

TREDENCE

— Beyond Possible —

THE MORAL AND MARKET ARGUMENT FOR **RESPONSIBLE AI**

Why Setting Standards and Using
Frameworks Is in Businesses' Best Interests



Ravindra Patil

Senior Director, Data Science,
Tredence Inc.



Table of Contents

Why Responsible AI Is More Important Than Ever	03
Challenges Implementing Responsible AI	05
Benefits of Operating a Best-in-Class Responsible AI Program	07
Critical Enablers of a Responsible AI Program	08
Why Work with Tredence	09
Enhance AI Responsibility. Talk to Tredence.	11

WHY RESPONSIBLE AI IS MORE IMPORTANT THAN EVER

Enterprises were making slow but steady progress with artificial intelligence (AI) and machine learning (ML), developing analytics, automating processes, and guiding decision-making. Then, OpenAI debuted ChatGPT in November 2022, ushering in a new era of generative AI. Within weeks, most enterprises were formally committed to exploring this exciting new technology, which can learn, understand, and create and promises to transform businesses everywhere.



Generative AI Is Galvanizing Growth

In 2022, companies had an average of 3.8 AI models in production.¹

Today, seven in 10 companies are experimenting with generative AI,² meaning that the number of AI models in production will skyrocket over the coming years.

Industry discussions around responsible AI have thus taken on greater urgency. Business leaders and teams understand that when they design and deploy new AI capabilities, they need to consider not only the business gains they'll create but also the societal impact of new AI-powered processes and outcomes.

Responsible AI is a governance framework that covers ethical, legal, safety, privacy, and accountability concerns.³

VentureBeat

Increasingly, AI will infuse experiences, not just solutions. That makes it more challenging to evaluate how humans and technology interact across channels and time. This paper shares why responsible AI is important, advocating that organizations implement frameworks and mature processes and sharing best practices to accomplish this critical goal.

¹ Michael Chui, Bryce Hall, Helen Mayhew, Alex Singla, and Alex Sukharevsky, The state of AI in 2022 – and a half decade in review, report, McKinsey, December 6, 2022, <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai-in-2022-and-a-half-decade-in-review>

² Carl Franzen, "More than 70% of companies are experimenting with generative AI, but few are willing to commit more funding," article, VentureBeat, July 25, 2023, <https://venturebeat.com/ai/more-than-70-of-companies-are-experimenting-with-generative-ai-but-few-are-willing-to-commit-more-spending/>

³ David Ellison, "Responsible AI will give you a competitive advantage," article, VentureBeat, January 15, 2022, <https://venturebeat.com/ai/responsible-ai-will-give-you-a-competitive-advantage/>



Questions about AI fairness abounded before the advent of generative AI. High-profile AI failures include facial recognition software that couldn't successfully detect women with darker skin, provoking claims of skin type and gender bias.⁴ A consumer technology giant's new credit card offered smaller credit lines to women than men,⁵ raising questions of gender bias. And healthcare patient data was shared with a hyperscaler without proper guardrails, breaching patient privacy.⁶ Unsurprisingly, consumers are becoming more suspicious of AI. In a recent survey, 60% of consumers said they were concerned about how organizations use AI today, while two-thirds (65%) have already lost trust in organizations due to their AI practices.⁷

Solving the Responsible AI Disconnect



More than 1 in 2

Have experienced a responsible AI failure.



Just 1 in 10

Are actively mitigating AI model risks across their company.⁸

Generative AI makes developing horizontal and vertical AI solutions even easier, unlocking real value for enterprises. However, generative AI can also create fake identities, imagery, and content; hallucinate outcomes; and demonstrate bias. As a result, enterprises will need to work even harder to set up and maintain guardrails to ensure new solutions are fair and non-discriminatory, transparent, explainable, and accountable. In addition, they will need to ensure data privacy and security and maintain human control over automated decisions. Ensuring the responsible deployment of AI solutions will become increasingly challenging as transformation speeds up and enterprise teams seek to get new solutions to market before competitors.

