Explaining the Principles to Practices Gap in AI

Daniel Schiff
Georgia Institute of Technology, Atlanta, GA 30332 USA

Bogdana Rakova
Happiness Alliance, Seattle, WA 98103 USA

Aladdin Ayesh
De Montfort University, Leicester LE1 9BH, U.K.

Anat Fanti
Bar-Ilan University, Ramat Gan 5290002, Israel

Michael Lennon
George Washington University, Washington, DC 20052 USA

As artificial intelligence (AI) permeates across social and economic life, its ethical and governance implications have come to the forefront. Active debates surround AI's role in labor displacement, autonomous vehicles, military, misinformation, healthcare, education, and more. As societies collectively grapple with these challenges, new opportunities for AI to proactively contribute to social good (AI4SG) and equity (AI4Eq) have also been proposed [1], [2], such as Microsoft's AI for Earth program. These efforts highlight the potential of AI to address global challenges and help achieve targets like the United Nation's sustainable development goals (SDGs) [3]. Yet, whether AI efforts are directed explicitly at social good and equity or not, there are many barriers that stand between aspirations to be responsible and the translation of these aspirations into concrete practicalities.

In this article, we review the principles-to-practices gap, particularly in the context of corporations—the entities responsible for the majority of AI development. We offer context on the current state of responsible AI and then review six potential explanations for the gap: 1) a misalignment of incentives; 2) the complexity of AI's impacts; 3) a disciplinary divide; 4) the organizational distribution of responsibilities; 5) the governance of knowledge; and 6) challenges with identifying best practices. We offer preliminary recommendations on what frameworks and practices are needed to help close the gap and emphasize impact assessment as one promising strategy, which we discuss in the context of AI used for forest ecosystem restoration. We argue that stakeholders interested in realizing AI's potential for good should attend to these issues when proposing pathways forward.

Digital Object Identifier 10.1109/MTS.2021.3056286
Date of current version: 3 June 2021.

Minding the gap: Principles without practices

As of 2020, dozens of firms along with nongovernmental organizations and governments [4] have produced frameworks, principles, guidelines, and policies related to the responsible development and use of AI. These documents are a response to the social and ethical issues that surround AI, ranging from labor displacement and algorithmic bias to privacy and human rights. Moreover, these documents are arguably a response to breaches of law and the public trust that have continued in the last few years, some associated directly with AI. In some cases, companies have come under serious scrutiny, facing significant media attention, customer criticism, and employee petitions, walkouts, and resignations [5].

Each document typically addresses a set of social and ethical concerns, proposes principles for ethical (or responsible or trustworthy) AI in response, and in some cases offers concrete reforms or internal governance strategies. Importantly, a common focus of the documents is the presentation of a set of high-level ethical principles that guide an organization's approach to responsible AI. For example, Google's AI principles include "Be socially beneficial," "Avoid creating or reinforcing AI bias," "Be built and tested for safety," and "Be accountable to people," among other principles [6]. OpenAI discusses its focus on "Broadly Distributed Benefits," "Long-Term Safety," "Technical Leadership," and its "Cooperative Orientation" [7]. These high-level responsible AI principles are often vague and may host a multitude of possible interpretations.

Yet, some companies have established relatively clear strategies following from their principles, such as changes to training, hiring, algorithm development frameworks and tools, and governance strategies. For example, Vodafone's AI Framework provides some detail on specific actions it will take, such as adhering to its Code of Conduct and privacy commitments [8]. SAP proposes as part of its Guiding Principles an AI Ethics Steering Committee and an AI Ethics Advisory Panel [9]. IBM's Everyday Ethics for AI provides a set of recommended actions and questions for its employees to address key concerns [10]. In contrast, some principles are yet to be accompanied by clear expressions of changes to practice. For example, documents from Tieto, Futurice, and Salesforce emphasize abstract principles and commitments.

The question then is why do aspirations not translate straightforwardly into effective and responsible practices? In fact, the process of translation is neither automatic nor obvious. According to Mittelstadt [11], "norms and requirements can rarely be logically deduced… without accounting for specific elements of the technology, application, context of use, or relevant local norms." Barring practical guidance and absent "empirically proven methods… in real-world development contexts," claims of responsible AI may amount to no more just that—claims.

In the best case, companies may still be in the process of working out the details or may have communicated their intended strategies in other venues. Nevertheless, remaining at the "mission statement" level and the lack of practical detail are worrisome, and engaging in responsible AI is, in any case, no simple task. The task of closing the principles-to-practices gap is critical for the future of AI, worthy of attention by companies developing AI, by governments that might procure and deploy AI systems, and by other stakeholders and the public more broadly. To assist with this effort, this article reviews possible explanations for the gap, depicted in Figure 1, along with preliminary recommendations. The six explanations we review are: 1) the incentives dilemma; 2) the complexity of AI's impacts; 3) the disciplinary divide; 4) the many hands problem; 5) the governance of knowledge; and 6) the over-abundance of tools.

The incentives dilemma

The first and perhaps most widely discussed explanation is that the values, motivations, and incentives that guide firms are not sufficiently aligned with responsible uses of AI. The strongest criticisms deeply impugn the motives of firms. For example, Greene et al. [12] argue that companies attempt to shift responsibility onto designers and experts to minimize scrutiny of business decisions. On this account, companies may be strategically promoting their principles merely to ameliorate customer trust and reputational concerns, thereby engaging in "ethics washing," "ethics shirking," "ethics shopping," or similar [13]. In this way, firms can appear actively engaged regarding AI's ethical risks in the public eye, but while framing issues so as to minimize genuine accountability.

Alternatively, organizations need not have the motives stated above, but may still be driven by an

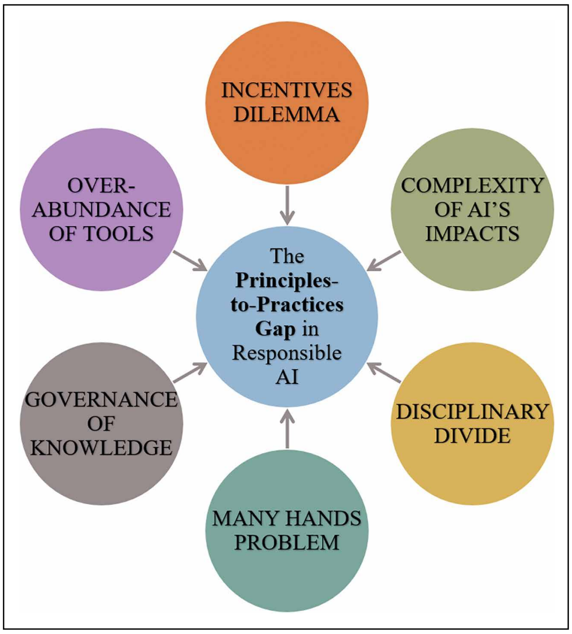

Figure 1. Six explanations for the principles-to-practices gap.

economic or financial logic that places profit-seeking and the mitigation of legal risks above pro-social concerns. When these motives shape the institutional context, firms may simply lack the infrastructure, processes, and culture to responsibly weigh social and ethical concerns against economic ones. For example, Hagendorff [14] argues that when companies are driven by an economic logic, "engineers and developers are neither systematically educated about ethical issues, nor are they empowered, for example by organizational structures, to raise ethical concerns."

