Towards Sparse Coding of Natural Movements for Neuroprosthetics and Brain–Machine Interfaces

Andreas A. C. Thomik[1], Sonja Fenske[1†] & A. Aldo Faisal[1,2,3]

Abstract— The correlation structure of natural hand & finger movements suggests that their motion is controlled in a lower–dimensional space than would be possible given their mechanical nature. Yet, it is unclear whether this low dimensional embedding is relevant to how the brain represents motor actions and how we can decode it for Brain-Machine Interface applications. We collected large data set of natural hand movement kinematics and analysed it using a novel sparse coding and dictionary learning approach – Sparse Movement Decomposition (SMD), which captures the embedding of the data in terms of spatial and temporal structure. We show that our sparse codes over natural movement statistics give a more parsimonious representation than the simple correlation structure. This suggest that, like V1 neuron receptive fields can be predicted from sparse code over natural image statistics, motor control may be encoded in such a manner. We further show how our sparse coding can help understand the temporal structure of behaviour, and thus our technique may be used for behavioural fingerprinting in diagnostics and for more naturalistic neuroprosthetic control.

I. INTRODUCTION

The search for a low–dimensional, easily interpretable description of motor outputs has attracted a considerable amount of scientific interest. As a result, numerous models have been suggested which aim to explain observed muscle activation patterns. The output of these models, termed *muscle synergies*, describe synchronous co–activation of muscles, time shifted co–activation or more complex time–varying activation patterns of muscles [1]. While very successful at capturing the observed muscle activations, the validity of these models is still widely debated [2].

Muscle synergies have been studied for arm [3], [4] and leg [5] movements, however, hand movements have received only limited attention, most likely because of the large number of muscles involved. Instead, the kinematics of hand movements during reach–to–grasp [6], tool use [7] and natural movement [8] have been analysed using principal component analysis (PCA), revealing a low–dimensional structure. While the true dimensionality of the data is unclear, the most complete set as recorded by Ingram *et al.* [8] suggests that 6 principal components (PCs) are sufficient to capture upwards of 80% of the variance in the data. The meaning of these PCs is, however, unclear. While the first

two PCs seem to be conserved across subjects [8], this is not the case for higher–order PCs.

A characteristic of these methods is that they seek a global dimensionality reduction across the entire data at their disposal. Yet, local dimensionality reduction, such as obtained by sparse encoding, has been shown to give important insights into the working of the sensory system. In particular, the analysis of natural scenes using a sparse code adapted to the statistics of real–world images has proven to provide encoding not unlike the one found in the visual cortex [9].

Inspired from the work on sparse coding in the sensory stream, we attempt to learn an over–complete representation of natural movements from raw data. We extract a dictionary of *eigenmotions* which correspond to elementary movements. Our method conserves the temporal structure of movement by only connecting data in a small temporal neighbourhood. Compared to PCA, we demonstrate increased dimensionality reduction over a large set of natural hand movements, as well as an application to behavioural fingerprinting.

II. METHODS

A. Data Acquisition

We acquired bimanual hand movement data from 10 healthy, right handed subjects going about their daily life. Digit movement was recorded using an 18-sensor Cyber-Glove [CyberGlove Systems LLC, San Jose, CA] on the left and a 22-sensor glove on the right hand. The gloves measure joint abduction using stretch sensors embedded in the material with a spatial resolution of <1 degree and a sampling rate of 90 Hz. We recorded 1-2 hours of data per subject, yielding more than 5 million data points per hand in total. All subjects gave written consent and the experimental procedure was approved by a local ethics committee.

B. Sparse Movement Decomposition

The algorithm's function is outlined in Figure 1. It is a (quasi–) deterministic, iterative procedure which seeks an optimal dictionary of *eigenmotions* to represent the input data by computing a short–time PCA on strategically chosen data windows. Details of the Sparse Movement Decomposition (SMD) methodology are described below. We chose to develop our own method instead of using a standard toolbox such as SPAMS [10] as the latter struggles with smoothness constraints on the latent representation.

