EXECUTIVE BRIEFING

Ethical AI

A Quick-Start Guide for Executives

BY SUSAN ETLINGER
MARCH 14, 2019

salesforce | ALTIMETER®
@PROPHET

TABLE OF CONTENTS

WHAT IS AI ETHICS?

While the idea of Artificial Intelligence (AI) — of machines that can replicate certain types of human capabilities — has been around for centuries, it's really only in the past several decades that it's evolved from science fiction to business reality. Back in 1950, Alan Turing, the English mathematician, computer scientist, cryptanalyst, and inventor of the eponymous "Turing Test," published Computing Machinery and Intelligence, a research paper that posed the question, "Can machines think?" Five years later, John McCarthy, an American computer scientist, coined the term "artificial intelligence," effectively introducing the term to the public and, in the process, creating a new field of study.

Today, the combination of massive amounts of data, inexpensive parallel processing, and continuously improving algorithms has made AI a reality — from powering search engines to image recognition to language-understanding algorithms that enable conversational agents like Google Home, Siri, and Alexa. But beyond the big technology companies, we're also seeing momentum in AI adoption among global companies. In its 2019 CIO Survey, Gartner noted that the number of Chief Information Officers (CIOs) that have deployed or plan to deploy AI has tripled in the past 12 months.

At the same time, the complexity of the technologies, the rate of change, and a lack of consensus on definitions makes it challenging to have a productive conversation about its implications. For our purposes, we define AI as a set of technologies that enable machines to reproduce certain types of human abilities; for example, the ability to see, listen, speak, move, reason, decide, predict, act, and — most importantly — learn from past experience.

Along with promising news of new AI deployments, however, we're also seeing a steady stream of news about the risks of AI and the need to use this extremely powerful set of technologies in a humane and responsible way. We've seen stories and books about how social and cultural bias, encoded into algorithms, can perpetuate and even amplify inequalities. We've seen the conversation escalate about authenticity and how and when bots and robots should disclose that they are not human.

We've seen the industry struggle with the implications of "deep fakes" — emerging technologies that enable machines to generate synthetic audio and video. We've seen more discussion of explanation — the ability for algorithms to reveal their decision-making processes and the duty that organizations have to put governance processes in place to manage AI responsibly.

What does all of this mean for business leaders? Simply that if we intend to unlock the potential of AI both to optimize and even transform organizations, we need to be clear-eyed about the issues we need to manage, the risks we will likely encounter, and the opportunities that lie on the other side.

As with many technologies of the past centuries, from the internal combustion engine to cloud computing, accelerating adoption tends to reveal both expected and unexpected issues in society and in business. With AI in particular, we must consider its impact on the education, healthcare, and criminal justice systems, as biased data models can lead to disparate and unfair outcomes for vulnerable populations.

From a business standpoint, the implications of AI are not merely technical; they span the organization. The most obvious areas are legal and operational risk. In fact, Alphabet recently warned investors that flawed AI could affect revenues and operating results, while Microsoft warned of potential "brand or reputational harm." Figure 1 lays out the primary AI ethics topics that organizations must consider:

Figure 1. Key AI Ethics Issues for Business

1	2	3	4	5
BIAS	**EXPLANATION**	**AUTHENTICITY & DISCLOSURE**	**GOVERNANCE**	**IMPACT**
The processes and techniques used to reveal and address societal and statistical bias in data models and algorithms that underlie products and services.	The ability to provide customers, consumers, citizens, patients, and other constituencies with explanations for decisions made by algorithms.	How organizations use new interaction models, such as chatbots, voice agents, and robots, in a trustworthy manner.	Policies and processes that govern the use of AI technology within and outside the organization.	The human and business impact of ethical programs, policies, and processes on the organization.

It's important to note that not all of these issues lie purely in the technical realm or in the realm of consumer technology companies, such as Facebook, Amazon, and Google. And while algorithmic bias and technical solutions to explanation have ignited entire fields of study within the technology and academic communities, questions of how organizations disclose their use of AI and bots, how they govern AI, and how they measure its impact are wide-ranging business decisions that

will affect any organization that uses these technologies.

