

Benchmarking cerebellar control

Patrick van der Smagt
Institute of Robotics and Mechatronics
German Aerospace Center / DLR Oberpfaffenhofen
P.O. Box 1116, 82230 Wessling, Germany
email smagt@dlr.de

October 13, 2000

Abstract

Cerebellar models have long been advocated as viable models for robot dynamics control. Building on an increasing insight in and knowledge of the biological cerebellum, many models have been greatly refined, of which some computational models have emerged with useful properties with respect to robot dynamics control.

Looking at the application side, however, there is a totally different picture. Not only is there not one robot on the market which uses anything remotely connected with cerebellar control, but even in research labs most testbeds for cerebellar models are restricted to toy problems. Such applications hardly ever exceed the complexity of a 2 DoF simulated robot arm; a task which is hardly representative for the field of robotics, or relates to realistic applications.

In order to bring the amalgamation of the two fields forwards, we advocate the use of a set of robotics benchmarks, on which existing and new computational cerebellar models can be comparatively tested. It is clear that the traditional approach to solve robotics dynamics loses ground with the advancing complexity of robotic structures; there is a desire for adaptive methods which can compete as traditional control methods do for traditional robots.

In this paper we try to lay down the successes and problems in the fields of cerebellar modelling as well as robot dynamics control. By analyzing the common ground, a set of benchmarks is suggested which may serve as typical robot applications for cerebellar models.

Keywords: robot dynamics, robot arm control, computational cerebellar models, neural networks

1 Introduction

For a long time, the field of robotics has been heavily dependent on advances in mechatronics; typically, heavy robot structures with strong motors were necessary for accurate positioning of payloads with relatively low weight. In recent years, however, the continuing development of new materials and of new drive concepts which are both strong and light-weight have enabled the construction of robot arms with an impressive force-to-weight ratio and which are very dextrous (e.g., (Hirzinger, 1996)).

The control of such robot structures, however, is problematic. Whereas traditional, rigid, robots can be stably controlled due to the fact that each joint can be treated as an

independent entity, stable high-speed control of a general dextrous robot arm with 6 or more strongly interdependent joints is, to date, highly problematic. Despite continued mini-successes in the analysis and modelling of the complex dynamic behaviour of such robot arms, it seems that a solution which does not use adaptive control will not solve all problems at hand.

Taking a look at nature, the cerebellum is a very likely candidate for solving these problems. In the cerebellum, sensor signals combined with motor plans¹ are used to generate motor signals. An example is the vestibulo-ocular reflex: when the head or body is moved, the cerebellum fine-tunes the motion commands for the eye muscles which make the eyes fixate on the same point. Early medical research (Holmes, 1917, 1939) has given clear indications that the cerebellum is used for stable vertebrate control; cerebellar lesions lead to instability of the motor system.

Cerebellar models were first applied to robot dynamics control after Albus' CMAC model was published in 1975. Especially the later implementations by Miller et al. (e.g., 1989, 1997) for control of a 4 DoF robot arm as well as biped control have demonstrated the power of this approach. Albus' model, however, is known to be an oversimplified version of the biological truth: based on the BOXES approach by Michie and Chambers (1968), it effectively implements an adaptive table lookup method with hashing-based output smoothing. The CMAC has often been considered a function approximation model only, and can be compared with other general function approximation models as such.

Although cerebellar modelling has come a long way since then, applications of such models remain limited to toy problems, not exceeding the complexity of (simulated) two-link arms. But is it at all reasonable to investigate such control systems for artificial robot arms? In order to evaluate the applicability of cerebellar models for robot control, and to suggest a better integration of cerebellar modelling and robotics, we investigate the feasibility and propose a set of benchmarks containing typical robotics problems.

This paper is organized as follows. In section 2 we describe the current-day understanding of the biological cerebellum. Section 3 gives an overview of the field of robotics, and the problem of dynamics control is more deeply discussed. In section 4 the two fields are brought together: which problems could be tackled with computational cerebellar models? Is the desire for such an integration at all valid? In this section, a number of benchmarks for cerebellar control are proposed. A discussion is found in section 5.

2 Biological motor control and the cerebellum

Throughout the ages the questions "What does the cerebellum do?" and "How does it do it?" have received ample attention among researchers in various fields. One of the earliest empirical investigations of the cerebellum were performed by the Greek surgeon Claudius Galenus of Pergamum (131–201 AD), who in his early career taught medicine in the Asclepion in Pergamum (currently Izmir), and later on worked as a gladiator physician and surgeon to Emperor Marcus Aurelius in Rome. From his surgical work he concluded that the cerebellum must be responsible for muscle control.

It was Andreas Vesalius (1514–1564), with his skeptical investigations of Galenic theory, who questioned ancient medical knowledge and published his *De Humani Corporis Fabrica Libri Septum* with detailed anatomical studies. From then on, the structure of the

¹A *motor plan* or *program* is a sequence of muscle commands that can be executed without feedback, using the correct timing, e.g., speech, playing golf, etc. In robotics, the term *desired joint trajectory* is commonly used.

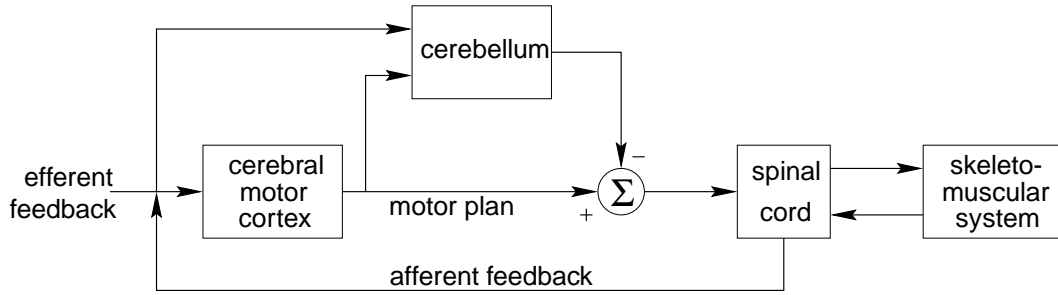


Figure 1: The cerebellum functions as a feedforward filter.

brain was gradually discovered. Detailed investigations of the cerebellum as the motor control centre began with Bell (1811), who proposed that motor fibers originate from the cerebellum; an assumption that was later verified by Flourens (1824). From various wartime studies, Holmes (1917, 1939) concluded that the cerebellum operates as a filter on motor commands originating from the cerebrum, correcting on previously issued motor commands. In his theory, the cerebellum has the function of a comparator. A very comprehensive study of the cerebellum, which subsequently formed the basis of many models, was published by Eccles *et al.* (1967). His work, and the subsequent book by Ito (1984), has influenced many of the subsequent cerebellar models.

2.1 Structure

The cerebellum functions as a feedforward control center for the motor commands originating in the cerebrum (Figure 1). Containing learned models of the skeletomuscular system, it provides timing control of opposing muscles, and force as well as stiffness control.

2.1.1 The spinal cord

The spinal circuitry provides independent control of muscle length and joint stiffness (Bullock & Contreras-Vidal, 1993). Fig. 2 depicts this part of the neuromuscular control system. The length of the muscle is sensed by neuromuscular spindles that consist of specialized intrafusal muscle fiber. This (Ia) afferent feedback Δf_s is sent to the α -neurons in the spinal cord as well as to the brain. The muscle is contractive in response to efferent feedback from the γ -neurons. The feedback enables the spindle to keep its tension within efficient operating limits to ensure optimal sensitivity.

The Golgi tendons consist of afferent fibers at the end of the muscle, and sense muscular tension (Ib afferent feedback). This feedback of the agonist inhibits α -efference of the agonist but excites the α -efference of the antagonist.

The muscle actuator is represented by the transfer function

$$\lambda(s) = G_\alpha(s)\Delta f_\alpha(s) + G_f(s)f(s). \quad (1)$$

The internal length λ of the muscle is a function which is related to the change of efferent nerve excitation Δf_α from the α -neuron and the force f acting on the muscle. The transfer function G_α and G_f depend on the muscle stiffness, the muscle damping, an elastic component serial to the muscle, the combined mass of the limb and interface element, and the spring constant of the interface.




