accenture

Responsible AI

From principles to practice

Contents

Artificial intelligence (AI) has the potential to add trillions to the global economy. And the benefits it can deliver to the enterprise are well documented. However, Accenture research found that to deliver on the promise of this technology, AI must be scaled across the organization.

Successfully scaling AI for business value requires a number of key factors. These include a focus on value, new-skilling and scaling to production. But often, a stumbling block that organizations struggle to overcome is the significant uncertainty and risks associated with AI. Trust inside and outside of the organization is a key component to getting to value from AI. And to get to trust, organizations must move beyond defining Responsible AI principles and put those principles into practice.

Accenture has worked with organizations worldwide to build this trust. How? By defining and implementing solutions across four Responsible AI pillars—moving from principles to practice. In this report we share what we have learned—from practitioners' pain points and how to address them, to case studies of what good looks like in the real world.

Responsible AI in practice—essential but not easy

A growing imperative

The potential value organizations can achieve through AI is clear. So is the danger of being left behind by competitors if the technology is not leveraged. But many businesses feel overwhelmed trying to determine how to address the risks associated with it. In a global survey of risk managers, 58% identify AI as the biggest potential cause of unintended consequences over the next two years. Only 11% describe themselves as fully capable of assessing the risks associated with adopting AI organization-wide.

One of the main reasons for this crisis of confidence? The way AI shifts the risk landscape in new and unexpected ways for organizations. As AI decisions increasingly influence and impact people's lives at scale, so responsibility on enterprises increases to manage the potential ethical and socio-technical implications of AI adoption.

In particular, bias, discrimination and fairness have emerged as areas of paramount concern, alongside explainability. Recent academic advances in algorithmic fairness focus on a variety of very specific definitions of these problem areas, which can be highly contextual, mutually exclusive and contradictory. Translating these into action is fraught with tough decisions, trade-offs and application-specific constraints, requiring multiple cross-domain perspectives to reach consensus.

Shareholders, regulators, media, employees and the public are increasingly aware of the positive impact that scaling AI can have. But it is also clear that the potential for significant damage exists if Responsible AI isn't included in an organization's approach.

In response, many enterprises have started to act (or in other words, to Professionalize their approach to AI and data). Those that have put in place the right structures from the start, including considering Responsible AI, are able to scale with confidence, achieving nearly three times the return on their AI investments when compared to those that have not.

However, there are also multiple examples of organizations which have taken initial steps to establish AI ethics principles and well-meaning proofs of concept but struggle to scale these concepts into their live processes. Which begs the questions: **why is it so hard to go from principles to practice? What are the common pitfalls and how can we address them?**

Responsible AI

We define Responsible AI as the practice of designing, building and deploying AI in a manner that empowers employees and businesses and fairly impacts customers and society.

Practitioners' insights—the realities of Responsible AI

To better understand why organizations struggle to move from principles to practice, Accenture conducted a global Responsible AI Practitioners' Survey—26 in-depth interviews with Responsible AI practitioners from 19 organizations, across four continents. We spoke with technological practitioners (data scientists and AI engineers), lawyers, industrial/organizational psychologists, and project managers, all of whom are charged with implementing Responsible AI practices in their organizations.

Our analysis indicates that given the embryonic nature of Responsible AI, some organizations have struggled to develop a systematic internal approach to convert their principles into practice. **Our experience shows that this is because they underestimate the technical complexity and scale of people and process change required.**

Approximately 75% of the Responsible AI projects referred to in these interviews were more than a year old. Yet, in almost all cases, efforts had stagnated or remained incomplete due to a range of problems.

In the report, we will analyze the problems we uncovered from the survey results, and provide recommendations for how to move forward.

Moving from principles to practice

The four pillars of Responsible AI

To successfully move from Responsible AI principles to practice, organizations need to tackle a central challenge: translating ethical principles and academic theories like algorithmic fairness into practical, measurable metrics and thresholds that are right for them. To embed these responsible approaches and metrics into everyday practices, organizations will also need to put in place the necessary organizational, technical, operational, and reputational scaffolding.

We've defined four pillars of successful Responsible AI implementations, based on what we've learned delivering Responsible AI solutions to organizations across the globe.

