



Pages 1–382

JOURNAL OF NEUROPHYSIOLOGY

JULY 2009 VOL. 101 NO. 1

Journal of Neurophysiology



Published by the American Physiological Society
Volume 101 Number 1 July 2009

Structured Variability of Muscle Activations Supports the Minimal Intervention Principle of Motor Control

Francisco J. Valero-Cuevas,^{1,3,*} Madhusudhan Venkadesan,^{3,4,*} and Emanuel Todorov^{5,6,*}

¹Department of Biomedical Engineering, University of Southern California and ²Division of Biokinesiology and Physical Therapy, University of Southern California, Los Angeles, California; ³Sibley School of Mechanical and Aerospace Engineering and ⁴Department of Mathematics, Cornell University, Ithaca, New York; ⁵Departments of Applied Mathematics and Computer Science and Engineering, University of Washington, Seattle, Washington; and ⁶Department of Cognitive Science, University of California San Diego, La Jolla, California

Submitted 1 March 2008; accepted in final form 8 April 2009

Valero-Cuevas FJ, Venkadesan M, Todorov E. Structured variability of muscle activations supports the minimal intervention principle of motor control. *J Neurophysiol* 102: 59–68, 2009. First published April 15, 2009; doi:10.1152/jn.90324.2008. Numerous observations of structured motor variability indicate that the sensorimotor system preferentially controls task-relevant parameters while allowing task-irrelevant ones to fluctuate. Optimality models show that controlling a redundant musculo-skeletal system in this manner meets task demands while minimizing control effort. Although this line of inquiry has been very productive, the data are mostly behavioral with no direct physiological evidence on the level of muscle or neural activity. Furthermore, biomechanical coupling, signal-dependent noise, and alternative causes of trial-to-trial variability confound behavioral studies. Here we address those confounds and present evidence that the nervous system preferentially controls task-relevant parameters on the muscle level. We asked subjects to produce vertical fingertip force vectors of prescribed constant or time-varying magnitudes while maintaining a constant finger posture. We recorded intramuscular electromyograms (EMGs) simultaneously from all seven index finger muscles during this task. The experiment design and selective fine-wire muscle recordings allowed us to account for a median of 91% of the variance of fingertip forces given the EMG signals. By analyzing muscle coordination in the seven-dimensional EMG signal space, we find that variance-per-dimension is consistently smaller in the task-relevant subspace than in the task-irrelevant subspace. This first direct physiological evidence on the muscle level for preferential control of task-relevant parameters strongly suggest the use of a neural control strategy compatible with the principle of minimal intervention. Additionally, variance is nonnegligible in all seven dimensions, which is at odds with the view that muscle activation patterns are composed from a small number of synergies.

INTRODUCTION

The nature of and causes for asymmetries in variability during motor performance is an integral part of the study of motor redundancy. This line of research has been fruitful in proposing several theories for computational mechanisms underlying sensorimotor function such as the uncontrolled manifold, minimal intervention, and optimal feedback control. A characteristic feature of these proposed neural control strategies is the presence of larger variance in task-irrelevant directions for a wide range of motor behaviors (Bernstein 1967; Li et al. 1998; Scholz and Schoner 1999; Todorov 2004; Todorov and Jordan 2002). This

observed asymmetry in motor variability is often quantified using the “uncontrolled manifold” method for comparing task-relevant and -irrelevant variance (Scholz and Schoner 1999). For example, there is compelling evidence that the nervous system uses a strategy that compensates for variability in forces produced by individual fingers to reduce variability in the task-relevant target of total force (e.g., Latash et al. 2001, 2002; Scholz and Schoner 1999). Such asymmetry in variability is also predicted by the “minimal intervention” principle whereby the neural controller corrects deviations from the average behavior only when they interfere with the task goals (Liu and Todorov 2007; Todorov 2004; Todorov and Jordan 2002). Although these neural control strategies are demonstrated in behavioral measures and not electrophysiological measures such as EMG (e.g., fingertip forces or during shifts in the center of pressure of a standing subject), the origin of behavioral variability presumably lies in muscle forces and EMGs.

Given that previous studies were mostly restricted to behavioral and kinematic measurements, a direct link to physiology remains to be made. The only exception in this regard is an uncontrolled manifold analysis of postural muscle activity during shifts in the center of pressure of a standing subject (Krishnamoorthy et al. 2003). However, that electromyographic (EMG) analysis was performed by first projecting the EMG data to a low-dimensional subspace, and so it remains unclear whether the full covariance in muscle activity was consistent with the hypothesis of task-relevant control. Furthermore, by not recording from all the muscles of the system, the nature and properties of the full motor command and the subspaces wherein its variability resided could not be studied. These limitations arise because it is practically impossible to record from all muscles involved in a sit-to-stand motion. Therefore we chose to examine the physiological causes of variability in motor performance in a simpler task in which all relevant muscle activities could be measured.

It is important to note that even when motor variability is convincingly shown to be structured, it can originate without any task-relevant control or it can be structured for a number of reasons. We now underscore several previously unaddressed confounds unrelated to task-relevant control. 1) Musculo-skeletal coupling, especially in the tendons of the hand (Valero-Cuevas et al. 1998, 2007), can induce complex correlations on the behavioral level without correlated drive to individual muscles. We avoid this confound by recording muscle activity, as approximated by fine-wire EMGs. 2) Motor noise is known

*All authors contributed equally to this work.

Address for reprint requests and other correspondence: F. J. Valero-Cuevas, The University of Southern California, 3710 McClintock Ave, RTH 404, Los Angeles, California 90089-2905 (E-mail: valero@usc.edu).

to be signal-dependent (Harris and Wolpert 1998; Sutton and Sykes 1967) and can therefore create structure in the variability that does not directly reflect the control law. For example, endpoint errors in reach are larger in the movement direction (Gordon et al. 1994) not because that direction is task-irrelevant but because muscles pulling along the movement axis are more active and therefore more strongly affected by signal-dependent noise. Here we rule out such confounds by showing that a signal-dependent noise model does not capture the variability pattern in our experimental data. 3) The motor system may purposefully vary task-irrelevant aspects of the movement from trial to trial, so as to minimize fatigue or explore different control strategies. Such trial-to-trial variability can inflate measures of task-irrelevant variability without having any origins in the control strategy. This type of confound is avoided here by analyzing the moment-to-moment fluctuations in motor output within a trial.

In this work we use fine-wire electrodes to simultaneously record the electrical activity in all muscles of a finger to test the hypothesis that the structure in the variability of muscle activations reduces variability in a task-relevant manner.

METHODS

Experimental design and data recording

Eight healthy volunteers (age 18–27 yr, 5 male, 3 female) first performed a finger-tapping experiment unrelated to the present study (Venkadesan and Valero-Cuevas 2008) followed by the experiment described here. This study was approved by Cornell University's Committee on Human Subjects. All analyses were performed using MATLAB (The MathWorks, Natick, MA). As shown in Fig. 1A, subjects grasped a metallic dowel fixed to ground with their thumb, middle, ring, and little finger and placed the fingertip of the index finger on center of the recording surface of the rigidly held force sensor with the distal phalanx in a vertical orientation. The force sensor was held by a robotic arm (AdeptSix 300, Adept Technologies) and adjusted for each subject to replicate the finger posture of 30° of flexion at the metacarpophalangeal joint, 30° of flexion at the proximal interphalangeal joints, and 15° of flexion of the distal phalangeal joints measured with a clinical goniometer. Defining the location of the hand and fingertip in three dimensions, plus the orientation of the distal phalanx in this way, suffices to replicate the posture of the fingers across trials in this isometric task.

EMGs

Fine-wire intramuscular electrodes were placed in all seven muscles acting on the index finger as described elsewhere (Burgar et al. 1997). The seven muscles actuating the index finger are flexor digitorum profundus, flexor digitorum superficialis, extensor indicis proprius, extensor digitorum communis, first lumbrical, first dorsal interosseous and first palmar interosseous. We recorded and digitally processed EMG using fine-wire intramuscular electrodes from all seven muscles using previously reported techniques (Valero-Cuevas et al. 1998; Venkadesan and Valero-Cuevas 2008). Amplified EMGs were each sampled at 2,000 Hz and, to avoid aliasing, band-pass filtered (20–800 Hz) using the filters in the preamplifiers (Neurodata Amplifier System, Model 15A, Grass-Telefactor, West Warwick, RI) before storing them for processing as described in the following section. Maximal voluntary contractions (MVC) of individual muscles were done immediately before and after fingertip force production with the index finger braced in the same posture used during the study. As in prior work, we asked subject to perform three maximal voluntary contraction trials to maximally activate each muscle while we provided EMG feedback as a sound delivered through speakers with

increasing sound level indicating an increasing level of EMG amplitude from the target muscle. Activating muscles maximally in isolation can be difficult for subjects given the complexity of the hand, and obtaining maximal EMG activity in each muscle does not necessitate eliminating EMG activity from other muscles. All subjects reported that their contractions were maximal, and we looked through all MVC trials to find the maximal activation of each muscle in any task as in prior work (Burgar et al. 1997; Valero-Cuevas et al. 1998; Venkadesan and Valero-Cuevas 2008).

