

Principles alone cannot guarantee ethical AI

Brent Mittelstadt^{1,2}

Artificial intelligence (AI) ethics is now a global topic of discussion in academic and policy circles. At least 84 public-private initiatives have produced statements describing high-level principles, values and other tenets to guide the ethical development, deployment and governance of AI. According to recent meta-analyses, AI ethics has seemingly converged on a set of principles that closely resemble the four classic principles of medical ethics. Despite the initial credibility granted to a principled approach to AI ethics by the connection to principles in medical ethics, there are reasons to be concerned about its future impact on AI development and governance. Significant differences exist between medicine and AI development that suggest a principled approach for the latter may not enjoy success comparable to the former. Compared to medicine, AI development lacks (1) common aims and fiduciary duties, (2) professional history and norms, (3) proven methods to translate principles into practice, and (4) robust legal and professional accountability mechanisms. These differences suggest we should not yet celebrate consensus around high-level principles that hide deep political and normative disagreement.

Over the past several years, a plethora of public-private initiatives have arisen globally to define values, principles and frameworks for the ethical development and deployment of AI¹. These initiatives can help focus public debate on a common set of issues and principles, and raise awareness among the public, developers and institutions of the ethical challenges that accompany AI².

To date, at least 84 such AI ethics initiatives have published reports describing high-level ethical principles, tenets, values or other abstract requirements for AI development and deployment³. Many envision these high-level contributions being translated into mid- or low-level design requirements and technical fixes, governance frameworks and professional codes⁴.

Existing initiatives to codify AI ethics are not without their critics. Many initiatives, particularly those sponsored by industry, have been characterized as mere virtue-signalling intended to delay regulation and pre-emptively focus debate on abstract problems and technical solutions^{4,5}. This view is difficult to dismiss: AI ethics initiatives have thus far largely produced vague, high-level principles and value statements that promise to be action-guiding, but in practice provide few specific recommendations⁵ and fail to address fundamental normative and political tensions embedded in key concepts (for example, fairness, privacy). Declarations by AI companies and developers committing themselves to high-level ethical principles and self-regulatory codes nonetheless provide policy-makers with a reason not to pursue new regulation^{5,6}.

Comparisons have recently been drawn between AI ethics initiatives and medical ethics⁷. A recent review found that many AI ethics initiatives have converged on a set of principles that closely resemble the four classic principles of medical ethics⁸. This finding has been subsequently endorsed by the Organisation for Economic Co-operation and Development⁹ and the European Commission's High-Level Expert Group on Artificial Intelligence, which proposed four principles to guide the development of 'trustworthy' AI: respect for human autonomy, prevention of harm, fairness and explicability¹⁰.

This convergence of AI ethics around principles of medical ethics is opportune, as it is historically the most prominent and well-studied approach to applied ethics. 'Principlism' emerged from medicine as a theoretical moral framework joining traditional ethical standards with the requirements of practitioners, research

ethics committees and medical institutions for practical ethical decision-making¹¹. Principlism proposes four core principles that require specification and balancing in different decision-making contexts¹². Whereas principlism in medical ethics provides a common language to identify and conceptualize ethical challenges^{13,14}, and provides guidance for setting health policy and clinical decision-making, a principled approach in AI ethics seems intended to embed normative considerations in technology design and governance. Both approaches address how to embed principles in professional practice. Principlism thus provides a helpful backdrop to assess the potential for AI ethics to enact real change in the development and deployment of AI.

Despite the initial credibility lent by the comparison with medical ethics, there are reasons to be concerned about the future impact of AI ethics. Important differences exist between medicine (and other traditional professions¹¹) and AI development that suggest a principled approach in the latter may not enjoy success comparable to the former.

This Perspective critically assesses the strategies and recommendations proposed by current AI ethics initiatives. Outputs of existing AI ethics initiatives were reviewed to determine their proposed strategy for embedding ethics into the development and governance of AI³. Prior work on the implementation and impact of principlism in medicine is used to critically assess the potential impact of a principled approach to AI ethics.

The challenges of a principled approach to AI ethics

Four characteristics of AI development suggest a principled approach may have limited impact on design and governance. Compared to medicine, AI development lacks (1) common aims and fiduciary duties, (2) professional history and norms, (3) proven methods to translate principles into practice, and (4) robust legal and professional accountability mechanisms.

Common aims and fiduciary duties. Medicine is broadly guided by a common aim: to promote the health and well-being of the patient¹⁵. It is a defining quality of a profession for its practitioners to be part of a 'moral community' with common aims, values and training^{16–18}. The pursuit of a common goal facilitates a principled approach to ethical decision-making¹¹. While there is much disagreement over

¹Oxford Internet Institute, University of Oxford, Oxford, UK. ²The Alan Turing Institute, British Library, London, UK. e-mail: brent.mittelstadt@oii.ox.ac.uk

the meaning of health, and how best to promote it in practice, the interests of patients and medical practitioners remain aligned at some fundamental level, which encourages solidarity and trust¹⁹. By some accounts, practitioners have a moral obligation to advocate for their patient's interests against institutional interests^{18–20}.

