
**Principles of selective stabilization
in high-dimensional state spaces: statistical
concepts, dynamic modeling and robotics**

Dissertation

zur Erlangung des Doktorgrades
der Naturwissenschaften
an der Fakultät für Mathematik
der Ruhr-Universität Bochum

vorgelegt von

Jan Hendrik Reimann

Thanks

My thanks go to everyone who supported me during my work on this thesis, especially Gregor Schöner for creating an extremely stimulating research environment and being a motivating and caring supervisor, and Holger Dette for his open-mindedness about interdisciplinary research.

Contents

1. Introduction	1
2. Analysis of behavioral data with the Uncontrolled Manifold method	5
2.1. Introduction	5
2.1.1. Task representations for motor control	5
2.1.2. The Uncontrolled Manifold method	7
2.1.3. Related work	8
2.2. Statistical testing of hypotheses related to the Uncontrolled Manifold	10
2.2.1. Basic definitions and underlying assumptions	10
2.2.2. The UCM basis	11
2.2.3. Estimators	11
2.2.4. Testing hypotheses	13
2.3. Implementation and results	18
2.3.1. Simulation Study	18
2.3.2. Comparison with standard techniques	18
2.4. Conclusions	20
3. A process model of quiet upright stance	24
3.1. Introduction	24
3.1.1. Related work	27
3.2. The model	29
3.2.1. System noise	30
3.2.2. Neural dynamics in the spinal cord	32
3.2.3. Torque generation in muscle-tendon systems	34
3.2.4. Biomechanics of the skeleton	35
3.2.5. Sensor data and neural representations	37
3.2.6. Neural dynamics in the brain	39
3.3. Parameters	43
3.3.1. Time delays	44
3.3.2. Muscle contraction	45
3.3.3. Passive effects	47
3.3.4. Free parameters	50
3.4. Results	51
3.4.1. Data acquisition and analysis	51
3.4.2. Trajectory examples	53
3.4.3. The functional role of spinal and higher feedback	53

3.4.4.	Comparison of variability measures	57
3.5.	Conclusions	60
3.5.1.	Differences between model and experimental data	61
3.5.2.	Effects of closing the eyes	64
3.5.3.	Outlook	66
4.	Movement generation for autonomous robots	67
4.1.	Introduction	67
4.1.1.	Constrained movement trajectories	67
4.1.2.	Selective stabilization of movement parameters	68
4.1.3.	Related work	70
4.2.	The movement generation system	71
4.2.1.	Target acquisition	72
4.2.2.	Obstacle avoidance	78
4.2.3.	Hand orientation	82
4.2.4.	Joint angle limits	84
4.2.5.	Homogeneous damping	86
4.2.6.	Behavioral organization	87
4.3.	Implementation and results	90
4.3.1.	Obstacle avoidance	91
4.3.2.	Complex tasks	94
4.3.3.	The complete system	98
4.4.	Conclusions	100
5.	Conclusions	103
5.1.	Scientific contribution	104
A.	Parameters of the bootstrap simulation study	107
B.	Parameters of the posture model	107
C.	Kinematic and biomechanic equations	108
D.	Calculating the time derivative of the manipulator Jacobian	110
E.	Parameters of the robotics movement generation system	112

1. Introduction

What happens in the brain when humans move? Even faced with a relatively simple movement task like switching on the light or pushing a button, the central nervous system (CNS) faces a large variety of options on several levels. Which effector should be used to push the button? How fast should I push, and with how much force? How should the arm and the rest of the body be held? Which muscles should be used to move the arm in a desired way? These questions represent selections from sets of possible solutions with the same functional result: the light is on, the button has been pressed. Each instance where one solution out of many possible ones is selected adds a degree of *redundancy* to the simple problem of switching on the light.

The problem of having more degrees of freedom available than necessary for a given task was first established by Bernstein (1967) as one of coordination between the abundant degrees of freedom. In the work presented here, we explore possible principles of how the CNS achieves the required coordination. We focus on one exemplary aspect of redundancy: for most common movements, the number of degrees of freedom available to the utilized effector is larger than the number of parameters needed to describe the movement goal. Take the example of switching on the light. At a single point in time, the configuration of a human arm has ten degrees of freedom: three in the sterno-clavicular joint, three in the shoulder, one in the elbow, one in the ulna-radial joint of the lower arm, and two in the wrist. Defining the desired position of the finger tip constrains only three of these ten available degrees of freedom. The sub-space of arm configurations in which the finger position is as desired still has dimension seven. This can be called a redundancy of *solutions*, in the sense that there are infinitely many arm configurations that solve the problem with the given constraints. Even if we assume for a moment that this redundancy of solutions has been resolved and a single solution has been selected, the CNS still has to *realize* that solution by bringing the arm into the chosen configuration. The set of trajectories that start at the current configuration and terminate in this desired one is once more infinite, and to a much larger degree. This latter case can be called a redundancy of *trajectories*.

Formally, the solutions to a movement task correspond to manifolds in a high-dimensional vector space, e.g. the ten-dimensional configuration space describing a human arm. Movements are completely described by trajectories through this high-dimensional space. On the other hand, movements can also be described by specifying trajectories for those variables that are relevant to the task, like the position of the finger tip. Which description does the CNS use to represent and generate movements? Which observable signatures would either mode of representation have that we can look for in behavioral data to answer this question?

Regardless of whether the movement goal is described in the high-dimensional configuration space or the usually low-dimensional task space, what principles can be used to generate a trajectory towards the goal state? One could pre-plan a complete path globally and then move along it, where the “control” ensures that the actual state does not deviate from the planned one. But what happens when the movement parameters change from the values used during planning? Or one could move in a locally optimal direction that brings the current state closer to the target. But is it guaranteed that a global solution will be found in this fashion?

These questions can be asked from two different points of view. The natural scientist wonders how the nervous systems of humans and other animals solve these problems. The engineer looks for any good solution that can be utilized in machines. These two angles are closely connected, though. On the one hand, humans are extremely adept movers: the degree of autonomy and dextrous manipulation we are capable of is unmatched in the animal domain and far beyond the capabilities of current technical solutions, so engineers obviously have a lot to learn from the human nervous system. How humans generate movement is, on the other hand, still not well understood. The CNS is an extremely complex system which is hard to study directly for both technical and ethical reasons. To understand how large scale brain functions like movement generation are structured, researchers depend heavily upon building models of contributing brain areas and the connections between these. But in models, one always makes simplifying assumptions, with the danger of disregarding some aspects that are integral for the functioning of the whole system. When transforming hypothetical principles of how the CNS generates movement into engineering solutions for generating movement for physical systems, such hidden assumptions that were ill-founded become overt by causing the system to fail.

The work presented here contributes to the understanding of movement generation schemes for complex systems at three different points. The first aspect we focus on in Chapter 2 is the question of movement representation: does the CNS represent goals and generate movement trajectories in the high-dimensional configuration space, solving the redundancy problem explicitly, or in the low-dimensional task space? One approach to decide questions of this type is to argue that those variables that are monitored and actively controlled by the CNS during movement should be more stable than other variables, and thus exhibit less variance over repeated trials. To compare the variance of different aspects of movement that are inherently not comparable because they have different units, Scholz and Schöner (1999) developed an analysis method that performs all measures in the high-dimensional configuration space. In Chapter 2, we develop a formal description of this method and the related hypotheses and make underlying assumptions explicit. This enables us to test such hypotheses statistically using the parametric bootstrap method. In contrast to the traditional comparison of variance measures on a population level, the bootstrap method analyzes the variance structure of single subjects and is thus capable of making more fine-grained statements in cases where the movement patterns of different subjects have large variation.

In Chapter 3, we present a model of how the CNS might actually coordinate the

available degrees of freedom to selectively stabilize those directions in configuration space that are relevant to the given task. The example we chose for this is the postural stabilization of upright stance in humans. This is a special case of movement generation in which the “movements” consist of short and relatively small modulations of muscle activation that counter sensed deviations from the stable, upright body configuration. Posture is particularly suited as an example because the lack of active, goal-directed movement largely excludes the redundancy of trajectories between an initial and a goal state. Instead, it focuses on the redundancy of solutions by isolating the process of stabilization and highlighting the question of how different directions of the high-dimensional state space are stabilized selectively. In the sensorimotor loop of balancing upright stance, sensory surfaces like the visual and vestibular systems detect deviations from the stable state. The motor cortex and related areas respond with activation patterns that are sent over descending pathways to the spinal cord, where they modulate the reflex loops that generate muscle contraction by activating motor neurons. The functional aspects of most stages of the sensorimotor loop are constrained by biophysical and neurophysiological data and biomechanics. Detailed models of these stages allow us to formalize hypotheses about the functional structure of the neural dynamics in the brain and include these into the model, thereby closing the modeled sensorimotor loop. The hypotheses can then be tested against by attempting to fit behavioral data via parameter adjustment of the neural dynamics. In our model, we postulate that the CNS stabilizes the body by a feedback term that counters sensed deviations of the low-dimensional task variables given by the body center of mass, head position and head orientation. The redundancy is resolved in the simple fashion of distributing the motor activation over all degrees of freedom that *can* contribute to it in a fashion that minimizes the total magnitude of the configuration change. The movement trajectories generated by this model capture the variability structure of human postural sway well.

The results of Chapter 2 support the notion that humans coordinate the available degrees of freedom to selectively stabilize variables that are relevant to the task at hand. The model presented in Chapter 3 proposes a principle of how the *solution redundancy* can be resolved during stabilization. To extend the scope of our work to the problem of *trajectory redundancy*, one could attempt to augment the model for quiet stance presented in Chapter 3 with more complex neural dynamics capable of generating goal-directed activation patterns with specific time profiles. As the neural dynamics increase in complexity, however, it is harder to argue that the whole system is well-constrained by the behavioral data. To make the point that such a model describes the functional aspects of the neural dynamics, one has to compare the time courses of the model variables with experimental data of actual activation patterns during movement. One way to make such a comparison is to make assumptions where the model variables are represented in the recorded neural activation patterns and estimate them from the experimental data. Another way is to formulate the model in a language of neural activation patterns, so that the variables are directly comparable to experimental data. Either of those ways relies heavily on simplifications and assumptions about the structure of the neural representations and dynamics, though.

To avoid the extreme increase in complexity implied by these comparisons, we switch the scope of our study to a different style of question that is possible to answer with different methods. Instead of making a concrete proposal of how the human CNS solves a problem, we formulate a general principle of how it solves a whole class of problems. To test the feasibility of this principle, we design an autonomous movement generation scheme for a robotic agent based on this principle and show that an agent using this system is capable of solving the tasks it is presented with. In accordance with the results from Chapters 2 and 3, the general principle we use is to represent a given task as desired states for several suitable, low-dimensional variables and coordinate the available degrees of freedom in a way that selectively counteracts detected deviations from the desired state. In Chapter 4, we present an architecture that combines several different tasks like bringing the hand to a target and orienting it in a way appropriate to grasp, while simultaneously avoiding collision with other objects. For each sub-task, a vector field over the relevant low-dimensional task variable is designed that has an attractor at the desired state. These vector fields are transformed into the configuration space of the robotic agent, resolving the redundancy in the same way as described before, and then superposed. The flow of the resulting vector field is used to generate movement trajectories through the high-dimensional configuration space. Instead of pre-planning an explicit solution, trajectories emerge from local stabilization of the currently relevant low-dimensional task variables in this movement generation scheme, indicating that the principle of selective stabilization is also capable of solving the problem of *trajectory redundancy*.

2. Analysis of behavioral data with the Uncontrolled Manifold method

2.1. Introduction

For many movement tasks, humans have more degrees of freedom available than necessary to fulfill the goal. The intuition is that we make use of this freedom in a way related to the concrete task at hand. Bernstein (1967) first described this relationship as redundancy of the motor system and proposed that the central nervous system coordinates the available degrees of freedom in a way that stabilizes the task but leaves other aspects of movement relatively uncontrolled.

The goal of this chapter is to develop a mathematical apparatus that allows us to express this intuition in a formal way. Tasks are represented by manifolds in a high-dimensional vector space that describes configurations of the motor system. These task manifolds are given by the fibers of mappings from the configuration space into another vector space that describes the state of a variable relevant to the task, e.g. the position of the finger tip in a pointing task. A task is redundant if the dimension of the task manifold is non-zero. Utilizing the available freedom in a motor task corresponds to not resisting movements within the task manifold.

This formalization allows us to interpret experimental data from repeated executions of a movement task relative to the task manifold. We can quantify to what degree humans use the available freedom by measuring how much of the variance of repeated solutions of a movement task lies within the task manifold. We can even formulate intuitions about principles of human movement generation as quantifiable hypotheses and use statistical methods to test them.

2.1.1. Task representations for motor control

How does the central nervous system generate goal-directed movement? The brain generates muscle activations and body movements based on afferent information from sensory surfaces and internal states of other brain regions. For any attempt to describe the dynamics of these neural activation patterns, a preliminary question must be asked and answered first: how does the CNS represent the relevant variables related to a movement?

For movements with redundant effector systems like human arms, there are two general answers to this question: the effector state and movement goal can be represented in task-related variables, e.g. the current and desired position of the finger

tip for a point-reaching task, or it can be represented in execution-related variables, which would be the complete configuration of the arm. These answers are not mutually exclusive. It is obvious that the CNS does have knowledge about the location of reaching targets and reach-relevant body parts in three-dimensional space, from e.g. the visual and auditory system. This has to be combined with proprioceptive information about the state of each muscle and joint, i.e. the configuration of the whole body or effector in question. And ultimately, the arm is moved by a set of muscles contracting due to activation of motorneurons, the specifics of which depend upon the current configuration of the arm.

At some point between the visual fixation of a movement target and the neural activation of muscle spindles the transformation between these two modes of representation has to be made. This could either be close to the sensory surfaces, implying that the visually sensed location of the movement target is transformed into a representation of a desired body configuration and the remainder of the neural processing is devoted to activating the motor system in an appropriate way to bring the body into this desired state and deal with perturbations that are encountered on the way. Or the transformation is made close to the motor surfaces, implying that the movement is planned and monitored in terms of the end-effector state relative to the target, and the details of how to realize a desired end-effector movement is left to the low-dimensional neural circuits of the periphery. It is also possible that both modes of representation are used in parallel during at least some stages of neural processing and different aspects of movement are processed in different reference frames (Saltzman, 1979; Saltzman & Kelso, 1987; Lacquaniti, 1989).

Both possible answers to this question are supported by experimental data to some degree. Soechting and Lacquaniti (1981) studied simple human point-reaching movements and reported that the ratio of angular velocities at the shoulder and elbow is invariant during the deceleratory phase of the movement and interpreted that as evidence for joint-level representation of movement states.

Feldman and colleagues argue that the CNS specifies positional references for each kinematic degree of freedom by descending commands, resulting in a state where external load and internal muscle force acting on a joint are at equilibrium. This so-called Equilibrium Point Hypothesis (EPH) implies that the neural processing is done in the reference frame of kinematic configurations at least in some part (Feldman, 1986; Feldman & Levin, 1995). In a different version of the EPH, however, it is conjectured that the CNS only specifies the equilibrium point for the end-effector, leaving the process of stabilizing it there to the peripheral motor circuits (Bizzi & Accornero, 1984; Flash, 1987).

Morasso (1981) was among the first to point out that in pointing movements, the hand trajectories are very reproducible while the joint trajectories exhibit a great deal of variance. Another invariance found at the end-effector level is the power law relationship between the speed and curvature of the pen trajectories in drawing and writing (Lacquaniti, Terzuolo, & Viviani, 1983), again contrasted by the lack of such an invariance on the joint level.

A theoretical study by Flash and Hogan (1985) makes another case for represen-

tation on the end-effector level. The authors conjecture that the CNS attempts to generate movements that are as smooth as possible, adhering to a minimum jerk principle. Although some hand trajectories show complex features such as varying curvature or multiple velocity peaks, these arise naturally under the minimum jerk constraint. Furthermore, movements that are maximally smooth in hand space predict the joint space trajectories measured experimentally surprisingly well.

This argument received criticism from Kawato and colleagues, who point out that the shape of hand trajectories between two points depend substantially upon the location of those two points in the work space of the hand, specifically that movements approaching the boundaries of the reachable work space are usually more curved (Uno, Kawato, & Suzuki, 1989). The minimum jerk model cannot predict that. The authors propose to minimize joint torques across the whole trajectory instead and show that the trajectories generated by this criterion also bear remarkable similarity to actual human movements.

2.1.2. The Uncontrolled Manifold method

The question of which variables related to the generation of movement are represented by the CNS was approached from a different angle by Scholz and Schöner (1999). They make the point that variables that are actively controlled by the CNS would be more stable than other variables during goal-directed movement. Stability is used in the dynamic systems sense as the capability of a system to react to transient perturbations and restore the desired state of a variable (Schöner, 1995).

The formation of movement trajectories is subject to perturbations from many different sources: the brain is an inherently noisy system, and no two patterns of neural activation are ever completely the same (see Section 3.2.1 for a more detailed treatment of neuronal noise). To achieve a movement goal reliably and with repeatability, the CNS must counteract the perturbations induced by noise and possible external influences by monitoring the relevant variables and changing the movement command according to detected deviations from the desired state. The stability of the actively controlled variables is thus expected to be higher than the stability of the uncontrolled variables.

One observable signature of stability of a dynamic variable at a fixed point is the variance of repeated measurements near that point (Schöner, 1990). If the CNS employs task-related variables to generate movements, like the hand position, we expect the variance of task-related variables to be high compared to variables that are not task-related, like the elbow joint angle. If a movement is generated using variables related to the motor system, no such difference in variance is expected. This is the line of reasoning that led Morasso to favor a task-level representation after observing that hand trajectories are substantially less variable than joint angle trajectories in a reaching task (Morasso, 1981).

This difference is hard to quantify, though. While one class of trajectories can be called more or less variable than another, this difference is not immediately measurable, because joint angles and hand trajectories have fundamentally different units

of measurement: the former are measured in radians, the latter in meters.

The solution to this problem proposed by Scholz and Schöner was to find a single set of generalized coordinates that describes the movement completely, called *configuration space*, usually the joint angles of the effector used in a given task. For any candidate of a task variable that is actively controlled by the CNS, a given desired state corresponds to a sub-manifold in the configuration space, called the *UnControlled Manifold* (UCM). Differences in configuration space that lie within this sub-manifold leave the candidate variable invariant, whereas differences that are orthogonal to it affect the candidate variable. If the candidate variable is indeed actively stabilized by the CNS, perturbations orthogonal to the UCM are expected to be resisted more strongly than perturbations within the UCM. Thus the variance of the configuration is expected to be larger along the UCM than orthogonal to it.

If a candidate variable is indeed actively controlled by the CNS, the resulting difference in variance along and orthogonal to the UCM should be testable by statistical methods. Scholz and Schöner initially applied this analysis to a sit-to-stand task, taking position of the body center of mass (CoM), head and hands in the sagittal plane as candidate variables (Scholz & Schöner, 1999). The configuration of the body was described by 8 angles between the body segments and the horizontal axis. The difference of the segment angle data to their mean was analyzed at three time slices by calculating the mean body configuration and linearly approximating the Uncontrolled Manifold at that mean. The segment angle differences were projected onto the linear approximation of the UCM and its orthogonal complement. The squared norm of the projections was divided by the dimension of the subspace to provide a measure of the total variance within each subspace.

To support the hypothesis that a candidate variable is indeed actively controlled by the CNS, the variance along the UCM should be larger than the variance orthogonal to it. The difference between these two measures was tested by an Analysis of Variance (ANOVA) and reported as significant in several combinations of condition and candidate variable.

Analyzing the difference between these projected variances relies on the assumption that the compared variables are distributed normally (Howell, 2010). This assumption is questionable, as the difference between the total variance within each subspace is a complex feature of the geometrical structure of the covariance in the configuration space. In the remainder of this chapter, we will deliver a more formal treatment of the UCM method that allows us to be more explicit about the hidden assumptions being made. We will then describe a parametric bootstrap test as an alternative to test statistical hypotheses about the variance structure of movement-related variables. Several different hypotheses are presented. The merit of the parametric bootstrap test is demonstrated by comparing it with the ANOVA method on a data set that was previously published in the literature.

2.1.3. Related work

The UCM method has been widely and successfully applied to the analysis of behavioral data from movement experiments. Studies on hand movements revealed that indeed those aspects of the arm configuration that are relevant to the current task are more stable than other, like the hand position in a point-reaching task (Tseng, Scholz, Schöner, & Hotchkiss, 2003) or the pistol orientation in a shooting task (Scholz, Schöner, & Latash, 2000). More recent studies showed that these effects depend upon the handedness of the subject (Tseng, Scholz, & Galloway, 2009) and decrease with old age (Verrel, Lövdén, & Lindenberger, 2012)

A challenging motor task that is almost unique to humans is bipedal upright quiet stance. While postural sway has been studied extensively, this is usually done by analyzing the patterns of low-level variables like the body center of mass or center of pressure on the support surface. The standing human body is modeled as a single-link inverted pendulum with one degree of freedom at the ankle, disregarding other joints like the knee, hip and vertebrae. A UCM analysis of postural sway in the sagittal plane in 6 degrees of freedom revealed very strong coordination effects in the joint angle trajectories: relevant variables such as position of the CoM or head are selectively stabilized (Hsu, Scholz, Schöner, Jeka, & Kiemel, 2007). Furthermore, when joint movements are induced by transient perturbations, those combinations of joint movements that affect these variables decay much faster than combinations that do not affect them, indicating active control by the CNS (Scholz et al., 2007). Even without artificially induced perturbations, a time-lagged auto-correlation analysis of variability along and orthogonal to the UCM revealed higher persistence of variations that leave important variables invariant (Verrel, Pradon, & Vuillerme, 2012). These effects are discussed in greater detail in Chapter 3.

Later the UCM method was applied more generally both methodologically and in terms of subject matter. Latash and colleagues used it to analyze the generation of ramp profiles from a combination of finger forces, finding significantly more variation in force combinations that leave the sum of forces invariant (Kang, Shinohara, Zatsiorsky, & Latash, 2004). Also on the force level, Chang and colleagues analyzed rhythmic hopping in humans, finding that joint torque combinations are coordinated to stabilize the vertical ground-reaction force at the time of landing and takeoff (Yen, Auyang, & Chang, 2009; Yen & Chang, 2010).

Other researchers have approached the same general question of inferring principles of how the CNS represents and processes information about movement states and goals from behavioral data using different techniques. Müller and Sternad approached the question by postulating that if a task variable is actively stabilized by coordination between the configuration variables that reduces the variance of the task variable, then removing that coordination by randomizing the configuration variables should increase the variance of the task variable (Müller & Sternad, 2003, 2004). This randomization method has the benefit of being independent of a metric in the configuration space and is better capable of dealing with strong non-linearities, but more sensitive to the choice of basis in the configuration space. The two approaches have

been compared in detail elsewhere (Schöner & Scholz, 2007; Verrel, 2011). Cusumano and Cesari (2006) also consider the dependency of the UCM variance analysis upon the choice of basis in the configuration space, specifically the alignment of the UCM relative to the basis vectors. This point, and the question of coordinate dependence in general, is also discussed by other researchers and still remains somewhat unclear in the community (Hogan, Sternad, Park, & Mu, 2010; Verrel, 2010).

2.2. Statistical testing of hypotheses related to the Uncontrolled Manifold

In this section, we develop the UCM analysis method formally and present a number of hypotheses about the covariance structure in repeated movement tasks. A parametric bootstrap method is used to test the hypotheses statistically, the applicability of which is appraised in a simulation study.

2.2.1. Basic definitions and underlying assumptions

We begin by defining basic concepts and making some assumptions explicit. The data we deal with are from behavioral experiments. Usually different subjects will repeat a task several times in different conditions. We ignore the range of subjects and conditions and for the time being focus on the repeated task executions of a single subject in a single condition. All data points are measured at the same relative time point during movement. The task is repeated N times and the result of the i -th trial is described by a random variable

$$X^{(i)} \in \mathbb{E}. \quad (2.1)$$

The *configuration space* or *execution space* \mathbb{E} is a real vector space of dimension n .

The general hypothesis is that the CNS cares about some directions of \mathbb{E} more than others, which has an effect on the covariance matrix of the random variable $X^{(i)}$. To specify such a hypothesis formally we make

Definition 1. *A task function or task variable is a surjective mapping*

$$f : \mathbb{E} \rightarrow \mathbb{T}_f \quad (2.2)$$

from the configuration space to some vector space \mathbb{T} of dimension k_{\perp} , called task space. Task functions for which $n > k_{\perp}$ are called redundant.

Definition 2. *For any $x \in \mathbb{E}$, the fiber*

$$M_f(x) = f^{-1}(f(x)) = \{x' \in \mathbb{E} : f(x') = f(x)\} \quad (2.3)$$

is called the UnControlled Manifold (UCM), or space of task-equivalent solutions of

f at x . The linear subspace

$$U_f(x) = \{x' \in T_x E : Df(x)(x') = 0\} \quad (2.4)$$

is called the space of null-changes of f at x .

$U_f(x)$ is the linearization of $M_f(x)$ at x , which we will usually abbreviate to M_f and U_f if the choice of x is unambiguous. We will also refer to the orthogonal complement

$$O_f(x) = \{x' \in T_x E : \langle x', x'' \rangle = 0 \forall x'' \in U_f(x)\} \quad (2.5)$$

of U_f , the space of changes that also affect the task variable f .

Assumption 1. For a single subject and a fixed condition, the trial data are i.i.d. with the same normal distribution

$$X^{(i)} \sim \mathcal{N}(\mu, \Sigma). \quad (2.6)$$

2.2.2. The UCM basis

Does the variability structure of the measured data confirm our hypothesis that the CNS cares about a specific task variable and stabilizes directions of the configuration space that affect it more than other directions? To answer this question, we express the data in a basis that corresponds to these directions. Without loss of generality, we can assume that our data set

$$X^{(i)}, \quad 1 \leq i \leq N \quad (2.7)$$

has mean $\mu = 0$.

Let f be any hypothetical task variable. To project the data points to the space of null changes U_f and its orthogonal complement O_f , let

$$J = Df(x)|_0 = U \cdot S \cdot V^T \quad (2.8)$$

be a singular value decomposition of the Jacobian of f at the origin, where $S \in \mathbb{R}^{k_\perp \times n}$ is a diagonal matrix with the non-negative singular values of J in decreasing order, $V \in \mathbb{R}^{n \times n}$ is orthonormal.

The column vectors v_i of V form an orthonormal basis, and assuming J has full rank, the first k_\perp vectors v_1, \dots, v_{k_\perp} span O_f , while the last k_\parallel vectors $v_{k_\perp+1}, \dots, v_n$ span U_f . Forming

$$E_\perp = (v_1 \ \dots \ v_{k_\perp}) \in \mathbb{R}^{n \times k_\perp}, \quad E_\parallel = (v_{k_\perp+1} \ \dots \ v_n) \in \mathbb{R}^{n \times k_\parallel}$$

gives the projection matrices

$$K_\perp = E_\perp E_\perp^T, \quad K_\parallel = E_\parallel E_\parallel^T = 1_{n \times n} - K_\perp \in \mathbb{R}^{n \times n}$$

to O_f and U_f .

2.2.3. Estimators

For each candidate task variable f , we have linearized the UCM at the mean, constructed a basis of the execution space \mathbb{E} that separates it into directions along the UCM and orthogonal to it, along with projection matrices to these sub-spaces. This allows us to analyze the structure of the covariance matrix Σ relative to the candidate task variable.

We estimate Σ by its maximum likelihood estimator

$$\widehat{\Sigma} = \widehat{\Sigma}(X) = \frac{1}{N} \sum_{i=1}^N X^{(i)} X^{(i)T}. \quad (2.9)$$

A well known result in multivariate statistics shows that

$$\widehat{\Sigma} \sim W_n(N, \Sigma), \quad (2.10)$$

where $W_n(N, \Sigma)$ denotes the multivariate Wishart Distribution with dimension n , N degrees of freedom and covariance matrix Σ (e.g., Muirhead, 1982, p. 87).

The total magnitude of the variance in each subspace can be measured by the trace of the covariance matrix projected to that subspace. For the orthogonal space, we estimate that quantity

$$\text{tr}(K_{\perp}^T \Sigma K_{\perp}) = \text{tr}(E_{\perp}^T \Sigma E_{\perp}) \quad (2.11)$$

by

$$\begin{aligned} V_{\perp} &= \frac{1}{k_{\perp}} \text{tr}(K_{\perp}^T \widehat{\Sigma} K_{\perp}) = \frac{1}{k_{\perp}} \frac{1}{N} \sum_{i=1}^N X^{(i)T} E_{\perp} E_{\perp}^T X^{(i)} \\ &= \frac{1}{k_{\perp}} \text{tr}(E_{\perp} E_{\perp}^T \widehat{\Sigma}) = \frac{1}{k_{\perp}} \text{tr}(E_{\perp}^T \widehat{\Sigma} E_{\perp}), \end{aligned}$$

where

$$E_{\perp}^T \widehat{\Sigma} E_{\perp} \sim W_{k_{\perp}}(N, E_{\perp}^T \Sigma E_{\perp}). \quad (2.12)$$

Similarly, we obtain

$$V_{\parallel} = \frac{1}{k_{\parallel}} \text{tr}(K_{\parallel}^T \widehat{\Sigma} K_{\parallel}) = \frac{1}{k_{\parallel}} \text{tr}(E_{\parallel}^T \widehat{\Sigma} E_{\parallel}) \quad (2.13)$$

with

$$E_{\parallel}^T \widehat{\Sigma} E_{\parallel} \sim W_{k_{\parallel}}(N, E_{\parallel}^T \Sigma E_{\parallel}) \quad (2.14)$$

as an estimate for the total variance along the UCM.

These quantities V_{\perp} and V_{\parallel} estimate the variance of X projected to O_f and U_f per

degree of freedom. They can also be written as

$$V_{\perp} = \frac{1}{k_{\perp}} \frac{1}{N} \sum_{i=1}^N \|K_{\perp} X^{(i)}\|^2, \quad V_{\parallel} = \frac{1}{k_{\parallel}} \frac{1}{N} \sum_{i=1}^N \|K_{\parallel} X^{(i)}\|^2, \quad (2.15)$$

which is closer to how they were introduced in (Scholz & Schönner, 1999).

The ratio of these measures, normalized by the dimensionality of the corresponding subspace,

$$S_f(X) = \frac{V_{\parallel}}{V_{\perp}} \quad (2.16)$$

is called the UCM-signature of the data set X , or f -signature when confusion about the task-variable has to be avoided.

2.2.4. Testing hypotheses

If the CNS monitors and actively stabilizes the candidate task variable f then we expect the variance of X along the UCM, V_{\perp} to be small compared to the variance of X perpendicular to it. This difference in magnitude is captured by the UCM-signature $S_f(X)$ – if the hypothesis about f holds, then the f -signature of the data should be larger than 1. In this section, we formalize this and other notions as statistically testable hypotheses and propose a parametric bootstrap method to derive a critical value for rejecting the corresponding null hypotheses.

Hypothesis 1. (*weak task-variable hypothesis*) For a given task function $f : \mathbb{E} \rightarrow \mathbb{T}_f$, the task-equivalent variance V_{\parallel} per degree of freedom is larger than the non-task-equivalent variance V_{\perp} per degree of freedom.

The corresponding null hypothesis

$$H_0 : \quad V_{\parallel} \leq V_{\perp} \quad (2.17)$$

states that the variance per degree of freedom in task-relevant directions (V_{\perp}) is larger or equal to the the variance in task-irrelevant directions (V_{\parallel}).

For a given data set, we need a statistical test that rejects the null hypothesis 2.17 when the f -Signature of the data is significantly larger than 1. As the distribution of the random variable $S_f(X)$ is comparatively complex, finding an analytical solution for the critical value $c_{1-\alpha}$ for the test which rejects the null hypothesis whenever

$$S_f = \frac{V_{\parallel}}{V_{\perp}} > c_{\alpha}, \quad (2.18)$$

we use the following parametric bootstrap procedure.

2.2.4.1. Bootstrap method

A statistical test provides an answer to the question: “How likely is it that an effect was observed just by chance and would disappear if more data was examined?” The parametric bootstrap method attempts to decide this by first asking: “If the effect was not there, how are the data expected to look like?” This is done by finding parameters for a distribution that fit the data set as closely as possible under the constraint that the observed effect is *not* present. A large number of data sets, the bootstrap data, is generated from this distribution. The original data set is then compared to these bootstrap data sets. If the effect measured in the original data is very extreme compared to what is expected from the bootstrap data sets, then it is very unlikely that the effect was observed just by chance. This line of reasoning is translated into a formal test in the following paragraphs.

Let $0 < \alpha < 1$ be the level of the test.

- (A) Estimate the covariance matrix Σ from the original data using the maximum likelihood estimator $\widehat{\Sigma}$ specified in Equation 2.9. Determine the matrix

$$\Sigma^* = \arg \min_A \left\{ \text{tr} \left((\widehat{\Sigma} - A)^T (\widehat{\Sigma} - A) \right) \left| \begin{array}{l} A = A^T \\ A > 0 \\ \frac{\text{tr}(K_{\perp}^T A K_{\perp})}{k_{\perp}} = \frac{\text{tr}(K_{\parallel}^T A K_{\parallel})}{k_{\parallel}} \end{array} \right. \right\}. \quad (2.19)$$

This is the covariance matrix that is closest to the estimated $\widehat{\Sigma}$ while satisfying the constraint that the total variances in the two relevant subspaces U_f and O_f are the same, i.e. that “just barely” satisfies the null hypothesis 2.17.

- (B) For $b = 1, \dots, B$, generate N data points

$$X^{(1)}(b), \dots, X^{(N)}(b)$$

from a $\mathcal{N}(0, \Sigma^*)$ distribution. Calculate the statistics $S_f(b)$ of this data set. This gives a series of B bootstrap data sets and corresponding values of $S_f(b)$ that we can compare with the original data set.

- (C) Sort the B values obtained in step (B) in ascending order, such that

$$S_{(1)} \leq S_{(2)} \leq \dots \leq S_{(B)}.$$

Use

$$\hat{c}_{1-\alpha} = S_{(\lfloor (1-\alpha)B \rfloor)}$$

as an approximation of the critical value $c_{1-\alpha}$. This is the UCM-signature of the bootstrap data set at the $(1 - \alpha)$ quantile, indicating that a value that is even higher can be regarded as extreme.

The bootstrap test is now defined by rejecting the null hypothesis (2.17) whenever

$$S_f > \hat{c}_{1-\alpha},$$

where S_f is calculated from the original data. One can prove that for $N \rightarrow \infty$, this test has level α , i.e.

$$P_{H_0}(S_f > \hat{c}_{1-\alpha}) \xrightarrow{N \rightarrow \infty} \alpha. \quad (2.20)$$

Alternatively, we can use the bootstrap method to compute the p -value for the given measure S_f , i.e. the probability that the hypothesis is true and a UCM-signature as extreme as S_f or larger is observed. The bootstrap method is used in the same way as detailed above. Then we count how many bootstrap samples have a UCM-signature S that is even larger than the one from the original data

$$b_p = B - \max_b S_{(b)} < S(X). \quad (2.21)$$

The ratio of the number of bootstrap samples with a UCM-signature even larger than what was measured over the total number of bootstrap samples yields the p -value

$$p = \frac{b_p}{B}. \quad (2.22)$$

2.2.4.2. Simulation study

The statement (2.20) is asymptotical, it does not say much about the probability for finite N . We use a simulation study to check whether the statement is also approximately correct for small sample sizes N . To test the limit case, we choose a feasible covariance matrix Σ for which the null hypothesis 2.17 is satisfied with equality. Then we create a large number of simulated data sets with sample size N similar to what is feasible in an experiment. For each data set, we run the bootstrap test as detailed above. Although the Null hypothesis 2.17 is satisfied, some of these tests will reject it, i.e. result in a false positive. If the ratio of these false positives much larger than the expected value α , the bootstrap method has to be rejected as not applicable for small sample sizes N . If, on the other hand, the ratio is close to the expected value α , the simulation study confirms that the bootstrap method is valid.

In detail, the simulation study consists of the following steps:

- (0) Choose a covariance matrix Σ for which

$$\frac{1}{k_{\parallel}} \text{tr}(K_{\parallel}^T \Sigma K_{\parallel}) = \frac{1}{k_{\perp}} \text{tr}(K_{\perp}^T \Sigma K_{\perp}). \quad (2.23)$$

- (1) Set $\tilde{f} = 0$. This variable counts the false positives.
 (2) For $s = 1, \dots, R$ do

- (a) Generate independent random variables $X^{(1)}, \dots, X^{(N)}$ from a $\mathcal{N}(0, \Sigma)$ distribution and calculate the statistics S_f .
- (b) Calculate the critical value $\hat{c}_{1-\alpha}$ by using the bootstrap procedure described above.
- (c) If $S_f > \hat{c}_{1-\alpha}$, then set $\tilde{f} = \hat{f} + 1$.

(3) Set

$$\hat{f} = \frac{\tilde{f}}{R}.$$

If for a realistic sample size N the error ratio \hat{f} of false positives is close to α , the critical value $c_{1-\alpha}$ is calculated approximately right by the bootstrap procedure. Several of these simulation studies were carried out, please see Section 2.3.1 for the results.

2.2.4.3. Other hypotheses

Often it is not clear what behavioral aspects the CNS really controls, a variety of different task functions is conceivable. Testing these for the weak task-variable hypothesis might not help much: a positive result just means a failure to falsify the assumption that the CNS cares about one specific task function f . Geometrically, this means that the extension of the covariance ellipsoid is larger in the task-irrelevant directions of execution space than in task-relevant directions, which is a rather weak statement and leaves much of the shape of the covariance ellipsoid unknown. The following hypothesis provides some more insight in the shape by comparing two different hypothetical task functions in how well they describe the covariance ellipsoid.

Hypothesis 2. (*task-variable comparison hypothesis*) Let $f_1 : \mathbb{E} \rightarrow \mathbb{T}_1, f_2 : \mathbb{E} \rightarrow \mathbb{T}_2$ be two task functions. The f_1 -signature of the data set is larger than the f_2 -signature.

To obtain the critical value c_α for which to reject the null hypothesis

$$H_0 : S_{f_1} \leq S_{f_2} \tag{2.24}$$

whenever

$$\frac{S_{f_1}}{S_{f_2}} = \frac{V_1^{\parallel} V_2^{\perp}}{V_1^{\perp} V_2^{\parallel}} > c_{1-\alpha}, \tag{2.25}$$

we use the same bootstrap procedure as described in section 2.2.4.1. In the case of Hypothesis 2 the covariance matrix of the critical case is given by

$$\Sigma^* = \arg \min_A \left\{ \text{tr} \left((\hat{\Sigma} - A)^T (\hat{\Sigma} - A) \right) \left| \begin{array}{l} A = A^T \\ A > 0 \\ (*) \end{array} \right. \right\}, \tag{2.26}$$

where the last constraint

$$(*) : \frac{\frac{1}{k_1^{\parallel}} \operatorname{tr}(K_1^{\parallel T} A K_1^{\parallel}) \frac{1}{k_2^{\perp}} \operatorname{tr}(K_2^{\perp T} A K_2^{\perp})}{\frac{1}{k_1^{\perp}} \operatorname{tr}(K_1^{\perp T} A K_1^{\perp}) \frac{1}{k_2^{\parallel}} \operatorname{tr}(K_2^{\parallel T} A K_2^{\parallel})} = 1 \quad (2.27)$$

means that Equation 2.24 is satisfied with equality.

So far we analyzed data that consisted of repetitions of the same behavioral task in the same, controlled condition. A standard experimental method is to probe a system by letting it perform the same task under *different* experimental conditions while keeping all other influences as constant as possible and observing how the change affects the behavior. In many cases, different conditions are set up by introducing perturbations like distractors or physical constraints. These perturbations usually increase the variance of the data (Schöner & Kelso, 1988). Often this increase in variance is not homogeneous, though. Under our basic assumption that the CNS stabilizes the task-relevant directions of execution space more than the task-irrelevant directions, we expect the increase in goal-equivalent variance V_{\parallel} to be larger than the increase in non-goal-equivalent variance V_{\perp} , as the latter is actively reduced by the CNS. This expectation translates into

Hypothesis 3. (*data comparison hypothesis*) Let $f : \mathbb{E} \rightarrow \mathbb{T}$ be a task function, X and Y two data sets. The UCM-signature of X is larger than the UCM-signature of Y .

As before, we obtain the critical value $c_{1-\alpha}$ for which to reject the null hypothesis

$$H_0 : S_f(X) \leq S_f(Y) \quad (2.28)$$

whenever

$$F_f''(X, Y) = \frac{S_f(X)}{S_f(Y)} = \frac{V_{\parallel}(X)V_{\perp}(Y)}{V_{\perp}(X)V_{\parallel}(Y)} > c_{1-\alpha}, \quad (2.29)$$

by the bootstrap procedure. We calculate

$$(\Sigma_X^*, \Sigma_Y^*) = \arg \min_{A_X, A_Y} \left\{ \begin{array}{l} \operatorname{tr} \left((\widehat{\Sigma}_X - A_X)^T (\widehat{\Sigma}_X - A_X) \right) \\ + \operatorname{tr} \left((\widehat{\Sigma}_Y - A_Y)^T (\widehat{\Sigma}_Y - A_Y) \right) \end{array} \middle| \begin{array}{l} A_X = A_X^T \\ A_Y = A_Y^T \\ A_X > 0 \\ A_Y > 0 \\ (*) \end{array} \right\}, \quad (2.30)$$

where the last constraint

$$(*) : \frac{\operatorname{tr}(K_X^{\parallel T} A_X K_X^{\parallel}) \operatorname{tr}(K_Y^{\perp T} A_Y K_Y^{\perp})}{\operatorname{tr}(K_X^{\perp T} A_X K_X^{\perp}) \operatorname{tr}(K_Y^{\parallel T} A_Y K_Y^{\parallel})} = 1 \quad (2.31)$$

means that these are the covariance matrices that best approximate $\widehat{\Sigma}_X = \widehat{\Sigma}(X)$ and

$\widehat{\Sigma}_Y = \widehat{\Sigma}(Y)$ while satisfying Equation 2.28 with equality.

This is just one example of a class of hypotheses that compare the variance structures of different data sets. A similar but weaker hypothesis is that the goal-equivalent variance increases between the two conditions

Hypothesis 4. (*data comparison hypothesis, U_{\parallel}*) Let $f : \mathbb{E} \rightarrow \mathbb{T}$ be a task function, X and Y two data sets. In the task-irrelevant subspace U_{\parallel} , the variance of X is larger than the variance of Y .

A critical value $c_{1-\alpha}$ for a test that rejects the null-hypothesis

$$H_0 : V_{\parallel}(X) \leq V_{\parallel}(Y) \quad (2.32)$$

is determined with the bootstrap method as before. This and another version that hypothesizes $V_{\perp}(X) > V_{\perp}(Y)$ will be used in the following section to compare the bootstrap method with conventional methods.

2.3. Implementation and results

The methods described here were tested on simulated data and data from human point-reaching experiments. All data sets consisted of joint angles $\theta_1, \dots, \theta_{10}$ for an upper arm with 10 degrees of freedom.

The bootstrap test was implemented in Matlab using mostly standard routines. In order to calculate Σ^* , i.e. determining the minima given in Equations 2.19, 2.26 and 2.30, the symmetric 10×10 matrices were parameterized by the 55 values of the diagonal and upper right triangle. The resulting symmetric matrix was transformed into a symmetric positive definite matrix by taking the exponential. It can be shown that the whole space of symmetric, positive matrices is parameterized by 55 values in this way (Baker, 2002). The Matlab function `fmincon` was then invoked to find the minimum that observes the last constraint (Equations 2.19, 2.27 and 2.31).