While some of these deeper criticisms may be true in part or for some organizations, we think a more multifaceted and charitable interpretation [15], [16] is both appropriate and likely to be beneficial toward seeking positive change. Indeed, significant literature has challenged the notion that firms or individuals within firms are characterized by purely "rational" economic behavior [17], [18]. We argue instead that organizations are best understood not as monolithic single actors, but as multiple coordinating and competing coalitions of individuals [19]. Individuals within a single organization may have multiple and conflicting preferences and roles. Organizational motives should therefore be considered a complex composition of genuine ethical concern, economic logic, signaling and framing strategies, and efforts to promote change both internally (i.e., within the firm) and externally (e.g., in regulation and public opinion) [20].

Regardless, the tension between economic and pro-social motives, and the incentive structure that this produces should be taken seriously. A key question then is what policies and practices can shape and enable firms to weigh competing values in a responsible way. There are many possible pathways here, both internal and external to organizations: industry collective self-regulation and formal governmental regulation, shifts in public and societal norms, changes to education and training of engineers, and adoption of a broader focus on stakeholder benefits along the lines on B Corps, benefit corporations, and a triple bottom line orientation. Organizations that already have a strong commitment to social good can also consider "value levers" [21] that facilitate responsible practice, as well as some of the recommendations we suggest below.

The complexity of AI's impacts

A second explanation for the principles-to-practices gap is that AI's impacts on well-being—positive or negative—are more complex than is sometimes assumed. Site-based research has identified that engineers are often focused on single products and the physical harm they may cause rather than broader kinds of social, emotional, or economic harms [22]. Even as conversations surrounding responsible AI increase, most work centers around a relatively small subset of issues, most often bias [23] and transparency [24] in particular AI models. This approach involves exposing and then attempting to mitigate bias in algorithms as well as trying to improve interpretability or explainability given the black-boxed nature of certain AI models which can make decision-making processes and their outcomes opaque. Other commonly emphasized issues include privacy, reliability, and safety.

Yet, these prominent issues most familiar to engineers and dominant in AI ethics research constitute only a subset of social and ethical risks and impacts related to AI. In place of such a narrow focus, AI should be understood to impact a wide variety of aspects of human and societal well-being, such as human rights, inequality, human–human relationships, social and political cohesion, psychological health, and more. AI can also impact natural ecosystems and animal life, impacts that are arguably both intrinsically and instrumentally important.

Moreover, many of these harms do not arise in a straightforward way from a single AI product, but from many AI systems influencing human social and economic life together over time. For example, algorithms on social media designed to steer consumers to entertaining video clips have also led to so-called filter bubbles that may foster political polarization, misinformation and propaganda, targeting of minority groups, and election interference [25]. AI as instantiated in autonomous vehicles has potentially massive implications for physical infrastructure, energy and environment, traffic fatalities, work productivity, urban design, and unemployment [26].

In short, addressing AI principles in full seriousness requires an expansive scope of attention to the full set of issues influencing human well-being. This requires looking well beyond a narrow set of topics such as bias, transparency, privacy, or safety and treating them as independent issues. Instead, the full range of topics and their complex interdependencies need to be understood. However, such a task can be enormously difficult.

The disciplinary divide

Another challenge related to AI's complexity is the plurality of professional disciplines with roles to play in shaping responsible AI, beyond engineers and computer scientists. Indeed, discourse on responsible AI has been advanced by ethicists, historians and philosophers of technology, journalists, policy analysts, political decision-makers, other social scientists, members of the public, and more. Yet these diverse stakeholders may differ in their technical and ethical education, their framing of problems and solutions, their attitudes and values toward responsible AI, and their norms of communication.

Consider attempts to apply the principle of fairness in attempting to minimize bias. Arguably, a thoughtful AI engineer today might identify a normative principle like "fairness," specified in a corporate responsible AI policy, pick a plausible fairness metric to instantiate it (noting there are ineliminable tradeoffs between different metrics [27]), apply it, and communicate these decisions transparently [28].

Yet approaching social issues like bias and fairness too narrowly leads to what Selbst et al. (2018) call category or abstraction errors [29]. Computer scientists developing AI systems may fail to consider how an AI system will be implemented in different social contexts, influence human behavior in those contexts, or lead to long-term ripple effects, all of which can threaten the assumptions on which the AI system is built. This is especially difficult as predicting a technology's usage and impact is known by historians of science and technology to be difficult [30].

Consider an algorithm used to inform a judge's decision about criminal sentencing. An algorithm designed and trained on test data from one jurisdiction may translate poorly to another. It may influence the judge's decisions in unexpected ways that are not often accounted for [31], as a judge may overtrust or undertrust the algorithm, or even hold values contrary to those reflected in the algorithm. The consequences for criminal justice outcomes when such a system is used in complex contexts is unclear, and may feed back in unexpected or problematic ways if an AI is trained on data the systems has itself helped to generate. To reiterate, there are many questions about responsible AI that cannot be straightforwardly addressed with a narrow technical lens.

On the other hand, social scientists may bring a lens that faces an inverse problem to that faced by engineers: frameworks for considering social and ethical consequences of AI more in line with the thinking of social scientists can be sufficiently broad but also unhelpfully complex and vague, making them difficult to translate into practices. For example, ethicists recognize that concepts like justice are complex, while political scientists know that values surrounding justice are politically contested [32]. Yet AI engineers must define some measure of justice to implement it.

In addressing issues like inequality, social scientists may propose large structural changes to economic and social systems, some of which are difficult to achieve (e.g., reforming the motives of corporations) and or even feel far-fetched (e.g., changing the structure of capitalism), changes which may be significantly outside of the scope of control of AI engineers. Conceptions of AI based on sweeping, overly futuristic, or unrealistic generalizations may also be unhelpful. In the best case, it is difficult to resolve the awkwardness of attempting to apply purely technical or social fixes to fundamentally sociotechnical problems. Something is lost in translation.

The many hands problem

It is clear that responsibly designing and applying AI is therefore both a technical challenge and a social one (implicating social, economic, and policy questions). For example, creating a facial recognition system for policing that minimizes racial bias (by some technical measure) is inseparable from questions on the legitimacy of the use of that system in a particular social and policy setting. However, the question of distributing accountability for addressing these multi-faceted choices remains open and contested.

Engineers and computer scientists may see their responsibility as focused on the quality and safety of a particular product rather than on larger -scale social issues and maybe unaware of this wider universe of implications [33]. Business managers and companies may view themselves as having a primary duty to shareholders which involves emphasizing revenue generation above broader societal impacts. This potentially creates holes in responsibility for addressing the key impacts of AI.

