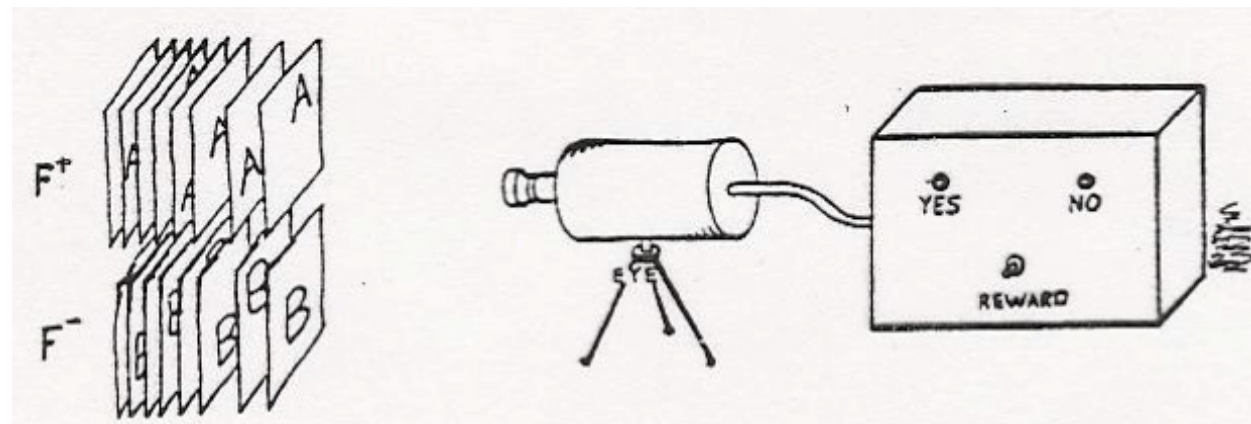
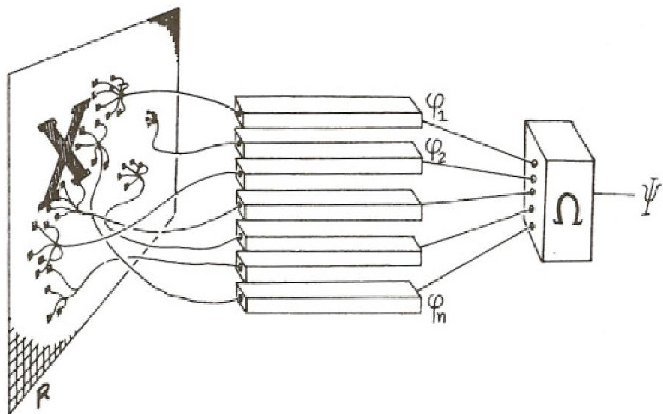


Mike's Brief History Of Machine Learning

1962

Frank Rosenblatt, *Principles of Neurodynamics: Perceptrons and the Theory of Brain Mechanisms*

Perceptron can learn anything you can program it to do.

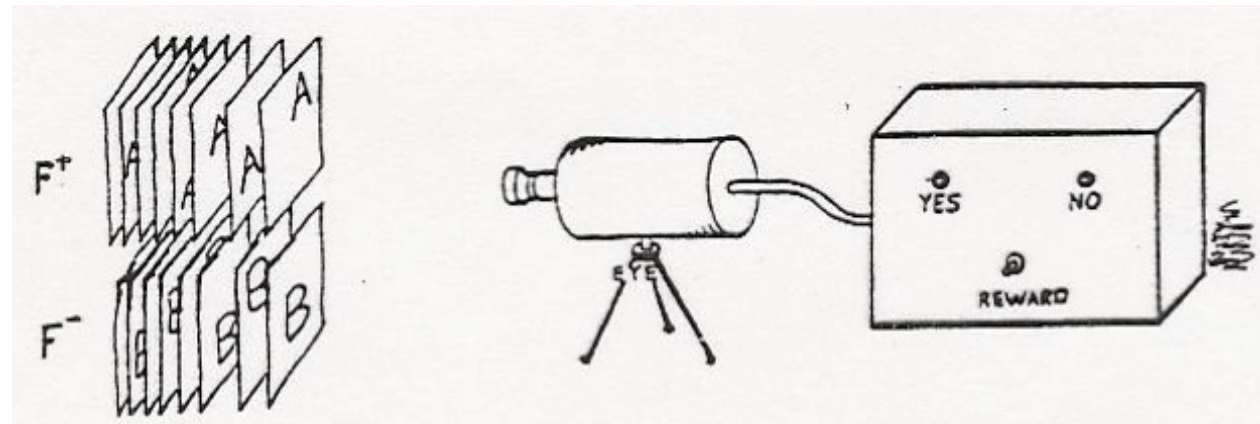


Mike's Brief History Of Machine Learning

1969

Minsky & Papert, *Perceptrons*

There are many things a perceptron can't in principle learn to do



Mike's Brief History Of Machine Learning

1970-1985

Attempts to develop symbolic rule discovery algorithms

1986

Rumelhart, Hinton, & Williams, *Back propagation*

Overcame many of the Minsky & Papert objections

1990-2000

Statisticians

Bayesian Optimization: From A/B Testing To A-Z Testing

**Robert V. Lindsey, Brett Roads,
Mohammad Khajah, Michael Mozer**

**Department of Computer Science
University of Colorado, Boulder**

Harold Pashler

**Department of Psychology
UC San Diego**

A/B Testing

**Randomly
assign web-
page visitors
to one of two
conditions,
A or B**

**Serve A or B
version of
web page
according to
condition**

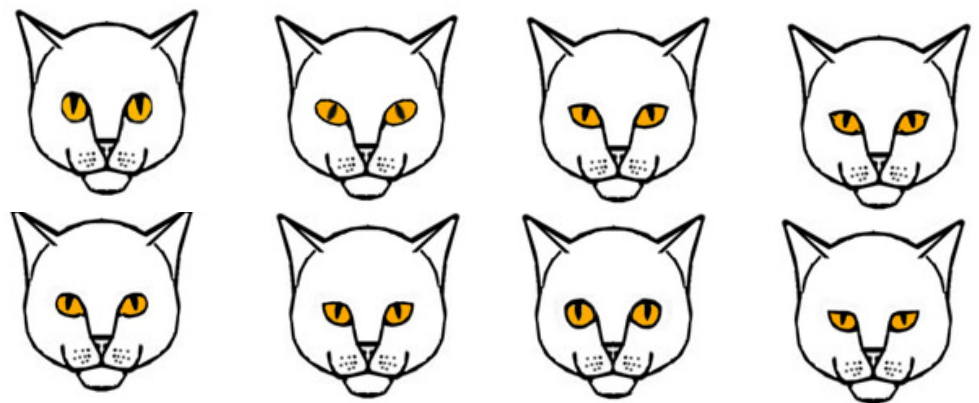
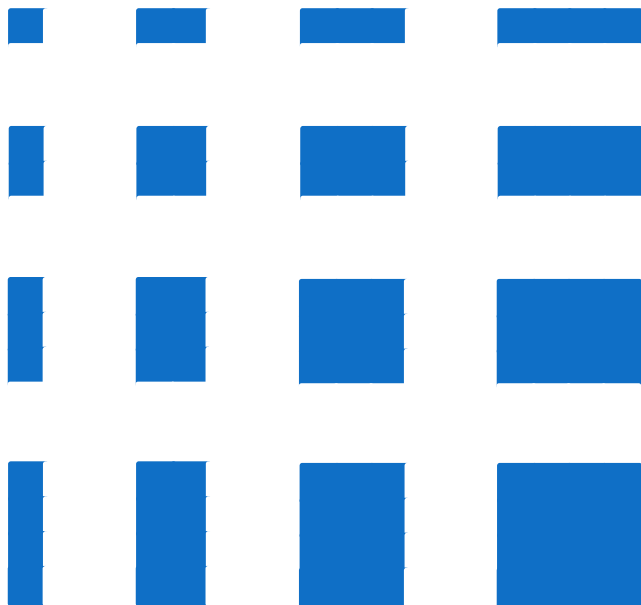
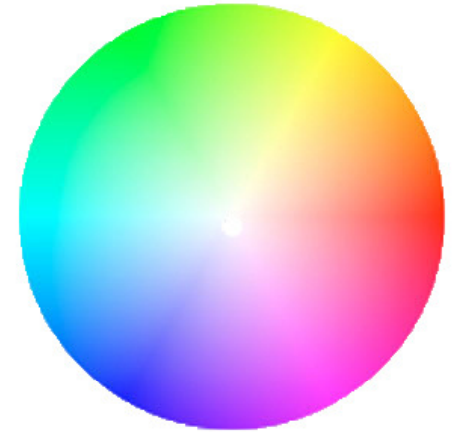
**Measure
which
condition
leads to
better
results**

A/B Testing On Steroids

Suppose we could compare not just two
or a small number of options...

But a continuum of options...

As efficiently as we compared 2.



From Your World To Mine

A/B testing isn't used just in marketing and high tech companies.

A/B testing is the core technique used in science.

- known as a *randomized controlled experiment*.

Randomized Controlled Experiments In Psychology

E.g., distributed-practice effect

massed vs. spaced practice

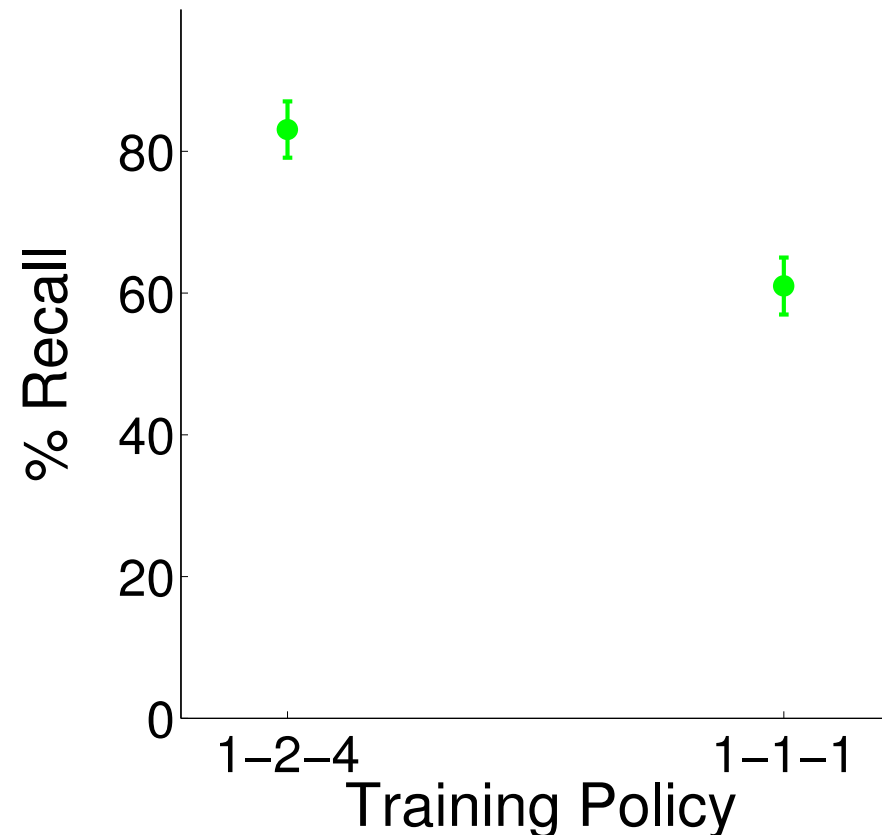
Propose several spaced conditions to compare

Equal: 1 – 1 – 1

Increasing: 1 – 2 – 4

Run many subjects in each
condition

Perform statistical analyses to
establish reliable difference
between conditions

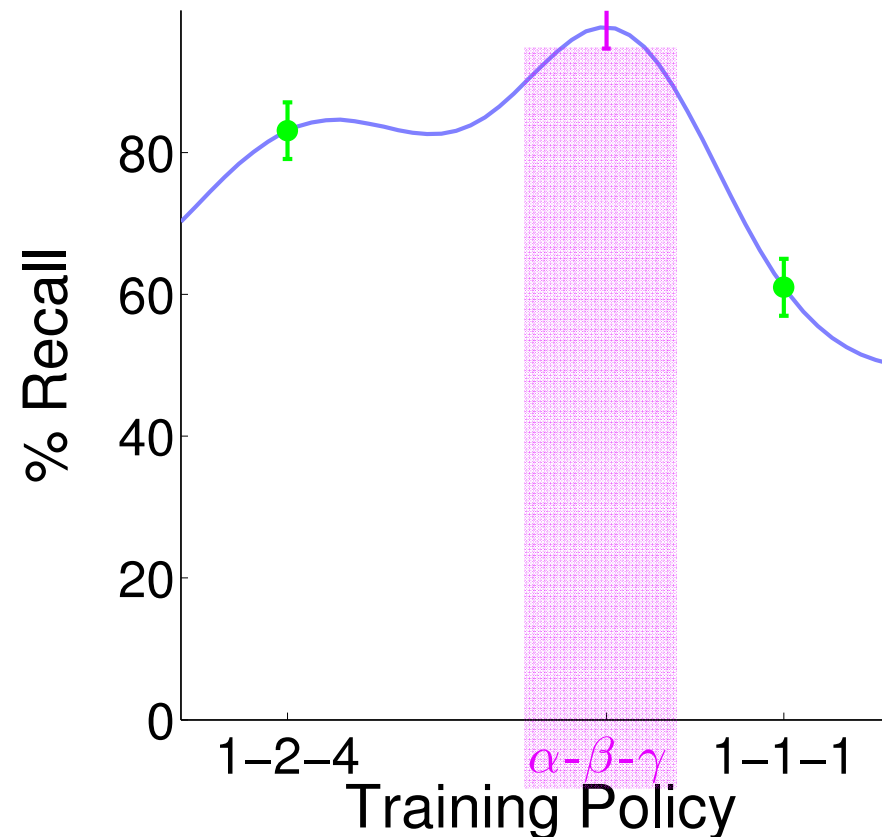


What Researchers Really Want To Do

Find the best study schedule (*training policy*)

