Active Learning

Paula Parpart
University College London

KLI 2017, MSc Cognitive and Decision Sciences

Outline

First Part

1. What is active learning and why is it important?
2. What makes a useful/informative query?
3. Active causal learning

Second Part

1. Active learning as a window on representation/ inference
2. Compensatory and Non-compensatory decision models
3. My research

What is active learning and why is it important?

Passive learning

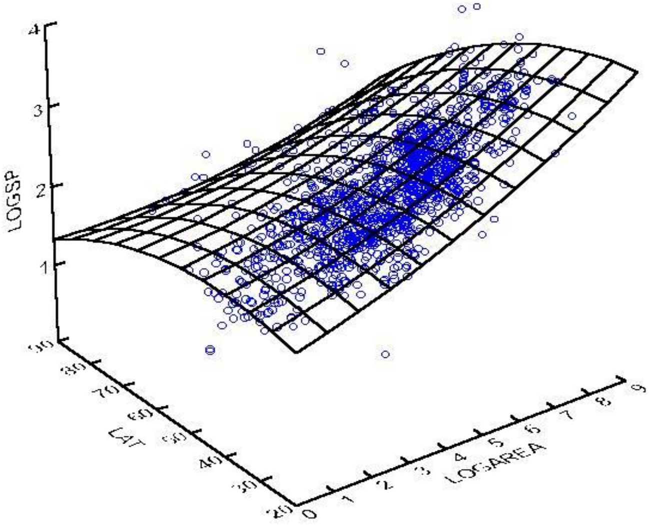

vs. Active learning

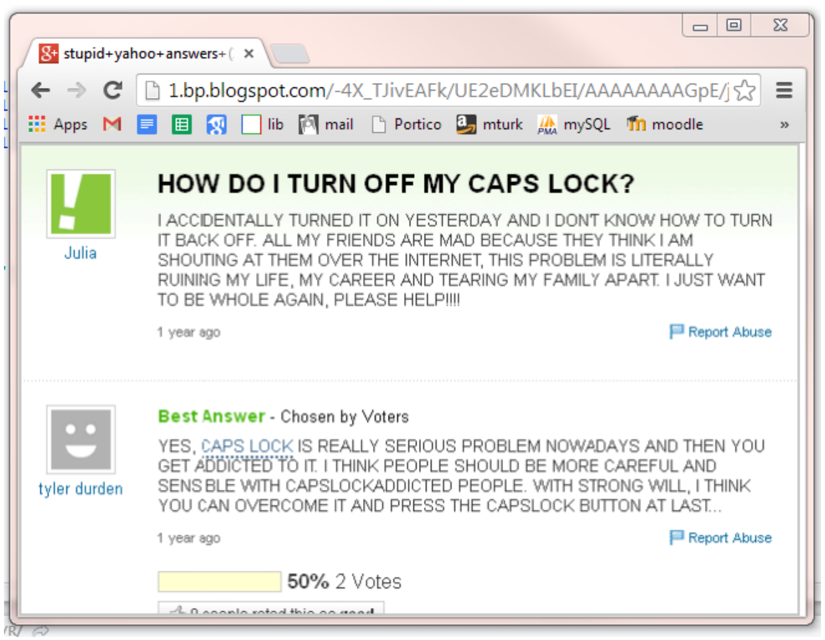

Active Learning

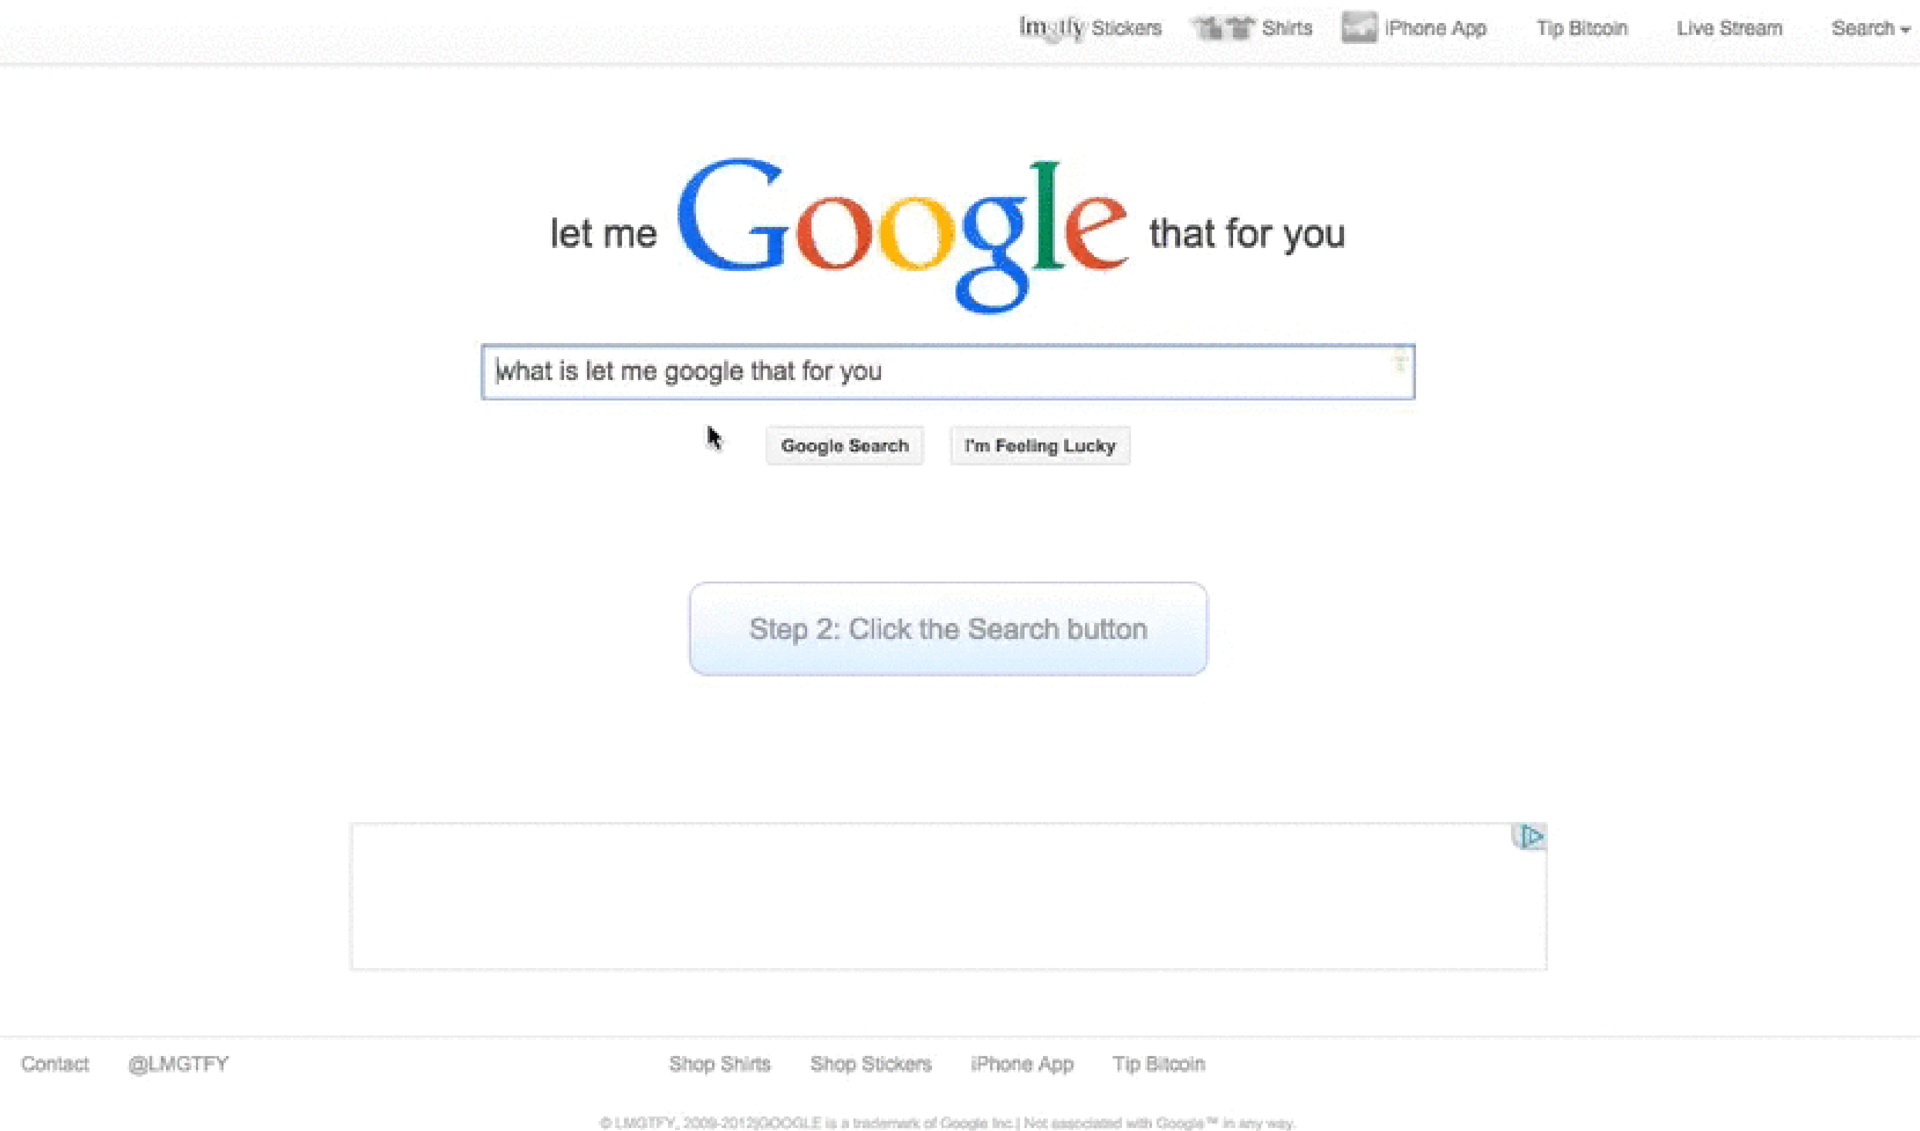

Active learning in cognition

- Smallest unit of active learning, a selection from a set or space of possible *queries* or *actions*
- Higher level cognition:
  - Choosing a test (e.g. medical diagnosis, fault finding)
  - Asking questions (e.g. point at an object and ask its category)
  - Designing an experiment (sometimes called *Optimal Experimental Design* - OED)
- Lower level cognition:
  - Moving the body (e.g. orient head to locate a sound)
  - Directing attention (e.g. sacading efficiently while resolving a scene)

From Gureckis & Markant (2012)

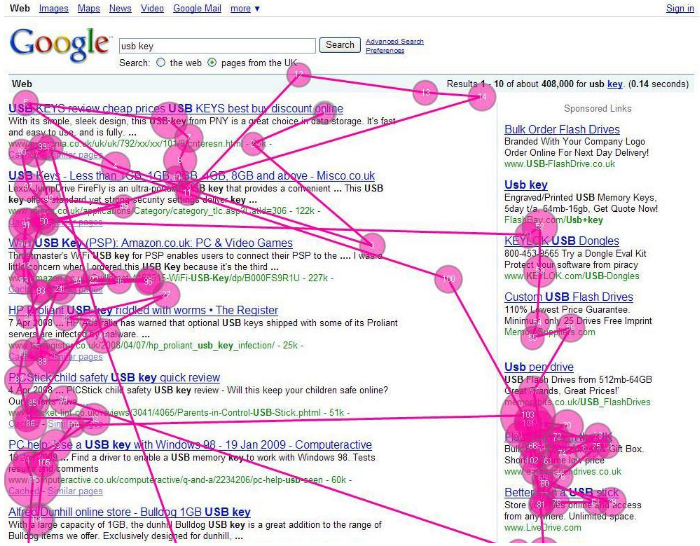

What is active learning not?

- ‘Being active while learning’ (e.g. Hillman et al 2008)
- Brain training (e.g. Ball et al 2002)

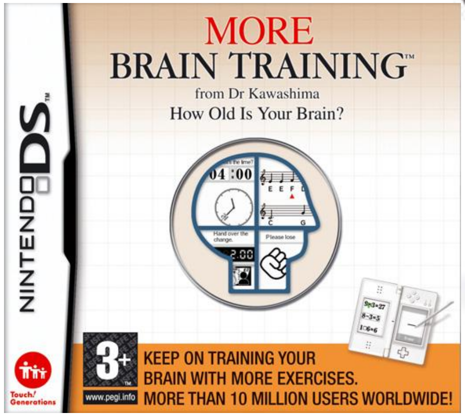

The ‘banana curve’

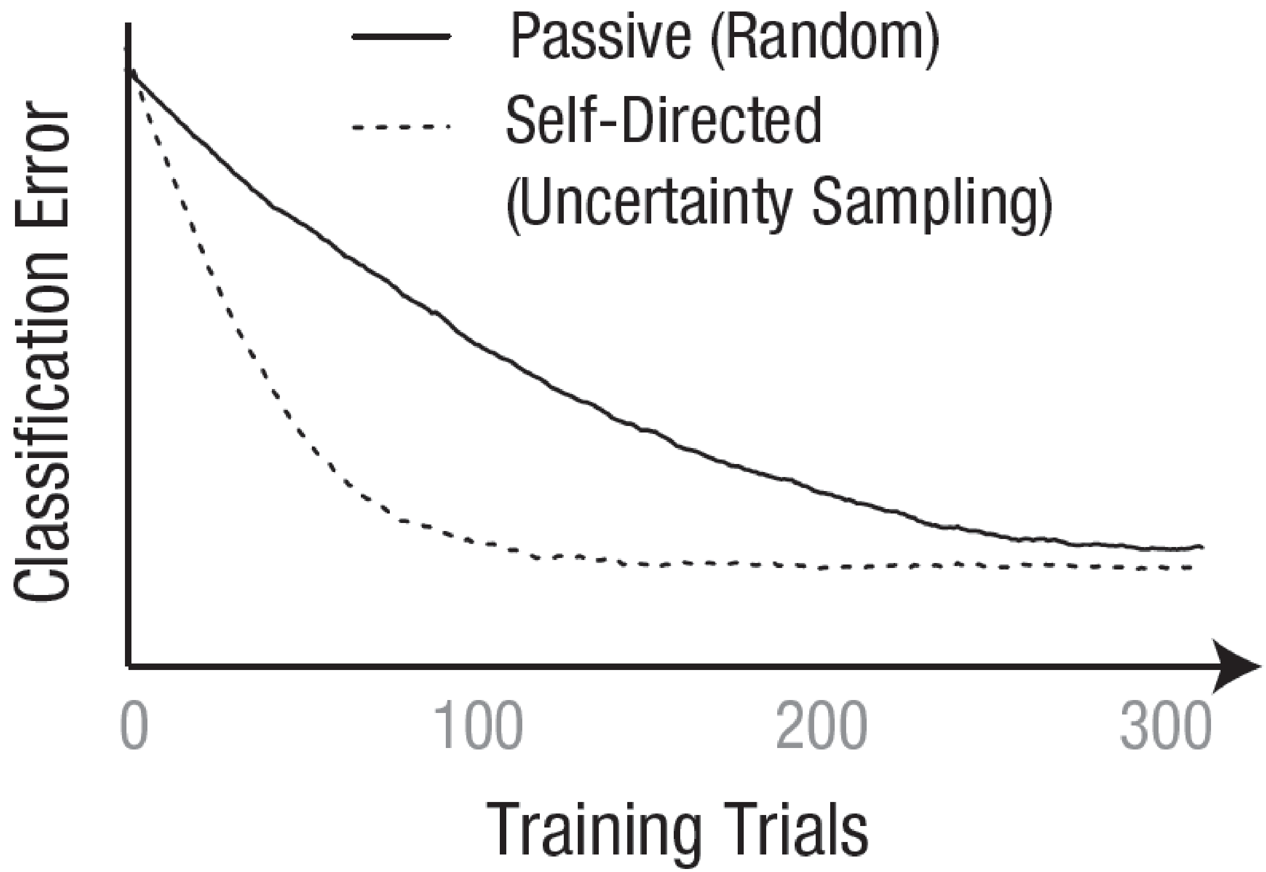

From Gureckis and Markant (2013)

