Movement Generation with Circuits of Spiking Neurons

Prashant Joshi
joshi@igi.tugraz.at
Wolfgang Maass
maass@igi.tugraz.at
Institute for Theoretical Computer Science, Technische Universität Graz,
A-8010 Graz, Austria

How can complex movements that take hundreds of milliseconds be generated by stereotypical neural microcircuits consisting of spiking neurons with a much faster dynamics? We show that linear readouts from generic neural microcircuit models can be trained to generate basic arm movements. Such movement generation is independent of the arm model used and the type of feedback that the circuit receives. We demonstrate this by considering two different models of a two-jointed arm, a standard model from robotics and a standard model from biology, that each generates different kinds of feedback. Feedback that arrives with biologically realistic delays of 50 to 280 ms turns out to give rise to the best performance. If a feedback with such desirable delay is not available, the neural microcircuit model also achieves good performance if it uses internally generated estimates of such feedback. Existing methods for movement generation in robotics that take the particular dynamics of sensors and actuators into account (embodiment of motor systems) are taken one step further with this approach, which provides methods for also using the embodiment of motion generation circuitry, that is, the inherent dynamics and spatial structure of neural circuits, for the generation of movement.

1 Introduction

Using biologically realistic neural circuit models to generate movements is not so easy, since these models are made of spiking neurons and dynamic synapses, which exhibit a rich inherent dynamics on several temporal scales. This tends to be in conflict with movement tasks that require sequences of precise motor commands on a relatively slow timescale. However, we show that without the construction of any particular circuit, training a linear readout to take a suitable weighted sum (with fixed weights after training) of the output activity of a fairly large number of neurons in a generic neural microcircuit model provides a very general paradigm for movement generation. It is obviously reminiscent of a number of experimental results (see, e.g.,

Wessberg et al., 2000) that show that a suitable weighted sum of the activity from a fairly large number of cortical neurons in monkeys predicts quite well the trajectory of hand positions for a variety of arm movements. Obviously the neural microcircuit model assumes here a similar role as a kernel for support vector machines in machine learning (for details, see Maass, Legenstein, & Bertschinger, 2004, and Maass, Natschläger, & Markram, 2004).

This letter demonstrates that controllers made from generic neural microcircuits are functionally generic in the sense that readouts from such circuits can learn to control the arm regardless of the model used to describe the arm dynamics, the type of feedbacks used (visual or proprioceptive), and the type of movements that are generated. This is shown here by teaching the same generic neural circuit to generate reaching movements for two different models with different kinds of feedbacks. The first model used (model 1) is the standard model of a two-joint robot arm described in Slotine and Li (1991). The other model (Todorov, 2000, 2003) comes from biology and relates the activity of neurons in the cortical motor area M1 to the kinematics of the arm (model 2).

It turns out that both the spatial organization of information streams, especially the population coding of slowly varying input variables, and the inherent dynamics of the generic neural microcircuit model have a significant impact on its capability to generate movements. In particular, it is shown that the inherent dynamics of neural microcircuits allows these circuits to cope with rather large delays for proprioceptive and sensory feedback. In fact, it turns out that the performance of this generic neurocontroller is optimal for feedback delays that lie in the biologically realistic range of 50 to 280 ms. Furthermore, it is shown that other readout neurons from the same neural microcircuit model can be trained simultaneously to estimate results of such feedback, and that in the absence of real feedback, the precision of reaching movements can be improved significantly if the circuit gets access to these estimated feedbacks.

This work complements preceding work where generic neural microcircuit models were used in an open loop for a variety of sensory processing tasks (Buonomano & Merzenich, 1995; Maass, Natschläger, & Markram, 2002; Maass, Natschläger, & Markram, 2004). It turns out that the demands on the precision of real-time computations carried out by such circuit models are substantially higher for closed-loop applications such as those considered in this article. The paradigm for movement generation discussed in this letter is somewhat related to preceding work (Ijspeert, Nakanishi, & Schaal, 2003), where a fixed parameterized system of differential equations was used instead of neural circuits, and to the melody generation and prediction of chaotic time series with artificial neural networks in discrete time of Jäger (2002) and Jäger and Haas (2004). In these other models, no effort is made to choose a movement generator whose inherent dynamics has a similarity to that of biological neural circuits. It has not yet been sufficiently

investigated whether feedback, especially feedback with a realistic delay, can have similarly beneficial consequences in these other models.

No effort was made in this article to make the process by which the neural circuit model (more specifically, the readouts from this circuit) learns to generate specific movement primitives biologically realistic. Hence, the results of this article provide evidence only that a generic neural microcircuit can hold the information needed to generate certain movement primitives and that it can generate a suitable slow dynamics with high precision.

The structure of this letter is as follows. Section 2 describes the neural microcircuit model. This is followed by the description of the robot arm model (model 1) in section 3. Sections 4, 5, and 6 present results of computer simulations for model 1. Section 7 repeats the experiment described in section 4 for the biologically motivated arm model (model 2). Finally, we discuss robustness issues related to our new paradigm for movement generation in section 8.

A preliminary version of some results from this letter (for movements of just one fixed temporal duration, and without model 2) was presented at a conference (Joshi & Maass, 2004).

2 Generic Neural Microcircuit Models

In contrast to common artificial neural network models, neural microcircuits in biological organisms consist of diverse components, such as different types of spiking neurons and dynamic synapses, that are each endowed with an inherently complex dynamics of its own. This makes it difficult to construct neural circuits out of biologically realistic computational units that solve specific computational problems, such as generating arm movements to various given targets. In fact, the generation of a smooth arm movement appears to be particularly difficult for a circuit of spiking neurons, since the dynamics of arm movements takes place on a timescale of hundreds of milliseconds, whereas the inherent dynamics of spiking neurons takes place on a much faster timescale. We show that this problem can be solved, even with a generic neural microcircuit model whose internal dynamics has not been adjusted or specialized for the task of creating arm movements, by taking as activation command for a muscle at any time t a weighted sum $\mathbf{w} \times \mathbf{z}(t)$ of the vector z that describes the current firing activity of all neurons in the circuit.[1] The weight vector $\mathbf{w}$, which remains fixed after training, is the only part that needs to be specialized for the generation of a particular movement task. Each component of $\mathbf{z}(t)$ models the impact that a particular neuron v may have on the membrane potential of a generic readout neuron. Thus, each spike of neuron v is replaced by a pulse of

---

[1] As usual, a constant component is formally included in $\mathbf{z}(t)$ so that the term $\mathbf{w} \times \mathbf{z}(t)$ may contain some fixed bias.

unit amplitude 1 that decays exponentially with a time constant of 30 ms. In other words, $\mathbf{z}(t)$ is obtained by applying a low-pass filter to the spike trains emitted by the neurons in the generic neural microcircuit model. Note that it is already known that hand trajectories of monkeys can be recovered from the current firing activity $\mathbf{z}(t)$ of neurons in motor cortex through the same types of weighted sums as considered in this article (Wessberg et al., 2000).

In principle, one can also view various parameters within the circuit as being subject to learning or adaptation, for example, in order to optimize the dynamics of the circuit for a particular range of control tasks. However, this has turned out to be not necessary for the applications described in this article, although it remains an interesting open research problem how unsupervised learning could optimize a circuit for motor control tasks. One advantage of viewing the weight vector $\mathbf{w}$ as being plastic is that learning is quite simple and robust, since it amounts to linear regression—in spite of the highly nonlinear nature of the control tasks to which this setup is applied. Another advantage is that the same neural microcircuit could potentially be used for various other information processing tasks (e.g., prediction of sensory feedback; see section 6) that may be desirable for the same or other tasks.

The generic microcircuit models used for the closed-loop control tasks described in this article were similar in structure to those used earlier for various sensory processing tasks in an open loop. More precisely, we considered circuits consisting of 600 leaky integrate-and-fire neurons arranged on the grid points of a $20 \times 5 \times 6$ cube in 3D (see Figure 1). Twenty percent of these neurons were randomly chosen to be inhibitory. Synaptic connections were chosen according to a biologically realistic probability distribution that favored local connections but also allowed some long-range connections. Biologically realistic models for dynamic synapses were employed instead of the usual static synapses of artificial neural network models. Parameters of neurons and synapses were chosen to fit data from microcircuits in rat somatosensory cortex (based on Gupta, Wang, & Markram, 2000, and Markram, Wang, & Tsodyks, 1998; see the appendix).

