



# A simple control policy for achieving minimum jerk trajectories

Mehrdad Yazdani<sup>a,\*</sup>, Geoffrey Gamble<sup>b</sup>, Gavin Henderson<sup>c</sup>, Robert Hecht-Nielsen<sup>a</sup>

<sup>a</sup> Electrical and Computer Engineering, University of California, San Diego, United States

<sup>b</sup> Computer Science and Engineering, University of California, San Diego, United States

<sup>c</sup> Mechanical and Aerospace Engineering, University of California, San Diego, United States

## ARTICLE INFO

### Article history:

Received 16 June 2011

Received in revised form 3 November 2011

Accepted 11 November 2011

### Keywords:

Jerk  
Minimum jerk  
Infinity norm  
Optimal control  
Manipulandum  
Ballistic movements  
Bell-shape velocity profiles  
Bang–bang  
Human movements  
Motor control  
Feedback effects

## ABSTRACT

Point-to-point fast hand movements, often referred to as ballistic movements, are a class of movements characterized by straight paths and bell-shaped velocity profiles. In this paper we propose a bang–bang optimal control policy that can achieve such movements. This optimal control policy is accomplished by minimizing the  $L_\infty$  norm of the jerk profile of ballistic movements with known initial position, final position, and duration of movement. We compare the results of this control policy with human motion data recorded with a manipulandum. We propose that such bang–bang control policies are inherently simple for the central nervous system to implement and also minimize wear and tear on the bio-mechanical system. Physiological experiments support the possibility that some parts of the central nervous system use bang–bang control policies. Furthermore, while many computational neural models of movement control have used a bang–bang control policy without justification, our study shows that the use of such policies is not only convenient, but optimal.

© 2011 Elsevier Ltd. All rights reserved.

## 1. Introduction

The process of evolution drives species to differentiate to produce some competitive advantage. Just as it is easy to observe that evolution has led to a great many modes of locomotion (Grillner & Jessell, 2009), it can be assumed that evolution has had effects on the control of behavior by the nervous system. Under this assumption, we can then ask which attributes of the nervous system's control policy might have adapted to provide a competitive advantage. It is reasonable to assume that a bio-mechanical system evolves to find an optimal control policy by optimizing over some cost or reward function. In this paper, we suggest a new cost function and discuss the control policy this cost function dictates.

How the CNS implements a control policy to achieve movements is not understood. Areas of cortex send axons to the spinal cord and generate movements when stimulated, which has led many to believe that cortical structures are responsible for most aspects of movement control (Georgopoulos, Schwartz, & Kettner,

1986; Lemon, 1988; Penfield & Boldrey, 1937). Determining the exact parameters by which the cortex is able to alter behavior has proved to be difficult. Many experiments have correlated cortical activity with various aspects of a behavior (Georgopoulos et al., 1986; Hammon, Makeig, Poizner, Todorov, & de Sa, 2008; Lemon, 1988; Penfield & Boldrey, 1937). For example, Graziano's work (Graziano, Taylor, & Moore, 2002) implies that the cortical description of a behavior may be limited to a high level goal state representation. Other evidence shows that the spinal cord plays a key role in the generation of behaviors (Giszter, Mussa-Ivaldi, & Bizzi, 1993; Pearson, 1976).

One method used to understand the nature of movements is to reduce them to simpler components. For this study we limit ourselves to examining simple, ballistic point-to-point reaching movements. Simple movements might be considered to make up a basis set of which more complicated movements are composed (Giszter, 1992; Hart & Giszter, 2010; Mussa-Ivaldi, Giszter, & Bizzi, 1994). By studying these simple movements, we may be able to gain insight into the control of more complex movements such as curved movements defined by via points (Flash & Hogan, 1985) or paths (Todorov & Jordan, 1998). Furthermore, it has been observed that this class of movements consistently follows bell-shaped velocity profiles (Abend, Bizzi, & Morasso, 1982; Flash & Hogan, 1985). This observation suggests that there exists a set of constraints placed on the dynamic system by a controller. Because

\* Correspondence to: 9500 Gilman Drive, Mail code 0409, La Jolla, CA 92093-0409, United States.

E-mail addresses: [myazdani@ucsd.edu](mailto:myazdani@ucsd.edu) (M. Yazdani), [ggamble@cs.ucsd.edu](mailto:ggamble@cs.ucsd.edu) (G. Gamble), [ghenders@ucsd.edu](mailto:ghenders@ucsd.edu) (G. Henderson), [r@ucsd.edu](mailto:r@ucsd.edu) (R. Hecht-Nielsen).

this set of constraints exists across a range of movements, it can be said that these movements result from a common control policy. We show here that the observable invariant parameters of movement suggest a neural control policy which is corroborated by neurophysiological experiments.

