

Ensuring fairness in credit scoring

Abstract

How to ensure that credit scoring models are fair and non-discriminatory presents an ongoing challenge to lenders. There is agreement between industry stakeholders that lenders must avoid making loan approval decisions based on these models that would exhibit undesirable bias, placing certain demographic or ethnic groups at a systematic disadvantage in terms of their access to credit. However, considerable uncertainty persists over how to best operationalise this notion, with some lenders choosing to scrutinise and, if necessary, adjust the model outputs in a post-development stage, whereas others are looking for solutions to incorporate fairness criteria into the actual model training. Furthermore, questions remain on how best to deal with historical bias in past application data, or the extent to which the practice of reject inference might risk amplifying past discriminatory bias (by assuming rejected cases are higher risk).

In this talk, we will discuss insights gained thus far from an ongoing project in which we set out to explore a series of different options for incorporating fairness into the various stages of credit scoring. In so doing, we seek to consider the role of reject inference and data debiasing, and trial different solutions on real-life application scoring data. In so doing, our goal is to formulate suggestions for analyses to include in the various steps of scorecard development and validation.

Authors & Affiliations

Mr Mark Thompson¹, Dr Huan Yu², Professor Christophe Mues²

¹Paragon Business Solutions, London, United Kingdom. ²University of Southampton, Southampton, United Kingdom