Interim summary

- Active learning research studies how people gather information
- Active learning appears ubiquitous in human cognition
- But is much less studied than passive learning
- Studying information-seeking behaviour can reveal cognitive representation and processes in ways studying passive learning cannot
- To study active learning, must understand the computational level problem, i.e. how to assess the ‘informativeness’ of different actions / queries

Psychological Theories of Active Learning

- Goal: find the true hypothesis out of many potential explanations

- Uncertainty about value of some variable: (Shannon, 1948)

$$H(O) = -\sum_{o} P(o) \log_2 P(o)$$

- Maximize an informational utility of the possible queries: e.g., ***information gain*** is the reduction in uncertainty due to seeing some data

$$I(C) = H(O) - H(O \mid C)$$

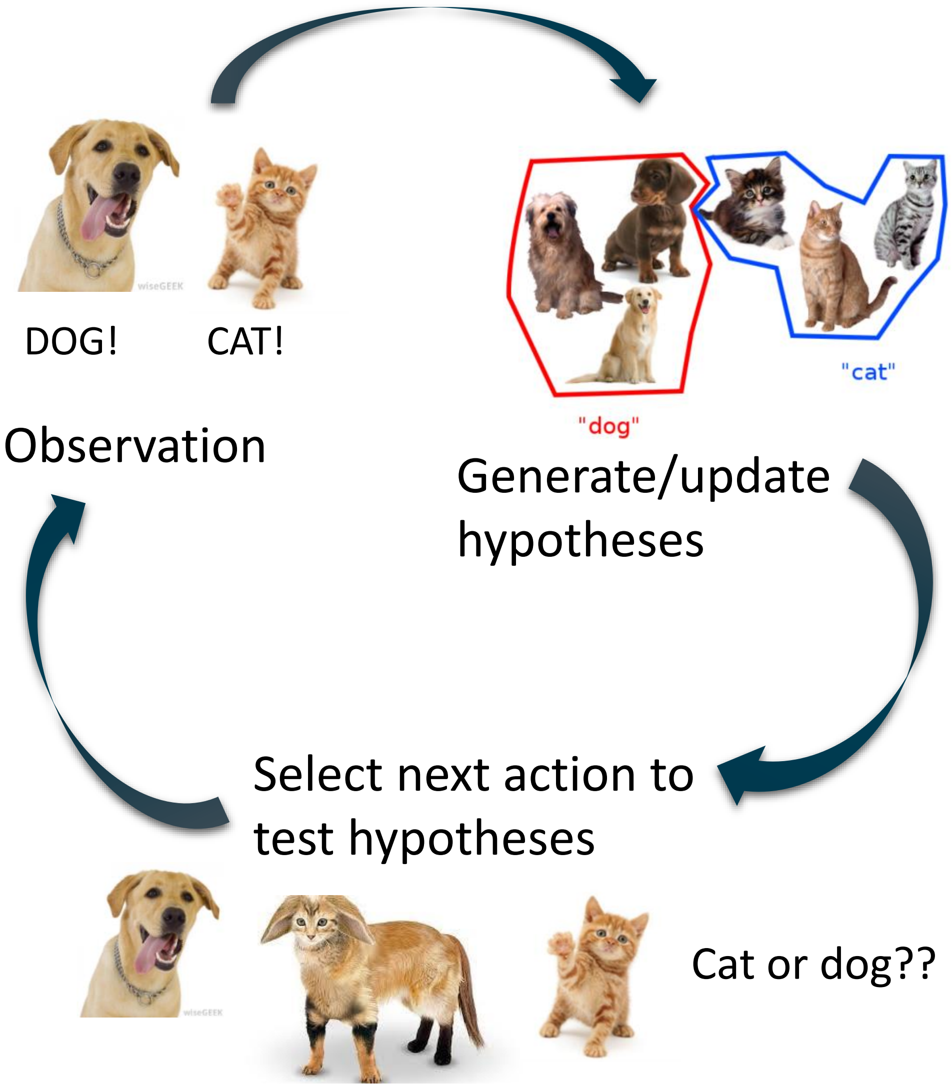

Technical stuff

- *Expected usefulness* is the average usefulness of the possible outcomes of an action weighted by their probability

$$E[U(R,A)] = \sum_{r \in R} U(r,a)p(r|a)$$

Other measures

Probability gain

$$\max_{s \in S} p(s|r,a)$$

- Which outcome gives best chance of guessing *S*
- Optimal if you must make a guess right after

Kullback-Leibler divergence

$$p(s|r,a)\log_2 \frac{p(s|r,a)}{p(s)}$$

- Widely used measure of *difference* between distributions
- Gives same answers as information gain in expectation
- …but individual values reflect amount of *belief change*, rather than *uncertainty reduction*

- See **Nelson (2005)** for an accessible introduction and benchmarking of different measures

What makes a query/action informative?

- Computational level characterisation of active learning:
  - Calculate the value of different actions through preposterior Bayesian analysis, choose the highest

- Can be applied to planning multiple actions/queries into future - but rapidly becomes computationally infeasible

- Can compare long run performance when applied “greedily” (i.e. to choose a series of actions/queries one after another)

Active Causal Learning - Intervention

- Active learning especially important for inferring causality
- Correlation ≠ causation
- Manipulating a system can reveal causal structure
- Pearl (2000) - can formalise interventions as ‘graph surgery’ on a Causal Bayes network (CBN), helping to reveal the true CBN for an encountered system

a)

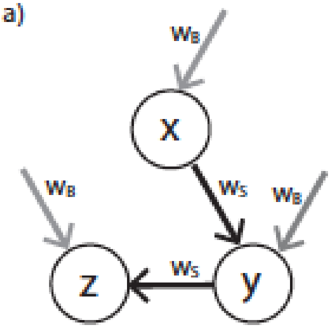

x	
0	$1-w_B$
1	w_B

y	x=0	x=1
0	$1-w_B$	$(1-w_B)(1-w_S)$
1	w_B	$1-(1-w_B)(1-w_S)$

z	y=0	y=1
0	$1-w_B$	$(1-w_B)(1-w_S)$
1	w_B	$1-(1-w_B)(1-w_S)$

$p(d|m) = p(x)\, p(y|x)\, p(z|y)$

b)

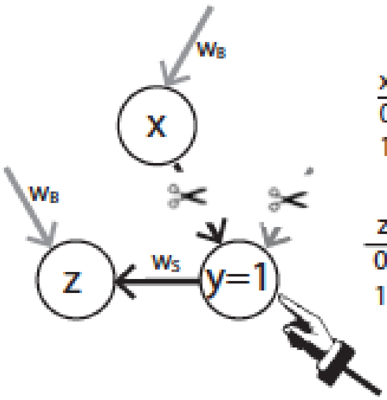

x	
0	$1-w_B$
1	w_B

z	y=1
0	$(1-w_B)(1-w_S)$
1	$1-(1-w_B)(1-w_S)$

$p(d|m) = p(x)\, p(z|y=1)$

Conservative Forgetful Scholars (Bramley et al, 2015)

1. Select an intervention
2. Observe the results
3. Update marked connections
4. Repeat

After 12 trials:

4. Get feedback, get paid according to correct/incorrect connections

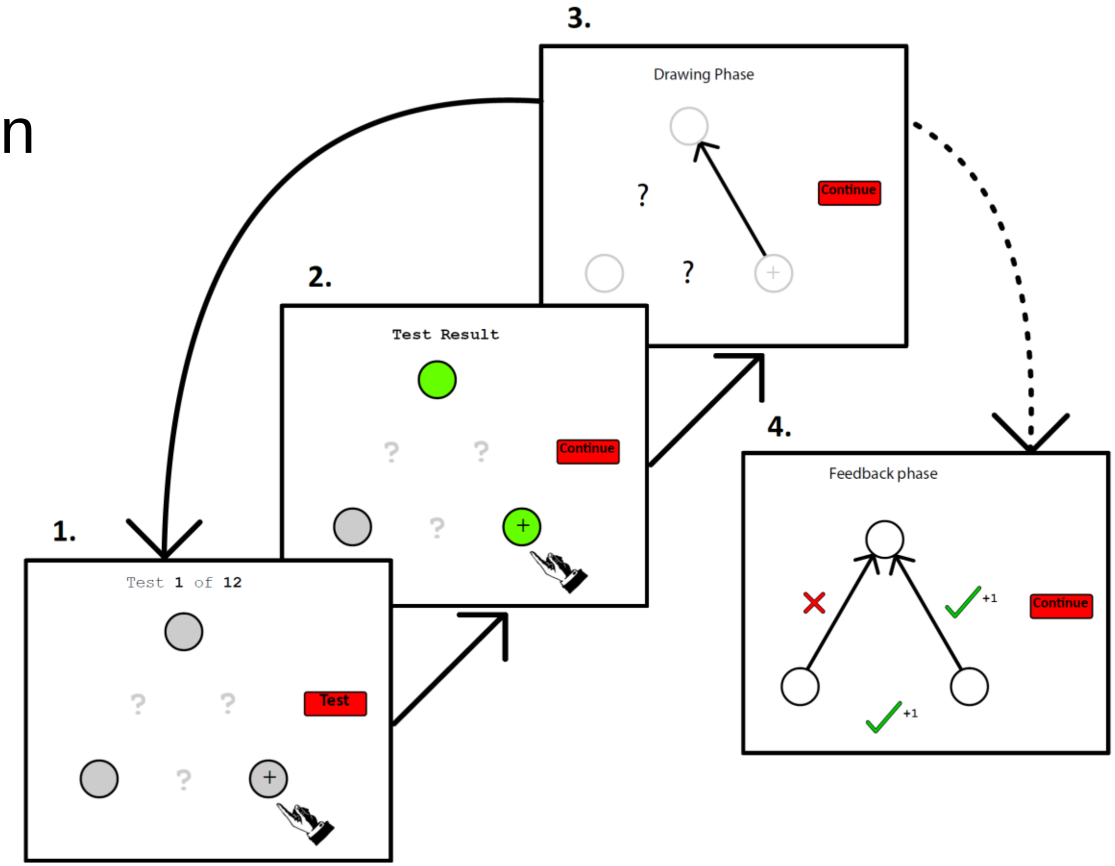

Conservative Forgetful Scholars (Bramley et al, 2015)

- Fit various models to participants sequences of interventions
- Participants better described as maximising ***information gain*** (‘scholar model’), than expected payment (‘utilitarian’ model), or *probability gain* (‘gambler model’)
- Participants also behaved like they were **highly *forgetful* about outcomes of previous tests**
- …but compensated by being *conservative* in their model changes, sticking close to their latest model, changing few connections at a time

Conservative Forgetful Scholars (Bramley et al, 2015)

- Participants also behaved greedily, i.e. better fit by models that optimised learning at the next time step rather than planning several steps into the future

Active learning as a new method for model discrimination

- Long-standing debate in decision making: Do people use non-compensatory (Take-The-Best Heuristic) or compensatory (Logistic Regression) decision strategies?

Compensatory versus Non-compensatory Strategies

or

X =  wins

✓ =  wins

(1) My flavour ? X

(2) Cost ? ✓

(3) Size ? ✓

(4) Texture ? ✓

(5) Appearance ? ✓

Compensatory strategy

Example: weighted-additive rules (WADD), linear/logistic regression

Compensatory versus Non-compensatory Strategies

or

(1) My flavour ? X

(2) Cost ? ✓

(3) Size ? ✓

(4) Texture ? ✓

(5) Appearance ? ✓

X =  wins

✓ =  wins

Non-compensatory strategy

Example: Take-The-Best Heuristic (Gigerenzer & Goldstein, 1996)

Experiment: Which Alien will win the fight?

or

?

4 Features

Antennae

Wings

Camouflage

Diamonds

Compensatory: Logistic Regression

or

Antennae	✓
Wings	–
Camouflage	✓
Diamonds	X

X	 wins
–	neither wins
✓	 wins

Compensatory: Logistic Regression

or

Antennae	✓	**1**
Wings	–	**0**
Camouflage	✓	**1**
Diamonds	X	**-1**

X	 wins
–	neither wins
✓	 wins

Compensatory: Logistic Regression

or

			Weights
Antennae	✓	1	.66
Wings	–	0	.81
Camouflage	✓	1	.98
Diamonds	X	-1	.30

X	 wins
–	neither wins
✓	 wins

Compensatory: Logistic Regression

			Weights
Antennae	✓	**1**	**.66**
Wings	–	**0**	**.81**
Camouflage	✓	**1**	**.98**
Diamonds	X	**-1**	**.30**

Compensatory strategy:
Logistic Regression

Non-compensatory: Take-The-Best

or

First Step: RANK ORDER THE CUES.

			Weights
(1) Camouflage	✓	1	.98
(2) Wings	–	0	.81
(3) Antennae	✓	1	.66
(4) Diamonds	X	-1	.30

Take-The-Best stops search as soon as it finds a cue that discriminates.

Non-compensatory: Take-The-Best

			Weights
(1) Camouflage	✓	1	.98
(2) Wings	–	0	.81
(3) Antennae	✓	1	.66
(4) Diamonds	X	-1	.30

Take-The-Best stops search as soon as it finds a cue that discriminates.

Non-compensatory strategies (Take-The-Best)	Compensatory strategies (Logistic Regression)
• Later cues cannot compensate for earlier cues. 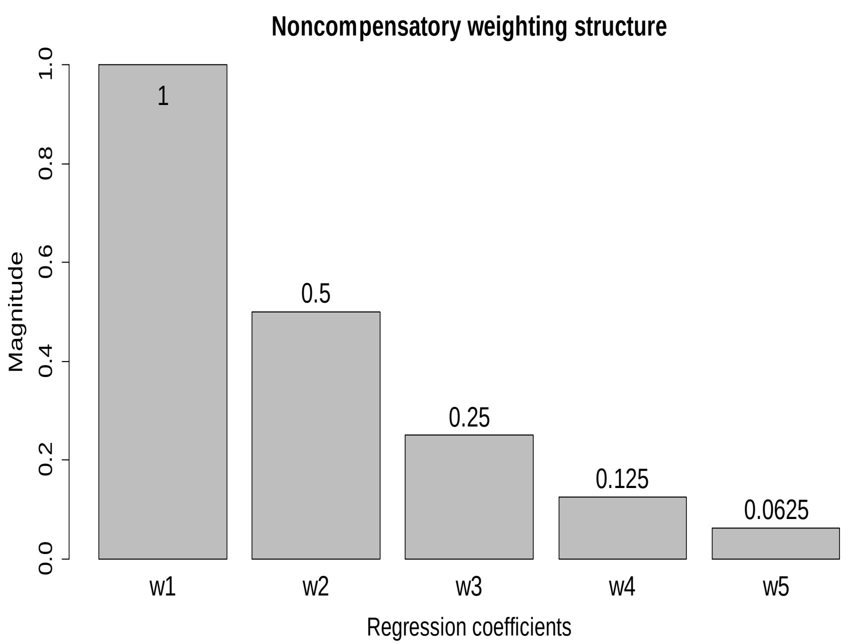	• Later cues can compensate for later cues. 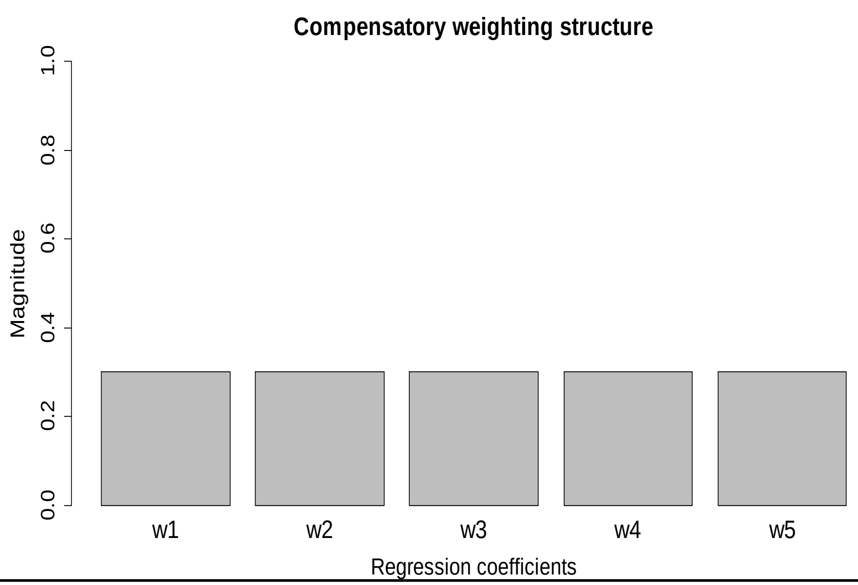

How to assess whether people rely on compensatory or non-compensatory strategies?

