Modelling Implicit Attitudes with Large Corpora: a comparison of linguistic co-occurrence models

Dermot Lynott
Lancaster University

Kerry O'Brien
Monash University

Louise Connell
Lancaster University

Nick Shryane
University of Manchester

Himanshu Kansal
University of Manchester

Michael Walsh
Universität Stuttgart

Abstract: People often have attitudes and biases of which they may not be consciously aware. The implicit association test (IAT) is one of the key tools used to investigate such attitudes, especially when it comes to controversial topics (e.g., racial prejudice). In an IAT, participants categorise stimuli (words/pictures) that are paired with either positive or negative attributes. The differences in response times to the positive/negative pairings provide a measure of the participant's overall implicit bias. We collated dozens of IATS from the literature, incorporating the data of over 10,000 participants and investigated whether the linguistic co-occurrence patterns of the lexical stimuli used in IATs could be used to predict the level of bias exhibited by participants. We used the lexical co-occurrence frequencies extracted from a large-scale corpus to construct a range of models which were then compared to the reported effect sizes of the IATs. We found that relatively simple models based on raw frequencies and conditional probabilities provided significant correlations with implicit effect sizes, while models incorporating log likelihood estimates yielded much poorer performance. It is possible that individuals' sensitivity to linguistic distributional information may influence responding in the IAT.