The four pillars of Responsible AI

01
Organizational

Democratize the new way of working and facilitate human + machine collaboration.

02
Operational

Set up governance and systems that will enable AI to flourish.

03
Technical

Help ensure systems and platforms are trustworthy and explainable by design.

04
Reputational

Articulate the Responsible AI mission and ensure it's anchored to your company's values, ethical guardrails, and accountability structure.

In our experience, many organizations begin their journey by focusing on one issue, such as algorithmic fairness or compliance. However, **those that are most successful understand the importance of investing in all four pillars that underpin Responsible AI from the very start.**

Pain points, recommendations and case studies

For each pillar of Responsible AI, we have summarized the key practitioner pain points from our survey, provided recommendations to address these pain points based on our implementation experience, and included a case study to demonstrate what good looks like.

01 | **Organizational**

Practitioner pain points

Practitioner interviewees regularly highlighted the need for appropriate performance metrics and recognition of their Responsible AI work. They indicated that little value was being placed today on risk mitigation, including prevention of reputational harm, while time pressures meant short-term product success was prioritized ahead of the long-term benefits of Responsible AI. This often came from the top, with interviewees highlighting an unwillingness by leadership to recognize Responsible AI failures and deeply engage with issues like lack of fairness, transparency, and the potential for discrimination.

Interviewees felt that as organizations create new Responsible AI duties and roles, a culture of confidence must be built across the enterprise, empowering employees to raise concerns and act ethically through the right combination of responsible product and individual success metrics and incentive structures.

What we recommend

Strong leadership is pivotal in empowering employees, elevating Responsible AI as a key business imperative. To democratize this new way of working, successful organizations recognize the need for new and changing roles, and actively upskill, re-skill, or hire.

In our experience, organizations should actively create and encourage an organizational culture that empowers individuals to raise doubts or concerns with AI systems, without stifling innovation.

Establishing clear success criteria, incentives and training helps to nurture these new roles and skills, diffusing a Responsible AI culture across the organization and cultivating trust in AI systems.

Case study
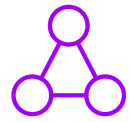
Putting **organizational** principles into practice

Leadership plays an essential role in the successful adoption of Responsible AI. A leading European financial services organization wanted to explore the use of algorithmic fairness across their model inventory. Strong buy-in and support at the executive level enabled a culture of openness and collaboration, where an internal, multidisciplinary team were free to work with Accenture to learn and explore algorithmic fairness, raising any concerns in a safe environment.

A strong focus was initially placed on education and upskilling, ensuring all team members fully understood algorithmic fairness, the choices and trade-offs that need to be made, and how to utilize Accenture's Algorithmic Assessment toolkit.

This foundation of support and knowledge removed many of the fears and orthodoxies that often exist around bias and fairness, creating an open, collaborative culture. Individuals were empowered to ask the hard questions, with potential sources of bias thoroughly investigated. In undertaking this process, the organization now has the training and skills needed to implement the latest advances in the algorithmic fairness space.

02 | **Operational**

Practitioner pain points

Practitioners interviewed as part of our research indicated that companies consistently struggled with stakeholder misalignment, frustrating bureaucracy, conflicting agendas and a lack of clarity on processes or ownership. For example, individuals operated in an ad hoc nature based on their own values and personal assessment of relative importance. A lack of clarity on governance and accountability structures, undue conflict and competing incentives across groups ultimately led to Responsible AI inertia and a reactive mindset.

What we recommend

Effective organizations establish transparent, cross-domain, governance structures. These structures identify roles, expectations and accountability to build internal confidence and trust in AI technologies. Throughout the project lifecycle, these organizations transform ethical principles into clear processes, procedures and chains of command that respect the context-specific needs of each application.

Accenture has found that the creation of a cross-domain ethics committee at an early stage has been invaluable for some organizations. In clearly defining roles and expertise, ways of working and authority to govern, procedures can be maintained on an ongoing basis, while also enabling on-demand responses as issues arise.

Putting operational principles into practice

Following the release of EU AI Trustworthy Guidelines, one global communications vendor asked Accenture to help develop their own internal ethical principles and translate them into operational actions, activities, and structures. This would give them a practical Responsible AI risk management foundation on which to build their future AI pipeline with confidence.

