EPFL

NX-414: Brain-like computation and intelligence

Alexander Mathis
alexander.mathis@epfl.ch

Lecture 10, April 30 2025

Reverse engineering neural circuits

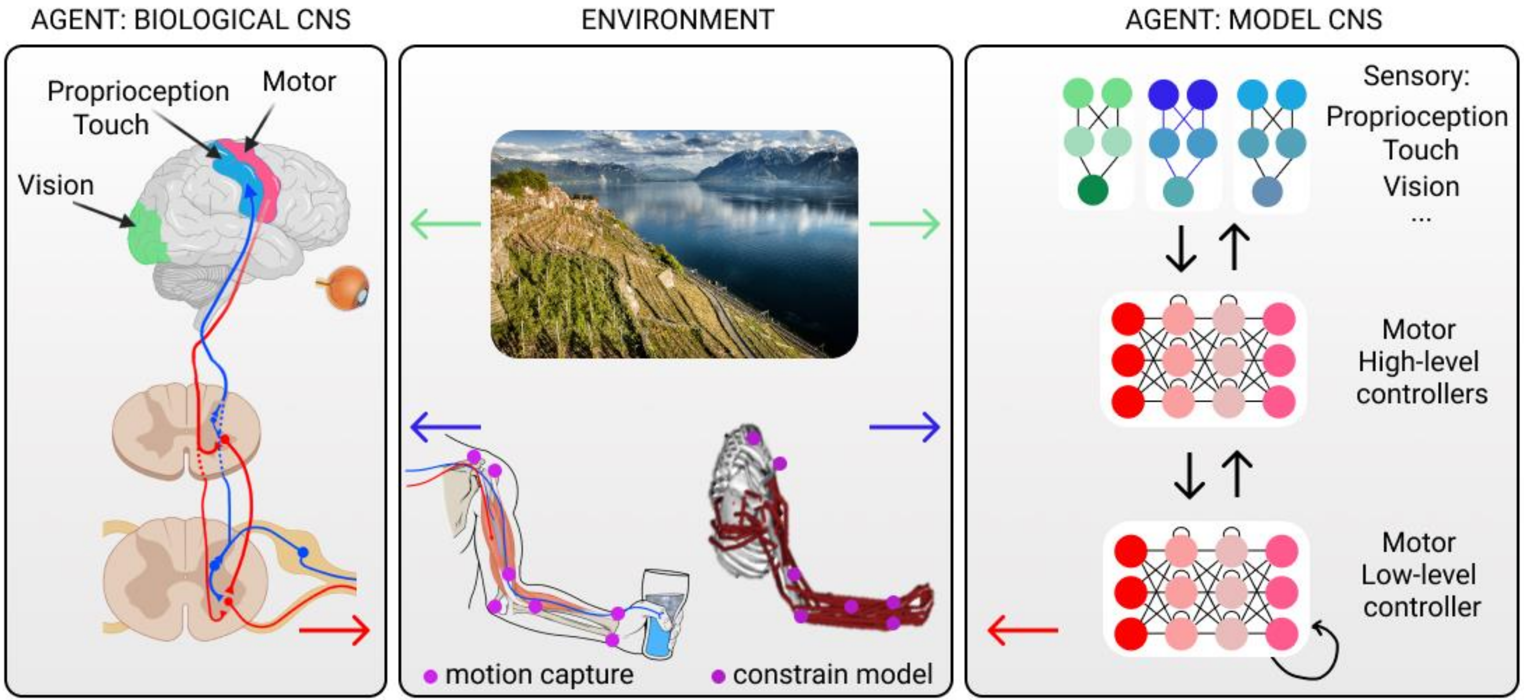

Hausmann*, Marin-Vargas*, Mathis**, Mathis**, Curr op. in Neuro., 2021

https://www.youtube.com/watch?v=fbqHK8i-HdA

Why is control hard for the brain?

- Unlike (typical) robots, animals live in **uncertain** environments
- Animals perform a **wide range** of behaviors
- Animals have **many degrees of freedoms** (e.g. human > 600 muscles)
- Biological sensors are **slower and noisier**
- Animal bodies change substantially over time (development, injury, fatigue, exercise,..). Our brain needs to **adapt** continuously
- Complex animals **learn** most of their behavioral repertoire, so the brain needs to not only control behavior, but also build control algorithms…

Behavior is hierarchical

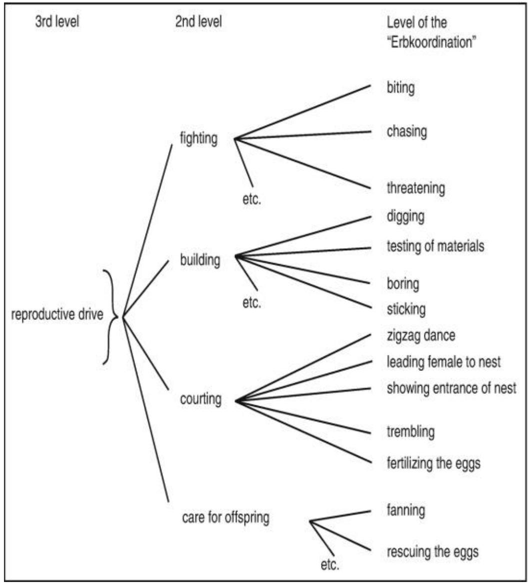

Tinbergen, 1942

Activities & Actions

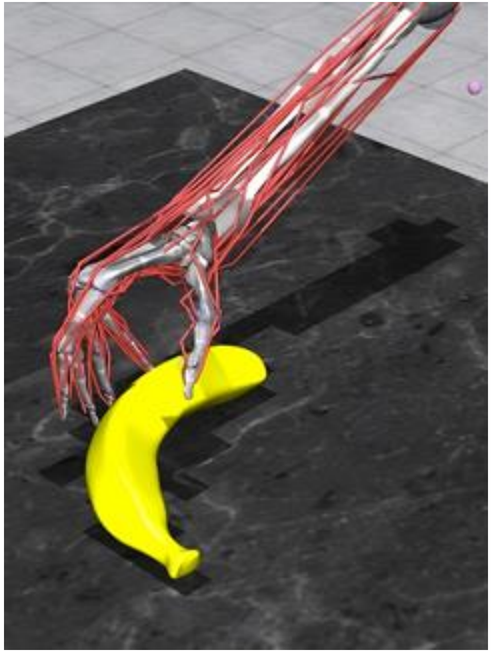

Movemes

Anderson & Perona Neuron 2014
Stoffl, Bonnetto, d'Ascoli, Mathis ECCV 2024

Hierarchical BehaveMAE (hBehaveMAE)

