Published in final edited form as:

Med Biol Eng Comput. 2013 February ; 51(0): 187–196. doi:10.1007/s11517-012-0983-8.

Practical limits on muscle synergy identification by non-negative matrix factorization in systems with mechanical constraints

Thomas J. Burkholder and

School of Applied Physiology, Georgia Institute of Technology, 555 14th St NW, Atlanta, GA 30332-0356, USA, thomas.burkholder@hps.gatech.edu

Interdisciplinary Bioengineering Program, Georgia Institute of Technology, 555 14th St NW, Atlanta, GA 30332-0356, USA

Keith W. van Antwerp

Interdisciplinary Bioengineering Program, Georgia Institute of Technology, 555 14th St NW, Atlanta, GA 30332-0356, USA

Abstract

Statistical decomposition, including non-negative matrix factorization (NMF), is a convenient tool for identifying patterns of structured variability within behavioral motor programs, but it is unclear how the resolved factors relate to actual neural structures. Factors can be extracted from a uniformly sampled, low-dimension command space. In practical application, the command space is limited, either to those activations that perform some task(s) successfully or to activations induced in response to specific perturbations. NMF was applied to muscle activation patterns synthesized from low dimensional, synergy-like control modules mimicking simple task performance or feedback activation from proprioceptive signals. In the task-constrained paradigm, the accuracy of control module recovery was highly dependent on the sampled volume of control space, such that sampling even 50 % of control space produced a substantial degradation in factor accuracy. In the feedback paradigm, NMF was not capable of extracting more than four control modules, even in a mechanical model with seven internal degrees of freedom. Reduced access to the low-dimensional control space imposed by physical constraints may result in substantial distortion of an existing low dimensional controller, such that neither the dimensionality nor the composition of the recovered/extracted factors match the original controller.

Keywords

Neural control; Muscle synergy; Non-negative matrix factorization; Mathematical methods; Biomechanics

1 Introduction

Both voluntary and involuntary movements exhibit reproducible patterns of coordinated muscle activity, and those patterns may reveal underlying neural structures. Coordinated or structured patterns of muscle activation can be evoked by spinal stimulation [16–18], suggesting the existence of spinally-defined motor modules or primitives. These modules can reconstruct relatively complex behaviors [10], and suggest a conceptual model in which the nervous system plans physical actions in a low-dimensional control space which is then

mapped into the high dimensional muscle space through these motor modules. Forces produced by these muscles are then mapped into low dimensional kinematic space by musculoskeletal anatomy. Similar low-dimensionality can be found by statistical decomposition of voluntary muscle activity [20, 23]. Synergies extracted in this way have been found common to diverse voluntary behaviors, which suggests that they represent a neural substrate from which behaviors are constructed [7, 8]. Several statistical methods appear to provide reasonable reconstruction of low dimensional controllers, if they exist, in synthetic data [23].

The use of non-negative matrix factorization (NMF) to identify correlated muscle activations has several conceptual advantages. NMF decomposes a data matrix, $\mathbf{a}$, into a synergy matrix, $\mathbf{W}$, and a command matrix, $\mathbf{k}$, such that $\mathbf{a} = \mathbf{W} \cdot \mathbf{k}$, where the components of $\mathbf{a}$, $\mathbf{W}$, and $\mathbf{k}$ are all non-negative. Muscle activations ($\mathbf{a}$) and neural outputs ($\mathbf{k}$), as measured by force or firing rate, are inherently non-negative, although the gains between neurons within a network may be negative. NMF may identify discrete “features” within a data space with intuitive meaning [15]. For example, Lee and Seung [15] were able to decompose 361-pixel images of faces into 49 very sparse component images suggesting eyes, mouth, and other discrete features of those faces. The identification of features depends strongly on the exact dataset, and on NMF converging to a set of highly sparse factors [12]. It is important to note that much of the machine learning literature supporting NMF is based on data sets with hundreds of measured variables and dozens of independent factors, frequently representing a ten fold or more reduction in dimensionality, where its application in motor control may have a dozen measured variables (muscle EMG) and a handful of independent factors often representing only a two to four fold reduction in dimensionality. NMF does suffer from degeneracy, in that the magnitude of individual factors is arbitrary, but the decomposition is unique if the data derive from a nonnegative generative process [9]. That is, if neural synergies actually have the structure assumed by NMF, then NMF will find the direction, but not magnitude, of those synergies.