There have been several attempts to mathematically model human point-to-point movements of an end effector (the hand) (Flash & Hogan, 1985; Todorov & Jordan, 1998; Viviani & Flash, 1995). A key study of this type is (Flash & Hogan, 1985). Flash and Hogan began their work on modeling movement by observing that short duration, straight line reaching movements (ballistic movements) exhibit a stereotypical bell-shaped velocity profile. Their work resulted in a model that was exceptionally good at reproducing the trajectory of a movement given the limited information of initial position, final position, and movement duration. This model was achieved by finding the trajectory that minimized the  $L_2$  norm of the 3rd derivative of the position trajectory. The 3rd derivative of position is known primarily as “jerk”, but is also known as “shock”, “jolt”, “surge” or “lurch”. We henceforth refer to their model minimizing the  $L_2$  norm of jerk as  $MJ_2$ .

Although the  $MJ_2$  is exceptionally good at reproducing trajectories from limited constraints, it remains unclear how the central nervous system would generate these trajectories. Various other models have been proposed that minimize other derivatives of position such as acceleration (Ben-Itzhak & Karniel, 2008). This recent work by Ben-Itzhak and Karniel has produced a model that not only generates accurate point-to-point movements but also suggests a control policy by which the CNS could be generating such movements. The work presented here expands upon those findings and suggests an alternative model that presents a simple control policy the CNS may implement.

We propose an optimal control policy for achieving ballistic movements based on minimizing the jerk of the trajectory of the end effector. We formulate the problem as an optimal control problem wherein the jerk is treated as the control signal. Our model minimizes the  $L_\infty$  norm of the jerk and shows that the optimal control policy is of a “bang–bang” type controller, a policy which simply switches a system between two states (Kirk, 2004). The appeal of such controllers is that they are inherently simple to implement. Furthermore, minimizing the  $L_\infty$  norm minimizes the maximum allowable jerk for the system, which can reduce wear and tear. Henceforth, we refer to our proposed model as  $MJ_\infty$ .

Flash and Hogan also solved for the trajectory that minimizes jerk; however, their cost function utilizes the  $L_2$  norm. While such a cost function yields bell-shaped velocity profiles and position profiles as observed in humans, the jerk profiles given by the model are not a simple bang–bang type controller. Ben-Itzhak and Karniel developed a model (MACC) that also yields accurate trajectories. In contrast to an  $MJ_2$  model, their model implies a simple bang–bang controller. They arrive at this model by minimizing the  $L_2$  norm of the acceleration with a free parameter constraining the maximum allowable jerk. While they show that their model improves error significantly, we argue that their usage of a free parameter is not needed and adds unnecessary complexity. The model presented here still yields a bang–bang control policy but is less complex in that no free parameter is required.

This paper is divided into the following sections. In Section 2, we describe our model used to describe realistic human movements as well as provide justification for its biological importance. Next, in Section 3, we explain the process by which the human movement data were collected and parsed for proper analysis and comparison. In Section 4 we discuss the results of the data comparison with our model,  $MJ_\infty$ , and the  $MJ_2$  model. The last section discusses the importance of this work and draws conclusions about the insights provided by modeling human movement.

## 2. Model description

The reasons the central nervous system minimizes the jerk of movements are not immediately apparent. Mechanical systems have maximum tolerances related to various dynamic variables (velocity, acceleration, jerk, etc.). Beyond these tolerance levels, components of the system may begin to fail. Biological systems are mechanical systems and therefore also have thresholds that, when exceeded, may lead to damage such as tearing of ligaments and muscles or breaking of bones. Jerk is one of the dynamic variables that bears directly on the well-being of a mechanical system. Mechanical engineers and roboticists have recognized the benefits of minimizing jerk and have incorporated this concept into their systems (Gasparetto & Zanotto, 2008; Kyriakopoulos & Saridis, 1988; Pattacini, Nori, Natale, Metta, & Sandini, 2010; Piazzi & Visioli, 2000). Optimizing animal movement by minimizing jerk is beneficial in that it can reduce stress on the mechanical components of the body.

It is not obvious what function of the instantaneous jerk should be minimized to match biological observation. The  $L_2$  norm (as used in the  $MJ_2$  model) measures the summation of squared jerk over the course of the movement while the  $L_\infty$  norm (as used in the  $MJ_\infty$  model) minimizes the maximum jerk value over the course of the movement. While the  $L_2$  norm metric penalizes high jerk values, it does not explicitly force the system to keep the maximum instantaneous jerk as low as possible. In the jerk profile figure shown in Fig. 1, notice that near time  $t = 0$  and  $t = 1$  s, the jerk resulting from the  $MJ_2$  model far exceeds the maximum jerk over the entire movement by the  $MJ_\infty$  model. In contrast to the  $L_2$  norm metric, limiting the maximum instantaneous magnitude of jerk via an  $L_\infty$  norm cost function reduces the possibility of the movement passing some critical jerk threshold, after which damage to the body may occur. This intuitive rationale helps justify why evolution may have minimized the maximum magnitude of instantaneous jerk ( $L_\infty$  norm) during a movement rather than the sum of squared jerk over the course of a movement ( $L_2$  norm).

