



THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Tractable Probabilistic Models for Ethical AI

Citation for published version:

Belle, V 2022, Tractable Probabilistic Models for Ethical AI. in T Braun, D Cristea & R Jäschke (eds), *Graph-Based Representation and Reasoning*. Lecture Notes in Computer Science, vol. 13403, Springer International Publishing, pp. 3-8, 27th International Conference on Conceptual Structures, Münster, North Rhine-Westphalia, Germany, 12/09/22. https://doi.org/10.1007/978-3-031-16663-1_1

Digital Object Identifier (DOI):

[10.1007/978-3-031-16663-1_1](https://doi.org/10.1007/978-3-031-16663-1_1)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

Graph-Based Representation and Reasoning

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.



Tractable Probabilistic Models for Ethical AI (Abstract of Invited Talk)[★]

Vaishak Belle

University of Edinburgh & Alan Turing Institute
vbelle@ed.ac.uk

Abstract. Among the many ethical dimensions that arise in the use of ML technology, three stand out as immediate and profound: enabling the interpretability of the underlying decision boundary, addressing the potential for learned algorithms to become biased against certain groups, and capturing blame and responsibility for a system’s outcomes. In this talk, we advocate for a research program that seeks to bridge tractable (probabilistic) models for knowledge acquisition with rich models of autonomous agency that draw on philosophical notions of beliefs, intentions, causes and effects.

1 Motivation

Machine learning (ML) techniques have become pervasive across a range of different applications, and are now widely used in areas as disparate as recidivism prediction, consumer credit-risk analysis, and insurance pricing [8, 25]. Likewise, in the physical world, machine learning models are critical components in autonomous agents such as robotic surgeons and self-driving cars. Among the many societal/ethical dimensions that arise in the use of ML technology in such applications, three stand out as immediate and profound. First, to increase trust and accommodate human insight, interpretability of the underlying decision boundary is essential. Second, there is the potential for learned algorithms to become biased against certain groups, which needs to be addressed. Third, in so much that the decisions of ML models impact society, both virtually (e.g., denying a loan) and physically (e.g., accidentally driving into a pedestrian), the enabling of blame and responsibility is a significant challenge.

Many definitions have been proposed in the literature for such ethical considerations [19, 2], but there is considerable debate about whether a formal notion is appropriate at all, given the rich social contexts that occur in human-machine interactions. Valid arguments are also made about the challenges of model building and deployment [10, 11]: everything from data collection to ascribing responsibility when technology goes awry can demonstrate and amplify abuse of power and privilege. Such issues are deeply intertwined with legal and regulatory problems [15, 32].

[★] This article is a written version of the keynote to be given at *27th International Conference on Conceptual Structures*, September 2022, Münster, Germany. A preliminary version of this work was also presented at the *Critical Perspectives on Artificial Intelligence Ethics Conference* in Edinburgh, UK, 2020. The author was supported by a Royal Society University Research Fellowship.

Be that as it may, what steps can be taken to enable ethical decision-making a reality in AI systems? Human-in-the-loop systems are arguably required given the aforementioned debate [34, 24], but such loops still need to interface with an automated system of considerable sophistication that in the very least reasons about the possible set of actions. In particular, simply delegating responsibility of critical decisions to humans in an ad hoc fashion can be problematic. Often critical actions can be hard to identify immediately and it is only the ramification of those actions that raise alarm, in which case it might be too late for the human to fix. Moreover, understanding the model’s rationale is a challenge in itself, as represented by the burgeoning field of explainable artificial intelligence [29, 14, 4]. So a careful delineation is needed as to which parts are automated, which parts are delegated to humans, which parts can be obtained from humans a-priori (i.e., so-called *knowledge-enhanced machine learning* [9]), but also how systems can be made to reason about their environment so that they are able to capture and deliberate on their choices, however limiting their awareness of the world might be. In the very least, the latter capacity offers an additional layer of protection, control and explanation before delegating, as the systems can point out which beliefs and observations led to their actions.

2 Two-pronged approach

In that regard, our view is that a two-pronged approach is needed in the least. On the one hand, we have to draw on philosophical notions and look to formalise them, as attempted by the knowledge representation community. Indeed this community has looked to capture beliefs, desires, intentions, and causality in service of formal notions that provide an idealised perspective on epistemology grounded in, say, a putative robot’s mental state [5, 16, 22, 21]. But the topic of knowledge acquisition, i.e., how the relevant propositions can be acquired from data is largely left open. Moreover, the topic of reasoning, i.e., of computing truths of acquired knowledge is a long-standing challenge owing to the intractability of propositional reasoning and the undecidability of first-order logic, and many higher-order logics.

On the other hand, although machine learning systems do successfully address acquisition from data, mainstream methods focus on atomic classification tasks, and not the kind of complex reasoning over physical and mental deliberation that humans are adept in. (There are exceptions from robotics and reinforcement learning, of course, but these all attempt some form of mental state modeling [1], and in the very least, reasoning about possible worlds [31].) Moreover, issues about robustness in the presence of approximate computations remain.

