

Seeking Ethical use of AI Algorithms: Challenges and Mitigations

Submission Type: Panel

Monideepa Tarafdar
Lancaster University
United Kingdom
m.tarafdar@lancaster.ac.uk

Mike H.M. Teodorescu
Boston College
Carroll School of Management
mike.teodorescu@bc.edu

Hüseyin Tanriverdi
University of Texas at Austin
McCombs School of Business
huseyin.tanriverdi@mcombs.utexas.edu

Lionel P. Robert Jr.
University of Michigan
School of Information
lprobert@umich.edu

Lily Morse
West Virginia University
John Chambers College of Business
and Economics
lily.morse@mail.wvu.edu

Abstract

This panel will discuss the problems of bias and fairness in organizational use of AI algorithms. The panel will first put forth key issues regarding biases that arise when AI algorithms are applied to organizational processes. We will then propose a socio-technical approach to bias mitigation. We will further share proposals for companies and policymakers on improving AI algorithmic fairness within organizations and bias mitigation. The panel will bring together scholars examining social and technical aspects of bias and its mitigation, from the perspective of information systems, ethics, machine learning, robotics, and human capital. The panel will end with an open discussion of where the field of information systems can step in to guide fairness and ethical use in AI algorithms in the coming years.

Keywords: machine learning; fairness; socio-technical; bias; protected attribute; ethics.

Introduction

This panel will discuss the implications of the emerging ethical issues around artificial intelligence (AI) algorithms. We define AI algorithms as those that extract insights and knowledge from big data sources. Computational and statistical techniques such as machine learning (ML) and deep learning embedded in such algorithms, aim to ‘teach’ computers the ability to do detect patterns in big data (e.g. Tarafdar et al 2019). We focus on ethical issues arising from AI algorithms as mathematical constructs (i.e. logic), algorithms as implementations (i.e. in business process execution), and algorithms as configurations (i.e. an artefact such as an IT application with an embedded algorithm) (Hill 2016).¹

¹ Traditionally, computer scientists defined the term ‘algorithm’ as “a fixed step-by-step procedure for accomplishing a given result; usually a simplified procedure for solving a complex problem, also a full statement of a finite number of steps” (Sippl and Sippl 1972). This definition is inclusive of any procedure or decision process, and hence, it covers a large set of artefacts. In this panel discussion we focus on AI algorithms. Specifically, we adopt Hill’s (2016) definition,

AI algorithms promise many benefits: e.g., better decision-making, improved productivity, streamlined processes that are ‘potentially’ free of human biases, etc. With the pervasive adoption and usage of the AI algorithms across a variety of industries and organizations, large and small, we have also witnessed the emergence of ethical issues (Ajunwa 2020) such as algorithmic biases. We define algorithmic bias as an algorithm’s *systematic* and *unfair discrimination* against certain individuals or groups of individuals in favor of others (Friedman and Nissenbaum 1996). The biases in AI algorithms are intangible and difficult to follow intuitively. The causes of algorithmic biases may not be easily discernible (Serra, 2018).

Firms applying such algorithms are thus susceptible to wrong business decisions, liability, loss of customers, and reputational damage (Ajunwa 2020). Just like cyber resources pose great liabilities for firms through cybersecurity breaches (Gordon et al. 2003; Gordon et al. 2014; Young 2012), the widespread use of ML without precautions for bias and fairness can yield great liabilities, even for large firms with presumably strong software engineering human capital and legal compliance capabilities.

It is thus clear that the same AI applications also lead to non-beneficial consequences such as amplifying existing biases, creating brand new biases, and thus generating new ethical dilemmas. *This panel will draw on the socio-technical approach to shine a light on the challenges and risks that organizations and societies face from the use of such applications and examine ways of their mitigation.*

Issues

Our first contention is that bias cannot be avoided in the use of big data algorithms and that it exists both in the technical as well as the social aspects. Bias, both technical and social, is an inherent part of the use of big data algorithms. Bias and fairness are complementary aspects. At the technical level, ‘fairness’ or an implicit lack of bias implies that the outputs of such algorithms should not discriminate against individuals based on a range of protected attributes such as race and age. Fairness is technically difficult to achieve because the predictive outputs may be biased toward certain attributes, if the historical datasets used to train them are so biased. There are tradeoffs involved in addressing bias, specifically that of fairness versus accuracy (Awwad et al. 2020). Furthermore, satisfying a fairness criterion based on one protected attribute is often at odds with satisfying fairness based on another attribute (for example, fairness across genders or fairness across ethnicities) as the two often create different optima, and achieving subgroup level fairness is not always feasible (Kearns et al. 2018; Dwork et al. 2018).