Figure 2: Simplified model of the upper limb control system (Reprinted with permission from (Koepppe, 2000)).

2.1.2 The cerebellum

The human cerebellum (see Figure 3) consists of about 10 million Purkinje cells (pc), each receiving about 150,000 excitatory synapses via the parallel fibers (pf) (Larsell & Jansen, 1972; Zagon, MacLaughlin, & Smith, 1977; Linden, 1996). The pf are the axons of the granule cells; these cells are excited by the mossy fibers (mf) originating from the spinal cord, the cerebrum, and the brainstem. Each pf synapses on about 200 Purkinje cells. A Purkinje cell receives further excitatory synapses from one single climbing fiber (cf); this can fire a cell when active. Basket cells, being activated by pf afferents but also inhibited by pc, can inhibit a Purkinje cell, thus ensuring activation of a single pc within a local neighbourhood, and are only found in birds and mammals (Ito, 1984). Stellate cells are similar in form and function to basket cells. Lugaro cells, the function of which is currently unknown, are found in the granular layer. Having elongated cell bodies, they receive input from pc, while their axons extend in the molecular layer. Finally, Golgi cells receive input from pf, mf, and cf. They inhibit granule cells.

The granule cells operate as pattern separators. The densely 'coded' patterns, originating from the spinal cord, have to be 'preprocessed' by the granule cells, such that the Purkinje cells can discriminate them. The output of a Purkinje cell is an inhibitory signal to the cerebellar nuclei.

The cerebellar cortex is divided in three layers: the outer synaptic layer (also called molecular layer), the Purkinje layer, and the inner receptive layer (the granular layer). The cortex appears to be organized in cortico-nuclear microzones (Figure 3); each of these microzones contains the parameters for a certain movement (Ito, 1984). The output of the microzones originates in the pc, and flows to the nuclear cells.

Mechanistic models of the cerebellum did not appear until the paper by Braitenberg and Atwood (1958). Braitenberg (1961) primarily interpreted the cerebellum as a timing




Figure 3: Major components of the structure of the cerebellum. An arrow indicates an excitatory connection, a termination (-) an inhibitory connection.

organ; in his theory, the pf have the functionality of delay lines.

Influenced by Eccles *et al.* (1967), two other early models by Marr (1969) and Albus (1971) view the cerebellum as a learning pattern recognition system. Their more detailed models, as well as a subsequent computational model by Albus (1975), have contributed to a wide acceptance of the pattern recognition theory.

2.2 Learning

A major contribution of the papers by Braitenberg, Marr, and Albus consisted of a theory of how learning in the cerebellum takes place. We know that the cerebellum learns sequences of voluntary movements as well as motor programs, but also adapts to external influences. Fortunately, the regular structure of the cerebellum has aided in a good insight in how learning takes place. The key issue is that learning is *context-driven*. In Marr's theory, pf synapses on pc are strengthened when simultaneous activity of the cf and pf occurs. Thus memory traces are stored at the pf synapses. Albus later suggested that the cerebellum functions as an adaptive pattern classifier, where the pf synapses become weaker; the cf signal is thus interpreted as an error signal. When active, the rate of change is given as

$$\Delta t \frac{dw_i}{dt} = -x_i (F^{\text{active}} - F^{\text{inactive}}) \quad (2)$$

where Δt is the time constant, w_i the value of the i th pc/pf synapse, x_i the firing incidence at the i th pc/pf synapse, F the firing frequency of the active or inactive pf. This hypothesis summarizes the combined Marr-Albus theory.

The adaptation of pf synapses is called *long-term depression* (LTD). Next time around that the same pattern appears at the pf, the Purkinje cell will *not* fire and thus *not* inhibit its corresponding deep nucleus. Verification of LTD at the pf synapses followed in the early eighties by Ito *et al.* (1982), who demonstrated that simultaneous stimulation of pf and cf result in LTD of the pf/pc synapses.

Learning through LTD occurs in many places of the brain, and is usually accompanied by the opposite process called *long-term potentiation* (LTP) (Bliss & Collingridge, 1993). Although Artola and Singer (1993) concluded that LTD learning in the cerebellum distinguishes itself through the absence of LTP, this conclusion is not generally accepted (Kano, 1996). It be noted that LTD cannot be the only learning mechanism for the pc/pf synapses,

since otherwise they would eventually be driven to zero responsiveness. Furthermore, LTP takes place both at the granule cells and at the deep nuclei (Hepp, 1999), a fact which has not yet been incorporated in any computational model.

Another cerebellar controversy concerns the vestibulo-ocular reflex (VOR). The function of this reflex is to stabilize an image on the retina while the head is subject to rotational or translational movements. Lesions in the flocculus, the oldest part of the cerebellum, prevent learning of the VOR, and Ito (1984) proposed the flocculus as an example of the Marr-Albus theory. Later experiments, however (Lisberger, Pavelko, & Broussard, 1994; Lisberger, Pavelko, Bronte-Stewart, & Stone, 1994; Lisberger, 1994) show that the same Purkinje cell can change its activity in the same as well as in the opposed direction of the VOR gain. Again, this phenomenon cannot be explained with the Marr-Albus theory.

Thirdly, there are conflicting reports on whether the cf signal is an error signal; other properties of the cf signals are also not explained by the Marr-Albus theory (De Schutter, 1997). Even the function of the cf signal is a mystery (Simpson, Wylie, & De Zeeuw, 1996): while some researchers assume that the cf signal leads to LTD at the pc/pf synapses, other conclude that it leads to short-term enhanced responsiveness at the pc/mf synapses. Also, the cf has been proposed to serve as internal timing signal.

A final problem is the solution of the *credit assignment problem*. This problem can be separated in the *temporal* credit assignment problem (e.g., (Sutton, 1984)) and the *structural* credit assignment problem (e.g., (Houk, Buckingham, & Barto, 1996)). This problem involves tracing back which signal from which unit causes an error in the output, such that only the properties of this unit should be changed for optimal learning.

Despite all the controversy, a few largely accepted facts of the cerebellum remain:

- the cerebellum is, among other things, responsible for the *coordination of movement* (Flourens, 1824)—although it's not completely clear what *that* means.
- the cerebellum learns models of the skeletomuscular system (Miall, Weir, Wolpert, & Stein, 1993; Kawato, 1995). Whether they be forward, inverse, or both still has to be verified.
- The learning process in the cerebellum is influenced by simultaneous activation of pf and cf at a pc.

2.3 Tasks of the cerebellum

Although there exist clear indications that the cerebellum is involved in noun-verb association, visual shape discrimination, mental rotation, attention, working memory, and IQ in general (Kawato, 1997), general areas of concern are classical conditioning tasks and motor control.

Conditioning. A large number of experiments has shown that the cerebellum is involved in conditioned motor reflexes (Bartha & Thompson, 1995). A traditional experiment is the eye blink reflex. This experiment works as follows: a subject is given a puff of air in the eye, before which a single tone is played. When the delay between the playing of the tone and the air puff is constant, the test subject learns to close the eyelid on hearing the tone, before the air puff is administered. This reflex, which can be learned within the range 100ms–1.5s, is known to be solved by the cerebellum. Models exist which explain this behaviour (e.g., (Thompson, 1990)).

Motor control. A second, related task of the cerebellum is the ‘coordination of movement’ (Flourens, 1824). Typical examples are the vestibulo-ocular reflex (Ito, 1984), limb motor control (Kawato, 1995), eye saccade movements (Houk et al., 1996), and so on. Many models have been proposed and described in detail (Albus, 1975; Houk, 1989; Paulin, 1989; Miller, Glanz, & Kraft, 1990; Kawato & Gomi, 1992; Miall et al., 1993; Buonomano & Mauk, 1994; Kawato, 1995; Schweighofer, 1995; Houk et al., 1996; De Schutter, 1997; Smagt, 1998).

2.4 Cerebellar lesions and agenesis

There is a popular belief that patients with cerebellar agenesis can eventually recover from this disability; supposedly, the cerebrum would be able to take over the functions of the cerebellum.

