Optimal feedback control, efference copy & learning in sensorimotor processing

Philip Sabes

Department of Physiology
Center for Integrative Neuroscience
UCSF

University of California
San Francisco

Definition of Dynamics

In the engineering literature, the "Dynamics" of a system is composed of two equations:

$\dot{x} = \mathrm{f}(x,u)$ next state equation

$y = \mathrm{g}(x)$ output equation (kinematics, sensory feedback)

The state x **is** the collection of variables that are needed to define next state equation.

Definition of Dynamics

$$\dot{x} = f(x, u)$$ next state equation

$$y = g(x)$$ output equation (kinematics, sensory feedback)

The physics:

$$I\ddot{\theta} = -b\dot{\theta} - k(\theta - \theta_0) + f$$

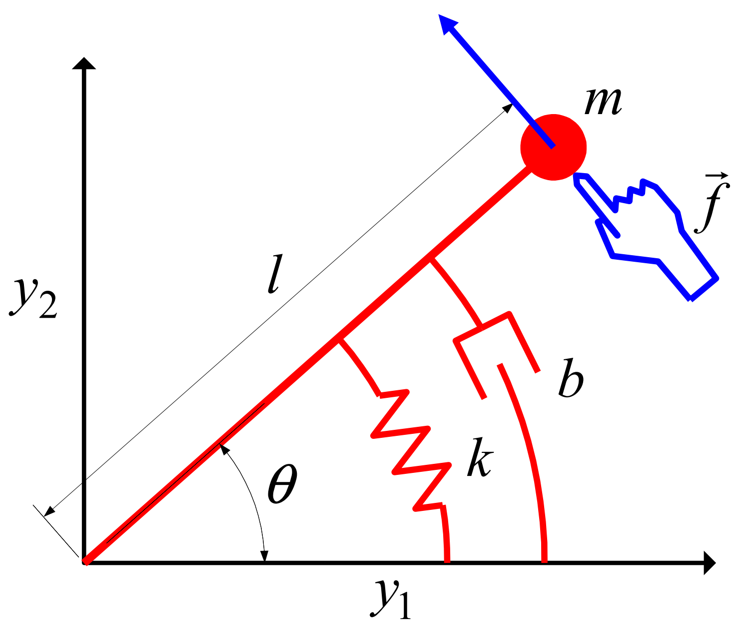

Next state equation:

$$\begin{bmatrix} \dot{\theta} \\ \ddot{\theta} \end{bmatrix} = \begin{bmatrix} 0 & 1 \\ -kI^{-1} & -bI^{-1} \end{bmatrix} \begin{bmatrix} \theta - \theta_0 \\ \dot{\theta} \end{bmatrix} + \begin{bmatrix} 0 \\ I^{-1} \end{bmatrix} f$$

$$\dot{x} = Ax + Bf$$

The state:

$$x = \begin{bmatrix} \theta - \theta_0 \\ \dot{\theta} \end{bmatrix}$$

Output equation:

$$y = \begin{bmatrix} l\cos(\theta) \\ l\sin(\theta) \end{bmatrix}$$

Dynamics of the arm

Analytic form of the dynamics for the *two-joint* arm:

$$\begin{bmatrix}\tau_1\\ \tau_2\end{bmatrix} = \begin{bmatrix} I_1 + I_2 + 2l_1m_2r_2\cos(q_2) + r_1^2m_1 + l_1^2m_2 + r_2^2m_2 & I_2 + r_2^2m_2 + r_2l_1m_2\cos(q_2) \\ I_2 + r_2^2m_2 + l_1m_2r_2\cos(q_2) & I_2 + r_2^2m_2 \end{bmatrix}\begin{bmatrix}\ddot{q}_1\\ \ddot{q}_2\end{bmatrix}$$

$$+ \begin{bmatrix} 0 & -r_2l_1m_2\sin(q_2) & -2r_2l_1m_2\sin(q_2) \\ l_1m_2r_2\sin(q_2) & 0 & 0 \end{bmatrix}\begin{bmatrix}\dot{q}_1^2\\ \dot{q}_2^2\\ \dot{q}_1\dot{q}_2\end{bmatrix} + \begin{bmatrix}B_{11} & B_{12}\\ B_{21} & B_{22}\end{bmatrix}\begin{bmatrix}\dot{q}_1\\ \dot{q}_2\end{bmatrix} \quad (1)$$

Scheidt and Rymer, 2000

The two-joint arm is highly non-linear

Generally, we write: $$\tau = \underbrace{D(\theta)\ddot{\theta}}_{\text{Interia, interaction torques}} + \underbrace{C(\theta,\dot{\theta})\dot{\theta}}_{\text{Coriolis and centripital forces}} + \underbrace{g(\theta)}_{\text{Gravity}}$$

Controlling a dynamical system

- Given a desired movement trajectory (position vs. time), what motor commands are required?
- Dynamics: $\dot{x} = f(x,u)$
- Inverse dynamics: $u = f^{-1}(\dot{x}^*, x)$
- Of course, f is not typically invertible
  - problem is "ill-posed"
- Optimal control can "regularize" problem:
  - optimal control given state $u = L(\hat{x} \mid \dot{x}^*)$
  - optimal estimation of state $\hat{x}(t) = K\left(y(t)\right)$

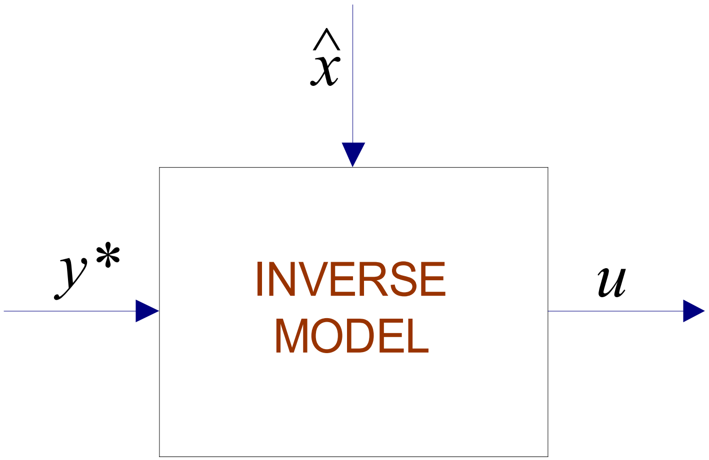

Possible Input:
- desired extrinsic location of ….
- desired perceptual state
- desired intrinsic state (x*)

Possible Output:
- motor unit activity
- signals to Brainstem (eye)
- signals to spinal cord (arm)

Optimal Feedback Control

- What is required to perform optimal control?
  - Optimal control given state
  - Optimal estimation of state
- No analytic solution holds for general optimal control problem
  - The two problems above are interlinked
  - Numerical solutions have been devised
- Model makes some testable behavioral predictions

Minimum intervention principle

Assuming signal dependent noise, a general feature of optimal control emerges

Task: set X_1, X_2 so that $X_1 + X_2 = X^*$

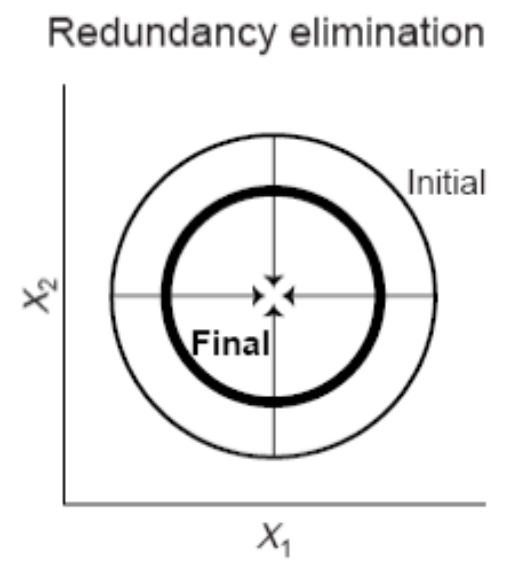

Todorov and Jordan (2002)

Minimum intervention principle

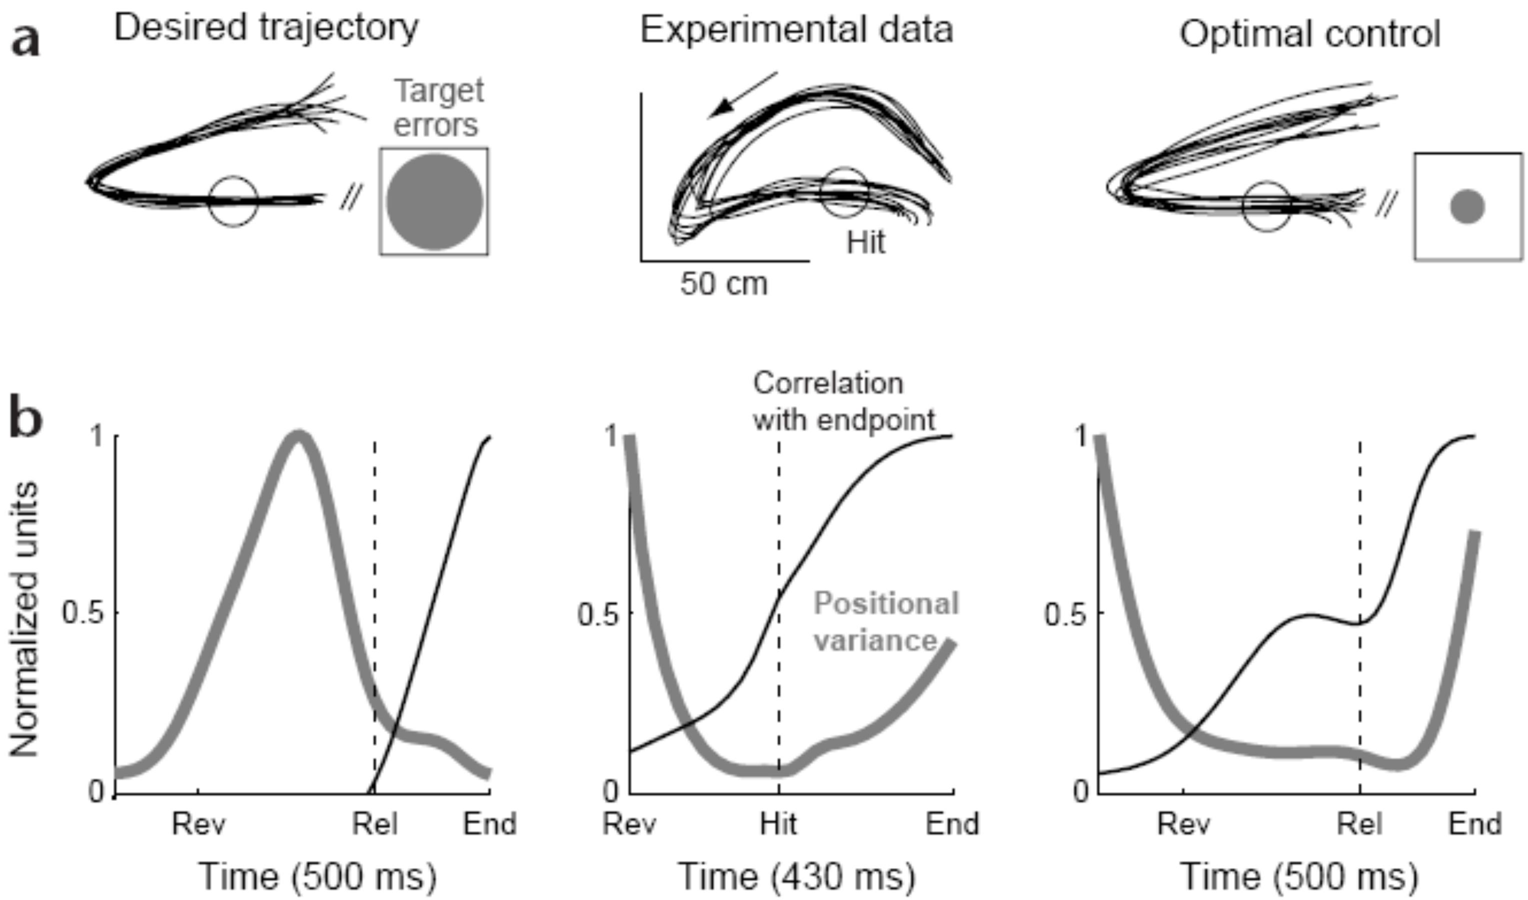

Note: the details of this “principle” rely heavily on the form of the noise model!

Todorov and Jordan (2002)

Minimum intervention principle

Endpoint precision with different target shapes

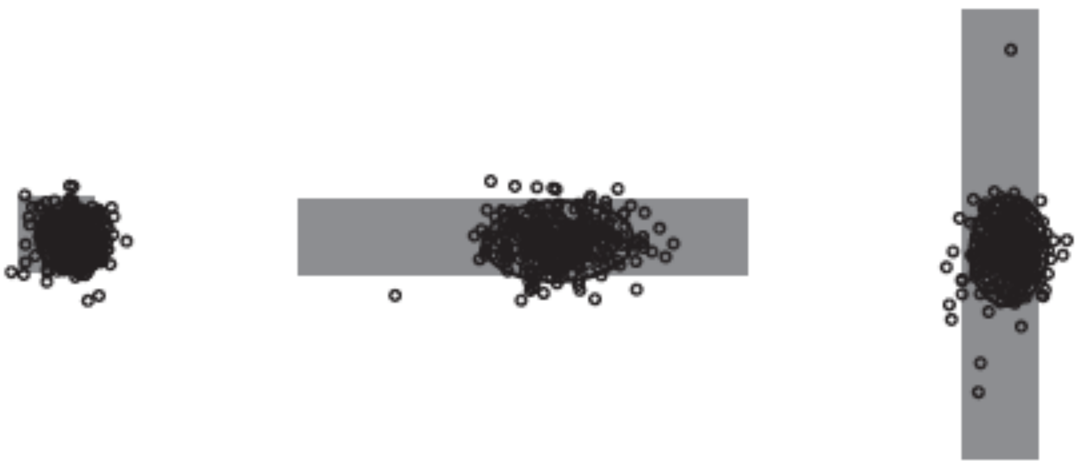

Knill, Bondada, and Manu Chhabra, Journal of Neuroscience (2011)

Minimum intervention principle

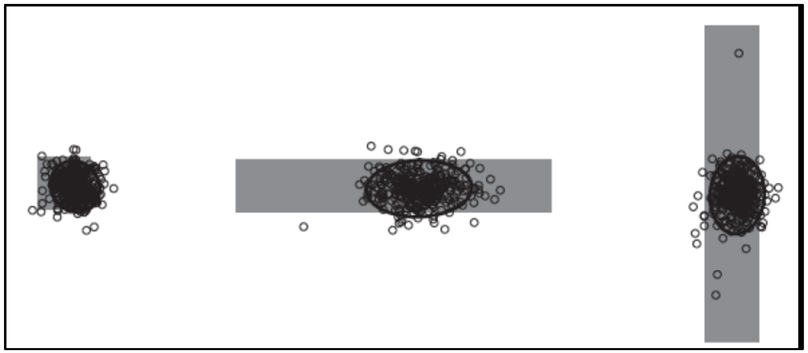

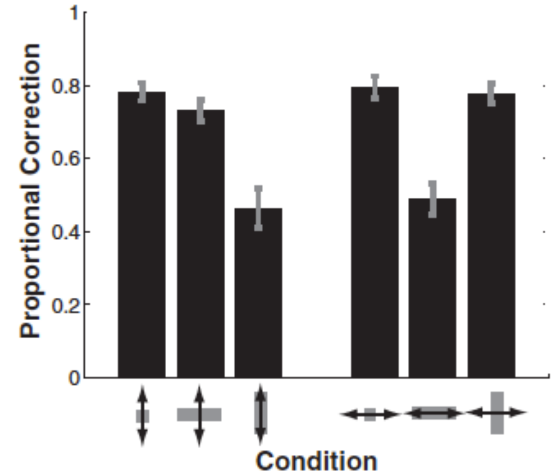

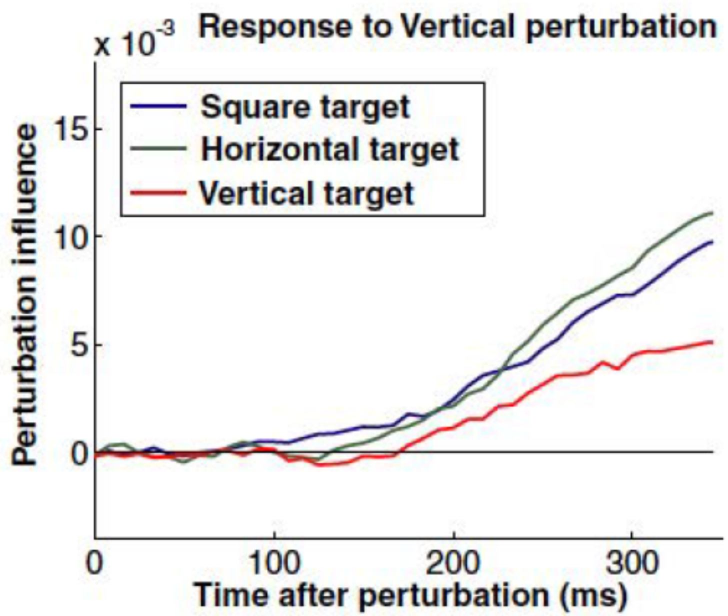

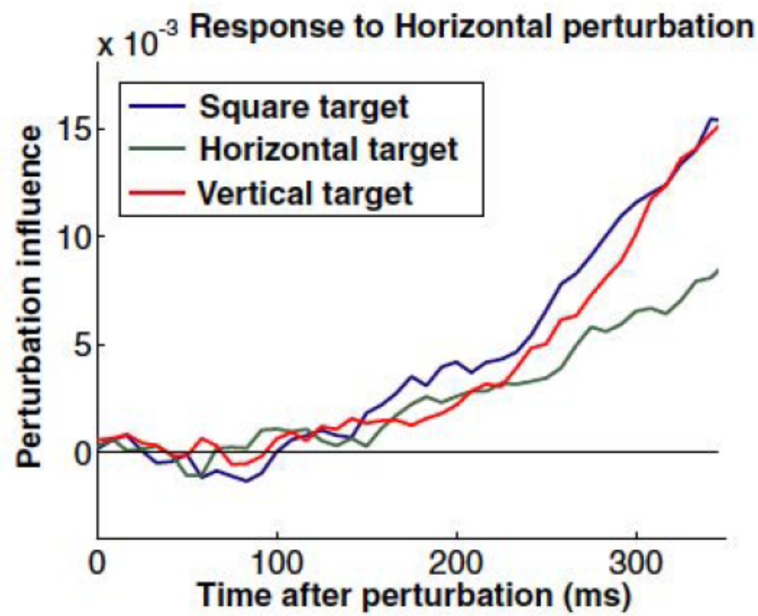

Knill, Bondada, and Manu Chhabra, Journal of Neuroscience (2011)

Optimal Feedback Control

- What is required to perform optimal control?
  - Optimal control given state
  - Optimal estimation of state
- No analytic solution holds for general optimal control problem
  - The two problems above are interlinked
  - Numerical solutions have been devised
- Model makes some testable behavioral predictions
  - Minimum intervention principle
- Solution does exist for a special case:
  - Linear quadratic Gaussian control (LQG)
  - **Linear** dynamical system
  - **Quadratic** cost function
  - Additive **Gaussian** noise

Optimal Feedback Control

Linear quadratic Gaussian control (LQG)

- For a linear system with Gaussian noise

$$\dot{x}_k = Ax_{k-1} + Bu_{k-1} + v_{k-1}$$
$$y_k = Cx_k + w_k$$

v_k, w_k are white noise, Gaussian processes

and a quadratic cost function:

$$J = \sum_{k=1}^{N} x_k^T Q_k x_k + u_k^T R_k u_k$$

- The ***Separation Principle*** holds:
    - State estimation and optimal control can be solved independently
    - Control is performed as if had complete knowledge of state
    - "Linear-Quadratic Regulator" and "Kalman Filter"

Linear-quadratic regulator

Finite-horizon, discrete-time LQR

For a discrete-time linear system described by [1]

$$x_k = Ax_{k-1} + Bu_k$$

with a performance index defined as

$$J = \sum_{k=0}^{N} \left(x_k^T Q x_k + u_k^T R u_k \right)$$

Cost function

the optimal control sequence minimizing the performance index is given by

$$u_k = -F_k x_{k-1}$$

Control law

where

$$F_k = (R + B^T P_k B)^{-1} B^T P_k A$$

and P_k is found iteratively backwards in time by the dynamic Riccati equation

$$P_{k-1} = Q + A^T \left(P_k - P_k B \left(R + B^T P_k B \right)^{-1} B^T P_k \right) A$$

from initial condition $P_N = Q$.

Note: no output equation or feedback signals – assumes that state is known perfectly

What does this mean??

