

AI Ethics – Too Principled to Fail?

Brent Mittelstadtⁱ

Abstract

AI Ethics is now a global topic of discussion in academic and policy circles. At least 63 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 in 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.

1 What medicine can teach us about ‘AI Ethics’ⁱⁱ

AI Ethics is now a global topic of discussion in academic and policy circles. Significant resources have been dedicated to public-private initiatives aiming to define universal high-level values or principles to guide ethical development and deployment of AI.¹ 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 AI.²

To date, at least 63 such ‘AI Ethics’ initiatives have published reports describing high-level ethical principles, tenets, values, or other abstract requirements for AI development and deployment.³ Many envision these high-level contributions being ‘translated’ into mid- or low-level design requirements and technical fixes, governance frameworks, and developer codes of ethics.¹

Existing initiatives to codify AI Ethics are not without their critics. Many initiatives are at least partially sponsored by industry, which has led some commentators to suggest they exist for disingenuous virtue signalling intended merely to delay regulation and pre-emptively shape the debate around abstract problems and technical solutions.^{1,4} This view undeniably has some merit: AI Ethics initiatives have thus far have produced vague, high-level principles and value statements which promise to be action-guiding, but in practice provide few specific recommendations⁵ and fail to address fundamental normative and political tensions embedded

ⁱ Oxford Internet Institute, University of Oxford, 1 St. Giles, Oxford, OX1 3JS, UK; the Alan Turing Institute, British Library, 96 Euston Road, London, NW1 2DB, UK. Correspondence via brent.mittelstadt@oii.ox.ac.uk. The author would like to thank Prof. Sandra Wachter, Prof. Barbara Prainsack, and Prof. Bernd Stahl for their insightful feedback that has immensely improved the quality of this work.

ⁱⁱ Funding: This work was supported by The Alan Turing Institute, EPSRC grant EP/N510129/1, and the British Academy, grant PF170151. Competing interests: The author has previously received travel reimbursement from funding provided by DeepMind Technologies Limited.

in key concepts (e.g. fairness, privacy).ⁱⁱⁱ 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.^{5,7}

Comparisons have recently been drawn between AI Ethics initiatives and medical ethics⁸ A recent review undertaken by the AI4People project found that AI Ethics initiatives have converged on a set of principles that closely resemble the four classic principles of medical ethics.⁹ This position was subsequently endorsed by the OECD¹⁰ and the European Commission's High Level Expert Group on Artificial Intelligence (HLEG),^{iv} which proposed four principles to guide the development of 'trustworthy' AI: respect for human autonomy, prevention of harm, fairness, and explicability.¹³

This convergence of AI Ethics around principles defined in medical ethics is opportune, as it is perhaps the most prominent and well-studied approach to applied ethics to date. 'Principlism' emerged from medicine as a theoretical moral framework grounded in the experiences of practitioners, research ethics committees, and medical institutions in grappling with ethics on a case-by-case basis.¹⁴ These experiences revealed that ethical decision-making on the ground often involves a trade-off or 'weighting' of interests to decide upon the ethically appropriate course of action. Principlism formalised these interests around four core principles that provide guidance and require balancing in different contexts.¹⁵ Whereas 'principlism' in medical ethics ensures a principled approach to setting health policy and clinical decision-making, a principled approach in AI Ethics seems intended to embed principles and normative concerns in technology design and governance. Both approaches address how to embed principles into professional practice. The convergence of AI Ethics around medical ethical principles thus provides a helpful backdrop to assess the current state and potential success of AI Ethics in terms of enacting real change in the development and deployment of AI.

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 in the latter may not enjoy success comparable to the former. Medical ethics thus provides an insightful mirror to critically examine the potential weaknesses of a principled approach to ethics in AI development and governance.^v

ⁱⁱⁱ There are, of course, a few exceptions; the IEEE 'Ethically Aligned Design' initiative,⁶ for instance, has been impressive in its scale, interdisciplinarity, and creation of subsequent projects to develop ethical standards and curriculum for AI.