### 2.1. Minimizing jerk as a control variable

In this section we formulate the problem of minimizing the jerk of a ballistic point-to-point movement as a control problem where the control signal is the jerk, the initial and final positions are known, and the duration of movement is also known. For ease of notation, here we restrict our problem formulation to one dimension and note that extensions to higher dimensions are straight-forward. The control signal that we seek to achieve a minimum jerk position trajectory  $x(t)$  is formulated as follows:


$$\begin{aligned} &\underset{u(t)}{\text{minimize}} \quad \|u(t)\|_p \\ &\text{subject to} \quad \dot{\mathbf{x}}(t) = A\mathbf{x}(t) + Bu(t) \end{aligned} \quad (1)$$

where  $A = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{bmatrix}$ ,  $B = \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}$ ,  $\mathbf{x}(t) = \begin{bmatrix} x(t) \\ \dot{x}(t) \\ \ddot{x}(t) \end{bmatrix}$ ,  $u(t) = \ddot{\ddot{x}}(t)$ , and where  $\|\cdot\|_p$  denotes the  $L_p$  norm.

The solution to Eq. (1) will determine the optimal control policy  $u(t)$ . The selection of the  $L_p$  norm can result in vastly different control policies. For  $1 \leq p < 2$ , the control policy will result in physiologically unrealistic movements and as a result these types of control policies are not discussed here. Instead we will pay close attention to cases where  $p = 2$  and  $p = \infty$ . For  $p = 2$ , we have the following policy as described in the following theorem:

**Theorem 1.** *The solution to Eq. (1) with  $p = 2$  is a straight line trajectory given by the following control policy:*

$$u(t) = \ddot{\ddot{x}}(t) = (x_f - x_i) \left( \frac{360}{T^5} t^2 - \frac{360}{T^4} t + \frac{60}{T^3} \right)$$



**Fig. 1.** The profiles corresponding to the  $L_2$  norm control policy and  $L_\infty$  norm control policy are shown in red and blue respectively. These plots are generated according to Theorems 1 and 2 respectively.

where  $x_i$  is the initial hand position at time  $t = 0$  and  $x_f$  is the final hand position at time  $t = T$ .

**Proof.** This was originally shown by Flash and Hogan (1985).  $\square$

The control policy corresponding to the above theorem has been shown to fit human data very well (Flash & Hogan, 1985). The next theorem shows that if the  $L_\infty$  norm of jerk is minimized (i.e.,  $p = \infty$ ) in Eq. (1) then we have a bang–bang control policy.

**Theorem 2 (Bang–Bang Control Theorem).** The solution to Eq. (1) with  $p = \infty$  is a straight line trajectory given by the following control policy:

$$u(t) = \ddot{x}(t) = \begin{cases} J & 0 \leq t < \frac{T}{4} \\ -J & \frac{T}{4} \leq t < \frac{3T}{4} \\ J & \frac{3T}{4} \leq t \leq T \end{cases} \quad (2)$$


with  $J = 32 \frac{x_f - x_i}{T^3}$  where  $x_i$  is the initial hand position at time  $t = 0$  and  $x_f$  is the final hand position at time  $t = T$ .

**Proof.** This was originally shown by Kyriakopoulos and Saridis (1988).  $\square$

Ben-Itzhak and Karniel (2008) proposed a similar bang–bang control policy for achieving ballistic point-to-point movements. Their control policy minimizes acceleration and also places a threshold on the jerk of the trajectory. This threshold is a free parameter in their model that controls the amount of allowable jerk. Here we show that achieving a bang–bang control policy can be done without introducing any free parameters simply by minimizing jerk as measured by the infinity norm.

### 3. Methods

The human arm movement data used in this study for comparison with our model were provided by Amir Karniel at the Ben Gurion University of the Negev, the same data he and Ben-Itzhak




**Fig. 2.** The boxes labeled A, B and C are targets which appeared on a screen in front of the subjects. The triangle at the end of the subjects arm represents the subjects hand and the manipulandum he is moving. The subject moves his hand between two of the points within a specified time window. The three targets allow for six distinct movement types indicated by the six arrows. The subject receives visual feedback indicating if he or she has successfully reached a target and if he or she has done so within the allotted time window.

