Lecture 2: Bayesian Games

Megan Peters
Larry Maloney

CoSMo 2018 Minneapolis, MN

Why are you so wonderful?

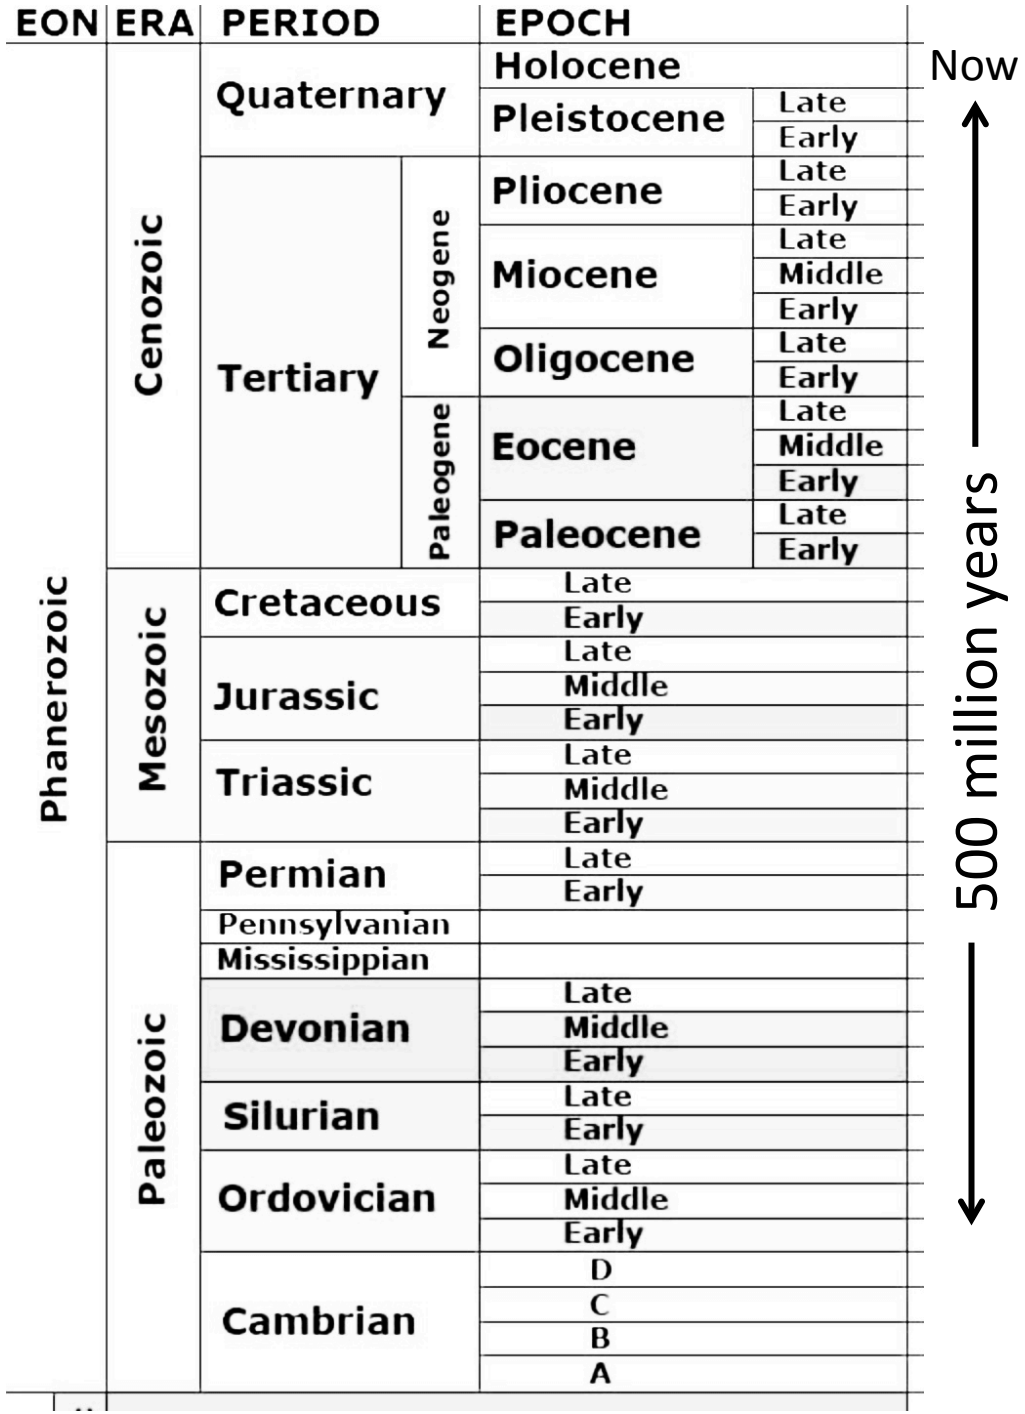

M F Land & D-E Nilson, (2002)
Animal Eyes. Oxford.

EON	ERA	PERIOD		EPOCH	
Phanerozoic	Cenozoic	Quaternary		Holocene	
				Pleistocene	Late
					Early
		Tertiary	Neogene	Pliocene	Late
					Early
				Miocene	Late
					Middle
					Early
				Oligocene	Late
					Early
			Paleogene	Eocene	Late
					Middle
					Early
				Paleocene	Late
					Early
	Mesozoic	Cretaceous		Late	
				Early	
		Jurassic		Late	
				Middle	
				Early	
		Triassic		Late	
				Middle	
				Early	
	Paleozoic	Permian		Late	
				Early	
		Pennsylvanian			
		Mississippian			
		Devonian		Late	
				Middle	
				Early	
		Silurian		Late	
				Early	
		Ordovician		Late	
				Middle	
				Early	
		Cambrian		D	
				C	
				B	
				A	

Now

500 million years

Cambrian explosion

anomalocaridid

eyes

M F Land & D-E Nilson, (2002)
Animal Eyes. Oxford.

EON	ERA	PERIOD		EPOCH	
Phanerozoic	Cenozoic	Quaternary		Holocene	
				Pleistocene	Late
					Early
		Tertiary	Neogene	Pliocene	Late
					Early
				Miocene	Late
					Middle
					Early
				Oligocene	Late
					Early
			Paleogene	Eocene	Late
					Middle
					Early
				Paleocene	Late
					Early
	Mesozoic	Cretaceous		Late	
				Early	
		Jurassic		Late	
				Middle	
				Early	
		Triassic		Late	
				Middle	
				Early	
	Paleozoic	Permian		Late	
				Early	
		Pennsylvanian			
		Mississippian			
		Devonian		Late	
				Middle	
				Early	
		Silurian		Late	
				Early	
		Ordovician		Late	
				Middle	
				Early	
		Cambrian		D	
				C	
				B	
				A	

Now

500 million years

Cambrian explosion

anomalocaridid

eyes

skeletons

M F Land & D-E Nilson, (2002)
Animal Eyes. Oxford.

Cambrian explosion

eyes
skeletons
muscles + levers

perception
movement

M F Land & D-E Nilson, (2002)
Animal Eyes. Oxford.

Cambrian explosion

eyes — perception

skeletons — movement

muscles + levers

"arms race"

M F Land & D-E Nilson, (2002)
Animal Eyes. Oxford.

Cambrian explosion

better eyes

better skeletons

muscles + levers

"arms race"

M F Land & D-E Nilson, (2002)
Animal Eyes. Oxford.

Cambrian explosion

perception

action

better eyes

better skeletons

muscles + levers

something missing …

M F Land & D-E Nilson, (2002)
Animal Eyes. Oxford.

Cambrian explosion

better eyes

better skeletons

muscles + levers

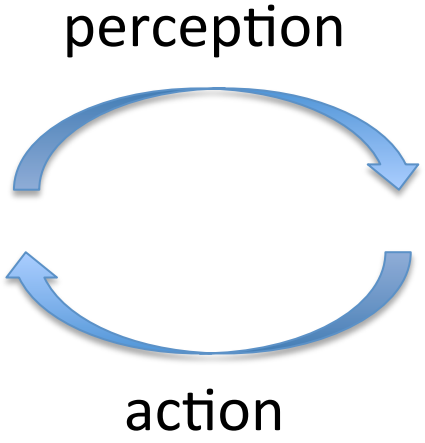

better action selection

M F Land & D-E Nilson, (2002)
Animal Eyes. Oxford.

action selection

movement planning

decision

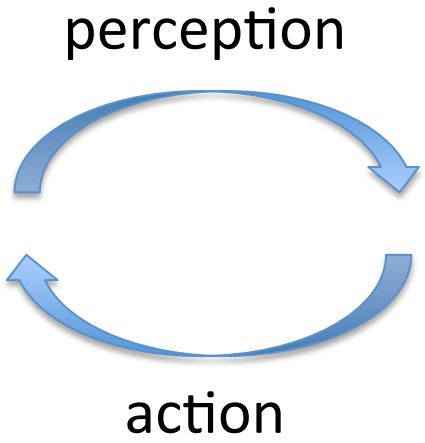

500 million years later …

Every life is a series of decisions ….

Anatomy of a Decision

Feasibility

Probability, Frequency,

Uncertainty, etc.

