G89.2223 Perception

Bayesian Decision Theory

Laurence T. Maloney

Last: Visual Tasks

Constantine Brancusi

Size?
Shape?
Distance?

Cue combination

Last: Visual Tasks

Constantine Brancusi

Cue combination tells us what to see, not what to do.

Planning of action.

Statistical Decision Theory

Abraham Wald

John von Neumann

David Blackwell

Oskar Morgenstern

M. A. Girschick

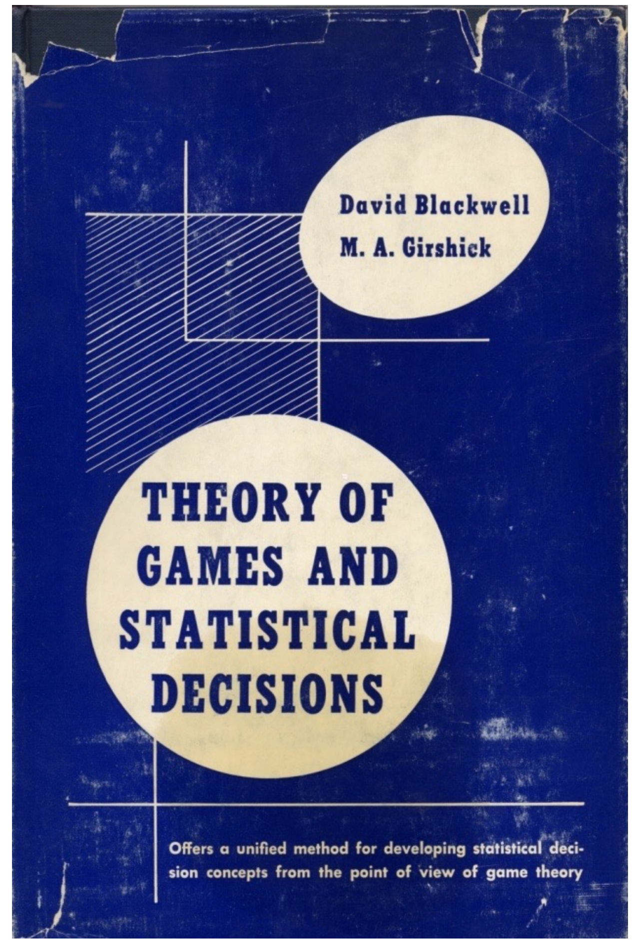

1954

Bayesian Decision Theory

Objective Information

Prior

Posterior

Likelihood

Value

Action

Internal Representations

Prior

Posterior

Likelihood

Value

Action

The organism has access to only these …

Bayesian Decision Theory

Internal Representations

Prior

Posterior

Likelihood

Value

Action

The organism has access to only these …

BDT with incorrect internal representations of probabilities, values.

The Three Elements of SDT

$$W = \{w_1, w_2, \ldots, w_m\}$$ *possible states of the world*

$$A = \{a_1, a_2, \ldots, a_p\}$$ *possible actions*

$$X = \{x_1, x_2 \ldots, x_n\}$$ *possible sensory events*

Bayesian Decision Theory (BDT)

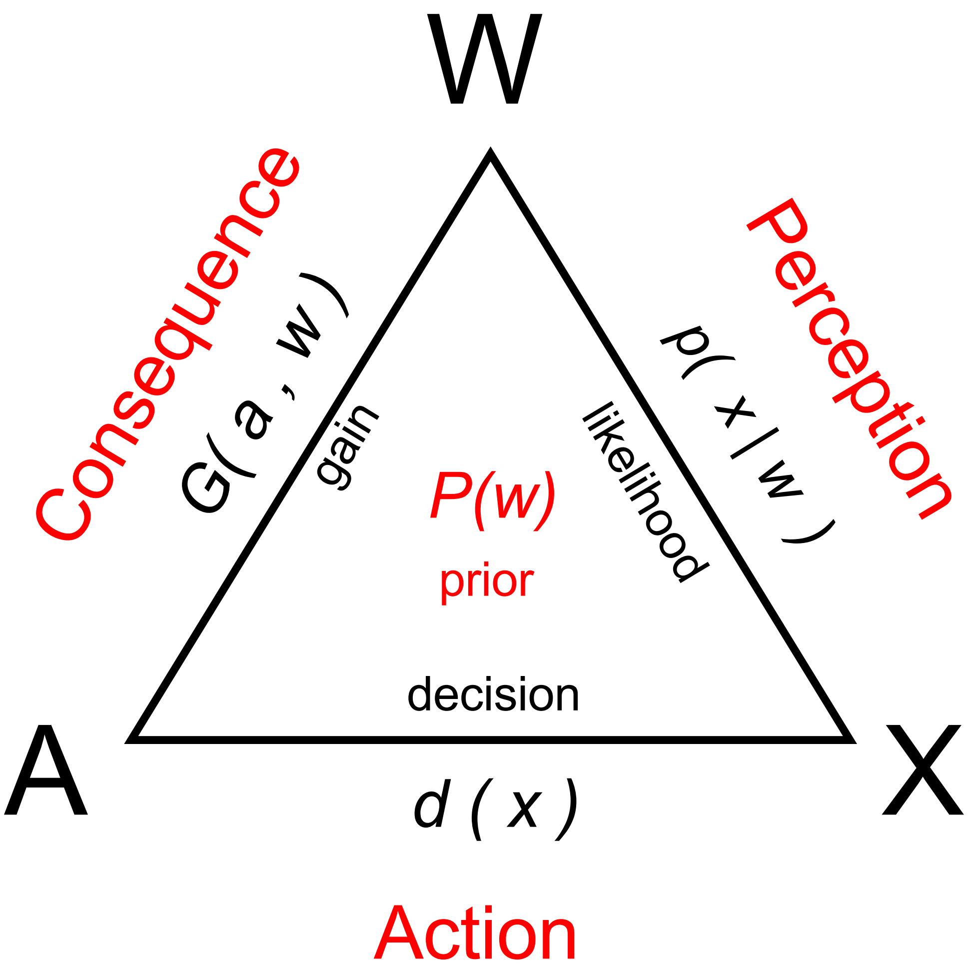

Fig. 1

Goal: select

$$d : X \rightarrow A$$

to maximize expected gain

Translated for

New Yorkers …

Time Out
New Yor
ARY 22–28, 2009 ISSUE 695
TIMEOUTNEWYORK.COM
Chances to earn $15 to $100K
Get-rich schemes from experts
Best of Chinatown
In time for the Lunar New Year
The Killers
Brandon Flowers talks underpants
Ski for cheap
7 pages of deals around NYC
57 ways to score extra cash—fast
MAKE MONEY

To help you make money….

Bayesian Decision Theory (BDT)

Maximize expected Bayes gain

$$EBG(d) \quad = \iint G\big(d(x),w\big)\; p(x \mid w)\,\pi(w)\,dx\,dw$$

by choice of a decision rule

$$\boxed{d : X \to A}$$

$$EBG(d) = \iint G(d(x), w)\, p(x \mid w)\, \pi(w)\, dx\, dw$$

Two stages: pick a random world: prior $\pi(w)$

Generate a random perception from that world:

likelihood $p(x \mid w)$

Maximize your expected gain over both random events.

$$EBG(d) = \iint G(d(x),w)\,\underbrace{\boxed{\overbrace{p(x \mid w)}^{\textit{likelihood}}\,\overbrace{\pi(w)}^{\textit{prior}}}}\,dx\,dw$$

$$EBG(d) \propto \iint G(d(x),w)\,\underbrace{\boxed{\tilde{p}(w \mid x)}}_{\textit{posterior}}\,dx\,dw$$

Bayes Theorem:

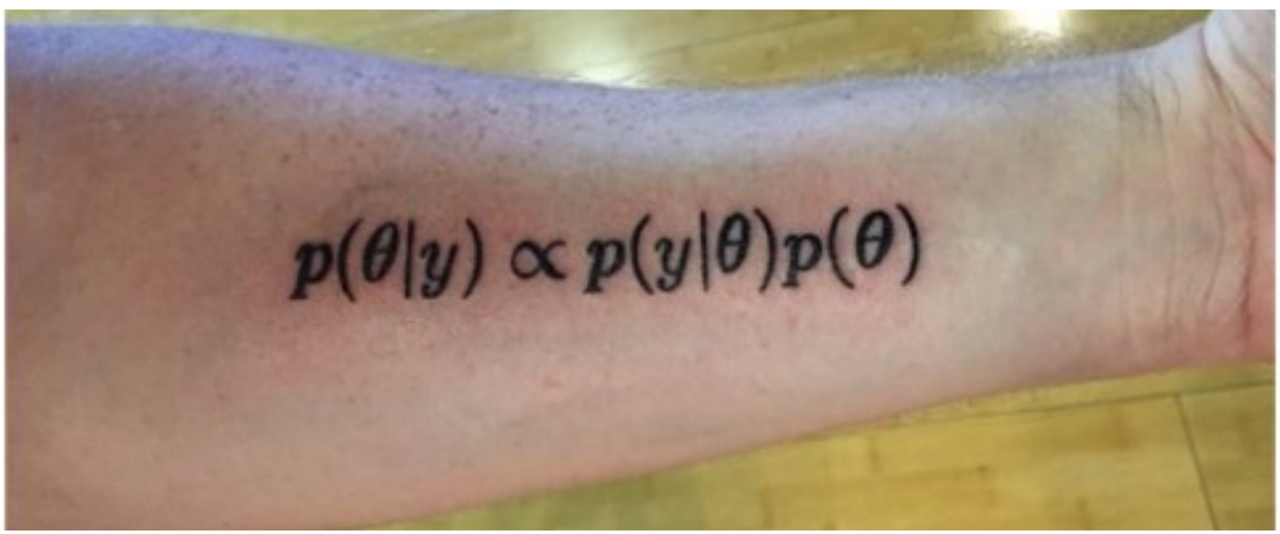

BDT, Perception and Action: A Brief History

Barlow(1950)
Geisler (1989)

$$\iint G\left(d\left(x\right),w\right)\; p\left(x \mid w\right)\; \pi\left(w\right)\, dx\, dw$$

BDT, Perception and Action: A Brief History

Barlow(1950)
Geisler (1989)

$$\iint G\big(d(x),w\big)\; p\big(x \mid w\big)\; \pi\big(w\big)\, dx\, dw$$

Knill & Richards (1996)
Yuille & Bulthoff
Tanenbaum &Griffiths
Many more

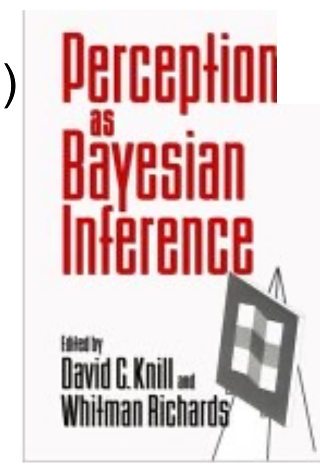

BDT, Perception and Action: A Brief History

Barlow(1950)
Geisler (1989)

$$\iint G\left(d\left(x\right),w\right)\, p\left(x\,|\,w\right)\pi\left(w\right)dx\,dw$$

Knill & Richards (1996)
Yuille & Bulthoff
Tanenbaum &Griffiths
Many more

The gain function is the problem posed by the world to the organism

Are *you* Bayesian?

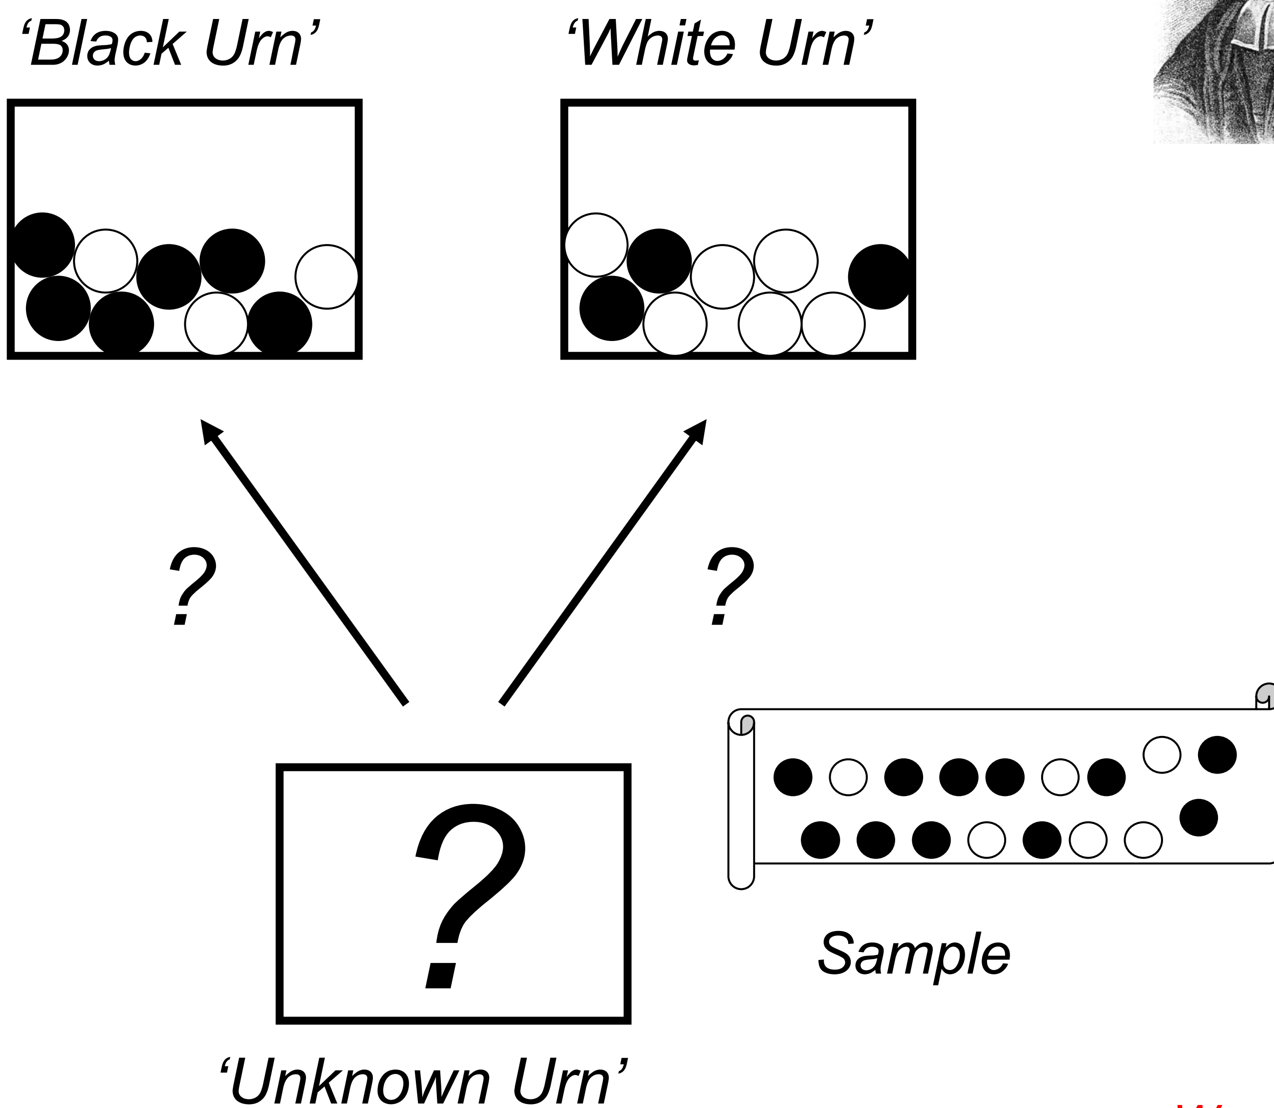

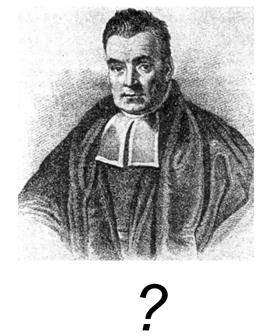

Ward Edwards

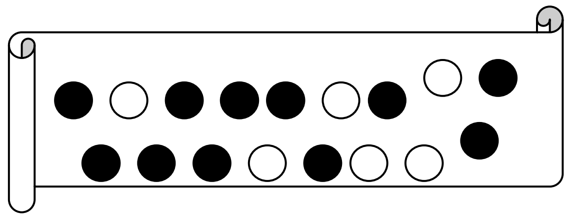

What is the probability that the unknown urn
Is the 'black' urn?

Please write down your estimate.

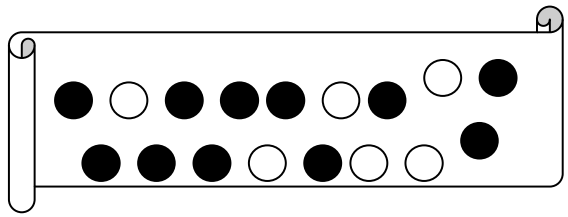

How can we estimate this probability using Bayesian methods?

Suppose we draw one black ball (b)

What is the probability now that the urn is the black one P[B|b]?

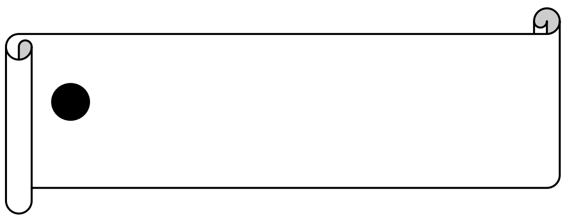

$$P[B|b] = \frac{P[b|B]P[B]}{P[b]}$$

Suppose we draw one black ball (b)

What is the probability now that the urn is the black one P[B|b]?

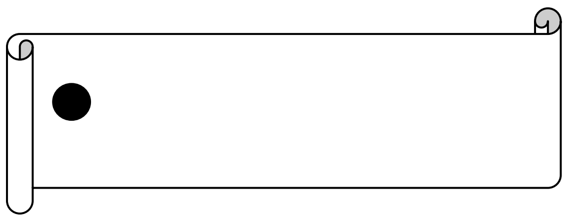

likelihood

$$P[B|b] = \frac{P[b|B]P[B]}{P[b]}$$

prior

posterior

?

Suppose we draw one black ball (b)

What is the probability now that the urn is the black one P[B|b]?

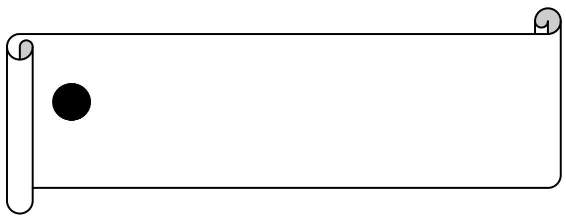

$$P[B|b] = \frac{P[b|B]P[B]}{P[b]}$$

$$P[W|b] = \frac{P[b|W]P[W]}{P[b]}$$

Suppose we draw one black ball (b)

What is the probability now that the urn is the black one P[B|b]?

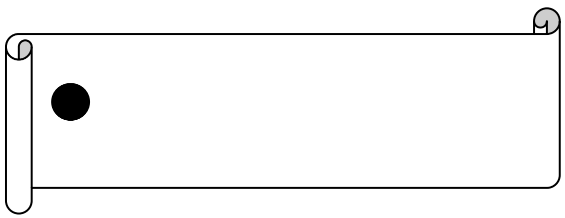

prior odds

posterior odds $\longrightarrow \dfrac{P[B|b]}{P[W|b]} = \dfrac{P[b|B]}{P[b|W]}\dfrac{P[B]}{P[W]}$

likelihood ratio

Suppose we draw one black ball (b)

What is the probability now that the urn is the black one P[B|b]?

