Robustness through sparsity: A comparison of decision heuristics

Paula Parpart (paula.parpart@nyu.edu), Eric Schulz (ericschulz@fas.harvard.edu), Bradley C Love (b.love@ucl.ac.uk)

New York University
Harvard University
University College London

INTRODUCTION

The Take-The-Best (TTB) heuristic's success may be due to its **deliberate ignorance of covariance** among cues, which leads to less overfitting (Gigerenzer & Brighton, 2009).

However, paradoxically, TTB is ecologically rational in environments with **high cue redundancy** (Todd & Gigerenzer, 2012; Dieckmann & Rieskamp, 2007; Hogarth & Karelia, 2005).

QUESTION

How come TTB was deliberately designed to ignore covariance, but does especially well when redundancy is high?

STRATEGIES

$$v = \frac{R}{R + W}$$

A

	v			*cue coding*
(1) League pos.	.90	☺	☹	+1
(2) Last game result	.81	😐	😐	0
(3) Home vs. away	.73	☹	☺	-1
(4) No. of goals	.54	☹	☺	-1

B

Take-The-Best
- only uses single cue
- differential weighting
- ignores covariance

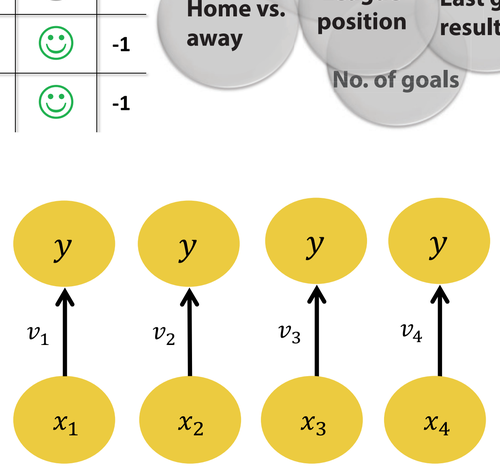

Logistic Regression
- uses all cues
- differential weighting
- estimates covariance

Tallying
- uses all cues
- unit-weighting
- ignores covariance

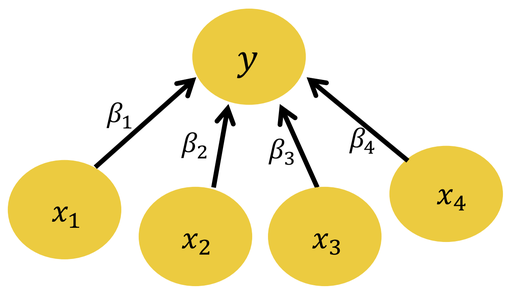

Naïve Bayes
- uses all cues
- differential weighting
- ignores covariance

SIMULATION STUDY

What is the effect of environmental covariance on strategy performance? (Dieckmann & Rieskamp, 2007; 2012)

The strategies' accuracies were evaluated by their generalizability using cross-validation (Pitt & Myung, 2002).

Method
- Systematically varied environmental covariance levels
- Holding cue validity constant: ***v*** = [.89,.82,.76,.69, .62, .56]
- Covariance was optimized with a brute-force, hill-climbing algorithm on the level of average inter-cue correlations
- True response variable was held constant
- 500 environments per covariance condition with the following parameters: *N* = 50, *m* = 6 cues, training/test sample size: 25

