Studies of Goal Directed Movements

by

Emanuel V. Todorov

B.S., West Virginia Wesleyan College (1993)

Submitted to the Department of Brain and Cognitive Sciences
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

at the

MASSACHUSSETTS INSTITUTE OF TECHNOLOGY

September 1998

© Massachussetts Institute of Technology 1998

Signature of Author ....
Department of Brain and Cognitive Sciences
September 14, 1998

Certified by ....
Michael Jordan
Associate Professor of Psychology
Thesis Supervisor

Certified by ....
Whitman Richards
Professor of Cognitive Science
Thesis Supervisor

Accepted by ....
Gerald Schneider
Chair, Department Graduate Committee

Studies of Goal Directed Movements

by

Emanuel V. Todorov

Submitted to the Department of Brain and Cognitive Sciences
on September 14, 1998, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Abstract

The principles underlying the production of goal directed movements are largely unknown, as evidenced by the lack of artificial systems matching the real-world performance of biological organisms. The ability to transform in real time a vaguely defined goal (skiing downhill) into a detailed motor act (synchronized movements at multiple joints) appropriate under a wide range of previously unencountered circumstances (details of the terrain, snow conditions, etc.) is something we take for granted (after a period of painful learning). This is reminiscent of the subjective ease with which we perceive visually cluttered environments, or follow a conversation in a noisy room.

This thesis argues that the classic view of motor control as a unidirectional process, starting with a planning stage (a learned mapping from goals into actions or a straightforward interpolation connecting a number of intermediate postures into a detailed trajectory) followed by execution of the motor plan, provides an oversimplified and inadequate account of the versatility of everyday human performance. We propose that this standard model of motor control as well as the standard bottom-up models of perception, are inadequate because both perception and motor control are actually "inverse" problems, and are better treated as such.

After some theoretical considerations of goal directed multijoint movements, we develop a computational model of sensory-motor processing in intermediate point tasks (including reaching) based on stochastic optimal control theory. In this model planning and execution are not separate stages, but are both integrated into a tight sensory-motor loop which constantly adjusts the ongoing movement to better achieve the specified goal. The model provides a natural account for a number of experimental findings reported here (as well as previously observed phenomena); taken together these experimental results strongly indicate that the motor system updates online its internal estimates of both the environment and the moving limb, and is always ready to modify its "plan" in favor of a different movement that better achieves the goal under the new circumstances. The observations include eye-hand synchronization patterns, corrections for undetected saccade-triggered visual perturbations and hand inertia anisotropies, effects of desired accuracy on hand kinematics, movement segmentation, speed-curvature relationships.

Thesis Supervisor: Michael Jordan
Title: Associate Professor of Psychology

Thesis Supervisor: Whitman Richards
Title: Professor of Cognitive Science

Contents

Chapter 1

Introduction

Movements produced in everyday life pursue a goal, i.e. a desired change in the state of the environment. That goal can be as trivial and well defined as touching an object with the finger, or as complex and vaguely defined as hitting a tennis ball in a way the opponent will find difficult to return. The main theme of this thesis is the dependence of biological movements on the goal they pursue, and the need to analyze carefully the relationship between movement details and parameters of the motor task which the movement was meant to accomplish.

Why is it necessary to emphasize such a seemingly obvious fact? Indeed, one would think that the motor control literature would provide numerous examples of such analyses. However that is not the case. What the literature provides instead are numerous descriptions of geometric properties of movement trajectories, sequences of muscle activations, etc., with rather sparse explanations of how these movement characteristics arise in the process of trying to achieve a certain goal. The reason of course is that such explanations are very difficult to construct. Strictly speaking, to understand why or how a human subject produces a movement with certain characteristics, we would first need a good understanding of how the environment and goal are perceived, what the motivational or emotional state of the subject is, what cognitive processes (and in particular reasoning about the physics of the world) are involved, what the "preferred" patterns of movement are, etc. Clearly motor control could not make much progress if researchers waited for satisfactory answers to these questions. The way around such complications is usually the following: postulate a rather strict separation between motor control proper and the rest of the processing involved in movement production, i.e. assume the rest of the brain is somehow figuring out in sufficient detail what needs to be done (constructing a plan), and the motor system is responsible for implementing (executing) that plan. In order to draw a line defining the domain of motor control one has to be very specific about the flow of information across that line, i.e. describe the "interface" or "format" of the commands that the motor system accepts and implements.

Some of the most influential ideas in motor control can be seen as attempts to define this format, and demonstrate that it is indeed sufficient to explain large classes of biological movements:

- Sherrington's idea that movements result from chaining and combinations of low level reflexes
- Synergies - which are nowadays seen as combinations of muscles activated simultaneously, while the concept seems to have its origin in Bernstein's coordination patterns
- Servo control - the idea (first proposed by Merton) that due to muscle elasticity, movements can be produced indirectly by controlling the equilibrium state of the musculo-skeletal system

- A variety of "building blocks" - primitive impulses or "strokes" [Morasso and Mussa-Ivaldi 82, Viviani and Flash 95] concatenated to form more complex movements

- Generalized motor programs [Schmidt et al. 79] - basic movement patterns that can be scaled in duration and amplitude to satisfy the requirements of the task.

While these ideas have had a great impact on the study of the motor system and stimulated a lot of fruitful investigations, they have invariably failed to demonstrate that the analysis of movement can be justifiably reduced to some fundamental building blocks. The problem is the "cognitive penetrability" [Pylyshyn 80, Jeannerod 88] of movements: it seems that cognitive expectations, goals, etc. can have a significant influence on all aspects of the movement, making a strict hierarchical separation very unlikely. Reflexes can be greatly modulated in task dependent ways [Lacquaniti et al. 93]. Activated muscle groups underlying very similar movements in the same subject can vary substantially [van Galen and de Jong 95], and the muscle recruitment patterns as a function of movement direction can be quite complex [Flanders et al. 94]. The gain (stiffness) of the stretch reflex necessary for servo control seems to be too low [Bennett et al. 92, Gomi and Kawato 96][1] and furthermore the majority of neurons in movement-related cortical areas represent force [Evarts 81] or velocity [Georgopoulos et al. 82] or some combination of both [Kalaska et al. 89] which should not be the case if the CNS controlled only the equilibrium position. The invariant relative timing suggested by the generalized motor program does not seem to exist [Gentner 87], and kinematic invariances are not as strong as previously thought [Wright 90]. While concatenations of strokes or splines can fit most hand trajectories (as well as many other smooth functions that are not hand trajectories), we have no direct evidence that such primitives are actually used by the motor system; also, the structure of such strokes (e.g. the speed profile) is significantly affected by the accuracy required in the task [Jeannerod 88].

New proposals for "building blocks" or simple sets of rules for movement generation are welcomed by the motor control community, because they resonate with the generally accepted notion that the problem of motor control is too difficult and the only way to build a functional motor system is to find ways of simplifying it. What is often missing however is a careful analysis of where exactly the difficulty lies, and consequently it is often unclear whether the proposed "simplification" is an actual simplification. If we believe that motor behavior is governed by detailed preexisting programs and that research in motor control should be restricted to the execution of such programs, the above ideas indeed represent simplifications to the problem of motor control (i.e. the proposed "building blocks" are usually "simple" control signals that should be easy to execute).

However, in the absence of direct evidence that the motor system is indeed isolated from the planning process, it seems appropriate to include planning under the general heading of motor control. In that case discovering the control signals necessary to achieve a certain goal becomes an integral part of motor control, and the geometric simplicity of their final form does not in itself imply simplicity in the process of constructing them (or selecting them from a large collection of possible control signals). In other words, the complexity of the problem may become dominated by the complexity of planning, and therefore the simplifications we consider (especially when they have deep implications about the overall structure of the motor system) should perhaps be aimed at simplifying the planning rather than the execution problem.

Planning is often equated with solving the degrees-of-freedom problem, i.e. the process of selecting one of the infinitely many possible movements that would satisfy the goal [Wolpert 97]. This definition is appropriate in cases where the family of possible movements is rather obvious (e.g. reaching to a target). In more complex situations however it is not at all clear what the possible movements are, i.e. that

[1] More on this in Chapter 3.

family first has to be constructed before we have a chance to select one of its members. This movement construction process [Bernstein 67] is likely to be rather complex, as evidenced by the difficulties that robotics researchers have encountered in the field of motion planning [Latombe 91, Mason 98].

1.1 Organization of the thesis

The main point of this thesis is that the motor system is more "intelligent" than the usual descriptions imply. Instead of being a hardwired set of rules somehow tuned to achieve good performance, it instead maintains a more explicit representation of the goal and uses some powerful mechanism that can rapidly translate that goal into appropriate movements under a wide range of previously unencountered circumstances.

Chapter 2 discusses the difficulties involved in building such a system, and outlines a general approach to producing goal-directed movements as well as some simplifications that could make planning more tractable. It also argues that a number of common observations of biological movement are consistent with this view. We are certainly not going to solve the general motion planning problem here; our main goal in this chapter is to convince the reader that the human motor system has some of the defining features of an active system performing a lot of computation immediately before and during the movement, and also provide a general framework for thinking about motor control as being more than a mere replay of a prerecorded movement template.

Chapter 3 derives a mathematical model consistent with the above views for a restricted family of experimental tasks, involving hand movements through small sets of targets. The model is based on stochastic optimal control theory, which in this limited context is equivalent to the framework developed in Chapter 2. The model has the unique feature that it includes spatially limited visual perception and thus can address visuo-motor coordination.

Chapter 4 presents a large set of experimental data, and simulations of the model from Chapter 3. The experiments include estimates of visual accuracy, characteristics of hand trajectories under varying accuracy constraints, eye-hand coordination in via-point tasks, within subject variability, movement segmentation, corrections for unpredictable visual perturbations. The model can address other experimental observations available in the literature, which are also discussed: speed-accuracy trade-offs, gradual specification of motor responses.

Chapter 5 focuses on some of the lower-level mechanisms in the model (automatic compensation for muscle viscosity) and shows how this explains functional corrections for anisotropic hand inertia that have been previously reported. Further evidence for such automatic compensation is provided by a number of neurophysiological observations from primary motor cortex, that are naturally explained by the model. The chapter also presents results from learning and adaptation experiments generally consistent with the framework outlined in Chapter 2.

Chapter 6 contains a related model to the one developed in Chapter 3, which is descriptive rather than being mechanistic, but accounts quite well for the relationships between movement paths and speed profiles that have been observed previously. The model is applied to data from two experiments, demonstrating significantly better performance compared to the two-thirds power law which is the currently accepted description of this phenomenon.

Chapter 7 concludes the thesis with some final remarks.

Chapter 2

Theoretical Considerations of Multijoint Goal-Directed Movement

2.1 Current view

Underlying most investigations of the motor system is the assumption that planning and execution are separate: in the planning stage the external goal is transformed into a detailed trajectory (or some motor plan/recipe/program/schema), which is then used in the execution stage to produce an actual movement. Planning has received relatively little attention. In the typical tasks studied in psychophysics and neurophysiology (such as reaching or other transport tasks) the goal is defined by a simple visual stimulus and the family of successful movements is rather obvious. Thus the proposed planning algorithms [Nelson 83, Flash and Hogan 85, Uno et al. 89] are quite straightforward, involving some smooth interpolation method applied directly to the sensory input (in the appropriate coordinate system). The main problem recognized at this stage is the so called "degrees-of-freedom" problem, i.e. how do we select the most appropriate out of many possible movements. The standard solution (incorrectly attributed to Bernstein) is to assume that redundancy is eliminated by adding internal constraints, evidenced by movement regularities not implied by the task itself (e.g. isochrony principle, minimum-jerk principle, 2/3 power law, piecewise-planar 3D hand paths, linearly related joint velocities, etc.) More complex behaviors have been studied in the skill acquisition community, where planning is not considered separately but becomes a part of motor learning: it is postulated that a "motor schema" [Schmidt 75] appropriate for the task has already been formed. Its formation relies on the existence of some general-purpose capability to learn sensory-motor associations and generalize from them correctly (inspired by perceptual learning experiments with dot-pattern prototypes - [Posner and Keele 70]).

The issues that have attracted most attention in the field are related to the execution of a pre-specified movement: impedance characteristics of the limb, servo control models, role of feedback in correcting deviations from the desired trajectory, synergies and pattern generators, recruitment order and muscle activation patterns, etc. While this level of analysis is certainly necessary for understanding how the biological system works, one often gets the impression that if only we had better actuators and faster sensors, motor control would be trivial. In other words, the problems that are commonly studied are not challenging on the computational level [Marr 82]; they have efficient solutions on digital computers and do not pose serious problems for robotics engineers. Instead, the main difficulty in robotics (other than building hardware and vision systems) is motion planning. The planning algorithms that have been proposed are suitable for restricted transport tasks [Latombe 91]; they scale poorly (at least exponentially) and become impractical for realistic dimensionality of the configuration space and complexity of the environment (e.g. movable objects, nonholonomic constraints, uncertainty). Adding a new type of goal or constraint usually requires a new algorithm; it is not clear how to even approach motion planning in manipulation tasks such as shuffling and dealing a deck of cards [Mason 98]. Furthermore, human motor behavior is a lot more flexible compared to the robotic tasks for which algorithms have been developed. We can produce movements that attain very diverse types of goals (or combinations of them), including various features of the limb trajectory or desired effects on the environment that can be quite broadly defined.

Just like visual perception or speech recognition, the generation of purposeful movements appears (subjectively) immediate and almost trivial. The difficulties involved can really be appreciated only after someone seriously attempts to build an artificial system matching human performance and fails. Thus the proper computational level definition of the problem of motor control is "produce some movement that satisfies a given set of goals and preferences", rather than "execute accurately a prespecified movement" - which seems to motivate most current research. While these definitions are very different when applied to real-life situations, that difference has been blurred by emphasizing restricted laboratory tasks with trivial planning requirements, and focusing on geometric details of the recorded trajectory in isolation from the environment where the movement took place and the goal it achieved.

2.2 Problems with the current view

Let us assume for the sake of argument that the overall picture above is correct: for a given goal the system selects/instantiates a detailed motor plan (using a mapping constructed via some unknown and very powerful learning process), and then executes that plan as accurately as possible. While a concrete model of motor control based on these ideas has not been proposed recently and the present formulation is too vague to disprove, we would argue that it has the wrong flavor. It accounts well for stereotyped mechanical movements (indeed, that is how robots usually work - substitute "all-powerful learning process" with "human engineer"). However, it is in conflict with some of the most interesting and unique characteristics of biological movements:

- A number of authors have repeatedly commented on the fact that biological movements are surprisingly variable while the task is being held constant. That variability is not completely random - it is restricted to the family of movements that achieve the desired goal (e.g. the only common point of several reaching movements produced by the same subject is the target - [Bernstein 67]). The above model can produce different trajectories on repeated trials if either planning or execution is affected by internal noise, but there is no reason why the only aspects of the trajectory affected by noise should be the ones irrelevant for the task. Furthermore, any artificial system affected by enough noise to match the variability of human movements will be dysfunctional; it is then not clear how a biological system using much slower feedback loops can ever succeed on any task.

- In the planning-execution model the role of sensory feedback during execution is to correct deviations away from the plan. Such a mechanism would imply "blind" corrections: since the execution system does not know why that concrete plan was formed (i.e. it doesn't have access to the goal), it has to correct all deviations (this is how PD control works for example). In contrast, the on-line corrections exhibited by biological systems are goal-directed: when our foot comes in contact with an obstacle we do not kick it as hard as possible; instead we lift the foot and readjust the whole body so as to maintain balance. It is possible to fix this flaw of the model by assuming a large numbers of "backup" plans that can handle all common errors, but that will multiply the number of necessary motor programs (which already has to be rather large considering the number of different tasks that an average human can perform).

- If the motor system is indeed a collection of a large number of task-specific motor programs (or algorithms), one would expect to find evidence of modularization on the level of physiology. The little we know about the neural processing underlying movement production points in the opposite direction. The most robust finding in the arm-related areas of neocortex [Georgopoulos et al. 82] is that the majority of recorded cells[1] are broadly tuned even in the simplest possible multijoint task (reaching), implying that most cells somehow participate in the execution of most tasks. Some specialization seems to exist (e.g. "grasping cells" - [Rizzolatti et al. 90]) but is does not appear to be a sufficient neural basis for thousands of motor programs, corresponding to all the things we can do with our hands. While it is possible to have functional modularity without anatomical modularity, physiological observations are more compatible with an integrated system performing some general computation that is applicable to most motor tasks.

- The kinematic regularities mentioned earlier have been interpreted as manifestations of planning mechanisms that solve the degree-of-freedom problem by selecting one of the many possible movements. In that case such regularities should be very robust laws of movement (such as Listing's law) that are almost impossible to violate. Instead, they apply only "on average" and tend to break even in simple extensions of the original tasks (such as reaching in a vertical instead of a horizontal plane - [Atkeson and Hollerbach 85]).

- Successful execution of familiar tasks should result from an appropriate motor program, that encodes the "right" movement. Imagine the following experiment: ask a novice basketball player to shoot from a significant distance, always from the same initial position. Record several ball trajectories, average them (presumably recovering the motor plan of the subject) and then move the basket so that the average ball trajectory passes through it. Now we have adjusted the goal so that it matches the hypothetical motor program, and therefore we should expect perfect performance - which is quite unlikely. A related point is that the features of the "right" movement should be shared by most experts on a given task. Instead, it has been very difficult to identify "signatures" of expert performance, other than movement features directly related to the goal [Brisson and Alain 96, Todorov et al. 97]. In other words, between subject variability is like within subject variability - substantial, but restricted to the family of successful movements.

- Finally, a novice should have a much simpler/smaller motor program compared to an expert. Simpler programs are usually easier to run; thus it is very difficult to see why unfamiliar tasks require a significant amount of concentration and effort, while well practiced tasks become "automatic" and subjectively much easier. It is certainly not because experts do not have to use sensory feedback (e.g. expert tennis players rarely ignore the flight of the ball). In general, if movements result from preexisting motor programs, their accuracy should only depend on the quality of that motor program. Thus changing the accuracy requirements of the task should have no effects on the actual accuracy of the movement, or its kinematic details - both of which are false as we know from studies of speed-accuracy trade-offs, reviewed in [Jeannerod 88].

[1] Physiological recording typically involves substantial experimental bias [Fetz 92], i.e. experimenters tend to isolate cells that "make sense" in the context of the task.

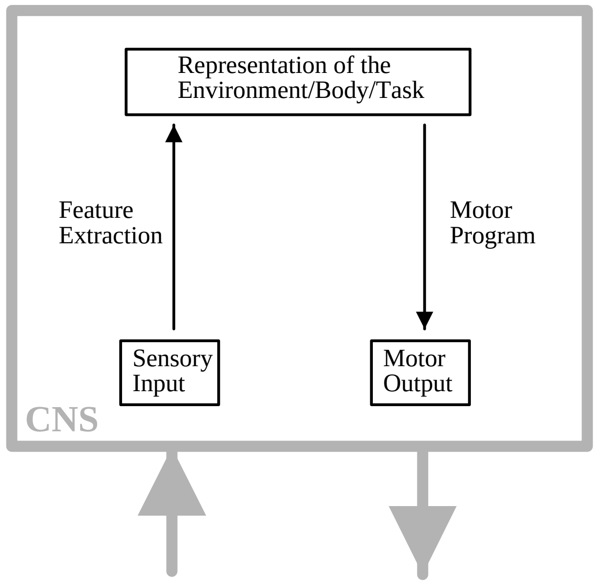

Figure 2-1: Schematic represenation of direct models of perception and motor control.

2.3 General approach

The computational level definition of motor control mentioned earlier is: given a goal or a set of goals and preferences G, generate a movement M satisfying G (i.e. "use all available means to achieve the desired effect"). Thus the input to the system is G, and the output is M. The current view of motor control postulates a solution involving a feedforward process that maps G into M in two steps (planning and execution). While this is the most natural (and at first glance the only) way to solve the problem, it is interesting that all difficulties discussed above can be traced to the direct, sequential nature of the proposed computation.

The approach advocated here is meant to address these problems and provide a more exciting picture of the motor system. While it has not been explicitly considered in biological motor control[2], very similar views have had a long history in the study of perception [Knill and Richards 96]. So we introduce the basic ideas through a comparison with visual perception, while paying specific attention to the fact that in many ways the two systems are opposites of each other.

The computational level definition of vision is: given a retinal image I, infer the state of the world S whose 2D projection on the retina is I. The traditional view of visual perception also involves a feedforward process that starts with I and sequentially extracts more and more complex features and builds higher and higher representations that eventually lead to S [Hubel and Wiesel 62, Marr 82]. Thus both vision and motor control are dominated by a "direct" theory, which claims that the input is mapped into the output via some sequential unidirectional process[3]. This is illustrated in Fig 2-1.

The trouble with direct theories is that it is not clear how to perform the required computation in

[2]But see section "Related work" for mathematical models that can be interpreted that way.

[3]To see this similarity we have to talk about inputs vs. outputs instead of top-down vs. bottom-up processing, since in the latter terminology the motor system is "upside-down".

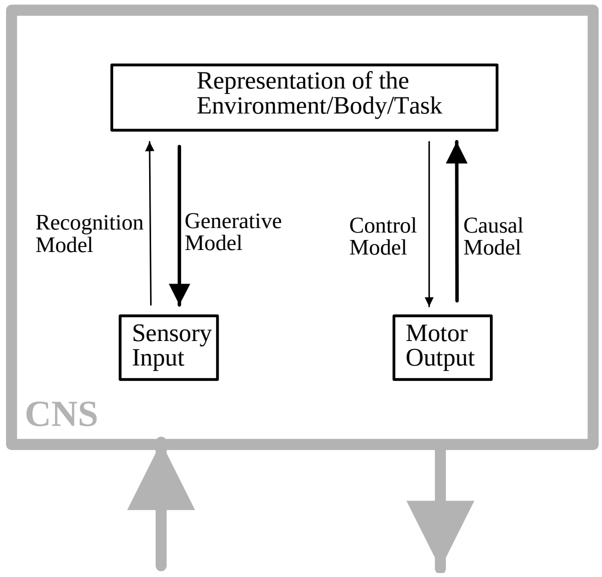

Figure 2-2: Schematic representation of inverse models of perception and motor control. Thin arrows correspond to the direction of flow that is desirable, but difficult to implement directly.

a principled way. In both cases the well-defined transformation is actually in the opposite direction. Given a state of the world S we can figure out what the retinal image I should look like (using a model of optics), i.e. there is a well-defined "generative model" and the only way to define the "recognition model" required for direct vision is as an inverse. Similarly, given a detailed description of a movement M we can figure out all potential goals that it achieves (by computing specific features of the trajectory or using a model of Newtonian mechanics to infer its effects on the environment), i.e. there is a well-defined "causal model" and the only way to define the "control model" necessary for direct motor control is as an inverse. This idea is illustrated in Fig 2-2.

There is no a priori reason why an inverse problem cannot be solved accurately by a direct algorithm. Indeed, many attempts have been made in AI (both in computer vision and robotics) to construct direct algorithms (typically involving a collection of heuristics) that work in carefully controlled environments [Russel and Norvig 95]. Such algorithms however tend to fail (or scale poorly) when the environment is made slightly more realistic. It is possible that the only general way to solve very hard inverse problems such as vision and motor control is to actually treat them as inverse problems, i.e. maintain a model of the well-defined transformation from outputs to inputs as well as priors over outputs, and use some (computationally intensive) algorithm for inverting that model. This corresponds to the "Bayesian inference" theory of vision [Knill and Richards 96]. In this view both visual perception and motor control are active processes, involving more on-line processing than a one pass feedforward computation.

That does not diminish the importance of "shortcuts" or pieces of direct transformations that can be incorporated in the inverse computation. None of the adherents of the Bayesian theory of vision would deny the existence of "edge detectors" in V1, and the fact that their activity is to a large extent generated by feedforward stimulation from the LGN [Hubel and Wiesel 62]. Still, even V1 edge detectors can be affected quite dramatically by stimuli outside their "classical" receptive fields [Somers et al. 98] - something that a well designed feature detector should not do. While the corresponding discovery in the motor system has not yet been made (orientation tuning in M1 is simply a correlation, not a mechanistic explanation), the motor equivalent of feature detectors is the concept of synergies, motor primitives, pattern generators: they can be seen as small pieces of a direct (control) model, assembled in various task-specific ways to produce complete movements.

This view of motor control may appear to correspond to the planning stage only, but it is not so. There is no reason to complete the entire computation before initiating the movement. It is better to finalize the results online, as sensory information becomes available. In fact, reflexes can be seen as shortcuts for inverting the causal model online. For example, the stretch reflex in many tasks produces useful corrections (i.e. it is usually appropriate to compensate for small deviations away from the expected movement); thus inverting the full model would often produce the same results as the stretch reflex, only slower. However the reflex is only a shortcut which can be modified when it becomes inappropriate (i.e. its gain can be greatly modulated in a task-specific manner - [Lacquaniti et al. 83]).

One may argue that computing a new movement for each trial is a waste of resources. Instead, why not store the movement just executed? This would be the optimal strategy if we consider the repetitive movements produced by robots and human subjects in laboratory experiments. However in the real world we rarely have to repeat exactly the same movement, i.e. there must exist a mechanism capable of computing the next movement while the present one is being produced. Once such a mechanism is in place it can handle repetitive tasks as well, so it may actually be wasteful to have a dedicated storage system for this rare case.

2.4 Mathematical modeling

The obvious candidate for a model of this kind is a belief network which encodes a causal model P(G|M) and priors P(M), and for a given goal G generates movements M with high posterior probability P(M|G). Here I discuss some of the features of such a model that are specific to the domain of motor control.

It was assumed for simplicity that the output of the system is the movement M. In fact, the output is a set of control signals C (the firing of motor neurons). The other quantities involved are the set of joint torques T, the state of the physical world W (which includes limb configuration, i.e. M), and the set of sensory inputs S. To create a Bayesian model, we need a model of muscle force production and moment arms that describes the transformation of C to T: P(T|C), a causal model of limb dynamics as well as mechanical properties of the environment that describes the effects of T on W: P(W|T), and a generative model describing how sensory inputs S arise from the physical world W: P(S|W). While studying P(S|W) falls in the domain of perception rather than motor control, we need to consider it here in order to incorporate reflexes (shortcuts) into the model, or study the effects of sensory limitations on movement.

The central part of a Bayesian model of motor control is the causal model P(W|C)[4]. It may not be feasible to build a very accurate model of Newtonian mechanics; even if we had one, it may not be useful since the physical properties of most objects in the environment are not known exactly. Thus is may be better to use some "qualitative" model. Such models have been studied in the field on qualitative physics [Weld and de Kleer 90] which can be further separated into qualitative dynamics and qualitative kinematics. The latter overlaps with the field of motion planning: most motion planning algorithms use some schematic representation of free space (e.g. silhouettes, roadmaps, cell decompositions) which is sufficient to plan a path in configuration space [Latombe 91]. While qualitative physics still seems to be

[4] Whether the causal model is decomposed as P(W|C) = P(W|T) P(T|C) or in some other way (or not at all) is not crucial for this discussion.

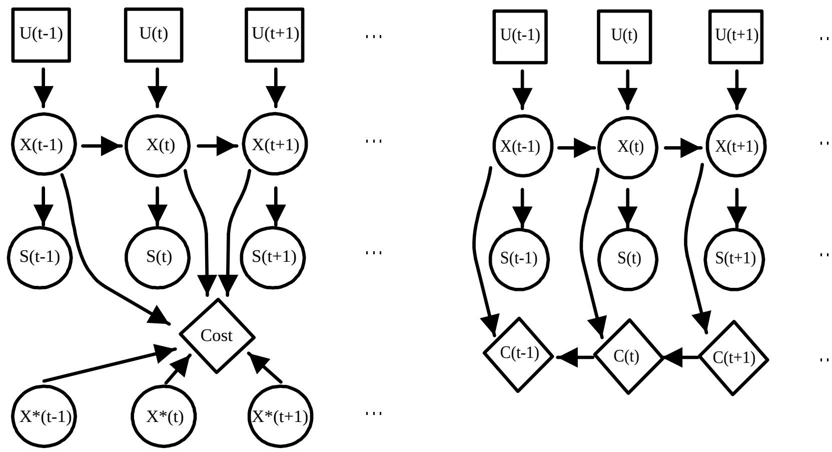

Figure 2-3: Influence diagrams. Squares correspond to control nodes (U), circles to random variables(X-state, S-sensory input), diamonds utility/cost function (C), arrows represent the causal flaw of information. Left: Influence diagram with a predetermined desired trajectory X* and single utility node. Right: No desired trajectory, sequence of "cost-to-go" utility nodes.

far from accomplishing its goals, it may be feasible to create models for restricted physical environments. This is sufficient for our purposes since the emphasis here is on the general principles of producing goal-directed movements in a Bayesian system (as opposed to a direct controller), which do not necessarily depend on all the details of physics.

A unique feature of the present model is that the state of the world W is not represented just because we want to know what it is (as assumed in perception), but because we want to act on it, i.e. achieve goals. If we were to use some compact representation of W (e.g. a set of generalized coordinates), it will be extremely difficult to represent the goal of the movement and develop control algorithms of sufficient generality. The problem is that most goals do not correspond to a unique state of the world, but instead to a rather large and vaguely defined family (e.g. "speed up towards the end of the movement", "write your name in a readable way", "move forward on two legs while maintaining balance"). Thus we need a massively redundant representation of W, in which most potential goals (e.g. desired effects, features of the movement, indices of performance) are explicitly represented. Once we have such a representation, we can simply instantiate the variables corresponding to the set of goals and use a general-purpose algorithm that transforms any set of explicitly represented goals into a detailed movement. The model should also include utility nodes (e.g. "do something as fast as possible"). Extending the representation in task-specific ways may be one of the changes that occur in the process of motor learning.

Another important feature of the model is that the represented quantities are functions of time (while one can learn a lot about vision by studying static images, studying posture alone in motor control is not that interesting). The present model is actually a mix of stochastic processes (the time-varying signals) and scalar variables (features of the entire movement). One possibility is to use dynamic belief nets, i.e. create copies of the network at discrete moments in time. Another possibility is to represent the time-varying quantities as combinations of component functions, whose coefficients become scalar variables (this is closely related to the popular concept of pattern generators or movement primitives).

Fig 2-3 illustrates the concept of influence diagrams, which are probabilistic models with control and value nodes added to them [Pearl 88]. In this framework time is represented by replicating the model at each time step. Arrows represent the causal flow of information: the next state of the system X(t+1) is affected by the current state X(t) and the next control signal U(t+1). The state X(t) affects the sensory reading S(t) generated at the current time step. The cost node is evaluated given the entire trajectory (sequence of states). While influence diagrams can be used as models of online execution only, i.e. after a desired trajectory has been computed (Fig 2-3 Left), it is also possible to incorporate both planning and execution (Fig 2-3 Right) in the same model, by defining the utility nodes appropriately. Using a sequence of utility nodes is closely related to the idea of computing a cost-to-go function in optimal control (see below) which is crucial for completing the necessary computations in a reasonable amount of time. The causal (possibly qualitative) model of mechanics and the rich representation of possible goals are hidden in the arrows among U's and X's and the C nodes respectively.

A fundamental difference between a Bayesian theory of vision and motor control is the treatment of priors. In vision, there is an actual state of the external world that we would like to infer, and we absolutely have to use informative priors in order to rule out the infinite number of alternative explanations consistent with the retinal image. In motor control it is not so obvious why informative priors have to be used: there is a specified goal and infinitely many realizable movements that achieve it, so in principle we can choose randomly any such movement without violating the computational-level requirements for the motor system. In other words, the "degrees-of-freedom" problem does not exist on the computational level of investigation. One possible function of priors is to generate efficient/comfortable/smooth movements[5]. However, the kinematic regularities found in arm movements (that have been explained as solutions to optimization problems) do not minimize quantities that are obviously advantageous to minimize - i.e. the derivative of acceleration or torque.

Alternatively, it is possible that the function of priors in motor control is not to introduce some arbitrary constraint, but to encode the family of movements that are most likely to be useful for a class of tasks (before the concrete goal is known). The function of such priors is not to lead to the "correct" output as in vision, but to speed up the process of finding a movement achieving the goal. This definition is more consistent with the usual meaning of "priors"; it can be learned by collecting statistics about the useful movements. It has been suggested for example [Hogan and Flash 87] that minimizing jerk could have a similar function, by making movements more "predictable". This is still not equivalent to the use of priors in vision, i.e. it is not necessary at the computational level as defined previously. However, if we modify the computational definition of motor control to require the production of a movement suitable for the task, in a *limited amount of time*, then it may actually become impossible to satisfy the computational level requirements without using such priors[6].

While a causal model can be inverted via some stochastic simulation process, it is clear that such an algorithm will not converge in time (especially for online specification of the remaining part of the movement) and at least a partial control model must exist. This suggests using some variant of a Helmholtz machine [Hinton et al. 95] which maintains both a causal and a control model. The causal model however should be used not only to construct a control model during learning but also in combination with the control model during movement production - otherwise we are back to a direct theory.

Finally, why are we assuming that the motor system uses any sort of probabilistic model? There are two important reasons to think that such models may be used in motor control (besides the obvious mathematical relevance to the types of computation we have in mind). First of all, both sensors and

[5] That can also be achieved by adding additional goals. The corresponding trick in vision would be to make changes to the retinal image, i.e. "hallucinate" extra information that will lead to the correct percept.

[6] This definition is not a mix of the computational and algorithmic levels, as it may appear. We are not making a statement about the speed of a specific algorithm, but limiting the running time of whatever algorithm is used - related to the computational complexity of the problem.

actuators in biological systems are affected by noise - and the optimal way to handle a system affected by noise is to have a model of that noise. More importantly, the internal models of the physics of the world as well as the effects of muscle commands are likely to be incomplete and/or slightly inaccurate. In fact, that has to be the case in manipulating objects with physical properties that we have no information about. Thus from the point of view of the motor system the behavior of the world (as inferred through sensation) is slightly different from the one that can be predicted using an efference copy of the motor output. Such inaccuracies are best represented as additional noise terms. Another possible concern is that correct probabilistic inference is not necessarily a good model of high-level human judgements and reasoning [Kahneman and Tversky 73, Osherson 90]. However, there is growing evidence that a number of phenomena in low- and mid-level vision and proprioception can be explained by the principles of Bayesian inference [Wolpert et al. 95b, Knill and Richards 96, Nakayama and Shimojo 92].

2.5 Supporting evidence

It is easy to see how a Bayesian theory of motor control can in principle resolve the problems with the current view discussed above:

- In such a system between-trial variability will result from a) the noise terms discussed above, and also by b) the necessity to terminate the computation before it has had sufficient time to converge to the optimal solution. On each trial a probabilistic system (using some stochastic simulation method) will produce a movement which is likely to achieve the goal in the given task; however every new movement will be an independent sample from a posterior probability distribution (conditional upon the specified goal). Thus we should expect to see a lot of trial-to-trial variability mostly constrained to the family of successful movements. Studying the space covered by observed movements may reveal something about the structure of the probabilistic model.

- Since sensory feedback is represented in the same model and there is no strict separation between planning and execution, the effect of online feedback will be to make successful movements more likely given what has already happened during the trial. Perturbations will have the effect of changing the distribution from which (parts of) movements are being sampled, so the system will react to perturbations in a goal-directed manner instead of blindly trying to return the limb to some predetermined trajectory. Note that in this model natural variability and online corrections have very similar bases.

- An important advantage of a probabilistic model is that achieving arbitrary combinations of sub-goals does not require qualitatively new mechanisms - one simply has to instantiate all desired effects of the movement, and the system can use the same general algorithm for sampling the posterior distribution. Furthermore, there is no need to create a completely new causal model for each task (although specific types of goals may be added). Thus a Bayesian motor system will be a lot more integrated than a large collection of motor programs.

- The kinematic regularities described above represent a restriction of the family of observed movements beyond that implied by conditioning upon the goal. In the present model they are emergent properties resulting from priors or task-specific shortcuts for inverting the causal model that have been incorporated in the general algorithm. In either case, we would expect regularities to hold approximately and be sensitive to small alterations of the goal.

- In the motor program view movements result from some motor plan that can be recovered by averaging a number of repeated trials, thus adjusting the goal should make the existing program successful. In contrast, averaging a few movements generated by a probabilistic system and adjusting the goal accordingly is not sufficient to achieve expert performance: we are guaranteed that the family of appropriate movements intersects the family assigned high posterior probability, but there is no reason to expect a large overlap between the two which is necessary for expert performance.

- In the model proposed here learning of a new motor program is replaced by learning a more accurate causal model (e.g. adding internal representations of task-specific goals that otherwise have to be approximated), or learning how to invert the model more efficiently (by creating shortcuts or adapting the priors). Thus producing movements in a novel task actually requires more computation, since it involves running the general algorithm without any task-specific speedup.

In the previous section we argued that vision and motor control are similar in the sense that both are inverse problems solved through a Bayesian inference process. A probabilistic system solving an inverse problem should have two general features distinguishing it from a direct computation: a large space of solutions will be sampled, and the solutions will be significantly affected by priors. One might expect then both the motor and the visual systems to exhibit both properties. Instead, the above evidence from motor control emphasizes variability and details of the output, while the typical arguments in vision involve demonstrations that prior knowledge about the world can affect the outcome of visual processing.

Is it possible to translate the evidence in favor of Bayesian theories of visual perception to the field of motor control (and vice versa)? This may be possible in principle, but very hard experimentally. The difference is that in vision the input (retinal image) is known and controlled by the experimenter, while the output (percept) is vaguely defined and hard to measure quantitatively. In contrast, in motor experiments we can record the output of the system in great detail, while the input (the goal) is vaguely defined - there is a specified motor task, but it is not clear how that is represented internally, and whether other subgoals/preferences are being added.

Thus vision experiments relying on carefully constructed inputs are hard to translate into motor control. For example there exist a number of visual illusion, i.e. stimuli that would be perceived correctly by a feedforward process but are in fact perceived incorrectly due to prior assumptions about the world (e.g. brightness illusions, size illusions, etc.) No motor illusions have been constructed so far, i.e. motor tasks where some of the (sub)goals are missed consistently due to priors over the space of movements (one possible exception is the phase locking in rapid repetitive movements - [Jekka and Kelso 89]). Another example from vision is the preference for specific interpretations of images that have infinitely many possible interpretations (e.g. the generic viewpoint assumption - [Knill and Richards 96]). This could correspond to kinematic regularities, if we find a way to express them as resulting from meaningful priors rather than just being arbitrary formulas.

2.6 Related work

2.6.1 Forward models and computational learning

The use of "forward" models has been proposed previously in motor control [Jordan and Rumelhart 92], and since then substantial evidence has accumulated in support of that notion [Wolpert 97]. This term corresponds to what we called a "causal" model above (and is completely unrelated to what we called a "direct" model); it is a model of muscle contraction as well as physics of the limb and the environment, than can be used to internally simulate or predict the outcome of currently considered motor outputs.

So far this idea has been used to model sensory estimation, and also to avoid the negative consequences of redundancy in learning a mapping from goals into actions (i.e. learning a "direct" controller). Such models have only been considered in the framework of planning a single desired trajectory for each possible goal, i.e. they have not previously been used in conjunction with an optimal controller to actually compute an appropriate trajectory given the goal of the movement. Another difference from previous work is the emphasis on rich representations of potential goals and effects on the environment, as opposed to equating the goal with a concrete sensory state.

To avoid a potential confusion of terminology, we mention "direct inverse modelling": the idea of applying control signals, measuring the outcome, and using the pair outcome-control as training data to some learning model that will after training be capable of mapping goals into actions. This type of learning has problems with redundancy: if several control signals achieve the same result they will be averaged, resulting in a control signal that may be unacceptable (if the equivalence class is nonconvex) - see [Jordan and Rumelhart 92]. Again, such a learning process will produce what we called a direct controller.

2.6.2 Stochastic optimal control

This is the mathematical framework most closely related to the above discussion - see Chapter 3 for details. It combines optimal estimation (i.e. inference about the most likely state of the environment, based on all available sensory data, previous control signals, and knowledge about the plant/sensory-motor delays/noise, etc.) with optimal control (i.e. selecting control signals that maximize expected performance, instead of tracking a specified trajectory). The optimal estimates are derived using a generative model that describes how sensory data depends on the state of the limb/environment. Optimal control signals are computed using a causal model describing how control signals affect the plant (plus a well-defined performance index). These two submodels correspond exactly to the thick arrows in Fig 2-2. The present optimal control framework has some very important limitations that restrict its usefulness in controlling a biological organism executing a nontrivial task. While it can handle the effects of noise and uncertainty, it relies heavily on an accurate model of system dynamics and a well defined cost function that is expressed in some (nonredundant) system of generalized coordinates. However, human beings can be quite skillful in tasks where accurate description of the objects being manipulated is unlikely, and the goal is rather vaguely defined and not obviously expressible in joint angles/hand location/etc. Another problem is that the existing algorithms require linear dynamics, quadratic cost functions, and Gaussian noise models (systems that violate those requirements have to be approximated in order to avoid an exponential increase in computation time). Still, it is possible to adapt the optimal control framework for studying limited classes of motor tasks (as in Chapter 3) where goals are defined in terms of desired hand locations, and no objects with unknown physical properties are being manipulated.

Optimal control models have been studied previously in the context of speed-accuracy trade-offs [Hoff 91, Meyer et al. 82]. These studies have ignored uncertainties in sensory estimation, and are limited to one-dimensional reaching movements. Still, they provide a very satisfactory fit to a large class of experimental data on the effects of desired accuracy on the timing and kinematics of reaching movements.

So far optimal estimation and optimal control have not been combined in one model in the study of biological motor control.

2.6.3 Reinforcement learning

Reinforcement learning is often proposed as a general solution to motor control problems. The model uses a discrete set of possible states of the environment/limb, and a discrete set of actions (or control signals) that can be applied in each state. The learning procedure (Q-learning) estimates the probabilities of making a transition from one state to another given that a certain action was applied (i.e. it learns a model of the system dynamics). It also learns an "optimal policy" - a prescription of what action to take in each state in order to maximize some performance index. The advantage of reinforcement learning is that it requires very little prior information, very little feedback, and is quite general (and also provides a fruitful ground for proving interesting mathematical results, but that fact in itself doesn't make it a more attractive candidate for modelling biological systems). The problem is that this model grows with the number of discrete states being represented. For control problems of realistic complexity, we have to discretize a (continuous) state space with rather large dimensionality, typically causing an exponential explosion (for example, to control a human hand manipulating an object, we probably need about 15 degrees of freedom; assuming a discretization of only 10 points in each dimension, the number of states is approximately equal to the number of synapses available in the CNS). Thus the model has only been applied to carefully constructed low-dimensional control problems. Unless this exponential explosion can be somehow eliminated, it is not clear how reinforcement learning can provide much insight into the organization of complex multi-joint movements. The same is of course true of dynamic programming algorithms that involve an exhaustive search of an exponentially large space. Although the optimal control problem (Chapter 3) is solved using dynamic programming, there are a number of practical cases where the solution is actually tractable - which explains its numerous applications in engineering, and our interest in it.

2.6.4 Situated autonomous agents

A strong reaction to the quest for provably exact planning algorithms that take forever to compute a single control signal is the work on situated autonomous agents [Brooks 86]. This work emphasizes the fact that motor control has to occur in real time, even if the resulting control signals are suboptimal. The solution originally proposed by Brooks is a layered architecture that is handcrafted to control a mobile robot. A more flexible algorithm that does not require that much manual programming is the action-selection network of Maes [Maes 91], in which possible actions are represented as nodes connected to each other according to the expected outcomes of each action and the preconditions it requires to be executed. The network is operated by passing "activation" corresponding to the current state and desired set of states, iterating it for a few cycles and executing the action with highest energy. This is an interesting example of a system that maintains a representation of the goal, and uses it in a rather direct way to select appropriate actions[7]. A modified version of this algorithm (combined with inspiration from ethology) has been used successfully to control animated characters interacting with human users in real time [Blumberg 97]. Unfortunately this type of action selection has limited applicability to the study of motor control because it relies on atomic/symbolic actions (but see Giszter 94) and it is not clear how to apply it to selection of control signals from a continuous control space. The same problem hinders the use of ethological observations in motor control - ethologists are typically interested in higher level "actions" that can be verbally described (e.g. a duck rolling an egg into its nest) and do not use the recording equipment necessary to document the fine movement details studied in motor control.

[7]Its problems are realted to the lack of information about the source of activation. Some attempts have been made to fix that problem by tagging activation signals with the index of the node where they originated, making the algorithm less elegant. Interestingly, this is very similar to the problems of double-counting of evidence in early expert systems, that are currently resolved by using beilef networks [Pearl 88]. Thus it may be interesting to reinterpret the action-selection algorithm in the belief network framework.

2.6.5 Simulated physical environments

A way to illustrate the theory outlined above would be to simulate a simplified model of a human arm, operating in a constrained environment obeying some simplified model of physics. We need to define a rich enough family of tasks that can be parameterized in a reasonably compact way, and see what types of behaviors the model will generate. The idea is to find parallels to the basic phenomena of human movement discussed above, rather than trying to replicate quantitatively the results of concrete experiments.

An inspiring example is the work of [Sims 94]: he simulated the interactions of simple virtual creatures with a simplified physical environment, defined a goal (e.g. swimming, jumping, walking) by a cost function (e.g. how much the center of mass moved forward or up) instead of a prespecified trajectory, and then used a CM5 to run a genetic algorithm producing creatures with successful behaviors. Although the creatures were made of blocks, the resulting behaviors look very "biological" in the sense that different creatures accomplish the goal in completely different ways. The limitation of Sims' approach is that each creature has just one hardwired behavior (which is a part of the genotype). However, if we think of the whole simulator as a model of the motor system, and constrain it to always use the same morphology of the virtual creature, then we have a system that can generate a whole family of movements that satisfy the same goal, but are very different in their details. In our terminology, this simulator creates a new control model for each trial (by searching the space of control models rather than inverting a causal model).

Another interesting example from the computer graphics literature is the work of [Terzopoulos et al. 94], who modeled quite realistically a fish body with about 60 degrees of freedom, controlled by 11 muscle-like actuators, and swimming in a reasonable simulation of water. This work represents a rare example of learning the "muscle synergies" optimal for a given organism and family of tasks - in this case swimming forward and turning. The authors used a simulated annealing algorithm to find the optimal time-varying muscle activation coefficients that resulted in the best swimming performance. Since their goal was believable animation rather than studying motor control, they were satisfied with the minimum possible number of synergies that could produce the desired behaviors. One can imagine however running a similar simulation with, say a model of a human hand trying to manipulate a number of objects and studying the types of movement patterns that emerge.

It should be emphasized that the success of these simulations seems to depend crucially on using a well defined and internally consistent (although not entirely realistic) model of physics, and specifying goals by numerical indices of performance rather than prescribing desired trajectories.

2.7 Motor psychophysics

The movements currently studied in motor psychophysics involve mostly reaching; while this is a very good starting point, focusing exclusively on these tasks is not likely to lead any further than the study of oriented bars has led in vision. Some very important aspects of realistic arm movements are eliminated in such tasks.

There is no manipulation of the environment (only transport), which is very restrictive since manipulation is the ultimate goal of most arm movements and requires a lot more complex planning than transport tasks. By focusing on the end-point (typically fingertip, or handle) we ignore the kinematic redundancy of the human arm and consequently a lot of the variability present in natural movements (this is particularly undesirable in the present framework). The visual stimuli are too simple and the demand of the task too low, leaving subjects with the feeling that they can execute the task in any way they want (and raising the possibility that arbitrary internal preferences are being added to the goal).

This makes the average motor control experiment a lot less convincing than a vision experiment where subjects are not free to choose what they see. More demanding motor tasks (such as playing computer games, or returning a difficult tennis ball) do not have that problem. Finally, studying simple discrete movements eliminates the need to continuously take into account the state of the environment and the demands of the task and to compute the next movement while producing the current one. This is another important feature of real-world performance, which a reasonable model of motor control should be able to deal with.

The observations of human movement that we used to argue against the motor program view mostly come from complex tasks [Bernstein 67] that are difficult to study rigorously. Such observations are less common in the simple tasks used in motor psychophysics, due to the limitations listed above. Thus it is desirable to design a family of controlled laboratory tasks (or a "metatask") that avoid these limitations. That would involve a virtual environment with some objects whose position/shape/properties/behavior can change, tracking more degrees of freedom than just the fingertip, and continuously manipulating objects (probably with some virtual manipulator simpler than a human hand). The metatask should have a relatively simple parameterization, and it should be possible to explain it to subjects so that showing a new stimulus configuration does not require new verbal instructions.

The most important feature of human arm movements in the present model is their variability, which has traditionally been ignored (by focusing on the average movement). In this framework it becomes very important to quantify not just the average movement but the whole family of movements being sampled, and analyze its overlap with the family of all successful movements. Furthermore, there is no reason to keep the task fixed but instead it can be varied on every trial, thus exploring a much larger subspace of possible movements and revealing more of the structure of the underlying system. That approach will also allow a better quantification of whatever kinematic regularities emerge in the task. An exciting possibility (suggested by recent results on one-trial learning of dynamic force fields - [Thoroughman and Shadmehr 97]) is that some of the variability normally observed is actually predictable given the sequence of recent movements and the corresponding sensory feedback (and a concrete model of how one-trial learning takes place).

Another important aspect of the present model is online computation, and balancing the available resources between current and future movements. Chapter 4 describes experiments involving undetected visual perturbations (synchronized with saccades) that are introduced transiently and should not affect the arm movement if a concrete plan was already formed and the execution stage only involved corrections away from that plan. Instead, such perturbations result in significant trajectory modifications. Also, simply studying the timing of eye movements while the arm moves through a sequence of targets yields evidence for appropriate allocation of processing time between current and future goals.

While the model outlined here is not very specific about learning new skills[8], it offers a new perspective. The problem with motor learning is the lack of a good definition of what exactly changes in the course of learning. Obviously, performance increases and that is somehow due to the feedback received on previous trials, but that says very little about the actual modifications in the motor system responsible for this improved performance[9]. In the present model, the state of the motor system is described not by the average movement, but by the whole family of movements (including those elicited by perturbations). It is possible that the important differences between novices and experts are to be found by

[8] More details are presented in Chapters 3 and 5.

[9] If we believe that observed movements are manifestations of detailed preexisting plans, then studyng changes in the movement trajectory is equivalent to studying changes in the state of the motor system. In that case the study of motor learning is reduced to simple observation. However, if the observed movements result from the functioning of a more complex system performing a significant amount of online computation, internal changes cannot be directly inferred from trajectory changes.

comparing the distributions of movements. In particular, it is possible that experts are better at using online information in a functional way, as opposed to applying some low-level correction. A clean way to test this hypothesis is to train subjects to different performance levels, and introduce perturbations some of which interfere with the goal and others do not. The prediction would be that novices will correct for all perturbations, while experts will ignore the ones that do not interfere with the goal. It is also possible that experts develop priors that are more appropriate for the task (that could be studied with a modified timed response paradigm - [Hening et al. 88]). Another possibility is that experts have a better internal representation of the goal. In general, it would be desirable to have an experimental method for determining what exactly the goal is for a particular subject; that vaguely corresponds to the methods for measuring psychometric functions in perceptual psychophysics.

Finally, it may be possible to find some equivalent of visual illusions. A complex but controlled motor task essentially gives the experimenter more control over the input to the motor system (i.e. the goal), which is one of the problems with constructing motor illusions (the other problem is that it is not clear what exactly the priors are). A less ambitious project may be to find situations where small (spatially localized) variations in the task lead to big differences in the family of resulting movements; this is reminiscent of arguments in favor of top-down effects in vision.

2.8 End-effector control

End-effector control is the idea of achieving control over a part of the limb, or possibly an external object, that is most directly related to the task criterion (thus the term "end-effector"). In support of such a control scheme, it has been demonstrated that the movements of the hand look "simpler" in Cartesian space (where manual tasks are usually defined) rather than joint space [Morasso 81], and that a number of kinematic invariances seem to be preserved in the relevant end-effector space. A closely related idea in robotics is that of operational space control [Khatib 87] or virtual model control [Pratt 95], where the central system assumes that it is controlling a much simpler manipulator (virtual model) defined in the appropriate task space and issues the corresponding control signals, which are then translated by some efficient low-level system into signals appropriate to drive the actual limb (this may involve cancellation of interaction torques, compensation for known damping, etc.)

This idea can be incorporated very naturally in the framework discussed above - see Fig 2-4.

Suppose we defined a set of virtual motor outputs, and a rule for translating them into actual motor outputs. Then the mapping from virtual to real motor outputs becomes a part of the causal model, i.e. we can now infer the expected achievement of the goal given the virtual motor output. If it were possible to design this mapping in such a way that virtual control signals have straightforward effects in end-effector space (i.e. we can easily compute how well the goal is satisfied given the virtual control signals) then it may be possible to ignore the complete causal model and replace it with a (task-specific) causal model that is much simpler and easier to invert. There is a tradeoff involved in designing such a virtual model. If we require the mapping to real motor outputs to be simple (e.g. a linear transformation) that would simplify learning/constructing the model, but may also limit its usefulness since a simple mapping to motor outputs may not be sufficient to achieve the desired effects in end-effector space. On the other hand, attempting to achieve perfect end-effector control may require a mapping from virtual to real motor outputs that is too complex to learn.

This tradeoff corresponds very closely to the different definitions of a "motor synergy" that have appeared in the literature. The original idea [Bernstein 67] was that synergies are general coordination patterns whose defining property is that they make execution of the task easy. The existence of such coordination patterns and their necessity for any skilled movement was obvious to that school of thought,

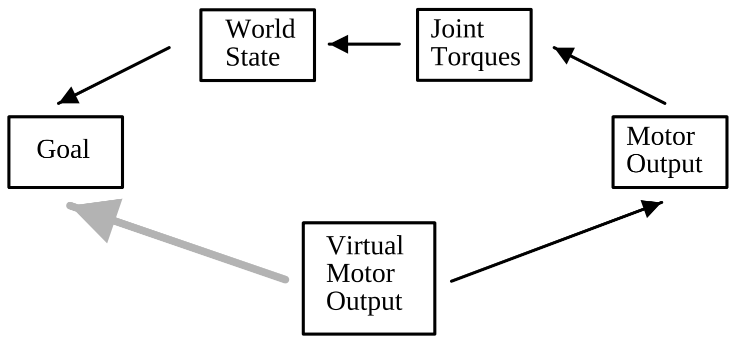

Figure 2-4: Schematic representation of using virtual control in a causal model. Solid arrows correspond to the original causal model. The virtual output is a vector that is translated into motor outputs, thus extending the causal model. The gray arrow corresponds to the end-effector causal model.

to the extent that observations of expert performance were interpreted as evidence for such coordination patterns [Aratyunyan et al. 69]. The problem with this definition of course is that it has never been possible to identify what these patterns are, and how exactly they can be formed. On the other extreme is the view of synergies as muscle synergies, i.e. groups of muscles that are activated simultaneously in fixed proportion. While it is clear how such synergies can be implemented, it is not that clear what they could be used for. Consider for example the simple problem of moving the hand in 3D space. An appropriate set of synergies in this case would be a set that at each point in space provides 3 directions of movement, along some reasonably defined axes. Thus it would be convenient to have 3 muscle synergies underlying all hand transport movements. If we were to activate a fixed set of muscles, however, the resulting directions of movement would not be very simple: the problem is caused by the variable length-tension curves of muscles, variable moment arms and geometry, non-commutative rotations in 3D, gravity, configuration-dependent inertia, etc. [Buneo et al. 97]. Unlike EMG data from frog preparations [D'Avella and Bizzi 98], the available EMG data from human subjects is not generally consistent with the existence of a small number of muscle combinations always used to transport the arm [Flanders et al. 94][10]. So if fixed combinations of muscles are being used as synergies, there must be a lot of them, and the system will have to perform a lot of additional computations to figure out when it is appropriate to use which synergy, i.e. the simplification of the control problem resulting from muscle synergies is not that obvious.

