The Bayesian Brain:
Ideal observer models for perceptual decisions

Jan Drugowitsch

Department of Neurobiology
Harvard Medical School

HARVARD MEDICAL SCHOOL

Neuro 140
February 6, 2024

We need theories to constrain our hypothesis space!

How do we develop useful theories?

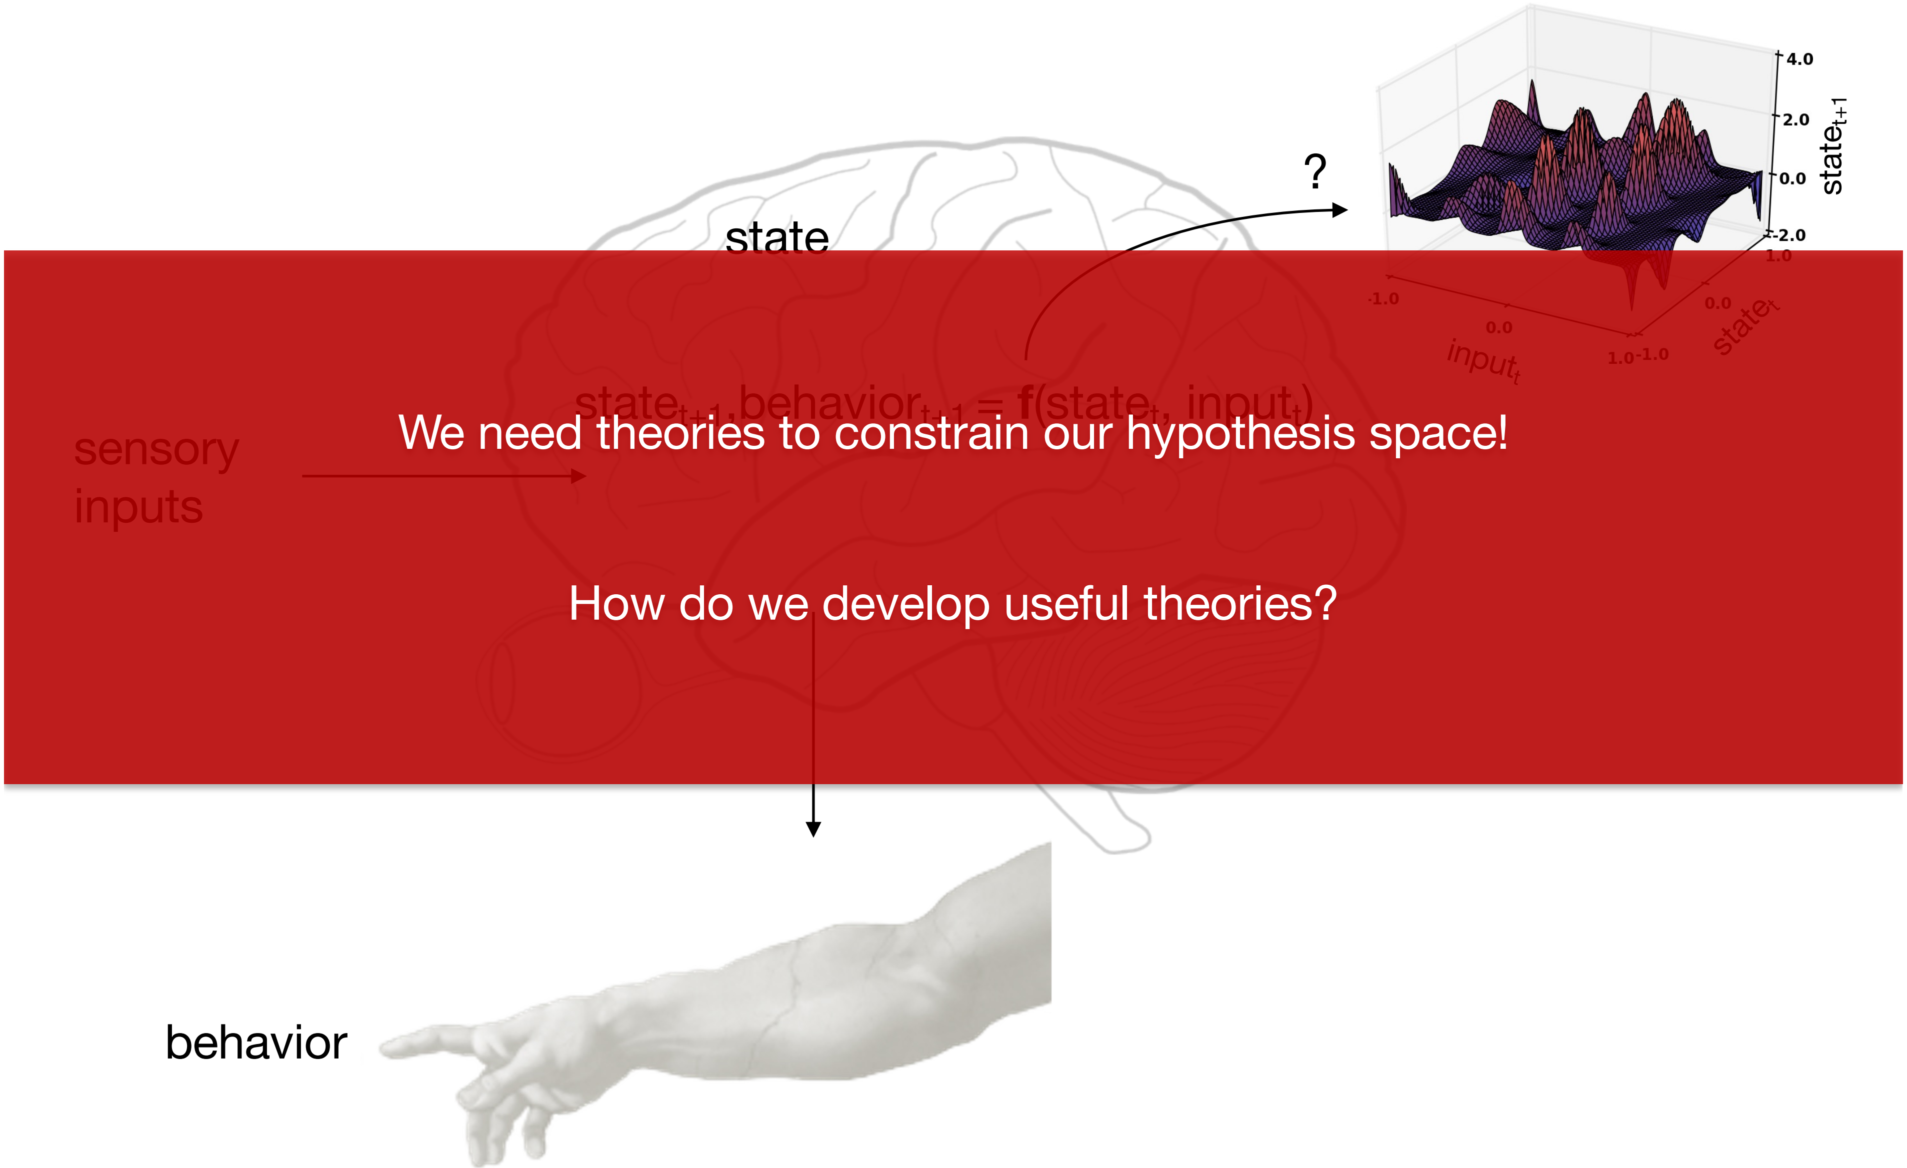

Roadmap

Ideal observer modeling:
Uncertainty as guiding principle
in AI and computational neuroscience

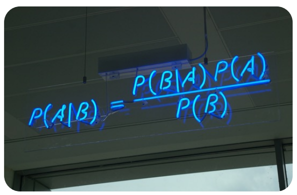

(Behavioral) evidence for handling uncertainty

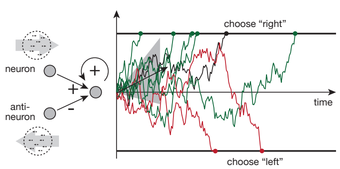

Example 1: ideal observer models for the speed/accuracy trade-off
in perceptual decision-making

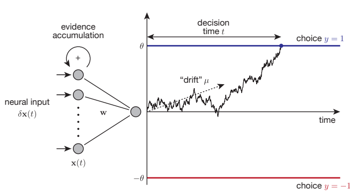

Example 2: Inference on a difference time-scale:
Decision confidence to improve decision strategies

Roadmap

Ideal observer modeling:
Uncertainty as guiding principle
in AI and computational neuroscience

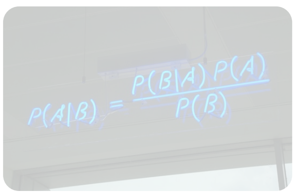

(Behavioral) evidence for handling uncertainty

Example 1: ideal observer models for the speed/accuracy trade-off
in perceptual decision-making

Example 2: Inference on a difference time-scale:
Decision confidence to improve decision strategies

Guiding principle: information is uncertain

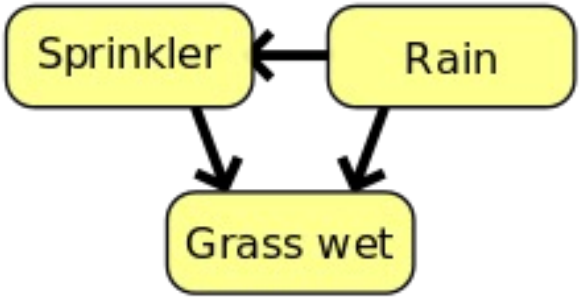

Uncertainty handling in artificial intelligence

(a few examples)

Boltzmann machines (stochastic Hopfield networks; Hinton & Sejnowski, 1983)

Bayesian networks (Pearl, 1985)

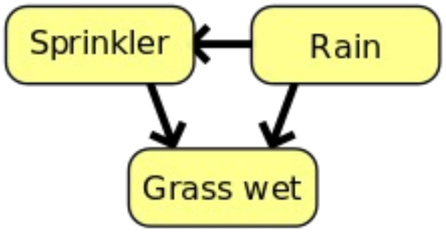

Statistical learning theory (Vapnik & Chervonenkis, 1971)
- brought us Support Vector Machines (Cortes & Vapnik, 1995)

Variational Bayes; MCMC; …

…

Deep learning (~2012): initially no uncertainty

…

Variational autoencoders (Kingma & Welling, 2014; Rezende et al., 2014)
- build statistical model of inputs

Distributional reinforcement learning (Bellemare et al., 2017)
- build statistical model of long-term rewards

Diffusion models, e.g., Stable Diffusion (Sohl-Dickstein et al., 2015; Ho et al., 2020)
- image generation by learning to revers stochastic diffusion process

Ideal observer modeling

Brain

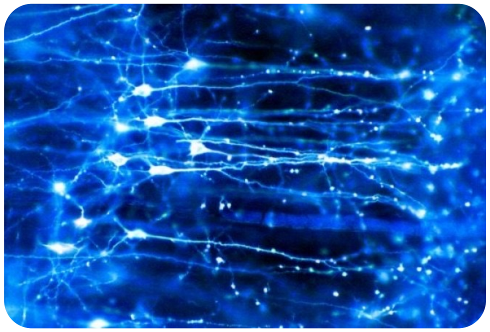

AI

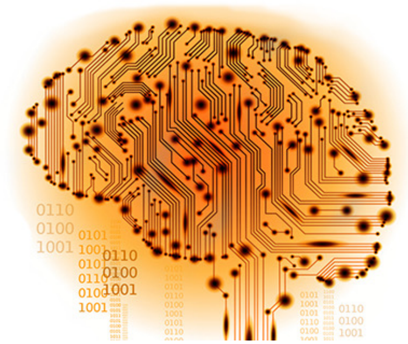

today's lecture

Needs to efficiently handle uncertain information

↓

"ideal observer" models

methods for (approximate) inference with uncertain information

↕

potential neural implementations

→

new algorithmic ideas (e.g., boltzmann machines, networks in general)

Principled way of handling uncertainty

Rev. Thomas Bayes
(1701-1761)

Using **Bayesian decision theory** to handle uncertainty

$$\underbrace{p\left(\begin{matrix}\text{state}\\\text{of world}\end{matrix}\;\middle|\;\begin{matrix}\text{sensory}\\\text{evidence}\end{matrix}\right)}_{\text{posterior}} \propto \underbrace{p\left(\begin{matrix}\text{sensory}\\\text{evidence}\end{matrix}\;\middle|\;\begin{matrix}\text{state}\\\text{of world}\end{matrix}\right)}_{\text{likelihood}} \underbrace{p\left(\begin{matrix}\text{state}\\\text{of world}\end{matrix}\right)}_{\text{prior}}$$

↓ $\mathit{loss}\left(\begin{matrix}\text{true}\\\text{state}\end{matrix}, \begin{matrix}\text{assumed}\\\text{state}\end{matrix}\right)$

decision

Pierre-Simon Laplace (*Théorie analytique des probabilités, 1812*):

"The most important questions of life are indeed, for the most part, really only problems of probability"

Cox's theorem: probabilities are the only principled way to handle uncertainty (Cox, 1946)

Roadmap

Ideal observer modeling:
Uncertainty as guiding principle
in AI and computational neuroscience

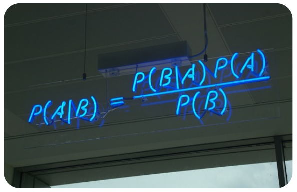

(Behavioral) evidence for handling uncertainty

Example 1: ideal observer models for the speed/accuracy trade-off
in perceptual decision-making

