Hierarchical Control of Reaching Movements Via Compositional Gain Modulation

Anonymous Author(s)
Affiliation
Address
email

Abstract

Reusing existing modules in novel settings via compositional generalization is the hallmark of intelligent behavior. While much research is dedicated to studying how to enable AI systems to learn and reuse modules effectively for better performance and increased computational efficiency, there is still a lack of consensus on which modules the brain leverages and how to identify them. To shed some light on this matter, here we investigate the modularity principles the brain uses to control the body efficiently. After briefly revisiting established models of domain-specific spatial and temporal motor modularity, we introduce a new, unifying computational model of compositional generalization in the motor system based on the Canonical Polyadic Decomposition (CPD) model. We show that the model — which leverages gain modulation — can simultaneously capture modularity in the spatial, temporal, and action domains with a lower number of parameters than established models. Furthermore, we show that the geometrical organization of the action modules the model isolates is not random but describes a smooth manifold that allows the zero-shot learning of muscle patterns for untrained movements. Taken together, our results suggest that the decomposition proposed here represents an effective compositional strategy the brain could leverage to control complex movements while saving computational resources.

1 Introduction

A large body of research is dedicated to studying the properties of compositional generalization of AI systems (e.g., Schug et al. [2023], Lippl and Stachenfeld [2024], Hupkes et al. [2020], Lake and Baroni [2018]), with the ultimate goal of encouraging intelligent behavior through the flexible reuse of previously learned modules (e.g., Duan et al. [2023], Liu et al. [2023], Berg et al. [2023]). But how does the brain accomplish compositional generalization? Can we identify the modules the brain uses to save computational resources and boost generalization? To take a step in this direction, in this work, we investigate generalization in the motor system.

Converging behavioral [Tresch et al., 1999, Ivanenko et al., 2004] and neurophysiological [Takei et al., 2017, Levine et al., 2014] evidence suggests that the motor system employs spatial [Tresch et al., 1999, Levine et al., 2014, Takei et al., 2017] and temporal [Ivanenko et al., 2004, Hart and Giszter, 2010, Takei et al., 2017] modules to simplify the control of movement. These modules are fixed across movements and allow the reuse of the spatiotemporal muscle activity patterns that are successful at moving the body purposefully. The adoption of this strategy relieves the motor system from the burden of computing such spatiotemporal patterns de novo for each movement and reduces the problem of computing appropriate muscle activation commands for new movements to the determination of scaling weights for such modules. If the motor system produces reaching movements by flexibly combining such fixed building blocks, one would expect activity patterns in

Submitted to 38th Conference on Neural Information Processing Systems (NeurIPS 2024). Do not distribute.

higher motor centers that largely invariant across reaching trajectories. Interestingly, this is consistent with what has been observed in primary motor cortex [Churchland et al., 2012], where the population dynamics tend to exhibit rotational structure that is invariant across conditions.