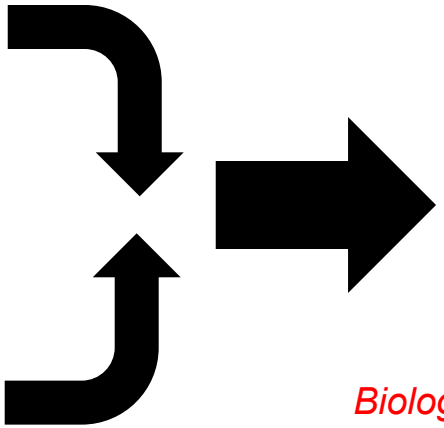

Action

Desireability

Reward, Utility, etc

Biology
Psychology
Economics
Political Science
Neuroscience

Stochastic Form

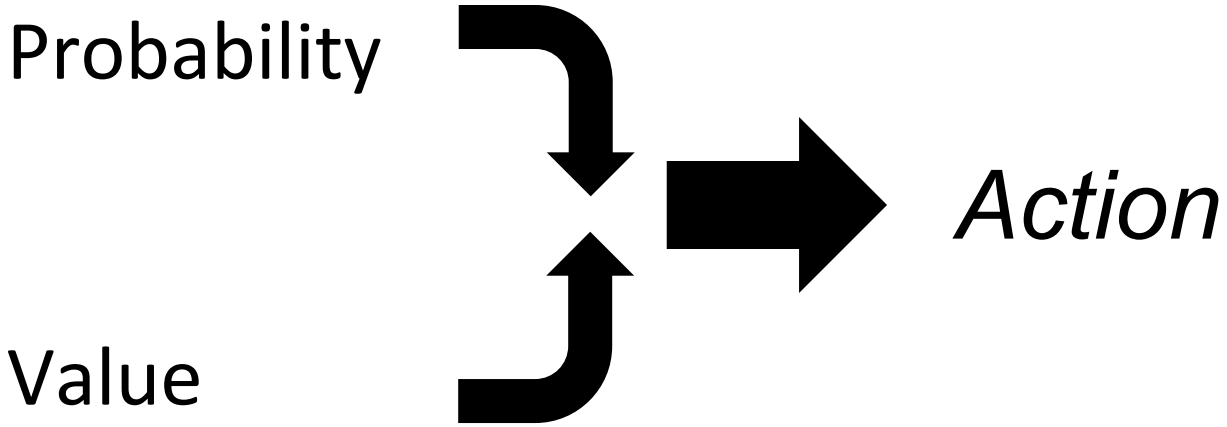

Statistical Decision Theory

Abraham Wald

John von Neumann

David Blackwell

Oskar Morgenstern

M. A. Girschick

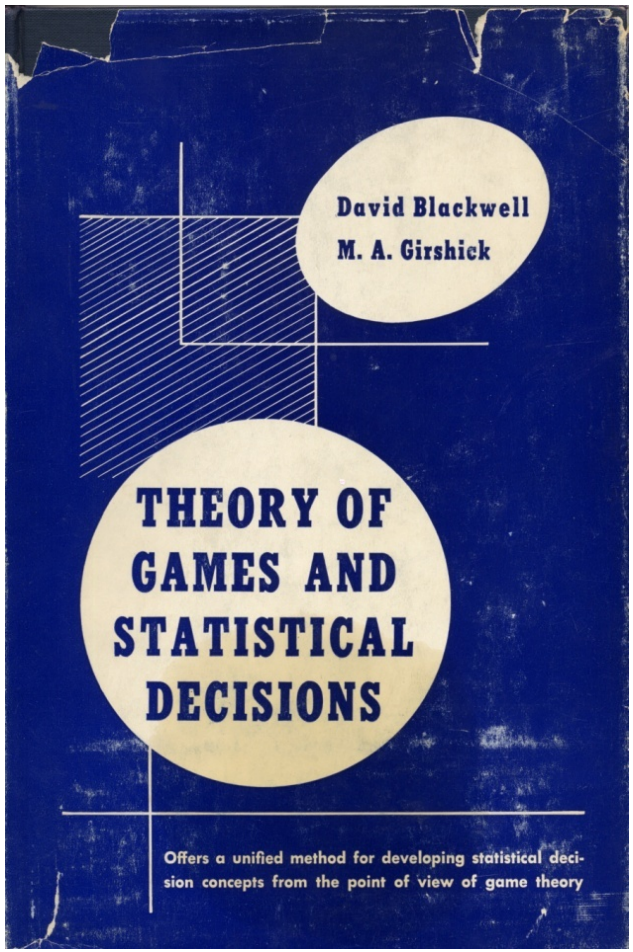

1954

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*

Three Functions of SDT

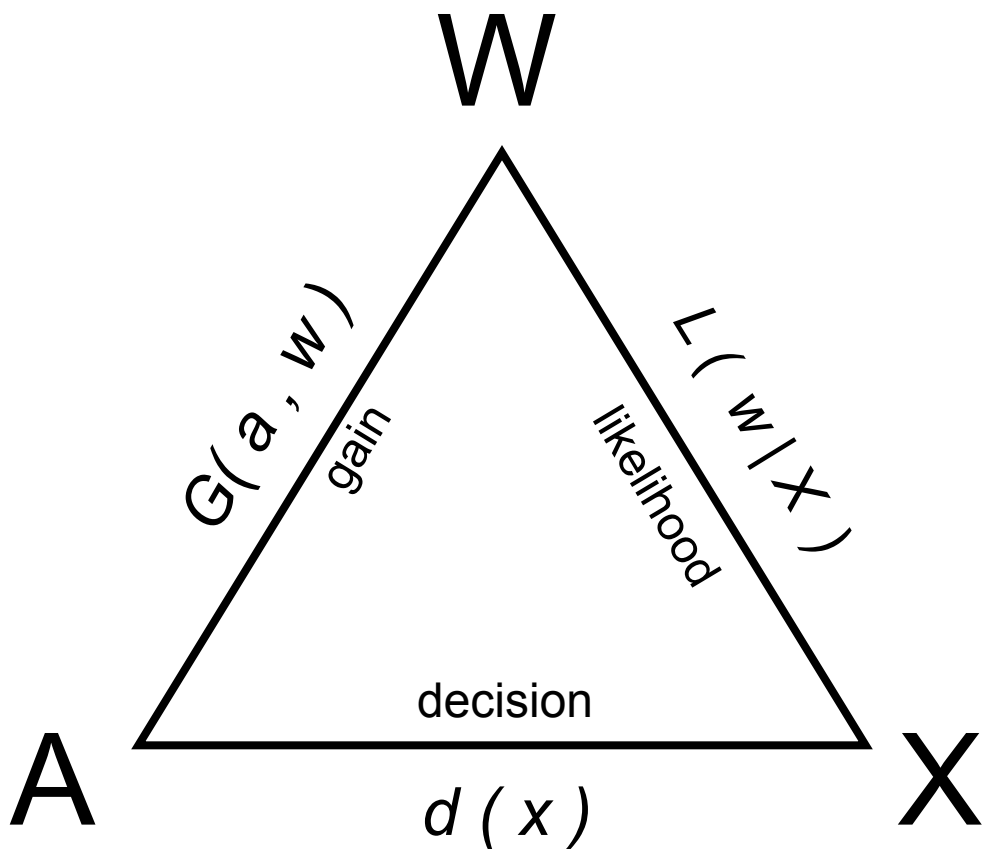

Three Functions of SDT

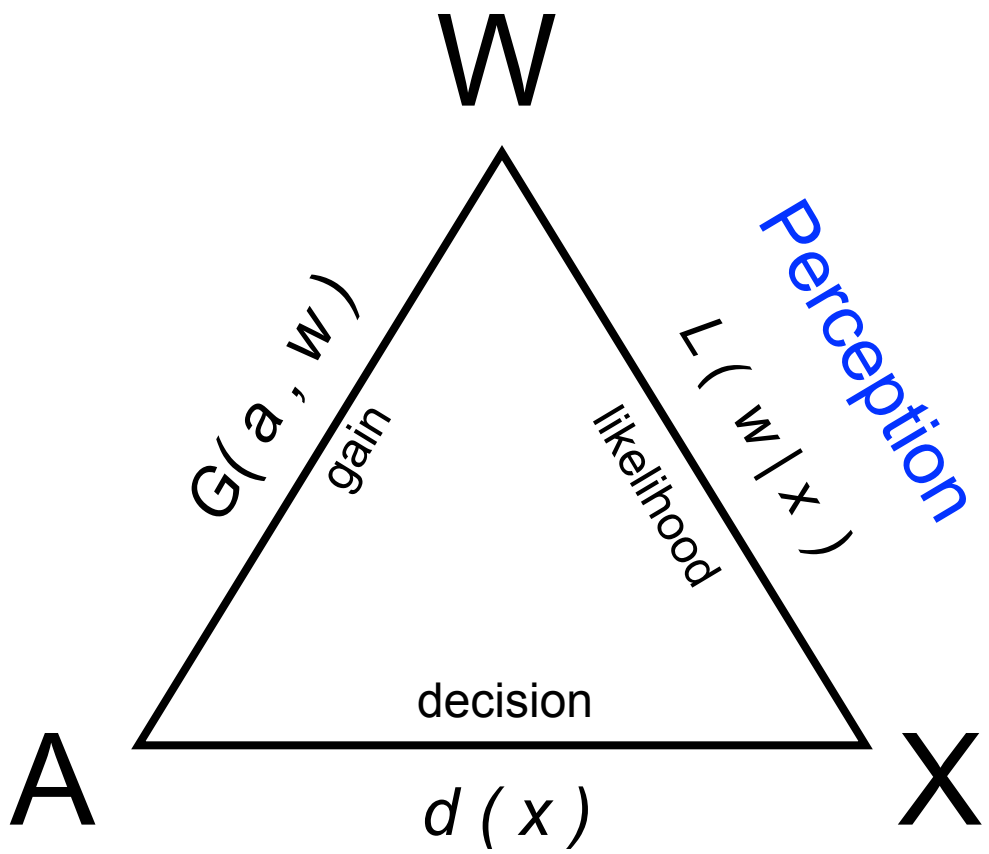

Three Functions of SDT

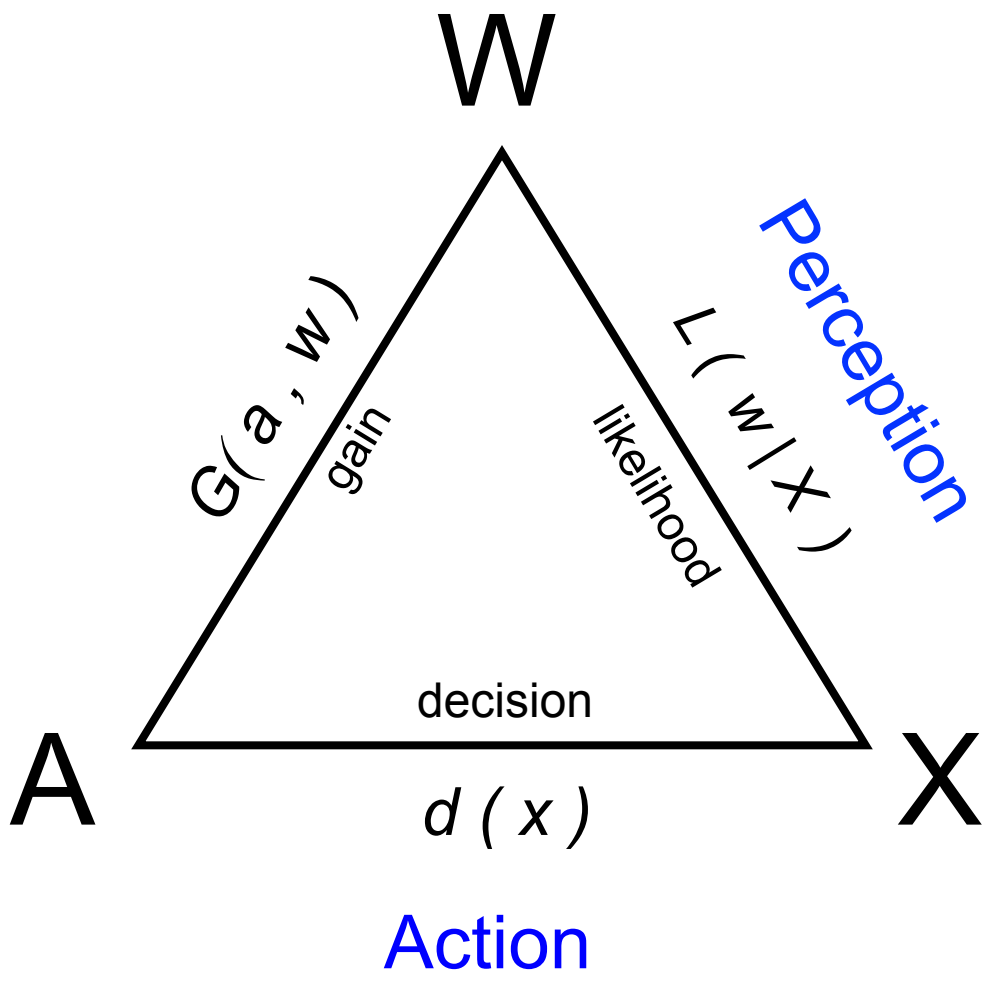

Statistical Decision Theory

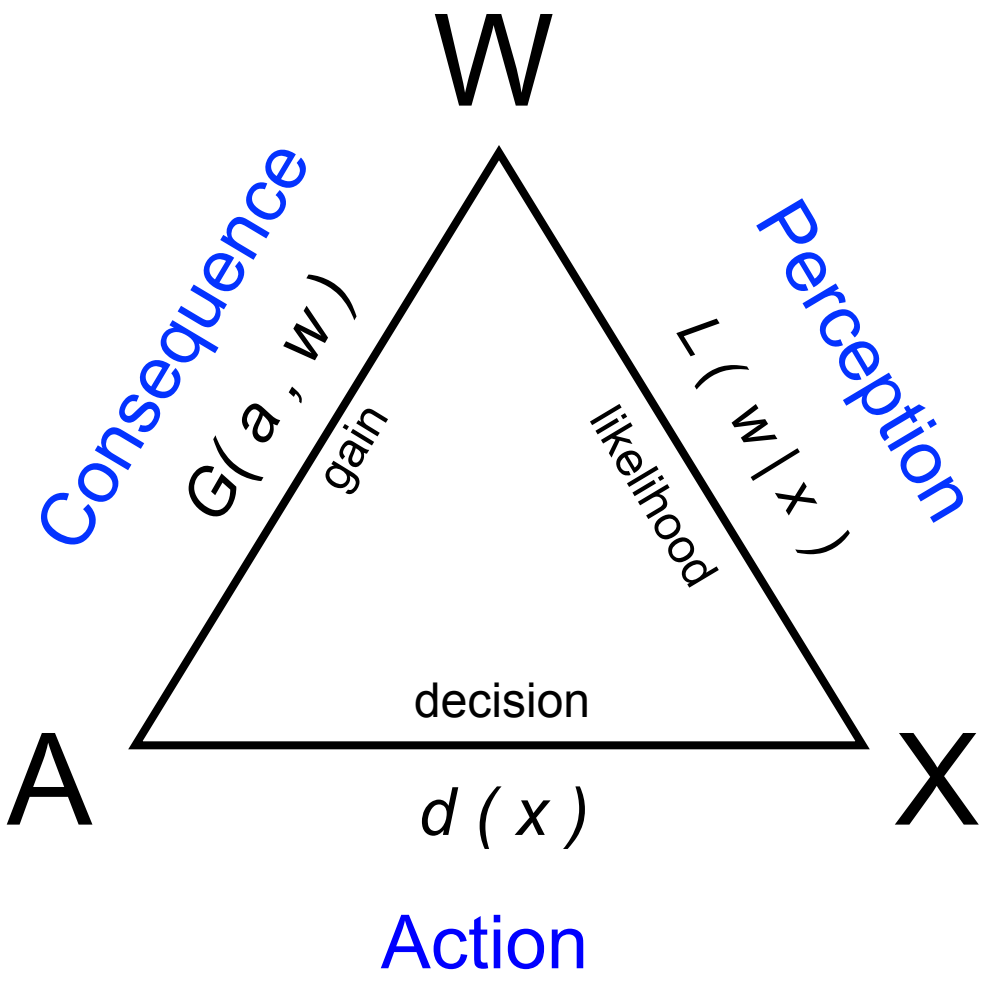

Statistical Decision Theory

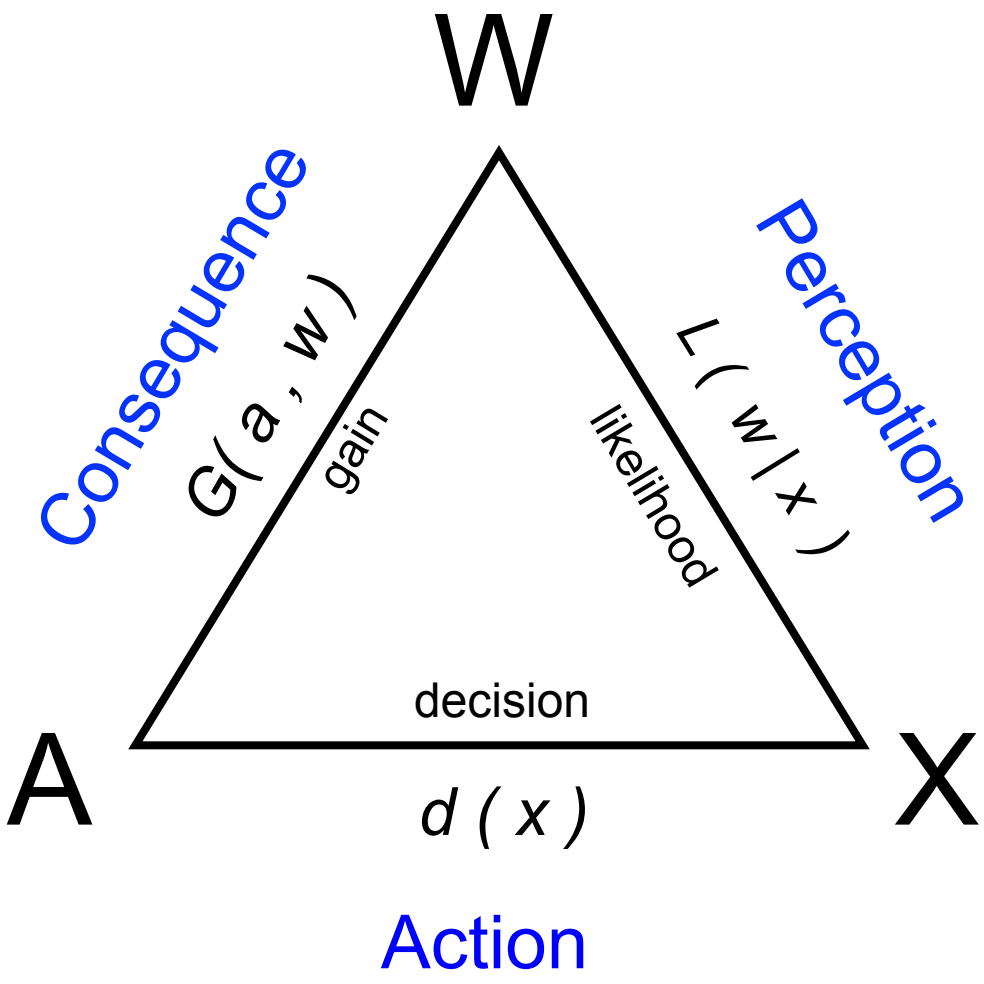

Bayesian Decision Theory

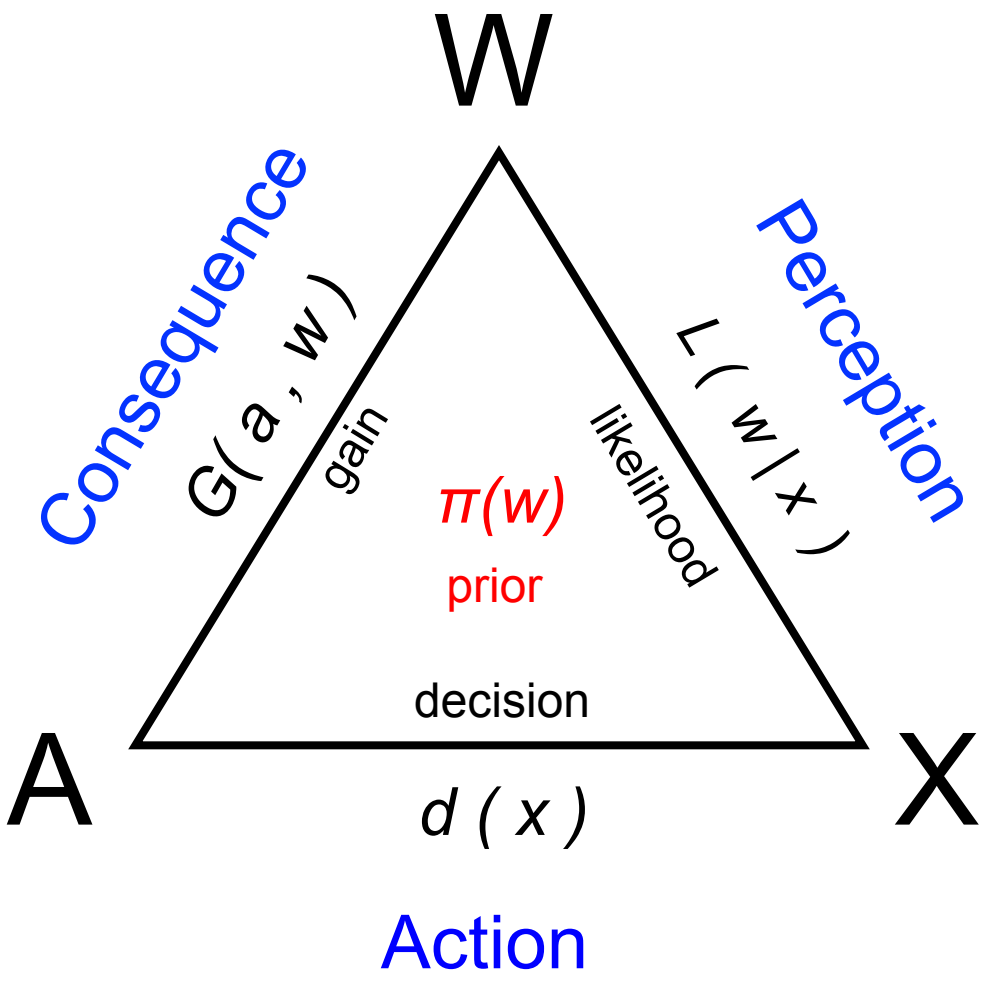

A Bayesian Problem

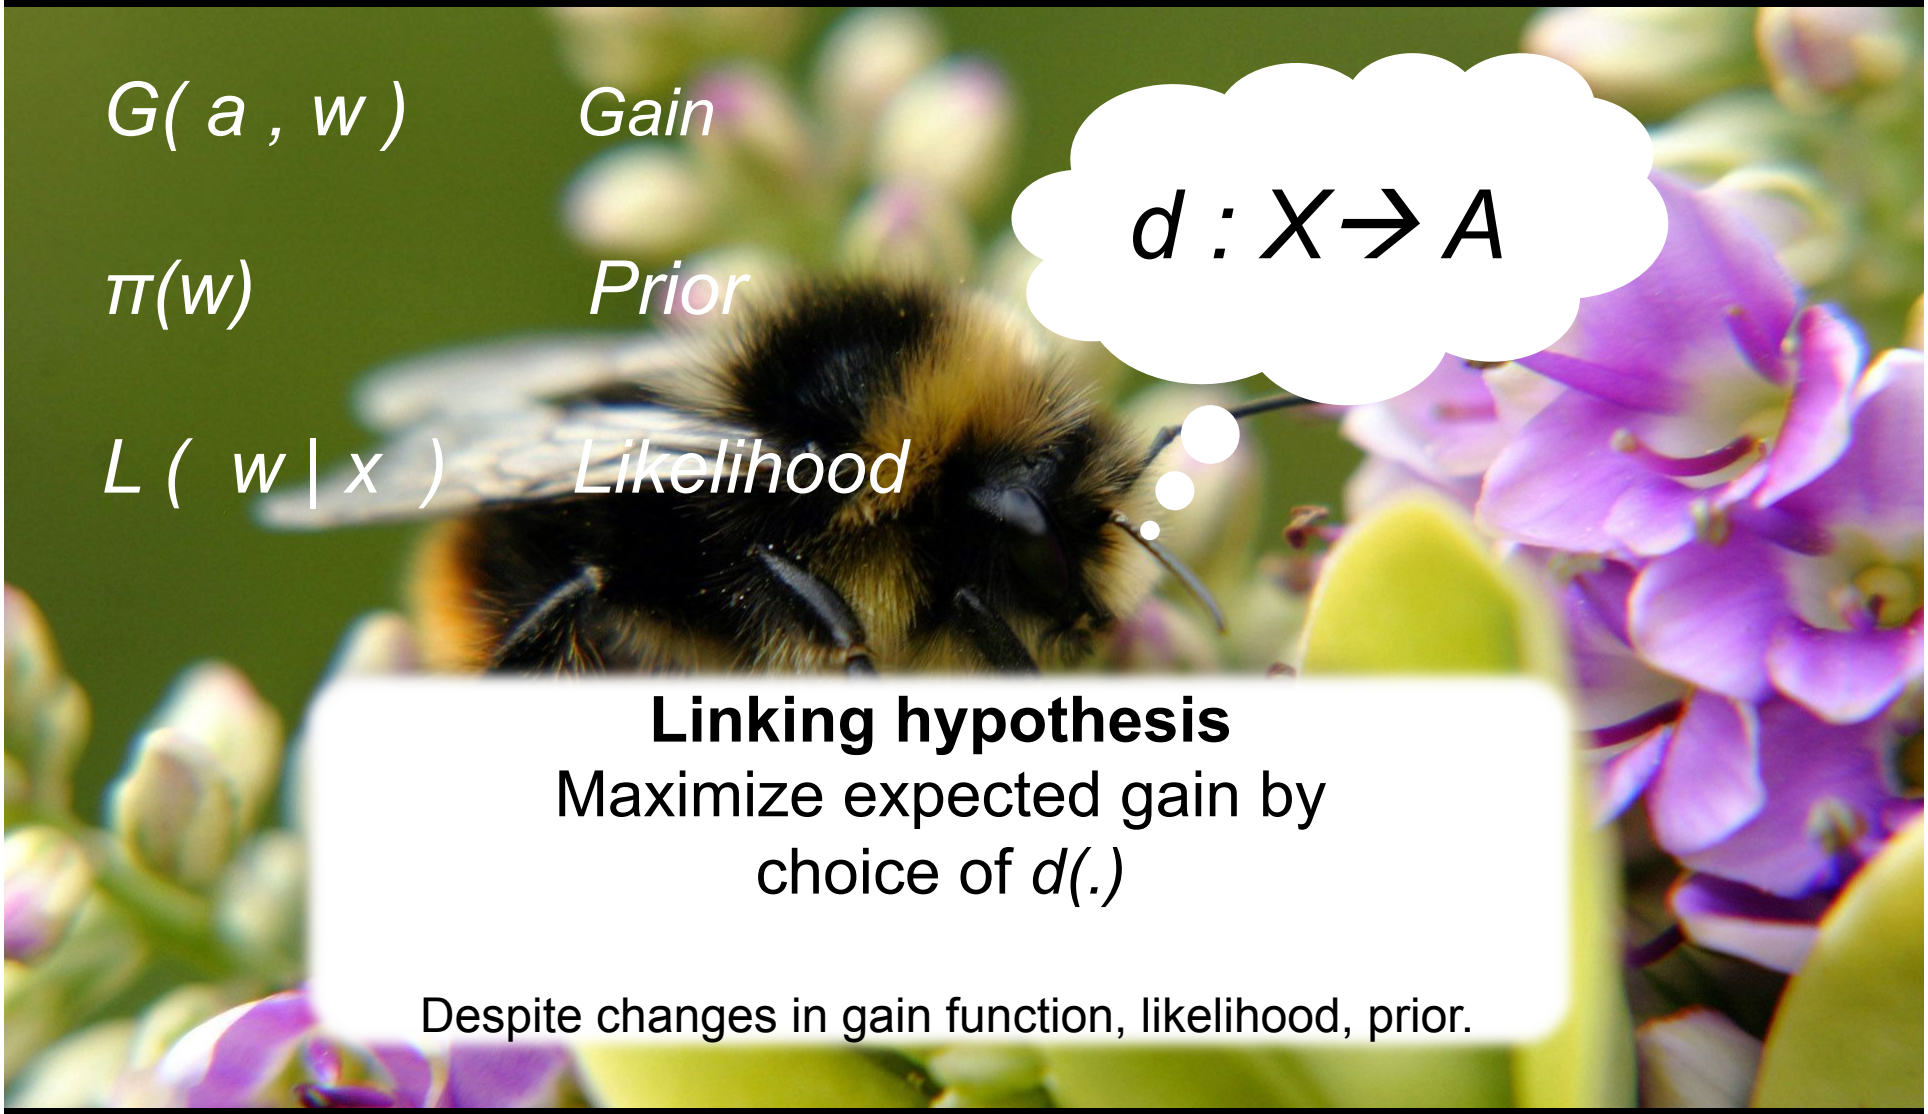

$G(a, w)$ Gain

$\pi(w)$ Prior

$L(w \mid x)$ Likelihood

$d : X \rightarrow A$

Linking hypothesis
Maximize expected gain by choice of *d(.)*

Despite changes in gain function, likelihood, prior.