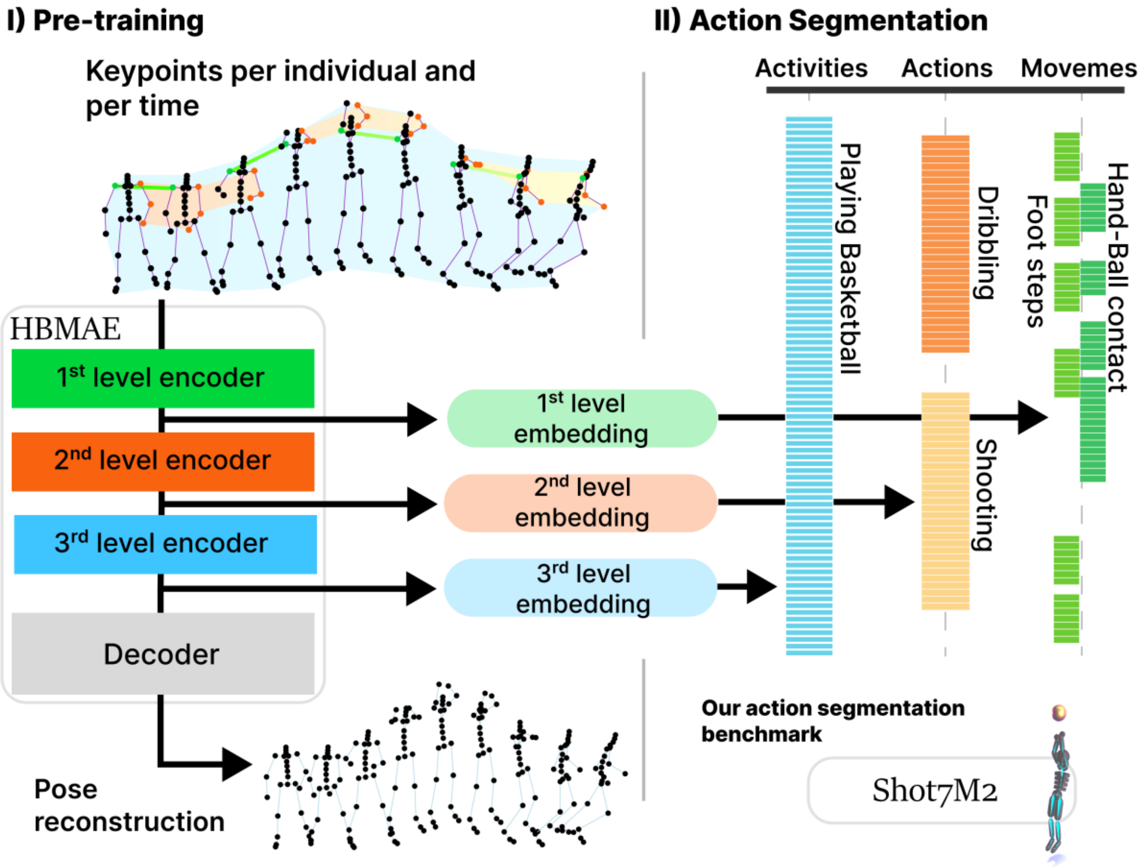

Stoffl, Bonnetto, d'Ascoli, Mathis ECCV 2024

Anatomy of motor control & pattern generators

Vertebrate motor control

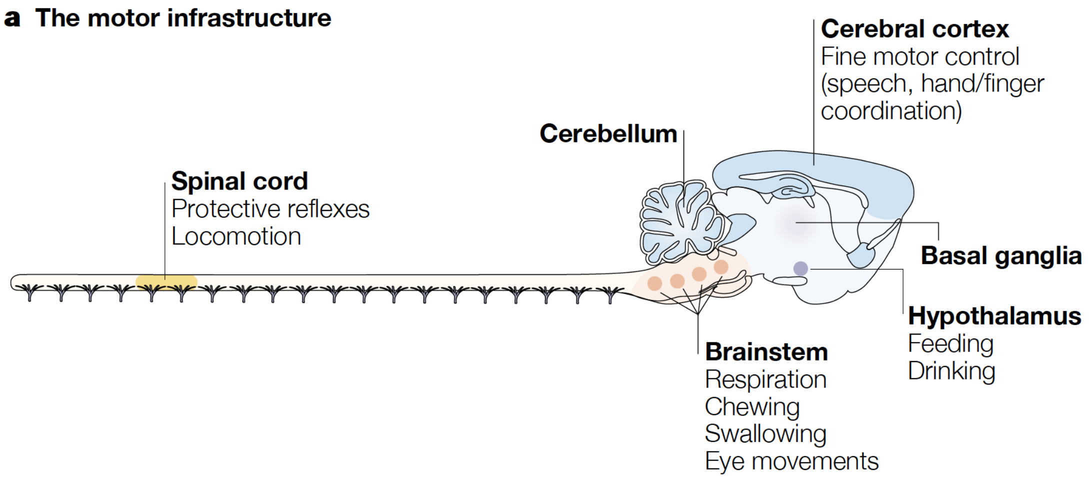

Grillner, Nature Review Neuroscience 2003

Vertebrate motor control

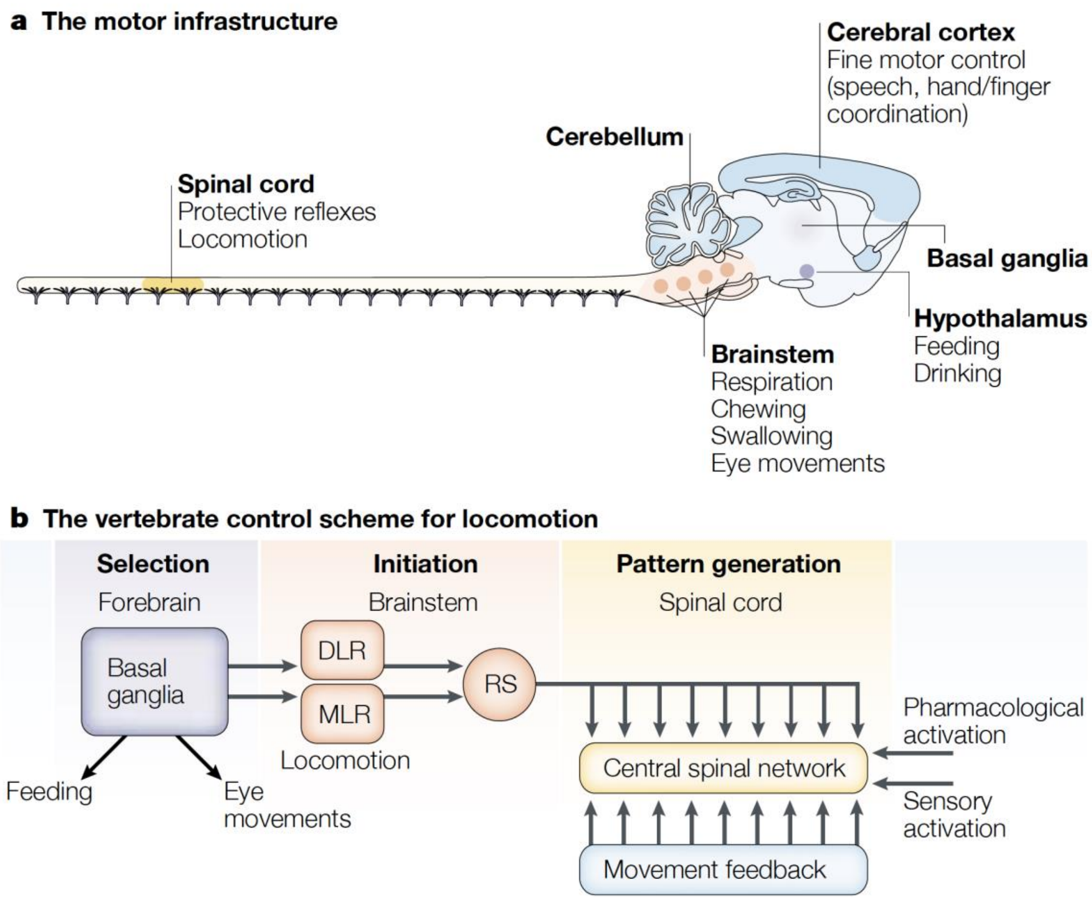

Grillner, Nature Review Neuroscience 2003

Pattern generation in the intact lamprey and an isolated spinal circuit

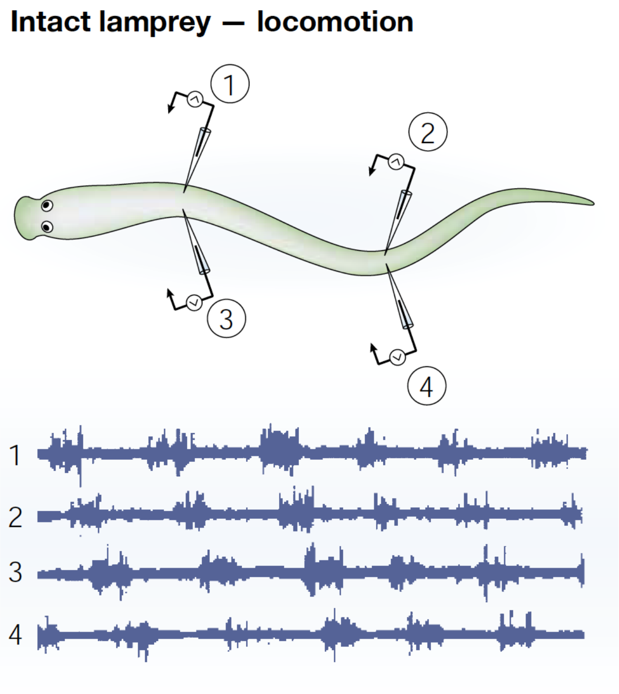

Note: alternation of 1/3 and 2/4 plus lag between 1 and 2.