Comparable solidarity cannot be taken for granted in AI development. AI is largely developed by the private sector for deployment in public (for example, criminal sentencing) and private (for example, insurance) contexts. The fundamental aims of developers, users and affected parties do not necessarily align. Developers often “work in an environment which constantly pressures them to cut costs, increase profit and deliver higher quality” systems, and face pressure from management to make decisions that prioritize company interests^{21,22}. While health professionals undoubtedly face similar organizational pressures, they are not equivalent in degree. Unlike medicine, AI development does not serve the equivalent of a patient whose interests are granted initial primacy in ethical decision-making. This lack of a common goal transforms ethical decision-making from a cooperative to a competitive process, which makes finding a balance between public and private interests more difficult in practice.

The implicit solidarity of medicine is formally recognized in professional codes of practice and regulatory frameworks that establish fiduciary duties towards patients²⁰. Formal professions are distinguished from other vocations by fiduciary duties derived from the client-practitioner relationship²³, which are mediated by common goals and values within the profession, and enforced through sanctions and self-governance (Box 1)^{16,24,25}. These characteristics facilitate a principled approach to ethical decision-making by requiring practitioners to promote their clients' best interests.

AI development is not a formal profession. Equivalent fiduciary relationships and complementary governance mechanisms do not exist for private sector AI developers⁵. AI developers do not commit to public service, which in other professions requires practitioners to uphold public interests in the face of competing business or managerial interests²⁴. For AI or software deployed in the public sector, such a commitment may be implicit in institutional or political structures. The same cannot be said for the private sector. Companies and their employees have principal fiduciary duties towards their shareholders. Public interests are not granted primacy over commercial interests.

It could be argued that many medical practitioners and institutions face similar pressures. Hospitals, for example, must ensure their model of healthcare delivery is sustainable and balances the interests of individual patients against public health goals. Likewise, many new therapies are developed in the private sector.

These apparent similarities are, however, superficial. Public and private medical institutions are required to operate within strict regulatory frameworks that ensure the health and well-being of patients and research participants are not subsumed in pursuit of sustainability or profit¹³. AI institutions are of course also subject to regulation in certain sectors; data protection and privacy law, for example, constrain processing of the personal data necessary to train models. However, the impact of such frameworks is variable and limited in scope^{26,27}. A unified regulatory framework that establishes clear fiduciary duties towards data subjects and users does not yet exist for AI. Should such a framework emerge from AI ethics, a principled approach could be deemed successful. But in the absence of strong regulation that establishes fiduciary duties or primacy for the vital interests of data subjects and users²⁸, a comparable degree of value alignment cannot be said to exist for AI.

The absence of a fiduciary relationship in AI means that users cannot trust that developers will act in their best interests when implementing ethical principles in practice. Reputational risks may push companies to engage with ethics, but these risks carry weight only as long as they remain in the public consciousness. Personal

Box 1 | Characteristics of a formal profession

1. Specialized education and training: members are expected to have undertaken extensive specialized education and training⁵¹, typically in accredited degree programmes³³.
2. Commitment to public service: professions involve a public declaration to provide a service to society or for the public good, making use of specialized, often privileged expertise⁶⁵ that takes precedence over individual gain⁵¹.
3. Higher standard of care: professionals commit to upholding higher ethical standards than would normally be expected in business relationships in service to both the client and the public⁵¹.
4. Enforcement and self-governance: often, these standards are recorded in an ethical code and enforced through a disciplinary system, administered by professional associations^{63,65}.
5. Licensing: entry to the profession is restricted by (government sanctioned) licensure to highly skilled individuals as a means to protect the public^{33,51}.

moral conviction may also push AI developers towards ‘good’ behaviour²⁹; recent examples of internal protests at Google provide some cause for hope^{30,31}. However, incentive structures that discourage whistleblowing or placing public interests before the company suggest that virtuous actions will often come at a high personal cost³². It is unacceptable that users and affected parties must rely on the personal convictions of developers, fear of reputational damage or public outcry for their vital interests in privacy, autonomy, identity and other areas to be taken seriously³³.

Professional history and norms. The second weakness of a principled approach to AI ethics is a relative lack of professional history and well-defined norms of good behaviour. Medicine benefits from a long history and shared professional culture with variation across cultures and specialties. Accounts of the moral obligations and virtues of health professions have developed over centuries¹⁴. In Western biomedicine, longstanding standards set in the Hippocratic oath, the Declaration of Geneva, the Declaration of Helsinki and other accounts of the good health professional have served as a basis for clinical decision-making and research ethics³⁴, and inspired the development of ethics and codes in other professions¹¹.

These standards and accounts of being a good doctor have not entirely prevented ethical negligence in clinical practice and medical research. Nonetheless, they provide a historically sensitive account of the obligations of the profession against which negligent conduct and practices can be identified¹⁹. As such failures have occurred across the profession's history, and as new technologies, treatments and changing social values have disrupted established norms of good behaviour, ethical standards have been revised.

Evidence of these historical lessons can be seen in modern codes of conduct and ethics; the American Medical Association's Code of Medical Ethics, for example, is a highly detailed document detailing opinions, behavioural norms and standards across a plethora of medical practices and technologies³⁵. The first version (1847) of the code focused solely on professional conduct and emphasized a paternalistic duty to maximize benefits and minimize harms to patients. Over time the profession's standards have shifted away from this myopic focus on beneficence and professional conduct, and towards other duties owed to patients, most notably a growing respect for autonomy^{14,34}. Particularly egregious failures have often provided the catalyst for change; atrocities in human experimentation, involuntary sterilization and euthanasia committed in Nazi Germany during World War II, for example, led to the formalization of a set of research ethics principled in the 1947 Nuremberg Code.