- Consider case Q=I, R=0
- Consider Q=I, R=scalar

http://en.wikipedia.org/wiki/Linear-quadratic_regulator

Kalman Filter

- For linear Gaussian dynamical process:

$$x_k = Ax_{k-1} + Bu_{k-1} + v_{k-1}$$
$$y_k = Cx_k + w_{k-1}$$
$$v_k \sim N(0,V), \quad w_k \sim N(0,W)$$

- The Kalman filter is given by

Predict

Predicted (*a priori*) state estimate $\hat{x}_{k|k-1} = A\ \hat{x}_{k-1|k-1} + B\ u_k$

Predicted (*a priori*) estimate covariance $P_{k|k-1} = A\ P_{k-1|k-1} A^T + V_k$

Update

Innovation (measurement residual) $\tilde{y}_k = y_k - C\ \hat{x}_{k|k-1}$

Innovation covariance $S_k = CP_{k|k-1}C^T + W_k$

Optimal Kalmal gain $K_k = P_{k|k-1} C S_k^{-1}$

Updated (*a posteriori*) state estimate $\hat{x}_{k|k} = \hat{x}_{k|k-1} + K_k \tilde{y}_k$

Updated (*a posteriori*) estimate covariance $P_{k|k} = (I - K_k C) P_{k|k-1}$

http://en.wikipedia.org/wiki/Kalman_filter

Kalman Filter

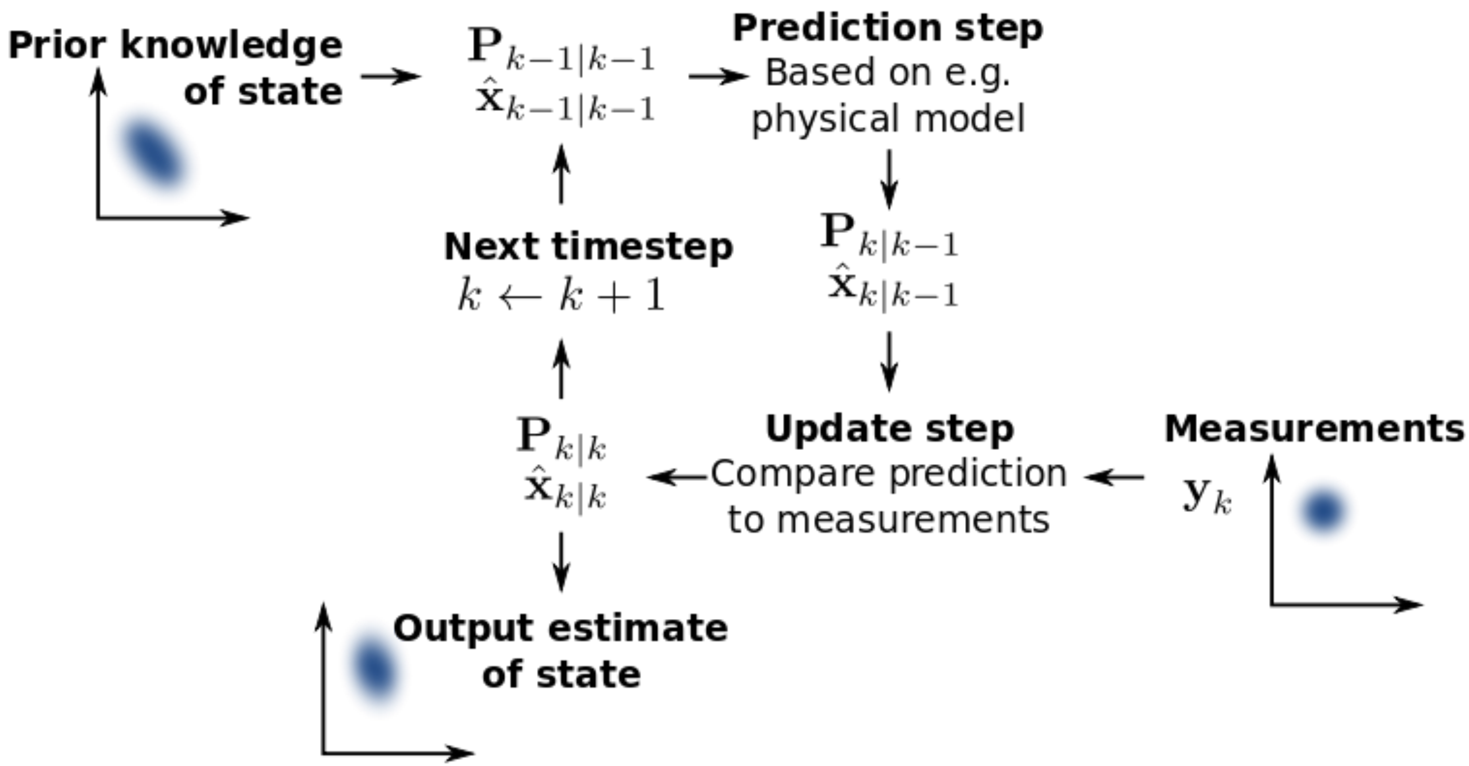

http://en.wikipedia.org/wiki/Kalman_filter

Kalman Filter

- For linear Gaussian dynamical process:
$$x_k = Ax_{k-1} + Bu_{k-1} + v_{k-1}$$
$$y_k = Cx_k + w_{k-1}$$
$$v_k \sim N(0,V), \quad w_k \sim N(0,W)$$

- The Kalman filter:
$$\hat{x}_{k|k-1} = A\hat{x}_{k-1|k-1} + Bu_{k-1}$$
$$\hat{x}_{k|k} = \hat{x}_{k|k-1} + K_k\left(y_k - C\hat{x}_{k|k-1}\right)$$

- An example of an "internal model":
  - "forward model" of the plant dynamics
  - "efference copy"

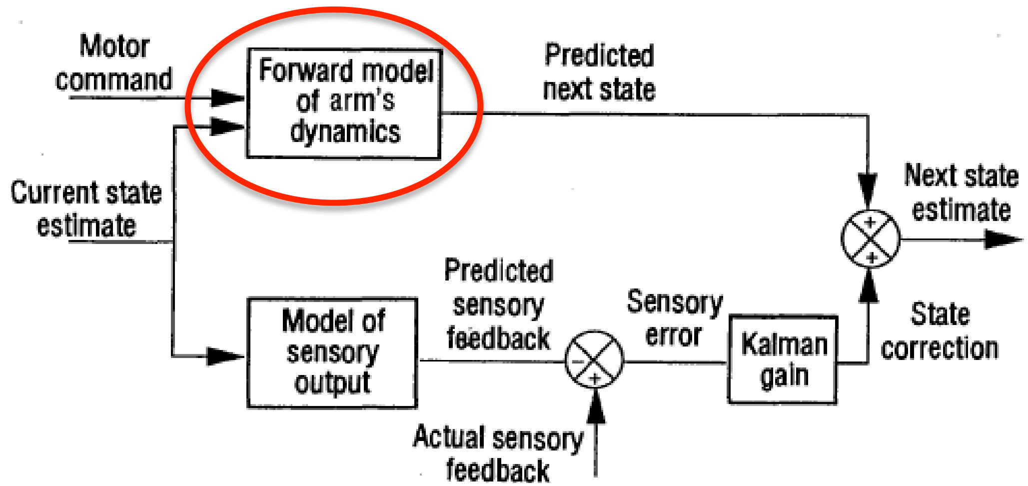

Sensory integration as Kalman Filter

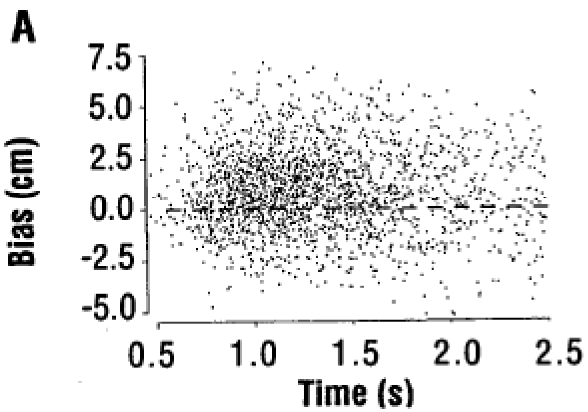

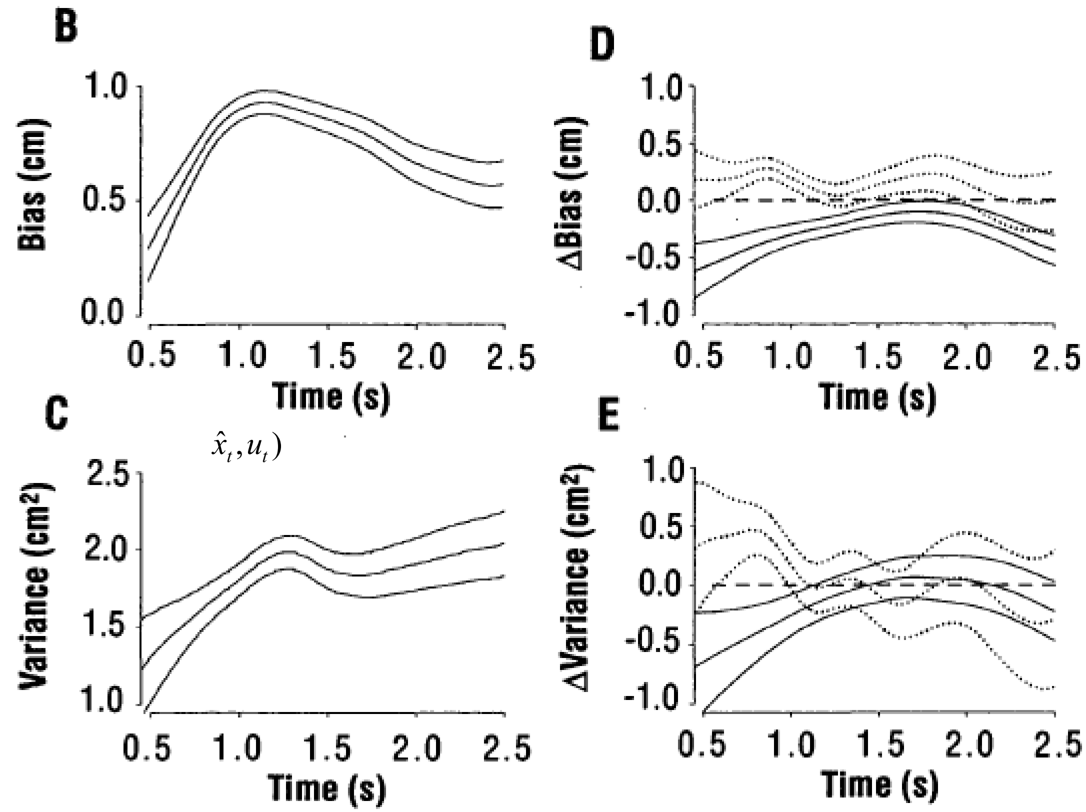

Fig. 1. The raw localization bias against movement duration is shown in (**A**) for all eight participants. There are few data points for short movement durations because of the reaction time of stopping in response to the tone. All graphs are therefore plotted from 0.5 s. (**B** through **E**) The main effect fits of the

Wolpert et al, 1995

Sensory integration as Kalman Filter

To get this to work, had to assume that force output is underestimated (??)

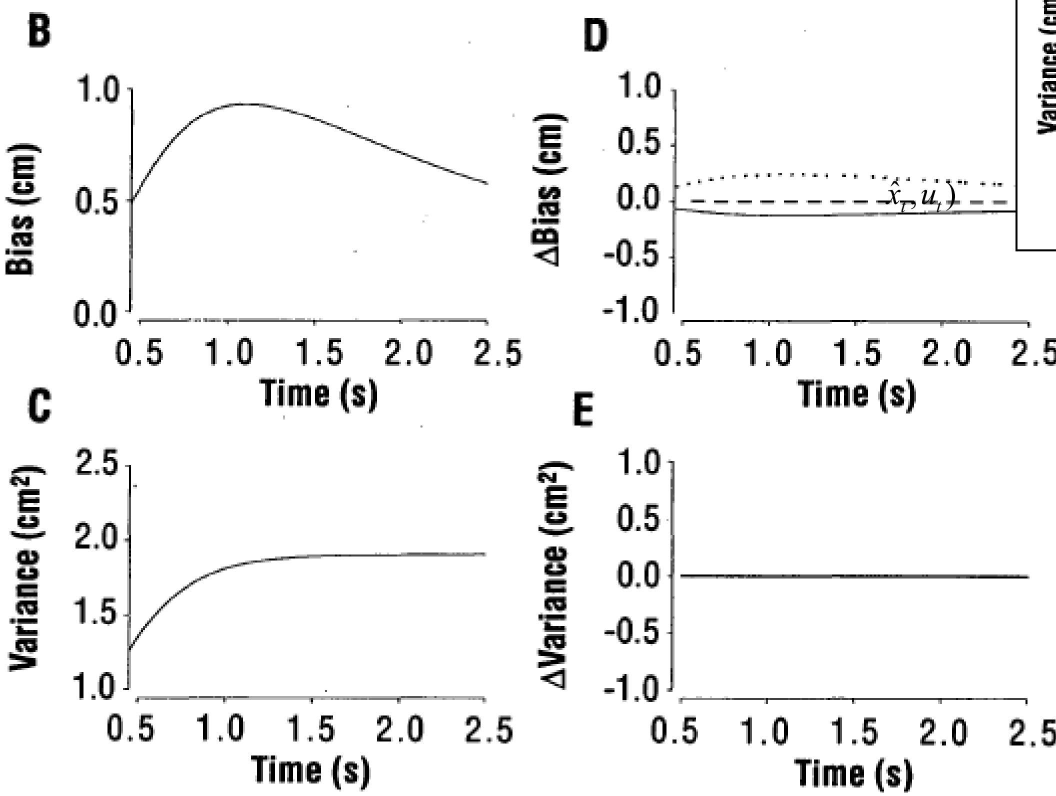

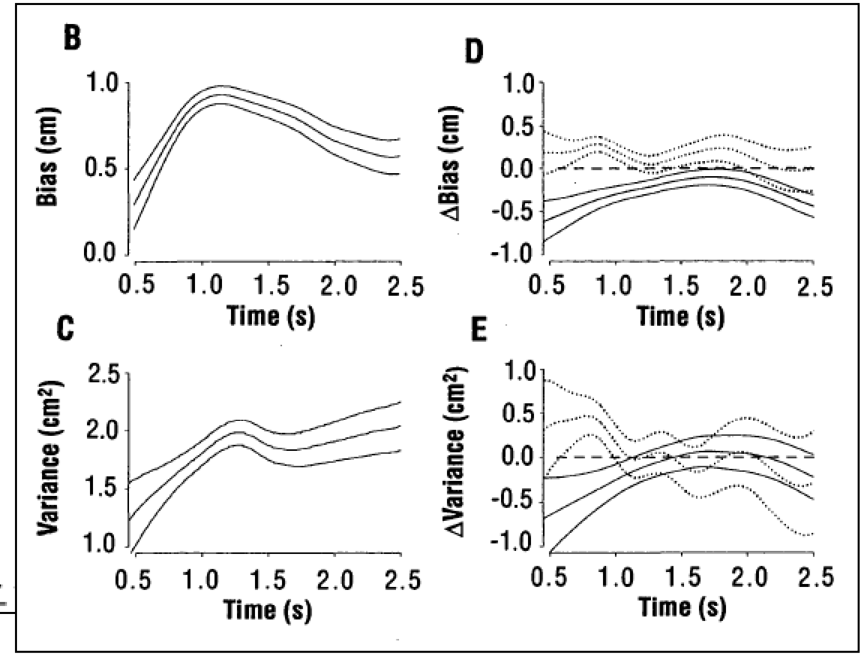

Wolpert et al, 1995

The role of proprioceptive feedback

The effect of Large Sensory Fiber Neuropathy

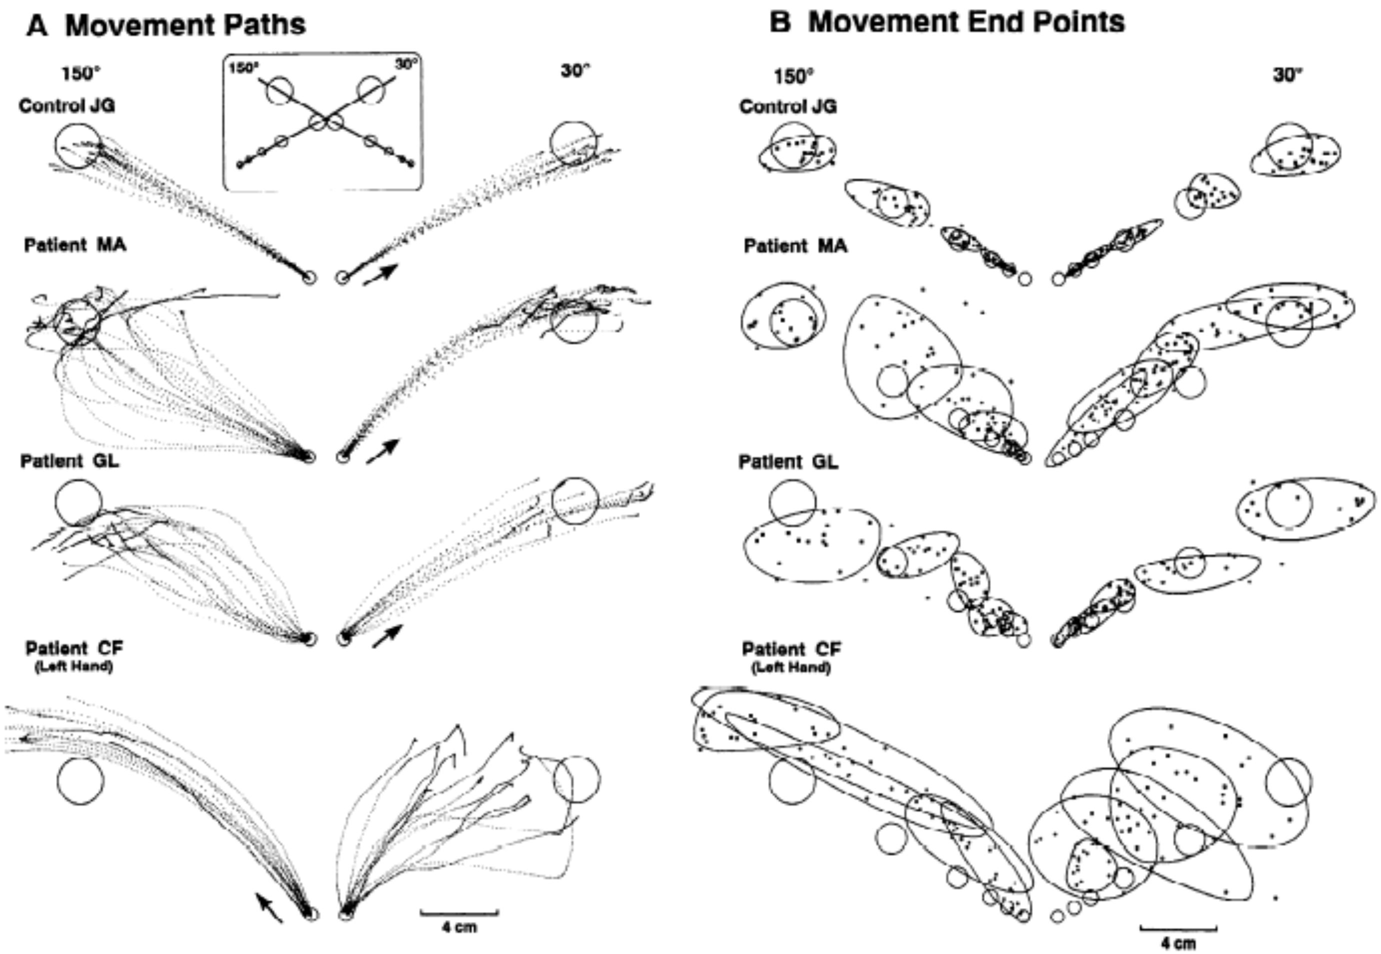

Gordon, Ghilardi and Ghez (1995)

Effects of inertia on initial movement

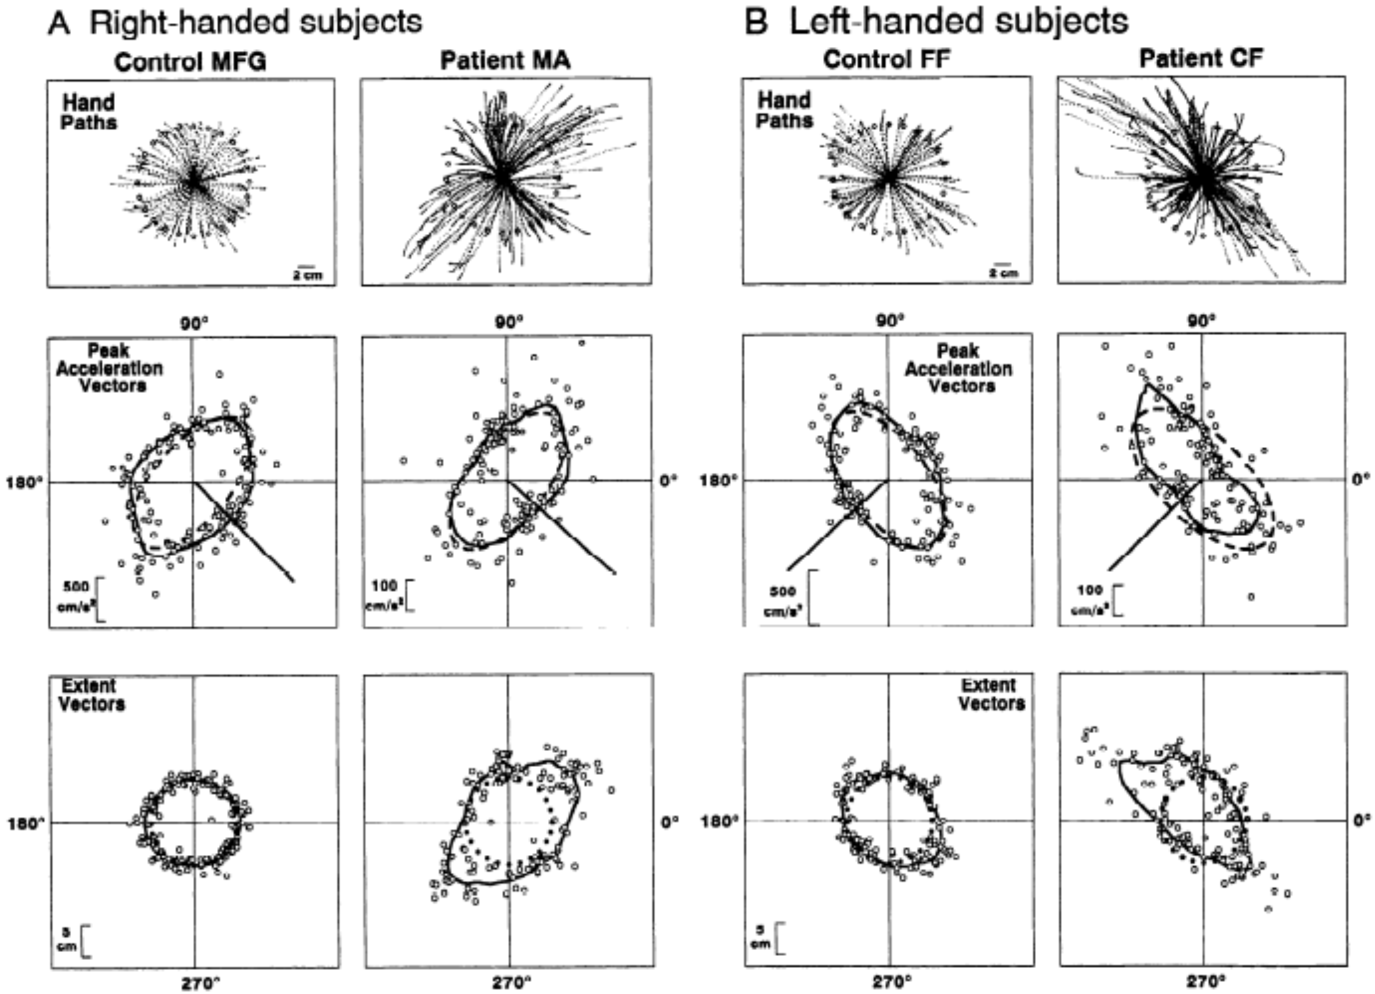

Normals compensate for inertia with movement time

Gordon, Ghilardi and Ghez (1995)

Why are deafferented patients impaired?

- Trouble reaching because they don' t know where their hands are
- What about when the have visual feedback?