In this talk, we advocate for a research program that seeks to apply tractable (probabilistic) models to problems in fair machine learning and automated reasoning of moral principles. Such models are compilation targets for propositional and finite-domain relational logic, and so can represent certain types of knowledge representation languages. They can also be learned from data. We report on a few preliminary results [7, 33, 17, 27, 26, 23, 28, 12, 35]. Firstly, we discuss results on studying causality-related properties in such models, and extracting counterfactual explanations from them. On the topic of fairness, it is shown that the approach enables an effective technique for determin-

ing the statistical relationships between protected attributes and other training variables. This could then be applied as a pre-processing step for computing fair models. On the topic of moral responsibility, it is shown how models of moral scenarios and blameworthiness can be extracted and learnt automatically from data as well as how judgements be computed effectively. In both themes, the learning of the model can be conditioned on expert knowledge allowing us to represent and reason about the domain of interest in a principled fashion.

3 Closing Remarks

We conclude with key observations about the interplay between tractability, learning and knowledge representation in the context of ethical decision-making. Among other things, we observe that the tractable model paradigm is in its early years, at least as far capturing a broad range of knowledge representation languages is concerned, and moreover, there is altogether less emphasis on mental modeling and agency. (First-order expressiveness is yet another dimension for allowing richness in specifications, as are proposals with an explicit causal theory such as [30]). In contrast, readers may want to consult discussions in [24, 23] on knowledge representation approaches where a more comprehensive model of the environment and its actors is considered, but where knowledge acquisition and learning are used in careful, limited ways.

Analogously, we observe that although many expressive languages [6, 18] are known to compile to tractable models, this is purely from the viewpoint of reasoning, or more precisely, probabilistic query computation. What is likely needed is a set of strategies for reversing this pipeline: from a learned tractable model, we need to be able to infer high-level representations. In the absence of general strategies of that sort, the more modest proposal is perhaps to interleave declarative knowledge for high-level patterns but allow low-level patterns to be learnt, which then are altogether compiled for tractable inference.

Overall, the discussed results can be seen occupying positions on a spectrum: the fairness result simply provides a way to accomplish de-biasing, but does not engage with a specification of the users or the environment in any concrete way. Thus, it is closer to mainstream fairness literature. The moral reasoning result is richer in that sense, as it explicitly accounts for actors and their actions in the environment. However, it does not explicitly infer how these actions and effects might have come about – these might be acquired via learning, for example – nor does it reason about what role these actions play amongst multiple actors in the environment. Thus, clearly, in the long run, richer formal systems are needed, which might account for sequential actions [3] and multiple agents [20]. However, this reverts the position back to the issues of tractability and knowledge acquisition not being addressed in such proposals. So, the question is this: can we find ways to appeal to tractable probabilistic models (or other structures with analogous properties) with such rich formal systems? As mentioned, it is known that certain probabilistic logical theories can be reduced to such structures, so perhaps gentle extensions to those theories might suggest ways to integrate causal epistemic models and tractable learning.

Beyond that technical front, much work remains to be done, of course, in terms of delineating automated decision-making from delegation and notions of accountability [13]. It is also worth remarking that computational solutions of the sort discussed in the previous section do make strong assumptions about the environment in which the learning and acting happens. In a general setting, even data collection can amplify positions of privilege, and moreover, there are multiple opportunities for failure and misspecification [10, 11]. Orchestrating a framework where this kind of information and knowledge can be communicated back to the automated system is not at all obvious, and is an open challenge. In that regard, the two-pronged approach is not advocated as a solution to such broader problems, and indeed, it is unclear whether abstract models can imbibe cultural and sociopolitical contexts in a straightforward manner. However, it at least allows us to specify norms for human-machine interaction, provide goals and situations to achieve, model the machine’s beliefs, and allow the machine to entertain models of the user’s knowledge. Ultimately, the hope is that the expressiveness argued for offers additional protection, control and explanation during the deployment of complex systems with machine learning components.