The truth probably lies somewhere in between the two extremes. Perhaps computational models can shed some light on this issue, by searching the space of virtual models that simplify the causal model and yet are easy to implement - such as [Terzopoulos et al. 94]. Note that implementing virtual model control in robotics is not so challenging, since a) the tasks robots perform are significantly simpler and involve less degrees of freedom, and b) the available sensors and actuators are faster and more reliable, and it is possible to estimate the physical properties of the plant quite accurately.

[10] Although it may be argued that a sufficient number of muscles has not been recorded from, or a sufficient range of motor tasks has not been explored.

Chapter 3

Stochastic Optimal Control of Eye-Hand Movements in Via-Point Tasks

3.1 Overview

In this chapter we restrict our attention to a small family of motor tasks which will allow us to instantiate the general theoretical framework developed above into a concrete model of motor control, and make quantitative comparisons to experimental data. We consider tasks where a subject is asked to move a point (that could coincide with the fingertip, or the tip of a rigid object held by the hand, or a simulated cursor controlled via isometric force application) to a specified target location, possibly passing through some intermediate targets (called "via-points"). The task parameters that can be varied include the number and location of targets, definition of the end-effector, timing and accuracy constraints, visual feedback conditions, and different visual or mechanical perturbations. Both hand and eye movements will be analyzed simultaneously, in a $2D$ plane[1].

Although restricted in many important ways to be discussed later, this family of tasks allows us to address a significant proportion of the experimental observations of arm movements: speed-accuracy trade-offs, kinematic features of planar movements, effects of both unpredictable and systematic perturbations, timing patterns of eye movements, reaction times, arm impedance characteristics, etc. Since these observations have been made by different researchers in different experiments, they have traditionally been considered in isolation, leading to different and not necessarily compatible theoretical accounts and emphasizing different aspects of sensory-motor processing. However all such phenomena would occur simultaneously in the same subject and the same experiment, if one were to manipulate all relevant task parameters and measure all relevant responses. While it is possible that the motor system contains a large collection of hardwired "rules", one for each empirical finding described in the literature, it is more likely that the execution of a simple task such as the one studied here is based on a few basic principles,

[1] While the actual arm movement has many more than 2 degrees of freedom, our assumption is that the motor system uses end-effector control of the type discussed in Chapter 2. Thus control signals will be assumed to affect the hand position directly. Inaccuracies due to remaining posture dependence will be modeled as noise, rather than being taken into account explicitly (as functions of joint configuration). We will use "hand" and "end-effector" interchangebly, although the end-effector may not always coincide with the hand.

giving rise to a number of emergent phenomena. Thus it is desirable to have a single model that could address, at least in principle, the majority of observations made within the context of via-point tasks.

The next section presents the rationale of our modelling approach, and its relation to the discussion from Chapter 2. We then summarize some classic results of stochastic optimal control theory, which serve as the starting point for a number of modifications necessary for a reasonably realistic model of eye-hand control. These modifications require substantial changes in the classic optimal control law, presented in the following section in a self-contained derivation of the algorithm for simulating our model. The methods used in the derivation are developed in detail in [Davis and Vinter 85, Bertsekas 87, Wilsky 94]. We conclude the chapter with a summary of the algorithm and some implementation details. The actual predictions of the model and their experimental verification is deferred to the next chapter. Note that the model proposed here is meant to be quite general, and it could be used to address a large number of phenomena not all of which will be discussed in this thesis.

3.2 Modelling Approach

In the spirit of the theoretical discussion from the previous chapter, we will assume that the sensory-motor system is simply trying to do what it was asked to do, as well as it can. In other words, it attempts to produce the "best" possible movement, within the specified time limits, using the available sensory input as well as prior information about its own properties and limitations. The goal here is to develop a *mechanistic* model of motor control for via-point tasks, that can be subjected to almost any experimental paradigm a human subject could participate in. Such a model would take as input the information that is available to the human subject, moment by moment, and produce a "real-time" response that can be compared to experimental data[2]. In contrast, *descriptive* models provide a parametric form of subjects' responses, and the fitting procedure involves estimating the values of the free parameters from observed data. While it is much easier to fit descriptive models to data and a number of models in the field fall in that category (see Chapter 6 for a discussion), classifying such models relative to the fine line between explanation and data fitting often involves a subjective judgement of how "surprising" it is that the data is fit with so "few" parameters.

It is clear from the above requirement for real-time optimal response that the model will have to use sensory information to guide the ongoing movement. At this point one may be surprised that we are dismissing so lightly the long-lasting debates about the role of feedback in movement production [Adams 71, Abbs and Winstein 90]. Such arguments have focused on rapid movements (since no one would dispute the importance of feedback for slower movements) and whether they are entirely ballistic, and for how long. Obviously any sensory input in a biological system has a non-negligible delay, and any serious model has to take that into account. However that fact in itself does not imply two different modes of operation - ballistic and sensory guided; what is proposed here is a single system that is capable of producing ballistic movements in the absence of sensory feedback and is also prepared to utilize feedback as soon as it becomes available.

Given our model requirements, a natural mathematical framework is provided by the theory of stochastic optimal control, which is concerned with maximizing some performance criterion by controlling a stochastic system with specified dynamics and noisy state observations. Before going into details, we

[2] By "real-time" we mean a response based only on information available up to this point in time. The actual time necessary to run the computer simulations is dependent on available hardware and our skill in developing efficient algorithms. Since both of these increase constantly, it is quite meaningless to make arguments about the "difficulty" or "biological plausibility" of certain computations based on our current ability to implement them in real time on digital computers. That comment obviously does not apply to provably intractable computations, but this is not the case here.

make a final point about the general properties of the model in relation to the discussion in the previous chapter. The solution to an optimal control problem takes the form of an optimal control law (or policy), which is a formula specifying what control signals to output at each moment in time as a function of all available information so far. In other words, the control law is a prescription of what to do in any conceivable situation that may arise. In some special cases (as the one in the next section) such a law can be computed before any feedback is received, which in some sense corresponds to a strict separation between planning and execution (where the "plan" contains a lot more than a single desired movement). However, as soon as we make the mathematical model slightly more complex and realistic, we have to resort to approximations to the control law that are performed in real time, after sensory information has become available (see below). Thus planning is not completed before movement onset, even if we allow the "plan" to be as sophisticated as a control law. To see why that is true in general, consider an extreme case such as playing computer chess. The typical algorithm for such games is a minimax algorithm, that for each move searches through the tree of all possible moves N steps ahead, and chooses the one that has the most desirable outcome. However, at depth N the algorithm has to stop (due to time limitations) and use some heuristic to evaluate the current situation. At the next move, it is possible to search the tree one step further, i.e. the heuristic will be applied to the next level down and that can potentially change the previous estimates of desirability (since the heuristic is not exact). Thus such an algorithm is essentially doing the same amount of planning at each time step (here opponent's moves play the role of sensory feedback, i.e. restrict the space of possible scenarios).

While simple tasks such as via-point movements are more similar to the tractable case where a control law can be computed at least approximately, these tasks are not representative of everyday goal-directed movements. In particular they ignore mechanical redundancy, dynamic manipulation of objects with uncertain physical properties, and in general any nontrivial planning. As soon as we leave the domain of $2D$ transport movements between targets, it becomes very unlikely that optimal control laws could be computed before movement onset (in fact, a number of proofs exist that exact spatial planning problems in robotics have exponential space complexity [Canny 88] - basically meaning "intractable").

One might then ask, what can we learn from a model of via-point tasks if it is clear that the same algorithm cannot be applied to more interesting tasks? The answer is that the computational model is different from the algorithm. Our goal here is to study the properties of a system that is attempting to perform a certain computation (maximize performance subject to task constraints). The fact that we can perform that computation accurately only under restricted conditions implies that the biological system may have to resort to approximations in more complex situations. While that would affect the behavior of the system in some implementation-dependent way, we would expect that effect to be of "second order" compared to the properties emerging on the computational level, especially in tasks where we know that significant approximations are not necessary.

3.3 Stochastic Optimal Control

In this section we summarize the classic results of stochastic optimal control theory, which will be modified later to suit our needs. Consider a system with state X_k where X is a vector of real numbers and k is a discrete time index. At each time step a control signal U_k is selected by a controller (which we are to design), an observation Y_k is obtained as a function of the current state, and the state makes a transition to a new state:

$$\begin{aligned} X_{k+1} &= AX_k + BU_k + C\eta_k \\ Y_k &= HX_k + G\omega_k \end{aligned} \tag{3.1}$$

where $\eta \sim N(0, \mathrm{I})$ and $\omega \sim N(0, \mathrm{I})$ are normally distributed random variables with zero mean and unit covariance, corresponding to system and observation noise respectively[3]. The controller can only observe the state of the system through noisy measurements, and infer its evolution from noisy dynamics. Given estimates of the mean $(\widehat{X}_1)$ and covariance (Σ_1) of the state at time 1, and all observations $Y^{k-1} \triangleq [Y_1, \ldots, Y_{k-1}]$ and controls $U^{k-1} \triangleq [U_1, \ldots, U_{k-1}]$ up to time $k-1$, the job of the controller is to find the functions $U_k^*\left(Y^{k-1}, U^{k-1}, \widehat{X}_1, \Sigma_1\right)$ that minimize the expected cost:

$$J = \mathrm{E}\left(X_N^{\mathrm{T}} Q_N X_N + \sum_{k=1}^{N-1} \left(X_k^{\mathrm{T}} Q_k X_k + U_k^{\mathrm{T}} R U_k\right)\right) \tag{3.2}$$

The duration of the movement N, as well as the matrices in the system and cost equations are known. This problem is called the Linear Quadratic Gaussian (LQG) regulator - since the system equations are linear, the cost function is quadratic, and the noise model is Gaussian.

Minimizing such a cost function requires at each time step a choice of U_k^* which minimizes expected cost by taking into account all possible sequences of future controls and observations:

$$\min_{U_k}\left(U_k^{\mathrm{T}} R U_k + \mathrm{E}\left(X_k^{\mathrm{T}} Q_k X_k\right) + \mathop{\mathrm{E}}_{Y_k}\left(\min_{U_{k+1}}\left(U_{k+1}^{\mathrm{T}} R U_{k+1} + \mathrm{E}\left(X_{k+1}^{\mathrm{T}} Q_{k+1} X_{k+1}\right) + \mathop{\mathrm{E}}_{Y_k+1}(\cdots)\right)\right)\right)$$

which in general amounts to an exhaustive search of an exponentially large space. Thus it is quite remarkable that in the LQG case the problem has an exact solution (which is a major success of contemporary mathematics and has given rise to an immense literature - see [Davis and Vinter 85] for a nice introduction). Intuitively, the problem is tractable because the expected cost at time k (called "cost-to-go") can be evaluated recursively without performing an exhaustive search (unlike computer chess for example). The solution to this problem is given by the following theorems:

Theorem 1 (Kalman Filtering Algorithm) *The estimates of the mean and covariance of the state are propagated forward in time according to:*

$$\begin{aligned}
\widehat{X}_{k+1} &= A\widehat{X}_k + BU_k + K_k\left(Y_k - H\widehat{X}_k\right) \\
K_k &= A\Sigma_k H^{\mathrm{T}}\left(H\Sigma_k H^{\mathrm{T}} + GG^{\mathrm{T}}\right)^{-1} \\
\Sigma_{k+1} &= A\Sigma_k A^{\mathrm{T}} + CC^{\mathrm{T}} - A\Sigma_k H^{\mathrm{T}}\left(H\Sigma_k H^{\mathrm{T}} + GG^{\mathrm{T}}\right)^{-1} H\Sigma_k A^{\mathrm{T}}
\end{aligned}$$

starting with the given estimates $\widehat{X}_1$ and Σ_1. K_k is called the Kalman gain matrix.

Theorem 2 (Optimal Control Law) *The optimal control law is obtained from the state estimates of the Kalman filter, and the constant matrices L and S propagated backwards in time:*

$$\begin{aligned}
U_k^* &= -L_k\widehat{X}_k \\
L_k &= \left(B^{\mathrm{T}} S_{k+1} B + R\right)^{-1} B^{\mathrm{T}} S_{k+1} A \\
S_k &= AS_{k+1}A^{\mathrm{T}} + Q_k - AS_{k+1}B^{\mathrm{T}}\left(BS_{k+1}B^{\mathrm{T}} + R\right)^{-1} BS_{k+1}A^{\mathrm{T}}
\end{aligned}$$

[3] The known matrices A, B, C, H, G are considered constant here for simplicity, but they become time/state dependent below.

starting with $S_N = Q_N$*.*

Theorem 3 (Expected Cost) *The expected cost of the movement is given by:*

$$J = \widehat{X}_1^{\mathrm{T}} S_1 \widehat{X}_1 + \sum_{k=1}^{N-1} \left(\mathrm{tr}\left(S_{k+1} R\right) + \mathrm{tr}\left(\left(Q_k - S_k + A^{\mathrm{T}} S_{k+1} A\right) \Sigma_k\right)\right)$$

A few properties of the solution are worth noting. The recursive equations for Σ_k and S_k are identical up to a time reversal, and are known as discrete-time Riccati equations. These quantities can be computed before the movement, i.e. they do not depend on the control signals U_k and observations Y_k. This convenient property of the LQG regulator is due to the fact that all sources of noise are Gaussian and additive, and thus the amount of uncertainty cannot be affected by the actions of the controller or the actual observations. The fact that the optimal control law is independent of any quantities related to the noise model, and the filtering algorithm is independent of any quantities specifying the cost function is called the *separation principle*. Note also that the control law is computed as if the current estimate $\widehat{X}$ were the true value of X (the optimal control law for a completely observable system has the same form). This is known as the *certainty-equivalence principle*. As we will see shortly none of these principles apply directly in the case of our model, which leads to some technical complications.

Recall our discussion of a Bayesian model of sensory-motor processing from the previous chapter. The main idea was that the sensory-motor system uses a causal model to predict the outcomes of hypothetical actions, and a generative model to predict the sensory stimuli arising from hypothetical states of the world. These two models are then "inverted" via some powerful computation. The mathematical model in this section does exactly that: the first line of (3.1) is the causal model describing the hand dynamics, and the second line is the generative model describing how sensory inputs arise. The above theorems "invert" the model analytically.

3.4 Adaptation to Eye-Hand Coordination in Via-Point Tasks

This section formulates the problem of eye-hand coordination in via-point tasks in the language of stochastic optimal control, and introduces the modifications necessary to make the LQG framework applicable to motor control. Rather than trying to remain as close as possible to the basic model at the price of unfounded assumptions, we prefer to develop a model that would actually reflect the important aspects of biological sensory-motor processing and then find a way to solve it.

3.4.1 State formulation

In the task considered here a subject is asked to move the hand through a sequence of specified targets, with some timing and accuracy constraints. Passage times through via-points, eye movement patterns, and kinematic details of the trajectory (other than passing through the targets) are not specified, thus we assume that they are being optimized so as to achieve the best possible performance (to be defined below). Subjects receive visual information about hand and target positions, and proprioceptive information about hand and eye fixation positions. All positions will be represented in body centered coordinates; assuming that the body is fixed relative to the experimental setup (which it is), world centered and body centered coordinates coincide[4].

The basic quantities that need to be represented at each point in time (and thus constitute the "state" X of the system) are hand position (H_x, H_y) and target positions $\left(T_x^p, T_y^p\right)_{p=1\ldots P}$, where the starting position is target 1 and the final position is target P. The control signal is the force applied at the hand. In order to evaluate the effect of force (which is equal to acceleration for unit mass), we need to keep track of both position and velocity $\left(\dot{H}_x, \dot{H}_y\right)$ of the hand. Defining

$$X \triangleq \begin{bmatrix} H_x \\ H_y \\ \dot{H}_x \\ \dot{H}_y \\ T_x^1 \\ T_y^1 \\ \vdots \\ T_x^P \\ T_y^P \end{bmatrix}, \quad U \triangleq \begin{bmatrix} \ddot{H}_x \\ \ddot{H}_y \end{bmatrix}, \quad A \triangleq \begin{bmatrix} 1 & 0 & dt & 0 & & & 0 \\ 0 & 1 & 0 & dt & & & \\ & & 1 & & & & \\ & & & & \ddots & & \\ & & & & & 1 & \\ 0 & & & & & & 1 \end{bmatrix}, \quad B \triangleq \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ dt & 0 \\ 0 & dt \\ \vdots & \vdots \\ 0 & 0 \end{bmatrix}$$

where dt is the discretization time step, we have system dynamics:

$$X_{k+1} = AX_k + BU_k \tag{3.3}$$

This is a matrix notation for the discrete-time equations $H\,(t+dt) = H\,(t) + \dot{H}\,(t)\,dt$, $\dot{H}\,(t+dt) = \dot{H}\,(t) + \ddot{H}\,(t)\,dt$, and $T(t+dt) = T(t)$ - the last equation corresponds to the fact that targets are assumed stationary. It is also convenient to propagate a constant (equal to 1) in the state vector X to allow the encoding of task constraints defined by constant quantities. Note that the state and control do not contain eye movement information. Eye movements will be handled differently.

3.4.2 Control model and biomechanical muscle properties

In the state formulation above control signals U are assumed to be equal to the net muscle force generated at the hand. This would be a natural assumption in controlling a robot arm driven by pure torque actuators. Muscles however are different from pure torque actuators in some important ways: the amount of force they produce is not only a function of neural activation, but is also affected by the length-tension and force-velocity characteristics of the muscle [Kandel et al. 91, McMahon 84]. Furthermore the motor neuronal pool driving muscles is not homogeneous, but can be separated in two components: α motor neurons driving extrafusal muscle fibers whose activation is analogous to pure force generation, and γ motor neurons driving intrafusal muscle fibers (containing spindle receptors) whose activation determines the equilibrium point for the stretch reflex. A number of servo control theories of biological movement have been proposed according to which the CNS controls directly the γ motor neurons (i.e.

[4] While there have been many arguments regarding the nature of the coordinates used by the CNS to represent space, the actual coordinate system would not affect the behavior of the model provided there is a one-to-one mapping to world-centered Cartesian coordinates (e.g. polar shoulder-centered 2D coordinates).

sets the equilibrium point for the musculo-skeletal system), and the actual movement results from the $\alpha - \gamma$ coactivation loop and spring-like mechanical properties of muscles [Bizzi et al. 92]. There are some major problems with such models:

- the observed activation of γ motor units does not precede the activation of α motor units by a sufficient amount of time [McMahon 84]
- the observed stiffness (gain of the stretch reflex plus muscle elasticity) appears too low to account for most rapid movements [Bennett et al. 92, Gomi and Kawato 96], and in general the "simplification" resulting from this type of control (see Chapter 2) does not seem to justify throwing away the possibility of direct force control[5]
- force/velocity (instead of equilibrium position alone) appear to be represented in the,ajority cortical motor areas studied so far [Georgopoulos et al. 82, Evarts 68, Johnson 92].

Thus we will not use pure equilibrium-point control. However it is clear that muscles have a significant amount of damping (which is not due to passive viscosity, but results from the shape of the force-velocity curve) and that stiffness (resulting from the slope of the muscle length-tension curve as well as the stretch reflex) is always present [Kandel et al. 91]. It is desirable to be able to combine force control with equilibrium-point control in the model. Indeed, the minimum-jerk and minimum-torque-change models were first developed on the level of the actual arm movement (implying that the control signal sent by the CNS is force or torque). Later both models were modified to take stiffness and damping into account, resulting in a better fit to experimental data. The improved accuracy of the modified models (assuming it is not due to the presence of extra parameters) implies that the CNS controls directly a higher-level "force", before the effects of arm impedance have been added.

Formally, assume that the acceleration of the hand results from a combination of PD control and force control:

$$\ddot{H} = U_a - b\dot{H} + k(U_e - H)$$

Here b and k are constants corresponding to damping and stiffness respectively, we are assuming mass $m = 1$, U_a is the acceleration part of the control signal, and U_e is an additional control signal specifying the equilibrium point for the stretch reflex. Note that U_e is not the "desired" hand position and U_a is not the "desired" hand acceleration, since the controller is aware of the actual hand dynamics. It may be interesting to study the optimal way to vary b and k by including them in the control signal, however this is beyond the scope of the present model (the system dynamics becomes nonlinear and requires an additional linearization to obtain the optimal control law).

Substituting $\ddot{H}$ for U in (3.3) and noting that H and $\dot{H}$ can be extracted from the current state X, we can redefine the system dynamics:

$$X_{k+1} = \widetilde{A}X_k + \widetilde{B}\widetilde{U}_k$$

[5] It has recently been shown [Gribble et al. 98] that a nonlinear impedance model can account for these findings, at the cost of extra assumptions that seem realistic but have not yet been verified experimentally. On the other hand, [Kawato 98] trained a neural network (capable of learning almost arbitrary impedance functions) on a large amount of perturbation data, and used the trained network to infer the equilibrium trajectory - confirming their previous conclusion that it is "more complex" than the actual trajectory. In either case, when a nonlinear impedance model is used to produce complex curved movements, it is not very clear how to move the equilibrium point of the arm (and set the other parameters of the model) in order to achieve the desired movement.

where

$$\widetilde{A} = A - kB\mathrm{I}_H - bB\mathrm{I}_{\dot{H}}, \quad \widetilde{B} = \left[\begin{array}{cc} B & kB \end{array} \right], \quad \widetilde{U} = \left[\begin{array}{c} U_a \\ U_e \end{array} \right]$$

and I_H and $\mathrm{I}_{\dot{H}}$ are identity matrices padded with 0's that extract H and $\dot{H}$ from X. Thus adding PD control to the direct force control model leaves the system dynamics in the same form. The remaining part of the derivation does not depend on the exact values of A and B, so PD control can be added independently of other modifications.

This model allows the controller to set U_a, U_e independently, so it can achieve the same endpoint force with different control signals. Fitting such a model to data is difficult when the additional costs (see below) associated with the different components of the control signal are unknown, e.g. if we assume that only U_a incurs an additional cost, the controller will set it to 0 and use U_e to move the hand, which is unrealistic. Also, the equilibrium point of the musculo-skeletal system cannot be controlled completely independently: while the system is free to set the activity of γ and α motor units independently, the stiffness depends on both muscle activation (U_a) on a short time scale and the stretch reflex on a longer time scale [Nicols and Houk 73]. Thus in the simulations we will use a simpler version, which preserves the damping and stiffness terms in the system dynamics $\widetilde{A}$ but combines the two control signals into a single U (i.e. we use the original B and U defined above).

3.4.3 Neuromotor noise

Comparing (3.1) to (3.3) we see that the only difference is the presence of noise in the LQG regulator. Noise is a very important concept in modelling biological sensory-motor processing. It is present both in the sensory input (see below) and the motor output of the system. The latter form of noise (called neuromotor noise, [Meyer et al. 82]) is used to describe the uncertainty of the actual force produced by the muscles for a given central signal. Thus it is not very realistic to model neuromotor noise as being additive, or dependent on the state of the arm - since the only unpredictable forces affecting the arm in the type of movements considered here are muscle forces. It has been established that the error in producing a desired force level by human subjects is proportional to the amount of force [Schmidt et al. 79] suggesting a noise term multiplying the control signal in our model:

$$X_{k+1} = AX_k + (B + \eta_k C)\, U_k + D\delta_k \tag{3.4}$$

where $\eta_k \sim N(0,1)$ is a Gaussian random variable, $\delta_k \sim N(0,\mathrm{I})$ has a multivariate normal distribution, $C = \sigma_U B$ where σ_U is the standard deviation of the control noise. We also preserve the additive noise term to be able to model other sources of noise, such as small body shifts relative to the external world (see below). If PD control is being used, the noise term will affect both the direct force component of the control signal as well as the equilibrium position component. Note that our model of control noise corresponds to correlated disturbances in all components of the control signal U, i.e. the force disturbance will always lie along the current force vector (which is reasonable since each muscle applies end-point force in one direction, and muscles that are more active are assumed more noisy).

As discussed in Chapter 2, the most likely source of noise is not true randomness (which is what the above model implies) but a failure to take into account all complexities of the controlled system. For example, muscles may produce a force different from the desired level because their current metabolic state, moment arm and geometry, recruitment level, position along the length-tension curve, etc. are not estimated accurately. Such inaccuracies are likely to fluctuate slowly over time. Thus it may be desirable to use a temporally correlated instead of a white noise model. A temporally correlated process has its own state or memory:

$$\begin{aligned} X_{k+1} &= AX_k + BU_k + \xi_k \\ \xi_{k+1} &= a\xi_k + \eta_k CU_k + D\delta_k \end{aligned}$$

where a corresponds to the amount of temporal correlation ($a = 0$ recovers the white noise process). Again we can rewrite the modified system in the form:

$$\widetilde{X}_{k+1} = \widetilde{A}\widetilde{X}_k + \left(\widetilde{B} + \eta_k \widetilde{C}\right) U_k + \widetilde{D}\delta_k$$

where

$$\widetilde{X} = \begin{bmatrix} X \\ \xi \end{bmatrix}, \quad \widetilde{A} = \begin{bmatrix} A & \mathrm{I} \\ 0 & a\mathrm{I} \end{bmatrix}, \quad \widetilde{B} = \begin{bmatrix} B \\ 0 \end{bmatrix}, \quad \widetilde{C} = \begin{bmatrix} 0 \\ C \end{bmatrix}, \quad \widetilde{D} = \begin{bmatrix} 0 \\ D \end{bmatrix}$$

Thus adding temporally correlated noise leaves the system in the form (3.4), so it can be added independently of other modifications[6].

3.4.4 Performance index and timing

To define a performance index similar to the cost function (3.2) we have to fix the matrices Q_k, R and the desired movement duration N. When movement duration is specified fixing N is straightforward. However, even in those cases the specified movement duration will not always be achieved exactly. Thus it may seem appropriate to add the difference between the specified and actual movement durations to the cost function, or add a penalty term that increases with movement duration (as in [Hoff 94]). In the present model this is very difficult since the actual movement duration is not explicitly represented in the state equations, so we will have to infer it from the trajectory. Evaluating an expression over the entire trajectory (after the movement is completed) makes it impossible to define a cost-to-go function, which eliminates all advantages of using the LQG framework. Instead of adding an explicit movement duration cost, we will incorporate it implicitly by requiring hand velocity to vanish at time N - thus movements with incorrect duration will be penalized. By varying the relative weight of the terminal velocity constraint we can indirectly vary the weight of the duration constraint. This approximation may not be appropriate when movement duration is the variable of interest, but it seems reasonable in cases where duration is either specified, or it has a simple scaling effect.

The Q_k matrices encode the requirement to pass through targets (and stop at the end). When intermediate targets are present, the passage times through them are not known in advance so we do not know which Q_k to set. Note that this indeterminacy is different from unknown movement duration: varying via-point passage times does not require an additional arbitrary weight, which is necessary with unknown duration to prevent infinitely slow and infinitely accurate movements. Our strategy is to solve the problem with fixed passage times $K = [1, k_2, \ldots, k_{P-1}, N]$ and then implement an outer minimization loop over the vector K. A good way to initialize K is to extract it from the minimum-jerk trajectory

[6] Strictly speaking the noise in this form is proportional to the previous control signal instead of the current one, but that is not likely to make a difference since consequtive controls are very similar.

(see Chapter 6).

Suppose that at time k_i the hand H is required to pass through target T^i. Define a matrix D_{ji} (unrelated to the additive noise term above) that extracts the difference between targets j and i (labeling the hand as target 0):

$$\begin{aligned} D_{0i} &= \begin{bmatrix} 1 & 0 & \cdots & -1 & 0 & \cdots \\ 0 & 1 & \cdots & 0 & -1 & \cdots \end{bmatrix}, \qquad D_{0i}X = \begin{bmatrix} H_x - T_x^i \\ H_y - T_y^i \end{bmatrix} \\ (D_{0i}X)^{\mathrm{T}} D_{0i}X &= X^{\mathrm{T}} D_{0i}^{\mathrm{T}} D_{0i} X = \left(H_x - T_x^i\right)^2 + \left(H_y - T_y^i\right)^2 \end{aligned}$$

suggesting that Q should be defined as:

$$Q_{k_i} \triangleq w_i D_{0i}^{\mathrm{T}} D_{0i}, \qquad Q_{k \notin K} \triangleq 0$$

where w_i corresponds to the width of target i, or the required accuracy at that target. To encode the terminal velocity constraint (with weight v) we have to add a term:

$$Q_N = Q_{k_P} + v D^{v\mathrm{T}} D^v, \qquad D^v = \begin{bmatrix} 0 & 0 & 1 & 0 & \cdots & 0 \\ 0 & 0 & 0 & 1 & \cdots & 0 \end{bmatrix}$$

It is also possible to encode constraints that enforce a velocity other than 0 at specified points in time, by using the constant 1 we are propagating in the state vector X.

Finally, the LQG cost function (3.2) includes a penalty for the control signal. There are two reasons for including this extra cost in engineering applications:

- the square of the control signal $U^{\mathrm{T}}RU$ corresponds to some physical quantity (e.g. energy) that we would like to minimize
- including a "regularizer" turns and ill-posed problem (one with infinitely many solutions) into a well-posed problem.

In biological motor control it could make sense to minimize metabolic activity, or muscle tension, or neuronal firing, etc. all of which correspond to an energy-like term in the cost function. However, comparisons of minimization models based on derivatives of different order strongly suggests that if anything is being minimized, it is the squared derivative of acceleration rather than squared acceleration itself. Thus we will include the term $(U_k - U_{k-1})^{\mathrm{T}} R (U_k - U_{k-1})$ rather than $U_k{}^{\mathrm{T}} R U_k$. This additional cost term represents a deviation from our minimalist principle of constructing a model based on task requirements only; the evidence for it will be discussed extensively later. Here we note that setting $R = 0$ eliminates this term, and setting to 0 a certain parameter that will appear later actually transforms this term into $U_k{}^{\mathrm{T}} R U_k$, so we can study different versions of the additional cost in the same general model. We use $R = \mathrm{I}r/dt^2$ which penalizes the length of the jerk vector with weight r. With the above definitions, the cost function for our model becomes:

$$J = \mathrm{E}\left(\sum_{k=1}^{N} X_k^{\mathrm{T}} Q_k X_k + \sum_{k=1}^{N} (U_k - U_{k-1})^{\mathrm{T}} R (U_k - U_{k-1})\right) \tag{3.5}$$

It is also possible to penalize jerk instead of acceleration by including acceleration in the system state, defining the control signal to be jerk, and using a ${U_k}^{\mathrm{T}} R U_k$ term as in the LQG cost (e.g. [Hoff 94]). This is somewhat less natural, and also leads to difficulties when combined with imperfect state observation and a noise model proportional to acceleration - the noise becomes proportional to the (uncertain) state, complicating matters later in the derivation.

3.4.5 Motor templates and learning

We argued in Chapter 2 against the "motor program" view according to which movements result from the execution of predefined templates, and in favor of a more robust model which always maintains the goal of the movement and attempts to find the control signals that maximize expected performance. The model developed thus far clearly belongs to latter category. If the plant is modeled accurately (i.e. all noise terms are unpredictable) and there is sufficient time to complete the required computations on-line, such a model would indeed be optimal and motor templates would not contribute anything to performance. However, it is likely that some errors are state-dependent and the motor system does not always have the time to complete all computations. In those cases it may be advantageous to store information about the control signals used on previous trials.

How should that prior information be represented and incorporated into the optimal control model? The above form of the cost function (3.5) allows an alternative interpretation of the additional control-dependent term, suggesting a natural way to use information compiled from previous trials. At time k we are penalizing deviations of the control signal U_k away from the "desired" value U_{k-1} (which is computed at the previous time step and therefore can be treated as a known constant). Thus U_{k-1} can be seen as a "hint" about the next control signal, and R as a measure of how seriously that hint should be taken. The reason we chose U_{k-1} as the desired value for U_k was to penalize jerk. Instead we could use a different hint, e.g. the control signal $\overline{U}_k$ compiled from previous successful trials. There are two disadvantages to using an explicit "template" $\overline{U}_{1\ldots N}$:

- Since the template is not state-dependent, it becomes unusable if the configuration of targets is translated or rotated or scaled (which is suboptimal from a theoretical point of view and also contradicts observations of kinematic invariances in planar movements)
- The "real" output of the optimal controller is the function $U^* = L\widehat{X}$ defined by the matrices L, while the actual control signal U is just the value of the optimal control law for a (possibly noisy) state estimate $\widehat{X}$.

Thus it may be preferable to store previously computed control laws $\overline{L}$ rather than control signals $\overline{U}$, and compute the hints as $\overline{U}_k = \overline{L}_k\widehat{X}$. As we will see below it is convenient (due to the sensory delay) to augment the state with the most recent control signals, and so the previous control U_{k-1} can be extracted from the augmented state $\mathcal{I}_k$: $U_{k-1} = \mathrm{I}_1\mathcal{I}_k$ where I_1 is a constant matrix to be defined later. Therefore the control cost term in (3.5) is already in "hint form", with "default" control law $\overline{L} = \mathrm{I}_1$. By using a different $\overline{L}$ we can make $\overline{U}$ dependent on $\widehat{X}$ as well as the most recent control signal. It may be interesting in future work to find a learning algorithm that compiles $\overline{L}$ from previous trials and combines it with the initial value I_1 in some "optimal" way.

The additional cost term above is not meant to change the optimal control law, but instead to "bias" the controller so that the optimal control law is computed faster and/or more accurately (whether that is indeed the case remains to be verified in simulations). This is reminiscent of numerical optimization methods (e.g. simulated annealing, relaxation methods, etc.) that find solutions to difficult problems by first approximating them with simpler ones, finding a solution, and then tracking it while the approximation is gradually eliminated [Press et al. 92]. Is it possible that the minimum-jerk cost (which takes the form of a default control law) has a similar function? As mentioned earlier, jerk is not a particularly meaningful quantity to minimize since it does not correspond to any physical resource that may be important to the motor system. It has been suggested [Hogan and Flash 87] that minimizing higher-order derivatives results in more "predictable" movements, which may be advantageous from a computational point of view (that notion has not been made more explicit since then). Although this suggestion is similar in spirit to the "bias" function proposed above, note that one can make movements predictable by adding an arbitrary state-dependent term to the cost function and so "predictability" in itself cannot explain why a particular derivative should be minimized.

In the above model it is easy to see why the default control law should minimize jerk. Suppose we were to start with arbitrary $\overline{L}_1$ and iterate the learning rule $\overline{L}_{T+1} = (1-\varepsilon)\,\overline{L}_T + \varepsilon L_T$ over a number of contexts (target configurations) while gradually decreasing the learning rate ε. What would $\overline{L}$ converge to? Due to symmetry, it cannot contain any specific "commands" such as turn left, turn right, speed up, slow down, etc. (e.g. for every target configuration requiring left turns there is a target configuration requiring right turns). $\overline{L}$ can only encode hints that are valid over the entire family of motor tasks. One obvious hint is that on average the control signal remains unchanged, i.e. the expected value of U_k conditional upon U_{k-1} (without any target information) is U_{k-1}. Thus the default control law will converge to a cost function requiring the control signal (whatever it is) to remain constant. The reason jerk appears to be minimized is a "side effect" of the fact that the control signal is proportional to acceleration.

3.4.6 Sensory model

Recalling (3.1), the sensory model specifies how the measurement vector depends on the current state and measurement noise: $Y = HX + G\omega$. We consider two channels of sensory input - vision and proprioception. An efference copy (all previous control signals) is also available to the controller, but that is not a part of the sensory model (note that the efference copy is also noisy, since the actual force produced can be different from the controlled force). As stated above, spatial information in the system state X is represented in a world/body-centered coordinate system (assumed to be Cartesian) so all sensory measurements have to be transformed into that coordinate system. In particular, proprioceptive information about hand position and eye fixation has to be integrated over several joints. There may be inaccuracies in these coordinate transformations, which we model as additional sensor noise. Thus the noise term G is the sum of "true" noise and transformation errors. The magnitude of the noise is known, i.e. the controller has estimates of the reliability of each sensor (these estimates may become inaccurate if an experimenter introduces feedback perturbations - see below).

It is well known that sensory-motor performance increases with processing time. For example, the experiments of [Hening et al. 88, Ghez et al. 97] show that visual information about a desired motor response is available well before the normal reaction time (RT), and as more processing time is allowed the response becomes more accurate: if subjects are asked to respond before the RT, the response is not completely random but instead approximates the specified response. The optimal control model can capture this phenomenon quite naturally. Consider a "pipeline" in which sensory inputs are being processed. Obtaining and processing one input takes Δ time steps ($\Delta < RT$) before it becomes available to the controller, and the variability of inputs arriving at the end of the pipeline is always the same (G). The optimal estimation procedure combines all available measurements to obtain an estimate whose accuracy increases over time. Thus a part of the controller is actually involved in sensory processing, corresponding to the Kalman filter above[7].

Now we turn to the model of visual processing and eye movements. Since the primary goal here is to study arm movements, we will be satisfied with a more phenomenological description of vision. In particular, we will assume that subjects always use saccades which seems to be the case (in fact, when targets were removed all subjects tested fixated somewhere in the middle of the workspace instead of using smooth pursuit eye movements to track their moving hand). Saccades are always directed at one of the stationary targets (see Chapter 4) and for our purposes are accurate (i.e. we combine possible corrective saccades in one "saccade" moving the eyes from one target to another). Exact fixation is achieved through some form of retinal error cancelation, such that no extra information is extracted from the eye movement signals (i.e. efference copy of eye movement is not used to control arm movement). Visual input is absent during the saccadic eye movement, assumed to be of fixed duration. This is modeled by adding a large constant to the vision-related noise terms in the measurement vector[8].

Thus the only parameters specifying eye movements are the times $[s_1, \dots s_{P-1}]$ of peak eye velocity (assuming targets are fixated in order, which is generally true after the hand passes the first target - see Chapter 4). Since target fixation is exact, proprioceptive information about eye position becomes proprioceptive information about the current fixation target (in absolute coordinates). The positions of all other targets (including the hand) are perceived relative to the fixation point, i.e. vision provides estimates about the vectors pointing from the fixation point to all other points. The accuracy of such estimates decreases with distance from the fovea (corresponding to decreasing visual acuity). Note that it is up to the controller to integrate relative retinal information into the body-centered coordinate system. Vision also provides information about hand velocity (there is no difference between relative and absolute velocity for a stationary fixation point). Motion sensitivity is known to extend much further into the periphery compared to spatial sensitivity, and in fact there is a bias towards higher stimulus velocities in peripheral vision [Jeannerod 88]. Thus we assume that the reliability of velocity measurements does not depend on distance from the fixation point.

Given the eye movement times, we can find the index of the fixation target $f(k)$ at each time step k. From the above discussion the measurement vector is:

$$\begin{aligned} Y_k &= H_{k-\Delta} X_{k-\Delta} + G_{k-\Delta}(X_{k-\Delta})\omega_k \qquad (3.6) \\ H_k X_k &= \begin{bmatrix} H \\ \dot{H} \\ T^{f(k)} \\ H - T^{f(k)} \\ T^1 - T^{f(k)} \\ \vdots \\ T^P - T^{f(k)} \end{bmatrix}, \quad H_k \triangleq \begin{bmatrix} \mathrm{I} & 0 & 0 & \cdots & 0 & \cdots & 0 \\ 0 & \mathrm{I} & 0 & \cdots & 0 & \cdots & 0 \\ 0 & 0 & 0 & \cdots & \mathrm{I} & \cdots & 0 \\ \mathrm{I} & 0 & 0 & \cdots & -\mathrm{I} & \cdots & 0 \\ 0 & 0 & \mathrm{I} & \cdots & -\mathrm{I} & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & -\mathrm{I} & \cdots & \mathrm{I} \end{bmatrix} \end{aligned}$$

where the first three rows contain absolute measurements and the last P rows contain relative measure-

[7] We will not consider different delays for vision and proprioception here. Note that increasing sensor delay has a very similar effect to increasing sensor noise (it increases the prediction period over which no measurements are available, and thus accumulates more prediction error resulting from system noise).

[8] In general, withdrawal of sensory input can be modeled as greatly increased noise in the corresponding terms of the measurement vector. We will use that approach to simulate conditions with limited visual feedback, or patients with loss of proprioception.

ments (there is no relative information corresponding to $T^{f(k)}$). Note that the observation model will change with every saccadic eye movement. The noise term $G_k(X_k)$ encodes the amount of variability in the corresponding measurements, and is a function of the state X since the amount of sensory noise depends on distances from the current fixation point. We will be more specific about G in the next section.

So far we have assumed that the eye movement times (and thus the function $f(k)$) are known. Since this is not the case, we will have to optimize over them as well, i.e. implement another outer minimization loop over the vector $[s_1, \ldots s_{P-1}]$. A reasonable initial estimate would be the vector of via-point passage times. The two timing minimizations can be performed simultaneously.

3.4.7 Memory limitations and decay

The optimal model of sensory processing has some problematic implications:

- If sensory feedback is withdrawn and no movements are made, the present estimates of the state and amount of uncertainty will remain unchanged forever[9]. Thus introducing a delay period after targets are presented transiently should not affect performance.

- For quantities that are assumed to be stationary (such as the targets and the hand in the absence of control signals), the estimates of the system improve with time[10]. Thus if subjects are allowed to view the targets for a sufficiently long period of time (i.e. acquire very accurate estimates), vision of the targets during movement should not affect performance.

It is well established that both of these predictions are false [Desmurget et al. 95, Rao and Ballard 95]. Instead, it appears that memory decays with time, and sensory inputs are sampled whenever possible. To account for the latter finding, [Ballard et al. 97] have proposed that spatial working memory capacity is limited (they consider memory to be an array of pointers that can bind to external locations) and motor strategies are optimized subject to that constraint. While this explains why the external world will be used as its own representation whenever possible, it is not clear why memory should decay when no new input is available - clearly that cannot be derived from a limited memory capacity, since the decaying information was already stored somehow.

One way to model these phenomena in the optimal control framework would be to add a cost term inversely related to the entropy of the probability distribution of the state, and artificially degrade the state estimate $\widehat{X}$ that the controller is propagating. However the controller has to be aware of these changes, otherwise they will not be reflected in the control law. The only way to make the controller aware of something is to include it in the system equations. Therefore we would have to augment the state vector with Σ, and modify the dynamics to include the Kalman filter plus some disturbance. This will make the system very complicated and would require substantial approximations.

Instead we propose the following simple modification that can intuitively account for both effects. Imagine that for some reason the controller is using inaccurate estimates of the noise terms in the system. In particular it is too confident in the sensory input (G is smaller than the actual amount of sensory noise) and too skeptical about the stability of the world (D is larger than the actual amount of additive system noise). Such a controller will use a Kalman gain K_k larger than the optimal value i.e. it will overcorrect based on sensory inputs; in particular, in the absence of inputs the gain will not become

[9]Formally, set $G = \infty$, $D = 0$, $U = 0$ corresponding to infinitely unreliable sensory input, no additive noise and no control signals. Substituting in Theorem 1, $K_k = 0$, $\widehat{X}_{k+1} = A\widehat{X}_k$, $\Sigma_{k+1} = A\Sigma_k A^{\mathrm{T}}$. Since $\det(A) = 1$ and the velocity component of X is 0, there is no change.

[10]This follows from the form of the Σ_k update in Theorem 1.

exactly 0 and thus the state estimate (memory) will gradually degrade. Furthermore, the estimated variance of the state Σ_k will be larger than its true value, thus from the point of view of the controller it will be advantageous to sample the sensory input as often as possible, and in particular look at the targets even if they have been present before the movement.

What could be the reasons for such biases? D may be set too large simply because under normal circumstances the body is not stationary relative to the external world, i.e. the experimental task is artificial in that respect and the sensory-motor system is not designed to, or doesn't need to take this into account. The inability to set $K = 0$ may result from a neural design in which the sensory pipeline feeds into the motor control system and since it usually provides useful information, there is no mechanism to disable it completely. Alternatively, it may be impossible to maintain a fixed pattern of neural firing in the absence of informative sensory input. Note that none of the two biases alone can explain both phenomena. Increasing D will not result in memory degradation - the controller will be convinced that its state estimate is inaccurate, but it will in fact be accurate (since in reality there is little additive noise in this task) and so performance will not be affected. Similarly, decreasing G will not result in higher uncertainty (quite the opposite) and thus the controller will not attempt to sample the targets whenever possible.

3.4.8 Sensory-motor adaptation

The final modification we consider deals with the possibility of sensory-motor adaptation. While the model presented so far has a limited capacity to "adapt" in the sense that it can estimate the positions of the hand and via-points differently as a result of an experimental perturbation, this adaptation will vanish as soon as the state estimate changes (i.e. it is valid within a single movement)[11]. Such a model would predict, for example, that aftereffects in prism adaptation can only last for one trial, which is clearly not the case [Welch 86]. Thus we need an extra mechanism for modeling adaptation. Intuitively, we need to estimate and store additional parameters describing the current displacement of vision or proprioception. These parameters are never observed directly, but they affect the sensory information about the hand and targets. We will only consider the case with a single displacement for each modality, although more complex forms of adaptation such as visuo-motor maps could also be handled (at the price of making the model nonlinear and linearizing the control law).

Assume that visual inputs are offset by a vector T_V and proprioceptive inputs about hand position are offset by T_H. Relative visual inputs will not be affected by that offset, so the only change is in the sensory input about the absolute locations of the hand and fixation target. The observation model becomes:

$$\widetilde{Y}_k \quad = \quad \widetilde{H}_{k-\Delta}\widetilde{X}_{k-\Delta} + \widetilde{G}_{k-\Delta}(\widetilde{X}_{k-\Delta})\omega_k$$

[11] Using a default control law $\overline{L}$ that corrects for systematic perturbations detected on previous trials can have effects quite similar to sensory adaptation - see Chapter 5.

$$\widetilde{H}_k \widetilde{X}_k = \begin{bmatrix} H - T_H \\ \dot{H} \\ T^{f(k)} - T_V \\ H - T^{f(k)} \\ T^1 - T^{f(k)} \\ \vdots \\ T^P - T^{f(k)} \end{bmatrix}, \quad \widetilde{H}_k \triangleq \begin{bmatrix} \mathrm{I} & 0 & 0 & \cdots & 0 & \cdots & 0 & -\mathrm{I} & 0 \\ 0 & \mathrm{I} & 0 & \cdots & 0 & \cdots & 0 & 0 & 0 \\ 0 & 0 & 0 & \cdots & \mathrm{I} & \cdots & 0 & 0 & -\mathrm{I} \\ \mathrm{I} & 0 & 0 & \cdots & -\mathrm{I} & \cdots & 0 & 0 & 0 \\ 0 & 0 & \mathrm{I} & \cdots & -\mathrm{I} & \cdots & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & -\mathrm{I} & \cdots & \mathrm{I} & 0 & 0 \end{bmatrix}$$

and the system dynamics is modified accordingly:

$$\begin{aligned} \widetilde{X}_{k+1} &= \widetilde{A}\widetilde{X}_k + \left(\widetilde{B} + \eta_k \widetilde{C}\right) U_k + \widetilde{D}\delta_k \\ \widetilde{X} &= \begin{bmatrix} X \\ T_H \\ T_V \end{bmatrix}, \quad \widetilde{A} = \begin{bmatrix} A & 0 & 0 \\ 0 & \mathrm{I} & 0 \\ 0 & 0 & \mathrm{I} \end{bmatrix}, \quad \widetilde{B} = \begin{bmatrix} B \\ 0 \\ 0 \end{bmatrix}, \quad \widetilde{C} = \begin{bmatrix} C \\ 0 \\ 0 \end{bmatrix}, \quad \widetilde{D} = \begin{bmatrix} D \\ 0 \\ 0 \end{bmatrix} \end{aligned}$$

Again, this modification leaves the system in its original form (only changing the values of the matrices) so it can be added independently of the other modifications. The new $\widetilde{A}$ encodes the assumption that the offsets are constant over time. A system with unknown parameters that are being estimated on-line is known as a self-tuning regulator, and augmenting the state with the unknown parameters (as we have done here) is one of the methods for solving it.

3.5 Optimal Control Law Derivation

Summarizing the preceding discussion, the model is:

$$\begin{aligned} X_{k+1} &= AX_k + (B + \eta_k C)\, U_k + D\delta_k \\ Y_k &= H_{k-\Delta} X_{k-\Delta} + G_{k-\Delta}\left(X_{k-\Delta}\right) \omega_k \\ J &= \mathrm{E}\left(\sum_{k=1}^{N} X_k^{\mathrm{T}} Q_k X_k + \sum_{k=1}^{N} (U_k - U_{k-1})^{\mathrm{T}} R (U_k - U_{k-1})\right) \end{aligned} \tag{3.7}$$

with state X_k, measurement Y_k, control U_k, cost function J, independent noise $\eta_k \sim N(0;1)$, $\omega_k \sim N(0;\mathrm{I})$, $\delta_k \sim N(0;\mathrm{I})$, fixed $A, B, C, D, H_k, G_k\left(X_k\right), Q_k, R, N, \Delta$, and $U_{0,-1,\cdots} \triangleq 0$.

Our goal is to find the optimal control law, i.e. the sequence of functions $U_k^*\left(Y^{k-1}, U^{k-1}, \widehat{X}_1, \Sigma_1\right)$ minimizing J. Assumptions about the values of the fixed matrices will be avoided to make the derivation applicable to any combination of modifications presented in the previous section.

3.5.1 State estimation

The first step is to derive the equations for propagating the estimated mean $\widehat{X}$ and covariance Σ of the state X forward in time, by taking into account all control signals and sensory inputs. Let $\widehat{X}_{p|k}$ and $\Sigma_{p|k}$

denote the expected mean and covariance of X_p given observations $Y_{1,\ldots,k}$, controls $U_{1,\ldots,k}$, and initial estimates $\widehat{X}_1$ and Σ_1. Because of the sensory delay $\Delta \geq 1$ it does not make sense to propagate $X_{k|k-1}$ and $\Sigma_{k|k-1}$ (as in the LQG case) since at time k the controller receives information about $X_{k-\Delta}$ rather than X_k. Instead we will propagate $\widehat{X}_{k-\Delta|k-1}$ and $\Sigma_{k-\Delta|k-1}$, and then use the system dynamics and the Δ most recent control signals to estimate the current state.

The derivation of the Kalman filter is based on the following lemmas:

Lemma 4 *Let X, Y have a joint normal distribution*

$$\begin{bmatrix} X \\ Y \end{bmatrix} \sim N\left(\begin{bmatrix} \widehat{X} \\ \widehat{Y} \end{bmatrix}; \begin{bmatrix} \Sigma_{XX} & \Sigma_{XY} \\ \Sigma_{YX} & \Sigma_{YY} \end{bmatrix}\right)$$

Then the conditional mean and covariance of X are given by

$$\begin{aligned} \mathrm{E}(X|Y) &= \widehat{X} + \Sigma_{XY}\Sigma_{YY}^{-1}\left(Y - \widehat{Y}\right) \\ \Sigma_{X|Y} &= \Sigma_{XX} - \Sigma_{XY}\Sigma_{YY}^{-1}\Sigma_{YX} \end{aligned}$$

Furthermore, if the joint distribution of X and Y is not normal, but the first and second moments are the same as above, then $\widehat{X} + \Sigma_{XY}\Sigma_{YY}^{-1}\left(Y - \widehat{Y}\right)$ is the optimal (i.e. unbiased, minimum variance) linear estimator of X given Y.

Lemma 5 *Let X have mean $\widehat{X}$ and covariance Σ. Then*

$$\begin{aligned} \mathrm{E}\left(X^{\mathrm{T}}QX\right) &= \widehat{X}^{\mathrm{T}}Q\widehat{X} + \mathrm{tr}\left(Q\Sigma\right) \\ \mathrm{Cov}\left(QX, QX\right) &= Q\Sigma Q^{\mathrm{T}} \\ \mathrm{Cov}\left(X, QX\right) &= \Sigma Q^{\mathrm{T}} \end{aligned}$$

Computation of the new estimates involves two steps: combining the new sensory input Y_k and the current estimates $\widehat{X}_{k-\Delta|k-1}$, $\Sigma_{k-\Delta|k-1}$ to obtain improved estimates $\widehat{X}_{k-\Delta|k}$, $\Sigma_{k-\Delta|k}$ and then using the system dynamics (3.7) and the corresponding control signal $U_{k-\Delta}$ to obtain estimates $\widehat{X}_{k+1-\Delta|k}$, $\Sigma_{k+1-\Delta|k}$ for the next time step. From (3.7) and lemma 5, the joint distribution of $X_{k-\Delta}$ and Y_k has mean

$$\begin{bmatrix} \widehat{X}_{k-\Delta|k-1} \\ H_{k-\Delta}\widehat{X}_{k-\Delta|k-1} \end{bmatrix}$$

and covariance

$$\begin{bmatrix} \Sigma_{k-\Delta|k-1} & \Sigma_{k-\Delta|k-1}H_{k-\Delta}^{\mathrm{T}} \\ H_{k-\Delta}\Sigma_{k-\Delta|k-1} & H_{k-\Delta}\Sigma_{k-\Delta|k-1}H_{k-\Delta}^{\mathrm{T}} + \mathrm{E}\left(G_{k-\Delta}\left(X_{k-\Delta}\right)G_{k-\Delta}^{\mathrm{T}}\left(X_{k-\Delta}\right)\right) \end{bmatrix}$$

We will choose (see below) a sensor noise model for which $\mathrm{E}(G(X)G^{\mathrm{T}}(X))$ can be expressed as a function of the mean and covariance of X which we will approximate with a function of the mean, i.e. $F_{k-\Delta}\left(\widehat{X}_{k-\Delta|k-1}\right)$. Note that the resulting joint distribution is no longer multinormal (which would be the case if G did not depend on X), thus lemma 4 only yields the optimal linear estimator of $X_{k-\Delta}$

given the new observation Y_k. This approximation is often made in practice, and we will use it here to obtain the "measurement update" equations:

$$\begin{aligned}
\widehat{X}_{k-\Delta|k} &= \widehat{X}_{k-\Delta|k-1} + \\
&\quad \Sigma_{k-\Delta|k-1} H_{k-\Delta}^{\mathrm{T}} \left(H_{k-\Delta} \Sigma_{k-\Delta|k-1} H_{k-\Delta}^{\mathrm{T}} + F_{k-\Delta} \right)^{-1} \left(Y_k - H_{k-\Delta} \widehat{X}_{k-\Delta|k-1} \right) \\
\Sigma_{k-\Delta|k} &= \Sigma_{k-\Delta|k-1} - \\
&\quad \Sigma_{k-\Delta|k-1} H_{k-\Delta}^{\mathrm{T}} \left(H_{k-\Delta} \Sigma_{k-\Delta|k-1} H_{k-\Delta}^{\mathrm{T}} + F_{k-\Delta} \right)^{-1} H_{k-\Delta} \Sigma_{k-\Delta|k-1}
\end{aligned}$$

Now we use the system dynamics (3.7) and the current control signal $U_{k-\Delta}$ to "predict" the state one step ahead:

$$\begin{aligned}
\widehat{X}_{k+1-\Delta|k} &= A \widehat{X}_{k-\Delta|k} + B U_{k-\Delta} \\
\Sigma_{k+1-\Delta|k} &= A \Sigma_{k-\Delta|k} A^{\mathrm{T}} + D D^{\mathrm{T}} + C U_{k-\Delta} U_{k-\Delta}^{\mathrm{T}} C^{\mathrm{T}}
\end{aligned}$$

Combining the two sets of equations above, we obtain the equivalent of the Kalman filter for our model.

Theorem 6 (Modified Kalman Filter) *The estimates of the mean and covariance of the state are propagated forward in time according to:*

$$\begin{aligned}
\widehat{X}_{k+1-\Delta|k} &= A \widehat{X}_{k-\Delta|k-1} + B U_{k-\Delta} + K_{k-\Delta} \left(Y_k - H_{k-\Delta} \widehat{X}_{k-\Delta|k-1} \right) \\
K_{k-\Delta} &= A \Sigma_{k-\Delta|k-1} H_{k-\Delta}^{\mathrm{T}} \left(H_{k-\Delta} \Sigma_{k-\Delta|k-1} H_{k-\Delta}^{\mathrm{T}} + F_{k-\Delta} \left(\widehat{X}_{k-\Delta|k-1} \right) \right)^{-1} \\
\Sigma_{k+1-\Delta|k} &= A \Sigma_{k-\Delta|k-1} A^{\mathrm{T}} + D D^{\mathrm{T}} + C U_{k-\Delta} U_{k-\Delta}^{\mathrm{T}} C^{\mathrm{T}} - \\
&\quad A \Sigma_{k-\Delta|k-1} H_{k-\Delta}^{\mathrm{T}} \left(H_{k-\Delta} \Sigma_{k-\Delta|k-1} H_{k-\Delta}^{\mathrm{T}} + F_{k-\Delta} \left(\widehat{X}_{k-\Delta|k-1} \right) \right)^{-1} H_{k-\Delta} \Sigma_{k-\Delta|k-1} A^{\mathrm{T}}
\end{aligned}$$

starting with $\Sigma_{1-\Delta|0} \triangleq \Sigma_1$ *and* $\widehat{X}_{1-\Delta|0} \triangleq \widehat{X}_1$.

