Hierarchical Control of Reaching Movements Via Compositional Gain Modulation

Alessandro Salatiello
Section for Computational Sensomotorics
Centre for Integrative Neuroscience
Hertie Institute for Clinical Brain Research
University Clinic Tübingen
72076 Tübingen, Germany
`alessandro.salatiello@uni-tuebingen.de`

Martin A. Giese
Section for Computational Sensomotorics
Centre for Integrative Neuroscience
Hertie Institute for Clinical Brain Research
University Clinic Tübingen
72076 Tübingen, Germany
`martin.giese@uni-tuebingen.de`

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 compositional 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

38th Conference on Neural Information Processing Systems (NeurIPS 2024).

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 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.

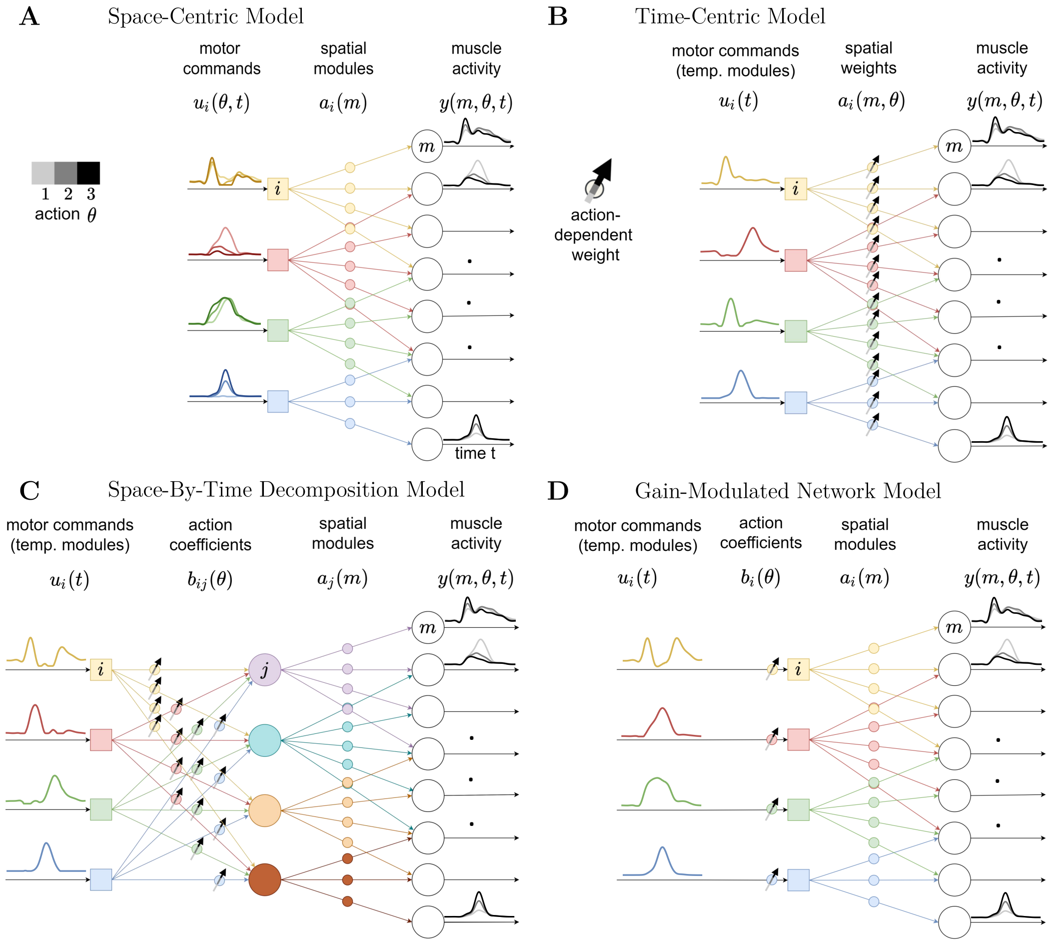

Figure 1: **Schematics of decomposition models.** Both the classical decomposition models (A-C) and the one proposed in this work (D) can be interpreted as simple linear feedforward neural networks. The space-centric model (A) is only able to capture the spatial invariances experimentally observed in 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