There are other potential sources of low-dimensional structure in patterns of muscle activity. In experimental application, muscle activation is constrained to those patterns that successfully perform some reference task or tasks, and patterns of coordination are also predicted by optimal control theory [22]. Experimental tasks are often presented in a very low dimensional form, such as moving a manipulandum to a point in space or maintaining a trajectory. Even behaviors as diverse as running and swimming are performed within a relatively low dimensional kinematic space that may restrict the ability of the nervous system to access the full control space and impose structure as a condition of inclusion [1, 13]. Performance optimization by higher centers may further restrict the use of low dimensional, but space-spanning control structures [22, 25]. Low dimensional coordination among muscles can also be found in feedback contexts [20], in which the musculoskeletal structure may impose a high degree of covariation among proprioceptive inputs to the feedback controller [13, 21]. Low-dimensional structure in muscle activation patterns may arise from true neural modules, from biomechanical structure, or from some combination.

The objective of this work was to determine, in synthetic data, whether NMF can accurately identify a low dimensional control structure within a dataset constrained by mechanical structure. That is, if a system is controlled through a synergy structure but subject to mechanical constraints, do the factors extracted by statistical decomposition represent the synergy structure? Accurate representation requires first accurate determination of the dimension of the control structure, and second accurate composition of individual synergy modules. We use synthetic datasets constructed from our feline hindlimb model [5, 26], using either three or seven internal degrees of freedom. We generate random controllers based on task-level structures, and evaluate the quality of factors extracted by NMF. We will

use “controller” to refer to the low-dimensional structure used to generate muscle activations, and “factor” to refer to the low-dimensional structure produced by NMF. The results indicate that the accuracy with which the factors match the controller depends strongly on sparseness and diversity of the controller modules and on uniformly sampling control space. Mechanical constraints distort the factors enough to make the originating synergy unrecognizable.

2 Methods

Three general methods for generating data are illustrated in Fig. 1. In all cases, a set of control vectors ($\mathbf{W}$) map a point (k) in the low-dimensional command space to the 31-dimensional muscle activation space (a), i.e., $a = \mathbf{W}k$. Muscle moment arms ($\mathbf{R}$) and the endpoint Jacobian ($\mathbf{J}$) map points in muscle space to 3-D Cartesian forces. NMF was performed on the activation data for comparison with the true control modules. In the feedforward paradigm (Fig. 1a), the low-dimensional command space was directly sampled, representing the best-case extraction as a definite reference. In the task-constrained paradigm (Fig. 1b), muscle activations were generated by the same process but the endpoint force produced by each activation pattern was determined, and only those activations that met the task criterion were included in NMF. This paradigm was used to evaluate how systematically undersampling the control space influenced the quality of factor extraction, and it is important to note that no optimization or inverse model was used to select points from command space. In the feedback paradigm (Fig. 1c), random sampling was performed in endpoint space which was used to determine muscle length changes and to generate signals in the low-dimensional command space. This paradigm was used to evaluate how strong structure within the sampling of command space influenced $\mathbf{W}$ extraction.

2.1 Mechanical models

To simplify the analysis, much of this work was performed using a 3 degree of freedom reduction of our 31-muscle cat hindlimb model [5, 26], allowing only hip ab-/ad-duction, knee flexion/extension and ankle flexion/extension. This kinematically constrained model was linearized in a stance-like posture, and was completely described by its endpoint Jacobian ($\mathbf{J}$) and moment arm matrix ($\mathbf{R}$). For feedback simulations, the lack of kinematic redundancy means that there is a one-to-one relationship between endpoint displacements, and $\mathbf{J}$ is invertible. The quasi-static endpoint force, $\mathbf{F}_E$, was calculated from muscle activations, $\mathbf{a}$, as $\mathbf{F}_E = \mathbf{J}^{-T} \mathbf{R} \mathbf{F} \mathbf{a}$, where $\mathbf{F}$ is a diagonal matrix of muscle force capacities. Changes in muscle length ($\mathbf{L}$) due to an endpoint displacement ($d\mathbf{x}$) were calculated as $d\mathbf{L} = \mathbf{R}^T \mathbf{J}^{-1} d\mathbf{x}$.

Further validation of feedback results was performed using a linearized version of the 7-DoF model [4]. Because the 3×7 Jacobian in this model is not invertible, we used muscle-induced joint stiffness to estimate the kinematic response to endpoint displacements [4]. In brief, joint stiffnesses, $\mathbf{K}_J$ were calculated from the partial derivatives of joint torques:

$$\mathbf{K}_J = \partial \mathbf{G} / \partial \theta + \partial (\mathbf{R} \mathbf{F} \mathbf{a}) / \partial \theta + \partial (\mathbf{J}^T \mathbf{F}_E) / \partial \theta$$

where θ is the vector of joint angles, $\mathbf{G}$ torques due to gravity, $\mathbf{R} \mathbf{F} \mathbf{a}$ muscle torques, and $\mathbf{J}^T \mathbf{F}_E$ torques due to an applied endpoint wrench. This estimate for the change in joint torques with respect to changes in joint angles ($\mathbf{K}_J$) is invertible, and allows calculation of changes in muscle length due to an endpoint displacement as $d\mathbf{L} = \mathbf{R}^T \mathbf{K}_J^{-1} \mathbf{J}^T (\mathbf{J} \mathbf{K}_J^{-1} \mathbf{J}^T)^{-1} d\mathbf{x}$. This approximation depends on the muscle activation pattern through $(\mathbf{R} \mathbf{F} \mathbf{a}) / \theta$, and

