Alexandre Pouget

Department of Basic Neuroscience
University of Geneva

Department of Brain and Cognitive Sciences
University of Rochester

What limits performance in decision making?

Conclusions

What limits performance in a fully attentive, well-trained animal/human?

Conclusions

What limits performance in a fully attentive, well-trained animal/human?

Is behavioral performance limited by:

Is behavioral performance limited by:

Noise in the brain

Is behavioral performance limited by:

Noise in the brain

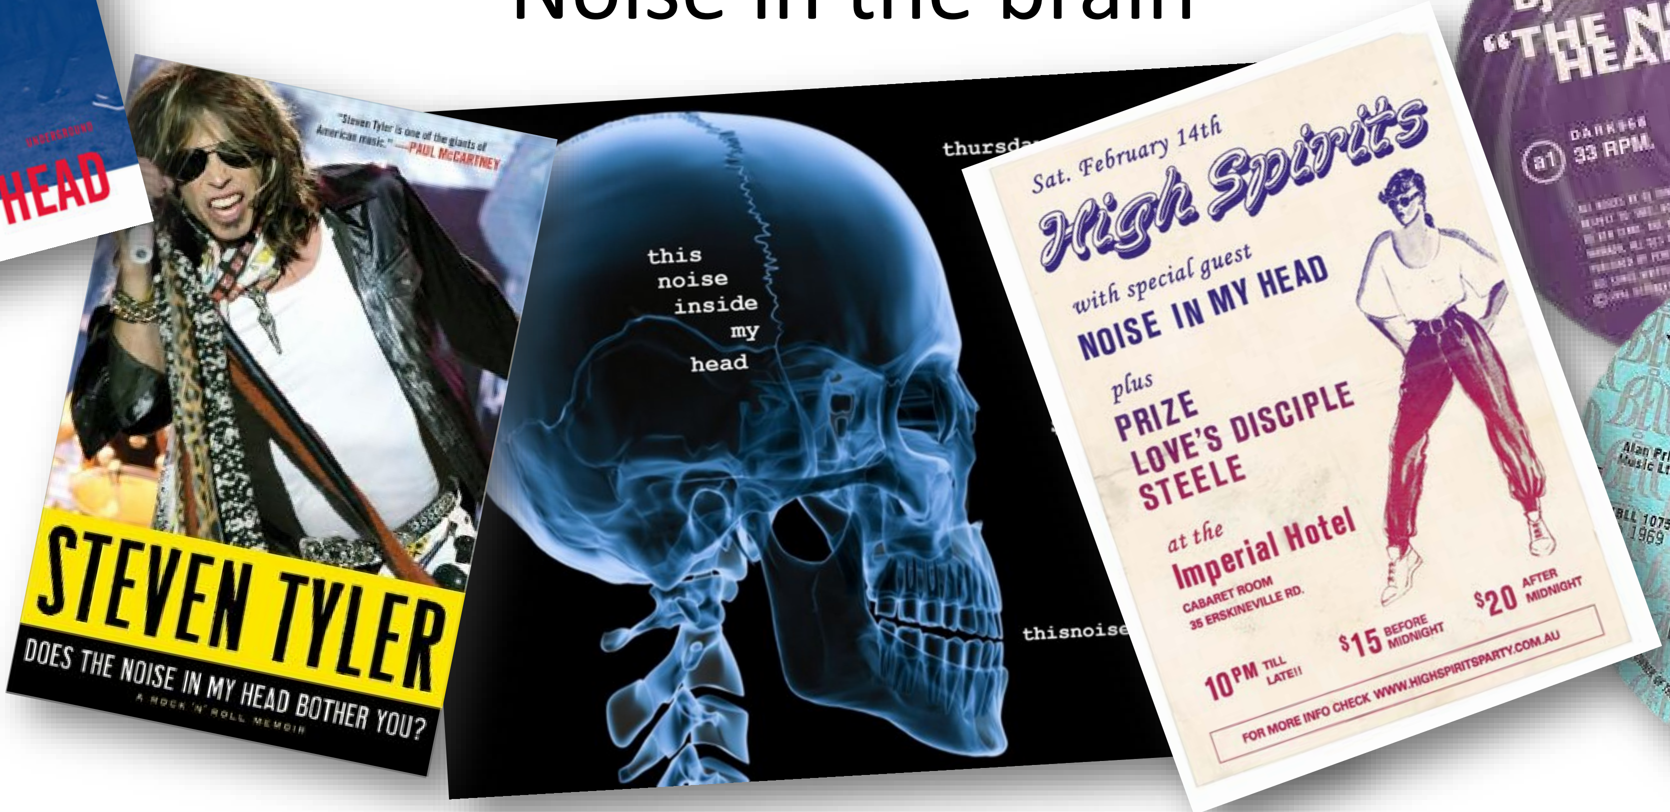

Is behavioral performance limited by:

Noise in the brain

Is behavioral performance limited by:

Noise in the brain

Is behavioral performance limited by:

Noise in the brain

Is behavioral performance limited by:

Noise in the brain

or

A combination of suboptimal inference and variability in the sensory inputs and sensors

Is behavioral performance limited by:

Noise in the brain

or

A combination of suboptimal inference and variability in the sensory inputs and sensors

Not Noisy, Just Wrong: The Role of Suboptimal Inference in Behavioral Variability

Jeffrey M. Beck,[1,5] Wei Ji Ma,[2,5] Xaq Pitkow,[1] Peter E. Latham,[3] and Alexandre Pouget[1,3,4,*] 2012

Roadmap

Roadmap

Experiment 1: what appears as noise can be suboptimal inference

Roadmap

Experiment 1: what appears as noise can be suboptimal inference

Experiment 2: noise only has a marginal impact on performance

Olfactory processing

- Olfactory detection and categorization

In collaboration with Mendonca, Vicente, DeWitt and Mainen

Olfactory processing

- Olfactory detection and categorization

Dual tasks

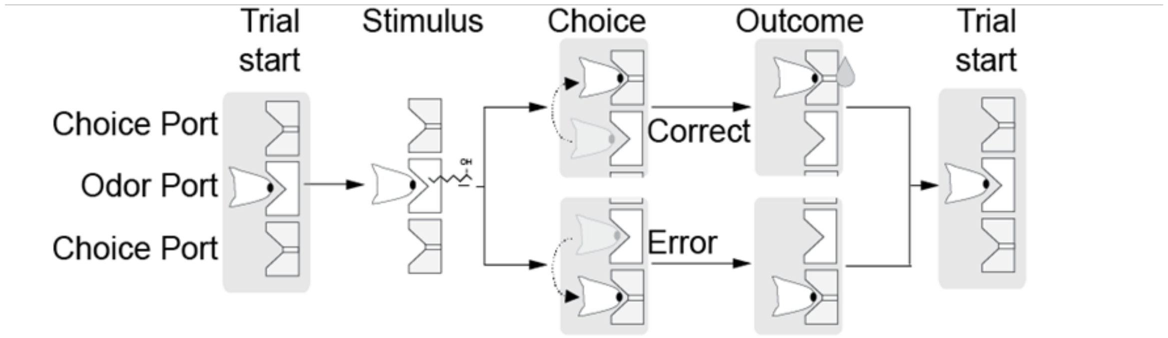

Trial start
Stimulus
Choice
Outcome
Trial start
Choice Port
Odor Port
Choice Port
Correct
Error

Dual tasks

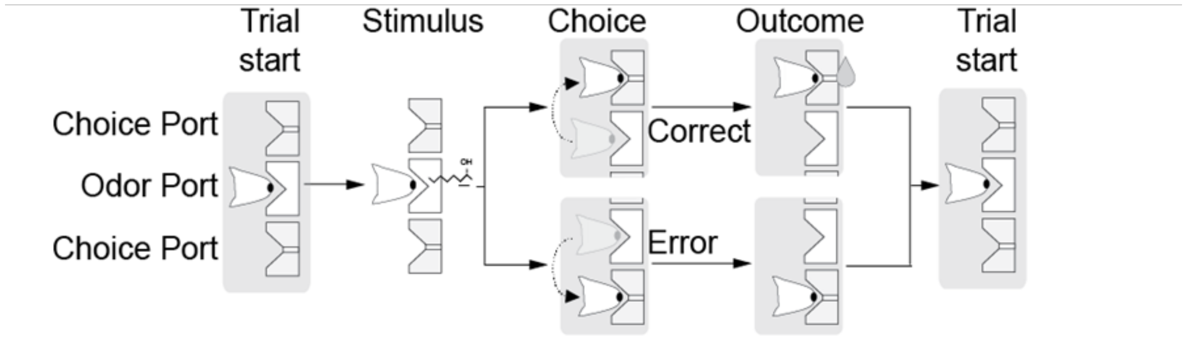

Trial start
Stimulus
Choice
Outcome
Trial start
Choice Port
Odor Port
Choice Port
Correct
Error

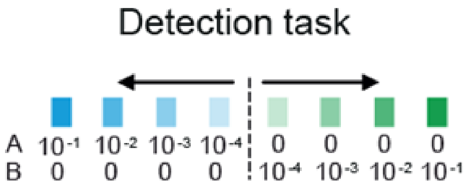

Detection task
A 10-1 10-2 10-3 10-4 0 0 0 0
B 0 0 0 0 10-4 10-3 10-2 10-1

Dual tasks

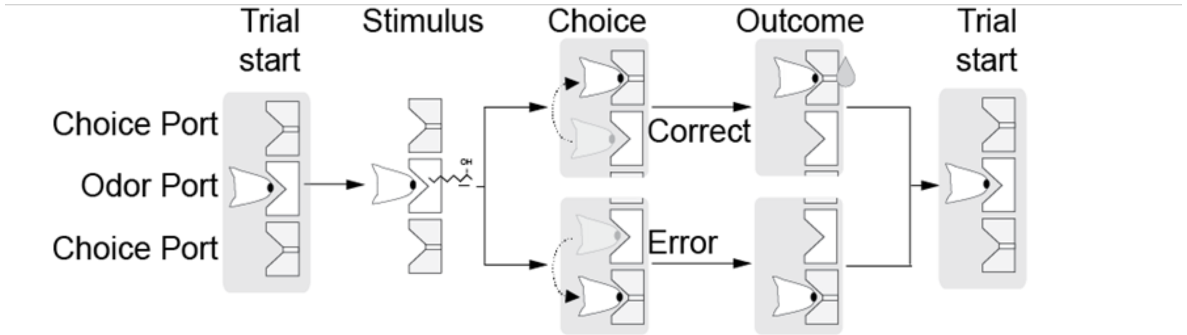
Trial start
Stimulus
Choice
Outcome
Trial start
Choice Port
Odor Port
Choice Port
Correct
Error

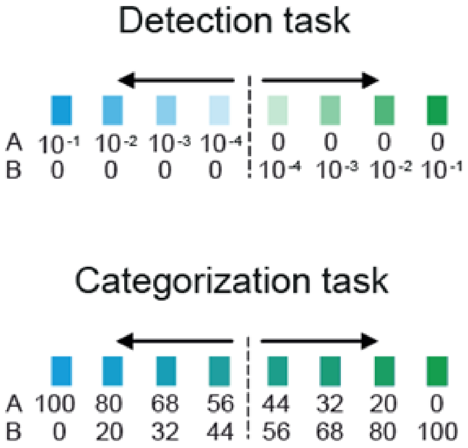
Detection task
A 10-1 10-2 10-3 10-4 0 0 0 0
B 0 0 0 0 10-4 10-3 10-2 10-1
Categorization task
A 100 80 68 56 44 32 20 0
B 0 20 32 44 56 68 80 100

Dual tasks

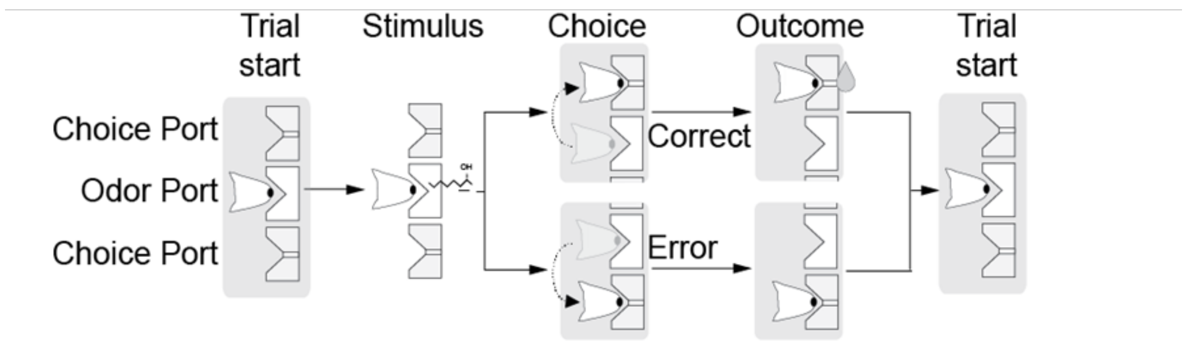
Trial start
Stimulus
Choice
Outcome
Trial start
Choice Port
Odor Port
Choice Port
Correct
Error

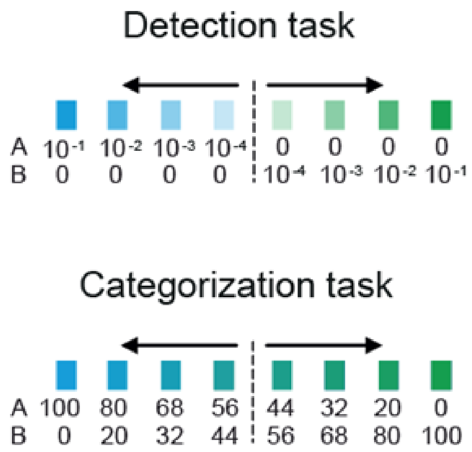
Detection task
A 10-1 10-2 10-3 10-4 0 0 0 0
B 0 0 0 0 10-4 10-3 10-2 10-1
Categorization task
A 100 80 68 56 44 32 20 0
B 0 20 32 44 56 68 80 100

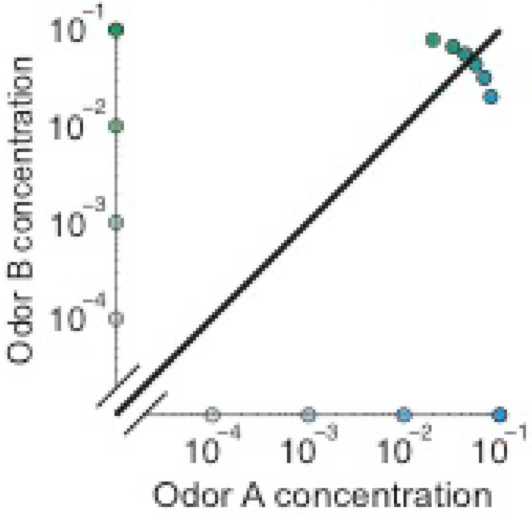
Odor B concentration
Odor A concentration
10-1
10-2
10-3
10-4

Dual tasks

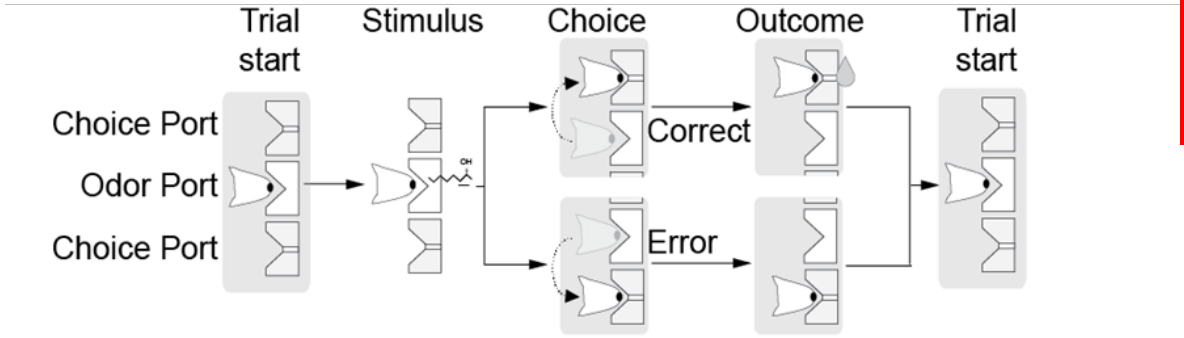

Trial start
Stimulus
Choice
Outcome
Trial start
Choice Port
Odor Port
Choice Port
Correct
Error

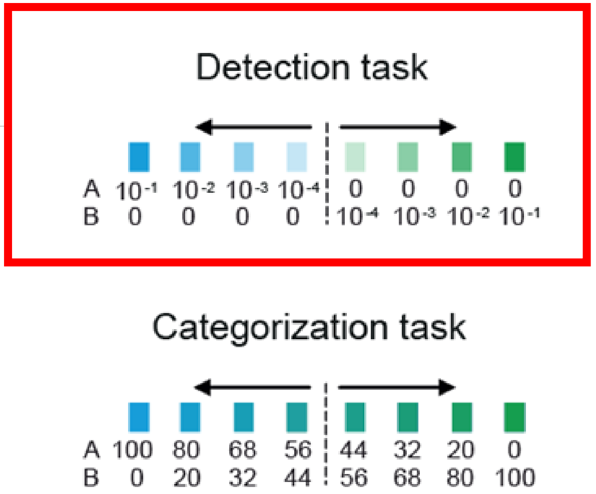

Detection task
A 10-1 10-2 10-3 10-4 0 0 0 0
B 0 0 0 0 10-4 10-3 10-2 10-1
Categorization task
A 100 80 68 56 44 32 20 0
B 0 20 32 44 56 68 80 100

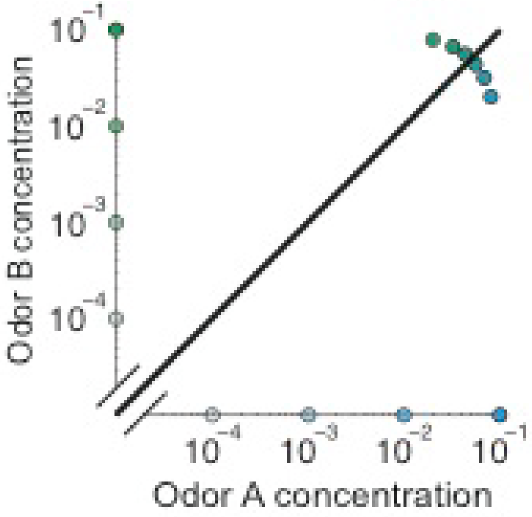

Odor B concentration
Odor A concentration

Dual tasks

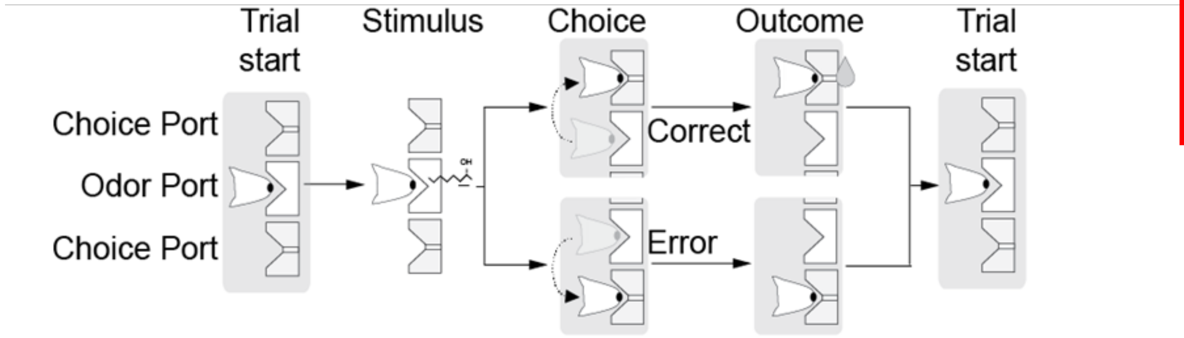

Trial start
Stimulus
Choice
Outcome
Trial start
Choice Port
Odor Port
Choice Port
Correct
Error

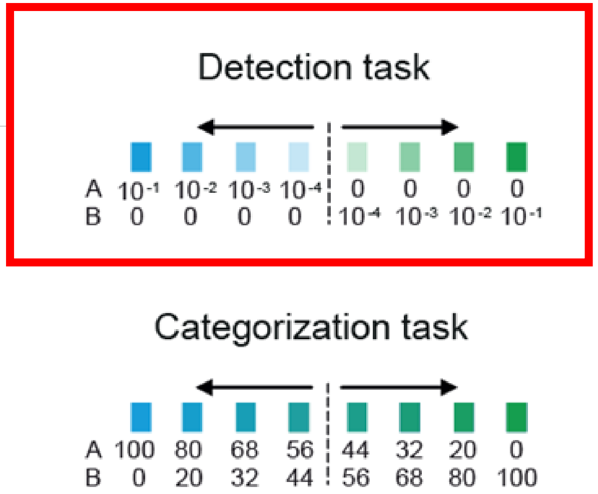

Detection task
A 10-1 10-2 10-3 10-4 0 0 0 0
B 0 0 0 0 10-4 10-3 10-2 10-1
Categorization task
A 100 80 68 56 44 32 20 0
B 0 20 32 44 56 68 80 100

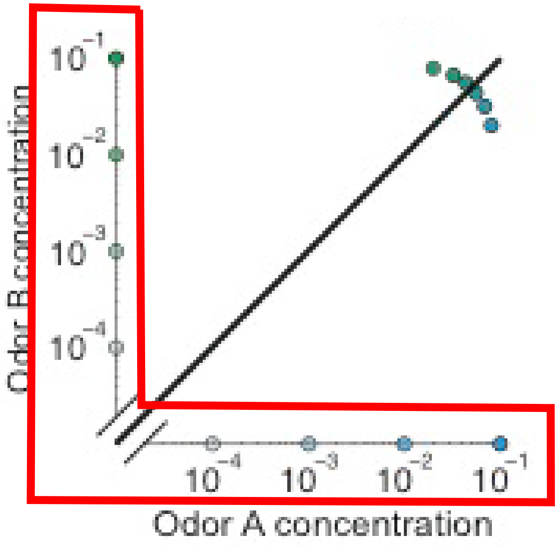

10-1
10-2
10-3
10-4
Odor B concentration
10-4 10-3 10-2 10-1
Odor A concentration

Dual tasks

Trial start
Stimulus
Choice
Outcome
Trial start
Choice Port
Odor Port
Choice Port
Correct
Error
Detection task
A 10-1 10-2 10-3 10-4 0 0 0 0
B 0 0 0 0 10-4 10-3 10-2 10-1
Categorization task
A 100 80 68 56 44 32 20 0
B 0 20 32 44 56 68 80 100
Odor B concentration
Odor A concentration