Similarly, in the United States the National Research Act (1974) and Belmont Report (1978) were drafted in response to abuses observed in the Tuskegee syphilis study and other human studies^{13,14,34}.

Principlism¹² subsequently arose to correct the “minimal and unsatisfactory professional morality” the Hippocratic tradition was perceived to have created¹⁴. To move beyond this tradition, “explicit recognition of basic ethical principles that could help identify various clinical practices and human experiments as morally questionable or unacceptable” was required¹⁴. By addressing concerns with conduct as well as practice, principlism has established a common moral language to identify and address problematic behaviours and practices. This framework has subsequently taken hold in curricula and training for medical students, practitioners and policy-makers, and continues to influence ethical codes and decision-making in medicine^{13,14,36}. Principlism thus has a strong, historically informed regulating influence on the behaviour and ethics of medical practitioners and institutions.

AI development does not have a comparable history, homogeneous professional culture and identity, or similarly developed professional ethics frameworks. The profession has not gone through comparable transformative moments in which its ethical obligations are clearly recognized and translated into specific, practical moral duties and best practices. Whereas AI can in principle be deployed in any context involving human expertise, medicine in comparison has narrower aims that facilitate development of standard practices and norms. Reflecting this, AI developers come from varied disciplines and professional backgrounds, which have incongruous histories, cultures, incentive structures and moral obligations^{11,24,25}. Reducing the field to a single vocation or type of expertise would be an oversimplification¹¹.

Software engineering, which is arguably the closest analogue, has historically not been legally recognized as a profession with fiduciary duties to the public^{11,37} due to the absence of licensure schemes and a well-defined professional standard of care³⁸. Reflecting this, a comparably rich account of what it means to be a good AI developer or software engineer does not exist; while the Institute of Electrical and Electronics Engineers (IEEE) and the Association of Computing Machinery (ACM), two of the field’s largest professional associations, have published and revised codes of ethics, these documents remain comparatively short, theoretical and lacking in grounded advice and specific behavioural norms¹¹.

Stronger professional standards (and supporting organizations) for AI development could, of course, be drafted in the future. Unfortunately, this will not be a simple task. Systems are often created by large, multidisciplinary and multinational teams. Whereas the effects of clinical decision-making are often (but not always) immediate and observable, the impact of decisions taken in designing, training and configuring AI systems for different uses may never become apparent to developers³⁹. This is worrying, as distance from potential victims has been shown to have a positive effect on unethical professional behaviour⁴⁰. The risks addressed by medical ethics largely arise from interventions performed (or not) on the physical body. In comparison, ethical risks in AI are continuous and not similarly bound, and may not be directly experienced by data subjects⁴¹. Systems are also often opaque in the sense that no single person will have a full understanding of the system’s design or functionality⁴², or be able to predict its behaviour. Even where problems are recognized, they can rarely be traced back to a single team member or action⁴³; responsibility must be assigned across a network of actors that influenced the system’s design, training and configuration. This inability to reliably predict the effects of development choices undermines the creation of standards to be a good AI developer or requirements for good AI.

AI ethics initiatives address this gap by defining broadly acceptable principles to guide the people and processes responsible for the development, deployment and governance of AI across

radically different contexts of use. At this level of abstraction meaningful guidance may be impossible^{44,45}. The great diversity of stakeholders and interests involved necessarily pushes the search for common values and norms towards a high level of abstraction²⁴. The results are statements of principles or values based on abstract and vague concepts—for example, commitments to ensure AI is fair or respects human dignity, which are not specific enough to be action-guiding⁷.

Statements reliant on vague normative concepts hide points of political and ethical conflict. Fairness, dignity and other such abstract concepts are examples of ‘essentially contested concepts’ with many possible conflicting meanings that require contextual interpretation through one’s background political and philosophical beliefs⁴⁶. These different interpretations, which can be rationally and genuinely held, lead to substantively different requirements in practice⁴⁷, which will only be revealed once principles or concepts are translated and tested in practice⁴⁸. At best, this conceptual ambiguity allows for context-sensitive specification of ethical requirements for AI. At worst, it masks fundamental, principled disagreement and drives AI ethics towards moral relativism. At a minimum, any compromise reached thus far around core principles for AI ethics does not reflect meaningful consensus on a common practical direction for good AI development and governance. We must not confuse high-level compromise with a priori consensus.

The truly difficult part of ethics—actually translating normative theories, concepts and values into good practices AI practitioners can adopt—is kicked down the road like the proverbial can. Developers are left to translate principles and specify essentially contested concepts as they see fit, without a clear roadmap for unified implementation. This process is likely to encounter incomensurable moral norms and frameworks that present true moral dilemmas that principlism cannot resolve^{49–51}. An established profession that can draw on a rich history, ethical culture and norms of good practice will be best placed to conceptualize and debate (if not resolve) these challenges^{11,52}. Unfortunately, AI development is left wanting in this regard.

Methods to translate principles in practice. The third weakness of a principled approach to AI ethics is the absence of proven methods to translate principles into practice. The prevalence of essentially contested concepts in AI ethics begs a question: how can normative disagreements over the correct specification of such concepts be resolved?