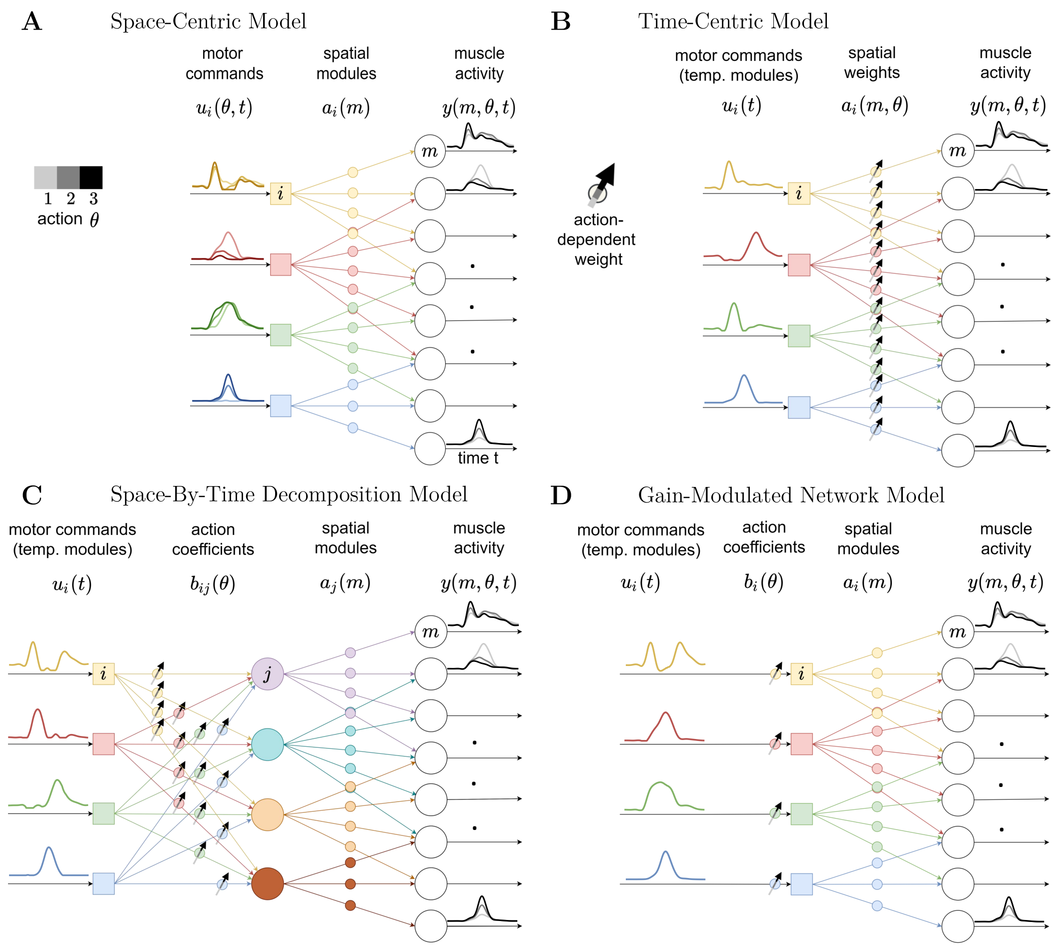

Figure 1: **Schematics of decomposition models.** Both the classical decomposition models (A-C) and the one proposed in this work (D) can be interpreted as simple linear feedforward neural networks. The space-centric model (A) is only able to capture the spatial invariances experimentally observed in the spatio-temporal muscle activity patterns; this is achieved at the cost of increasing the model complexity in the temporal domain where the model assumes the existence of action-specific motor commands. The time-centric model (B) is only able to capture the temporal invariances; this is achieved at the cost of increasing the model complexity in the spatial domain, where the model assumes the existence of action-and-muscle-specific weights. The space-by-time decomposition model (C) is able to capture both spatial and temporal invariances; this is achieved at the cost of increasing the complexity at the network-level, as the model introduces an additional hidden layer that can completely change the routing of the motor commands to the downstream layers. The gain-modulated network model (D) can also capture both spatial and temporal invariances without introducing network-level complexity; the model assumes that new actions can be generated by only modulating the gains of the input neurons.

Despite the large body of evidence supporting the existence of both spatial and temporal motor modules, the classical methods used to extract such modules from muscle activity signals — Non-negative Matrix Factorization and Principal Component Analysis — are intrinsically matrix decomposition methods and can only identify either spatial or temporal motor modules, but not both [Chiovetto et al., 2022]. This leads to potentially overparameterized models that, rather than providing a plausible account of the mechanism the brain uses to simplify the control of movement, capture the regularities

