Properties of Synergies Arising From a Theory of Optimal Motor Behavior

Manu Chhabra (mchhabra@cs.rochester.edu)
Department of Computer Science
University of Rochester
Rochester, NY 14627 USA

Robert A. Jacobs (robbie@bcs.rochester.edu)
Department of Brain & Cognitive Sciences
University of Rochester
Rochester, NY 14627 USA

Abstract

We consider the properties of motor components, also known as synergies, arising from a computational theory (in the sense of Marr, 1982) or theory of optimal motor behavior. We study this topic in the domain of optimal control of a two-joint arm for reaching and via-point tasks. Studies of the motor synergies revealed several interesting findings. First, optimal motor actions for a new reaching or via-point movement 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. Lastly, 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.

Keywords: action; motor control; robotics; learning

Introduction

Marr (1982) hypothesized that a complex information processing device could be analyzed at three levels. The top level, referred to as the computational theory, examines what the device does and why. The analysis at this level provides an explanation for why a device does what it does by studying the device's goals. Although there are 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 paper studies motor synergies that arise from computational theories (in the sense of Marr, 1982) of 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). 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 and, thus, the degrees of freedom are typically redundant (Jordan & Rosenbaum, 1989). For example, consider the problem of touching the tip of your nose. The location of your nose has 3 degrees of freedom (its x, y, and z position in Cartesian coordinates), but the joints of your arm have 7 degrees of freedom (the shoulder has 3 degrees of freedom, and the elbow and wrist each have 2). As a result, there are several 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 synergy might be a coupling of the motions of your shoulder and elbow. Synergies provide two types of benefits to motor systems. First, synergies ameliorate the problem of redundancy; e.g., synergies 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. Importantly, they provide basic units of motor behavior that can be linearly combined to form more complex units of behavior.

This paper overviews the properties of synergies arising from a theory of optimal motor behavior. A more detailed description of this work, and additional results, can be found in Chhabra & Jacobs (2006). The motor behavior that we studied is the control of a nonlinear dynamical system, namely a simulated two-joint arm that resembles a human arm. We have created an optimal controller for this arm that formalizes goals as mathematical constraints, and searches for control signals that optimize the constraints. This has been done both for reaching tasks (move an end-effector from one point to another) and for via-point tasks (move from one point

to another while passing through an intermediate point). This has also been done both for an arm controlled directly by torques applied at the joints, and for an arm in which forces are applied by muscles (though this paper only reports the results for the arm controlled directly by torques applied at the joints). We then derived synergies from the optimal control signals using an extension to non-negative matrix factorization (d'Avella, Saltiel, & Bizzi, 2003), and studied the properties of these synergies.

Computing the optimal control signals

We simulated a two-joint arm that can be characterized as a second-order nonlinear dynamical system:

$$\mathcal{M}(\theta)\ddot{\theta} + \mathcal{C}(\theta, \dot{\theta}) + \mathcal{B}\dot{\theta} = \tau \quad (1)$$

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 & Todorov (2004).

We studied both a reaching and a via-point task. In the reaching task, the arm must be controlled so that its end-effector moves from a start location to a target location. The via-point task is identical except that there is an additional requirement that the end-effector 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 the optimal control signal is the signal that minimizes this function.

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

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

where k_1 and k_2 are constants (we used the same values as Todorov & Li, 2005, namely $k_1 = 0.001$ 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 & Todorov (2004; see also Todorov & Li, 2005).

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

$$J(\tau(t)) = \frac{1}{2} \parallel e(T) - e^* \parallel^2 + \frac{1}{2} \parallel e(T/2) - e_v^* \parallel^2 + k_1 \parallel \dot{e}(T) \parallel^2 + \frac{k_2}{2} \int_0^T \tau(t)^T \tau(t) dt \quad (3)$$

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 & Todorov (2004; see also Todorov & Li, 2005). In brief, the iLQR linearly approximates the system's dynamics and quadratically approximates the cost function at each moment in time, and then solves the approximate control problem. We have found that the iLQR works well on both reaching and via-point tasks when using the two-joint arm.

Obtaining optimal synergies

As discussed above, motor synergies are dependencies among dimensions of a motor system. They are useful because they ameliorate the problem of redundancy, and 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 non-negative matrix factorization developed by d'Avella, Saltiel, & Bizzi (2003). This algorithm requires two inputs. One input is the number of synergies, denoted N. The other input 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. 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:

$$\mathbf{m}(t) = \sum_{i=1}^{N} c_i \mathbf{w}_i (t - t_i) \quad (4)$$

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 *non-negative* control signals (each element of a control vector must be non-negative). 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 anti-clockwise and clockwise torques for the first joint, and the last two elements provide the same information for the second joint.

Simulation Results

This section reports the results of four experiments. All experiments used the same collection of reaching and via-point tasks. Three hundred twenty instances of each task were created 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 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.

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 five-fold 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, non-negative 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. Non-negative 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 1, 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 6-7 synergies were used. This is an important result 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.