Dual tasks

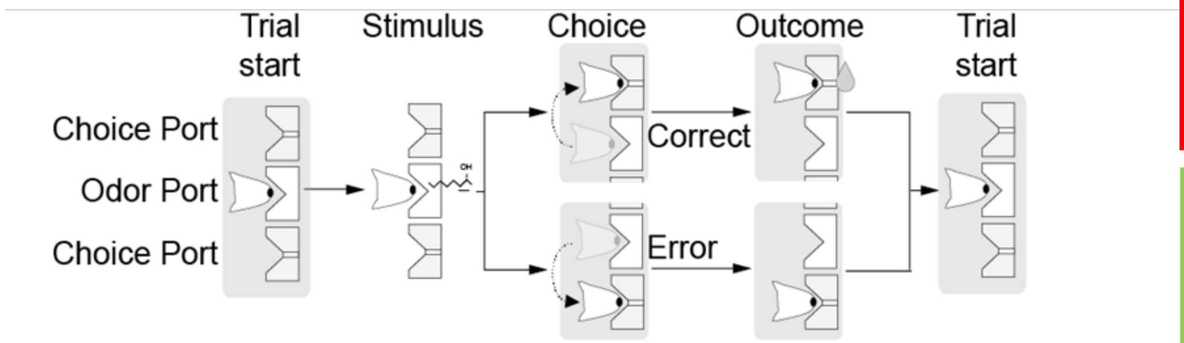
Trial start
Stimulus
Choice
Outcome
Trial start
Choice Port
Odor Port
Choice Port
Correct
Error

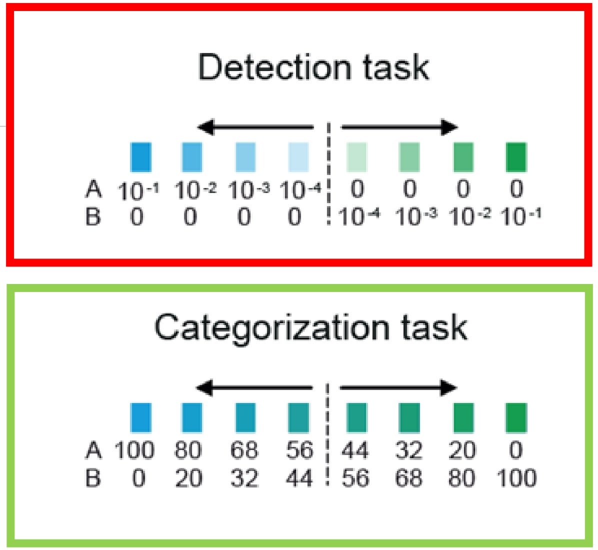
Detection task
A 10-1 10-2 10-3 10-4 0 0 0 0
B 0 0 0 0 10-4 10-3 10-2 10-1
Categorization task
A 100 80 68 56 44 32 20 0
B 0 20 32 44 56 68 80 100

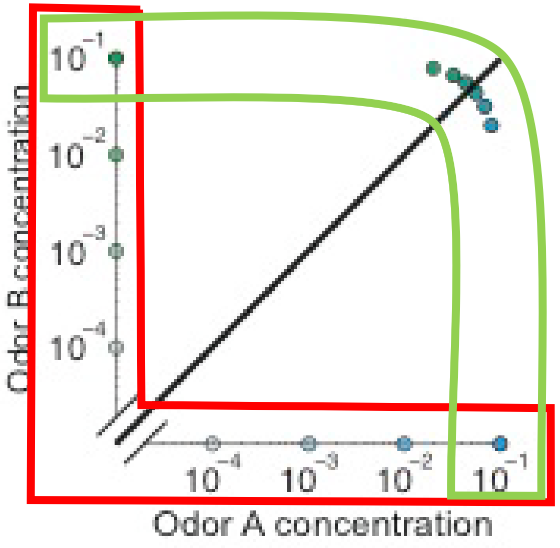
Odor B concentration
Odor A concentration
10-1
10-2
10-3
10-4

Standard Drift Diffusion Model (DDM)

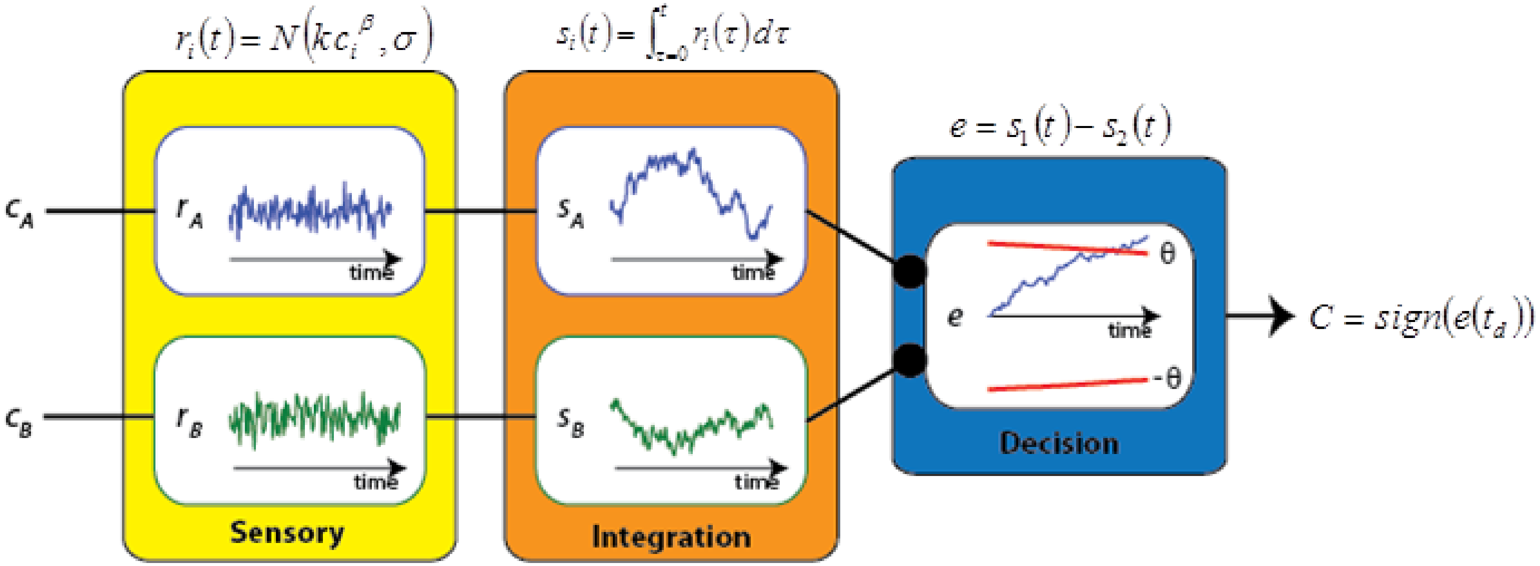

Parameters

1 Sensitivity - k

2 Exponent - β

3 Non-decision tir

4 Bound height -

5 Collapse rate - (

6 Lapse rate - l_r

Predicting Categorization

Detection

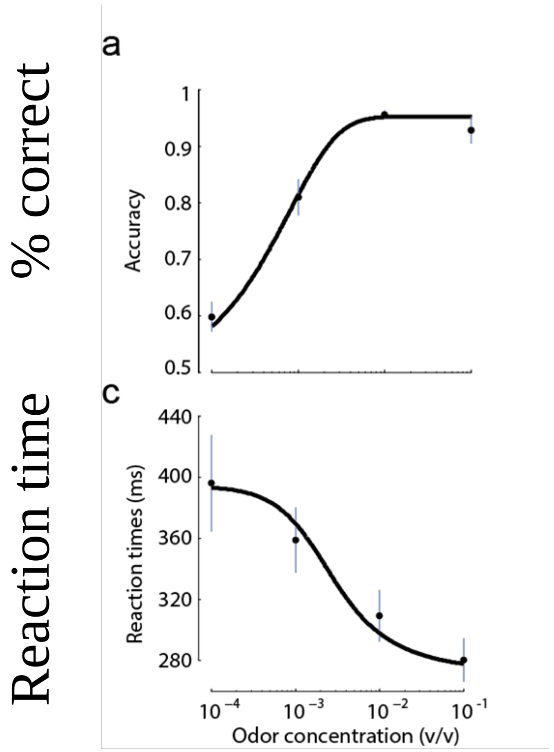

Predicting Categorization

Detection

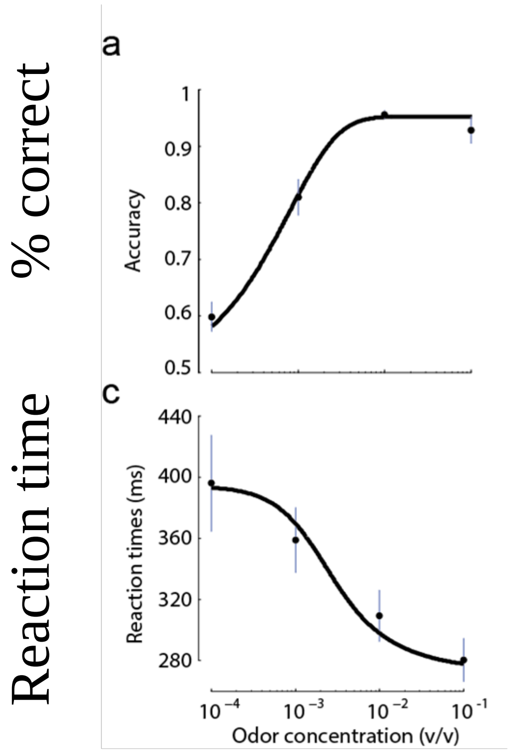

Predicting Categorization

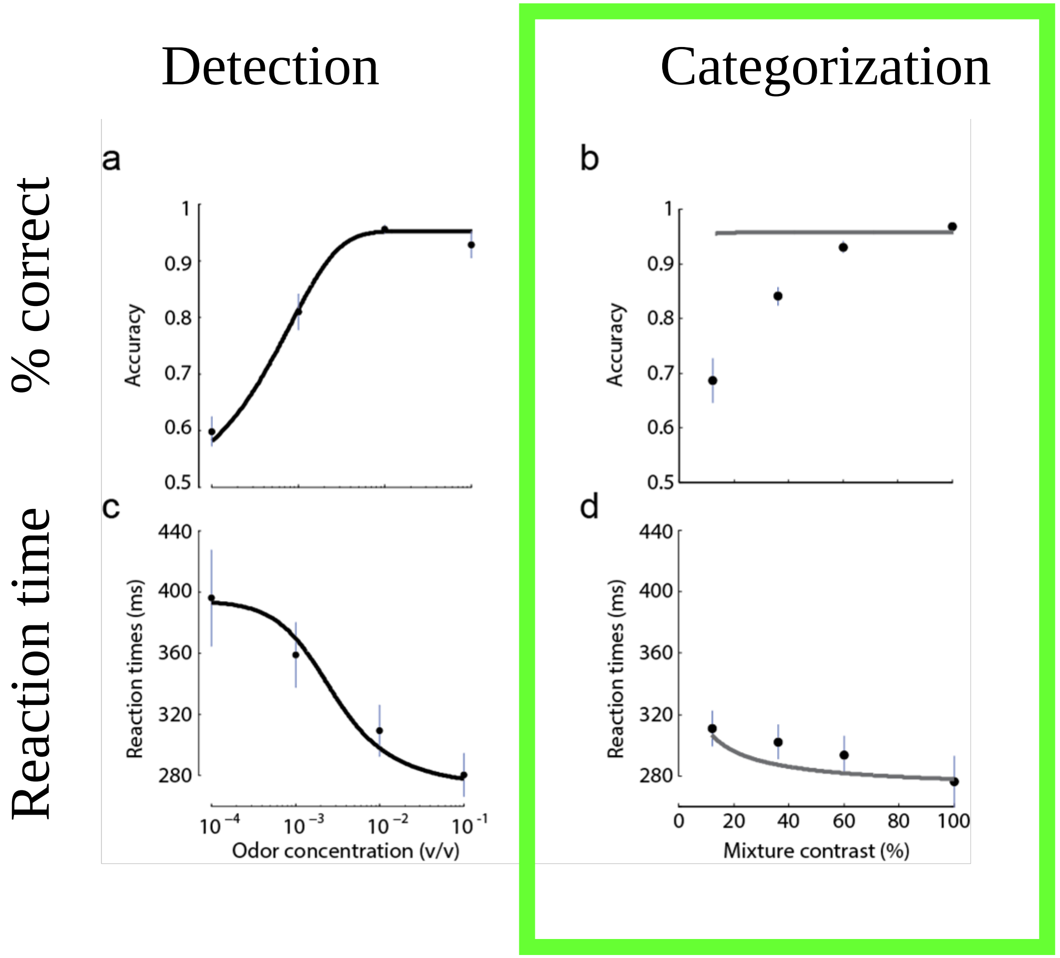

Predicted

Possible Explanation

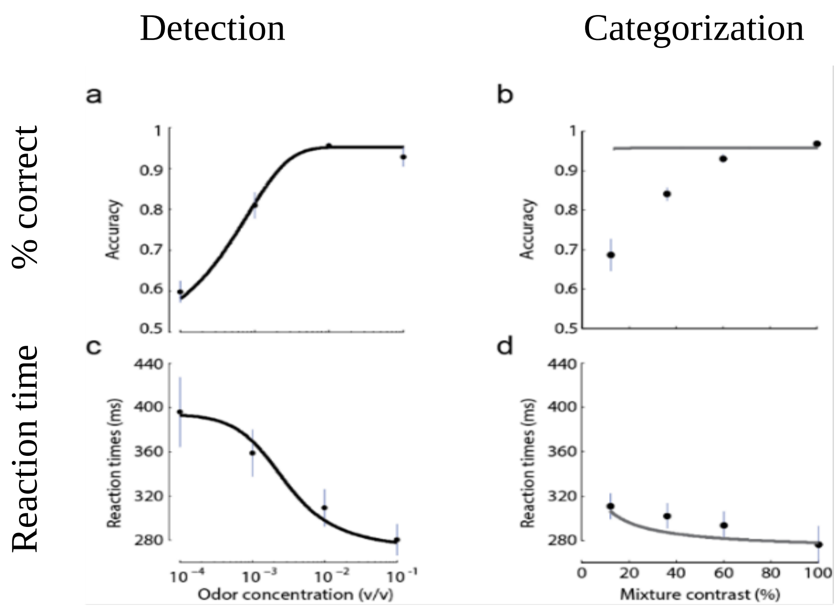

Possible Explanation

- There is an extra source of variability in the categorization task.

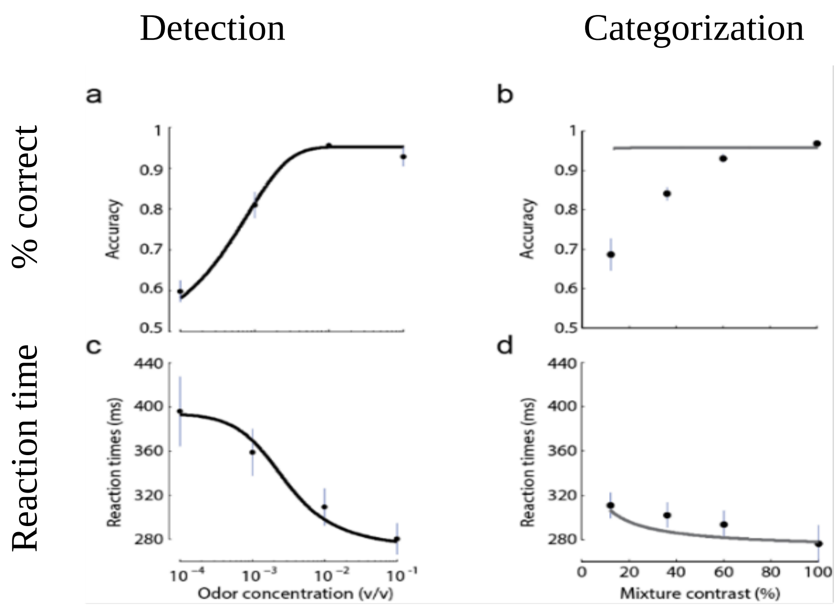

Possible Explanation

- There is an extra source of variability in the categorization task.
- Maybe the animal can't quite remember the 50% boundary, i.e., the memory of the boundary is variable.

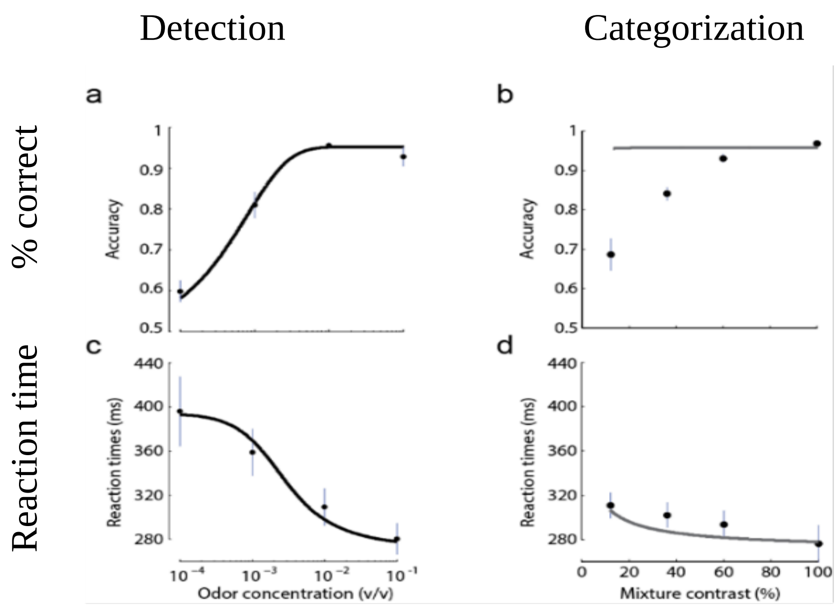

Possible Explanation

- There is an extra source of variability in the categorization task.
- Maybe the animal can't quite remember the 50% boundary, i.e., the memory of the boundary is variable.

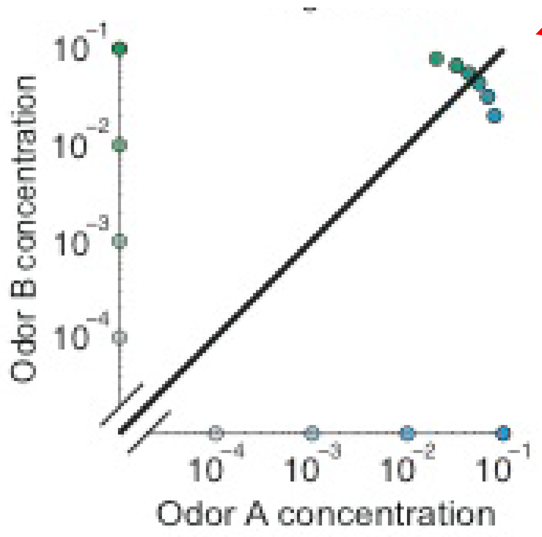

50% categorization boundary

Possible Explanation

- There is an extra source of variability in the categorization task.
- Maybe the animal can’t quite remember the 50% boundary, i.e., the memory of the boundary is variable.

50% categorization boundary

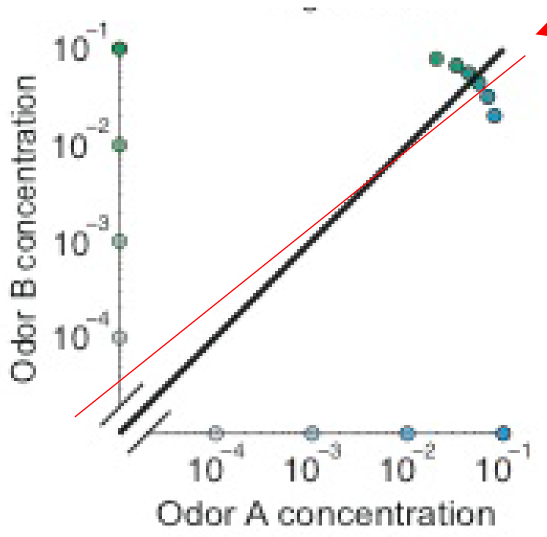

Possible Explanation

- There is an extra source of variability in the categorization task.
- Maybe the animal can’t quite remember the 50% boundary, i.e., the memory of the boundary is variable.

50% categorization boundary

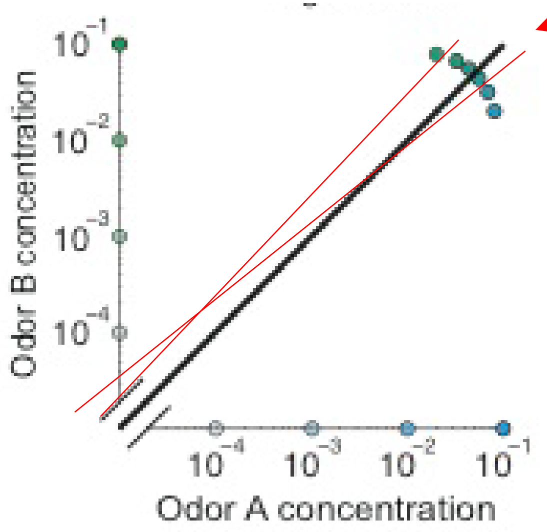

Possible Explanation

- There is an extra source of variability in the categorization task.
- Maybe the animal can't quite remember the 50% boundary, i.e., the memory of the boundary is variable.

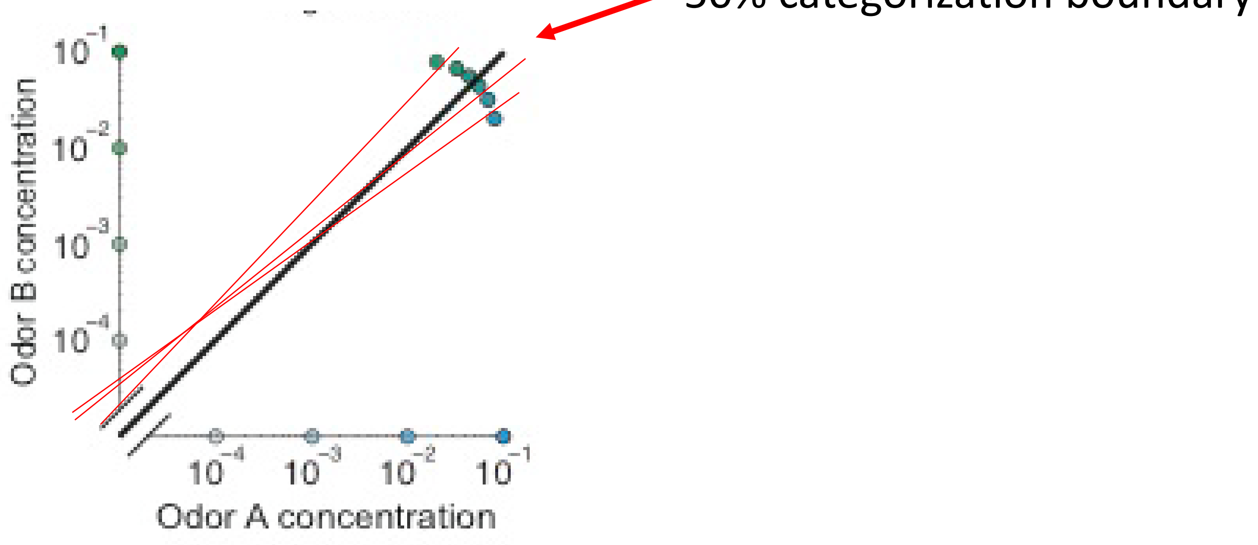

Possible Explanation

- There is an extra source of variability in the categorization task.
- Maybe the animal can’t quite remember the 50% boundary, i.e., the memory of the boundary is variable.