1,000 different muscle activation patterns satisfying $\mathbf{F}_E = (0.91 \ 7.93 \ -0.54)^T \text{ N}$, an experimentally derived postural force [4, 20], were tested.

2.2 Control models: controller selection

There are many possible strategies for organizing muscles into coherent structures, but for this investigation we have focused on combinations of muscle that generate unique endpoint force directions. Controller definition started with 2–9 points maximally dispersed over the unit sphere [19]. This set was rotated randomly as a group to form the Cartesian basis set ($\mathbf{B}_C$). $\mathbf{B}_C$ was mapped into muscle space by multiplying with $\mathbf{R}^T \mathbf{J}$, and non-negative components randomized to form the activation seeds. Because of the randomization, the seeds do not generate force in the desired Cartesian directions, and quadratic programming was used to find final unit controllers close to the random seed in muscle space and generating force in the direction of the $\mathbf{B}_C$ vector. Quadratic programming minimized the distance between the final controller ($\mathbf{W}$) and the random seed, subject to the constraint that cross product of the muscle-imposed joint torques with the endpoint torques due to the target force was zero: $\mathbf{J}^T \mathbf{B}_{ci} \times \mathbf{R} \mathbf{F} \mathbf{W}_i = 0$, where i indicates matching vectors of $\mathbf{B}_C$ and $\mathbf{W}$. Pilot trials revealed that factor extraction was strongly dependent on the angle between control modules, and this complex process was used to ensure that directional controllers were sparse, substantially different in muscle space, and span Cartesian space.

2.3 Control models: feedforward

Activation data were generated by randomly generating the command vector (k) with components from a uniform distribution. Across repeated trials, this results in uniform sampling of a unit hypercube in command space. Muscle activations could exceed 1.0 under this scheme, and no normalization or truncation of activations was performed. Scaling throughout this work was chosen for mathematical convenience and preservation of manifold shapes, that is, the shape of the low-dimensional command cube projected into the high dimensional muscle space was considered more important than specific numeric limits. In the specific case of NMF, the extraction algorithm is most sensitive to the direction of its edge rays near the origin, and clipping at the top end is not likely to distort the results, while the magnitude of extracted factors is algorithmically arbitrary. Thus, direction is much more relevant than magnitude or scaling. Sample sets using 1,000 or more samples produced essentially identical outcomes, and all results reported here use synthetic data sets with 1,000 samples yielding an activation matrix, a , of $31 \times 1,000$.

2.4 Control models: task-constrained

This paradigm was intended to represent the conceptual case where the nervous system exerts no systematic control to achieve a specific goal, but chooses at random from the entire task null space. Directional controllers for the task constraint analysis were the same as the feedforward set, and the “task” was to produce force within a cone aligned with the negative X axis and having apex angle between 6° and 60° . Due to the limb posture, anatomy and moment arms, the feasible force set is concentrated near the X axis, and randomly generated forces cluster close to this direction. Muscle activations were determined by randomly sampling the command space, as in feedforward. The endpoint force produced by each muscle activation sample was determined, and only those that produced a force within the target zone were retained (Fig. 1b, red, hollow points). This process was repeated until the sample set contained at least 1,000 data points, and the fraction of points retained provides an estimate of the allowed volume of control space or of the stringency of the task constraint.

2.5 Control models: feedback

Muscle activations in feedback simulations were generated by applying endpoint displacements, determining muscle length changes, transforming those length changes by a sensor law into command space and finally by a chosen controller back into muscle activation space (Fig. 1c). Conceptually, this represents a nervous system constructing an internal model based solely on muscle-derived proprioceptive information, then using that model to implement a response. Mathematically, $a = \mathbf{W} * \mathbf{P} * \{\mathbf{R}^T \mathbf{J}^{-1} \mathbf{x}\}$, where $\mathbf{P}$ represents the sensory transformation and $\{\}$ denote that negative values were set to zero or half-wave rectified. Sensor-structures were chosen with dimensions between 2 and 9 by a similar process as the directional controllers, and the directions of $\mathbf{W}$ and $\mathbf{P}$ were not aligned. The limiting case, in which both sensor and control matrices were the identity matrix (31×31), was also considered. This model was termed “Autogenic,” and represents independent, linear activation of each muscle in proportion with its change in length.