The differences from the Kalman filter in Theorem 1 are two: the covariance update depends on the control signal (due to the presence of noise proportional to the control signal in the system equations), and the gain matrix and covariance update depend on the state estimate (through $F_{k-\Delta}$). Thus the sequences Σ_k, K_k can no longer be precomputed before the onset of the movement. However it is still the case that given the control signals, the state estimation is independent of the cost function.

3.5.2 Square root filter

The numerical solution of Riccati equations (necessary to propagate Σ_k) is known to be unstable - due to the subtraction in the update equation, Σ_k may not always be positive definite (which is theoretically impossible, but happens due to numerical errors) and that may cause problems with the inversion in the update equation for K_k. The problem is avoided by using a square root filter, which propagates a square root of Σ_k rather than Σ_k itself. This reduces by half the dynamic range necessary to represent small floating point numbers, and also guarantees that Σ_k remains positive definite (since it is now represented as the square of another matrix). Recall that for a positive definite matrix M a square root $M^{1/2}$ (which is unique modulo orthonormal transformations) satisfies $M^{1/2} M^{\mathrm{T}/2} = M$ and therefore $M^{-\mathrm{T}/2} M^{-1/2} = M^{-1}$. If M is also symmetric then there exists a symmetric square root (which is the one computed in Matlab for example). We will also need to perform QR decomposition, which for a given matrix M produces matrices Q, R such that $M = QR$ where Q is an orthonormal matrix and R is an upper triangular matrix of the same dimension as M. Both matrix square roots and QR decomposition can be computed efficiently.

Given estimates $\Sigma^{1/2}_{k-\Delta|k-1}$, $\widehat{X}_{k-\Delta|k-1}$ and (temporarily) defining $M \triangleq DD^{\mathrm{T}} + CU_{k-\Delta}U^{\mathrm{T}}_{k-\Delta}C^{\mathrm{T}}$ we can form the matrix on the left hand side below, and use QR decomposition to obtain the matrices on the right:

$$\begin{bmatrix} H_{k-\Delta}\Sigma^{1/2}_{k-\Delta|k-1} & F^{1/2}_{k-\Delta} & 0 \\ A\Sigma^{1/2}_{k-\Delta|k-1} & 0 & M^{1/2} \end{bmatrix} = \begin{bmatrix} V^{1/2}_{k-\Delta} & 0 & 0 \\ \widetilde{K}_{k-\Delta} & S^{1/2}_{k+1-\Delta} & 0 \end{bmatrix} Q$$

By taking the transpose of that equation, and multiplying from the right (and using the fact that $QQ^{\mathrm{T}} = \mathrm{I}$) we have:

$$\begin{bmatrix} H_{k-\Delta}\Sigma_{k-\Delta|k-1}H^{\mathrm{T}}_{k-\Delta} + F_{k-\Delta} & \cdots \\ \cdots & A\Sigma_{k-\Delta|k-1}A^{\mathrm{T}} + M \end{bmatrix} = \begin{bmatrix} V_{k-\Delta} & \cdots \\ \cdots & \widetilde{K}_{k-\Delta}\widetilde{K}^{\mathrm{T}}_{k-\Delta} + S_{k+1-\Delta} \end{bmatrix}$$

where the $\cdots$ entries are irrelevant. Equating the corresponding entries of the matrix equation and comparing to the Kalman filter above, we see that:

$$\begin{aligned} S_{k+1-\Delta} &= \Sigma_{k+1-\Delta|k} \\ \widetilde{K}_{k-\Delta} &= A\Sigma_{k-\Delta|k-1}H^{\mathrm{T}}_{k-\Delta}\left(H_{k-\Delta}\Sigma_{k-\Delta|k-1}H^{\mathrm{T}}_{k-\Delta} + F_{k-\Delta}\right)^{-\mathrm{T}/2} \\ K_{k-\Delta} &= \widetilde{K}_{k-\Delta}V^{-1/2}_{k-\Delta}, \qquad V_{k-\Delta} = H_{k-\Delta}\Sigma_{k-\Delta|k-1}H^{\mathrm{T}}_{k-\Delta} + F_{k-\Delta} \end{aligned}$$

therefore the QR decomposition has computed the square root of the state covariance, and by multiplying $\widetilde{K}$ and the inverse of $V^{1/2}$ we can recover the Kalman gain matrix $K_{k-\Delta}$. Instead of recovering $K_{k-\Delta}$, it is more convenient to define the "innovation" $\upsilon_k \triangleq V^{-1/2}_{k-\Delta}\left(Y_k - H_{k-\Delta}\widehat{X}_{k-\Delta|k-1}\right)$ which has zero mean and unit covariance (since the covariance of the $Y - HX$ term is exactly V). Now the update equations for the expected mean and covariance become

$$\begin{aligned} \widehat{X}_{k+1-\Delta|k} &= A\widehat{X}_{k-\Delta|k-1} + BU_{k-\Delta} + \widetilde{K}_{k-\Delta}\upsilon_k \\ \Sigma_{k+1-\Delta|k} &= A\Sigma_{k-\Delta|k-1}A^{\mathrm{T}} + DD^{\mathrm{T}} + CU_{k-\Delta}U^{\mathrm{T}}_{k-\Delta}C^{\mathrm{T}} - \widetilde{K}_{k-\Delta}\widetilde{K}^{\mathrm{T}}_{k-\Delta} \end{aligned} \tag{3.8}$$

3.5.3 Information vector and innovations form

Next we collect all information available up to time $k-1$ into an "information" vector, and rewrite the system dynamics and cost function using this information vector. The expected mean $\widehat{X}_{k-\Delta|k-1}$ incorporates all sensory inputs up to $k-1$ and all control signal up to $k-\Delta-1$. Thus the information vector is:

$$\mathcal{I}_k \triangleq \begin{bmatrix} \widehat{X}_{k-\Delta|k-1} \\ U_{k-\Delta} \\ \vdots \\ U_{k-1} \end{bmatrix}$$

and we can extract its components via multiplication by the matrices

$$\begin{aligned}\widehat{X}_{k-\Delta|k-1} &= \mathrm{I}_X \mathcal{I}_k, \quad \mathrm{I}_X = \begin{bmatrix} \mathrm{I} & 0 & \cdots & 0 \end{bmatrix} \\ U_{k-p} &= \mathrm{I}_p \mathcal{I}_k, \quad \mathrm{I}_p = \begin{bmatrix} 0 & \cdots & \mathrm{I} & \cdots \end{bmatrix}\end{aligned}$$

Using the innovation defined above, we can rewrite the system dynamics in terms of the information vector $\mathcal{I}$:

$$\mathcal{I}_{k+1} = \mathcal{A}\mathcal{I}_k + \mathcal{B}U_k + \mathcal{K}_k(\mathcal{I}_k)\upsilon_k \tag{3.9}$$

where

$$\mathcal{A} = \begin{bmatrix} A & B & & & 0 \\ & & \mathrm{I} & & \\ & & & \ddots & \\ & & & & \mathrm{I} \\ 0 & & & & 0 \end{bmatrix}, \quad \mathcal{B} = \begin{bmatrix} 0 \\ 0 \\ \vdots \\ 0 \\ \mathrm{I} \end{bmatrix}, \quad \mathcal{K}_k(\mathcal{I}_k) = \begin{bmatrix} \widetilde{K}_{k-\Delta}(\mathrm{I}_X \mathcal{I}_k) \\ 0 \\ 0 \\ \vdots \\ 0 \end{bmatrix}$$

This is called the innovations form of the system. Notice that the new "state" $\mathcal{I}$ is completely observable, and the sensor $(F_k\omega_k)$ and system $(\eta_k CU_k, D\delta_k)$ noise terms from (3.7) have been combined into one noise term $\mathcal{K}_k\upsilon_k$.

Before the cost function can be written in terms of $\mathcal{I}$, we need to compute the covariance at time k by integrating the dynamics for Δ time steps using the control signals encoded in $\mathcal{I}$:

$$\Sigma_{k|k-1} = A^{\Delta}\Sigma_{k-\Delta|k-1}A^{\mathrm{T}\Delta} + \sum_{p=1}^{\Delta} A^{p-1}\left(DD^{\mathrm{T}} + CU_{k-p}U_{k-p}^{\mathrm{T}}C^{\mathrm{T}}\right)A^{\mathrm{T}(p-1)}$$

Expanding the brackets and substituting $\mathrm{I}_p\mathcal{I}_k$ for U_{k-p}, the covariance at time k becomes:

$$\begin{aligned}\Sigma_{k|k-1} &= A^{\Delta}\Sigma_{k-\Delta|k-1}A^{\mathrm{T}\Delta} + \mathcal{D} + \sum_{p=1}^{\Delta} \mathcal{C}_p \mathcal{I}_k \mathcal{I}_k^{\mathrm{T}} \mathcal{C}_p^{\mathrm{T}} \\ \mathcal{D} &\triangleq \sum_{p=1}^{\Delta} A^{p-1}DD^{\mathrm{T}}A^{\mathrm{T}(p-1)}, \quad \mathcal{C}_p \triangleq A^{p-1}C\mathrm{I}_p\end{aligned}$$

We also need to predict X_k from the information vector $\mathcal{I}$:

$$\widehat{X}_{k|k-1} = \mathcal{M}\mathcal{I}_k, \quad \mathcal{M} = A^{\Delta}\mathrm{I}_X + \sum_{p=1}^{\Delta} A^{p-1}B\mathrm{I}_p$$

Finally we are ready to rewrite the cost function. Let $J_k^*\left(\mathcal{I}_k, \Sigma_{k-\Delta|k-1}\right)$ denote the minimum cost that can be achieved starting at time k (this is known as the "cost-to-go" function). From Bellman's principle of optimality, we have:

$$\begin{aligned} J_k^*\left(\mathcal{I}_k, \Sigma_{k-\Delta|k-1}\right) &= \underset{X_k}{\mathrm{E}}\left(X_k^{\mathrm{T}} Q_k X_k\right) + \\ & \min_{U_k}\left\{(U_k - U_{k-1})^{\mathrm{T}} R (U_k - U_{k-1}) + \underset{\upsilon_k}{\mathrm{E}}\left(J_{k+1}^*\left(\mathcal{I}_{k+1}, \Sigma_{k+1-\Delta|k}\right)\right)\right\} \end{aligned}$$

Clearly $J = J_1^*\left(\mathcal{I}_1, \Sigma_{1-\Delta|0}\right)$, where $\mathcal{I}_1 \triangleq \begin{bmatrix} \widehat{X}_1 \\ 0 \end{bmatrix}$, $\Sigma_{1-\Delta|0} \triangleq \Sigma_1$. Evaluating the expectation over X, grouping terms dependent on $\mathcal{I}$, and using $\mathrm{tr}\,(ABC) = \mathrm{tr}\,(CAB)$, $\mathrm{tr}(x^{\mathrm{T}} A x) = x^{\mathrm{T}} A x$, the linearity of the tr operator, and lemma 5 leads to:

$$\begin{aligned} J_k^* &= \mathcal{I}_k^{\mathrm{T}}\left(\mathcal{M}^{\mathrm{T}} Q_k \mathcal{M} + \sum_{p=1}^{\Delta} \mathcal{C}_p^{\mathrm{T}} Q_k \mathcal{C}_p\right)\mathcal{I}_k + \mathrm{tr}\left(A^{\mathrm{T}\Delta} Q_k A^{\Delta} \Sigma_{k-\Delta|k-1}\right) + \mathrm{tr}\left(Q_k \mathcal{D}\right) \\ &+ \min_{U_k}\left\{(U_k - \mathrm{I}_1 \mathcal{I}_k)^{\mathrm{T}} R (U_k - \mathrm{I}_1 \mathcal{I}_k) + \underset{\upsilon_k}{\mathrm{E}}\left(J_{k+1}^*\left(\mathcal{I}_{k+1}, \Sigma_{k+1-\Delta|k}\right)\right)\right\} \end{aligned} \tag{3.10}$$

Note that the last expression no longer depends on X.

3.5.4 Control law

The optimal control law U^* is computed recursively, using mathematical induction. We will define the general closed form of J_k^*, show that for $k = N$ the cost-to-go function is indeed in that form (which will provide initialization values for the parameters in the closed form) and then compute J_k^* under the assumption that J_{k+1}^* is already available in the proposed closed form.

Assume the cost-to-go function and optimal control law at time k are in the form

$$J_k^*\left(\mathcal{I}_k, \Sigma_{k-\Delta|k-1}\right) = \mathcal{I}_k^{\mathrm{T}} S_k \mathcal{I}_k + \mathrm{tr}\left(P_k \Sigma_{k-\Delta|k-1}\right) + \alpha_k, \qquad U_k^* = L_k \mathcal{I}_k$$

Since the cost-to-go function after time N is by definition 0, we can perform the minimization at time N explicitly

$$J_N^*\left(\mathcal{I}_N, \Sigma_{N-\Delta|N-1}\right) = \mathcal{I}_N^{\mathrm{T}} S_N \mathcal{I}_N + \mathrm{tr}\left(P_N \Sigma_{N-\Delta|N-1}\right) + \alpha_N$$

therefore at time N the cost-to-go function is in the proposed form, the initial values for the parameters are

$$S_N \triangleq \mathcal{M}^{\mathrm{T}} Q_N \mathcal{M} + \sum_{p=1}^{\Delta} \mathcal{C}_p^{\mathrm{T}} Q_N \mathcal{C}_p, \qquad P_N \triangleq A^{\mathrm{T}\Delta} Q_N A^{\Delta}, \qquad \alpha_N \triangleq \mathrm{tr}\left(Q_N \mathcal{D}\right)$$

and the optimal control law is

$$U_N^* = L_N \mathcal{I}_N, \qquad L_N = \mathrm{I}_1$$

The last step is to compute J_k^* explicitly. Since $J_{k+1}^*$is now in closed form (i.e. does not involve a minimization over J_{k+2}^*), we can use that expression to evaluate J_k^* and perform the minimization over U_k. First we have to evaluate the expectation term $\underset{\upsilon}{\mathrm{E}}\left(J^*\right)$ in (3.10). By substituting the information vector and covariance at the next time step according to (3.9) and (3.8), grouping terms and simplifying we get:

$$\underset{\upsilon_k}{\mathrm{E}}\left(J^*_{k+1}\right) = (\mathcal{A}\mathcal{I}_k + \mathcal{B}U_k)^{\mathrm{T}} S_{k+1}(\mathcal{A}\mathcal{I}_k + \mathcal{B}U_k) + \mathcal{I}_k^{\mathrm{T}}\left(\mathrm{I}_\Delta^{\mathrm{T}} C^{\mathrm{T}} P_{k+1} C \mathrm{I}_\Delta\right)\mathcal{I}_k + \alpha_{k+1} + \tag{3.11}$$
$$\mathrm{tr}\left(A^{\mathrm{T}} P_{k+1} A \Sigma_{k-\Delta|k-1}\right) + \mathrm{tr}\left(S_{k+1}\mathcal{K}_k\mathcal{K}_k^{\mathrm{T}} - P_{k+1}\widetilde{K}_{k-\Delta}\widetilde{K}_{k-\Delta}^{\mathrm{T}}\right) + \mathrm{tr}\left(P_{k+1} D D^{\mathrm{T}}\right)$$

The minimization over U_k in (3.10) can now be performed by setting the derivative to 0:

$$\frac{\partial J_k}{\partial U_k} = 2R(U_k - \mathrm{I}_1\mathcal{I}_k) + 2\mathcal{B}^{\mathrm{T}} S_{k+1}(\mathcal{A}\mathcal{I}_k + \mathcal{B}U_k) = 0$$

Solving the last equation for U_k yields the optimal control law:

$$U_k^* = L_k\mathcal{I}_k, \qquad L_k = \left(R + \mathcal{B}^{\mathrm{T}} S_{k+1}\mathcal{B}\right)^{-1}\left(R\mathrm{I}_1 - \mathcal{B}^{\mathrm{T}} S_{k+1}\mathcal{A}\right) \tag{3.12}$$

Finally, we can substitute the optimal control U_k^* in (3.10, 3.11) to obtain the function J_k^* in closed form, and find the recursive equations for the parameters. From the definition (3.9) $\mathcal{K}_k\mathcal{K}_k^{\mathrm{T}} = \begin{bmatrix} \widetilde{K}_{k-\Delta}\widetilde{K}_{k-\Delta}^{\mathrm{T}} & 0 \\ 0 & 0 \end{bmatrix}$ and therefore $\mathrm{tr}\left(S_{k+1}\mathcal{K}_k\mathcal{K}_k^{\mathrm{T}}\right) = \mathrm{tr}\left(S_{k+1}^X\widetilde{K}_{k-\Delta}\widetilde{K}_{k-\Delta}^{\mathrm{T}}\right)$ where S_{k+1}^X denotes the upper-left submatrix with dimensions equal to that of the original state vector X. The expression for J_k^* becomes:

$$\mathcal{I}_k^{\mathrm{T}}\left(\mathcal{A}^{\mathrm{T}} S_{k+1}\mathcal{A} + \mathcal{M}^{\mathrm{T}} Q_k\mathcal{M} + \sum_{p=1}^{\Delta}\mathcal{C}_p^{\mathrm{T}} Q_k\mathcal{C}_p + \mathrm{I}_\Delta^{\mathrm{T}} C^{\mathrm{T}} P_{k+1} C\mathrm{I}_\Delta + \mathrm{I}_1^{\mathrm{T}} R\mathrm{I}_1 - \left(R\mathrm{I}_1 - \mathcal{B}^{\mathrm{T}} S_{k+1}\mathcal{A}\right)^{\mathrm{T}} L_k\right)\mathcal{I}_k$$
$$+\,\mathrm{tr}\left(\left(A^{\mathrm{T}} P_{k+1} A + A^{\mathrm{T}\Delta} Q_k A^{\Delta}\right)\Sigma_{k-\Delta|k-1}\right)$$
$$+\,\mathrm{tr}\left(\left(S_{k+1}^X - P_{k+1}\right)\widetilde{K}_{k-\Delta}\widetilde{K}_{k-\Delta}^{\mathrm{T}}\right) + \mathrm{tr}\left(P_{k+1} D D^{\mathrm{T}}\right) + \mathrm{tr}\left(Q_k\mathcal{D}\right) + \alpha_{k+1}$$

At this point one might be tempted to read off the update equations for S, P, α from the first, second and third line of the above expression respectively. In the LQG case that would be possible. Note also that in the LQG control law (Theorem 2) the matrices P multiplying the covariance matrix Σ are not propagated explicitly (instead the corresponding terms are absorbed in the constants α). The reason we had to propagate P here is that the noise covariance is no longer independent of the control signals U and observations Y (the complication is due to the multiplicative system noise and state-dependent observation noise).

Recall that $\widetilde{K}_{k-\Delta}$ depends on both $\widehat{X}$ and Σ. If we were to take into account that dependence exactly, J_k^* would become a nonlinear function of $\widehat{X}$ and Σ (involving a matrix inversion and some arbitrary noise covariance $F(X)$) and it would be impossible to obtain a closed form expression for U^*. If we only had to compute U^* for one time step that would not be a major disaster (since we could use some efficient minimization algorithm), but the problem is that without an explicit expression for U^* we can no longer avoid the minimization operator in J_k^* - meaning that at each time step we would have to minimize over all possible control sequences until the end of the movement, which is clearly intractable. The way around this problems is to linearize the $\mathrm{tr}\left(\left(S_{k+1}^X - P_{k+1}\right)\widetilde{K}_{k-\Delta}\widetilde{K}_{k-\Delta}^{\mathrm{T}}\right)$ term around an estimate $\overline{X}, \overline{\Sigma}$, which yields additional terms in the update equations for S, P, α. The actual linearization is presented in the next subsection. To complete the derivation of the control law, assume we have:

$$\operatorname{tr}\left(\left(S_{k+1}^X - P_{k+1}\right)\widetilde{K}_{k-\Delta}\widetilde{K}_{k-\Delta}^{\mathrm{T}}\right) \approx \overline{g}_{k-\Delta}\left(\overline{X},\overline{\Sigma}\right) + \mathcal{I}_k^{\mathrm{T}}\mathcal{L}_{X_{k-\Delta}}\left(\overline{X},\overline{\Sigma}\right)\mathcal{I}_k + \operatorname{tr}\left(\mathcal{L}_{\Sigma_{k-\Delta}}\left(\overline{X},\overline{\Sigma}\right)\Sigma_{k-\Delta|k-1}\right)$$

Now we can read off the update equations:

Theorem 7 (Modified control law) *The (approximate) optimal control law is*

$$U_k^* = L_k\mathcal{I}_k, \qquad L_k = \left(R + \mathcal{B}^{\mathrm{T}}S_{k+1}\mathcal{B}\right)^{-1}\left(R\mathrm{I}_1 - \mathcal{B}^{\mathrm{T}}S_{k+1}\mathcal{A}\right)$$

where the constant matrices are propagated backwards in time according to

$$\begin{aligned}
S_k &= \mathcal{A}^{\mathrm{T}}S_{k+1}\mathcal{A} + \left(\mathcal{M}^{\mathrm{T}}Q_k\mathcal{M} + \sum_{p=1}^{\Delta}\mathcal{C}_p^{\mathrm{T}}Q_k\mathcal{C}_p + \mathrm{I}_\Delta^{\mathrm{T}}C^{\mathrm{T}}P_{k+1}C\mathrm{I}_\Delta + \mathrm{I}_1^{\mathrm{T}}R\mathrm{I}_1 + \mathcal{L}_{X_{k-\Delta}}\left(\overline{X},\overline{\Sigma}\right)\right) \\
&\quad - \left(R\mathrm{I}_1 - \mathcal{B}^{\mathrm{T}}S_{k+1}\mathcal{A}\right)^{\mathrm{T}}\left(R + \mathcal{B}^{\mathrm{T}}S_{k+1}\mathcal{B}\right)^{-1}\left(R\mathrm{I}_1 - \mathcal{B}^{\mathrm{T}}S_{k+1}\mathcal{A}\right) \\
P_k &= A^{\mathrm{T}}P_{k+1}A + A^{\mathrm{T}\Delta}Q_kA^{\Delta} + \mathcal{L}_{\Sigma_{k-\Delta}}\left(\overline{X},\overline{\Sigma}\right) \\
\alpha_k &= \alpha_{k+1} + \operatorname{tr}\left(P_{k+1}DD^{\mathrm{T}}\right) + \operatorname{tr}\left(Q_k\mathcal{D}\right) + \overline{g}_{k-\Delta}\left(\overline{X},\overline{\Sigma}\right)
\end{aligned}$$

with $S_N = \mathcal{M}^{\mathrm{T}}Q_N\mathcal{M} + \sum_{p=1}^{\Delta}\mathcal{C}_p^{\mathrm{T}}Q_N\mathcal{C}_p$*,* $P_N = A^{\mathrm{T}\Delta}Q_NA^{\Delta}$*,* $\alpha_N = \operatorname{tr}\left(Q_N\mathcal{D}\right)$*,* $L_N = \mathrm{I}_1$*.*

Note the similarity with the corresponding equations in the LQG case. S_k is still updated according to a Riccati equation, but the term in the brackets is more complicated and involves the linearization part $\mathcal{L}_{X_{k-\Delta}}$. The P_k update is new (as explained above). The $R\mathrm{I}_1$ term in the definition of L_k appeared because of the modified cost function penalizing squared difference of control signals. Recalling the above discussion of using a default control law, replacing $\mathrm{I}_1$by an arbitrary constant matrix does not affect the derivation. If we were to add an additional cost term $\left(U_k - \overline{L}_k\mathcal{I}_k\right)^{\mathrm{T}}R_L\left(U_k - \overline{L}_k\mathcal{I}_k\right)$ in (3.10) the only difference is that in the final result (Theorem 7) we have to add a $\overline{L}_k, R_L$ term every time a R, I_1 term appears. L_N becomes $\left(R + R_L\right)^{-1}\left(R_L\overline{L}_N + R\mathrm{I}_1\right)$ - since before simplification it was $R^{-1}R\mathrm{I}_1$ rather than I_1.

The expected cost is equal to J_1^* where the α_1 term has accumulated all the constant terms from $J_{2\ldots N}^*$.

Theorem 8 (Modified expected cost) *The (approximate) expected cost of the movement is*

$$\begin{aligned}
J &= \mathcal{I}_1^{\mathrm{T}}S_1\mathcal{I}_1 + \operatorname{tr}\left(P_1\Sigma_{1-\Delta|0}\right) + \alpha_1 \\
&= \widehat{X}_1^{\mathrm{T}}S_1^X\widehat{X}_1 + \operatorname{tr}\left(P_1\Sigma_1\right) + \sum_{k=1}^{N-1}\left(\operatorname{tr}\left(Q_k\mathcal{D}\right) + \operatorname{tr}\left(P_{k+1}DD^{\mathrm{T}}\right) + \overline{g}_{k-\Delta}\left(\overline{X},\overline{\Sigma}\right)\right) + \operatorname{tr}\left(Q_N\mathcal{D}\right)
\end{aligned}$$

There are several possible strategies for obtaining the estimates $\left(\overline{X}_k, \overline{\Sigma}_k\right)_{k=1\ldots N}$ necessary for the linearization. We could just set them to some constant, or obtain them once at the beginning of the movement (e.g. from the minimum-jerk trajectory), or compute the optimal control law iteratively by starting with some estimates and replacing them at each step with the new optimal trajectory. Any of those strategies can be used just once at the first time step, or at each time step depending on the numeric accuracy of the approximation.

3.5.5 Noise covariance and linearization

Before proceeding with the linearization we have to define the observation noise covariance $F_{k-\Delta}\left(\widehat{X}_{k-\Delta|k-1}\right)$. Recall from (3.6) that the first three entries of the measurement vector contain positions in absolute (body-centered) coordinates obtained through proprioception and visual movement detection, that we assume are independent of the fixation point. The remaining P entries contain difference vectors from the current fixation target to the hand and the remaining targets, and their reliability decreases with distance. Thus the matrix F will be a diagonal matrix with the first three entries being constants, and the remaining proportional to the squared distance from the fixation target T^f:

$$F\left(\widehat{X}\right) = \text{diag}\begin{bmatrix} \sigma_H^2 & \sigma_{\dot{H}}^2 & \sigma_f^2 & \gamma\left\|\widehat{H}-\widehat{T}^f\right\|^2 & \gamma\left\|\widehat{T}^1-\widehat{T}^f\right\|^2 & \cdots & \gamma\left\|\widehat{T}^P-\widehat{T}^f\right\|^2 \end{bmatrix}$$

i.e. standard deviation of visual position estimates increases linearly with distance from the fovea. Ideally F would be a function of X rather than $\widehat{X}$ but the resulting complications do not seem to be justified (the difference is not likely to be substantial, and using X instead of $\widehat{X}$ makes linearization w.r.t Σ much more difficult). Note that according to this model the variance of some observations may become exactly 0. This is not a problem since even if we have exact relative measurements, transforming them into body-centered coordinates requires a proprioceptive estimate of eye fixation which is always noisy ($\sigma_f^2 > 0$).

Suppressing time indices, from the definition of $\widetilde{K}$ (3.8 and Theorem 6) we have:

$$\begin{aligned} g(X,\Sigma) &= \text{tr}\left(\left(S^X - P\right)\widetilde{K}\widetilde{K}^{\mathrm{T}}\right) = \text{tr}\left(Z\Sigma H^{\mathrm{T}}\left(H\Sigma H^{\mathrm{T}} + F(X)\right)^{-1} H\Sigma\right) \\ Z &\triangleq A^{\mathrm{T}}\left(S^X - P\right)A \end{aligned}$$

Since F depends on $\widehat{X}$ only through the squared distance terms, we will expand g around the vector of squared distances: $\widehat{V}_i \triangleq \left(\widehat{T}_x^i - \widehat{T}_x^f\right)^2 + \left(\widehat{T}_y^i - \widehat{T}_y^f\right)^2 = \widehat{X}^{\mathrm{T}} D_{if}^{\mathrm{T}} D_{if}\widehat{X}$ (defining the hand position as T^0, and D_{if} as the matrix that extracts the difference between targets i and f as before): $F(V) = \text{diag}\begin{bmatrix} \sigma_H^2 & \sigma_{\dot{H}}^2 & \sigma_f^2 & \gamma V_1 & \cdots & \gamma V_P \end{bmatrix}$. To eliminate the matrix inverse in g we use the following

Lemma 9 *If A is a nonsingular matrix, then*

$$\frac{\partial\, \text{tr}\left(BA^{-1}C\right)}{\partial A} = -\left(A^{-1}CBA^{-1}\right)^{\mathrm{T}}.$$

Applying the lemma to the definition of g, and using the fact that only diagonal elements of F depend on X, we obtain the first order term in the expansion of g w.r.t V and express it as a function of X:

$$\begin{aligned} \left(\widehat{V}-\overline{V}\right)^{\mathrm{T}} \left.\frac{\partial g}{\partial V}\right|_{V=\overline{V}} &= \widehat{X}^{\mathrm{T}}\left(\sum_{i\neq f}\gamma W_i D_{if}^{\mathrm{T}} D_{if}\right)\widehat{X} - \overline{X}^{\mathrm{T}}\left(\sum_{i\neq f}\gamma W_i D_{if}^{\mathrm{T}} D_{if}\right)\overline{X} \\ W\left(\overline{X},\overline{\Sigma}\right) &\triangleq \text{diag}\left(-\left(H\overline{\Sigma}H^{\mathrm{T}} + F\left(\overline{X}\right)\right)^{-1} H\overline{\Sigma}Z\overline{\Sigma}H^{\mathrm{T}}\left(H\overline{\Sigma}H^{\mathrm{T}} + F\left(\overline{X}\right)\right)^{-1}\right)_{4\cdots P+3} \end{aligned}$$

Linearizing g w.r.t Σ can be accomplished using the

Lemma 10 (Matrix inversion lemma) *For a non-singular matrix A and vectors b, c*

$$\left(A + bc^{\mathrm{T}}\right)^{-1} = A^{-1} - \frac{A^{-1}bc^{\mathrm{T}}A^{-1}}{1 + c^{\mathrm{T}}A^{-1}b}$$

We represent $\Sigma = \sum_{ij} s_{ij}\mathbf{e}_i\mathbf{e}_j^{\mathrm{T}}$(where $\mathbf{e}_i$ is the unit vector with 1 in position i), and simplify the expression for $\frac{\partial g}{\partial s_{ij}}\Big|_{\Sigma=\overline{\Sigma}}$ by identifying $A \leftarrow H\overline{\Sigma}H^{\mathrm{T}} + F\left(\overline{X}\right)$, $b \leftarrow \sqrt{s_{ij}}H\mathbf{e}_i$, $c \leftarrow \sqrt{s_{ij}}H\mathbf{e}_j$ and using the above lemma. Skipping some tedious algebra, the first order term in the expansion of g w.r.t Σ is:

$$\mathrm{tr}\left(\left(\Sigma - \overline{\Sigma}\right)\left(T\overline{\Sigma}Z + Z\overline{\Sigma}T - T\overline{\Sigma}Z\overline{\Sigma}T\right)\right), \qquad T\left(\overline{X}, \overline{\Sigma}\right) \triangleq H^{\mathrm{T}}\left(H\overline{\Sigma}H^{\mathrm{T}} + F\left(\overline{X}\right)\right)^{-1} H$$

The linearization is summarized in the following

Theorem 11 (Linearization) *The function $g(X, \Sigma) \triangleq \mathrm{tr}\left(Z\Sigma H^{\mathrm{T}}\left(H\Sigma H^{\mathrm{T}} + F(X)\right)^{-1} H\Sigma\right)$ is approximated up to first order by*

$$\begin{aligned}
g(X, \Sigma) &\approx \overline{g}\left(\overline{X}, \overline{\Sigma}\right) + X^{\mathrm{T}}\mathcal{L}_X\left(\overline{X}, \overline{\Sigma}\right)X + \mathrm{tr}\left(\mathcal{L}_\Sigma\left(\overline{X}, \overline{\Sigma}\right)\Sigma\right) \\
\overline{g}\left(\overline{X}, \overline{\Sigma}\right) &= g\left(\overline{X}, \overline{\Sigma}\right) - \overline{X}^{\mathrm{T}}\mathcal{L}_X\left(\overline{X}, \overline{\Sigma}\right)\overline{X} - \mathrm{tr}\left(\mathcal{L}_\Sigma\left(\overline{X}, \overline{\Sigma}\right)\overline{\Sigma}\right) \\
\mathcal{L}_X\left(\overline{X}, \overline{\Sigma}\right) &= \sum_{i \neq f} \gamma W_i\left(\overline{X}, \overline{\Sigma}\right) D_{if}^{\mathrm{T}} D_{if} \\
\mathcal{L}_\Sigma\left(\overline{X}, \overline{\Sigma}\right) &= T\left(\overline{X}, \overline{\Sigma}\right)\overline{\Sigma}Z + Z\overline{\Sigma}T\left(\overline{X}, \overline{\Sigma}\right) - T\left(\overline{X}, \overline{\Sigma}\right)\overline{\Sigma}Z\overline{\Sigma}T\left(\overline{X}, \overline{\Sigma}\right)
\end{aligned}$$

The form of $\mathcal{L}_X$ allows an intuitive interpretation of the effects of our spatial visual acuity model[12]. The $\gamma W_0\left(\overline{X}, \overline{\Sigma}\right) D_{0f}^{\mathrm{T}} D_{0f}$ term (which is the only one that depends on controllable quantities - i.e. hand position) is equivalent to the definition of the Q_{k_i} matrices penalizing deviations from targets (the required accuracy w_i is replaced by $\gamma W_0\left(\overline{X}, \overline{\Sigma}\right)$). $\mathcal{L}_X$ appears in Theorem 7 as an extra cost being added to S_k. Thus taking into account the state dependence of observation noise is similar to including extra penalty for distance between the hand and current fixation target. This is to be expected since keeping the hand close to the target results in better visual estimates of hand location, which leads to better expected performance.

3.6 Summary and Implementation

This sections summarizes the model and outlines the steps involved in its implementation. We begin with the parameters that have to be specified in order to construct the optimal controller. The first group of parameters are derived directly from the task description (and thus do not represent free parameters):

[12]In the preceding discussion $\mathcal{L}_X$ multiplies $\mathcal{I}$ rather than $\widehat{X}$, so we use $\mathcal{L}_X \leftarrow \begin{bmatrix} \mathcal{L}_X & 0 \\ 0 & 0 \end{bmatrix}$.

Parameter	Description
N	Number of time steps in the simulation
P	Number of targets, excluding starting position
dt	Duration of one time step, in seconds
$Dsac$	Average saccade duration, in time steps (measured directly)
$w_{1 \ldots P}$	Weights for each target (desired spatial accuracy)
v	Weight of terminal velocity (desired timing accuracy)
$\widehat{X}_1$	Initial state estimate (hand position, velocity, target positions)
Σ_1	Variance of the initial state estimate

The next group of parameters describe the internal properties of the sensory-motor system and have to be estimated in order to instantiate the model. Independent estimates for most of them can be found in the literature:

Parameter	Description
$\sigma_H^2, \sigma_{\dot{H}}^2, \sigma_f^2$	Variances of absolute measurements
γ	Variance of relative measurements $= \gamma \left\|\text{target-fixation}\right\|^2$
σ_U^2	Variance of control noise
σ_D^2	Variance of positional additive noise (velocity noise $= 0$)
k, b	Stiffness and damping coefficients
r	Weight of differences between consecutive controls
Δ	Sensory delay, in number of time steps

The controller can be logically separated in two parts, one derived from the stochastic optimal control framework and the other being responsible for setting additional variables by using the first part as an inner minimization loop. These additional variables are:

Variable	Description
$k_1, \ldots, k_{P-1}$	Via-point passage times
$s_1, \ldots, s_{P-1}$	Saccade times (peak eye velocity)
$\left(\overline{X}, \overline{\Sigma}\right)_{1 \ldots N}$	Expected trajectory and covariance for linearization
$\overline{L}_{1 \ldots N}$	Default control law (initially equal to I_1 - min.jerk)

All these quantities can be initialized from the minimum-jerk trajectory (see Chapter 6). The timing variables can be optimized by computing the expected cost for different settings. We assumed above that eye-movement and passage times are discrete (represented in number of simulation time steps). This is a problem if their optimal values are to be computed by some minimization procedure, since most numerical minimization procedures will stop when they encounter a perfectly flat region in parameter space. Interpolation of the cost function between integer parameter values is not a very good idea since function evaluations are quite expensive, and the parameter space is large. Instead we will make the cost function continuous in the timing parameters by assuming that visual input is removed gradually around saccade onset/offset time, and that the target cost is distributed among Q_k and Q_{k+1} when the passage time is between k and $k+1$.

$\overline{X}, \overline{\Sigma}, \overline{L}$ cannot be optimized since they affect the definition of the cost function; instead they are modified as discussed in the previous sections. These additional variables can be computed either once at the beginning of the movement, or fine-tuned at each time step. If they are fixed at the beginning of the movement, the optimal control law can also be computed in advance. Whenever the additional variables change, the fixed quantities $Q, H, \mathcal{L}_X, \mathcal{L}_\Sigma$ used in the optimal controller have to be recomputed.

Finally, the optimal controller computes (according to the above theorems) and propagates through time (k) the following variables:

Variable	**Description**
$\widehat{X}_{k-\Delta\|k-1}$	Estimate of the system state
$\Sigma^{1/2}_{k-\Delta\|k-1}$	Square root covariance of the state estimate
$\widetilde{K}_{k-\Delta}$	Modified kalman gain
S_k	Cost-to-go term multiplying $\mathcal{I}$
P_k	Cost-to-go term multiplying Σ
α_k	Constant cost-to-go term
L_k	Optimal control law
U^*_k	Optimal control signal

Assuming that $\overline{X}, \overline{\Sigma}$ are being modified during the movement, the algorithm involves the following steps at each time k:

1. Start with estimates $\widehat{X}_{k-\Delta|k-1}$ and $\Sigma^{1/2}_{k-\Delta|k-1}$.
2. Propagate S_k and P_k backwards starting from time N, to obtain L_k and $U^*_k = L_k \mathcal{I}_k$. Possibly iterate this step to improve $\overline{X}, \overline{\Sigma}$. Compute cost-to-go α_k if it is needed to optimize timing variables.
3. Using the new measurement Y_k, control U^*_k and $\widetilde{K}_{k-\Delta}$, compute $\widehat{X}_{k+1-\Delta|k}$ and $\Sigma^{1/2}_{k+1-\Delta|k}$ and form $\mathcal{I}_{k+1}$.

Recall that we had to make some approximations in deriving the optimal control law. In particular, we used the optimal linear estimator of X_k given Y_k instead of the optimal estimator (which is nonlinear due to our model of spatial vision). Also, we used a linear approximation to the term $\mathrm{tr}\left(\left(S^X_{k+1} - P_{k+1}\right) \widetilde{K}_{k-\Delta} \widetilde{K}^{\mathrm{T}}_{k-\Delta}\right)$ around an expected trajectory. To assess the accuracy of these approximations, we ran the model on 5 different parameter settings (close to the ones used in the next chapter) to obtain the optimal control law and expected cost, and then simulated 50 trials using the corresponding optimal control law to generate control signals. The optimal control law was computed once before

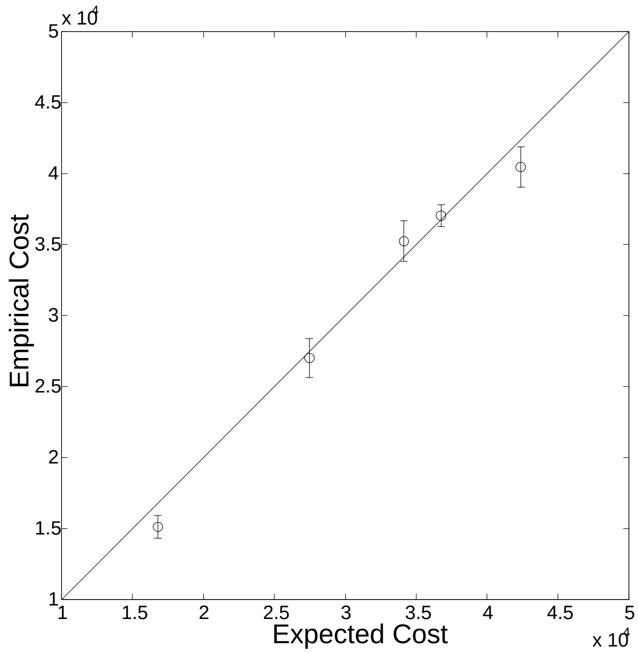

Figure 3-1: Expected vs. empirical cost, computed for 5 simulation runs of 50 trials each (shown are standard error bars).

movement onset and remained constant during the movement. Target passage times and eye movement times were set equal to the target passage times obtained from the minimum jerk model (Chapter 6). The expected trajectory $\overline{X}_k$ was initialized to the minimum jerk trajectory, and the expected covariance $\overline{\Sigma}_k$ was computed from the minimum jerk trajectory by running the Kalman filter.

First we assessed how close the expected cost is to the empirical cost obtained by evaluating the cost function (3.5) directly on the simulated trajectories. The difference was within 2% of the cost value, as shown in Fig 3-1.

We also simulated the model on the same parameter sets as before without the linearization (this can be done by simply setting $\mathcal{L}_X = 0, \mathcal{L}_\Sigma = 0$). The error (difference between expected and empirical cost) increased by about 50% and the empirical cost increased by 10%, demonstrating that the linearization a) results in a better control law and b) the controller has a better estimate of how good that control law is.

The linearization parameters (and thus the control law) typically converged after 1 update. This is shown in Fig 3-2.

We also computed the difference between the actual value of the $\mathrm{tr}\left(\left(S^X_{k+1} - P_{k+1}\right)\widetilde{K}_{k-\Delta}\widetilde{K}^{\mathrm{T}}_{k-\Delta}\right)$ and the one computed from the linearization, at every time step. It is not surprising that the error vanishes after $\overline{X}_k$ has converged to the expected trajectory obtained from the corresponding optimal control law (i.e. we are computing the cost at the linearization point). Note however that the error is rather small even before the first iteration, when $\overline{X}_k$ is equal to the minimum jerk trajectory. Typical results are shown in Fig 3-3.

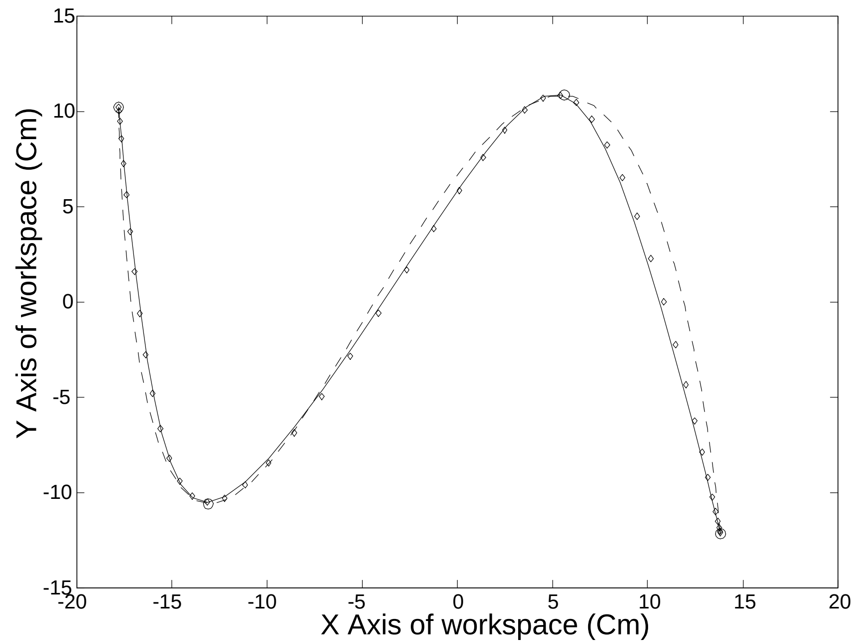

Figure 3-2: Minimum jerk trajectory (- -), expected trajectory defore (diamond) and after (solid) first iteration. Circles correspond to target locations.

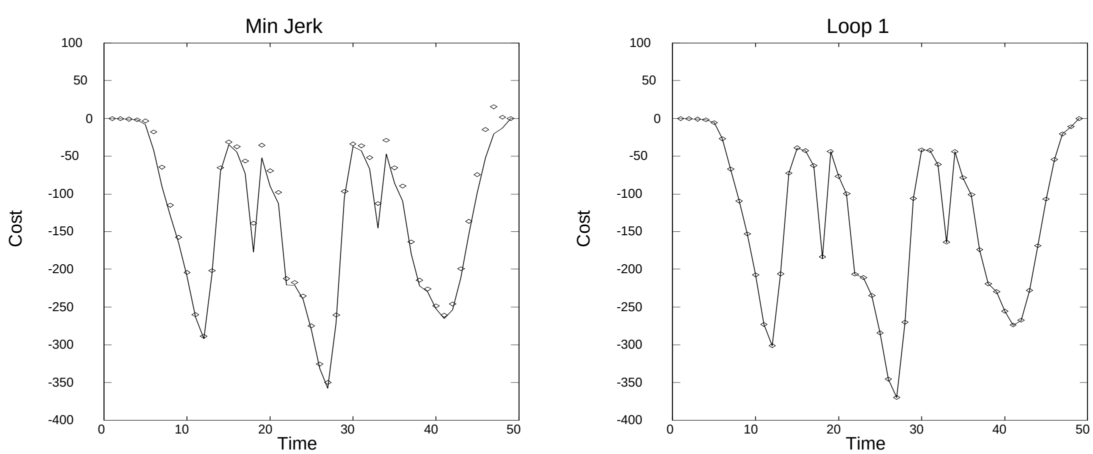

Figure 3-3: Linearized (diamond) vs actual (solid) value of the tr(.) term at every time step, before and after the first iteration.

Chapter 4

Evidence for online goal-directed control in via-point tasks

4.1 Overview

In this chapter we describe a number of experimental results indicating that even in simple tasks such as reaching or moving through a small set of sequential targets, the sensory-motor system constantly updates its internal estimate of the environment (including the moving hand) and modifies the ongoing movement accordingly. Presently we cannot rule out the possibility that the phenomena described below result from the execution of preexisting motor programs. This we think is mostly due to the lack of an experimentally falsifiable definition of "motor program". In the absence of such a definition, all we can show is that a number of interesting effects are better explained in the framework developed in the previous chapters, and using a prerecorded motor template in addition to the processing postulated here would not account for any additional findings.

As will become clear shortly, most of the experimental results discussed in this chapter are replicated by our model for rather obvious reasons. Thus we do not "explain" anything with emergent properties of the model that cannot be deduced easily from Chapter 3. While this type of explanation is sometimes preferred in modelling work (since it clearly demonstrates that the desired outcomes were not built into the model) it is not our goal here. On the contrary, the model was directly inspired by some of the data presented below.

The experiments were performed using a projection system that enabled us to display virtual targets on the surface of a horizontal table on which subjects were asked to move their hand. In all experiments the room was dark, i.e. the only visual feedback subjects received about their hand position was the one we chose to display. Hand position was displayed by using an LED attached to a small pointer held by the subject (in experiments that did not require visual perturbation of hand location), or by a computer-generated cursor calibrated to match the physical location of the arm (in experiments where hand location was displaced unknown to the subject). Hand position was recorded with an Optotrack 3020 running at 120Hz. In some experiments eye movements were also recorded using an Iscan infrared tracker, running at 120Hz. We estimated the delay of the eyetracker relative to the Optotrack (by attaching a positional sensor on the eye camera, tapping on the headmount while a subject was fixating, and comparing the readings of the two sensors) and compensated for it - it was one video frame.

4.2 Reaching movements

4.2.1 Experiment 4.1 - Spatial accuracy of peripheral vision

The goal of this experiment was to estimate the spatial accuracy of peripheral vision, i.e. the accuracy with which subjects can localize targets presented transiently in the peripheral visual field. This is clearly important for validation of the model which makes specific assumptions about the decay of visual accuracy with distance from the fixation point. Note that our estimates are a function of distance along the horizontal table rather than visual angle.

Methods

Six subjects participated in the experiment, which lasted about 15 minutes. Subjects made 300 pointing movements according to the following protocol: while the hand was stationary and the subject was fixating at the LED attached to the hand, a computer generated target appeared for 700 ms and then disappeared. The subject was instructed to maintain fixation at the LED while the target was present, and to make a pointing movement to its (remembered) location when it disappeared. Subjects were not required to respond as fast as possible, and the duration of the pointing movement was not limited - i.e. we wanted to allow the best possible localization of the remembered target location. Once the pointing movement terminated (detected automatically using a small velocity threshold), the next target appeared with a delay of 700 ms. Target locations were selected randomly, such that the distance from the current hand position was uniformly distributed between 30 and 280 mm, in random directions (with the restriction that targets had to remain within the limits of the workspace).

Results

Endpoint errors from a typical subject are shown in Fig 4-1 - note that the error along the movement direction is suppressed on this graph. It is clear that the spread of the errors increases linearly with distance. This is confirmed in the summary plot in Fig 4-2. Overall endpoint variable error (standard deviation) increases remarkably linearly with distance. There was also a small bias - movements tended to undershoot by 6 mm on the average (over all subjects) - corresponding to a range effect. That bias has of course been subtracted in computing the variance. The variable errors along the direction of movement were slightly bigger than the ones in the orthogonal direction (13mm vs. 10.5mm) suggesting that some of the localization error may be due to the fact that the hand is used as a pointer, and hand reaching movements are known to have larger endpoint errors along the direction of movement [Gordon et al. 94].

We used this data to fit the parameters of the model corresponding to visual perception, and fixed these values throughout all simulations. The model also contains sensory noise terms (position and velocity noise) corresponding to proprioception. Increasing those terms has very similar effects to decreasing the amount of system noise. Since it is difficult to design experiments that isolate the accuracy of proprioception (that would involve temporarily deafferenting the subjects), we fixed those parameters so that proprioception was significantly less reliable than vision (in the vicinity of the target), yet resulted in improved performance.

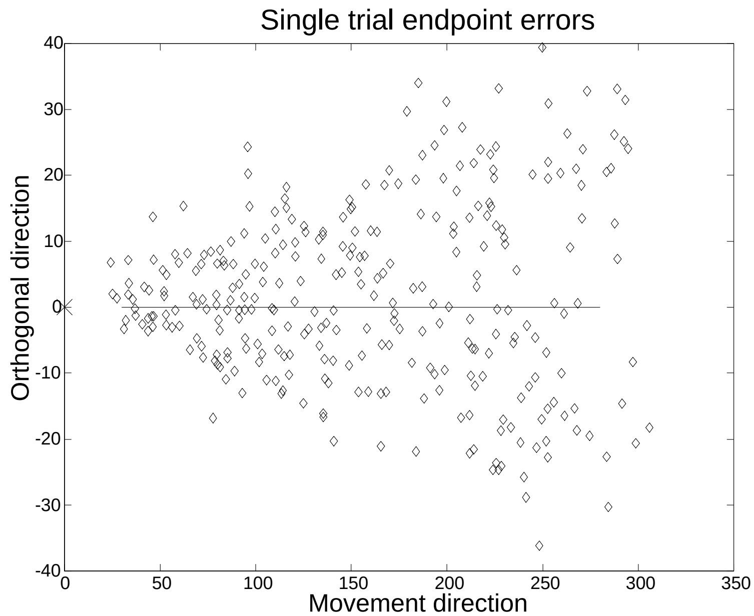

Figure 4-1: Endpoint errors of a typical subject. Targets were sampled from the solid line, the starting position is marked with the cross, all movements have been aligned with the horizontal direction. Note different scale on X and Y axes.

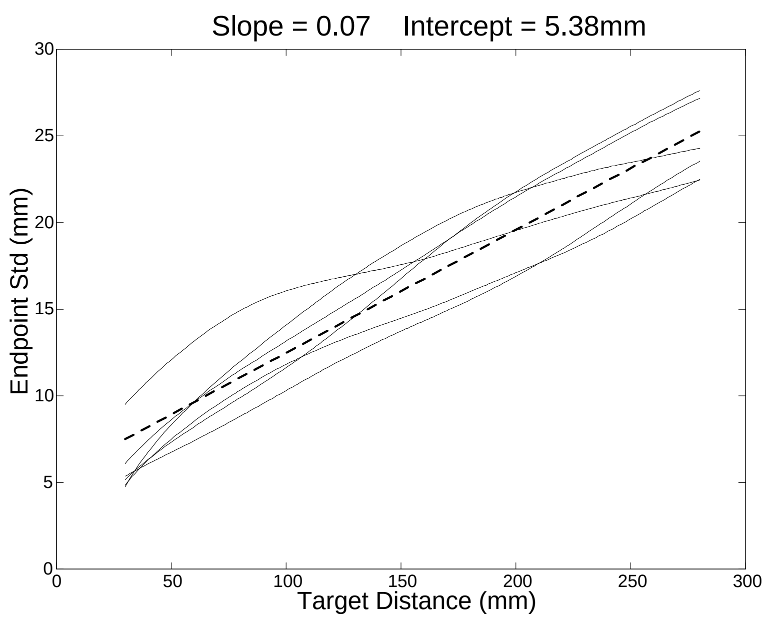

Figure 4-2: Standard deviation of distances away from the mean, computed in overlapping bins of 50 ms and smoothed using a cubic spline. Shown is data for individual subjects, and the linear regression line for the combined data for all subjects.

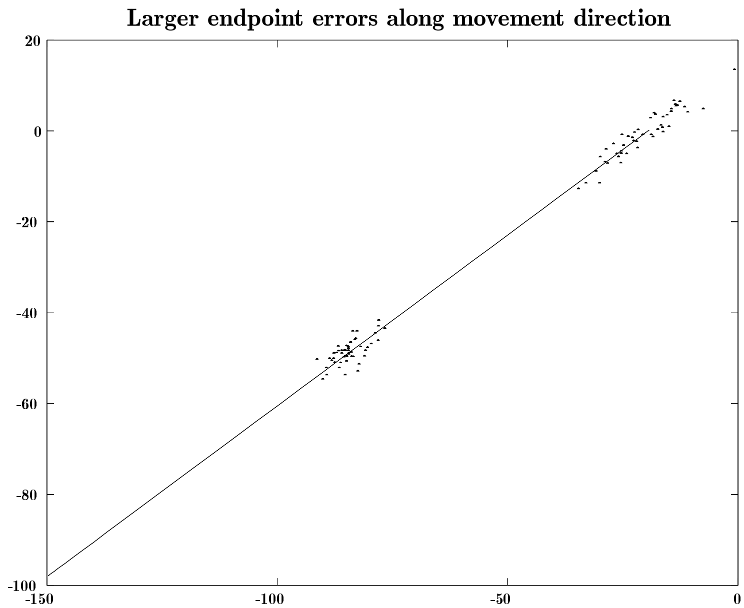

Figure 4-3: Simulation: reaching movements of equal duration to targets at different distances without visual feedback of hand location. Endpoint error increases with distance.