1) Todd & Gigerenzer, 2000: non-compensatory strategies are simpler and require less computational capacity and are therefore more plausible

Traditional model testing approaches

1. **Model fitting** to human behaviour in a highly controlled, passive experiment

   - Model fitting often does not distinguish between models. Danger of mimicry of strategies (e.g., Czerlinski et al., 1999; Chater et al., 2003)

2. **Cross-validation**: Pitting decision making models against each other in a computer simulation to compare their predictive accuracy (i.e., generalization performance)

→ Nevertheless, just because one class of models can beat another with better predictions, **it does not follow that this class is necessarily a better psychological representation of what people actually do.**

Traditional model testing approaches

3. **Novel Approach**: **Active Learning** as a way to differentiate among different decision making models.

- Decision making models
- Category learning models
- …

Argument (Parpart et al., 2015):

If a cognitive agent has evolutionarily developed to prefer a certain class of models as her/his means to learn a cognitive representation in a particular environment, then the way he/she actively selects information should reflect this representation.

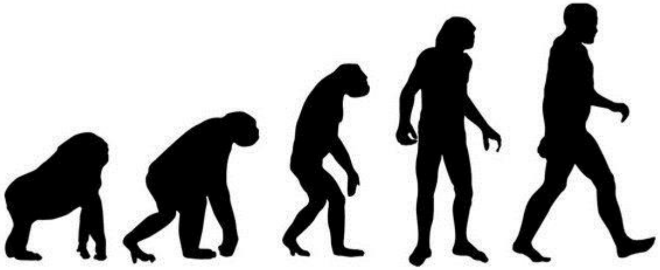

Argument (Parpart et al., 2015):

For example, if an agent has come to apply a non-compensatory strategy (e.g., TTB), then –intuitively she should try to establish a rank order among cues first as this will decrease her uncertainty maximally. That is, she will try figure out what is the first best cue, the second best, third best and so on.

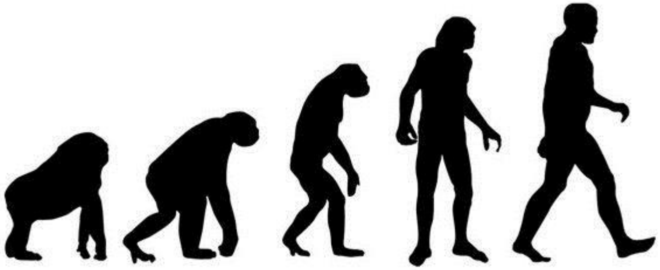

Model-based active learning (Parpart et al., 2015, submitted):

- We introduce model-based active learning as a method to compare psychological models.

- Model-based active learning relies on the assumption that an agent's information gathering behaviour reflects the psychological model that best describes the agent’s cognitive strategy.

Model-based active learning (Parpart et al., 2015, submitted):

- We formalize this assumption as a generalized way in which psychological models can be defined and empirically tested and show how model-based active learning distinguishes better between candidate models than either pure fitting or cross-validation

- Given that psychology seems to be in dire need for better ways to test its candidate models, we believe that our approach is a valuable addition to its methodological tool kit.

Theories of Active Learning

- **Information gain** = uncertainty reduction after seeing some evidence:

$$I(T) = H(x) - H(x|T)$$

- Prior uncertainty (Shannon entropy, 1948):

$$H(x) = -\sum_x P(x)\log_2 P(x)$$

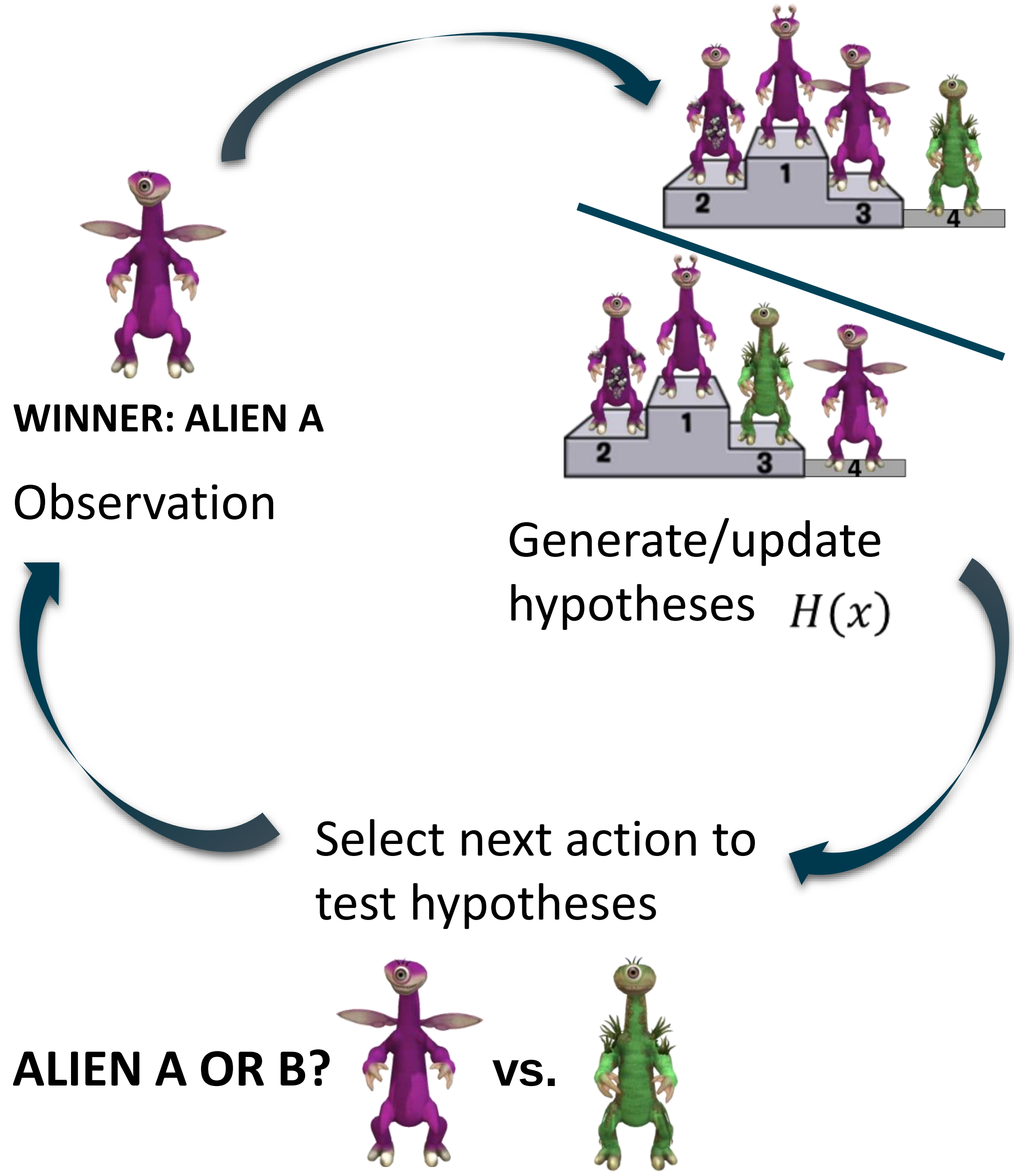

Question

- Active Learning Question: Do people learn with respect to **cue weights** (Regression) or **cue orders** (TTB)?

- There are no active algorithms yet for heuristics, so we developed an active learning algorithms for the TTB Heuristic as well as Logistic Regression.

Model 1: Active Take-The-Best

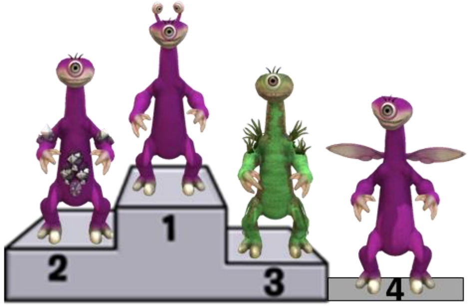

- Tries to find the underlying **cue rank order.**

- Uniform prior over all possible cue orders of available cues.

- Computes expected entropy for each comparison and chooses best.

- posterior entropy updated every time a binary comparison is made.

- Greedy algorithm

Model 1: Active TTB

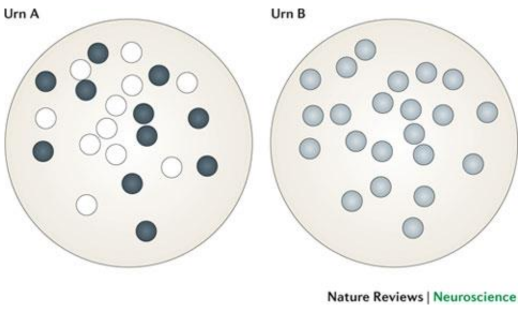

- We put a uniform prior over all possible cue orders
- Calculate prior Shannon entropy

$$O_0 = \sum_i p(o)_i \log p(o)_i$$

- For every comparison c (e.g., Alien 1 versus Alien2) calculate $p(y=w)$ and $p(y=l)$
- Calculate posterior expected uncertainty

$$\mathrm{E}\left[O \mid c\right] = O(y=W) \times p(y=W) + O(y=L) \times p(y=L)$$

- Choose $c^* = \arg\max \left\{ O_0 - \mathrm{E}\left[O \mid o\right] \right\}$

Model 2: Active Logistic Regression

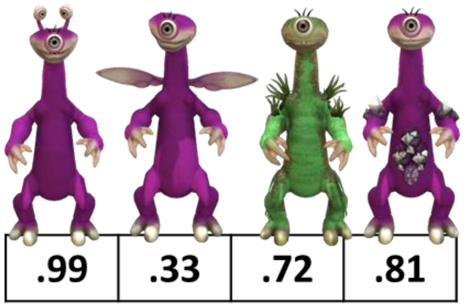

- Tries to learn underlying **cue weights** as precisely as possible.

- Entropy over **cue weights** as the sum of the coefficients' uncertainty

$$S = \sum_k V(\beta_k)$$

- Computes expected entropy for each comparison and chooses best.

- Updates the posterior expected entropy every time a binary comparison is made.
- Greedy algorithm

Model 2: Active Logistic Regres

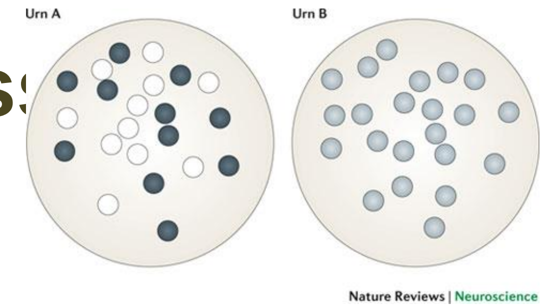

- Given a Bayesian variant of logistic regression:

$$f(x) = \frac{1}{1+\exp(-(\beta_0+\sum_k \beta_k x_k))}$$

- Calculate sum of coefficients' uncertainty $S = \sum_k V(\beta_k)$
- For every comparison c, calculate $p(y=W\,|\,c)$ and $p(y=L\,|\,c)$
- For every comparison and outcome calculate $S\,|\,c, y=L$ and $S\,|\,c, y=W$
- Calculate posterior expected uncertainty