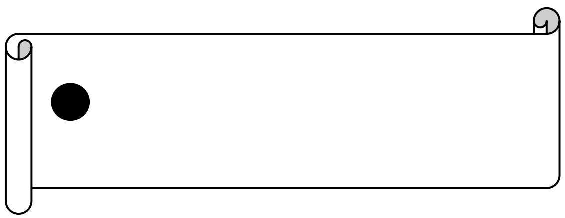

$$\frac{P[B|b]}{P[W|b]} = \frac{P[b|B]}{P[b|W]}\frac{P[B]}{P[W]}$$

$$\frac{P[B|b]}{P[W|b]} = \frac{2/3}{1/3}\ \frac{1/2}{1/2} = 2$$

Suppose we draw one black ball (b)

What is the probability now that the urn is the black one P[B|b]?

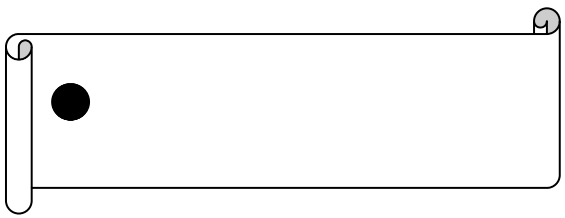

log likelihood ratio

log posterior odds

$$log\frac{P[B|b]}{P[W|b]} = log\,\frac{P[b|B]}{P[b|W]} + log\,\frac{P[B]}{P[W]}$$

log prior odds

$$\log\frac{P[B|b]}{P[W|b]} = 1 + 0$$

Next we draw a white ball (w)

What is the probability ***now*** *that the urn is the black one P[B|bw]?*

log prior after b

$$log\frac{P[B|bw]}{P[W|bw]} = log\,\frac{P[w|B]}{P[w|W]} + log\,\frac{P[B|b]}{P[W|b]}$$

$$log\,\frac{P[B|bw]}{P[W|bw]} = -1 + 1 = 0$$

Log Odds

$$\log_2 \frac{P\left[B \mid d_1 \cdots d_n\right]}{P\left[W \mid d_1 \cdots d_n\right]} = \sum_{i=1}^{n} \log_2 \frac{P\left[d_i \mid B\right]}{P\left[d_i \mid W\right]} + \log_2 \frac{P\left[B\right]}{P\left[W\right]}$$

$$d_i = b,w$$

log posterior odds *log likelihood ratio(s)* *log prior odds*

+1	*black*	*0*
-1	*white*	

Log Odds

$$\log_2 \frac{P\left[B \mid d_1 \cdots d_n\right]}{P\left[W \mid d_1 \cdots d_n\right]} = \sum_{i=1}^{n} \log_2 \frac{P\left[d_i \mid B\right]}{P\left[d_i \mid W\right]} + \log_2 \frac{P\left[B\right]}{P\left[W\right]}$$

$$d_i = b, w$$

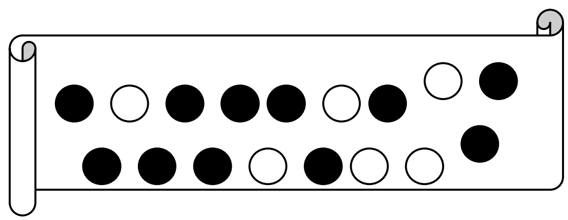

Sample

6 white
11 black
5 difference

$$32:1 \Rightarrow \frac{32}{33}$$

Probability that The urn is the black urn

Only the difference matters

Log Odds

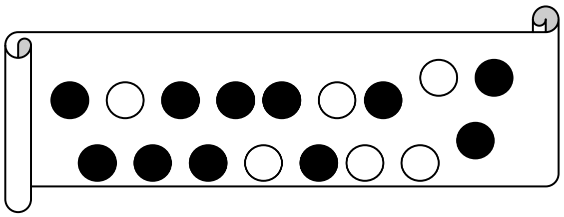

Sample

$$+5 \Rightarrow 32:1 \Rightarrow {}^{32}\!/_{33}$$

Probability that The urn is the black urn

Only the difference matters

Is that your intuition?

B	***W***
5	0
11	6
95	90

Are *you* Bayesian?

Probably not.

People tend to pick odds closer to 1:1 than the correct odds. This error is an example of human tendency to distort probability.

Conservatism [Ward Edwards]

BDT in Action:

Signal Detection Theory

Origin of SDT: WW2 radar operator

- Are the blobs enemy aircraft? Or just noise (e.g. clouds)?

- Decision has consequences:
  - If you miss an aircraft, people might get killed
  - If you mistake “noise” for an aircraft, fuel, time & resources are wasted

Radar screen

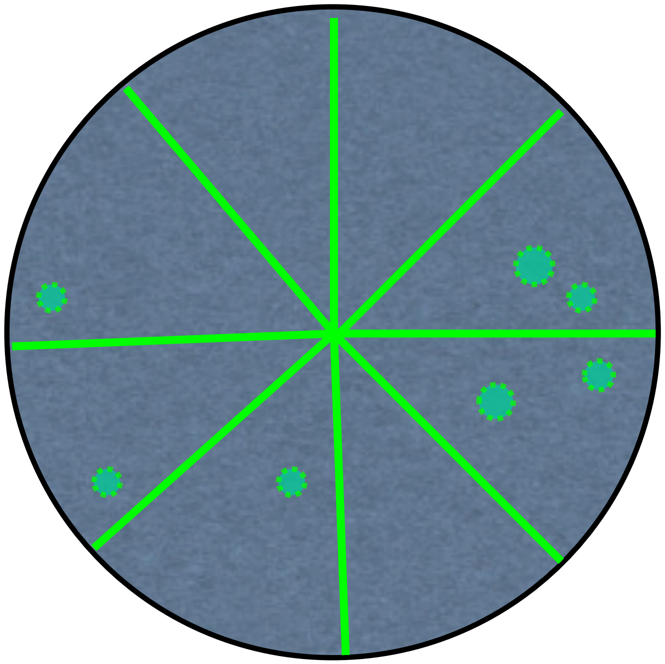

$W = \{S, \bar{S}\}$ *World States*

$A = \{Y, N\}$ *Actions*

$X = (-\infty, \infty)$ *Stimulus Intensity*

Decision outcomes

SIGNAL: are the blobs real enemy aircraft?

DECISION: should you alert the air force?

	S	$\bar{S}$
Y	**Hit**	**False alarm**
N	**Miss**	**Correct reject**

$$W \quad = \left\{ S, \bar{S} \right\}$$

$$A \quad = \left\{ Y, N \right\} \qquad p\left[X \mid S \right], p\left[X \mid \bar{S} \right] \quad \textit{likelihood}$$

$$X \quad = \left(-\infty, \infty \right)$$

$$\begin{array}{cc} & \begin{array}{cc} S & \bar{S} \end{array} \\ \begin{array}{c} Y \\ N \end{array} & \begin{bmatrix} V_{YS} & -V_{Y\bar{S}} \\ -V_{NS} & V_{N\bar{S}} \end{bmatrix} \end{array} \qquad \textit{gain}$$

Likelihood Functions

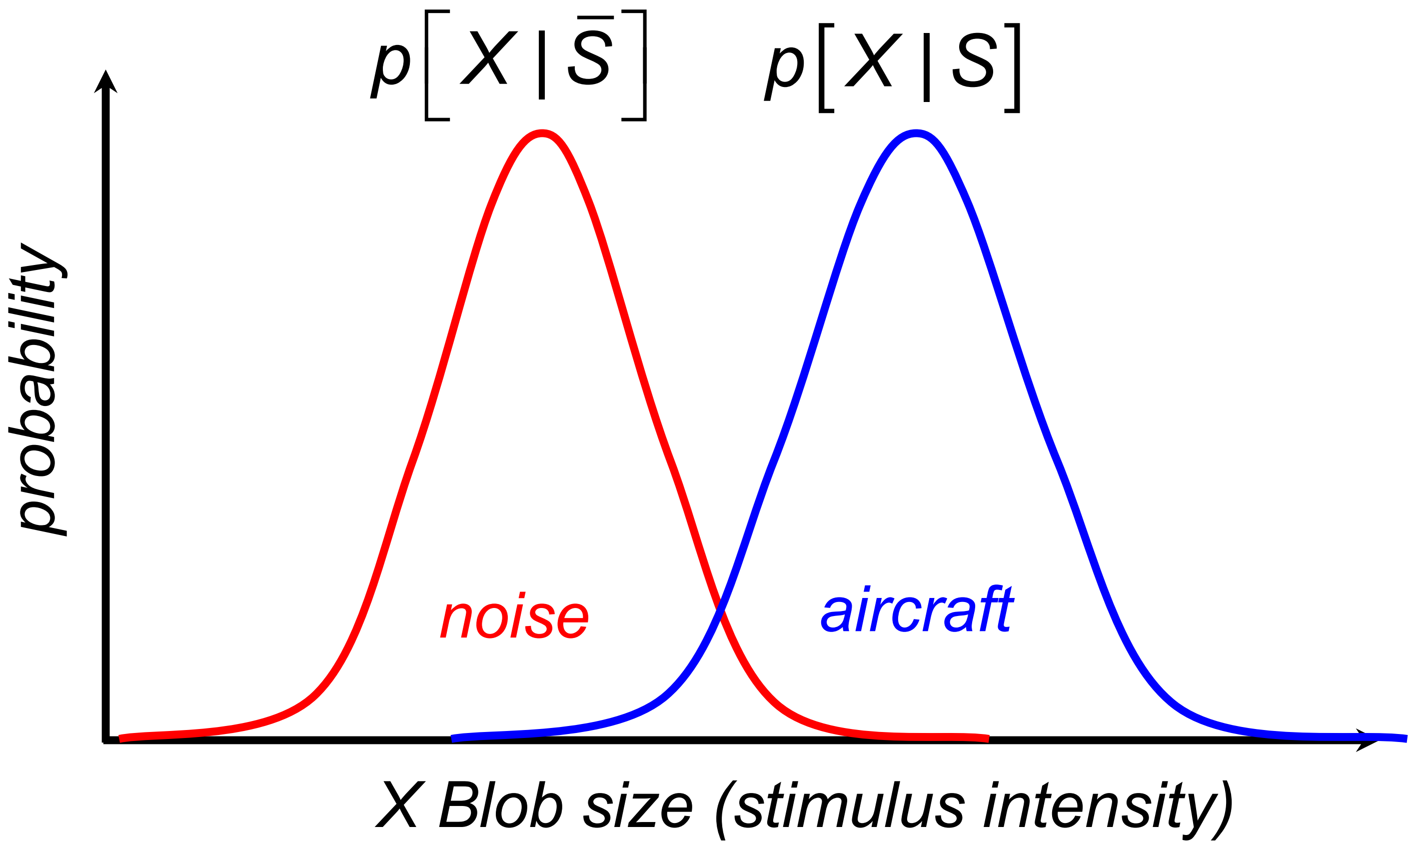

Computing Expected Bayes Gain

$$EBG\left(Y\,|\,X\right) \quad = V_{YS}\,p\left[X\,|\,S\right]\pi\left(S\right) - V_{Y\bar{S}}\,p\left[X\,|\,\bar{S}\right]\pi\left(\bar{S}\right)$$

$$EBG\left(N\,|\,X\right) \quad = -V_{NS}\,p\left[X\,|\,S\right]\pi\left(S\right) + V_{N\bar{S}}\,p\left[X\,|\,\bar{S}\right]\pi\left(\bar{S}\right)$$

$$RULE: \quad \text{"}Say\ Y\text{"} \quad \Leftrightarrow \quad EBG\left(Y\,|\,X\right) > EBG\left(N\,|\,X\right)$$

$$\frac{p\left[X\,|\,S\right]}{p\left[X\,|\,\bar{S}\right]} \quad > \quad \frac{V_{Y\bar{S}} + V_{N\bar{S}}}{V_{YS} + V_{NS}} \times \frac{\pi\left(\bar{S}\right)}{\pi\left(S\right)}$$

likelihood ratio

$$\frac{p\left[X \mid S\right]}{p\left[X \mid \bar{S}\right]} > \frac{V_{Y\bar{S}} + V_{N\bar{S}}}{V_{YS} + V_{NS}} \times \frac{\pi\left(\bar{S}\right)}{\pi\left(S\right)}$$

criterion

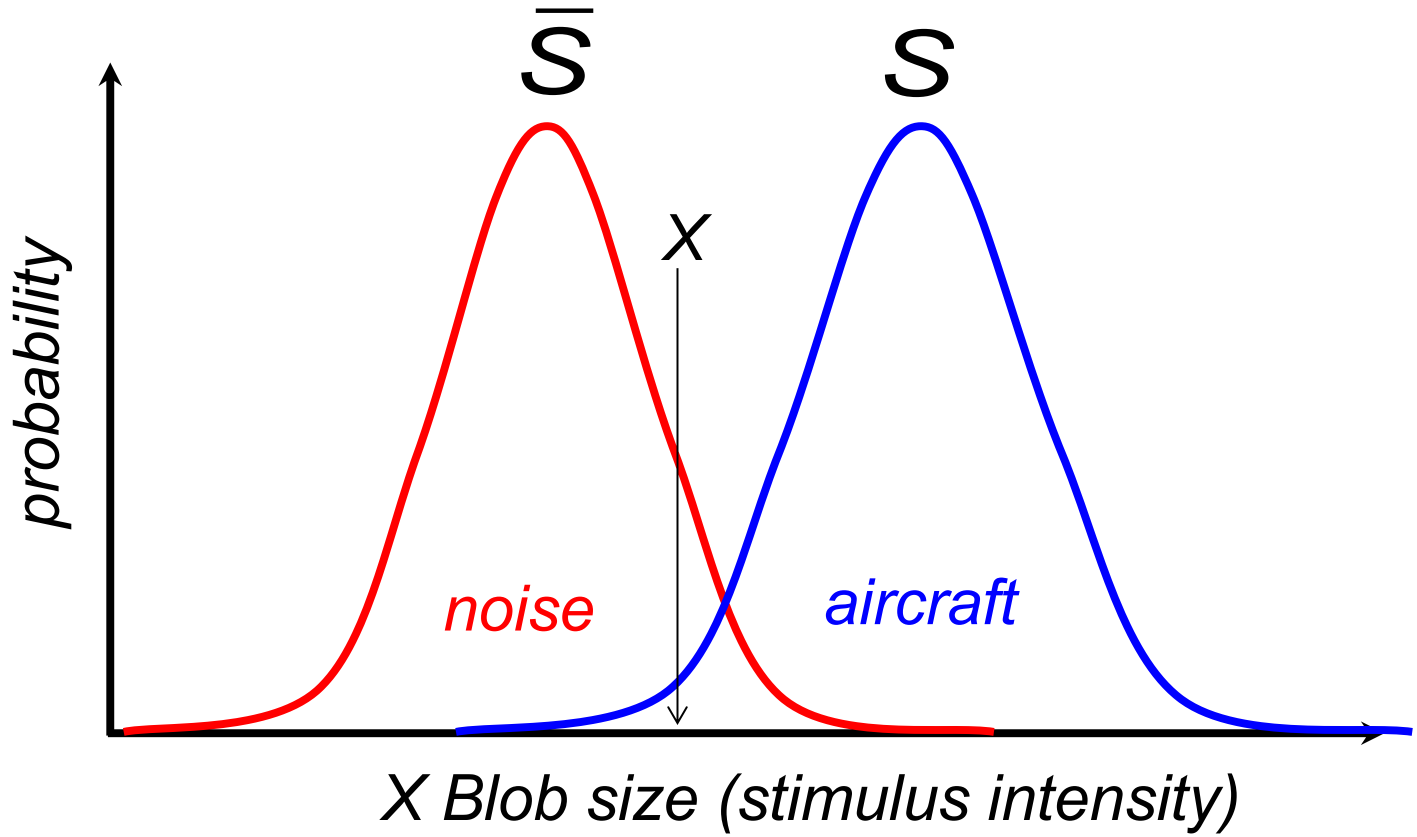

How should we set criterion?

likelihood ratio *prior odds*

$$Y \Leftrightarrow \frac{p[X \mid S]}{p[X \mid \bar{S}]} > \frac{\pi(\bar{S})}{\pi(S)}\left[\frac{V_{YS}+V_{NS}}{V_{Y\bar{S}}+V_{N\bar{S}}}\right]$$

posterior odds

Bayes Theorem

$$Y \Leftrightarrow \frac{p[S \mid X]}{p[\bar{S} \mid X]} > \frac{V_{YS}+V_{NS}}{V_{Y\bar{S}}+V_{N\bar{S}}}$$

Given the stimulus X are the posterior odds large enough to motivate a Yes response?

How should we set criterion?

$$Y \quad \Leftrightarrow \quad \log \frac{p[X \mid S]}{p[X \mid \bar{S}]} \quad > \quad \log \frac{\pi(\bar{S})}{\pi(S)} \quad + \quad \log\left[\frac{V_{YS} + V_{NS}}{V_{Y\bar{S}} + V_{N\bar{S}}}\right]$$

$$Y \quad \Leftrightarrow \quad \log \frac{p[X \mid S]}{p[X \mid \bar{S}]} \quad > \log \beta$$

log likelihood ratio

Compare the LLR to a criterion, log β

This is equivalent to $X > c$ for the right choice of c.

How well do people do?

We can estimate log beta [optimal] and compare it to the log beta people choose.