Vision helps, and the benefit persists

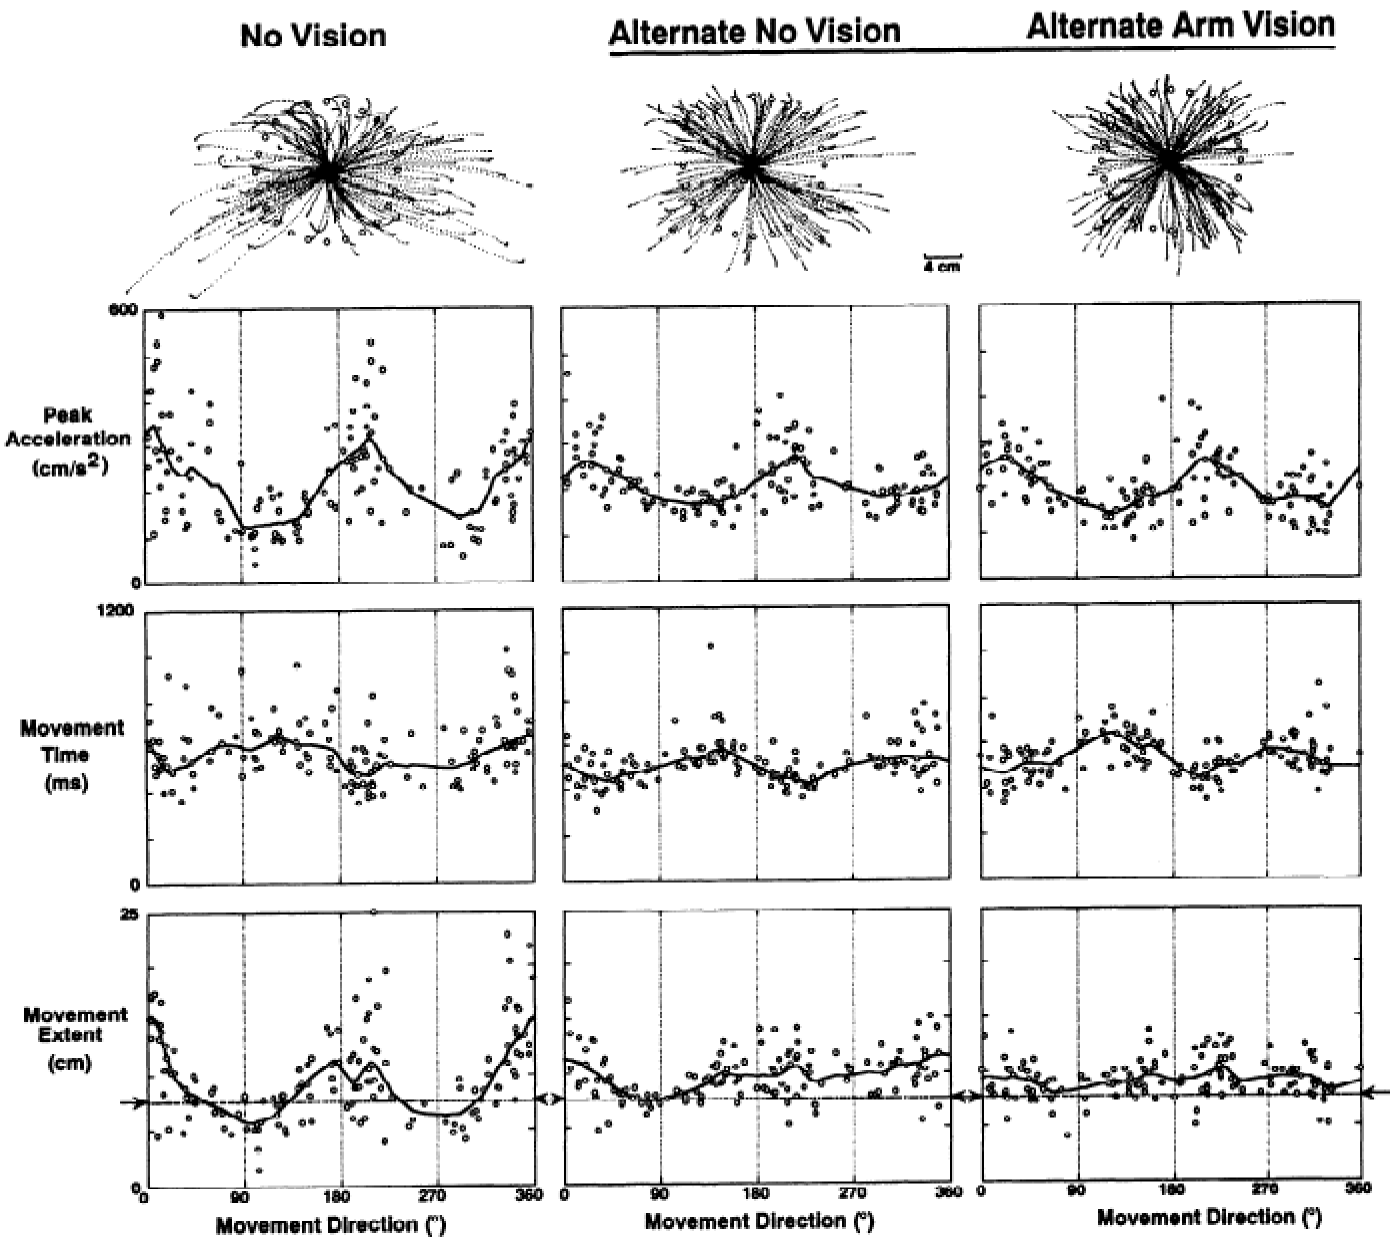

Ghez, Gordon, and Ghilardi (1995)

Vision helps, and the benefit persists

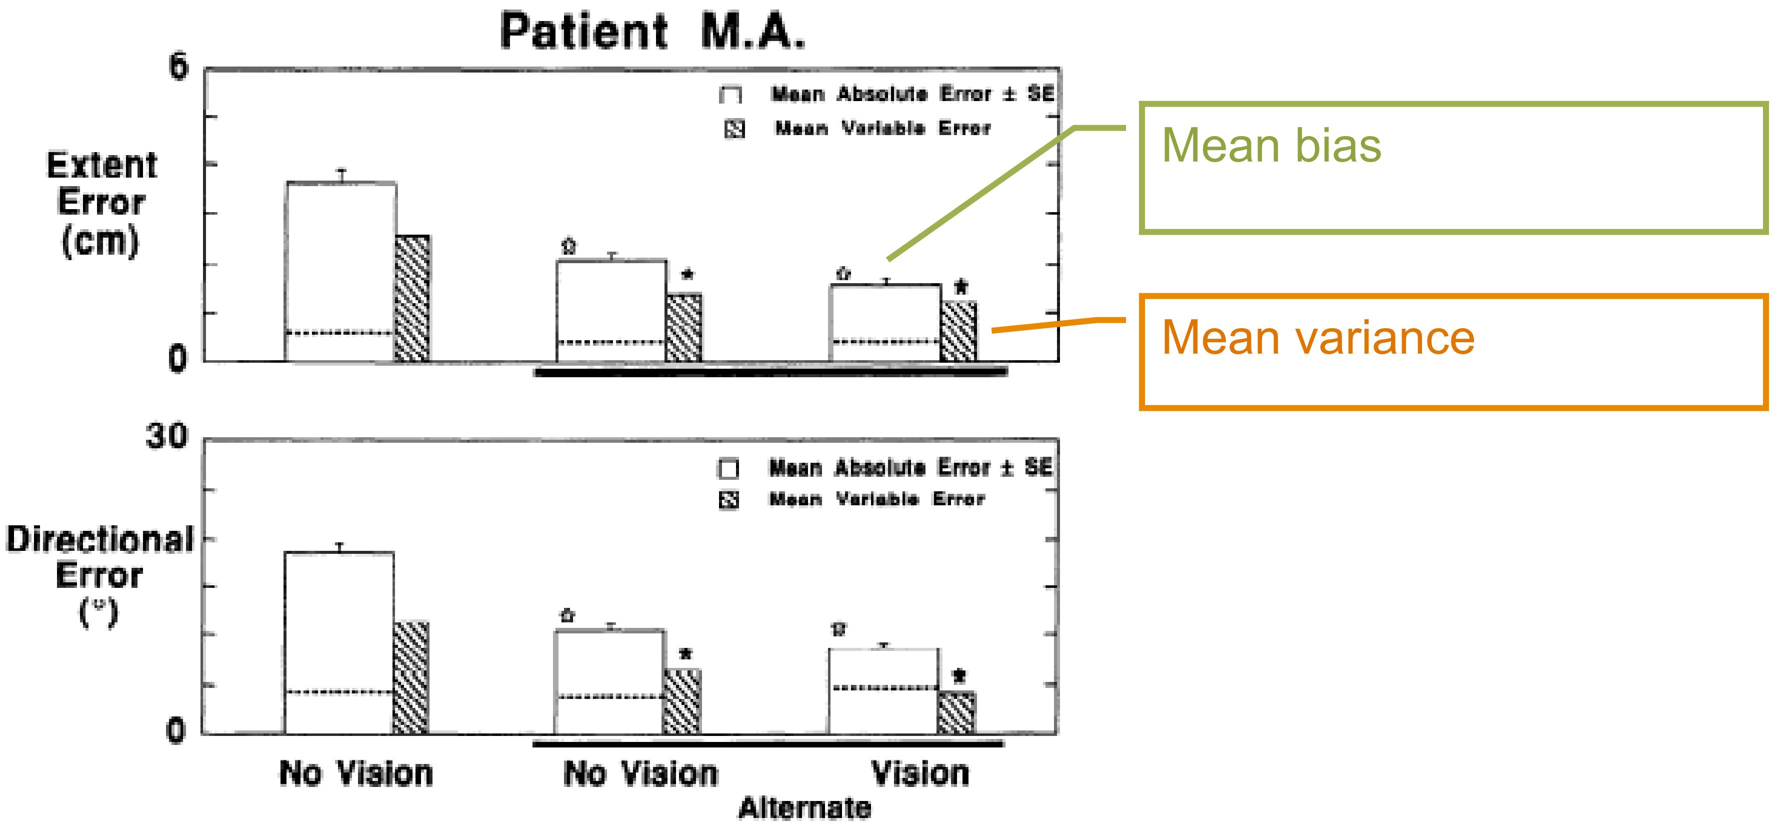

Ghez, Gordon, and Ghilardi (1995)

Why does prior vision help?

- For linear Gaussian dynamical process:
$$x_k = Ax_{k-1} + Bu_{k-1} + v_{k-1}$$
$$y_k = Cx_k + w_{k-1}$$
$$v_k \sim N(0,V), \quad w_k \sim N(0,W)$$

- The Kalman filter:
$$\hat{x}_{k|k-1} = A\hat{x}_{k-1|k-1} + Bu_{k-1}$$
$$\hat{x}_{k|k} = \hat{x}_{k|k-1} + K_k\left(y_k - C\hat{x}_{k|k-1}\right)$$

- An example of an "internal model":
  - "forward model" of the plant dynamics
  - "efference copy"

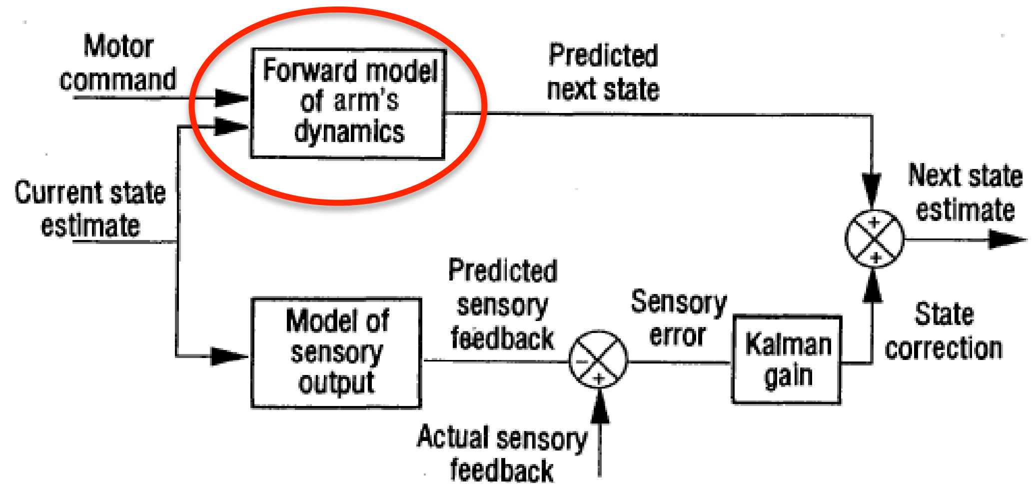

Why does prior vision help?

- For linear Gaussian dynamical process:
$$x_k = Ax_{k-1} + Bu_{k-1} + v_{k-1}$$
$$y_k = Cx_k + w_{k-1}$$
$$v_k \sim N(0,V), \quad w_k \sim N(0,W)$$

- The Kalman filter:
$$\hat{x}_{k|k-1} = A\hat{x}_{k-1|k-1} + Bu_{k-1}$$
$$\hat{x}_{k|k} = \hat{x}_{k|k-1} + K_k\left(y_k - C\hat{x}_{k|k-1}\right)$$

- What does the Kalman gain do??

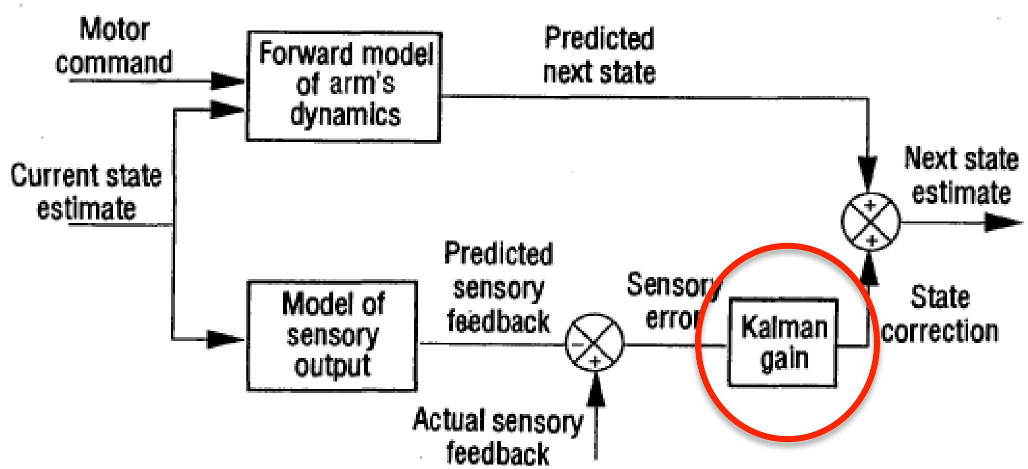

Optimal Estimation

How should multiple sensory inputs be combined – "sensory integration"?

If goal is *accurate perception*, it is reasonable consider the variability of each input modality:

$\hat{x}_{\text{vis}}$ has variance σ^2_{vis}

$\hat{x}_{\text{prop}}$ has variance σ^2_{prop}

$$\Rightarrow \quad \hat{x}_c = \frac{\sigma^2_{\text{prop}}}{\sigma^2_{\text{vis}} + \sigma^2_{\text{prop}}} \hat{x}_{\text{vis}} + \frac{\sigma^2_{\text{vis}}}{\sigma^2_{\text{vis}} + \sigma^2_{\text{prop}}} \hat{x}_{\text{prop}}$$

Optimal Estimation

$\hat{x}_{\text{vis}}$ has variance σ^2_{vis}

$\hat{x}_{\text{prop}}$ has variance σ^2_{prop}

$$\Rightarrow \quad \hat{x}_c = \frac{\sigma^2_{\text{prop}}}{\sigma^2_{\text{vis}} + \sigma^2_{\text{prop}}} \hat{x}_{\text{vis}} + \frac{\sigma^2_{\text{vis}}}{\sigma^2_{\text{vis}} + \sigma^2_{\text{prop}}} \hat{x}_{\text{prop}}$$

"Minimum Variance" is also "Maximum Likelihood" for the case of independent Gaussian signals

$$K_k = P_{k|k-1} C \left(C P_{k|k-1} C^T + W_k \right)^{-1} \approx \frac{\text{variance of } \hat{\text{x}}_{k|k-1}}{\left(\text{variance of } \hat{\text{x}}_{k|k-1} + \text{variance of } y \right)}$$

The Kalman gain represents the "maximum likelihood" way to combine **current sensory information** with predictions from past experience (i.e., **prior information**) … Bayesian integration

Summary so far

- Optimal feedback control
  - Minimum intervention principle
- Special case: LQG
  - Separation principle
  - Kalman filter

Learning

- Optimal control, optimal estimation require learning
- Optimal control:
  - Movements that minimize cost functions
  - Need to learn "plant", costs, etc
- Sensory integration
  - Combine signals according to their statistics
  - Need to learn statistics
- Bayesian integration
  - Need time to acquire "prior information"
- A brief survey of sensorimotor learning

Optimality requires learning

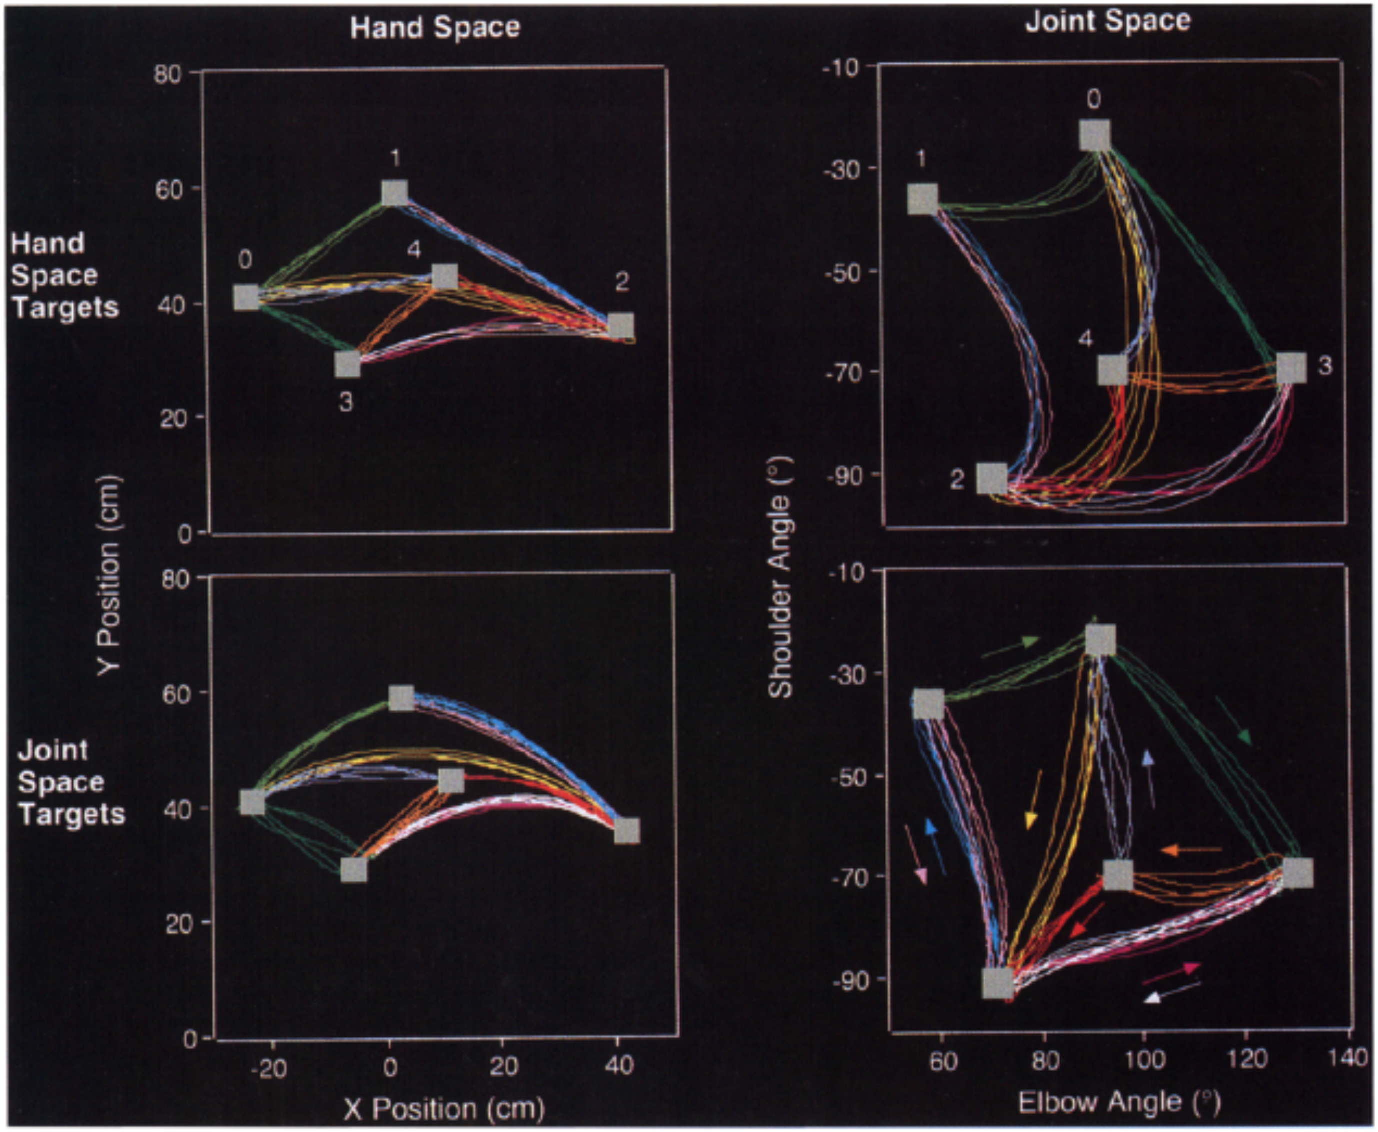

Flanagan and Rao (1995)

Optimality requires learning

Most well-known example: reaching in a "curl-viscous" force field:

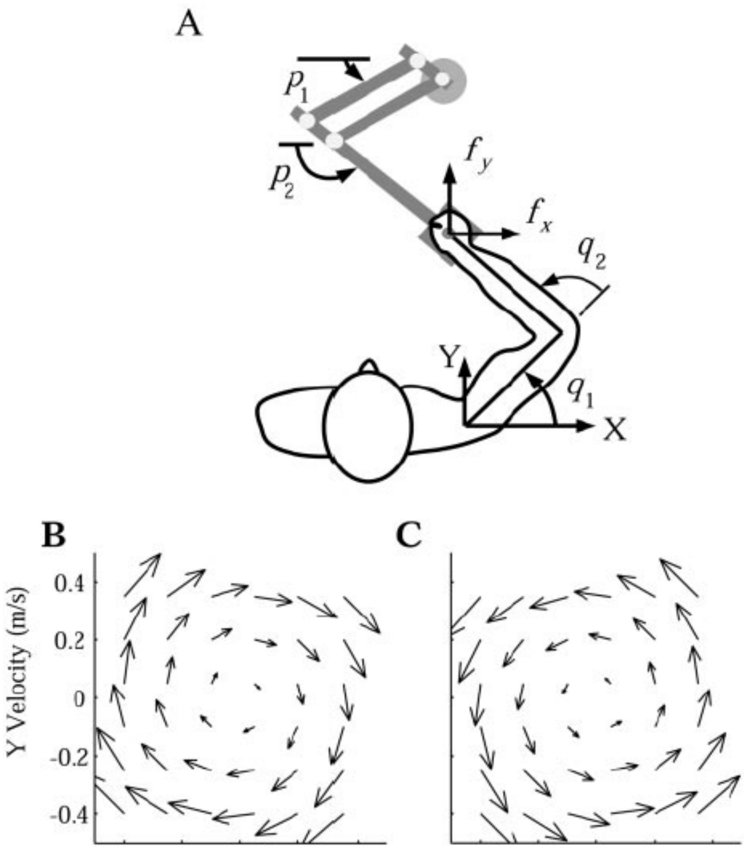

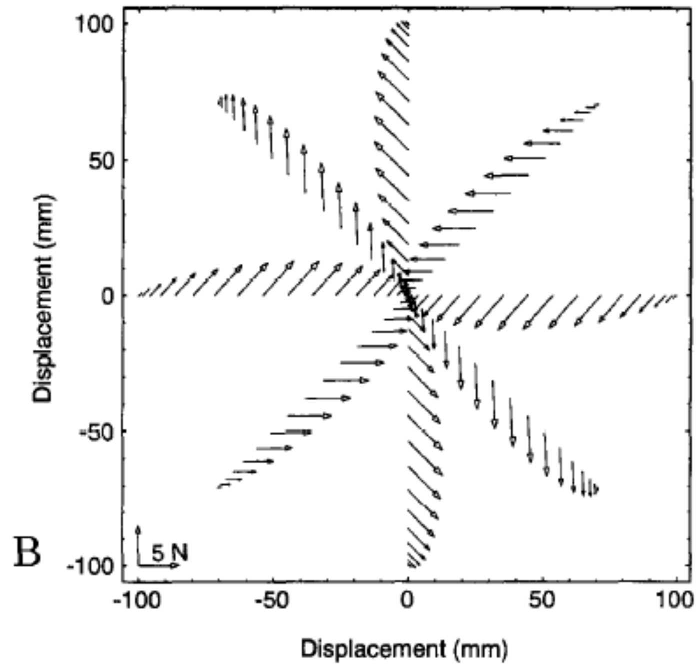

Shadmehr and Mussa Ivaldi(1994)

Prism Adaptation

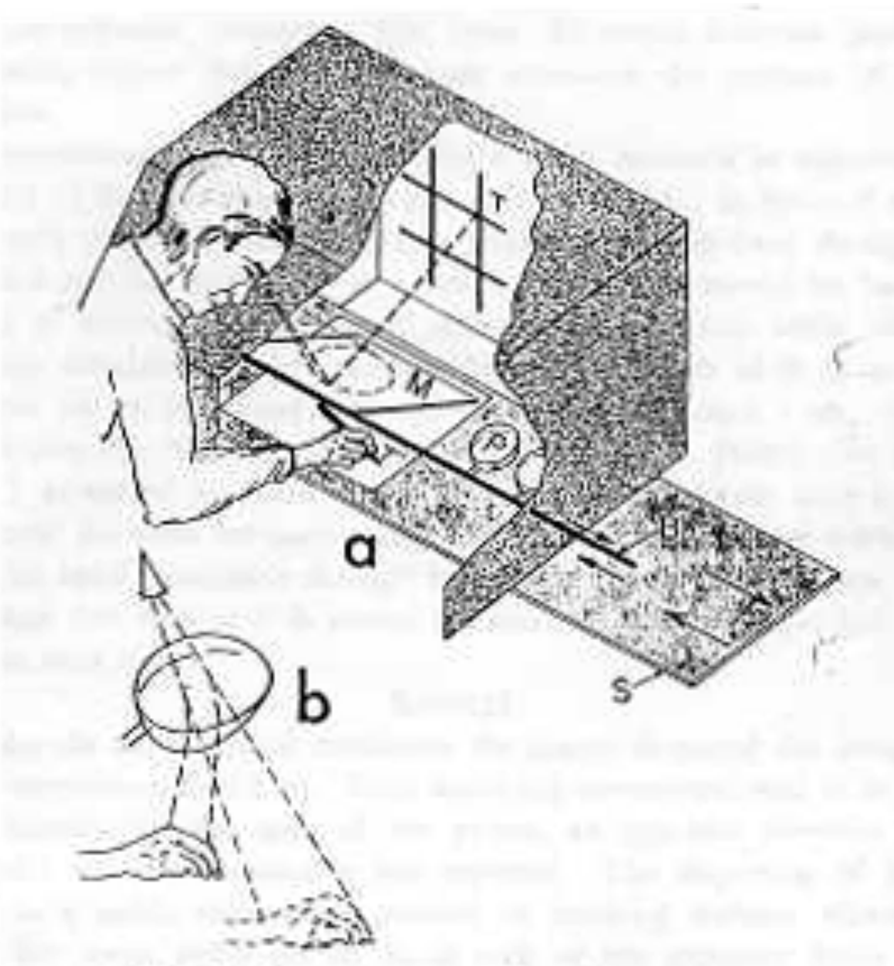

FIG. 1. Schematic representation of the apparatus

Held and Gottlieb (1958)

What drives learning?

