Spatial and spatio-temporal characterization of movement for the analysis of actions and actors

Leonardo Claudino and Yiannis Aloimonos
Department of Computer Science
University of Maryland at College Park
{claudino,aloimonos}@cs.umd.edu

Abstract

Movement data is high-dimensional but often redundant, meaning there is certainly a lower dimensional subspace that spans most of the body configurations within an action performance. We propose that one such representation can be achieved through a decomposition method that explores the existence of key configurations and temporal correlations of those configurations that are typical of action matrices. The approach is compatible with computational models of motor synergies based on matrix factorizations, and it builds upon a method that was earlier proposed in the context of biological motion perception. Our experiments show that vertical jump trials collected from children and young adults can be consistently reconstructed from the resulting representation. We also observe that a subset of that same representation suggests differences among populations of jumpers based on their trials, which serves to illustrate the potential of the method as a tool to analyze both actions and actors.

1 Introduction

Motion capture technology is getting cheaper, more diverse and achieving higher throughput. While raw movement data is very high-dimensional by nature, it is also highly redundant, both at the level of body configurations (e.g. limb positions, postures or joint angles) and the occurrences of such configurations along the timeline of the action; in other words: in both space and time. Therefore, to allow for meaningful computational exploration of hidden patterns of activity, these data must first be shaped into some lower dimensional representation to be manipulated later.

Along these lines, we propose that a suitable compact representation can be obtained by decomposing an action matrix of size T (time instants) $\times$ J (body configurations) with a method that consists of (1) finding a vector space corresponding to key configurations and (2) fitting a mixture model to the projections of the matrix onto each of the individual vectors found, that is, imposing a parametric model to the temporal correlations of each of the body configurations represented by the vectors. Throughout the document, we refer to the computed vector space as the *spatial basis* of the action and we call the parameter set of the mixture model the *spatio-temporal* representation. The approach was tested on optical motion capture data of adults, typically developing children and children with Developmental Coordination Disorder (DCD), all performing vertical jumps. Results suggested it is indeed an adequate action representation for it was observed to allow for proper reconstruction of the jumps and also helped in the inspection of both similarities and differences among populations of jumpers based on their trial distributions.

Our work is related to research in computational models of *motor synergies*, more specifically the models involving matrix factorizations. The method we describe builds upon an

existent computational model of *biological motion perception*, proposed by vision psychologists. In the next section, we sketch some of the ideas behind these two related disciplines – one motor, another perceptual – so that the scope of the work is properly situated. In sequence, we formalize the computation of spatial and spatio-temporal representations and present the experiments with the jump data set, along with our main findings.

2 Related work: compact representations in the motor and visual spaces

Synergies, *motor programs* or *action primitives* are only a few of the many ways to denote a hypothetical set of pre-existing modules of effector activation that would be combined by the central nervous system to produce action. Many believe this is the way the brain cuts down dimensionality when controlling and coordinating multiple degrees of freedom in space and time, the so called "Degrees of freedom (DOF) problem" and it came out of the first round of investigations of Bernstein's work in control and coordination, as once posed by Turvey [1]. This problem has been recently revisited by Latash et al. [2, 3] who discuss the related "principle of abundance", which refers to the fact that a task demands less degrees of freedom than what is available to be controlled. See Flash et al. [4] for a very interesting summary of findings around the nature of motor primitives at behavioral, muscle, neural, and computational levels.

Previous electrophysiological experiments in spinalized animals have indeed presented strong evidence supporting the existence of basic modules of movement that would be additively combined to produce behavior [6]. These modules also seem to be connected to how the same activities are perceived, which is referred by Turvey [1] as "simultaneous organization of afferentiation and efferentiation". The vast applicability of modeling movement signals based on a compact set of primitives makes the quest for motor synergies an active research topic across many different communities, namely cognitive and humanoid robotics, kinesiology and movement psychology. Theories around the nature of synergies have been proposed in terms of spinal force-fields [6, 7, 8], time-varying synergies of muscle forces (TVMS) [9, 10, 11], joint-angle configurations [14], and uncontrolled manifolds [2, 3, 5], among others.

Assuming that these action modules actually exist, the way they would be obtained from collected data is still a mystery. The SB-ST decomposition method that will be covered in the next section has aspects in common with attempts to look for synergies based on matrix factorizations, some of which have been experimented and reported successful in both artificial and real motor data, at least in terms of reconstruction power and consistency of factors and coefficients [16]. Ivanenko et al. [12] tried different factorization methods on EMG data from 8 subjects performing walking actions and walking plus voluntary behaviors, such as kicking a ball or overcoming an obstacle along the way, and found factors to be consistent for the different methods. When they inspected the spatio-temporal profiles of factors (they call it activation components) in ascending order of peak activation time, they realized that all voluntary behaviors agreed upon the same five first components, which happened to be very similar to the first five components of walking. Later, Ivanenko, Popelle et al. [13] presented a spatio-temporal map of 32 muscles along the rostrocaudal axis of the spinal cord (as an indirect measure of the firing of spinal α-motorneurons innervating those muscles) for a single walking cycle, and again noted five factors with main peaks that coincided with major kinematic and kinetic events of walking. But the main finding of the two accounts

was a sixth profile that was distinct among all voluntary behaviors, with main peaks at different stages of the cycle. Authors then proposed an additive model that coordinates locomotion and voluntary actions by the superimposition of the discovered extra activation program onto the five walking programs. Ivanenko, Popelle et al. [13] observed that activation programs should comprise a single Gaussian each, because when single Gaussians with standard deviations fixed at 6% of the cycle duration were fit to the activation programs of each of the factors, they were shown to account for about 90% of the data variance. They later concluded: "It appears that the question of muscle activation during locomotion can be divided into two parts, the timing of activation and the specification of which muscles to activate. In fact the necessity for such a dual motor pattern representation has given rise to a nontraditional definition of muscle synergies: a time-varying muscle synergy". In other words, both the muscle ensembles (which we extrapolate here to the consequent body configurations and refer to as spatial basis) and their time profiles (we refer to as spatio-temporal basis, since these are profiles linked to a particular body configuration) are parametrized. This distinction is critical when one wants to analyze spatio-temporal sensitive movement phenomena such as coordination disorders. SB-ST also shares common aspects with the referred time-varying muscle synergy (TVMS) model, as will be further discussed.

Compact representations of movement have also been pursued by vision psychologists while trying to computationally model the visual phenomenon referred to as *biological motion perception* – a term coined to express the ability of humans to perceive moving dots from point-light displays as coherent articulated rigid bodies that give rise to the perception of classes of activities [17, 18, 19]. Of particular relevance, Troje [20] has offered a computational method that produces walking patterns and it is able to discriminate between male and female walks from point-light displays coming from 3-D motion capture positional data. He also breaks up the action analysis into spatial and spatio-temporal representations, although he does not actually use this terminology or address the decomposition in a more general way. The spatial basis described in this paper is analogous to his *eigenpostures*, which represented the 4 first principal components of a single-walker data matrix of dimension $T \times 3J$, with T corresponding to the length of the walk trial (rows) and $3J$ being the joint-$\{x, y, z\}$ coordinates (columns). He modeled the temporal occurrences of the eigenpostures with a family of sine functions, for which he determined the 4 fundamental frequencies and the relative phases of components 2, 3 and 4 with respect to the first. His sine functions correspond to the basis functions in the spatio-temporal representation presented here.

