Introduction to computational motor control

Olivier White, PhD, Ir

Associate Professor
INSERM - U1093 Cognition, Action, and Sensorimotor Plasticity

Dijon, 2012

This lecture is based on Reza Shadmehr's, Gunnar Blohm and Jörn Diedrichsen's work
Johns Hopkins University, USA
Queens University, CA
University College London, UK

Motor control

A research field that studies:

- How sensory information from the body and the environment are integrated and used to control actions?

- How does an individual use these information to select an appropriate action for a given objective?

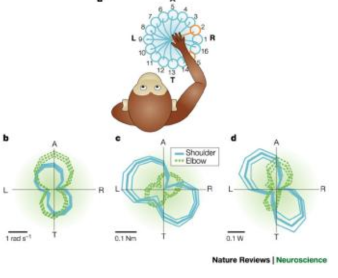

Computational Motor control

A research field that studies:

- How sensory information from the body and the environment are integrated and used to control actions?

- How does an individual use these information to select an appropriate action for a given objective?

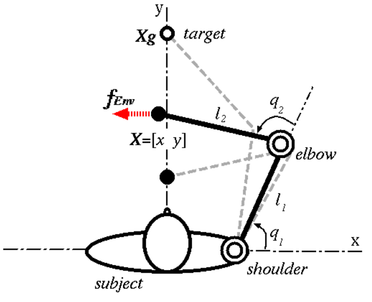

Control theory

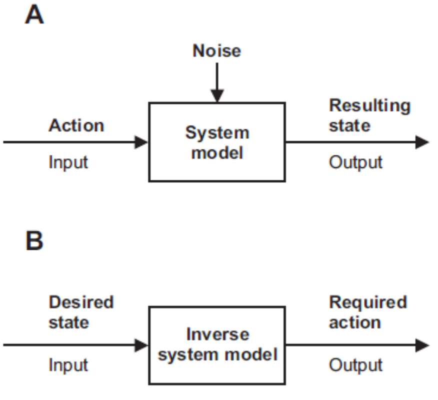

Description in terms of a set of differential equations

$$\frac{dx}{dt} = f(x, u, \eta)$$

We need to have the knowledge of the system model (implicit or explicit)

Links engineering ↔ neuroscience

• Cybernetics
• (Optimal) **control theory**

→ *Many concepts put forward:*

- Negative feedback
- Feedforward control
- Stochastic control
- Control of over/under actuated systems
- State estimation
- Movement planning (with optimization criteria)
- Adaptive control

How far can these concepts inspire our understanding of (biological) motor control?

Links engineering ↔ neuroscience

An <u>interdisciplinary</u> field: psychology, medicine,… **engineering!**

Links engineering ↔ neuroscience

How far can these concepts inspire our understanding of biological motor control?

Not a clear answer…

• Complex multi-layer and distributed systems:

“box charts” approach not always feasible

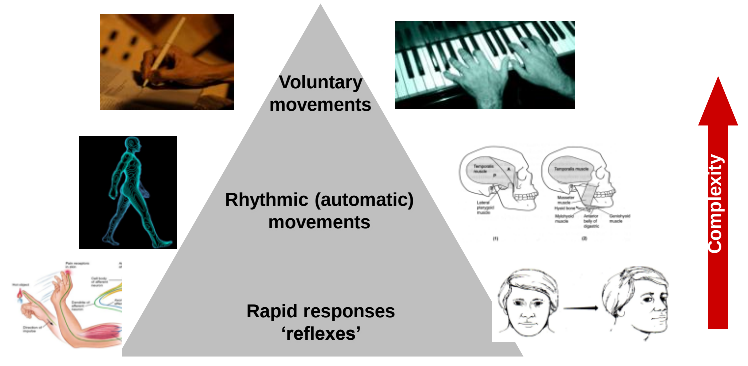

Mixture of experts

Hierarchy in the nervous system:

→ how to choose the right combination according to contexts?

- Divide and conquer
- Mixture of experts

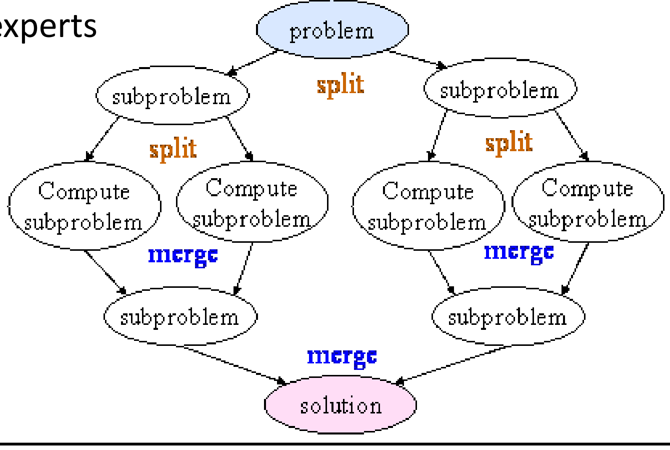

Links engineering ↔ neuroscience

How far can these concepts inspire our understanding of biological motor control?

Not a clear answer…

- Complex multi-layer and distributed systems:

    “box charts” approach not always feasible

10^{10} neurons
10^{14} synapses
100 000 miles of dendrites
different potential tracks:
1 followed by 7 millions miles of 0’s

Our brain is capable of generating more ideas than atoms in the Universe

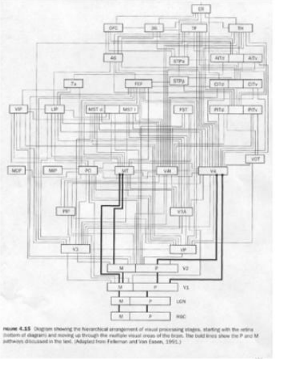

Links engineering ↔ neuroscience

How far can these concepts inspire our understanding of biological motor control?

Not a clear answer…

- Complex multi-layer and distributed systems:

    “box charts” approach not always feasible

- **Several seconds**: complex movement
- **1 s**: elementary actions *(motor cortex, basal ganglia)*
- **500 ms**: decomposition in trajectories
- **10 ms**: decomposition in positions, velocities and muscle forces *(CB)*
- **1 ms**: impulsion trains to motoneurons

Links engineering ↔ neuroscience

How far can these concepts inspire our understanding of biological motor control?

- **Evolution** came up with strategies developed in a long time scale not yet transferred to robotics/engineering

- A striking similarity between control theory and motor control:
  Time-delayed systems lead to the notion of **internal models**

- **Inspiration** of what approaches are possible to solve a certain problem

- Provides a **baseline** of what optimal performance can be achieved in theory

- **Direct applications** are possible in some fields (e.g. oculomotor)

Model

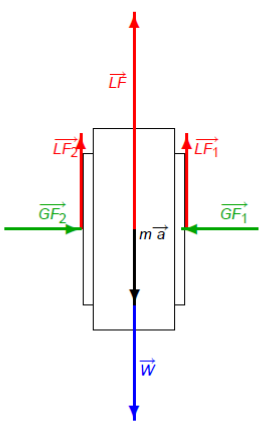

Model

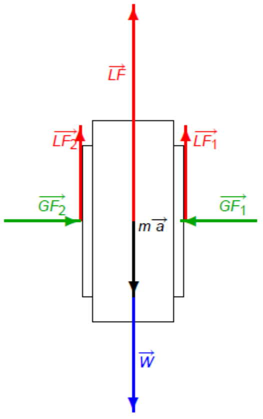

- $\overrightarrow{GF}$: Grip Force
- $\overrightarrow{LF}$: Load Force
- $\vec{W}$: Weight

Relationships

- $\|\overrightarrow{LF}\| \leq \mu \cdot \|\overrightarrow{GF}\|$
- $\vec{W} = m \cdot \vec{g}$
- $\overrightarrow{LF} = \vec{W} + m \cdot \vec{a}$

Typical result derived from a simple approach

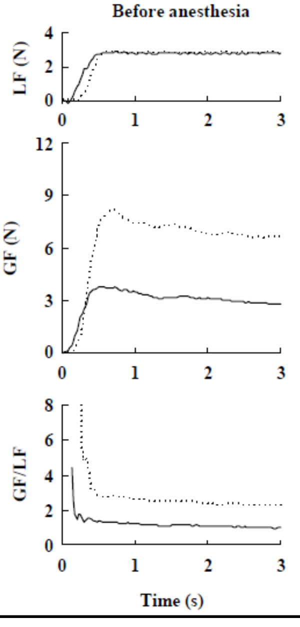

Typical result derived from a simple approach

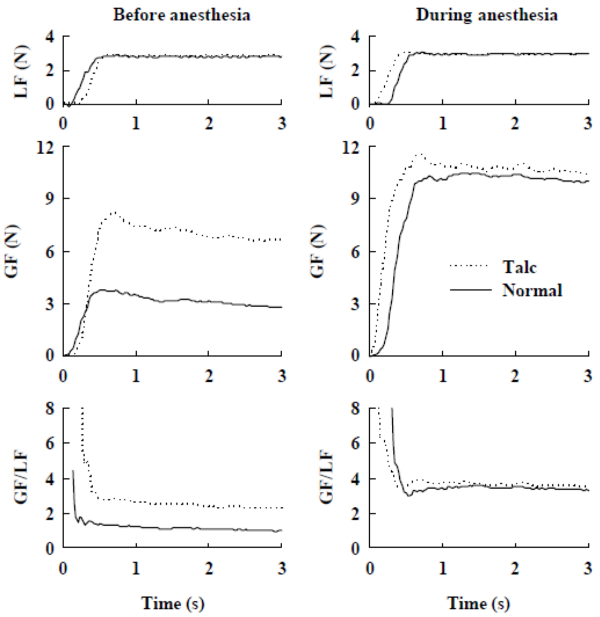

Model => (over)simplification

Tribology:

Research field that studies friction between bodies (in particular, non elastic).

