

On Fairness and Interpretability

Deepak P¹ Sanil V² Joemon M. Jose³

¹Queen’s University Belfast ²IIT Delhi ³University of Glasgow
deepaksp@acm.org sanil@hss.iitd.ernet.in Joemon.Jose@glasgow.ac.uk

Abstract

Ethical AI spans a gamut of considerations. Among these, the most popular ones, fairness and interpretability, have remained largely distinct in technical pursuits. We discuss and elucidate the differences between fairness and interpretability across a variety of dimensions. Further, we develop two principles-based frameworks towards developing ethical AI for the future that embrace aspects of both fairness and interpretability. First, *interpretability for fairness* proposes instantiating interpretability within the realm of fairness to develop a new breed of ethical AI. Second, *fairness and interpretability* initiates deliberations on bringing the best aspects of both together. We hope that these two frameworks will contribute to intensifying scholarly discussions on new frontiers of ethical AI that brings together fairness and interpretability.

1 Introduction

AI algorithms are being increasingly used for decision making within scenarios with social and political consequences (e.g., benefit eligibility, recidivism) as opposed to traditional automation scenarios (e.g., handwriting recognition). This has rightly spurred recent interest in *Ethical AI*. The broad umbrella of ethical AI or responsible AI [Dignum, 2019] involves considerations such as fairness [Chouldechova and Roth, 2020], interpretability¹ [Molnar, 2020], privacy [Mohassel and Zhang, 2017] and trustworthiness [Toreini *et al.*, 2020]. Of these, fairness and interpretability are arguably the two considerations that have been explored quite heavily. Certain pairs of considerations, such as interpretability and trustworthiness, may be seen as apparently synergistic. There is much less understanding of how technological efforts across other pairs of considerations, such as fairness and interpretability as we consider here, can work together.

Our Contributions: In this position paper, we first discuss the conceptual distinctions between fairness and interpretability as it applies to AI/ML. Next, we propose two frontiers of AI research in which efforts towards fairness and in-

¹Interpretability, as we use in this paper, is quite related to, and often interchangeable with the notion of explainability.

terpretability can be meaningfully blended towards advancing ethical AI in unique and novel ways. Where appropriate, we will use example scenarios from unsupervised data-driven AI to illustrate the arguments. This is motivated partly by the fact that unsupervised learning is relatively less explored within ethical AI, making it malleable to fresh thought leads.

2 Fairness & Interpretability: Distinctions

Fairness and Interpretability have largely been, within scholarly literature, seen as two distinct and different frontiers under the broader *Ethical AI* umbrella. The distinctions may be characterized under various dimensions as we discuss below.

First, the family of fairness principles [Narayanan, 2018] are normative values that relate to the politics of AI-driven decision making [Wong, 2019]. On the other hand, interpretability considers user’s ability to understand decisions, and lies at the interface between humans and AI. In other words, technological efforts towards deepening fairness would benefit from better grounding on political philosophy, whereas interpretability would have an analogous relationship with behavioral sciences. In fact, much interpretable and explainable AI work has appeared within HCI-related computing avenues. *Second*, there is a tension between fairness and accuracy (or any other utilitarian metric, say, efficiency), and similarly, there is a tension between interpretability and accuracy. However, these tensions are in *different directions*, as recently outlined in [Kleinberg and Mullainathan, 2019]. The authors illustrate that a simple and interpretable model can be strictly improved on both fairness and accuracy by making it more complex (thus reducing interpretability). In other words, there exists a tension between fairness and interpretability. *Third*, a system that produces interpretable results enhances user’s trust in the AI in ways that a fair AI may not (at least, in the short-term). This means that interpretability is arguably likely to be more ‘popular’, and thus would be prioritized over fairness by the private sector which is where most AI is developed. *Fourth*, interpretability can be assessed at the level of individual decisions made by the AI, whereas fairness assessments need to go much beyond analyzing individual decisions. Fairness assessment often involves a normative evaluation of the process and/or the distribution of decisions made. *Fifth*, there is a distinct contrast between the subjectivity of fairness and interpretability. Fairness is subjective at a normative level, and the subjectivity

