

Data Bias: The Hidden Risk of AI

What it is and what you need to be doing about it

RESEARCH STUDY

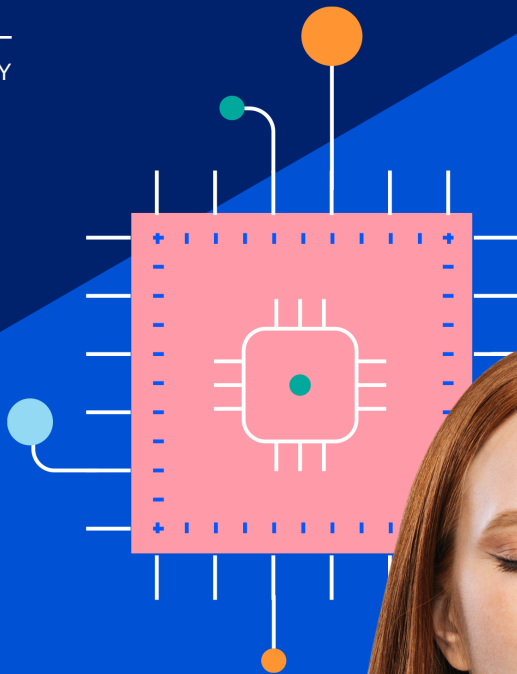


Table of Contents

Introduction / 3

Methodology / 3

Study Highlights / 4

Background / 4

Study Findings / 7

Conclusion / 12



Introduction

When it comes to artificial intelligence (AI) and machine learning (ML), the algorithms are only as good as the data used to create them. If data sets are flawed—or worse, biased—incorrect assumptions will be baked into every resulting decision. Sound farfetched? It's not. Instances of data bias are well documented, but how big is the problem? And what, if anything, are businesses doing to correct it?

Those are the questions Progress sought to answer in an extensive study of AI and ML data bias. As a company focused on helping its customers make smart use of data to drive business outcomes, bias is a key consideration that can't be ignored. Companies looking to implement a sound digital ethics strategy must first understand and scope the size of the AI data bias issue. Only then will they be equipped to aggressively weed it out. This study can be an essential wakeup call for businesses still struggling to understand the issue and the associated risks of letting it fester unaddressed.



Methodology

Progress commissioned a worldwide study from Insight Avenue, a research firm in the United Kingdom specializing in business-to-business technology research. The team conducted 640 interviews with business and IT professionals, director level and above, who “use data to make decisions and are using or plan to use AI and ML to support their decision making.” Respondents from 12 countries participated, hailing from Europe, Asia, Latin America and the United States. The sample specifically surveyed organizations with over 500 employees to better understand the overall awareness of data bias, how it was impacting businesses and what companies were doing to address it.



Study Highlights

66%

of organizations anticipate becoming more reliant on AI/ML decision making

78%

believe data bias will become a bigger concern as AI/ML use increases

65%

of respondents believe there is currently data bias in their organization

77%

believe they need to be doing more to address data bias

13%

are currently addressing data bias and have an ongoing evaluation process to weed it out

51%


consider lack of awareness and understating of biases as a barrier to addressing it



Background

What is Data Bias?

Stories about data bias are all around us, but what is it? In short, data bias happens when human creators inject decision biases into systems by training them using biased data, or with rules that intrinsically reflect their personal biases. Biased data can include flawed data sets, blind assumptions, and models that wrongly discriminate against under-represented groups.



AI/ML vs. Business Rules

Data bias can enter anywhere. While this survey focuses on artificial intelligence and machine learning, those coding business rules must be equally vigilant. While we haven't expressly referenced business rules throughout, these data bias learnings apply equally to decision logic.

When assessing data bias, organizations should think expansively because unconscious bias can creep in at any point during the AI and app development lifecycle. Key tech and process vulnerabilities include:

- Data sets used to create and test analytical models or used by business logic that are not truly representative.
- Algorithms or business logic developed, trained or tested with flawed logic or assumptions which are exacerbated by algorithms lacking transparency.
- Human or automated applications developed with flawed logic or assumptions.
- Human assessment of analytics that take a limited or biased view.
- Human or automated testing protocols that are not appropriately inclusive.
- Model drift due to unfactored changes that occur, such as with data, logic, environment/usage patterns, etc.