Over the course of seven weeks, Accenture helped to define a governance framework, with four key deliverables:

Principles and playbook	Responsible AI committee	Training	Monitoring
Responsible AI principles defined and detailed in an actionable playbook.	New committee created with clearly defined roles, ways of working and authority to govern.	Personalized training plans built for all team members based on their roles and Responsible AI responsibilities.	Structures and channels established to safely escalate risks and ways to actively foster dissent.

In particular, the playbook enabled our client to navigate the difficult path of principles to practice. For each key activity along a project lifecycle, the playbook makes clear the activities to be completed, the team members required to take action and the stakeholders responsible for successful completion. Each of these activities are aligned with one or more of the client's ethical principles, ensuring that the right principles are being considered and acted on accordingly, at each stage. The committee supports day-to-day activity, resolving any alerts or escalations, while also running check-ins at each phase of the AI lifecycle.

03 | **Technical**

Practitioner pain points

One of the biggest barriers that practitioner's organizations faced was a lack of expertise in defining and measuring ethical use and algorithmic impact of data, models and model outcomes on an ongoing basis. Without established technical methods to identify, mitigate and monitor these risks, organizations cannot be confident that the system is fair and safe.

Interviewees also indicated that their organizations struggled to integrate academic metrics like algorithmic fairness, finding them very different from traditional benchmarks and KPIs. Responsible AI cannot be measured in revenue generation or click-through rates, yet organizations still relied on performance metrics like these to define the success or failure of Responsible AI practitioners.

What we recommend

Successful organizations architect and deploy AI models, systems and platforms that are trustworthy, fair and explainable *by design*. Using proven qualitative and quantitative techniques to assess potential risks, they are better placed to reach cross-domain consensus on mitigation strategies. They also clearly define measurable performance metrics and establish techniques for ongoing monitoring, control and re-assessment.

It is important for organizations to invest time in fully understanding the sources of bias in their various systems. In our experience, this process leads to better informed resolution strategies that match the organization and the application. Having the right set of tools to thoroughly investigate sources of bias and understand the trade-offs and impacts of fairness decisions is invaluable in reaching common cross-domain consensus.

Case study

Putting **technical** principles into practice

Allied Irish Bank (AIB), wanted to ensure they were ahead of the industry and bring their data-science teams rapidly up to speed on the latest developments in algorithmic fairness, while further enhancing the integration of algorithmic fairness assessment in the models used to aid their decision-making.

Working with the bank's data science team, we leveraged Accenture's Algorithmic Assessment toolkit to assess fairness and actions needed to mitigate bias in two new models that were in development.

At relevant points in the model development workflow, analyses were surfaced to a multidisciplinary group. By breaking this complex problem down into manageable, understandable "chunks", the tools gave data scientists and business executives a deeper understanding of their data and model outcomes from a fairness perspective. They then directed further investigation into areas of potential bias. This methodology informed and improved decision-making during the model-build process.

