3. The space inside

In the previous two chapters we have considered how the brain forms maps of the environment as the body moves in it. Now and in the following chapters, we will discuss how the brain forms maps of the body in relation to the external world. As we move and act on the environment, our brain is like a crane operator that must act on a myriad of levers for commanding the motions of the limbs. Even the simplest actions, like reaching for a glass of water, require the coordinated control of about 35 arm and hand muscles. These muscles are divided in smaller compartments, each compartment operated by single motor neuron. In the end there are thousands of little levers that our body-crane operator must act upon successfully for conducting the simple gestures. The body-crane operator receives information about the state of motion of each limb, and about the forces applied by the muscles. There are two geometries for this scenario: the geometry of the body (of its motors and sensors), and the geometry of the world outside. We already encountered this duality in the navigation problem, where the geometry of the images projected on the eyes is different from the geometry of the surrounding space. Here, as well as in the navigation, what matters is that the two geometries are mutually connected.

3.1 Geometry v. Dynamics

Consider the situation depicted in Fig. 3.1A. It describes an experiment carried out by Pietro Morasso (1981) while he was visiting the laboratory of Emilio Bizzi at MIT. The sketch is the top view of a subject holding the handle of a planar mechanism that was designed by Neville Hogan to study the reaching movements of the hand. We can describe the subject's geometry in two alternative ways. We can focus on his hand or we can focus on his arm. In the first case, the position of the hand on the plane is determined by the two coordinates (x, y), of its center, taken with respect to a Cartesian system. An alternative way (not the only one) is to look at the two joint angles, the shoulder angle, ϕ, and the elbow angle θ. These different points of view are related to the different perspective that one may have. Traditional motor physiologists, before Morasso's time, focused on the movements of individual body segments. And, most often, of one body segment at a time, for example the upper arm motions about the elbow. On the other hand, robotics and artificial intelligence placed the focus on the goals of behaviors, for example, moving objects. Morasso came from the school of robotics and artificial intelligence. In his experiments, he paid attention to the motion of the hand.

In the planar configuration of Fig. 3.1A, the hand coordinates (x, y) and the joint coordinates (ϕ,θ) are connected by simple trigonometry:

$$\begin{cases} x = l_1 \cos(\phi) + l_2 \cos(\phi+\theta) \\ y = l_1 \sin(\phi) + l_2 \sin(\phi+\theta) \end{cases} \tag{2.1}$$

where l_1 and l_2 are the lengths of the two arm segments. What is the geometrical interpretation of this relation between angular coordinates of the joints and Cartesian coordinates of the hand? First, it is a non-linear relation. Suppose that we need to move the arm from an initial configuration (ϕ_I, θ_I) to a final configuration (ϕ_F, θ_F) of joint angles. A simple way to do so, is by setting

$$\begin{cases} \phi = \phi_I + \Delta\phi \cdot u(t) \\ \theta = \theta_I + \Delta\theta \cdot u(t) \end{cases} \tag{2.2}$$

Here, $\Delta\phi = \phi_F - \phi_I$, $\Delta\theta = \theta_F - \theta_I$ and $u(t)$ is an arbitrary function of time that varies from 0 to 1 between the start and the end of the movement. In other words, we may drive the two joint synchronously from start to end locations. In this case it is easy to see – by eliminating $u(t)$ - that the two angles are linearly related to each other:

$$\phi = m\theta + q \quad \text{with} \quad \begin{cases} m = \dfrac{\Delta\phi}{\Delta\theta} \\ q = \phi_I - \dfrac{\Delta\phi}{\Delta\theta}\theta_I \end{cases} \tag{2.3}$$

We can apply the same logic to the movements of the hand in Cartesian coordinates and derive a rectilinear motion from $\left(x_I, y_I\right)$ to $\left(x_F, y_F\right)$.

The two rectilinear motions are not generally compatible. This is illustrated in Fig. 3.1B, where we see that straight line in hand coordinates correspond to curved lines in joint coordinates and vice-versa. The diagram in this figure, however, is somewhat misleading. We know that it is appropriate to use Cartesian axes to describe the ordinary space in which the hand moves. We have seen in Chapter 2 that Cartesian coordinates capture the metric properties of Euclidean space and, in particular, the independent concepts of length and angle. If we have two points P, Q, their Cartesian coordinates may change, depending on the origin and on the orientation of the coordinate axes. But the distance computed with the sum-of-squares rule remains invariant.

What can we say about using Cartesian coordinates for the joint angles, ϕ and θ? If angles are represented by real numbers what prevents us from placing these numbers on two lines and call them Cartesian coordinates of an "angle space"? In principle, this could be done (and is routinely done by scientists). However, by doing so, one neglects the critical fact that 0, 2π, 4π, etc. are not really different angles. The numbers are different, but the angles are not. Because of their cyclical character, angular variables are better represented over circles than over lines. Angles have curvature! But then how do we represent two angular variables, as in our case of shoulder and elbow joint coordinates? As shown in Fig. 3.2, the natural geometrical structure for describing a two-link planar mechanism, like a double pendulum or our simplified arm, is a doughnut-shaped object known as a torus.

A torus, like a sphere is a two-dimensional geometrical object. A fundamental theorem by John Nash (1956) establishes that this type of objects can be placed inside an Euclidean space of higher dimension by an operation called an isometric embedding[i] . By this operation, we can represent the two angles of the arm in Fig. 3.1A as coordinates over a torus inside a three-dimensional Euclidean space (Fig. 3.2).

Figure 3.1. Arm kinematics. Top: Sketch of a subject in the experiment by Morasso. The subject holds the handle of a manipulandum that is free to move over the horizontal plane. The subject maintains the arm on the same plane. The manipulandum records two Cartesian coordinates, x and y, that describe the position of the hand with respect to the shoulder. Since the lengths of

the subject's upper arm and forearm are known, we can derive the joint angle of the shoulder (Φ) and of the elbow (ϑ) corresponding to the Cartesian coordinates of the hand. Bottom panels: straight solid lines traced by the subject hand (left) map onto curved solid line in joint angle coordinates (right). Vice versa straight dashed lines in angle coordinates map onto curved dashed line traced by the hand. Note that there is one exception. The line that intersects the subject's shoulder is straight in both coordinate systems. (Modified from Morasso, 1981)

Figure 3.2. Riemannian structure. The configuration space of a two-joint planar mechanism, analogous to the arm in Morasso's experiments, forms a torus (left), a 2-dimensional curved manifold embedded in the Euclidean 3D space, reminiscent of the surface of a doughnut. The light gray mesh forms orthogonal geodesics (minimum path length excursions) spanning the torus. The four panels on the right contain trajectories of the endpoint of the arm's endpoint, marked by the thick black lines. The letters of each panel correspond to the trajectories on the torus.

