NX-435: Systems neuroscience

Alexander Mathis
alexander.mathis@epfl.ch

Guest lecture, May 15th

Reminder: Anatomy of motor control & pattern generators

Vertebrate motor control

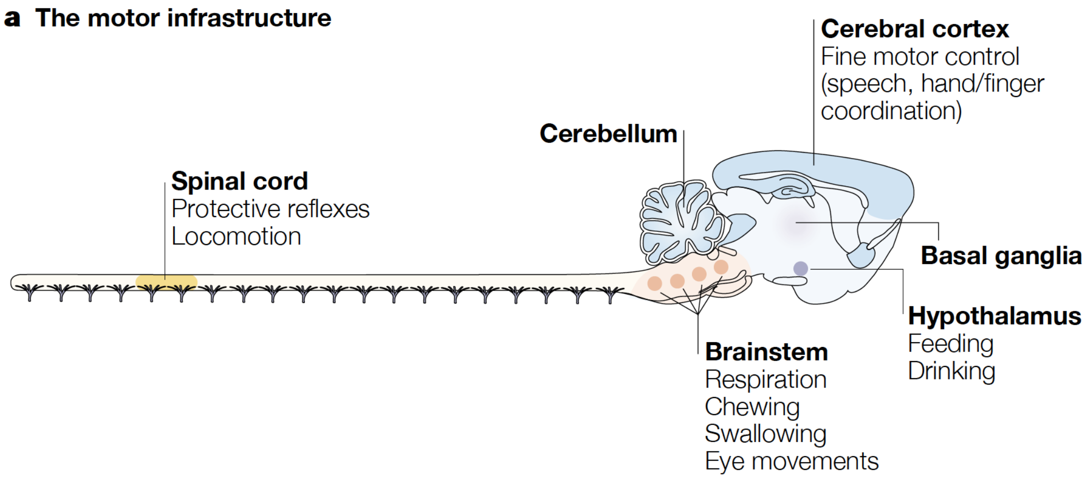

Grillner, Nature Review Neuroscience 2003

Vertebrate motor control

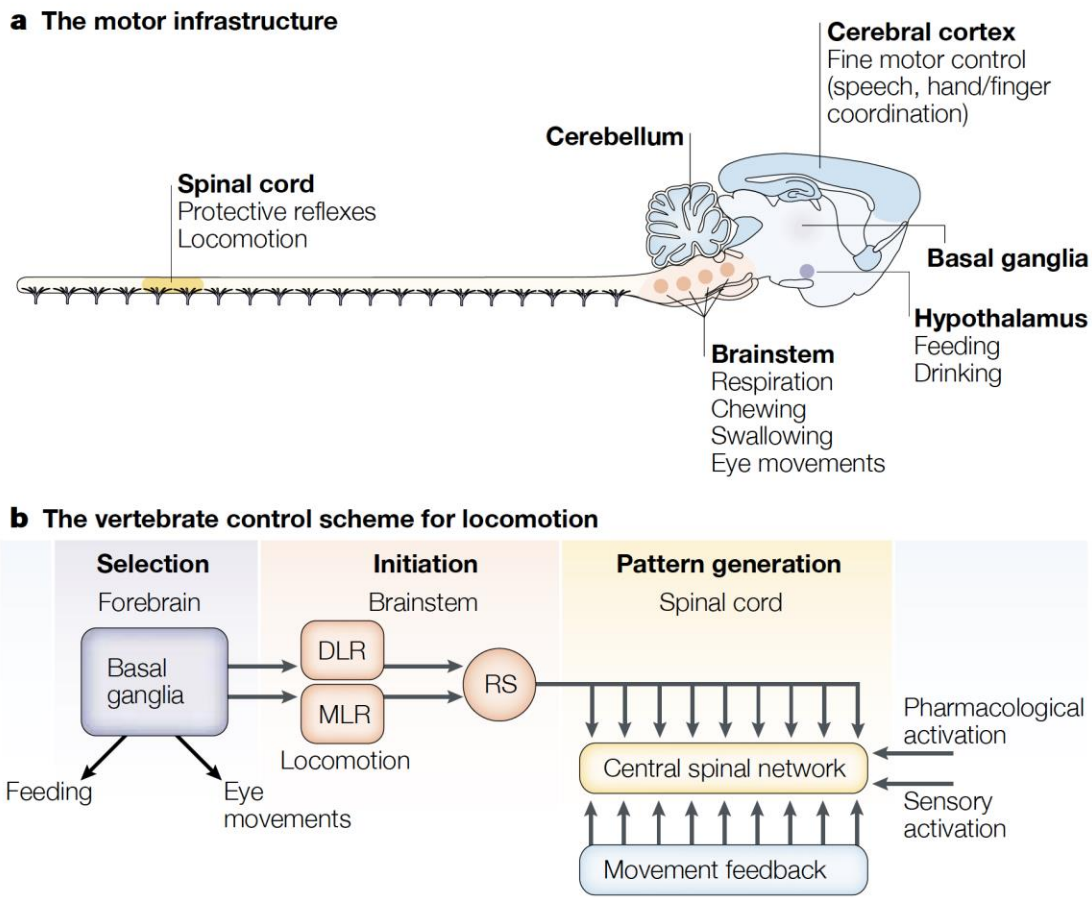

Grillner, Nature Review Neuroscience 20

Pattern generation in the intact lamprey and an isolated spinal circuit

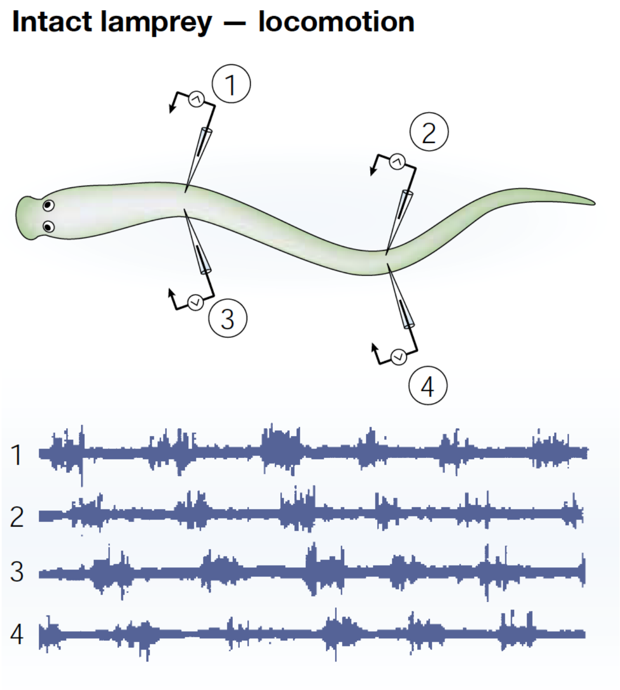

Note: alternation of 1/3 and 2/4 plus lag between 1 and 2.

Grillner, Nature Review Neuroscience 2003

Pattern generation in the intact lamprey and an isolated spinal circuit

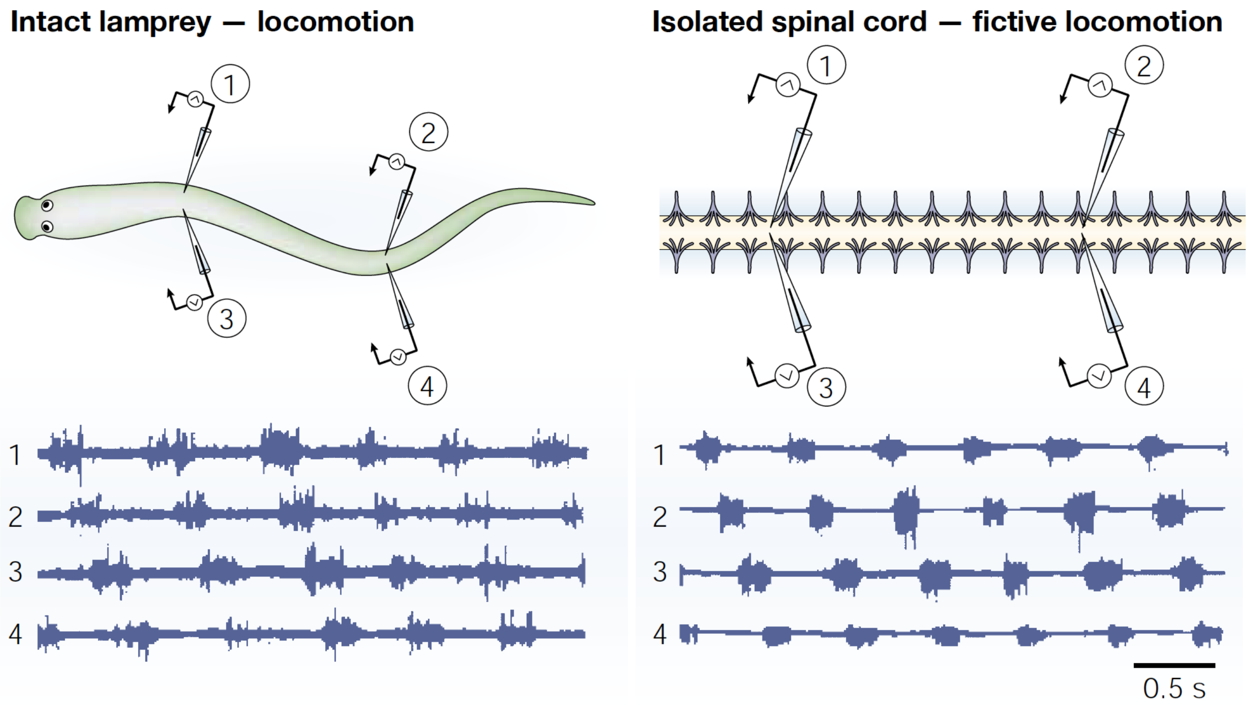

superfusion of glutamate agonists

Grillner, Nature Review Neuroscience 2003

Brain stem circuits to control locomotion

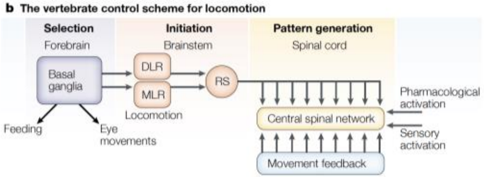

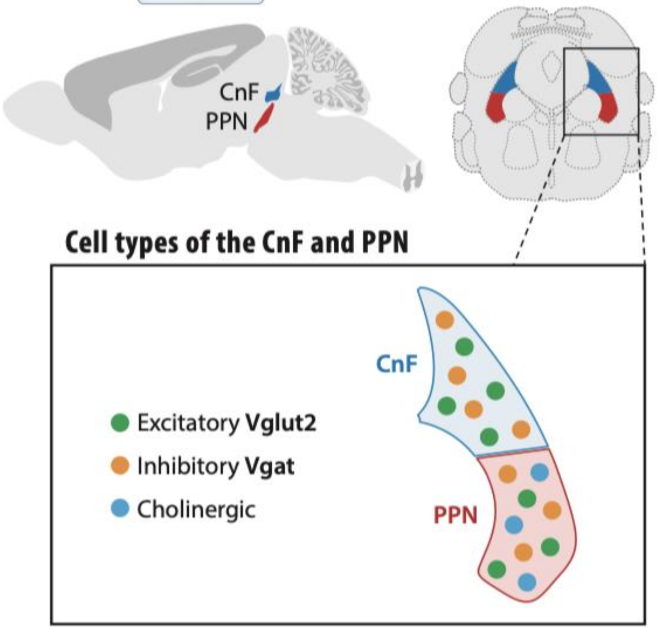

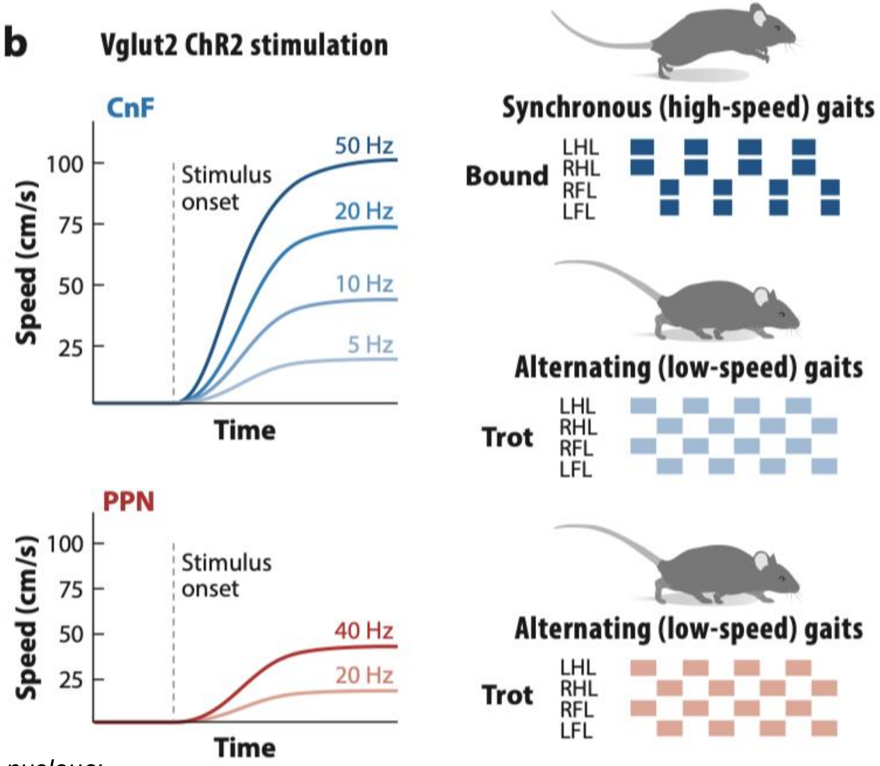

- *MLR, mesencephalic locomotor region; PPN, pedunculopontine nucleus; RFL, right forelimb; RHL, right hindlimb.*

Leiras, Cregg & Kiehn Annual Review of Neuroscience 2021

Brain stem circuits to control reaching & handling

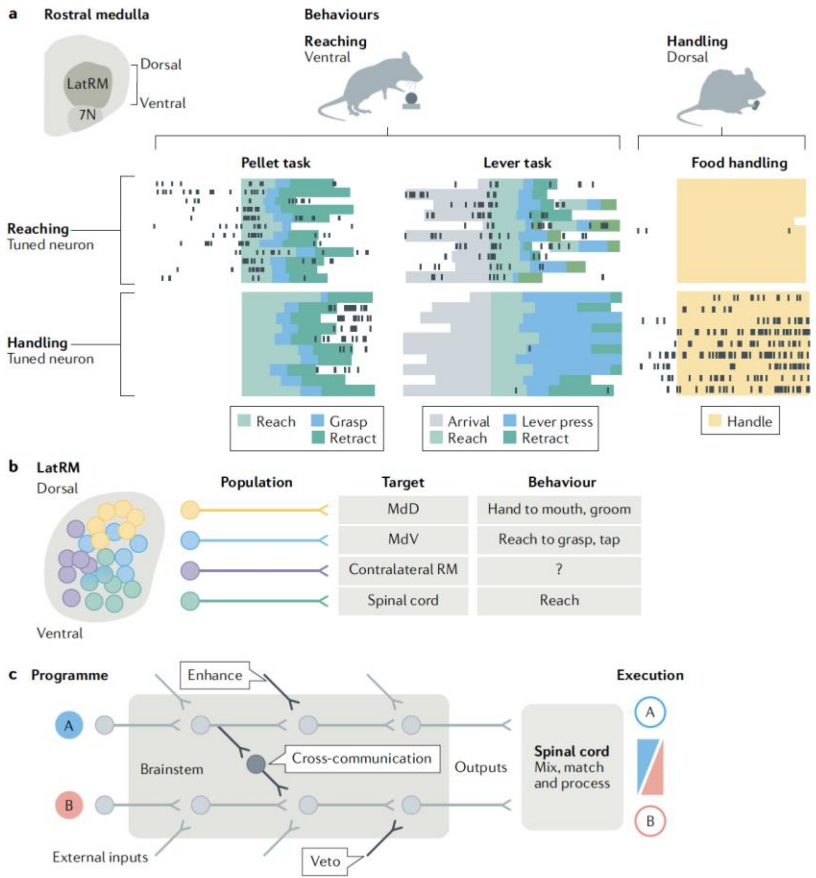

Arber & Costa Nature Review Neuroscience 2022

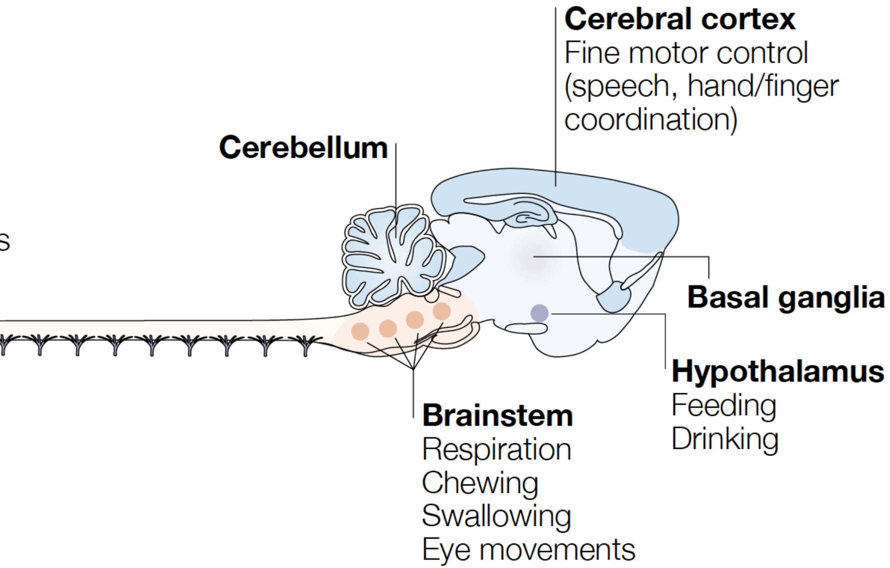

How are the many degrees of freedom tamed?