^{iv} It is worth noting that a second body in the European Commission, the European Group on Ethics in Science and New Technologies (EGE), also issued a statement in 2018 on AI Ethics that does not align with the approach advocated by the HLEG.¹¹ The EGE later wrote an open letter¹² to President Jean-Claude Juncker drawing attention to shortcomings of the Commission's current approach to AI Ethics, most notably criticising "technological mastery" being pursued as an end in itself and the lack of emphasis placed on operationalisation of principles.

^v The comparisons drawn in the remainder of the paper will hold even if AI Ethics initiatives were to move away from the four principles borrowed from medical ethics. The target of analysis here is a principled approach to ethics in AI development and governance, not of utility or justifiability of the four principles themselves. Many of the points addressed apply to professions featuring codes of conduct or codes of ethics that significantly shape practice within the profession and are not limited to medicine.¹⁴

2 The challenges of a principled approach to AI Ethics

When compared with medical ethics, four potential weaknesses of AI Ethics become apparent that suggest a principled approach may have limited impact on AI development. 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.

2.1 *Common aims and fiduciary duties*

Medicine is broadly guided by a common aim: to promote the health and well-being of the patient.¹⁶ It is a defining quality of a profession for its practitioners to be part of a ‘moral community’ with common aims, values, and training.^{17–19} While there is much disagreement over how best to promote health and well-being in practice, the interests of patients, medical professionals, and medical institutions remain aligned at a fundamental level which encourages solidarity and trust.^{vi} The pursuit of a common goal facilitates a principled approach to ethical decision-making.

Solidarity cannot be taken for granted in AI development. AI is largely developed by the private sector for deployment in public (e.g. criminal sentencing) and private (e.g. 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 benefit the company at the cost of user interests.^{20,21} While health professionals undoubtedly face organisational pressures similar in nature, 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 much more difficult in practice.

The implicit solidarity of medicine is formally recognised in professional codes of practice and medical law that establish fiduciary duties towards patients who must trust that health professionals will act in their best interests.²² Formal professions are defined by trust between clients and practitioners, mediated by common goals and values within the profession, and enforced through sanctions and self-governance.^{17,23,24} The fiduciary duties derived from the client-practitioner relationship separate ‘professions’ from other vocations,²⁵ and facilitate a principled approach to ethical decision-making by requiring practitioners to promote their clients’ best interests.

AI development is not a formal profession. An equivalent fiduciary relationship does not exist for the private sector practitioners and institutions that develop and deploy AI; nor do mechanisms with comparable weight to fiduciary duties.⁵ 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 interests.²³ For AI 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 have principal fiduciary duties towards their

^{vi} This is true of other professions and practices where a principled approach to ethics has been successful. Filipović et al. provide an overview of how this characteristic manifests in other professions.¹⁴

shareholders which can conflict with the interests of users and affected parties. Public interests are not granted primacy over commercial interests.

The lack of a fiduciary relationship means that users cannot trust that AI developers will act in their best interests when implementing ethical principles in practice. Affected parties without a ‘seat at the table’ cannot take for granted that their interests will be given serious consideration. Reputational risks may push companies to engage with ethics, but they carry weight only as long as they remain in the public consciousness. Personal moral conviction may also push AI developers towards ‘good’ behaviour;²⁶ recent examples of internal protests at Google provide some cause for hope.^{27,28} However, incentive structures discourage placing public interests before the company, and may include sanctions against ‘whistle-blowers’ (e.g. job loss, lack of advancement; see recent reports of backlash against protest organisers at Google).²⁹ Virtuous behaviour may therefore come at a high personal cost. This situation is unacceptable; users and affected parties should not need to 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 seriously.³⁰