References

1. S. V. Albrecht and P. Stone. Autonomous agents modelling other agents: A comprehensive survey and open problems. *Artificial Intelligence*, 258:66–95, 2018.
2. C. Allen, I. Smit, and W. Wallach. Artificial morality: Top-down, bottom-up, and hybrid approaches. *Ethics and Information Technology*, 7(3):149–155, 2005.
3. V. Batusov and M. Soutchanski. Situation calculus semantics for actual causality. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.
4. V. Belle and I. Papantonis. Principles and practice of explainable machine learning. *arXiv preprint arXiv:2009.11698*, 2020.
5. R. J. Brachman, H. J. Levesque, and R. Reiter. *Knowledge representation*. MIT press, 1992.
6. G. V. d. Broeck, I. Thon, M. v. Otterlo, and L. D. Raedt. Dtpblog: A decision-theoretic probabilistic prolog. In *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence*, AAAI’10, pages 1217–1222. AAAI Press, 2010.
7. Y. Choi, M. Dang, and G. V. d. Broeck. Group fairness by probabilistic modeling with latent fair decisions. *arXiv preprint arXiv:2009.09031*, 2020.
8. A. Chouldechova. Fair prediction with disparate impact: A study of bias in recidivism prediction instruments. In *Big Data, Volume 5, No. 2: 153–163*, 2017.
9. F. G. Cozman and H. N. Munhoz. Some thoughts on knowledge-enhanced machine learning. *International Journal of Approximate Reasoning*, 2021.
10. K. Crawford. *The Atlas of AI*. Yale University Press, 2021.
11. K. Crawford. The hidden costs of ai. *New Scientist*, 249(3327):46–49, 2021.
12. A. Darwiche. Causal inference using tractable circuits. *arXiv preprint arXiv:2202.02891*, 2022.
13. V. Dignum. *Responsible artificial intelligence: how to develop and use AI in a responsible way*. Springer Nature, 2019.
14. F. Doshi-Velez, M. Kortz, R. Budish, C. Bavitz, S. Gershman, D. O’Brien, K. Scott, S. Schieber, J. Waldo, D. Weinberger, et al. Accountability of ai under the law: The role of explanation. *arXiv preprint arXiv:1711.01134*, 2017.
15. A. Etzioni and O. Etzioni. Incorporating ethics into artificial intelligence. *The Journal of Ethics*, 21(4):403–418, 2017.

16. R. Fagin, Y. Moses, J. Y. Halpern, and M. Y. Vardi. *Reasoning about knowledge*. MIT press, 2003.
17. G. Farnadi, B. Babaki, and L. Getoor. Fairness in relational domains. In *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society*, pages 108–114, 2018.
18. D. Fierens, G. Van den Broeck, I. Thon, B. Gutmann, and L. De Raedt. Inference in probabilistic logic programs using weighted CNF's. In *Proceedings of UAI*, pages 211–220, 2011.
19. S. A. Friedler, C. Scheidegger, and S. Venkatasubramanian. On the (im) possibility of fairness. *arXiv preprint arXiv:1609.07236*, 2016.
20. H. Ghaderi, H. Levesque, and Y. Lespérance. Towards a logical theory of coordination and joint ability. In *Proceedings of the 6th international joint conference on Autonomous agents and multiagent systems*, pages 1–3, 2007.
21. J. Y. Halpern. *Actual causality*. MIT Press, 2016.
22. J. Y. Halpern. *Reasoning about uncertainty*. MIT press, 2017.
23. L. Hammond and V. Belle. Learning tractable probabilistic models for moral responsibility and blame. *Data Mining and Knowledge Discovery*, 35(2):621–659, 2021.
24. S. Kambhampati. Challenges of human-aware ai systems. *AI Magazine*, 41(3), 2020.
25. A. Khandani, J. Kim, and A. Lo. Consumer credit-risk models via machine-learning algorithms. In *Journal of Banking and Finance* 34, 2010.
26. I. Papantonis and V. Belle. Interventions and counterfactuals in tractable probabilistic models. In *NeurIPS Workshop on Knowledge Representation & Reasoning Meets Machine Learning*, 2019.
27. I. Papantonis and V. Belle. Closed-form results for prior constraints in sum-product networks. *Frontiers in Artificial Intelligence*, 4, 2021.
28. I. Papantonis and V. Belle. Principled diverse counterfactuals in multilinear models. *arXiv preprint arXiv:2201.06467*, 2022.
29. C. Rudin. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, 1(5):206–215, 2019.
30. B. Salimi, H. Parikh, M. Kayali, L. Getoor, S. Roy, and D. Suciu. Causal relational learning. In *Proceedings of the 2020 ACM SIGMOD International Conference on Management of Data*, pages 241–256, 2020.
31. S. Sardina, G. De Giacomo, Y. Lespérance, and H. J. Levesque. On the limits of planning over belief states under strict uncertainty. *KR*, 6:463–471, 2006.
32. J. Stilgoe. Machine learning, social learning and the governance of self-driving cars. *Social studies of science*, 48(1):25–56, 2018.
33. M. Varley and V. Belle. Fairness in machine learning with tractable models. *Knowledge-Based Systems*, 215:106715, 2021.
34. F. M. Zanzotto. Human-in-the-loop artificial intelligence. *Journal of Artificial Intelligence Research*, 64:243–252, 2019.
35. M. Zečević, D. Dhami, A. Karanam, S. Natarajan, and K. Kersting. Interventional sum-product networks: Causal inference with tractable probabilistic models. *Advances in Neural Information Processing Systems*, 34, 2021.