Properties of Synergies Arising from a Theory of Optimal Motor Behavior

Manu Chhabra
Department of Computer Science, University of Rochester, Rochester, NY 14627, U.S.A.

Robert A. Jacobs
Department of Brain and Cognitive Sciences, University of Rochester, Rochester, NY 14627, U.S.A.

We consider the properties of motor components, also known as synergies, arising from a computational theory (in the sense of Marr, 1982) of optimal motor behavior. An actor's goals were formalized as cost functions, and the optimal control signals minimizing the cost functions were calculated. Optimal synergies were derived from these optimal control signals using a variant of nonnegative matrix factorization. This was done using two different simulated two-joint arms—an arm controlled directly by torques applied at the joints and an arm in which forces were applied by muscles—and two types of motor tasks—reaching tasks and via-point tasks.

Studies of the motor synergies reveal several interesting findings. First, optimal motor actions can be generated by summing a small number of scaled and time-shifted motor synergies, indicating that optimal movements can be planned in a low-dimensional space by using optimal motor synergies as motor primitives or building blocks. Second, some optimal synergies are task independent—they arise regardless of the task context—whereas other synergies are task dependent—they arise in the context of one task but not in the contexts of other tasks. Biological organisms use a combination of task-independent and task-dependent synergies. Our work suggests that this may be an efficient combination for generating optimal motor actions from motor primitives. Third, optimal motor actions can be rapidly acquired by learning new linear combinations of optimal motor synergies. This result provides further evidence that optimal motor synergies are useful motor primitives. Fourth, synergies with similar properties arise regardless if one uses an arm controlled by torques applied at the joints or an arm controlled by muscles, suggesting that synergies, when considered in "movement space," are more a reflection of task goals and constraints than of fine details of the underlying hardware.

1 Introduction

Marr (1982) defined three levels of analysis of a complex information processing device. The top level, known as the computational theory, examines what the device does and why. A distinguishing feature of this level is that it provides an explanation for why a device does what it does by studying the device's goals. Although there may be many different ways of developing a computational theory of aspects of human behavior, an increasingly popular way is through optimal models that formalize goals as mathematical constraints or criteria, search for behaviors that optimize the criteria, and compare the optimal behaviors with human behaviors. If there is a close match, then it is hypothesized that people are behaving as they do because they are efficiently satisfying the same goals as were built into the optimal model.

This approach is commonplace in the study of human motor behavior (see Todorov, 2004, for a review). Flash and Hogan (1985), for example, proposed an optimal model of how people plan trajectories for reaching movements. This model emphasizes that trajectories should be smooth—the model searches for trajectories that minimize the jerk of a movement (i.e., the third derivative of position with respect to time). It is able to explain the fact that reaches tend to move along straight lines and tend to have bell-shaped velocity profiles. Harris and Wolpert (1998) developed an optimal model of motor control that attempts to minimize the variance of the point reached at the end of a movement despite motor noise whose magnitude is dependent on the size of the control signals. They showed that this model explains several aspects of both eye movements and hand reaches.

This article is concerned with motor synergies arising from a computational theory (in the sense of Marr, 1982) of optimal motor behavior. To understand motor synergies, it is helpful to first understand the degrees of freedom problem (Bernstein, 1967). Biological motor systems typically have many degrees of freedom, where the degrees of freedom in a system are the number of dimensions in which the system can independently vary (Rosenbaum, 1991). Because the number of degrees of freedom of a system carrying out a task often exceeds the number of degrees of freedom needed to specify the task, the degrees of freedom are typically redundant (Jordan & Rosenbaum, 1989). Consider, for example, the problem of touching the tip of your nose. The location of your nose has three degrees of freedom (its x, y, and z position in Cartesian coordinates), but the joints of your arm have seven degrees of freedom (the shoulder has three degrees of freedom, and the elbow and wrist each have two). Consequently, there are many different settings of your arm's joint positions that all allow you to touch your nose. Which setting should you use?

A solution to this problem is to create motor synergies, which are dependencies among dimensions of the motor system. For example, a motor synergy might be a coupling of the motions of your shoulder and elbow.

Motor synergies provide two types of benefits to motor systems. First, synergies ameliorate the problem of redundancy—they can constrain the set of possible shoulder, elbow, and wrist positions that allow you to touch your nose. Second, synergies reduce the number of degrees of freedom that must be independently controlled, thereby making it easier to control a motor system (Bernstein, 1967). Because synergies make motor systems easier to control, they are often hypothesized to serve as motor primitives, building blocks, or basis functions: they provide fundamental units of motor behavior that can be linearly combined to form more complex units of behavior.

Investigators of motor control are attempting to develop a comprehensive understanding of biological motor synergies. Typically, these researchers analyze neuroscientific or behavioral data using mathematical techniques in order to derive the motor synergies used by an organism. Sanger (1995) analyzed people's cursive handwriting using principal component analysis (PCA) to discover their motor synergies. He showed that linear combinations of these synergies closely reconstructed human handwriting. Thoroughman and Shadmehr (2000) studied people's motor learning behaviors to derive motor synergies based on gaussian radial basis functions. They showed that linear combinations of these synergies matched people's behaviors when adapting to new environmental conditions. Mussa-Ivaldi, Giszter, and Bizzi (1994) identified frogs' motor synergies by stimulating sites in their spinal cords and verified that stimulation of two sites leads to the vector summation of the forces generated by stimulating each site separately.

A possible confusion in the motor control literature is that synergies derived from neuroscientific or behavioral data using mathematical techniques are sometimes referred to as "optimal." For example, Sanger (1995) derived synergies from human behavioral data using PCA, a linear optimal dimensionality-reduction technique, and referred to the results as "optimal movement primitives." It is important to keep in mind, however, that these synergies arise from optimal analysis of people's actions and are not necessarily the same ones as would arise from a computational theory (again, in the sense of Marr, 1982) of optimal motor behavior. Based on the discussion above, a computational theory might involve a model that formalizes the actor's goals as mathematical criteria and searches for the actions that optimize the criteria. An optimal analysis of the optimal actions could then derive the motor synergies. Synergies discovered in this way would be "optimal" in the sense that they arise from a computational theory of optimal motor behavior.

To date, we know of only one study of motor synergies that arise from a computational theory. Todorov and Jordan (2002) proposed a computational theory that uses an optimal feedback controller as a model of motor coordination and noted that this controller produces motor synergies. In brief, the controller implements the "principle of minimal intervention"—it

does not attempt to control a system along dimensions that are irrelevant for a task. Because the system's degrees of freedom are controlled along some task dimensions but not others, couplings or synergies among the degrees of freedom arise. Todorov and Jordan thereby explained the emergence of synergies.

