

A Consumers' Guide to How to Question an Algorithmic Decision

Last year, a bank rejected my mortgage application. This came as a surprise to me, because I had thought my credit rating was good, and I was borrowing less than 70% of the value of the property. I asked for an explanation, and the answer was even more of a surprise. I was told that the blocker was that I was currently defaulting on an existing mortgage with them. This was a surprise because I didn't have any mortgages with that bank! The story ended well. My data was corrected, and the bank issued me a mortgage. I am now the proud owner of a new home.

AI and algorithms are increasingly making decisions about our lives. When you do a Google search, an algorithm chooses which web pages to list. Netflix and Amazon use AIs to suggest TV series, movies, books, and products for you to view or purchase. While these decisions have only a minor impact upon us, algorithms are increasingly making decisions that can significantly affect our quality of life. Banks use AIs to decide whether to accept a loan application and what interest rate to charge. Courts use algorithms to predict the likelihood that a person will re-offend, which in turn influences the sentencing decision. Healthcare providers use algorithms to prioritize patients for treatment protocols and care programs.

The Emerging Importance of AI Disclosure Practices

With so much at stake, it can be frustrating when a computer decision about you doesn't feel right. Regulators are increasingly responding to the evolving needs of consumers. In Europe, the [General Data Protection Regulation](#) (GDPR) gives the right to appeal an algorithmic decision. Many countries have data protection regulations that give consumers the right to check and correct data that is held about them. In the USA, the proposed [Algorithmic Accountability Act of 2019](#) implies a right for consumers to access the results of the system, and correct or appeal decisions.

Over the past 12 months, governments around the world have been publishing guidelines for AI ethics principles such as the [OECD Principles on Artificial Intelligence](#). The ideas expressed in these documents are converging; there is a growing consensus around the world about the principles of ethical AI. Among these key ethical principles is the importance of disclosure to stakeholders, providing human-friendly explanations for decisions that affect them, especially when an adverse decision is made about them.

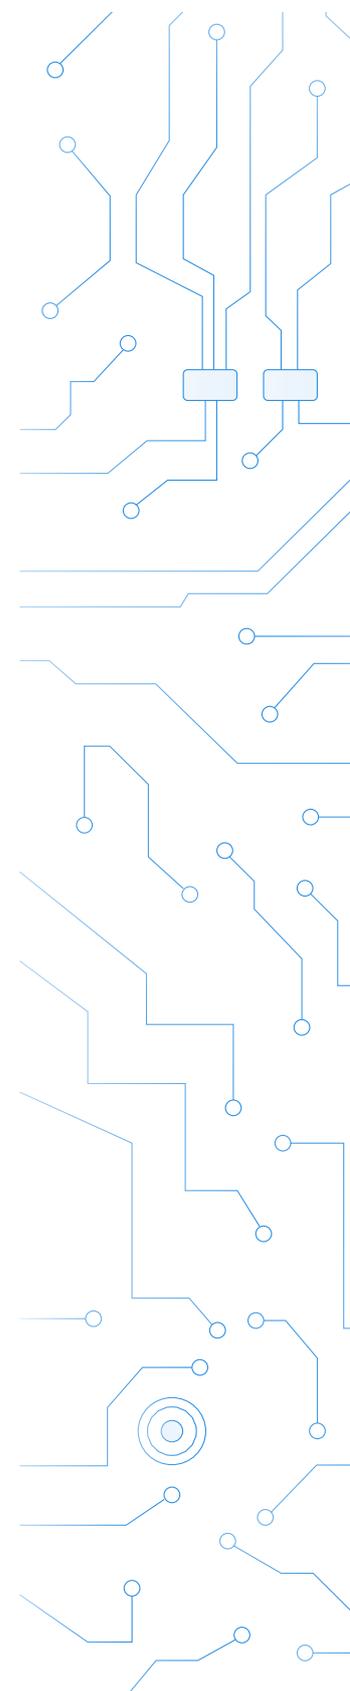
It can be good business for organizations to disclose reasons for algorithmic decisions. Published research shows that human-friendly explanations and disclosures build trust with customers, improving brand value. Human-friendly explanations for algorithmic decisions can help staff to identify and remove unfair bias. The explanation details provided to customers empowers them to detect and correct data errors, reversing incorrect decisions. Best practice disclosure can also inform customers when to change their behavior to earn favorable decisions in future.