Task-specific learning: an error corrective learning model

Let y_t be the reach endpoint on trial t, and the "state" x_t be the mean endpoint:

$$y_t = x_t + w_t, \quad \text{where } w_t \sim N(0,R)$$

Learning can be modeled as incremental error correction driven by error ε:

$$x_{t+1} = Ax_t + B\varepsilon_t + u_t, \quad \text{where } u_t \sim N(0,Q)$$

The model is a linear dynamical system
We can fit parameters to trial-by-trial experimental data

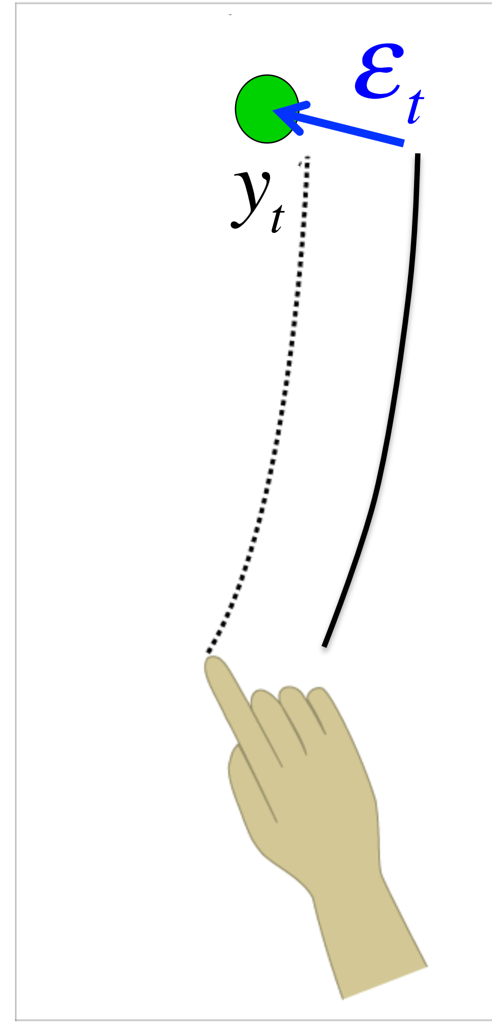

Cheng and Sabes (2006,2007)

Sample Model Fits

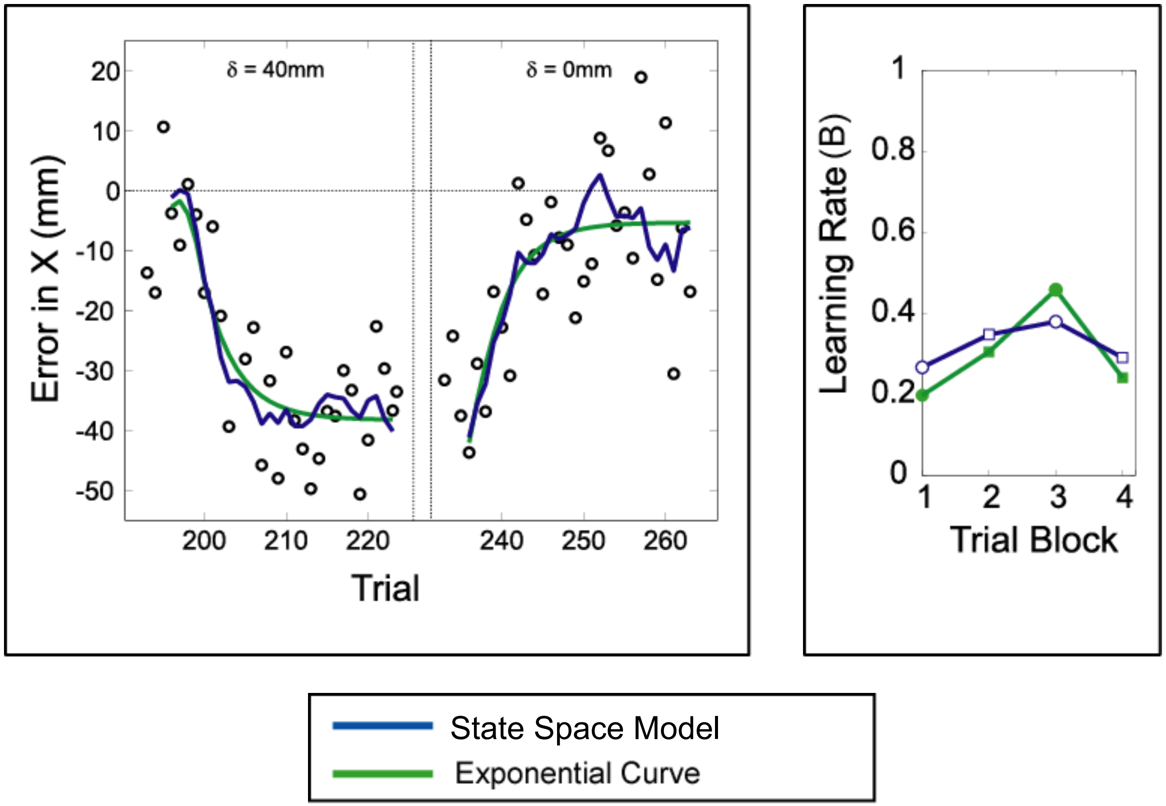

Cheng and Sabes (2006,2007)

Sensory integration is also learned

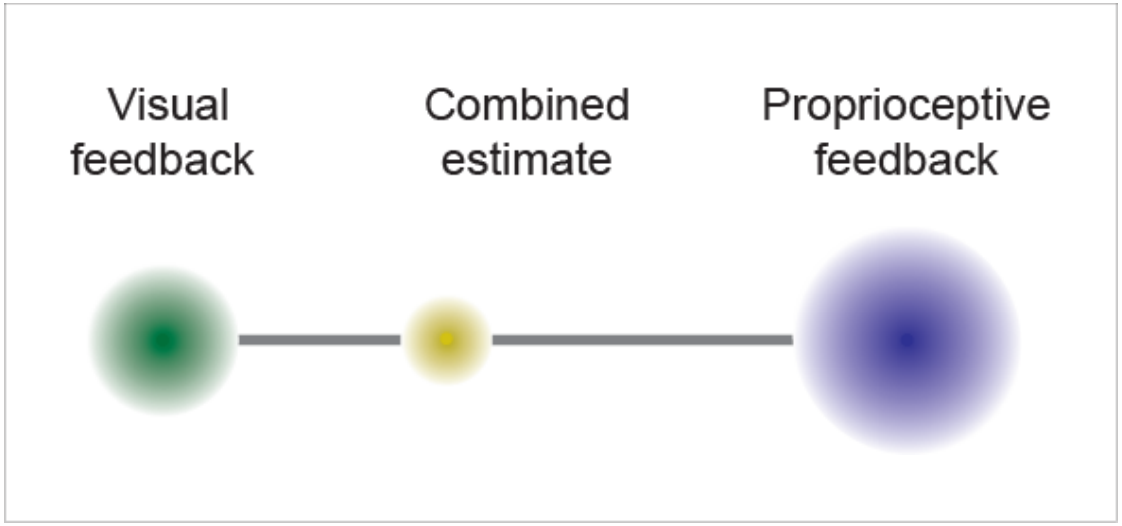

Integration

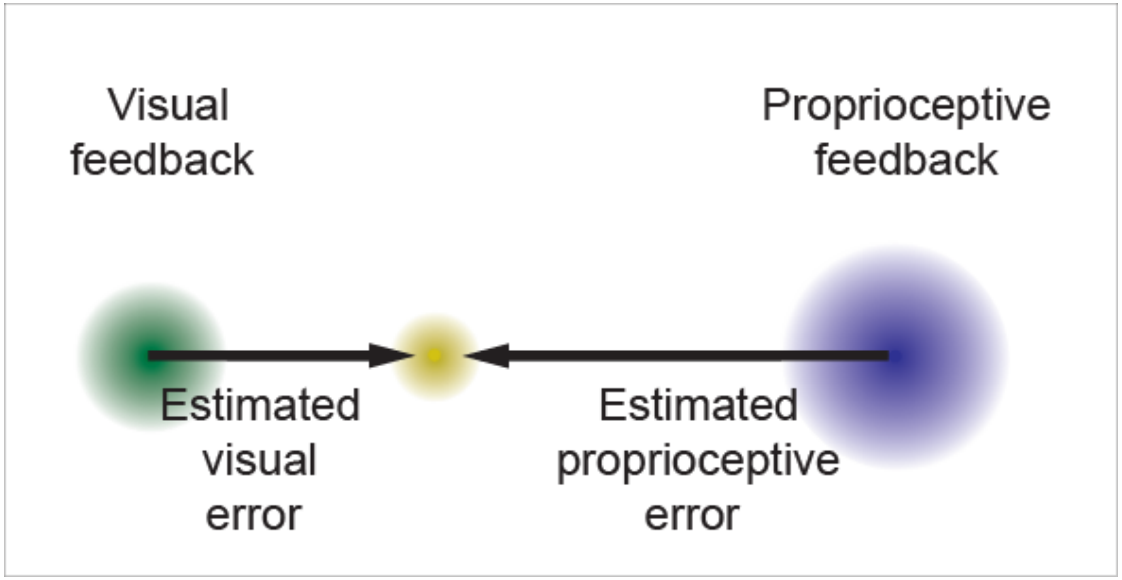

Adaptation

Maximum likelihood sensory integration

Sensory integration

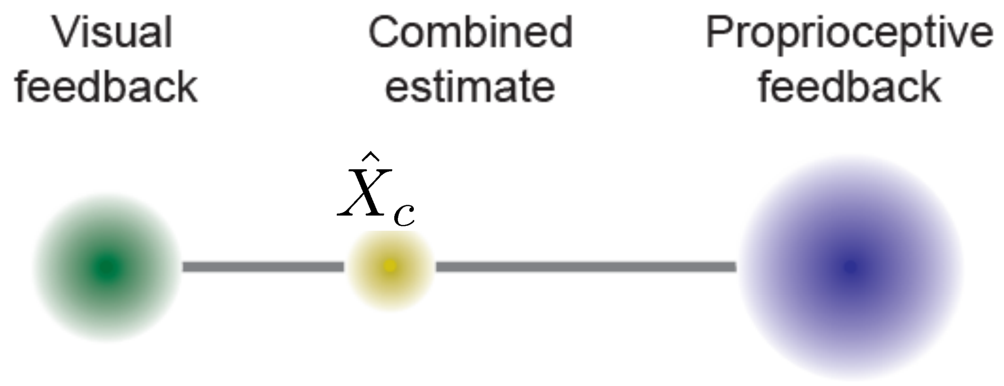

Maximum likelihood sensory integration

Sensory integration

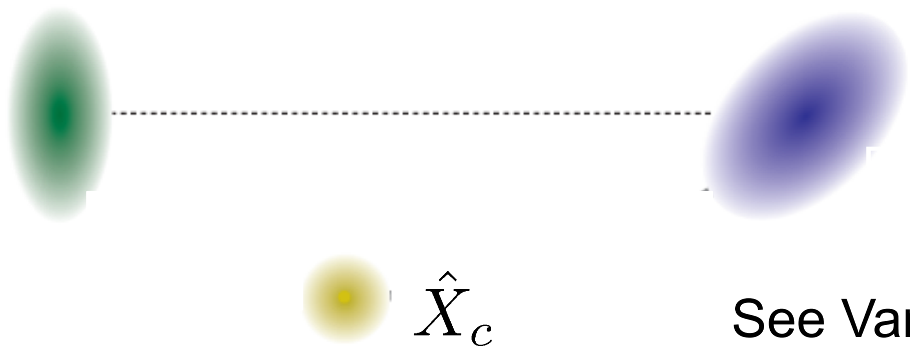

See VanBeers et al., 1999

- Where does discrepancy come from?
- What happens if discrepancy persists?
  - Sensory adaptation, “recalibration” (Simani, McGuire and Sabes, 2007)
- How should the sensory modalities adapt?
  - A “credit assignment” problem

Integration and adaptation

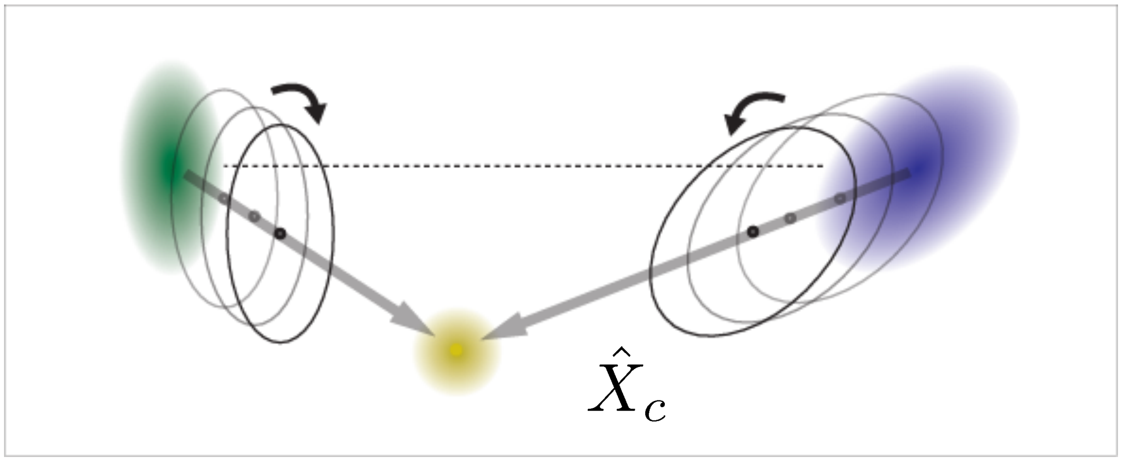

- An iterative version of Expectation-Maximization (EM) algorithm (Ghaharamani, 1995)
- Can also formulate as a Kalman filter (brain is trying to estimate "bias")
- Somewhere in-between supervised and unsupervised learning
- Again, this is the **normal state**

Integration and adapatation

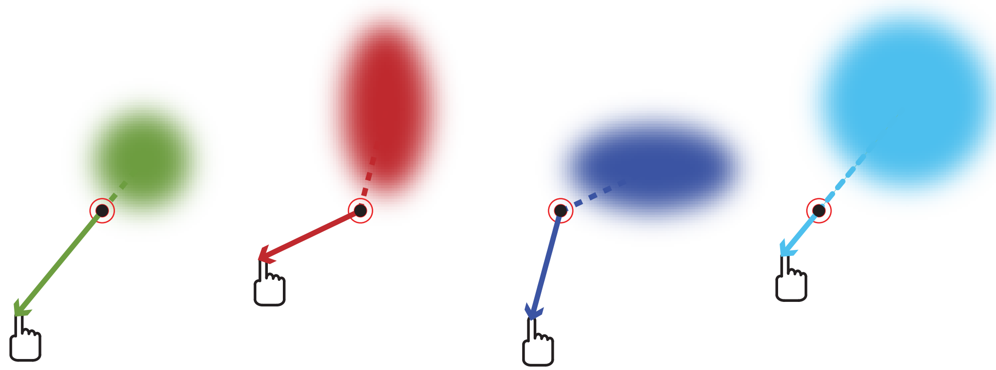

Visual shifts with different uncertainty ….

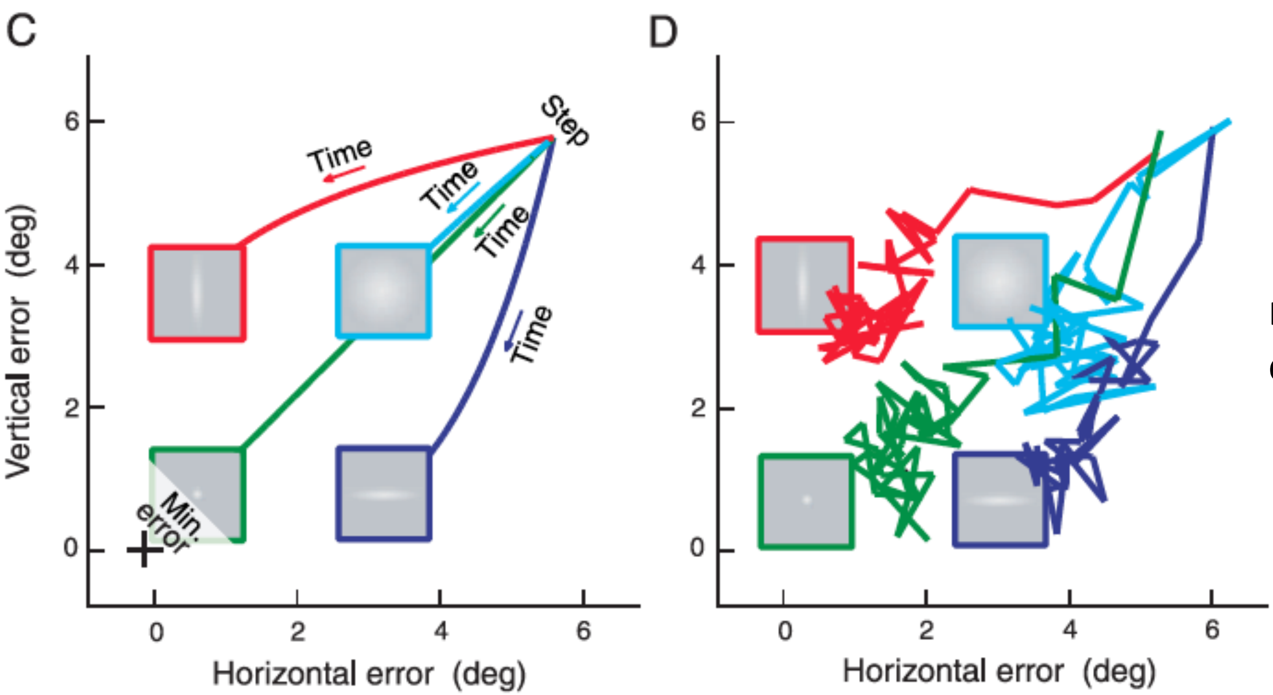

result in different pattern of adapation

Burge and Banks, J Vision (2007)

Adaptive sensory integration

- Network performs ML integration

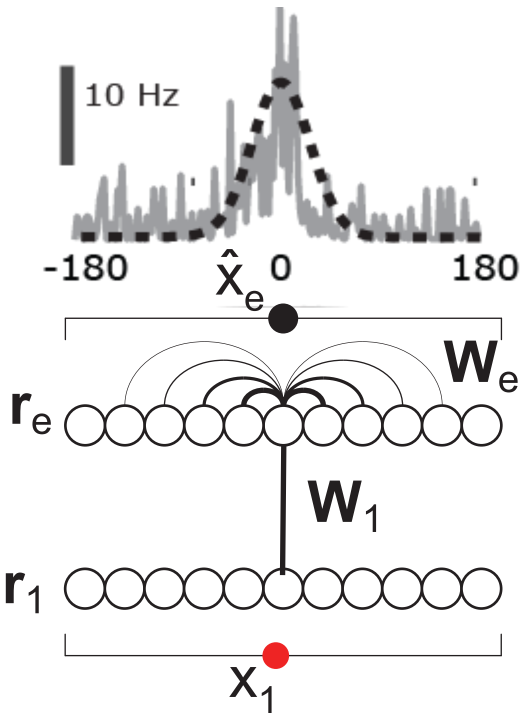

Deneve, Latham, Pouget (1999)

