Published in final edited form as:

Neural Comput. 2014 October ; 26(10): 2163–2193. doi:10.1162/NECO_a_00639.

Hierarchical Control Using Networks Trained with Higher-Level Forward Models

Greg Wayne¹ and L.F. Abbott^{1,2}

¹Department of Neuroscience, Columbia University College of Physicians and Surgeons New York, NY 10032-2695 USA

²Department of Physiology and Cellular Biophysics Columbia University College of Physicians and Surgeons New York, NY 10032-2695 USA

Abstract

We propose and develop a hierarchical approach to network control of complex tasks. In this approach, a low-level controller directs the activity of a “plant,” the system that performs the task. However, the low-level controller may only be able to solve fairly simple problems involving the plant. To accomplish more complex tasks, we introduce a higher-level controller that controls the lower-level controller. We use this system to direct an articulated truck to a specified location through an environment filled with static or moving obstacles. The final system consists of networks that have memorized associations between the sensory data they receive and the commands they issue. These networks are trained on a set of optimal associations that are generated by minimizing cost functions. Cost function minimization requires predicting the consequences of sequences of commands, which is achieved by constructing forward models, including a model of the lower-level controller. The forward models and cost minimization are only used during training, allowing the trained networks to respond rapidly. In general, the hierarchical approach can be extended to larger numbers of levels, dividing complex tasks into more manageable sub-tasks. The optimization procedure and the construction of the forward models and controllers can be performed in similar ways at each level of the hierarchy, which allows the system to be modified to perform other tasks, or to be extended for more complex tasks without retraining lower-levels.

Keywords

Optimal Feedback Control; Internal Models; Neural Networks; Bi-Level Optimization

Introduction

A common strategy used by humans and machines for performing complex, temporally extended tasks is to divide them into sub-tasks that are more easily and rapidly accomplished. In some cases, the sub-tasks themselves may be quite difficult and time-consuming, making it necessary to further divide them into sub-sub-tasks. We are interested in mimicking this strategy to create hierarchical control systems. The top-level controller in such a hierarchy receives an external command that specifies the overarching task objective, whereas the bottom-level controller issues commands that actually generate actions. At each

level, a controller receives a command from the level immediately above it describing the goal it is to achieve and issues a command to the controller immediately below it describing what that controller is supposed to do. In this approach, optimization of a global cost function is abandoned in favor of a novel, more practical, but approximately equivalent approach of optimizing cost functions at each level of the hierarchy, with each level propagating cost-related information to the level below it.

As an example of this hierarchical approach, we solve a problem that requires two levels of control, what we call lower-level and a higher-level controllers. The basic problem is to drive a simulated articulated semi-truck backward to a specified location that we call the final target location (the truck is driven backward because this is harder than driving forward). The backward velocity of the truck is held constant, so the single variable that has to be controlled is the angle of the truck's wheels. This problem was first posed and solved by Nguyen and Widrow (Nguyen and Widrow, 1989), and their work is an early example of the successful solution of a nonlinear control problem by a neural network. We make this problem considerably harder by moving the final target location quite far away from the truck and, inspired by the swimmer of Tassa, Erez, and Todorov (Tassa et al., 2011), distributing a number of obstacles across the environment. Although the lower-level controller can drive to a nearby location when no obstacles are in the way, it cannot solve this more difficult task. Thus, we introduce a higher-level controller that feeds a series of unobstructed, closer locations that we call sub-targets to the lower-level controller that generates the wheel-angle commands. The job of the higher-level controller is to generate a sequence of sub-targets that lead the truck to its ultimate goal, the final target location, without hitting any obstacles. Thus, we divide the problem into lower-level control of the truck and higher-level navigation.

The controllers at both levels of the hierarchy we construct are neural networks. The specific form of these networks is not unique and is unlikely to generalize to other tasks, so we discuss their details primarily in the Methods section. We focus, instead, within the text, on general principles of their operation and construction. The tasks we consider are dynamic and ongoing, so commands must be computed by the network controllers at each simulation time step. To realize the speed for this computational intensive requirement, the network controllers are constructed to implement complex lookup tables. Each controller receives input describing the goal it is to achieve and, in addition, "sensory" input providing information about the environment relevant to achieving this goal. Its output is the command specifying the goal for the controller one level down in the hierarchy.

The network controllers are trained to implement the appropriate lookup table by back-propagation on the basis of optimal training data. The training data consist of input-output combinations computed to optimize a cost function defined for each hierarchical level. At each level, optimization is achieved with the aid of an additional neural network that implements a forward model of the controller being trained. The forward model is only used for optimization during learning; the fully-trained model consists only of the controller networks. For the higher-level controller, the training procedure involves what is effectively a control-theory optimization in which the "plant" being controlled is actually the lower-level controller. This approach thus extends ideas about forward models and optimization

from the problem of controlling a plant to that of controlling a controller. Once the optimal output commands are determined for a large set of input commands and sensory inputs, these are used as training data for the controller, which effectively “memorizes” them.

We begin by describing how the hierarchical approach, consisting of lower- and higher-level controllers, operates after both of these networks have been fully trained. We do this sequentially, first showing the lower-level controller operating the truck when its sub-target data are generated externally (by us), rather than by the higher-level controller. We then discuss how the higher-level controller generates a sequence of sub-target locations to navigate through the environment. To allow the higher-level controller to detect and locate obstacles, we introduce a sensory grid system. We present and analyze the complete hierarchical system with the two controllers working together and compare its operation with that of a more conventional optimal controller. After we have shown the system in operation and analyzed its performance, we present the procedures used for training, including the cost functions and forward models used for this purpose at each level.

Results

All of the networks we consider run in discrete time steps, and we use this step as our unit of time, making all times integers. Distance is measured in units such that the length of the truck cab is 6, the trailer is 14, and both have a width of 6. In these units, the backward speed of the truck is 0.2. The final target for the truck and the obstacles it must avoid have a radius of 20. Distances from the initial position of the truck to the final target are typically in the range of 100 to 600.

Driving the Truck

Our hierarchical model for driving the truck (Figure 1) starts with a lower-level controller that sends out a sequence of commands $u(t)$ that determines the angle of the wheels of the truck. This controller is provided with “proprioceptive” sensory information, namely the cosine and sine of the angle between the cab and trailer of the truck, $[\cos(\theta_{rel}), \sin(\theta_{rel})]$, and a sub-target location toward which it is supposed to direct the truck (to ensure continuity and promote smoothness, we process all angles by taking their cosines and sines). The target information is provided as a distance from the truck to the sub-target, d_{st}/L ($L=100$ is a scale factor), and the cosine and sine of the angle from the truck to the sub-target, $[\cos(\theta_{st}), \sin(\theta_{st})]$. This sub-target information is provided by a higher-level controller that receives the same proprioceptive input from the truck as the lower-level controller but also receives sensory information about obstacles in the environment (to be described later). In addition, the higher-level controller is provided with external information about the distance from the truck to the final target location and also the cosine and sine of the angle from the truck to this location, $[\log(1+d_{ft}/L), \cos(\theta_{ft}), \sin(\theta_{ft})]$. The task of the higher-level controller is to provide a sequence of sub-targets to the lower-level controller that lead it safely past a set of obstacles to the final target location. Note that the higher-level controller receives the logarithm of the distance to the final target, $\log(1 + d_{ft}/L)$, rather than d_{ft}/L itself. This allows for operation over a larger range of distances without saturating the network activities. The

logarithm is not needed for d_{st}/L because the distance to the sub-target is maintained within a constrained range by the higher-level controller.

The Lower-Level Controller—The job of the lower-level controller is to generate a sequence of wheel angles, $u(t)$, given the proprioceptive data, $[\cos(\theta_{rel}(t)), \sin(\theta_{rel}(t))]$, and a sub-target location specified by $[d_{st}(t), \cos(\theta_{st}(t)), \sin(\theta_{st}(t))]$ (Figure 1). The proprioceptive information is needed by the controller not only to move the truck in the right direction but also to avoid jackknifing. The lower-level controller is a 3-layer basis-function network with the 5 inputs specified above, 100 Gaussian-tuned units in a hidden-layer, and 1 output unit that reports u as a linear function of its input from the hidden layer (Methods). Figure 2A shows an example in which the sub-target location is held fixed and the lower-level controller directs the truck along the backward path indicated by the curved line.

At this point, we are showing the lower-level controller working autonomously with the sub-targets specified by us, but when the truck is directed by the higher-level controller, it will be given a time-dependent sequence of sub-targets. To test whether it can deal with sequential sub-targets, we switched the sub-target we provide every time the truck got close to it, using a rather fanciful sequence of sub-targets (Figure 2B). This indicates that the lower-level control is up to the job of following the directions that will be provided by the higher-level controller.

