

Visuomotor Optimality and its Utility in Parametrization of Response

Michael Sherback, Raffaello D'Andrea

Abstract—We present a method of characterizing visuomotor response by inferring subject-specific physiologically meaningful parameters within the framework of optimal control theory. The characterization of visuomotor response is of interest in the assessment of impairment and rehabilitation, the analysis of man-machine systems, and sensorimotor research. We model visuomotor response as a Linear Quadratic Gaussian (LQG) controller, a Bayesian optimal state estimator in series with a linear quadratic regulator. Subjects used a modified computer mouse to attempt to keep a displayed cursor at a fixed desired location despite a Gaussian random disturbance and simple cursor dynamics. Nearly all subjects' behavior was consistent with the hypothesized optimality. Experimental data was used to fit an LQG model whose assumptions are simple and consistent with other sensorimotor work. The parametrization is parsimonious and yields quantities of clear physiological meaning: noise intensity, level of exertion, delay, and noise bandwidth. Inferred control cost and noise intensity varied significantly across subjects. Response variations were consistent with changes in exerted effort. This is a novel example of the role of optimal control theory in explaining variance in human visuomotor response. We also present technical improvements on the use of LQG in human operator modeling.

Index Terms—Sensorimotor, visuomotor, human operator, HITL, LQG.

I. INTRODUCTION

We seek a parametric model of human visuomotor response in a simple feedback task. Models of visuomotor behavior are of value in the assessment of neuromuscular health during rehabilitation, in human operator modeling, and in sensorimotor research. In our experiment subjects used a modified computer mouse to attempt to keep a displayed cursor at a fixed desired location despite a Gaussian random disturbance and simple cursor dynamics. We parametrized response in terms of physiologically meaningful quantities by fitting a subject specific optimal control model. Our parametrization is parsimonious and rests on accepted priors and assumptions. We demonstrate that optimality is invariant across nearly all subjects despite variance in the criteria of optimality. The methods of this paper are automated and subject-specific, and thus potentially useful in clinical evaluation of neuromuscular health.

This paper is organized as follows: Section I provides background on theory and past work. The hypothesized model is given in Sec. II. Experimental, spectral, and data fitting methods follow in Sec. III. Results and analysis are in Sec. IV, and discussion in Sec. V.

A. Assumptions, control theory background, and past work

The key hypothetical assumptions are quadratic optimality and Gaussian endogenous noise. Optimality as an organizing




Fig. 1. Generalized Feedback Model. The plant P has control input v , measurement output y , disturbance input w , and performance output z . It is controlled by K .

principle for descriptions of animal behavior has proven useful and robust [1], [2]. Gaussian noise is a standard assumption with experimental justification [3]. The utility of quadratic cost models has been demonstrated in related sensorimotor contexts [4].

Linear Quadratic Gaussian (LQG) control is well known and widely used. We omit details not relevant here (see textbooks, e.g. [5]) and adapt notation. Given the following conditions:

- 1) a linear system P to be controlled as shown in Fig. 1, with a control input v , disturbance input w , measurement output y , and performance measure z (all generally vectors),
- 2) the objective of minimizing a quadratic form of the expected performance, $E(z(t)^T z(t))$ (in a simple scalar case this is often $E(y^2(t) + \rho v^2(t))$ with constant ρ being a design choice),
- 3) stationary white Gaussian disturbances w ,

then the optimal control v is given by the LQG feedback controller K . The LQG controller is the series combination of a Kalman filter (a Bayesian optimal estimator of the state of the system, see [5]) and a Linear Quadratic Regulator that sets v to be a linear function of the estimated state. An LQG controller operating at steady state to reject a stationary Gaussian disturbance is a linear time invariant (LTI) system. The Fourier transforms of input and output data of an LTI system are related in relatively simple ways to the system parameters (Sec. III), allowing for system identification.

Relevant past work for the purposes of this study is in both the sensorimotor and early controls literature. Optimality and Bayesian estimation in sensorimotor response are the subject of much recent work [6], [7], [8]. Todorov's 2004 survey of sensorimotor optimality [1] was exhaustive up to that date and in particular includes LQG applications [9], [10]. Past sensorimotor work has focused on invariant features of sensorimotor response rather than development of subject specific models, and on functionally relevant tasks rather than

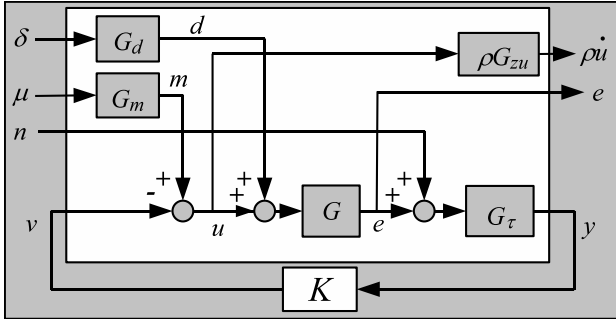


Fig. 2. Detailed view of the model described in Sec. II.

tasks design to elicit LTI response well suited to analysis. We are also interested in existing tests of neuromuscular health and skill, in particular those involving upper extremities. Among these, the Fugl-Meyer [11], the MAS Motor Assessment Scale [12], and the DASH [13] tests are prominent. Thirteen scales measuring upper body function are compared in [13]. In the field of control engineering, through the 1950s and 1960s a large body of work was produced on the behavior of pilots as surveyed in [14]. The most relevant work from this field is [15], [16], clarified by [17], which proposes an LQG model of pilot response. This is similar to the current paper in that human behavior was compared to that of a controller design based on LQG. However, the approach of [15], [16], and [17] is suboptimal in that motor noise autocorrelation is neglected during control synthesis. In addition, this paper describes a method to automatically fit data using subject-specific parameters, rather than presenting a single controller and noting that it resembles the aggregate behavior of three trained pilots.