Source: Diagram adapted from Ben-Itzhak and Karniel (2008).

used in their 2008 paper outlining their model of fast arm reaching movements (Ben-Itzhak & Karniel, 2008). The data originated from a 2002 paper by Karniel and Mussa-Ivaldi (2002) investigating the nervous system's ability to adapt to perturbations. An abbreviated description of the data collection techniques is given below. For a complete description, see Karniel and Mussa-Ivaldi (2002).

Seated subjects held a robotic manipulandum which was restricted to two dimensional movements corresponding to the horizontal plane of the subjects. They watched a screen which displayed the position of their hand (and the manipulandum) in relation to three positional markers A, B and C. The markers were positioned to form an equilateral triangle (see Fig. 2). The



**Fig. 3.** A typical recording of the position profile. The oscillatory behavior at the end of the movement corresponds to the overshoot and correctional effects described in the text and are discarded since they are not part of the ballistic portion of the movement. The units of time and position are not relevant and are not shown.

subjects were instructed to move the on-screen representation of the manipulandum from one target to another. The distance between the targets was 10 cm. This motion was to occur within one third of a second,  $\pm 50$  ms. Feedback was given to the subjects indicating if they had reached the target and if they did so within the appropriate time window. Position profiles were recorded for all six possible movement types for five subjects over the course of four days. The original data included a subset of trials in which the arm was perturbed during movement from one marker to another. This subset was excluded from our analysis. Only unperturbed movements were analyzed. See Fig. 3 for an example of a typical movement.


The data included uninteresting aspects such as near stationary positional information before a subject began moving and after a subject reached his or her goal and stopped moving. Various methods have been used for movement onset detection (Botzer & Karniel, 2009; Staude, 2001; Staude, Wolf, Appel, & Dengler, 1996). Unfortunately, there is no consensus regarding which technique is best for choosing the relevant portion of a movement as the definition of what is relevant may change from study to study or from one movement type to another. We employ a simple method to determine the start and end times of each movement. We start by finding the onset of the movement. To do so, we compute the energy of a moving window of five time steps over the velocity profile from a given trial:

$$E = \sum_{i=1}^5 v^2[i] \quad (3)$$

where  $v[i]$  is the velocity of the manipulandum at time step  $i$  of the current window. The window starts at the beginning of the recorded data of a trial (time steps 1–5) and moves forward in one time step increments (e.g. 2–6, 3–7, etc.). If  $E$  is not greater than a threshold,  $\delta$ , the window continues moving forward and the test is repeated.

When the window moves over a portion of the velocity profile where the manipulandum is both stationary and close to the starting position (i.e. before the movement begins),  $E$  is low. As the window moves over a portion of the velocity profile which is increasing,  $E$  becomes greater. We define the starting time of the movement to be the beginning time step of the window at the first window position where  $E \geq \delta$ .

Finding the time at which the ballistic portion of the movement ends is difficult due to corrective movements made by the subjects



**Fig. 4.** Histogram of the ratio of average starting and stopping velocities to peak velocity for each trial. Ideally, the trials should have start and end velocities as close to rest as possible in order to conform to both models' assumptions.

after they reach their target, e.g. the correction of perceived target overshoot. Recall that we are only interested in ballistic movements. The corrective portions of the movement fall outside the ballistic portion of the movement. We define the middle ( $T/2$ ) and end ( $T$ ) times of a movement in the same way as done in Ben-Itzhak and Karniel (2008). End effector velocity profiles for ballistic point-to-point movements are known to have a symmetric bell shape (Abend et al., 1982; Flash & Hogan, 1985). In order to determine the end time of the movement, we first define the middle position of the movement ( $T/2$ ) to be the point of maximum velocity, i.e. the top of the symmetric bell. We then simply double this value to find the end time  $T$ .


Like other methods, this heuristic technique is not guaranteed to find the ideal onset and end of the recorded movements. Both the  $MJ_2$  and  $MJ_\infty$  models assume an initial and final position at which velocity is zero. For this reason, it is appropriate to filter the trials, keeping those for which our start/end detection algorithm has chosen points which most closely meet the zero velocity start and end point assumptions of the models. By definition these are the only trials that are relevant to the models.