The Higher-Level Controller—The higher-level controller is a five-layer, feedforward network with a bottleneck architecture. It has 205 inputs (3 specifying the final target location, 2 the angle between the cab and trailer of the truck, 199 describing the state of the sensory grid described below, and a bias input; Figure 1), hidden layers consisting of 30, 20 and 30 units, and three command outputs providing the sub-target information for the lower-level network (Methods). The bottleneck layer with 20 units ensures that the network responds only to gross features in the input that reliably predict the desired higher-level command. When we initially trained the higher-level controller without any form of bottleneck, it did not generalize well to novel situations.

In the absence of any obstacles, the job of the higher-level controller is to provide a sequence of sub-targets to the lower-level controller that lead it to the location of the final target, which is specified by the variables $[d_{ft}, \cos(\theta_{ft}), \sin(\theta_{ft})]$ that the higher-level controller receives as external input. The higher-level controller also receives a copy of the proprioceptive input provided to the lower-level controller (Figure 1). Figure 3 shows a trajectory generated by the higher-level controller and the motion of the truck as directed by the lower-level controller, leading to a target just beyond the bottom-left corner of the plot. Note the sequence of target locations that lead the truck along the desired path. Although this example shows that the higher-level controller is operating as it should and that the lower-level controller can follow its lead, this task is quite simple and could be handled by the lower-level controller alone. To make the task more complex so that it requires hierarchical control, we introduced obstacles into the environment.

The obstacles are discs with the same radius as the final target scattered randomly across the arena (Figure 4B). These are soft obstacles that do not limit the movement of the truck, but

during training we penalize commands of the higher-level controller that cause the truck to pass too close to them (see below). Making this environmental change requires us to introduce a sensory system that provides the higher-level controller with information about the locations of the obstacles. Just as the final and sub-target locations are provided in “truck-centric” (egocentric) coordinates (distances and angles relative to the truck), we construct this sensory system in a truck-centric manner (Figure 4A).

Specifically, we construct a hexagonal grid of points around the truck (Figure 4). The grid is a lattice of equilateral triangles with sides of length 20 units. One grid point lies at the back of the trailer, and the most distant grid points are 150 units away from this point. In total, there are 199 grid points. These points move with the truck and align with the longitudinal axis of the trailer (Figure 4A). If a grid point lies inside an obstacle, we consider it to be activated; otherwise, it is inactive. The state of the full grid is specified by an 199-component binary vector g with component i specifying whether grid point i is active ($g_i = 1$) or inactive ($g_i = 0$). Neither topological closeness nor Euclidean distance information is explicit in this vector representation. The grid vector is provided as additional input to the higher-level controller (Figure 1).

Operation of the Full System—We now show how the full system operates when the higher-level controller provides the lower-level controller with sub-targets as they drive the truck together through a field of obstacles to the final target (Figure 5A). Figure 5B shows a number of guided trajectories through an obstacle-filled arena.

To quantify the performance of the system, we executed 100 trials in 100 different environments with 20 obstacles. The hierarchical controller avoids the obstacles on each trial; the minimum distance to an obstacle never decreases below one obstacle radius (20 units; Figure 6A). As the number of environmental obstacles is increased (Figure 6B), the probability of obstacle collision grows slowly. This occurs even though the controllers were trained with only 20 obstacles in the environment. The control system directs the truck to the final target along short paths (Figure 6C) that are comparable in length to the straight line distance between the initial position and the nearest edge of the final target, with deviations when the truck must execute turning maneuvers or circumnavigate obstacles.

The trained higher-level controller continuously generates sub-targets based on the sensory information it receives (Video 1). Because all the contingencies are memorized, it needs very little time to compute these plans. Thus, the higher-level controller should be able to respond quickly to changes in the environment. To illustrate this, we tested the system with obstacles that moved around, even though it was trained with stationary obstacles. The obstacle motions were generated as random walks (Video 2). The probability of collision grows slowly with increasing diffusion constant of the random walk (Figure 6D). At high rates of diffusion, the obstacles move significantly farther than the truck for small numbers of time steps. For example, when the diffusion coefficient is 1 (units [L^2/T]), the obstacles typically diffuse (but can diffuse farther than) the width of the trailer within 9 time steps. It takes the truck 30 time steps to travel the same distance.

Comparison to Optimal Control—The hierarchical system cannot be described as optimal with respect to a single cost criterion for several reasons. First, the responses of the networks are memorized and therefore only approximate the responses of optimization computations. Second, the networks do not observe all state variables in the environment exactly; they observe mappings of those state variables through a sensory system. Third, the temporal horizon, or the amount of planning foresight granted during training, is less than the total duration of typical trials. Fourth, and most important, the controller networks are trained on separate cost functions and then coupled. We therefore compare the results obtained from the hierarchical network with those from an optimal control calculation. It should be stressed that the optimal control approach is not a practical way to solve the truck problem we considered because it is far too slow without speed-tweaks and, as we will see, it fails to find reasonable paths a fair fraction of the time. Nevertheless, comparing paths produced by the hierarchical controller, at an expense of tens of milliseconds per path, with paths constructed by an optimization program over many seconds per path provides a way to judge the success of the hierarchical. To compare our results with those of an optimal control calculation, we generated 150 random environments and computed solution trajectories from identical initial conditions using either the network hierarchy or an optimal control calculation computed by differential dynamic programming, a commonly used algorithm for hard optimal control calculations (Methods). We compared the two solutions using the non-constraint portions of the cost function used in the optimal control calculation (Methods). Note that this means that we are judging the hierarchical model using a cost function that was not used in its construction.

Figure 7A shows a comparison of paths generated by the hierarchical network and the optimal control algorithm. Many of the paths are quite similar. Other paths are clearly equivalent although they differ by going around opposite sides of an obstacle. In some cases, the optimal control algorithm has clearly failed to find a good solution, and it produces paths with loops in them. These cases appear as a long negative tail on the distribution of cost-function differences shown in Figure 7B. Importantly, there are very few cases in which the cost of the solution provided by the hierarchical network is significantly worse than that of the optimal control path (the small positive tail in Figure 7B). Our conclusions are that the majority of paths constructed by the hierarchical network come close to being optimal with respect to the cost function used for this comparison and that the hierarchical model produces fewer and less disastrously bad paths when it fails to approximate optimality.

The Training Procedure

The controllers at each level of the hierarchy work because they have been trained to generate commands (sub-targets or wheel angles for the truck) that are approximately optimal for the input they are receiving at a given time. Recall that this input consists of the sub-target received from the upstream controller and whatever sensory information is provided. The coupling between the two levels of the hierarchy requires an extension of the methodology of model-based optimal control. The key modification has two connected parts. Input from the higher-level act not only as a command, but also serves to specify the cost function that the lower-level controller is trained to minimize. As a result, the lower-level controller is actually trained to minimize a family of cost functions parameterized by

the value of the higher-level command. The higher-level control problem is to choose a sequence of higher-level commands to send to the lower-level controller that will minimize a higher-level cost function. To find the appropriate higher-level command, a higher-level forward model is trained to predict the the feedback that will result from propagating commands. Because the higher-level optimization acts on the higher-level forward model, which is trained only after the lower-level controller has been trained, the optimization problems at both levels are decoupled. This is a virtue because we do not need to execute nested optimizations to train the higher-level controller.

Tasks like the one we consider are difficult because a significant amount of time may elapse before the cost associated with a particular command strategy can be determined. In our lower-level example, it takes a while for the truck to move far enough to reveal that the wheels are not at a good angle. Typically, as tasks get more complex, this delay gets longer. For example, it takes longer to evaluate whether a sub-target issued by the higher-level controller is going to get the truck closer to the final target without leading it into an obstacle. One consequence of this delay is that we cannot assess the cost associated with a single command, we must evaluate the cost of a sequence of commands. This, in turn, requires that we predict the consequences of issuing a command, which we do at each level using a forward model. Once we have minimized the cost function by choosing an optimal sequence of commands, we train a controller network at each level to memorize the optimal commands given particular inputs.

To deal with the hierarchy of timescales that are associated with a hierarchy of control levels, we introduce two time scales per level. The first is associated with the temporal scale over which a process needs to be controlled. There is no point in issuing commands that change more rapidly than the dynamics of the object being controlled. In the truck example, wiggling the wheels back and forth rapidly is not an intelligent way to drive the truck, and swinging the sub-target around wildly is not a good way to guide the lower-level controller. At level l of the hierarchy, we call this dynamic timescale T_l . In general, the choice of T_l is governed by the dynamics of the system being controlled by level l . For the truck problem, we take $T_1 = 6$ and $T_2 = 72$ time steps.