If we take two points, A and B, there is a unique line on the torus that has the smallest length. This is called a *geodesic* line. The linear joint interpolation of Equation (2.2) defines such a geodesic line. Since the concept of a manifold is a generalization of the concept of a surface, the Euclidean plane is also a manifold, whose geodesics are straight lines. Summing up, the Cartesian coordinates of the hand (x, y) and the angular coordinates (ϕ, θ) describe two geometrical domains with different geodesic properties. If we move the hand between two points over a plane, what type of geodesic do we tend to follow? Morasso addressed this question by asking his subject to move their hand between targets on the plane and he found that movements tend to occur along straight lines- that is along Euclidean geodesics (Fig. 3.3). Tamar Flash and Neville Hogan (1985) interpreted this finding as evidence that the nervous system seeks to optimize the smoothness of hand movements by moving along trajectories that minimize the integrated amplitude of the third time derivative of hand position, or "jerk":

$$\int_0^T \left\{ \left(\frac{d^3x}{dt^3} \right)^2 + \left(\frac{d^3y}{dt^3} \right)^2 \right\} dt \qquad (2.4)$$

Figure 3.3. Reaching movements. Left, hand trajectories between the targets shown on the top panel of Figure 3.1. Left. Joint angles (Top, e: elbow, s: shoulder) and angular velocities (Middle) corresponding to the trajectories labeled c, d and e on the left panel. The traces in the bottom panel are the speed of the hand calculated from the sum of squares of the x- and y- components of hand velocity. Note how the variety of shapes in joint angle profile corresponds to similar bell-shaped curves for the hand speed. Also note that movements D and E had reversals of joint motions. Movement C is the only one that does not show any reversal. This is consistent with the observation (see Fig. 3.1) that straight lines of the hand that pass by the shoulder joint map to straight lines in joint coordinates. (From Morasso, 1981)

The idea that there is a separation between the geometry of movement – the "kinematics" – and its dynamical underpinnings did not go unchallenged and is still an object of controversy. Yoji Uno, Mitsuo

Kawato and Suzuki (1989) suggested that, instead of optimizing the smoothness of hand motions, the nervous system was concerned with minimizing variables that are more directly connected with effort. Then, they considered, as cost function, the integral of the net torque-changes over each movement:

$$\int_0^T \sum_{i=1}^n \left(\frac{d\tau_i}{dt} \right)^2 dt \tag{2.5}$$

The term τ_i represents the joint torque produced by one of the *n* arm muscles. One can intuitively see the similarity between the two formulations. Consider a point mass. By Newton's law, the point mass accelerates in proportion to the applied force. If one takes one more time derivative, one obtains a relation between the jerk and the rate of change of the force. This explains why the two optimization criteria yield relatively similar predictions of arm trajectories. There is however an important difference between kinematic and dynamic criteria. The outcome of the dynamical optimization depends on properties such as the inertia of the moving arm. The minimization of the cost integral (2.5) requires knowing the relation between forces and motions, as expressed by the limb's dynamical equations. These equations are quite complex and grow in complexity as more degrees of freedom are introduced. The reader can find explicit forms for two-joint dynamics in standard robotics handbooks. Here, we do not want to get into excessive detail and we limit ourselves the describe dynamics as a second-order non-linear ordinary differential equation (ODE), whose general form[ii] is

$$\mathbf{D}(\mathbf{q}, \dot{\mathbf{q}}, \ddot{\mathbf{q}}) = \mathbf{Q}(t) \tag{2.6}$$

The generalized coordinate $\mathbf{q} = [\phi, \theta]^T$ is the configuration of the arm and the corresponding generalized force is a joint-torque vector $\mathbf{Q} = [\tau_\phi, \tau_\theta]^T$. Thus, with the vector notation of Equation (2.6) we are representing a system of two coupled non-linear ODEs . The optimization by Uno and colleagues requires us to calculate the torque vector based on these ODEs. On the other hand the kinematic approach of Flash and Hogan is strictly geometrical and does not involve forces (i.e., dynamics). Unlike dynamic criteria, the geometrical optimization predicts that neuromuscular activities are coordinated by the brain to generate a rectilinear motion of the hand or of whatever other element plays the role of a controlled "endpoint element" and this result will not be affected by changes in the dynamical structure controlled by the muscles.

Randy Flanagan and Ashwini Rao (1995) performed a simple and elegant variant of Morasso's experiment. They asked a group of volunteers to control a cursor by moving the hand over a plane. However, in this case, the position of the cursor was not an image of the position of the subjects' hand but was in direct relation to their shoulder and elbow angle. This task was difficult and somewhat confusing for the subjects. However after some practice, they learned to move the cursor onto the targets. Most remarkably, as they learned to do so, they gradually but consistently and spontaneously learned to produce rectilinear movements of the cursor, at the expenses of more curved motions of the hand (Fig. 3.4). This result is incompatible with the notion that the nervous system attempts to minimize a dynamic criterion, such as the minimum-torque-change of Eqn. (2.5).

Figure 3.4. Subjects move in straight lines over the visual space. Motion paths represented in hand (left) and joint (right) coordinates for movements between targets presented in either hand

coordinates (top) or joint coordinates (bottom). In both cases, subjects learned to organize coordination so as to produce straight line motions path in the observed space. (from Flanagan and Rao, 1995).

3.2 Does the brain compute dynamics equations?

Until the late 1970s, the study of neurobiology was mostly limited to the control of single muscles or of muscle pairs acting on a single joint. A large volume of work was dedicated to the control of eye movements, where six extrinsic muscles are responsible to move the eye and to control the direction of the gaze. The engineering frameworks for oculomotor control and for the control of limb movements derived from the theory of linear control systems and of feedback control in particular. Computational models were abundant in "box diagrams" such as the one in Fig. 3.5. Robotics research and, in particular, the theory of robot manipulators brought to the neurobiology of motor control a new awareness of the dynamical complexity of multiarticular structures such as the human arm. Let us compare the inertial dynamics of the upper arm about the elbow joint and the inertial dynamics of the two-joint arm of Fig. 3.1. The first are described by the equation

$$I\ddot{q} = Q\,. \tag{2.7}$$

Figure 3.5. Block diagram based on control engineering to describe the operation of the oculomotor system in response to a visual target. (From David Zee, Lance Optican , Jay Cook, David Robinson and King Engel, 1976)

This is a straightforward translation of $ma = F$ in angular terms. Here, I , is the moment of inertia of the forearm about the axis of rotation of the elbow, $\ddot{q}$ is the angular acceleration about the same axis and Q is the corresponding joint torque vector. In the two-joint case, we have two angles that we call q_1 and q_2 (instead of ϕ and θ, for notational convenience) and two corresponding torques, Q_1 and Q_2. The inertial dynamics look like this:

$$\begin{aligned}
&\left(I_1 + I_2 + m_2 l_1 l_2 \cos(q_2) + \frac{m_1 l_1^2 + m_2 l_2^2}{4} + m_2 l_1^2\right)\ddot{q}_1 + \left(I_2 + m_2 \frac{l_1 l_2}{2}\cos(q_2) + m_2 \frac{l_2^2}{4}\right)\ddot{q}_2 + \\
&\qquad - \left(m_2 \frac{l_1 l_2}{2}\sin(q_2)\right)\dot{q}_2^2 - \left(m_2 l_1 l_2 \sin(q_2)\right)\dot{q}_1\dot{q}_2 = Q_1 \\
&\left(I_2 + m_2 \frac{l_1 l_2}{2}\cos(q_2) + m_2 \frac{l_2^2}{4}\right)\ddot{q}_1 + \left(I_2 + m_2 \frac{l_2^2}{4}\right)\ddot{q}_2 + \left(m_2 \frac{l_1 l_2}{2}\sin(q_2)\right)\dot{q}_1^2 = Q_2
\end{aligned} \tag{2.8}$$

The reader who is not acquainted with this equation can find its derivation at www.shadmehrlab.org/book/dynamics.pdf. Here, we wish to focus the attention not on the detailed form of these equations but on a few relevant features. First, note how the transition from Eq. (2.7) to Eq. (2.8) has led to a disproportionate increase in size. In the first we have a single multiplication, whereas in the second one counts 15 additions, 42 multiplications and a couple of trigonometric functions. It is not a pretty sight. If we were to write the detailed equations for the entire arm in 3D, we would fill a few pages of text. To put some logical order in these equations, we may start by noticing that some terms contain accelerations and some contain velocities. The terms that contain acceleration are the first two in each equation. They may be written in vector-matrix form as the left end term of Eq. (2.7), i.e.