A number of phenomena called speed/accuracy trade-offs have been observed (see [Plamondon and Alimi97] for a recent review) that indicate a strong relationship between movement speed (or duration) and the accuracy that can be achieved. Such trade-offs are observed by either specifying desired accuracy and asking subjects to move as quickly as possible (as in Fitts' law) or specifying movement duration and asking subjects to be as accurate as possible (as in Schmidt's law) - [Jeannerod 88, Schmidt et al. 79]. While a number of explanations for these phenomena have been considered [Plamondon and Alimi97] the most parsimonious explanation so far seems to be the stochastic optimal control model of [Meyer et al. 82], and more recently of [Hoff 94, Hoff 91]. There are three main differences between these prior models and the one developed here:

- we consider two-dimensional as opposed to one-dimensional movements (which leads to some new results presented in the next section);
- we have combined a model of sensory estimation with the optimal controller, while prior models assumed perfect state observations;
- the system noise in our model is proportional to acceleration/force rather than velocity - which seems physiologically more realistic, and is also supported by psychophysical data [Schmidt et al. 79].

Below we show examples of the two classes of speed/accuracy trade-offs. Fig **??** shows how endpoint errors increase with distance (keeping movement duration constant), consistent with the observations of [Gordon et al. 94] and related to Schmidt's law. The reason for this effect is the system noise proportional to acceleration - to produce a two times larger movement with the same duration, the model has to use two times higher acceleration which results in more noise.

The next figure 4-4 demonstrates the basic phenomenon underlying Fitts' law: when desired accuracy is increased, the decceleration phase of the movement increases (which accounts for increased movement

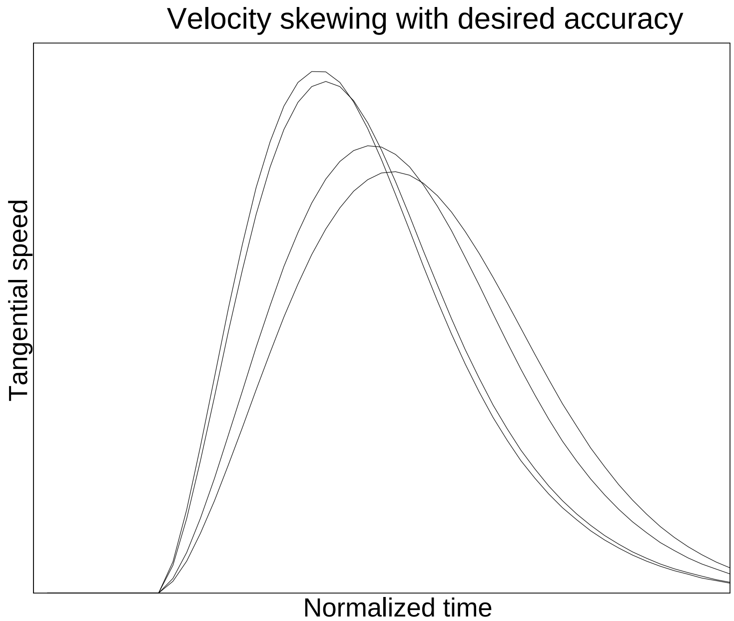

Figure 4-4: Simulation: when desired accuracy is increased (by increasing the weight of the spatial error), the speed profiles become more skewed.

duration when duration is not controlled) - [MacKenzie et al. 87, Milner and Ijaz 90]. This occurs in the model for two reasons: slowing down at the end of the movement results in less system noise (as the previous optimal control models have shown), but also allows the hand to remain in the region of highest visual sensitivity for a longer period of time.

Fig 4-5 illustrates the optimal estimation procedure which is a part of the model. The estimation procedure starts with a prior probability (determined by the range of possible target locations - in this case only two); this prior remains unaffected for Δ time steps (which is the delay of the sensory pipeline), and then begins to shift towards the correct target. Note that such a sensory model accounts quite naturally for the observations of [Hening et al. 88, Ghez et al. 97] who found that responses forced to occur before the RT have a distribution that is equal to the prior for less than 80 ms, and gradually shifts towards the correct target and becomes more peaked[1].

4.3 Multiple via-point tasks

4.3.1 Experiment 4.2 - Eye-hand coordination in passing through multiple targets

The objective of this experiment was to describe the patterns of eye movements that accompany hand movements through multiple targets, and the way this pattern is affected by task parameters. We were

[1] We cannot rule out the possibility - suggested by the authors - that the effect is due to a gradual specification of the motor response, rather than a gradual improvement of the sensory estimate.

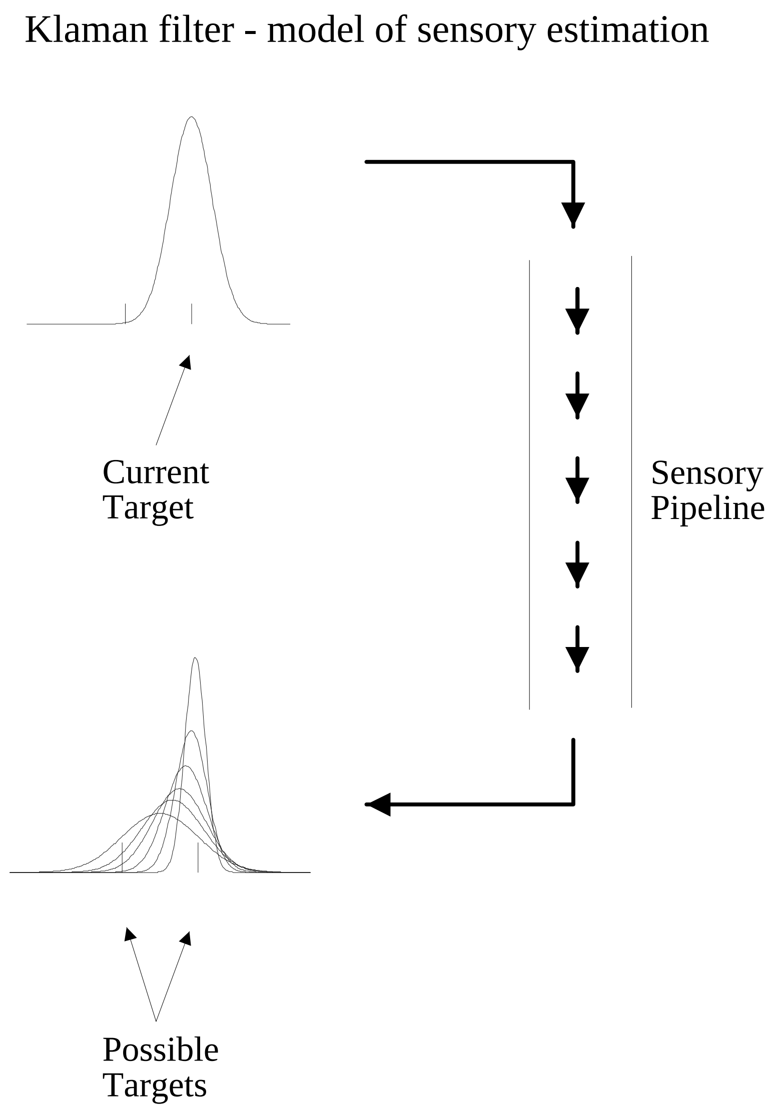

Figure 4-5: Simulation: A schematic description of the sensory processing model. Inputs are generated continuously with a probability distribution determined by the target (Top). The inputs reach the estimator with a fixed delay. With time the optimal estimate shifts towards the correct value, and its variance decreases (Bottom).

also interested in the effects of different feedback conditions and accuracy requirements on the kinematics of the hand trajectory.

Methods

Seven subjects participated in this experiment, which lasted about 20 min. They were asked to move a pointer through sequences of consecutively numbered targets, with a time limit. Each target had a radius, and subjects were informed that the computer would signal an error if they failed to pass inside that radius. At the beginning of each trial a starting box appeared and the subject had to move inside it and stop. At that point the starting box disappeared and the sequence of targets appeared. The subject was free to move when ready. However, once the movement started (detected automatically and signalled with a beep) there was a time limit within which the movement had to be completed. This was enforced by sounding a second beep 300 ms before the end of the allowed period, and signalling an error if the subject was still moving. Errors were signalled by crossing out the missed targets, and showing an arrow in the last target if movement speed was above the acceptable threshold at the end of the trial period. The error display remained on the screen for 1 sec. Target configurations were presented in blocks of 15, with a warning message printed at the beginning of each block. Within a block the only change from trial to trial was a small random displacement of the target set, uniformly sampled along the X and Y axis between -35 and 35 mm (this was done to reduce overtraining). Also, at the beginning of each trial the targets appeared in exactly the same position relative to the hand location, i.e. even if the subject made a repositioning error it was eliminated by the program.

The different conditions included 2 different target configurations, presented with and without visual feedback about hand location (in the no feedback conditions the LED was turned off when the targets appeared, and turned on again at the end of the movement). We also varied the target sizes. The default radius was 8 mm, which in some conditions (and some targets) was reduced to 4 mm. The allowed duration of the movements was 1200 ms. Subjects generally complied with this time limit.

Results

The pattern of eye movements used by all subjects on the large majority of trials (above 90%) was the following: the eyes fixated on a target that the hand had not yet reached, and when the hand was in the vicinity of the target, a saccade was made to the following target. Before the beginning of the hand movement most subjects made saccades between the targets. Two of the subjects actually tended to initiate the hand movement while still scanning, but the eyes returned to target 1 before the hand reached it, and from that point on the eye movement pattern was the same as described above. Thus the important variable to analyze is the timing of the eye movements, relative to the time when the hand passes the current fixation target (defined as the point of nearest approach) or the time elapsed from the beginning of the movement. In this analysis the eye movement time is defined as the time of peak eye velocity (since it can be estimated more reliably, and is also the parameter used in the model). The onset of the eye movement is on the average 25 ms earlier.

We observed that in the presence of visual feedback about hand position, fixation was maintained on the average 100 ms after the hand passed the current fixation target. A histogram of this eye-hand delay for 2000 saccades from all subjects and target configurations is shown in Fig 4-6 - Top. This delay would be unnecessary if the role of vision in this task was to correct deviations between the hand at the target that is currently being approached.

This delay was not fixed, but instead was easily modifiable by altering the task - see Fig 4-6 Bottom. The largest effect was observed when we removed feedback about hand position during the movement

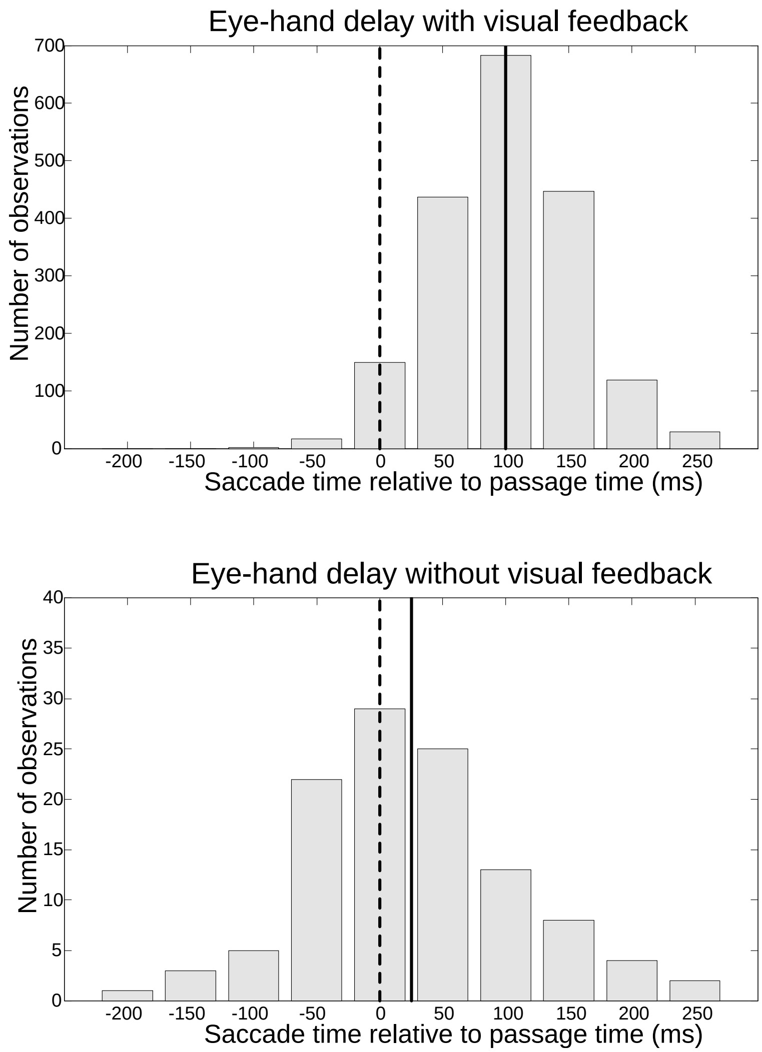

Figure 4-6: Histograms of time of peak eye velocity relative to target passage time. Dashed line marks 0, solid line is the median value. Saccade onset occurs around 25ms earlier. Top: saccades with visual feeback of hand location. Bottom: visual feedback removed.

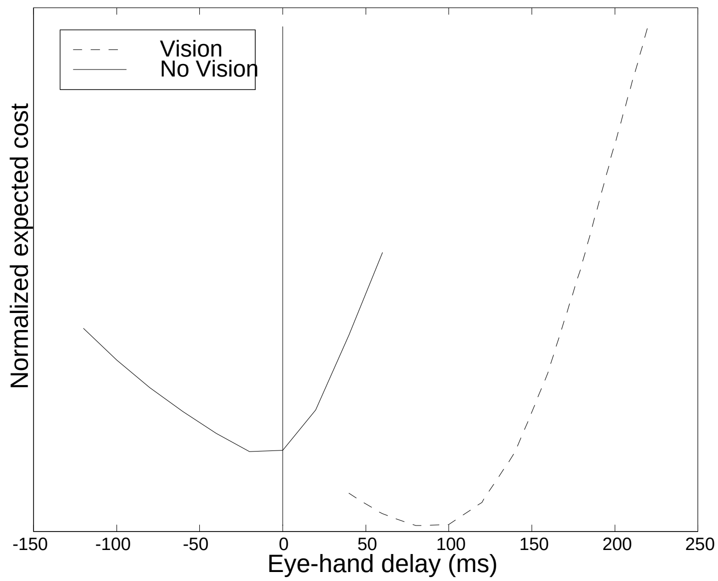

Figure 4-7: Simulation: For a movement with a single intermediate target and fixed passage time, we varied the eye movement time and computed the expected cost for the resulting optimal control law. When vision is removed (by greatly increasing visual noise in the model), the otimal eye movement time moves closer to the passage time.

(the targets were always displayed). In this case the eye-hand delay was reduced by about 70 ms, i.e. if we computed the saccade onset time instead of the time of peak eye velocity, that delay would be indistinguishable from zero. We observed other systematic differences in eye-hand delay - between different targets along a movement, or at the same target as some of the remaining targets changed. These differences were smaller (on the order of 30 ms) but significant. It is presently not clear how to describe such effects (other than presenting the specific target configurations where they occur) so we will not discuss them any further. However, their existence suggests that eye-hand delay is an important variable that is tightly controlled, and studying it in more detail may lead to other interesting findings.

The same effect (eliminating the eye-hand delay in the absence of visual feedback) was observed in the simulations - Fig 4-7. In the model it is advantageous to maintain fixation at the current target after the hand passes through it, because this is the only time when the hand is close to the center of fixation and thus its position/velocity can be estimated more accurately. Removing visual feedback of hand location eliminates that advantage, and the optimal eye movement time moves closer to the passage time (in fact, in this model there is no reason to make the eye movement after the passage time without visual feedback). Also, the minimum in the cost function in the presence of visual feedback is much more pronounced (not visible on the graph since we have shifted and rescaled the curves to facilitate comparison). This may explain why the eye movement histogram with visual feedback (Fig 4-6) shows significantly less spread.

We also asked whether it is more appropriate to analyze eye movement times relative to the hand passage time, or the beginning of the movement (which would be preferred if we believed that eye movement times are programmed before hand movement onset). To this end we computed the within subject variability of absolute saccade times, absolute passage times, and saccade times relative to the

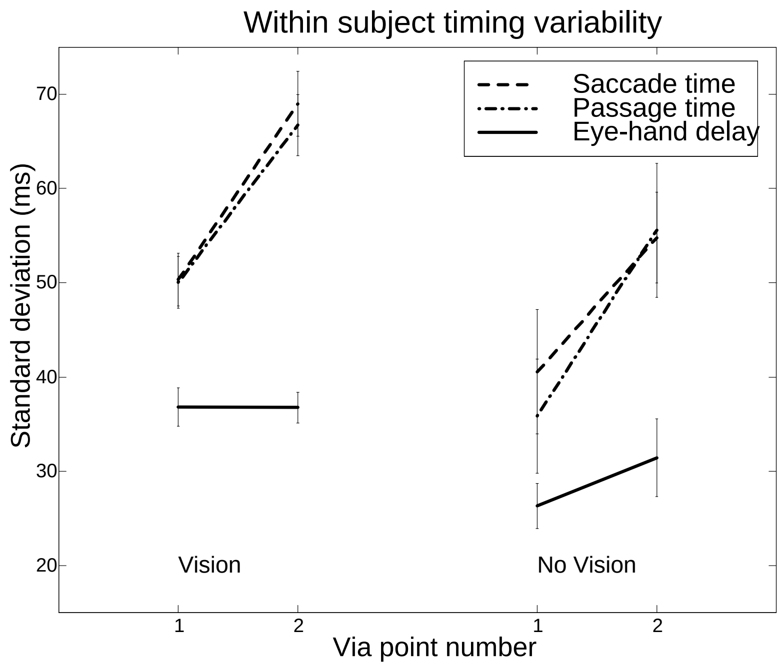

Figure 4-8: Variability of absolute saccade and passage times, and relative saccade times for different via points and feedback conditions.

corresponding passage times. This computation was performed in blocks of 15 identical trials for single subjects, and the results were averaged over subjects. As Fig 4-8 shows relative eye-hand delay was significantly less variable than saccade or passage times measured from movement onset. Also, it was about the same for the first and second intermediate points, while the variability of absolute times increased substantially for the second intermediate point. This excludes the possibility that passage times and saccade times are preprogrammed before movement onset - if that were the case and we subtracted the two, we would expect the variability of their difference (i.e. eye-hand delay) to be bigger than the variability in any of the absolute times. Thus the eye and hand movements are synchronized online. However, this is not done (entirely) through visual feedback - note that in the absence of visual feedback about hand location the result is the same, and in fact eye-hand variability even decreases at the first intermediate point. Therefore the eye movement is synchronized online with the estimated location of the hand (that estimate can be obtained through proprioception, or integration of motor commands, or most likely both).

Next we turn to the analysis of hand kinematics. Fig 4-9 shows the hand paths for all subjects in the baseline conditions (all target sizes equal to 8 mm radius). Most of these median trajectories overlap and turn at the via points, but there are a couple of exceptions where the paths turn either before or after the via point. The bottom part of the figure shows the effects of increased accuracy - the segments of the movement become straighter, and the curvature at the via point increases. This effect of increased accuracy has not been documented before since previous studies have focused on single reaching movements [MacKenzie et al. 87, Milner and Ijaz 90].

Changes in accuracy also have a significant effect on the speed profiles. Fig 4-10 shows that there is a significant decrease in (normalized) speed when subjects are moving towards a point with a smaller

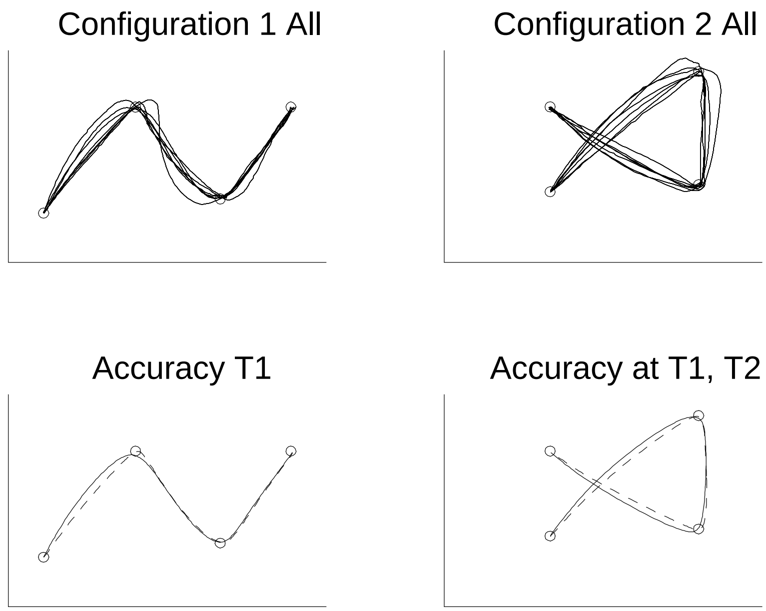

Figure 4-9: Median hand paths for all subjects (top) and average (over all subjects) paths for baseline and increased accuracy conditions.

radius. This corresponds closely to the finding from the literature on speed-accuracy trade-offs that speed profiles become more skewed when accuracy increases [Meyer et al. 82, Milner and Ijaz 90]. In our case movement duration is specified, so we observe a redistribution of speed.

Similar effects are observed in the model when we vary the relative weights of the jerk cost and the spatial accuracy cost. Increasing required accuracy (equivalent to decreasing jerk cost) results in less smooth trajectories, essentially connecting the targets with straight lines - Fig 4-11. It is interesting that the main source of between subject variability in the experimental data is also the amount of curvature at the targets. Thus it is possible that different subjects interpret differently the accuracy requirements of the task (the instructions were to pass through the centers of the targets if possible, and avoid missing the targets altogether).

The remaining part of this analysis focuses on the within subject trial-to-trial hand trajectory variability over a block of 15 identical trials, where movement time has been normalized. Fig 4-13 conveys several interesting findings. The positional variance (Top) is elongated along the movement direction between targets, but shrinks to almost a circle when passing through the target. The area of the position variance ellipse (Middle) correlates with movement speed. Also, the variability of instantaneous speed correlates with the absolute value of acceleration (Bottom).

This variability pattern is not fixed, but changes with the accuracy constraints and visual feedback conditions - Fig 4-12. The effect of removing visual feedback is to increase (by a factor of 2) the baseline variability with visual feedback. More interesting is the effect of increased accuracy requirements. What we see is basically that subjects are capable of trading-off variability: when the variability at the specified via point decreases, the variability at the other via point increases. Note that the trade-off most likely results from the limit on movement duration: if subjects were allowed to increase movement duration for target configurations with increased accuracy requirements, presumably they would (as in Fitts' law).

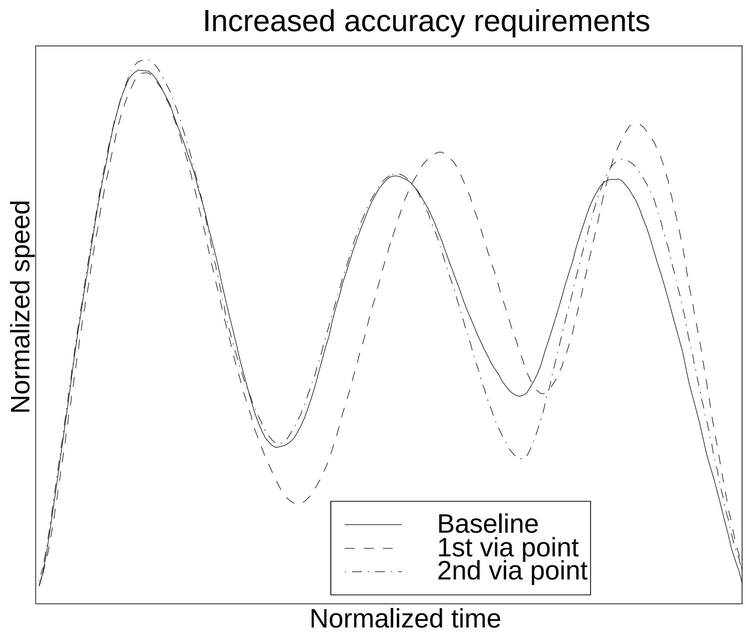

Figure 4-10: Speed profiles in the baseline condition and the two increased accuracy conditions (first vs. second via point).

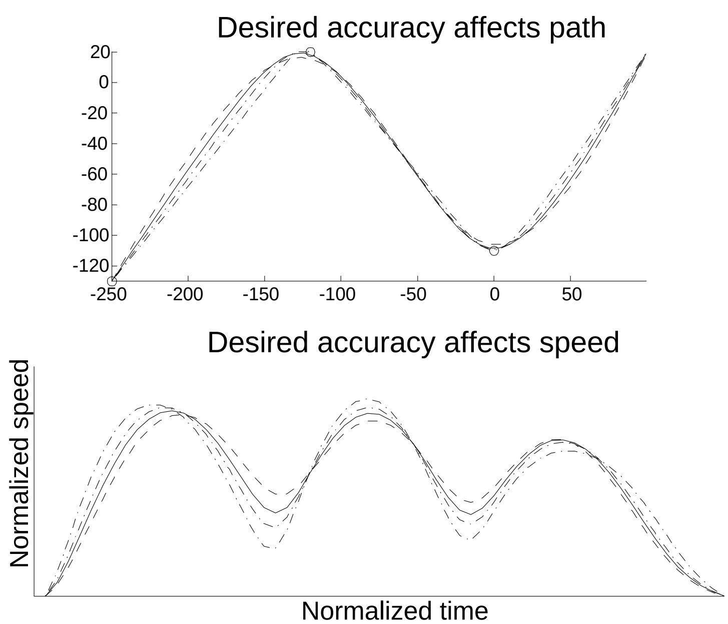

Figure 4-11: Simulation: Increasing the relative weight of the jerk cost results in smoother trajectories. If jerk is eliminated completely from the model, the trajectories become straight and velocity approaches zero at the intermediate targets.

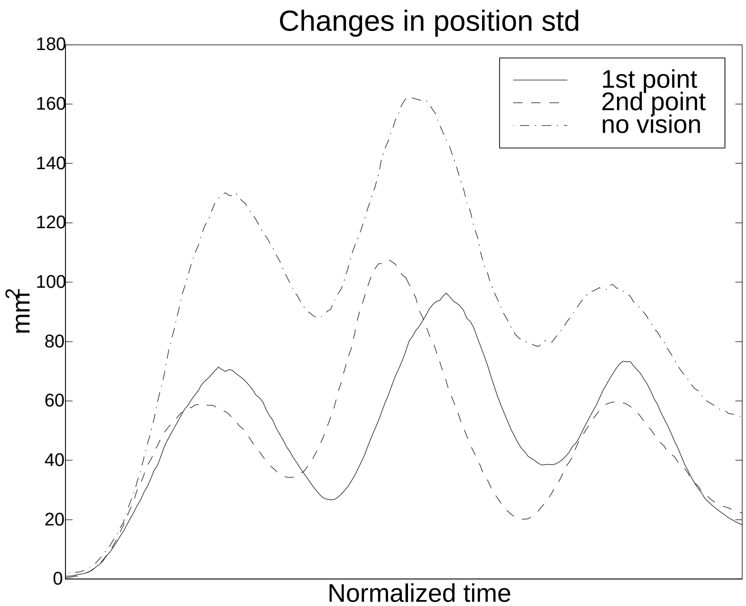

Figure 4-12: Positional variance increases in the absense of visual feedback, and decreases with desired accuracy.

4.4 Movement segmentation

Here we examine in more detail the optimal control law computed by the model, and how it can give rise to observations that seem to suggest movement segmentation.

In Fig 4-14 we have plotted some of the coefficients of the matrices L defining the optimal control law over time (we have chosen a linear arrangement of targets to simplify the plot). The optimal control law is smoothly interpolating between the minimum jerk principle (bottom plot - if the graph reaches 1 that corresponds to setting jerk to zero) and proportional control (top plot - the symmetry corresponds to pushing the hand towards the estimated target position). Note that immediately before the hand reaches a target the accuracy constraint takes precedence and the model uses predominantly positional control, while at the beginning of each "segment" when corrections are not that important it is advantageous to operate closer to the minimum jerk principle. Although it appears that the control law is a concatenation of very similar pieces (corresponding to each pair of consecutive targets), recall that it is being computed as a solution to the optimal control problem (3.7) that makes no explicit reference to movement segmentation. The apparent segmentation arises only because of the particular form of the cost function (the Q matrices encode the existence of multiple targets). Exactly the same computational mechanism however would produce control laws without any hint of segmentation if the cost function is different (e.g. tracing a continuous curve, or reaching for a single target).

It has been known for a while [Woodworth 1899] that reaching to small targets results in a prolonged decceleration phase of the movement. More careful kinematic analysis has revealed that on some (but not all) trials the speed profiles exhibit multiple peaks towards the end of the reaching movement [Milner and Ijaz 90], which has been interpreted as evidence that endpoint visually-guided corrections consist of discrete stereotypical movements (with a straight path and a symmetric bell-shaped speed profile) superimposed on each other. This interpretation leaves some open questions: a) multiple peaks

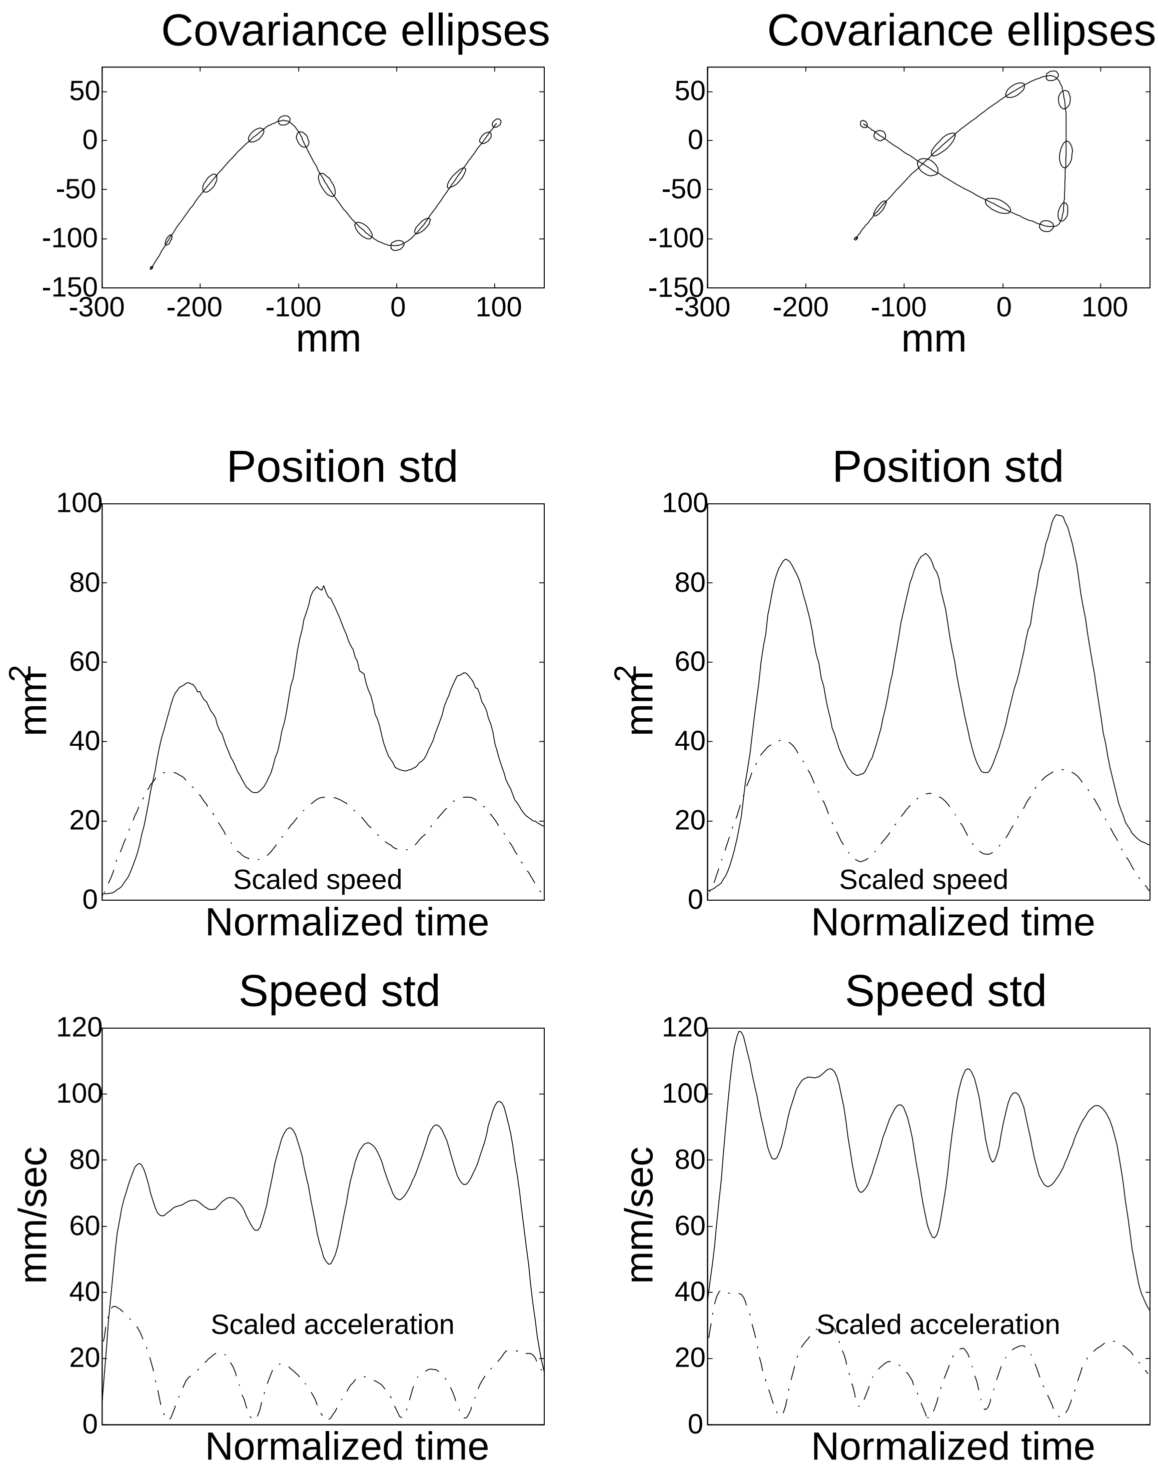

Figure 4-13: Positional variance per spatial location (Top), Positional variance (area of the corresponding ellipse) over time - compared to the (arbitrarily scaled) speed profile (Middle), Speed variance compared to scaled absolute acceleration (Bottom).

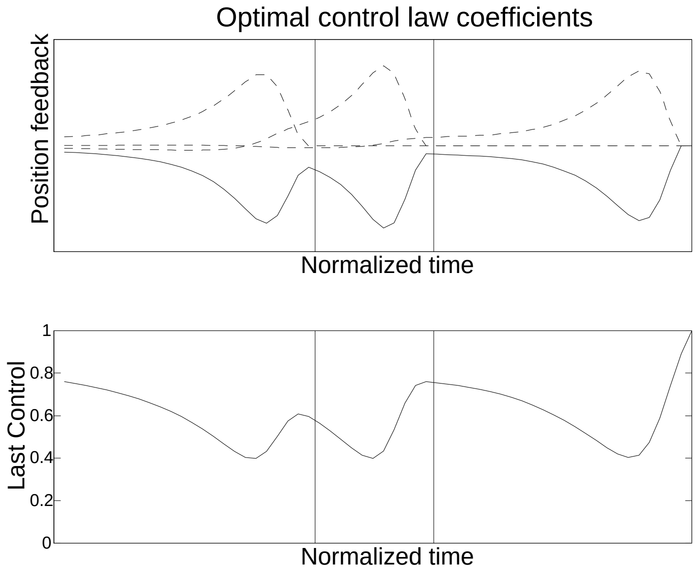

Figure 4-14: Simulation: Coefficients of the optimal control law over time. Vertical lines correspond to target passage times. Dashed lines in the top plot represent the coefficients multiplying target location, solid line - estimated hand location. The temporal evolution of the damping terms (not shown) is similar to that of the positional control terms.

in the speed profile also occur in the absence of visual feedback [Jeannerod 88], b) on many trials the decceleration phase has a rather smooth and continuous speed profile and although it can be fit with a set of basis functions, it is a matter of speculation whether that fit reflects the underlying structure of the movement. Still, the existence of a variable number of secondary peaks in the speed profile is a serious problem for existing optimal control models of Fitts' law, which either predict a fixed number of peaks [Meyer et al. 82] or a continuous skewing of the speed profile [Hoff 91].

Our model is most similar to the continuous-time model of [Hoff 91], however it avoids the above problem because here movements are two-dimensional. In Fig 4-15 we have plotted the expected speed profile corresponding to the optimal control law for a reaching movement with high accuracy requirements (relative to the jerk cost). The expected speed profile is skewed and the expected path perfectly straight (not shown), as in the model of [Hoff 91]. However, the simulated trajectories are not perfectly straight (due to the noise terms in the model), instead they are more curved towards the end of the movement and frequently exhibit multiple peaks in the speed profile. In the model this phenomenon does not arise from concatenation of discrete predefined movement templates, but is a consequences of the high positional error gain at the end of the movement. Similar behavior has recently been observed in a two-dimensional PD control model with more realistic arm dynamics [Bhushan and Shadmehr 98].

Movement segmentation may also be inferred from the following analysis: compute the pairwise correlations between features of the trajectory at different points in time, and look for time intervals within which the points along the trajectory are strongly correlated with each other and the correlation drops significantly at the boundary of the temporal interval. We have found such evidence in a number of experiments. We computed the pairwise correlations between hand positions in different points in time (average of r_x^2 and r_y^2, thresholded at 0.5), after the duration of the trajectories was normalized. The block-diagonal structure of the correlation matrix in Fig 4-16 can be seen as evidence for segmentation:

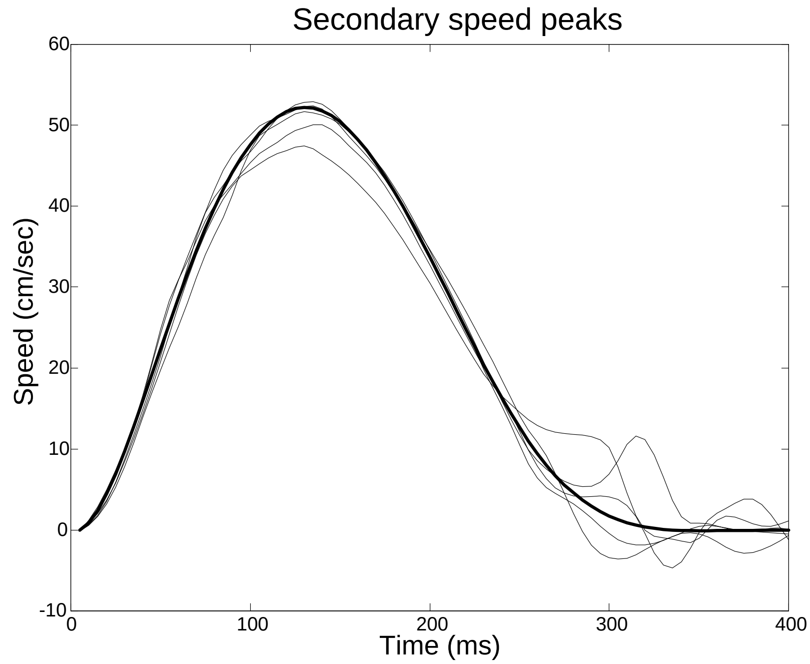

Figure 4-15: Simultion (reacing movement to a single target): Thick line - speed profile of expected trajectory. Thin lines - single trial speed profiles, affected by noise.

points in the same segment are correlated with each other, but not with points in other segments. It is interesting that we have obtained a very similar correlation matrix by analyzing the 3D hand paths of subjects hitting a table tennis ball (the experiment is described Chapter 5). The resulting matrix - Fig 4-17 - has a very similar structure, although in this experiment there were no explicit spatial targets. The apparent segmentation points correspond to a) the transition between the preparatory phase and the fast forward swing of the paddle, and b) the point of impact with the ball. Thus the second segmentation point is in the middle of a relatively straight movements with a bell-shaped speed profile.

While these findings (especially in the case of table tennis) can be interpreted as rather strong evidence for movement segmentation, the same phenomenon occurs in the model which does not explicitly segment the movement - Fig 4-18. In general, feedback regulation has a decorrelating effect: if the controller was attempting to maintain a certain posture, or otherwise had a "goal" defined at each point in time we would expect to see an almost diagonal correlation matrix. In our case however positional constraints are only defined at the passage times. While we can compute an expected trajectory from the control law, it is not a "desired" trajectory, i.e. the controller does not prefer that particular trajectory over other trajectories that have similar costs. As a result, it only makes corrections when required by the goal constraints (i.e. near the targets - as seen in Fig 4-14) so noise effects are allowed to accumulate inside each segment, resulting in block-diagonal correlation matrices. Again, this is an example where a seemingly segmented movement is generated by a computational mechanism that has no explicit segmentation built into it, but is trying to accomplish a task which dictates a segmented form of control.

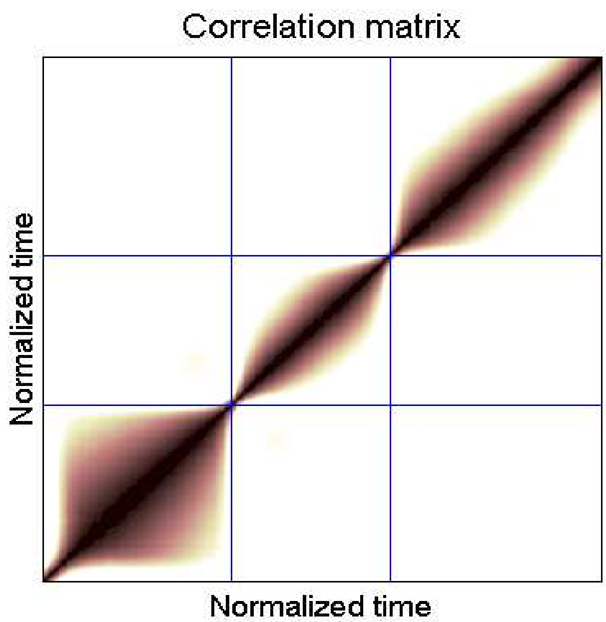

Figure 4-16: Correlation matrix, showing how predictable the hand location at one point in time is given the location at another point in time. The solid lines mark the pasage times through the targets.

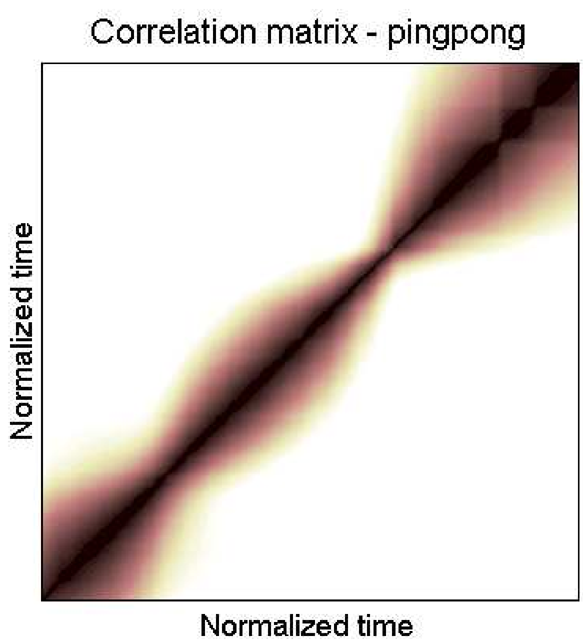

Figure 4-17: Correlation matrix of hand position in a table tennis stroke.

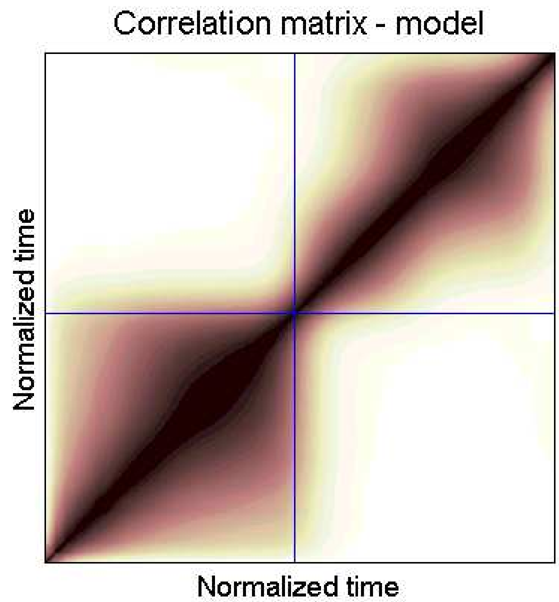

Figure 4-18: Simulation: correlation matrix of hand position in the model for a single via-point (passage time maker with solid lines).

4.5 Visual perturbations

The experiments in this section explore the effects of visual perturbations during the movement. In the first two experiments we displaced one of the targets, while in the third one we introduced a prism shift (i.e. displaced the hand location) in a limited region of the workspace. Both perturbations result in rapid, appropriate compensations, even though subjects are unaware of the existence of the perturbations. Of particular interest is the finding that subjects corrected for future "errors" that could only be inferred from the visual display, suggesting that the motor system is actively monitoring the state of the environment and making appropriate adjustments online, as opposed to following a predetermined movement plan.

4.5.1 Experiment 4.3 - Saccade-triggered target displacements

Methods

Six subjects were asked to move through a sequence of 4 targets (plus a starting position, numbered 1) - see Fig 4-19. Eye movements were monitored, and targets were displaced around the time of peak eye velocity to avoid the possibility of subjects being aware of the perturbations. Recall that eye movements occur roughly when the hand passes the fixation target. Each displacement was 2 cm either up or down (chosen randomly). Trials from the three conditions shown in the figure were mixed randomly; on 50% of the trials no perturbation was introduced. Subjects were asked to pass through the targets in sequence and stop at the final target. Hand location was always displayed through a computer-generated cursor. In this and all remaining experiments we did not display spatial error information; movement duration was enforced by printing a warning message whenever the trial duration exceeded the allocated time (1400 ms). As before, subjects were free to start the movement when ready, and movement onset was detected automatically.

Results

We show the effects of target displacements on the average movement path in each condition, for displacement at an intermediate target - Fig 4-20 and the final target - Fig 4-21. The most important result is the significant difference between up- and down- perturbed trials in condition 1 (Top). In this case target 4 (in Fig 4-20) remained displaced while the hand was traversing the segment between targets 2 and 3, and the eyes were fixated on target 3. If it were the case that a movement plan was available and the system was only monitoring deviations away from that plan, there is no reason why transient undetected displacements of targets that the hand has not yet reached should affect the movement, and yet they do. The effect is clearly consistent with a model that continuously samples the visual input, forms new estimates of target locations, and adjusts the hand movement accordingly. The results from the other two conditions correspond to what one might expect, i.e. if the perturbation is introduced earlier the correction is larger (Middle), but there is also a substantial correction when the perturbation is introduced after the hand begins the movement to that target (Bottom). This makes the result from condition 1 even more interesting, because clearly the visuo-motor system has time to make adjustments during segment 3-4, yet the effect of the target perturbation while tracing segment 2-3 must have been big enough so that it cannot be canceled during segment 3-4 (the system should attempt to cancel it since in condition 1 the target is back at its original location during tracing of segment 3-4).

4.5.2 Experiment 4.4 - Randomized target displacements

The drawback of the experimental paradigm above is that saccades occur at relatively constant times, thus synchronizing target displacements with saccade times does not allow us to explore the effects of introducing perturbations at different times. To obtain that measurement, we informed subjects that perturbations will occur and asked them to correct as much as possible while still completing the movement in time.

Methods

Five subjects were asked to move their hand through consecutively numbered targets. Here the configurations were chosen so that the staring and ending target overlapped - to avoid the time necessary for repositioning and thus speed up the experiment. There were 4 intermediate targets, plus the start/end target. The last three intermediate targets were displaced (randomly) to one of two possible locations, roughly 2 cm away from the average movement paths in a direction orthogonal to the path. Displacements occurred at random times before the target was reached, and the target returned to its original location at another random time before it was reached (the program collected a baseline for each subject to estimate when (s)he was likely to pass through the targets, and used that to distribute the perturbation times uniformly in a -500:-100 ms window before passage time). The perturbation was present for at least 50 ms. Subjects were instructed that displacements will occur, and asked to pass through the most recent target location while still completing the movement in the allocated time (enforced as before). On each trial a target and a direction of perturbation were chosen randomly. On 50% of the trials there was no perturbation.

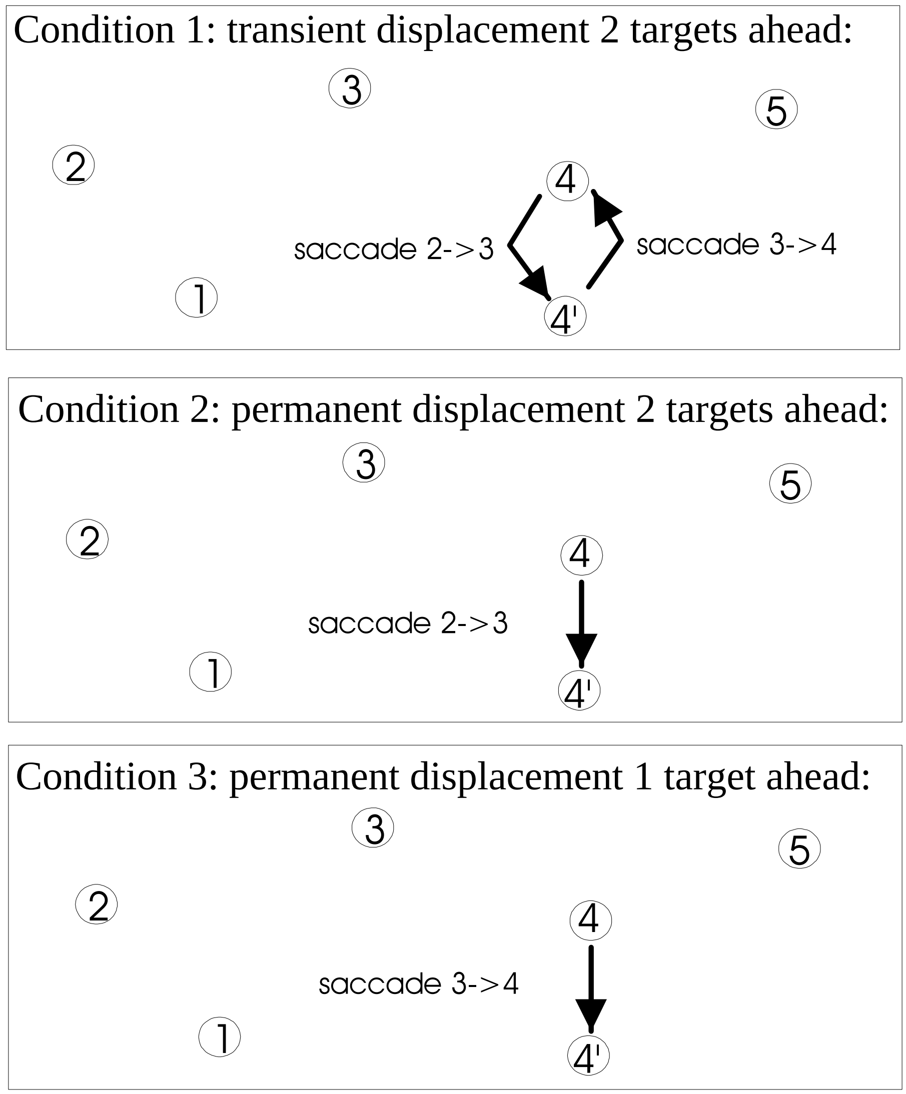

Figure 4-19: Schematic represenation of the 3 conditions. Top: target is displaced when the hand is two targets ahead, and displaced back when the hand is one target ahead. Middle: target is displaced when the hand is two targets ahead and remains in the new position throughout the movement. Bottom: target is displaced when the hand is one target ahead and remains in the new position throughout the movement.

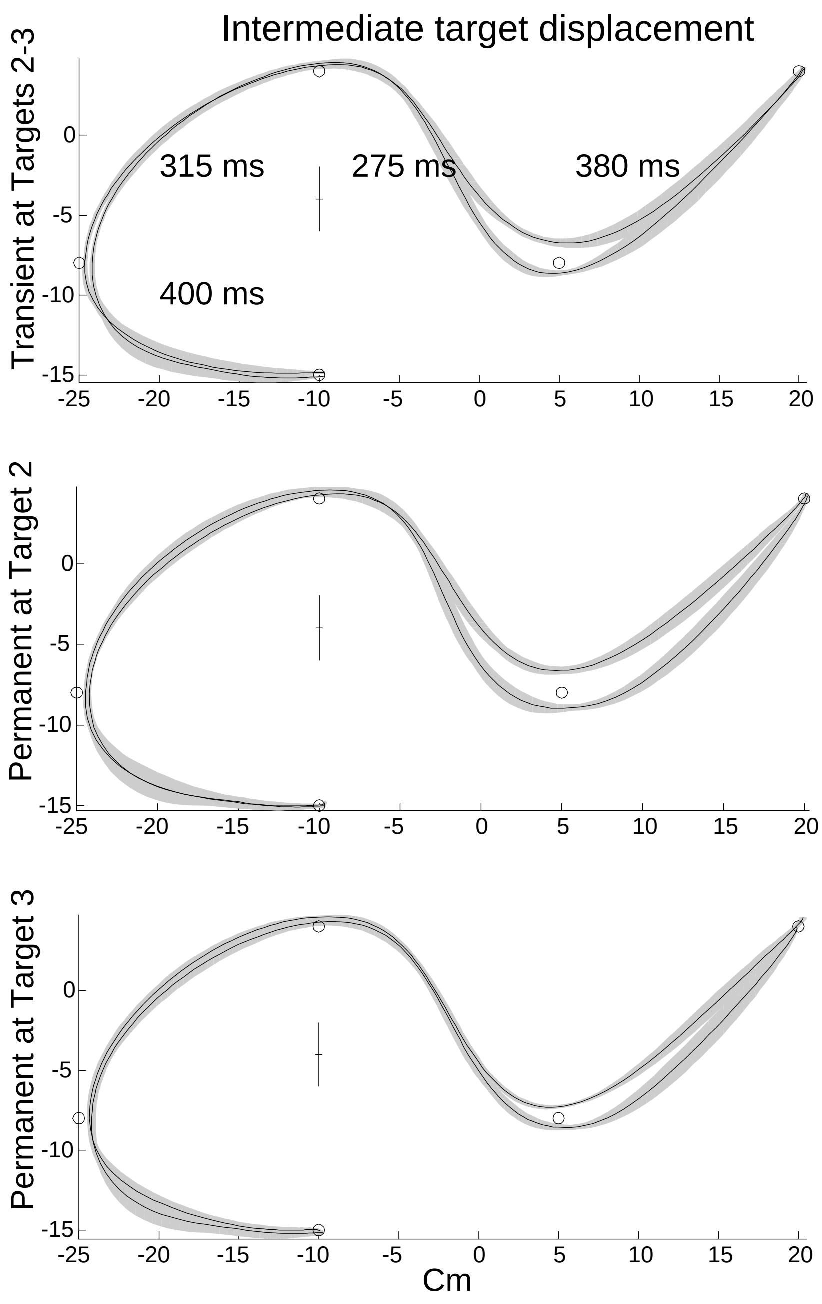

Figure 4-20: For each subject we computed the median path, and then averaged the results over subjects (grey regions correspond to error bars). The vertical mark is twice the magnitude of the displacement (2 cm). Shown also are the average durations for traversing each segment.