In order to test the noise robustness of movement generation by the neural microcircuit model, the initial condition of the circuit was randomly drawn (initial membrane potential for each neuron drawn uniformly from the interval [13.5 mV, 14.9 mV], where 15 mV was the firing threshold). In addition a substantial amount of noise was added to the input current of each neuron throughout the simulation at each time step; a new value for the noise input current with mean 0 and SD of 1 nA was drawn for each neuron and added (subtracted) to its input current.

In the case of the arm model that is considered in sections 3 through 6 (model 1), the neural circuit receives analog input streams from six sources (from eight sources in the experiment with internal predictions discussed
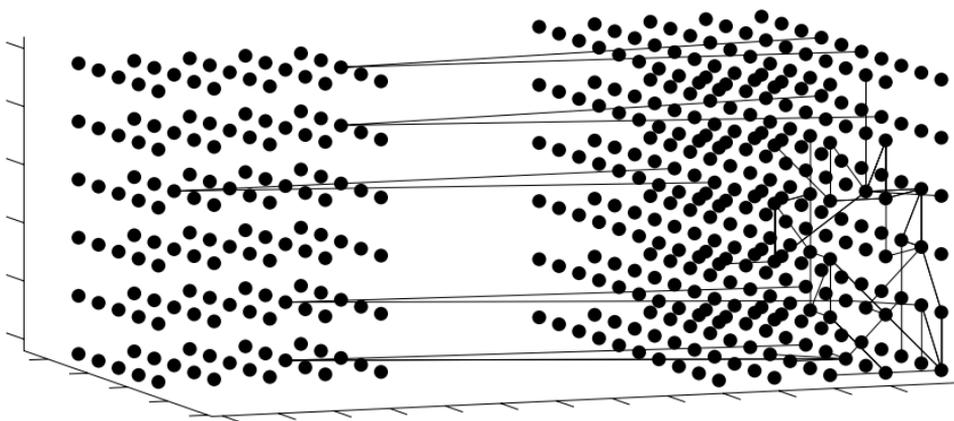
Figure 1: Spatial arrangement of neurons in the neural microcircuit models considered in this letter. The neurons in the six layers on the left-hand side encode the values of the six input-feedback variables x_{dest}, y_{dest}, $\theta_1(t - \Delta)$, $\theta_2(t - \Delta)$, $\tau_1(t)$, $\tau_2(t)$ in a standard population code. Connections from these six input layers (shown for a few selected neurons), as well as connections between neurons in the subsequent six processing layers, are chosen randomly according to a probability distribution discussed in the text (a typical example is shown).

in Figures 8 and 9). A critical factor for the performance of these neurocontrollers is the way in which these time-varying analog input streams are fed into the circuit. The outcomes of the experiments discussed in this article would have been all negative if these analog input streams were fed into the circuit as time-varying input currents. Apparently the variance of the resulting spike trains was too large to make the information about the slowly varying values of these input streams readily accessible to the circuit. Therefore, we employed instead a standard form of population coding (Pouget & Latham, 2003). Each of the six time-varying input variables was mapped onto an array of 50 symbolic input neurons with bell-shaped tuning curves (see the appendix). Thus, the value of each of the six input variables is encoded at any time by the output values of the associated 50 symbolic input neurons (of which at least 43 neurons output at any time the value 0). The neurons in each of these six input arrays are connected[2] with one of the six layers consisting of 100 neurons in the circuit of 100×6 $((20 \times 5) \times 6)$ integrate-and-fire neurons, providing a time-varying input current to a randomly selected subset of integrate-and-fire neurons on that layer (see Figure 1).

---

[2] With a value of 3.3 for λ in the formula for the connection probability given in the appendix.
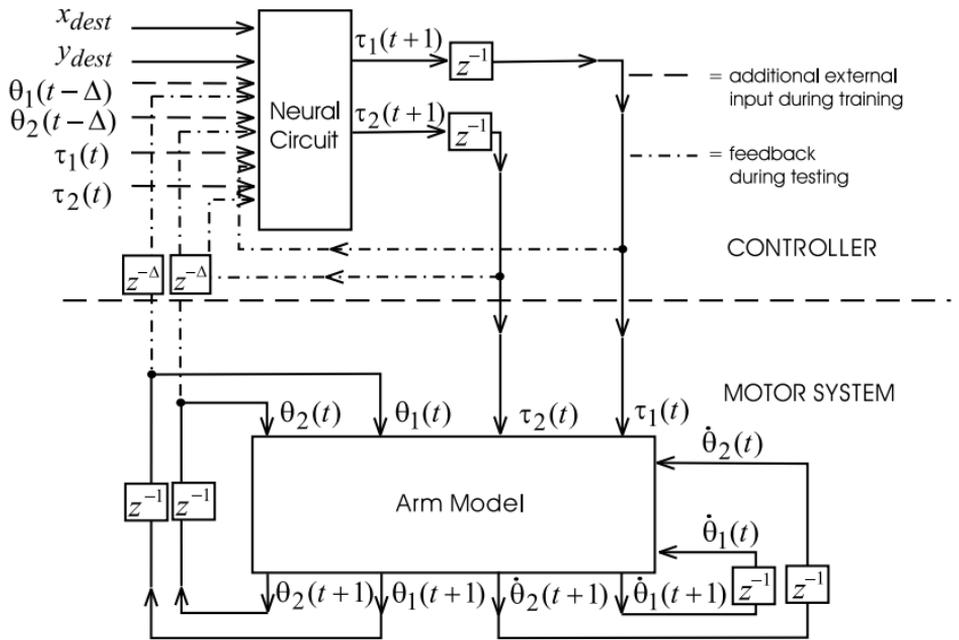
Figure 2: Closed-loop application of a generic neural microcircuit. The weight vectors of the linear readouts from this circuit that produce the next motor commands $\tau_1(t + 1)$, $\tau_2(t + 1)$ are the only parameters that are adjusted during training. After training, the neural circuit receives in this closed loop as inputs a target position x_{dest}, y_{dest} for the tip of the robot arm (in Cartesian coordinates; these inputs remain constant during the subsequent arm movement) as well as feedback $\theta_1(t - \Delta)$, $\theta_2(t - \Delta)$ from the arm representing previous values of joint angles delayed by an amount Δ, as well as "efferent copies" $\tau_1(t)$, $\tau_2(t)$ of its preceding motor commands. All the dynamics needed to generate the movement is then provided by the inherent dynamics of the neural circuit in response to the switching on of the constant inputs (and in response to the dynamics of the feedbacks). During training of the readouts from the generic neural circuits, the proprioceptive feedbacks $\theta_1(t - \Delta)$, $\theta_2(t - \Delta)$ and the efferent copies of previous motor commands $\tau_1(t)$, $\tau_2(t)$ are replaced by corresponding values for a target movement, which are given as external inputs to the circuit ("imitation learning").

3 A Two-Joint Robot Arm as a Benchmark Nonlinear Control Task

We first trained a generic neural microcircuit model (see Figures 1 and 2) to control a standard model for a two-joint robot arm (model 1; see Figure 3). This model is used in Slotine and Li (1991) as a standard reference model for a complex nonlinear control task (see in particular sections 6 and 9). It is assumed that the arm is moving in a horizontal plane, so that gravitational forces can be ignored.
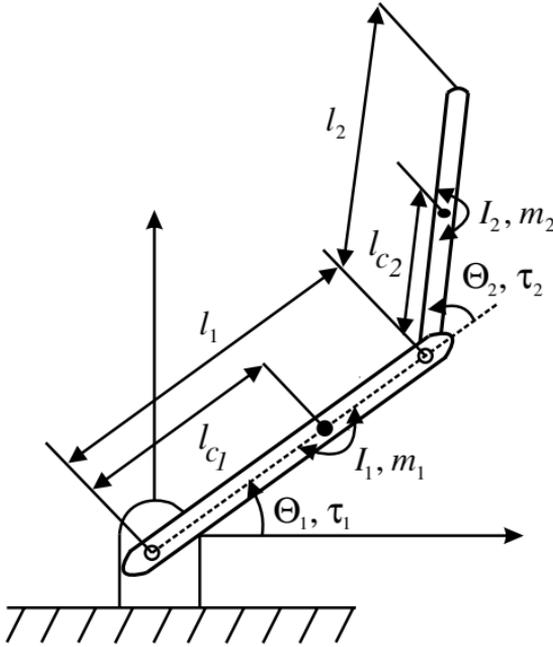
Figure 3: Standard model of a two-joint robot arm.