Bias also arises due to causes that are social in nature (Chouldechova and Roth 2018). For example, if the historical data sets used to train the ML models are biased (e.g. salary disparity among genders), the cause is usually attributable to organizational decisions that have been taken in the past, which are thus biased.

Our second contention is that biases get reinforced in an ongoing way. The legal frameworks for addressing discrimination based on race, religion, gender, national origin, marital status, and socioeconomic status pre-dates the creation of algorithms. For example, in the United States, laws that prevent social discrimination have to do with labor (Federal Equal Employment Opportunity), civil rights (Civil Rights Act 1964), disability rights (Americans with Disabilities Act of 1990), and credit lending (Equal Credit Opportunity Act of 2008). Thus, the notion of bias in business and civic processes is not new. However, the unleashing of ML algorithms on data from past organizational and social processes has brought such biases to the surface in a very prominent way (Robert 2020a). Further, it has raised the very real possibility of further reinforcing such biases in ongoing computations, leading to their exacerbation if unsupervised.

Our third contention is that the human user of the algorithm does not have a ‘voice’ in bias mitigation. AI systems have been found to be unfair to workers – in relation to trust and ethics. Despite this, little attention has been paid to proposing a theoretical and systematic approach to designing fairer AI systems. This is particularly problematic when we consider that in a recent survey of 1,770 managers from 14 counties, 86% of managers stated that they planned to use AI systems for managing their workers and 78% of the

which refers to an algorithm as a mathematical construct with “a finite, abstract, effective, compound control structure, imperatively given, accomplishing a given purpose under given provisions” (Hill 2016). Rather than discussing all ethical issues related to all algorithms, we focus on ethical issues arising from AI algorithms.

managers said they trust decisions made by AI systems (Kolbjørnsrud et al, 2016). Therefore, it is vital for IS scholars to develop safeguards to ensure that applications based on AI algorithms adhere to fairness.

Our fourth contention is that while there are a number of technical approaches to addressing bias, they do not go the full distance. As of the date of writing of this manuscript, there are over twenty different criteria for determining ‘fairness’ in the computer science literature (Verma and Rubin 2018), often with contradictory results (Chouldechova 2017). For instance, statistical metrics are defined based on calculations using the cells of the confusion matrix, where a fair outcome would equalize a measure across categories of the protected attribute: genders have the same positive predictive value (predictive parity), or the same accuracy (accuracy equality), or the same false negative rate (equality of opportunity), to name a few. There are several issues with this seemingly benign approach of picking an additional metric to calculate and check – these group-level fairness criteria are at often at odds with one another (Kleinberg et al 2016). Therefore, the answer of which fairness criterion an algorithm developer must abide to is not something that can be picked by theory in advance, but a case-by-case choice the organizational leader or policy maker must decide.

Another problem is that ensuring fairness at the group level does not guarantee fairness at the individual level (Teodorescu et al. 2020): by equalizing at the group level certain otherwise qualified individuals for the positive outcome (e.g., admissions, or credit approval) would end up disqualified just so that the algorithm can satisfy the group fairness criterion (assuming base rates differ across the groups). This is an issue irrespective of which group fairness criterion one chooses. Furthermore, the issue of ensuring fairness on just one subgroup is that the algorithm may perform very unfairly on other subgroups, lest the user tests for every combination of the protected attributes (Kearns et al. 2018).

Socio-Technical Approach to Addressing Bias

It is clear from the above that when applied to organizational processes, AI leads to a plethora of process errors and ethical dilemmas. This has been seen across a range of fields such as law, medicine, retail, government policies and banking. Yet, bias that leads to lack of fairness is extremely hard to detect and address because it is insidious. We put forth below a few considerations in the panel for how it can possibly be done, which will go into in depth, subject to final selection.

