

Why AI Needs a Social License

FEBRUARY 22, 2022

By [François Cadelon](#), [Rodolphe Charme di Carlo](#), and [Steven Mills](#)

Six years ago, in March 2016, Microsoft Corporation launched an experimental AI-based chatbot, TayTweets, whose Twitter handle was @TayandYou. Tay, an acronym for “thinking about you,” mimicked a 19-year-old American girl online, so the digital giant could showcase the speed at which AI can learn when it interacts with human beings. Living up to its description as “AI with zero chill,” Tay started off replying cheekily to Twitter users and turning photographs into memes. Some topics were off limits, though; Microsoft had trained Tay not to comment on societal issues such as Black Lives Matter.

Soon enough, a group of Twitter users targeted Tay with a barrage of tweets about controversial issues such as the Holocaust and Gamergate. They goaded the chatbot into replying with racist and sexually charged responses, exploiting its repeat-after-me capability. Realizing that Tay was reacting like IBM’s

Watson, which started using profanity after perusing the online Urban Dictionary, Microsoft was quick to delete the first inflammatory tweets. Less than 16 hours and more than 100,000 tweets later, the digital giant shut down Tay. Although Microsoft is one of the pioneers and adherents of the principles of “Responsible AI” in algorithm development, Tay was a public relations disaster. And critics, ominously, saw the problem as “AI at its very worst—and only [the] beginning.”

Two years earlier, Amazon quietly built an AI algorithm that could review and rate job applications on a five-point scale. The objective was to screen the enormous number of resumes it received and identify the most promising candidates. The retailer created 500 models to analyze applicants for each job by location, and it taught the algorithm to recognize more than 50,000 terms that had appeared in the applications and resumes it had received in the past. The process helped AI learn to assign a low weight to generic skills, such as the number of computer languages a programmer knew.

Over time, however, Amazon woke up to the fact that AI wasn’t rating candidates in a neutral way; it recommended mostly males. The company had trained the algorithm on applications submitted to the company over a decade, and most of those applicants had been men. The AI therefore learned that male candidates were preferable. It penalized resumes that included words such as “woman” or “women”—as in “women’s college” or “women’s gold medal winner.” The AI development team edited the program to make it gender-neutral, but it couldn’t guarantee that the algorithm wouldn’t find other ways of discriminating against women. Although Amazon is a member of the Partnership on AI and thought it had built an algorithm that was fair and inclusive, it had to terminate the experiment by 2016.

In January 2020, Clearview AI suddenly shot into the limelight for all the wrong reasons. A low-profile US facial-recognition company, it flew under the radar until The New York Times published an exposé titled “The Secretive Company That Might End Privacy as We Know It.” It turned out that Clearview was providing software to companies, law enforcement agencies, universities, and individuals, with its algorithm matching human faces to a database of more than three billion images it had indexed from the internet.

That kicked off a fierce global debate about the use of AI-based facial recognition by governments and law enforcement agencies. Most people called for a ban on Clearview’s AI because it had created its database by mining the internet and social media websites. In January 2020, Twitter sent the firm a cease-and-delete letter, and YouTube and Facebook followed. When the COVID-19 pandemic erupted in March 2020, Clearview pitched its technology for use in contact tracing in an effort to regain its credibility and gain social acceptance. Although Clearview’s AI could have helped tackle the pandemic by working as it was designed, the manner in which the firm had gathered data raised a social firestorm that prevented its deployment during the global crisis.

*Even as AI drives the **digital transformation** sweeping business today, it is becoming the digital elephant in the room. Companies are turning to AI to cope with an ambiguous, complex,*

uncertain, and volatile future in the aftermath of the pandemic. In addition to enabling automation, AI helps business forecast the future, improve decision making by providing data-driven insights, and prepare for the unexpected by developing complex scenarios. With market leaders trying to use [AI at scale](#), it's difficult to think of an industry that AI doesn't seem likely to upend. Yet business faces a foundational challenge in getting started with AI.