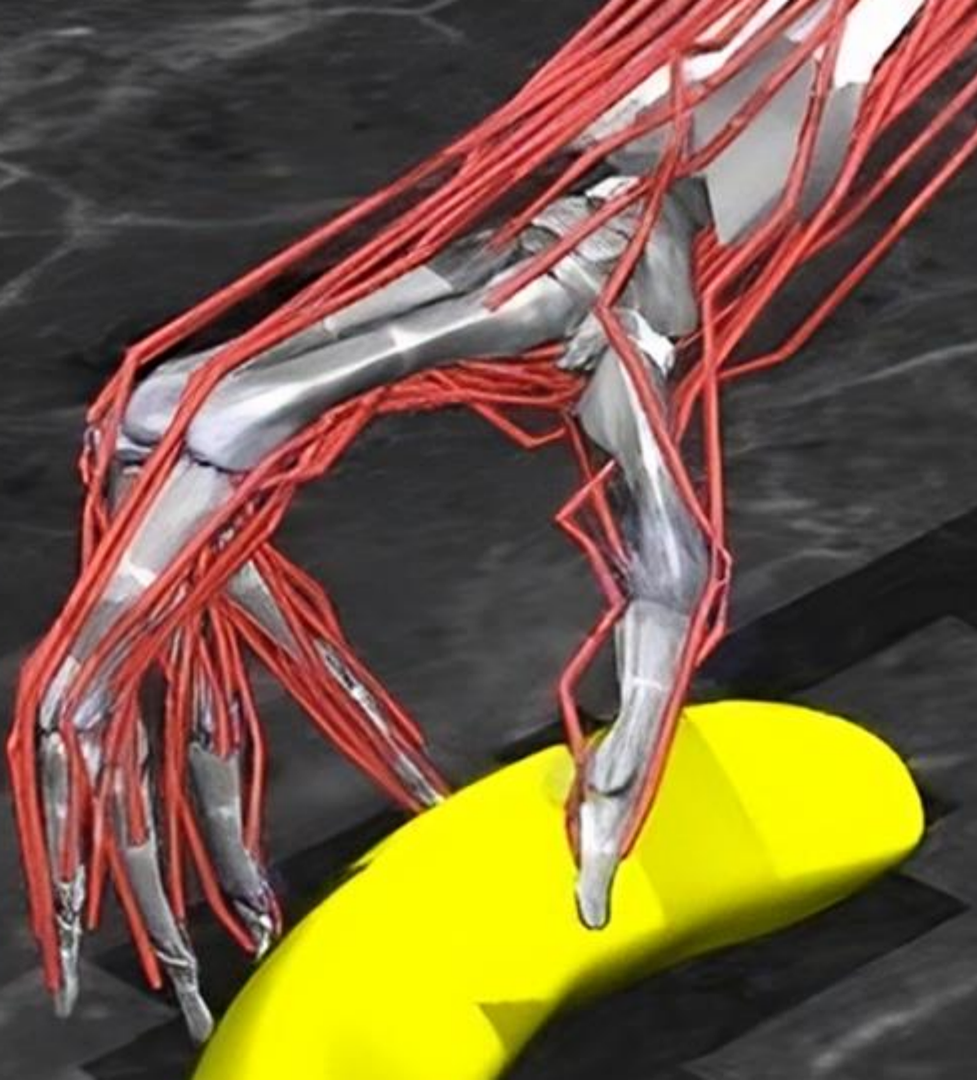

How many muscle states are there?

q^600

(for 600 muscles assuming q states per muscle)

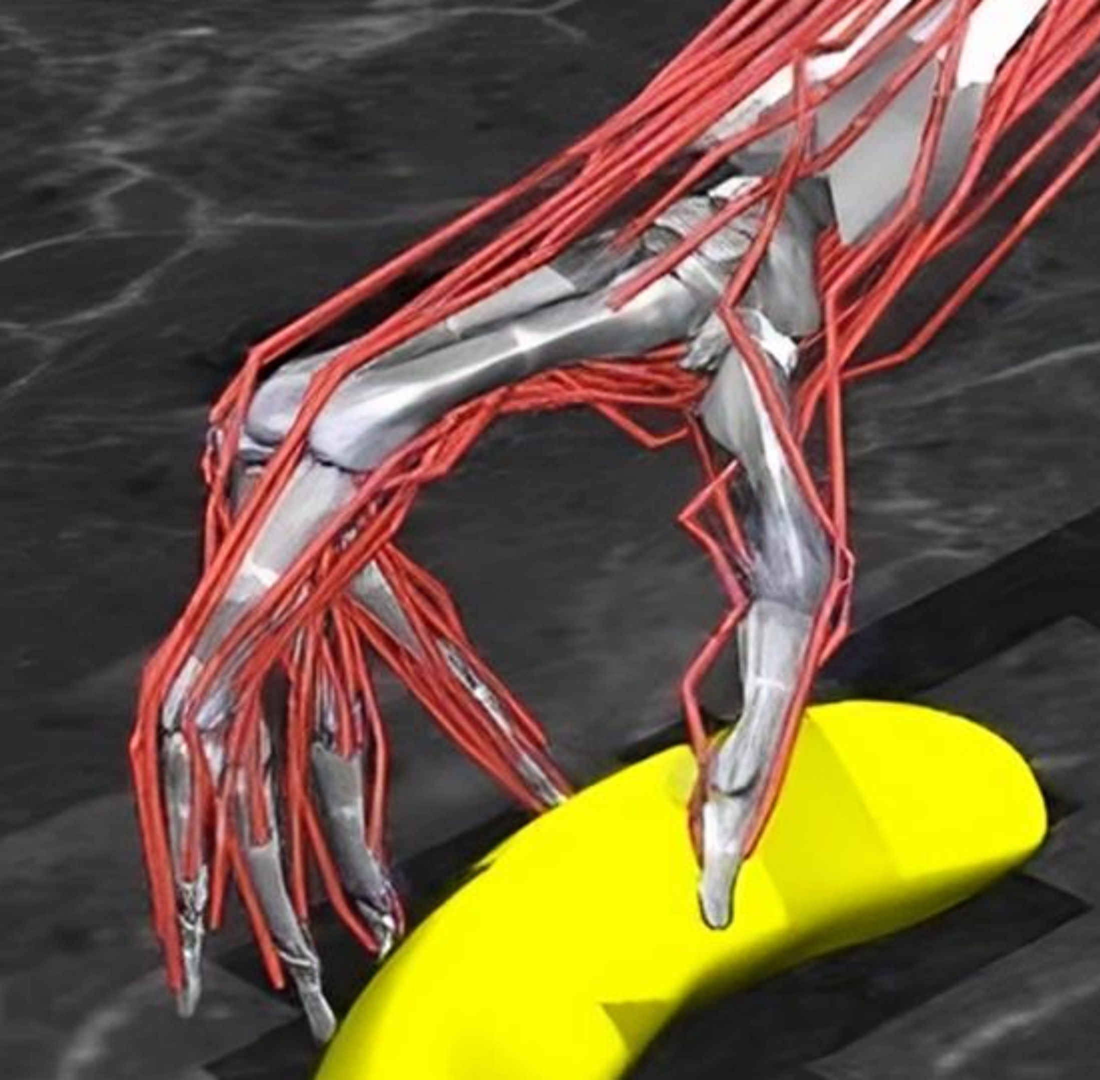

Similarly, there are many kinematic states

Boading balls: an example skill

Source: reddit

How do humans control the hand?

PC 1

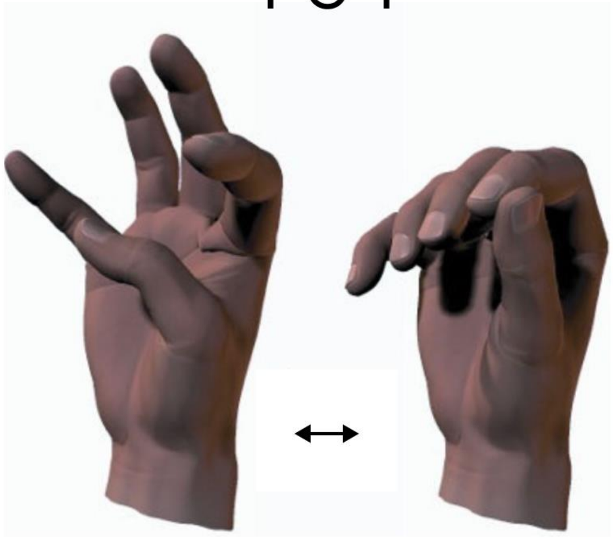

PC 2

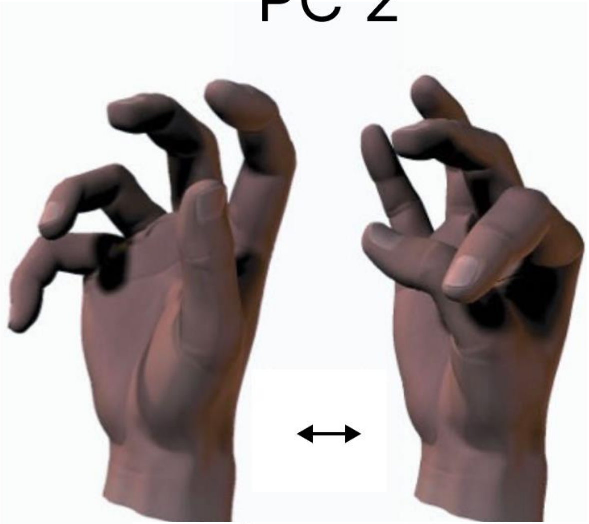

Human

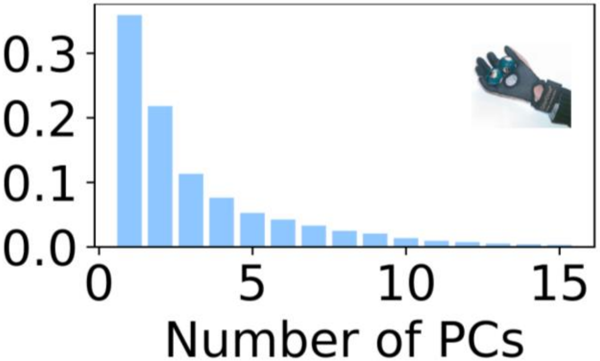

Classic result: Bernstein, Bizzi, D'Avella, …

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

Muscle synergies as principle for motor control

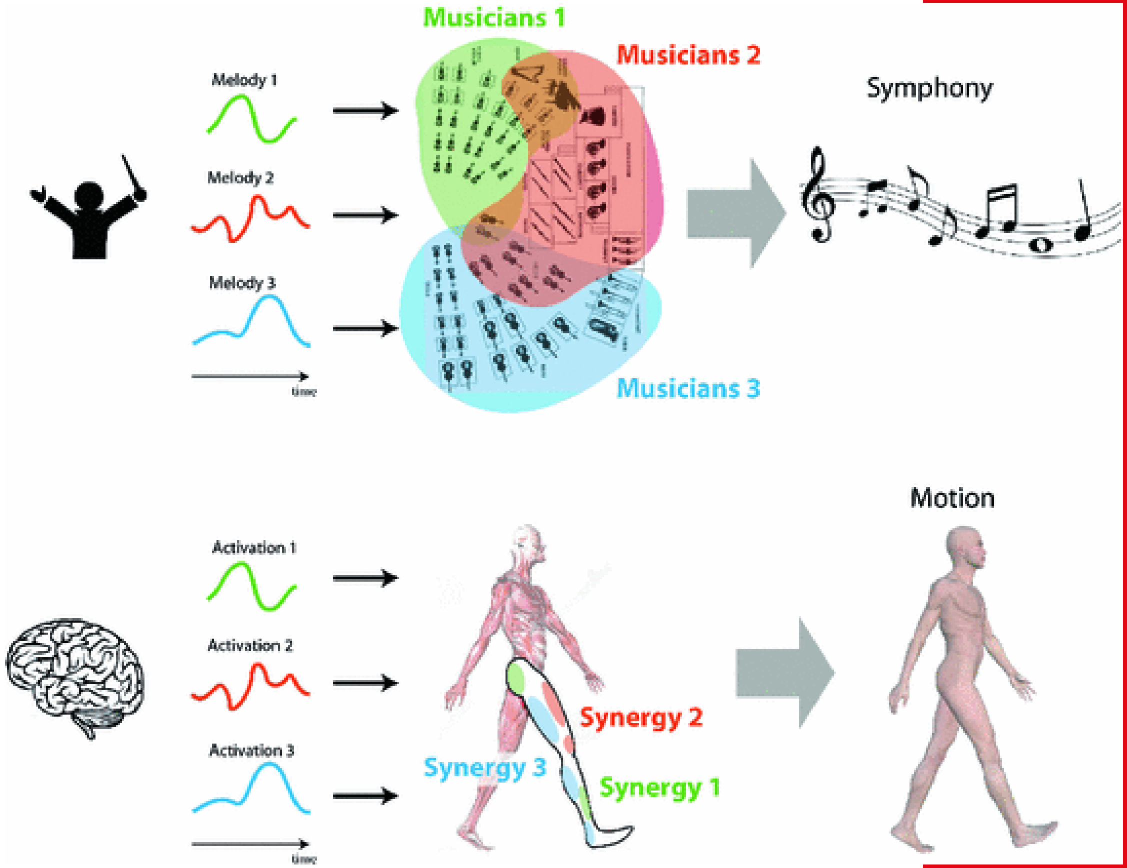

Torricelli et al., Emerging Therapies in Neurorehabilitation 2015

Integration of feedback?

From (open-loop) pattern generation to control theory

Note: Feedback is also present in spinal cord/brain stem examples!