2.2 Professional history and norms

The second weakness of a principled approach to AI Ethics is the lack of professional history and well-defined norms of ‘good’ behaviour. Medicine has a long history and shared professional culture with variation across regions and specialties. Norms of ‘good behaviour’ and codes of conduct have developed over centuries in these contexts. 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 ethics, and inspired the development of ethics and codes in other professions.¹⁴ As new technologies and treatments disrupt established norms, the profession can look back on a long ethical tradition to update best practices and its account of a ‘good’ health professional. Evidence of this historical development of an account of being a ‘good’ health professional can be seen in the American Medical Association’s Code of Medical Ethics, a lengthy document detailing opinions, behavioural norms, and standards across a plethora of medical practices and technologies.³¹ Over time, Western medical ethics has converged around key principles formalised in the ‘principilism’ moral framework, which defines four principles to guide ethical decision-making.¹⁵ This culture, and the values underlying it, has a strong regulating influence on the behaviour of medical practitioners.

AI development does not have a comparable history, homogenous professional culture, uniform professional identity, or a similarly developed professional ethics. Whereas AI can in principle be deployed in any context involving human expertise, medicine in comparison has narrower aims (see: Section 2.1) which facilitates development of standard practices and norms. Likewise, AI development includes practitioners from varied disciplines and professions such as software engineering and data science, which have incongruous histories, cultures, incentive structures, and ethics.^{14,23,24} Software engineering, which is the closest analogue, has historically not been legally recognised as a profession with fiduciary duties to the public^{14,32} due to a perceived lack of licensure schemes and a well-defined ‘standard of care’ for the profession.³³ Reflecting this, a comparably rich account of what it means to be a ‘good’ AI developer or software engineer does not exist; while the IEEE and ACM, two of the field’s largest professional associations, have published codes of ethics, these documents are

comparatively short and relatively lacking in grounded advice and specific behavioural norms.^{vii} As a result, it remains unclear what it means, in practice, to be a ‘good’ AI developer.

Developing such an account will not be easy. Systems are often created by large, multi-disciplinary and multi-national teams. Whereas the effects of clinical decision-making are often immediate and obvious, the impact of decisions taken in designing, training, and configuring an AI system for different uses may never become apparent to development teams.³⁴ This is worrying, as distance from potential victims has been shown to have a positive effect on unethical professional behaviour.³⁵ The risks addressed in medical ethics arise in reference to interventions performed on a ‘physical body’ (e.g. clinical or ‘bed-side interventions’). In comparison, the risks of AI and data ethics are continuous and not similarly bound, and may not be directly experienced by data subjects.³⁶ Systems are also often ‘opaque’ in the sense that no single person will have a full understanding of the system’s design or behaviour,³⁷ or be able to predict its behaviour. Even where problems are recognised, they can rarely be traced back to a single team member or action;³⁸ responsibility must be assigned across the collective network of actors and choices that influenced the system’s design, training, and configuration.^{viii} This uncertainty undermines the development of standards to be a ‘good’ AI developer, or a set of behaviours which will ultimately lead to ‘good’ or trustworthy AI, because the effects of development decisions cannot be reliably predicted at the time they are made.

AI Ethics initiatives aim to 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. This may be an impossible task.^{39,40} The great diversity of stakeholders and interests involved necessarily pushes the search for common values and norms towards a high level of abstraction.²³ 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’, or enables ‘human flourishing’, which are not specific enough to be action-guiding.⁸

Statements reliant on vague normative concepts hide points of political and ethical conflict. ‘Fairness’, ‘dignity’, ‘flourishing’ and other such abstract concepts are instances of “essentially contested concepts,” or concepts with many possible conflicting meanings requiring contextual interpretation through one’s background political and philosophical beliefs.⁴¹ These different interpretations, which can be rationally and genuinely held, lead to substantively different requirements in practice.⁴² These incompatible interpretations will only be revealed once principles or concepts are translated and tested in practice. 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

^{vii} It is worth noting, however, that the 2018 update to the ACM’s code of ethics placed greater emphasis on discussion of specific behaviours. However, even with this encouraging shift in focus, the code remains comparatively brief when contrasted with codes in other professions.¹⁴

^{viii} It is worth noting that these challenges arguably apply to individual medical professionals’ understanding of the ‘medical system’. However, clearly bounded clinical or ‘bed side’ interventions have much clearer and immediate impact.