The issue, our studies show, isn't technological; it's human. Because of our history of technophobia—probably dating back to Socrates (470-399 BC), who warned against writing because it would “lead to forgetfulness and weaken the mind”—when humans see technology that behaves like them or mimics their decision-making skills, they worry. As any number of movies also suggest, our fears and suspicions stem from either the idea that AI will evolve beyond its creators' expectations, escape, and eradicate the human race, or the possibility that some humans will develop AI for evil purposes. “Mark my words, AI is far more dangerous than nukes. I am really quite close to the cutting edge in AI, and it scares the hell out of me,” admitted Tesla and SpaceX founder Elon Musk four years ago at the South by Southwest Conference.



When humans see technology that behaves like them or mimics their decision-making skills, they worry.

Using AI, which generates as much hope as it does horror, therefore poses a conundrum for business. It's compounded by the fact that many companies already follow the principles of responsible AI while developing the technology, and they assume that's more than enough. Responsible AI is a framework that, if adhered to, should enable companies to develop AI systems that work for the good, going beyond algorithmic fairness and bias to identify the potential effects of the technology on safety, privacy, and society. However, following the principles of responsible AI is not enough to ensure that AI's use yields only acceptable outcomes. For instance, Georgetown University's Center for Security and Emerging Technology and Partnership on AI¹ recently published a study documenting 1,200 publicly reported cases of AI failures over just three years.

Some CEOs, such as those of digital giants like Microsoft and Google, have called for governments to regulate AI. Even as policymakers step up to the challenge of doing so, the development of trustworthy AI will still be up to each company and the scientists who write the

algorithms. As a result, efforts to prevent AI-caused negative outcomes and unforeseen consequences vary by company, business, and function, which is far from ideal. If business wants to use AI at scale, it needs to go beyond responsibility in AI development; it must obtain society's explicit approval to deploy it. In other words, companies have no option but to acquire a social license for AI.

Defining a Social License to Operate

The concerns of rational people about AI extend across the spectrum—from the algorithmic institutionalization of income, gender, racial, and geographic prejudices to privacy and political issues. That's why companies are struggling to come to terms with the gulf between possessing the legal right to use AI, which they have, and the social right to do so, which they may not—in a nutshell, between what business can do and what it should do.

Like every technology before it, AI will deliver benefits and have costs; how companies apply the [technology](#) will determine whether or not the benefits outweigh the costs. To quell fears, business must ensure that the benefits are maximized and the costs minimized. Surveys show that most people believe that there should be oversight of AI, and that companies should adhere to strict codes of conduct while using AI. Three years ago, the Center for the Governance of AI at the University of Oxford found that 82% of respondents believed that AI should be carefully managed.²

Tough regulations for AI are on the way. In May 2021, the European Union proposed rules for AI usage that it hopes will become the de facto global standard, just as the EU's data privacy regulation, the General Data Protection Regulation, became the global privacy standard after it came into effect in May 2018. Around the same time, the US Federal Trade Commission issued a rare guidance making it clear that using an algorithm that results in discrimination would constitute “unfair or deceptive practices” that the FTC Act prohibits. It warned companies to refrain from gathering training data for algorithms in a misleading way. “Keep in mind that if you don’t hold yourself accountable, the FTC may do it for you,” the agency warned bluntly.

It's becoming evident that responsible AI, which is primarily a methodology to manage AI's technical failings, is far from enough. Those principles, as we at BCG have [discussed elsewhere](#), entail developing AI that integrates human empathy, creativity, and care to ensure that it works for the greater good. It starts with embedding accountability into all levels of an organization, as well as across all stages of an algorithm's life cycle. Maintaining human control is central to responsible AI; the risks of AI failures are greatest when timely human intervention isn't possible. It also demands tempering business performance with safety, security, and fairness.



The risks of AI failures are greatest when timely human intervention isn't possible.