In addition to uncertainty regarding one's scope of professional accountability, engineers and computer scientists who focus on the design of systems may have limited influence within their organizations, especially if they focus on a very narrow aspect of AI system development. They may expect top executives, liability officers, or corporate social responsibility staff to assess broader social and ethical issues. Even social scientists and ethicists tapped specifically for these issues may find themselves similarly handicapped, perhaps in an external advisory role without a real say. The result is the "many hands" problem, where responsibility for managing AI is distributed and muddled [34].

This question of responsibility is intimately tied with how organizations structure job responsibilities and workflow related to AI. A major concern is that computer scientists and engineers may be functionally separated from other staff roles likely to be tasked with thinking about an AI system's broader implications—such as higher-level business managers, the C-suite, and corporate social responsibility and compliance staff.

Such an organizational structure exacerbates the disciplinary divide, as inevitably, functional separation of technical and nontechnical experts in organizations limits the potential to communicate effectively, understand issues robustly, and respond to considerations of AI's impacts on well-being [21]. Note that such a problem may persist even for companies that have proposed AI ethics advisory or governance boards, if functional separation remains. Figuring out how to distribute responsibility for AI's impacts may involve challenging long held assumptions and re-envisioning organizational processes and roles, no easy task.

The governance of knowledge

To that end, consideration of knowledge management practices may constitute one avenue toward reimagining organizational processes. An effective knowledge management system requires the creation of knowledge, its storage and retrieval, transfer of tacit and explicit knowledge within organizations, and the how-to of application itself [35]. When an organization does not have a proper methodology for governing knowledge about AI's possible ethical implications along any of these dimensions, the likelihood of responsible and effective AI practice decreases.

For example, organizations that fail to consider the complexity of AI and its possible impacts broadly enough are unlikely to create the knowledge necessary to guide decision-making and action. Further, even if such knowledge is collected by some individuals or teams within an organization, it must be effectively stored (e.g., in documents, and organizational memory) in a way that is easily retrievable by the appropriate teams. Moreover, such knowledge must be carefully translated during retrieval and transfer, so that the teams that need to make decisions are able to process the information and learn. The disciplinary divide suggests the transfer process cannot be taken for granted. Finally, organizations need processes for aggregating preferences and making decisions using this information, as well as tools to guide action, which we discuss in the next section.

Without these systems in place, various teams within the organization and in different phases of development may fail to become aware of possible ethical dimensions or problems, and may be incapable of effective decision-making and informing appropriate stakeholders. For example, if there is no method of "flagging" specific issues other than the "usual" suspects in software development, such as bugs and incomplete versions, various important questions might be disregarded or simply left unasked.

The challenge of governing knowledge, therefore, runs across organizational structure, culture, and technology [36]. Structural elements include software and engineering development processes [37], their integration with larger business development cycles, and organizational incentives and hierarchy. Organizational culture can be more intangible and intractable, and is shaped by corporate vision, employee interaction, and the values that permeate business and engineering practices. Finally, technology (construed broadly) addresses the tools used to create, store, transfer, and apply knowledge. Carefully chosen tools must be in place to support each process toward better governing knowledge.

Addressing these issues is a long-standing organizational challenge, but it reinforces the need for thinking holistically. An organizational statement about AI ethics principles or a tool for mitigating algorithmic bias may be a valuable starting point, but these components must be conceived of and incorporated into a broader organizational governance system. For example, new staff positions and processes can be created that challenge traditional disciplinary silos. Organizational leadership can incentivize ethical concern and interdisciplinary interaction to build a better culture, and the stages of software engineering can be mapped to knowledge management practices [38] that are attuned explicitly to issues of responsible AI.

The abundance of tools

The final explanation for the principles-to-practices gap that we discuss is, ironically, the potential over-proliferation of solutions. Indeed, an increasing number of researchers who have noticed the principles-to-practices gap have begun proposing strategies [39], often aimed at AI practices within industry. These proposals include changes to software mechanisms (such as audit trails), hardware mechanisms (such as secure hardware enclaves), and institutional mechanisms (such as red team exercises) [37]. This work highlights that it is not only technical practices that must adapt, but also organizational ones.

Among the most comprehensive work assessing the principles-to-practices gap is the 2019 review by Morley et al. [40], which systematically explores existing responsible AI tools and methodologies mapped against seven components of the AI development lifecycle: 1) business and use-case development; 2) design phase; 3) training and test data procurement; 4) building; 5) testing; 6) deployment; and 7) monitoring. They identify 106 tools and methodologies.

Some such methods are relatively narrower in scope, such as those surrounding explainable AI [41], bias [42], or procurement (e.g., the AI-RFX Procurement Framework). Other methodologies adopt a broader scope of focus, including impact assessments like the ISO 26000 Framework for Social Responsibility [43] and the IEEE 7010 Recommended Practice for Assessing the Well-Being Implications of Autonomous and Intelligent Systems [44]. Still more relevant methods and approaches proposed for responsible AI also come from outside of the AI domain and include privacy-by-design [45], value-sensitive design [46], the Responsible Research and Innovation (RRI) approach [47], and numerous others.

Yet while creating more and better such tools and methodologies is a worthy pursuit, an over-abundance problem makes it difficult for individuals to sort through and assess the utility of a given tool, or to weigh it against the many other available tools. As a result, individuals and organizations may fail to take advantage of the useful tools and methodologies that already exist. Even tools that do exist have arguably not been tested sufficiently to demonstrate which are most effective and in which contexts [11]. A further problem is that many tools and methodologies do not contain sufficient instructions to apply, customize, or troubleshoot, much less in particular organizational contexts and use cases [40].

Closing the principles-to-practices gap

Criteria of an effective framework for responsible AI

Given the proposed explanations above, how can we begin to close the principles-to-practices gap? We think an overarching framework for responsible AI development can help to streamline practice and leverage existing tools and methodologies. What would be the desiderata of such a framework for responsible AI?[1] As a starting point and based on the identified gaps, we suggest the following (see Figure 2):

- *Broad*: It should consider AI's impacts expansively, across many different ethical issues and aspects of social and economic life.
- *Operationalizable*: It should enable users to cast conceptual principles and goals into specific strategies that can be implemented in real-world systems.
- *Flexible*: It should be able to adapt to a wide variety of AI systems, use cases, implementation contexts, and organizational settings.
- *Iterative*: It should be applied throughout the lifecycle of an AI system and repeatedly.
- *Guided*: It should be easy to access and understand, with sufficient documentation.
- *Participatory*: It should incorporate the perspectives and input from stakeholders from a range of disciplines as well as those that may be impacted by the AI system, especially the public.

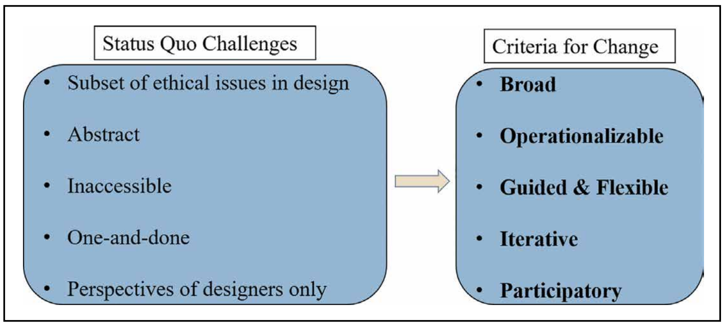

Figure 2. Criteria of a framework for responsible AI development.