Adaptive sensory integration

- Network performs ML integration

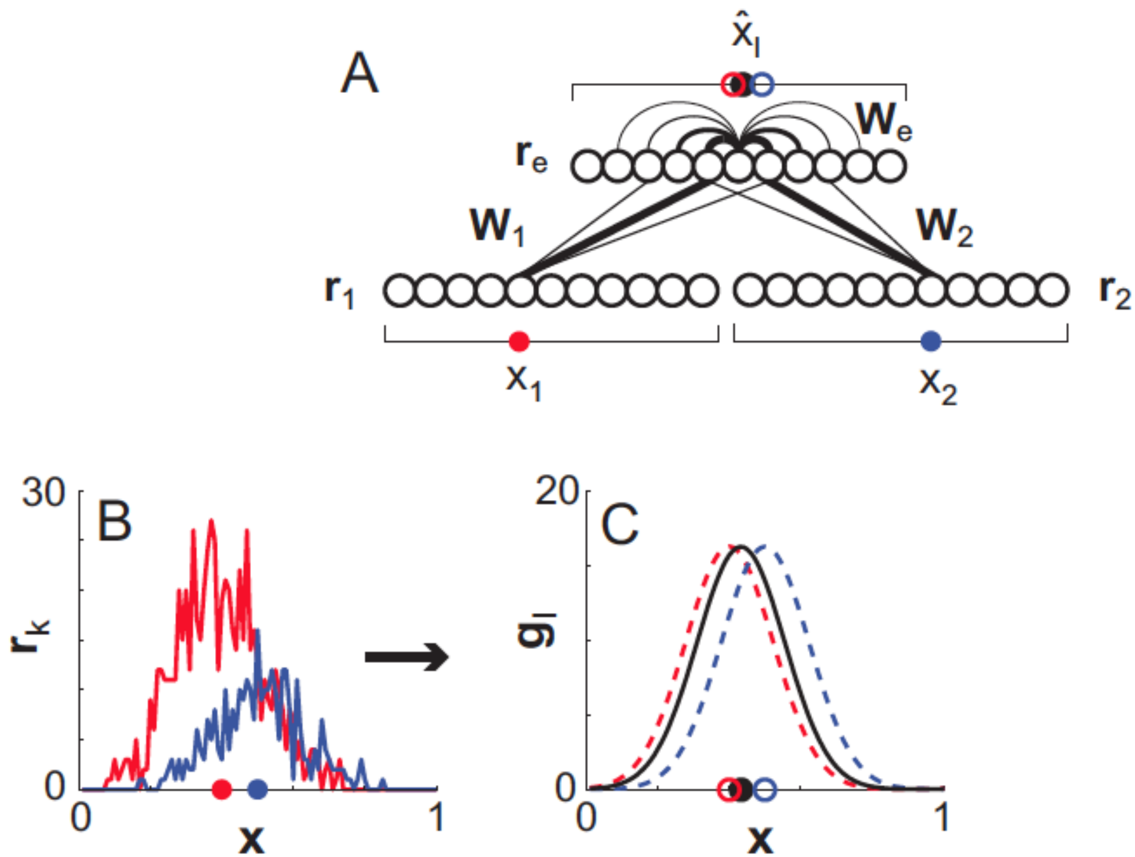

Deneve, Latham, Pouget (1999)

Biddle Snead and Sabes, *unpub.*

Adaptive sensory integration

- Network adapts, using same scheme as EM model (integrated estimate as "teacher")
  - Hebbian learning at the inputs

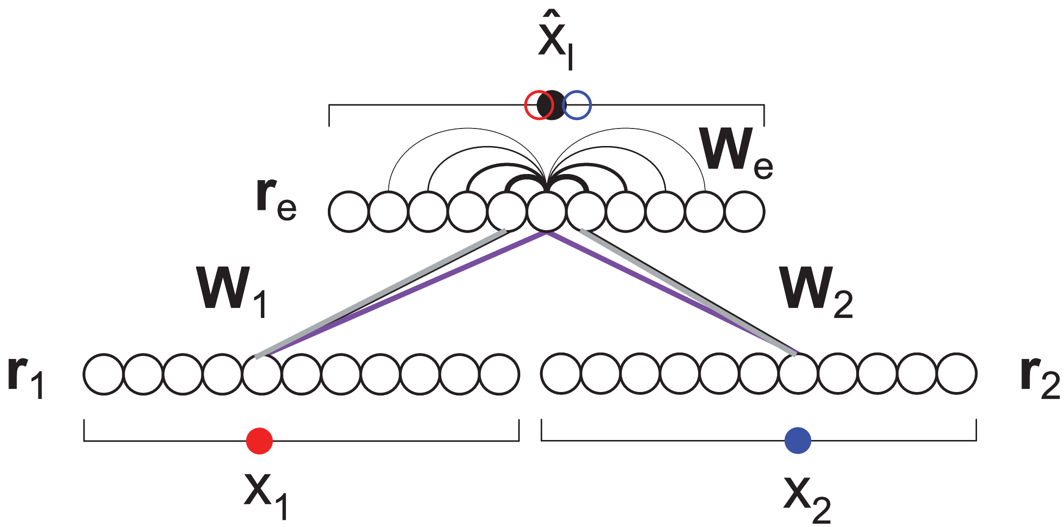

Biddle Snead and Sabes, *unpub.*

Adaptive sensory integration

Burge and Banks, J Vision (2007)

Models capture temporal dynamics of learning

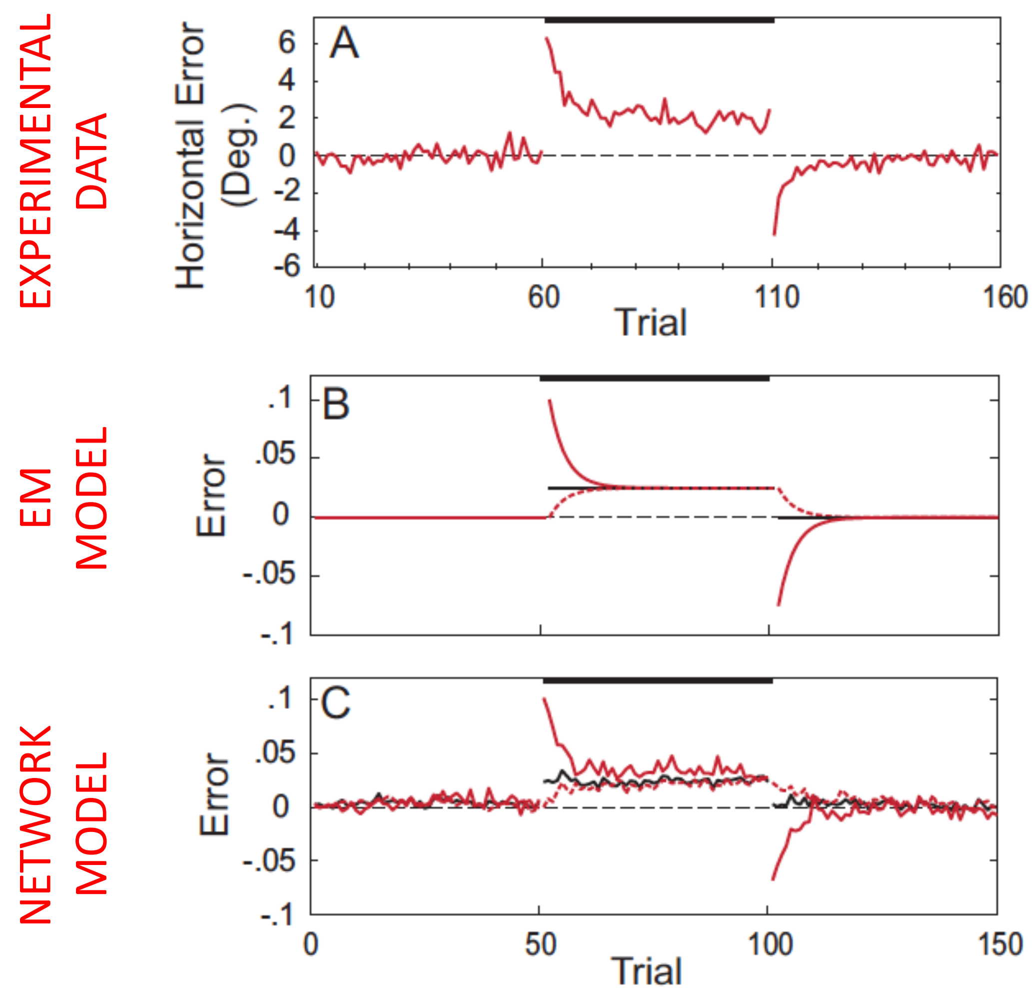

Biddle Snead and Sabes, *unpub.*

Adaptive sensory integration

Models capture spatial dynamics of learning, learning spatial anisotropy *de novo*

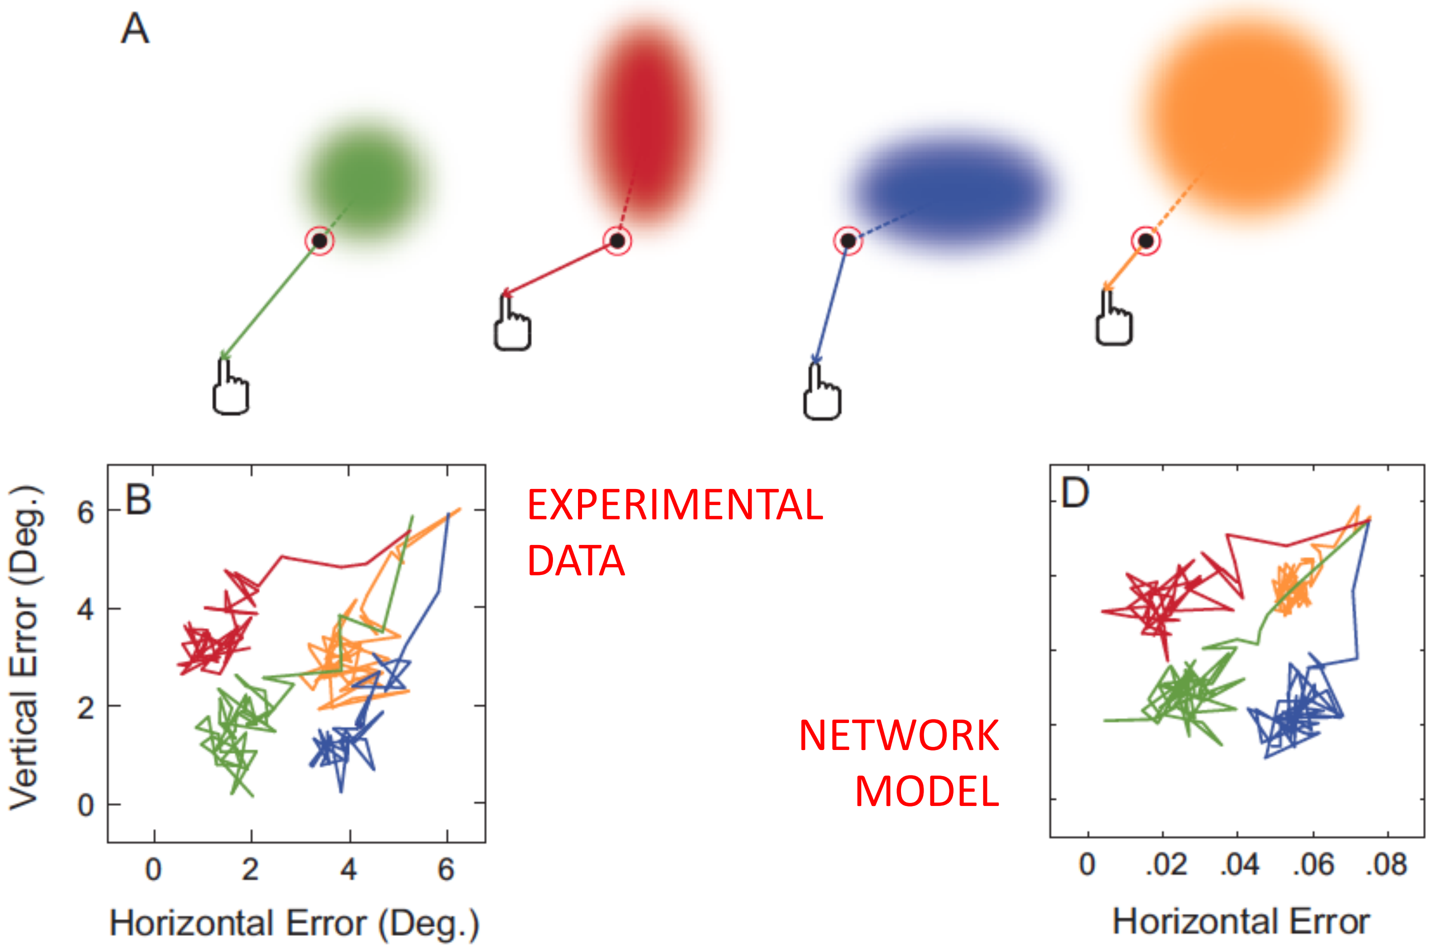

Burge and Banks, J Vision (2007)

Biddle Snead and Sabes, *unpub.*

Incorporating recent experience through learning

Tim Verstynen

Incorporating recent experience through learning

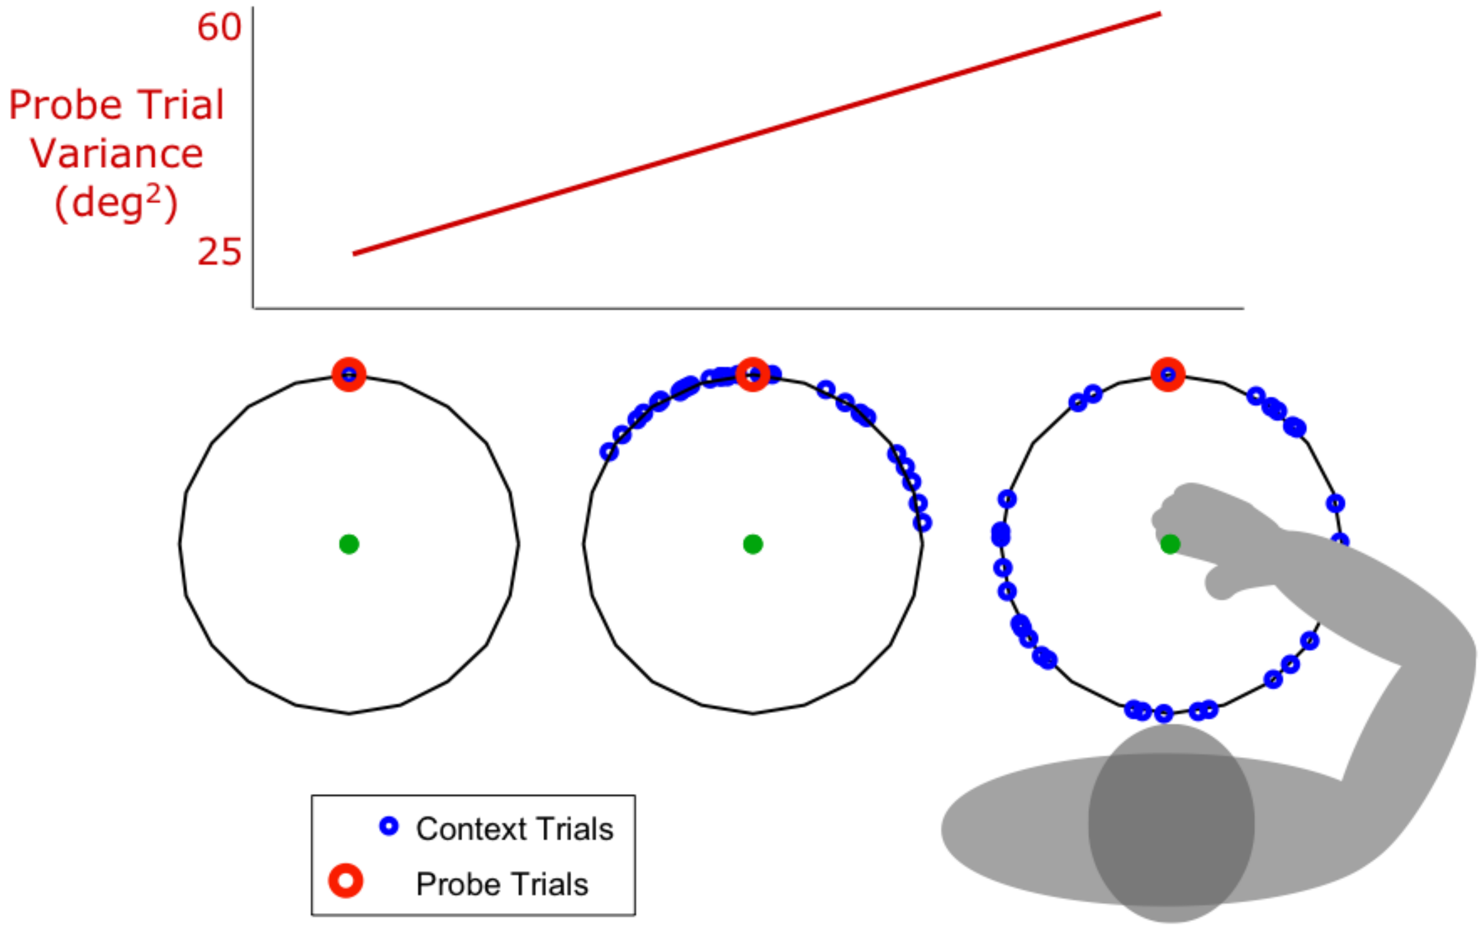

Practice makes perfect on a short timescale … but at a cost.

Incorporating recent experience through learning

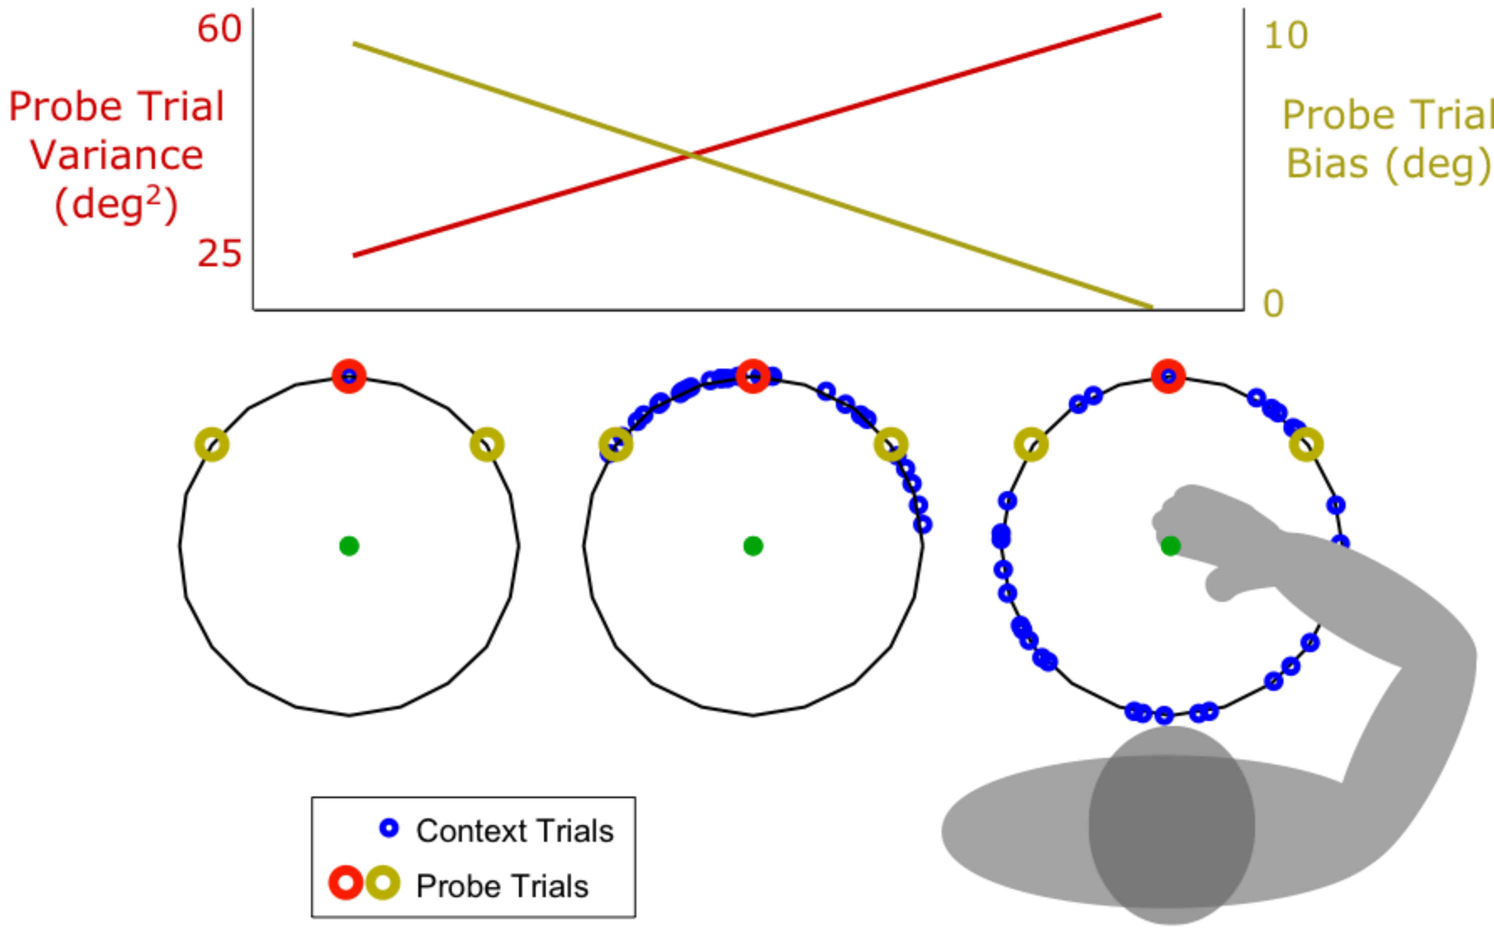

There is a bias-variance tradeoff

Incorporating recent experience through learning

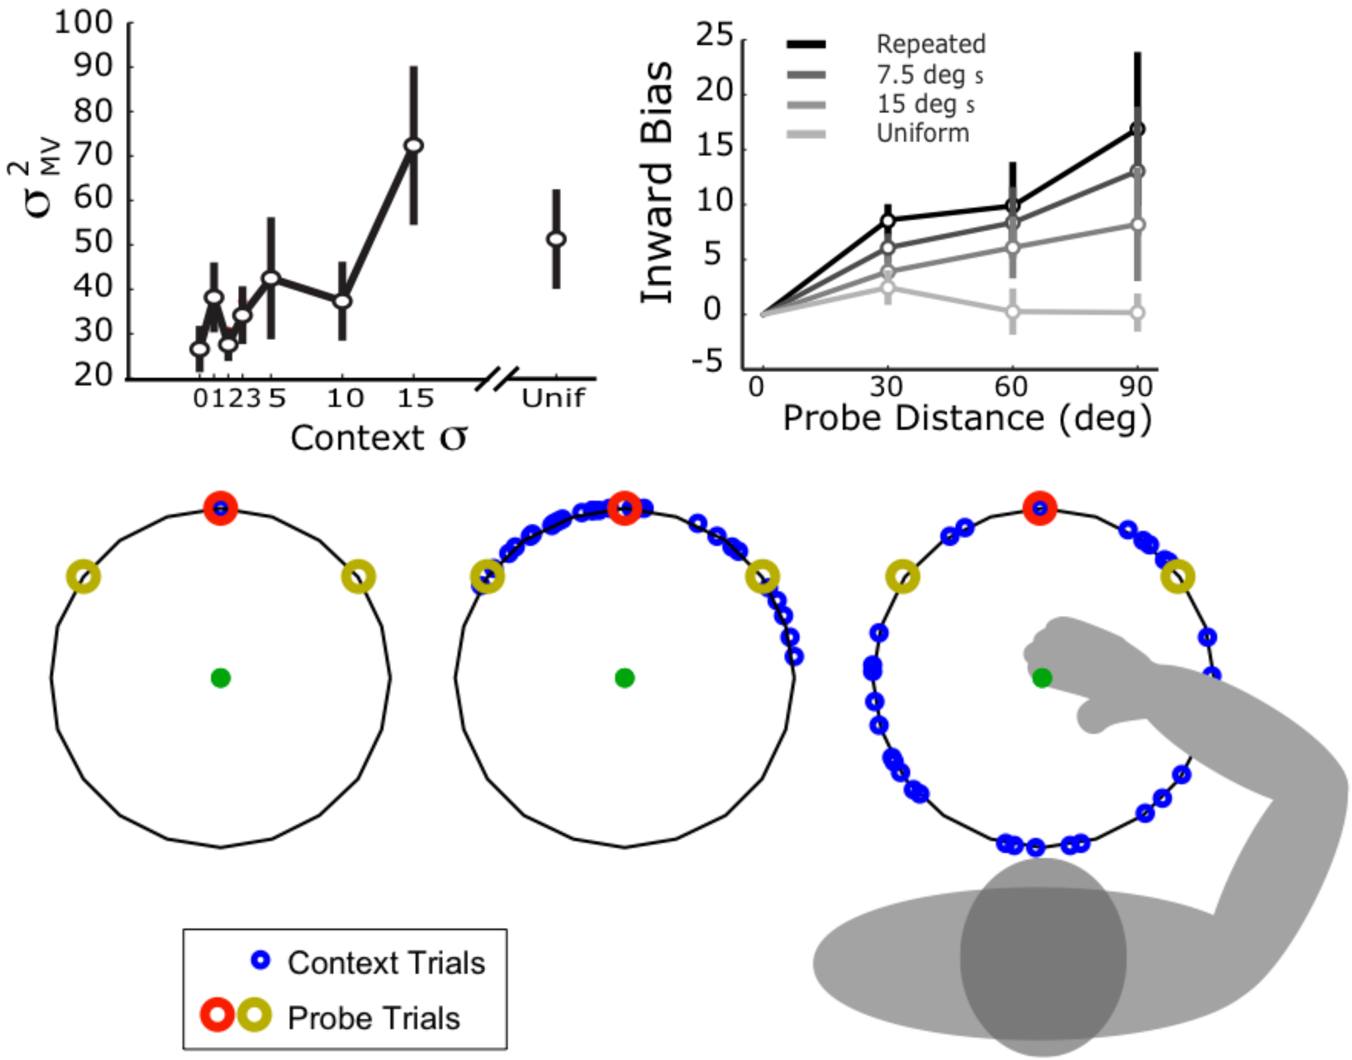

What causes this bias-variance tradeoff?