50% categorization boundary

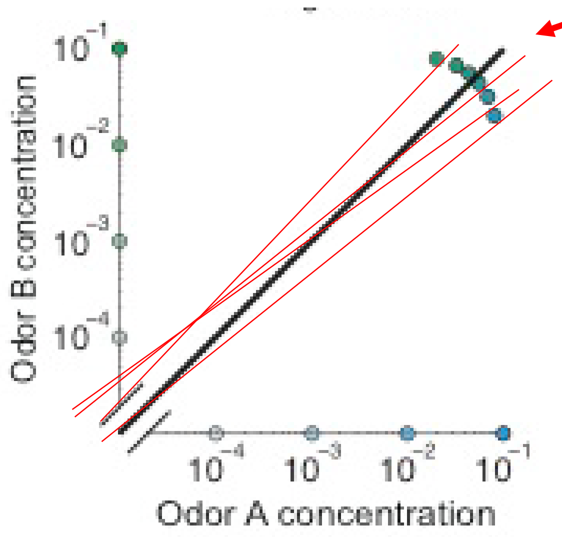

Possible Explanation

- There is an extra source of variability in the categorization task.
- Maybe the animal can't quite remember the 50% boundary, i.e., the memory of the boundary is variable.

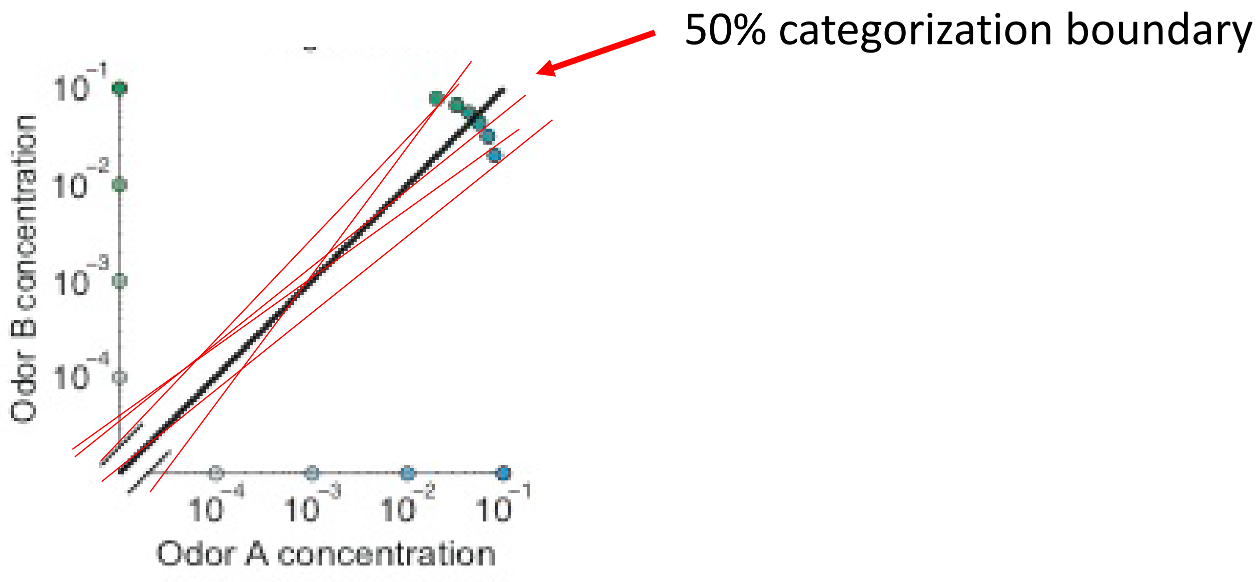

Possible Explanation

- There is an extra source of variability in the categorization task.
- Maybe the animal can't quite remember the 50% boundary, i.e., the memory of the boundary is variable.

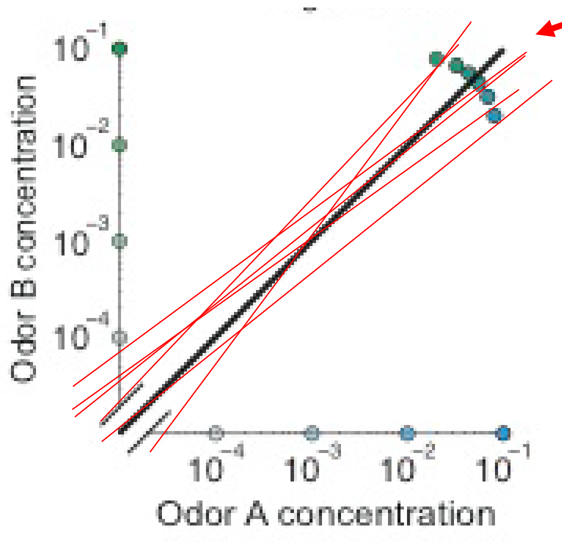

50% categorization boundary

Possible Explanation

- There is an extra source of variability in the categorization task.
- Maybe the animal can't quite remember the 50% boundary, i.e., the memory of the boundary is variable.

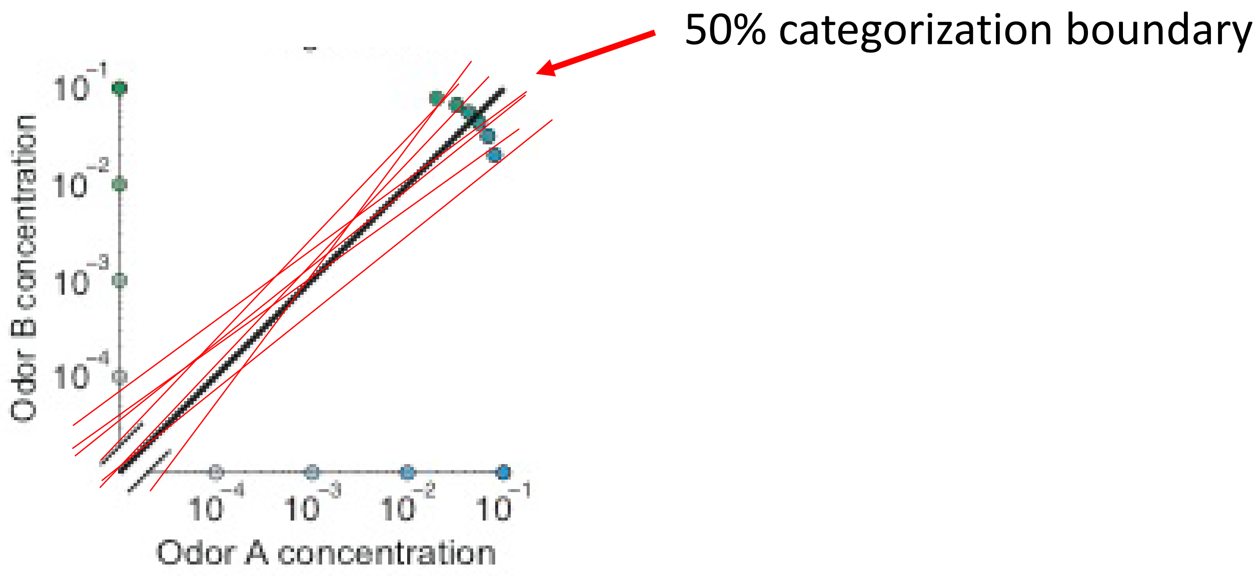

Possible Explanation

- There is an extra source of variability in the categorization task.
- Maybe the animal can’t quite remember the 50% boundary, i.e., the memory of the boundary is variable.

Possible Explanation

- There is an extra source of variability in the categorization task.
- Maybe the animal can’t quite remember the 50% boundary, i.e., the memory of the boundary is variable.

- But why would it be variable?

Possible Explanation

- There is an extra source of variability in the categorization task.
- Maybe the animal can’t quite remember the 50% boundary, i.e., the memory of the boundary is variable.

- But why would it be variable?

- Bad hardware

Possible Explanation

- There is an extra source of variability in the categorization task.
- Maybe the animal can't quite remember the 50% boundary, i.e., the memory of the boundary is variable.

- But why would it be variable?

- Bad hardware

or

Possible Explanation

- There is an extra source of variability in the categorization task.
- Maybe the animal can't quite remember the 50% boundary, i.e., the memory of the boundary is variable.

- But why would it be variable?

- Bad hardware

or

- Wrong assumption about the environment (suboptimal inference)

Wrong World Model

- Perhaps the animal wrongly assumes that the task changes over time

Wrong World Model

- Perhaps the animal wrongly assumes that the task changes over time
- This would lead the animal to adjust the categoration boundary even though it should remain the same once properly learned

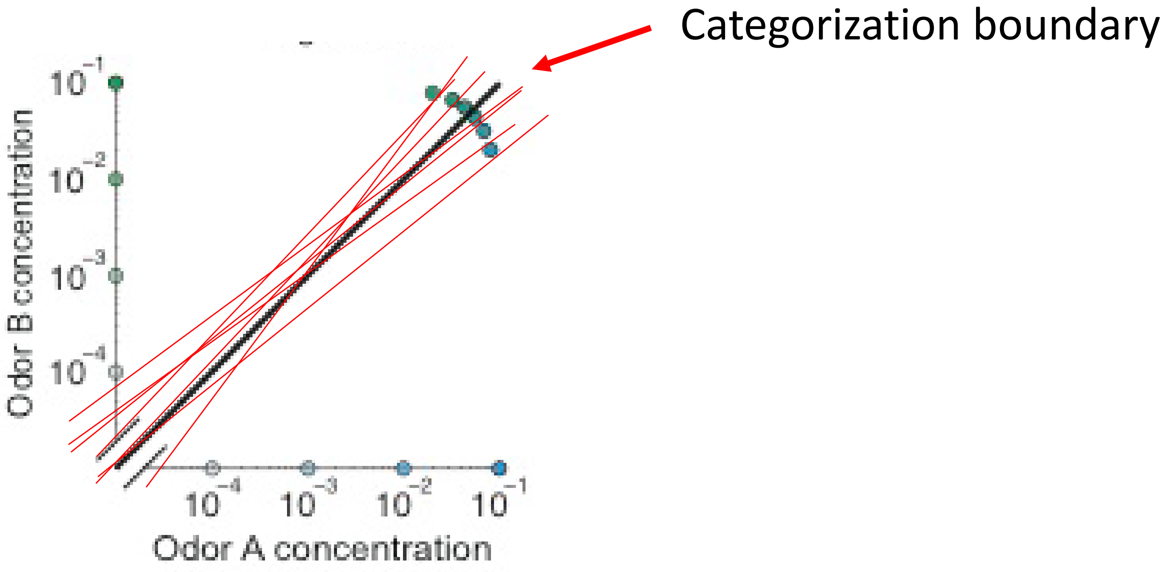

Choice Biases

- Decisions are biased toward the previous choice

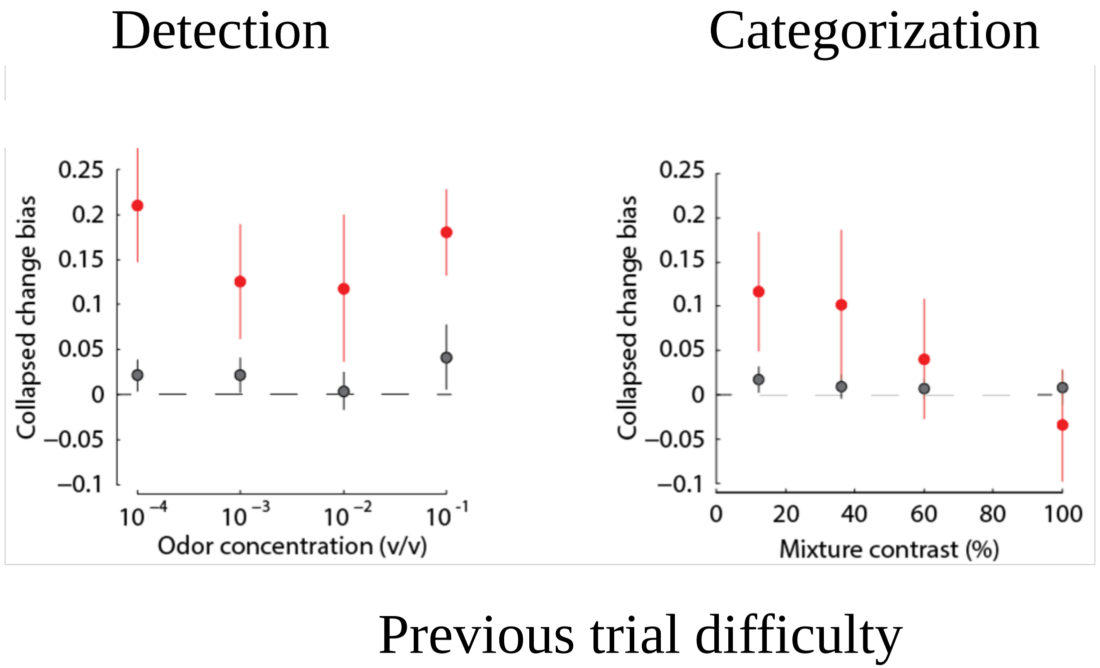

Current trial difficulty:
- hard
- easy

Bad hardware? Stochastic boundary?

Bad hardware? Stochastic boundary?

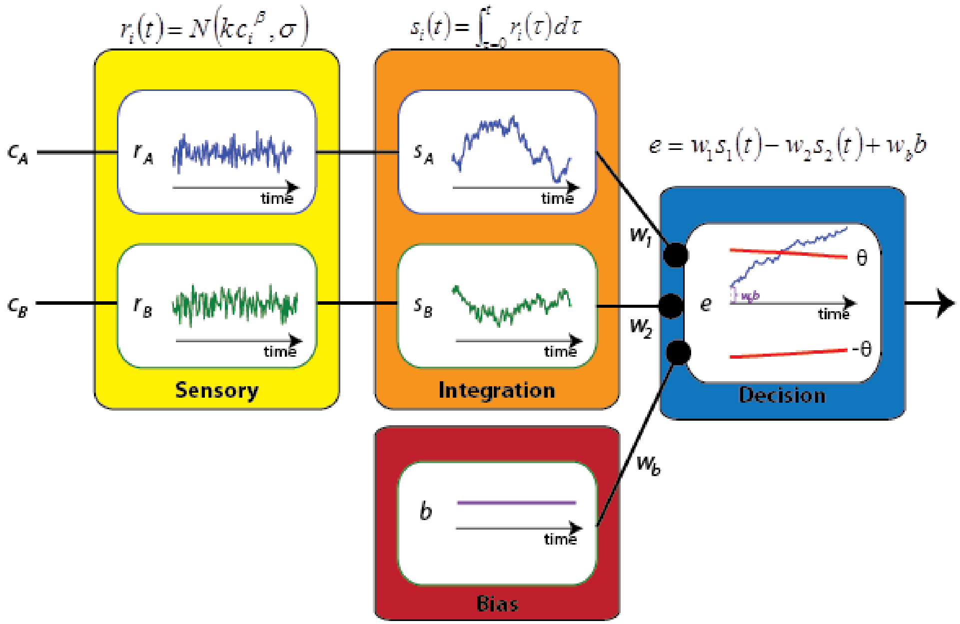

Bad hardware? Stochastic boundary?

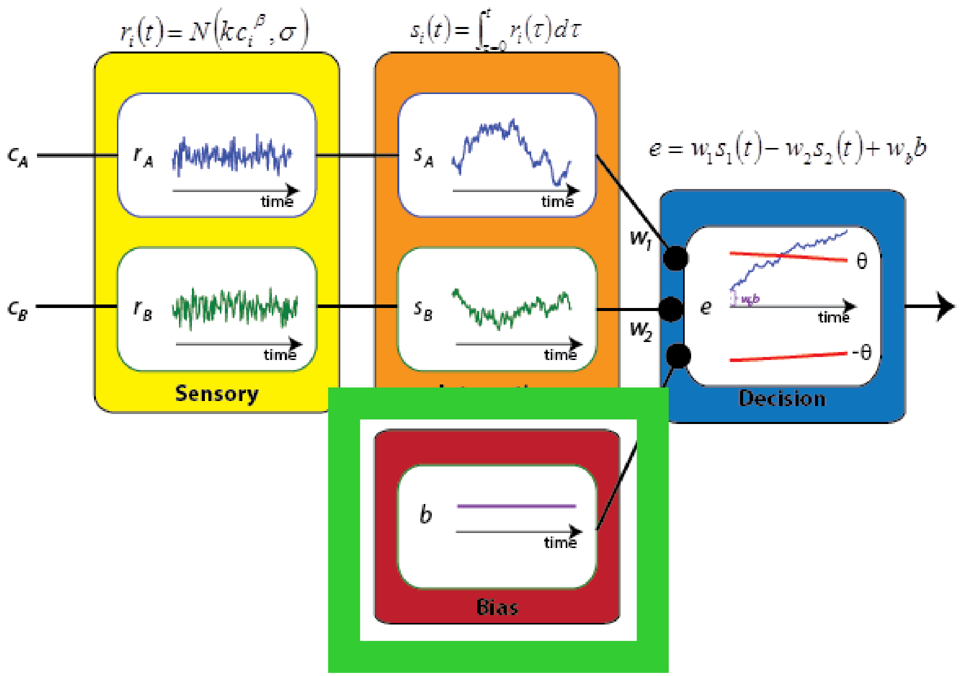

Bad hardware? Stochastic boundary?

- Noise added to the weight

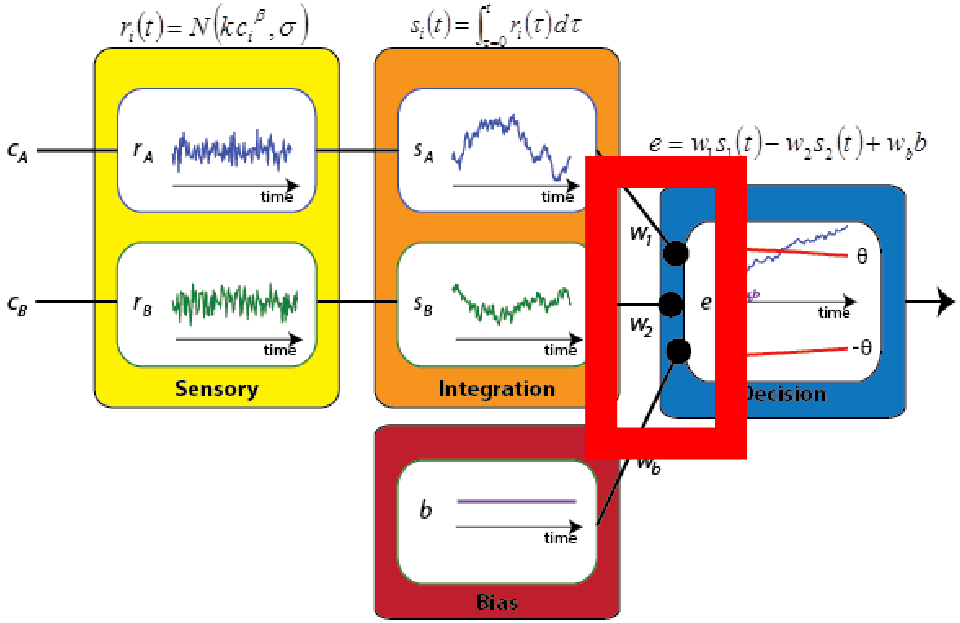

Bad hardware? Stochastic boundary?

- Noise added to the weight

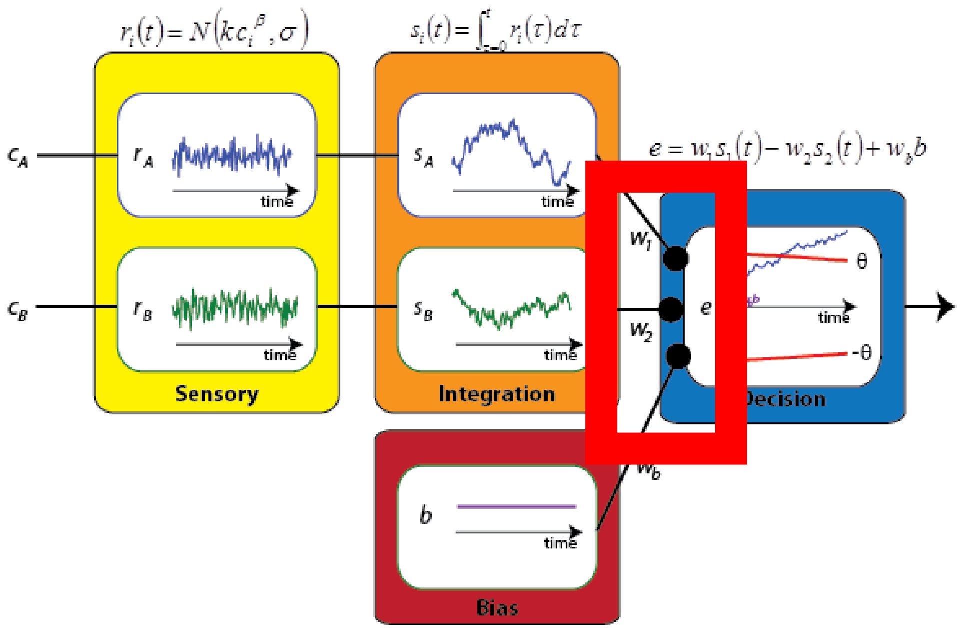

- Weights determine categorization boundary

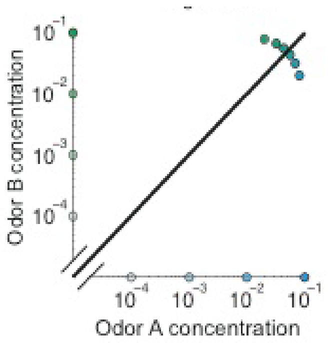

Bad hardware? Stochastic boundary?

- Noise added to the weight

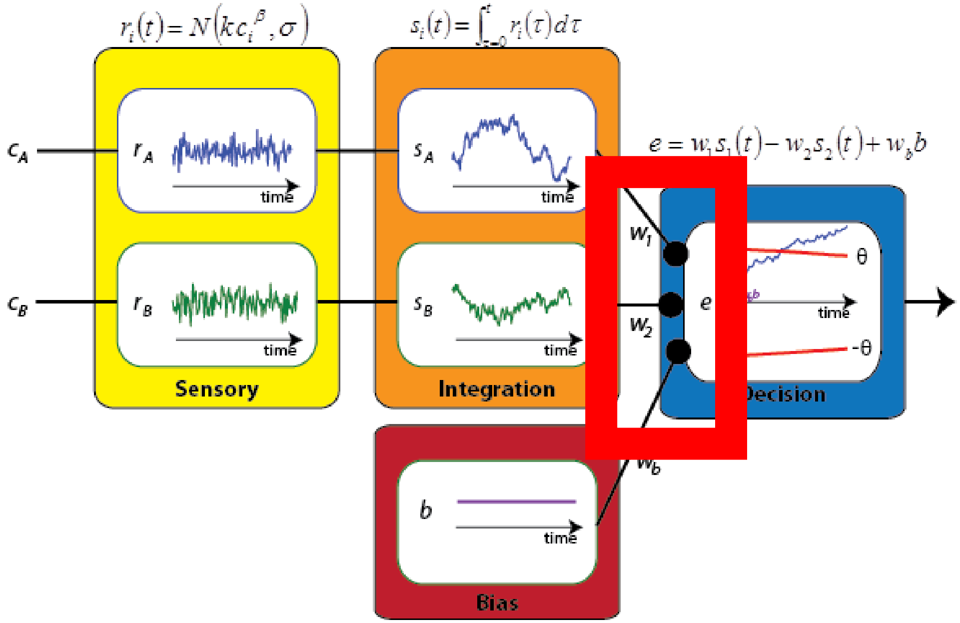

- Weights determine categorization boundary
- Weights of (1,-1,0) correspond to black diagonal

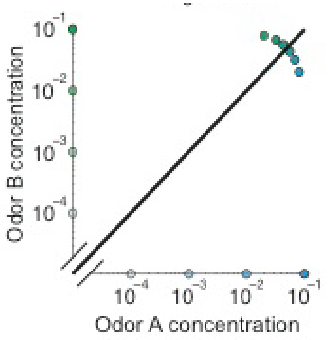

Bad hardware? Stochastic boundary?