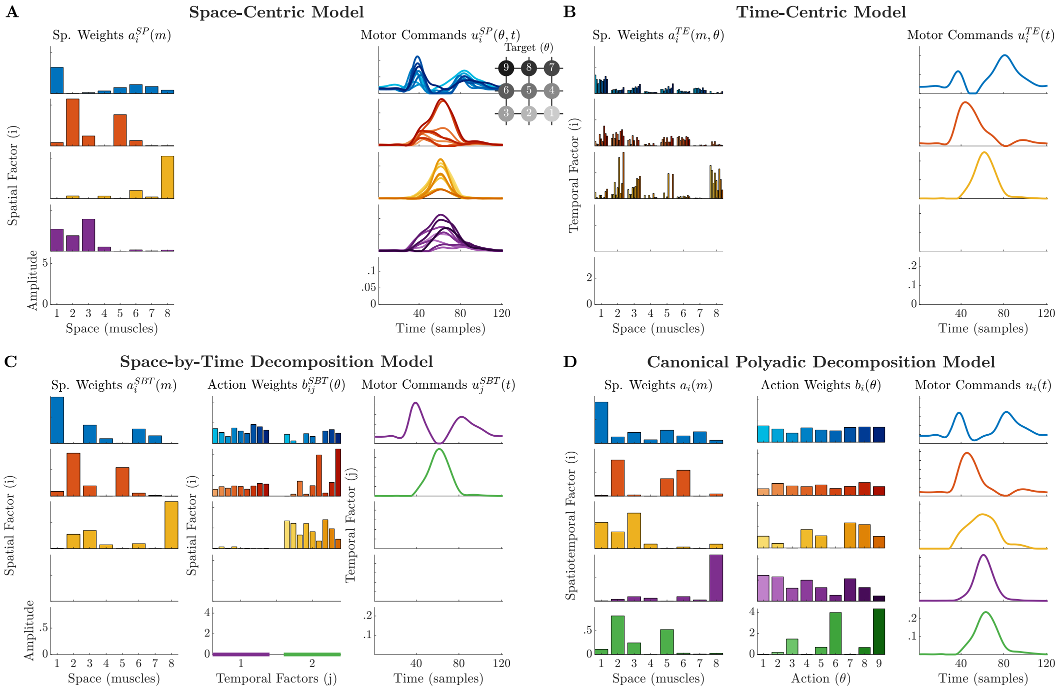

Figure 2: **Decomposition of muscle activity of representative participant.** (A) Space-centric decomposition into spatial modules (left) and time-dependent coefficients. (B) Time-centric decomposition into temporal modules (right) and space-dependent coefficients. (C) Space-by-time decomposition into 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.

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 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. In particular, we show that the CPD model can effectively factorize 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:

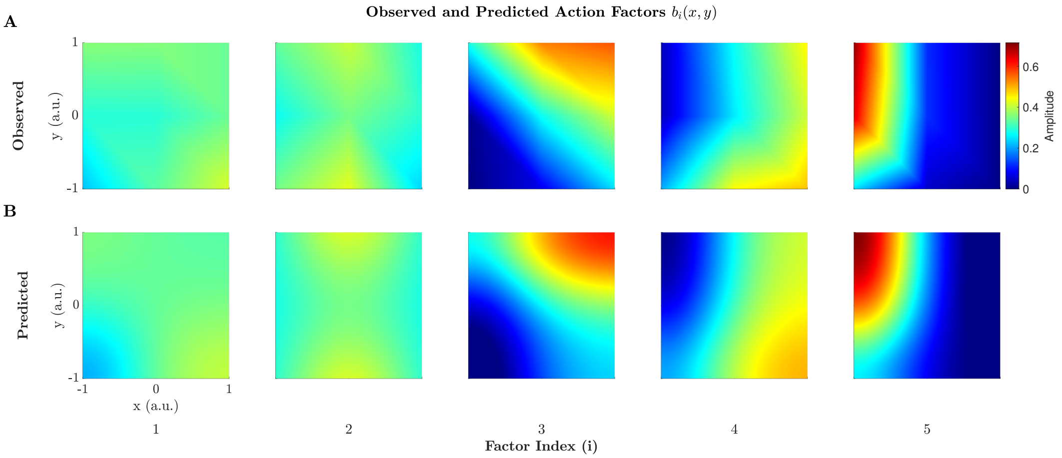

Figure 3: **Observed and reconstructed action manifold of gain-modulated network model** (A) Observed factors. (B) Factors predicted with radial basis function network.