Grillner, Nature Review Neuroscience 2003

Pattern generation in the intact lamprey and an isolated spinal circuit

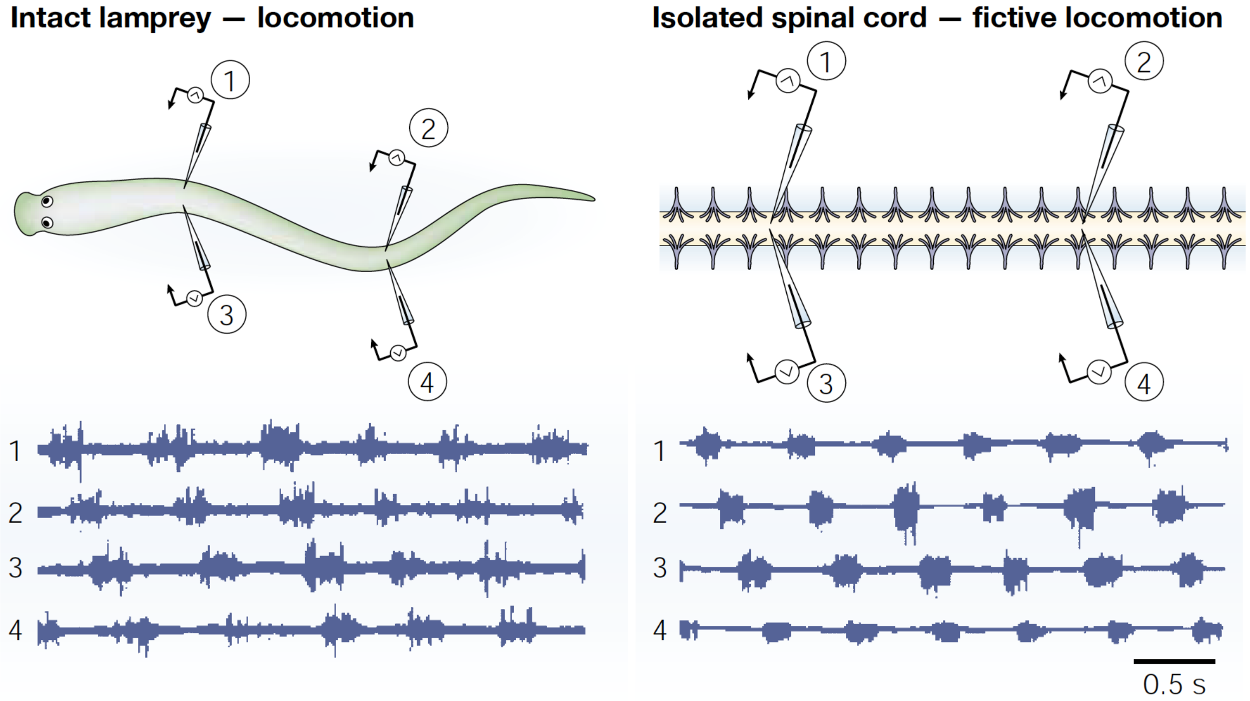

superfusion of glutamate agonists

Grillner, Nature Review Neuroscience 2003

Brain stem circuits to control locomotion

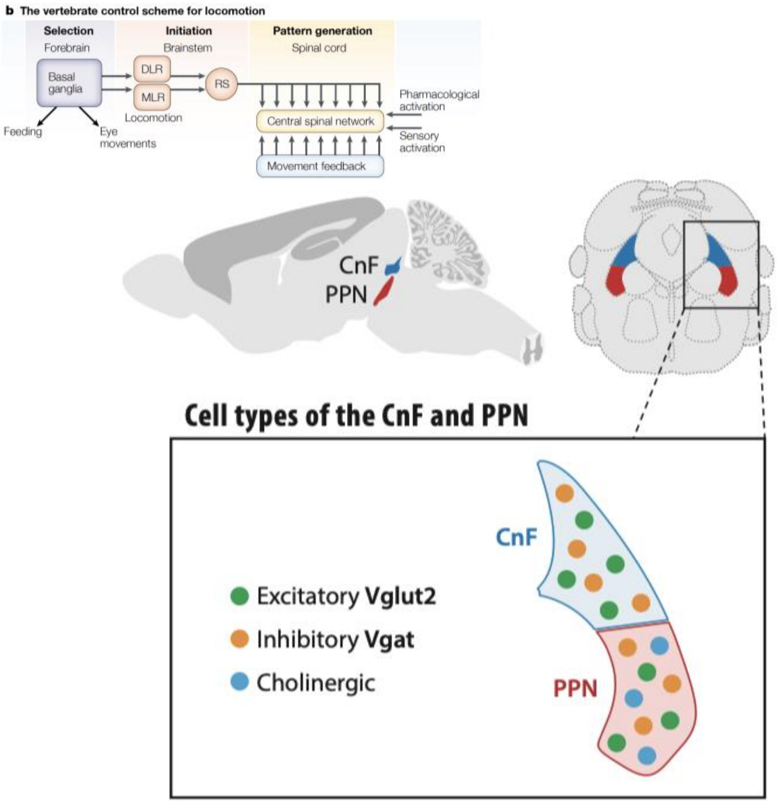

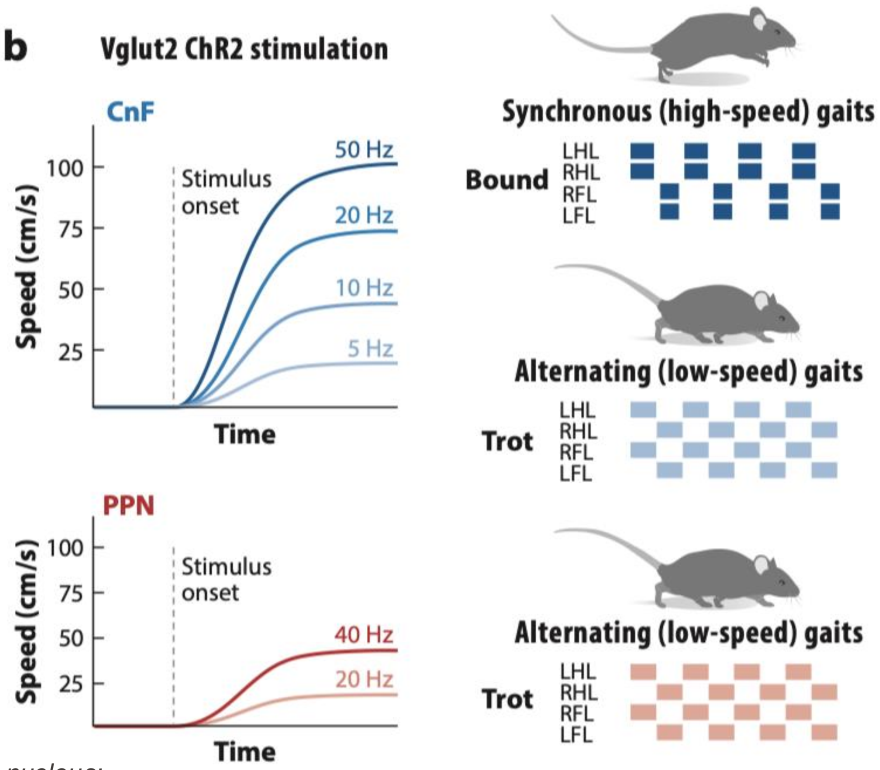

- *MLR, mesencephalic locomotor region; PPN, pedunculopontine nucleus; RFL, right forelimb; RHL, right hindlimb.*

Leiras, Cregg & Kiehn Annual Review of Neuroscience 2021

Brain stem circuits to control reaching & handling

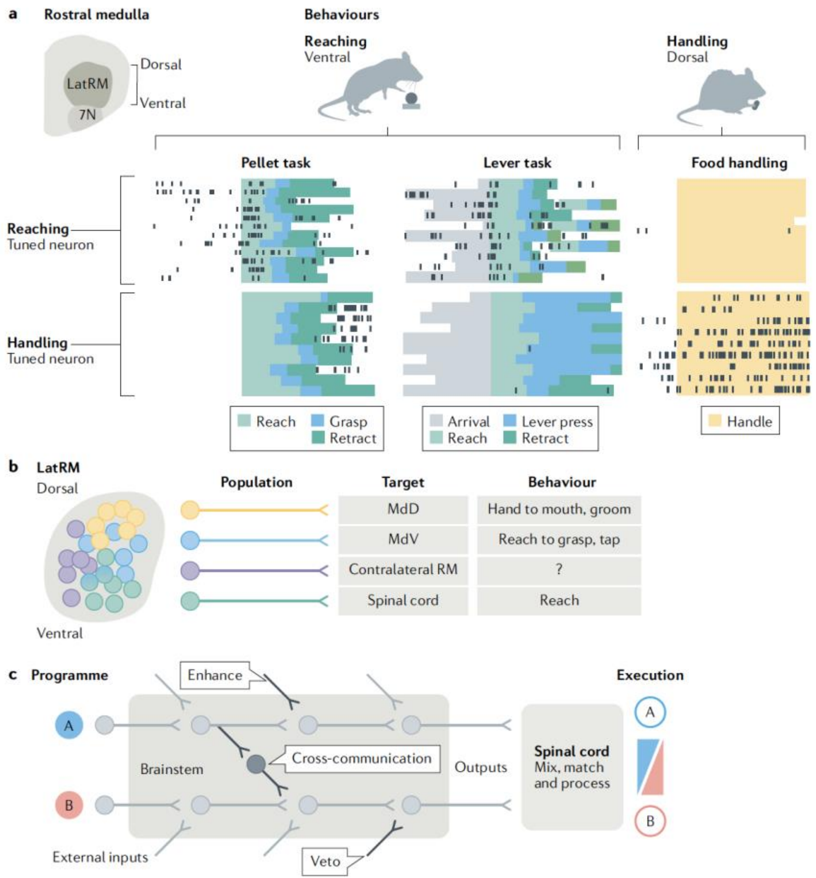

Arber & Costa Nature Review Neuroscience 2022

Cortical control

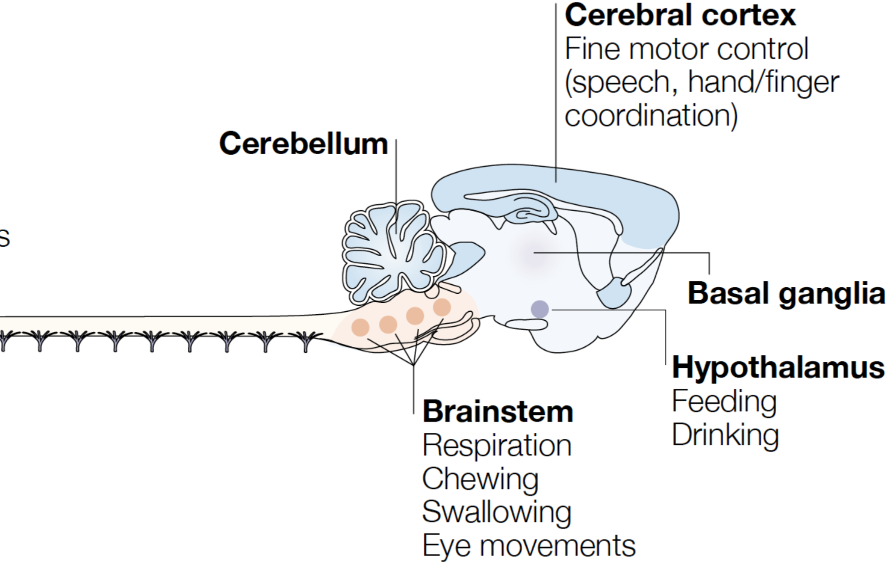

Reminder: Coding for the direction of movement

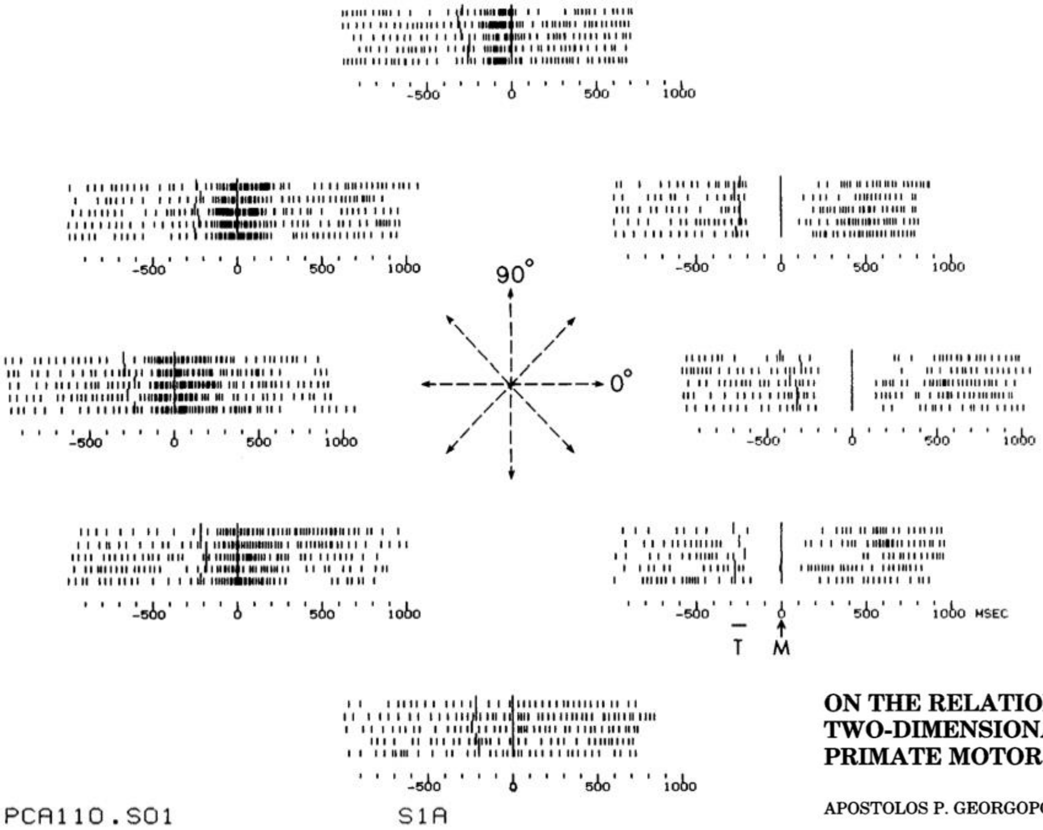

ON THE RELATIONS BETWEEN THE DIRECTION OF TWO-DIMENSIONAL ARM MOVEMENTS AND CELL DISCHARGE IN PRIMATE MOTOR CORTEX[1]

APOSTOLOS P. GEORGOPOULOS,[2] JOHN F. KALASKA,[3] ROBERTO CAMINITI,[4] AND JOE T. MASSEY[5]

Departments of Physiology and Neuroscience, The Johns Hopkins University School of Medicine, Baltimore, Maryland 21205