is often grounded in political positions; supporters of *individual fairness* are likely to be to the right of *group fairness*, within the left-right political spectrum. On the other hand, interpretability is often highly subjective in *politically neutral* ways; the same result or explanation may be regarded as less interpretable by one user, whereas it could be rated as more interpretable by another. This implies that interpretability is amenable to personalization (potentially through data-driven A/B testing, when user feedback is available) at a much more extensive level than fairness is. *Lastly*, it may be argued that different application domains of AI have different priorities between interpretability and fairness. Interpretability may be considered very important in fields involving high-bandwidth human-technology interaction such as robotics and HCI. On the other hand, fairness may be of prime importance in societally relevant applications such as automation of applicant screening for jobs, policing and automated decision making on benefit and healthcare eligibility.

The distinctions discussed above do not just apply to fairness vs. interpretability. While other theories from ethical philosophy share many features with fairness as discussed above, dimensions such as privacy, trustworthiness and transparency share several characteristics with interpretability.

3 Frontiers of Synergy

We develop two frontiers of synergy for technical efforts towards interpretability and fairness. Specifically, the directions we outline require that technological building blocks situated within either of the two (interpretability and fairness) work together to achieve meaningful advances within ethical AI. We neither target nor accomplish a conceptual unification of the concepts of interpretability and fairness.

'Justice must not only be done, but must be seen to be done' - Chief Justice Hewart (UK High Court, 1924)

4 Interpretability for Fairness

Motivation: Consider a decision that is output by a system that is purportedly *'fair'*. How can somebody at the receiving end of a decision from such a system be sure that it is fair, according to the notion of fairness used in the system? This question may be instantiated based on the specific notion of fairness used, as follows. For a system that claims to use Rawlsian fairness, how can we ascertain that the present outcome is reasonable to ensure that the system functions in a way that is most advantageous to the least disadvantaged (Ref. *difference principle* [Rawls, 1971]). Or in the case of *demographic parity* as fairness, how can one be sure that the decision is a natural fallout of a process that is designed to achieve parity across sensitive groups. Particularly, users are likely to ask these questions when they find themselves at the receiving end of what they perceive as a *bad* decision (e.g., denial of welfare support). In the best interests of explainability and in making sure that *fairness is perceived to be ensured* (alluding to the 1924 quote above), we (i) ought to answer these *'how is this fair?'* questions, and (ii) do so without reference to other decisions for other specific individuals (for privacy).

Background: Having motivated that some sort of fairness explanations are necessary, we now consider ethical theories in the space. *Accountability for Reasonableness (AFR)* [Daniels and Sabin, 2008] is an ethical framework designed for healthcare scenarios, in particular, when fairness is to be accounted for in scenarios involving allocation of scarce healthcare resources. Of particular interest to us is AFR's *relevance* condition that suggests that decisions are explained by appealing to rationales that are reasonable enough to be accepted by *fair-minded* people who are disposed to finding justifiable terms of co-operation. Badano [Badano, 2018] generalizes this to require acceptance by *each* reasonable person (aka *full acceptability*), which implicitly requires that those subject to most adverse decisions also be convinced. Recently, Wong [Wong, 2019] has argued that AFR could provide directions towards addressing the political dimensions of algorithmic fairness.

Interpretability for Fairness: We propose a novel framework, *Interpretability for Fairness (IFF)*, drawing inspiration from AFR. IFF blends AFR with the *design for values* (Ref. [Dignum, 2019] Sec 4.4) approach to formulate a set of principles targeted at using interpretability as a pathway to enhance acceptability of fair AI. The two IFF principles are:

- **Fairness Publicity Condition:** The fairness value(s) that are sought to be achieved by the AI system must be laid out clearly in layman-friendly language as comprehensively as possible. If a trade-off between values is intended (as often sought, such as a balance between utilitarianism and demographic parity fairness), the relative importance between the values in the mix should be exemplified.
- **Values to Decision Interpretability Condition:** The system should strive to produce a layman-friendly and simple interpretation of each decision substantiating how it relates to the mix of values embodied in the system, as outlined in the publicity condition. This interpretation should be acceptable to any reasonable person who is disposed to finding mutually justifiable terms of co-operation.