Abcissa: space of all training policies

Performance function defined
over policy space



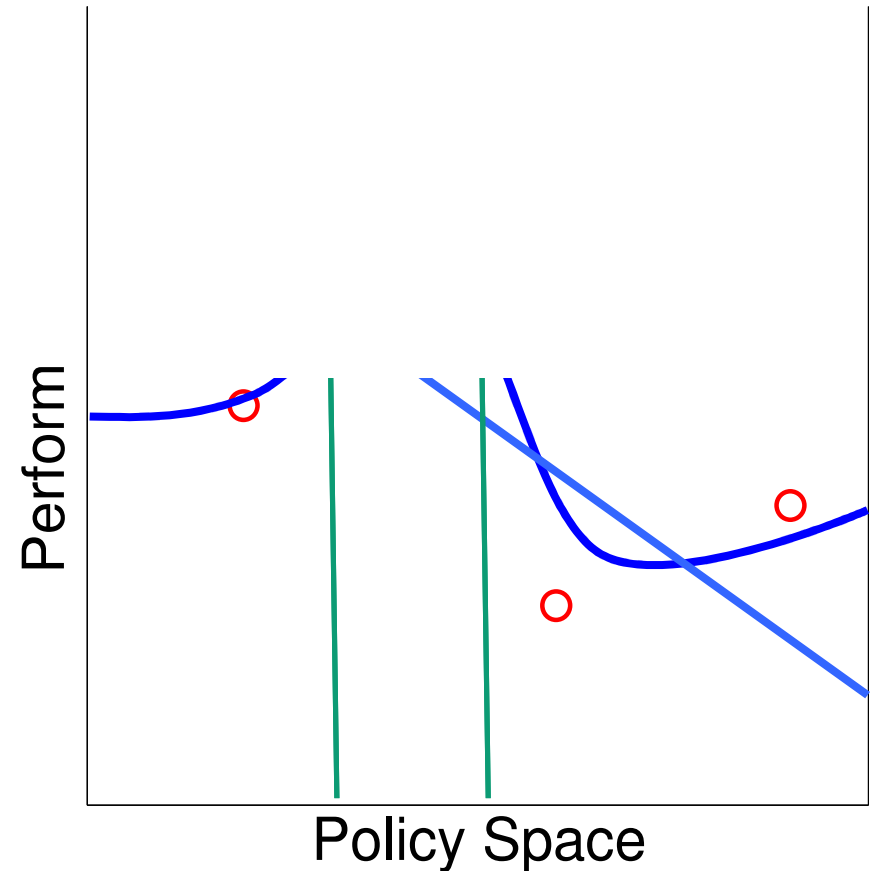
Approach

Perform single-subject experiments at selected points in policy space (o)

Use curve fitting (function approximation) techniques to estimate shape of the performance function

Given current estimate, select *promising* policies to evaluate next.

- promising = has potential to be the optimum policy



Gaussian Process Regression

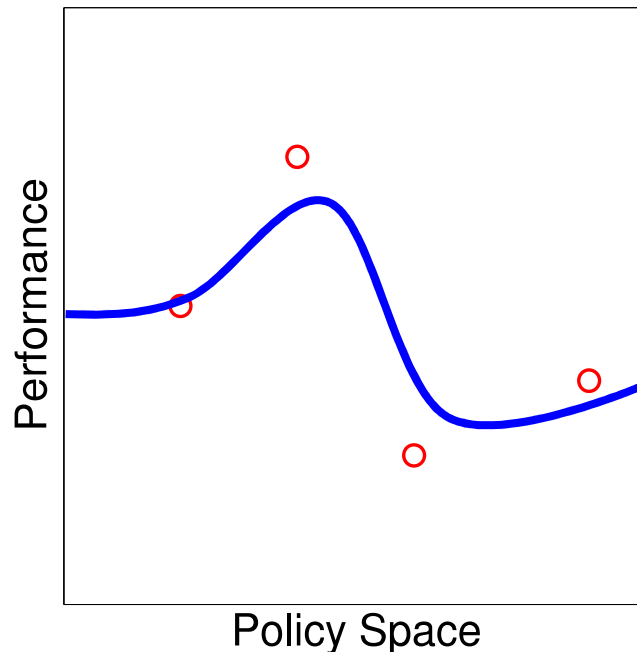
Assumes only that functions are smooth

How smooth is determined by the data

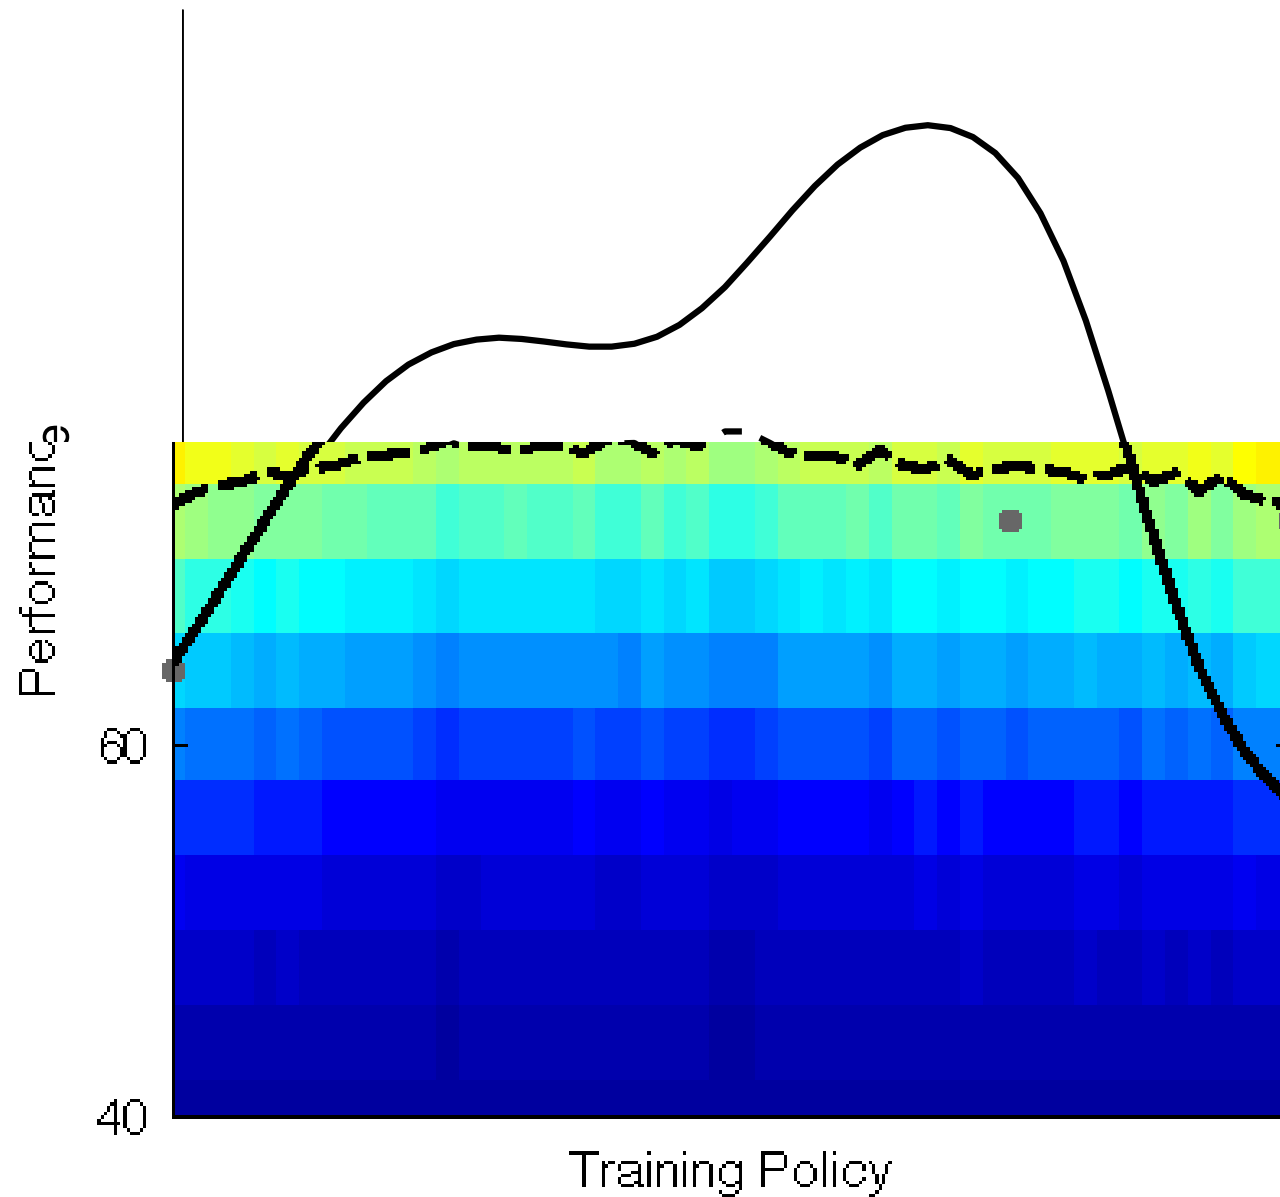
Uses data efficiently

Accommodates noisy data

Produces estimates of both function shape and uncertainty



Simulated Experiment



1 Policy Selection Heuristic

I propose the following heuristic for choosing the next training policy to evaluate. Let the random variable π_x be the population average performance at policy x ,

$$\pi_x = .5 + .5 \frac{1}{1 + \exp(-f_x)} \quad (1)$$

We can calculate the expectation $\mu_x = \mathbb{E}[\pi_x]$ using samples from the posterior predictive distribution of the GP \tilde{f} . I propose we choose the next training policy \hat{x} based on

$$\hat{x} = \arg \max_x \mathbb{E} \left[(m\mu_x - \pi_x)^2 \right] \quad (2)$$

where $0 \leq m \leq 1$. Pure exploitation ($m = 0$) and pure exploration ($m = 1$) are the extreme cases of this:

$$\hat{x} = \arg \max_x \mathbb{E}[\pi_x^2] = \arg \max_x \mathbb{E}[\pi_x] \quad \text{when } m=0 \quad (3)$$

$$\hat{x} = \arg \max_x \text{Var}[\pi_x] \quad \text{when } m=1 \quad (4)$$

Thinking of this expectation as a weighted sum of squared-distances between $m\mu_x$ and π (summed across possible x 's), m lets us manipulate the magnitude of the distances while keeping the weights fixed. If m is near 0, the distances are at their largest for large π_x values. Hence, even if small π_x are more highly weighted (ie more probable), the big squared-distance values of the less probable larger π_x 's will matter the most. Conversely, an m near 1 places less emphasis on the squared-distances and more emphasis on their weights.

Assuming this is a sensible approach, I prefer it to a selection policy that uses the raw GP function estimates because of previously discussed issues associated with large uncertainty at extremely high or low GP values not mattering. Also, having a policy selection heuristic that's based loosely on our prediction of an observable variable seems better than using a prediction of an unobservable variable.

Also note that $\text{Var}[\pi_x]$ goes to 0 as we run more and more experiments at policy x . We shouldn't get stuck choosing the same policy over and over again when $m > 0$.