Bayesian integration of prior data

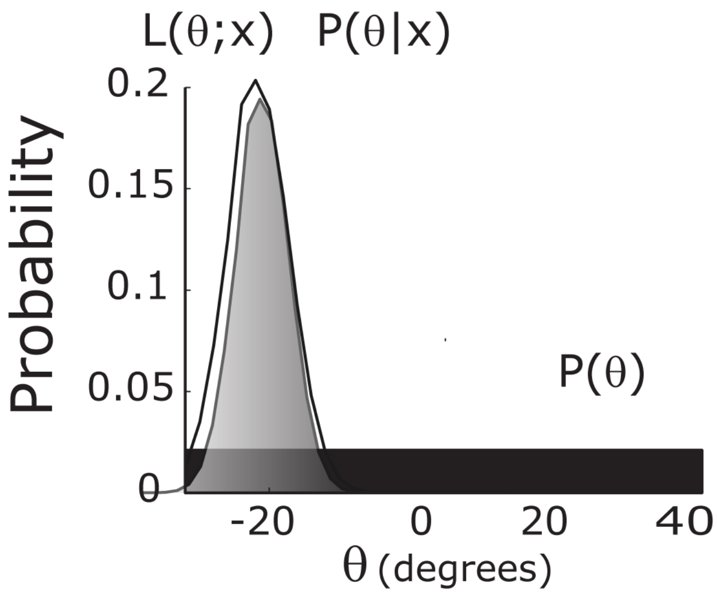

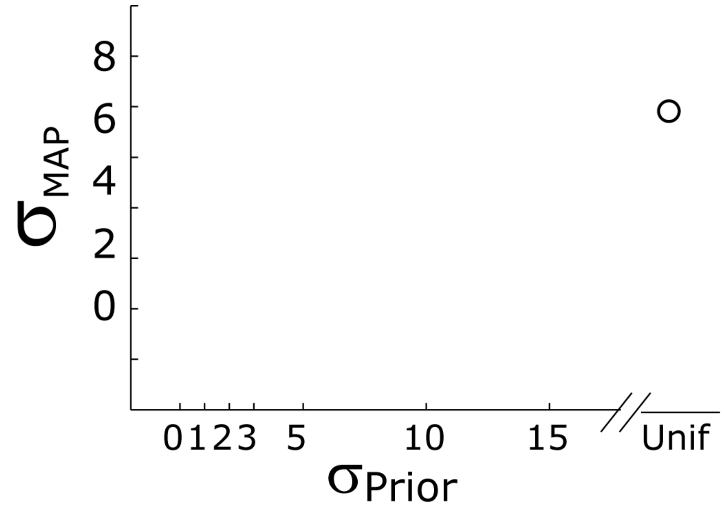

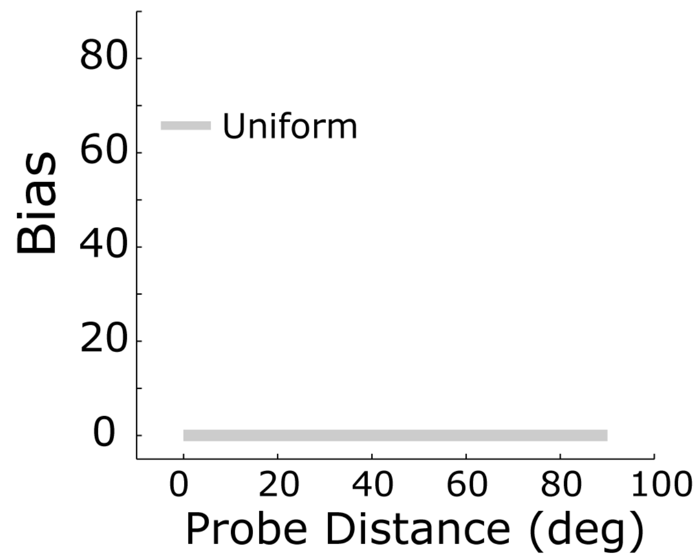

Posterior: $P(\theta|x) \propto L(\theta;x)P(\theta)$

Bayesian integration of prior data

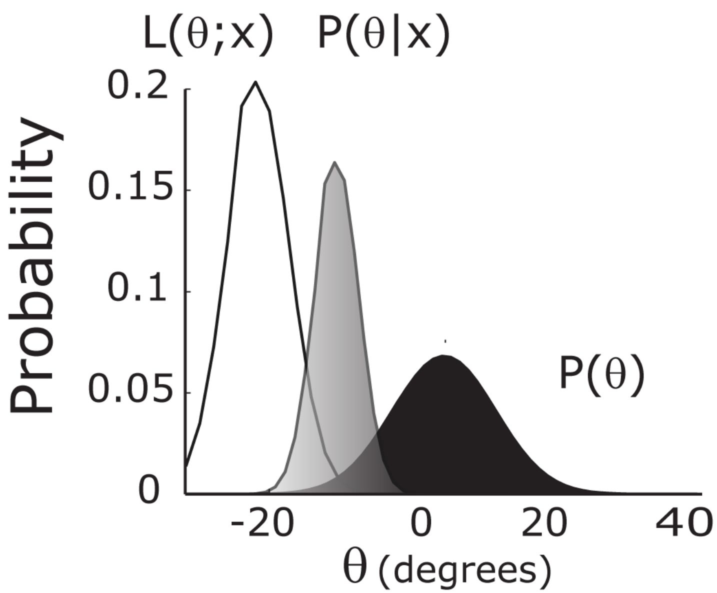

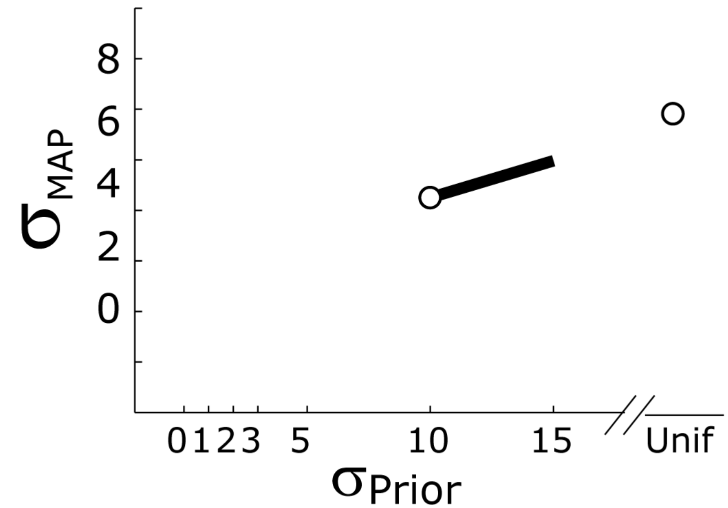

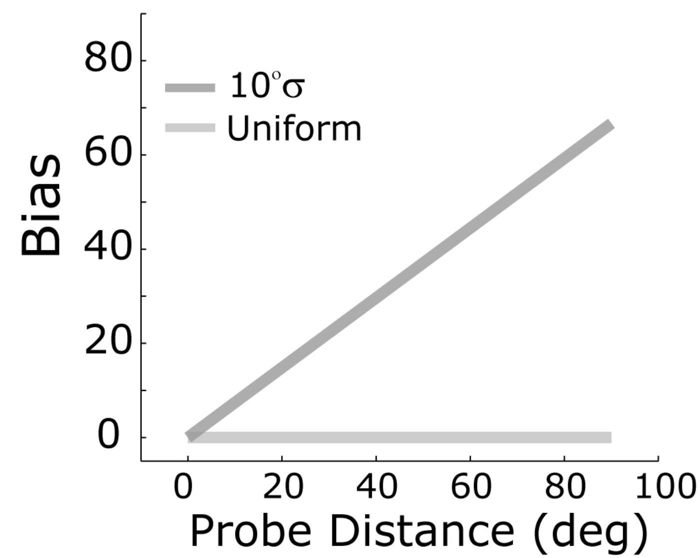

Posterior: $P(\theta|x) \propto L(\theta;x)P(\theta)$

Bayesian integration of prior data

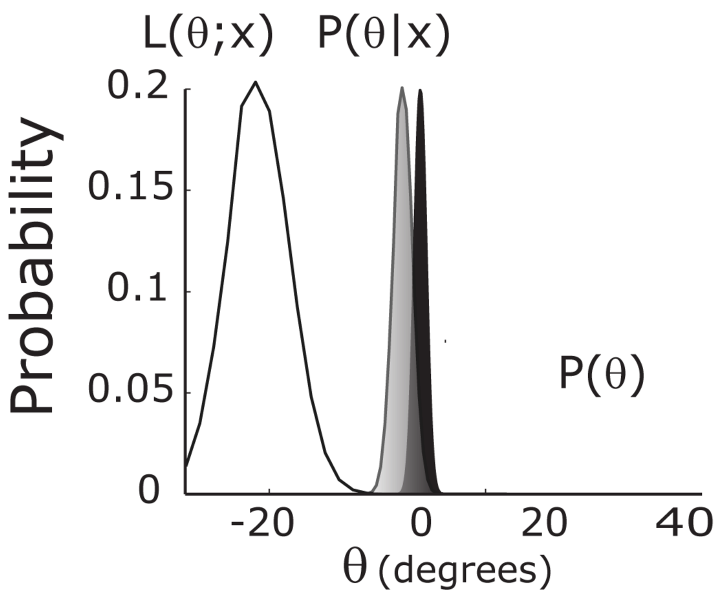

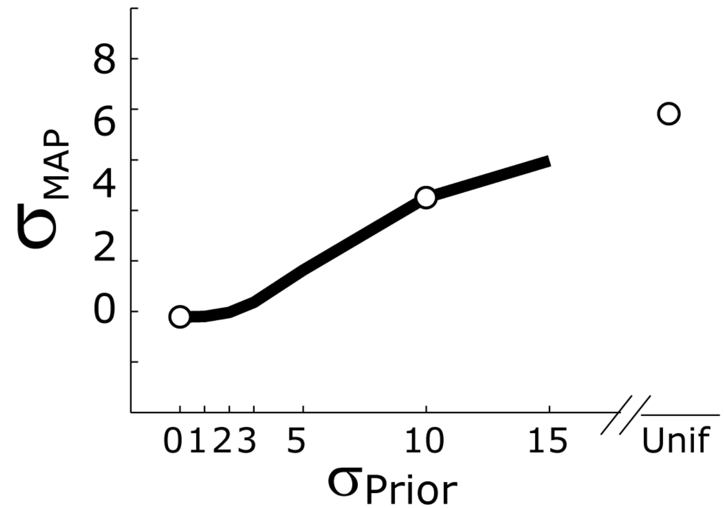

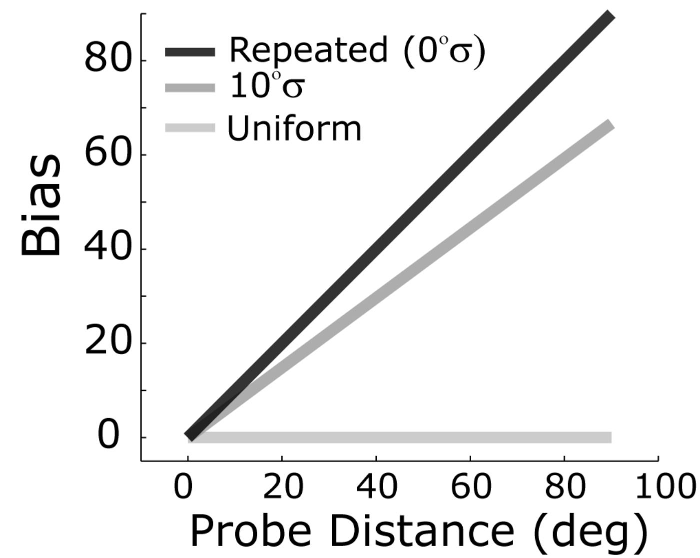

Posterior: $P(\theta|x) \propto L(\theta;x)P(\theta)$

Comparison of data and Bayesian model

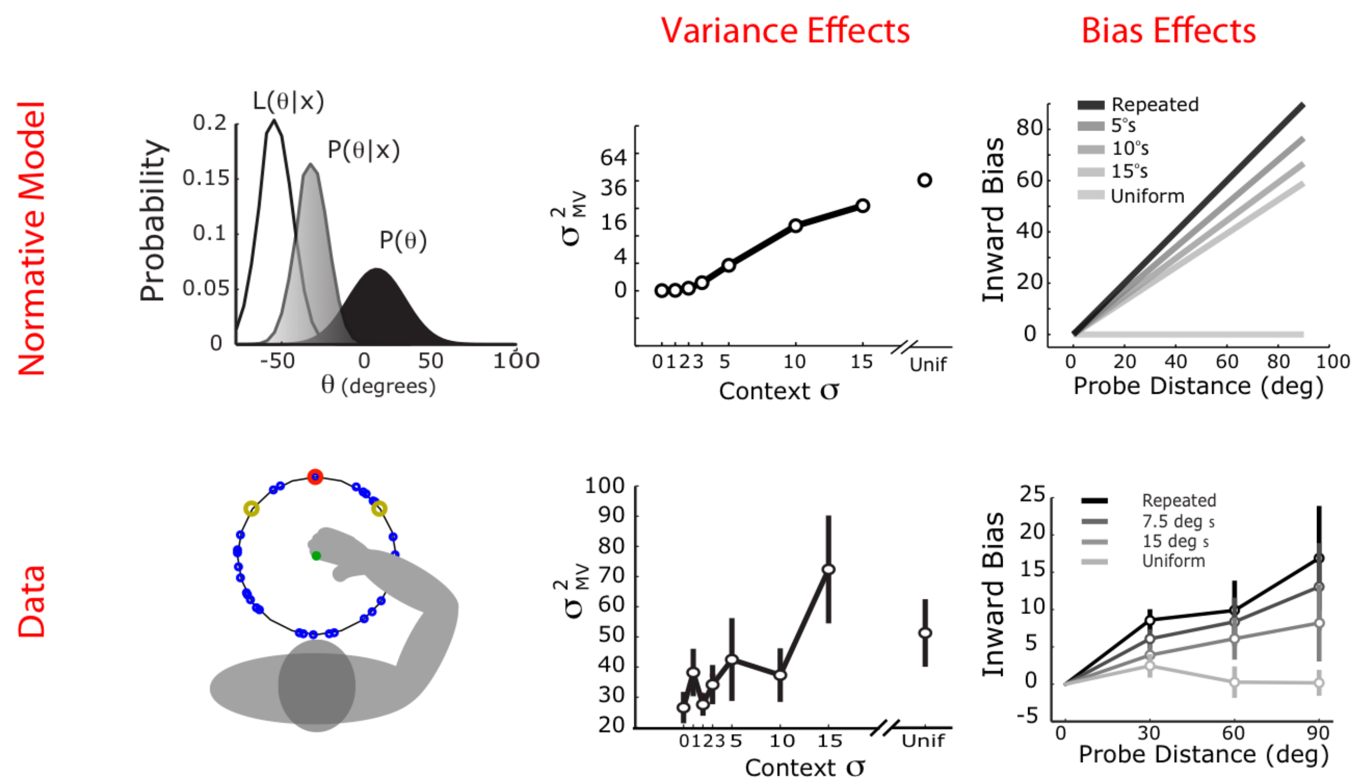

- Experimentally observed bias-variance tradeoff is **qualitatively** consistent with Bayesian integration
- Not **quantitatively** accurate
- But it doesn't take learning into account

Comparison of data and adaptive model

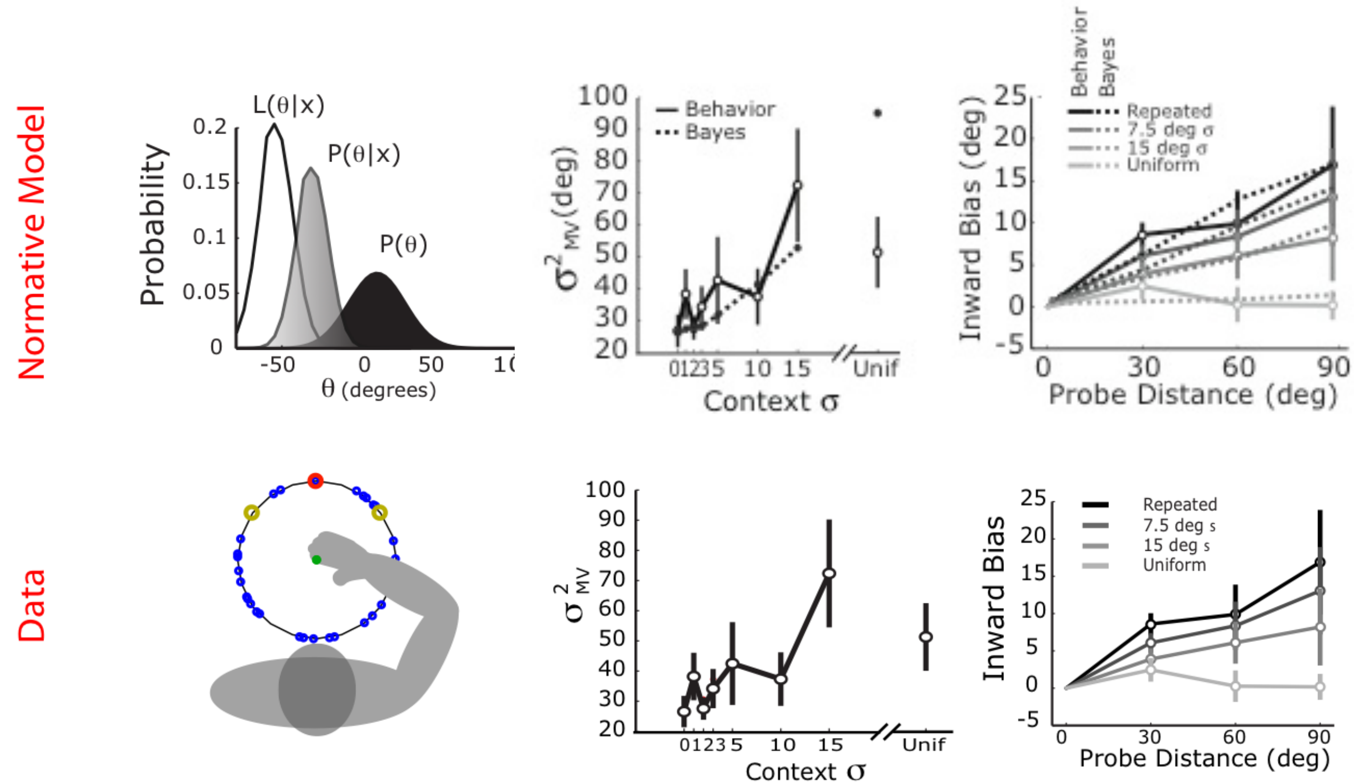

- Simple adaptive Bayesian model:

$$\mu_t = (1-\alpha)\mu_{t-1} + \alpha\theta$$
$$\sigma_t^2 = (1-\alpha)\sigma_{t-1}^2 + \alpha(\mu-\theta)^2$$

- Simple **(normative)** adaptive model captures data fairly well
- How can you distinguish this model from alternatives?