Holmes (1939), possibly profiting from the availability of abundant research material, records a large number of motor control impairments due to partial cerebellar lesions. All his experiments demonstrate an instability in the motor programs of affected patients. He concludes that “a striking feature of cerebellar injury [is] that its symptoms gradually decrease in intensity and may in time disappear . . . by compensation by intact parts of the cerebellum.” However, “in acute stages of *disease* the disturbances may be to some extent reduced.” Through a careful analysis of medical case histories, Glickstein (1994) finds that patients with full cerebellar defect do not recover from their impairment. Clearly, the mammalian brain requires a cerebellum for stable motor control.

3 Control

But can robots use a cerebellar model for stable control? Besides the fact that the majority of robot arms, that are subject of academic research, have an anthropomorphic structure, there is little resemblance between human and robot arms. A first, important, difference is the material that they consist of. The rigid and heavy-weight construction material of industrial robot arms implies the requirement of actuators (typically DC motors) which can generate the necessary forces. The robot motion is incurred through the forces or torques that are exerted by the actuators.

3.1 Robot motion tasks

Usually, the task that has to be performed by a robot arm with up to six degrees of freedom can be described by a desired Cartesian trajectory $(\hat{\mathbf{x}}[t], \hat{\boldsymbol{\phi}}[t])$, where $\hat{\mathbf{x}}$ is the desired (3D) position and $\hat{\boldsymbol{\phi}}$ the desired (3D) orientation, that must be followed by its end effector. In sensor-based robotics, this trajectory may not be explicitly available before the whole motion has been completed, but when computed would uniquely describe the movement at hand. Also, in some applications only the final position of the end effector is of importance; in that case, **trajectory planning** or **path planning** is used to determine the trajectory from the current robot position to the desired position.

In traditional robotics, the instantiation of this trajectory is performed in a sequence of steps:

Inverse kinematics. For a robot with revolute joints (e.g., an anthropomorphic robot), this process translates the end effector trajectory into a desired trajectory $\hat{\boldsymbol{\theta}}[t]$ in

joint space. The forwards mapping, the **kinematics**, is in fact determined by (the mechanical structure of) the robot.

In many instances, the inverse kinematics is not a function but has multiple solutions. First, when the dimensionality (degrees of freedom) of the robot exceeds that of the task, there are infinitely many solutions. But also when these two are equal (typically 6), there may be more than one arm configuration to reach a point, e.g., elbow-up and elbow-down. Additional constraints, such as restrictions in allowed acceleration, can be used to choose favourite solutions.

When the inverse kinematics solution is only used to determine a target end effector position, this step must be followed by trajectory planning. At set intervals Δt (typically in the range 1–20ms) a joint position must be computed which the robot arm should move to. Since the interval Δt is constant, this determines the full kinematics of the motion.

Dynamic control or tracking. When a positional increment is available at each Δt , the joint servo control computes the necessary forces or torques at the motor side to realize the requested motion. In traditional industrial robotics, a PD (*Proportional-Derivative*) controller is customarily used. For a rotary joint, this controller sets the torque to a proportional constant k_p times the joint position error, plus a derivative constant k_v times the joint velocity error. For joint i :

$$\hat{\tau}_i(t) = k_{p,i} \left[\hat{\theta}_i(t) - \theta_i(t-1) \right] + k_{v,i} \left[\dot{\hat{\theta}}_i(t) - \dot{\theta}_i(t-1) \right]. \quad (3)$$

Clearly, all the joints are independently controlled, and centrifugal and Coriolis forces are assumed to be nonexistent. This control law is known as *servo control*.

An important property of the PD control rule is its proven stability. Arimoto and Miyazaki (1984) showed that the joint error goes to 0 when using PD control, it was later on shown that it decays at least exponentially.

Motor control. In the study of biological control systems, the term *motor control* refers to the control of movement in general. In robotics, *motor control* refers to the mechanism which controls the DC motors in order to generate a desired torque $\hat{\tau}$. Assuming that the joint servo control problem is adequately solved, the motor control problem can be tackled on a joint-independent basis. A fast local feedback loop, using the measured motor current i as steering signal, suffices here to control the motor.

From section 2 it can be concluded that the applicability of cerebellar models to robot control problems can be found in the improvement of adaptive dynamic control methods. But where are such methods required?

3.2 Robot dynamics

A robot arm consists of the following parts:

- (1) **the links** or arm segments. We consider these to be rigid bodies, which is realistic for most small and medium-sized robots, as well as for skeletomuscular systems.
- (2) **the actuators.** A tendency exists towards using DC motors or step motors for generating the required force; however, other types of actuators such as pneumatic artificial muscles (Chou & Hannaford, 1996) have also received considerable attention.

(3) **the connection between the actuators and the links** (e.g., gear boxes). With a tendency towards light-weight robot arms, for DC or step motor based robot arms it is customary to use high-ratio gear boxes such that the motors used can be kept small and light. On the downside, however, is a considerable elasticity, such that both the rotation at the motor side and at the link side must be measured. *Direct drive* robots are also under consideration; yet, the motors have a very low force-to-weight ratio, and are therefore not suitable for light-weight robots.

Under the assumption that the links and joints do not mechanically deform, the forces that the robot structure exerts at the actuators is given by

$$\boldsymbol{\tau} = M(\boldsymbol{\theta})\ddot{\boldsymbol{\theta}} + C(\boldsymbol{\theta}) \left[\dot{\boldsymbol{\theta}}\dot{\boldsymbol{\theta}} \right] + D(\boldsymbol{\theta}) \left[\dot{\boldsymbol{\theta}}^2 \right] + F(\boldsymbol{\theta}, \dot{\boldsymbol{\theta}}) + G(\boldsymbol{\theta}) \quad (4a)$$

$$\equiv M(\boldsymbol{\theta})\ddot{\boldsymbol{\theta}} + V(\boldsymbol{\theta}, \dot{\boldsymbol{\theta}}) + G(\boldsymbol{\theta}) \quad (4b)$$

where $\boldsymbol{\tau}$ is an N -vector of torques exerted by the links, and $\boldsymbol{\theta}$, $\dot{\boldsymbol{\theta}}$, and $\ddot{\boldsymbol{\theta}}$ are N -vectors denoting the positions, velocities, and accelerations of the N joints. $[\dot{\boldsymbol{\theta}}\dot{\boldsymbol{\theta}}]$ and $[\dot{\boldsymbol{\theta}}^2]$ are vectors

$$\left[\dot{\boldsymbol{\theta}}\dot{\boldsymbol{\theta}} \right] = \left[\dot{\theta}_1\dot{\theta}_2, \dot{\theta}_1\dot{\theta}_3, \dots, \dot{\theta}_{N-1}\dot{\theta}_N \right]^T, \quad (5)$$

$$\left[\dot{\boldsymbol{\theta}}^2 \right] = \left[\dot{\theta}_1^2, \dot{\theta}_2^2, \dots, \dot{\theta}_N^2 \right], \quad (6)$$

$M(\boldsymbol{\theta})$ is the matrix of inertia (the mass matrix), $C(\boldsymbol{\theta})$ is the matrix of Coriolis coefficients, $D(\boldsymbol{\theta})$ is the matrix of centrifugal coefficients, $F(\boldsymbol{\theta}, \dot{\boldsymbol{\theta}})$ is a friction term, and $G(\boldsymbol{\theta})$ is the gravity working on the joints. Eq. (4b) is a simplified notation of Eq. (4a).

Agonist/antagonist drive. Biological arms, as well as some experimental robots, use an agonist/antagonist drive principle. An example robot arm using this principle was the Bridgestone rubbertuator arm, equipped with two or four McKibben (Chou & Hanaford, 1996) artificial muscles per joint, connected via sprockets (Katayama & Kawato, 1991; Hesselroth, Sarkar, Smagt, & Schulten, 1994; Smagt, Groen, & Schulten, 1996). However, complex dynamic properties as well as extreme temperature sensitivity and costly maintenance have lead to the production stop of this robot.

The advantage of an agonist/antagonist drive concept is that the *stiffness* of a joint can be very intuitively set: while the joint angle depends on the difference in exerted forces, the joint stiffness depends on the sum of those forces. There are several reasons, however, for mechanical systems not to use this scheme; the key one being the fact that two drives per joint are required, resulting in a significantly higher weight and more power consumption. These resources are better invested in a robot arm construction with a higher force-to-weight ratio.