As a result, the stakeholder list for AI ethics is broad and deep, comprising research, engineering, product management, design, legal, human resources, communications, public policy, ecosystems, human rights organizations, and beyond. The following section lays out a "quick start" guide to help your organization plan its approach to responsible AI.
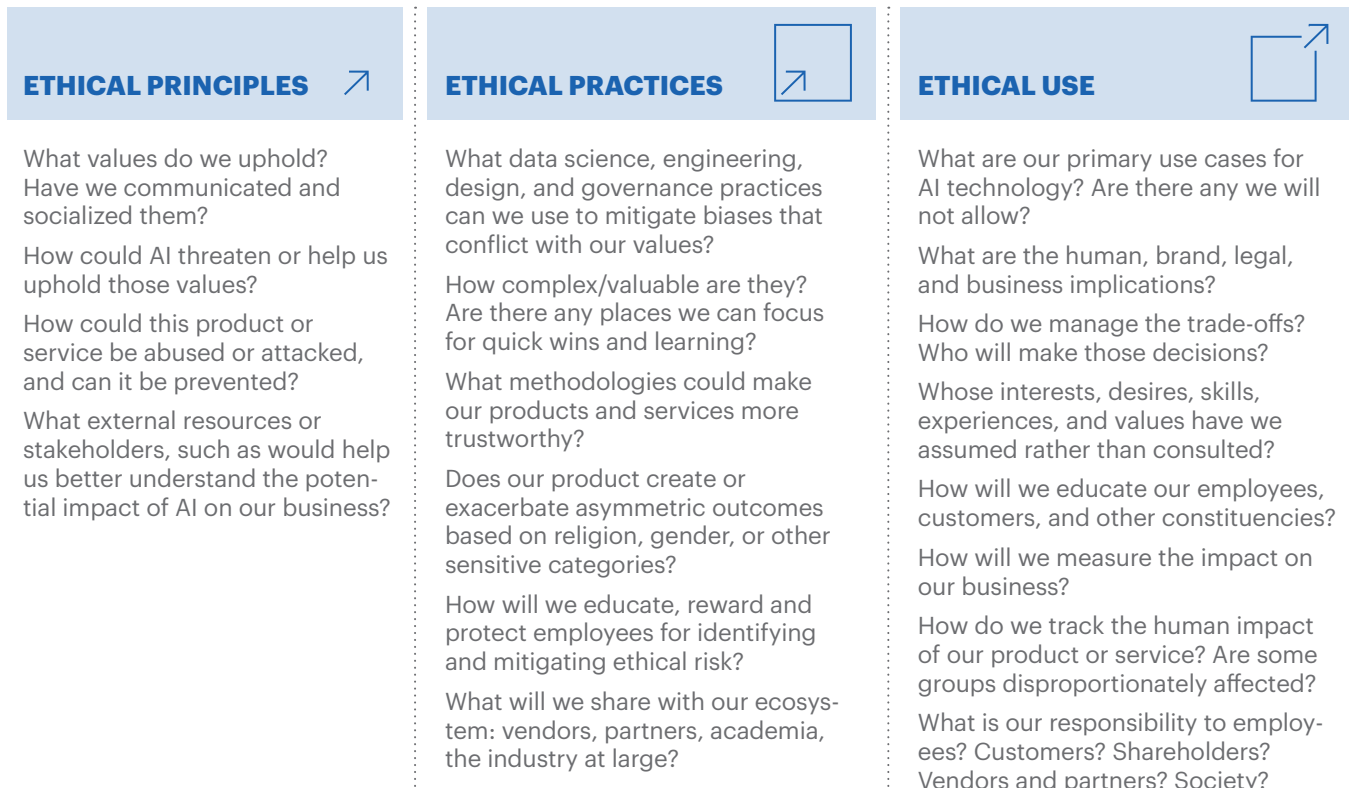
BUILDING AI ETHICS INTO YOUR ORGANIZATION

SETTING A STRATEGY

One of the most important strategic decisions you can make at the outset of an AI ethics program is to determine your strategy: what you will do, what you won't do, and in what order. It's critical to remember that this is a complex and dynamic set of technologies with a complex and dynamic set of implications, so focus early on (with an eye to the horizon) is your friend.

This may mean that you focus on a single issue to start, such understanding your data:

Is it plentiful, clean, and labeled? If you have already taken that step, a next natural step would be to focus on bias in your data set, model, or application, or start with a single stakeholder group, such as research or governance. Whatever your starting point, it's important to think both at a micro level (What will we do first?) and a macro level (What issues are likely to emerge?). Figure 2 lays out a set of simple "thought-starter" questions to help you craft an initial action plan.

Figure 2. Scoping an Ethical AI Action Plan: Principles, Practices, Use

ETHICAL PRINCIPLES ↗

What values do we uphold? Have we communicated and socialized them?

How could AI threaten or help us uphold those values?

How could this product or service be abused or attacked, and can it be prevented?

What external resources or stakeholders, such as would help us better understand the potential impact of AI on our business?