Although Todorov and Jordan (2002) explained the emergence of synergies; they did not study the specific properties of these synergies. In contrast, this article details the properties of synergies arising from a theory of optimal motor behavior. We have created an optimal controller for a nonlinear system that formalizes goals as mathematical constraints and searches for control signals that optimize the constraints. This was done using two different simulated two-joint arms—an arm controlled directly by torques applied at the joints and an arm in which forces are applied by muscles—and two types of motor tasks—reaching tasks (move an end effector from one point to another) and via-point tasks (move from one point to another while passing through an intermediate point). In all cases, we derived synergies from the optimal control signals using an extension to nonnegative matrix factorization (d'Avella, Saltiel, & Bizzi, 2003) and studied the properties of these synergies.

Our studies of the resulting motor synergies reveal several interesting findings. First, optimal motor actions can be generated by summing a small number of scaled and time-shifted motor synergies, indicating that optimal movements can be planned in a low-dimensional space by using optimal motor synergies as motor primitives or building blocks. Second, some optimal synergies are task independent—they arise regardless of the task context—whereas other synergies are task dependent—they arise in the context of one task but not in the contexts of other tasks. Biological organisms use a combination of task-independent and task-dependent synergies. Our work suggests that this may be an efficient combination for generating optimal motor actions from motor primitives. Third, optimal motor actions can be rapidly acquired by learning new linear combinations of optimal motor synergies. This result provides further evidence that optimal motor synergies are useful motor primitives. Fourth, synergies with similar properties arise regardless if one uses an arm controlled directly by torques applied at the joints or an arm controlled by muscles, suggesting that synergies, when considered in "movement space," are more a reflection of task goals and constraints than of fine details of the underlying hardware.

2 Computing the Optimal Control Signals

We simulated a two-joint arm that can be characterized as a second-order nonlinear dynamical system (e.g., Hollerbach & Flash, 1982):

$$\mathcal{M}(\theta)\ddot{\theta} + \mathcal{C}(\theta, \dot{\theta}) + \mathcal{B}\dot{\theta} = \tau \qquad (2.1)$$

Table 1: Parameter Values Used in the Simulation of a Two-joint Arm.

Parameter	Value	Parameter	Value
b_{11}	0.05 kgm^2 s^{-1}	b_{22}	0.05 kgm^2 s^{-1}
b_{21}	0.025 kgm^2 s^{-1}	b_{12}	0.025 kgm^2 s^{-1}
m_1	1.4 kg	m_2	1.0 kg
l_1	0.30 m	l_2	0.33 m
I_1	0.025 kgm^2	I_2	0.045 kgm^2
s_1	0.11 m	s_2	0.16 m

Note: The parameters $\{b_{ij}\}$ denote the (i, j)th elements of a joint friction matrix; m_i, l_i, and I_i denote the mass, length, and moment of inertia of the ith link, respectively; and s_i denotes the distance from the ith joint to the ith link's center of mass.

where τ is a vector of torques, θ is a vector of joint angles, $\mathcal{M}(\theta)$ is an inertial matrix, $\mathcal{C}(\theta, \dot\theta)$ is a vector of coriolis forces, and $\mathcal{B}$ is a joint friction matrix. We used the same parameter values for the arm as Li and Todorov (2004). These values are listed in Table 1.

We studied two types of tasks: reaching tasks and via-point tasks. In a reaching task, the arm must be controlled so that its end effector moves from a start location to a target location. A via-point task is identical except that there is an additional requirement that the end effector also move through an intermediate location known as a via-point.

For any reaching or via-point task, there are many time-varying torque vectors $\tau(t)$ that will move the arm so that it successfully performs the task. As discussed above, this multiplicity of control solutions is due to redundancy in the two-joint arm and is known as the degrees-of-freedom problem. How do we choose a particular solution? According to the optimality framework, an actor's goals are formalized as mathematical constraints that are combined in a cost function, and an optimal control signal is a signal that minimizes this function.

For the reaching task, we used the following cost function,

$$J(\tau(t)) = \frac{1}{2} \parallel e(T) - e^* \parallel^2 + k_1 \parallel \dot e(T) \parallel^2 + \frac{k_2}{2} \int_0^T \tau(t)^T \tau(t)dt, \qquad (2.2)$$

where k_1 and k_2 are constants (we used the same values as Todorov & Li, 2005: $k_1 = 0.001$ and $k_2 = 0.0001$), T is the duration of the movement, $e(T)$ is the end-effector location at time T, and e^* is the target location at time T. The first term penalizes reaches that deviate from the target location, the second term penalizes reaches that do not have a zero velocity at the end of the movement, and the third term penalizes reaches that require large torques (or "energy"). This cost function has previously been used by Li and Todorov (2004; see also Todorov & Li, 2005). Minimization of this function results in control signals that produce reaches with several properties of

natural movements, including bell-shaped velocity profiles, lower velocities at higher curvatures, and near-zero velocities at the beginnings and ends of movements.

For the via-point task, we modified the above cost function to also penalize movements that do not pass through the via-point midway through the movement. The cost function has the form

$$J(\tau(t)) = \frac{1}{2} \parallel e(T) - e^* \parallel^2 + \frac{1}{2} \parallel e(T/2) - e_v^* \parallel^2$$

$$+ k_1 \parallel \dot{e}(T) \parallel^2 + \frac{k_2}{2} \int_0^T \tau(t)^T \tau(t) dt, \tag{2.3}$$

where e_v^* is the via-point or desired end-effector location at the middle of the movement. This function penalizes reaches that deviate from the via-point at time $T/2$.

To find the optimal control signal for a reaching or via-point task, the corresponding cost function must be minimized. Unfortunately, when using nonlinear systems such as the two-joint arm described above, this minimization is computationally intractable. Researchers typically resort to approximate methods to find locally optimal solutions. We used one such method, known as the iterative linear quadratic regulator (iLQR), developed by Li and Todorov (2004; see also Todorov & Li, 2005). We now briefly summarize this method in a generic setting.

A continuous-time linear dynamical system is given by

$$\dot{x}(t) = f(x(t), u(t)), \tag{2.4}$$

where x is the state of the system and u is the input control signal. For the two-joint arm described above, the state x is $(\theta, \dot{\theta})^T$, and the control u is τ. Consider a cost function of the form

$$J(u(t)) = \sum_i h_i(x(t_i)) + \int_0^T l(x(t), u(t)) dt. \tag{2.5}$$

Note that the cost functions for the reaching and via-point tasks are of this form. In the cost function for the reaching task, for example, the two discrete penalties (deviation of the end-effector location from the target location at time T and deviation of the end-effector velocity from zero at time T) correspond to the first term on the right-hand side of equation 2.5, and a continuous energy-like cost for large torques corresponds to the second term.