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, left to developers to translate principles and specify essentially contested concepts as they see fit, without a clear roadmap for unified implementation. For an established profession with a rich history, ethical culture, and norms of ‘good’ practice to draw on to translate principles into practice, this ambiguity would be less concerning. Unfortunately, this is not the case for AI development.

2.3 *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’ interpretation of such concepts be resolved in practice?

Principles do not automatically translate into practice.²⁰ Throughout its history, medicine has developed effective ways of translating high-level commitments and principles into practical requirements and norms of ‘good’ practice.¹⁹ Professional societies and boards, ethics review committees, accreditation and licensing schemes, peer self-governance, codes of conduct and other mechanisms help determine the ethical acceptability of day-to-day practice by assessing difficult cases, identifying negligent behaviour, and sanctioning bad actors.^{43,44} The formal codes and informal norms that now govern medical practice have been extensively tested, studied, and revised over time, with their recommendations and norms (and underlying principles) evolving to remain relevant.

AI development does not have comparable empirically proven methods to translate principles into practice in real-world development contexts. This is a multi-faceted methodological challenge. Translation involves the specification of high-level principles into mid-level norms and low-level requirements. Norms and requirements can rarely be logically deduced directly from principles without accounting for specific elements of the technology, application, context of use, or relevant local norms.^{45,46ix} Rather, normative decisions must be made at each stage of translation.⁴⁶ 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.

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 justification, the prominence of essentially contested concepts in AI Ethics, and the field’s relative lack of a binding professional history (see: Section 2.2), conflicting practical requirements will likely

^{ix} Reflecting this, medical ethics committees rarely factor high-level principles into their deliberation; rather, the facts of the case, precedents and relevant intuitions take priority.⁴⁵ This finding accords with principlism, by which ethical decision-making ideally involves deliberation between case-specific facts and intuitions, precedents, and relevant background principles and theories to reach a consensus, or ‘reflective equilibrium’.¹⁵

emerge across the diverse sectors and contexts in which a principled approach to AI Ethics is used.

One other methodological challenge remains. Thus far, we have assumed that norms and normative practical requirements can be successfully procedurally embedded in development, and functionally implanted in design requirements. Neither can be taken for granted. Prior work in computer ethics and science and technology studies points to the difficulty of embedding ethical values and principles in technology design and the development cycle.^{1,46-48} Many such methods exist,^x including participatory design, reflective design, Values@Play, and Value-Sensitive Design,¹ but thus far they have largely been implemented and studied in academic contexts which are more receptive to normative concerns than commercial settings.^{21,49,50} Value-conscious methods are also largely procedural, not functional. Generally speaking, they introduce values, normative issues, and relevant stakeholders into the development process.⁴⁶ 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 specifications of essentially contested concepts.^{47xi}

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, resolving conflicts between different specifications of essentially contested concepts, and similar methods add additional work and costs to the research and development process. Unsurprisingly, ethical considerations may be discarded when they conflict with commercial incentives.²¹ It cannot be assumed that value-conscious frameworks will be meaningfully implemented in commercial processes that value efficiency, speed, and profit.

2.4 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 which uphold professional standards and provide patients with redress for negligent behaviour, including malpractice law, licensing and certification schemes, ethics committees, and professional medical boards.^{43,44} Legally supported accountability mechanisms provide an external impetus for health professionals to fulfil their fiduciary duties, amplifies the impact of complementary forms of self-governance by establishing a clear link between ‘bad’ behaviour and professional sanctions (e.g. losing one’s license to practice),¹⁴ mandates a professional standard of care, and allows patients to make claims against negligent members of the profession.