ETHICAL PRACTICES ↗

What data science, engineering, design, and governance practices can we use to mitigate biases that conflict with our values?

How complex/valuable are they? Are there any places we can focus for quick wins and learning?

What methodologies could make our products and services more trustworthy?

Does our product create or exacerbate asymmetric outcomes based on religion, gender, or other sensitive categories?

How will we educate, reward and protect employees for identifying and mitigating ethical risk?

What will we share with our ecosystem: vendors, partners, academia, the industry at large?

ETHICAL USE ↗

What are our primary use cases for AI technology? Are there any we will not allow?

What are the human, brand, legal, and business implications?

How do we manage the trade-offs? Who will make those decisions?

Whose interests, desires, skills, experiences, and values have we assumed rather than consulted?

How will we educate our employees, customers, and other constituencies?

How will we measure the impact on our business?

How do we track the human impact of our product or service? Are some groups disproportionately affected?

What is our responsibility to employees? Customers? Shareholders? Vendors and partners? Society?

IDENTIFYING STAKEHOLDERS: AI IS A TEAM SPORT

Depending on the nature of your business, you will likely have a wide-ranging group of internal and external stakeholders. Figure 3 lists potential stakeholders and the key issues to address.

Figure 3. Scoping an Ethical AI Action Plan: Principles, Practices, Use

STAKEHOLDER GROUP	CONSIDERATIONS
Internal Stakeholders	
Research and Data Science	• Consider impacts of ethical and societal implications of AI models, and disclose them in research papers. This is standard practice in academia but emerging in business.

Product Development (Product Management, Engineers, User Experience Researchers and Designers)	• Consider impact of using a protected data class in a data model, as that may lead to a biased decision.
	• Look for collinearity (e.g., highly correlated variables, such as zip code and race) and data leakage, both of which can result in biased outcomes.
	• Consider places in digital products where it is possible to alert users to low confidence or potentially ambiguous results.
	• Conduct an ethics pre-check as a way to discuss whether a proposed product, service, or feature could potentially result in biased or unintended outcomes. If using agile methodologies:
	  – Include the ethics pre-check in the "definition of ready" and "definition of done."
	  – Build in a retrospective to identify what went well and what should be done differently next time.
	• If your organization uses agile methodologies, adding ethics as an attribute of readiness, completeness, and success will scale.

Customers	• Work with customers to educate and empower them:
	  – If you are a business-to-consumer company, this may consist of disclosures in digital products and services, provisions in terms of use, and other digital content to educate consumers about ethical AI.
	  – If you are a business-to-business company, this may consist of the above, plus white papers, executive briefings, and consultations with customers if they have questions or concerns.

Legal	• Partner with your legal department to:
	  – Better understand data privacy laws, such as GDPR and others, to ensure that you are not unintentionally exposing the company to risk.
	  – Craft acceptable use policies that cover how customers and consumers may use your products and services in situations where there is no legal precedent.
	• Educate yourself on emerging guidelines and regulations on how AI may be used within governments or countries.

Figure 3, continued. Scoping an Ethical AI Action Plan: Principles, Practices, Use

STAKEHOLDER GROUP	CONSIDERATIONS
Internal Stakeholders	
Human Resources	• Engage and partner to understand and update employee processes (onboarding, performance reviews, incentive structures) and resources (intranet, managers) to account for issues and opportunities raised by AI ethics.
Diversity and Inclusion	• Diverse teams make scenario planning and bias identification easier and reduce the likelihood that a similar group of people will view a situation in the same way. Engage with D&I teams to identify places where ethics and diversity and inclusion programs can support each other.
Communications	• Work with marketing and PR to understand how they engage with customers at the point of purchase and to develop clear, trustworthy messages about how you do (and don't) use AI.
Other Stakeholders	
Non-Governmental Organizations	• Engage with NGOs or other expert advocacy groups like World Economic Forum, Omidyar, Partnership on AI to ensure you are listening to and learning best practices, as well as sharing what you have learned.
Ethicists	• Engage with applied ethicists at academic institutions to preview key decisions and collect insight on potential ethical issues and solutions.
Government	• Engage and maintain dialogue with think tanks and policy advisors on issues related to using artificial intelligence inclusively, fairly and for the public good.
Shareholders	• Begin to build, and continually test, a way to evaluate the impact of ethical technology use on key metrics such as brand preference, risk mitigation, customer and employee loyalty, and other key performance indicators (KPIs).