Example 2: Inference on a difference time-scale:
Decision confidence to improve decision strategies

Combining uncertain evidence from multiple sources

e.g. visual/auditory for object localization
visual/vestibular for self-motion
visual/haptic for bar width estimation

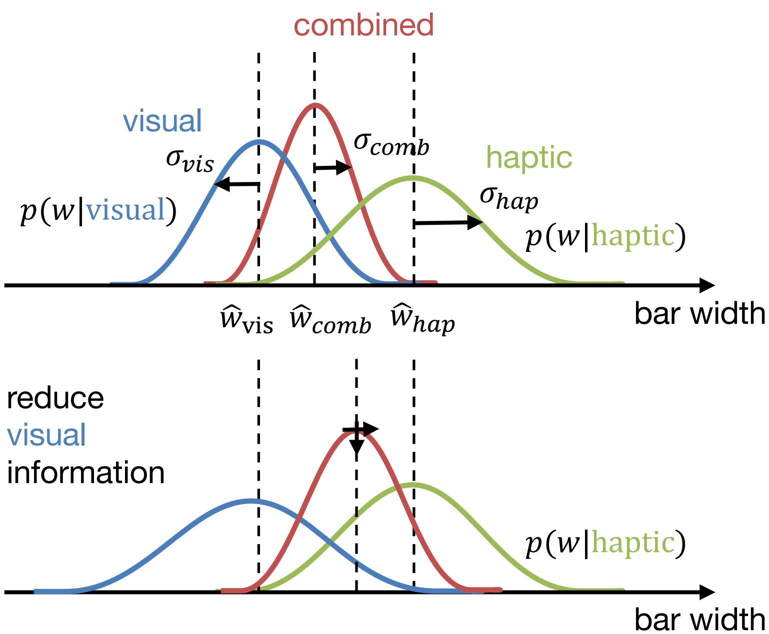

Cue combination using the laws of probability

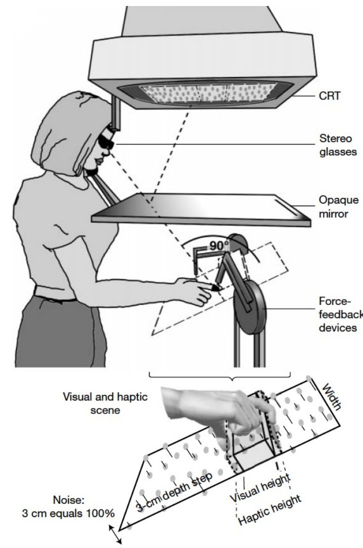

(Ernst & Banks, 2002)

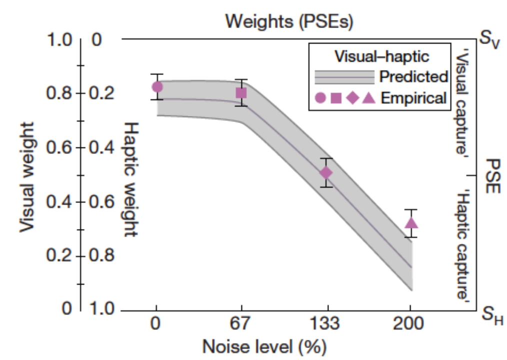

Rely on prior information

Prior = state of the world in absence of evidence

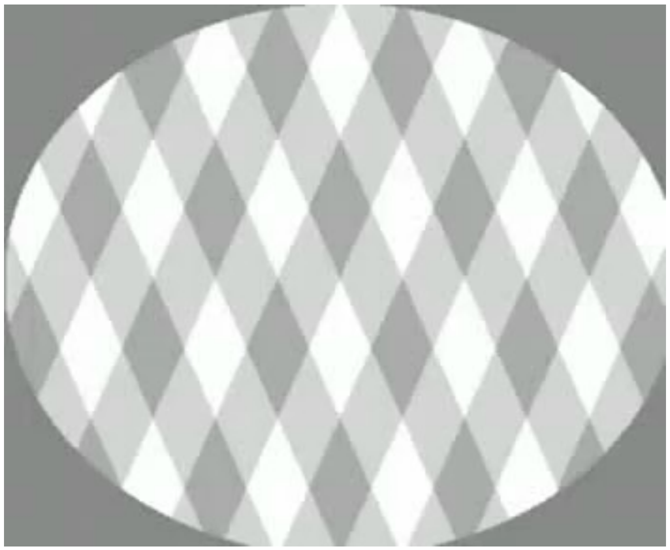

https://www.youtube.com/watch?v=g_sn0WtHK1g

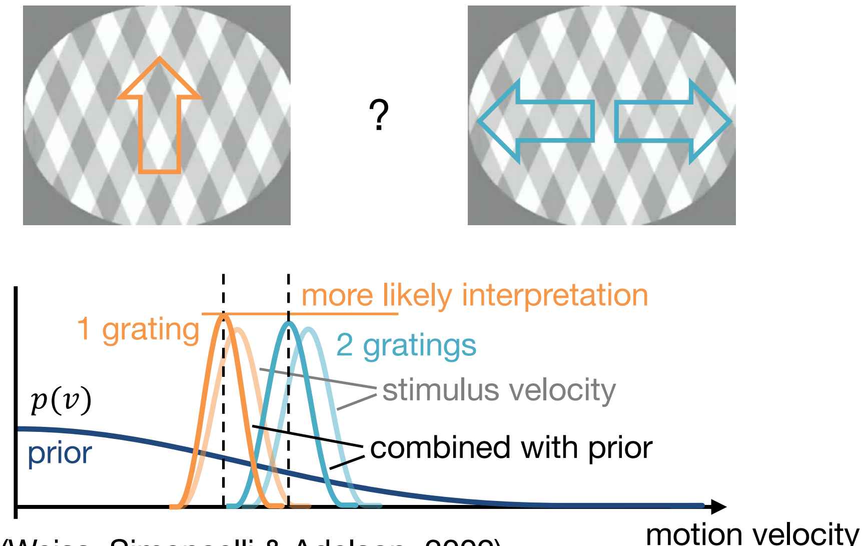

(Weiss, Simoncelli & Adelson, 2002)

Real-world: underestimating speed in bad weather

vs.

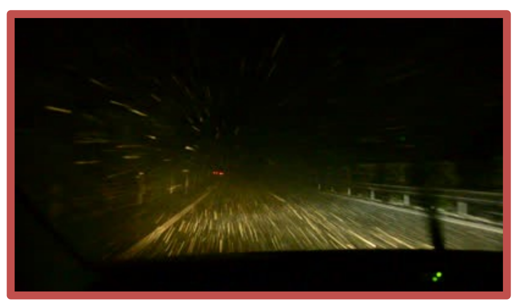

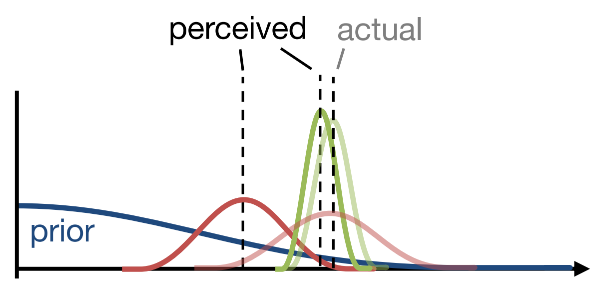

Sensitivity to rewards/losses

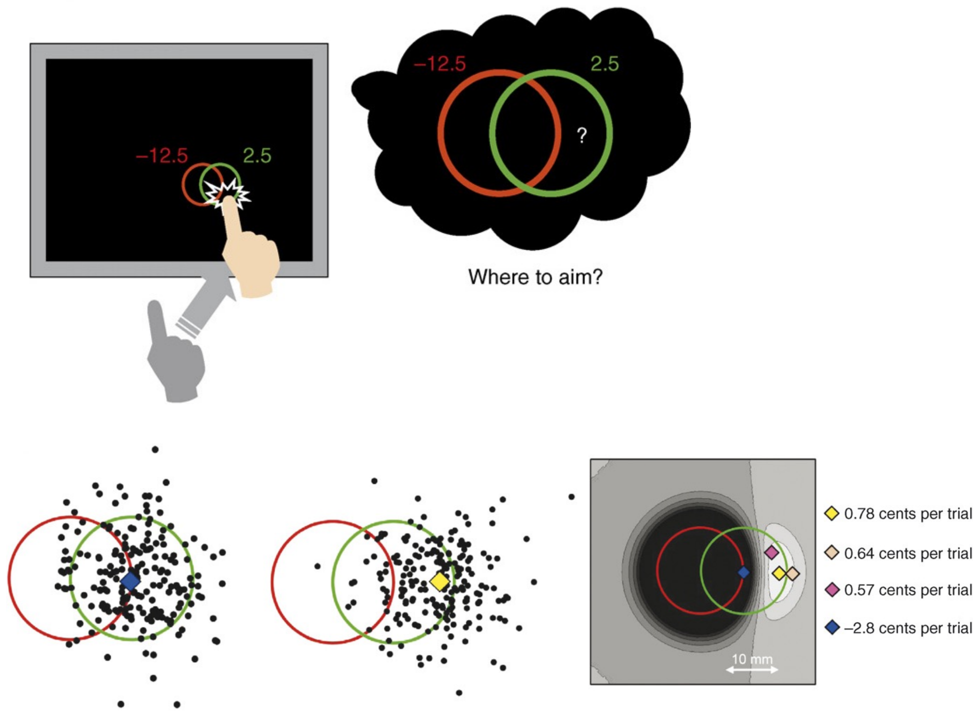

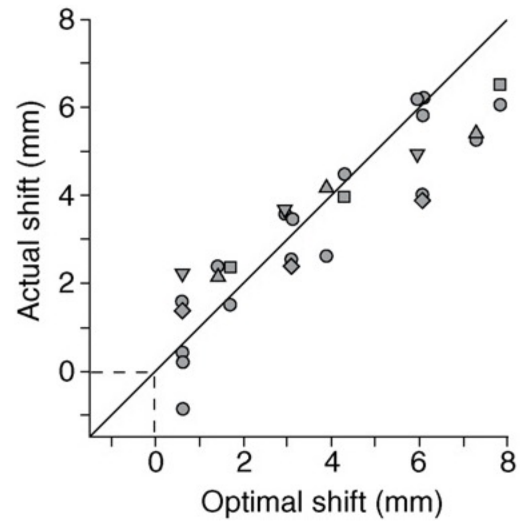

(Trommerhäuser, Maloney & Landy, 2008)

…has also been used to reverse-engineer the reward/loss function

e.g. Körding & Wolpert (2004); Drugowitsch et al. (2012)

Recap: Bayesian decision theory

Rev. Thomas Bayes
(1701-1761)

$$\underbrace{p\left(\text{state of world} \mid \text{sensory evidence}\right)}_{\text{posterior}} \propto \underbrace{p\left(\text{sensory evidence} \mid \text{state of world}\right)}_{\text{likelihood} \checkmark} \underbrace{p\left(\text{state of world}\right)}_{\text{prior} \checkmark}$$

posterior ↓ *loss* $\left(\text{true state}, \text{assumed state}\right)$ ✓

decision

Roadmap

Ideal observer modeling:
Uncertainty as guiding principle
in AI and computational neuroscience

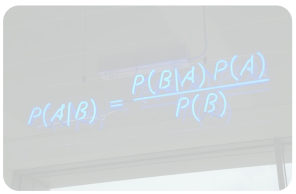

(Behavioral) evidence for handling uncertainty

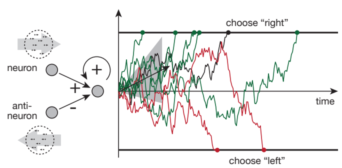

Example 1: ideal observer models for the speed/accuracy trade-off
in perceptual decision-making

Example 2: Inference on a difference time-scale:
Decision confidence to improve decision strategies