MyoChallenge: Baoding Balls

Inaugural NeurIPS Challenge 2022

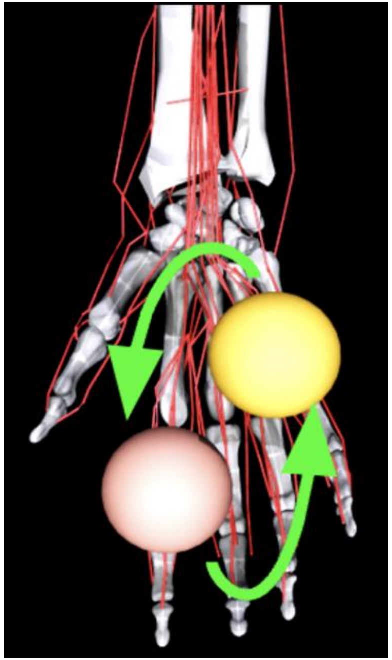

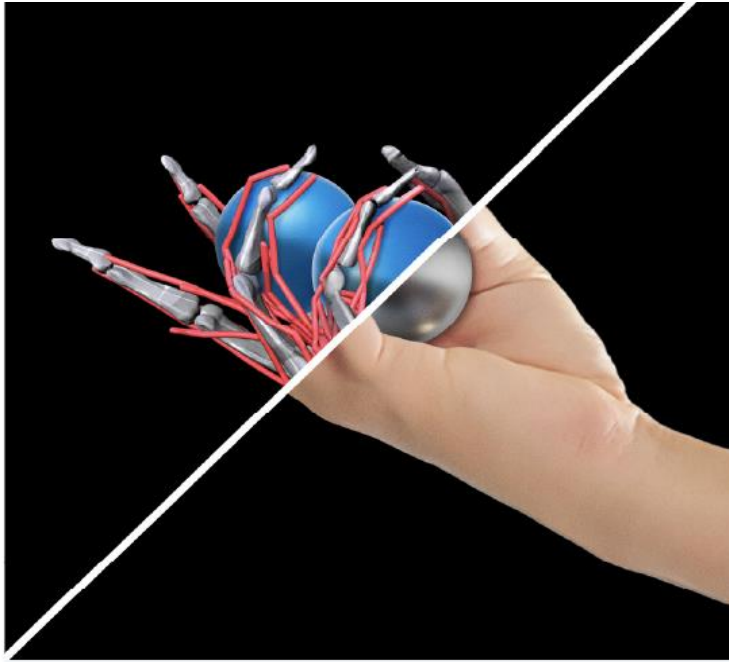

Caggiano et al. Proceedings of the NeurIPS 2022 Competitions Track, PMLR 220:233-250

Recent Successes in Reinforcement Learning (RL)

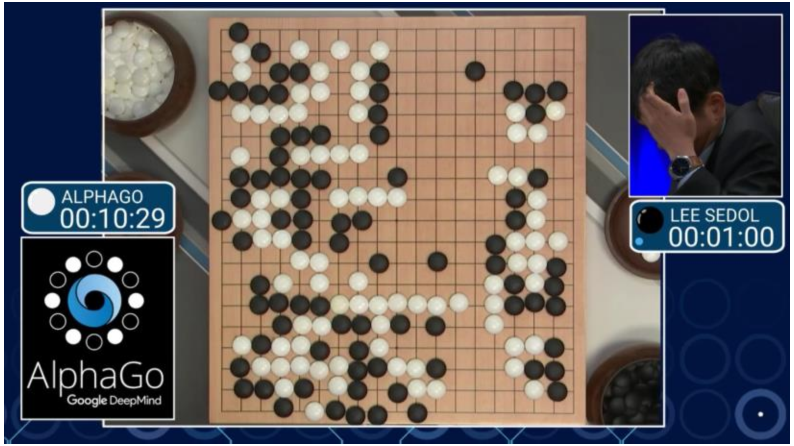

SCIENCE ROBOTICS | RESEARCH ARTICLE

ANIMAL ROBOTS

Learning quadrupedal locomotion over challenging terrain

Joonho Lee[1]*, Jemin Hwangbo[1,2], Lorenz Wellhausen[1], Vladlen Koltun[3], Marco Hutter[1]

Reinforcement learning is a natural framework for skill learning

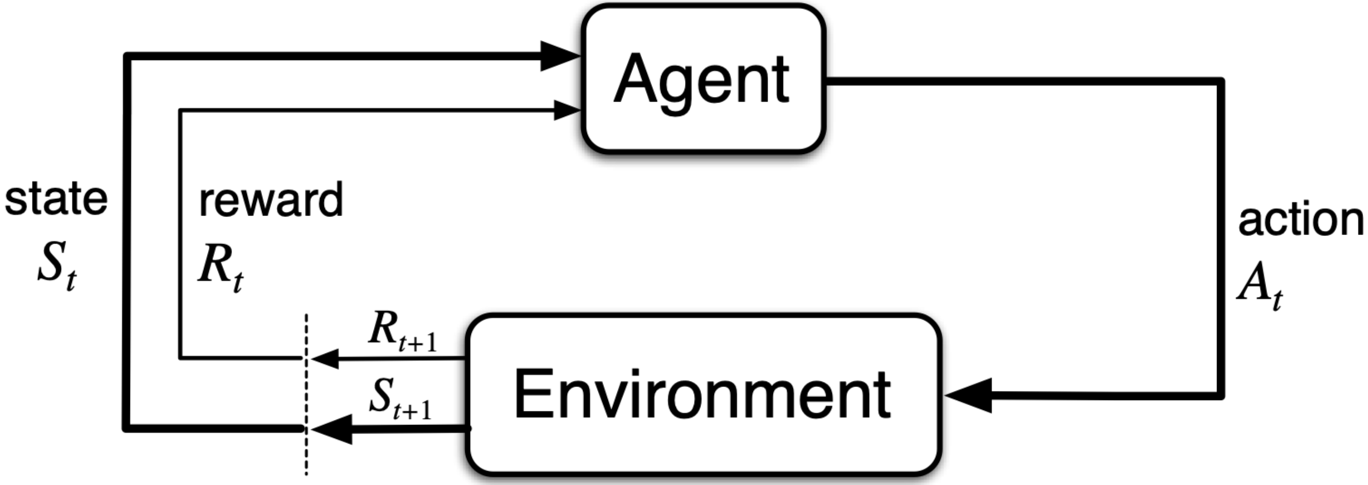

$$J(\theta) = v_{\pi_\theta}(s_0) = \mathbb{E}\left[\sum_{t=0}^{\infty} \gamma^t R_t \middle|\, S_0 = s_0\right]$$

Sutton & Barto, http://incompleteideas.net/book/RLbook2020.pdf

MyoChallenge: Baoding Balls

Inaugural NeurIPS Challenge 2022

q^39

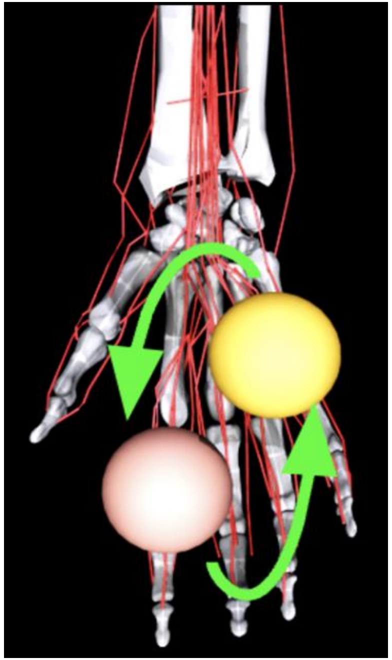

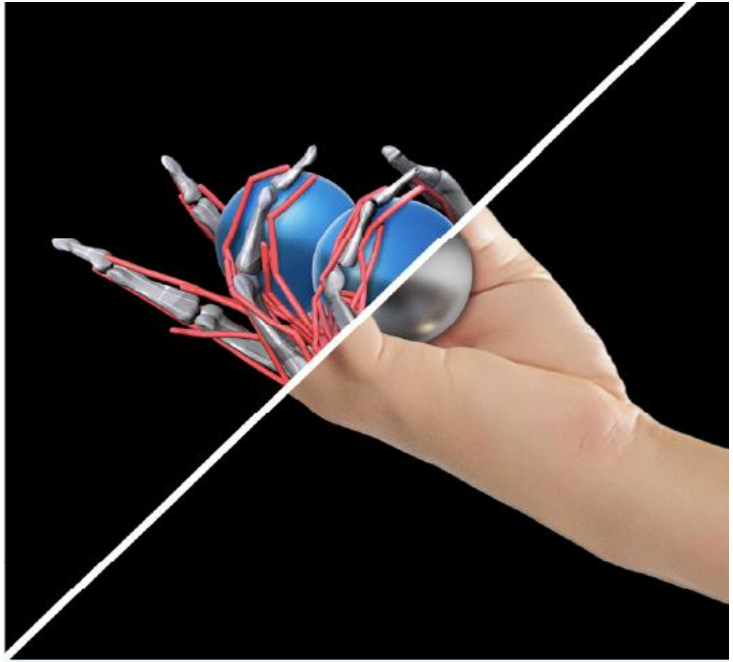

SOTA RL:

Phase 1	Phase 2
41%	0%

Caggiano et al. Proceedings of the NeurIPS 2022 Competitions Track, PMLR 220:233-250

Better exploration?

Learning sensorimotor skills

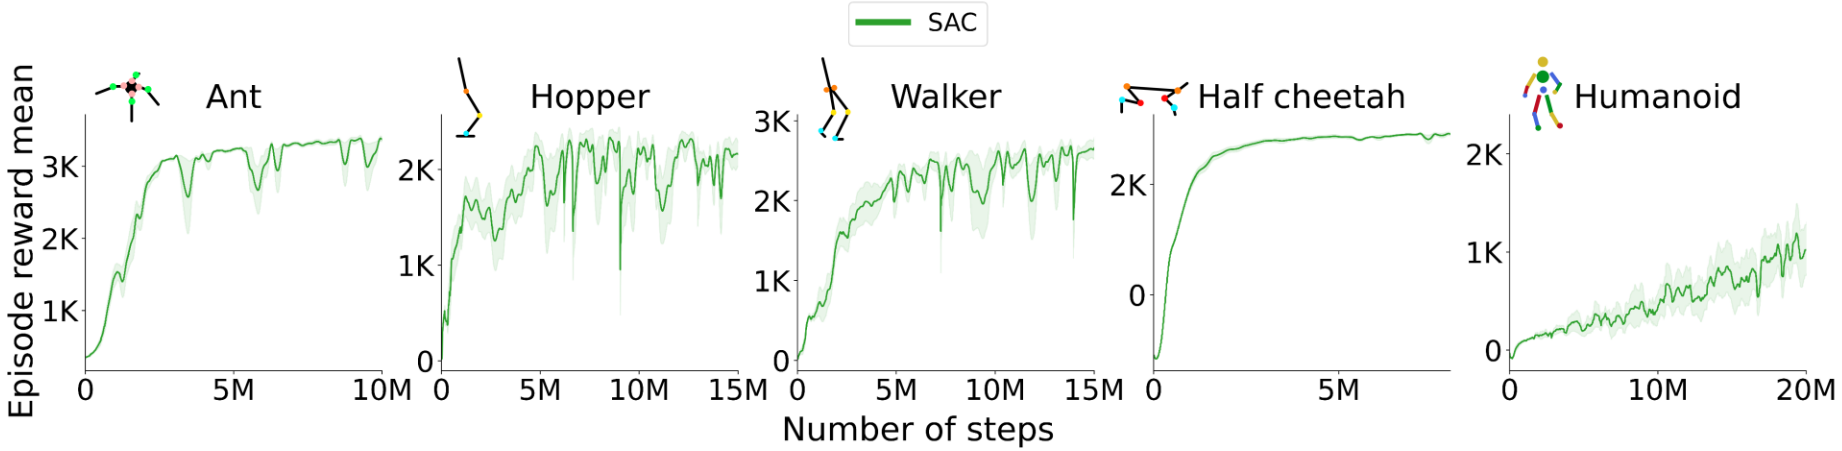

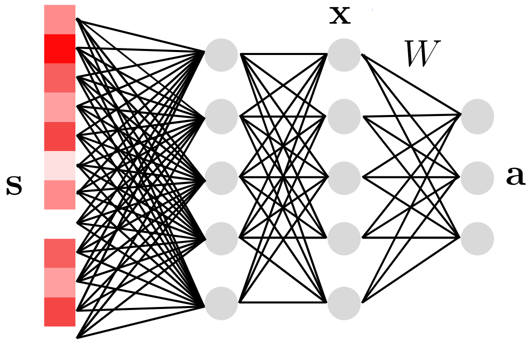

Default $\mathbf{a} = W\mathbf{x} + \epsilon$

Isotropic exploratory noise

Chiappa, Marin Vargas, Huang, Mathis, NeurIPS 2023

Basic intuition for better exploration

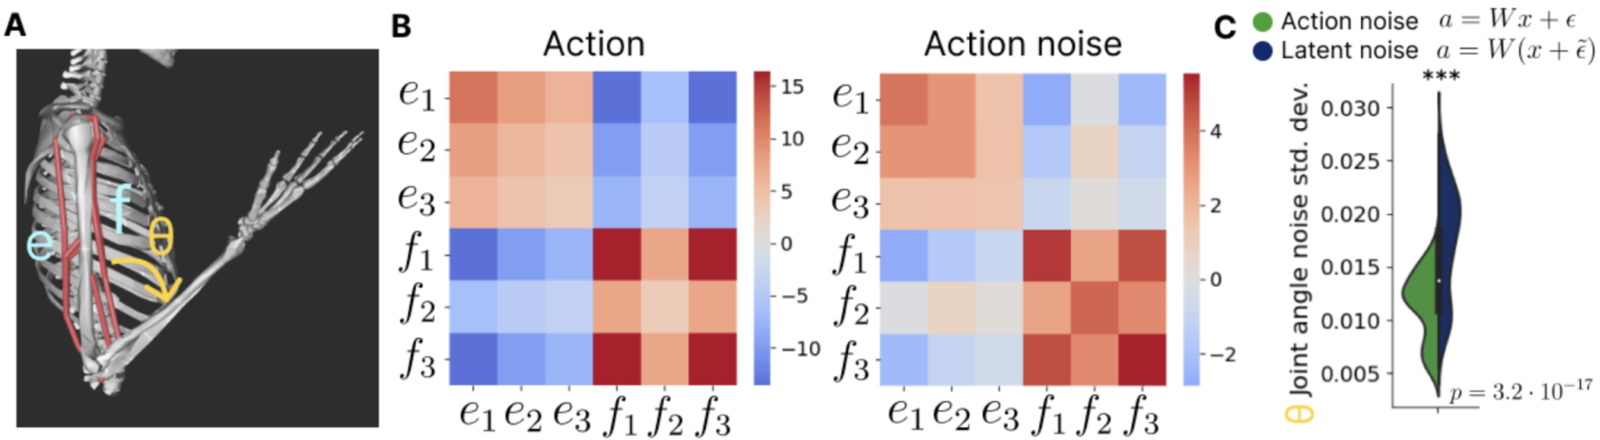

Chiappa, Marin Vargas, Huang, Mathis, NeurIPS 2023

Latent time-correlated exploration

LATTICE - **LAT**ent **TI**me-**C**orrelated **E**xploration

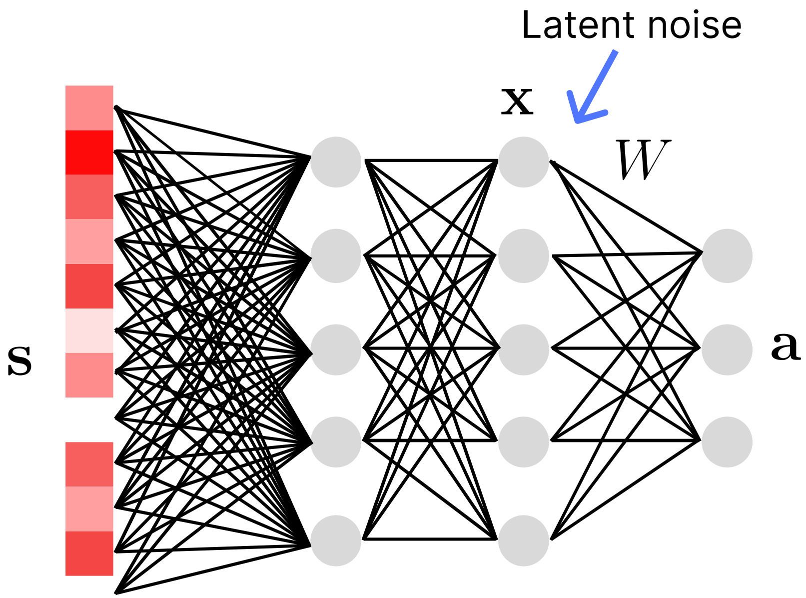

Perturbation matrices

$N_{\mathbf{a}} \times N_{\mathbf{x}}$: $P_{\mathbf{a}}$ $\quad (P_{\mathbf{a}})_{i,j} \sim \mathcal{N}\left(0, (S_{\mathbf{a}})_{i,j}\right)$

$N_{\mathbf{x}} \times N_{\mathbf{x}}$: $P_{\mathbf{x}}$ $\quad (P_{\mathbf{x}})_{i,j} \sim \mathcal{N}\left(0, (S_{\mathbf{x}})_{i,j}\right)$