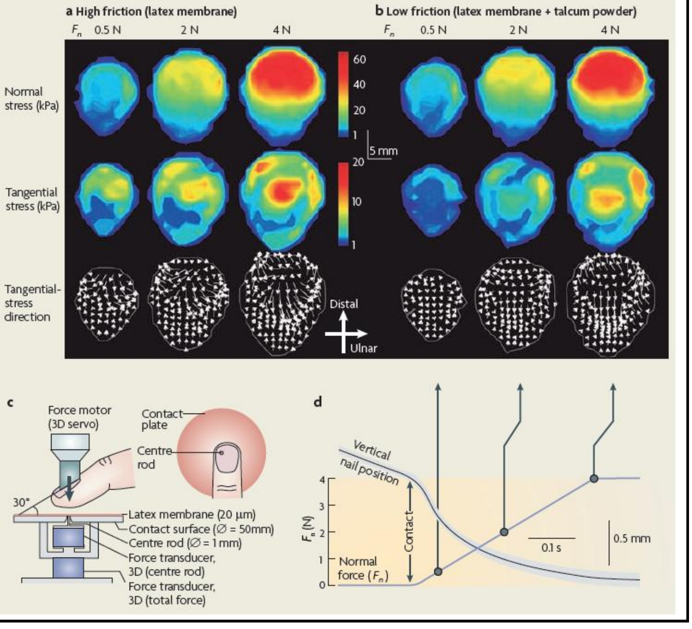

Computational Motor control

Huge interest!

Formalizing mechanisms underlying action can help to shed light on:

- What drives actions?
- How does (re)learning work?
- How do we consolidate what we've learned?
- How to design optimal rehab protocols?
- Predictive tool

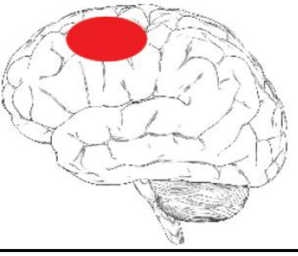

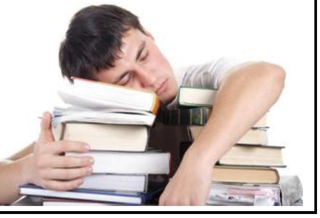

So far…

- Heritage from control theory
- Real system is complex…
- … as are fine models!

- But great interest

Contents

1. **Introduction and intuition of computational motor control**
   - From engineering to neuroscience, and back
   - Intuition
   - Box chart approach and some physiological evidence
   - Internal models
   - Learning mechanisms
   - Bayesian brain
   - Open questions

2. **Kalman filtering**
   - Rationale
   - (simple) statistics refresher
   - Derivation of the Kalman Filter
   - Examples in sensorimotor control

3. **Stochastic Optimal Feedback Control**

600 millions years ago…

What separates plants and animals is that animals can move. To control movement, multi-cell organisms developed a nervous system.

Development of the nervous system began when multi-cell organisms began to move.

The sea squirt: In larval form, is free swimming and is equipped with a brain-like structure of 300 cells.

Upon finding a suitable substrate, it buries its head and starts absorbing most of its own Brain and looses its ability to move.

Box chart

Figure 1

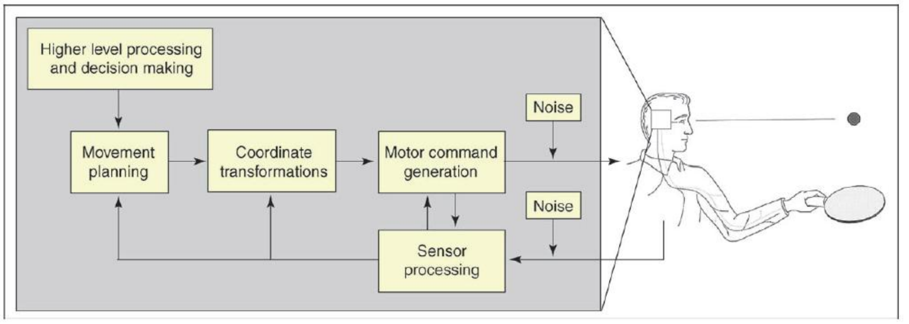

Sketch of a generic motor control diagram, typically used in robotics research, that can also function as a discussion guideline for biological motor control.

1. Decision making
2. Motor planning
3. Coordinate transformation
4. Plant to motor commands
5. Processing of sensory information

Three examples of computational models

1. **Feedback control**: stretch reflex with <0 feedback

- Muscle spindle report muscle extension to the spinal cord
- That info is fed back in the motor command (1! Synapse)

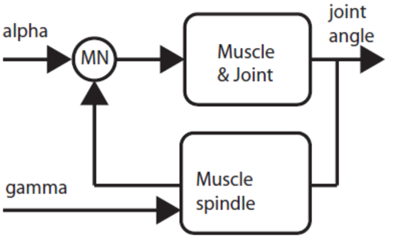

Three examples of computational models

2. Optimality approach: minimum jerk

Flash & Hogan, J Neurosci, 1987

- Observation: straight movements in extrinsic space
  => bell-shaped velocity profiles

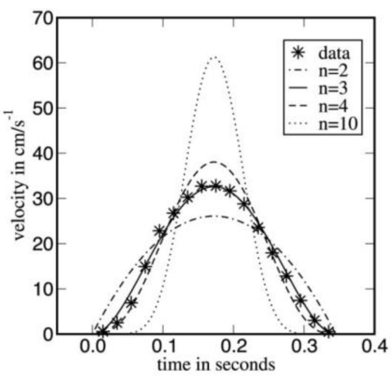

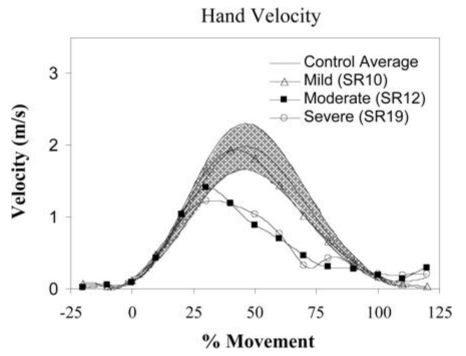

- System is optimal in some sense: $J = \int_0^T \left(\frac{d^3x}{dt^3}\right)^2 dt$

Three examples of computational models

3. Error-based learning

- Observed error (feedback) drives learning
- Gradient descent to minimize – even cancel – error signal
- Trial by trial adaptation
- Typical experimental approach:
  Visuomotor rotations or reaching in force fields

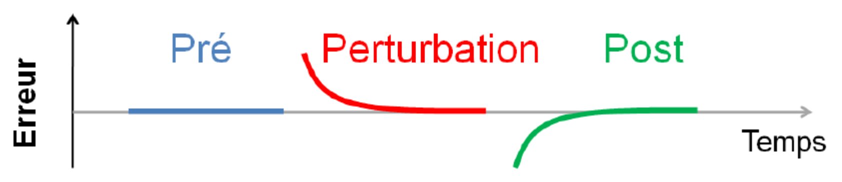

- Will be detailed later

Three examples of computational models

3. Error-based learning

- Typical experimental approach: visuomotor rotations

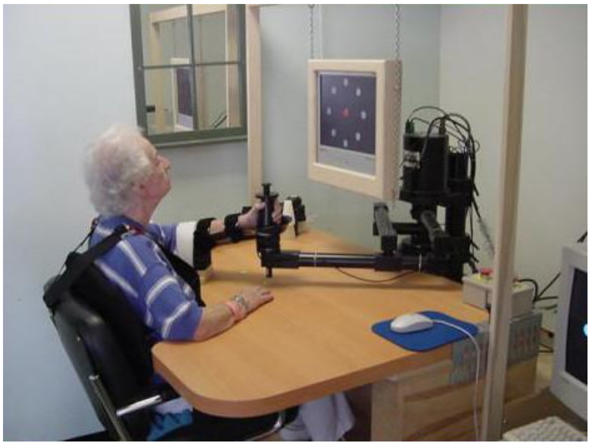

Three examples of computational models

3. Error-based learning

- Typical experimental approach: visuomotor rotations

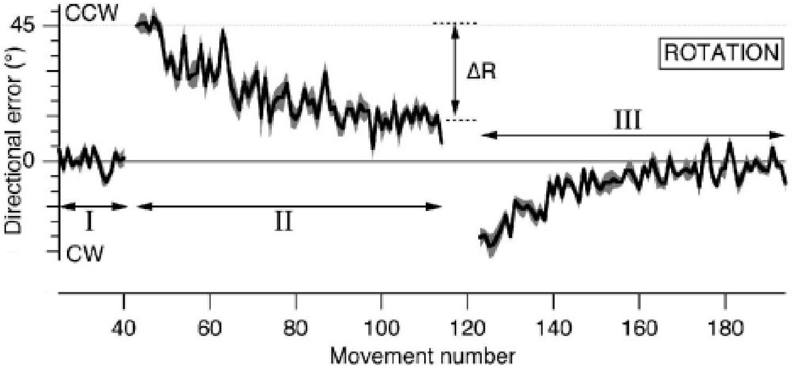

- Notion of catch trials

Again: physiologically plausible?

Internal models are implemented in the primary motor cortex and in the cerebellum

A. Forward (direct) models
B. Inverse models

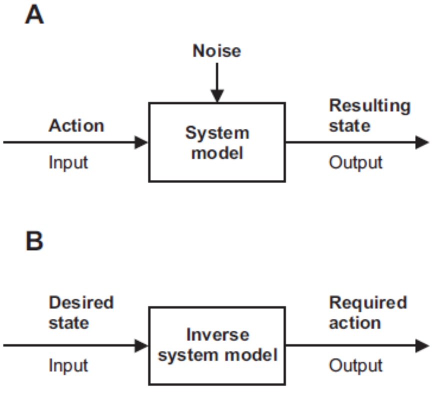

Behavioral evidence: forward model

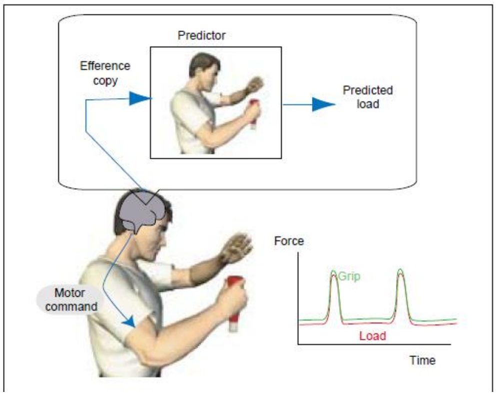

Behavioral evidence: forward model

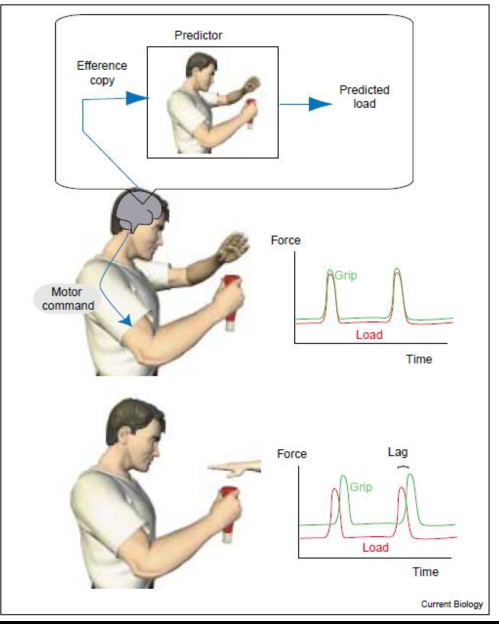

How to deal with the many contexts?