3.2.1 Rigid robots

The simplest kind of robot arm consists of rigid links which are connected by rigid joints. This assumption is true enough for industrial robots; the construction of the robot arm is thus that any yield in the links as well as the joints can be neglected. In this case control is done by linearization of the control equation. Eq. (4b) can be simplified in order to obtain:

$$\boldsymbol{\tau} = Y(\boldsymbol{\theta}, \dot{\boldsymbol{\theta}}, \ddot{\boldsymbol{\theta}})w \quad (7)$$

where w are the parameters which have to be estimated.

In industrial robots, the actuators are typically strong enough such that the diagonal elements of the mass matrix M and the centrifugal matrix D are prevalent, while C is approximately 0. Additionally, all matrices are approximately constant, i.e., independent of θ and its derivatives. These simplifications result in $\tau_i = m_i\ddot{\theta} + d_i\dot{\theta}^2 + f_i$, where i is the joint number: the joints can be independently controlled.

Such simplifications usually do not hold for light-weight robot arms, which are subject of research in many robotics research labs. Due to weight and space limitations, the actuators that are employed are not powerful enough to eliminate the influence of gravity, friction, and Coriolis and centrifugal forces. This means that, apart from having to take the full matrices M , C , and D into account, these and the F and G matrices are parameterized by the joint positions and velocities; Eq. (4a) cannot be simplified anymore.

3.2.2 Flexible links

Fortunately, material science allows the use of high-carbon wire in the production of robots: a strong light-weight material which, although still costly to produce and especially to process, is used as an alternative to aluminium because of its strength. The little research that is currently being done on flexible link robots is generally restricted to two-link arms (see, e.g., Talebi, Khorasani, & Patel, 1998). We consider this topic to be outside the scope of this paper.

3.2.3 Compliance

As mentioned before, light-weight robot arms are mostly equipped with actuators which can exert only a limited torque, and therefore require high-ratio gear boxes. The disadvantage of this approach is the yield in such gear boxes, which has to be taken into account in the control law.

The effect of a yielding gear box is that the torque at the motor side differs from the torque at the link side. The resulting actuator is modeled by a motor-link pair connected by a spring. To measure the spring properties, angle sensors both at the drive (measuring θ_m) and link (measuring θ_l) sides have to be available. Equation (4b) changes as follows:

$$\tau_l = M(\theta_l)\ddot{\theta}_l + V(\theta_l, \dot{\theta}_l) + G(\theta_l), \quad (8a)$$

$$\tau_m = J(\dot{\theta}_m)\ddot{\theta}_m + \tau_l \quad (8b)$$

where τ_m is the torque at the drive side and, using a linear spring model, $\tau_l \equiv k(\theta_m - \theta_l)$ the torque at the link side. J can generally be assumed to be a diagonal matrix.

Naturally, biological skeletomuscular systems are always compliant, and it is this aspect in which cerebellar control may be able to play an important role.

3.2.4 Delays

A robot arm is generally controlled in various feedback loops: using imperfect measurements, some model is used to generate new action commands which are subsequently executed by the robot. During action, the actual motion is guarded and corrected where necessary (see Figure 4).

Feedback loops are necessary since the models on which control commands are based are never in perfect correspondence with the real world. Clearly, the faster the feedback loop works, the less precise the robot models that are used need to be. Therefore the

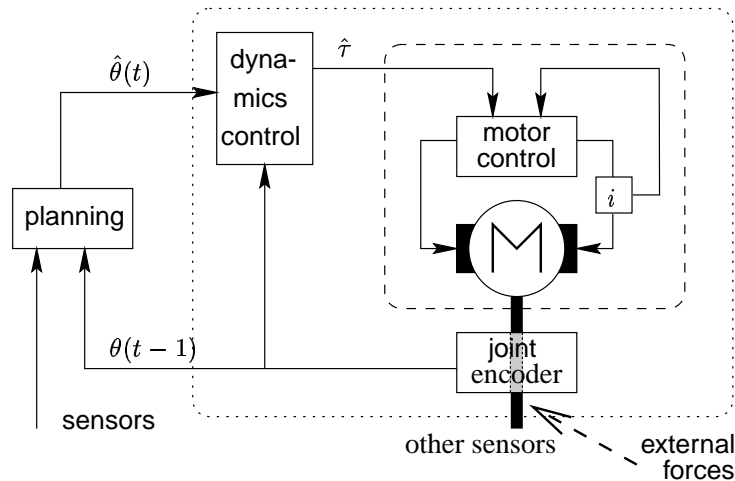


Figure 4: Feedback loops in a robot control system. At the lowest level, the motor controller uses the measured electric current i to effectuate a desired torque $\hat{\tau}$. Above that, the joint servo control uses the joint angle measurements $\theta(t-1)$ to generate $\hat{\tau}$ such as to reach a desired joint position $\hat{\theta}(t)$.

tendency in robotics is to decrease the joint servo feedback delays. For light-weight robot arms, delays of 1ms are not uncommon, whereas industrial robots typically have delays of around 10ms, necessary to obtain fast and stable control.

The situation is radically different in biological systems. Feedback delays can be as much as 110–150ms for proprioceptive control (Cole & Abbs, 1987) and 200–250ms for visuomotor control (Miall et al., 1993). Stable control in the presence of such extensive delays can partly be explained by the fact that skeletomuscular systems have an apparent passive behaviour (Hogan, 1990).

3.3 Traditional control

To control a robot structure, a control law has to be devised which computes torques which, when applied to the joints, make the robot arm stably follow a trajectory $\theta[t]$, taking the dynamics equation (4b) or Eqs. (8a) and (8b) into consideration.

It is customary to simplify the problem via *control law partitioning*: the control law is dissected in a model based part and an error based part. If $\hat{\tau}$ is the input torque to the plant, we write $\hat{\tau} = \alpha \hat{\tau}' + \beta$ where $\hat{\tau}'$ is the torque applied to the unit mass system (see Figure 5). By setting $\alpha = M(\theta)$ and $\beta = V(\theta, \dot{\theta}) + G(\theta)$ (Eq. (4b)), the control law is

$$\hat{\tau}' = \hat{\ddot{\theta}} + \mathbf{k}_v(\hat{\dot{\theta}} - \dot{\theta}) + \mathbf{k}_p(\hat{\theta} - \theta), \quad (9)$$

i.e., a simple servo control method (cf. Eq. (3)). This control system is also known as *computed torque control*.

Clearly, the key question remains: how do we find the optimal α and β ? There exist various robotics techniques which address this problem; however, it is far from solved, and requires substantial modeling of robot structures. In PD control, the simplified model $\alpha \equiv I$ and $\beta \equiv 0$ is chosen: the model-based component is eliminated. Furthermore, the elements of the diagonal matrices \mathbf{k}_v and \mathbf{k}_p are chosen as large as possible (depending on how strong the actuators are), such as to eliminate the errors introduced by the neglected Coriolis and centrifugal forces. If this approach is reasonable, a proven stable control

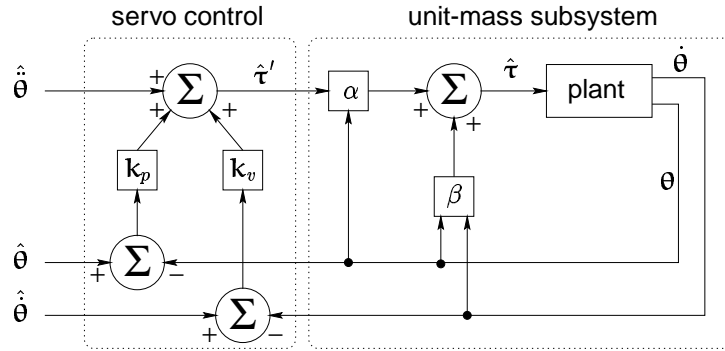


Figure 5: Easier control by partitioning. The unit-mass subsystem can be controlled using a simple error-based controller; the ‘only’ prerequisite here is the availability of the model-based parts α and β .

method results (Arimoto & Miyazaki, 1984). In other cases, however, α and β can, to date, only be found by extensive modeling of the robot.