Since none of the trials have exactly zero velocity start and end points, some degree of tolerance must be allowed. Furthermore, some metric must be employed to define the degree of closeness to zero velocity for a given start/end point. We define close to zero velocity for start and end points on a trial by trial basis by computing a ratio between the velocities in question and the peak velocity:

$$c = \frac{v_s + v_e}{2v_p} \quad (4)$$

where  $c$  is a unitless indicator of closeness to zero,  $v_s$  is starting velocity,  $v_e$  is the ending velocity and  $v_p$  is the peak velocity of the movement. Applying this metric to all trials yields a distribution between 0 and 1 (see Fig. 4). As the  $c$  value gets closer to zero, the corresponding trial increasingly conforms to the zero start and end assumptions of the  $MJ_2$  and  $MJ_\infty$  models and is therefore more appropriate for comparison against the models. We included all trials with a  $c$  value less than the harmonic mean of the entire distribution. That is, we favored trials which most closely conform to the assumptions of the models while still leaving enough trials to properly gauge statistical significance. The use of the harmonic mean as a level of tolerance of deviation from zero velocity is somewhat arbitrary. What is important is that the trials with





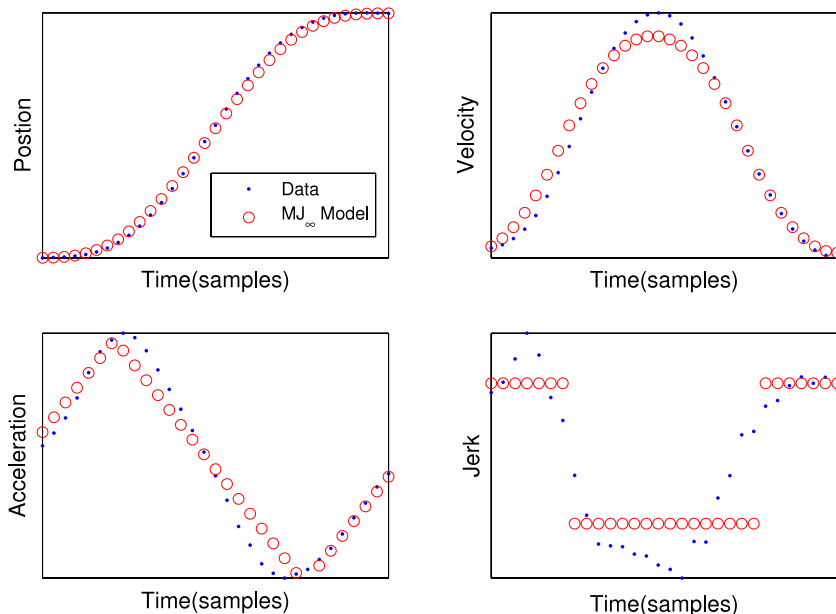
**Fig. 5.** Bars A–E show the comparison of the average MSE of position profiles predicted by  $MJ_2$  and  $MJ_\infty$  vs. human trial data along with standard error bars. “All Subjects” is an aggregate of the trials in A–E. The  $MJ_\infty$  model performs better than the  $MJ_2$  model in all cases. The results for A, B, E are significant with  $p < 0.05$  by a Wilcoxon rank-sum test. The “All Subjects” aggregate results are extremely significant with  $p < 0.001$  by a Wilcoxon rank-sum test. This test was utilized due to the non-normality of the data as is typically done in this situation (Ben-Itzhak & Karniel, 2008).

large  $c$  values (that do not conform to the models’ assumptions) are discarded while keeping enough trials to maintain statistical significance. In all, our filtered data set included 406 movement trials.

#### 4. Results

To assess the performance of the minimum  $L_2$  norm jerk model ( $MJ_2$ ) in comparison to the minimum  $L_\infty$  norm jerk model ( $MJ_\infty$ ), we compute the time-series mean-squared error between the model’s predicted position trajectory with the human subjects trajectory. Our intention is to show that  $MJ_\infty$ ’s much simpler control policy can fit data at a high accuracy. In fact, we show that our model significantly exceeds the accuracy of the  $MJ_2$  model for this data set. Formally, the mean-squared error is computed as follows:

$$MSE_k = \frac{1}{t_k} \sum_{i=1}^{t_k} (x_k[i] - p_k[i])^2 \quad (5)$$



**Fig. 6.** Overlay of  $MJ_\infty$  model with human movement data of a single trial.

where  $k$  refers to the trial,  $t_k$  the number of samples for the  $k$ th trial,  $x_k$  is the model position profile for the  $k$ th trial, and  $p_k$  is the recorded data position data for the  $k$ th trial.

Fig. 5 shows the comparison of the average MSE trajectories between  $MJ_2$  and  $MJ_\infty$ .  $MJ_\infty$  has a smaller MSE in all cases. We therefore can conclude that the trajectory estimates of  $MJ_\infty$  are just as good or significantly better than those of  $MJ_2$ . Furthermore, their control policies differ significantly and we suspect that the simple control policy (the bang–bang controller of  $MJ_\infty$ ) would be favored by a biological system.

