
# Section: Chiappa et al. 2024

# Background Recap

EPFL

# Muscle synergies as principle for motor control

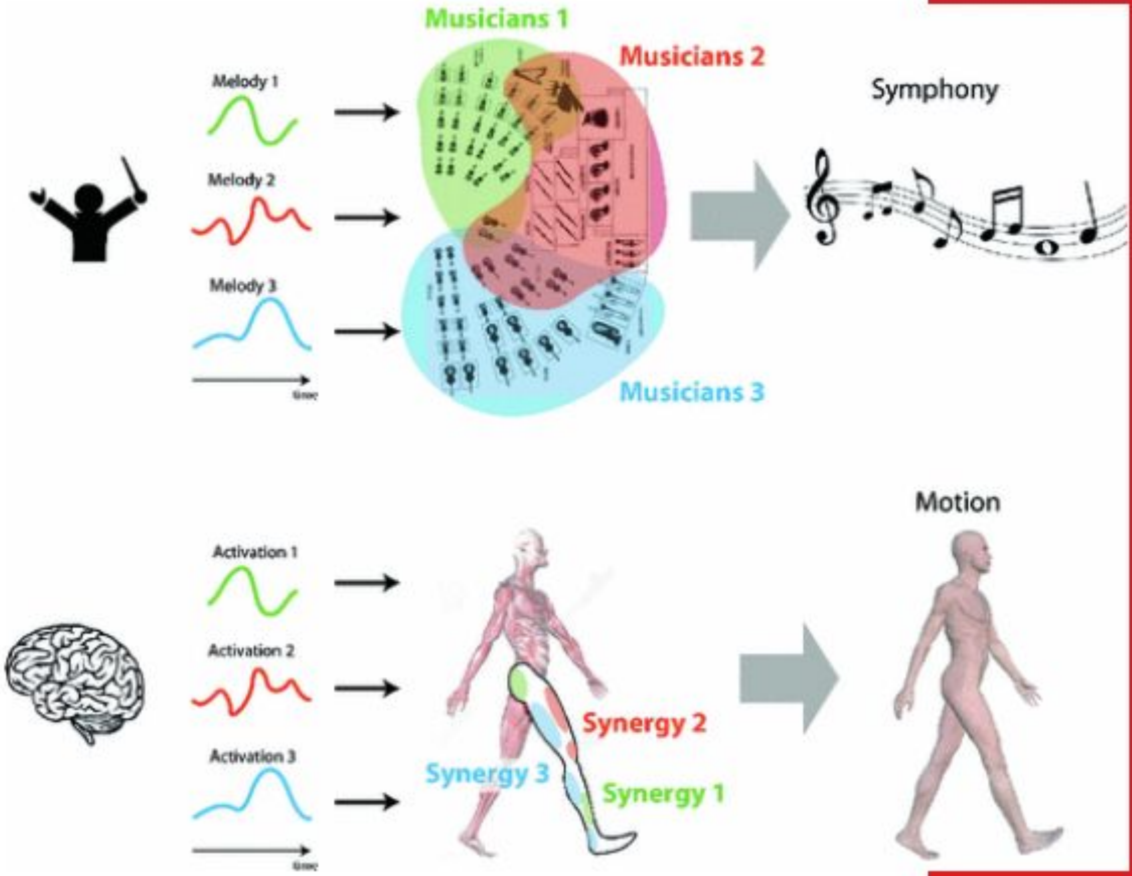


Torricelli et al., Emerging Therapies in Neurorehabilitation 2015