Perceptual decision-making & the speed/accuracy trade-off

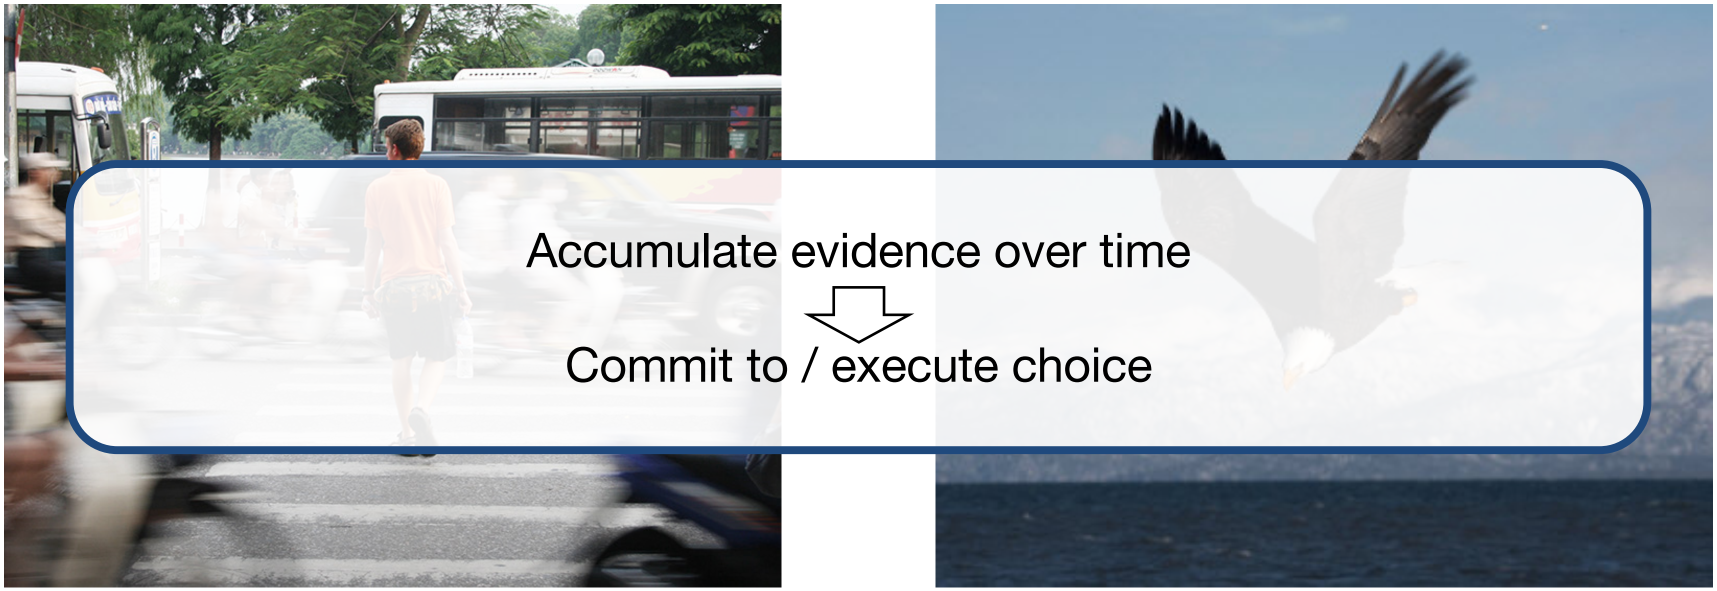

fast choices ← speed/accuracy trade-off → **slow choices**

inaccurate	accurate
low cost of accumulating evidence (e.g. attention, loss of time)	high cost

In the lab: the random-dot motion task

(e.g., *Newsome, Britten, Movshon & Shadlen, 1989; Roitman & Shadlen, 2002*)

"respond as quickly and accurately as possible"

51.2% coherence

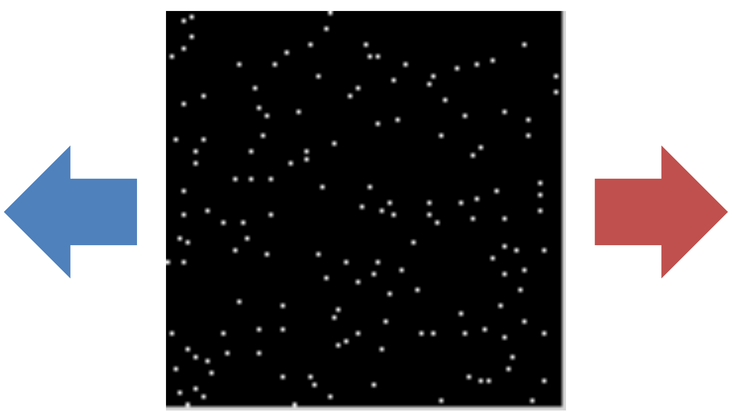

12.8% coherence

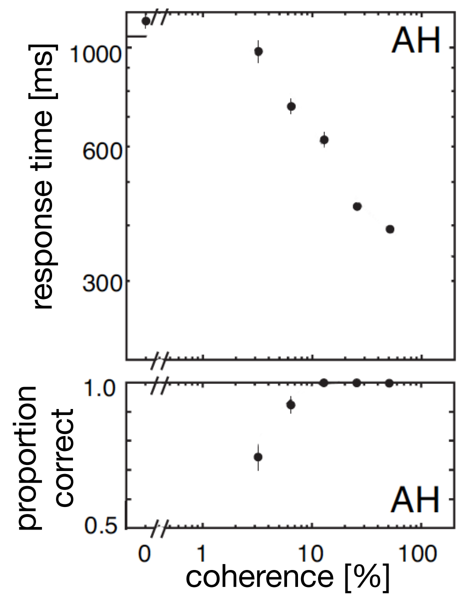

Palmer, Huk & Shadlen (2005)

Formalizing evidence accumulation

Latent state $\mu = \{-\mu_0, \mu_0\}$

Noisy evidence per δt $\delta x_n | \mu \sim \mathrm{N}(\mu \delta t, \sigma^2 \delta t)$

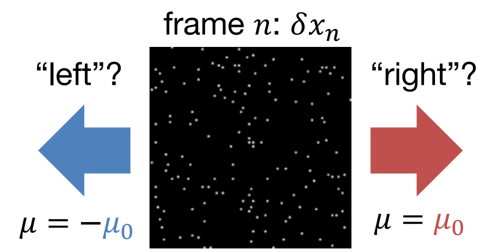

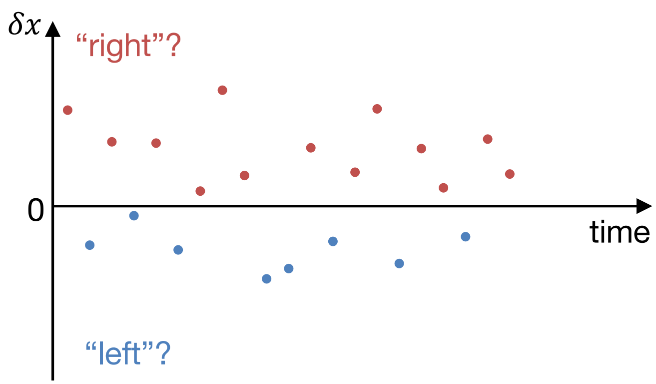

Optimal evidence accumulation: Bayes' rule

$$\underbrace{p(\mu | \delta x_{1:N})}_{\text{posterior}} \propto \underbrace{p(\delta x_{1:N} | \mu)}_{\text{likelihood}} \underbrace{p(\mu)}_{\text{prior}}$$

Posterior belief about motion being "right-ward" ($\mu > 0$)

$$g \equiv p(\mu = \mu_0 | \delta x_{1:N}) = \frac{1}{1 + e^{-\frac{2\mu_0 x(t)}{\sigma^2}}}$$

accumulated evidence ↘

$$x(t) = \sum_{n=1}^{N} \delta x_n$$

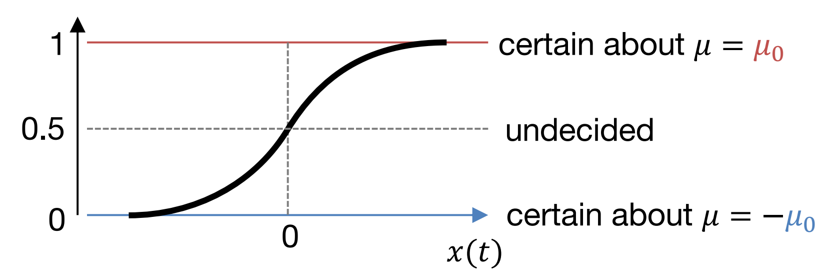

Drugowitsch et al. (2012)

Evidence accumulation by diffusion

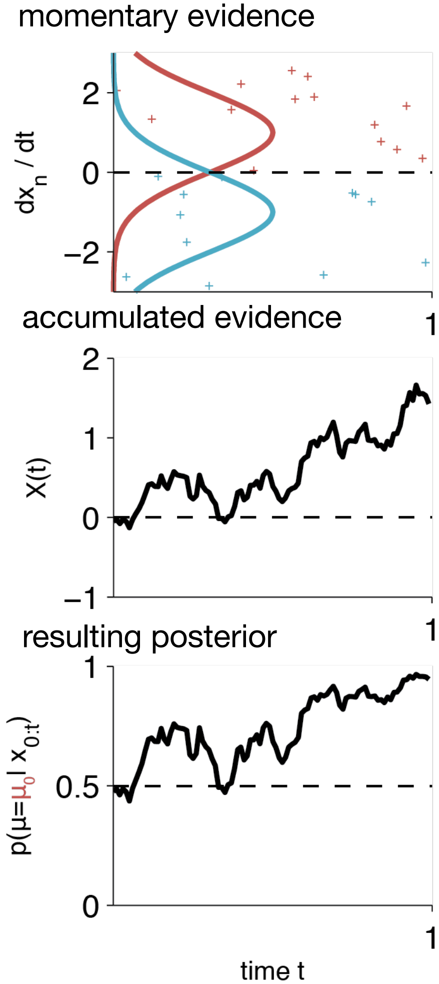

$$\delta x_n | \mu \sim \mathrm{N}(\mu \delta t, \sigma^2 \delta t)$$

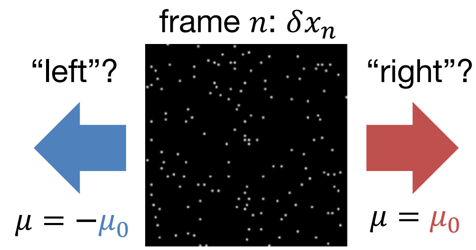

$$x(t) = \sum_{n=1}^{N} \delta x_n$$

$$\frac{\mathrm{d}x}{\mathrm{d}t} = \mu + \sigma \eta(t)$$

choose "right"

choose "left"

make choices directly
In space of
accumulated evidence

choose "right"

choose "left"

Diffusion decision models (DDMs)

(Ratcliff, 1978)

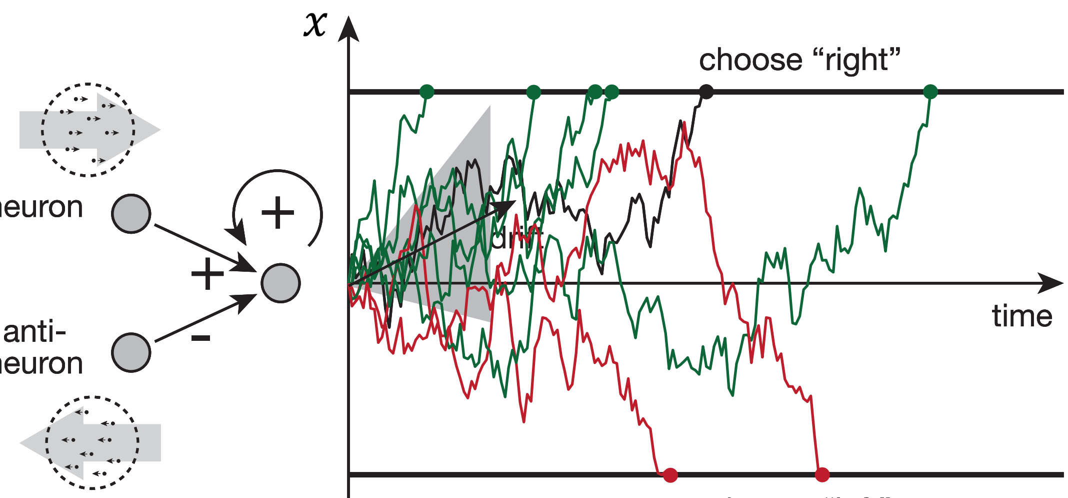

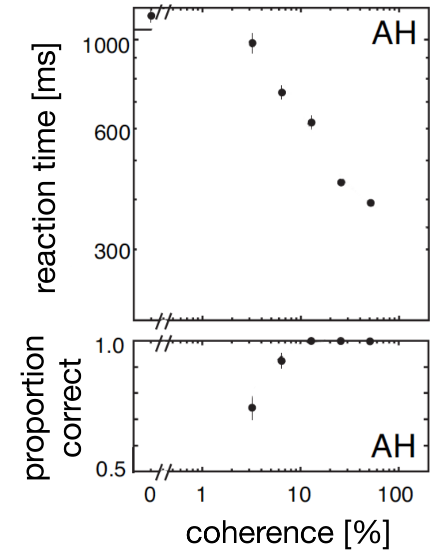

(Palmer, Huk & Shadlen, 2005)

$$\frac{\mathrm{d}x}{\mathrm{d}t} = \overbrace{kc}^{\text{drift } \mu} + \overbrace{\sigma\eta(t)}^{\text{diffusion}}$$

“coherence” white noise process

Works surprisingly well for, fast (<1.5s), single-stage decisions, e.g.,

- Word/non-word judgments (e.g., Ratcliff & Gomez, 2004)
- Numerosity judgments (e.g., Ratcliff & McKoon, 2018)
- Recognition memory (e.g., Ratcliff, 1978)
- …

Deciding when to decide: decision boundaries

Free evidence: accumulate forever!