A Bayesian Game

:

speeded reaching

A Bayesian Problem

Statistical decision theory and the selection of rapid, goal-directed movements

Julia Trommershäuser, Laurence T. Maloney, and Michael S. Landy

Department of Psychology and Center for Neural Science, New York University, New York, New York 10003

Received October 1, 2002; revised manuscript received January 30, 2003; accepted February 3, 2003

We present two experiments that test the range of applicability of a movement planning model (MEGaMove) based on statistical decision theory. Subjects attempted to earn money by rapidly touching a green target region on a computer screen while avoiding nearby red penalty regions. In two experiments we varied the magnitudes of penalties, the degree of overlap of target and penalty regions, and the number of penalty regions. Overall, subjects acted so as to maximize gain in a wide variety of stimulus configurations, in good agreement with predictions of the model.

OCIS codes: 330.4060, 330.7310.

1. INTRODUCTION

Motor responses have consequences. In the first semifinal of the 2002 World Cup, Germany met host South Korea, and the match was decided in the 75th minute by the only goal of the match. Oliver Neuville passed the ball to Germany's play maker Michael Ballack, who scored in his

small target region, the subject cannot be certain that a movement aimed at the center of the reward region will not, instead, end up outside the reward region, possibly in the penalty region. To summarize, the subject, in each trial, must select among possible actions and must do so very rapidly. There are explicit monetary penalties associated with the outcome of the action selected, but the

Experimental Task

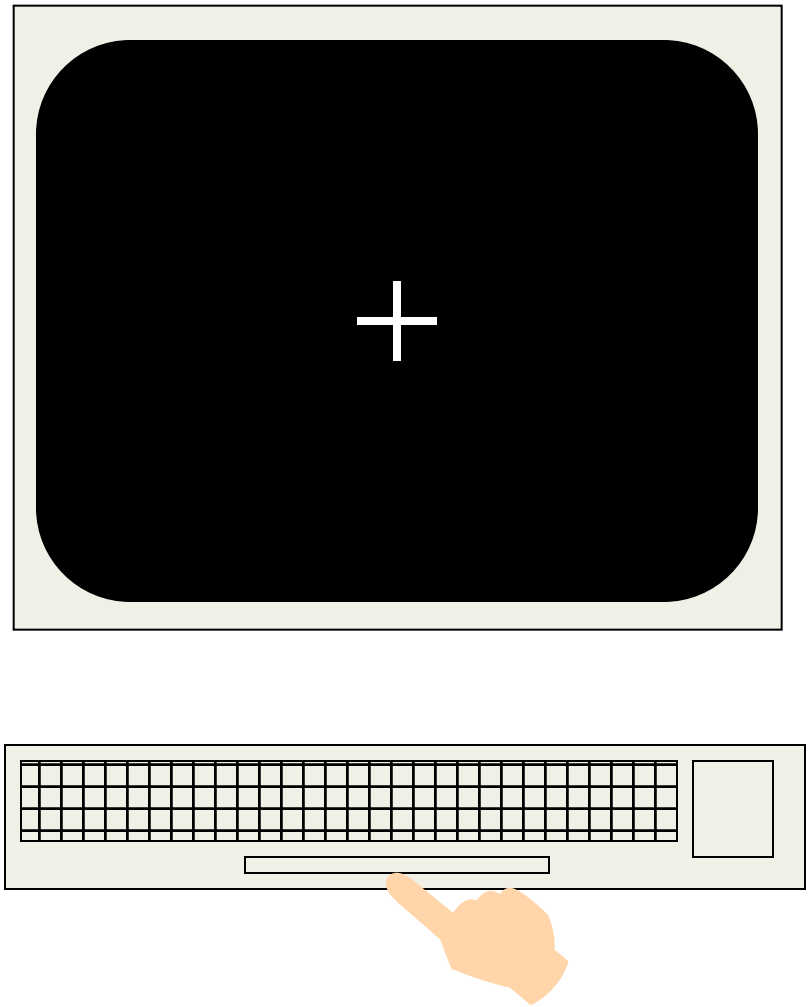

Start of trial:
display of fixation
cross (1.5 s)

Experimental Task

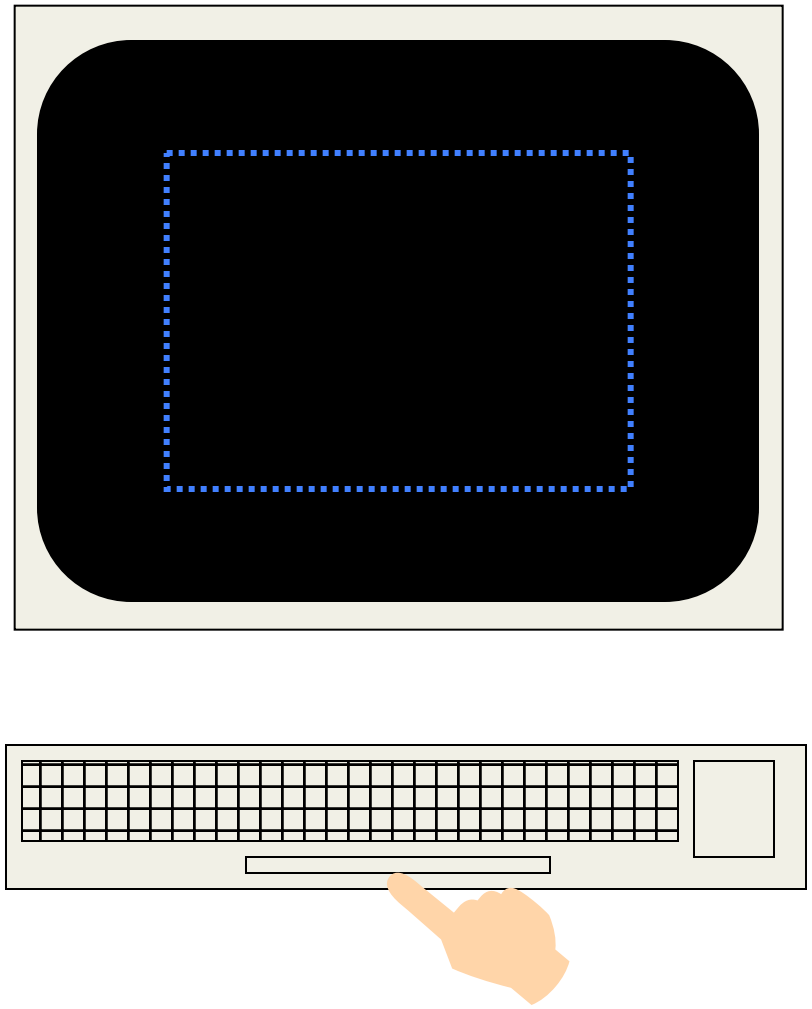

Display of response area,
500 ms before
target onset
(114.2 mm x 80.6 mm)

Experimental Task

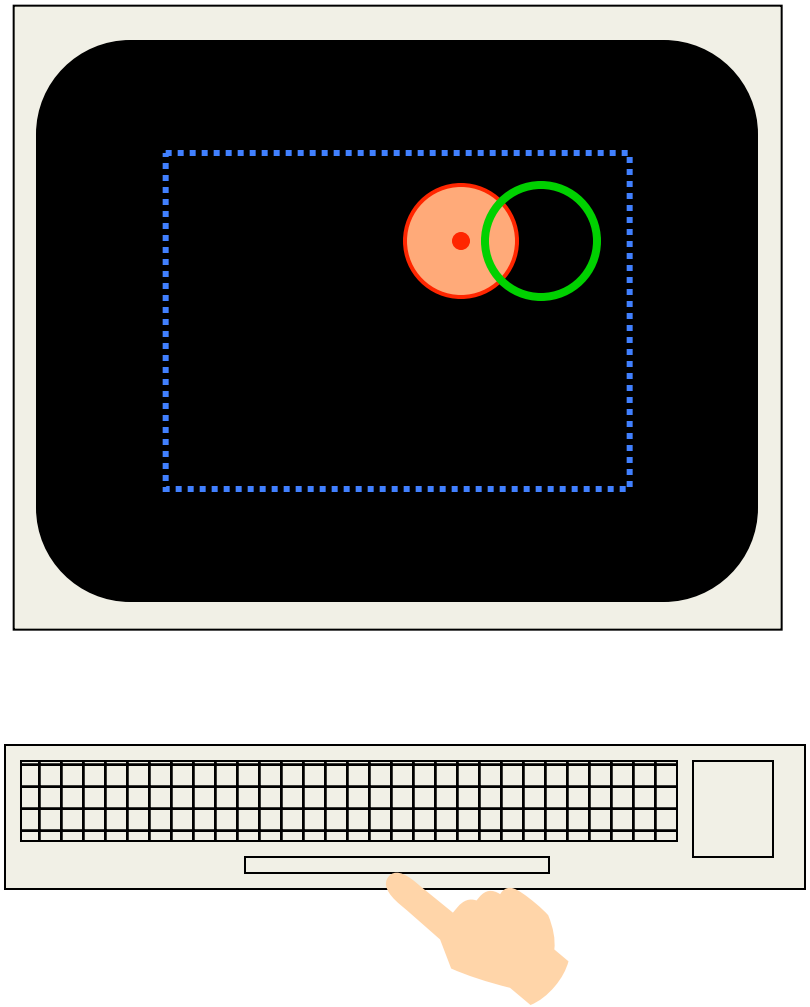

Target display (700 ms)