First, we recommend that a socio-technical, rather than a technical approach is necessary. The data science and computer science disciplines have focused on finding technical/analytical solutions, focusing mostly on the algorithmic bias problem. For example, in the tradeoffs in satisfying a fairness criterion based on one selected protected attribute (for example, fairness across genders versus fairness across ethnicities) are ethical matters because technically speaking, not all criteria for fairness can be simultaneously satisfied. From the algorithmic/technical perspective, it is not clear which tests would satisfy the laws in place, especially because different fairness definitions can satisfy the disparate impact and treatment definitions used in the laws related to socioeconomic statutes. Further, many social biases that exist in the broader community are finding their way into AI algorithms, such as in those used for screening applicants, potentially disadvantaging minorities. The algorithms ‘learn’ these biases from past hiring data and then use the learning to make their own recommendations, further perpetuating the bias. Thus, a socio-technical approach that focuses on ongoing interaction between the technical actions and the ethical frameworks and implicit organizational assumptions that govern the selection parameters and data is needed.

Second, it is necessary to develop an organizational governance framework for IS that brings together the complex activities associated with developing, deploying, testing, managing the learning and maintaining the big data algorithms (Tarafdard et al. 2019). In other words, ethically conscious practices need to be embedded in the organization’s treatment of algorithms, particularly those procedures which bestow voice to the human. The management discipline has noted important ways that organizations may consider ethical influences in concert with analytical influences. Of particular note, Tenbrunsel and colleagues (2003) have called for firms to extend beyond more common ethical “fixes”, such as codes of conduct, mission statements, and written standards, which represent the “tip of the ethical infrastructure”, and instead focus on highlighting the unofficial values they wish to convey. Although informal systems are less visible than explicit documents, they are more closely tied to the actual behavior of employees and thus offer greater potential for mitigating social bias. Further, AI accountability is getting less attention but is vital to addressing any questions related to ethical dilemmas (Robert et al. 2020a; Robert et al. 2020b). The

issue of accountability is complex because organizations may outsource business processes (e.g. resume screening) to third party companies, who may use AI algorithms in the outsourced process. In situations such as these it is difficult to answer the question: Who should be held accountable for the ‘unfair’ actions of an AI algorithm? The problem of AI accountability is slowly being acknowledged by policy makers (World Economic Forum, 2020). Finally, AI audits are a way to demonstrate that the AI complies with laws, regulations, or policies (Robert et al. 2020a).

Third, both companies and policy makers need to be involved. Both companies and policymakers need to think through the consequences of bias. Current corporate thinking does not acknowledge the reality that algorithmic flaws can lead to wide-ranging discrimination and unethical behaviors, at scale. Organizations, knowingly or unknowingly using such algorithms are thus potentially subject to far greater penalties than those applicable to smaller scale, individual decision-maker driven discrimination. As many countries (as well as the United Nations) have laws protecting against discrimination of individuals based on attributes such as race, gender, age, socioeconomic status, marital status, nationality, and ethnicity (known as ‘protected’ attributes), organizations need to ensure that the widespread algorithms they use in AI applications, do not discriminate. Penalties for discrimination can be severe, ranging from financial to reputational damages. Further, there are also deeper societal issues because biases from AI algorithms can lead to unfair hiring and recruitment and loan processes, potentially resulting in displaced workers and affordable-housing shortage. Such challenges are leading to further inequalities in our society.

Fourth, we caution that ethical issues relating to AI are context dependent. For example, they are not only pertinent to large companies, they also matter to small firms in developing countries (Awwad et al., 2020). Further, the very definitions of bias may differ across geographies and cultures (Robert et al. 2020b). These challenges are also global in nature. This means the solutions may not be scalable across borders. Most current discussions on the topic are focused on North America because the uptake of AI is most prominent in North American companies. However, the same issues may not hold for other societies. For example, women, minorities, and people of color may be protected groups in one country but may not be so in another. Protected groups are not universal and often do not translate well across borders and there are different ethical frameworks among different nations. To address issues such as these, local values need to be considered.

The view of this panel therefore is that ethical use of algorithms is a multi-layered, complex, and socio-technical problem, calling for mitigations that encompass multiple aspects and dimensions. These include actions by IT and business leaders, technical experts, and those in government and policy. The panel will thus discuss the following topics – (1) socio-technical challenges and risks in applying algorithms to organizational processes in the private and public sectors; (2) technical points of breakdown in achieving fairness in machine learning; (3) socio-technical solutions and mitigations for the ethical use of algorithms.

Panelist Positions

Collectively, the panelists will cover the points of view outlined above. Their specific positions are outlined below.

