(\theta) \in R^{2\times 2}$ is a positive definite symmetric inertia matrix, $\mathcal{C}(\theta, \dot{\theta}) \in R^2$ is a vector centripetal and Coriolis forces, $\mathcal{G}(\theta)$ is the gravity, and $\tau_{mus} \in R^2$ is the joint torque.

The tension $T(a, l, v)$ produced by a muscle obviously depends on the muscle activation a, but also varies substantially with the length l and velocity v of that muscle

$$T(a, l, v) = aF_{vl}(l(\theta), v(\theta, \dot{\theta})) + F_p(l(\theta)). \tag{10}$$

The joint torque generated by the muscles is given by

$$\tau_{mus} = M(\theta)\; T(a, l(\theta), v(\theta, \dot{\theta})), \tag{11}$$

where $M(\theta)$ denotes muscle moment arms.

Muscle activation a is not equal to instantaneous neural input u_x, but is generated by passing u_x through a filter that describes calcium dynamics. It is modelled as the following

$$\dot{a} = \alpha(u_x - a), \tag{12}$$

where α is the muscle decay constant. Combining (9)-(12), the low-level dynamic system can be obtained where the state is given by $x = (\theta_1\ \theta_2\ \dot{\theta}_1\ \dot{\theta}_2\ a_1\ \cdots\ a_6)^T$.

Equation (7) represents the fundamental relationship between the net muscle force f and the joint torque τ_{mus} consistent with the end-effector and arm dynamic equations. With (7), (10), (11) and the assumption $\dot{a} \gg \theta, \dot{\theta}$, we can derive $\dot{f} = J^+(\theta) M(\theta) F_{vl}(\theta, \dot{\theta})\ \dot{a}$, where $J^+(\theta)$ denotes the pseudo inverse. Based on (12) and the assumption $\dot{f} = -\alpha(f - c) + u_y$, we obtain

$$u_y = \alpha Q u_x, \qquad Q = J^+(\theta) M(\theta) F_{vl}(\theta, \dot{\theta}). \tag{13}$$

Now we can design the low-level controller by solving a constrained quadratic optimization problem

$$\min_{u_x} \frac{1}{2} u_x^T H u_x + b^T u_x \tag{14}$$

subject to

$$0 \le u_x \le 1, \tag{15}$$

where $H = \alpha^2 Q^T Q + r_2 I,\ b = -\alpha Q^T u_y$, and $r_2 = 0.001$. Therefore we can find the low-level control u_x that activate the muscles in a way that achieves the specified change of force acting on the hand.

3.2. Simulation Results

We chose different pairs of starting postures and targets, and applied the hierarchical control scheme described above. Hand trajectories are shown in Fig 2. The black curves are the actual trajectories of the hand, that result from the coupling of the two-level hierarchy with the detailed arm model. The gray curves are the trajectories that would have resulted from applying the feedback control law to the virtual dynamical system. Note that before learning the "virtual trajectories" are straight, because we do not have nonlinearities in the initial virtual model. However, after the system identification stage the virtual model is improved, and it now contains nonlinear terms. As a result, both the virtual and real trajectories become curved, and more importantly, they get closer to each other. Over the set of movements we simulated, the average distance between the

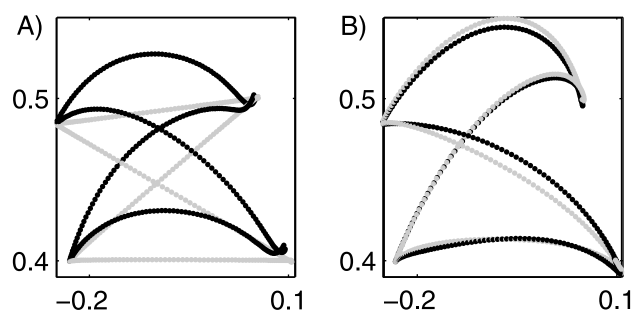

Fig. 2. Trajectories in hand space: (A) before learning, (B) after learning. Gray lines – trajectories obtained by applying the high-level feedback controller to the virtual dynamics. Black lines – trajectories obtained by applying the hierarchical control scheme to the real plant.