RECOMMENDATIONS

1. Find your community of practice.

One of the common pitfalls of beginning an AI ethics program is the assumption that you need to start from scratch, but many organizations, such as the World Economic Forum, The Partnership on AI, Omidyar Network, The AI Now Institute, and others, are excellent resources. Says Paula Goldman, Chief Ethical Use Officer at Salesforce, "There is a community of practice around these issues; it's really important not to reinvent the wheel. It's complicated, because some of these issues are complex and confidential, so it's sometimes hard for companies to share what they are doing, but these problems can't be solved as islands. If we're going to shift norms in the industry, we have to work together."

2. Listen before doing.

Goldman counsels organizations beginning an AI ethics program to listen first. "These issues are really complex, and opinions differ in and across cultures, so we must be much more inclusive around finding solutions from civil society and beyond in a really meaningful way." This means actively engaging with outside and internal stakeholder groups, listening, collaborating on solutions, and building accountability.

3. Identify stakeholders and core issues.

There are so many AI technologies, so many use cases, so many ways of using them, and so many implications that there is no single hard-and-fast rule for where to begin. Once you've assembled a community of practice and done some initial discovery, work with stakeholders to collectively agree on a few key priorities to tackle first. This may mean clarifying your responsible AI principles (Accenture is a good example), looking at your product development process, or updating terms of use to start.

4. Reinforce AI ethics throughout the business — not simply as a compliance exercise.

One of the potential pitfalls of an AI ethics program is that it is seen merely as a compliance exercise and not as a fundamental to success. Rather than creating "the ethics police," it's critical to build relationships from the ground up and raise awareness internally. As part of this, says Kathy Baxter, Ethical AI Architect at Salesforce, "you have to cultivate an environment of psychological safety, where the people who are closest to customers feel safe finding issues. They're the ones most likely to find issues early, so why wouldn't you listen to that?"

5. Create Organizational Support for Ethics.

One of the most fundamental drivers of ethical AI and ethical technology overall is a diverse and empowered workforce. Says Lily Jampol, Ph.D., Social Scientist at ReadySet, a diversity and inclusion consultancy, "AI is not really distinct from other technologies in that the more diversity of perspective and backgrounds you have, the more likely you are to spot problems you wouldn't have seen otherwise." It's also critical to create incentive structures that reward ethical behavior, which may run counter to standard incentive metrics, such as user engagement. Additionally, it's important to ensure your operational or business processes support the responsible creation of AI by adding ethical questions and discussions into product development and software development lifecycle. If you are an enterprise company, provide incentives and protections for sales staff to raise a red flag if they believe a potential customer will use your products in a harmful way.

6. Track your impact.

One of the core issues in AI ethics is the question of how organizations evaluate the effectiveness of AI ethics programs on the business, customers and stakeholders. While there is as yet no generally accepted set of standard ethics metrics, it's important to start to put mechanisms in place to understand and track impact. Are your recommendations being implemented internally? Is ethics a differentiator for your product? Are customers mentioning ethics as an indicator of trust in your brand?