In optimal control problems, the cost of a command is only considered in the context of an entire sequence of commands. What is good to do now depends on what will be done later. The second time scale is therefore the length of the sequence of commands used to compute the cost of a trajectory. At level l , we denote this number by K_l . In other words, it requires K_l commands, spaced apart by T_l time steps, to determine the cost of a particular command strategy. In the case of the truck, we set $K_1 = 15$ and $K_2 = 10$. An explanation of these particular settings is provided in the Methods.

The Cost Functions—We denote a command given at time t by the level l controller by the vector $\mathbf{m}_l(t)$. For the lower-level controller in the truck example, $\mathbf{m}_1(t) = u(t)$, and for the upper-level controller $\mathbf{m}_2(t) = [d_{st}(t), \cos(\theta_{st}(t)), \sin(\theta_{st}(t))]$. We also define a vector $\mathbf{s}_l(t)$ that represents the sensory input to layer l at time t upon which the decision to issue the command sequence is based. For the case of the truck, $\mathbf{s}_1(t) = [\cos(\theta_{rel}), \sin(\theta_{rel}), d_{st}, \cos(\theta_{st}), \sin(\theta_{st})]$ and $\mathbf{s}_2 = [\cos(\theta_{rel}), \sin(\theta_{rel}), \mathbf{g}, \log(1 + d_{ft}/L), \cos(\theta_{ft}), \sin(\theta_{ft})]$ (Figure 1).

The cost of a sequence of commands $[\mathbf{m}_l(t), \mathbf{m}_l(t + T_l), \mathbf{m}_l(t + 2T_l), \dots, \mathbf{m}_l(t + K_l T_l)]$ takes the general form

$$\mathcal{J}_l = \sum_{k=0}^{K_l} \mathcal{L}_l(s_l(t + (k+1)T_l), \mathbf{m}_l(t + kT_l); \mathbf{m}_{l+1}(t)). \quad (1)$$

The functions $\mathcal{L}_l$ for the lower- ($l = 1$) and higher- ($l = 2$) level controllers are specified below. It is important to observe that the cost function at level l depends parametrically on the command from the level above, $\mathbf{m}_{l+1}(t)$. This command is provided at time t and, during training, is fixed until time $t + K_l T_l$.

We would like the truck to drive toward the target along a straight angle-of-attack, without articulating the link between the cab and trailer too much, and using minimal control effort. A cost function satisfying these requirements can be constructed from

$$\mathcal{L}_1 = \alpha_1 d_{st} + \beta_1 \theta_{st}^2 + \gamma_1 (|\theta_{rel}| - \theta_{max})^2 \Theta(|\theta_{rel}| - \theta_{max}) + \zeta_1 u^2. \quad (2)$$

The third term makes use of the Heaviside step function, which is 1 if $x \geq 0$ and 0 otherwise. We use the convention that all angles are in radians, and we center all angles around 0 so they fall into the range between $\pm\pi$. The parameters α_1 , β_1 , γ_1 , ζ_1 , and θ_{max} are given in Table 1 of the Methods.

The higher-level cost function is divided into three parts:

$\mathcal{L}_2 = \mathcal{L}_2^{\text{sensory}} + \mathcal{L}_2^{\text{command}} + \mathcal{L}_2^{\text{obstacle}}$. All the parameters in these cost functions are given in Table 1 of the Methods. The sensory cost contains a distance-dependent term, but we no longer need to penalize large cab-trailer angles because the lower-level controller takes care of this on its own, so

$$\mathcal{L}_2^{\text{sensory}} = \alpha_2 \log(1 + d_{ft}/L). \quad (3)$$

The higher-level motor command portion of the cost is given by

$$\begin{aligned} \mathcal{L}_2^{\text{command}} = & \beta_2 \left(\cos(\theta_{st})^2 + \sin(\theta_{st})^2 - 1 \right)^2 + \gamma_2 \left((d_{st} - d_{min})^2 \Theta(d_{min} - d_{st}) \right. \\ & \left. + (d_{st} - d_{max})^2 \Theta(d_{st} - d_{max}) \right) \end{aligned} \quad (4)$$

The first term in this equation may look strange because the sum of the squares of a cosine and a sine is always 1. However, the optimization procedure does not generate an angle θ_{st} and take its cosine and sine. Instead, it generates values for the cosine and sine directly without any constraint requiring that these obey the laws of trigonometry. As a result, this constraint needs to be included in the cost function. The distance-dependent terms in equation 4 penalize sub-target distances that are either too short or too long.

The final term in the higher-level cost function, $\mathcal{L}_2^{\text{obstacle}}$ penalizes truck positions that are too close to an obstacle. We are not concerned with the distance from the truck to every single obstacle; rather, we care primarily if the truck is too near a single obstacle, the closest one. Thus, we choose the smallest distance to an obstacle, $d_{min}^{\text{obstacle}}(t)$, and impose a

Gaussian penalty for proximity to this closest obstacle with standard deviation equal to the disc's radius, σ_{disc} . We also add a smaller, flatter penalty with a larger standard deviation, σ_{areola} , as a warning signal to prevent the truck from wandering nearby the obstacle. The resulting cost function is

$$\mathcal{L}_2^{\text{obstacle}} = \rho_2 \exp\left(-\frac{(d_{\min}^{\text{obstacle}})^2}{2\sigma_{\text{disc}}^2}\right) + v_2 \exp\left(-\frac{(d_{\min}^{\text{obstacle}})^2}{2\sigma_{\text{areola}}^2}\right). \quad (5)$$

Recall that the obstacles are detected by the sensory grid system show in Figure 4, and thus the distances to obstacles are not directly available. We solve this problem by introducing a network that evaluates the cost function 5 directly from the sensory grid information $\mathbf{g}(t)$. We call this network the “obstacle critic” because it serves the same role as critic networks in reinforcement learning (Widrow et al., 1973) (Sutton and Barto, 1998): predicting the cost of sensory data, ultimately to train another network. It has 199 grid inputs and one bias and a single output, representing the estimated cost of the sensor reading (Methods). We train this network to predict the obstacle cost from grid data by creating a large number of measurement scenarios and computing the true cost function.

The complete higher-level cost function is the sum of the costs given in equations 3, 4, and 5. The trajectories that minimize the complete higher-level cost function are goal-seeking and obstacle-avoiding.

The Forward Models—In equation 1, the cost function depends not only on the sequence of commands but also on the entire sequence of sensory consequences of those commands. In other words, $\mathcal{J}_l$ depends on $\mathbf{s}_l(t) \dots \mathbf{s}_l(t + K_l T_l)$ but only $\mathbf{s}_l(t)$ is provided. To predict the future sensory data resulting from the command sequence, we build a forward model at each level of the hierarchy. To do this, we sample the space $(\mathbf{s}_l(t)), \mathbf{m}_l(t)$ and record the resultant states $\mathbf{s}_l(t + T_l)$. We then build a network that predicts the resultant sensory data.

The forward model for level l predicts the future sensory data $\mathbf{s}_l(t + T_l)$ that result from passing the command $\mathbf{m}_l(t)$ to the level below it. We use the convention that the forward model is named after the level that issues the commands and receives the sensory data, not the level that follows those commands and causes the sensory data to change. Thus, the higher-level forward model is actually modeling the sensory consequences of sending a command to the lower-level controller, and the lower-level forward model is modeling the consequences of sending a command to the truck. Because we have a full kinematic description of the truck, we could use the truck itself, rather than a forward model of it, to optimize the lower-level cost function. Nevertheless, we use a forward model so that we can treat and discuss the lower and higher levels in a similar manner.

The lower-level forward model, which is a single network, is trained by randomly choosing a set of sensory data, $\mathbf{s}_1(t)$, consisting of the distance and cosine and sine of the angle to a target, and a command $\mathbf{m}_1(t)$, defining a wheel angle, within their allowed ranges. We then simulate the motion of the truck for a time T_1 and determine the sensory data $\mathbf{s}_1(t + T_1)$, indicating the articulation of the cab with respect to the trailer and where the truck lies in

relation to the sub-goal. We continue to gather data for the lower-level forward model by applying another command to the truck and taking a new sensory measurement after the delay. In this way, we generate several command sequences along a single trajectory to gather more data. To create a diversity of training cases for the forward model, we periodically terminate a trajectory and start a new trial from a random initial condition. On the basis of several thousand such measurements, we train the lower-level forward model network to predict the motion and articulation of the truck over the full range of initial conditions and motor commands.