II. HYPOTHESIZED MODEL

In this section we propose the model of the complete experimental feedback loop shown in Fig. 2 and describe its components. The plant P of Fig. 1 (the system to be controlled) and its inputs and outputs are decomposed as shown.

We turn first to partitioning the white disturbance vector w into an exogenous disturbance δ and endogenous vision and motor noises n and μ . The exogenous disturbance δ is a fixed sequence applied during the experiment. The vision noise intensity S_{nn} was set to yield noise with a standard deviation of a pixel width, as all subjects demonstrated an ability to read text small enough that visual acuity was not relevant. This intensity is sufficiently small that results are essentially insensitive to this value. Motor noise intensity S_{mm} is treated as a subject-specific parameter to be inferred. The noises are modeled as additive, but there is evidence that motor noise is multiplicative, that is, its amplitude scales with the magnitude of the corrupted signal [18], [19]. It is shown in [20] that optimal stationary control in the presence of multiplicative noise results in an alteration to the control cost, and we neglect non-stationary effects.

Two of the subsystems within the plant P are fixed by the experimental design, the plant G and disturbance dynamics

G_d . The plant G discussed in this paper is a single integrator

$$\dot{e}(t) = u(t) + d(t) \quad (1)$$

The disturbance filter G_d is a first order filter at 0.3 rad/sec. Limited disturbance bandwidth is necessary in order to make the task possible and this filter was chosen to yield reasonable difficulty and to allow direct comparison with [15]. The transfer functions for the subsystems of P fixed by the experiment are

$$G = 1/s, \quad G_d = 0.3/(s + 0.3) \quad (2)$$

Three subsystems of P model processes within the subject, and so must include assumed or inferred parameters. We have a motor noise filter G_m that shapes the spectrum of the motor noise, G_τ implementing a lumped delay, and G_{zu} implementing differentiation of the control input for inclusion in z . There is no filter on vision noise n because is hypothesized to be white within the frequency range of interest. The motor noise filter can be thought of as a musculoskeletal filter on a white noise disturbance. Pilot modeling studies typically represented musculoskeletal dynamics with first order lags at approximately 12 rad/s [15], [21], but second order actuator dynamics at approximately 17 rad/s are more consistent with results in [22], [23]. We used a second order Butterworth filter whose adjustable cutoff frequency in rad/s is denoted ω_m . This avoids the addition of a quantitative assumption at the expense of increased parametrization dimension. This approach yields the transfer function

$$G_m = \omega_m^2 / (s^2 + \omega_m \sqrt{2}s + \omega_m^2) \quad (3)$$

Delays of all sources are lumped into a single delay G_τ at the vision output without loss of generality. Delay cannot be modeled by a finite dimensional linear system. In order to avoid an intractable synthesis problem or an unnecessarily complicated two-stage state estimator, we approximate the delay with the well known Padé approximation [24]. For a typical 200 ms delay a fourth order approximation is accurate to < 1 degree of phase error at 20 rad/s, sufficient for our purposes. The value of the delay is a free parameter. The transfer function for a given delay τ is obtained with the standard formula given in Table 1 of [24]

$$G_\tau = \frac{1680 - 840\tau s + 180(\tau s)^2 - 20(\tau s)^3 + (\tau s)^4}{1680 + 840\tau s + 180(\tau s)^2 + 20(\tau s)^3 + (\tau s)^4} \approx e^{\tau s} \quad (4)$$

The performance vector z includes the error e and an approximation to $\rho \dot{u}$, the weighted derivative of the control signal u .

$$z = [e, \rho G_{zu} u] \approx [e, \rho \dot{u}] \quad (5)$$

Differentiation across the frequency range of interest is accomplished by G_{zu} . The value of the control weighting scalar ρ is a free parameter. Weighting the derivative is reasonable in the sense that we pay for movement, rather than for occupying an off-center location, and has a straightforward energy interpretation. This parameter reflects the level of exertion. The constants in the following transfer function set the frequency

range over which the differentiation is accomplished to be 0.01 to 100 rad/s, a sufficient range that results are completely insensitive to changes in these constants. These constants are required to avoid technical problems during LQG synthesis.

$$G_{zu} = \frac{s + 0.01}{0.01s + 1} \quad (6)$$

Our approach avoids modeling the nonlinearities and internal feedback loops in the human by using the assumptions of Sec. I-A to infer that the optimally performing central nervous system (CNS) should emit signals that cause the unified CNS/musculoskeletal system to have LTI input/output characteristics when controlling a linear plant. Linear system behavior throughout an operating range despite nonlinear components is a common goal in engineering, for example in the creation of linear amplifiers from nonlinear transistors by appropriate internal feedback loops. Experiments designed to isolate and identify delays and noise within these loops are described in [25].