2.6 Evaluation

NMF potentially provides two useful features of a large data set: dimensionality and structure. The dimensionality or number of factors to extract was determined by calculating variance accounted for (VAF) using different numbers of factors and selecting the number of factors that maximized VAF. In our noise-free datasets, in contrast to experimental data, VAF has a local maximum, and decreases as the excess factors result from convergence to inappropriate, local minima. Structure was evaluated by comparing the components of extracted factors with the generating control vectors. Extracted factors were matched to control vectors by least-squares, and squared differences in components, normalized to the mean of non-zero component magnitudes, were calculated. These errors follow a steep Chi-squared distribution, in which the mean is about $2.5 \times$ the median, and the 95th percentile is about $10 \times$ the median. The 95th percentile is reported. For each case, we generated 10,000 different control structures, generated an array of 1,000 muscle activation samples for each control structure, and extracted synergies using non-NMF. NMF was performed in Matlab R2011a, using the function `nnmf` from the Statistics toolbox to minimize the sum-squared errors in muscle activations using the alternating least squares algorithm. Pilot experiments indicated that TolX (residual convergence criterion) and TolFun (step size) settings needed to be decreased to $1e-6$, and that 12 replicate runs were required to find a consistent, high quality extraction. Qualitative results were similar for both multiplicative and alternating least squares algorithms, for minimizing either squared errors or “divergence” cost functions [14], with or without an explicit sparseness cost [12].

3 Results

3.1 Feedforward command

When the command space is uniformly sampled and the data are free of noise, the variance accounted for (VAF) by extracted factors has a local maximum where the number of factors matches the number of generating controllers (Fig. 2a). With feedforward sampling, NMF almost always identified the control dimensionality correctly, and extracted factors matched generating vectors with high accuracy (Fig. 2b). A total of 80,000 test sets were considered (10,000 at each of eight dimensions), with only 10 cases where dimensionality incorrectly identified. Both the mean-squared residual and the factor component error increase with the number of dimensions in the controller, as a consequence of step sizes decreasing with extraction dimension. Within one controller dimension, the accuracy of extraction depended on the angle among the control modules. Comparing two sets of six random vectors without any optimization, results in a 95th percentile component difference of 0.54, representing the worst-case extraction. Purely random, 6-D controllers had an angle between vectors of $49^\circ \pm 6^\circ$ and were extracted with 95th percentile component error of 0.02. Six-D controllers

composed of directional control vectors, had an angle between vectors of $67^\circ \pm 2^\circ$. These vectors were extracted with component error of 0.005 (Fig. 2b), so in 6-D controllers, errors of 0.005 were the best solutions, errors of 0.02 reflect controllers with greater directional similarity, and errors of 0.5 are random guesses.

3.2 Task constraint

Because of the random directional structure of the controllers, it was not always possible to find a suitable set of commands to satisfy the task constraint, and command sets in which less than 0.02 % of command space satisfied the task constraint were excluded. This means that, although we started with 10,000 controllers of each dimension, not all of them contributed to this dataset. The reason for this is most apparent for the 2 and 3 dimensional controllers, where accessible forces were constrained to a line or a plane: when that plane lies outside the task cone, it is impossible to satisfy the task constraint. Only 2 % of 2-D and 10 % of 3-D directional control sets were included in the most stringent task, compared with 80 % of 9-D control sets. Fifty-five percent of 3-D controllers were included in the least stringent task, compared to 99 % of 9-D controllers. The 0.02 % of command space exclusion criterion may induce a sampling bias within the lower dimensional controllers.

Restricting the sampling of command space to meet a task constraint does not substantially alter the determination of dimensionality, but does limit the ability to resolve the underlying control structure (Fig. 3). When controller dimension was six or higher, control dimension was sometimes underestimated, and this increased with both task stringency and controller dimension. The 4-dimensional controller was always correctly identified. In the 9-dimensional controller, the 6° task cone retained only 0.6 ± 0.9 % of the command space, and undersampling resulted in an average 1.5-D underestimation with correct dimension identified in only 2.8 % of cases. For the same controller set, the 29° task cone sampled 13 % of command space and dimension was correctly identified for 72 % of cases. Composition of the extracted factors was substantially compromised by the task restriction, with the 60° task cone (37 ± 27 % of command space) resulting in errors about ten fold greater than the unconstrained sampling (Fig. 2). These increased another five fold when control space was restricted below 1 %, but were still twice as accurate as guessing.