2 Marginal Likelihood

2.1 Lemma

Given

$$p \mid \alpha, \beta \sim \text{Beta}(\alpha, \beta) \quad (5)$$

$$\tilde{p} = .5 + .5p \quad (6)$$

$$n_c \mid p, n. \sim \text{Binomial}(\tilde{p}, n.) \quad (7)$$

where, in our case, $n.$ is the number of test questions, n_c is the number of correct responses made, \tilde{p} is the subject's mean recall probability *corrected for chance guessing*. The marginal likelihood is

$$P(n_c \mid \alpha, \beta) = 2^{-n} \binom{n}{n_c} \sum_{i=0}^{n_c} \binom{n_c}{i} \frac{B(\alpha + i, n. + \beta - n_c)}{B(\alpha, \beta)} \quad (8)$$

2.2 Proof

The chance-corrected likelihood equation is

$$L(n_c \mid n., p) = \binom{n.}{n_c} \tilde{p}^{n_c} (1 - \tilde{p})^{n. - n_c} \quad (9)$$

$$= \binom{n.}{k} .5^{n.} (1 + p)^{n_c} (1 - p)^{n. - n_c} \quad (10)$$

The beta prior is

$$\pi(p \mid \alpha, \beta) = \frac{1}{B(\alpha, \beta)} p^{\alpha-1} (1-p)^{\beta-1} \quad (11)$$

where B is the beta function. The marginal likelihood is defined as

$$P(n_c \mid \alpha, \beta, n.) = \int_0^1 L(n_c \mid n., p) \pi(p \mid \alpha, \beta) dp \quad (12)$$

$$= 2^{-n.} \frac{1}{B(\alpha, \beta)} \binom{n.}{n_c} \int_0^1 (1+p)^{n_c} p^{\alpha-1} (1-p)^{\beta-1+n.-n_c} dp \quad (13)$$

Because n_c is an integer, we can apply the binomial theorem

$$P(n_c \mid \alpha, \beta, n.) = 2^{-n.} \frac{1}{B(\alpha, \beta)} \binom{n.}{n_c} \int_0^1 \sum_{i=0}^{n_c} \binom{n_c}{i} p^i p^{\alpha-1} (1-p)^{\beta-1+n.-n_c} dp \quad (14)$$

$$= 2^{-n.} \frac{1}{B(\alpha, \beta)} \binom{n.}{n_c} \sum_{i=0}^{n_c} \binom{n_c}{i} \int_0^1 p^{\alpha+i-1} (1-p)^{\beta-1+n.-n_c} dp \quad (15)$$

The integral in the summation is over an unnormalized Beta($\alpha + i, n + \beta - k$) density. Therefore,

$$P(n_c \mid \alpha, \beta, n.) = 2^{-n.} \binom{n.}{n_c} \sum_{i=0}^{n_c} \binom{n_c}{i} \frac{B(\alpha + i, n. + \beta - n_c)}{B(\alpha, \beta)} \quad (16)$$

3 Inference

3.1 Model

The model we assume is

$$\mathbf{f} \sim \text{GP}(m(x), \Sigma(x, x')) \quad (17)$$

$$p_s \mid \alpha, f_s \sim \text{Beta}(\alpha, \alpha \exp(-f_s)) \quad (18)$$

$$\tilde{p}_s = .5 + .5p_s \quad (19)$$

$$n_{cs} \mid p_s, n_{.s} \sim \text{Binomial}(\tilde{p}_s, n_{.s}) \quad (20)$$

where s is a subject index, α is a free parameter controlling inter-subject variability. Before performing inference, we analytically marginalize p_s via Equation 8.

The model likelihood can be written as

$$\mathcal{L} = P(\mathbf{n}_c \mid \mathbf{f}) = \prod_s 2^{-n_{.s}} \binom{n_{.s}}{n_{cs}} \sum_{i=0}^{n_{cs}} \binom{n_{cs}}{i} \frac{B(\alpha + i, n_{ws} + \alpha e^{-f_s})}{B(\alpha, \alpha e^{-f_s})} \quad (21)$$

where n_{ws} is the number of wrong responses made by subject s . The prior follows a MVN density.

3.2 Gradient and Hessian

Let

$$z \equiv \alpha \left(\sum_{i=0}^{n_c} \binom{n_c}{i} B(\alpha + i, n. + \beta + n_c) \right)^{-1} \quad (22)$$

We have

$$\frac{\partial}{\partial \mathbf{f}_s} \log \mathcal{L} = z e^{-f_s} \Gamma(n_{ws} + \alpha e^{-f_s}) \sum_{i=0}^{n_c} \binom{n_c}{i} \Gamma(\alpha + i) \frac{\Psi(n_{ws} + i + \alpha + \alpha e^{-f_s}) - \Psi(n_{ws} + \alpha e^{-f_s})}{\Gamma(n_{ws} + i + \alpha + \alpha e^{-f_s})} \quad (23)$$

where Ψ is the digamma function.

$$\frac{\partial^2}{\partial \mathbf{f}_s^2} \log \mathcal{L} = \dots \quad (24)$$

3.3 Laplace Approximation

Unfinished section:

We can approximate the model's posterior distribution via a Gaussian centered at the mode

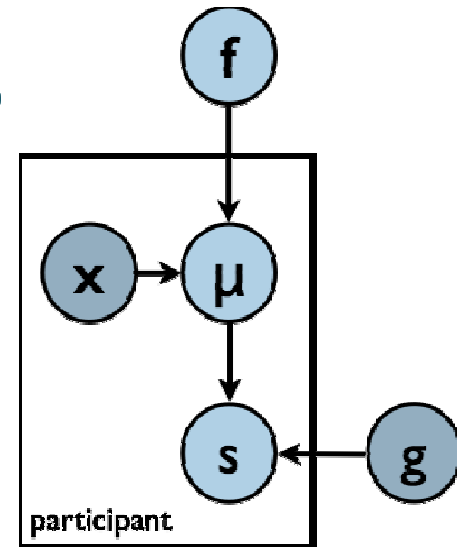
$$p(\mathbf{f} \mid \mathbf{n}_c) \sim q(\mathbf{f} \mid \mathbf{n}_c) = \mathcal{N}(\hat{\mathbf{f}}, (K^{-1} + W)^{-1}) \quad (25)$$

where K is the covariance matrix of the data, $W \equiv -\nabla \nabla L$ is the (diagonal) Hessian, $\hat{\mathbf{f}}$ is the mode (maximum likelihood) found via Newton's method using the gradient (Eqn. 23).

Model Of Human Behavior

Skill level achieved by a training policy

$$f(\mathbf{x}) \sim \mathcal{GP}(m(\mathbf{x}), k(\mathbf{x}, \mathbf{x}')) \quad -\infty \rightarrow +\infty$$



Chance-
corrected
beta
binomial

Parameter fitting

Model has a couple of free parameters

- **how much variability in performance is there across individuals?**
- **how smooth is the function?**

Free parameters fit to data via hierarchical Bayesian inference

Fact Learning Experiment

Associate each person with the name of their favorite sports team

Six training faces

30 seconds of training

Each face shown for duration d ms

→ each face shown $5000/d$ times

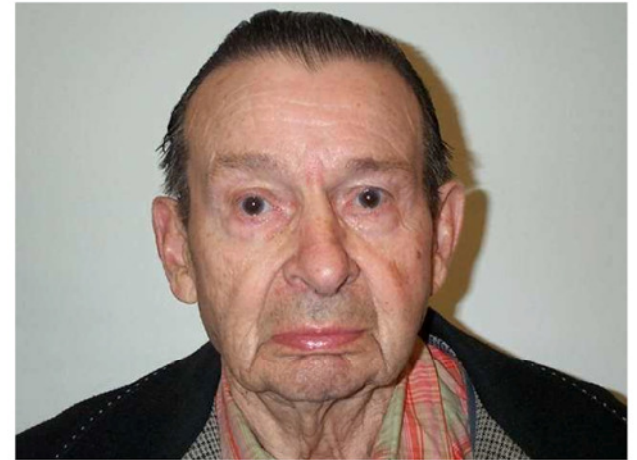
Immediate 2AFC testing following training

Demos

$d = 250$ ms

$d = 5000$ ms

Jets Fan



Fact Learning Experiment

What is the optimal presentation duration?

$d = 250 \text{ ms}$

$d = 5000 \text{ ms}$

20 presentations /
face

1 presentation /
face

more presentations is better
(with diminishing returns)

more time to process is better
(with diminishing returns)

Trade off

Fact Learning Experiment: Details

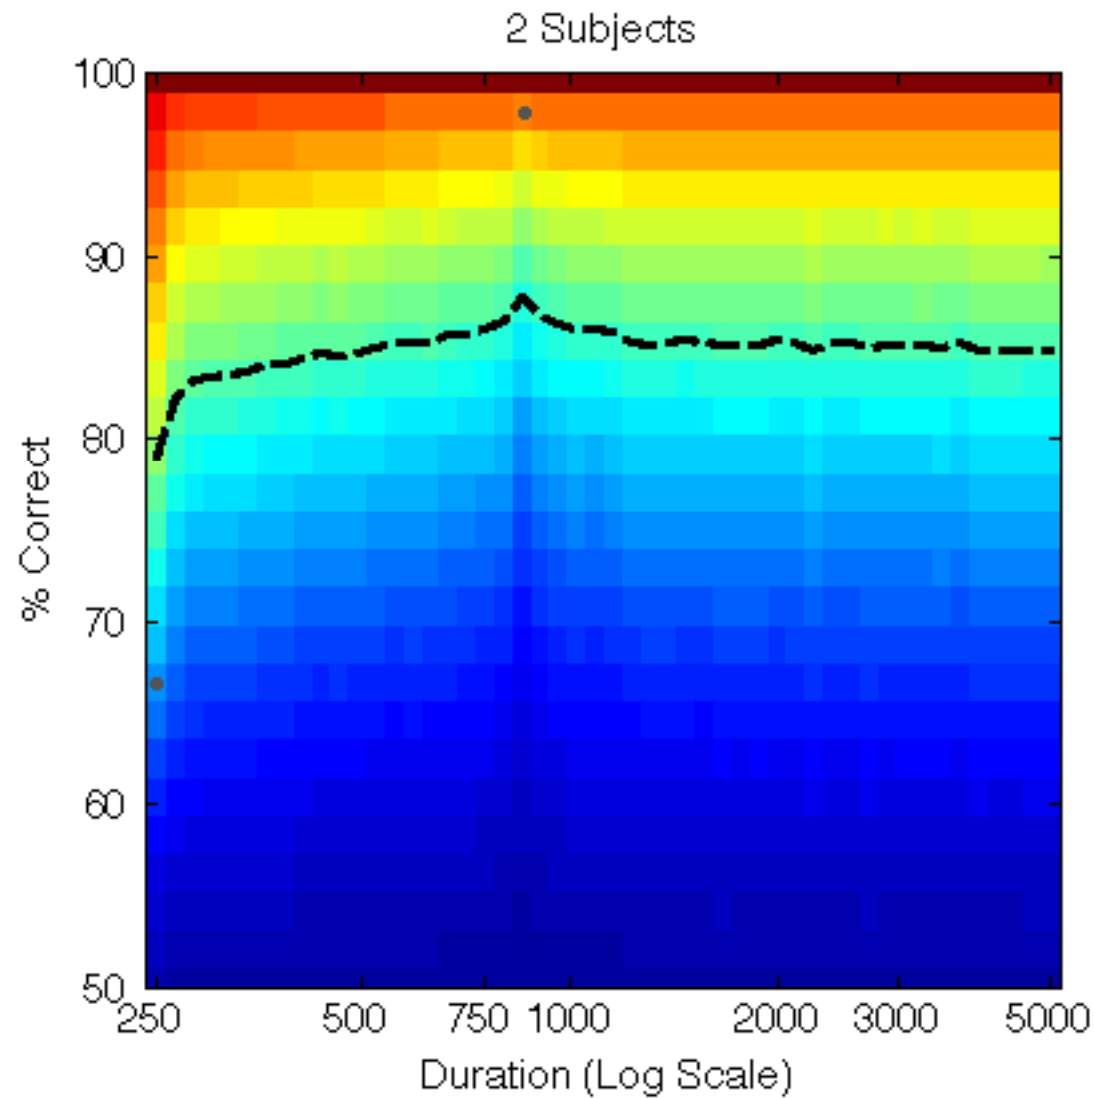
8 training/testing blocks with different faces