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

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 [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 curve is well approximated by a line. The resulting 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 leverages 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 generation 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; we then fitted recruitment models to the estimated action-dependent coefficients; we finally 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. Specifically, recruitment models are obtained by fitting a real-valued function $f_i(x, y)$ that maps the target coordinates to the action-coding values isolated by the modularity models for each action-coding module i. Note that only the CPD model completely isolates the action-coding modules. Therefore, while for the CPD model we only need N_s recruitment models (i.e., one recruitment model per module), we need $N_s^{SP} \cdot N_T$ for the space-centric model, $N_s^{TE} \cdot N_M$ for the time-centric model, and $N_{ss}^{ST} \cdot N_{ts}^{ST}$ for the space-by-time model.

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.

6 Discussion

A fundamental research question in motor control is how the human motor system enables smooth behavior by effortlessly controlling the over 600 muscles that actuate the human body [Stewart and Cooley, 2009]. Modularity has been theorized [Sherrington, 1910] and identified [Tresch et al., 1999] early on as a fundamental organizing principle exploited by the motor system to both simplify the control problem and save computational resources. Nevertheless, it has never been clear how the system could exploit modularity in the spatial and temporal domains simultaneously. Here, we showed that a simple model – the CPD model [Harshman et al., 1970] – provides a potential solution to this problem. According to this model, new movements θ can be composed by simply specifying a vector of gains $\mathbf{b}(\theta)$, which determines the recruitment of *paired* spatial and temporal modules. Thus, the motor system can generate new movements through the principle of *compositional gain modulation*.

Gain modulation is an established neural computational mechanism [Salinas and Thier, 2000, Ferguson and Cardin, 2020] that allows modulating the input-output behavior of a neuron without altering other aspects of its response, such as its tuning properties. This kind of modulation appears to be critical for functions including coordinate transformation (e.g., from retinocentric to head-centric) [Salinas and Abbott, 1995], attention allocation [Treue and Trujillo, 1999], and context-dependent processing [Salinas, 2004]. Importantly, gain modulation has been observed not only in cortical but also in spinal networks, which are largely believed to be the locus of action-invariant motor modules [Tresch et al., 1999, Hart and Giszter, 2010, Levine et al., 2014, Takei et al., 2017]. Indeed, a few studies have found consistent evidence that the gain of spinal motoneurons can be modulated with subsecond precision [Hultborn et al., 2004, Wei et al., 2014, Vestergaard and Berg, 2015] to strike the right balance between movement precision and vigor for the intended movement.

The CPD model proposed here provides a simple gain-based mechanism to compose new movements by recruiting pre-existing motor modules. We have shown that, for simple movements like frontal reaches, action-coding gains describe a smooth manifold such that reaches to similar directions require similar gain vectors, facilitating the generation of reaches to new directions. The gains necessary for reaches to new directions, which can be considered as variations of the same basic movement, can be thus performed leveraging smooth recruitment models or simple similarity-based mechanisms. But how can such gains be learned for new classes of movements? Additionally, do these out-of-repertoire movements require learning new motor modules? First, gain modulation has been previously linked to learning [Eldar et al., 2013, Williams et al., 2018, Ferguson and Cardin, 2020], while lack of gain-modulating inputs has been shown to affect motor skill learning [Molina-Luna et al., 2009] and retention [Leow et al., 2013, Marinelli et al., 2017]. Thus, it appears that neural gain modulation is a mechanism the motor system can leverage to learn new movements. Second, although there is abundant evidence that different classes of movements can be performed via the adaptive recruitment of a fixed set of motor modules [d'Avella and Bizzi, 2005, Torres-Oviedo et al., 2006, Al Borno et al., 2020], new modules can be learned [Kargo and Nitz, 2003, Berger et al., 2013, Sawers et al., 2015]. Thus, out-of-repertoire movements can potentially be performed by reusing existing existing modules or learning new ones. However, a unifying computational theory of learning and recruitment of spatial and temporal modules is still lacking and warrants further research.