The higher-level forward model is composed of three networks: proprioceptive, goal-related, and obstacle-related (Methods). Each of these networks receives the command $\mathbf{m}_2(t)$. The proprioceptive network also receives the proprioceptive information, $[\cos(\theta_{\text{rel}}(t)), \sin(\theta_{\text{rel}}(t))]$, and predicts $[\cos(\theta_{\text{rel}}(t + T_2)), \sin(\theta_{\text{rel}}(t + T_2))]$. The goal-related forward model receives these proprioceptive data and the goal-related information, $[\log(1 + d_{\text{ft}}/L), \cos(\theta_{\text{ft}}), \sin(\theta_{\text{ft}})]$, and predicts the goal-related variables at time $t + T_2$. The obstacle-related forward model is the most complicated. It receives the proprioceptive information and the obstacle grid data $\mathbf{g}(t)$. Unlike the other forward model networks we have described, which make deterministic predictions, the predictions of the goal-related forward model are probabilistic. This is necessary because the grid predictions are underdetermined. For example, at time t the grid input cannot provide any information about obstacles outside its range, but one of these may suddenly appear inside the grid at time $t + T_2$. The goal-related forward model cannot predict such an event with certainty. Therefore, we ask the obstacle-related network to predict the probability that each grid point will be occupied at time $t + T_2$, given a particular grid state and command issued at time t . This is, obviously, a number between 0 and 1, in contrast with the true value of $g_i(t + T_2)$, which would be either 0 or 1. We explain how this is done in the Methods.

When we use forward model networks to predict sensory information along a trajectory, $\mathbf{s}_l(t), \mathbf{s}_l(t+T_l), \mathbf{s}_l(t+2T_l), \dots$, we simply iterate, using the predicted sensory data time $t + kT_l$ to generate a new prediction at $t + (k + 1)T_l$. This allows us to compute the summed cost functions of equation 1 for both controller levels (Figure 8).

Computing Optimal Commands—Our procedure for generating optimal command sequences is schematized in Figure 8. We begin by initializing states of the environment and truck randomly (within allowed ranges) and measure $\mathbf{s}_l(t)$. We then choose an initial (or “nominal”) sequence of commands $[\mathbf{m}_l(t), \mathbf{m}_l(t + T_l), \dots, \mathbf{m}_l(t + K_l T_l)]$. For the lower-level controller, we choose $\mathbf{m}_1(t + kT_1) = u(t + kT_1) = 0$ for all k , indicating that pointing the wheels straight is our first guess for the optimal command. For the higher-level controller, we choose the initial command sequence to consist of identical commands placing a sub-target 50 length units directly behind the truck. We apply these commands sequentially to calculate an entire truck trajectory.

The command sequences are then optimized by using the Dynamic Optimization algorithm described in the Methods (also known as Pontryagin’s Minimum Principle (Stengel, 1994) or Backpropagation-through-Time (LeCun, 1988)). Briefly, we calculate the effect that a small change to each command will have on the total trajectory cost and determine the

gradient of the cost function with respect to the parameters defining the commands. This gradient information is used to change the commands according to a variable-metric update accomplished with the limited-memory Broyden-Fletcher-Goldfarb-Shanno (L-BFGS) method in “minFunc” (Schmidt, 2013). L-BFGS uses successive gradient computations to compute an approximation of the inverse Hessian of the cost function, which speeds up optimization considerably compared to steepest descent. We iterate the optimization until performance does not improve (more details are provided in the Methods).

The above procedure generates a single command sequence that is optimal for the initial sensory data $\mathbf{s}_I(t)$. However, we can use the sensory data after the first command has generated its sensory consequences (at time $t + T_I$), to optimize a new trajectory starting from the new sensory data. This process can obviously be repeated at subsequent times. We periodically terminate the trajectory and begin a new trial to obtain sufficient variety in the initial sensory data. When this process is completed, the pairings of initial sensory data and optimal commands constitute a training set for the controller being trained, which learns to associate each sensory measurement with the appropriate command.

Computing an optimal trajectory of K_1 steps at the lower-level takes on average 1.3 seconds on our computer (averaged over 100 such optimizations). Computing an optimal trajectory of K_2 steps at the higher-level takes on average 12.47 seconds (averaged over 10 such optimizations). Once the lower- and higher-level controllers have been trained, only 0.001 seconds is needed to run the two controllers in series. Clearly, memorization provides a tremendous advantage in speed over online optimization.

Training the Networks—Having drawn random initial sensory data $\mathbf{s}_I(0)$ and solved the optimal control problem defined by equation 1, subject to the dynamics of the forward model, we then train the controller network to predict the first motor command in the sequence from the sensory data, i.e., $\mathbf{m}_I(0)$ from $\mathbf{s}_I(0)$. We do not use the full sequence of optimal commands for training the controller; we only use the full sequence to evaluate the cost function. The rationale is that all the future commands beyond the first one in the sequence are based on sensory inputs predicted by the forward model. Training the controller to associate these predicted sensory states with their paired commands in the sequence introduces errors into the controller because the sensory predictions of the forward model are not entirely accurate.

The training procedure for each controller is straightforward. We apply a particular input to the controller network and use backpropagation to modify its parameters so as to minimize the squared difference between the output command given by the controller and the optimal command for that particular set of inputs. After a sufficient number of such trials (Methods), the controller network learns to produce the desired command in response to a particular input. Furthermore, if the network is properly designed it will generalize to novel inputs by smoothly interpolating among the trained examples.

Discussion

Biological Evidence for Hierarchical Control and Forward Models

Although our networks are certainly not models of specific biological neural systems, we cannot resist drawing various connections to physiology. Biological evidence for hierarchical control is diverse, although we lack a clear understanding of its architectural logic. Lashley (1930) and Bernstein (1967) independently advocated hierarchical theories of motor control after reasoning from psychophysical experiments and intuition. Lashley's *principle of motor equivalence* paralleled Bernstein's *degrees of freedom* problem: namely, there is more than one movement that will accomplish a task goal. It is, for example, possible to write with one's left or right hand. Raibert (1977) studied the problem of cursive writing while varying properties of the utilized effector. He wrote, "Able was I ere I saw Elba," using both hands, with an immobilized wrist, and even with his teeth, demonstrating a relative invariance of the basic form of the orthography, suggesting that musculo-skeletal control is only loosely coupled to the problem of goal-directed planning.

Anatomically, of course, projections from motor cortex to the spinal cord is hierarchical in that control of muscle contraction is indirect (Loeb et al., 1999). This was perhaps first grasped by Hughlings Jackson in 1889 (Jackson, 1889). Therefore, from the perspective of cortical control centers, the problem of movement corresponds to the problem of sending commands to the spinal cord in such a way that will satisfy higher-level intentions. In the octopus, large networks of ganglia, interposed between the central nervous system and the arm muscles, are themselves necessary and sufficient for the production of complicated movements; when the axial nerve cords of the peripheral nervous system are stimulated electrically, severed arms propagate a bend that results in qualitatively normal extension (Sumbre et al., 2001). Analogously, in zebra finch singing behavior, the nucleus HVc controls song timing through its projections to the motor region RA, which is ultimately responsible for the control of syllable vocalization by syringeal muscles (Albert and Margoliash, 1996).

A long-standing hypothesis is that the cerebellum implements a forward model that predicts the delayed sensory consequence of a motor command (Miall et al., 1993; Shadmehr et al., 2010). Intriguingly, a recent study suggests that neurons in Clarke's column in the spinal cord simultaneously receive input from cortical command centers and from proprioceptive sensory neurons, and the authors speculate that this joining of sensory and motor inputs could function as a predictor, akin to a forward model (Hantman and Jessell, 2010). Experimentalists have therefore already proposed multiple loci for forward models, functioning at different levels of the motor system. This connection suggests that motor behavior may be produced by hierarchies of controllers that are trained by forward models at each level of the hierarchy.

Relationship to Other Approaches

Research on hierarchical control is dispersed among several different fields with varying agendas and, in fact, conceptions of the word "hierarchy." Robotics has embraced hierarchical mechanisms whole-heartedly, and most current systems for complicated tasks

are now hierarchical. A typical mobile robot will segregate low-level control from planning and navigation, a functional division that we have borrowed. This is the case in Stanley, the self-driving car that won the DARPA Grand Challenge (Thrun et al., 2006), which was modeled on the “three-layer architectures” of Gat (1998). In these architectures, however, planning is entirely separated from lower-level motor control; in particular, higher levels do not model the result of commanding lower levels, so no notion of feasibility is available to the higher-level system. As a workaround, the navigational models in these systems are heuristically instructed to generate paths that are smooth and free from obstruction. We feel that endowing higher-level systems with explicit models of the results of propagating commands to lower levels can make robots more flexible and better able to cope with or exploit the particularities of their lower-level controller and plant dynamics.