Received October 30, 1981; Revised April 30, 1982; Accepted May 21, 1982

Goal-driven modeling of motor control

Different colors correspond to different targets/conditions

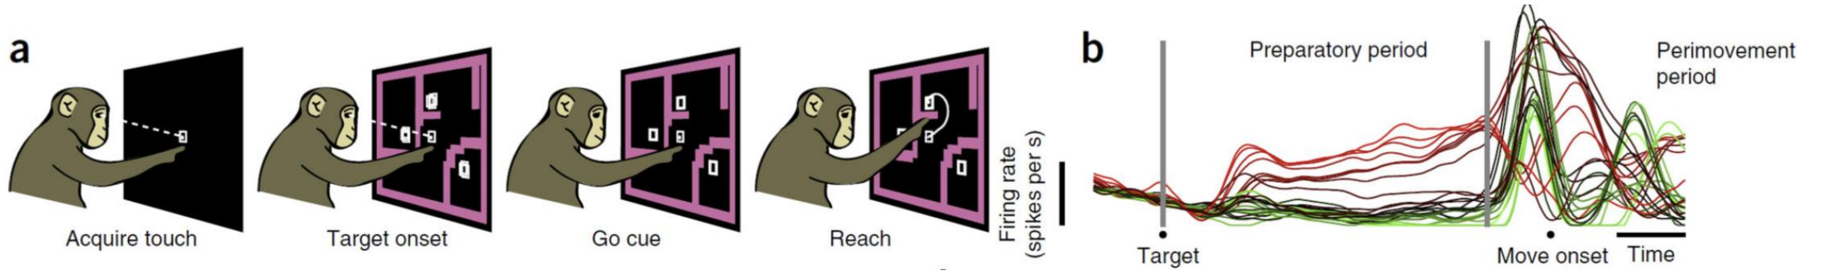

Susillo et al. Nature Neuro 2015

Goal-driven modeling of motor control

Technically data-driven as they fit EMG

Different colors correspond to different targets/conditions

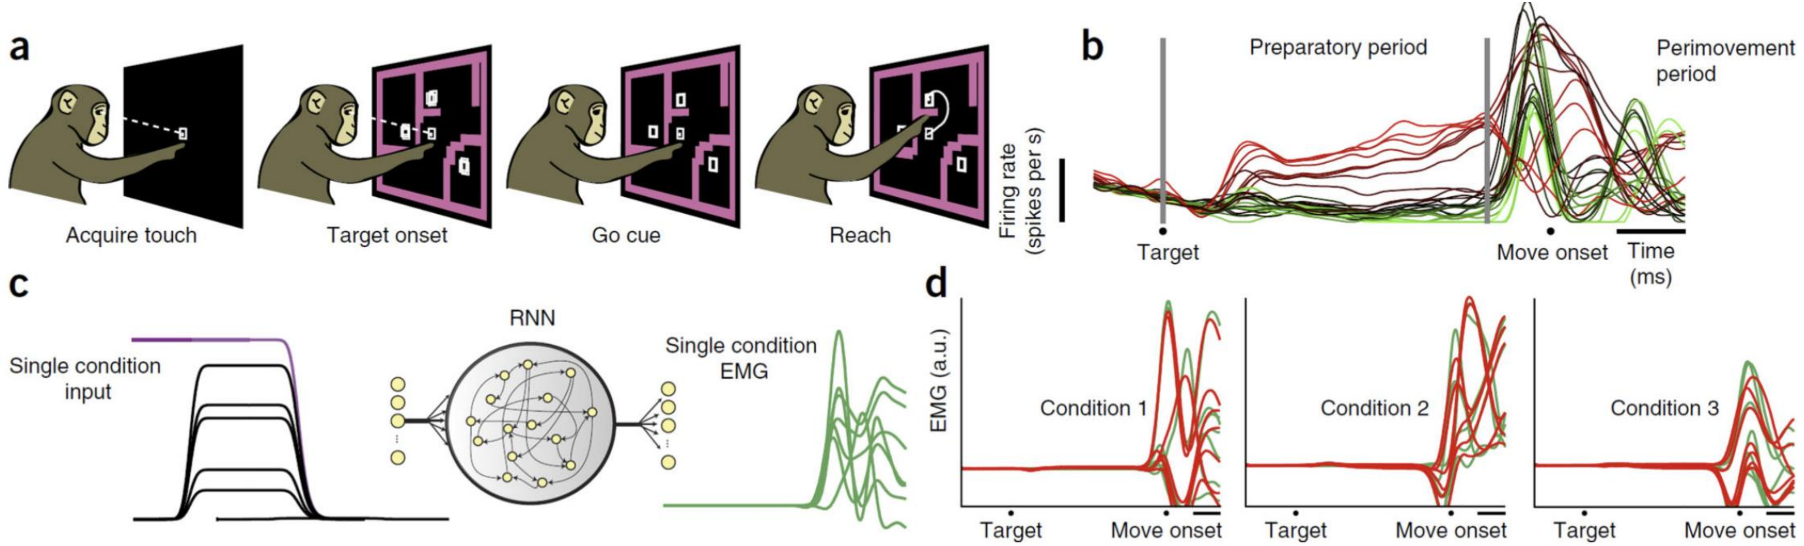

Susillo et al. Nature Neuro 2015

Example Data

Example M1 units

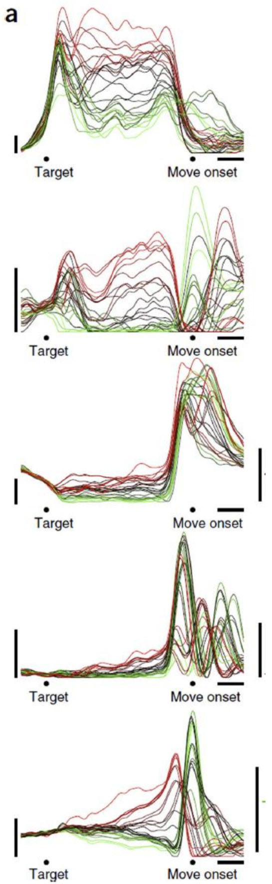

Susillo et al. Nature Neuro 2015

Model predictions

Example M1 units

Model units w/wo regularization

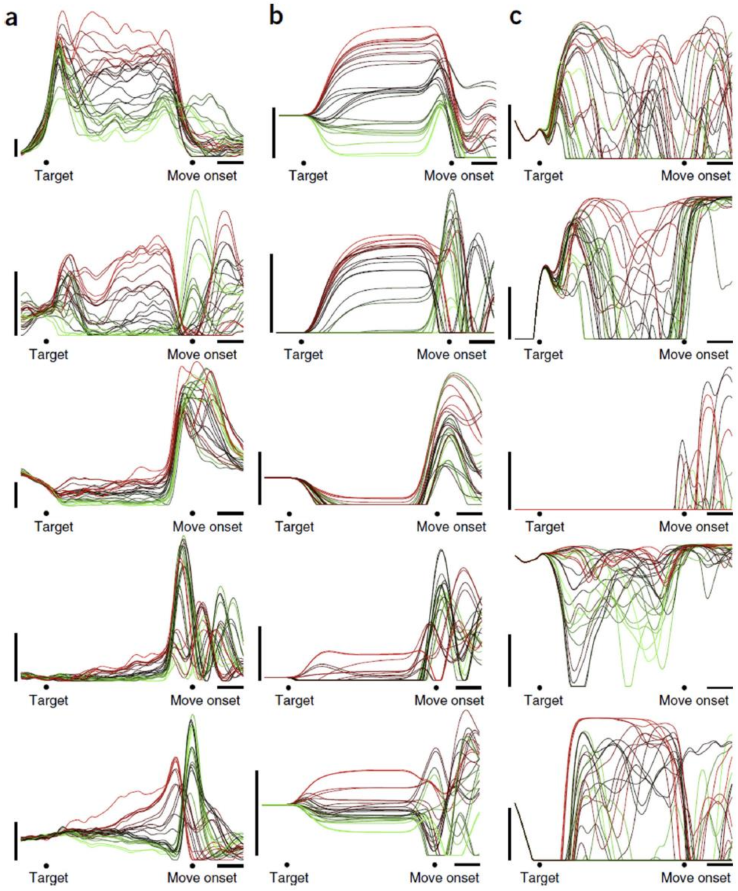

Susillo et al. Nature Neuro 2015

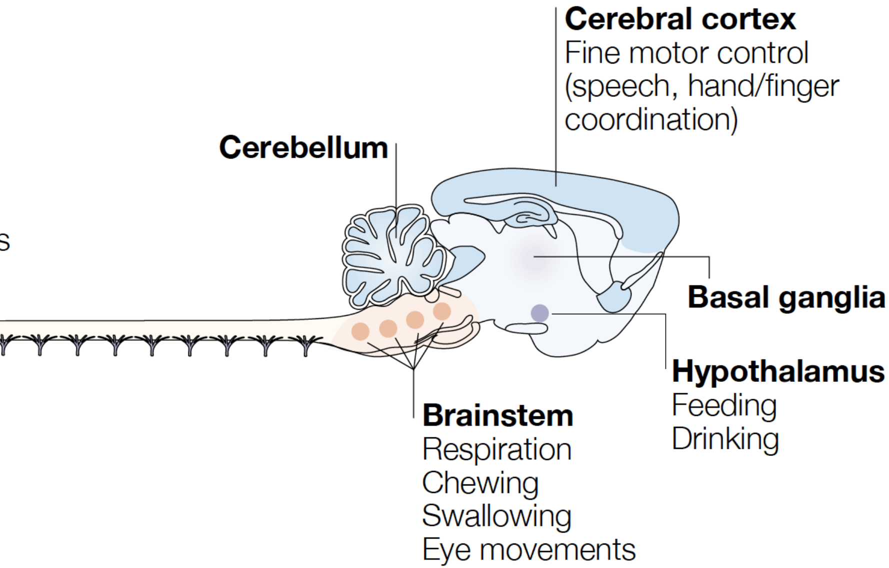

How are the many degrees of freedom tamed?