Vary prior odds
Estimate criterion β

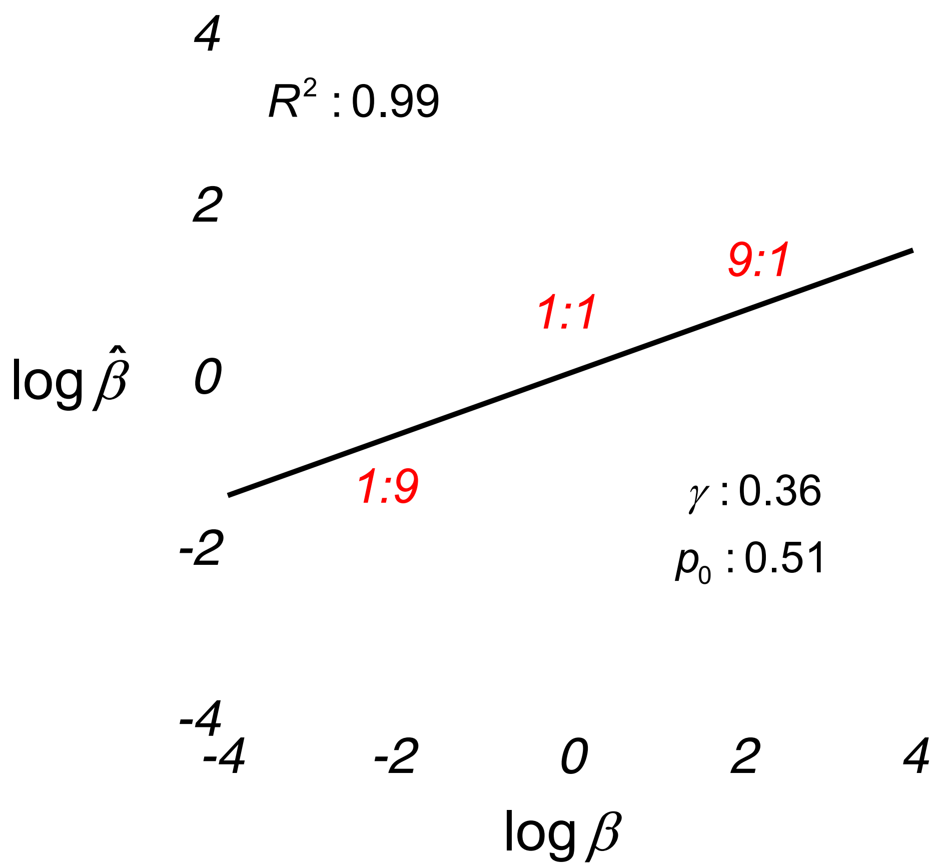

Tanner, Swets, & Green (1956)

Conservatism

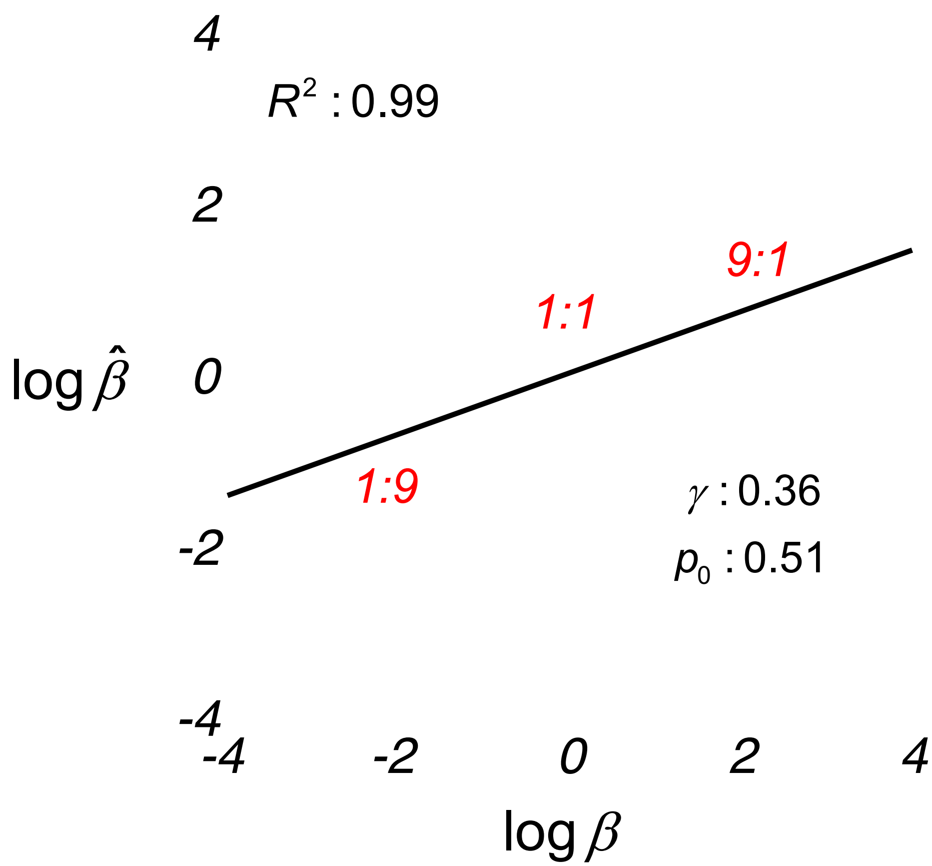

Tanner, Swets, & Green (1956)

Themes

Managing uncertainty to maximize gain
is the central task of a biological organism.

The use of explicit cost and rewards allow us to probe a much wider range of behavior than previously explored (Trommershäuser et al, 2003, 2008)

We can test SDT/BDT as a framework for modeling perception and action (Maloney & Mamassian, 2009)

Aside

The gain function is the organism's link to the environment.

It represents a problem, posed by the environment, a problem that can rapidly change.

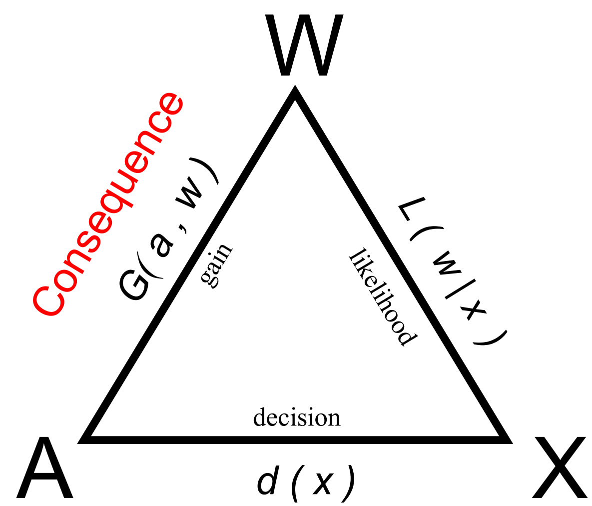

Only the luckiest organism can choose its gain function.

Planning Actions
Maximizing
Expected Gain

One example

Experimental Task

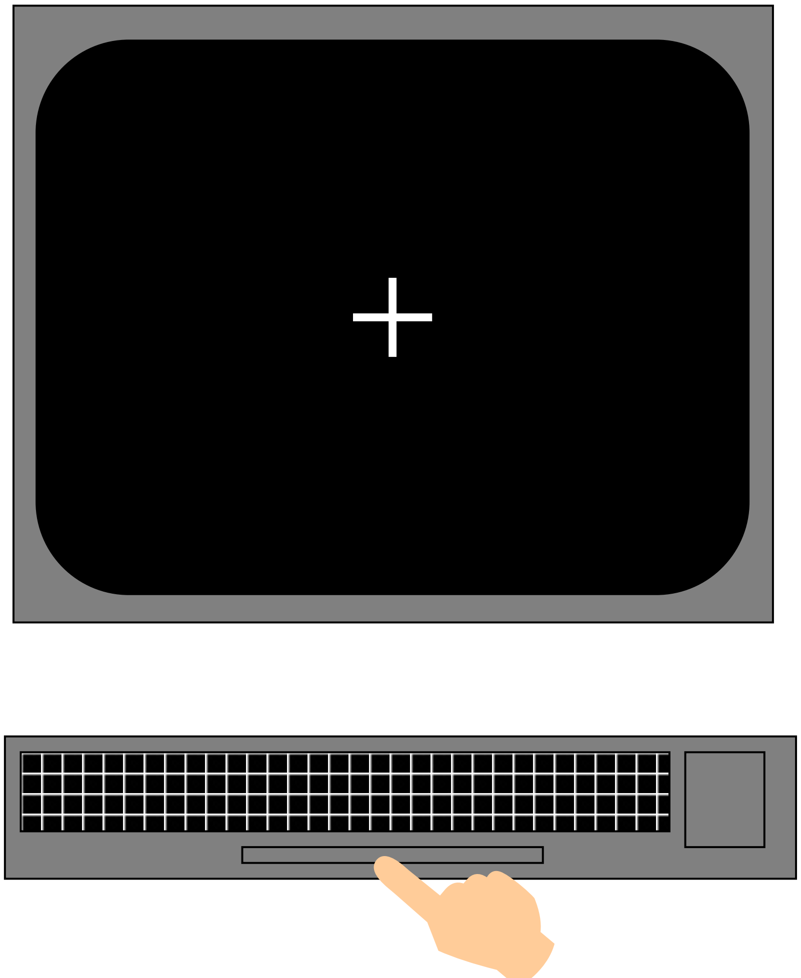

Start of trial:
display of fixation cross (1.5 s)

Experimental Task

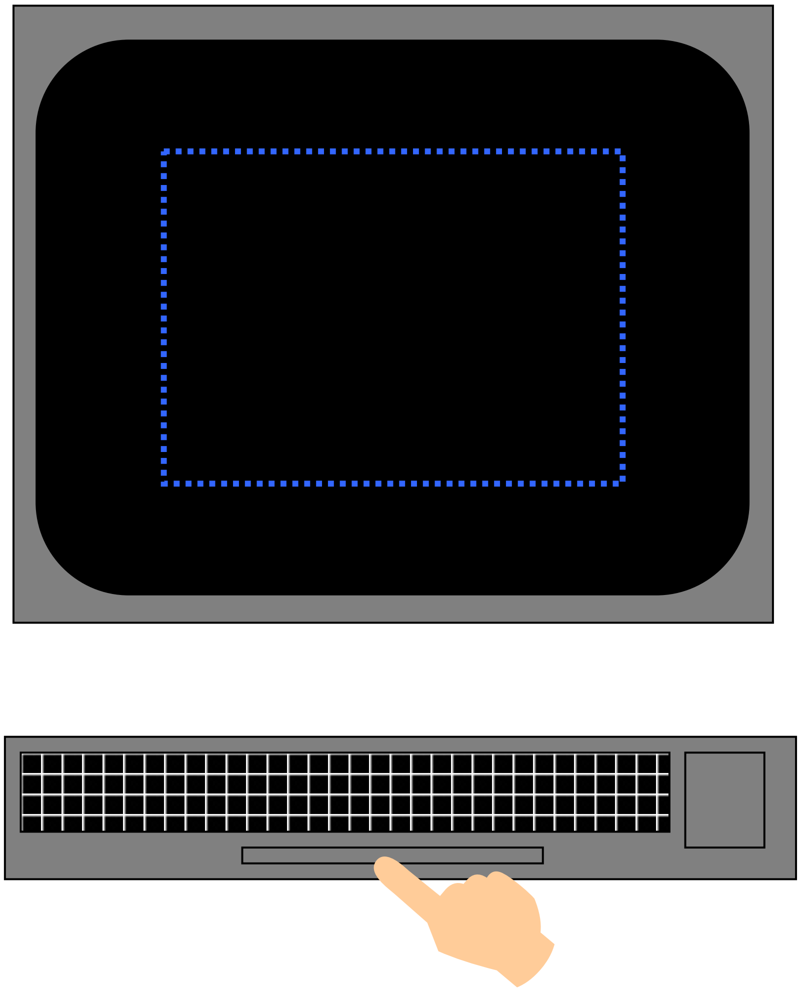

Display of response area,
500 ms before
target onset
(114.2 mm x 80.6 mm)

Experimental Task

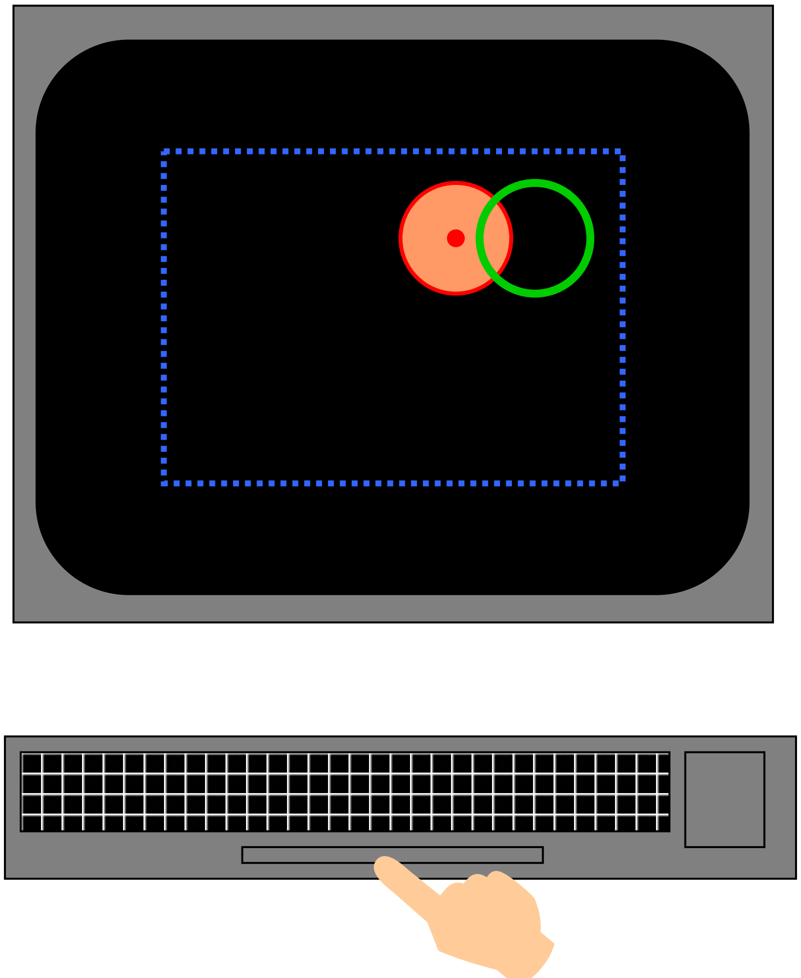

Target display (700 ms)