6 faces per block

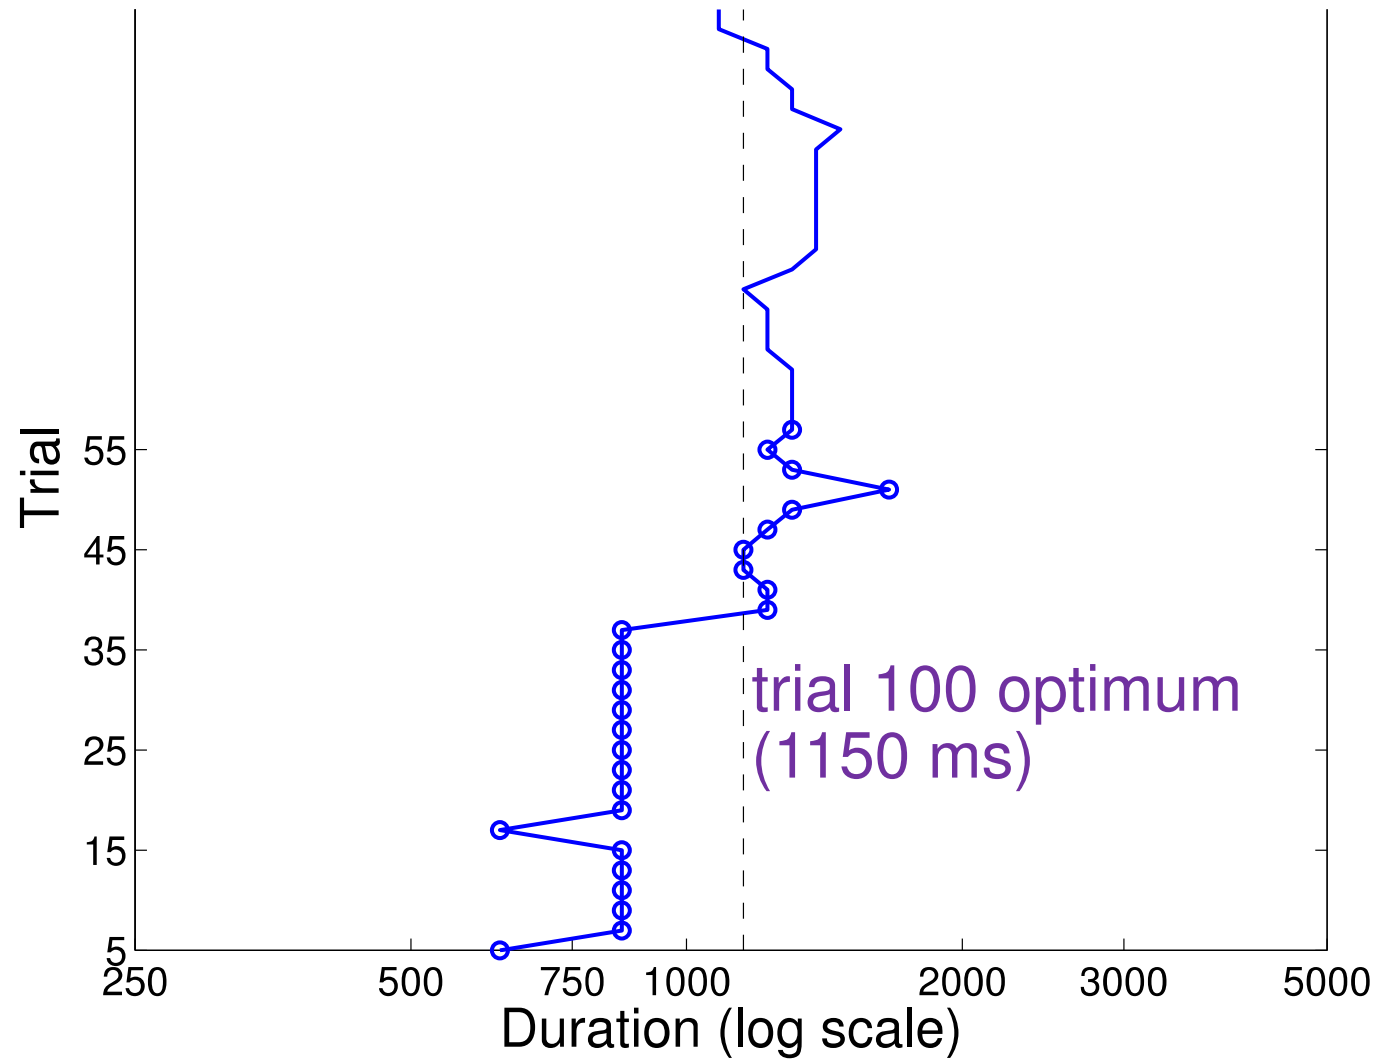
run on Mechanical Turk

30 cents/subject

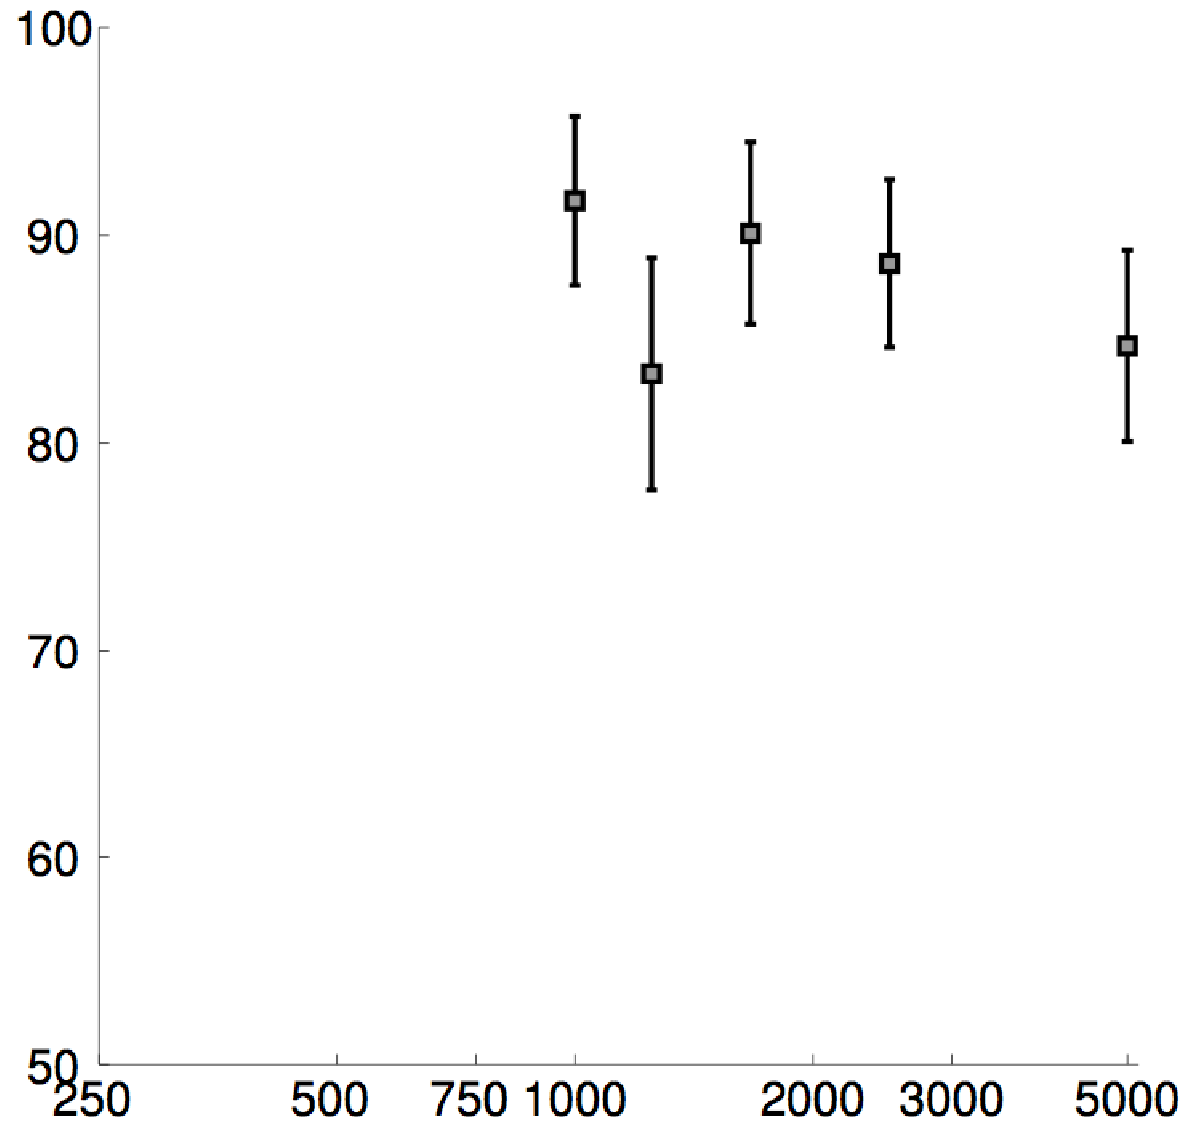
Fact Learning Experiment: Optimization



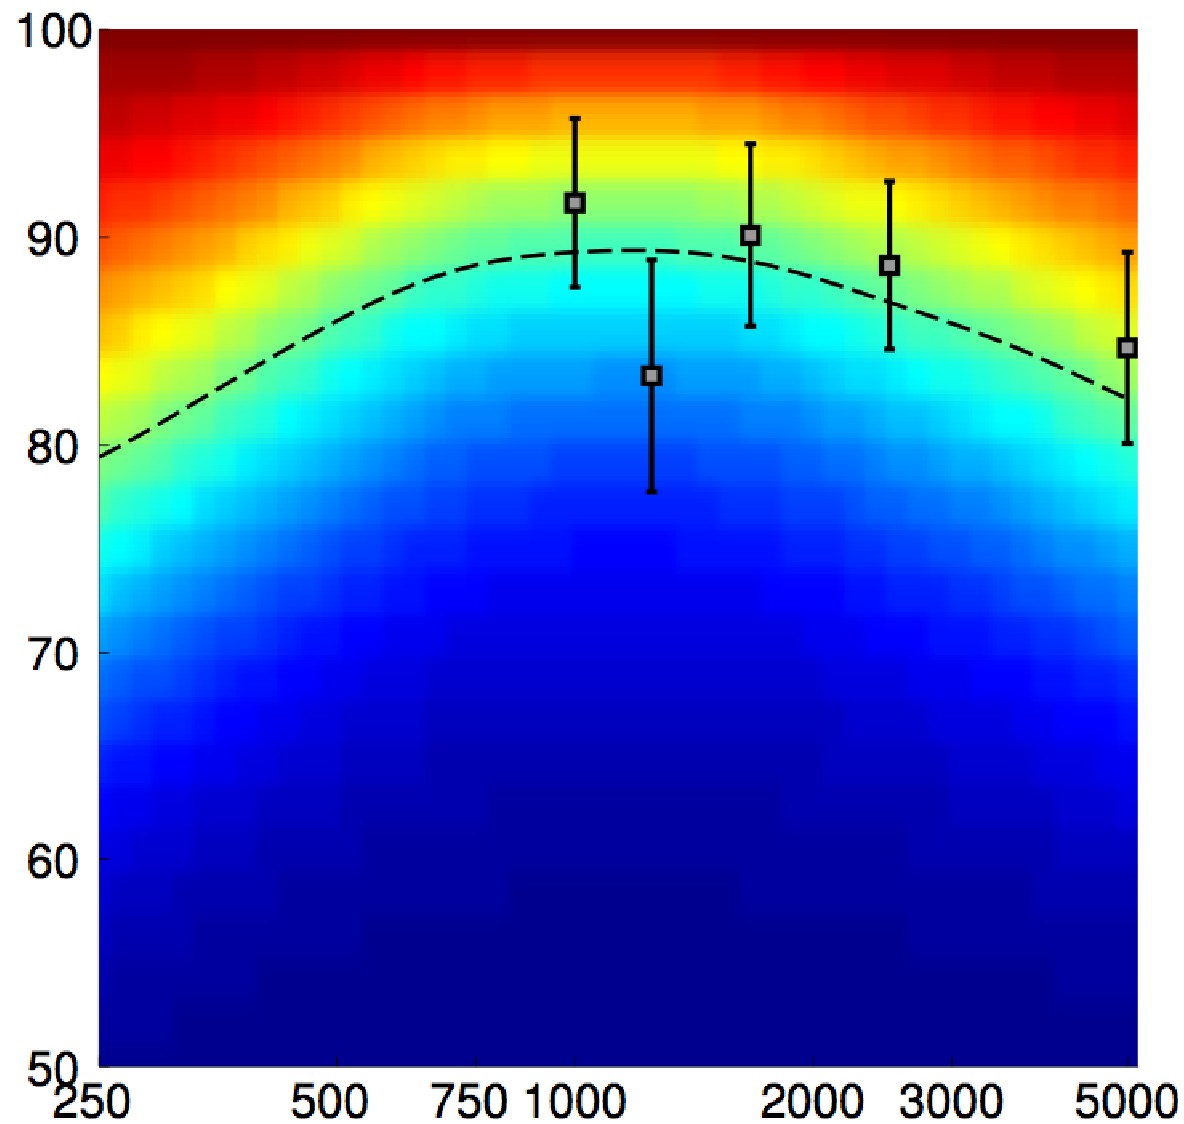
Convergence



Comparison With Traditional Experiment



Comparison With Traditional Experiment



Concept Learning Experiment

Amazon Mechanical Turk

Amazon Mechanical Turk

+

Instructions

Imagine that you have encountered aliens from the Andromeda Galaxy who want to teach you their language. In the next few minutes, they will teach you the meaning of the GLOPNOR. GLOPNOR is a word that describes a set of objects. Some objects are GLOPNOR, other objects are not GLOPNOR. In the past, aliens have taught you words that mean 'breakable', 'bendable', 'larger than a toaster oven', and 'able to be used by two or more people at once.'

The aliens will show you a sequence of objects. For each object, you are to determine whether it is GLOPNOR or not. Initially, the aliens will give you feedback to tell you if your guess was correct. After these examples with feedback, the aliens will test your understanding of GLOPNOR by asking you to judge additional objects.

You must complete the entire series of objects to receive payment. You can only participate once. This should take 5-10 minutes.

Begin



Is this GLOPNOR?

No

Perhaps no

Don't know

Perhaps yes

Yes



Is this GLOPNOR?

No

Perhaps no

Don't know

Perhaps yes

Yes

Wrong! This is GLOPNOR.



(GLOPNOR)

Is this GLOPNOR?

No

Perhaps no

Don't know

Perhaps yes

Yes



(GLOPNOR)

Is this GLOPNOR?

No

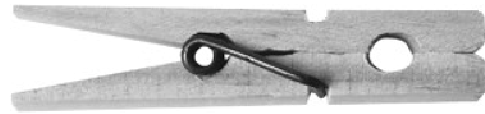
Perhaps no

Don't know

Perhaps yes

Yes

Correct! This is not GLOPNOR.



(NOT GLOPNOR)

Is this GLOPNOR?