Assume: time/evidence is costly

speed/accuracy trade-off

		fast choices ⟷	**slow choices**
cost	linear in time, ct	cheap	expensive
reward	1 for correct, 0 for incorrect	inaccurate	accurate

Optimal trade-off: dynamic programming (Bellman, 1960s)

After accumulating for some time t: expected "return" $V\big(g(t)\big)$
(recall, $g(t) \equiv p(\mu = \mu_0 | \delta x_{1:t})$)

	expected return
Choosing $-\mu_0$ or μ_0	$1-g$ or g
Accumulating another δt	$\langle V(g+\delta g)\rangle - c\delta t$

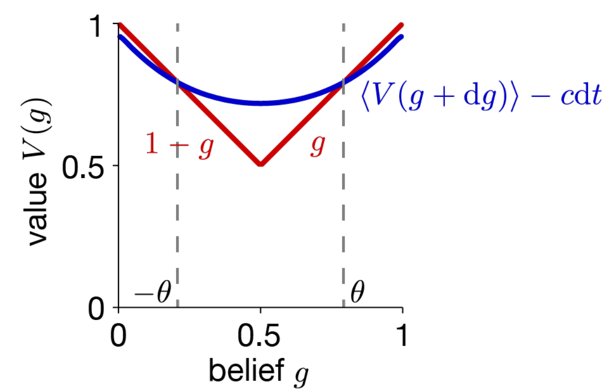

Bellman's equation

$$V(g) = \max\{1-g, g, \langle V(g+\delta g)\rangle - c\delta t\}$$

DDMs & Sequential probability ratio test

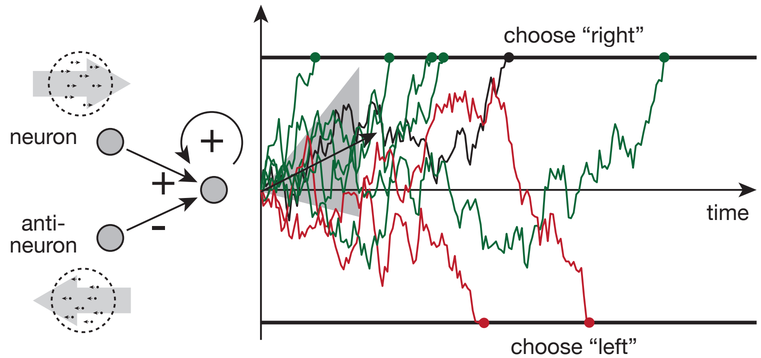

Diffusion Bayes-optimal evidence accumulation

Decision-boundaries Optimal speed-accuracy trade-off

Sequential probability ratio test (Turning, 1940s; Wald & Wolfowitz, 1948)

accumulate → decide

$$p_{ij} = \prod_{k=1}^{j} p_i(x_k).$$

If

$$\frac{p_{1n}}{p_{0n}} \geq A^*, \quad \text{the hypothesis } H_1 \text{ is accepted,}$$

if

$$\frac{p_{1n}}{p_{0n}} \leq B^* \text{ the hypothesis } H_0 \text{ is accepted.}$$

Neural correlates of evidence accumulation

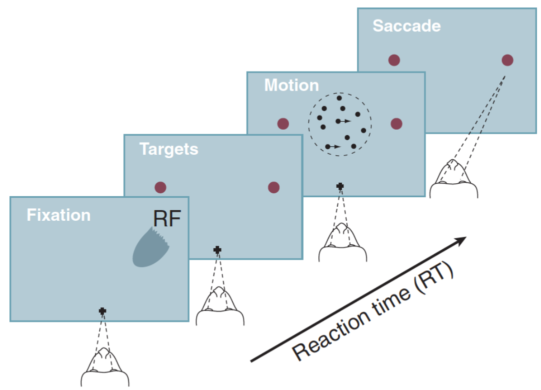

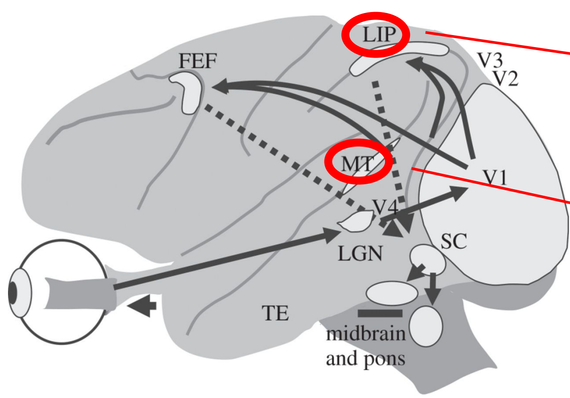

Wurtz (2015)

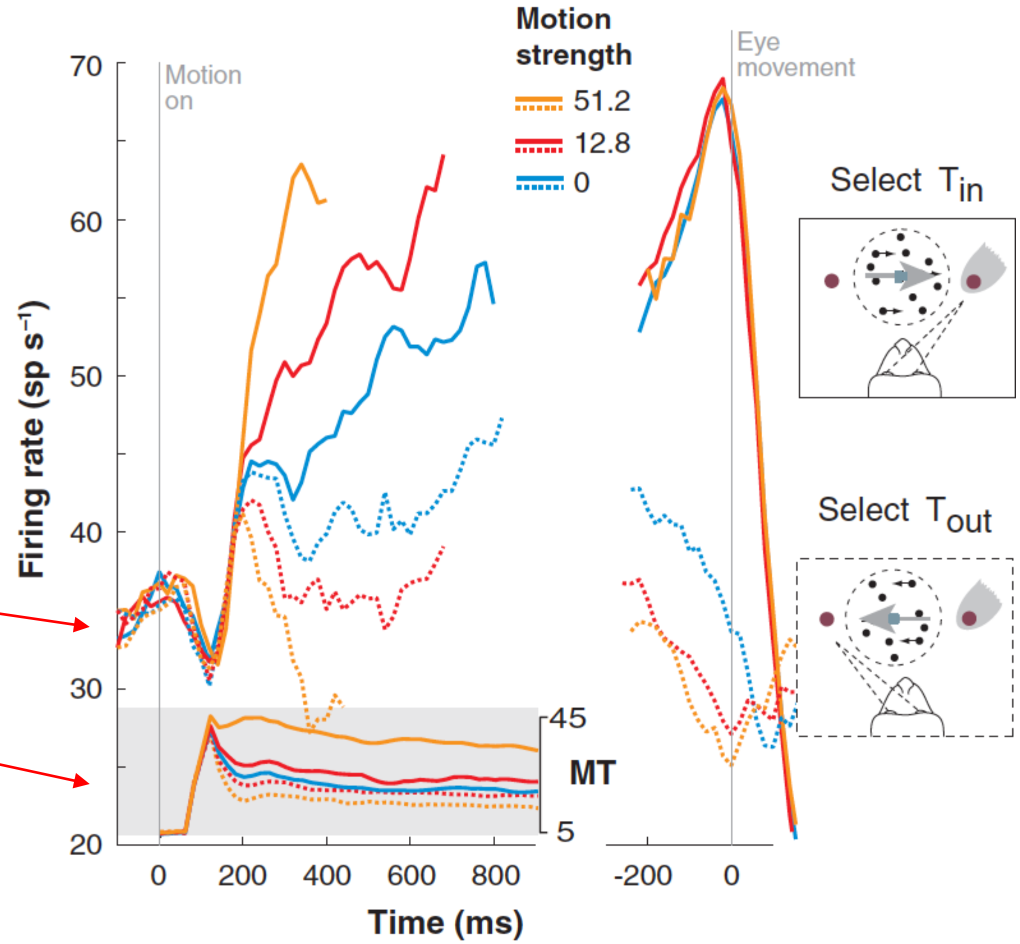

Gold & Shadlen (2007);
LIP data from *Roitman & Shadlen (2002)*;
MT data from *Britten (1992)*

Roadmap

Ideal observer modeling:
Uncertainty as guiding principle
in AI and computational neuroscience

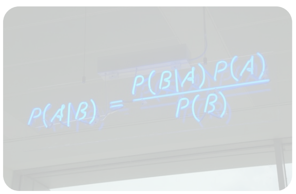

(Behavioral) evidence for handling uncertainty

Example 1: ideal observer models for the speed/accuracy trade-off
in perceptual decision-making

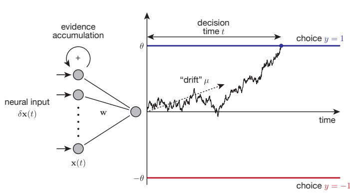

Example 2: Inference on a difference time-scale:
Decision confidence to improve decision strategies

Extending diffusion model to higher-dimensional inputs

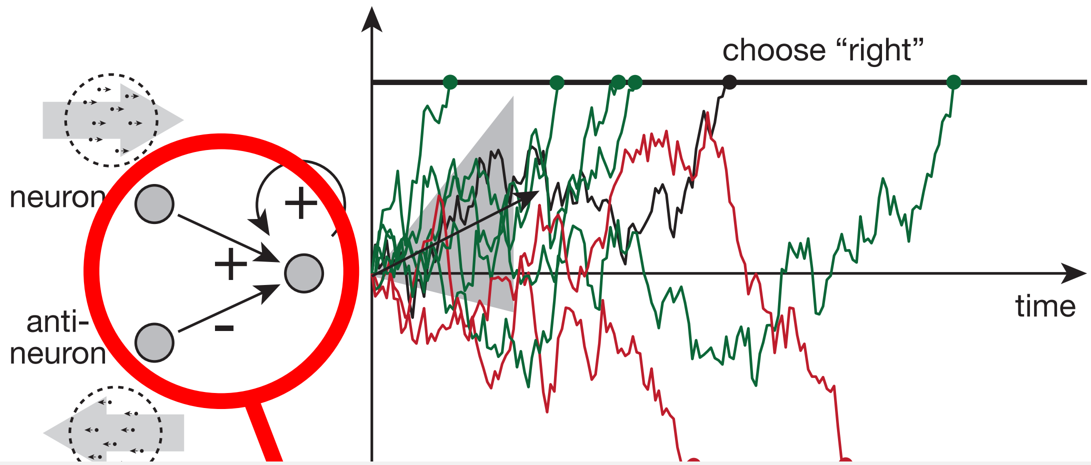

usually
only 2 inputs
known input weights

more realistically
input is larger neural population
inputs weights need to be learned

Max

fixed, known difficulty (e.g., fixed coherence)
(Turing, 1941; Wald & Wolfowitz, 1948)

difficulty that varies across trials
(Drugowitsch et al., 2012)

difficulty that varies within trials, multiple sources of information
(Drugowitsch et al., 2014; Drugowitsch, Moreno Bote & Pouget, 2014)

A larger input population

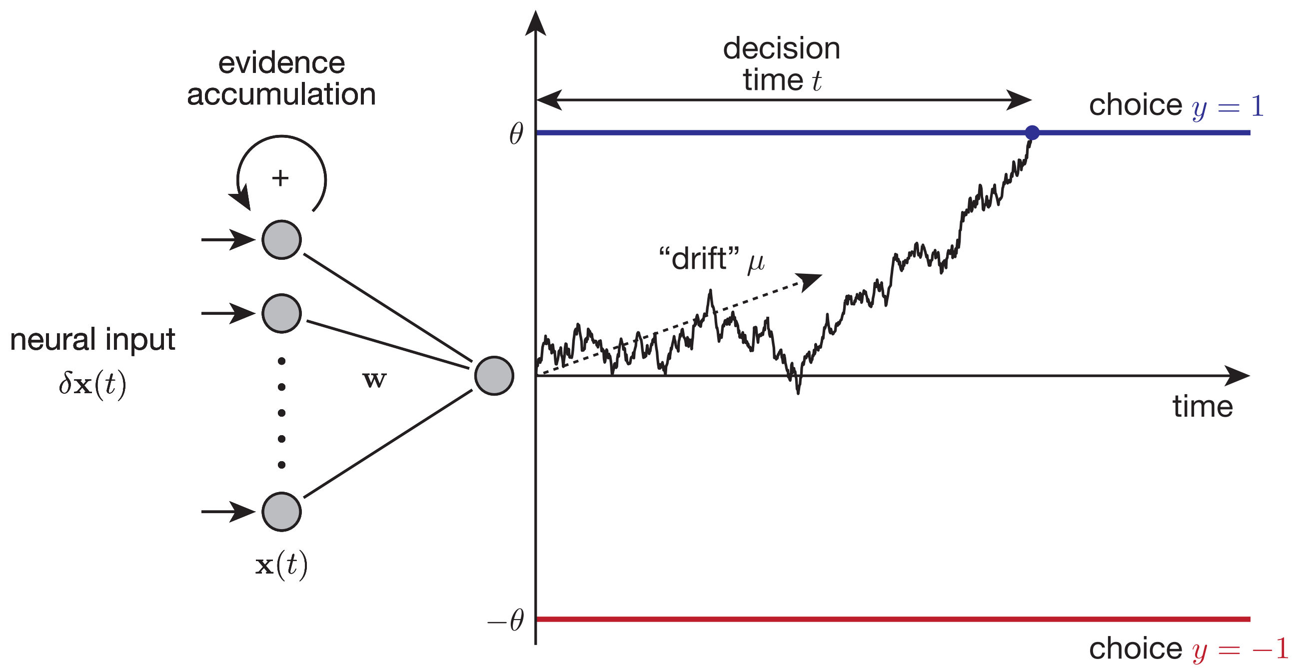

Learning input weights after feedback y^*

$$p\left(\mathbf{w}|\mathbf{x}, t, y^*\right) \propto p\left(y^*|\mathbf{w}, \mathbf{x}, t\right) p\left(\mathbf{w}\right)$$

posterior belief after feedback

likelihood of weights given feedback

belief prior to feedback

$$\mathbf{x}, t \xrightarrow{\mathbf{w}} y^*$$

Approximating Bayes rule

$$p\left(\mathbf{w}|\mathbf{x},t,y^*\right) \propto p(y^*|\mathbf{w},\mathbf{x},t)\ \ p\left(\mathbf{w}\right)$$

INTRACTABLE

from generative model:
cumulative Gaussian likelihood

Assumed density filtering

Minimizing $KL\left(p||q\right)$ between true posterior p and Gaussian approximation q

$$q\left(\mathbf{w}|\mathbf{x},t,y^*\right) \underset{\sim}{\propto}\ \ p(y^*|\mathbf{w},\mathbf{x},t)\ \ q(\mathbf{w})$$

assumed Gaussian

becomes Gaussian

Learning of input weights: tracking mean and covariance of Gaussian

The Bayes-(near)-optimal learning rule

With prior belief before feedback, $\mathbf{w} \sim \mathcal{N}\left(\mu_w, \beta_w \bar{\boldsymbol{\Sigma}}_w\right)$

monotonic function of decision confidence

accumulated evidence

$$\Delta\mu_w = \frac{y^* C_w}{\sigma\sqrt{t + \frac{\sigma^2}{\sigma_\mu^2} + \beta_w \frac{\mathbf{x}^T \bar{\boldsymbol{\Sigma}}_w \mathbf{x}}{\sigma^2}}} \beta_w \bar{\boldsymbol{\Sigma}}_w \mathbf{x}$$

decision time

learning rate α_{ADF}

weight uncertainty

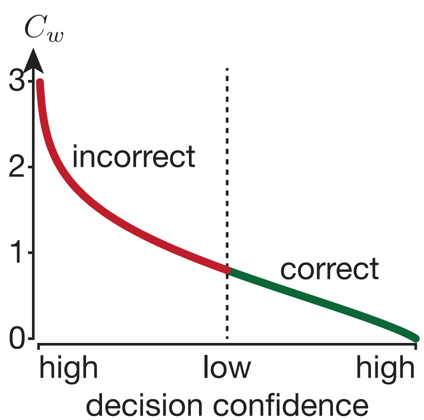

What determines decision confidence?