$$\mathrm{E}\left[S\,|\,c\right] = S(y=W)\times p(y=W) + S(y=L)\times p(y=L)$$

- Choose $c^* = \arg\max\left\{S_0 - \mathrm{E}\left[S\,|\,c\right]\right\}$

Passive logistic

Passive Logistic Regression

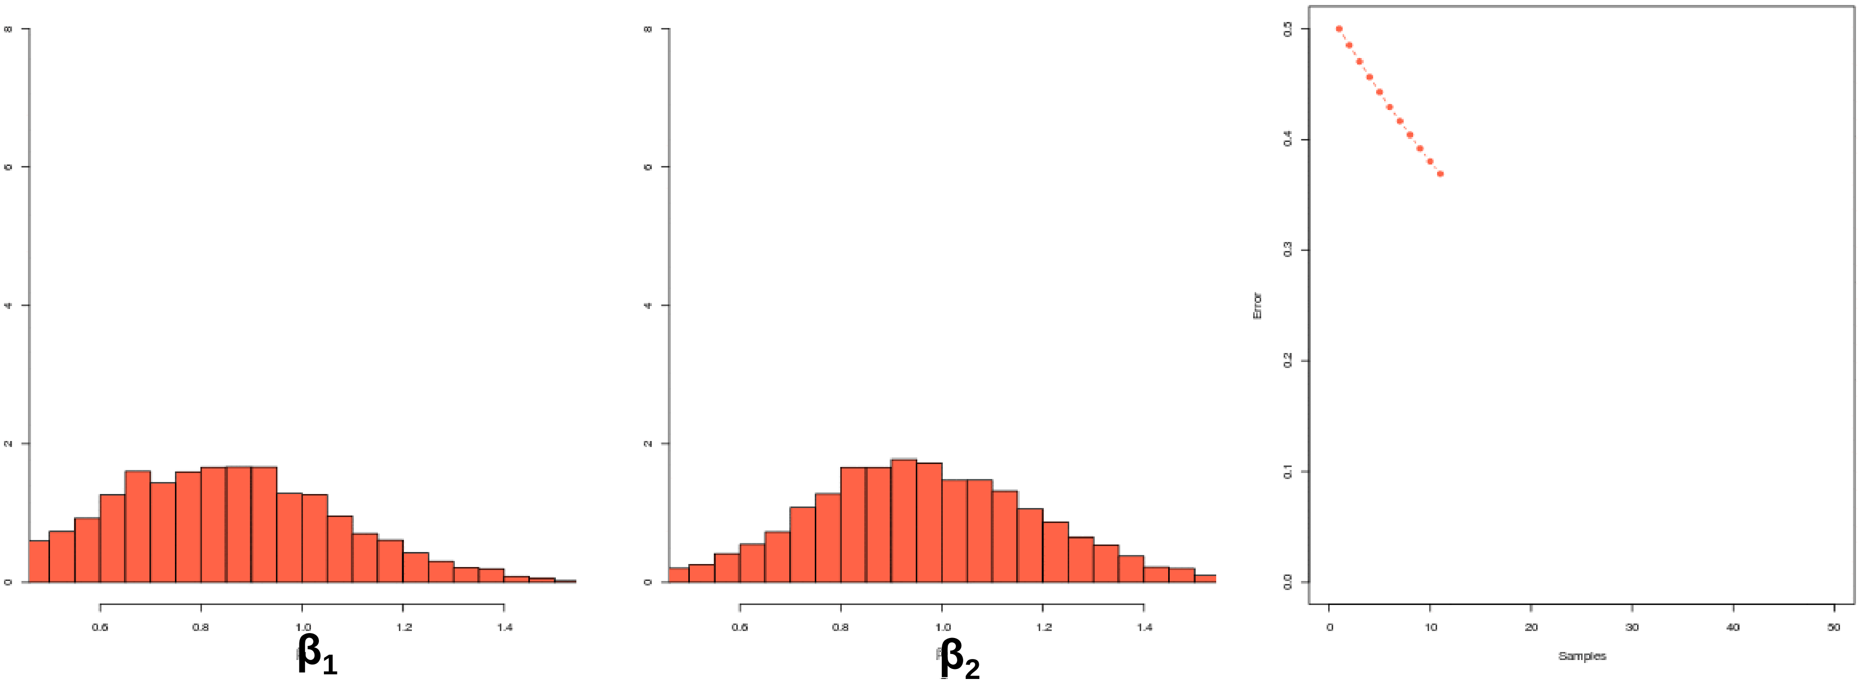

Active Logistic Regression

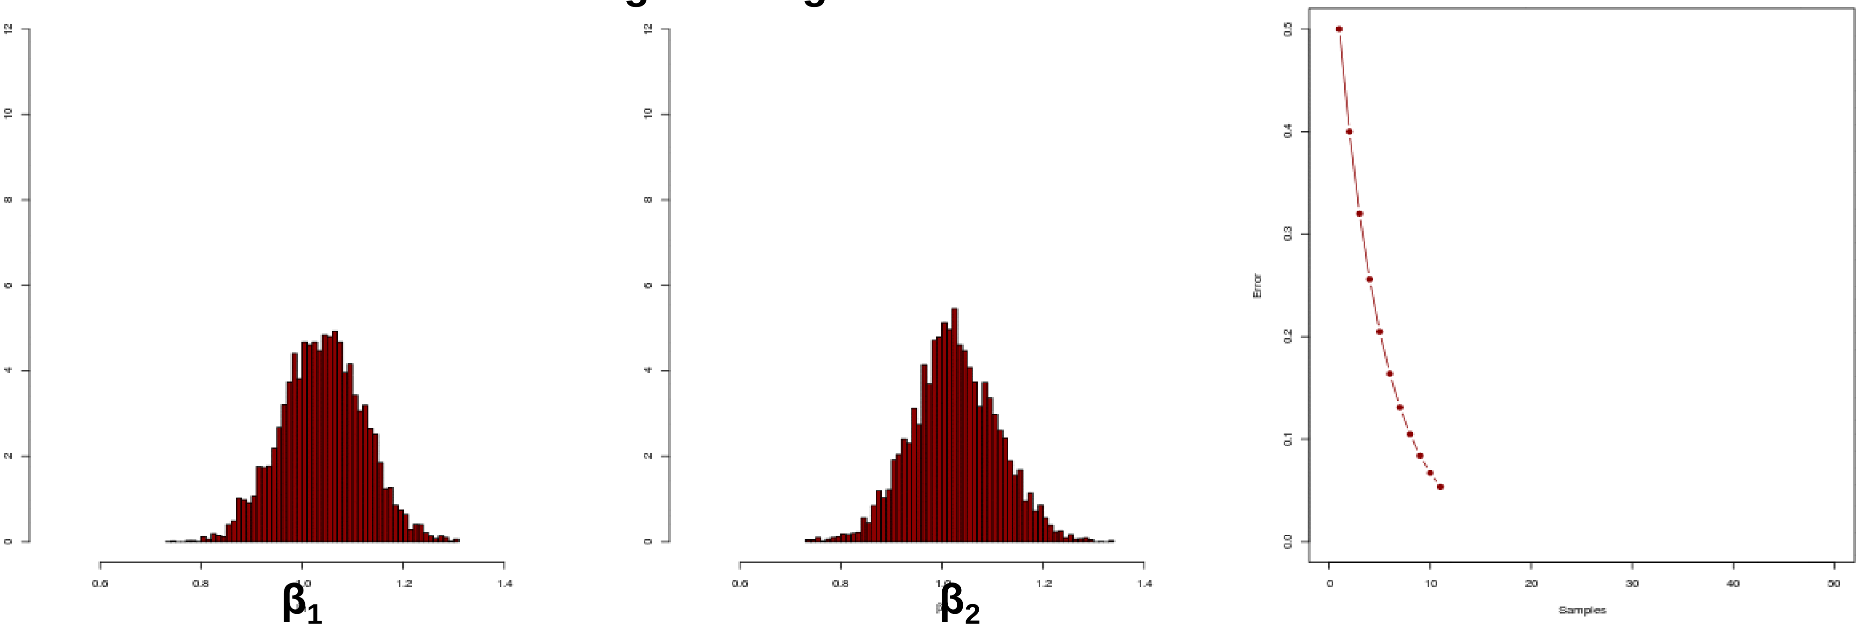

Alien Olympics: Participants

- 264 participants recruited via Mechanical Turk
- Participants were paid $0.50 for participation
- Bonus as reward for test performance (between $0 and $0.5)

Alien Olympics: Design

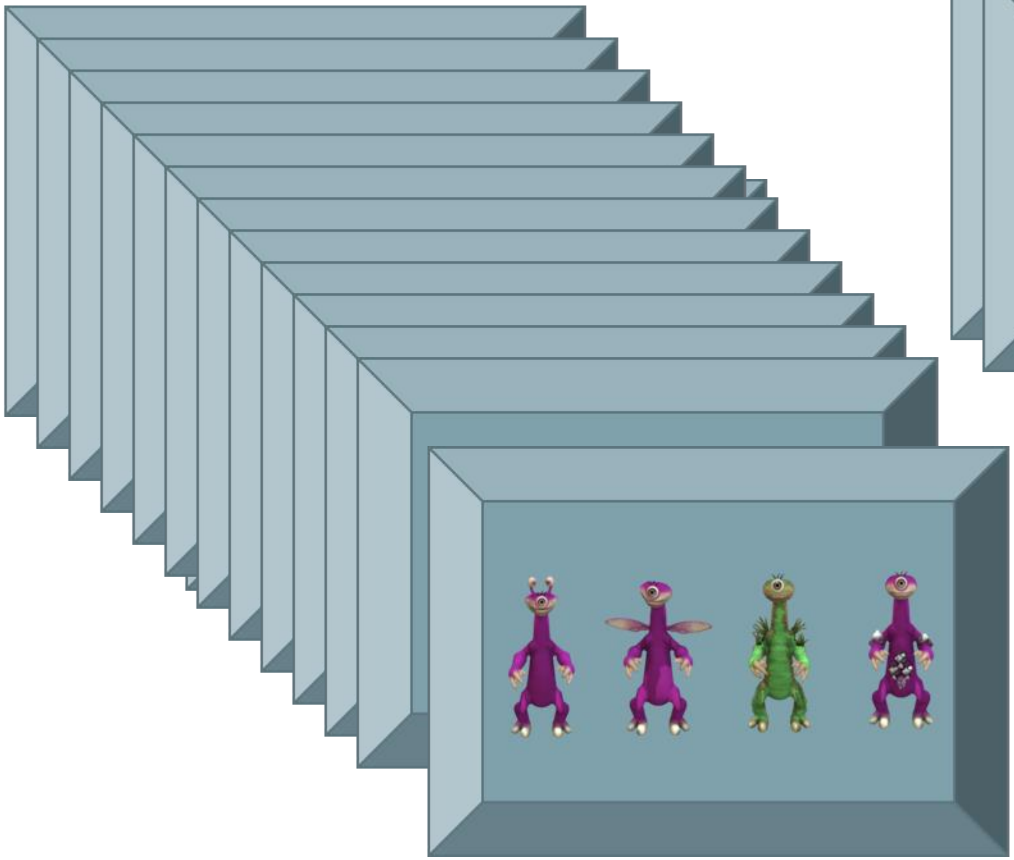

30 Learning Trials

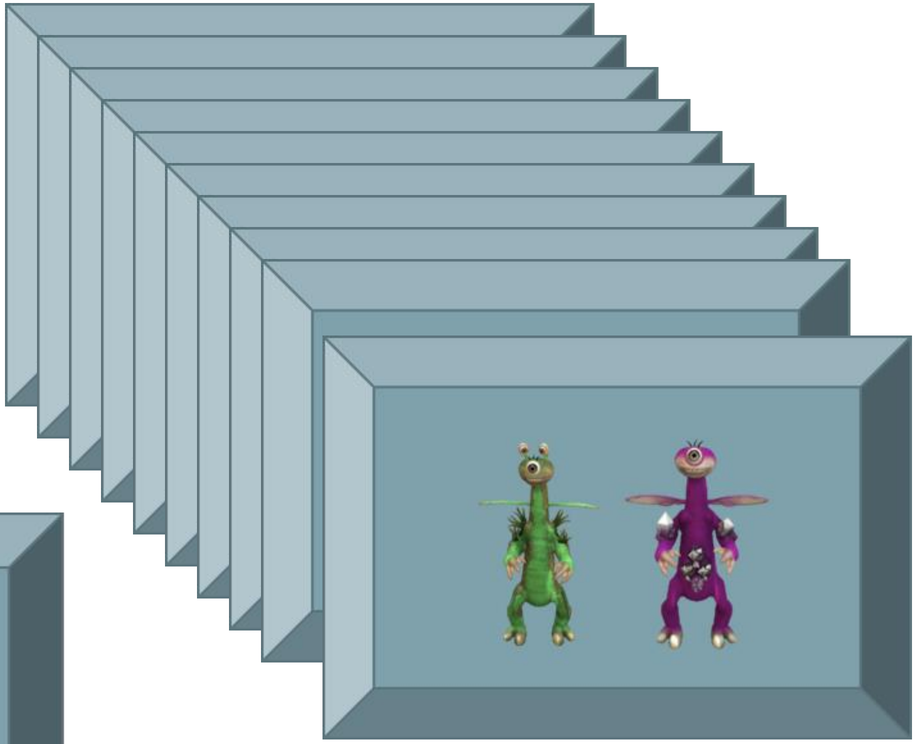

10 Test Trials

Alien Olympics: Design

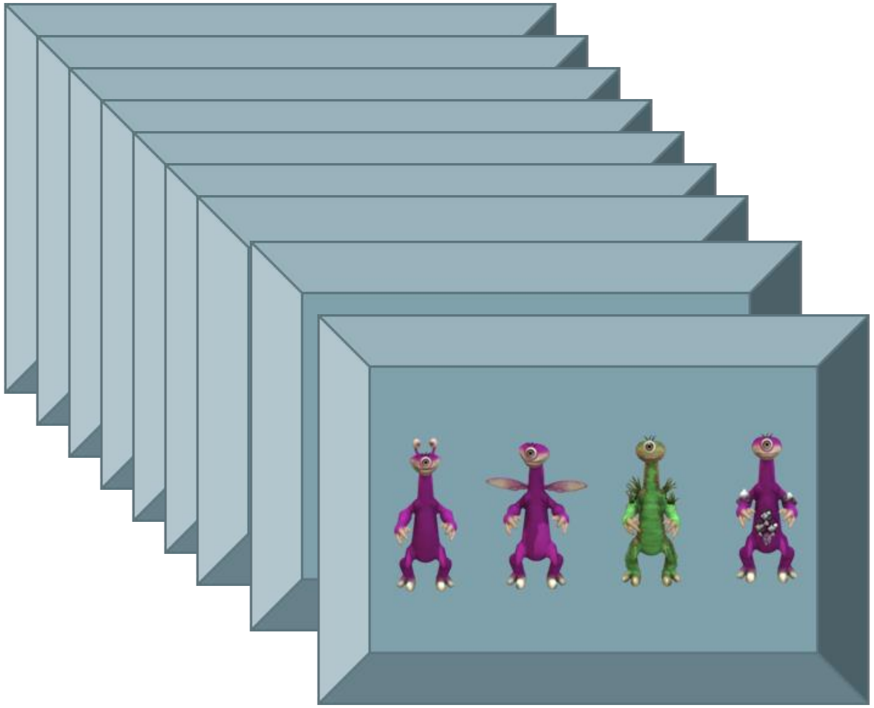

30 Learning Trials

INSTRUCTIONS

It will be your task to choose **2 out of the 4 Aliens** to compete with each other. You should choose the 2 Aliens such that you can learn as much as possible about the importance of their characteristics for their strength.

That means you should choose your Aliens wisely by selecting informative comparisons out of the 4 presented Aliens. Later in the experiment you will need this feature knowledge in order to correctly answer some questions.

Alien Olympics: Design

INSTRUCTIONS

You will be presented with 2 different Aliens representing the candidates for your Olympic Team. Having learned which characteristics make an Alien strong, it will be your task to always **select the Alien you consider to be stronger.**

Your overall payment will depend on the number of times you made the correct choice.

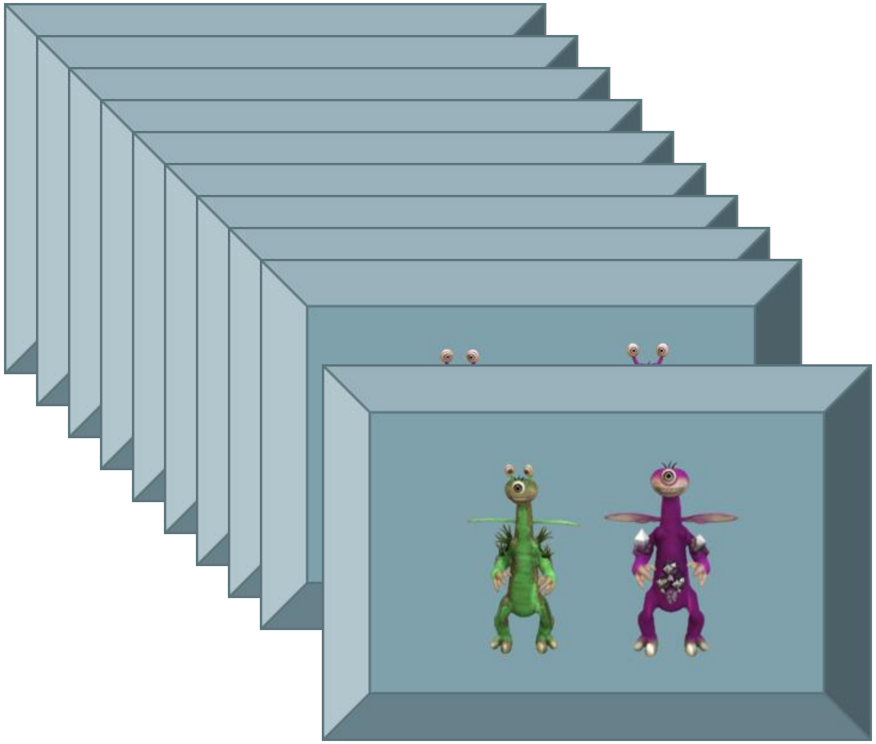

10 Test Trials

LEARNING TRIAL EXAMPLE

ALIEN OLYMPICS

Learning Stage: Please choose **2** Aliens to compete with each other.

Guidelines:

I. Below you see 4 different Aliens. The Aliens are described by 4 different characteristics. These characteristics influence how strong they are.

II. It is your task to learn how the different characteristics influence an Alien's strength by always choosing two Aliens to compete against each other. You therefore have to make this choice wisely/informatively, i.e. by creating comparisons with an outcome that tells you something about the effect of the different characteristics.

III. Once you click on an Alien it will be marked by a black rectangle.

IV. After you have chosen two Aliens, please press the "Compete"-Button and you will see which of the two Aliens has won the competition. You cannot choose more or less than 2 Aliens for a competition.

V. Remember that just as in any sport, sometimes a weaker Alien can win against a stronger. This can happen. Nevertheless you should try to find out which factors influence whether an Alien lost or won and how strong these effects are.

VI. After each competition, you have to click "Next trial" to continue.

Hide guidelines

Number of trials left: 28

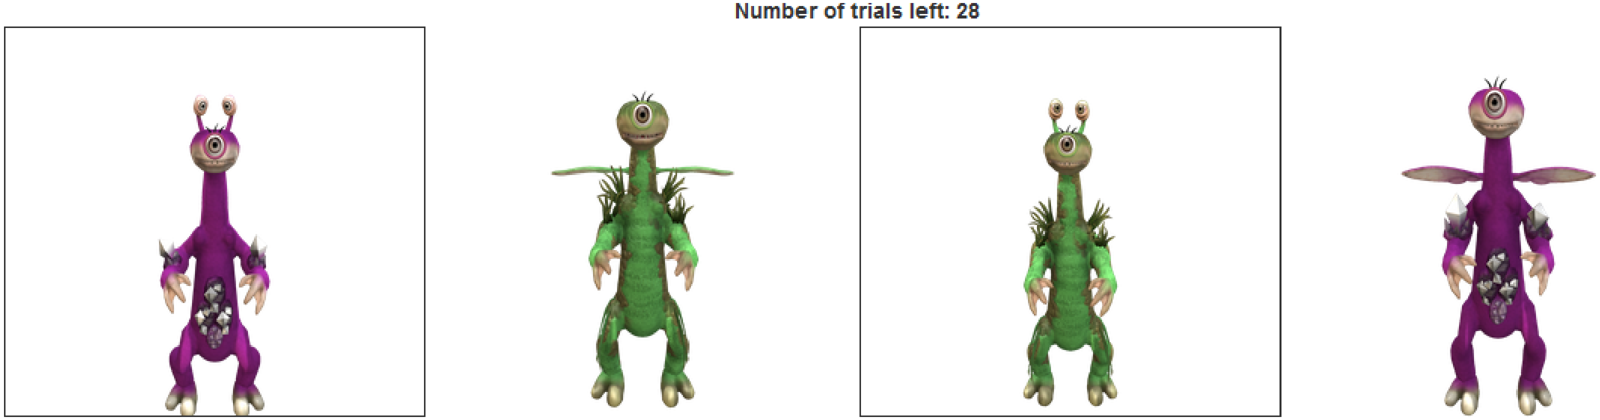

Compete

LEARNING TRIAL FEEDBACK

ALIEN OLYMPICS

Learning Stage: Please choose **2** Aliens to compete with each other.