In summary, our assumptions are: Gaussian disturbances with reasonable spectra, the existence of delay, and a control strategy based on minimizing a weighted sum of squared tracking error e and squared velocity at the hand \dot{u} . Model subsystems and disturbances are fixed except for the motor noise intensity, the control cost weighting scalar, the delay, and the cutoff frequency of the motor noise spectrum, compactly written as

$$\gamma := [S_{mm}, \rho, \tau, \omega_m] \quad (7)$$

These are the parameters we used to characterize the subjects.

III. METHODS

In this section we describe the experimental procedure and apparatus, the methods used to obtain spectral data, and the methods used to fit models to the experimental spectral data. Experimental data and the software used to execute the methods of this section are available at [26].

A. Experimental method

Subjects used a computer mouse mounted on a low-friction cart to provide input u that altered the behavior of a displayed cursor according to Eqn. (1). Their task was to minimize the displayed error e relative to a fixed desired cursor location.

The equipment consisted of an optical mouse input whose position was sensed as u , a hood, and a computer as shown in Fig. 3. The screen was basically empty except for a 230 mm long region within which the 2 mm wide and 15 mm tall rectangular cursor moved horizontally. The software sampled user input u , added disturbance d , and updated error e at $R = 100$ samples/s. The display monitor had a hood to ensure that subjects were undistracted and did not observe their hands. The Windows pointing cursor was hidden during the trial. The mouse was mounted on a custom ball-bearing cart to reduce static friction effects, and cursor enhancement nonlinearities in Windows were disabled. Subjects were healthy students at Cornell University between 20 and 32 years of age. They completed consent forms approved by




Fig. 3. Experimental setup with display, hood, and mouse input.




Fig. 4. Typical time series data: All subjects, trial 7. Note identical $d(t)$.

the University Committee on Human Subjects, brief health questionnaires, and Edinburgh handedness surveys. They then read the instructions:

Moving the mouse from side to side will affect how the error indicator moves. An unseen disturbance will also cause the error to move. TRY TO KEEP THE ERROR AS CLOSE TO ZERO AS POSSIBLE.

The subjects' interpretation of these instructions is an uncontrollable part of the experiment. Subjects were allowed to position themselves as they felt comfortable, as long as their elbow rested on the table. There were twenty 60 second trials. The first 9 seconds of data from each test were removed. Tests started every two minutes to allow for an average of 60 seconds rest. The first ten trials used the plant dynamics given in (1). Trials 11-14 used proportional dynamics in which $e(t) = u(t) + d(t)$, and 15-20 used a double integrator $\ddot{e}(t) = u(t) + d(t)$. This paper analyzes the single integrator results but the others are briefly qualitatively discussed in Sec. V. Subjects all had the same disturbance sequence on any given trial, as can be seen in Fig. 4. Each trial had a different disturbance. The variance of the displayed error as well as the values of the inferred parameters converged to typical behavior within the first three trials with the exception of subject 10.

The first five subjects received verbal cues and displayed instructions, and the last 11 solely had displayed instructions.

It became apparent that the presence or absence of a competitive attitude due to unintentional cues from the experimenter might alter the results. The data from the first five subjects was removed for this reason. The remaining sample had nine males and two females.

B. Obtaining spectral data

We obtain spectral data at discrete frequencies ω_k from time series $d(t_{i=1:6000}), u(t_{i=1:6000})$ using the Blackman-Tukey procedure [27] with the first nine seconds of data removed. For real valued signals the estimated correlation at lag k is given by

$$\hat{r}_{xy}(k) := \frac{1}{N} \sum_{i=0}^{N-1-k} x(t_i)y(t_{i+k}) \quad k = 0, 1, \dots, N-1 \quad (8)$$

and $\hat{r}_{xy}(-k) := \hat{r}_{xy}(k)$. Estimated cross spectra are given by

$$\hat{P}_{xy}(\omega_k) := \sum_{k=-(N-1)}^{N-1} W(k) \hat{r}_{xy}(k) e^{-j\omega_k k} \quad (9)$$

where ω_k is a vector of $2N-1$ frequencies evenly spaced from $-\pi(N-1)/N$ to $\pi(N-1)/N$ rad, and $W(k)$ is the smoothing window. We used a Parzen smoothing window of width $N/5$ [27]. Smoothing is a standard technique in spectral estimation used to reduce the effect of noise on the data. It yields an empirical estimate of the component of the response attributable to an LTI system.

The resolution of \hat{P}_{xy} scaled to continuous frequencies is approximately $2\pi R/(N/5) \approx 0.5$ rad/s [27]. This introduces a small amount of bias for a wide-band signal such as we encounter.

Un-smoothed data is denoted with bars and defined by simply removing the window as follows

$$\bar{P}_{xy}(\omega_k) := \sum_{k=-(N-1)}^{N-1} \hat{r}_{xy}(k) e^{-j\omega_k k} \quad (10)$$

This has one-fifth the bias but much higher variance due to noise.

For the purpose of comparing predicted and experimental closed loop behavior it is useful to define the following:

$$\bar{T}(\omega_k) := \bar{P}_{ud}(\omega_k) / \bar{P}_{dd}(\omega_k) \quad (11)$$

$$\hat{T}(\omega_k) := \hat{P}_{ud}(\omega_k) / \hat{P}_{dd}(\omega_k) \quad (12)$$

C. LQG parameter inference method

In this subsection we show how to obtain the optimal controller K for a parameter set γ from Eqn. (7), how to use this K to make predictions that may be compared to spectral data, and how this is iteratively used to infer a parameter set that fits the observed data.