In all cases, $B = 500$ bootstrap samples were generated.

2.3.1. Simulation Study

To check whether the critical values $\hat{c}_{1-\alpha}$ generated by the bootstrap test are actually close to the real values $c_{1-\alpha}$ for the rather low sample size of $N = 20$, the simulation study described in section 2.2.4.2 was carried out with $S = 10000$ runs. For this simulation study, a roughly anthropomorphic arm with spherical joints at the sterno-clavicular junction and shoulder and paired hinge joints at the elbow and wrist was simulated. For each run of the simulation study, a trial data set with $N = 20$ data points was generated by drawing pseudo-random numbers from a $\mathcal{N}(\mu, \Sigma)$ distribution using the Matlab function `randn`. The mean μ was chosen in a way that the arm configuration roughly corresponded to a human reaching for some object in front of him. The covariance matrix was proportional to the identity matrix. See

Hyp.	task variable	data distribution	ratio of false positives
1	$\mathbf{p}(\theta) \in \mathbb{R}^3$	$X \sim \mathcal{N}(\mu_0, 0.1 \cdot I_{10 \times 10})$	$\hat{f} = 0.0573$
2	$f_1 = \mathbf{p}(\theta) \in \mathbb{R}^3,$ $f_2 = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \end{pmatrix} \mathbf{p}(\theta) \in \mathbb{R}^2$	$X \sim \mathcal{N}(\mu_0, 0.1 \cdot I_{10 \times 10})$	$\hat{f} = 0.0471$
3	$\mathbf{p}(\theta) \in \mathbb{R}^3$	$X \sim \mathcal{N}(\mu_0, 0.1 \cdot I_{10 \times 10})$ $Y \sim (\mu_0, 0.2 \cdot I_{10 \times 10})$	$\hat{f} = 0.0554$

Table 2.1.: Results of the simulation study for the weak task variable hypothesis (1), the comparative task variable hypothesis (2) and the data comparison hypothesis (3). The finger tip position $\mathbf{p}(\theta)$ depends upon the random joint configuration $\theta = X, Y$.

Appendix A for details and parameters. Different simulation studies were carried out for the hypotheses 1, 2 and 3. In each case, the null hypotheses 2.17, 2.24 and 2.28 respectively were satisfied with equality by the covariance matrix Σ . The candidate task variables were components of the cartesian finger tip position $\mathbf{p}(\theta)$.

If the bootstrap test generates feasible values, we expect the ratio of false positives \hat{f} in the simulation study to be close to the level of the test α . This was indeed the case in all three studies. Table 2.1 provides a summary of the tested hypotheses, the hypothetical task variables, the varying parameters and the results. The test level was $\alpha = 0.05$ in both cases. The ratio of false positives \hat{f} is close to the test level in all three cases.

2.3.2. Comparison with standard techniques

To demonstrate the benefits of the parametric bootstrap analysis presented in this chapter, we apply the method to reanalyze a previously published study by Sandra Freitas and colleagues (de Freitas, Scholz, & Stehman, 2007). We give a summary of this study and the relevant results.

In this experiment, subjects were sitting at a table and reached for a target that was presented on a touch screen in front of them. In the random condition (RD), the target had a $\frac{1}{3}$ chance either to jump to the left at the onset of movement, jump to the right, or stay fixed. In the blocked condition (BL), the target was always fixed. The two conditions were not mixed and the subjects were always aware which condition they were currently in.

The research question underlying this study was that the CNS would control the movement trajectory less tightly if the target was liable to jump. To answer this question, de Freitas and colleagues measured the 10 DoF joint angles with an optical motion capture system and projected the variances of these data sets onto the subspaces U_f and O_f for two different task variables *movement direction* and *movement*

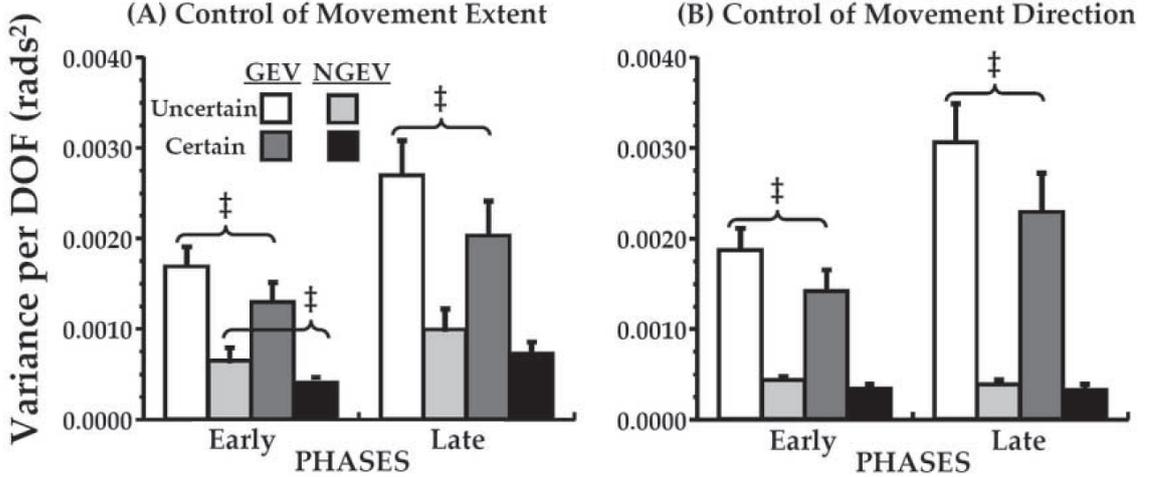


Figure 2.1.: Variance per degree of freedom, averaged across subjects and combined across the first half (early) and second half (late) of the movement. Task variables are movement extent (left) and direction (right). Error bars depict standard errors, ‡ indicates significant differences.

extent for each subject. Then they averaged the variance data over two movement phases, *early* and *late*, separated by the time of peak velocity for each single trial. On the averaged data, separate repeated measures ANOVAs was carried out for each hypothetical task variable to examine the effects of target condition (RD vs. BL) and variance component (V_{\parallel} vs. V_{\perp}). The ANOVAs showed that V_{\parallel} in the RD condition was significant larger than in the BL condition, in both movement phases and for both task variables. In contrast, the increase in V_{\perp} was significant only for movement extent in the *early* movement phase. Figure 2.1, reproduced here from (de Freitas et al., 2007) with permission of the authors, illustrates these results.

As the averaging over movement phases is at odds with the underlying assumptions of the bootstrap method, we reanalyze the data to make the results of the ANOVA comparable with the results of the bootstrap method. First, we do not analyze movement extent and movement direction separately, but make one analysis with the three-dimensional pointer position \mathbf{p} as hypothetical task variable. Secondly, we do not average over time, but analyze at two separate points of time, peak velocity (t_1) and movement termination (t_2).

Separate repeated measures ANOVAs were run for each of these four time points to test for differences in V_{\parallel} and V_{\perp} between the two conditions RD and BL. The average across-subjects results for V_{\parallel} and V_{\perp} are presented in Figure 2.2. At peak velocity and movement termination, the magnitude of both V_{\parallel} and V_{\perp} was significantly higher for the RD condition compared to the BL condition ($V_{\parallel}, t_2 : F_{[1,10]} = 6.0, p < 0.05, \eta^2 = 0.376$), ($V_{\parallel}, t_4 : F_{[1,10]} = 6.9, p < 0.05, \eta = 0.409$), ($V_{\perp}, t_2 : F_{[1,10]} = 5.1, p < 0.05, \eta = .340$), ($V_{\perp}, t_4 : F_{[1,10]} = 6.5, p < 0.05, \eta = .392$).

The ANOVA asserts the statistical significance of the difference in the variance

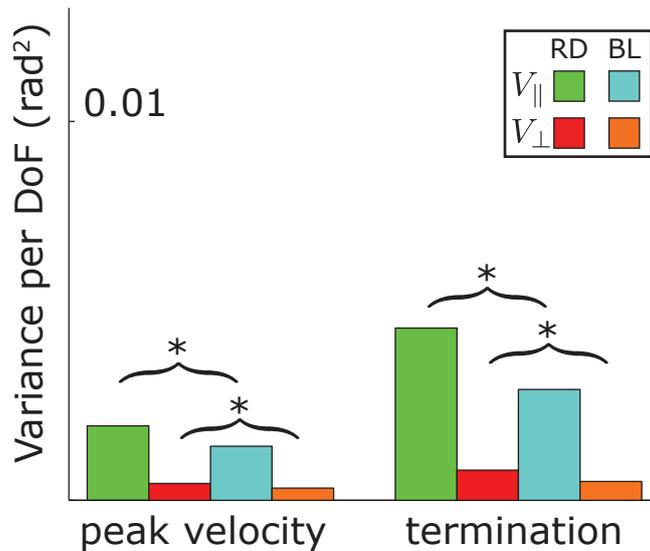


Figure 2.2.: Mean across subjects of the pointer tip variance per degree of freedom along the task manifold (V_{\parallel}) and orthogonal to it (V_{\perp}) at peak velocity (left) and movement termination (right). Significant differences reported by the ANOVA are indicated by *.

structure for the test population. The bootstrap method allows us to test the same hypotheses on the level of single subjects. We ran bootstrap tests for the *data comparison hypothesis* with the two test hypotheses $V_{\parallel}(\text{RD}) > V_{\parallel}(\text{BL})$ and $V_{\perp}(\text{RD}) > V_{\perp}(\text{BL})$ at the same time points of peak velocity and movement termination. Figure 2.3 shows these data for each of the eleven participants. The differences that were classified as significant by the bootstrap test are marked with an asterisk. The confidence level used in these tests was $\alpha = 0.05$.

2.4. Conclusions

We presented a mathematical formalism that represents movement tasks as manifolds in a vector space describing the configuration of the motor system. This enabled us to express intuitive notions about movements being task-related as formal statements. We used statistical methods to transform these geometric concepts of task representation into a scientific strategy to analyze experimental data. Several hypotheses about the stochastic properties of movements in relation to a given task were formalized. A parametric bootstrap method was described to determine the critical value for rejecting the corresponding null hypotheses.

The mathematical description of tasks we developed enabled us to test hypotheses about the principles of movement generation used by single human subjects. Traditionally, such hypotheses were tested on a level of a whole population of subjects

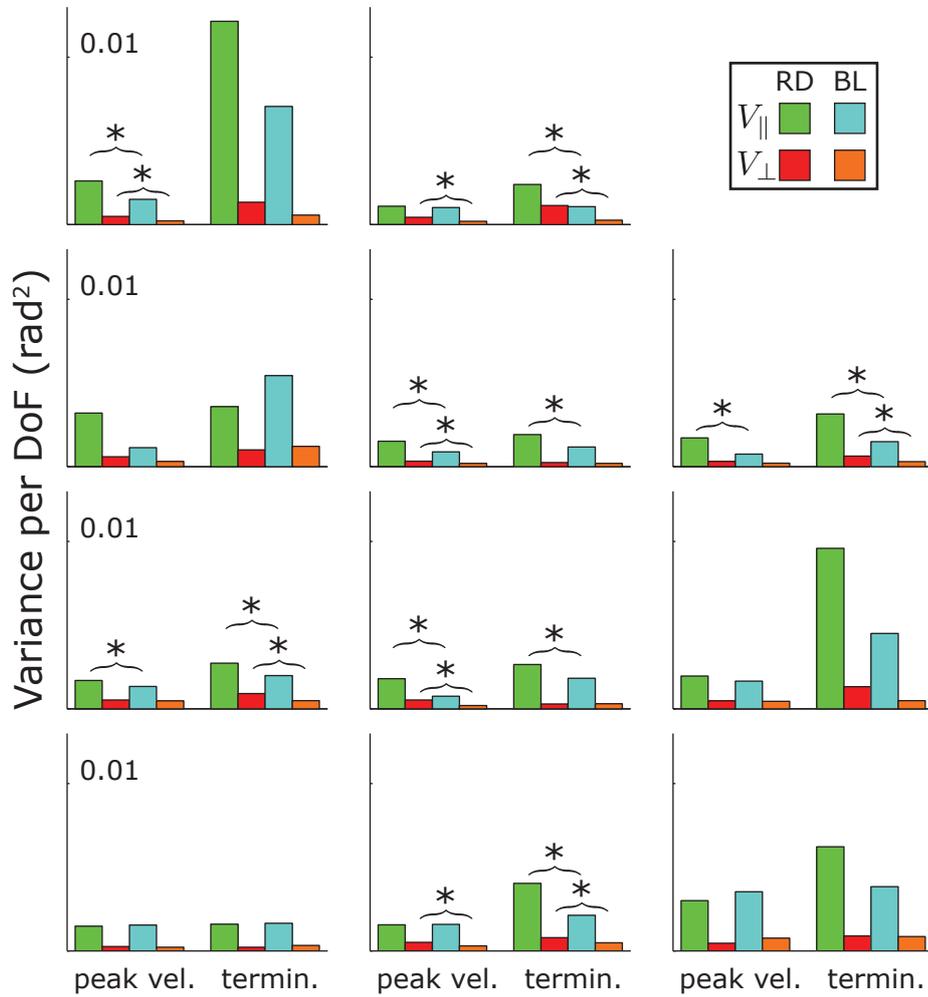


Figure 2.3.: Single subject pointer tip variance per degree of freedom along the task manifold (V_{\parallel}) and orthogonal to it (V_{\perp}) at peak velocity (left) and movement termination (right). Significant differences reported by the parametric bootstrap test are indicated by *.

Hypothesis	number of subjects for which null hypothesis was rejected	p -level reported by ANOVA
$V_{\parallel}(\text{RD}, t_1) > V_{\parallel}(\text{BL}, t_1)$	5	0.034
$V_{\perp}(\text{RD}, t_1) > V_{\perp}(\text{BL}, t_1)$	5	0.047
$V_{\parallel}(\text{RD}, t_2) > V_{\parallel}(\text{BL}, t_2)$	6	0.025
$V_{\perp}(\text{RD}, t_2) > V_{\perp}(\text{BL}, t_2)$	4	0.029

Table 2.2.: Comparison between ANOVA and parametric bootstrap test. Total number of subjects was $N = 11$.

using ANOVAs with the implicit assumption that the variance measures of the subjects conform to a normal distribution. Our formalization showed that under the more plausible assumption that the repeated task executions of single subjects are distributed normally, this assumption does not hold on the population level.

To examine the effects of this discrepancy, we analyzed the variance structures in a data set from a reaching study with 11 participants both with an ANOVA on the population level and with a parametric bootstrap analysis on the single-subject level. On the population level, the ANOVA resulted in significant differences between V_{\parallel} and V_{\perp} at two points in time. Analyzing the single subjects with the parametric bootstrap test reveals that each of these four differences is significant for at best only about half of the test subjects. Table 2.2 makes this comparison explicit. In general, one is interested in effects that are prevalent throughout a whole population of subjects sharing similar characteristics. This comparison implies that statements about population level effects based on ANOVA have to be treated with care and that analysis of single subjects with the parametric bootstrap method can provide more insight of how prevalent an effect is within a population.

The detailed analysis emphasizes the large differences between the movement patterns of single subjects. The fact that single subjects exhibit unique and highly repeatable movement patterns has been acknowledged informally within the motor control community. The lack of approaches to quantify these patterns in a meaningful, task-related way has inhibited formal research of this phenomenon. The mathematical apparatus presented here makes the quantitative comparison of unique movement patterns possible and should enable researchers to design experiments to better understand the origins of these patterns.

3. A process model of quiet upright stance

3.1. Introduction

So far we developed a formal language to express movement tasks as manifolds in the configuration space of the motor system and described a statistical method to test whether the central nervous system coordinates the available degrees of freedom in a way that is consistent with principles of selective stabilization relative to the specific task manifold. In the present chapter we will explore possible mechanisms of how this coordination might be realized by the CNS. We approach this problem by presenting a dynamic process model of the sensorimotor loop of postural stabilization in quiet, upright stance. We chose postural stabilization as an example because this allows us to focus on the redundancy of selecting a single solution to a movement task and factors out the more complex problem of generating a trajectory.

Dealing with the postural state of the whole body also involves additional problems. A human standing upright is inherently unstable. The only contact points are the feet on the ground, and most of the body mass is concentrated at the upper half of the body, far away from the ground. Small movements of the body center of mass (CoM) lead to comparatively large gravitational torques. Without any stabilizing forces from muscles to counter the pull of gravity, the body would collapse very quickly.

Mechanically, the human body consists of a number of rigid bodies, the bones, held together by muscles and ligaments that constrain the movement of the bones relative to each other. For all means and purposes, we will restrict our discussion to the sagittal plane. Traditionally, this complex system has been approximated by modeling it as an inverted pendulum with a single rotational joint at the ankle, standing upon the support surface of the foot soles. As long as the center of mass of the body as a whole is over this support surface, the feet remain standing on the ground. Any small deviation from the upright configuration results in a gravitational torque upon the ankle joint that increases that deviation. If the muscle torques counter the gravitational torques in some way, the whole system is stable. How is this stabilizing muscle activation structured? What is the role of the central nervous system (CNS) in modulating it? What sensory modes are used in detecting deviations from the stable state?

One possible answer to this group of questions is to assume that the CNS does not play an active role in the stabilization of posture at all, but completely delegates

the task to the reflex loops in the spinal cord. Winter and colleagues (Winter, Patla, Prince, Ishac, & Gielo-Perczak, 1998) proposed a “stiffness control model” for quiet stance: the CNS makes the ankle joint stiff by setting the muscle tone so high that for any deviation from upright stance, the stabilizing pull of the muscle being stretched is larger than the destabilizing effect of gravity, so that the system is stable on the whole. The stabilizing force from the muscles is ascribed to both the tonic stretch reflex and passive properties, but explicitly *not* to modulation from the CNS – aside from initially setting the muscle tone, the CNS does not participate at all in stabilization of upright stance in this model.

The stiffness control model has been criticized for a variety of reasons. As an immediate response, Morasso and Schieppati (1999) point out that while experimental measures of ankle stiffness during quiet stance are not conclusive, none of the values reported in the literature is sufficiently high to actually stabilize the body in quiet stance. This evaluation was supported by subsequent studies by other groups designed specifically to answer that question (see Section 3.3.3 for more details).

Another problem that the stiffness control model shares with the majority of models for quiet upright stance in the literature is that it approximates the biomechanics of the body by a single-joint inverted pendulum. This simplification rests upon the assumption that movement in other mechanical degrees of freedom along the body, e.g. knee and hip joints, is very small compared to movement in the ankle joint. This assumption is at odds with data from a recent study by Scholz and colleagues (Hsu et al., 2007) showing that during quiet stance there is *more* movement in proximal joints along the body than in the ankle joint, rather than less.

These two points of criticism were compounded in a theoretical study that explored the feasibility of stiffness control in a multi-joint model (Rozendaal & Van Soest, 2008). The authors showed that assuming pure stiffness control, the threshold for stability in a multi-segment inverted pendulum is even higher than in the traditional single-joint model, because displacements of the center of mass from additional joints add up. Furthermore, if stiffness is finite, the actual resting state for such a system is not in the upright configuration, but always includes some lean. Even for the maximally attainable stiffness values in full voluntary contraction of the leg muscles, this lean would be significant – though somewhat reduced if taking into account muscles that span multiple joints.

These experimental and theoretical findings indicate that the CNS is actively contributing to the stabilization of the body in quiet upright stance. But what exactly is this contribution? As has been discussed in Chapter 2, answering this question experimentally is not straightforward. The neural substrates responsible for this contribution are in the brain, which is hard to access directly in human subjects. Comparatively little is known about how neural activity in the brain relates to activity in the spinal cord related to motor control. Sensor signals play a role, but how exactly the activation patterns of the sensory surfaces are used to modulate brain activity, possibly after being transformed into neural representations of the body in space, is unknown.

These difficulties of directly analyzing in detail the role of the brain in stabilizing

quiet upright stance make it necessary to resort to indirect approaches. One way to do so is to probe the system by subjecting it to perturbations. These can target the sensory surfaces, e.g. modifying the visual input using projections (Kiemel, Oie, & Jeka, 2002; Kiemel, Zhang, & Jeka, 2011) or vibrating tendons to change proprioceptive signals (Goodwin, McCloskey, & Matthews, 1972; Polónyová & Hlavacka, 2001). Or they can target the state of the motor system directly, by displacing or rotating the support surface or applying a force at some other point on the body (Pai, Rogers, Patton, Cain, & Hanke, 1998; Kluzik, Horak, & Peterka, 2005). As the perturbations are targeted at the sensory surfaces, though, they allow conclusions regarding how the brain processes the sensory data and generates representations about the body state from them. Insights about how the motor areas transform these representations into descending motor commands are harder to infer.

Mechanical perturbations of the motor system have similar problems. Different parts of the complete system respond to such perturbations differently. There is a purely mechanical resistance to any perturbation from the inertia of the body and the passive elastic properties of the connective tissue. Low-level reflexes lead to fast muscle contractions opposing the muscle stretch after about 30-50 ms. If the perturbation is expected, pre-programmed reactions can be measured after ≈ 70 ms, while voluntary responses occur as late as 150 ms after the perturbation (Latash, 2008). At the time when voluntary or pre-programmed reactions are observed, the passive effects and reflex responses are still active, of course. The passive effects can be estimated and isolated by modeling the musculo-skeletal system, or bypassed by measuring muscle activation directly using electromyography.

Separating the effects of the perturbation on the higher brain areas from the reflex response is much harder and less accessible to isolating one from the other by either modeling or measuring the interface between them. The spinal reflex loops have been modeled on different levels of detail (Feldman, 1972; Raphael, Tsianos, & Loeb, 2010). One class of these models, the Equilibrium Point Hypothesis (see section 3.2.2 for details), has led to a variety of studies that measure kinematic or EMG data from behavioral experiments and use a model of the tonic stretch reflex to estimate the trajectories of the descending commands, with a wide range of results (Laboissière, Ostry, & Feldman, 1996; Gomi & Kawato, 1996; Ostry, Gribble, Levin, & Feldman, 1997; St-Onge, Adamovich, & Feldman, 1997; Gribble, Ostry, Sanguineti, & Laboissière, 1998; Micheau, Kron, & Bourassa, 2003; Pilon & Feldman, 2006). Many of these studies were part of a controversy of whether the particular model of the integration between the spinal reflexes and the descending motor commands assumed by the EPH is feasible at all (Kistemaker, Van Soest, & Bobbert, 2007). Far from being decided, this controversy underlines the limited role of the currently available models of the spinal reflex loops: they describe the functional aspects of the modeled neural circuitry well enough, but it is highly questionable whether the neurophysiological detail is sufficient to be used for the isolation of descending motor commands.

In the present study, we chose to approach the role of the CNS in the stabilization of quiet upright stance using a process model that describes the complete sensori-

motor loop. The parts of the model describing the estimation of the body in space from sensory data, the spinal reflexes, the muscle contraction dynamics, and the biomechanics are taken from the literature. The original contribution of our study is combining them into a single model and closing the loop by proposing a mechanism of how the brain generates descending motor commands from estimates of the body in space. We suggest a combination of different feedback terms based on the sensed movement of the body center of mass and head position in anterior-posterior direction, and the deviation of the head orientation around the media-lateral axis from the horizontal.

As explained above, no direct constraints are available for a model of the neural dynamics for movement generation within the brain. The model as a whole, on the other hand, is constrained by experimental data from behavioral studies. If the complete model describes the behavioral data well, we can infer that the model of the neural dynamics is an appropriate functional description of the role of the higher brain areas in the stabilization of quiet upright stance – if this part was wrong, while all other parts are right, the whole model could not be correct. As all other parts of the model are constrained directly and locally by data from the literature, the brain dynamics model is thus indirectly constrained by behavioral data.

What kind of behavioral data is appropriate to constrain a model such as this? Stabilization of a mechanically unstable system means reacting to perturbations that drive the system away from the stable state. In quiet stance, all of these perturbations are quasi-random. Specific movement trajectories of postural sway during quiet stance are essentially a random walk that is kept within the bounds of stability by the CNS. As it is impossible to estimate the random perturbations throughout that the CNS reacted to, constraining the model by fitting specific movement trajectories is not applicable.

While we cannot measure or reconstruct the random influences in single movement trajectories, the general characteristics of the perturbations can be captured statistically. We have chosen to constrain our model by three different measures of the variability of quiet stance trajectories. The joint excursion variance (JEV) measures the general magnitude of the postural sway. The geometrical structure of the sway trajectories is revealed by an uncontrolled manifold analysis (UCM, see Chapter 2 for details). The temporal structure of the sway is captured by the Power Spectral Density (PSD) of the trajectories.

3.1.1. Related work

Stabilizing the body in quiet, upright stance can be divided into the two sub-problems of estimating the deviations from the stable state based on the sensory signals and generating appropriate motor commands that counter these deviations. Most modeling work has focused on one of these points.

Models about the sensory estimation aspect of upright stance describe how the different sensory surfaces related to the body posture are integrated into a single estimate of the current deviation from the stable state. The methods used to model the

integration process vary. A comparatively simple weighted addition of independent channels to a common torque output (Peterka, 2002; Maurer, Mergner, & Peterka, 2006) is similar to what we will use in our model. More sophisticated models use Kalman filters with parameters that are adapted to fit experimental data (Kiemel et al., 2002; Kuo, 2005). All these models use torque as the motor output. Some take into account the additional passive torques generated by the elastic and viscous properties of the tissue (Peterka, 2002), but do not account for the dynamics of the spinal reflex loops. One could argue that these are models of motor command generation in the CNS and the spine is part of the CNS, so they include these aspects implicitly. Still, modeling the motor command by a simple torque feedback might ignore the stabilizing properties of modularity between the brain and the spinal reflex loops.

On the motor side, models are concerned with identifying the control strategies used by the CNS to transform sensory information into motor commands. There is some agreement that purely feedforward controllers like the stiffness control model proposed by Winter et al. (1998) are neither feasible nor necessary (Morasso & Sanguineti, 2002; Lakie, Caplan, & Loram, 2003; Masani, Vette, & Popovic, 2006). Even taking into account considerable time delays, feedback control models have successfully reproduced many aspects of the sway trajectories of one-dimensional variables like the center of pressure, center of mass or ankle joint angle (Qu, Nussbaum, & Madigan, 2007).

The origin of other phenomena are still debated, like the mechanisms responsible for the two distinct components of sway: Trajectories of sway-related variables like the center of pressure have a high-frequency component that oscillates at ≈ 2 Hz. The center of this oscillation is not fixed, but oscillates itself with a lower frequency of ≈ 0.2 Hz (Zatsiorsky & Duarte, 1999). Using systems analysis approaches, Kiemel, Oie, and Jeka (2006) have shown that the slow oscillation is part of a neural feedback loop instead of a feedforward component. Zatsiorsky and Duarte (1999) have named the reference of the fast oscillation the “instant equilibrium point” and speculated that the fast component is related to spinal stretch reflexes, while the slow component represents changes of the physiological threshold of these reflexes. They interpret this as support for the Equilibrium Point Hypothesis of motor control, postulating that the fast oscillations are the result of the stretch reflexes stabilizing the system around the point of equilibrium where tonic muscle activation and gravitational forces are equal, while the slow oscillations represent shifts in the underlying muscle tone. A study by Micheau et al. (2003) made this line of reasoning explicit, modeling the tonic stretch reflex based on the EPH and inverting it to calculate the trajectories of the muscle tone threshold parameters from experimentally observed joint angle trajectories. That study was purely descriptive, though, and did not propose any mechanism for how these changes in the threshold parameters might be generated based on sensory data.

Another ongoing debate is about whether the neural dynamics that stabilize the body in upright stance are active continuously or only intermittently (Loram, Gollee, Lakie, & Gawthrop, 2011). The feedback models inspired by control theory were traditionally using continuous feedback (Mergner, 2010). This is at odds with observed

patterns of muscle activation, as pointed out by Lakie and Loram (2006): in quiet upright stance, muscle activation is not continuously changing but comes in short, temporally well separated bursts. A model trying to replicate these patterns of intermittent control based on suppressing feedback when the estimations of relevant variables are within a region of uncertainty due to sensory noise was proposed by Bottaro, Yasutake, Nomura, Casadio, and Morasso (2008). In a broad comparison study, van der Kooij and de Vlugt (2007) try to replicate data from perturbation experiments with different models using either continuous or intermittent feedback. The continuous models provide a better fit for the experimental data, though the authors report some remnants of low-frequency body sway that could not be explained by the purely continuous models and might be due to intermittent effects. A purely theoretical analysis by Asai et al. (2009) compares a continuous PD-controller with an intermittent dynamic controller that switches between two response functions when the estimated state exits a “dead zone” around the nominal equilibrium state. They show theoretically that the stability region in parameter space is much larger for the intermittent controller and demonstrate that it is capable of resisting larger perturbations in a simulation study. In another theoretical study, Gawthrop, Loram, Lakie, and Gollee (2011) show that the movement pattern generated by an intermittent controller with small activation thresholds cannot be distinguished from those generated by a purely continuous controller. According to the authors, this explains why data from some experimental conditions support the continuous control hypothesis even if the general control strategy is intermittent.

Most of the modeling work quoted so far describes the upright body as an inverted pendulum with a single rotational joint at the ankle. This is an appropriate simplification in general. The existence of stabilization strategies utilizing a combination of hip and ankle joints has long been acknowledged (Nashner & McCollum, 1985). It was generally assumed that this “hip strategy” is predominantly used to cope with sudden perturbations, while in unperturbed quiet stance the control patterns accord to the “ankle strategy”, where active, central control feeds into the ankle joint and motions in the hip joint are minimal (Horak & Nashner, 1986). That view has been challenged by studies using a wider range of perturbation magnitudes. The results suggest that the ankle strategy and hip strategy are two extremes of a continuous range of postural responses rather than discrete choices (Runge, Shupert, Horak, & Zajac, 1999; Creath, Kiemel, Horak, Peterka, & Jeka, 2005). Investigating the relationship between these two patterns and modulation of sensory information, Zhang, Kiemel, and Jeka (2007) report that addition or removal of sensory modes predominantly affects the sway component related to the ankle strategy. Subsequent work by the same group indicates that the CNS applies a single control strategy to activate all muscles related to postural control (Kiemel et al., 2011) and that the two distinct sway patterns are largely a result of the interplay between the inertial properties of the body segments and the experimental conditions (Alexandrov, Frolov, & Massion, 2001). Kuo (2005) has shown that a proportional-derivative feedback controller is capable of stabilizing a two-link inverted pendulum and can reproduce the geometrical properties of the sway trajectories in the space spanned by the two joint angles.

An experimental analysis of these patterns in six degrees of freedom using the UCM method was carried out by Hsu et al. (2007), revealing a high degree of coordination between the joint angles that selectively stabilizes important task variables.

3.2. The model

Stabilizing the human body in upright stance is a closed action-perception loop. Sensor systems estimate the configuration of the body in space and how it changes. The brain generates adequate motor signals that counter the sensed deviations and sends them along descending neural pathways to the motor-neurons in the spine. The spinal reflex loops transform the descending signals into α -motorneuron activation, factoring in the current state of the muscles and joints. The motorneuron activation is transformed into muscle force by biophysical processes that make the muscles contract or relax. The muscle force changes the configuration of the body in space, closing the loop. An overview of the whole loop is given in Figure 3.1.

In the following sections, we will present each of these phases in detail. At different points, different influences act as random perturbations on the system state, so we begin with a treatment of how these are modeled (3.2.1). Because the loop is closed and each phase depends upon the previous one, there is no natural place to begin describing it. As the original contribution of this work modeled the neural dynamics between sensory information and spinal reflex loops, we have chosen to organize the presentation in a way that starts with the spinal dynamics (3.2.2) and then follows the loop to muscle contraction (3.2.3), the biomechanical equations of motion (3.2.4) and forming representations of the body in space from sensory data (3.2.5). The loop is closed by describing our proposed model for the neural dynamics transforming the represented body state into descending motor commands 3.2.6.

3.2.1. System noise

The elementary information processing unit in the nervous system is a single neuron. Information is transmitted between neurons by spikes, electrical action potentials traveling along axons and across synapses between neurons. The activation level of a neuron can be captured by the firing rate, measured in spikes per second.

An action potential is a probabilistic event, and a model of neural activation processing must take this randomness into account. There are different methods of how to do that appropriately, though. According to Smith and Ratcliff (2004), any model trying to explain behavioral data on a neural basis needs to account for noise on three levels of analysis: the individual neurons, populations of neurons, and behavioral data. Accounting for the noise of behavioral data is what constrains our model, so the model as a whole deals with this level. For a single neuron, the spiking probability depends largely upon the membrane potential, and the spike trains can be modeled by a Poisson process (Shadlen & Newsome, 1994). As we model neural activity on a relatively abstract, functional level that does not capture single neurons,

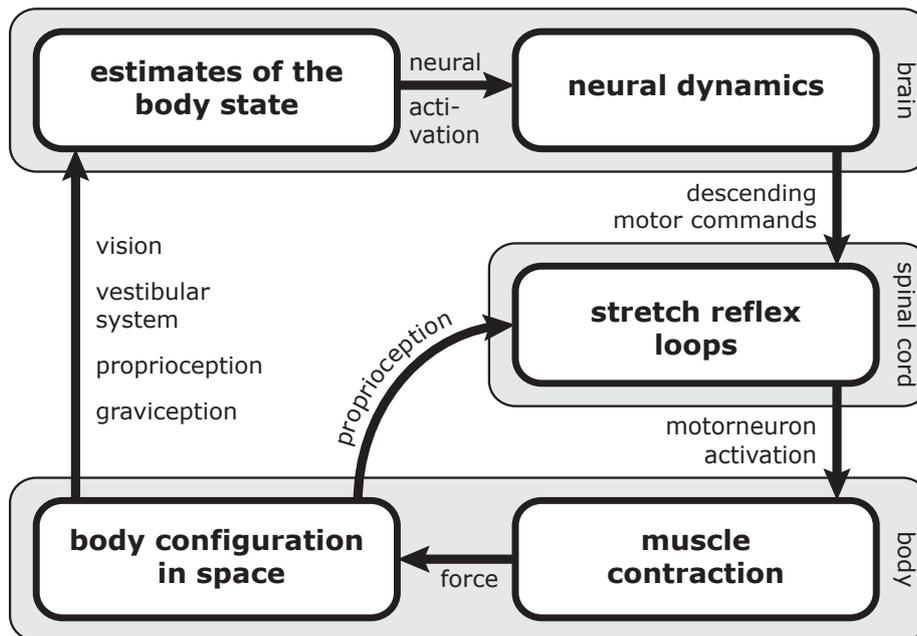


Figure 3.1.: Overview of the complete sensorimotor loop for balancing the body in quiet, upright stance.

we can disregard this level of detail and focus on randomness at the population level.

Firing rates of individual neurons are often weakly correlated ($r \approx 0.15 - 0.2$) when they are physically close to each other in the cortex (Zohary, Shadlen, & Newsome, 1994; Bair, Zohary, & Newsome, 2001). One cannot disregard this noise by assuming that the fluctuations in individual neuron firing rate cancel out over whole populations. In an in-vivo study of the cat visual cortex, Arieli and colleagues compare real-time optical imaging data to local field potentials and single neuron firing rate (Tsodyks, Kenet, Grinvald, & Arieli, 1999). They were able to predict the seemingly random response in single trials from the fluctuations in neural firing rates and the deterministic response, and conclude that the neural activity is an important source of the variability. A later study by the same group provides further evidence of the correlation between single neuron firing rate and population activity, both in spontaneous and stimulus-driven cases.

The spatio-temporally correlated variability of neural population activation can be modeled by an Ornstein-Uhlenbeck process (Smith, 2010; Ricciardi & Sacerdote, 1979; Lánský & Sacerdote, 2001). This is expressed mathematically by the solution of the stochastic differential equation

$$\dot{\eta} = -\alpha_{\eta}\eta + \xi, \quad (3.1)$$

where η is the colored noise, α_{η} the relaxation time parameter and ξ is Gaussian white noise with zero mean and parameterized by the variance of the integrated noise after

one second

$$\text{var} \left(\int_0^1 \xi dt \right) = \sigma^2. \quad (3.2)$$

We will use random signals $\eta(t)$ as noise for several different neural activation variables throughout our model. It is usually an acceptable simplification to use additive noise for neural activation variables (Faisal, Selen, & Wolpert, 2008). For motor-neural activation, though, the noise level has been empirically shown to depend upon the mean activation level (Matthews, 1996). A behavioral consequence of this is the speed-accuracy trade-off known as Fitt's law, which says that faster movements are less accurate – because they require larger muscle forces which are noisier (Fitts, 1954). Harris and Wolpert (1998) postulated that this signal-dependency of motor noise is the deciding factor in the shape of goal-directed movement trajectories and show that the trajectories of both saccadic eye movements and upper extremity reaching movements can be predicted by assuming that the CNS attempts to minimize the end-point variance of the movement. We follow this line of reasoning by using multiplicative noise for the activation level of motor neurons.

3.2.2. Neural dynamics in the spinal cord

The spinal cord is the interface between the descending neural commands and the actual motor output in the form of muscle activations. While much is known about the connectivity of spinal cord neurons and the feedback loops, many functional aspects of how these low-level neural dynamics interact with high-level commands and how meaningful movement is generated is still debated. For the model of upright stance presented here, we chose to follow the Equilibrium Point Hypothesis (EPH).

The Equilibrium Point Hypothesis

One well-understood function of the spinal cord is the tonic stretch reflex (Latash, 2008). This is a monosynaptic reflex loop in which Ia-afferents from the muscle spindle excite α -motorneurons of the same muscle. The Ia-afferents are sensory fibers sensitive to changes in muscle length. When these sensory afferents signal an increase in muscle length, they excite the α -motorneuron, which leads to contraction of muscle fibers, generating force to counteract the muscle stretch. The amount of force generated depends upon the amount of stretch via the activity of the Ia-afferents.

This dependency has been measured experimentally in some detail (Feldman, 1972; Ostry et al., 1997; Adamovich, Levin, & Feldman, 1997). It can be described as an exponential force-length relationship:

$$F = \left[e^{\alpha(L-L^0)} \right]^+ - 1, \quad (3.3)$$

where L is the muscle length, L^0 a threshold length below which no force is generated, α a physiological form parameter and F the generated force. The *shape* of the curve

describing this relationship depends upon α , it is more or less constant for each muscle and is sometimes referred to as the *invariant characteristic* of the muscle. The *position* of the curve along the length-axis, on the other hand, is variable and depends upon the threshold parameter L^0 . (Latash, 2008)

The Equilibrium Point Hypothesis states that this reflex loop is essential for both stabilizing posture and generating movement. Consider a single joint actuated by a pair of agonist and antagonist muscles, e.g. the elbow with biceps and triceps brachii, flexed at 90° . The state of this system can be described by the joint angle θ as a single parameter, as muscle length is uniquely determined by the joint angle and muscle force uniquely determines the joint torque T

$$T = T_{AG} - T_{AN} \quad (3.4)$$

$$T_{AG} = \left[e^{\alpha_{AG}(\theta - \theta_{AG}^0)} \right]^+ - 1, \quad T_{AN} = \left[e^{-\alpha_{AN}(\theta - \theta_{AN}^0)} \right]^+ - 1. \quad (3.5)$$

A perturbation in the form of a small elbow extension generates a resisting torque from activation of the stretched muscles via the reflex loop, while a perturbation in the opposite direction is resisted by other muscles spanning the same joint. We will refer to such opposing muscle groups as agonist/antagonist pairs. The combination of the responses from these pairs leads to a negative feedback loop with a single equilibrium point.

This system has one stable equilibrium point. Its position depends upon the threshold parameters θ_{AG}^0 and θ_{AN}^0 of the agonist and antagonist acting upon the joint. For $\alpha_{AG} = \alpha_{AN}$ and in the absence of external torques, the equilibrium point is given by

$$\lambda = \frac{\theta_{AG}^0 + \theta_{AN}^0}{2}. \quad (3.6)$$

Fixing the equilibrium point λ leaves one free variable, which is defined as

$$\rho = \frac{\theta_{AG}^0 - \theta_{AN}^0}{2}, \quad (3.7)$$

describing the co-contraction of the muscle pair. High values of ρ imply high forces of the individual muscles at the equilibrium point, while the resulting joint torque is still zero. This allows us to reformulate the motorneuron activation in terms of λ and ρ .

$$E_{AG} = e^{[\alpha_E(\theta - \lambda + \rho)]^+} - 1, \quad E_{AN} = e^{[-\alpha_E(\theta - \lambda - \rho)]^+} - 1. \quad (3.8)$$

So far, this model of α -motorneuron activation takes only position information into account, i.e. muscle length or joint angles. The α -motorneurons receive afferent information from both type Ia and type II sensory fibers, though. While the latter are sensitive to the length of a the static muscle, the former primarily fire when the muscle length changes, i.e. they encode velocity information. De Lussanet and colleagues have shown that it is beneficial to model this velocity dependence as being relative to the rate of change of the threshold, i.e. depending upon $(\dot{\theta} - \dot{\lambda})$ instead of

just on the absolute velocity $\dot{\theta}$ (de Lussanet, Smeets, & Brenner, 2002). We add the velocity term and make the activation dependent upon the proprioceptive signals of the state variables $\hat{\theta}$ for θ and $\hat{\dot{\theta}}$ for $\dot{\theta}$ (see below, Section 3.2.5) to get

$$\begin{aligned} E_{AG} &= e^{[\alpha_E(\hat{\theta}-\lambda+\rho+\mu(\hat{\dot{\theta}}-\lambda))]^+} - 1, \\ E_{AN} &= e^{[-\alpha_E(\hat{\theta}-\lambda-\rho+\mu(\hat{\dot{\theta}}-\lambda))]^+} - 1. \end{aligned} \quad (3.9)$$

Adding up the activation for agonist and antagonist motorneurons with the appropriate sign and including signal-dependent noise (see Section 3.2.1) gives an expression for the idealized total motorneuron activation

$$E = (-E_{AG} + E_{AN}) \eta_m \in \mathbb{R}^3, \quad (3.10)$$

which is proportional to the total torque generated from active muscle contraction in both agonists and antagonists.

3.2.3. Torque generation in muscle-tendon systems

Activation of motor neurons is transformed into force by muscle contraction. Muscles are connected to different bones in the skeleton via elastic tendons that can store and release energy. The muscle itself has elastic properties as well (Van Soest & Bobbert, 1993; Brown, Scott, & Loeb, 1996). We model the muscle-tendon complex as a contractile element in parallel with a viscoelastic element.

The physical characteristics and dynamics of muscle force generation have been modeled on various levels of detail. Many optimal control models of motor control assume that the central nervous system can directly generate joint torques or even accelerations (Peterka, 2000; Todorov & Jordan, 2002; Kiemel et al., 2002). This is physiologically implausible, due to the spinal reflex loops described in the previous section. The equilibrium point hypothesis does take these spinal circuits seriously, assuming that the control variables available to descending commands are the thresholds of the stretch reflex (see Section 3.2.2). Other researchers have modeled the spinal reflex loop in even more detail (Mileusnic, Brown, Lan, & Loeb, 2006; Raphael et al., 2010), but the level of complexity of these models is beyond the scope of the present study.