The good news is that there is widespread acceptance of the need for responsible AI. More than half of organizations already champion AI ethics. However, only one in five have implemented comprehensive programs with frameworks, governance, and guardrails to oversee AI model development and proactively identify and mitigate risks, as indicated by the graphic below.⁹

⁴ Larry Hordesty, "Study finds gender and skin-type bias in commercial artificial-intelligence systems," MITNews, February 11, 2018, <https://news.mit.edu/2018/study-finds-gender-skin-type-bias-artificial-intelligence-systems-0218>

⁵ Will Knight, "The Apple Card Didn't 'See' Gender—and That's the Problem," article, Wired, November 19, 2019, <https://www.wired.com/story/the-apple-card-didnt-see-gender-and-thats-the-problem/>

⁶ Natasha Lomas, "Google confirms it's pulling the plug on Streams, its UK clinician support app," TechCrunch, August 26, 2022, <https://techcrunch.com/2022/08/26/google-confirms-its-pulling-the-plug-on-streams-its-uk-clinician-support-app/>

⁷ James Coker, "#DataPrivacyWeek: Consumers Already Concerned About AI's Impact on Data Privacy," InfoWeek, January 24, 2023, <https://www.infosecurity-magazine.com/news/consumers-ai-impact-data-privacy/>

⁸ Jim Boehm, Liz Grennan, Alex Singla, and Kate Smaje, "Why digital trust truly matters," report, McKinsey, September 12, 2022, <https://www.mckinsey.com/capabilities/quantumblack/our-insights/why-digital-trust-truly-matters>


⁹ AI ethics in action, report, IBM, page 6, <https://www.ibm.com/downloads/cas/4DPJK92W>

Organizations Have Good Intentions with AI Ethics, But Governance Lags

A recent study by IBM found that only around 20% of organizations have broad-based responsible AI programs set up to govern AI development and ensure models are transparent, free from bias, and protect consumers' sensitive data.¹⁰

The Intention-action gap

Organizations are endorsing AI ethics principles—but are still catching up on implementing them



CHALLENGES IMPLEMENTING RESPONSIBLE AI

Many enterprises are struggling to source responsible AI talent, train staff on new processes, and ensure governance across fast-growing AI programs. In addition, responsible AI spans up to 20 different business functions, increasing process and decision-making complexity.¹¹

¹⁰ AI ethics in action, report, IBM, page 6, <https://www.ibm.com/downloads/cas/4DPJK92W>
¹¹ IBM, ibid.

Responsible AI teams must work with key stakeholders, including leadership; business owners; and data, AI, and IT teams to:



Build AI solutions that are fair and free from bias

AI solutions should be fair and free from bias. AI can reflect bias in multiple ways, including historical biases against individuals or groups based on race, ethnicity, gender, sexual identity, religion, disability, and other factors. However, unintentional biases can also be introduced when systems are trained on data that doesn't effectively represent key populations, when data is preprocessed incorrectly, or algorithms weigh certain features more heavily than others.

Teams and partners can use different techniques, such as exploratory data analysis, to identify and mitigate potential biases before developing solutions—that way, models are built with fairness in mind from the start.

Teams and partners can also review the data used in preprocessing, algorithm design, and postprocessing to ensure that it is representative and balanced. In addition, they can use group and individual fairness techniques to ensure that algorithms treat different groups and individuals fairly. And counterfactual fairness approaches model outcomes if certain factors are changed, helping identify and address biases.



Promote AI transparency and explainability

AI transparency means it is easy to understand how AI models work and make decisions. Explainability means these decisions can be easily communicated to others in non-technical terms. Using common terminology, holding regular discussions with stakeholders, and creating a culture of AI awareness and continuous learning can help achieve these goals.



Ensure data privacy and security

AI models use mountains of data. Companies are leveraging first- and third-party data to feed models. They also use privacy-preserving learning techniques, such as creating synthetic data to overcome sparsity issues. Leaders and teams will want to review and evolve data privacy and security safeguards to ensure that confidential and sensitive data is still protected as it is used in new ways. For example, synthetic data should emulate customers' key characteristics but not be traceable back to individuals.