Using the well-known Lagrangian equation in classical dynamics, the dynamic equations for this arm model are given by equation 3.1:

$$\begin{bmatrix} H_{11} & H_{12} \\ H_{21} & H_{22} \end{bmatrix} \begin{bmatrix} \ddot{\theta}_1 \\ \ddot{\theta}_2 \end{bmatrix} + \begin{bmatrix} -h\dot{\theta}_2 & -h(\dot{\theta}_1 + \dot{\theta}_2) \\ h\dot{\theta}_1 & 0 \end{bmatrix} \begin{bmatrix} \dot{\theta}_1 \\ \dot{\theta}_2 \end{bmatrix} = \begin{bmatrix} \tau_1 \\ \tau_2 \end{bmatrix}, \tag{3.1}$$

with $\theta = [\theta_1 \theta_2]^T$ being the two joint angles, $\tau = [\tau_1 \tau_2]^T$ being the joint input torques to the two joints, and

$$H_{11} = m_1 l_{c_1}{}^2 + I_1 + m_2 \big[l_1{}^2 + l_{c_2}{}^2 + 2l_1 l_{c_2} \cos \theta_2 \big] + I_2$$

$$H_{12} = H_{21} = m_2 l_1 l_{c_2} \cos \theta_2 + m_2 l_{c_2}{}^2 + I_2$$

$$H_{22} = m_2 l_{c_2}{}^2 + I_2$$

$$h = m_2 l_1 l_{c_2} \sin \theta_2.$$

Equation 3.1 can be compactly written as

$$H(\theta)\ddot{\theta} + C(\theta, \dot{\theta})\dot{\theta} = \tau,$$

where H represents the inertia matrix and C represents the matrix of Coriolis and centripetal terms. I_1, I_2 are the moments of inertia of the two joints. The

values of the parameters that were used in our simulations were $m_1 = 1$, $m_2 = 1$, $l_{c_1} = 0.25$, $l_{c_2} = 0.25$, $I_1 = 0.03$, and $I_2 = 0.03$.

The closed-loop control system that we used is shown in Figure 2. During training of the weights of the linear readouts from the generic neural microcircuit model, the circuit was used in an open loop with target values for the output torques provided by equation 3.1 (for a given target trajectory $\{\theta_1(t), \theta(t)\}$), and feedbacks from the plant replaced by the target values of these feedbacks for the target trajectory. The delay Δ of the proprioceptive or sensory feedback is assumed to have a fixed value of 200 ms, except for section 6, where we study the impact of this value for the precision of the movement. For each such target trajectory, 20 variations of the training samples were generated, for which at each time step[3] t a different noise value of $10^{-5} \times \rho$ was added to each of the input channels where ρ is a random number drawn from a gaussian distribution with mean 0 and SD 1, multiplied by the current value of that input channel. The purpose of this extended training procedure was to make the readout robust with regard to deviations from the target trajectory caused by faulty earlier torque outputs given by the readouts from the neural circuit (see section 8). Each target trajectory had a time duration of 500 ms.

4 Teaching a Generic Neural Microcircuit Model to Generate Basic Movements

As a first task, the generic neural microcircuit model described in section 2 was taught to generate with the two-joint arm described in section 3 the four movements indicated in Figure 4. In each case, the task was to move the tip of the arm from point A to point B on a straight line, with a biologically realistic bell-shaped velocity profile. The two readouts from the neural microcircuit model were trained by linear regression to output the joint torques required for each of these movements.[4]

---

[3] All time steps were chosen to have a length of 2 ms, except for the experiment reported in Figure 6, where a step size of 1 ms was used to achieve a higher precision.

[4] Training data were generated as follows. For a given start point $\langle x_{start}, y_{start} \rangle$ and target end point $\langle x_{dest}, y_{dest} \rangle$ of a movement (both given in Cartesian coordinates), an interpolating trajectory of the tip of the arm was generated according to the following equation given in Flash and Hogan (1965):

$$x(t) = x_{start} + (x_{start} - x_{dest}) \cdot (15\tau^4 - 6\tau^5 - 10\tau^3)$$
$$y(t) = y_{start} + (y_{start} - y_{dest}) \cdot (15\tau^4 - 6\tau^5 - 10\tau^3),$$

where $\tau = t/MT$ and MT is the target movement time (in this case, $MT = 500$ ms). From this target trajectory for the end point of the robot arm, we had generated target trajectories of the angles Θ_1, Θ_2 of the robot arm by applying standard equations from geometry (see, e.g., Craig, 1955). From these, the target trajectories of the torques were generated according to equation 3.1.
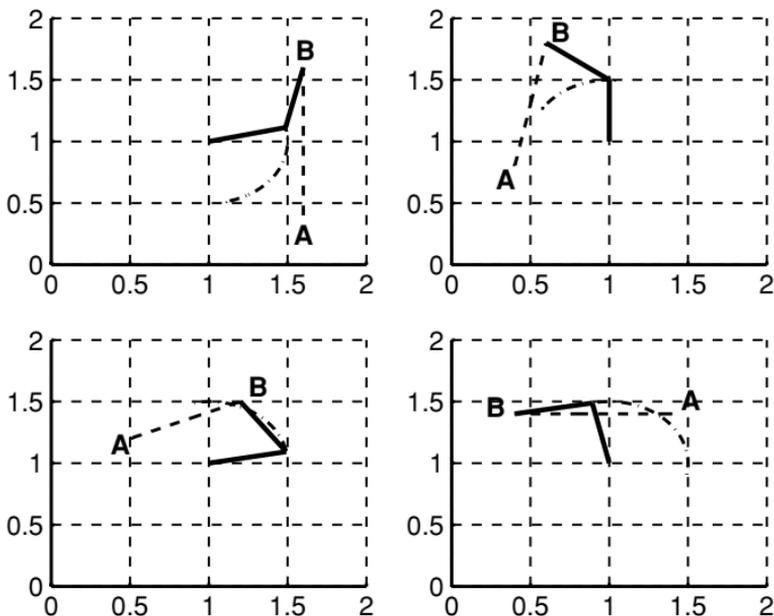
Figure 4: Initial position A and end position B of the robot arm (model 1) for four target movements, scaled in meters. The target trajectory of the tip of the robot arm and of the elbow are indicated by dashed and dashed-dotted lines. One sees clearly that even simple linear movements of the tip to the arm require quite nonlinear movements of the elbow.

Twenty noisy variations of each of the four target movements were used for the training of the two readouts by linear regression, as specified in section 3. Note that each readout is simply modeled as a linear gate with weight vector **w** applied to the liquid state **x**(t) of the neural circuit. This weight vector is fixed after training, and during validation all four movements are generated with this fixed-weight vector at the readout.

The performance of the trained neural microcircuit model during validation in the closed loop (see Figure 2) is demonstrated in Figure 5. When the circuit receives as input the coordinates $\langle x_{dest}, y_{dest} \rangle$ of the end point B of one of the target movements shown in Figure 4, the circuit autonomously generates in a closed loop the torques needed to move the tip of the two-joint arm from the corresponding initial point A to this end point B.[5]

---

[5] In these experiments, no effort was made to stabilize the end point of the arm at or near the target position. Rather, the movement was externally halted at the end of the allotted time period of 500 ms. Hence, the neural circuit model acts as a movement generator rather than as a controller. However, we are not aware of a fundamental obstacle that would make it impossible to teach a circuit to stabilize the arm once it has reached the target position (which the circuit receives as an extra input).
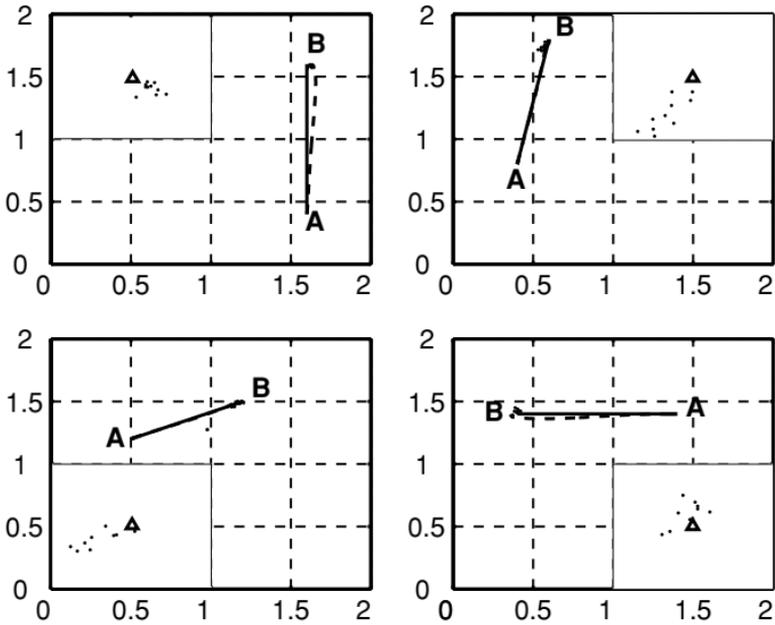
Figure 5: Target trajectories of the tip of the robot arm as in Figure 4 (solid) and resulting trajectories of the tip of the robot arm in a closed loop for one of the test runs (dashed line). The dots around the target end points show the end points of the tip of the robot arm for 10 test runs for each of the movements (enlarged inserts show a 20 cm × 20 cm area with target end point B marked by a black open triangle). Differences are due to varying initial conditions and simulated inherent noise of the neural circuit. Nevertheless, all movement trajectories converged to the target, with an average deviation from the target end point of 4.72 cm and the SD of 0.85 cm (scale of figures in m).