For a given parameter set γ , the plant P and the statistics of the disturbance w of Fig. 1 are fully defined. Therefore, given LQG optimality and γ , K is fully defined [5]. The LQG synthesis process is automated in commercial software [28] and briefly presented here for completeness. By the separation principle [5] LQG synthesis can be treated as two separate

problems, a deterministic Linear Quadratic Regulator (LQR) problem and a Kalman filter state estimation problem. We solve the “output weighted cost function” steady state Linear Quadratic Regulator problem [29] and the Kalman Filter synthesis problem, each of which involves the solution of an Algebraic Riccati Equation (ARE). We thereby obtain the state equation of the LQG optimal controller K corresponding to the parameter set γ .

After K is defined the entire system is defined and expected spectral properties may be computed. We work in continuous time and frequency and derive asymptotic results for infinite duration experiments. Operators characterizing input-output behavior are assumed to be LTI and are Laplace transformed [5]. Without loss of generality, we will only be interested in behavior along the imaginary axis and therefore to simplify notation all transfer functions and variables are assumed to be functions of $j\omega$. Asterisks denote complex conjugates. Let $K_H := KG_r$ and define the sensitivity gain $S := 1/(1 + GK_H)$ and complex sensitivity $T := 1 - S$. It can be shown that

$$u = S[G_m\mu - K_H Gd - K_H n] \quad (13)$$

$$e = S[G G_m\mu + Gd - G K_H n] \quad (14)$$

For infinite time the periodogram estimate of the cross spectra is unbiased [27] and we can express the expected power spectrum of u as $P_{uu} := E(uu^*)$. The variables μ and n are zero mean, Gaussian, and uncorrelated and therefore all cross terms have an expected value of zero. The quantities $E(\mu\mu^*)$ and $E(nn^*)$ are the white noise intensities S_{mm} and S_{nn} , yielding

$$P_{uu} := |S|^2 [|G_m|^2 S_{mm} + |G K_H|^2 d d^* + |K_H|^2 S_{nn}] \quad (15)$$

We can similarly compute the expected cross spectrum of u and d , normalized by $-dd^*$ in order that its expectation will conveniently be T .

$$-\frac{E(ud^*)}{dd^*} = S G K_H = T \quad (16)$$

For comparison to experimental data we evaluate the expected P_{uu} from Eqn. (15) and T from (16) at the experimental frequencies $R\omega_k$ where R is the sample rate of 100 samples/sec.

The above technique was applied iteratively in order to obtain our estimate of γ . We started by guessing a parameter set γ , computing K , and finding expected spectral data. These expected quantities were compared to the experimentally observed $\hat{T}(\omega_k) := -\hat{P}_{ud}/\hat{P}_{dd}$ and $\bar{P}_{uu}(\omega_k)$ from Sec. III-B. We used a commercially available [28] Nelder-Mead optimization function to repeat this process systematically to find a parameter set γ that minimized the following objective function.

$$\min_{\gamma} \sum_{k=A}^B (|\omega_k T(R\omega_k, \gamma) - \omega_k \hat{T}(\omega_k)|^2 + 10|\omega_k^2 P_{uu}(R\omega_k, \gamma) - \omega_k^2 \bar{P}_{uu}(\omega_k)|) \quad (17)$$

We used the smoothed \hat{T} because it reduced the variance of the inferred parameters. No advantage was found to using smoothed \bar{P}_{uu} , and in fact data is actually better fit without smoothing due to precise knowledge of the disturbance d . The




Fig. 5. The observed and fitted closed loop transfer function T for one subject on one trial. Squares denote the smoothed \hat{T} (see Sec. IV), and lines the predicted T from our method.




Fig. 6. Observed and fitted P_{uu} for one subject on one trial. Circles denote the un-smoothed \hat{P}_{uu} (see Sec. IV), and lines the predicted P_{uu} from our method. Note that this is not a transfer function. The jagged appearance of the predicted spectrum is due to the jagged spectrum of the known excitation d , and does not represent overfitting.

summation limits A and B were chosen to limit the frequency range of interest to be 0.3 to 20 rad/s. This range was chosen because lower frequency cross spectral data is unreliable, and above this frequency range the response is dominated by noise. The frequency weight emphasizes higher frequencies where the effects of the parameters to be estimated γ are strongly expressed, reducing intra-subject variance. The factor of 10 is used to avoid over-fitting T at the expense of P_{uu} . This constant may be varied from 3 to 30 without substantially altering the results. In particular, mean inferred properties are insensitive to this constant. At low values of this constant, the variance of the inferred motor noise properties increases, and at higher values outlying inferred delays are observed.

IV. RESULTS AND ANALYSIS

In this section we show the results of the fitting process, verify the presumed optimality, benchmark our method against a standard linear system ID method, and analyze the inferred

parameters. Performance in trials repeated eight months later is used to demonstrate that subjects exercised their ability to alter the level of effort.

A. Results of the fitting process

Figs. 5 and 6 are the result of the fitting process on a subject trial. The method is able to fit behavior well despite using only four parameters. When the response is dominated by the effects of the known disturbance, the power spectrum of Fig. 6 can be fit with great accuracy. The majority of fits were of similar quality, with poor fits often seen with subject 1 and 10. The agreement of the empirically obtained cross-spectral estimates with the LQG fitted model confirms that the component of the response due to d is well modeled.