Experimental Task

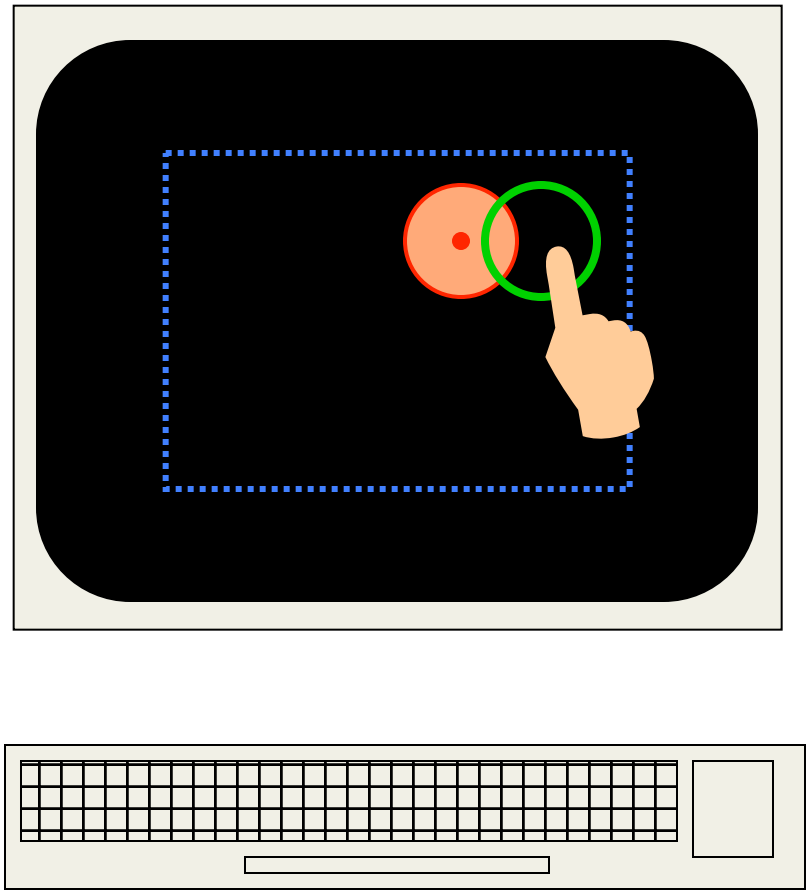

Experimental Task

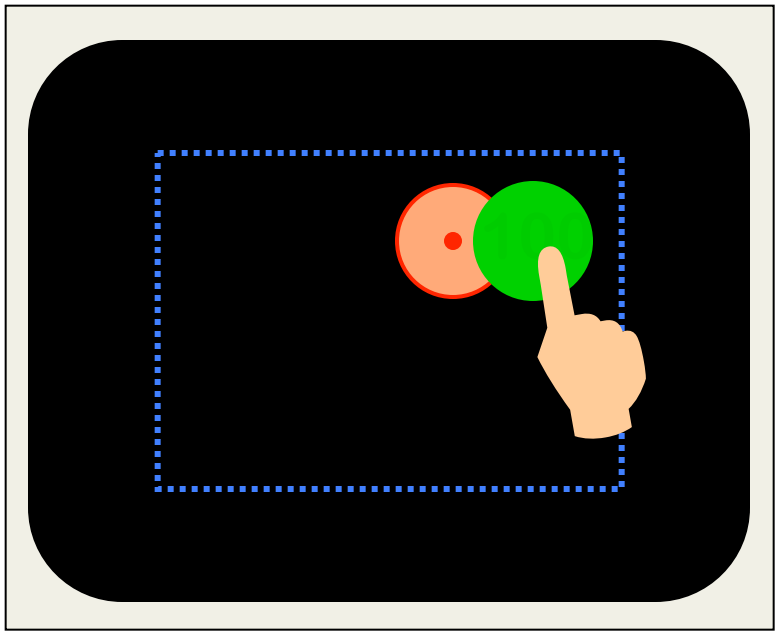

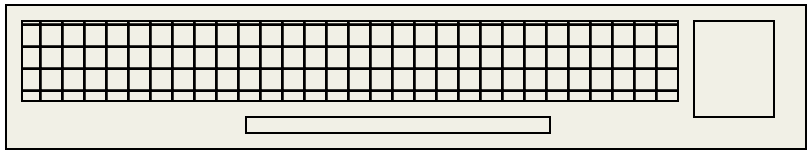

The green target is hit:
+100 points

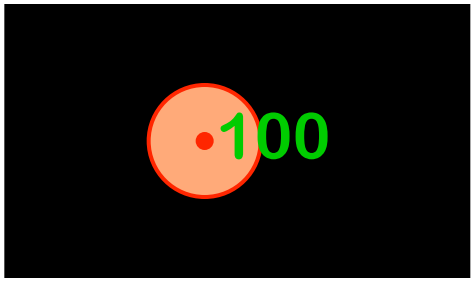

Experimental Task

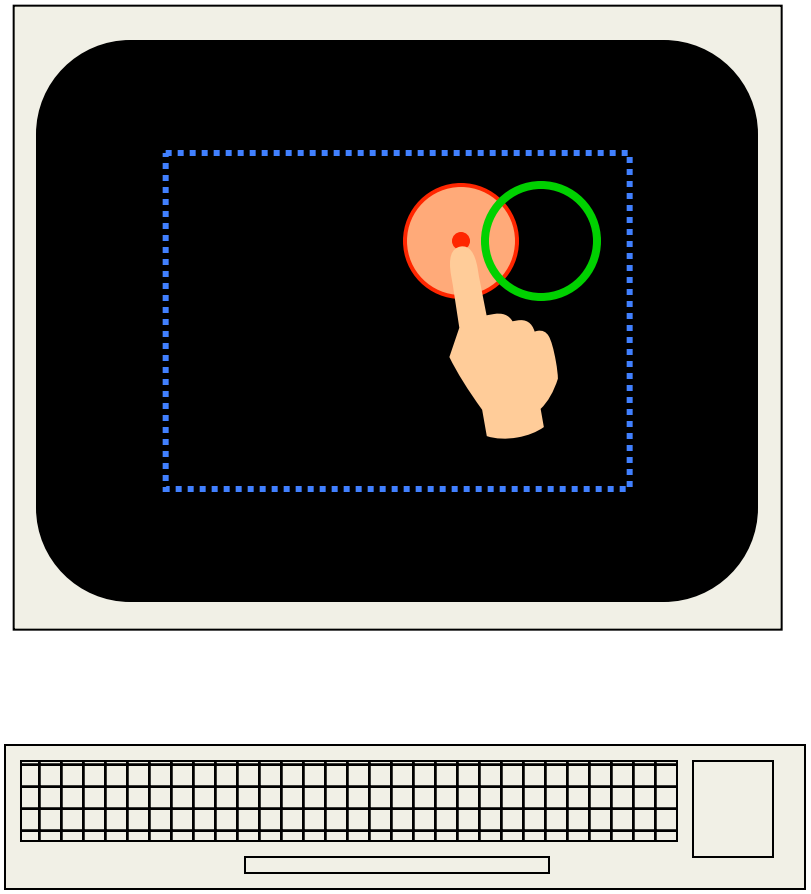

Experimental Task

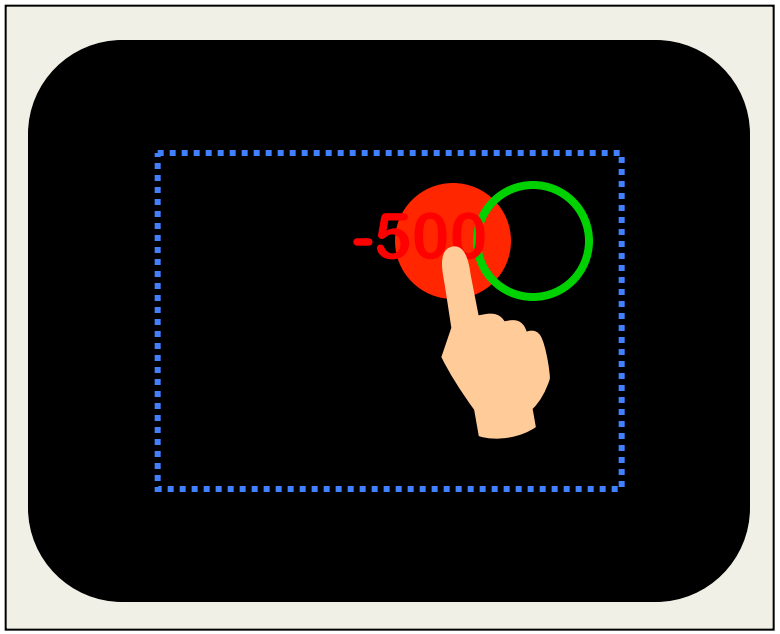

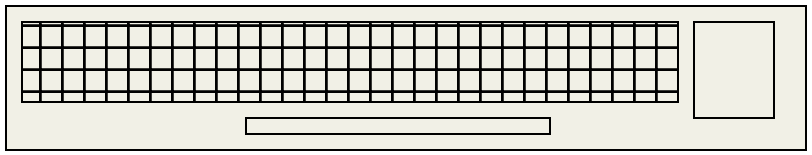

The red target is hit:
-500 points

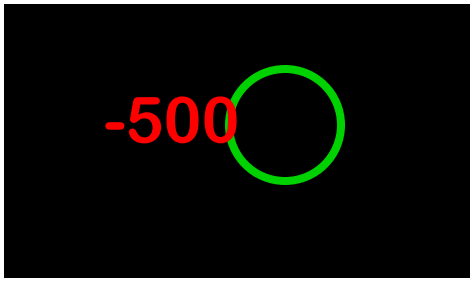

Experimental Task

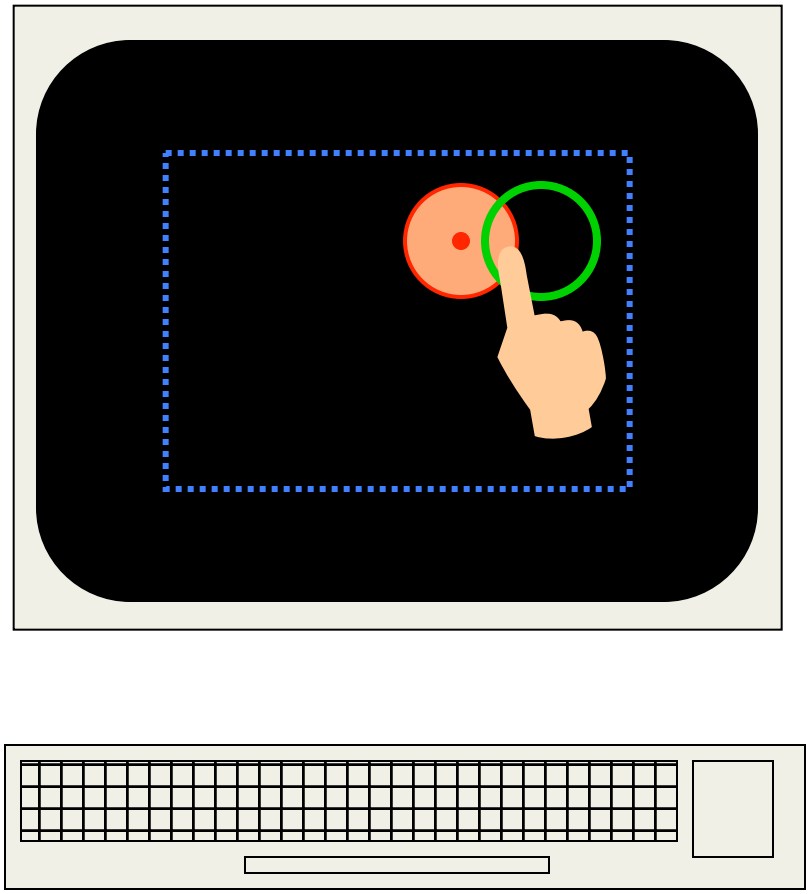

Experimental Task

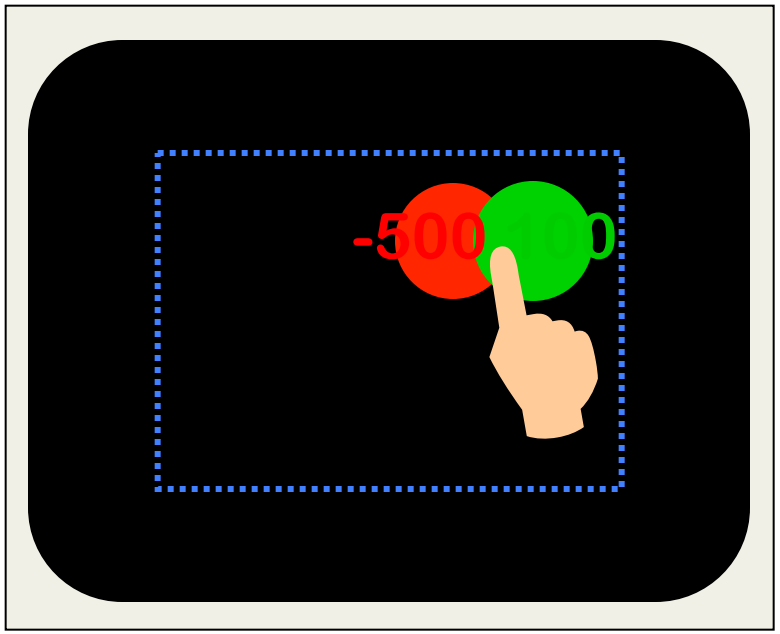

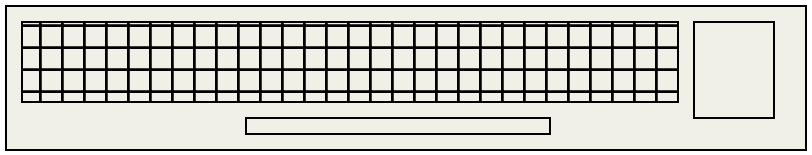

Scores add if both targets are hit:

Experimental Task

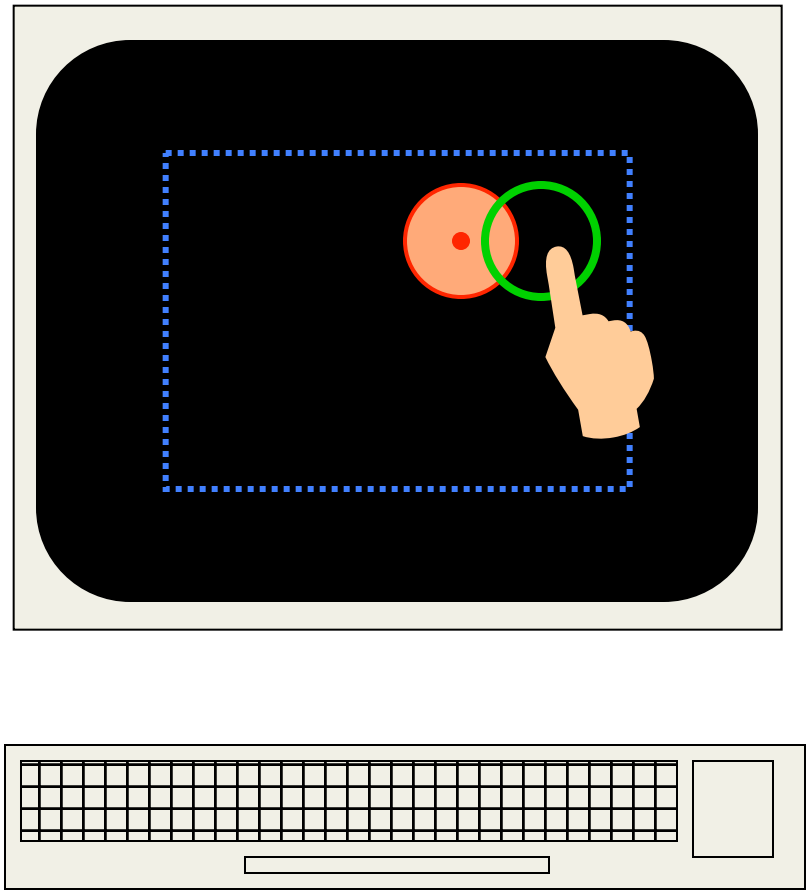

Experimental Task

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

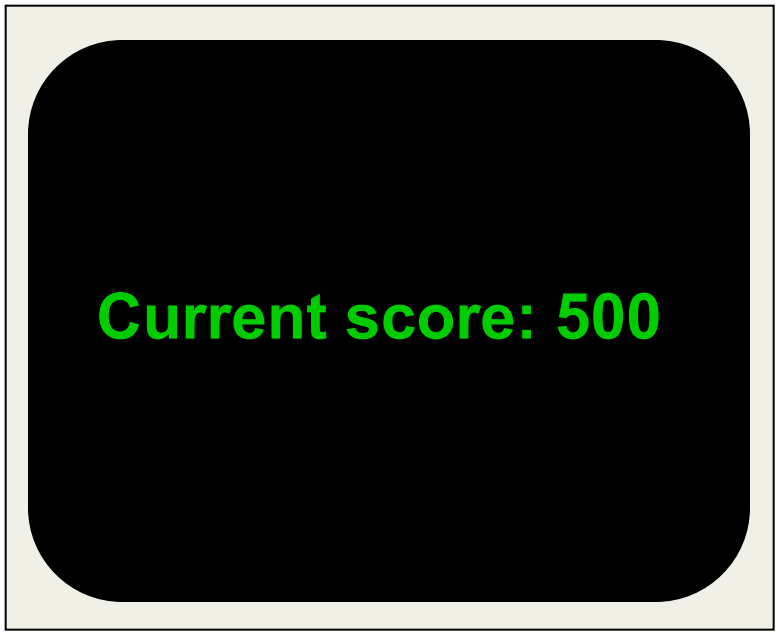

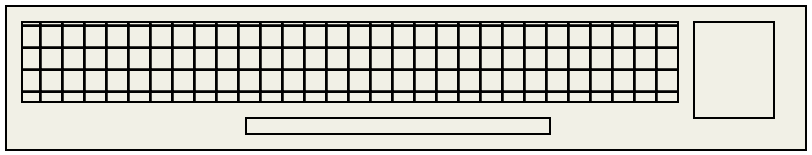

End of trial

Choice among Movement Strategies

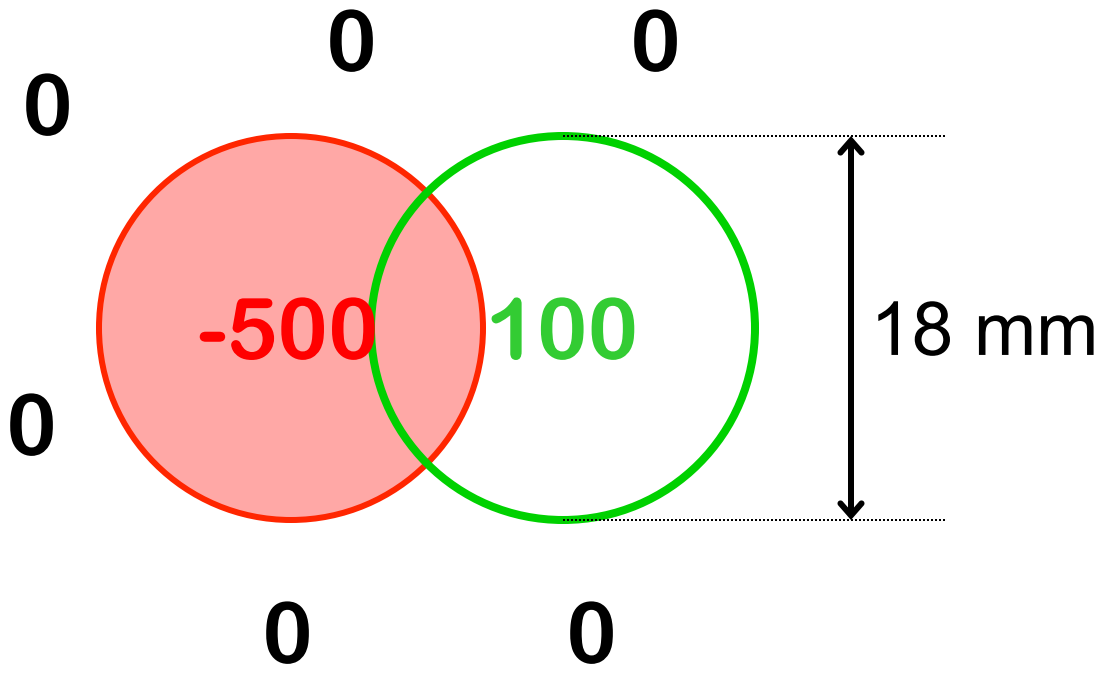

What should Paulina do?

Distribution of movement end points

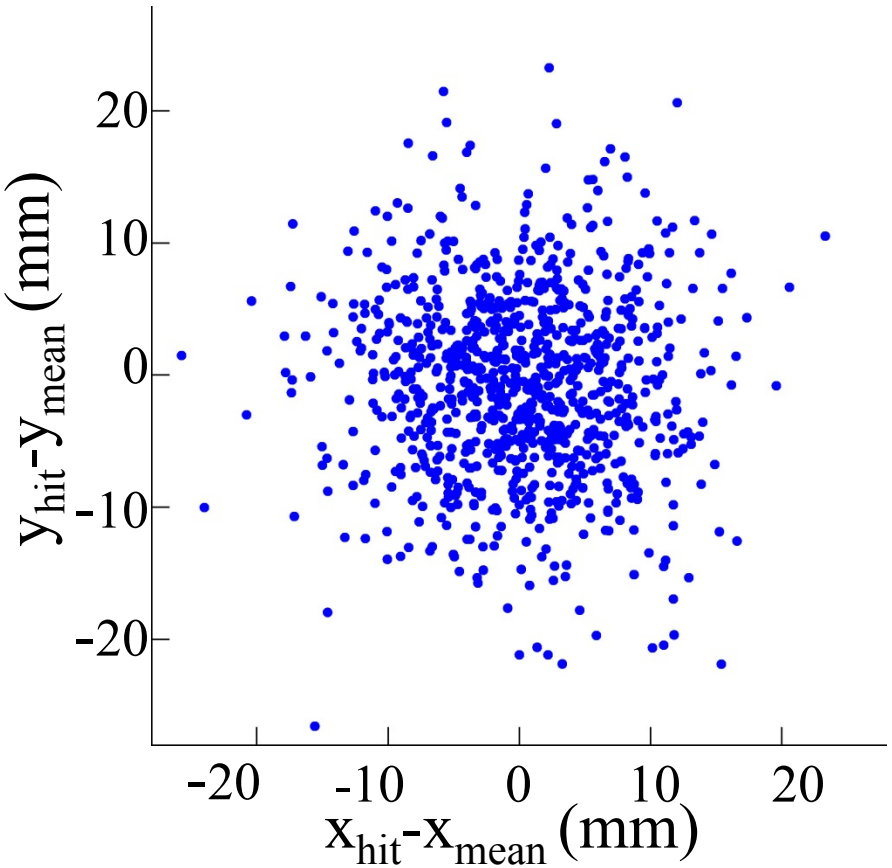

Q-Q Plot

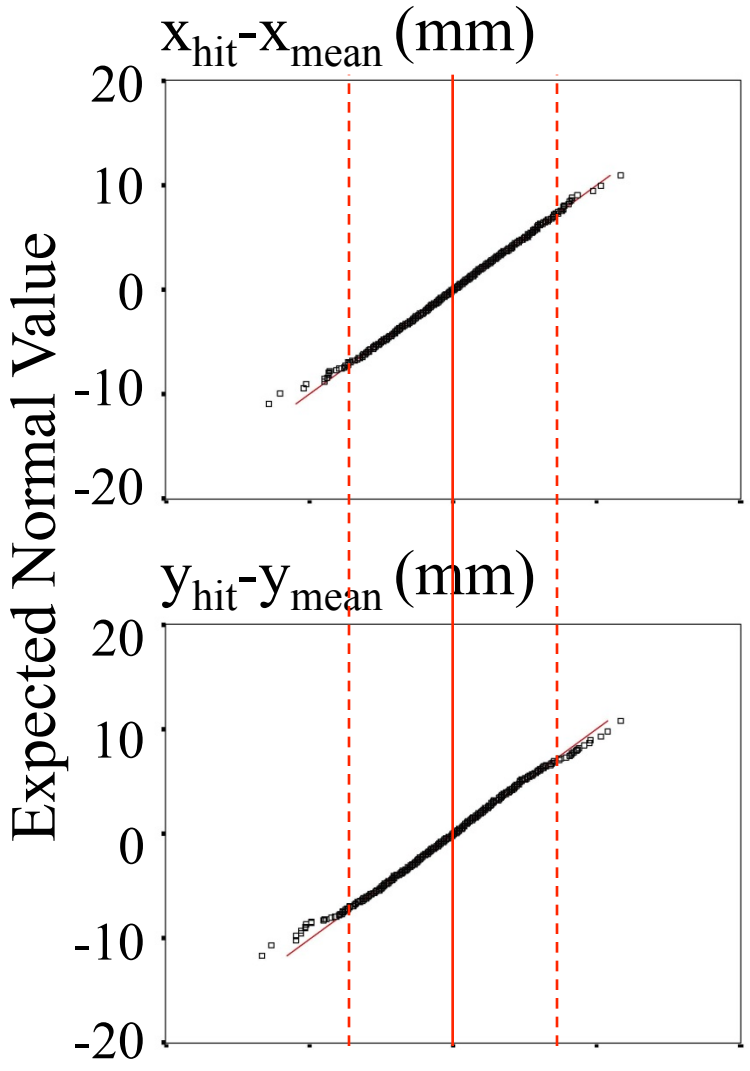

Subject S4, [illegible] = 3.62 mm,
72x15 = 1080 end points

If there were no red penalty circle ….