Hierarchical reinforcement learning (Dietterich, 1998; Barto and Mahadevan, 2003; Sutton et al., 1999) considers a different notion of hierarchy, which we might term a “recursive hierarchy.” Whereas our notion of hierarchy is structural and embedded in a hierarchy of circuits, hierarchical reinforcement learning systems are representatives of the same form of hierarchy as procedural programs: procedures that call sub-procedures recursively. The procedures in hierarchical reinforcement learning correspond to tasks, which can be recursively decomposed into sub-tasks. Termination conditions within each sub-task return the system back to the caller task. Our major borrowing from this literature is the idea of a sub-task or sub-goal, but our implementation of this idea is only superficially related to the models in reinforcement learning.

Mathematically, our design procedure is closest to work in inverse optimal control. This field concerns itself with the problem of inferring the cost function that an observed agent is obeying (Ratliff et al., 2009; Abbeel and Ng, 2004; Kalakrishnan et al., 2013). Within inverse optimal control, work connected to bi-level programming (Albrecht et al., 2010; Mombaur et al., 2010) is the closest to ours, as it explicitly considers the question of how to choose parameters for a lower-level cost function to minimize a higher-level cost function. In most studies using inverse optimal control, the higher-level cost function is a measure of deviation between the experimentally measured movements of a subject and the outputs of an optimal control model optimizing the lower-level cost function, defined by the higher-level parameters. The aim of this work is different from ours in that the goal is to understand the behavior of an experimental subject. Additionally, in implementation, these inverse optimal control studies have not used higher-level forward models to decouple the optimization across levels, nor have they memorized the results of optimization to construct feedback controllers.

Other work on hierarchical control in the context of theoretical modeling of human motor control has sought to reduce the dimensionality of the state space, sending the higher-level a compressed description of the state of the plant, to simplify the computation of optimal commands (Liu and Todorov, 2009). While this may often prove useful, it is orthogonal to our approach. In our model, the state dimensionality at the higher level is in fact almost two orders of magnitude larger than at the lower level because it includes new sensory data unavailable to the lower level. This massive dimensionality expansion is, for example, consistent with the observation that the visual and motor information for visuo-motor

control are first merged cortically, even though reflexive movements, for example, triggered by noxious stimuli, are processed entirely within the spinal cord. Thus, higher-level control can serve roles beyond dimensionality reduction and can take into account altogether new channels of information.

Work on dynamical movement primitives (Ijspeert et al., 2013; Schaal et al., 2005) aims to construct lower-level control systems obeying attractor or limit cycle dynamics that simplify the production and planning of smooth, feedback-sensitive movements. In this research program, a “canonical” low-dimensional dynamical system with intrinsic dynamics generates control outputs, and the goal is specified by sending a set of parameters to the movement primitive that specify the end-goal, for example. This work is quite closely connected to our own in its purpose. The primary difference is that work in this field studies a somewhat limited set of dynamical systems of a specific kind, and the higher-level control of these systems has not used higher-level forward models but instead sampling-based reinforcement learning and supervised learning from demonstration.

The idea of controlling a physical system by building a model of it is quite old, entering the connectionist literature with Nguyen and Widrow (1989) and Jordan and Rumelhart (1992), but existing in a prior incarnation as “model-reference adaptive control.” More recently, model-based control has been identified as a very data-efficient means to train controllers (Deisenroth and Rasmussen, 2011). Neural networks have also received renewed attention in recent years as generic substrates for feedback controllers (Huh and Todorov, 2009; Sutskever, 2013). To the best of our knowledge, the idea of using a network to model the feedback from another network and then to use that model to control the modeled network is novel. We expect it to have ramifications that exceed the traditional framework of motor control. Forward models could be constructed either by motor babbling or by a more intelligent experimentation procedure. Once the controllers have been trained, they exhibit *automaticity*, an ability to generate answers without extensive computation. Automating complex computations by “caching” has been proposed before (Kavukcuoglu et al., 2008; Dayan, 2009) in different contexts, but using cached circuits as lower-level substrates for higher-level control circuits has not.

Conclusions

Our work consists of three innovations: 1) Dividing the task hierarchically by designing a higher level that propagates cost-related information to a lower-level; 2) Introducing higher-level forward models to train higher-level controllers; and 3) Caching computed optimal commands in network controllers.

Apart from navigation problems, the hierarchical scheme and training procedure using forward models should be useful for tackling any problem requiring simultaneous application of sensorimotor and cognitive skills. Consider, for example, a robotic gripper moving blocks. It would be quite a challenge to construct a unitary network that directs the gripper to pick up, move, and drop blocks and that decides how to arrange them into a prescribed pattern at the same time. It is easier to separate the problem into a manual coordination task and a puzzle-solving task, and a natural vehicle for this separation is the network architecture itself. More generally, the approach we have described is suited to

problems that can be formulated and solved by division into easier sub-problems. Many interesting tasks have this structure, so this strategy should be quite widely applicable. Surprisingly, despite the intuitiveness of this idea, it has received little attention.

The construction of the hierarchy is recursive or self-similar. This makes training considerably easier because each controller in the hierarchy is basically doing the same thing: receiving and issuing commands describing goals. Thus, we can apply the same training procedure at each level. Another advantage of this approach is that lower-level controllers do not need to be retrained if the overall task changes. In addition, the hierarchy can be extended by adding more levels if the task gets more difficult, again without requiring retraining of the lower levels.

Although we have trained the forward-model and controller networks sequentially, we do not preclude the attractive possibility of training them at the same time, maybe even at multiple levels of the hierarchy simultaneously. To do so, we expect it would be crucial to account for errors in the forward models; we could generate cautious commands by penalizing those commands that a forward model is unlikely to predict accurately, or we could also generate commands that attempt to probe the response properties of the level below to improve the model.

This study presents a hard result, a hierarchical neural network, and a method for training it based on higher-level forward models, but, perhaps equally importantly, it bolsters a soft view. To create intelligent behaviors from neuron-like components, we need to embed those neurons in modules that perform specific functions and operate to a large extent independently. These modules should have protocols for interfacing one to another, protocols which effectively hide the complexity of the full computation from the constituents. In experimental neuroscience, studying how such modules interact may require us to move beyond single-area recordings to understand the causal interactions between connected regions and to identify the goals of the computations performed by each region.

Methods

A. Parameters

The lower-level dynamic time scale T_1 was chosen to be as many time steps as possible without causing the emergence of jackknifing of the truck during optimization. The length of the lower-level control sequence K_1 was chosen to be relatively short, subject to the requirement that the lower-level controller needed to be able to turn around and approach a nearby target. The performance of the lower-level controller was not strongly sensitive to these particular values for K_1 and T_1 . The upper-level dynamic time scale T_2 was chosen to be shorter than the amount of time it takes for the truck to cross directly through an obstacle. Otherwise, the optimization could command the truck through an obstacle, but, because the forward model is only showing a “before-and-after” snapshot of the location of the truck, the whole system would be blind to its error. The length of a higher-level command sequence K_2 was chosen so that a truck that drives straight for $K_2 T_2$ time steps would not reach beyond the edge of the obstacle grid observed at time $t = 0$. That is, the sensory system contains most of the information needed for planning $K_2 T_2$ steps. Again, small variations in

these numbers do not change the behavior of the system significantly, but the numbers are intelligently chosen to make the system work as well as possible.

B. Network Structure and Training

The lower-level makes use of radial-basis function (RBF) neural networks. RBF networks possess a strong bias toward generating smoothly-varying outputs as a function of their inputs and can train very quickly on low-dimensional input data, two qualities that are useful at the lowest level of the hierarchy. The functional targets for the higher-level networks were more complicated and higher-dimensional, so we had to develop “deep” (or many-layered) networks to approximate them accurately.

The architectures of the networks are:

Network	Dimensions and Activation Functions
Lower-Level Forward	$6 \times [G]150 \times [L]5$
Lower-Level Controller	$5 \times [G]100 \times [L]1$
Higher-Level Forward (Proprio.)	$5 \times [T]40 \times [T]20 \times [T]20 \times [L]2$
Higher-Level Forward (Goal)	$8 \times [T]40 \times [T]20 \times [T]20 \times [L]3$
Higher-Level Forward (Obstacle)	$205 \times [T]300 \times [T]200 \times [T]200 \times [T]300 \times [Si]199$
Higher-Level Critic (Obstacle)	$199 \times [T]300 \times [T]100 \times [T]100 \times [So]1$
Higher-Level Controller	$205 \times [So]30 \times [So]20 \times [So]30 \times [L]3$

The bracketed letters indicate the activation functions used for the units. [G] is a normalized Gaussian radial basis function for input x ,

$$\frac{\exp\left(-\|\mathbf{x} - \mu_i\|^2 / (2\sigma_i^2)\right)}{\sum_j \exp\left(-\|\mathbf{x} - \mu_j\|^2 / (2\sigma_j^2)\right)},$$

with basis function centers μ_i and standard deviations σ_i . The radial basis function centers were chosen by randomly selecting exemplars from the input data. They were not further adapted in training. The RBF standard deviations were initialized to scale linearly with the number of input dimensions. To avoid division-by-zero, we actually computed using the inverse standard deviations $1/\sigma_i$.