Correlation among the commanded activations imposed by the task constraint is the most reasonable source for factor distortion. Correlation coefficients (R^2) between command vectors in the 4-dimensional controller averaged 10^{-5} in the feedforward data sets, 10^{-3} in the least restrictive task set and 0.4 in the most restrictive task set. However, correlations were smaller in the higher-dimension controllers, while factorization errors were greater. In the 9-dimensional controller, the most restrictive task set resulted in correlation coefficient of 0.04 and 95th percentile error of 0.3, while the R^2 of 0.4 at 4-D resulted in errors of 0.2. Errors induced by a restrictive task more strongly distort lower-dimension controllers, and obscure dimensionality of higher-dimension controllers.

3.3 Feedback controllers

In the feedback paradigm, NMF identified six synergies from the 31-dimensional Autogenic controller in the 3-DoF mechanical model, consistent with six non-negative Cartesian dimensions. In the presence of a feedback controller with 4 or fewer dimensions, NMF determined the correct dimensionality (Fig. 4a). For controllers between 5 and 9 dimensions, NMF underestimates the true dimensionality reaching a plateau of 4.1 factors, meaning that NMF most often identified four factors regardless of the true controller dimension. The peak VAF in the feedback extractions was not different from the peak VAF in either feedforward or task extractions, so this underestimation is not reflected in an objectively poor quality of fit.

The accuracy of statistical models derived from the feedback control structure declined sharply in controllers with more than four dimensions. Component errors were calculated for matched controller-factor pairs, and no penalty was assigned for unmatched controllers. For example, in a controller with eight modules that was fit with only four factors, component errors were calculated only between the four factors and their matching controllers. The 95th percentile error jumps from 0.027 in the 4-D controller to 0.16 in the 5-D controller, and 0.29 in the 9-D controller. The largest increase in error, relative to the feedforward solution, was the 5-D controller at $77 \times$ the feedforward error, while errors at higher dimensions were comparable to the most restrictive task extractions. It is difficult to estimate the volume of command space sampled by the feedback controllers because the scaling of muscle length, sensor, and controller transformations are arbitrary. For the 3-D controller, the volume of a convex hull enclosing the commands was 0.004 ± 0.002 (0.4 % of the unit command cube). It is probably more relevant to note that this hull filled 25 ± 6 % of the minimal enclosing cube, and the 4-D controller hull filled 5 ± 2 % of the enclosing hypercube. Convex hulls for five and higher dimension controllers were impractical to calculate, but are certainly smaller than the most restrictive task-constraint sampling volumes by either metric.

To determine whether the limited dimensionality of the feedback model resulted from the Cartesian endpoint space in which perturbations were applied or to the joint space in which the system deformed, the analysis was repeated using seven internal degrees of freedom (3 hip, 2 knee, 2 ankle). In the Autogenic controller, NMF consistently identified seven factors regardless of background activation pattern. This dimensionality is not consistent with six non-negative dimensions of endpoint displacement, nor with 14 non-negative dimensions of joint displacement. Finding greater than 6 dimensions does indicate that the displacement-rectified length change-synergy structure can exceed the dimensionality of its Cartesian input. Muscles can be clustered into seven groups based on pulling direction [3] using the joint-stiffness method to relate endpoint displacement to muscle length change, which may be the source of the 7-D limit. When applied to low-dimensional feedback controllers, NMF again correctly identified the dimensionality of controllers with dimension of 4 or less, but the maximum number of extracted factors was less than 5 (Fig. 4b). Although the plateau is slightly higher with 7 DoF mechanics (4.4 ± 0.6) than with 3 DoF mechanics (4.1 ± 0.6), the limit of 4–5 dimensions is similar to that reported in a number of experimental studies [6, 20, 24]. Factor errors were slightly smaller than with 3 DoF mechanics at dimension 4–6, but increased faster above there, becoming similar to random for the 9-D controller.

4 Discussion

The goal of this study was to determine, using best-case data, whether NMF can extract a low dimensional controller convoluted through a low dimensional mechanical structure. In the absence of mechanical structure, NMF identifies the correct dimensionality and reconstructs the generating controller with high accuracy (error < 0.5 %). When sampling of control space is limited by mechanical constraints, the composition of the controller is distorted by even modest constraints (error < 10 %), while the dimensionality remains identifiable until access to control space is severely limited. Finally, incorporating feedback of a Cartesian signal to drive activation of a low dimensional prevents NMF from identifying control structures with more than four dimensions. There was no post hoc metric to distinguish an accurately identified 4-D controller from a poorly identified 5-D controller.