1) Data Preprocessing: Prior to analysis we smooth the data with a 2^{nd} order Savitzky-Golay filter on a 5 sample window to remove the digitisation noise induced by the
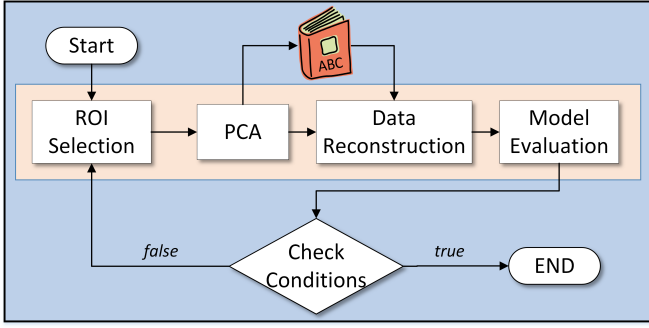
Fig. 1: Sparse Movement Decomposition algorithm flowchart. After pre-processing of the data, a small chunk of data is selected and eigenmotions (principal components) are extracted from it and added to the dictionary. The data is then reconstructed using the learnt dictionary and the accuracy of the model is evaluated, before reiterating the procedure. See text for details.

analogue–to–digital (A/D) converter. To obtain an approximately stationary signal, the position data is then differentiated to obtain the velocity time–series. Finally, we identify data points in the time–series in which there is activity as all those points for which the instantaneous absolute velocity exceeds one standard deviation of the overall velocity data. To avoid zero-crossings to be counted as immobile or noise as movement, a morphological closing operation followed by morphological opening operation is applied to the binary classification (moving = 1; immobile = 0) using a linear structuring element. This effectively removes very short movement epochs and closes the gap at zero crossings.

2) Region-of-Interest Selection: The key step to identifying new *eigenmotions* is to select regions of interest (ROI) in the time-series which can provide us with novel information about the latent representation. We consider two different cases: (1) the first iteration and (2) all subsequent iterations. Selection criteria for the ROI vary as follows:

- When the algorithm is confronted with a new time-series for the first time, the ROI is randomly selected from the data.
- During the reconstruction stage of the algorithm, the new ROI is selected as the data around the time-point with the lowest reconstruction score S (see II-C).

When selecting the size of the ROI, it is essential to take into account two potentially conflicting needs: (1) time-scale of the underlying dynamics and (2) numerical stability of the estimated statistics. It should be noted that in all cases, the ROI selected is then discarded to ensure that it does not get selected on multiple occasions. We chose a ROI of 61 samples (ca. $680ms$).

3) Eigenmotion Extraction: The eigenmotions we seek can be considered as local principal components (which are equivalent to atoms of our dictionary). They can thus be estimated from the data within the ROI selected using classic PCA. The number of significant PCs extracted for each ROI is automatically determined using the approach described by Cheung and colleagues [11] with a threshold set at 10^{-4}. Because of the limited amount of data present, this estimate may, however, be inaccurate. This inaccuracy is partially mitigated by bootstrapping the sample covariance matrix used as input to the PCA.

4) Termination: At the end of each iteration, the following conditions are checked:

- Is the reconstruction criteria met?
- Has the maximal number of iterations been reached?
- Is there data left to analyse?
- Has there been any significant improvement in reconstruction quality over the last five iterations?

If any of those conditions fail, the algorithm terminates and returns the learnt dictionary. For our dataset, we find that on average 1.25% of the data was analysed for new eigenmotions before the method terminates.

C. Reconstruction Score

To identify regions in which the current model is not well suited for the data, we compute a local goodness of fit metric:

$$S(t) = 1 - \frac{\text{VAR}[\boldsymbol{X}_W(t) - \hat{\boldsymbol{X}}_W(t)]}{\text{VAR}[\boldsymbol{X}_W(t)]} \tag{1}$$

where $\boldsymbol{X}_W(t)$ and $\hat{\boldsymbol{X}}_W(t)$ are the data and its estimate taken from a window of size W around the time-step t and VAR is the variance operator. This is advantageous over using the RMSE as it is independent of variations in data amplitude and has an immediate interpretation in terms of variance explained.