$$\begin{bmatrix} I_1 + I_2 + m_2 l_1 l_2 \cos(q_2) + \dfrac{m_1 l_1^2 + m_2 l_2^2}{4} + m_2 l_1^2 & I_2 + m_2 \dfrac{l_1 l_2}{2}\cos(q_2) + m_2 \dfrac{l_2^2}{4} \\ I_2 + m_2 \dfrac{l_1 l_2}{2}\cos(q_2) + m_2 \dfrac{l_2^2}{4} & I_2 + m_2 \dfrac{l_2^2}{4} \end{bmatrix} \begin{bmatrix} \ddot{q}_1 \\ \ddot{q}_2 \end{bmatrix} = \mathbf{I}(\mathbf{q})\ddot{\mathbf{q}}$$

This is called the inertial term. The acceleration appears linearly, as it is typical of all mechanical systems derived from Newton's equation. However, the inertia matrix I depends on arm configuration, unlike the inertial term of Eq. (2.7), which is constant. In addition to this term, the two-link arm equation also contains terms that depend on the angular velocities. These are the "centripetal" terms, depending on the squares of the velocities and the "Coriolis" term depending on the product of the velocities. We may collect these terms in a single vector $\mathbf{G}(\mathbf{q},\dot{\mathbf{q}})$ and the dynamical equation assumes the much less intimidating form

$$\mathbf{I}(\mathbf{q})\ddot{\mathbf{q}} + \mathbf{G}(\mathbf{q},\dot{\mathbf{q}}) = \mathbf{Q} \tag{2.9}$$

Figure 3.6. The dynamics of a reaching movement. A: Shoulder and elbow angle traces corresponding to the hand movement in B. Different components of the shoulder (C) and elbow (D) torque were derived by applying the equation of motions to the observed kinematics. Solid lines: net torque at each joint. Dotted line: shoulder inertial torque. Dot-dashed line: elbow inertial torque. Dashes: centripetal torques. Two-dots with a dash: Coriolis torque at the shoulder. (From Hollerbach and Flash, 1982)

While this simple algebraic representation provides a valid description for more complex structures, such as a typical robotic manipulator or the whole human arm, it highlights the most relevant difference between the dynamics of a single joint and the dynamics of multiarticular limb and points to some computational challenges. One such challenge stems from the dependence of the torque on one joint from the state of motion of another. These dynamics cannot be simplified by treating each dimension independently of the others. How important are these cross-dependencies in natural arm movements? Equations (2.8) provide us with a simple and direct way to answer this question. It is sufficient to record the trajectory of a reaching arm movement, derive the corresponding angular motion of the shoulder and elbow joints, take the first and second derivatives and plug these data in the equations. John Hollerbach and Tamar Flash (1982) performed this experiment, using the same apparatus and task of Morasso. When they calculated the different contributions to the net torques at the shoulder and elbow, they were able to conclude that the interaction components that are summarized by the non-linear term $\mathbf{G}(\mathbf{q},\dot{\mathbf{q}})$ are quite substantial compared to the pure inertial terms. As an example, consider the trajectory in Fig. 3.6A and B, which is similar to the movement labeled as "e" in Fig. 3.3. The movement time is about half a second, corresponding to a relatively rapid but not unusual speed. The different components of the torque profiles at the shoulder and elbow joint are plotted in Fig 3.6C and D respectively. Note the amount of toque that is due to centripetal and Coriolis terms, and also the amount of inertial torque at the shoulder, due to the acceleration at the elbow and vice versa. Most importantly, the same motion of the hand between two different targets would lead to qualitatively different torque profiles. This evidence was sufficient to conclude that arm movements cannot be controlled by some linear feedback mechanism, such as the one depicted in Fig. 3.5 for the oculomotor system.

3.3 The engineering approach

Let us briefly switch our perspective from neuroscience to robotics. Control engineers are primarily concerned with the concept of stability. Loosely speaking, a system is stable if, given some unexpected perturbation it will eventually go back to its unperturbed state. However, "eventually" is not strong enough of a requirement for practical applications. A stronger form of stability requires the system to converge exponentially in time to the unperturbed state. One of the objectives, if not the main objective, of feedback control is to insure stability against uncertainties that arise from limited knowledge of the environment and of the controlled "plant". Importantly, robotic engineers have some good prior knowledge of their own devices. So, for example, they would know to good accuracy the parameters in Equation (2.9). Given this knowledge, we now want to write the right-hand term, the torque vector, as the driving control input to the arm. As we are dealing with feedback control, we think of $\mathbf{Q}$ as a function of the state of motion – i.e. the vector $[\mathbf{q},\dot{\mathbf{q}}]^T$ - and of time. A physical system whose dynamics equations do not depend upon time is said to be *autonomous*. For such system the future is entirely determined by a measure of the state at some initial instant. For example, the equation of a simple oscillator, like a mass attached to a spring is

$$m\ddot{q} = -kq \tag{2.10}$$

The solution of this equation is

$$q = A\cos\left(\left(\sqrt{\frac{k}{m}}\right)t + \phi\right) \tag{2.11}$$

The phase and the amplitude are determined by the initial position $q_0 = q(t_0)$ and velocity $v_0 = \dot{q}(t_0)$. Therefore, by setting (or measuring) these initial conditions, we know how the oscillator will behave for all future times. However, if we add an external input, as in

$$m\ddot{q} = -kq + u(t) \tag{2.12}$$

then, the trajectory of the system is no longer determined by the initial state alone and depends upon the future values of $u(t)$ as well[iii].

Now, consider the stability problem. Suppose that we have a robotic arm, governed by a feedback controller that attempts to track a desired trajectory, $\hat{\mathbf{q}}(t)$. Then, the combined system+controller equation takes the form

$$\mathbf{I}(\mathbf{q})\ddot{\mathbf{q}} + \mathbf{G}(\mathbf{q},\dot{\mathbf{q}}) = \mathbf{Q}(\mathbf{q},\dot{\mathbf{q}},\hat{\mathbf{q}}(\mathbf{t})) \tag{2.13}$$

Our goal is to design the control function $\mathbf{Q}(\mathbf{q},\dot{\mathbf{q}},\hat{\mathbf{q}}(\mathbf{t}))$ so that the resulting movement is exponentially stable about the desired trajectory. As we stated before, a robotics engineer would have a good model of the controlled mechanical arm. The model would contain the two terms

$$\begin{gathered}\hat{\mathbf{I}}(\mathbf{q}) \approx \mathbf{I}(\mathbf{q}) \text{ and} \\ \hat{\mathbf{G}}(\mathbf{q},\dot{\mathbf{q}}) \approx \mathbf{G}(\mathbf{q},\dot{\mathbf{q}})\end{gathered} \tag{2.14}$$

We can use these terms to design the control function as

$$\mathbf{Q}(\mathbf{q},\dot{\mathbf{q}},\hat{\mathbf{q}}(\mathbf{t})) = \hat{\mathbf{I}}(\mathbf{q})\cdot\mathbf{a}(\mathbf{q},\dot{\mathbf{q}},\mathbf{t}) + \hat{\mathbf{G}}(\mathbf{q},\dot{\mathbf{q}}) \tag{2.15}$$

Substituting this in Equation (2.13) with the assumptions of Equation (2.14) we get the simple double-integrator system:

$$\ddot{\mathbf{q}} \approx \mathbf{a}(\mathbf{q},\dot{\mathbf{q}},t) \tag{2.16}$$