Four Simple Questions That Consumers Can Ask

Consumers are increasingly seeking to engage with organizations using AI, to better understand how and why AIs have made decisions about them. However, with AI use becoming widespread only relatively recently, many consumers may not be sure what questions to ask. So, we've created a step by step guide for you.

QUESTION 1: Do you provide the option for a human to review my case?

Check whether the organization has an appeals process, or the ability to opt out of an AI decision and have a human review your case. Under some circumstances, an organization may have a regulatory obligation to provide this process, such as those that must comply with the [GDPR right to appeal or explanation](#) for an algorithmic decision. Some organizations with higher customer service standards offer this option even when not required by law.



**QUESTION 2: May I see the data values you hold about me that were used to make the decision?**

This question aims to check whether the organization is using correct and appropriate data to make decisions about you, and then make corrections as necessary. Many countries have data protection regulations that require organizations to give consumers access to data held about them and the right to have that data corrected. In addition to regulatory requirements, it is best practice to provide these rights. For example, International Standard [ISO/IEC 27701](#) includes the requirement, "The organization should implement policies, procedures and/or mechanisms to meet their obligations to PII principals to access, correct and/or erase their PII."

QUESTION 3: What is different about the data you hold about me that caused you to make a different decision for me versus others?

Ask this question to determine whether the logic of the algorithmic decision is reasonable. Even if the data held about you is correct, an algorithmic decision may be incorrect if it is inconsistent with regulatory requirements for fairness, or if it is inconsistent with common sense. For example, if the answer is that you were rejected for a loan because of your gender, or because you had a higher credit rating than most loan applicants, then you may conclude that the decision was not reasonable. The latest generation of AI systems provides human-friendly explanations, such as those based upon SHAP or DataRobot's [XEMP](#) prediction explanations, that a consumer can easily understand.

QUESTION 4: What could I change to obtain a favorable decision, and how much change is required?

The purpose of this question is to determine whether you can change a decision by changing your behavior. This approach to explaining algorithmic decisions is known as [counterfactual explanations](#) and makes for a more productive experience than a simple rejection. For example, just because a person is currently scored as a poor credit risk, doesn't mean they will always be. By paying down credit card balances, and saving money, they may improve their credit score, and their next loan application may prove successful.

Conclusion

It is good practice to ask questions whenever an AI makes a decision that doesn't seem correct. Check that the data is accurate and that the logic used by the algorithm is sensible. It's best practice for organizations to answer your questions, and the technology exists for them to quickly and easily provide answers. Note, though, that while organizations can answer your key questions, they aren't likely to tell you every detail of how an algorithm works. They still need to protect their intellectual property and protect against malicious operators trying to attack their systems.

Best practice for questioning an algorithmic decision follows the same approach to questioning a human decision: ask for a second opinion, check the data, check the reasoning, and find ways to improve the outcome next time.

**About the Author**

Colin Priest is VP, AI Strategy for DataRobot, where he advises businesses on how to build business cases and successfully manage data science projects. Colin has held several CEO and general management roles, where he has championed data science initiatives in financial services, healthcare, security, oil and gas, government, and marketing.

With a passion for how humans and AIs can work together, Colin has contributed to projects on AI ethics, AI governance, and the future of work. Colin is a strong supporter of the science of healthcare and does pro bono work on cancer research.

Contact Us

DataRobot
225 Franklin Street, 13th Floor
Boston, MA 02110, USA

www.datarobot.com
info@datarobot.com



© 2020 DataRobot, Inc. All rights reserved. DataRobot and the DataRobot logo are trademarks of DataRobot, Inc. All other marks are trademarks or registered trademarks of their respective holders.