Statistical decomposition, including NMF, can potentially provide three useful functions. The first is to simplify the technical description of large, multivariate datasets. This does not provide any insight into the mechanisms underlying any extracted correlations, but does provide a quantitative, formulaic method for describing data. For non-negative data, NMF can provide substantial simplification of datasets and may reveal features with intuitive

interpretations to guide further mechanistic study [15]. In general application to arbitrary data, NMF is not guaranteed to provide a unique decomposition, and extraction of features may be strongly dependent on the specific dataset [9, 12]. For data actually generated by non-negative combination of a set of non-negative basis vectors, NMF will return a set of factors parallel to the original basis, unique within a scaling factor [9]. The present simulations demonstrate that, by showing that the components of unit vectors produced in feedforward extractions match the unit generating vectors within 0.5 %. The more interesting functions of NMF potentially provide insight into neuro-mechanical structures that generated the data: dimensionality and structure.

NMF was generally successful at identifying the lowest dimensional structure in the generated data, although this was not necessarily the dimension of the activity controller. Two features deserve special attention: first, the dimensionality was correctly identified under task constraint, and second, the dimensionality in the feedback paradigm was less than 6. In order to satisfy the task constraint, the combination of control weights had to produce force in the assigned Cartesian direction. Because the control vectors were chosen to produce forces individually distributed throughout Cartesian space, this necessarily imposes correlations among the controllers. The magnitude of these correlations increases with task stringency, and decreases with controller dimension. For example, in the limiting case of producing force of arbitrary magnitude exactly along a specified line, a 4-D controller has a null space of 1 (magnitude only) and all four commands would have to be perfectly correlated. However, the dimensionality of higher dimension controllers was more difficult to accurately identify, with their dimension being gradually reduced towards the Cartesian dimension. The intended application of this study was to address whether it is possible, in experimental application, to extract a hidden neural controller, or whether the kinematic limitations necessarily obscure or distort such controllers. The “task” chosen for this work was different from experimental tasks involving physical motion, such as walking or swimming [8, 24], but was conceptually similar to postural tasks [20, 25]. The task constraint had to be very stringent before substantial reduction of dimensionality occurred, which suggests that it is possible to identify the dimensionality of a neural synergy structure, if one exists, from experimental data.

Within the feedback paradigm, factor dimensionality was restricted to less than 5. The non-negative Cartesian input to the system was dimension 6, was found in the 3-DoF Autogenic controller, and represents a limiting value for dimension extraction. The convolution of mechanical structure with a low dimensional control structure restricted the apparent complexity of the system to 4-D. In the feedback models used here, muscle length changes are half-wave rectified before transformation by the sensor model (P). Without rectification, sensor responses would have cosine tuning relative to their preferred endpoint directions, and the Cartesian sphere would be an ellipsoid in command space. With rectification, this ellipsoid is projected onto the non-negative sector of control space, and the majority of the ellipsoids lie outside of this sector. This sensory structure was chosen because it incorporates the asymmetric length response observed in muscle spindles [11] and the cosine tuning observed in spinal interneurons [2], but this specific sensory structure does strongly affect the factorization outcomes. In application to experimental interpretation, the result suggests that experimental paradigms that incorporate a strong proprioceptive component may have synergy-like structures imposed by anatomical (mechanical) structures, and that anatomical structures may obscure existing neural complexity. It is interesting to note that several studies report 4–5 synergies as sufficient to describe human and animal motor control [6, 20, 24].

The final potential product of statistical decomposition is the actual identity or composition of a hidden neural structure, which might reveal the higher-level framework in which the

nervous system plans movements. Results presented here indicate that NMF can reproduce the structure of a set of non-negative generating controls under ideal conditions, but that the accuracy of those extractions is compromised under the least restrictive task conditions. When NMF extracted fewer than the true control vectors, it did not extract a subset of the true controllers but rather distortions of those controllers. Thus, this outcome extends recent report that autogenic feedback produces the appearance of low-dimensional control [13], by demonstrating that mechanical structure also distorts any low-dimensional control structure that is present. Qualitatively, these matches were not good. In component-by-component examination factor, errors of 0.2 and greater resulted in about half of the large (top 10 % magnitude) contributors to the true controller having small (bottom 50 % magnitude) components in the extracted factors, and about half of the large factor components derived from small controller components. In experimental application, where scaling between recorded EMG and descending command is somewhat speculative, qualitatively identifying the major contributors seems like a reasonable result, and NMF does not meet this goal in the presence of proprioceptive feedback.