Monideepa Tarafdar will chair and moderate the panel. She will introduce the panel with an overall socio-technical framework of ethical use of applications-based AI algorithms.

Mike Teodorescu: will be introducing the concept of protected attributes, typical definitions of fairness in machine learning, performance metrics for machine learning algorithms, and real-life tradeoffs which make fairness at odds with performance. He will bring examples from fairness applied to credit lending, fairness applied to criminal justice, and fairness applied to hiring contexts. He will share how fairness criteria are often at odds with the optimum algorithm performance and discuss challenges at the forefront of the field including tradeoffs algorithm developers may be forced to make. He will draw on examples from the computer science literature, as well as some of his own work with industry.

Hüseyin Tanriverdi will discuss organizational-level governance and control mechanisms that prove effective in mitigating different types of algorithmic biases and the damages caused by the biases. Hüseyin will first present a typology of algorithmic biases observed across a variety of industries. Then, he will present a typology of organizational mitigation mechanisms developed for addressing the most commonly observed algorithmic biases in organizations. Building on an empirical study of 115 biased ML and AI

algorithms across a variety of industries, Hüseyin will discuss which organizational governance and control mechanisms serve as mitigation mechanisms in reducing the emergence of the algorithmic biases and the damages associated with them.

Lionel Robert will discuss the issues that underlie ethics and unfairness in the context of robots and autonomous vehicles. The introduction of autonomous vehicles has been lauded as one way to make roads safer while reducing our carbon footprint. Autonomous vehicles constitute one of the fastest growing areas of investment in the automotive industry, worldwide. Yet, the introduction of autonomous vehicles brings to the forefront new ethical challenges and problems associated with societal unfairness. Lionel will outline these challenges and emerging mitigation strategies drawing on his research on fair AI systems for managing employees in organizations.

Lily Morse will discuss the psychology of ethical decision making for organizations in AI settings. Her part of the talk will unpack why employees and organizations may continue to behave unethically despite the reduction of algorithmic biases in machine learning tools. She will speak on how overreliance on automation may lead to ethical fading and other ethical blind spots and will offer a empirically-based recommendations for nudging companies toward designing more efficient and aware ethical infrastructures relating to AI algorithms.

Panel Structure

The first part of the panel will have each panelist speaking for 6 minutes, i.e. a total of about 30 minutes. This will be followed by the second part of the panel, with 30 minutes of moderated discussion driven by audience questions. The audience can ask questions in two ways: as the audience in the panel auditorium or through a twitter hashtag, which, if possible, we will display on the screen via a laptop. We will setup an online forum to collect questions and issues from the AIS community on the topic. We intend to use these inputs to shape the debate during the panel discussion.

Participation Statement

All participants have made a commitment to register and attend the conference virtually and serve on the panel if the panel is accepted.

Biographies of the Panelists in Alphabetical Order

Dr. Lily Morse is an Assistant Professor at the John Chambers College of Business & Economics at West Virginia University. Her research examines unethical behavior in the workplace, which she has investigated through the lens of moral character and prosocial deviance. Specifically, she has examined the nature of moral character and how moral character traits manifest across various occupational settings, such as negotiation and independent auditing. Her work has also explored how interpersonal relationships, including relationships with artificial intelligence tools, undermine ethical decision making and behavior. Dr. Morse has published in journals such as *Organizational Behavior and Human Decision Processes*, *Academy of Management Perspectives*, *Journal of Research in Organizational Behavior*, *Journal of Personality and Social Psychology*, and *Journal of Research in Personality*. She received her PhD in Organizational Behavior and Theory from Carnegie Mellon University.

Dr. Lionel Robert Jr. is an Associate Professor in the University of Michigan School of Information and core faculty at the University of Michigan Robotics Institute. His researcher focuses broadly on collective action through technology and human collaboration with autonomous systems. Lionel is the director of the Michigan Autonomous Vehicle Research Intergroup Collaboration (MAVRIC), an affiliate of the Center for Hybrid Intelligence Systems and the National Center for Institutional Diversity all at the University of Michigan and the Center for Computer-Mediated Communication at Indiana University. He is currently on the editorial board of *MIS Quarterly*, *Journal of the AIS*, *ACM Transactions on Social Computing* and the *AIS Transactions on Human-Computer Interaction*. Dr. Robert has published in journals such as *ISR*, *JMIS* and *JAIS*. He has also published in premier HCI conferences such as *CHI*, *CSCW*, *Group*, *HRI* as well as premier data and computational science conferences such as *ICWSM*, *WSDM* and the *Web Conference WWW*. His research has been sponsored by the U.S. Army, Toyota Research Institute, MCity, Lieberthal-Rogel Center for Chinese Studies, AAA Foundation, and the National Science Foundation. He has appeared in print, radio and/or television for such outlets as ABC, CNN, CNBC, and Michigan Radio.