Typical inferred loop gain magnitudes for all subjects are shown in Fig. 9. The inferred loop gain phases are in excellent agreement with results based only on empirical smoothing per Eqn. (9), with an example given in Fig. 9. The change in the slope of the loop gain magnitude as one progresses from low to high frequency can be understood as an optimal response to delay and is in agreement with classical approaches to the control of plants with delay [30]. Accurately modeling this aspect of response requires high order models. With the LQG parametrization this is accomplished without loss of parsimony in the parametrization.

B. Evidence for optimality in most subjects

The most basic prediction of the hypothesized optimality is that subjects will have the freedom to trade effort against performance. Under the assumption that all young subjects are in similar condition, they should therefore fall along a Pareto front in a plot of performance against exertion, with the location on the front parametrized by control cost ρ . This was observed for all but subject 10. The RMS tracking error e normalized by the RMS perturbation d is plotted against RMS input velocity du/dt in Fig. 7. In conjunction with good agreement between the spectral properties of the fitted LQG controllers and the observed behaviors, these results support the hypothesis that LQG optimality is an invariant feature of typical healthy proficient human visuomotor response in a simple feedback task. This occurs despite the absence of any instructions regarding u .

The hypothesized optimality is also supported for all but subjects 1 and 10 by comparing the fitted models' predicted RMS tracking error e with experimental observations as shown in Fig. 8. This quantity is meaningful in that the fitting method makes no use of e , but given a fitted model, the RMS e may be predicted. We interpret this quantity as follows: If the power spectrum of u and the cross spectrum of u and d are well fit, the response has been separated into that attributed to the effects of endogenous noises and that attributed to an optimal response to the known disturbance d in the presence of these endogenous noises. That is, the best fit to an optimal model is obtained, and the remaining components of u are attributed to noise. Under the hypothesized model the effects of noises m and n should increase RMS e by a predictable amount. If the response not attributed to an optimal model serves to reduce




Fig. 7. Normalized RMS error e against RMS velocity of input u . Note that subjects fall along a Pareto optimal curve except for subjects 10 and possibly 4.




Fig. 8. This boxplot [31] gives the ratio of the predicted RMS displayed error e from the fitted model to the observed value. Values near one are consistent with a well modeled trial. See Sec. IV-B for details.




Fig. 9. The upper plot gives the inferred loop gain magnitudes $|GK_H|$ for all subjects on trial 7. The solid line in the lower plot gives the inferred loop gain phase for the subject displayed with a bold dotted line in the upper plot. The dashed line in the lower plot is the empirical estimate of experimental loop phase, obtained using $\hat{T}/(1 - \hat{T})$. This agreement is typical with the exception of outlying subject 1. Empirical estimates of experimental loop gain magnitude are omitted because they are biased in the presence of noise.

e and is thus not well described as noise, the model will over-predict the RMS e . In this case, the component of response not fit to an optimal control model is systematic and functional. If the response not attributed to an optimal model increases e by more than the amount expected for the hypothesized noise, the model will under-predict the RMS e . In this case, the component of response not fit to an optimal control model is systematically dysfunctional. Recall that these interpretations of the data are contingent on a good fit to the power spectrum of u .

Observed RMS e was generally consistent with that predicted, supporting the hypothesized optimality. Experimental data was noisy because low frequency (< 2 rad/s) components dominate e , and tests were limited to 60 seconds to avoid fatigue. Unusual behavior by subjects 1 and 10 highlight behaviors not consistent with the optimal control model, and the criterion of a good fit to the power spectrum of u was violated. In the case of subject 10, the RMS e was over-predicted. The model did not fit power spectrum of u well in the neighborhood of 1-3 rad/s. In the case of subject 1, experimental power spectra of u had large components in the 2-10 rad/s range that could not be fit to an optimal model.

C. Comparison to N4sid

In this subsection we assess the quality of fit of our technique in the class of system identification techniques yielding linear models. We benchmark our method against a current standard, N4sid [32]. N4sid is a well regarded general purpose black box system ID method. We establish that our technique is as good or better than this standard technique because it suggests the non-existence of linear models substantially better than those of this paper.

We obtained N4sid estimates of T of order 1, 2, 4, and 7 for each data set's u and d . A seventh order model was used because the result of our method can be shown to have at most seven significant states (see Sec. V). The following cost was evaluated within each trial:

$$\sum_{k=A}^B (|T(R\omega_k) - \hat{T}(\omega_k)|^2) \quad (18)$$

Our method had mean decreases in cost of 65%, 44%, 45%, and 30% relative to N4sid. This was accomplished despite a 25, 63, and 79% reduction in the dimension of the parametrization for orders 2, 4, and 7, and a 50% increase for the first order fit. We also tried using un-smoothed data \bar{T} in the cost function

$$\sum_{k=A}^B (|T(R\omega_k) - \bar{T}(\omega_k)|^2) \quad (19)$$

In this case the results are dominated by noise, so that the improvements due to our fit are only 6%, 3%, 3%, and 1%. We conclude that if the user only seeks a model of closed loop behavior such that noise is the principal source of un-modeled response, our methods offer slight advantage. Our method is advantageous when one seeks to infer physiologically meaningful quantities. This calls for precise and structured




Fig. 10. Inferred delay. Lateral displacement of markers within a subject indicates chronological order.

identification of the deterministic component of the response regardless of the relative amount of noise.