No

Perhaps no

Don't know

Perhaps yes

Yes

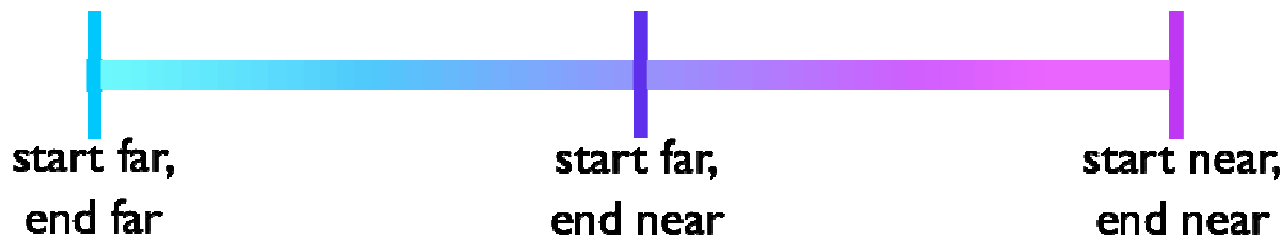
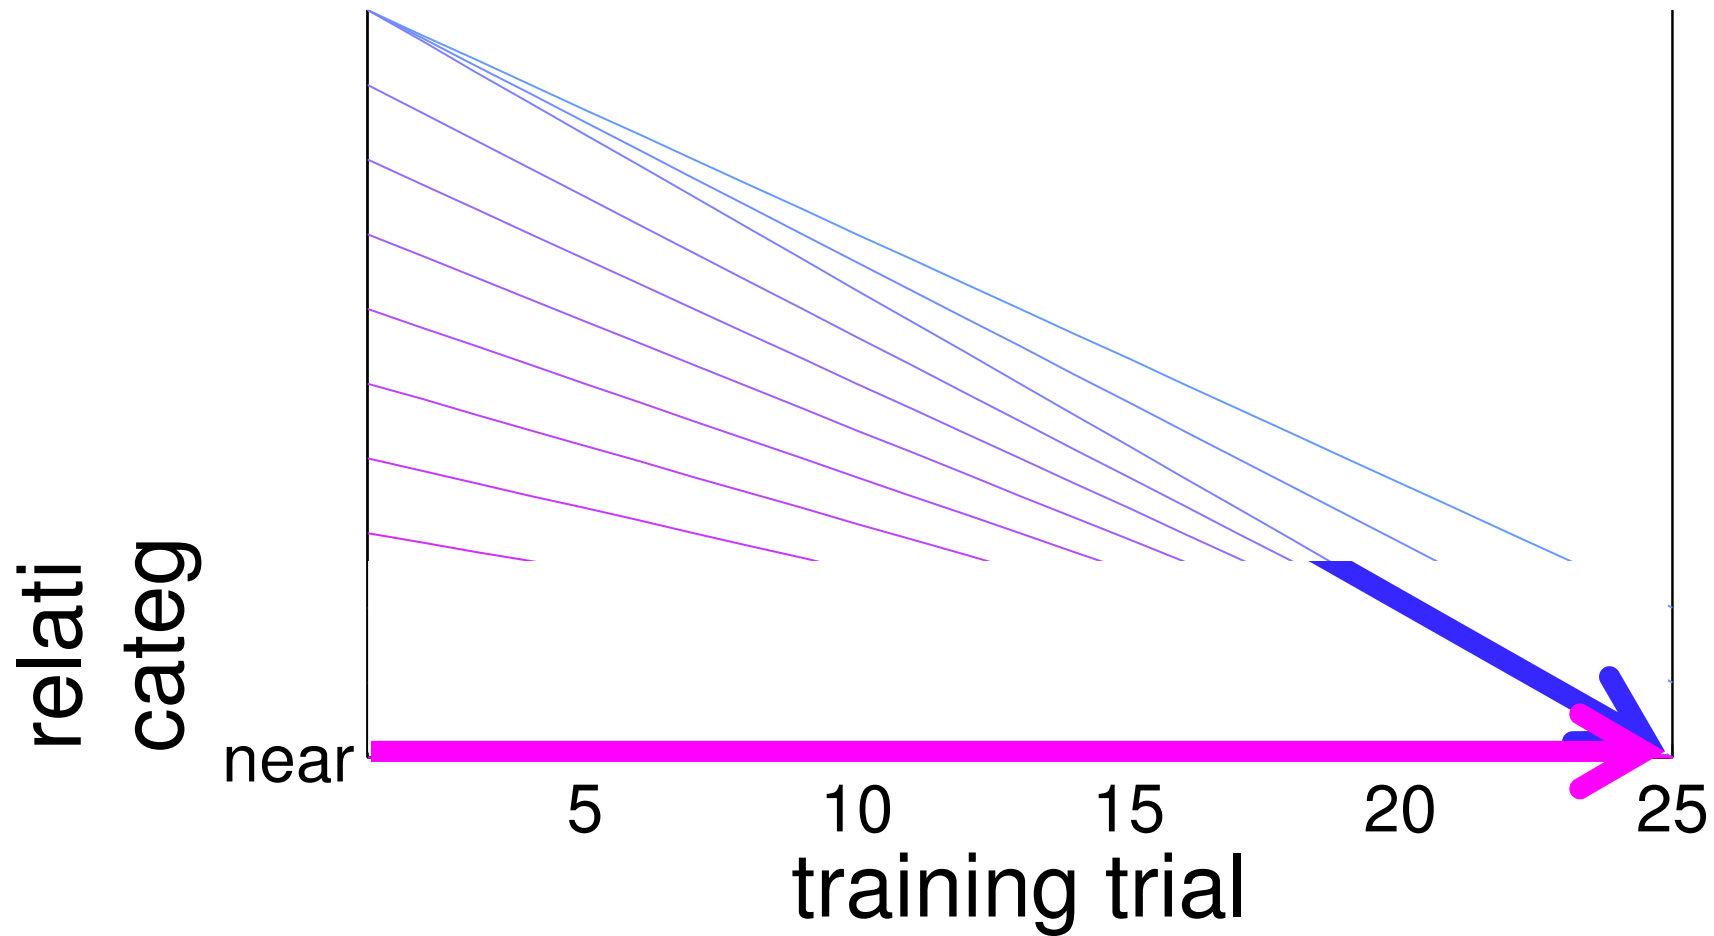
GLOPNOR = Graspability

Ease of picking up & manipulating object with one hand

Based on norms from Salmon, McMullen, & Filliter (2010)



Fading



Blocking

vs.