Not just guessing

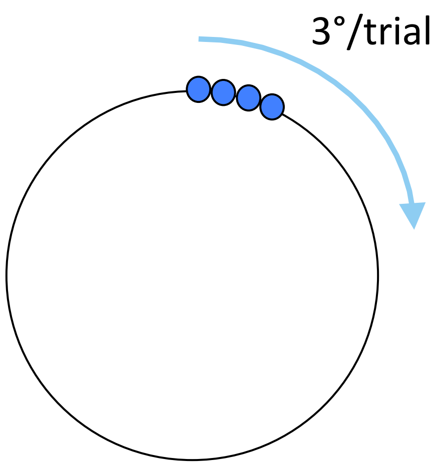

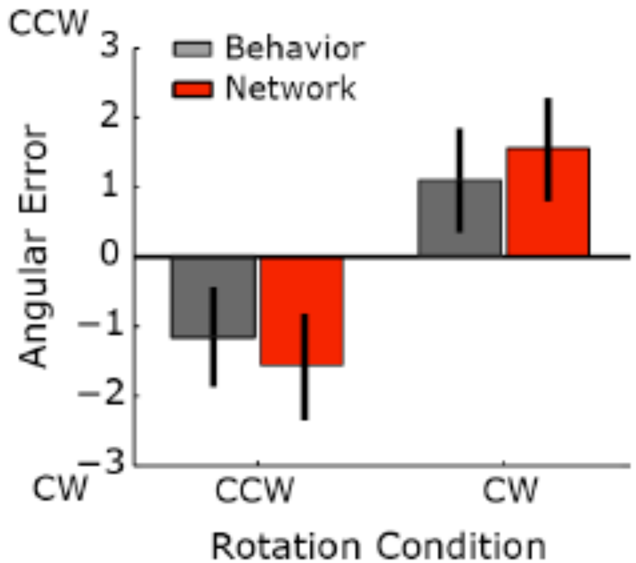

- Not just a "cognitive strategy", e.g. guessing where the target is
- Another experiment where guessing makes a different prediction than an "automatic" process

Network model: associative learning as adaptive prior

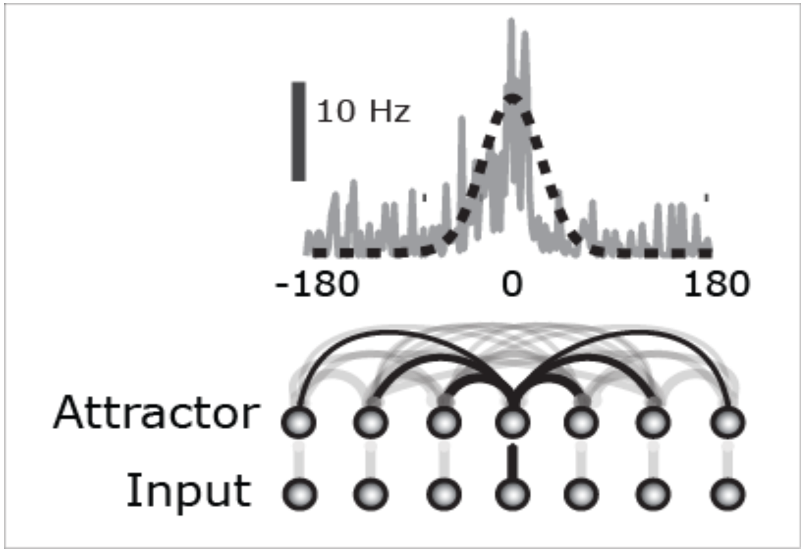

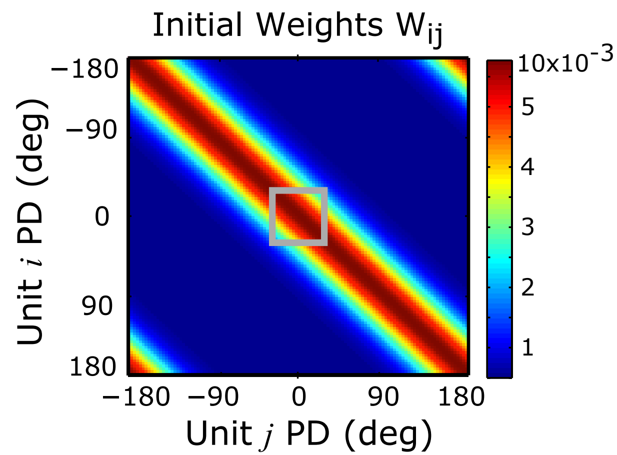

Deneve, Latham, Pouget (1999)

- Associative (Hebbian) learning in lateral connections
  - Recent activity patterns (experiences) reshape connectivity
  - New connectivity reshapes "energy landscape"
  - We show that network acts Bayesian:
    - Train with one contest, test with various input gains
    - Get variance and bias as predicted
    - See also Wu and Amari (2005)

Network model: associative learning as adaptive prior

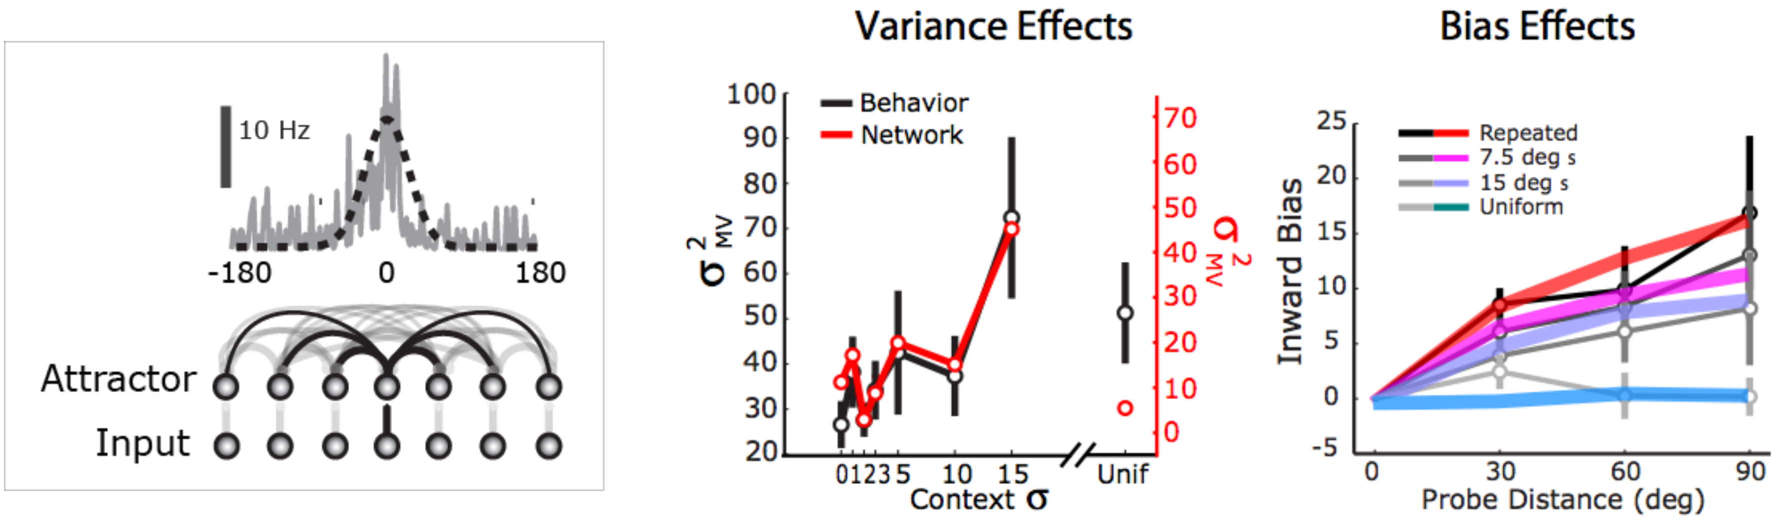

- Good match to experimental results when network experiences same trial sequence as subjects
- Associative (Hebbian) learning in lateral connections
  - Recent activity patterns (experiences) reshape connectivity
  - New connectivity reshapes "energy landscape"
  - Network acts Bayesian

Learning due to changes in network weights

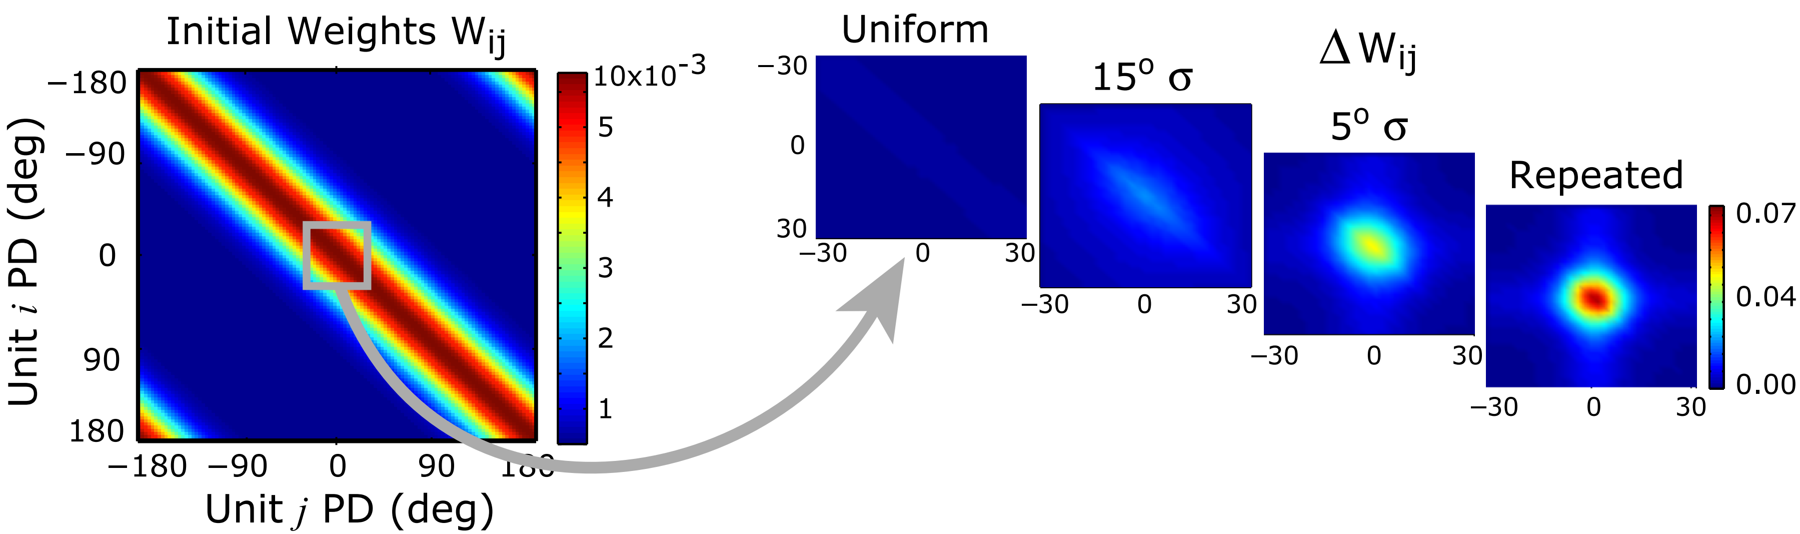

Network mimics Bayesian estimation

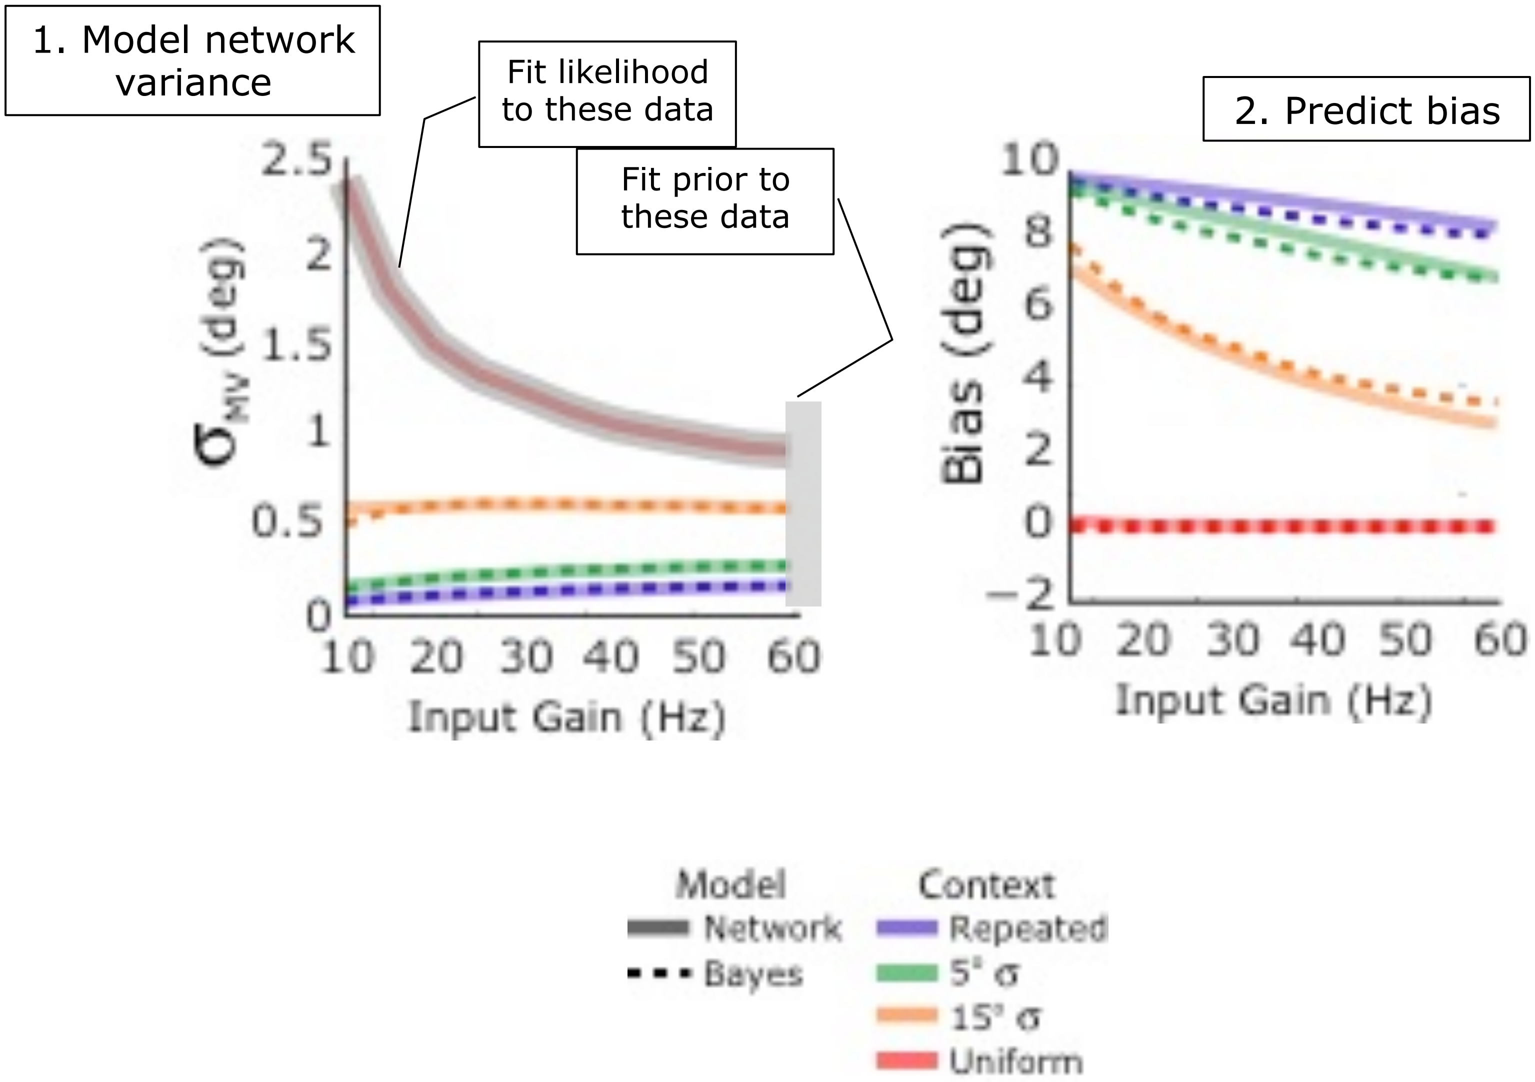

Network implementation sensorimotor processing

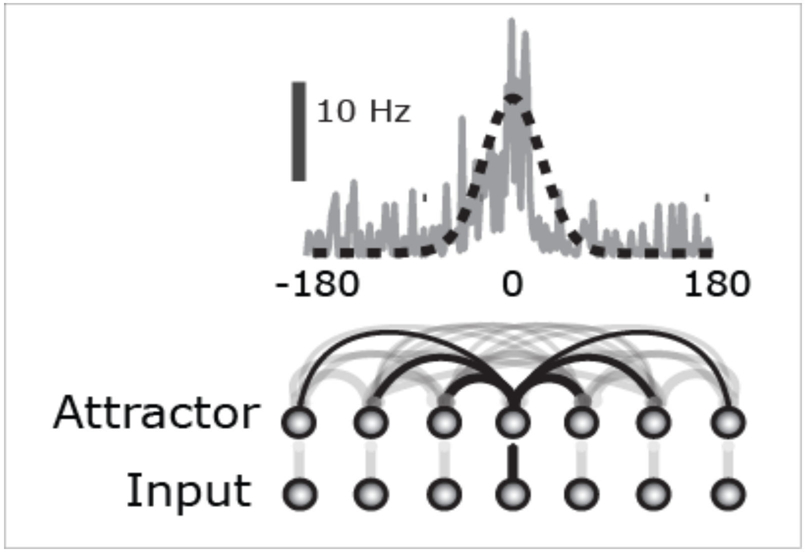

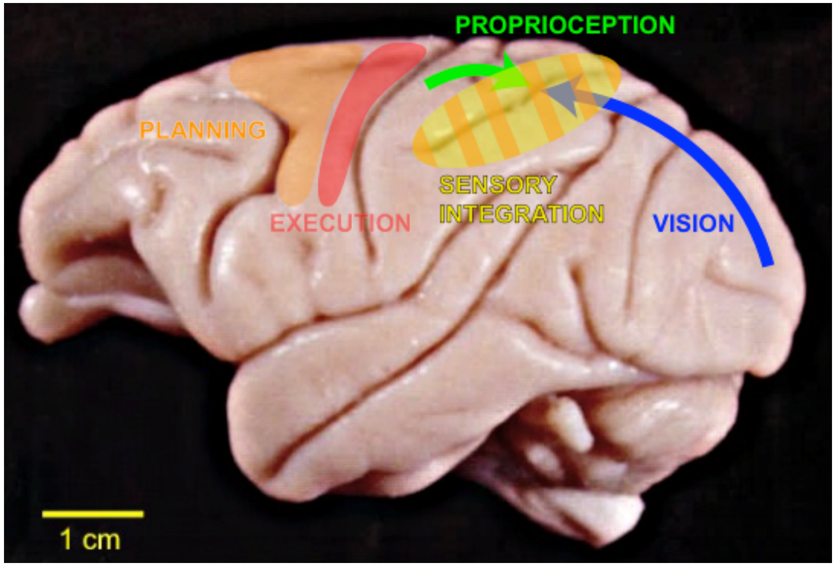

- Good news: simple line-attractor (with Hebbian learning) can implement key elements of sensorimotor processing: sensory integration, Bayesian inference, and coordinate transformations
- Bad news: these models are too simple to explain real fronto-parietal circuits
  - Hand-wired to work with simple representations

Network implementations of the building blocks

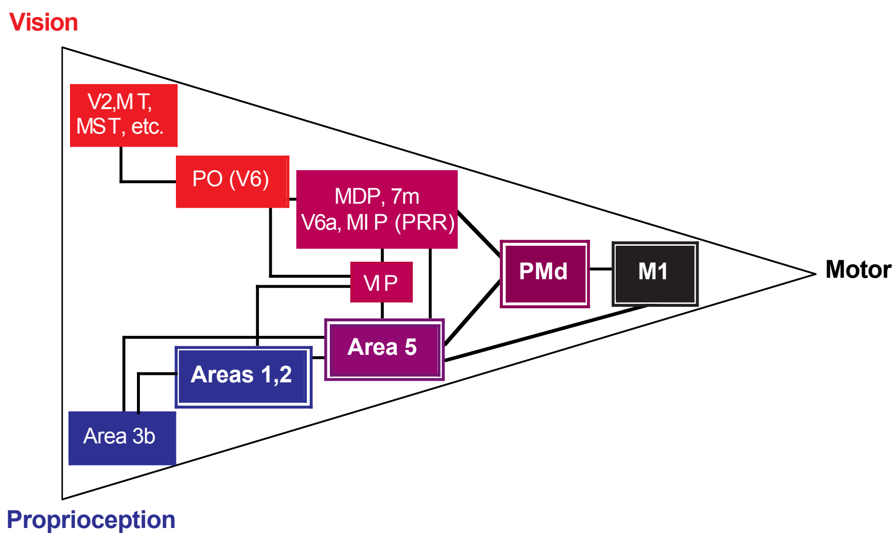

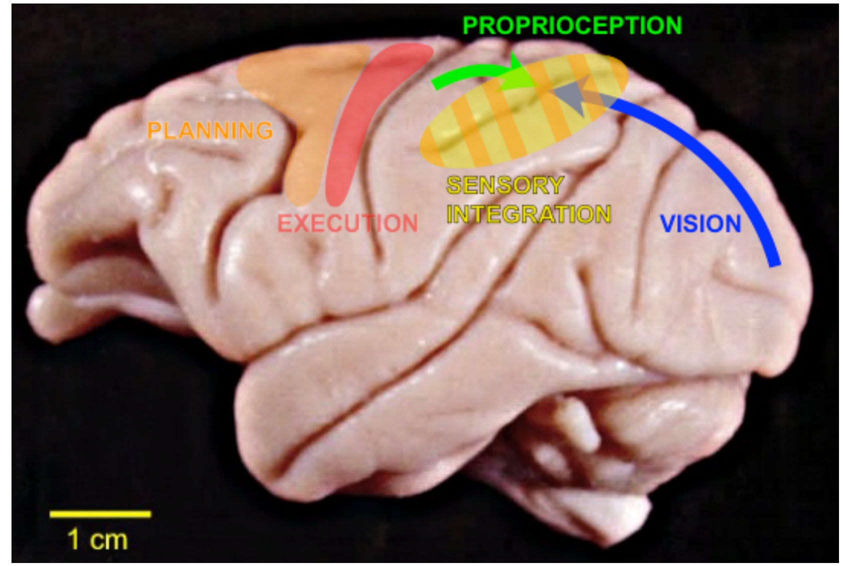

- Good news: simple line-attractor (with Hebbian learning) can implement key elements of sensorimotor processing: sensory integration, Bayesian inference, and coordinate transformations
- Bad news: these models are too simple to explain real fronto-parietal circuits
  - Hand-wired to work with simple representations
  - No preservation of statistical information – can't combine hierarchically
- Perhaps these limitations can be overcome with more sophisticated architectures and learning algorithms

Sensory integration as statistical model-building

Joseph Makin

Matthew Fellows

Sensory integration as statistical model-building

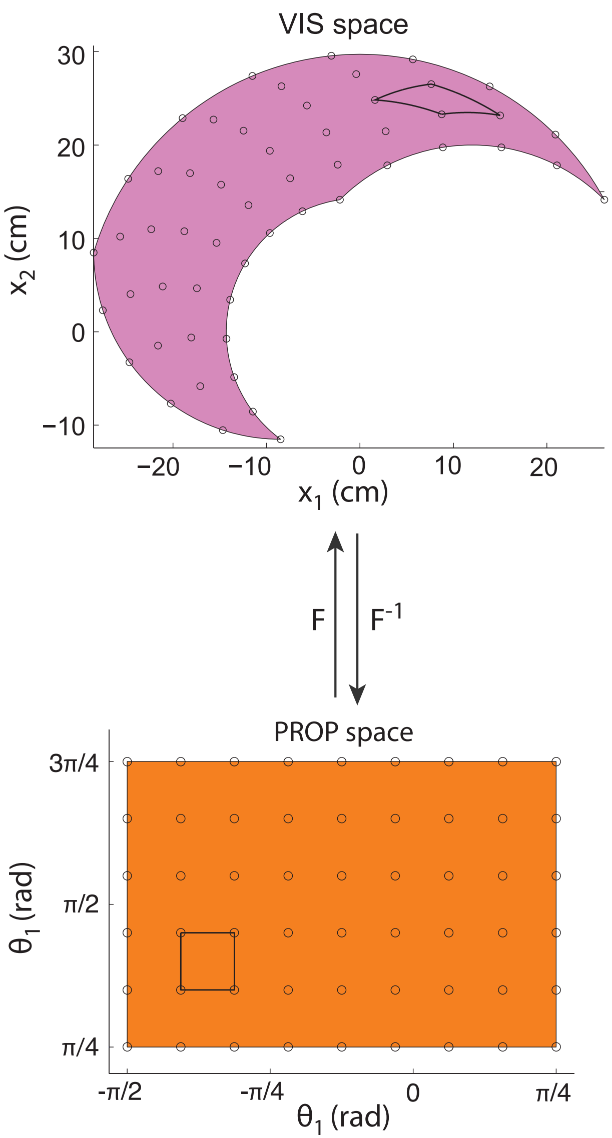

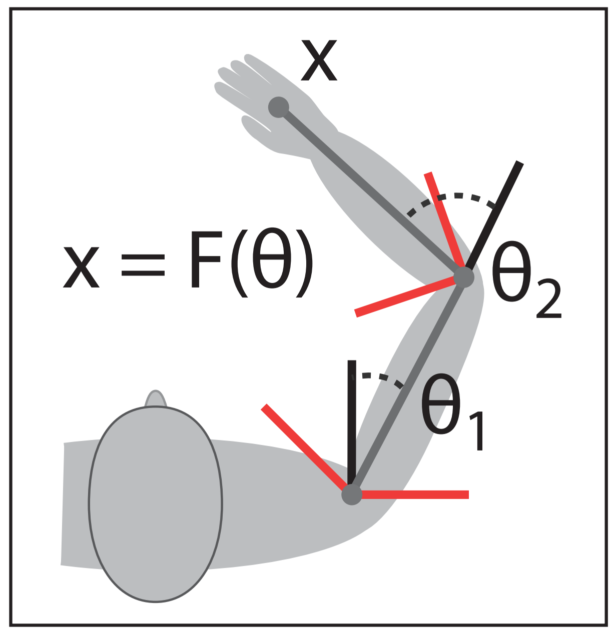

Makin, Fellows, and Sabes, *in revision*

Sensory integration as statistical model-building

- Restricted Boltzmann Machine
- Contrastive divergence (Hinton et al)