It is important to observe that the symbol $\approx$ is not the ordinary equal sign, but an approximate equality. We will come back to this in a little while but for the next few lines we will use the standard equality. We begin by giving a particular form to the term $\mathbf{a}(\mathbf{q},\dot{\mathbf{q}},t)$. Starting from the desired trajectory, we calculate the desired velocity, $\dot{\hat{\mathrm{q}}}(t)$, and acceleration, $\ddot{\hat{\mathrm{q}}}(t)$, by taking the first and second time derivatives of $\hat{\mathrm{q}}(t)$. Then, we set

$$\mathbf{a}(\mathbf{q},\dot{\mathbf{q}},t) = \ddot{\hat{\mathbf{q}}}(t) - \mathbf{K_P}(\mathbf{q} - \hat{\mathbf{q}}_d(t)) - \mathbf{K_D}(\dot{\mathbf{q}} - \dot{\hat{\mathbf{q}}}_\mathbf{d}(t)) \tag{2.17}$$

and

$$\boldsymbol{\eta}(\mathbf{q},t) \equiv \mathbf{q} - \hat{\mathbf{q}}(t) \tag{2.18}$$

The last equation defines the *tracking error* as a function of the current position and time. The two matrices, $\mathbf{K}_P$ and $\mathbf{K}_D$ contain gain coefficients that multiply the position and velocity errors[iv]. By combining Equations (2.18), (2.17) and (2.16) we finally obtain a linear second order differential equation for the error function:

$$\ddot{\boldsymbol{\eta}} + \mathbf{K_D}\dot{\boldsymbol{\eta}} + \mathbf{K_P}\boldsymbol{\eta} = 0. \tag{2.19}$$

From calculus we know that by setting

$$\begin{aligned}\mathbf{K_P} &= \begin{bmatrix}\omega_1^2 & 0 & \cdots & 0 \\ 0 & \omega_2^2 & \cdots & 0 \\ \cdots & \cdots & \cdots & \cdots \\ 0 & 0 & \cdots & \omega_N^2\end{bmatrix} = diag\left(\omega_1^2,\omega_2^2,\ldots,\omega_N^2\right) \\ \mathbf{K_D} &= 2diag\left(\omega_1,\omega_1,\ldots,\omega_N\right)\end{aligned} \tag{2.20}$$

we can re-write Equation (2.19) as N decoupled equations

$$\ddot{\eta}_i + 2\omega_i \dot{\eta}_i + \omega_i^2 = 0, \tag{2.21}$$

whose general solution is

$$\eta_i = (c_{1,i} + c_{2,i} t) e^{-\omega_i t}. \tag{2.22}$$

Therefore, with the appropriate choice of the two gain matrices, we can insure that the tracking error will go to zero in exponential time. This type of dynamic stability was observed in reaching movements by Justin Won and Neville Hogan (1995). They asked subjects to execute reaching movements of the hand while holding a planar manipulandum (Fig. 3.7). Most of the times, the manipulandum allowed for free motion of the hand in the plane and the subjects moved along straight lines from start to end target. In some random trials, one of the joints of the manipulandum was blocked at the onset of the movement, by an electromagnetic brake. As a consequence in these test trials the movement of the hand was initially constrained to remain on a circular arc (dotted line). Shortly after start, the brake was released. The resulting motion of the hand converged rapidly to the unperturbed path, thus demonstrating a dynamic stability of the planned movement trajectory, similar to the exponential stability discussed above.

But now remember that in Equations (2.21) and (2.22) instead of = we should have used $\approx$ as the model of the arm dynamics can at most approximate the actual value of the parameters. However, for small deviations, the uncertainty of the model can be regarded as an internal perturbation that does not critically compromise stability. The essential point is that if we have an internal representation of the dynamics, a control system can generate behaviors that are stable, in the face of uncertainties concerning the environment, and the controlled structure.

Figure 3.7. Dynamic stability. Right: Subjects practiced moving the hand from A to B while holding the handle of a planar manipulandum. In most trials the manipulandum could be moved freely over the plane and the subjects moved on straight paths (dashed line). In some trials, however, electromagnetic brake blocked the motion of one manipulandum joint at the onset of the reaching movement. In these perturbed trials, the hand was constrained to remain on a circular path (dotted line) until brake was suddenly released. Right. Constrained and released hand trajectories. The arrow indicates approximately where the release took place. The hand moved almost instantaneously toward the planned straight trajectory from A to B. (From Won and Hogan, 1995)

3.4 Does the brain represent force?

The idea of model-based control was born as an engineering idea. What can it tell us about the brain? One of the fathers of computer technology, the mathematician John von Neumann, was driven by the desire to create an artificial brain. And some of the earliest computational neuroscientists, like Warren Mc Culloch and Walter Pitts, described neurons as logical gates. But it is hard to capture the intricacies of biology within the confines of engineering and mathematics. And the idea that the brain computes something like the dynamics equations for the limb it controls has encountered – and still encounters, to this day - some strong resistance. One of the strongest counterarguments to the idea of explicit computations comes from the observation that ordinary physical systems behave in ways that may be described by equations but do not require solving equations with calculus-like rules. If you drop a stone, the stone will follow a simple and regular trajectory that you can derive by integrating the equations of motion in the gravitational field. Yet the stone does not compute integrals! In the same vein, an important concept in science is to seek for principles that can explain observations based on simple rules of causality.

Consider another example. Take the two opposing springs and dampers connected to a mass (Fig. 3.8A). These are physical elements that generate forces acting on the point mass. The two springs define a static equilibrium point. If an external agent places the point mass away from this equilibrium, the physical system will eventually bring it back there, because the equilibrium is stable. This behavior can be readily implemented by analog or digital elements (Fig. 3.8B) that produce forces in response to measured position and velocity of the point mass. Thus, we have a physical system and a computational system (which is nothing but another type of physical system!) that do the same job. In 1978 Andres Polit and Emilio Bizzi published a finding that physiologists of that time found hard to digest. They trained monkeys to reach a visual target by a movement of the upper arm about the elbow. The monkeys could not see their arm as it was moving under an opaque surface and were rewarded when they pointed correctly to the visual target. Not surprisingly, the monkey learned the task. However, the monkeys were able to perform the task after they were surgically deprived of all sensory information coming from their arm. And they successfully reached the targets even if the arm was unexpectedly displaced by a motor before starting the movement. If one were to think of movement control according to the scheme drawn in Fig. 3.8B, this result would not make any sense, because now the lines carrying the information about position and velocity were cut! But, the scheme of Fig. 3.8A is perfectly consistent with the finding of Polit and Bizzi: if the muscles act upon the arm, as opposing viscoelastic elements, then all what the nervous system has to do is to set their equilibrium point to the target. No need to monitor the position and state of the arm in order to correct for errors. Instead, in order to move a limb the brain can set the equilibrium point so as to follow a desired movement trajectory.

This idea had been pioneered in 1966 by Anatol Feldman. Then, as well as in more recent works, Feldman proposed that we move our bodies by setting the equilibrium postures specified by reflexes and muscle mechanical properties. Does this make dynamic computations unnecessary? Answering this question is not quite as simple as it seems.

Figure 3.8. Spring-like behavior in physical (A) and computational (B) systems. The forces generated by springs on a mass are reproduced by a control systems that multiplies by a gain (K) the difference between the measured position of the mass (M) and a commanded equilibrium position (u). The output of the K is then added to another feedback signal proportional to the measured velocity of the mass and the resulting force is applied to M. The presence of delays along the feedback lines in the system B may cause unstable oscillations that are not observed in A.