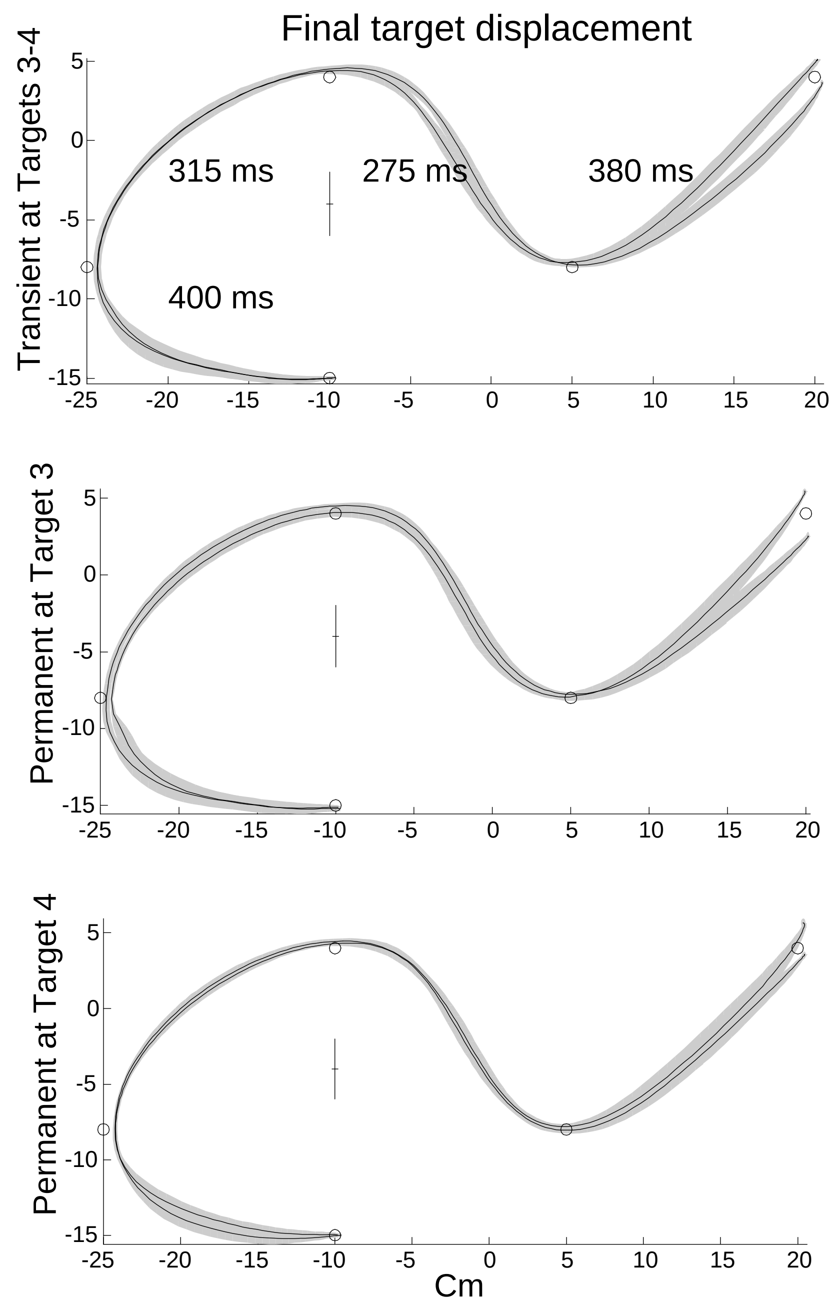

Figure 4-21: The figure format is the same as in Fig 4-12, except we are showing effects for displacements of the final target.

Results

Fig 4-22 shows the templates and the largest path deformations resulting from the perturbations. The PCA constrained to the subspace orthogonal to the average path (to avoid mixing of velocity and path fluctuations) may be a useful method for studying trajectory deformations, but we only use it here to illustrate the main effect.

The surface in Fig 4-23 summarizes the magnitude of target related corrections that occurred for different onset and offset times of the perturbation. For each trial we found the intersection of the path with the line along which the target of interest was displaced, and transformed that in %, where 100% corresponded to passing through the displaced location, and 0% to passing through the original location (note that at the time of passage the target had already returned to its original location). We pooled the perturbation trials (about 1000) from all subjects, and plotted the amount of correction for each trial as a function of the onset and offset times. We then used bicubic interpolation to a regular grid, followed by bicubic spline smoothing to obtain the surface in the figure.

The result is exactly what we would expect from a system using all available information about future targets to adjust the movement. If the perturbation occurs early, and is removed early (e.g. onset -450 ms, offset -350 ms) there is very little effect: presumably the target location is updated during the displacement, but then there is enough time to update it again to the original location - thus we don't see any effect at the time of passage through the target. On the other extreme, if the onset is too late (e.g. -100 ms) there is insufficient time to update the internal estimate of target location and correct the movement. However, if the onset is about -350 ms and the offset about -250 ms (middle of the surface), we see a substantial effect of about 50% - i.e. the 100 ms during which the target was displaced were sufficient to update the estimate of its location and correct the hand movement. Part of that effect is probably canceled during the remaining 250 ms, so the values shown in the figure are actually lower bounds on how much the internal estimates of target location change.

It is clear why the model behaves in the same way - Fig 4-24. While fixating at target 2, it still processes all available visual information, and in particular the displaced location of target 3. This perturbed feedback shifts the internal estimate of target 3, and since the movement is "planned" online[2] it immediately affects the hand path. After the target is moved back to its original location the internal estimate also starts moving in the opposite direction; however that process takes some time and thus the hand path is not restored completely to what it would have been if the target were not displaced.

4.5.3 Experiment 4.5 - Local prism shifts

In this experiment we explore the effects of visual displacement of hand location (i.e. prism shift) on the hand path. Note that our perturbations are randomized (i.e. the direction of the displacement is unpredictable, and on 50% of the trials there is no perturbation), thus this is not an adaptation experiment (we study adaptation in the next chapter). The corrections observed here (as in the above experiments) are all single trial corrections, caused by sensory input received in that trial.

Methods

As before, subjects (8) were asked to move their hand through consecutively numbered targets. During the movement, hand position was displayed only in a region in the middle of the workspace - the periphery

[2] Although the optimal control law (and thus the expected trajectory) exist before movement onset, the actual trajectory results from applying the control law to the online estimates of the state of the limb and the environment. As discussed above, in more complex tasks it may not be feasible to compute even the optimal control law before movement onset.

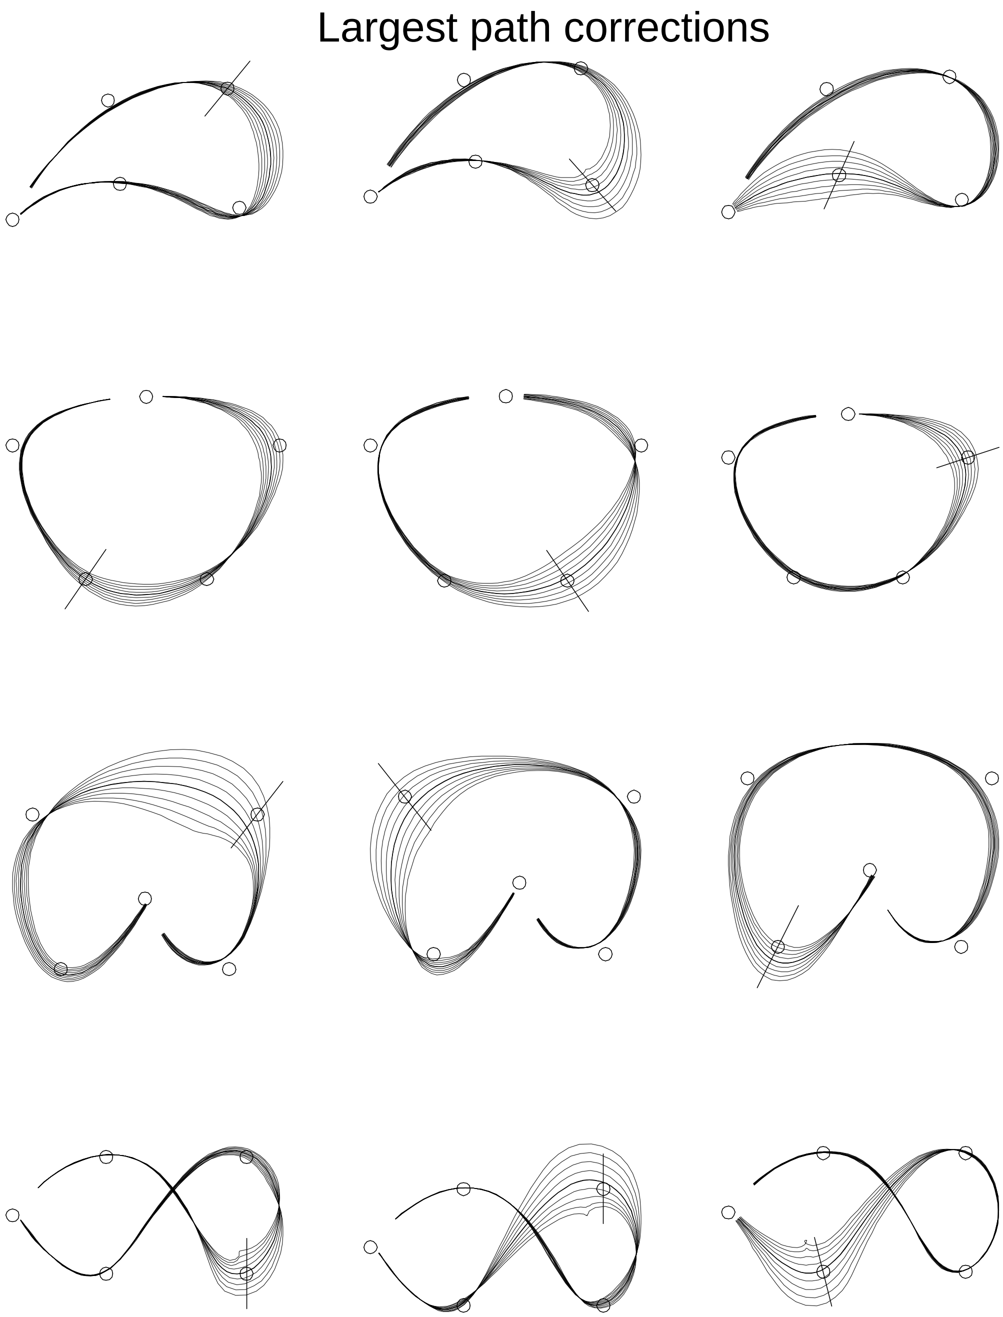

Figure 4-22: Target configurations, average movement paths (over all subjects), and the displaced target are shown. The multiple curves correspond to the largest shape deformation resulting from the perturbation, computed with PCA constrained to the direction orthogonal to the average path - i.e. the first PCA, added to the mean with different weights.

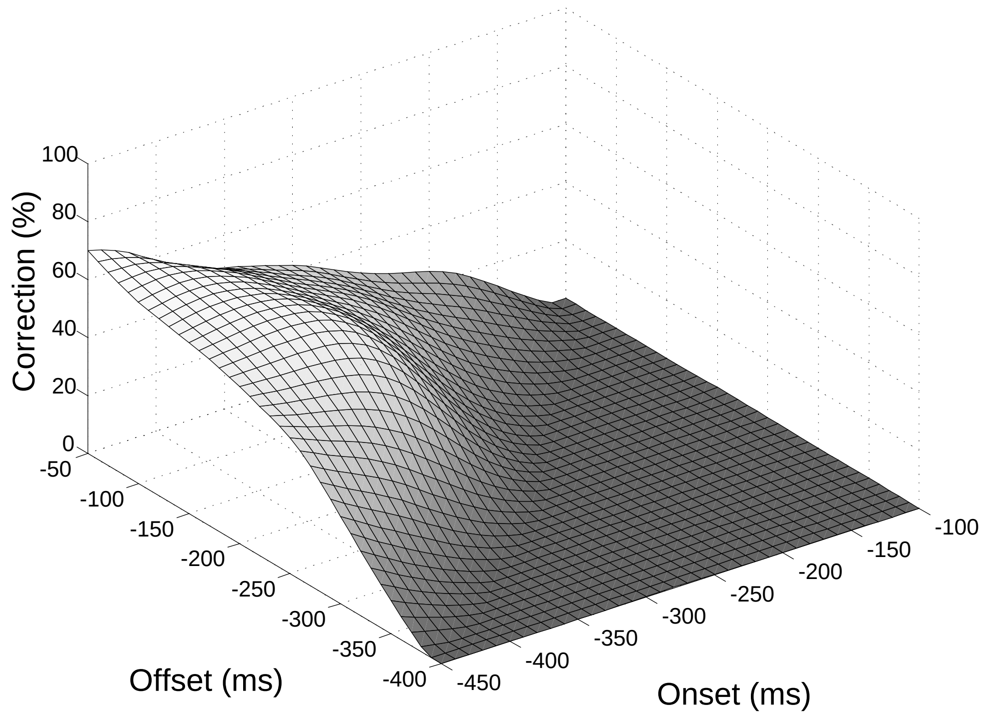

Figure 4-23: Percent correction (in the direction of the displacement) as a function of (onset, offset) time.

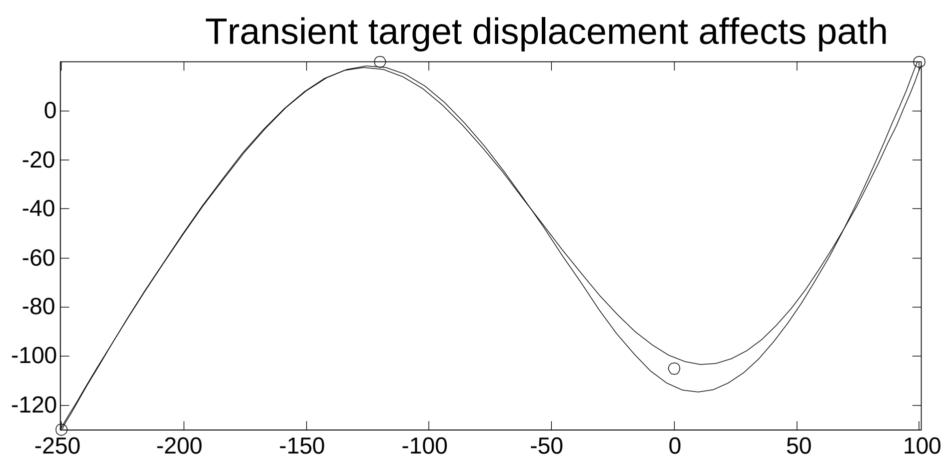

Figure 4-24: Simulation: Transient displacements of target 3 while fixation is maintained at targets 2 and the hand is moving from 1 to 2 affect the path significantly.

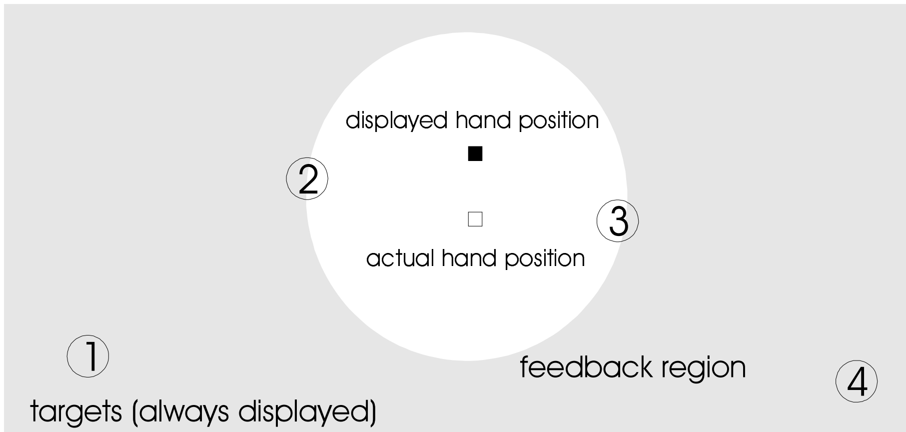

Figure 4-25: Local prism shift paradigm. Hand position is displayed only in the light region, and may be displaced in a random direction.

of the "opening" passed through the middle two targets - see Fig 4-25 (the actual target configuration was different from the figure - see below). After a trial was completed, targets disappeared, and the hand cursor appeared when the subject began moving towards the starting position, i.e. there was no prism displacement on the return, and subjects could see their hand position while repositioning. This paradigm was chosen to allow us to introduce undetected hand displacements. If we had displaced the hand cursor at the beginning of the movement (when subjects know that they are positioned correctly) the perturbation would be rather obvious. However, once subjects began moving it was easy to convince them that their hand was at a different position - all subjects remained unaware of the displacement (2 cm), and interpreted the perceived errors as movement errors.

Results

The result of the undetected prism shift was a rapid correction to the ongoing movement, in the appropriate direction - see Fig 4-26. The exposure to the prismatic shift is about 200 ms, and towards the end of it we begin to see a change. Interestingly, the correction increases in the remaining part of the movement where no visual feedback was presented. This cannot occur if the motor system was using only currently available feedback to correct for deviations away from a desired trajectory. The result suggests that the estimate of hand location obtained during the exposure period was internally extrapolated into the last portion of the movement. Note also that displacements along the movement resulted in larger correction. This observation indicates that the motor system relies on visual feedback to estimate hand location along the movement more than in the direction orthogonal to it. Since the visual accuracy in these two direction is likely to be the same (see Experiment 4.1), the only explanation is that the uncertainty associated with the internal estimate of hand location is bigger along the movement direction (and thus perturbed sensory input has a bigger effect). This agrees with the observation from experiment 4.2 that variability along the movement is larger [Gordon et al. 94]. The present result suggests that the visuo-motor system is actually aware of that fact. Notice also that there is a bias (undershoot) in reaching to the final target, which is to be expected in the absence of visual feedback in that region (from a range effect).

We have performed an earlier experiment where subjects were informed that hand location will be displaced, and asked after every movement (horizontal) to indicate the direction of the displacement (in a forced choice paradigm - either up-down or left-right). The threshold for reliable detection of the displacement along the movement was about two times higher than the threshold along the orthogonal direction, suggesting again that the uncertainty of the sensory-motor system about its internal estimate of hand location along the movement is significantly larger.

We also varied the instructed movement time, which consequently increased the time spent in the prism shift region. The effect of movement duration is shown in Fig 4-27.

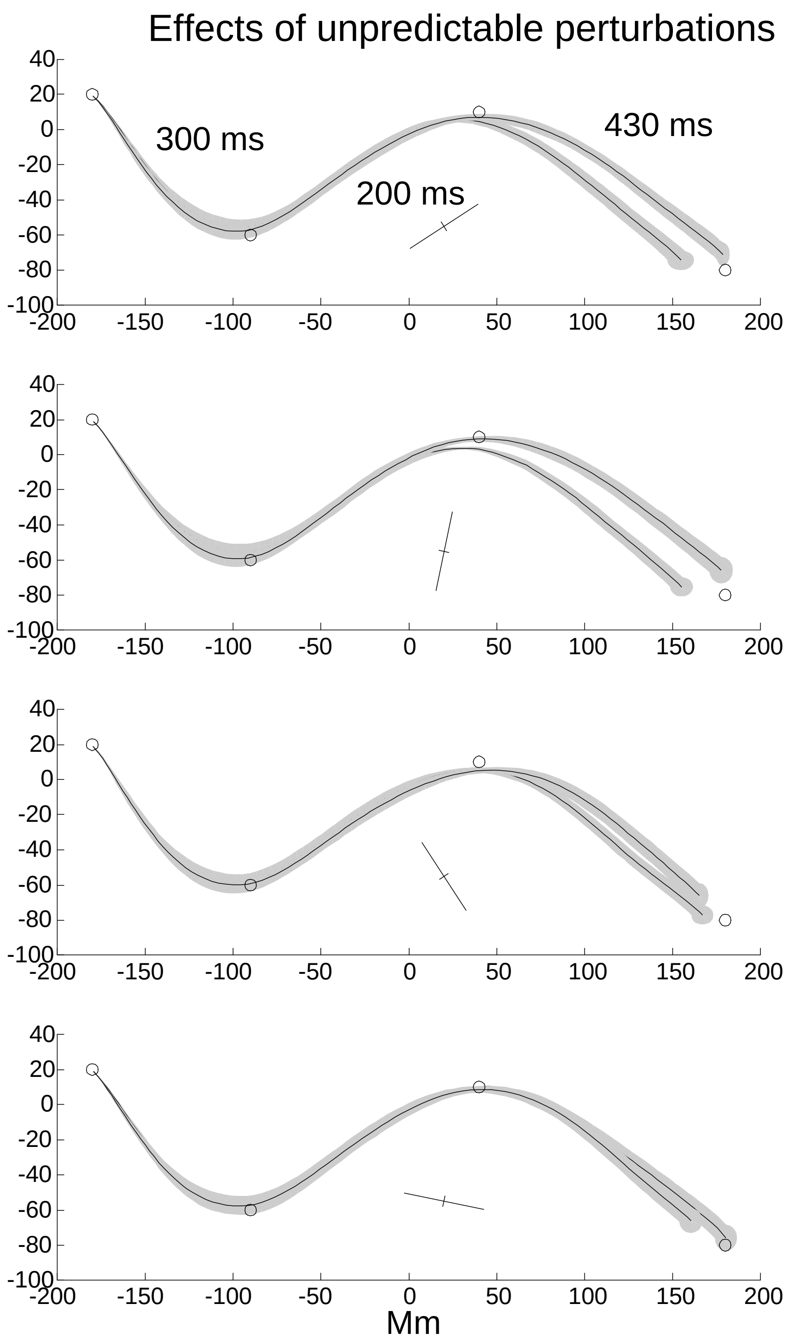

Figure 4-26: Average paths (over all subjects) with corresponding error bars, for perturbations in opposite directions (perturbation direction and magnitude are shown).

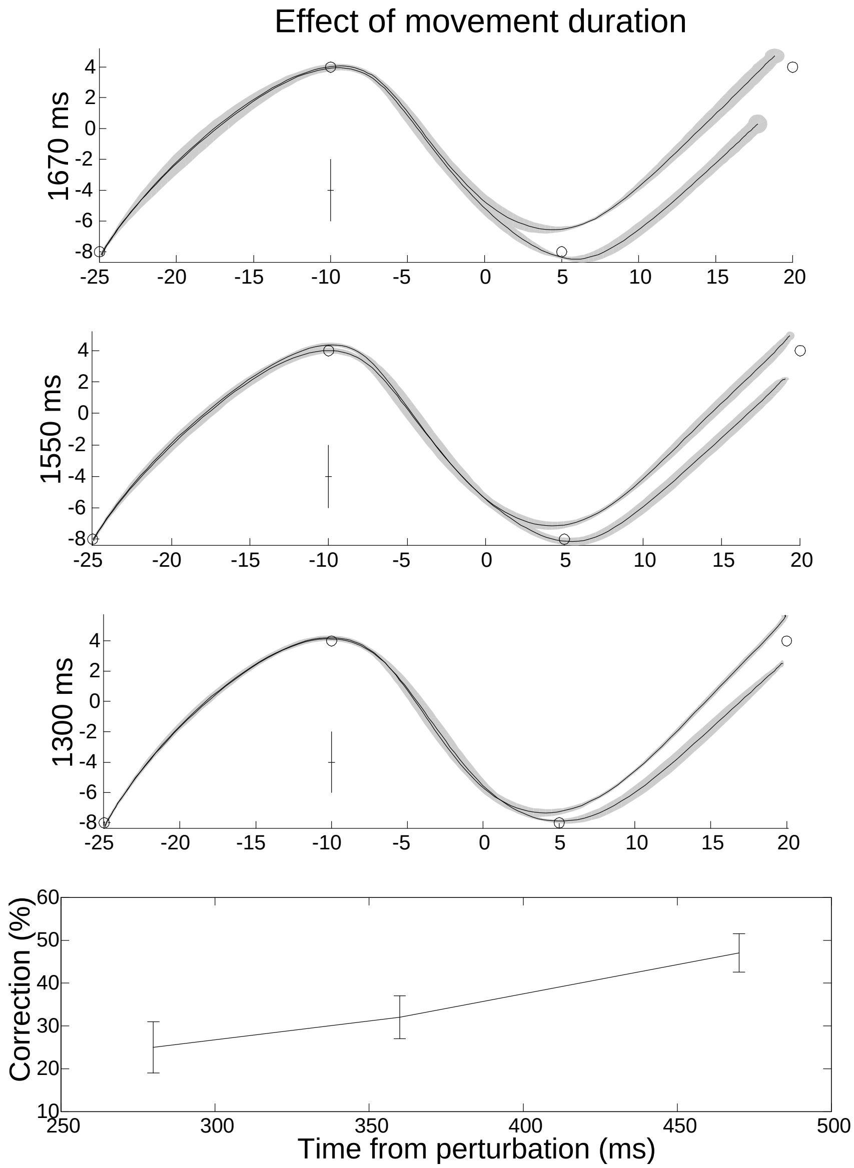

Figure 4-27: Effects of local prism shifts for different movement durations. Bottom - amount of correction at target 3.

Chapter 5

Model inaccuracies and correction strategies

When we introduced our model in Chapter 3, it was assumed that the stochastic optimal controller computes the signals driving a set of "virtual outputs" - Fig 2-4 - that provide "simple" control in end-effector space. In particular, we assumed that the control signals are the X- and Y- net muscle forces applied at the hand, considered as a point mass moving on a planar surface. The possible inaccuracies in implementing end-effector control (i.e. making the hand behave as a point mass) were modeled as noise whose magnitude matched the magnitude of those inaccuracies. However, in reality the errors resulting from imperfect implementation of end-effector control would be state dependent. The sensory part of the model may also be imperfect. Even if the model was exactly accurate, it may not be feasible (in more complex tasks) to compute a control law in real time. In all these cases, it would be possible to improve the performance of the optimal controller by augmenting it with some correction strategies.

In the first section we review experimental evidence that the CNS indeed seems to use a simplified model of end-effector control, and show that our model (in particular the compensation for muscle viscosity included in the control signal) provides a very parsimonious explanation for some types of automatic correction that human subjects exhibit. In the next section we consider a number of (rather odd) neurophysiological results that can be explained by our model. The following section interprets some adaptation results as evidence of correction strategies meant to improve the performance of the controller (rather than directly compensate for the experimentally introduced perturbation). Finally we present data from two learning experiments that are not closely related to the rest of the chapter, but provide a good illustration of some points made in Chapter 2 about the distinctions between novice and expert performance.

5.1 Evidence for simplifications of system dynamics

This section is based on the findings of [Gordon et al. 94b]. These authors asked subjects to make rapid reaching movements in 24 different directions, without vision of the hand. They found a number of systematic effects of movement direction, none of which should exist if perfect end-effector control/kinematic invariance was present. The data was consistent with the hypothesis that the CNS fails to take into

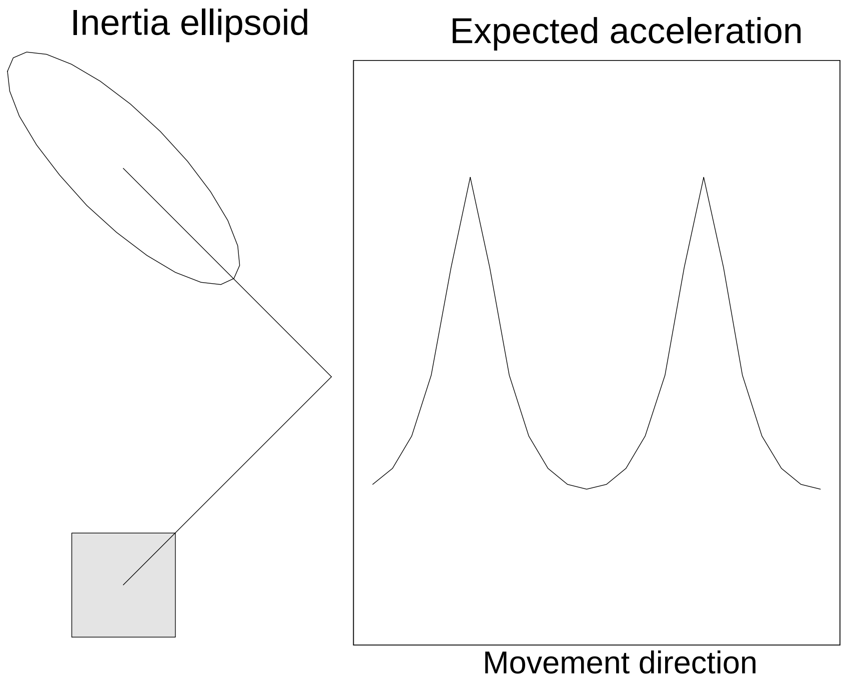

Figure 5-1: Left: Schematic 2-link arm (box=shoulder), with inertia ellipsoid at the hand. Right: acceleration resulting from applying the same end-point force in all directions.

account the anisotropic inertia at the hand, i.e. it applies the same net muscle force in all directions - the prediction of this hypothesis is shown in Fig 5-1. The data is in Fig 5-2 Top - the acceleration ellipsoid is oriented orthogonal to the inertia ellipsoid [1]. If the CNS were to ignore the anisotropic inertia throughout the movement (i.e. output a feedforward force profile), we would expect movements with larger extent in directions with smaller inertia, and equal movement duration in all directions.

Instead, subjects achieved almost constant movement extent (Third row) by reducing movement duration in directions where the initial acceleration was greater (Bottom). The authors argued that it is difficult to explain this compensation with feedback mechanisms. First of all, movements were quite rapid and performed in the absence of visual feedback. More importantly, evidence of compensation was present from the very beginning of the movement - Fig 5-3. Even as early as the time of peak acceleration, movements in directions with larger initial acceleration were already advanced in time. Thus it appears that: a) the CNS is unaware of the inertial anisotropy when it initiates a reaching movement, b) the differences in initial acceleration are compensated by differences in movement duration so that the extent of the movement (which is what the subjects are asked to control) remains constant, and c) that compensation is present from the very beginning of the movement, as if the CNS were aware of its "mistake" while it is making it[2].

Here we will show how these observations may result from the type of control signal used to drive

[1]The figures show simulation results from the model, which are very similar to the original figures.

[2]The authors also consider the possibility that these results arise from equilibrium point control, and reject it through simulations - briefly, equilibrium point control cannot account for these results because the stiffness ellipse is oriented differently.

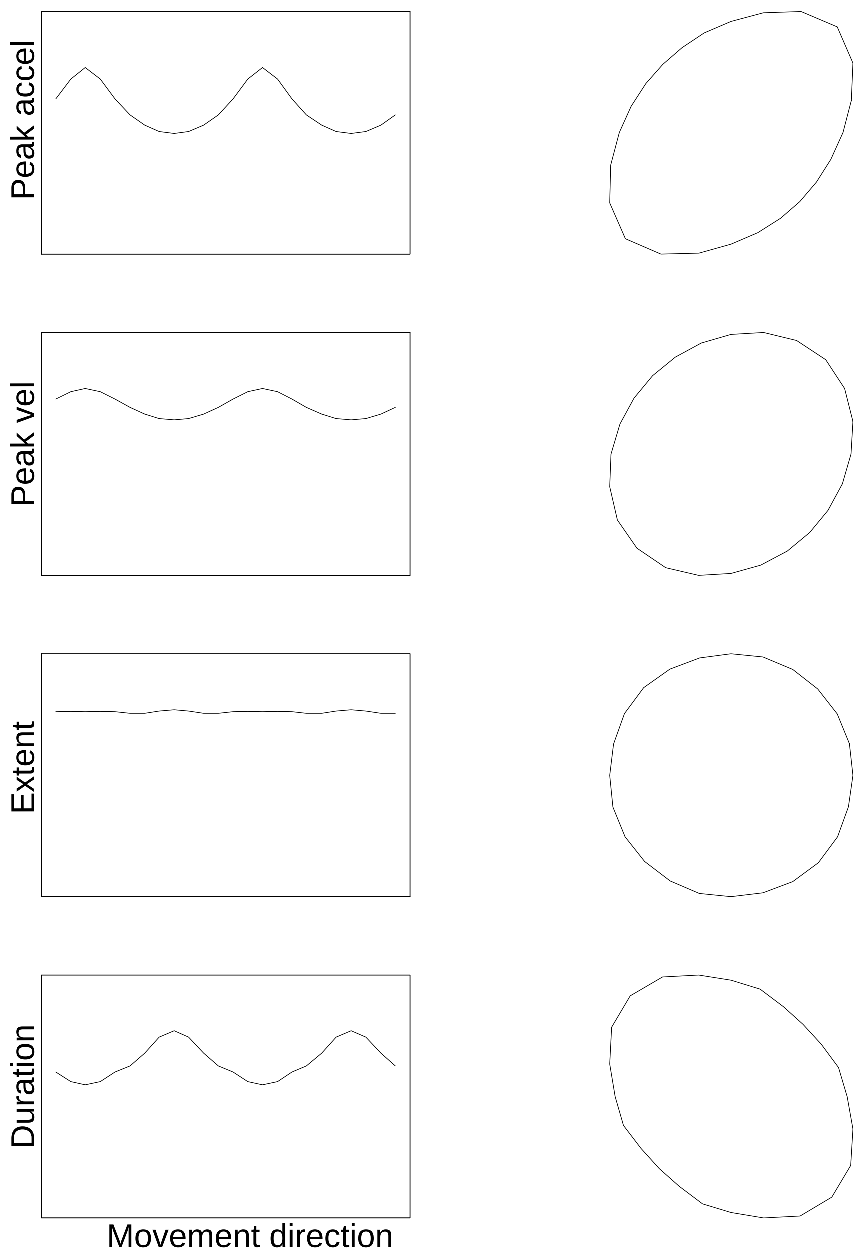

Figure 5-2: Simulation: Peak acceleration (top), peak velocity (second row), movement extent (third row), and movement duration (bottom) as a function of movement direction (left), also shown as a radial plot (right).

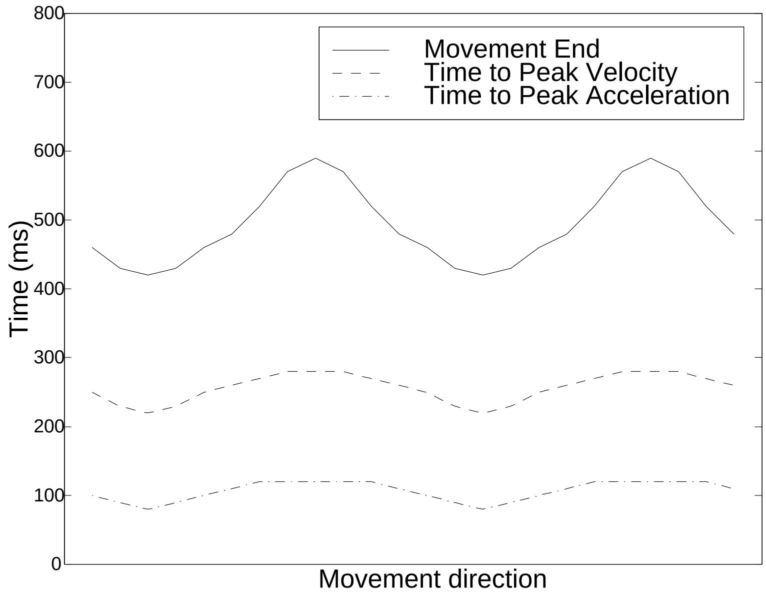

Figure 5-3: Simulation: time of occurence of different movement landmarks (end, peak velocity, peak acceleration) as a function of movement direction.

the hand in our model. In particular, the compensation for muscle viscosity built into the control signal produces exactly the behavior observed by [Gordon et al. 94b]. To simplify the discussion in this and the following section we will only consider a feedforward control signal corresponding to a minimum-jerk trajectory[3]. We model the hand as a point mass moving in a viscous environment, and compute the (feedforward) force profile that would produce a minimum-jerk trajectory in the absence of noise or model errors. We averaged the measurements of [Tsuji et al. 95] for all subjects and directions to obtain estimates of inertia and viscosity at the hand - $m = 1kg$ and $b = 10Ns/m$ respectively [4]. Then the control signal necessary to drive the hand was computed as $\mathbf{f}(t) = m\mathbf{a}(t) + b\mathbf{v}(t)$, where the vectors $\mathbf{a}(t), \mathbf{v}(t)$ denote the acceleration and velocity corresponding to a minimum-jerk trajectory $\mathbf{x}(t)$ in the appropriate direction. If we were to apply this control signal (force) to a point with mass m and viscosity b we would recover $\mathbf{x}(t)$ exactly.

What if the mass is not m in all directions, but instead varies according to the inertia ellipsoid at the hand - Fig 5-1? In some directions (orthogonal to the forearm) the mass will be smaller than m, and so the initial acceleration and velocity will exceed those used to compute the control force. As a result the viscous force will also exceed $b\mathbf{v}(t)$ and movement duration will consequently decrease. This is shown in Fig 5-2. While the changes in initial acceleration and movement duration in this model are always reciprocal (regardless of parameter settings), the two effects do not cancel out to produce a relatively constant movements extent for any parameter settings. Thus it is encouraging that parameters very close to experimental measurements of hand impedance produce this cancelation. Since the "compensation" for different accelerations is due to corresponding changes in the viscous forces (and does not require

[3]This is a close approximation to the behaviour of our model for rapid movements executed without vision and low accuracy constraints.

[4]A small stiffness term (20 N/m) was also used in the simulations, but it does not affect the results in this chapter significantly.

feedback mechanisms), the effect is evident from the very beginning of the movement - Fig 5-3.

Thus we have an example where the internal model of limb dynamics used by the CNS appears to be simplified, but the impedance of the hand combined with appropriate control signals is sufficient to eliminate reaching errors. Of course such low-level mechanisms may not be sufficient in executing more complex movements with higher accuracy constraints. It has been hypothesized previously that such a combination of direct and compensatory control signals may allow the CNS take advantage of the visco-elastic properties of the limb and reduce movement errors [Hinton 84]. Compensation for damping in combination with proportional control rather than direct force control has also been considered in the equilibrium point framework [McIntyre and Bizzi 93].

5.2 Evidence for direct force control combined with damping compensation

In this section we consider a number of electrophysiological observations of primary motor cortex (M1) activity associated with planar hand movements, and argue that they have a natural explanation once we assume that the output of M1 is a control signal with the properties described above.

It has been known for a while that M1 is a crucial component of the cortical network responsible for the control of movement (see [Johnson 92] for a review). Despite years of intense study, it is still not clear what the descending signal[5] means and how exactly it controls the ongoing movement. The idea that M1 is a "final common pathway" which collects all movement-related signals from cortex and transmits them to the spinal cord [Evarts 81] has lost credibility due to the finding that premotor and supplementary motor areas contain about the same numbers of PTNs [Dum and Strick 91]; still, lesions in M1 result in paralysis, while lesions in other movement-related cortical areas have much more subtle effects [Johnson 92] suggesting a fundamental role of M1 activity in the production of all voluntary movements. We can separate the controversies surrounding the representation in M1 it two categories: what is the format (coordinate system) being used, and what is the quantity being encoded in that format.

Support for an extrinsic (possibly Cartesian) coordinate system has come from studies [Georgopoulos et al. 82, Schwartz 94] showing that the population activity in M1 can be used to reconstruct reasonably well the hand kinematics in Cartesian space, and that the firing of individual neurons has a broad cosine tuning with respect to movement direction. However, it is known that under mild assumptions the same results are to be expected even if M1 neurons encoded quantities in joint coordinates, or activation levels of muscle groups [Mussa-Ivaldi 88, Sanger 94]. On the other hand, it has been shown that the population vector[6] rotates around the shoulder when movements with the same extrinsic direction are performed in different areas of the workspace [Caminiti et al. 90]. Furthermore, changing the posture of the arm while a monkey is making reaching movements identical at the level of hand kinematics leads to significant changes in the firing rates/preferred directions in the majority of recorded M1 neurons [Scott and Kalaska 97]. Thus it seems that whatever the format of the representation is, it is closely related to joint coordinates[7]. It has also been argued that there may not exist a unique "format" in

[5] The descending signal is transmitted via the long axons of pyramidal tract neurons - PTNs. While early experimenters were careful to identify PTNs through antidromic stimulation, this is not done routinely anymore. This complicates matters even further since it is not clear what cell types different groups record from.

[6] The population vector is the hand velocity predicted from neuronal mean firing rates under a linear encoding model - [Salinas and Abbot 94].

[7] Interestingly, the overwhelming majority of researchers (see most recent Neuroscience Abstracts) are still focusing on

M1, i.e. one can find evidence for representations in any coordinate system one chooses to look for [Thach 85, Fetz 92]. A mathematical model using coordinate-free representations (basis functions) has been proposed [Pouget and Sejnowski 97].

The nature of the quantities encoded in M1 (which is what we are interested in here) is even less clear than the coordinate system used to encode them[8]. Historically M1 was thought to be involved in higher motor functions (since it is a part of neocortex) and Penfield's stimulation studies seemed to confirm that. Evarts argued that it is a lot closer to the periphery (thus the "final common pathway" idea) by showing strong correlations between firing rates and forces rather than displacement [Evarts 68]. Along the same lines, spike-triggered averaging and microstimulation studies confirmed that PTNs have rather direct effects on EMG activity (however M1 can be active in the absence of EMG changes) [Fetz 92]. On the other extreme, Georgopoulos and collaborators have argued that M1 exerts higher-level control than simply specifying force or muscle contractions - it was shown that movement direction as well as more abstract quantities (e.g. mental rotations, moving objects that have to be intercepted with the hand, changes in force rather than net force) are well represented on the population level.

While some of the apparent contradictions may be explained by the fact that different researchers may be recording from different cell types, it has been shown that the two basic findings - encoding of force under isometric conditions and velocity during movement - coexist in the same population of cells. [Kalaska et al. 89] studied arm movements made under external loads in the same neurons and the same task, and demonstrated that both load- and velocity- related signals are present in the firing rates, and their effects are almost perfectly additive[9]. Furthermore, it was shown that the isometric force population vector has very similar properties to the velocity population vector, including the rotation with the shoulder [Sergio and Kalaska 97].

Why would the firing pattern in M1 encode additively both instantaneous force and velocity, and how could lower levels in the CNS separate those quantities and use them to control the ongoing movement? Decoding is not typically seen as a problem, maybe because experimenters can afford to collect data over long periods of time and arrange so that only one parameter varies and the others are eliminated simply by subtracting a baseline. The spinal circuitry however cannot do that since it has to operate in real time. For example, when two variables (load direction and movement direction) change in the same experiment and the population vector is applied without attempting to average out one of them, it no longer points in the direction of either the movement velocity or the external load [Kalaska et al. 89]. Even if separation of these two signals in real time was somehow possible, we are still left with a fundamental problem. Since the physical arm is an inertial mechanism that converts effort (force, torque) into flow (velocity, momentum) [Hogan and Flash 87], these variables cannot be controlled independently. To go around that problem, different approaches have been tried in robotics, such as hybrid force/position control, impedance control, etc. One can either attempt to control force along some axes and movement along others [Khatib 87], or control both indirectly by specifying an equilibrium state and arm impedance [Hogan and Flash 87]. In either case, the output of such controllers never represents the desired velocity and force (and position) simultaneously along the same axis. It is then not clear why both signals should be embedded in the descending cortical command.

Note that the M1 output seems to correspond to the psychological effort involved. In isometric force production, or load compensation, the perception of effort matches the force exerted at the hand. This

hand kinematics and not even recording joint angles, despite the wide availability of 3D motion analysis systems.

[8] We will assume that whatever the format of the respresentation is, the population vector is a useful description of M1 activity (in a particular coordinate system) - based on the analysis of [Mussa-Ivaldi 88].

[9] A positional gradient has also been observed repeatedly [Georgopoulos et al. 84], and its effect appears to be additive as well. That effect is smaller and we will not consider it here; it may be interpreted as specifying the reference (equilibrium) point for a proportional control term.

is no longer true during free arm movements. Consider a reaching movement forward: since the arm has little passive damping, the accelerating forces at the beginning of the movement are very similar in magnitude and opposite in sign to the deccelerating forces at the end of the movement. Yet the perceived amount of "pushing" at the beginning seems to be much greater than the amount of "pulling" at the end (if the latter is felt at all). From this point of view the output of M1 corresponds more closely to some higher level control mechanism rather than a muscle controller, but the question still remains what this mechanism is trying to control.

The model from the previous section seems to be the simplest way to reconcile the experimental findings discussed above, and interpret the firing pattern in M1 as a meaningful control signal driving a physical arm. Assume that the population activity corresponds to the force exerted at the hand, including a compensation for the expected viscous forces (due to active muscle damping): $\text{Control}\,(t) = m\mathbf{a}\,(t) + b\mathbf{v}\,(t)$. Clearly under isometric conditions that signal equals the force produced by the hand (since there is no movement and thus no damping), while during movement the signal is a sum of the expected velocity and acceleration vectors, weighted by the inertia and viscosity coefficients respectively.

Since extensive measurements of arm impedance in monkeys have not been done, we will use the corresponding numbers from human subjects [Tsuji et al. 95], noting that for the effects discussed here it is the ratio of inertia and damping that matters, not their absolute values. In Fig 5-4 we have plotted the speed and acceleration profiles of a minimum jerk trajectory, and their weighted sum (thick curve). Compared to the speed profile, the control signal is skewed to the left and in particular during the decceleration phase falls off faster than the speed profile. This phenomenon is visible in the data of [Moran and Schwartz 98] who compared the actual velocity profile and the one reconstructed from the population vector, averaged over a huge number of trials from several monkeys and several different experiments. Unfortunately those authors terminated kinematic data recording before the actual end of the movement so the complete decceleration phase of the control signal is not visible, but the part that was recorded clearly shows a steeper decline in the cortical firing rates compared to the speed profile. It may be interesting to verify that observation (and the ones discussed below) using the data of [Benda et al. 97] who studied reaching movements under variable force fields, allowing a better dissociation of movement- and force- related signals.

It has been observed that the M1 population vector roughly corresponds to instantaneous velocity not only during straight reaching movements, but also during more complex drawing movements [Schwartz 94]. Not surprisingly, the cortical signal leads the hand velocity vector by a "prediction interval" (PI) that could result from the delays in neural transmission. It has been observed however that the PI is not constant, but varies with the curvature of the trajectory - the time lag is significantly larger for regions with higher curvature, leading [Schwartz 94] to postulate different types of processing in the transition regions between relatively straight segments. While this explanation cannot be excluded, our model provides a simpler explanation - see Fig 5-5. The control signal is a linear combination of velocity and acceleration. In relatively linear segments of the trajectory both of these vectors point in the movement direction, thus the PI appears to be smaller. When the trajectory begins to turn however, the acceleration vector always turns before the velocity vector, and adding the two results in a control vector that leads hand velocity by a longer PI. Note that both of these phenomena will occur to different degrees regardless of the exact parameter setting of the model, or the trajectory being traced.

While the initial observations of movement-related M1 activity in reaching were described as velocity signals that do not reverse towards the end of the movement (the way a force-related signal would reverse), more careful temporal analysis has revealed reversals of the population vector [Kalaska et al. 92]. It is interesting that such reversals were only seen in side-to-side movements and not in forward-backward movements. This effect was interpreted as evidence that the population vector is partially related to the forces exerted by the hand, but it was somewhat puzzling that reversals only occurred in the movement

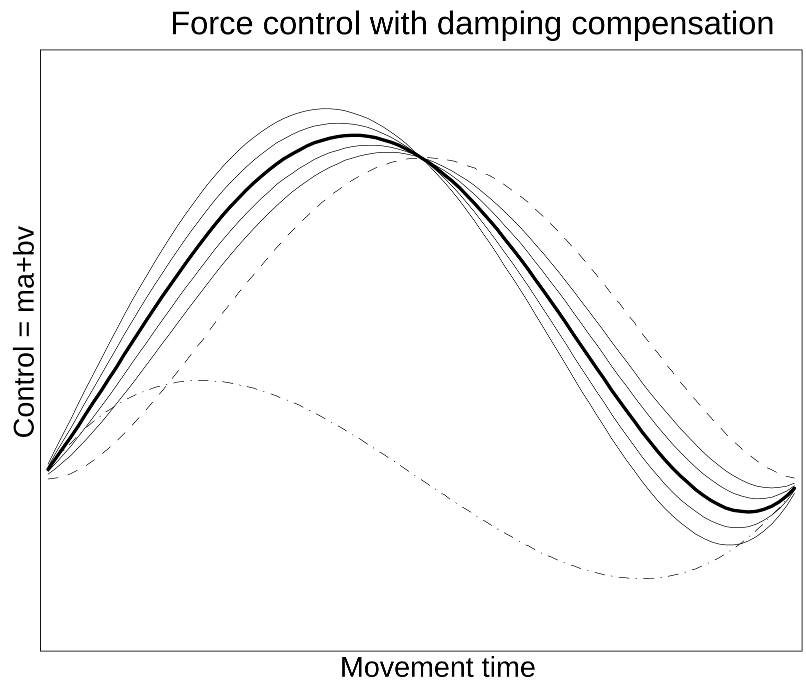

Figure 5-4: Simulation: Thick line - the control signal for our model trying to produce a minimum-jerk trajectory. Thin lines - changes in the control signal for 25% variations in mass. Dashed - velocity profile, Dash-dot - acceleration profile of minimum-jerk trajectory.

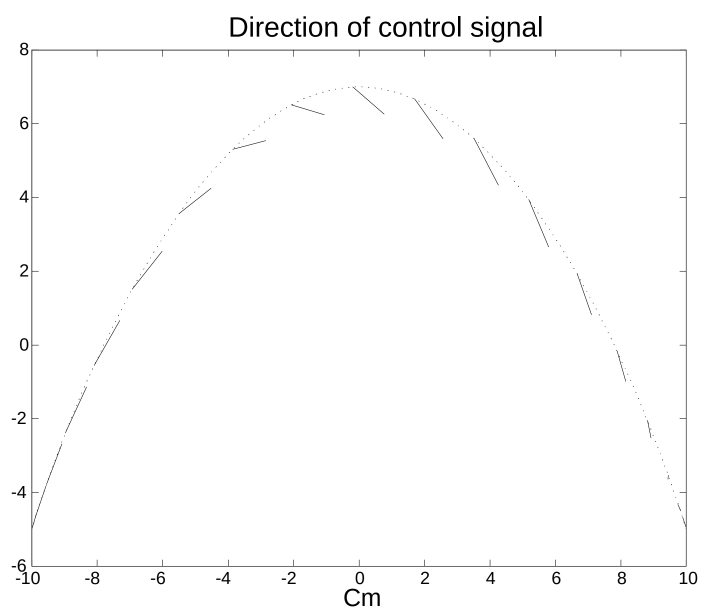

Figure 5-5: Simultion: the control signal (vector) plotted for the corresponding points along a minimum-jerk trajectory through a single via point.

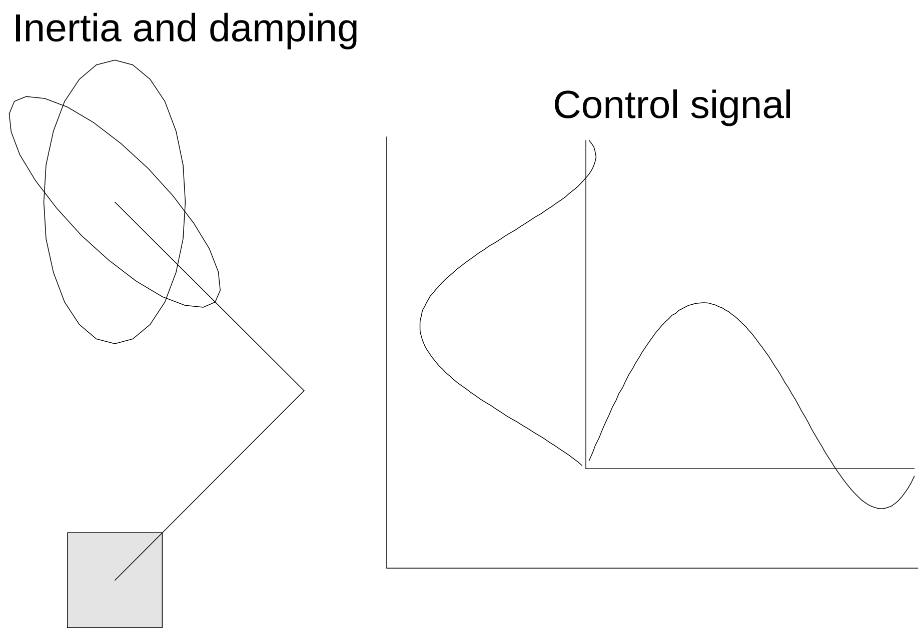

Figure 5-6: Simulation: Left - inertia and viscosity ellipsoids (scaled to fit on the same plot) for averge human subjects. Right - the control signals corresponding to a minimum-jerk trajectory in the two directions.

direction where they would be least expected if one only takes into consideration the inertia ellipsoid. In our model however the presence of reversals in the control signal is determined not by inertia alone, but by the inertia/viscosity ratio. The orientation of the two ellipsoids in a canonical arm posture is shown in Fig 5-6 Left. The viscosity ellipse is known to point away from the shoulder, i.e. the largest viscosity will be experienced during forward-backward movements. This is sufficient to produce significantly bigger reversals of the control signal for side-to-side movements - Fig 5-6 Right. While this effect depends on the exact parameter values, the viscosity ellipses in humans and monkeys probably have similar orientations (and the inertia ellipses are always oriented along the forearm).

Finally, in the study of reaching movements with similar hand kinematics and different postures [Scott and Kalaska 97] the authors attempted to identify the nature of the control signals encoded in M1 by building different models and comparing the predicted amount of postural dependence to their empirical observations. The outcome of their analysis was that neurons encoding joint velocities would be less sensitive, while neurons encoding joint torque would be more sensitive to changes in arm configuration compared to M1 firing rates. This suggests a control signal somewhere "in between" velocity and torque - which is what our model assumes.

5.3 Sensory-motor adaptation

Adaptation - a systematic change in motor behavior resulting from repeated sensory or motor perturbations - has been studied extensively in human motor psychophysics [Shadmehr and Mussa-Ivaldi 94, Welch 86, Wolpert 97]. The simplest and most efficient recipe for adaptation would seem to be the following: detect any systematic deviations from the expected sensory input during movement, identify the source of the deviations as accurately as possible, and introduce a corresponding correction (in either the sensory or motor part of the sensory-motor loop) that cancels the perturbation. This section discusses a number of examples where this simple adaptation scheme does not seem to be used. The most widely studied example is prism adaptation - reaching to visual targets while wearing prism glasses that shift the visual world. If the system were to discover the source of perturbations (i.e. a visual shift) we would expect to see a compensation whereby all visually perceived locations are offset in the opposite direction. Furthermore this is about the simplest perturbation one could imagine, and in addition subjects are aware that their motor errors are caused by the prism glasses. Thus if the sensory-motor system was designed to discover the true sources of perturbation and cancel them out, prism adaptation is one scenario where we would expect to see evidence of that. In particular, we would expect prism adaptation to transfer between limbs - which is not the case [Welch 86]. Instead the resulting generalization patterns have been quite difficult to interpret.

If we insist that adaptation involves an inference about the source of perturbations followed by an appropriate compensation, we would have to conclude that even in the simplest possible case the adapting system is incapable of doing its job. Yet it is a very powerful and robust system, that requires only a few perturbed trials to restore motor performance to its level from before perturbations were introduced. Even fewer trials without perturbations are sufficient to eliminate the adaptation effect [Welch 86, Shadmehr and Mussa-Ivaldi 94, Sabes 96]. What would such a rapid plasticity be needed for in the real world, outside the overconstrained and oversimplified laboratory tasks we typically study? The common answer to this question is calibration for changes in the sensors or actuators, e.g. changes in eye geometry, arm growth, changes in muscle strength, etc. However, if we take the ethological point of view somewhat seriously, such an account of adaptation seems quite absurd. The biological changes that adaptation is supposed to calibrate for occur on the time scale of months or years, and are usually irreversible. Why would we use an adaptation system whose time scale is off by a few orders of magnitude (i.e. seconds), which is ready to jump to erroneous conclusions about the nature of the perturbations based on a few trials and is willing to erase any trace of adaptation following an even smaller number of nonperturbed trials?

We would suggest that the numerous rapid sensory-motor adaptation effects that have been observed do not arise from a system whose goal is to calibrate for slow developmental changes. Instead it must exist for the purpose of very rapid corrections of motor behavior, that are only useful if they can be accomplished based on information from a few trials, and also become irrelevant if they are not useful for even fewer trials. What could necessitate such rapid corrections? Clearly it is not the need to quickly compensate for prism glasses that we accidentally put on. It is then not surprising that the adaptation system does not seem to discover the nature of the prism shifts and compensate accordingly (by having perfect transfer between limbs) - from its point of view a prism shift is not a "simple" (i.e. likely) event, and being forced to make an inference from a few trials the system prefers some other explanation.