The MACC (Minimum Acceleration Criterion with Constraints) model proposed by Ben-Itzhak and Karniel (2008) also implies a simple bang–bang controller; however, their approach requires a free parameter that places a cap on the maximum allowable jerk (manifested as a constraint on the control signal). Although changing this parameter can change the switching times of the bang–bang controller, it is not clear how to select an appropriate value for this parameter. Ben-Itzhak and Karniel choose the value for this parameter by performing a grid search and selecting the best value that fit the data for each trial. Even though they showed that using this method of selecting the parameter value results in trajectories that have MSE significantly smaller than the  $MJ_2$ , these results depend upon the model’s ability to tune this parameter on a trial by trial basis. This may explain the improvement over other models they have used for comparison. While it is feasible for the CNS to implement additional parameters, free parameters add unnecessary complexity for achieving bang–bang control. It is worth noting that while our model does not require any free parameters, it is a special, but important, case of the MACC model where the free parameter is chosen such that the infinity-norm of the jerk profile is minimized and hence is equivalent to the  $MJ_\infty$  jerk profile.

Our suggestion of a bang–bang control policy is based upon the fact that minimizing jerk with an  $L_\infty$  norm measure results in a two state jerk profile. Since the jerk is an observable characteristic of the movement, experimental data should be able to confirm the true nature of the jerk profile. Unfortunately, numerical computations of jerk by derivative approximation (as done in this paper) amplify noise inherent in the original recording. This makes drawing direct conclusions about the nature of the jerk profile difficult (see jerk approximation in Fig. 6). To overcome

this difficulty we decided to evaluate the error of the models with respect to the positional data. Future work on this hypothesis should include recordings of acceleration, which should reduce noise amplification and possibly make the true jerk profile apparent.

## 5. Discussion and future work

In this paper we studied an optimal control policy for achieving point-to-point ballistic movements using the minimum-jerk criterion. We focused on minimizing the  $L_\infty$  norm of jerk ( $MJ_\infty$ ) to achieve a simple bang–bang control policy as opposed to using the  $L_2$  norm ( $MJ_2$ ). We compared the two policies with human motion data recorded with a manipulandum and showed that the  $MJ_\infty$  outperforms the  $MJ_2$  at predicting ballistic human arm movements.

Determining the precise contributions of the various components of the CNS to the control of movement is difficult, since observations of the motor system's neural activity in behaving animals are hard to obtain. However, measurements of external motor behavior are much easier to record. It is natural then to attempt to leverage the movement data we have to explain what the control policy used by the CNS may be. Since the movements were performed by well trained individuals in an unperturbed workspace, it is reasonable to assume that feedback due to movement errors would be minimal. Furthermore, experimental studies on deaf-ferented animals have demonstrated that trajectory planning for fast point-to-point movements is not disrupted (Bizzi, Accornero, Chapple, & Hogan, 1984) and that proprioceptive or cutaneous feedback is not necessary for the execution of such movements. Because of the lack of feedback involved with these movements, we can model these movements as a feed forward control problem.

With this in mind, our proposed optimal control problem was solved once, and the solution was used for each trial (as done in Flash and Hogan (1985)). We remind the reader that we are not proposing that the biological system is performing an optimization computation each time a movement occurs. Instead, we are suggesting that evolution has already performed the optimization process via some cost function and arrived at a neuromechanical system (the human body) with a construction intrinsically built to minimize jerk.

The selection of a cost function to optimize is crucial and can result in vastly different control policies. If we assume a cost function that evolution has used to optimize animal movement, we can extrapolate a corresponding control policy. In addition, we can draw inferences regarding the nature of the biological mechanisms that might implement such a system. Since minimizing the  $L_\infty$  norm of jerk results in a bang–bang control policy, we can hypothesize that simple two-state step functions are utilized to control a biomechanical system. These two-state step functions are desirable because binary control is simple.

In addition, utilizing two-state control policies have been shown to be effective in computational models of movement. Recent computational models of spinomuscular control require only step functions representing supraspinal inputs in order to drive a network to achieve human like movements (Bullock & Grossberg, 1988; Raphael, Tsianos, & Loeb, 2010). Other models have shown that central pattern generators can be driven via step inputs (Buchanan & McPherson, 1995).

Similarly, on/off control policies have been observed both *in vivo* and *in vitro* in multiple vertebrates. Complex movements such as walking can be activated by gross on/off stimulation of groups of neurons in the brain stem or the spinal cord. The experiments reported in Pearson (1976) induced various patterns of locomotion in a spinally transected cat by administering a simple step-like electrical stimulation of the lower region of the cat's spinal cord. The experiments in Hagglund, Borgius, Dougherty, and Kiehn

(2010) showed that fictive locomotor patterns could be induced with the use of step-like excitation to either the brain stem or spinal cord of mice. Classes of neurons in the brain stem and lumbar regions of the spinal cord of the mice were genetically engineered to contain channelrhodopsin light gated ion channels. Using this technique, light stimulation (or lack thereof) served as an on/off switch for the genetically modified motor system neurons. Gross “on” stimulation to either the brain stem or the lumbar region of the spinal cord activated a class of neurons in those regions and induced fictive locomotion patterns. The gross “on” stimulation was a control signal driving the generation of locomotor patterns.