One may ask if to move your arm your brain must represent forces, as defined by Newton. Consider the equation that represents the behavior of a spring-mass-damping system, such as the one shown in Fig 3.8:

$$M\ddot{x}+B\dot{x}+K(x-u(t))=0 \tag{2.23}$$

We derive this by substituting the expression of the forces generated by the spring and damper in $F=m\ddot{x}$. As a result the equation does not contain force at all. It has only the state (position and velocity) its derivative (acceleration) and an external input, $u(t)$, that sets instant by instant the equilibrium point of the springs. The situation is not different from the one described in the previous paragraph to derive the controller for a robotic arm. We may not need to represent forces to derive a trajectory produced by the input function, but we need to know how the system reacts to applied forces. These are the terms that appear in Equation (2.23). This equation allows us to derive the trajectory $x(t)$ that results from the commanded equilibrium-point trajectory $u(t)$. And the same equation allows us to derive the equilibrium-point trajectory from the desired movement. The only condition for this to be possible is $K\neq 0$. Then, given a desired trajectory $\hat{x}(t)$ with first and second time derivatives $\dot{\hat{x}}(t)$ and $\ddot{\hat{x}}(t)$ the commanded equilibrium point-trajectory is

$$u(t)=K^{-1}M\ddot{\hat{x}}(t)+K^{-1}B\dot{\hat{x}}(t)+\hat{x}(t) \tag{2.24}$$

Figure 3.9. Movement simulations. A and C: Movements of the arm obtained by shifting the static equilibrium point from left to right along a straight line with two different speeds. The equilibrium point motions follow a bell-shaped speed profile with durations of 2 seconds (slow) and .8 seconds (fast). B and D: movements of the equilibrium points that would produce a straight line motion of the hand with the slow and fast speed profiles.

This is not substantially different from what we have done in deriving the engineering controller of Equation (2.17). The only new element here is the focus on representing the control signal as a moving equilibrium point that drives the controlled "plant" along the desired trajectory. Intuitively, if K is sufficiently large and if the motion is smooth and slow, the first two terms in Equation (2.24) may be small enough to be ignored and then all what one needs to do is to move the equilibrium point along the desired trajectory. This is the attractive promise of equilibrium-point control: there is no need to compute dynamics. The mechanical properties of the body and, in particular, the viscoelastic properties – responsible for the K and B terms - of the neuromuscular system may provide a way to avoid complex computations altogether. Interestingly, this viewpoint emerged vigorously at the same time in which the robotics viewpoint on inverse dynamic computations was becoming influential in the neuroscience of the motor system. But is it true that dynamics computations may be avoided?

Consider the simple movement of the right arm illustrated schematically in Fig. 3.9. This is a simulation of a simplified model of the arm, based on Equation (2.8). The numerical parameters are listed in the box. We re-write the dynamics equations to include the viscoelastic terms and make it formally similar to Equation (2.23):

$$\mathbf{I}(\mathbf{q})\ddot{\mathbf{q}}+\mathbf{G}(\mathbf{q},\dot{\mathbf{q}})+\mathbf{B}\dot{\mathbf{q}}+\mathbf{K}(\mathbf{q}-\mathbf{u}(t))=0 \tag{2.25}$$

An appealing idea is to simply move the equilibrium point along the desired trajectory of the hand. The panels on the left illustrate the outcome of this approach. If the movement is sufficiently slow, for example if one takes two seconds to move the equilibrium point about 40 cm across, the hand moves in a nearly rectilinear path. However the deviation from the desired trajectory increases as the planned motion becomes faster (0.8 sec) and when the movement is more proximal to the body (Fig. 3.9C). Then, we may ask what would it take to change the equilibrium trajectory, $\mathbf{u}(t)$, in such a way to obtain the desired straight motion of the hand. We can certainly move our hand in a nearly straight 40 cm line within less than one second. To find the equilibrium point trajectory that achieves this goal, we begin by transforming the desired rectilinear trajectory of the hand, $\mathbf{r}(t) = \left[x(t), y(t)\right]^T$ into the trajectories of the shoulder and elbow joints, $\mathbf{q}(t) = \left[q_1(t), q_2(t)\right]^T$ with their first and second time derivatives. Then, we solve Equation (2.25)algebraically to obtain

$$\mathbf{u}(t) = \mathbf{K}^{-1}\left(\mathbf{I}(\mathbf{q}(t))\ddot{\mathbf{q}}(t) + \mathbf{G}(\mathbf{q}(t), \dot{\mathbf{q}}(t)) + \mathbf{B}\dot{\mathbf{q}}(t)\right) + \mathbf{q}(t) \tag{2.26}$$

The results of this simulation are shown in the left panels, B and D for the distal and proximal arm configurations. Only a small correction of the equilibrium point trajectory is sufficient to “straighten” the movement at low speed (2 sec). However, as the movement becomes faster (0.8 sec) the equilibrium point must take a quite distorted path in order to keep the hand on track. It is evident that we could not derive this complex equilibrium trajectory by performing a simple linear manipulation on the errors – shown on the left panels - that we would make by moving the equilibrium on the straight path. Here, to derive the motion of the equilibrium point that produces a desired actual movement of the hand we have explicitly solved an inverse dynamics problem. Our brain is likely using some different method. But, as we are indeed capable to move our limbs with dexterity, this method must effectively solve the same inverse dynamics problem.

3.5 Adapting to predictable forces

So far, we have described the basic computational idea of an internal model. Is there a more stringent empirical basis for this idea? Perhaps the first empirical evidence came from an experiment that the authors of this book performed in the early 90s, while we both were in the laboratory of Emilio Bizzi at MIT. We asked volunteers to move their arm while holding the handle of a planar manipulandum (Fig. 3.10A). They performed reaching movements of the hand, as in Morasso’s experiments. The robot, designed by Neville Hogan, was a light weight two-joint device, equipped with position sensors and with two torque motors. The motors were “backdrivable”, meaning that when they were turned off, they generated only minimal resistance to motion. And the manipulator itself was made of light aluminum bars, so that the whole thing did not impede appreciably the free motion of the hand.

Figure 3.10. Learning to move in a force field. A: A subject holds the handle of a manipulandum. The computer monitor displays a cursor indicating the position the hand and presents targets to the subjects. Two torque motors are programmed to apply forces that depend upon the velocity of the handle, as shown in B. C: When the motors are inactive, subjects perform straight reaching movements to the targets. D: In random trials the force field shown in C is unexpectedly activated. The movement trajectories are strongly perturbed and display a pattern of characteristic hooks. At the end, the movement terminate in the proximity of the targets. (From Shadmehr and Mussa-Ivaldi, 1994)

Therefore, with the motors turned off the hand trajectories were rectilinear and smooth (Fig. 3.10B). After acquiring unperturbed movement data, the motors were programmed to generate a force, $\mathbf{F}$, that depended linearly upon the velocity of the hand, $\dot{\mathbf{r}}$:

$$\mathbf{F} = \mathbf{B} \cdot \dot{\mathbf{r}} \tag{2.27}$$

This force field was equivalent to a controlled change in the dynamics of the subjects' arm. During the initial phase of the experiment, the subjects dynamics were as in Equation (2.13). Without making any assumption about the control function, $\mathbf{Q}(\mathbf{q},\dot{\mathbf{q}},t)$ that a subject used to perform a particular trajectory $\mathbf{q}_{\mathbf{A}}(t)$ the dynamics equations are reduced to an algebraic identity over this trajectory. If the force field (2.27) is introduced without changing the control function, then the new limb dynamics are

$$\mathbf{I}(\mathbf{q}) \cdot \ddot{\mathbf{q}} + \mathbf{G}(\mathbf{q},\dot{\mathbf{q}}) + \mathbf{J}^{\mathbf{T}}(\mathbf{q}) \cdot \mathbf{B} \cdot \mathbf{J}(\mathbf{q}) \cdot \dot{\mathbf{q}} = \mathbf{Q}(\mathbf{q},\dot{\mathbf{q}},t) \tag{2.28}$$

where $\mathbf{J}(\mathbf{q})$ is the Jacobian of the kinematic map from joint angles to hand position in Cartesian coordinates. With $\mathbf{r} = [x, y]^T$ and $\mathbf{q} = [q_1, q_2]^T$ the Jacobian is a 2x2 matrix of partial derivatives:

$$\mathbf{J}(\mathbf{q}) = \begin{bmatrix} \frac{\partial x}{\partial q_1} & \frac{\partial x}{\partial q_2} \\ \frac{\partial y}{\partial q_1} & \frac{\partial y}{\partial q_2} \end{bmatrix} \tag{2.29}$$