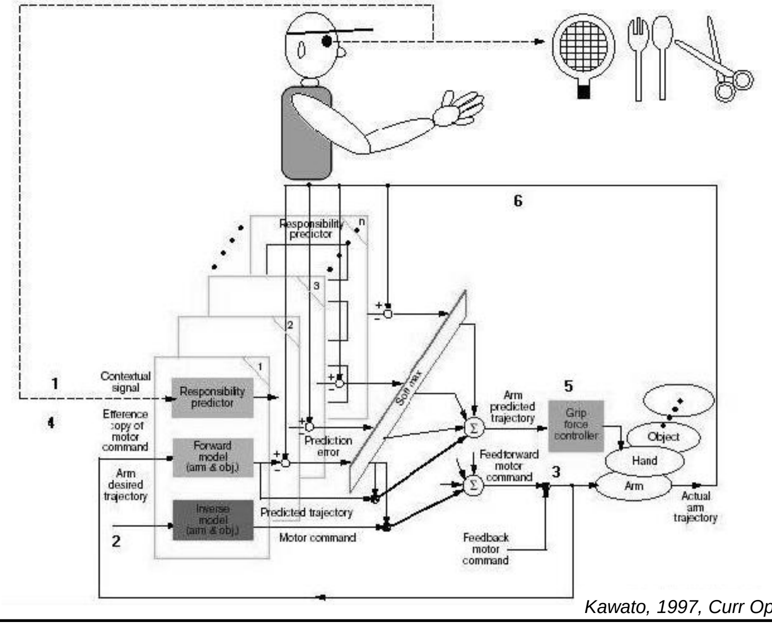

Kawato, 1997, Curr Opin Neuro

Noise, noise, noise

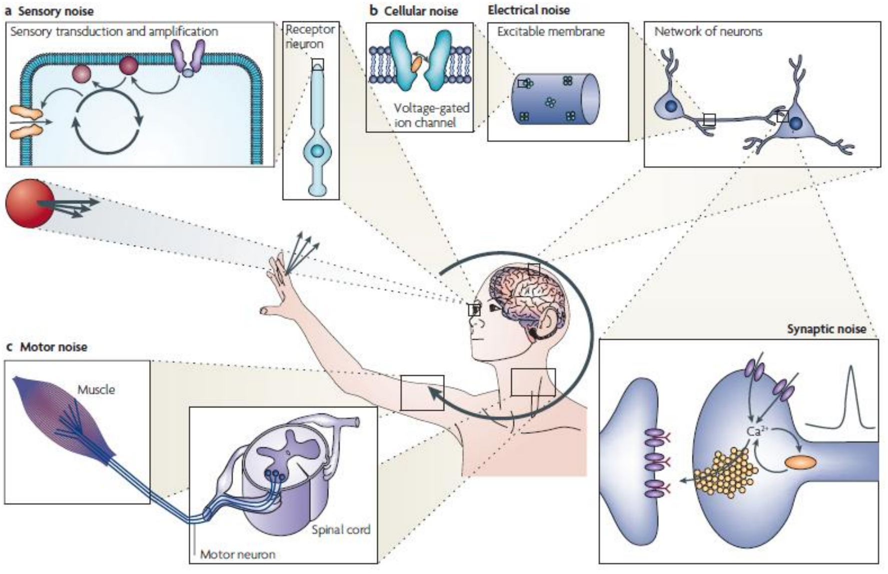

Noise, noise, noise

Strategies developed by the motor system to minimize the effects of noise on movements?

- Generate smooth movements!
  → signal-dependent noise

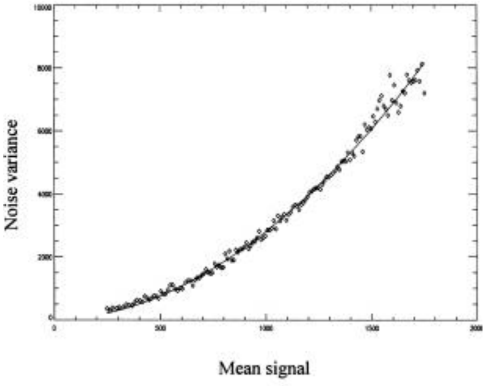

(b)

Evidence of sensory noise

The SD of noise grows with mean force in an isometric task. Participants produced a given force with their thumb flexors. In one condition (“voluntary”), the participants generated the force, whereas in another condition (“NMES”) the experimenter electrically stimulated their muscles to produce force. To guide force production, the participants viewed a cursor that displayed thumb force, but the experimenters analyzed the data during a 4-s period in which this feedback had disappeared. **A**. Force produced by a typical participant. The period without visual feedback is marked by the horizontal bar in the 1st and 3rd columns (top right) and is expanded in the 2nd and 4th columns. **B**. When participants generated force, noise (measured as the SD) increased linearly with force magnitude. Abbreviations: NMES, neuromuscular electrical stimulation; MVC, maximum voluntary contraction. Jones et al. 2002 JNP.

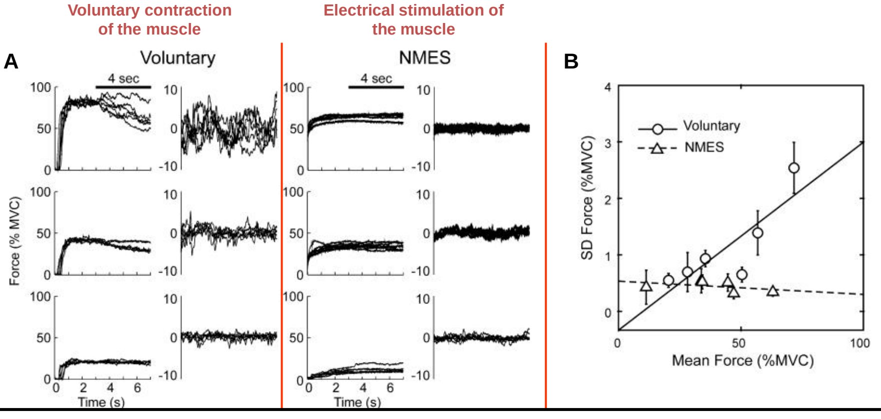

Integrated approach

Stochastic optimal control theory

- Promising but complex (not established for non linear systems)
- Model plant, environment and noise
- Account for a large body of experimental data
- Exists limitations… why does variability decrease if co-contraction?

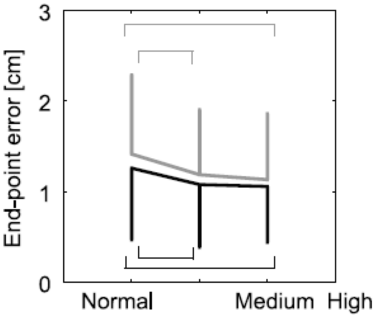

Osu et al., JNP, 2002

Optimal control: core equations

System

$$\begin{aligned} x_{k+1} &= Ax_k + Bu_k + \xi_k + \sum_{i=1}^{n_c} \varepsilon_{i,k} C_i u_k, \\ y_k &= Hx_k + \omega_k + \sum_{i=1}^{n_d} \delta_{i,k} D_i x_k, \end{aligned}$$

Cost to go

$$J_k(x_k, u_k) = x_k^T Q_k x_k + u_k^T R u_k, \qquad k = 1, 2, \ldots, N-1,$$

Total cost

$$J = E\left[J_N(x_N) + \sum_{k=1}^{N-1} J_k(x_k, u_k) \right].$$

$$\boxed{\begin{aligned} u_k &= -\left(B^T U_{k+1} B + R + \Gamma_{k+1}\right)^{-1} B^T U_{k+1} A \hat{x}_k, \\ &:= -L_k(x_k - e_k). \end{aligned}}$$

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

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

$$S^{\mathrm{x}}_t = Q_t + A^{\mathsf{T}} S^{\mathrm{x}}_{t+1} (A - BL_t) + \sum_i D_i^{\mathsf{T}} K_t^{\mathsf{T}} S^{\mathrm{e}}_{t+1} K_t D_i; \qquad S^{\mathrm{x}}_n = Q_n$$

$$S^{\mathrm{e}}_t = A^{\mathsf{T}} S^{\mathrm{x}}_{t+1} B L_t + (A - K_t H)^{\mathsf{T}} S^{\mathrm{e}}_{t+1} (A - K_t H); \qquad S^{\mathrm{e}}_n = 0$$

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

How to choose a movement?

State i

? **?**

State i+1 **State i+1**

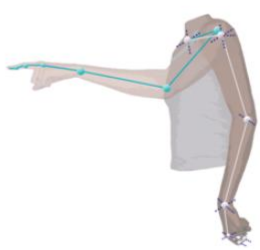

Effort!

State i

State i+1

State i+1

Reward!

State i

State i+1

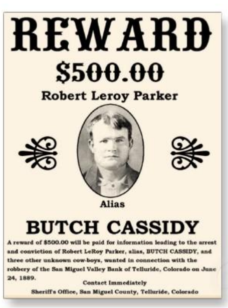

State i+1

Types of rewards

Vegetative needs of individual subject
- Food
- Liquid

Reproduction of genes
- Sex

Higher and mental rewards
- Money
- Novelty and challenge
- Taste, pleasantness, beauty
- Acclaim and power
- Altruistic punishment
- Territory and Security

How to achieve a more rewarding state?