# Curriculum learning

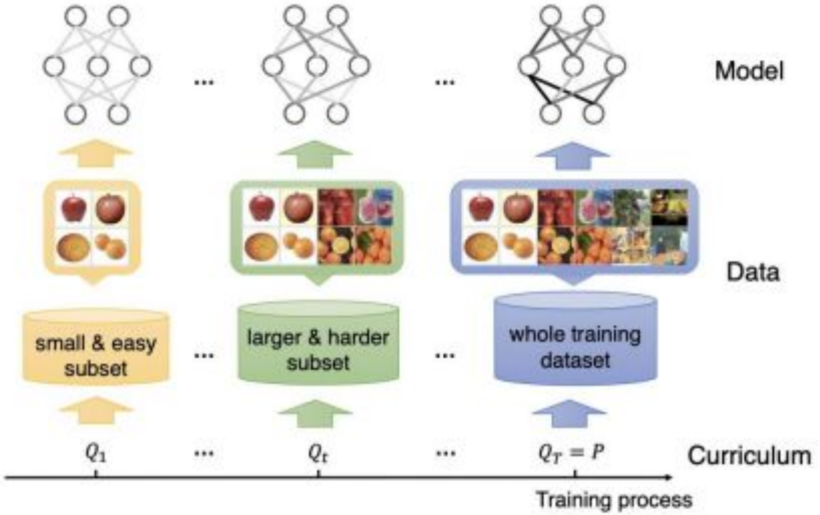


Wang, Chen, Zhu IEEE TPAMI, 2021

# Inspiration from coaching: part-to-whole practice

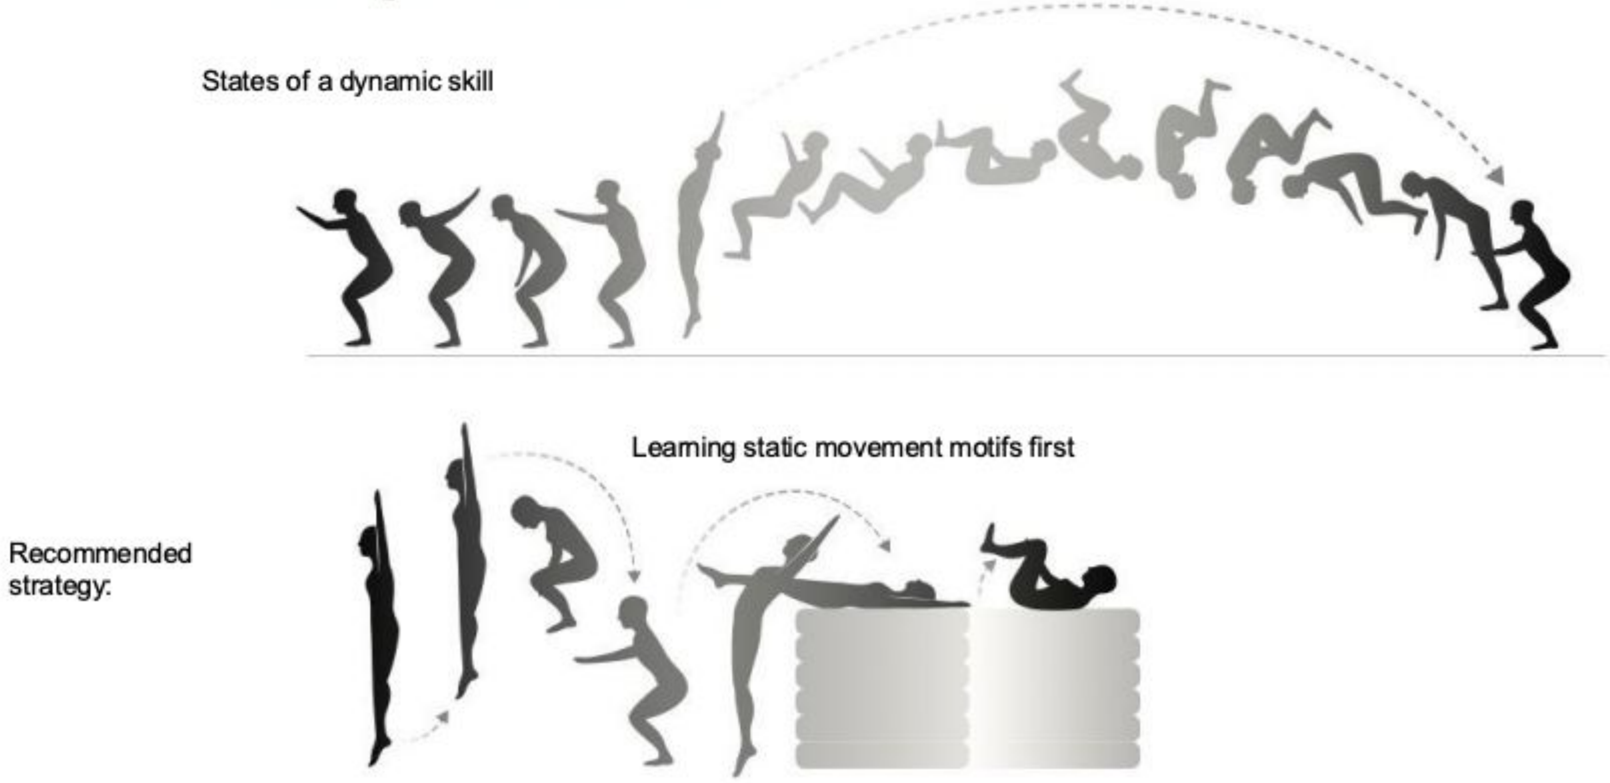


Chiappa*, A. S., Tano*, P., Patel*, N., Ingster, A., Pouget, A., & Mathis, A. Neuron 2024

