



**Does your company need
a Chief AI Ethics Officer,
an AI Ethicist, AI Ethics
Council, or all three?**

Positioning your organization
for success on AI Ethics

Authored by Beena Ammanath,
Executive Director, Deloitte AI Institute

The role of AI Ethicist is becoming a hot topic as businesses grow more and more reliant on AI—and as AI systems become increasingly sophisticated and autonomous. Yet too many companies mistakenly believe they need to fill this crucial new role with just one person.

The ethical issues that arise from AI are complex and multi-dimensional. As such, they require expertise and insights across a wide range of disciplines, from technology and regulatory compliance to philosophy, psychology, and sociology—and everything in between.

Trying to find one person with credible experience and knowledge in all these areas is practically impossible. Instead, companies should take a team approach to AI ethics, achieving the required multi-disciplinary capabilities and experience by learning how to seamlessly orchestrate and integrate insights from a variety of specialized experts.

Contents

01 The rising importance of AI ethics

4

02 From afterthought to strategic business issue

6

03 Solving the problem: it takes a team

7

04 Specific responsibilities and essential capabilities

9

05 Moving forward on AI ethics

10



01

02

03

04

05

The rising importance of AI ethics

Deloitte's Trustworthy AI™ framework highlights the importance of addressing the challenges related to AI ethics and governance. And our most recent "State of AI in the Enterprise" survey ranked the need for managing ethical risks in AI as a top priority. Businesses today are rapidly expanding the scale and scope of their AI systems.

This trend, which has been accelerated by the COVID-19 crisis and need for social distancing, gives companies powerful new capabilities to improve how they operate. However, it also exposes them to heightened risk of AI behaving in ways that are unethical and inappropriate. And in the *Age of With*™, when humans and machines are increasingly working together, the risks and challenges related to AI ethics become even more important and complex.

Just like their human counterparts in the workforce, AI systems are expected to adhere to social norms and ethics, and to make fair decisions in ways that are consistent, transparent, explainable, and unbiased. Of course, figuring out what is ethical and socially acceptable isn't always easy—even for human workers.

Systemic bias remains a difficult and persistent challenge for humans and society in general. And unethical behavior has always been a risk in business. However, AI increases those problems exponentially.

When humans and machines are increasingly working together, the risks and challenges related to AI ethics become even more important and complex.



01

02

03

04

05

The rising importance of AI ethics

With human workers and person-to-person interactions, the scope and impact of unethical behavior is typically limited by a person's reach. But the reach of AI systems can be millions of times greater. Also, AI currently lacks the extraordinary general intelligence required to apply common sense to decision-making, and to autonomously understand and respect complex social norms. This can lead to AI actions that are technically correct but socially unacceptable.



The AI social dilemma

AI learns from the datasets used to train it, and if its programming and training data are biased it can amplify and propagate that bias at digital speed on a global scale—affecting vast numbers of people in the blink of an eye.

For example, a financial services company that uses AI to screen mortgage applications might find its algorithm unfairly discriminating against people based on factors that are not socially acceptable, such as race, gender, and age. Similarly, an AI system that decides on-the-fly where to place online job ads might unfairly target ads for higher paying jobs at a website's male visitors because historical data shows men typically earn more than women.



01

02

03

04

05

From afterthought to strategic business issue

Unethical or misbehaving AI can have severe consequences, including lawsuits, regulatory fines, angry customers, reputation damage, and destruction of shareholder value. However, those tangible consequences are just the tip of the iceberg.

Ultimately, the most compelling reason for your company to boost its capabilities for handling AI ethics issues effectively is that it has no choice. AI is quickly becoming an essential business capability and strategic differentiator, and organizations that don't figure out how to use AI ethically will likely find themselves constantly running into roadblocks that make it hard for them to use it at all.

Yet, despite this emerging strategic imperative, most companies currently lack effective mechanisms for developing, deploying, and operating AI that is ethical and trustworthy. Instead, responsibility for AI ethics is scattered across a variety of roles, with all those roles having other responsibilities that in practice are treated as far more important and pressing.

For example, AI programmers are primarily tasked with developing powerful AI systems as quickly as possible, and in their excitement over all the amazing things they can get AI to do it's easy for them to overlook ethical issues. Similarly, businesspeople involved in AI are often so focused on AI's power to create extraordinary

business value that they fail to see the potential ethical downsides for their customers and/or society at large. And when AI ethics issues do arise, they are often viewed from the perspective of regulators and lawyers.

Today, most companies don't have anyone whose primary responsibility is to identify and address AI ethics issues across the enterprise. To make matters worse, the people involved with AI tend to be experts in subjects such as technology, business, law, and regulatory compliance and lack expertise in subjects such as psychology, sociology, and philosophy that are essential to tackling ethical issues effectively.

AI ethics is scattered across a variety of roles, with all those roles having other responsibilities that in practice are treated as far more important and pressing.



01

02

03

04

05

Solving the problem: it takes a team

The most obvious way to fill the AI Ethicist role is to hire one person with expertise in all the required areas and then make that person responsible for ensuring all the organization's AI ethics issues get addressed. However, unless your business is just looking to “check the box” on AI ethics, there are at least two reasons why this approach won't work.

First, it is practically impossible to find an individual with credible levels of expertise in all the required areas. This high level of credibility is essential because the requirements of ethical AI often conflict with what AI developers and businesspeople would choose to do on their own (which is why an AI Ethicist is needed in the first place). Yet, without credible levels of expertise, it's easy for technical and business specialists to simply dismiss or ignore what the AI Ethicist is saying.