Principles do not automatically translate into practice²¹. Throughout its history, medicine has developed effective ways of translating high-level commitments and principles into practical requirements and norms of good practice^{4,18}. Professional societies and boards, ethics review committees, accreditation and licensing schemes, peer self-governance, codes of conduct, and other mechanisms supported by strong institutions help determine the ethical acceptability of day-to-day practice by assessing difficult cases, identifying negligent behaviour and sanctioning bad actors^{53,54}. The formal codes and informal norms that govern medical practice have been extensively tested, studied and revised over time, with their recommendations and norms (and underlying principles) evolving to remain relevant. High-level principles rarely feature explicitly in clinical decision-making, which is instead bounded by institutional policies that include principled concerns¹³. On the ground, case-relevant precedents or specifications that use the moral language provided by principlism are more common^{36,55}. Taken together, institutional and clinical decision-making are effectively a coherence approach accounting for both high-level and grounded considerations¹².

AI development does not have comparable empirically proven methods to translate principles into practice in real-world development contexts. This is a multifaceted methodological challenge.

Translation involves the specification of high-level principles into mid-level norms and low-level requirements. Norms and requirements cannot be deduced directly from mid-level principles without accounting for specific elements of the technology, application, context of use or relevant local norms^{51,55,56}. Normative decisions must be made at each stage of translation, and coherence must be sought between principles, specified norms or rules and the facts of the case^{12,56}. It follows that the justification arising from widespread consensus on a set of common principles does not transfer to the mid-level norms and low-level requirements derived from them. Each stage of translation and specification must be independently justified. A universally accepted hierarchy of principles to prioritize competing norms does not exist^{49,57}.

This observation reveals the scope of work that remains for AI ethics. High-level consensus is encouraging, but it has little bearing on the justification of norms and practical requirements proposed within specific contexts of use. Due to the necessity of local specification and justification, the prominence of essentially contested concepts in AI ethics, and the field's relative lack of a binding professional history (see the 'Professional systems and norms' section), conflicting practical requirements will almost certainly emerge across the diverse sectors and contexts in which a principled approach to AI ethics is used.

One other methodological challenge remains. Normative practical requirements must somehow be embedded in development processes, and functionally implanted in design requirements. Prior work points to the difficulty of embedding ethical values and principles in technology design and the development cycle^{1,56,58,59}. Many such methods exist⁶⁰, including participatory design, reflective design, Values at Play and value-sensitive design¹, but thus far they have largely been implemented and studied in academic contexts, which are more receptive to normative concerns than commercial settings^{22,60,61}. Value-conscious methods are also largely procedural, not functional. Generally speaking, they introduce values, normative issues and relevant stakeholders into the development process⁵⁶. They do not, however, allow for particular values to be injected into system design, and struggle to capture the degree to which the resulting artefact reflects particular values or specifications⁵⁸.

Value-conscious design frameworks face additional challenges in commercial development processes. Ethics has a cost. AI is often developed behind closed doors without public representation. Gathering the views of relevant stakeholders, embedding an ethicist with the development team and resolving conflicts between different specifications of essentially contested concepts create additional work and costs. Unsurprisingly, ethical considerations may be discarded when they conflict with commercial incentives²². It cannot be assumed that value-conscious frameworks will be meaningfully implemented in commercial processes that value efficiency, speed and profit.

Legal and professional accountability. The fourth weakness of a principled approach to AI ethics is the relative lack of legal and professional accountability mechanisms. Medicine is governed by legal and professional frameworks that uphold professional standards and provide patients with redress for negligent behaviour, including malpractice law, licensing and certification schemes, ethics committees, and professional medical boards^{53,54}. Medical institutions subject to regulation help ensure these standards are upheld¹³. Legally supported accountability mechanisms provide an external impetus for health professionals to fulfil their fiduciary duties, amplifies complementary forms of self-governance by establishing a clear link between bad behaviour and professional sanctions (for example, losing one's licence to practice)¹¹, mandates a professional standard of care and allows patients to make claims against negligent members of the profession.

Box 2 | Key questions to assess AI ethics value and principle statements

1. Who wrote it, and how?
2. Who is it intended for, and what is its purpose?
3. Why should I follow it?
4. How do I follow or implement it?
5. How should I resolve conflicting interpretations of essentially contested concepts?
6. How will you know I am following it?
7. What happens if I fail to follow it?
8. How can I raise disagreements or questions for clarification?

Excluding certain types of risks (for example, privacy violations governed by data protection law), AI development does not have comparable professionally or legally endorsed accountability mechanisms⁶². This is a problem. Serious, long-term commitment to self-regulatory frameworks cannot be taken for granted^{4,5}.

Prior research on the impact of codes of ethics on professional behaviour has revealed mixed results. Codes are often followed in letter rather than spirit, or as a checklist rather than as part of a critical reflexive practice^{24,50,63,64}. A recent study of the ACM Code of Ethics revealed that it has little effect on the day-to-day decision-making of software engineering professionals and students⁶⁵. Other studies of corporate and professional codes of ethics outside computing have reported similar results^{66–68}. A recent meta-analysis of evidence on the impact of codes on professional behaviour found that the mere existence of a code has no discernible effect on unethical behaviour; rather, an effect is only found when codes (and their underlying principles) are embedded in organizational culture and actively enforced^{29,69,70}. Norms must be clearly defined and highly visible if they are to influence practitioners¹¹ and inspire peer self-governance. Current governance structures in AI companies are insufficient in this regard².