LATTICE $\mathbf{a} = (W + P_{\mathbf{a}} + W P_{\mathbf{x}})\mathbf{x}$

State-Dependent Exploration

gSDE $\mathbf{a} = (W + P_{\mathbf{a}})\mathbf{x}$

Default $\mathbf{a} = W\mathbf{x} + \epsilon$

Time

Time + Action

Chiappa, Marin Vargas, Huang, Mathis, NeurIPS 2023

Benchmarking learning to locomote

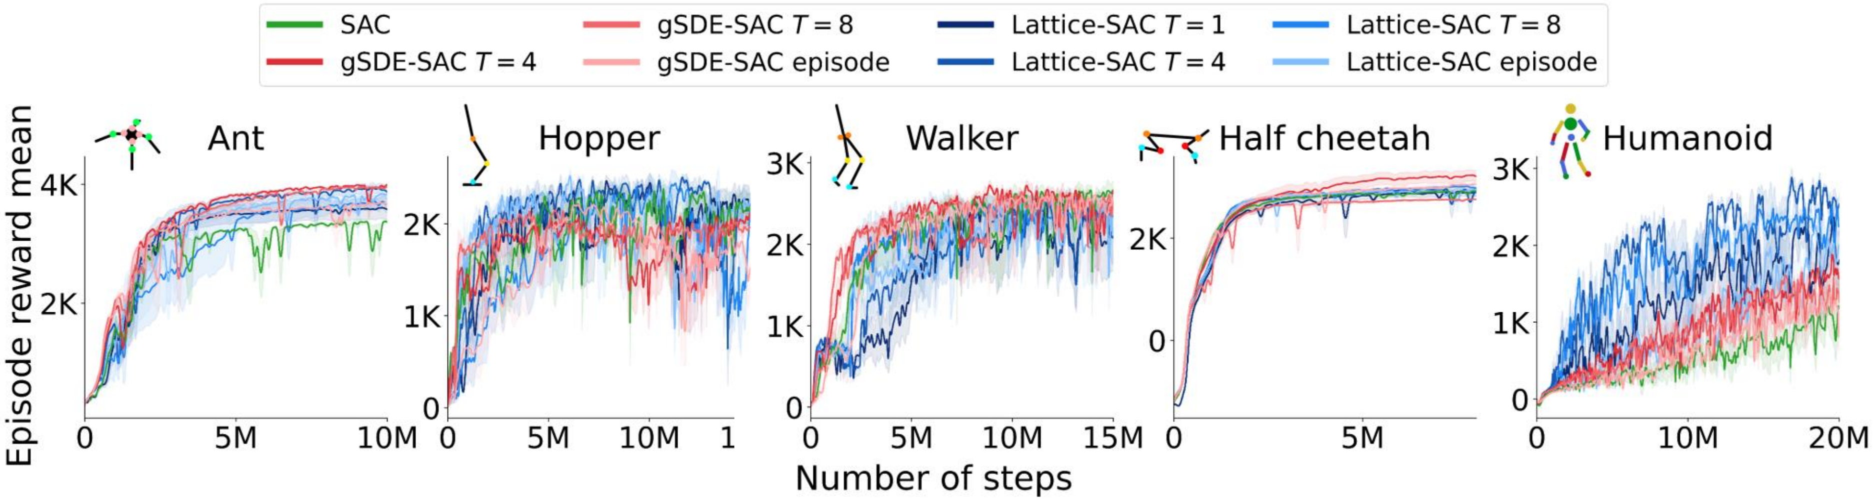

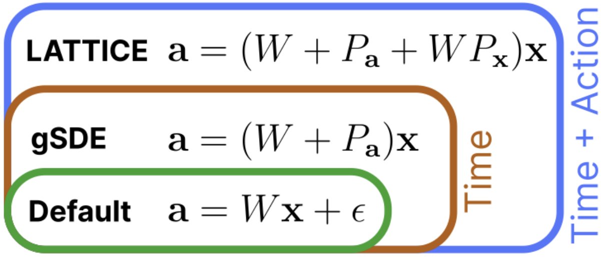

Chiappa, Marin Vargas, Huang, Mathis, NeurIPS 2023

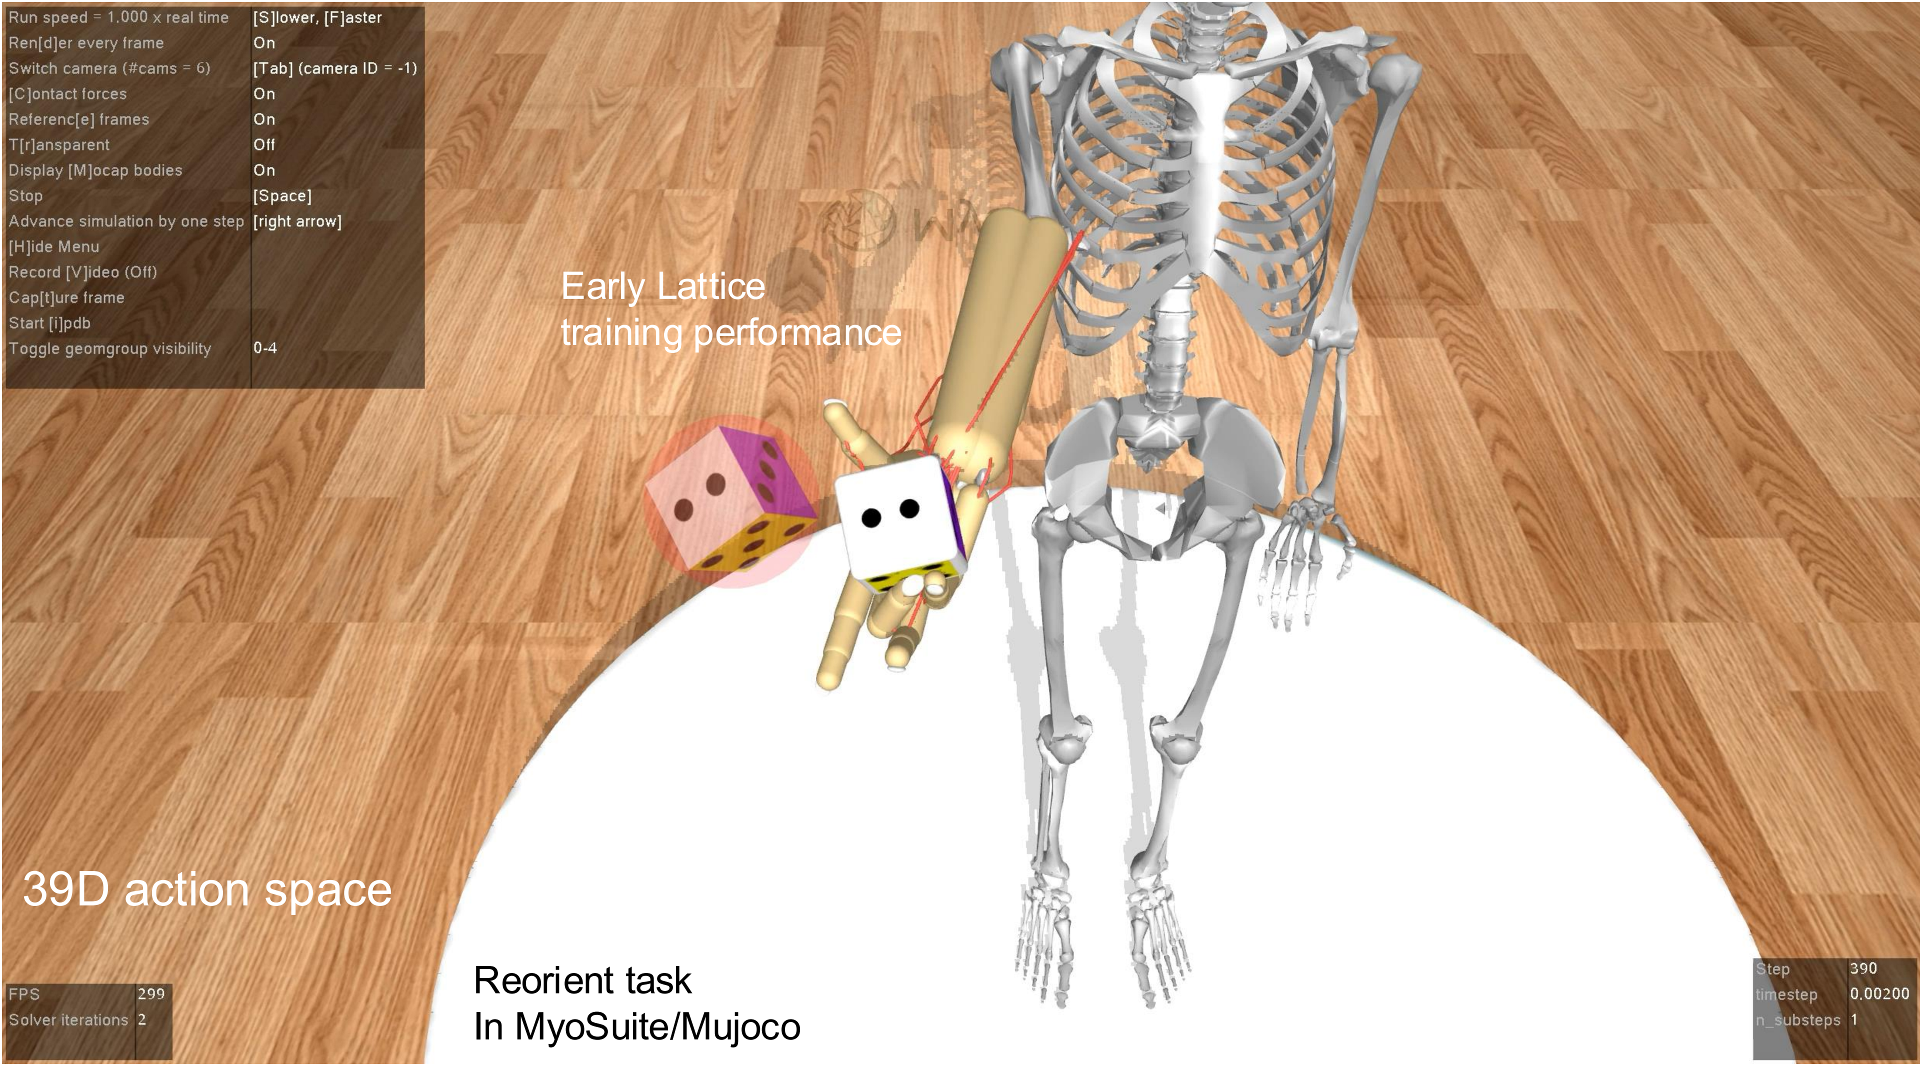
Run speed = 1.000 x real time [S]lower, [F]aster
Ren[d]er every frame On
Switch camera (#cams = 6) [Tab] (camera ID = -1)
[C]ontact forces On
Referenc[e] frames On
T[r]ansparent Off
Display [M]ocap bodies On
Stop [Space]
Advance simulation by one step [right arrow]
[H]ide Menu
Record [V]ideo (Off)
Cap[t]ure frame
Start [i]pdb
Toggle geomgroup visibility 0-4
Early Lattice
training performance
39D action space
Reorient task
In MyoSuite/Mujoco
FPS 299
Solver iterations 2
Step 390
timestep 0.00200
n_substeps 1

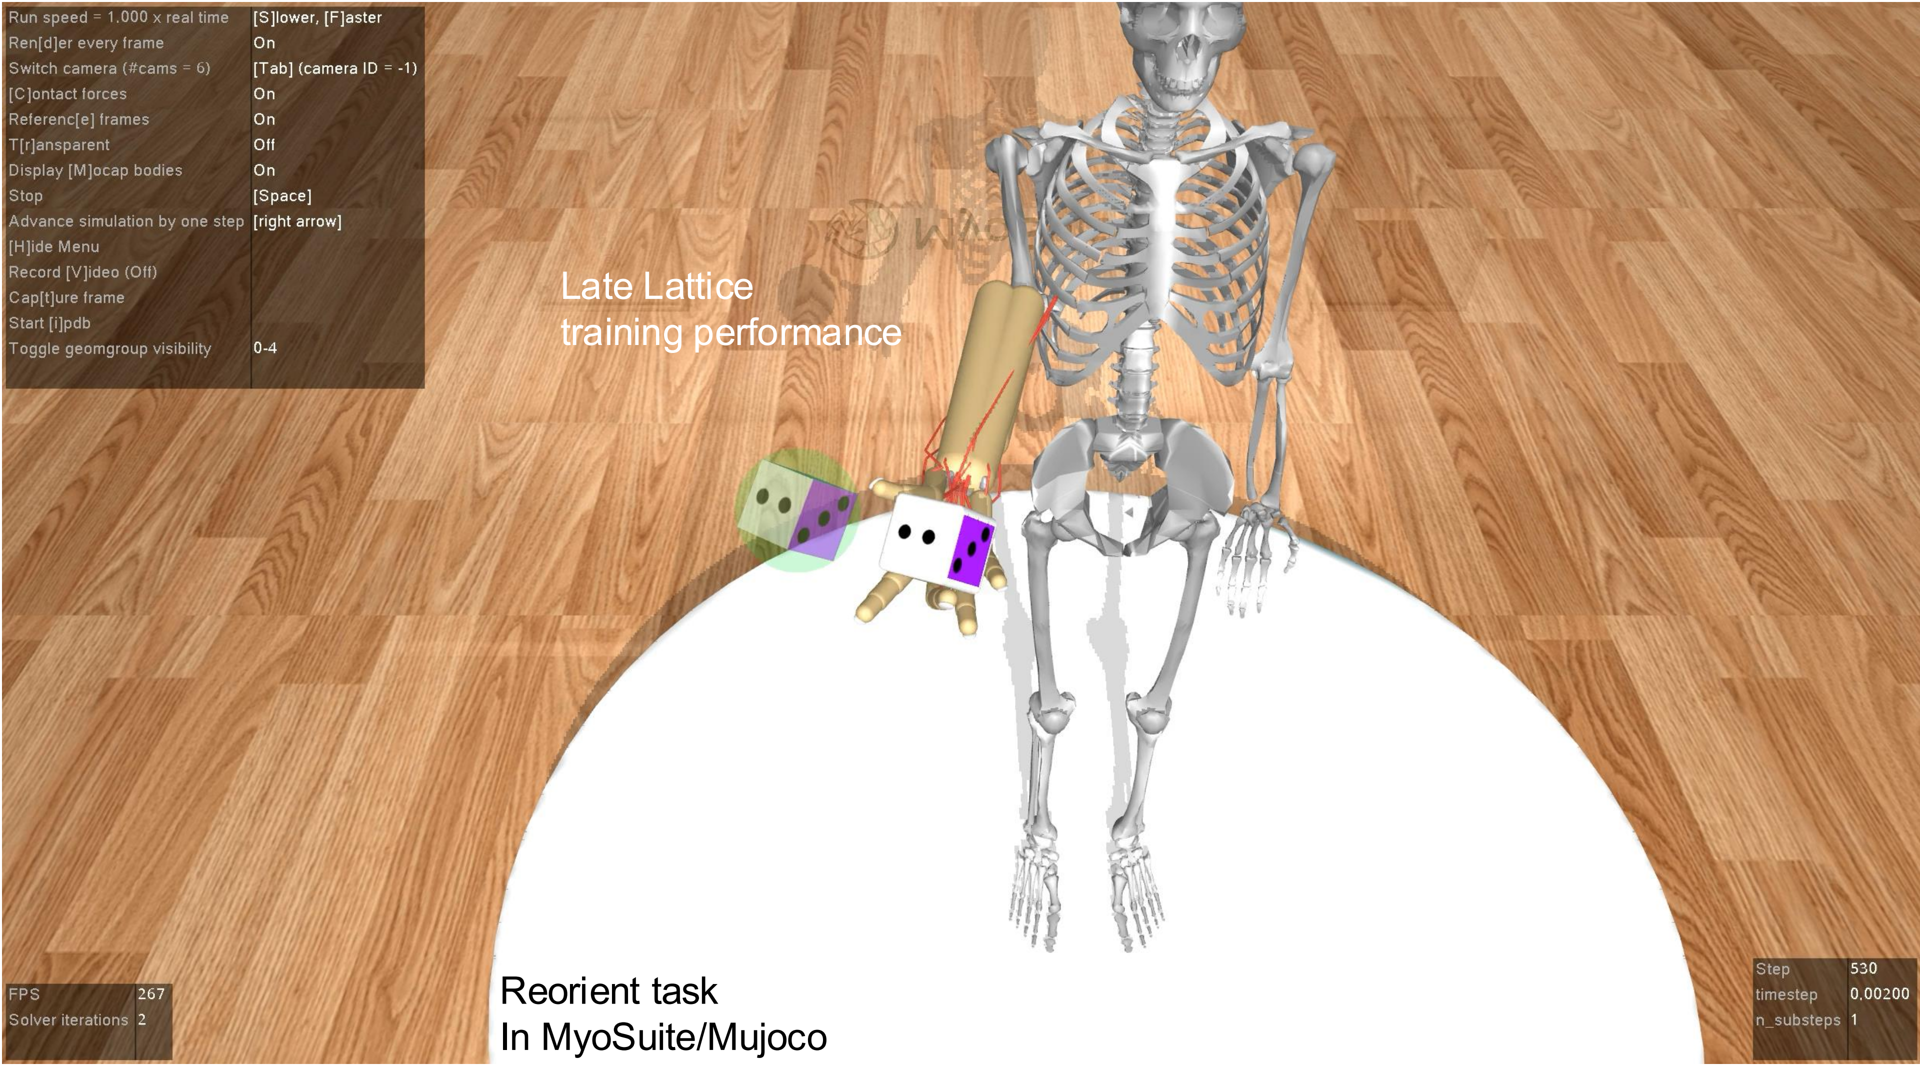

Run speed = 1.000 x real time
[S]lower, [F]aster
Ren[d]er every frame
On
Switch camera (#cams = 6)
[Tab] (camera ID = -1)
[C]ontact forces
On
Referenc[e] frames
On
T[r]ansparent
Off
Display [M]ocap bodies
On
Stop
[Space]
Advance simulation by one step
[right arrow]
[H]ide Menu
Record [V]ideo (Off)
Cap[t]ure frame
Start [i]pdb
Toggle geomgroup visibility
0-4
Late Lattice
training performance
FPS
267
Solver iterations
2
Reorient task
In MyoSuite/Mujoco
Step
530
timestep
0.00200
n_substeps
1