Implement governance

Enterprise experience with AI ranges from those just piloting solutions to those who are scaling them enterprise-wide. Governance will vary based on corporate AI maturity. However, companies should set AI principles and policies from the start. As their AI model use increases, they can appoint AI officers; implement frameworks; create accountability and reporting mechanisms; and develop feedback loops and continuous improvement programs.

BENEFITS OF OPERATING A BEST-IN-CLASS RESPONSIBLE AI PROGRAM

Many companies have transitioned responsible AI from IT to the business, indicating that they now view this issue as an organizational imperative. More than one in four CEOs (28%) oversee their companies' AI ethics programs, with the board (10%), the general counsel (10%), the privacy officer (8%), and the risk and compliance officer (6%) also involved.¹²

As a result, these companies are dedicating time and investment to building and strengthening their programs. Here are some benefits of maturing responsible AI programs that accrue to companies and teams:



Communicating openly and transparently about AI

Responsible AI programs provide the terminology to discuss developing AI solutions that benefit businesses and society. Leaders feel comfortable sharing metrics-driven commitments and progress publicly. Teams have the definitions, guidelines, and processes to develop and maintain AI solutions, considering responsibility throughout models' lifecycles.



Ensuring governance and accountability

Leaders who take ownership of responsible AI are more likely to champion the need to mature this important capability. They'll serve as highly visible sponsors of the program; provide investment for staffing, training, and tools; and oversee efforts to implement governance and adopt responsible AI workflows. They'll also maintain funding for these programs during times of economic austerity, believing that using responsible AI as a source of strength fosters product innovation, reduces risks, and provides other benefits that exceed the rewards of short-term cost cutting. And if failures occur, leaders and teams share lessons learned and improve processes to strengthen programs rather than shirking from responsibility.



Proactively mitigating risks

By mandating responsible AI policies and processes, enterprises can mitigate many AI model development risks. They can understand how data choices create or harm fairness and provide robust oversight of all development and maintenance processes. One easy win is to improve model monitoring to prevent drift, make ongoing corrections, and keep trust in predictions high.



Maintaining a positive brand in the marketplace

Enterprises that adopt responsible AI frameworks, tools, and metrics are accountable for AI solutions and outcomes, communicate decisions, and reduce risks before they impact employees, customers, and members of the public. By so doing, they increase trust in their AI solutions.

¹² IBM, *ibid.*

Abiding by all relevant regulations



Enterprises already abide by data privacy and security regulations that govern their industries and regions. More than 60 countries have developed AI policies.¹³ The US and European Union are working on AI regulations, while China has already passed an AI regulation governing online recommendation systems.¹⁴ By developing responsible AI programs now, companies can gain experience with frameworks and processes, preparing for an era of increasing regulation globally.

Being a top partner and employer in the market



In the U.S., 76% of employees want to work for a company that positively impacts the world. Talented staff feels similarly around the world.

Creating and publicizing a responsible AI program enables human resources teams to recruit top talent, like data scientists and engineers. It also helps engage and retain workers who want to on programs that deliver societal good.

CRITICAL ENABLERS OF A RESPONSIBLE AI PROGRAM

So, what differentiates companies that are responsible AI leaders from others? They:



Create a vision and goals for AI

Leaders communicate their vision and goals for AI and how it benefits the company, customers, and society.



Set expectations

Senior leaders set the right expectations with teams to build responsible AI solutions from the ground up rather than trying to tailor solutions after they're completed.



Implement a framework and processes

Partners can provide responsible AI frameworks with transparent processes and guardrails. For example, data privacy, fairness, and bias checks should be built into initial data preparation, model development, and ongoing monitoring.



Ensure team adoption and accountability

Teams are trained on new policies and processes. Leaders may include responsible AI behaviors in teams' personal development plans. In addition, enterprises audit teams for compliance with key policies.



Use the right metrics to quantify results

Leaders and teams use benchmarks and other metrics to demonstrate how responsible AI contributes business value to keep stakeholder engagement high.