The λ -model specifies how the α -motorneuron activation depends upon the proprioceptive signals encoded by the activity of the sensory afferents – the E in Equation 3.10 signifies a correspondence to electromyographic data (EMG). The relationship between the motorneuron activation and the actual force or torque generated by the muscle is not trivial, though (Kandel, Schwartz, Jessell, Siegelbaum, & Hudspeth, 2012). The activity of the α -motorneuron causes calcium to be released in the muscle fiber. The calcium facilitates the sliding of actin against myosin layers within the fibers. The calcium is transported back out of the muscle fiber during this process. The sliding process continues as long as calcium is available, i.e. the length of

the muscle contraction depends on the amount of calcium initially released, which depends upon the activity of the α -motorneurons. This transformation of neural activity into force takes time. It is usually modeled as a second order low-pass filter. There are different versions available in the literature, for the present model we choose to follow the model of Gribble and colleagues (Gribble et al., 1998):

$$\tilde{T}_{\text{act}} = AE \in \mathbb{R}^3, \quad (3.11)$$

$$\tau_m^2 \ddot{T}_{\text{act}} + 2\tau_m \dot{T}_{\text{act}} + T_{\text{act}} = \tilde{T}_{\text{act}} \in \mathbb{R}^3, \quad (3.12)$$

where \tilde{T}_{act} is the steady state torque, τ_m a time constant and T_{act} the instantaneous torque generated by the active contraction process. The physiological parameter A describes the relationship between motorneuron activity or EMG and generated torque.

In addition to the active feedback loops that counteract stretch by the reflex loop described above, muscles also have viscoelastic properties that resist stretch passively. In addition to that, the tendon is a purely viscoelastic element that cannot be actively modulated by the nervous system. Joint torques generated by these passive elastic properties of the muscles and tendons along the leg have been measured experimentally by Riener and Edrich (1999). These researchers found that the passive torques can be described well by a double exponential curve of the general form

$$T_{\text{ela},j} = \exp(a_{j0} + \sum_{i=1}^3 a_{ji}\theta_i) - \exp(b_{j0} + \sum_{i=1}^3 b_{ji}\theta_i) + c_{ji}, \quad (3.13)$$

where j indicates the joint. For the knee joint, an additional exponential term accounts for the steep increase in torque when the knee is fully extended (see Section 3.5.1 for further discussion). We adopt this formulation for our model.

The passive viscous properties of muscles and joints are difficult to determine for lack of experimental data. The viscosity of muscles and tendons has been modeled by a linear (Hatsopoulos, 1994; Flash, 1987; Hogan, 1984) or nonlinear (Barto, Fagg, Sitkoff, & Houk, 1999; Gribble et al., 1998; Loeb, Brown, & Cheng, 1999; Tee, Burdet, Chew, & Milner, 2004) damping element. As the range of movement in quiet stance is not large enough for the non-linearity to be significant, we chose to model viscosity by a linear term

$$T_{\text{vis}} = -B\dot{\theta} \in \mathbb{R}^3. \quad (3.14)$$

The total force generated by the muscle-tendon complex is given by the sum of the active torques generated by muscle contraction and the elastic and viscous passive torques

$$T = T_{\text{act}} + T_{\text{ela}} + T_{\text{vis}}, \quad (3.15)$$

which corresponds to a contractile element in parallel with a viscoelastic element.

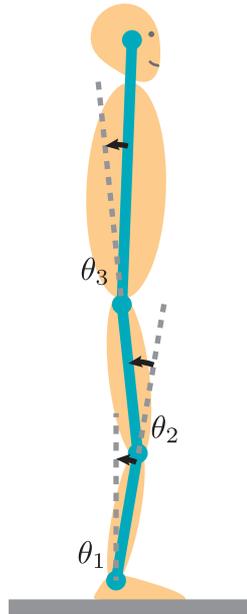


Figure 3.2.: Sketch of the kinematic model of the body in upright stance as an inverted pendulum with three segments. The body configuration is described by the three angles θ_i of the ankle, knee and hip joints.

3.2.4. Biomechanics of the skeleton

The configuration of the body in space can be described by specifying the configuration of each bone as a rigid body and neglecting the shifts of muscle, skin and other tissues around the bones. The bones do not move freely, though, but are connected to each other as joints, by structures of muscles, tendons, ligaments and cartilage.

The exact structure of the connection between two bones varies from joint to joint. A common feature is that parts of two or more bones moving against each other without losing the surface connection. The two surfaces are usually rounded, implying that if described in an appropriate coordinate frame, the motion is mostly rotational and the translational components can be neglected. The relative configuration of the femur (thigh) and the tibia (shank), e.g., is specified by defining a single axis of rotation somewhere between the lateral and medial condyles of the femur and an angle of rotation around that axis relative to some arbitrary reference configuration.

The configuration of the whole body in the sagittal plane can be described by a number of segments connected by rotational joints. The lower body consists of the leg segments foot, shank and thigh, connected by the ankle joint between foot and shank and the knee joint between shank and thigh. The upper body is more complicated: the spine consists of 24 articulate vertebrae and the fused vertebrae of the sacrum. The head is attached to the spine by the atlanto-occipital joint, describing movement between the uppermost vertebra of the cervical spine (atlas, or C1) and the base of the skull (occipital bone). We have chosen to mostly neglect this complexity in the current model: The movement of two single vertebrae against each

other is quite restricted. Furthermore, adding 24 degrees of freedom would make the model computationally infeasible. Instead, we model the trunk and head as a single rigid body, linked to the lower body at the hip joint. The body configuration is thus described by a set of three generalized coordinates θ_i : the angles of the ankle (θ_1), knee (θ_2) and hip joints (θ_3), as illustrated in Figure 3.2.

The equations describing the kinematics and dynamics of the body can be derived explicitly using basic trigonometry and mechanics. For the three degrees of freedom model, the equations are provided in Appendix C. For more degrees of freedom, the equations become too long to be derived explicitly.

While the equation for the position of the head in anterior-posterior direction p is comparatively simple, the equations for the dynamic terms are more complex and depend upon large numbers of parameters, most notably the locations of the joint axes and segment centers of mass and the weight and moments of inertia of the segments. These parameters were derived for an ideal male subject of 1.8 m height and 80 kg weight using the methods specified by Winter (Winter, 1990). The values of the biomechanical parameters are listed in Appendix B.

The configuration of the body in space changes according to torques acting on the joints, both from muscle-tendon complexes and gravity. The equations of motion relating the torques to accelerations are given by

$$M(\theta)\ddot{\theta} + C(\theta, \dot{\theta})\dot{\theta} + N(\theta) = T, \quad (3.16)$$

where M is the inertia matrix, C a matrix representing Coriolis and centrifugal forces, N the vector of gravitational forces and T the vector of torques generated by the muscle-tendon complexes defined in Equation 3.15.

3.2.5. Sensor data and neural representations

Having modeled how the state of the body in space changes depending upon internal and external forces, we now move on along the sensorimotor loop to how the CNS senses these changes of body configuration and derives estimates of variables that are important for body stability from those sensations. For upright stance, the most important sensory surfaces are the eyes, the vestibular system, proprioceptive muscle spindles and pressure sensors in the soles of the feet. Among these sensor modes, the role of proprioception is unique, because besides contributing to the formation of central estimates of the body in space, the activation of the proprioceptive muscle spindles also play a vital role for the stretch reflexes in the spinal cord. We will first treat this special role of proprioception in the activation of α -motorneurons, then move on to deal with the formation of central estimates of the body in space from fusing multiple different sensory channels, including perception.

Proprioception is of paramount importance for postural stability. Nevertheless, the term is loosely defined and refers to a collection of several different sensory surfaces (Taylor, 2009). There are three different types of afferent: Ia, Ib and II. Type Ia and II afferents terminate in the muscle fibers. Type Ia afferents mostly fire when the

muscle length changes and can be said to encode the velocity of the muscle stretch. Type II afferents mostly fire when the muscle remains still, they encode the muscle length itself and fire in static muscles.¹ Type Ib afferents terminate in Golgi tendon organs. The firing rate in the Ib afferents depends upon the stretch of the tendon, which is directly determined by the total force currently generated by the muscle, so Ib afferents can be said to encode force.

Physiologically, the neural activation in these pathways depends upon muscle length, rate of change, and force. The important aspect of information about the body in space is not so much the length of the skeletal muscles, though, but rather the configuration of the joints. Muscles are connected to joints via elastic tendons (Van Soest, Bobbert, & van Ingen Schenau, 1994). The relationship between the joint angle and the muscle length is not unique, but depends upon the tendon length. Kistemaker and colleagues have shown recently that it is feasible to assume that the tendon length is known to the proprioceptive system: the relationship between tendon length and muscle force is unique, and muscle force is coded by activity of the Ib afferents (Kistemaker, Van Soest, Wong, Kurtzer, & Gribble, 2013). It is thus reasonable to represent the processes and activation variables described above as estimates of the joint angles $\hat{\theta}$ and their rate of change $\dot{\hat{\theta}}$ which we used to describe the stretch reflex in Equation 3.9.

The other sensory channels included in our model are much further removed from neurophysiological data. Loosely speaking, vision and the vestibular system provide information about the kinematic state of the body in space, while the pressure sensors in the sole of the feet provide kinetic information (Peterka, 2002). We assume that the brain processes the information into single estimates for kinematic states of three variables that describe aspects of the body in space: the center of mass position in anterior-posterior direction, the head position in anterior-posterior direction, and the orientation of the head around the media-lateral axis. This process of combining different, possibly conflicting sensory estimates for the same or related values into a single representation is known as *multisensory fusion* or *multisensory integration* (van der Kooij, Jacobs, Koopman, & Grootenboer, 1999; Maurer et al., 2006). Models of multisensory fusion often adopt Kalman filters to optimally combine the different sensory modes into a single estimate by adapting weights based on the reliability of each channel (van der Kooij et al., 1999; Kiemel et al., 2002). Instead of modeling the multisensory fusion process explicitly, we assume that it results in an unbiased but delayed estimate of the variable in question.

All these signaling processes are not instantaneous. The creation of action potentials in the muscle spindles and Golgi tendon organs and the transmission along the axons and the synapse to the α -motorneuron take time. The neural processes of fusing the different sensory modes to a unified estimate of important variables takes even more time. Again, we cannot account for these processes in detail, but capture them by adding a fixed time delay to the sensory estimates, as is common practice in

¹ This difference is one of preferred firing rather than a complete division of labor: both type Ia and II are active in both static and changing muscles.

functional models on this level (van der Kooij et al., 1999; Maurer & Peterka, 2005; Asai et al., 2009).

After these considerations, we can state the sensory estimates of our model as

$$\begin{aligned}
\widehat{\theta}(t) &= \theta(t - d_\theta) + \eta_\theta, & \widehat{\dot{\theta}}(t) &= \dot{\theta}(t - d_\theta) + \eta_{\dot{\theta}}, \\
\widehat{p}(t) &= p(t - d_p) + \eta_{\dot{p}}, & \widehat{\ddot{p}}(t) &= \ddot{p}(t - d_p) + \eta_{\ddot{p}}, \\
\widehat{c}(t) &= c(t - d_c) + \eta_{\dot{c}}, & \widehat{\ddot{c}}(t) &= \ddot{c}(t - d_c) + \eta_{\ddot{c}}, \\
&& \widehat{o}(t) &= o(t - d_o) + \eta_o,
\end{aligned} \tag{3.17}$$

where $\theta \in \mathbb{R}^3$ is the vector of joint angles, $p, c \in \mathbb{R}$ the head and CoM positions in the anterior-posterior direction, and o is the head orientation around the media-lateral axis. The η_* are random processes as described in Section 3.2.1.

Different sensory modes that estimate the same functional variable have different levels of accuracy (Fitzpatrick & McCloskey, 1994). Adding or blocking a sensory channel during quiet stance has a significant effect upon the magnitude of the postural sway. A well known phenomenon is that the postural sway increases when subjects close their eyes (Nashner, Black, & Wall, 1982; Kiemel et al., 2002; Krishnamoorthy, Yang, & Scholz, 2005; Hsu et al., 2007). If, on the other hand, an additional sensor mode is provided by lightly touching a fixed reference object with a finger, sway is reduced (Zhang et al., 2007; Wing, Johannsen, & Endo, 2011). We model the loss of reliability when closing the eyes by an increase in the noise for estimates of variables where vision plays a role, i.e. the position and orientation of the head.

3.2.6. Neural dynamics in the brain

How does the central nervous system utilize the available sensor information to generate descending commands that generate appropriate muscle activations to stabilize the body in space? While all other parts of the motor loop described in the preceding sections are constrained by anatomical or physiological data to some degree, these constraints apply much less to the dynamics of the brain areas involved in movement generation.

The interfaces for these neural dynamics are the sensory signals \widehat{c} , $\widehat{\dot{c}}$, \widehat{p} , $\widehat{\dot{p}}$ and \widehat{o} , on one side and the descending motor commands $\dot{\lambda}$ on the other side. The neural dynamics then formalize as any function

$$\dot{\lambda}(\widehat{c}, \widehat{\dot{c}}, \widehat{p}, \widehat{\dot{p}}, \widehat{o}) \in \mathbb{R}^n. \tag{3.18}$$

The role of the brain dynamics is to detect deviations from the stable state of upright stance and counter them with appropriate motor commands. The stable state is usually defined as any state where the center of mass position in anterior-posterior direction is within the support surface, i.e. between the toes and the heel of the feet when standing on normal ground. This is a region in the three-dimensional state space defined by the joint angles. Regardless of where exactly within that region

the system is at any given point in time, if it moves in any direction for a substantial length of time, it will leave the stability region. One approach to stabilizing the system is thus to detect any movement and generate a descending command that reduces that movement to zero.

Movement is indicated by the estimates of the velocity and acceleration of the head and center of mass. Even if a velocity is zero, a non-zero acceleration results in a non-zero velocity over time. For any variable v , a negative feedback law of the form

$$\ddot{v} = -\alpha_1 v - \alpha_2 \dot{v} \quad (3.19)$$

with positive gain parameters α_i is a system with a single asymptotically stable fixed point (Perko, 1991). We can translate that into feedback terms

$$f_c = -\alpha_{\hat{c}} \hat{c} - \alpha_{\dot{c}} \dot{\hat{c}}, \quad f_p = -\alpha_{\hat{p}} \hat{p} - \alpha_{\dot{p}} \dot{\hat{p}} \quad (3.20)$$

for the variables head velocity and center of mass velocity. If the brain could generate descending signals resulting in

$$\ddot{c} = f_c, \quad \ddot{p} = f_p, \quad (3.21)$$

i.e. the equation 3.19 holding, then this feedback law would reduce any movement of the center of mass and head to zero.

Realizing this feedback directly as specified by equation 3.21 is not possible, though. The available descending commands are modulations $\dot{\lambda}$ of the activation thresholds of the spinal reflex loops. The nervous system has to transform the desired feedback for center of mass and head movement states into appropriate shifts in these thresholds λ . In the following paragraphs, we propose one solution to this transformation problem for the center of mass.

The first step is to account for the kinematic structure of the body. The effect of joint motion on center of mass displacement is different for each joint. The relationship between these is captured in the Jacobian matrix $J_c = \frac{dc}{d\theta}$. The relationship between the movement states is given by

$$\begin{aligned} \dot{c} &= J_c \dot{\theta}, \\ \ddot{c} &= J_c \ddot{\theta} + \dot{J}_c \dot{\theta}, \\ \ddot{c} &= J_c \ddot{\theta} + 2\dot{J}_c \dot{\theta} + \ddot{J}_c \theta. \end{aligned} \quad (3.22)$$

In the regime of quiet stance, though, the changes in the Jacobian matrix are very small, so it is safe to assume that

$$\dot{J}_c = \ddot{J}_c = 0, \quad \text{implying} \quad (3.23)$$

$$\ddot{c} = J_c \ddot{\theta}. \quad (3.24)$$

The descending motor commands $\dot{\lambda}$ affect θ directly and c only indirectly by its

dependency on θ , so we have to solve this equation for $\ddot{\theta}$. As the Jacobian J_c is not square, we cannot invert it. We can use a right inverse to get the relationship we want, though. We choose the Moore-Penrose pseudo-inverse (Siciliano & Khatib, 2008), given by

$$J_c^+ = J_c^T (J_c J_c^T)^{-1}. \quad (3.25)$$

This specific choice of right inverse has the property of minimizing the summed squares of the resulting solution in joint space.

Using this right-inverse, we arrive at the implication

$$\ddot{\theta} = J_c^+ f_c \quad \implies \quad \ddot{c} = J_c \ddot{\theta} = J_c J_c^+ f_c = f_c. \quad (3.26)$$

For any desired center of mass jerk \ddot{c} , we can now calculate a joint jerk $\ddot{\theta}$ that will result in the desired center of mass jerk.

How can the brain generate this joint jerk vector? Deriving the equation of motion 3.16 by time yields

$$M\ddot{\theta} + \dot{M}\dot{\theta} = \dot{T} - C\ddot{\theta} - \dot{C}\dot{\theta} - \dot{N}. \quad (3.27)$$

Again we can assume that during quiet stance, the inertia matrix M is constant, so $\dot{M} = 0$. The term of velocity-dependent forces is so small that we neglect it as well, assuming $C = \dot{C} = 0$. The changes in the gravitational force matrix depend nonlinearly upon $\dot{\theta}$. While these changes are not small, it is not feasible to assume that the CNS can estimate them fast and accurately enough to actually benefit from doing so. Instead, we can assume that the changes in N are treated as a quasi-random perturbation that has to be stabilized against: we set $\dot{N} = 0$ as well. The equation then simplifies to

$$M\ddot{\theta} = \dot{T}, \quad (3.28)$$

which can be used to transform a desired joint jerk vector into a desired torque change vector.

From the point of view of the brain, each joint can be seen as a damped mass-spring system that can be influenced by shifting its threshold parameter λ . How should the threshold parameters be changed in order to get a desired change in torques? Again, we need several simplifications to approach this question. First, we neglect changes in the passive stiffness – while they are not zero, estimating them is not straightforward, so we assume they are treated as unpredictable perturbations similar to changes in gravitational force. Furthermore, we neglect the time delay introduced by the transformation of motorneuron activation into muscle force (Equation 3.12). With these simplifications and equation 3.11, we get

$$T = \tilde{T}_{\text{act}} = AE. \quad (3.29)$$

Deriving this by time and applying the chain rule yield

$$\dot{T} = A \frac{d}{dt} E = A \frac{dE}{d\lambda} \frac{d\lambda}{dt} = AR\dot{\lambda}, \quad (3.30)$$

where the active stiffness matrix

$$R = \frac{dE}{d\lambda}. \quad (3.31)$$

describes the relationship between changes in lambda and changes in muscle activation.

Calculating the partial derivatives that make up R is rather straightforward: The muscle activation is the sum of agonist and antagonist activation, as specified in equation 3.10. The diagonal terms of R are given by

$$\frac{\partial(E_{AG})_i}{\partial\lambda_i} = -\alpha_E \left[e^{\alpha_E(\hat{\theta}_i - \lambda_i + \rho + \mu(\hat{\theta}_i - \lambda_i))} \right]^+ = -\alpha_E (E_{AG} + 1) \quad (3.32)$$

and analogously

$$\frac{\partial(E_{AN})_i}{\partial\lambda_i} = \alpha_E (E_{AN} + 1), \quad (3.33)$$

whereas the off-diagonal terms simply vanish.

Combining equations 3.28 and 3.30 provides a relationship

$$M\ddot{\theta} = AR\dot{\lambda} \quad (3.34)$$

between joint jerks and threshold changes. This allows us to refine the implication 3.26 to

$$\dot{\lambda} = R^{-1}A^{-1}MJ_c^+ f_c \implies \ddot{c} = f_c. \quad (3.35)$$

With the simplification we made, this implication holds because

$$\ddot{c} \stackrel{3.24}{=} J_c \ddot{\theta} \quad (3.36)$$

$$\stackrel{3.28}{=} J_c M^{-1} \dot{T} \quad (3.37)$$

$$\stackrel{3.30}{=} J_c M^{-1} AR \dot{\lambda} \quad (3.38)$$

$$\stackrel{\text{premise}}{=} J_c M^{-1} AR R^{-1} A^{-1} M J_c^+ f_c. \quad (3.39)$$

Using the actual feedback term we stated in equation 3.20 we arrive at the formula

$$\dot{\lambda} = F_c = R^{-1}A^{-1}MJ_c^+ \left(-\alpha_{\hat{z}}\hat{c} - \alpha_{\hat{z}}\hat{c} \right) \in \mathbb{R}^3 \quad (3.40)$$

describing the neural dynamics that stabilize upright stance by reducing sensed movement of the center of mass.

3.2.6.1. Sensory integration

The brain has more sensory information available than just about the movement state of the center of mass. Sensory integration is the process of combining different sensory channels into a coherent percept of the body in space. There are two layers to the problem of sensory fusion. The first problem is to integrate two or more

sensory channels that provide different, possibly conflicting estimates of the *same* state variable. As stated in section 3.2.5, we do not deal with this problem explicitly, but assume it is being solved somehow by the central nervous system. The second problem is how to combine information from sensory channels that estimate *different* state variables. In our concrete case, how do we combine the feedback term for center of mass position 3.40 with feedback for the other state variables head position and orientation?

For head position, our solution is simple and straightforward. First we calculate a similar feedback term depending on sensed head movement states. Analogous to the derivations above, the visual feedback term is given by

$$F_p = R^{-1}A^{-1}MJ_p^+ \left(-\alpha_{\hat{p}}\hat{p} - \alpha_{\dot{\hat{p}}}\dot{\hat{p}} \right) \in \mathbb{R}^3. \quad (3.41)$$

As they both provide modulations of λ , we can add the center of mass feedback and the head feedback terms as two contributions to a single solution

$$\dot{\lambda} = F_c + F_p \in \mathbb{R}^3. \quad (3.42)$$

This simple version of sensory integration works because the directions given by the pseudo-inverses of the Jacobians for center of mass and head position are more or less collinear in joint space, so the two terms mostly support each other, and errors in one contribution are mitigated by the other one. For head orientation, the situation is different. The pseudo-inverse of the head orientation Jacobian is

$$J_o^+ = \frac{1}{3} \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix} \in \mathbb{R}^3, \quad (3.43)$$

i.e. changes of head orientation are distributed equally across all joints. As the Moore-Penrose pseudo-inverse minimizes joint changes, a little change in head orientation can and usually does induce a rather large change in the other relevant state variables center of mass and head position. This is obviously not desirable, therefore we suggest a different right-inverse of J_o to transform changes of head orientation into changes of joint angles. To prevent the undesirable effect, this right-inverse should have the property of leaving the center of mass position invariant. This can be achieved by the augmented Jacobian technique (Siciliano, 1990), setting

$$\tilde{J}_o = \begin{pmatrix} J_o \\ J_c \end{pmatrix} \quad (3.44)$$

and defining

$$F_o = R^{-1}A^{-1}M\tilde{J}_o^+ \begin{pmatrix} f_o \\ 0 \end{pmatrix}, \quad (3.45)$$

where the

$$f_o = -\alpha_o \hat{o} \quad (3.46)$$

and 0 in the rightmost term correspond to the desired changes in orientation and center of mass. The orientation feedback depends upon the current head orientation with a gain factor α_o , instead of on the movement state like the center of mass and head position feedback terms. This is mostly because for head orientation, the state itself is available from the sensory system, in contrast to the other two variables.

Again, we can integrate the feedback contributions from the different sensory modes by adding them up. The final equation for the neural dynamics in the brain transforming sensory information into descending motor commands is

$$\dot{\lambda} = F_c + F_p + F_o + \eta_\lambda. \quad (3.47)$$

The noise term η_λ represents the random factors of the neural processes corresponding to the transformations expressed in the equations above.

3.3. Parameters

In this section, we discuss the choice of some of the model parameters and provide constraints on them from the literature. For a complete overview of all model parameters, please see Appendix B.

3.3.1. Time delays

Many of the variables used in our model depend upon time. Some of the processes described by transformations between variables are instantaneous, while others are abstract descriptions of processes that need time to perform an action. Time delays are a general problem for error feedback systems in general and the stabilization of quiet upright stance in particular (Mergner, 2010).

One tractable case is the second-order low pass filter describing how muscle force is generated in response to motorneuron activation by calcium kinetics (Equation 3.12). This process has been studied in great detail and is theoretically well understood. The relaxation speed of the current torque to the steady state torque specified by the motorneuron activation is described by the time parameter τ_m . Feldman and colleagues used a value of $\tau_m = 10$ ms (St-Onge et al., 1997), while Gribble and colleagues use a slightly higher value of $\tau_m = 15$ ms (Gribble et al., 1998). Micheau and colleagues use an even higher value of $\tau_m = 20$ ms, but approximate the low pass filter by a fixed time delay of $2\tau_m = 40$ ms (Micheau et al., 2003). In our model, we chose the middle ground with $\tau_m = 15$ ms.

The spinal reflex loops transforming deviations sensed by muscle spindles into motor-neuron activity also take time. We do not model the dynamics of these processes explicitly. Instead, we assume a fixed time delay d_θ for the sensor signal to arrive in the spinal cord and be transformed into motorneuron activation there. The

value of d_θ is constrained by experimental data. Values for different physiological reflex loops vary. The fastest response times are found in monosynaptic reflex loops. The latency between the electrical stimulation of human calf muscles (triceps surae) and motorneuron activation is between 30–35 ms, the range depending mostly upon the distance between the muscle spindles and the motorneurons, i.e. the subject’s leg (Rothwell, 1987; Latash, 2008). Polysynaptic reflexes have a longer delay time, due to the additional time requirements of the neural processing involved. Depending upon the specific reflex loop and the muscle involved, the reported delay time lies within 30–70 ms (Matthews, 1994; Latash, 2008). We choose the lower end of this interval as time delay d_θ in our model. This value well within the range of time delays used in similar models of motor control. Gribble and colleagues use 25 ms (Gribble et al., 1998) for upper extremity reaching. For postural control, van Soest and colleagues use 35 ms (Van Soest, Haenen, & Rozendaal, 2003). Most other models of quiet upright stance do not differentiate between the short delays of proprioception in the spinal feedback loop and the long delays of afferent sensory information in the brain.

The time requirements of the neural processes in the brain are much less constrained by neurophysiology. One constraint is the measured latency of responses to movement-related perturbations that are processed in the brain. In one experiment, Smeets and colleagues studied hand movements towards targets that could be identified among other objects by a number of different characteristics like color, shape or orientation. During the movements, the target sometimes swapped location with another object unexpectedly. The time after which the perturbed movements started to differ from the unperturbed ones varied between ≈ 120 – 220 ms, depending on the identification characteristics (Veerman, Brenner, & Smeets, 2008). These numbers are delay times for the complete visuomotor loop, of course, so they include the time delays that we already accounted for explicitly, most notably the processing time of the spinal reflexes and the force generation of muscles. Other researchers have fitted the time-delay in feedback models to experimental data. Time delays in the range of ≈ 160 – 200 ms between sensory stimuli and observable change in the behavioral sway are reported by Peterka (2002) and Maurer et al. (2006). The models used in these studies do not distinguish between time delays from neural processing in the brain and spine, and force generation from the calcium kinetics in the muscle, so reported values have to be interpreted as the sum of these three effects. As we already accounted for the other two delays with values that sum up to ≈ 70 – 80 ms, we choose the value $d_c = 120$ ms to reach a total delay in the appropriate range.

All other processes are modeled by differential equations, so the dependency on time is explicitly included.

3.3.2. Muscle contraction

The equation we have used for the λ -model of the Equilibrium Point Hypothesis was used in the same or similar form in several modeling studies over the last years (Laboissière et al., 1996; St-Onge et al., 1997; Gribble et al., 1998; Micheau et al.,

2003; Kistemaker, Soest, & Bobbert, 2006; Martin, Scholz, & Schöner, 2009). Most of these models deal with different motor systems like arms or jaw, though, and none of them uses exactly the same equations as our model, which makes parameter comparison difficult.

The model that is closest to our case is the one used by Micheau and colleagues in a simulation study of upright stance (Micheau et al., 2003). In this study, the gain between muscle spindle and motorneuron activation α_E varies between 30–55 rad^{-1} for α_E , with the time constant μ set to 90 ms. The transformation parameter between motorneuron activation and joint torque is $A = 5 \text{ Nm}$. One difference between their model and ours is that they do not include passive stiffness and damping of the muscle-tendon complex, nor co-contraction.

A study of single-joint movements was done by the group of Anatol Feldman (St-Onge et al., 1997). They simulate elbow movements by assuming time-shifting descending commands for λ , ρ and μ . They use values of $\alpha_E = 0.05 \text{ deg}^{-1} \approx 2.86 \text{ rad}^{-1}$ with μ varying between 0.05 and 65 ms and $A = 1.2\text{--}1.4 \text{ Nm}$.

Many studies model on a level of single muscles rather than single joints. Paul Gribble and colleagues used a model of the upper arm with two joints and six muscles to fit the parameters of a λ -model to experimental data (Gribble et al., 1998). For the static case, they used fixed values of $\alpha_E = 0.112 \text{ mm}^{-1}$, $\mu = 60 \text{ ms}$ and $A = 1 \text{ N cm}^{-2}$. Transforming these into joint angles from muscle length results in values similar to those of St-Onge et al.

Are these parameter settings used by different groups for different effector systems consistent? The difference between the form parameter α_E in the two models is difficult to understand. Gribble and colleagues state that α_E is a form parameter that is the same for all muscles (Gribble et al., 1998). This statement holds for actual muscles though, not joint-muscles that model combinations of several actual muscles acting on one joint. The transformation between parameter values relative to joint angles and values relative to muscle lengths is done via the moment arm of the muscle (Shadmehr & Arbib, 1992). For the elbow joint, a moment arm of 0.02 m transforms the value $\alpha_E = 0.122 \text{ mm}^{-1}$ (Gribble et al., 1998) into 2.44 rad^{-1} , which is close to the 2.86 rad^{-1} used by St-Onge et al. For the ankle joint, a moment arm of 0.044 m results in $\alpha_E = 5.36 \text{ rad}^{-1}$, which is still very different from the range 30–55 rad^{-1} reported by Micheau et al.

The method used by Micheau and colleagues to determine α_E consisted of a perturbation from a ball swinging on a pendulum that hit the subject into the lower back during quiet stance with a force of well-defined, fixed magnitude (Kron, 1997). These perturbations resulted in resisting ankle torques that peaked at about 10 Nm and vanished after 2 s. The authors fitted a linearized version of the λ -model invariant characteristic (compare equations 3.9, 3.10) to these responses to determine α_E , assuming that λ is constant over the measured interval and having estimated μ and A using other methods.

Two effects are disregarded in this method. First, the assumption that λ is constant during the relevant interval after the perturbation is disputable. As discussed above in section 3.3.1, a reasonable estimate of the interval between a perturbation and

neural response in the brain is around 120–220 ms, not 2 s. As demonstrated by other researchers, it is not feasible to assume that the spinal reflex loops alone can stabilize the body in upright stance (Morasso & Schieppati, 1999; Morasso & Sanguineti, 2002; Van Soest et al., 2003). We must assume that Kron’s estimate for α_e includes shifts in the activation threshold λ that generate additional torque. The second factor neglected by Kron is torque generated by passive elasticity of the tissue. Considering these two factors, we conclude that the range of 30–55 rad^{-1} for α_E is a significant overestimation.

During the development of the model it turned out that both the value range used by Micheau and colleagues, $\alpha_E = 30\text{--}55\text{ rad}^{-1}$, and the transformed value from St-Onge and colleagues, $\alpha_E = 5.36\text{ rad}^{-1}$ do not work. For small values of α_E , strong neural feedback gains α_p, α_c and α_o can prevent the body from falling down. For our case of stabilizing quiet stance of a multi-segment inverted pendulum however, we were unable to generate a parameter setting for which the feedback holds the body in a stable configuration. Substantial oscillations with low frequency persisted for a low parameter value of $\alpha_E = 6\text{ rad}^{-1}$. For a high value of $\alpha_E = 30\text{ rad}^{-1}$, the body inevitably started oscillating with high frequency of $\approx 5\text{ Hz}$ and increasing amplitude in a mode where ankle and hip were in-phase and the knee phase shifted by almost exactly 180° , even in the absence of feedback in the brain. For these reasons, we chose an intermediate value of $\alpha_E = 12\text{ rad}^{-1}$ for our model.

The difference in the parameter A between the models of Micheau and St-Onge is easier to explain. The parameter A describes the capabilities of a muscle system to generate force or torque relative to motorneuron activation. It depends upon the physiological cross-sectional area (PCSA) of the actuating muscle(s). The PCSA is a combination of the anatomical cross-sectional area (ACSA) and the pennation angle of the muscle fibers (Winters & Woo, 1990). Gribble et al. define A relative to the PCSA of the modeled muscle system and use a value of 1 N cm^{-2} . The total PCSA of the ankle plantarflexors is estimated as 124.3 cm^2 (Ward, Eng, Smallwood, & Lieber, 2009). Transformation from muscle forces to joint torques is done via the moment arm of the muscles in question (Shadmehr & Arbib, 1992). The main ankle plantarflexors are the soleus and the gastrocnemius. Both muscles share a tendon that is fixed to the heel behind the ankle, with a moment arm of 0.044 m (Van Soest et al., 2003). This yields a value of $A = 1\text{ N cm}^{-2} \cdot 124.3\text{ cm}^2 \cdot 0.044\text{ m} = 5.47\text{ N m}$, which is very close to the value used by Micheau and colleagues.

What are the values A for the knee and hip joints? Ward and colleagues report the combined PCSA of the knee extensors as 88.4 cm^2 and of the hip extensors as 73.4 cm^2 (Ward et al., 2009). Moment arms of the major muscles in each group are 0.042 m at the knee for the vasti and 0.062 m for the glutei at the hip (Van Soest et al., 2003). This results in parameter values 3.71 N m at the knee and 4.55 N m at the hip.

The main contributors in each muscle group are mono-articular, i.e. they act only on one joint. Nonetheless, we have to account for the presence of bi-articular muscles that span the ankle and knee or the knee and hip. The relevant bi-articular muscle actions are the gastrocnemius, semitendinosus and semimembranosus on the knee

	PCSA (cm ²)	moment arms (m)		
		ankle	knee	hip
gastrocnemius	30.8	+0.044	+0.018	
hamstrings	23.2		+0.026	+0.077
rectus femoris	13.5		-0.042	-0.035

Table 3.1.: Physiological cross-sectional area and moment arms for the relevant bi-articular muscles in the leg. PCSA according to Ward et al. (2009), moment arms according to Van Soest et al. (2003). Gastrocnemius includes medial and lateral head, hamstring stands for semitendinosus and semimembranosus.

and the rectus femoris on the hip. Their PCSA and moment arms are summed up in table 3.1. The presence of these bi-articular muscles is captured by the off-diagonals in the matrix A of our model. The action of the gastrocnemius on the knee results in a value of 0.55 N m, the semitendinosus and semimembranosus on the knee in 0.60 N m and the rectus femoris on the hip in 0.47 N m. Combining these with the single-joint values from above, and multiplying by two for having two legs, results in the parameter matrix

$$A = \begin{pmatrix} 10.94 & 1.1 & 0 \\ 0 & 7.43 & 1.2 \\ 0 & 0.94 & 9.10 \end{pmatrix} \text{ N m.} \quad (3.48)$$

3.3.3. Passive effects

Passively generated torques consist of elastic torques depending on the joint angle configuration and viscous torques depending upon the velocity of joint angle motions. Passive elastic effects at the ankle joint during quiet stance have been extensively studied using a wide range of methods. Usually it is formulated as a linear stiffness parameter, relating displacements in joint angle to the magnitude of ankle torque in the opposite direction generated by the passive structures around the joint. We have chosen to adopt the explicit dependency of passive torque on joint angles proposed by Riener and Edrich, along with the parameters estimated in their study (Riener & Edrich, 1999). To compare the elastic torques to stiffness parameters in other models, we linearize these equations around the initial configuration used in our simulation study (see below, 3.50), which yields an ankle stiffness of $K = 52.3420 \text{ N m rad}^{-1}$.

The starting point for extensive research into passive ankle joint stiffness during quiet stance was a study by Winter and colleagues (Winter et al., 1998). The authors argue that long delays of more than 150 ms in the postural control cycle would be evident in the temporal relationship between center of pressure and center of mass trajectories. As experimental data do not exhibit such delays, control cannot be reactive. Instead, Winter and colleagues proposed a “stiffness control model”, hypothesizing that the CNS stabilizes the body as a single-link inverted pendulum by

setting a constant muscle tone at the ankle and then leaving the control to the fast spinal feedback loops. By fitting the amplitude spectrum of the mechanical system to experimental data, they estimated a stiffness parameter value of $K = 802 \text{ N m rad}^{-1}$ in normal conditions.

In a subsequent paper, Morasso and Schieppati (1999) argue against the “biological inevitability” of that model by stating that the short time-lag between CoP and CoM is simply a result of the biomechanical features of the intrinsically unstable inverted pendulum. They further argue that the stiffness of the muscles from the spinal feedback loop is not large enough to stabilize the inverted pendulum, citing estimates from several studies that are reproduced and extended in table 3.2.

The difference of methods used in this study makes it hard to compare the results directly. Four studies out of six attempted to separate active effects from purely passive torques. These estimates are all larger than the value from the purely passive experiment with relaxed muscles by Riener and Edrich. It is possible that taut muscles also exhibit some passive elastic properties that are responsible for this difference. Some studies report that the passive elastic stiffness depends upon the baseline muscle force (Hof, 1998; Blanpied & Smidt, 1992; Weiss, Hunter, & Kearney, 1988; Mirbagheri et al., 2000), while others state that it does not (Loram & Lakie, 2002), though the increase in baseline torque analyzed in the latter study was small compared to the others. Another source of the difference in results might be that stiffness depends upon the magnitude and velocity of perturbations (Joyce, Rack, & Westbury, 1969).

We have decided to use the formulas from Riener and Edrich, for several practical reasons, although they likely underestimate the torques by neglecting the increased stiffness from taut muscles. First, when modeling the passive elastic effects as a rotational spring with a stiffness parameter, one must specify the resting angle of that spring. While the stiffness parameter is estimated in the cited studies, the resting angle is not specified. Second, while there are several studies that estimate stiffness at the ankle, there are almost no similar studies that estimate stiffness at the knee and hip joints (but see Whittington, Silder, Heiderscheit, and Thelen (2008); Silder, Heiderscheit, and Thelen (2008) for studies in walking). The off-diagonal elements of the stiffness matrix that represent passive elastic torques from bi-articular muscle-tendon complexes have not been estimated experimentally, to the best of our knowledge. While it is possible to determine these parameters making some broad generalizations, the effects they describe are well captured by the model by Riener and Edrich. Lastly, it turns out that with the model as presented here, the magnitude of the passive elastic effects does not have a major effect on the characteristics of the generated movement trajectories.

Passive viscous effects of the leg joints have received much less attention than elastic torques in the literature. Some of the studies examining parameters for ankle stiffness have also estimated damping parameters, as summarized in table 3.2. Again, the range of reported parameters varies by an order of magnitude. Passive viscosity is often assumed to be non-linear, being large for small velocities and then decreasing as velocity magnitudes get larger (Tee et al., 2004; Loeb et al., 1999; Barto et al.,

stiffness (N m rad ⁻¹)	viscosity (N m rad ⁻¹)	type	perturbation type	perturbation time	reference
802 ± 304	–	act. + pass.	–	–	Winter et al., 1998
500–800	–	passive	large displacement	fast	Hof, 1998
240–700	–	act. + pass.	large displacement	slow	Blanpied & Smidt, 1992
596 ± 138	5.73	passive	small displacement	fast	Loram & Lakie, 2002
92 ± 29	25 ± 11	passive	stochastic	continuous	Peterka, 2002
824 ± 128	1.4 ± 0.16	passive	stochastic	fast, intermittent	Mirbagheri, Barbeau, & Kearney, 2000

Table 3.2.: Results of different studies estimating the stiffness of the ankle joint.

1999). This is a potential explanation for the discrepancies in the viscosity estimates between different studies.

As the velocities are always small during quiet stance, we have neglected the non-linearity of the viscosity and choose the magnitude parameter for the ankle joint from the upper end of the reported values. For the upper joints and the off-diagonal entries of the viscosity matrix, we determine values in a way such that B is proportional to the muscle PCSA matrix A , resulting in

$$B = \begin{pmatrix} 25 & 2.5137 & 0 \\ 0 & 16.9790 & 2.7422 \\ 0 & 2.1481 & 20.7952 \end{pmatrix} \text{ N m s rad}^{-1}. \quad (3.49)$$

3.3.4. Free parameters

The parameters discussed in Section 3.3.1–3.3.3 describe physiologically well understood relationships and are strongly constrained by the literature. The model of the how representations of the body in space are established using sensory data and the neural dynamics transforming these into descending motor commands are functional and not constrained by the physiological data. They are considered free parameters for the process of fitting the model to experimental data. These are the parameters

- feedback gain parameters: $\alpha_{\dot{c}}, \alpha_{\ddot{c}}, \alpha_{\dot{p}}, \alpha_{\ddot{p}}, \alpha_o$
- noise magnitude parameters: $\sigma_{\dot{c}}, \sigma_{\ddot{c}}, \sigma_{\dot{p}}, \sigma_{\ddot{p}}, \sigma_o, \sigma_{\dot{\lambda}}, \sigma_m, \sigma_{\theta}, \sigma_{\dot{\theta}}$

These 14 parameters are constrained by the data obtained from behavioral experiments on quiet upright stance. As the concrete trajectories of the postural fluctuations are a result of how the CNS reacts to estimation imperfections and random perturbation within the neural processes and motor systems, reproducing the measured trajectories explicitly would be equivalent to estimating these random factors, instead of providing insight of how the CNS deals with them (though Micheau et al. (2003) did fit λ -trajectories of a similar model to joint angle trajectories). Instead of fitting explicit trajectories, we attempt to reproduce the variability patterns shown by these data sets. We choose three different classes of variability measures: the variance within each joint, the variance relative to different task variables as defined by a UCM analysis, and the frequency dependence of the fluctuations (see Section 3.4.1 below). The free model parameters listed above were fitted by hand to reproduce these measures reasonably well.

Though only loosely linked to the neurophysiology, these parameters still have a meaning and cannot be chosen arbitrarily. One phenomenon we are trying to explain is the difference in variability data between conditions with eyes open (EO) and eyes closed (EC). This difference has to be reflected in the parameters for the sensory channels containing vision. We do that by assuming that the accuracy of the estimates for head movement and orientation suffers when the eyes are closed, i.e.

	EO	EC	unit		value	unit
$\sigma_{\dot{c}}$	0.02		m s^{-1}	$\alpha_{\dot{c}}$	1.5	$\text{rad}^{-1}\text{s}^{-2}$
$\sigma_{\ddot{c}}$	0.02		m s^{-2}	$\alpha_{\ddot{c}}$	0.5	$\text{rad}^{-1}\text{s}^{-1}$
$\sigma_{\dot{p}}$	0.015	0.02	ms^{-1}	$\alpha_{\dot{p}}$	9	$\text{rad}^{-1}\text{s}^{-2}$
$\sigma_{\ddot{p}}$	0.015	0.02	ms^{-2}	$\alpha_{\ddot{p}}$	3	$\text{rad}^{-1}\text{s}^{-1}$
σ_o	0.02	0.03	rad	α_o	60	$\text{rad}^{-1}\text{s}^{-3}$
σ_θ	0.002		rad			
$\sigma_{\dot{\theta}}$	0.002		rad s^{-1}			
σ_λ	0.001		rad s^{-1}			
σ_m	0.02		–			

Table 3.3.: Model parameters used in the simulations. Noise parameters may vary by experimental condition (left), gain parameters are fixed (right).

the noise magnitude parameters $\sigma_{\dot{p}}$, $\sigma_{\ddot{p}}$ and σ_o are larger for condition EC than for EO. Table 3.3 sums up the values of all model parameters used.