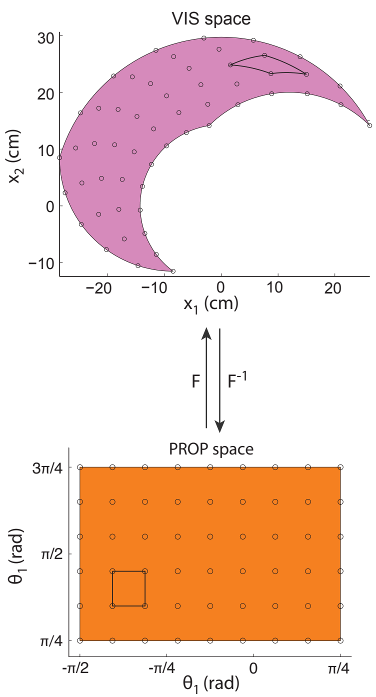

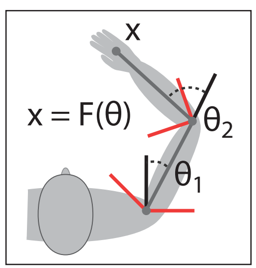

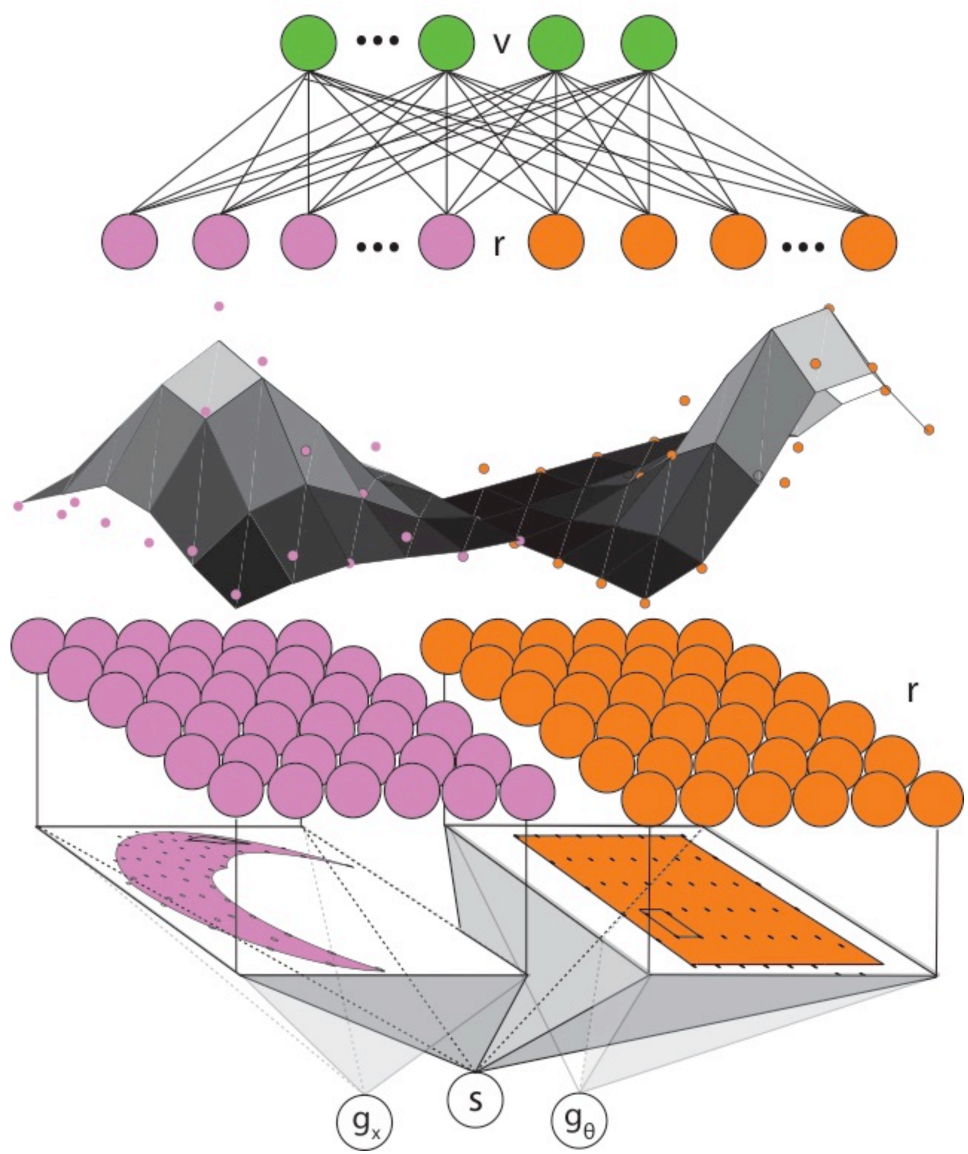

Sensory integration as statistical model-building

Network learns *de novo* the statistical relationships between the inputs:

- Integration

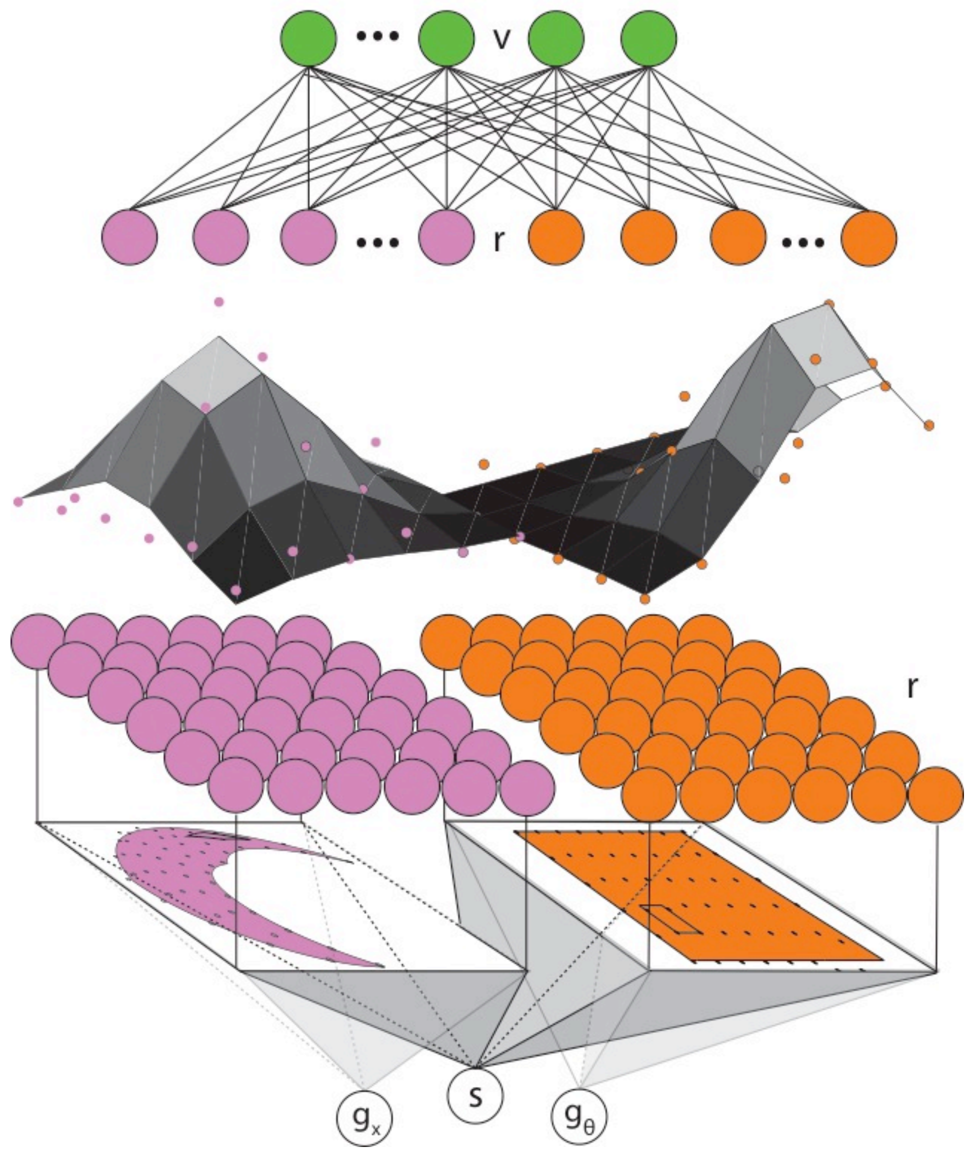

Maximum Likelihood sensory integration

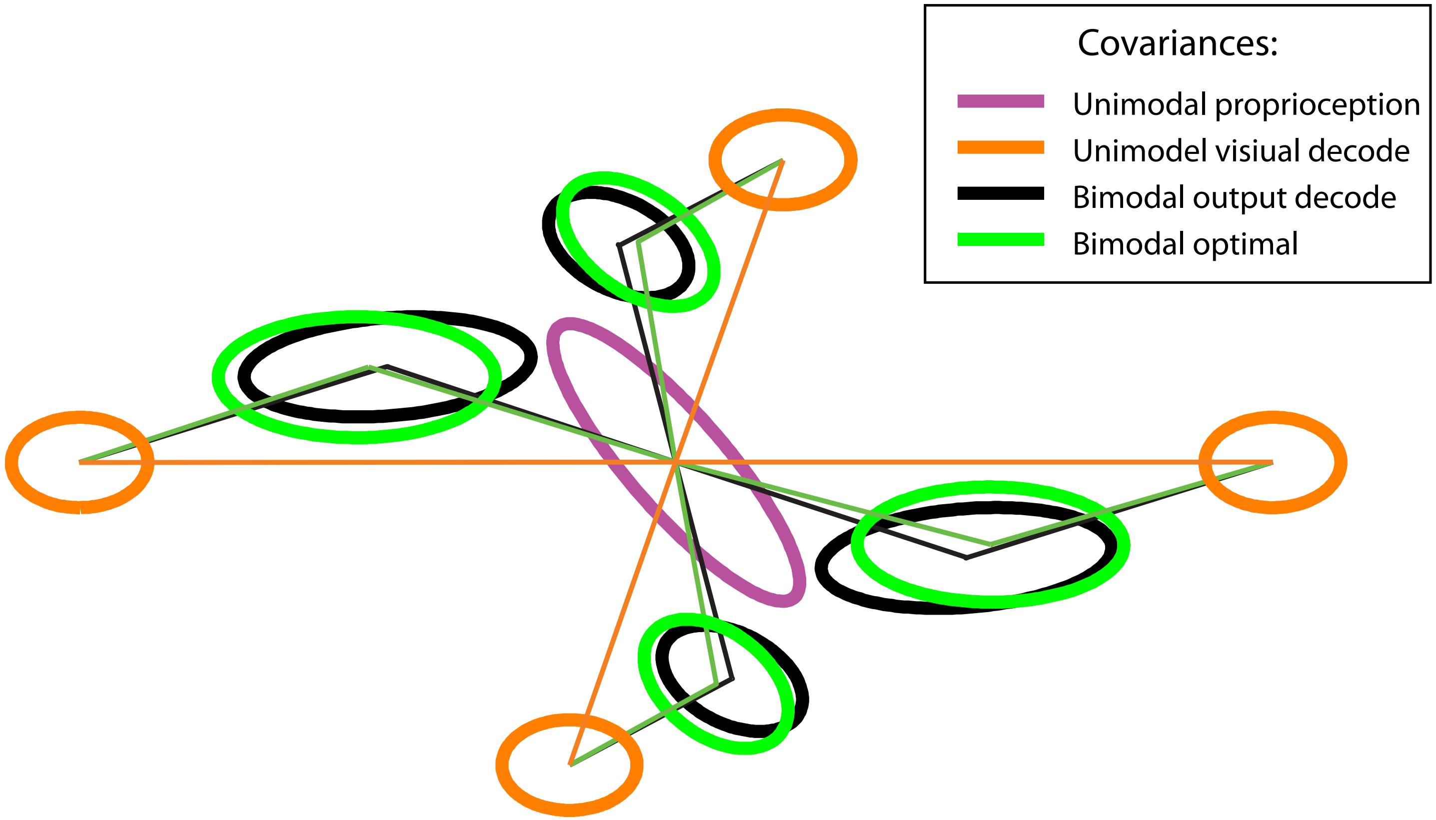

Makin, Fellows, and Sabes, *in revision*

Sensory integration as statistical model-building

Network learns *de novo* the statistical relationships between the inputs:

- Integration
- Sensory transformation
- Bayesian priors
- Hierarchical combination

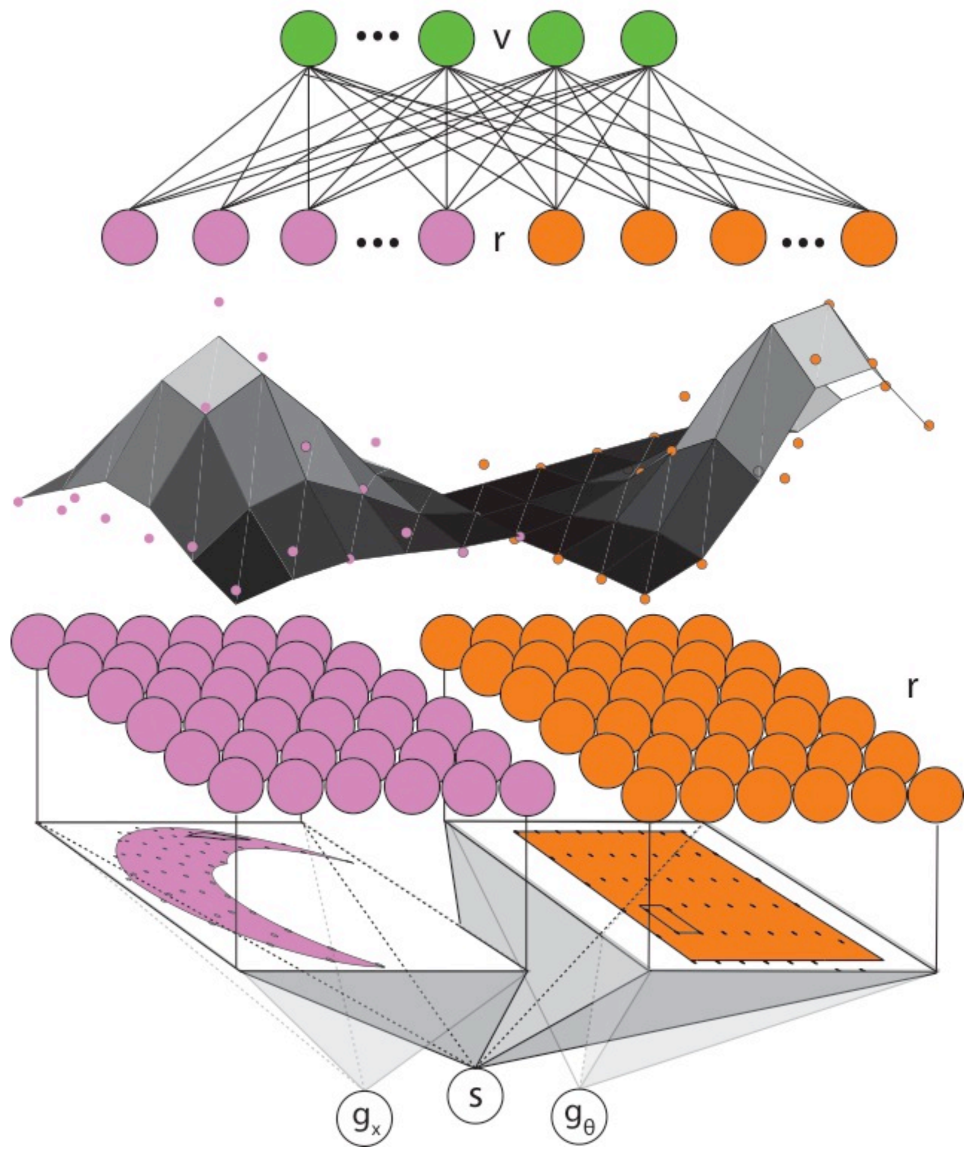

Hierarchical composition

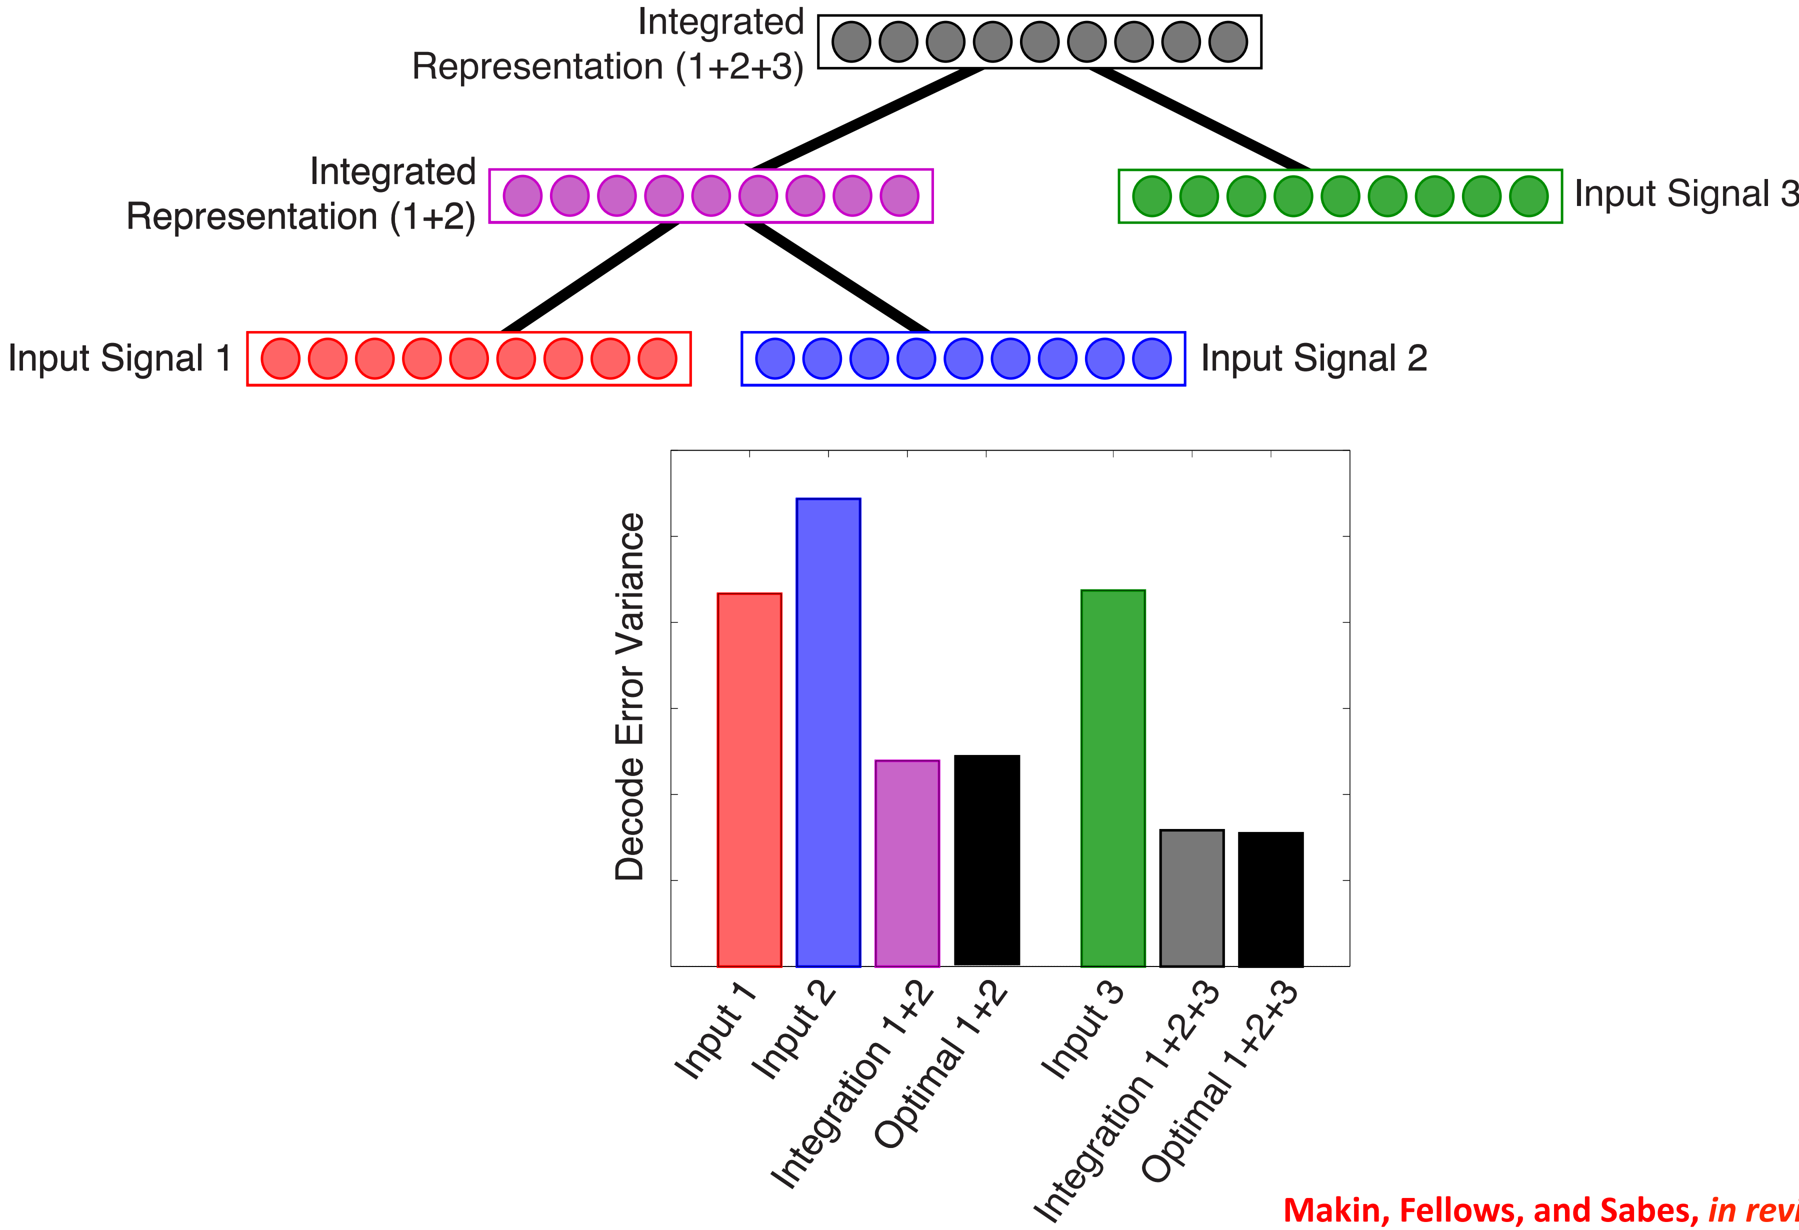

Makin, Fellows, and Sabes, *in revision*

Sensory integration as statistical model-building

- Results:
  - Integration of signals related via a non-linear transformation
  - Learns priors
  - Learns additive transformations
  - Can be composed hierarchically
  - See also Yildirim and Jacobs *Vis. Res.* (2012)
- Other learning algorithms are also likely to work
  - Key is the ability to learn from the local statistics in incoming signals
- Composition and the ability to handle more complex representations – candidate schema for the posterior parietal cortex?

Makin, Fellows, and Sabes, *in revision*