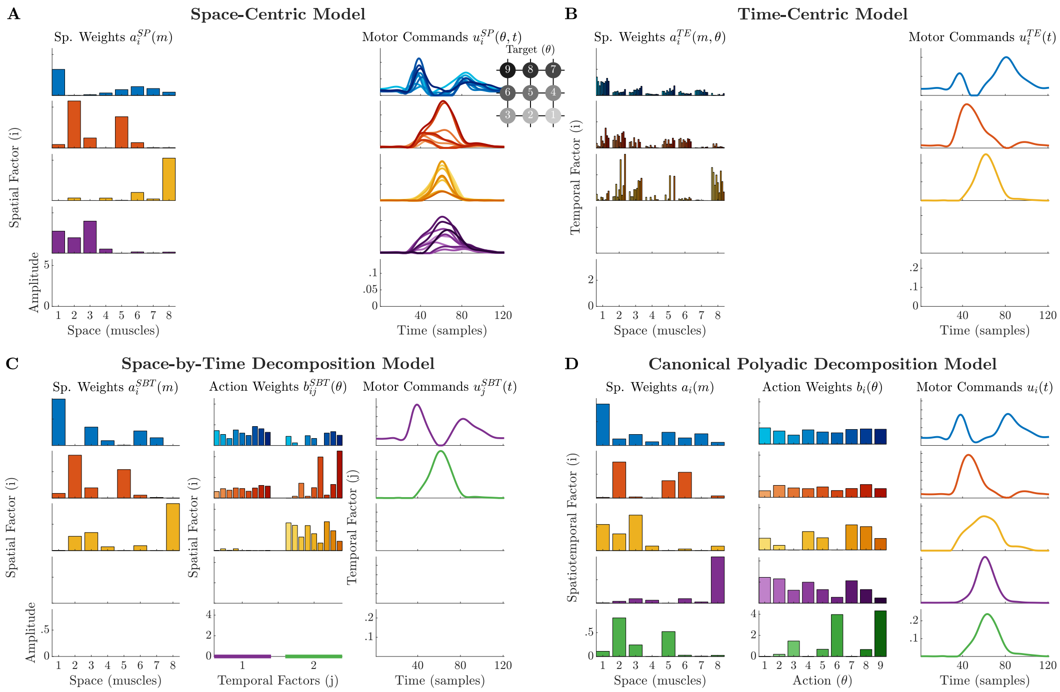

Figure 2: **Decomposition of muscle activity of representative participant.** (A) Space-centric decomposition into spatial modules (left) and time-dependent coefficients. (B) Time-centric decomposition into temporal modules (right) and space-dependent coefficients. (C) Space-by-time decomposition into spatial modules (left), temporal modules (right), and action coefficients (center); (D) Canonical polyadic decomposition into spatial (left), temporal (right) and action (center) modules. Note: the time-varying coefficients in (A) and the temporal modules in (B), (C), and (D) can be interpreted as motor commands sent to the muscles from higher motor centers.

in a single target domain while increasing the computational burden in the non-target domain. For example, models based on spatial modules [Tresch et al., 1999] simplify the control problem in the spatial (i.e., muscle) domain at the cost of complicating it in the temporal domain, where they assume the existence of time-varying coefficients that are specific to each action (Fig. 1A). According to this view, to specify the temporal activation patterns for a new movement, the motor system would need to find a way to determine a completely new waveform for each spatial module — in general, not a trivial problem. Likewise, models based on temporal modules [Ivanenko et al., 2004] assume the existence of muscle- and action-specific coefficients that need to be computed for every movement (Fig. 1B).

To meet the challenge of simultaneous identification of spatial and temporal modules, we propose a decomposition of muscle signals based on the Canonical Polyadic Decomposition (CPD) [Harshman et al., 1970] — a higher-order tensor decomposition method. The associated model factorizes muscle activity into fixed spatial and temporal modules that are flexibly modulated by space- and time-independent coefficients, depending on the movement to perform (Fig. 1D).

2 Modularity models for the space, time, and action domains

If the motor system generates muscle activation commands employing both spatial and temporal modules, it should be possible to approximate the muscle activity signals y recorded from muscle m during movement θ at time t as the sum of the product of the contributions of N_s spatial modules $\boldsymbol{a_i}$ and temporal modules $\boldsymbol{u_i}$, weighted by action-coding coefficients $\boldsymbol{b_i}$. That is:

$$y(m, \theta, t) \approx \sum_{i=1}^{N_s} a_i(m) \cdot b_i(\theta) \cdot u_i(t) \tag{1}$$