These principles are intended to be meaningful to a data scientist equipped with a reasonable understanding of the nuances and social aspects of the domain in which the AI is being designed to operate in, as opposed to abstract ones that pose a *'translation'* challenge. In particular, IFF is at a lower level of abstraction due to instantiating accountability as interpretability and reasonableness as fairness-oriented reasonableness. IFF also keeps process governance aspects (e.g., appeals as in AFR) out of its scope and is focused on the technical design of the AI. While the first IFF condition draws from the analogous condition in AFR, the second condition is inspired by the *'design for values'* [Aldewereld *et al.*, 2015] maxim on linking values to concrete software functionalities. The second condition requires *fairness interpretability*, or explaining a decision based on the fairness values and any trade-offs with non-fairness values used in the system. We note here that the layman-friendliness requirement in the first condition entails exclusion of socio-technical terms such as *structural discrimination* [Pincus, 1999] and *intersectionality* [Carastathis, 2014], whereas that in the second condition would require exclusion of AI-terminology such as *bayesian inference* [Ghosal and Van der Vaart, 2017].

IFF and Current Solutions: IFF requires interpretability that is distinctly different from interpretability in the sense it is used in current Fair ML literature. Within our focus area of unsupervised learning, interpretability is dominated by rule-learning where features for rules are drawn from the data [Balachandran *et al.*, 2012] or auxiliary features [Sambaturu *et al.*, 2020] such as tags. While these enable human-understandable descriptions of the outcomes, they do not satisfy the IFF requirement of explaining the fairness. Recent work on fair representation learning [He *et al.*, 2020] has posited that fairness-based re-engineered versions of original features may be considered interpretable as long as they remain attached to the semantic labels (e.g., *maths marks*, *annual income*) associated with features. However, IFF requires that the *re-engineering process be interpretable on the basis of fairness values*; for example, if the re-engineering transforms annual income for an individual from \$40k to \$35k prior to further downstream processing, IFF requires that this transformation be explained as a fallout of a reasonable fairness-seeking process. We haven't come across Fair ML work that may be argued to be conformant to IFF.

IFF-aligned AI: IFF is construed, much like AFR, as a set of guiding principles, and is not prescriptive as to *how* conformance may be achieved. By way of an illustrative example, a feature re-engineering method that *corrects the dimensions of achievement of socially discriminated demographics* (e.g., *backward castes*) *upward* may be acceptable by reasonable people as a fallout of a process targeting demographic parity, as long as the re-engineering process is interpretable. This is similar in spirit to differentiated age and attempt limits as enforced in India's affirmative action process².

Why IFF? As indicated in the motivation, IFF has the potential to deepen user confidence and trust in fairness-seeking algorithms through explicitly illustrating the conformance to fairness. A detailed treatment of the various aspects of the two IFF principles is not feasible due to space constraints.

5 Fairness and Interpretations

Motivation: Consider using fair AI *along with* state-of-the-art approaches for interpretability that explain the outcomes using data or auxiliary information (not the IFF fairness interpretability). The absence of a connection between the two may lead to a dissonance between their outputs, especially for fair AI that operationalizes notions of group fairness. The fairness-agnostic search for user-friendly explanations could lead to manifestly unimpressive explanations. For example, the membership of an individual within a group could be motivated by demographic fairness considerations, but an explanation referring to a sensitive attribute such as ethnicity may be unacceptable, since that may be perceived as demeaning individuality. On the other hand, excluding sensitive attributes from the realm of explanations may lead to lower quality (thus, unacceptable) explanations. Either of the above could undermine user's trust in the AI, given recent research linking trust and explanation quality [Kunkel *et al.*, 2019].