But there are other bias areas as well. For instance, user experiences that don't meet W3C accessibility standards can be considered a form of bias against those with disabilities. Combining process and technical considerations only scratches the surface for rooting out bias. Organizations must also begin rethinking hiring practices and how their teams are constructed.

Cultural Homogeneity and AI

According to the Harvard Business Review¹, most AI development is happening by a small pool of extremely technical people with advanced PhDs. Most of those working on these advanced systems are largely male and culturally homogenous, leaving the data vulnerable to inaccuracies based on the limited worldview, similar training and non-diverse perspectives of those designing it.

¹Kriti Sharma, "Can We Keep Our Biases from Creeping into AI?," Harvard Business Review, February 9, 2018, <https://hbr.org/2018/02/can-we-keep-our-biases-from-creeping-into-ai>.

Examples of Data Bias: Business Risk

Business practices based on biased AI data can have severe consequences for those negatively affected.

- **Retail:** One famous retailer found a flawed hiring algorithm exclusively put forth men for open technology roles, excluding otherwise qualified women candidates.
- **Finance:** A financial institution found itself wrongly rejecting qualified loan candidates because of a flawed AI tool discriminating by applicant zip code.
- **Healthcare:** A company using AI to assign healthcare eligibility wrongly assigned lower health risk status to Black patients, denying them the proper care they were entitled to, and leading to adverse medical outcomes.

Left unchecked, biased AI has the capacity to do serious business harm. But how much is really understood about it? And are businesses ready to take action?





Study Findings


Concern Growing About Artificial Intelligence (AI) Bias

A full 66% of respondents anticipated their reliance on AI would grow in the next five years, though 78% believed data bias would become a larger problem as reliance on AI/ML increases. This interest was consistent across regions, which left us wondering how concerned these businesses were about current levels of data bias in their organizations?

While 57% indicated their business decision makers were concerned about data bias, 65% believed there was currently data bias in their organization. The types of decisions impacted included:


- 44% Finance
- 40% IT/Digital
- 33% Operations
- 32% Customer Acquisition

“Significant” concerns centered on the business impact of biased decisions which included both poor customer experience and potential legal exposure. Moderate concerns included unfair stereotyping and negative impacts to inclusion and diversity efforts.



Data Bias Maturity Varies Across Businesses

Given these concerns, it was interesting to note that an overwhelming majority of business respondents were only just beginning to understand and address data bias. The analysis categorized their maturity into four groupings:



- Novice (6%):** Had yet to begin investigating data bias
- Emerging (36%):** Began learning about and starting to address data bias
- Improving (45%):** Started to implement tools, training and policy to address data bias
- Advanced (13%):** Committed to constantly evaluating how they use tools, training and culture to address data bias


For those working to combat data bias, effective measures were found to include education and training; improved transparency and traceability of algorithms and data; more time spent model training, building and evaluation; and using tools to help source bias within data sets. Despite some successes, 77% said their organizations still needed to do more to understand data bias. While the effective measures they identified primarily included skills, practices and training, 65% saw technology/tools as the most urgent need to better combat data bias. This was followed by more training (59%) and adjusting their strategy/vision (49%).

Baby Steps: Addressing Data Bias

Ironically many organizations enter their remediation efforts at a structural disadvantage. This is because they had yet to address bias in other areas in their organizations, despite having set plans to do so. These sub-optimal practices included:

- Inaccessible web, mobile and desktop experiences
- Digital design limitations that impact user experience
- Non-diverse hiring, team and leadership dynamics
- Non-inclusive usability and testing efforts

Yet, the maturity data shows positive steps were being taken. Sixty-seven (67%) believed their organization had evaluated technology to tackle data bias, 40% said data bias was a consideration when evaluating AI/ML vendors, and a full 76% recognized that data bias was best tackled centrally across the organization rather than relying on siloed departments to handle. Based on the study's findings, it will likely take a comprehensive approach that combines people, tools, training and ongoing policy vigilance to ensure data bias is meted out within AI/ML practices.



Leading the Change

When asked which leader was best equipped to own company initiatives designed to tackle data bias, opinions varied, but a clear majority named the Chief Information Officer/Chief Technology Officer:

39%

Chief Information Officer/Chief Technology Officer

16%

Chief Data Officer

13%

Chief Operations Officer

9%

Chief Human Resource Officer/Human Resource Director



Barriers to Addressing Data Bias

Slow progress is understandable when one considers the significant barriers businesses must overcome to make gains. These included:

- Lack of awareness of potential biases (51%)
- Lack of understanding about how to identify bias (43%)
- Lack of expert resources, such as having access to data scientists (31%)

Transparency: The Antidote to Bias

For AI to deliver unbiased results, it needs unbiased data. Achieving this requires an agile, transparent, rules-based data platform where data can be ingested, harmonised and curated for the AI tool. Data lineage features allow the human expert to track any changes made to the data, including back to the moment humans introduced bias.