Troje's work illustrates that a lower dimensional generative action model is able to both produce the action efficiently and discriminate among performing actors, and his results provided us with the necessary inspiration to re-address his ideas in a more generally applicable model. His choice of modeling the temporal activation of spatial basis vectors as a sum of sinusoids is more suitable to pick patterns that appear throughout the whole timeline of the action but will miss localized events that can reveal coordination differences across populations of performers. Instead, we approach the problem as a matter of fitting a mixture model to the time series and we foster the use of a powerful non-linear least-squares regression tool, Variable Projection or VARPRO [21], which has been used in many different scenarios after almost 40 years of its existence [22], but (to the best of our knowledge) it has not been applied in the context of movement data decomposition. The idea of fitting a mixture model to the temporal activations is consistent with the locomotory model of Ivanenko, Popelle et al. although they suggest a single Gaussian-shaped activation component per

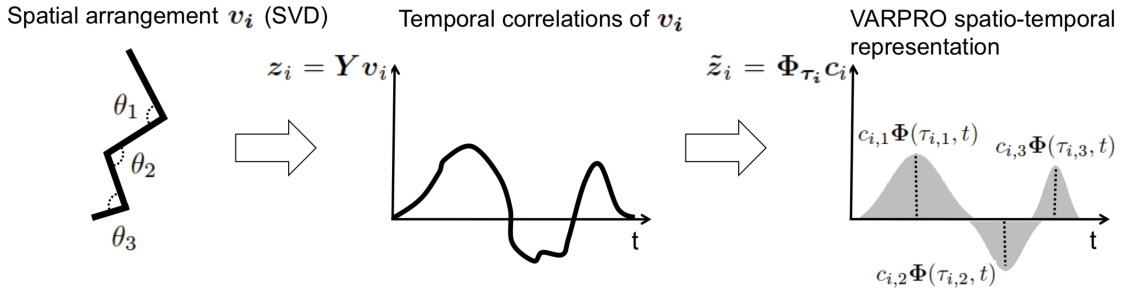

Figure 1: In the example above, J-dimensional spatial basis vector $\boldsymbol{v}_i$ encodes a linear combination of joint angles θ_1, θ_2 and θ_3 computed with SVD, as shown by the leftmost figure. The projection $\boldsymbol{z}_i = \boldsymbol{Y}\boldsymbol{v}_i$ of action matrix $\boldsymbol{Y}_{T\times J}$ onto $\boldsymbol{v}_i$ results into an often smooth (therefore differentiable) temporal series of correlations that represents the activity of that particular spatial arrangement (posture) along the timeline of the action (center figure). We use VARPRO to produce a compact parametric representation for this temporal behavior by fitting a mixture of $\tilde{\boldsymbol{z}}_i = \boldsymbol{\Phi}_{\boldsymbol{\tau_i}}\boldsymbol{c}_i$ to $\boldsymbol{z}_i$ (right figure) which results in parameter vectors $\boldsymbol{\tau}_i = \{\tau_{i,1}, \tau_{i,2}, \tau_{i,3}\}$ and $\boldsymbol{c}_i = \{c_{i,1}, c_{i,2}, c_{i,3}\}$. An action matrix is therefore fully characterized by each spatial basis vector $\boldsymbol{v}_i$ and corresponding set of spatio-temporal parameter vectors $\boldsymbol{\tau}_i$ and $\boldsymbol{c}_i$. See text for more details.

factor-cycle. For the particular case of vertical jumps, we have observed these activations to reveal more than one mode, so a mixture appeared more appropriate and conceptually more general. Besides, even though we will present the decomposition based on Gaussians, the VARPRO framework can accommodate other functions. Figure 1 summarizes the method, which is covered in the next section.

3 SB-ST action decomposition

3.1 Computing a spatial basis and spatio-temporal representations

Let $\boldsymbol{Y}_{T\times J}$ be a multi-dimensional action signal, for example, a T-length sequence of J-joint angle configurations. The k-th order approximation of that signal by SVD, in matrix notation is:

$$\hat{\boldsymbol{Y}}_{T\times J} = \boldsymbol{z}_1\boldsymbol{v}_1^\top + \boldsymbol{z}_2\boldsymbol{v}_2^\top + \ldots + \boldsymbol{z}_k\boldsymbol{v}_k^\top,$$

where $\boldsymbol{v}_i$ is one of the top k right singular vectors of $\boldsymbol{Y}$, therefore spanning the column (postural) space of that matrix, and projection $\boldsymbol{z}_i = \boldsymbol{Y}\boldsymbol{v}_i$ corresponds to the one-dimensional time series that expresses the correlations of the particular spatial configuration represented by $\boldsymbol{v}_i$ along the timeline of the action[1].

For each i, let $\{\boldsymbol{\Phi}(\tau_{i,j}, t) : j = 1 \ldots \boldsymbol{N}_i\}$ be a family of $\boldsymbol{N}_i$ Gaussians. Also, let the mean vector $\boldsymbol{\tau}_i = \{\tau_{i,1}, \tau_{i,2} \ \ldots \ \tau_{i,\boldsymbol{N}_i}\}$ be the only relevant set of parameters, i.e. let the Gaussian functions have fixed standard deviations. Consider $\boldsymbol{\Phi}_{\boldsymbol{\tau_i}}$ to be the corresponding $T \times \boldsymbol{N}_i$

[1]Note that, for right singular vector $\boldsymbol{v}_i$, $\boldsymbol{Y}\boldsymbol{v}_i = \sigma_i\boldsymbol{u}_i$, with σ_i being the i-th singular value and $\boldsymbol{u}_i$ the i-th left singular vector. We chose to use $\boldsymbol{Y}\boldsymbol{v}_i$ rather than $\boldsymbol{u}_i$ through the text just to emphasize that vector $\boldsymbol{z}_i$ expresses a time series of correlations between the data matrix $\boldsymbol{Y}$ and the particular spatial configuration $\boldsymbol{v}_i$.

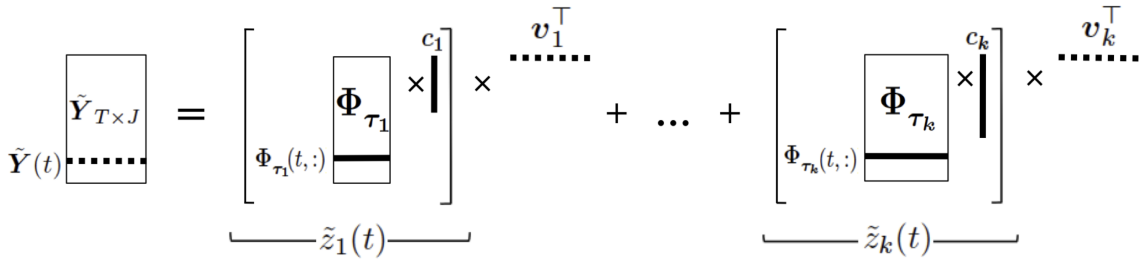

Figure 2: Approximating $\boldsymbol{Y}(t)$ as a linear combination of spatial basis vectors $\boldsymbol{v}_1, \boldsymbol{v}_2 \ \ldots \ \boldsymbol{v}_k$ (dashed lines), as in Equation 1. Coefficients $\tilde{z}_i(t)$ of each vector $\boldsymbol{v}_i$ are the product of the t-th time row of its spatio-temporal matrix $\boldsymbol{\Phi}_{\boldsymbol{\tau_i}}$ and respective linear parameter vector $\boldsymbol{c}_i$ (solid lines).

matrix such that each function is sampled at T instants and it becomes a column of that matrix. We will now model $\boldsymbol{z}_i$ by fitting a linear combination of the columns of $\boldsymbol{\Phi}_{\boldsymbol{\tau_i}}$ with linear parameters $\boldsymbol{c}_i = \{c_{i,1}, c_{i,2} \ \ldots \ c_{i,\boldsymbol{N}_i}\}$:

$$\tilde{\boldsymbol{Y}}_{T\times J} = \underbrace{(\boldsymbol{\Phi}_{\boldsymbol{\tau_1}}\boldsymbol{c}_1)}_{\tilde{\boldsymbol{z}}_1}\boldsymbol{v}_1^\top + \underbrace{(\boldsymbol{\Phi}_{\boldsymbol{\tau_2}}\boldsymbol{c}_2)}_{\tilde{\boldsymbol{z}}_2}\boldsymbol{v}_2^\top + \ldots + \underbrace{(\boldsymbol{\Phi}_{\boldsymbol{\tau_k}}\boldsymbol{c}_k)}_{\tilde{\boldsymbol{z}}_k}\boldsymbol{v}_k^\top,$$

and we have $\tilde{\boldsymbol{z}}_i = \boldsymbol{\Phi}_{\boldsymbol{\tau_i}}\boldsymbol{c}_i$. Equivalently, the posture produced by the model at time t is:

$$\tilde{\boldsymbol{Y}}(t) = \tilde{z}_1(t)\boldsymbol{v}_1^\top + \tilde{z}_2(t)\boldsymbol{v}_2^\top + \ldots + \tilde{z}_k(t)\boldsymbol{v}_k^\top, \tag{1}$$

where $\tilde{z}_i(t) = c_{i,1}\boldsymbol{\Phi}(\tau_{i,1}, t) + c_{i,2}\boldsymbol{\Phi}(\tau_{i,2}, t) + \ldots + c_{i,\boldsymbol{N}_i}\boldsymbol{\Phi}(\tau_{i,\boldsymbol{N}_i}, t)$. The schema in Figure 2 illustrates how a posture $\tilde{\boldsymbol{Y}}(t)$ is generated.