Obviously temporal integration capabilities of the controller are needed for the control of many types of movements. The next experiment was designed to test explicitly this capability of neurocontrollers constructed from generic circuits of spiking neurons. Figure 6 shows results for the case where the readouts from the neural microcircuit have been trained to generate an arm movement with an intermediate stop of all movement from 225 to 275 ms (see the velocity profile at the bottom of Figure 6). The initiation of the continuation of the movement at time $t = 275$ ms has to take place without any external cue, just on the basis of the inherent temporal integration capability of the neural circuit. For demonstration purposes, we chose for the experiment reported in Figure 6 a feedback delay of just 1 ms, so that all circuit inputs are constant during 49 ms of the 50 ms while the controller has to wait, forcing the readouts to decide only on the basis of the inherent circuit dynamics when to move on. Nevertheless, the average deviation of the tip of the robot arm for 20 test runs (with noisy initial conditions and noise
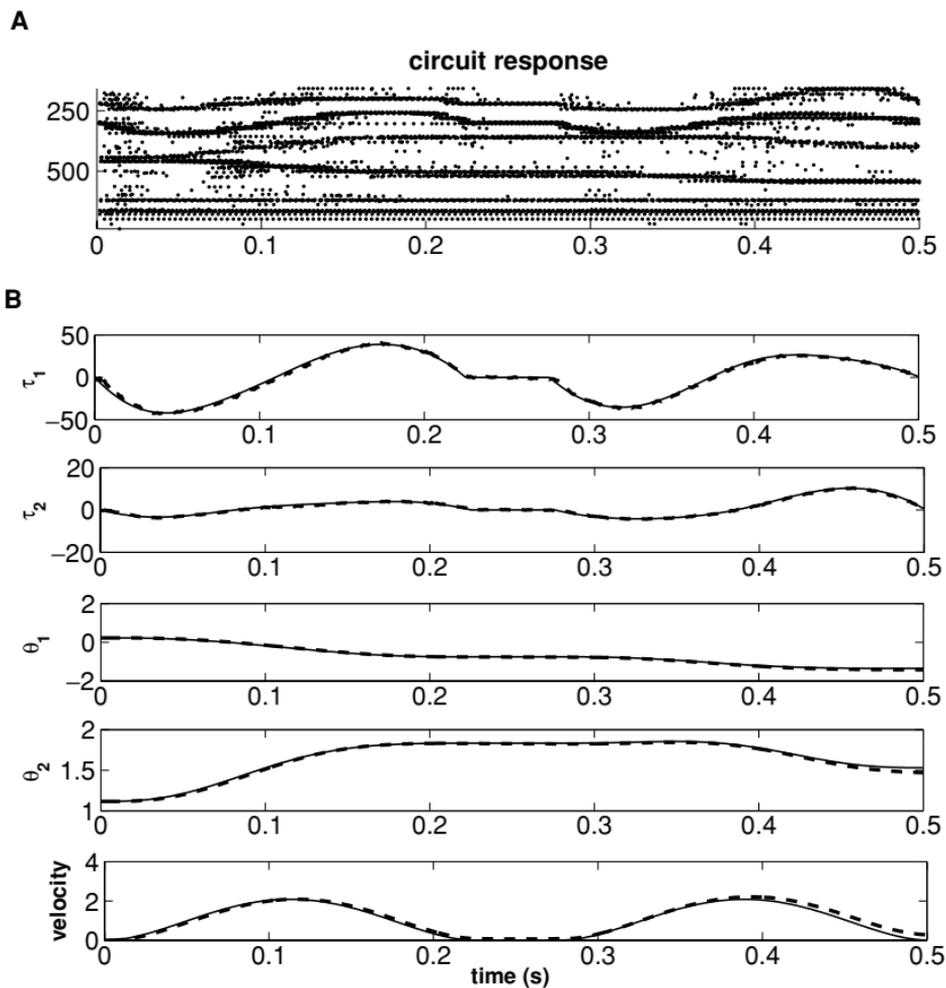
A

B

Figure 6: Demonstration of the temporal integration capability of the neural controller. The data shown are for a validation run for a circuit that has been trained to generate a movement that requires an intermediate stop and then autonomous continuation of the movement after 50 ms. (A) Spike raster of the 600 neurons on the right-hand side of Figure 1. Note that the readout neurons receive at time t only information about the last few spikes before time t (more precisely, they receive at time t the liquid state $x(t)$ of the circuit as their only input). (B) Target time courses of the joint angles θ_1, θ_2, joint torques τ_1, τ_2, and the velocity of the tip of the robot arm are shown as a solid line; actual time courses of these variables during a validation run in closed loop are shown as dashed lines.

on feedbacks as before) was just 6.86 cm, and the bottom part of Figure 6 shows (for a sample test run) that the tip of the robot arm came to a halt during the period from 225 to 275 ms, and then autonomously continued to move.
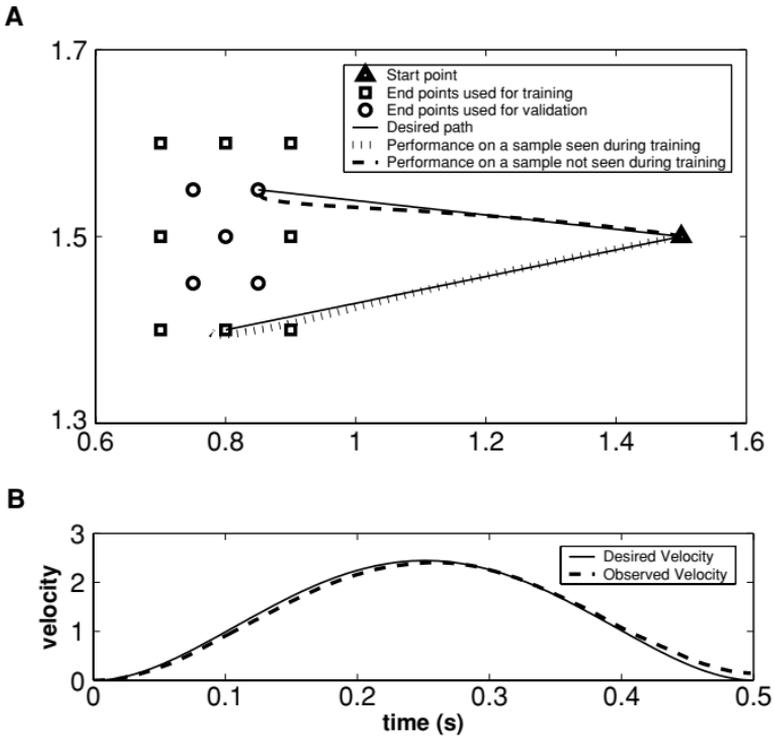
Figure 7: Generation of reaching movements to new target end points that lie between end points used for training. (A) Generalization of movement generation to five target end points (small circles) that were not among the eight target end points (small squares) that occurred during training. Movement to a new target end point was initiated by giving its Cartesian coordinates as constant inputs to the circuit. Average deviation for 15 runs with new target end points: 10.3 cm (4.8 cm for target end points that occurred during training). (B) The velocity profile for one of the movements to a new target end point the solid line is the ideal bell-shaped velocity profile; the actual profile is a dashed line.

5 Generalization Capabilities

The trained neurocontroller (with the weights of the linear readouts being the only parameters that were adjusted during training) had some limited capabilities to generate arm-reaching movements to new targets. For the experiment reported in Figure 7, the circuit was trained to generate from a common initial position reaching movements to eight different target positions, given in terms of their Cartesian coordinates as constant inputs $\langle x_{dest}, y_{dest} \rangle$ to the circuit. After training, the circuit was able to generate with fairly high-precision reaching movements to other target points never used during training, provided that they were located between target points used

for training. The autonomously generated reaching movements moved the tip of the robot arm on a rather straight line with a bell-shaped velocity profile, just as for those reaching movements to targets that were used for training.

