

Confirmation Bias Trumps Performance Optimization in Overt Active Learning

Yul Kang

University of Cambridge, Cambridge, United Kingdom

Daniel Wolpert

Columbia University, New York, New York, United States

Mate Lengyel

University of Cambridge, Cambridge, United Kingdom

Abstract

When gathering information, different sources typically have distinct levels of informativeness. Therefore, it is optimal to actively select the source of information to learn from (i.e., perform active learning). It has been debated whether humans optimize task performance in active learning or use a simple heuristic of seeking information that confirms their beliefs. Critically, depending on ones subjective beliefs, confirmation bias can in fact be optimal. Thus, without measuring subjective beliefs, previous approaches were unable to distinguish between these alternatives. Using a perceptual decision-making task, we measured participants subjective beliefs before and after a new piece of information was presented. We then characterized confirmation-based and performance optimizing strategies with respect to these subjective beliefs. We found that participants strategy was dominated by confirmation bias, modulated only weakly by the performance optimization. We discuss potential reasons that may limit performance optimization in active learning.