Translating goals into motor commands

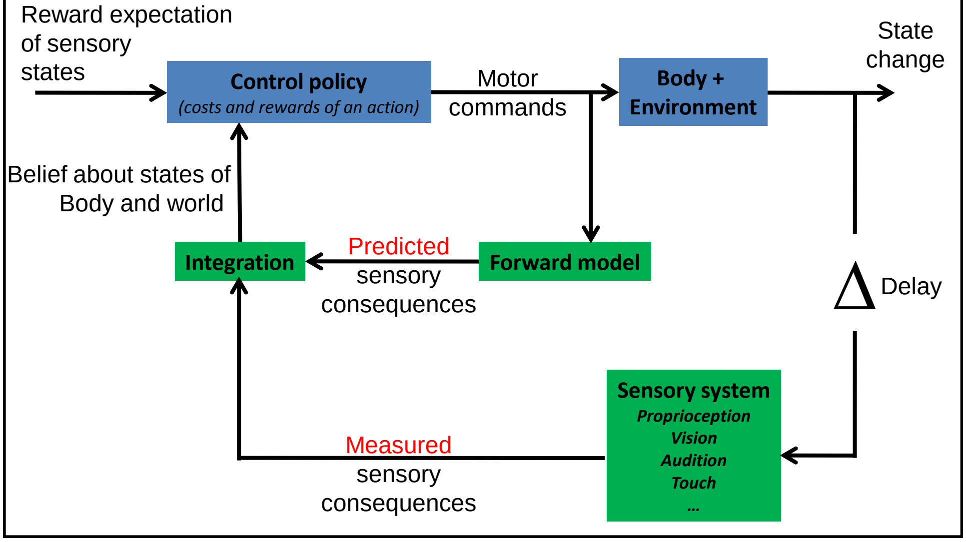

What to do - locations

Selection of action based on a value function associated with locations on a spatial map
Associating reward to a location on a spatial map depends on the **hippocampus**

Mouse is released into a pool of water from any starting point. A platform is positioned in a specific location just below the water line. The platform is always at the same location.

The normal rat can learn to locate that position with respect to the cues that surround the pool. This requires learning a spatial map of where the platform is located with respect to the surrounding visual cues.

With repeated swims, the animal learns a spatial map and find the platform regardless of where he is released into the water. If the platform is removed, the normal animal will spend most of his time searching in the quadrant where the platform should be.

Learning of this sort of spatial map depends on the **hippocampus**. If a genetically altered rat with a malfunctioning hippocampus is given the same training, he will not learn the spatial map and will spend equal time in each quadrant.

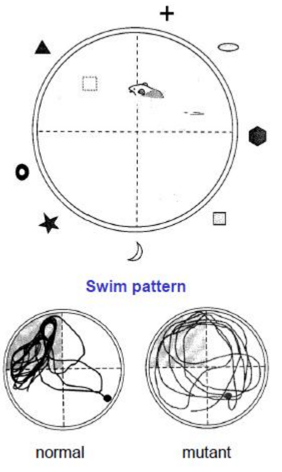

What to do - objects

Selection of action based on a value function associated with objects
Associating reward to stimuli regardless of their location depends on the **basal ganglia**

In this task, there are two platforms. One that is large enough for the mouse to mount, and one that is too small. Both have a visual cue associated with them. The platforms may be positioned in any quadrant. Animal performs 8 swims per day for 15 days.

The animal needs to learn that the green ball, and not the other ball or surrounding cues, is important and that it indicates location of the platform. He needs to ignore the memory of the spatial position of the platform in the previous trial.

Every time the animal tries to mount the platform below the gray ball, an error is recorded.

Lesion in the caudate severely disrupts the ability of the animal to recognize that across repeated trials, the only cue that consistently predicted platform location was the green ball.

Lesion in the hippocampus has no effect. Because the spatial cues are irrelevant to finding the platform, the animal behaves normally.

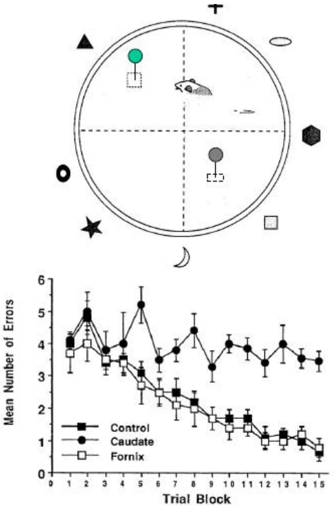

Dopamine neurotransmitter

Dopamine releasing neurons have their cell bodies in the brain stem.

<u>Neurotransmitter</u>: chemical that sends signals from one cell to another cell.
Reward increases dopamine level in the brain.

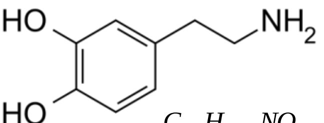

$C_8H_{11}NO_2$

Five dopamine receptors: D1 to D5.

They project to 3 main areas:
- the **striatum** (nigrostriatal tract)
- the **hippocampus** (mesolimbic tract)
- **prefrontal cortex** (mesocortical tract)

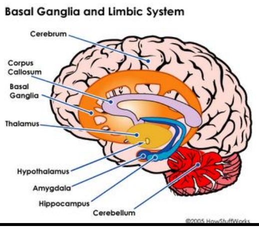

Where dopamine is manufactured

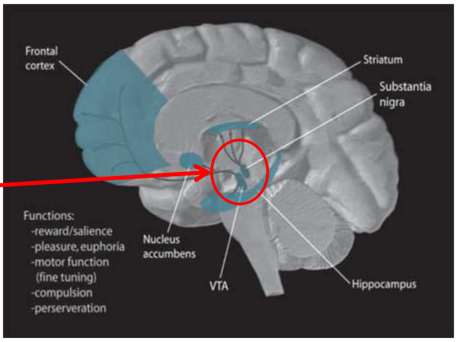

Dopamine neurotransmitter

Dopamine is the main teaching signal for the basal ganglia.
80% of the brain's dopamine is in the basal ganglia.

This figure shows a dopamine neuron in substantia nigra that responded to unexpected rewards that occurred in association with a visual cue.

As the probability of reward increased, the cell's response after the reward decreased, responding instead to the visual cue, which now predicted the reward.

Cells respond to reward, when probability=0.
Cells respond to visual cue, when probability=1
(visual cue is now associated to reward).

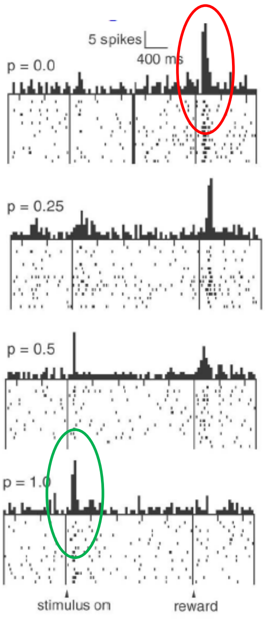

A word on eye movements

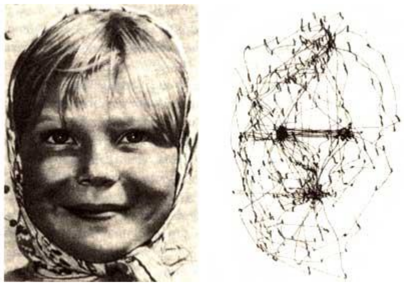

Saccades

Smooth pursuit

Rewarding behaviors

Reward implies faster movements and increased activity in the caudate nucleus of the basal ganglia.

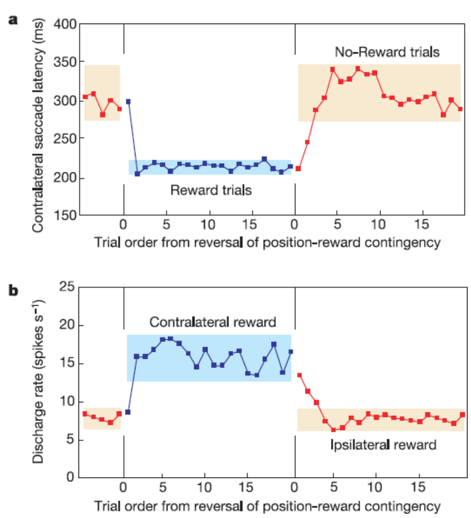

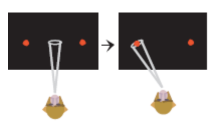

Activity of a neuron in the caudate nucleus as a monkey made saccades to a leftward target.

At the beginning of the plot, the leftward saccade did not produce reward, but a rightward saccade did. The monkey reacted to the onset of the left target with a response latency of approximately 300ms, and the cell discharged at about 8 impulses per second.

On the trial following the left vertical line (the first 0 on the *x-axis),* the leftward saccades produced a reward. After the monkey experienced this “contingency” once, its reaction time quickened by nearly 100 ms and the cell’s discharge rate nearly doubled.

Generating motor commands

Choosing the best movement that produces the <u>largest reward</u> while <u>minimizing motor cost</u> (effort)

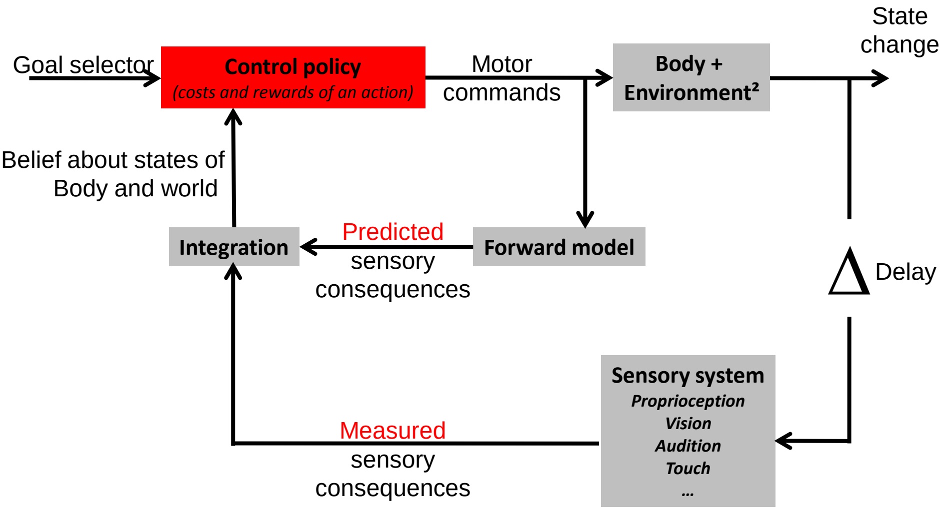

Evolution of control policies

1928
1.59m

1936
2.03m

1968
2.24m

Ethel Catherwood (Canada)
gold medal winner

Cornelius Johnson (USA)
gold medal winner

Dick Fosbury (USA)
gold medal winner

Eye movements

Two types of eye movements:
smooth pursuit and **saccades.**

Main sequence for saccades: robust relationship between
- Amplitude
- Duration
- Peak velocity

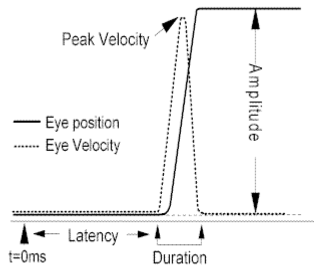

Linear

$$D = 2.21A + 21$$

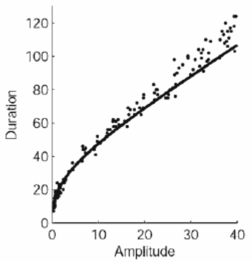

Exponential

$$PV = 134.6A^{0.38}$$

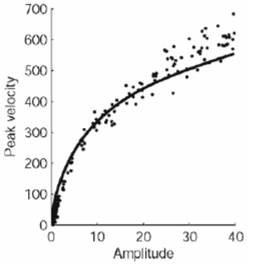

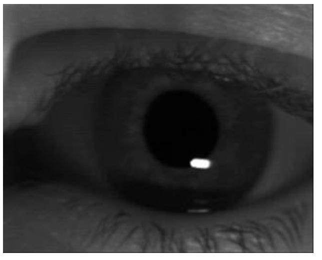

Eye movements

Costs involved in making a saccade:

1. There is something interesting off to one side of my fovea. I want to fixate this interesting thing, and I incur a <u>cost by not looking at it</u>.