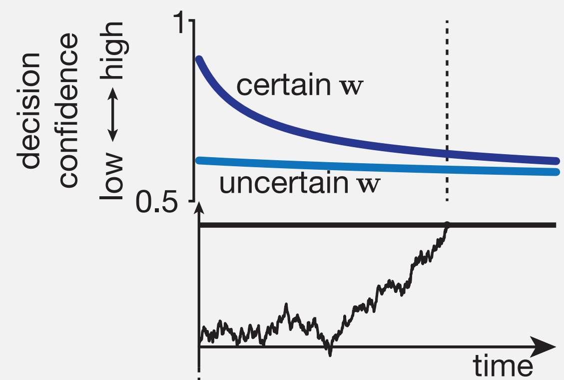

less confident due to

- slower choices (usually harder)
- less certain about weights

The Bayes-(near)-optimal learning rule

With prior belief before feedback, $\mathbf{w} \sim \mathcal{N}\left(\mu_w, \beta_w \bar{\boldsymbol{\Sigma}}_w\right)$

$$\Delta\mu_w = \frac{y^* C_w}{\sigma\sqrt{t + \frac{\sigma^2}{\sigma_\mu^2} + \beta_w \frac{\mathbf{x}^T \bar{\boldsymbol{\Sigma}}_w \mathbf{x}}{\sigma^2}}} \beta_w \bar{\boldsymbol{\Sigma}}_w \mathbf{x}$$

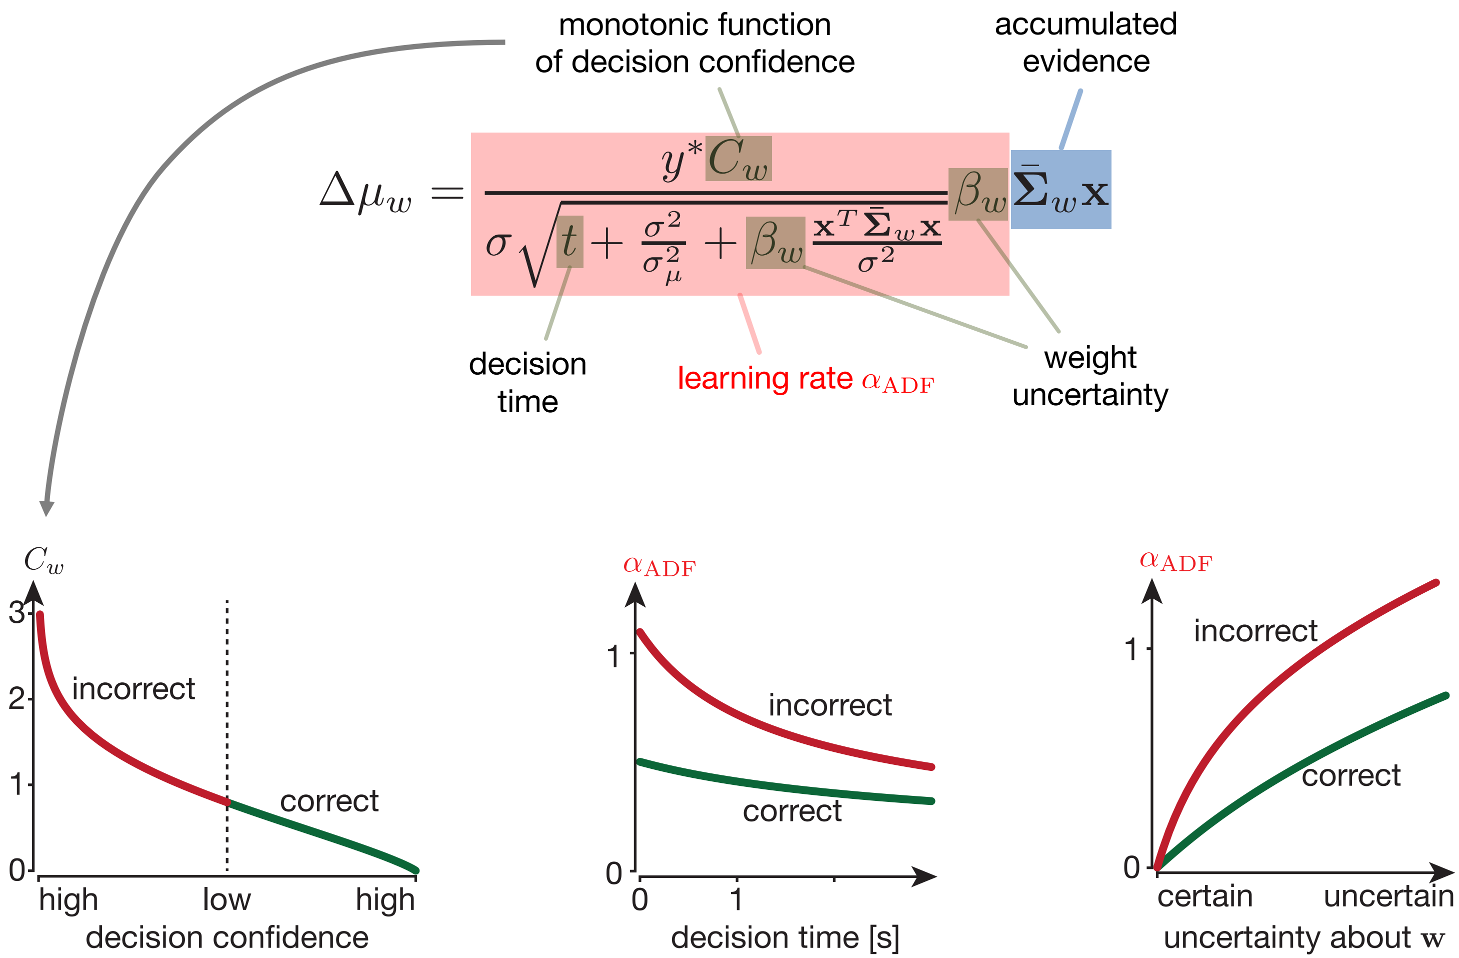

How good is the approximation?

Compare performance to optimal (Gibbs sampling) solution

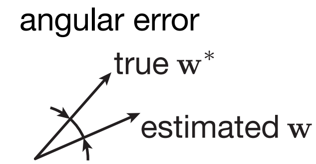

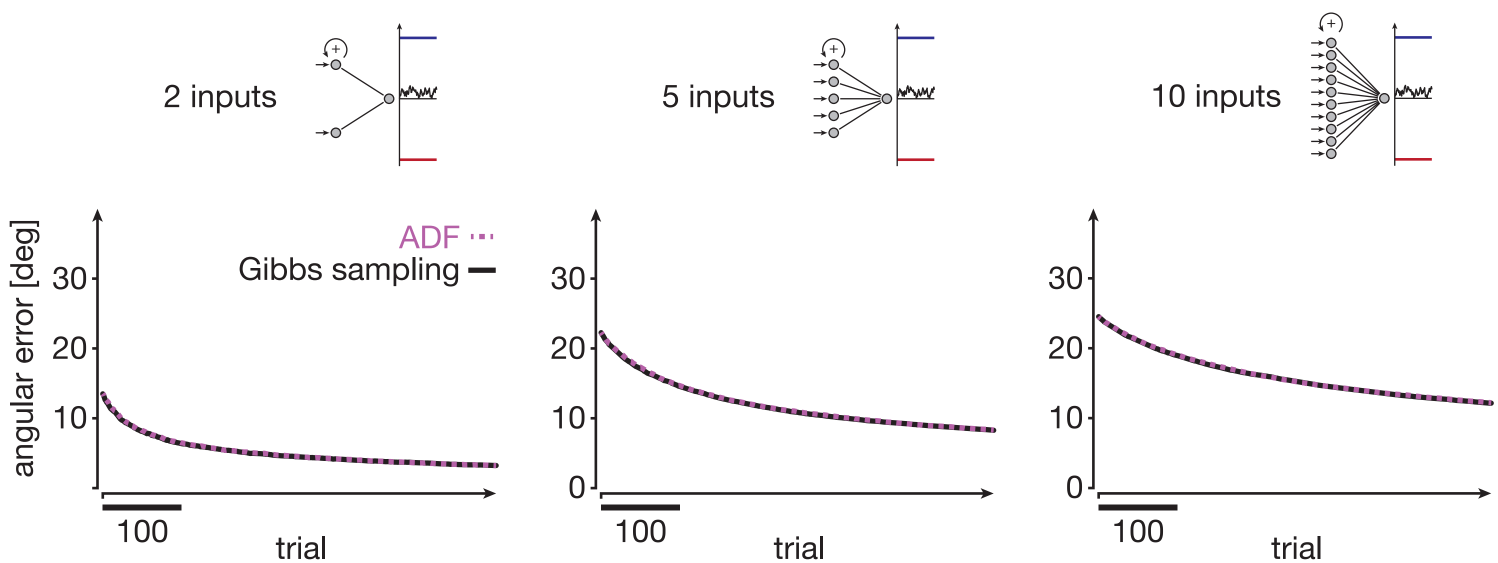

Do we need to be probabilistic? Simpler heuristics

Simple delta rule

minimizes distance between chosen and correct bound

$$\Delta \mathbf{w} = \alpha \underbrace{\left(y^* \theta - \mathbf{w}^T \mathbf{x}(t)\right)}_{\text{either } 2\theta \text{ or } 0} \mathbf{x}(t)$$

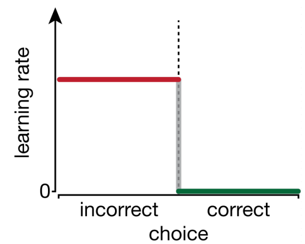

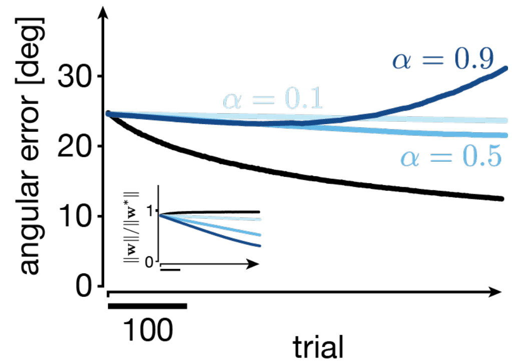

Normalized delta rule

delta rule with normalization, $\mathbf{w} \leftarrow \frac{\mathbf{w}}{\|\mathbf{w}\|} \langle \|\mathbf{w}^*\| \rangle$

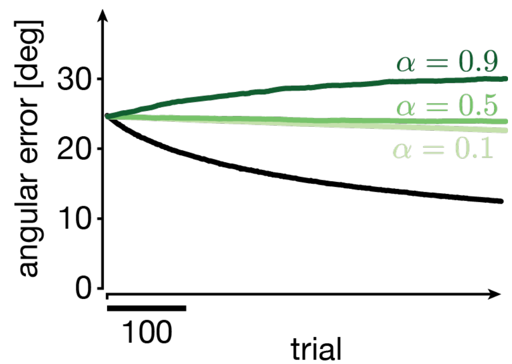

Stochastic gradient ascent

on log-likelihood

$$\Delta \mathbf{w} = \alpha \nabla_{\mathbf{w}} \ln p\left(y^* | \mathbf{w}, \mathbf{x}(t), t\right)$$

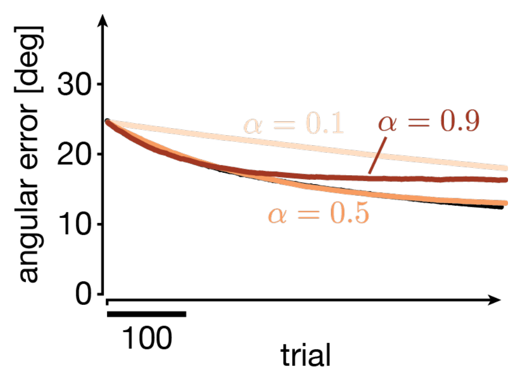

Steady-state performance

So far	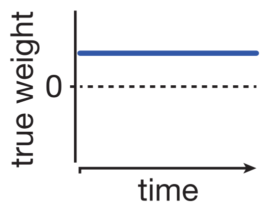	stationary weights	aim to learn weights
Now	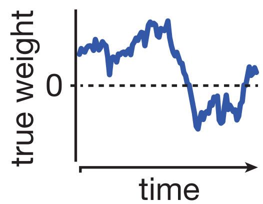	weights follow AR(1) process	aim to track changes

