

The GDPR's Version of Algorithmic Accountability

Author : Margot Kaminski

Date : August 16, 2018

Lilian Edwards and Michael Veale, *Slave to the Algorithm? Why a 'Right to an Explanation' is Probably Not the Remedy You are Looking For*, 16 **Duke L. & Tech. Rev.** 18 (2017), available at [SSRN](#).

Scholarship on whether and how to regulate algorithmic decision-making has been proliferating. It addresses how to prevent, or at least mitigate, error, [bias and discrimination](#), and [unfairness](#) in algorithmic decisions with significant impacts on individuals. In the United States, this conversation largely takes place in a policy vacuum. There is no [federal agency for algorithms](#). There is no algorithmic due process—no notice and opportunity to be heard—not for [government decisions](#), nor for [private companies](#)'. There are—as of yet—no required [algorithmic impact assessments](#) (though there are some [transparency requirements](#) for government use). All we have is a tentative piece of proposed legislation, the [FUTURE of AI Act](#), that would—gasp!—establish a committee to write a report to the Secretary of Commerce.

Europe, however, is a different story. The General Data Protection Regulation (GDPR) went into direct effect on EU Member States on May 25, 2018. It contains a hotly debated provision, [Article 22](#), that may impose a version of due process on algorithmic decisions that have significant effects on individuals. For those looking to understand how the GDPR impacts algorithms, I recommend Lilian Edwards' and Michael Veale's *Slave to the Algorithm? Why a 'Right to an Explanation' is Probably Not the Remedy You are Looking For*. Edwards and Veale have written the near-comprehensive guide to how EU data protection law might affect algorithmic quality and accountability, beyond individualized due process. For U.S. scholars writing in this area, this article is a must-read.

Discussions of algorithmic accountability in the GDPR have, apart from this piece, largely been limited to the debate over [whether](#) or [not](#) there is an individual “right to an explanation” of an algorithmic decision. Article 22 of the GDPR places restrictions on companies that employ algorithms without human intervention to make decisions with significant effects on individuals. Companies can deploy such algorithmic decision-making only under certain circumstances (when necessary for contract or subject to explicit consent), and even then only if they adopt “suitable measures to safeguard the data subject’s rights and freedoms and legitimate interests.” These “suitable measures” include “at least the right to obtain human intervention on the part of the controller, to express his or her point of view and to contest the decision.” They also arguably include a right to obtain an explanation of a particular individualized decision. The debate over this right to an explanation centers on the fact that it appears in a [Recital](#) (which, in brief, serves as interpretative guidance), and not in the GDPR’s actual text. The latest [interpretative document](#) on the GDPR appears to agree with scholars who [argue](#) that a right to an explanation does exist, because it is necessary for individuals to contest algorithmic decisions. This suggests that the right to explanation will be oriented towards individuals, and making algorithmic decisions understandable by (or [legible to](#)) an individual person.

Edwards and Veale move beyond all of this. They do engage with the debate about the right to an explanation, pointing out both potential loopholes and the limitations of individualized transparency. They helpfully add to the conversation about the kinds of explanations that could be provided: (A) model-centric explanations that disclose, for example, the family of model, input data, performance metrics, and how the model was tested; and (B) subject-centric explanations that disclose, for example, not just [counterfactuals](#) (what would I have to do differently to change the decision?) but the characteristics of others similarly classified, and the confidence the system has in a particular individual outcome. But they worry that an individualized right to an explanation would in practice prove to be a “transparency fallacy”—giving a false sense of individual control over complex and far-reaching systems. They valuably add that the GDPR contains a far broader toolkit for getting at many of the potential problems with algorithmic decision-making. Edwards and Veale

observe that the tools of omnibus data protection law—which the U.S. lacks—are tools that can also work in practice to govern algorithms.

First, they point out that the GDPR consists of far more than Article 22 and [related transparency rights](#). This is an important point to make to a U.S. audience, which might otherwise come away from the right to explanation debate believing that in the absence of a right to an explanation, algorithmic decision-making won't be governed by the GDPR. That conclusion would be wrong. Edwards and Veale point out that the GDPR contains other individual rights—such as the right to erasure, and the right to data portability—that will affect data quality and allow individuals to contest their inclusion in profiling systems, including ones that give rise to algorithmic decision-making. (I was surprised, given concerns over algorithmic error, that they did not also discuss the GDPR's related [right to rectification](#)—the right to correct data held on an individual—which has been included in calls for algorithmic due process by U.S. scholars such as [Citron & Pasquale](#) and [Crawford & Schultz](#).) These individual rights potentially give individuals control over their data, and provide transparency into profiling systems beyond an overview of how a particular decision was reached. But there remains the question of whether individuals will invoke these rights.

Edwards and Veale identify that the GDPR goes beyond individual rights to “provide a societal framework for better privacy practices and design.” For example, the GDPR requires something like privacy by design ([data protection by design and by default](#)), requiring companies to build data protection principles, such as data minimization and purpose specification, into developing technologies. For high-risk processing, including algorithmic decision-making, the GDPR requires companies to perform (non-public) [impact assessments](#). And the GDPR includes a system for formal co-regulation, nudging companies towards [codes of conduct](#) and [certification](#) mechanisms. All of these provisions will potentially influence design and best practices in algorithmic decision-making. Edwards and Veale argue that these provisions—aimed at building better systems at the onset, and providing ongoing oversight over systems once deployed—are better suited to governing algorithms than a system of individual rights.

Edwards and Veale are not GDPR apologists. They recognize significant limitations in the law, including the lack of a true class-action mechanism, even where the GDPR contemplates [third-party actions](#) by NGOs. They acknowledge that data-protection authorities are often woefully underfunded and understaffed. And, [like others](#), they point out mismatches between the GDPR's language and current technological and social practices—asking, for example, whether behavioral advertising constitutes an algorithmic “decision.” But they helpfully move the conversation about algorithmic accountability away from the “right to an explanation” and towards the broader regulatory toolkit of the GDPR.

Where the piece falters most is in its almost offhand dismissal of individualized transparency. Some form of transparency will be necessary for the regulatory system that they describe to work—a complex co-regulatory system involving impact assessments, codes of conduct, and self-certification. Without public oversight of some kind, that system may be subject to capture, or at least devoid of important feedback from both civil society and public experts. And, as the ongoing [conversation](#) about [justifiability](#) shows, both the legitimizing and the dignitary value of individualized decisional transparency cannot be dismissed so lightly.

I wish this piece had a different title. In dismissing the value of an individual right to explanation, the title obscures the valuable work Edwards and Veale do in charting other regulatory approaches in the GDPR. However the right to an explanation debate plays out, they show that unlike in the United States, algorithmic decision-making is in the regulatory crosshairs in the EU.

Cite as: Margot Kaminski, *The GDPR's Version of Algorithmic Accountability*, JOTWELL (August 16, 2018) (reviewing Lilian Edwards and Michael Veale, *Slave to the Algorithm? Why a 'Right to an Explanation' is Probably Not the Remedy You are Looking For*, 16 **Duke L. & Tech. Rev.** 18 (2017), available at SSRN), <https://cyber.jotwell.com/the-gdprs-version-of-algorithmic-accountability/>.