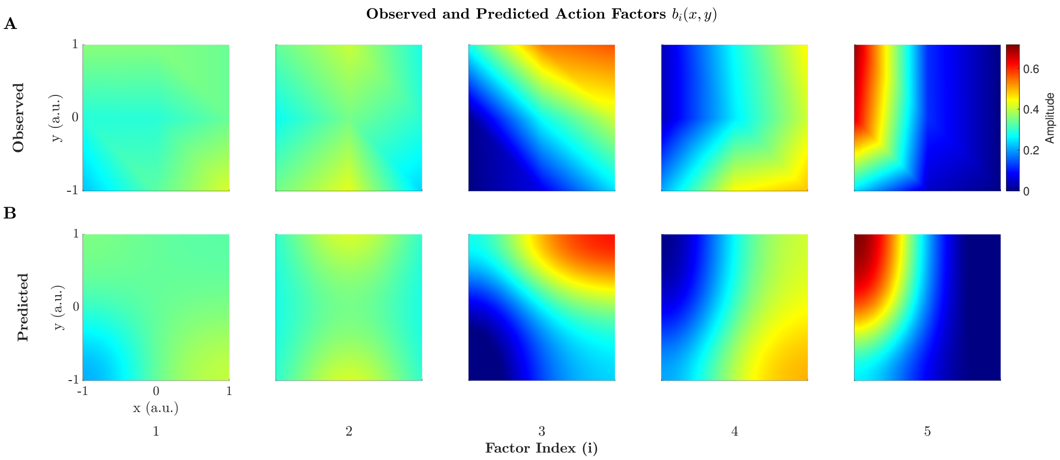

Figure 3: **Observed and reconstructed action manifold of gain-modulated network model** (A) Observed factors. (B) Factors predicted with radial basis function network.

To estimate the parameters of this model, we organize the muscle activations $y(m, \theta, t)$ into a third-order tensor $Y \in \mathbb{R}^{N_m X N_\theta X N_T}$ and fit a non-negative *Canonical Polyadic Decomposition* (CPD) model [Harshman et al., 1970], which approximates the original tensor as the sum of N_s rank-one tensors, with the non-negative randomized hierarchical alternating least squares (HALS) algorithm [Erichson et al., 2018]. Importantly, this decomposition differs from popular space-centric [Tresch et al., 1999] and time-centric [Ivanenko et al., 2004] decomposition models, which can only isolate either spatial or temporal modules, respectively.

Specifically, the *space-centric* decomposition is given by:

$$y(m, \theta, t) \approx \sum_{i=1}^{N_s^{SP}} a_i^{SP}(m) \cdot u_i^{SP}(\theta, t) \tag{2}$$

On the other hand, the *time-centric* decomposition is given by:

$$y(m, \theta, t) \approx \sum_{i=1}^{N_s^{TE}} a_i^{TE}(m, \theta) \cdot u_i^{TE}(t) \tag{3}$$

To estimate the parameters of the *space-centric* model, we first organize the muscle activation data into a matrix $Y^{SP} \in \mathbb{R}^{N_m X (N_\theta N_T)}$ — where the signals related to different movements are concatenated along the temporal dimension — and then apply non-negative matrix factorization [Lee and Seung, 1999]. Similarly, to estimate the parameters of the *time-centric* model, we apply non-negative matrix factorization to the matrix $Y^{TE} \in \mathbb{R}^{N_m N_\theta X N_T}$ — where the signals related to different movements are concatenated along the spatial dimension.

More recently, Delis et al. [2014] proposed a *space-by-time* decomposition model that has the potential to isolate both spatial and temporal modules underlying muscle signal activations. Compared to the CPD, the space-by-time decomposition accommodates different numbers of spatial and temporal modules, at the cost requiring the specification of action coefficient for each combination of spatial and temporal modules. Specifically, this decomposition is given by:

$$y(m, \theta, t) \approx \sum_{i=1}^{N_{ss}^{ST}} \sum_{j=1}^{N_{ts}^{ST}} a_i^{ST}(m) \cdot b_{ij}^{ST}(\theta) \cdot u_j^{ST}(t) \tag{4}$$

To fit this model, we applied the sample-based non-negative matrix tri-factorization algorithm (sNM3F — Delis et al. [2014]).

All of the above decomposition models require the a priori specification of the number of underlying spatial and/or temporal modules. Following standard practice (e.g., d'Avella et al. [2006]), to identify a plausible number of modules we first fit, for each decomposition model, models with a linearly increasing number of modules N. We then analyze how the coefficient of determination R^2 varies with N, and identified the elbow of the curve by locating the number of modules N^* from which the R^2 curve is well approximated by a line. The resulting R^2 curves we used to select the number of modules are reported in Fig.4. We note that, for the space-by-time decomposition model, we computed the R^2 curve by including, at each step, the temporal or spatial module that increased R^2 the most.

3 Dataset

To isolate the spatial and temporal invariances underlying muscle activity patterns during reaching movements, we analyzed the surface electromyographic data [Israely et al., 2018] recorded from healthy participants during the execution of reaches in the frontal plane. The reaches were performed towards nine targets arranged on a rectangular grid in front of the participants. The data were preprocessed with a custom pipeline and averaged across trials to isolate the condition-specific muscle patterns.