Signal processing

Our analysis required estimating the instantaneous tension generated by each muscle that was then used to estimate the contribution of that muscle to fingertip force. Using either the raw or full-wave rectified and unit-normalized EMG voltage by itself would distort the relationship between that estimated neural command and fingertip force because small and large muscles appear to contribute equally to fingertip force. Therefore we focused on estimating “normalized muscle tension” (i.e., the percentage of maximal force each muscle is producing), where muscle tension is known to be a low-pass filtered version of EMG due to activation-contraction dynamics (Zajac 1989). To arrive at estimates of normalized muscle tension, the raw EMG voltages collected at 2,000 Hz and band-pass filtered (20–800 Hz) were full-wave rectified then each channel was normalized by the peak value of the band-passed EMG voltage recorded during MVC trials of that muscle. This normalization is an accepted means to compensate across EMG channels for differences in amplifiers setting, proximity to the active motor unit pools, electrode impedance, etc. It provides a signal that is scaled to the maximal output of that muscle (Burgar et al. 1997; Valero-Cuevas 2000; Valero-Cuevas et al. 1998; Venkadesan and Valero-Cuevas 2008). To emulate the low-pass filtering effect of activation-contraction dynamics, we then passed the full-wave-rectified, band-pass filtered, normalized EMG signals through a fourth-order Butterworth causal filter with time constants 0.03 and 0.23 s. These two time constants were found by an algorithm minimizing the mismatch between measured fingertip forces and fingertip forces predicted from the processed EMGs (the prediction method is described in the following text). We also used cross-correlation to analyze the time lag of the force predictions and found that the lag was indistinguishable from zero for the chosen time constants. Note that the 0.23 s time constant is unusually large for isolated muscles. However, finger muscles tend to have long tendons relative to their fiber lengths, which are known to increase the effective time constant of the musculo-tendon actuator (Zajac 1989). We then down-sampled from 2,000 to 100 Hz by averaging within nonoverlapping bins of 0.01 s and expressed in % MVC (Fig. 1C). The fingertip force data were also down-sampled to 100 Hz by averaging within 0.01 s bins. Henceforth “EMG” and “force” refer to these 100-Hz processed data. Obtaining this envelope of the normalized muscle tension by regressing an optimally low-pass filtered version of the EMG signal on the fingertip force produces the best input-output mapping (in the least-squares sense) that suffices to study the structure of motor variability. This “black-box” approach circumvents the problems caused by introducing large numbers (~50) of unknown parameters necessary to perform the same analysis using a detailed physiological model.

Fingertip force measurement and analysis

The index fingertip was fitted with a thimble that had a small polished Teflon sphere (~3 mm diam) attached to it. The sphere made contact with a force-sensing plate which recorded the three-dimensional (3D) isometric force vector. The plate was covered with 300-grit sand paper, explicitly defining an isometric force production task with a friction cone of ~30° half-angle. Seven EMGs (see *Signal processing*) and three fingertip force components were recorded at

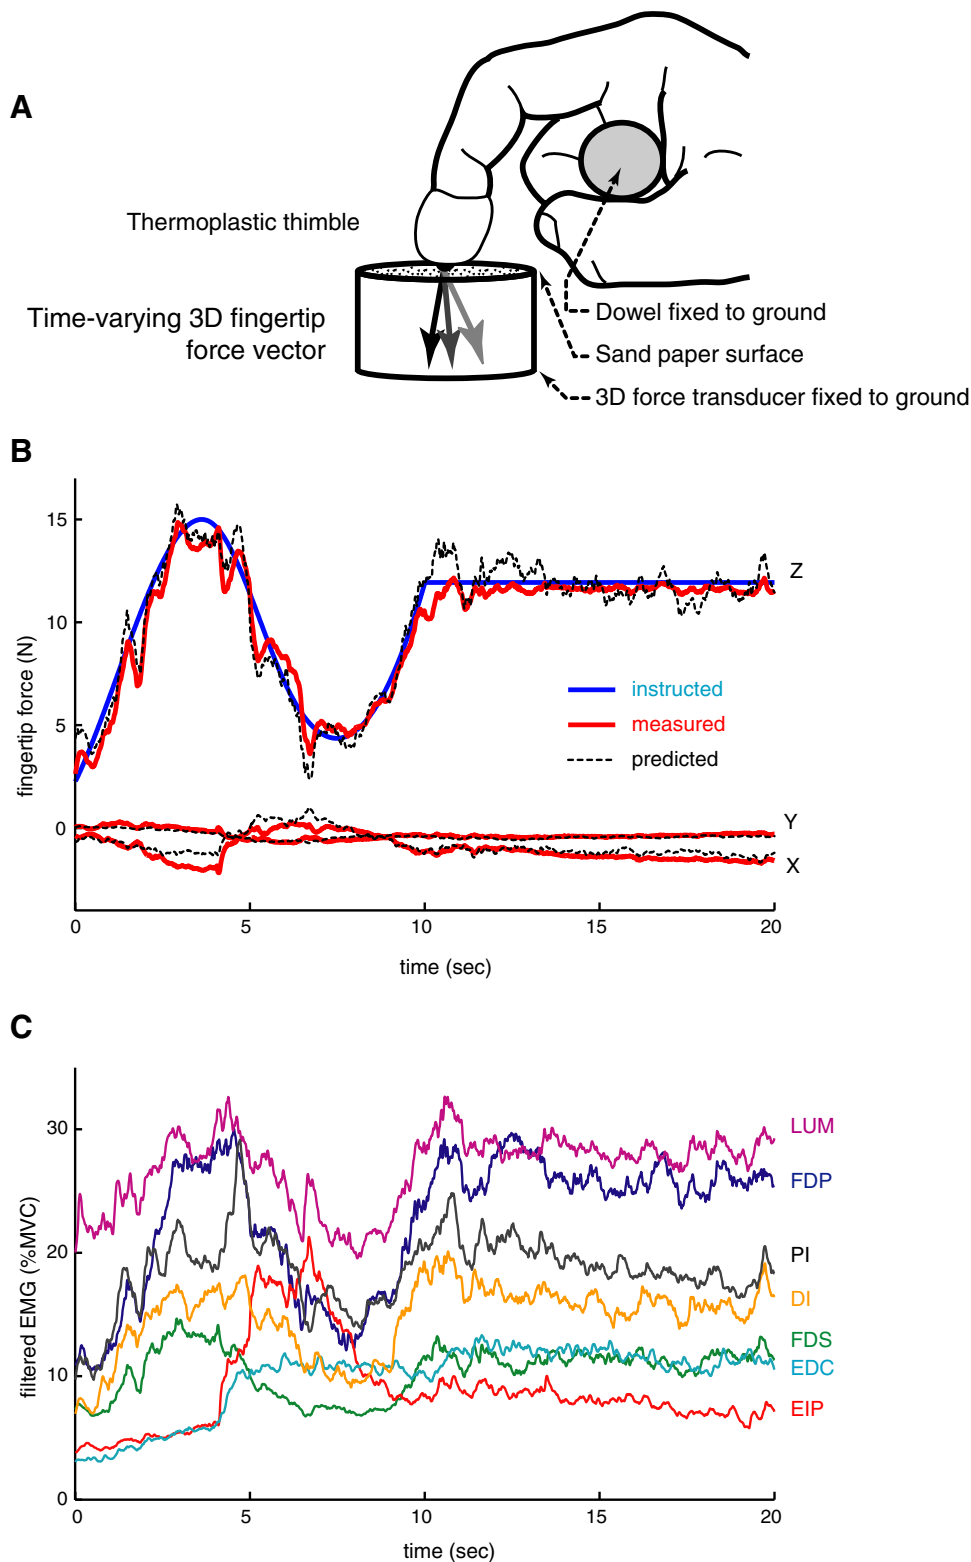


FIG. 1. *A*: experimental paradigm. As in Venkadesan and Valero-Cuevas (2008), the subject grasped a metallic dowel fixed to ground with their thumb, middle, ring, and little finger and placed the fingertip of the index finger on the center of the recording surface of the rigidly held force sensor, with the distal phalanx in a vertical orientation. Defining the location of the hand and fingertip in 3 dimension (3D), plus the orientation of the distal phalanx, in this way suffices to replicate the posture of the fingers across trials in this isometric task. *B*: measured fingertip forces (thick, red), predicted forces from the linear normalized muscle tension-to-force model (dotted, black), and instructed normal force (thin, blue) for a typical trial. Each trial started with a brief preloading phase which is not shown (and is not used in the analyses). *C*: processed normalized muscle tensions for the 7 muscles acting on the index finger, for the same trial shown in subplot *B*. FDP, flexor digitorum profundus (slip to the index finger); FDS, flexor digitorum superficialis (slip to the index finger); EIP, extensor indicis proprius; EDC, extensor digitorum communis (slip to the index finger); LUM, 1st lumbrical; DI, 1st dorsal interosseous; and PI, 1st palmar interosseous.