These claims should be subject to a number of caveats. All synthetic datasets were generated from a single posture, and it is often argued that observing either the same endpoint force or the same pattern of coordinated muscle activity, across different postures, with different muscle moment arms and different Jacobians, is the hallmark of neural synergies [10, 21]. The test controllers were selected based on generation of endpoint forces, but there are many strategies the nervous system might use to organize synergies. Much of the distortion we report is due to the mapping of endpoint into muscle space, and a sensory-motor system based in joint coordinates would not be subject to that distortion. The mapping of endpoint perturbations to muscle length changes in the 7-DoF mechanical model used a linearized model that ignored both inertia and viscosity-like muscle properties, and used “background” muscle activation patterns with a high degree of cocontraction. While this may distort the relationship between endpoint perturbations and the “proprioceptive input” to the controller, we have previously shown that the linearized model accurately represents the full, nonlinear system response for at least 300 ms [4], and found comparatively little difference in joint displacements among the background activation patterns. Both the sensor and synergy matrices were chosen arbitrarily, were not coordinated with each other, and may not represent biologically meaningful structures. These transformations were selected to mimic observed proprioceptive characteristics, and to provide the greatest potential for accurate synergy extraction.

In summary, the major constraint on accurately extracting legitimate, low dimensional control structures is completeness of sampling the control space. In the case where control modules exist, a restrictive task or task set may force covariance among those control modules that distorts their appearance. It is possible that this phenomenon accounts for some reports of synergies “splitting” or “merging” under experimental or pathological conditions, where posture or mechanics are altered. In the case of feedback control structures, a kinematically constrained system (i.e., 3 kinematic DoF and 3 endpoint DoF) imposes an upper limit of 6 non-negative dimensions on structured variability among muscle activations, and a practical limit of four on the number of identifiable control modules. Relaxing the kinematic constraint by adding internal kinematic degrees of freedom does not increase these limits. Finally, although this work focused on NMF, as best matched to our synthetic control structures, it seems likely that these limitations apply to all statistical decomposition methods because they reflect the intersection of a synergy-imposed manifold in muscle space with a mechanically-imposed manifold in command space.

Acknowledgments

This work was supported by the National Institutes of Health (HD046922, HD032571). The NIH had no role in the design, performance or interpretation of the study, nor in the decision to publish.

References

1. Berniker M, Jarc A, Bizzi E, Tresch MC. Simplified and effective motor control based on muscle synergies to exploit mus-culoskeletal dynamics. *Proc Natl Acad Sci USA*. 2009; 106:7601–7606. [PubMed: 19380738]
2. Bosco G, Poppele RE. Broad directional tuning in spinal projections to the cerebellum. *J Neurophysiol*. 1993; 70:863–866. [PubMed: 8410178]
3. Bunderson NE, Burkholder TJ, Ting LH. Reduction of neuromuscular redundancy for postural force generation using an intrinsic stability criterion. *J Biomech*. 2008; 41:1537–1544. [PubMed: 18374342]
4. Bunderson NE, McKay JL, Ting LH, Burkholder TJ. Directional constraint of endpoint force emerges from hindlimb anatomy. *J Exp Biol*. 2010; 213:2131–2141. [PubMed: 20511528]
5. Burkholder TJ, Nichols TR. Three-dimensional model of the feline hindlimb. *J Morphol*. 2004; 261:118–129. [PubMed: 15164372]
6. Clark DJ, Ting LH, Zajac FE, Neptune RR, Kautz SA. Merging of healthy motor modules predicts reduced locomotor performance and muscle coordination complexity post-stroke. *J Neurophysiol*. 2010; 103:844–857. [PubMed: 20007501]
7. d'Avella A, Bizzi E. Shared and specific muscle synergies in natural motor behaviors. *Proc Natl Acad Sci USA*. 2005; 102:3076–3081. [PubMed: 15708969]
8. d'Avella A, Saltiel P, Bizzi E. Combinations of muscle synergies in the construction of a natural motor behavior. *Nat Neurosci*. 2003; 6:300–308. [PubMed: 12563264]
9. Donoho D, Stodden V. When does non-negative matrix factorization give a correct decomposition into parts? *Adv Neur In*. 2003; 16:1141–1148.
10. Giszter SF, Mussa-Ivaldi FA, Bizzi E. Convergent force fields organized in the frog's spinal cord. *J Neurosci*. 1993; 13:467–491. [PubMed: 8426224]
11. Houk JC, Rymer WZ, Crago PE. Dependence of dynamic response of spindle receptors on muscle length and velocity. *J Neurophysiol*. 1981; 46:143–166. [PubMed: 6455505]
12. Hoyer PO. Non-negative matrix factorization with sparseness constraints. *J Mach Learn Res*. 2004; 5:1457–1469.
13. Kutch JJ, Valero-Cuevas FJ. Challenges and new approaches to proving the existence of muscle synergies of neural origin. *PLOS Comput Biol*. 2012; 8(5):e1002434. [PubMed: 22570602]
14. Lee DD, Seung HS. Algorithms for non-negative matrix factorization. *Adv Neural Inf Process Syst*. 2001; 13(13):556–562.
15. Lee DD, Seung HS. Learning the parts of objects by nonnegative matrix factorization. *Nature*. 1999; 401:788–791. [PubMed: 10548103]
16. Lemay MA, Galagan JE, Hogan N, Bizzi E. Modulation and vectorial summation of the spinalized frog's hindlimb endpoint force produced by intraspinal electrical stimulation of the cord. *IEEE Trans Neural Syst Rehabil Eng*. 2001; 9:12–23. [PubMed: 11482358]
17. Lemay MA, Grill WM. Modularity of motor output evoked by intraspinal microstimulation in cats. *J Neurophysiol*. 2004; 91:502–514. [PubMed: 14523079]
18. Saltiel P, Wyler-Duda K, D'Avella A, Tresch MC, Bizzi E. Muscle synergies encoded within the spinal cord: evidence from focal intraspinal NMDA iontophoresis in the frog. *J Neurophysiol*. 2001; 85:605–619. [PubMed: 11160497]
19. Sloane, NJA.; Hardin, RH.; Smith, WW. [22 Mar 2012] Tables of spherical codes. 2000. <http://www.research.att.com/~njas/packings/>
20. Ting LH, Macpherson JM. A limited set of muscle synergies for force control during a postural task. *J Neurophysiol*. 2005; 93:609–613. [PubMed: 15342720]
21. Ting LH, McKay JL. Neuromechanics of muscle synergies for posture and movement. *Curr Opin Neurobiol*. 2007; 17:622–628. [PubMed: 18304801]