External sanctions for breaching a code are also key for adherence and effective self-governance¹¹. Compared with medicine, information professions lack sanctions that can impact the professional's livelihood²⁴. While software engineering degrees can now be accredited, a licence is not required to practise⁷¹, with some national exceptions^{72–74}. Professional bodies such as the IEEE and ACM lack formal sanction powers beyond expulsion from the organization, which, absent licensing, does not impact the ability to practise.

While stronger legal and professional accountability mechanisms could be adopted, this seems unlikely in the near term. AI development is not a unified profession with a longstanding history and harmonized aims. AI developers do not formally provide a public service, meaning public interests need not be given primacy⁷⁵. AI does not operate in a single sector, meaning any new legal or professional mechanisms must account for myriad potential benefits and harms, and integrate with existing sector-specific law. Finally, proposals to introduce professional sanctions and licensing schemes for computing professionals are also not new, but have thus far seen limited uptake^{72,74,76,77}.

These weaknesses in existing legal and professional accountability mechanisms for AI raises a difficult question: is it enough to define good intentions and hope for the best? Without complementary punitive mechanisms and governance bodies to step in when self-governance fails, a principled approach runs the risk of merely providing false assurances of ethical or trustworthy AI⁷⁸.

Where should AI ethics go from here?

While principlism has undoubtedly significantly influenced medical ethics owing to the four characteristics analysed above, it has not been an unqualified success. The comparative shortcomings of AI

development along these lines thus gives cause for concern. First, AI development is not a formal profession with aims that align with public interests. Second, developers are not governed by a historically validated account of what it means to be a good AI developer. Third, outside of academic contexts AI development lacks proven methods to translate principles into practice. And finally, when a developer falls afoul of these vaguely defined requirements, few sanction mechanisms and channels for redress to set things right remain. Signing up to self-regulatory codes lacking clearly defined and enforceable obligations costs developers nothing, but can have immediate benefits in terms of trustworthiness and reputation. Together, these shortcomings point towards significant challenges facing the implementation of AI ethics.

We must therefore hesitate to celebrate consensus around high-level principles that hide deep political and normative disagreement. Shared principles are not enough to guarantee trustworthy or ethical AI in the future. Without a fundamental shift in regulation, translating principles into practice will remain a competitive, not cooperative, process. This is a problem, as principles remain vacuous until tested, at which point the true costs and value of a principled approach to AI ethics will be revealed. Conflicting prescriptions of essentially contested concepts are to be expected. Resolving these conflicts is where the real work starts for AI ethics. A key question remains: how can this essential work be supported by government, industry and civil society?

Clearly define sustainable pathways to impact. A principled approach requires cooperative oversight to ensure translated norms and requirements remain fit for purpose and impactful over time. Going forward, the long-term aims and pathways to impact of principled initiatives must be more clearly defined (see Box 2 for key questions). Binding and highly visible accountability structures as well as clear implementation and review processes are needed at a sectoral and organizational level²⁵. Professional and institutional norms can be established by defining clear requirements for inclusive design, transparent ethical review¹⁶, documentation of models and datasets, and independent ethical auditing.

Support bottom-up AI ethics in the private sector. A top-down approach to AI ethics is uniquely difficult due to the diversity of technologies described as AI. In such a diverse field, generalist top-down approaches must be complemented by bottom-up case studies of production AI systems. Local practices can be collaboratively assessed to specify principles and define precedents to move professional standards forward^{11,52}. Novel cases reveal new challenges for AI ethics, which are desperately needed to move the field beyond well-worn cases^{79,80} and develop sector- and case-specific guidelines, technical solutions and an empirical knowledge base detailing the impact and harms of production AI technologies. Much current bottom-up work pursues technological solutions and metrics for ethical concepts amenable to quantification (for example, fairness)^{5,81}, and occurs primarily in an academic environment. Increased support and access to development settings should be made available to support multidisciplinary bottom-up research and development in AI ethics, particularly in commercial development contexts currently closed to external scrutiny.

License developers of high-risk AI. To encourage long-term recognition of ethical commitments, it may be necessary to formally establish AI development as a profession with equivalent standing to other high-risk professions. It is a regulatory oddity that we license professions providing a public service, but not the profession responsible for developing technical systems to augment or replace human expertise and decision-making within them. The risks of licensed professions have not dissipated, but rather been displaced to AI. To unpack the significant challenges facing licensing such a

diverse set of practitioners¹¹, initiatives could initially target developers of systems with elevated risk or built for the public sector, such as facial recognition systems designed for policing.

Shift from professional ethics to organizational ethics. The outputs of many AI ethics initiatives resemble professional codes of ethics that address design requirements and the behaviours and values of individual professions¹. The legitimacy of particular applications and their underlying business and organizational interests remain largely unquestioned^{1,82}. This approach conveniently steers debate towards the transgressions of unethical individuals, and away from the collective failure of unethical organizations and business models²⁷. Developers will always be constrained by the institutions that employ them. To be truly effective, the ethical challenges of AI cannot be conceptualized as individual failures. Going forward, AI ethics must become an ethics of AI businesses and organizations as well.