Guidelines:

I. Below you see 4 different Aliens. The Aliens are described by 4 different characteristics. These characteristics influence how strong they are.

II. It is your task to learn how the different characteristics influence an Alien's strength by always choosing two Aliens to compete against each other. You therefore have to make this choice wisely/informatively, i.e. by creating comparisons with an outcome that tells you something about the effect of the different characteristics.

III. Once you click on an Alien it will be marked by a black rectangle.

IV. After you have chosen two Aliens, please press the "Compete"-Button and you will see which of the two Aliens has won the competition. You cannot choose more or less than 2 Aliens for a competition.

V. Remember that just as in any sport, sometimes a weaker Alien can win against a stronger. This can happen. Nevertheless you should try to find out which factors influence whether an Alien lost or won and how strong these effects are.

VI. After each competition, you have to click "Next trial" to continue.

Hide guidelines

Number of trials left: 28

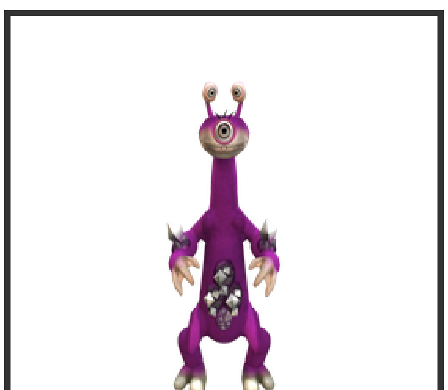


Alien 1 has won the comparison.

Next trial

TEST TRIAL EXAMPLE

ALIEN OLYMPICS

Assessment Stage: Please decide which of the two Aliens you would like for your Olympics Team.

Guidelines:

I. Below you see 2 different Aliens. The Aliens are described by 4 different characteristics as before. These characteristics influence how strong they are.

II. It is now your task to choose the Alien you consider to be the stronger of the two.

III. Once you click on an Alien it will be marked by a black rectangle.

IV. After you have chosen the Alien, please press the "Select"-Button and this Alien will become a member of your team. You can only chose one Alien at a time.

V. Remember that just as in any sport, sometimes a weaker Alien can win against a stronger. This can happen. It is your task to choose the Alien (out of the two candidates) you consider to be stronger.

VI. You have 12 choices in total and your final reward will depend on the quality of your choices.

VII. After you have chosen an Alien, you have to click on the "Next trial"-Button to continue with the next trial.

Hide guidelines

Number of trials left: 9

Choose

Underlying weights of 4 features

Participants were randomly assigned to **1 of 4 compensatoriness conditions**:

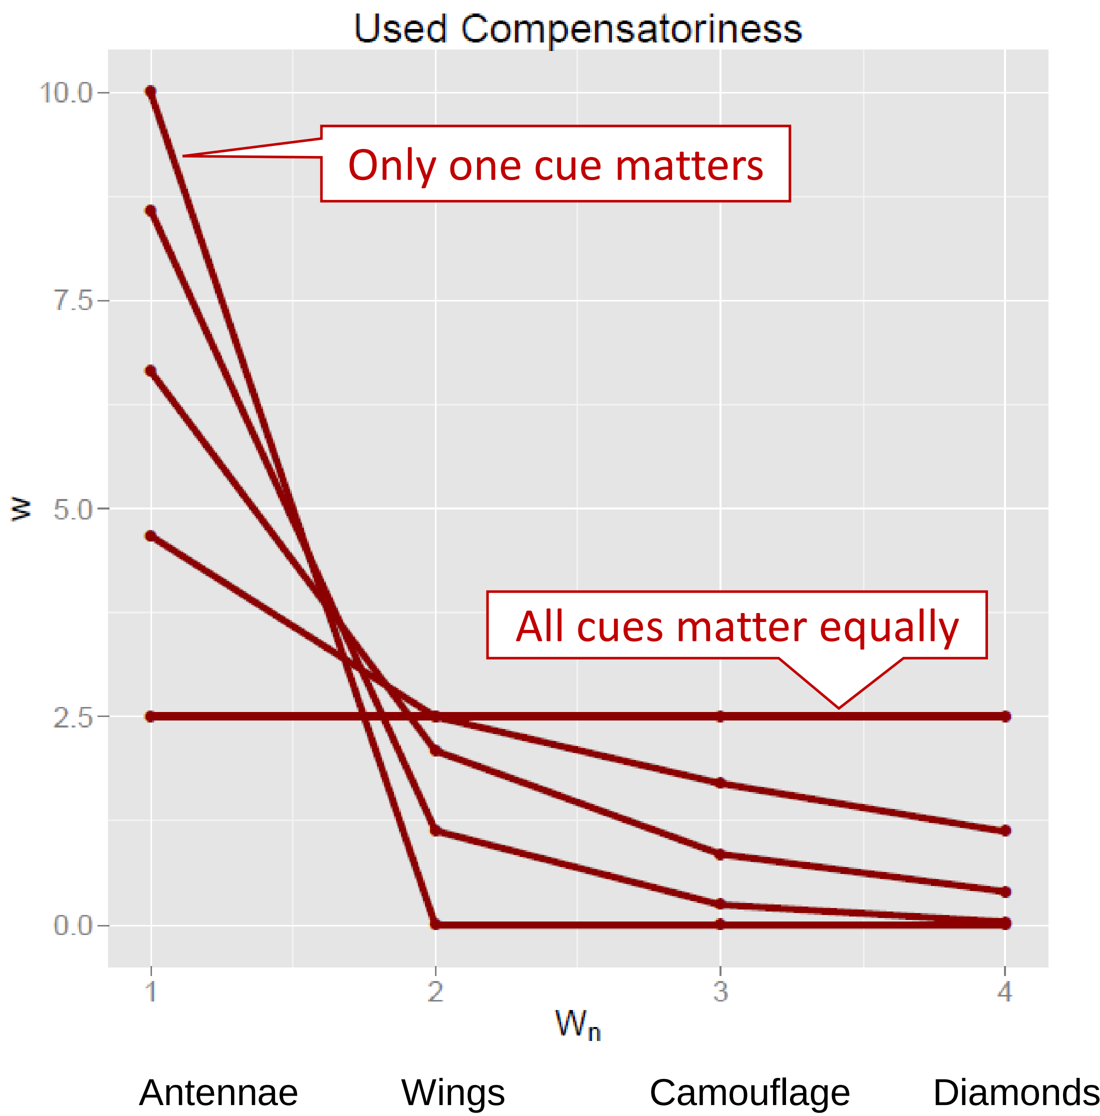

Underlying weights of 4 features

Hypotheses: Are people adaptive to the underlying weight structure in the task? It is possible, that:

- TTB is better throughout
- Logistic is better throughout
- Depending on compensatoriness, people alternate between TTB and Logistic.

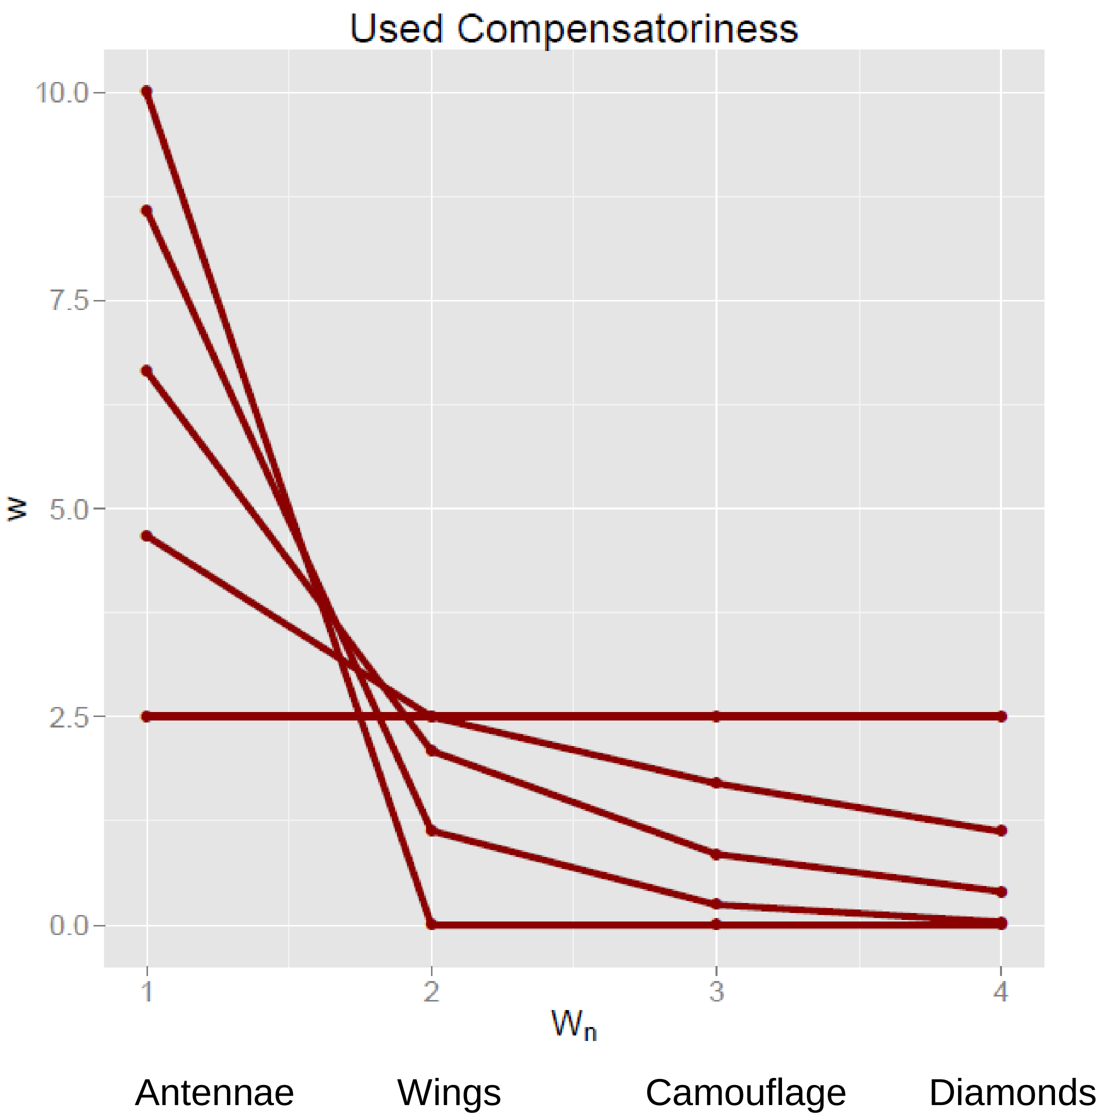

Alien Olympics: Results from Test phase

→ As the environmental structure gets more non-compensatory, performance drops.

→ makes draws more likely and informative comparisons less likely

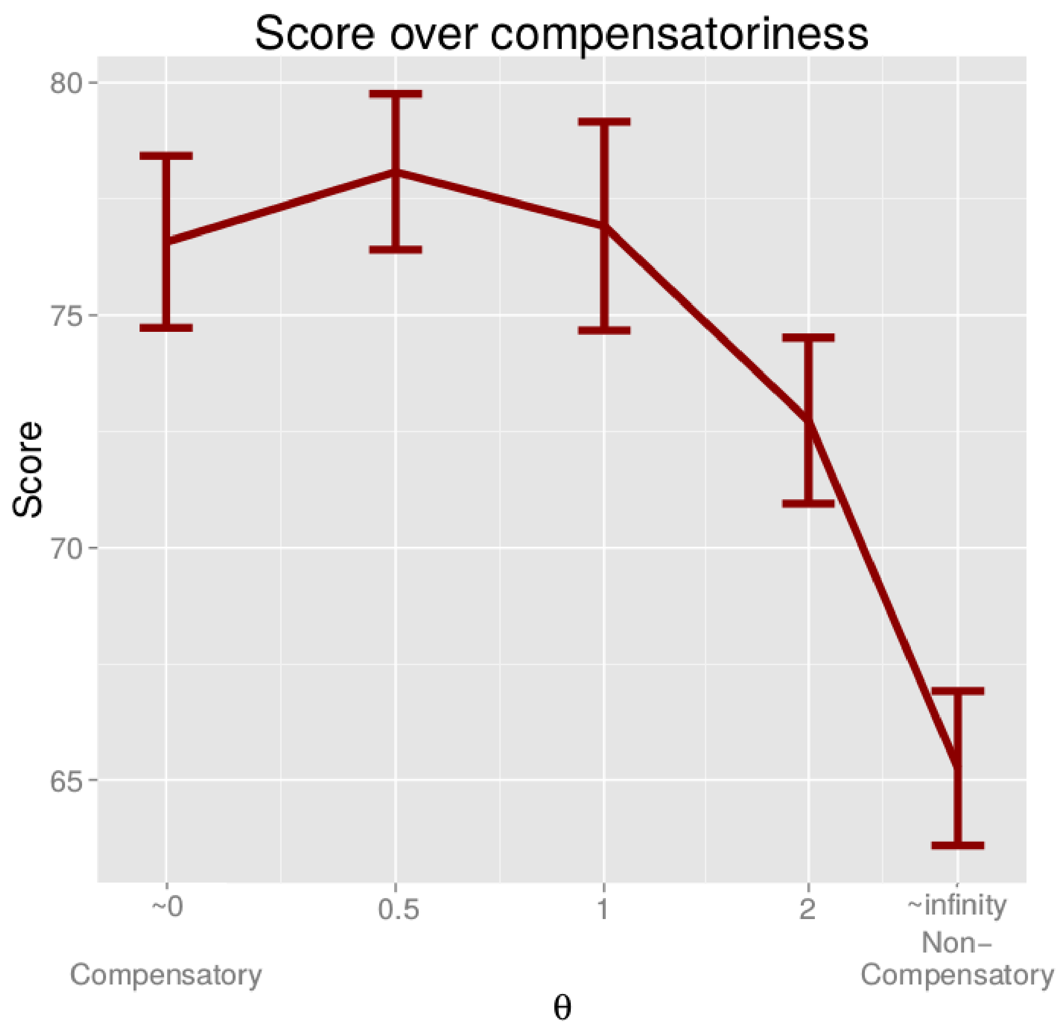

Results: Aggregated frequency of queries

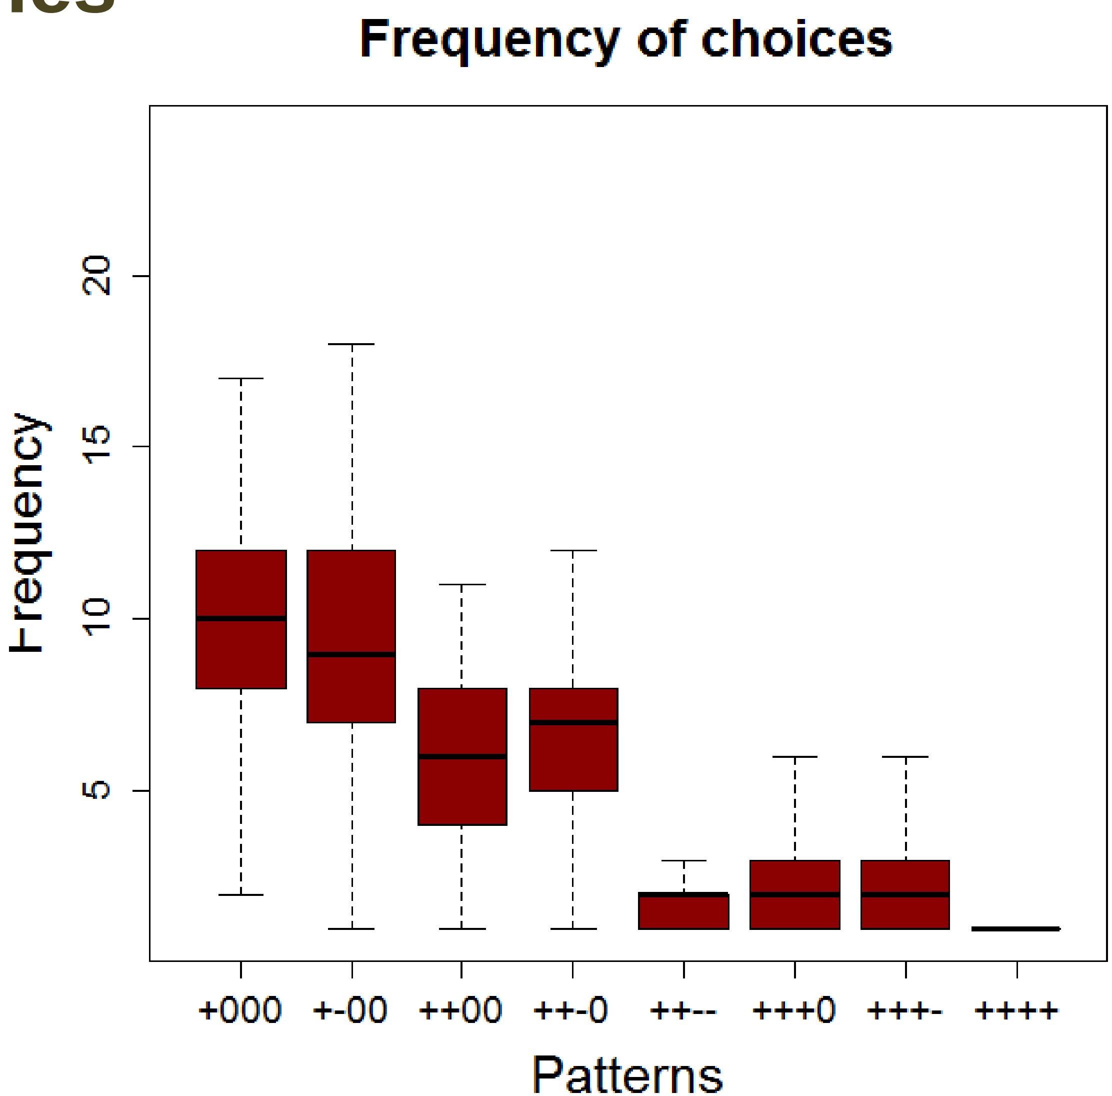

+	Has the feature
-	Lacking the feature
0	Draw

Results: Active Learning

- People rarely chose comparisons where it was unclear what feature was responsible, e.g., +++0 or ++++.