3.4. Results

The model presented in section 3.2 was implemented and used in several simulation studies. We explore the role of different factors for the whole system by selectively turning them off. Then we compare the structure of the postural sway generated by the model to experimental data.

3.4.1. Data acquisition and analysis

We reanalyze a data set of a quiet stance experiment where 10 human subjects stood upright with arms crossed on a normal support surface for 5 minutes (Hsu et al., 2007). Nine infrared markers 1cm in diameter were attached to the subjects' body (for details, please refer to Hsu et al., 2007). The marker positions were recorded using a VICON visual motion measurement system (Oxford Metrics) at 120 Hz. For comparison with the model, the marker data were transformed into joint angles for the ankle, knee and hip joints.

The present model encompasses stabilizing feedback on a short and medium time scale. Drifts in the threshold positions λ over a long time scale might still lead to configurations that are unstable. It is assumed that the CNS has additional feedback cycles to identify and counter these slow drifts, but these are not part of the current model. For that reason, the experimental data was separated into episodes of 30 seconds each. Each trial yielded 8 such episodes, the first one starting 10 seconds after trial start and each subsequent one starting where the previous one stopped. This resulted in a total of 240 episodes from 10 subjects with 3 trials each.

The variance of the data episodes was analyzed with three different measures. To

calculate total joint excursion variance (JEV), the variance of each joint trajectory in each episode was calculated separately, then the mean over the episodes was taken to get a JEV measure of each joint. In addition to that, a UCM analysis of each episode was performed with respect to the anterior-posterior position of center of mass and head position and the head orientation around the media-lateral axis as candidate task variables. Again, the mean over all episodes was taken to get measures V_{\parallel}^* , V_{\perp}^* for $* = c, p, o$ (see Chapter 2).

One important measure of postural fluctuations during quiet upright stance is the power spectrum of the relevant variables (Collins & De Luca, 1994; Zatsiorsky & Duarte, 1999; Saffer, Kiemel, & Jeka, 2008). We estimated the power spectrum of the center of mass position and the joint angles for each trial with the Matlab PSD function that implements Welch’s averaged, modified periodogram spectral estimation method, using a 20s Hamming window with 50% overlap (Bendat & Piersol, 1993). Because this method already uses a time-window of medium length that is moved in steps over the data, we analyzed the whole trials instead of the separated episodes, resulting in 30 PSD estimates. A population estimate was formed by taking the mean at each frequency over all subjects and trials.

3.4.1.1. Model data

The model described in Section 3.2 was implemented in Matlab. The differential equations were solved with the Euler method, using a time step of 2 ms.

The initial kinematic state of the body was chosen to be the same as in the model of Van Soest et al. (2003)

$$\theta^{(0)} = \begin{pmatrix} -0.1 \\ 0.2 \\ -0.2 \end{pmatrix}, \quad \dot{\theta}^{(0)} = 0, \quad (3.50)$$

resulting in an initial center of mass position of $c \approx 3$ cm in front of the ankle joint. The initial values of the threshold parameters λ were chosen such that the sum of passive and active torques exactly cancelled out the gravitational torques at each joint. For a fixed θ , the force-length relationship 3.9 is monotonous in lambda and thus invertible, so $\lambda^{(0)}$ is uniquely determined by 3.50 and requiring $\ddot{\theta}^{(0)} = 0$. To find the actual value for $\lambda^{(0)}$, we first calculated the required initial motorneuron activation pattern $E^{(0)}$ analytically from Equation 3.11, 3.13, 3.15 and 3.16 (assuming $\tilde{T}_{\text{act}} = T_{\text{act}}$). The vector of threshold values actually generating this motorneuron activation via the neural dynamics of the spinal reflexes given in Equation 3.9–3.10 was found by calculating

$$\lambda^{(0)} = \arg \min_{\lambda} \|E(\lambda) - E^{(0)}\|^2 \quad (3.51)$$

using the Matlab function `fminunc`. The initial conditions of all sensor estimates were set to the actual values of the estimated variables. Time delays d were decreased to

the current simulation time t for $t < d$. To prevent effects of these artificial choices on the model results, the first 5 s of each trial were disregarded.

Simulations of the model using different settings were carried out. For each setting, several trials of 35 s length were simulated. Excluding the initial 5 s resulted in a number 30 s episodes for each setting. The variability characteristics JEV, V_{\parallel}^* , V_{\perp}^* and PSD were calculated for the model data in the same fashion as for the experimental data.

3.4.2. Trajectory examples

The model as presented in Section 3.2 was capable of stabilizing the simulated body against fluctuations from estimation errors, processing noise and the destabilizing effects of gravity. This section provides example trajectories of several different important variables over a single episode in the eyes open condition. Example trajectories from a single human trial (subject VK, trial 1, EO) are also presented for visual comparison.

Figure 3.3 shows the joint angle trajectories. Visual inspection already indicates that the variability is similar to the experimental data: both the sway magnitude and the oscillation patterns are not too different. The same holds for the center of mass trajectories shown in Figure 3.3.

Figure 3.4 shows the joint torque trajectories for the ankle joint during the same simulation trial of the model, separated into active torque T_{act} generated by muscle contraction and passive torques T_{ela} , T_{vis} from elastic and viscous properties. The gravitational torque $-N$ is also plotted, with inverse sign for easier comparison. The trajectories show that while the level of passive elastic torque is significant and sometimes exceeds the magnitude of the active torque, its change in magnitude is minimal. Its contribution to counter the destabilizing effects of gravitation and random fluctuations is small. The contribution of the viscous properties is even smaller.

3.4.3. The functional role of spinal and higher feedback

The role of the feedback dynamics of the spine and brain can be assessed by simulating the model in their absence. Removing the higher feedback of center of mass motion, head motion and head orientation is equivalent to setting $\dot{\lambda}(t) = 0 \forall t$. Removal of the spinal feedback loop was modeled by making the motorneuron activation E constant except for noise.

In both conditions, the model fails to hold the body upright. As these failure trials leave the regime of quiet stance, some of the assumptions we made do not apply anymore. This is not critical, because in these simulations we do not attempt to model the system accurately over the whole duration of movement, but intend to highlight the role of certain parts of the model by analyzing when and how it fails to balance the body. One aspect that we address nonetheless is the generation of unrealistically high torques in situations where the joint angles are very far from

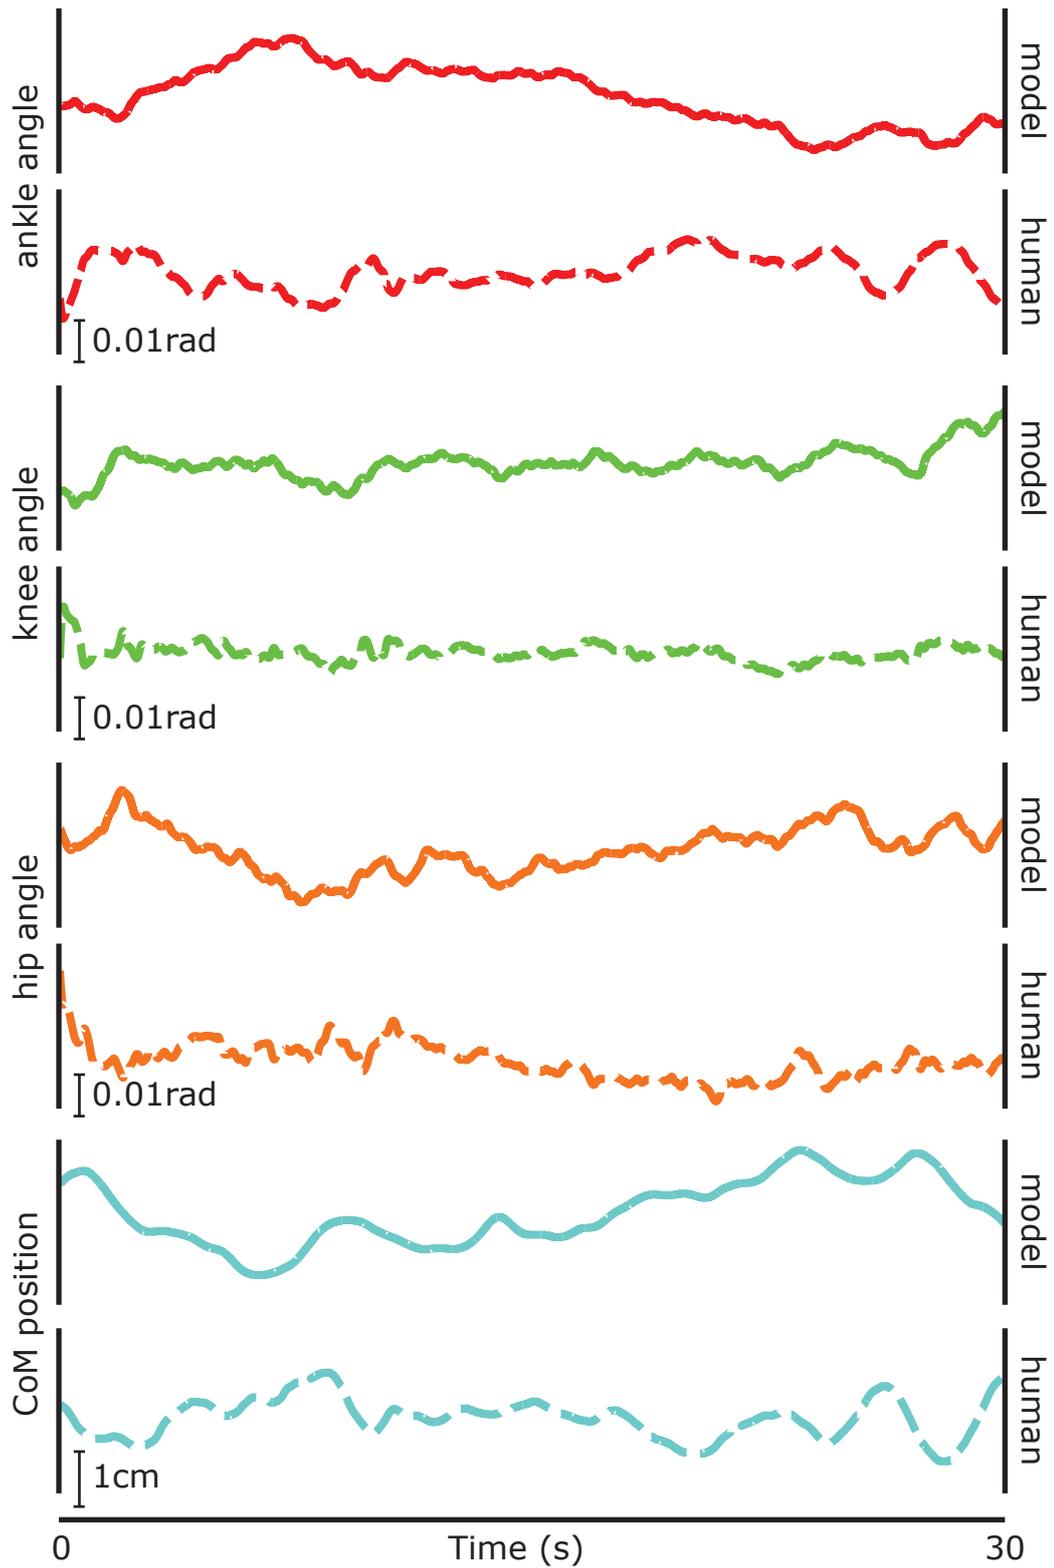


Figure 3.3.: Example trajectories of the joint angles and anterior-posterior center of mass position from one model simulation (solid) and one human trial (dashed).

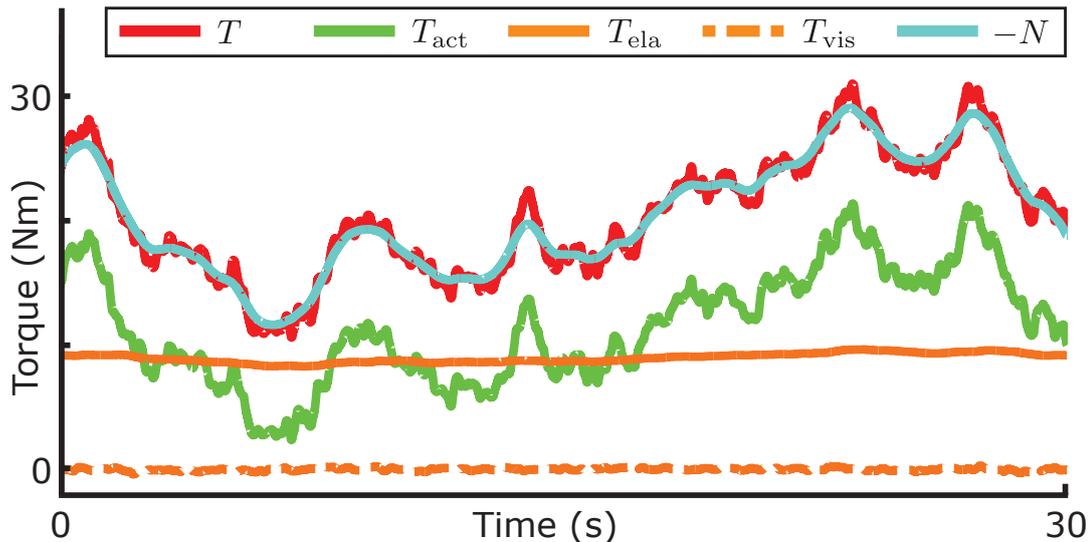


Figure 3.4.: Trajectories of the torques acting on the ankle joint from the same model simulation as in Figure 3.3. The total torque T from the muscle-tendon system is the sum of torque from active muscle contraction (T_{act}) and passive elastic and viscous (T_{ela} , T_{vis}) torques. On average, the torques generated by the muscle-tendon system cancel out the gravitational torque N .

the threshold values λ . To avoid this, we introduce a boundary E_{max} for the motoneuron activation $E_{\text{AG}}, E_{\text{AN}}$. At a point ϕ_0 where the combined activation of the proprioceptive pathways

$$\phi = \hat{\theta} - \lambda + \rho + \mu(\hat{\theta} - \lambda) \quad (3.52)$$

pushes the motoneuron activation

$$E_{\text{AG}} = e^{[\alpha_E \phi]^+} - 1 \quad (3.53)$$

to 75% of that boundary, we switch from the exponential function 3.9 to a hyperbolic function

$$E'_{\text{AG}} = \frac{c_1}{\alpha_E \phi + c_2} + c_3, \quad (3.54)$$

with c_1, c_2 and c_3 chosen in a way that the function is continuous at the switching point, has continuous derivative, and converges to E_{max} for $\phi \rightarrow \infty$. The maximal motoneuron activation was set to $E_{\text{max}} = 10$ resulting in a maximal torque at the ankle joint close to the maximal isometric torque measured in humans (Hasson, Miller, & Caldwell, 2011).

Figure 3.5 illustrates the effect of removing higher feedback: at first the body is partly stabilized by the spinal feedback and does not move much. But the gain of the spinal feedback loop is not sufficient to counter the destabilizing effects of gravity:

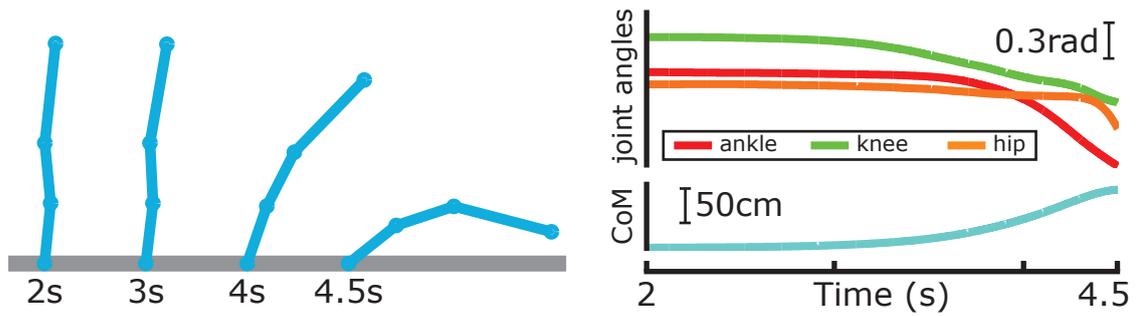


Figure 3.5.: Results of removing the higher feedback dynamics. The simulated body falls over and hits the floor after ≈ 5 s. The left panel shows a series of body configurations during falling. The right panel gives the time course of the joint angles and the anterior-posterior position of the center of mass.

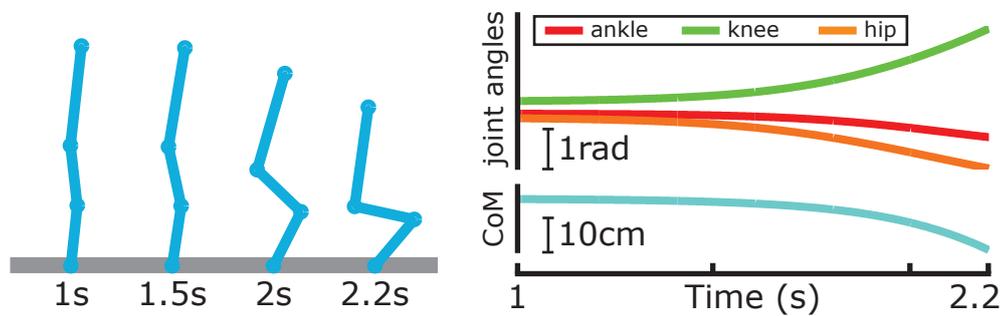


Figure 3.6.: Results of removing the spinal reflex loop and assuming the activation level of the motorneurons to be constant. The simulated body buckles and falls down after ≈ 2 s. The left panel shows a series of body configurations during falling. The right panel gives the time course of the joint angles and the anterior-posterior position of the center of mass.

after about 3 seconds the body starts falling forward. All joint angles increase under the gravitational pull, until the body hits the floor. The center of mass in anterior-posterior shows the same movement pattern as the joint angles. Note that the final part of the movement is not realistic, as the heel would lift off the floor at some point.

It is worth comparing this to the case where spinal feedback is also taken away, as shown in Figure 3.6. In this case, instead of toppling over at the ankle with all joint angles decreasing, the body buckles: the knee angle starts increasing, while ankle and hip angle decrease, resulting in a folding movement. The center of mass movement in anterior-posterior direction is much smaller than for the case with spinal and without higher feedback.

These movement patterns are persistent. Table 3.4 summarizes the results of 1000 trials each with no higher feedback and no spinal feedback by giving the percentage of trials that are currently in the toppling mode at a given point in time, i.e. for which

condition	% of toppling trials after			
	1s	2s	3s	4s
no higher feedback	16.5	85.3	98.4	99.9
no spinal feedback	0.1	0	0	0

Table 3.4.: Percentage of toppling trials.

all joint angles have moved in the same direction since the beginning of the trial. Without any feedback, that number is essentially zero at any time. With spinal but no higher feedback, that number starts relatively low, then converges to 100% as time increases. These results highlight the functional role of the two neural feedback loops. The faster spinal stabilizes against buckling, the slower higher feedback stabilizes against toppling.

3.4.4. Comparison of variability measures

To assess whether the structure of the random walk generated by the simulated model is similar to that of the human data, an extensive simulation study was carried out. Like the human experiment, the study consisted of two conditions, eyes open (EO) and eyes closed (EC). A total number of 240 trials of 35 s length was simulated for each condition. For the analysis, the first 5 s were disregarded. From this data set, we calculated the variance measures V_{\parallel} and V_{\perp} for the task variables c , p and o , the power spectral density of θ_i and c , and the joint excursion variability.

The UCM results are presented in Figure 3.7. They are a good qualitative match of the experimental data. The main effect that is captured by the model is that both V_{\parallel} and V_{\perp} increase when eyes are closed for all three task variables. The increase in variance is not homogeneous. V_{\parallel} increases slightly stronger than V_{\perp} for all task variables, as indicated by the increase in the UCM signature, shown in the bottom of Figure 3.7. Again, this same inhomogeneity is also present in the experimental data. As mentioned before, this phenomenon is counter-intuitive. We explore possible reasons for this effect in Section 3.5.2. One difference between the experimental data and the model results shows up in the plots comparing the UCM signatures: the model has smaller c -signatures than the experimental data, but larger o -signatures. The center of mass is not stable enough in the model, whereas the head orientation is too stable, compared to humans. This effect is relatively small, though.

Figure 3.8 shows the mean power spectral density (PSD) of the experimental and the simulation data for the three joint angles and the body center of mass in the eyes-open condition. In general, the model is a good qualitative match of the experimental data. Especially the center of mass power of the model is close to the experimental data across all analyzed frequencies. The characteristic peak, or change of slope, in the CoM spectrum around 0.2–0.3 Hz is present in the model, though slightly less pronounced than in humans (Collins & De Luca, 1994). In the frequency range above 0.5 Hz, the ankle joint sway has less power than the knee and hip joint. While the

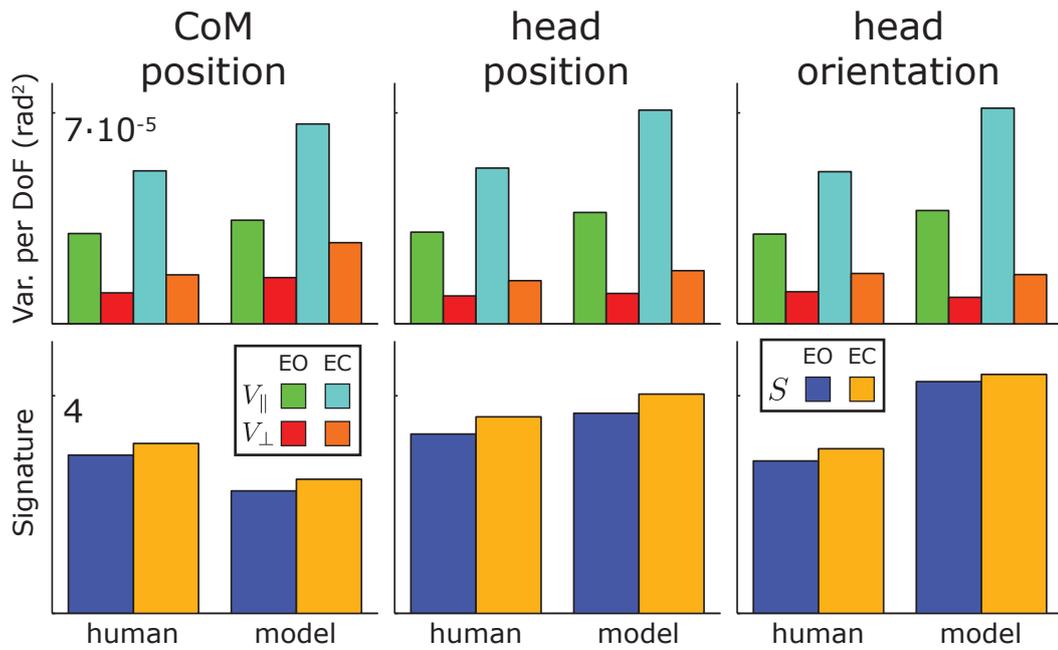


Figure 3.7.: Comparison of the geometrical structure of the postural sway patterns generated by the model with those of humans. The top row shows the mean variance per degree of freedom along the UCM (V_{\parallel}) and orthogonal to it (V_{\perp}), with eyes open (EO) and eyes closed (EC). The bottom row shows the UCM signature $S = V_{\parallel}/V_{\perp}$. The three different task variables are anterior-posterior position of the center of mass (left) and head (middle) and orientation of the head around the media-lateral axis (right).

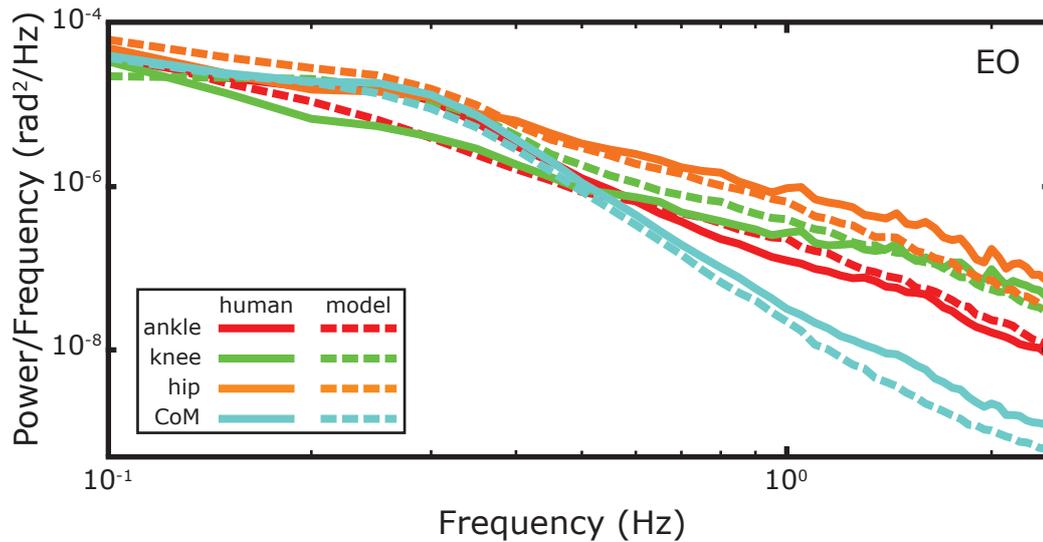


Figure 3.8.: Power spectral density of the joint angles and anterior-posterior center of mass position during quiet stance with eyes open. Solid lines show mean data from human subjects, dashed lines show mean data from the model simulations.

model captures these differences in the high frequency range, the pattern in the range below 0.5 Hz is different. While human postural sway shows a peak in the ankle joint power around 0.2–0.3 Hz, in the model this pattern is shifted to the knee joint and absent in the ankle joint. Again, these differences are discussed below in Section 3.5.2.

A comparison of sway power between the eyes-open and eyes-closed conditions is made in Figure 3.9, showing the PSD in both conditions from both experimental and model data. In general, the slight increase in power across all frequencies is captured well by the model, with approximately the right magnitude. There is a noticeable change in this power difference depending on the frequency: the difference between conditions starts to disappear for frequencies above 1 Hz in the human data. For the joint angles, this effect is reproduced by the model. For the center of mass, the model still exhibits the same difference for high frequencies.

The mean joint excursion variability in both conditions is plotted in Figure 3.10. The general magnitude of the human JEV is captured well by the model. The increase in JEV between conditions is similar for the different joints, in accordance with the human data. The distribution of the variance across joints is different though: compared to the experimental data, the model exhibits more variance in the ankle and slightly less variance in the knee joint.

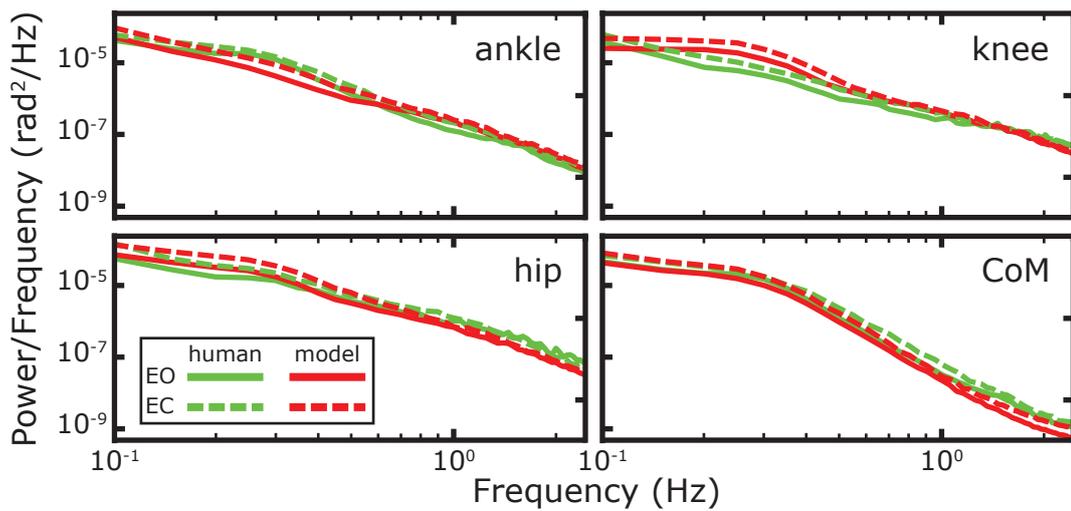


Figure 3.9.: Comparison of the PSD between the conditions eyes-open (EO, solid) and eyes-closed (EC, dashed). The difference between the two conditions is similar for experimental data from humans (green) and the model simulation (red).

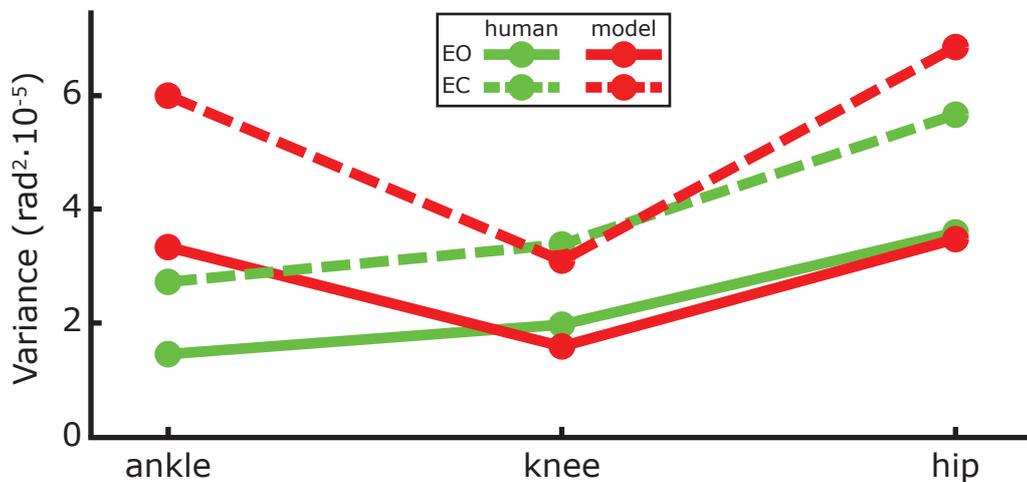


Figure 3.10.: Joint excursion variance means over all subjects and trials for human subjects (green) and over all episodes of the model simulation (red). As with the PSD shown in Figure 3.9, the difference between eyes open (EO, solid) and eyes closed (EC, dashed) is similar in humans and the model.

3.5. Conclusions

We presented a mechanistic model of the neural dynamics within the brain responsible for balancing the body in quiet, upright stance. The model generates a simple negative feedback signal based on sensed deviations from the stable state. This negative feedback is transformed into desired changes of all joint angles along the body in a way that minimizes the total motor activation while still realizing the desired feedback command. The sensory estimates representing the kinematic state of the body in space used by the system are the velocity and acceleration of the center of mass and head position in anterior-posterior direction, and the orientation of the head around the media-lateral axis. The descending motor commands generated muscle activity indirectly by modulating the parameters of the spinal reflex loops.

This negative feedback system was capable of stabilizing the simulated body against random fluctuations from sensory misestimations and neural processing noise. The geometrical and temporal properties of the body sway were reproduced by the model, as indicated by similar patterns of the PSD and variance in UCM space. Removing vision was modeled by increasing the sensory noise for the estimates of the head movement state and orientation. The homogeneous increase of V_{\parallel} and V_{\perp} observed in human subjects when removing vision was reproduced by the model.

3.5.1. Differences between model and experimental data

Though the data produced by the model simulations were a good fit of the experimental data, some substantial differences remain. The mean joint excursion variance (JEV) pattern of the human subjects shows a monotonic increase from distal to proximal joints: the knee varies more than the ankle, the hip varies more than the knee. In the model, the knee has substantially less variance than the ankle joint, and the hip only slightly more than the ankle.

To examine this difference further, we take a closer look at the experimental data. Figure 3.11 shows the mean JEV for each of the 10 subjects separately (the means are taken over the 24 episodes of 30 s length that the data from each subject were blocked in). These plots reveal that the differences *between* subjects are much stronger than the differences *within* the subjects between the EO and EC conditions. The JEV patterns for EO and EC of each subject are similar to each other, indicating that each subject has a unique pattern of postural sway that is consistent across sensory conditions. Also, many of these patterns are as different from the mean pattern shown in Figure 3.10 as the pattern generated by the model simulations. We conclude that while the JEV pattern of the model does not describe the mean pattern of the human subjects very closely, it fits well within the wide range of patterns exhibited by the single subjects.

Another difference between the experimental and model data is the power spectral density. The peak or change in slope of the log-log-plotted PSD of the center of mass at $\approx 0.2\text{--}0.3\text{ Hz}$ is captured by the model. This is the slow component of postural sway that is hypothesized to be the result of the higher feedback loop (Zatsiorsky

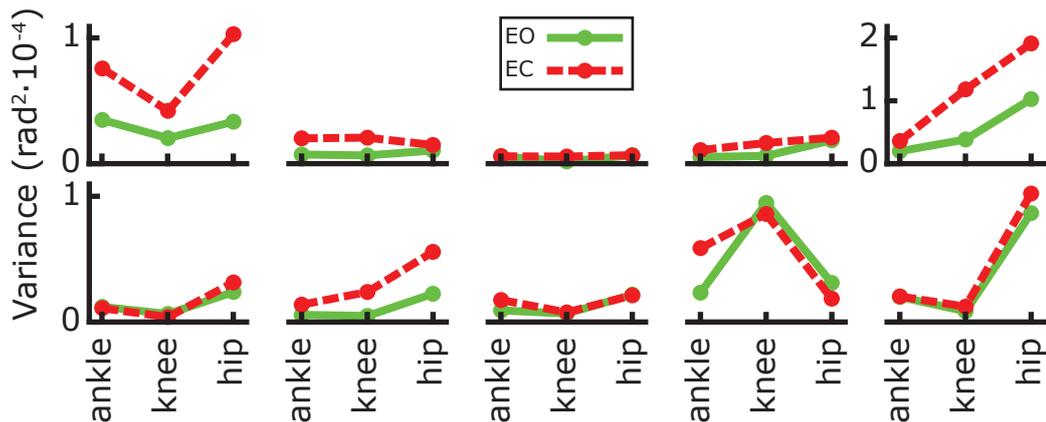


Figure 3.11.: Joint excursion variance means for each single human subject with eyes open (EO, solid green) and eyes closed (EC, dashed red). Note the different scale for an outlier with extremely large variance in the top right.

& Duarte, 1999, see Section 3.1.1). The simulated model exhibits this effect at the knee joint instead of the ankle joint, though.

One possible reason for this discrepancy is that the principle by which the neural dynamics in the brain distribute the error feedback signal among the joints is not appropriate. Maybe the feedback signal is indeed transformed into ankle joint motion exclusively, and the fluctuations of the knee and hip joints are simply side-effects from imperfect control. To test this hypothesis, we altered the feedback distribution among the joints given in Equation 3.41 to a different version

$$\tilde{F}_p = R^{-1}A^{-1}M \begin{pmatrix} \left(\frac{\partial p}{\partial \theta_1}\right)^{-1} \\ 0 \\ 0 \end{pmatrix} \left(-\alpha_{\hat{p}}\hat{p} - \alpha_{\hat{\ddot{p}}}\hat{\ddot{p}}\right) \in \mathbb{R}^3, \quad (3.55)$$

that exchanges the pseudo-inverse of the head position Jacobian with the inverse of the partial derivative of the head position by the ankle joint. This means that only the ankle joint receives descending motor commands to counter sensed deviations. Figure 3.12 compares the PSD results of a simulation ($N = 48$ trials) of this alternate model with the original and the human experimental data. The comparison shows that the dependence of both the ankle and hip angle PSD on the choice of feedback distribution is minimal. Furthermore, the qualitative fit of the experimental data seems to get worse when using the alternative model: the peak in the sway power at $\approx 0.2\text{--}0.3\text{ Hz}$ is visibly less pronounced.

Could the differences between model and human PSD be a result of inappropriately chosen parameters for the brain dynamics? The shape of the PSD is affected by changes in both the noise magnitude and the gain parameters we used to fit the model to the experimental data. During the manual fitting process, however, we were unable

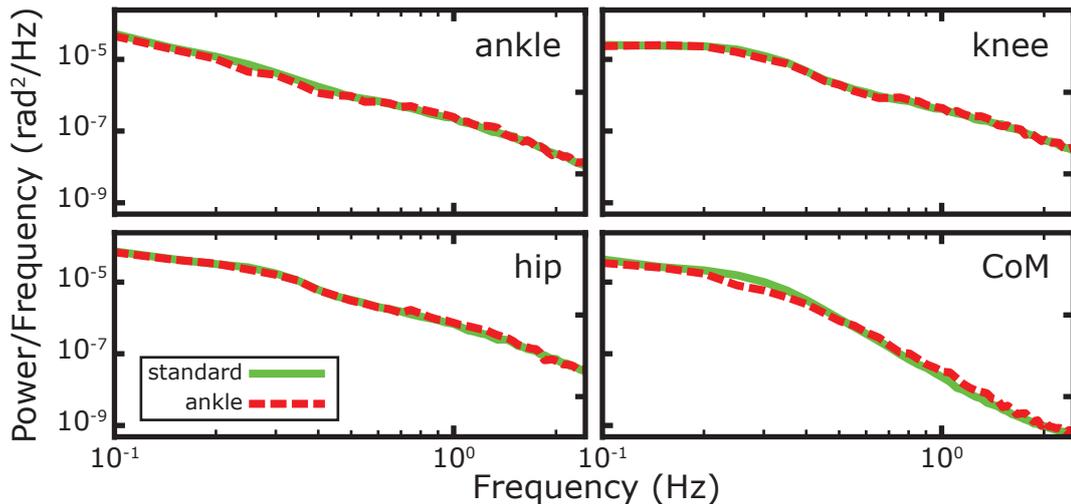


Figure 3.12.: Comparison of the PSD generated by the model with the standard equations (solid green) and when using the alternative equation for sending feedback only to the ankle joint (dashed red).

to find a parameter combination that changed the problematic difference that the peak in the center of mass PSD reappears in the PSD of the knee but not of the ankle. One indicator of this is that increasing the noise magnitude for the head position and orientation estimates results in an increase in PSD that is relatively homogeneous across all joints and does not affect the shape of the curves in a substantial way (see Figure 3.9).

We speculate that the source of the problematic difference lies in the models of the spinal reflex loops and muscle-tendon systems, particularly those acting on the knee. The knee joint is near its maximal extension at quiet stance. This fact is partly accounted for by an additional term in Equation 3.13 describing passive elastic torque, which results in a sharper increase in elastic torque when approaching the maximal knee extension. For extreme knee extensions, this is almost certainly an underestimation: Figure 3.13 plots the passive-elastic knee torque for a range of knee angles near the extension limit according to Equation 3.13 for ankle and hip angles fixed $\theta_1 = 0.1 \text{ rad}$, $\theta_3 = 0 \text{ rad}$. The gravitational torque for the knee joint is also plotted. This comparison reveals that the passive elastic torque is large enough to counter the gravitational torque only for knee angles below $\approx -3.5 \text{ rad} \approx -20^\circ$. Knee extensions of this magnitude are anatomically extremely unrealistic. We infer from these arguments that the physical constraints of extreme knee extensions in quiet stance are not described properly by Equation 3.13.

Yet the limits of knee extension play a role in balancing the body. Consider a contraction of the calf muscles that generates a positive ankle torque. Due to the inertial properties of the body, this would result in a negative interaction torque at the knee joint. But near maximal extension, this interaction torque is negated by the joint limit. This means that the muscle contraction generates movement mostly

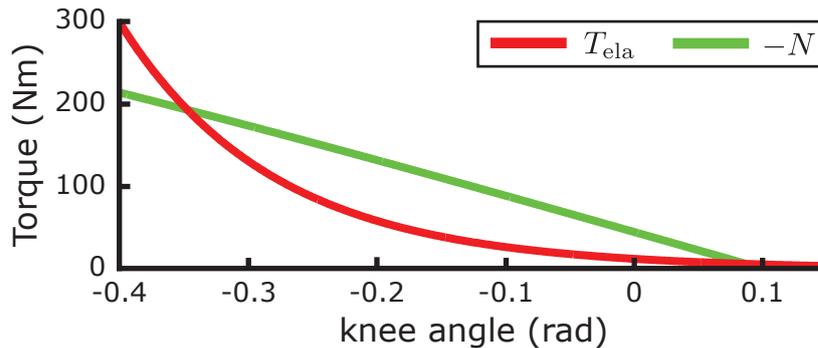


Figure 3.13.: Torques at the knee for extreme joint extensions. The passive elastic torque (red) exceeds the gravitational torque (green) only at extreme values.

in the ankle joint.² By underestimating the forces generated by the physical limits to knee extensions, the model fails to account for this effect and generates movement patterns in the knee that would otherwise have occurred in the ankle. One such pattern could be the oscillatory component around $\approx 0.2 - 0.3$ Hz.

3.5.2. Effects of closing the eyes

Closing the eyes leads to increased postural sway. The UCM analysis revealed that this increase is of similar magnitude in V_{\parallel} and V_{\perp} for all three task variables c , p and o , though V_{\parallel} increases slightly more than V_{\perp} . This is puzzling, because one might intuitively expect that making the estimate of the head movement state less reliable would lead to erroneous corrective signals due to the decreased signal-to-noise ratio. This would imply that V_{\perp} increases much more than V_{\parallel} .

These considerations apply to both head position and orientation, though. The (linear approximations of the) Uncontrolled Manifolds for these two variables are substantially different from each other: the angle between them in quiet, upright stance is ≈ 0.3 rad $\approx 17^{\circ}$. Is it possible that the additional noise in the head movement state representations generates variance in the orientation-UCM, and vice versa? To test this possibility, we conducted simulations ($N = 240$) of the model with two different parameter settings. In one setting, TR, we increased the noise magnitude for the head movement estimates $\sigma_{\dot{p}}$ and $\sigma_{\dot{o}}$ to the values used for the eyes-closed condition (see Table 3.3). In the other setting, OR, we did the same for the head orientation noise level σ_o . In other words, we selectively removed the sense of head translation in TR and the sense of head orientation in OR.

The results of this simulation study are plotted in Figure 3.14, along with the data from the previous simulations for the EO and EC conditions. Intriguingly, the effects of both increases in noise magnitude is quite similar. For all three candidate task

² We neglect the hip joint in these considerations.