The Jacobian is a function of the arm configuration (i.e. the joint angles) as the kinematic map is non-linear. We derived Equation (2.28) from the force field equation (2.27) and from the expression of the Jacobian, by combining the transformation from joint to hand velocity,

$$\dot{\mathbf{r}} = \mathbf{J}(\mathbf{q}) \cdot \dot{\mathbf{q}} \tag{2.30}$$

with the transformation from hand force to joint torque

$$\boldsymbol{\tau} = \mathbf{J}^{\mathbf{T}}(\mathbf{q}) \cdot \mathbf{F} \tag{2.31}$$

We can therefore regard the field as an unexpected change on the function $\mathbf{G}(\mathbf{q},\dot{\mathbf{q}})$:

$$\mathbf{G}'(\mathbf{q},\dot{\mathbf{q}}) = \mathbf{G}(\mathbf{q},\dot{\mathbf{q}}) + \mathbf{J}^{\mathbf{T}}(\mathbf{q}) \cdot \mathbf{B} \cdot \mathbf{J}(\mathbf{q}) \cdot \dot{\mathbf{q}} \tag{2.32}$$

The solution of Equation (2.28) is a new trajectory $\mathbf{q}_{\mathbf{B}}(t) \neq \mathbf{q}_{\mathbf{A}}(t)$. Figure 3.10C shows the response of a subject to this sudden dynamic change. It is important to observe that the force field and, accordingly, the change to G vanishes when the hand is at rest. This has two consequences. First, the subjects could not know if the field is on or off before starting to move and, second, the field does not alter the static equilibrium at the end of each movement. Therefore, we were not surprised to see that the effect of the unexpected force field on the trajectories was to cause a hook-shape deviation, with the hand landing eventually on the planned target. While further studies revealed the presence of online corrections, these are not needed to maintain the final position unchanged by the force field. In the second part of the experiment, subjects were continuously exposed to the force field, with only sporadic "catch trials" in which the field was unexpectedly suppressed.

Figure 3.11. Learning the field. Left panels. Average and standard deviation of hand trajectories while a subject was training in the force field of Figure 10B. Performance plotted during the first, second, third and final 250 targets (F1, F2, F3 and F4). Left panel. After effects were observed when the field was unexpectedly removed during the four training sets. The four panels (C1 to C4) show the average and standard deviation of the hand trajectory while moving in the absence of perturbing forces. As the trajectories recovered the initial rectilinear shape in the field, they developed increasing after-effects in the catch trials. (From Shadmehr and Mussa-Ivaldi, 1994)

Figure 3.11 shows the evolution of motor performance across four consecutive blocks of trials. The four trajectories on the left – F1 to F4 – were produced by a subject moving in the force field. The trajectories on the right – C1 to C4 – were obtained from the same subject in the same four periods, but when the force field was unexpectedly removed. These are the catch trials, where one can observe the after-effect of learning, that is how the motor command evolved as the subject learned to move inside the force field. If one compares the first set of movements in the field (F1 and Fig 3.10D), to the last set of after-effects (C4) one can immediately observe that the shape of the after-effect at the end-of learning is qualitatively the mirror image of the initial effect of the exposure to the field. This is easily understood in the dynamical framework. As the subjects learn to move in the field, they compensate for the forces that are generated by it along the desired trajectory. We represent this process by adding a term to the original controller:

$$\mathbf{I(q)}\cdot\ddot{\mathbf{q}}+\mathbf{G(q,\dot{q})}+\mathbf{J^T(q)}\cdot\mathbf{B}\cdot\mathbf{J(q)}\cdot\dot{\mathbf{q}}=\mathbf{Q}(\mathbf{q},\dot{\mathbf{q}},t)+\mathbf{\Delta}(\mathbf{q},\dot{\mathbf{q}},t) \tag{2.33}$$

For this equation to admit the original trajectory as a solution it is sufficient that

$$\mathbf{J^T(q)}\cdot\mathbf{B}\cdot\mathbf{J(q)}\cdot\dot{\mathbf{q}}=\mathbf{\Delta}(\mathbf{q},\dot{\mathbf{q}},t) \tag{2.34}$$

along $\mathbf{q_A}(t)$. This is equivalent to perform a local approximation of the external field. In this case, if the external field is suddenly removed after learning, this approximation – i.e. the *internal model* of the field – becomes a perturbation with the opposite sign:

$$\mathbf{I(q)}\ddot{\mathbf{q}}+\mathbf{G(q,\dot{q})}-\mathbf{\Delta}(\mathbf{q},\dot{\mathbf{q}},t)\approx\mathbf{I(q)}\cdot\ddot{\mathbf{q}}+\mathbf{G(q,\dot{q})}-\mathbf{J^T(q)}\cdot\mathbf{B}\cdot\mathbf{J(q)}\cdot\dot{\mathbf{q}}=\mathbf{Q}(\mathbf{q},\dot{\mathbf{q}},t) \tag{2.35}$$

This, however, is not the only way for our motor system to compensate for the external force. We discuss briefly two alternatives: stiffening up, or create a motor tape. The first consists in increasing the rigidity of the moving arm by co-activating opposing muscles. This is what you would do if you want to resist movement when someone shakes your hand. You activate the biceps and the triceps simultaneously. The two muscles oppose each other and the result is that the forearm becomes more rigid. Our brains chose this approach whenever facing an environment that is hard or impossible to predict. For example when holding the rudder of a sailboat and you need to keep the heading as steady as possible. In our force-field experiment, the presence of after-effects is sufficient to rule out this type of control response: subjects were not merely becoming stiffer but they learned to counteract forces that – by design – were not random but predictable. The second possibility, instead, would be consistent with our observation of after-effects. Equation (2.34) does not imply that $\mathbf{\Delta}$ must be a representation of the force field on the left-hand side. The applied forces were designed to depend of the state of motion of the arm, $\left[\mathbf{q}^T, \dot{\mathbf{q}}^T\right]^T$. However, to recover the initial trajectory it would be sufficient for Δ to be a pure function of time,

$$\mathbf{\Delta}(t) = \mathbf{J}^{\mathbf{T}}(\mathbf{q}_{\mathbf{A}}(t)) \cdot \mathbf{B} \cdot \mathbf{J}(\mathbf{q}_{\mathbf{A}}(t)) \cdot \dot{\mathbf{q}}_{\mathbf{A}}(t) \tag{2.36}$$

This is simply "motor tape" that plays back the forces encountered along the desired trajectory. The motor tape does not contain any information about the dependence of the external force upon the state variables. A second experiment was needed to address this possibility.