[1]Akin to what Dignum calls "Ethics in Design" [48].

A framework that meets these criteria balances the need for technical specificity with an equally important need for conceiving of AI's impacts on human well-being in their full breadth and complexity, understanding we cannot fully predict all of AI's possible ramifications. That is, while some prominent strategies for responsible AI assume there are only a small set of issues to address, such as bias, transparency, and privacy, we have argued that AI's impacts are more complex.

Impact assessments are one promising strategy toward achieving many of these criteria [49]. Impact assessments have been used historically in human rights [50], in regulatory contexts [51], and more recently to study the impact of AI or algorithms [52], [53]. Measuring impact has several benefits: it enables monitoring of risks, allows for a more secure environment for investment, promotes accountability and transparency, and overall enhances the prospects of pro-social innovation [3].[2] We focus on the recently published IEEE 7010 standard as an exemplar [44], [53] created specifically to assess AI's impacts on human well-being.[3] We argue below that impact assessments like the well-being impact assessment from the IEEE 7010 standard could be adopted by companies pursuing responsible AI development as well as incorporated into the institutions which train future practitioners.

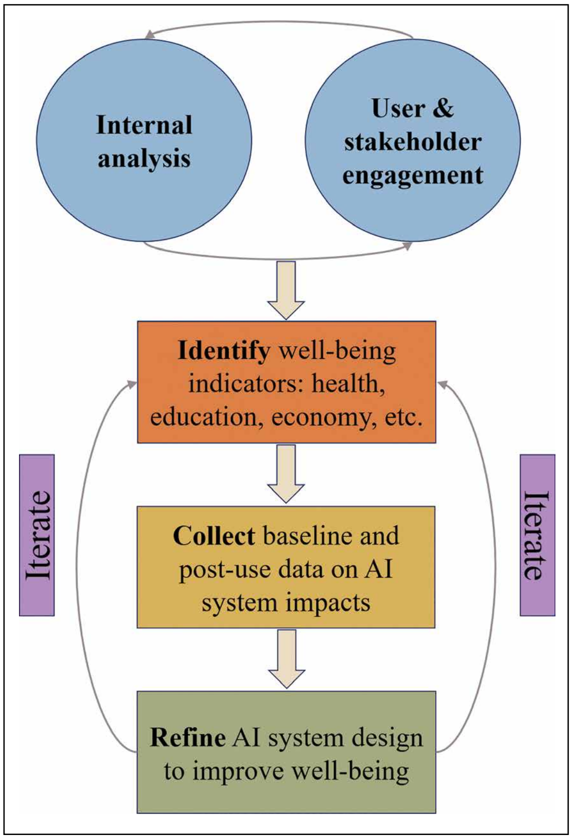

Figure 3. Well-being impact assessment process.

Impact assessments for responsible AI

According to the IEEE 7010 standard, a well-being impact assessment is an iterative process that entails: 1) internal analysis; 2) user and stakeholder engagement; and 3) data collection, among other activities (see Figure 3). The internal analysis involves broadly assessing the possible harms, risks, and intended and unintended users and uses of an AI system. Here, developers and managers of an AI system carefully consider a wide range of an AI system's potential impacts on human well-being, not limited to prominent issues like privacy, bias, or transparency.

Critically, assessing impacts requires not just speculating about impacts, but also measuring them. Therefore, the user and stakeholder engagement stages of the assessment include learning from users of AI systems as well as others more indirectly impacted to determine how the system impacts their well-being. When developers have access to users, this may include asking them about possible or actual psychological impacts, economic impacts, changes to relationships, work-life balance, or health. This is again in contrast to strategies that focus solely on technical fixes to issues like bias or privacy during the design stage alone and fail to account for the broader universe of well-being implications.

Beyond initial stakeholder engagement, data collection based on the continuous assessment of the identified possible impact areas is key. Data can be collected through user surveys, focus, groups, publicly-available data sources, or directly as system outputs. In sum, we propose that adherence to—and rigorous documentation of—impact assessments can contribute to the continuous improvement of AI systems by ensuring that organizations are better able to understand and address AI's many impacts on human well-being.

Importantly, not all tools entitled "impact assessment" meet our definition. Many existing tools consider only a small scope of possible impacts.

[2] Also see [3] for a discussion of challenges with impact measurement.
[3] While the authors of this article were involved in helping to develop IEEE 7010, this article reflects the individual views of the authors and not an official position of the IEEE.

Some fail to measure impacts at all, instead focusing on anticipating impacts *assumed* to be important and applying best practices to avoid associated harms, such as bias. Inversely, some tools that are not labeled "impact assessments" might be classified as such under our definition, such as the European Commission's Ethics Guidelines for Trustworthy AI [54]. Notably, some frameworks have been proposed by the public sector (i.e., governments) and others by nongovernmental organizations and companies.

We are hopeful that more scholars and organizations focused on responsible uses of AI will adopt an impact assessment approach or alternative frameworks that meet the criteria identified above. A key aspect of enabling effective impact assessment or other frameworks is creating supportive structures, such as by adopting new practices in institutions of higher education as well as organizational changes in firms. We turn to these issues briefly.

Supportive practices in higher education and industry

Educational systems have undertaken meaningful efforts aimed at increasing ethical sensitivity and decision-making, but have not yet made the changes needed to support responsible AI practice. Of around 200 AI/ML/data science courses reviewed by Saltz et al. (2019), little more than 1 in 10 mentioned ethics in their syllabus or course description. Those that did focused overwhelmingly on bias, fairness, and privacy [55]. Although courses focused specifically on AI ethics cover a wider set of issues including consequences of algorithms, technically tractable issues like bias and privacy are still dominant [56]. We suggest that AI ethics education focus not solely on a few prominent or technically tractable issues nor on general awareness-building alone, but also on impact assessment as an overarching framework to understand AI's impacts on human well-being. Further, educational institutions should seek to serve diverse groups of students through coursework, extracurricular clubs, and even contests to encourage cross-disciplinary learning and practice.

Companies developing or deploying AI should move toward the integration of technical and nontechnical teams rather than functional separation of roles, for reasons discussed previously. These integrated teams could include technical developers as well as other individuals tasked with considering impacts of an AI system who may have social science, humanities, business, law, or ethics expertise, or who can represent user and stakeholder interests effectively. Such a change requires establishing new organizational and knowledge management practices that are integrated with traditional engineering and software lifecycles. Already, organizations have proposed including a residential nontechnical thinker tasked with responsible AI—an "ethics engineer" or "responsible AI champion" [57]. Organizations could also engage in interdisciplinary and interdepartmental cross-training, using red team exercises [37] or hypothetical case studies that draw on the impact assessment approach.

