Peter Flach, Professor of Artificial Intelligence
www.cs.bris.ac.uk/~flach

University of BRISTOL
Jean Golding Institute

- Main research interests
  - Machine learning
  - Data streams
  - Calibration and uncertainty
  - Context change
  - AI for science

Making data work for everyone

Beta calib

Meelis
University
University

ELS

$a = 0.2$

Reframing in context: A systematic appr
for model reuse in machine learning

José Hernández-Orallo
DSIC, Universitat Politècnica de València, Spain, jorallo@dsic.upv.es

Adolfo Martínez-Usó
DSIC, Universitat Politècnica de València, Spain, admarus@dsic.upv.es

Ricardo B.C. Prudêncio
Centro de Informática, Universidade Federal de Pernambuco, Recife (PE), Brazil, rbcp@cin.ufpe.br

Meelis Kull
Department of Computer Science, University of Bristol, UK, meelis.kull@bristol.ac.uk

Peter Flach
Department of Computer Science, University of Bristol, UK, peter.flach@bristol.ac.uk

Chowdhury Farhan Ahmed
ICube, Université de Strasbourg, France, cfahmed@unistra.fr

Nicolas Lachiche
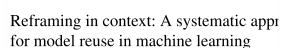ICube, Université de Strasbourg, France, nicolas.lachiche@unistra.fr

DOI:10.1145/2979672

New tools tackle an age-old practice.

BY SIMON PRICE AND PETER A. FLACH

Computational Support for Academic Peer Review: A Perspective from Artificial Intelligence

PEER REVIEW IS the process by which experts in some discipline comment on the quality of the works of others in that discipline. Peer review of written works is firmly embedded in current academic research practice where it is positioned as the gateway process and quality control mechanism for submissions to conferences, journals, and funding bodies across a wide range of disciplines. It is probably safe to assume that peer review in some form will remain a cornerstone of academic practice for years to come, evidence-based criticisms of this process in computer science[22,32,45] and other disciplines[23,28] notwithstanding.

While parts of the academic peer review process have been streamlined in the last few decades to take technological advances into account, there are many more opportunities for computational support that are not currently being exploited. The aim of this article is to identify such opportunities and describe a few early solutions for automating key stages in the established academic peer review process. When developing these solutions we have found it useful to build on our background in machine learning and artificial intelligence: in particular, we utilize a feature-based perspective in which the handcrafted features on which conventional peer review usually depends (for example, keywords) can be improved by feature weighting, selection, and construction (see Flach[17] for a broader perspective on the role and importance of features in machine learning).

Twenty-five years ago, at the start of our academic careers, submitting a paper to a conference was a fairly involved and time-consuming process that roughly went as follows: Once an author had produced the manuscript (in the original sense, that is, manually produced on a typewriter, possibly by someone from the university's pool of typists), he or she would make up to seven photocopies, stick all of them

» **key insights**

- State-of-the-art tools from machine learning and artificial intelligence are making inroads to automate parts of the peer-review process; however, many opportunities for further improvement remain.
- Profiling, matching, and open-world expert finding are key tasks that can be addressed using feature-based representations commonly used in machine learning.
- Such streamlining tools also offer perspectives on how the peer-review process might be improved: in particular, the idea of profiling naturally leads to a view of peer review being aimed at finding the best publication venue (if any) for a submitted paper.
- Creating a more global embedding for the peer-review process that transcends individual conferences or conference series by means of persistent reviewer and author profiles is key, in our opinion, to a more robust and less arbitrary peer-review process.