Dr. Hüseyin Tanriverdi is an Associate Professor of information, risk, and operations management at the Red McCombs School of Business at the University of Texas at Austin. His research focuses on firm-level risk/return implications of digital technologies. On the risk side, he studies causes, consequences, and mitigation mechanisms of digital technology related risks such as cybersecurity, privacy, and algorithmic risks (e.g., biases in ML and AI algorithms). On the return side, he studies strategic uses of digital technologies to survive and improve firm performance in complex, hypercompetitive, disruptive business ecosystems. Hüseyin teaches courses on cybersecurity, IT governance and controls for enterprise risk management and regulatory compliance, and strategic IT management. His research has been published in information systems journals such as *Information Systems Research*, *MIS Quarterly*, *Journal of the Association for Information Systems*, and *European Journal of Information Systems*, and management journals such as *Academy of Management Journal* and *Strategic Management Journal*. His publications received Best Published Paper Awards from the Organizational Communications and Information Systems Division of the Academy of Management and the *Telemedicine Journal*.

Dr. Monideepa Tarafdar is a Professor of Information Systems at Lancaster University (Management School) in the UK, where she also co-directs the Centre for Technological Futures. She is a Lever Hulme Research Fellow (UK) and a Research Affiliate at the MIT Sloan Center for Information Systems Research (US). She is currently the Principal Investigator of a grant from the UK's Economic and Social Science Research Council (ESRC) on AI related biases and discrimination in HRM processes and labor markets. She serves as Senior Editor at *Information Systems Journal*, Associate Editor at *Information Systems Research*, and Editorial Review Board member at *Journal of the AIS* and *Journal of Strategic Information Systems*. She is a member of the Digital Skills and Research Working Group, a policy advisory group under the UK Government's *Department for Digital, Culture, Media & Sport*. Her research has been covered by the BBC, Boston Globe, MIT Technology Review, The Medium, Hindustan Times, and Tsinghua News

Dr. Mike Teodorescu is an Assistant Professor of Information Systems at Boston College and an affiliate of D-Lab at Massachusetts Institute of Technology where he studies fairness in machine learning. He teaches Data Analytics at the MBA and college levels. His research is on applying machine learning to innovation as well as criteria for fair use of machine learning within organizations. During his doctoral years at Harvard Business School he studied the innovation ecosystem in the United States and strategic uses of patents in technology firms. He is an award-winning inventor of several technologies.

Acknowledgements

Monideepa Tarafdar acknowledges funding from the UK Research and Innovation Council, the Economic and Social Research Council (UK), the Canadian Institutes of Health Research (CIHR), the Natural Sciences and Engineering Research Council (Canada) and the Social Sciences and Humanities Research Council (SSHRC) for funding for the project 'BIAS: Responsible AI for Labour Market Equality' (Ref: ES/To12382/1).

Mike Teodorescu thanks United States Agency for International Development (USAID) Grant AID-OAA-A-12-00095 "Appropriate Use of Machine Learning in Developing Country Contexts" (awarded to MIT D-Lab and funded partly this author) and Boston College which funded research referenced in this panel, as well as Dr. Sam Ransbotham, Dr. Robert Fichman, Dr. Gerald Kane (all Boston College), Dr. Daniel Frey, Kendra Leith, Nancy Adams, Amit Gandhi (all MIT), Dr. Aubra Anthony, Dr. Shachee Doshi, Dr. Craig Jolley, Dr. Amy Paul, Dr. Maggie Linak (all USAID).

Opinions, findings, conclusions, or recommendations expressed here are of the authors and do not necessarily represent the views of any of the above funding bodies.