EPFL

# How do humans achieve this task?

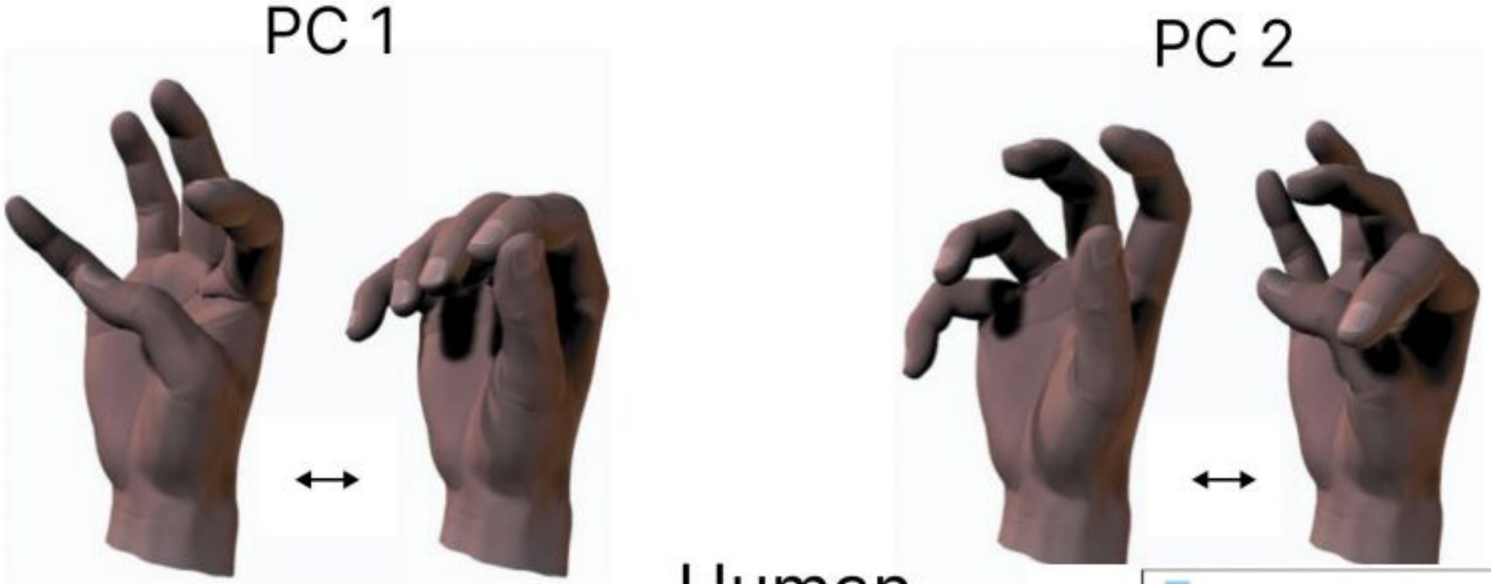


Human

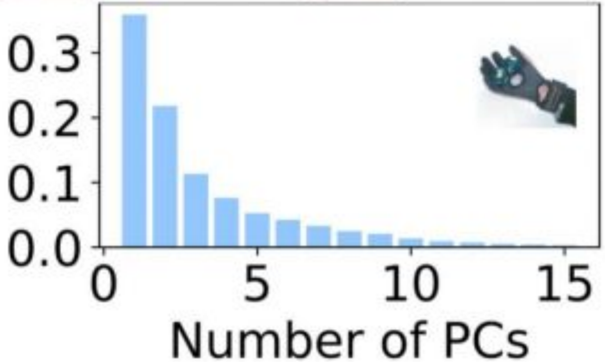


- *Todorov, & Ghahramani,*
  *Annual International Conference of the IEEE Engineering in Medicine and Biology 2004*