We enabled AIB to integrate a data-driven assessment of the complex problem of algorithmic fairness in the end-to-end model lifecycle. The bank's data science teams are now self-sufficient on the tool, and are able to independently use it in their ongoing work to affirm confidence and a deeper understanding of their models.
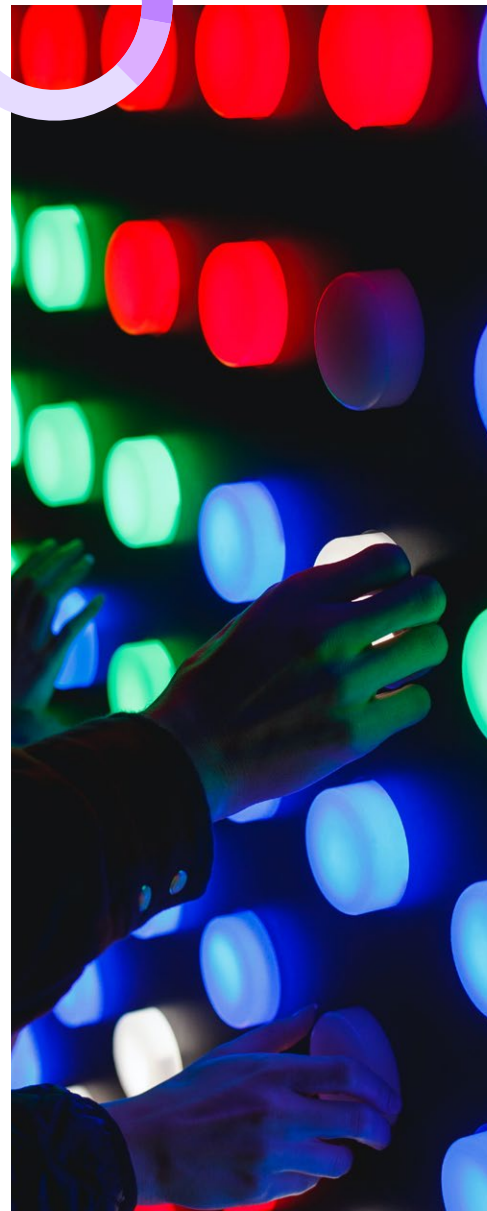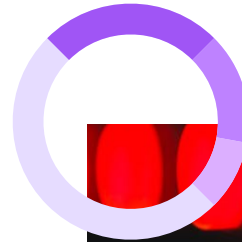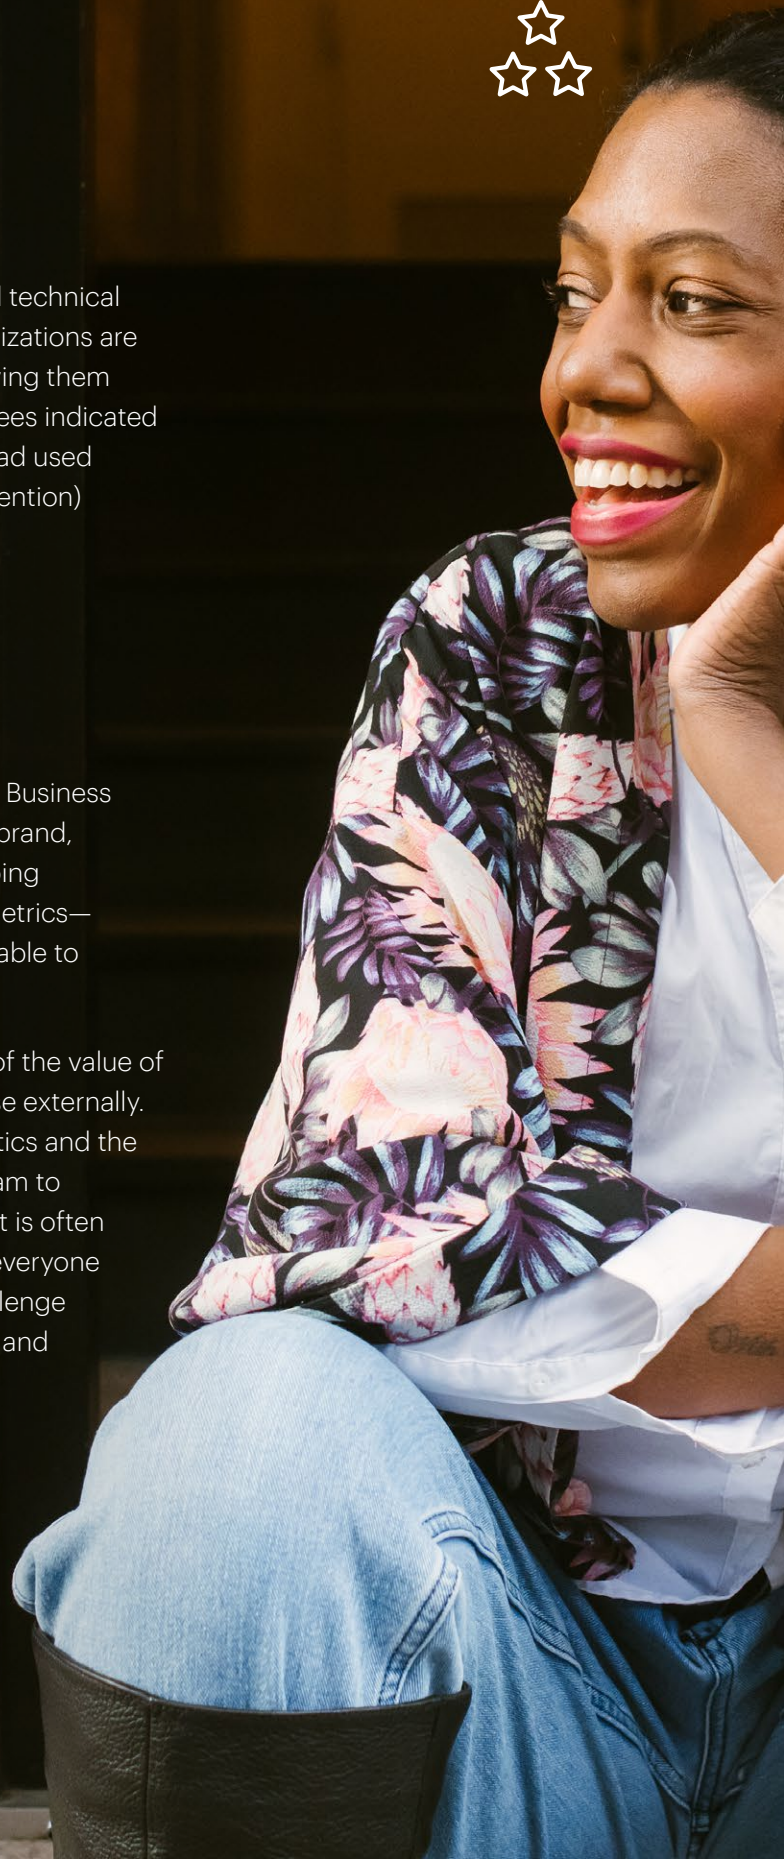
04 | **Reputational**