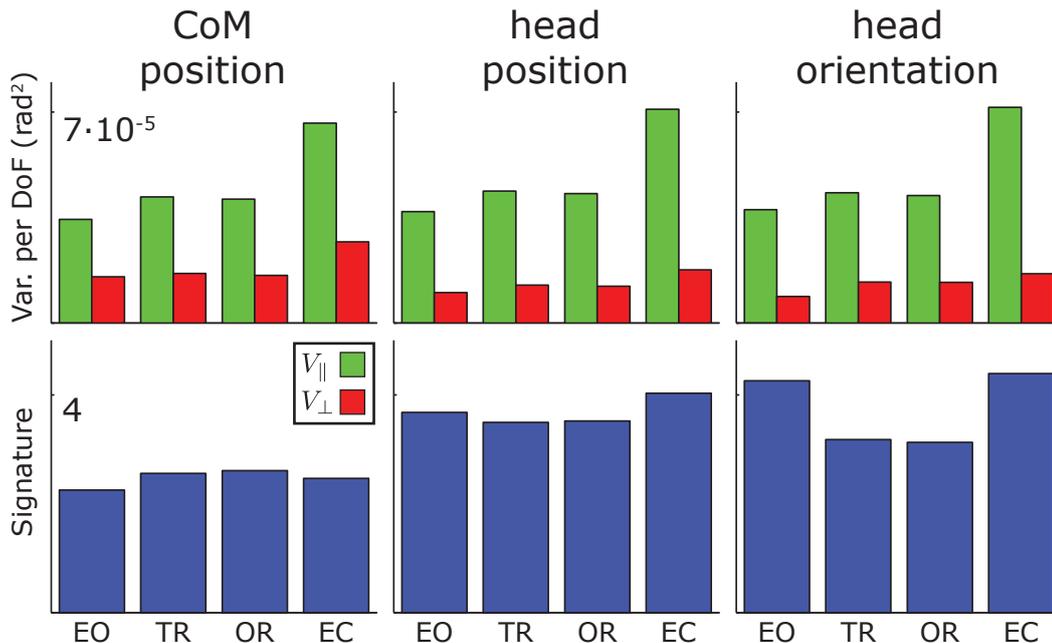


Figure 3.14.: UCM results of the model with translation (TR) and orientation (OR) sense selectively removed. Results of with eyes open (EO) and eyes closed (EC) are reproduced for comparison. The top row shows the mean variance per degree of freedom along the UCM ($V_{||}$) and orthogonal to it (V_{\perp}). The bottom row shows the UCM signature $S = V_{||}/V_{\perp}$. The three different task variables are anterior-posterior position of the center of mass (left) and head (middle) and orientation of the head around the media-lateral axis (right).

variables, increasing either noise level leads to increases in both $V_{||}$ and V_{\perp} . For p and o , the increase in V_{\perp} is slightly larger than the increase in $V_{||}$, as indicated by the decrease in the signature S . This is contrary to the expectation that less reliability of an estimate should lead to more variance in that variable due to erroneous corrective movements. Another unexpected effect is that the increases in variance from adding noise to the head position and orientation sensory estimates separately in conditions TR and OR is small compared to the increase from adding noise to both estimates simultaneously in condition EC. Also surprising is that adding noise to both estimates somehow reverses the effect on the UCM signature, making it slightly larger than in EO. We conclude that the effects of adding noise to different sensory estimates on the variance structure of the sway trajectories are complicated and non-linear, and understanding them requires more detailed analysis.

3.5.3. Outlook

We presented a mechanistic model of the neural dynamics in the brain responsible for balancing the body in upright stance. A simulation study of the whole sensorimotor loop including the neural dynamics of the brain reproduced the characteristic signatures of postural sway. This is a first step in explaining how task-related movement patterns are generated by the CNS.

The model raises interesting question about the neurophysiological foundations of these mechanisms. Electrophysiological data from cyclic movement tasks in primates indicates that most of the behavioral variability in the motor system can be explained by variability in the neural activation patterns of the motor and pre-motor cortex (Churchland, Afshar, & Shenoy, 2006). Comparing the dynamics of our model to neurophysiological data could supply further insight into the mechanisms used by the brain to generate movement by linking the activation variables used in the functional description of our model to the measured activation patterns. Even a failure to establish such a link would be valuable by indicating where and how the model breaks down.

4. Movement generation for autonomous robots

4.1. Introduction

We have established a formal description of tasks as manifolds in the body configuration space and described behavioral signatures of selective stabilization principles relative to these manifolds in Chapter 2. A possible mechanism of how the CNS might organize the available degrees of freedom in a way that generates these signatures in a postural task was presented in Chapter 3. But why does the CNS achieve this coordination? What is the functional benefit of this task-related redundancy resolution?

Inquiries about the functional role of specific elements in a system are generally hard to answer from a natural science perspective. To explore the benefits that a movement generation system might receive from applying principles of selective stabilization, we change our angle of approach from describing and analyzing a system to the engineer's point of view as a designer of a system. In this chapter, we will describe a movement generation system for autonomous robotic agents based on principles of selective stabilization relative to task-manifolds. The engineering approach allows us to methodically list the capabilities we expect from a movement generation system as formal constraints and goals and explore the functional benefits of applying selective stabilization principles to fulfill these goals.

4.1.1. Constrained movement trajectories

Achieving a movement goal means leaving the current body configuration and finding a new state that fulfills the task at hand. A movement trajectory between the current state and such a goal state has to be selected out of the infinite number of possible trajectories. On one hand, this selection process has a large number of degrees of freedom: in most situations, the exact movement direction or speed do not matter and small local variations will not prevent fulfilling the global movement goal. On the other hand, there are clear constraints that have to be observed both during the movement and at movement termination, like avoiding collision with obstacles or orienting the hand in a way that allows grasping the target in a manipulation task. How can the redundancy of trajectories be resolved without violating these constraints?

We present an approach to this problem for autonomous robotic agents based

on principles of selective stabilization. In the postural case, the current configuration is always close to a desired state and a feedback law that selectively counters task-relevant deviations is sufficient for stabilization, as shown in Chapter 3. In a movement task, the current configuration is by definition far away from achieving the goal, though. In simple cases, a similar feedback term on the task level can be sufficient to achieve the movement goal. But how can we ensure that the resulting trajectory does not violate additional constraints like obstacle avoidance?

One answer to this question is provided by the Artificial Potential Field Approach (APFA) for robotic manipulators proposed by Khatib (1986). To move the end-effector of a manipulator to a desired location, an artificial potential field is constructed in the work space of the robot with a minimum at the target location. The gradient of this potential is then applied to the end-effector as a pseudo-force. A viscous term ensures that the movement stops when the target is reached. This is essentially a PD-controller for the end-effector position, reducing the error between current and desired location to zero. Collision constraints are introduced by constructing additional potentials with a maximum or pole at the obstacle location, so that the gradient points away from the obstacle. The superposition of these potentials should then generate movements that arrive at the target without colliding with obstacles on the way.

The movement trajectories generated by the APFA are not unproblematic. In the presence of multiple obstacles, the superposed potential fields often contain minima at locations other than the desired target, called “spurious minima”. If the trajectory enters the basin of attraction of such a minimum, the movement terminates without achieving the behavioral goal (McLean & Cameron, 1995). Even if the target is successfully reached, the trajectories generated by descending the gradient of the potential often have unattractive characteristics like unnecessary curves near obstacles (Fajen, Warren, Temizer, & Kaelbling, 2003). There are analytic solutions to these problems that construct the potentials in a way that avoids spurious minima, but they are technically demanding and require complete knowledge of the whole workspace configuration (Rimon & Koditschek, 1992).

We postulate that the source of these problems is that the APFA applies what is essentially a stabilization principle to position-related variables. We propose that during movement generation, the state of movement-related variables is more important: when reaching for a target, controlling the location of my hand matters less than controlling whether it is actually moving towards the target. In this chapter, we propose a movement generation scheme that applies stabilization principles to *movement*-related variables. The results of this approach have been published elsewhere before (Reimann, Iossifidis, & Schöner, 2010a, 2010b, 2011). Here we give a formal treatment of the underlying principles and how they relate to the methods. The results are summarized and presented in an integrated and concise fashion.

4.1.2. Selective stabilization of movement parameters

Any movement goal can be broken up in a number of different sub-tasks, e.g. bringing the hand to a target, orienting it appropriately for a grasp and avoiding collisions with obstacles. For each sub-task, we choose a relevant task-variable and formulate the sub-task as a desired value of that variable. We distinguish between a *movement* phase, during which the variable is far away from this desired state, and a *postural* phase, during which the task is essentially fulfilled and the controller has to stabilize against perturbations driving the task variable out of this desired state. During this postural state, we will apply feedback control similar to the APFA and the model of neural dynamics in the brain described in Section 3.2.6.

In the movement phase, we choose *another* variable that describes how the task-variable changes relative to its desired value. This can simply be the rate of change of the position-level task variable, or something more complex like the direction of the hand movement relative to the target location. The desired state for this movement-level variable is one in which the position variable moves towards the target. This can mean simply ensuring that the change of a currently positive variable with desired value zero is negative. Or that the angle between the hand movement direction and the direction towards the target is zero. Having defined a movement-level variable and a desired state for it, we can once more apply a feedback controller to reduce the deviation from the desired state to zero. This is done by defining a vector field over the movement variable with an attractor at the desired state. When the movement-level variable is in the desired state, the position-level variable is *moving towards* its own desired state. Once it reaches it, we switch to the position-level postural controller. This selective stabilization of movement states that contribute to the completion of behavioral goals resolves the trajectory redundancy. Each task contribution will usually results in the movement state of that task variable being stabilized at the desired value.

In some cases a state that is beneficial to the completion of one task is detrimental for another task. Consider reaching for a cup of coffee that is standing behind a bottle of milk: the movement state that is desired for reaching the cup conflicts with the desired state for avoiding collision with the bottle, because the direction towards the cup and towards the bottle are identical. One contribution pulls the movement state towards that direction, the other contribution pushes away from it. If the relative strength of the two contributions is chosen appropriately, the hand will first turn sideways, away from that direction. But after it moved some distance to the side, the direction towards the cup and towards the bottle will not be the same anymore, and the desired value of the reaching task, “moving towards the cup”, can be achieved simultaneously with the desired state of the collision avoidance task, “not moving towards the bottle”. The result is a trajectory that gently curves around the bottle and, once having cleared that obstacle, moves directly towards the cup.

The redundancy of trajectories is resolved by stabilizing low-dimensional movement variables. The stabilizing vector fields have to be transformed into the high-dimensional configuration space. We do this by identifying each sub-task with a task

manifold (see Chapter 2) and constructing a vector field in configuration space for which the task manifold is asymptotically stable. This presents another redundant problem, as there are infinitely many vector fields in joint space vectors that solve this problem. This redundancy of solutions is resolved in the same way as in Chapter 3 by using a generalized inverse of the task Jacobian. Khatib did essentially the same on a level of kinetics instead of kinematics by expressing constraints as forces in the “operator space” and transforming them into joint space using the *dynamically consistent pseudo-inverse* (Khatib, 1995).

4.1.3. Related work

Classical manipulators used to be designed in a way that largely avoided redundancy of solutions by having exactly as many degrees of freedom as necessary, which left a finite number of solutions for a specific task, often only one. This solution was calculated analytically by solving the inverse kinematic problem, yielding e.g. a set of joint angles for which the end-effector was at a desired location. Trajectories can then be generated by moving to this desired location in a straight line in joint space (Siciliano & Khatib, 2008). The inverse kinematics can also be solved locally by inverting the task Jacobian, yielding a movement direction in joint space that moves the task variable closer to the desired state. Using a pseudo-inverse, this technique can be generalized to redundant problems (Whitney, 1972). Secondary problems can be solved using the redundant degrees of freedom in the null-space of the primary task solution, like avoiding obstacles with the proximal manipulator segments while following a trajectory with the end-effector (Maciejewski & Klein, 1985).

The Artificial Potential Field Approach was among the first approaches that resolved the redundancy of trajectories implicitly using a gradient descent. This facilitated the inclusion of dynamic constraints like obstacle avoidance in the generation of the end-effector trajectory, because the specifics of how the obstacle is avoided did not have to be made explicit but emerged from a comparatively simple superposition of vector fields (Khatib, 1986). Although it was initially presented for manipulator control, the APFA was more widely applied to vehicle path planning (Arkin, 1989). The problems of the APFA with spurious attractors have been recognized and addressed by the robotics community (Rimon & Koditschek, 1992; McLean & Cameron, 1995). Still, the APFA is not applied widely outside of laboratory environments.

The attractor dynamics approach was also initially designed for autonomous vehicle path planning. Having less degrees of freedom, movement generation for vehicles is more suitable for analytic treatment, enabling a detailed analysis of the system bifurcations between dynamic regimes and how they interact with behavior generation (Schöner, Dose, & Engels, 1995). A low-level version connected input from active infra-red sensors and passive photo-resistors to the motor output for the wheels, the only intermediate step was the calculation of the vector field contribution from the current sensor state (Bicho & Schöner, 1997). The first extension of the approach to redundant manipulators was restricted to end-effector motion, combining it with a closed-form solution of the inverse kinematics problem (Iossifidis & Schöner, 2006).

A different perspective on trajectory redundancy was introduced by the extensive use of machine learning techniques in robotics. Instead of selecting trajectories according to some design principle, the movements are demonstrated by human teachers and then imitated (Schaal, 1997). One successful method are the Dynamic Movement Primitives (DMP, Ijspeert, Nakanishi, & Schaal, 2002). Observed movement trajectories are represented by a mixture of Gaussians relative to the target location and an implicit timing variable. This allows to generalize the trajectory locally by small variations in the target location. Larger regions can be covered by sampling the space with movements to different target locations and combining neighboring samples for movements to new targets in between the learned ones. This approach has been successfully applied to a wide range of robotic problems from biped walking (Nakanishi et al., 2004) to table tennis (Mülling, Kober, Kroemer, & Peters, 2012). Because DMP essentially realize pre-planned trajectories, attempts to combine them with additional constraints like obstacle avoidance have had only limited success (Stulp, Oztop, Pastor, Beetz, & Schaal, 2009; Hoffmann, Pastor, Park, & Schaal, 2009).

Another approach to trajectory generation using dynamical systems is proposed by Billard and colleagues. The robot generalizes movements demonstrated by a human teacher using statistical methods like Gaussian Process Regression or Gaussian Mixture Regression and constructs vector fields that can be used to imitate the demonstrated trajectories (Khansari-Zadeh, Kronander, & Billard, 2012). The authors emphasize that in scenarios with robot-human interaction, the movement generation system must be highly reactive to unpredicted perturbations and adapt to both temporal and spatial changes of the desired movement parameters to ensure fulfilling the movement goal and prevent harmful collisions with objects and humans. Temporal adaptation can be achieved by coupling between the different dynamical systems, either internally to guarantee that two processes like reaching and grasping finish at the same time, or externally, e.g. to catch a moving object (Shukla & Billard, 2012). Obstacles, on the other hand, are accounted for by spatial modification of the vector fields in a way that the flow moves around a convex obstacle region instead of through it (Khansari-Zadeh & Billard, 2012).

4.2. The movement generation system

We construct a movement out of several sub-tasks that can be activated or ignored depending upon the current state and behavioral goals. The behavioral sub-tasks treated here are reaching for a target (Target acquisition, Section 4.2.1), avoiding collision with obstacles (Obstacle avoidance, Section 4.2.2), orienting the hand appropriately for grasping an object (Hand orientation, Section 4.2.3) and avoiding the physical limits of the joint actuators (Joint angle limits, Section 4.2.4). In addition to these behavior-oriented sub-tasks, we include one sub-task that dampens out residual velocities, a side-effect of second order approaches (Homogeneous damping, Section 4.2.5).

The configuration of a generic robotic manipulator is described by a set of joint angles

$$\theta = (\theta_1, \dots, \theta_n) \in \mathbb{R}^n.$$

the movement state is given by the joint velocities

$$\dot{\theta} = (\dot{\theta}_1, \dots, \dot{\theta}_n) \in \mathbb{R}^n.$$

For our present purposes, we can assume that the configuration and movement states describe the complete kinematic state of the manipulator and all variables that are relevant for movement generation are functions of $(\theta, \dot{\theta})$. The robot is controlled by passing a vector of joint accelerations $\ddot{\theta}$ to the hardware interface.

4.2.1. Target acquisition

Similar to a vehicle steering towards a goal, the target acquisition is pursued by changing the heading direction of the end-effector towards the direction of the target. Ideally, the movement direction and the direction toward the target are identical. The behaviorally relevant variable is thus the angular difference between these two directions.

Let $\mathbf{p} = \mathbf{p}(\theta) \in \mathbb{R}^3$ be the position of the end-effector. The movement state of \mathbf{p} is given by its velocity $\mathbf{v} = \dot{\mathbf{p}} = J_{\mathbf{p}}\dot{\theta}$, where $J_{\mathbf{p}} = \frac{d\mathbf{p}}{d\theta}$ is the Jacobian matrix of \mathbf{p} . Let $\mathbf{g} \in \mathbb{R}^3$ be the location of the target and $\mathbf{k} = \mathbf{g} - \mathbf{p}$ the vector from \mathbf{p} to the target. We define the *target deviation angle* as

$$\phi = \phi(\theta, \dot{\theta}, \mathbf{g}) = \arccos\left(\frac{\mathbf{v}^T \mathbf{k}}{|\mathbf{v}| |\mathbf{k}|}\right). \quad (4.1)$$

Throughout this section, we will assume that both $\mathbf{v} \neq 0$ and $\phi \neq 0$, i.e. the end-effector is moving, but not directly towards the target. These trivial cases will be treated separately.

The behavioral variable ϕ expresses to what degree the movement state is beneficial for reaching the target. The desired value for it is $\phi = 0$, implying that the end-effector moves directly towards the target. We define a vector field $f_{\mathbf{p}, \mathbf{g}}^{(\text{dev})}$ over the deviation angle ϕ by setting

$$f_{\mathbf{p}, \mathbf{g}}^{(\text{dev})} = -\alpha_{\phi} \sin \phi, \quad (4.2)$$

where α_{ϕ} is a gain factor. We will usually drop the subscripts \mathbf{p}, \mathbf{g} where no confusion can arise about which points are indicated. This vector field has a single attractor at the desired state $\phi = 0$, meaning the flow of $f^{(\text{dev})}$ turns the movement of the hand towards the target direction.

To realize this flow for the dependent variable ϕ , we need to find a flow for the state variables $\theta, \dot{\theta}$ that implies $\dot{\phi} = f^{(\text{dev})}$. This means resolving the redundancy of solutions by selecting one out of the infinite set of joint velocity vectors $\dot{\theta}$ for which this equation is satisfied. To achieve this transformation from a low-dimensional flow of ϕ into a high-dimensional flow of $\dot{\theta}$, we make an intermediate step to Cartesian

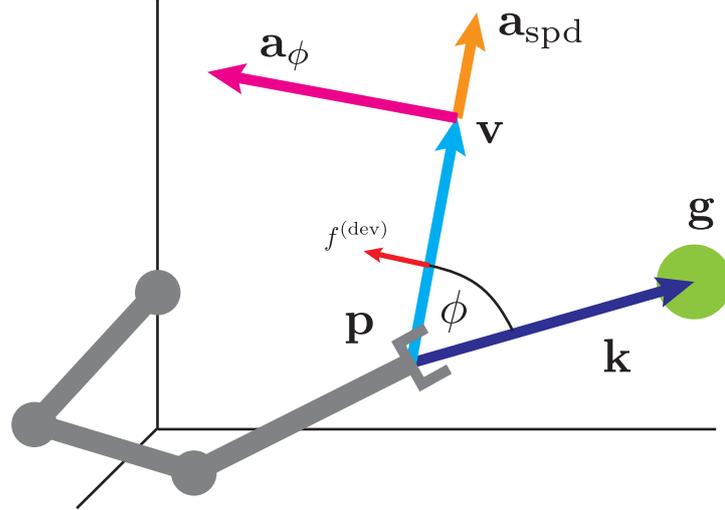


Figure 4.1.: The target deviation angle ϕ measures how directly the end-effector \mathbf{p} is moving towards the target \mathbf{g} . A desired change $f^{(\text{dev})}$ of ϕ can be realized by accelerating the end-effector in direction \mathbf{a}_ϕ . Accelerating the end-effector towards \mathbf{a}_s changes the movement speed of the end-effector.

workspace. For any given movement state, define the Cartesian acceleration direction \mathbf{a}_ϕ for as

$$\mathbf{a}_\phi = e^{\hat{\omega} \frac{\pi}{2}} \mathbf{v}, \quad \text{with } \omega = \frac{\mathbf{k} \times \mathbf{v}}{|\mathbf{k} \times \mathbf{v}|}, \quad (4.3)$$

where the exponential function $e^{\hat{\omega}\theta}$ of the skew symmetric matrix

$$\hat{\omega} = \begin{pmatrix} 0 & -\omega_3 & \omega_2 \\ \omega_3 & 0 & -\omega_1 \\ -\omega_2 & \omega_1 & 0 \end{pmatrix} \quad (4.4)$$

with the factor θ is a rotation matrix around the axis given by ω by the angle θ (Murray, Li, & Sastry, 1994). The Cartesian acceleration vector \mathbf{a}_ϕ is perpendicular to the velocity vector \mathbf{v} and lies in the plane spanned by \mathbf{v} and \mathbf{k} . It solves the first part of the solution redundancy problem by selecting one acceleration for the Cartesian end-effector position out of the infinite number that results in a desired deviation angle change $\dot{\phi}$. Figure 4.1 illustrates the relationship between changes in ϕ and accelerations of \mathbf{p} in direction \mathbf{a}_ϕ . The following Theorem 1 shows the relationship between the acceleration direction \mathbf{a}_ϕ and changes of the deviation angle ϕ . Subsequently, we show the optimality of this solution in the sense of having minimal norm in Theorem 2.

Theorem 1. Let $f^{(\text{dev})} \in \mathbb{R}$ be any desired velocity of ϕ . Then

$$\dot{\mathbf{v}} = \mathbf{a}_\phi (f^{(\text{dev})} - \bar{\phi}) \implies \dot{\phi} = f^{(\text{dev})} \quad (4.5)$$

where

$$\bar{\phi} = \frac{-1}{\sqrt{1 - \cos^2 \phi}} \frac{\mathbf{v}^T \dot{\mathbf{k}} |\mathbf{v}| |\mathbf{k}| - \mathbf{v}^T \mathbf{k} |\mathbf{v}| \frac{\mathbf{k}^T \dot{\mathbf{k}}}{|\mathbf{k}|}}{|\mathbf{v}|^2 |\mathbf{k}|^2}. \quad (4.6)$$

Proof. For ease of notation, we choose a basis such that $\mathbf{k} = \begin{pmatrix} k_1 \\ 0 \\ 0 \end{pmatrix}$, $\mathbf{v} = \begin{pmatrix} v_1 \\ v_2 \\ 0 \end{pmatrix}$,
 $\omega = \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}$, with $k_1, v_2 \geq 0$. In these coordinates,

$$\mathbf{a}_\phi = e^{\omega \frac{\pi}{2}} \mathbf{v} = \begin{pmatrix} 0 & -1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} v_1 \\ v_2 \\ 0 \end{pmatrix} = \begin{pmatrix} -v_2 \\ v_1 \\ 0 \end{pmatrix}. \quad (4.7)$$

Now we calculate the temporal derivative

$$\dot{\phi} = \frac{d\phi}{dt} = \frac{d\phi}{d \cos \phi} \frac{d \cos \phi}{dt} = \frac{d\phi}{d \cos \phi} \frac{d \frac{\mathbf{v}^T \mathbf{k}}{|\mathbf{v}| |\mathbf{k}|}}{dt} \quad (4.8)$$

$$= \frac{-1}{\sqrt{1 - \cos^2 \phi}} \frac{(\dot{\mathbf{v}}^T \mathbf{k} + \mathbf{v}^T \dot{\mathbf{k}}) |\mathbf{v}| |\mathbf{k}| - \mathbf{v}^T \mathbf{k} \left(\frac{d|\mathbf{v}|}{dt} |\mathbf{k}| + |\mathbf{v}| \frac{d|\mathbf{k}|}{dt} \right)}{|\mathbf{v}|^2 |\mathbf{k}|^2} \quad (4.9)$$

$$= \frac{-1}{\sqrt{1 - \frac{(\mathbf{v}^T \mathbf{k})^2}{|\mathbf{v}|^2 |\mathbf{k}|^2}}} \left(\frac{\dot{\mathbf{v}}^T \mathbf{k} |\mathbf{v}| |\mathbf{k}| - \mathbf{v}^T \mathbf{k} \overbrace{\frac{\mathbf{v}^T \dot{\mathbf{v}}}{|\mathbf{v}|} |\mathbf{k}|}^{=0}}{|\mathbf{v}|^2 |\mathbf{k}|^2} + \frac{\mathbf{v}^T \dot{\mathbf{k}} |\mathbf{v}| |\mathbf{k}| - \mathbf{v}^T \mathbf{k} |\mathbf{v}| \frac{\mathbf{k}^T \dot{\mathbf{k}}}{|\mathbf{k}|}}{|\mathbf{v}|^2 |\mathbf{k}|^2} \right) \quad (4.10)$$

$$= \frac{-|\mathbf{v}| |\mathbf{k}|}{\sqrt{|\mathbf{v}|^2 |\mathbf{k}|^2 - (\mathbf{v}^T \mathbf{k})^2}} \frac{\dot{\mathbf{v}}^T \mathbf{k}}{|\mathbf{v}| |\mathbf{k}|} + \bar{\phi} \quad (4.11)$$

$$= \frac{-\dot{\mathbf{v}}^T \mathbf{k}}{\sqrt{|\mathbf{v}|^2 |\mathbf{k}|^2 - (\mathbf{v}^T \mathbf{k})^2}} + \bar{\phi} \quad (4.12)$$

$$= \frac{-(f^{(\text{dev})} - \bar{\phi}) \mathbf{a}_\phi^T \mathbf{k}}{\sqrt{|\mathbf{v}|^2 |\mathbf{k}|^2 - (\mathbf{v}^T \mathbf{k})^2}} + \bar{\phi} \quad (4.13)$$

$$= \frac{-(f^{(\text{dev})} - \bar{\phi})(-v_2 k_1)}{\sqrt{(v_1^2 + v_2^2) k_1^2 - v_1^2 k_1^2}} + \bar{\phi} \quad (4.14)$$

$$= \frac{(f^{(\text{dev})} - \bar{\phi})(v_2 k_1)}{\sqrt{v_2^2 k_1^2}} + \bar{\phi} \quad (4.15)$$

$$= f^{(\text{dev})} - \bar{\phi} + \bar{\phi} \quad (4.16)$$

$$= f^{(\text{dev})} \quad (4.17)$$

□

Theorem 2. Let $f^{(\text{dev})} \in \mathbb{R}$ be any desired velocity of ϕ . Let $\mathbf{a} \in \mathbb{R}^3$ be any vector that satisfies the implication

$$\dot{\mathbf{v}} = \mathbf{a}(f^{(\text{dev})} - \bar{\phi}) \implies \dot{\phi} = f^{(\text{dev})} \quad (4.18)$$

and is orthogonal to \mathbf{v} . Then

$$|\mathbf{a}_\phi| \leq |\mathbf{a}|. \quad (4.19)$$

Proof. We express \mathbf{a} as the solution we already know plus some difference vector

$$\mathbf{a} = \mathbf{a}_\phi + \tilde{\mathbf{a}} \quad (4.20)$$

and then show that the difference vector $\tilde{\mathbf{a}}$ necessarily increases the norm. Assuming without loss of generality that $\dot{\mathbf{v}} = \mathbf{a}(f^{(\text{dev})} - \bar{\phi})$, we know from line 4.12 in the proof of Theorem 1 above that

$$\dot{\phi} = \frac{-\dot{\mathbf{v}}^T \mathbf{k}}{\sqrt{|\mathbf{v}|^2 |\mathbf{k}|^2 - (\mathbf{v}^T \mathbf{k})^2}} + \bar{\phi} \quad (4.21)$$

$$= \frac{-(f^{(\text{dev})} - \bar{\phi}) \mathbf{a}^T \mathbf{k}}{\sqrt{|\mathbf{v}|^2 |\mathbf{k}|^2 - (\mathbf{v}^T \mathbf{k})^2}} + \bar{\phi} \quad (4.22)$$

$$\stackrel{(4.20)}{=} \underbrace{\frac{-(f^{(\text{dev})} - \bar{\phi}) \mathbf{a}_\phi^T \mathbf{k}}{\sqrt{|\mathbf{v}|^2 |\mathbf{k}|^2 - (\mathbf{v}^T \mathbf{k})^2}}}_{= f^{(\text{dev})}} + \bar{\phi} + \frac{-(f^{(\text{dev})} - \bar{\phi}) \tilde{\mathbf{a}}^T \mathbf{k}}{\sqrt{|\mathbf{v}|^2 |\mathbf{k}|^2 - (\mathbf{v}^T \mathbf{k})^2}} \quad (4.23)$$

$$\stackrel{(4.13)}{=} f^{(\text{dev})} - \frac{(f^{(\text{dev})} - \bar{\phi}) \tilde{\mathbf{a}}^T \mathbf{k}}{\sqrt{|\mathbf{v}|^2 |\mathbf{k}|^2 - (\mathbf{v}^T \mathbf{k})^2}} \quad (4.24)$$

As \mathbf{a} satisfies (4.18) and $f^{(\text{dev})}$ is arbitrary so in general $f^{(\text{dev})} \neq \bar{\phi}$, this equation implies that $\tilde{\mathbf{a}}^T \mathbf{k} = 0$. In the basis of Theorem 2, we get

$$\tilde{\mathbf{a}}^T \mathbf{k} = \tilde{a}_1 k_1 = 0 \implies \tilde{a}_1 = 0, \quad (4.25)$$

and from the orthogonality of $\tilde{\mathbf{a}}$ and \mathbf{v}

$$\tilde{\mathbf{a}}^T \mathbf{v} = \tilde{a}_2 v_2 = 0 \implies \tilde{a}_2 = 0, \quad (4.26)$$

so we know that

$$\mathbf{a} = \mathbf{a}_\phi + \tilde{\mathbf{a}} \stackrel{(4.7)}{=} \begin{pmatrix} -v_2 \\ v_1 \\ 0 \end{pmatrix} + \begin{pmatrix} 0 \\ 0 \\ \tilde{a}_3 \end{pmatrix} = \begin{pmatrix} -v_2 \\ v_1 \\ \tilde{a}_3 \end{pmatrix} \quad (4.27)$$

This yields the inequality

$$|\mathbf{a}| = v_2^2 + v_1^2 + \tilde{a}_3^2 = |\mathbf{a}_\phi| + \tilde{a}_3^2 \leq |\mathbf{a}_\phi| \quad (4.28)$$

we were looking for. \square

The second part of the solution redundancy is to find a vector of joint accelerations $\ddot{\theta}$ that realize a desired end-effector acceleration. As before in Chapter 3, this is done using the Moore-Penrose pseudo-inverse of the end-effector Jacobian $J_{\mathbf{p}} = \frac{d\mathbf{p}}{d\theta}$. Analogous to the first step, this is optimal in the sense that the resulting solution has minimal norm (Greville, 1959). The complete transformation from a low-dimensional flow of ϕ into a high-dimensional flow of $\dot{\theta}$ is described in the following theorem.

Theorem 3. *Let $f^{(\text{dev})} \in \mathbb{R}$ be any desired velocity of ϕ . Then*

$$\ddot{\theta} = J_{\mathbf{p}}^+ \left(\mathbf{a}_{\phi}(f^{(\text{dev})} - \bar{\phi}) - \dot{J}_{\mathbf{p}}\dot{\theta} \right) \implies \dot{\phi} = f^{(\text{dev})} \quad (4.29)$$

Proof. The product rule yields

$$\dot{\mathbf{v}} = J_{\mathbf{p}}\ddot{\theta} + \dot{J}_{\mathbf{p}}\dot{\theta}, \quad (4.30)$$

which together with Equation 4.29 implies that

$$\dot{\mathbf{v}} = \mathbf{a}_{\phi}(f^{(\text{dev})} - \bar{\phi}). \quad (4.31)$$

This satisfies the precondition of Theorem 2, completing the proof. \square

Calculating the temporal derivative of the task Jacobian used in 4.30 in practice is not trivial. A formula for how this can be done using screw calculus is given in Appendix D.

Theorem 3 provides a mapping from desired changes $\dot{\phi}$ of the low-dimensional task variable to changes $\ddot{\theta}$ of the high-dimensional state variables. This allows us to transform the vector field $f_{\mathbf{p},\mathbf{g}}^{(\text{dev})}$ over the deviation angle into a vector field $F_{\mathbf{p},\mathbf{g}}^{(\text{dev})}$ over the joint velocities

$$F_{\mathbf{p},\mathbf{g}}^{(\text{dev})} = J_{\mathbf{p}}^+ \left(\mathbf{a}_{\phi}(f^{(\text{dev})} - \bar{\phi}) - \dot{J}_{\mathbf{p}}\dot{\theta} \right) \in \mathbb{R}^n. \quad (4.32)$$

Choosing $f^{(\text{dev})}$ as given in Equation 4.2, the flow of this vector field results in the end-effector movement direction relaxing towards the target direction. Theorem 3 directly implies the task manifold for $\phi = 0$ is asymptotically stable in the joint velocity space.

4.2.1.1. Movement speed

The direction is just one aspect of the movement state. To get the end-effector moving to reach the target location, we also need to control the speed $s = |\mathbf{v}|$ of the hand. For the present purpose, we do this in a very simple manner, by arbitrarily setting a reasonable desired value s_{des} and letting the speed relax towards it under the flow of the vector field

$$f_{\mathbf{p}}^{(\text{spd})} = -\alpha_s(s - s_{\text{des}}) \quad (4.33)$$

over the low-dimensional behavioral variable s with a single attractor at s_{des} . The transformation into a vector field over joint velocities is given by

$$F_{\mathbf{p}}^{(\text{spd})} = J_{\mathbf{p}}^+ \left(\mathbf{a}_s f_{\mathbf{p}}^{(\text{spd})} - \dot{J}_{\mathbf{p}} \dot{\theta} \right) \in \mathbb{R}^n \quad (4.34)$$

where the acceleration direction is just the normed velocity $\mathbf{a}_s = \frac{\mathbf{v}}{|\mathbf{v}|}$ (see Figure 4.1). Analogous to Theorem 3, we get

$$\ddot{\theta} = F_{\mathbf{p}}^{(\text{spd})} \implies \dot{s} = \frac{d|\mathbf{v}|}{dt} = \frac{\mathbf{v}^T \dot{\mathbf{v}}}{|\mathbf{v}|} = \frac{\mathbf{v}^T \frac{\mathbf{v}}{|\mathbf{v}|}}{|\mathbf{v}|} f_{\mathbf{p}}^{(\text{spd})} = f_{\mathbf{p}}^{(\text{spd})}, \quad (4.35)$$

i.e. the flow of the vector field $F_{\mathbf{p}}^{(\text{spd})}$ over the joint velocities induces the desired flow of $f_{\mathbf{p}}^{(\text{spd})}$ that lets $s = |\dot{\mathbf{p}}|$ relax towards s_{des} .

4.2.1.2. Stabilization near the target location

The deviation angle is a suitable task variable for generating movements towards a target and stabilizing the movement direction at a state beneficial for reaching the target. In close vicinity of the target, this choice of task variable is less appropriate. When the distance to the target $|\mathbf{k}|$ is small in relation to $|\mathbf{v}|$, small changes in the end-effector velocity \mathbf{v} potentially correspond to large changes in the deviation angle ϕ . That means that perturbations from other behavioral vector fields affecting the hand position can have a large destabilizing effect on ϕ . For small distances $d = |\mathbf{k}|$, we switch to stabilizing the postural state instead of the movement state.

As relevant variables for postural stabilization, we use the Cartesian position \mathbf{p} and velocity \mathbf{v} of the end-effector. The desired state is given by $\mathbf{p} = \mathbf{g}$ and $\mathbf{v} = 0$. A vector field with an attractor at this state is given by the damped harmonic oscillator

$$f_{\mathbf{p},\mathbf{g}}^{(\text{pos})} = -\alpha_{\mathbf{v}}(\mathbf{v} - \alpha_{\mathbf{p}}(\mathbf{g} - \mathbf{p})). \quad (4.36)$$

Once more, this is transformed to joint space by

$$F_{\mathbf{p},\mathbf{g}}^{(\text{pos})} = J_{\mathbf{p}}^+ \left(f_{\mathbf{p},\mathbf{g}}^{(\text{pos})} - \dot{J}_{\mathbf{p}} \dot{\theta} \right). \quad (4.37)$$

The switch between the angle-dependent vector field for directing the hand movement at long-range and the position and velocity dependent vector field for stabilizing the hand in the proximity of the target is done smoothly after the distance d falls below a certain threshold $d_1^{(\text{tar})}$. Once a second threshold $d_2^{(\text{tar})}$ is reached only the postural stabilization vector field is used. This is done by setting

$$F_{\mathbf{p},\mathbf{g}}^{(\text{tar})} = (1 - \sigma(d)) \left(F_{\mathbf{p},\mathbf{g}}^{(\text{dev})} + F_{\mathbf{p}}^{(\text{spd})} \right) + \sigma(d) F_{\mathbf{p},\mathbf{g}}^{(\text{pos})}, \quad (4.38)$$

where the sigmoid switching function $\sigma(d) = \sigma_{d_1^{(\text{tar})}, d_2^{(\text{tar})}}(d)$ is defined as

$$\sigma_{a,b}(x) = \begin{cases} 0 & : x \leq a, \\ -\frac{1}{2} \cos\left(\frac{x-a}{b-a}\pi\right) + \frac{1}{2} & : a < x < b, \\ 1 & : b \leq x. \end{cases} \quad (4.39)$$

In the special case where $\mathbf{v} = 0$, we initiate movement towards the target by setting $\mathbf{F}_{\mathbf{p}, \mathbf{g}}^{(\text{tar})} = -\alpha_p J_p^+ \mathbf{k}$. The singularity of the mapping between $f^{(\text{dev})}$ and $F^{(\text{dev})}$ at $\phi = 0$ can easily be filled by setting $F^{(\text{dev})} = 0$ in that case.

4.2.2. Obstacle avoidance

In order to prevent collision of manipulator link segments with obstacles in the scene, we change the movement vectors of link segments away from directions in which obstacles are positioned. To describe the vector fields that achieve this, we first define the direction in which the repelling force acts, and then its magnitude depending on the distance and current movement states.

For the present study, we assume that all link segments and obstacles are enclosed in spheres or cylindrical bounding volumes topped off with half-spheres at the ends. This provides an acceptable approximation of many different types of objects. The same approach also works with other geometric types, like the planar surface of a table, but to present the principles of the approach and analyze it, cylinders and spheres provide sufficient complexity.

The obstacle cylinders are always oriented towards the z -axis. For any link segment S and obstacle O , let \mathbf{s} and \mathbf{o} be the points on their respective bounding volumes with minimal distance to each other, and $\mathbf{v}_S = \dot{\mathbf{s}}$ the velocity vector of the segment point. This velocity vector $\mathbf{v}_S = \dot{\mathbf{s}}$, relative to the obstacle configuration, is the relevant variable for avoiding collision of the segment S with the obstacle O . We will construct the vector field that achieves this collision avoidance by first specifying the direction and then the magnitude of the vector field for a given scene configuration and movement state.

4.2.2.1. Avoidance direction

Analogous to what we did for target acquisition, we want to minimize the overall effect of the vector field on the movement, so we look for a vector that is perpendicular to the movement vector of the segment point $\dot{\mathbf{s}} = \mathbf{v}_S$ so that accelerating the segment point in this direction only changes the direction of the movement state of \mathbf{s} , not the speed. The vector we look for should also point away from the obstacle, but as the obstacle is a cylinder instead of a single point like the target, defining what “away” means is more complicated. We approach this problem by considering the normal plane of the velocity vector, representing all vectors that are perpendicular to \mathbf{v} . By projecting the shape of the obstacle to this plane, we can select one vector that “optimally points away” from the obstacle.

Let N be the plane that is normal to \mathbf{v}_S . An orthonormal base $\mathbf{u}_1, \mathbf{u}_2$ of N is given by

$$\mathbf{u}'_2 = \mathbf{e}_3 - \frac{\langle \mathbf{v}_S, \mathbf{e}_3 \rangle}{\langle \mathbf{v}_S, \mathbf{v}_S \rangle} \mathbf{v}_S, \quad \mathbf{u}'_1 = \mathbf{v}_S \times \mathbf{u}_2, \quad \mathbf{u}_i = \frac{\mathbf{u}'_i}{|\mathbf{u}'_i|},$$

in which the second base vector is the projection of \mathbf{e}_3 , the up-axis of the world coordinate frame, onto N .

Let \mathbf{a} and \mathbf{b} be the projections of the center points of the half-spheres on the top and bottom of the obstacle bounding volume to N . Due to obstacles always standing upright and the choice of the base $\mathbf{u}_1, \mathbf{u}_2$, the projected points \mathbf{a} and \mathbf{b} lie on a vertical line in N . If \mathbf{a} is the upper point, i.e. $\langle \mathbf{a}, \mathbf{u}_2 \rangle > \langle \mathbf{b}, \mathbf{u}_2 \rangle$, the point on that line segment with minimal distance to the origin of N is given by

$$\mathbf{q} = \begin{cases} \mathbf{a} & : \langle \mathbf{a}, \mathbf{u}_2 \rangle < 0, \\ \mathbf{b} & : \langle \mathbf{b}, \mathbf{u}_2 \rangle > 0, \\ \langle \mathbf{b}, \mathbf{u}_1 \rangle \mathbf{u}_1 & : \text{otherwise,} \end{cases} \quad (4.40)$$

and the distance of the obstacle bounding volume projection to the origin of N is

$$h = |\mathbf{q}| - r, \quad (4.41)$$

where r is the radius of the obstacle bounding volume. Figure 4.2 shows the plane N and how an obstacle is projected.

The “best” avoidance direction, i.e. the one changing the movement direction directly away from the obstacle, is given by $-\mathbf{q}$. For some situations, though, this is not a suitable choice: when the link segment is moving horizontally in a plane that intersects the cylinder, then \mathbf{q} will also lie in that plane, resulting in a horizontal change of link segment movement. If the link segment is near the base of the kinematic chain, this is unlikely to result in a movement in which the whole manipulator avoids the obstacle, because either the immobile base or the rest of the arm up to the end-effector are still on one side of the obstacle, while the link segment in question is trying to avoid it by moving around the other side.

To prevent this kind of deadlock, instead of choosing \mathbf{q} itself as avoidance direction, we use a vector that also lies in N , but is rotated towards \mathbf{u}_2 by an amount depending upon the location of the link segment along the kinematic chain. The angle of $-\mathbf{q}$ with \mathbf{u}_2 is given by

$$\nu = \arccos \left(\frac{\langle -\mathbf{q}, \mathbf{u}_2 \rangle}{|\mathbf{q}|} \right). \quad (4.42)$$

For the j -th link segment of kinematic chain with n joints, we decrease that angle by applying the ad-hoc function

$$\nu' = \frac{3j-2}{4n-2} \nu, \quad (4.43)$$

and set

$$\mathbf{a}_o = \begin{cases} -\sin \nu' \mathbf{u}_1 + \cos \nu' \mathbf{u}_2 & : \langle \mathbf{q}, \mathbf{u}_1 \rangle \geq 0, \\ +\sin \nu' \mathbf{u}_1 + \cos \nu' \mathbf{u}_2 & : \langle \mathbf{q}, \mathbf{u}_1 \rangle < 0, \end{cases} \quad (4.44)$$

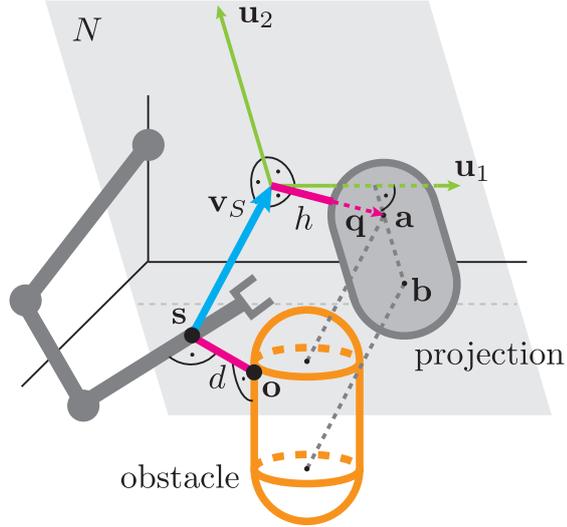


Figure 4.2.: Projection of an obstacle O to the plane N . N is normal to the velocity \mathbf{v}_S of the segment point s that has minimal distance d to the obstacle. N is spanned by the orthonormal base $\mathbf{u}_1, \mathbf{u}_2$, the projection of the obstacle center line segment is the line segment between \mathbf{a} and \mathbf{b} in N . The smallest distance of any point on this line segment to the projection of \mathbf{v}_S is given by \mathbf{q} , with h indicating the distance to the obstacle projection boundary.

which is $-\mathbf{q}$ rotated towards \mathbf{u}_2 so that $\angle(\mathbf{a}_o, \mathbf{u}_2) = \nu'$. Figure 4.3 illustrates this dependency of avoidance direction upon the link segment index in the chain.