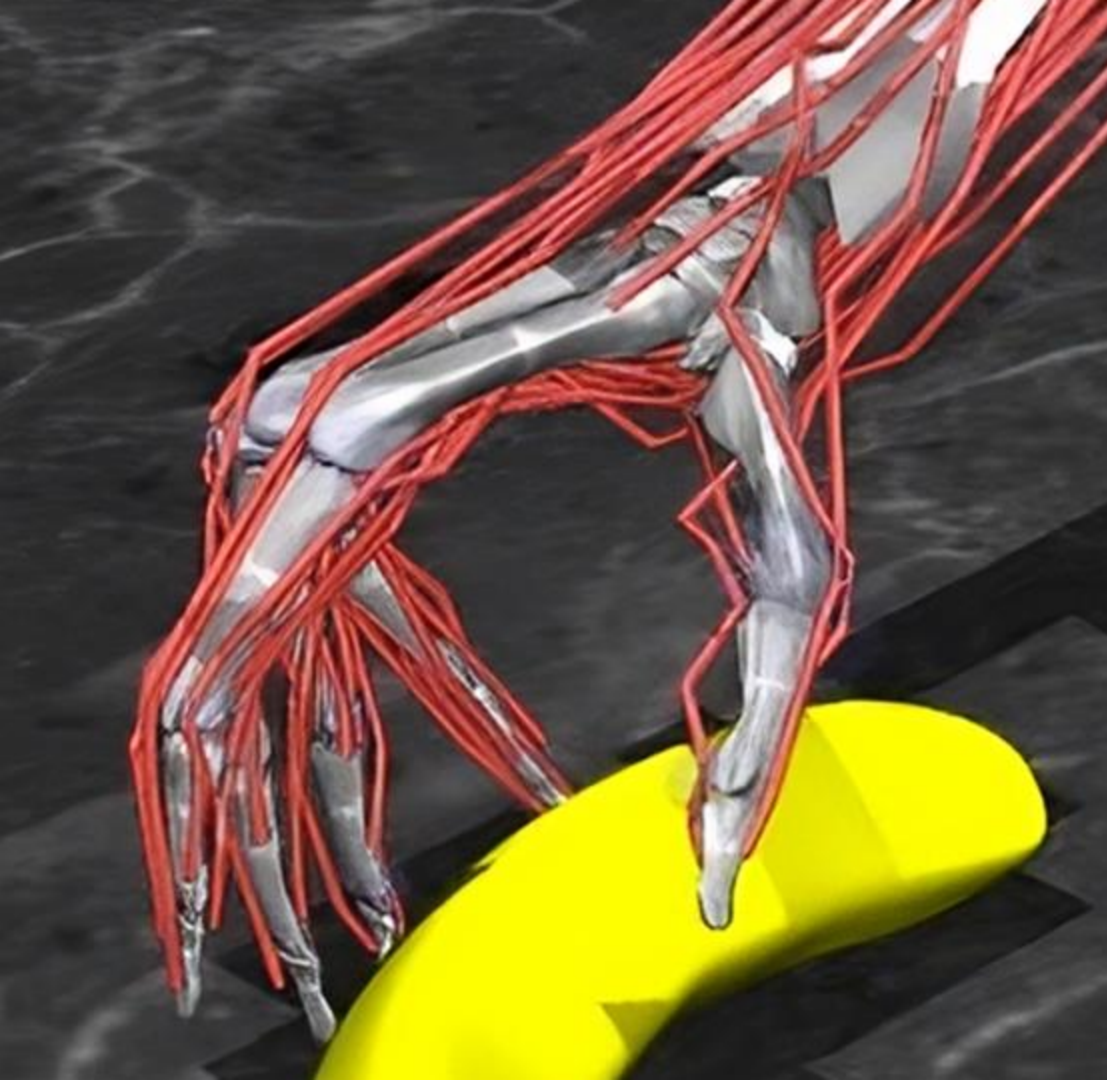

How many muscle states are there?

q^{600}

(for 600 muscles assuming q states per muscle)

EPFL

Muscle synergies as principle for motor control

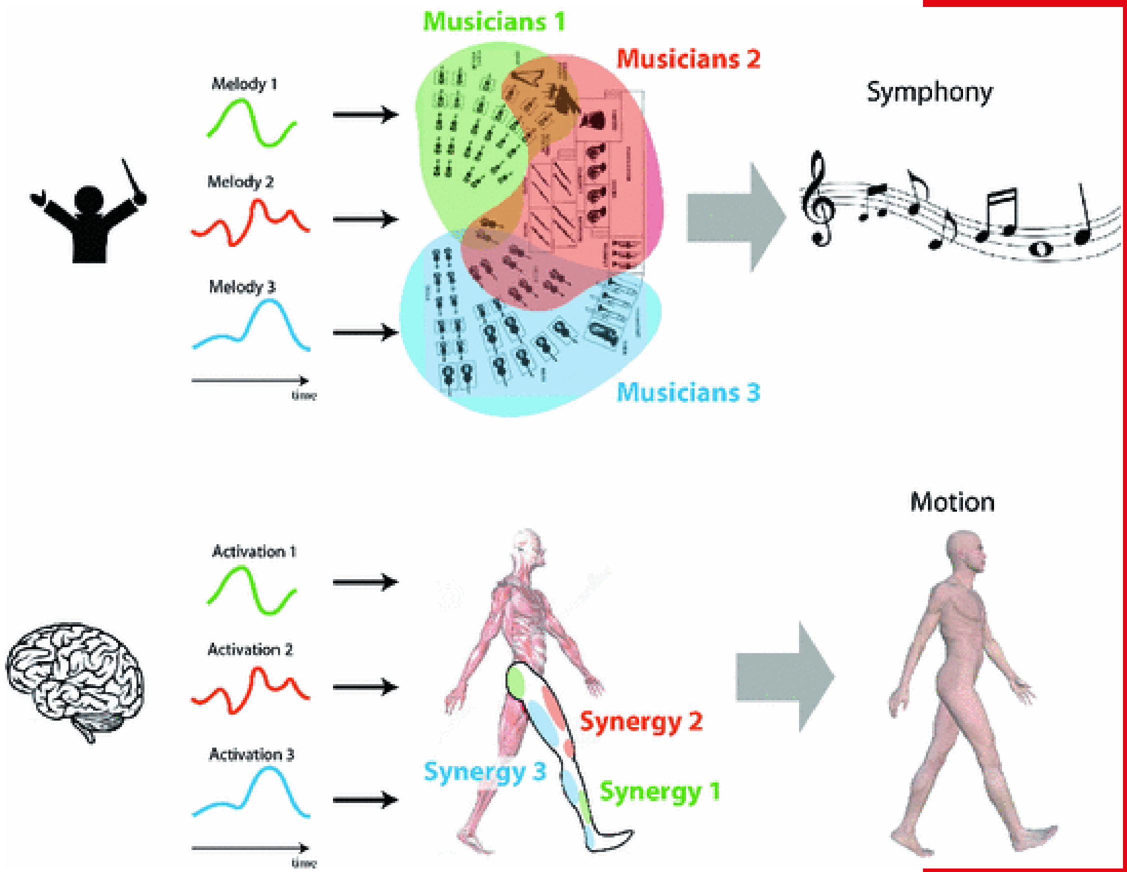

Torricelli et al., Emerging Therapies in Neurorehabilitation 2015

Integration of feedback?

From (open-loop) pattern generation to control theory

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

Motor skills need sensory feedback

AGENT: BIOLOGICAL CNS

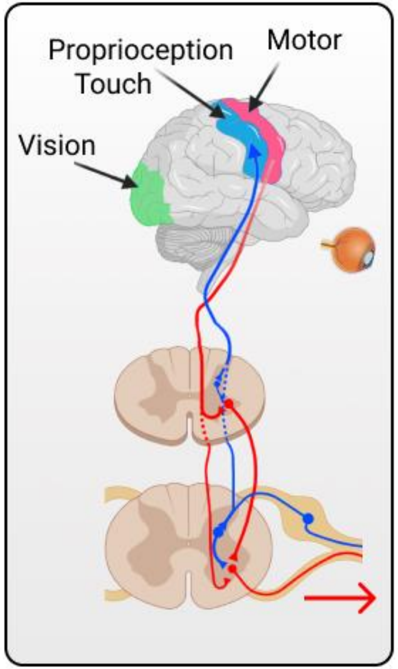

Sensory feedback

Motor commands

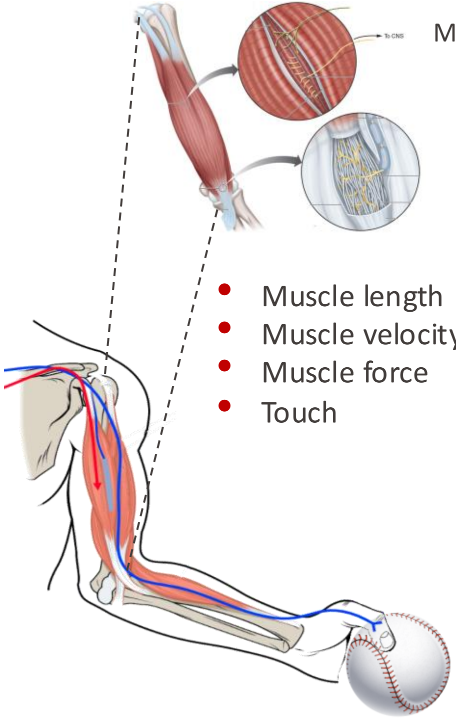

- Muscle length
- Muscle velocity
- Muscle force
- Touch

Simple skills require feedback

- Video taken from Roland Johansson Lab - Department of Integrative Medical Biology. Umea University, Sweden

Direct model of perception and control

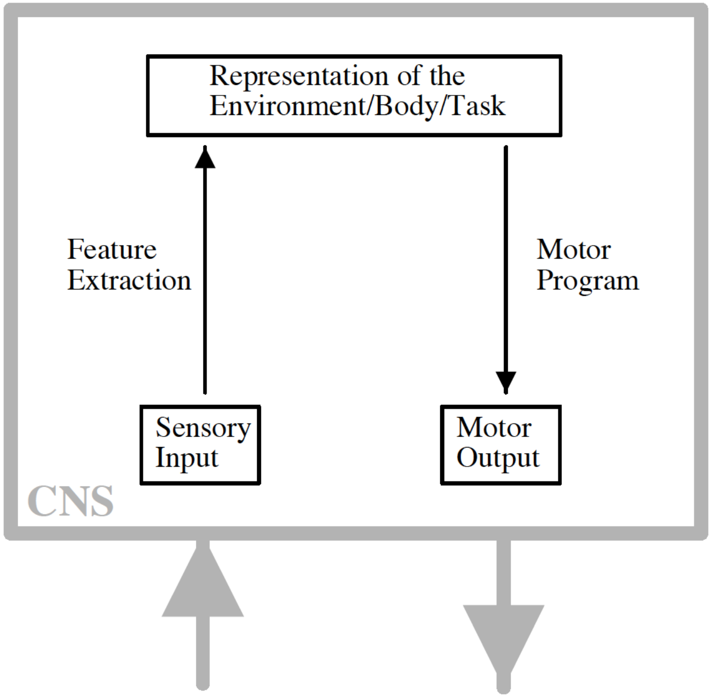

Todorov, 1998

Inverse models of perception and motor control

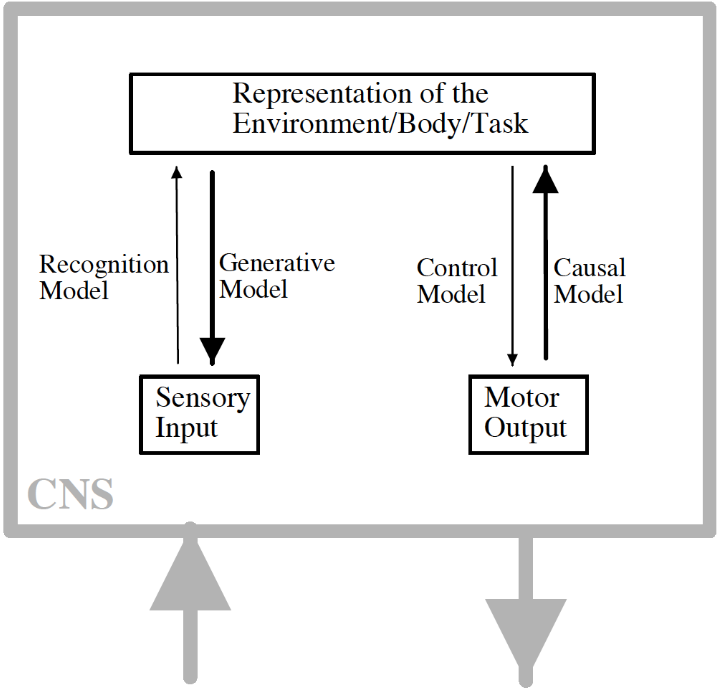

$P(G|M)$ causal model

$P(M)$ Movement prior

$$P(M|G)$$

Objective: find motor commands with high posterior probability!