Several stakeholders—technology firms such as Microsoft and Google, research institutes and think tanks such as The Ethical Institute, and consulting firms—have proposed ways in which business can develop AI responsibly. Unfortunately, the number of principles that have been published tends to confuse the issue, especially because there is significant variance over what constitutes responsible AI. It's also tough for companies to figure out how to bridge the gap between the principles of responsible AI and the actions they must take. As [a recent article](#) in MIT-SMR pointed out: “But even if there is broad agreement on the principles underlying responsible AI, how to effectively put them into practice remains unclear. Organizations are in various states of adoption, have a wide range of internal organizational structures, and are often still determining the appropriate governance frameworks to hold themselves accountable.”

Besides, responsible AI represents [a mainly technology-based approach](#) to the issue. In fact, scientists usually focus on the technical challenge of building goodness and fairness into AI, which, logically, is impossible to accomplish unless all humans are good and fair. Moreover, the answer to whether AI’s response to a problem is ethical is usually “It depends.” That’s the classic trolley problem,³ where you can do nothing and kill five people on one track—or divert the trolley and kill one person on another track. It just isn’t obvious which option is better.

What’s necessary is to adopt a general, human-focused approach to AI. As the MIT-SMR article cited earlier says: “When asked what an ideal future state would look like, interviewees preferred an approach that anticipates rather than reacts to risk. In order to achieve that, organizations need standard processes, communication, and transparency.” That’s why we believe the concept of a social license for AI is appropriate.

Coined to symbolize community approval of projects that affect the environment—as is commonplace in mining, forestry, fishing, and energy—the term began as a metaphor. Jim Cooney, an executive with the Canadian mining company Placer Dome—whose tailings dam at a gold mine in the Philippines collapsed and released toxic mud that buried a village—may have used it for the first time in 1997 at a World Bank meeting. He pointed out that if mining companies lost their social licenses, local and national communities would show little hesitation in shutting them down.

To avoid that, before a company in an environment-destroying industry invests in a project, it must understand the first-order and second-order ramifications. While regulations govern the management of the direct consequences—land acquisition, environmental pollution, and water use—they aren't sufficient. The business must also invest in the local community's physical infrastructure, such as roads, electricity, and telecommunication services; improve people's access to education and health care; and foster economic activity, so local people benefit from its presence in more ways than one. Only then will the company earn, and retain, the community's goodwill to operate the project.

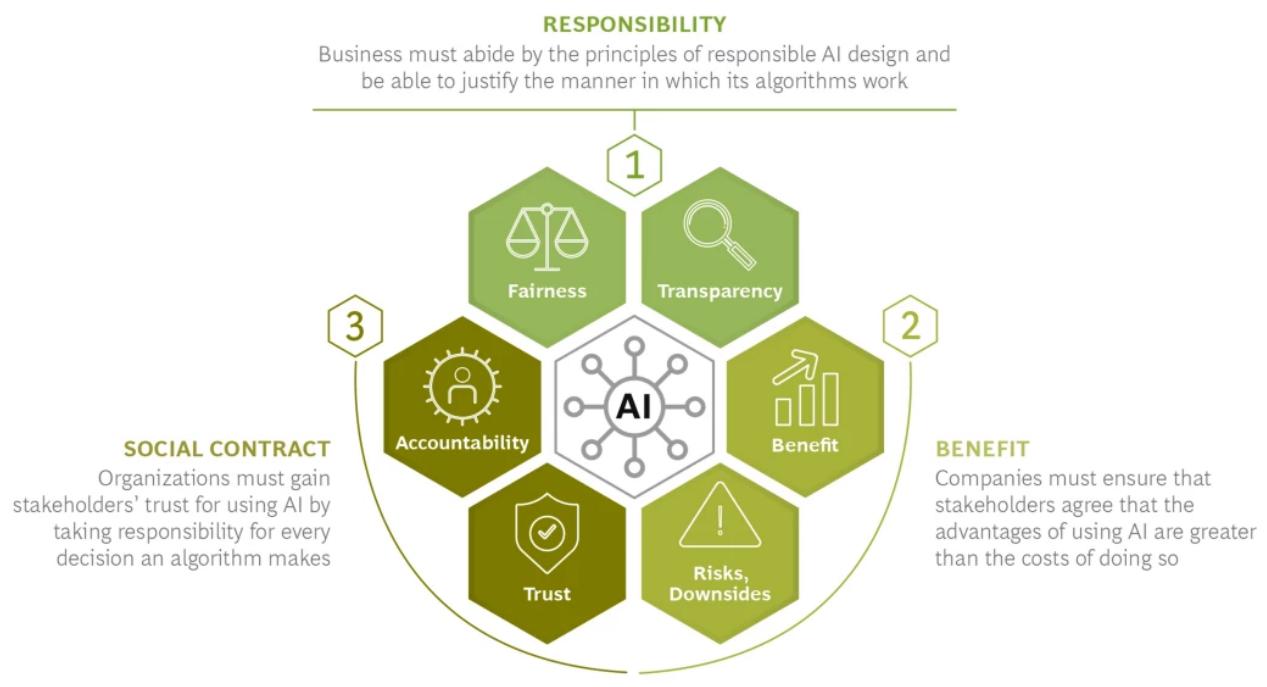
Keep in mind that the social license isn't a document like a government license; it's a form of acceptance that companies must gain through consistent and trustworthy behavior and stakeholder interactions. Thus, a social license for AI is a common perception in society that a company has the right to use the technology for specific purposes in the markets in which it operates. Companies cannot award themselves social licenses; they must win them by proving they can continue to be trusted, as John Morrison argued in 2014.⁴ And losing the social license to operate can have dire consequences, as the energy companies Shell (accused of polluting the Niger Delta) and BP (after the Deepwater Horizon oil spill in the Gulf of Mexico in 2010) learned.