The iLQR starts with an initial guess of the optimal control signal and iteratively improves it. From the control signal $u_i(t)$ at iteration i, the trajectory $x_i(t)$ is computed using a standard Euler approximation. The algorithm

uses three steps to find the control signal for the next iteration. It starts by linearly approximating the dynamical system given in equation 2.4 and quadratically approximating the cost function given in equation 2.5. These approximations are made around $(x_i(t), u_i(t))$ at each time step t. Using these approximations, a modified linear quadric gaussian (LQG) control problem is then formulated in the $(\delta x, \delta u)$ space, where $x + \delta x$ and $u + \delta u$ are improved approximations to x and u, respectively. This formulation is valid only where these approximations are accurate: a small region around $(x_i(t), u_i(t))$. Finally, the optimal correction to control $u_i(t)$ at iteration i, denoted $\delta u_i^*(t)$, is computed by solving this modified LQG problem. This step requires the solution to a modified Riccatti-like set of equations. Fortunately, finding this solution is computationally efficient. Once the optimal corrections have been obtained, the algorithm sets $u_{i+1}(t) = u_i(t) + \delta u_i^*(t)$ and proceeds to the next iteration. The algorithm stops if there is no significant improvement in the trajectory. The end result is a locally optimal trajectory x^* and locally optimal control signal u^*.

We have found that the iLQR works well on both reaching and via-point tasks when using the two-joint arm. Figure 1 shows examples of optimal reaching (top row) and via-point (bottom row) movements computed by the iLQR. The graphs in the left column show the movement of the end effector (horizontal and vertical axes give the x and y coordinates of the end effector in Cartesian space), whereas the graphs in the right column show the velocity profiles (horizontal axes represent time, and vertical axes represent velocity of the end effector). Clearly, the iLQR produces smooth movements with bell-shaped velocity profiles. In addition, the velocity profile for the via-point movement (bottom-right graph) indicates that end-effector velocity decreases with increasing path curvature.

3 Obtaining Optimal Synergies

As discussed above, motor synergies are dependencies among dimensions of a motor system. They are useful because they can ameliorate the problem of redundancy and because they reduce the number of degrees of freedom that must be independently controlled, thereby making it easier to control a motor system. Synergies are often hypothesized to serve as motor primitives, building blocks, or basis functions.

Researchers have used a variety of methods to compute motor synergies. We used a variant of nonnegative matrix factorization developed by d'Avella et al. (2003). This algorithm requires two inputs. One input is the number of synergies, denoted N. The other is a matrix of control signals, where each control signal is a $2 \times T$ matrix of optimal torques computed by the iLQR for a given task (this matrix has $2 \times T$ elements because torques are applied to both joints of the two-joint arm at each time step of a movement, and there are T time steps per movement). The input matrix of control signals is a vertical stack of individual control signal matrices.
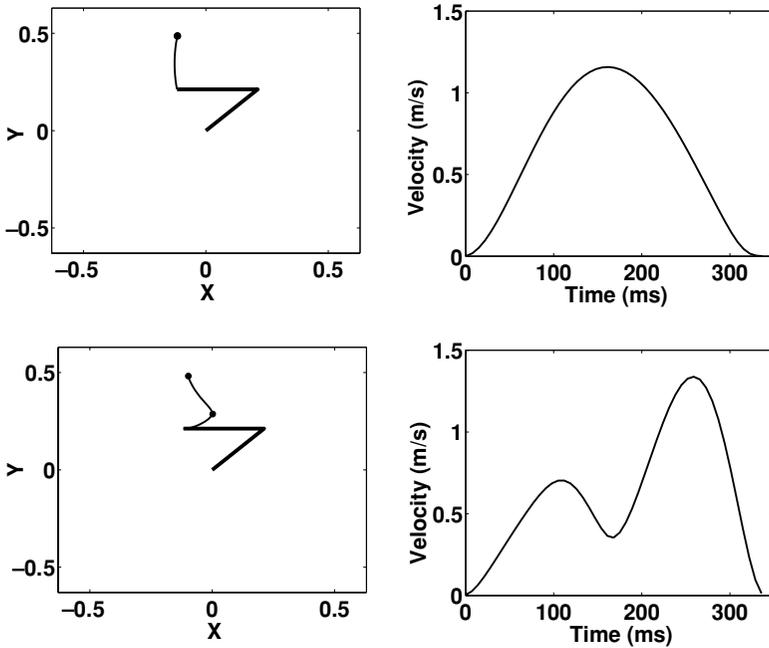
Figure 1: Examples of optimal reaching (top row) and via-point movements (bottom row) computed using the iLQR. The graphs in the left column show the movement of the end effector: the horizontal and vertical axes represent the *x* and *y* coordinates of the end effector in Cartesian space. The thick lines show the orientation of the two-joint arm at the start of the movement, and the thin lines show the path of the end effector. The graphs in the right column show the velocity profiles: the horizontal axes represent time, and the vertical axes represent velocity of the end effector.

For example, if the iLQR was used to find the optimal control signals for 500 reaching tasks (tasks with different initial configurations of the arm or different target locations) and each reach had a duration of 400 time steps, then the matrix would consist of 1000 rows where each block of two rows is a 2×400 element matrix giving the optimal torques for each joint at each time step of a reach. As its output, the algorithm seeks a set of synergies such that every control signal can be expressed as a sum of scaled and time-shifted synergies. Mathematically, it seeks a set of N synergies, denoted $\{\mathbf{w}_i, i = 1, \ldots, N\}$, such that control signal $\mathbf{m}$ can be written as follows,

$$\mathbf{m}(t) = \sum_{i=1}^{N} c_i \mathbf{w}_i(t - t_i), \qquad (3.1)$$

where $\{c_i, i = 1, \ldots, N\}$ is a set of coefficients that scale the synergies, and $\{t_i, i = 1, \ldots, N\}$ is a set of times that time-shift the synergies. The algorithm searches for the synergies, scaling coefficients, and time shifts that minimize the sum of squared errors between the actual control signals and the reconstructed signals.

A technical detail is that the algorithm requires a set of nonnegative control signals (each element of a control vector must be nonnegative). In our case, a torque vector might have negative elements. We overcame this problem in a manner inspired by biological motor systems' use of agonist and antagonist muscles to apply torques at joints. We recoded a 2×1 torque vector as a 4×1 vector in which the first two elements give the anticlockwise and clockwise torques for the first joint (shoulder), and the last two elements provide the same information for the second joint (elbow).[1] For example, if torque $(2, -1)^T$ is applied to the joints, it means that a $+2$ torque is applied to the first joint in the anticlockwise direction, and a $+1$ torque is applied to the second joint in the clockwise direction. We recoded this torque vector to the nonnegative vector $(2, 0, 0, 1)^T$.

4 Simulation Results

This section reports the results of seven experiments. The first four experiments used the two-joint arm described above in which torques were applied at the joints. The last three experiments used the same arm, except forces were applied by muscles.