2. During the eye movement, I am effectively <u>blind</u>. The eye movement should complete as soon as possible.

3. A fast movement requires large motor commands. The larger the motor command, the larger the <u>noise</u> in those commands. Noisy motor commands produce inaccurate movements.

<u>Policy</u>: Try to find a way to move the eyes to the target as soon as possible, while minimizing the motor commands.

Internal models

Predicting consequences of motor commands: internal models

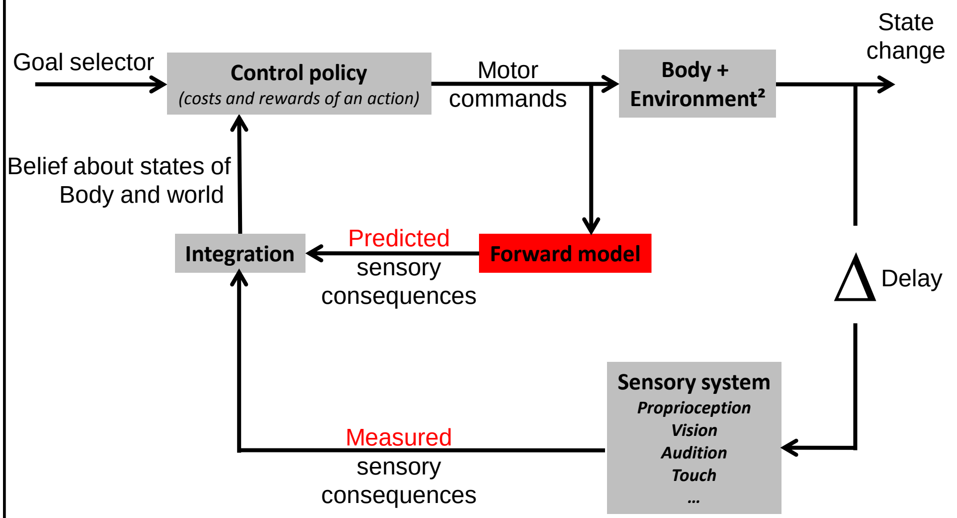

Internal models

Prediction

Reaction time

Internal models: eye/hand coordination

When subjects are asked to use their eyes to track their hand during an active movement, the eyes **look ahead** by about 200ms.

Active trials: subject move their hand but cannot see it.

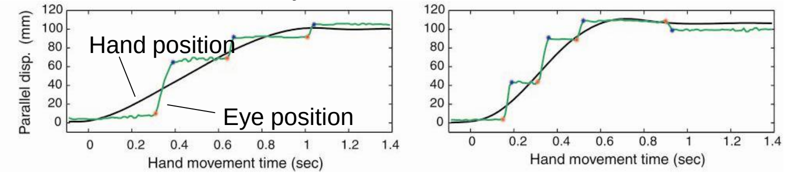

About 200ms after a saccade, the hand reaches where the eyes are looking.

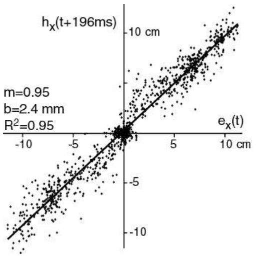

Internal models: eye/hand coordination

When subjects are asked to use their eyes to track their hand during a passive movement, the **eyes lag behind** the hand.

Passive trials: robot moves the hand.

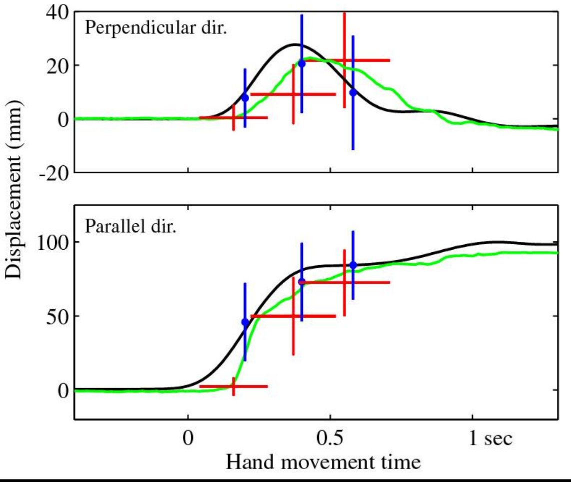

Internal models: cerebellum

Predicting the sensory consequences of motor commands depends on the cerebellum

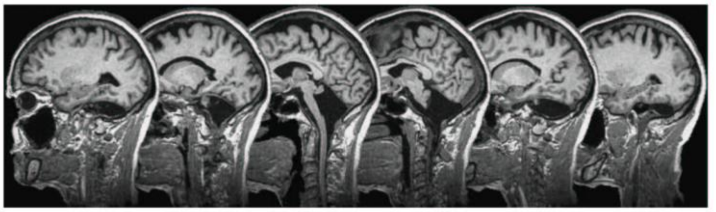

Patient HK (cerebellar agenesis)

Experimental conditions:

1) Experimenter drops the ball at random times. This tests the sensory **feedback** pathways.
2) The subject holds the ball and drops it. This tests the **predictive** pathways.

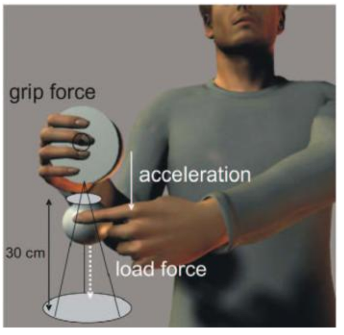

Nowak et al. (2007) Neuropsychologia

Internal models: cerebellum

Feedback: normal motor responses to perturbation

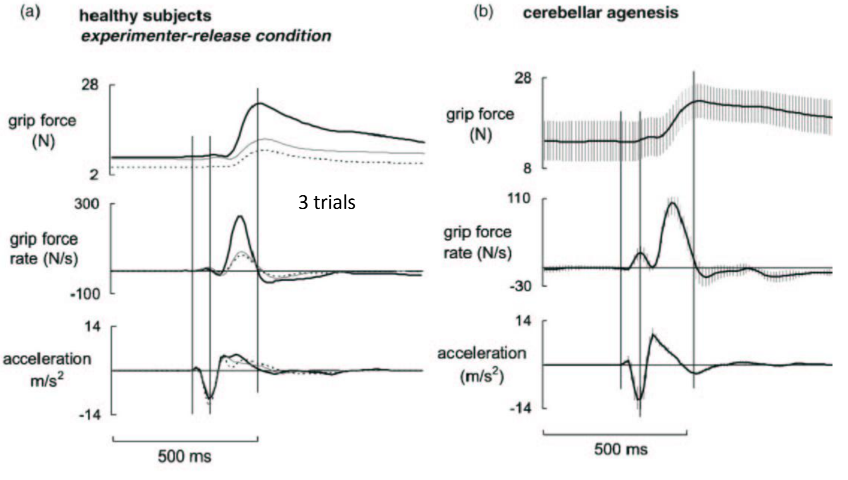

Nowak et al. (2007) Neuropsychologia

Internal models: cerebellum

Feedforward: **no ability to predict** sensory consequences of self-generated motor commands

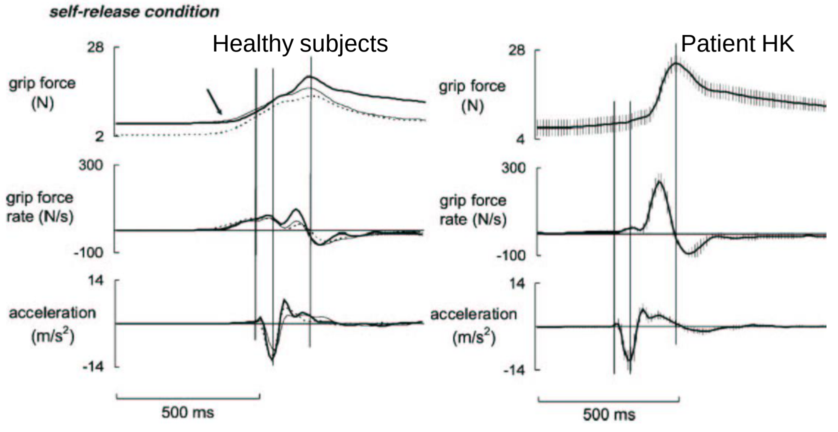

Nowak et al. (2007) Neuropsychologia

… so far…

- Implementation of smart mechanisms need a CNS
- Simple (jerk) to complex (OFC) approaches
- How to deal with the myriad contexts, noise, delays…
- Behavioral and clinical evidence of IM
- What drives action (reward, effort)
- Evolution of control policies

- But still: how do we learn??

Motor learning

Internal models must cope with changes all the time

How do we learn??