So, what is the adaptation system trying to compensate for? We do not think that the experiments necessary to answer that question have been designed yet. The point of the above discussion is that we should perhaps approach adaptation experiments with the understanding that we are not yet doing the right experiment, i.e. the perturbations we introduce and the ways we describe the corresponding changes in motor behavior are not necessarily "prototypical" from the point of view of the adapting system. Here is another example besides prism adaptation. Experiments with viscous force fields imposed by a robot arm while the subject is reaching to visual targets have indicated that the motor system builds an internal model of the velocity-dependent force field and uses that model to cancel the external forces

[Shadmehr and Mussa-Ivaldi 94].[10] However, closer examination of this adaptation effect has revealed a more complex picture. For example, if the external force field is introduced only in the second half of the movement, subjects begin to deviate from their baseline trajectories from the very beginning of the movement, where no forces are present. The deviation is in the direction opposite to the force field, so that in the second half of the movement the force pushes the hand back towards the target [Matsuoka 98]. Another aspect of subject's behavior not consistent with a simple model of adaptation is the following observation (made independently by Matsuoka and Thoroughman - personal communication): the overall shape of the trajectories (reflecting the hypothetical internal model of the force field) tends to converge within very few trials, and after that the only changes appear to be in the final phase of the movement when subjects attempt to acquire the target. These observations indicate that adaptation to force fields results from some (difficult to describe) changes in the motor strategy used for reaching, rather than the acquisition of an internal model of the perturbation. It is possible that the primary objective of the adaptation process is restoring performance on the task (i.e. acquire the target within the time limit) rather than restoring the shape of the baseline trajectory - the latter being an epiphenomenon.

The following two experiments provide more examples of adaptation effects where the motor strategy appears to be modified so as to recover performance on the motor task, while the "true" source of the perturbation is not canceled.

5.3.1 Experiment 5.1 - Cursor displacement adaptation

Methods

We studied adaptation to repeated hand perturbations triggered by saccadic eye movements. As before, subjects were asked to pass through a sequence of intermediate targets; the cursor corresponding to hand position was displaced at the time of peak eye velocity from target 2 to target 3, and remained displaced relative to the hand (horizontally, by +2 or -2 cm) for the duration of the movement. The perturbation was systematic, in blocks of 10 trials followed by 5 trials without visual feedback about hand position. We also interleaved blocks with no perturbation.

Results

Subjects adapted in response to the perturbation, so that the hand after adaptation appeared to pass through the targets - Fig 5-7 Top. The speed profiles however adapted in a rather unexpected way - instead of scaling the speed locally in the segment of the movement whose length changed (as the isochrony principle, or minimum-jerk model predicts), the entire speed profile was scaled. When (a different group of) subjects were asked to make nonperturbed movements through offset target configurations resulting in the same paths as the adaptation experiment, the speed profile was now "normal", i.e. it exhibited the characteristics described in the literature. This effect on the speed profiles should not occur if the system was executing "default" movements through a remapped space. Instead it appears that a different motor strategy (control law) is being used, which can be described verbally as "do what you did before, but speed up and pass through the targets".

[10] This is similar to the accout of prism adaptation as discovery of the visual shift and a corresponding offset of visual perception.

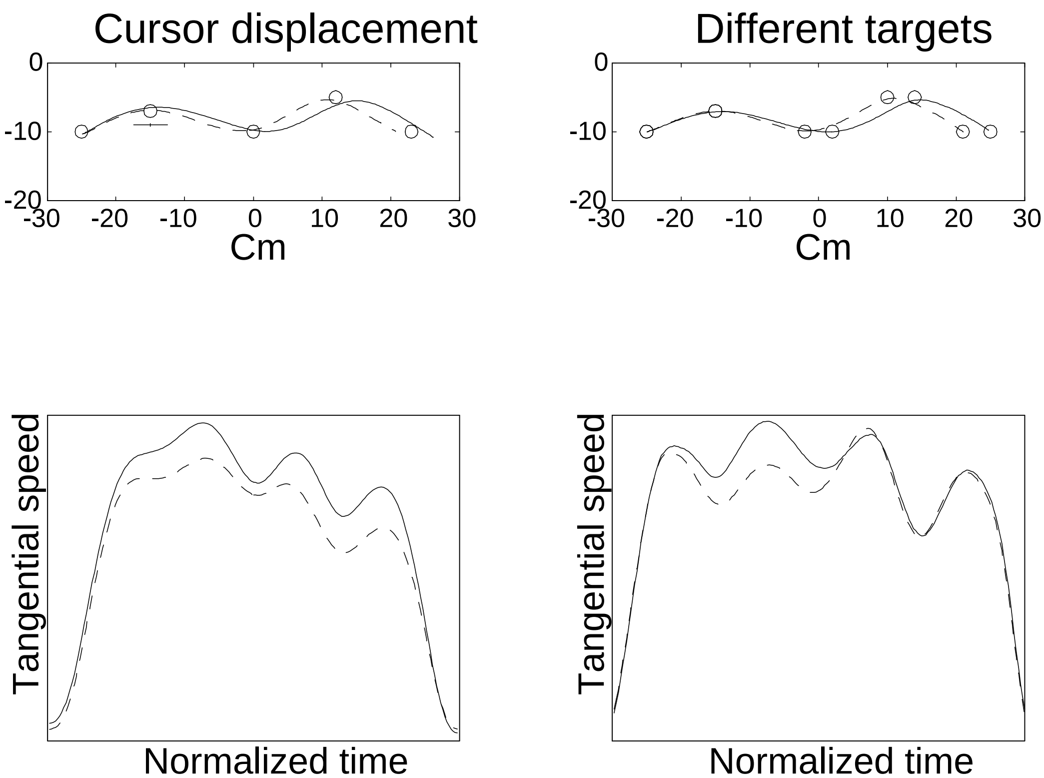

Figure 5-7: Left: aftereffect of adaptation to 2 cm hand displacement left or right, and normalized speed profiles after adaptation for the two displacements. Right: no perturbations - subjects are asked to move through targets offset by +2 or -2 cm, resulting in the same paths but different speed profiles compared to adaptation.

5.3.2 Experiment 5.2 - Local prism adaptation

Methods

This experiment is identical to the local prism shift experiment above (a new group of 10 subjects was recruited), except that here the same perturbation was applied in blocks of 10 trials, followed by 5 trials without visual feedback about hand position. Again, on the return to the starting box the correct hand position was displayed.

Results

Fig 5-8 shows the aftereffect (i.e. trajectories on the no-feedback trials at the end of each block) of the systematic displacement. Subjects corrected their movements appropriately, and retained that adaptation after the displacement was removed. This result is interesting for two reasons. First, the adaptation cannot be explained by a global shift of the visual world, since at the starting position subjects are aware of the actual hand location. Thus we have no adaptation at the starting position, and comparable amounts of adaptation in the exposed and the nonexposed region at the end of the movement. Also, recall that after each trial the hand passes through the same area of space where the displacement was present. But that does not interfere with adaptation during the movement.

Taken together these observations indicate that in the process of adaptation, the sensory-motor system does not really adapt to the systematic perturbation it was exposed to, but instead finds a new way of achieving the goal of the movement. Earlier we raised the question: what is the system constantly trying to adapt to? Maybe it is not really trying to adapt to any systematic perturbations, but instead it is constantly trying to improve its performance on the motor task. One way to do that (as in the model from Chapter 3) would be to modify the default control law it uses.

5.4 Learning new motor skills

This section presents results from two learning experiments. Although not directly related to the model, they illustrate the argument developed in the introduction that learning new motor skills does not involve the acquisition of an identifiable motor template. In the first experiment we demonstrate a form of learning that does not improve performance in the absence of visual feedback (which a motor template should do), yet as soon as visual feedback becomes available it has a significant effect. The second experiment emphasizes the point that performance may increase substantially in the course of a few experimental sessions, without being correlated to identifiable kinematic features of the movements. This observation is particularly interesting since we trained subjects by asking them to replicate a prerecorded expert's movement - yet subjects' performance after training was not related to how well they managed to replicate that reference movement.

5.4.1 Experiment 5.3 - Learning to trace continuous curves

Methods

Ten subjects were asked to trace continuous curves presented on a vertical computer screen, by moving the handle of a planar manipulandum whose position in space was always displayed on the screen. The same curve was presented in a block of 15 trials. We studied different feedback conditions, involving a) full feedback - curve visible throughout the movement; b) partial feedback - curve displayed before

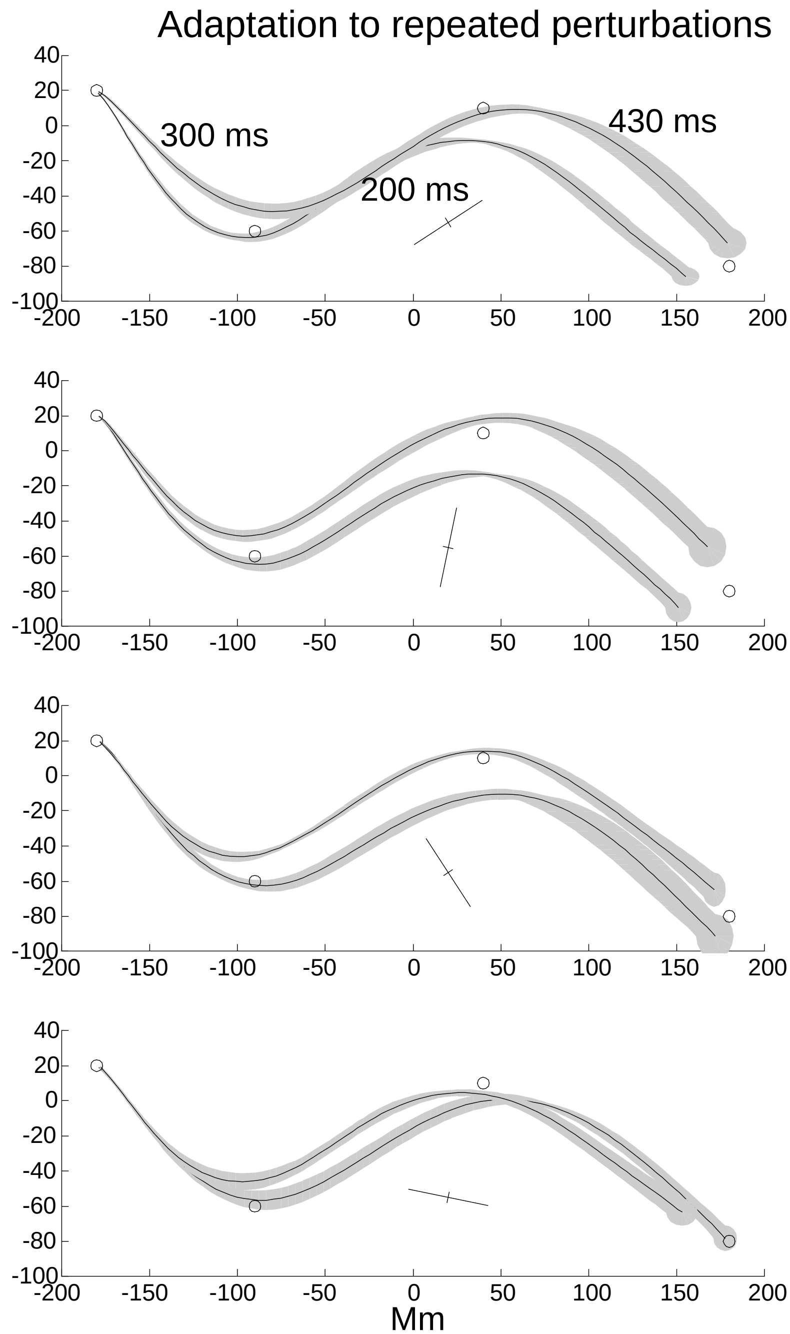

Figure 5-8: Average paths in no-feedback trials following adaptation (aftereffects), shown for all displacement directions.

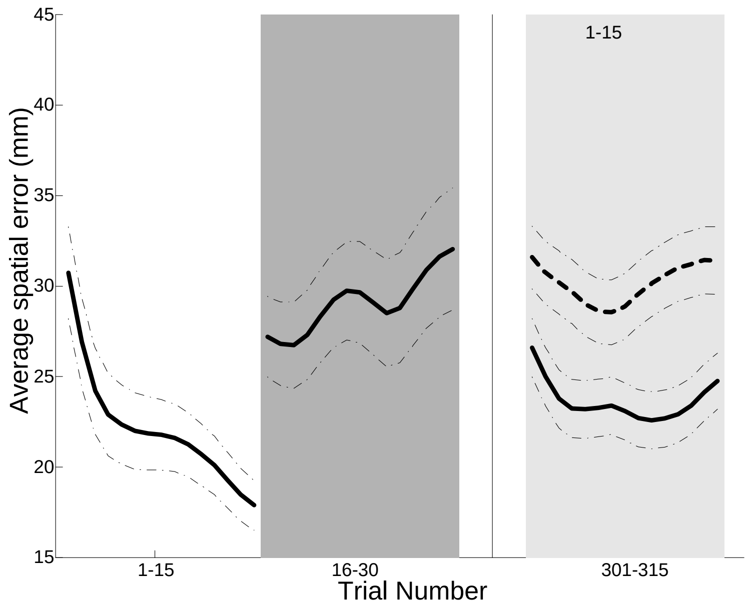

Figure 5-9: Performance (average distance) in curve tracing experiment, for two sets of curves (solid/dashed), presented in full feedback-memory-partial feedback and partial feedback conditions respectively.

movement onset and invisible during the movement; c) curve is never shown in the entire block, and subjects are asked to trace it from memory (memory blocks obviously followed feedback blocks on which the same curve was presented). The set of curves was divided in two subsets. The first subset was presented with full feedback, then no feedback, and 300 trials later (after other curves have been traced) was presented again with partial feedback. The other subset was first presented with partial feedback.

Results

Fig 5-9 shows average performance as a function of trial number - performance was defined as average hand deviation from the desired curve. When a curve is traced with full feedback (left), performance gradually improves. If immediately after that the subject is asked to produce it from memory (middle) with no feedback performance decreases substantially on the very next trial, and keeps deteriorating slowly. The same curve is presented again with partial feedback (i.e. removed from the screen at the beginning of the movement) 300 trials and 20 min later - solid, right plot. Compared to a curve that is presented for the first time with partial feedback (dashed line) performance on the previously practiced curve is significantly better, therefore some learning of that curve has taken place and is preserved during the intervening trials. However, what is retained is not a motor program that specifies what movement to make - if that were the case performance from memory (middle plot) should not decay back to the levels corresponding to a previously unseen curve. Thus the motor system retains an ability to trace the curve better if transient visual information is available, but that does not contribute to performance in the complete absence of visual feedback.

5.4.2 Experiment 5.4 - Learning a table tennis stroke

In this experiment we used a method for training complex skills by displaying subject's movements in real time in a virtual environment, superimposed on an animation of a "teacher" movement that subjects were asked to replicate. This work is published in [Todorov et al. 97]. It is relevant here because it provides an example of motor learning where performance improves quite significantly, yet it is not related to any specific features of the movement that could be interpreted as a motor template or a motor program. This is even more surprising given the fact that we actually trained subjects to execute a specific movement.

Methods

The experimental setup is shown in Fig 5-10. We used a black wooden table (85 cm by 245 cm) and a standard net (height of 15 cm) which was attached 165 cm away from the edge closest to the subject, along with a black tape 15 cm above the net. A square target of size 15 cm was placed horizontally on the table 20 cm behind the net. White table tennis balls were dropped (manually by the experimenter) through a transparent tube that was attached to a tilted platform on the left side of the table. The balls came out of the tube at a velocity and angle typical for an intermediate table tennis shot, and bounced off the table in a highly consistent way. The subject was standing at a convenient distance in front of the table, holding a paddle in the left hand. All subjects in this study were right handed. Therefore, our intention was to train subjects on a difficult task with their non-dominant hand. The task was to let the ball bounce once, hit it with the paddle, and send it through the horizontal opening between the net and the tape such that the ball landed on the target. After every trial the experimenter recorded the score by pressing a button. The interval between two trials was usually in the range of 5 to 10 seconds.

An electro-magnetic sensor (IsoTrack II, Polhemus Corp.) was used to track the position and orientation of the paddle. It included a small (2x2x2cm) receiver attached to the paddle, and a transmitter placed under the table. Both the receiver and the transmitter were connected to a sensor box with a thin cable. The cable from the receiver was placed around the shoulders of the subject and did not restrict movements in any way. The sensor sampled the three dimensional position and orientation of the paddle at 60Hz.

During training the computer displayed a realistic three dimensional simulation of the environment - Fig 5-11. The simulated environment consisted of a graphical representation of the experimental setup, the subject's paddle, the teacher's paddle, and the ball. Illumination, occlusion, and perspective projection were used to provide depth cues. The two paddles and the ball were the only moving objects in the virtual scene. Appropriate sound effects were included in an attempt to replace the missing sensory information at the moment of impact between the paddle and the virtual ball. The position of the subject's paddle in the simulator was updated at 40 Hz, and the sensor delay was less than 20ms.

In training the teacher paddle executed the desired movement repeatedly (waiting for the simulated ball to come out of the tube and bounce, hitting it and sending it to the target), while the subject was trying to move his/her paddle with that of the teacher. After every movement a score reflecting a measure of similarity (in space and time) between the teacher and the subject movements was displayed. The teacher movement was a replay of a recording of an expert player hitting a real ball and sending it to the target. In this replay, the trajectory of the ball was not recorded but constructed from a mathematical model of the dynamics of the ball. This was possible because the ball was always delivered at the same position and velocity. This model was accurate enough so that the real movement recorded from the expert player hit the simulated ball and sent it in the direction of the target.

Twenty subjects were randomly divided to a control group and a training group. Each subject participated in the experiment for 3 consecutive days, 2 sessions per day (separated by a short break). In

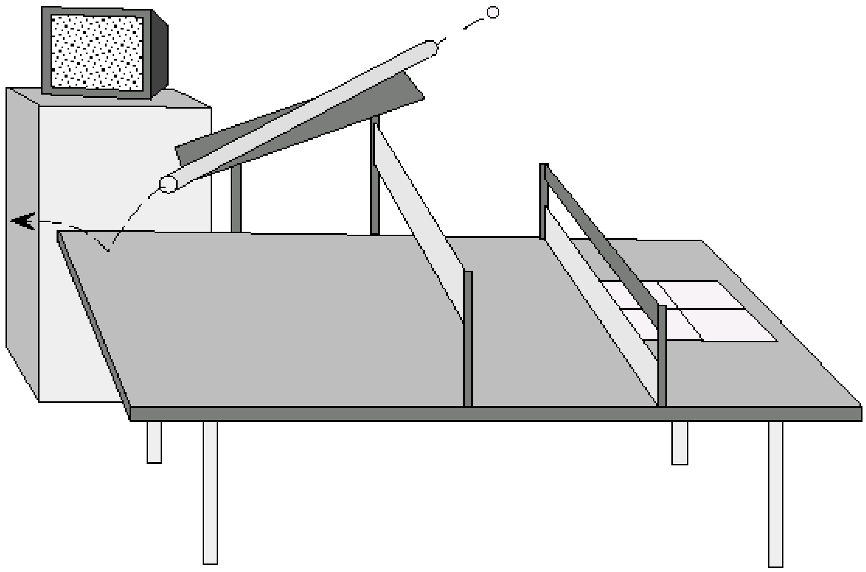

Figure 5-10: Experimental setup. Subjects held a paddle in their left hand, and hit balls that were dropped through the transparent tube. The task was to send the ball between the net and the tape above it, and hit the target behind the net. During training the simulation was shown on a conveniently located computer monitor; a change in body position from training to practice was not necessary.

each session control subjects played 30 balls which were not scored, followed by 60 balls which were scored thus extra practice replaced coaching in this experiment. The experimenter provided verbal feedback on gross errors he noticed (for control subjects only). Training subjects started each session (including the first one) with training in the simulator, and then played 60 balls. Training was terminated after the similarity score between the subject's movements and the teacher saturated at a predefined acceptable value. As a result of this termination criterion experimental subjects were trained longer in the first session than in any other session. The time it took the control subjects to play the first (unscored) 30 balls in each session was matched to the average training time per session for the experimental subjects.

In summary, the control and training subjects spent the same amount of time practicing. However, the control subjects hit 50% more real balls than the training subjects. The remaining practice time for the training subjects was spent in the virtual environment.

Results

The performance of both groups is presented in Fig 5-12. Analysis of variance showed that performance was increasing for both groups ($F(5,80)=4.0$, $p<0.01$). The difference in performance between the two groups over the whole experiment was not within accepted significance levels (main effect of group was $F(1,16)=3.2$, $p=0.09$). The reason is that we are looking at learning curves which are sampled often, so we should not expect to see a main effect of training, but rather a difference at the end of the experiment. Indeed, the score of the training group was significantly higher than the score of the control group on the last session (t-test, $t(16)=2.41$, $p=0.02$). Also, the improvement of the training group was bigger than the improvement of the control group (t-test, $t(16)=2.1$, $p=0.05$). Interestingly, the training group was better on the very first session (t-test, $t(16)=2.2$, $p=0.04$) but the difference disappeared on the

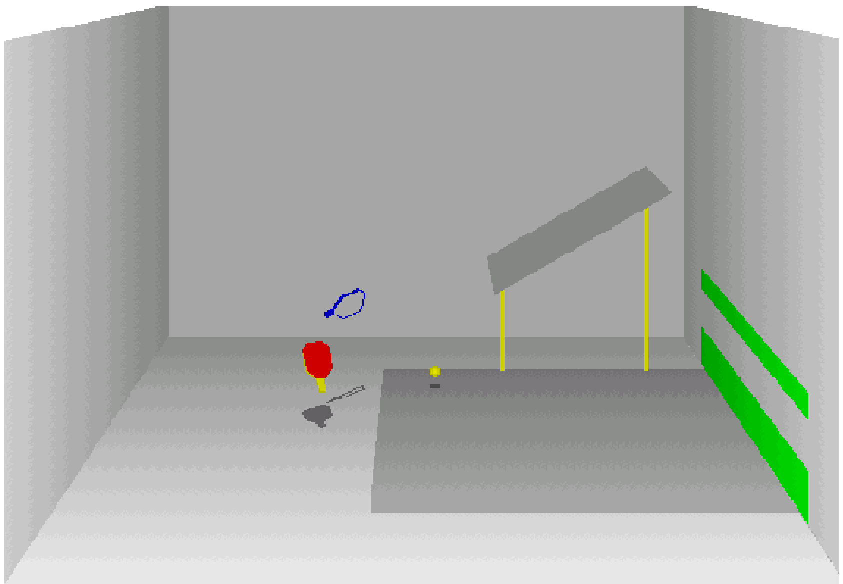

Figure 5-11: Virtual environment. A 15" SVGA monitor displayed a simulation of the teacher paddle hitting the ball, superimposed on the movement of the subject paddle (tracked by a magnetic sensor). Different colors and intensities were used for the elements in the scene, which was rendered in 648x480 pixels, 256 color mode.

next session.

Even though the training group as a whole executed movements similar to the teacher trajectory, we observed that individual subjects modified the teacher movement in subject-specific ways. As an example, Fig 5-13 shows the evolution of the movements of one training subject over the 3 days of the experiment. The general shape is similar to the teacher, particularly during the forward swing, but an extra twist in depth is introduced, and it becomes more pronounced with time. Other subjects deviated consistently in the other two dimensions, indicating that these subject-specific modifications are not caused by poor perception of the teacher movement in depth.

We could not identify any trajectory features (within the training group) correlated with performance.

A similar form of online feedback (in a much simpler experimental paradigm, involving a single-joint elbow movement - introduced by [Schmidt and Young 91]) has been studied by [Brisson and Alain 96]. These authors also found that tracking in real time a prespecified trajectory resulted in a very significant increase in performance. They also attempted to correlate performance (between subjects) with a number of kinematic features, and concluded that no such correlations exist. An interesting conclusion of that study is that the "quality" of the prerecorded movement is not crucial - training with the best movement of the best subject vs. a movement chosen from baseline performance of each individual subject resulted in equal improvement.

Taken together these experiments suggest that even when learning is clearly a result of attempting to reproduce a specific movement template, how well the template is acquired does not affect performance and furthermore the exact template being tracked is not crucial. Why is it then that performance improves significantly, compared to equal amounts of practice? The only possible explanation seems to be that the process of tracking itself contributes to learning, and it doesn't really matter what template

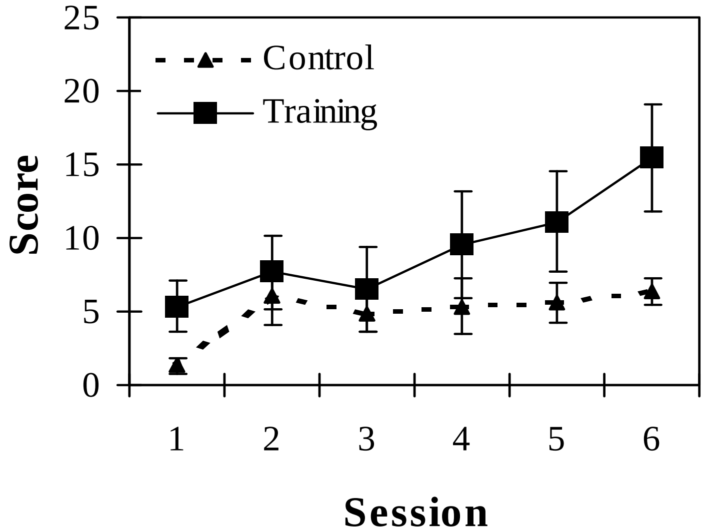

Figure 5-12: Performance. Scores averaged over subjects in each group and session, with standard error bars, are shown. The score corresponds to the number of target hits (60 trials per session).

is tracked and whether it is eventually acquired. This is somewhat related to the result from the previous experiment (5.3) where we saw that tracing a continuous curve has a learning effect, but it cannot be revealed in the absence of visual feedback, i.e. it is not due to the acquisition of a particular movement trajectory.

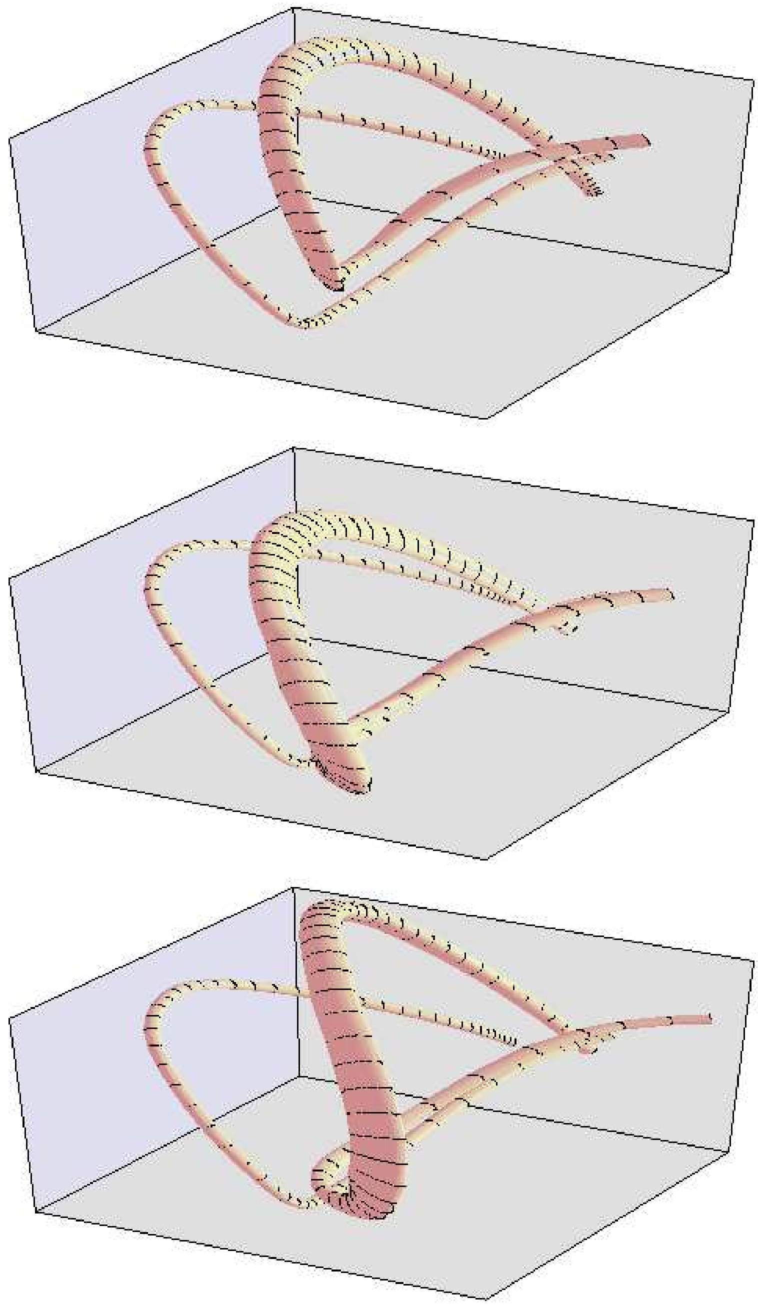

Figure 5-13: Summarized trajectories for one training subject executing the real task on day 1 (A), day 2 (B), and day 3 (C), superimposed on the teacher trajectory.

Chapter 6

Kinematic Regularities of Spatially Constrained Hand Movements

6.1 Overview

The speed profiles of arm movements display a number of regularities, including bell-shaped speed profiles in straight reaching movements and an inverse relationship between speed and curvature in extemporaneous drawing movements (described as a 2/3 power law). Here we propose a new model that simultaneously accounts for both regularities, by replacing the 2/3 power law with a smoothness constraint. For a given path of the hand in space, our model assumes that the speed profile will be the one that minimizes the third derivative of position (or "jerk"). Analysis of the mathematical relationship between this smoothness constraint and the 2/3 power law revealed that in both two and three dimensions, the power law is equivalent to setting the jerk along the normal to the path to zero; it generates speed predictions that are similar, but clearly distinguishable from the predictions of our model. We have assessed the accuracy of the model on a number of motor tasks in two and three dimensions, involving discrete movements along arbitrary paths, traced with different limb segments. The new model provides a very close fit to the observed speed profiles in all cases. Its performance is uniformly better compared to all existing versions of the 2/3 power law, suggesting that the correlation between speed and curvature may be a consequence of an underlying motor strategy to produce smooth movements. Our results indicate that the relationship between the path and the speed profile of a complex arm movement is stronger than previously thought, especially within a single trial. The accuracy of the model was quite uniform over movements of different shape, size, and average speed. We did not find evidence for segmentation, yet prediction error increased with movement duration, suggesting a continuous fluctuation of the "tempo" of discrete movements. The implications of these findings for motor planning and online control are discussed.

The model presented in this section is a descriptive one, in the sense that it describes geometric properties of hand trajectories rather than suggesting an actual mechanism for their generation (unlike the optimal control model developed in Chapter 3). As we will see, it provides a very close fit to experimental data from spatially constrained hand movements[1]. Of particular interest here is the finding that the model describes single trial data better than average trajectories, suggesting that it may be closely related to the model from the previous chapter. At this point however the mathematical similarity between the two models is not very clear - the difficulty in phrasing this model in the optimal control framework is that it prescribes an explicit relationship between features of inidivual movements (i.e. speed and local path shape) none of which is being represented explicitly in the optimal control model. The work presented in this chapter is published in [Todorov and Jordan 98].

6.2 Reconstruction of speed profiles from hand paths via smoothness maximization

The majority of human arm movements produced in everyday life are underconstrained: the desired effect on the external environment can be achieved in a large number of ways, thus the exact movement executed at a particular time is not completely determined by the motor task. We effortlessly plan and execute underconstrained arm movements that achieve the desired effect, and are at the same time smooth and graceful. Furthermore, we can produce appropriate movements in tasks that involve almost arbitrary combinations of constraints: spatial locations that the hand should pass through, objects that have to be avoided, forces that have to be exerted with very accurate timing, requirements on the speed at certain points along the path, desired durations of the entire movement or smaller segments of it, etc. This impressive ability to take into account a diverse set of task constraints and rapidly "fill in" the missing details of the movement suggests that the motor system has very general and efficient mechanisms for solving ill-posed problems. Consequently, a fruitful line of research in the field of motor control has been to identify, mostly through observational studies, regularities in biological movement that are not in any way implied by the task. Once identified, such regularities can be used to infer the structure of the underlying motor system that produces them. Examples of this approach include: Listings law, Fitts law, the isochrony principle (Viviani and Schneider, 1991), linearly related joint velocities (Soechting and Terzuolo, 1986), piece-wise planar $3D$ hand paths (Soechting and Terzuolo, 1987), straight paths and bell-shaped speed profiles of reaching movements (Morasso, 1981), the inverse relationship between instantaneous tangential speed and curvature of the hand path (Jack, 1895; Lacquaniti et al., 1983).

In this work we focus on the relationships observed between the hand paths and speed profiles of arm movements that are spatially constrained. We begin with a brief summary of existing versions of the minimum-jerk model (Flash and Hogan, 1985) and the 2/3 power law (Lacquaniti et al., 1983), and compare their implications for the biological mechanisms underlying trajectory formation. We then develop a new model that essentially combines the appealing features of these prior models; it simultaneously accounts for both the bell-shaped speed profiles of straight reaching movements and the inverse relationship between curvature and speed of curved hand movements. In a series of experiments we demonstrate that the new model provides an accurate description of a number of discrete arm movements with arbitrary hand paths in both two and three dimensions. Finally, we discuss the implications of our observations for the generation of complex hand trajectories.

[1] It does not take into account sensory-motor noise and desired accuracy, so phenomena related to the speed-accuracy tradeoff cannot be modeled in this framework.

6.3 Descriptive Models of Trajectory Formation

In this section we distinguish three groups of models (Fig 6-1) based on the levels of biological processing they parallel. Our classification emphasizes the implications of different models for motor control, rather than mathematical similarities among them.

6.3.1 Complete models of trajectory formation

The most ambitious group of models predict the outcome of the processing in the entire CNS in a restricted set of tasks. Such models provide a complete "recipe" for trajectory generation, starting with the externally specified task description, and leading to an observable trajectory. It would be ideal if one could construct a model of this kind that agrees closely with a wide range of experimental data; unfortunately, this has proven to be difficult.

A general method for constructing such models is provided by optimization theory: one defines a "cost" functional, and attempts to show that the movements typically produced in a given task are the ones that minimize cost, subject to the constraints imposed by the task. A number of models based on optimization theory have been considered (Nelson, 1983; Uno et al., 1989; Wann et al., 1988; Flash and Hogan, 1985), the cost being energy, force, impulse, peak speed, peak acceleration, duration, or torque change. The model which seems to account for the largest body of data, while at the same time being attractively simple, is the "minimum-jerk" model (Flash and Hogan, 1985). The cost is defined as the squared jerk (third derivative of hand position over time), integrated over the entire movement; it is minimized over all trajectories that have the same initial point, final point, and possibly pass through a small set of intermediate (via) points specified in the task itself. In reaching tasks the minimum-jerk model predicts straight paths and bell-shaped speed profiles that match experimental observations rather well (Flash and Hogan, 1985). There are cases however where subjects systematically produce non-straight reaching movements (Atkeson and Hollerbach, 1985). Since this model is purely kinematic, it predicts invariance with respect to translation, rotation, and spatial or temporal scaling. It has been shown that planar rotation of the desired movement can significantly affect the hand path in via-point (Uno et al., 1989), as well as obstacle avoidance (Sabes, 1997) tasks. We show more examples of this phenomenon later.

6.3.2 Models of intrinsic regularities

On the other extreme, we have models that do not explicitly assume anything about the motor task, or the flow of information processing in the CNS. They simply quantify intrinsic relationships between features of the observed trajectories that are not related a priori (in this case between the path and speed profile). Note that a causal relationship between the two features is not necessarily implied. Such models are consistent with any proposed mechanism of trajectory formation that produces the observed relationship either through an explicit computation, or as an emergent property of some other, possibly unrelated process.

It has been proposed (Lacquaniti et al., 1983) that for a large class of movements with curved hand paths, the speed $v(t)$ is related to the curvature $\kappa(t)$ through a power law: $v(t) = \gamma\kappa(t)^{-1/3}$, where γ is a constant gain factor. The validity of this law is typically demonstrated through a correlation between v and $\kappa^{-1/3}$, or between $\log(v)$ and $\log(\kappa)$. While the resulting correlation coefficients are rather high (Viviani and Cenzato, 1985), it is not immediately obvious from such analyses how much of the details of the speed profiles are captured, apart from the general trend that speed decreases with increasing curvature. It has been argued, for example, that the actual speed profiles can be skewed relative to the

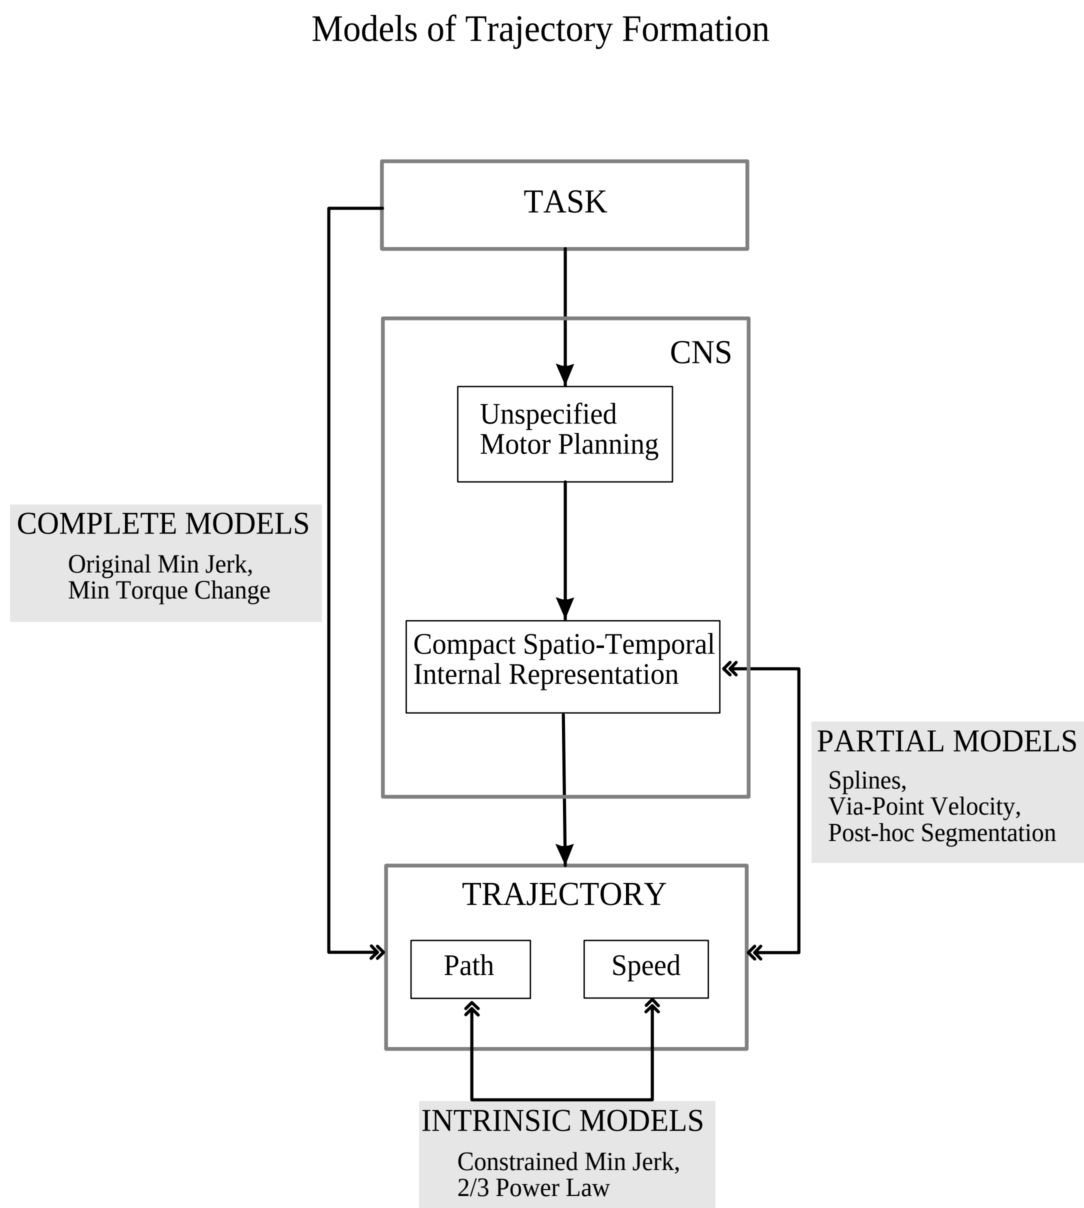

Figure 6-1: Schematic representation of our model classification scheme. Grey boxes correspond to groups of models, and dashed arrows are the predictions these models make. Solid arrows represent the actual flow of information in the CNS (the internals of the "CNS" box are of course hypothetical). The observed trajectory can be further separated into a priori independent features, in this case speed and path.

power law prediction, and the correlation between $\log(v)$ and $\log(\kappa)$ still remains high (Wann et al., 1988). The analysis that could resolve this issue, which thus far has been absent from the literature, is a direct comparison superimposing the actual and predicted speed profiles. In its original form the power law has two limitations: a) it does not apply to paths that have inflection points or straight segments (i.e. for $\kappa = 0$ the predicted speed is infinite), and it is inaccurate at low curvature; b) the fact that $v(t)$ decreases towards 0 at the endpoints of discrete movements cannot be captured. To fix the first problem, Viviani and Schneider (1991) suggested a modified power law: $v(t) = \gamma(\kappa(t) + \epsilon)^{\beta}$, where ϵ depends on the average speed and β is a free parameter estimated from the actual data.

Despite these limitations, the power law seems to provide a surprisingly good description of the speed profiles of complex movements, which has motivated recent attempts to find an explanation of this regularity. Within the framework of equilibrium-point control (Bizzi et al., 1992), Gribble and Ostry (1996) have shown that if the equilibrium point is moved along an ellipse with a constant speed, the hand traces a similar ellipse obeying the power law. A potential problem with this explanation is that the validity of the equilibrium-point hypothesis itself has recently been questioned in studies arguing that the required stiffness is much larger than experimental measurements under dynamic conditions (Bennett et al., 1992; Gomi and Kawato, 1996)[2]. Another proposed explanation (Pollick and Sapiro, 1996) is the mathematically interesting fact that a point moving according to the power law maintains constant "affine velocity"; it remains to be clarified how affine velocity is related to arm movements, and why it might be advantageous to keep it constant. An alternative suggestion, that we will consider in the present paper, is that maximally smooth (i.e. minimum-jerk) movements have speed profiles described by the power law (Wann et al., 1988). Viviani and Flash (1995) tested this empirically by applying both the 2/3 power law and their modified minimum-jerk model (see below) to the same experimental data. These authors concluded that the speed-curvature relationship described by the power law is not implicit in the minimum-jerk hypothesis.

6.3.3 Partial models of trajectory formation

Finally, a number of intermediate models have been proposed, attempting to bridge the gap between providing a complete recipe for trajectory formation and simply quantifying intrinsic regularities of observed trajectories. Such models assume that the motor task is transformed (via some unspecified planning process) into a compact spatio-temporal representation, involving a small number of parameters that are sufficient to generate the observed trajectory.

Morasso and Mussa-Ivaldi (1983) suggested that complex hand trajectories are composed of partially overlapping linear "strokes" (modeled as B-splines with bell-shaped speed profiles), that can be scaled and positioned as to form a desired trajectory. The parameters of the model (timing, scale, and position of each segment) were initially extracted from the experimental data, and then adjusted iteratively until a good fit to the observed trajectories was obtained.

Viviani and Flash (1995) proposed a modification of the minimum-jerk model, which assumes that the movement is internally represented as a set of salient points, and the velocity (speed direction and magnitude) at these points. The authors extracted the velocities from the observed hand trajectories, and used standard optimization methods to solve for the resulting minimum-jerk path and speed profile (the solution is a concatenation of 5^{th} order polynomial splines). This model has also been studied by Yalov (1991), who found that overall it is more accurate than versions of the minimum-jerk model utilizing only the position (and tangent to the path) of a via-point.

[2] Recent arguments regarding the nonlinearities of hand impedance and whether they can accout for this low stiffness [Gribble et al. 98, Kawato 98] have focused exclusively on rapid reaching movements so we will not consider them here.

Another possible model in this category is "post-hoc segmentation": start with any model that predicts a local speed profile given a portion of the path, define a segmentation rule which somehow select segment boundaries given the complete path, and apply the model to each segment. Since the local speed profiles generated by this procedure will have to be scaled appropriately, one has to measure the actual time it takes the subject to traverse each segment of the path, and use those times in fitting the model.

Since the hypothetical internal representation is not directly observable, and cannot be derived from the task description (otherwise the model would belong to the first category) it has to be inferred from the observed trajectory itself. Thus the above models are validated by extracting their parameters from the experimental data, and demonstrating that these parameters are sufficient to reconstruct the same data. Strictly speaking, this procedure shows that the family of trajectories produced by the model contains the family of human arm trajectories; in itself it is not a proof that the proposed internal representation is actually being used by the motor system. To support the latter claim and make detailed comparisons to models in the other two categories, we need a way of quantifying how much of the information present in the data is being extracted in fitting the parameters of the model. This may appear to be a problem with the power law as well, since it predicts a speed profile given a path extracted from the observed trajectory. The crucial difference is that the path of a trajectory is not a priori related to the speed profile, while in the above examples the parameters extracted from the observed trajectory contain information about all aspects of the trajectory (i.e. both spatial and temporal information).

6.4 A constrained minimum jerk model

The main advantage of the power law is that in principle it can be applied to almost any movement: given an arbitrary path it always predicts a speed profile. The problems with it are related to the particular formula used to predict speed from path. In contrast, all known problems with the original minimum-jerk model are in the path prediction: it can fail to predict the hand path given the configuration of via-points, but there are no known cases where the path was predicted accurately while the speed profile was not.

Our model combines the appealing features of the two models: it preserves and extends the generality of the power law, while replacing the particular relationship between curvature and speed with a smoothness constraint. We propose that for an observed path of the hand in space, the speed profile along that path will be the one that minimizes jerk (in general "smoothness" is synonymous to "having small high-order derivatives"; the measure of smoothness we use is one of many possibilities). This constrained principle, like the power law, quantifies an intrinsic relationship between the path and speed of the observed trajectory, rather than suggesting a concrete mechanism for trajectory formation (i.e. we do not suggest that speed is being computed explicitly as a function of hand path). Note that we use the path observed on a particular trial to predict the speed profile on that same trial, which is a natural way to model the significant variability in speed profiles in a block of identical trials (see below). In contrast, the original minimum-jerk model predicts a single optimal movement for a given configuration of via-points, and thus can only be applied to average data.

6.4.1 Formal definition

Let $\mathbf{r}(s) = (x(s), y(s), z(s))$ be a $3D$ curve, describing the path of the hand during a particular trial, where s is the curvilinear coordinate (distance along the path), and the tangential speed is $\dot{s}(t)$ ($\dot{s}$ is a time derivative, $\mathbf{r}'$ is the derivative with respect to s, and boldface signifies vector quantities). Our model assumes that the temporal profile of the movement is the scalar function $s(t)$ that minimizes:

$$J = \int_0^T \left\| \frac{d^3}{dt^3}\mathbf{r}(s(t)) \right\|^2 dt \tag{6.1}$$

This differs from the original formulation of the minimum jerk hypothesis in that here the path $\mathbf{r}$ is given, and we are only minimizing over the space of speed profiles. Flash and Hogan (1985) computed both the path and the speed given a set of via points. Let us denote the term inside the integral by L:

$$L = \left\| \frac{d^3}{dt^3}\mathbf{r}(s) \right\|^2 = \left\| \frac{d^2}{dt^2}\mathbf{r}'(s)\dot{s} \right\|^2 = \left\| \frac{d}{dt}\left(\mathbf{r}''(s)\dot{s}^2 + \mathbf{r}'(s)\ddot{s}\right) \right\|^2$$

$$L = \left\| \mathbf{r}'''(s)\dot{s}^3 + 3\mathbf{r}''(s)\dot{s}\ddot{s} + \mathbf{r}'(s)\dddot{s} \right\|^2 \tag{6.2}$$

This minimization problem is invariant to rotations and translation of the path $\mathbf{r}$. Our goal is to transform the above expression into a form that makes that invariance explicit. We know from differential geometry (e.g. Prakash, 1981) that a $3D$ curve can be defined uniquely, up to translations and rotations, by its curvature $\kappa(s)$ and its torsion $\tau(s)$. It is not possible to simply express $\mathbf{r}$ in terms of κ and τ and substitute in (6.2). However, the derivatives of $\mathbf{r}$ can be expressed as functions of the curvature and torsion. The path $\mathbf{r}$ satisfies Frenet's formulas:

$$\mathbf{t}' = \kappa\mathbf{n}\,, \quad \mathbf{n}' = \tau\mathbf{b} - \kappa\mathbf{t}\,, \quad \mathbf{b}' = -\tau\mathbf{n}$$

where $\mathbf{t}(s)$ is the unit tangent, $\mathbf{n}(s)$ is the unit normal, and $\mathbf{b}(s)$ is the unit binormal vector to the curve $\mathbf{r}(s)$ at point s. Using the fact that the derivative of the path with respect to the curvilinear coordinate s is the unit tangent vector, we obtain:

$$\mathbf{r}' = \mathbf{t}\,, \quad \mathbf{r}'' = \kappa\mathbf{n}\,, \quad \mathbf{r}''' = \kappa'\mathbf{n} + \kappa\left(\tau\mathbf{b} - \kappa\mathbf{t}\right)$$

We can now substitute in (6.2):

$$L = \left\| \mathbf{n}\left(\kappa'\dot{s}^3 + 3\kappa\dot{s}\ddot{s}\right) + \mathbf{t}\left(\dddot{s} - \kappa^2\dot{s}^3\right) + \mathbf{b}\left(\dot{s}^3\kappa\tau\right) \right\|^2 \tag{6.3}$$

and since the tangent, normal, and binormal are orthogonal,

$$L = \left(\kappa'\dot{s}^3 + 3\kappa\dot{s}\ddot{s}\right)^2 + \left(\dddot{s} - \kappa^2\dot{s}^3\right)^2 + \left(\dot{s}^3\kappa\tau\right)^2 \tag{6.4}$$

6.4.2 Relation to the power law

Now that we have an explicit relation between curvature and jerk, we analyze the constrained minimum jerk principle to find how it relates to the power law. We define a function (dependent on $s(t)$) that is equal to the term multiplying the torsion τ in (6.4):

$$A_s = \dot{s}^3\kappa(s)$$

Taking derivatives with respect to time on both sides, we obtain:

$$\kappa'(s)\dot{s}^4 + 3\dot{s}^2\ddot{s}\kappa(s) = A'_s\dot{s}$$

$$\kappa'(s)\dot{s}^3 + 3\dot{s}\ddot{s}\kappa(s) = A'_s$$

Note that the term on the left is exactly the term multiplying the normal vector in (6.3). We now substitute in (6.3):

$$L = \left\|\mathbf{n}\left(A'_s\right) + \mathbf{t}\left(\dddot{s} - A_s\kappa\right) + \mathbf{b}\left(A_s\tau\right)\right\|^2 \tag{6.5}$$

$$L = A'^2_s + \left(\dddot{s} - A_s\kappa\right)^2 + A_s^2\tau^2$$

From the definition of A_s, we have:

$$\dot{s}(t) = A_s^{1/3}\kappa^{-1/3}$$

The power law states that $A_s^{1/3} = const$, or $A'_s = 0$. Thus the power law is equivalent to setting the normal component of the instantaneous jerk to 0, and the binormal component proportional to τ, everywhere along the movement.

For $2D$ curves $\mathbf{r}(t) = (x(t), y(t))$, this can be shown more directly. The power law $v(t) = const.\kappa(t)^{-1/3}$ in $2D$ can be written as:

$$(\dot{x}^2 + \dot{y}^2)^{1/2} = const\left(\frac{\sqrt{(\dot{x}\ddot{y} - \dot{y}\ddot{x})^2}}{(\dot{x}^2 + \dot{y}^2)^{3/2}}\right)^{-1/3}$$

$$\dot{x}\ddot{y} - \dot{y}\ddot{x} = const$$

Taking derivatives, and canceling terms, we get:

$$\dot{x}/\dot{y} = \dddot{x}/\dddot{y}\ ,\ \ \dot{\mathbf{r}}||\dddot{\mathbf{r}}$$

The jerk vector points along the tangent, which is orthogonal to the normal; therefore the jerk along the normal is zero.

Before analyzing experimental data, we can ask whether the power law and the constrained minimum jerk model predict similar speed profiles for discrete movements along arbitrary synthetic paths. In Fig 6-2 we have plotted three synthetic paths, for which we have computed the speed profiles predicted by our model (solid) and the power law (dashed). Also shown is the magnitude and direction of the instantaneous jerk vector at several positions along the path. Note that in most cases the normal jerk is close to zero, and the two speed profiles are rather similar (the deviation near the end points is due to the fact that the power law cannot predict movements that start and stop). There are differences, however: in example B where the curvature is almost constant the power law predicts a speed profile which is flatter than the minimum jerk prediction. For most of the paths we have studied, the two predictions are similar in the sense that they go up and down together, but clearly distinguishable (i.e. example A is not typical).