The similarity of Eqn. (18) to the Eqn. (17) used during fitting may be questioned. We fit the model using smoothed data in order to reduce the variance of inferred parameters. If we use un-smoothed data during the fitting process and then evaluate Eqn. (18) we still obtain mean decreases in cost of 58%, 32%, 32%, and 13% relative to N4sid.

D. Inferred properties

The technique yields three inferred quantities that are repeatable across trials 3-10 for most subjects, and vary significantly across subjects as shown in Figs. 10, 11, and 12. Inferred motor noise bandwidths varied widely and no sensible interpretation is apparent.

In two subjects, significant anomalies in inferred or observed quantities were coincident. Mean inferred noise levels were large in subject 10. Subject 10 is also off the Pareto optimality front in Fig. 7. The highly variable delays inferred for subject 1 are believed to result from abnormally low bandwidth creating a lack of information at higher frequencies where delay is salient. The low bandwidth is associated with low RMS control input velocities \dot{u} and high mean control costs ρ .

E. Repeated testing

Repeated testing was performed on eight of the eleven subjects including outlier subjects 1 and 10. Both outlier subjects behaved in less unusual ways during repeated testing, and results were better fit. We determined whether the subjects had significantly changed their behavior using multivariate ANOVA [31] with two data groups per subject, where each data group contains the inferred properties in γ . Four out of eleven total subjects and three out of nine optimal subjects exhibited significant differences in their inferred parameters with a false alarm level of 0.05. The most salient difference was in the inferred control cost ρ , as plotted in Fig. 13. Performing t-tests on each subject's two sets of inferred ρ ,




Fig. 11. Inferred control cost. Lateral displacement of markers within a subject indicates chronological order. Squares indicate the last trial. All but subject 7 had minimal control cost on their last trial, often to a degree that makes the trial an outlier. However, because all subjects had the same disturbance sequence on any given trial, no conclusions can be drawn.

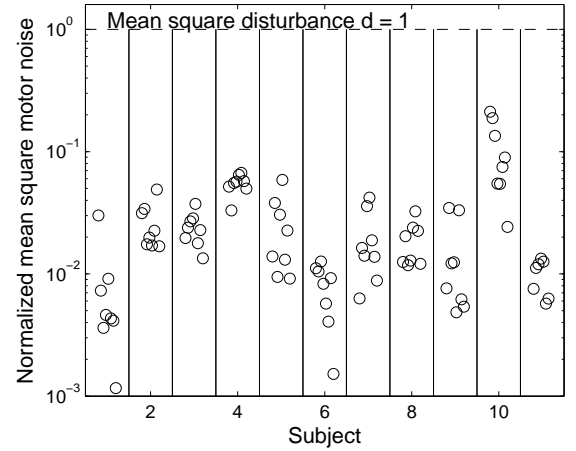


Fig. 12. Normalized inferred motor noise power levels. Lateral displacement of markers within a subject indicates chronological order.

the same three out of nine optimal subjects have significantly different ρ at a false alarm rate of 0.05. This indicates that an altered willingness to expend effort is common and detectable. Differences in delay were negligible. Differences in motor noise intensity were small with the exception of subjects 1, 4, and 10, which shifted towards typical levels.

The two significantly suboptimal subjects behaved in a way more consistent with optimality during the repeated test as shown in Fig. 14.

V. DISCUSSION

Results were consistent with the LQG model for nine of eleven subjects, and outlying behavior decreased in repeated testing. The accuracy of the LQG model is demonstrated by comparing the fit to that obtained by model-free empirical fitting of spectral data in Figs. 5 and 6, and to established model fitting techniques in Sec. IV-C. We demonstrated that fitting a model with four LQG parameters instead of fourteen




Fig. 13. Inferred control cost for repeated trials.




Fig. 14. This boxplot [31] gives the ratio of the predicted RMS displayed error e from the fitted model to the observed value in the repeated trials. Values near one are consistent with a well modeled trial. Sec. IV-B gives an explanation of the meaning of this ratio.

parameters carries no loss of accuracy, and gives a large gain in parsimony.

Variance in sensorimotor control has traditionally been approached in terms of variability in time histories for repetitive tasks. Optimal response to noise has been proposed as a means of explaining its origins and expression [33], [18]. In the traditional approach, the control strategy itself is treated as invariant. We differ significantly with [15], [16] in that we do not find the control strategies themselves to be invariant. We also do not assign musculoskeletal meaning to control cost ρ , something tentatively proposed in [15], [16]. In contrast, our method addresses variations in control strategies that correspond to altered levels of expended effort, that is, variance in the objective function of the optimality. From this perspective it is optimality itself that is invariant. Indeed optimality was found to be an invariant feature of response for nine of eleven young healthy subjects, while significant variance in the objective function was observed in four of the nine optimal subjects upon repeated testing. We conclude that the invariant features of response by nearly all subjects in this

task are delay and optimality itself.

Inferred parameters were significantly subject specific and physiologically reasonable. The inferred delays are typically somewhat larger than delays found in simple reaction time studies [34]. This experiment differs from those studies in that the task is more complex and delay is inferred as a separate phenomenon from musculoskeletal lag. The spectral techniques of [35] are able to separate lag and delay in continuously perturbed postural control, and inferred delays are comparable to ours.