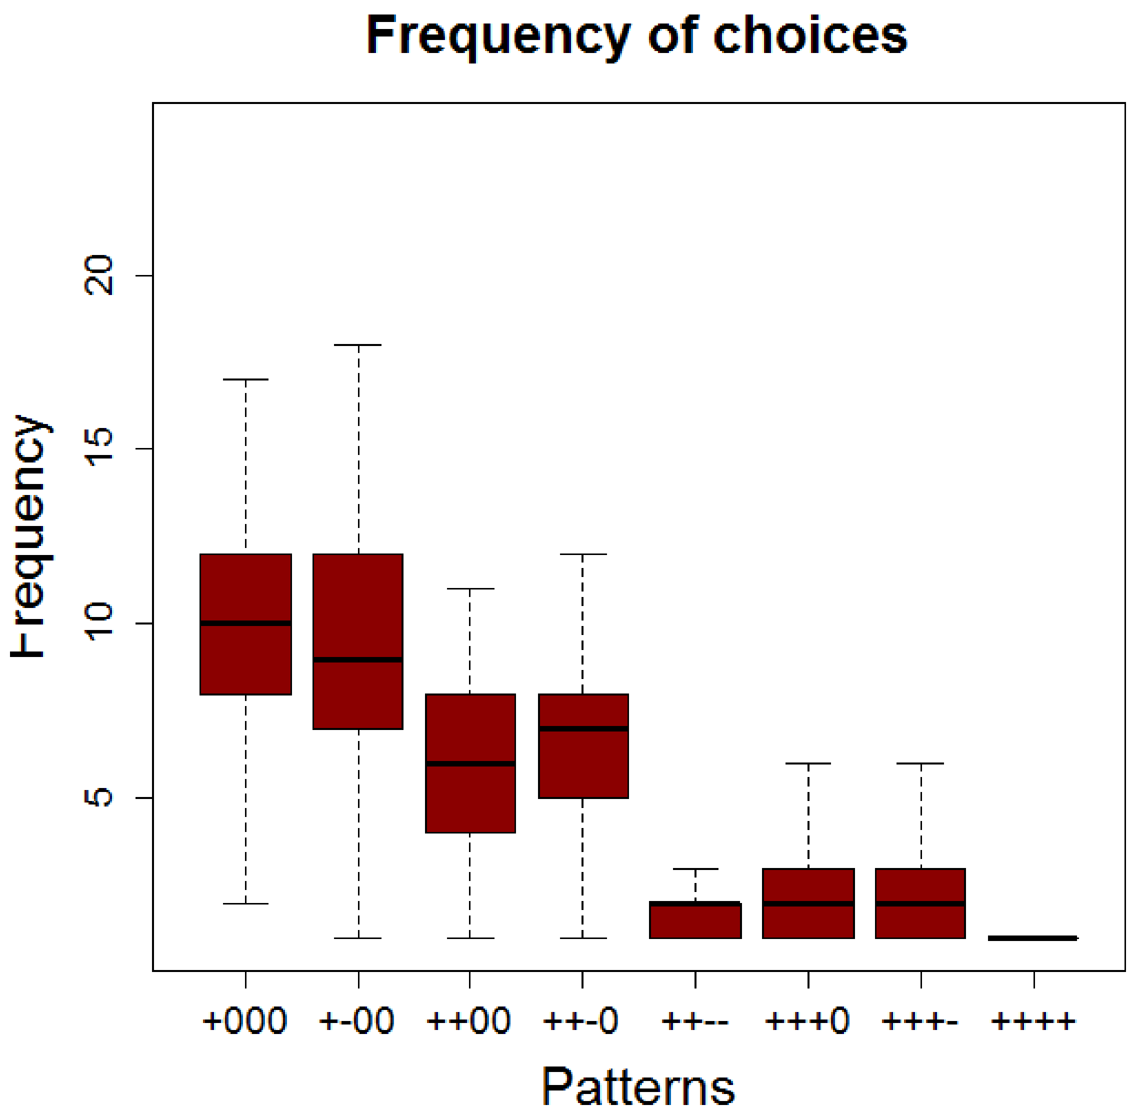

Results: Active Learning

- Participants performed **simpler, more controlled** queries than suggested by optimal algorithms:

+-00 = tests relative effect of a feature in comparison to another
→ **rank-order query**

+000 = assesses whether feature improves outcome
→ **Importance query**

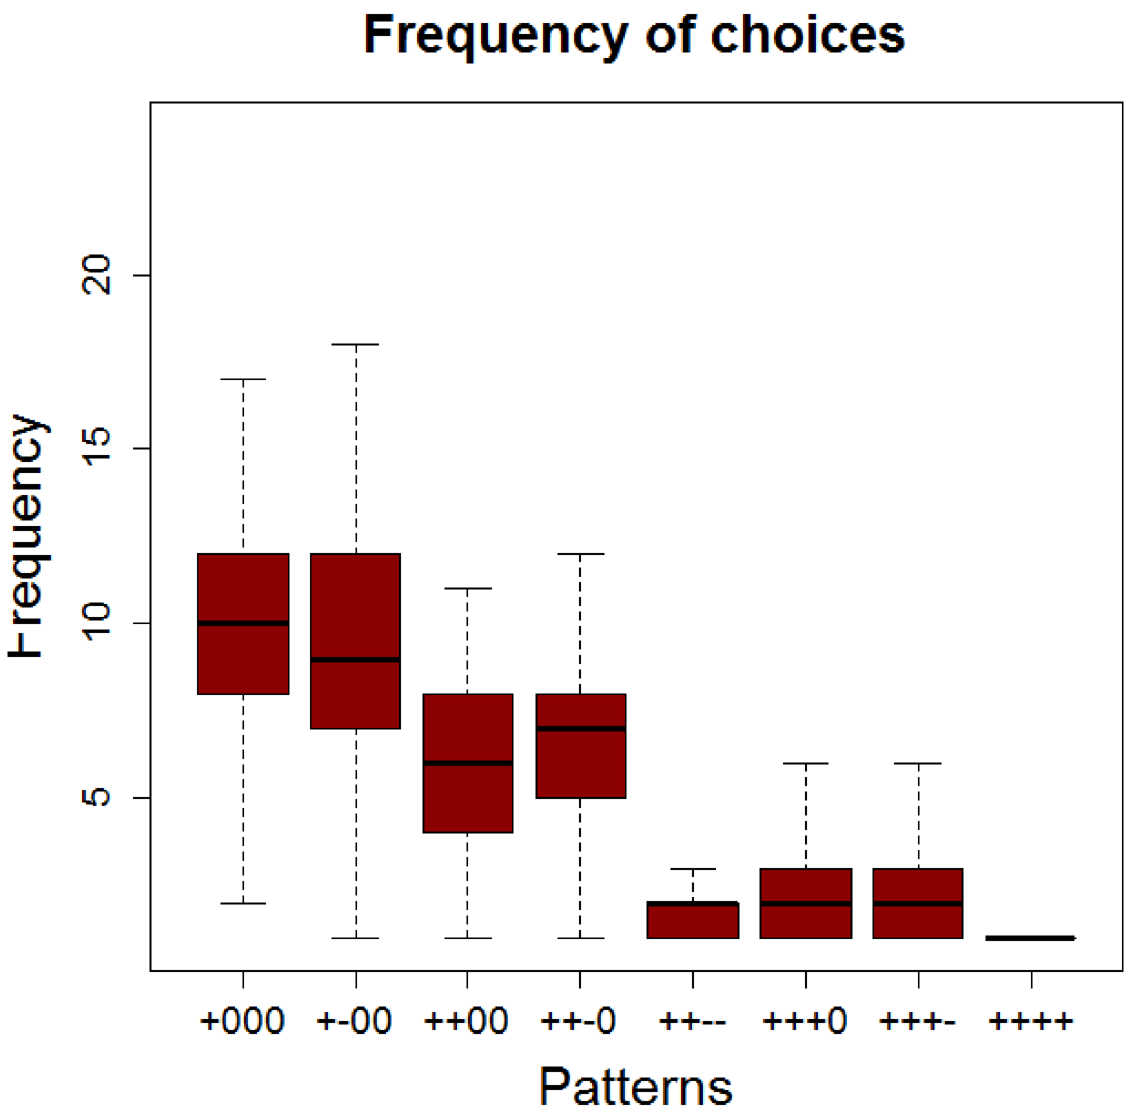

Results: Behavioral Predictions (Test)

→ **This is where most psychology experiments stop. Predicting test data (cross-validation).**

→ Hard to distinguish between models based on model fitting in test phase alone. -> easily visible by how much the error bars overlap!

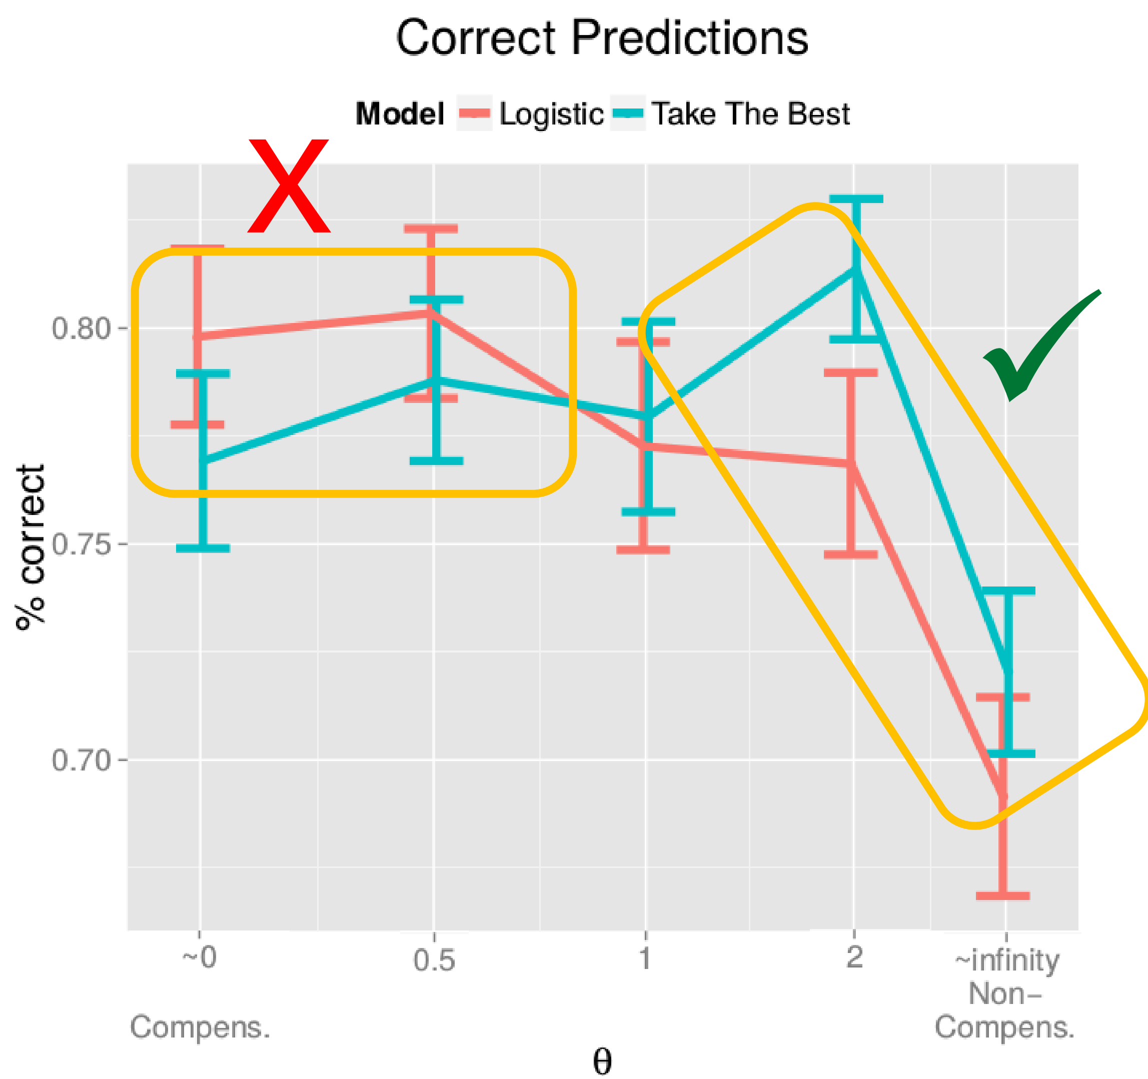

Results: Behavioral Predictions (Test)

✓ **Logistic Regression > Take-The-Best**

Behavioural model fits indicate that Logistic Regression is slightly better at predicting people's test choices than TTB.

→ Evidence alone is not very strong though.