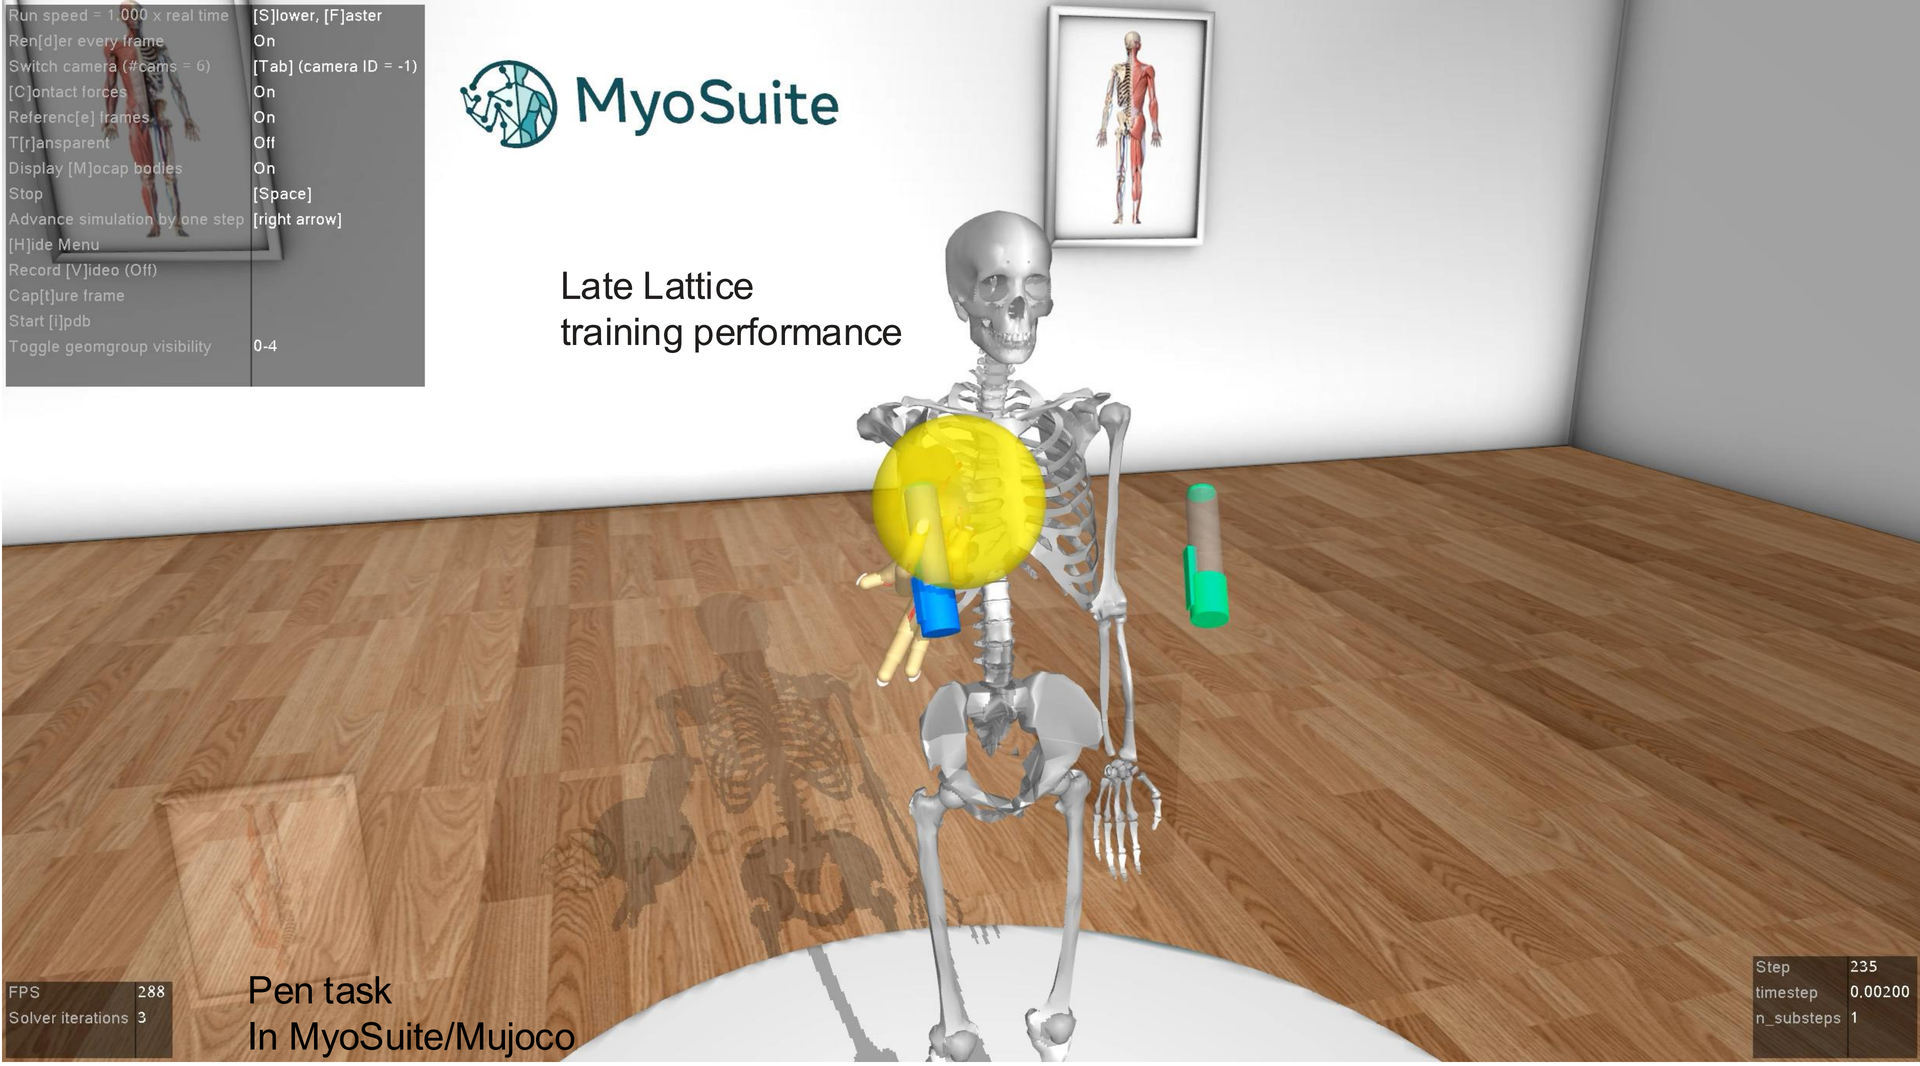
Run speed = 1.000 x real time
[S]lower, [F]aster
Ren[d]er every frame
On
Switch camera (#cams = 6)
[Tab] (camera ID = -1)
[C]ontact forces
On
Referenc[e] frames
On
T[r]ansparent
Off
Display [M]ocap bodies
On
Stop
[Space]
Advance simulation by one step
[right arrow]
[H]ide Menu
Record [V]ideo (Off)
Cap[t]ure frame
Start [i]pdb
Toggle geomgroup visibility
0-4
MyoSuite
Late Lattice
training performance
FPS
288
Solver iterations
3
Pen task
In MyoSuite/Mujoco
Step
235
timestep
0.00200
n_substeps
1

Object manipulation learning curves

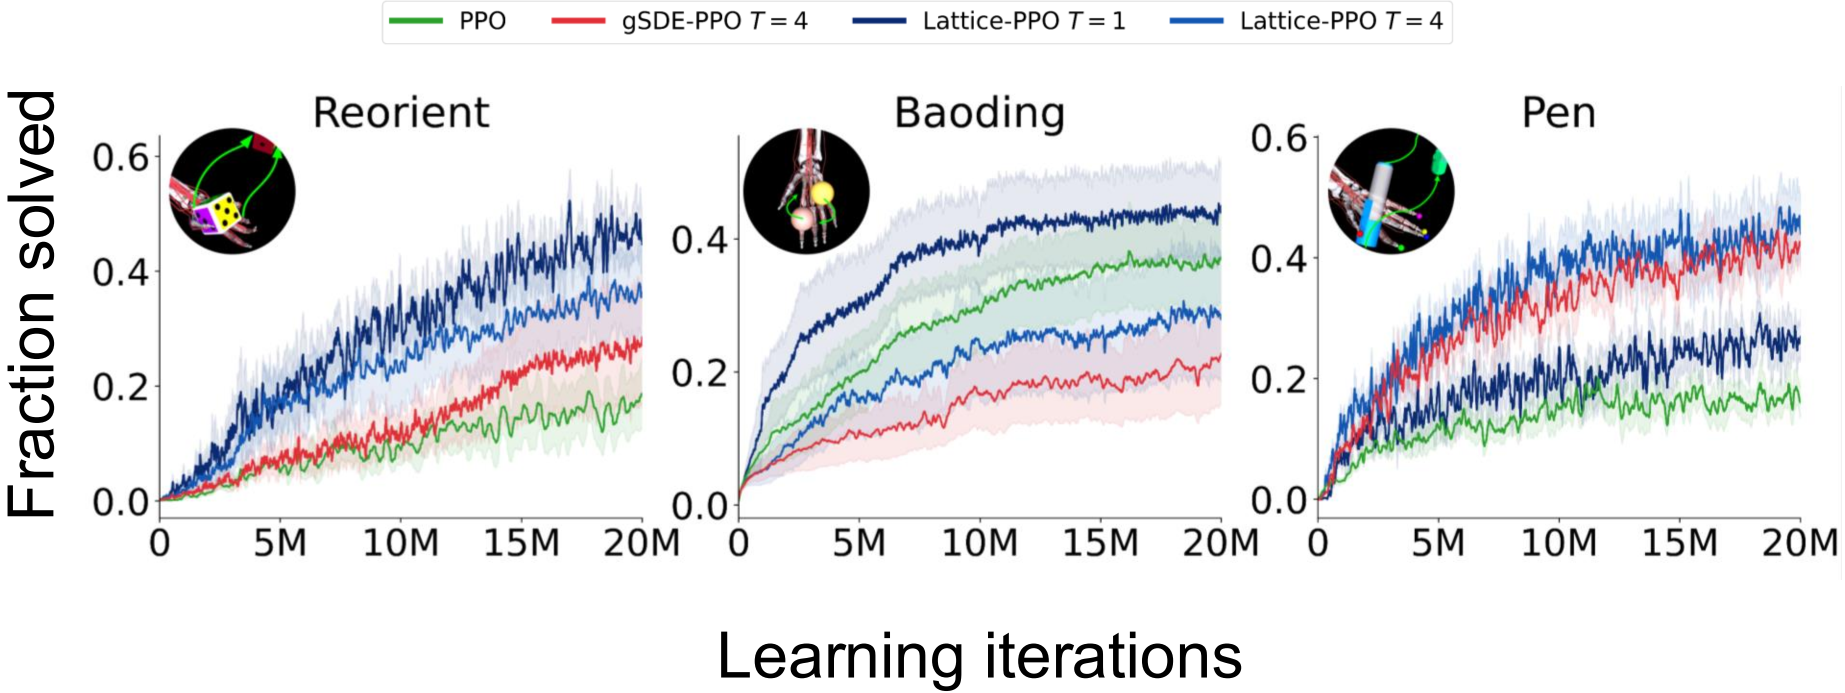

Chiappa, Marin Vargas, Huang, Mathis, NeurIPS 2023

Lattice learns more energy efficient solutions

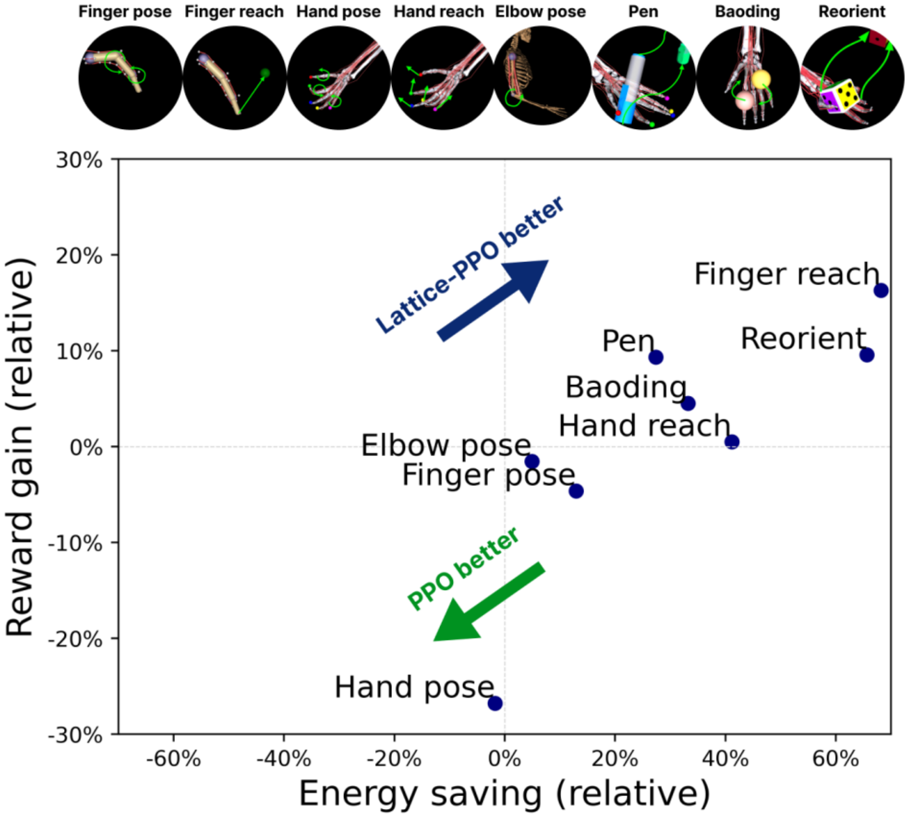

Chiappa, Marin Vargas, Huang, Mathis, NeurIPS 2023

Can we model complex behavior using high-dimensional musculoskeletal body?