Practitioner pain points

Without the necessary organizational, operational, and technical foundations in place, interviewees reported that organizations are forced into a reactive approach to Responsible AI, leaving them exposed to significant reputational damage. Interviewees indicated that in the absence of clear legal requirements, they had used reputational risk (in the form of catastrophic media attention) as a way of incentivizing change.

What we recommend

Leaders in this field clearly articulate their Responsible Business mission, anchored in their principles and informed by brand, public risk assessments and guidance. Accurate, ongoing measurement and monitoring of key Responsible AI metrics—such as fairness—ensures they are managing risk and able to communicate with confidence and transparency.

In some cases, internal stakeholders can be skeptical of the value of ethical principles or risk averse in communicating these externally. Successful organizations embrace these internal skeptics and the fresh perspective they bring, encouraging the core team to pressure-test the principles they're defining. The result is often a more considered, precise set of principles in which everyone has confidence. Responsible AI is a cross-domain challenge and consensus is vital to build a culture of confidence and enable trust in the technology.

Case study

Putting **reputational** principles into practice

Establishing a Responsible AI approach that is robust, fair and maintained on an ongoing basis can also enable organizations to communicate and collaborate with confidence.

A major government agency wanted to apply an "ethics-by-design" approach to the algorithms it utilizes to serve citizens. The agency wanted to have a clear view of potential bias present in the data, model, and output to identify improvement points and make sure the algorithms it used were treating all people in a fair way.

Working together, we established an end-to-end approach to Responsible AI, combining our Algorithmic Assessment toolkit with strong governance structures, ensuring continued monitoring and evaluation of fairness over time. A multidisciplinary team was also established to ensure adherence to ethical principles, as well as serving as a mechanism for escalation of ethical issues, should they ever arise. In taking a Responsible AI approach, the agency was also able to establish new ethics roles to create internal and external awareness, foster trust and encourage communication and collaboration.
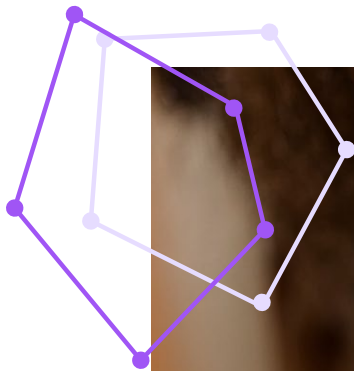
The next step? From practice to proof

The value of AI is clear. But it can bring with it new, dynamic, ethical and social issues. A failure to manage these issues can have a significant impact at a human and societal level, leaving organizations exposed to financial, legal, and reputational repercussions.