- Noise added to the weight

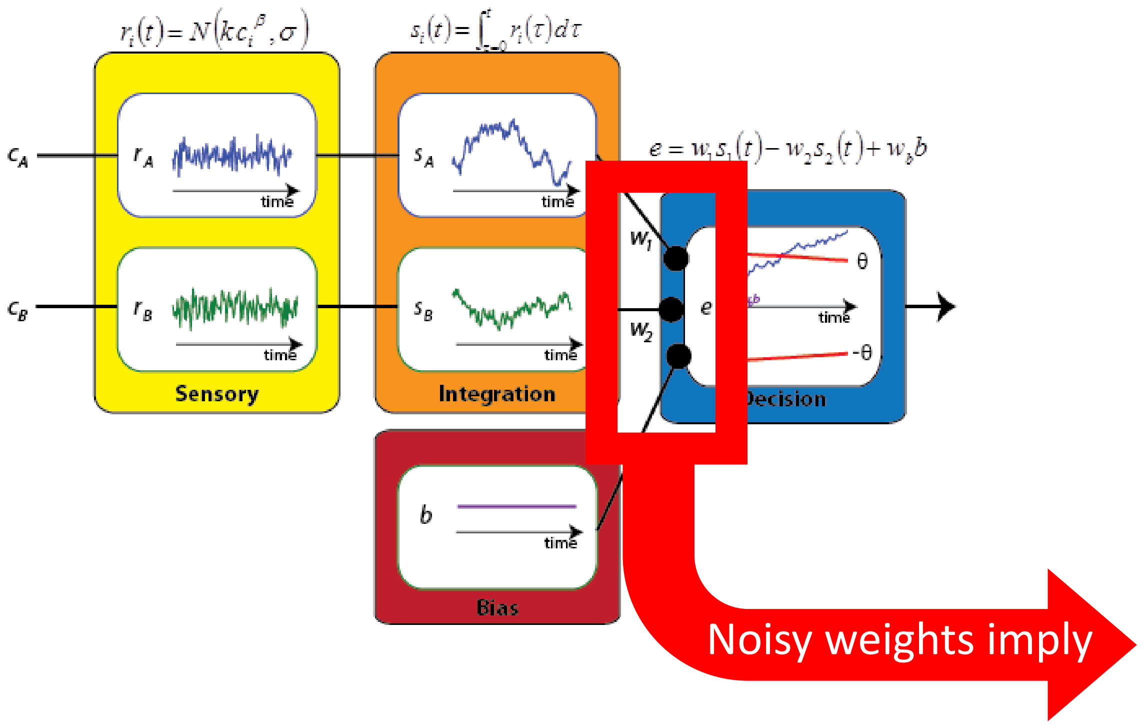

- Weights determine categorization boundary
- Weights of (1,-1,0) correspond to black diagonal

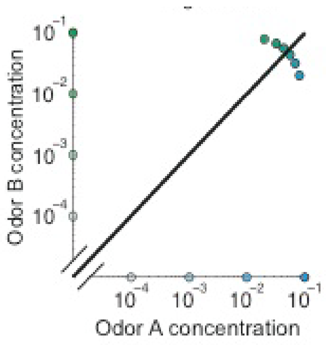

Bad hardware? Stochastic boundary?

- Noise added to the weight

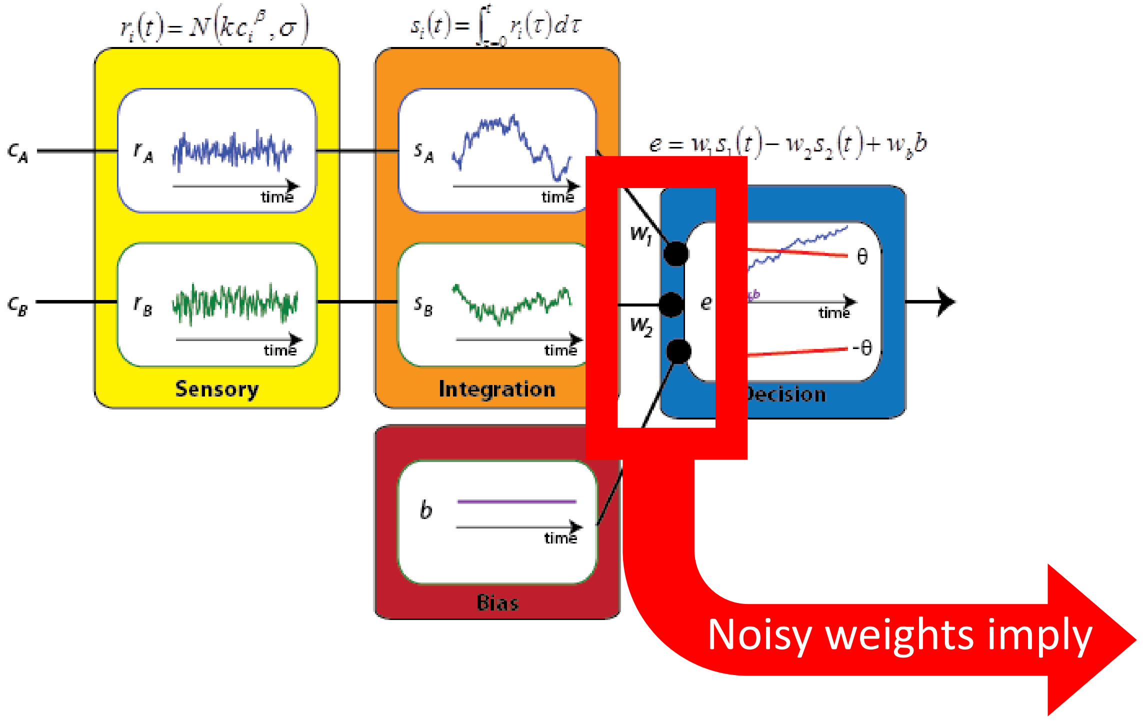

- Weights determine categorization boundary
- Weights of (1,-1,0) correspond to black diagonal