4 Integrating cerebellar systems in robot control

In describing the benchmarks for a cerebellar controller, there are two issues that have to be taken care of. First, what are the (dynamical) properties of the system on which the controller is tested? Section 4.2 provides with a set of systems with increasing dynamical properties. Second, on what tasks are these complete systems then tested? Three tasks are given in section 4.3. By combining from both pools, up to 18 different benchmarks can be selected.

4.1 Where can cerebellar models be applied?

The long delays that are present in the biological motor system require a highly accurate control mechanism. Since many movements (such as swinging a golf club) have a duration of well below one second and therefore use very little feedback, the control has to be mainly model based. It is proposed by many authors that such models are learned by the cerebellum, be they feedback (e.g., Atkeson, 1989; Kawato & Gomi, 1992) or feedforward (e.g., Miall et al., 1993).

But how can a cerebellar model be used in a control loop? Figure 6 shows three possibilities. First, using the cerebellum as a model in a negative feedback loop (Fig. 6a). In this case, the cerebellum is used as an exact copy of the plant, and mimics its input-output behaviour. The model is used by the control loop to update its internal behaviour, and can correct its output without receiving direct feedback from the plant. Feedback from the plant can be used to update the forward model.

A plant model used with positive feedback is shown in Fig. 6b. In this case, the model is used to mimic the plant as a unit mass system, such that the control law can be restricted to a simple servo law. As discussed above, this is a traditional robotics approach known as computed torque control (Fig. 5), and the major problem lies in finding the correct model parameters. A fast feedback loop is required to control the robot, in order to have the servo control loop lead to stable control.

In Fig. 6c an model is used which directly translates the motor plan $\hat{\theta}$ to torques $\hat{\tau}$, which are fed into the plant. Feedback is required where the inverse model is imprecise.

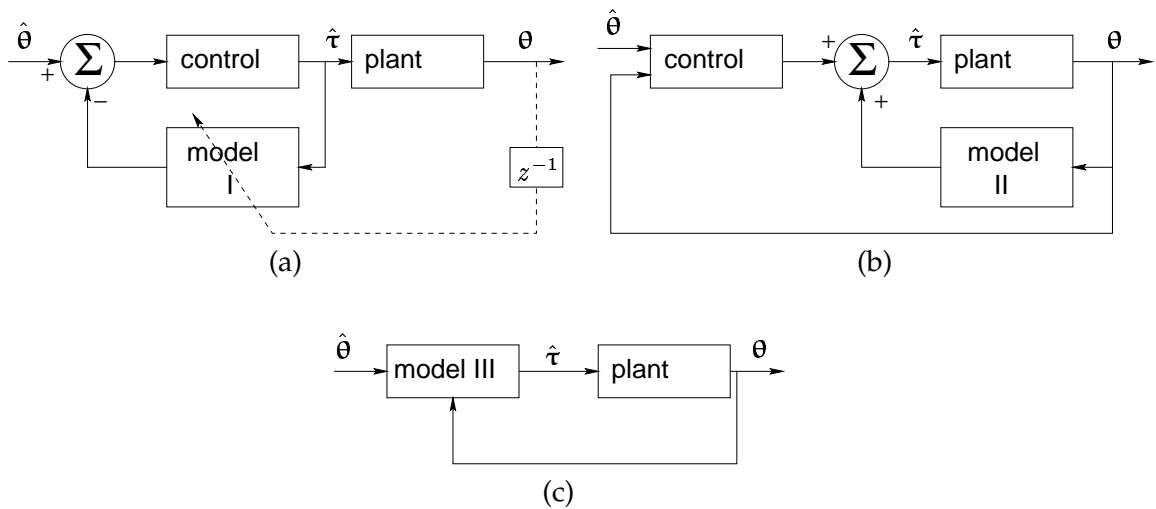


Figure 6: The use of a model of the plant in a control loop. When negative feedback (a) is used, the model mimics the action of the plant, and can be used to internally optimize the controller (Miall et al., 1993). Using a model with positive feedback (b), the control law is a servo law (cf. Fig. 5). When a model (c) is used, the desired state can be directly translated to a motor command.

Miall et al. (Miall et al., 1993), choosing model (a) as optimal solution because of the more restricted motor feedback requirements, go a step further and also model delays in the system; their resulting theory describes a model based on the Smith predictor (Smith, 1959), a well-known engineering concept which is used to model systems with extensive delays.

4.1.1 About the validity of integration

It is clear that the kinematic and dynamic properties of biological and robot arms are radically different. Similarly, differences in material properties make the respective control systems very different. Are such differences prohibitive for applying biological control methods to robotic tasks?

This problem can be largely circumvented by simulating many of the known properties of such biological systems. This approach, although technically unsound from a current-day point of view, is used in many settings where biological control principles are applied to (simulated) robotics. In this paper, however, we prefer to take existing robotics concepts as a basis, in order to investigate how biological control principles can help us there,

To compare the biological and robotics control principles, we discuss three basic differences.

Delays. An important difference is the length of delays present in the control loop. For instance, human proprioceptive feedback can be as slow as 110–150ms. In robots, this feedback can be well below 1ms, while control loops are in the range 1–10ms. In exteroceptive delays, this difference is up to an approximate factor 10.

On the other hand, robotic systems are typically capable of performing precision tasks at very high velocities and accelerations, which of course requires such fast feedback. It must be noted, however, that stable control of robot arms in the presence of delays of around 100ms is difficult.

Bandwidth of the actuators. Clearly, the bandwidth of a skeletomuscular system is much lower than that of a robot arm, by an approximate order of magnitude. This lower bandwidth, however, suffices for stable control, since skeletomuscular systems have an apparent passive behaviour (Hogan, 1990). Although partly due to the fact that muscles exhibit a spring-like behaviour, it is not entirely clear how this passivity is otherwise obtained. At any rate, neural feedback pathways exist which relate motions at one joint to muscle activity at another (Eccles, Eccles, & Lundberg, 1957).

Stiffness. As discussed in the introduction of this paper, industrial robots are designed as completely stiff mechanical systems. Although being varied by the α -neurons, due to the structure and the compliance of the muscles the stiffness of the skeletomuscular system is of course much lower than that of industrial robots. Modern-day lightweight robots, however, have stiffness characteristics similar to those of biological arms (Hirzinger, 1996).

In conclusion, it seems plausible that the cerebellum is responsible for the apparent stability of skeletomuscular systems. Clinical studies (Holmes, 1917; Glickstein, 1994) seem to confirm this assumption. When considering the above arguments, there appears to be no reason that cerebellar models could not be applied to robotic systems.

4.2 The dynamics

When applied to robotics, where fast feedback loops are common, all three approaches depicted in Fig. 6 are feasible control approaches. Note that, for biological systems, there exists an additional advantage of using feedforward instead of feedback models: the internal prediction of the actions can be used in other parts of the brain, before feedback of the central neural system is available.

In order to test a cerebellar model in such a feedback loop, it should be tested on typical robotics applications. We have prepared a number of typical benchmarks stemming from traditional robot dynamics modelling, which can serve as the “plant” in Fig. 6. Each of these benchmarks contains a simulated model of a robot arm.

When restricting the control system to the feedback loop Fig. 6a, it suffices in theory to have the cerebellum learn the robot dynamics benchmarks proposed in this section. In realistic applications, however, the forward dynamics may not be available, such that the problem of generating learning samples without the availability of a model still has to be solved. When successfully completed, the resulting model can be plugged in the Fig. 6a model.

The simulations are based on the iterative Newton-Euler algorithm for robot dynamics (see, e.g., Craig, 1986 for a lengthy discussion). The principle of this algorithm is as follows: in the *forward* iterations, the joint velocities and accelerations are computed starting at the base joint and moving outward to the last joint. Then, in the *backward* iterations, the resulting torques are computed back towards the base of the robot.

Although these computations can be done symbolically, the required number of computations is humongous even for a robot with only three degrees of freedom (whereas 6 are required for complete dexterity). It is more efficient, therefore, to compute the forward and backward equations numerically, and get the results from there.

The benchmarks which we suggest in this paper are all based on such simulations².

²These benchmarks are available from the author as computer program, and can currently be obtained via <http://www.robotic.dlr.de/Smagt/research/cerebellum/>.