References

- Ajunwa, I., 2020. "The Paradox of Automation as Anti-Bias Intervention". *Cardozo Law Review* (41).
- Awwad, Y., Fletcher, R., Frey, D., Gandhi, A., Najafian, M., Teodorescu, M. 2020. "Exploring Fairness in Machine Learning for International Development". *MIT D-Lab | CITE Report*. Cambridge: MIT D-Lab, URI: <https://hdl.handle.net/1721.1/126854>.
- Chouldechova, A., 2017. "Fair prediction with disparate impact: A study of bias in recidivism prediction instruments". *Big data* 5(2), pp.153-163.

- Chouldechova, A. and Roth, A., 2018. "The frontiers of fairness in machine learning". *arXiv preprint*, arXiv:1810.08810.
- Dwork, C. and Ilvento, C., 2018. "Group fairness under composition". In *Proceedings of the 2018 Conference on Fairness, Accountability, and Transparency*.
- Friedman, B. and Nissenbaum H. 1996. "Bias in Computer Systems". *ACM Transactions on Information Systems* 14(3): 330–347.
- Gordon, L.A., Loeb, M.P. and Sohail, T., 2003. "A framework for using insurance for cyber-risk management". *Communications of the ACM* 46(3): 81-85.
- Gordon, L.A., Loeb, M.P., Lucyshyn, W. and Zhou, L., 2014. "Externalities and the magnitude of cyber security underinvestment by private sector firms: a modification of the Gordon-Loeb model". *Journal of Information Security* 6(01), p.24.
- Hill, R. K. 2016. "What an Algorithm Is". *Philosophy & Technology* 29: 35-59.
- Kearns, M., Neel, S., Roth, A. and Wu, Z.S., 2018. "Preventing fairness gerrymandering: Auditing and learning for subgroup fairness". In *Proceedings of the 35th International Conference on Machine Learning* 80:2564-2572.
- Kleinberg, J., Mullainathan, S. and Raghavan, M., 2016. "Inherent trade-offs in the fair determination of risk scores". *arXiv preprint* arXiv:1609.05807.
- Kolbjørnsrud, V., Amico, R., & Thomas, R. J., 2016. "How artificial intelligence will redefine management. Harvard Business Review". Retrieved from <https://hbr.org/2016/11/how-artificial-intelligence-will-redefine-management>
- Robert, L. P., Pierce, C., Marquis, E., Kim, S., Alahmad, R. 2020a. "Designing Fair AI for Managing Employees in Organizations: A Review, Critique, and Design Agenda", *Human-Computer Interaction*, pp. 1-31, DOI: <https://doi.org/10.1080/07370024.2020.1735391>
- Robert, L. P., Gaurav, B., & Lütge, C. (2020b). "ICIS 2019 SIGHCI Workshop Panel Report: Human–Computer Interaction Challenges and Opportunities for Fair, Trustworthy and Ethical Artificial Intelligence." *AIS Transactions on Human-Computer Interaction* 12(2): 96-108, DOI: 10.17705/1thci.00130.
- Serra, J. 2018. "When the State of the Art Is Ahead of the State of Understanding: Unintuitive Properties of Deep Neural Networks," *Mètode Science Studies Journal*, pp. 127-132. doi:10.7203/metode.9.11035.
- Sippl, C. J. and C. P. Sippl 1972. *Computer Dictionary and Handbook*. Indianapolis, Indiana, Howard W. Sams & Co., Inc.
- Tarafdar, M., Beath, C.M., Ross, J.W. 2019, "Using AI to Enhance Business Operations", *MIT Sloan Management Review*, June 2019.
- Teodorescu, M. H., Morse, L., Awwad, Y., Kane, J. 2020a, "A Fairness Framework for Machine Learning in Organizations", In *Academy of Management Proceedings* 2020(1): 16889. Briarcliff Manor, NY 10510: Academy of Management.
- Tenbrunsel, A.E., Smith-Crowe, K. and Umphress, E.E., 2003. "Building houses on rocks: The role of the ethical infrastructure in organizations". *Social justice research* 16(3): 285-307.
- Verma, S. and Rubin, J. 2018. "Fairness definitions explained". *2018 IEEE/ACM International Workshop on Software Fairness (FairWare)*, pp. 1-7. IEEE.
- Young, S., 2012. "Contemplating corporate disclosure obligations arising from cybersecurity breaches". *J. Corp. L.* 38, p.659.
- World Economic Forum. 2020. "Model artificial intelligence governance framework and assessment guide. World Economic Forum." Retrieved from <https://www.weforum.org/projects/model-ai-governance-framework>