All subjects with all plants had a “dead band”, that is, a preference for remaining motionless that is not predicted by the hypothesized optimality. The relative significance of this effect can be assessed by inspecting time domain data in Fig. 4. Our method effectively treats this tendency as a source of noise, that is, an aspect of behavior that can not be fit to an optimal control strategy. Thus it is assumed to not contribute to disturbance rejection. The utility of this approach is demonstrated through prediction of the RMS displayed error e in Sec. IV-B.

The method was also applied to other plants G . The method was able to fit data with G altered to $e(t) = u(t) + d(t)$ and similar noise powers and delay were observed. It was not able to consistently fit $\ddot{e}(t) = u(t) + d(t)$, and inferred parameters for the \ddot{e} plant were erratic and unrealistic. Additional non-Gaussian data inconsistent with linear models of behavior with the \ddot{e} plant is found in [21]. Dead band behavior was more pronounced. All of this is consistent with the time-optimal bang-bang behavior observed by [36]. It cannot be ruled out that with more practice and possibly coaching people might come to resemble LQG. The results of [15] and [16] with the double integrator do not include the full frequency range given for other plants and are given for fewer subjects. Another abnormality with the \ddot{e} plant was that before becoming skilled, subjects entered into input-saturating limit cycles resembling category II pilot induced oscillation (PIO) as described in [37].

A key measure of parametrization quality is parsimony, the ability to characterize processes of large order or dimension with few parameters. We claim that parsimony is a matter of parametrization dimension rather than state dimension. For example, parsimonious and precise models of beams, heat conduction, fluid mechanics, gas dynamics, radiation, etc. are of infinite order and perfectly parametrized by small sets of constants. Methods of modeling such systems in lumped parameter form can be of arbitrary order, and the problems associated with order are computational and do not reflect a lack of parsimony. We are far from model orders that would cause computational problems. The optimal control parametrization allows us to fit a model with only four parameters, far fewer than a general linear model of similar accuracy. Alternatively, the advantage can be exploited by fitting the data better than a general purpose method yielding equal or greater parameter dimension. Both advantages were demonstrated by comparison to N4sid in Sec. IV-C.

The large number of states used to represent the system in Fig. 2 may be reduced. This is of computational interest but is irrelevant to parsimony as discussed previously. Inspection of the Hankel Singular Values (HSVs) [5] of the estimate

of K_H shows a mean normalized sixth HSV of 0.19 and a negligible seventh HSV. This indicates that despite the large order of the model used in the fitting method, its input-output characteristics are almost perfectly approximated by a sixth order system.

A. Engineering context

The methods of this paper perform closed loop identification of systems hypothesized to contain optimal feedback controllers. The advantage is parsimony.

For the purpose of modeling operator behavior the LQG feedback model works well except for the case of a double integrator with untrained novice subjects. The large and statistically significant inferred parameter variations observed across subjects caution against generalizations from studies based on pilots. Technical improvements are made within the process of LQG controller synthesis as compared to the approach described in [15], [16], and [17]. First, the frequency weighting technique presented here avoids observability [5] problems. Second, the approach in [15], [16], and [17] is not truly optimal in that LQR synthesis is performed in a way that neglects dynamics later attributed to m .

B. Future work and acknowledgements

Having demonstrated the methods on healthy young subjects, a follow-on study comparing them to healthy aged subjects is underway. We are also interested in analyzing the effects of Parkinson's disease. In particular, we seek to provide a quantitative assessment of the bradykinesia (sluggishness) currently used to diagnose incipient Parkinson's.

We thank Francisco Valero-Cuevas, Madhu Venkadesan, Oliver Purwin, and Simone Bortolami, as well as the reviewers for helpful comments. This work was supported by an NSF Graduate Student Fellowship.

All data and code used to generate the results of this paper is available to reviewers [26].