6 On the Role of Feedback Delays and Autonomously Generated Feedback Estimates

Our model assumes that the neural circuit receives as inputs, in addition to the constant target end points and efferent copies $\tau_1(t)$, $\tau_2(t)$ of its movement commands with very little delay, proprioceptive or visual feedback that provides at time t information about the values of the angles of the joints at time $t - \Delta$. Whereas it is quite difficult to construct circuits or other artificial controllers for imitating movements that can benefit significantly from feedback (e.g., with the approach of Ijspeert et al., 2003), especially if this feedback is significantly delayed, we show in Figure 8 that neurocontrollers built from generic neural microcircuit models are able to generate and control movements for feedback with a wide range of delays. In fact, Figure 8 shows that the smallest deviation between the target end point $\langle x_{dest}, y_{dest} \rangle$ and the actual end point of the tip of the robot arm is not achieved when this delay Δ has a value of 0, but for a range of delays between 50 and 280 ms. In order to make sure that this surprising result is not an artifact of some particular randomly drawn neural microcircuit model or a particular arm movement, it has been tested on each of 10 randomly drawn neural microcircuits with 12 different movements (four trajectories as shown in Figure 4, each created at three different speeds, resulting in movement times of 300, 500, and 700 ms). The results of these statistical experiments are reported in Figure 8. The right-most point on each of the three curves shows the performance achieved without any feedback (since for this point, the delay of the feedback is as large as the duration of the whole movement). Compared with that, feedback with a suitable delay reduces the imprecision of the movement by at least 50%. Altogether, these data show that the best values for the feedback delay lie in the range of 50 to 280 ms. The upper bound for this interval depends somewhat on the duration of the movement. A possible explanation for the fact that feedback with a delay of less than 50 ms is less helpful is that in this case, the current target circuit output is very similar to the currently arriving feedback, and hence it is more difficult for the circuit to learn the map from current feedback to current target output in a noise-robust fashion. In addition, a delayed feedback complements the inherent temporal integration property of the neural microcircuit model (see Maass, Natschläger, & Markram, 2004), and therefore tends to enlarge the time constant for the fading of memory in the closed-loop system. Hence these neurocontrollers perform best for a range of feedback delays that contain typical values of actual delays for proprioceptive and visual feedback measured in a variety of species (e.g., 120 ms for proprioceptive feedback
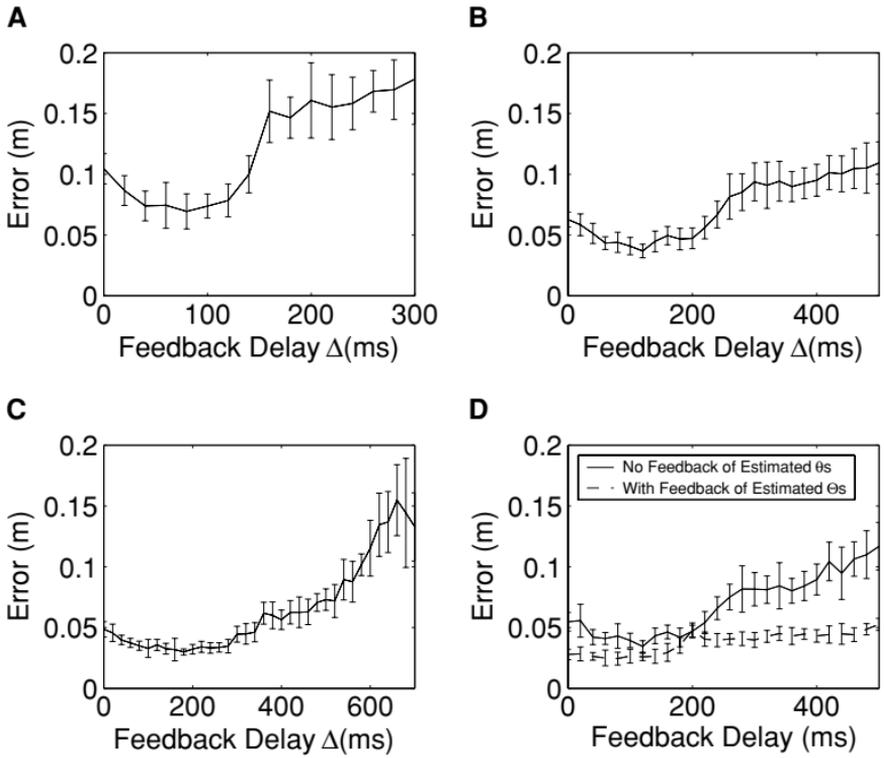
Figure 8: Influence of feedback delay Δ on movement error. Error is defined as the difference of the desired and observed end point of movement. The delay Δ is for proprioceptive feedbacks $\theta_1(t - \Delta)$, $\theta_2(t - \Delta)$. The curves show the averages, and the vertical bars show the SD of the data achieved for 400 movements for each value of Δ (four different movements as shown in Figure 4 repeated 10 different times with different random initial conditions of the circuit and different online noise for each of 10 randomly drawn generic neural microcircuit models). (A–C) Data for three movement durations: 300, 500, and 700 ms. (D) In the upper curve, results for a slightly larger neural circuit (consisting of 800 instead of 600 early integrate-and-fire neurons). The lower (dashed) curve in D shows the performance of the same circuits when internally generated estimates of proprioceptive feedbacks (for a delay of 200 ms) were fed back as additional inputs to the neural circuit. Note that the use of such internally estimated feedbacks not only improves the movement precision for all values of the actual feedback delay Δ expected for $\Delta = 200$ ms, but also reduces the SD of the precision achieved for different circuits considerably.

and 200 ms for visual feedback is reported in van Beers, Baraduc, & Wolpert, 2002).

In another computer experiment, we examined the potential benefit of using estimated feedback for the neurocontroller under consideration. Estimation of feedback is very easy for such neural architecture, since the generic
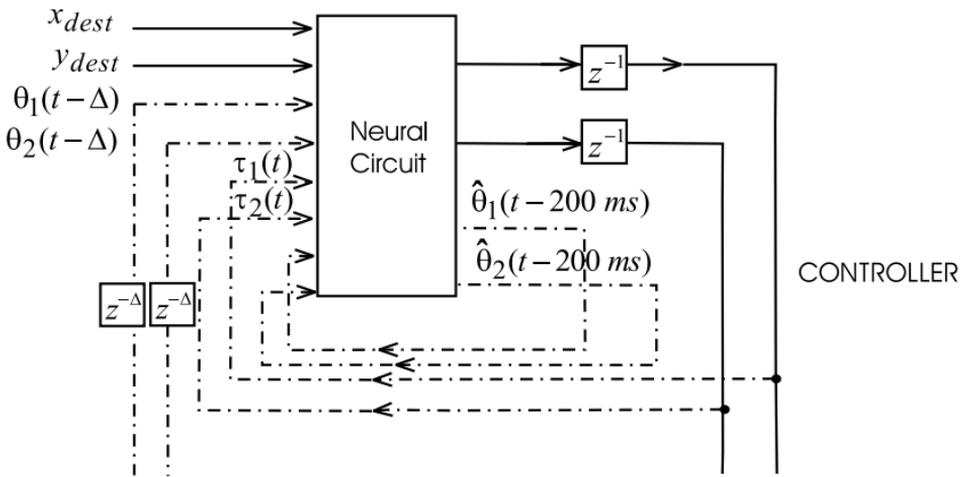
Figure 9: Information flow for the case of autonomously generated estimates $\hat{\theta}(t - 200\ \text{ms})$ of delayed feedback $\theta(t - 200\ \text{ms})$. The rest of the circuit is as in Figure 2.

neural microcircuit model that generates (via suitable readouts) the movement commands has not been specialized in any way for this movement generation task and can simultaneously be used as information reservoir for estimating feedback. More precisely, two additional readouts were added and trained to estimate at any time t the values of the joint angles θ_1 and θ_2 at time $t - 200$ ms, or 200 ms earlier. These delayed values were chosen as targets for these two additional readouts during training, since the previously reported results (see, in particular, Figure 8B) show that feedback of the actual values of θ_1 and θ_2 with a delay of 200 ms is quite beneficial for the precision of the movement that is generated. After training, the weights of these two additional readouts were frozen (for the first two readouts, which produce the movement commands).[6] The outputs of these two additional readouts were also fed back into the circuit (without delay; see Figure 9). Compared with the architecture shown in Figure 2, the neural circuit now receives two additional time-varying inputs. These were fed into the circuit in the same way as the other six inputs (described in section 2). Thus, two additional arrays consisting of 50 neurons each were used for a population coding of these time-varying input variables, and 2 "columns" consisting of

---

[6] Since neither the training of the readouts for movement commands nor the training of the readouts for retrodiction of sensory feedback changes the neural circuit itself, it does not matter whether these readouts are trained sequentially or in parallel. In our experiments both types of readouts were trained simultaneously, while the target values of both $\theta_1(t - \Delta)$, $\theta_2(t - \Delta)$ and $\theta_1(t - 200)$, $\theta_2(t - 200)$ were given to the circuits as additional inputs during training (where Δ is the assumed actual feedback delay plotted on the x-axis in Figure 8D).