Interleaving

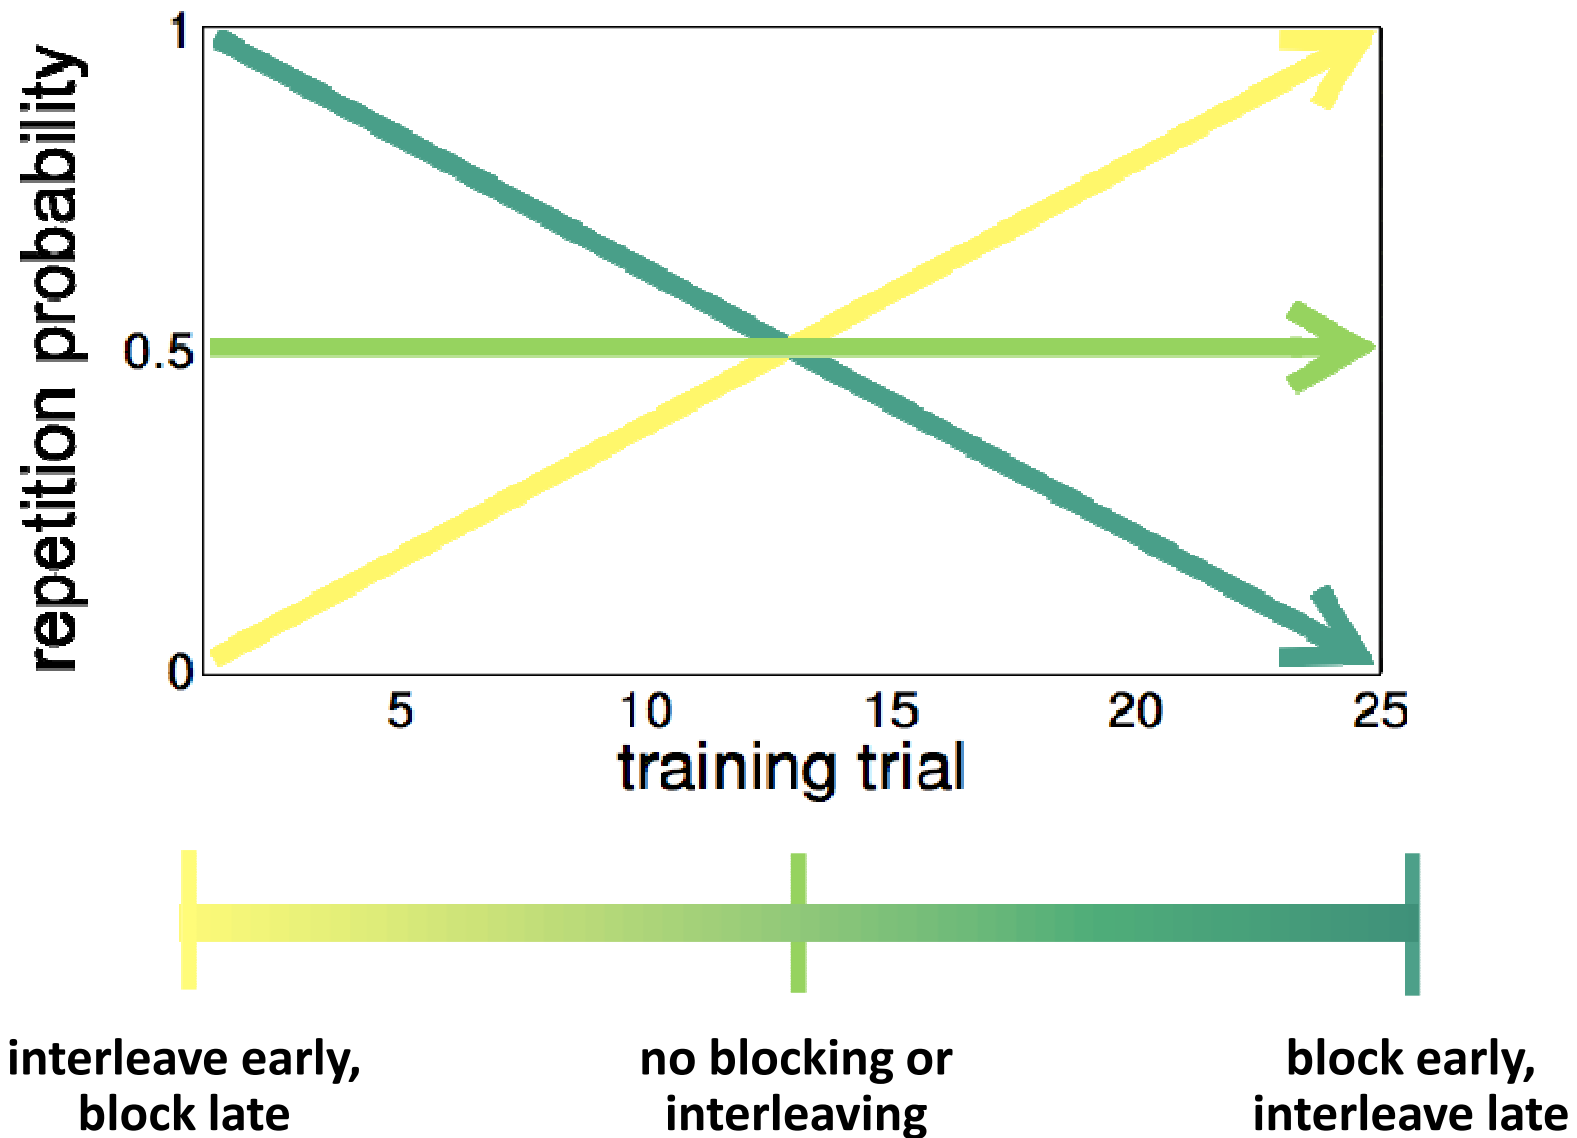
+ + + + - - - -

+ - + - + - + -

**mostly
repetitions**

**mostly
alternations**

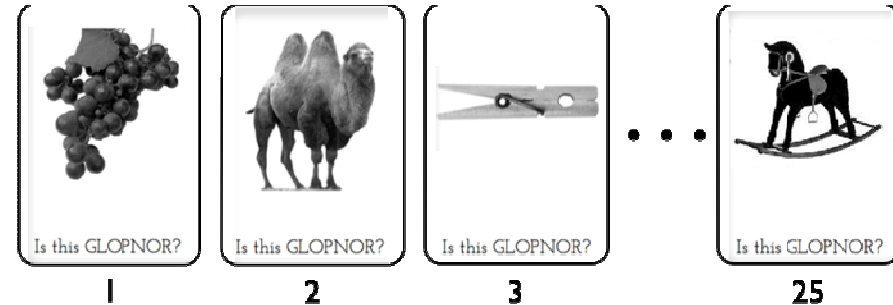
Blocking vs. Interleaving



Concept Learning Experiment

Training

- 25 trial sequence generated by chosen policy



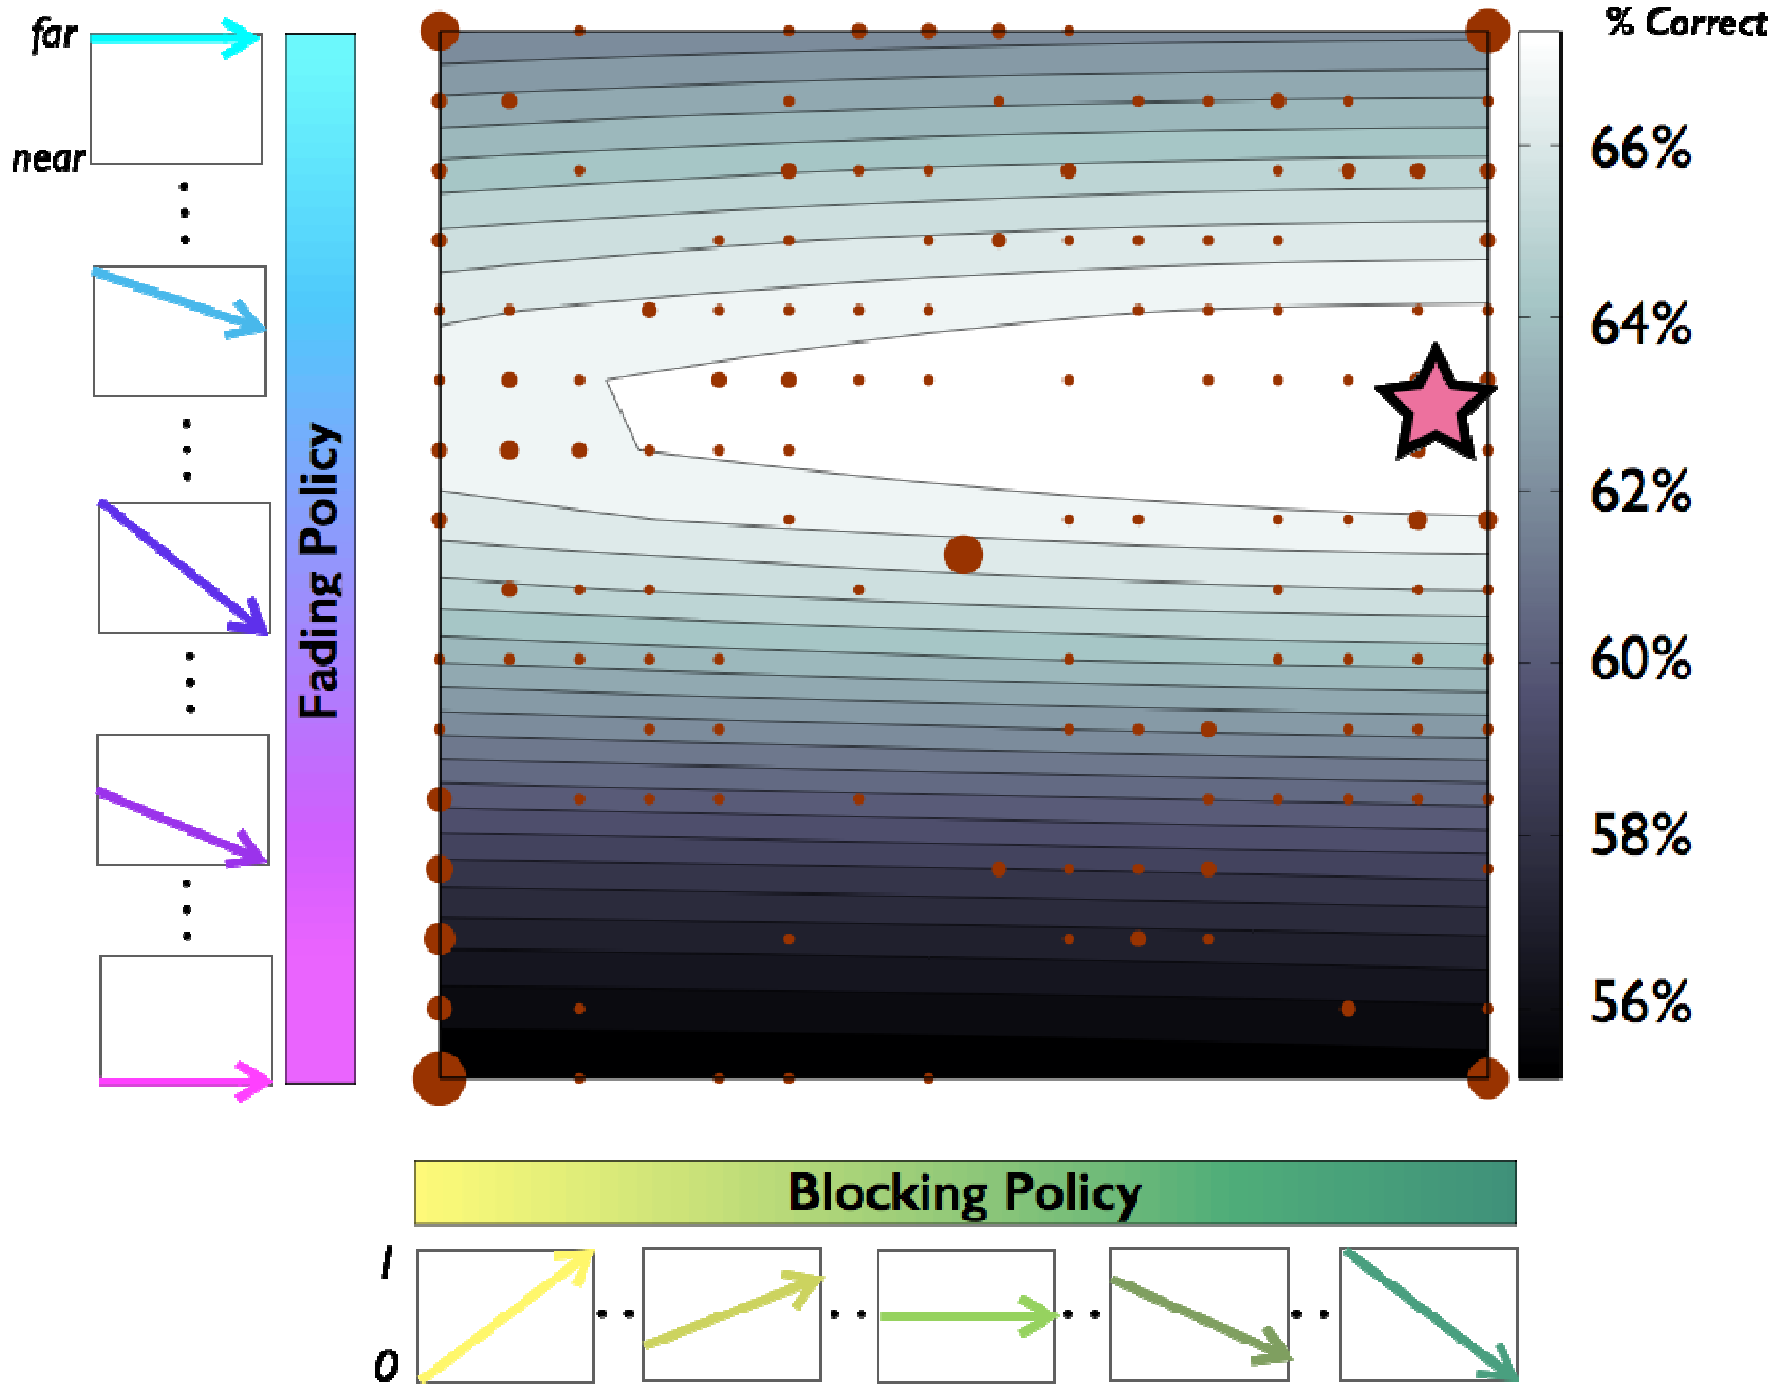
Testing

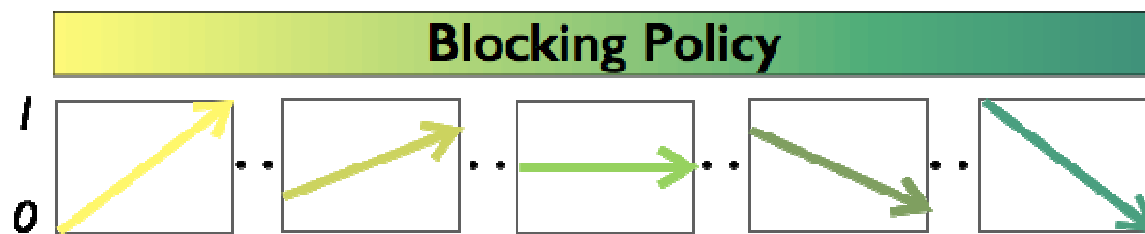
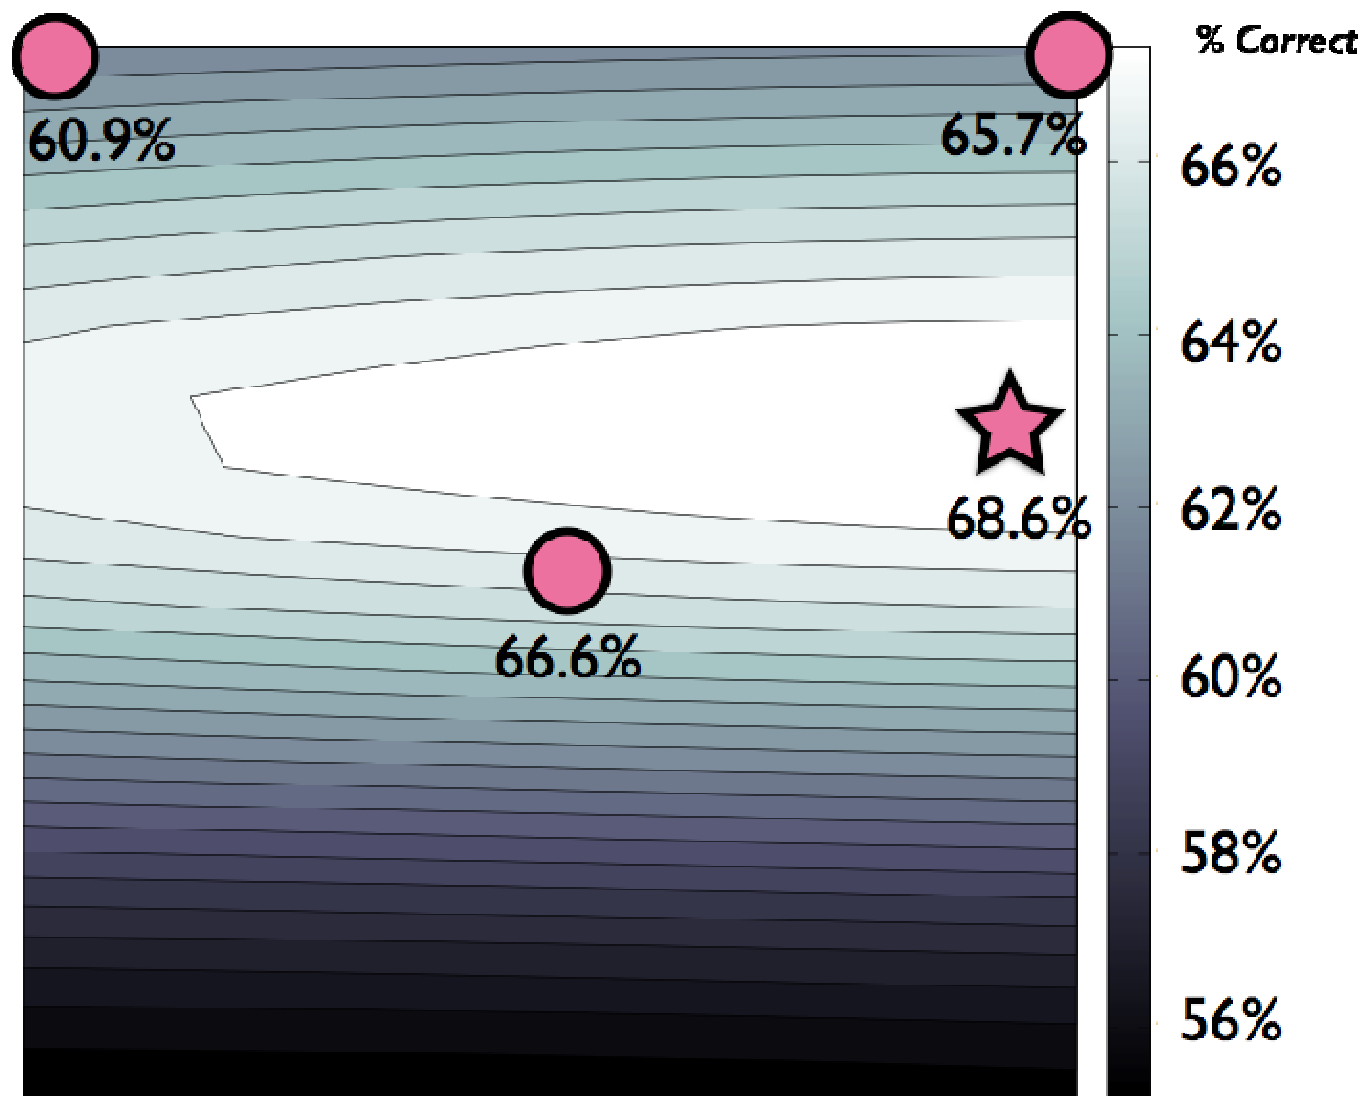
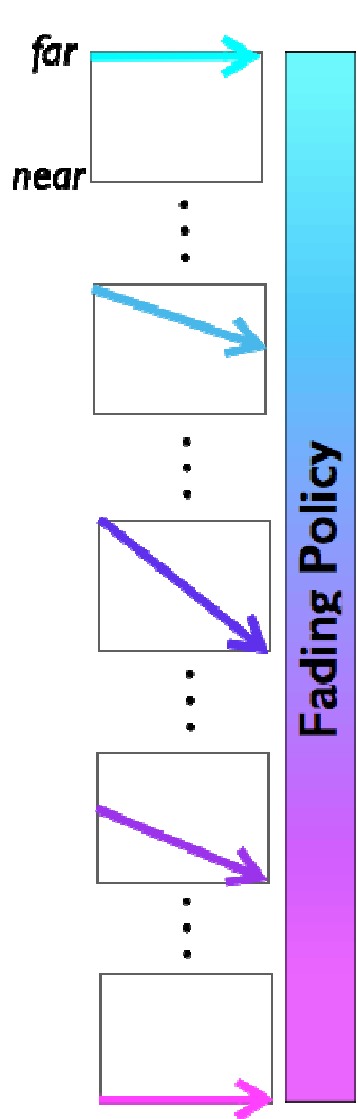
- 24 test trials, ordered randomly
- No feedback, forced choice

