DOI:10.1145/3231715

James Larus and Chris Hankin

Regulating Automated Decision Making

DISDAIN FOR REGULATION is pervasive throughout the tech industry. In the case of automated decision making, this attitude is mistaken. Early engagement with governments and regulators could both smooth the path of adoption for systems built on machine learning, minimize the consequences of inevitable failures, increase public trust in these systems, and possibly avert the imposition of debilitating rules.

Exponential growth in the sophistication and applications of machine learning is in the process of automating wholly or partially many tasks previously performed only by humans. This technology of automated decision making (ADM) promises many benefits, including reducing tedious labor as well as improving the appropriateness and acceptability of decisions and actions. The technology also will open new markets for innovative and profitable businesses, such as self-driving vehicles and automated services.

At the same time, however, the widespread adoption of ADM systems will be economically disruptive and will raise new and complex societal challenges, such as worker displacement; autonomous accidents; and, perhaps most fundamentally, confusion and debate over what it means to be human.

From a European perspective, this is a strong argument for governments to take a more active role in regulating the use of ADM. The European Union has already started to grapple with privacy concerns through the General Data Protection Regulation (GDPR), which regulates data protec-

tion and requires explanation of automated decisions involving people. However, widespread use of ADM will raise additional ethical, economic, and legal issues. Early attention to these questions is central to formulating regulation for autonomous vehicles. The German Ministry for Transport and Digital Infrastructure created an Ethics Commission, which identified 20 key principles to govern ethical and privacy concerns in automated driving.^a

To raise these concerns more broadly, a group assembled by Informatics Europe and EUACM, the policy committee of the ACM Europe Council, recently produced a report entitled “When Computers Decide.”^b The white paper makes 10 recommendations to policy leaders:

1. Establish means, measures, and standards to assure ADM systems are fair.
2. Ensure ethics remain at the forefront of, and integral to, ADM development and deployment.
3. Promote value-sensitive ADM design.
4. Define clear legal responsibilities for ADM’s use and impacts.
5. Ensure the economic consequences of ADM adoption are fully considered.
6. Mandate that all privacy and data acquisition practices of ADM deployers be clearly disclosed to all users of such systems.
7. Increase public funding for non-commercial ADM-related research significantly.
8. Foster ADM-related technical education at the university level.

9. Complement technical education with comparable social education.

10. Expand the public’s awareness and understanding of ADM and its impacts.

Systems built on an immature and rapidly evolving technology such as machine learning will have spectacular successes and dismaying failures. Especially when the technology is used in applications that affect the safety and livelihood of many people, these systems should be developed and deployed with special care. Society must set clear parameters for what uses are acceptable, how the systems should be developed, how inevitable trade-offs and conflicts will be adjudicated, and who is legally responsible for these systems and their failures.

Automated decision making is not just a scientific challenge; it is simultaneously a political, economic, technological, cultural, educational, and even philosophical challenge. Because these aspects are interdependent, it is inappropriate to focus on any one feature of the much larger picture. The computing professions and technology industries, which together are driving these advances forward, have an obligation to start a conversation among all affected disciplines and institutions whose expertise is relevant and required to fully understand these complex issues.

Now is the time to formulate appropriately nuanced, comprehensive, and ethical plans for humans and our societies to thrive when computers make decisions. □

James Larus, a professor and Dean of the School of Computer and Communication Sciences at EPFL, is on the board of Informatics Europe.

Chris Hankin, chair of ACM Europe Council, is co-director of the Institute for Security Science and Technology and a professor of computing science at Imperial College London.

Copyright held by authors.

a <https://www.bmvi.de/SharedDocs/EN/publications/report-ethics-commission.html>

b <https://dl.acm.org/citation.cfm?id=3185595>