Michael Conditt, Francesca Gandolfo and Sandro Mussa-Ivaldi (1997) tested the hypothesis of a motor tape with a simple experiment. They trained a group of volunteers to make reaching movements in a force field. The duration of each movement was typically less than one second. Therefore, after learning to reach targets in the field, the motor tapes would have the typical duration of each movement and would be "specialized" to reproduce the selectively the force sequences encountered in the training phase. The hypothesis then predicts that subjects would not be able to compensate for forces encountered while performing different movements, like drawing a circle, a triangle or figure-eight pattern. Even if these drawing movements take place in the same region of space and velocity ranges as the reaching movements, the states of motions are traversed in different temporal sequences. And the drawing movements last at least twice as long as the reaching movements. Nevertheless, the subjects of this experiment, after training to reach in the force field, were perfectly capable to compensate for the disturbing forces while drawing circles and other figures. A second group of subjects, instead of practicing reaching movements, practiced drawing circles in the force field. The drawing movements of these subjects after training were not distinguishable from the drawing movements of subjects in the first group, who trained with reaching. These findings are sufficient to rule out the idea that subjects learned a motor tape and provide instead additional support to the concept that the brain construct, through learning a computational representation of the dynamics as a dependence of the actions to be produced by the control system – in this case the function $\mathbf{\Delta}(\mathbf{q}, \dot{\mathbf{q}})$ - upon the state of motion of the limb. While the represented entity does not have to be a force in the Newtonian sense, the resulting behavior is consistent with the representation of movement dynamics as formulated by classical laws.

3.6 Another type of state-based dynamics: motor learning

The force-field experiments demonstrate the keen ability of the brain to deal with dynamics and with the concept of "state". In mechanics, the state is a minimal set of variables that is sufficient to predict the future evolution of a system. In a single motor act, such as reaching for an object, the state in question is the mechanical state of the limb. Because movement is governed by the laws of classical mechanics, the state is a collection of position and velocity variables. There is another kind of state that is relevant to motor behavior. Learning experiments have demonstrated that it is possible to predict how knowledge and control evolve through time based on past experience and of incoming information. Learning itself can therefore be observed as a dynamical system. This perspective will be developed in the remainder of this book. Here, we briefly introduce some of the key concepts of this approach. Let us start again from the classical Newton's equation for a point mass, $F = m\ddot{x}$. We derive the velocity of the point mass by integrating once this equation:

$$\dot{x}(t) = \dot{x}(t_0) + \int_{t_0}^{t} \frac{F}{m} dt' \tag{2.37}$$

Instead of looking at the velocity as a continuous function of time, we may take samples at discrete instants of time. Then, provided that these instants are separated by small time intervals, δt, we may use Newton's equation to derive as an approximation the velocity at the next instant $\dot{x}^{(n+1)}$ from the velocity $\dot{x}^{(n)}$ and the external force $F^{(n)}$ at the current instant. This is the *difference equation*:

$$\dot{x}^{(n+1)} = \dot{x}^{(n)} + \frac{F}{m}\delta t \tag{2.38}$$

If the mechanical work in an infinitesimal displacement is an exact differential, then the force is the gradient of a potential energy function,

$$F(x,t) = -\frac{\partial U(x,t)}{\partial x} \tag{2.39}$$

We now can re-write the difference equation for a point mass in a potential field as

$$\dot{x}^{(n+1)} = \dot{x}^{(n)} - \frac{\partial U^{(n)}}{\partial x} \cdot \alpha \,. \tag{2.40}$$

with $\alpha = \dfrac{\delta t}{m}$.

Opher Donchin, Joseph Francis and Reza Shadmehr (2003) proposed to describe motor learning with a differential equation that has a similar form. However, in this case, the state variable in question is not the velocity or the position of the arm, but the internal model of the force field in which the arm is moving. They represented the internal model as a collection of basis functions over the state space of the arm. For example, each basis function could be a Gaussian centered on a preferred velocity. But this is certainly not the only possible choice. Referring to Equation (2.34), let us represent the force field model in generalized coordinates as

$$\mathbf{\Delta}(\mathbf{q},\dot{\mathbf{q}}) = \mathbf{W} \cdot \mathbf{g}(\mathbf{q},\dot{\mathbf{q}}) \tag{2.41}$$

where $\mathbf{g} = \left[g_1(\mathbf{q},\dot{\mathbf{q}}), g_2(\mathbf{q},\dot{\mathbf{q}}),\ldots, g_m(\mathbf{q},\dot{\mathbf{q}})\right]^T$ is a collection of m scalar basis functions and W is a $n \times m$ matrix of coefficients. For a two-joint arm $n = 2$. But, more generally, n is the dimension of the configuration space or, equivalently, the number of degrees of freedom under consideration. Then, the matrix **W** can also be written as an array of m n-dimensional vectors, one for each basis function

$$\mathbf{W} = \left[\mathbf{w_1}, \mathbf{w_2}, \ldots, \mathbf{w_m}\right] \tag{2.42}$$

and the basis functions are scalars that modulate the amplitude of the w vectors over regions of the workspace, in a way similar to the receptive field of a sensory neuron. As we perform a series of reaching movements in the force field, $\mathbf{q}^{(1)}(t), \mathbf{q}^{(2)}(t), \ldots, \mathbf{q}^{(k)}(t), \ldots$ we experience a corresponding sequence of "force trajectories"

$$\mathbf{D}^{(i)}(t) = \mathbf{J^T}(\mathbf{q}^{(i)}(t)) \cdot \mathbf{B} \cdot \mathbf{J}(\mathbf{q}^{(i)}(t)) \cdot \dot{\mathbf{q}}^{(i)}(t) \tag{2.43}$$

Suppose that on the i-th movement, we intended to move along the trajectory $\mathrm{q_A}(t)$. Then, based on the internal model that we had at that point, we expected to experience a force

$$\mathbf{\Delta}^{(i)}(t) = \mathbf{\Delta}^{(i)}(\mathbf{q_A}(t), \dot{\mathbf{q}}_\mathbf{A}(t)) = \mathbf{W}^{(i)} \cdot \mathbf{g}(\mathbf{q_A}(t), \dot{\mathbf{q}}_\mathbf{A}(t)) \tag{2.44}$$

Here, we have that the variable portion of the internal model is the coefficient matrix W. We assume that the basis function remain fixed through time. This, of course is for conceptual convenience. More complex (and also more realistic) models of learning would allow for variations in all terms of $\mathbf{\Delta}$. Most importantly, at the end of the i-th movement we have a discrepancy – an error – between what we experienced and what we expected to experience. This is a function

$$\mathbf{e}^{(i)}(t) = \mathbf{D}^{(i)}(t) - \mathbf{\Delta}^{(i)}(t) \tag{2.45}$$

If we integrate the Euclidean norm of this function along the movement, we obtain a scalar quantity, a real number

$$E^{(i)}(\mathbf{W}^{(i)}) = \int_0^T \mathbf{e}^{(i)T}(t) \cdot \mathbf{e}^{(i)}(t)\, dt\,. \tag{2.46}$$

This is an overall error that we have experienced across the movement and this error depends on the parameters of the internal models. The purpose of learning is to get better at what we are doing. In this case it means to make the expectation error as smallest as possible. A straightforward way to do so is by changing the parameters of the internal model so as to minimize the error on the next trial. If we expand the argument of the integral (2.46), we see that it is a quadratic form of the parameters:

$$\mathbf{e}^{(i)T} \cdot \mathbf{e}^{(i)} = \left(\mathbf{D}^{(i)} - \mathbf{\Delta}^{(i)}\right)^T \cdot \left(\mathbf{D}^{(i)} - \mathbf{\Delta}^{(i)}\right) = \mathbf{D}^{(i)T} \cdot \mathbf{D}^{(i)} - 2 \cdot \mathbf{D}^{(i)T} \cdot \mathbf{W}^{(i)} \cdot \mathbf{g} + \mathbf{g}^T \cdot \mathbf{W}^{(i)T} \cdot \mathbf{W}^{(i)} \cdot \mathbf{g} \tag{2.47}$$

Therefore, the error function $E^{(i)}(W^{(i)})$ has a global minimum that we reach by changing W along the gradient :

$$\mathbf{W}^{(i+1)} = \mathbf{W}^{(i)} - \frac{\partial \mathbf{E}^{(i)}}{\partial \mathbf{W}^{(i)}} \cdot \alpha \tag{2.48}$$

This is a difference equation that tells us how the internal model of the field changes in time. We compare it with equation (2.40) that describes the motion of a point-mass in a potential field. The two equations have different order. Newtonian mechanics is of the second order while this simple learning model is of the first. Both equations tell us how a state variable changes in time under the action of an external input. In the case of the mass, the input is the mechanical force, which the gradient of a potential energy function. In the case of learning, the "driving force" is the gradient of the prediction error. Perhaps, the most important common feature of Newtonian mechanics and the description of learning as a dynamical system is that both allow us to connect theory with experiment. This will be the leading thread of the chapters that follow.