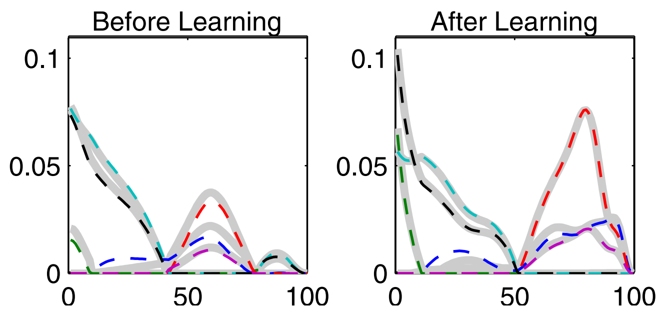

Fig. 3. Comparison of the muscle control sequences generated by our hierarchical controller (dashed lines) vs. the non-hierarchical optimal controller (thick gray lines).

virtual and actual trajectory was 2.9 cm before learning, and only 0.8 cm after learning. Additionally, the cost achieved by the hierarchical control system after learning was 21% lower compared to its value before learning.

Finally, we compared the behavior of the hierarchical control scheme to that of an optimal (non-hierarchical) controller designed for the real plant. Note that the plant we are studying, although quite complex, is still amenable to iterative methods for nonlinear optimal control. We used the control sequences generated by our control scheme as an initial control law, and then improved that control law via gradient descent on the performance criterion. Both before and after learning, we found that the controls generated by our method are close to a local minimum of the unconstrained problem (see Fig 3). After learning, the distance to the nearest local minimum was about 50% less than it was before learning.

4. DISCUSSION

Here we presented a general approach to approximately-optimal hierarchical feedback control of redundant systems, and illustrated its application in the context of reaching with a realistic model of the human arm. Imposing any predefined control structure is likely to result in suboptimal performance. However, our simulations demonstrate that the suboptimality due to the hierarchical structure is negligible (especially after learning). At the same time, the computational demands are much lower than what is required for designing an optimal feedback controller. The performance of the new method in more complex control problems remains to be established. It will be particularly interesting to identify the behavioral differences between the present control scheme and the (unstructured) optimal controller, and examine those differences in light of experimental data.

While this may be the first comprehensive approach to hierarchical sensorimotor control, the main computational ideas underlying it are not entirely new. Feedback transformations that create easier-to-control virtual plants have been used in robotics; the Operational Space Framework [2] and Virtual Model Control [4] are particularly relevant. The new computational aspects of our work are: (i) application of optimal feedback control to the virtual plant; (ii) principles for designing the virtual model that make it suitable for optimal control; (ii) continuous improvement of the internal model of virtual dynamics, compensating for unavoidable transformation errors.

The present approach is also related to our recent work [6], where we pursued the idea of simplifying feedback transformation combined with high-level optimal control. Our previous feedback transformation was ”extracted” from the plant dynamics using unsupervised learning, rather than being designed to capture the task-relevant variables as we did here. As a result we were forced to use inefficient policy-gradient methods for controller optimization, rather than the efficient model-based iLQG method [7] employed here. Both approaches have advantages, that we hope to combine in future work. The basic idea is to make the low-level itself hierarchical: starting with a task-independent virtual model that captures the properties of the plant (and is acquired via unsupervised learning), and then creating higher-level virtual models suitable for specific tasks.

5. REFERENCES

[1] N.A. Bernstein. *On dexterity and its development*. Lawrence Erlbaum Associates, 1996.

[2] O. Khatib. A unified approach for motion and force control of robot manipulators: The operational space formulation. *IEEE J. of Robotics and Automation*, Vol. RA-3, No.1, 1987.

[3] G.E. Loeb, I.E. Brown and E.J. Cheng. A hierarchical foundation for models of sensorimotor control. *Exp Brain Res* 126:118, 1999.

[4] J. Pratt, C.M. Chew, et al. Virtual model control: An intuitive approach for bipedal locomotion. *Int. J. of Robotics Research*, Vol. 20, No. 2, pp. 129-143, 2001.

[5] E. Todorov and M. Jordan. Optimal feedback control as a theory of motor coordination. *Nat Neurosci*, Vol.5, No.11, 1226-1235, 2002.

[6] E. Todorov and Z. Ghahramani. Unsupervised learning of sensory-motor primitives. *Proceedings of the 25th IEEE EMBS Conference*, 2003.

[7] E. Todorov and W. Li. A generalized iterative LQG method for locally-optimal feedback control of constrained nonlinear stochastic systems. Submitted to the *43rd IEEE Conf on Decision and Control*, 2004.