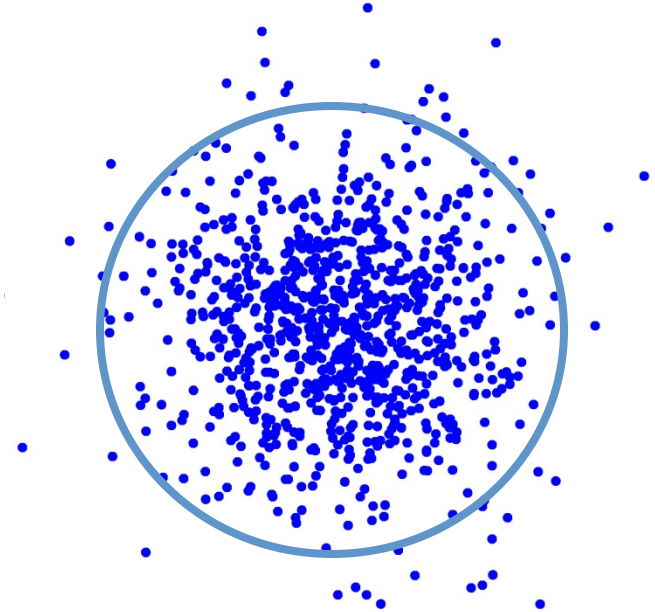

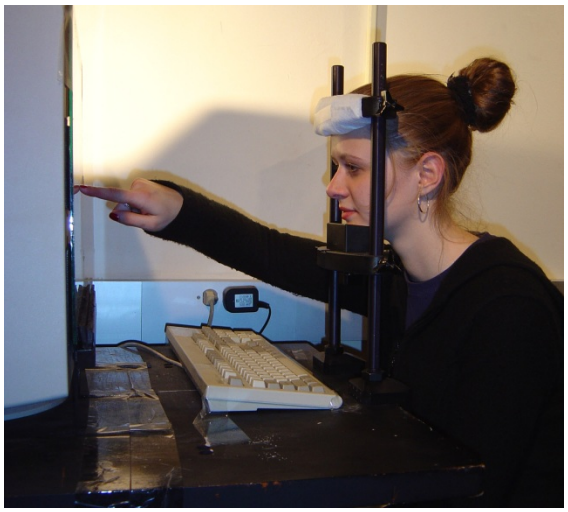

Aim for center
Select perceptual-motor strategies that
minimize variance

Harris & Wolpert (1998)

Choice among Movement Strategies

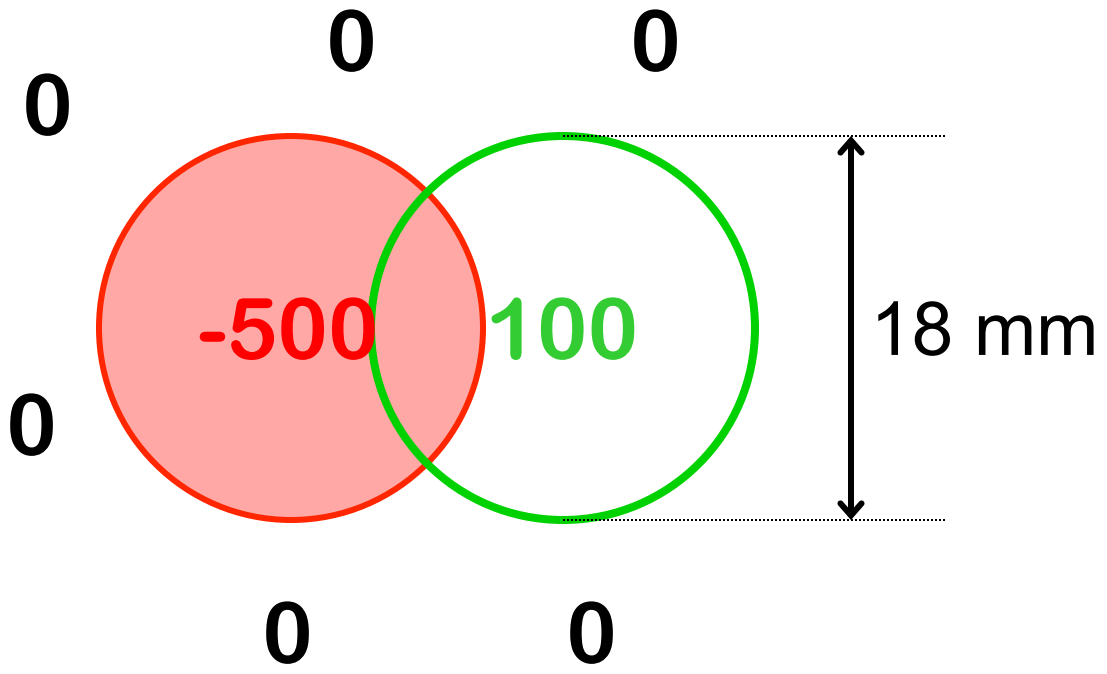

What should Paulina do?

Thought Experiment

: -500 : 100 points (2.5 ¢)

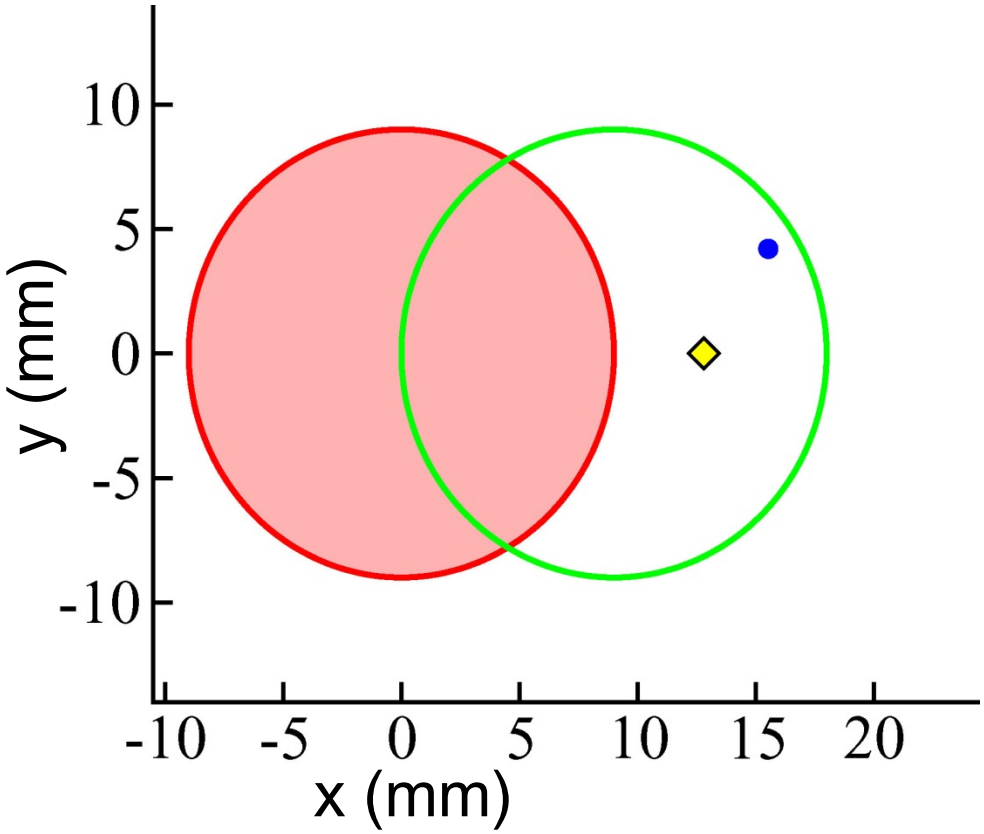

100 points

= 4.83 mm

Thought Experiment

: -500 : 100 points (2.5 ¢)

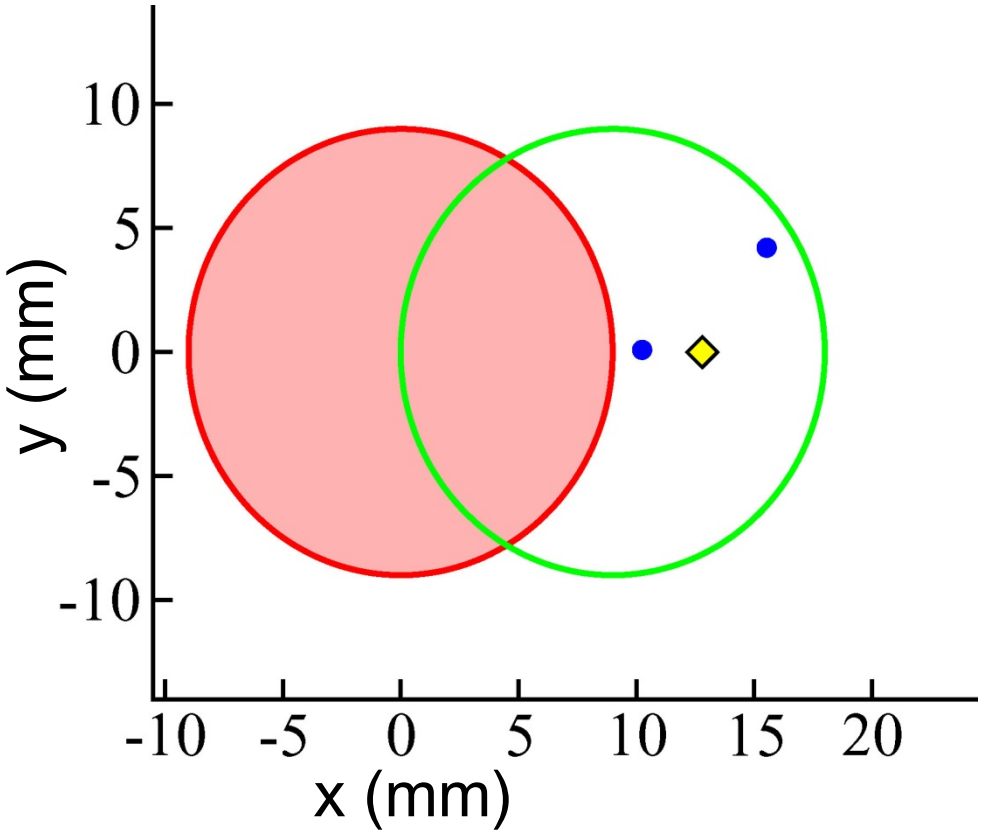

100 points
100 points

200 points

= 4.83 mm

Thought Experiment

: -500 : 100 points (2.5 ¢)

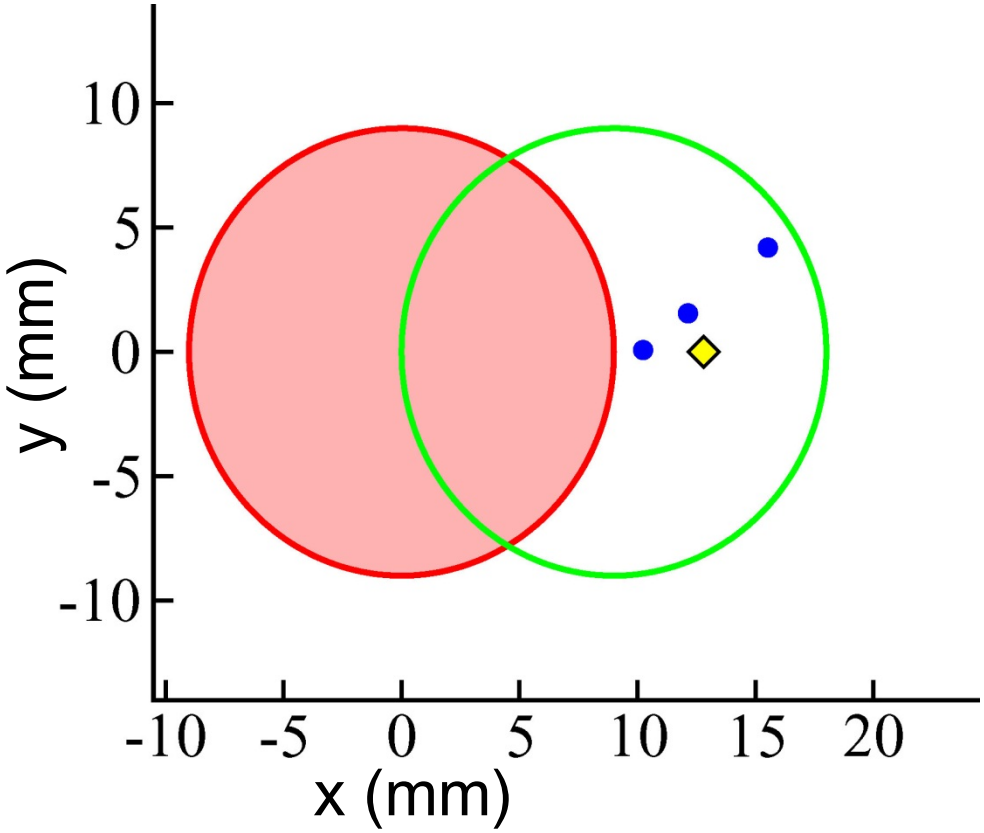

100 points
100 points
100 points

300 points

= 4.83 mm

Thought Experiment

: -500 : 100 points (2.5 ¢)

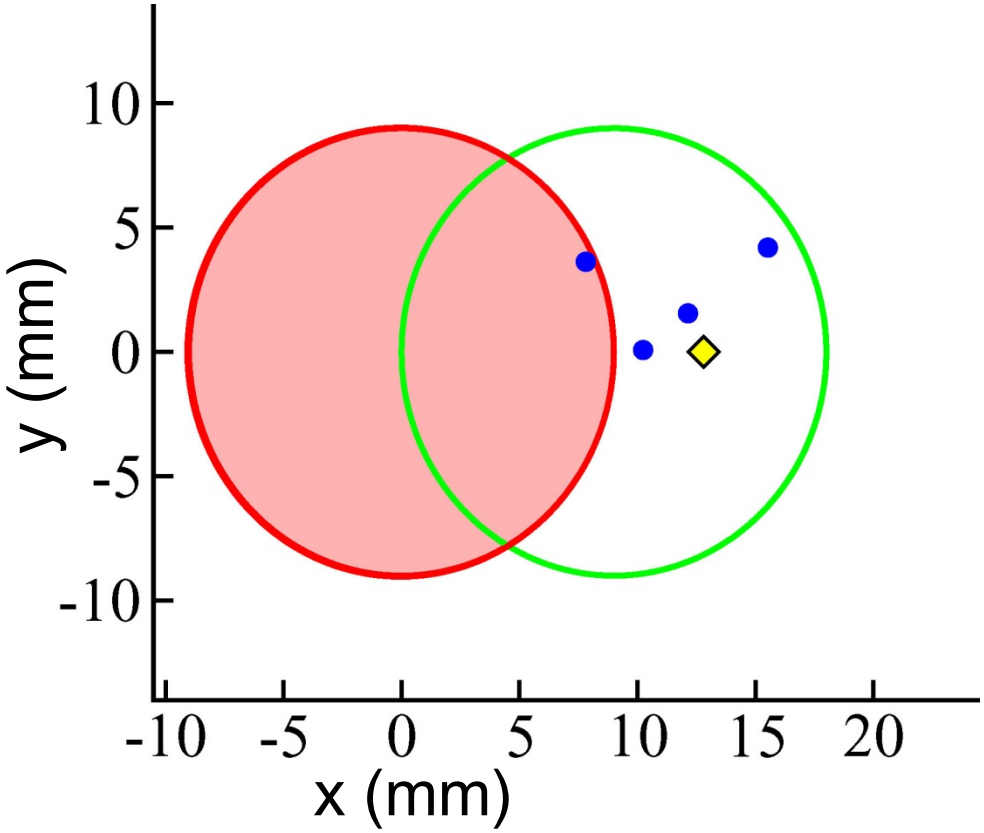

100 points
100 points
100 points
-400 points

-100 points

= 4.83 mm

Thought Experiment

: -500 : 100 points (2.5 ¢)

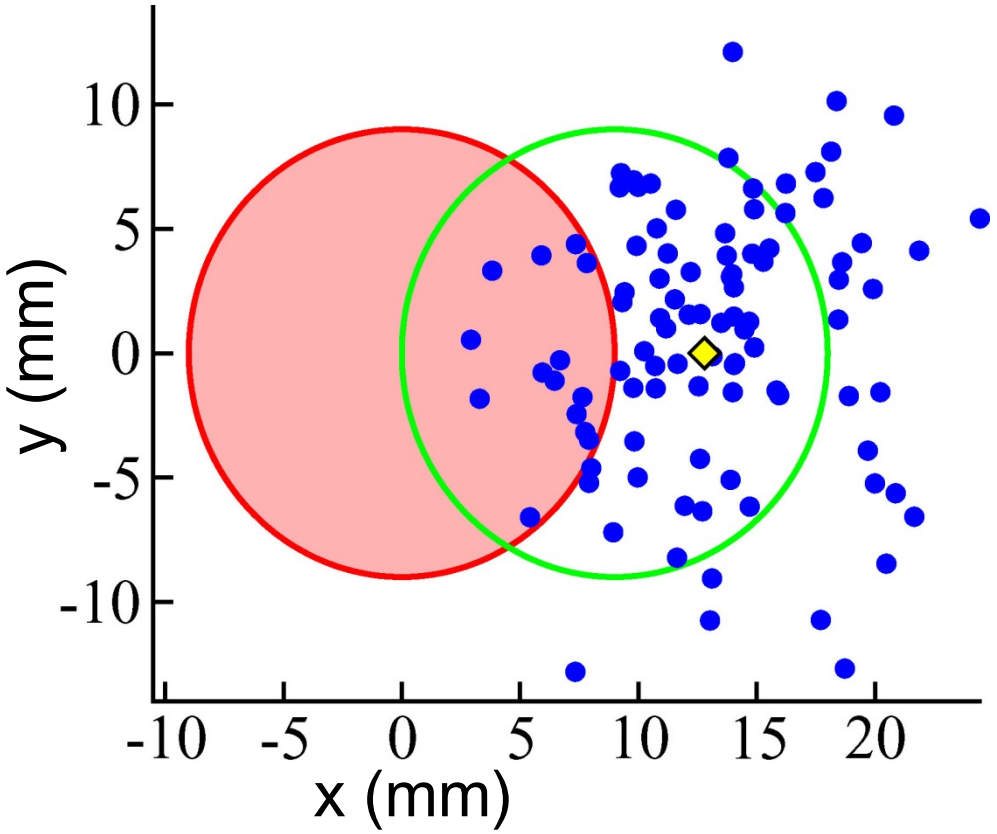

100 points
100 points
100 points
-400 points
.
.
.
.