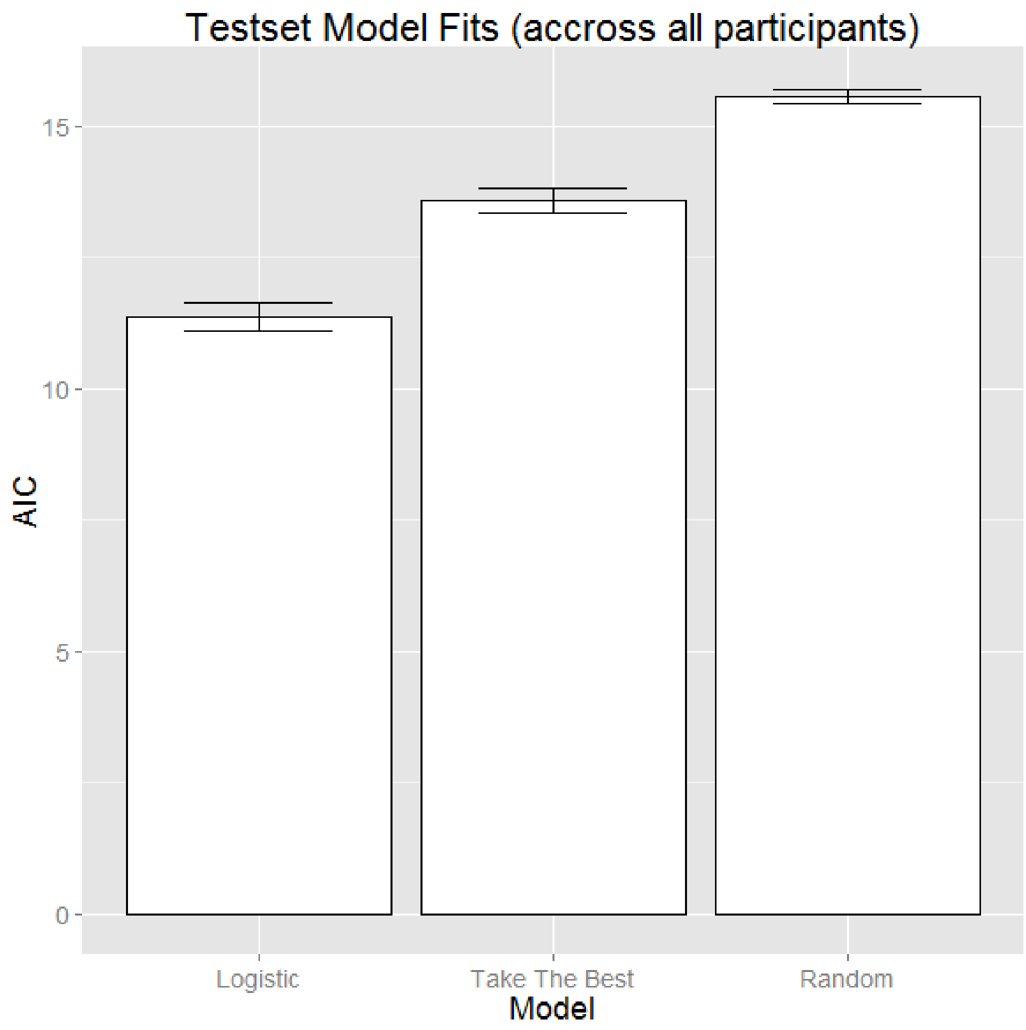

Results: Active Learning

✓ **Logistic Regression > Take-The-Best**

Combining AIC evidence from both active learning results and passive model fitting gives best results

→ Active learning AIC's seem to discriminate better between cognitive models (however different data!)

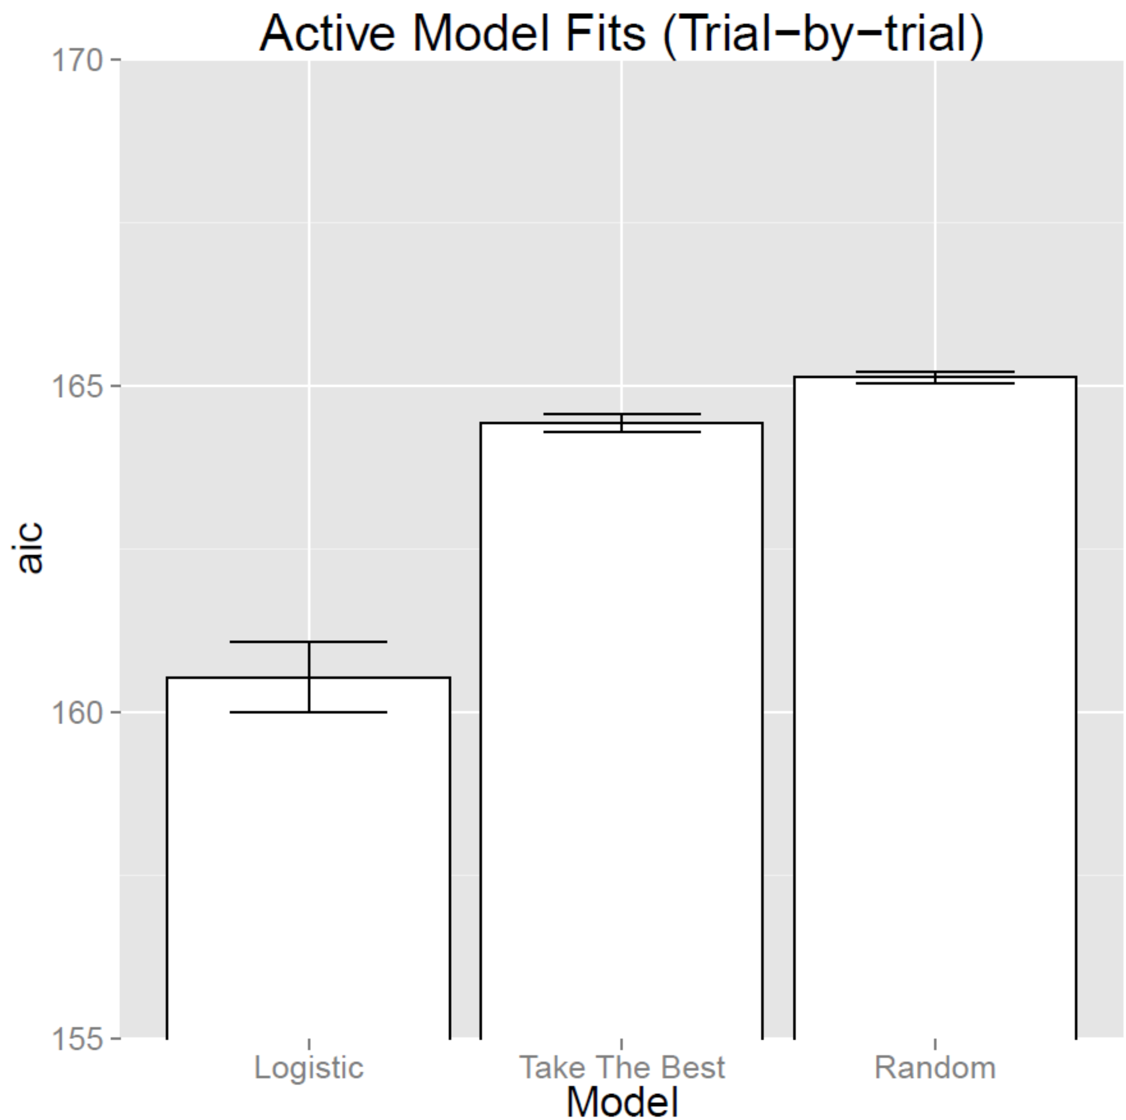

Results: Active Learning

- ✓ **Logistic Regression> Take-The-Best**
- ✓ The active Logistic algorithm captured peoples‘ queries best in more compensatory environments.
- ? In compensatory environments, certain queries were preferred.

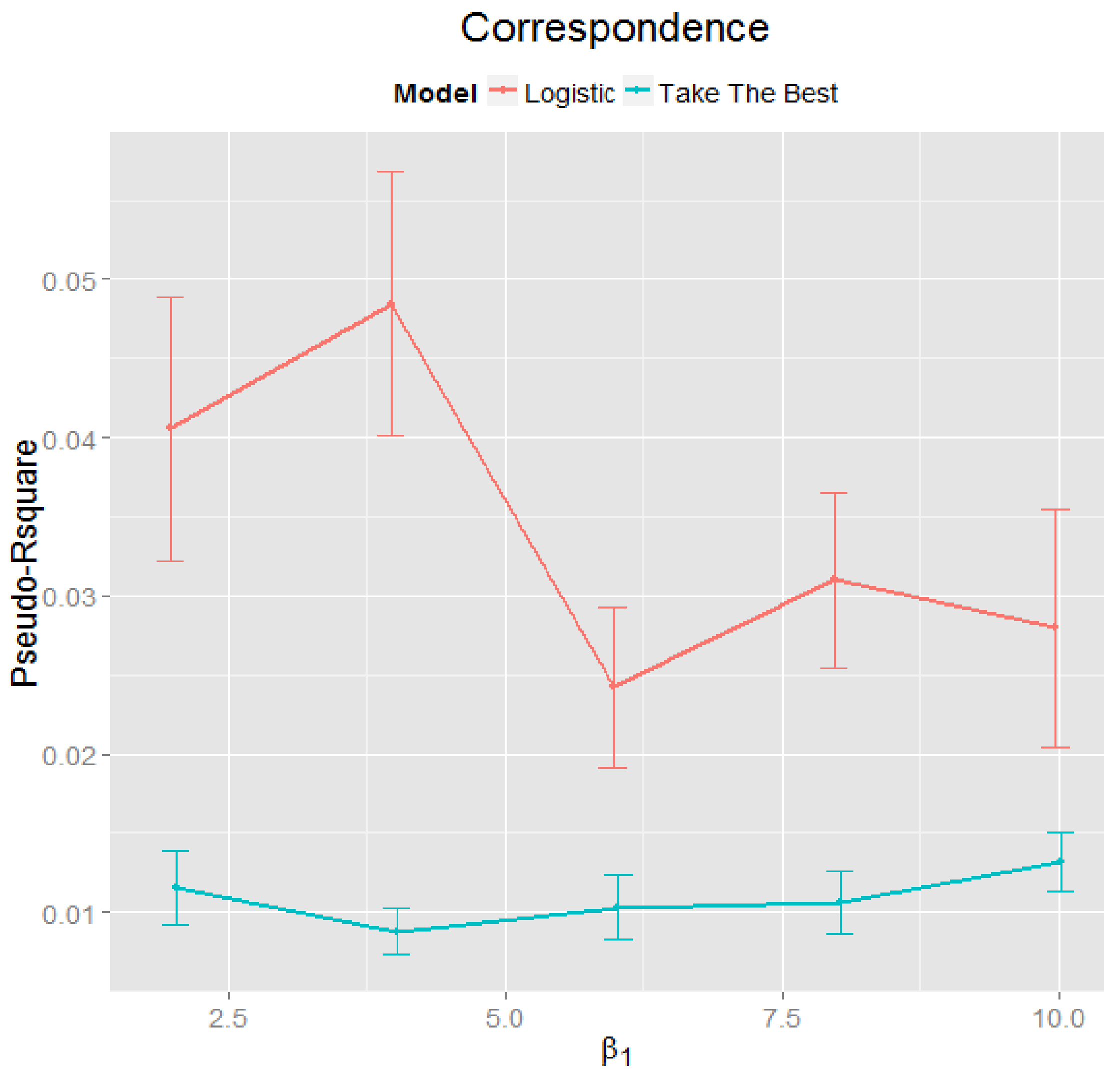

Compensatory **Non-Compensatory**

Results: Behavioral Predictions (Test)

→ **This is where most psychology experiments stop. Predicting test data (cross-validation).**

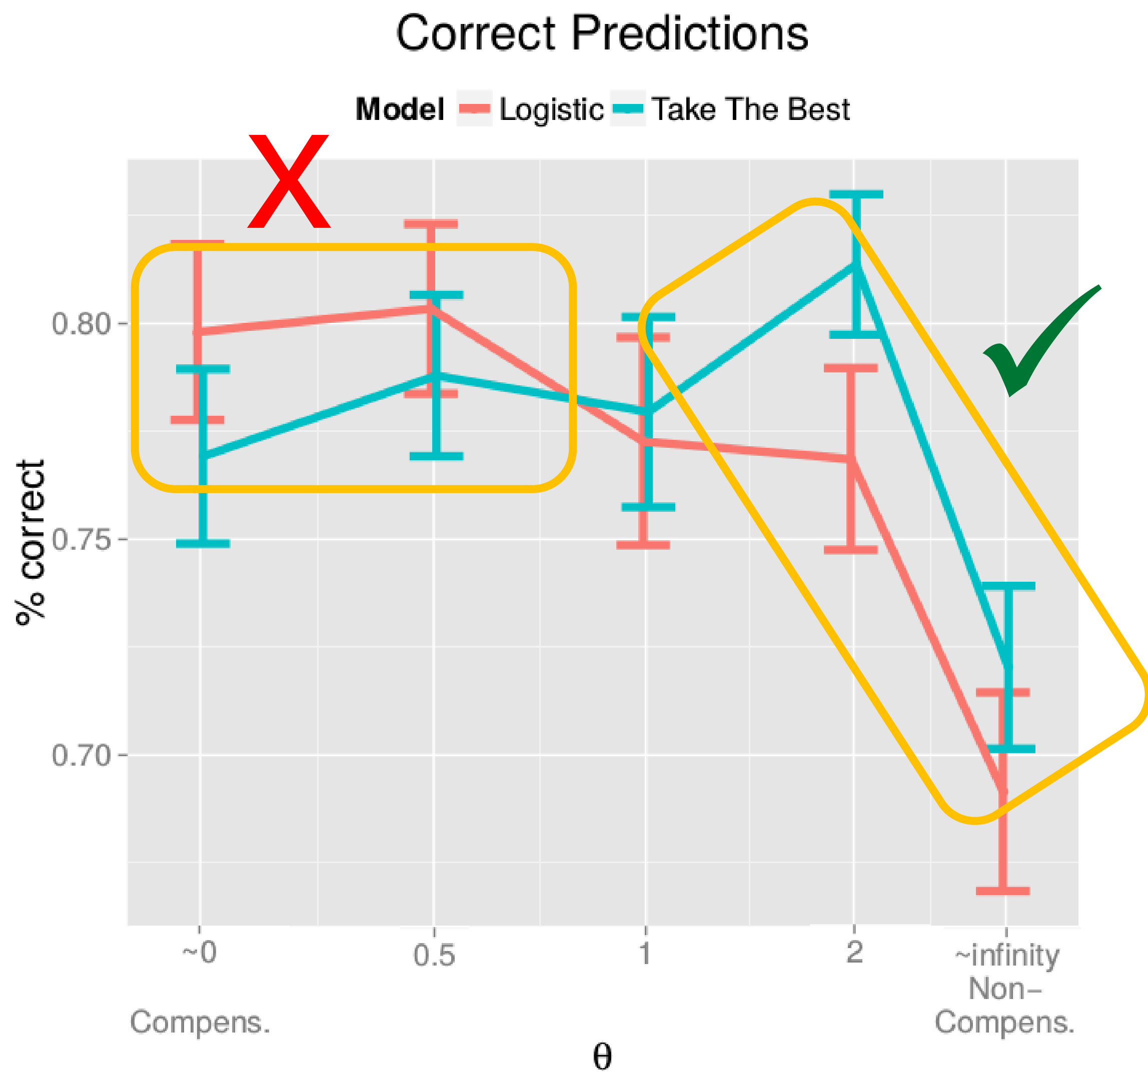

→ Hard to distinguish between models based on model fitting in test phase alone. -> easily visible by how much the error bars overlap!

Results: Active vs. Passive Learning

✓ **Logistic Regression > Take-The-Best**

→ Active model testing is more insightful than passive model testing.

→ Active model comparison allows to discriminate among decision models better.