2,000 Hz. The task was to produce an instructed pattern of force, normal to the sensor surface, without deviating from a prescribed, comfortable finger posture (neutral abduction, proximo-distal joint angles of 30, 45, and 15° flexion, respectively). The instructed force pattern was displayed on a monitor as a function of time. It was randomly generated for every subject and trial and consisted of three phases: an initial hold phase of 2 N for 4 s ($t = 0-4$ s), a random, slowly varying phase for 10 s ($t = 4-14$ s), and a final hold phase for

10 s ($t = 14-24$ s). The slowly varying phase was generated as follows. We drew five random numbers from the uniform distribution over the interval [0,1] and formed a sequence of six numbers by prefixing the five random numbers with 0. These six numbers represent normalized force values (between 0 and 1) at every 2 s in the interval $t = 4-14$ s. Then we used a cubic spline interpolation to calculate a normalized force value at every millisecond for $t = 4-14$ s (i.e., 10,000 samples). Finally, we scaled this smooth normalized time-series to lie between 2 and 15 N.

The value of this slowly varying force at $t = 14$ s was set as the force-level for the final hold-phase. For an example of one such randomly generated target-force time series, see Fig. 1B. The instantaneous normal force produced by the subject was displayed as a small cross moving over the instructed pattern. This enabled subjects to perform the task accurately (a typical trial is shown in Fig. 1B). Note that while the normal force (Z) had to match the instructed pattern, the forces tangential to the sensor surface (X, Y) had to remain within a large friction cone defined by the thimble on sand paper to avoid slip. The friction cone was large enough that we never saw slips. Only the component normal to the surface was shown to the subject. The tangential components were only constrained by the friction cone. Thus the subject was required to control all three force vector components, the normal component to track the instructed pattern—while maintaining the tangential components low enough to stay within the friction cone. In recent study using a similar finger-pressing task, Cole (2006) has shown that young adults will control the direction of the fingertip force to be close to the surface normal even though there was no feedback provided on the tangential components. Contrast this with a hypothetical rigid coupling between the fingertip and the force sensor where only the normal force would have to be controlled (we did not test this case). We collected five to six usable trials per subject (for a total of 44 trials across all subjects). The usable trials were defined as trials without recording artifacts (identified by visual inspection of the raw EMG voltages). Then we performed a test for MVC in the same posture as the experiment by asking subjects to push against the experimenter's hand as hard as possible in different directions chosen to elicit maximal activity in each muscle.

Prediction of force given normalized muscle tension

The prediction of fingertip force given EMG was based on the linear model

$$\mathbf{f}_n(t) = \mathbf{A}_n * \mathbf{e}_n(t) + \mathbf{b}_n$$

where $\mathbf{f}_n(t)$ is the measured 3D fingertip force, $\mathbf{e}_n(t)$ is the 7D vector of measured normalized muscle tensions, \mathbf{A}_n is a 3×7 time-invariant matrix mapping normalized muscle tensions to fingertip force, \mathbf{b}_n is a time-invariant 3D force bias vector (to account for the finger's weight, preloading, etc.), t is the time index within the trial (varying from 0 to 20 s in discrete steps of 0.01 s), and the subscript n denotes the trial number. Recall that normalized muscle tension refers to the estimates of instantaneous percentage of maximal force generated by each muscle as inferred from EMGs. Linearity and time invariance of the biomechanical structure of the finger as well as force-length and -velocity properties of muscle are valid assumptions as long as the muscle moment arms or posture are not changing, which is the case here (Valero-Cuevas 2000; Valero-Cuevas et al. 1998). The unknown model parameters \mathbf{A}_n and \mathbf{b}_n should be constant for a given subject. However, to avoid over-fitting, we computed these parameters using leave-one-out cross-validation—that is, we used data from all trials performed by the same subject except for the trial n being analyzed. The variations in \mathbf{A}_n and \mathbf{b}_n were small, and repeating all analyses with the average values \mathbf{A} and \mathbf{b} produced results that were almost identical to those shown in the paper. The same results were also found using a Bayesian analysis (see Supplementary notes¹). As in our previous work, the linear model predicted the measured fingertip force rather well (Valero-Cuevas et al. 1998). For each trial, we computed the R^2 between the measured and predicted normal force. The median over trials was 0.91, meaning that the measured normalized muscle tension variance explained 91% of the variance in the measured normal fingertip force. This is not surprising in principle given that changes

in force are caused by muscle activation. However, this remarkable predictive ability of measured normalized muscle tension is not typical in the literature because EMG recordings often do not include all muscles, are not selective enough (especially surface recordings), or are too noisy to allow such accurate prediction.

Analysis of variability patterns

The goal of this analysis was to test whether the variance in the normalized muscle tension across all seven muscles had a structure indicative of task-relevant preferential attenuation. In this work, we use the qualitative term “variability” to mean the temporal changes in normalized muscle tension, and the quantitative term “variance” to mean the statistical metric of dispersion (i.e., SD squared). We performed this analysis of variability separately for the constant (Fig. 2, A–C) and time-varying phases (D–F). In each condition, we tested three alternative metrics of variance: “full,” which considers interactions within and across muscles, calculated by the covariance matrix; “diagonal,” which only considers within-muscle interactions, calculated by the autocovariance matrix (the multidimensional version of autocorrelation with diagonal elements identical to the full covariance and all off-diagonal elements set to 0); and “signal-dependent noise” (SDN), which estimates the variance attributable to mean muscle activation level. These covariance matrices were calculated as follows.

In the constant phase, we simply took the normalized muscle tension data for each trial (\mathbf{e}_n) and computed three trial-specific 7×7 matrices: First, the covariance matrix P_n (i.e., full variance); second, the matrix D_n (i.e., diag variance) the diagonal elements of which were the same as in P_n but the off-diagonal elements were set to 0. D_n is therefore the covariance of normalized muscle tensions that are individually as variable as the recorded signals but are uncorrelated with each other (i.e., diagonal variance or auto-covariance); finally, we formed a diagonal matrix S_n (i.e., SDN variance) with diagonal elements proportional to the mean muscle activations squared. This is the covariance matrix expected from signal-dependent noise—a well-documented characteristic of the motor system (Harris and Wolpert 1998; Sutton and Sykes 1967). This proportionality constant was determined separately for each subject, by fitting a linear model that predicts the SDs of the normalized muscle tensions given their means. Averaged over subjects, the SD of the normalized muscle tension was 8% of the mean.

In the time-varying phase, the calculation of these covariance matrices was complicated by the fact that the mean was time varying and could not be measured directly (the instructed patterns were not repeated so as to avoid rote motor pattern repetition by the subjects and explore a richer set of motor commands). However, it was possible to estimate the time-varying mean using the fact that, as we have shown in the same experimental paradigm (Valero-Cuevas 2000; Venkadesan and Valero-Cuevas 2008), subjects scale the activity of all muscles together to modulate the force. For each muscle, we computed what the “ideal” pattern should be according to this scaling strategy. To do so, we fitted another linear model predicting each muscle's normalized tension as a function of the instructed force. That model had two unknown scalar parameters computed separately for each subject and muscle. The predicted normalized muscle tension was then subtracted from the actual normalized muscle tension, and the residuals were used to compute the matrices P_n , D_n , S_n .

The second step in our analysis consisted of calculating how the output force was affected by the variance in the muscle coordination pattern estimated using the P_n , D_n , S_n matrices. More specifically, our hypothesis is expressed mathematically as testing whether the variance in normalized muscle tension (as quantified P_n , D_n , S_n matrices) is preferentially channeled into the null space of the biomechanical transformation (i.e., the matrix \mathbf{A}_n). See Valero-Cuevas (2009) and Strang (1980) for a review of these concepts. Briefly, if the variability of muscle activations is structured in such a way to reduce its effect on the relevant elements of the task (in this case, the magnitude of the

¹ The online version of this article contains supplemental data.

