

Anti-bias test implemented by Amazon

A test for detecting bias in AI and machine learning systems, developed by researchers at the Turing, has been adopted by Amazon for its cloud computing platform, Amazon Web Services (AWS).

The test, called 'Conditional Demographic Disparity', was first proposed in a [2020 paper](#) by Sandra Wachter, Brent Mittelstadt (both Turing Fellows at the time) and Chris Russell (former Group Leader in Safe and Ethical AI at the Turing). It is a metric that gives a measure of inequality within a dataset, and so can flag up discrimination in, for example, job recruitment processes, automated loan approval, healthcare access and university admissions. The strength of the test is that it incorporates the standards

of fairness used in European courts of law, and it also accounts for underlying factors which might be driving the bias, making it useful for detecting 'intersectional' discrimination where multiple factors are at play.

Amazon has included the test as part of its SageMaker Clarify service, which provides machine learning developers using AWS with tools to detect and measure biases in their datasets and models, helping them to understand their models' predictions and pinpoint issues of inequality. It's a major success for the researchers behind this test, as their work is now in the hands of those who are developing the AI systems of the future.



“The paper that proposes this test is a delight to read, and it clearly lays out the legal and ethical foundations of the work. Machine learning researchers using Amazon Web Services are now using this test to help them identify bias in their datasets and model predictions.”

Sanjiv Das

Amazon Scholar at AWS and Terry Professor of Finance
and Data Science at Santa Clara University