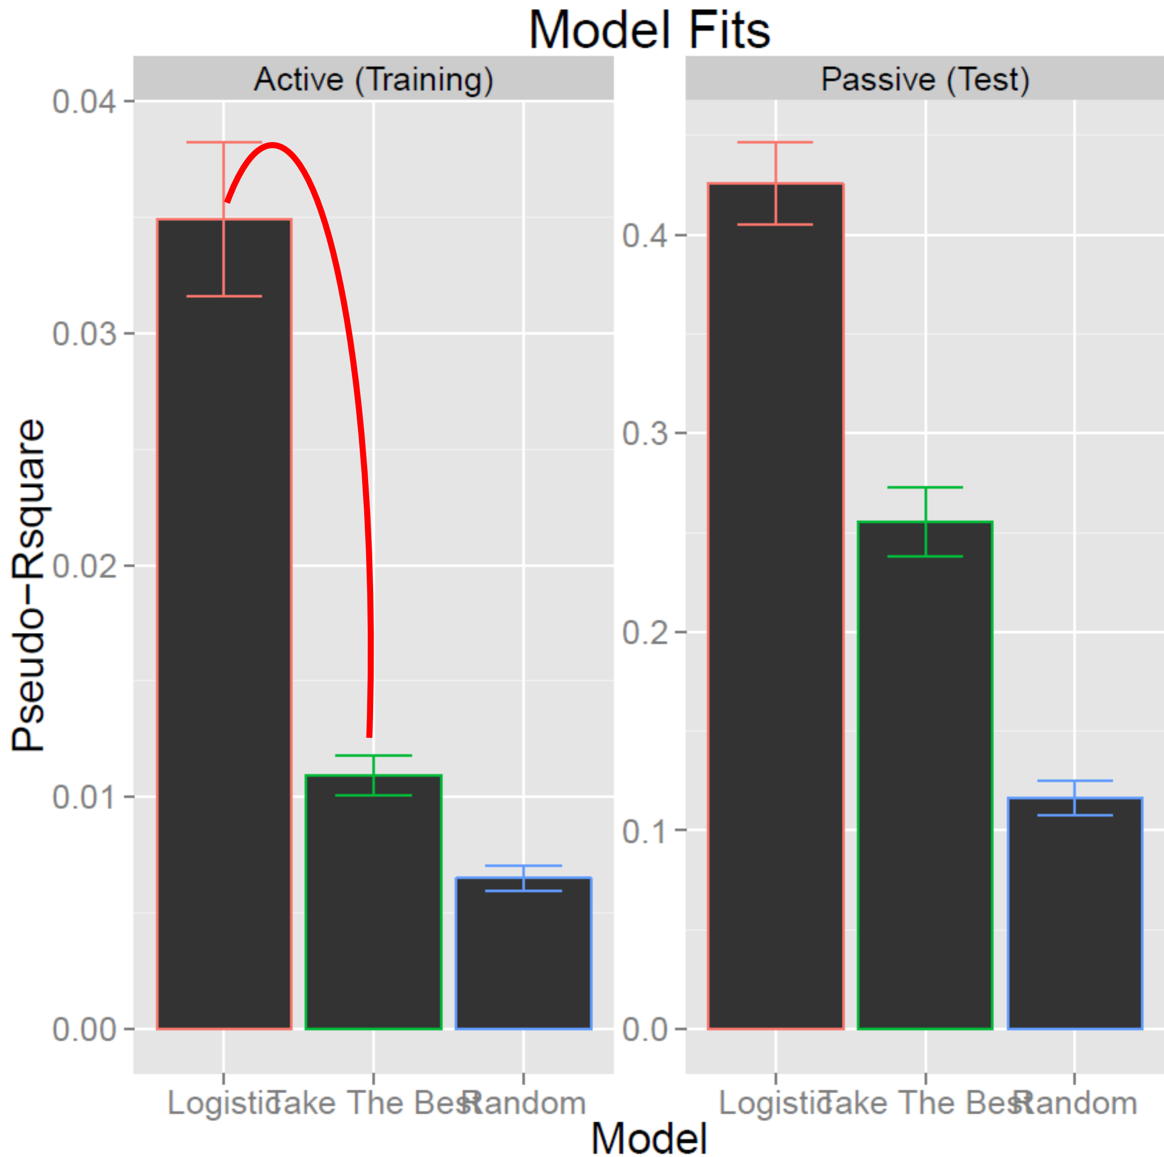

Results: Frequency of choosing query that maximally reduces uncertainty

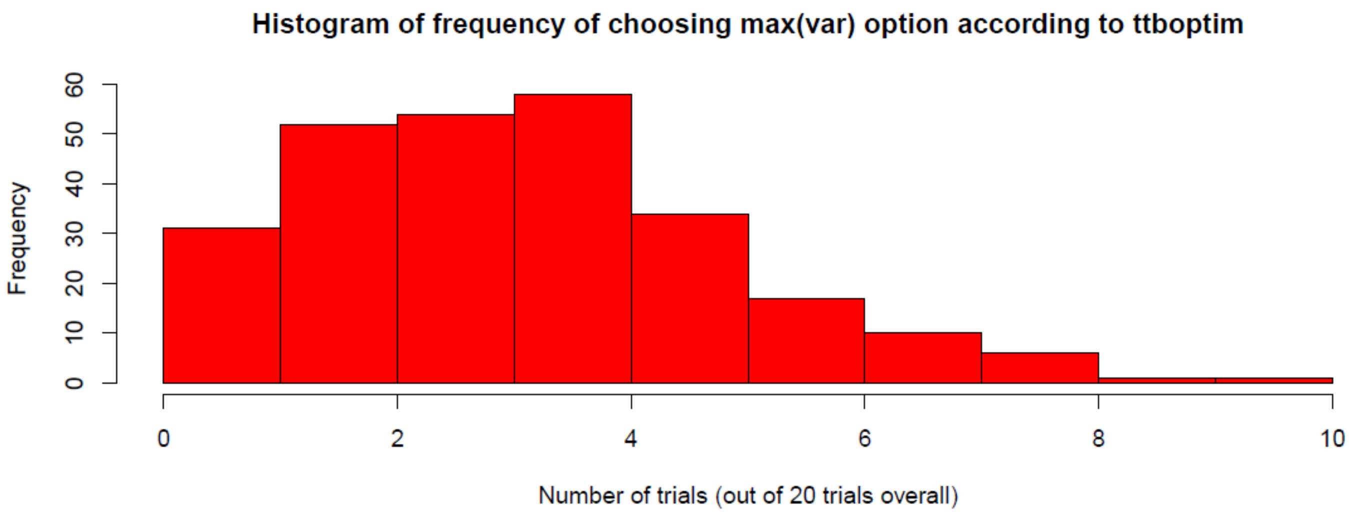

TTB active algorithm

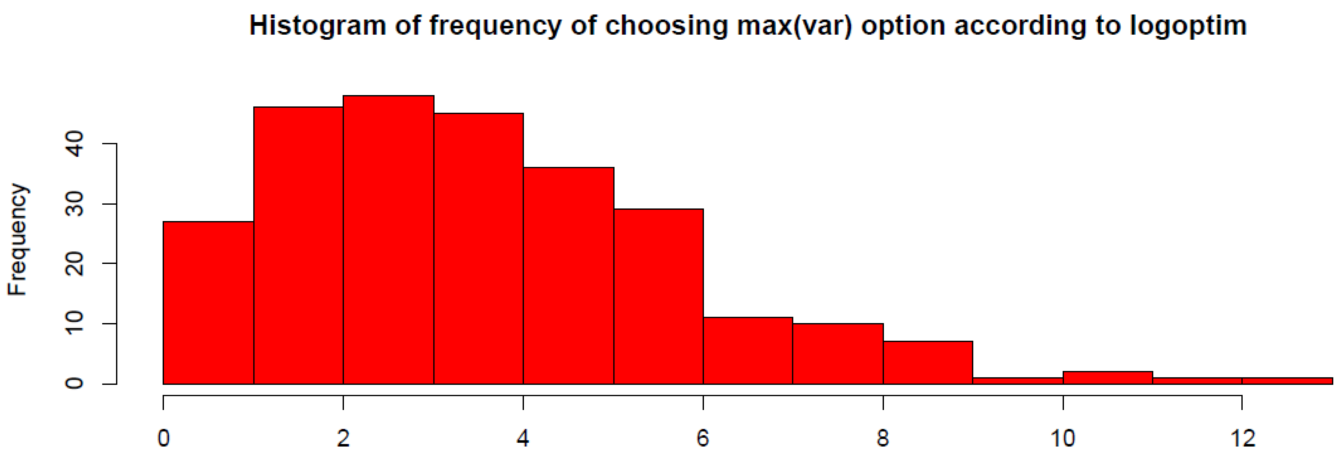

Logistic active algorithm

Interim Summary

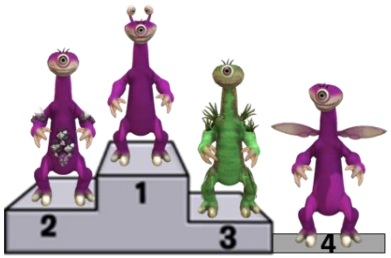

- Active Logistic Regression was better than Active TTB at capturing peoples active learning behaviour.
- The fact that people preferred much simpler queries could reflect a preference to perform confirmatory tests.
- In compensatory environments, people preferred certain queries over uncertain queries: learning where they already know more/feel more certain (**confirmatory testing or a different strategy**? Markant & Gureckis, 2012)

Take-Home Messages

- Active Learning Question: Do people learn with respect to **cue weights** (Regression) or **cue orders** (TTB)?

Answer from 1st Experiment: **cue weights > cue orders**

Take-Home Messages

I hope to have convinced you that, ….

Active Learning is a powerful tool to discriminate among decision making models, with the potential to resolve long-standing debates in psychology.

→ Active learning provides a window on representation/inference

Take-Home Messages

People generally preferred **simpler**, e.g., importance-based **queries** (positive testing) in comparison to the optimal active algorithms.

This could have to do with our limited resources.

→ Current research on creating plausible, heuristic models of human active learning (Bramley et al., 2015; Markant & Gureckis, 2012) that allow ‘forgetful’ models.

Interim summary

- People might be performing a combination of information gain and a *confirmatory test heuristic (i.e., positive testing)*
- As suggested by the ‘*positive test heuristic*’ in Coenen et al. (2015) and Bramley et al.’s (in press) Neurath’s ship model of causal learning

Coenen et al (2015)

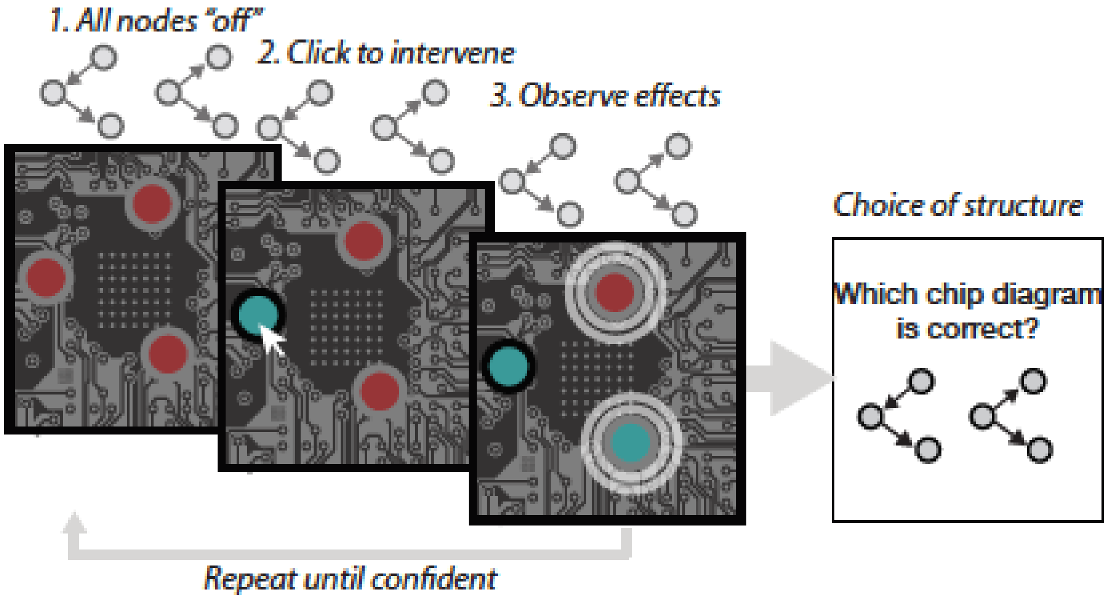

- Find in 3 experiments and sophisticated analyses participants intervention choices not well described by information gain alone
- Propose a mixture of *information gain* and "positive test heuristic" (i.e. tendency to turn on the root-cause node regardless of whether this is informative)
- Level of heuristic use increased under time pressure

Neurath's ship (Bramley et al, 2015, in press)

- Explored interventional causal learning in complex situations (up to 4-variables and 543 possible true models)

- Participants judgments **better described by a piecemeal local update process than global Bayesian updating** – i.e. changing one connection at a time trying to improve the fit with the latest evidence

- This captures apparent *forgetting* and *conservatism*
- In parallel – interventions better described driven by several varieties of "local focus" rather than a singe global focus
  - *Confirmation* – Focusing on confirming/disconfirming the current hypothesis
  - *Effects* – Focusing on the effects of a single component
  - *Edges* – Focusing on a single connection

“Local” information (Gureckis et al, 2015)

- Version of TV aerial task where participants must distinguish between 3 possibilities
- Find that people select tests that resolve uncertainty between 2 hypotheses at a time
- Propose that active learning is local in the sense that generally attempt to distinguish only a small set of hypotheses at a time rather than the whole space

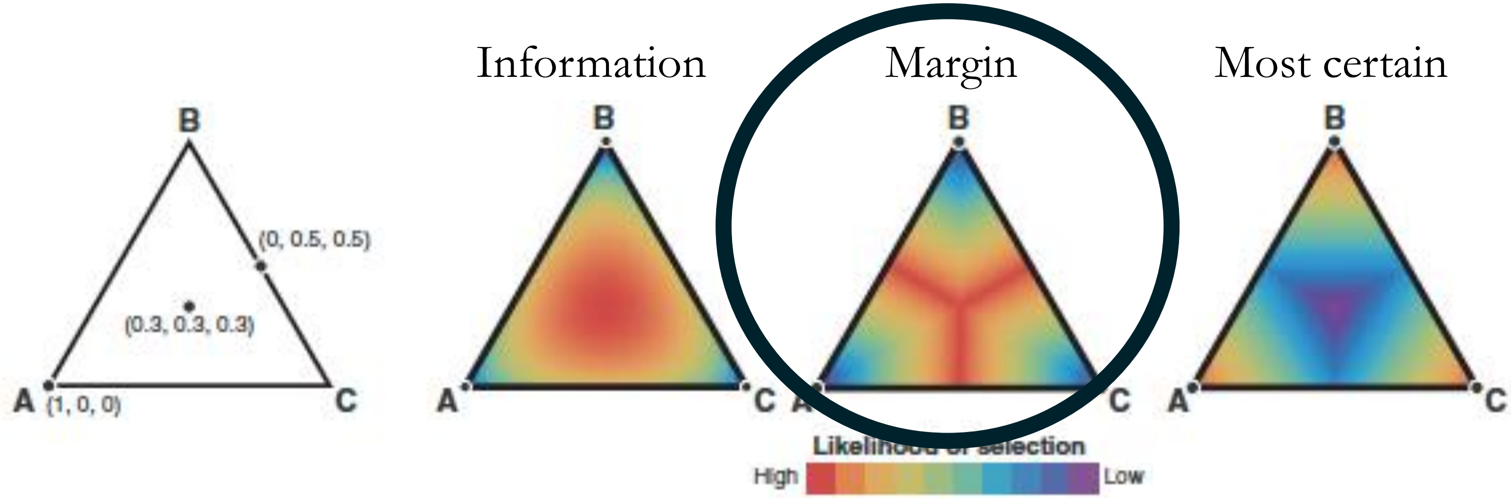

Follow-up studies

<u>1. Alien experiment with forced choice between two alien pairs, optimally designed:</u> Alien pairs are optimised to distinguish between the two models as much as possible.

- Generate comparisons that maximally reduce uncertainty with respect to the two active models: Entropy(TTB) – Entropy(Logistic)
- Comparisons are presented to participants with probability of discriminability among decision models
- Downside: less active and more constrained by experimenter

vs.

or

vs.

Follow-up studies

<u>2. Alien experiment with forced choice between two alien pairs:</u>
Participants always see one alien on the top and two aliens on the bottom. The aliens are generated randomly. They have to choose the pair they want to compare.

compare with:

Follow-up studies

<u>3. Alien experiment where participants can design their own aliens:</u>
Participants get to design the aliens with a set of possible features:

+

-Antennas
-Diamonds
-Wings
-Camouflage

?

→ Participants design <u>both aliens</u> they want to let fight.

→ Participants choose their own alien competitions.

→ This experiment is much more **active** as participants choose everything from stimulus design to comparison quries

Experiment 2

2. Alien experiment with forced choice between two alien pairs:
Participants always see one alien on the top and two aliens on the bottom. The aliens are generated randomly. They have to choose the pair they want to compare.

- Less degrees of freedom
- Less comparisons to choose from
- Only 2 possible comparisons instead of 6 (Exp. 1)

compare with:

Experiment 2

We expected to be able to better discriminate among models than in Experiment 1. The baseline probability of people’s choices matching the model’s predictions is higher. (1/2 = 50%)

- Less degrees of freedom
- Less comparisons to choose from
- Only 2 possible comparisons instead of 6 (Exp. 1)

compare with:

Thank you!

UK PhD Centre for Financial Computing & Analytics

The Leverhulme Trust

Collaborators:

Eric Schulz

Maarten Speekenbrink

Bradley C. Love