All experiments used the same collection of reaching and via-point tasks. We created 320 instances of each task as follows. Ten initial positions of the arm were randomly generated by uniformly sampling the first joint angle from the set $[-\pi/4, \pi/2]$ and the second joint angle from the set $[0, 3\pi/4]$. For each initial position, 32 target locations were generated. A target was generated by randomly selecting a movement distance (sampled uniformly from the range 10–50 cm) and an angle of movement (sampled uniformly from the range 0–2π). For the via-point task, a via-point was placed at a random angle (sampled uniformly from the set $[-\pi/3, \pi/3]$) from the line joining the initial and target locations. The via-point's distance from the

---

[1] Below we present results in which we consider the number of motor synergies required to reconstruct optimal movements with small error when an arm is controlled by torques applied directly to its joints. It is possible that the operation of mapping a two-dimensional vector with real values to a four-dimensional vector with nonnegative values introduces a bias into the estimate of this number. Nonetheless, our use of the mapping is justified as follows. We wish to compare synergies obtained when an arm is controlled by torques applied directly to its joints with synergies obtained when an arm is controlled by forces applied by muscles. Therefore, it is necessary to use the same representational format and dimensionality-reduction algorithm for obtaining synergies in both cases. When an arm is controlled by muscles, synergies are extracted on the basis of muscle activations, which must be nonnegative values.
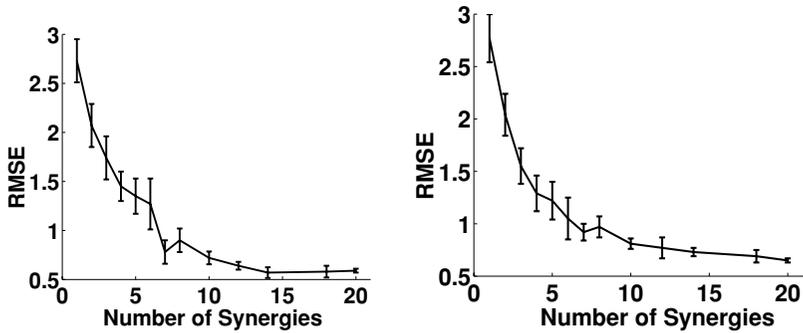
Figure 2: The graphs plot the root mean squared error (RSME) between actual and reconstructed test items for reaching (left graph) and via-point (right graph) tasks as a function of the number of synergies used in the reconstructions. The error bars give the standard errors of the means.

initial location was selected randomly to be between one-third and two-thirds of the distance between initial and target locations. The duration of a movement was 350 msec, and new torques were applied every 7 msec.

4.1 Experiment 1: A Small Set of Synergies Can Reconstruct Optimal Movements. The first experiment evaluated whether optimal reaching or via-point control signals can be expressed as a sum of a small number of scaled and time-shifted synergies. If so, then the synergies can be regarded as useful motor primitives.

For each type of task, the iLQR was applied to each instance of the task to generate 320 optimal control signals. These signals were divided into five equal-sized sets, which were then used by a fivefold cross-validation procedure to create training and test data items. Four sets of control signals were used for training, and the remaining set was used for testing. This was repeated for all five such combinations of training and test sets. During training, nonnegative matrix factorization was used as described above to discover a set of synergies. During testing, these synergies were time-shifted and linearly combined to reconstruct the test control signals. Nonnegative matrix factorization was used to find the time shifts and linear coefficients.

The results for the reaching and via-point tasks are shown in the left and right graphs of Figure 2, respectively. The horizontal axes give the number of synergies. The vertical axes give the root mean squared error (RMSE) between actual and reconstructed test control signals. The error bars show the standard errors of the means. With both reaching and via-point tasks, the error is near its minimum when relatively few synergies (about six or seven) were used. For our purposes, this is an important result
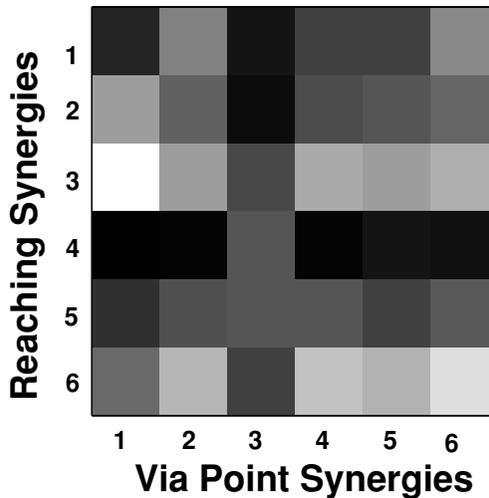
Figure 3: The similarity matrix when six synergies were obtained from the reaching task and the via-point task. The lightness of the square at row i and column j gives the cosine of the angle between the ith reaching-task synergy vector and the jth via-point task synergy vector: white is a value of 1, black is a value of 0, and intermediate gray-scale values represent intermediate values.

because it means that the synergies are useful motor primitives: optimal movements can be planned in a relatively low-dimensional space by time-shifting and linearly combining a small number of synergies. Furthermore, the fact that the error curves for the reaching and via-point tasks are similar suggests that these tasks have similar task complexity. This is surprising because generating optimal via-point movements intuitively seems more complicated than generating optimal reaches.

4.2 Experiment 2: Task-Independent and Task-Dependent Synergies. The second experiment evaluated whether optimal motor synergies are task independent or task dependent. This issue is interesting due to recent neurophysiological findings. d'Avella and Bizzi (2005), for example, recorded electromyographic activity from 13 muscles of the hind limbs of frogs performing jumping, swimming, and walking movements. An analysis of the underlying motor synergies revealed that some synergies were used in all types of movements, whereas other synergies were movement dependent.

Figure 3 shows the similarity matrix when six synergies were obtained for the reaching task and six synergies were obtained for the via-point task. The lightness of the square at row i and column j gives the cosine of the angle between the ith reaching-task synergy vector and the jth via-point
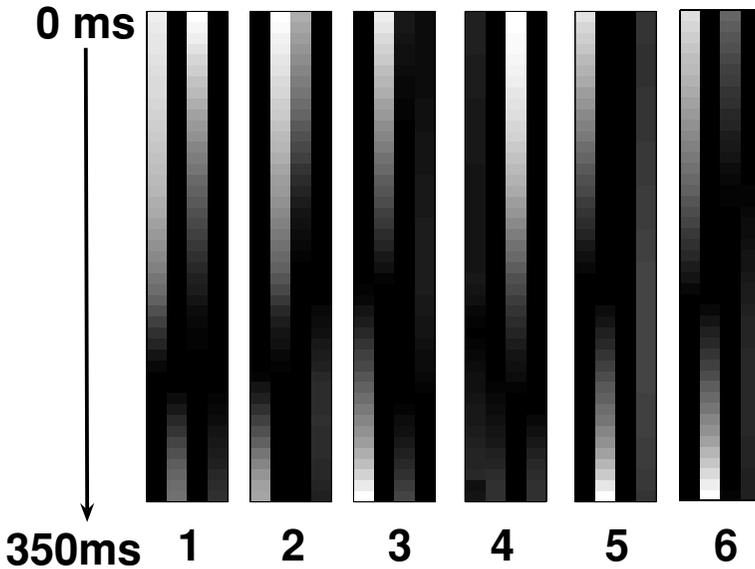
Figure 4: The six synergies obtained for the reaching task. Each synergy is represented by four columns; the first two columns represent the anticlockwise and clockwise torques for the first joint (shoulder), whereas the second two columns represent this same information for the second joint (elbow). Torques were linearly scaled to the interval [0, 1]. White indicates a torque of 1, black indicates a torque of 0, and intermediate shades of gray represent intermediate values.

task synergy vector: white is a value of 1, black is a value of 0, and intermediate gray-scale values represent intermediate values. Some synergies, such as the third reaching-task synergy and the first via-point task synergy, are highly similar, indicating that these synergies are task independent. In contrast, other synergies, such as the fourth reaching-task synergy or the third via-point task synergy, are dissimilar from all other synergies, indicating that they are task dependent. This result suggests that the combination of task-independent and task-dependent synergies found in biological organisms (e.g., d'Avella & Bizzi, 2005; Jing, Cropper, Hurwitz, & Weiss, 2004) may be efficient for generating optimal motor actions from motor synergies.