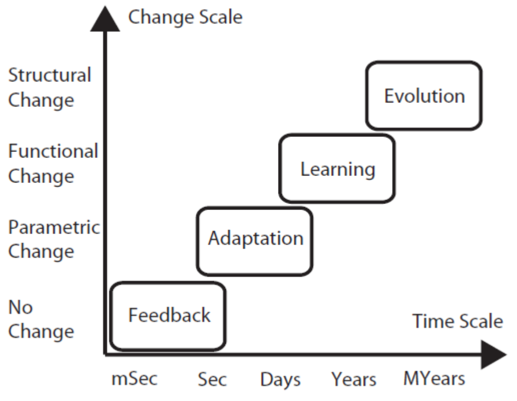

Karniel, 2011, Front in Neurosci

Motor learning

Internal models must cope with changes all the time

How do we learn flexibly and efficiently??

Generalization **Transfer** **Consolidation**

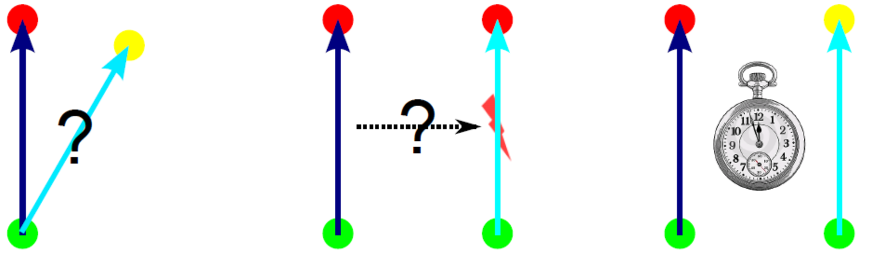

Generalization formalized with RBF: Purkinje cells fire at preferred spatial directions

Motor learning

Internal models must cope with changes all the time

How do we learn flexibly and efficiently??

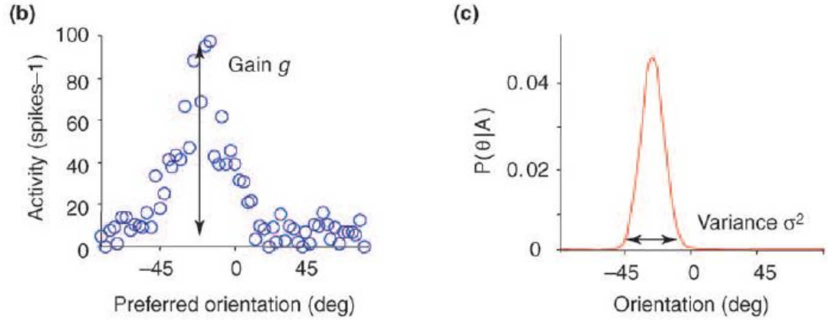

Generalization formalized with RBF: Purkinje cells fire at preferred spatial directions

Knill and Pouget, Trends Neurosci, 2004

Motor learning

Again, ambiguity because of redundancy

- Coordinate transformation from Cartesian space to muscle space
- #ext dof << #int dof (muscles: 100:1, joints: 30:1)
- Degree-of-freedom problem: Bernstein 1967

- How to interpret an error signal?

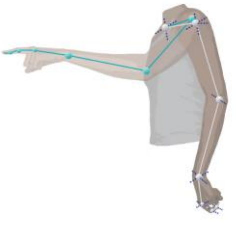

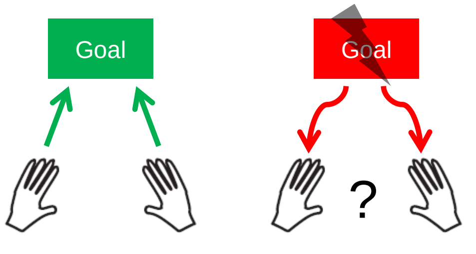

Credit assignment problem

- This question can be translated experimentally in a redundant task

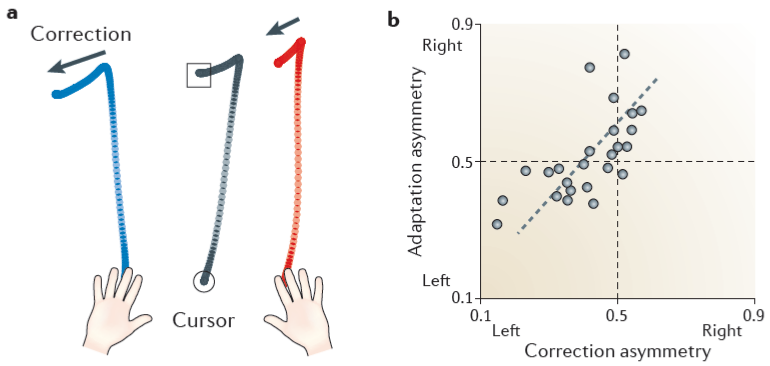

White and Diedrichsen., Curr Biol, 2010.

Credit assignment problem

- In <u>complex high level movements</u>, causality is sometimes hard to infer: Reinforcement learning may be a better strategy

- If a tennis player systematically misses her serve, which is responsible?
  - fatigue?
  - Racket too lose?
  - Hit too early
  - …

Reinforcement learning

- In complex high level movements, causality is sometimes hard to infer: Reinforcement learning may be a better strategy
- Associate an action to a reward provided by the system
- Goal: max discounted reward over time (dopamine!)
- Needs other information such that relative success/failure of movement
- Because RL is inherently unsigned, exploration is important and slows the learning process down

Motor learning

Again, ambiguity because of redundancy

- **Freeze or sleeve**: decrease #dof by linking unrelevant dof together
  But… cannot explain why on average the variability in internal coordinates > variability in task-space coordinates

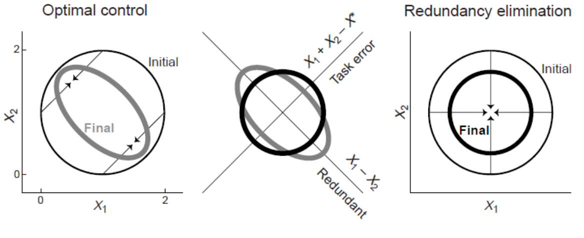

Fig. 1. Redundancy exploitation. The system described in the text (with $X^* = 2$, $a = \sigma = 0.8$) was initialized 20,000 times from a circular two-dimensional Gaussian with mean (1, 1) and variance 1. The control signals given by the two control laws were applied, the system dynamics simulated, and the covariance of the final state measured. The plots show one standard deviation ellipses for the initial and final state distributions, for the optimal (left) and desired-state (right) control laws. The arrows correspond to the effects of the control signals at four different initial states (scaled by 0.9 for clarity).

Todorov and Jordan, Nat Neurosci, 2002

Motor learning

Again, ambiguity because of redundancy

- **Operational space control** and **Inverse kinematics control**
  Desired movement in task space can be transformed in desired movement in internal space
  $\rightarrow$ Need Jacobian of the kinematics of the movement system

Partition of dof:

- A subset of (chosen) dof are constrained
- Unconstrained manifold: dof used to optimize other criteria (energy expenditure, joint limits etc)

Error based learning

Generic approach:

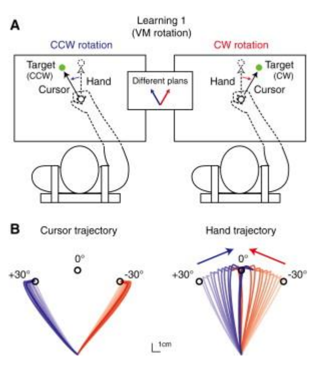

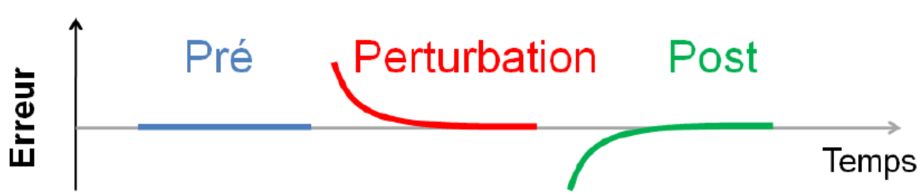

Error based learning

- Importance of error magnitude and direction
- Need to estimate gradient of error wrt each component of the motor command

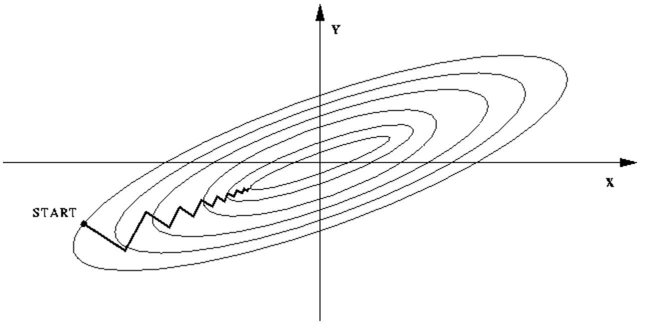

- So powerful that adaptation occurs trial by trial even if explicit instruction NOT to adapt