While many organizations have taken the first step and defined AI principles, translating these into practice is far from easy, especially with few standards or regulations to guide them. Our global Responsible AI Practitioner's Survey identified a range of organizational, operational, technical and reputational challenges that hold well-intentioned organizations back. While the initial focus is often on ethical and legal requirements, success is also a function of an organization's ability to modify its traditional ways of working to support Responsible AI—and AI more broadly.

In our experience, successful organizations understand the importance of taking a systematic approach from the start, addressing these challenges in parallel, while others underestimate the scale and complexity of change required. A systematic approach requires proven tools, frameworks and methodologies, enabling the organizations to move from principles to practice with confidence and supporting the [professionalization of AI](#).

In undertaking this process organizations also establish the structures needed to demonstrate the long-term value of Responsible AI by scaling it across the organization, enabling the essential move from "practice to proof."

We use a set of 25 questions to help our clients to benchmark their motivators and challenges, together with their maturity in terms of people, process and technology against their peers. Where are you on your Responsible AI journey?

Authors

Ray Eitel-Porter

Managing Director, Applied Intelligence
Global Lead for Responsible AI
ray.eitel-porter@accenture.com

Medb Corcoran

Managing Director, Accenture Labs
Global Responsible AI Lead for Tech Innovation
medb.corcoran@accenture.com

Patrick Connolly

Research Manager
Responsible AI
patrick.connolly@accenture.com

Contributors

Anika Mahajan

Responsible AI – Technical & ASEAN Lead

Bogdana Rakova

Artificial Intelligence Consultant, Responsible AI

About Accenture

Accenture is a global professional services company with leading capabilities in digital, cloud and security. Combining unmatched experience and specialized skills across more than 40 industries, we offer Strategy and Consulting, Interactive, Technology and Operations services—all powered by the world's largest network of Advanced Technology and Intelligent Operations centers. Our 514,000 people deliver on the promise of technology and human ingenuity every day, serving clients in more than 120 countries. We embrace the power of change to create value and shared success for our clients, people, shareholders, partners and communities. Visit us at www.accenture.com.

About Applied Intelligence

Applied Intelligence is Accenture's approach to scaling AI for our clients. We embed AI-powered data, analytics and automation capabilities into business workflows to accelerate time to value. Our expertise in defining end to-end strategy, combined with deep data infrastructure capabilities, cognitive services and industrialized accelerators help smooth clients' path to AI adoption, extending human capabilities and supporting clients in scaling AI responsibly. Recognized as a leader by industry analysts, we collaborate with a powerful global alliance, innovation and delivery network to help clients deploy and scale AI within any market and industry. Visit us at www.accenture.com/appliedintelligence.

Follow @AccentureAI

About Accenture Labs

Accenture Labs incubates and prototypes new concepts through applied R&D projects that are expected to have a significant impact on business and society. Our dedicated team of technologists and researchers work with leaders across the company and external partners to imagine and invent the future. Accenture Labs is located in seven key research hubs around the world: San Francisco, CA; Sophia Antipolis, France; Washington, D.C.; Shenzhen, China; Bangalore, India; Herzliya, Israel and Dublin, Ireland; and 25 Nano Labs The Labs collaborates extensively with Accenture's network of nearly 400 innovation centers, studios and centers of excellence located in 92 cities and 35 countries globally to deliver cutting-edge research, insights and solutions to clients where they operate and live. For more information, please visit www.accenture.com/labs.

About Accenture Research

Accenture Research shapes trends and creates data driven insights about the most pressing issues global organizations face. Combining the power of innovative research techniques with a deep understanding of our clients' industries, our team of 300 researchers and analysts spans 20 countries and publishes hundreds of reports, articles and points of view every year. Our thought-provoking research—supported by proprietary data and partnerships with leading organizations, such as MIT and Harvard—guides our innovations and allows us to transform theories and fresh ideas into real-world solutions for our clients. For more information, visit www.accenture.com/research.