4 Validation

If the motor system makes use of spatial and temporal modules, their recruitment should not be random, but systematically vary with reaching direction. To quantify the regularity in the recruitment strategy, we fitted linear and radial basis function network *recruitment models* to the action-depended coefficients of the four decomposition models considered in this work, and quantified the reconstruction error. To further assess the robustness of the decomposition models and the regularity in the recruitment strategy, we measured the ability of the decomposition models to facilitate the zero-shot learning of muscle patterns for untrained reaching directions. To achieve this, we first fitted decomposition models on reduced datasets that excluded the data for one of the reaching directions; subsequently, we fitted recruitment models to the estimated action-dependent coefficients; finally, we used such recruitment models to estimate the action coefficients corresponding to the left-out reaching direction, and with these, the full set of muscle signals.

5 Results

We found that, compared with classical decomposition models [Tresch et al., 1999, Ivanenko et al., 2004], CPD identifies qualitatively similar spatial and temporal modules (Fig.2), explains a comparable amount of data variance (Fig. 4), and requires a lower number of parameters. Furthermore, we found that the space-by-time decomposition model [Delis et al., 2014], despite having a similar number of action-coding coefficients, tends to underfit the data (Fig. 4). Moreover, we found that the geometrical organization of the action coefficients in all models is not random, but describes a smooth manifold that is well approximated by simple recruitment models (e.g., see Fig.5 for the reconstruction quality of the action coefficients of the considered models, and Fig.3 for observed and estimated action coefficients of the CPD model). The smoothness of the action manifolds allows the zero-shot generation of muscle activity patterns for untrained reaching directions that closely resemble those experimentally recorded (Fig.6). However, the reconstruction quality obtained by fitting recruitment models on the space-by-time action coefficients tends to be worse than those of the other models (Fig. 5), consistently with what we observed when fitting the recruitment model to the muscle activity data (Fig. 4). Taken together, our results suggest that the decomposition proposed here represents a biologically plausible hierarchical organization of the control of reaching movements that the brain could leverage to control the body efficiently via compositional generalization.

References

Cameron Berg, Vittorio Caggiano, and Vikash Kumar. Sar: Generalization of physiological agility and dexterity via synergistic action representation. *arXiv preprint arXiv:2307.03716*, 2023.

Enrico Chiovetto, Alessandro Salatiello, Andrea d'Avella, and Martin A Giese. Toward a unifying framework for the modeling and identification of motor primitives. *Frontiers in Computational Neuroscience*, 16:926345, 2022.

Mark M Churchland, John P Cunningham, Matthew T Kaufman, Justin D Foster, Paul Nuyujukian, Stephen I Ryu, and Krishna V Shenoy. Neural population dynamics during reaching. *Nature*, 487 (7405):51–56, 2012.

Andrea d'Avella, Alessandro Portone, Laure Fernandez, and Francesco Lacquaniti. Control of fast-reaching movements by muscle synergy combinations. *Journal of Neuroscience*, 26(30): 7791–7810, 2006.

Ioannis Delis, Stefano Panzeri, Thierry Pozzo, and Bastien Berret. A unifying model of concurrent spatial and temporal modularity in muscle activity. *Journal of neurophysiology*, 111(3):675–693, 2014.

Sunny Duan, Mikail Khona, Adrian Bertagnoli, Sarthak Chandra, and Ila R. Fiete. See and copy: Generation of complex compositional movements from modular and geometric rnn representations. In Sophia Sanborn, Christian Shewmake, Simone Azeglio, Arianna Di Bernardo, and Nina Miolane, editors, *Proceedings of the 1st NeurIPS Workshop on Symmetry and Geometry in Neural Representations*, volume 197 of *Proceedings of Machine Learning Research*, pages 388–400. PMLR, 03 Dec 2023. URL `https://proceedings.mlr.press/v197/duan23a.html`.

N Benjamin Erichson, Ariana Mendible, Sophie Wihlborn, and J Nathan Kutz. Randomized nonnegative matrix factorization. *Pattern Recognition Letters*, 104:1–7, 2018.

Richard A Harshman et al. Foundations of the parafac procedure: Models and conditions for an" explanatory" multimodal factor analysis. 1970.

Corey B Hart and Simon F Giszter. A neural basis for motor primitives in the spinal cord. *Journal of Neuroscience*, 30(4):1322–1336, 2010.

Dieuwke Hupkes, Verna Dankers, Mathijs Mul, and Elia Bruni. Compositionality decomposed: How do neural networks generalise? *Journal of Artificial Intelligence Research*, 67:757–795, 2020.