4.3 Experiment 3: Visualizing Synergies. In experiment 3, we obtained synergies for the purpose of visualizing the movements induced by these synergies. Using our collections of instances of each type of task, six synergies for the reaching task and six synergies for the via-point task were calculated as described above. The scaling coefficients for the reaching-task
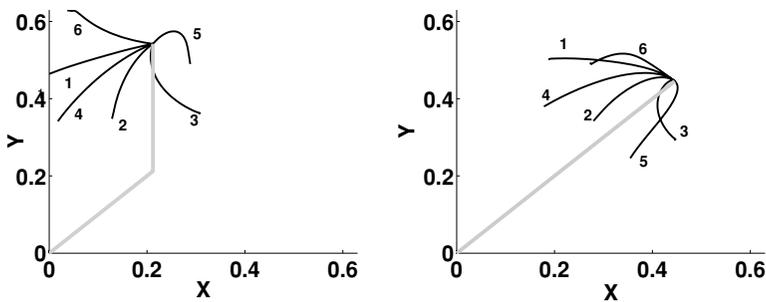
Figure 5: Movements induced by six synergies obtained for the reaching task. The horizontal and vertical axes of each graph give the x and y coordinates of the end effector in Cartesian space, the gray lines show the initial configuration of the arm, the black lines show the movements of the end effector, and the number next to each movement indicates the synergy that was applied (using the same labels as Figure 4). The left and right graphs illustrate induced movements when the initial configuration of the arm was near the center of the work space or at a far edge of the work space, respectively.

synergies or the via-point task synergies were set to their average values over the collection of reaching tasks or via-point tasks, respectively. The time-shift parameters were set to zero.

The six synergies obtained for the reaching task are illustrated in Figure 4. The horizontal axis labels the synergies, and the vertical axis depicts time. Each synergy is represented by four columns; the first two columns represent the anticlockwise and clockwise torques for the first joint, whereas the second two columns represent this same information for the second joint. Torques were linearly scaled to the interval [0, 1]. White indicates a torque of 1, black indicates a torque of 0, and intermediate shades of gray represent intermediate values.

Figure 5 illustrates movements based on these synergies. The left graph shows the induced movements when the initial arm configuration was near the center of the workspace. The horizontal and vertical axes of the graph give the x and y coordinates of the end-effector in Cartesian space, the gray lines show the initial configuration of the arm, the black lines show the movements of the end effector, and the number next to each movement indicates the synergy that was applied (using the same labels as Figure 4). The induced movements tend to be relatively straight (though some are curved) and tend to cover a wide range of directions. The right graph of Figure 5 shows the induced movements when the initial arm configuration was at a far edge of the work space. Again, the movements tend to be relatively straight. As should be expected, movements in this case are directed toward the center of the work space. Figure 5 demonstrates that synergies tend to broadly cover all possible directions of motion.

0 ms

350ms **1**   **2**   **3**   **4**   **5**   **6**
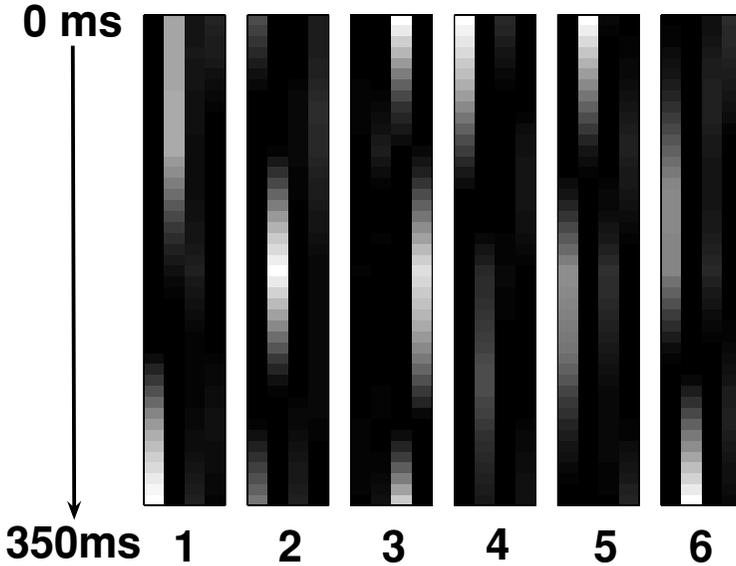
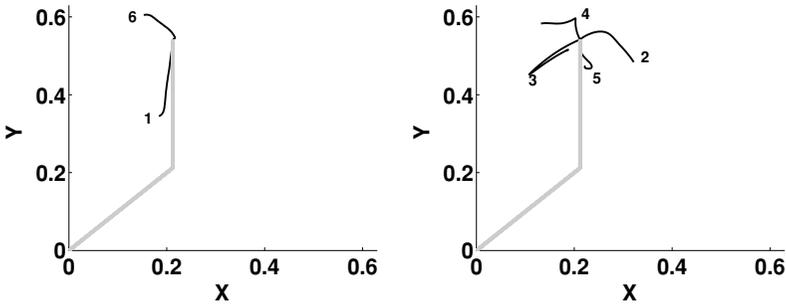Figure 6: The six synergies obtained for the via-point task.

Figure 7: Movements induced by synergies obtained from the via-point task. The left and right graphs illustrate movements induced by task-independent and task-dependent synergies, respectively. The number next to each movement indicates the synergy that was applied (using the same labels as Figure 6).