Experimental perturbations

Force fields stimuli allow simulation of new objects

- Elastic $\vec{F} = k\vec{x}$

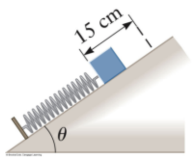

- Viscous $\vec{F} = k\dot{\vec{x}}$

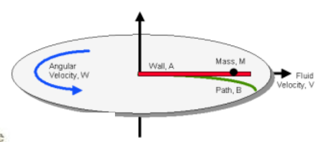

- Inertial $\vec{F} = k\ddot{\vec{x}}$

- Gravitationnal $\vec{F} = f(g)$

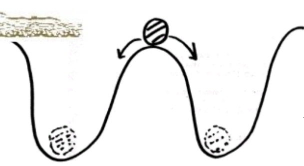

- Unstable $\vec{F} = f(\cdots)$

- Combinations of these types

Error based learning

- Problems of local optimums…

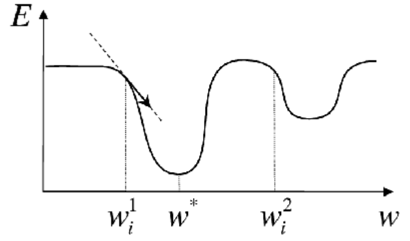

- Explains many behavioral results
  Visuomotor, force fields, grip force

- The CB and neo-cortex play a crucial role in error based learning

Impedance control

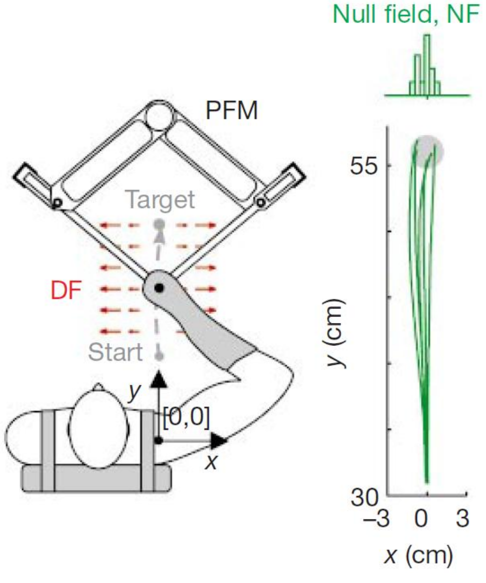

Impedance control

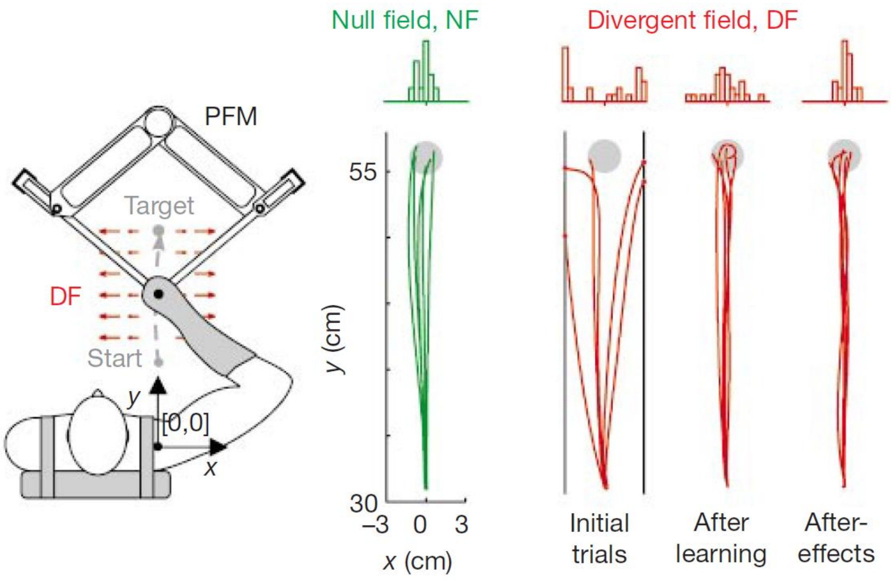

Use-dependent learning

→ **How to improve further??**

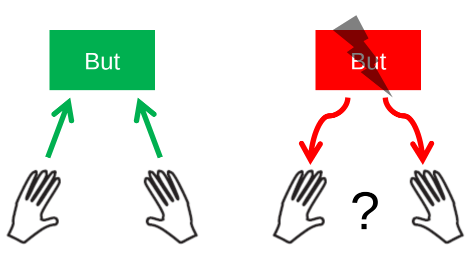

2 Parallel mechanisms

- **Error-based**
- **Use-dependent**

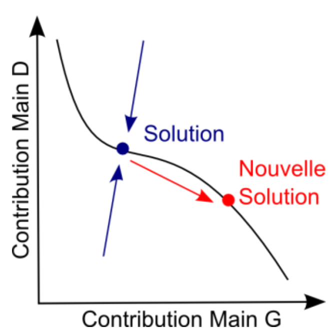

Use-dependent learning

- State of the motor system can change through pure repetitions even in absence of outcome

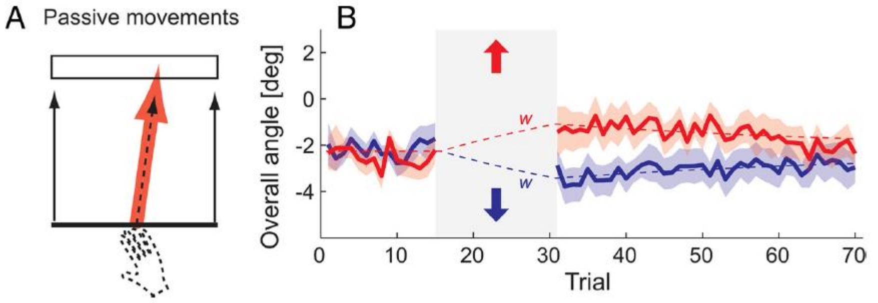

Diedrichsen et al., JN, 2010.

Use-dependent learning

- Combinations of different learning mechanisms occur

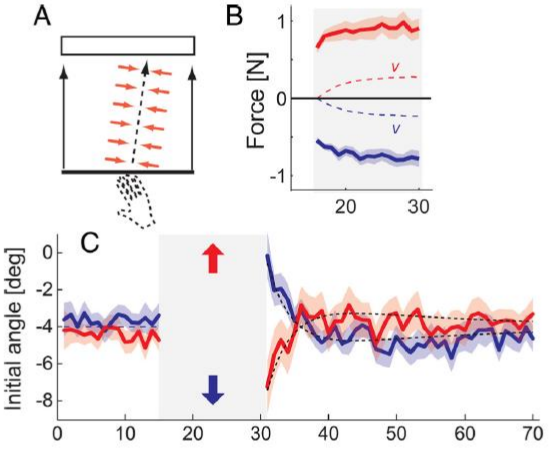

Diedrichsen et al., JN, 2010.

Structural learning

- Combinations of different learning mechanisms occur

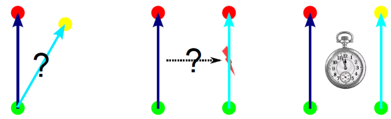

- <u>Parameters</u> of a specific task can be learned quickly through e.g. Error-based learning
- Expose someone to RND perturbation sharing the same structure

→ ***What about the structure of the task?***

A change of gravity?

Ground-based facilities

Table A.1: Comparison between microgravity platforms.

Platform	μ-g level (g)	Duration	Volume (m^3)	Control
Drop towers	10^{-3} - 10^{-6}	< 5 s	< 1	indirect
Parabolic flights	10^{-2} - 10^{-3}	20-25 s	> 10	direct
Sounding rockets	10^{-4} - 10^{-5}	5 - 13 min	< 1	indirect
Recoverable capsules	$\leq 10^{-5}$	weeks	> 1	indirect
Manned orbital platform (ISS)	10^{-2} - 10^{-5}	weeks - years	> 1	direct

A good approach to test structural learning

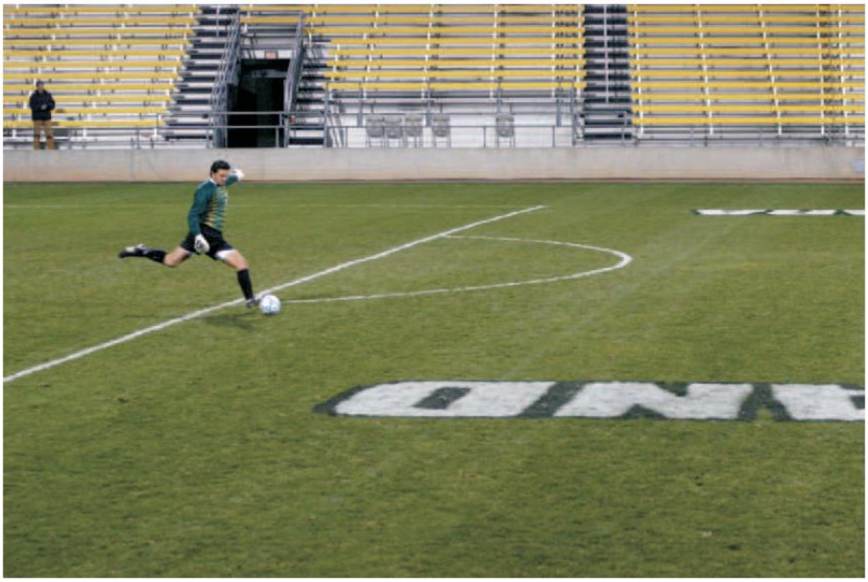

A good approach to test structural learning

Parabolic flights

Plane	weight	W	$=$	mg
	thrust	T	$=$	ma
Aerodynamics	drag	D	$=$	$\rho \frac{v^2}{2} A C_D(\alpha)$
	lift	L	$=$	$\rho \frac{v^2}{2} A C_L(\alpha)$