A. Marin Vargas

MyoChallenge'23

Towards Human-Level Dexterity and Agility

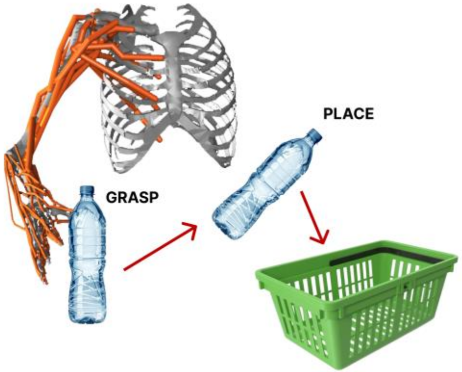

Phase 2			
Rank	**Team**	**Score**	**Effort**
1	**Lattice**	0.34343	0.05220
2	GaitNet	0.30303	0.05506
3	CarbonSiliconAI	0.20202	0.07620
4	NUABILITY	0.14141	0.07492
5	NUR-TEAM	0.10101	0.04470

NEURAL INFORMATION PROCESSING SYSTEMS

Winning solution – Lattice Team:

- Best score
- Best energy efficiency

Chiappa A. S., Marin Vargas A., Huang A., & Mathis A., NeurIPS (2023)

Caggiano V., Durandau G., Wang C., Kwang Tan C., Schumacher P., Wang H., Chiappa A., Marin Vargas A., Mathis A., Park J., Won J., Park G., Shin B., Kim M., Koo S., Yang Z., Dang W., Cai H., Song J., Song S., Sartori M., Kumar V., ArXiv (2024)

Despite Lattice, the performance is not so high…

How can one reach a higher fraction?

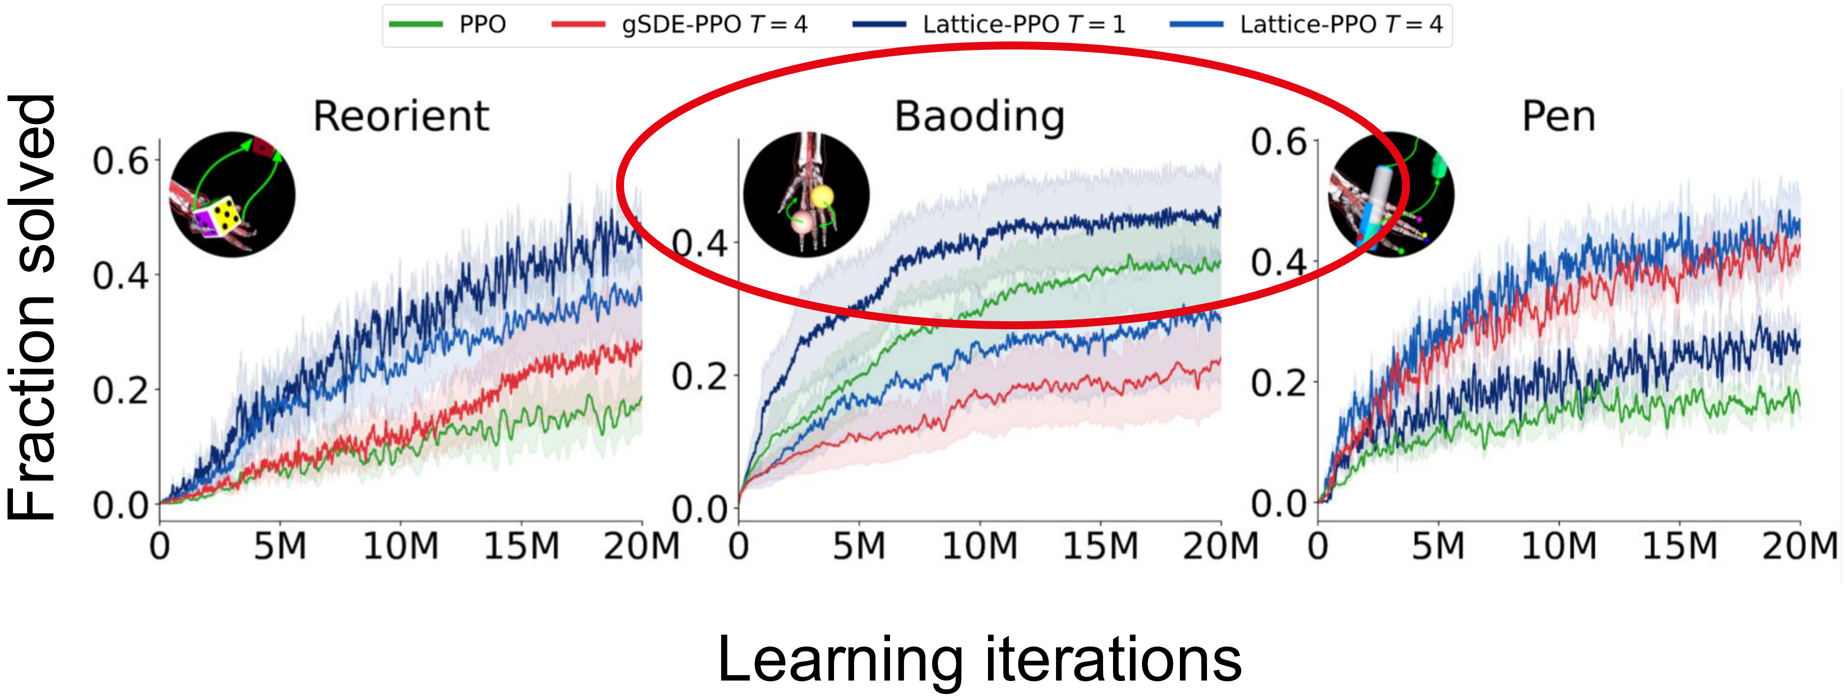

Chiappa, Marin Vargas, Huang, Mathis, NeurIPS 2023

Curriculum learning

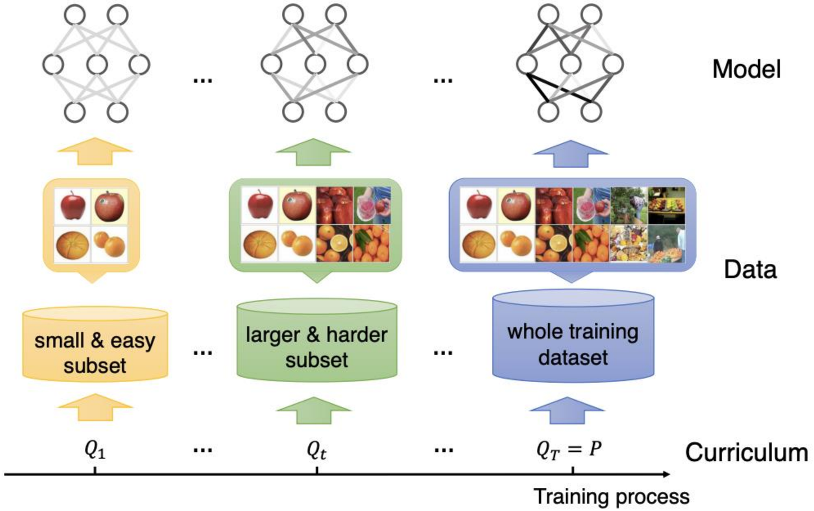

Wang, Chen, Zhu IEEE TPAMI, 2021

EPFL

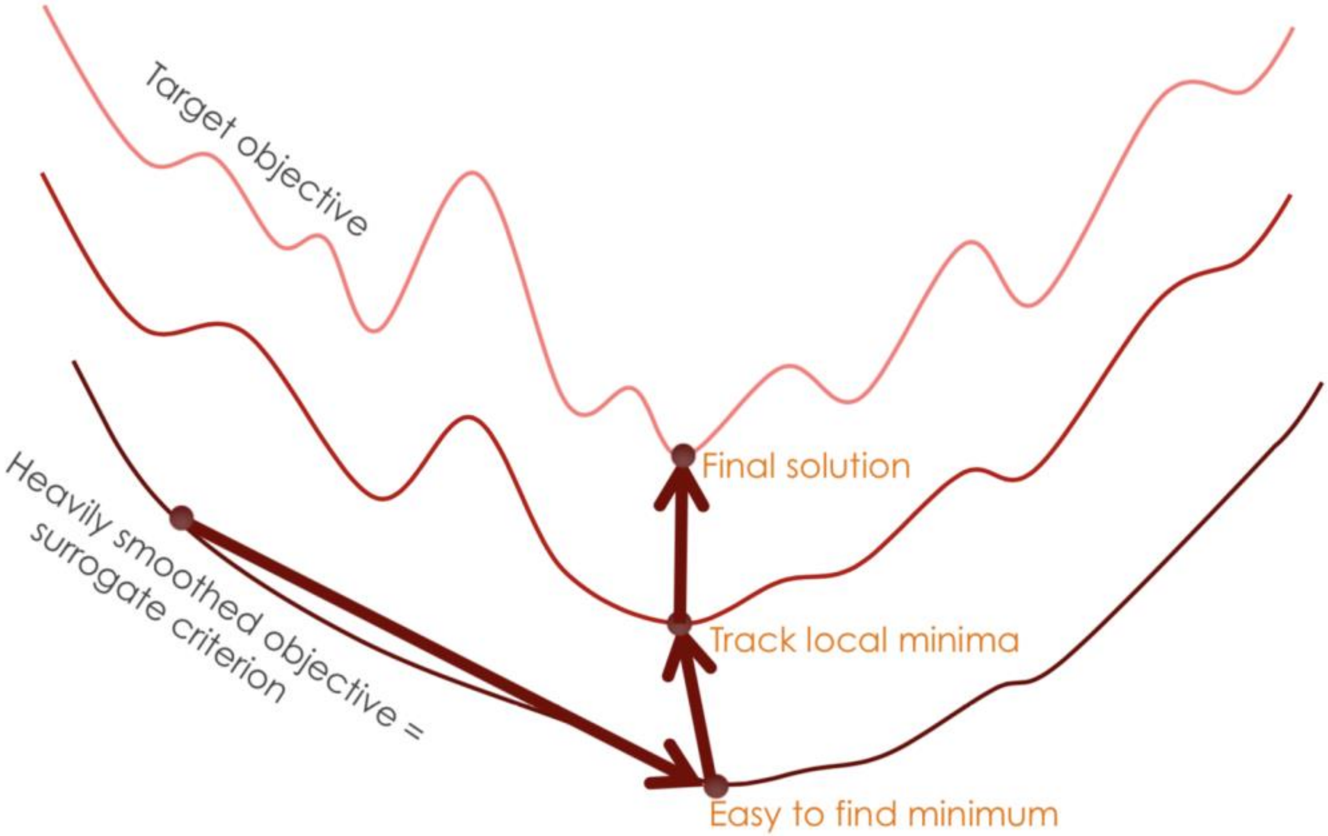

Wang, Chen, Zhu IEEE TPAMI, 2021

Curriculum learning in biology

Teaching in Wild Meerkats

Alex Thornton* and Katherine McAuliffe

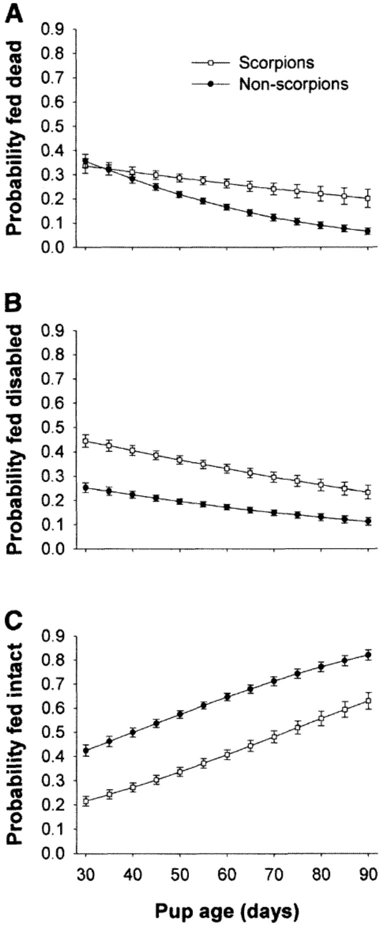

Science, 2006

Inspiration from coaching: part-to-whole practice

States of a dynamic skill

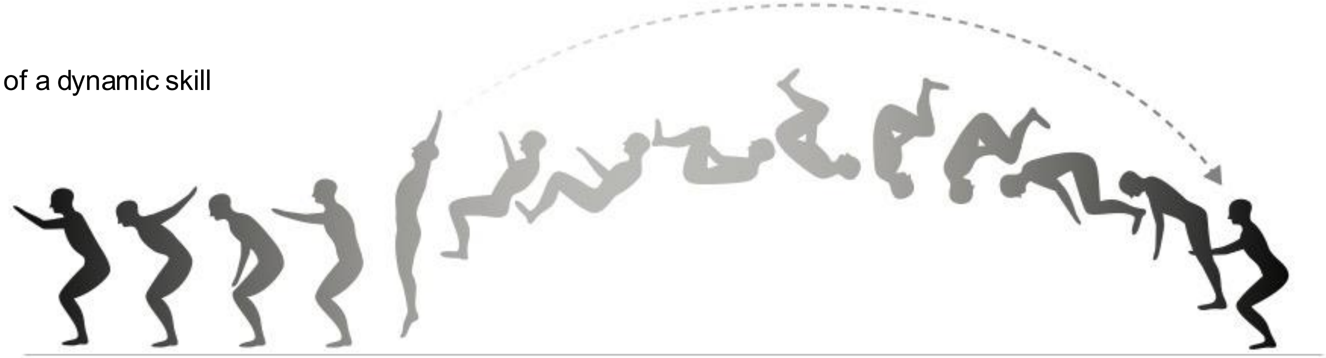

Learning static movement motifs first

Recommended strategy:

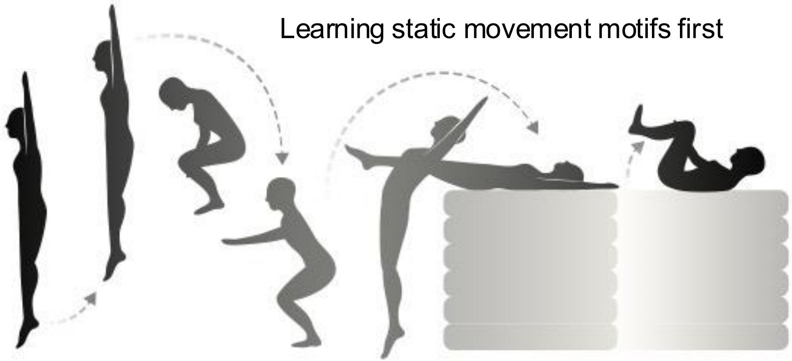

Chiappa, A. S., Tano*, P., Patel*, N., Ingster, A., Pouget, A., & Mathis, A. Neuron 2024*

Curriculum learning

- Static to Dynamic Stability (SDS)
  - SDS creates stability at desired states *before* learning a policy that reaches them
  - A curriculum gradually transforms static stability into dynamic movement motifs

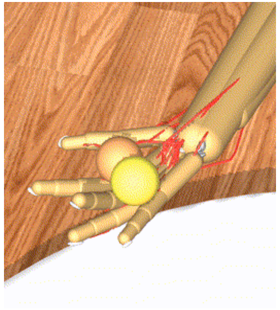

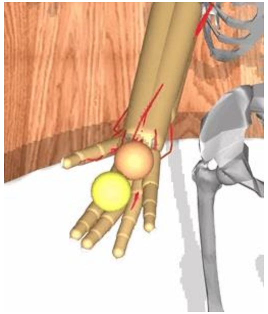

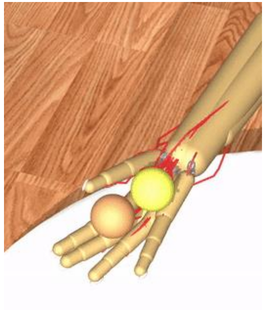

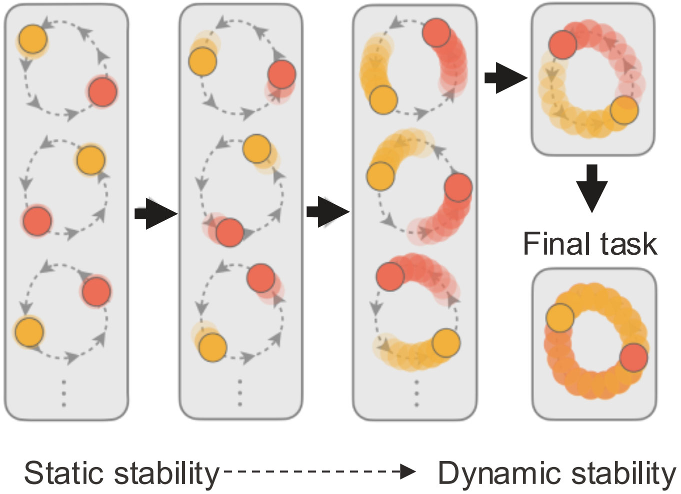

- *Chiappa*, A. S., Tano*, P., Patel*, N., Ingster, A., Pouget, A., & Mathis, A. Neuron 2024*
  Caggiano et al. Proceedings of the NeurIPS 2022 Competitions Track, PMLR 220:233-250