The six synergies obtained for the via-point task are illustrated in Figure 6. It uses the same format as Figure 4. Figure 7 illustrates movements based on these synergies. The left graph illustrates movements induced by two synergies that were highly similar to synergies obtained from the reaching task—that is, these are task-independent synergies. The induced movements are relatively straight. Consequently, the underlying
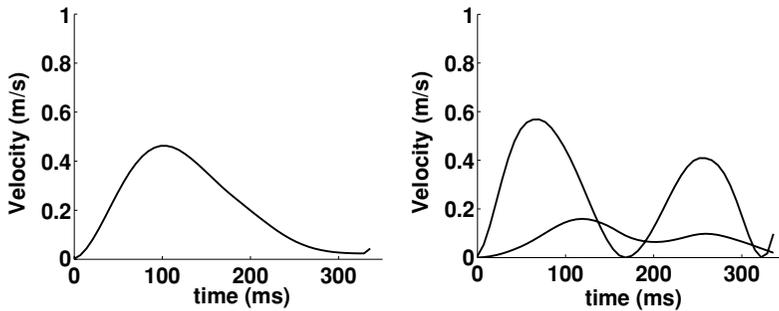
Figure 8: Velocity curves for induced movements. The left graph plots the velocity curve for a movement based on a synergy obtained from the reaching task. The right graph plots the velocity curves for two movements based on two task-dependent synergies obtained from the via-point task.

synergies are useful for both reaching and via-point tasks. The right graph illustrates movements based on four synergies that are task dependent; these synergies were not similar to synergies obtained from the reaching task. The induced movements tend to be almost piecewise linear, with a region of large curvature near the middle of the movement that is preceded and followed by regions of relatively straight motion.

Figure 8 shows the velocity curves (velocity at each moment in time) for induced movements. The left graph plots the velocity curve for a movement based on a synergy obtained from the reaching task. This curve has a bell-shaped profile, which is commonly found for reaching movements. The right graph plots the velocity curves for two movements based on two task-dependent synergies obtained from the via-point task. The shapes of these curves are typical for via-point movements.

In summary, we find that the synergies for reaching and via-point movements have intuitive forms. Movements based on synergies obtained from the reaching task tend to be straight, to broadly cover the directions available to the arm based on its initial configuration, and to have bell-shaped velocity profiles. Movements based on task-independent via-point synergies tend to have these same properties. In contrast, movements based on task-dependent via-point synergies tend to have a piecewise-linear shape with a region of high curvature near the middle of the movement and have velocity profiles with two bell shapes.

4.4 Experiment 4: Learning with Synergies. Experiment 4 evaluated whether the use of optimal motor synergies makes it easier to learn to perform new optimal motor actions. If motor synergies are useful motor primitives, then this ought to be the case.

The task was to learn to generate a reaching movement starting from an initial configuration of the arm so that the arm's end point reached a randomly selected target location. When synergies were used, control signals were expressed as linear combinations of synergies (to minimize computational demands, we did not time-shift synergies), meaning that the parameter values that needed to be learned were the linear coefficients. When synergies were not used, the values that needed to be learned were the torques applied to each joint at each moment in time.

From a collection of 320 instances of the reaching task, fivefold cross validation was used to create training and test sets. Policy gradient, a type of reinforcement learning algorithm, was used to learn estimates of the relevant parameter values (Sutton, McAllester, Singh, & Mansour, 2000). This algorithm was applied for 300 iterations. Learning with synergies occurred as follows. We calculated the optimal movements for each instance in a training set using the iLQR, and obtained four motor synergies using nonnegative matrix factorization. The policy gradient algorithm was then used to learn to perform each instance of the reaching task in the test set. At each iteration of the learning process, we numerically computed the derivatives of the reaching-task cost function (see equation 2.2) with respect to the linear coefficients used in the linear combination of synergies and performed gradient descent with the constraint that the coefficients had to be nonnegative. When learning without synergies, we computed the derivatives of the reaching-task cost function with respect to the torques at each joint and at each time step and performed gradient descent. Step sizes or learning rates that produced near-optimal performance were used when performing gradient descent with and without synergies.

The results for a typical instance of a reaching task from a test set are shown in Figure 9. The graph on the left shows the learning curves for learning with and without motor synergies. The horizontal axis gives the iteration number, and the vertical axis gives the value of the reaching-task cost function. Whereas learning without synergies was slow and never achieved good performance, learning with synergies was rapid and achieved excellent performance. Indeed, learning with synergies achieved roughly the same cost as the iLQR. The graph on the right shows the movements learned with and without synergies in Cartesian coordinates and the movement calculated by the iLQR. The movement learned without synergies never reached the target location, whereas the movement learned with synergies did. Overall, the results indicate that optimal synergies are useful motor primitives or building blocks in the sense that their use in linear combinations leads to rapid and accurate acquisition of new optimal motor actions.

4.5 Experiment 5: Motor Synergies When Forces Are Applied by Muscles. Whereas experiments 1 to 4 simulated a two-joint arm controlled directly by torques applied at the joints, experiments 5 to 7 simulated the
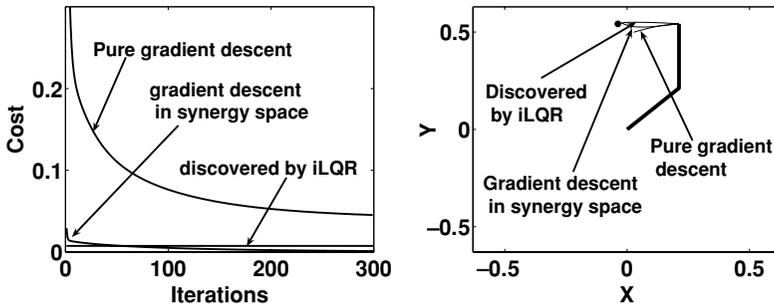
Figure 9: The graph on the left shows the learning curves for learning with and without motor synergies on a typical instance of a reaching task. The horizontal axis gives the iteration number, and the vertical axis gives the value of the reaching-task cost function. The graph on the right shows the movements learned with and without synergies in Cartesian space and the movement calculated by the iLQR.

same arm except forces were applied by muscles. We conducted experiments with muscles so that we could verify that the results reported above are also valid when simulating more complex and biologically realistic systems such as an arm controlled by muscles.

We used the muscle model developed by Todorov and Li (2005; see also Brown, Cheng, & Leob, 1999). In brief, this model uses six muscles that apply forces to a two-joint arm. The control signal is the neural input to the muscles. This input passes through a nonlinear low-pass filter to produce muscle activations. The tension of a muscle is a function of the muscle's current activation, length, and length velocity. The tension produces forces on the arm's links that, in turn, produce joint torques. Note that this system is significantly more complicated than the arm in which torques are applied directly at the joints. This system has a six-dimensional control space (neural input to each of six muscles), a 10-dimensional state space (six muscle activations and the angular position and velocity of each joint), muscle activations that might saturate, and dynamics with temporal delays (due to the low-pass filtering of neural input).