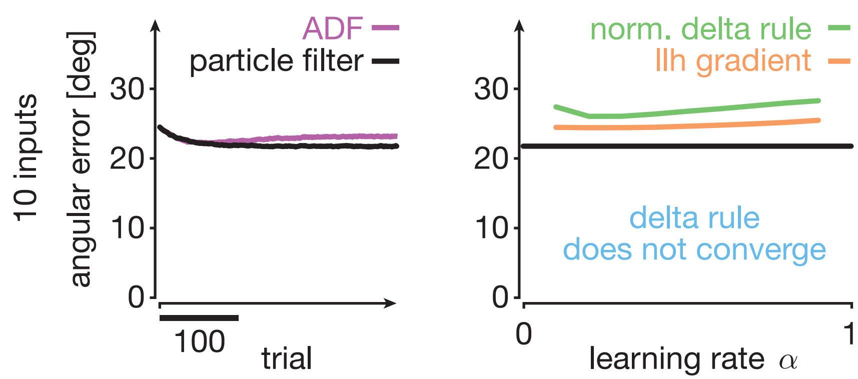

Continual learning predicts sequential choice dependencies

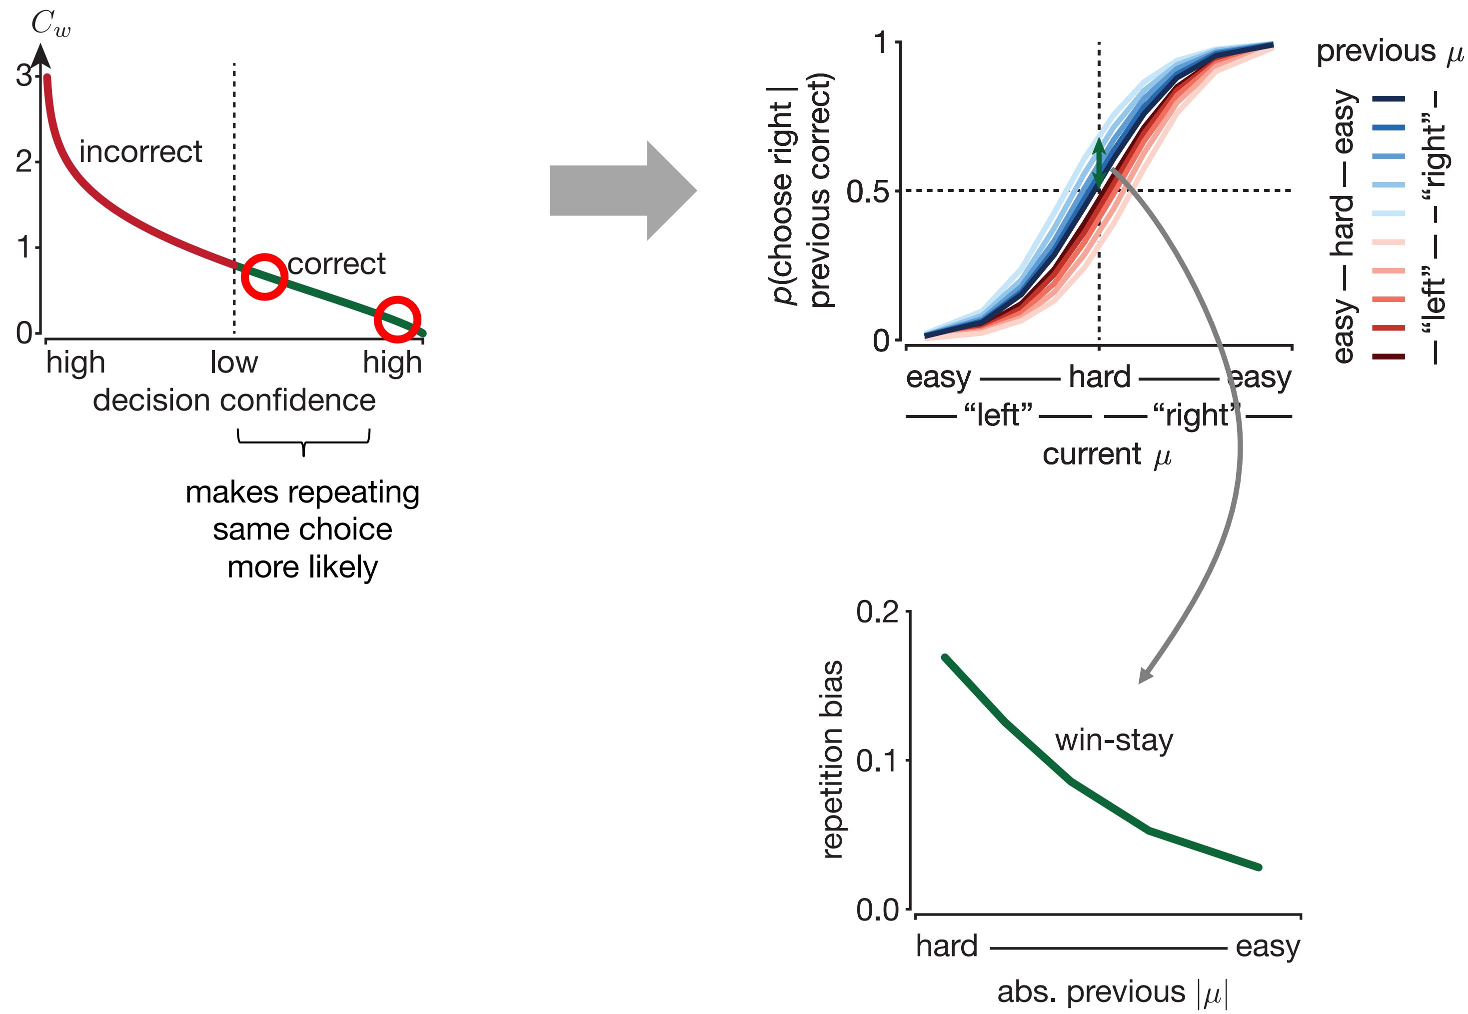

Odor categorization/identification task

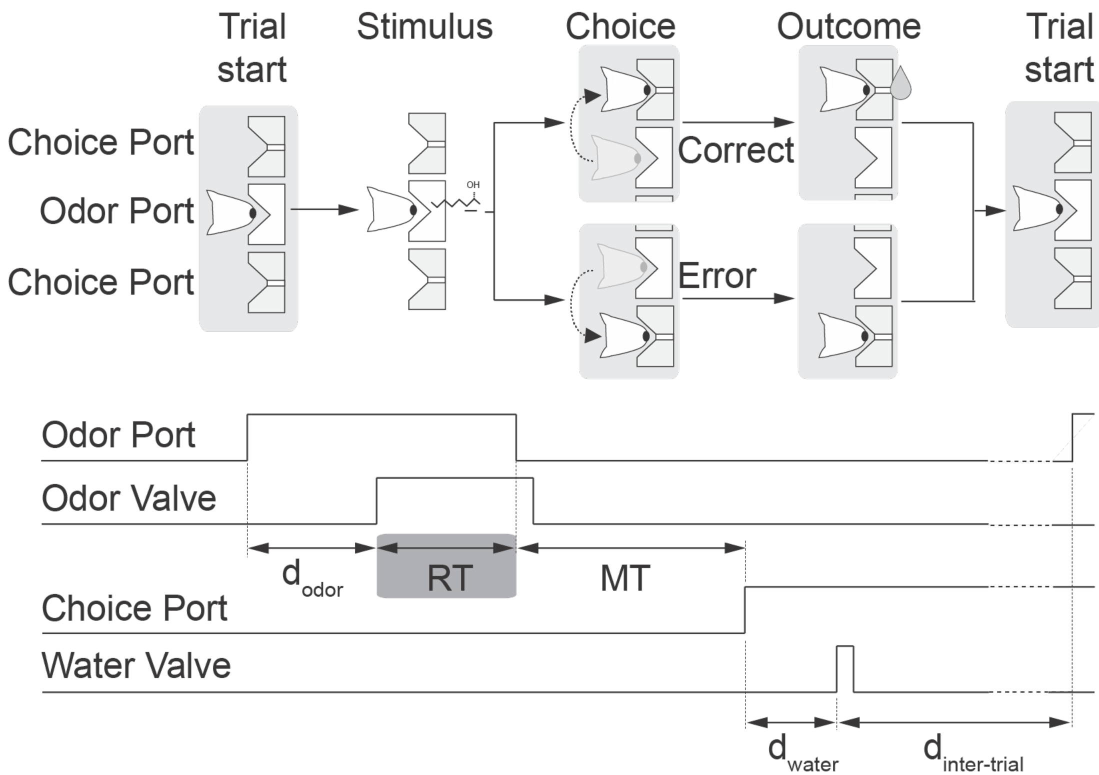

Task conditions

Odor identification condition

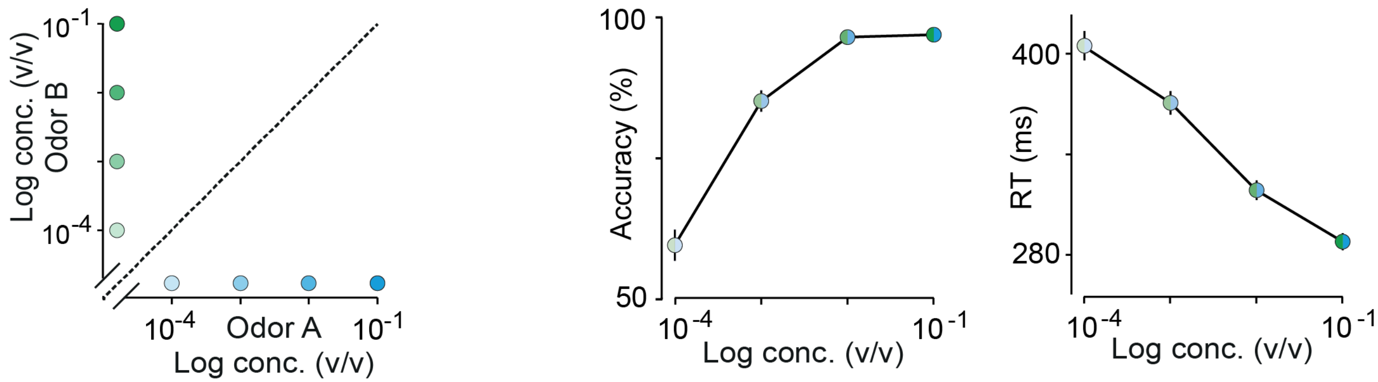

Odor categorization condition

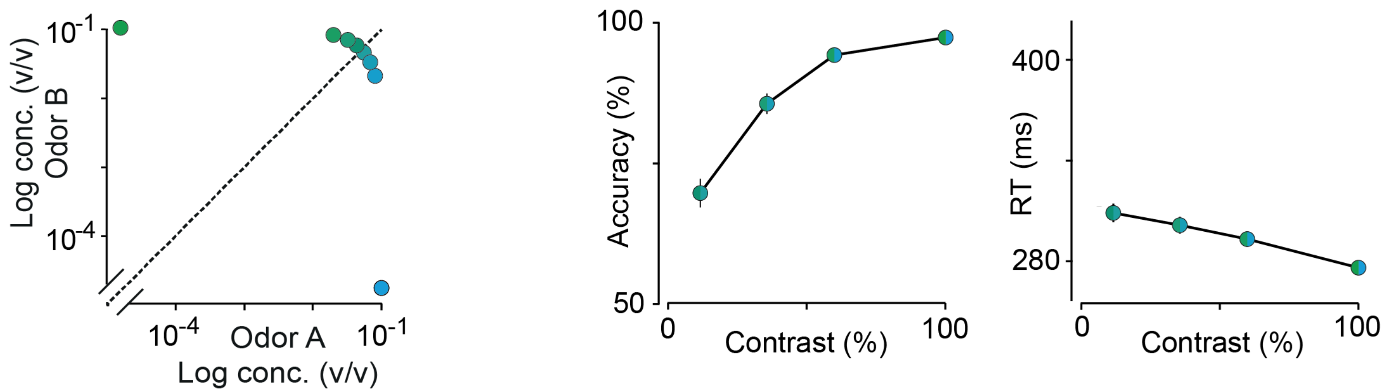

Vanilla diffusion models can't fit both conditions

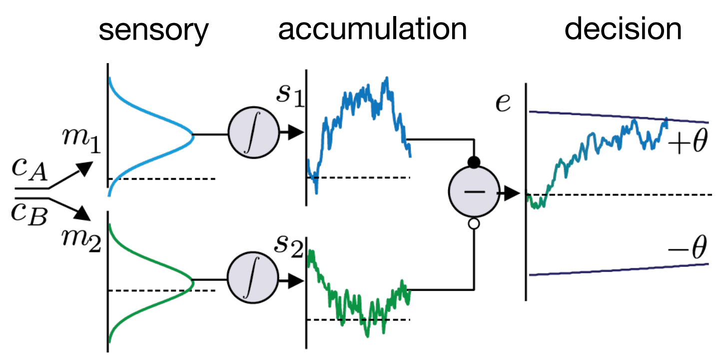

Odor identification condition

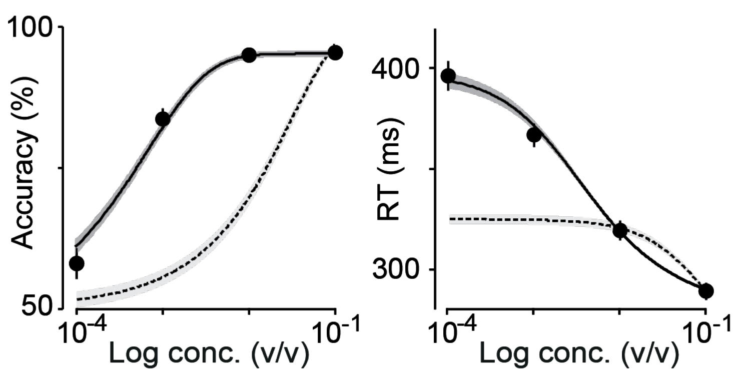

Odor categorization condition

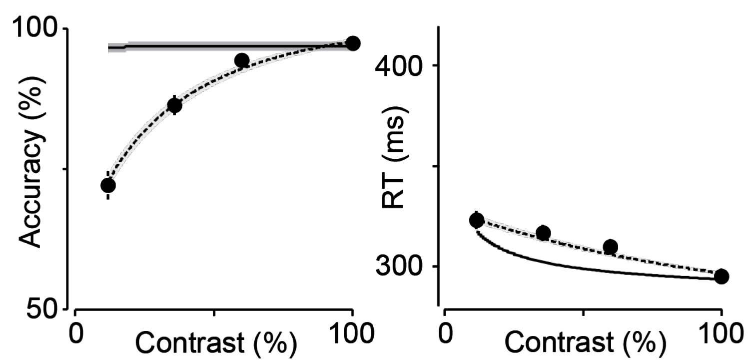

…but a learning model can

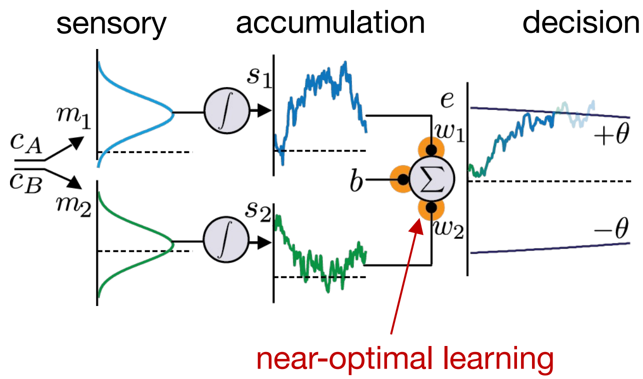

Odor identification condition

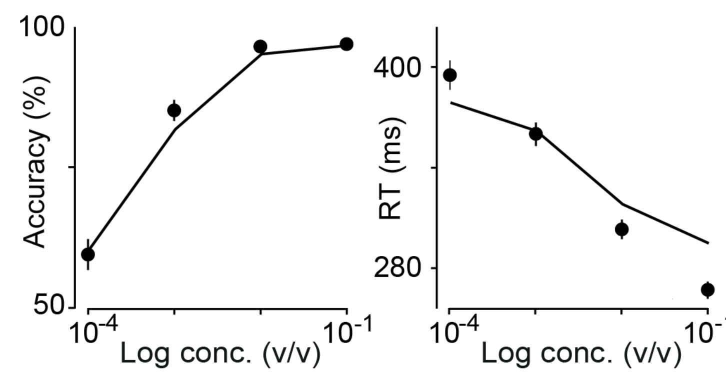

Interleaved condition

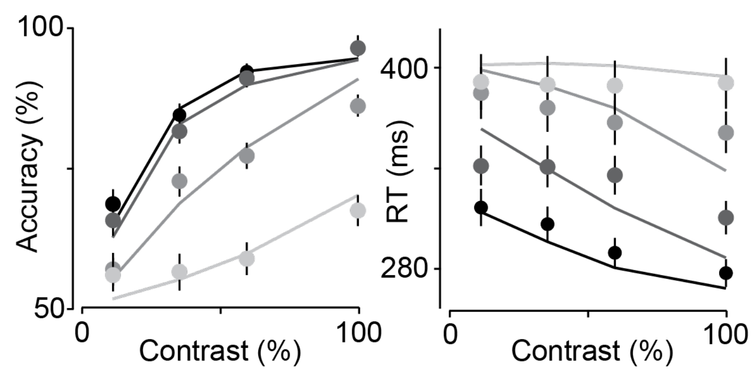

Odor categorization condition

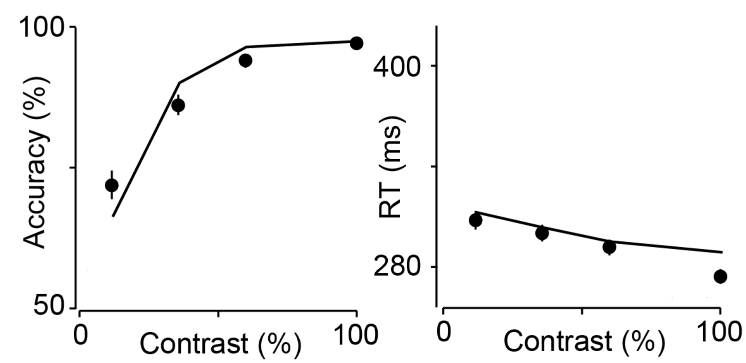

Sequential effects

Sequential effects are not fitted, but predicted

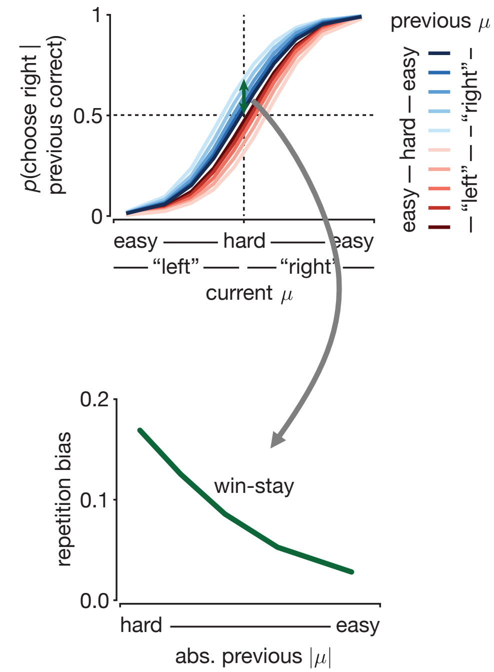

Odor identification condition

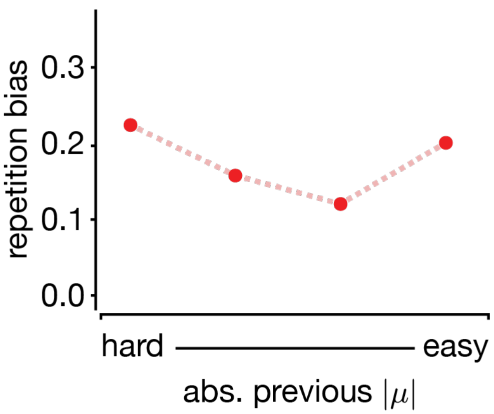

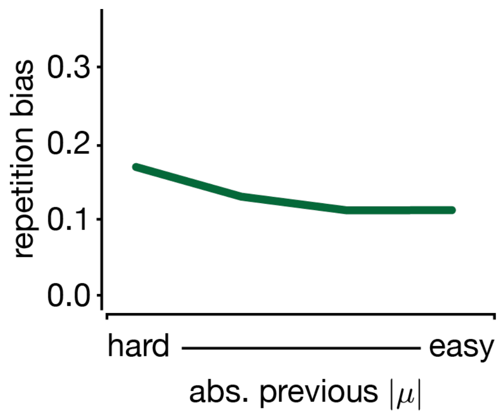

Odor categorization condition

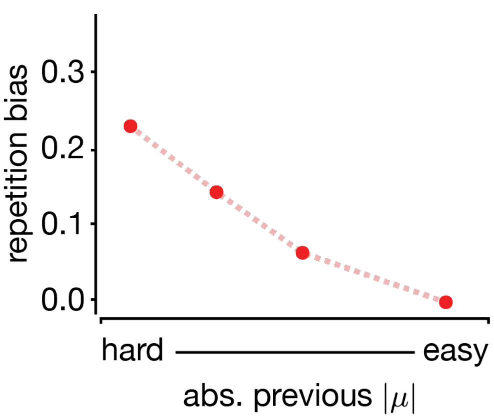

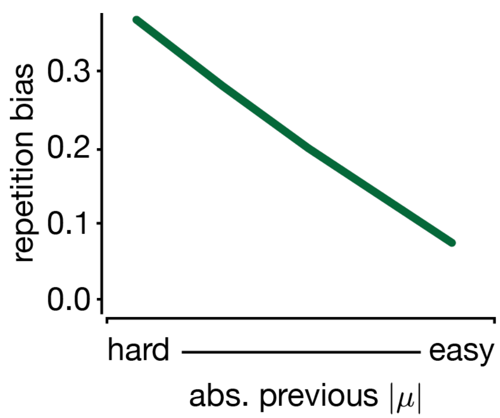

Simpler models fit data less well

Model comparison on psychometric/chronometric curves & sequential effects