REFERENCES

- [1] E. Todorov, "Optimality principles in sensorimotor control," *Nat. Neuroscience*, vol. 7, no. 9, pp. 907–915, 2004.
- [2] R. M. Alexander, *Principles of Animal Locomotion*. Princeton, NJ, USA: Princeton University Press, 2003.
- [3] H. P. Clamann, "Statistical analysis of motor unit firing patterns in a human muscle," *Biophysical Journal*, vol. 9, pp. 1233–1251, 1969.
- [4] K. P. Koerding and D. M. Wolpert, "The loss function of sensorimotor learning," *PNAS*, vol. 101, no. 26, pp. 9839–9842, 2004.
- [5] S. Skogestad, *Multivariable Feedback Control*. West Sussex, England: Wiley, 2005.
- [6] K. P. Koerding and D. M. Wolpert, "Bayesian decision theory in sensorimotor control," *Trends in Cognitive Sciences*, vol. 10, no. 7, pp. 319–326, 2006.
- [7] E. Todorov, "Stochastic optimal control and estimation methods adapted to the noise characteristics of the sensorimotor system," *Neural Computation*, vol. 17, pp. 1084–1108, 2005.
- [8] D. Wolpert and Z. Ghahramani, "Computational principles of movement neuroscience," *Nat. Neuroscience*, vol. 3, pp. 1212–1217, 2000.
- [9] J. He, W. S. Levine, and G. E. Loeb, "Feedback gains for correcting small perturbations to standing posture," *IEEE Trans. Automat. Contr.*, vol. 36, no. 3, pp. 322–332, 1991.
- [10] A. D. Kuo, "An optimal control model for analyzing human postural balance," *IEEE Trans. Biomed. Eng.*, vol. 42, no. 1, pp. 87–101, 1995.
- [11] A. R. Fugl-Meyer, L. Jsk, I. Leyman, S. Olsson, and S. Steglind, "The post-stroke hemiplegic patient. 1. a method for evaluation of physical performance," *Scandinavian Journal of Rehabilitative Medicine*, no. 1, pp. 13–31, 1975.
- [12] J. H. Carr, R. B. Shepherd, L. Nordholm, and D. Lynne, "Investigation of a new motor assessment scale for stroke patients," *Physical Therapy*, vol. 65, no. 2, pp. 175–180, 1985.
- [13] P. L. Hudak, P. C. Amadio, and C. Bombardier, "Development of an upper extremity outcome measure: the dash (disabilities of the arm, shoulder and hand)," *American Journal of Industrial Medicine*, vol. 29, no. 6, pp. 602–608, 1996.
- [14] D. McRuer, "Human dynamics in man-machine systems," *Automatica*, vol. 16, no. 3, pp. 237–253, 1980.
- [15] D. Kleinman, S. Baron, and W. H. Levison, "A control theoretic approach to manned-vehicle systems analysis," *IEEE Trans. Automat. Contr.*, vol. AC-16, no. 6, pp. 824–832, Dec. 1971.
- [16] —, "An optimal control model of human response," *Automatica*, vol. 6, pp. 357–369, 1970.
- [17] J. B. Davidson and D. K. Schmidt, "Modified optimal control pilot model for computer-aided design and analysis," NASA, Tech. Rep. TM-4384, 1992.
- [18] C. M. Harris and D. M. Wolpert, "Signal-dependent noise determines motor planning," *Nature*, vol. 394, no. 6695, pp. 780–784, 1998.
- [19] R. A. Schmidt, H. Zelaznik, B. Hawkins, J. S. Frank, and J. T. Quinn, "Motor-output variability: A theory for the accuracy of rapid motor acts," *Psychological Review*, vol. 86, no. 5, pp. 415–450, 1979.
- [20] D. Kleinman, "Optimal stationary control of linear systems with control-dependent noise," *IEEE Trans. Automat. Contr.*, vol. AC-14, no. 6, pp. 824–832, Dec. 1969.
- [21] D. McRuer, "Mathematical models of human pilot behavior," AGARD, Tech. Rep. AG-188, 1974.
- [22] A. Fagergren, O. Ekeberg, and H. Forssberg, "Precision grip force dynamics: A system identification approach," *IEEE Trans. Biomed. Eng.*, vol. 47, no. 10, pp. 1366–1375, 2000.
- [23] F. A. Mussa-Ivaldi, N. Hogan, and E. Bizzi, "Neural, mechanical, and geometric factors subserving arm posture in humans," *Neuroscience*, vol. 5, no. 10, pp. 2732–2743, 1985.
- [24] M. Vajta, "Some remarks on pade-approximations," in *Proc. Third TEMPUS-INTCOM Symp.*, 2000.
- [25] M. Vendadesan, J. Guckenheimer, and F. J. Valero-Cuevas, "Manipulating the edge of instability," *Journal of Biomechanics*, vol. 40, no. 8, pp. 1653–1661, 2007.
- [26] M. Sherback. (2006) Supporting materials. [Online]. Available: <http://control.mae.cornell.edu/sherback/HITL/index.htm>
- [27] S. Kay, *Modern Spectral Estimation*. Englewood Cliffs, NJ, USA: Prentice-Hall, 1988.
- [28] *MATLAB Version 7.2*. Natick, MA, USA: The Mathworks Inc., 2006.
- [29] R. F. Stengel, *Optimal Control and Estimation*. New York, NY: Dover, 1986.
- [30] J. Doyle, B. Francis, and A. Tannenbaum, *Feedback Control Theory*. Macmillan, 1990.
- [31] J. Neter, M. Kutner, C. J. Nachtsheim, and W. Wasserman, *Applied Linear Statistical Models, 4th Edition*. Boston, MA, USA: WCB McGraw-Hill, 1996.
- [32] P. V. Overschee and B. D. Moor, "N4sid: subspace algorithms for the identification of combined deterministic-stochastic systems," *Automatica*, vol. 30, no. 1, pp. 75–93, 1994.
- [33] E. Todorov and M. Jordan, "Optimal feedback control as a theory of motor coordination," *Nat. Neuroscience*, vol. 5, no. 11, pp. 1226–1235, 2002.
- [34] R. A. Schmidt, *Motor Control and Learning*. Champaign, IL: Human Kinetics Publishers, Inc., 1982.
- [35] R. J. Peterka, "Sensorimotor integration in human postural control," *J. Neurophysiology*, vol. 88, no. 10, pp. 1097–1118, 2002.
- [36] H. Inooka and T. Koitabashi, "Experimental studies of manual optimization in control tasks," *IEEE Control Syst. Mag.*, vol. 10, no. 5, pp. 20–23, 1990.
- [37] D. McRuer, "Pilot-induced oscillations and human dynamic behavior," NASA, Tech. Rep. CR-4683, 1995.