Why is it that the power law, which predicts instantaneous speed only using the local curvature, comes close to solving the global minimization problem (6.1)? One intuitive answer is the following: since the total jerk is a sum of two (in $3D$ of three) nonnegative terms, and the power law minimizes one of them (normal jerk is set to zero), it is an efficient approximation to the solution of (6.1) (suggesting that in $3D$ the difference between the two speed predictions will be larger). It seems however that this is not always the complete answer. In panel D of Fig 6-2 we have plotted the tangential and normal

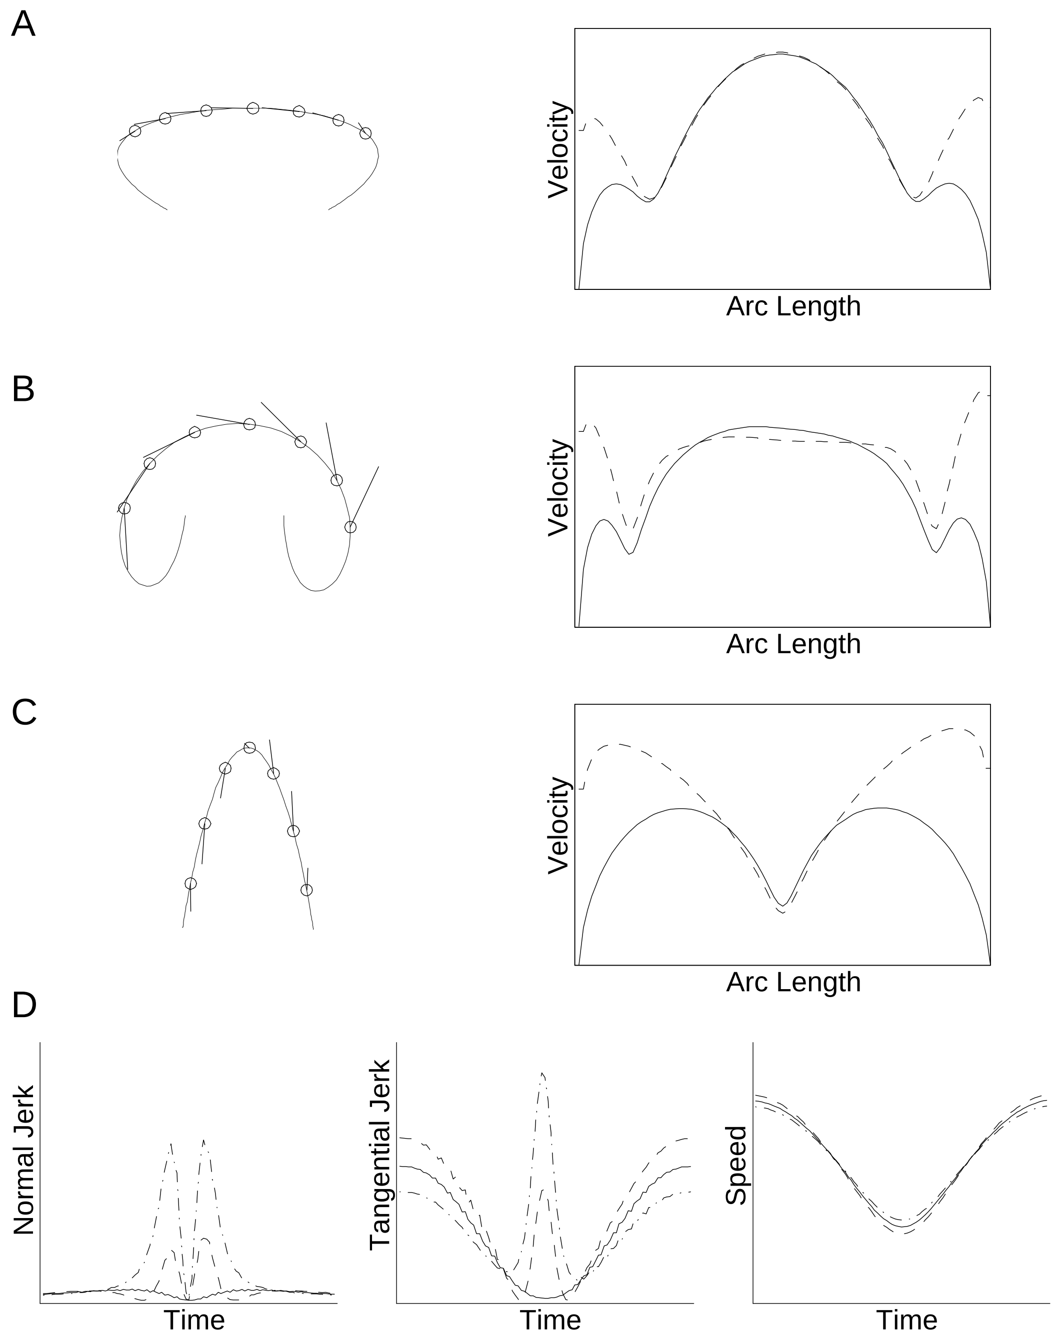

Figure 6-2: A,B,C - Examples of the constrained minimum jerk (solid) and power law (dashed) speed profile predictions for the synthetic paths on the left. The jerk vector is plotted at several points along the path (for the minimum jerk speed profile). D - The tangential and normal jerk components for the middle portion of example C. Two small perturbations are applied to the minimum jerk speed profile (solid), and the corresponding effects on the jerk components are shown.

components of the jerk over the middle portion of example C. We can see that small perturbations around the optimal speed profile increase both terms, implying that in this case the optimal speed profile simultaneously minimizes both terms, and not just their sum (this happens around points of high curvature). In summary, the constrained minimum jerk model and the power law predict similar speed profiles for a given path. There seem to exist a family of paths for which the two are exactly equivalent, and identifying that family may provide further insights, but this is outside the scope of the present paper.

6.5 Experiment 6.1

6.5.1 Methods

To test our model, and compare it to the power law, we studied 4 tasks in which a path was specified and subjects were asked to execute an arm movement along that path. We used discrete movements of short duration and paths of varying degree of complexity. To avoid drawing conclusions from an isolated movement, we used several paths in each task. The same 8 right handed subjects participated in all 4 tasks. Short breaks were given between tasks. The entire experiment lasted about 45 minutes. End-point position was recorded at 100 Hz with an Optotrak 3020 infrared system.

Task 1 - hand/finger movements

Subjects stood in front of a horizontal table, to which we had taped a white sheet of paper (letter format) with the templates shown in Fig 6-3 printed on it. Subjects held a pointer in their right hand; the task was to position the pointer near the circle on the current template, wait to hear a beep signaling that the tip of the pen was in the starting area, and then trace the specified path. After 10 trials a beep of different pitch signaled the subject to move to the next template. No instructions were given regarding the duration or speed of the movement. We observed that subjects used predominantly hand and finger movements in this task.

Task 2 - $2D$ arm movements

Subjects sat on an adjustable chair and made arm movements on a table positioned slightly below shoulder height, so that the movements (made by sliding the arm on the table) occurred in a horizontal plane at shoulder height. We used a VGA projector to display computer generated images on a tilted screen, which the subject viewed in a semitransparent mirror. The system was calibrated so that the projected images appeared to be on the table (see Wolpert et al., 1995 for details). Subjects held a small cursor (tracked by the infrared cameras), which had an LED attached to it that provided endpoint visual feedback during the movement. The room was dark so subjects could not see their arm. The same eight templates from task 1, scaled up by a factor of 8, were projected one at a time. The task was again to position the cursor at the starting point of the template, wait for a beep confirming a correct initial position, and then trace the path. Each path was traced ten times. After each trial we displayed a positional error score, which was the maximum deviation (in millimeters) between the desired and actual paths. Again, subjects were allowed to make movements of any speed. This task involved predominantly shoulder and elbow movements.

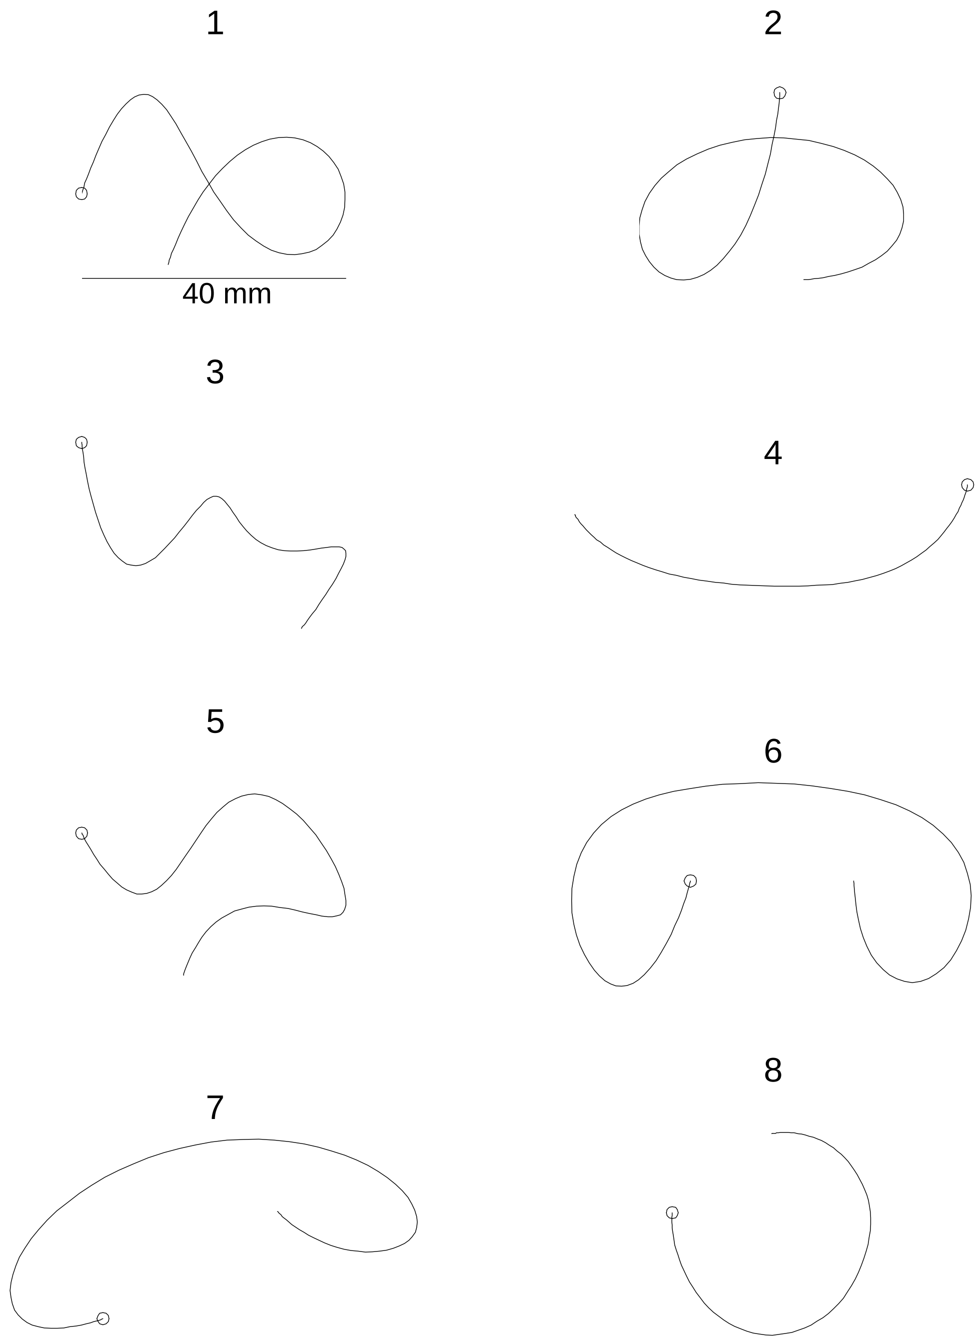

Figure 6-3: Templates used in tasks 1 and 2, drawn to scale. The circle marks the starting point on each template. The order of presentation is row-first.

Task 3 - via points

We used the same setup as in task 2. The display contained a start point, an end point, and four intermediate points, numbered 1 through 4. The subject was asked, after moving to the start point and hearing a confirmation beep, to make a movement that passed through the intermediate points in the specified order (without stopping), and finished at the end point. The configurations we used are shown in Fig 6-4. There were a total of 6 configurations, every other configuration being a 90 degree rotation of the preceding one. Subjects executed 10 trials in each block. As before, movement speed was not specified. Note that the configurations used do not correspond to the continuous paths from the previous two tasks.

Task 4 - $3D$ arm movements

Subjects stood in front of the table, holding the cursor in their right hand, and a $3D$ wireframe model in their left hand. The models were approximately 10x10x10 cm, and represented smooth $3D$ curves. The task was to first study the model, and then try to make a movement with the right hand that had the same shape (not necessarily the same size). The starting point on the model was marked. There was no specified starting position for the arm, since the movement could not be superimposed on the wireframe model. At the beginning of each trial subjects had to stop moving the cursor and wait for a beep. The experimenter waited for the subject to reposition, and then manually activated a procedure that checked for zero speed prior to movement onset. Each wireframe model was used in 10 trials, there were a total of three models. Movement speed and duration were not specified.

6.5.2 Data analysis

The data from each trial for each subject and task were analyzed separately, i.e. no averaging over trials was performed. The first step was to smooth the raw data. We used cubic spline smoothing (De Boor, 1978) applied to $x(t)$, $y(t)$, and $z(t)$. The parameter λ (which determines the amount of smoothing) was set adaptively for each trial, so that the maximum deviation between any smoothed sample point and the corresponding raw data point was within a pre specified radius. That radius was set to 0.002 of the extent of the path - the high accuracy of the Optotrak system obviated the need for substantial smoothing. Once we obtained the cubic spline representation of the data for each trial, all subsequent steps (i.e. extracting speed and curvature) were applied directly, without any further filtering.

We then computed a speed profile prediction from the original power law, the modified power law, and the constrained minimum jerk method (see below). To compute the difference between each prediction and the actual speed profile, we expressed all speed profiles as a function of arc length s and scaled them so that the average speed was one and the path lengths were equal. Scaling of path length was necessary, since subjects did not follow the specified templates exactly, thus the paths on different trials had slightly different length. Alignment along s does not affect the power law predictions which are local; it is necessary for our model, since it is a global method and prediction errors early in the movement accumulate and cause misalignment towards the end if speed is expressed as a function of time.

We defined four measures of deviation between two speed profiles $v1(s)$ and $v2(s)$ that are already scaled to have unit average speed; $\langle . \rangle$ signifies averaging over s, $r(.,.)$ is a correlation coefficient.

Maximum deviation:	$MAX(v1, v2) = max(\|v1(s) - v2(s)\|)$
Mean absolute deviation:	$MAD(v1, v2) = \langle \|v1(s) - v2(s)\| \rangle$
Standard deviation:	$STD(v1, v2) = \sqrt{\langle (v1(s) - v2(s))^2 \rangle}$
Unexplained variance:	$VAR(v1, v2) = 1 - r(v1(s), v2(s))^2$

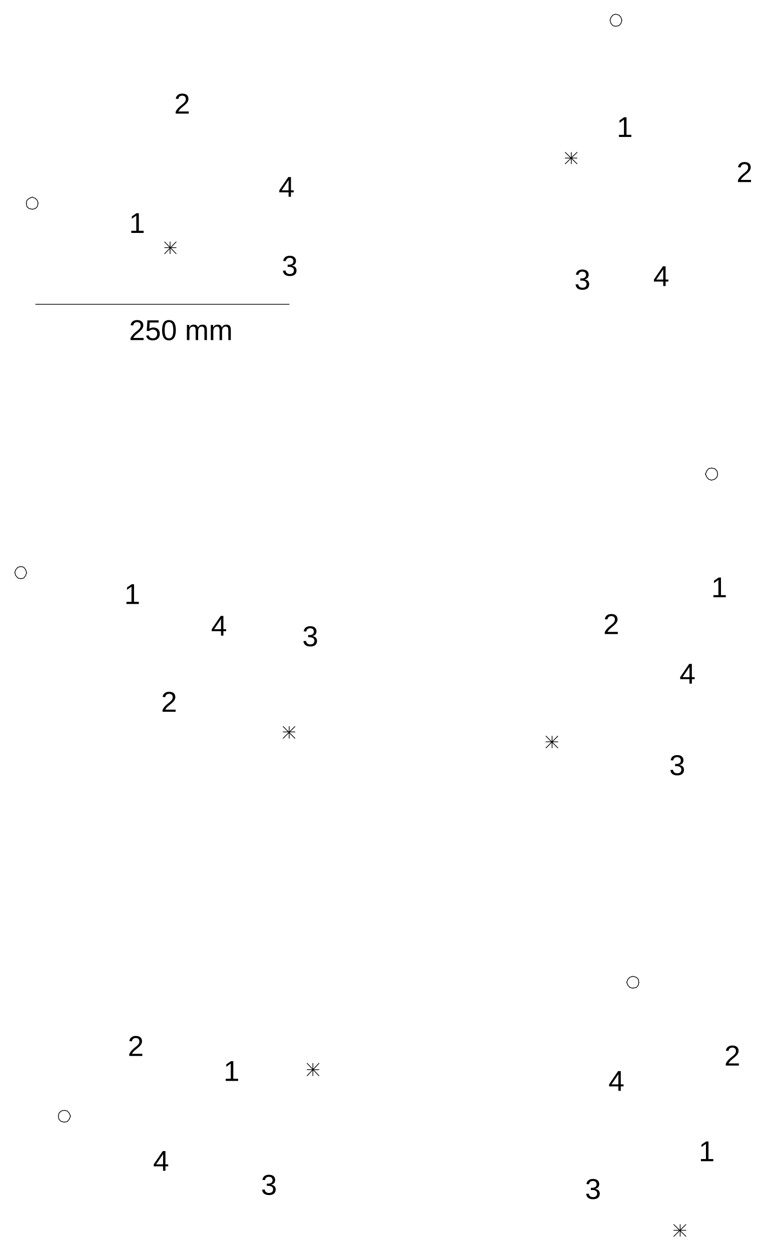

Figure 6-4: Via-point configurations used in task 3, drawn to scale. The circle marks the starting point, the star marks the end point. Every other configuration (right column) is a 90 degree rotation of the preceding one.

The first three quantities indicate deviations between instantaneous speed predictions, in units relative to the average speed. Thus they are sensitive to the amount of fluctuations in the observed speed profiles (i.e. relatively flat speed profiles will be easier to fit according to these measures). If the speed profiles are not scaled to have unit average speed, then these measures will be sensitive to the average speed. The last measure corresponds to the percent variance (fluctuations around the mean) in the observed speed profile not explained by the predicted speed profile, and is affected by difference in shapes rather than scaling factors.

For each model we computed the predictions over the entire movement, and over the middle 60% of the movement (i.e. starting at 20% and ending at 80% along the path). The latter was done in order to assess the accuracy of the power law away from the end points, where it is likely to be inaccurate since it doesn't have a mechanism for enforcing 0 speed at the initial and final points of discrete movements. We examined a number of speed profiles to ensure that the 20% of the movement eliminated at each end was enough to cover the parts of the speed profile that seemed affected by initial acceleration and final deceleration.

The prediction for the original power law was obtained directly from the equation $v(s) = \gamma\kappa(s)^{-1/3}$. For paths with inflection points, we added a small positive constant to the curvature, that was held constant for all trials. Thus the original power law has no free parameters: it takes the path, and the duration of the movement, and predicts a speed profile. The modified power law $v(s) = \gamma(\kappa(s) + \epsilon)^{\beta}$ has two free parameters (ϵ, β). We optimized them (using a nonlinear simplex method) separately for each trial and each measure of deviation defined above. Although in the original formulation ϵ depended on average speed, we estimated it from the data to obtain the most accurate fit the modified power law could provide. The possibility of movement segmentation (Viviani and Cenzato, 1985) is addressed later.

The prediction of the constrained minimum-jerk model was computed using nonlinear optimization methods for variational problems (see below). In its present form, the minimization requires the speed and acceleration at the endpoints (in principle this could be avoided, by minimizing over those parameters as well). For a discrete movement, both the speed and acceleration at the endpoints are zero; since this is know a priori, and not extracted from the experimental data, the model has no free parameters when applied to the entire movement. In order to apply the model to the middle 60% of the movement, we used the actual endpoint speed and acceleration measured experimentally. Note however that we are predicting rather complex speed profiles, thus the two endpoints of a sub movement contain relatively very little information. In contrast, the modified power law uses the entire speed profile to extract the values of its free parameters.

Constrained minimization procedure

Minimizing the integral of $L = (\kappa'\dot{s}^3 + 3\kappa\dot{s}\ddot{s})^2 + (\dddot{s} - \kappa^2\dot{s}^3)^2 + (\dot{s}^3\kappa\tau)^2$ with respect to the temporal profile $s(t)$ is a variational problem, which can be approached using a number of techniques. It turns out that the standard techniques for solving variational problems do not work very well in this case.

From the Euler-Lagrange principle, the solution $s(t)$ satisfies the 6^{th}-order nonlinear ordinary differential equation $\frac{\partial L}{\partial s} - \frac{d}{dt}\frac{\partial L}{\partial s'} + \frac{d^2}{dt^2}\frac{\partial L}{\partial s''} - \frac{d^3}{dt^3}\frac{\partial L}{\partial s'''} = 0$, subject to the constraints on $s(t)$ and its derivatives at the two endpoints of the movement. This problem is very difficult numerically, because the solution depends on high order derivatives of the path, which is recorded experimentally and thus has some noise in it. We found that both shooting and relaxation methods for solving boundary value problems (Press et al., 1992) fail.

Another possibility is to minimize (6.1) directly, by representing $s(t)$ as a linear combination of N fixed basis functions $f_i(t)$: $s(t) = \sum_{i=1}^{N} c_i f_i(t)$, and using gradient descent with respect to the coefficients

c_i. We used sixth-order B-splines as basis functions, and a preconditioned conjugate gradient method to find the optimal coefficients (Gill and Murray, 1981). This minimization always converged, but frequently found local minima, and it was necessary to restart it a number of times before a good solution could be found. An additional problem is that it is not clear what a good solution is.

The method we used here is less direct, but has the advantage that it is faster, and seems to converge to the same minimum regardless of the initial conditions. Rather than using the complete path $\mathbf{r}(s)$, we chose a set of 10 intermediate points, equally spaced along the path (i.e. their positions contain no temporal information). We then found the minimum jerk movement (in the sense of Flash and Hogan, 1985) through those points, given the velocity and acceleration at the two end points. If the resulting minimum jerk movement has a path very close to $\mathbf{r}(s)$, we are guaranteed that its speed profile is the solution to our original problem (i.e. we are computing a minimum over a given set by minimizing over a superset and ensuring that the solution is in the original set). We describe the minimization procedure below.

It has been shown (Flash and Hogan, 1985) that for given passage times T, positions x, velocities v, and accelerations a at the end-points of one segment, the minimum-jerk trajectory is a 5^{th}-order polynomial in t, whose coefficients can be easily determined using the end-point constraints. It is then possible to integrate the squared jerk analytically, and sum it over all segments. The jerk of the optimal trajectory is:

$$
\begin{aligned}
J_{T,x,v,a}(T_{1\ldots N}, x_{1\ldots N}, v_{1\ldots N}, a_{1\ldots N}) = \sum_{i=1}^{N-1} \\
3\frac{3a_i^2 - 2a_i a_{i+1} + 3a_{i+1}^2}{(T_{i+1}-T_i)^4} + 24\frac{3a_i v_i - 2a_{i+1}v_i + 2a_i v_{i+1} - 3a_{i+1}v_{i+1}}{(T_{i+1}-T_i)^3} \\
+48\frac{4v_i^2 + 7v_i v_{i+1} + 4v_{i+1}^2}{(T_{i+1}-T_i)^2} + 120\frac{(a_i - a_{i+1})(x_i - x_{i+1})}{(T_{i+1}-T_i)^2} \\
+720\frac{(v_i + v_{i+1})(x_i - x_{i+1})}{T_{i+1}-T_i} + 720(x_i - x_{i+1})^2
\end{aligned}
$$

Note that the above expression is quadratic in all v and a, and thus can be minimized with respect to those parameters by setting the gradient to 0 and solving (numerically) the resulting system of linear equations. Using this minimization, for given passage times T (and intermediate point positions x), we can compute the jerk for the optimal trajectory: $J_{T_{1\ldots N},x_{1\ldots N}} = \sum_{x=X,Y,Z} \min_{v_{1\ldots N},a_{1\ldots N}} J_{T,x,v,a}(T_{1\ldots N}, x_{1\ldots N}, v_{1\ldots N}, a_{1\ldots N})$ Finally, the function $J_{T_{1\ldots N},x_{1\ldots N}}$ is minimized over all possible $T_{2\ldots N-1}$ using a nonlinear simplex method. We found that the resulting path is equivalent to $\mathbf{r}(s)$, and the speed profiles obtained in this way were indistinguishable from the speed profiles obtained with the direct method above, when it was run enough times to find a global minimum.

6.5.3 Results

In the first two tasks, for which the complete path was specified and the movement was superimposed on it, subjects were able to follow the specified paths accurately. In the $3D$ task, the majority of subjects were not able to replicate the specified shapes well, and did not converge to a very stable movement strategy in 10 trials. After the experiment subjects reported that they were aware of their poor performance, but did not know how to perform the task better. Although this is an important issue to consider, both the power law and the constrained minimum jerk model predict a speed profile given the actual, not the desired path of the hand, so we can still apply them and assess the accuracy of their predictions. It was also observed that the torsion τ varied rather smoothly, instead of being nonzero only at a discrete set of points as suggested by Morasso (1983) and Soechting and Terzuolo (1987). Deviations from piecewise-planar drawing in 3D have also been observed by Gusis (1995); we do not analyze this discrepancy further since it is not the focus of the present study.

In the via-point task we specified four intermediate points, and subjects were free to choose any path that passed through them in the correct order. Our constrained minimum jerk method cannot make a prediction about the actual path subjects choose in this task, but the original model (Flash and Hogan, 1985) predicts that the path will be the one that minimizes jerk. While this is often the case, we found numerous systematic deviations. Fig 6-5 shows data from six different subjects; each pair of plots represents the last five out of the ten trials (dotted) that the subject made, on the original and the rotated configuration of via-points. The solid line is the minimum jerk path, which remains invariant to rotation. Note that the deviations are not just isolated trials, but stable strategies utilized by the subjects; also, the actual paths depended on the orientation of the template, contrary to the model prediction. While some of these deviations may be the result of arm dynamics affecting the path, the large differences (e.g. 5A) between the paths adopted by different subjects for the same via-point configuration suggest that multiple via-point tasks may have a significant cognitive component, i.e. subjects have to study the configuration and decide what general path they want to take. In fact we observed that on the initial trials in a new configuration, subjects sometimes stopped and corrected their movement, which is unlikely to happen if they have already planned a smooth path that passes through the via-points in the specified order. Whatever the reason for these discrepancies, we conclude that the original minimum-jerk model cannot account for the significant between subject variability on multiple via-point tasks.

We now turn to a quantitative comparison of our model and the power law. Fig 6-6 shows summary error statistics, for all tasks and all error measures, over the complete (left) and middle 60% of the movement (right). For each subject, we computed the median prediction error in each block of 10 trials and then averaged over all templates in a given task. The constrained minimum jerk model performed uniformly better than the modified power law, which in turn was uniformly better than the original power law (the latter obviously has to be the case). Analysis of variance showed that in the majority of the comparisons (the $3D$ task/complete movements on the MAD and STD measures being the exception) the differences are significant at the $p < 0.05$ level. Furthermore, this ordering holds for every subject we tested. The difference is equally large over the complete movement and the middle 60%, therefore it is not explained by the discrepant predictions of the power law at the end points of discrete movements.

To analyze the dependence of the model performance on the template being traced, we averaged the data over subjects and compared the three prediction methods for each template. Fig 6-7 (arm task - complete movements) shows that the ordering holds for all templates we tested, and not just the ones that include inflection points. This was true for all tasks. It is interesting that the constrained minimum jerk model shows a rather uniform level of performance, both over different templates and over different tasks. We expected that there might be an effect of trial number, resulting from learning the pattern in the course of the 10 trials, but there was no evidence of that. It is still possible that in more difficult tasks the performance of the model changes with experience, but there was simply very little learning in our experiments. Indeed the spatial error did not change significantly over trials.

There are some differences between the four error measures we defined. The maximum deviation penalizes the original power law on templates with inflection points, since the predicted speed at 0 curvature is very large (it would be infinite if we didn't add a small constant to the curvature). The mean absolute deviation and the standard deviation are very similar, up to a scaling factor. The unexplained variance seems to be the most sensitive measure - it is the only one that detects significant differences in the behavior of our model on different tasks and templates. Thus it will be the measure used in the

A

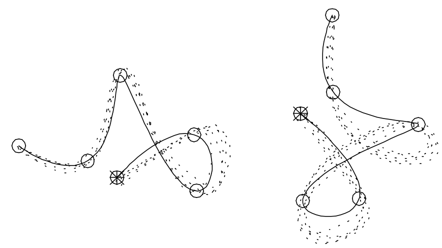

B

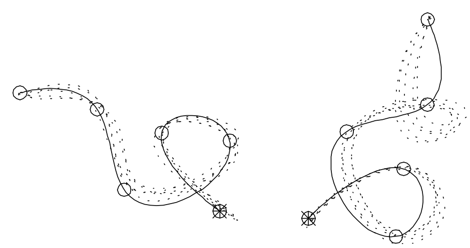

C

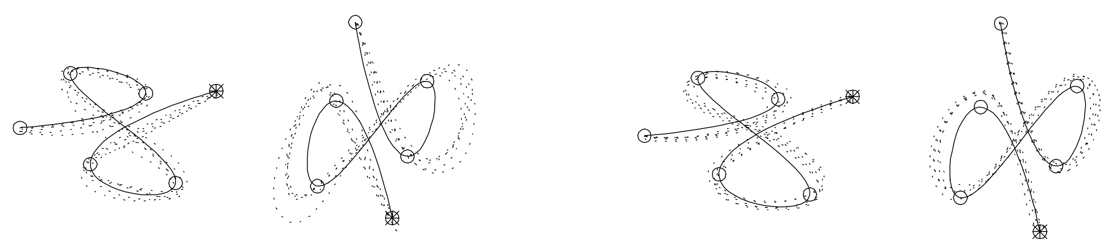

Figure 6-5: Examples of paths traced in task 3. Each pair of plots is data from one subject (six different subjects shown). The solid line is the minimum jerk path (which was not shown to the subjects), and the dotted lines are the actual paths on the last 5 trials of the 10 trial block. The filled circle marks the end point.

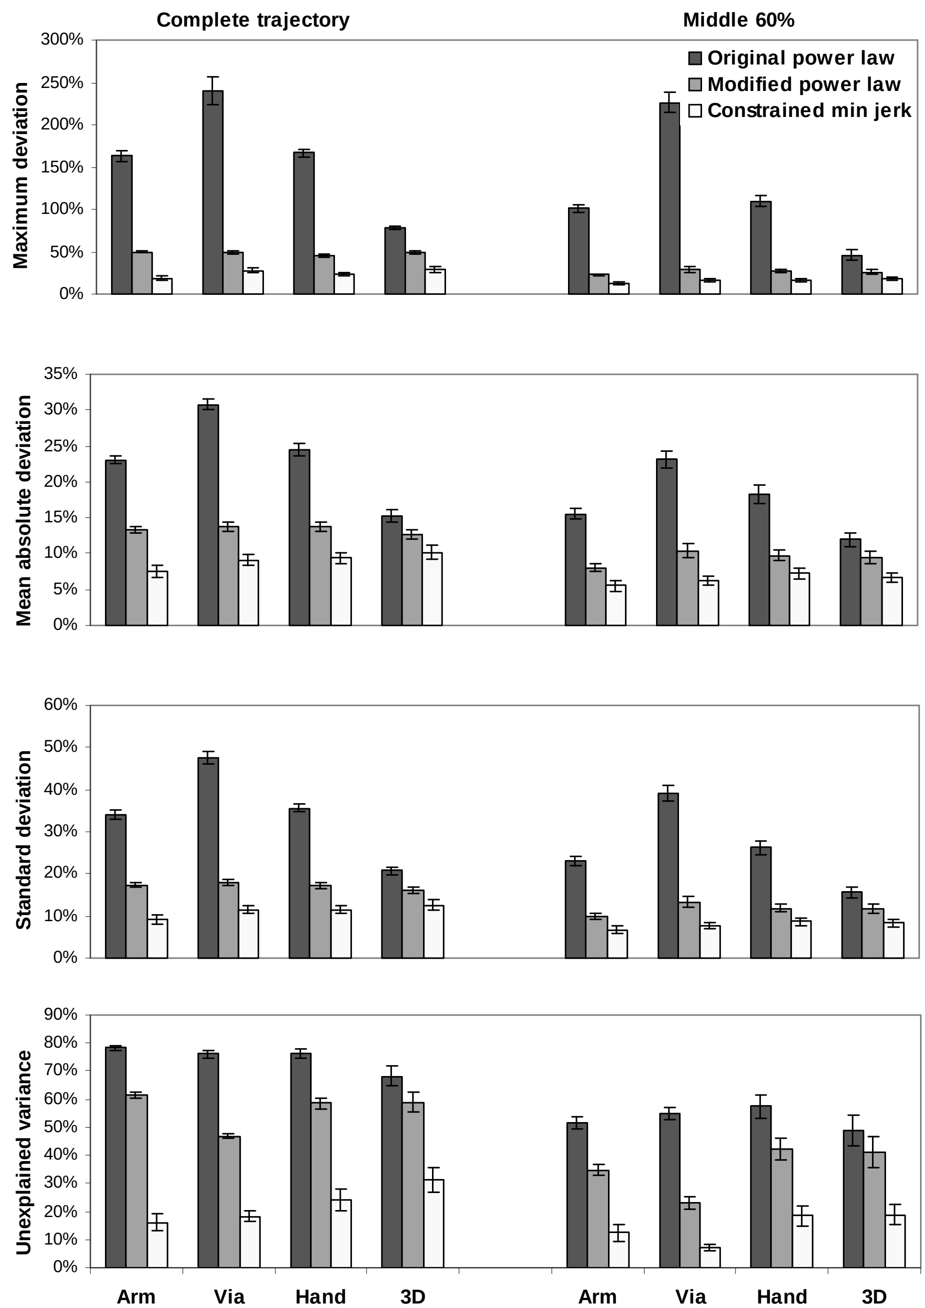

Figure 6-6: Summary error statistics with standard error bars for the three models, computed over the entire movement (left) and the middle 60% (right), for all tasks and error measures. For each subject, the median error score in each block of 10 trials was computed, and the results were averaged over all templates in the task.

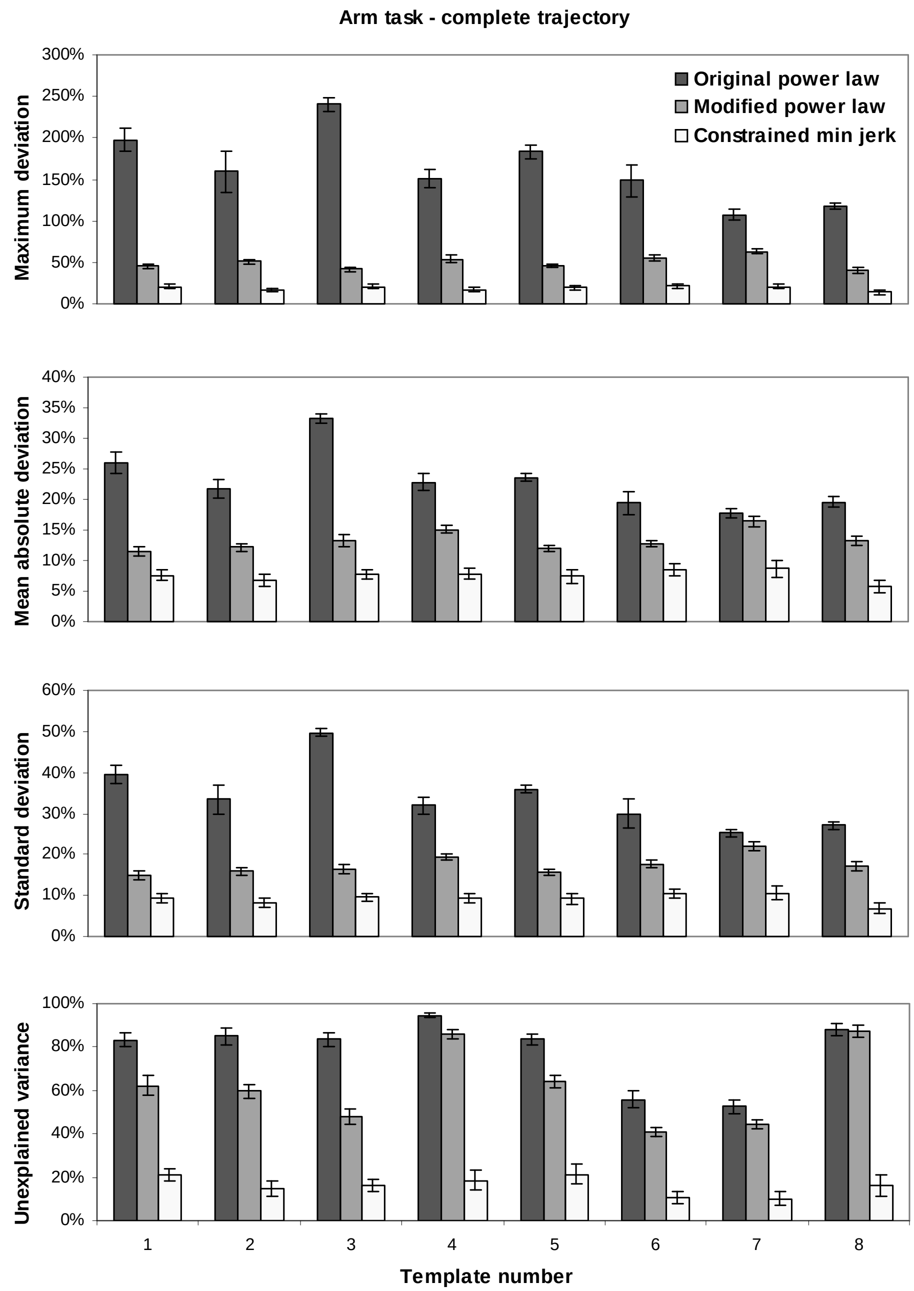

Figure 6-7: Prediction errors for all error measures and all templates in task 2 - complete movements, averaged over subjects.

remaining part of the paper, where we study how the model depends on the details of the task. All subsequent analyses have also been performed with the mean absolute deviation measure, yielding very similar results.

6.5.4 Analysis of speed profiles

Thus far we have considered statistics based on the average prediction errors. In this section we perform a more detailed comparison of the actual speed profiles. Fig 6-8 shows speed profiles of all 8 subjects tracing templates 1,2,3,5 from Fig 6-3. This data is from the next experiment, where a short movement time (1.5 sec) was enforced - these are the conditions under which both our model and the modified power law are most accurate (see below). We have scaled the last 5 speed profiles (complete movements) in each block for each subject, and averaged them. Solid lines are the actual speed profiles, the dashed lines are the constrained minimum jerk predictions, and the dotted lines are the modified power law predictions. Since scaling does not align the speed profiles perfectly, averaging is not necessarily a meaningful operation, but it helps identify systematic prediction errors. Note that we show averages over the predictions for individual trials, not a prediction for the average trial (which would be less meaningful since both prediction methods are nonlinear).

It is evident that the modified power law prediction has significant systematic error, which is not restricted to the end points of the movements. The flat speed profiles it predicts in template 2 correspond to a region of constant curvature (as was noted in Fig 6-2). We conclude that the power law captures very well the maxima and minima of the actual speed profiles, but otherwise it is not a satisfactory model of the details present in the experimental data. While some of the errors are due to skewing (as Wann et al., 1988 argued), the predominant error is in the predicted values at the speed extrema.

The prediction of the constrained minimum jerk method is much more accurate, and most of the remaining error seems to be due to between-subject variability. Still, there is some systematic error left. The first speed maximum on template 1, and the second maximum on template 3 are both slightly underestimated for all subjects. Overall, the predictions of the model are rather close to the experimentally observed speed profiles, and in some cases (last 2 subjects in template 2 for example) it is hard to imagine that any model would be more accurate.

Next we examine data from individual trials, since the averaging procedure is likely to obscure some details. For templates 1,2,3,5,8, we pooled all trials together, and found the one for which both error scores were as close as possible to the corresponding medians. The actual (solid), modified power law (dotted), and constrained minimum jerk (dashed) speed profiles for the selected trials (middle 60%) are shown in Fig 6-9. We can see that there are large errors in the power law prediction that are not related to end point acceleration (deceleration). Note the unnatural sharp peaks, corresponding to inflection points (they were smoothed out in Fig 6-8) - although the modified power law adds a positive constant to the curvature (optimized for the particular trial), this mechanism is apparently not sufficient to deal with inflection points. We might expect that if the paths included straight segments, the performance of the modified power law would be even worse. The inflection points are not the only problem however–note the deviations in the top right panel, which correspond to a path with strictly positive curvature.

6.5.5 Tempo fluctuations and segmentation

Both the constrained minimum-jerk model and the modified power law assume that the tempo, or gain, or "psychological speed" of the movement is held constant, and the variation in speed is only related to the path. Clearly that doesn't have to be the case; e.g. we can ask the subjects to speed up towards the end of the movement. It is conceivable that even in the absence of explicit instructions to control

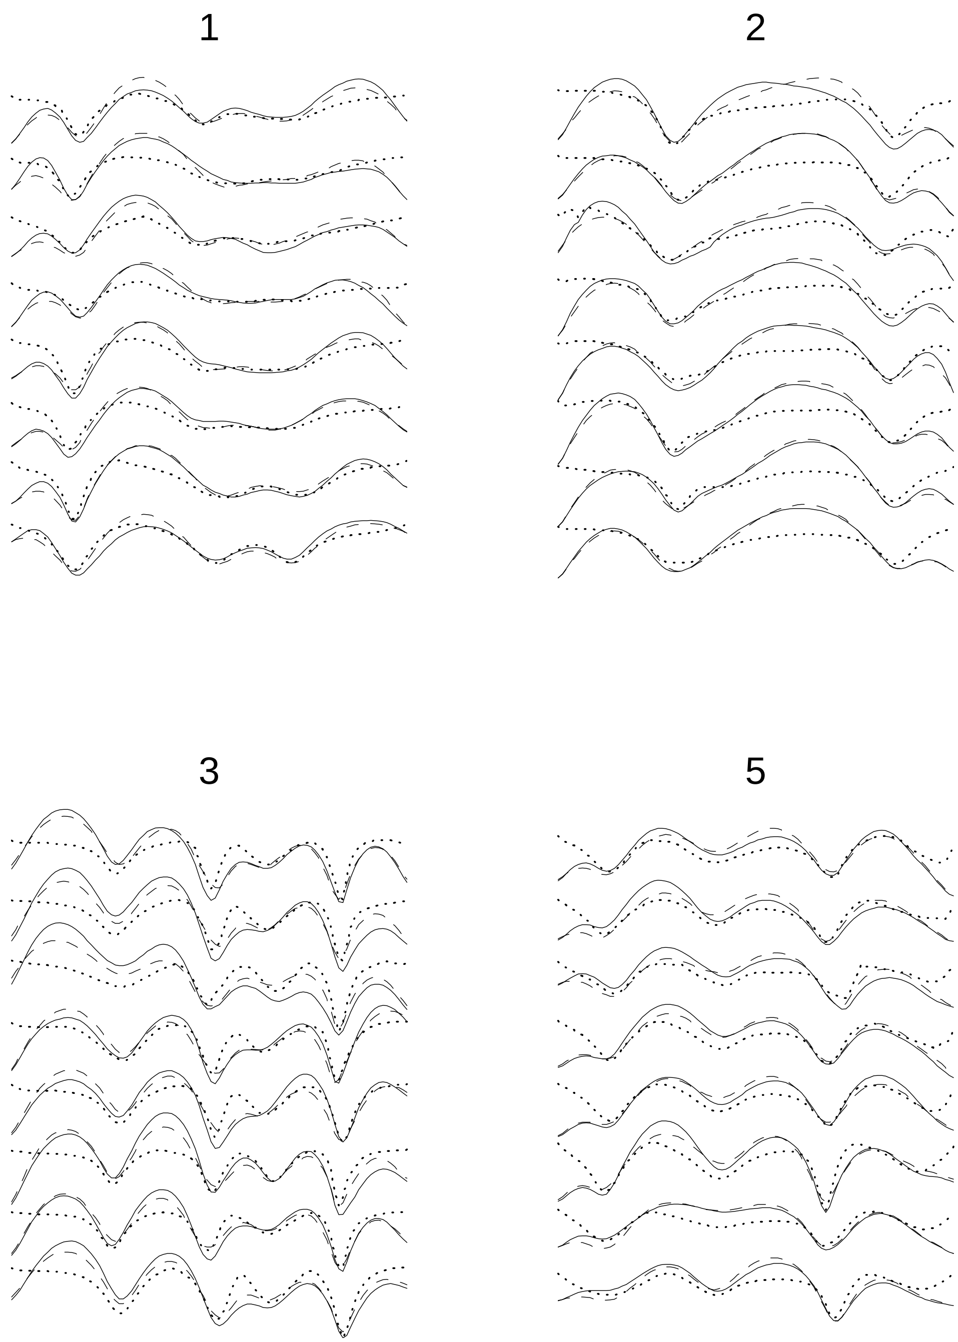

Figure 6-8: Average speed profiles for all subjects, small 1.5 seconds movement in experiment 2 (4 of the 5 templates shown, numbers correspond to template positions in figure 3). Actual speed profile - solid, constrained minimum jerk prediction - dashed, power law - dotted. All speed profiles were scaled before averaging.

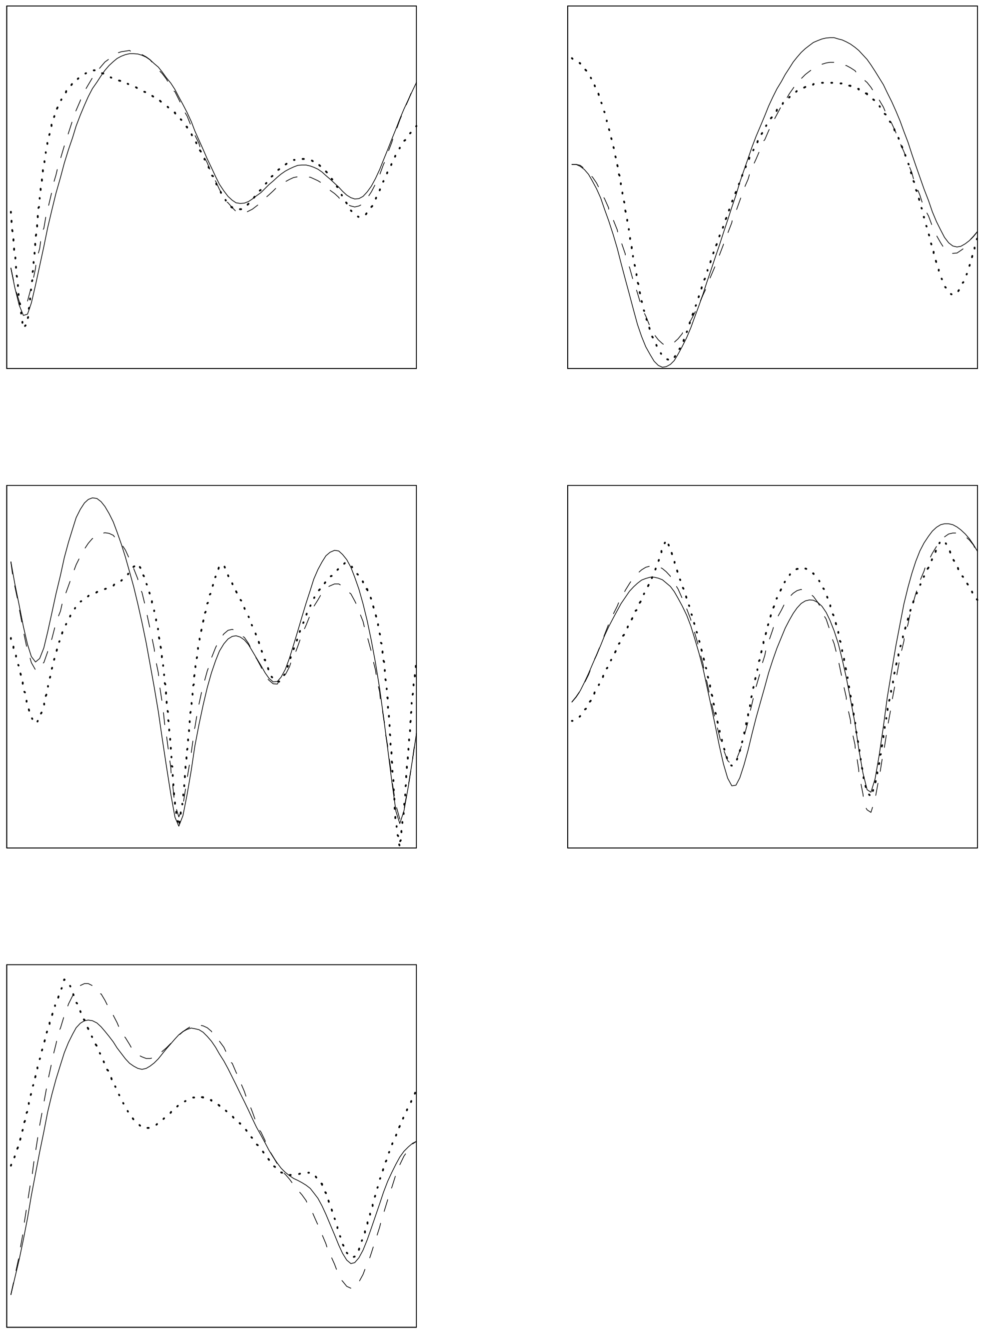

Figure 6-9: Typical examples of single trial speed profiles (solid), constrained minimum jerk prediction (dashed) and power law (dotted) on single trials - experiment 2, middle 60% of the movement.

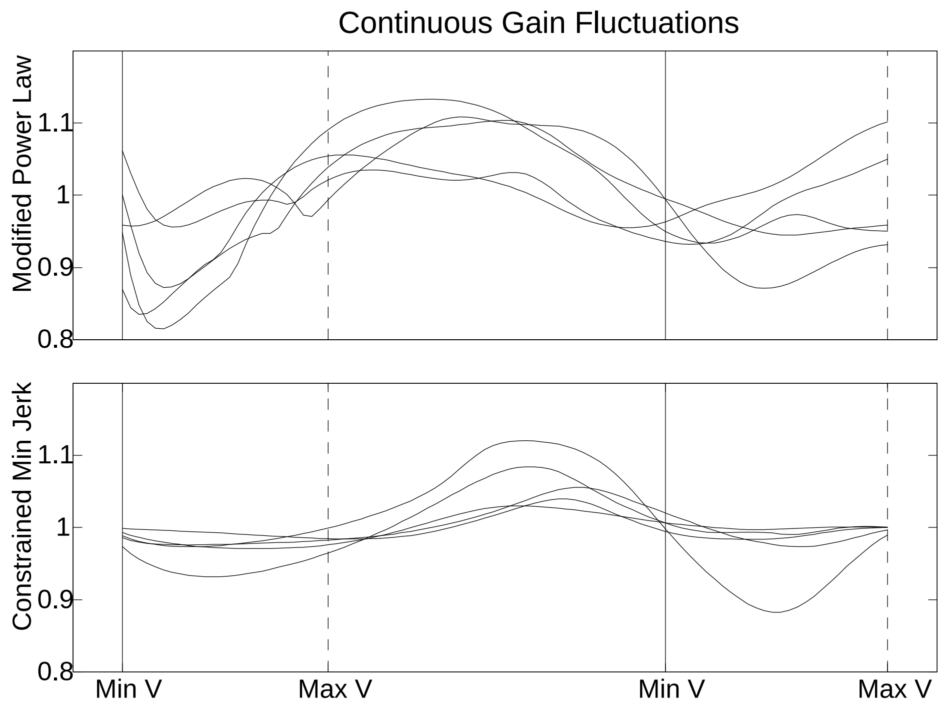

Figure 6-10: Gain fluctuations for the modified power law and our model (consecutive movements of one subject tracing template 1). Vertical lines correspond to points of maximum and minimum speed.

speed, the tempo of the movement is variable. For each model, we define the gain (or tempo) as the ratio between the actual and predicted speed profiles at each point in time (the definition is model-dependent by necessity). Viviani and Cenzato (1985) suggested that the gain is a piecewise-constant function of time: in each segment it is proportional to a power function of the corresponding path length: $\gamma = KP^{\alpha}$ (α here corresponds to the parameter γ defined by Viviani and Cenzato).

We computed the gain factor for a number of discrete movements for both models, and did not find any case in which it was convincingly piece-wise constant. Fig 6-10 shows a typical plot–five consecutive trials for one subject tracing template 1. We then applied the segmented version of the modified power law to the data from all subjects on the arm task - subset of templates used in Experiment 2 (see below). Each movement was segmented at the speed maxima (corresponding to curvature minima). The portions between the end points and the closest segment boundary were eliminated. We applied the modified power law with segmentation with either fixed $\alpha = 0.25$ (which was found to be the optimal fixed value) or α being optimized for each individual trial of each subject (making the comparison to our model somewhat unfair).

Fig 6-11 shows the speed prediction error, relative to the prediction error of the non-segmented modified power law, for the 2 different segmentation methods. The improvement in all cases is much smaller compared to the improvement of the (non-segmented) constrained minimum-jerk method over the modified power law. This segmentation method was also applied at speed minima (corresponding to curvature maxima - as in the model of Morasso and Mussa-Ivaldi, 1983) yielding very similar results. Note that the optimal fixed value of $\alpha = 0.25$ is significantly smaller than the value of $\alpha = 0.6$ observed by Viviani and Cenzato (1985). Since $\alpha = 0$ corresponds to a lack of segmentation, this is another indication

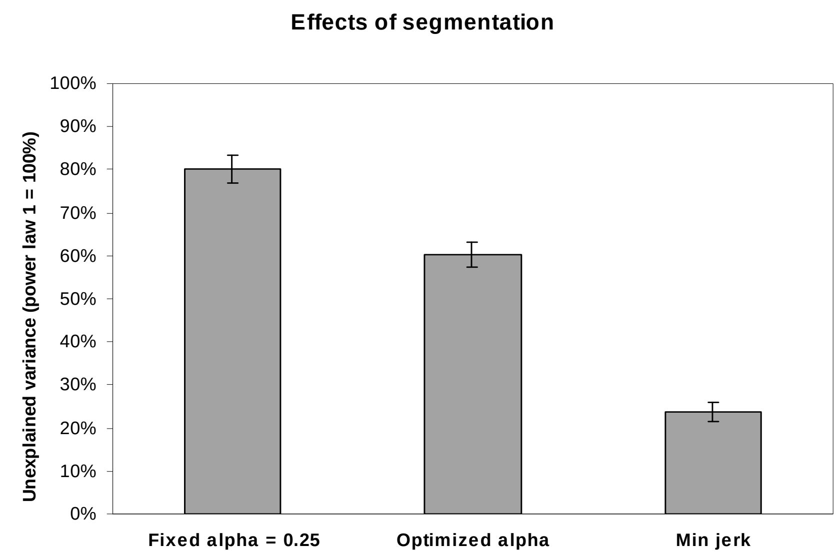

Figure 6-11: For each segmentation method (fixed $\alpha = 0.25$ or α optimized separately for each trial) and our model, we show the prediction error relative to the prediction error of the unsegmented modified power law, applied over the same portion of the path as the one selected by the segmentation scheme.

that segmentation is simply not present in our data, and the tempo fluctuations are continuous and at present unpredictable. The fluctuations in Fig 6-10 are systematic since the trials are from the same subject, but they become unpredictable across subjects (as can be inferred from Figure 6-8).

While the details of the gain fluctuations cannot be explained, their magnitude (directly related to prediction error) could vary systematically with some other variable. We already saw that it is not very sensitive to the task, or the template being traced. A comparison of prediction error over subjects however revealed a strong correlation between how rapidly subjects moved on the average, and how well the constrained minimum jerk model predicted their speed profiles (a weaker correlation was observed for the modified power law). Fig 6-12 shows a scatter plot of mean movement duration (over all templates) and mean prediction error for our model in all tasks. There are several possible interpretations of this result. It may be that the tempo fluctuations accumulate with movement time, thus movements that take longer have less predictable speed profiles. It is also possible that at high average speed, movements are for some reason more regular. Since we did not vary the size of the templates systematically, duration was correlated with speed (Wann et al. 1988 also observed that the power law holds better for faster movements, and attributed the effect to speed instead of duration). Another possibility is that the effect is a spurious between-subject correlation, i.e. some subjects naturally produce more predictable speed profiles, and the same subjects tend to make faster movements when duration is not specified. The next experiment was designed to distinguish between these possibilities.