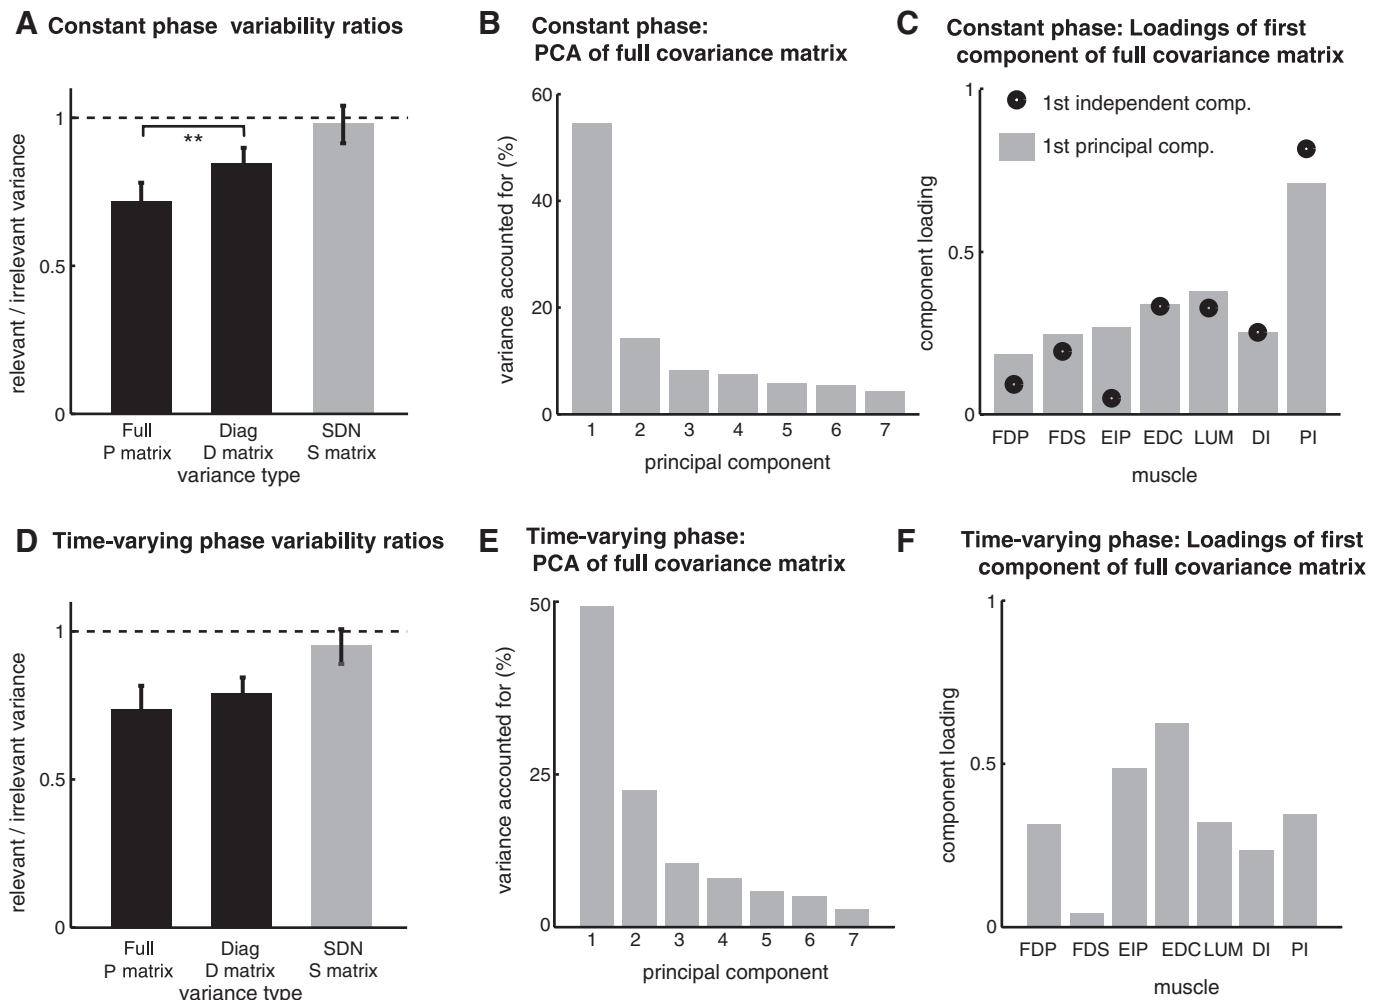


FIG. 2. *A*: ratio of task-relevant to -irrelevant variance indices, computed from the full covariance (full, *P*), diagonal covariance (diag, *D*), and signal-dependent noise covariance (SDN, *S*) in the constant phase (last 10 s of each trial). The bars show the grand means \pm SE computed from the means from the 8 individual subjects. Black, significantly different from 1 (*t*-test, $P < 0.01$); gray, not significantly different from 1. *B*: variance accounted for by each principal component. Principal-component analysis (PCA) is applied to the full covariance matrix averaged over all trials for the constant phase. *C*: loadings of the 1st principal component on all muscles (■) for the constant phase. Note that all values are positive, indicating positive correlation among muscles or “coactivation.” ●, the loadings of the 1st independent component, extracted via independent-component analysis (ICA) from the pooled data (all trials, constant phase). Trial-specific means were subtracted before pooling the data. *D–F*: same as in subplots *A–C* but for the time-varying phase of each trial (1st 10 s after the constant preloading phase).

three components of force), then the majority of the variability in normalized muscle tension is part of the family of motor commands that do not cause a change in the force normal to the surface (i.e., they belong to the nullspace of the biomechanical transformation A_n). The nullspace can be described intuitively as the task-irrelevant subspace: the set of all possible normalized muscle tension combinations that *do not* produce a fingertip force output. Thus normalized muscle tension variations in that task-irrelevant subspace do not affect the fingertip force. Once the P_n , D_n , S_n matrices for the constant and time-varying phases were computed, we compared the projected variance per dimension in the *task-relevant* subspace (row space of the A_n matrix) versus *task-irrelevant* subspace (i.e., nullspace of the A_n matrix) of the 7D normalized muscle tension space, separately for each trial. The task-relevant subspace is found directly as the subspace spanned by the three 7D row vectors of the matrix A_n . Recall that A_n maps 7D normalized muscle tension signals into the three components of the fingertip force vector (1 normal and 2 tangential to the surface). Note that this task-relevant 3D subspace² is by definition embedded in

the 7D normalized muscle tension space, and the three row vectors define the set of all possible normalized muscle tension combinations that produce a fingertip force output (Valero-Cuevas 2009; Valero-Cuevas et al. 1998). Thus this is the subspace in which normalized muscle tension variations cause variations in fingertip force. The *task-irrelevant* subspace is the 4D complement of the 3D task-relevant subspace, namely, the 4D null space³ of A_n . Last we compared how the variances (P_n , D_n , or S_n) project onto the task-relevant versus the task-irrelevant subspaces. For each subspace, we formed an orthonormal coordinate frame, computed the projection of the variance of the normalized muscle tension data projected on each axis of this coordinate frame, and averaged over the number of axes (3 for the task-relevant and 4 for the task-irrelevant subspaces) to obtain a scalar variability index for the corresponding subspace as done in Scholz and Schoner (1999). The projected variance is calculated by the formula: $\mathbf{u}^T * V * \mathbf{u}$, where V is a covariance matrix (P_n , D_n , or S_n) and \mathbf{u} is a unit vector from the orthonormal bases (i.e., 1 of the axes of the coordinate frame). An important mathematical fact is that the variability index computed this way does not depend on the choice of

² By 3D we do not mean Cartesian space. Rather that three 7D vectors (each a row vector of the A matrix) define a subspace in 7D space whose dimensionality is 3 (Strang, 1980).

³ By 4D we mean that four 7D vectors (each a basis vector of the nullspace of the A matrix) define a subspace in 7D space whose dimensionality is 4.

coordinate frame for a given subspace but only on the subspace itself; in other words, the variability index is coordinate-free. Finally, we divided the task-relevant variability index by the task-irrelevant index. The resulting dimensionless *variability ratio* should be statistically smaller than 1 to support the hypothesis that the nervous system preferentially controls task-relevant parameters.

RESULTS

Variability structure reveals preferential control of task-relevant parameters

The results in Fig. 2, *A* and *D*, confirm our hypothesis for both the constant and time-varying phases, respectively. The graphs show the means and SE of the ratio defined in METHODS for each type of covariance matrix. The mean ratios for the eight individual subjects were averaged to produce the grand means and SE shown in Fig. 2, *A* and *D*; ■ are cases where the ratio was significantly lower than 1 (*t*-test, $P < 0.01$), □ are cases where the difference was not statistically significant. In agreement with our hypothesis, both the full covariance P_n and the diagonal covariance D are structured such that variance per dimension is smaller in task-relevant compared with task-irrelevant subspaces. This is not the case for the signal-dependent noise covariance matrix S , therefore the task-relevant reduction in variance is not a trivial consequence of signal-dependent noise. In the constant phase (Fig. 2*A*), the ratio (0.73 ± 0.08) was significantly smaller for the full covariance compared with the diagonal covariance (paired *t*-test, $P < 0.01$), therefore correlations among muscles contributed to the effect. The latter difference was not statistically significant in the time-varying phase (Fig. 2*D*). But recall that our calculation was based on additional assumptions in the time-varying phase and is therefore less reliable. Nevertheless this result shows the same trend for both constant and time-varying phases and confirms that the variability ratios for both phases are statistically lower than 1. Figures 2, *C* and *F*, and 3, *A* and *B*, show the loadings for all principal components for both the constant and time-varying phases. The loadings for each principal component indicate only the trends in variance among muscles (but not on their mean values in the coordination pattern) in order of decreasing importance (from 1st through 7th). Note that both the variance explained by each principal component agree in spite of the need to make additional assumptions to perform this analysis on the time-varying phase (see METHODS). The loadings of the first principal component also agree for both phases, with some differences in magnitude (but not in sign) for mostly the first palmar interosseous, the extensor digitorum communis and flexor digitorum superficialis. Last, Fig. 4 shows the histogram composed of the ratio of task-relevant:task-irrelevant variance for the full covariance matrix from all trials. This figure shows that the variability ratio was <1 for 80% of the trials. Thus we can conclude that our hypothesis holds for the neural control of both constant and time-varying force production.