In order to obtain a social license to use AI for decision making, companies must work closely with stakeholders—employees, software developers, consumers, and shareholders, among others—to initiate a dialogue. These two-way conversations will catalyze the expression of reactions, motivations, positions, and objections from diverse groups. That will allow company executives to develop a shared understanding with stakeholders about each AI application and the guardrails that must fence its application.

The Three Pillars

To use AI, companies will depend on a legal license, in the form of regulatory permits and statutory obligations; an economic license, which stems from the demands of shareholders and executives; and a social license, or the people's demands. Only then will society sanction the sustained use of AI. Our studies show that a social license for AI rests on three pillars: responsibility, benefit, and social contract. (See the exhibit.)

The Three Pillars of an AI Social License



Source: BCG analysis.

Responsibility

If business is to be answerable to society about using AI, it must be able to justify the manner in which AI algorithms work. Companies will be able to do that if, to begin with, people perceive that the AI algorithms are ethical. Society will hold that opinion as long as companies design algorithms that, to the extent possible, are perceived to be fair and transparent when they work.

AI will be deemed to be fair if the outcomes generated by using the technology don't vary because of demographic factors, such as age, gender, or location, or because of changes in the economic, social, or political context. For example, a company that uses an AI-based recruitment system must be able to demonstrate that all the candidates who provided the same or similar responses to a question posed by the machine received the same rating or score.

Society must perceive the AI application to be transparent in terms of its working and results. Business should be as open as it can be about an algorithm's design—just as the open-source community makes public all the code of the cutting-edge software it develops. In addition, executives must be able to explain causality, the mechanics of the algorithm's decision making and logic. That may not be easy. In August 2020, for instance, the British government's decision to grade school students taking their "A level" examinations with an AI-based system evoked howls of protest. It happened because of a lack of transparency about how the algorithm

made grading decisions. Eventually, the British government had no option but to advise the country's universities to ignore the AI-determined grades when granting college admissions.



Society must perceive the AI application to be transparent in terms of its working and results.

Human decisions may not always be fair or transparent; every decision has conscious and unconscious biases shaping it. But society seems to expect that AI, even though humans create it and feed it data that they directly or indirectly produce, will always make decisions that are fairer and show greater transparency. As Gill Pratt, CEO of the Toyota Research Institute, said in 2017: “Remember elevators from the past? They used to be staffed, but over time, automatic elevators became widely accepted. We must ensure that our products are safe and that automated vehicles perform significantly better than a human driver if we want a mother to trust her children’s life with an automated car.” Thus, society’s standards for accepting algorithmic decisions will tend to be higher than those for accepting human decisions.