## Section Paper:

# Acquiring musculoskeletal skills with curriculum-based reinforcement learning

Alberto Silvio Chiappa [1,2,4] · Pablo Tano [3,4] · Nisheet Patel [3,4] · Abigaïl Ingster [1,2] · Alexandre Pouget [3] · Alexander Mathis [1,2,5,6]

Efficient musculoskeletal simulators and powerful learning algorithms provide computational tools to tackle the grand challenge of understanding biological motor control. Our winning solution for the inaugural NeurIPS MyoChallenge leverages an approach mirroring human skill learning. Using a novel **curriculum learning approach**, we trained a recurrent neural network to control a realistic model of the human hand with 39 muscles to rotate two Baoding balls in the palm of the hand. In agreement with data from human subjects, **the policy uncovers a small number of kinematic synergies**, even though it is not explicitly biased toward low-dimensional solutions. However, selectively inactivating parts of the control signal, **we found that more dimensions contribute to the task performance than suggested by traditional synergy analysis**. Overall, our work illustrates the emerging possibilities at the interface of musculoskeletal physics engines, reinforcement learning, and neuroscience to advance our understanding of biological motor control.

# Figure 1: Definition of the SDS curriculum, performance benchmarks, and ablation study

A The musculoskeletal hand

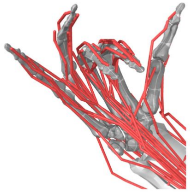

## Explain the task tackled by the authors in the MyoChallenge?

*Boading Balls: make balls rotate in tandem along a circular trajectory in a hand (39 muscles, 23 joints) to follow a pair of moving targets.*

## Why is a model-free method like PPO failing on that task according to the authors? (in text)

*Sparse reward function: insufficient to develop effective policies.*
*Gradient of the task objective needs to be inferred from the reward function.*

B Task variations in Phase II

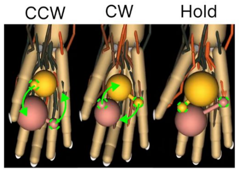


## How is the SDS curriculum enhancing the model-free optimization process?

*structured learning pathway → 1: hold the balls, learns multiple static configuration, 2: dynamical transition, gradually merge the confs via incr. faster dynamic traj, 3: random target initialization / domain randomization. Note: no parallelization possible (400h training).*

F Schematic of the SDS curriculum

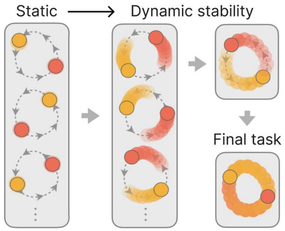


I MyoChallenge Phase II curriculum

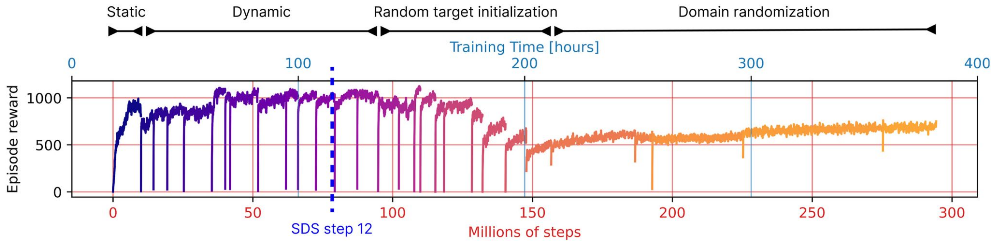

# Figure 2: Kinematic and dynamic motifs within one full rotation of each Baoding ball for one episode of the SDS policy, after completion of the training curriculum