Monitor AI systems

Partners can provide model monitoring services, solving problems proactively and ensuring they deliver trusted results.

WHY WORK WITH TREDENCE

Tredence solves last-mile challenges with data, analytics, and AI, closing the gap between creating insights and unlocking their total value. We provide the expertise, services, tooling, and partnerships enterprises need to succeed with responsible AI. Tredence:



Provides full-spectrum AI skills

Tredence offers deep expertise across the entire AI and generative AI value chain, including large-language model (LLM) engineering, development, operations, and platform engineering capabilities.



Offers domain and industry-specific expertise

Tredence combines data and AI strategy-setting and execution with customer analytics, marketing technology, supply chain, and other capabilities. We leverage these domain and industry insights to help enterprises deploy new AI solutions that will provide significant value.



Provides a responsible AI framework

Tredence develops and deploys responsible AI frameworks at pace, enabling companies to rapidly improve maturity without hindering innovation.




Offers access to our ATOM.AI ecosystem

ATOM.AI, our intelligently engineered accelerator ecosystem, reduces development time for generative AI pilot projects, responsibly. We leverage partnerships with hyperscalers to develop vertical solutions that speed time to value and speed to scale.



Responsible AI



ENHANCE AI RESPONSIBILITY. TALK TO TREDENCE.

If your company is accelerating AI innovation, you likely need a responsible AI program. Move proactively to reduce risks, mature programs, and processes, and demonstrate accountability to stakeholders.

By partnering with Tredence, you gain access to the skill sets, frameworks, tools, and partnerships you need to unlock business value with responsible AI. Deploy models that are fair and free from bias, enforce controls, and increase compliance with company requirements while preparing for forthcoming regulations.