Experimental Task

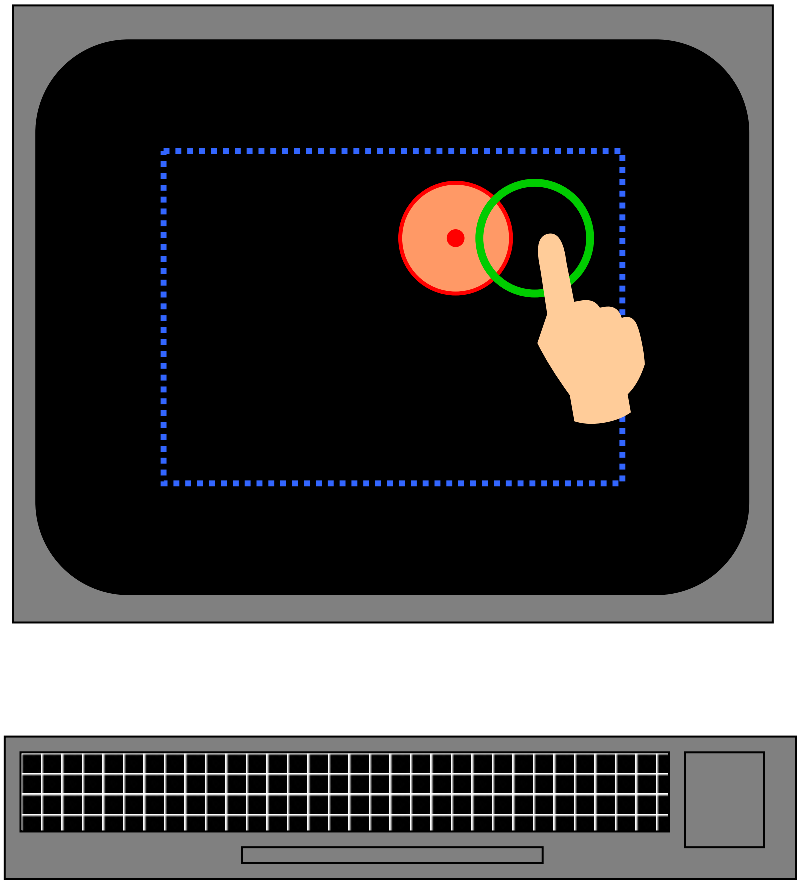

Experimental Task

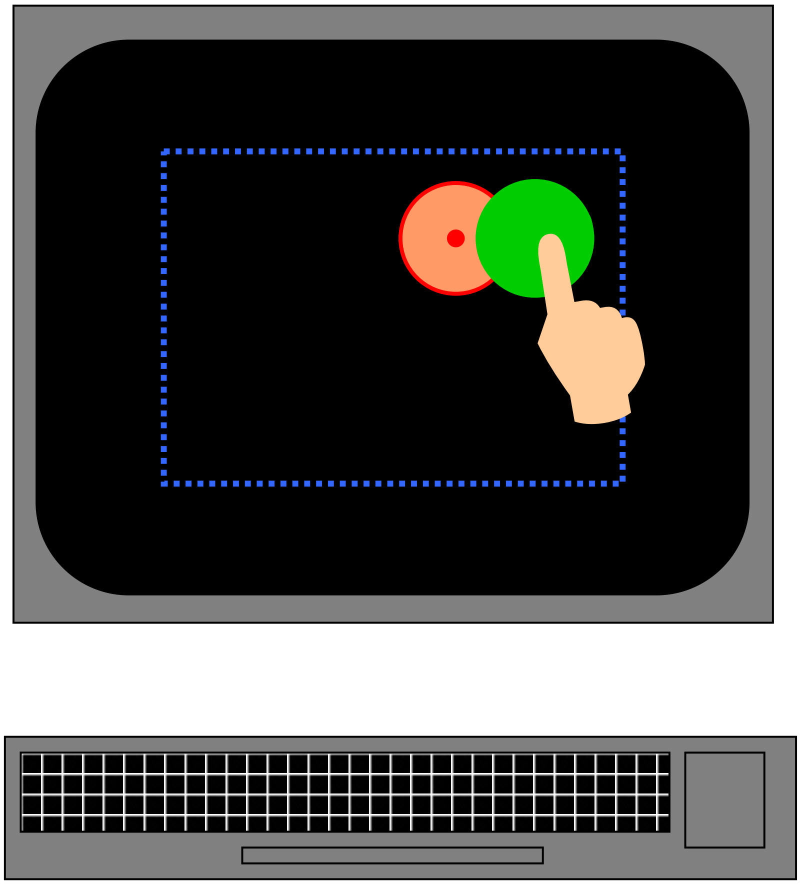

The green target is hit:
+100 points

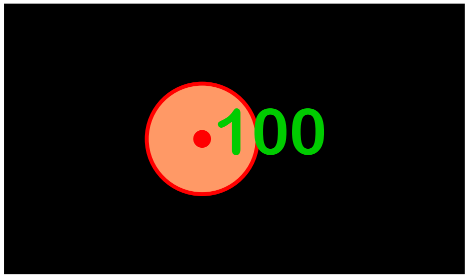

Experimental Task

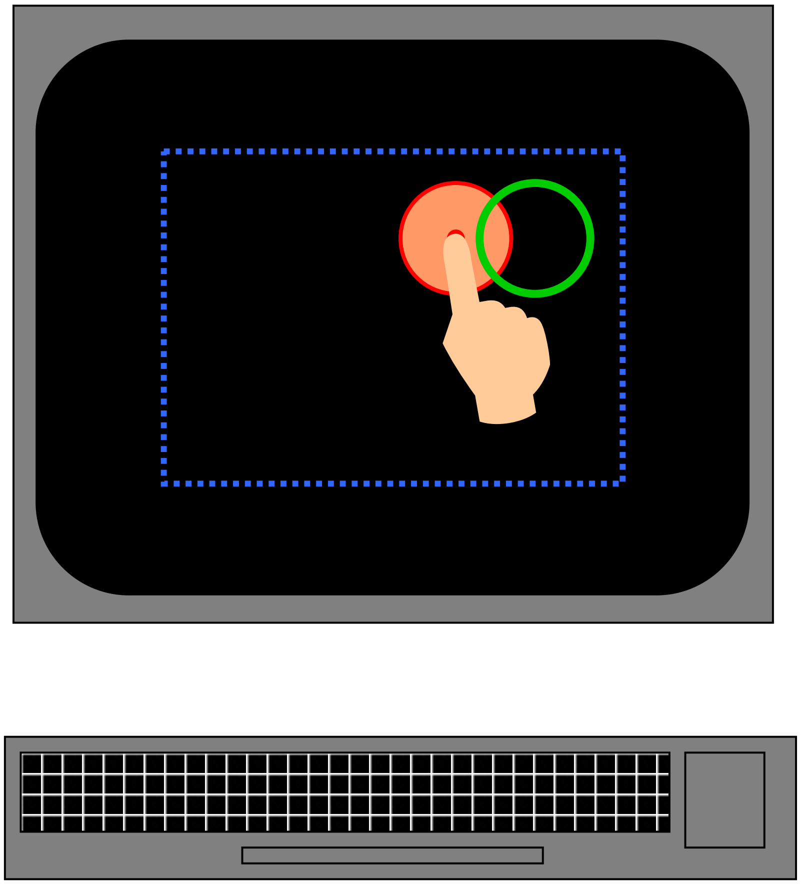

Experimental Task

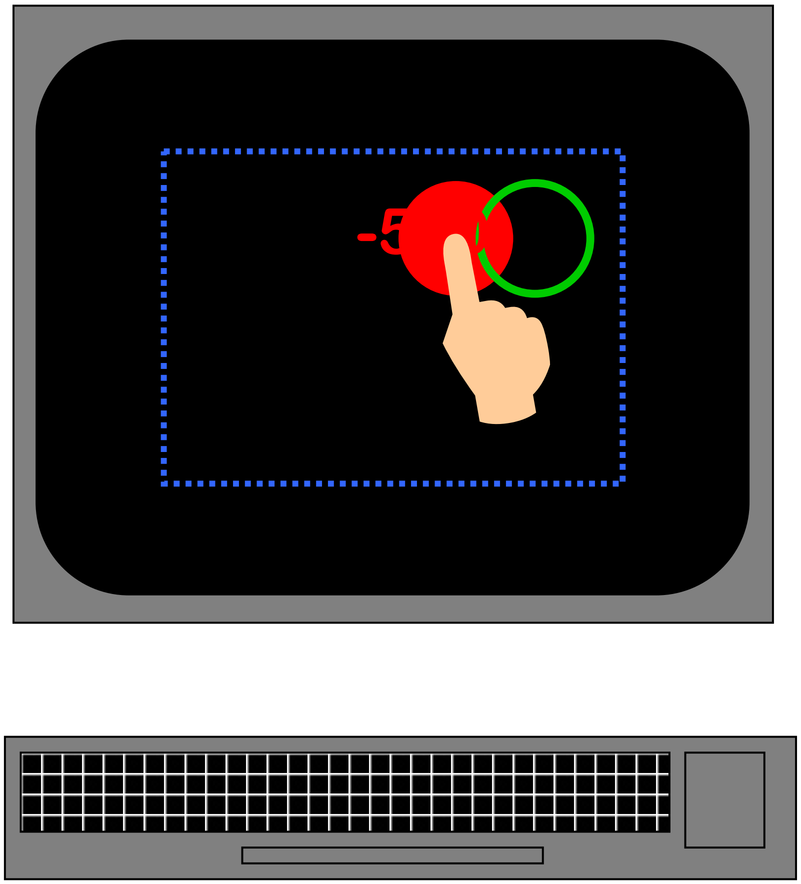

The red target is hit:
-500 points

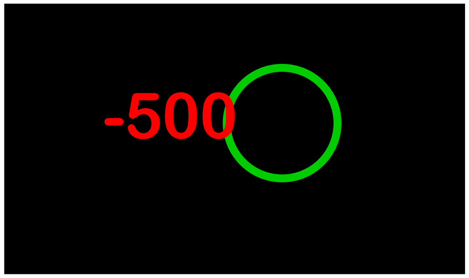

Experimental Task

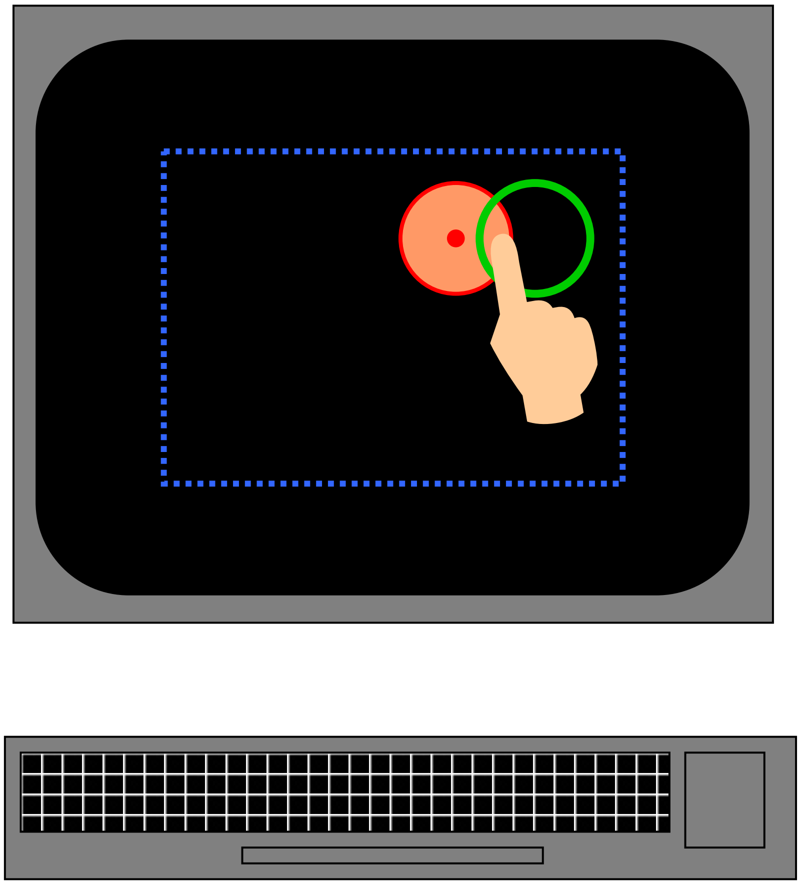

Experimental Task

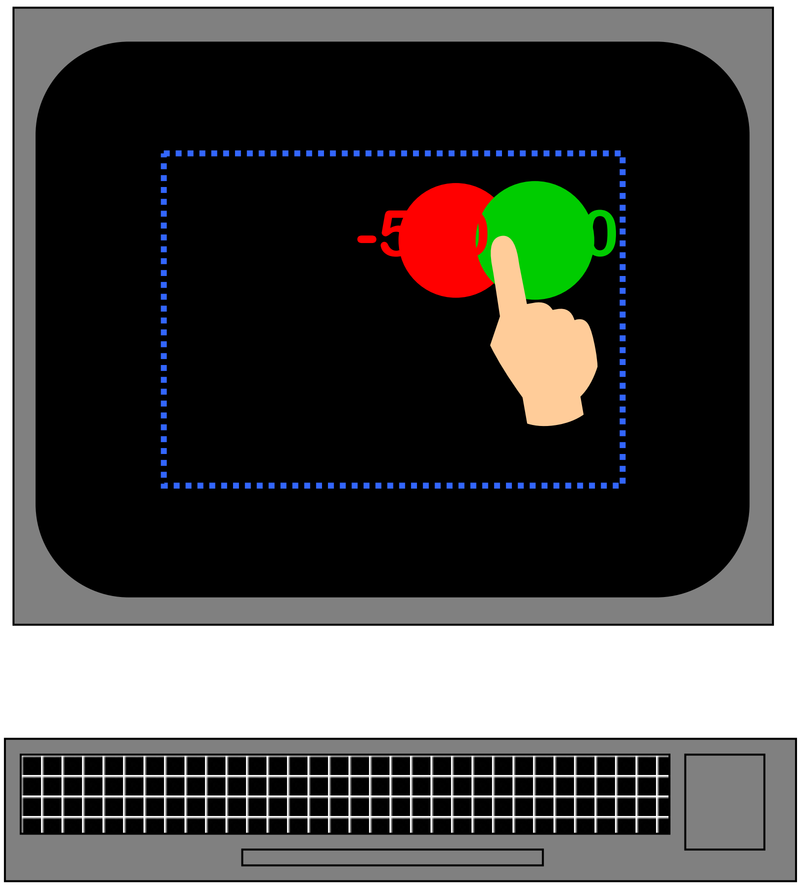

Scores add if both targets are hit:

Experimental Task

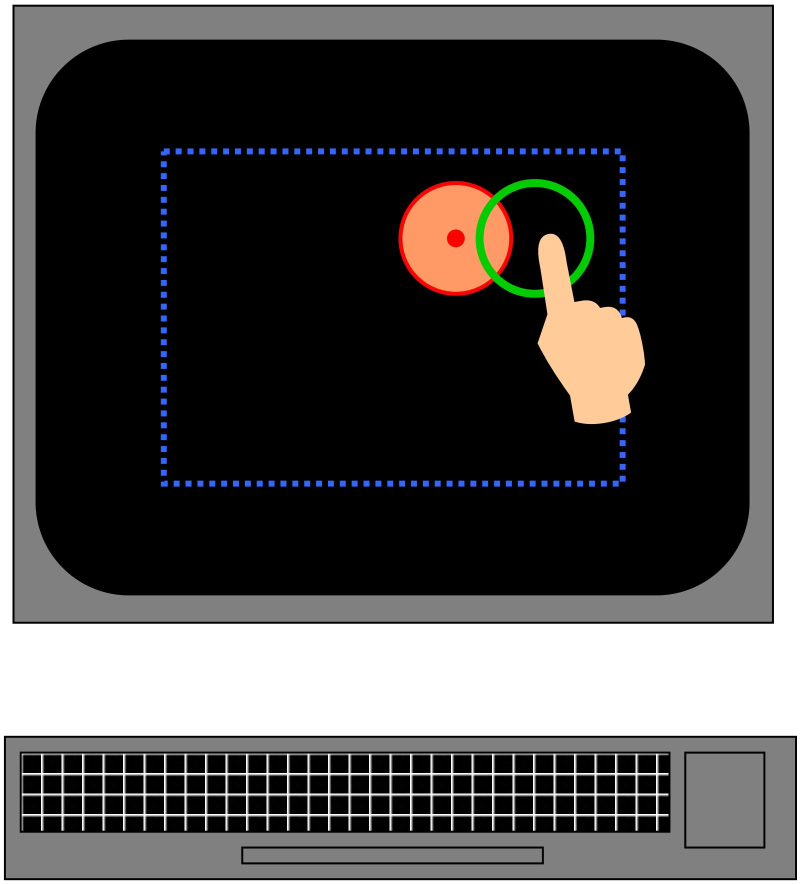

Experimental Task

The screen is hit later than 700 ms after target display: -700 points.

If you are on time but Miss the targets, 0.

Experimental Task

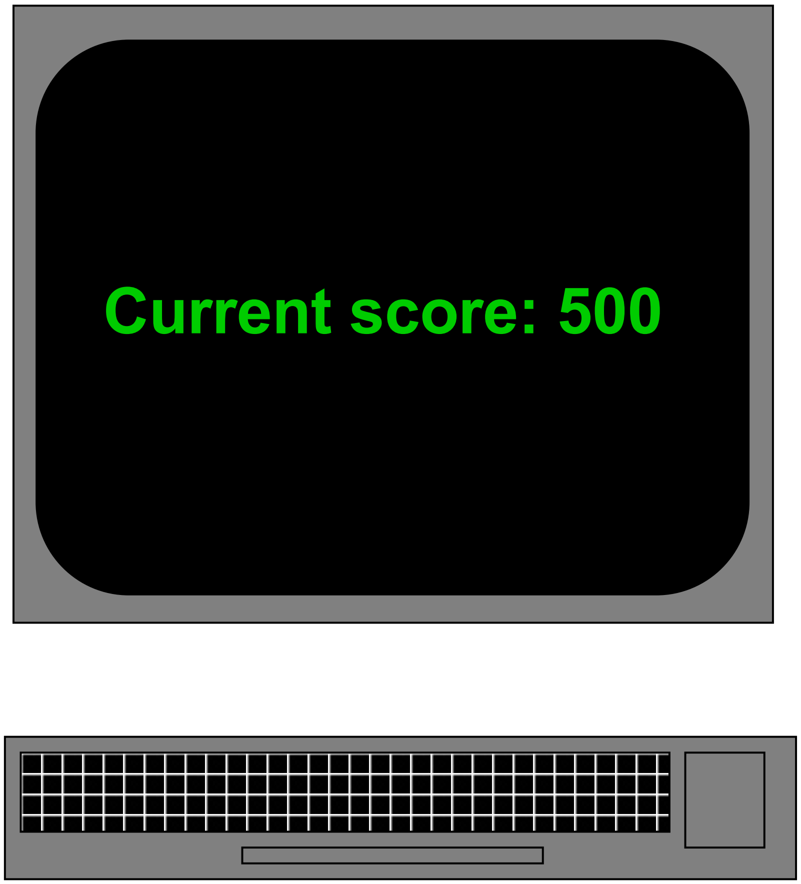

End of trial

Choice among Movement Strategies

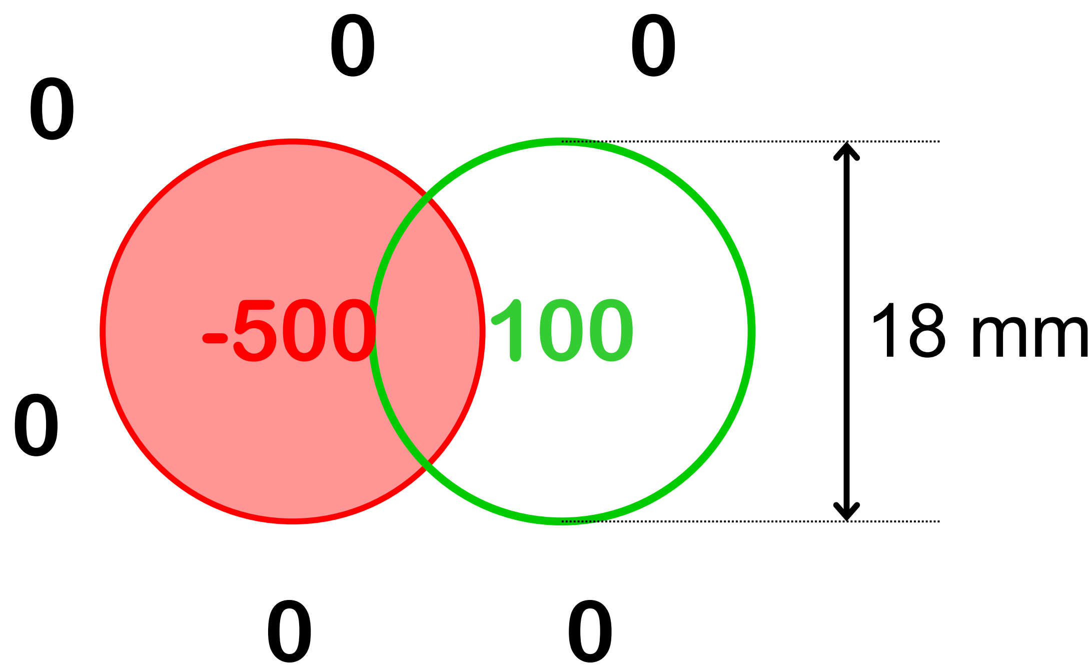

What should Paulina do?

Distribution of movement end points

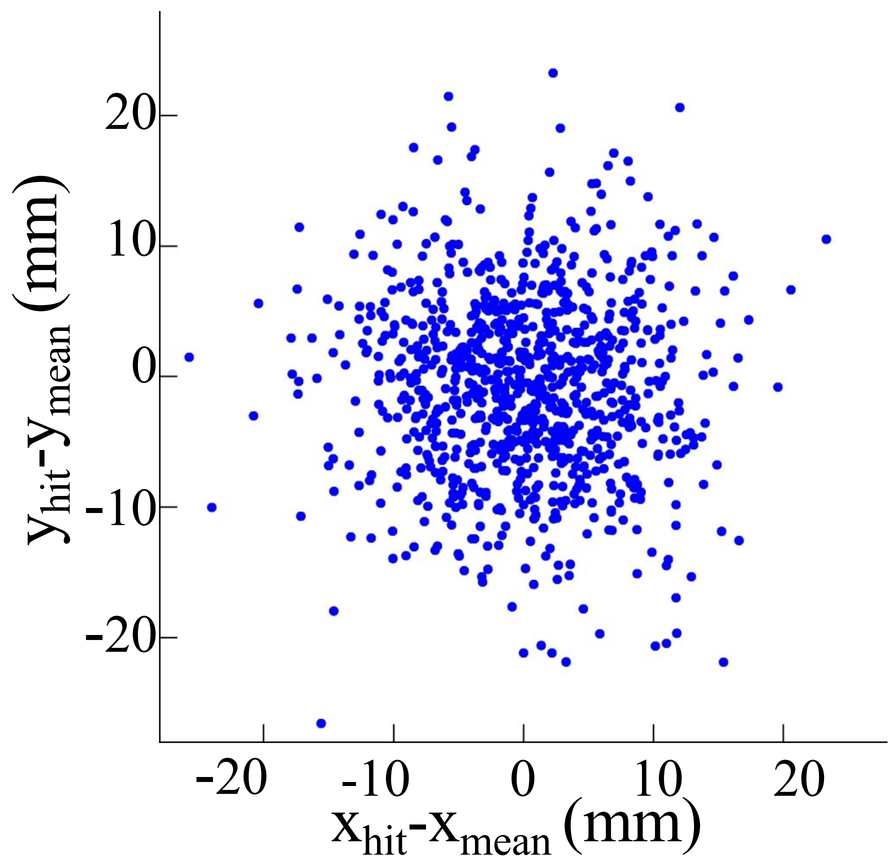

Subject S4, σ = 3.62 mm,
72x15 = 1080 end points

Q-Q Plot

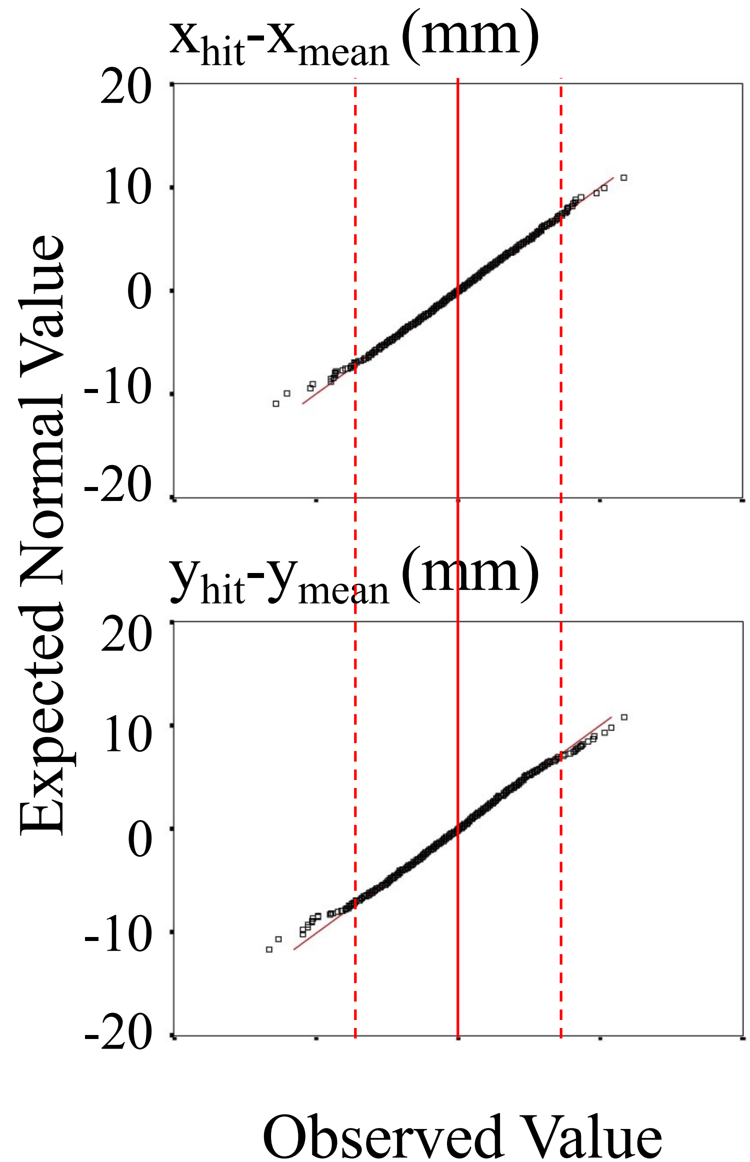

If there were no red penalty circle ….