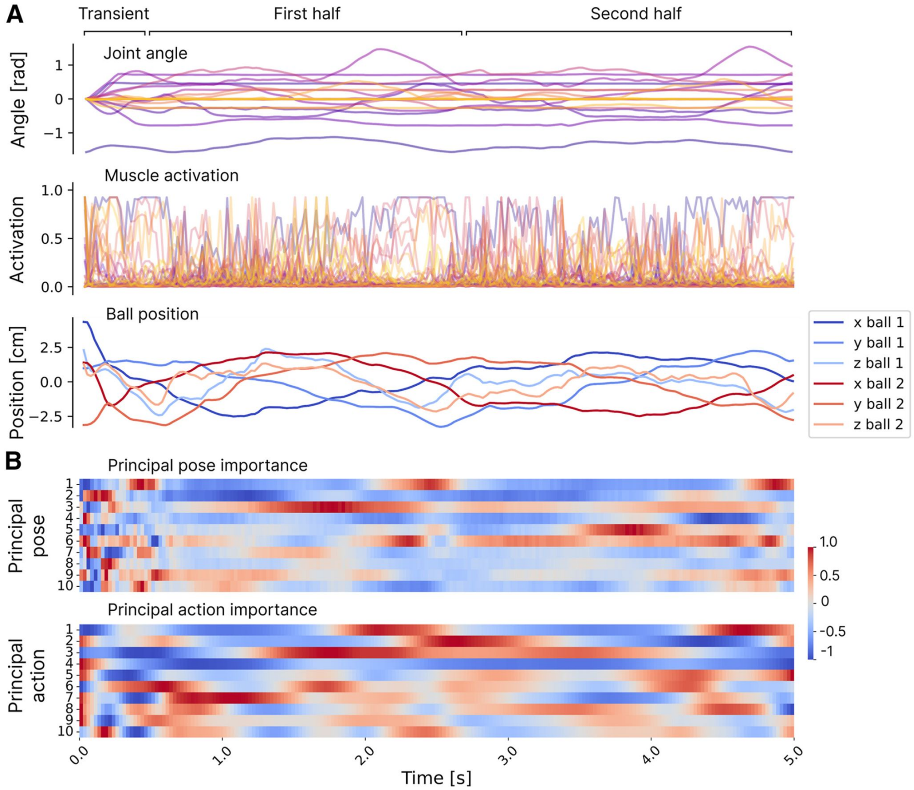


**What is their hypothesis regarding motor synergies of their artificial agent compared to biological agents?**

*They learn to operate in a reduced kinematic (pose) and dynamic (muscles) space, similar to the coordinated patterns of joint angles in low dim subspace observed in motor control studies.*

**How do they confirm it? (Panel B – first row)**

*PCs of the hand poses (joint angles) and policy's actions (muscle activation) and evaluate their significance. Same PC reused to swap the positions of the balls.*

# Figure 3: Dimensionality of the control policy and SDS policy in the Baoding balls task

**What can you say about the number of synergies used to control the hand? Comment on tasks and artificial/biological subjects (Panels A-B-C)**

- *Only a few synergies to capture most of the variance in posture (joint angle) and muscle space (muscle activations).*
- *Dim of hand poses during Baoding lower than control task for both SDS and human.*
- *Dim of hand poses really similar between SDS and human.*
- *Difference of dim for angular velocity of the joints between SDS and human: expected as instructions on speed were different.*
- *Muscle activations dim higher for baoding: more complex co-activation of antagonist muscles.*