RESULTS

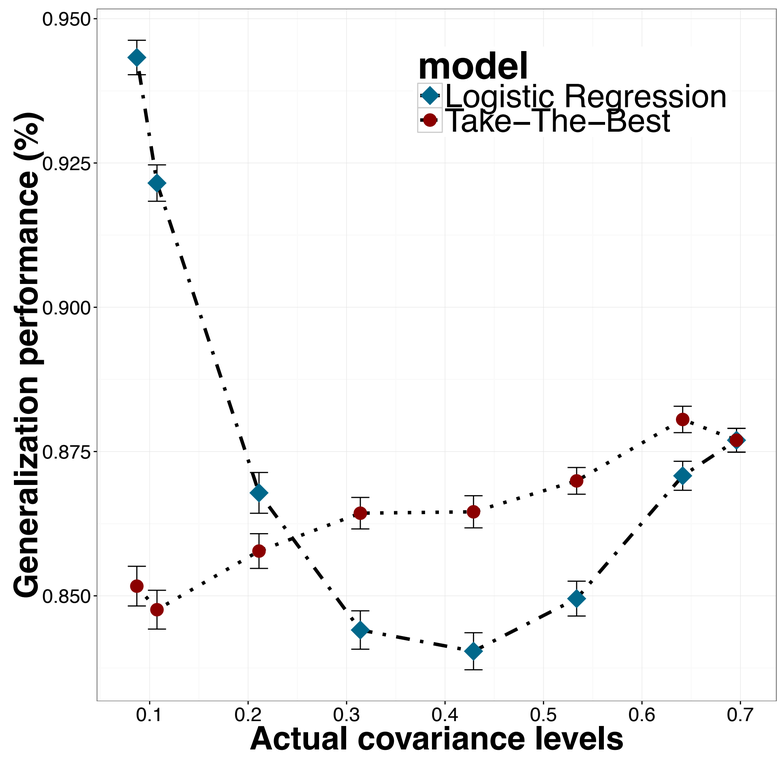

TTB is insensitive to covariance, but succeeds in high-covariance environments, TTB > LOG.

Why is TTB superior in high-redundancy environments?

- Tallying and TTB both ignore covariance, but Tallying fails.
- TTB and Naïve Bayes differ in only one aspect, yet TTB has the advantage, while Naïve Bayes does not.

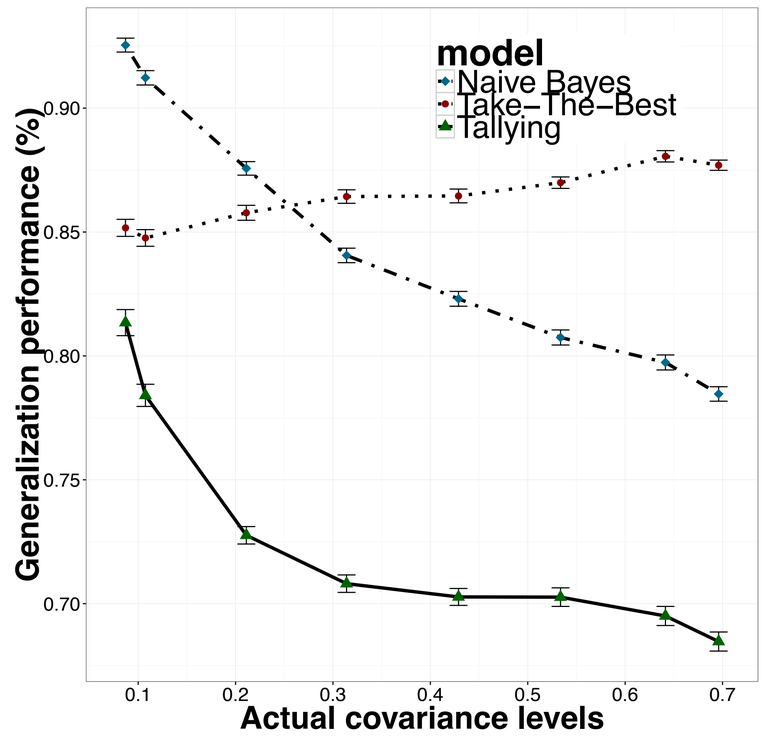

CONCLUSIONS

In contrast to previous assumptions (Gigerenzer & Brighton, 2009), ignorance of covariance alone is not sufficient to explain TTB's success in high-redundancy environments.

Instead, results indicate that TTB's robustness advantage stems from its cue sparsity.

IMPLICATIONS

Only when one knows covariance levels, it becomes possible to judge the ecological rationality of a heuristic (Tallying or TTB) (Parpart et al., 2018).

Can models do well when they have the wrong model of the world?
i.e., possibly rank order/search rule are less important than previously thought?

Research into natural environments that are predictive and contain inter-correlated information should look into **feature extraction and dimensionality reduction** in combination with heuristics**.**

Dieckmann, A., & Rieskamp, J. (2007). The influence of information redundancy on probabilistic inferences. *Memory & Cognition*, *35*(7), 1801-1813.
Gigerenzer, G., & Brighton, H. (2009). Homo heuristicus: why biased minds make better inferences. *Topics in Cognitive Science*, 1(1), 107-43.
Hogarth, R. M., & Karelaia, N. (2005). Ignoring information in binary choice with continuous variables: When is less more? *Journal of Mathematical Psychology*, 49(2), 115–124
Parpart, P., Jones, M., & Love, B. C. (2018). Heuristics as Bayesian inference under extreme priors. *Cognitive Psychology*, 102, 127-144.
Pitt, M. A., & Myung, I. J. (2002). When a good fit can be bad. *Trends Cogn Sci*, 6(10), 421-425.
Rieskamp, J., & Dieckmann, A. (2012). Redundancy: Environment structure that simple heuristics can exploit. In *Ecological rationality: Intelligence in the world* (p. 187-215). New York: Oxford University Press.