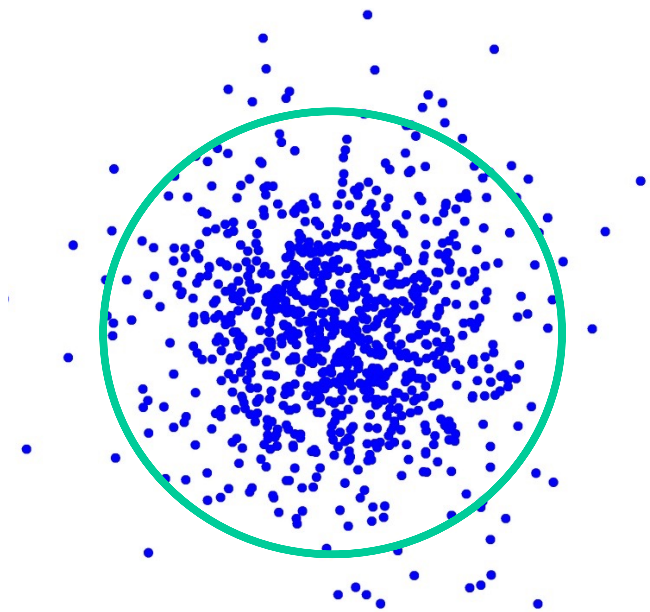

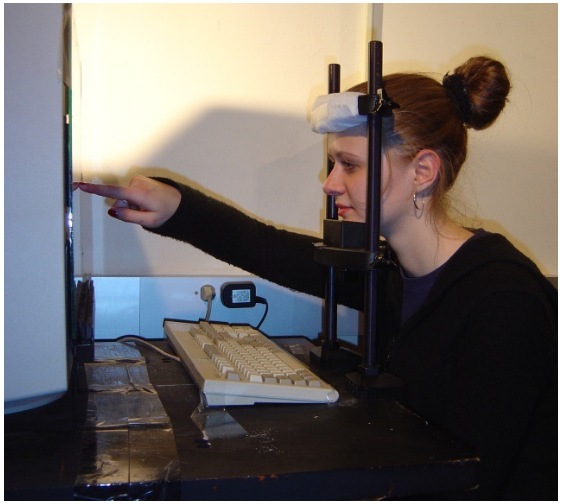

Aim for center
Select perceptual-motor strategies that
minimize variance

Harris & Wolpert (1998)

Choice among Movement Strategies

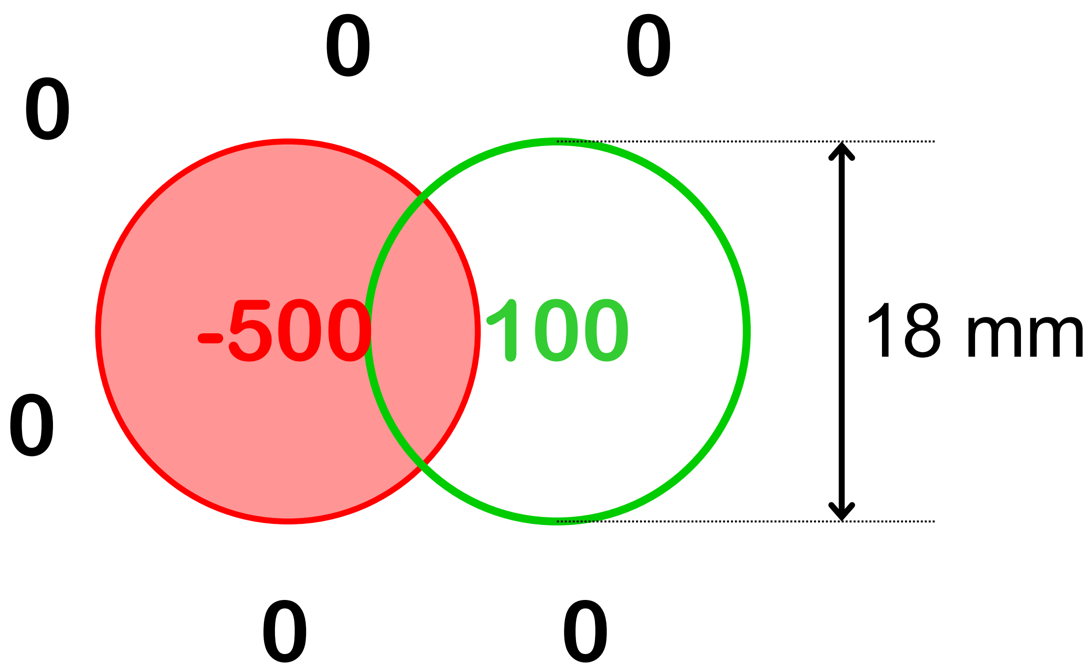

What should Paulina do?

Thought Experiment

: -500 : 100 points (2.5 ¢)

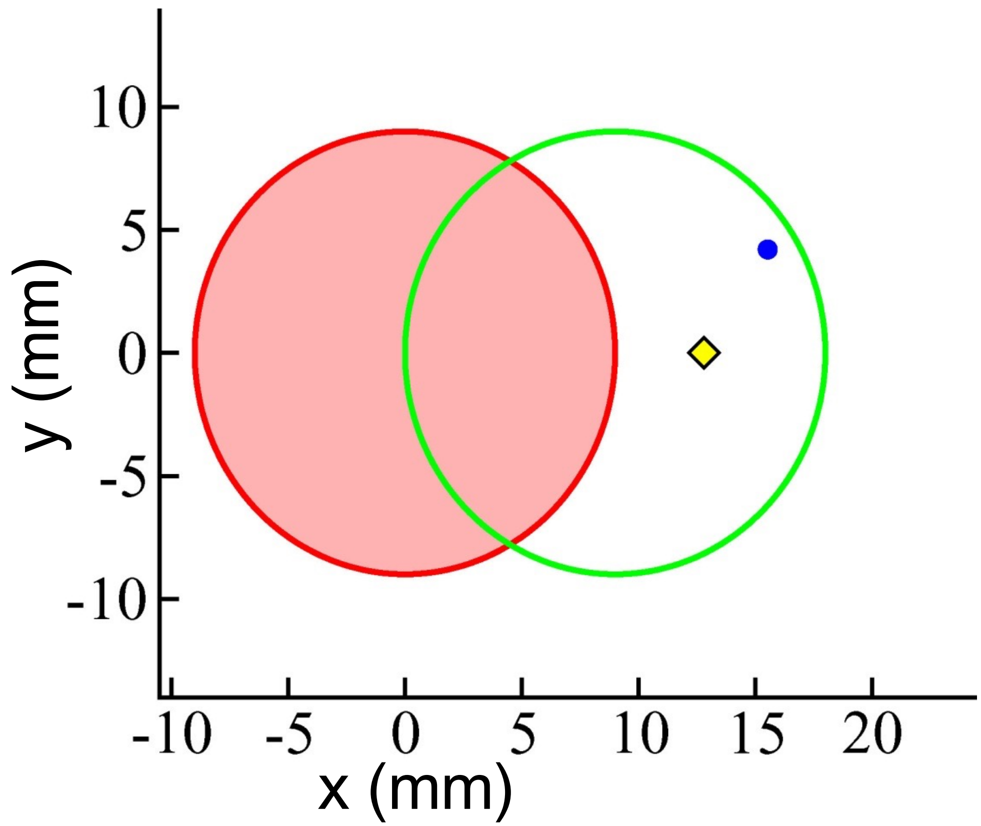

100 points

σ = 4.83 mm

Thought Experiment

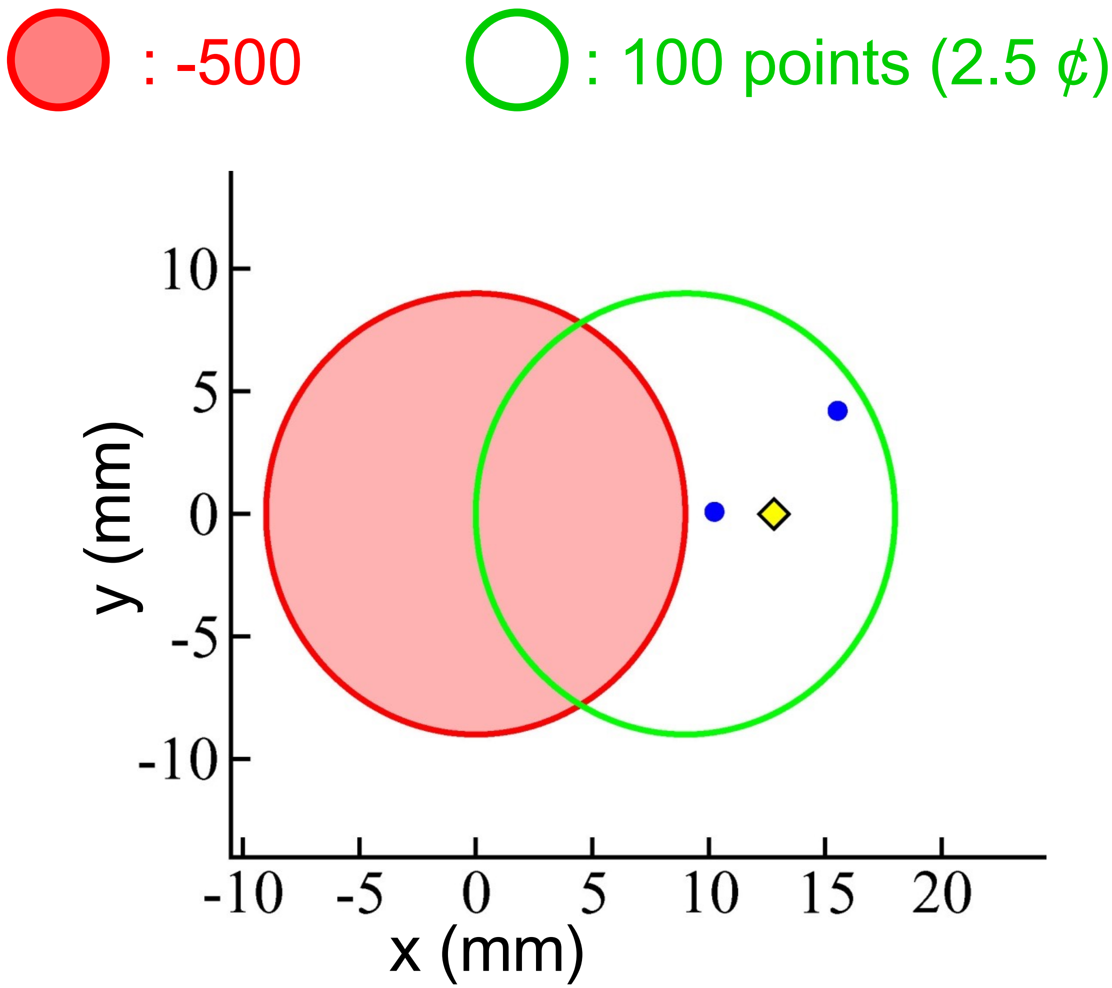

100 points
100 points

200 points

σ = 4.83 mm

Thought Experiment

: -500 : 100 points (2.5 ¢)

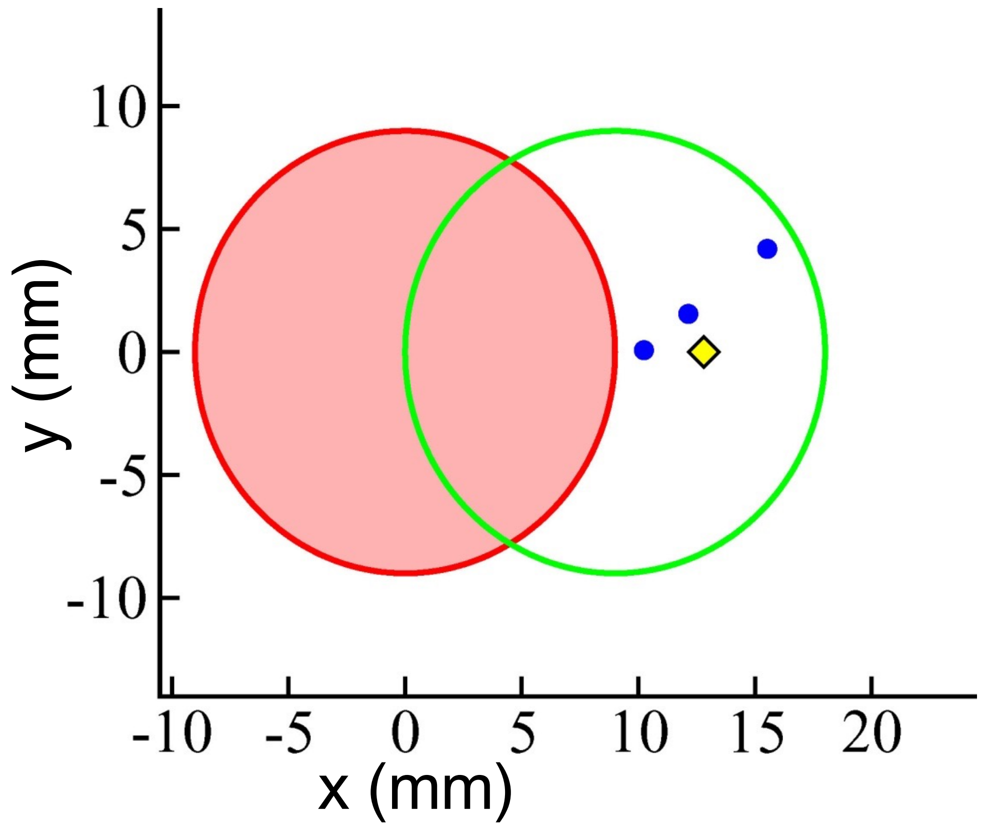

100 points
100 points
100 points

300 points

σ = 4.83 mm

Thought Experiment

: -500 : 100 points (2.5 ¢)

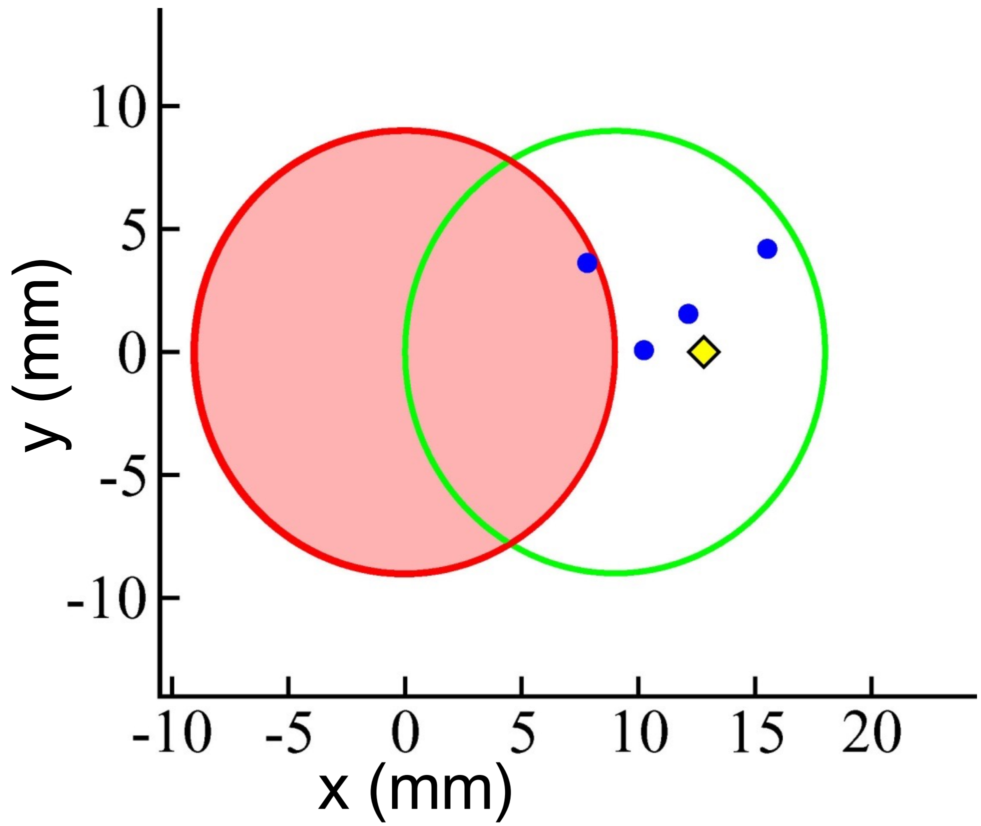

100 points
100 points
100 points
-400 points

-100 points

σ = 4.83 mm

Thought Experiment

: -500 : 100 points (2.5 ¢)

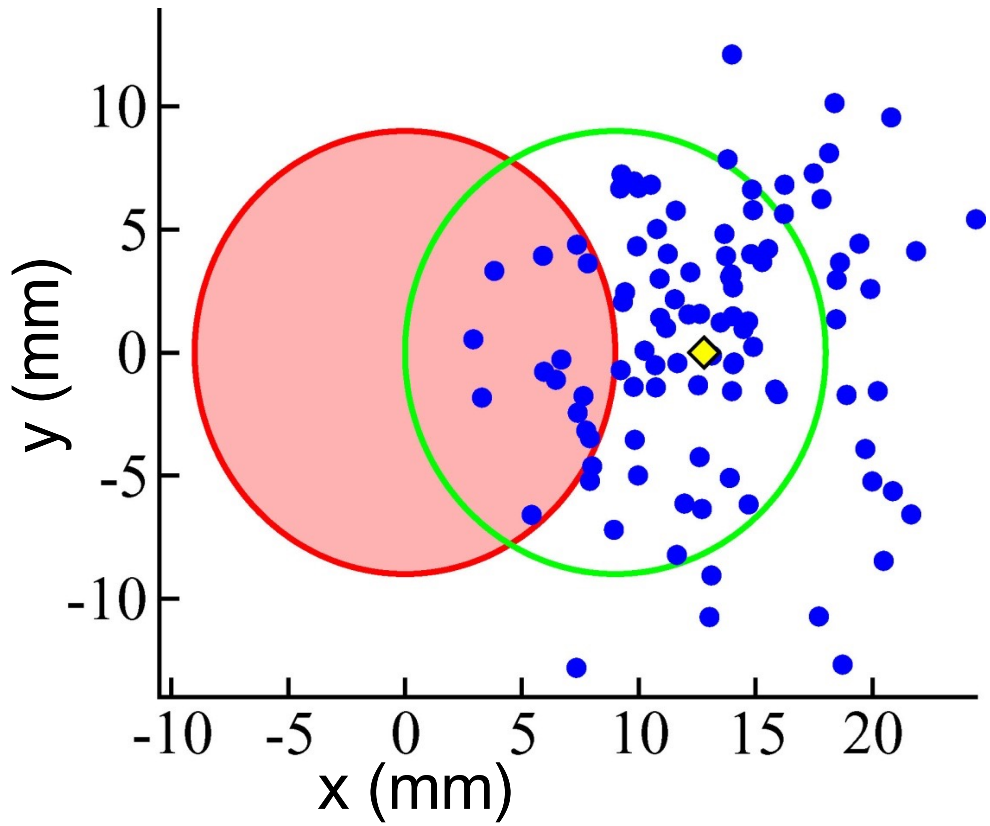

100 points
100 points
100 points
-400 points
.
.
.
.