Fluctuations in muscle “coactivation” are the largest source of variability

We also performed principal-component analysis (PCA) on the full covariance matrix (P_n) averaged over all trials in the constant phase of force production. We found one large com-

ponent (Fig. 2*B*) accounting for $> 50\%$ of the variance, a second for 14%, followed by five smaller components each accounting for $\sim 7\%$ on average (Fig. 3*A*). Figure 2*C* shows the loading of the first principal component on the seven muscles. The fact that all loadings are all the same sign suggests that a substantial source of variability is positively correlated fluctuations—a form of muscle “coactivation.” Note however that hand muscles do not have a simple agonist-antagonist arrangement (Valero-Cuevas 2005, 2009; Valero-Cuevas et al. 1998), and furthermore the loadings in Fig. 2*C* are far from being equal, so the term coactivation should not be taken literally to mean simply joint stiffening via agonist-antagonist co-contraction. Rather the loadings of the first principal component for both the constant and time-varying phases indicate that all muscles contribute to the overall variance, and this variance has a structure favoring its prevalence in the task-irrelevant subspace (Fig. 2, *A* and *D*). There is, of course, some amount of joint stiffening naturally due to this muscle coactivation, but as we have shown before (Valero-Cuevas 2005, 2009; Valero-Cuevas et al. 1998), the biomechanical arrangement of finger muscles is such that, for example, coactivation of flexors and extensors is biomechanically necessary to direct the fingertip force well and *not* indicative of a “stiffening strategy.”

Our analysis using PCAs was based on covariance matrices that can at most capture second-order correlations (i.e., in the off-diagonal elements). It is possible, however, that the variance of muscle activations has more complex statistical structure. Thus we also applied independent-component analysis (FastICA with *tanh* nonlinearity) (Hyvarinen et al. 2001), which looks for linear projections of the data that are statistically independent and not merely uncorrelated. PCA and ICA are based on different mathematical models and tend to find different solutions—which is why ICA has become so popular recently despite the increased computational requirements (Hyvarinen et al. 2001). In our data, however, the first PCA and first ICA components were very similar (Fig. 2*C*), indicating that the dominant component we found represents not only the direction of maximal variance but also an independent source of variation. There are differences between the two methodologies, mostly for extensor indices proprius, but the loadings of the first component for the two methods were of the same sign. The agreement of PCA and ICA on higher components was not as good, but those components explain comparatively little variance.

For completeness, we also report the magnitude and direction of the fingertip force vectors. For the constant phase, the mean normal force was 8.88 ± 0.42 N, and the mean tangential force was 0.83 ± 0.12 N. The angular deviation of the fingertip force vector from the vertical was $2.77 \pm 0.59^\circ$. The mean maximal angular deviation was 5.72° . For the time-varying phase, the mean normal force was 9.20 ± 3.15 N (the large SD reflects the time-varying target), and the mean tangential force was 0.75 ± 0.31 N. The angular deviation of the fingertip force vector from the vertical was $2.34 \pm 0.67^\circ$. The mean maximal angular deviation was 4.92° . Note that in all cases, subjects stayed well within the large friction cone of $\sim 30^\circ$, the half angle afforded by the 300-grit sand paper covering the force sensing surface.

DISCUSSION

Our hypothesis of smaller variance in the task-relevant subspace was confirmed, and the possibility that the phenom-



A Constant phase: Loadings of 2nd-7th component of full covariance matrix**B Time-varying phase: Loadings of 2nd-7th component of full covariance matrix**

FIG. 3. Loadings of the 2nd through 7th principal component for constant (A) and time-varying (B) phases. This figure shows the strong similarity in the percent variance explained and the loadings obtained for both phases despite the assumptions made for the time-varying analysis.

enon is a trivial consequence of signal-dependent noise was ruled out. Unlike previous behavioral measurements where the structure of the variability is affected by biomechanical coupling, such coupling was not an issue here. The possibility that trial-to-trial variability in the task-irrelevant subspace is inflated by fatigue-related or exploratory strategies was not an issue either because we computed variance within trials. Therefore this work is the first direct physiological evidence for preferential control of task-relevant parameters that strongly suggest the use of a neural control strategy compatible with the principle of minimal intervention.

The main result of the present study is an extension of a similar preliminary finding originally detected in data from our previous experiment (Valero-Cuevas 2000). That finding was reported in a conference format (Valero-Cuevas and Todorov,

Neural Control of Movement 13 2003) and served as motivation to conduct this expanded study. The previous experiment only included a constant phase, with shorter duration (~ 2 s); however, the number of trials was larger and the directions of instructed force were varied. In that preliminary report, the variance ratio for the full covariance was 0.62 ± 0.08 (mean \pm SE), which was significantly different from 1 (t -test, $P < 0.01$). The corresponding ratio in the present study is given by the leftmost bar in Fig. 2A; which is 0.73 ± 0.08 .

Our experiment has some limitations, which point to directions for future work but do not challenge the validity of our results. First, the limited time window of opportunity, characteristic of experiments with fine-wire electrodes, prevented us from repeating the dynamic templates. This complicated the variance analysis in the time-varying phase. Second, subjects




FIG. 4. Histogram and cumulative histogram of the ratio of task-relevant to -irrelevant variance for both constant and time-varying phases for the Full covariance matrix. Note that 80% of the trials have a ratio <1 , in support of our hypothesis.

had to control the fingertip forces in all three directions accurately, and so the task-relevant 3D subspace was quite generic. This issue could be addressed in future work by rigidly fixing the fingertip to the target surface—making the fine control of tangential forces unnecessary, and thus creating a less generic 1D task-relevant subspace. However, investigating motor variability in unloaded finger movements and during object manipulation task would be stronger tests of whether this approach to understanding variability is useful in tasks of daily life. Third, in isometric tasks, it is difficult to separate state variables from control variables—a separation that is essential if feedback control analyses are to be applied (recall that a feedback control law is a mapping from states to controls). Movement tasks would be more suitable for such analyses. Near-isometric tasks or posture maintenance tasks in which small postural fluctuations are recorded are also suitable and may allow estimates of stiffness modulation. Last, the assumptions in the analysis of the time-varying phase could be relaxed by embedding an identical time-varying pattern within a sequence of random patterns. However, such an approach would be open to the criticism that our results may not apply in general (i.e., to a rich variety of random force patterns as done here), and would unavoidably lengthen these invasive experiments. These are all extensions that we hope to explore in future work. More generally speaking, this work also underscores the need to develop methods to study physiological causes of motor output variability for general motor tasks without the need to record from all relevant muscles. The

limitations of the use of EMG recordings are well known, and we have discussed them in this context before (Valero-Cuevas et al. 1998; Venkadesan and Valero-Cuevas 2008). For example, if the estimate of normalized muscle tension from fingertip force has a bias there could be an offset in the normalized muscle tension estimate that would artificially increase the apparent variance. However, EMG recordings are the best tool we have at the moment to estimate descending drive to muscles and our results are strong enough to suggest that the limitations of EMG did not overwhelm the effect we detected.

On a related note, one might think that correlated drive to motor units of hand muscles (for a review of this extensive literature, see Schieber and Santello 2004) is a confound in this study. However, this is not so. Correlated drive is a mechanistic explanation of the same phenomenon we are trying to explain in computational terms. Every computational model must have an underlying neural mechanism. Indeed if the control laws used by the sensorimotor system were task-specific in the way we envision (Liu and Todorov 2007; Todorov 2004), their neural implementation would involve the kind of correlated drive that has been reported—with the caveat that the correlations would have to be task-specific. So these two explanations are complementary: one tells us what the control law is and why, the other tells us how that control law is implemented.