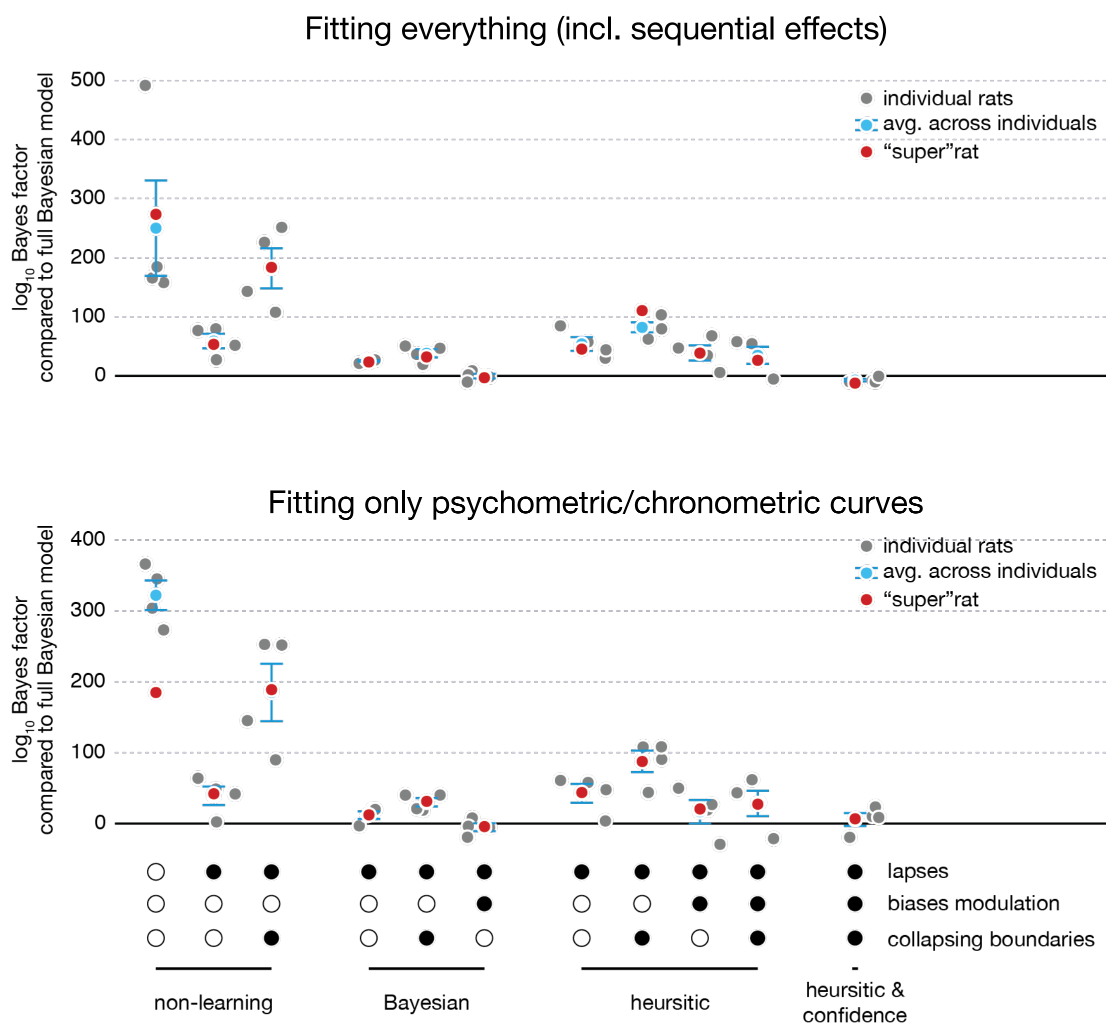

Using BIC; qualitatively same results for AIC & AICc

Summary

Ideal observer modeling:
Uncertainty as guiding principle
in AI and computational neuroscience

Handling uncertainty by Bayesian decision theory

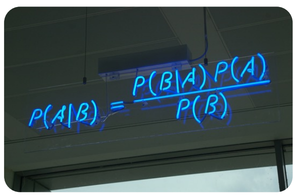

(Behavioral) evidence for handling uncertainty

Reliability-weighted cue combination
Use of prior information for uncertain evidence
Loss-sensitive decision-making

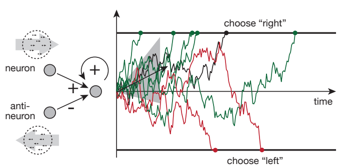

Example 1: ideal observer models for the speed/accuracy trade-off
in perceptual decision-making

Optimal speed-accuracy trade-off by diffusion models

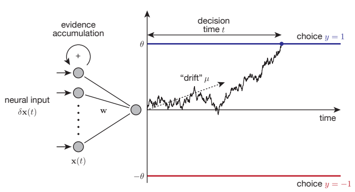

Example 2: Inference on a difference time-scale:
Decision confidence to improve decision strategies

Bayes-optimal learning is confidence-weighted
Provides computational role for decision confidence
Predicts sequential choice dependencies for continual learning

Drugowitsch et al. (2012). The Journal of Neuroscience
Drugowitsch et al. (2019). PNAS
Mendonça et al. (2020). Nature Communications

Ideal observer models as hypothesis generators

Brain is too high-dimensional to
fully explore by experiments

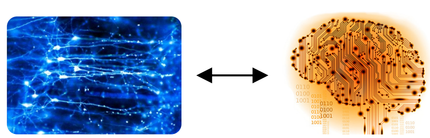

Use of ideal observer models to
generate hypotheses of brain function

Further questions?