In experiments 1 to 4, nonnegative matrix factorization was applied to the optimal control signals to obtain motor synergies. In contrast, this factorization was not applied to the control signals—the neural input—in experiments 5 to 7; rather, it was applied to the optimal muscle activations. We found that the factorization procedure was significantly more robust when applied to the muscle activations due to the smoothness of their values (recall that the activations are low-pass accumulations of the neural inputs). Factorization of muscle activations was also conducted by d'Avella et al. (2003).
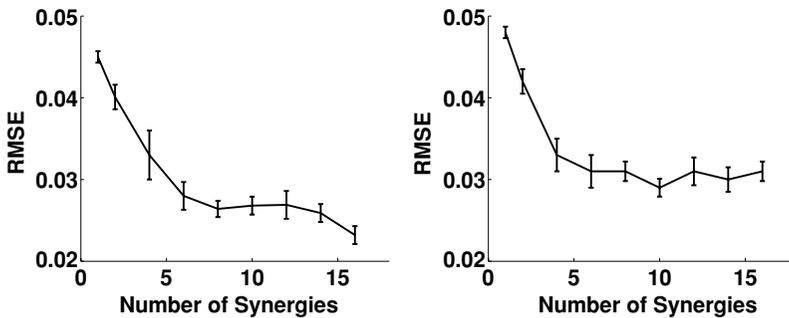
Figure 10: The graphs plot the root mean squared error (RMSE) between actual and reconstructed test items for reaching (left graph) and via-point (right graph) tasks when using the arm controlled by muscles as a function of the number of synergies used in the reconstructions. The error bars give the standard errors of the means.

Data for experiments 5 to 7 were collected in a similar manner as for experiments 1 to 4. We created 320 instances of the reaching task and the via-point task. Fivefold cross validation was used to create training and test sets of task instances. The iLQR was used to calculate the optimal sequence of neural inputs for each training instance (the cost functions for the reaching and via-point tasks given above were suitably modified by replacing the torque vector—the control input in experiments 1 to 4—with the neural input vector—the control input in experiments 5 to 7). Optimal muscle activations were created from the optimal neural inputs by low-pass filtering. Nonnegative matrix factorization was applied to the optimal muscle activations to generate optimal synergies. Based on these synergies, nonnegative matrix factorization was also used to perform task instances from test sets by computing optimal sums of scaled and time-shifted synergies.

Experiment 5 parallels experiment 1 in the sense that it evaluated whether optimal muscle activations can be expressed as a linear combination of a small number of time-shifted synergies. The results for the reaching and via-point tasks are shown in the left and right graphs of Figure 10, respectively. With both reaching and via-point tasks, the error is near its minimum when relatively few synergies were used. We conclude that synergies are useful motor primitives because optimal movements can be planned in a relatively low-dimensional space by summing a small number of scaled and time-shifted synergies.

4.6 Experiment 6: Task-Independent and Task-Dependent Synergies When Forces Are Applied by Muscles. Experiment 6 parallels experiment 2 in the sense that it evaluated whether optimal motor synergies are task independent or task dependent (d'Avella & Bizzi, 2005; Jing et al., 2004).
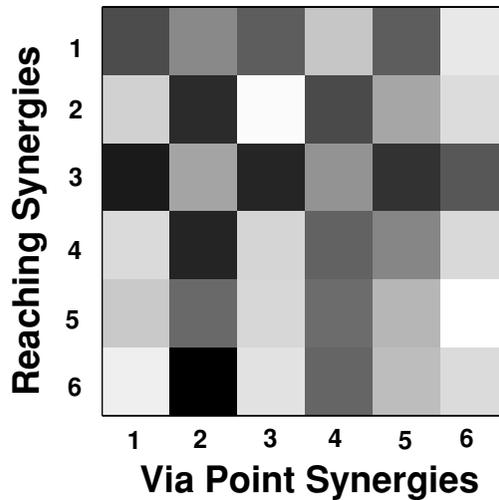
Figure 11: The similarity matrix for the six synergies obtained from the reaching task and the six synergies obtained from the via-point task using the arm controlled by muscles. The lightness of the square at row i and column j gives the cosine of the angle between the ith reaching-task synergy vector and the jth via-point task synergy vector—white is a value of 1, black is a value of 0, and intermediate gray-scale values represent intermediate values.

Figure 11 shows the similarity matrix for the six synergies obtained from the reaching task and the six synergies obtained from the via-point task. Some synergies, such as the second reaching-task synergy and third via-point task synergy are highly similar, indicating that these synergies are task independent. In contrast, other synergies, such as the third reaching-task synergy and the second via-point task synergy, are dissimilar from all other synergies, indicating that they are task dependent. A combination of task-independent and task-dependent synergies was also found in experiment 2, which used an arm controlled directly by torques. This result suggests that the combination of task-independent and task-dependent synergies found in biological organisms may be efficient for generating optimal motor actions from motor primitives.

4.7 Experiment 7: Visualizing Synergies When Forces Are Applied by Muscles. In experiment 7, we obtained synergies for the purpose of visualizing the movements induced by these synergies when forces are applied by muscles. Consequently, experiment 7 parallels experiment 3 and was conducted in an analogous manner.

Six synergies were obtained for the reaching task when the arm was controlled by forces applied by muscles. These synergies are illustrated in
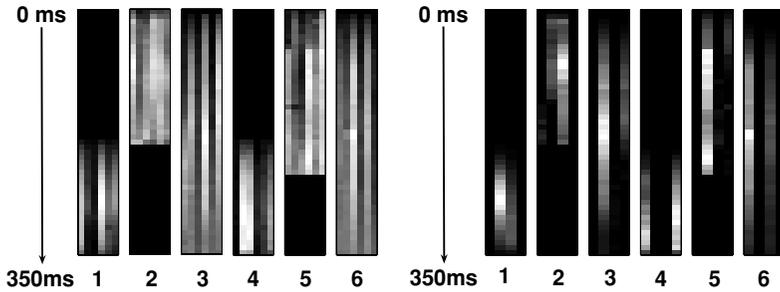
Figure 12: Six synergies were obtained for the reaching task when the arm was controlled by forces applied by muscles. The left graph shows the muscle activations. For each synergy, there are six columns corresponding to the activations of the six muscles. The right graph shows the torques generated by the synergies. For each synergy, the four columns indicate the anticlockwise and clockwise torques for joints 1 and 2, respectively.

Figure 12. The left graph shows the muscle activations. For each synergy, there are six columns corresponding to the activations of the six muscles.[2] Muscle activations were linearly scaled to the interval [0, 1]. White indicates an activation of 1, black indicates an activation of 0, and intermediate shades of gray indicate intermediate activation values. The right graph shows the torques generated by the synergies. For each synergy, the four columns indicate the anticlockwise and clockwise torques for joints 1 and 2, respectively.

   Figure 13 illustrates the movements induced by these synergies. The left graph shows the induced movements when the initial arm configuration was near the center of the work space. These movements tend to be relatively straight (though some are curved) and tend to cover a wide range of directions. The right graph shows the induced movements when the initial arm configuration was at a far edge of the work space. The movements tend to be relatively straight and are directed toward the center of the work space. For our purposes, a notable feature of these induced movements is that they closely resemble the movements induced by reaching-task synergies obtained when the arm was controlled by torques applied directly at the joints (see experiment 3 above). These data are consistent with the idea that synergies, when considered in "movement space," are more a reflection of task goals and constraints than of fine details of the underlying hardware.