Using this \mathbf{a}_o as avoidance direction is suitable for most situations. Only for the segment nearest to the base, this does not make sense, because it only has one single degree of freedom. For this segment only, we set the avoidance direction to

$$\mathbf{a}_o = -\frac{\mathbf{v}_S}{|\mathbf{v}_S|}, \quad (4.45)$$

essentially just braking the segment when it approaches an obstacle.

4.2.2.2. Magnitude of the repelling force

Defining the strength of the avoidance action for a given situation boils down to deciding how likely an obstacle collision is in that situation. Three factors play a role: the distance between the link segment and the obstacle, the movement direction of the link segment, and how fast it is moving in that direction. For each of these factors, we define a weight factor. The product of these weight factors will give the magnitude of the repelling force.

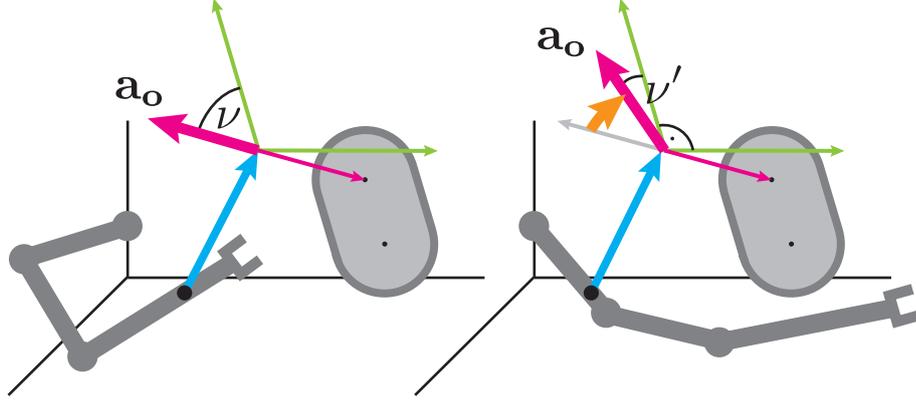


Figure 4.3.: The avoidance direction depends upon the link index. Links that are far away from the base can avoid an obstacle in the direction $\mathbf{a}_o = -\mathbf{q}$ that points directly away from the obstacle (left). For link segments close to the base, the avoidance direction \mathbf{a}_o is tilted upwards (right).

For the dependency on distance $d = |\mathbf{s} - \mathbf{o}|$, we define

$$w_d = (1 - \sigma_{d_1^{(\text{obs})}, d_2^{(\text{obs})}}(d)) \frac{d_1^{(\text{obs})}}{d}, \quad (4.46)$$

which is zero for distances $d > d_2^{(\text{obs})}$, 1 for $d = d_1^{(\text{obs})}$ and grows towards $+\infty$ for $d \rightarrow 0$.

For the movement direction dependency, let \mathbf{m} be the point on the line $\mathbf{s} + \lambda \mathbf{v}_s$, $\lambda \in \mathbb{R}$, that is closest to the obstacle, and \mathbf{o}_m the point on the obstacle bounding volume with minimal distance to \mathbf{m} . Define the obstacle angle as

$$\psi = \text{atan}_2(|\mathbf{m} - \mathbf{o}_m|, |\mathbf{m} - \mathbf{s}|), \quad (4.47)$$

which is the minimal angle between \mathbf{v}_s and any vector going through or touching the obstacle bounding volume. Set

$$w_\psi = \begin{cases} 0 & : \langle \mathbf{v}_s, \mathbf{o} - \mathbf{s} \rangle \leq 0, \\ 1 - \sigma_{\psi_1, \psi_2}(\psi) & : \text{otherwise,} \end{cases} \quad (4.48)$$

which vanishes if \mathbf{v}_s is zero or ψ is too large.

Finally, as a dependency on the movement speed we just take that value itself as

$$w_v = |\mathbf{v}_s|. \quad (4.49)$$

With these three weight factors depending on distance, movement direction and speed, we can define the magnitude of the repelling force as of the obstacle O on the

segment S as

$$f_{S,O}^{(\text{obs})} = \alpha_o \cdot w_d \cdot w_\psi \cdot w_v. \quad (4.50)$$

4.2.2.3. The obstacle vector field

Having defined an avoidance direction in \mathbb{R}^3 and a magnitude f_O^S , what remains is to define a corresponding vector field in joint space. Let

$$J_s = \left(\frac{\partial \mathbf{s}_i}{\partial \theta_j} \right)_{i,j} \quad (4.51)$$

be the Jacobian of the segment point \mathbf{s} . Then $J_s^{\mathbf{a}_o} = \mathbf{a}_o^T J_s$ gives the change along \mathbf{a}_o by changes of θ . We use the pseudo inverse of this to define

$$F_{O,S}^{(\text{obs})} = f_{O,S}^{(\text{obs})} \left(J_s^{\mathbf{a}_o} \right)^+, \quad (4.52)$$

which is a vector in joint space that realizes the desired change in direction \mathbf{a}_o with magnitude $f_{O,S}^{(\text{obs})}$.

For the complete vector field of one obstacle, we just sum up the vector fields over all obstacles O and link segments S , getting

$$F^{(\text{obs})} = \sum_{O,S} F_{O,S}^{(\text{obs})}. \quad (4.53)$$

4.2.3. Hand orientation

Manipulating an object imposes certain constraints upon the orientation of the hand. To formulate these constraints, one usually chooses coordinates for the space of hand orientations and specifies desired orientations as vectors in these coordinates. Topologically, the space of hand orientations $SO(3)$ is isomorphic to the real projective space \mathbb{RP}^3 (Murray et al., 1994). This means that there exists no global coordinate system that describes all hand orientations, so for each orientation-related task, a local coordinate frame must be chosen. The standard approach is to use systems of Euler angles, where the specific choice of angles can vary depending on the task (Siciliano & Khatib, 2008). Because of the different topologies, describing orientations with Euler angles always includes singularities. The approach of finding relevant, low-dimensional variables for a given behavioral goal and selectively stabilizing them allows us to bypass this problem. Instead of having to choose coordinates that completely describe the hand orientation, we only parameterize those aspects that are relevant for the given task.

For most tasks, some aspects of the hand orientation are important, while others are not. To lay a grasped pen down on a table, the pen should be parallel to the table, but the rotation of the hand in the table plane is irrelevant. Similarly, whether a liquid is poured from a bottle depends upon the inclination angle of the bottle,

but not the rotation of the bottle around its long axis. We capture this variety of possibly relevant aspects of hand orientation by a class of behavioral variables.

Definition 3. Let \mathbf{w} be any unit vector fixed to the end-effector, called **hand vector**. Let $\mathbf{k} = \mathbf{o} - \mathbf{p}$ be the vector from the hand position to the position \mathbf{o} of any relevant object in the scene, called **object vector**. The orientation angle is defined as

$$\gamma = \gamma_{\mathbf{p}, \mathbf{w}, \mathbf{o}} = \arccos \left(\frac{\mathbf{w}^T \mathbf{k}}{|\mathbf{k}|} \right). \quad (4.54)$$

For any given desired orientation angle γ_{des} , the vector field

$$-\alpha_\gamma(\dot{\gamma} - \beta_\gamma(\gamma_{\text{des}} - \gamma)), \quad (4.55)$$

is a damped harmonic oscillator with a single fixed point at $\gamma = \gamma_{\text{des}}, \dot{\gamma} = 0$ similar to what we used to stabilize the hand position (c. Equation 4.36). To avoid overly fast reorientations of the hand that might go along with large motions of the upper arm, we impose an upper bound on the absolute value of the instantaneous velocity attractor $\beta_\gamma(\gamma_{\text{des}} - \gamma)$, by choosing a limit ρ_{des} for the desired magnitude of the orientation change and setting

$$\rho = \begin{cases} -\rho_{\text{des}} & : & \beta_\gamma(\gamma_{\text{des}} - \gamma) \leq -\rho_{\text{des}}, \\ \beta_\gamma(\gamma_{\text{des}} - \gamma) & : & -\rho_{\text{des}} < \beta_\gamma(\gamma_{\text{des}} - \gamma) < \rho_{\text{des}}, \\ \rho_{\text{des}} & : & \rho_{\text{des}} \leq \beta_\gamma(\gamma_{\text{des}} - \gamma). \end{cases} \quad (4.56)$$

Then we define the vector field $f^{(\text{ori})}$ as

$$f^{(\text{ori})} = f_{\gamma, \gamma_{\text{des}}}^{(\text{ori})} = -\alpha_\gamma(\dot{\gamma} - \rho). \quad (4.57)$$

To transform this low-dimensional vector field over γ into a high-dimensional vector field over $\dot{\theta}$, we state a general relationship between vector fields over orientation angles and vector fields over joint angles, analogous to Theorem 3.

Theorem 4. Let $f^{(\text{ori})} \in \mathbb{R}$ be any desired acceleration of γ .

$$\begin{aligned} \ddot{\theta} &= J_\gamma^+ \left(f^{(\text{ori})} - \dot{J}_\gamma \dot{\theta} \right) \\ \implies \ddot{\gamma} &= f^{(\text{ori})}, \end{aligned} \quad (4.58)$$

where $J_\gamma = \frac{d\gamma}{d\theta}$ is the orientation angle Jacobian.

The proof of Theorem 4 is a direct implication of the product rule. The difficulty lies with calculating the Jacobian and its temporal derivative.

Lemma 1. The derivatives of the orientation angle with respect to \mathbf{w} and its time derivative are given by

$$\frac{d\gamma}{d\mathbf{w}} = \frac{-1}{\sqrt{1 - \cos^2 \gamma}} \frac{\mathbf{k}^T}{|\mathbf{k}|}, \quad (4.59)$$

$$\frac{d}{dt} \left(\frac{d\gamma}{d\mathbf{w}} \right) = \frac{\dot{\gamma} \sin \gamma \cos \gamma}{(1 - \cos^2 \gamma)^{\frac{3}{2}}} \frac{\mathbf{k}^T}{|\mathbf{k}|} + \frac{-1}{\sqrt{1 - \cos^2 \gamma}} \frac{\dot{\mathbf{k}}^T |\mathbf{k}| - \mathbf{k}^T \frac{\dot{\mathbf{k}}^T \mathbf{k}}{|\mathbf{k}|}}{|\mathbf{k}|^2}. \quad (4.60)$$

Proof. The first equation follows directly from the chain rule

$$\frac{d\gamma}{d\mathbf{w}} = \frac{d\gamma}{d \cos \gamma} \frac{d \cos \gamma}{d\mathbf{w}}.$$

The second equation also follows from straightforward application of derivation rules, using $\frac{d|\mathbf{k}|}{dt} = \frac{\mathbf{k}^T \dot{\mathbf{k}}}{|\mathbf{k}|}$. \square

We can now calculate the orientation angle Jacobian and its time derivative

$$J_\gamma = \frac{d\gamma}{d\mathbf{w}} \frac{d\mathbf{w}}{d\theta}, \quad \dot{J}_\gamma = \frac{d}{dt} \left(\frac{d\gamma}{d\mathbf{w}} \right) \frac{d\mathbf{w}}{d\theta} + \frac{d\gamma}{d\mathbf{w}} \frac{d}{dt} \left(\frac{d\mathbf{w}}{d\theta} \right). \quad (4.61)$$

We know the derivatives of γ by \mathbf{w} from lemma 1, calculating the other derivatives is straightforward using screw calculus (Murray et al., 1994).

We have used the Moore-Penrose pseudo-inverse to transform the low-dimensional vector field over γ into a high-dimensional vector field over θ . The resulting vector of joint accelerations has minimal norm, meaning that the accelerations are distributed over all joints as much as possible. For the generation of reaching movements, this solution is often impractical: While it is minimal in joint space, the effect on hand position can be quite large: to change the orientation of the hand frame z -axis, the hand would move in a curve instead of rotating on the spot. A small change in the orientation angle thus leads to a large perturbation in other relevant variables. This can be prevented by choosing a different transformation to joint space, with the added constraint that the hand position remains invariant. Define the augmented behavioral variable, vector field and Jacobian

$$\bar{\gamma} = \begin{pmatrix} \gamma \\ \mathbf{p} \end{pmatrix} \in \mathbb{R}^4, \quad \bar{f}^{(\text{ori})} = \begin{pmatrix} f^{(\text{ori})} \\ 0 \end{pmatrix} \in \mathbb{R}^4, \quad J_{\bar{\gamma}} = \begin{pmatrix} J_\gamma \\ J_{\mathbf{p}} \end{pmatrix} \in \mathbb{R}^{4 \times n}. \quad (4.62)$$

where the added $0 \in \mathbb{R}^3$ in the vector field $\bar{f}_\gamma^{(\text{ori})}$ corresponds to the desired acceleration of the hand position \mathbf{p} . The transformation to joint space of $f^{(\text{ori})}$ that does not affect the end-effector position \mathbf{p} is now given by

$$F^{(\text{ori})} = F_{\gamma, \gamma_{\text{des}}}^{(\text{ori})} = J_{\bar{\gamma}}^+ \begin{pmatrix} f_{\gamma, \gamma_{\text{des}}}^{(\text{ori})} - \dot{J}_\gamma \dot{\theta} \\ 0 \end{pmatrix}. \quad (4.63)$$

This is similar to an augmented Jacobian (e.g., Seraji, 1989), only that in our case the resulting task still does not have a unique solution and is transformed by the pseudo-inverse instead.

If choosing f_{ori} as defined in Equation 4.57, Theorem 4 implies that the task manifold for $\gamma = \gamma_{\text{des}}$ is asymptotically stable.

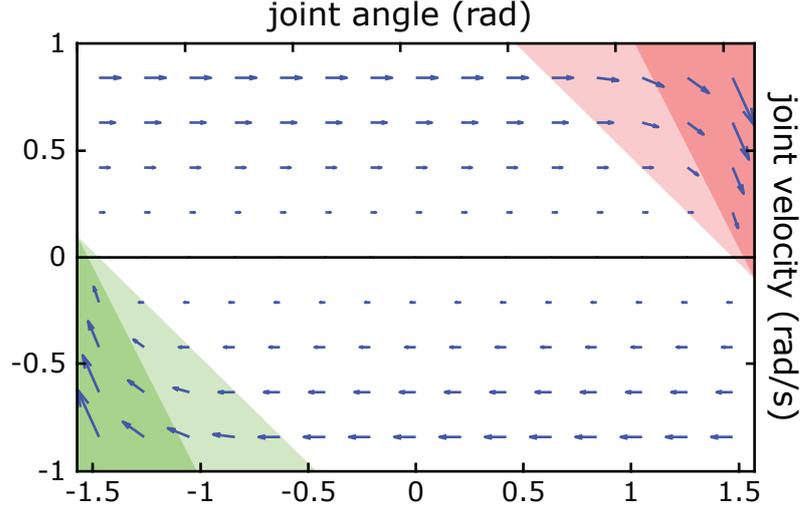


Figure 4.4.: Vector field for joint limit avoidance. The joint angle limits are at $\pm \frac{\pi}{2}$. Rate of change for θ is $\dot{\theta}$, rate of change for $\dot{\theta}$ is $f^{(\text{lim})} = f^{(\text{lower})} + f^{(\text{upper})}$ (see Equation 4.67). The upper joint limit vector field $f^{(\text{upper})}$ is fully active in the upper right region (dark red), the weight factor w_γ rises smoothly (light red region). The same areas are indicated by green for the lower joint limit vector field $f^{(\text{lower})}$.

4.2.4. Joint angle limits

Part of any movement generation system for robotic manipulators is accounting for the hardware limits imposed on the single joints. These are defined as the bounds of the interval $[\theta_i^{(\text{lower})}, \theta_i^{(\text{upper})}]$ of values the i -th joint can be moved to. When the interface between the movement generation system and the hardware is simply a vector of desired joint angles sent to the hardware, the joint limits can already be guaranteed by reducing or increasing values outside $[\theta_i^{(\text{lower})}, \theta_i^{(\text{upper})}]$ to the nearest value in the interval. If a joint limit is encountered, this will lead to a joint angle trajectory that sticks to the boundary of the workspace defined by the joint limit, resulting in a sharp change. Also, the relation between directions in joint space and changes of a behavioral variable is misrepresented by the task Jacobian, as the directions can point towards the joint limit, out of the allowed workspace.

To avoid the boundaries of the work space, we construct vector fields that repel from the joint angle limits. The relevant behavioral variables for this task are the distances

$$d^{(\text{upper})} = \theta^{(\text{upper})} - \theta, \quad d^{(\text{lower})} = \theta - \theta^{(\text{lower})} \quad (4.64)$$

from the joint angle limits.

One way to construct the repelling vector fields is to simply accelerate away from the joint limits when they are approached by setting

$$\tilde{f}_i^{(\text{upper})} = \alpha_{\text{lim}} \sigma(d^{(\text{upper})}, a_{\text{lim}}, b_{\text{lim}}), \quad (4.65)$$

where a_{lim} and b_{lim} are tolerance limits.

Defining $\tilde{f}_i^{(\text{lower})}$ analogously, the total joint limit avoidance vector field is then given by

$$\tilde{F}^{(\text{lim})} = \sum_i (\tilde{f}_i^{(\text{lower})} + \tilde{f}_i^{(\text{upper})}) \mathbf{e}_i, \quad (4.66)$$

where \mathbf{e}_i is the i -th unit vector.

This design suffers from two problems: if a joint limit is approached with a high velocity, the repelling contribution might not be sufficiently strong to reduce the movement to zero before the limit is reached. If on the other hand the joint limit is approached slowly, the magnitude of the repellation might be much higher than necessary. These two problems are really two faces of the same problem: determining how undesired the state of the variable is without taking into account its current rate of change. A value close to a joint limit can be tolerable as long as the distance from the limit is not reduced further, i.e. the joint only moves slowly towards the limit, or not at all. On the other hand a value far away from a joint limit can still be dangerous if the distance from the limit is rapidly shrinking.

To make a meaningful evaluation of how dangerous the current kinematic state $\theta_i, \dot{\theta}_i$ is in terms of hitting a joint limit that accounts for these effects, we consider the remaining time to contact with the limit, given by $\frac{\dot{\theta}}{d^{(\text{upper})}}$ and define a vector field using a hyperbolic function

$$f_i^{(\text{upper})} = -\alpha_{\text{lim}} w_d^{(\text{upper})} \frac{\dot{\theta} + c^{(\text{upper})}}{d^{(\text{upper})}} \quad (4.67)$$

that becomes large when the time to contact, adjusted by a safety offset $c^{(\text{upper})}$, is small. The factor

$$w_d^{(\text{upper})} = 1 - \sigma \left(d^{(\text{upper})}, \frac{\dot{\theta} + c^{(\text{upper})}}{2}, \dot{\theta} + c^{(\text{upper})} \right), \quad (4.68)$$

is a weight that vanishes when the time to contact becomes larger than 1 second, including cases where the joint is moving *away* from the limit. Figure 4.4 illustrates the resulting vector field for $\theta_i, \dot{\theta}_i$ and sketches the role of the single factors.

A vector $f_i^{(\text{lower})}$ is defined for the lower joint limit analogously. The complete joint limit avoidance contribution is defined as

$$F^{(\text{lim})} = \sum_i (f_i^{(\text{lower})} + f_i^{(\text{upper})}) \mathbf{e}_i. \quad (4.69)$$

4.2.5. Homogeneous damping

Each behavioral task erects a vector field over $\dot{\theta}$. The flow of each vector field contribution accelerates the joint angles in a direction that changes the low-dimensional variable relevant to that task in the desired way. The direction of the contribution in joint space changes over a movement, because of the non-linearities of the map-

ping between the state variables $(\theta, \dot{\theta})$ and the task-variable. Because of this change, movement states that comply with a task at one time might violate it at another.

One consequence of this non-linearity in second order approaches is the well-known problem of residual velocities building up in the null-space of a task (Hollerbach & Suh, 1987). As the task-relevant direction in joint spaces changes over a movement, velocity components in other directions come to lie in the null space of the task and are unaffected by the controlling vector fields. Over time, this can lead to substantial null-space velocities, called self-motion.

One way to counter this effect is to introduce a small homogeneous damping term

$$- \alpha_{\text{damp}} \dot{\theta} \quad (4.70)$$

that slowly reduces velocities in directions that do not affect any currently relevant behavioral variable. This homogeneous damping also restricts the freedom in directions that are currently relevant for behavioral goals, though. To avoid this, the damping term is only activated when the accuracy requirements on the behavioral variables are high, i.e. when the behavioral variables of target acquisition are near their desired values, as defined in Equation 4.37.

$$\mathbf{F}_{\text{damp}} = -\alpha_{\text{damp}} w^{(\text{tar})} w^{(\text{ori})} \dot{\theta}, \quad (4.71)$$

with

$$w^{(\text{tar})} = \left(1 - \sigma_{d_1^{(\text{tar})}, d_2^{(\text{tar})}}(d^{(\text{tar})})\right), \quad w^{(\text{ori})} = \left(1 - \sigma_{d_1^{(\text{ori})}, d_2^{(\text{ori})}}(d^{(\text{ori})})\right),$$

and $d^{(\text{tar})} = |\mathbf{g} - \mathbf{p}|$, $d^{(\text{ori})} = |\gamma - \gamma_{\text{des}}|$. This leaves maximal freedom by not reducing movement in any direction during transport, but increases the stability of the whole system by reducing movement in the null-space of the behavioral variables during stabilization.

4.2.6. Behavioral organization

As we want to demonstrate our approach not only with simple single movements, but also with more complex object manipulation tasks requiring a sequence of action, we now describe a simple system of behavioral organization that can accomplish this. We partition complex movements into sequences of behavioral phases, e.g. approaching an object or closing the hand to grasp it. A simple, neurally inspired dynamical system is used to switch between these phases.

Each behavioral phase is represented by one Amari neuron (Amari, 1977). Formally, this is a dynamical system

$$\dot{u} = -u + \alpha_u \sigma_{u_0}(u) + h + \text{input}, \quad (4.72)$$

where u is the dynamical variable, h a resting level and $\sigma = (1 + \exp(-\beta(u - u_0)))^{-1}$ a sigmoid function with threshold u_0 and slope β . The input comes from sensor information or activation of other neurons. With appropriate parametrization, this

dynamical system is bistable, with two attractors that can be interpreted as *on* and *off* (Schöner, 2008). Switching between these two stable states can be induced by short-term inputs that briefly destabilize one of the states.

Consider one such neuron B receiving moderate excitatory input from another neuron A , but not enough to switch it to the *on* state. Add another excitatory connection from a sensory channel to B , also of moderate strength. With this connectivity, B will activate when both its pre-condition A is active and the appropriate sensory condition is given. This can be used to build a sequence of neurons representing elementary phases, where the completion of one phase activates the next phase. The previous phase neuron can then be deactivated by adding a strong inhibitory connection from B to A .

These dynamic networks of Amari neurons allow to code simple sequences (and also trees or more complex graphs) of goal-directed behaviors without the use of algorithmic if-then structures. The whole movement system thus consists of interconnected dynamical systems, increasing the similarity to how living organisms generate behavior. For more on this, see (Sandamirskaya & Schöner, 2010).

We use this basic approach to dynamic behavior organization in two examples: (1) picking up a pen standing on a table and laying it down, and (2) picking up a bottle containing a soft drink and pouring it into a glass. The different elementary behavioral phases used to model this are: *approach* the current target, *settle* on the target, *retreat* from the target, *close the hand* to grasp an object, *open the hand* to release an object, *pour liquid* from the grasped object, and *return* to the resting configuration.

Each behavioral phase has a number of elementary behavioral goals assigned to it, e.g. in the *approach* phase the hand should be brought to an appropriate point from which to approach the target object (see below) and the hand should be oriented in a way that allows to grasp it. This is accomplished with three movement generating vector fields, one for target acquisition, one for orienting the hand opening \mathbf{x}_{hand} and one for orienting the grasp axis \mathbf{z}_{hand} (see Figure 4.6 below). Each of these vector fields F is activated or deactivated by multiplying it with the output $\sigma(u_F)$ of the linked behavioral phase u_F .

Movement is then generated by superposing all vector fields described in this section, weighted with the appropriate behavioral activation

$$F = \sum_i \sigma(u_i) F_i. \quad (4.73)$$

The flow

$$\ddot{\theta} = F \quad (4.74)$$

of this vector field is used to generate a trajectory for $\dot{\theta}$ and θ .

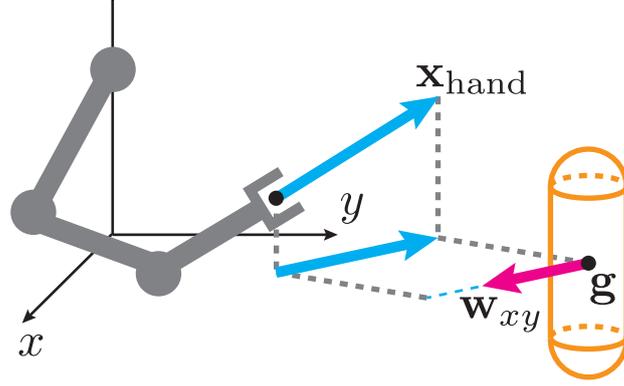


Figure 4.5.: The direction of the hand opening \mathbf{x}_{hand} is projected to the xy -plane of the object to determine the offset \mathbf{w}_{xy} that is added to the object location \mathbf{g} to determine the goal position for approaching to grasp.

4.2.6.1. The grasp and return phases

Constraints on the hand position are different for the behavioral phases of approaching a target, settling on the target, and retreating from it. Settling on the target for a successful grasp requires the direction of hand opening to coincide with the direction towards the target. Lifting a grasped object or retreating from it after setting it down also requires moving in a certain direction depending on object and hand positions and hand opening.

As an intermediate step in the fulfilling of these constraints, we use different desired hand positions in the different phases. When settling on the object, the desired hand position \mathbf{g} is the position of the object. When approaching or retreating from the object, an offset vector \mathbf{w} is added. Depending on the situation, this offset can lie in the plane orthogonal to the object cylinder axis, or simply in direction of the world frame z -axis. The former is given by $\mathbf{w}'_{xy} = P\mathbf{x}_{\text{hand}}$, $\mathbf{w}_{xy} = -c_{xy} \frac{\mathbf{w}'_{xy}}{|\mathbf{w}'_{xy}|}$, where P is the matrix that projects onto the xy -plane of the object coordinate frame, \mathbf{x}_{hand} is the x -vector of the hand coordinate frame and c_{xy} is the fixed length of the offset. The offset in direction of the world frame z -axis \mathbf{z}_{base} is simply $\mathbf{w}_z = c_z \mathbf{z}_{\text{base}}$, where c_z is again the fixed length of the offset. See Figure 4.5 for an illustration of how these offsets are defined.

Using these offsets, the hand first approaches the object by moving onto the horizontal circle around its center, where the exact goal position on this circle $\mathbf{o} + \mathbf{w}_{xy}$ is determined by the current direction of the hand opening. Then it settles upon the object by moving towards the object \mathbf{o} itself.

When the object is reached, it can be grasped by closing the hand. For all grippers we use, we can represent the hand opening by a scalar value between 1 and 0, where 1 stands for a fully opened and 0 for a completely closed hand. As the hand opening is completely independent of the other state variables, we use this value $\eta \in [0, 1]$

as task variable to generate the hand opening and closing movement. A vector field over η with a linear attractor at the desired value η_0 is given by

$$f^{(\text{grasp})} = -\alpha_\eta (\eta - \eta_0), \quad (4.75)$$

the transformation into joint space is trivial, so the vector field for grasping and releasing is

$$F^{(\text{grasp})} = F_{\eta, \eta_{\text{des}}}^{(\text{grasp})} = f^{(\text{grasp})}. \quad (4.76)$$

After a movement is completed successfully, the robot arm simply returns to a resting configuration $\theta^{(\text{rest})}$. Using the joint velocities as variables, the vector field for this is given by

$$F^{(\text{rest})} = -\alpha_{\dot{\theta}} (\dot{\theta} - \alpha_\theta (\theta^{(\text{rest})} - \theta)) \quad (4.77)$$

4.3. Implementation and results

The movement generation system described in the previous section was implemented in C++ and tested in with combinations of tasks for a variety of scenes. In all experiments, the scene configuration was pre-defined and known to the system a-priori, i.e. information about location and size of all objects was given and did not have to be gathered from sensor information.

Two different robotic agents were used in the experiments, CoRA and the Aldebaran NAO. CoRA (Cooperative Robotic Agent, see Figure 4.6 and Iossifidis et al., 2002) has an anthropomorphic seven degrees of freedom arm mounted on a one degree of freedom trunk. CoRA is assembled from modular robotic parts, each module is servo-controlled and communicates via a CAN-bus interface with the controlling PC. Above the trunk CoRA has a head consisting of a two DoF pan/tilt unit carrying a stereo color camera system and microphones. The arm of NAO consists of a *pitch-roll* joint at the shoulder, a *yaw-roll* joint at the elbow, and a *yaw* joint at the wrist, for a total of five degrees of freedom affecting the hand, with a single additional degree of freedom for the opening of the hand (see Figure 4.6). For CoRA, the arm and trunk were used for a total of 8 DoF. For NAO, only the 5 DoFs of the arm were used.

In all experiments, a trajectory was generated by calculating the vector field given in Equation 4.74 at each time step and passing it as a joint acceleration to the robot interface.

4.3.1. Obstacle avoidance

To examine the obstacle avoidance scheme, a series of experiments was run with a movement generation system that consisted only of the vector fields for target acquisition (Equation 4.38) and obstacle avoidance (Equation 4.53). If the end-effector was near the target, the homogeneous damping field was also activated to reduce residual motion in the null-space of the end-effector position, as indicated in Equation 4.71.

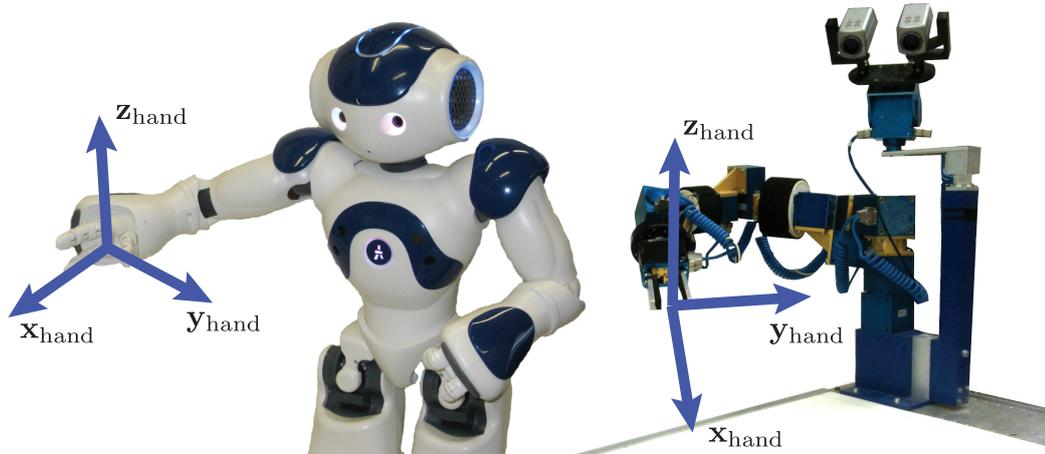


Figure 4.6.: The anthropomorphic robots CoRA (right) and NAO (left). The vectors x_{hand} and z_{hand} of the hand coordinate frame are relevant for grasping and manipulating cylindrical objects.

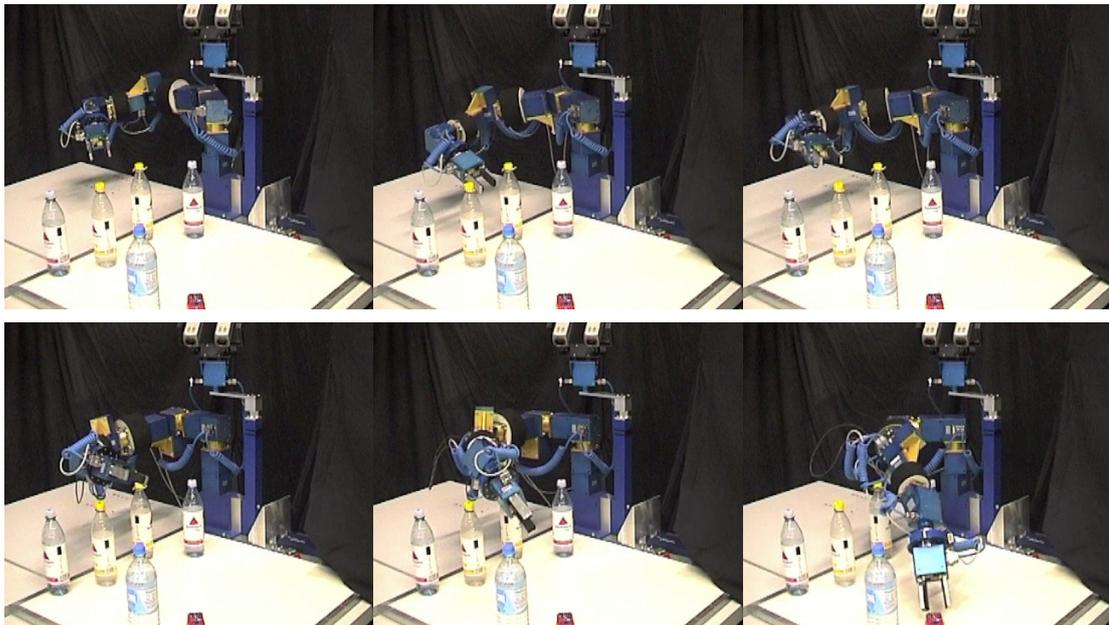


Figure 4.7.: Sequence of CoRA successfully reaching for a toy car while avoiding collision with five obstacles in the scene.

The behavior generation scheme was capable of producing satisfying movement trajectories for CoRA for scenes with several obstacles. Figure 4.7 shows CoRA reaching for a red toy car in a scene where the direct path to the target is obstructed by several plastic bottles. The manipulator successfully avoids all obstacles and reaches the target with the tip of the gripper. The experiment was stopped shortly before the car was actually touched.

For a more systematic examination of the behavior generation scheme, the method was tested in a sequence of randomly generated scenes in the software simulator of CoRA. In each trial, the initial configuration was randomized around a reference configuration ($\sim \mathcal{N}(\theta_{\text{init}}, 0.1 \text{ rad})$). The target was placed randomly in a predetermined area, and a varying number of obstacles was also placed randomly in an area ranging from the manipulator to around the target (uniform distributions). No obstacle was placed closer than 9cm to the target, allowing a minimal leeway of 3cm for the end-effector, enclosed in a 6cm radius bounding volume, to reach it. Radius and height of the obstacles was also randomized, while the manipulator starting configuration was fixed at θ_{init} (see Figure 4.8).

Table 4.1 shows the results of the experiment. A successful trajectory to the target was found in the majority of all trials, overwhelmingly so for small numbers of obstacles and decreasing only significantly as the scenes got cluttered. An investigation of the failed trials revealed that the reasons for failing to find a path to the target fall in fairly distinct categories, which are described below, listed with the reference letter used in the table.

Proximate obstacle (P): The most prevalent reason was a single high obstacle near the base of the manipulator that prevented the link segments close to the base to move towards the target, reducing the effectively reachable workspace significantly. This resulted in failures when the target was located far out in the workspace, though for closer targets successful trajectories could still be found (an example of this situation is shown in Figure 4.8).

Workspace boundaries (W): In a small number of cases, the manipulator reached the end of the workspace while moving around an obstacle, and then failed to find a way back.

Obstructed target (O): For certain combinations the target was so obstructed by the obstacles that a configuration that reached the target without collision was nonexistent or very hard to find. Failures of this kind did not occur for scenes with small numbers of obstacles, and only became frequent in very cluttered scenes.

Cancellation (C): When one link segment simultaneously approached two or more obstacles on different sides, the avoidance directions partly cancelled each other out, until the obstacle distance had become very small and the corresponding factor w_δ in the repelling vector field so large that the simulation became numerically unstable.

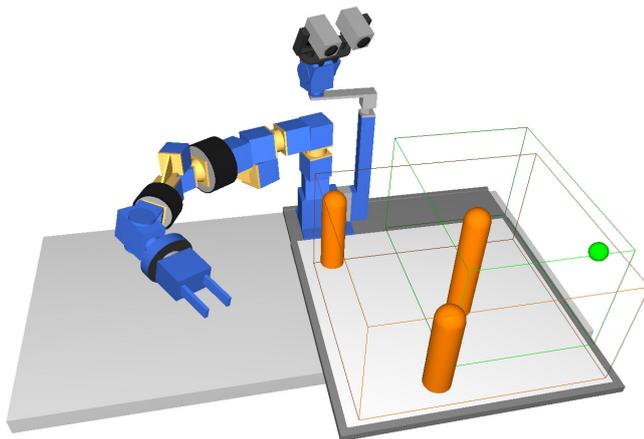


Figure 4.8.: Randomly generated scene with three obstacles. The manipulator is in the reference configuration θ_{init} . The target was randomly placed in the green box, the obstacles distributed over the orange box.

Null-space drift (D): When the manipulator avoided an obstacle with a link segment close to the base, this resulted in persisting motion in the null space of the end-effector even after the obstacle has been passed. In situations where only a tight path towards the target was available, that null space drift prevented the manipulator from successfully taking that path.

The first three categories of failures (P, W, O) are instances of the general situation where a path that reaches the target without collision might exist, but can only be found by first realizing a significant change of manipulator configuration. As the behavior generation scheme presented here is a local approach, this shortcoming is to be expected to some degree. A possible way to prevent the cancellation problems (C) would be to include a more sophisticated method of regulating the absolute velocity than the one given by Equation (4.33), with the additional function to reduce the overall velocity in the vicinity of obstacles. This would allow more time to change the movement direction away from possible collision paths. This is not the focus of the present work, though. The persistence of null-space drift after obstacles have been cleared (D) is problematic, but a rather academic one, as the small number of occurrences indicates. One possible solution is to identify occurrences of strong velocity components in the null-space of the main tasks and selectively increasing the damping in these directions.

4.3.1.1. Special cases

In addition to the randomized scenes, we set up two special cases with obstacle configurations for which it is particularly hard to find a successful trajectory as qualitative demonstration of the practical applicability of our presented scheme. In the first case, the target is encircled within several large obstacles, leaving only a

No. of obst.	N	success (%)	failure (%)	Failure reasons (%)				
				P	W	O	C	D
1	1000	99.7	0.3	33.3	67.7	–	–	–
3	1000	99.1	0.9	77.8	11.1	–	11.1	–
6	1000	95.4	4.6	54.4	6.5	19.6	15.2	4.3
10	1000	92.7	7.3	47.9	6.8	27.4	9.6	8.2
15	500	87.0	13.0	49.2	3.0	35.4	9.2	3.0
20	200	83.0	17.0	41.1	–	47.0	8.8	2.9

Table 4.1.: Results of experiments with randomized scenes. Failure reasons are given in percent of the total number of failures.

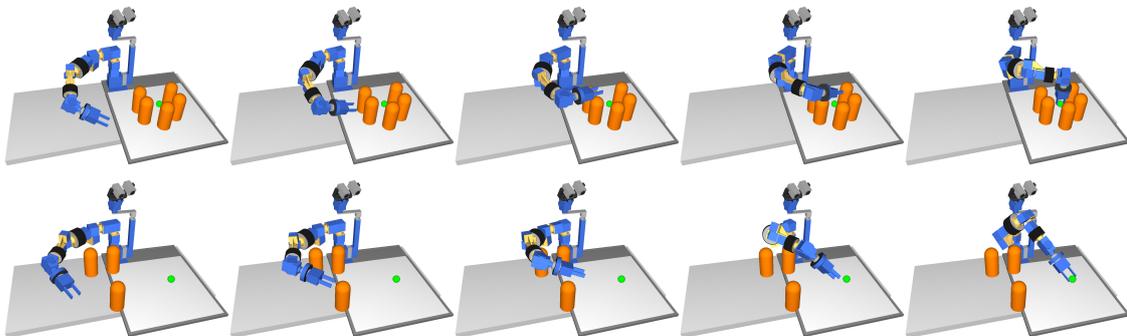


Figure 4.9.: Two sequences of simulated movements of CoRA reaching for a target. The obstacles were manually arranged to pose interesting challenges.

narrow path from above for the manipulator to reach through (Figure 4.9, upper panel). The second scene is a borderline case of the failed random trials with a large obstacle near the first link segments (P), but chosen in a way that the manipulator has barely enough freedom to find a path above it (Figure 4.9, lower panel). In both cases the behavior generation scheme finds viable trajectories that reach the target and avoid obstacle collision.

4.3.2. Complex tasks

Our next experiment was designed to test the orientation contribution and whether the system is capable of generating complex sequences of actions. To this end, two tasks were set up for the NAO, both containing the grasping and manipulation of a cylindrical object that required appropriate orientation of the hand. In the first movement, NAO picks up a marker pen standing on the table in front of the robot and lays it down at another place on the table. In the second experiment, NAO takes a straight plastic container filled with a soft drink, pours the contents into a glass and puts the container back on the table.

To generate the sequence of behavioral phases for these movements, each phase was represented by an Amari neuron as described in section 4.2.6. The next neuron in the sequence was activated when (a) the previous neuron was active and (b) all active behavioral variables were near the desired values.

4.3.2.1. Pick-and-place movement

The pick-and-place movement consists of two similar parts. In the first part the target object is approached and settled upon, then the hand is closed and the object lifted. In the second part, the target position for the object is approached and settled upon with the object in hand, then the object is released by opening the hand, which is then removed. Once clear of the object, the arm returns to the resting configuration.

In the first part of the movement, to successfully grasp the marker pen, the z -axis of the hand frame \mathbf{z}_{hand} must be aligned with the long axis of the pen. When the pen is standing on the table, this is the z -axis of the world coordinate frame $\mathbf{z}_{\text{world}}$. The corresponding orientation angle is

$$\gamma_z = \angle(\mathbf{z}_{\text{hand}}, \mathbf{z}_{\text{world}}) = \arccos(\mathbf{z}_{\text{hand}}^T \mathbf{z}_{\text{world}}), \quad (4.78)$$

and the desired state is $\gamma_z = 0$. When placing the pen on the table, the z -axis of the hand, which is now equivalent to the long axis of the pen, must be parallel to the table surface, corresponding to the desired state $\gamma_z = \frac{\pi}{2}$.