Learning curve for our policy

Sport science terminology: Part to whole practice Deliberate practice

ML terminology:

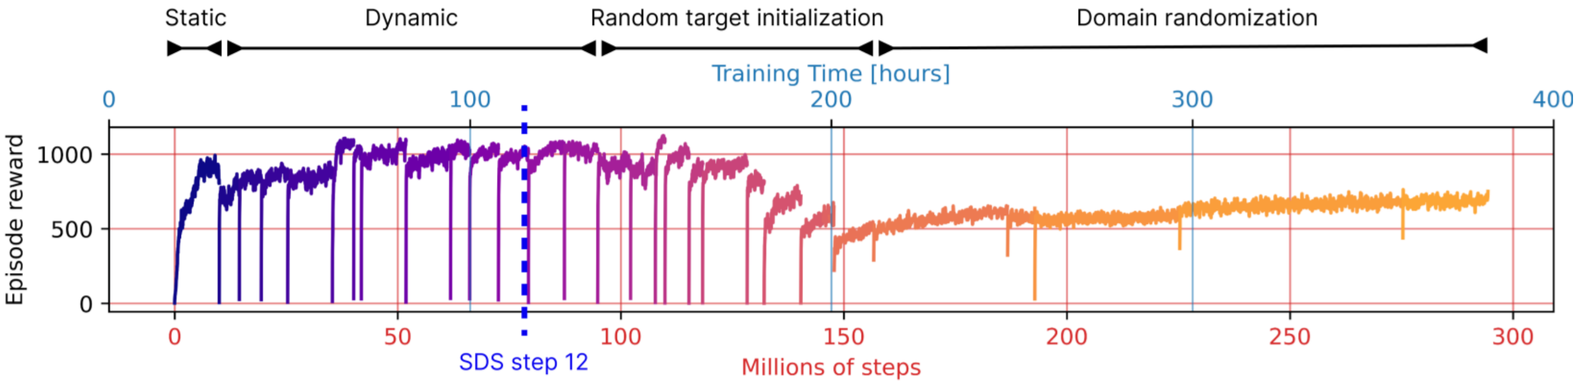

EPFL

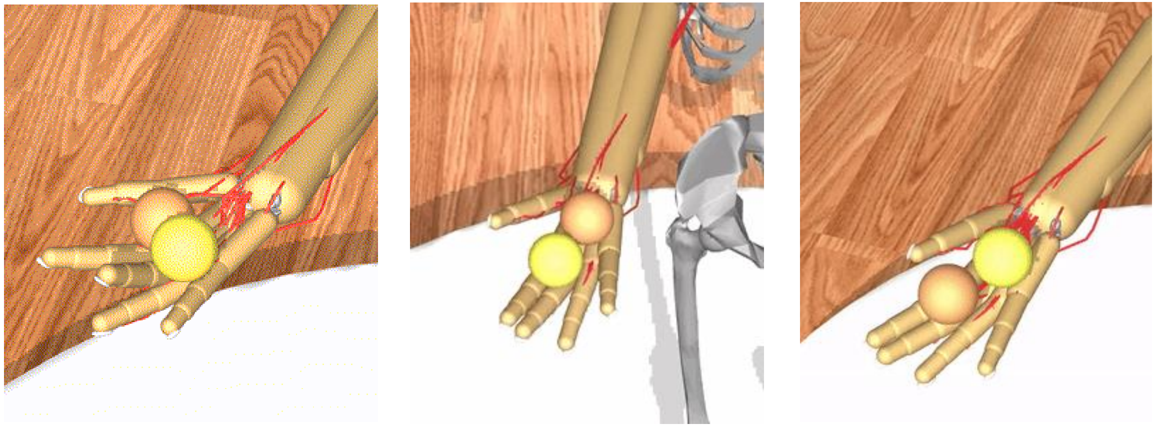

Curriculum	Phase 1	Phase 2
None	41%	0%
Location only	42%	4%
Speed only	45%	0%
SDS (ours)	100%	55%

Team	Performance
SDS (ours)	55%
AI4Muscles	41%
IARAI-JKU	15%
pkumarl	14%

Chiappa, A. S., Tano*, P., Patel*, N., Ingster, A., Pouget, A., & Mathis, A. Neuron 2024*
Caggiano et al. Proceedings of the NeurIPS 2022 Competitions Track, PMLR 220:233-250

How do humans achieve this task?

PC 1

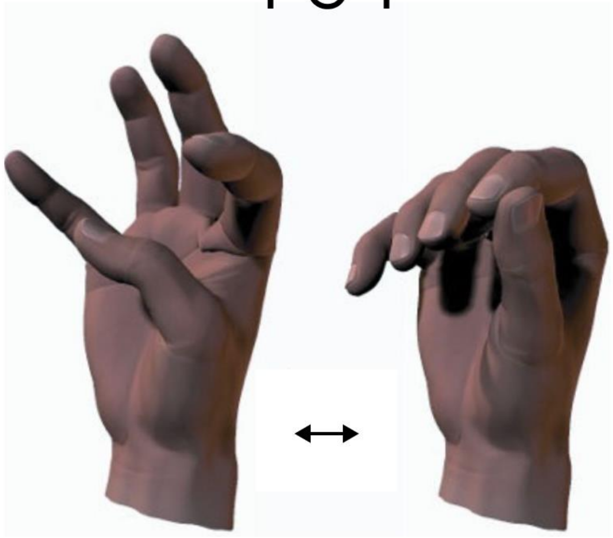

PC 2

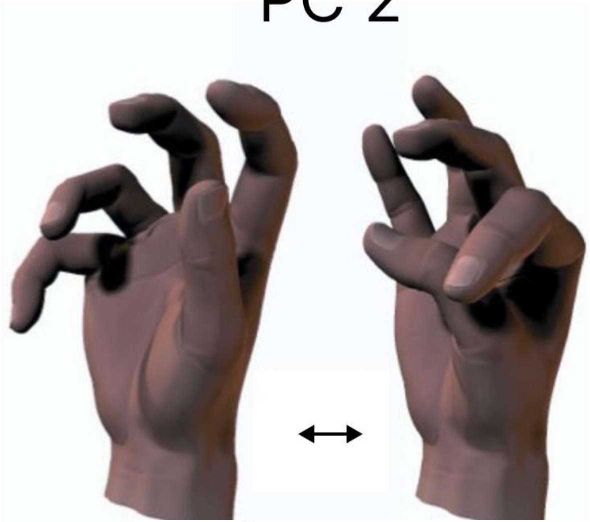

Human

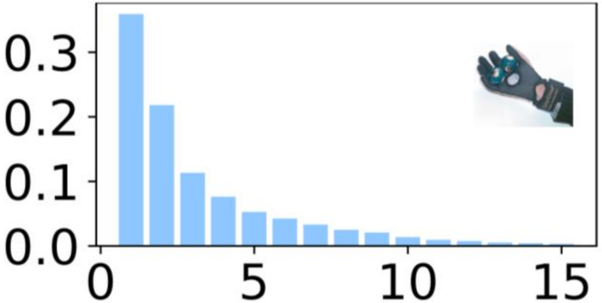

Todorov, & Ghahramani,
Annual International Conference of the IEEE Engineering in Medicine and Biology 2004

SDS also discovers a low-dimensional control space

SDS

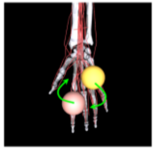

Human

		Position	Muscle act.
SDS	Baoding	4.5	12
	Control	8	7
Human	Baoding	5	
	Control	8.5	

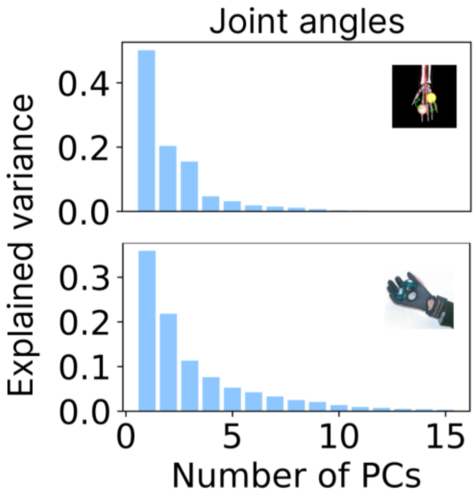

This notion of muscle/kinematic synergy is purely based on reconstruction error!

Physics engine allows causal experiments with “muscle synergies”

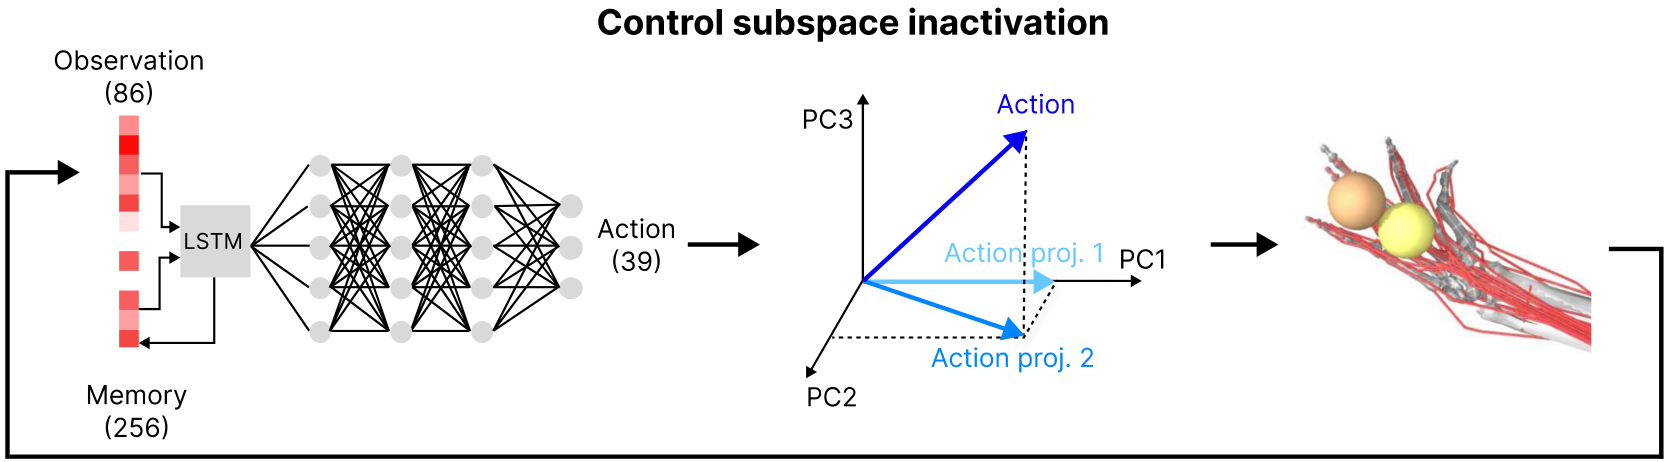

Signal reconstruction underestimates necessary DoF

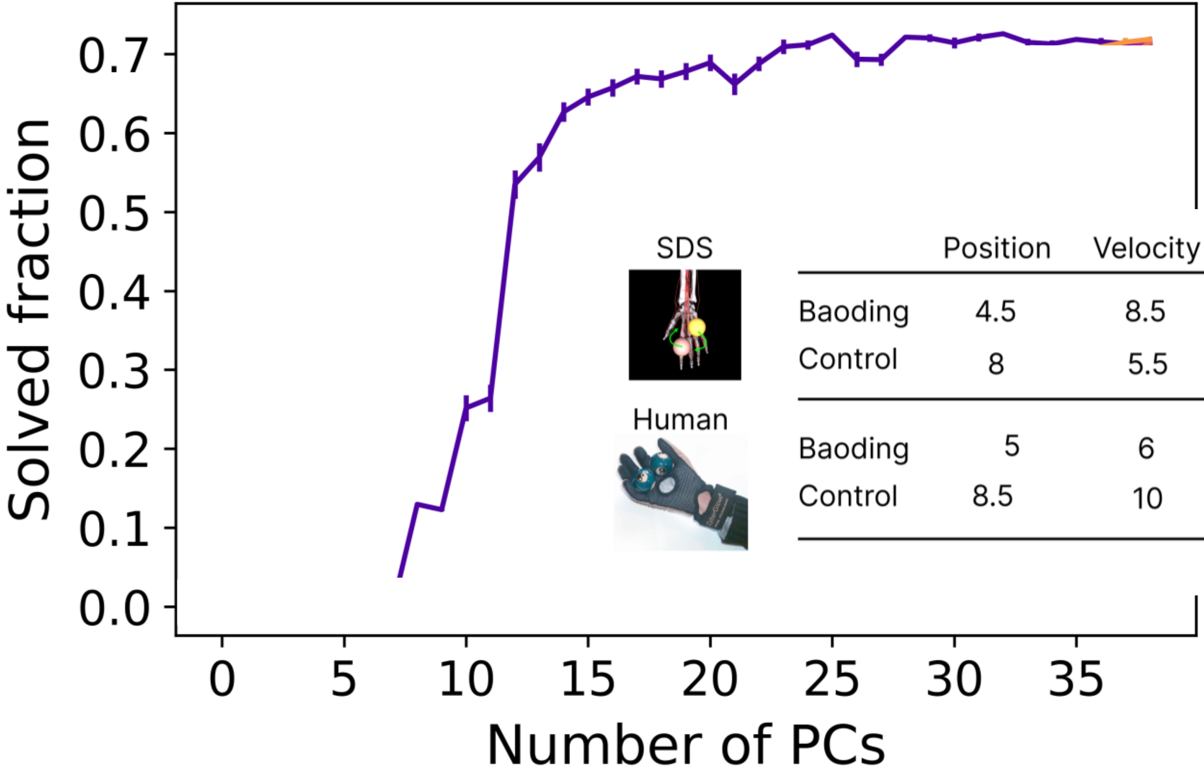

SDS

Human

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

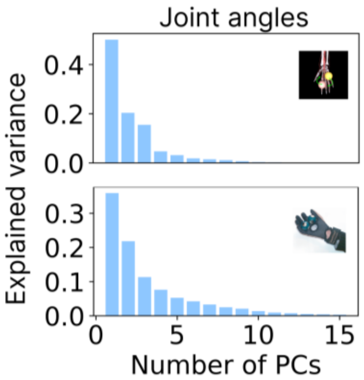


Late Lattice training performance
Chiappa et al. NeurIPS 2023
Run speed = 1.000 x real time [S]lower, [F]aster
Ren[d]er every frame On
Switch camera (#cams = 6) [Tab] (camera ID = -1)
[C]ontact forces On
Referenc[e] frames On
T[r]ansparent Off
Display [M]ocap bodies On
Stop [Space]
Advance simulation by one step [right arrow]
[H]ide Menu
Record [V]ideo (Off)
Cap[t]ure frame
Start [i]pdb
Toggle geomgroup visibility 0-4
FPS 267
Solver iterations 2
Reorient task
In MyoSuite/Mujoco
Step 530
timestep 0.00200
n_substeps 1