All the other multiplication signs in Table 1 imply matrix multiplication followed by an activation function. [L] is a linear activation function, [T] is $\tanh(x)$, [Si] is a logistic sigmoid $1/(1 + \exp(-x))$, and [So] is a “soft-rectification” function $\log(1 + \exp(x))$. We chose these activation functions using a mixture of prior knowledge and experimentation. The soft-rectification in the obstacle critic imposed the constraint that costs are positive. The logistic function in the obstacle grid forward model bounds the outputs between 0 and 1 and allows us to interpret them as the probabilities that the grid points will be occupied. The soft-rectification in the higher-level controller alleviated over-fitting.

A weight matrix W from a layer of size M to another layer was initialized to have independent Gaussian entries of mean 0 and standard deviation $1/\sqrt{M}$. This makes the networks “balanced” in that the mean input to every unit is 0 and the response variance of all units is self-consistently $\mathcal{O}(1)$. This assures a constant response variance from one layer to the next. All multi-layer networks included a single additional bias unit in their inputs.

We chose to interpret an obstacle grid model output as the probability that the corresponding grid element would be occupied after movement. We consequently used the “crossentropy” cost function to train this network. We can derive this cost function from a maximum-likelihood argument. Suppose we have a data set of patterns $(\mathbf{x}_k, \mathbf{y}_k)_{k=1}^{N_{\text{patterns}}}$ and a neural network with M outputs whose i -th output unit as a function of the k -th input vector given by $z_i(\mathbf{x}_k)$ represents the binomial probability that $y_i(\mathbf{x}_k)=1$. The likelihood of the data is binomial, so the log-likelihood is

$$\text{Log-Likelihood} = \sum_{k=1}^{N_{\text{patterns}}} \sum_{i=1}^M y_i(\mathbf{x}_k) \log z_i(\mathbf{x}_k) + (1 - y_i(\mathbf{x}_k)) \log(1 - z_i(\mathbf{x}_k)).$$

When negated, this gives the cross-entropy cost function.

Optimization of the network parameters was accomplished using batch training with the quasi-Newton optimization method L-BFGS in “minFunc” (Schmidt, 2013). To train the lower-level forward model or the lower-level controller, at the beginning of each trial a random angle θ_{st} was drawn along with a random distance d_{st} in the range between 0 and 500. The trailer angle θ_{trailer} was similarly drawn uniformly. The cab angle was initialized to be within $\pm(\pi/2 - \pi/64)$ radians of the trailer angle.

The higher-level proprioceptive and goal-related models were trained to predict not the values of their targets but the difference between the values of their targets before and after movement. This reduced training time because the interesting predictions of many forward models are the deviations from the identity.

C. Equations for the Truck

The truck is a kinematic model of a cab and trailer (Figure 9), first defined by Nguyen and Widrow (Nguyen and Widrow, 1989). The cab is connected to the trailer by a rigid linkage.

The wheels are connected to the front of the cab and translate backward by distance r in one time step. Note that the wheels drive the cab the same way that the linkage drives the trailer, so we can solve for the motion of the cab and trailer in a similar way.

We begin by decomposing the motion of the front of the cab caused by the wheels into a component orthogonal to the front of the cab, defined as $A = r \cos(u(t))$, and a component parallel to the front of the cab, $C = r \sin(u(t))$. Only the orthogonal component, A , gets transferred through the linkage to the trailer. Performing a similar decomposition of the motion of the front of the trailer, we find an orthogonal component $B = A \cos(\theta_{\text{cab}}(t) -$

$\theta_{\text{trailer}}(t)$) and a parallel component $D=A \sin(\theta_{\text{cab}}(t) - \theta_{\text{trailer}}(t))$. In one time step, the center of the front of the trailer (Figure 8) therefore moves by

$$\begin{aligned} x'(t+1) &= x'(t) - B \cos(\theta_{\text{trailer}}(t)) + D \sin(\theta_{\text{trailer}}(t)) \\ y'(t+1) &= y'(t) - B \sin(\theta_{\text{trailer}}(t)) - D \cos(\theta_{\text{trailer}}(t)). \end{aligned}$$

The back of the trailer is constrained to move straight backward, so

$$\begin{aligned} x(t+1) &= x(t) - B \cos(\theta_{\text{trailer}}(t)) \\ y(t+1) &= y(t) - B \sin(\theta_{\text{trailer}}(t)). \end{aligned}$$

If the length of the trailer is L_{trailer} , $x'(t) = x(t) + L_{\text{trailer}} \cos(\theta_{\text{trailer}}(t))$ and $y'(t) = y(t) + L_{\text{trailer}} \sin(\theta_{\text{trailer}}(t))$. The tangent of the angle of the trailer is equal to $(y' - y)/(x' - x)$, so at time $t + 1$, we have

$$\tan(\theta_{\text{trailer}}(t+1)) = \frac{L_{\text{trailer}} \sin(\theta_{\text{trailer}}(t)) - D \cos(\theta_{\text{trailer}}(t))}{L_{\text{trailer}} \cos(\theta_{\text{trailer}}(t)) + D \sin(\theta_{\text{trailer}}(t))}.$$

An identical argument applied to the cab yields

$$\tan(\theta_{\text{cab}}(t+1)) = \frac{L_{\text{cab}} \sin(\theta_{\text{cab}}(t)) - C \cos(\theta_{\text{cab}}(t))}{L_{\text{cab}} \cos(\theta_{\text{cab}}(t)) + C \sin(\theta_{\text{cab}}(t))}.$$

Consolidating all of the equations, we have

$$\begin{aligned} A &= r \cos(u(t)) \\ B &= A \cos(\theta_{\text{cab}}(t) - \theta_{\text{trailer}}(t)) \\ C &= r \sin(u(t)) \\ D &= A \sin(\theta_{\text{cab}}(t) - \theta_{\text{trailer}}(t)) \\ x(t+1) &= x(t) - B \cos(\theta_{\text{trailer}}(t)) \\ y(t+1) &= y(t) - B \sin(\theta_{\text{trailer}}(t)) \\ \theta_{\text{cab}}(t+1) &= \tan^{-1} \left(\frac{L_{\text{cab}} \sin(\theta_{\text{cab}}(t)) - C \cos(\theta_{\text{cab}}(t))}{L_{\text{cab}} \cos(\theta_{\text{cab}}(t)) + C \sin(\theta_{\text{cab}}(t))} \right) \\ \theta_{\text{trailer}}(t+1) &= \tan^{-1} \left(\frac{L_{\text{trailer}} \sin(\theta_{\text{trailer}}(t)) - D \cos(\theta_{\text{trailer}}(t))}{L_{\text{trailer}} \cos(\theta_{\text{trailer}}(t)) + D \sin(\theta_{\text{trailer}}(t))} \right). \end{aligned}$$

It is worth mentioning that $(x, y, \theta_{\text{cab}}, \theta_{\text{trailer}})$ is a four-dimensional state vector with a one-dimensional control variable, u . Control theorists call problems in which the control vector is lower-dimensional than the state vector “under-actuated”; such problems are typically more difficult than “fully-actuated” problems because it may take more than one step to modify a given state variable, and it may be impossible to drive the system to an arbitrary point in the state space (a concept known as “controllability”). The truck is also a nonlinear system and an unstable one because setting $u=0$ amplifies any angular deviation of the cab from the trailer.

D. Minimizing the Cost Functions

In this section we describe how to minimize the cost functionals with respect to the command parameters. We are minimizing a cost functional of the form

$$\mathcal{J} = \sum_{k=0}^K \mathcal{L}(\mathbf{s}(k+1), \mathbf{m}(k)) \quad (6)$$

as in equation 1, but to streamline the notation, we have dropped the time t and the temporal scale factor T that appear in equation 1. This means that we have shifted the time variable to the starting at time t and are measuring time in units of T . The original equations can be recovered by shifting and scaling back. We have also dropped the subscripts l because the same procedure is applied at each level.

The sensory vector $\mathbf{s}$ is estimated by a forward model, and we denote the output of the forward model by $\mathbf{F}(\mathbf{s}, \mathbf{m})$, so that $\mathbf{s}(k+1)$ is estimated as $\mathbf{F}(\mathbf{s}(k), \mathbf{m}(k))$. To implement this constraint, we introduce Lagrange multipliers for every component of $\mathbf{m}$ and at every moment in time and minimize

$$\mathcal{J}_{\text{constrained}} = \sum_{k=0}^K \mathcal{L}(\mathbf{s}(k+1), \mathbf{m}(k)) + \lambda(k)^\top (\mathbf{F}(\mathbf{s}(k), \mathbf{m}(k)) - \mathbf{s}(k+1)). \quad (7)$$

For convenience, we define $\lambda(K+1)=0$.