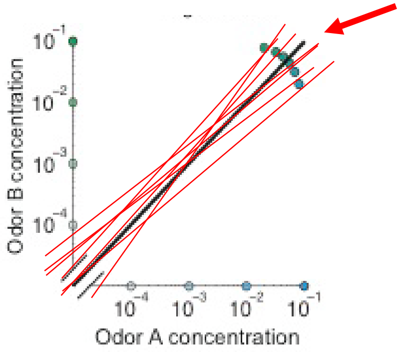

Noise in categorization boundary

Stochastic boundary is not enough

Detection

Categorization

% correct

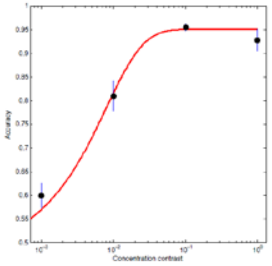

Reaction time

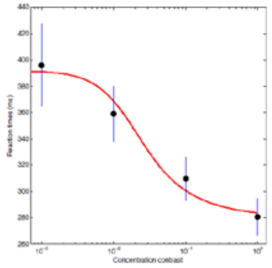

Stochastic boundary is not enough

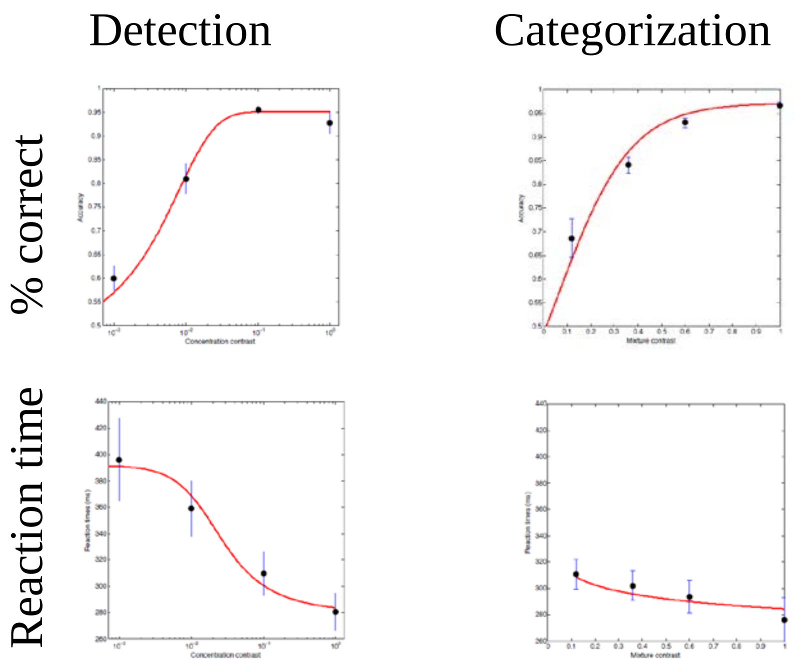

Stochastic boundary is not enough

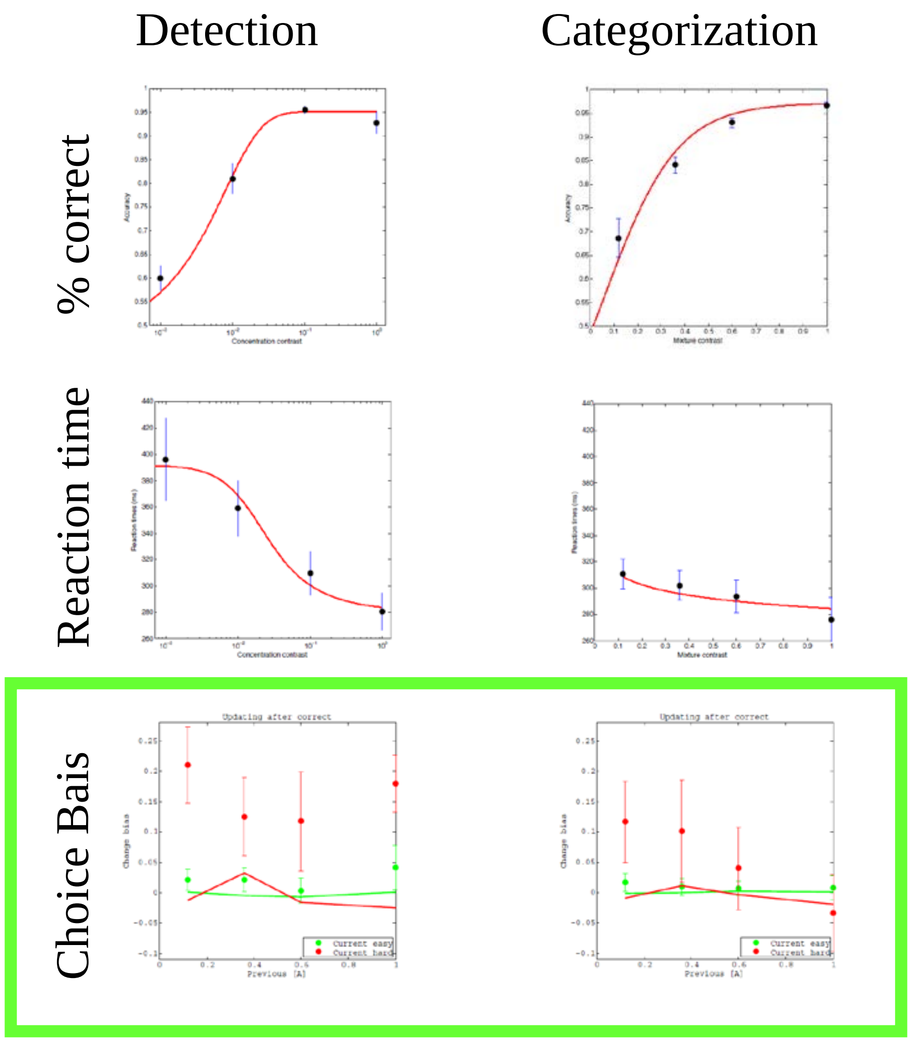

Predicted

Adaptive DDM

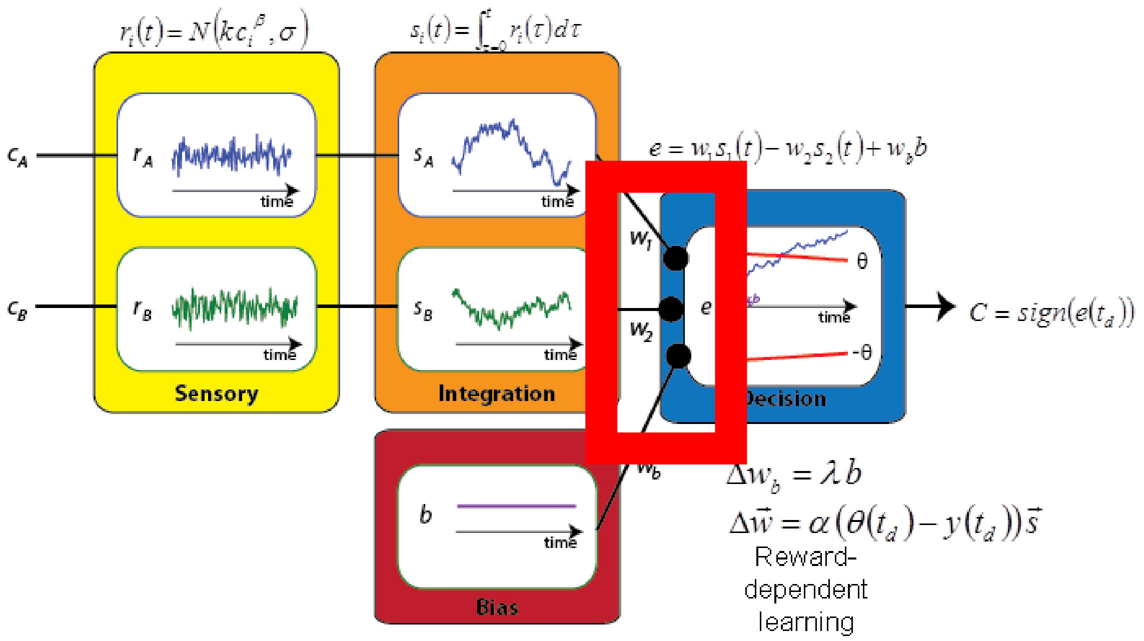

Adaptive DDM

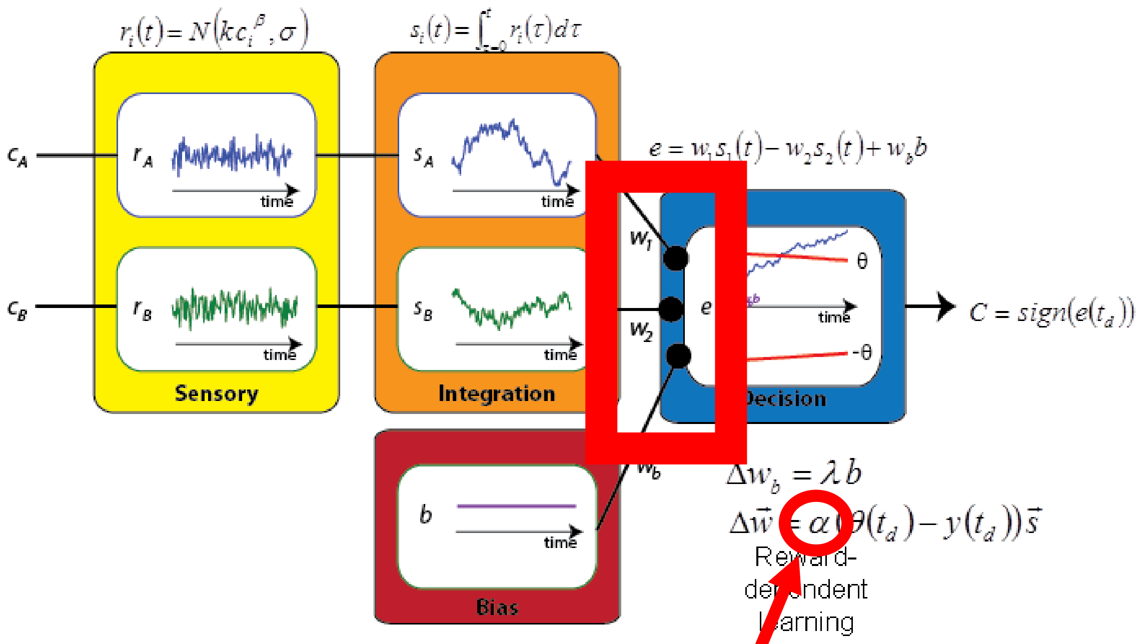

Learning rate

Fits with Adaptive DDM

Detection

Categorization

% correct

Reaction time

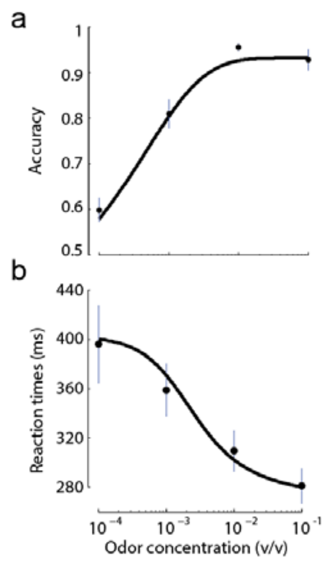

Fits with Adaptive DDM

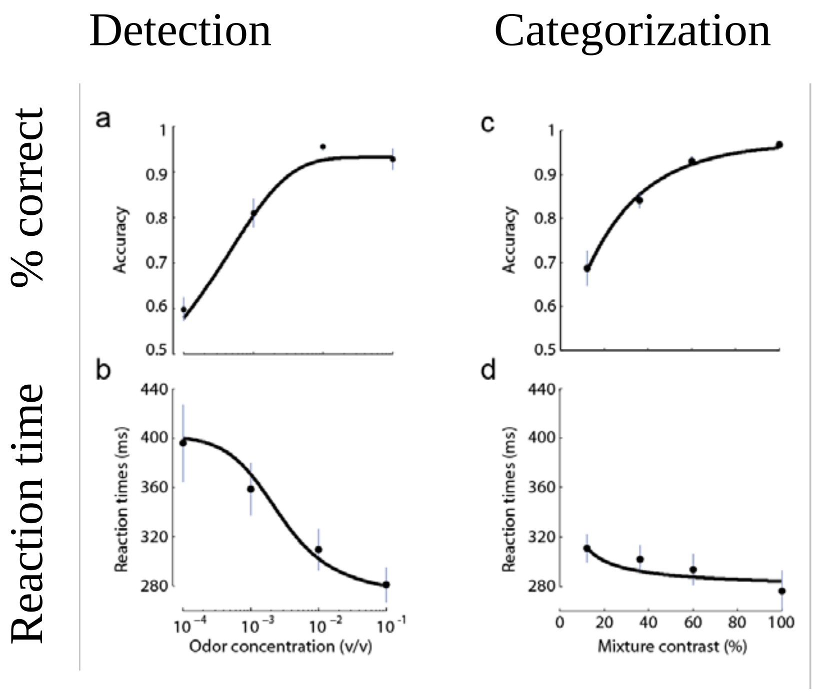

Fits with Adaptive DDM

Detection | Categorization

% correct

Reaction time

Choice Bais

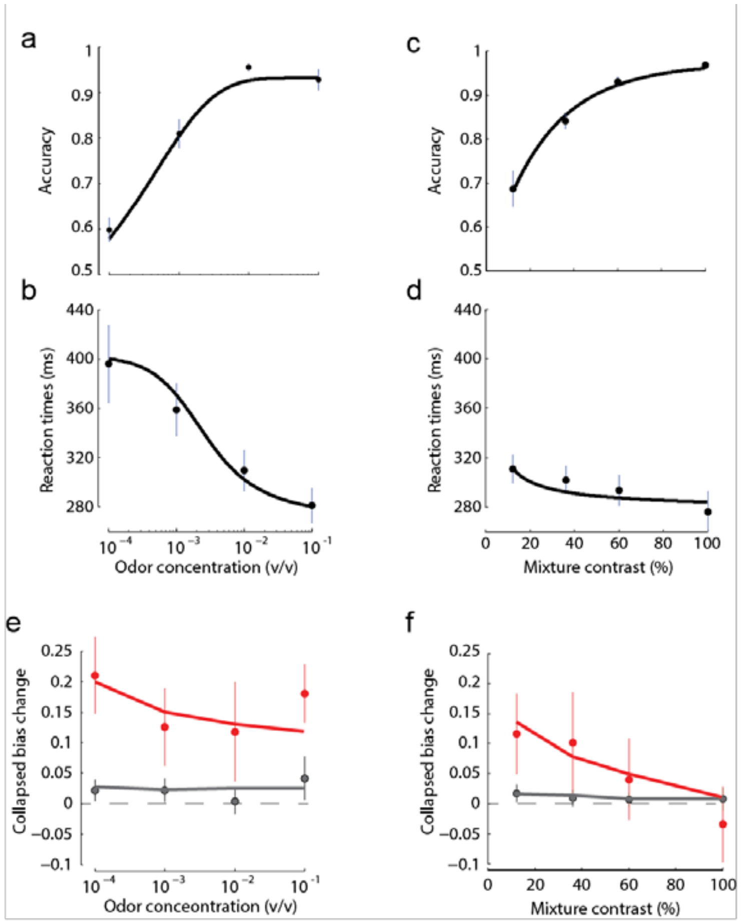

Fits with Adaptive DDM

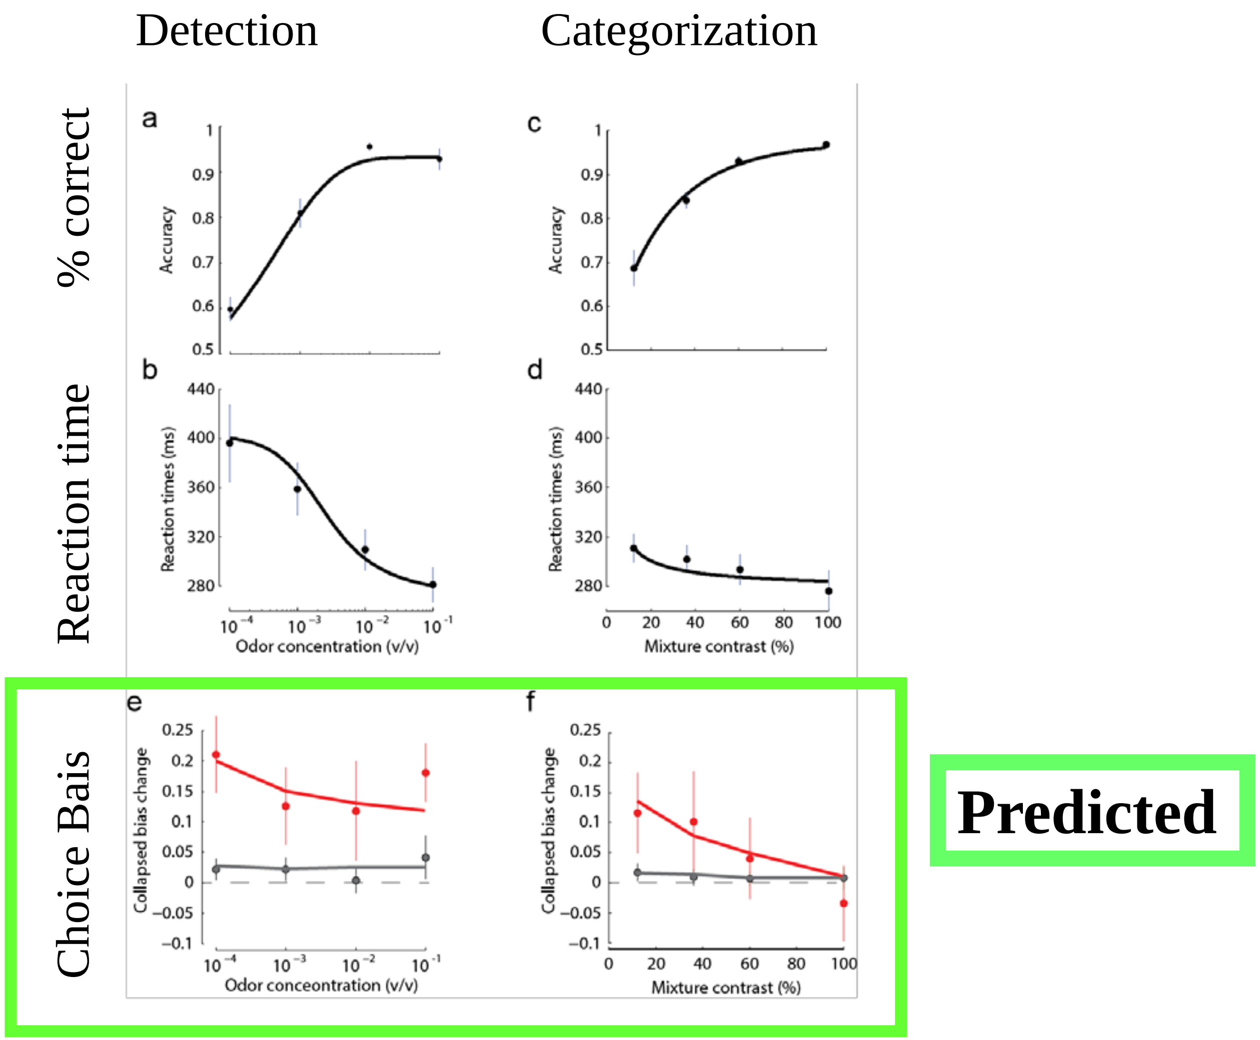

Prediction For All Mixtures

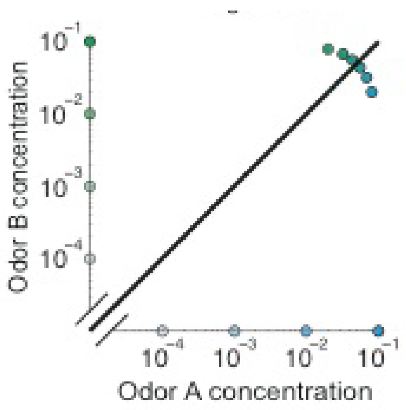

Prediction For All Mixtures

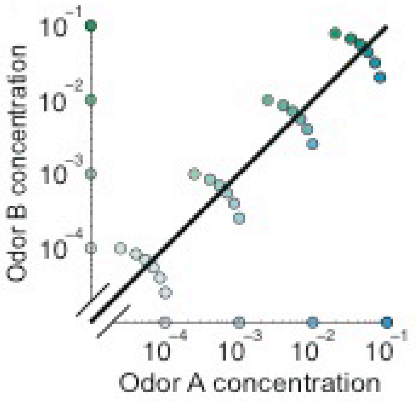

Prediction For All Mixtures

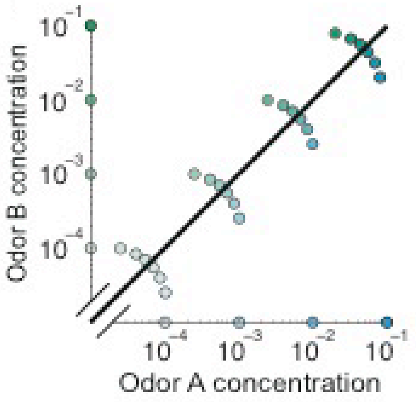

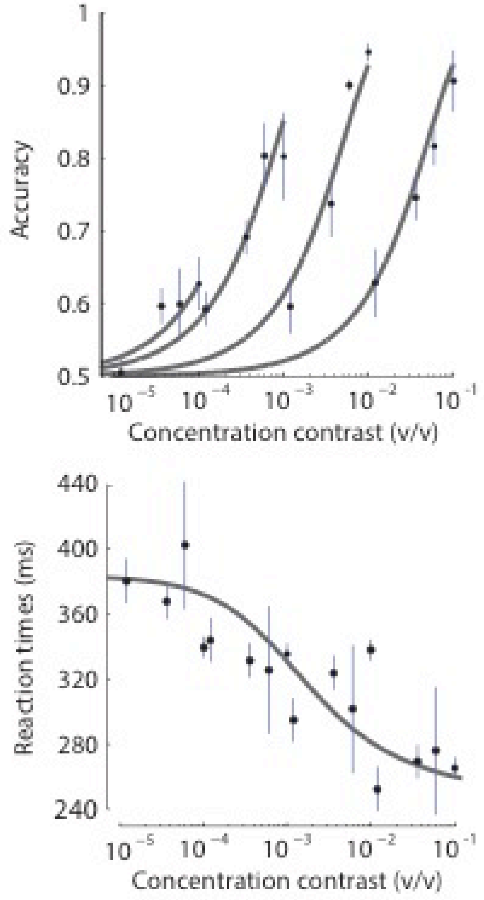

General Principle

General Principle

- If you don’t know how the data were generated, behavior cannot be optimal, which leads to extra variability.

General Principle

- If you don’t know how the data were generated, behavior cannot be optimal, which leads to extra variability.

- For most problems of interest, wrong assumptions about the data generating process is the main source of variability

General Principle

- If you don't know how the data were generated, behavior cannot be optimal, which leads to extra variability.

- For most problems of interest, wrong assumptions about the data generating process is the main source of variability

Not Noisy, Just Wrong: The Role of Suboptimal Inference in Behavioral Variability

Jeffrey M. Beck,[1,5] Wei Ji Ma,[2,5] Xaq Pitkow,[1] Peter E. Latham,[3] and Alexandre Pouget[1,3,4,*]

Roadmap

Experiment 1: what appears as noise can be suboptimal inference

Experiment 2: noise only has a marginal impact on performance

Noise and Decoding

In collaboration with Pitkow, Lakshminarasimhan, DeAngelis and Angelaki

Noise and Decoding

In collaboration with Pitkow, Lakshminarasimhan, DeAngelis and Angelaki

Noise and Decoding

Does internal noise affect behavioral performance?

In collaboration with Pitkow, Lakshminarasimhan, DeAngelis and Angelaki

Noise and Decoding

Does internal noise affect behavioral performance?

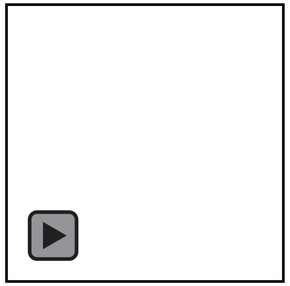

In collaboration with Pitkow, Lakshminarasimhan, DeAngelis and Angelaki

Noise and Decoding

Does internal noise affect behavioral performance?

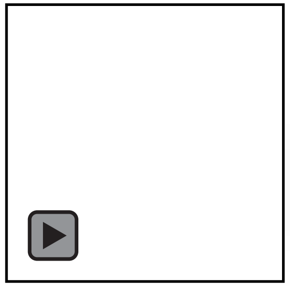

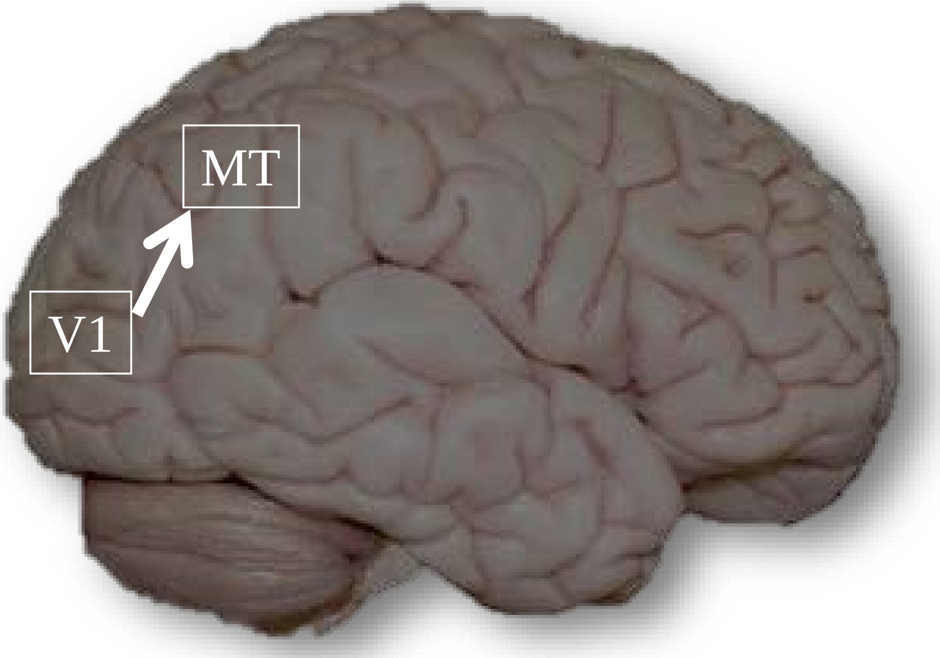

In collaboration with Pitkow, Lakshminarasimhan, DeAngelis and Angelaki

Noise and Decoding

Does internal noise affect behavioral performance?

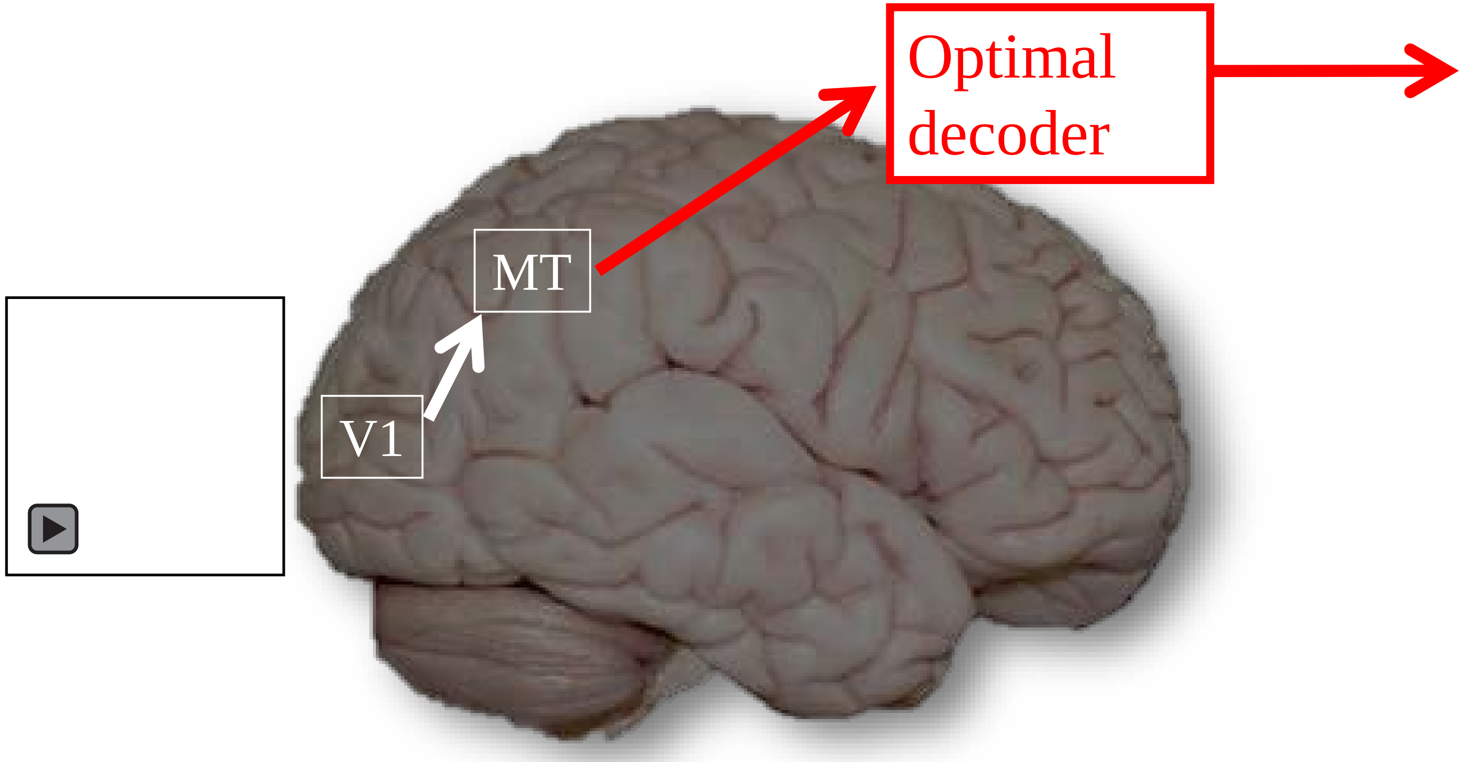

Decoder performance

In collaboration with Pitkow, Lakshminarasimhan, DeAngelis and Angelaki

Noise and Decoding

Does internal noise affect behavioral performance?

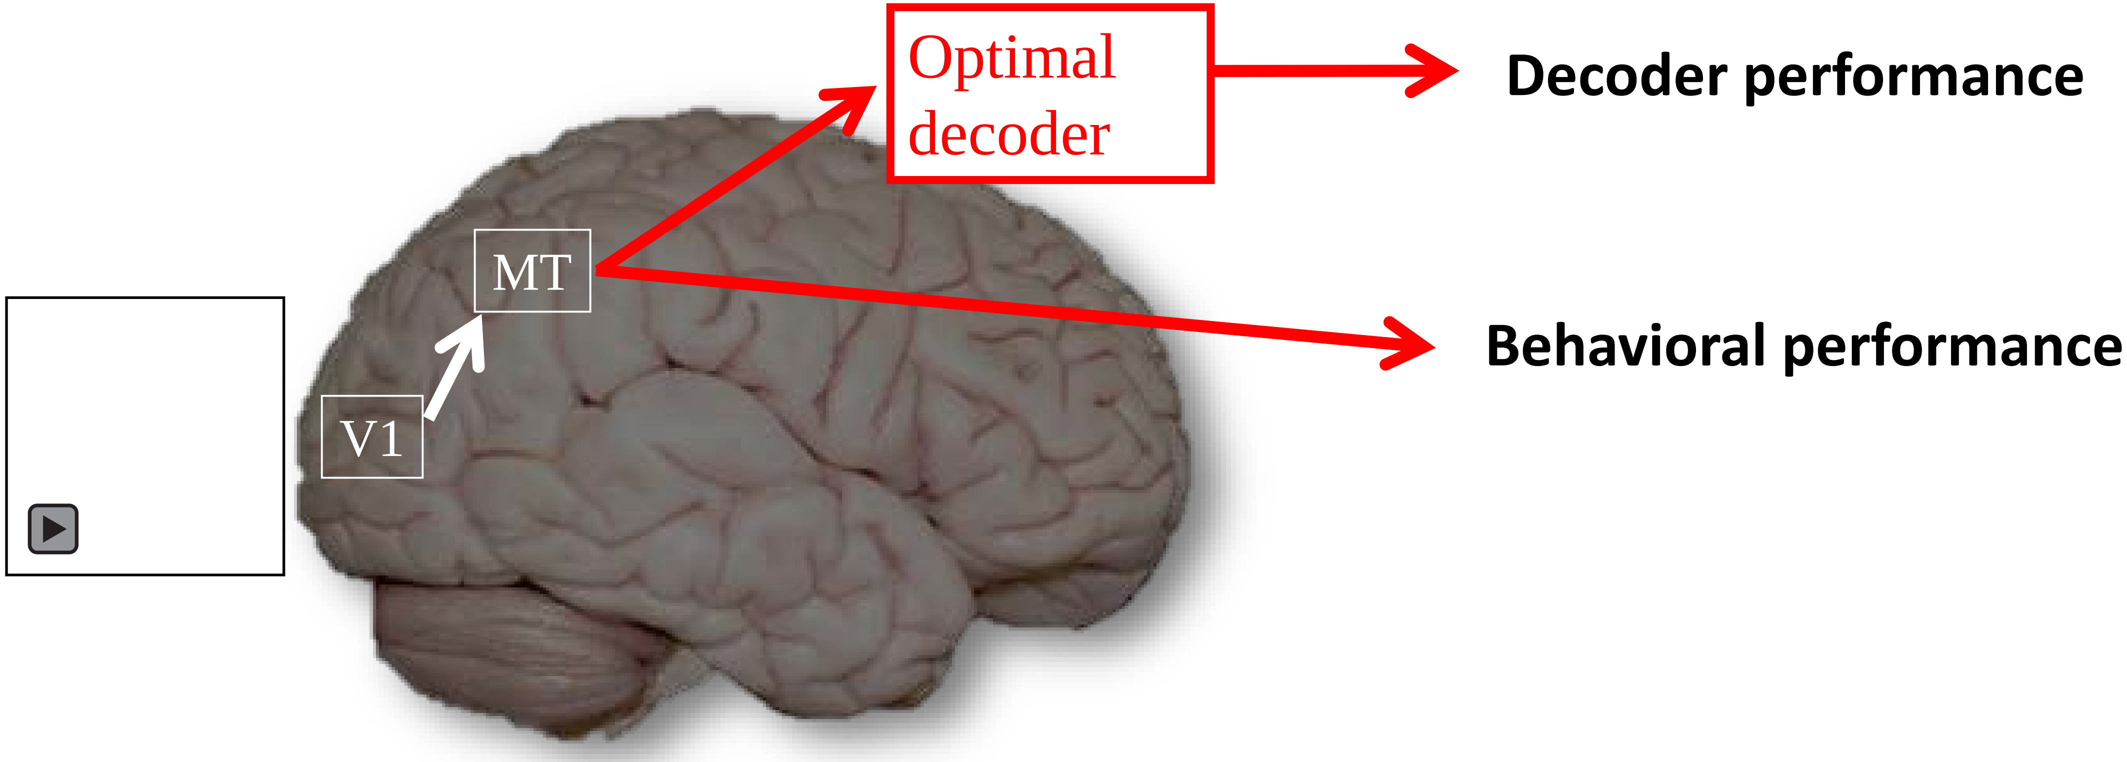

In collaboration with Pitkow, Lakshminarasimhan, DeAngelis and Angelaki

Noise and Decoding

Does internal noise affect behavioral performance?

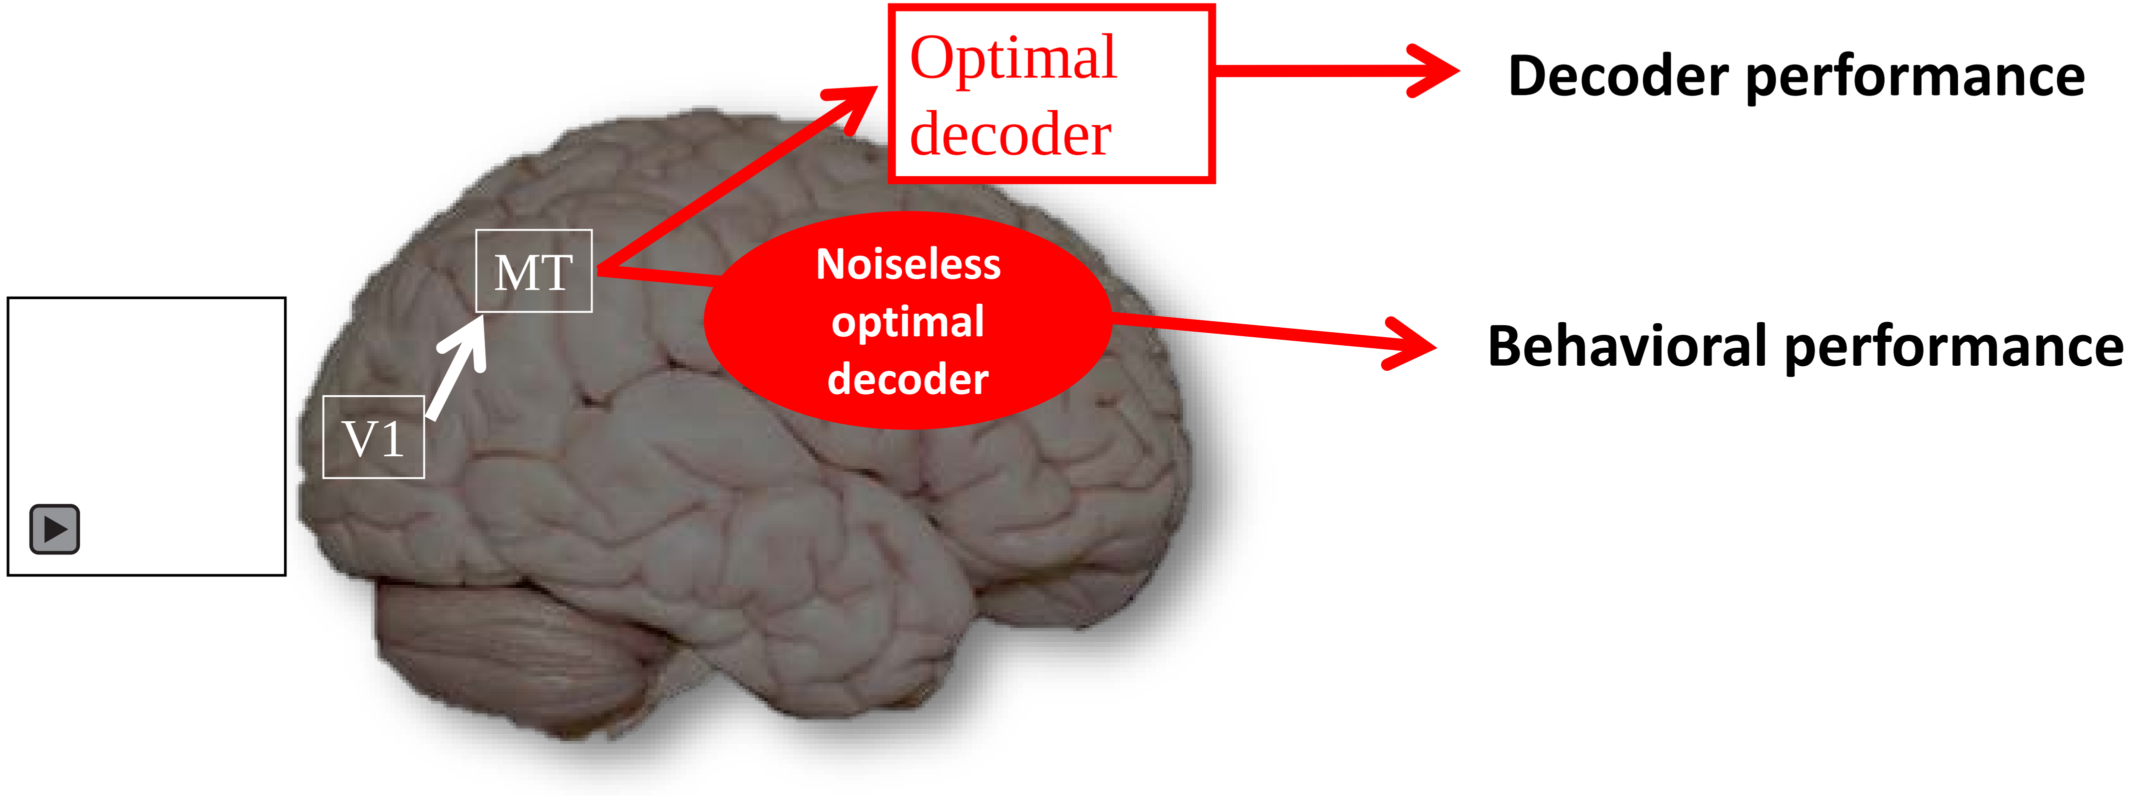

In collaboration with Pitkow, Lakshminarasimhan, DeAngelis and Angelaki

Noise and Decoding

Does internal noise affect behavioral performance?