22. Todorov E. Optimality principles in sensorimotor control. *Nat Neurosci.* 2004; 7:907–915. [PubMed: 15332089]
23. 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.* 2006; 95:2199–2212. [PubMed: 16394079]
24. Tresch MC, Saltiel P, Bizzi E. The construction of movement by the spinal cord. *Nat Neurosci.* 1999; 2:162–167. [PubMed: 10195201]
25. Valero-Cuevas FJ, Venkadesan M, Todorov E. Structured variability of muscle activations supports the minimal intervention principle of motor control. *J Neurophysiol.* 2009; 102:59–68. [PubMed: 19369362]
26. van Antwerp KW, Burkholder TJ, Ting LH. Inter-joint coupling effects on muscle contributions to endpoint force and acceleration in a musculoskeletal model of the cat hindlimb. *J Biomech.* 2007; 40:3570–3579. [PubMed: 17640652]

Fig. 1.

Illustrative data in five relevant coordinate systems. For spaces with more than three dimensions, three of those dimensions were arbitrarily chosen for presentation. *Column headings* give the name used for each space and terms in the *bottom row* give the mapping between spaces. Feedforward (a) and task-constrained (b) paradigms directly sample the low-dimensional command space and produce high-dimensional muscle activations by linear transformation. The hypercube of command samples can be seen, stretched and rotated, in muscle activation. Activations are mapped into endpoint forces through the Jacobian, moment arm matrix, and maximum muscle force scaling. The task constraint is easily seen in force space as a *wedge of hollow, red points*. In muscle and command space, these points appear intermingled with task-excluded points because of the projection of higher dimensions into the 3-D plot. Feedback paradigm (c) samples endpoint displacement, which uniquely determined muscle length changes. Length changes determine low-dimensional, task-relevant command signals through the linear mapping P , which are used to produce high-dimensional muscle activations as in feedforward and task-constrained

Fig. 2.

a Variance accounted for (VAF) as a function of the number of factors extracted from data sets generated using either six (*circle*) or nine (*triangle*) dimensional controllers (Mean $\pm$ SD). **b** Factor extraction metrics for 2–9 directional synergies extracted from uniformly sampled control space

Fig. 3.

Factor error increases with task stringency. For controllers with 4 (a), 6 (b) or 9 (c) dimensions, control space was randomly sampled, and only those commands that produced force within a target cone were retained. Circles (left hand scale) represent the mean ($\pm$ SD) dimension of the resulting muscle activations, determined by maximum-VAF criterion and the mean ($\pm$ SD) volume fraction of the sampled control space, for cone angles ranging from 6° to 60° . Apex angles of the task cones are shown below corresponding points (a) and are the same in (b) and (c). Lines (right hand scale) represent the 95th percentile component error as a fraction of the mean component value

Fig. 4.

The optimum number of factors (Mean $\pm$ SD), as determined by maximum VAF, extracted from the feedback paradigm for controllers with dimension 2–9 and the 95th percentile error in the extracted factor components, using a mechanical model with three internal degrees of freedom (a) or with seven internal degrees of freedom (b). Factor number matches controller dimension up to 4, but reaches a plateau of 4.1 ± 0.6 with 3 DoF mechanics or 4.4 ± 0.6 with 7 DoF mechanics. Factor component errors are large with five or more dimensional controllers