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. For example, d'Avella & Bizzi (2005) recorded electromyographic activity from 13 muscles of the hind limb 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 2 shows the similarity matrix when 6 synergies were obtained for the reaching task and 6 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 i^{th} reaching-task synergy vector and the j^{th} via-point task synergy vector—white is a value of one, black is a value of zero, and intermediate gray-scale values represent intermediate values. Some synergies, such as the 4^{th} reaching-task synergy and the 2^{nd} via-point task synergy are highly similar, indicating that these synergies are task-independent. In contrast, other synergies, such as the 2^{nd} reaching-task synergy or the 3^{rd} 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 an efficient combination for generating optimal motor actions from motor primitives.

Experiment 3: Visualization of 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 or 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.

Figure 3 illustrates movements based on the six synergies obtained for the reaching task. 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, and the black lines show the movements of the end-effector. The induced movements tend to be relatively straight (though some are curved), and tend to cover a wide range

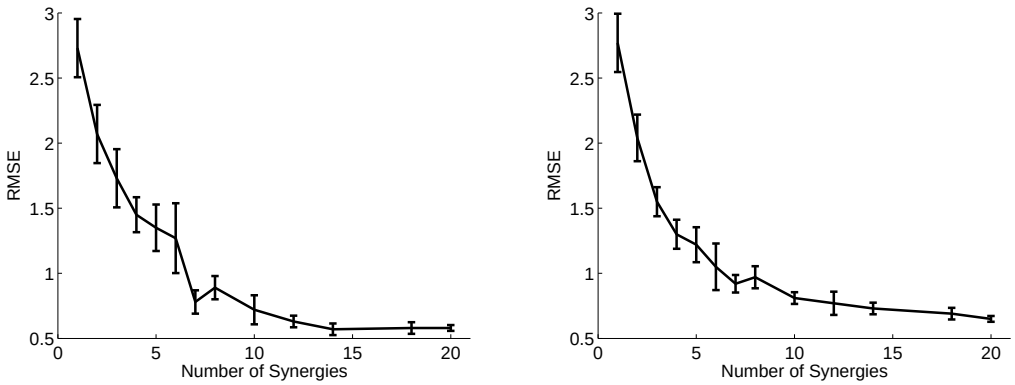

Figure 1: The graphs plot the root mean squared error (RMSE) between actual and reconstructed test items for reaching (left graph) and via-point (right graph) tasks as a function of the number of synergies used in the reconstructions. The error bars give the standard errors of the means.

of directions. The right graph of Figure 3 shows the induced movements when the initial arm configuration was at a far edge of the workspace. Again, the movements tend to be relatively straight. As should be expected, movements in this case are directed toward the center of the workspace. Figure 3 demonstrates that synergies tend to broadly cover all possible directions of motion.

Figure 4 illustrates movements based on the six synergies from the via-point task. 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 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 which is proceeded and followed by regions of relatively straight motion.

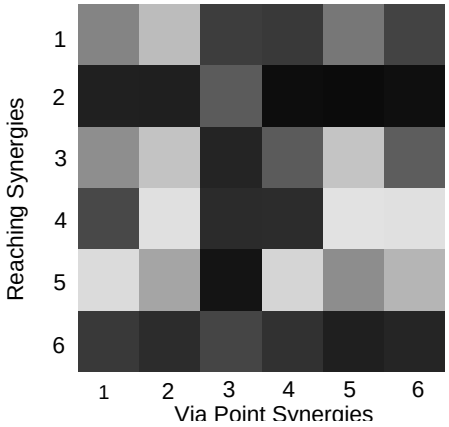

Figure 2: Similarity matrix when 6 synergies were obtained from the reaching task and from the via-point task. See text for details.

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 endpoint 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, five-fold 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 non-negative 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 (Equation 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 non-negative. When learning without synergies, we numerically computed the derivatives of the reaching-task cost function with respect to the torques at each joint and at each time step, and per-

formed 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 5. 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 that achieved by the movement calculated by 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. 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.

Summary and Conclusions

In summary, this paper has considered the properties of synergies arising from a computational theory (in the sense of Marr, 1982) or theory of optimal motor behavior. Studies of the motor synergies revealed 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. 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. Lastly, 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.

Future work will need to address shortcomings of our experiments. Our findings were obtained with simple motor tasks and a two-joint arm. We conjecture that our basic results will still be found with more complex tasks (note that many complex movements can be regarded as combinations of simpler reaching and via-point movements) and more complex arms (we obtained similar results with a two-joint arm controlled by forces applied by muscles; see Chhabra & Jacobs, 2006). Computationally, an obstacle to using more complex tasks and arms is the fact that the calculation of optimal controls for nonlinear systems with many degrees of freedom is typically not possible with current computer technology.

Our findings were also obtained using specific mathematical techniques, such as the iLQR optimization method and the non-negative 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 non-negative 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

We thank E. Todorov for help with the iLQR optimal control algorithm. This work was supported by NIH research grant R01-EY13149.

References

Bernstein, N. (1967). *The Coordination and Regulation of Movements*. London: Pergamon.