Do rodents learn both weights *and* biases?

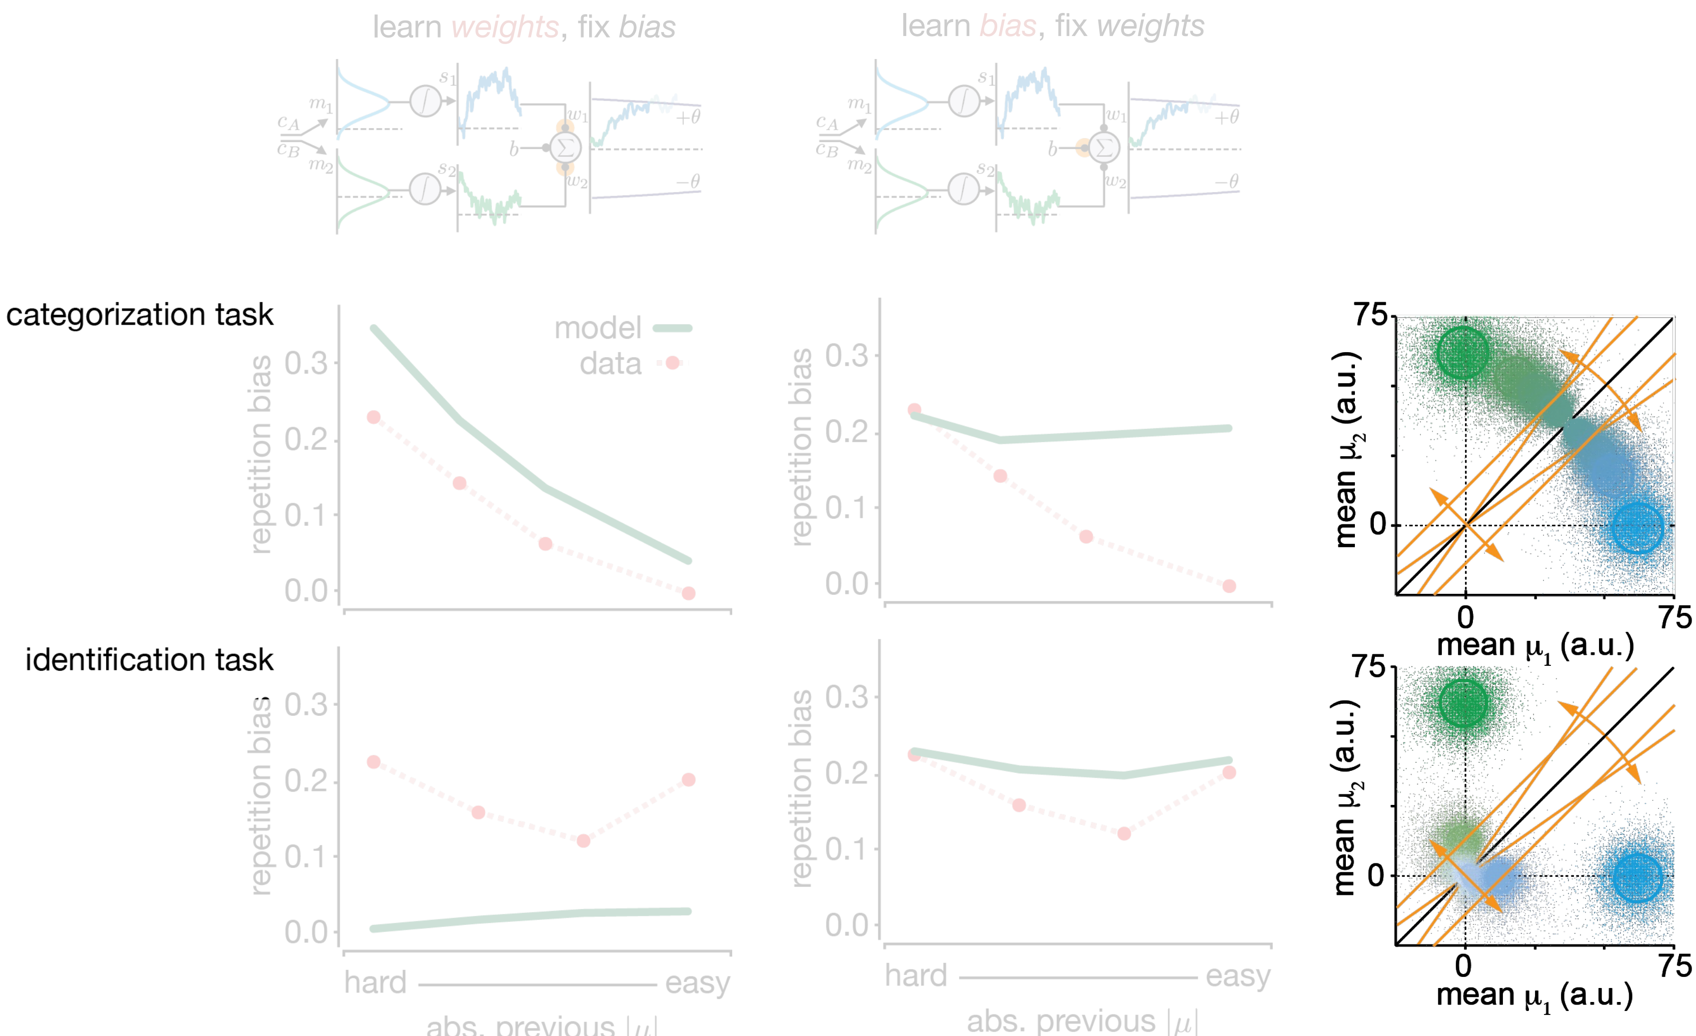

Simple parameter adjustment across blocked conditions?

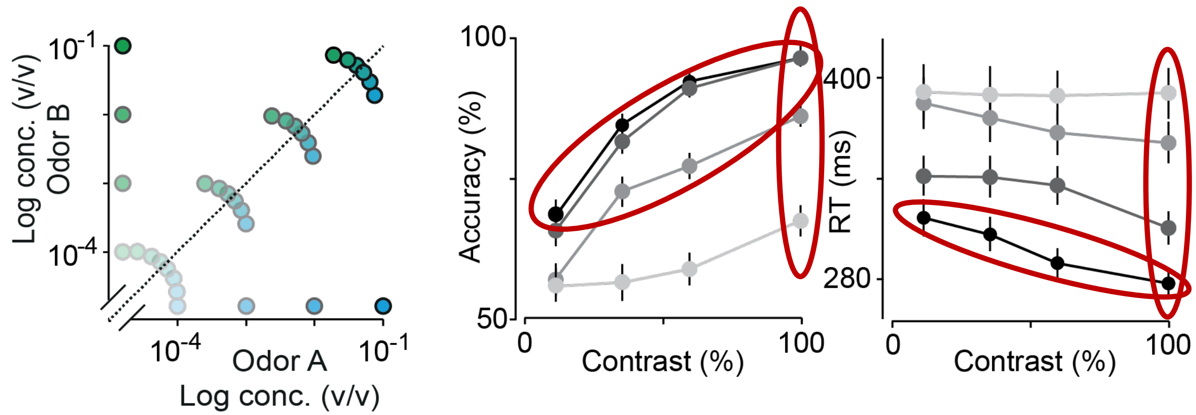

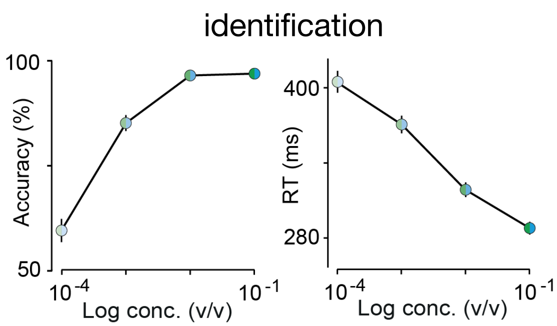

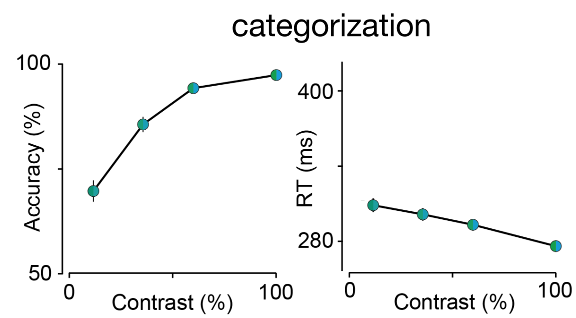

Confidence-weighting only for diffusion models?

For diffusion models

$$p\left(\mathbf{w}|\mathbf{x},t,y^*\right) \propto \; p\left(y^*|\mathbf{w},\mathbf{x},t\right) \; p\left(\mathbf{w}\right)$$

More generally, following choice y after observing $\mathbf{x}$

after correct choices, $y = y^*$

$$p\left(\mathbf{w}|\mathbf{x},\text{choice}\right) \propto \quad p(y = \text{choice}|\mathbf{w},\mathbf{x},\text{choice}) \quad p\left(\mathbf{w}\right)$$

after incorrect choices, $y \neq y^*$

$$p\left(\mathbf{w}|\mathbf{x},\text{choice}\right) \propto \; (1 - p(y = \text{choice}|\mathbf{w},\mathbf{x},\text{choice})) \; p\left(\mathbf{w}\right)$$

decision confidence

What about *N*-AFC with *N*>2?

Feedback is *correct/incorrect*: again decision confidence

Feedback is *correct choice*: need full posterior $p\left(y|\mathbf{w},\mathbf{x}\right)$ for all y

Confidence trumps imperfect feedback

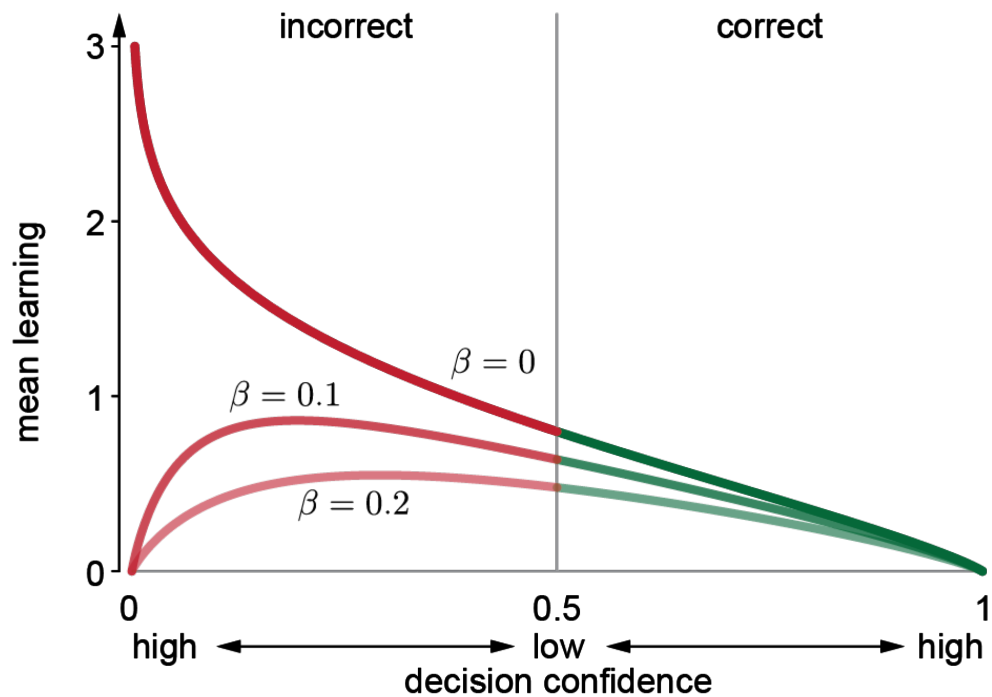

β = probability of inverted feedback

Further questions?