In summary, we have argued that impact assessments are a promising strategy to address the gaps between principles and effective practices for responsible AI. However, applying an impact assessment might feel like an abstract exercise to those who have not done it. To demonstrate how closing the principles-to-practices gaps with an impact assessment might occur, we move now to a demonstrative example case study.

Case study: Impact assessments to support responsible AI for forest ecosystem restoration

This section provides an example case study of AI's use in forest ecosystem restoration efforts to demonstrate the ideas and approach proposed in this article. Further, it helps to highlight future work needed for practical adoption of impact assessments or similar responsible AI frameworks.

Case study background

Forest ecosystem restoration is essential for many reasons. Forests have the most species diversity on the planet, with some 80% of land-based species. Forests also reduce the risk of natural disasters such as floods, droughts, and landslides and help protect watersheds [58]. Further, forests are critical for mitigating land-based carbon emissions by increasing carbon sequestration, critical for climate change prevention goals [59]. Project Drawdown, for example, has calculated that the restoration and protection of tropical forests could lead to 61.23 gigatons of carbon reduction by 2050 [60].

Achieving these goals requires the restoration of forest ecosystems through the cultivation of trees, known as afforestation [61]. Applied afforestation projects typically involve three stages - planning, execution, and monitoring of ecosystem restoration. Several AI technologies have been used in afforestation efforts and their use is increasing. During planning, AI systems have been used to predict forest carbon sequestration potential through the use of satellite and drone image data [62], [63]. AI can also facilitate execution of afforestation through computer vision algorithms used in identifying appropriate planting sites, monitoring plant health, and analyzing trends [64]. Lastly, in the monitoring stage of restoration projects, AI can be used to identify where deforestation may have been conducted illegally [65], [66], as well as assess risks due to fire, disease, insects, or other causes [67].

Current challenges

According to international governance efforts like the UN SDGs, the UN Forum on Forests, Agenda 21, and the Future We Want (the outcome document of the Rio+20 Conference) there is a need for holistic, multi-stakeholder engagement to address forest ecosystem restoration adequately [68], [69]. This is due to the existence of multiple groups with critical interests in forest ecosystems. For example, local communities and businesses may engage in harvesting timber and farming as part of the local economy, while natives living off the land may depend on hunting animals and harvesting plants, and policymakers must also worry about carbon sequestration and climate change efforts.

Though the goals of these groups are not always in conflict, they can bring different perspectives and have competing priorities. Therefore, AI-driven systems used for afforestation that do not take into account these "multiple ecological, economic, social and cultural roles" important to various stakeholders [58] may lead to blind spots and unintended harms. For example, an AI system that uses imaging data to determine carbon sequestration potential could optimize climate change goals in a narrow sense, but fail to account for social-ecological aspects of the land important to indigenous groups, or ignore endangered species important to conservationists.

As a result, carbon sequestration targets optimized in the short term could fall short in the long term as afforestation progress fails to translate into a sustainably managed multi-stakeholder effort. Failing to develop and implement AI systems for ecosystem restoration in a participatory fashion is thus an example of how the laudable goal of improving environmental well-being can fail to translate into responsible and effective practices.

Applying impact assessment

As Rolnick et al. [70] state, that "Each stakeholder has different interests, and each often has access to a different portion of the data that would be useful for impactful [machine learning] applications. Interfacing between these different stakeholders is a practical challenge for meaningful work in this area." Therefore, stakeholders, such as landowners, policymakers, local communities, and others need to have a voice in the application of AI to forest ecosystem restoration. Impact assessment can help address this gap (see Figure 4).

As discussed in the "Disciplinary divide" section, the impact assessment process involves a broad internal analysis by the organizations developing AI systems for forest ecosystem restoration. For an AI developer engaged in forest ecosystem restoration, this would involve trying to understand the variety of possible stakeholders and intended or unintended impacts of their products. A company that develops AI to identify target areas for afforestation given carbon sequestration potential might begin to identify previously unrecognized impacts on species diversity, the local economy, and the well-being of native groups.

To have a more accurate understanding of broader impacts—as well as to build consensus among stakeholders—the company would then begin the user and stakeholder engagement process. Critically, this would involve soliciting the input of

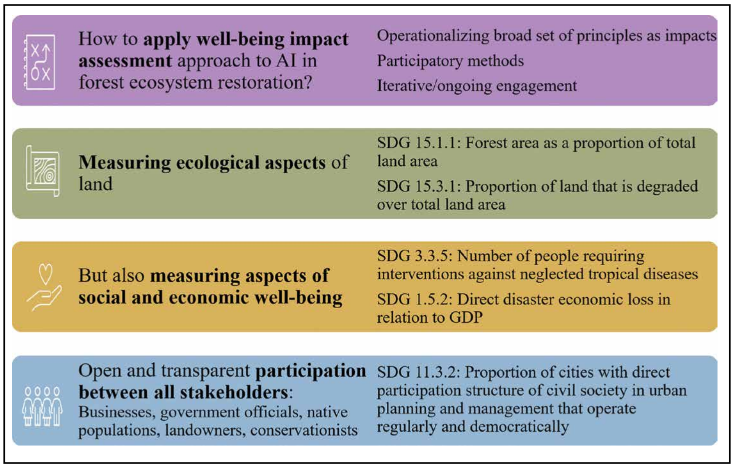

Figure 4. Applying impact assessment to AI in forest ecosystem restoration.

the numerous stakeholders, such as conservation groups, landowners, scientists, government officials, local businesses, and native populations. A method like participatory action research or other participatory design methods [71] could be used to facilitate this engagement.

This process, which should be ongoing and iterative throughout the management of the forest ecosystem, should surface a number of clear concerns about possible implications of the afforestation efforts. For example, the company may have originally been optimizing a target through their AI system such as SDG indicators 15.1.1, "Forest area as a proportion of total land area," or 15.3.1, "Proportion of land that is degraded over total land area." However, the impact assessment process should lead to the flexible identification of new indicators critical to having a broader understanding of the relevant social, economic, and ecological context, such as SDG indicators 3.3.5, "Number of people requiring interventions against neglected tropical diseases," or 1.5.2, "Direct disaster economic loss in relation to global gross domestic product." These new indicators would therefore operationalize possible dimensions and impacts of the forest ecosystem management effort such as diseases and natural disasters as specific measurable indicators.

Finally, the company would have new ways to apply their deeper understanding of the well-being implications of their AI system. One such approach could be embedding expert domain knowledge garnered from the participatory process into the architecture of the AI system itself [72]. For example, an AI system that previously optimized carbon sequestration potential as part of its objective function could incorporate new data regarding tropical diseases or natural disasters as additional constraints or targets in the optimization of its model. Finally, though the impact assessment process and identification of solutions could initially feel unfamiliar and complex, the company would gradually develop best practices and guidance toward a more responsible application of its AI system for forest ecosystem restoration.