D. Data Reconstruction with an Over–Complete Dictionary

Given a dataset $\mathbf{X}$ and an over-complete and/or non-orthogonal dictionary $\mathbf{B}$, determining the optimal selection of atoms to represent a given data point is a non-trivial task. The possibly simplest solution to the problem, termed Matching Pursuit (MP) [12], projects each $\mathbf{x}_i$ onto $\mathbf{B}$ and selects greedily the best matching atom. This process is then repeated with the residuals until a pre-determined number of atoms has been selected or the residuals are smaller than a given threshold. Orthogonal Matching Pursuit (OMP) [13] takes the same approach as MP but restricts the selection of additional atoms to those which have not yet been chosen. These methods are, however, very sensitive to noise in the data and yield unstable representations which may change significantly depending on the signal-to-noise ratio. Moving away from greedy methods, Basis Pursuit (BP) [14] and the Least Absolute Shrinkage and Selection Operator (LASSO) [15] both consider the task of finding the best matching bases as a classical optimisation problem to be solved by any valid optimisation algorithm, where the cost function to be minimised is:

$$C = \|\mathbf{X} - \mathbf{BZ}\|_2^2 + \lambda \sum_i \|\mathbf{z}_i\|_1 \tag{2}$$

where $\mathbf{Z}$ is the latent representation, $\mathbf{z}_i$ the latent representation at time point i and λ a regularisation parameter.

The SMD algorithm internally reconstructs the data using the LASSO. However, to reduce the dimensionality of the data to a specific number we use OMP. Both algorithms are implemented in the SPAMS Toolbox by Julien Mairal [10].
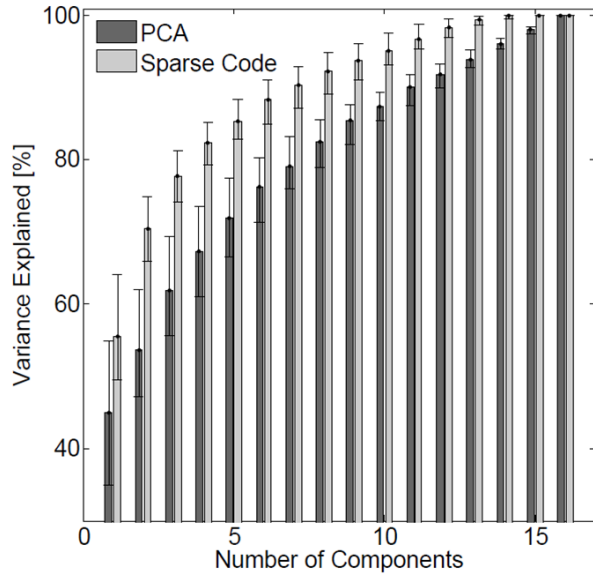

Fig. 2: Comparison of variance explained for sparse coding (light grey) and PCA (dark grey). For the same number of components, sparse codes explain a significantly superior amount of variance.

III. RESULTS

We recorded more than 15 hours of hand movement data from 10 subjects freely behaving in a natural environment, preparing food, working and interacting with other people. The data is analysed using both PCA and a custom algorithm for dictionary learning and sparse coding.

A. Low-Dimensional Encoding of Natural Movement

We wanted to know whether natural human movement has sparse, latent representation. To this end, we compared the variance explained by a sparse code with N non-zero activations to the variance explained by the first N PCs of classic PCA. If there is no sparse structure underlying the observed data, we would expect that the components obtained from PCA and SMD each explain similar amount of variance. Conversely, if there is a sparse structure underlying the data, then we would expect each component of the sparse code to capture a larger amount of variance than its PCA counterpart. The results of our comparison are shown in Figure 2. We find that representing the data using a sparse code invariably explains a significantly larger proportion of the variance than PCA, thus suggesting that there is indeed a sparse latent representation underlying natural hand movement.