At the same time, many decisions are acceptable to society, even if they’re made in an idiosyncratic fashion, if the decision maker can articulate a rationale. The same logic should apply to AI algorithms, even though people will hold AI to higher standards. For example, people who entrust their savings to an investment company must understand the firm’s investment philosophy and its strategy, and they must communicate their risk-tolerance preferences to their financial advisors. It shouldn’t matter thereafter whether human traders or AI algorithms execute trades on a customer’s behalf, as long as the investor’s objectives are met.

Benefit

Companies must ensure that stakeholders share their perception that the advantages of using AI systems are greater—or, at least, no less—than the costs of doing so. They can measure the tangible and intangible tradeoffs at the individual, company, and societal levels by comparing the benefits of the AI-delivered outcomes—such as increased health, convenience, and comfort, in the case of health care—with the potential downsides, which usually relate to security, privacy, and safety.

Some AI-based systems will demonstrate their benefits easily. For example, the AI application Google Translate offers the advantage of being able to instantly translate text to and from more than 100 languages with increasing accuracy. No downsides are obvious, although some professional translators and translation firm employees may have lost their jobs as a consequence. In 2020, Google Translate translated more than 100 billion words a day and boasted 500 million users. In fact, the use of AI-based translation is rising; according to a 2019 Mordor Intelligence report, the industry will grow at 11% a year until 2025, when it will reach \$1 billion.



Society's verdict will not always favor AI's use, and business should be prepared for that.

In other cases, AI's benefits may not be overwhelmingly positive compared with the costs, so stakeholders will need to be persuaded. Consider the German delivery company DHL, which installed an AI-based system two years ago. It scans all the pallets that employees load in the company's planes and determines which ones should be loaded on top because they are fragile, and which ones can go below them. In the beginning, AI couldn't beat the expertise of human decision making, but it became better over time. AI saved many pallets from damage, and it allowed DHL's employees to do their work faster, which gained AI their grudging acceptance.

Society's verdict will not always favor AI's use, and business should be prepared for that. For instance, when the COVID-19 pandemic began, health care experts mooted the idea of merging all the patient databases in the EU to hasten the creation of an algorithm to help find a cure. However, it wasn't immediately obvious to policymakers and experts if the benefits of finding a solution faster would be greater than the costs of breaching patients' privacy and allowing business access to confidential, and coveted, health-related data. And there the matter ended.

Social Contract

Finally, society must accept that companies that want to develop AI can be trusted with its use, as well as the acquisition and analysis of real-time data to feed their algorithms, and that they will be accountable for the decisions made by AI systems. Trust is critical for social acceptance,

especially in cases where AI can act independent of human supervision and have an impact on human lives.

That's one reason why society has been slow to approve unrestricted use of self-driving automobiles, for example. Autonomous cars will probably be safer than human-driven ones; according to the US NHTSA, human error caused 93% of crashes on US roads in 2020. Although people trust the incumbent automobile manufacturers to create reliable mechanical vehicles, they don't believe in their ability to develop digital technologies. For instance, Fiat Chrysler discovered in June 2015 that the automobiles it had equipped with touchscreens could be attacked by hackers, who could even switch off the engines of moving vehicles. The carmaker had to recall 1.4 million vehicles to update the software and plug the vulnerability. In addition, public trust in autonomous vehicles has been eroded by the fact that the telecommunication companies' networks, which provide the 5G infrastructure for communications between man and machine, have been frequently hacked for decades.

Autonomous vehicles also suffer from a lack of accountability. Vehicle owners don't know who will be responsible for the decisions that their autonomously driven cars make on the road. And, in the event of an accident, who will bear the liability? Governments haven't drawn up regulations that will allow self-driving vehicles to operate at scale, so manufacturers are finding it tough to enter into a contract for AI-driven automobiles with society.