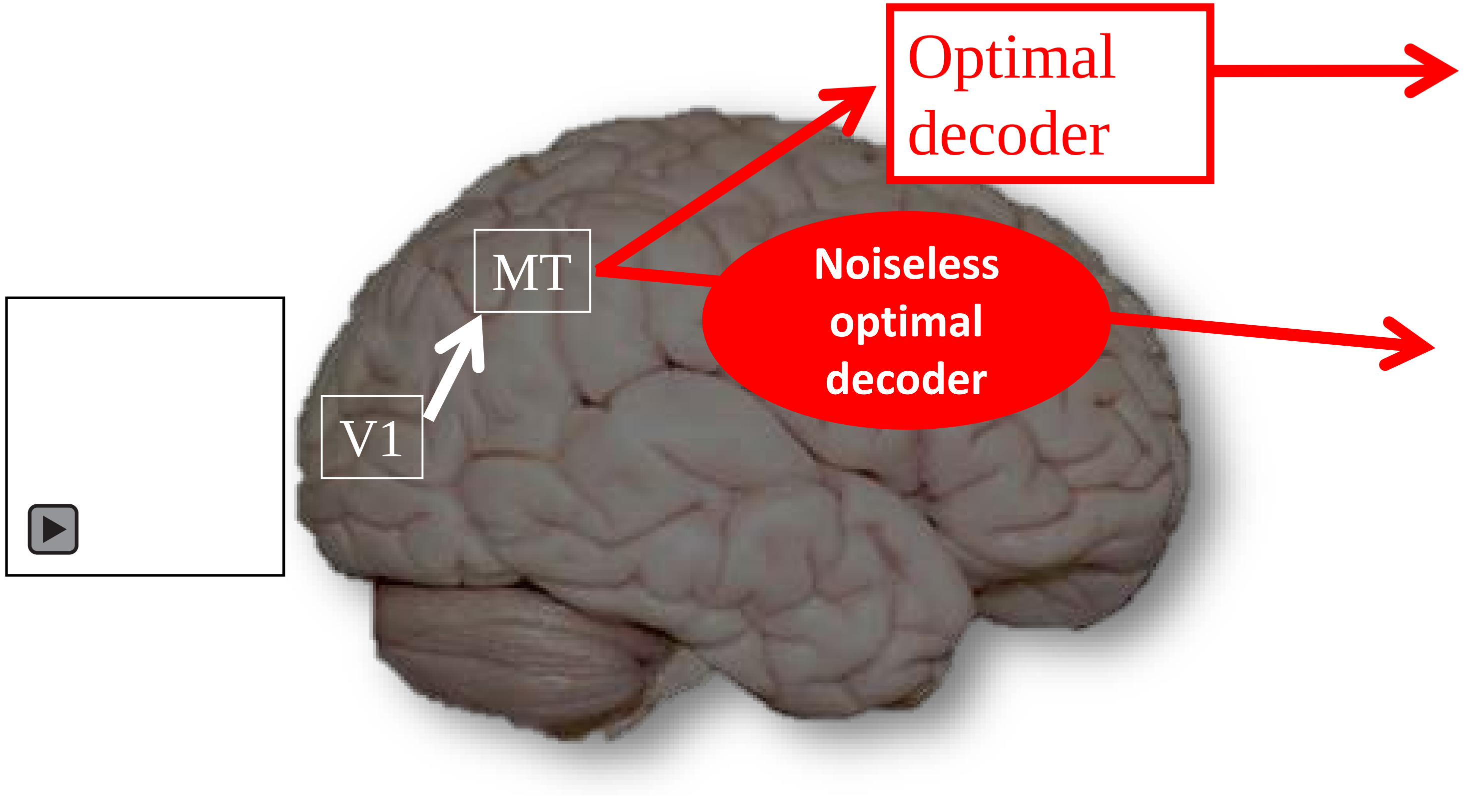

Decoder performance

=

Behavioral performance

In collaboration with Pitkow, Lakshminarasimhan, DeAngelis and Angelaki

Noise and Decoding

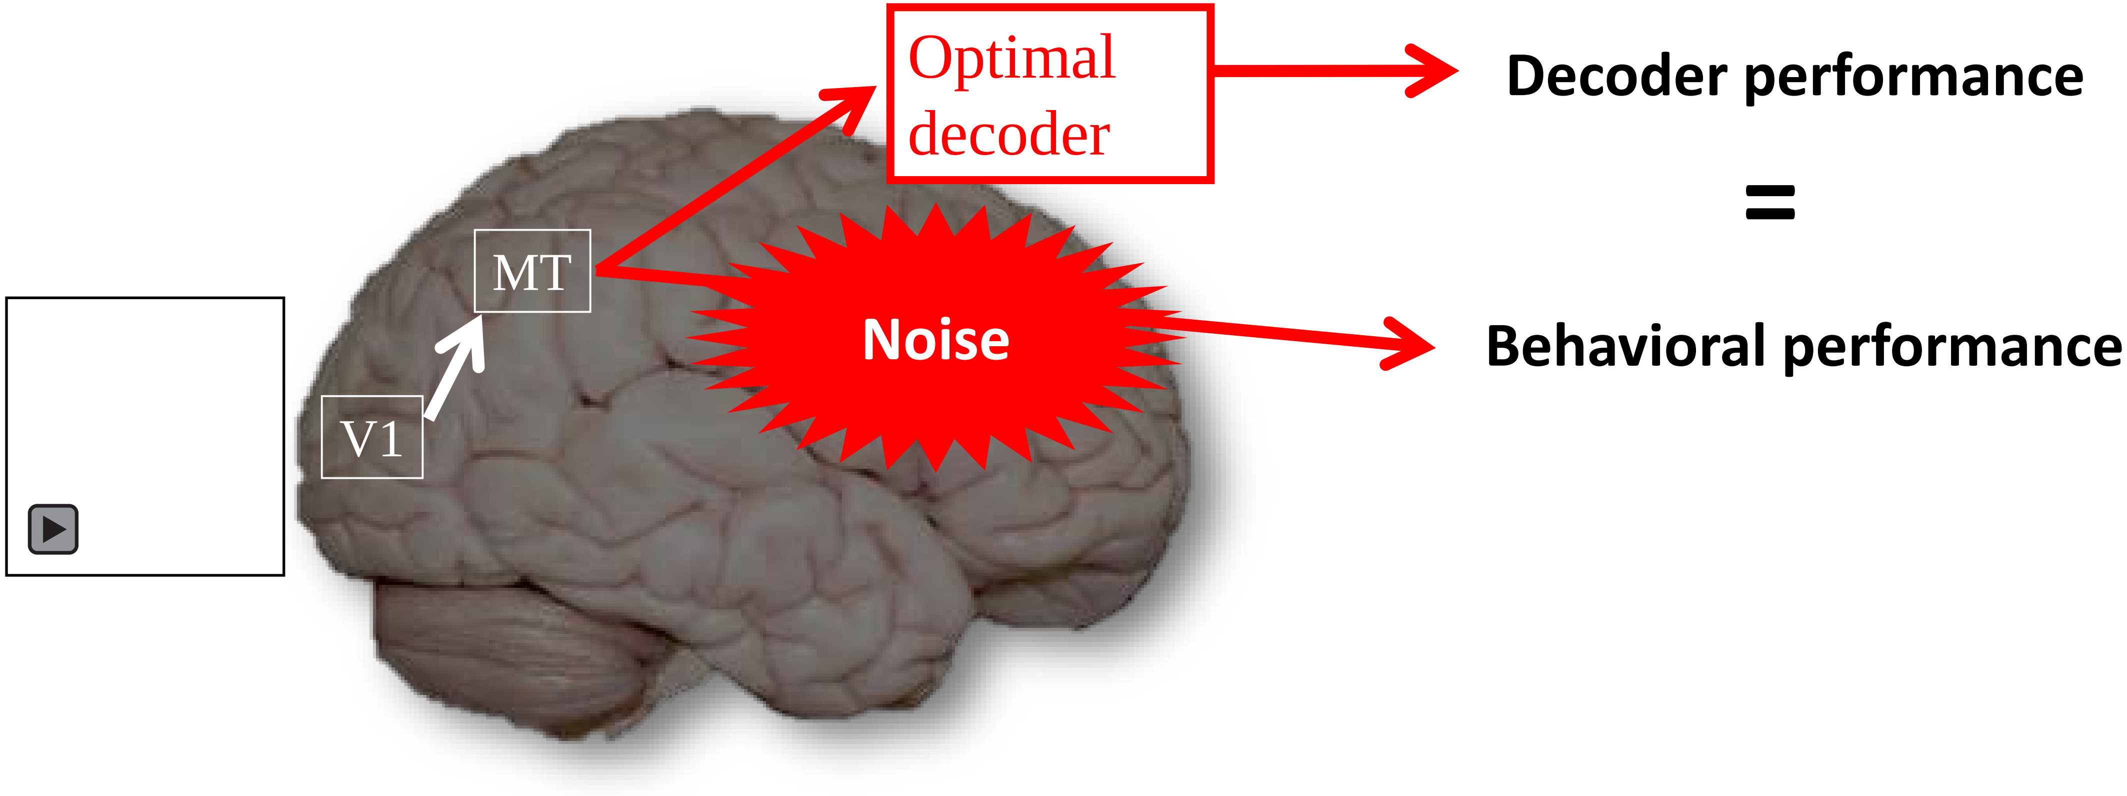

In collaboration with Pitkow, Lakshminarasimhan, DeAngelis and Angelaki

Noise and Decoding

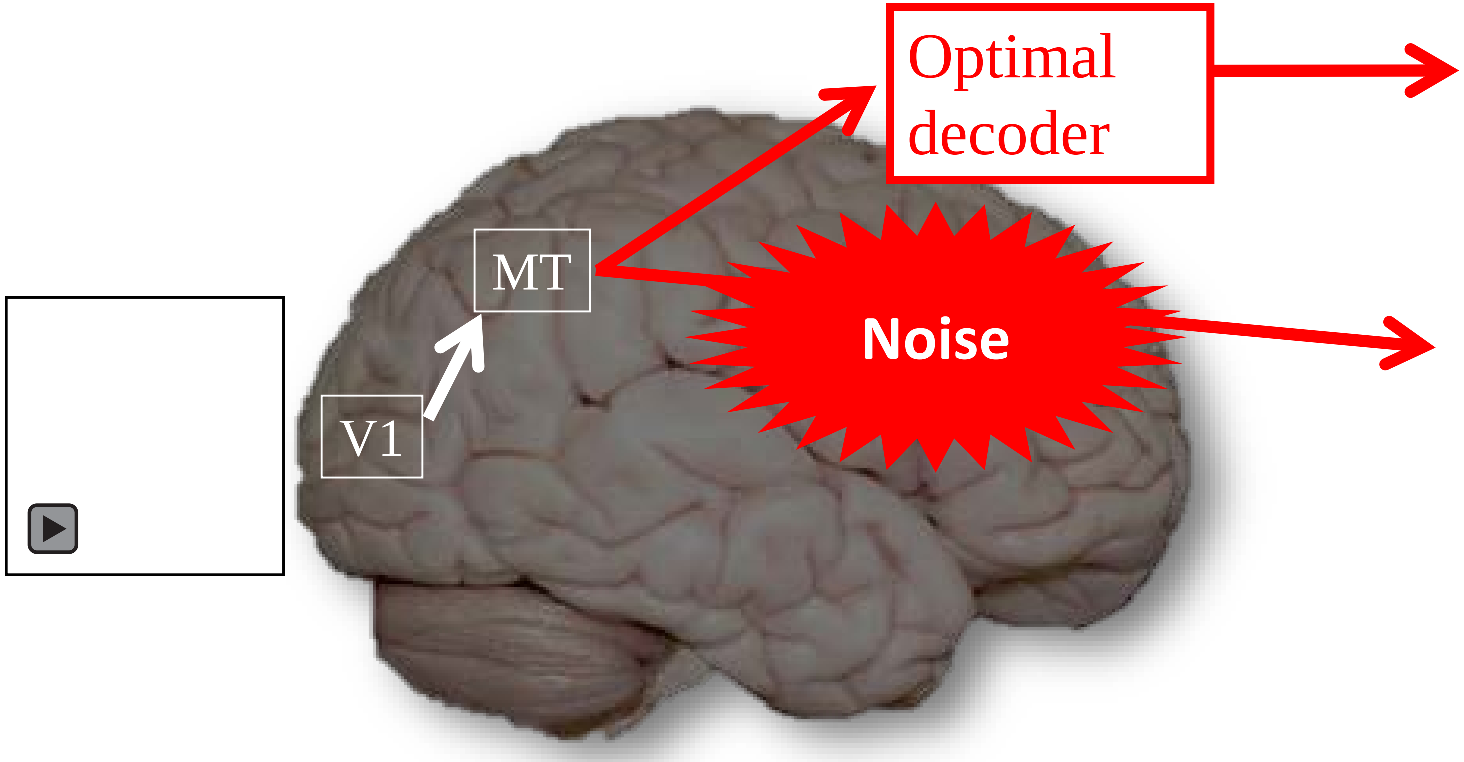

Decoder performance

V

Behavioral performance

In collaboration with Pitkow, Lakshminarasimhan, DeAngelis and Angelaki

Noise and Decoding

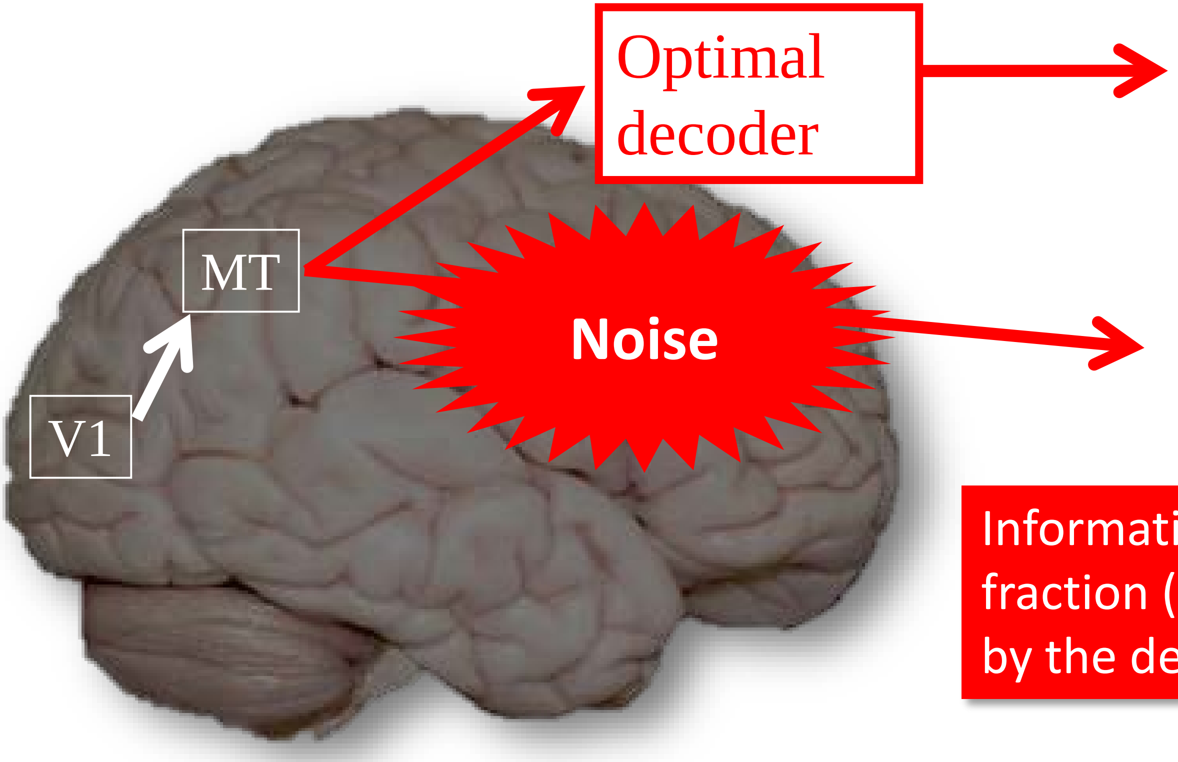

Decoder performance

>

Behavioral performance

Information in behavior should be a tiny fraction (<5%) of the information recovered by the decoder

Heading Discrimination

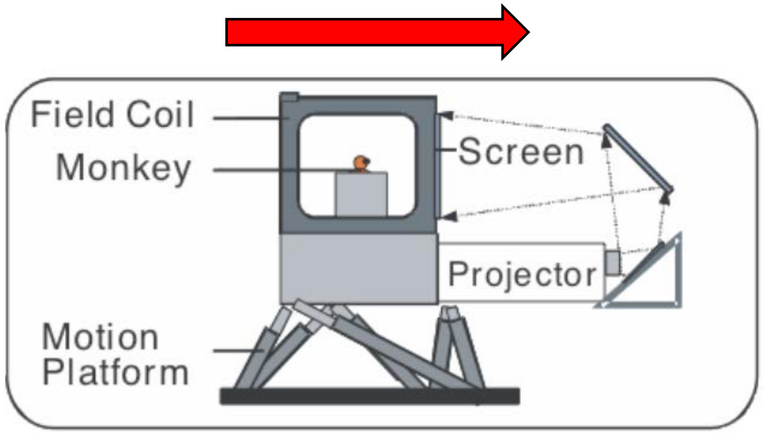

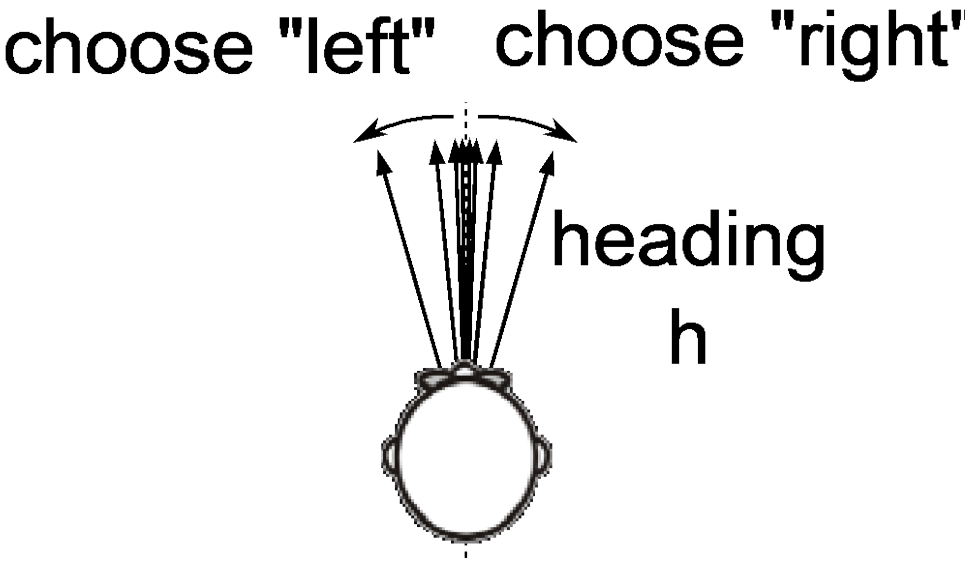

Pitkow, DeAngelis, Angelaki and Pouget. In press.

Choice Correlations

Pitkow, DeAngelis, Angelaki and Pouget. In press.

Choice Correlations

- If the read-out is optimal, and the optimal decoder is linear, choice correlations should follow

Choice Correlations

- If the read-out is optimal, and the optimal decoder is linear, choice correlations should follow

Correlations between cell responses and behavioral choices

$$C_k^{opt} = \frac{\theta}{\theta_k}$$

Pitkow, DeAngelis, Angelaki and Pouget. In press.

Choice Correlations

- If the read-out is optimal, and the optimal decoder is linear, choice correlations should follow

Correlations between cell responses and behavioral choices

Behavioral threshold

$$C_k^{opt} = \frac{\theta}{\theta_k}$$

Choice Correlations

- If the read-out is optimal, and the optimal decoder is linear, choice correlations should follow

Correlations between cell responses and behavioral choices

Behavioral threshold

$$C_k^{opt} = \frac{\theta}{\theta_k}$$

Threshold for neuron *k*

Pitkow, DeAngelis, Angelaki and Pouget. In press.

Choice Correlations

- If the read-out is optimal, and the optimal decoder is linear, choice correlations should follow

Correlations between cell responses and behavioral choices

Behavioral threshold

$$C_k^{opt} = \frac{\theta}{\theta_k}$$

Threshold for neuron k

$C_k^{Measured}$

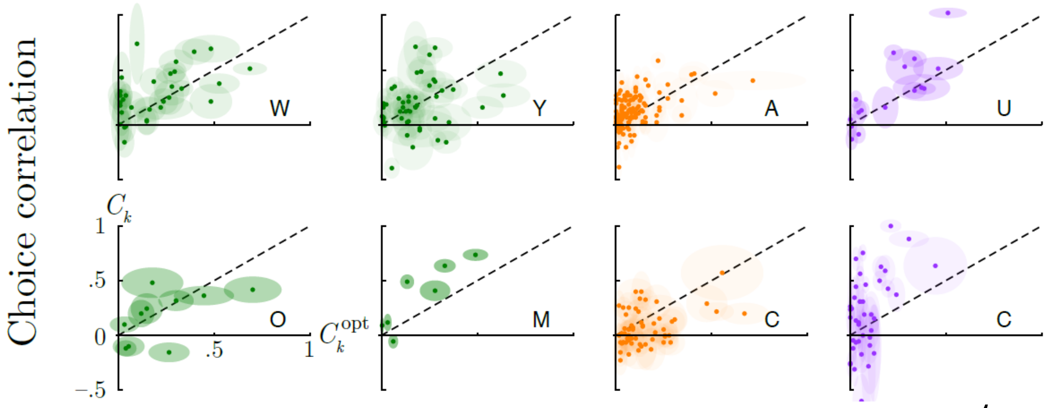

Optimal choice correlations C_k^{opt}

Pitkow, DeAngelis, Angelaki and Pouget. In press.