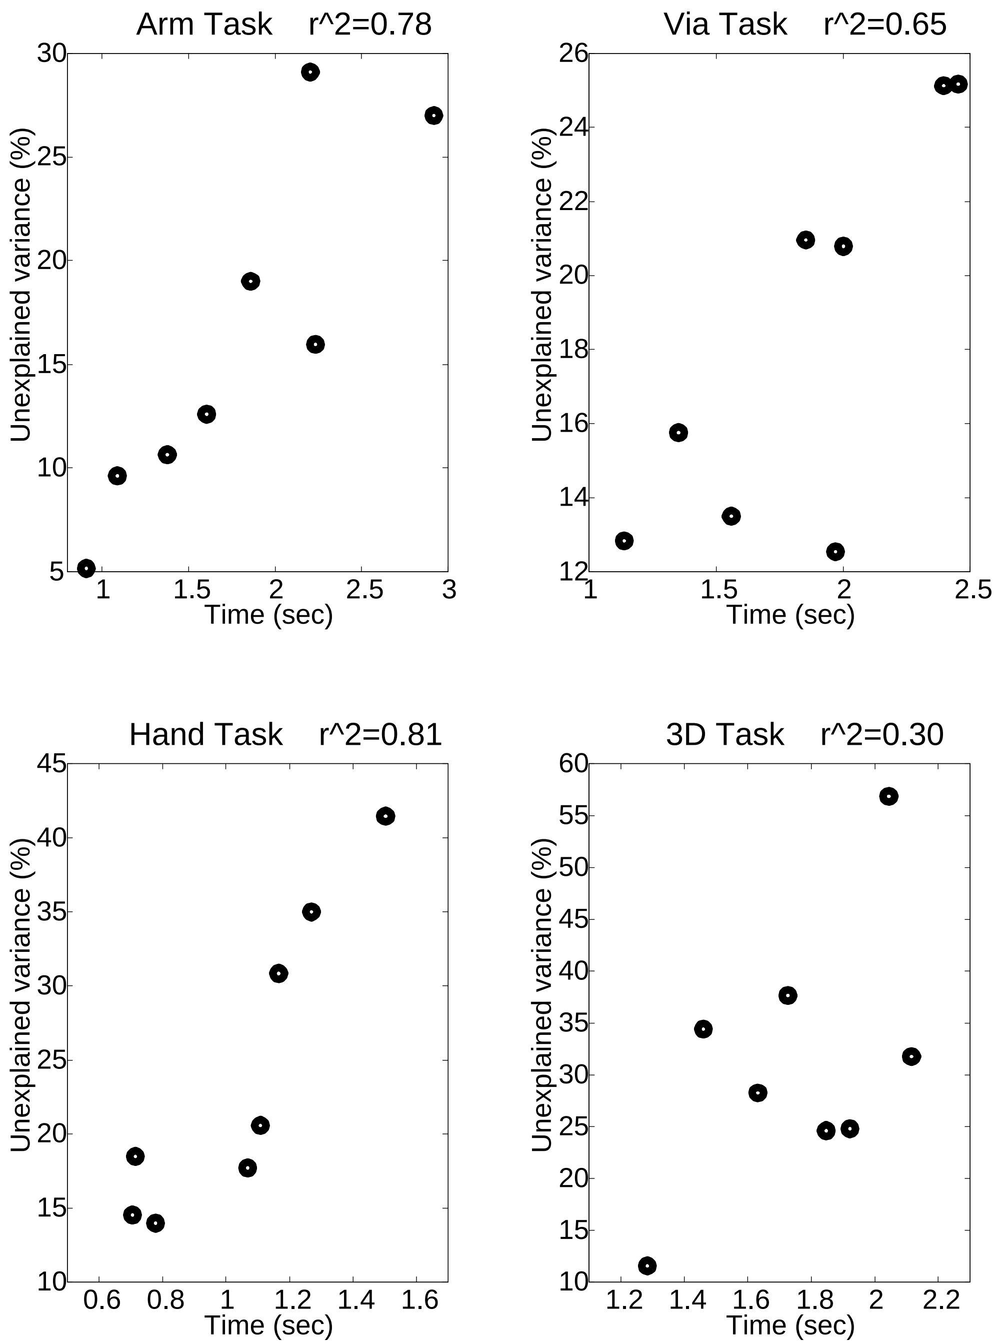

Figure 6-12: Scatterplot of movement duration vs. prediction error for the constrained minimum jerk model, all tasks. Each data point represents one subject, the results were averaged over all trials the subject made in the corresponding task.

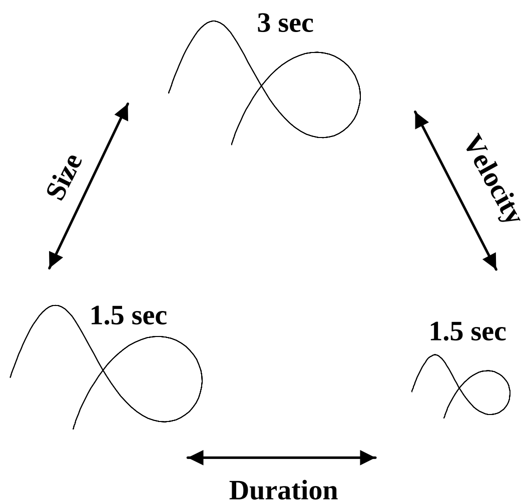

Figure 6-13: Schematic representation of the 3 conditions in the experiment.

6.6 Experiment 6.2

To eliminate the possibility that between-subject variability can account for this correlation, the experiment had a within-subject design. We now specified the duration of the movement, and also varied the spatial scale, which allowed us to distinguish between the effects of movement duration and average speed (and size). A new group of 8 subjects was recruited for this experiment.

6.6.1 Methods

The setup was the same as in task 2 - subjects were asked to make arm movements on a horizontal table, by tracing a continuous curve projected on the table. To reduce experimental time, we used an arbitrary subset of 5 templates (1,2,3,5,8 in Fig 6-3), each presented in 3 blocks of 10 identical trials. In block 1, the template was large (same size as in experiment 1), and desired movement duration was 3 seconds. In block 2, the template was also large, and duration was 1.5 seconds. In block 3, the template was scaled down by a factor of 2, and desired duration was 1.5 seconds. Thus, blocks 1 and 2 were matched for size, blocks 2 and 3 for movement duration, and blocks 1 and 3 for average speed (assuming of course that subjects produced movements of the specified durations, which they did) - see Fig 6-13. At the beginning of each trial, the subject positioned the end-point in the starting area and waited for the start circle to disappear, confirming a correct initial position. The computer sounded a beep when the subject started moving and a second beep when the subject stopped moving. After each trial, the actual duration was shown (in units of 1/100 seconds). The desired duration was printed on the screen throughout the experiment. Subjects were told what the numbers meant and were instructed to trace the paths for the specified durations.

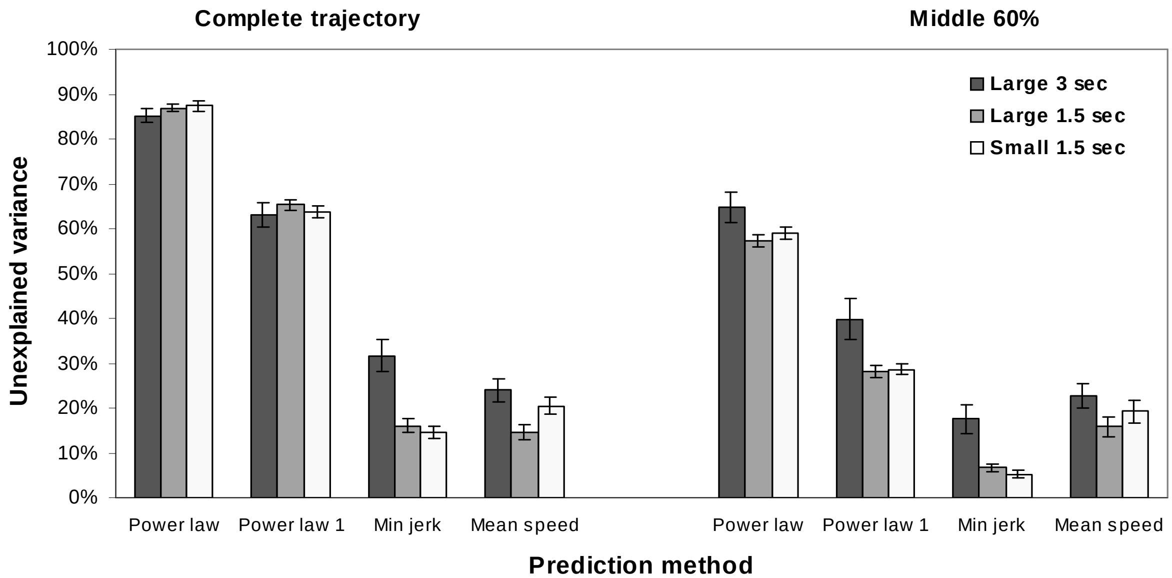

Figure 6-14: Summary error statistics for experiment 2, complete movements (left) and middle 60% (right). The results are computed in the same way as in figure 6. The last prediction method is the average speed profile over a block of 10 consecutive trials, used to predict each individual trial.

6.6.2 Results

The data was analyzed in exactly the same way as in experiment 1. All subjects traced the paths accurately, and were able to adjust the duration of their movements in the first 2-3 trials in each block. Prediction error statistics, averaged over templates, are shown in Fig 6-14. Again, we found that the constrained minimum jerk method predicts speed profiles closer to the ones observed experimentally, compared to both versions of the power law. In this case, the constrained minimum jerk prediction was about three times better than the modified power law, and four times better than the original power law. The difference is still present when speed profiles are computed for the middle 60% of the movements.

It can be seen in Fig 6-14 that the prediction errors for the large fast movements are always similar to the small ones and different from the errors for the large slow movements. Analysis of variance revealed significant differences for the constrained minimum jerk predictions, and the modified power law applied to the middle 60% of the movements; the prediction error for the large 3 seconds movement was larger than the error for both the large 1.5 seconds, and the small 1.5 seconds movement ($p < 0.01$). This result indicates that the accuracy of the prediction depends on the duration, not the average speed or spatial scale of the movement. We can also conclude that the correlation in experiment 1 was not entirely due to systematic between-subject variability, since the second experiment yielded the result within subjects. The higher accuracy of our model when applied to the middle 60% instead of the complete movement may be due to the fact that the model is being applied to a movement of shorter duration. However, this may also be caused by using the experimentally observed values of speed and acceleration at the endpoints of the middle segment.

It is possible that the effect of movement duration on model accuracy is due to the fact that observed speed profiles in short duration movements are somehow different/simpler, and any model would yield better results. To assess this alternative, we have included in Fig 6-14 another prediction method related to the amount of trial-to-trial variability present in a block of identical trials. We computed the mean speed profile over each block of 10 trials (separately for each subject), and used it to predict the speed profile for each of the individual trials. The pattern of prediction errors we see in Fig 6-14 is very different: it appears that trial-to-trial variability decreases with speed rather than duration, and is essentially the same over the entire movement and the middle 60%. Thus the higher model accuracy for shorter durations results from a stronger coupling between path and speed, rather than a difference in the speed profiles themselves.

6.6.3 Trial-to-trial variability

It is often assumed that the motor system constructs a plan of the desired trajectory for the current task, and instantiates that plan on each particular trial. Although there is no widespread agreement on what exactly such a plan might contain, it seems reasonable to assume that the actual trajectories produced on consecutive trials on the same task are variations around the same plan, where the variability is due to unspecified sources of noise. Thus the average trajectory over a block of identical trials corresponds to the movement that would result from the plan in the absence of variability. In support of this notion, it has been observed that the speed profiles of reaching movements of different length and duration can be aligned almost perfectly after straightforward rescaling, and complete models of trajectory formation such as the original minimum jerk model or the minimum torque-change model (Uno et al. 1988) have been applied to the average trajectory.

Under these assumptions, we can attempt to identify the level of processing where the relationship between path and speed emerges. In particular, the smoothness observed in experimental data may be already present in the plan, in which case our model should apply better to the average trajectory than to single-trial trajectories (note that averaging by itself does not necessarily produce smoother trajectories, i.e. if we average several optimal trajectories we will obtain a suboptimal one - because the model is nonlinear). Furthermore, the speed profile predicted from the average path should be more accurate in explaining the single trial speed profiles than the predictions obtained from the single trial paths. In the notation of Fig 6-15 (where $MJ(.)$ is the speed predicted by the model for a given path, $\langle . \rangle$ signifies averaging over a block of trials, and : is the error between actual and predicted speed profiles), $MJ(path) : speed > MJ(\langle path \rangle) : speed > MJ(\langle path \rangle) : \langle speed \rangle$. Alternatively, if smoothness is inherent in the (unknown) sources of variability, we should expect the single trial predictions to be more accurate than both average methods: $MJ(path) : speed < MJ(\langle path \rangle) : speed < MJ(\langle path \rangle) : \langle speed \rangle$.

Since this analysis is most likely to yield clear results when the model works best, we focused on the data from all 1.5 sec movements in experiment 2. As Fig 6-15 shows, none of the above orderings hold. It is true that the constrained minimum jerk model describes the average trajectory better than individual trajectories: $MJ(path) : speed > MJ(\langle path \rangle) : \langle speed \rangle$, therefore smoothness is present in the plan and adding variability decreases the planned smoothness. However, the speed profile predicted from the average path explains the individual speed profiles a lot worse than the single trial predictions: $MJ(path) : speed < MJ(\langle path \rangle) : speed$. In fact, even the average speed profile (which is the best possible constant predictor) performs worse than the single trial prediction method: $MJ(path) : speed < \langle speed \rangle : speed$. Therefore, the trial-to-trial variability observed in speed profiles is partially predictable given the trial-to-trial variability in the path. Note that speed predictions based on the template rather than the average path are a lot less accurate, suggesting that the spatial deviations from the specified template may be present in the plan.

In summary, the picture that emerges from this analysis is the following: the relationship between

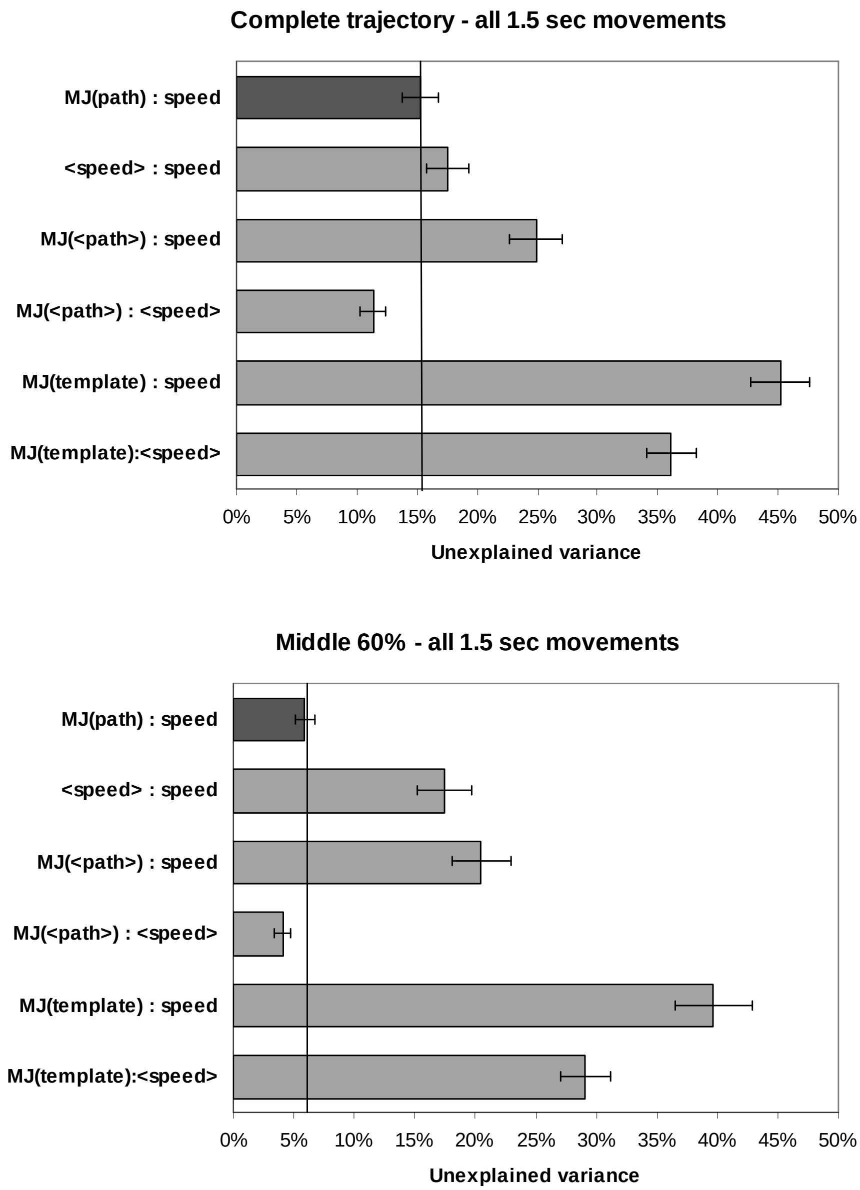

Figure 6-15: Comparison of prediction errors of our model (on all 1.5 sec movements in experiment 2) and possible reconstructions of a "central plan". $MJ(path)$: $speed$ - our model, applied on a trial-by-trial basis (dark bar). $\langle speed \rangle$: $speed$ - the average speed profile over a block of 10 trials, used to predict each individual speed profile. $MJ(\langle path \rangle)$: $speed$ - the speed profile generated by our model for the average path, and used to predict each individual speed profile. $MJ(\langle path \rangle)$: $\langle speed \rangle$ - the speed profile generated by our model for the average path, and used to predict the average speed profile. $MJ(template)$: $\langle speed \rangle$ - the speed profile generated by our model for the specified template, and used to predict the average speed profile. $MJ(template)$: $speed$ - the speed profile generated by our model for the specified template, and used to predict each individual speed profile.

speed and path is present in the plan, and is stronger than that observed on individual trials. However, the trial-to-trial variability is not at all "random"; the path and the speed profile represented in the central plan are modified together in a way that preserves the planned smoothness. It is of course possible that a strict separation between planning and execution does not exist, and thus the above analysis yields mixed results.

6.7 Discussion

This paper has presented a constrained minimum-jerk model of the relationship between hand path and speed profile of complex arm movements. We applied the model to a wide range of spatially constrained motor tasks involving discrete movements of short durations and arbitrary paths. The new model proposed here unifies two sets of experimental observations: the characteristic bell-shaped speed profile of straight reaching movements, and the path-speed relationship in extemporaneous movements with complex paths. The speed predictions of our model are equivalent to those of the original minimum-jerk model when the path is straight, and more accurate when the Flash and Hogan (1985) model fails in the path prediction. While our model uses the entire path (i.e. a continuous curve) to predict speed, it is possible instead to represent the path as a large number of intermediate points and possibly tangents at those points - if such a representation contains sufficient information to reconstruct the continuous path through some interpolation method, this becomes equivalent to our model. In practice, for the family of templates studied here we have found that the entire path can be reconstructed from about 10 points sampled at equal distances. This does not imply any internal representation consisting of 10 discrete points - it is simply a consequence of the mathematical fact that curves as smooth as the ones used here can easily be reconstructed through interpolation.

We showed that although mathematically the constrained minimum-jerk model yields predictions that are generally similar to the power law, its performance is significantly better than both the original and modified versions of the power law, for all tasks and movement templates we studied. Furthermore, the model naturally extends the path-speed relationship previously observed in extemporaneous movements to include the acceleration and deceleration phases of discrete movements, which are much more common. Thus, it can be argued that the local speed-curvature relationship expressed by the power law, to the extent that it fits experimental data, can be derived from a more global smoothness constraint that relates the path to the speed profile.

In comparing any two models, we have to look not only at their competence in accounting for experimental data, but also at the model complexity (Occam's razor). We emphasize that model complexity has nothing to do with the complexity of the mathematics involved; instead it is related to how many and what free parameters are used in fitting the data. In that sense, the constrained minimum-jerk model and the original power law are equivalent - they both postulate that the speed profile of a movement is a well defined function of hand path, and do not require any extra parameters. The modified power law has higher model complexity since it involves a free parameter (β) extracted from the observed trajectories. Furthermore, all versions of the power law require extra mechanisms for dealing with regions of zero curvature, and with zero speed at the endpoints of a discrete movement. Thus we argue that the statistical complexity of the model presented here, when applied to discrete movements with inflection points or straight segments, is no greater than any version of the 2/3 power law. We also have to compare the difficulties involved in the possible biological implementation of each model. From that point of view the power law may appear more appealing since it is analytically tractable (i.e. speed can be obtained directly from curvature, without any numerical approximation). If we believe that the biological system actually uses the concrete mathematical formula, this is certainly a big advantage.

However we are dealing with models that only quantify intrinsic relationships between path and speed, instead of specifying a recipe for trajectory formation. For example, it has been suggested that a 2/3 power law may emerge out of the complex nonlinear interactions of a large number of direction-tuned neurons in primary motor cortex (Lukashin and Georgopoulos, 1993) - such an implementation does not use the analytical tractability of the power law, eliminating any advantage it may have over the model presented here.

Contrary to previous investigations (Viviani and Cenzato, 1985) we did not find evidence for segmentation; instead the "tempo" seemed to fluctuate rather smoothly and unpredictably. We think this is due to the different nature of the movements studied here. Previous work on segmentation has focused on repetitive movements of long duration, which implied (at least visually) a rather obvious segmentation pattern. Instead we used movement durations of one-two seconds, with templates that are not obviously composed of a small number of simple curves concatenated together. It is possible that discrete movements of such short duration are not internally segmented, even when they have complex paths. Alternatively, the decomposition into segments may be variable from trial to trial, preventing us from finding any evidence for it.

The high accuracy of the constrained minimum-jerk model allowed us to make two important observations that have not received previous attention:
a) the relationship between path and speed is stronger for movements of shorter duration, and is not affected by spatial scale or average speed (it is also rather uniform across tasks and movement paths);
b) the path of the hand on a particular trial contains significant information about the speed profile on that same trial; this information is lost if we only study the average path over a block of identical trials.

Presently the model is applied to the complete path, yielding a complete speed profile. It is unlikely that the motor system maintains the path-speed relationship globally, especially over movements of very long duration. Our finding that prediction error increases with movement duration (and not average speed or spatial scale) suggests that the "sliding window" over which the model applies best may be rather small, i.e. about 1 sec. A more local relationship between path and speed is also consistent with observations that the 2/3 power law applies in tracking tasks where little advance planning of maximally smooth movements is possible. Thus it is desirable to modify the procedure for fitting the model so that it can be applied to segments of the path, without requiring endpoint speed and acceleration (only duration for that segment). This can be done by minimizing jerk over these parameters as well.

Our examination of trial-to-trial variability suggests that smoothness is centrally planned, and the variability of the hand path and the speed profile are coupled in a way that maintains the planned path-speed relationship. Since the movements we studied are not executed at maximum speed, one likely source of variability is the effect of online corrections. Such online corrections, however, have to be incorporated smoothly in the movement, in a manner that is consistent with the overall path-speed relationship described by the constrained minimum jerk model. Evidence for such error correction mechanisms is provided by the double-step paradigm, where a reaching movement is smoothly corrected to end at a perturbed target location if subjects are unaware of the perturbation (Pellison et al., 1986). It is also possible that low-level correction/stability mechanisms (e.g. some version of equilibrium point control) combined with a source of noise may give rise to variations around the desired trajectory that preserve its intrinsic properties. The desired trajectory itself does not have to be generated by a servo control mechanism, but only the low-level online corrections to it.

Finally, we cannot rule out the possibility that the trial-to-trial variability we studied was actually planned. It is quite difficult to identify the levels of processing on which different regularities emerge, based on observational studies alone. The problem is that once a movement is executed, it is impossible to determine whether the subject had planned to produce exactly that movement or the outcome was significantly affected by online processes, without perturbative experimental manipulations.

Chapter 7

Conclusion

The main point we tried to make in this thesis was that everyday movements are goal-directed, and this fact makes it impossible to understand motor control without explicitly analyzing how the motor task determines the concrete movement trajectories produced by the motor system. While few people would disagree with this general statement, there are not too many examples of studies that address the relationship between movement patterns and the goal they pursue. This is perhaps due to the difficulty of studying biological motor planning (both experimentally and theoretically), and to the limitations of existing recording techniques that can only quantify the representations used by the CNS and the kinematic patterns of the resulting movements, but not the underlying computations. Intellectual justification for focusing on the execution of already formed plans while largely disregarding the plan formation process[1] is provided by the "motor program view", which essentially postulates that such a simplification is reasonable. Assuming a strict separation between planning and execution, we can indeed hope to understand the building blocks or primitives that motor plans are made of and the way they are executed, independent of the environmental goals they pursue.

But what if, as we claimed in Chapter 2, this simplification is unreasonable? If we agree that the strict separation between planning and execution is not directly supported by experimental data and instead contradicts a number of important characteristics of biological movement, then we should consider the possibility that "planning" is an integral part of the sensory-motor loop and thus the specification of the motor task can affect all movement details moment-by-moment. In Chapter 2 we argued that the motor system does not need a "magic" set of primitives and corresponding "simple" rules for translating all possible motor tasks into those primitives, without ever having to analyze the task and figure out what movements would accomplish it[2]. Instead the problem of motor control could be solved as an inverse problem (which is what it actually is), by maintaining a representation of what needs to be achieved and a causal model of what movements achieve what goals, and using some powerful computation to invert that.

In Chapter 3 we developed a model of eye-hand movements in via-point tasks (based on stochastic optimal control theory) that illustrates how planning and execution can be integrated into the same

[1] Bernstein emphasized repeatedly that "movement construction" is very complex and we have to go a long way before we understand how it occurs. Although his pioneering work has been highly praised and cited, it has not really been followed in that regard.

[2] It would of course be very nice if such a set of primitives (corresponding to a set of feature detectors in the study of perception) existed and could be found, but the repeated failures to do so in designing artificial perception/control systems seem to indicate otherwise.

system producing optimal movements online. The model was used in Chapter 4 to simulate a number of psychophysical findings on eye-hand coordination, speed/accuracy trade-offs, and hand kinematics in via-point tasks. Independent of the model, these findings (and especially the perturbation data) strongly indicate that the motor system is constantly updating its internal estimates of both the environment and the moving limb, and is ready to modify its "plan" in favor of a different movement that better achieves the goal under the new circumstances. In particular, we found a number of cases where trajectory modifications caused by visual perturbations cannot be explained by any model that is trying to execute a preexisting plan as accurately as possible and uses *currently* available sensory input to correct for detected deviations.

Chapter 5 focused on one aspect of the model - automatic damping compensation - which seems to account for a range of psychophysical and physiological data. We also discussed learning and adaptation in the Bayesian framework of motor control. In Chapter 6 we presented a descriptive model of hand kinematics that captures well the relationship between the path and speed profile of drawing movements. Is in interesting in the present context since that relationship turned out to hold better for individual trials rather than average trajectories, which should not be the case if between-trial variability resulted from noise polluting the central plan.

In summary, we argued that the traditional motor-program view is inadequate for studying everyday motor acts of realistic complexity, and we demonstrated that the alternative Bayesian framework proposed here can be instantiated into a concrete model that captures a number of psychophysical phenomena from simple laboratory tasks (where on first glance a more sophisticated model is not needed). However, the most challenging and interesting aspect of this work - demonstrating that our framework is indeed applicable to realistic motor tasks, involving a high degree of redundancy and nontrivial planning requirements - remains to be done in the future.

Bibliography

[Abbs and Connor 89] Abbs, J.H., Connor, N.P. (1989) Motor coordination for functional human behaviors: perspectives from a speech motor data base. In *Perspectives on the Coordination of Movement*, Wallace ed. North-Holland.

[Abbs and Winstein 90] Abbs, J.H., Winstein, C.J. (1990) Functional contributions of rapid and automatic sensory-based adjustments to motor output. In *Attention and Performance XIII: Motor Representation and Control*, Jeannerod ed. Lawrence Erlbaum Associates Inc, Publishers.

[Adams 71] Adams, J.A. (1971) A closed-loop theory of motor behavior. *Journal of Motor Behavior*, 3: 111-149.

[Aratyunyan et al. 69] Aratyunyan, G.H., Gurfinkel, V.S., Mirskii, M.L. (1969) Investigation of aiming at a target. *Biophysics*, 13: 536-538.

[Atkeson and Hollerbach 85] Atkeson, C.G., Hollerbach, J.M. (1985) Kinematic features of unrestrained vertical arm movements. *Journal of Neuroscience*, 5: 2318-2330.

[Ballard et al. 97] Ballard, D.H., Hayhoe, M.M., Pook, P.K., Rao, R.P.N. (1997) Deictic codes for the embodiment of cognition. *Behavioral and Brain Sciences*, 20(4): 723.

[Benda et al. 97] Benda, B.J., Gandolfo, F., Li, C.-S.R., Tresch, M.C., DiLorenzo, D., Bizzi, E. (1997) Neuronal activities in MI of a macaque monkey during reaching movements in a viscous force field. *Society for Neuroscience Abstracts*, 607.12, 23: 1556.

[Bennett et al. 92] Bennett, D.J., Hollerbach, J.M., Xu Y., Hunter, I.W. (1992) Time-varying stiffness of human elbow joint during cyclic voluntary movement. *Experimental Brain Research*, 88: 433-442.

[Bertsekas 87] Bertsekas, D.P. (1987) *Dynamic Programming: Deterministic and Stochastic Models*. Prentice-Hall N.J.

[Blumberg 97] Blumberg, B.M. (1997) *Old Tricks, New Dogs: Ethology and Interactive Creatures*. Ph.D. Thesis, MIT, Cambridge MA.

[Bizzi et al. 92] Bizzi, E., Hogan, N., Mussa-Ivaldi, F.A., Giszter, S. (1992) Does the nervous system use equilibrium-point control to guide single and multijoint movements. *Behavioral and Brain Sciences*, 15: 603-613.

[Bernstein 67] Bernstein, N.I. (1967) *The Coordination and Regulation of Movements.* Pergamon Press.

[Brisson and Alain 96] Brisson, T.A., Alain, C. (1996) Should common optimal movement patterns be identified as the criterion to be achieved. *Journal of Motor Behavior*, 28(3): 211-223.

[Brooks 86] Brooks, R. (1986) A robust layered control system for a mobile robot. *IEEE Journal of Robotics and Automation*, RA(2)-1.

[Bhushan and Shadmehr 98] Bushan, N., Shadmehr, R. (1998) *Neural Control of Movement*.

[Buneo et al. 97] Buneo, C.A., Soechting, J.F., Flanders, M. (1997) Postural dependence of muscle actions: implications for neural control. *Journal of Neuroscience*, 17(6): 2128-2142.

[Caminiti et al. 90] Caminiti, R., Johnson, P.B., Urbano, A. (1990) Making arm movements within different parts of space: dynamic aspects in the primate motor cortex. *Journal of Neuroscience,* 10(7): 2039-2058.

[Canny 88] Canny, J. (1998) *The Complexity of Robot Motion Planning*. MIT Press, Cambridge MA.

[Crossman and Goodeve 63] Grossman, E.R.F.W., Goodeve, P.J. (1963) *Feedback Control of Arm Movement and Fitts' law*. Experimental Psychology Society, Oxford, England.

[D'Avella and Bizzi 98] D'Avella, A., Bizzi, E. (1998) Low dimensionality of supraspinally induced force fields. *Proceedings of the National Academy of Science*, 95: 7711-7714.

[Davis and Vinter 85] Davis, M.H.A., Vinter R.B. (1985) *Stochastic Modelling and Control: Monographs on Statistics and Applied Probability*. Chapman and Hall.

[De Boor 78] De Boor, C. (1978) *A Practical Guide to Splines.* New-York: Springer-Verlag.

[Desmurget et al. 95] Desmurget, M., Rossetti, Y., Prablanc, C., Stelmach, G., Jeannerod, M. (1995) Representation of hand position prior to movement and movement variability. *Canadian Journal of Physiology and Pharmacology*, 72: 262-272.

[Evarts 68] Evarts, E.V. (1968) Relation of pyramidal tract activity to force exerted during voluntary movement. *Journal of Neurophysiology*, 31: 14-27.

[Evarts 81] Evarts, E.V. (1981) Role of motor cortex in voluntary movements in primates. In: Brooks V.B. (ed) *Handbook of Physiology, Section I Nervous System, Volume 2, Part 2, Motor Control*. Williams and Wilkins, Baltimore, 1083-1120.

[Fetz 92] Fetz, E.E. (1992) Are movement parameters recognizably coded in the activity of single neurons? *Behavioral and Brain Sciences*, 15: 679-690.

[Flash and Hogan 85] Flash, T., Hogan, N. (1985) The coordination of arm movements: an experimentally confirmed mathematical model. *Journal of Neuroscience*, 5(7): 1688-1703.

[Flanders et al. 94] Flanders, M., Pellegrini, J.J., Soechting, J.F. (1994) Spatial/temporal characteristics of a motor pattern for reaching. *Journal of Neurophysiology*, 71(2): 811-813.

[Giszter 94] Giszter, S. (1994). Reinforcement tuning of action synthesis and selection in a virtual frog. In *From Animals to Animats III.* Cliff, Husbands, Meyer and Wilson, eds. MIT Press, Cambridge MA.

[Gribble and Ostry 96] Gribble, P., Ostry, D. (1996) Origins of the power law relationship between movement velocity and curvature: modeling the effects of muscle mechanics and limb dynamics. *Journal of Neurophysiology*, 76(5): 2853-60.

[Gribble et al. 98] Gribble, P.L., Ostry, D.J., Sanguineti, V., Laboissiere, R. (1998) Are complex control signals required for human arm movement? *Journal of Neurophysiology*, 79(3): 1409-1424.

[Gomi and Kawato 96] Gomi, H., Kawato, M. (1996) Equilibrium-point control hypothesis examined by measured arm stiffness during multijoint movement. *Science*, 272: 117-120.

[Gusis 95] Gusis, A. (1995) *Arm Trajectory Planning in 3D Space.* M.Sc. Thesis, Weizmann Institute of Science.

[Gentner 87] Gentner, D.R. (1987) Timing of skilled motor performance: tests of the proportional duration model. *Psychological Review*, 94(2): 255-276.

[Georgopoulos et al. 82] Georgopoulos, A.P., Kalaska, J.F., Caminiti, R., Massey, J.T. (1982) On the relationships between the direction of two-dimensional arm movements and cell discharge in primate motor cortex. *Journal of Neuroscience*, 2(11): 1527-1537.

[Georgopoulos et al. 84] Georgopoulos, A.P., Caminiti, R., Kalaska, J. (1984) Static spatial effects in motor cortex and area 5: quantitative relations in a two-dimensional space. *Experimental Brain Research*, 54: 446-454.

[Ghez et al. 90] Ghez, C., Hening, W., Favilla, M. (1990) Parallel interacting channels in the initiation and specification of motor response features. In *Attention and Performance XIII: Motor Representation and Control*, Jeannerod ed. Lawrence Erlbaum Associates Inc, Publishers.

[Ghez and Sainburg 95] Ghez, C., Sainburg, R. (1995) Proprioceptive control of interjoint coordination. *Canadian Journal of Physiological Pharmacology*, 73(2): 305-315.

[Ghez et al. 97] Ghez, C., Favilla, M., Ghilardi, M.F., Gordon, J., Bermejo, R., Pullman, S. (1997) Discrete and continuous planning of hand movements and isometric force trajectories. *Experimental Brain Research*, 115: 217-233.

[Gordon et al. 94] Gordon, J., Ghilardi, M.F., Ghez, C. (1994) Accuracy of planar reaching movements I. Independence of direction and extent variability. *Experimental Brain Research*, 99: 97-111.

[Gordon et al. 94b] Gordon, J., Ghilardi, M.F., Cooper, S.E., Ghez, C. (1994) Accuracy of planar reaching movements II. Systematic extent errors resulting from inertial anisotropy. *Experimental Brain Research*, 99: 112-130.

[Hening et al. 88] Hening, W., Favilla, M., Ghez, C. (1988) Trajectory control in targeted force impulses. V. Gradual specification of response amplitude. *Experimental Brain Research*, 71: 116-128.

[Hinton 84] Hinton, G. (1984) Some computational solutions to Bernstein's problems. In Whiting, H.T.A (Ed) *Human Motor Actions: Bernstein Reassessed.* North-Holland, 1984.

[Hinton et al. 95] Hinton, G.E., Dayan, P., Frey, B.J., Neal, R. (1995) The wake-sleep algorithm for unsupervised neural networks. *Science*, 268: 1158-1161

[Hogan and Flash 87] Hogan, N., and Flash, T. (1987) Moving gracefully: quantitative theories of motor coordination. *Trends in Neuroscience*, 10(4): 170-174.

[Hoff 94] Hoff, B. (1994) A model of duration in normal and perturbed reaching movement. *Biological Cybernetics*, 71: 481-488.

[Hoff 91] Hoff, B. (1991) Ph.D. Thesis. USC.

[Hollands et al. 95] Hollands, M.A., Marple-Horvat, D.E., Henkes, S., Rowan, A. (1995) Human eye movements during visually guided stepping. *Journal of Motor Behavior*, 27(2): 155-163.

[Hubel and Wiesel 62] Hubel, D.H. and Wiesel, T.N. (1962) Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *Journal of Physiology*, 165: 559-568.

[Jack 1895] Jack, W.R. (1895) On the analysis of voluntary muscular movements by certain new instruments. *Journal of Anatomy and Physiology*, 29: 473-478.

[Jekka and Kelso 89] Jeka, J.J., Kelso, J.A.S (1989) The dynamic pattern approach to coordinated behavior: a tutorial review. In *Perspectives on the Coordination of Movement*, Wallace ed. North-Holland.

[Jeannerod 88] Jeannerod, M. (1988) *The Neural and Behavioral Organization of Goal-Directed Movements.* Oxford Science Publications.

[Johansson and Westling 90] Johansson, R.S., Westling, G. (1990) Tactile afferent signals in the control of precision grip. In *Attention and Performance XIII: Motor Representation and Control*, Jeannerod ed. Lawrence Erlbaum Associates Inc, Publishers.

[Johnson 92] Johnson, P.B. (1992) Toward an understanding of the cerebral cortex and reaching movements: a review of recent approaches. In Caminiti, R., Johnson, P.B., Burnod, Y. (ed) *Control of Arm Movement in Space: Neurophysiological and Computational Approaches*. Berlin; New York: Springer-Verlag, 1992.

[Jordan and Rumelhart 92] Jordan, M.I., Rumelhart, D.E. (1992) Forward models: supervised learning with a distal teacher. *Cognitive Science*, 16: 307-354.

[Kandel et al. 91] Kandel, E.R., Schwartz, J.H., Jessell, T.M. (1991) *Principles of Neural Science*. Appleton & Lange.

[Kalaska et al. 89] Kalaska, J.F., Cohen, D.A., Hyde, M.L., Prud'homme, M. (1989) A comparison of movement direction-related versus load direction-related activity in primate motor cortex, using a two-dimensional reaching task. *Journal of Neuroscience*, 9(6): 2080-2102.

[Kalaska et al. 92] Kalaska, J.F., Crammond, D.J., Cohen, D.A.D., Prud'homme, M., Hyde, M.L. (1992) Comparison of cell discharge in motor, premotor, and parietal cortex during reaching. In Caminiti, R., Johnson, P.B., and Burnod, Y. (ed) *Control of Arm Movement in Space: Neurophysiological and Computational Approaches*. Berlin; New York: Springer-Verlag, 1992.

[Kahneman and Tversky 73] Tversky, Kahneman (1973) On the psychology of prediction. *Psychological Review*, 80: 237-251.

[Kawato 98] Kawato, M. (1998) Talk presented in "Coming to grips with internal models" session, *8th Annual Neural Control of Movement Meeting*.

[Keele et al. 90] Keele, S.W., Cohen, A., Ivry, R. (1990) Motor programs: concepts and issues. In *Attention and Performance XIII: Motor Representation and Control*, Jeannerod ed. Lawrence Erlbaum Associates Inc, Publishers.

[Khatib 87] Khatib, O. (1987) A unified approach to motion and force control of robot manipulators: the operational space formulation. *IEEE Journal of Robotics and Automation*, 3(1): 43-53.

[Krakauer et al. 96] Krakauer, J.W., Pine, Z.M., Ghez, C. (1996) Difference in generalization of adaptation to altered gains and display rotations in reaching movements. *Society for Neuroscience Abstracts*, 354.12.

[Knill and Richards 96] Knill, D., Richards, W. (1996) *Perception as Bayesian Inference*. Cambridge University Press.

[Lacquaniti et al. 83] Lacquaniti, F., Terzuolo, C., Viviani, P. (1983) The law relating the kinematic and figural aspects of drawing movements. *Acta Psychologica*, 54: 115-130.

[Lacquaniti et al. 93] Lacquaniti, F., Carrozzo, M., Borghese, N.A. (1993) Time-varying mechanical behavior of multijointed arm in man. *Journal of Neurophysiology*, 69(5): 1443-1464.

[Latombe 91] Latombe, J.C. (1991) *Robot Motion Planning.* Kluwer Academic Publishers.

[Laszlo 92] Laszlo, J.I. (1992) Motor control and learning: how far do the experimental tasks restrict our theoretical insight? In *Approaches to the Study of Motor Control and Learning*, Summers ed. North-Holland.

[Lukashin and Georgopoulos 93] Lukashin, A.V., Georgopoulos, A.P. (1993) A dynamical neural network model for motor cortical activity during movement: population coding of movement trajectories. *Biological Cybernetics*, 69: 517-524.

[MacKenzie et al. 87] MacKenzie, C.L., Marteniuk, R.G., Dugas, C., Liske, D., Eickmeier, B. (1987) Three-dimensional movement trajectories in Fitts' task: implications for control. *The Quaterly Jounral of Experimental Psychology*, 39A: 629-647.

[MacKenzie and Van Eerd 90] MacKenzie, C.L., Van Eerd, D.L. (1990) Rhythmic precision in the performance of piano scales: motor psychophysics and motor programming. In *Attention and Performance XIII: Motor Representation and Control*, Jeannerod ed. Lawrence Erlbaum Associates Inc, Publishers.

[Maes 91] Maes, P. (1991) *Designing Autonomous Agents*, MIT Elsevier 1991.

[Mason 98] Mason, M. (1998) *Mechanics of Manipulation*, Lecture notes, CMU.

[McIntyre and Bizzi 93] McIntyre, J., Bizzi, E. (1993) Servo hypotheses for the biological control of movement. *Journal of Motor Behavior,* 25(3): 193-202.

[Matsuoka 98] Matsuoka, Y. (1998) *Models of Generalization in Motor Control*. Ph.D. Thesis, MIT, Cambridge MA.

[Marr 82] Marr, D. (1982) *Vision.* W.H. Freeman and Company, New York.

[McMahon 84] McMahon, T. (1984) Reflexes and motor control. In *Muscles, Reflexes and Locomotion*. Princeton University Press, Princeton NJ.

[Meyer et al. 82] Meyer, D.E., Smith, J.E.K., Wright, C.E. (1982) Models for the speed and accuracy of aimed movements. *Psychological Review*, 89: 449-482.

[Milner and Ijaz 90] Milner, T.E., Ijaz, M.M. (1990) The effect of accuracy constraints on three-dimensional movement kinematics. *Neuroscience*, 35(2): 365-374.

[Moran and Schwartz 98] Moran, D.W., Schwartz, A.B. (1998) Motor cotrical representation of speed and direction during reaching. *Jounral of Neurophysiology*, in press.

[Morasso 81] Morasso, P. (1981) Spatial control of arm movements. *Experimental Brain Research*, 42: 223-227.

[Morasso and Mussa-Ivaldi 82] Morasso, P., and Mussa-Ivaldi, F.A. (1982) Trajectory formation and handwriting: a computational model. *Biological Cybernetics*, 45: 131-142.

[Morasso 83] Morasso, P. (1983) Three dimensional arm trajectories. *Biological Cybernetics*, 48: 187-194.

[Mussa-Ivaldi 88] Mussa-Ivaldi, F.A. (1988) Do neurons in the motor cortex encode movement direction? An alternative hypothesis. *Neuroscience Letters*, 91: 106-111.

[Nakayama and Shimojo 92] Nakayama, K., Shimojo, S. (1992) Experiencing and perceiving visual surfaces. *Science*, 257: 1357-1363.

[Nelson 83] Nelson, W.L. (1983) Physical principles for economies of skilled movements. *Biological Cybernetics*, 46: 135-147.

[Newel and Emmerik 89] Newel, K.M., van Emmerik, R.E.A. (1989) The acquisition of coordination: preliminary analysis of learning to write. *Human Movement Science*, 11:113-123.

[Newel et al. 78] Newel, K.M., Kugler, P.N., van Emmerik, R.E.A., MacDonald, P.V. (1989) Search strategies and the acquisition of coordination. In *Perspectives on the Coordination of Movement*, Wallace ed. North-Holland.

[Newel 78] Newel, K.M. (1978) Some issues on action plans. In *Information Processing in Motor Control and Learning.* Stelmach ed. Academic Press.

[Nicols and Houk 73] Nicols, T.R., Houk, J.C. (1973) Reflex compensation for variations in the mechanical properties of a muscle. *Science*, 181:182-184.

[Osherson 90] Osherson, D. (1990) Judgement. In Osherson and Smith, ed. *An Invitation to Cognitive Science. Volume 3: Thinking*. MIT Press, Cambridge MA.

[Pouget and Sejnowski 97] Pouget, A., Sejnowski, T. (1997) Spatial transformations in the parietal cortex using basis functions. *Journal of Cognitive Neuroscience*, 9(2):222-237.

[Pratt 95] Pratt, J.E. (1995) *Virtual Model Control of a Biped Walking Robot.* ME Thesis, MIT, Cambridge MA.

[Pearl 88] Pearl, J. (1988) *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann, San Mateo CA.

[Pellison et al. 86] Pellison, D., Prablanc, C., Goodale, M.A., Jeannerod, M. (1986) Visual control of reaching movements without vision of the limb. II. Evidence of fast unconscious processes correcting the trajectory of the hand to the final position of a double-step stimulus. *Experimental Brain Research*, 62: 303-311.

[Perlin 97] Perlin, K. (1997) *Real Time Responsive Animation with Personality.* Online techical report, Media Research Laboratory, NYU.

[Plamondon and Alimi97] Plamondon, R., Alimi, A.M. Speed/accuracy trade-offs in target-directed movements. *Behavioral and Brain Sciences*, 20: 279-349.

[Pollick and Shapiro 97] Pollick, F.E.,, Sapiro, G. (1997) Constant affine velocity predicts the 1/3 power law of planar motion perception and generation. *Vision Research*, 37(3): 347.

[Posner and Keele 70] Posner, M.I., Keele, S.W. (1970) Retention of abstract ideas. *Journal of Experimental Psychology*, 83: 304-308.

[Prablanc and Pellison 90] Prablanc, C., Pelisson, D. (1990) Gaze saccade and hand pointing are locked to their goal by quick internal loops. In *Attention and Performance XIII: Motor Representation and Control*, Jeannerod ed. Lawrence Erlbaum Associates Inc, Publishers.

[Press et al. 92] Press, W., Teukolsky, S.A., Vetterling, W.T., Flannery, B.P. (1992) *Numerical Recipes in C*. Cambridge University Press.

[Prakash 81] Prakash, N. (1981) *Differential Geometry. An Integrated Approach.* TATA McGraw-Hill Publishing Company Limited.

[Pylyshyn 80] Pylyshyn, Z. (1980) Computational models and empirical constraints. *Behavioral and Brain Sciences*, 1: 93-128.

[Rao and Ballard 95] Rao, R.P.N., Ballard, D.H. (1995) An active vision architecture based on iconic representations. *Artificial Intelligence*, 78: 461-505.

[Rizzolatti et al. 90] Rizzolatti, G., Gentilucci, M., Camarda, R.M., Gallese, V., Luppino, G., Matelli, M., Fogassi, L. (1990) Neurons related to reaching-grasping arm movements in the rostral part of area 6 (area $6\alpha\beta$). *Experimental Brain Research*, 82: 337-350.

[Russel and Norvig 95] Russel, S. J., Norvig, P. (1995) *Artificial Intelligence, a Modern Approach.*

[Sabes and Jordan 97] Sabes, P. and Jordan, M. (1997) Obstacle avoidance and a perturbation sensitivity model for motor planning. *Journal of Neuroscience*, 1997 17: 7119-7128.

[Sabes 96] Sabes, P. (1996) *The Planning of Visually Guided Arm Movements: Feedback Perturbation and Obstacle Avoidance Studies*. Ph.D. Thesis, MIT.

[Salinas and Abbot 94] Salinas, E., Abbott, L.F. (1994) Vector reconstruction from firing rates. *Journal of Computational Neuroscience,* 1:89-107.

[Sanger 94] Sanger, T.D. (1994) Theoretical considerations for the analysis of population coding in motor cortex. *Neural Computation,* 6: 29-37.

[Schmidt 75] Schmidt, R.A. (1975) A schema theory of discrete motor skill learning. *Psychological Review*, 82(4): 225-260.

[Schmidt et al. 79] Schmidt, R.A., Hawkins, B., Franck, J.S., Quinn, J.T. (1979) Motor-output variability: a theory for the accuracy of rapid motor tasks. *Psychological Review*, 86: 415-451.

[Schmidt and Young 91] Schmidt, R.A., Young, D.E. (1991) Methodology for motor learning: a paradigm for kinematic feedback. *Journal of Motor Behavior*, 23: 13-24.

[Schwartz 94] Schwartz, A. (1994) Direct cortical representation of drawing. *Science*, 265: 540-542.

[Schwartz and Moran 98] Schwartz, A.B., Moran, D.W. (1998) Motor cortical activity during drawing movements. 4: Population activity during lemniscate drawing. *Journal of Neurophysiology*, in press.

[Scott and Kalaska 97] Scott, S.H., Kalaska, J.F. (1997) Reaching movements with similar hand paths but different arm orientations. I. Activity of individual cells in motor cortex. *Journal of Neurophysiology*, 77: 826-852.

[Sergio and Kalaska 97] Sergio, L.E., Kalaska, J.F. (1997) Systematic changes in directional tuning of motor cortex cell activity with hand location in the workspace during generation of static isometric forces in constant spatial directions. *Journal of Neurophysiology*, 1170-1174.

[Shadmehr and Mussa-Ivaldi 94] Shadmehr, R., Mussa-Ivaldi, F. (1994) Adaptive representation of dynamics during learning of a motor task. *Journal of Neuroscience*, 14(5): 3208-3224.

[Shadmehr and Thoroughman 97] Learning and memory formation of arm movements. In: *Biomechanics and Neural Control of Movement*, Winters, J.M. and Crago, P. E. eds., Springer-Verlag, in press.

[Sims 94] Sims, K. (1994) Evloving 3D morphology and behavior by competition. *Artificial Life IV Proceedings*, Brooks, R. and Maes, P. ed. MIT Press.

[Summer 92] Summers, J.J. (1992) Motor behavior: a field in crisis? In *Approaches to the Study of Motor Control and Learning*, Summers ed. North-Holland.

[Soechting and Terzuolo 86] Soechting, J.F., Terzuolo, C.A. (1986) An algorithm for the generation of curvilinear wrist motion in an arbitrary plane in three dimensional space. *Neuroscience*, 19: 1393-1405.

[Soechting and Terzuolo 87] Soechting, J.F., Terzuolo, C.A. (1987) Organization of arm movements in three-dimensional space. Wrist motion is piecewise planar. *Neuroscience*, 23(1): 53-61.

[Soechting et al. 95] Soechting, J.F., Buneo, C.A., Herrmann, U., Flanders, M. (1995) Moving effortlessly in three dimensions: does Donder's law apply to arm movement? *Journal of Neuroscience*, 15(9): 6271.

[Somers et al. 98] Somers, D.C., Todorov, E.V., Siapas, A.G., Toth, L.J., Kim, D., Sur, M. (1998) A local circuit approach to understanding integration of long-range inputs in primary visual cortex. *Cerebral cortex*, 8(3): 204.

[Dum and Strick 91] Dum, R.P., Strick, P.L. (1991) The origin of corticospinal projections from premotor areas in the frontal lobe. *Journal of Neuroscience*, 11: 667-689.

[Terzopoulos et al. 94] Terzopoulos, D., Tu, X., Grzeszczuk, R. (1994) Artificial fishes: autonomous locomotion, perception, behavior, and learning in a simulated physical world. *Artificial Life*, 1(4): 327-351.

[Thach 85] Thach, W.T. (1985) Correlation of neuronal discharge with pattern and force of muscular activity, joint position, and direction of intended next movement in motor cortex and cerebellum. *Journal of Neurophysiology*, 41: 654-676.

[Thoroughman and Shadmehr 97] Thoroughman, K., Shadmehr, R. (1997) Influence of after-effects on learning of an internal model for reaching movements. *Society for Neuroscience Abstracts*, 23:203.

[Todorov et al. 97] Todorov, E., Shadmehr, R., Bizzi, E. (1997) Augmented feedback presented in a virtual environment accelerates learning of a difficult motor task. *Journal of Motor Behavior*, 29(2): 147-158.

[Todorov and Jordan 98] Todorov, E., Jordan, M.I. (1998) Smoothness maximization along a predefined path accurately predicts the speed profiles of complex arm movements. *Journal of Neurophysiology*, 80(2): 696.

[Tsuji et al. 95] Tsuji, T., Morasso, P.G., Goto, K., Ito, K. (1995) Human hand impedance characteristics during maintained posture. *Biological Cybernetics*, 72: 475-485.

[Tuller and Kelso 90] Tuller, B., Kelso, J.A.S. (1990) Phase transitions in speech production and their perceptual consequences. In *Attention and Performance XIII: Motor Representation and Control*, Jeannerod ed. Lawrence Erlbaum Associates Inc, Publishers.

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

[van Emmerik et al. 89] van Emmerik, R.E.A., Brinker, B.P.L.M., Vereijken, B., Whiting, H.T.A. (1989) Preferred tempo in the learning of a gross cyclical action. *Quaterly Journal of Experimental Psychology*, 41A: 251-262.

[van Galen and de Jong 95] van Galen, G.P, de Jong, W.P. (1995) Fitts' law as the outcome of a dynamic noise filtering model of motor control. *Human Movement Science*, 14: 539-571.

[Vindras et al. 96] Vindras, P., Viviani, P., Pellison, D. (1996) Transfer of gain adaptation: evidence of vectorial coding in goal-directed movements. *Society for Neuroscience Abstracts*, 169.8.

[Viviani and Cenzato 85] Viviani, P., Cenzato, M. (1985) Segmentation and coupling in complex movements. *Journal of Experimental Psychology: Human Perception and Performance*, 11: 828-845.

[Viviani and Schneider 91] Viviani, P., Schneider, R. (1991) A developmental study of the relationship between geometry and kinematics of drawing movements. *Journal of Experimental Psychology: Human Perception and Performance*, 17: 198-218.

[Viviani and Flash 95] Viviani, P., Flash, T. (1995) Minimum-jerk, two-thirds power law, and isochrony: converging approaches to movement planning. *Journal of Experimental Psychology: Human Perception and Performance*, 21(3): 2-53.

[Wann et al. 88] Wann, J., Nimmo-Smith, I., Wing, A.M. (1988) Relationship between velocity and curvature in movement: equivalence and divergence between a power law and a minimum-jerk model. *Journal of Experimental Psychology: Human Perception and Performance*, 14: 622-637.

[Weld and de Kleer 90] Weld, D.S., de Kleer, J. (1990) *Readings in Qualitative Reasoning About Physical Systems*. Morgan Kaufmann Publishers.

[Welch 86] Welch, R. (1986) Adaptation of space perception. In *Handbook of Human Perception and Performance*, Boff, Kaufman and Thomas, eds. Wiley, New York.

[Weiss 98] Weiss, Y. (1998) *Bayesian Motion Estimation and Segmentation*. Ph.D. Thesis, MIT, Cambridge MA.

[Wilsky 94] Wilsky, A.S. (1994) *Recursive Estimation*. Lecture notes, MIT.

[Wolpert et al. 95] Wolpert, D., Ghahramani, Z., Jordan, M.I. (1995) Are arm trajectories planned in kinematic or dynamic coordinates? An adaptation study. *Experimental Brain Research*, 103: 460-470.

[Wolpert et al. 95b] Wolpert, D., Ghahramani, Z., Jordan, M.I. (1995) An internal model for sensorimotor integration. *Science*, 269(5232): 8036-8075.

[Wolpert 97] Wolpert, D. (1997) Computational approaches to motor control. *Trends in Cognitive Science*, 1(6): 209-216.

[Winter 89] Winter, D.A. (1989) Coordination of motor tasks in human gait. In *Perspectives on the Coordination of Movement*, Wallace ed. North-Holland.

[Wolfe et al. 92] Wolfe, J.M., Yee, A., Friedman-Hill S.R. (1992) Curvature is a basic feature for visual search tasks. *Perception*, 21(4): 468-480.

[Woodworth 1899] Woodworth, R.S. (1899) The accuracy of voluntary movement. *Psychological Review*, 3: 1-114.

[Wright 90] Wright, C.E. (1990) Generalized motor programs: reexamining claims of effector independence in writing. In *Attention and Performance XIII: Motor Representation and Control*, Jeannerod ed. Lawrence Erlbaum Associates Inc, Publishers.

[Yalov 91] Yalov, S. (1991) *Computatinal Models of Hand Trajectory Planning in Drawing Movements*. M.Sc. Thesis, Weizmann Institute of Science.