Winning a Social License

Although the paths that business can take to obtain a social license for AI are not yet clear, the first guideposts are becoming visible. For one thing, the approach will need to vary by the problem that AI is tackling and the number of stakeholders that are involved in the process. Identifying the stakeholders who will influence the grant of a social license is therefore a key prerequisite. The stakeholder groups may be few or many, with divergent or convergent objectives. In our earlier example, DHL had to worry about obtaining a social license for its pallet-sorting algorithm only from employees, trade unions, and shareholders. By contrast, Tesla must work with a wide number of stakeholders—such as automobile owners and drivers (consumers); city, state, and federal governments (regulators); and civil society (critics and advocates)—to obtain a social license for its self-driving technology.

As a rule of thumb, business leaders must keep two things in mind while dealing with stakeholders. On the one hand, a company must identify the least influential stakeholders' lowest conditions for the use of AI; that's the threshold for each pillar. On the other hand, it must identify the highest expectations of the most influential stakeholders, which establishes the ceilings. Executives must then strike a complex balance between the thresholds and the ceilings for the social and economic pillars. DHL, for instance, had to balance the tension

between augmenting employees' jobs without replacing them (the threshold) and maximizing shareholder returns (the ceiling).

Moreover, every organization operates in a different context depending on its history, geography, and current capabilities. Once executives figure out their company's starting point, they can take seven steps to obtain a social license.

Communicate all the costs and benefits. At the outset, companies must explain to stakeholders all the possible benefits of AI, as well as all the potential downsides. It's important to reach out and talk to each group rather than burying the information in fine print at the bottom of a website.

Businesses shouldn't shy away from openly describing the risks of AI, but at the same time they must explain how they will tackle them. Google, for example, shares an organization-wide commitment to ensuring the safe, fair, and unbiased use of AI by publishing and periodically updating its document "Policies & Principles of AI at Google." One of the first companies to start the practice, Google publicly outlines the ethical boundaries surrounding the use of AI in its products. Doing so provides a public reference for anyone who wants to criticize the company if they believe that Google isn't complying with its own rules. It's what Google intended: "They [*"Policies & Principles"*] don't—and shouldn't—allow us to sidestep hard conversations," but to kickstart them, says the company.



Businesses shouldn't shy away from describing the risks of AI, but they must also explain how they will tackle them.

BCG followed the same logic in being specific about the workings of CO2 AI, an AI-based solution that helps organizations simulate, measure, and track their environmental footprints.

Open the black box. Most companies haven't developed the expertise to build AI in ways that explain to users exactly how the algorithms work without giving away too much to rivals. CEOs should consider partnering with innovators—such as AI startups, digital giants, think tanks, and scientists—to develop ecosystems that will help them understand the cutting edge of AI development.

For example, scientists are already trying to create self-explainable AI—that is, algorithms that can provide a decision as well as a supporting explanation, but without sacrificing the accuracy of the former for the sake of the latter. The focus isn’t on the correlation; instead, the search is for the “because”—an explanation of why the decision made sense to the AI application. Similarly, causal AI identifies the factors in an algorithm that lead to particular behaviors or outcomes, and it tests what will change them. In a 2020 *Stanford Social Innovation Review* paper,⁵ the authors highlight the fact that causal AI can help avoid the mistakes that arise when people, and the AI applications they create, ignore contexts or mistake correlation for causation. It’s the kind of research from which business can benefit, so companies must develop AI ecosystems that will help them access it from the outset.

Develop override mechanisms. Unexpected outcomes that have negative fallout are a possibility when companies depend on AI, so they should act preemptively to prevent them. Recognizing the limits of algorithms, business must carefully define the playing field. Doing so will ensure that AI learns to flag exceptions that it cannot process and to recommend human over-the-loop interventions to deal with them. It would be impossible for that to happen without built-in override mechanisms.