Vector $\boldsymbol{v}_i$ corresponds to the i-th *spatial basis* (SB) vector of action matrix $\boldsymbol{Y}$ or SB-i. Basis functions $\boldsymbol{\Phi}(\tau_{i,j}, t)$ (and, equivalently, its matrix version $\boldsymbol{\Phi}_{\boldsymbol{\tau_i}}$) together with the mean vector $\boldsymbol{\tau}_i$ and the linear parameter vector $\boldsymbol{c}_i$ constitute what we call the i-th *spatio-temporal representation* (ST) of $\boldsymbol{Y}$ or ST-i. We are now left with the task of solving for ST-i parameters $\boldsymbol{\tau}_i$ and $\boldsymbol{c}_i$.

3.2 Solving for ST-i parameters with VARPRO

Because $\boldsymbol{\Phi}(\tau_{i,j}, t)$ was chosen to be a family of single-parameter Gaussians, this problem turns out to be a separable least-squares regression problem, which allows us to solve for $\boldsymbol{\tau}_i$ and $\boldsymbol{c}_i$ using variable projection (VARPRO) proposed in the early 1970's by Golub & Pereira [21]. The method exploits the linear substructure of this particular case of nonlinear least squares (NLLS) regression: if you fix the set of non-linear parameters $\boldsymbol{\tau}_i$, the problem turns out to be linear in $\boldsymbol{c}_i$ and can be solved for the latter using linear least squares (LLS) technology. In other words, parameter $\boldsymbol{c}_i$ becomes a function of parameters $\boldsymbol{\tau}_i$ and so, instead of solving:

$$\min_{\boldsymbol{\tau}_i, \boldsymbol{c}_i} ||\boldsymbol{z}_i - \tilde{\boldsymbol{z}}_i(\boldsymbol{\tau}_i, \boldsymbol{c}_i)||_2^2,$$

we solve:

$$\min_{\boldsymbol{\tau}_i} ||\boldsymbol{z}_i - \tilde{\boldsymbol{z}}_i(\boldsymbol{c}_i(\boldsymbol{\tau}_i))||_2^2.$$

Note that this is now a less parametrized problem, a clear advantage of the VARPRO framework. In the LLS stage, the pseudo-inverse solution for $\boldsymbol{c}_i$ is:

$$\tilde{\boldsymbol{c}}_i = [\boldsymbol{\Phi}_{\boldsymbol{\tau_i}}]^{\dagger} \boldsymbol{z}_i. \tag{2}$$

Recall that $\boldsymbol{z}_i = \boldsymbol{Y}\boldsymbol{v}_i$ is computed by projecting data matrix $\boldsymbol{Y}$ onto spatial basis vector $\boldsymbol{v}_i$ computed with SVD, and vector $\tilde{\boldsymbol{z}}_i$ is the current VARPRO approximation. The solution can be expressed in terms of the truncated SVD of $\boldsymbol{\Phi}_{\boldsymbol{\tau_i}}$:

$$\tilde{\boldsymbol{c}}_i = \boldsymbol{V}\tilde{\boldsymbol{\Sigma}}^{\boldsymbol{-1}}\boldsymbol{U}^{\top}\boldsymbol{z}_i. \tag{3}$$

The LLS solution is then directly embedded in the calculation of the Jacobian of $\tilde{\boldsymbol{z}}_{\boldsymbol{i}}(\boldsymbol{c}_i(\boldsymbol{\tau}_i))$ for the NLLS part of the optimization. As in [23], the Jacobian can be expressed as a sum of two matrices:

$$\boldsymbol{J} = -(\boldsymbol{A} + \boldsymbol{B}), \tag{4}$$

where each of their $\boldsymbol{N}_i$ columns are:

$$\begin{aligned} \boldsymbol{a}_j &= \boldsymbol{D}_j\tilde{\boldsymbol{c}}_i - \boldsymbol{U}(\boldsymbol{U}^{\top}(\boldsymbol{D}_j\tilde{\boldsymbol{c}}_i)), \\ \boldsymbol{b}_j &= \boldsymbol{U}(\boldsymbol{\Sigma}^{-1}(\boldsymbol{V}^{\top}(\boldsymbol{D}_j^{\top}\boldsymbol{r}))). \end{aligned} \tag{5}$$

Here, $\boldsymbol{D}_j$ is a matrix with zeros at all columns but j, which will have the partial derivatives of the j-th Gaussian $\boldsymbol{\Phi}(\tau_{i,j}, t)$ (or the j-th column of matrix $\boldsymbol{\Phi}_{\boldsymbol{\tau_i}}$) w.r.t. its mean $\tau_{i,j}$, evaluated at all time instants t.

Matrices $\boldsymbol{U}$, $\tilde{\boldsymbol{\Sigma}}^{\boldsymbol{-1}}$ and $\boldsymbol{V}$ come from the truncated SVD of $\boldsymbol{\Phi}_{\boldsymbol{\tau_i}}$ as in Equation 2. Vector $\boldsymbol{r}$ is the residual $\boldsymbol{z}_i - \tilde{\boldsymbol{z}}_i$. Operations were grouped so that only matrix-vector product multiplications are required, as in O'Leary & Rust [23], who also propose modifications to the way both the partial derivatives and the Jacobian are stored to exploit sparseness. The presented SB-ST decomposition and our VARPRO implementation are summarized in Algorithms 1 and 2, respectively. MATLAB® sample source codes are attached to this submission as supplemental materials.

3.3 Relationship with time-varying muscle synergies model

Time-varying muscle synergies (TVMS) is a compositional model that produces an action as a linear combination of asynchronous, short-length time series of muscle assemblies. It was proposed by d'Avella et al. [9, 10, 11] to characterize relative muscle coordination in both space and time, after their findings from frog natural behavior data. It is similar to SB-ST, in it also approximates the temporal evolution of a multi-dimensional action vector with k components, which according to our notation would be:

$$\hat{\boldsymbol{Y}}(t) = z_1\boldsymbol{v}_1(t-\tau_1)^{\top} + z_2\boldsymbol{v}_2(t-\tau_2)^{\top} + \ldots + z_k\boldsymbol{v}_k(t-\tau_k)^{\top}, \tag{6}$$

where $\boldsymbol{v}_i(t-\tau_i)$ is a J-dimensional vector that derives from one of N *synergy* matrices $\boldsymbol{V}_i$ of dimension J(muscles) $\times$ Q(discrete time units). It is equal to the null vector whenever $t - \tau_i < 0$ (t is a point in time earlier than when synergy $\boldsymbol{V}_i$ is supposed to be active) or $t - \tau_i \geq T_i$ (t is past $\boldsymbol{V}_i$'s activation).

The latest edition of TVMS [11] is a more elaborated version of the previous two [9, 10], and it determines $\boldsymbol{v}_i(t-\tau_i)$ and z_i through a non-negative factorization of $\boldsymbol{Y}$ similar to what was also proposed in [24]. They find delays τ_i by using a greedy technique that resembles

Algorithm 1 : $[\boldsymbol{v}_i,\, \boldsymbol{c}_i,\, \boldsymbol{\tau}_i,\, \tilde{\boldsymbol{Y}}]$ = SB_ST_decomposition($\boldsymbol{Y}$, k, $\boldsymbol{N}_i$)

Compute $[\mathbf{U}, \boldsymbol{\Sigma}, \mathbf{V}]$ = SVD of $\boldsymbol{Y}$
for $i = 1$ to the first k columns $\boldsymbol{v}_i$ of $\mathbf{V}$ (SB-i vectors) **do**
  Form $\boldsymbol{z}_i = \boldsymbol{Y}\boldsymbol{v}_i$
  Form approximation $\tilde{\boldsymbol{z}}_{\boldsymbol{i}}$ by:
  1. finding optimal $\boldsymbol{\tau}_i$ with a NLLS solver that calls $[\boldsymbol{r},\, \tilde{\boldsymbol{c}}_i,\, \boldsymbol{J}]$ = VARPRO_loop($\boldsymbol{\tau}_i$, $\boldsymbol{z}_i$), with $\boldsymbol{\tau}_i$ initialized at random (the solver should minimize the $\boldsymbol{r}^2$ using Jacobian $\boldsymbol{J}$),
  2. computing matrix (family of functions) $\boldsymbol{\Phi}_{\boldsymbol{\tau_i}}$ from optimal $\boldsymbol{\tau}_i$ and fixed standard deviations,
  3. letting $\boldsymbol{c}_i = \tilde{\boldsymbol{c}}_i$ and making $\tilde{\boldsymbol{z}}_i = \boldsymbol{\Phi}_{\boldsymbol{\tau_i}}\boldsymbol{c}_i$ (ST-i vectors).
  Update approximation $\tilde{\boldsymbol{Y}}_{T\times J} \leftarrow \tilde{\boldsymbol{Y}}_{T\times J} + \tilde{\boldsymbol{z}}_i\boldsymbol{v}_i^\top$
end for
Return $\boldsymbol{v}_i$, $\boldsymbol{c}_i$, $\boldsymbol{\tau}_i$ ($i = 1\ \ldots\ k$) and $\tilde{\boldsymbol{Y}}$