M Motor command
G goal/target

The thin arrows correspond to the the directions that are desirably but *harder to implement*!

The thick arrows correspond to well-defined (relatively simpler transformations);
e.g., - generative model of vision: given the state of the world, predicting the retinal image (Optics,…)

- - causal model: given a motor command we can predict how it will change the world (Newtonian physics, ..)

Todorov, 1998

Motor responses can correct “online”

Nashed, Crevecoeur & Scott J. Neuro 2012

Corrective motor responses are tuned to goals

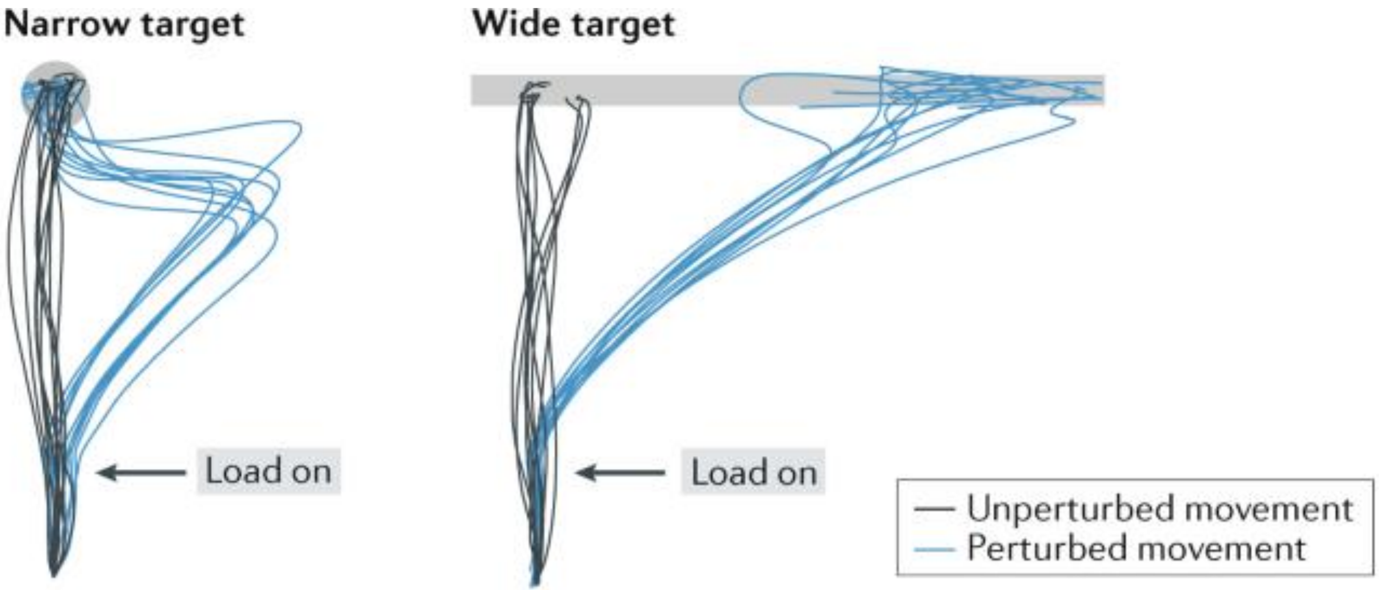

Nashed, Crevecoeur & Scott J. Neuro 2012

EPFL

How does the motor system work?

The motor system shows two apparently conflicting features:

- the ability to accomplish high-level goals reliably and repeatedly,
- despite high variability on the level of movement details

This is fundamentally incompatible with models that enforce a **strict separation between trajectory planning and trajectory execution.**

Optimal feedback control theory (OFC) postulates that the motor system approximates the best possible control scheme for a given task. It is updated online based on available information.

Todorov & Jordan, Nature Neuroscience 2001

Optimal feedback control (OFC) theory

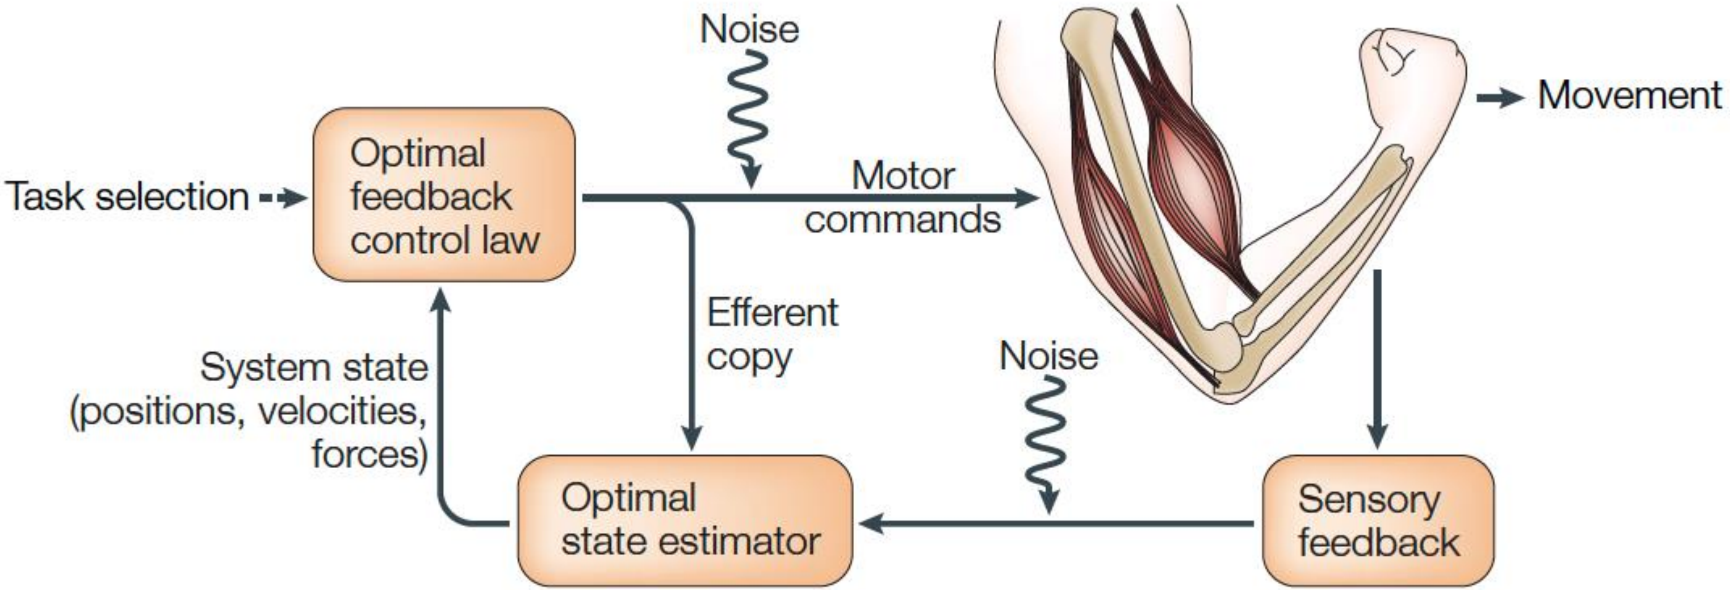

Scott, Nature Reviews Neuroscience 2004

The computational problem of reaching

We assume that reaching (picking motor commands)
is a consequence of maximizing rewards and minimizing costs.

Problem statement

Consider a linear dynamical system with state x, control u and feedback y in discrete time t:

Dynamics $$\mathbf{x}_{t+1} = A\mathbf{x}_t + B\mathbf{u}_t + \boldsymbol{\xi}_t + \sum_{i=1}^{c} \varepsilon_t^i C_i \mathbf{u}_t$$

Feedback $$\mathbf{y}_t = H\mathbf{x}_t + \boldsymbol{\omega}_t + \sum_{i=1}^{d} \epsilon_t^i D_i \mathbf{x}_t$$

Cost per step $$\mathbf{x}_t^T Q_t \mathbf{x}_t + \mathbf{u}_t^T R \mathbf{u}_t$$

Σ_t covariance of x_t

Note: this dyn. model is a simple causal model!

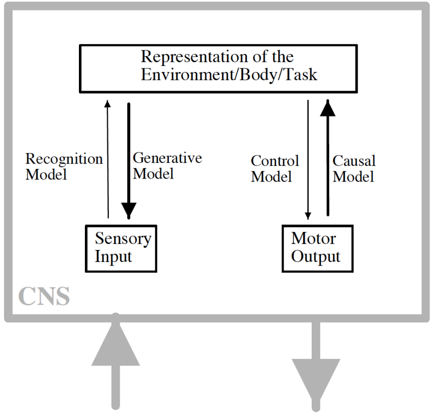

Todorov, Neural Computation 2005

EPFL

What is needed?

Dynamics	$\mathbf{x}_{t+1} = A\mathbf{x}_t + B\mathbf{u}_t + \boldsymbol{\xi}_t + \sum_{i=1}^{c} \varepsilon_t^i C_i \mathbf{u}_t$
Feedback	$\mathbf{y}_t = H\mathbf{x}_t + \boldsymbol{\omega}_t + \sum_{i=1}^{d} \epsilon_t^i D_i \mathbf{x}_t$
Cost per step	$\mathbf{x}_t^T Q_t \mathbf{x}_t + \mathbf{u}_t^T R \mathbf{u}_t$

1) The controller can only observe the state through noisy observations & needs to infer the state from noisy dynamics

2) Given state estimates, and past observations and controls one needs to compute

$$U_k^*\left(Y^{k-1}, U^{k-1}, \widehat{X}_1, \Sigma_1\right)$$

that minimizes $J = \mathrm{E}\left(X_N^{\mathrm{T}} Q_N X_N + \sum_{k=1}^{N-1} \left(X_k^{\mathrm{T}} Q_k X_k + U_k^{\mathrm{T}} R U_k\right)\right)$

Y^{k-1}: $[Y_1, \ldots, Y_{k-1}]$ $\quad$ U^{k-1}: $[U_1, \ldots, U_{k-1}]$

3) This amounts to minimizing expected costs by taking into account possible sequences of future controls & observations

$$\min_{U_k}\left(U_k^{\mathrm{T}} R U_k + \mathrm{E}\left(X_k^{\mathrm{T}} Q_k X_k\right) + \mathop{\mathrm{E}}_{Y_k}\left(\min_{U_{k+1}}\left(U_{k+1}^{\mathrm{T}} R U_{k+1} + \mathrm{E}\left(X_{k+1}^{\mathrm{T}} Q_{k+1} X_{k+1}\right) + \mathop{\mathrm{E}}_{Y_k+1}(\cdots)\right)\right)\right)$$

In general, this amounts to an exhaustive search in an exponentially large space.

Yet, with additive noise (next slide) there is an analytical solution!

