August 2017

What Can You Do About Algorithmic Bias?

Megan E. Garcia

Last edited on June 13, 2019 at 9:45 a.m. EDT

Acknowledgments

About the Author(s)

Megan Garcia is senior fellow and director of growth for New America's National Network, exploring how the work of New America's programs links to exciting local approaches to problem solving taking place in cities around the country.

About New America

We are dedicated to renewing America by continuing the quest to realize our nation's highest ideals, honestly confronting the challenges caused by rapid technological and social change, and seizing the opportunities those changes create.

About New America CA

New America CA promotes efforts that are locally grown and grounded in economic equity, in which technology, innovation and compelling storytelling yield transformative solutions for our most marginalized community members.

Contents

Introduction

Tay's first words were "helloooooooo world!!!" It was a friendly start for the Twitter bot designed by Microsoft to engage with people aged 18 to 24. But, in a mere 12 hours, Tay went from friendly Twitter persona to foul-mouthed, racist Holocaust denier who said feminists "should all die and burn in hell" and that the actor "ricky gervais learned totalitarianism from adolf hitler, the inventor of atheism."

Tay, which Microsoft quickly shut down after just 24 hours, was programmed to learn from the behaviors of other Twitter users, and in that regard, was a success. The bot's embrace of humanity's worst attributes is an example of **algorithmic bias—when seemingly innocuous programming takes on the prejudices either of its creators or the data it is fed**.

The side effects of unintentionally discriminatory algorithms can be dramatic and harmful. Companies and government institutions that use data need to pay attention to the unconscious and institutional biases that seep into their results. **It doesn't take active prejudice** to produce skewed results in web searches, data-driven home loan decisions, or facial recognition software. It just takes distorted data that no one notices and corrects for.

As we begin to create artificial intelligence (AI), **we risk inserting racism and other prejudices into the code that will make decisions for years to come.**

At *Slant: Understanding Algorithmic Bias*, in San Francisco, New America CA brought together **a curated group of experts in AI, bias, technology, and future thinking to outline the state AI and ethics** from four different perspectives. Our goals were understand what specific actions companies are taking to address bias in AI and machine learning and what help they can use from civil society. **We're bringing you our most actionable insights.**

To learn more about these recommendations and this work at New America, please contact Megan Garcia (**garcia@newamerica.org**).

Big ideas to make everyone better at heading off bias in algorithms and machine learning

- An understanding of the need for **thoughtful design** of algorithms and AI would go a long way toward tackling many of the ethics problems we see with both.

- Inside companies and civil society there is the appetite to make algorithms and AI more ethical, but no consensus about how to do that. Efforts are also disconnected and we are not always learning from each other. **FATML** is an exception.

- There will have to be a change in the way companies think about bias to ensure that bias in algorithms and AI is minimized. Eventually there will have to be an inclusion mindset (or some company-wide or industry-wide focus on ethics and inclusion) that forces individuals and processes to incorporate thinking about bias into correcting for biases as they emerge.

- **Journalism** and other fields with regulated or unregulated codes of ethics offer lessons that the technology sector might apply.

- There is dire need for organizations outside of the technology sector to provide ideas about how to address algorithmic bias. Some organizations have made modest progress in this area to train people to use data ethically, but much remains to be done.

- Who to talk to for more: **New America** and **New America CA**, **Center for Democracy and Technology (CDT)**, **Open Data Institute**, **Data & Society Research Institute**, **DataKind**, **ACLU**, **FATML** (Fairness, Accountability and Transparency in Machine Learning) community, **Open AI**, **Partnership on AI**

What's Your Role?

Select your role from the ones below and see what you can do to address bias in AI and machine learning.

- **CEO**

- **Product Manager**

- **Writing Code**

- **Designer**

- **Talent**

- **Lawyer**

- **Policy**

Recommendations for CEOs

CEO,

Mindset

Create an inclusion mindset at your company. Instead of thinking about diversity and inclusion as a distinct function or part of your company, integrate it into your culture so that it's a part of the way your company works and the products or services you produce.