Pursue ethics as a process, not technological solutionism. Many initiatives suggest ethical challenges can best be addressed through “technical and design expertise”, and address concepts for which technical fixes seem feasible (for example, privacy, fairness)¹, but rarely propose technical definitions or explanations⁵. Exceptions such as the IEEE’s Ethically Aligned Design initiative do exist, although the impact and uptake of this work in commercial environments remains to be seen⁸³. Nonetheless, the rationale seems to be as follows: insufficient consideration of ethics leads to poor design decisions which create systems that harm users.

This attitude is misguided. The promise of AI largely owes to its apparent capacity to replace or augment human expertise. This malleability means AI inevitably becomes entangled in the ethical and political dimensions of vocations and practices in which it is embedded. AI ethics is effectively a microcosm of the political and ethical challenges faced in society. Framing ethical challenges in terms of design flaws ensures they remain “fundamentally technical, shielded from democratic intervention”¹. It is foolish to assume that very old and complex normative questions can be solved with technical fixes or good design alone. The risk is that complex, difficult ethical debates will be oversimplified to make the concepts at hand computable and implementable in a straightforward but conceptually shallow manner⁸⁰.

Ethics is not meant to be easy or formulaic. Intractable principled disagreements should be expected and welcomed, as they reflect both serious ethical consideration and diversity of thought. They do not represent failure, and do not need to be ‘solved’. Ethics is a process, not a destination. The real work of AI ethics begins now: to translate and implement our lofty principles, and in doing so to begin to understand the real ethical challenges of AI.

Received: 23 May 2019; Accepted: 9 October 2019;
Published online: 4 November 2019

References

1. Greene, D., Hoffmann, A. L. & Stark, L. Better, nicer, fairer: a critical assessment of the movement for ethical artificial intelligence and machine learning. In *Proc. 52nd Hawaii International Conference on System Sciences* 2122–2131 (HICSS, 2019).
2. Whittaker, M. et al. *AI Now Report 2018* (AI Now Institute, 2018).
3. Jobin, A., Ienca, M. & Vayena, E. The global landscape of AI ethics guidelines. *Nat. Mach. Intell.* **1**, 389–399 (2019).
4. Nemitz, P. Constitutional democracy and technology in the age of artificial intelligence. *Philos. Trans. R. Soc. A* **376**, 20180089 (2018).
5. Hagendorff, T. The ethics of AI ethics -- an evaluation of guidelines. Preprint at <https://arxiv.org/abs/1903.03425> (2019).
6. Calo, R. Artificial intelligence policy: a primer and roadmap. *UC Davis Law Rev.* **51**, 399–436 (2019).
7. Whittlestone, J., Nystrup, R., Alexandrova, A., Dihal, K. & Cave, S. *Ethical and Societal Implications of Algorithms, Data, and Artificial Intelligence: A Roadmap For Research* 59 (Nuffield Foundation, 2019).