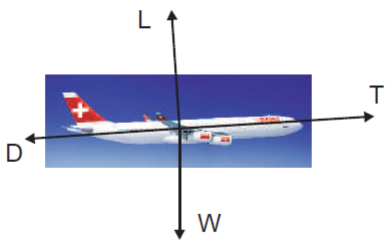

Parabolic flights

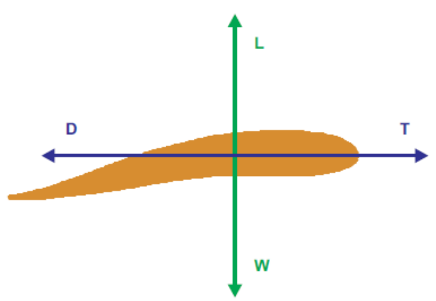

$$\vec{W} + \vec{L} + \vec{D} + \vec{T} = \vec{0}$$

Parabolic flights

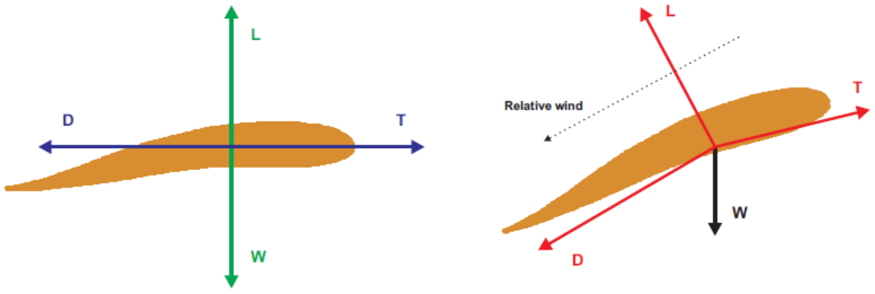

$$\vec{W} + \vec{L} + \vec{D} + \vec{T} = \vec{0}$$

$$\vec{L} + \vec{D} + \vec{T} = \vec{0}$$

Parabolic flights

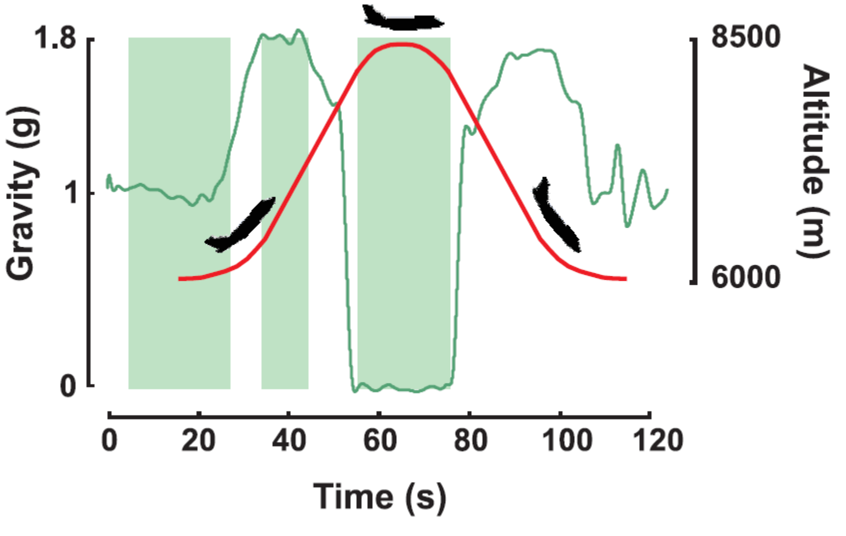

PARABOLAS SEQUENCE

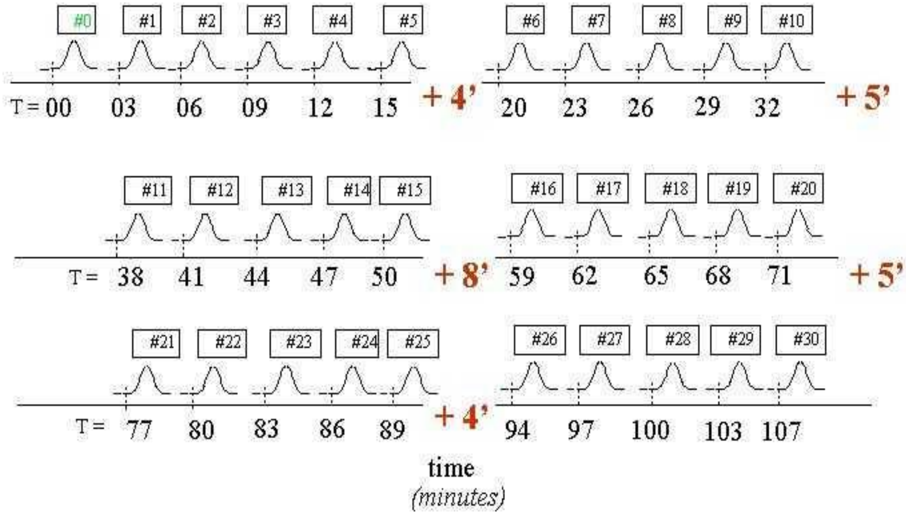

Range of gravity explored

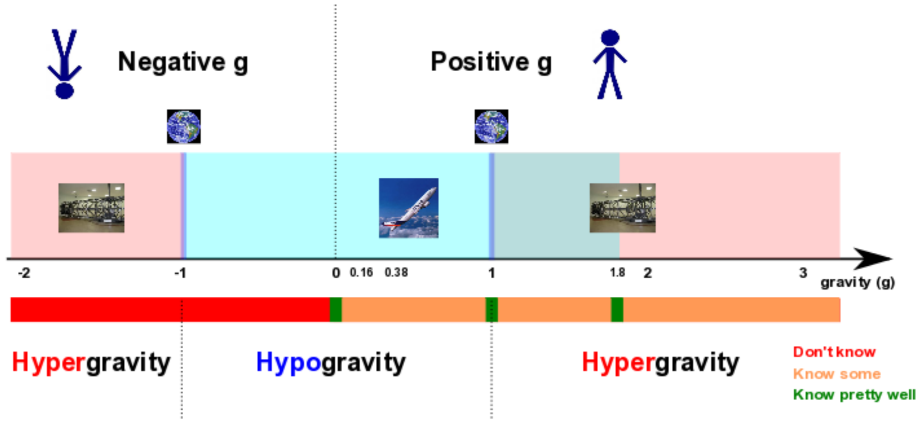

Centrifuges

Can we learn anything?

- There are conditions in which learning is impossible (or very difficult)
- <u>Examples</u>: conflicting force fields, force fields varying explicitly on time

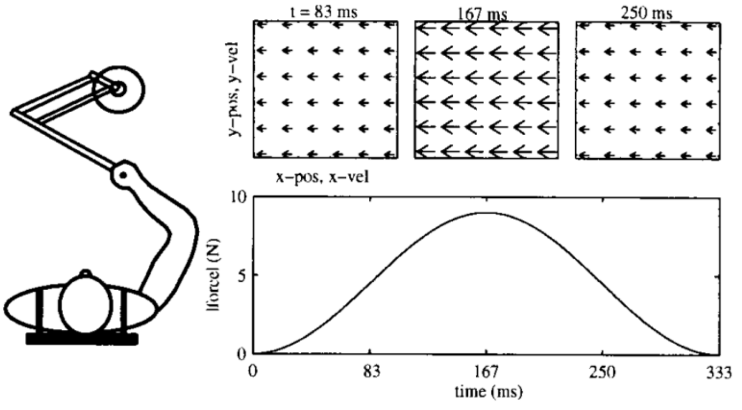

Can we learn anything?

- Relevant cue for conflicting FF learning

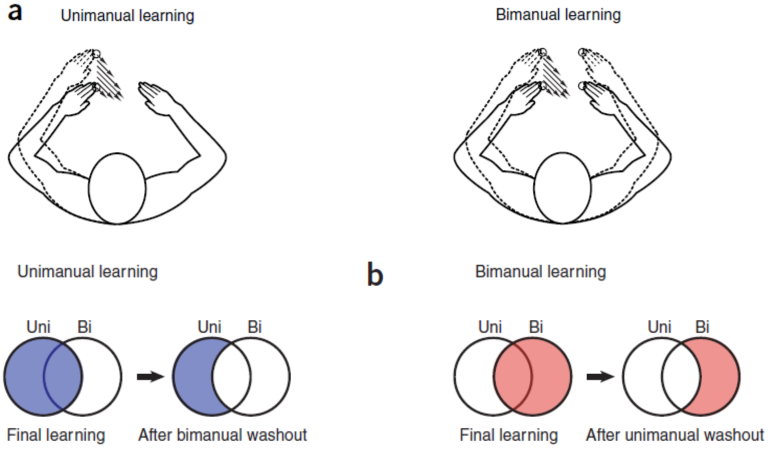

The Bayesian brain

- Sensorimotor memory is a inestimable source of knowledge
- How to update that *database*?
- Real world is uncertain, feedbacks are in different modalities, guesses are necessary… probabilities!

- **Bayesian inference**: optimal integration of prior knowledge and sensed noisy information

$$\boxed{P(x|y) = \frac{P(y|x)P(x)}{P(y)}}$$

- Many artificial and biological neural nets can be interpreted as Bayes optimal signal processing systems, despite possessing neither explicit knowledge of Bayes rule nor knowledge of the probability distribution of the involved variables.

Humans behave in a Bayes optimal way

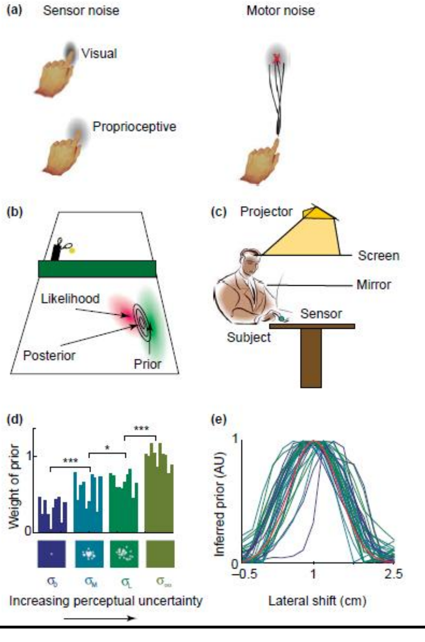

(A) Sensory noise results from different uncertainties. Motor noise will induce variability on the target.

(B) Illustration in tennis: integration of likelihood with priors.

(C) Reliance on priors increase with uncertainty about the environment.

Koerding et al., 2006

Mathematical view of motor control

Sensory state of our body and the world we interact with

Motor command

Motor noise

$$x_{k+1} = Ax_k + B\left(u_k + w_k\right)$$

What we can observe about the state

$$y_k = Hx_k + v_k$$

Sensory noise

Total cost to minimize

$$J = \sum_{k=0}^{p-1} \left(\underbrace{\boxed{y_{k+1}^T Q_{k+1} y_{k+1}}}_{\text{Tracking cost}} + \underbrace{\boxed{u_k^T R_k u_k}}_{\text{Control cost}} \right)$$

Feedback control policy

$$u_k = -L_k \hat{x}_k$$

Belief about state

$$\hat{x}_{k+1} = \hat{A}\hat{x}_k + \hat{A}K_k\left(y_k - \hat{y}_k\right) + \hat{B}u_k$$

Measured sensory consequences

Predicted sensory consequences

Mathematical view of motor control

Parkinson disease

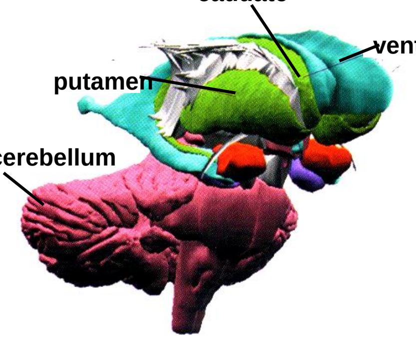

In PD, there is a degeneration of dopaminergic neurons in the substantia nigra. This results in severe loss of dopamine in the basal ganglia, especially the putamen.

These patients exhibit very **slow movements** (bradykinesia), very low voice, and very small writing (micrographia).

Loss of dopamine can be viewed as a **loss of expected reward**, which in turn increases the relative costs of the motor commands with respect to the expected reward.

$$J = \sum_{k=0}^{p-1} \left(\underbrace{\boxed{y_{k+1}^T Q_{k+1} y_{k+1}}}_{\text{Expected reward of a movement}} + \underbrace{\boxed{u_k^T R_k u_k}}_{\text{Motor cost for a movement}} \right)$$