---

[2] The muscles are (1) biceps long, brachialis, brachioradialis; (2) triceps lateral, anconeus; (3) deltoid anterior, coracobrachialis; (4) deltoid posterior; (5) biceps short; and (6) triceps long. See Li and Todorov (2004) for details.
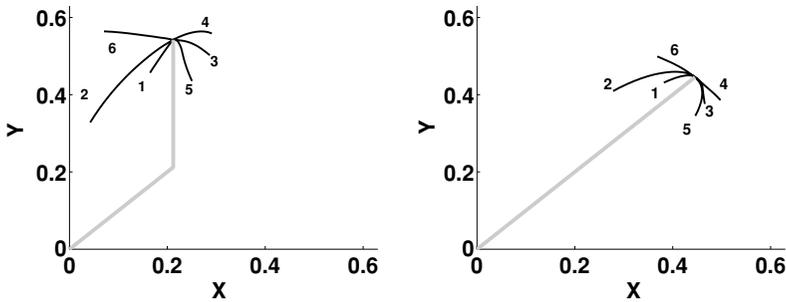
Figure 13: Movements induced by six synergies obtained for the reaching task when the arm was controlled by forces applied by muscles. The left and right graphs illustrate induced movements when the initial configuration of the arm was near the center of the work space or at a far edge of the work space, respectively. The number next to each movement indicates the synergy that was applied (using the same labels as Figure 12).

5 Discussion

In summary, this letter has considered the properties of synergies arising from a computational theory (in the sense of Marr, 1982) of optimal motor behavior. An actor's goals were formalized as cost functions, and the optimal control signals minimizing the cost functions were calculated by the iLQR. Optimal synergies were derived from these optimal control signals using a variant of nonnegative matrix factorization. This was done for both reaching and via-point tasks and for a simulated two-joint arm controlled by torques applied at the joints as well as an arm in which forces were applied by muscles. In brief, studies of the motor synergies revealed several interesting findings: (1) optimal motor actions can be generated by summing a small number of scaled and time-shifted motor synergies; (2) some optimal synergies are task independent, whereas other synergies are task dependent; (3) optimal motor actions can be rapidly acquired by learning new linear combinations of optimal motor synergies; and (4) synergies with similar properties arise regardless if one uses an arm controlled by torques applied at the joints or an arm controlled by muscles.

Future work will need to address shortcomings of our experiments. Our findings were obtained using simple motor tasks and a simple two-joint arm. We used reaching and via-point tasks because these are commonly performed movements and are frequently studied in the literature. We used a two-joint arm because it is computationally tractable. We conjecture that our basic results will still be found even with more complex tasks. This hypothesis is based on the fact that many complex movements can be regarded as combinations of simpler reaching and via-point movements.

We also conjecture that our results will still be found with more complex arms. This hypothesis is based on the fact that we obtained similar results regardless of whether we used a simple arm—a two-joint arm controlled by torques applied at the joints—or a more complex arm—a two-joint arm controlled by forces applied by muscles. Computationally, an obstacle to using more complex tasks and arms is the need to calculate optimal control signals. Using current computer technology, the calculation of optimal controls for nonlinear systems with many degrees of freedom is typically not possible.

Our findings were also obtained using specific mathematical techniques, such as the iLQR optimization method and the nonnegative matrix factorization method. We believe that our choices of mathematical techniques were reasonable. Again, this is an area in which important computational issues will need to be addressed before future studies can consider more complex motor tasks and arms. In particular, there is a need to develop improved dimensionality-reduction techniques for obtaining synergies. For example, the nonnegative matrix factorization method, like other methods, cannot be applied when movements have widely different durations and, thus, control signals have widely different dimensions. Future work will need to address this and many other unsolved problems.

Acknowledgments

References

Bernstein, N. (1967). *The coordination and regulation of movements*. London: Pergamon.

Brown, I. E., Cheng, E. J., & Leob, G. E. (1999). Measured and modeled properties of mammalian skeletal muscle; II: The effects of stimulus frequency on forcelength and force-velocity relationships. *Journal of Muscle Research and Cell Motility*, *20*, 627–643.

d'Avella, A., & Bizzi, E. (2005). Shared and specific muscle synergies in natural motor behaviors. *Proceedings of the National Academy of Sciences USA*, *102*, 3076–3081.

d'Avella, A., Saltiel, P., & Bizzi, E. (2003). Combinations of muscle synergies in the construction of a natural motor behavior. *Nature Neuroscience*, *6*, 300–308.

Flash, T., & Hogan, N. (1985). The coordination of arm movements: An experimentally confirmed mathematical model. *Journal of Neuroscience*, *5*, 1688–1703.

Harris, C. M., & Wolpert, D. M. (1998). Signal-dependent noise determines motor planning. *Nature*, *394*, 780–784.

Hollerbach, J. M., & Flash, T. (1982). Dynamic interactions between limb segments during planar arm movement. *Biological Cybernetics*, *44*, 67–77.

Jing, J., Cropper, E. C., Hurwitz, I., & Weiss, K. R. (2004). The construction of movement with behavior-specific and behavior-independent modules. *Journal of Neuroscience*, *24*, 6315–6325.

Jordan, M. I., & Rosenbaum, D. A. (1989). Action. In M. I. Posner (Ed.), *Foundations of cognitive science*. Cambridge, MA: MIT Press.

Li, W., & Todorov, E. (2004). Iterative linear-quadratic regulator design for nonlinear biological movement systems. In *Proceedings of the First International Conference on Informatics in Control, Automation, and Robotics* (pp. 222–229).

Marr, D. (1982). *Vision*. New York: Freeman.

Mussa-Ivaldi, F. A., Giszter, S. F., & Bizzi, E. (1994). Linear combination of primitives in vertebrate motor control. *Proceedings of the National Academy of Sciences USA*, *91*, 7534–7538.

Rosenbaum, D. A. (1991). *Human motor control*. San Diego: Academic Press.

Sanger, T. D. (1995). Optimal movement primitives. In G. Tesauro, D. S. Touretzky, & T. K. Leen (Eds.), *Advances in neural information processing systems, 7*. Cambridge, MA: MIT Press.

Sutton, R. S., McAllester, D., Singh, S., & Mansour, Y. (2000). Policy gradient methods for reinforcement learning with function approximation. In S. A. Solla, T. K. Leen, & K.-R. Müller (Eds.), *Advances in neural information processing systems, 12*. Cambridge, MA: MIT Press.

Thoroughman, K. A., & Shadmehr, R. (2000). Learning of action through adaptive combination of motor primitives. *Nature*, *407*, 742–747.

Todorov, E. (2004). Optimality principles in sensorimotor control. *Nature Neuroscience*, *7*, 907–915.

Todorov, E., & Jordan, M. I. (2002). Optimal feedback control as a theory of motor coordination. *Nature Neuroscience*, *5*, 1226–1235.

Todorov, E., & Li, W. (2005). A generalized iterative LQG method for locally-optimal feedback control of constrained nonlinear stochastic systems. *Proceedings of the 2005 American Control Conference* (Vol. *1*, pp. 300–306).