Finally, we note that the spatial and temporal modules recruited by the model to compose new movements are assumed to be pre-defined; thus, how they are implemented remains unclear. While spatial modules are likely to be encoded by synaptic weight patterns of effectively *feedforward* networks in the spinal cord [Tresch et al., 1999, Levine et al., 2014, Takei et al., 2017], temporal modules are likely to be generated by *recurrent* networks in the spinal cord [Grillner and El Manira, 2019, Lindén et al., 2022], primary motor cortex [Churchland et al., 2012, Russo et al., 2018], or motor thalamocortical networks [Guo et al., 2017, Logiaco et al., 2021]. Recent modeling work [Duan et al., 2023, Lindén et al., 2022, Salatiello and Giese, 2020, Stroud et al., 2018, Sussillo et al., 2015, Hennequin et al., 2014] was able to effectively model important features of the recurrent dynamics underlying the generation of movement-specific temporal patterns. Notably, some models elucidated how such temporal patterns can be learned without interference and chained together in time [Logiaco et al., 2021, Duan et al., 2023], and how they can be learned with a biologically plausible mechanism while supporting speed [Stroud et al., 2018, Lindén et al., 2022] and amplitude [Lindén et al., 2022] modulation; nevertheless, these models do not offer a clear separation of temporal and spatial modules, a core feature of compositional motor control.

References

Mazen Al Borno, Jennifer L Hicks, and Scott L Delp. The effects of motor modularity on performance, learning and generalizability in upper-extremity reaching: a computational analysis. *Journal of the Royal Society Interface*, 17(167):20200011, 2020.

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

Denise J Berger, Reinhard Gentner, Timothy Edmunds, Dinesh K Pai, and Andrea d'Avella. Differences in adaptation rates after virtual surgeries provide direct evidence for modularity. *Journal of Neuroscience*, 33(30):12384–12394, 2013.

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 and Emilio Bizzi. Shared and specific muscle synergies in natural motor behaviors. *Proceedings of the national academy of sciences*, 102(8):3076–3081, 2005.

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`.

Eran Eldar, Jonathan D Cohen, and Yael Niv. The effects of neural gain on attention and learning. *Nature neuroscience*, 16(8):1146–1153, 2013.

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

Katie A Ferguson and Jessica A Cardin. Mechanisms underlying gain modulation in the cortex. *Nature Reviews Neuroscience*, 21(2):80–92, 2020.

Sten Grillner and Abdeljabbar El Manira. Current principles of motor control, with special reference to vertebrate locomotion. *Physiological reviews*, 2019.

Zengcai V Guo, Hidehiko K Inagaki, Kayvon Daie, Shaul Druckmann, Charles R Gerfen, and Karel Svoboda. Maintenance of persistent activity in a frontal thalamocortical loop. *Nature*, 545(7653): 181–186, 2017.

Richard A Harshman et al. Foundations of the parafac procedure: Models and conditions for an "explanatory" multi-modal factor analysis. *UCLA working papers in phonetics*, 16(1):84, 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.

Guillaume Hennequin, Tim P Vogels, and Wulfram Gerstner. Optimal control of transient dynamics in balanced networks supports generation of complex movements. *Neuron*, 82(6):1394–1406, 2014.

Hans Hultborn, Robert B Brownstone, Tibor I Toth, and Jean-Pierre Gossard. Key mechanisms for setting the input–output gain across the motoneuron pool. *Progress in brain research*, 143:75–95, 2004.

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.

William J Kargo and Douglas A Nitz. Early skill learning is expressed through selection and tuning of cortically represented muscle synergies. *Journal of Neuroscience*, 23(35):11255–11269, 2003.

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.