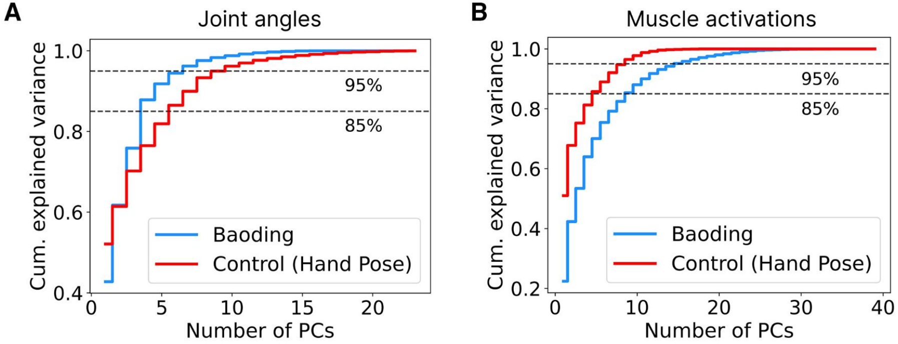


C

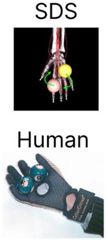


| | | Position | Velocity | Muscle act. |
|---|---|---|---|---|
| SDS | Baoding | 4.5 | 8.5 | 12 |
| | Control | 8 | 5.5 | 7 |
| Human | Baoding | 5 | 6 | |
| | Control | 8.5 | 10 | |

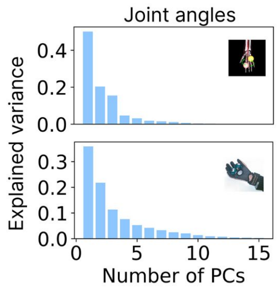

# Figure 4: Task relevance of the low-variance PCs of the SDS policy

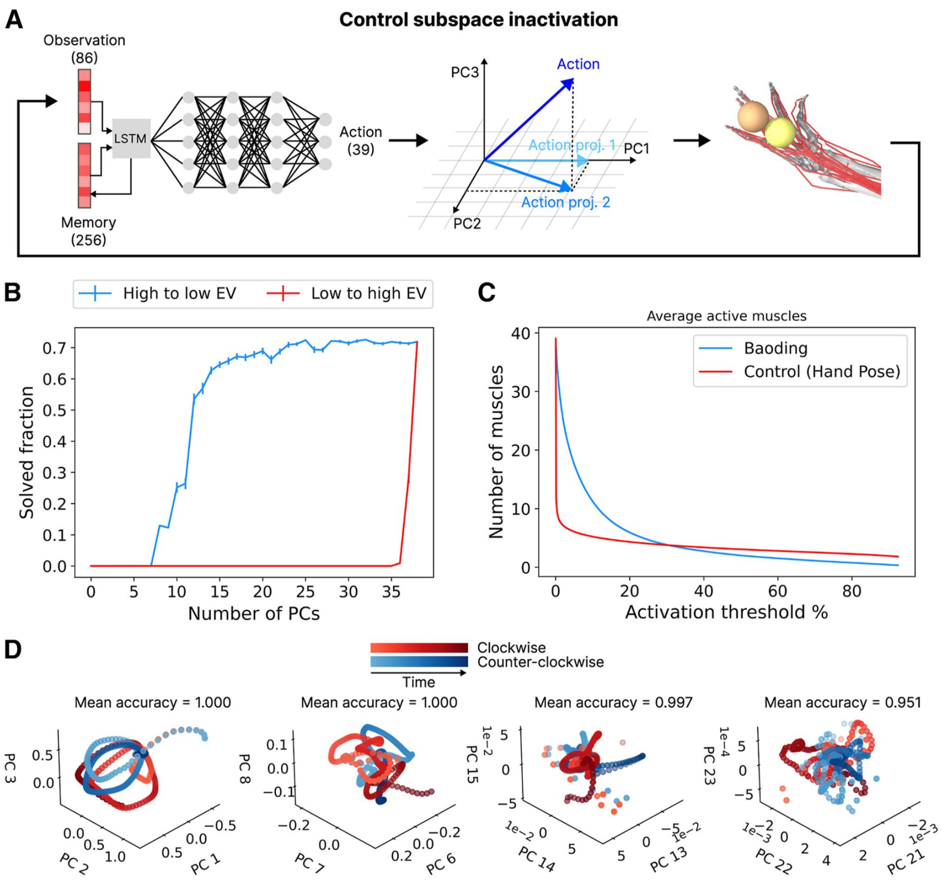


**Explain the control subspace inactivation (CSI) procedure proposed by the authors. What is it used for? What are the results ?**