An additional constraint when placing the pen on the table is that the direction of the hand opening, the x -axis of the hand frame \mathbf{x}_{hand} , must point roughly towards the table, so the pen actually drops to the table when the hand opens. The corresponding orientation angle is

$$\gamma_x = \angle(\mathbf{x}_{\text{hand}}, \mathbf{z}_{\text{world}}) = \arccos(\mathbf{x}_{\text{hand}}^T \mathbf{z}_{\text{world}}). \quad (4.79)$$

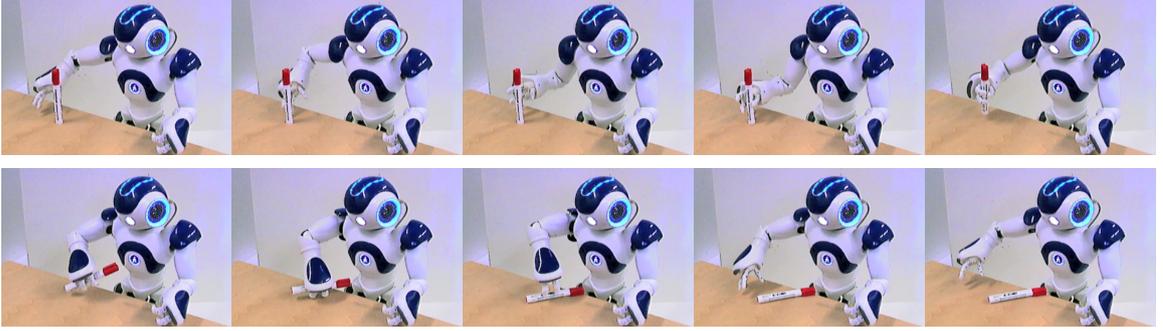


Figure 4.10.: Sequence of NAO picking up a marker pen, rotating it parallel to the surface and placing it on the table.

Table 4.2.: Pick-and-place movement phase sequence

vector field attractor	$F_{\mathbf{p},\mathbf{x}}^{(\text{tar})}$ $\mathbf{x} =$	$F_{\gamma_z, \gamma_{z,\text{des}}}^{(\text{ori})}$ $\gamma_{z,\text{des}} =$	$F_{\gamma_x, \gamma_{x,\text{des}}}^{(\text{ori})}$ $\gamma_{x,\text{des}} =$	$F^{(\text{grasp})}$ $\eta_{\text{des}} =$
<i>approach</i>	$\mathbf{q}_i + \mathbf{w}_{xy}$	0	-	1
<i>settle</i>	\mathbf{q}_i	0	-	1
<i>close hand</i>	$\mathbf{q}_i + \mathbf{w}_{xy}$	0	-	0
<i>retreat</i>	$\mathbf{q}_i + \mathbf{w}_z$	0	-	0
<i>approach</i>	$\mathbf{q}_t + \mathbf{w}_z$	$\pi/2$	$\geq \pi - 0.3$	0
<i>settle</i>	\mathbf{q}_t	$\pi/2$	$\geq \pi - 0.3$	0
<i>open hand</i>	\mathbf{q}_t	$\pi/2$	$\geq \pi - 0.3$	1
<i>retreat</i>	$\mathbf{q}_t + \mathbf{w}_z$	$\pi/2$	$\geq \pi - 0.3$	1

As the exact value of this orientation angle γ_x does not matter as long as it is large enough, a whole region $\gamma_x \geq \pi - 0.3$ was used as desired states. This was realized by a vector field that had a one-sided point attractor at the border of the desired state region (Equation 4.57) and vanished inside it.

Besides the orientation constraints, there are constraints on the hand position \mathbf{p} throughout the movement. The first desired position is the location \mathbf{q}_i of the marker, adjusted by an approach offset \mathbf{w}_{xy} as described in section 4.2.6.1 above. These locations are set as the desired state of the hand position for the target acquisition vector field contribution (Equation 4.37).

Table 4.2 gives an overview of the sequence for this movement. For each behavioral phase, the movement generating vector fields contributing to the total vector field are listed. The entries represent the desired values of the behavioral variable, i.e. the attractor of the behavioral dynamical system f_* that the joint velocity vector field F_* corresponds to. A dash means that the respective behavioral variable is irrelevant in that phase and the vector field does not contribute. The movement is finished with a phase where the arm returns to the predefined resting configuration θ_{res} , not

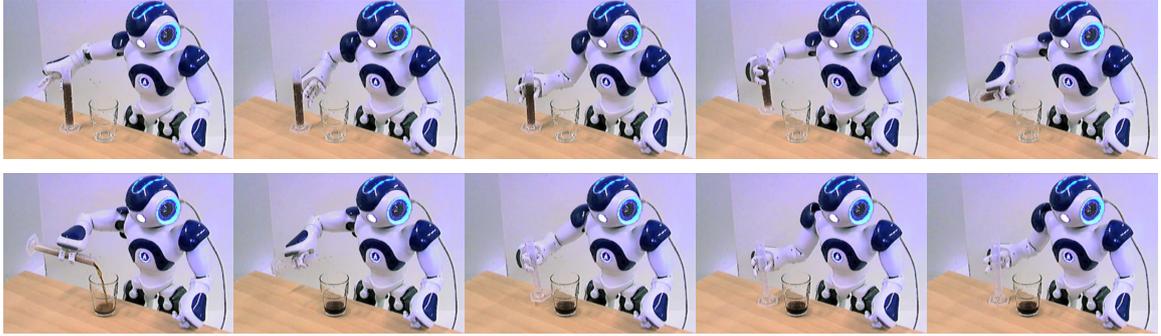


Figure 4.11.: Sequence of NAO taking a container holding a soft drink and pouring it into a drinking glass, then putting the container back on the table.

shown in the table. During this last phase, the vector field contributions for all other movement generators are set to zero.

NAO was able to repeatedly pick up the pen and place it on the table with this movement generation system. An example of a successful movement is shown as a series of stills in Figure 4.10. The range of initial positions of the pen relative to the robot was rather limited, though we did not test this systematically. One reason for this is the small number of available degrees of freedom in relation to the task constraints. Incorporating locomotion and trunk movement might alleviate this problem.

4.3.2.2. Pour liquid movement

The pouring movement consists of three parts: picking up the soft drink container, pouring its content into the glass, and putting the container back on the table. The container is picked up at the position \mathbf{q}_c and placed back there in the same manner as in the first movement, except that the container is placed on the table upright instead of horizontally. For these parts, the same orientation angle γ_z was used as behavioral variable with the appropriate desired value $\gamma_z = 0$.

For the middle part of the movement, pouring the soft drink from the container into a glass on the table, the mouth of the container must be kept over the glass in order for the liquid to be poured into the glass instead of on the table. In this phase, the contribution of the container position vector field $F_{\mathbf{c}, \mathbf{q}_g}^{(\text{tar})}$ becomes active, with the location of the glass \mathbf{q}_g plus a vertical offset as desired value for the container mouth position \mathbf{c} . The hand position becomes irrelevant in this phase, and the contribution from that vector field is turned off. The pouring movement is realized by setting a sufficiently large desired value $\gamma_z = \frac{3}{4}\pi$ for the behavioral variable γ_z , meaning the long axis of the container, which is collinear to the z -axis of the hand, is pointing down towards the table at an angle of 45° . Again, an overview of the behavioral phase sequence for this movement is given in table 4.3.

With the movement generation system presented here, NAO successfully picked up the container, poured the soft drink into the glass and put the container down on

Table 4.3.: Pour liquid movement phase sequence

vector field attractor	$F_{\mathbf{p},\mathbf{x}}^{(\text{tar})}$ $\mathbf{x} =$	$F_{\gamma_z, \gamma_{z,\text{des}}}^{(\text{ori})}$ $\gamma_{z,\text{des}} =$	$F_{\mathbf{c},\mathbf{x}}^{(\text{tar})}$ $\mathbf{x} =$	$F^{(\text{grasp})}$ $\eta_{\text{des}} =$
<i>approach</i>	$\mathbf{q}_c + \mathbf{w}_{xy}$	0	-	1
<i>settle</i>	\mathbf{q}_c	0	-	1
<i>close hand</i>	\mathbf{q}_c	0	-	0
<i>retreat</i>	$\mathbf{q}_c + \mathbf{w}_z$	0	-	0
<i>pour liquid</i>	-	$\frac{3}{4}\pi$	$\mathbf{q}_g + \mathbf{w}_z$	0
<i>approach</i>	$\mathbf{q}_c + \mathbf{w}_z$	0	-	0
<i>settle</i>	\mathbf{q}_c	0	-	0
<i>open hand</i>	\mathbf{q}_c	0	-	1
<i>retreat</i>	$\mathbf{q}_c + \mathbf{w}_{xy}$	0	-	1

the table without spilling. Figure 4.11 illustrates this action sequence with a series of stills.

4.3.3. The complete system

A last series of experiments was carried out on CoRA to examine the complete system with all vector field contributions active. The scene consisted of a target and a single obstacle on the table, both cylindrical. The active vector field contributions were target acquisition, hand orientation, obstacle avoidance, joint limit avoidance and homogeneous damping. Orientation angle and attractor were the same as the ones used for NAO grasping the pen in the previous section.

4.3.3.1. Systematic survey

To explore the capabilities of successfully generating movements that adhere to all the constraints described in Section 4.2, a systematic experiment was conducted in a simulation study. The manipulator started in one of 6 initial configurations, roughly similar to how a human would hold the right arm when manipulating something on a table. The gripper was located in front of the base or 40cm to the left or right, 15 or 35cm above the work table. The target was in one of 28 different positions: 20 or 40cm above the table, 60 or 70cm in front of the base, and from 60cm to the left to 60cm to the right of the base, in steps of 20cm. The obstacle was a cylinder with a diameter of 8cm and a length of 30cm. It was placed in one of 27 positions, the locations of which depended upon the initial position of the gripper \mathbf{p}_{init} and the target position \mathbf{g} : one location was halfway between the gripper and the target, the other locations were on the corners and the centers of the edges and sides of a cube around the first one. The side length of the cube was one sixth of the distance

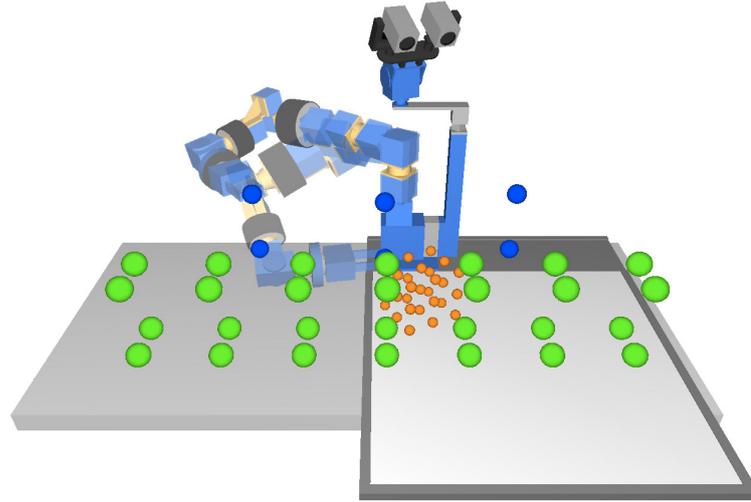


Figure 4.12.: Sketch of the experimental setup. The 6 initial gripper positions are marked blue, the 28 target positions green. For one example combination of gripper and target position, the 27 obstacle positions are indicated by small orange markers. Two example initial configurations of the whole manipulator are also shown.

$|\mathbf{g} - \mathbf{p}_{\text{init}}|$; the cube was aligned with the z -axis and the direction from \mathbf{p}_{init} to \mathbf{g} . See Figure 4.12 for an overview of all possible combinations.

Out of these $6 \cdot 28 \cdot 27 = 4536$ movements, 4423 were completed with all constraints met throughout the movement, a success rate of 97.5%.

4.3.3.2. Effect of damping dependency on the behavioral state

The effect of the homogeneous damping term is to prevent velocities from building up in the null-space of currently relevant task variables. This becomes evident in movements where different task contributions are opposing each other. In these situations, even small differences between the vector field contributions can build up over time, leading to a way out of the conflict. Homogeneous damping is a hindrance to this, as illustrated in Figure 4.13, where joint angle profiles for two movements with the same setup are shown, one with damping depending on the behavioral state as defined in Equation 4.71, the other with damping permanently active.

The scene for this experiment was set up in a way such that after avoiding the obstacle with the last link segment and turning towards the target, reaching the target would lead to the back of the “wrist” link segment colliding with the obstacle. One way to avoid this is to turn the wrist and reach the target from a different direction. This solution emerged almost immediately in the undamped case, after about 10s the movement is completed successfully. With permanent damping, the manipulator does not leave the configuration range where the contribution for wrist obstacle avoidance and the contribution for target reaching oppose each other for a

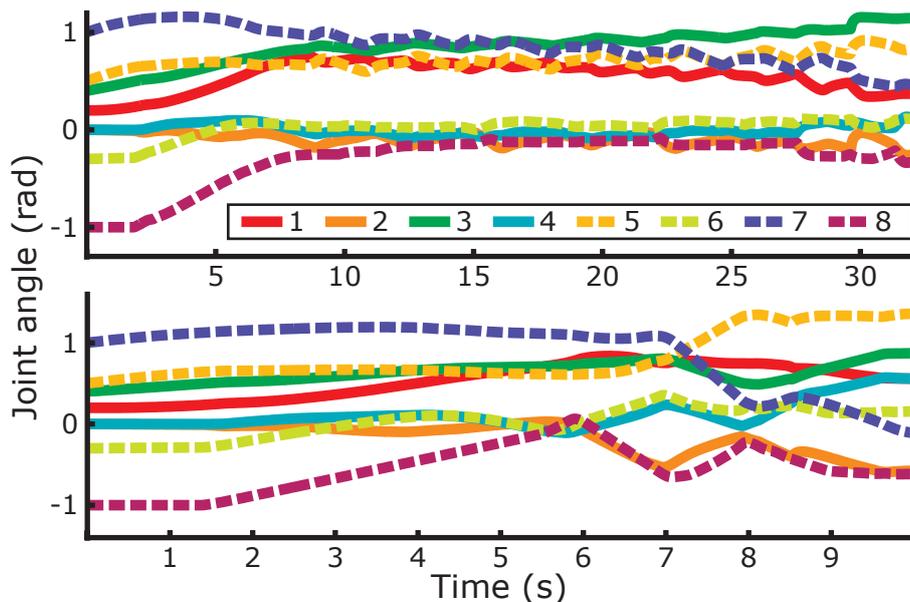


Figure 4.13.: Joint angle profiles for a movement with damping depending on the behavioral state as described in Equation 4.71 (upper panel), and damping permanently tuned on (lower panel).

long time – for about 20 seconds, the arm oscillates in the same general region of state space, until a slow underlying movement component finally lets it escape and reach the target. The solution of the permanently damped system is similar to the one where the damping depends upon the behavioral state, but the system takes three times as long to find it.

4.3.3.3. Joint limit avoidance

The effects of different versions of joint limit avoidance are demonstrated in an experiment where the final configuration is very close to the joint limit for the 5th joint, the “elbow”, due to obstacle placement. Figure 4.14 shows the 5th joint angle for three different avoidance modes.

When the limit is simply guaranteed by the low-level motor system, the target is reached, but the joint velocity goes through a very sharp drop as the limit is encountered. With the avoidance strategy of accelerating the joint away from the limit when it comes too close, given in Equation 4.66, the movement fails, as the manipulator enters a cycle of approaching the target, then being forced away from it by the joint limit vector field. Using the vector field depending on the relation of the current distance from the joint angle limit and rate of change towards it, given in Equation 4.69, the target is reached with a much smoother curve than without the repelling vector field.

In order to isolate the effects of the different joint limit vector fields and minimize influences from other behavioral goals, the hand orientation contribution was dis-

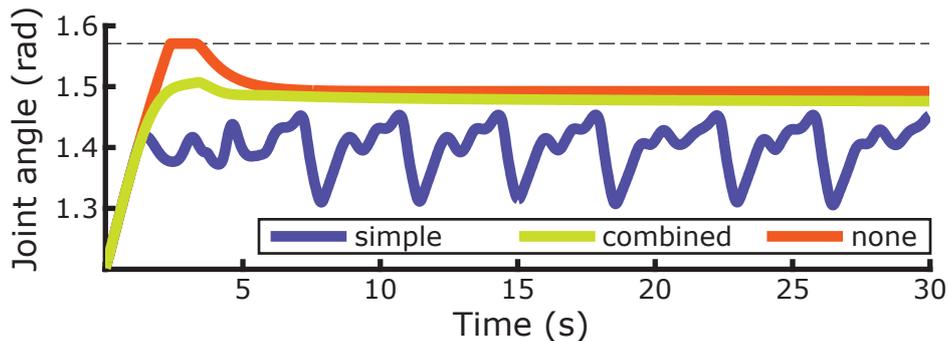


Figure 4.14.: Angle profile for the 5th joint (“elbow”) for the same movement with three different modes of joint limit avoidance. In the first movement (orange), excessive joint angle commands $\theta > \theta^{(\text{upper})}$ were simply reduced to the limit $\theta^{(\text{upper})}$. The second movement (blue) used the simple angle-dependent repellation from the limit given by Equation 4.66, which resulted in a cycle. The last movement (green) used the combined angle-velocity dependent vector field given by Equation 4.69. The joint limit $\theta^{(\text{upper})} = \pi/2$ is indicated by the dashed line.

regarded and the damping term activated throughout the whole movement in this experiment.

4.4. Conclusions

We presented an approach to generate goal-directed movement for autonomous robots based on principles of selective stabilization of task-relevant variables. Each goal-directed sub-task was expressed as the desired state of a relevant low-dimensional task variable. To move the task variable to its desired state, the relevant aspects of its movement states were represented with a second set of task variables. By specifying desired values for these second-order task variables and stabilizing them at these values, we ensured that the first-order task variable was changing in the right direction. This movement generation scheme was phased over to a postural stabilization scheme once the task-relevant variable came close to its desired state.

The proposed movement generation system was capable of solving a wide variety of goal-directed movement tasks for different robotic manipulators. The system successfully integrated the large number of constraints and solved the complex task based on purely local planning. The performance achieved in the systematic experiments was surprisingly strong for a system without capabilities to compare and choose from possible trajectories globally. Failures to fulfill the behavioral goals or globally suboptimal trajectories can be explained by problems inherent to local approaches. Drawbacks were mostly dealt with heuristically: a class of problems was parameterized, allowing the relevant behavioral task variables and corresponding vector fields

to be modified depending on that parameter.

Our approach is based on the insight that the concrete value of the movement state is less important than the general direction of movement. Even less important during movement is the current value of the first-order task variable. During movement, we selectively stabilize the general direction of the task variable change at values that bring it closer to the goal state. Because it is sufficient if the movement direction is only generally appropriate, errors introduced by perturbations from other task contributions are not critical. Instead of correcting the errors immediately with high effort and possible violation of other constraints, they are remedied over time: the manipulator moves into a configuration where the two sub-tasks are not in conflict anymore. The trajectories emerging from the local application of selective stabilization principles are usually quite smooth and appear more natural than trajectories generated by more classical robotic approaches (Fajen et al., 2003).

Another benefit that has only been hinted at in this study is the compatibility of our system with dynamic approaches to behavioral organization. To generate complex sequences of movements, we essentially hard-coded when each contribution should be activated with which parameters. This kind of decision can also be made autonomously by more complex dynamical systems that are capable of making decisions based on both sensory data and internal states (Richter, 2012). Instead of following a pre-defined sequence, these systems can dynamically select an appropriate action from a range of several possibilities. Selections and other decisions are represented by dynamic neural fields in this system. The neural activation of these fields can be connected to the weight with which the vector field of a sub-task contributes to the movement generation. Using such a complex system of behavioral organization instead of the simple one described in Section 4.2.6 would provide the robotic agent with another degree of autonomy.

One aspect of movement generation we did not address here is the time profile of the trajectory. The solution we used was to set an attractor for the end-effector speed at an arbitrarily chosen value. While this was necessary to resolve the redundancy of trajectories, we did not attempt to benefit from this freedom. The speed of movement and the danger of collision with obstacles or joint limits are closely related. We took this relationship into account in one direction by increasing the magnitude of the avoidance vector fields large for small ratios of speed over distance to the obstacle. But collision was still not prevented successfully in some rare cases where contributions from two different obstacles cancelled each other out. This could be prevented by also coupling the obstacle distance to the movement speed dynamics in the other direction. If the movement speed is reduced in proximity of an obstacle, there is more time to turn away from it. The small contributions from the two obstacles that do not cancel each other out are then more likely to be sufficient to prevent collision. This temporal aspect of trajectory generation is something that might be addressed in future work.

5. Conclusions

Interacting with the world in a meaningful way requires moving the body in a goal-directed fashion. Usually there are infinitely many possible solutions to a given behavioral task. Any process that resolves this redundancy has to coordinate the available degrees of freedom in a way that all task-specific constraints are fulfilled. This selection of one solution out of the infinite set of movement trajectories that reach the behavioral goal has to be robust and flexible enough to deal with noisy state sensors and unexpected perturbation.

Redundancy in movement generation has both local and global aspects. Any desired instantaneous change of the task variable can be realized by infinitely many changes in the body state. To form a global movement trajectory for the task variable, different sequences of local changes can terminate in the same state and are globally equivalent from a task-perspective.

We proposed a general principle of how this redundancy can be resolved: variables that are currently relevant to a given task are stabilized, while others that are not relevant at the moment are released from control. This notion was applied directly to the local aspect of redundancy by selecting a change of the body state that is minimal while still generating the desired change of the task variable. To resolve the global redundancy of trajectory formation, we posited that at any point in time during a movement, those variables are relevant that describe whether the system is currently moving in the right direction.

The feasibility of this principle was examined from three different angles. First we asked whether humans use selective stabilization principles when generating movements. We presented a method to quantify signatures of selective stabilization in behavioral data and formalized hypotheses about whether these signatures are present. A statistical test was developed to examine the verity of this type of hypothesis in data from single human subjects and applied to an exemplary data set.

The second question we asked was what mechanisms the human nervous system might use to achieve this selective stabilization. We designed a feedback controller for the postural stabilization of the body in quiet, upright stance. Sensed deviations in task-relevant variables are transformed into negative feedback and distributed among the available degrees of freedom in a way that minimizes the total muscle activation. This feedback system is integrated into a complete model of the sensorimotor loop for quiet stance. The feasibility of our design was demonstrated by successfully reproducing the characteristic variance patterns of body sway in quiet stance, among them the behavioral signatures of selective stabilization presented previously.

To investigate the global aspects of redundancy in movement generation, we utilized embodied robotic agents. We designed a system that formulates a behavioral

goal as the desired value of a task variable. To generate a trajectory, a vector field is constructed that stabilizes the movement state of that variable in a region where it is approaching the desired value. Less relevant aspects like the exact value of the movement state and how the rest of the robotic agent moves are released from control. The flow of the superposed vector fields for different behavioral goals generated trajectories that fulfilled each goal simultaneously in a large number of tested scenarios.

Based on the successful results of these three studies, we conclude that selective stabilization is indeed a viable strategy of movement generation and that there is substantial evidence that humans apply similar strategies to resolve the inherent redundancy in postural stabilization and trajectory formation.

5.1. Scientific contribution

The scientific contribution of this work is twofold. On the one hand, the concrete work described in Chapters 2–4 each extend the knowledge in the specific fields of research. On the other hand, the general approach of describing the generation of movement in high-dimensional state space with dynamical systems has the potential to provide a unifying language for research questions that have traditionally been analyzed by separate communities using disjoint methods.

In Chapter 2 we formalize the UCM method and make underlying assumptions about the distribution of the elementary variables explicit. We formulate hypotheses about the structure of the variance in a way that makes them testable with the parametric bootstrap method. Similar tests have been made on the population level before using repeated measures ANOVA. Our work allows to test these effects on the level of individual subjects instead of whole populations. An exemplary application of the developed test showcases possible problems in the population-level analysis.

The process model of quiet upright stance presented in Chapter 3 is one among only a few other such models that account for the existence of multiple degrees of freedom along the body instead of reducing the analysis to a single joint at the ankle. To the best of our knowledge, the model presented here is the first closed-loop process model of quiet upright stance that includes the dynamics of the spinal reflex loops. While other researchers have pointed out that the existence of these reflex loops might reduce the complexity of the control problem for the higher parts of the central nervous system (Van Soest et al., 2003), our model is the first to actually support this claim by showing that in the presence of these low-level contributions a relatively simple feedback dynamics is sufficient to stabilize the body. Furthermore, the model provides additional evidence that for effector systems with many degrees of freedom, a motor system that represents and monitors movement states on a task level rather than the high-dimensional configuration space is feasible.

This claim is extended by the results of Chapter 4, where we show that a movement generation system for autonomous robotic agents that is based on principles of selectively stabilizing task-relevant low-level variables at desired states is capa-

ble of solving a wide variety of motor tasks with multiple simultaneous constraints. In the field of robotics, this can be interpreted as a generalization of the widely-used Artificial Potential Field Approach (APFA, Khatib, 1986). While the vector fields erected by the APFA are induced by potential functions and thus restricted to position-dependent variables only, the attractor dynamics approach presented here allows freedom in the choice of relevant variables. This allows using vector fields over velocity-dependent variables, an essential part of the kinematic state of the robot. The notion of stability as a design principle for robotics had been proposed in the attractor dynamics approach to vehicle path planning before (Schöner et al., 1995). In the present work, we extended this notion to the principle of selective stabilization in redundant tasks and applied it to reaching, grasping and collision avoidance for redundant manipulators.

On a more general level, this work makes a contribution by proposing dynamic systems as a common language to describe and generate movement in high-dimensional state spaces in and across different disciplines. Each academic discipline has a collection of classical, time-proven methods, and problems that these classical methods are ill-suited to deal with have a tendency to be marginalized.

The field of robotics is dominated by engineers and computer scientists. This has resulted in a preference for actuators and systems that can be modeled with great accuracy, like joints actuated by micro-controlled servo-motors that have been analyzed in detail by control theoretical approaches. For movement generation, there is a tendency to model the kinematics and dynamics of the robotic systems and solve the movement generation problem by inverting these forward models. This is driven to the limits of what is possible in grasping problems, where the kinematic state of some robotic hands is described by 20 or more degrees of freedom. Possible grasps for any given object correspond to sub-manifolds, but these are sometimes so sensitive that small errors in position estimation already lead to failure. Though there has been considerable progress in dealing with these problems, in recent years increasingly many roboticists have acknowledged that by using actuators with inherent stability properties like elastic tendons, many of these problems are greatly alleviated: with elastic tendons, one can just close the hand and the fingers will generate an appropriate amount of force against the object. The exact magnitude of force does not matter in many cases. Doing the same for servo-controlled joints will usually result in broken motors. Yet the robotics community is slow to adapt hardware with elastic properties, because modeling the kinematics of these accurately is almost impossible. As we show in Chapter 3 of this thesis, for a control approach based on dynamic systems with stable attractors such low-level elastic properties can actually reduce the complexity of the high-level control problem.

In the field of human motor control, many researchers have acknowledged the beneficial effects of redundancy and elastic properties. Additional degrees of freedom are not seen as something problematic because the controller *has to* select a solution out of an infinite set, but as an advantage because they add more freedom of ways in which a specific movement goal *can* be achieved. To reflect this different view, many researchers have begun to talk about additional degrees of freedom as *motor*

abundance instead of redundancy. The “softness” induced by elastic components and low-level reflex loops is treated with more ambiguity, though. One group of researchers embraces it and postulate that a major part of the work is solved by the spinal reflex loops and the control problem for the higher motor areas is as simple as specifying a desired position for a relevant body part (Feldman, 1986). Many other researchers approach the motor control problem from a very different angle, using engineering techniques like optimal control to support hypotheses about trajectory formation adhering to optimality principles (Todorov, 2004). While this direction has had considerable success in exposing optimality characteristics of mean trajectories and thus pointing out aspects that the CNS might care about in movement generation, no satisfying explanation for how these optimality constraints might actually be used in the neural dynamics that generate descending movement commands has been proposed yet (Loeb, 2012). One overt problem is that optimal control approaches often assume that the descending motor commands generated by the CNS encode muscle force or joint torque directly, ignoring both the elastic properties of the muscles and tendons and the dynamics of the spinal reflex loops. A more general problem is that these approaches generate globally optimal trajectories, requiring pre-planning of a complete movement in advance, and re-planning every time the movement parameters change in some unpredicted way. In other words, the movements generated by optimal controllers fail to stabilize the movement goal against perturbations.

We propose that these problems can be overcome by adopting dynamical systems as a general language of description for the representation and generation of movement. In human motor control, describing the dynamics of the neural activation patterns in the brain by differential equations allows the introduction of local optimality principles. The principle of selective stabilization is closely related in spirit to the “principle of minimal intervention” proposed by members of the optimal control community (Valero-Cuevas, Venkadesan, & Todorov, 2009), and could even be called a dynamic version of it. Models that use dynamical systems as a unifying language have the additional benefit that the level of detail used in describing sub-modules of a complex system can be varied relatively freely, as the interface between subsequent sub-modules are always neural activation variables. It is, for instance, conceivable to replace the abstract, functional description of the neural dynamics we presented in Chapter 3 by a detailed model using dynamic neural fields (Schöner, 2008), which would allow the comparison of the neural activation patterns generated by the model with experimental data. In robotics, designing movement generation systems with stability properties can simplify the high-level control problem and interface well with soft actuators that have elastic properties. Across these disciplines, a common language of dynamic systems thinking can provide a basis for academic exchange.

A. Parameters of the bootstrap simulation study

Joint and pointer positions in reference configuration

$$\mathbf{j}_1 = \mathbf{j}_2 = \mathbf{j}_3 = \begin{pmatrix} 0 \\ 0 \\ 0.3 \end{pmatrix} \text{ m}, \quad \mathbf{j}_4 = \mathbf{j}_5 = \mathbf{j}_6 = \begin{pmatrix} 0 \\ 0.2 \\ 0.3 \end{pmatrix} \text{ m},$$
$$\mathbf{j}_7 = \mathbf{j}_8 = \begin{pmatrix} 0 \\ 0.4 \\ 0.3 \end{pmatrix} \text{ m}, \quad \mathbf{j}_9 = \mathbf{j}_{10} = \begin{pmatrix} 0 \\ 0.6 \\ 0.3 \end{pmatrix} \text{ m}, \quad \mathbf{p} = \begin{pmatrix} 0 \\ 0.75 \\ 0.3 \end{pmatrix} \text{ m}$$

Joint axes in reference configuration

$$\omega_1 = \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}, \quad \omega_2 = \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix}, \quad \omega_3 = \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}, \quad \omega_4 = \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}, \quad \omega_5 = \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix}, \quad \omega_6 = \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix},$$
$$\omega_7 = \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}, \quad \omega_8 = \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix}, \quad \omega_9 = \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}, \quad \omega_{10} = \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}$$

Normal distribution parameters

$$\boldsymbol{\mu} = (0 \ 0 \ 0 \ 0 \ -0.25 \ -1.4 \ 1.25 \ 0 \ 0 \ 0)^T \text{ rad}, \quad \boldsymbol{\Sigma} = 0.1 \cdot I_{10 \times 10} \text{ rad}$$

B. Parameters of the posture model

Time delays

$$d_\theta = 30 \text{ ms}, \quad d_c, d_p, d_o = 120 \text{ ms}$$

Noise

$$\sigma_*^2 = \text{variance of integrated white noise in the Ornstein-Uhlenbeck process after 1s}$$
$$\sigma_\theta = 0.002 \text{ rad}, \quad \sigma_{\dot{\theta}} = 0.002 \text{ rad s}^{-1}$$

$$\begin{aligned} \sigma_{\dot{c}} &= 0.02 \text{ m s}^{-1}, \quad \sigma_{\ddot{c}} = 0.02 \text{ m s}^{-2} \\ \sigma_{\dot{p}} &= 0.015 \text{ m s}^{-1} \text{ (EO)}, \quad \sigma_{\dot{p}} = 0.02 \text{ m s}^{-1} \text{ (EC)} \\ \sigma_{\ddot{p}} &= 0.015 \text{ m s}^{-2} \text{ (EO)}, \quad \sigma_{\ddot{p}} = 0.02 \text{ m s}^{-2} \text{ (EC)} \\ \sigma_o &= 0.02 \text{ rad (EO)}, \quad \sigma_o = 0.03 \text{ rad (EC)} \\ \sigma_{\dot{\lambda}} &= 0.001 \text{ rad s}^{-1}, \quad \sigma_m = 0.02 \\ \text{time parameter of the Ornstein-Uhlenbeck process: } \alpha_\eta &= 5 \text{ s}^{-1} \end{aligned}$$

Spinal reflex loop

$$\alpha_E = 12 \text{ rad}^{-1}, \quad \rho = 0.01 \text{ rad}, \quad \mu = 0.1 \text{ s}$$

Muscle-tendon complex

$\tau_m = 15 \text{ ms}$ – time constant of the calcium kinetics low pass filter

$$A = \begin{pmatrix} 10.94 & 1.1 & 0 \\ 0 & 7.43 & 1.2 \\ 0 & 0.94 & 9.10 \end{pmatrix} \text{ N m – muscle activation to torque factor}$$

$$B = \begin{pmatrix} 25 & 2.5137 & 0 \\ 0 & 16.9790 & 2.7422 \\ 0 & 2.1481 & 20.7952 \end{pmatrix} \text{ N m s rad}^{-1} \text{ – passive damping matrix}$$

Biomechanic parameters – 3 joints

Link segment lengths

$$l_1 = 0.4428 \text{ m}, \quad l_2 = 0.4410 \text{ m}, \quad l_3 = 0.7308 \text{ m}$$

Link segment masses

$$m_1 = 7.4400 \text{ kg}, \quad m_2 = 16 \text{ kg}, \quad m_3 = 54.2400 \text{ kg}$$

Link segment center of mass distance to distal joint

$$r_1 = 0.2511 \text{ m}, \quad r_2 = 0.2500 \text{ m}, \quad r_3 = 0.3245 \text{ m}$$

Link segment moments of inertia around media-lateral axis

$$I_1 = 0.1330 \text{ kg m}^2, \quad I_2 = 0.3246 \text{ kg m}^2, \quad I_3 = 3.5860 \text{ kg m}^2$$

Neural dynamics

$$\begin{aligned} \alpha_{\dot{c}} &= 1.5 \text{ rad}^{-1} \text{ s}^{-2}, \quad \alpha_{\ddot{c}} = 0.5 \text{ rad}^{-1} \text{ s}^{-1} \\ \alpha_{\dot{p}} &= 9 \text{ rad}^{-1} \text{ s}^{-2}, \quad \alpha_{\ddot{p}} = 3 \text{ rad}^{-1} \text{ s}^{-1} \\ \alpha_o &= 60 \text{ rad}^{-1} \text{ s}^{-3} \end{aligned}$$

C. Kinematic and biomechanic equations

Abbreviations

$$s_1 = \sin(\theta_1), \quad s_{12} = \sin(\theta_1 + \theta_2), \quad s_{123} = \sin(\theta_1 + \theta_2 + \theta_3)$$
$$c_1 = \cos(\theta_1), \quad c_{12} = \cos(\theta_1 + \theta_2), \quad c_{123} = \cos(\theta_1 + \theta_2 + \theta_3)$$

Head position in anterior-posterior direction

$$p = -l_1 s_1 - l_2 s_{12} - l_3 s_{123},$$

Head position Jacobian

$$J_p = \frac{dp}{d\theta} = (-l_1 c_1 - l_2 c_{12} - l_3 c_{123} \quad -l_1 c_1 - l_2 c_{12} \quad -l_1 c_1)$$

Center of mass position in anterior-posterior direction

$$c = \frac{-m_1 r_1 s_1 - m_2 (l_1 s_1 + r_2 s_{12}) - m_3 (l_1 s_1 + l_2 s_{12} + r_3 s_{123})}{m_1 + m_2 + m_3}$$

Center of mass position Jacobian

$$(J_c)_1 = \frac{-m_1 r_1 c_1 - m_2 (l_1 c_1 + r_2 c_{12}) - m_3 (l_1 c_1 + l_2 c_{12} + r_3 c_{123})}{m_1 + m_2 + m_3}$$
$$(J_c)_2 = \frac{-m_2 r_2 c_{12} - m_3 (l_2 c_{12} + r_3 c_{123})}{m_1 + m_2 + m_3}, \quad (J_c)_3 = \frac{-m_3 r_3 c_{123}}{m_1 + m_2 + m_3}$$

Terms of the equations of motion

$$\begin{aligned}
M_{11} &= I_1 + I_2 + I_3 + m_1 r_1^2 + m_2 (l_1^2 + 2l_1 r_2 c_2 + r_2^2) \\
&\quad + m_3 l_1^2 + m_3 l_2^2 + m_3 r_3^2 + m_3 (2l_1 l_2 c_2 + 2l_1 r_3 c_{23} + 2 * l_2 r_3 c_3) \\
M_{12} &= I_2 + I_3 + m_2 (r_2^2 + l_1 r_2 c_2) + m_3 (l_2^2 + r_3^2 + l_1 l_2 c_2 + l_1 r_3 c_{23} + 2l_2 r_3 c_3) \\
M_{13} &= I_3 + m_3 (r_3^2 + l_1 r_3 c_{23} + l_2 r_3 c_3) \\
M_{22} &= I_2 + I_3 + m_2 r_2^2 + m_3 (l_2^2 + 2l_2 r_3 c_3 + r_3^2) \\
M_{23} &= I_3 + m_3 (r_3^2 + l_2 r_3 c_3) \\
M_{33} &= I_3 + m_3 r_3^2 \\
C_{11} &= \dot{\theta}_2 ((-m_2 l_1 r_2 - m_3 l_1 l_2) s_2 - m_3 l_1 r_3 s_{23}) + \dot{\theta}_3 (-m_3 l_1 r_3 s_{23} - m_3 l_2 r_3 s_3) \\
C_{12} &= (\dot{\theta}_1 + \dot{\theta}_2) ((-m_2 l_1 r_2 - m_3 l_1 l_2) s_2 - m_3 l_1 r_3 s_{23}) + \dot{\theta}_3 (-m_3 l_1 r_3 s_{23} - m_3 l_2 r_3 s_3) \\
C_{13} &= (\dot{\theta}_1 + \dot{\theta}_2 + \dot{\theta}_3) (-m_3 l_1 r_3 s_{23} - m_3 l_2 r_3 s_3) \\
C_{21} &= \dot{\theta}_1 ((m_3 (l_1 l_2 s_2 + l_1 r_3 s_{23})) + l_1 m_2 r_2 s_2) - \dot{\theta}_3 \left(\frac{1}{2} (m_3 (2l_2 r_3 s_3 + l_1 r_3 s_{23})) - (l_1 m_3 r_3 s_{23}) \right) \\
C_{22} &= -l_2 m_3 r_3 \dot{\theta}_3 s_3 \\
C_{23} &= -\dot{\theta}_1 \left(\frac{1}{2} (m_3 (2l_2 r_3 s_3 + l_1 r_3 s_{23})) - (l_1 m_3 r_3 s_{23}) \right) - l_2 m_3 r_3 \dot{\theta}_2 s_3 - l_2 m_3 r_3 \dot{\theta}_3 s_3 \\
C_{31} &= (\dot{\theta}_1 + \dot{\theta}_2) (m_3 l_2 r_3 s_3) + \dot{\theta}_1 (m_3 l_1 r_3 s_{23}) \\
C_{32} &= (\dot{\theta}_1 + \dot{\theta}_2) (m_3 l_2 r_3 s_3) \\
C_{33} &= 0 \\
N_1 &= -g ((m_3 (r_3 s_{123} + l_2 s_1^2 + l_1 s_1)) + (m_2 (r_2 s_1^2 + l_1 s_1)) + (m_1 r_1 s_1)) \\
N_2 &= -g (m_3 (r_3 s_{123} + l_2 s_1^2)) + (m_2 r_2 s_1^2) \\
N_3 &= -g (m_3 r_3 s_{123})
\end{aligned}$$

D. Calculating the time derivative of the manipulator Jacobian

We follow the notation of Murray et al. (1994). The joint twists of a manipulator in reference configuration are $\xi_i, i = 1, \dots, n$. In an arbitrary joint angle configuration $\theta = \theta_1, \dots, \theta_n$, the spatial manipulator Jacobian is given by

$$J^s(\theta) = [\xi_1 \quad \xi_2 \quad \cdots \quad \xi_n], \quad (\text{D.1})$$

with

$$\xi'_i = \text{Ad}_{(e^{\hat{\xi}_1\theta_1} \dots e^{\hat{\xi}_{i-1}\theta_{i-1}})} \xi_i. \quad (\text{D.2})$$

The i -th column of the spatial Jacobian ξ'_i is the i -th joint twist, transformed to the current manipulator configuration.