Li-Ann Leow, Aymar De Rugy, Andrea M Loftus, and Geoff Hammond. Different mechanisms contributing to savings and anterograde interference are impaired in parkinson's disease. *Frontiers in human neuroscience*, 7:55, 2013.

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.

Henrik Lindén, Peter C Petersen, Mikkel Vestergaard, and Rune W Berg. Movement is governed by rotational neural dynamics in spinal motor networks. *Nature*, 610(7932):526–531, 2022.

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.

Laureline Logiaco, LF Abbott, and Sean Escola. Thalamic control of cortical dynamics in a model of flexible motor sequencing. *Cell reports*, 35(9), 2021.

Lucio Marinelli, Angelo Quartarone, Mark Hallett, Giuseppe Frazzitta, and Maria Felice Ghilardi. The many facets of motor learning and their relevance for parkinson's disease. *Clinical Neurophysiology*, 128(7):1127–1141, 2017.

Katiuska Molina-Luna, Ana Pekanovic, Sebastian Röhrich, Benjamin Hertler, Maximilian Schubring-Giese, Mengia-Seraina Rioult-Pedotti, and Andreas R Luft. Dopamine in motor cortex is necessary for skill learning and synaptic plasticity. *PloS one*, 4(9):e7082, 2009.

Abigail A Russo, Sean R Bittner, Sean M Perkins, Jeffrey S Seely, Brian M London, Antonio H Lara, Andrew Miri, Najja J Marshall, Adam Kohn, Thomas M Jessell, et al. Motor cortex embeds muscle-like commands in an untangled population response. *Neuron*, 97(4):953–966, 2018.

Alessandro Salatiello and Martin A Giese. Recurrent neural network learning of performance and intrinsic population dynamics from sparse neural data. In *International Conference on Artificial Neural Networks*, pages 874–886. Springer, 2020.

Emilio Salinas. Fast remapping of sensory stimuli onto motor actions on the basis of contextual modulation. *Journal of Neuroscience*, 24(5):1113–1118, 2004.

Emilio Salinas and Larry F Abbott. Transfer of coded information from sensory to motor networks. *Journal of Neuroscience*, 15(10):6461–6474, 1995.

Emilio Salinas and Peter Thier. Gain modulation: a major computational principle of the central nervous system. *Neuron*, 27(1):15–21, 2000.

Andrew Sawers, Jessica L Allen, and Lena H Ting. Long-term training modifies the modular structure and organization of walking balance control. *Journal of neurophysiology*, 114(6):3359–3373, 2015.

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.

Charles Scott Sherrington. Flexion-reflex of the limb, crossed extension-reflex, and reflex stepping and standing. *The Journal of physiology*, 40(1-2):28, 1910.

Gregory Stewart and Denton A Cooley. *The skeletal and muscular systems*. Infobase Publishing, 2009.

Jake P Stroud, Mason A Porter, Guillaume Hennequin, and Tim P Vogels. Motor primitives in space and time via targeted gain modulation in cortical networks. *Nature neuroscience*, 21(12): 1774–1783, 2018.

David Sussillo, Mark M Churchland, Matthew T Kaufman, and Krishna V Shenoy. A neural network that finds a naturalistic solution for the production of muscle activity. *Nature neuroscience*, 18(7): 1025–1033, 2015.

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.

Gelsy Torres-Oviedo, Jane M Macpherson, and Lena H Ting. Muscle synergy organization is robust across a variety of postural perturbations. *Journal of neurophysiology*, 96(3):1530–1546, 2006.

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

Stefan Treue and Julio C Martinez Trujillo. Feature-based attention influences motion processing gain in macaque visual cortex. *Nature*, 399(6736):575–579, 1999.

Mikkel Vestergaard and Rune W Berg. Divisive gain modulation of motoneurons by inhibition optimizes muscular control. *Journal of Neuroscience*, 35(8):3711–3723, 2015.

Kunlin Wei, Joshua I Glaser, Linna Deng, Christopher K Thompson, Ian H Stevenson, Qining Wang, Thomas George Hornby, Charles J Heckman, and Konrad P Kording. Serotonin affects movement gain control in the spinal cord. *Journal of Neuroscience*, 34(38):12690–12700, 2014.