Though this case study—the use of AI for forest ecosystem restoration—is based on real uses of AI and associated real-world challenges, the specific actions taken by the company and indicators suggested are hypothetical, and the impact assessment only responds to some of the possible challenges faced. We do not mean to suggest that there are not companies or governments already taking thoughtful approaches to multistakeholder governance in this area. However, to the best of the authors' knowledge, current sustainability efforts have not yet incorporated impact assessments of AI-driven technological solutions applied to ecosystem restoration. We hope this case study helps to demonstrate how impact assessments are a promising tool to close the principles-to-practices gap toward responsible AI.

In this article, we reviewed and synthesized explanations for the gap between high-level responsible AI principles and the capacity to implement those principles in practice. These explanations are neither mutually exclusive nor exhaustive. Scholars should continue to probe—conceptually and empirically—what barriers stand in the way of responsible AI practice within organizations. We suggest this research agenda will benefit from collaboration between computing, engineering, organizational, business, and other scholars. As the AI community moves forward in its ambitions to promote socially responsible AI, attention to the nature of the principles-to-practices gap can help guide the identification of best practices and promising solutions such as impact assessment. ■

References

[1] J. Cowls et al., "Designing AI for social good: Seven essential factors," *Social Sci. Res. Netw.*, Rochester, NY, USA, SSRN Scholarly Paper 3388669, May 2019.

[2] A. Manjarres et al., "AI4Eq: for a true global village not for global pillage," in *Proc. Eur. Conf. AI*, Aug. 2020. [Online]. Available: https://underline.io/events/24/sessions/464/lecture/2212-ai4eq-for-a-true-global-village-not-for-global-pillage

[3] S. Pitt et al. (Aug. 2020). *No App Is an Island: Collective Action and Sustainable Development Goal-Sensitive Design*. [Online]. Available: https://underline.io/events/24/sessions/464/lecture/2215-no-app-is-an-island-collective-action-and-sustainable-development-goal-sensitive-design

[4] A. Luccioni and Y. Bengio, "On the morality of artificial intelligence [Commentary]," *IEEE Technol. Soc. Mag.*, vol. 39, no. 1, pp. 16–25, Mar. 2020, doi: 10.1109/MTS.2020.2967486.

[5] S. Shane and D. Wakabayashi, "'The business of war': Google employees protest work for the Pentagon," *New York Times*, Apr. 2018.

[6] Google, "AI at Google: Our principles," Google, Mountain View, CA, USA, Tech. Rep., Jun. 2018. [Online]. Available: https://www.blog.google/technology/ai/ai-principles/

[7] OpenAI, "OpenAI charter," OpenA, San Francisco, CA, USA, Tech. Rep., Apr. 2018.

[8] Vodafone Group Plc, "Vodafone's AI framework," Vodafone, Newbury, U.K., Tech. Rep., Jun. 2019.

[9] SAP, "SAP's guiding principles for artificial intelligence," SAP, Waldorf, Germany, Tech. Rep., Sep. 2018.

[10] IBM, "Everyday ethics for artificial intelligence," IBM, New York, NY, USA, Tech. Rep., Sep. 2018.

[11] B. Mittelstadt, "Principles alone cannot guarantee ethical AI," *Nature Mach. Intell.*, vol. 1, no. 11, pp. 501–507, Nov. 2019, doi: 10.1038/s42256-019-0114-4.

[12] D. Greene, A. L. Hoffmann, and L. Stark, "Better, nicer, clearer, fairer: A critical assessment of the movement for ethical artificial intelligence and machine learning," in *Proc. Crit. Ethical Stud. Digit. Social Media*, 2019, p. 10.

[13] K. S. Gill, "Prediction paradigm: The human price of instrumentalism," *AI Soc.*, vol. 35, no. 3, pp. 509–517, Sep. 2020, doi: 10.1007/s00146-.020-01035-6.

[14] T. Hagendorff, "The ethics of AI ethics: An evaluation of guidelines," *Minds Mach.*, vol. 30, no. 1, pp. 99–120, Feb. 2020, doi: 10.1007/s11023-020-09517-8.

[15] R. Abebe et al., "Roles for computing in social change," in *Proc. Conf. Fairness, Accountability, Transparency (FAT)*. Barcelona, Spain: ACM, Jan. 2020, pp. 252–260, doi: 10.1145/3351095.3372871.

[16] E. Bietti, "From ethics washing to ethics bashing: A view on tech ethics from within moral philosophy," in *Proc. Conf. Fairness, Accountability, Transparency (FAT)*. Barcelona, Spain: ACM, Jan. 2020, pp. 210–219, doi: 10.1145/3351095.3372860.

[17] V. L. Nielsen and C. Parker, "Mixed motives: Economic, social, and normative motivations in business compliance," *Law Policy*, vol. 34, no. 4, pp. 428–462, 2012, doi: 10.1111/j.1467-9930.2012.00369.

[18] L. J. Spence and R. Rutherfoord, "Social responsibility, profit maximisation and the small firm owner-manager," *J. Small Bus. Enterprise Develop.*, vol. 8, no. 2, pp. 126–139, Jun. 2001, doi: 10.1108/EUM0000000006818.

[19] J. G. March, "The business firm as a political coalition," *J. Politics*, vol. 24, no. 4, pp. 662–678, Nov. 1962, doi: 10.1017/S0022381600016169.

[20] D. Schiff et al., "What's next for AI ethics, policy, and governance? A global overview," in *Proc. AAAI/ACM Conf. AI, Ethics, Soc. (AIES)*. New York, NY, USA: ACM, Feb. 2020, pp. 153–158, doi: 10.1145/3375627.3375804.

[21] K. Shilton, "Values levers: Building ethics into design," *Sci., Technol., Hum. Values*, vol. 38, no. 3, pp. 374–397, May 2013, doi: 10.1177/0162243912436985.

[22] V. Vakkuri et al., "Ethically aligned design of autonomous systems: Industry viewpoint and an empirical study," 2019, *arXiv:1906.07946*. [Online]. Available: http://arxiv.org/abs/1906.07946

[23] R. K. E. Bellamy et al., "AI fairness 360: An extensible toolkit for detecting, understanding, and mitigating unwanted algorithmic bias," 2018, *arXiv:1810.01943*. [Online]. Available: http://arxiv.org/abs/1810.01943

[24] A. Adadi and M. Berrada, "Peeking inside the black-box: A survey on explainable artificial intelligence (XAI)," *IEEE Access*, vol. 6, pp. 52138–52160, 2018, doi: 10.1109/ACCESS.2018.2870052.

[25] D. DiFranzo and K. Gloria-Garcia, "Filter bubbles and fake news," *XRDS: Crossroads, ACM Mag. Students*, vol. 23, no. 3, pp. 32–35, Apr. 2017, doi: 10.1145/3055153.

[26] S. A. Bagloee et al., "Autonomous vehicles: Challenges, opportunities, and future implications for transportation policies," *J. Modern Transp.*, vol. 24, no. 4, pp. 284–303, Dec. 2016, doi: 10.1007/s40534-016-0117-3.