Algorithm 2 : $[\boldsymbol{r},\, \tilde{\boldsymbol{c}}_i,\, \boldsymbol{J}]$ = VARPRO_loop($\boldsymbol{\tau}_i$, $\boldsymbol{z}_i$)

Compute matrix (family of functions) $\boldsymbol{\Phi}_{\boldsymbol{\tau_i}}$ from $\boldsymbol{\tau}_i$ and fixed standard deviations
Compute truncated $[\mathbf{U}, \boldsymbol{\Sigma}, \mathbf{V}]$ = SVD of $\boldsymbol{\Phi}_{\boldsymbol{\tau_i}}$
Make $\tilde{\boldsymbol{c}}_i = \boldsymbol{V}\tilde{\boldsymbol{\Sigma}}^{\boldsymbol{-1}}\boldsymbol{U}^\top\boldsymbol{z}_i$
Compute current approximation $\tilde{\boldsymbol{z}}_{\boldsymbol{i}} = \boldsymbol{\Phi}_{\boldsymbol{\tau_i}}\tilde{\boldsymbol{c}}_i$ and residual (or error) $\boldsymbol{r} = \boldsymbol{z}_i - \tilde{\boldsymbol{z}}_i$
for $j = 1$ to $\boldsymbol{N}_i$ Gaussians of $\boldsymbol{\Phi}_{\boldsymbol{\tau_i}}$ **do**
  Form matrix with partial derivatives $\boldsymbol{D}_j = \frac{\partial \boldsymbol{\Phi}(\tau_{i,j},t)}{\partial \tau_{i,j}}$ (see last paragraph of Section 3.2)
  Make $\boldsymbol{a}_j = \boldsymbol{D}_j\tilde{\boldsymbol{c}}_i - \boldsymbol{U}(\boldsymbol{U}^\top(\boldsymbol{D}_j\tilde{\boldsymbol{c}}_i))$ and $\boldsymbol{b}_j = \boldsymbol{U}(\boldsymbol{\Sigma}^{-1}(\boldsymbol{V}^\top(\boldsymbol{D}_j^\top\boldsymbol{r})))$
  Add $\boldsymbol{a}_j$ and $\boldsymbol{b}_j$ and form the j-th column of Jacobian $\boldsymbol{J}$ as in Equation 4
end for
Return $\boldsymbol{r}$, $\tilde{\boldsymbol{c}}_i$ and $\boldsymbol{J}$

matching pursuit. Please refer to the papers for further details. Figure 7 shows the resulting factors of a sample decomposition.

In our context, TVMS spatio-temporal synergy vectors $\boldsymbol{v}_i(t-\tau_i)$ would each correspond to a time-sequence of postures occurring only at instant (time-shift) τ_i, and scaled by constant magnitude z_i. In contrast, our spatial basis vectors $\boldsymbol{v}_i$ are individual postures that can occur at anytime throughout the trial, but with *time-varying scaling magnitudes* instead (recall from Equation 1).

4 Experiments and results

The goals of our experiments were (1) to validate the decomposition approach, by checking whether SB-ST parameters would allow for successful reconstruction of movements performed by different people; (2) to illustrate how the parameters of the model can be used to provide important insights related to both the action and actors involved. Although any kind of action could have been chosen, here we decided to look at *vertical jumps*, a non-trivial behavior that requires strength, coordination and balance.

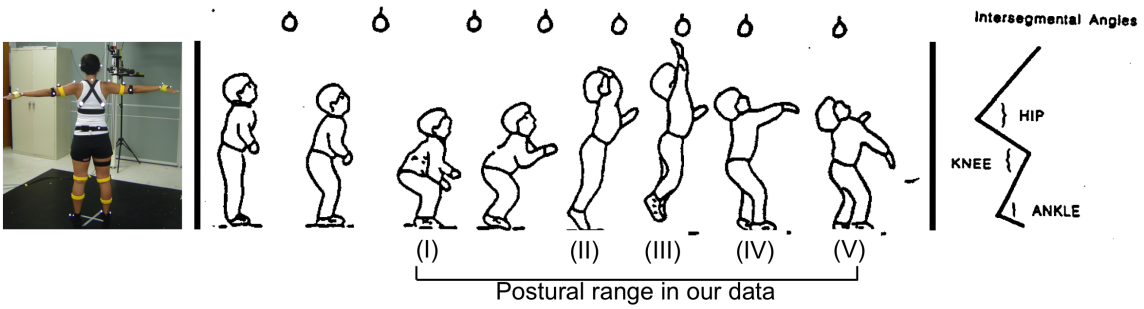

Figure 3: From left to right: subject setup in our lab, vertical jump experiment and intersegmental joint angles used in the experiments. Second and third pictures were adapted from [25].

4.1 Experiment setup and data description

Figure 3 shows the task our subjects performed: participants were instructed to jump vertically as high as possible trying to reach for a visual target. Written informed consent was obtained from all subjects/parents/legal representatives after a careful explanation of the testing procedures.

Each vertical jump resulted in a data matrix $\boldsymbol{Y}_{T\times J}$, of which a single row corresponds to a joint configuration or posture at a certain time. Collected jump trial matrices have $T \approx 80$ rows (about 0.8 seconds) and $J = 6$ columns encoding six joints: left and right hips, left and right knees plus left and right ankles. We only used the flexion/extension intersegmental joint angles. The data set described in this section has a total of 358 vertical jumps collected with optical motion capture technology. These jumps come from 4 different populations, totalizing 37 participants: 9 typically developing female children (98 jumps), 6 adult females (61 jumps), 10 typically developing male children (88 jumps), 5 adult males (52 jumps) and 7 children diagnosed with DCD [25] (59 jumps). Children were in the broad age range of 6 to 14 years old. Adults were all in their early 20's. Trials were manually segmented by an expert in the vertical jump movement, so that they span the same postural range: all poses captured within the initial and final peak knee flexions, as also in Figure 3. The raw data is attached in BVH (Biovision Hierarchy) format as part of the supplemental contents of this paper.

4.2 Vertical jump reconstruction

In our reconstruction experiments, all jump trials were decomposed into a spatial basis of 3 vectors SB-i ($i = \{1, 2, 3\}$) with varying (depending on the particular test being conducted) $\boldsymbol{N}_i$ pairs of basis functions/ST parameters $\boldsymbol{\tau}_i$ and $\boldsymbol{c}_i$. In particular, standard deviations were fixed as $\boldsymbol{\sigma}_i = \{1/(2 \cdot 1), 1/(2 \cdot 2) \ \ldots \ 1/(2 \cdot \boldsymbol{N}_i)\} \times T$. Note that we do not need to require all $\boldsymbol{N}_i$ to be the same, but we opted to do so in our experiments to simplify the analysis. We refer to these values as N from now on.

Prior to parameter estimation, each $\boldsymbol{z}_i$ was normalized into a unit vector. For the main loop of VARPRO, we used MATLAB® `lsqnonlin()` with $\boldsymbol{\tau}_i$ subject to being within $[0, 1]$, while no constraints were applied to $\boldsymbol{c}_i$. Figure 4 shows the statistics of the coefficients of determination R^2 for each individual joint series: the method was able to successfully

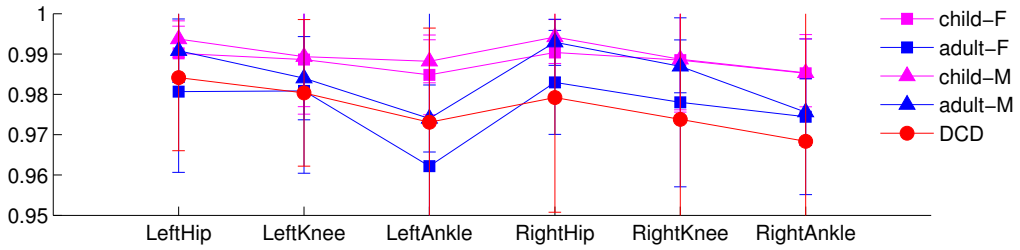

Figure 4: Mean ± std coefficients of determination R^2 per joint series approximation over all jumps in our dataset, grouped by population. Overall high average coefficients of determination ($\geq$0.95) reflect the successful reconstruction ability of the decomposition.

reconstruct all jumps in the dataset for all 5 populations considered, with average R^2 not less than 0.95. We also looked at how the reconstruction results are affected by the number of basis functions/pairs of ST parameters N. Figure 5 presents R^2 scores averaged by joints and populations versus $N = \{2 \ldots 20\}$(blue solid bar plots). For the given data, believe the best trade-off between dimensionality and quality of reconstruction was achieved when $4 \leq N \leq 8$.