B. Behavioural Fingerprinting

Unlike PCA, where the importance of components is ranked by the amount of variance they explain, atoms in a sparse code tend to explain similar amounts of variance in the data overall, but are only relevant at specific moments in time. This has the benefit of giving insights not only into the spatial, but also temporal dynamics of task execution. In particular, we take the approach of finding the extremely sparse representation by allowing only one eigenmotion to be active at any given moment in time. While this is not sufficient for accurate data reconstruction, it allows us to identify which eigenmotion is most relevant for any given data point. Thus, complex behaviour can be reduced to a one–dimensional, discrete time–series where each number indicates the eigenmotion active at that moment in time. We exemplify this procedure in Figure 3. The top panel shows a short extract of the data representing a subject turning three pages in a magazine. Using Eigenmotion Decomposition on
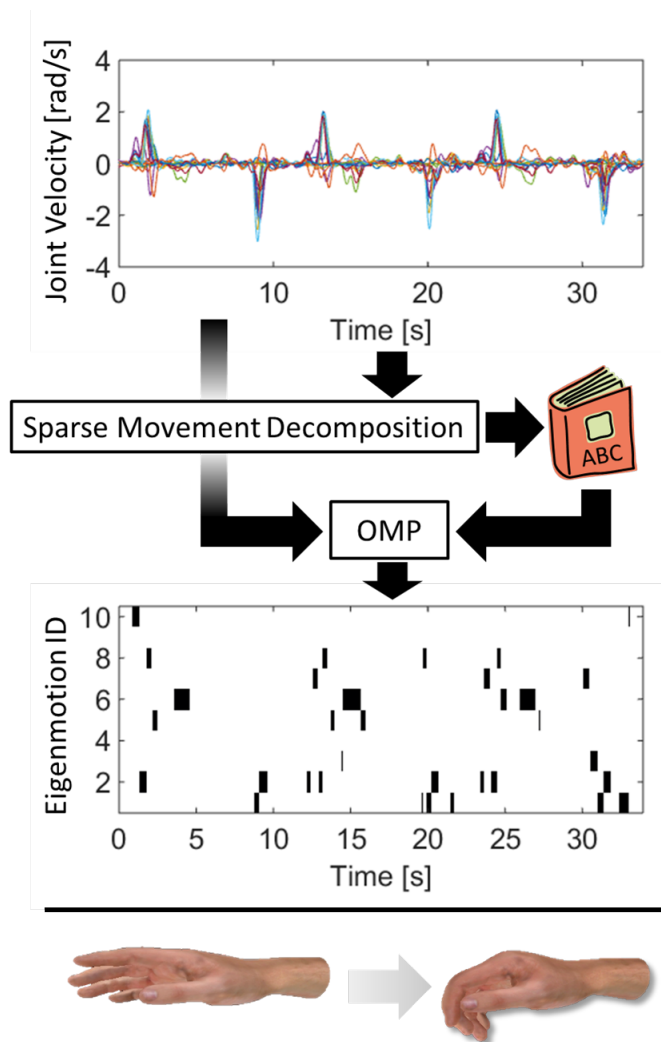

Fig. 3: *Top* Using the dictionary of eigenmotions learnt with the SMD algorithm and the raw data, we find a temporal encoding of the behaviour by identifying at each moment in time the most relevant eigenmotion. This yields a "behavioural barcode". *Bottom* Representation of a single *eigenmotion* producing a grasp.

a larger data set, we generate a dictionary of eigenmotions. Together with the raw data, this dictionary is used as input to the OMP algorithm to determine the most important at each moment in time, thus effectively developing a bar-code of behaviour. The specific example shown in Figure 3 thus nicely exposes the repetitive structure of the task despite the variability of the input data.

IV. DISCUSSION

The brain needs to translate inputs from the sensory system into a code with which it can operate and transform into actions. Many different encodings are possible, each optimal for a specific situation. While it is unclear what criteria of optimality the brain uses, it has been suggested [16] that sparse coding may be a good candidate. We asked whether the motor output may also result from such a sparse code. The results reported in Figure 2 suggest that this may indeed be the case. The brain may thus selects from a dictionary of primitive movements which are scaled and combined to create the observed motor output [17]. These findings are consistent with previous studies which found neurons coding specifically for different grasp types in the ventral premotor area F5 of the macaque cortex [18]. We postulate that similar to language grammar there may be something like a "movement grammar", indicating which sequences of eigenmotions are possible and meaningful. The addition of temporal structure has a direct application to prosthetic control: We previously showed that learning the likelihood of future movement sequences in natural behaviour allows to better decode movement intention by combining the estimate from the measurements on the body (e.g. EMG or MMG [19], [20]) with the known temporal structure of the behaviour [21], [22].

REFERENCES

[1] A. d'Avella and E. Bizzi, "Shared and specific muscle synergies in natural motor behaviors," *PNAS*, vol. 102, no. 8, pp. 3076–3081, 2005.

[2] M. C. Tresch and A. Jarc, "The case for and against muscle synergies," *Curr Opin Neurobiol*, vol. 19, no. 6, pp. 601–607, 2009.