In March 2020, for example, the AI-based cybersecurity mechanisms of one of Britain’s premier online grocery stores suddenly shut down the entire website. It had confused an unprecedented COVID-19-related surge in demand with a distributed denial-of-service attack. Engineering a way of monitoring and managing the system’s kneejerk response would have helped the British retailer better cope with the situation, generate more revenues, and not, as it did that day, antagonize all its customers.

Prioritize risk management. Using AI is risky, just like every other business decision. While 88% of respondents expressed confidence in their ability to address AI-related risks, 46% had taken no action, according to a Clifford Chance survey of board members three years ago. Companies must learn to mitigate AI-related risk by mapping and evaluating it systematically. They must examine the elements inherent to the risk—such as its severity and probability—as well as the context, such as the regulatory environment and their general **risk management** strategies.

Managing AI-related risk doesn’t differ very much from handling other risks, such as human error or employee fraud, so chief risk officers should determine the levels of legal, financial, reputational, and physical risk that their companies are willing to take. They must do so early in the AI utilization process. As the CRO of Zalando, the Berlin-based e-commerce company, wrote in September 2020: “Recognizing the relevance of proper risk management for AI and machine learning as a condition sine qua non of durable success, the CRO will therefore strive to get his teams involved at a very early stage in the discussions around the AI tools and

solutions that the organization is contemplating, because it's only by being part of the solution (early on) that risk management will not be perceived as part of the problem (later on)."

Own the workforce implications. At the very outset, companies must anticipate the employment-related issues that will crop up when they use AI at scale. That's the only way they will be able to preemptively deal with the social consequences. They must identify the gaps and the surpluses that will be created in the workforce by the organization's use of AI and develop forecasts for the job functions most affected by the AI rollout. Business must upskill and reskill affected employees whenever possible in order to fill new positions, while being transparent about the implications that AI's use will have for compensation. They must also look for positions outside the organization for employees who cannot be retrained and whose jobs will disappear.



In the AI age, companies must serve as employees' providers of continuous education.

In the AI age, companies have a new moral obligation: They must serve as employees' providers of continuous education. Given the speed at which the technology changes, even newly minted graduates will soon need to upgrade their AI skills. Traditional education methods are unlikely to suffice, so companies will need to play a key role in ensuring that their workforces remain up to speed. For instance, Morocco's OCP Group—one of the world's biggest phosphate rock miners, phosphoric acid manufacturers, and phosphate fertilizer producers—has partnered with the country's Mohamed VI Polytechnique University, as well as Ecole des Mines, Ecole Polytechnique, MIT, Columbia University, and École Polytechnique Fédérale de Lausanne, to open several graduate and executive programs for managing AI in the country. It has also created a new generation of coding schools in Benguerir and Khouribga, where OCP's employees and locals can hone their digital skills.

Spread the word. Companies that use AI must educate all their employees about the technology. The level of understanding currently differs significantly across functions, which poses a major challenge. According to the BCG-MIT 2020 AI survey, 67% of employees who don't understand AI don't trust AI-based decisions, while the number is just 23% among employees who do come to grips with the technology. Realizing the need to teach employees about AI, the Spanish energy company Repsol has created an in-house data school, where the

curricula on AI and data is tailored to each employee. The company is fast becoming a data-driven organization and is trying to reskill and upskill all its employees to ensure that none of them is left behind.

Because trust influences the ability to successfully interact with AI, educating society will be a key success factor for using AI at scale. In order to increase the likelihood of obtaining a social license for AI, companies obviously need to go beyond employees and educate customers, shareholders, civil society, and the general population about the advantages and disadvantages of AI.

Help shape a world of AI. Digital technologies, spearheaded by AI, seem likely to change the world and the future of work. According to the World Economic Forum's *Future of Jobs Report 2020*: "In the next two years (by 2022), 42% of the core skills required to perform existing jobs are expected to change" and "85 million jobs will be displaced." If that happens, it will prove to be an economic, technological, and social disruption on a scale similar to that of the First and Second Industrial Revolutions of the 18th century. Instead of blindly pushing ahead with AI and focusing only on maximizing profits with the technology, business would do well to play an active role in helping governments, people, and society understand the technology and cushion its impact on humans.