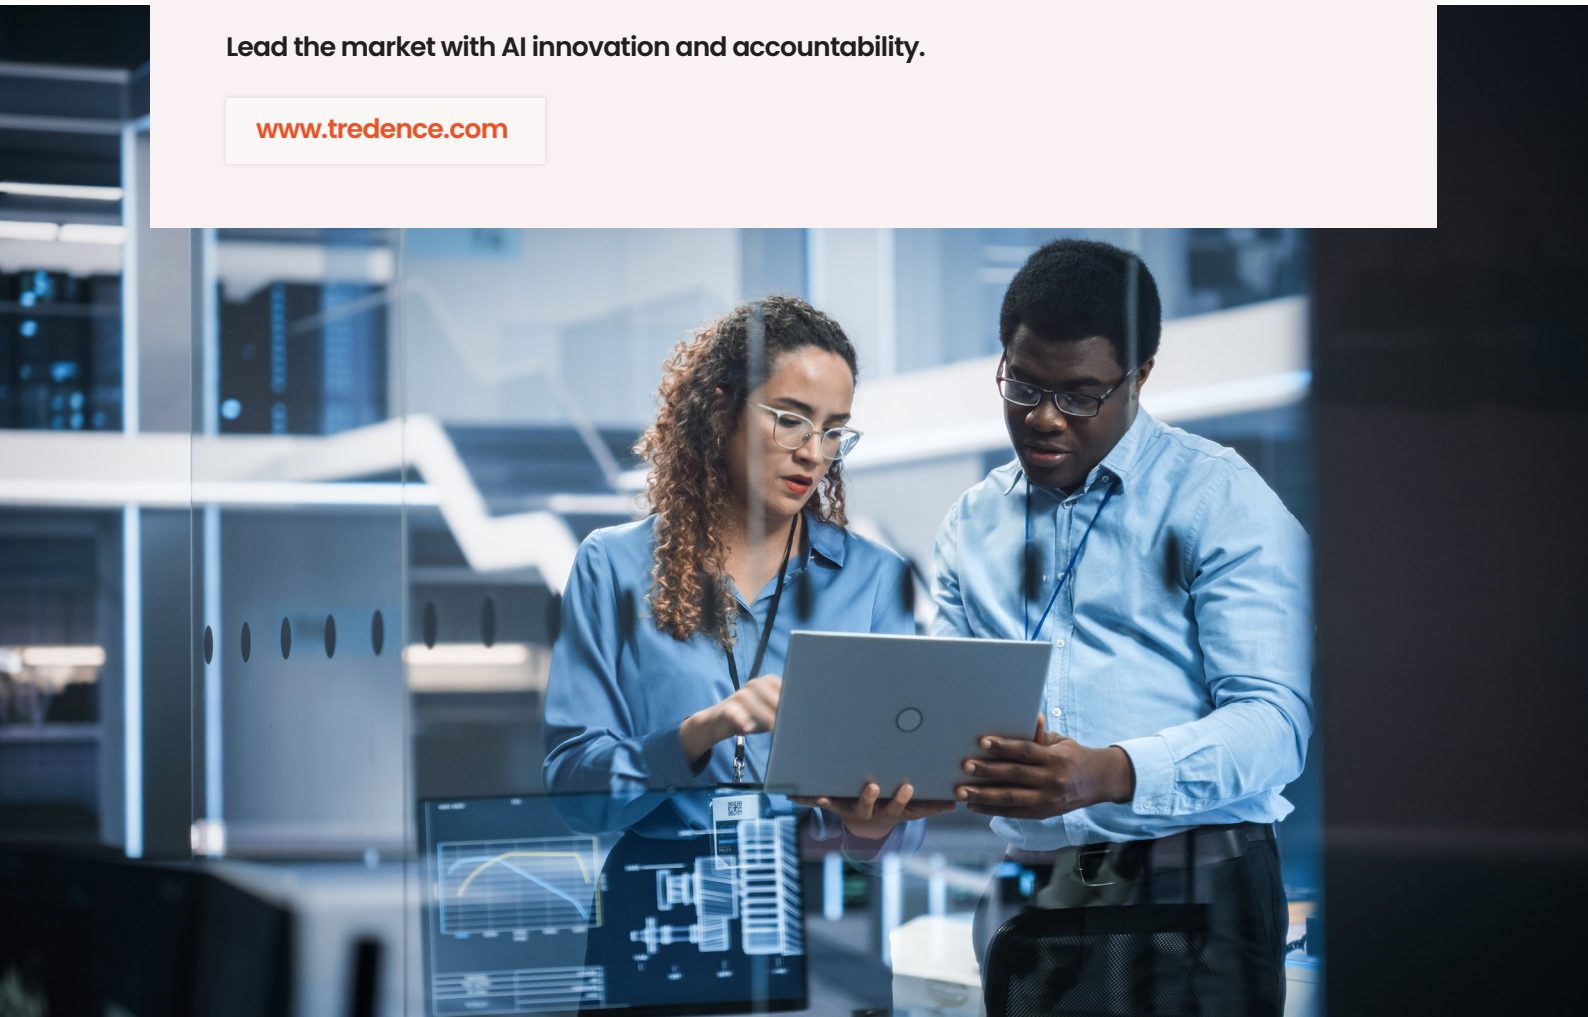
Seize This No-Risk Offer

Ready to get started?

Ready to get started? Contact us to schedule a 60-minute discovery call, where we'll learn about your needs, discuss how to implement a responsible AI framework, and provide a roadmap for moving forward.

Lead the market with AI innovation and accountability.

www.tredence.com



ABOUT THE AUTHOR



Ravindra Patil

Senior Director, Data Science,
Tredence Inc.

As the Senior Director of Data Science at Tredence Inc., Ravindra specializes in creating Data and AI solutions that solve complex customer challenges. With advanced degrees, including a Ph.D. from the University of Maastricht and a Master's from IIT Madras, he has contributed significantly to the field, filing multiple patents and publishing over 25 research papers. A sought-after speaker in Data and AI circles, Ravindra excels in machine learning, deep learning, and computer vision, continually pushing the boundaries of innovation and strategy.



Want to know more about us?

Please visit: www.tredence.com

Follow us at: [in](#) [twitter](#) [youtube](#) [facebook](#)