- *The control subspace inactivation means the projection of the pattern of muscle activity (the action) onto specifics sets of PCs*
- *It is used to measure how the task performance varies as a function of the dimensionality of the enforced space.*
- *The highest variance PCs are necessary to solve the tasks, and the lowest variance PCs also helps to improve performance, as the lowest 12 PCs are sufficient for 50% performance.*
- *All PCs are task dependents based on the decoding accuracy.*

# Figure 5: Transfer of muscle synergies from other tasks to Baoding

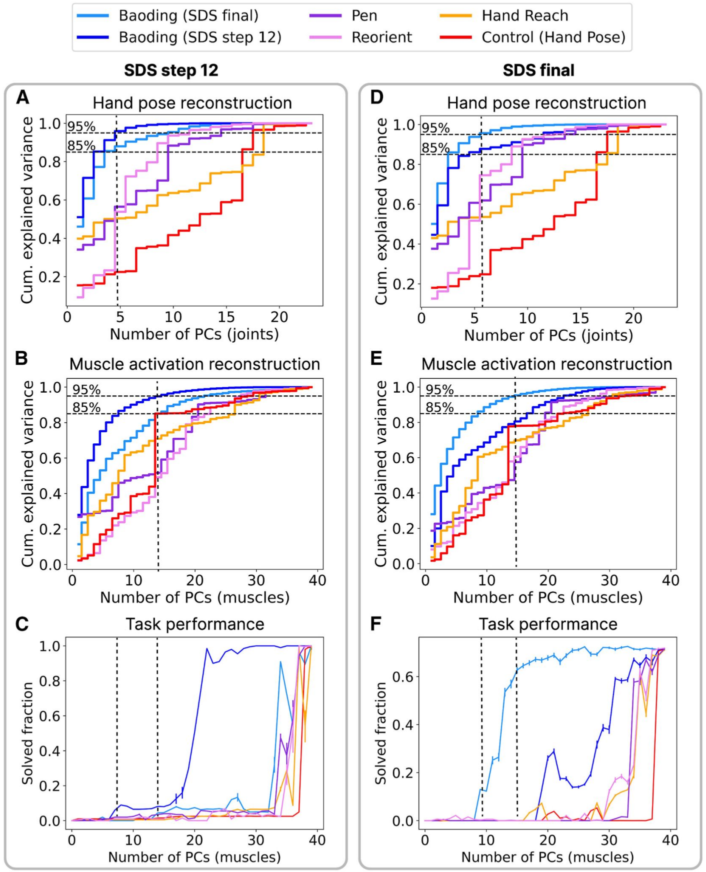


## What are the differences between SDS step 12 and SDS final ?

- *The PCA of the hand kinematics shows that the final SDS policy and the SDS step 12 policy are the very similar, while in the muscle-activation space, the SDS step 12 and the final SDS are more different*
- *Require different muscle synergy spaces for high performance*

## Can policies transfer from one task to another ?

- *Not unless at least 30 PCs are used.*

# Figure 6: Population activity of the SDS policy network

**What can be said about the differences across layers ?**

- *Memory subspace separates tasks better than observation (as each tasks require separate motor plans).*
- *Action space does separate as similar motor synergies are used on both tasks.*