4.2.1 Comparative reconstruction performance

Comparing with Troje-inspired decomposition

We have compared our VARPRO-based, Gaussian mixture decomposition of $\boldsymbol{z}_i$ with a Troje-inspired[2]/Fourier-based decomposition, by inspecting coefficients of determination versus number of basis functions utilized. For a certain N, the corresponding Fourier decomposition consisted in selecting the top-N responding harmonics (via FFT of $\boldsymbol{z}_i$) and using only these harmonics to reconstruct the original $\boldsymbol{z}_i$ (via IFFT). Note that, given a certain N, the number of parameters needed to reconstruct $\boldsymbol{z}_i$ is the same in both cases making these methods comparable: SB-ST fits a mixture of N Gaussians of fixed scale, therefore resulting in N pairs $\boldsymbol{\tau_i}, \boldsymbol{c}_i$ (ST-i parameters) while Troje's uses pairs of N selected Fourier harmonics along with respective responses.

The bottom plot of Fig. 5 clearly shows that our fitted mixture (blue) outperformed Troje-inspired approximation (red) of $\boldsymbol{z}_i$ from $4 \leq N \leq 7$, which could be considered the range with the best trade-off between the number of parameters used versus reconstruction error (note the change of slope in both methods when N moves from 3 to 4, as well as the dramatic decrease in R^2 cross-trial variances). Figure 6 also shows superior qualitative performance: for the same $N = 8$, our method fits the local details of $\boldsymbol{z}_i$ better than its competitor.

VARPRO regression

To assess VARPRO regression in particular, we replaced it with the default interior-point (IP) implementation of MATLAB® `fmincon()` and compared their reconstruction perfor-

[2]In his work, Troje [20] fits the temporal series of each of his eigenpostures with a single fundamental harmonic, which he finds sufficient to model the walking action. A natural extension to non-periodic actions as jumps is to select as many harmonics as necessary to obtain good approximations. This is what we call a *Troje-inspired* decomposition here.

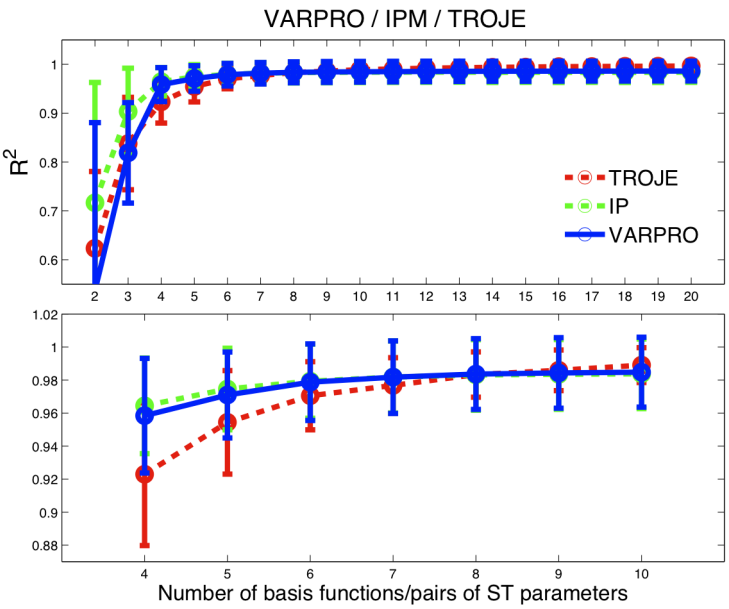

Figure 5: R^2 is plotted against the number N of basis functions/pairs of ST parameters, for different decomposition methods: Troje-inspired (dashed red), SB-ST decomposition with VARPRO regression (solid blue) and SB-ST with interior-point (IP) regression (dashed-green). Each point in the plot corresponds to the mean coefficient of determination $R^2 \pm$ std. Here, the coefficient of determination of a trial is itself the average of the coefficients of determination across all joints, since a jump trial is a multi-dimensional time series. A high mean R^2 (good reconstruction) together with a low std (good generalization across trials/populations) reveal $4 \leq N \leq 8$ to be an adequate trade-off, as can be noted from the bottom plot. See text for more details.

mances. As with VARPRO, $\boldsymbol{\tau}_i$ was constrained to the $[0, 1]$ interval. In other words, we kept the same SB-ST decomposition framework, but adopted a different regression method to estimate $\boldsymbol{\tau_i}, \boldsymbol{c}_i$ (ST-i parameters).

The top plot in Figure 5 shows that, in terms of R^2, VARPRO is outperformed by IP (red) when $N < 4$ and produced virtually as good average reconstruction errors as IP otherwise. However, both methods are only practically acceptable when $N \geq 4$, that is, when the cross-trial variances in R^2 drop to values that are relatively low. Within the range of $4 \leq N \leq 7$, the two performances were almost equivalent, as depicted by the bottom graph of Fig. 5. IP does slightly better when $N = 4$, which is the first time when both methods' R^2 exceed 0.95. Qualitatively, both VARPRO and IP behaved very well, as can be seen from the fits of the sample jump in Figure 6, but again IP did a little better. However, it is important to note that VARPRO is much faster than IP: for instance, when $N = 8$, VARPRO's per-trial average $\pm$ std execution time was 0.34 ± 0.16 s, while IP took 7.65 ± 0.61 s. The testing platform was an AMD Phenom™ 9750 quad-core 2.41 GHz processor, with 8GB of RAM, running 64-bit Microsoft® Windows XP Professional Version 2003/Service Pack 2.

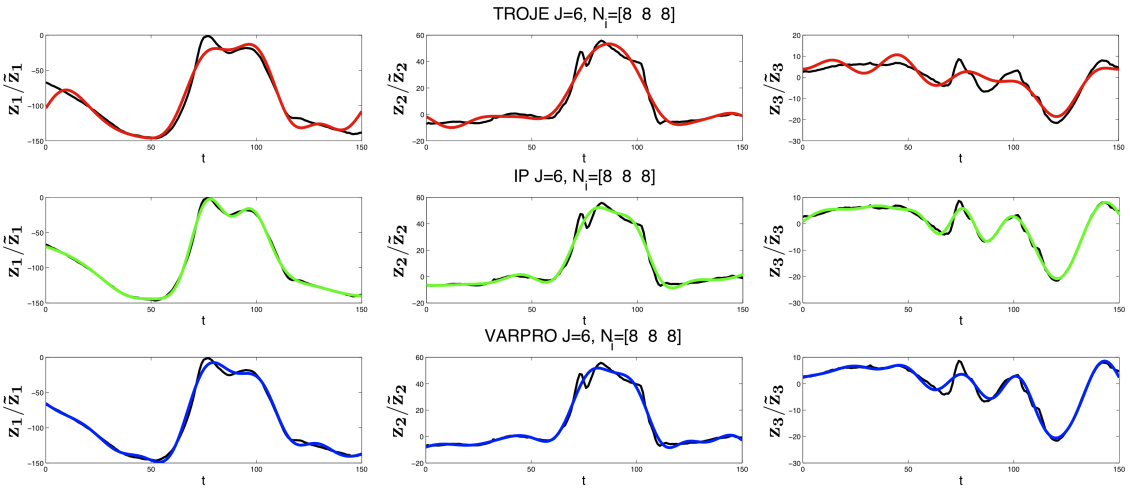

Figure 6: Fits of $\boldsymbol{z}_1$, $\boldsymbol{z}_2$ and $\boldsymbol{z}_3$ for one of the jump trials. Actual data are shown as black lines, and the results of Troje-inspired, VARPRO and IP approximations appear as red, blue and green lines, respectively.

Comparing with the TVMS model

As shown in Section 3.3, TVMS is a decomposition approach that shares similarities with the method described in this account, so we looked at how it would perform on our data: Figure 7 (a) shows the result of factoring one of our jumps into $N = 3$ synergies of length $Q = 60$ time units: most of the movement loads on a single synergy matrix $\boldsymbol{V}_2$, with hips-ankles and knees acting as agonists and antagonists, respectively. As will be seen in Section 4.3, this is consistent with the statistics of our spatial basis SB-1 coefficients. Moreover, Figure 7 (b) reveals that the second synergy seems to cover the central portion of the jump motion (note time-shifts z_2 and z_3 in Figure 7 (b)), from some point in between postures (I) or (II) and posture (IV) of Figure 3. The other two synergies took care of initial and final sections of the motion.