Reasoning from a robotics point of view, we distinguish three levels of complexity. First, the dynamics of a robot arm are considered with no external forces working on it. This type of robot is relatively easy to control with servo control only. When we add gravity to the robot, an extra gravity compensation or integral servo term must be added. Finally friction, which is highly problematic in light-weight robot arms, is added.

For all types of benchmarks we distinguish between a stiff robot, and a robot with flexible joints. For a stiff robot, the simulation can easily be performed as follows. Starting from an initial joint position and velocity, we can simulate the moving robot by computing

$$\ddot{\theta} \leftarrow M(\theta)^{-1} [\hat{\tau} - V(\theta, \dot{\theta}) - G(\theta)]. \quad (10)$$

The torque $\hat{\tau}$, generated by the external controller, influences the joint acceleration as described in this equation.

In the case of yielding joints, the simulation becomes somewhat more complex. This time starting from an initial θ_l , $\dot{\theta}_l$, θ_m , and $\dot{\theta}_m$, by looking at Eqs. (8a) and (8b), we find

$$\ddot{\theta}_l \leftarrow M(\theta_l)^{-1} [k(\theta_m - \theta_l) - V(\theta_l, \dot{\theta}_l) - G(\theta_l)], \quad (11)$$

$$\ddot{\theta}_m \leftarrow J(\dot{\theta}_m)^{-1} [\hat{\tau}_m - k(\theta_m - \theta_l)]. \quad (12)$$

The following three sections describe the six benchmarks. Note that the (b) benchmarks are all more complex to control than the (a) benchmarks.

4.2.1 Benchmark I: N DoF robot arm with internal forces

For the first benchmark we consider the dynamics of a robot arm where the gravity and friction forces are neglected. The basic equation (4a) is simplified:

$$\tau = M(\theta)\ddot{\theta} + C(\theta) [\dot{\theta}\dot{\theta}] + D(\theta) [\dot{\theta}^2]. \quad (13)$$

(a) For this benchmark we assume a robot arm with stiff joints. In the case that the motors are strong enough (true for most industrial robot arms), the Coriolis forces can be neglected. Such a robot arm can be optimally controlled using a standard PD controller, of which the stability is proven.

A cerebellar control method should be able to control such a robot for $N \geq 6$. Both the simplified decoupled case as the case where cross influences are to be considered have to be taken into account. The robot simulation equation can now be written as

$$\ddot{\theta} \leftarrow M(\theta)^{-1} \left\{ \hat{\tau} - C(\theta) [\dot{\theta}\dot{\theta}] - D(\theta) [\dot{\theta}^2] \right\}. \quad (14)$$

(b) A robot arm with yielding joints. For the simulation we use

$$\ddot{\theta}_l \leftarrow M(\theta_l)^{-1} \left\{ k(\theta_m - \theta_l) - C(\theta) [\dot{\theta}\dot{\theta}] - D(\theta) [\dot{\theta}^2] \right\}, \quad (15)$$

$$\ddot{\theta}_m \leftarrow J(\dot{\theta}_m)^{-1} [\hat{\tau}_m - k(\theta_m - \theta_l)]. \quad (16)$$

4.2.2 Benchmark II: N DoF robot arm with external forces A

For this benchmark we only add the gravity component:

$$\boldsymbol{\tau} = M(\boldsymbol{\theta})\ddot{\boldsymbol{\theta}} + C(\boldsymbol{\theta}) \left[\dot{\boldsymbol{\theta}}\dot{\boldsymbol{\theta}} \right] + D(\boldsymbol{\theta}) \left[\dot{\boldsymbol{\theta}}^2 \right] + G(\boldsymbol{\theta}). \quad (17)$$

For PD servo control this type of control is more complex, since the steady-state forces which are introduced by the gravity cannot be compensated.

These external forces are traditionally compensated by adding an integral error term $k_i \int_T \boldsymbol{\theta}(t) - \hat{\boldsymbol{\theta}}(t) dt$ to the servo control law, resulting in an extra Integral component added to the PD law (known as a PID law). The stability of the resulting control law has not been proven, however.

(a) Again we assume a robot arm with stiff joints. When the gravity component is added, the robot arm can be optimally controlled using a standard PID controller. Alternatively, the gravity part is solved by a separate controller (such as, e.g., the ISM model (Katayama & Kawato, 1991; Kawato & Gomi, 1992)) to which a controller of section 4.2.1(a) is added.

Again, cerebellar control method should be able to control such a robot for $N \geq 6$. Both the simplified decoupled case as the case where cross influences are to be considered have to be taken into account. The robot simulation equation is now written as

$$\ddot{\boldsymbol{\theta}} \leftarrow M(\boldsymbol{\theta})^{-1} \left\{ \hat{\boldsymbol{\tau}} - C(\boldsymbol{\theta}) \left[\dot{\boldsymbol{\theta}}\dot{\boldsymbol{\theta}} \right] - D(\boldsymbol{\theta}) \left[\dot{\boldsymbol{\theta}}^2 \right] - G(\boldsymbol{\theta}) \right\}. \quad (18)$$

(b) A robot arm with yielding joints. For the simulation we use

$$\ddot{\boldsymbol{\theta}}_l \leftarrow M(\boldsymbol{\theta}_l)^{-1} \left\{ k(\boldsymbol{\theta}_m - \boldsymbol{\theta}_l) - C(\boldsymbol{\theta}) \left[\dot{\boldsymbol{\theta}}\dot{\boldsymbol{\theta}} \right] - D(\boldsymbol{\theta}) \left[\dot{\boldsymbol{\theta}}^2 \right] - G(\boldsymbol{\theta}) \right\}, \quad (19)$$

$$\ddot{\boldsymbol{\theta}}_m \leftarrow J(\dot{\boldsymbol{\theta}}_m)^{-1} [\hat{\boldsymbol{\tau}}_m - k(\boldsymbol{\theta}_m - \boldsymbol{\theta}_l)]. \quad (20)$$

4.2.3 Benchmark III: N DoF robot arm with external forces B

Finally we add the most complex force: friction. The dynamics of the robot has to be fully considered:

$$\boldsymbol{\tau} = M(\boldsymbol{\theta})\ddot{\boldsymbol{\theta}} + C(\boldsymbol{\theta}) \left[\dot{\boldsymbol{\theta}}\dot{\boldsymbol{\theta}} \right] + D(\boldsymbol{\theta}) \left[\dot{\boldsymbol{\theta}}^2 \right] + F(\boldsymbol{\theta}, \dot{\boldsymbol{\theta}}) + G(\boldsymbol{\theta}). \quad (21)$$

There are two problems with modelling the friction. First, the friction component cannot be determined from the model. Second, the nonlinear Coulomb component of the friction, described by $\tau_F = c \operatorname{sgn}(\dot{\theta})$, makes accurate control at low velocity virtually impossible, due to the jagged form at $\dot{\theta} = 0$.

Friction plays a prominent role in light-weight robot arms with high-ratio gear boxes; it typically accounts for over 50% of the power consumption. It is therefore not an effect which can be neglected in robot modeling.

(a) We first assume a robot arm with stiff joints. There exist no general stable control methods which can control such robots without extensive modeling. Standard methods are often computed torque-based. The robot simulation equation is now written as

$$\ddot{\boldsymbol{\theta}} \leftarrow M(\boldsymbol{\theta})^{-1} \left\{ \hat{\boldsymbol{\tau}} - C(\boldsymbol{\theta}) \left[\dot{\boldsymbol{\theta}}\dot{\boldsymbol{\theta}} \right] - D(\boldsymbol{\theta}) \left[\dot{\boldsymbol{\theta}}^2 \right] - F(\boldsymbol{\theta}, \dot{\boldsymbol{\theta}}) - G(\boldsymbol{\theta}) \right\}. \quad (22)$$

(b) A robot arm with yielding joints. For the simulation we use

$$\ddot{\theta}_l \leftarrow M(\theta_l)^{-1} \left\{ k(\theta_m - \theta_l) - C(\theta) [\dot{\theta}\dot{\theta}] - D(\theta) [\dot{\theta}^2] - F(\theta, \dot{\theta}) - G(\theta) \right\}, \quad (23)$$

$$\ddot{\theta}_m \leftarrow J(\dot{\theta}_m)^{-1} [\hat{\tau}_m - k(\theta_m - \theta_l)]. \quad (24)$$

4.3 Application

Having laid out the possible variations in dynamics, we finalize the benchmarks by describing three applications that are typically solved by a controller: trajectory tracking, compliance control, and exteroceptive feedback control.