-0.3 pts. per trial

= 4.83 mm

Thought Experiment

: -500 : 100 points (2.5 ¢)

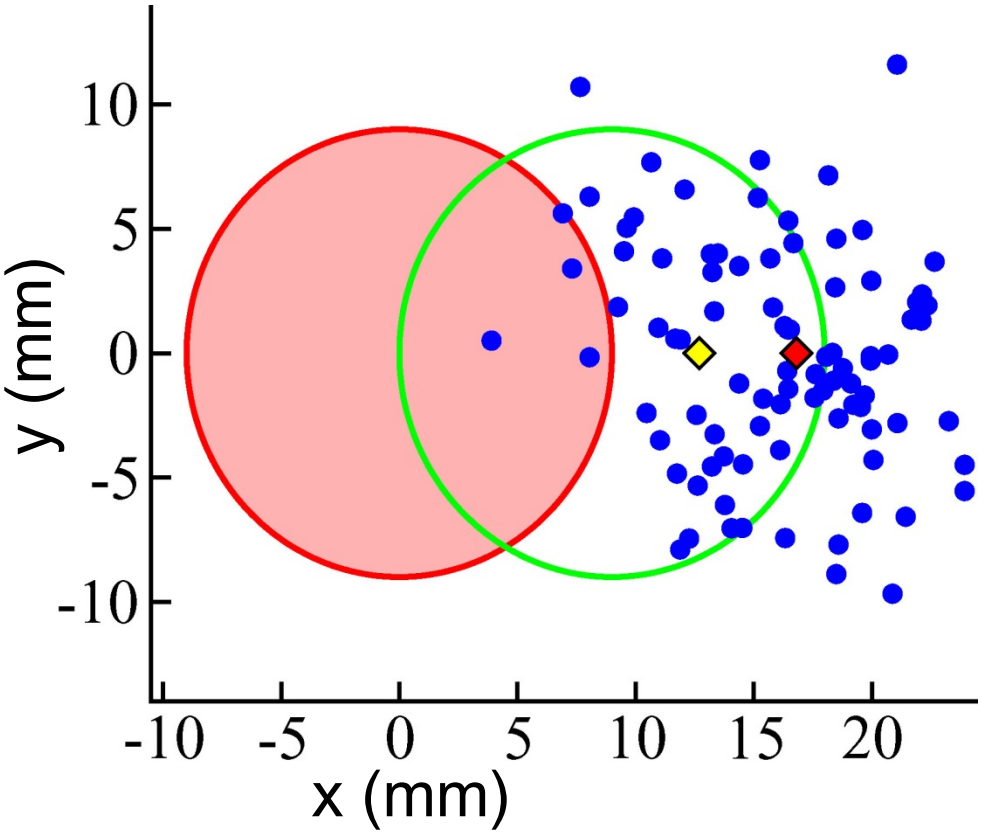

- 0.3 pts. per trial

30.7 pts. per trial

= 4.83 mm

Thought Experiment

: -500 : 100 points (2.5 ¢)

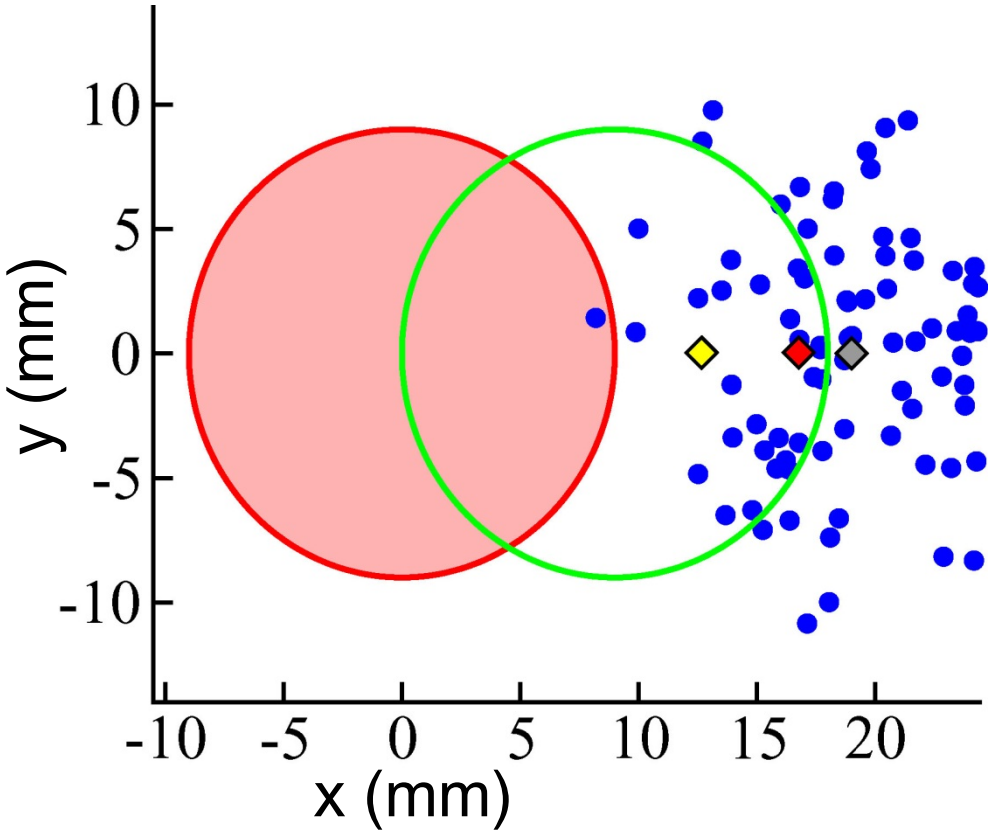

- 0.3 pts. per trial

30.7 pts. per trial

25.5 pts. per trial

= 4.83 mm

Thought Experiment

: -500

: 100 points (2.5 ¢)

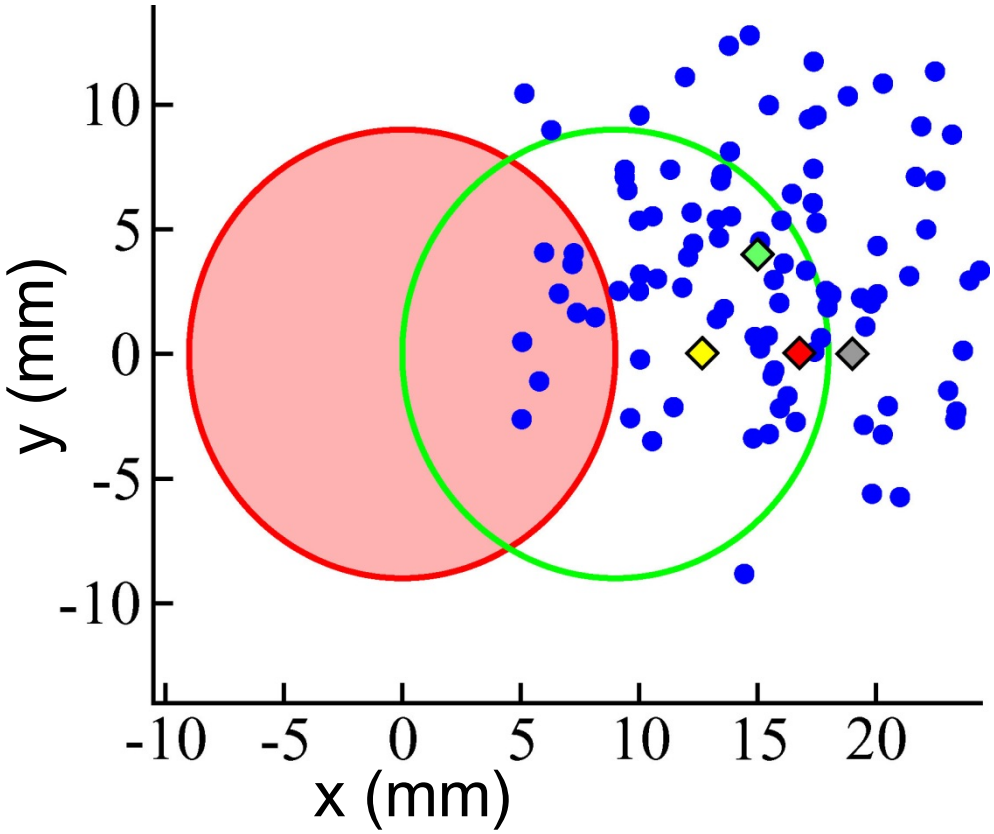

- 0.3 pts. per trial

30.7 pts. per trial

25.5 pts. per trial

22.6 pts. per trial

= 4.83 mm

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

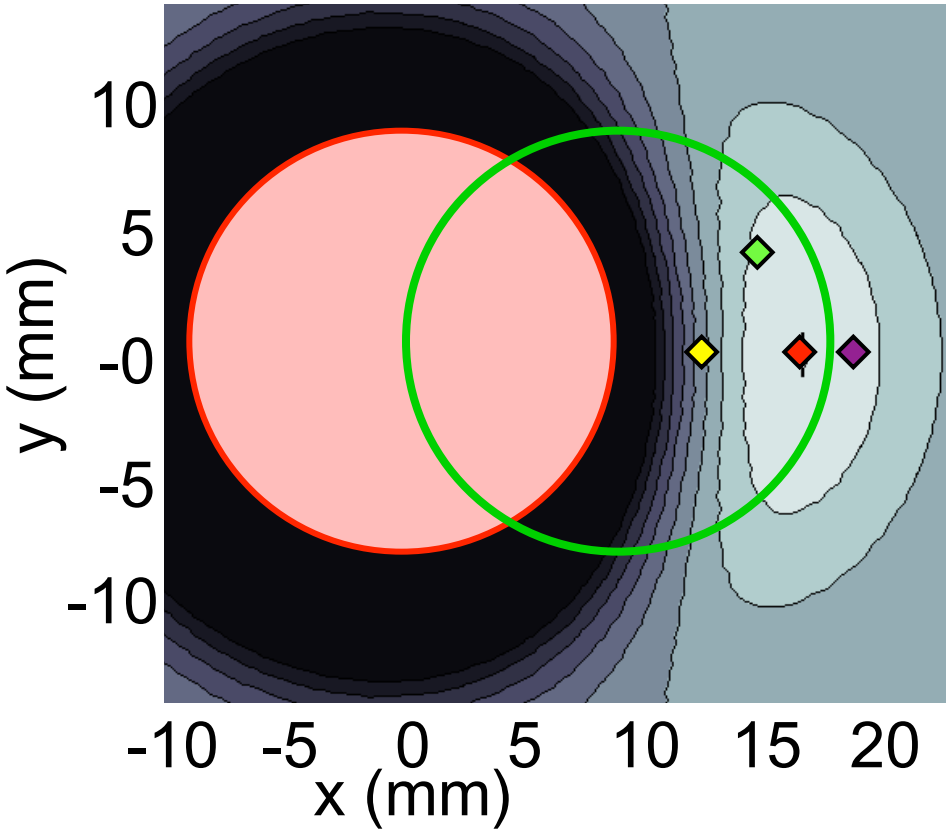

points per trial

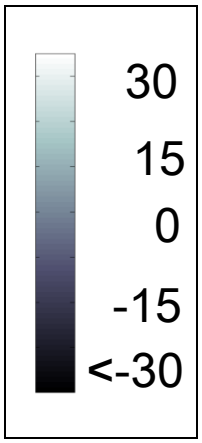

target: 100
penalty: -500

= 4.83 mm

Thought Experiment

penalty: 0 penalty: 100 penalty: 500

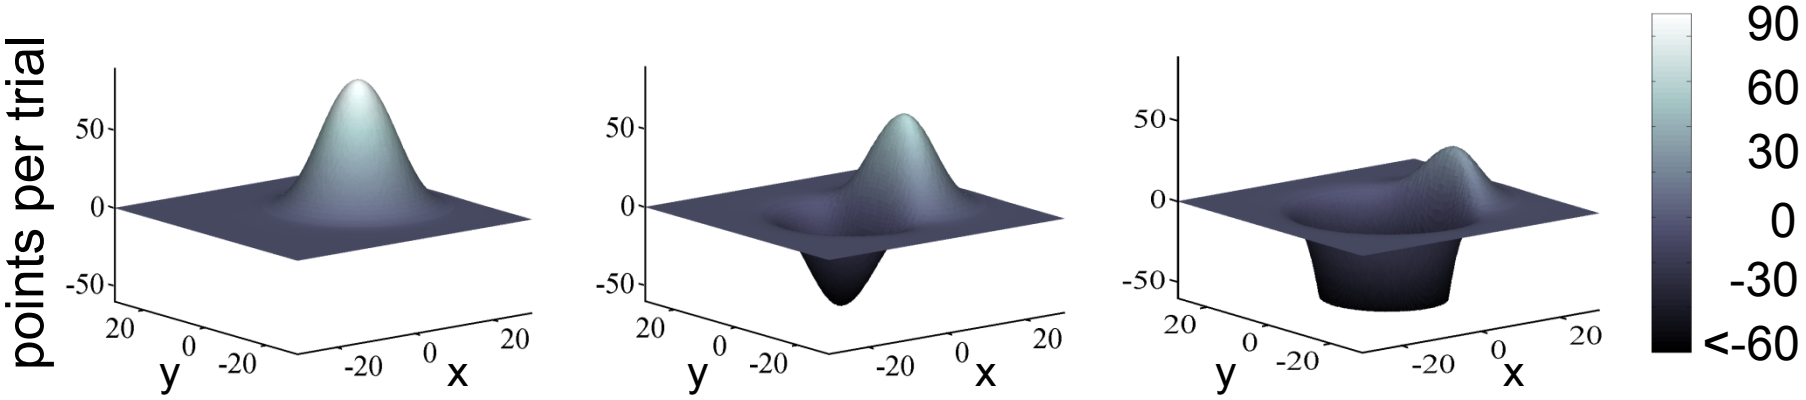

x, y: mean movement end point [mm]

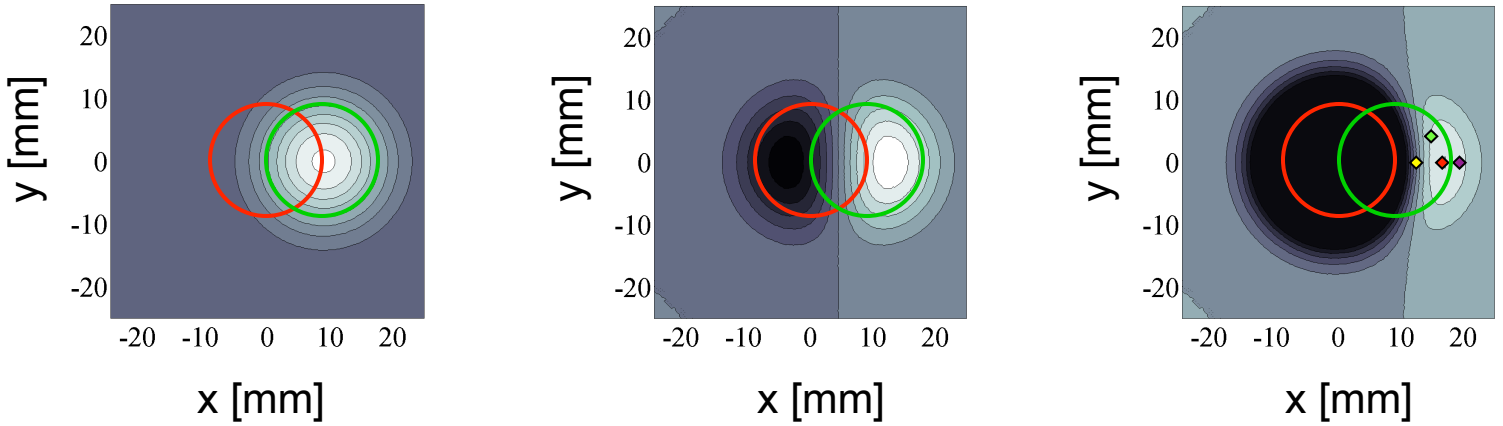

= 4.83 mm

Movement plans as lotteries

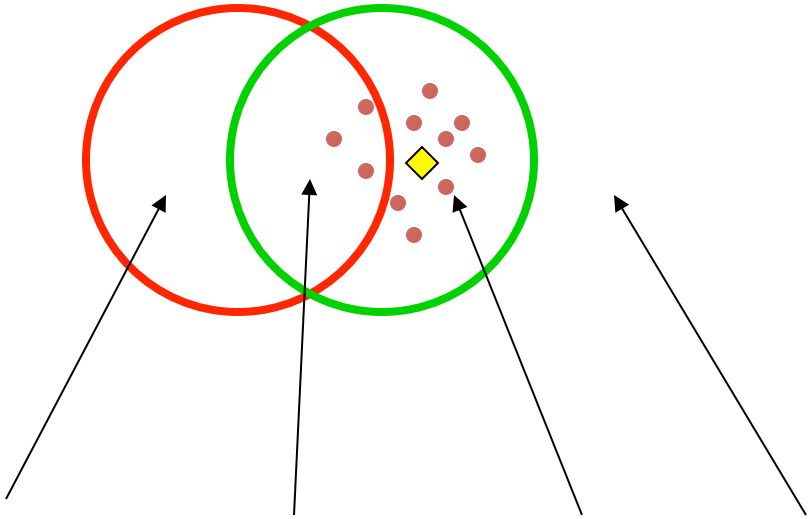

$(p_1, -500; p_2, -400; p_3, 100; p_4, 0)$

Movement plans as lotteries

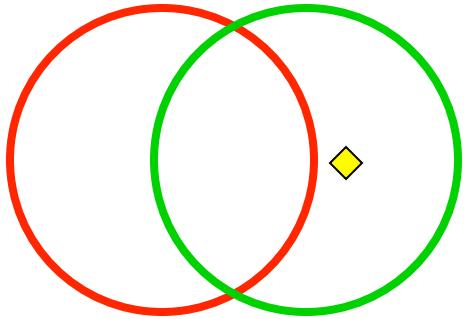

= 4 mm

Lottery:

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

Movement plans as lotteries

Optimal aim point: lottery with MEV

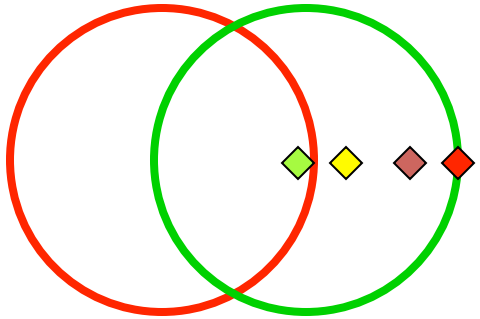

= 4 mm

(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)

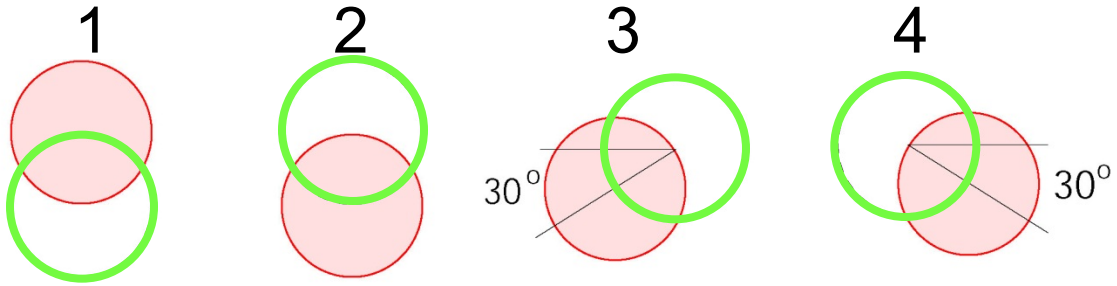

R = 9 mm

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.*

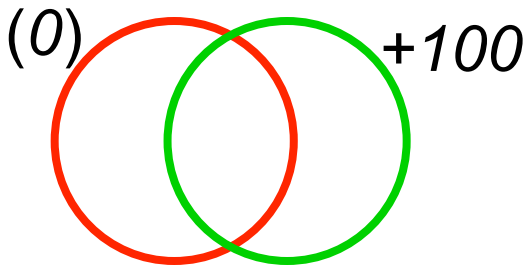

General Methods: Training

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.*

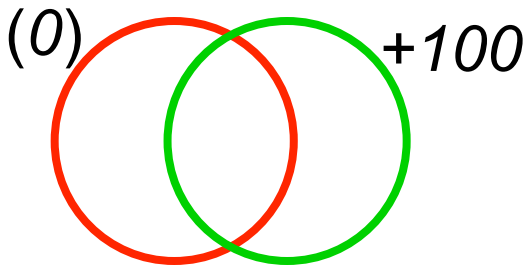

Results: Experiment 1

Model prediction:

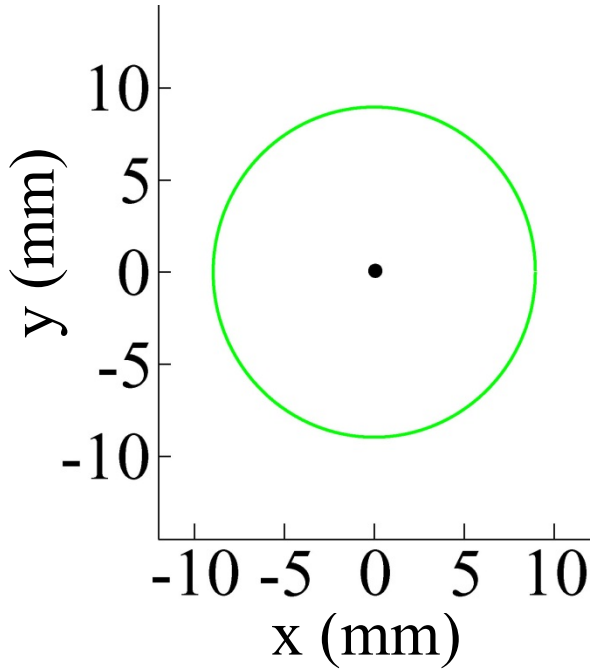

- model, penalty = 0

Subject S5, [illegible] = 2.99 mm

Results: Experiment 1

Model prediction: configuration 1

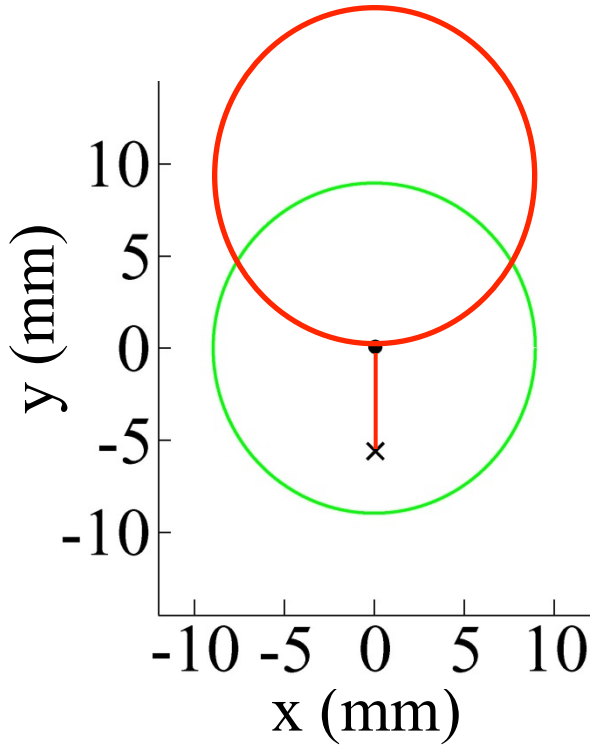

• model, penalty = 0
x model, penalty = 500

Subject S5, [illegible] = 2.99 mm

Results: Experiment 1

Model prediction: configuration 2

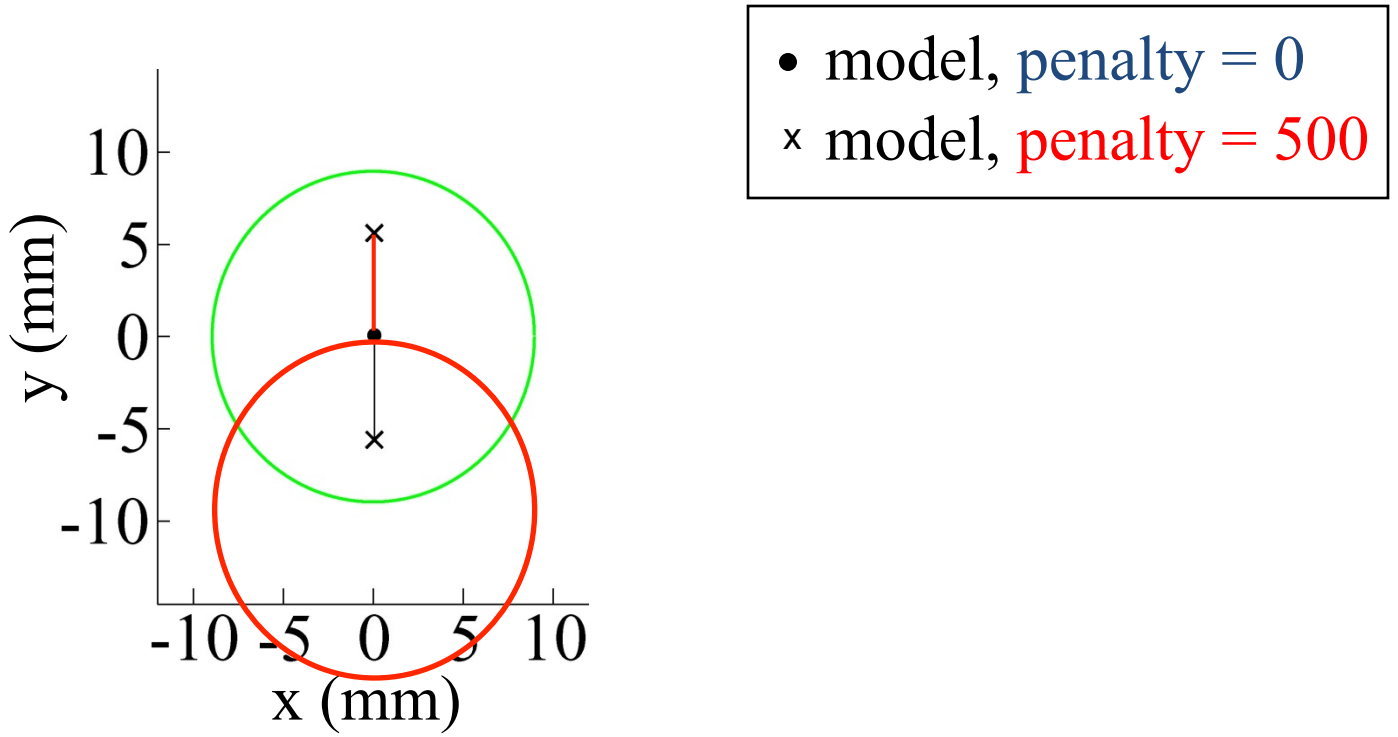

Subject S5, [illegible] = 2.99 mm

Results: Experiment 1

Model prediction: configuration 3

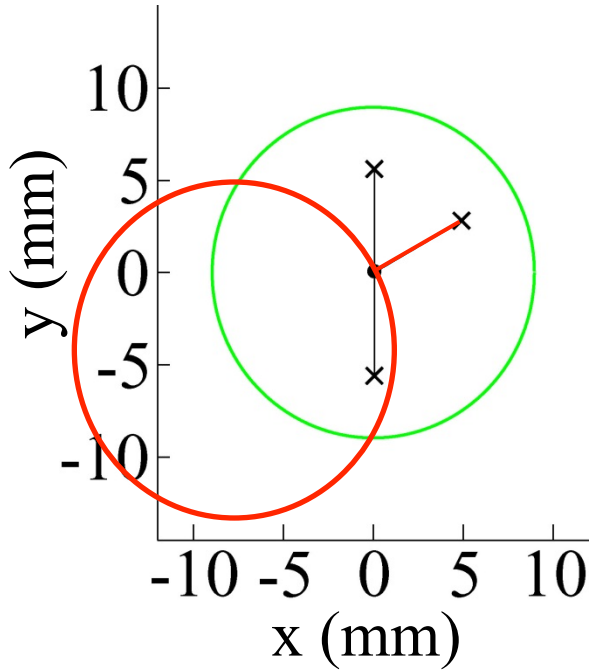

- model, penalty = 0
- x model, penalty = 500

Subject S5, [illegible] = 2.99 mm

Results: Experiment 1

Model prediction: configuration 4

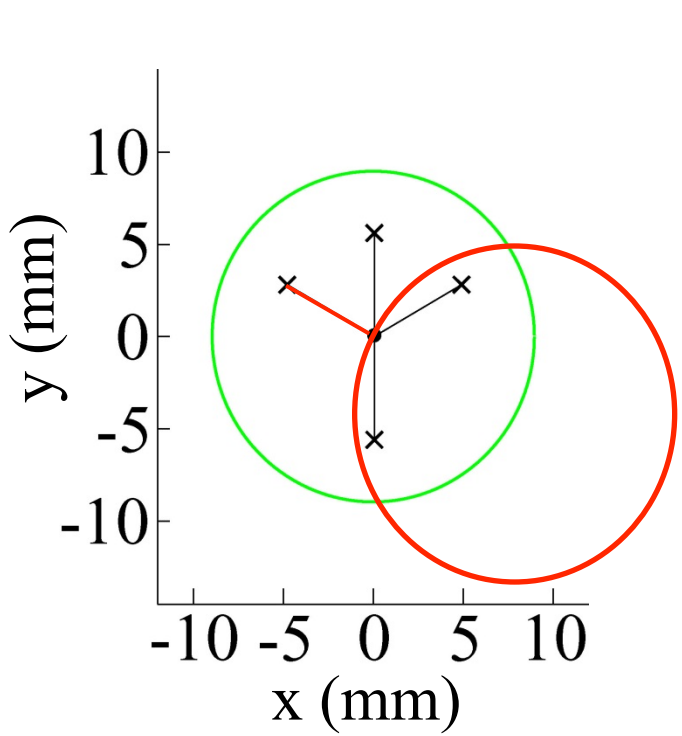

- model, penalty = 0
- x model, penalty = 500

Subject S5, [illegible] = 2.99 mm

Results: Experiment 1

Comparison with experiment

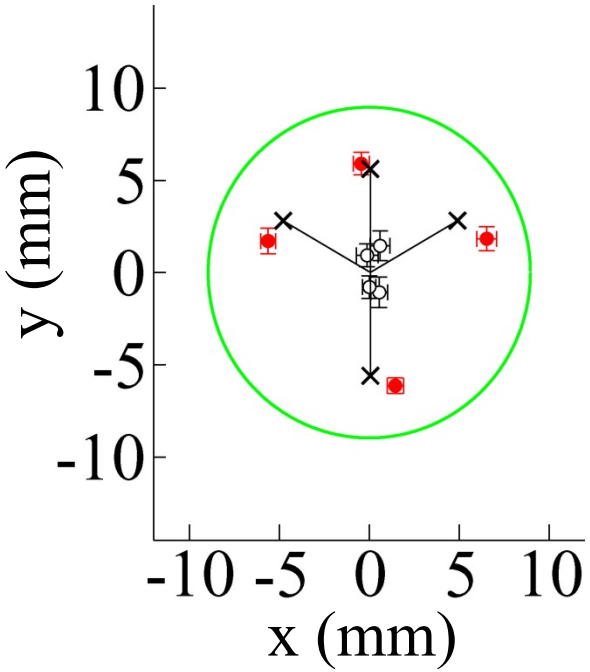

- ○ exp., penalty = 0
- • exp., penalty = 500
- x model, penalty = 500

Subject S5, [illegible] = 2.99 mm

Results: Experiment 1

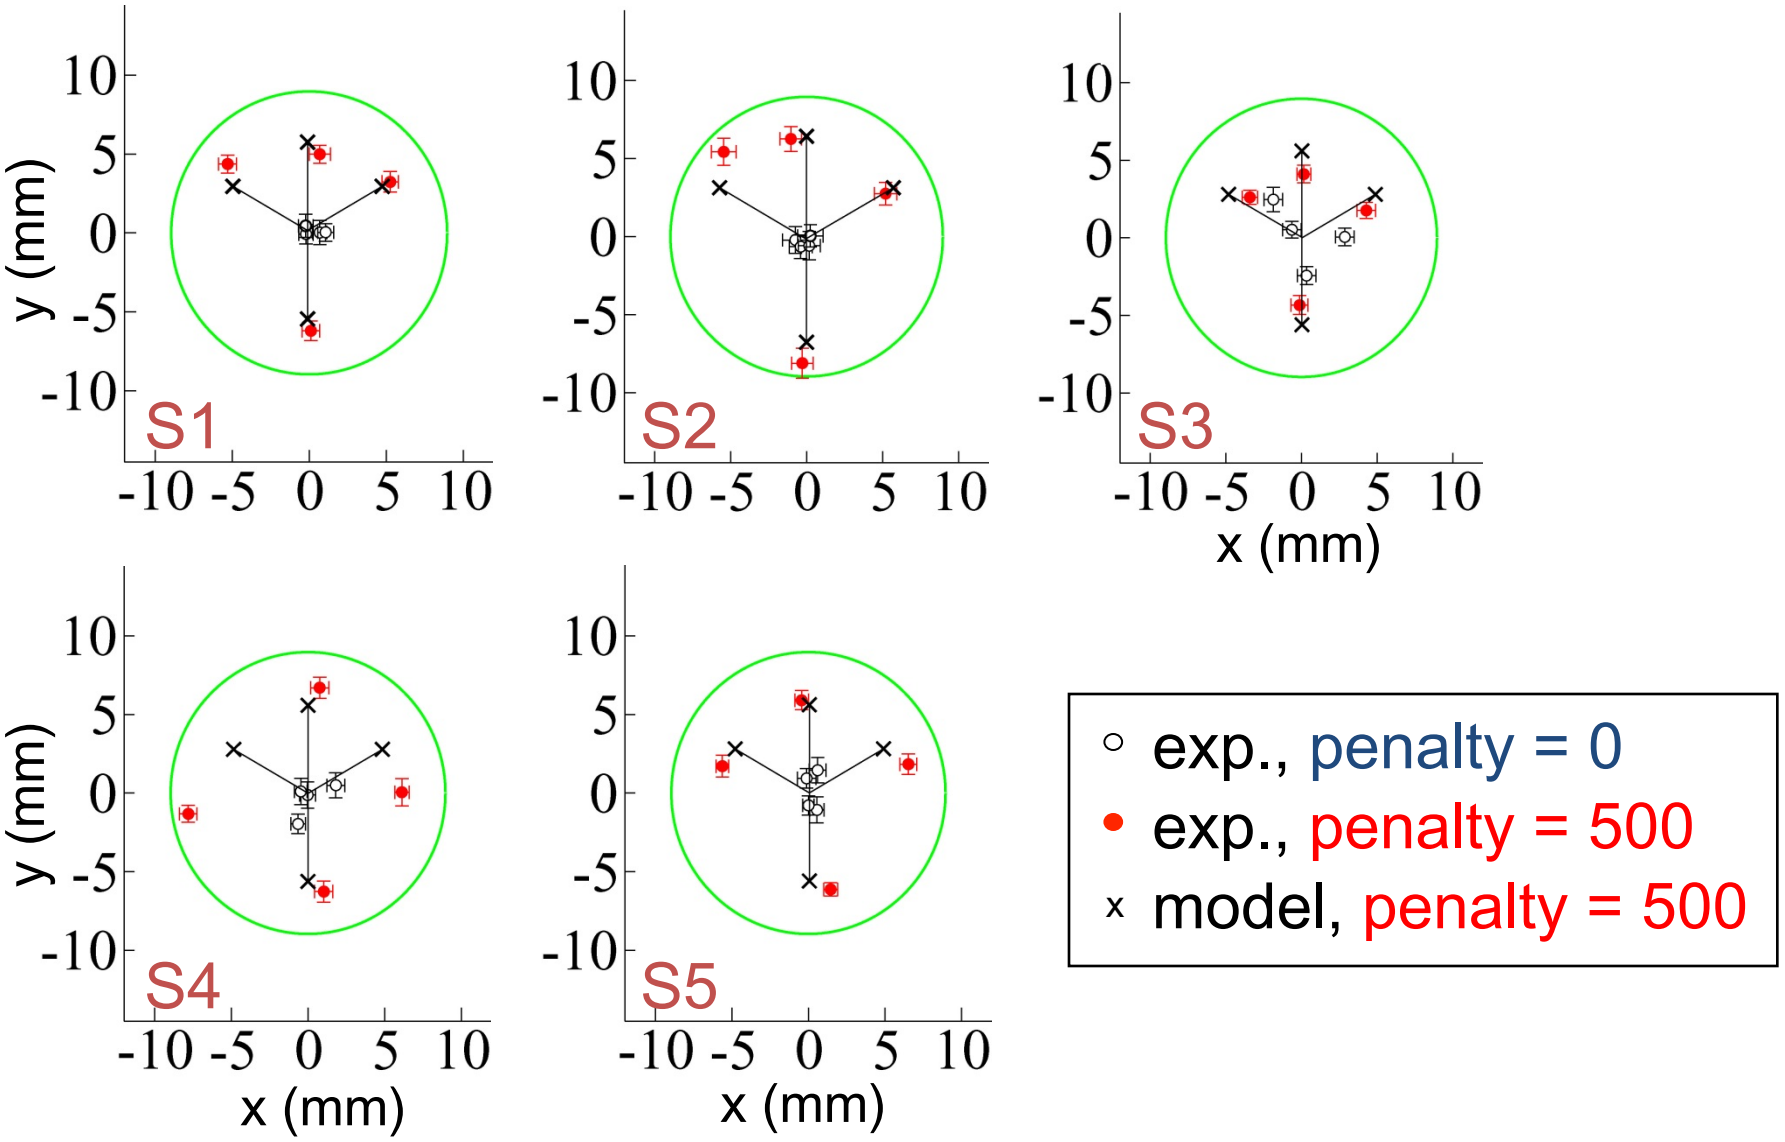

Results: Experiment 1

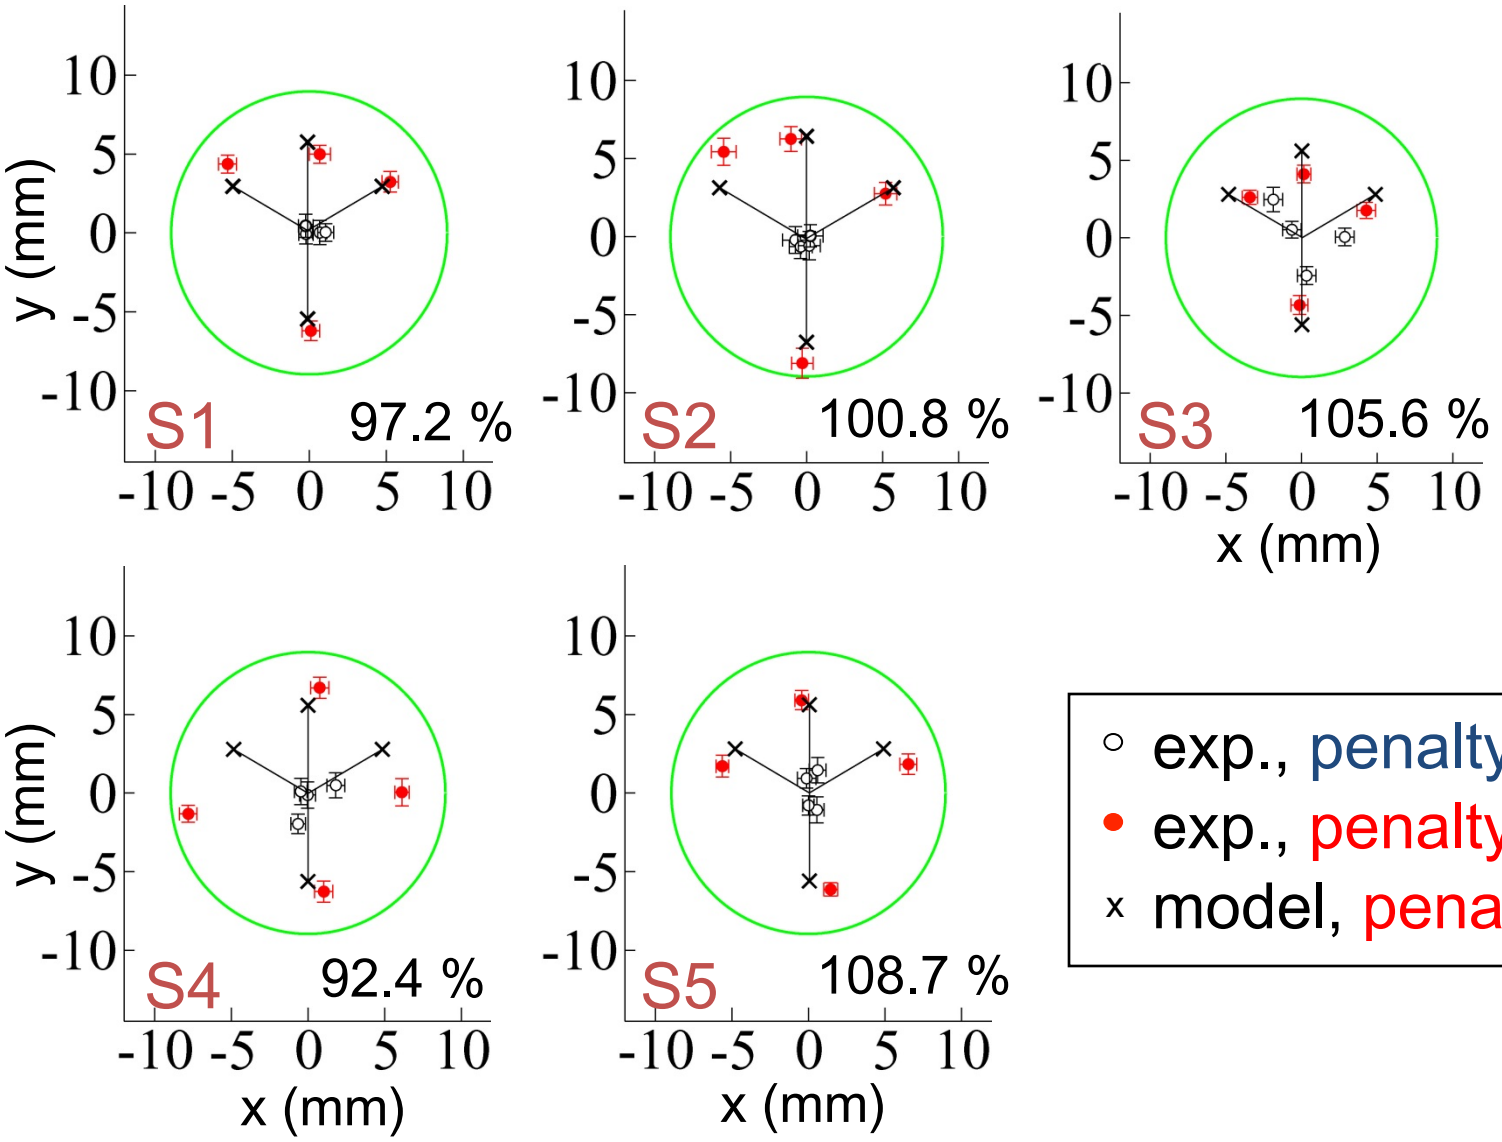

- ○ exp., penalty = 0
- • exp., penalty = 500
- x model, penalty = 500

C

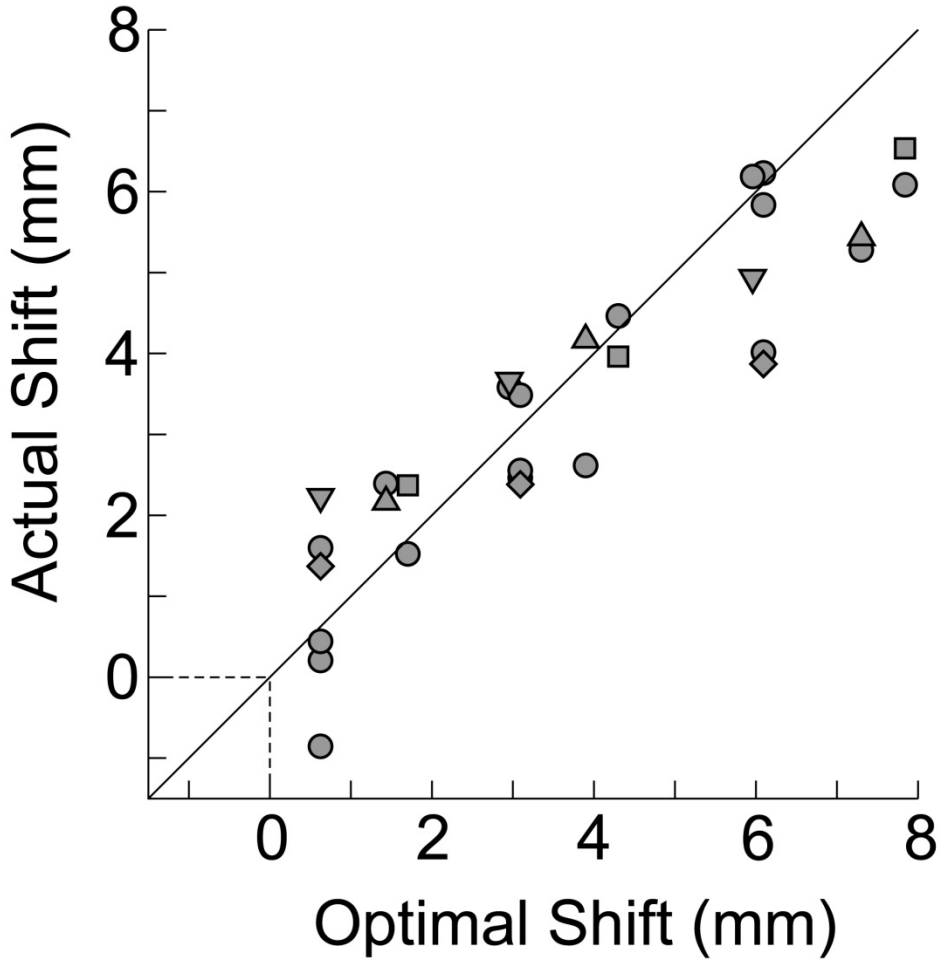

b

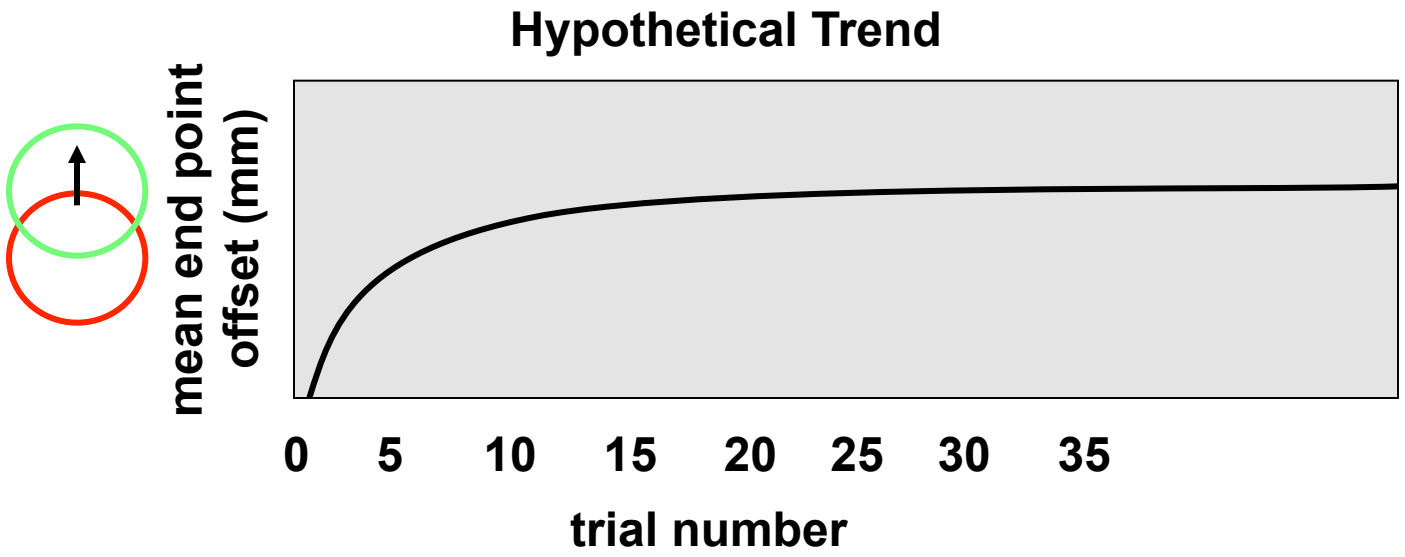
Hypothetical Trend
mean end point offset (mm)
0
5
10
15
20
25
30
35
trial number

a

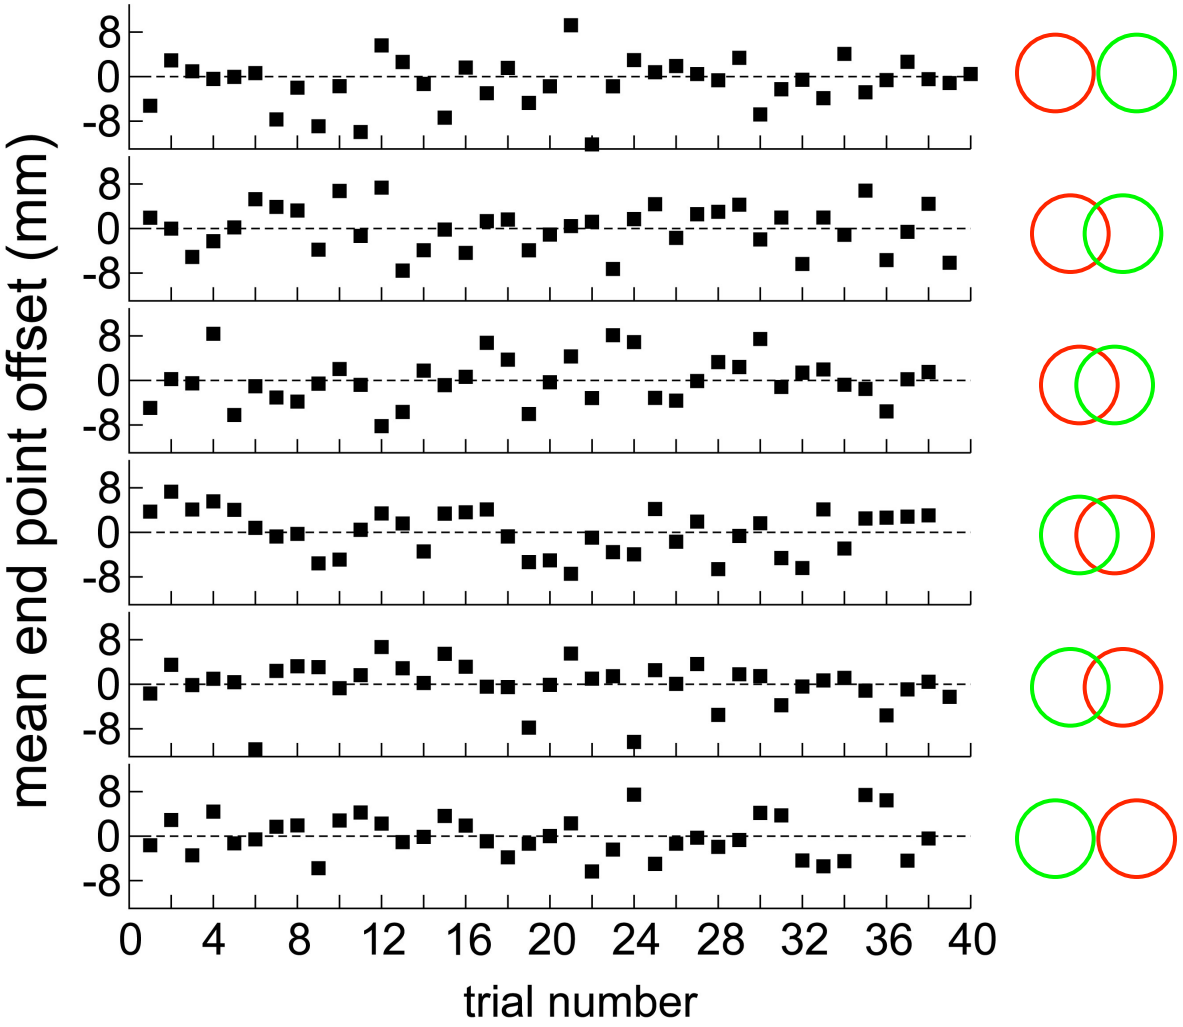
mean end point offset (mm)
8
0
-8
0
4
8
12
16
20
24
28
32
36
40
trial number

Movement Planning

Trommershäuser, Maloney & Landy (2003), JOSA A, 20, 1419-1433.

Trommershäuser, Maloney & Landy (2003), Spatial Vision, 16, 255-275.

Trommershäuser, Gepshtein, Maloney, Landy & Banks, (2005), J. Neuroscience 25, 7169-7178.

Wu, Trommershäuser, Maloney & Landy (2006), JOV, 6, 53-63.

Trommershäuser, Mattis, Maloney & Landy (2006), EBR, 175, 276-284.

Trommershäuser, Landy & Maloney (2006), Psych. Science, 17, 981-988.

Hudson, Maloney & Landy (2007), J. Neurophysiology, 98, 3034-3046.

Dean, Wu & Maloney (2006), JOV, 7(5):10,1-12.

Trommershäuser, Maloney & Landy (2008), TiCS, 12(8), 291-297.

Hudson, Maloney & Landy (2008), PLoS Comp Biol, 4(7):e100130,1-9.

and more …