Todorov, 1998

LQR solution (for only additive noise)

Remarkably there is an analytical solution:

Kalman Filter

$$\hat{\mathbf{x}}_{t+1} = A\hat{\mathbf{x}}_t + B\mathbf{u}_t + K_t(\mathbf{y}_t - H\hat{\mathbf{x}}_t)$$

$$K_t = A\Sigma_t H^\top (H\Sigma_t H^\top + \Omega^\omega)^{-1}$$

$$\Sigma_{t+1} = \Omega^\xi + (A - K_t H)\Sigma_t A^\top$$

Linear-Quadratic Regulator

$$\mathbf{u}_t = -L_t \hat{\mathbf{x}}_t$$

$$L_t = (R + B^\top S_{t+1} B)^{-1} B^\top S_{t+1} A$$

$$S_t = Q_t + A^\top S_{t+1}(A - BL_t)$$

Dynamics	$\mathbf{x}_{t+1} = A\mathbf{x}_t + B\mathbf{u}_t + \xi_t + \sum_{i=1}^{c} \varepsilon_t^i C_i \mathbf{u}_t$
Feedback	$\mathbf{y}_t = H\mathbf{x}_t + \omega_t + \sum_{i=1}^{d} \epsilon_t^i D_i \mathbf{x}_t$
Cost per step	$\mathbf{x}_t^T Q_t \mathbf{x}_t + \mathbf{u}_t^T R \mathbf{u}_t$

Note: **This hard computation inverts the thick arrow!!**

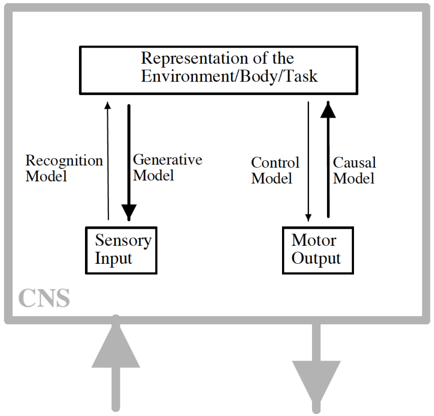

Todorov, Neural Computation 2005

LQR solution (for only additive noise)

Dynamics	$\mathbf{x}_{t+1} = A\mathbf{x}_t + B\mathbf{u}_t + \xi_t + \sum_{i=1}^{c} \varepsilon_t^i C_i \mathbf{u}_t$
Feedback	$\mathbf{y}_t = H\mathbf{x}_t + \omega_t + \sum_{i=1}^{d} \epsilon_t^i D_i \mathbf{x}_t$
Cost per step	$\mathbf{x}_t^T Q_t \mathbf{x}_t + \mathbf{u}_t^T R \mathbf{u}_t$

The (optimally) estimated state propagates forward according to the Kalman Filter

Kalman Filter

Internal state estimate $$\widehat{\mathbf{x}}_{t+1} = A\widehat{\mathbf{x}}_t + B\mathbf{u}_t + K_t(\mathbf{y}_t - H\widehat{\mathbf{x}}_t)$$

Kalman gain $$K_t = A\Sigma_t H^{\mathsf{T}}(H\Sigma_t H^{\mathsf{T}} + \Omega^{\omega})^{-1} \quad (3.2)$$

Internal state covariance $$\Sigma_{t+1} = \Omega^{\xi} + (A - K_t H)\Sigma_t A^{\mathsf{T}}$$

Sensory feedback covariance

Motor noise covariance

Todorov, Neural Computation 2005

LQR solution (for only additive noise)

Dynamics	$\mathbf{x}_{t+1} = A\mathbf{x}_t + B\mathbf{u}_t + \xi_t + \sum_{i=1}^{c} \varepsilon_t^i C_i \mathbf{u}_t$
Feedback	$\mathbf{y}_t = H\mathbf{x}_t + \omega_t + \sum_{i=1}^{d} \epsilon_t^i D_i \mathbf{x}_t$
Cost per step	$\mathbf{x}_t^T Q_t \mathbf{x}_t + \mathbf{u}_t^T R \mathbf{u}_t$

The optimal control law is obtained from the state estimate (of the KF), and *constant* matrices propagated backwards in time!

Linear-Quadratic Regulator

Control Law $\quad \mathbf{u}_t = -L_t \widehat{\mathbf{x}}_t$

$$L_t = (R + B^\top S_{t+1} B)^{-1} B^\top S_{t+1} A$$

$$S_t = Q_t + A^\top S_{t+1}(A - BL_t)$$

Does not depend on noise covariances; only KF depends on it (*Separation principle*)

Note: the optimal control law is computed from the optimal estimate, not the actual state!

- This is called the ***certainty-equivalence principle***.

Todorov, Neural Computation 2005

LQR solution (for only additive noise)

Dynamics	$\mathbf{x}_{t+1} = A\mathbf{x}_t + B\mathbf{u}_t + \xi_t + \sum_{i=1}^{c} \varepsilon_t^i C_i \mathbf{u}_t$
Feedback	$\mathbf{y}_t = H\mathbf{x}_t + \omega_t + \sum_{i=1}^{d} \epsilon_t^i D_i \mathbf{x}_t$
Cost per step	$\mathbf{x}_t^T Q_t \mathbf{x}_t + \mathbf{u}_t^T R \mathbf{u}_t$

Kalman Filter

$$\widehat{\mathbf{x}}_{t+1} = A\widehat{\mathbf{x}}_t + B\mathbf{u}_t + K_t(\mathbf{y}_t - H\widehat{\mathbf{x}}_t)$$

$$K_t = A\Sigma_t H^\mathsf{T}(H\Sigma_t H^\mathsf{T} + \Omega^\omega)^{-1}$$

$$\Sigma_{t+1} = \Omega^\xi + (A - K_t H)\,\Sigma_t A^\mathsf{T}$$

Linear-Quadratic Regulator

$$\mathbf{u}_t = -L_t\widehat{\mathbf{x}}_t$$

$$L_t = (R + B^\mathsf{T} S_{t+1} B)^{-1} B^\mathsf{T} S_{t+1} A$$

$$S_t = Q_t + A^\mathsf{T} S_{t+1}(A - BL_t)$$

Note: **This hard computation inverts the thick arrow!!**

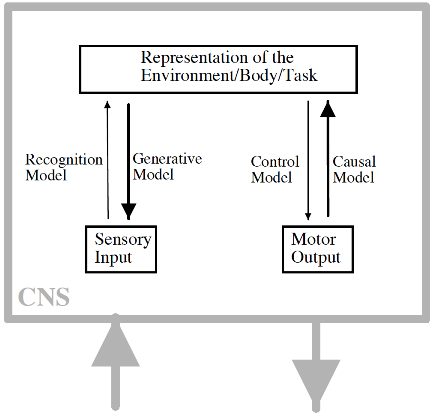

Todorov, Neural Computation 2005

The full case...

Dynamics	$\mathbf{x}_{t+1} = A\mathbf{x}_t + B\mathbf{u}_t + \boldsymbol{\xi}_t + \sum_{i=1}^{c} \varepsilon_t^i C_i \mathbf{u}_t$
Feedback	$\mathbf{y}_t = H\mathbf{x}_t + \boldsymbol{\omega}_t + \sum_{i=1}^{d} \epsilon_t^i D_i \mathbf{x}_t$
Cost per step	$\mathbf{x}_t^T Q_t \mathbf{x}_t + \mathbf{u}_t^T R \mathbf{u}_t$

$$\widehat{\mathbf{x}}_{t+1} = A\widehat{\mathbf{x}}_t + B\mathbf{u}_t + K_t(\mathbf{y}_t - H\widehat{\mathbf{x}}_t) + \eta_t$$

Controller $\mathbf{u}_t = -L_t \widehat{\mathbf{x}}_t$

$$L_t = \left(R + B^{\mathsf{T}} S_{t+1}^{\mathbf{x}} B + \sum_i C_i^{\mathsf{T}} (S_{t+1}^{\mathbf{x}} + S_{t+1}^{\mathbf{e}}) C_i\right)^{-1} B^{\mathsf{T}} S_{t+1}^{\mathbf{x}} A$$

$$S_t^{\mathbf{x}} = Q_t + A^{\mathsf{T}} S_{t+1}^{\mathbf{x}} (A - BL_t) + \sum_i D_i^{\mathsf{T}} K_t^{\mathsf{T}} S_{t+1}^{\mathbf{e}} K_t D_i; \qquad S_n^{\mathbf{x}} = Q_n$$

$$S_t^{\mathbf{e}} = A^{\mathsf{T}} S_{t+1}^{\mathbf{x}} B L_t + (A - K_t H)^{\mathsf{T}} S_{t+1}^{\mathbf{e}} (A - K_t H); \qquad S_n^{\mathbf{e}} = 0$$

$$s_t = \mathrm{tr}\left(S_{t+1}^{\mathbf{x}} \Omega^{\xi} + S_{t+1}^{\mathbf{e}} \left(\Omega^{\xi} + \Omega^{\eta} + K_t \Omega^{\omega} K_t^{\mathsf{T}}\right)\right) + s_{t+1}; \qquad s_n = 0.$$