The computational and animal experiments explained above indicate that one or more bang–bang type controllers may exist somewhere in the nervous system. It is still an open question as to how and where their neural implementations exist. The evidence cited here suggests that these controllers may exist in supraspinal centers (Graziano et al., 2002; Raphael et al., 2010), within the spinal cord (Giszter et al., 1993; Pearson, 1976), or both in the brain stem and spinal cord (Hagglund et al., 2010). Furthermore, it is plausible that populations of bursting neurons could implement a bang–bang control signal (Izhikevich, 2007). This minimum jerk-based bang–bang control signal could then be converted to a signal representing any lower order derivative (acceleration, velocity, etc.) via integration of the signal. Certain populations of neurons are known to perform functions akin to time integration (Goldman, Compte, & Wang, 2009). Integration of a neural signal encoding velocity would lead to a signal encoding position as happens in the ocular-motor system (Robinson, 1989; Seung, Lee, Reis, & Tank, 2000). We stress that these suggestions for our control policy's neural implementation are merely hypotheses backed by neuroscientific evidence.

We suggest several important extensions to this work. As discussed earlier, to gauge a more accurate jerk profile than that attained by numerical approximations of higher derivatives, we propose using a manipulandum device where the jerk profile can be recorded directly with high bit precision. As shown in Fig. 1 the jerk profile that we propose has step-like features and discontinuities. If the jerk profile of human arm reaching movements has such features, then care must be taken to acquire and digitize the jerk signal appropriately. Only then can we effectively compare the  $MJ_\infty$  jerk profile with the acquired jerk profile from human subjects.

Another interesting extension would be to investigate how well  $MJ_\infty$  models curved movements as done by  $MJ_2$  and other models (Edelman & Flash, 1987; Todorov & Jordan, 1998). Although the  $MJ_\infty$  performs very well at modeling straight point-to-point ballistic movements, and it is likely to perform well with curved movements, it is possible that more exotic control policies might be needed to explain more complex movements. With this in mind, we plan on making simultaneous use of multiple control policies by switching between them or blending them depending on the nature of the task at hand. In addition, to achieve a model that replicates a wider array of human movements (such as perturbed movements), future research directions should extend this feed-forward model to include feedback (environmental and proprioceptive). Finally, to acquire a deeper understanding of the central nervous system, we propose modeling ballistic movements with neuronal networks in order to test the control signal schemes outlined in this work and their potential neural implementations.

## Acknowledgments

The authors are most grateful to Amir Karniel for providing the data that made this work possible. We thank ONR for its long-standing support of this research and William Lennon Jr., Rupert Minett, and Andrew T. Smith for useful discussions. We are also grateful to the anonymous reviewers for their suggestions and comments.