We avoid discussing the individual off-diagonal entries (i.e., pair-wise interactions) or rows (i.e., muscle groups) of the full covariance matrix. This is because it is imperative to note that considering the low-dimensional trends in data represented by those entries (such as pair- or group-wise correlations) can lead to overinterpretation of the causality behind muscle interactions. Even if such correlations are detected on a restricted subset of the entire data, the causal reasons for them cannot be interpreted in a mechanically sound manner. This is because the complicated musculo-skeletal structure of the hand obviates certain traditional notions of agonist-antagonist pairs, flexor versus extensor muscle groups, co-contraction, etc. (Valero-Cuevas 2005, 2009; Valero-Cuevas et al. 1998, 2000). For example, vector contribution of action of the extrinsic flexors (FDP and FDS) to fingertip force in a flexed posture is actually not in the direction normal to the surface but rather tilted proximally, whereas the intrinsic muscles have lines of action better aligned with the surface normal (Cole 2006; Valero-Cuevas 2005, 2009; Valero-Cuevas et al. 2000). The fact that the flexors do not have the highest loadings in Fig. 2, C and F, is therefore biomechanically reasonable. These complex biomechanical interactions, however, are difficult to interpret via the loadings in the principal components or individual entries of the covariance matrices. Figure 5 further reinforces this idea using a representative trial. The data in Fig. 5 correspond to the same trial as in Fig. 1, but instead of showing the 7D normalized muscle tension space as spanned by each muscle, it shows the projection of the 7 normalized muscle tensions onto the task-relevant (top 3 traces) and task-irrelevant (bottom 4 traces) subspaces. These subspaces can only be calculated using the full covariance matrix. Therefore these are not the traces of variability in individual muscles but of variability within two disjoint subspaces that are defined by specially oriented coordinate systems in the 7D normalized muscle tension space. The variability in these subspaces either will (top 3 traces) or will not (bottom 4 traces) affect the

Projection of EMG variance (arbitrary units) onto the:




FIG. 5. Projection of the normalized muscle tension variance onto the task-relevant (top 3 red traces) vs. task-irrelevant (bottom 4 traces) subspaces. The data are for the same trial as shown in Fig. 1. Note that this figure highlights the fact that the multidimensional interactions across muscles is a complex phenomenon that is not readily detectable by studying individual muscle signals or, more importantly, when one records only from a subset of muscles as was done in prior studies. An important motivation for this study is that, to our knowledge, this is the first time this analysis is performed on electromyographic signals collected simultaneously from all muscles of a musculoskeletal system, in this case a finger.

fingertip force output. Even though specific off-diagonal elements in the full covariance represent the coactivation of specific muscle pairs, the true nature of the variability that is embedded in a 7D space cannot be captured unless analyzing all elements of the full covariance matrix simultaneously. The same applies when considering only restricted muscle groups, i.e., specific rows of the full covariance matrix. Moreover, even in Fig. 5, it is difficult to assess differences in variability across the two subspaces by comparing across traces. The full calculation of variance per dimension (see METHODS) is necessary to detect the effect by calculating the ratio of task-relevant:task-irrelevant variance, which in this particular trial is strong at 0.48 (i.e., much lower than 1).

Lack of evidence for dimensionality-reducing muscle “synergies” in within-trial variability

We find that our results speak to motor control issues beyond our hypothesis—particularly to the issue of motor synergies. The idea that related motor behaviors may be constructed by recombining a small set of synergies has been around for a long time (Bernstein 1967). Recently it has been instantiated on the muscle level by using linear decomposition methods like PCA on the normalized muscle tension signals to identify dimensionality-reducing synergies for the control of redundant musculature (e.g., d’Avella et al. 2006; Krishnamoorthy et al. 2003; Li et al. 1998; Soechting and Lacquaniti 1989; Torres-Oviedo and Ting 2007; Tresch et al. 2006). For example, it was shown that four or five different muscle synergies can explain most of the EMG variance in 3D reaching movements to/from a variety of targets in 3D (d’Avella et al. 2006).

The existing literature on synergies and dimensionality reduction has focused on explaining between-trial variability. In contrast, the focus of our study is within-trial variability. If we assume that all muscle activity (including within-trial variability) is generated by recombining a small set of synergies, then the variability structure we observed during the production and regulation of this task is inconsistent with the notion of synergies, for the following reasons. If we were to count the number of hypothetical muscle synergies by drawing a cut-off line somewhere in Fig. 2, *B* and *E*, the only logical place is after the first principal component. This would leave >40% of the variance unaccounted for, and would imply that either there is only one synergy, making it impossible for the CNS to utilize the compositionality property of multiple synergies. Furthermore, given that all coefficients of the first principal component have the same sign (which to our knowledge has never been observed in previous studies), this first principal component may simply reflect overall modulation of fingertip force as well as stiffness. Or it could imply there was no dimensionality reduction because each of the seven principal components explains a nontrivial amount of variance.

This finding can be reconciled with the existing literature by noting that there may be differences between open- and closed-loop control as well as between planning and execution. Previous studies, focusing on between-trial variability that is mostly driven by changes in task parameters, have emphasized planning and open-loop control. Even though these movements were executed under closed-loop control, averaging over multiple trials in the same condition is likely to eliminate within-trial variability. In contrast, our analysis emphasizes variability within a trial, where the task conditions are kept constant and

the only fluctuations are internally generated—presumably reflecting noise as well as closed-loop corrections. So it may be that planning and open-loop control rely on synergies while execution and closed-loop control do not. Of course it is also possible that our task is too simple, and the closed-loop controller is not as rich as it may be in other tasks where synergies might be revealed. However, to the extent that synergies are low-level mechanisms that are mostly task-independent, we should see evidence for them in every task.

Both the open- and closed-loop point of view of synergies are in principle valid, and a more thorough comparison between the two is likely to be illuminating. Because prior studies have almost exclusively emphasized the former approach, here we present some arguments in favor of the latter with the hope of stimulating a more balanced treatment in future work. In analyses of between-trial variability, it is difficult to dissociate the unavoidable consequences of task variation and musculoskeletal structure from the intrinsic properties of the neural controller. For example, consider the EMG patterns during a center-out reaching task. Suppose for a moment that there is no inherent variability in the sensorimotor system and all variability is imposed by the task—meaning that reaches in different directions require different EMG patterns. Suppose also that the control strategy is such that small changes in target location correspond to small changes in EMG and the mapping between the two is smooth. Then the observed EMG patterns will lie on a 1D manifold embedded in the high-dimensional EMG space simply because the reach targets lie on a 1D manifold (i.e., a circle). An ideal and necessarily nonlinear dimensionality reduction algorithm will be able to explain all EMG data in this hypothetical experiment with a single muscle synergy. Linear dimensionality reduction algorithms such as PCA are not ideal, so we should expect them to find a subspace with more than one dimension—but still it should be a very low-dimensional subspace. The same reasoning applies to tasks like locomotion, where the behavior is very complex but nevertheless remains close to a 1D limit cycle embedded in some high-dimensional space. The findings from such studies are useful in the sense that they tell us what the low-dimensional space is, but the fact that the space is low-dimensional is hardly surprising given the low-dimensionality inherent in the task. This confound is largely avoided in analyses of within-trial variability.

ACKNOWLEDGMENTS

Present address of M. Venkadesan: School of Engineering and Applied Sciences, Harvard University, Cambridge, MA.

GRANTS

This material is based on work supported by National Science Foundation Grant 0237258 to F. J. Valero-Cuevas and National Institutes of Health Grants AR-050520 and AR-052345 to F. J. Valero-Cuevas, and NS-045915 and NS-058633 to E. Todorov. M. Venkadesan was supported by NSF Grant 0425878 during manuscript preparation.

REFERENCES

- Bernstein N.** The problem of interrelation of coordination and localization. In: *The Coordination and Regulation of Movements*. New York: Pergamon, 1967, p. 15–59.
- Burgar CG, Valero-Cuevas FJ, Hentz VR.** Fine-wire electromyographic recording during force generation. Application to index finger kinesiologic studies. *Am J Phys Med Rehabil* 76: 494–501, 1997.
- Clewley RH, Guckenheimer JM, Valero-Cuevas FJ.** Estimating effective degrees of freedom in motor systems. *IEEE Trans Biomed Eng* 55: 430–442, 2008.
- Cole KJ.** Age-related directional bias of fingertip force. *Exp Brain Res* 175: 285–291, 2006.
- d'Avella A, Portone A, Fernandez L, Lacquaniti F.** Control of fast-reaching movements by muscle synergy combinations. *J Neurosci* 26: 7791–7810, 2006.
- Gordon J, Ghilardi MF, Ghez C.** Accuracy of planar reaching movements. I. Independence of direction and extent variability. *Exp Brain Res* 99: 97–111, 1994.
- Harris CM, Wolpert DM.** Signal-dependent noise determines motor planning. *Nature* 394: 780–784, 1998.
- Hyvarinen A, Karhunen J, Oja E.** *Independent Component Analysis*. New York: Wiley, 2001.
- Krishnamoorthy V, Latash ML, Scholz JP, Zatsiorsky VM.** Muscle synergies during shifts of the center of pressure by standing persons. *Exp Brain Res* 152: 281–292, 2003.
- Latash ML, Scholz JF, Danion F, Schoner G.** Finger coordination during discrete and oscillatory force production tasks. *Exp Brain Res* 146: 419–432, 2002.
- Latash ML, Scholz JF, Danion F, Schoner G.** Structure of motor variability in marginally redundant multifinger force production tasks. *Exp Brain Res* 141: 153–165, 2001.
- Li ZM, Latash ML, Zatsiorsky VM.** Force sharing among fingers as a model of the redundancy problem. *Exp Brain Res* 119: 276–286, 1998.
- Liu D, Todorov E.** Evidence for the flexible sensorimotor strategies predicted by optimal feedback control. *J Neurosci* 27: 9354–9368, 2007.
- Schieber MH, Santello M.** Hand function: peripheral and central constraints on performance. *J Appl Physiol* 96: 2293–2300, 2004.
- Scholz JP, Schoner G.** The uncontrolled manifold concept: identifying control variables for a functional task. *Exp Brain Res* 126: 289–306, 1999.
- Soechting JF, Lacquaniti F.** An assessment of the existence of muscle synergies during load perturbations and intentional movements of the human arm. *Exp Brain Res* 74: 535–548, 1989.
- Strang G.** *Linear Algebra and its Applications*. Orlando, FL: Academic, 1980.
- Sutton GG, Sykes K.** The variation of hand tremor with force in healthy subjects. *J Physiol* 191: 699–711, 1967.
- Todorov E.** Optimality principles in sensorimotor control. *Nat Neurosci* 7: 907–915, 2004.
- Todorov E, Jordan MI.** Optimal feedback control as a theory of motor coordination. *Nat Neurosci* 5: 1226–1235, 2002.
- Torres-Oviedo G, Ting LH.** Muscle synergies characterizing human postural responses. *J Neurophysiol* 98: 2144–2156, 2007.
- Tresch MC, Cheung VC, d'Avella A.** Matrix factorization algorithms for the identification of muscle synergies: evaluation on simulated and experimental data sets. *J Neurophysiol* 95: 2199–2212, 2006.
- Tresch MC, Saltiel P, Bizzi E.** The construction of movement by the spinal cord. *Nat Neurosci* 2: 162–167, 1999.
- Valero-Cuevas FJ.** An integrative approach to the biomechanical function and neuromuscular control of the fingers. *J Biomech* 38: 673–684, 2005.
- Valero-Cuevas FJ.** A mathematical approach to the mechanical capabilities of limbs and fingers. *Adv Exp Med Biol* 629: 619–633, 2009.
- Valero-Cuevas FJ.** Predictive modulation of muscle coordination pattern magnitude scales fingertip force magnitude over the voluntary range. *J Neurophysiol* 83: 1469–1479, 2000.
- Valero-Cuevas FJ, Towles JD, Hentz VR.** Quantification of fingertip force reduction in the forefinger following simulated paralysis of extensor and intrinsic muscles. *J Biomechan* 33: 1601–1609, 2000.
- Valero-Cuevas FJ, Yi JW, Brown D, McNamara RV, 3rd, Paul C, Lipson H.** The tendon network of the fingers performs anatomical computation at a macroscopic scale. *IEEE Trans Biomed Eng* 54: 1161–1166, 2007.
- Valero-Cuevas FJ, Zajac FE, Burgar CG.** Large index-fingertip forces are produced by subject-independent patterns of muscle excitation. *J Biomechan* 31: 693–703, 1998.
- Venkadesan M, Valero-Cuevas FJ.** Neural control of motion-to-force transitions with the fingertip. *J Neurosci* 28: 1366–1373, 2008.
- Zajac FE.** Muscle and tendon: properties, models, scaling, and application to biomechanics and motor control. *Crit Rev Biomed Eng* 17: 359–411, 1989.