8. Floridi, L. et al. AI4People—an ethical framework for a good AI society: opportunities, risks, principles, and recommendations. *Minds Mach.* **28**, 689–707 (2018).
9. Forty-two countries adopt new OECD Principles on Artificial Intelligence *OECD* <http://www.oecd.org/science/forty-two-countries-adopt-new-oecd-principles-on-artificial-intelligence.htm> (2019).
10. High Level Expert Group on Artificial Intelligence *Ethics Guidelines for Trustworthy AI*. (European Commission, 2019).
11. Filipović, A., Koska, C. & Paganini, C. *Developing a Professional Ethics for Algorithmists: Learning from the Examples of Established Ethics* (Bertelsmann Stiftung, 2018).
12. Beauchamp, T. L. & Childress, J. F. *Principles of Biomedical Ethics* (Oxford Univ. Press, 2009).
13. Bosk, C. L. Bioethics, raw and cooked: extraordinary conflict and everyday practice. *J. Health Soc. Behav.* **51**, S133–S146 (2010).
14. Beauchamp, T. L. & DeGrazia, D. in *Handbook of Bioethics: Taking Stock of the Field from a Philosophical Perspective* (ed. Khushf, G.) 78, 55–74 (Springer, 2006).
15. Marshall, T. H. The recent history of professionalism in relation to social structure and social policy. *Can. J. Econ. Polit. Sci.* **5**, 325–340 (1939).
16. Frankel, M. S. Professional codes: why, how, and with what impact? *J. Bus. Ethics* **8**, 109–115 (1989).
17. *Black's Law Dictionary* 2nd edn (West, 2009).
18. MacIntyre, A. *After Virtue: A Study in Moral Theory* (Gerald Duckworth & Co, 2007).
19. Gillon, R. Do doctors owe a special duty of beneficence to their patients? *J. Med. Ethics* **12**, 171–173 (1986).
20. Pellegrino, E. D. & Thomasma, D. C. *The Virtues in Medical Practice* (Oxford Univ. Press, 1993).
21. Van den Bergh, J. & Deschoolmeester, D. Ethical decision making in ICT: discussing the impact of an ethical code of conduct. *Commun. IBIMA* **2010**, 127497 (2010).
22. Manders-Huits, N. & Zimmer, M. Values and pragmatic action: the challenges of introducing ethical intelligence in technical design communities. *Int. Rev. Inform. Ethics* **10**, 37–44 (2009).
23. McDowell, B. *Ethical Conduct and the Professional's Dilemma: Choosing Between Service and Success* (Quorum, 1991).
24. Iacovino, L. Ethical principles and information professionals: theory, practice and education. *Aust. Acad. Res. Libr.* **33**, 57–74 (2002).
25. Boddington, P. *Towards a Code of Ethics for Artificial Intelligence Research* (Springer, 2017).
26. Zarsky, T. Incompatible: the GDPR in the age of big data. *Seton Hall Law Rev.* **47**, 995–1010 (2017).
27. Wachter, S. & Mittelstadt, B. D. A right to reasonable inferences: re-thinking data protection law in the age of big data and AI. *Columbia Bus. Law Rev.* **2019**, 494–620 (2019).
28. Balkin, J. M. Information fiduciaries and the first amendment. *UCDL Rev.* **49**, 1183–1234 (2015).
29. Kish-Gephart, J. J., Harrison, D. A. & Treviño, L. K. Bad apples, bad cases, and bad barrels: meta-analytic evidence about sources of unethical decisions at work. *J. Appl. Psychol.* **95**, 1–31 (2010).
30. Wakabayashi, D. & Shane, S. Google will not renew Pentagon contract that upset employees. *The New York Times* <https://www.nytimes.com/2018/06/01/technology/google-pentagon-project-maven.html> (2018).
31. Conger, K. & Wakabayashi, D. Google employees protest secret work on censored search engine for China. *The New York Times* <https://www.nytimes.com/2018/08/16/technology/google-employees-protest-search-censored-china.html> (2018).
32. Wong, J. C. Demoted and sidelined: Google walkout organizers say company retaliated. *The Guardian* <https://www.theguardian.com/us-news/2019/apr/22/google-mass-protests-employee-retaliation> (2019).
33. Parker, D. B. *Ethical Conflicts in Computer Science and Technology* (AFIPS Press, 1981).
34. Carrese, J. A. & Sugarman, J. The inescapable relevance of bioethics for the practicing clinician. *Chest* **130**, 1864–1872 (2006).
35. Brotherton, S., Kao, A. & Crigger, B. J. Professing the values of medicine: the modernized AMA code of medical ethics. *J. Am. Med. Assoc.* **316**, 1041–1042 (2016).
36. Greenfield, B. & Jensen, G. M. Beyond a code of ethics: phenomenological ethics for everyday practice. *Physiother. Res. Int.* **15**, 88–95 (2010).
37. Panensky, S. A. & Jones, R. Does IT go without saying? *Professional Times Magazine* 30–33 (Summer, 2018).
38. Perlman, D. T. Who pays the price of computer software failure notes and comments. *Rutgers Comput. Technol. Law J.* **24**, 383–416 (1998).
39. Mittelstadt, B., Allo, P., Taddeo, M., Wachter, S. & Floridi, L. The ethics of algorithms: mapping the debate. *Big Data Soc.* <https://doi.org/10.1177/2053951716679679> (2016).
40. Jones, T. M. Ethical decision making by individuals in organizations: an issue-contingent model. *Acad. Manage. Rev.* **16**, 366–395 (1991).
41. Metcalf, J. & Crawford, K. Where are human subjects in big data research? The emerging ethics divide. *Big Data Soc.* <https://doi.org/10.1177/2053951716650211> (2016).
42. Burrell, J. How the machine 'thinks': understanding opacity in machine learning algorithms. *Big Data Soc.* <https://doi.org/10.1177/2053951715622512> (2016).
43. Floridi, L. Faultless responsibility: on the nature and allocation of moral responsibility for distributed moral actions. *Philos. Trans. R. Soc. A* **374**, 20160112 (2016).
44. Awad, E. et al. The Moral Machine experiment. *Nature* **563**, 59–64 (2018).
45. Ess, C. Ethical pluralism and global information ethics. *Ethics Inf. Technol.* **8**, 215–226 (2006).
46. van den Hoven, J. Computer ethics and moral methodology. *Metaphilosophy* **28**, 234–248 (1997).
47. Gallie, W. B. Essentially contested concepts. *Proc. Aristot. Soc.* **56**, 167–198 (1955).
48. Richardson, H. S. Specifying norms as a way to resolve concrete ethical problems. *Philos. Public Aff.* **19**, 279–310 (1990).
49. Turner, L. Bioethics in a multicultural world: medicine and morality in pluralistic settings. *Health Care Anal.* **11**, 99–117 (2003).
50. Rhodes, R. Good and not so good medical ethics. *J. Med. Ethics* **41**, 71–74 (2015).
51. Clouser, K. D. & Gert, B. A critique of principlism. *J. Med. Philos.* **15**, 219–236 (1990).
52. Degrazia, D. Moving forward in bioethical theory: theories, cases, and specified principlism. *J. Med. Philos.* **17**, 511–539 (1992).
53. Orentlicher, D. The influence of a professional organization on physician behavior symposium on the legal and ethical implications of innovative medical technology. *Albany Law Rev.* **57**, 583–606 (1993).
54. Papadakis, M. A. et al. Disciplinary action by medical boards and prior behavior in medical school. *N. Engl. J. Med.* **353**, 2673–2682 (2005).
55. Toulmin, S. How medicine saved the life of ethics. *Perspect. Biol. Med.* **25**, 736–750 (1982).
56. van de Poel, I. in *Philosophy and Engineering: Reflections on Practice, Principles and Process* (eds Michelfelder, D. P., McCarthy, N. & Goldberg, D. E.) 253–266 (Springer, 2013).
57. Gillon, R. Defending the four principles approach as a good basis for good medical practice and therefore for good medical ethics. *J. Med. Ethics* **41**, 111–116 (2015).
58. Friedman, B., Hendry, D. G. & Borning, A. A survey of value sensitive design methods. *Found. Trends Hum. Comp. Interact.* **11**, 63–125 (2017).
59. Friedman, B. & Kahn, P. H. Human agency and responsible computing: Implications for computer system design. *J. Syst. Softw.* **17**, 7–14 (1992).
60. Morley, J., Floridi, L., Kinsey, L. & Elhalal, A. From what to how: an overview of AI ethics tools, methods and research to translate principles into practices. Preprint at <https://arxiv.org/abs/1905.06876> (2019).
61. Shilton, K. Values levers: building ethics into design. *Sci. Technol. Hum. Values* **38**, 374–397 (2013).
62. MacKinnon, K. S. Computer malpractice: are computer manufacturers, service bureaus, and programmers really the professionals they claim to be. *St. Clara Rev.* **23**, 1065 (1983).
63. Bynum, T. W. in *Computer Ethics and Professional Responsibility* (eds Bynum, T. W. & Rogerson, S.) 60–87 (Blackwell, 2004).
64. Ladd, J. in *Ethical Issues in the Use of Computers* (eds Johnson, D. G. & Snapper, J. W.) 8–13 (Wadsworth, 1985).
65. McNamara, A., Smith, J. & Murphy-Hill, E. Does ACM's code of ethics change ethical decision making in software development? In *Proc. 2018 26th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering* (ACM, 2018).
66. Brief, A. P., Dukerich, J. M., Brown, P. R. & Brett, J. F. What's wrong with the treadway commission report? Experimental analyses of the effects of personal values and codes of conduct on fraudulent financial reporting. *J. Bus. Ethics* **15**, 183–198 (1996).
67. Helin, S. & Sandström, J. An inquiry into the study of corporate codes of ethics. *J. Bus. Ethics* **75**, 253–271 (2007).
68. McCabe, D. L., Trevino, L. K. & Butterfield, K. D. The influence of collegiate and corporate codes of conduct on ethics-related behavior in the workplace. *Bus. Ethics Q.* **6**, 461–476 (1996).
69. Jin, K. G., Drozdenko, R. & Bassett, R. Information technology professionals' perceived organizational values and managerial ethics: an empirical study. *J. Bus. Ethics* **71**, 149–159 (2007).
70. Shilton, K. "That's not an architecture problem!": techniques and challenges for practicing anticipatory technology ethics. In *iConference 2015 Proceedings* 7 (iSchools, 2015).
71. Goertzel, K. M. Legal liability for bad software. *CrossTalk* 23–28 (September/October, 2016).
72. Laplante, P. A. Licensing professional software engineers: seize the opportunity. *Commun. ACM* **57**, 38–40 (2014).