Choice Correlations

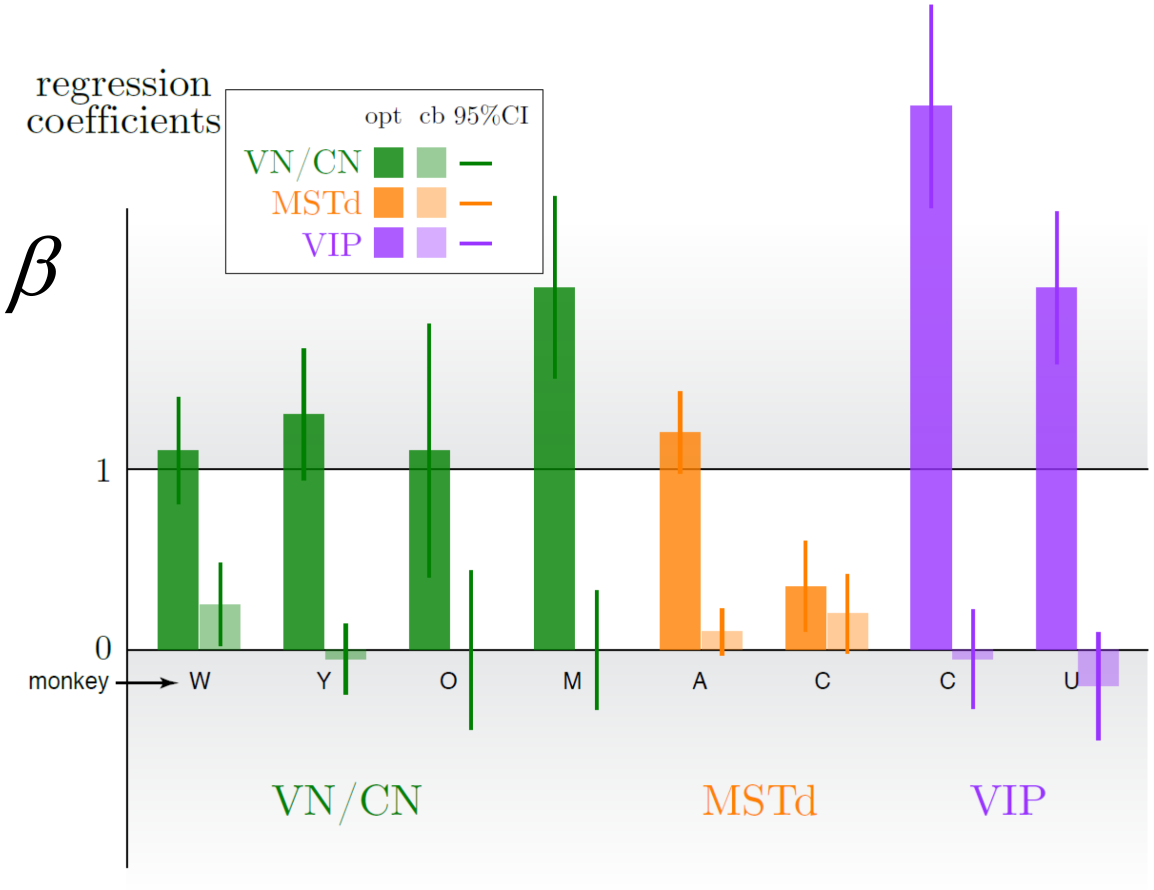

$$C_k = \beta \frac{\theta}{\theta_k}$$

Pitkow, DeAngelis, Angelaki and Pouget. In press.

Choice Correlations

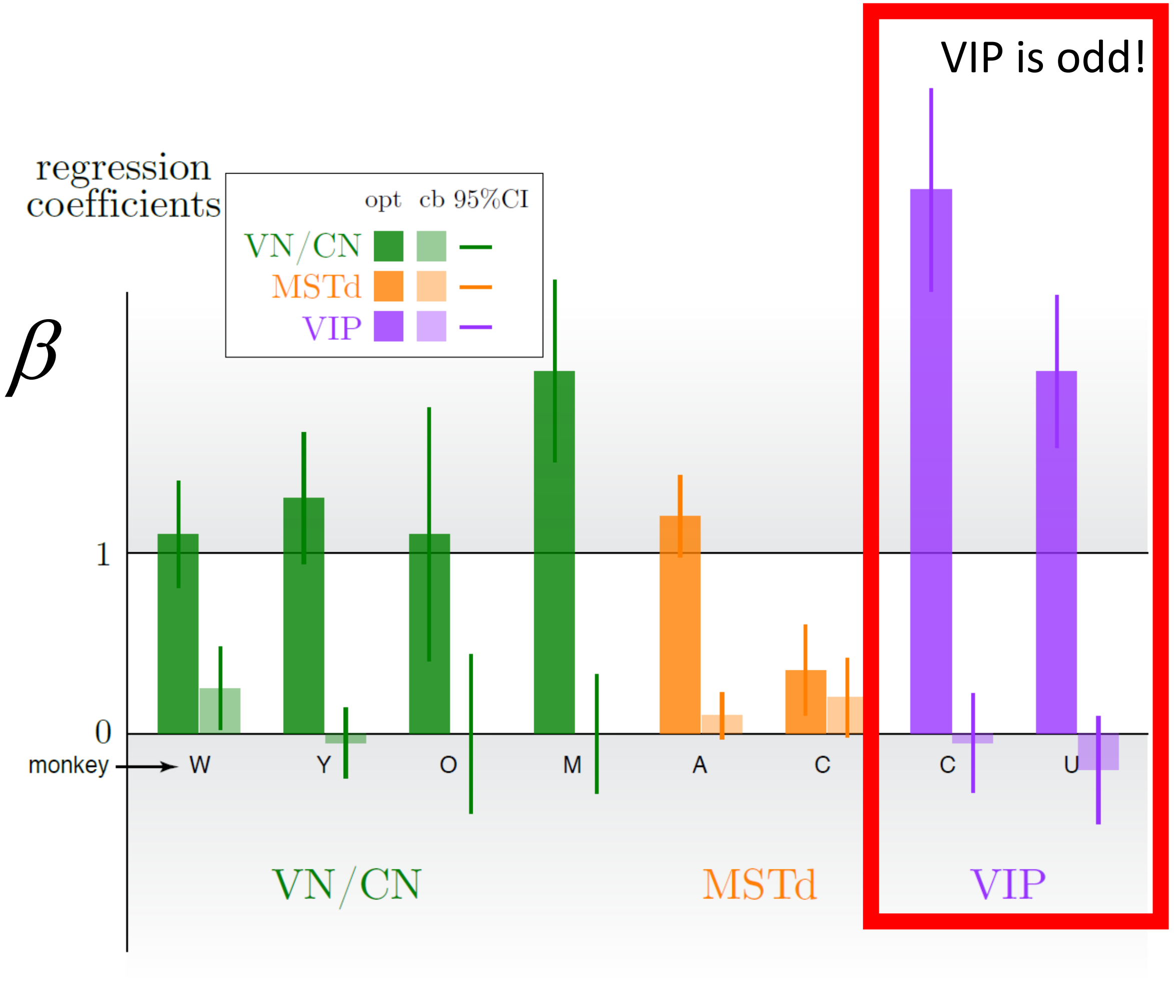

$$C_k = \beta \frac{\theta}{\theta_k}$$

Pitkow, DeAngelis, Angelaki and Pouget. In press.

Choice Correlations

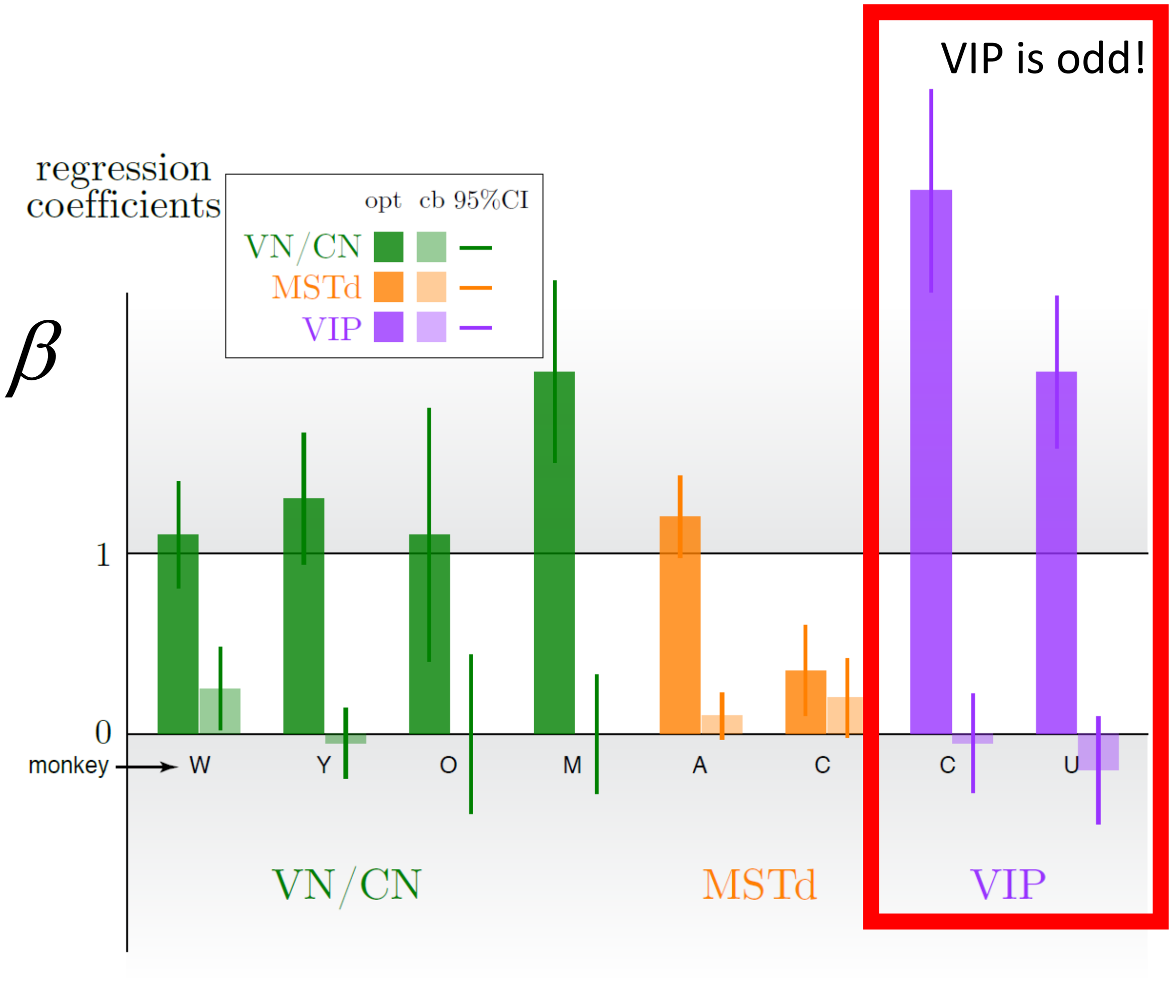

$$C_k = 3\frac{\theta}{\theta_k}$$

Choice Correlations in VIP

- In VIP, some neurons are choice correlations near 1!

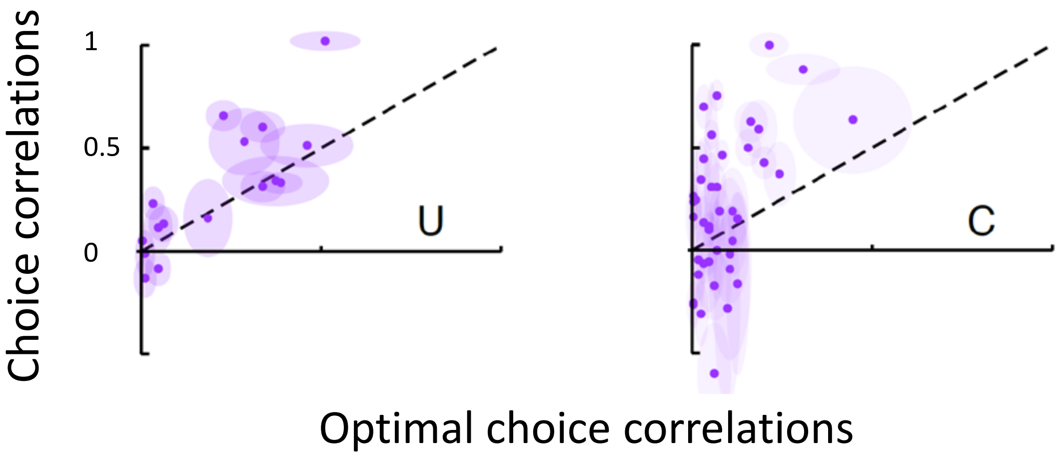

Choice Correlations in VIP

- In VIP, some neurons are choice correlations near 1!

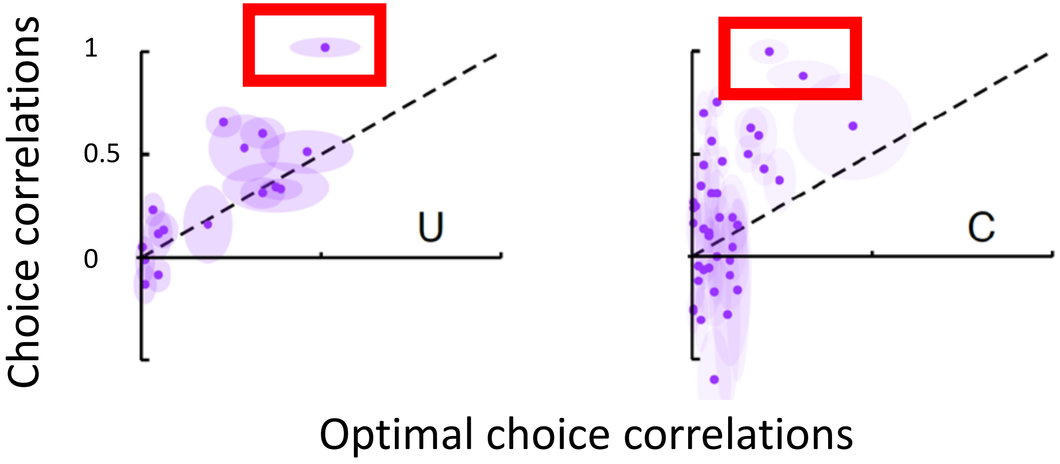

Pitkow, DeAngelis, Angelaki and Pouget. In press.

Choice Correlations

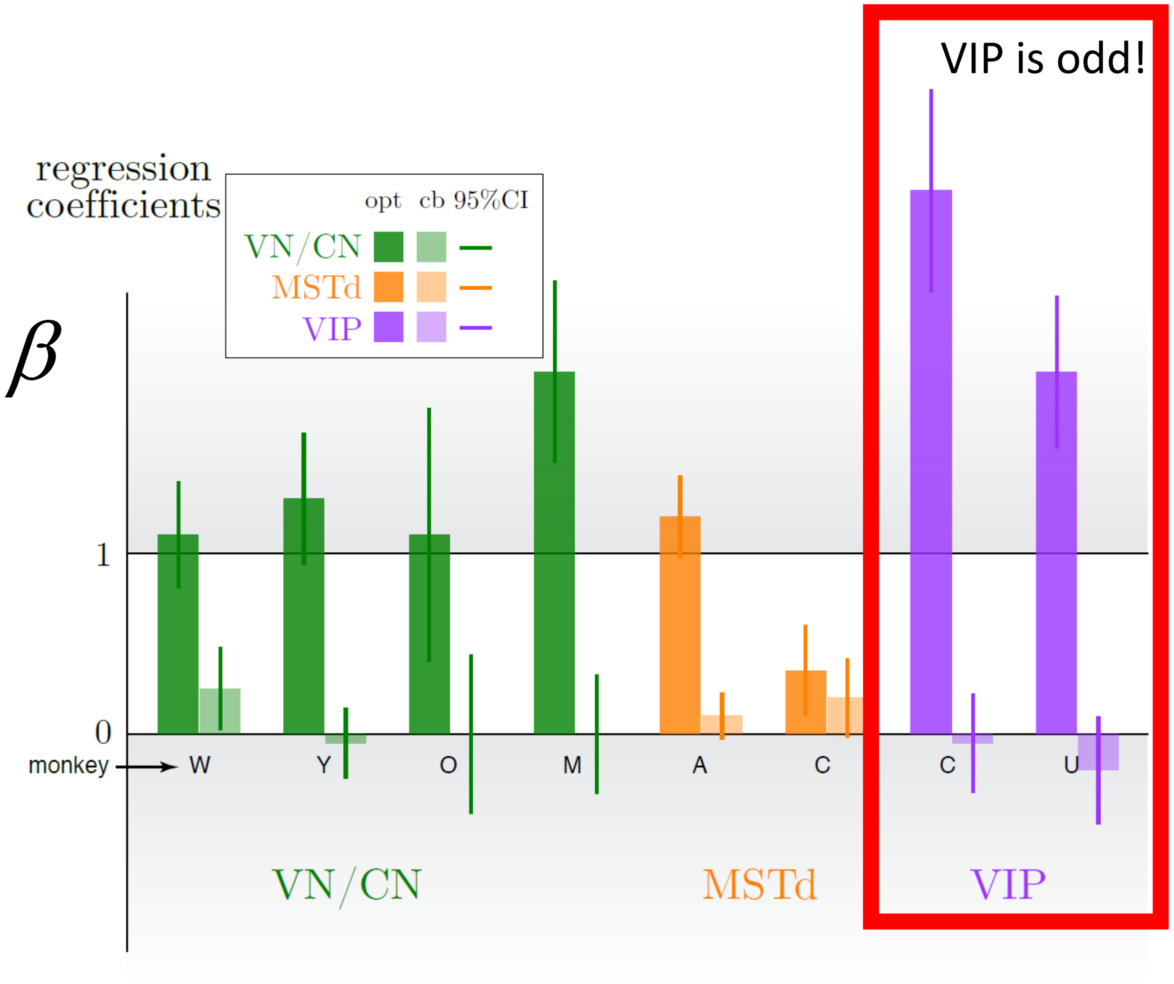

$$C_k^{opt} = 3\frac{\theta}{\theta_k}$$

Pitkow, DeAngelis, Angelaki and Pouget. In press.

Choice Correlations

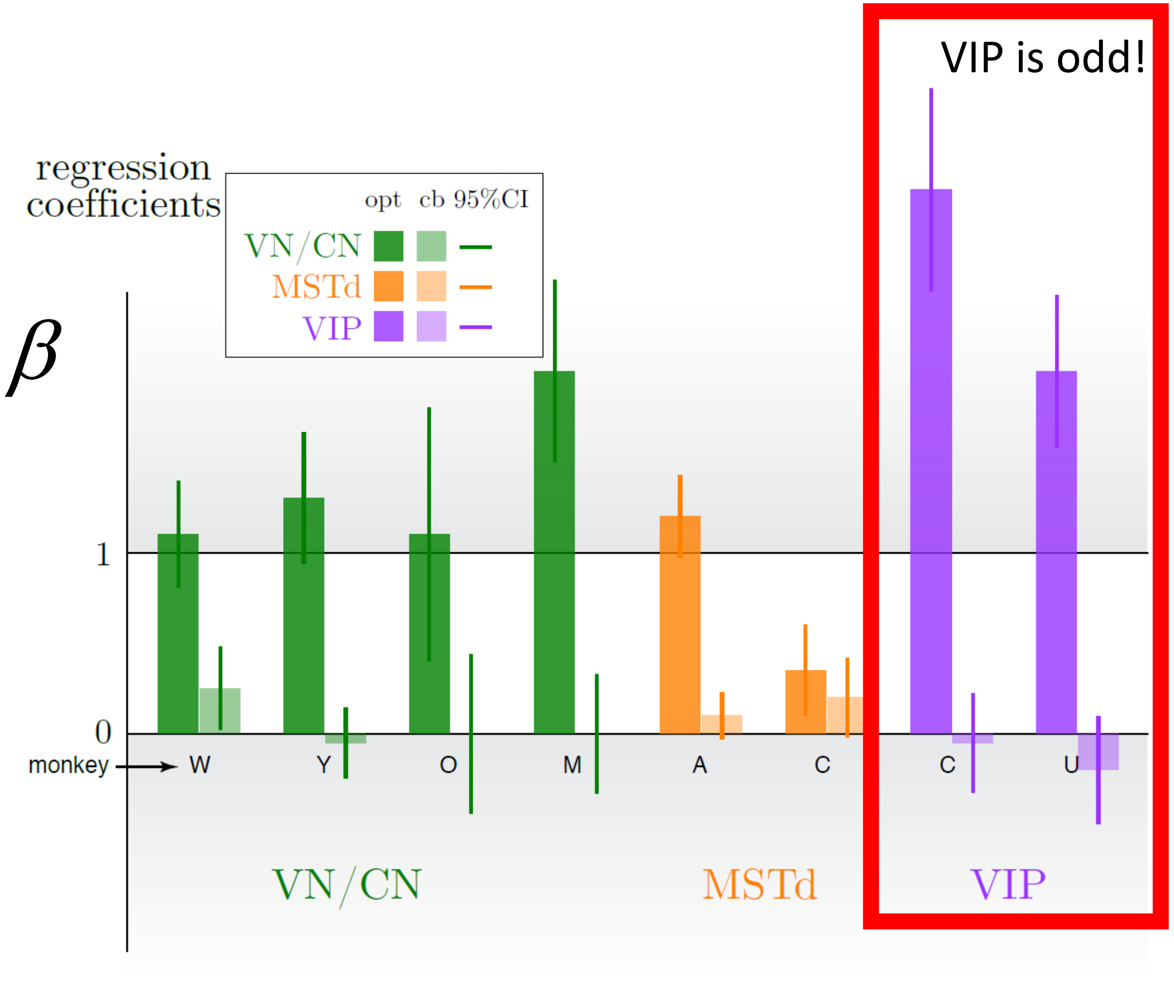

$$C_k^{opt} = 3\frac{\theta}{\theta_k}$$

Pitkow, DeAngelis, Angelaki and Pouget. In press.