[3] A. d'Avella, P. Saltiel, and E. Bizzi, "Combinations of muscle synergies in the construction of a natural motor behavior," *Nat Neurosci*, vol. 6, no. 3, pp. 300–308, 2003.

[4] A. d'Avella *et al.*, "Control of fast-reaching movements by muscle synergy combinations," *J Neurosci*, vol. 26, no. 30, pp. 7791–7810, 2006.

[5] G. Torres-Oviedo and L. H. Ting, "Muscle synergies characterizing human postural responses," *J Neurophysiol*, vol. 98, no. 4, pp. 2144–2156, 2007.

[6] C. Mason, J. Gomez, and T. Ebner, "Hand synergies during reach-to-grasp," *J Neurophysiol*, vol. 86, no. 6, pp. 2896–2910, 2001.

[7] M. Santello, M. Flanders, and J. Soechting, "Postural hand synergies for tool use," *J Neurosci*, vol. 18, no. 23, pp. 10 105–10 115, 1998.

[8] J. N. Ingram *et al.*, "The statistics of natural hand movements," *J Neurophysiol*, vol. 188, no. 2, pp. 223–236, 2008.

[9] B. A. Olshausen *et al.*, "Emergence of simple-cell receptive field properties by learning a sparse code for natural images," *Nature*, vol. 381, no. 6583, pp. 607–609, 1996.

[10] J. Mairal *et al.*, "Online learning for matrix factorization and sparse coding," *JMLR*, vol. 11, pp. 19–60, 2010.

[11] V. Cheung *et al.*, "Central and sensory contributions to the activation and organization of muscle synergies during natural motor behaviors," *J Neurosci*, vol. 25, no. 27, pp. 6419–6434, 2005.

[12] S. G. Mallat and Z. Zhang, "Matching pursuits with time-frequency dictionaries," *IEEE T Signal Proces*, vol. 41, no. 12, pp. 3397–3415, 1993.

[13] Y. C. Pati, R. Rezaiifar, and P. Krishnaprasad, "Orthogonal matching pursuit: recursive function approximation with applications to wavelet decomposition," in *Signals, Systems and Computers, 1993. 1993 Conference Record of The Twenty-Seventh Asilomar Conference on*, IEEE, 1993, pp. 40–44.

[14] S. S. Chen, D. L. Donoho, and M. A. Saunders, "Atomic decomposition by basis pursuit," *SIAM J Sci Comput*, vol. 20, no. 1, pp. 33–61, 1998.

[15] R. Tibshirani, "Regression shrinkage and selection via the lasso," *J Roy Stat Soc B Met*, pp. 267–288, 1996.

[16] B. A. Olshausen and D. J. Field, "Sparse coding with an overcomplete basis set: a strategy employed by v1?" *Vision Res*, vol. 37, no. 23, pp. 3311–3325, 1997.

[17] G. Rizzolatti and M. Gentilucci, "Motor and visual-motor functions of the premotor cortex," *Neurobiology of Neocortex*, pp. 269–284, 1988.

[18] V. Raos *et al.*, "Functional properties of grasping-related neurons in the ventral premotor area f5 of the macaque monkey," *J Neurophysiol*, vol. 95, no. 2, pp. 709–729, 2006.

[19] C. Gavriel and A. A. Faisal, "A comparison of day-long recording stability and muscle force prediction between bsn-based mechanomyography and electromyography," in *Wearable and Implantable Body Sensor Networks (BSN), 2014 11th Intl Conf*, IEEE, 2014, pp. 69–74.

[20] S. Fara *et al.*, "Prediction of arm end-point force using multi-channel mmg," in *Wearable and Implantable Body Sensor Networks (BSN), 2014 11th Intl Conf*, IEEE, 2014, pp. 27–32.

[21] A. A. Thomik, D. Haber, and A. A. Faisal, "Real-time movement prediction for improved control of neuroprosthetic devices," in *Neural Engineering (NER), 2013 6th Intl IEEE Conf*, IEEE, 2013, pp. 625–628.

[22] D. Haber, A. A. Thomik, and A. A. Faisal, "Unsupervised time series segmentation for high-dimensional body sensor network data streams," in *Wearable and Implantable Body Sensor Networks (BSN), 2014 11th Intl Conf*, IEEE, 2014, pp. 121–126.