We found worth mention that it was not always the case that we were able to get satisfactory reconstruction of our data using TVMS. Figure 7 (c) shows a typical problem: a significant part of the signal (top) is not covered by the resulting synergies/parameters, resulting in a very low coefficient of determination i. e. high reconstruction error. As an effort to rule out the cause of the problem to be the poor selection of N and Q, we ran TVMS on the whole dataset using different combinations of these quantities, and the poor performance still persisted, as illustrated by the mean and standard deviations of Figure 7 (d). As a result, we discontinued the analysis based on that method.

We then conjecture that the reconstruction problems of TVMS on our data should result from not using more than a single trial to compute synergies and other parameters. As a matter of fact, D'Avella et al. designed TVMS under the assumption that there exists latent repertoires of synergies plus control and coordination parameters that span both multiple behaviors [10] and others that are behavior-dependent [11], and thus enforced their optimization to obtain factors that are faithful to these assumptions, in other words, synergies and parameters are supposed to be obtained from minimizing reconstruction errors across several trials (or episodes as worded by its authors) of one of more behaviors, meaning

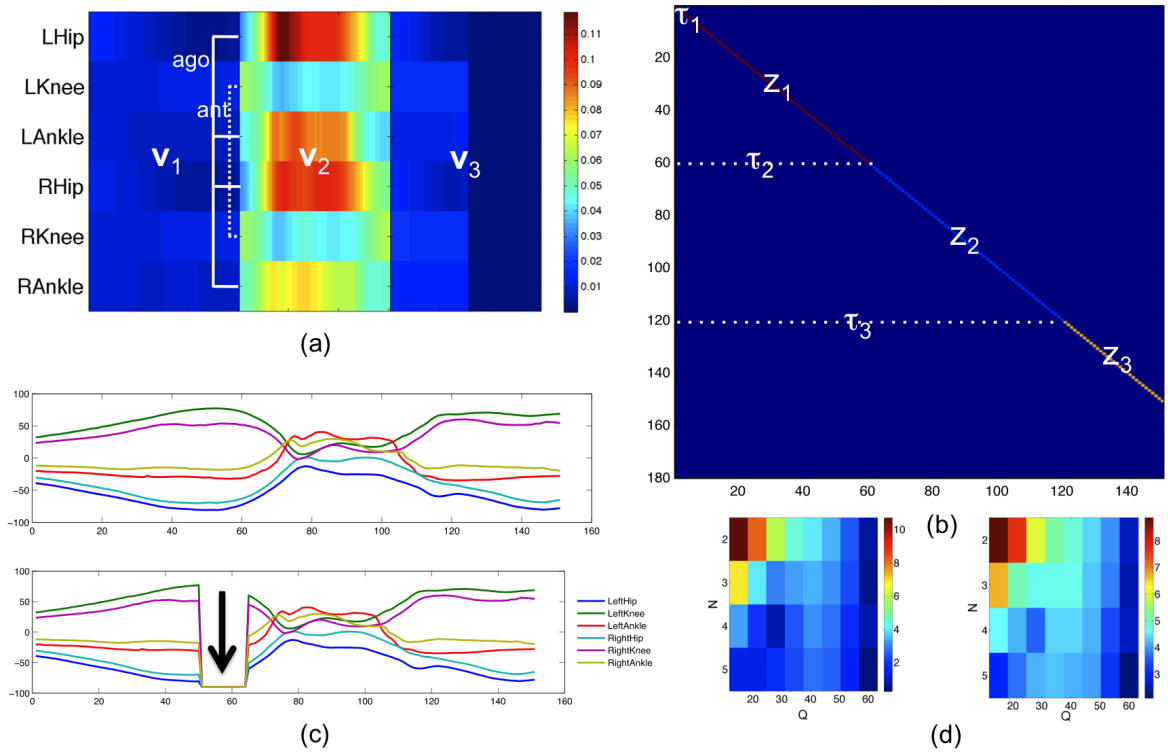

Figure 7: Single-trial based time-varying muscle synergies (TVMS) decomposition adapted to our jump data. It is a non-negative factorization that breaks up the trial into (a) a block synergy matrix $[\boldsymbol{V}_1|\boldsymbol{V}_2|\ ...\ |\boldsymbol{V}_N]$ and (b) a sparse matrix that shifts and scale each i-th synergy by τ_i and z_i, respectively (dark blue corresponds to zero). Matrices $\boldsymbol{V}_i$ are of size J (joints) by Q (time length), and their product recreates the input trial (see Equation 6 for the time-indexed notation). Note from (a) that most of the movement loads on $\boldsymbol{V}_2$, with agonist hips and ankles and antagonist knees (dashed line). Although the sample reconstruction based on the two matrices here displayed produced a near-perfect input signal such as the one at the top of figure (c), in many cases TVMS failed to recreate the jump trials and produced distorted results such as the one at the bottom of the same figure, where certain portions of the signal are zeroed (note the black arrow). This type of error was widely observed through the whole dataset for different values of N and Q attempted, and clearly reflected on the computed (d) means and standard deviations of absolute coefficients of determination. See text for further discussion.

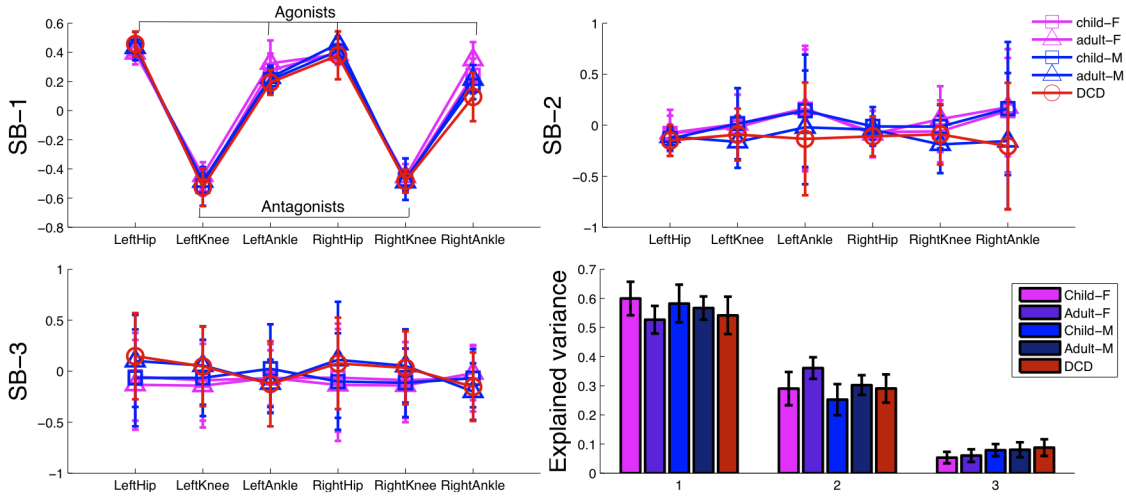

Figure 8: Mean ± std coefficients of vectors SB-1 (top-left), SB-2 (top-right) and SB-3 (bottom-left). At the bottom right, note that the mean ± std explained variances per vector are consistent across all populations in the dataset. The distribution of SB-1 coefficients reveals the major component of the vertical jump, apparently common to all populations, which can be noted by the overlapping means and low standard deviations. See text for more details.

several trials should be expected to be available to TVMS. Our study, on the other hand, *is not assuming the existence of common factors, but rather trying to observe if that should be the case by looking at the joint behavior of factors that come out of individual trials*. Hence, to allow for comparisons between TVMS and our SB-ST, we had to run the former on a single-trial basis. The corresponding MATLAB® source code of our TVMS implementation is supplied as part of supplemental materials.