Amazon Mechanical Turk

- \$0.25 / subject

Results





Color Aesthetics

Karen Schloss, Brown University

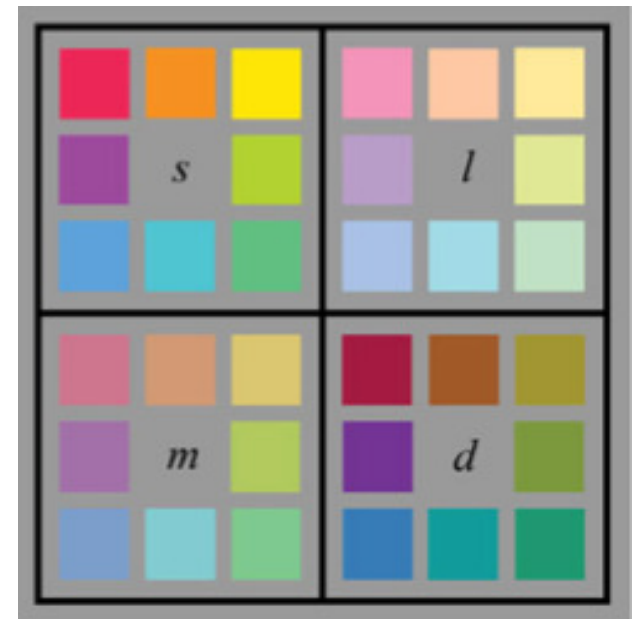
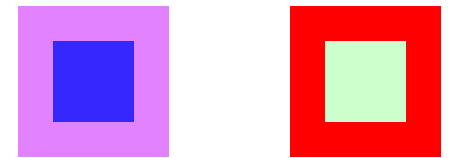
- the perception of color combinations
- how experience shapes preferences
- how preferences influence cognition and decision making



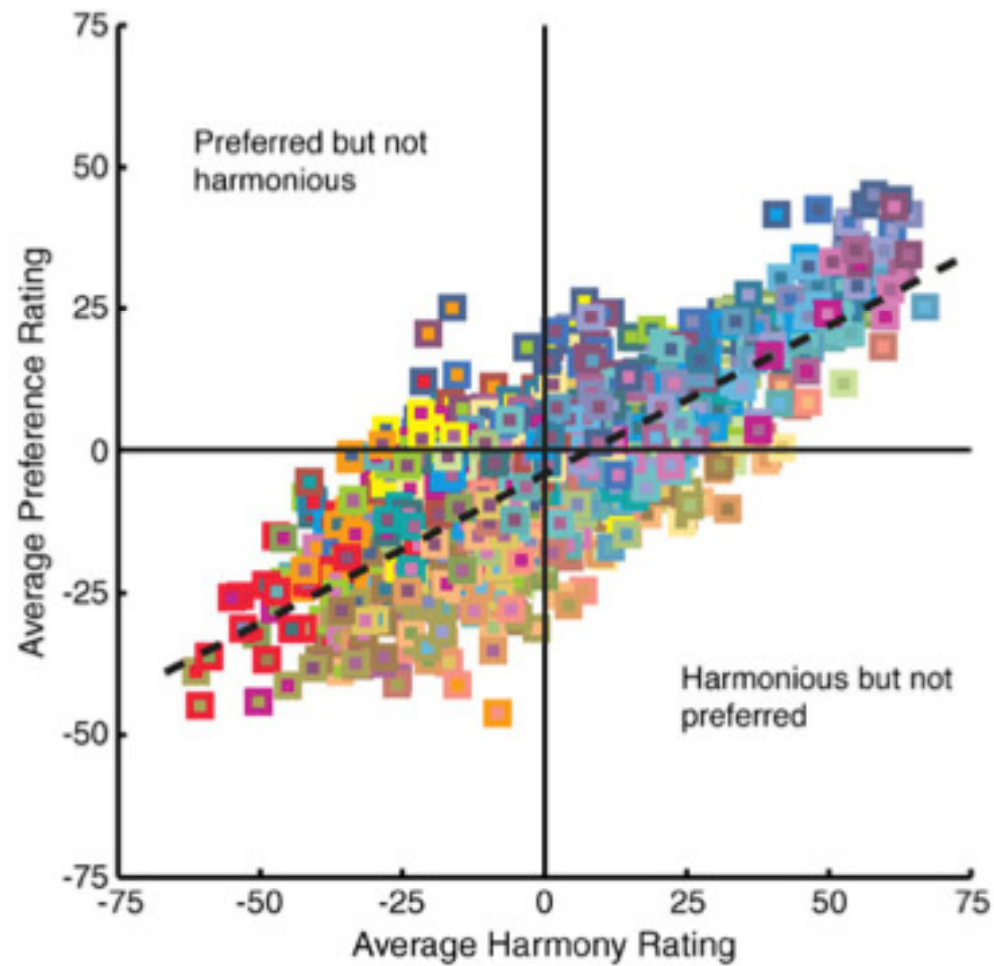
Color Preferences

Schloss and Palmer (2011)

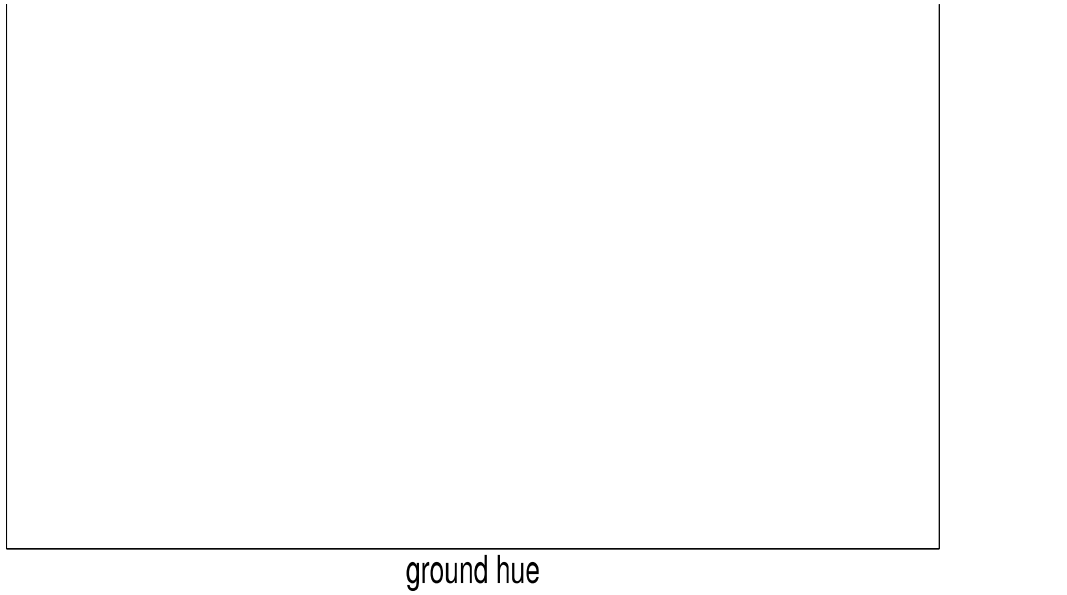
- present a wide variety of color pairs
figure against a background
- asked 48 participants to rate how well the colors go together using a slider
- 32 x 31 color pairs =
992 ratings per participant




Color Preferences



Most And Least Preferred Combinations



Charitable Giving

 University of Colorado
Foundation

GIVING TO CU NEWS & INFO ABOUT JOBS CONTACT

Home » Give Now

Give Now

MAKING A GIFT IS AS EASY AS 1, 2, 3

- 1 Choose where you would like your gift to go
- 2 Enter the amount you would like to give
- 3 Add comments about this gift and add to cart

CU Faculty and Staff may give via payroll deduction. [More information here.](#)

1 CHOOSE WHERE YOU WOULD LIKE YOUR GIFT TO GO

Write-in the name of the fund
you'd like to support here:

OR

Or browse for a fund within a campus:

[Anschutz Medical Campus »](#)
[Boulder Campus »](#)
[Colorado Springs Campus »](#)
[Denver Campus »](#)
[University of Colorado »](#)

2 ENTER GIFT AMOUNT

would like to give:

- ☐ \$
- ☐ \$5000
- ☐ \$1000
- ☐ \$500
- ☐ \$100
- ☐ \$50

☐ This is a recurring gift. [\(more info\)](#)

3 ADD TO CART

- ☐ This gift is part of my pledged amount. [\(more info\)](#)
- ☐ This is an honorary or memorial gift.

To make an honorary or memorial gift to a fund that is not a named honorary or memorial fund, please complete the forms below so we can contact the honoree or next of kin. If you are making a gift to a fund with the honoree's name in the fund title, this information is not necessary.

- ☐ In honor of (for a living person)
- ☐ In memory of (for a deceased person)

Add Comments
on this Gift:

Optimizing Donation Anchors

Turk participants do a bogus task and get paid 5 cents.

Then taken to donation page:

We will give you a 10 cent bonus. You may donate some or all of this bonus to the Red Cross for disaster relief. How much would you like to donate?