AI is still an emerging discipline, and no one has all the answers. Listening, educating yourself and your organization, and, as with AI, agile experimentation are key. But while this is a complex topic, it is by no means a zero-sum game that pits growth and ethics against each other. In fact, trust in intelligent technologies and the products and services they enable will be key to growth. Says Salesforce's Goldman, "The organizations that go out of their way to establish and maintain trust are going to win."
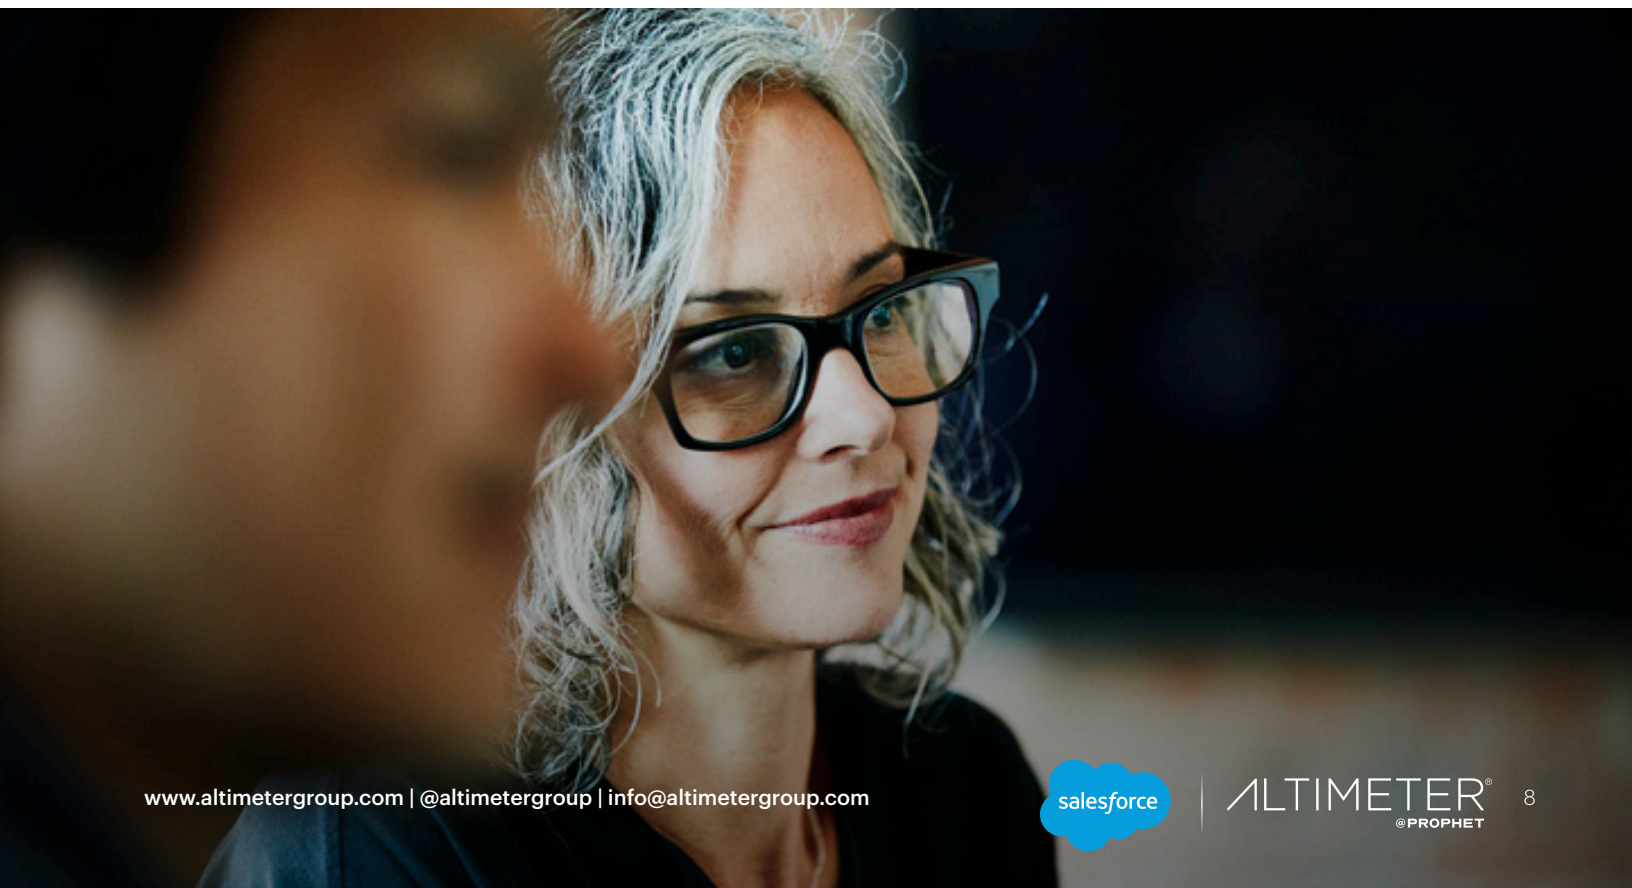
ABOUT US

SUSAN ETLINGER, INDUSTRY ANALYST

Susan Etlinger (@setlinger) is an industry analyst with Altimeter, a Prophet Company, where she publishes research and works with clients on issues related to data strategy, artificial intelligence, digital ethics and trust, and the impact of emerging technologies on business and society. Etlinger was named a 2016 LinkedIn top Voice for Technology, and her TED talk on Big Data has been viewed over 1.25 million times. She has been quoted in media outlets, including *The Wall Street Journal*, *The New York Times*, and *BBC*.

ALTIMETER, A PROPHET COMPANY

Altimeter is a research and consulting firm owned by Prophet Brand Strategy that helps companies understand and act on technology disruption. We give business leaders the insight and confidence to help their companies thrive in the face of disruption. In addition to publishing research, Altimeter analysts speak and provide strategy consulting on trends in leadership, digital transformation, social business, data disruption, and content marketing strategy.

SALESFORCE

Salesforce, the global CRM leader, empowers companies to connect with their customers in a whole new way. For more information about Salesforce (NYSE: CRM), visit: **www.salesforce.com**.