Is covariance ignorance responsible for the success of heuristics?

Paula Parpart
New York University, New York City, New York, United States

Eric Schulz
Harvard, Boston, Massachusetts, United States

Abstract

Previous work proposes that heuristics, such as Take-The-Best, may succeed because of deliberate ignorance of covariance in their cue weight estimates as opposed to full-information models (logistic regression). Other studies find that Take-The-Best performs particularly well compared to full-information models in high covariance as opposed to low covariance environments. This poses the question of whether heuristics perform well when there is a mismatch between their covariance prior and the covariance in the environment? We test this by gradually manipulating solely the level of covariance among cues. Indeed, Take-The-Best performs better as average covariance increases, while tallying, nave Bayes and logistic regression worsen. Since both nave Bayes and tallying also disregard covariance but integrate across cues, this indicates the competitive advantage of Take-The-Best stems from relying on a single cue when redundancy is high. We extend previous work by Rieskamp and Dieckmann (2012) and imply a reinterpretation of past Take-The-Bests successes.