Sharon Israely, Gerry Leisman, Chay C Machluf, and Eli Carmeli. Muscle synergies control during hand-reaching tasks in multiple directions post-stroke. *Frontiers in computational neuroscience*, 12:10, 2018.

Yuri P Ivanenko, Richard E Poppele, and Francesco Lacquaniti. Five basic muscle activation patterns account for muscle activity during human locomotion. *The Journal of physiology*, 556(1):267–282, 2004.

Brenden Lake and Marco Baroni. Generalization without systematicity: On the compositional skills of sequence-to-sequence recurrent networks. In *International conference on machine learning*, pages 2873–2882. PMLR, 2018.

Daniel D Lee and H Sebastian Seung. Learning the parts of objects by non-negative matrix factorization. *Nature*, 401(6755):788–791, 1999.

Ariel J Levine, Christopher A Hinckley, Kathryn L Hilde, Shawn P Driscoll, Tiffany H Poon, Jessica M Montgomery, and Samuel L Pfaff. Identification of a cellular node for motor control pathways. *Nature neuroscience*, 17(4):586–593, 2014.

Samuel Lippl and Kim Stachenfeld. When does compositional structure yield compositional generalization? a kernel theory. *arXiv preprint arXiv:2405.16391*, 2024.

Ziming Liu, Eric Gan, and Max Tegmark. Seeing is believing: Brain-inspired modular training for mechanistic interpretability. *Entropy*, 26(1):41, 2023.

Simon Schug, Seijin Kobayashi, Yassir Akram, Maciej Wołczyk, Alexandra Proca, Johannes Von Oswald, Razvan Pascanu, João Sacramento, and Angelika Steger. Discovering modular solutions that generalize compositionally. *arXiv preprint arXiv:2312.15001*, 2023.

Tomohiko Takei, Joachim Confais, Saeka Tomatsu, Tomomichi Oya, and Kazuhiko Seki. Neural basis for hand muscle synergies in the primate spinal cord. *Proceedings of the National Academy of Sciences*, 114(32):8643–8648, 2017.

Matthew C Tresch, Philippe Saltiel, and Emilio Bizzi. The construction of movement by the spinal cord. *Nature neuroscience*, 2(2):162–167, 1999.

A Appendix

A.1 Model fitting procedure

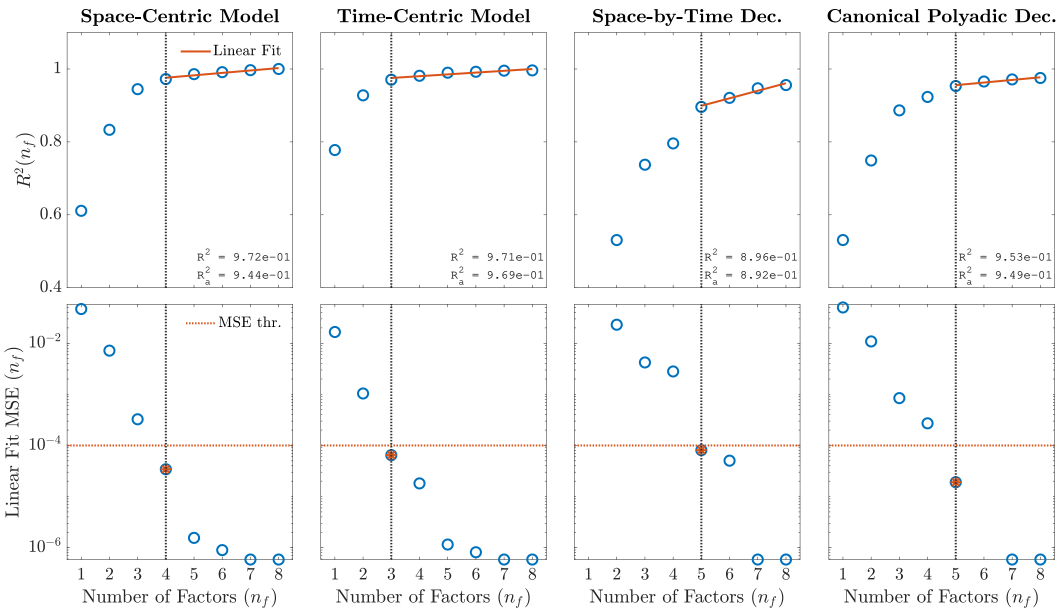

Figure 4: **$\mathbf{R^2}$ curves of the model fitting procedure used to determine the number of modules.**