Only 9% mentioned not seeing data bias as an issue, meaning inaction is more due to planning and execution than a failure to recognize the imminent threat bias presents.

But the longer data bias goes ignored, the more risks respondents recognized as potentially damaging to their businesses. In priority order, these included:

- Security and governance
- Unreliable decisions that impact revenue
- Reputational damage
- Lost customer trust
- Ethical exposure from automated decision making

There was remarkable consistency across regions in terms of the priorities of perceived risk, though our Advanced maturity respondents were the only ones to elevate “lost financial opportunities” and “unreliable business decisions” ahead of governance. One might conclude that their deeper level of experience with AI data bias has exposed this mature group to the deeper business implications of letting bias go unchecked.




Benefits to Tackling AI Data Bias

Respondents saw great opportunities in working to address data bias. This included lessening risk, making better decisions, advancing market opportunities, becoming an employer of choice for data scientists, improving company reputation and not ironically—gaining peace of mind. At a larger level, 76% believed there were wider societal impacts at play if enterprises were collectively unable to adequately address data bias.

Broad-Based Industry Priorities by Region

The complex issue of AI data bias will require a more broad-based approach to resolve than one organization could possibly tackle alone. Here, the desired marketplace way forward varied a bit by region. Latin American respondents sought increased investment in multidisciplinary research into bias. APAC and the US respondents favored more exploration into how humans and machines can best work together, while European respondents slightly favored more collaboration and best-practice sharing within industries.





Conclusion

Driving Essential Change Forward

As use of AI grows and companies aspire for sustainable value, use of artificial intelligence and machine learning will only increase. More data scientists, line-of-business practitioners and programmers will dive into datasets and produce ever-more algorithms. The challenge then becomes carefully considering all aspects of a project to avoid the consequences of unconscious bias, which includes:

- **Hiring and team diversity:** For AI to be sustainable over time, the pool of those developing these algorithms must expand. This not only needs to include those across the racial and gender spectrums, but should also include those with less-advanced degrees and those who hail from a broader cross-section of professions and backgrounds.
- **Training:** Cultural training on inclusiveness isn't enough. Companies must provide technical training on dataset management so expert practitioners can develop protocols to detect, remediate and avoid creating biased algorithms.
- **Technology:** Every touchpoint within the entire tech or development stack and process must consider and factor in the reality of data bias. This includes all aspects of data selection and preparation, business logic development and analytical models, testing and results analysis. Only a continuous commitment to assessment and removal will ensure bias doesn't seep in over time.

Every day, bias impacts day-to-day business, from security and governance and bad business decisions, to lost customer trust and potential legal and ethical exposure. And these risks don't even begin to cover the profound consequences experienced by victims of data bias, including those who suffer adverse outcomes resulting from intrinsically biased AI algorithms.

Eliminating AI data bias will take a combination of technology, training and practices to prevent it from entering the development process. But as our world grows more dependent on machines to make vital decisions that impact lives, it's up to those leading these efforts to ensure their work is a force for good.








About Progress

Dedicated to propelling business forward in a technology-driven world, [Progress](#) (NASDAQ: PRGS) helps businesses drive faster cycles of innovation, fuel momentum and accelerate their path to success. As the trusted provider of the best products to develop, deploy and manage high-impact applications, Progress enables customers to build the applications and experiences they need, deploy where and how they want and manage it all safely and securely. Hundreds of thousands of enterprises, including 1,700 software companies and 3.5 million developers, depend on Progress to achieve their goals—with confidence. Learn more at www.progress.com

2023 Progress Software Corporation and/or its subsidiaries or affiliates. All rights reserved.
Rev 2023/03 RITM0188507

Worldwide Headquarters

Progress Software Corporation
15 Wayside Rd, Suite 400, Burlington, MA 01803, USA
Tel: +1-800-477-6473

-  facebook.com/progresssw
-  twitter.com/progresssw
-  youtube.com/progresssw
-  linkedin.com/company/progress-software
-  [progress_sw_](https://instagram.com/progress_sw_)