Alex H Williams, Tony Hyun Kim, Forea Wang, Saurabh Vyas, Stephen I Ryu, Krishna V Shenoy, Mark Schnitzer, Tamara G Kolda, and Surya Ganguli. Unsupervised discovery of demixed, low-dimensional neural dynamics across multiple timescales through tensor component analysis. *Neuron*, 98(6):1099–1115, 2018.

A Appendix

A.1 Model fitting procedure

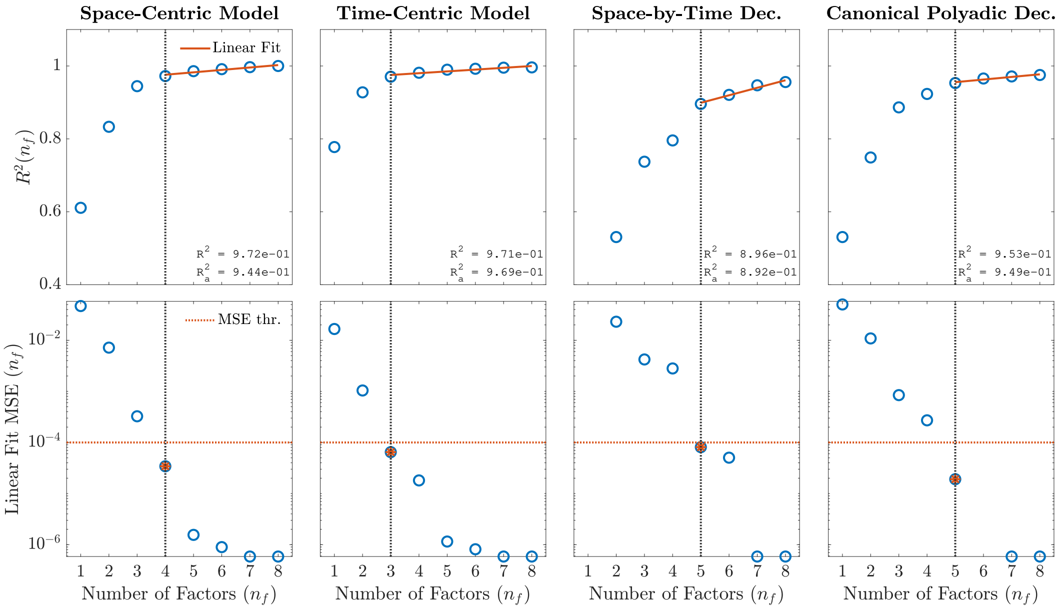

Figure 4: $\mathbf{R^2}$ **curves of the model fitting procedure used to determine the number of modules.** To select the optimal number of motor modules n_f for each model, we followed the standard practice [d'Avella et al., 2006] of identifying the elbow of the $R^2(n_f)$ curve, defined as the point from which the curve is well approximated by a line (MSE < 10^{-4}). For the space-by-time decomposition model, we computed the curve by including, at each step, the temporal or spatial module that increased R^2 the most [Delis et al., 2014]. The top panel shows the R^2 curves, while the bottom panel shows the MSE of the linear fits.