Supplementary Note on Bayesian analysis

Structured variability of muscle activations supports the minimal intervention principle of motor control

Francisco J. Valero-Cuevas^{1,2,3,*} Madhusudhan Venkadesan^{3,4,†}
Emanuel Todorov⁵

¹*Department of Biomedical Engineering, University of Southern California, Los Angeles, CA*

²*Division of Biokinesiology & Physical Therapy, University of Southern California, Los Angeles, CA*

³*Sibley School of Mechanical & Aerospace Engineering, Cornell University, Ithaca, NY*

⁴*Department of Mathematics, Cornell University, Ithaca, NY*

⁵*Department of Cognitive Science, University of California San Diego, La Jolla, CA*

1 Introduction

In the main text of this manuscript, we estimated the linear mapping from recorded electromyograms (EMGs) to fingertip force outputs using a least squares algorithm (i.e., linear regression). Throughout this supplement and the main text we denote this mapping by A , a 3×7 matrix that transforms muscle EMGs into fingertip forces. Subsequent analysis of variance in muscle EMGs was carried out using this matrix A . Specifically, the row-space and null-space of A were used to parse the variability in EMGs into task-relevant and task-irrelevant portions. However, estimating A using a least squares process could be biased. In particular if the EMGs for some muscles are more noisy than others, the corresponding entries in A could be smaller than the actual moment arms, which in turn could affect our conclusions regarding variability structure. Yet another reason for bias in using a least squares estimate of A arises from the fact that a least-squares procedure specifically tries to find A that minimizes errors in estimated fingertip

forces, i.e., noise in EMGs may be selectively dumped into the null-space of A . However, it is possible to work around these forms of bias in estimation by not only estimating the best-fit A , but finding a probability distribution for A . Here, we use Bayesian analysis to estimate the probability distribution for A using a numerical sampling scheme calling the Metropolis algorithm. This is in contrast to the point-estimate we found in the main text. We find that our conclusions about the structure of variability in muscle activations remain the same as when using a point-estimate of A . Therefore, we conclude that the least-squares estimation procedure presented in the main text did not bias the results of this study. We relegated the Bayesian analysis to this supplementary note for ease of reading the main text given greater familiarity with linear regression than with Bayesian techniques for several readers.

2 Methods

We pool together all trials for a given subject, drop the trial index, and also subtract the mean so that the offset b introduced earlier (in the main text) is

*Corresponding author

†Currently at School of Engineering & Applied Sciences, Harvard University, Cambridge, MA, USA

no longer needed. The procedure described below is repeated separately for every subject.

2.1 Bayesian Estimation

Instead of taking the data at face value, we consider the force ($\mathbf{f}(t)$) and measured EMG ($\mathbf{e}(t)$) as noisy versions of some underlying signals which correspond to the true fingertip force ($\bar{\mathbf{f}}(t)$) and true muscle activation ($\bar{\mathbf{e}}(t)$):

$$\mathbf{f}(t) = \bar{\mathbf{f}}(t) + \boldsymbol{\gamma}(t) \quad (1a)$$

$$\mathbf{e}(t) = \bar{\mathbf{e}}(t) + \boldsymbol{\omega}(t) \quad (1b)$$

where the noise terms $\boldsymbol{\gamma}(t)$ and $\boldsymbol{\omega}(t)$ are independent, zero-mean, multivariate, Gaussian white noise variables with diagonal covariance matrices Q and R. These two matrices along with the matrix A constitute the model parameters which we will estimate from the data. Note that A now relates the underlying signals ($\bar{\mathbf{f}}$, $\bar{\mathbf{e}}$) and not the measured signals (\mathbf{f} , \mathbf{e}):

$$\bar{\mathbf{f}}(t) = A\bar{\mathbf{e}}(t) \quad (2)$$

We do not know the underlying signals, but fortunately it is not necessary to know them. Indeed, multiplying the two noise models (Equation 1) by A and subtracting yields:

$$A\mathbf{e}(t) - \mathbf{f}(t) = A\boldsymbol{\omega}(t) - \boldsymbol{\gamma}(t) \quad (3)$$

The expression on the left (i.e., the residual) can be computed from the data for an assumed (guessed) A. The expression on the right however, is a Gaussian white random variable with zero mean and covariance matrix $ARA^T + Q$. Thus using the formula for the probability density function of a multivariate normal distribution, the likelihood of the measured residual at time t for given (or guessed) model parameters A, R, Q is

$$p_t(\boldsymbol{\theta}) = \frac{e^{(-\frac{1}{2}(A\mathbf{e}(t) - \mathbf{f}(t))^T (ARA^T + Q)^{-1} (A\mathbf{e}(t) - \mathbf{f}(t)))}}{(2\pi)^{\frac{3}{2}} \sqrt{|ARA^T + Q|}} \quad (4)$$

where $\boldsymbol{\theta}$, the parameters to be estimated using the data, are simply the elements of A, Q and R strung out as a long vector. Because, A has 21 elements

(3×7), Q has 3 elements (a 3×3 diagonal matrix), and R has 7 elements (a 7×7 diagonal matrix), the total number of parameters to be estimated (i.e., size of $\boldsymbol{\theta}$) is 31. Given that our data are in the form of discrete measurements, the total likelihood of the measured data time-series is simply the product over all t of the likelihood given by Equation 4:

$$p(\boldsymbol{\theta}) = \prod_t p_t(\boldsymbol{\theta}) \quad (5)$$

2.1.1 Numerical sampling of the posterior distribution

In the absence of a prior probability distribution¹ for $\boldsymbol{\theta}$ (i.e., a uniform prior), the posterior distribution² for $\boldsymbol{\theta}$ (the product of $p(\boldsymbol{\theta})$ and the uniform prior) is then proportional to the likelihood ($p(\boldsymbol{\theta})$). The key idea in Bayesian analyses is to use the posterior distribution instead of a single parameter estimate. When this distribution is too complex to handle analytically, as is the case here, one can use Markov Chain Monte Carlo (MCMC) sampling. The specific algorithm we use is the Metropolis algorithm, which uses the following iterative scheme to generate sample the posterior distribution. Let $\boldsymbol{\theta}^k$ denote the k^{th} parameter sample. Generate a candidate new sample $\boldsymbol{\theta}' = \boldsymbol{\theta}^k + \boldsymbol{\epsilon}$, where $\boldsymbol{\epsilon}$ is drawn from a zero-mean multivariate Gaussian distribution. Then, we compute the probability ratio:

$$a = \frac{p(\boldsymbol{\theta}')}{p(\boldsymbol{\theta})} \quad (6)$$

where p is given by Equation 5 and is proportional to the target distribution (posterior) from which we want to sample. If $a \geq 1$, make the deterministic update $\boldsymbol{\theta}^{k+1} = \boldsymbol{\theta}'$. Otherwise make the stochastic

¹The prior probability distribution or simply, the *prior*, is the name given to the uncertainty in the variable we want to estimate before any data are taken into account.