189 **A.2 Fitting of action manifolds for zero-shot learning**

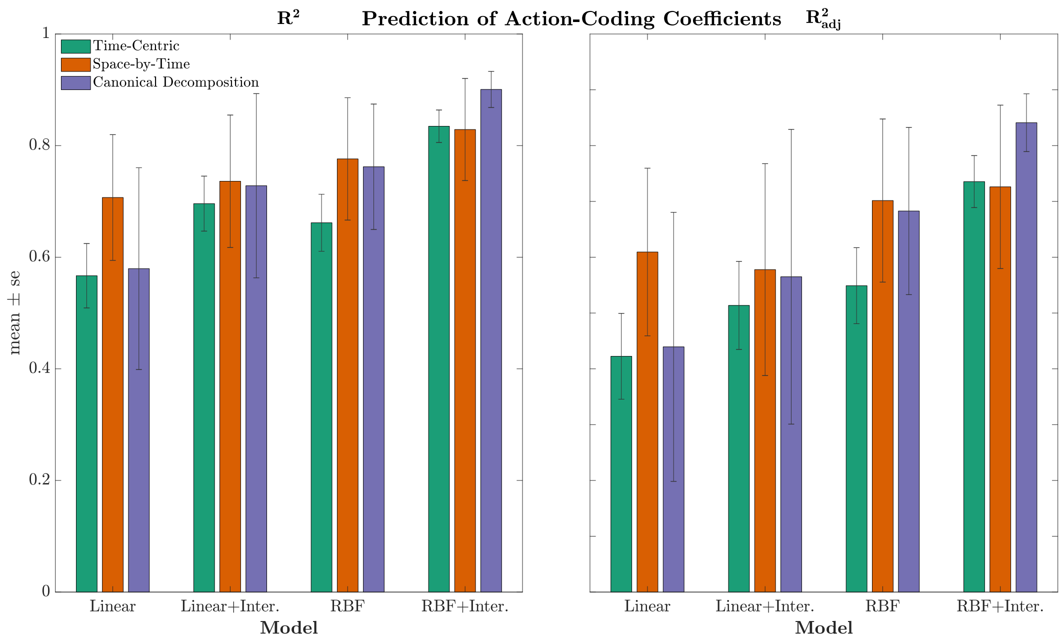

Figure 5: $\mathbf{R^2}$ **bars of the action-coefficient models**

A.3 Zero-shot generation of muscle activity

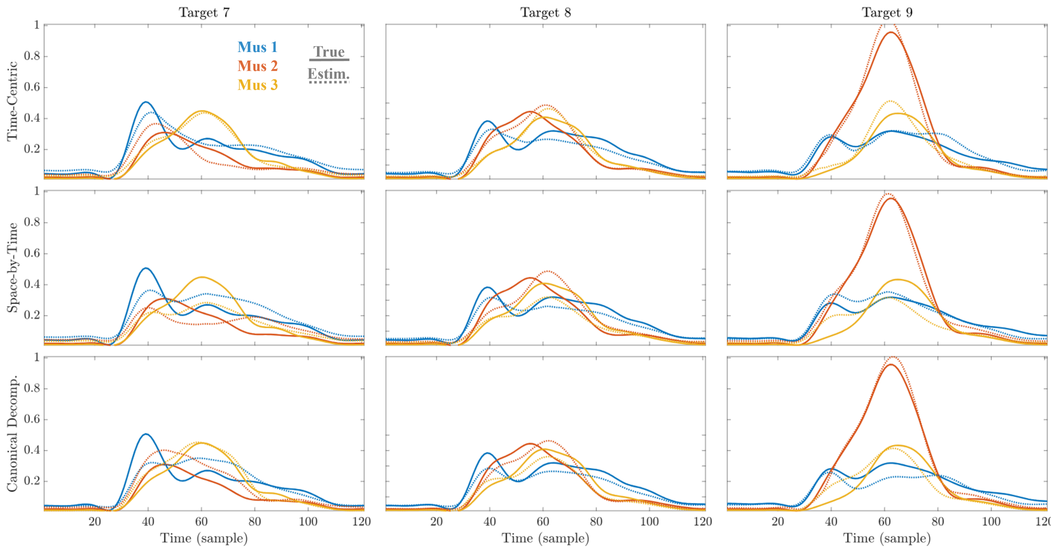

Figure 6: **True and zero-shot estimated muscle patterns.**