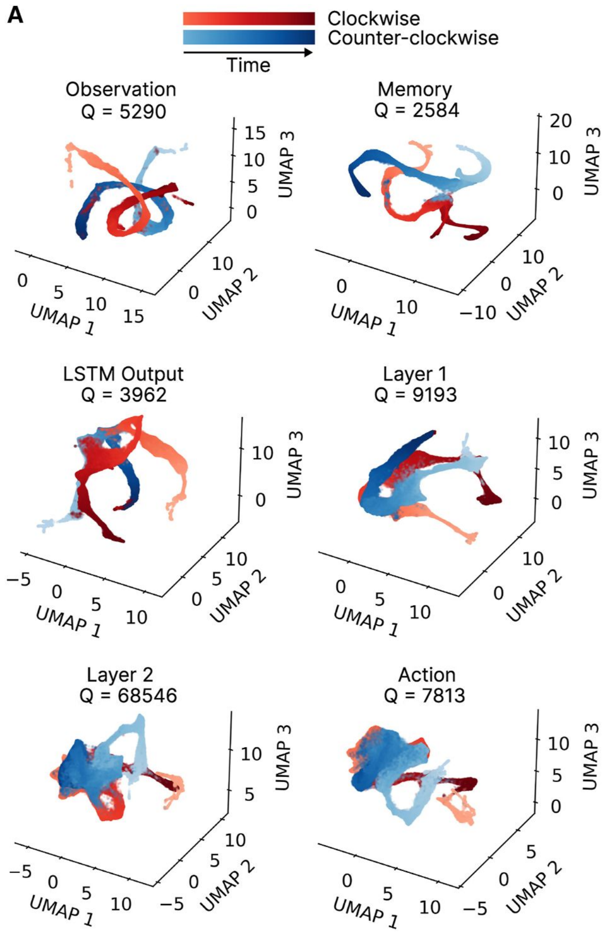

# Paper round-up

- They succeeded in training a musculoskeletal model on an object-manipulation task.
- They propose a static to dynamic stabilization (SDS) curriculum, inspired by coaching practice.
- They show that, akin to experimental data, SDS learns low-dimensional kinematic and kinetic spaces.
- They show that muscle synergies are highly task specific and thus generalize poorly.
- They found that more dimensions contributed to the task performance than suggested by traditional synergy analysis.
- They found lower tangling of the dynamics in the controller state space than in the action state space, consistent with previous observations that motor cortical dynamics avoiding tangling more than muscle dynamics.

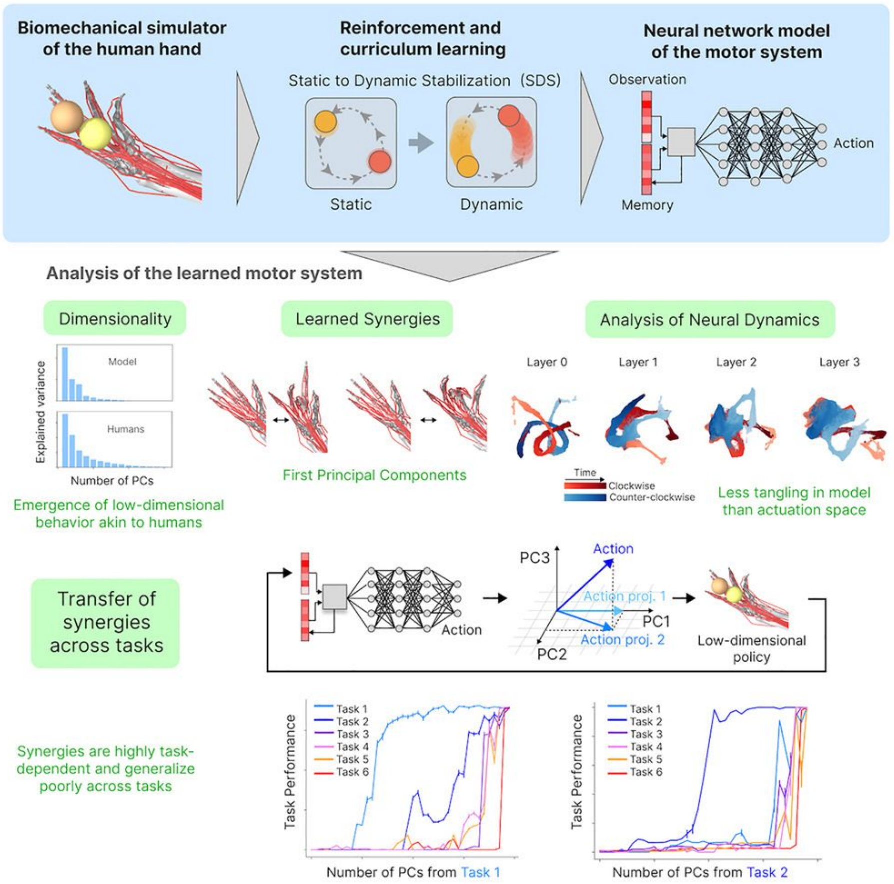

# What did we learn? What questions do we have?

- **What points do they make in the discussion?**

- **Is anything unclear?**

- **What would you do next if you had to design an experiment?**
  - *Add haptic feedback in the model.*
  - *Implement curriculum-based RL on other tasks.*