4.3 Looking at jumps and jumpers based on the model parameters

From Figure 8, note that spatial basis SB-1 coefficient statistics suggest that over 50% of the (trial-averaged) explained variances in the vertical jump consists of 2 main groups of rotations: hips and ankles (top coefficient values in the range of 0.4 to 0.6) together with knee rotations (bottom coefficients within -0.6 to -0.4). Moreover, overlapping lines show these distributions seem to generalize across all populations examined. In fact, SB-1 works by clustering leg joints into the two existing agonist and antagonist motions, which is also clear from the picture. The same figure also reveals that SB-2 coefficients are almost zero-centered and have high variances, in special, left and right ankle coefficients. SB-3 coefficients are also mostly zero-centered, have even higher variances than SB-2 coefficients and less agreement across populations. With average SB-2 and SB-3 coefficients close to zero and no clear interpretation in the context of the jump action, the remaining discussion will focus only on spatio-temporal aspects of SB-1, that is, the statistics of ST-1's $\boldsymbol{\tau}_1 = \{\tau_{1,1} \ \ldots \ \tau_{1,N}\}$ and $\boldsymbol{c}_1 = \{c_{1,1} \ \ldots \ c_{1,N}\}$ estimated by VARPRO.

We looked at parameter distributions resulting from decompositions with $N = 4$ and $N = 5$ (see Gaussians G1 ... G5 on the left of Figure 9) to note that when the number of basis function changes, so does the distribution of parameters, thus the observed differences

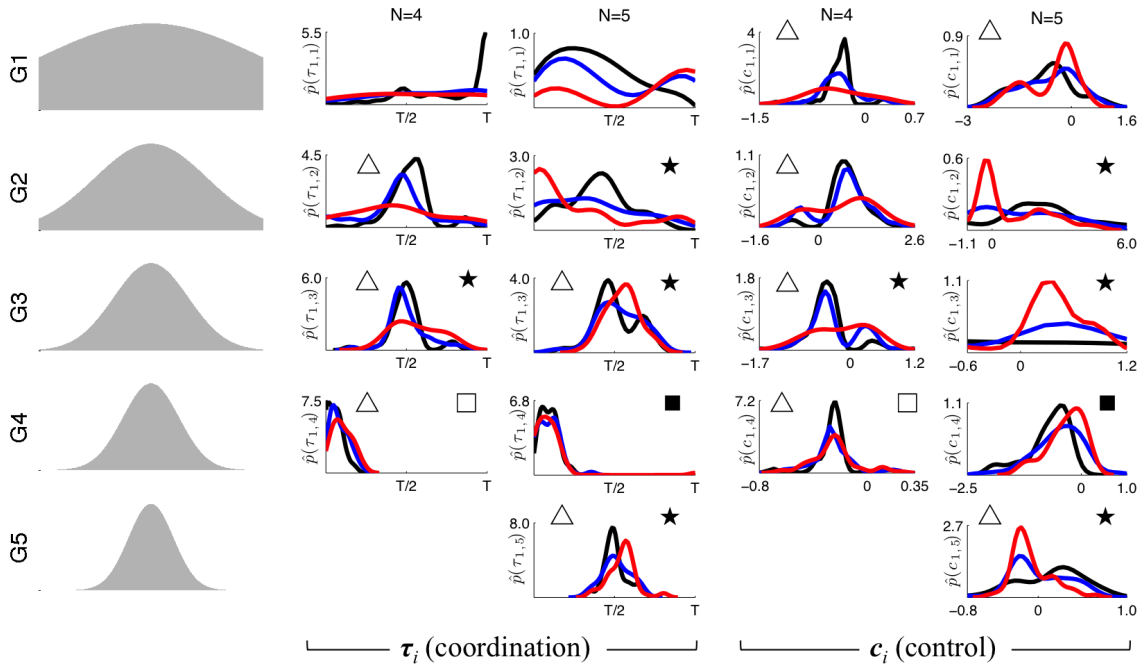

Figure 9: ST-1 parameter statistics for each specific developmental population in our dataset. To the left, the Gaussian basis functions G1 ... G5 of ST-1. Following, each row displays the distributions of $\boldsymbol{\tau}_i$ (coordination) and $\boldsymbol{c}_i$ (control) parameters for its corresponding Gaussian on the left. Parameters were computed for $N = 4$ and $N = 5$ were placed side-by-side to show that distributions may vary with the choice of N. Blue, red and black curves refer to data from typically developing children (TD), children with Development Coordination Disorder (DCD) and young adults (AD), respectively. Distributions were approximated with MATLAB® `ksdensity()` function, which was set to sample the data range at 50 points and to use a Gaussian kernel for smoothing. Bandwidths were automatically computed by that function, and varied across parameter distributions. To illustrate the extent to which these distributions agreed between populations, we labeled some of the plots with □ (AD, TD and DCD agree on both coordination and control), upper-left △ (anywhere AD and TD agree), ■ (DCD agrees with AD and TD on coordination, but not control) and ★ (DCD completely disagrees with AD and TD). See text for further discussion.

will also depend on the choice of N. The level to which this variation occur will be a consequence of the selection of the basis functions and/or the range of scale parameters utilized (i. e. standard deviations, when using Gaussian functions). Here, we are not assuming the existence of a *right* decomposition, but instead arguing that many decompositions are possible, and some may be useful for the movement analyst to uncover interesting differential features.

After smoothing all distributions with the MATLAB® `ksdensity()` kernel density estimator function, we looked at how jumpers of different developmental stages agreed in terms of the distributions of their ST-1 parameters. For simplicity, from now on we will call $\boldsymbol{\tau}_1$ and $\boldsymbol{c}_1$ *coordination* and *control* parameters respectively, because the former places each of the Gaussians along the timeline, so they match the local features of the spatial-temporal profile of posture SB-1, while the latter scales these Gaussians in accordance to the intensities of SB-1 activation. To be considered to agree, two distributions should have either similar shape or approximately the same peak abscissa, on the basis of visual inspection.

As seen from Figure 9, the frequent agreements between AD and TD (△) and simultaneous partial (■) and full (★) disagreements of these populations with DCD (mostly when $N = 5$) suggest these parameters are capable of discriminating these actors. Although qualitative interpretation of these parameters is not easy at first glance, i. e. it is hard to tell what is the meaning of fitting G4 to part of SB-1's spatio-temporal profile, interesting insights can still be derived. Let us consider the situations when coordination parameters could peak at approximately the same points of the timeline, have similar shape but differ in the distribution of control parameters: from Figure 9, when $N = 5$, the 3 populations seem to be recruiting G4 early in the timeline, since $\tau_{1,4}$ appears to be consistent for TD, AD and DCD. However, $c_{1,4}$ distributions do not agree among the 3 populations, which can be noted by the different peaks and shapes. We can then conjecture, based on the chosen decomposition, that there could be inter-population discrepancies related to the spatial configuration encoded by SB-1 taking place somewhat early in the course of the vertical jump, perhaps between the onset of the motion and time of takeoff (Figure 3). To discern what exactly these differences mean would require a more thorough analysis, and surpasses the scope of this work. Still, we believe that the framework outlined here would be able to assist in such investigation.

5 Conclusions

This paper describes the SB-ST decomposition method, which breaks the $\boldsymbol{Y}_{T\times J}$ action data matrix into spatial and spatio-temporal representations. Relevant spatial configurations of joints are identified by SVD and form a spatial basis (SB) for the action. Temporal correlation series of SB vectors, that is, the projection of the data matrix onto SB, are approximated by combinations of Gaussians through a non-linear least-squares regression method, VARPRO. These functions form the basis functions of the spatio-temporal representation (ST) of the action.

Concerning dimensionality, the action matrix is fully described by k (J-dimensional) SB vectors $\boldsymbol{v}_1, \boldsymbol{v}_2 \ \ldots \ \boldsymbol{v}_k$ plus another $2k$ ($\boldsymbol{N}_i$-dimensional) ST vectors $\boldsymbol{\tau}_1, \boldsymbol{\tau}_2 \ \ldots \ \boldsymbol{\tau}_k$ and $\boldsymbol{c}_1, \boldsymbol{c}_2 \ \ldots \ \boldsymbol{c}_k$, totalizing $kJ + 2k\sum_{i=1}^{k} \boldsymbol{N}_i$ parameters. Although the number of parameters increases linearly with the number of effectors J, it does not depend directly on the length of the trial. Dimensionality also depends on the number of basis functions/pairs of ST

parameters $\boldsymbol{N}_i$, but these numbers tend to be relatively small: recall that with $N = 4$, the method already presented good reconstruction with low inter-trial variance, as in Figure 5.

Based on SB/ST representations, we were able to reconstruct vertical jump trials with high accuracy. When $4 \leq N \leq 7$, it was shown to outperform a Troje-inspired decomposition. Within a similar range, VARPRO was also shown to perform almost as well as MATLAB® default interior-point (IP) implementation, but much faster. The method also appears to be more suitable to decompose single trials than the time-varying muscle synergies model (TVMS).