Companies must work with policymakers to build a business ecosystem with AI at its core. They must grasp every opportunity to help design regulations that will expand the use of AI and mediate its impact. Regulators always find it difficult to keep pace with technology; they will welcome the help of business leaders to identify AI-related policies that balance innovation, creative destruction, its economic impact, and societal fairness. It will result in a win-win; companies will benefit from AI that conforms to future regulation, and regulators will help accelerate the pace at which business earns its social license for AI.

For a little over two decades, since the Enron, WorldCom, and Tyco scandals erupted, business has been steadily losing its credibility. That's not surprising given companies' lack of social responsibility and the environmental damage they have wrought. Many people, especially in the developed economies, have become convinced that because of its relentless profit-maximizing actions, business is close to losing its social legitimacy—the seal of approval that society confers on companies, accepting the notion that their profit-seeking actions are broadly appropriate and desirable given human values, norms, and beliefs. Despite the long history of capitalism, they believe that the day is not far off when every company will need to regain its social license if it wants to stay in business.

In the same way, corporations must realize that using AI, no matter how responsibly it is designed and how rigorously it is tested, will not be accepted automatically by society. It won't

be enough to meet government-enacted laws and regulations; that is just a hygiene factor. Business may enjoy the legal right to use AI, but it must obtain a social license from all its stakeholders if it is to deploy AI at scale.



The BCG Henderson Institute is Boston Consulting Group's strategy think tank, dedicated to exploring and developing valuable new insights from business, technology, and science by embracing the powerful technology of ideas. The Institute engages leaders in provocative discussion and experimentation to expand the boundaries of business theory and practice and to translate innovative ideas from within and beyond business. For more ideas and inspiration from the Institute, please visit our [website](#) and follow us on [LinkedIn](#) and [Twitter](#).

Authors



François Candelier

MANAGING DIRECTOR & SENIOR PARTNER; GLOBAL DIRECTOR, BCG HENDERSON INSTITUTE

Paris



Rodolphe Charme di Carlo

ALUMNUS



Steven Mills

**MANAGING DIRECTOR & PARTNER
CHIEF AI ETHICS OFFICER, BCG GAMMA**

Washington, DC

- ¹ Sean McGregor, "When AI Systems Fail: Introducing the AI Incident Database," November 2020 (<https://www.partnershiponai.org/aiincidentdatabase/>).
- ² Baobao Zhang and Allan Dafoe, "Artificial Intelligence: American Attitudes and Trends," Center for the Governance of AI, Future of Humanity Institute, University of Oxford, 2018.
- ³ "The Problem of Abortion and the Doctrine of the Double Effect," Oxford Review, 1967.

- 4 The Social License: How to Keep Your Organization Legitimate, Palgrave Macmillan.
- 5 Sema K. Sgaier, Vincent Huang, and Grace Charles, “The Case for Causal AI,” Stanford Social Innovation Review, Summer 2020.

ABOUT BOSTON CONSULTING GROUP

Boston Consulting Group partners with leaders in business and society to tackle their most important challenges and capture their greatest opportunities. BCG was the pioneer in business strategy when it was founded in 1963. Today, we work closely with clients to embrace a transformational approach aimed at benefiting all stakeholders—empowering organizations to grow, build sustainable competitive advantage, and drive positive societal impact.

Our diverse, global teams bring deep industry and functional expertise and a range of perspectives that question the status quo and spark change. BCG delivers solutions through leading-edge management consulting, technology and design, and corporate and digital ventures. We work in a uniquely collaborative model across the firm and throughout all levels of the client organization, fueled by the goal of helping our clients thrive and enabling them to make the world a better place.

© Boston Consulting Group 2023. All rights reserved.

For information or permission to reprint, please contact BCG at permissions@bcg.com. To find the latest BCG content and register to receive e-alerts on this topic or others, please visit bcg.com. Follow Boston Consulting Group on [Facebook](#) and [Twitter](#).