To calculate the columns of $\dot{J}^s = \frac{d}{dt} J^s$, we have to derive these ξ'_i by time. Using Lemma 2.13 of Murray et al. (1994), we switch to the matrix form of the twist, getting

$$\frac{d}{dt} \left(\widehat{\xi}'_i \right) = \frac{d}{dt} \left((e^{\hat{\xi}_1\theta_1} \dots e^{\hat{\xi}_{i-1}\theta_{i-1}}) \widehat{\xi}_i (e^{\hat{\xi}_1\theta_1} \dots e^{\hat{\xi}_{i-1}\theta_{i-1}})^{-1} \right) \quad (\text{D.3})$$

$$= \sum_{k=1}^{i-1} \left(e^{\hat{\xi}_1\theta_1} \dots e^{\hat{\xi}_{k-1}\theta_{k-1}} \frac{d}{dt} (e^{\hat{\xi}_k\theta_k}) e^{\hat{\xi}_{k+1}\theta_{k+1}} \dots e^{\hat{\xi}_{i-1}\theta_{i-1}} \right. \\ \left. \cdot \widehat{\xi}_i \cdot e^{-\hat{\xi}_{i-1}\theta_{i-1}} \dots e^{-\hat{\xi}_1\theta_1} \right) \quad (\text{D.4})$$

$$+ \sum_{k=1}^{i-1} \left(e^{\hat{\xi}_1\theta_1} \dots e^{\hat{\xi}_{i-1}\theta_{i-1}} \cdot \widehat{\xi}_i \cdot \right. \\ \left. e^{-\hat{\xi}_{i-1}\theta_{i-1}} \dots e^{-\hat{\xi}_{k+1}\theta_{k+1}} \frac{d}{dt} (e^{-\hat{\xi}_k\theta_k}) e^{-\hat{\xi}_{k-1}\theta_{k-1}} \dots e^{-\hat{\xi}_1\theta_1} \right) \\ = \sum_{k=1}^{i-1} \dot{\theta}_k \left(e^{\hat{\xi}_1\theta_1} \dots e^{\hat{\xi}_{k-1}\theta_{k-1}} \widehat{\xi}_k e^{\hat{\xi}_k\theta_k} e^{\hat{\xi}_{k+1}\theta_{k+1}} \dots e^{\hat{\xi}_{i-1}\theta_{i-1}} \right. \\ \left. \cdot \widehat{\xi}_i \cdot e^{-\hat{\xi}_{i-1}\theta_{i-1}} \dots e^{-\hat{\xi}_1\theta_1} \right) \quad (\text{D.5})$$

$$- \sum_{k=1}^{i-1} \dot{\theta}_k \left(e^{\hat{\xi}_1\theta_1} \dots e^{\hat{\xi}_{i-1}\theta_{i-1}} \cdot \widehat{\xi}_i \cdot \right. \\ \left. e^{-\hat{\xi}_{i-1}\theta_{i-1}} \dots e^{-\hat{\xi}_{k+1}\theta_{k+1}} \widehat{\xi}_k e^{-\hat{\xi}_k\theta_k} e^{-\hat{\xi}_{k-1}\theta_{k-1}} \dots e^{-\hat{\xi}_1\theta_1} \right).$$

As the derivative operator commutes with both the \wedge - and the \vee -operator, we can now calculate the i -th column of \dot{J}^s as

$$\frac{d}{dt} \xi'_i = \left(\left(\frac{d}{dt} \xi'_i \right)^\wedge \right)^\vee = \left(\frac{d}{dt} \widehat{\xi}'_i \right)^\vee. \quad (\text{D.6})$$

Knowing the time derivative of the manipulator Jacobian \dot{J}^s allows us to calculate the accelerations of arbitrary points or vectors \mathbf{q} attached to the manipulator (in homogeneous world coordinates). As stated by Murray et al. (1994), the time derivative of \mathbf{q} is given by

$$\mathbf{v} = \frac{d}{dt} \mathbf{q} = \left(J^s \dot{\theta} \right)^\wedge \mathbf{q}. \quad (\text{D.7})$$

In a similar fashion, the second time derivative of \mathbf{q} is given by

$$\mathbf{a} = \frac{d^2}{dt^2} \mathbf{v} = \frac{d}{dt} \left(\left(J^s \dot{\theta} \right)^\wedge \right) \mathbf{q} + \left(J^s \dot{\theta} \right)^\wedge \frac{d}{dt} \mathbf{q} \quad (\text{D.8})$$

$$= \left(J^s \ddot{\theta} + \dot{J}^s \dot{\theta} \right)^\wedge \mathbf{q} + \left(J^s \dot{\theta} \right)^\wedge \mathbf{v}. \quad (\text{D.9})$$

E. Parameters of the robotics movement generation system

Target acquisition

$$\begin{aligned} \alpha_\phi &= 10, & s_{\text{des}} &= 150 \text{ mm s}^{-1}, & \alpha_s &= 15 \\ \alpha_{\mathbf{p}} &= 5, & \alpha_{\mathbf{v}} &= 25 \\ d_1^{(\text{tar})} &= 5 \text{ mm}, & d_2^{(\text{tar})} &= 15 \text{ mm} \end{aligned}$$

Obstacle avoidance

$$\begin{aligned} \alpha_o &= 50 \\ d_1^{(\text{obs})} &= 15 \text{ mm}, & d_2^{(\text{obs})} &= 50 \text{ mm} \\ \psi_1 &= 0.25 \text{ rad}, & \psi_2 &= 1.5 \text{ rad} \end{aligned}$$

Gripper orientation

$$\begin{aligned} \alpha_\gamma &= 10, & \beta_\gamma &= 15, & \rho_{\text{des}} &= -0.3 \text{ rad s}^{-1} \\ d_1^{(\text{ori})} &= \gamma_{\text{des}}, & d_2^{(\text{ori})} &= \gamma_{\text{des}} \pm 0.01 \text{ rad} \end{aligned}$$

Joint limit avoidance

$$\begin{aligned} \alpha_{\text{lim}} &= 10 \text{ (simple)}, & \alpha_{\text{lim}} &= 0.2 \text{ (combined)} \\ a_{\text{lim}} &= 0 \text{ rad}, & b_{\text{lim}} &= 0.2 \text{ rad}, & c^* &= \pm 0.1 \text{ rad} \end{aligned}$$

Homogeneous damping

$$\alpha_{\text{damp}} = 10$$

References

- Adamovich, S. V., Levin, M. F., & Feldman, A. G. (1997). Central Modifications of Reflex Parameters May Underlie the Fastest Arm Movements. *Journal of Neurophysiology*, *77*, 1460–1469.
- Alexandrov, A. V., Frolov, A. A., & Massion, J. (2001, June). Biomechanical analysis of movement strategies in human forward trunk bending. I. Modeling. *Biological Cybernetics*, *84*(6), 425–34. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/11417054>
- Amari, S.-i. (1977). Dynamics of Pattern Formation in Lateral-Inhibition Type Neural Fields. *Biological Cybernetics*, *27*, 77–87.
- Arkin, R. (1989). Motor schema-based mobile robot navigation. *The International Journal of Robotics Research*, *8*(4), 92–112. Retrieved from <http://ijr.sagepub.com/content/8/4/92.short>
- Asai, Y., Tasaka, Y., Nomura, K., Nomura, T., Casadio, M., & Morasso, P. (2009, January). A model of postural control in quiet standing: robust compensation of delay-induced instability using intermittent activation of feedback control. *PloS one*, *4*(7), e6169. Retrieved from <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2704954&tool=pmcentrez&rendertype=abstract>
- Bair, W., Zohary, E., & Newsome, W. T. (2001, March). Correlated firing in macaque visual area MT: time scales and relationship to behavior. *The Journal of Neuroscience*, *21*(5), 1676–97. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/11222658>
- Baker, A. (2002). *Matrix groups: an introduction to Lie group theory*. Springer Verlag.
- Barto, A. G., Fagg, A. H., Sitkoff, N., & Houk, J. C. (1999, April). A cerebellar model of timing and prediction in the control of reaching. *Neural computation*, *11*(3), 565–94. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/10085421>
- Bendat, J. S., & Piersol, A. G. (1993). *Engineering applications of correlation and spectral analysis*. J. Wiley.
- Bernstein, N. (1967). *The co-ordination and regulation of movements*.
- Bicho, E., & Schöner, G. (1997). The dynamic approach to autonomous robotics demonstrated on a low-level vehicle platform. *Robotics and Autonomous Systems*, *21*, 23–35.
- Bizzi, E., & Accornero, N. (1984). Posture control and trajectory formation during arm movement. *The Journal of Neuroscience*, *4*(11), 2738–2744. Retrieved from <http://www.jneurosci.org/content/4/11/2738.short>
- Blanpied, P., & Smidt, G. L. (1992, January). Human plantarflexor stiffness to multiple single-stretch trials. *Journal of Biomechanics*, *25*(1), 29–39. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/1733982>
- Bottaro, A., Yasutake, Y., Nomura, T., Casadio, M., & Morasso, P. (2008, June). Bounded stability of the quiet standing posture: an intermittent control model.

-
- Human Movement Science*, 27(3), 473–95. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/18342382>
- Brown, I. E., Scott, S. H., & Loeb, G. E. (1996). Mechanics of feline soleus: II. Design and validation of a mathematical model. *Journal of Muscle Research and Cell Motility*, 17(2), 221–233. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/8793724>
- Churchland, M. M., Afshar, A., & Shenoy, K. V. (2006). A central source of movement variability. *Neuron*, 52(6), 1085–1096.
- Collins, J. J., & De Luca, C. (1994). Random walking during quiet standing. *Physical Review Letters*. Retrieved from <http://link.aps.org/doi/10.1103/PhysRevLett.73.764>
- Creath, R., Kiemel, T., Horak, F. B., Peterka, R. J., & Jeka, J. J. (2005, March). A unified view of quiet and perturbed stance: simultaneous co-existing excitable modes. *Neuroscience Letters*, 377(2), 75–80. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/15740840>
- Cusumano, J. P., & Cesari, P. (2006). Body-goal variability mapping in an aiming task. *Biological Cybernetics*, 94(5), 367–379.
- de Freitas, S., Scholz, J. P., & Stehman, A. (2007). Effect of Motor Planning on Use of Motor Abundance. *Neuroscience Letters*, 417(1), 66–71. Retrieved from <http://www.ncbi.nlm.nih.gov/pmc/articles/PMC1950341/>
- de Lussanet, M. H. E., Smeets, J. B. J., & Brenner, E. (2002, April). Relative damping improves linear mass-spring models of goal-directed movements. *Human Movement Science*, 21(1), 85–100. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/11983435>
- Faisal, A. A., Selen, L. P. J., & Wolpert, D. M. (2008, April). Noise in the nervous system. *Nature Reviews. Neuroscience*, 9(4), 292–303. Retrieved from <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2631351&tool=pmcentrez&rendertype=abstract>
- Fajen, B. R., Warren, W. H., Temizer, S., & Kaelbling, L. P. (2003, October). A neural model of visually guided steering, obstacle avoidance, and route selection. *International Journal of Computer Vision*, 54, 13–34. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/19803653>
- Feldman, A. G. (1972). The Influence of Tonic Different Stretch Descending Reflex Systems on the in the Cat. *Experimental Neurology*, 37, 481–494.
- Feldman, A. G. (1986). Once more on the equilibrium-point hypothesis (lambda model) for motor control. *Journal of Motor Behavior*, 18(1), 17–54. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/15136283>
- Feldman, A. G., & Levin, M. (1995). The origin and use of positional frames of reference in motor control. *Behavioral and Brain Sciences*, 18(4), 723–744. Retrieved from <http://e.guigon.free.fr/rsc/article/FeldmanLevin95.pdf>
- Fitts, P. M. (1954, September). The information capacity of the human motor system in controlling the amplitude of movement. *Journal of Experimental Psychology*, 47(6), 381–391. Retrieved from <http://www.ncbi.nlm.nih.gov/>

-
- pubmed/1402698
- Fitzpatrick, R., & McCloskey, D. I. (1994, July). Proprioceptive, visual and vestibular thresholds for the perception of sway during standing in humans. *The Journal of Physiology*, *478* (Pt 1), 173–86. Retrieved from <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=1155655&tool=pmcentrez&rendertype=abstract>
- Flash, T. (1987). The control of hand equilibrium trajectories in multi-joint arm movements. *Biological Cybernetics*, *27*(4), 257–274. Retrieved from <http://link.springer.com/article/10.1007/BF00338819>
- Flash, T., & Hogan, N. (1985). The Coordination of Arm Movements: An Experimentally Confirmed Mathematical Model. *Journal of Neuroscience*, *5*(7), 1688–1703.
- Gawthrop, P. J., Loram, I. D., Lakie, M., & Gollee, H. (2011, February). Intermittent control: a computational theory of human control. *Biological Cybernetics*, *104*(1-2), 31–51. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/21327829>
- Gomi, H., & Kawato. (1996, April). Equilibrium-point control hypothesis examined by measured arm stiffness during multijoint movement. *Science (New York, N.Y.)*, *272*(5258), 117–20. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/8600521>
- Goodwin, G. M., McCloskey, D. I., & Matthews, P. B. C. (1972). The contribution of muscle afferents to kinaesthesia shown by vibration induced illusions of movement and by the effects of paralysing joint afferents. *Brain*, *95*, 705–748.
- Greville, T. (1959). The Pseudoinverse of a Rectangular or Singular Matrix and its Application to the Solution of Systems of Linear Equations. *SIAM review*, *1*(1), 38–43. Retrieved from <http://epubs.siam.org/doi/abs/10.1137/1001003>
- Gribble, P. L., Ostry, D. J., Sanguineti, V., & Laboissière, R. (1998, March). Are complex control signals required for human arm movement? *Journal of Neurophysiology*, *79*(3), 1409–24. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/9497421>
- Harris, C., & Wolpert, D. (1998). Signal-dependent noise determines motor planning. *Nature*, *394*(August), 780–784. Retrieved from <http://www.nature.com/nature/journal/v394/n6695/abs/394780a0.html>
- Hasson, C. J., Miller, R. H., & Caldwell, G. E. (2011, January). Contractile and elastic ankle joint muscular properties in young and older adults. *PloS one*, *6*(1), e15953. Retrieved from <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3019216&tool=pmcentrez&rendertype=abstract>
- Hatsopoulos, N. G. (1994, January). Is a virtual trajectory necessary in reaching movements? *Biological Cybernetics*, *70*(6), 541–51. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/8068769>
- Hof, A. L. (1998, September). In vivo measurement of the series elasticity release curve of human triceps surae muscle. *Journal of Biomechanics*, *31*(9), 793–800. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/9802779>

-
- Hoffmann, H., Pastor, P., Park, D.-h., & Schaal, S. (2009). Biologically-inspired dynamical systems for movement generation: automatic real-time goal adaptation and obstacle avoidance. In *International Conference on Robotics and Automation (ICRA)* (pp. 2587–2592).
- Hogan, N. (1984). An organizing principle for a class of voluntary movements. *Journal of Neuroscience*, *4*(11), 2745–2754.
- Hogan, N., Sternad, D., Park, S.-W., & Mu, H. (2010). Coordinate Dependence of Variability Analysis. *PLoS Computational Biology*, *6*(4).
- Hollerbach, J. M., & Suh, K. C. (1987). Redundancy resolution of manipulation through torque optimization. *IEEE Transactions on Robotics and Automation*, *3*, 308–316.
- Horak, F., & Nashner, L. M. (1986). Central programming of postural movements: adaptation to altered support-surface configurations. *Journal of Neurophysiology*, 1369–1381. Retrieved from <http://jn.physiology.org/content/55/6/1369.short>
- Howell, D. C. (2010). *Statistical methods for psychology*. Cengage Learning.
- Hsu, W.-l., Scholz, J. P., Schöner, G., Jeka, J. J., & Kiemel, T. (2007). Control and Estimation of Posture During Quiet Stance Depends on Multijoint Coordination Control and Estimation of Posture During Quiet Stance Depends on Multijoint Coordination. *Journal of Neurophysiology*, *97*, 3024–3035.
- Ijspeert, A. J., Nakanishi, J., & Schaal, S. (2002). Learning attractor landscapes for learning motor primitives. In *Advances in Neural Information Processing Systems (NIPS)* (pp. 1547–1554).
- Iossifidis, I., Bruckhoff, C., Theis, C., Grote, C., Faubel, C., & Schöner, G. (2002). Cora: An anthropomorphic robot assistant for human environment. In *Robot and human interactive communication (ro-man)* (pp. 392–398). Retrieved from http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=1045654
- Iossifidis, I., & Schöner, G. (2006). Dynamical Systems Approach for the Autonomous Avoidance of Obstacles and Joint-limits for an Redundant Robot Arm. In *Intelligent Robots and Systems (IROS)* (pp. 580–585).
- Joyce, G. C., Rack, P. M. H., & Westbury, D. R. (1969). The mechanical properties of cat soleus muscle during controlled lengthening and shortening movements. *Journal of Physiology*, *204*(2), 461–474. Retrieved from <http://jp.physoc.org/cgi/content/abstract/204/2/461>
- Kandel, E., Schwartz, J., Jessell, T., Siegelbaum, S., & Hudspeth, A. J. (2012). *Principles of Neural Science, Fifth Edition*. McGraw-Hill Education.
- Kang, N., Shinohara, M., Zatsiorsky, V. M., & Latash, M. L. (2004, August). Learning multi-finger synergies: an uncontrolled manifold analysis. *Experimental Brain Research*, *157*(3), 336–50. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/15042264>
- Khansari-Zadeh, S. M., & Billard, A. (2012, March). A dynamical system approach to realtime obstacle avoidance. *Autonomous Robots*, *32*(4), 433–454. Retrieved from <http://link.springer.com/10.1007/s10514-012-9287-y>
- Khansari-Zadeh, S. M., Kronander, K., & Billard, A. (2012, December). Learning

-
- to Play Minigolf: A Dynamical System-Based Approach. *Advanced Robotics*, 26(17), 1967–1993.
- Khatib, O. (1986). Real-time obstacle avoidance for manipulators and mobile robots. *International Journal of Robotics Research*, 5(1), 90–98.
- Khatib, O. (1995, February). Inertial Properties in Robotic Manipulation: An Object-Level Framework. *The International Journal of Robotics Research*, 14(1), 19–36. Retrieved from <http://ijr.sagepub.com/cgi/doi/10.1177/027836499501400103>
- Kiemel, T., Oie, K. S., & Jeka, J. J. (2002, October). Multisensory fusion and the stochastic structure of postural sway. *Biological Cybernetics*, 87(4), 262–77. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/12386742>
- Kiemel, T., Oie, K. S., & Jeka, J. J. (2006, March). Slow dynamics of postural sway are in the feedback loop. *Journal of Neurophysiology*, 95(3), 1410–8. Retrieved from <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2717851&tool=pmcentrez&rendertype=abstract>
- Kiemel, T., Zhang, Y., & Jeka, J. J. (2011, January). Visual flow is interpreted relative to multisegment postural control. *Journal of Motor Behavior*, 43(3), 237–46. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/21534025>
- Kistemaker, D. A., Soest, A. K. J. V., & Bobbert, M. F. (2006). Is Equilibrium Point Control Feasible for Fast Goal-Directed Single-Joint Movements? *Journal of Neurophysiology*, 95, 2898–2912.
- Kistemaker, D. A., Van Soest, A. K. J., & Bobbert, M. F. (2007, September). Equilibrium point control cannot be refuted by experimental reconstruction of equilibrium point trajectories. *Journal of Neurophysiology*, 98(3), 1075–82. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/17615122>
- Kistemaker, D. A., Van Soest, A. K. J., Wong, J. D., Kurtzer, I., & Gribble, P. L. (2013, February). Control of position and movement is simplified by combined muscle spindle and Golgi tendon organ feedback. *Journal of Neurophysiology*, 109(4), 1126–39. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/23100138>
- Kluzik, J., Horak, F. B., & Peterka, R. J. (2005, May). Differences in preferred reference frames for postural orientation shown by after-effects of stance on an inclined surface. *Experimental Brain Research*, 162(4), 474–89. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/15654594>
- Krishnamoorthy, V., Yang, J.-F., & Scholz, J. P. (2005, July). Joint coordination during quiet stance: effects of vision. *Experimental Brain Research*, 164(1), 1–17. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/15841397>
- Kron, A. (1997). *Comparaison dynamique et électromyographique de l'homme en posture debout d'un modèle de pendule inverse sous contrôle lambda*. Master thesis, Université de Sherbrooke.
- Kuo, A. D. (2005, September). An optimal state estimation model of sensory integration in human postural balance. *Journal of Neural Engineering*, 2(3), S235–49. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/16135887>
- Laboissière, R., Ostry, D. J., & Feldman, A. G. (1996). The control of multi-muscle

-
- systems: human jaw and hyoid movements. *Biological Cybernetics*(74), 373–384.
- Lacquaniti, F. (1989). Central representations of human limb movement as revealed by studies of drawing and handwriting. *Trends in Neurosciences*, 12(8), 287–291. Retrieved from <http://www.sciencedirect.com/science/article/pii/0166223689900088>
- Lacquaniti, F., Terzuolo, C., & Viviani, P. (1983). The law relating the kinematic and figural aspects of drawing movements. *Acta Psychologica*, 54, 115–130.
- Lakie, M., Caplan, N., & Loram, I. D. (2003, August). Human balancing of an inverted pendulum with a compliant linkage: neural control by anticipatory intermittent bias. *The Journal of physiology*, 551(Pt 1), 357–70. Retrieved from <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2343154&tool=pmcentrez&rendertype=abstract>
- Lakie, M., & Loram, I. D. (2006, November). Manually controlled human balancing using visual, vestibular and proprioceptive senses involves a common, low frequency neural process. *The Journal of Physiology*, 577(Pt 1), 403–16. Retrieved from <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2000668&tool=pmcentrez&rendertype=abstract>
- Lánský, P., & Sacerdote, L. (2001). The Ornstein–Uhlenbeck neuronal model with signal-dependent noise. *Physics Letters A*, 285(July), 132–140.
- Latash, M. L. (2008). *Neurophysiological basis of movement*. Human Kinetics.
- Loeb, G. E. (2012, December). Optimal isn't good enough. *Biological Cybernetics*, 106(11-12), 757–65. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/22895830>
- Loeb, G. E., Brown, E., & Cheng, J. (1999). A hierarchical foundation for models of sensorimotor control. *Experimental Brain Research*, 126, 1–18.
- Loram, I. D., Gollee, H., Lakie, M., & Gawthrop, P. J. (2011, January). Human control of an inverted pendulum: is continuous control necessary? Is intermittent control effective? Is intermittent control physiological? *The Journal of Physiology*, 589(Pt 2), 307–24. Retrieved from <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3043535&tool=pmcentrez&rendertype=abstract>
- Loram, I. D., & Lakie, M. (2002, November). Direct measurement of human ankle stiffness during quiet standing: the intrinsic mechanical stiffness is insufficient for stability. *The Journal of Physiology*, 545(3), 1041–1053. Retrieved from <http://www.jphysiol.org/cgi/doi/10.1113/jphysiol.2002.025049>
- Maciejewski, A. A., & Klein, C. A. (1985). Obstacle Avoidance for Kinetically Redundant Manipulators in Dynamically Varying Environments. *International Journal of Robotics Research*, 4(3), 109–117.
- Martin, V., Scholz, J. P., & Schöner, G. (2009). Redundancy, Self-Motion, and Motor Control. *Neural Computation*, 21(5), 1371–1414. Retrieved from <http://www.mitpressjournals.org/doi/pdf/10.1162/neco.2008.01-08-698>
- Masani, K., Vette, A. H., & Popovic, M. R. (2006). Controlling balance during quiet standing: Proportional and derivative controller generates preceding mo-

-
- tor command to body sway position observed in experiments. *Gait and Posture*, 23, 164–172.
- Matthews, P. B. C. (1994, December). The simple frequency response of human stretch reflexes in which either short- or long-latency components predominate. *The Journal of Physiology*, 481(Pt 3), 777–98. Retrieved from <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=1155918&tool=pmcentrez&rendertype=abstract>
- Matthews, P. B. C. (1996, April). Relationship of firing intervals of human motor units to the trajectory of post-spike after-hyperpolarization and synaptic noise. *The Journal of Physiology*, 492(Pt 2), 597–628. Retrieved from <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=1158851&tool=pmcentrez&rendertype=abstract>
- Maurer, C., Mergner, T., & Peterka, R. J. (2006, May). Multisensory control of human upright stance. *Experimental Brain Research*, 171(2), 231–50. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/16307252>
- Maurer, C., & Peterka, R. J. (2005, January). A new interpretation of spontaneous sway measures based on a simple model of human postural control. *Journal of Neurophysiology*, 93(1), 189–200. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/15331614>
- McLean, A., & Cameron, S. (1995, August). The Virtual Springs Method: Path Planning and Collision Avoidance for Redundant Manipulators. *The International Journal of Robotics Research*, 15(4), 300–319. Retrieved from <http://ijr.sagepub.com/cgi/doi/10.1177/027836499601500401>
- Mergner, T. (2010, December). A neurological view on reactive human stance control. *Annual Reviews in Control*, 34(2), 177–198. Retrieved from <http://linkinghub.elsevier.com/retrieve/pii/S1367578810000404>
- Micheau, P., Kron, A., & Bourassa, P. (2003, September). Evaluation of the lambda model for human postural control during ankle strategy. *Biological Cybernetics*, 89(3), 227–36. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/14504941>
- Mileusnic, M. P., Brown, I. E., Lan, N., & Loeb, G. E. (2006, October). Mathematical models of proprioceptors. I. Control and transduction in the muscle spindle. *Journal of Neurophysiology*, 96(4), 1772–88. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/16672301>
- Mirbagheri, M., Barbeau, H., & Kearney, R. (2000, December). Intrinsic and reflex contributions to human ankle stiffness: variation with activation level and position. *Experimental Brain Research*, 135(4), 423–436. Retrieved from <http://www.springerlink.com/openurl.asp?genre=article&id=doi:10.1007/s002210000534>
- Morasso, P. (1981). Spatial control of arm movements. *Experimental Brain Research*. Retrieved from <http://link.springer.com/article/10.1007/BF00236911>
- Morasso, P., & Sanguineti, V. (2002, October). Ankle muscle stiffness alone cannot stabilize balance during quiet standing. *Journal of Neurophysiology*, 88(4), 2157–62. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/12364538>

-
- Morasso, P., & Schieppati, M. (1999, September). Can muscle stiffness alone stabilize upright standing? *Journal of Neurophysiology*, *82*(3), 1622–6. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/10482776>
- Muirhead, R. (1982). *Aspects of multivariate statistical theory*. Wiley.
- Müller, H., & Sternad, D. (2003, July). A randomization method for the calculation of covariation in multiple nonlinear relations: illustrated with the example of goal-directed movements. *Biological Cybernetics*, *89*(1), 22–33. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/12836030>
- Müller, H., & Sternad, D. (2004, February). Decomposition of variability in the execution of goal-oriented tasks: three components of skill improvement. *Journal of Experimental Psychology: Human Perception and Performance*, *30*(1), 212–33. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/14769078>
- Mülling, K., Kober, J., Kroemer, O., & Peters, J. (2012). Learning to select and generalize striking movements in robot table tennis. *International Journal of Robotics Research*, *32*(3), 263–279. Retrieved from <http://www.aaai.org/ocs/index.php/FSS/FSS12/paper/download/5602/5884>
- Murray, R. M., Li, Z., & Sastry, S. S. (1994). *A Mathematical Introduction to Robotic Manipulation*. CRC PressINC.
- Nakanishi, J., Morimoto, J., Endo, G., Cheng, G., Schaal, S., & Kawato, M. (2004, June). Learning from demonstration and adaptation of biped locomotion. *Robotics and Autonomous Systems*, *47*(2-3), 79–91. Retrieved from <http://linkinghub.elsevier.com/retrieve/pii/S0921889004000399>
- Nashner, L. M., Black, F., & Wall, C. (1982). Adaptation to altered support and visual conditions during stance: patients with vestibular deficits. *The Journal of Neuroscience*, 536–544. Retrieved from <http://www.jneurosci.org/content/2/5/536.short>
- Nashner, L. M., & McCollum, G. (1985). The organization of human postural movements: A formal basis and experimental synthesis. *Behavioral and Brain Sciences*, *8*(01), 135–150. Retrieved from http://journals.cambridge.org/article_S0140525X00020008
- Ostry, D. J., Gribble, P. L., Levin, M. F., & Feldman, A. G. (1997, September). Phasic and tonic stretch reflexes in muscles with few muscle spindles: human jaw-opener muscles. *Experimental Brain Research*, *116*(2), 299–308. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/9348128>
- Pai, Y.-C., Rogers, M. W., Patton, J., Cain, T. D., & Hanke, T. A. (1998). Static versus dynamic predictions of protective stepping following waist—pull perturbations in young and older adults. *Journal of Biomechanics*, *31*, 1111–1118.
- Perko, L. (1991). *Differential Equations and Dynamical Systems*. Heidelberg: Springer.
- Peterka, R. J. (2000, April). Postural control model interpretation of stabilogram diffusion analysis. *Biological Cybernetics*, *82*(4), 335–43. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/10804065>
- Peterka, R. J. (2002, September). Sensorimotor integration in human postural control. *Journal of Neurophysiology*, *88*(3), 1097–1118. Retrieved from

-
- <http://www.ncbi.nlm.nih.gov/pubmed/13679407>
- Pilon, J.-F., & Feldman, A. G. (2006, September). Threshold control of motor actions prevents destabilizing effects of proprioceptive delays. *Experimental Brain Research*, *174*(2), 229–39. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/16676171>
- Polónyová, A., & Hlavacka, F. (2001, January). Human postural responses to different frequency vibrations of lower leg muscles. *Physiological Research*, *50*(4), 405–10. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/11551147>
- Qu, X., Nussbaum, M. A., & Madigan, M. L. (2007, January). A balance control model of quiet upright stance based on an optimal control strategy. *Journal of Biomechanics*, *40*(16), 3590–7. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/17628566>
- Raphael, G., Tsianos, G. A., & Loeb, G. E. (2010, July). Spinal-like regulator facilitates control of a two-degree-of-freedom wrist. *The Journal of Neuroscience*, *30*(28), 9431–44. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/20631172>
- Reimann, H., Iossifidis, I., & Schöner, G. (2010a). Generating collision free reaching movements for redundant manipulators using dynamical systems. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (pp. 5372–5379).
- Reimann, H., Iossifidis, I., & Schöner, G. (2010b). Integrating orientation constraints into the attractor dynamics approach for autonomous manipulation. In *IEEE-RAS International Conference on Humanoid Robots* (pp. 294–301).
- Reimann, H., Iossifidis, I., & Schöner, G. (2011). Autonomous movement generation for manipulators with multiple simultaneous constraints using the attractor dynamics approach. In *IEEE International Conference on Robotics and Automation (ICRA)* (pp. 5470–5477).
- Ricciardi, L., & Sacerdote, L. (1979). The Ornstein-Uhlenbeck Process as a Model for Neuronal Activity. *Biological Cybernetics*, *9*, 1–9. Retrieved from <http://link.springer.com/article/10.1007/BF01845839>
- Richter, M. (2012). A robotic architecture for action selection and behavioral organization inspired by human cognition. *IEEE/RSJ International Conference on Intelligent Robots and Systems*. Retrieved from http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=6386153
- Riener, R., & Edrich, T. (1999, May). Identification of passive elastic joint moments in the lower extremities. *Journal of Biomechanics*, *32*(5), 539–44. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/10327008>
- Rimon, E., & Koditschek, D. (1992). Exact robot navigation using artificial potential functions. *IEEE Transactions on Robotics and Automation*, *8*(5), 501–518. Retrieved from <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=163777>
- Rothwell, J. C. (1987). *Control of human voluntary movement*. Croom Helm.
- Rozendaal, L. A., & Van Soest, A. K. J. (2008, October). Stabilization of a multi-segment model of bipedal standing by local joint control overestimates the

-
- required ankle stiffness. *Gait and Posture*, 28(3), 525–7. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/18403206>
- Runge, C. F., Shupert, C. L., Horak, F. B., & Zajac, F. E. (1999, October). Ankle and hip postural strategies defined by joint torques. *Gait and Posture*, 10(2), 161–70. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/10502650>
- Saffer, M., Kiemel, T., & Jeka, J. (2008). Coherence analysis of muscle activity during quiet stance. *Experimental Brain Research*, 185(2), 215–226. Retrieved from <http://link.springer.com/article/10.1007/s00221-007-1145-3>
- Saltzman, E. (1979). Levels of sensorimotor representation. *Journal of Mathematical Psychology*, 163, 91–163. Retrieved from <http://www.sciencedirect.com/science/article/pii/0022249679900208>
- Saltzman, E., & Kelso, J. (1987). Skilled actions: A task dynamic approach. *Psychological Review*. Retrieved from http://www.haskins.yale.edu/sr/sr076/SR076_01.pdf
- Sandamirskaya, Y., & Schöner, G. (2010). Serial order in an acting system: a multi-dimensional dynamic neural fields implementation. In *International Conference on Development and Learning*.
- Schaal, S. (1997). Learning from demonstration. In M. Mozer, M. Jordan, & T. Petsche (Eds.), *Advances in Neural Information Processing Systems*.
- Scholz, J. P., & Schöner, G. (1999, June). The uncontrolled manifold concept: identifying control variables for a functional task. *Experimental Brain Research*, 126(3), 289–306. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/10382616>
- Scholz, J. P., Schöner, G., Hsu, W.-l., Jeka, J. J., Horak, F. B., & Martin, V. (2007, June). Motor equivalent control of the center of mass in response to support surface perturbations. *Experimental Brain Research*, 180(1), 163–79. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/17256165>
- Scholz, J. P., Schöner, G., & Latash, M. L. (2000, November). Identifying the control structure of multijoint coordination during pistol shooting. *Experimental Brain Research*, 135(3), 382–404. Retrieved from <http://www.springerlink.com/openurl.asp?genre=article&id=doi:10.1007/s002210000540>
- Schöner, G. (1990). A dynamic theory of coordination of discrete movement. *Biological Cybernetics*, 270, 257–270. Retrieved from <http://link.springer.com/article/10.1007/BF00203449>
- Schöner, G. (1995). Recent developments and problems in human movement science and their conceptual implications. *Ecological Psychology*.
- Schöner, G. (2008). Dynamical Systems Approaches to Cognition. In R. Sun (Ed.), *The Cambridge Handbook of Computational Psychology*. Cambridge University Press.
- Schöner, G., Dose, M., & Engels, C. (1995). Dynamics of behavior: theory and applications for autonomous robot architectures. *Robotics and Autonomous Systems*, 16, 213–245. Retrieved from <http://www.sciencedirect.com/science/article/pii/0921889095000496>
- Schöner, G., & Kelso, J. a. (1988, March). Dynamic pattern generation in behavioral

-
- and neural systems. *Science*, 239(4847), 1513–20. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/3281253>
- Schöner, G., & Scholz, J. P. (2007, July). Analyzing variance in multi-degree-of-freedom movements: uncovering structure versus extracting correlations. *Motor Control*, 11(3), 259–75. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/17715459>
- Seraji, H. (1989). Configuration control of redundant manipulators: Theory and implementation. *Robotics and Automation, IEEE Transactions on*, 5(4), 472–490. Retrieved from http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=88062
- Shadlen, M. N., & Newsome, W. T. (1994, August). Noise, neural codes and cortical organization. *Current Opinion in Neurobiology*, 4(4), 569–79. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/7812147>
- Shadmehr, R., & Arbib, M. A. (1992). A mathematical analysis of the force-stiffness characteristics of muscles in control of a single joint system. *Biological Cybernetics*, 66, 463–477.
- Shukla, A., & Billard, A. (2012, March). Coupled dynamical system based arm–hand grasping model for learning fast adaptation strategies. *Robotics and Autonomous Systems*, 60(3), 424–440.
- Siciliano, B. (1990). Kinematic control of redundant robot manipulators: A tutorial. *Journal of Intelligent and Robotic Systems*, 3(3), 201–212. Retrieved from <http://www.springerlink.com/index/10.1007/BF00126069>
- Siciliano, B., & Khatib, O. (2008). *Springer Handbook of Robotics*. Springer.
- Silder, A., Heiderscheit, B., & Thelen, D. G. (2008, January). Active and passive contributions to joint kinetics during walking in older adults. *Journal of Biomechanics*, 41(7), 1520–7. Retrieved from <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2702713&tool=pmcentrez&rendertype=abstract>
- Smith, P. L. (2010, April). From Poisson shot noise to the integrated Ornstein–Uhlenbeck process: Neurally principled models of information accumulation in decision-making and response time. *Journal of Mathematical Psychology*, 54(2), 266–283. Retrieved from <http://linkinghub.elsevier.com/retrieve/pii/S0022249609001576>
- Smith, P. L., & Ratcliff, R. (2004, March). Psychology and neurobiology of simple decisions. *Trends in Neurosciences*, 27(3), 161–8. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/15036882>
- Soechting, J., & Lacquaniti, F. (1981). Invariant characteristics of a pointing movement in man. *The Journal of Neuroscience*. Retrieved from <http://www.jneurosci.org/content/1/7/710.short>
- St-Onge, N., Adamovich, S. V., & Feldman, A. G. (1997, July). Control processes underlying elbow flexion movements may be independent of kinematic and electromyographic patterns: experimental study and modelling. *Neuroscience*, 79(1), 295–316. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/9178885>

-
- Stulp, F., Oztop, E., Pastor, P., Beetz, M., & Schaal, S. (2009, December). Compact models of motor primitive variations for predictable reaching and obstacle avoidance. *2009 9th IEEE-RAS International Conference on Humanoid Robots*, 589–595. Retrieved from <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=5379551>
- Taylor, J. L. (2009). Proprioception. In M. Binder, N. Kirokawa, & U. Windhorst (Eds.), *Encyclopedia of Neuroscience* (pp. 1143–1149). Springer.
- Tee, K. P., Burdet, E., Chew, C. M., & Milner, T. E. (2004, May). A model of force and impedance in human arm movements. *Biological Cybernetics*, *90*(5), 368–75. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/15221397>
- Todorov, E. (2004, September). Optimality principles in sensorimotor control. *Nature Neuroscience*, *7*(9), 907–15. Retrieved from <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=1488877&tool=pmcentrez&rendertype=abstract>
- Todorov, E., & Jordan, M. I. (2002, November). Optimal feedback control as a theory of motor coordination. *Nature Neuroscience*, *5*(11), 1226–35. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/12404008>
- Tseng, Y.-W., Scholz, J. P., & Galloway, J. C. (2009, February). The organization of intralimb and interlimb synergies in response to different joint dynamics. *Experimental Brain Research*, *193*(2), 239–54. Retrieved from <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3122082&tool=pmcentrez&rendertype=abstract>
- Tseng, Y.-W., Scholz, J. P., Schöner, G., & Hotchkiss, L. (2003, April). Effect of accuracy constraint on joint coordination during pointing movements. *Experimental Brain Research*, *149*(3), 276–88. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/12632230>
- Tsodyks, M., Kenet, T., Grinvald, A., & Arieli, A. (1999, December). Linking spontaneous activity of single cortical neurons and the underlying functional architecture. *Science (New York, N.Y.)*, *286*(5446), 1943–6. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/10583955>
- Uno, Y., Kawato, M., & Suzuki, R. (1989). Formation and control of optimal trajectory in human multijoint arm movement. *Biological Cybernetics*, *101*, 89–101. Retrieved from <http://link.springer.com/article/10.1007/BF00204593>
- Valero-Cuevas, F. J., Venkadesan, M., & Todorov, E. (2009). Structured Variability of Muscle Activations Supports the Minimal Intervention Principle of Motor Control. *Journal of Neurophysiology*, 59–68.
- Van Soest, A. K. J., & Bobbert, M. F. (1993). The contribution of muscle properties in the control of explosive movements. *Biological Cybernetics*, *69*, 195–204.
- Van Soest, A. K. J., Bobbert, M. F., & van Ingen Schenau, G. J. (1994). A Control Strategy for the Execution of Explosive Movements from Varying Starting Positions. *Journal of Neurophysiology*, *71*(4), 1390–1402.
- Van Soest, A. K. J., Haenen, W. P., & Rozendaal, L. A. (2003, April). Stability of bipedal stance: the contribution of cocontraction and spindle feedback. *Biological Cybernetics*, *88*(4), 293–301. Retrieved from <http://www.ncbi.nlm.nih>

-
- .gov/pubmed/12690488
- van der Kooij, H., & de Vlugt, E. (2007, August). Postural responses evoked by platform perturbations are dominated by continuous feedback. *Journal of Neurophysiology*, *98*(2), 730–43. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/17460106>
- van der Kooij, H., Jacobs, R., Koopman, B., & Grootenboer, H. (1999, May). A multisensory integration model of human stance control. *Biological Cybernetics*, *80*(5), 299–308. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/10365423>
- Veerman, M. M., Brenner, E., & Smeets, J. B. J. (2008, May). The latency for correcting a movement depends on the visual attribute that defines the target. *Experimental Brain Research*, *187*(2), 219–28. Retrieved from <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2335293&tool=pmcentrez&rendertype=abstract>
- Verrel, J. (2010, August). Distributional properties and variance-stabilizing transformations for measures of uncontrolled manifold effects. *Journal of Neuroscience Methods*, *191*(2), 166–70. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/20599556>
- Verrel, J. (2011, September). A formal and data-based comparison of measures of motor-equivalent covariation. *Journal of Neuroscience Methods*, *200*(2), 199–206. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/21513738>
- Verrel, J., Lövdén, M., & Lindenberger, U. (2012, January). Normal aging reduces motor synergies in manual pointing. *Neurobiology of Aging*, *33*(1), 200.e1–10. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/20724037>
- Verrel, J., Pradon, D., & Vuillerme, N. (2012, January). Persistence of motor-equivalent postural fluctuations during bipedal quiet standing. *PloS one*, *7*(10), e48312. Retrieved from <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3482199&tool=pmcentrez&rendertype=abstract>
- Ward, S. R., Eng, C. M., Smallwood, L. H., & Lieber, R. L. (2009, April). Are current measurements of lower extremity muscle architecture accurate? *Clinical Orthopaedics and Related Research*, *467*(4), 1074–82. Retrieved from <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2650051&tool=pmcentrez&rendertype=abstract>
- Weiss, P. L., Hunter, I. W., & Kearney, R. E. (1988, January). Human ankle joint stiffness over the full range of muscle activation levels. *Journal of Biomechanics*, *21*(7), 539–44. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/3410857>
- Whitney, D. E. (1972). The Mathematics of Coordinated Control of Prosthetic Arms and Manipulators. In *Conference on manual control* (pp. 207–220).
- Whittington, B., Silder, A., Heiderscheid, B., & Thelen, D. G. (2008, May). The contribution of passive-elastic mechanisms to lower extremity joint kinetics during human walking. *Gait and Posture*, *27*(4), 628–34. Retrieved from <http://www.pubmedcentral.nih.gov/articlerender.fcgi>

-
- ?artid=2505349&tool=pmcentrez&rendertype=abstract
- Wing, A. M., Johannsen, L., & Endo, S. (2011, November). Light touch for balance: influence of a time-varying external driving signal. *Philosophical Transactions of the Royal Society of London. Series B, Biological sciences*, 366(1581), 3133–41. Retrieved from <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3172607&tool=pmcentrez&rendertype=abstract>
- Winter, D. A. (1990). *Biomechanics and Motor Control of Human Movement*. Wiley.
- Winter, D. A., Patla, A. E., Prince, F., Ishac, M., & Gielo-Perczak, K. (1998, September). Stiffness control of balance in quiet standing. *Journal of Neurophysiology*, 80(3), 1211–21. Retrieved from <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3007648&tool=pmcentrez&rendertype=abstract>
- Winters, J. M., & Woo, S. L. Y. (1990). *Multiple muscle systems: biomechanics and movement organization*. Springer-Verlag.
- Yen, J. T., Auyang, A. G., & Chang, Y.-H. (2009, July). Joint-level kinetic redundancy is exploited to control limb-level forces during human hopping. *Experimental Brain Research*, 196(3), 439–51. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/19495732>
- Yen, J. T., & Chang, Y.-H. (2010, May). Rate-dependent control strategies stabilize limb forces during human locomotion. *Journal of the Royal Society Interface*, 7(46), 801–10. Retrieved from <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2874235&tool=pmcentrez&rendertype=abstract>
- Zatsiorsky, V. M., & Duarte, M. (1999). Instant Equilibrium Point and its Migration in Standing Tasks: Rambling and Trembling Components of the Stabilogram. *Motor Control*, 3, 28–38.
- Zhang, Y., Kiemel, T., & Jeka, J. (2007, July). The influence of sensory information on two-component coordination during quiet stance. *Gait and Posture*, 26(2), 263–71. Retrieved from <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2041859&tool=pmcentrez&rendertype=abstract>
- Zohary, E., Shadlen, M. N., & Newsome, W. T. (1994). Correlated neuronal discharge rate and its implications for psychophysical performance. *Nature*, 370, 140–143. Retrieved from <http://www.nature.com/nature/journal/v370/n6485/abs/370140a0.html>

Lebenslauf

Persönliche Daten

Name	Hendrik Reimann
Geburtsdatum	20. November 1978
Geburtsort	Dortmund
Staatsangehörigkeit	deutsch
Familienstand	verheiratet, zwei Kinder

Ausbildung und beruflicher Werdegang

1985–1989	Kerschensteiner Grundschule, Dortmund
1989–1998	Phoenix-Gymnasium, Dortmund
1998–1999	Zivildienst am Leibnitz-Institut für Arbeitsforschung an der TU Dortmund
1999–2006	Studium der Mathematik und Philosophie an der Ruhr-Universität Bochum
2003–2004	Auslandsaufenthalt an der Peking University, Beijing, China
September 2006	Abschluss des Studiums mit Diplom
2007–heute	Wissenschaftlicher Mitarbeiter am Institut für Neuroinformatik an der Ruhr-Universität Bochum