100 neurons each were added of the neural circuit that received the outputs of these two additional input arrays.

The top solid line in Figure 8D shows the result, computed in the same way as in the other panels of Figure 8 for the case when the values of the estimates of $\theta_1(t - 200\,\text{ms})$ and $\theta_2(t - 200\,\text{ms})$ produced by the two additional readouts were not fed back into the circuit. The bottom dashed line shows the result when these estimates were available to the circuit via feedback. Although this additional feedback does not provide any new information to the circuit, but only collects and redistributes information within the neural circuit, it significantly improved the performance of the neurocontroller for all values of the actual delay Δ of feedback about the values of θ_1 and θ_2 (except for $\Delta = 200\,\text{ms}$). The value on the right-most point of the lower curve for $\Delta = 500\,\text{ms}$ shows the improvement achieved by using estimated sensory feedback when no feedback arrives at all, since the total movements lasted 500 ms. Altogether, the use of internally estimated feedback improved the precision of the movement by almost 50% for most values of the delay of the actual feedback.

7 Application to a Biological Model for Arm Control

Whereas in the preceding section we focused on a model for a robot arm as a standard example for a highly nonlinear control task, we demonstrate in this section that the same paradigm for movement generation can also be applied to a well-known model for cortical control of arm movements in primates (Todorov, 2000, 2003). This model proposes a direct relationship between the firing rate c_j of individual neuron j in primary motor cortex M1 (relative to some baseline firing rate C) and the kinematics (in cartesian coordinates) and end point force f_{ext} of the hand, which is viewed here simply as the tip of a two-joint arm:

$$c_j(t - d) = \frac{\mathbf{u}_j^T}{2}(F^{-1}\mathbf{f}_{ext}(t) + m\ddot{\mathbf{x}}(t) + k\mathbf{x}(t)) + b\lfloor \mathbf{u}_j^T\dot{\mathbf{x}}(t)\rfloor. \tag{7.1}$$

The vector $\mathbf{u}_j$ denotes the direction in which the end point force is generated due to activation of muscles by neuron j (assuming cosine tuning of neurons). In our simulations we simply took four unit vectors $\mathbf{u}_j$ pointing up, down, left, and right. $\mathbf{f}_{ext}(t)$ is the end point force that the hand applies against external objects. $\mathbf{x}$, $\dot{\mathbf{x}}$, and $\ddot{\mathbf{x}}$ are the position, velocity, and acceleration of the hand, respectively (we usually write $\langle x, y\rangle$ for the hand position x in a two-dimensional space).

Although the precise relationship between the activity of neurons in motor cortex and the activation of individual muscles is extremely complicated and highly nonlinear, a derivation given in in Todorov (2000) suggested that equation 7.1 provides a quite good (almost linear) local approximation to

multijoint kinematics over a small work space. As a consequence, we have applied this model only for arm movements when the hand moves on the boundaries of a 28.28 cm × 28,28 cm square.

Since we are concerned only with the movement of the hand in its work space and do not require the hand to exert an end point force on external world objects, $\mathbf{f}_{ext}(t)$ can be set to 0.

For simplicity, we have also set the transmission delay d from cortex to muscles to 0 (but our model would work just as well for other values of d). This simplifies the model to

$$c_j(t) = \frac{\mathbf{u}_j^T}{2}(m\ddot{\mathbf{x}}(t) + k\mathbf{x}(t)) + b\lfloor \mathbf{u}_j^T \dot{\mathbf{x}}(t)\rfloor. \tag{7.2}$$

In our computer experiments, we applied this model with the parameter values $b = 10$ Ns/m, $k = 50$ N/m, and $m = 1$ kg, suggested in Todorov (2000).

In order to produce a paradigm for arm control by cortical circuits, we took a generic cortical microcircuit model consisting of 800 neurons as described in section 2. Four readouts that received inputs from all neurons in this microcircuit model were trained by linear regression to assume the role of these four neurons in motor cortex that directly control arm muscles resulting in hand movements according to equation 7.2.[7] We trained these readout neurons to produce four different hand movements along the edges of a 28.28 cm × 28.28 cm square whose diagonals were parallel to the x- and y-axis, respectively.

The inputs to the neural microcircuit were the coordinates $\langle x_{dest}, y_{dest} \rangle$ of the desired target end point of the hand, efferent copies of the outputs $c_1(t), \ldots, c_4(t)$ of motor neurons $1, \ldots, 4$, and feedback $x(t - 200)$ ms, $y(t - 200)$ ms about preceding hand positions with a delay of 200 ms that is biologically realistic for visual feedback into motor cortex. The values of these eight inputs were fed into the 800 neuron microcircuit model in the same way as for the eight-input circuit discussed at the end of the preceding section. The results for this experiment are shown in Figure 10. The average deviation over 40 runs of the tip of the arm from the desired end point was 0.13 cm with a SD of 7.2295×10^2 cm.

It is interesting to note that the generic neural microcircuit can also learn to generate movements for this quite different arm model. Another point of interest is that the control performance of the generic neural microcircuit is independent of the kind of feedback that it is receiving (cf. the angles in the earlier model and position coordinates in this model).

---

[7] Target trajectories of the end point of the arm were generated for training as described in note 4. Target outputs $c_j(t)$ for the readouts were generated from these trajectories by equation 7.2.
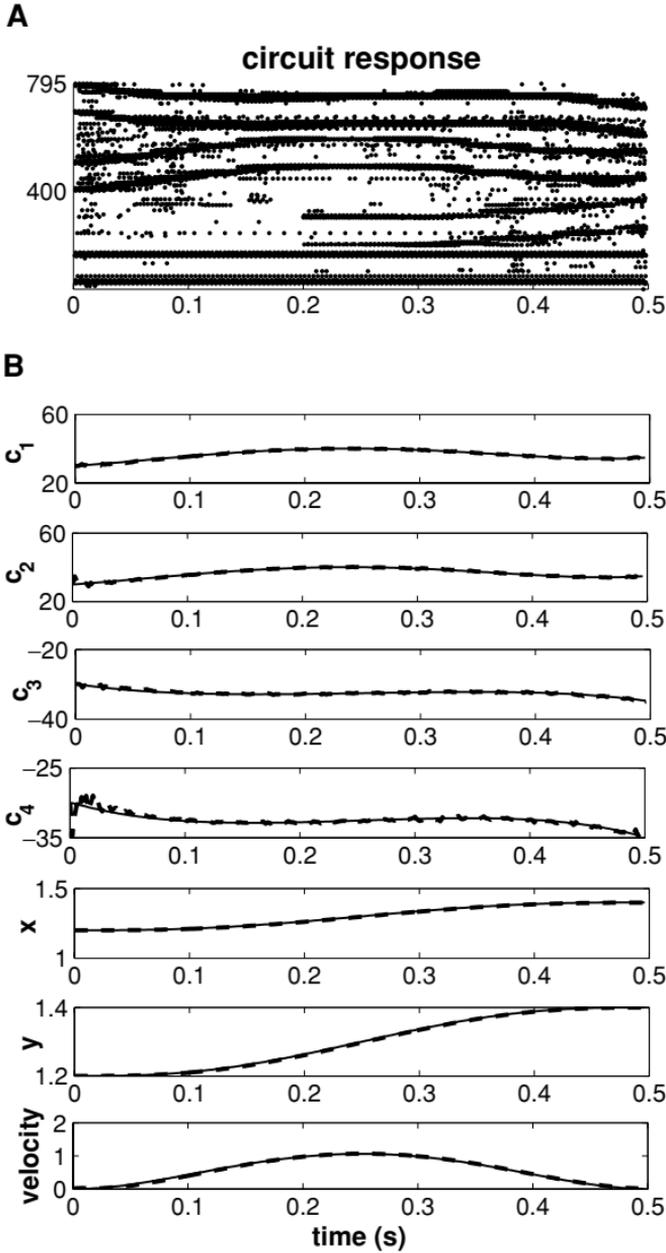
Figure 10: Generation of an arm movement for biological model for cortical control of muscle activations. (A) Spike raster analogous to Figure 6. (B) Solid lines denote target values, and dashed lines show the performance of simulated readouts $c_1, \ldots, c_4$ from a simulated microcircuit in motor cortex that receives significantly delayed information about earlier hand positions as feedback (simulating visual feedback to motor cortex). Scales for $c_1, \ldots, c_4$ in N, for x, y in m, for the velocity of the hand in m/s.