73. Pour, G., Griss, M. L. & Lutz, M. The push to make software engineering respectable. *Computer* **33**, 35–43 (2000).

74. Seidman, S. B. in *E-Government ICT Professionalism and Competences Service Science* (eds Mazzeo, A., Bellini, R. & Motta, G.) 280, 59–67 (Springer, 2008).

75. Abbott, A. Professional ethics. *Am. J. Sociol.* **88**, 855–885 (1983).

76. O'Connor, J. E. Computer professionals: the need for state licensing. *Jurimetr. J.* **18**, 256–267 (1978).

77. Gotterbarn, D. The ethical computer practitioner—licensing the moral community: a proactive approach. *SIGCSE Bull.* **30**, 8–10 (1998).

78. Suchman, L. Corporate accountability. *Robot Futures* <https://robotfutures.wordpress.com/2018/06/10/corporate-accountability/> (2018).

79. Angwin, J., Larson, J. & Kirchner, L. Machine Bias. *ProPublica* <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing> (2016).

80. European Group on Ethics in Science and New Technologies *Statement on Artificial Intelligence, Robotics and 'Autonomous' Systems* (European Commission, 2018).

81. Holstein, K., Vaughan, J. W., Daumé, H. III, Dudík, M. & Wallach, H. Improving fairness in machine learning systems: what do industry practitioners need? In *Proc. 2019 CHI Conference on Human Factors in Computing Systems* 600 (2018).

82. Benkler, Y. Don't let industry write the rules for AI. *Nature* **569**, 161–161 (2019).

83. Chatila, P. & Havens, J. C. in *Robotics and Well-Being* (eds Aldinhas Ferreira, M. I. et al.) 11–16 (Springer, 2019).

Acknowledgements

The author would like to thank S. Wachter, B. Prainsack and B. Stahl for their insightful feedback, which has immensely improved the quality of this work. Financial support for this work was provided by the Alan Turing Institute (EPSRC) and the British Academy.

Competing interests

The author has previously received reimbursement for conference-related travel from funding provided by DeepMind Technologies Limited.

Additional information

Correspondence should be addressed to B.M.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

© The Author(s), under exclusive licence to Springer Nature Limited 2019