Instead of having ethics or fairness relegated to one team or department, encourage everyone working for your business to take responsibility for fairness. **Research shows** that, increasingly, instilling this kind of culture will help build strong relationships with employees and customers.

Today

Consider developing operating principles to address bias. Many tech leaders report that they don't take action to address implicit bias that is likely to be transferred to code because they worry that taking any action on bias will be criticized. One option to get started is to create operating principles on race and bias that outline how everyone in the company will make decisions about reducing bias and provide common language to use. One example is Code2040's **Principles on Race**.

Think about involving professional societies for engineering, computer science, and other related fields in developing an ethics code for your company. The Institute of Electrical and Electronics Engineers (IEEE) is starting to work on **AI ethics** and may want to help.

Longer-term

Be transparent about the goal of making AI and products that use AI as ethical as possible. Don't keep it to yourself. Let your customers, suppliers and

partners know that your goal is reduce the bias in your AI. Technology trends indicate that creating ethical AI may become a strong selling point in the future.

Recommendations for Coders

Writing Code,

Mindset

Recognize that **you can't resolve fairness just by looking at observational criteria.** Get curious about your data.

Today

Think about the data you're feeding algorithms and ways the data itself might be biased. How might you use less biased data or account for those biases?

Longer-term

Work with social scientists to look at assumptions about data over time. Gather those inside your company to figure out how to use less biased data for product creation.

Recommendations for Designers

Designer,

Mindset

As the creative person in the room, your questions about **who is engaged in the design process (and who isn't)** can make the product you're working on less biased and more effective.

Today

Consider having a conversation among cross-disciplinary product thinkers about machine learning and bias. Get the designers, coders, and product managers you work with thinking about how to design for a broader set of users. We know that including diverse voices will decrease the likelihood that the AI is biased and could increase the value of the product.

Longer-term

Think about how you might build healthy and sustainable collaborations with AI, including product narratives that focus more on the 50th use than the first.

Recommendations for Talent Managers

Work in Talent,

Mindset

Create an inclusion mindset at your company. Instead of thinking about diversity and inclusion as a distinct function or part of your company, integrate it into your culture so that it's a part of the way your company works and the products or services you produce.

Instead of having ethics or fairness relegated to one team or department, encourage everyone working for your business to take responsibility for fairness. **Research shows** that, increasingly, instilling this kind of culture will help build strong relationships with employees and customers.

Today

Educate recruiters on the role they play and hold them accountable. Encourage recruiters to make ethical design of AI and machine learning a recruitment selling point. Measure how effectively they are sourcing diverse talent for your product teams. We know that having diverse teams work on AI will decrease the likelihood that the AI is biased and could increase the value of the product.

Longer-term

Urge your company's leadership to develop operating principles on bias (for example, Code2040 recently developed such **operating principles**) in an effort to move beyond the fear that taking any action on bias will be criticized.

Recommendations for Lawyers

Lawyer,

Mindset

How might legal code and computer code interact to yield products that are more ethical and less biased?

Today

Think about how you might communicate to computer scientists and engineers that "public data" should not necessarily be freely used for machine learning. What questions might they ask before they use data to ensure they account for biases or minimize them?

Longer-term

Think about how to define practices for data collection so that the data coders feed to algorithms is as unbiased as possible. And think about whether bias auditing makes sense. **FairML** is doing bias auditing on a small scale right now.

Recommendations for Policymakers

Work on Policy,

Mindset

We are at a pivotal moment in the development of AI and machine learning in which we can either create systems that are as unbiased as possible or continue to ignore the biases we accidentally embed in technology. Some regions, like the European Union, are already beginning to regulate the ethics of AI through vehicles like the General Data Protection Regulation (GDPR). The policies we put in place now inside companies could have far-reaching impacts.

Today

Think about how you can create policies inside your company that emphasize the importance of minimizing bias in algorithms and machine learning.

Longer-term

Start algorithmic auditing. We can't assume that companies will be able to notice and root out their own bias. Instead consider inviting **outside auditors** to establish ways to reduce bias at your company.

Notes

NEW AMERICA