Choice Correlations

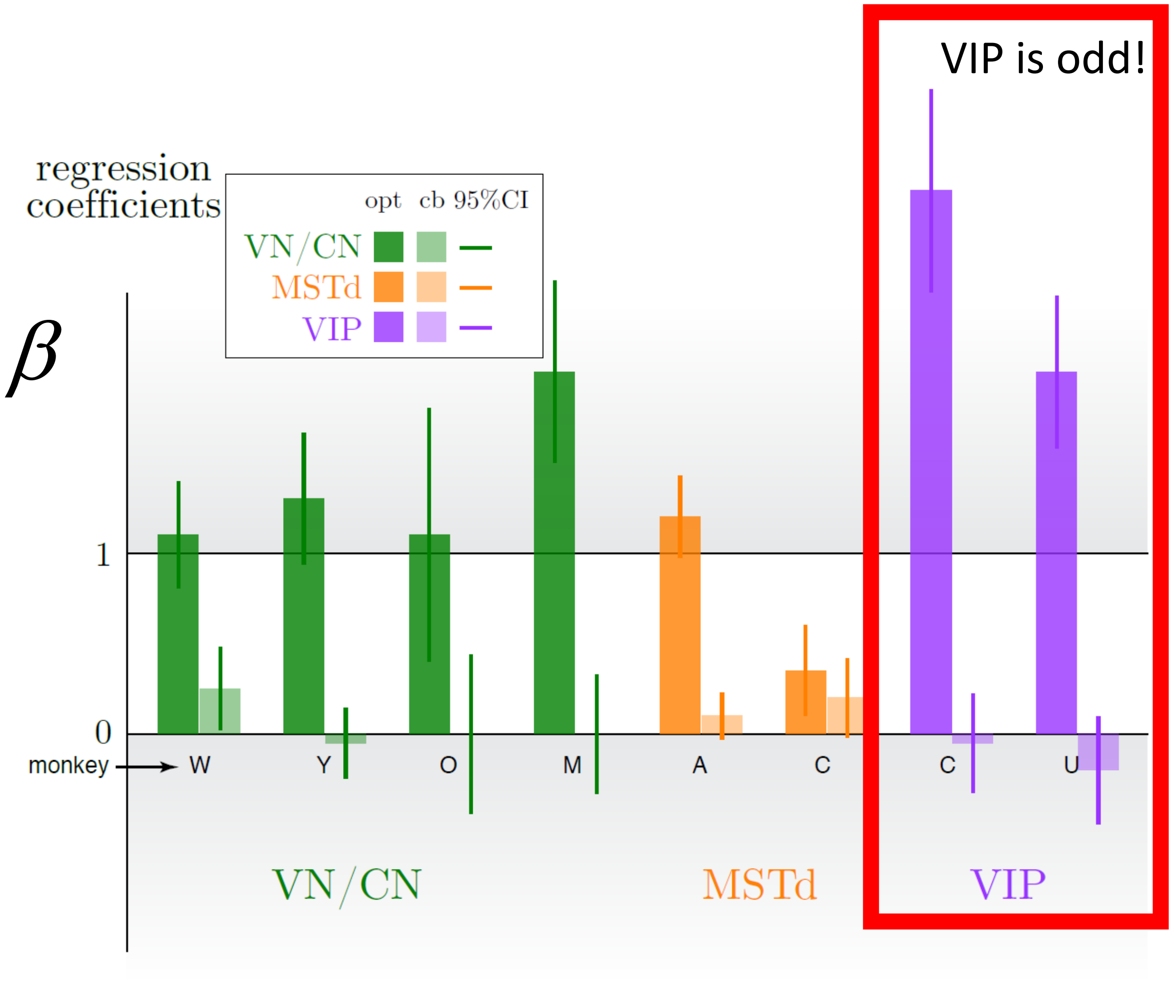

$$C_k^{opt} = 3\frac{\theta}{\theta_k}$$

Choice Correlations

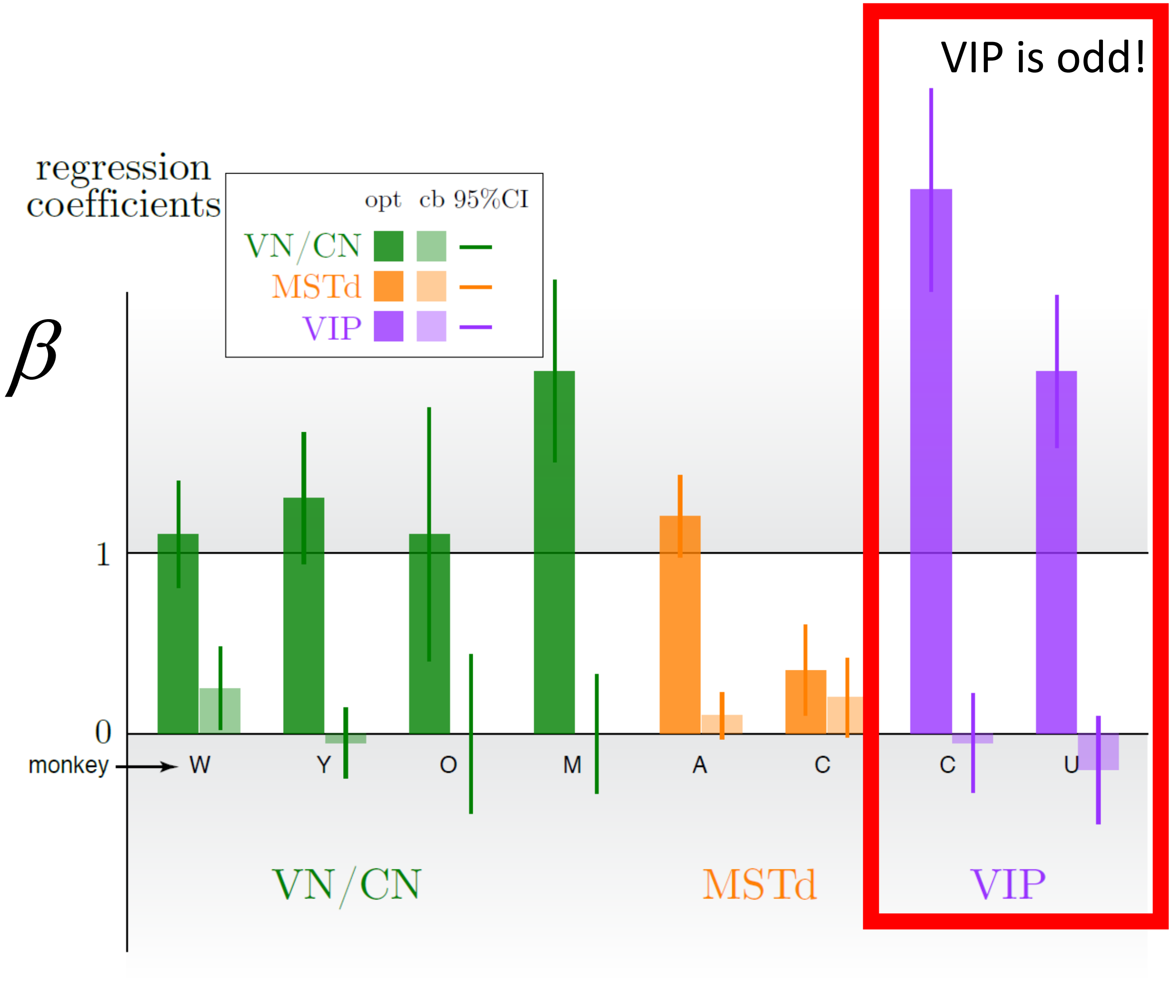

$$C_k^{opt} = 3\frac{\theta}{\theta_k}$$

This is only possible if

Pitkow, DeAngelis, Angelaki and Pouget. In press.

Choice Correlations

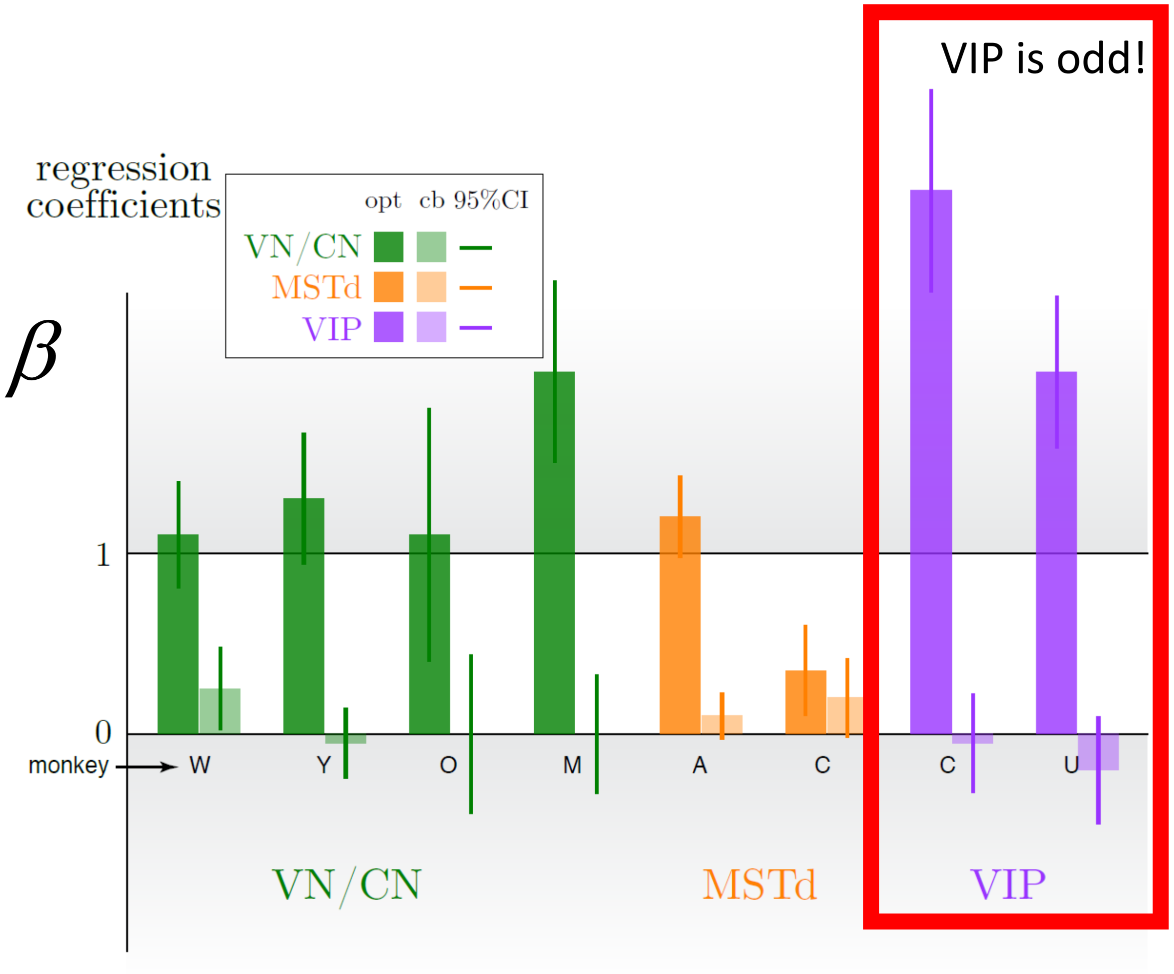

$$C_k^{opt} = 3\frac{\theta}{\theta_k}$$

This is only possible if
1- VIP is highly redundant with other areas

Choice Correlations

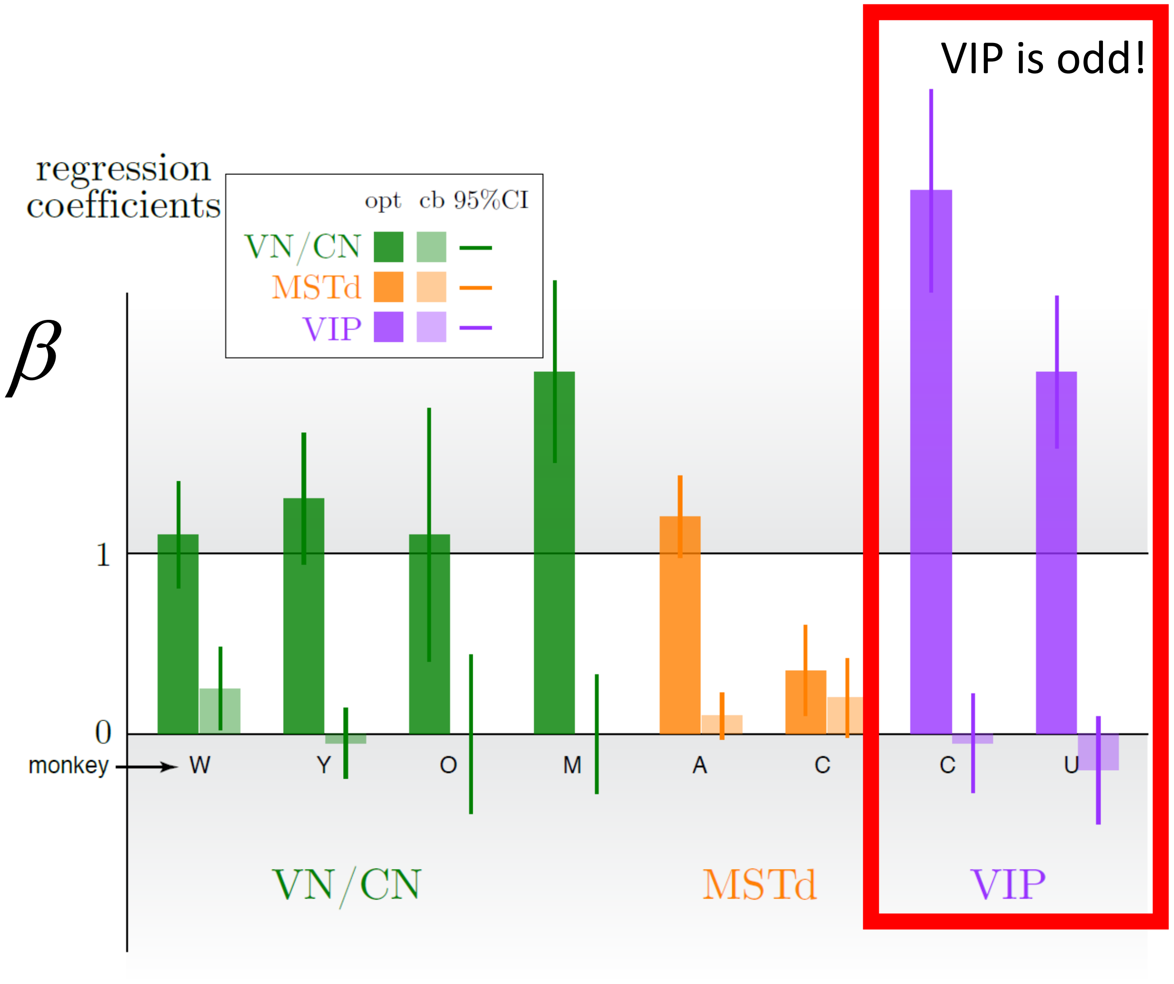

$$C_k^{opt} = 3\frac{\theta}{\theta_k}$$

This is only possible if

1- VIP is highly redundant with other areas

2- it's not read out

VIP and MSTd inactivation

- Inactivating VIP does not affect performance

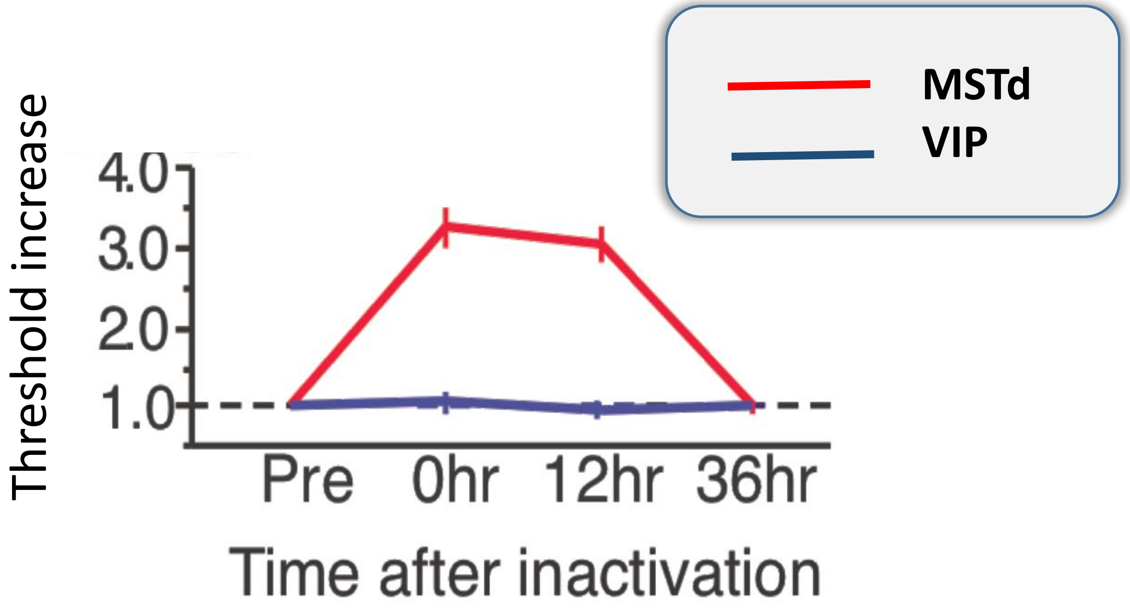

Pitkow, DeAngelis, Angelaki and Pouget. In press.

Choice Correlations

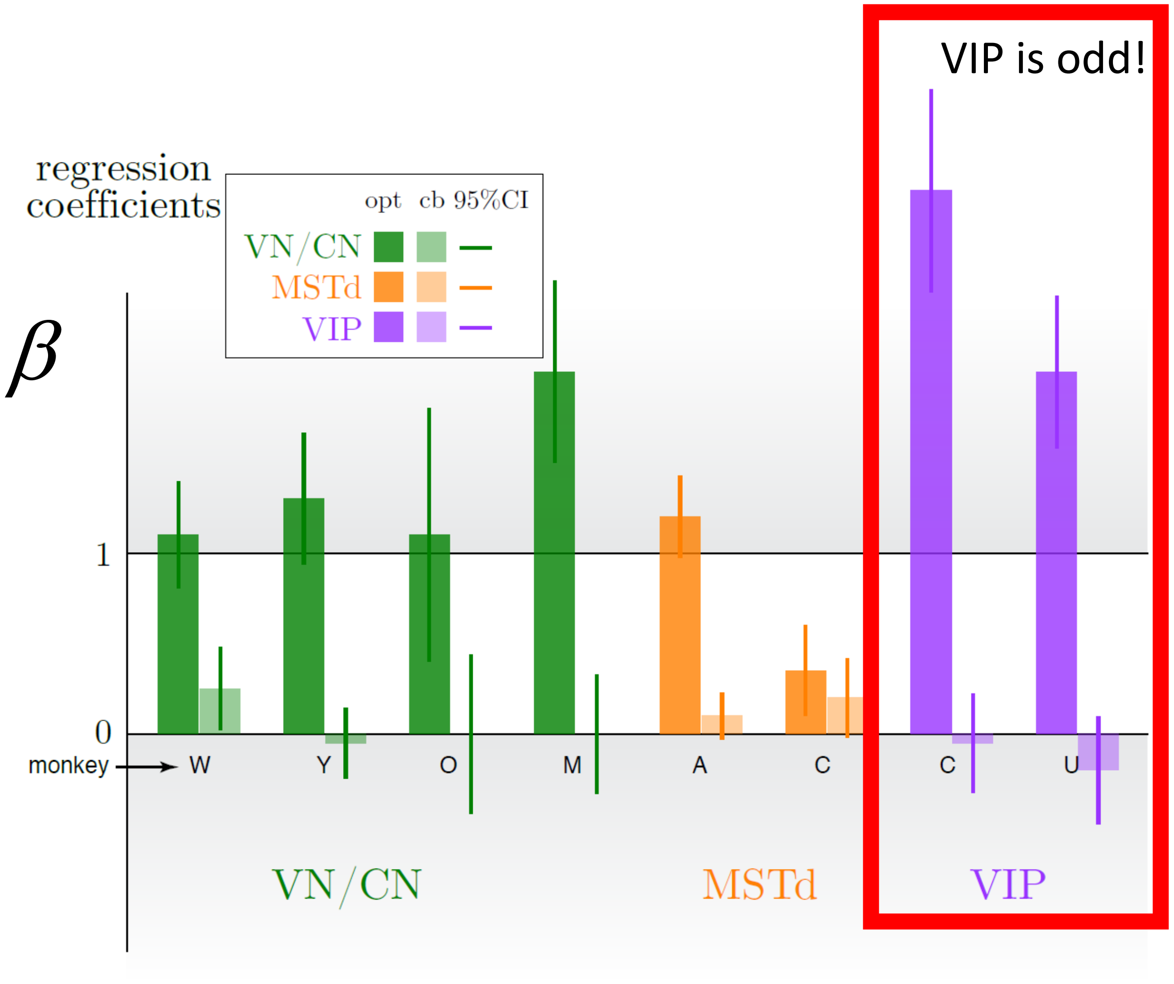

$$C_k^{opt} = 3\frac{\theta}{\theta_k}$$

This is only possible if
1- VIP is highly redundant with other areas
2- it’s not read out

Noise and Decoding

Noise and Decoding

- About 80% of the information available in MSTd and VIP is reflected in behavioral performance. (Lakshminarasimhan, Liu, Gu, Pouget, DeAngelis, Angelaki and Pitkow. Submitted)

Noise and Decoding

- About 80% of the information available in MSTd and VIP is reflected in behavioral performance. (Lakshminarasimhan, Liu, Gu, Pouget, DeAngelis, Angelaki and Pitkow. Submitted)
- Noise contribution is 20%, at most, but it could be zero. Decoding may be suboptimal.

Conclusions

Conclusions

Conclusions

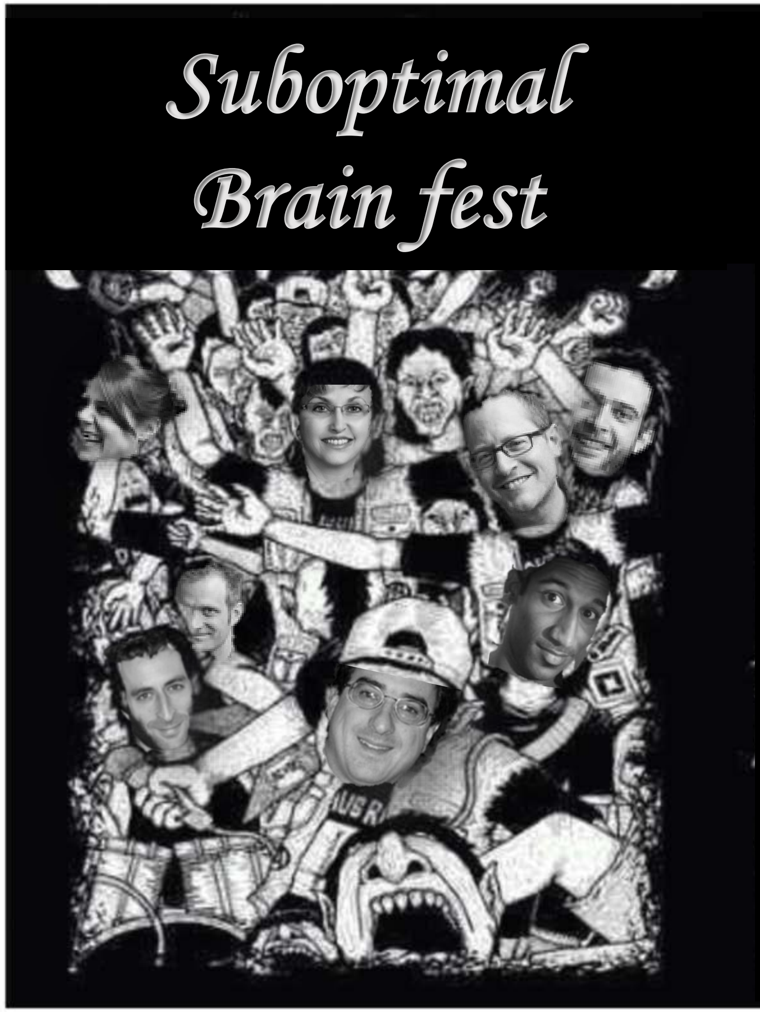

Suboptimal
Brain fest

With:

Mainen	Pitkow
Mendonca	Angelaki
Vicente	Lakshminarasimhan
DeWitt	De Angelis

Funded by:
Swiss National Funds
NSF
McDonnell Foundation
Simons Foundation