Second, even if a company is lucky enough to find someone with the required breadth and depth of expertise to effectively fulfill the role of AI Ethicist, the scope of AI ethics will likely outgrow that individual's capabilities in the very near future as AI becomes increasingly sophisticated and important in business.

We've seen a similar situation play out in the field of data science. Just a few years ago, many companies were scrambling to hire data scientists, and the requirements of the role were conceptually very similar to the AI Ethicist role in the sense that it demanded a diverse mix of technical and non-technical capabilities that are nearly impossible to find in one person. Also, while companies could clearly see the need to hire data scientists—given the growing importance of data and analytics—many companies weren't actually sure what to do with their data scientists once they hired them. Fast forward to today and you'll find that as the field of data science has grown in scope and importance, the role of data scientist at many companies has been replaced by a coordinated team approach that features separate experts in specialized areas such as dataset preparation, data engineering, machine learning, and model testing/deployment.



01

02

03

04

05


Solving the problem: it takes a team

It might seem appealing—or at least expedient—to designate one person as the sole champion for AI ethics. However, such an approach is likely to fall short. With ethical AI becoming a strategic business issue, everyone involved with AI needs to be responsible for AI ethics and must make it a higher priority in their day-to-day activities.

While there is clearly a need to make someone in the company formally responsible for overseeing AI ethics, that person should be viewed as just one element in a larger set of AI ethics resources.

A leader at the C-suite level—presumably the Chief Trust Officer or Chief AI Ethics Officer, for companies that have one—would be a logical choice to lead the charge. The dedicated role of AI Ethicist could add value to the process; however, the broader role of overseeing and championing AI ethics across the enterprise will likely require more organizational authority than that in order to be effective.

The broader role of overseeing and championing AI ethics across the enterprise will likely require more organizational authority than that in order to be effective.



Specific responsibilities and essential capabilities



01


02

03

04

05

A key responsibility of an AI ethicist or champion would be to improve the engineering approach to AI by adding ethical, social, and political perspectives to the design, development, and deployment of AI systems. Other key responsibilities include advising on ethical AI practices, protecting against unintended consequences of misbehaving AI, and ensuring accountability for AI-related decisions and actions. Fulfilling these responsibilities requires a broad range of expertise, skills, and capabilities, including:



Technical knowledge

Since technology is the foundation of AI, the AI ethics champion and team must have a robust working knowledge of AI technologies. They can't credibly advise on AI ethics issues if they don't fully understand the capabilities and limitations of the technology. Although they won't be doing the actual coding of AI systems, they need a deep understanding of technical concepts such as the difference between supervised and unsupervised learning, how user consent is obtained, and whether the machine learning model being used supports transparency—all the small but important details that affect how an AI system is designed, developed, and deployed.



Regulatory knowledge

AI ethics involves much more than regulatory compliance. In fact, the pace of AI innovation is so rapid that related laws and regulations are almost always lagging what the technology can do (and how social norms and public perceptions are shifting). However, regulatory knowledge is still essential. This includes not just knowledge of current regulations related to AI, such as data privacy laws, but also a forward-looking strategic perspective on what might arise in the future based on existing regulations and what ethics-related challenges are currently top of mind for lawmakers and regulators.



Business savvy and industry knowledge

AI ethics in a business context can't just be a philosophical exercise. Policies, frameworks, and other guidance related to AI ethics need to be usable in the real world. This requires practical business knowledge and experience. Without a clear understanding of the industry and business—and how AI ethics might affect various processes, systems, and stakeholder groups—AI ethics requirements can easily become so burdensome that they are simply ignored. An awareness of how the industry is addressing AI ethics is also crucial.



Communication skills and the ability to work across organizational boundaries

AI ethics issues are by their very nature complex and hard to understand. Also, they often require people to do things that are unintuitive or potentially burdensome. Helping people understand the issues and persuading them to change requires strong communications skills. This includes active listening, personal empathy, and the ability to clearly articulate what needs to be done—and why. It also requires the ability to operate effectively across organizational boundaries, up and down the corporate hierarchy as well as across functions and business units—often without direct authority over the people involved.

Moving forward on AI ethics

Reasonable people can disagree about the best way to fill the role of AI Ethicist. However, no one can deny the fact that AI is quickly becoming a fundamental business capability, and that AI systems need to behave in ethical and appropriate ways. As such, ignoring the issue of AI ethics is no longer an option. Whether you are leaning toward an individual- or team-based approach, the key is to pick an approach and get started—then adjust as needed based on the obstacles and opportunities you encounter.

The time to start is now.

Author

Beena Ammanath
Executive Director
of Deloitte AI Institute,
Managing Director

Deloitte Consulting LLP

Tel: +1 415 783 4562

Email: bammanath@deloitte.com





About Deloitte

Deloitte refers to one or more of Deloitte Touche Tohmatsu Limited, a UK private company limited by guarantee (“DTTL”), its network of member firms, and their related entities. DTTL and each of its member firms are legally separate and independent entities. DTTL (also referred to as “Deloitte Global”) does not provide services to clients. Please see www.deloitte.com/about for a detailed description of DTTL and its member firms. Please see www.deloitte.com/us/about for a detailed description of the legal structure of Deloitte LLP and its subsidiaries. Certain services may not be available to attest clients under the rules and regulations of public accounting.

This publication contains general information only and Deloitte is not, by means of this publication, rendering accounting, business, financial, investment, legal, tax, or other professional advice or services. This publication is not a substitute for such professional advice or services, nor should it be used as a basis for any decision or action that may affect your business. Before making any decision or taking any action that may affect your business, you should consult a qualified professional advisor.

Deloitte shall not be responsible for any loss sustained by any person who relies on this publication.

Copyright © 2021 Deloitte Development LLC. All rights reserved.