8 How to Make the Movement Generation Noise Robust

Mathematical results from approximation theory (see the appendix of Maass et al., 2002, and Maass & Markram, 2004, for details) imply that a sufficiently large neural microcircuit model (which contains sufficiently diverse dynamic components to satisfy the separation property) can in principle (if combined with suitable static readouts) uniformly approximate any given time-invariant fading memory filter F.

Additional conditions have to be met for successful applications of neural microcircuit models in closed-loop movement generation tasks, such as those considered in this article. First, one has to assume that the approximation target for the neural microcircuit, some movement generator F for a plant P, is a time-invariant fading memory filter (if considered in an open loop). But without additional constraints on the plant or target movement generator F, one cannot guarantee that neural microcircuits L that uniformly approximate F in an open loop can successfully generate similar movements of the plant P. Assume that F can be uniformly approximated by neural microcircuit models L, that is, there exists for every $\varepsilon > 0$ some neural microcircuit model L so that $||(Fu)(t) - (Lu)(t)|| \leq \varepsilon$ for all times t and all input functions $u(\cdot)$ that may enter the movement generator. Note that the feedback f from the plant has to be subsumed by these functions $u(\cdot)$, so that $u(t)$ is in general of the form $u(t) = \langle u_0(t), f(t) \rangle$, where $u_0(t)$ are external movement commands[8] and $f(t)$ is the feedback (both $u_0(t)$ and $f(t)$ are in general multidimensional). Assume that such microcircuit model L has been chosen for some extremely small $\varepsilon > 0$. Even if the plant P has the common bounded input–bounded output (BIBO) property, it may magnify the differences $\leq \varepsilon$ between outputs from F and outputs from L (which may occur even if F and L receive initially the same input u) and produce for these two cases feedback functions $f_F(s)$, $f_L(s)$ whose difference is fairly large. The difference between the outputs of F and L for these different feedbacks $f_F(s)$, $f_L(s)$ as inputs may become much larger than ε, and hence the outputs of F and L with plant P may eventually diverge in this closed loop. This situation does in fact occur in the case of a two-joint arm as plant P. Hence, the assumption that L approximates F uniformly within ε cannot guarantee that $||(Fu_F)(t) - (Lu_L)(t)|| \leq \varepsilon$ for all t (where $u_F(t) := \langle u_0(t), f_F(t) \rangle$ and $u_L(t) := \langle u_0(t), f_L(t) \rangle$), since even $||(Fu_F)(t) - (Fu_L)(t)||$ may already become much larger than ε for sufficiently large t.

This instability problem can be solved by training the readout from the neural circuit L to create an "attractor" around the trajectory generated by F in the noise-free case. This is possible because the current liquid

---

[8] In our experiments $u_0(t)$ was a very simple two-dimensional function with value $\langle 0, 0 \rangle$ for $t < 0$ and value $\langle x_{dest}, y_{dest} \rangle$ for $t \geq 0$. All other external inputs to the circuit were given only during training.

state of the circuit depends not just on the most recent feedback to the circuit, but also on the preceding stream of feedbacks (therefore, the liquid state also contains information about which particular part of the movement has to be currently carried out) as well as on the target end position $\langle x_{dest}, y_{dest}\rangle$. If one trains the readout from circuit L to ensure that $||(Lu_F)(t) - (Lu_L)(t)||$ stays small when $u_F(\cdot)$ and $u_L(\cdot)$ did not differ too much at preceding time steps, one can bound $||(Fu_F)(t) - (Lu_L)(t)||$ by $||(Fu_F)(t) - (Lu_F)(t)|| + ||(Lu_F)(t) - (Lu_L)(t)|| \leq \varepsilon + ||(Lu_F)(t) - (Lu_L)(t)||$ and thereby avoid divergence of the trajectories caused by F and L in the closed-loop system.

This makes clear why it was necessary to train the readouts of the neural microcircuit models L to produce the desired trajectory not just for the ideal feedback $u_F(t)$ but also for noisy variations of $u_F(t) = \langle u_0(t), f_F(t)\rangle$ that represent possible functions $u_L(t)$ that arise if the approximating circuit L is used in the closed loop.

9 Discussion

Whereas traditional models for neural computation focused on constructions of neural implementations of Turing machines or other off-line computational models, more recent results have demonstrated that biologically more realistic neural microcircuit models consisting of spiking neurons and dynamic synapses are well suited for real-time computational tasks (Buonomano & Merzenich, 1995; Maass et al., 2002; Maass, Natschläger, & Markram, 2004; Natschläger & Maass, 2004). Previously, only sensory processing tasks such as speech recognition or visual movement analysis (Buonomano & Merzenich, 1995; Maass et al., 2002; Legenstein, Martram, & Maass, 2003) were considered in this context as benchmark tests for real-time computing. In this letter, we have applied such generic neural microcircuit models for the first time in a biologically more realistic closed-loop setting, where the output of the neural microcircuit model directly influences its future inputs. Obviously closed-loop applications of neural microcircuit models provide a harder computational challenge than open-loop sensory processing, since small imprecisions in their output are likely to be amplified by the plant to yield even larger deviations in the feedback, which is likely to increase even further the imprecision of subsequent movement commands. This problem can be solved by teaching the readout from the neural microcircuit during training to ignore smaller recent deviations reported by feedback, thereby making the target trajectory of output torques an attractor in the resulting closed-loop dynamical system. After training, the learned reaching movements are generated completely autonomously by the neural circuit once it is given the target end position of the tip of the robot arm as (static) input.

We have demonstrated that the capability of the neural circuit to generate reaching movements generalizes to novel target end positions of the tip of

the arm that lie between those that occurred during training (see Figure 7). The velocity profile for these autonomously generated new reaching movements exhibits a bell-shaped velocity profile, as for the previously taught movements. We propose to view the basic arm movements that are generated in this way as possible implementations of muscle synergies, that is, of rather stereotypical movement templates (d'Avella, Saltiel, & Bizzi, 2003). In this interpretation, the learning of a larger variety of arm movements requires superposition of time-shifted versions of several different basic movement templates of the type as are considered in this letter. Such learning on a higher level is a topic of currrent research.

Surprisingly, the performance of the neural microcircuit model for generating movements not only deteriorates if the (simulated) proprioceptive feedback is delayed by more than 280 ms or if no feedback is given at all, but also if this feedback arrives without any delay. Our computer simulations suggest that the best performance of such neurocontrollers is achieved if the feedback arrives with a biologically realistic delay in the range of 50 to 280 ms. If the delay assumes other values or is missing altogether, a significant improvement in the precision of the generated reaching movements can be achieved if additional readouts from the same neural microcircuit models that generate the movements are taught to estimate the values of the feedback with an optimal delay of 200 ms, and if the results of these internally generated feedback estimates are provided as additional inputs to the circuit (see Figure 8D).

Apart from these effects resulting from the interaction of the inherent circuit dynamics with the dynamics of externally or internally generated feedbacks, also the spatial organization of information streams in the simulated neural microcircuit plays a significant role. The capability of such a circuit to generate movement is quite poor if information about slowly varying input variables (such as externally or internally generated feedback) is provided to the circuit in the form of a firing rate of a single neuron (not shown) rather than through population coding (see the description in section 2) as implemented for the experiments reported in this article.

Another interesting point is that our model for motor control can learn to control the arm movement regardless of the model that is used to describe the dynamics of the arm movement and the types of feedback that the circuit is receiving. One of the two arm models that was tested (see section 7) is a model for cortical control of muscle activation. Hence, our model also provides a new hypothesis for the computational function of neural circuits in the motor cortex.

The results presented in this letter may be viewed as a first step toward an exploration of the role of the embodiment of motion generation circuitry, that is, of concrete spatial neural circuits and their inherent temporal dynamics, in motor control. This complements existing work on the relevance of the embodiment of actuators to motor control (Pfeifer, 2002).