Fairness and Interpretations: Towards addressing the above conundrum, we propose a layered paradigm, called

Fairness and Interpretations (F&I). *First*, we propose that an AI system be constrained to conform to **both**: (i) fairness, in accordance with the fairness values it targets, and (ii) reasonable interpretability, i.e., be able to provide reasonable explanations for its decisions. This rules out fair AI that is not reasonably interpretable, and vice versa. There could be several ways of characterizing reasonableness in explanations. One way would be to characterize reasonableness as individual fairness; in other words, the space of reasonable explanations may be characterized as being a space where the associated outcomes are *smooth*. Others ways could be to choose explanation paradigms (of which there are many [Binns *et al.*, 2018]) that enhance user's perception of system fairness (an empirical study appears at [Dodge *et al.*, 2019]). *Second*, for scenarios where both fairness and reasonable interpretability are hard to satisfy together (as could be the case where the data has high degrees of bias), we propose that the user be informed that no interpretable explanation can be supplied for the decision, and that the result be accompanied by *fairness explanations* as outlined in IFF. Thus, an F&I-compliant AI is one that adheres to *fairness* (as designed for), and provides either *reasonable decision interpretability* or *only fairness explanations*. Additionally, the two F&I principles are ordered lexically; unless there is a good reason that the first principle cannot be satisfied, the second does not come into play. The appeals process for such a system needs to be cognizant of whether the decision was accompanied by *reasonable explanations* or just *fairness explanations*; a higher appeal rate for the latter may be expected and planned for.

Why F&I? F&I is intended as a paradigm that will meaningfully bring both *fairness* and *interpretability* together without artifacts of one dampening the other. While a detailed analysis of F&I consequences is infeasible here due to space, F&I-aligned AI is expected to be able to fuse normative and user-oriented aspects towards enhancing ethical AI. F&I is well-aligned with and intends to further operationalization of the *right to explanations* enshrined within GDPR [Selbst and Powles, 2018] recommendations of the European Union.

Technical Challenges: In contrast to IFF, F&I lies in a technically pristine space and entails crisp technical challenges. First, the notion of *reasonable explanations* needs to be technically instantiated, and computational approaches to determine reasonableness effectively and efficiently needs to be developed. Second, the twin constraints (fairness and reasonable explanations) would need to be achieved together, requiring novel multi-criteria optimization methods. Third, a decision procedure to determine when to fall back to IFF explanations needs to be developed. While these may sound simple to state, domain-specific nuances would entail different domain-specific technical pathways for achieving F&I.

6 Concluding Notes

We considered the distinctions between fairness and interpretability, and outlined two principles-based frameworks that entail technical challenges where fairness and interpretability can meaningfully work together. We hope that these will contribute to deepening the scholarly debate towards enhancing ethical AI in meaningful ways.

²e.g., [https://en.wikipedia.org/wiki/Civil.Services.Examination_\(India\)#Age](https://en.wikipedia.org/wiki/Civil.Services.Examination_(India)#Age)