Control spaces are highly task-dependent & transfer poorly

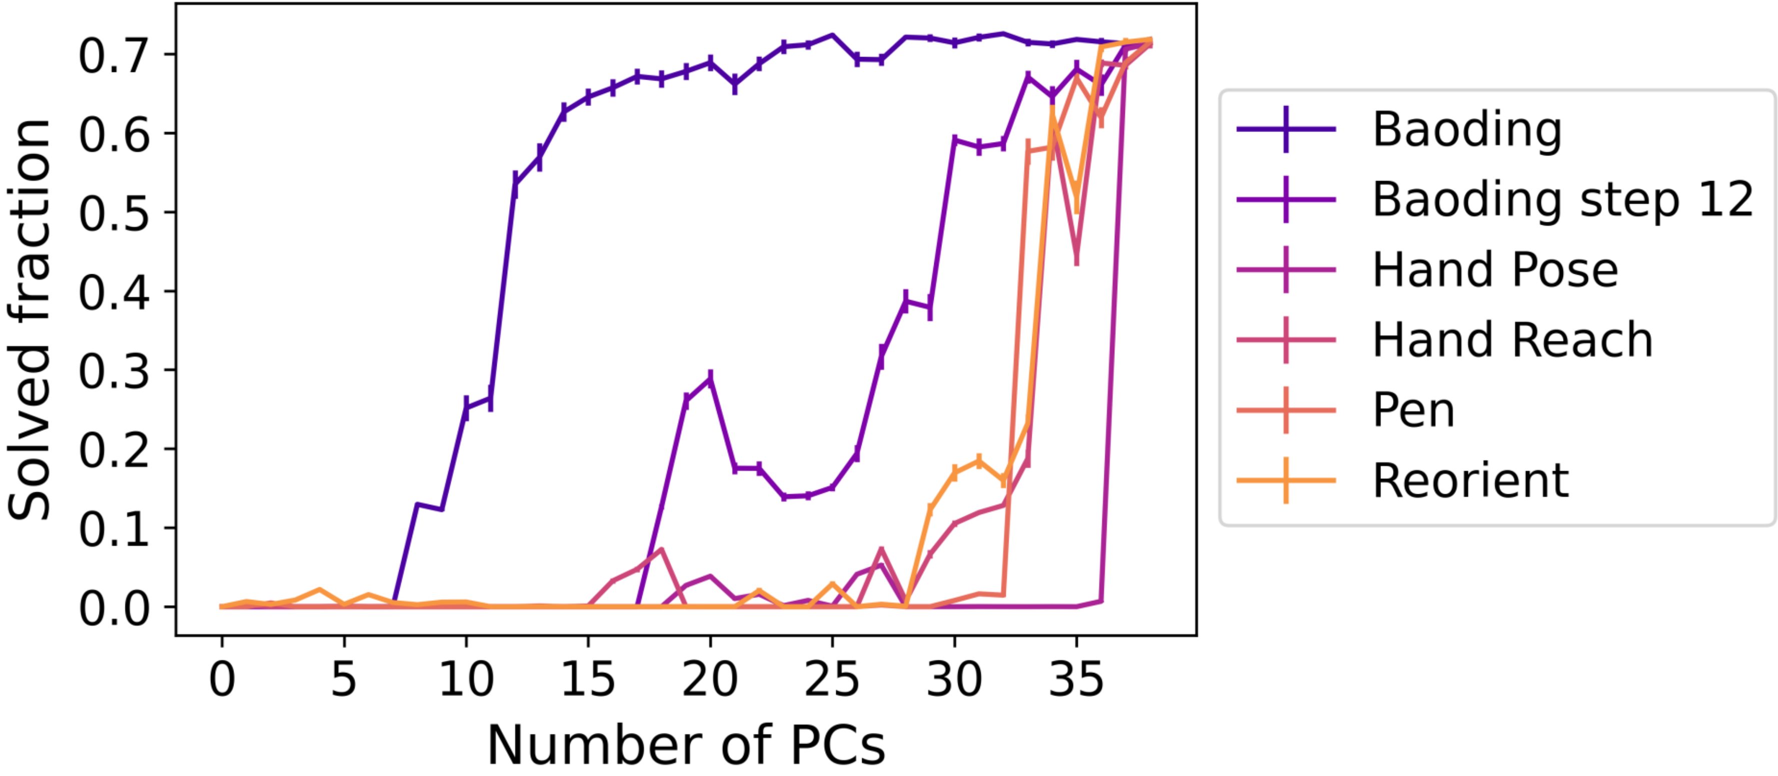

Intermediate take-home messages

- Muscle synergies have been proposed as a key principle for motor control
- Yet, low-dimensional nature might be **underestimated with existing techniques!**
- For the hand -- learned muscle synergies are highly task-specific, and thus generalize poorly
- This suggests that low-dimensional control is an emergent property (of the task/biomechanics/distributed circuits) rather than the mechanism of control (not a simplifying strategy)
- Neural networks are ideal for taming complex biomechanics

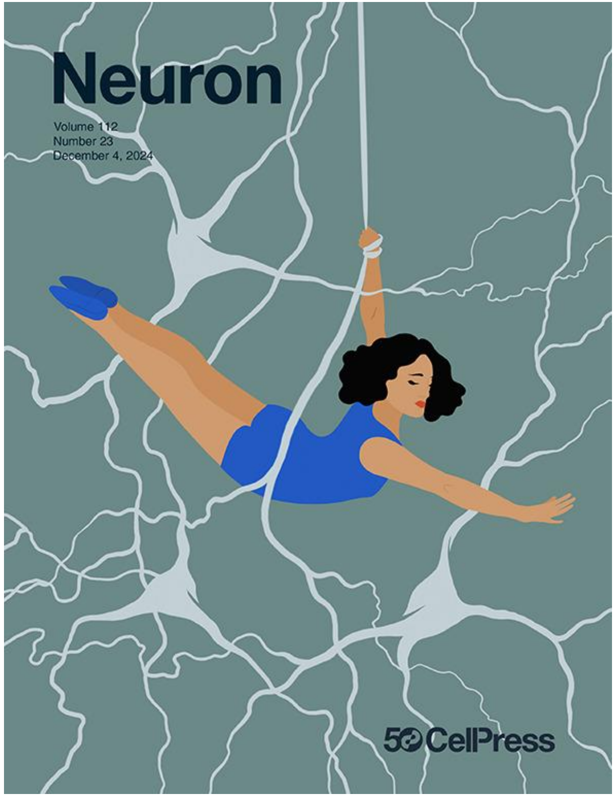

Curriculum learning in human motor control?

Part-to-whole practice, deliberate practice & challenge point learning

Part-to-whole practice

- Video suggested by Prof. Nicola Hodges

Challenge point: a framework for conceptualizing the effects of various practice conditions in motor learning

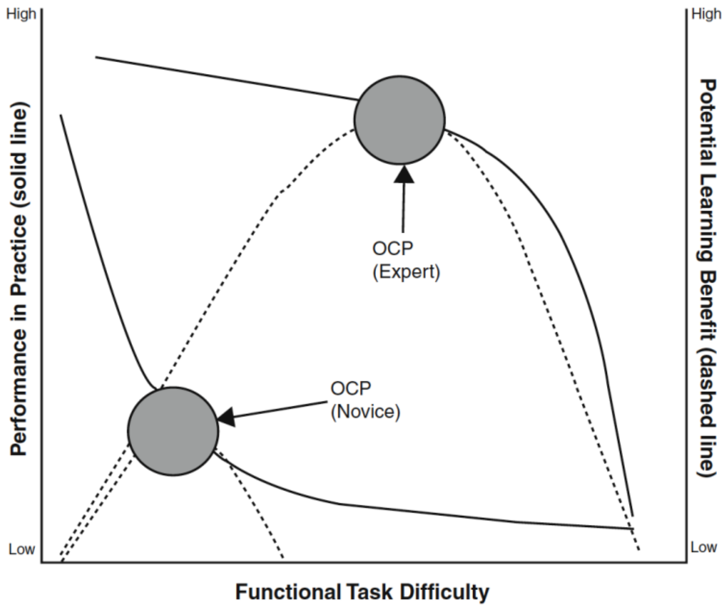

FIGURE 4. The relation between learning curves, performance curves, and the optimal challenge point (OCP) related to 2 performers of different skill levels.

Deliberate practice

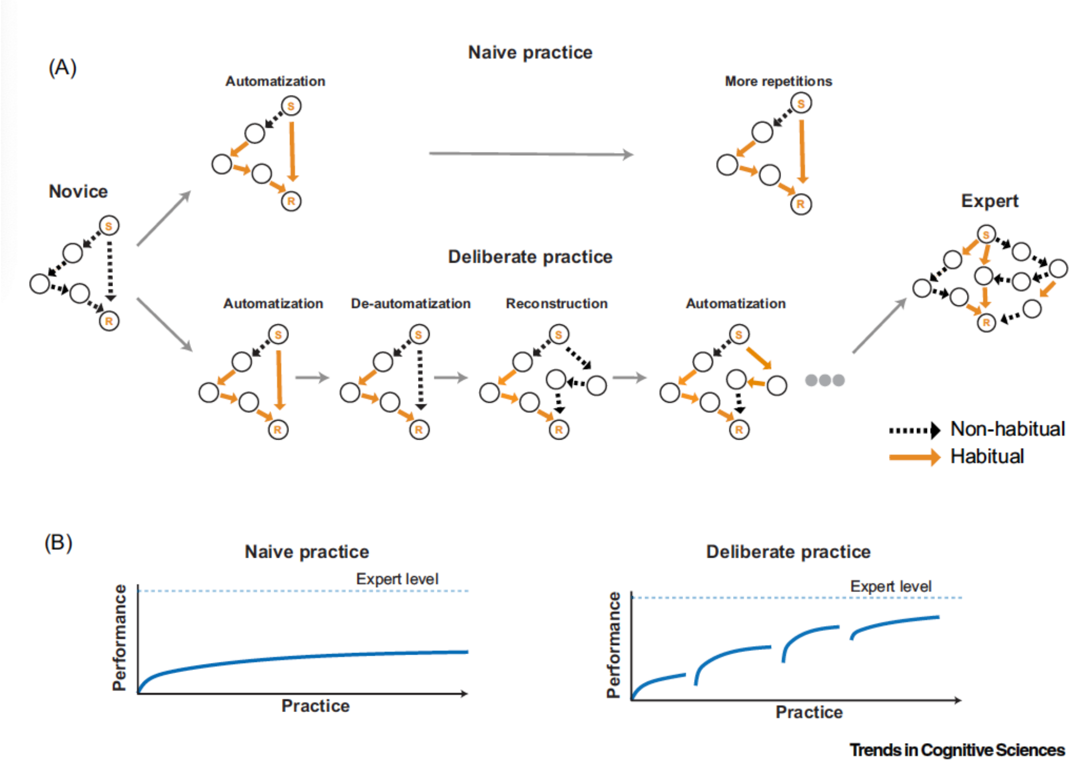

Du, Krakauer, Haith Trends in Cog Sci 2022

Deliberate practice & the SDS learning curve

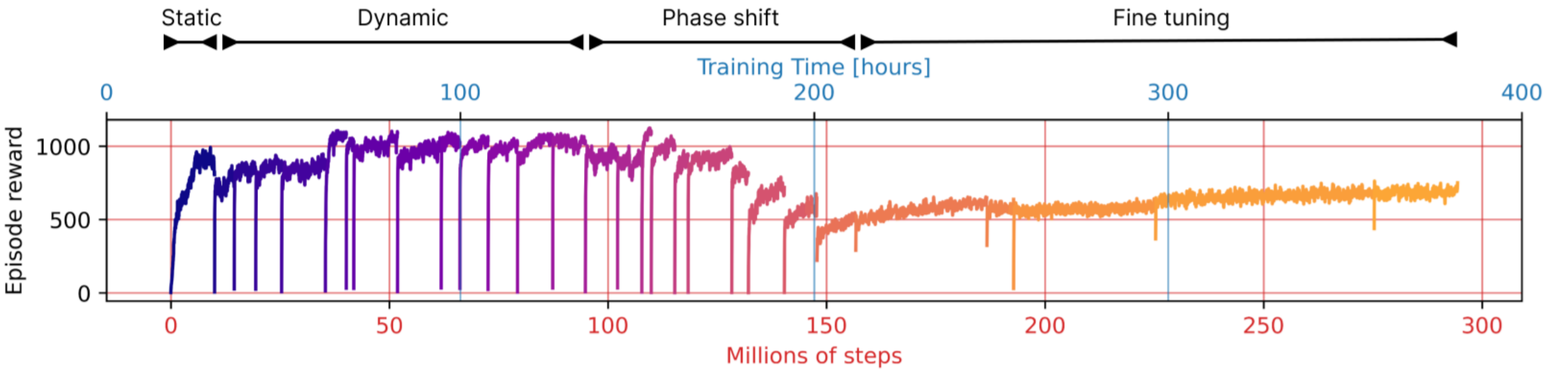

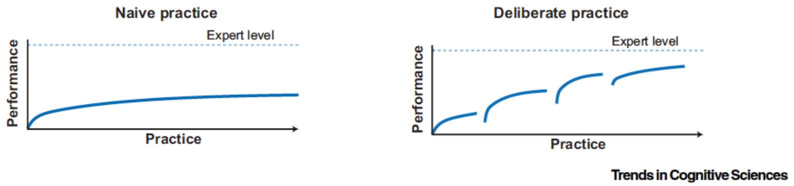

Du, Krakauer, Haith Trends in Cog Sci 2022

Reinforcement learning

Chess: 10^{120} Go: 3^{361}

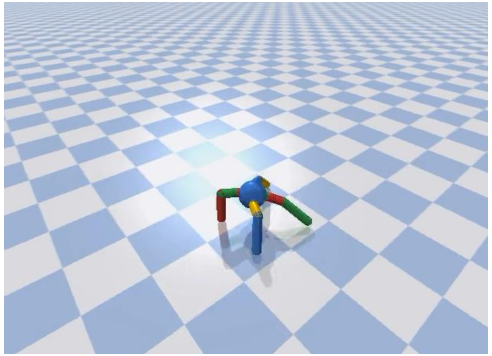

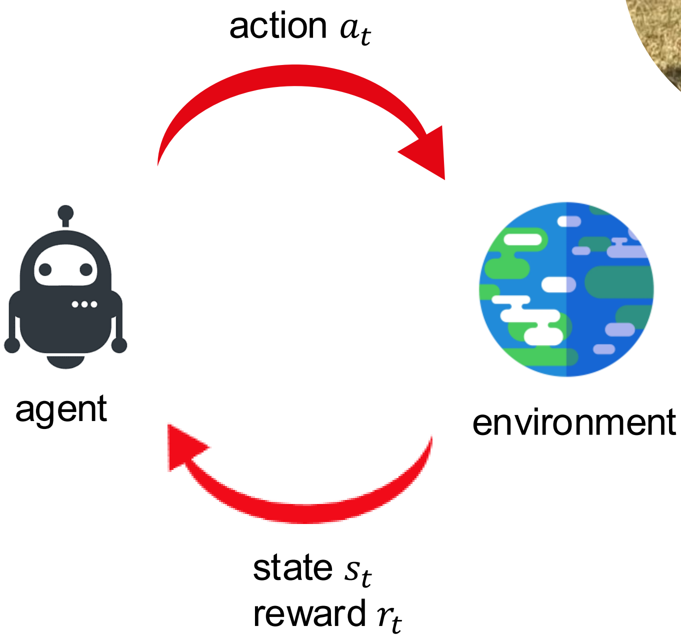

Nath*, Mathis* et al.
Nature Protocols 2019

M. Mathis Lab

https://www.youtube.com/watch?v=8vNxjwt2AqY

https://www.forbes.com/

enormous gap

What is missing?

- Internal models
- Inductive biases (innate architecture)
- *Better exploration*
- Baked in reward functions (which we don't know…)
- Using language
- *Curriculum learning*
- *Deliberate practice*
- ….

While we do not know their contributions or even the necessity of either one of those claims, I will show preliminary evidence for each to give you an idea.

There is a lot of research to be done to close this gap & figure out what actually matters…

A counter point – the bitter lesson

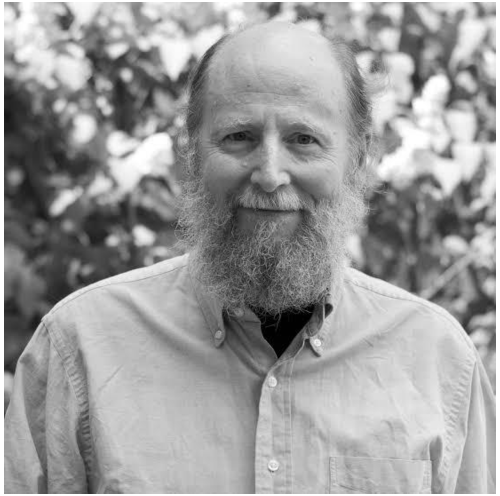

Richard Sutton

"The bitter lesson is based on the historical observations that 1) AI researchers have often tried to build knowledge into their agents, 2) this always helps in the short term, and is personally satisfying to the researcher, but 3) in the long run it plateaus and even inhibits further progress, and 4) breakthrough progress eventually arrives by an opposing approach based on scaling computation by search and learning. *The eventual success is tinged with bitterness, and often incompletely digested, because it is success over a favored, human-centric approach."*

http://www.incompleteideas.net/IncIdeas/BitterLesson.html