Appendix: Specification of Generic Neural Microcircuit Models _____

Neuron parameters: Membrane time constant 30 ms, absolute refractory period 3 ms (excitatory neurons), 2 ms (inhibitory neurons), threshold 15 mV (for a resting membrane potential assumed to be 0), reset voltage drawn uniformly from the interval [13.8, 14.5 mV] for each neuron, constant nonspecific background current I_b uniformly drawn from the interval [13.5 nA, 14.5 nA] for each neuron, noise at each time-step I_{noise} drawn from a gaussian distribution with mean 0 and SD of 1nA, input resistance 1 MΩ.

For each simulation, the initial conditions of each integrate-and-fire neuron, that is, the membrane voltage at time $t = 0$, were drawn randomly (uniform distribution) from the interval [13.5 mV, 14.9 mV]. The probability of a synaptic connection from neuron a to neuron b (as well as that of a synaptic connection from neuron b to neuron a) was defined as $C \cdot \exp(-D^2(a, b)/\lambda^2)$, where $D(a, b)$ is the Euclidean distance between neurons a and b and λ is a parameter that controls both the average number of connections and the average distance between neurons that are synaptically connected (we set $\lambda = 1.2$). Depending on whether the pre- or postsynaptic neurons were excitatory (E) or inhibitory (I), the value of C was set according to Gupta et al. (2000) to 0.3 (EE), 0.2 (EI), 0.4 (IE), 0.1 (II).

We modeled the (short-term) dynamics of synapses according to the model proposed in Markram et al. (1998), with the synaptic parameters U (use), D (time constant for depression), and F (time constant for facilitation) randomly chosen from gaussian distributions that model empirically found data for such connections.

Depending on whether a and b were excitatory (E) or inhibitory (I), the mean values of these three parameters (with D, F expressed in seconds, s) were chosen according to Gupta et al. (2000) to be .5, 1.1, .05 (EE), .05, .125, 1.2 (EI), .25, .7, .02 (IE), .32, .144, .06 (II). The SD of each parameter was chosen to be 50% of its mean. The mean of the scaling parameter A (in nA) was chosen to be 70 (EE), 150 (EI), −47 (IE), −47 (II). In the case of input synapses, the parameter A had a value of 70 nA if projecting onto a excitatory neuron and −47 nA if projecting onto an inhibitory neuron. The SD of the A parameter was chosen to be 70% of its mean and was drawn from a gamma distribution. The postsynaptic current was modeled as an exponential decay $\exp(-t/\tau_s)$ with $\tau_s = 3$ ms ($\tau_s = 6$ ms) for excitatory (inhibitory) synapses. The transmission delays between neurons were chosen uniformly to be 1.5 ms (EE) and 0.8 ms for the other connections.

We applied the following input convention. Each input variable is first scaled into the range [0, 1]. This range is linearly mapped onto an array of 50 symbolic input neurons. At each time step, one of these 50 neurons, whose number $n(t) \in \{1, \ldots, 50\}$, reflects the current value $i_n(t) \in [0, 1]$, which is the normalized value of input variable $i(t)$ (e.g., $n(t) = 1$ if $i_n(t) = 0$, $n(t) = 50$ if $i_n(t) = 1$). The neuron $n(t)$ then outputs at time t the value of $i(t)$. In addition, the three closest neighbors on both sides of neuron $n(t)$ in this

linear array get activated at time t by a scaled-down amount according to a gaussian function (the neuron number n outputs at time step t the value $i(t) \cdot \frac{1}{\sigma\sqrt{2\pi}} e^{\frac{-(n-n(t))^2}{2\pi^2}}$, where $\sigma = 0.8$).

Acknowledgments

References

Buonomano, D. V., & Merzenich, M. M. (1995). Temporal information transformed into a spatial code by a neural network with realistic properties. *Science*, *267*, 1028–1030.

Craig, J. J. (1955). *Introduction to robotics: Mechanics and control* (2nd ed.). Reading, MA: Addison-Wesley.

d'Avella, A., Saltiel, P., & Bizzi, E. (2003). Combination of muscle synergies in the construction of a natural motor behavior. *Nature Neuroscience*, *6*(3), 300–308.

Flash, T., & Hogan, N. (1965). The coordination of arm movements: An experimentally confirmed mathematical model. *Journal of Neuroscience*, *5*(7), 1699–1703.

Gupta, A., Wang, Y., & Markram, H. (2000). Organizing principles for a diversity of GABAergic interneurons and synapses in the neocortex. *Science*, *287*, 273–278.

Ijspeert, A. J., Nakanishi, J., & Schaal, S. (2003). Learning attractor landscapes for learning motor primitives. In S. Becker, S. Thrun, & K. Obermayer (Eds.), *Advances in neural information processing systems*, *15* (pp. 1547–1554). Cambridge, MA: MIT Press.

Jäger, H. (2002). *Tutorial on training recurrent neural networks, covering BPPT, RTRL, EKF and the "echo state network" approach* (GMD Report 159). German National Research Center for Information Technology.

Jäger, H., & Haas, H. (2004). Harnessing nonlinearity: Predicting chaotic systems and saving energy in wireless communication. *Science*, *304*, 78–80.

Joshi, P., & Maass, W. (2004). Movement generation and control with generic neural microcircuits. In A. J. Ijspeert, M. Murata, & N. Wakamiya (Eds.), *Biologically inspired approaches to advanced information technology* (pp. 258–273).

Legenstein, R. A., Markram, H., & Maass, W. (2003). Input prediction and autonomous movement analysis in recurrent circuits of spiking neurons. *Reviews in the Neurosciences (Special Issue on Neuroinformatics of Neural and Artificial Computation)*, *14*(1–2), 5–19. Available at http://www.igi.tugraz.at/maass/publications.html.

Maass, W., Legenstein, R. A., & Bertschinger, N. (2004). Methods for estimating the computational power and generalization capability of neural microcircuits. Submitted for publication. Available online: http://www.igi.tugraz.at/maass/publications.html.

Maass, W., & Markram, H. (2004). On the computational power of recurrent circuits of spiking neurons. *Journal of Computer and System Sciences*, *69*, 593–616. Available online: http://www.igi.tugraz.at/maass/publications.html.

Maass, W., Natschläger, T., & Markram, H. (2002). Real-time computing without stable states: A new framework for neural computation based on perturbations. *Neural Computation*, *14*(11), 2531–2560. Available online: http://www.igi.tugraz.at/maass/publications.html.

Maass, W., Natschläger, T., & Markram, H. (2004). Computational models for generic cortical microcircuits. In J. Feng (Ed.), *Computational neuroscience: A comprehensive approach* (pp. 575–605). Boca Raton, FL: Chapman & Hall/CRC. Available online: http://www.igi.tugraz.at/maass/publications.html.

Markram, H., Wang, Y., & Tsodyks, M. (1998). Differential signaling via the same axon of neocortical pyramidal neurons. *Proc. Natl. Acad. Sci.*, *95*, 5323–5328.

Natschläger, T., & Maass, W. (2004). Information dynamics and emergent computation in recurrent circuits of spiking neurons. In S. Thrun, L. Saul, & B. Schölkopf (Eds.), *Advances in neural information processing systems*, *16* (pp. 1255–1262). Cambridge, MA: MIT Press. Available online: http://www.igi.tugraz.at/maass/publications.html.

Pfeifer, R. (2002). On the role of embodiment in the emergence of cognition: Grey Walter's turtles and beyond. In *Proc. of the Workshop "The Legacy of Grey Walter*." Bristol. Available online: http://www.ifi.unizh.ch/groups/ailab/.

Pouget, A., & Latham, P. E. (2003). Population codes. In M. A., Arbib (Ed.), *The handbook of brain theory and neural networks* (pp. 893–897). Cambridge, MA: MIT Press.

Slotine, J. J. E., & Li, W. (1991). *Applied nonlinear control*. Englewood Cliffs, NJ: Prentice Hall.

Todorov, E. (2000). Direct control of muscle activation in voluntary arm movements: A model. *Nature Neuroscience*, *3*(4), 391–398.

Todorov, E. (2003). On the role of primary motor cortex in arm movement control. In M. L. Latash & M. Levin (Eds.), *Progress in motor control III* (pp. 125–166). Champaign, IL: Human Kinetics.

van Beers, R. J., Baraduc, P., & Wolpert, D. M. (2002). Role of uncertainty in motor control. *Phil. Trans. R. Soc. Lond. B*, *357*, 1137–1145.

Wessberg, J., Stambaugh, C. R., Kralik, J. D., Beck, P. D., Laubach, M., Chapin, J. K., Kim, J., Biggs, S. J., Srinivasan, M. A., & Nicolelis, M. A. (2000). Real-time prediction of hand trajectory by ensembles of cortical neurons in primates. *Nature*, *408*(6810), 361–365.