References

- [Aldewereld *et al.*, 2015] Huib Aldewereld, Virginia Dignum, and Yao-hua Tan. Design for values in software development. *Handbook of Ethics, Values, and Technological Design: Sources, theory, values and application domains*, Berlin: Springer, 2015.
- [Badano, 2018] Gabriele Badano. If you're a rawlsian, how come you're so close to utilitarianism and intuitionism? a critique of daniels's accountability for reasonableness. *Health Care Analysis*, 26(1):1–16, 2018.
- [Balachandran *et al.*, 2012] Vipin Balachandran, P Deepak, and Deepak Khemani. Interpretable and reconfigurable clustering of document datasets by deriving word-based rules. *Knowledge and information systems*, 32(3):475–503, 2012.
- [Binns *et al.*, 2018] Reuben Binns, Max Van Kleek, Michael Veale, Ulrik Lyngs, Jun Zhao, and Nigel Shadbolt. 'it's reducing a human being to a percentage' perceptions of justice in algorithmic decisions. In *Proceedings of the 2018 Chi conference on human factors in computing systems*, pages 1–14, 2018.
- [Carastathis, 2014] Anna Carastathis. The concept of intersectionality in feminist theory. *Philosophy Compass*, 9(5):304–314, 2014.
- [Chouldechova and Roth, 2020] Alexandra Chouldechova and Aaron Roth. A snapshot of the frontiers of fairness in machine learning. *Communications of the ACM*, 63(5):82–89, 2020.
- [Daniels and Sabin, 2008] Norman Daniels and James E Sabin. Accountability for reasonableness: an update. *Bmj*, 337, 2008.
- [Dignum, 2019] Virginia Dignum. *Responsible Artificial Intelligence: How to Develop and Use AI in a Responsible Way*. Springer Nature, 2019.
- [Dodge *et al.*, 2019] Jonathan Dodge, Q Vera Liao, Yunfeng Zhang, Rachel KE Bellamy, and Casey Dugan. Explaining models: an empirical study of how explanations impact fairness judgment. In *Proceedings of the 24th International Conference on Intelligent User Interfaces*, pages 275–285, 2019.
- [Ghosal and Van der Vaart, 2017] Subhashis Ghosal and Aad Van der Vaart. *Fundamentals of nonparametric Bayesian inference*, volume 44. Cambridge University Press, 2017.
- [He *et al.*, 2020] Yuzi He, Keith Burghardt, and Kristina Lerman. A geometric solution to fair representations. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, pages 279–285, 2020.
- [Kleinberg and Mullainathan, 2019] Jon Kleinberg and Sendhil Mullainathan. Simplicity creates inequity: implications for fairness, stereotypes, and interpretability. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pages 807–808, 2019.
- [Kunkel *et al.*, 2019] Johannes Kunkel, Tim Donkers, Lisa Michael, Catalin-Mihai Barbu, and Jürgen Ziegler. Let me explain: Impact of personal and impersonal explanations on trust in recommender systems. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, pages 1–12, 2019.
- [Mohassel and Zhang, 2017] Payman Mohassel and Yupeng Zhang. Secureml: A system for scalable privacy-preserving machine learning. In *2017 IEEE Symposium on Security and Privacy (SP)*, pages 19–38. IEEE, 2017.
- [Molnar, 2020] Christoph Molnar. *Interpretable Machine Learning*. Lulu. com, 2020.
- [Narayanan, 2018] Arvind Narayanan. 21 fairness definitions and their politics. In *Proc. Conf. Fairness Accountability Transp.*, New York, USA, volume 1170, 2018.
- [Pincus, 1999] Fred L Pincus. From individual to structural discrimination. *Race and ethnic conflict: Contending views on prejudice, discrimination, and ethnoviolence*, pages 120–124, 1999.
- [Rawls, 1971] John Rawls. *A theory of justice*. Harvard university press, 1971.
- [Sambaturu *et al.*, 2020] Prathyush Sambaturu, Aparna Gupta, Ian Davidson, SS Ravi, Anil Vullikanti, and Andrew Warren. Efficient algorithms for generating provably near-optimal cluster descriptors for explainability. In *AAAI*, pages 1636–1643, 2020.
- [Selbst and Powles, 2018] Andrew Selbst and Julia Powles. “meaningful information” and the right to explanation. In *Conference on Fairness, Accountability and Transparency*, pages 48–48. PMLR, 2018.
- [Toreini *et al.*, 2020] Ehsan Toreini, Mhairi Aitken, Kovila Coopamootoo, Karen Elliott, Carlos Gonzalez Zelaya, and Aad van Moorsel. The relationship between trust in ai and trustworthy machine learning technologies. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, pages 272–283, 2020.
- [Wong, 2019] Pak-Hang Wong. Democratizing algorithmic fairness. *Philosophy & Technology*, pages 1–20, 2019.