Moreover, the representation was able to generalize over 300 jumps coming from a broad range of jumpers: children from 6-14 years old with and without coordination disorders and young adults. Both genders were represented in the data. We also note that a single SB-1 vector seems to be sufficient to characterize the most important spatial aspect of the jump trials (based on lower-body, intersegmental flexion/extension angles), which helped to narrow down the analysis into 4 or 5 pairs of spatio-temporal parameters. From the statistics of the means and linear coefficients of ST-1, we also observed differences among populations of jumpers to become evident from our low dimensional representation of their jumps. Although the results presented here come from a specific case study, we believe the method serves as framework for a variety of movement analyses that can be carried on in the future.

6 Discussion: limitations and final remarks

For ST parameter estimation, we chose to work with fixed standard deviations and with a family of N Gaussians. In principle, this makes it simpler to compare populations by just looking at parameter statistics of Gaussians having corresponding standard deviations, as we did. Given the comparison results, this choice appeared to be adequate to the type of data we are looking at right now, but the pre-selection of the standard deviations for each Gaussian imposes a limitation in optimizing reconstruction. A potential improvement would be to add standard deviations to the parameter set, estimate and quantize it into a number of bins computed over all trials in the dataset. Population distributions could then be compared on the basis of these discovered bins rather than the current imposed standard deviations.

Furthermore, our family of functions $\boldsymbol{\Phi}_{\boldsymbol{\tau_i}}$ does not explicitly accommodate spatio-temporal profiles that are periodic or that have more than one mode of the same type, because each Gaussian will fit only a certain piece of the $\boldsymbol{z}_i$ profile. This could be attenuated if $\boldsymbol{\Phi}_{\boldsymbol{\tau_i}}$ was designed in a way that it included more than one Gaussian with the same standard deviation. Trying to cope with a similar problem, D'Avella et al. proposed to include repeated occurrences of synergies in their latest edition of TVMS, with overlaps constrained by imposed refractory periods. Nevertheless, to repeat either ST functions or synergies is a mere workaround, since it will only account for a fixed number of repetitions. Moreover, if we allowed multiple occurrences of the same function, there would be a computational impact on how we compare populations based on parameter distributions along the lines we proposed here. For example, the method could end up fitting three instances of G2 to the spatial-temporal profile $\boldsymbol{z}_i$ of a certain trial A and only one G2 to the corresponding profile of trial B, leaving the solution to matching B's G2 instances to A's undefined. For periodic actions, an approach like the Troje-inspired one (first case of Section 4.2.1) may be more suitable. Allowing for a short digression, it could be the case that the brain deals with

generating and interpreting cyclic and non-cyclic actions in different ways, and an artificial system that wanted mimic its behavior might have to figure out its type before parsing the signals.

Finally, this is work is part of a project that attempts to bridge cognitive science, psychology of movement, kinesiology and computer science. A long-term goal is to develop a non-invasive tool that could aid in the diagnosis of behavioral disorders based on the automatic detection of inter-population differences, by using SB-ST parameters as classification features. This could be extremely useful when discrepancies are not immediately evident to the eyes of the health care professional, or when there is disagreement in the diagnostic. The tool could even be extended to integrate parameters from the decomposition of a larger pool of behaviors, and statistical models for motor disorders such as DCD could be defined in terms of the co-distributions of control and coordination parameters.

Acknowledgements

This work was supported by NIH grant 5R21DA24323-3. We want to acknowledge Dr. Melissa Pangelinan for organizing schedules, recruiting subjects, running them through the test protocol, and taking pictures of the marker setup (e. g. Figure 3). Dr. Brad King, for manually labeling key snapshots of each jump trial and sharing his expertise in kinematics and biomechanics of vertical jumps. Prof. Jane Clark, for the relevant references and technical advices. Dr. Anne Jorstad and João Soares, for carefully reviewing previous versions of the manuscript.

References

[1] Michael T. Turvey, Coordination. *American Psychologist* 45(8):938–953, 1990.

[2] Mark L. Latash, John P. Scholz and Gregor Schöner, Toward a New Theory of Motor Synergies. *Motor Control* 11:276–308, 2007.

[3] Mark L. Latash, Mindy F. Levin, John P. Scholz and Gregor Schöner, Motor Control Theories and Their Applications. *Medicina (Kaunas)* 46(6):382–392, 2010.

[4] Tamar Flash and Binyamin Hochner, Motor primitives in vertebrates and invertebrates. *Current opinion in neurobiology* 15:660–666, 2005.

[5] John P. Scholz and Gregor Schöner, The uncontrolled manifold concept: identifying control variables for a functional task. *Exp Brain Res* 126:289–306, 1999.

[6] Ferdinando A. Mussa-Ivaldi, Simon F. Giszter and Emilio Bizzi, Linear combination of primitives in vertebrate motor control. *Proceedings of the Natural Academy of Sciences* 91: 7534–7538, 1994.

[7] Ferdinando A. Mussa-Ivaldi, Nonlinear force fields: a distributed system of control primitives for representing and learning movements. *Proceedings of Computational Intelligence in Robotics and Automation* 84–90, 1997.

[8] Ferdinando A. Mussa-Ivaldi and Emilio Bizzi, Motor learning through the combination of primitives. *Philosophical Transactions of the Royal Society of London B* 355:1755–1769, 2000.

[9] Andrea d'Avella and Matthew C. Tresch, Modularity in the motor system: decomposition of muscle patterns as combination of time-varying synergies. *Proceedings of Advances in Neural Information Processing Systems* 14:141–148, 2001.

[10] Andrea d'Avella, Philippe Saltiel and Emilio Bizzi, Combinations of muscle synergies in the construction of a natural motor behavior. *Nature Neuroscience* 3(6):300–308, 2003.

[11] Andrea d'Avella, Philippe Saltiel and Emilio Bizzi, Shared and specific muscle synergies in natural motor behaviors. *Proceedings of the National Academy of Science* 8(102):3076–3081, 2005.

[12] Yuri P. Ivanenko, Germana Capellini, Nadia Dominici, Richard E. Poppele and Francesco Lacquaniti, Coordination of Locomotion with Voluntary Movements. *The Journal of Neuroscience*25(31):7238–7253, 2005.

[13] Yuri P. Ivanenko, Richard E. Poppele and Francesco Lacquaniti, Motor Control Programs and Walking. *The Neuroscientist*12(4):339–348, 2006.

[14] Marco Santello, Martha Flanders and John F. Soechting, Postural Hand Synergies for Tool Use. *The Journal of Neuroscience* 18(23):10105–10115, 1998.

[15] Martin A. Giese, Tomaso Poggio, Morphable models for the analysis and synthesis of complex motion patterns. *International Journal of Computer Vision* 38(1): 59–73, 2000.

[16] Matthew C. Tresch, Vincent C. Cheung and Andrea d'Avella, Matrix factorization algorithms for the identification of muscle synergies: evaluation on simulated and experimental data sets. *Journal of Neurophysiology* 95(4):2199–2112, 2006.

[17] Gunnar Johansson, Visual perception of biological motion and a model for its analysis. *Perception and Psychophysics* 14:201–211, 1973.

[18] Gunnar Johansson, Spatio-Temporal Differentiation and Integration in Visual Motion Perception. An Experimental and Theoretical Analysis of Calculus-Like Functions in Visual Data Processing. *Psychological Research* 38(4):379–393, 1976.

[19] Winand H. Dittrich, Action categories and the perception of biological motion. *Perception* 22(1):15–22, 1993.

[20] Nikolaus F. Troje, Decomposing biological motion: A framework for analysis and synthesis of human gait patterns. *Journal of Vision* 2:371–387, 2002.

[21] Gene Golub and Victor Pereyra, The differentiation of pseudo-inverses and non-linear least squares problems whose variables separate. *SIAM Journal on Numerical Analysis* 10:413–432, 1973.

[22] Gene Golub and Victor Pereyra, Separable Nonlinear Least Squares: the Variable Projection Method and its Applications. *Inverse Problems* 19(2):1–26, 2002.

[23] Dianne P. O'Leary and Bert W. Rust, Variable Projection for Nonlinear Least Squares Problems. *Unpublished manuscript*, 2010, Available at `http://www.cs.umd.edu/users/oleary/software/varpro.pdf`.

[24] Daniel D. Lee, H. Sebastian Seung, Algorithms for Non-negative Matrix Factorization. *Proceedings of Advances in Neural Information Processing Systems* 13:556–562, 2000.

[25] Jody L. Jensen, Sally J. Phillips and Jane E. Clark, For Young Jumpers, Differences Are in the Movement's Control, Not Its Coordination. *Research Quarterly for Exercise and Sport* 65:258–268, 1994.