- ☐ 1 cent
- ☐ 3 cents
- ☐ 7 cents
- ☐ __ cents

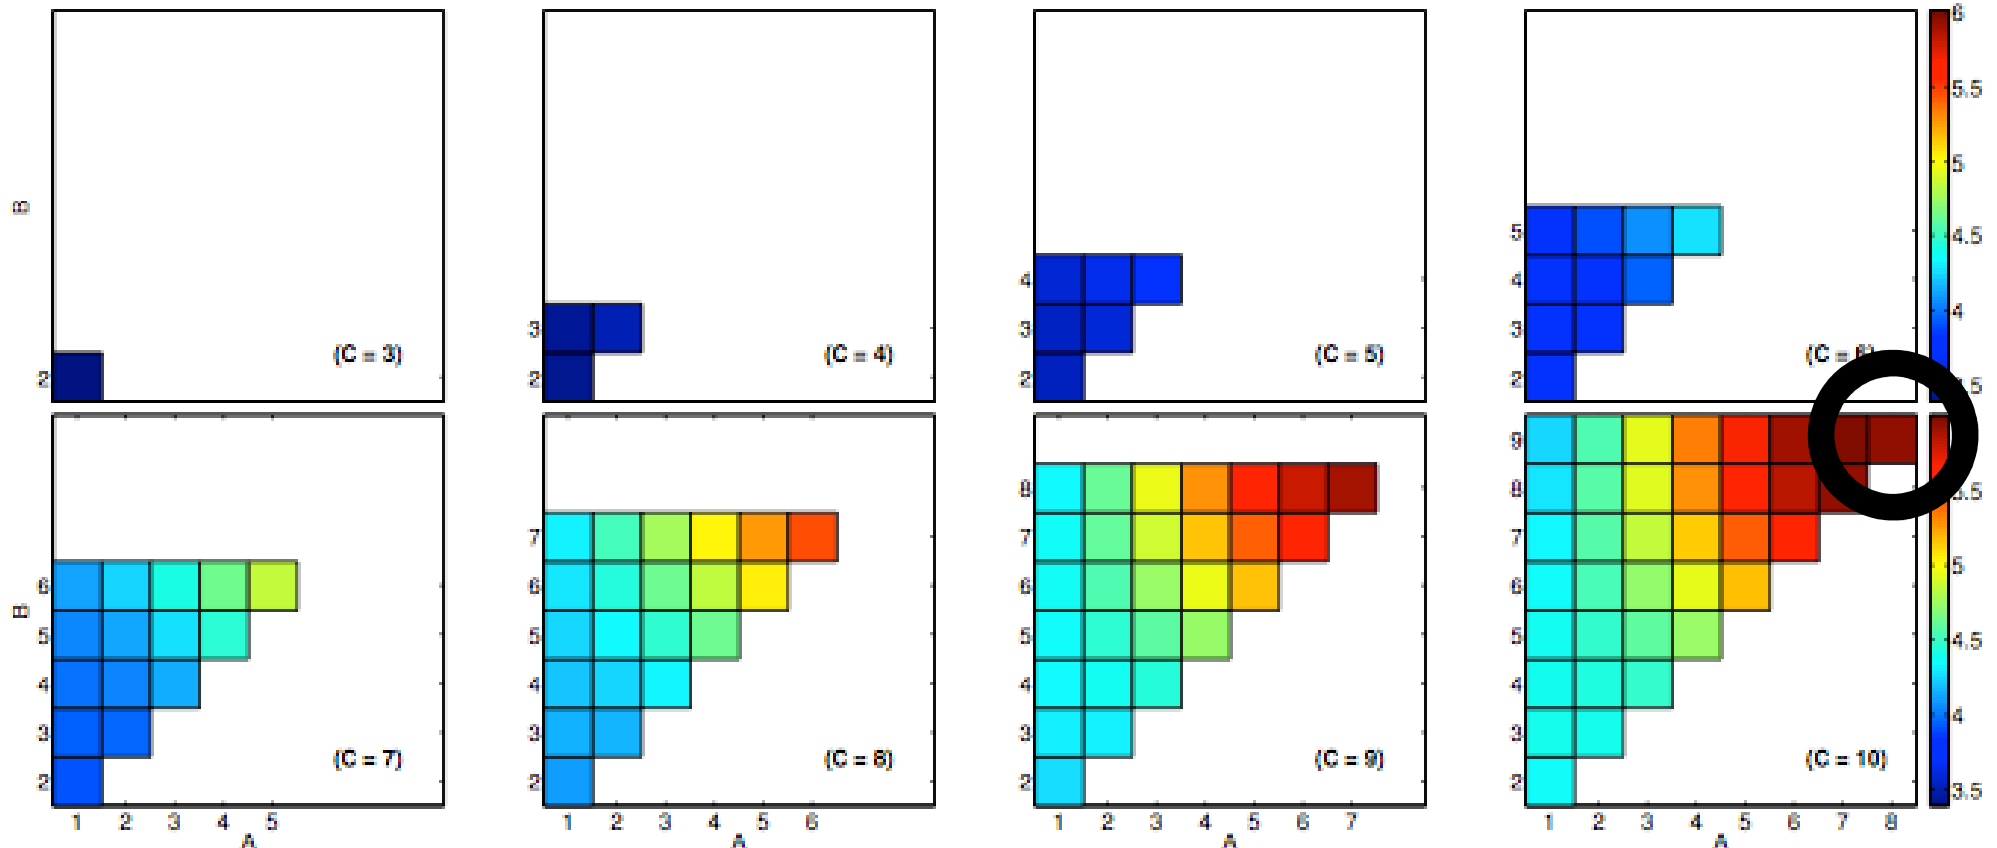
Optimizing Donation Anchors

Anchor triples: (A, B, C)

$A \in \{1, 2, \dots, 10\}$

$B \in \{A+1, \dots, 10\}$

$C \in \{B+1, \dots, 10\}$



Optimum at (8, 9, 10)

New Donation Experiment

- Boring task for 20 trials
- Option to donate more time

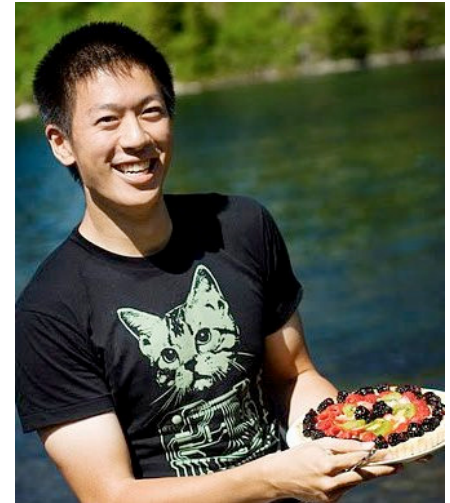
You've earned 5 cents now. We can't pay you any more, but for every additional 20 trials you pledge to do, we'll donate 1 cent to the Red Cross for disaster relief. If you do not complete your pledge, we will not donate.

How much would you like to donate?

- ☐ 1 cent
- ☐ 3 cents
- ☐ 7 cents
- ☐ ___ cents

Making Games Engaging

Yun-En Liu, University of Washington



Treefrog Treasure

- educational game
- solve number line problems, learn fractions
- many variants of game
2 x 2 x 2 x 4 configurations

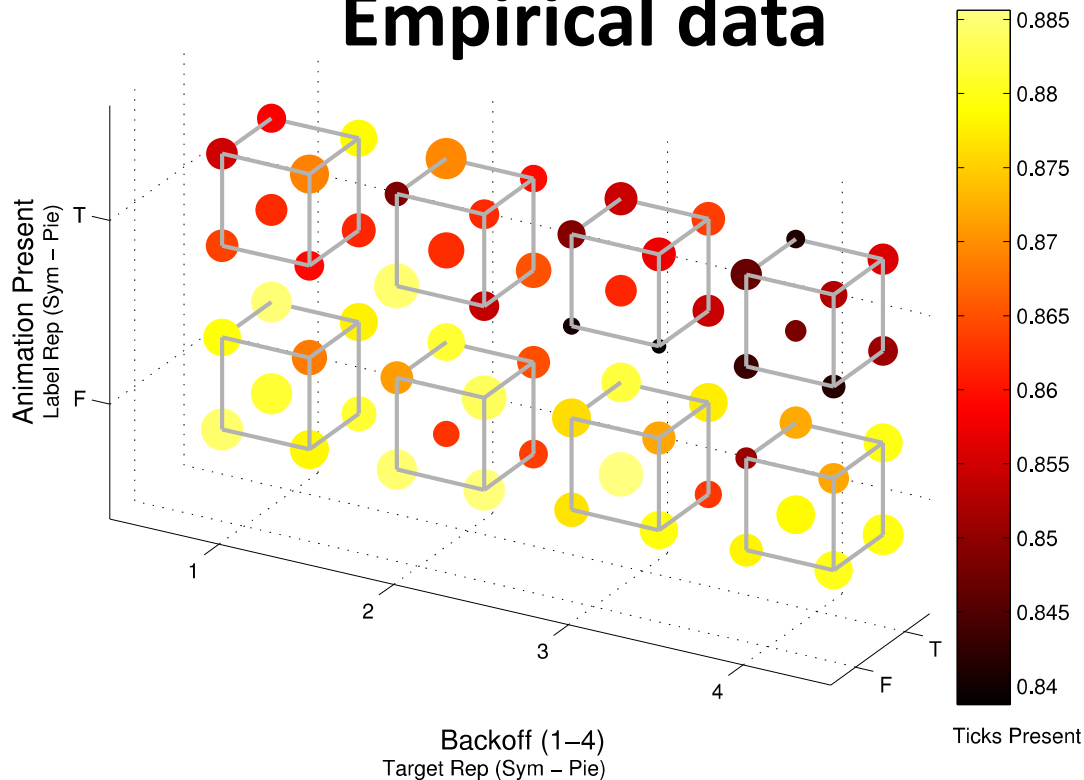


Making Games Engaging

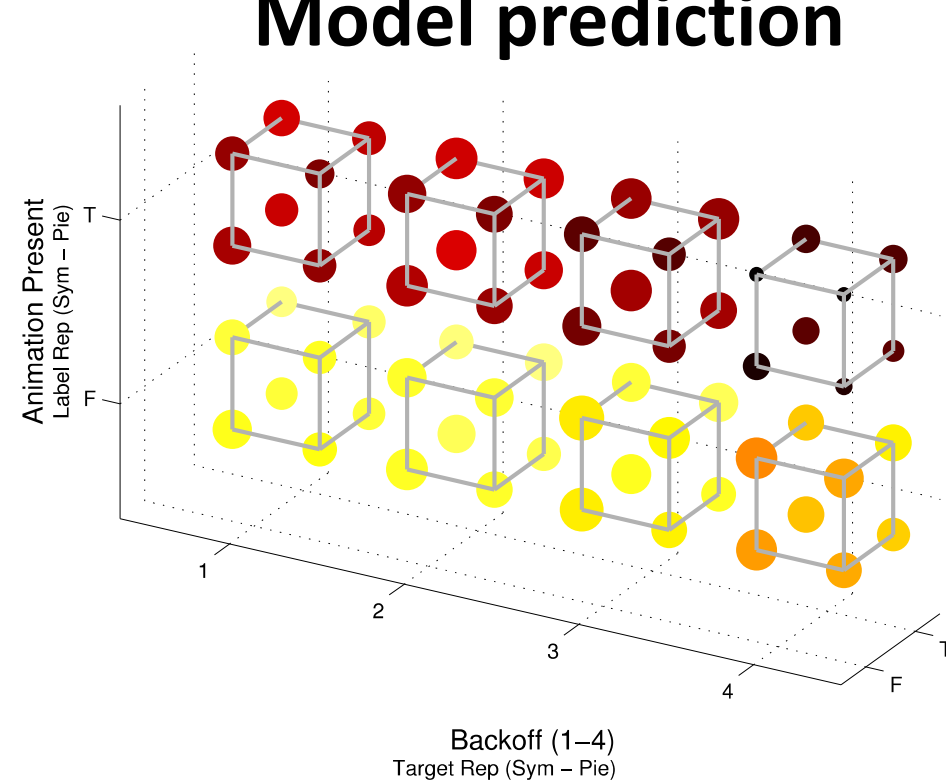
Which game configurations are more/less likely to cause student to quit playing?

- 360k trials, randomly assigned to 64 configurations

Empirical data



Model prediction

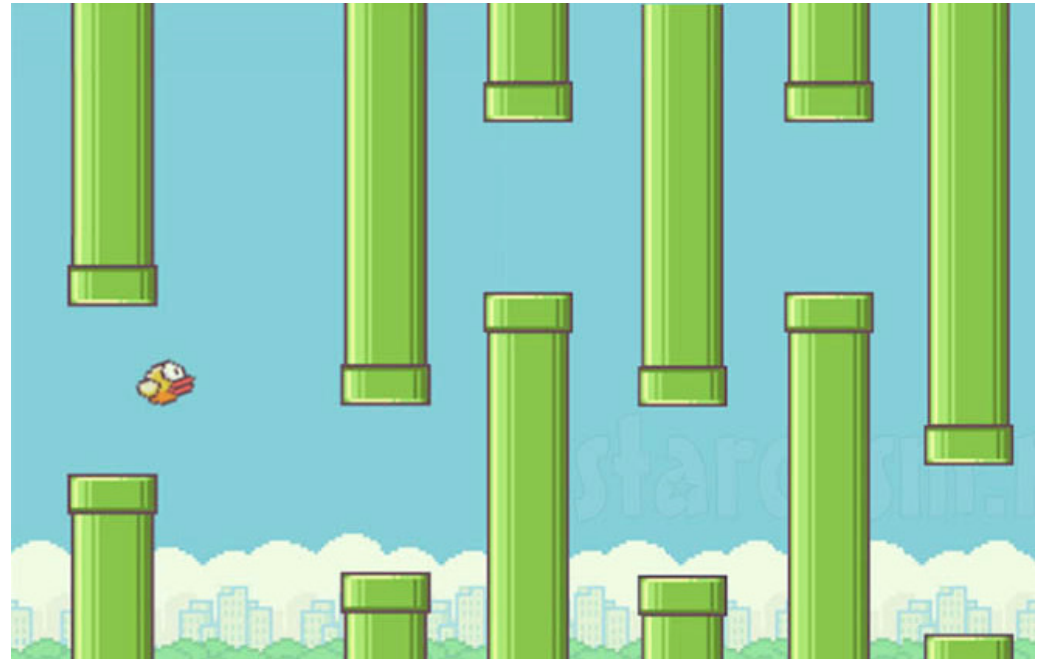


Making Games Engaging II

Flappy bird

Many constants

- gap between pipes
- distance between pipes
- gravitational constant
- wing strength



Can we determine the optimum settings to make game more engaging for a novice?

Bayesian Optimization: A-Z Testing

Alternative to traditional A/B testing

Allows us to efficiently search over a continuum of alternatives to discover an optimum

Machine learning techniques allow us to make stronger inferences from very noisy data.

Do we need this kind of smarts?

- Isn't there an infinite supply of guinea pigs on the web?

Why We Need Bayesian Optimization

More efficient search leads to

- less bad press from running large experiments

-  student learning

- **Furor Erupts Over Facebook's Experiment on Users**

Almost 700,000 Unwitting Subjects Had Their Feeds Altered to Gauge Effect on Emotion

-  ne to individuals, not populations

**Why OKCupid's 'Experiments'
Were Worse Than Facebook's**

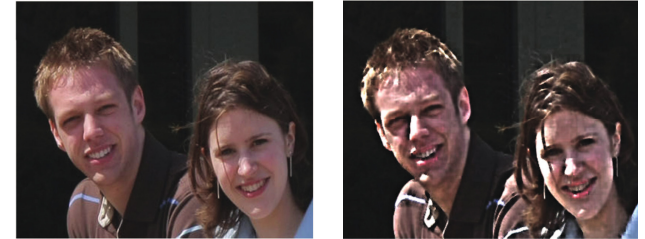
Thank you!

Other Domains

Determine optimal image transform to assist analysts and visually impaired

Learn user-specific relationships

- e.g., Donation anchors as a function of # years since graduation



Satgunam et al. (2012)