The gradient of this function with respect to the sensory vector is

$$\frac{\delta \mathcal{J}_{\text{constrained}}}{\delta \mathbf{s}(k)} = \mathcal{L}_{\mathbf{s}}(\mathbf{s}(k), \mathbf{m}(k-1)) + \mathbf{F}_{\mathbf{s}}(\mathbf{s}(k), \mathbf{m}(k))^\top \lambda(k) - \lambda(k-1),$$

where the subscript $\mathbf{s}$ indicates a derivative with respect to that variable. At a minimum of the cost functional we find the backward equation for the Lagrange multipliers

$$\lambda(k-1) = \mathcal{L}_{\mathbf{s}}(\mathbf{s}(k), \mathbf{m}(k-1)) + \mathbf{F}_{\mathbf{s}}(\mathbf{s}(k), \mathbf{m}(k))^\top \lambda(k). \quad (8)$$

It is easy to find an extremum of the cost functional with respect to the sensory variables and the Lagrange multipliers. It is more difficult to minimize with respect to the command variables, where the relevant gradient is

$$\frac{\delta \mathcal{J}_{\text{constrained}}}{\delta \mathbf{m}(k)} = \mathcal{L}_{\mathbf{m}}(\mathbf{s}(k+1), \mathbf{m}(k)) + \mathbf{F}_{\mathbf{m}}(\mathbf{s}(k), \mathbf{m}(k))^\top \lambda(k).$$

This is done using the following algorithm:

Dynamic Optimization

Input: Initial state $\mathbf{s}(0)$ and “nominal” control tape $\{\mathbf{m}(k)\}_{k=0}^K$

repeat

for $k := 0$ to K **do**

$\mathbf{s}(k+1) := \mathbf{F}(\mathbf{s}(k), \mathbf{m}(k));$

end for

$\lambda(K+1) := 0;$

for $k := K+1$ down to 1 **do**

$\lambda(k-1) := \mathbf{L}_s(\mathbf{s}(k), \mathbf{m}(k-1)) + \mathbf{F}_s(\mathbf{s}(k), \mathbf{m}(k))^\tau \lambda(k);$

$\frac{\delta \mathcal{S}_{\text{constrained}}}{\delta \mathbf{m}(k-1)} = \mathbf{L}_m(\mathbf{s}(k), \mathbf{m}(k-1)) + \mathbf{F}_m(\mathbf{s}(k-1), \mathbf{m}(k-1))^\tau \lambda(k);$

    Update $\mathbf{m}(k-1)$ using a gradient-based method (steepest descent, L-BFGS, etc.)

    provided with $\frac{\delta \mathcal{S}_{\text{constrained}}}{\delta \mathbf{m}(k-1)}.$

end for

until convergence

E. Differential Dynamic Programming Comparison

We used an implementation of Differential Dynamic Programming (DDP) based on the description in the Ph.D. dissertation of Tom Erez (Erez, 2011). The DDP cost function was closely related to the two cost functions used in the hierarchical optimization, $\mathcal{L}_1$ and $\mathcal{L}_2$, but we had to modify a few of the terms because DDP appeared more sensitive to discontinuities of the derivatives of the original costs. Our overall cost function could be decomposed as $\mathcal{L}_{\text{DDP}} = \mathcal{L}_{\text{DDP}}^{\text{obstacle}} + \mathcal{L}_{\text{DDP}}^{\text{goal}} + \mathcal{L}_{\text{DDP}}^{\text{constraints}}.$

Here, $\mathcal{L}_{\text{DDP}}^{\text{constraints}}(\mathbf{s}_2) = u^2/2 + 20(1 - \cos(\theta_{\text{rel}}))^4$. These two terms roughly accomplish the same purpose as the constraint terms in the lower-level cost function but are $\mathcal{C}^\infty$ smooth.

$\mathcal{L}_{\text{DDP}}^{\text{goal}}(\mathbf{s}_2) = \log(1 + d/L)$, as in equation 3. The obstacle cost function was identical to equation 5, except that the minimum operation used to compute $d_{\min}^{\text{obstacle}}$ was replaced by a smooth approximation: $\min(x_1, x_2, \dots, x_n) \approx \left(\sum_{i=1}^n x_i^p\right)^{1/p}$, with $p=10$.

We executed DDP as a receding-horizon controller in which trajectories of length $K_2 T_2$ steps were planned at once. The first step of the plan was executed, and the entire sequence of optimized motor commands was shifted back by 1 step. This was used to “warm-start” DDP on the next iteration of planning. The motor command sequence at initialization of each trial was chosen to be the same as the command sequence that the hierarchical controller selected.

When comparing the costs of trajectories, we computed $\mathcal{L}_{DDP}^{\text{obstacle}} + \mathcal{L}_{DDP}^{\text{goal}}$ for both trajectories. Therefore, the hierarchical controller was evaluated on a cost function it was not exactly trained to minimize.

Acknowledgments

We are grateful to Yann LeCun, Sara Solla, John Krakauer, and Peter Dayan for ideas, support, and criticism. We thank Yashar Ahmadian, Loïc Matthey, Mattia Rigotti, Ann Kennedy, Darcy Wayne, and Saul Kato for helpful discussions. Research supported by the Gatsby, Swartz, Kavli, and National Science Foundations and by NIH grant MH093338.

References

- Abbeel P, Ng AY. Apprenticeship learning via inverse reinforcement learning. Proceedings of the twenty-first international conference on Machine learning. 2004
- Albert CY, Margoliash D. Temporal hierarchical control of singing in birds. *Science*. 1996; 273(5283): 1871–1875. [PubMed: 8791594]
- Albrecht S, Sobotka M, Ulbrich M. A bilevel optimization approach to obtain optimal cost functions for human arm-movements. Preprint in preparation, Fakultät für Mathematik, TU München. 2010
- Barto AG, Mahadevan S. Recent advances in hierarchical reinforcement learning. *Discrete Event Dynamic Systems*. 2003; 13(4):341–379.
- Bernstein, NA. The co-ordination and regulation of movements. Pergamon Press Ltd.; 1967.
- Dayan P. Goal-directed control and its antipodes. *Neural Networks*. 2009
- Deisenroth, M.; Rasmussen, CE. Pilco: A model-based and data-efficient approach to policy search. Proceedings of the 28th International Conference on Machine Learning (ICML-11); 2011. p. 465–472.
- Dietterich TG. The maxq method for hierarchical reinforcement learning. In ICML. 1998:118–126. Citeseer.
- Erez, T. Doctoral Dissertation. Washington University in St. Louis. 2011. Optimal control for autonomous motor behavior.
- Gat E. On Three-Layer Architectures. *Artificial Intelligence and Mobile Robots*. 1998
- Hantman AW, Jessell TM. Clarke's column neurons as the focus of a corticospinal corollary circuit. *Nature Neuroscience*. 2010
- Huh D, Todorov E. Real-time motor control using recurrent neural networks. *Adaptive Dynamic Programming and Reinforcement Learning*. 2009
- Ijspeert AJ, Nakanishi J, Hoffmann H, Pastor P, Schaal S. Dynamical movement primitives: learning attractor models for motor behaviors. *Neural computation*. 2013; 25(2):328–373. [PubMed: 23148415]
- Jackson JH. On the comparative study of diseases of the nervous system. *British Medical Journal*. 1889; 2:355–62.
- Jordan MI, Rumelhart DE. Forward Models: Supervised Learning with a Distal Teacher. *Cognitive Science*. 1992
- Kalakrishnan, M.; Pastor, P.; Righetti, L.; Schaal, S. Learning objective functions for manipulation. *Robotics and Automation (ICRA)*. 2013 IEEE International Conference on; 2013. p. 1331–1336.
- Kavukcuoglu K, Ranzato M, LeCun Y. Fast inference in sparse coding algorithms with applications to object recognition. Technical Report CBLL-TR-2008-12-01: Computational and Biological Learning Lab, Courant Institute, NYU. 2008
- Lashley KS. Basic Neural Mechanisms in Behavior. *Psychological Review*. 1930
- LeCun Y. A Theoretical Framework for Backpropagation. Proceedings of the 1988 Connectionist Models Summer School. 1988