Chhabra, M. & Jacobs, R. A. (2006). Properties of synergies arising from a theory of optimal motor behavior. *Neural Computation*, in press.

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.

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.

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. *Proceedings of the First International Conference on Informatics in Control, Automation, and Robotics*, pp. 222-229.

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

Rosenbaum, D. A. (1991). *Human Motor Control*. San Diego: Academic 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, and K.-R. Müller (Eds.), *Advances in Neural Information Processing Systems 12*. Cambridge, MA: MIT Press.

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*, 1, 300-306.

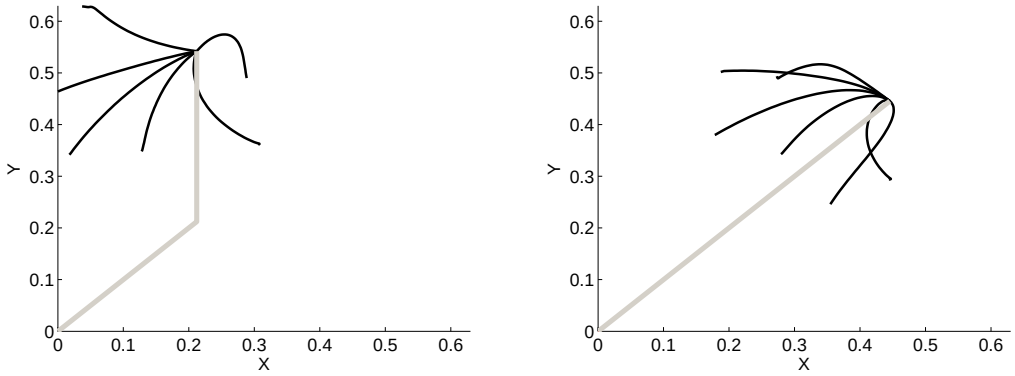

Figure 3: Movements induced by six synergies obtained for the reaching task. The left and right graphs illustrate induced movements when the initial configuration of the arm was near the center of the workspace or at a far edge of the workspace, respectively.

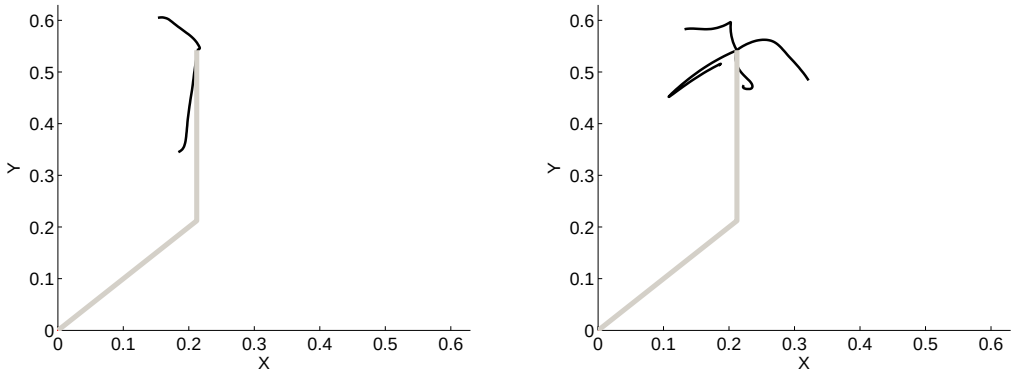

Figure 4: Movements induced by synergies obtained from the via-point task. The left and right graphs illustrate movements induced by task-independent and task-dependent synergies, respectively.

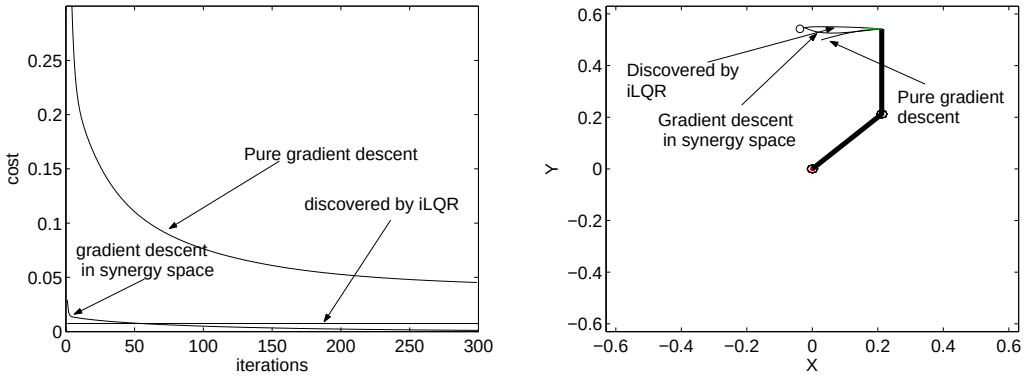

Figure 5: The graph on the left shows the learning curves (value of the reaching-task cost function as a function of iteration number) for learning with and without motor synergies on a typical instance of a reaching task. The graph on the right shows the movements learned with and without synergies in Cartesian space, and the movement calculated by the iLQR.