[27] A. Chouldechova,"Fair prediction with disparate impact: A study of bias in recidivism prediction instruments," 2017, *arXiv:1703.00056*. [Online]. Available: http://arxiv.org/abs/1703.00056

[28] D. Leben, "Normative principles for evaluating fairness in machine learning," in *Proc. AAAI/ACM Conf. AI, Ethics, Soc. (AIES)*. New York, NY, USA: ACM, Feb. 2020, pp. 86–92, doi: 10.1145/3375627.3375808.

[29] A. D. Selbst et al., "Fairness and abstraction in sociotechnical systems," *Social Sci. Res. Netw.*, Rochester, NY, USA, SSRN Scholarly Paper 3265913, Aug. 2018.

[30] L. Winner, "Do artifacts have politics?" *Daedalus*, vol. 109, no. 1, pp. 121–136, 1980.

[31] J. Salgado and C. Fernandez, "A wide human rights approach to AI regulation in Europe," in *Proc. Eur. Conf. AI*, Aug. 2020. [Online]. Available: https://underline.io/events/24/sessions/464/lecture/2214-a-wide-human-rights-approach-to-artificial-intelligence-regulation-in-europe

[32] R. Binns, "Fairness in machine learning: Lessons from political philosophy," in *Proc. Conf. Fairness, Accountability Transparency*. McKees Rocks, PA, USA: PMLR, Jan. 2018, pp. 149–159.

[33] W. Faulkner, "'Nuts and bolts and people': Gender-troubled engineering identities," *Social Stud. Sci.*, vol. 37, no. 3, pp. 331–356, Jun. 2007, doi: 10.1177/0306312706072175.

[34] L. Floridi, "Distributed morality in an information society," *Sci. Eng. Ethics*, vol. 19, no. 3, pp. 727–743, Sep. 2013, doi: 10.1007/s11948-012-9413-4.

[35] M. Alavi and D. E. Leidner, "Review: Knowledge management and knowledge management systems: Conceptual foundations and research issues," *MIS Quart.*, vol. 25, no. 1, pp. 107–136, Mar. 2001, doi: 10.2307/3250961.

[36] A. H. Gold, A. Malhotra, and A. H. Segars, "Knowledge management: An organizational capabilities perspective," *J. Manage. Inf. Syst.*, vol. 18, no. 1, pp. 185–214, May 2001, doi: 10.1080/07421222.2001.11045669.

[37] M. Brundage et al., "Toward trustworthy AI development: Mechanisms for supporting verifiable claims," 2020, *arXiv:2004.07213*. [Online]. Available: http://arxiv.org/abs/2004.07213

[38] J. B. de Vasconcelos et al., "The application of knowledge management to software evolution," *Int. J. Inf. Manage.*, vol. 37, no. 1, pp. 1499–1506, Feb. 2017, doi: 10.1016/j.ijinfomgt.2016.05.005.

[39] J. C. Newman, "Decision points in AI governance: Three case studies explore efforts to operationalize AI principles," CLTC, Berkeley, CA, USA, White Paper, May 2020. [Online]. Available: https://cltc.berkeley.edu/wp-content/uploads/2020/05/Decision_Points_AI_Governance.pdf

[40] J. Morley et al., "From what to how: An initial review of publicly available AI ethics tools, methods and research to translate principles into practices," *Sci. Eng. Ethics*, vol. 26, no. 4, pp. 2141–2168, Dec. 2019, doi: 10.1007/s11948-019-00165-5.

[41] R. Guidotti et al., "A survey of methods for explaining black box models," *ACM Comput. Surv.*, vol. 51, no. 5, pp. 93:1–93:42, Aug. 2018, doi: 10.1145/3236009.

[42] T. Bolukbasi et al., "Man is to computer programmer as woman is to homemaker? Debiasing word embeddings," in *Advances in Neural Information Processing Systems 29*, D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, Eds. New York, NY, USA: Curran Associates, Inc., 2016, pp. 4349–4357.

[43] W.-W. Zhao, "Improving social responsibility of artificial intelligence by using ISO 26000," in *Proc. IOP Conf. Ser., Mater. Sci. Eng.*, vol. 428, Oct. 2018, Art. no. 012049, doi: 10.1088/1757-899X/428/1/012049.

[44] *IEEE Recommended Practice for Assessing the Impact of Autonomous and Intelligent Systems on Human well-Being,* IEEE Standard 7010-2020, May 2020. [Online]. Available: https://ieeexplore.ieee.org/servlet/opac?punumber=9084217

[45] M. C. Oetzel and S. Spiekermann, "A systematic methodology for privacy impact assessments: A design science approach," *Eur. J. Inf. Syst.*, vol. 23, no. 2, pp. 126–150, Mar. 2014, doi: 10.1057/ejis.2013.18.

[46] B. Friedman, D. G. Hendry, and A. Borning, "A survey of value sensitive design methods," *Found. Trends® Human–Comput. Interact.*, vol. 11, no. 2, pp. 63–125, Nov. 2017, doi: 10.1561/1100000015.

[47] R. V. Schomberg, "A vision of responsible research and innovation," in *Responsible Innovation*. Hoboken, NJ, USA: Wiley, 2013, pp. 51–74.

[48] V. Dignum, "Ethics in artificial intelligence: Introduction to the special issue," *Ethics Inf. Technol.*, vol. 20, no. 1, pp. 1–3, Mar. 2018, doi: 10.1007/s10676-018-9450-z.

[49] R. A. Calvo, D. Peters, and S. Cave, "Advancing impact assessment for intelligent systems," *Nature Mach. Intell.*, vol. 2, no. 2, p. 8991, Feb. 2020, doi: 10.1038/s42256-020-0151-z.

[50] M. Latonero, "Governing artificial intelligence: Upholding human rights & dignity," Data Soc., New York, NY, USA, Tech. Rep., Oct. 2018. [Online]. Available: https://datasociety.net/wp-content/uploads/2018/10/DataSociety_Governing_Artificial_Intelligence_Upholding_Human_Rights.pdf

[51] C. M. Radaelli, "Measuring policy learning: Regulatory impact assessment in Europe," *J. Eur. Public Policy*, vol. 16, no. 8, pp. 1145–1164, Dec. 2009, doi: 10.1080/13501760903332647.

[52] D. Reisman et al., "Algorithmic impact assessments: A practical framework for public agency accountability," AI Now, New York, NY, USA, Tech. Rep., Apr. 2018. [Online]. Available: https://ainowinstitute.org/aiareport2018.pdf

[53] D. Schiff et al., "IEEE 7010: A new standard for assessing the well-being implications of artificial intelligence," in *Proc. IEEE SMC*, Oct. 2020, pp. 2746–2753.

[54] European Commission, High-Level Expert Group on Artificial Intelligence (AI HLEG). (Apr. 2019). *European Commission: Ethical Guidelines for Trustworthy AI*. [Online]. Available: https://ec.europa.eu/newsroom/dae/document.cfm?doc_id=58477

[55] J. Saltz et al., "Integrating ethics within machine learning courses," *ACM Trans. Comput. Edu.*, vol. 19, no. 4, pp. 32:1-32:26, Aug. 2019, doi: 10.1145/3341164.