## References

- Abend, W., Bizzi, E., & Morasso, P. (1982). Human arm trajectory formation. *Brain*, 105(Pt. 2), 331–348.
- Ben-Itzhak, S., & Karniel, A. (2008). Minimum acceleration criterion with constraints implies bang–bang control as an underlying principle for optimal trajectories of arm reaching movements. *Neural Computation*, 20(3), 779–812.
- Bizzi, E., Accornero, N., Chapple, W., & Hogan, N. (1984). Posture control and trajectory formation during arm movement. *Journal of Neuroscience*, 4(11), 2738–2744.
- Botzer, L., & Karniel, A. (2009). A simple and accurate onset detection method for a measured bell-shaped speed profile. *Frontiers in Neuroscience*, 3(61), 1–8.
- Buchanan, J. T., & McPherson, D. R. (1995). The neuronal network for locomotion in the lamprey spinal cord: evidence for the involvement of commissural interneurons. *Journal of Physiology—Paris*, 89(4–6), 221–233.
- Bullock, D., & Grossberg, S. (1988). Neural dynamics of planned arm movements—emergent invariants and speed accuracy properties during trajectory formation. *Psychological Review*, 95(1), 49–90.
- Edelman, S., & Flash, T. (1987). A model of handwriting. *Biological Cybernetics*, 57, 25–36.
- Flash, T., & Hogan, N. (1985). The coordination of arm movements—an experimentally confirmed mathematical-model. *Journal of Neuroscience*, 5(7), 1688–1703.
- Gasparetto, Alessandro, & Zanotto, Vanni (2008). A technique for time-jerk optimal planning of robot trajectories. *Robotics and Computer-Integrated Manufacturing*, [ISSN: 0736-5845] 24, 415–426.
- Georgopoulos, A. P., Schwartz, A. B., & Kettner, R. E. (1986). Neuronal population coding of movement direction. *Science*, 233(4771), 1416–1419.
- Giszter, S. (1992). Spinal movement primitives and motor programs—a necessary concept for motor control. *Journal of Behavioral and Brain Science*, 15(4), 744–745.
- Giszter, S. F., Mussa-Ivaldi, F. A., & Bizzi, E. (1993). Convergent force fields organized in the frog's spinal cord. *Journal of Neuroscience*, 13(2), 467–491.
- Goldman, M. S., Compte, A., & Wang, X. J. (2009). Neural integrator models. In Larry R. Squire (Ed.), *Encyclopedia of neuroscience* (pp. 165–178). Oxford: Academic Press, Editor in Chief.
- Graziano, M. S. A., Taylor, C. S. R., & Moore, T. (2002). Complex movements evoked by microstimulation of precentral cortex. *Neuron*, 34(5), 841–851.
- Grillner, S., & Jessell, T. M. (2009). Measured motion: searching for simplicity in spinal locomotor networks. *Current Opinion in Neurobiology*, 19(6), 572–586.
- Hagglund, M., Borgius, L., Dougherty, K. J., & Kiehn, O. (2010). Activation of groups of excitatory neurons in the mammalian spinal cord or hindbrain evokes locomotion. *Nature Neuroscience*, 13(2), 246–252.
- Hammon, P. S., Makeig, S., Poizner, H., Todorov, E., & de Sa, V. R. (2008). Predicting reaching targets from human EEG. *IEEE Signal Processing Magazine*, 25(1), 69–77.
- Hart, C. B., & Giszter, S. F. (2010). A neural basis for motor primitives in the spinal cord. *Journal of Neuroscience*, 30(4), 1322–1336.
- Izhikevich, Eugene (2007). *Dynamical systems in neuroscience: chapter 9*. Cambridge, MA: MIT Press.
- Karniel, A., & Mussa-Ivaldi, F. A. (2002). Does the motor control system use multiple models and context switching to cope with a variable environment? *Experimental Brain Research*, 143(4), 520–524.
- Kirk, Donald E. (2004). *Optimal control theory: an introduction*. Mineola, NY: Dover Publications.
- Kyriakopoulos, K. J., & Saridis, G. N. (1988). Minimum jerk path generation. In 1988 IEEE international conference on robotics and automation, 1988. *Proceedings. April. Vol. 1* (pp. 364–369).
- Lemon, R. (1988). The output map of the primate motor cortex. *Trends in Neurosciences*, 11(11), 501–506.
- Mussa-Ivaldi, F. A., Giszter, S. F., & Bizzi, E. (1994). Linear-combinations of primitives in vertebrate motor control. *Proceedings of the National Academy of Sciences of the United States of America*, 91(16), 7534–7538.
- Pattacini, U., Nori, F., Natale, L., Metta, G., & Sandini, G. (2010). An experimental evaluation of a novel minimum-jerk cartesian controller for humanoid robots. In *Intelligent robots and systems. IROS. 2010 IEEE/RSJ international conference on. October* (pp. 1668–1674).
- Pearson, K. (1976). Control of walking. *Scientific American*.
- Penfield, W., & Boldrey, E. (1937). Somatic motor and sensory representation in the cerebral cortex of man as studied by electrical stimulation. *Brain*, 60, 389–443.
- Piazzi, A., & Visioli, A. (2000). Global minimum-jerk trajectory planning of robot manipulators. *IEEE Transactions on Industrial Electronics*, 47(1), 140–149.
- Raphael, G., Tsianos, G. A., & Loeb, G. E. (2010). Spinal-like regulator facilitates control of a two-degree-of-freedom wrist. *Journal of Neuroscience*, 30(28), 9431–9444.
- Robinson, D. A. (1989). Integrating with neurons. *Annual Review of Neuroscience*, 12, 33–45.
- Seung, H. S., Lee, D. D., Reis, B. Y., & Tank, D. W. (2000). Stability of the memory for eye position in a recurrent network of conductance-based model neurons. *Neuron*, 26, 259–271.
- Staudte, G. H. (2001). Precise onset detection of human motor responses using a whitening filter and the log-likelihood-ratio test. *IEEE Transactions on Biomedical Engineering*, 48(11), 1292–1305.
- Staudte, G. H., Wolf, W. M., Appel, U., & Dengler, R. (1996). Methods for onset detection of voluntary motor responses in tremor patients. *IEEE Transactions on Biomedical Engineering*, 43(2), 177–188.
- 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), 696–714.
- Viviani, P., & Flash, T. (1995). Minimum-jerk, 2/3-power law, and isochrony-converging approaches to movement planning. *Journal of Experimental Psychology: Human Perception and Performance*, 21(1), 32–53.