4.3.1 Tracking

In robotics applications, accurate trajectory following of a robot arm is often a minimal requirement for a controller. We suggest three types of trajectories on which a controller should be tested: a smooth trajectory, a pulse response trajectory, and a standard robot controller test trajectory.

Smooth trajectory. The choice of a trajectory on which a control method is tested depends on the kinematic and dynamical properties of the robot arm, as well as the target area of application. For instance, for a 2-link robot arm (of which the applicability is very limited), a typical task would be to track a sinusoidal curve within the plane of the arm in both Cartesian dimensions, i.e., a circle by its end-effector (when the phase difference of both sinusoids equals $\pi/2$), such that the angular velocity of the circle remains constant. Although nonlinear in the robots joint angles, this trajectory is relatively easy to follow since the joint velocities as well as accelerations remain within small bounds. These bounds can be varied with the radius of the circle.

The same principle, i.e., of following a sinusoidal curve in all Cartesian dimensions, can be used as a trajectory for a 3 DoF robot arm (see Figure 7a).

Pulse response. Difficult trajectories are those where the joint velocity profiles are nearly-discontinuous, i.e., joint accelerations are close to their mechanical maximum. To test the performance of a controller in these extreme situations, the pulse response of the system can be measured. The desired velocity of each joint should follow a step function as closely as possible. Possible pulse responses are shown in Figure 7b. Here, an overdamped system will reach the desired velocity too slowly, while an underdamped system will swing. Although a critically damped system will have the fastest response, a controller which exhibits underdamping is often used, since in that case the desired joint position $\hat{\theta}$ can also be reached (the integral of the dashed curve equals the integral of the pulse).

The Schmidt trajectory. There exist several standard methods to test such controllers, one of which is the *Schmidt trajectory* (Fig. 8). This trajectory consists of a 2D curve which must be accurately followed by the end-effector of a robot. The whole trajectory should be followed with constant velocity, or rather maximum acceleration at the corners of the trajectory.

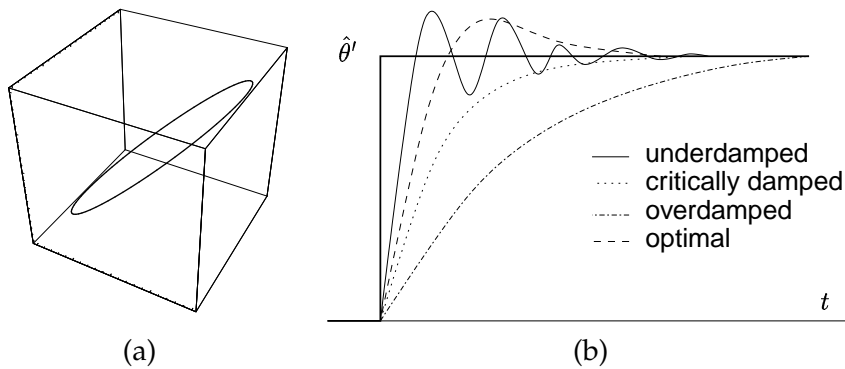


Figure 7: (a) Circle trajectory $(x, y, z) = (\sin(t), \cos(t), \sin(t))$ for a 3 DoF robot arm. (b) Possible controller pulse response.




Figure 8: The Schmidt trajectory for testing robot arm controller. This 2D curve describes the path that should be followed by the end-effector of a robot arm. The dimensions of the trajectory depend on the dimensions of the manipulator, such that the maximum reach space of the robot is explored.

4.3.2 Compliance control

In the case that a robot manipulator is in physical contact with its environment, the problem of compliancy has to be solved. Inability particular, the *stiffness* of the manipulator has to be adapted to the stiffness of the object. For instance, if the robot arm is moving freely, its stiffness can be maximal, thus allowing for faster and more accurate control. If in contact with a rigid object, however, its stiffness has to be reduced immediately, in order to prevent instability.

Although being varied by the α -neurons, the stiffness of the skeletomuscular system is of course much lower than that of industrial robots. Also, compliance is typically solved by the kinematic structure of arms, e.g., grasping is typically done in such a fashion that the arm position allows for compliance.

4.3.3 Exteroceptive feedback delays

As a final task, control with exteroceptive (e.g., visual) feedback has to be considered. The delays in such setups are much more extensive than with proprioceptive feedback alone. Typically, this task involves the tracking of a moving object in sensor space, but in an environment where sensor data preprocessing is assumed to be solved.

5 Discussion

Clearly, there exists an enormous gap between the applicability of cerebellar models to robot dynamics control, and existing applications of such. Nevertheless, the emerging interest in cerebellar models demonstrate an enormous possibility. Although many cerebellar models exist which can be applied to robot dynamics control, the applications of such models are almost exclusively restricted to 2-link simulated robot arms.

To advance the merger of the two, we advocate the use of a set of benchmarks which clearly demonstrate the current-day problems in robot dynamics control.

It should never be forgotten that the real test of any robot control method can only be done on a robot, and not with a simulated model. The benchmarks should therefore only be used as a final step before integration on a robot arm is started.

Acknowledgements

The expertise of Alin Albu-Schäffer in the field of robot dynamics and modelling was key for many parts of this paper. Discussions with Ralf Koeppel concerning learning and muscle models contributed largely to the corresponding parts of the paper. The author kindly acknowledges the continuing support of Gerd Hirzinger, who leads the Department of Robotic Systems at the German Aerospace Center in Oberpfaffenhofen.

References

- Albus, J. (1971). A theory of cerebellar function. *Mathematical Bioscience*, 10, 25–61.
- Albus, J. (1975). A new approach to manipulator control: The cerebellar model articulation controller (CMAC). *Journal of Dynamic Systems, Measurement and Control, Transactions of the ASME, series G*, 97, 270–277.

- Arimoto, S., & Miyazaki, F. (1984). Stability and robustness of PID feedback control for robot manipulators of sensory capability. In M. Brady & R. P. Paul (Eds.), *Proceedings of the First International Symposium on Robotics Research* (pp. 783–801). MIT Press.
- Artola, A., & Singer, W. (1993). Long-term depression of excitatory synaptic transmission and its relationship to long-term potentiation. *Trends in Neuroscience*, *16*, 480–487.
- Atkeson, C. G. (1989). Learning arm kinematics and dynamics. *Annual Review of Neuroscience*, *12*, 157–183.
- Bartha, G. T., & Thompson, R. F. (1995). Cerebellum and conditioning. In M. A. Arbib (Ed.), *The Handbook of Brain Theory and Neural Networks* (pp. 169–172). Cambridge, MA: MIT Press.
- Bell, C. (1811). *Idea of a New Anatomy of the Brain; Submitted for the Observations of His Friends*. London: Strahan and Preston.
- Bliss, T. V. P., & Collingridge, G. (1993). A synaptic model of memory: Long-term potentiation in the hippocampus. *Nature*, *361*, 31–39.
- Braitenberg, V. (1961). Functional interpretation of cerebellar histology. *Nature*, *190*, 539–540.
- Braitenberg, V., & Atwood, R. P. (1958). Morphological observations on the cerebellar cortex. *Journal of Comparative Neurology*, *109*, 1–34.
- Bullock, D., & Contreras-Vidal, J. (1993). How spinal neural networks reduce discrepancies between motor intention and motor realization. In K. Newell & D. Corcos (Eds.), *Variability and motor control* (pp. 183–221). Champaign, Illinois: Human Kinetics Press.
- Buonomano, D. V., & Mauk, M. D. (1994). Neural network model of the cerebellum: Temporal discrimination and the timing of motor responses. *Neural Computation*, *6*, 38–55.
- Chou, C.-P., & Hannaford, B. (1996). Measurement and modeling of McKibben pneumatic artificial muscles. *IEEE Transactions on Robotics and Automation*, *12*(1), 90–102.
- Cole, K. J., & Abbs, J. H. (1987). Kinematic and electromyographic responses to perturbation of a rapid grasp. *Journal of Neurophysiology*, *57*, 1498–510.
- Craig, J. J. (1986). *Introduction to Robotics*. Addison-Wesley Publishing Company.
- De Schutter, E. (1997). A new functional role for cerebellar long term depression. *Progress in Brain Research*, *114*, 529–542.
- Eccles, J. C., Eccles, R. M., & Lundberg, A. (1957). The convergence of monosynaptic excitatory afferents onto many different species of alpha motoneurons. *Journal of Physiology*, *137*, 227–252.
- Eccles, J. C., Ito, M., & Szentagothai, J. (1967). *The Cerebellum as a Neuronal Machine*. Berlin: Springer-Verlag.
- Flourens, P. (1824). *Recherches expérimentales sur les propriétés et les fonctions du système nerveux, dans les animaux vertébrés*. Paris: Cervot.

