

Paul Daugherty

Chief Technology & Innovation Officer
Accenture

Responsible AI Means Raising It Right



As AI takes a more central role in our lives, it needs to be deployed responsibly—with accountability and transparency, stewardship and security.

Artificial Intelligence is arguably the single most disruptive technology the world has experienced since the Industrial Revolution and is set to have a wide-ranging impact on every part of our lives and society. We are already witnessing how AI applications and insights are affecting how we work and live—helping us drive cars, diagnose medical conditions, and provide financial advice—and these applications are growing exponentially.

Businesses are significantly ramping up investment in AI to capitalize on new growth opportunities and reinvent the way things are done. In fact, according to research conducted by Accenture, 81 percent of executives believe that in the next two years AI will work alongside humans in their organizations. This collaboration between humans and machines will change the nature of work and drive competitive advantage. Consider Accenture modeling that estimates human-machine collaboration could boost revenues for companies by 38 percent between 2018 and 2022.

However, because AI is such a powerful tool, as it takes a more central role in our lives, it needs to be deployed responsibly—with accountability and transparency, stewardship and security—to engender trust and maximize its potential. Potentially negative consequences of the technology are likely to emerge unless appropriately addressed during development. Consider the race to achieve autonomous vehicles, which has provided more than one example of how AI can go terribly wrong. There are other significant issues if AI is not applied responsibly and thoughtfully, with a history of recent incidents including AI chatbots that developed racial and gender bias.

These challenges can and must be addressed. AI is what we make of it, and if we design our AI algorithms to reflect business and societal norms of responsibility, fairness, and transparency, there's no reason we can't enjoy the benefits of the technology with few, or even no, disadvantages.



From Programming to Learning

A responsible approach to AI can be achieved if we “raise” the technology right. We don’t expect our children to act ethically without guidance, so we educate and nurture them. This is exactly the approach we should take with AI. These carefully raised AIs will not only be able to scale operations, but also adapt to new needs via feedback loops from other deployed models—in the same way as continuing education enables employees to adapt to new tasks.

The parallels between human and machine education don’t stop there. As my colleague Jim Wilson and I explain in *Human + Machine: Reimagining Work in the Age of AI*, raising AI to be responsible means addressing many of the challenges we tackle through human development and education, such as fostering an understanding of right and wrong; imparting knowledge without bias; and building self-reliance while emphasizing the importance of collaborating with others. In human development these

skills are acquired in three steps: (1) learning how to learn; (2) acquiring the ability to rationalize and explain thoughts and actions; and (3) accepting responsibility for decisions. Companies should look to apply these same measures of learning and responsibility when developing their AI systems.

Building an AI Curriculum

Based on this learning paradigm, companies can start training their AI. The first place to start is with data—as much of it as possible. Data is to machines what language is to humans: without a full taxonomy of language humans cannot scale their learning; nor can machines without huge volumes of data. Ultimately, businesses with the richest datasets to teach their AI will create the best systems.

Companies must also ensure there is an accepted background of understanding for AI and others it will be communicating with, whether customers, employees, or intelligent systems. It’s here that great care is needed. The wrong taxonomies of training data won’t just affect

scale, they could also unintentionally introduce bias into the system. Data scientists need to be mindful of this eventuality and mitigate it at all costs. Companies that get this right will be able to build the strongest libraries of AI models, ready for reuse.

The Responsible AI Imperative

As AI plays an increasingly important role in our society, there’s a clear moral obligation for businesses to ensure that their systems are raised to be responsible. And this isn’t only a moral imperative; it also makes good business sense.

The leaders of tomorrow will be those companies that tackle the ethical challenges of AI head-on and set the standards for what it means to create a responsible, explainable AI system. As a result, they will build greater levels of trust with customers and employees and will be more likely to win these important audiences over to the dramatic changes that lie ahead.