-0.3 pts. per trial

σ = 4.83 mm

Thought Experiment

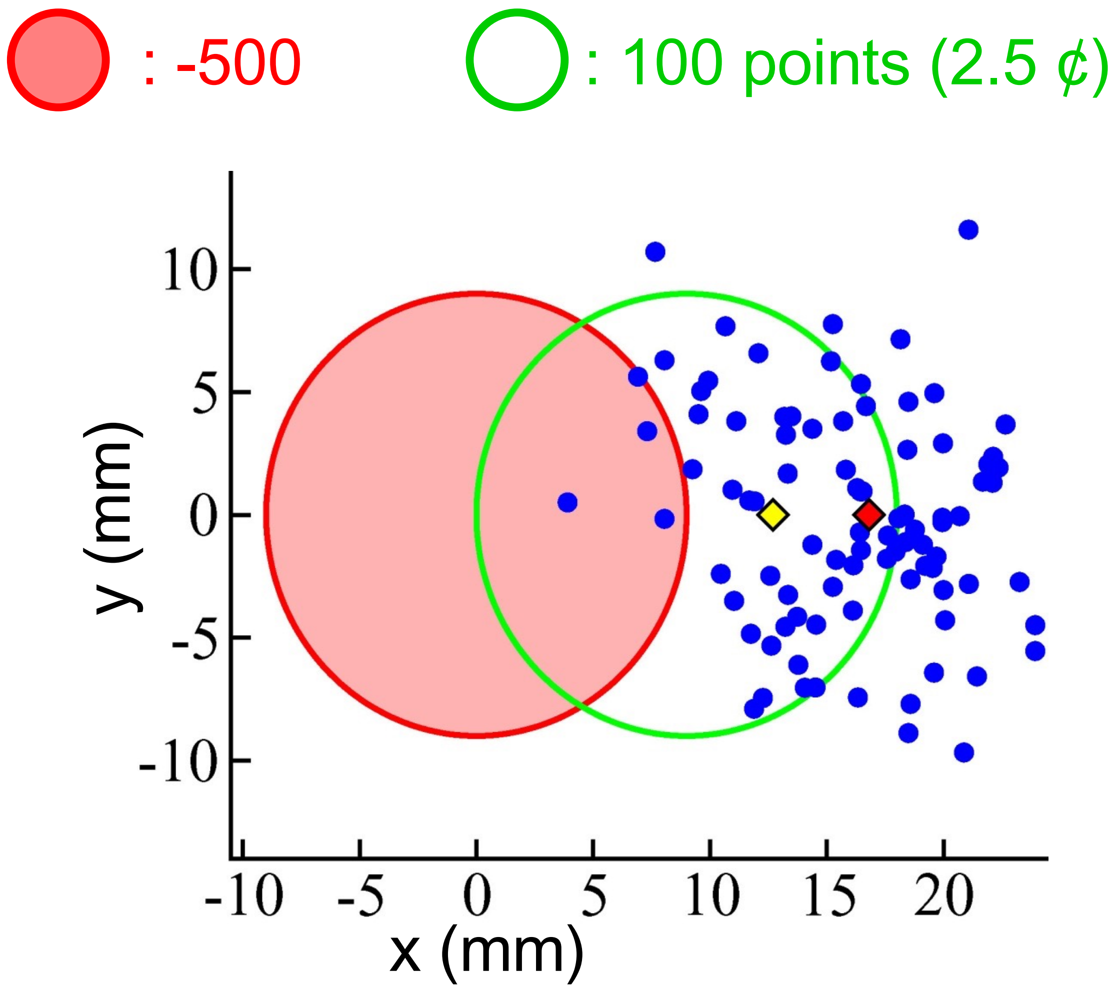

- 0.3 pts. per trial

30.7 pts. per trial

σ = 4.83 mm

Thought Experiment

: -500

: 100 points (2.5 ¢)

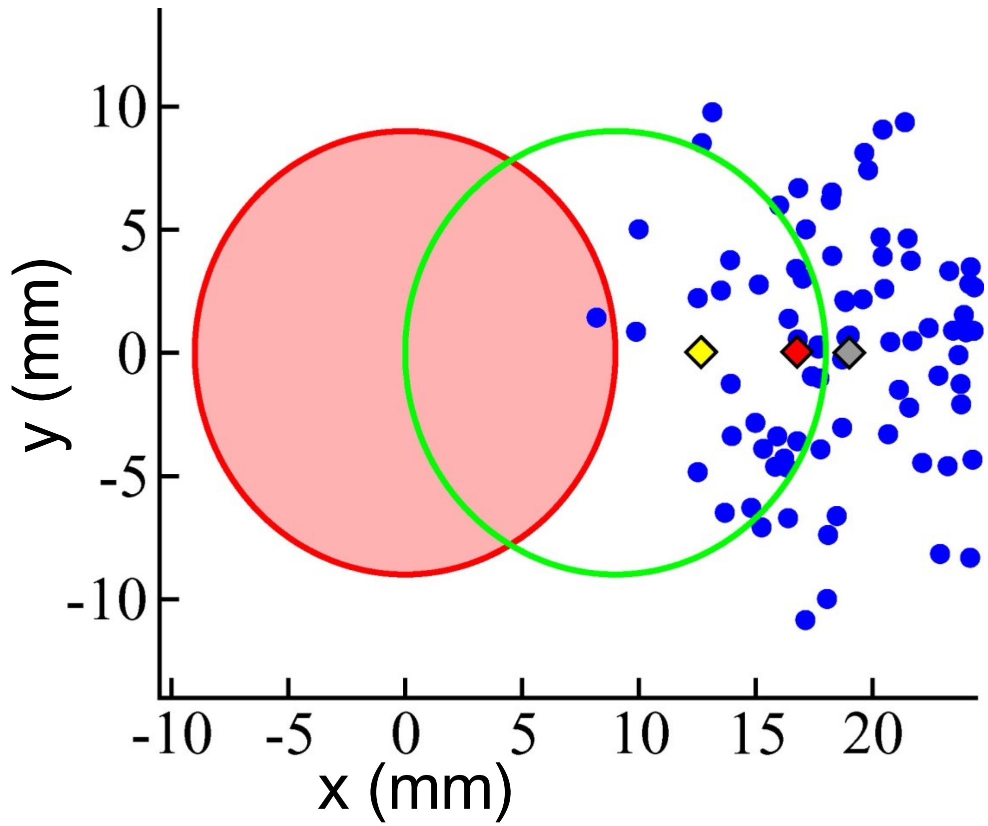

- 0.3 pts. per trial

30.7 pts. per trial

25.5 pts. per trial

σ = 4.83 mm

Thought Experiment

: -500 : 100 points (2.5 ¢)

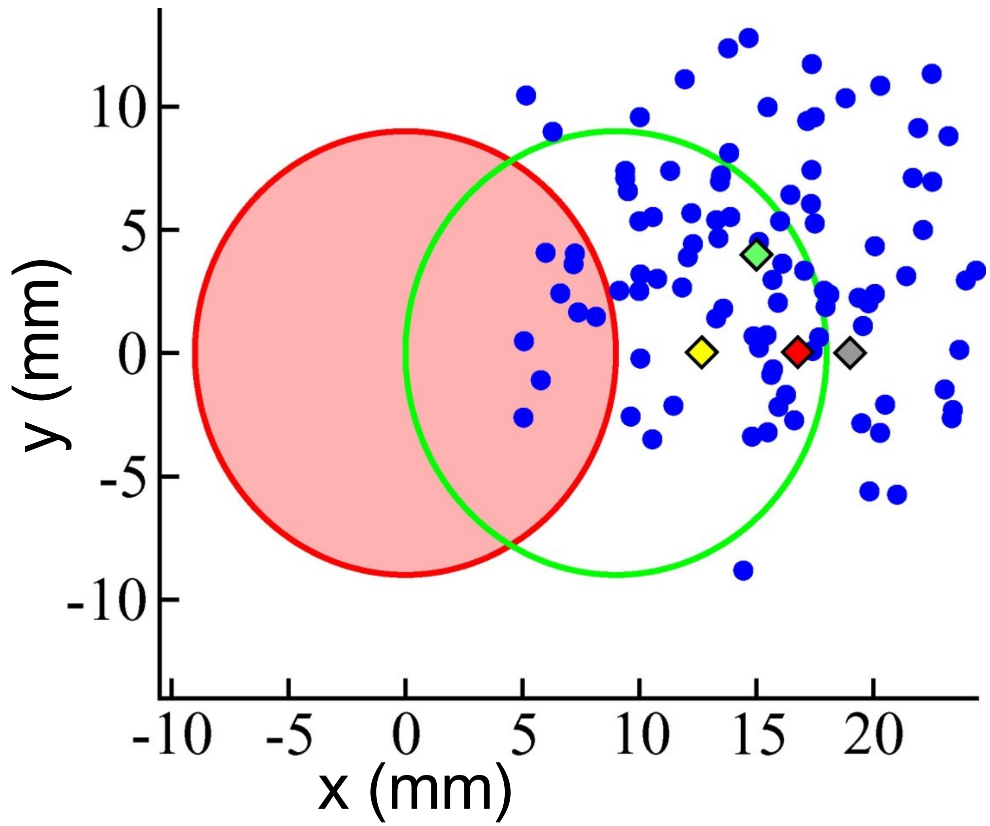

- 0.3 pts. per trial

30.7 pts. per trial

25.5 pts. per trial

22.6 pts. per trial

σ = 4.83 mm

Expected value as function of mean movement end point (x,y):

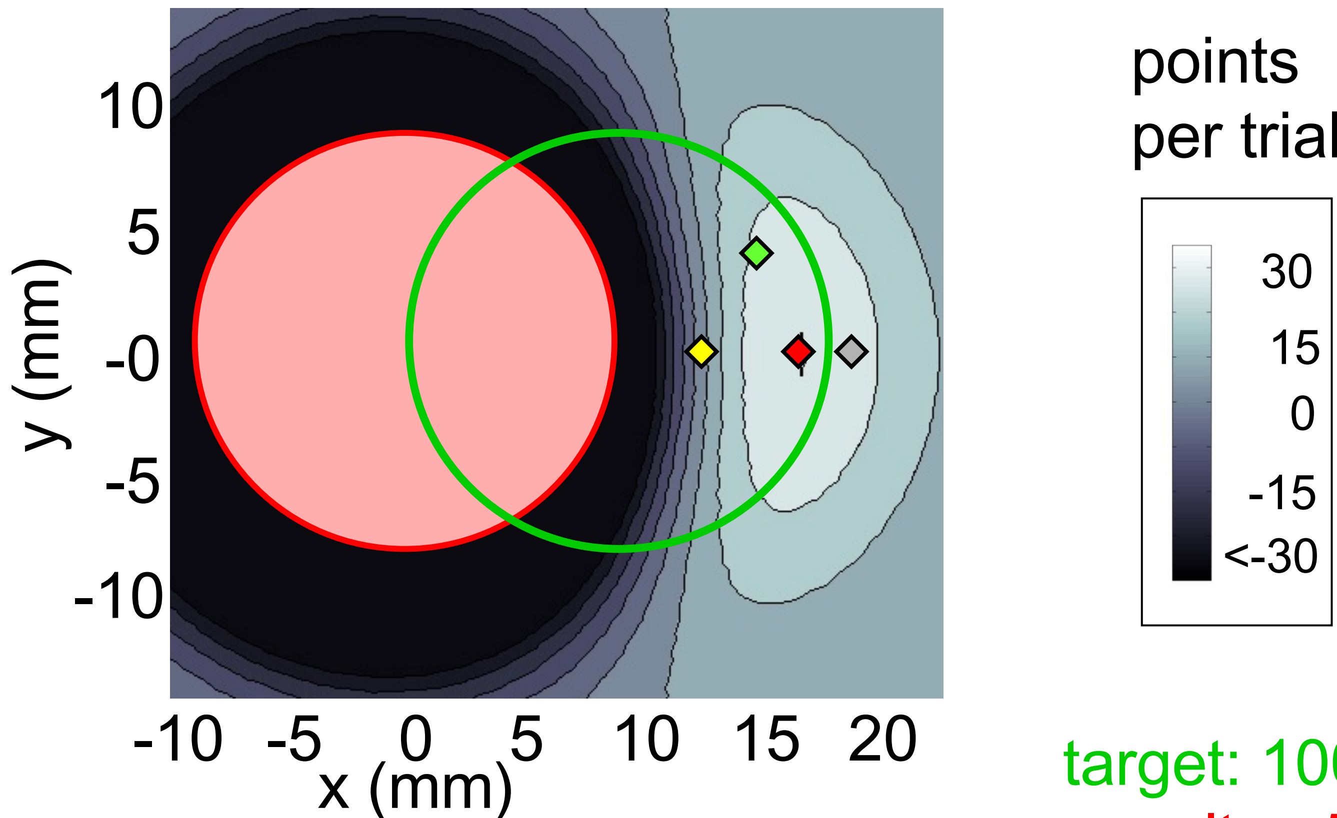

σ = 4.83 mm

target: 100

penalty: -500

Thought Experiment

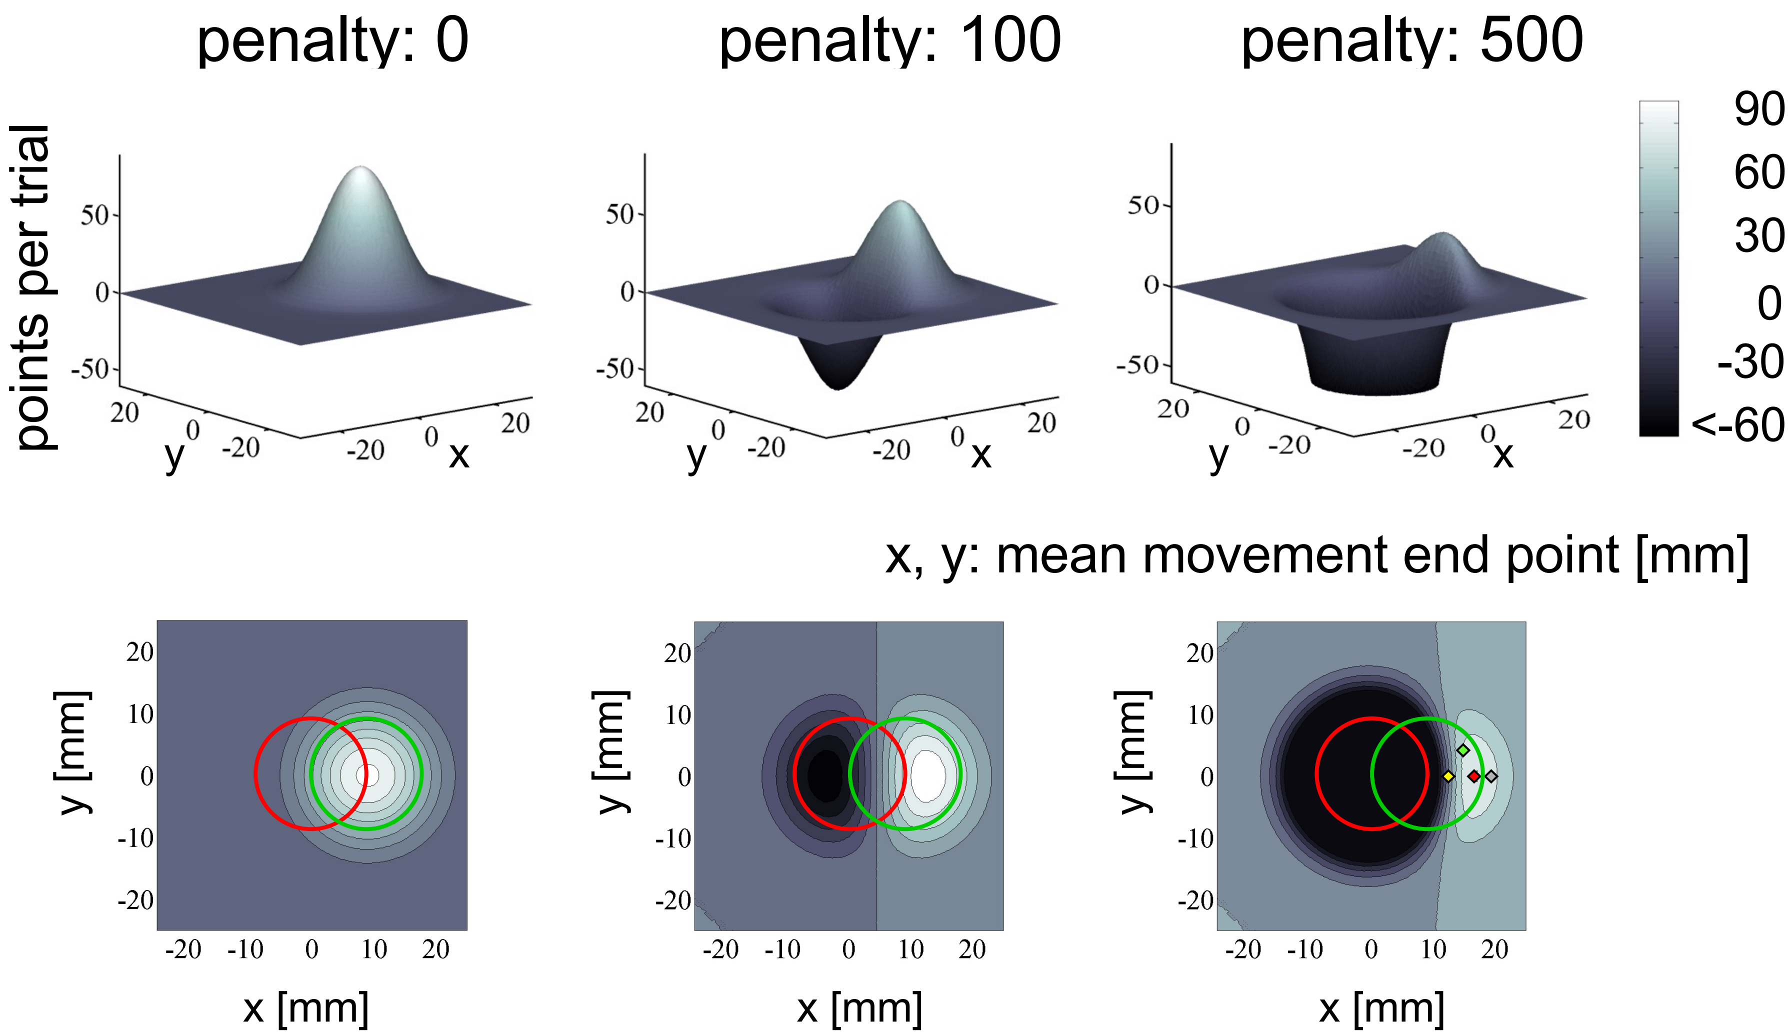

σ = 4.83 mm

Movement plans as lotteries

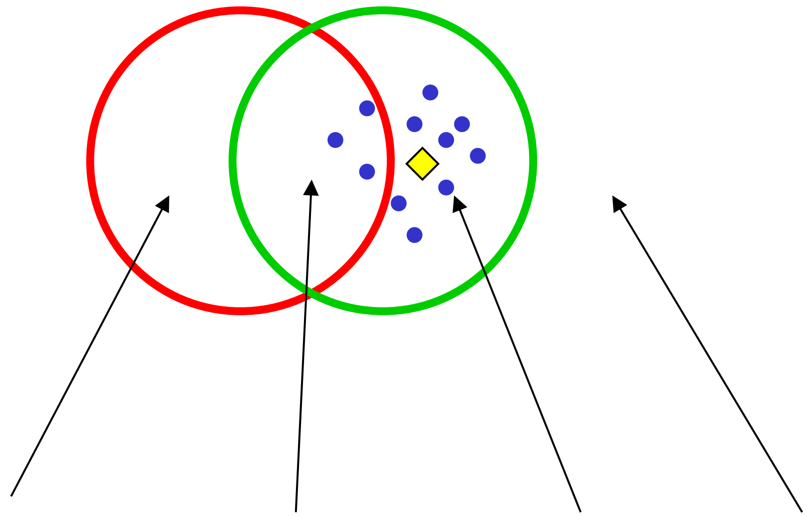

$(p_1, -500;\ p_2, -400;\ p_3, 100;\ p_4, 0)$

Movement plans as lotteries

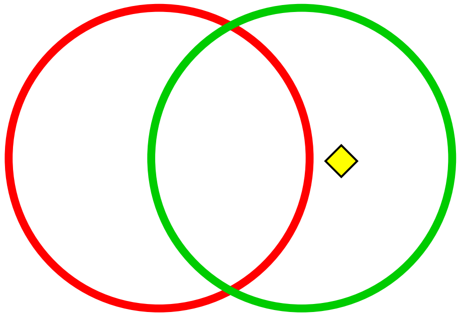

σ = 4 mm

Lottery:

(1.3%, -500;   30.3%, -400;   60.9%, 100;   7.5%, 0)

Movement plans as lotteries

Optimal aim point: lottery with MEV

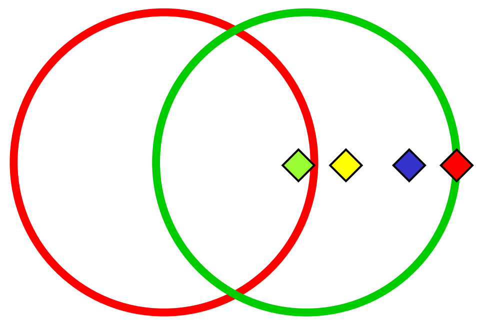

σ = 4 mm

(6.6%, -500; 52.3%, -400; 37.0%, 100; 4.0%, 0)

(1.3%, -500; 30.3%, -400; 60.9%, 100; 7.5%, 0)

(0%, -500; 4.6%, -400; 62.6%, 100; 32.8%, 0)

(0%, -500; 0.7%, -400; 37.6%, 100; 61.7%, 0)

Experiment 1

Reaching with Asymmetric Gain/Loss

Julia Trommershäuser

Trommershäuser, Maloney, Landy (2003) JOSA A

Test of the model: Experiment 1

4 stimulus
configurations:
(varied within block)

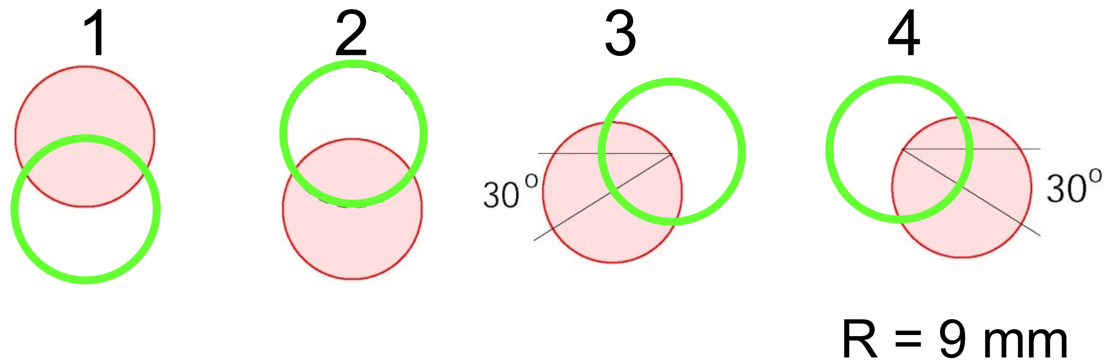

2 penalty conditions:
0 and -500 points (varied between blocks)

5 “practiced movers”
1 session of data collection: 360 trials
24 data points per condition

General Methods: Training

For all experiments:

- *All subjects practice the task for 360 trials or more until their variance stabilizes.*
- *The timeout limit is gradually decreased to 700 ms during training.*

- *There are no penalties during training (the concept is never mentioned).*
- *We verify that each subject's movement variance has stabilized.*

- *They are told only to make money.*

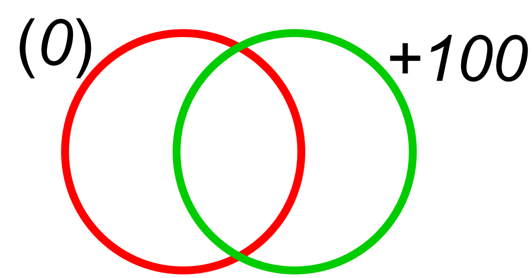

Results: Experiment 1

Model prediction:

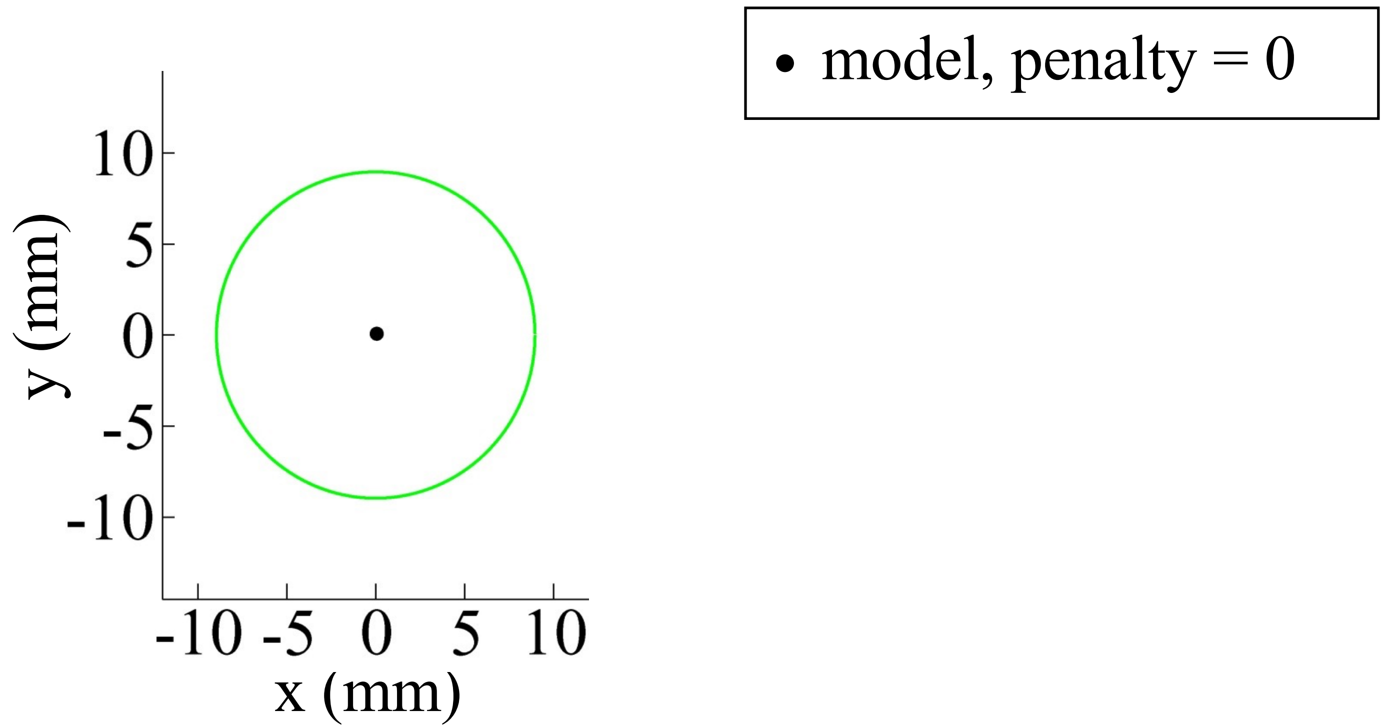

Subject S5, σ = 2.99 mm

Results: Experiment 1

Model prediction: configuration 1

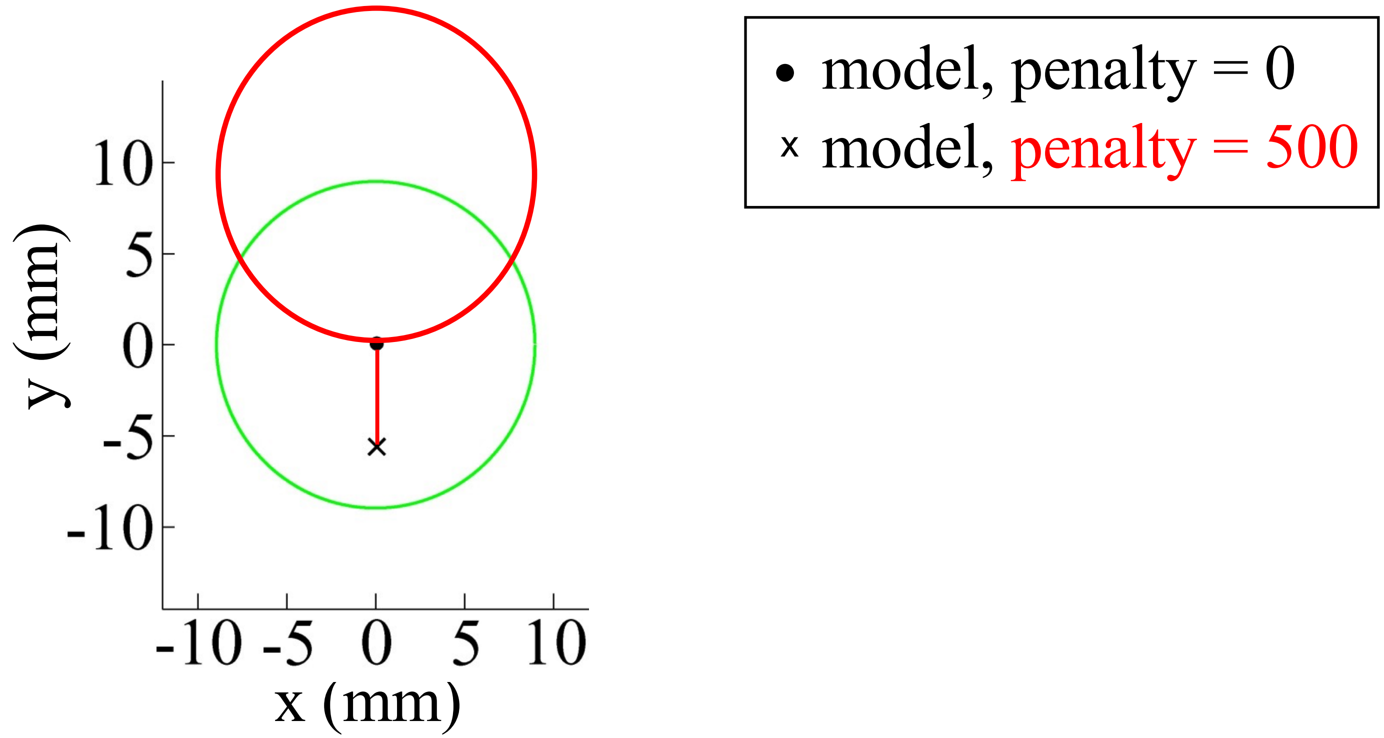

Subject S5, σ = 2.99 mm

Results: Experiment 1

Model prediction: configuration 2

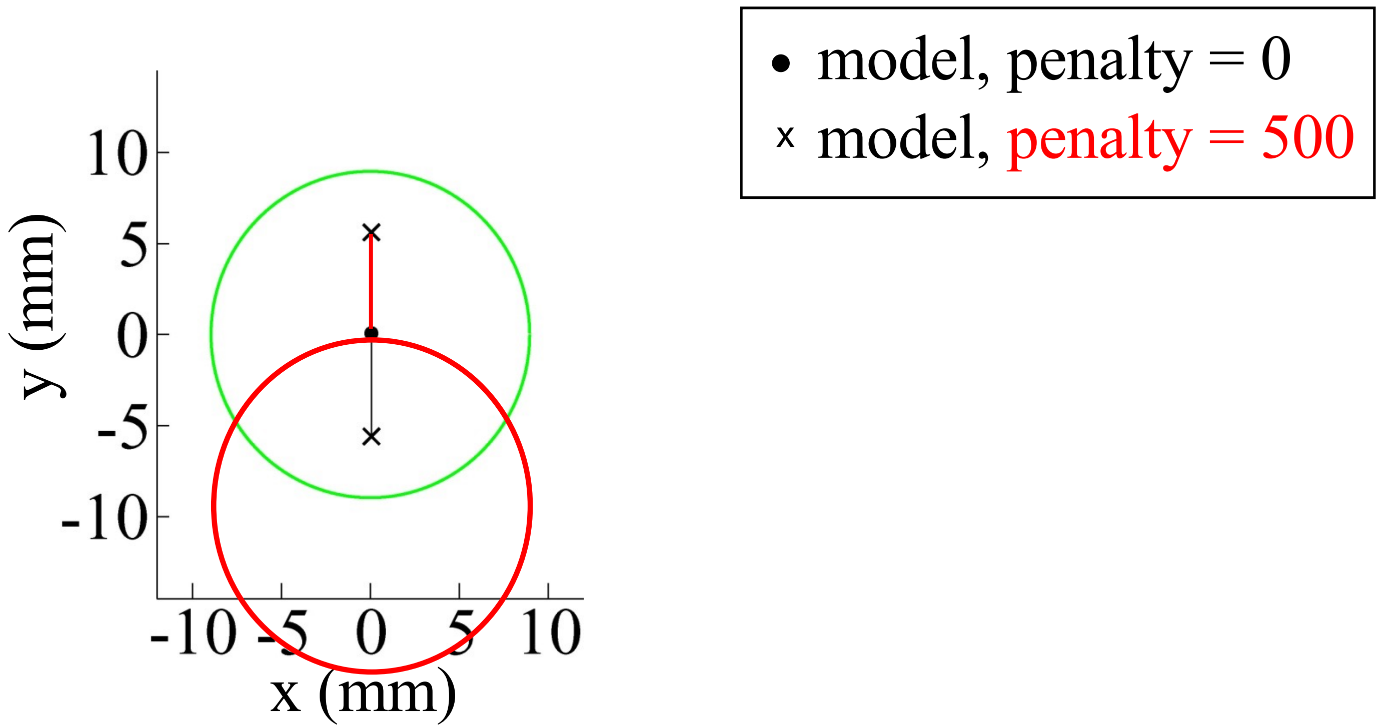

Subject S5, σ = 2.99 mm

Results: Experiment 1

Model prediction: configuration 3

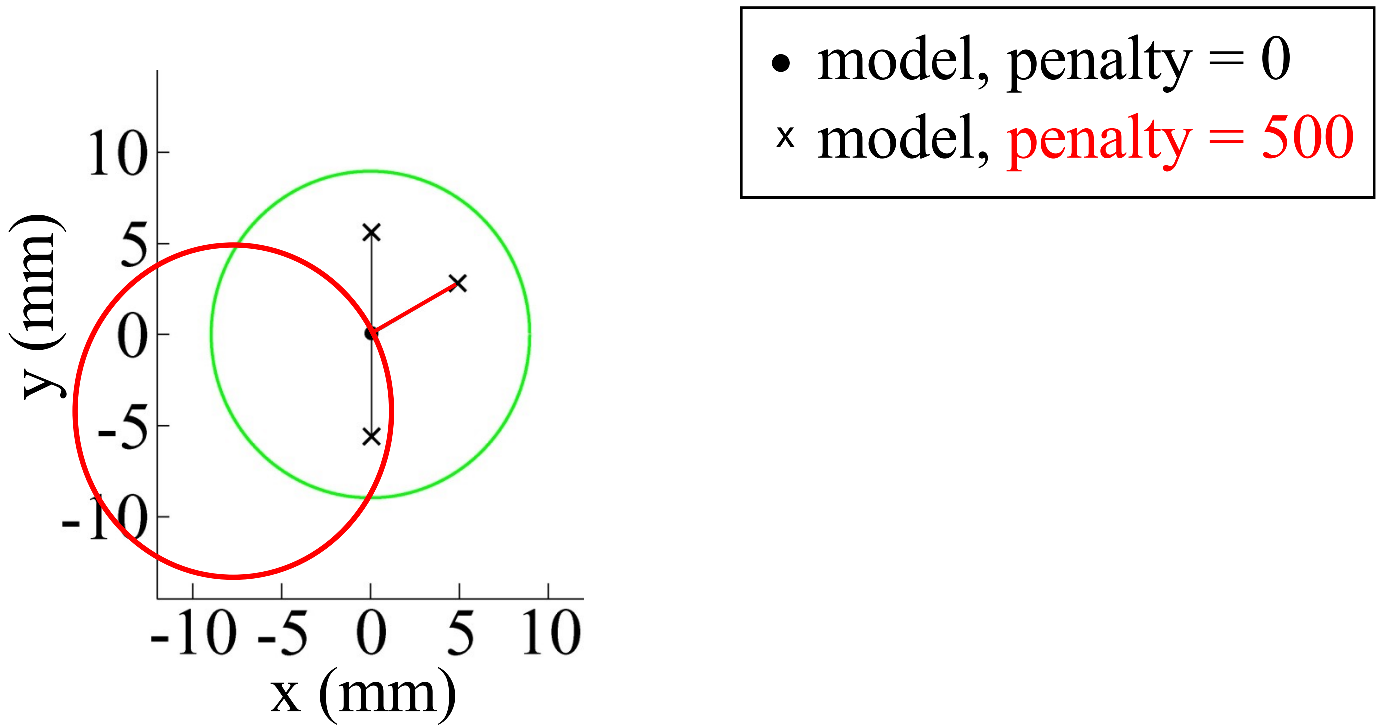

Subject S5, σ = 2.99 mm

Results: Experiment 1

Model prediction: configuration 4

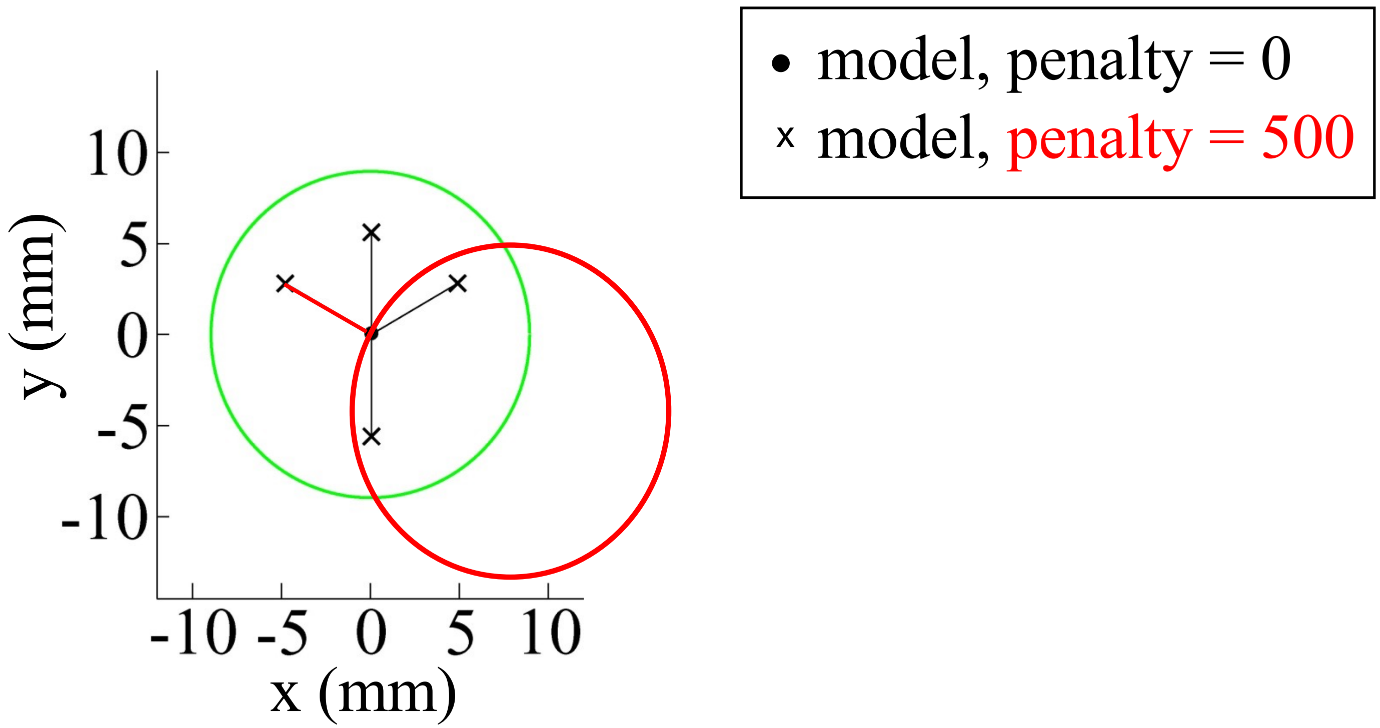

Subject S5, σ = 2.99 mm

Results: Experiment 1

Comparison with experiment

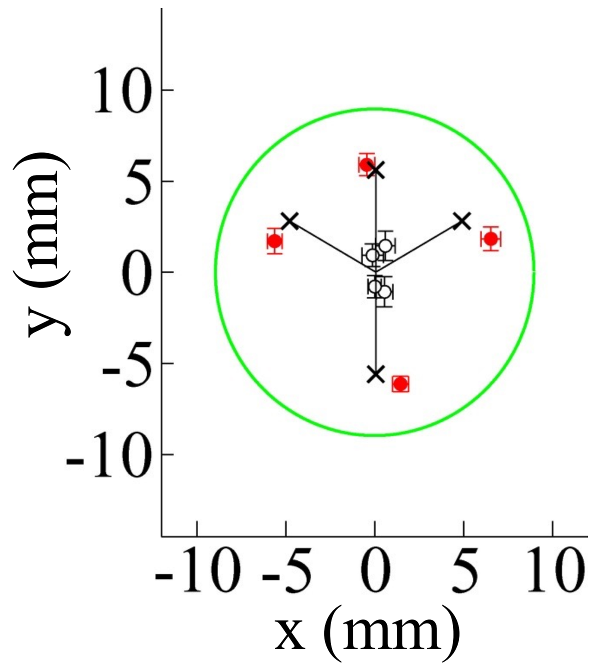

○ exp., penalty = 0
● exp., penalty = 500
x model, penalty = 500

Subject S5, σ = 2.99 mm

Results: Experiment 1

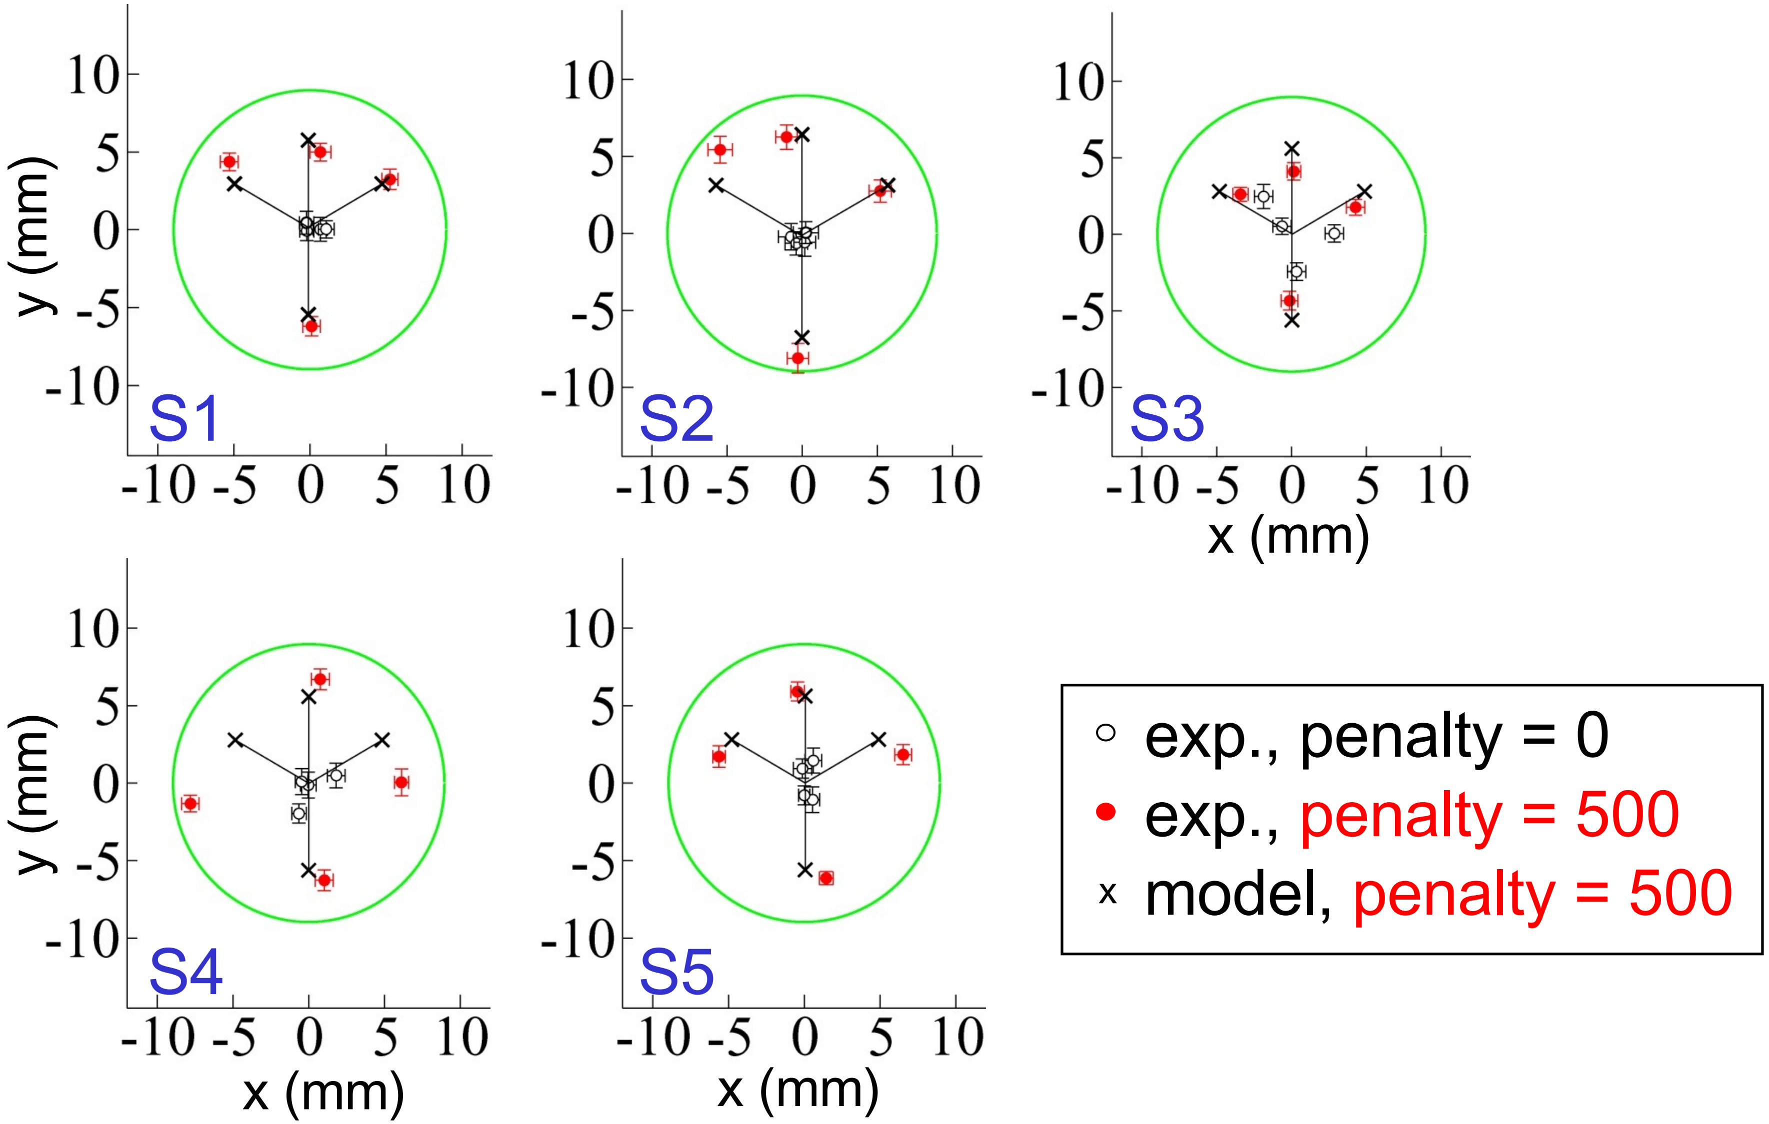

Results: Experiment 1

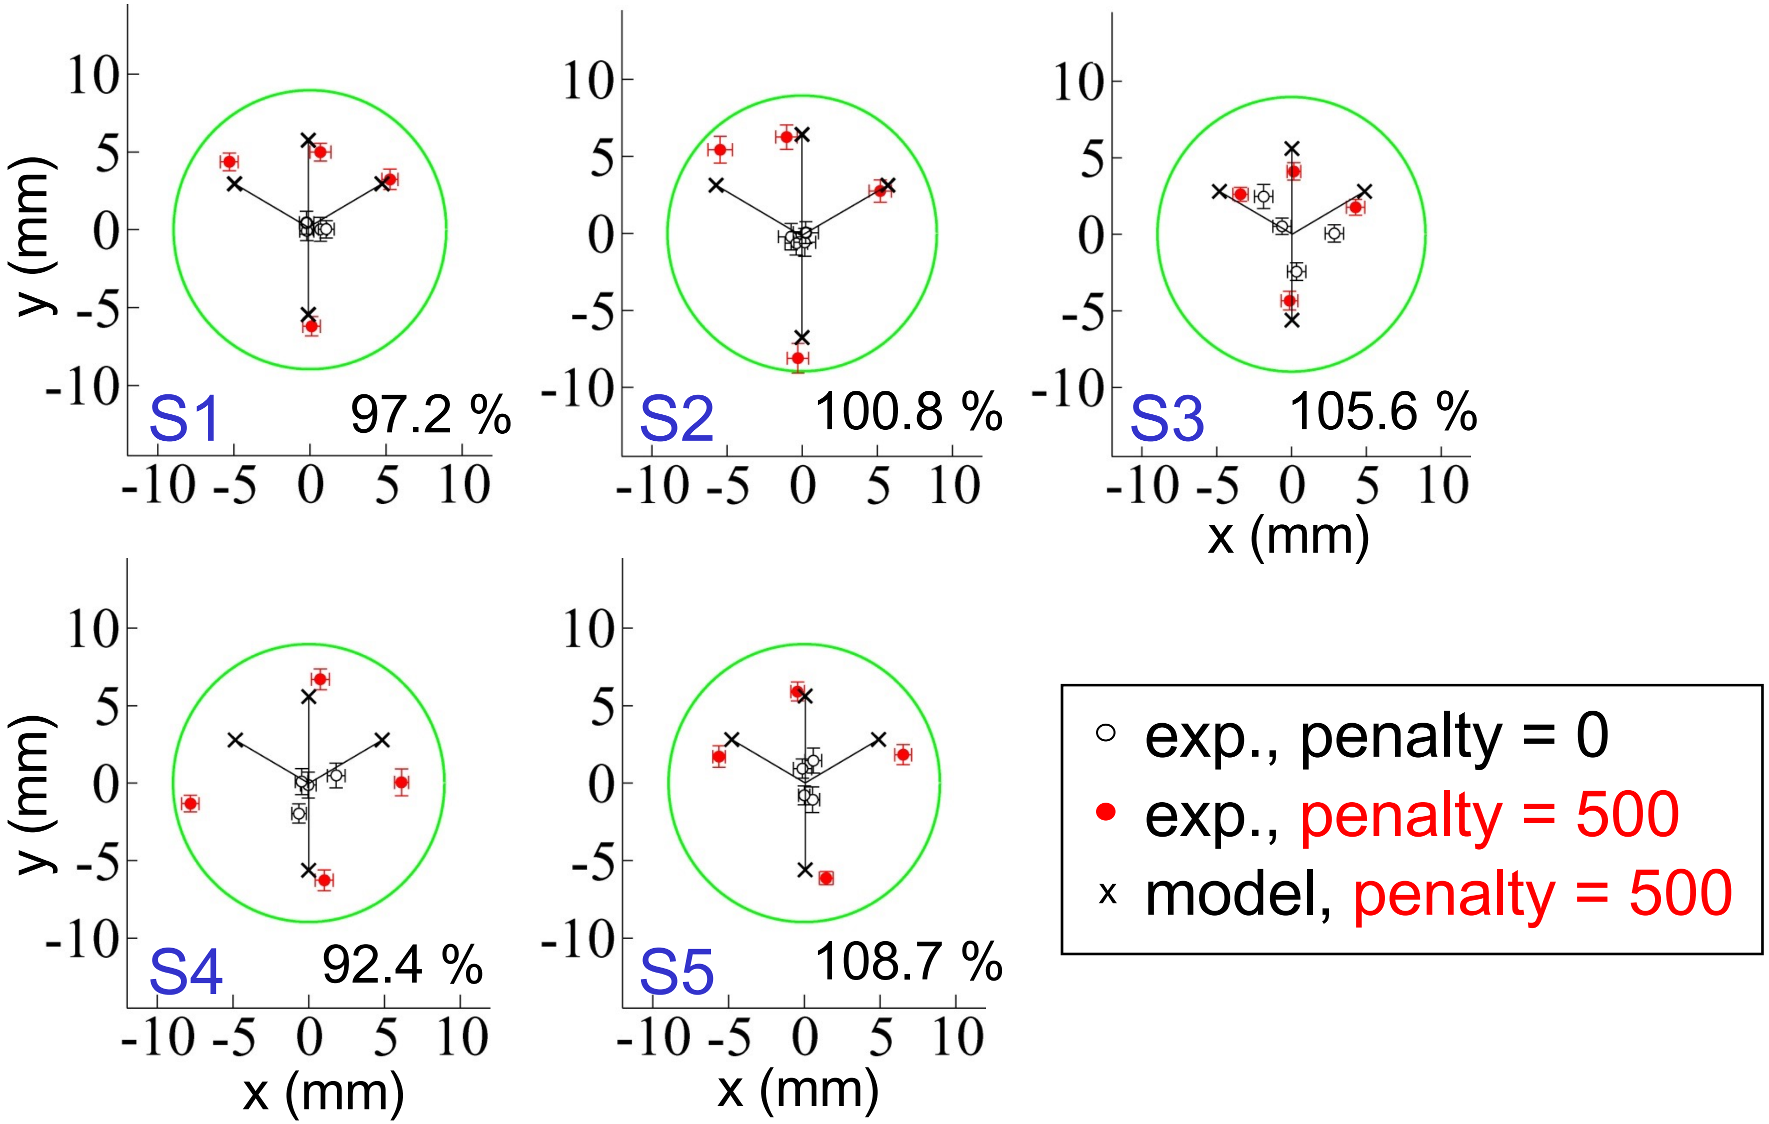

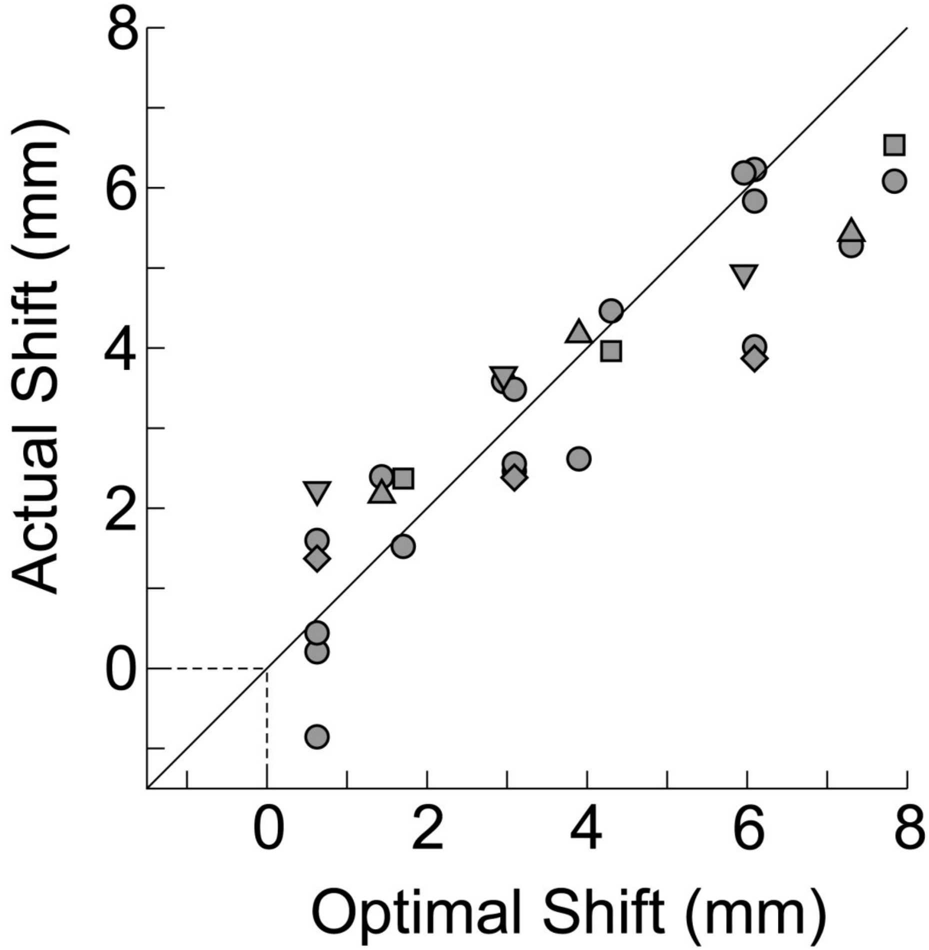

8
6
4
2
0
Actual Shift (mm)
0
2
4
6
8
Optimal Shift (mm)

Conclusions

BDT is a promising mode; of movement planning

The use of explicit cost and rewards allow us to probe a much wider range of behavior than previously explored.

Managing uncertainty is the central task of a biological organism.