Todorov, Neural Computation 2005

Linking brain areas to computations

Modeling predictions

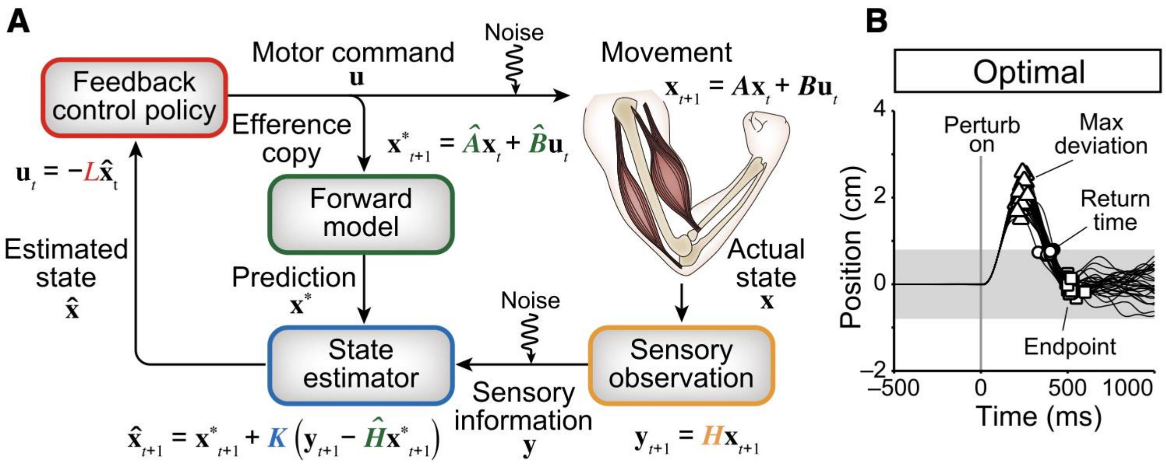

Takei et al., Curr Biology 2021

Modeling predictions

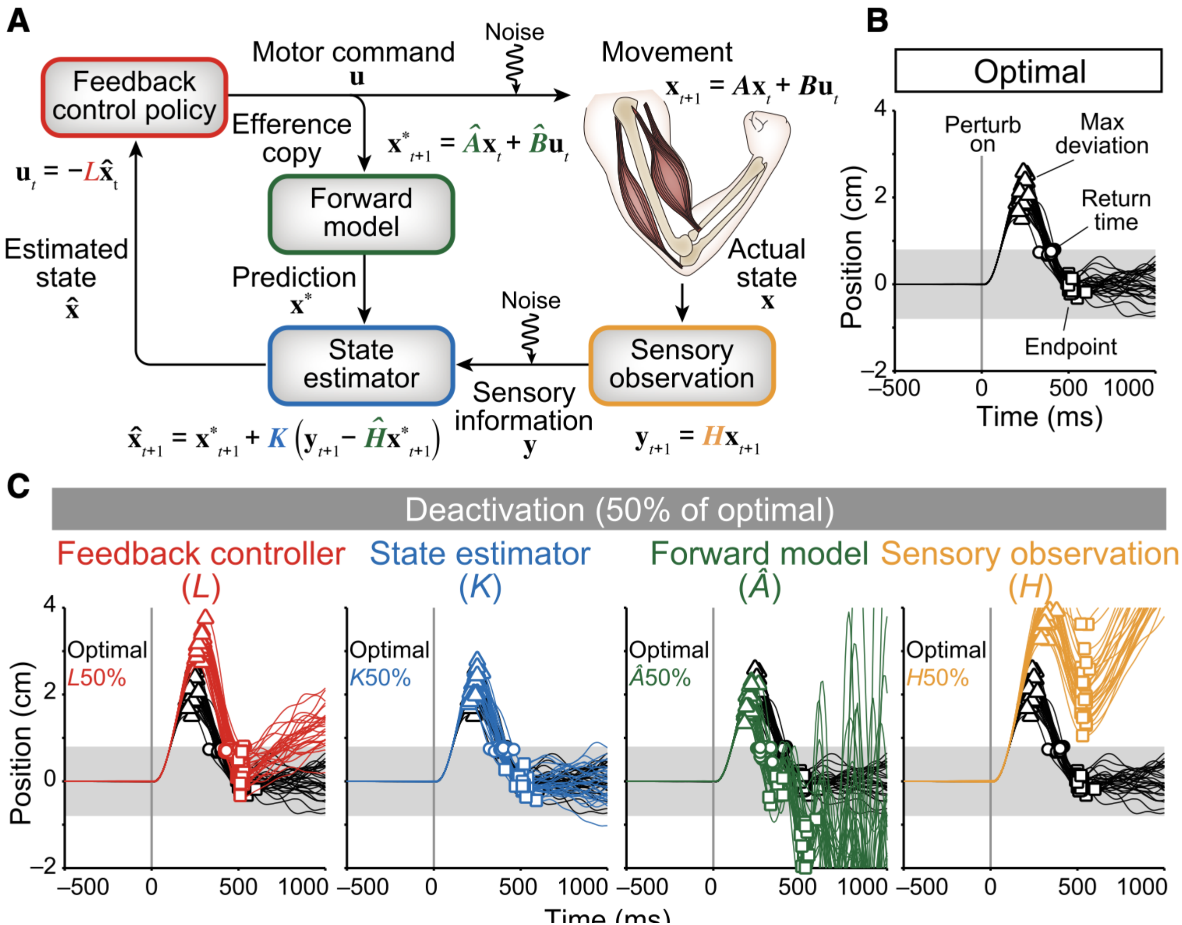

$$x_{t+1} = Ax_t + Bu_t$$

$$x^*_{t+1} = \hat{A}x_t + \hat{B}u_t$$

$$u_t = -L\hat{x}_t$$

$$\hat{x}_{t+1} = x^*_{t+1} + K\left(y_{t+1} - \hat{H}x^*_{t+1}\right)$$

$$y_{t+1} = Hx_{t+1}$$

Takei et al., Curr Biology 2021

Cooling the brain

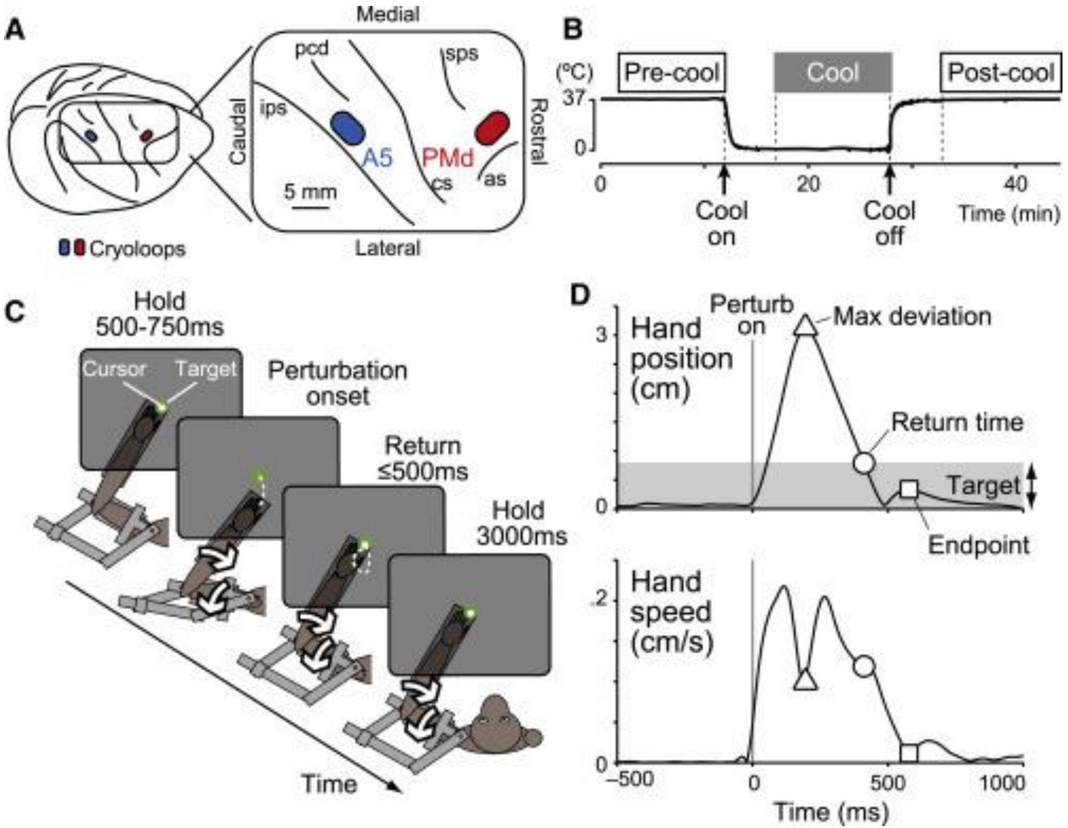

Takei et al., Curr Biology 2021

Experimental results for PMd and A5 cooling

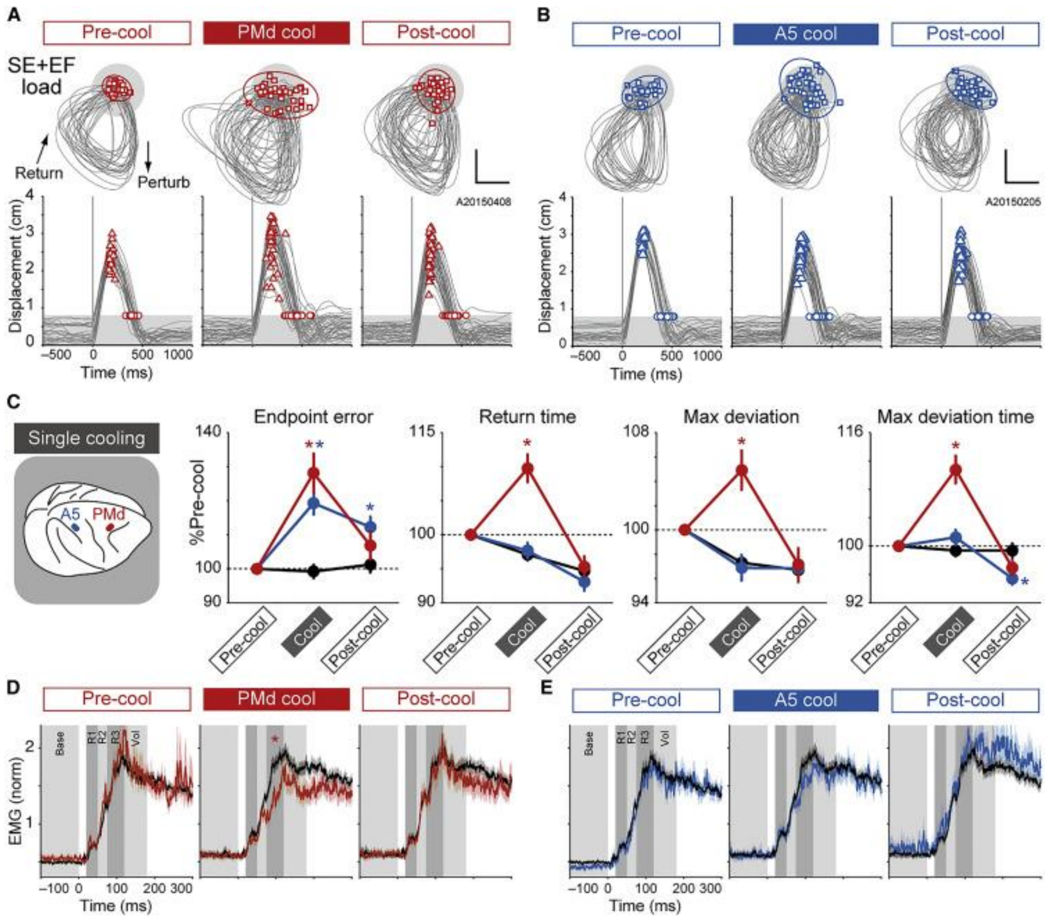

These results are consistent with our hypothesis that PMd cooling impairs the feedback control policy.

When A5 was cooled, the endpoint variability was increased similarly to the PMd cooling. However, in contrast to PMd cooling, during A5 cooling monkey was still able to return to the target quickly within 500 ms

Takei et al., Curr Biology 2021

EPFL

Take-home messages

- Many behaviors are generated by pattern generators (locomotion, breathing, …)
- Feedback plays a key role for skilled behavior
- Optimal feedback control theory explains two apparently conflicting features: high accuracy and high variability
- OFC *shows that in the face of uncertainty the optimal strategy is to allow variability in redundant (task-irrelevant) dimensions*
- *From this framework, task-constrained variability, goal-directed corrections, motor synergies … emerge [Todorov 2001 & 2004]*
- Somatosensory, parietal and premotor areas play key roles in feedback base control and motor adaptation – and we have evidence which areas are involved in specific components.
- However, how neural circuits implement these computations remains an open question