²The posterior probability distribution or simply, the *posterior*, is the name given to the conditional probability distribution of the variable we want to estimate after taking into account the data and a known prior. According to Bayes' theorem, this is simply the product of the prior and the likelihood of observing the measured data.

update

$$\theta^{k+1} = \begin{cases} \theta' & \text{with probability } a \\ \theta^k & \text{with probability } 1 - a \end{cases} \quad (7)$$

After an initial “burn-in” period, the sequence of samples generated in this way is guaranteed to match the target distribution. To ensure that the Markov chain has enough time to converge we allowed a large number of updates (10 million), and furthermore repeated the entire process 3 times with different initial values (θ^0). Visual inspection confirmed that for 7 out of the 8 subjects, the chain always converged in less than 4 million updates. Thus we discarded the first half of each chain. To speed up processing we kept only 1 out of each 100 consecutive samples – which is motivated by the fact that MCMC generates samples that are correlated over time. Combining the results from the 3 chains, we thus analyzed 150,000 samples from the posterior distribution over the parameters for each subject. For one of the 8 subjects we did not observe convergence. Furthermore for this subject the matrix A estimated with linear regression had substantially larger elements compared to the other subjects. Thus we decided to exclude this subject from the present analysis, even though this subject’s (unreliable) results were actually consistent with our hypothesis. This left us with a total of 41 trials from 7 subjects.

2.2 Variability ratio

We briefly recall how we quantified the ratio of task-irrelevant to task-relevant variability in EMGs. A detailed description is found in the main text. From the trial EMG data, we calculated three covariance matrices: Λ – the full covariance matrix, D – the diagonal covariance matrix, with diagonal elements identical to Λ and all off-diagonal elements set to zero, and finally S – the diagonal signal dependent noise covariance matrix such that the diagonal terms were proportional to the mean muscle activations squared. For a given A, the three rows of A span the task-relevant subspace of A and the four basis vectors of the nullspace of A span the task-irrelevant subspace of A. For every basis vector $\hat{\mathbf{u}}$

and covariance matrix V (where V is one of Λ , D , or S), the projected variance is given by the scalar quantity, $\hat{\mathbf{u}}^T V \hat{\mathbf{u}}$. Then, the variability index for each of the task-relevant and -irrelevant subspaces is just the average of the variance projected onto each of their respective basis vectors. The ratio of the variability index of the task-relevant to that of the task-irrelevant subspace is the quantity of interest (we call it the *variability ratio*). In the main text we found that for a least squares estimation of A, this ratio was smallest for Λ and specifically, it was smaller than 1 (both results were statistically significant). Here we repeat the same analysis using the 150,000 samples of A generated by the Metropolis algorithm. In addition to the average value of this ratio, the Bayesian method yields the posterior distribution of the variability ratio for a single trial, thus enabling single-trial hypothesis testing.

3 Results and discussion

We found that samples of the posterior distributions of A, R and Q, and therefore, the posterior distribution of the variability ratio, all converged. In other words, all three chains of the MCMC led to very similar distributions for the estimated mapping from EMGs to forces (A) (see Figure 1a) and the variability index (see Figure 1c). The estimated noise variances (R for EMG and Q for force) however, were typically non-Gaussian and differed to a greater degree between chains (see Figure 1b). This is not surprising because variances are typically harder to estimate. Finally, we found the linear regression estimates to be surprisingly close to the Bayesian estimates.

Using single-trial hypothesis testing we found that for 30 out of 41 trials, the variability ratio was statistically significantly smaller than 1 (filled, blue data points in Figure 2a). Recall that our hypothesis is that the variability index is below 1. For each trial, the Bayesian method gives us a sample from the posterior distribution of the variability ratio. Thus, hypothesis testing is performed by simply counting how many samples agree with the hypothesis and checking that the percentage is above a desired significance level (say 0.95). Of the 41 trials analyzed, in 30 cases

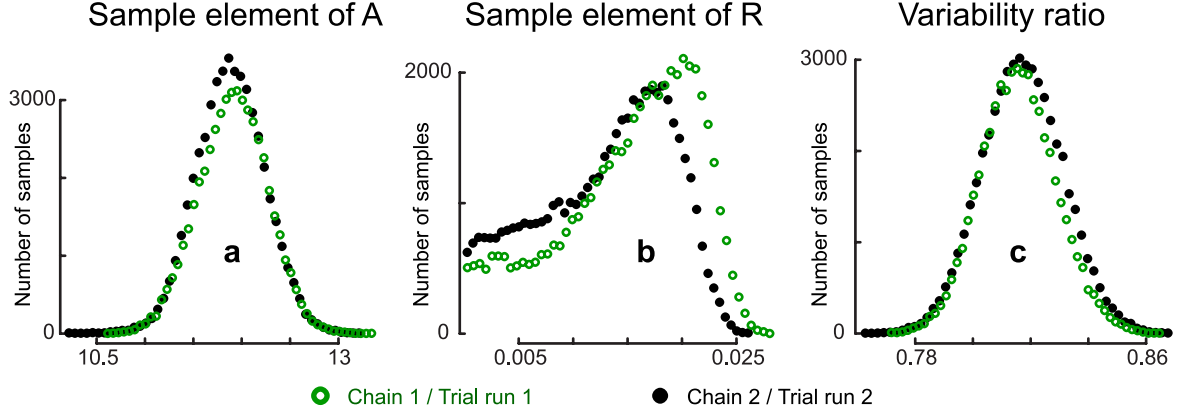


Figure 1: Examples of posterior distributions of A, R (EMG variance), and the calculated variability ratio. (a) Histograms generated using 50,000 samples (number of samples per chain) of one element of A for one subject and two different chains (different initial values) of the Metropolis algorithm. Note that the two distributions are almost identical and near-Gaussian. This was typical for all 21 elements of A. For this element, the linear regression estimate was 11.4, which is close to the mean of the histograms (11.9). (b) Histograms from two different chains for one of the EMG noise magnitudes (expressed as standard deviation rather than variance) for one subject, one trial. Here the distributions are no longer Gaussian and tend to differ more between chains. (c) Histograms of the variability ratio for one trial in one subject for two different chains of the sampling algorithm. Again we have similar and near-Gaussian shapes. For this trial/subject, the entire posterior distribution lies below 1, supporting our hypothesis about task-relevant:task-irrelevant variability.

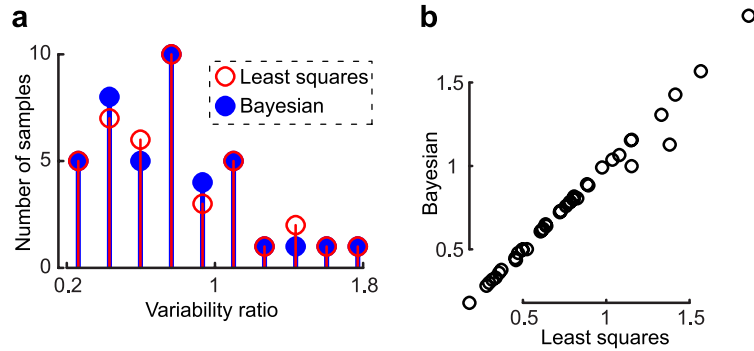


Figure 2: Variability ratios for Bayesian vs. least squares estimates of A. (a) Histogram of the variability index from 41 trials for the Bayesian (filled-in, blue circles) and the least squares (open, red circles) estimates. In case of the Bayesian estimates, the histogram represents the mean of the posterior distribution of the variability ratio. (b) Scatter plot comparing the variability ratio found using least squares vs. that found using the Bayesian method.

the index was significantly smaller than 1, in 9 cases it was significantly larger than 1, and in the remaining 2 cases the result was not significant (meaning that the distribution was centered near 1 and the spread was too large to reach significance in either direction). Note that compared to standard statistical tests, Bayesian tests such as this are much more intuitive as well as accurate. See [Mackay \(2003\)](#) for a discussion of sampling and Bayesian hypothesis testing.

Figure 2b shows a scatter plot comparing the variability ratio computed using least squares and the Bayesian method. Note the almost perfect agreement. The mean was 0.775 for the least squares method and 0.765 for the Bayesian method. This rules out the potential confounds of biased least-squares estimation. To be sure the estimated noise magnitudes differed, but these differences did not interact with our hypothesis. The average (RMS) noise magnitude was 3% MVC for EMGs and 0.18 N for fingertip forces.

References

- D. Mackay. *Information theory, inference and learning algorithms*. Cambridge University Press, Cambridge, U.K., 2003. URL <http://www.inference.phy.cam.ac.uk/mackay/itprnn/book.pdf>.