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

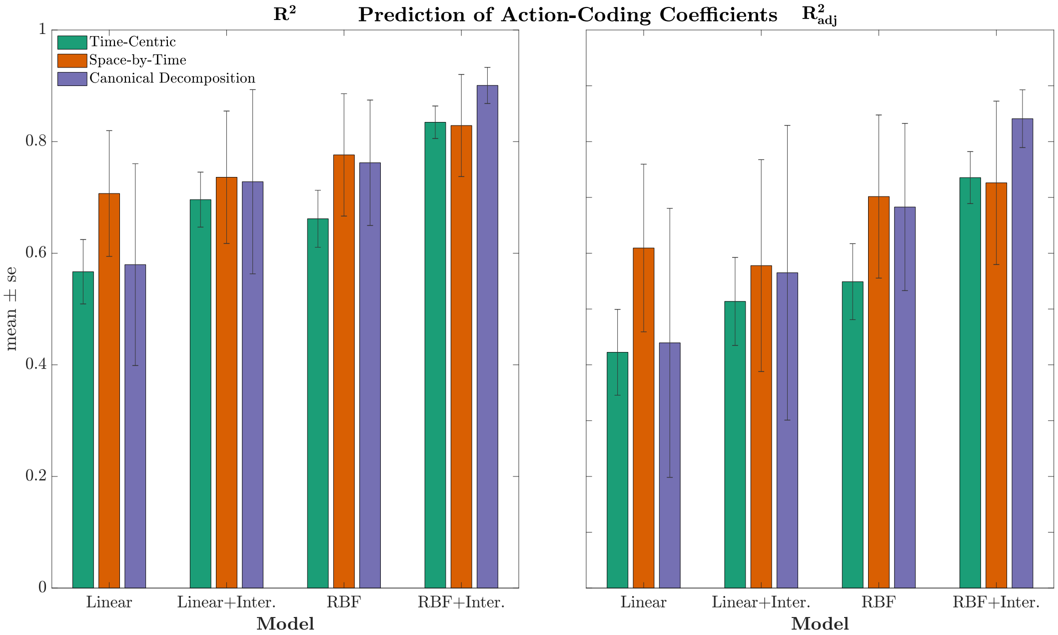

Figure 5: **$\mathbf{R^2}$ bars of the recruitment models**. Recruitment models are obtained by fitting a real-valued function $f_i(x, y)$ of two variables that maps the target coordinates x, y to the action-coding values isolated by the modularity models for each action-coding module i. For some recruitment models (Linear+Inter. and RBF+Inter), we also consider an interaction term $z = xy$.

A.3 Zero-shot generation of muscle activity

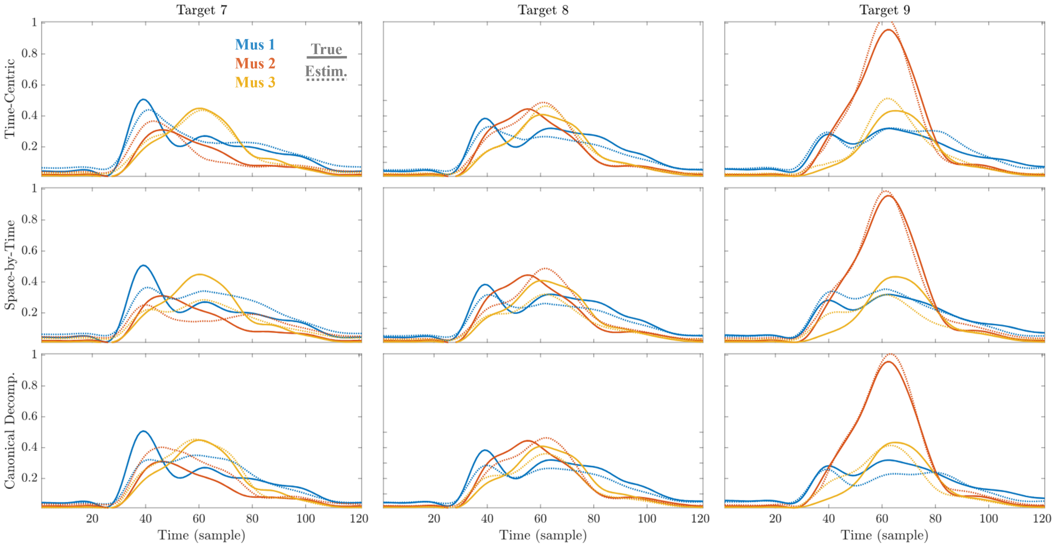

Figure 6: **True and zero-shot estimated muscle patterns.** Each row shows the zero-shot predictions of one modularity model for three example reaching targets (7, 8, and 9) and three example muscles (1, 2, and 3). Model predictions are represented by dotted lines, while ground-truth data are represented by solid lines. For each reaching target direction, the predictions are based on paired modularity and recruitment models trained on the remaining 8 target directions. The recruitment model used to generate this figure is the RBF+Inter. model.