- Liu, D.; Todorov, E. Adaptive Dynamic Programming and Reinforcement Learning, 2009. ADPRL'09. IEEE Symposium on. IEEE; 2009. Hierarchical optimal control of a 7-dof arm model; p. 50-57.
- Loeb G, Brown I, Cheng E. A hierarchical foundation for models of sensorimotor control. *Experimental brain research*. 1999; 126(1):1–18. [PubMed: 10333003]
- Miall RC, Weir DJ, Wolpert DM, Stein JF. Is the Cerebellum a Smith Predictor? *Journal of Motor Behavior*. 1993
- Mombaur K, Truong A, Laumond J-P. From human to humanoid locomotion an inverse optimal control approach. *Autonomous robots*. 2010; 28(3):369–383.
- Nguyen, D.; Widrow, B. The Truck Backer-Upper: An Example of Self-Learning in Neural Networks. *Proceedings of the International Joint Conference on Neural Networks*; 1989.
- Raibert, MH. Doctoral Dissertation. Massachusetts Institute of Technology. 1977. Motor control and learning by the state space model.
- Ratliff N, Ziebart B, Peterson K, Bagnell JA, Hebert M, Dey AK, Srinivasa S. Inverse optimal heuristic control for imitation learning. *AISTATS*. 2009
- Schaal S, Peters J, Nakanishi J, Ijspeert A. Learning movement primitives. *Robotics Research*. 2005:561–572.
- Schmidt, M. minFunc. 2013. <http://www.di.ens.fr/~mschmidt/Software/minFunc.html>
- Shadmehr R, Smith MA, Krakauer JW. Error Correction, Sensory Prediction, and Adaptation in Motor Control. *Annual Review of Neuroscience*. 2010
- Stengel RF. *Optimal Control and Estimation*. Dover publications. 1994
- Sumbre G, Gutfreund Y, Fiorito G, Flash T, Hochner B. Control of octopus arm extension by a peripheral motor program. *Science*. 2001; 293(5536):1845–1848. [PubMed: 11546877]
- Sutskever, I. Doctoral Dissertation, University of Toronto. 2013. Training Recurrent Neural Networks.
- Sutton, RS.; Barto, AG. *Reinforcement Learning: An Introduction*. The MIT Press; 1998.
- Sutton RS, Precup D, Singh S. Between mdps and semi-mdps: A framework for temporal abstraction in reinforcement learning. *Artificial intelligence*. 1999; 112(1):181–211.
- Tassa, Y.; Erez, T.; Todorov, E. Fast model predictive control for reactive robot swimming. 2011. <http://www.cs.washington.edu/homes/todorov/papers/MPCswimmer.pdf>
- Thrun S, Montemerlo M, et al. Stanley: The Robot that Won the DARPA Grand Challenge. *Journal of Field Robotics*. 2006
- Widrow B, Gupta NK, Maitra S. Punish/reward: learning with a critic in adaptive threshold systems. *IEEE Transactions on Systems, Man, and Cybernetics*. 1973

Figure 1.

Flow diagram of the hierarchical control system. Commands that control the wheel angle of the truck are issued by the lower-level controller, which receives information about a sub-target direction toward which the truck should be driven from the higher-level controller. Both controllers receive proprioceptive information about the angle between the cab and trailer of the truck, and the higher-level controller also receives information about obstacles in the environment from a grid of sensors. In addition, the higher-level controller receives input about the final target that the truck is supposed to reach.

Figure 2.

A) The lower-level controller directs the truck to a sub-target (white square). The black trace shows the path of the back of the truck. B) The lower-level controller directs the truck to trace the constellation *Ursa Major* (white line is the path of the truck) by approaching sub-targets at the locations of the stars. The sub-targets appear one at a time; when the back of the truck arrives close to the current sub-target, it is replaced by the next sub-target. Photo by Akira Fujii.

Figure 3.

The truck following a sequence of sub-targets provided by the higher-level controller. The sub-targets are indicated by black/grey/white squares with darker colors representing earlier times in the sequence. The trajectory that the truck follows in pursuing the sub-targets is shown in black. A connecting line indicates the sub-target that is active when the truck reaches particular trajectory points. The background coloration indicates the distance to the final target, located off the lower-left corner.

Figure 4.

A) An egocentric coordinate system surrounds the truck, composed of grid points. During movement, the grid shifts with the truck. Only a small fraction of the grid points is shown here. B) The full set of grid points in an environment with obstacles (yellow circles). Those points that lie within an obstacle are blackened, indicating that the grid element is activated.

Figure 5.

A) The truck is directed to avoid the obstacles and reach the final target. The white square indicates the first sub-target; note that it is not at a position the truck actually reaches. The higher-level controller merely uses this to indicate the desired heading to the lower-level controller. B) With 50 static obstacles, more than the 20 that were present during training, the higher-level controller steers the truck around all of the obstacles to the final target on each of 50 consecutive trials. The black lines show the paths taken by the back of the trailer.

Figure 6.

Performance measures. A) Obstacle Avoidance. Black dots show the minimum distance between the truck and any target averaged over 50 runs to the final goal with 20 obstacles in the environment. B) Collisions versus Obstacles. The black line shows the probability of a collision with an obstacle per trip to the final target as a function of the number of obstacles. The shaded regions are 95% confidence intervals. As the number of obstacles increases from 5 to 75, the probability of a collision grows slowly, despite high obstacle densities. C) Target-directedness. The black dots show the lengths of paths taken to the final target, averaged over 50 trials, in an environment with 20 obstacles. The dashed line shows the straight-line distance from the initial location of the truck to the nearest edge of the final target. D) Brownian Obstacle Motion. The black line shows the probability of a collision with an obstacle per trip as a function of the diffusion constant of the obstacle motion. The gray region is as in B. Although we did not explicitly train the controllers to handle obstacle movement, the controller can frequently navigate to the goal without collision in an environment of 20 obstacles undergoing Brownian motion.

Figure 7.

Comparison with Optimal Control. A) In 25 trials in a single environment, we can see that the hierarchical controller's (continuous line) and optimal control solver's (dashed line) trajectories often overlap. On some trials, the trajectories diverge symmetrically around an obstacle. On other trials, the optimal control solver finds poor local minima in which the trajectories develop loops. This behavior is extremely rare for the hierarchical controller. B) The per-trial cost differences over 150 trials in random environments show that the hierarchical controller infrequently performs worse than the optimal control solver (positive tail in the distribution), often performs equivalently (peak near 0), and sometimes perform much better when the optimal control calculation fails (negative tail). The average cost per path for either scheme is approximately 300.

Figure 8.

Schematic showing the iterated use of the forward model for computing an optimal command sequence. The system is provided with sensory input at time t through the vector $\mathbf{s}_l(t)$. The forward model is used repeatedly to generate predictions of the sensory vector at times $t + T_l, \dots, t + (K_l + 1)T_l$. The command sequence $\mathbf{m}_l(t), \mathbf{m}_l(t + T_l), \dots, \mathbf{m}_l(t + K_l T_l)$ is an input to both the cost-function computation and the forward model, and it is optimized.

Figure 9.

Variables describing the truck. The position of the center of the back of the trailer is given by (x, y) , and the center of the front of the trailer is at (x', y') . The angle of the cab with respect to the x-axis is $\theta_{\text{cab}}(t)$ and for the trailer, $\theta_{\text{trailer}}(t)$. The angle by which the front wheels deviate from straight ahead is u , and the relative angle between the cab and trailer is $\theta_{\text{rel}} = \theta_{\text{cab}} - \theta_{\text{trailer}}$.

Tabel 1

Parameters for (from left to right) the truck and obstacles, the lower- and higher- level cost function, and the lower- and higher-level training. $N_{c,l}$ denotes the number of example patterns used to train the controller and $N_{fm,l}$ the number for the forward model. k_l^{total} refers to the number of steps taken when optimizing trajectories before restaring the truck in a new environment.

truck & obstacles				lower-level cost function				higher-level cost function	lower-level training	higher-level training
Parameter	Value	Parameter	Value	Parameter	Value	Parameter	Value	Parameter	Value	
r (speed)	0.2	a_1	1	a_2	1	T_1	6	T_2	$12 \cdot T_1$	
l_{cab} (length)	6	β_1	0.5	β_2	100	k_l^{model}	15	k_2^{model}	1	
w_{cab} (width)	6	γ_1	0.5	γ_2	100	$k_l^{\text{optimization}}$	15	k_2^{total}	10	
l_{trailer}	14	ζ_1	0.5	d_{min}	0.1	k_l^{total}	30	k_2^{total}	10	
w_{trailer}	6	θ_{max}	$\frac{\pi}{2} - \frac{\pi}{6}$	d_{max}	0.9	$N_{c,1}$	2×10^4	$N_{c,1}$	5×10^5	
σ_{disc} (radius)	20			σ_{disc}	20	$N_{fm,1}$	5×10^4	$N_{fm,2}$	2×10^6	
				σ_{areola}	30					
				ρ_2	2					
				ν_2	0.5					