Summary

Space maps are not only relevant to navigation, but to movements in general. Here, we consider the maps that our brains must form to manipulate the environment. The kinematics maps relate the motions of the hand to the motions of the arm and come in two forms: direct kinematics, from arm's joint angles to hand position, and inverse kinematics, from hand position to arm's joint angles. These maps are non-linear and, as a consequence, a straight line in hand space maps into a curve in arm space and vice-versa. This reflects the fact that a space of angular variables is inherently curved, while the Euclidean space where the hand moves is flat. Observation of human movements over a plane demonstrated that the hand tends to follow rectilinear trajectories.

Observations with a altered kinematics and with force perturbations demonstrated that the kinematics of arm movement is not a collateral effect of the brain attempting to economize or simplify the generation of muscle forces. Instead, to produce naturally observed movements, the brain must effectively solve a complex problem of dynamics. The solution of this problem does not necessarily involve any explicit representation of Newtonian forces, but it requires the ability to relate the motor commands that reach the muscles to the consequences in terms of changes in state of motion of the arm. This connection is represented in the arm's dynamics equations, which contain terms representing the inertial properties of the arm and terms representing the viscoelastic properties of the muscles. The motor commands needs not to be a force and may represent instead the static equilibrium point expressed by the neuromuscular system. This is the position at which the arm would eventually settle under the influence of opposing muscle forces. Controlling the motion of this equilibrium point for producing a desired movement of the hand is equivalent to solving an inverse dynamic problem.

The first experimental evidence for internal models of dynamics derived from observing how subjects learned to perform reaching movements in the presence of deterministic force fields. These fields apply to the hand a force that depends upon the hand's state of motion, i.e. its position and velocity. The observations were consistent with the subjects gradually developing the ability to predict and cancel the applied forces along the movement trajectories. The sudden removal of the field results in after-effects of learning, that is in a perturbation that mirrors the initial effect of the unexpected application of the field. The after-effects, as well as the pattern of learning generalizations demonstrate that the subjects learn to predict the functional form of the perturbation, which effectively corresponds to a change in the controlled dynamical system.

We introduce the concept that motor learning is a dynamical process whose state is the associated internal model of the body and its environment. In its simplest form, learning to move in a force field is described by a first order differential equation whose state variable is the model of the field and whose input is the gradient of the prediction errors. In this respect, movement and learning are two concurrent processes operating over different time scales.

References

Conditt, M. A., F. Gandolfo, et al. (1997). "The motor system does not learn the dynamics of the arm by rote memorization of past experience." Journal of Neurophysiology **78**(1): 554-560.

Donchin, O., J. T. Francis, et al. (2003). "Quantifying generalization from trial-by-trial behavior of adaptive systems that learn with basis functions: theory and experiments in human motor control." Journal of Neuroscience **23**(27): 9032.

Feldman, A. (1966). "Functional Tuning of the Nervous System During Control of Movement or Maintenance of a Steady posture. II. Controllable Parameters of the Muscles. III. Mechanographic Analysis of the Execution by Man of the Simplest Motor Task." Biophysics **11**: 565-578, 766-755.

Flanagan, J. R. and A. K. Rao (1995). "Trajectory adaptation to a nonlinear visuomotor transformation: evidence of motion planning in visually perceived space." Journal of Neurophysiology **74**: 2174-2178.

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

Hollerbach, J. M. and T. Flash (1982). "Dynamic interactions between limb segments during planar arm movements." Biological Cybernetics **44**: 67-77.

McCulloch, W. and W. Pitts (1943). "A logical calculus of the ideas immanent in nervous activity." Bulletin of mathematical biology **5**(4): 115-133.

Morasso, P. (1981). "Spatial control of arm movements." Experimental Brain Research **42**: 223-227.

Nash, J. (1956). "The imbedding problem for Riemannian manifolds." Annals of Mathematics **63**(1): 20-63.

Polit, A. and E. Bizzi (1978). "Processes controlling arm movements in monkeys." Science **201**(4362): 1235.

Shadmehr, R. and F. A. Mussa-Ivaldi (1994). "Adaptive representation of dynamics during learning of a motor task." Journal of Neuroscience **14**(5): 3208-3224.

Spong, M. W., S. Hutchinson, et al. (2005). Robot Modeling and Control. Hoboken, NJ, John Wiley and Sons.

Uno, Y., M. Kawato, et al. (1989). "Formation and control of optimal trajectory in human multijoint arm movement." Biological Cybernetics **61**: 89-101.

Won, J. and N. Hogan (1995). "Stability properties of human reaching movements." Experimental Brain Research **107**: 125-136.

Zee, D., L. Optican, et al. (1976). "Slow saccades in spinocerebellar degeneration." Archives of Neurology **33**(4): 243.

Index

NOTES

[i] An isometric embedding is a transformation from a mathematical object to a space that contains the objects and preserves the object's metric properties. An example is the embedding of a sphere in the Euclidean 3D space. We can locate a point over the sphere by two coordinates, such as latitude and longitude. Alternatively, we can establish three Cartesian axes and describe the point by a triplet x, y, z. Importantly, the first type of description is non-Euclidean. Over the sphere, parallel lines may intersect and Pythagoras' theorem is violated. However, the second description is Euclidean. Looking at the earth from outer space one would see that the meridians of longitude indeed meet at the poles, but they are not parallel lines. They are closed curves.

[ii] We adopt the standard notation in classical mechanics, where the lower-case q denotes a position and the upper-case Q represent a force in a system of "generalized coordinates". Generalized coordinates reflect the effective movement space of a mechanical system. Classical mechanics assumes that the essential law that governs the movement of any system is Newton's law, $F = ma$, applied to each of its constituents particles. However, the great majority of those particles are bound to stay at fixed distances from one another. While a simple pendulum contain billions of molecules, the motion of this immense set of particles is described by a single variable. This is called the "degree of freedom" of the pendulum. The idealized two-joint arm of Fig. 3.1 has only two degrees of freedom. The variables that describe each degree of freedom are called generalized coordinates. The generalized forces are the forces in that particular system of coordinates. For example, for an angular coordinate, the generalized force is a torque. For a linear coordinate, the generalized force is an ordinary force.

[iii] To see this, consider a simple second order system, such as a spring and a mass or a pendulum. You may drive the system by applying an external force that will have the system reaching several times the same state, $s = [x, \dot{x}]^T$ s (e.g. a given position at zero velocity) and leaving this state on different trajectories. Then, knowing only that the system is at s is no longer sufficient to know its future.

[iv] According to standard engineering terminology, Equation (2.17) describes a PD control system, where P stand for "Proporional" and D for "Derivative".