- Glickstein, M. (1994). Cerebellar agenesis. *Brain*, 117, 1209–1212.
- Hepp, K. (1999). *Personal communication*.
- Hesselroth, T., Sarkar, K., Smagt, P. van der, & Schulten, K. (1994). Neural network control of a pneumatic robot arm. *IEEE Transactions on Systems, Man, and Cybernetics*, 24(1), 28–38.
- Hirzinger, G. (1996). Towards a new robot generation for space, terrestrial, and medical applications. In *Proc. Second ECPD Int. Conference on Advanced Robotics, Intelligent Automation, and Active Systems*.
- Hogan, N. (1990). Mechanical impedance of single- and multi-articular systems. In J. M. Winters & S. L.-Y. Woo (Eds.), *Multiple Muscle Systems: Biomechanics and Movement Organization* (pp. 149–163). New York: Springer-Verlag.
- Holmes, G. (1917). The symptoms of acute cerebellar injuries due to gunshot injuries. *Brain*, 40, 461–535.
- Holmes, G. (1939). The cerebellum of man. *Brain*, 62, 1–30.
- Houk, J. C. (1989). Cooperative control of limb movements by the motor cortex, brainstem, and cerebellum. In R. M. J. Cotterill (Ed.), *Models of Brain Function* (pp. 309–325). Cambridge University Press.
- Houk, J. C., Buckingham, J. T., & Barto, A. G. (1996). Models of the cerebellum and motor learning. *Behavioral and Brain Sciences*, 19(3), 368–383.
- Ito, M. (1984). *The Cerebellum and Neural Control*. New York: Raven Press.
- Ito, M., Sakurai, M., & Tongroach, P. (1982). Climbing fiber induced depression of both mossy fiber responsiveness and glutamate sensitivity of cerebellar Purkinje cells. *Journal of Physiology*, 324, 113–134.
- Kano, M. (1996). Long-lasting potentiation of GABAergic inhibitory synaptic transmission in cerebellar Purkinje cells: Its properties and possible mechanisms. *Behavioral and Brain Sciences*, 19, 354–361.
- Katayama, M., & Kawato, M. (1991). Learning trajectory and force control of an artificial muscle arm by parallel-hierarchical neural network model. In R. Lippmann, J. Moody, & D. Touretzky (Eds.), *Advances in Neural Information Processing Systems* (Vol. 3, pp. 436–442). San Mateo: Morgan Kaufmann.
- Kawato, M. (1995). Cerebellum and motor control. In M. A. Arbib (Ed.), *The Handbook of Brain Theory and Neural Networks* (pp. 172–178). Cambridge, MA: MIT Press.
- Kawato, M. (1997). *Multiple Internal Models in the Cerebellum*. (In van der Smagt & Bullock, pp. 17–19, NIPS*97)
- Kawato, M., & Gomi, H. (1992). The cerebellum and VOR/OKR learning models. *Trends in Neurosciences*, 15, 445–453.
- Koeppe, R. (2000). *Sensorimotor Skill Transfer of Compliant Motion Assembly Tasks*. Unpublished doctoral dissertation, DLR / Deutsches Zentrum für Luft- und Raumfahrt, Oberpfaffenhofen.

- Larsell, O., & Jansen, J. (1972). The comparative anatomy and histology of the cerebellum. III. The human cerebellum. In *Cerebellar Connections and Cerebellar Cortex*. Minneapolis: University of Minnesota Press.
- Linden, D. J. (1996). Cerebellar long-term depression as investigated in a cell culture preparation. *Behavioral and Brain Sciences*, 19, 339–346.
- Lisberger, S. G. (1994). Neural basis for motor learning in the vestibuloocular reflex of primates. III. Computational and behavioral analysis of the sites of learning. *Journal of Neurophysiology*, 72, 974–998.
- Lisberger, S. G., Pavelko, T. A., Bronte-Stewart, H. M., & Stone, L. S. (1994). Neural basis for motor learning in the vestibuloocular reflex of primates. II. Changes in the responses of horizontal gaze velocity purkinje cells in the cerebellar flocculus and ventral paraflocculus. *Journal of Neurophysiology*, 72, 954–973.
- Lisberger, S. G., Pavelko, T. A., & Broussard, D. M. (1994). Neural basis for motor learning in the vestibuloocular reflex of primates. I. Changes in the responses of brain stem neurons. *Journal of Neurophysiology*, 72, 928–953.
- Marr, D. (1969). A theory of cerebellar cortex. *Journal of Physiology*, 202, 437–470.
- Miall, R. C., Weir, D. J., Wolpert, D. M., & Stein, J. F. (1993). Is the cerebellum a Smith predictor? *Journal of Motor Behavior*, 25, 203–216.
- Michie, D., & Chambers, R. A. (1968). BOXES: An experiment in adaptive control. In E. Dale & D. Michie (Eds.), *Machine Intelligence* (Vol. 2, pp. 137–152). Edinburgh: Oliver and Boyd.
- Miller, W. T., Glanz, F. H., & Kraft, L. G. (1990). CMAC: An associative neural network alternative to backpropagation. *Proceedings of the IEEE, Special Issue on Neural Networks*, 78, 1561–1567.
- Miller III, W. T. (1989). Real-time application of neural networks for sensor-based control of robots with vision. *IEEE Transactions on Systems, Man, and Cybernetics*, 19(4), 825–831.
- Miller III, W. T., & Kun, A. L. (1997). Neural systems for robotics. In O. Omidvar & P. van der Smagt (Eds.), *Dynamic Balance of a Biped Walking Robot*. Boston, MA: Academic Press.
- Paulin, M. G. (1989). A Kalman filter theory of the cerebellum. In M. A. Arbib & S. Amari (Eds.), *Dynamic Interactions in Neural Networks* (pp. 241–259). Heidelberg-Berlin: Springer-Verlag.
- Schweighofer, N. (1995). *Computational models of the cerebellum in the adaptive control of movements*. Unpublished doctoral dissertation, University of Southern California.
- Simpson, J. I., Wylie, D. R., & De Zeeuw, C. I. (1996). On climbing fiber signals and their consequence(s). *Behavioral and Brain Sciences*, 19, 384–398.
- Smagt, P. van der. (1998). Cerebellar control of robot arms. *Connection Science*, 10, 301–320.
- Smagt, P. van der, Groen, F., & Schulten, K. (1996). Analysis and control of a rubbertuator arm. *Biological Cybernetics*, 75(5), 433–440.

- Smith, O. J. M. (1959). A controller to overcome dead time. *ISA Journal*, 6, 28–33.
- Sutton, R. S. (1984). *Temporal Credit Assignment in Reinforcement Learning*. Unpublished doctoral dissertation, University of Massachusetts, Amherst, MA.
- Talebi, H. A., Khorasani, K., & Patel, R. V. (1998). Neural network based control schemes for flexible-link manipulators: Simulations and experiments. *Neural Networks*, 11(7–8), 1357–1377.
- Thompson, R. F. (1990). Neural mechanisms of classical conditioning in mammals. *Philosophical Transactions of the Royal Society London: Biological Sciences*, 329, 161–170.
- Zagon, I. S., MacLaughlin, P. J., & Smith, S. (1977). Neuronal population in the human cerebellum: Estimations from isolated cell nuclei. *Brain Research*, 127, 279–282.