[56] N. Garrett, N. Beard, and C. Fiesler, “More than `If time allows’: The role of ethics in AI education,” in *Proc. AAAI/ACM Conf. AI, Ethics, Soc.* New York, NY, USA: ACM, Feb. 2020, pp. 272–278, doi: 10.1145/3375627.3375868.

[57] T. O’Brien et al. (Jan. 2020). *A Model for Ethical Artificial Intelligence*. [Online]. Available: https://www.weforum.org/agenda/2020/01/tech-companies-ethics-responsible-ai-microsoft/

[58] United Nations. (2020). *Forests*. [Online]. Available: https://sustainabledevelopment.un.org/topics/forests

[59] S. Nilsson and W. Schopfhauser, “The carbon-sequestration potential of a global afforestation program,” *Climatic Change*, vol. 30, no. 3, pp. 267–293, Jul. 1995, doi: 10.1007/BF01091928.

[60] P. Drawdown. (Feb. 2020). *Tropical Forest Restoration*. [Online]. Available: https://www.drawdown.org/solutions/tropical-forest-restoration

[61] (Feb. 2020). *Tree Plantations (on Degraded Land)*. [Online]. Available: https://www.drawdown.org/solutions/tree-plantations-on-degraded-land

[62] A. M. M. Rasol, “An agent-based framework for forest planning,” Ph.D. dissertation, Dept. Technol., De Montfort Univ., Leicester, U.K., 2007.

[63] M. A. Obaidy and A. Ayesh, “Energy efficient algorithm for swarmed sensors networks,” *Sustain. Comput., Informat. Syst.*, vol. 5, pp. 54–63, Mar. 2015, doi: 10.1016Zj.suscom.2014.09.004.

[64] *Drones Planting Trees: An Interview With Biocarbon Engineering*, Biocarbon Eng., IMPAKTER, London, U.K., May 2019. [Online]. Available: https://impakter.com/biocarbon-engineering/

[65] M. G. Hethcoat et al., “A machine learning approach to map tropical selective logging,” *Remote Sens. Environ.*, vol. 221, Feb. 2019, Art. no. 569582.

[66] C. D. Lippitt et al., “Mapping selective logging in mixed deciduous forest,” *Photogramm. Eng. Remote Sens.*, vol. 74, no. 10, pp. 1201–1211, Oct. 2008.

[67] D. L. Schmoldt, *Application of Artificial Intelligence to Risk Analysis for Forested Ecosystems* (Managing Forest Ecosystems). Dordrecht, The Netherlands: Springer, 2001, pp. 49–74.

[68] D. Sitarz. (Jan. 1993). *Agenda 21: The Earth Summit Strategy to Save Our Planet*. [Online]. Available: https://www.osti.gov/biblio/6289330

[69] B. Carolee and C. Heather, “Shaping the future we want: UN decade of education for sustainable development,” UNESCO, Paris, France, Final Rep., Nov. 2014. [Online]. Available: https://unesdoc.unesco.org/ark:/48223/pf0000230171

[70] D. Rolnick et al., “Tackling climate change with machine learning,” 2019, *arXiv:1906.05433*. [Online]. Available: http://arxiv.org/abs/1906.05433

[71] D. Schuler and A. Namioka, *Participatory Design: Principles and Practices*. Boca Raton, FL, USA: CRC Press, Mar. 1993.

[72] U. Kursuncu, M. Gaur, and A. Sheth, “Knowledge infused learning (K-IL): Towards deep incorporation of knowledge in deep learning,” Feb. 2019, *arXiv:1912.00512*. [Online]. Available: http://arxiv.org/abs/1912.00512

Daniel Schiff received the master’s degree in social policy from the University of Pennsylvania, Philadelphia, PA, USA, in 2013. He is currently pursuing the Ph.D. degree with the School of Public Policy, Georgia Institute of Technology, Atlanta, GA, USA, where he studies the social, ethical, and policy implications of artificial intelligence. He studied philosophy at Princeton University, Princeton, NJ, USA, focusing on robotics and intelligent systems.

Mr. Schiff served as the Secretary for the IEEE 7010 Recommended Practice for Assessing the Impact of Autonomous and Intelligent Systems on Human Well-Being. He has published in venues such as *AI & Society*, *Journal of Engineering Education,* and *AMA Journal of Medical Ethics,* and *Nanotechnology.*

Bogdana (Bobi) Rakova is currently a Data Scientist, part of the Responsible AI Team at Accenture, as well as a Contributor to the IEEE 7010 Recommended Practice for Assessing the Impact of Autonomous and Intelligent Systems on Human Well-Being. She has researched in the fields of fairness, accountability, and transparency of machine learning (ML) and artificial intelligence (AI) for social good. She was part of the Assembly: Ethics and Governance of AI Program in 2018, a collaboration between the Berkman Klein Center for Internet and Society at Harvard Law School, Cambridge, MA, USA, and the MIT Media Lab, Cambridge. She is passionate about social change through responsible systems that empower everyone to participate.

Aladdin Ayesh received the M.Sc. degree from Essex University, Colchester, U.K., in 1995 and the Ph.D. degree from Liverpool John Moores University (LJMU), Liverpool, U.K., in 2000.

He is currently a Professor of artificial intelligence with De Montfort University, Leicester, U.K. He has over 140 publications. His current research focuses on computational cognition and machine learning.

Dr. Ayesh is a member of several IEEE standards working groups and a contributor to IEEE 7010-2020—IEEE Recommended Practice for Assessing the Impact of Autonomous and Intelligent Systems on Human Well-Being. He is also a founding editor of four international journals. He has chaired several international conferences.

Anat Fanti received the M.A. degree in organizational psychology from Columbia University, New York, NY, USA, in 1995. She is currently pursuing the Ph.D. degree in science and technology studies (STS) with Bar-Ilan University, Ramat Gan, Israel.

She practiced as an organizational consultant and held various executive positions in global human resources. She teaches in the course "Policy for Progress and Quality of Life" at the Government and Sustainability School, Herzliya Interdisciplinary College, Herzliya, Israel. Her research traces the history of happiness, well-being, and quality-of-life measures over the past two decades.

Michael Lennon is currently a Senior Fellow at the Center for Excellence in Public Leadership, George Washington University, Washington, DC, USA, and an academic, a technologist, and an activist translating insights from psychosocial sciences into technology and methods for prosocial collective intelligence. His past clients include the World Bank, IBM, and PriceWaterHouseCoopers. He contributes to global groups animating systemic change, such as TheRSA.org, the IEEE 7010 Recommended Practice for Assessing the Well-Being Implications of A/IS, as well as community groups practicing regenerative localism, such as www.ActionforHappiness.org, www.positivestepwork.org, and TransitionNet-work.org.

■ Direct questions and comments about this article to Daniel Schiff, School of Public Policy, Georgia Institute of Technology, Atlanta, GA 30332 USA; schiff@gatech.edu.