Excluding certain types of risks (e.g. privacy violations governed by data protection law), AI development does not have comparable legally and professionally endorsed accountability

^x Morley et al. provide an up-to-date overview of tools, methods, and research to translate AI Ethics principles into practice.⁴⁹

^{xi} Other value-specific methods, such as ‘privacy by design’ or ‘security by design’ are more successful in this regard, in part because they address technically tractable normative concepts.

mechanisms.^{xii} 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 turned up mixed results. Codes are often followed in letter rather than spirit, or as a ‘checklist’ rather than as part of a critical reflexive practice.^{23,52,53} A recent study of the Association for Computing Machinery’s (ACM) Code of Ethics revealed that it has little effect on the day-to-day decision-making of software engineering professionals and students.⁵⁴ Other studies of corporate and professional codes of ethics outside computing have reported similar results.^{55–57} A recent meta-analysis of evidence on the impact of codes on professional behaviour^{xiii} 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 organisational culture and actively enforced.^{26,58,59} Norms must be clearly defined and highly visible if they are to influence practitioners¹⁴ and inspire peer self-governance. Current governance structures in AI companies are insufficient in this regard.²

External sanctions also play an important role. Studies of other professions have shown that the existence of sanctions for breaching a code of ethics is key for adherence and effective self-governance.¹⁴ Compared with medicine, information professions lack sanctions that can impact the professional’s livelihood.²³ While software engineering degrees can now be accredited, a license is not required to practice,⁶⁰ with some national exceptions.^{61–63} Professional bodies such as the IEEE and ACM lack formal sanction powers beyond expulsion from the organisation. As licensing is not provided by either organisation, their impact on a developer’s livelihood or the ability to practice is limited.

The lack of empirical evidence linking codes of ethics to actual impact on developer behaviour raises a difficult question: is it enough to define ‘good intentions’ and hope for the best? Without complementary punitive mechanisms and bodies providing appeal and redress that can step in’ when self-governance fails, a principled approach runs the risk of merely providing false assurances of ethical or trustworthy AI.⁶⁴

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 long-standing history and harmonised aims (see: Sections 2.1 and 2.2). Professional accountability mechanisms are created to protect clients and the public,⁶⁵ but AI developers do not provide a public service and predominantly work in commercial institutions. As a result, public interests do not need to be given provisional primacy over competing commercial interests. AI also does not operate in a single sector, meaning any new legal or professional mechanisms will need to account for many different types of possible benefits and harms, and integrate with existing sector-specific law and policy. Finally, proposals to introduce professional sanctions and

^{xii} While contract law can provide legal accountability in some cases, it is not directly comparable to medical malpractice or tort which define negligence against a heightened ‘professional standard of care’ to be upheld by individual practitioners. Such a standard has not been recognised historically for computing professionals.⁵¹

^{xiii} It should be noted that the effect studied is on reducing unethical behaviour, which is a distinct phenomenon from the positive encouragement of ethical behaviour and embedding of values in technology design and governance which is pursued in AI Ethics.

licensing schemes for computing professionals are also not new, but have thus far seen limited uptake.^{61,63,66xiv}

3 Where should AI Ethics go from here?

The four critical challenges facing AI Ethics are now clear. First, AI development is not a formal profession with aims that align with public interests (Table 1 reviews five criteria traditionally required for legal recognition).^{xv} Fiduciary duties are owed principally to shareholders, not the public. 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. It remains unclear how developers should specify essentially contested concepts into practical behavioural norms and system requirements. And finally, when a developer falls afoul of these vaguely defined requirements, there remain few sanction mechanisms and channels for redress to set things right. 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.

Table 1 - Characteristics of a Formal Profession

1. **Specialised education and training** - Members are expected to have undertaken extensive specialised education and training,⁵¹ typically in accredited degree programmes.³³
2. **Commitment to public service** - Professions involve a public declaration to provide a service to society or for the public good making use of specialised, often privileged expertise⁶⁵ which takes precedence over individual gain.⁵¹
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 public.⁵¹
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}

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 AI’ or ‘Ethical AI’ in the future. Without a fundamental shift in regulation,

^{xiv} In particular, see the work of Donald Gotterbarn from the 1990’s and early 2000’s.⁶⁷

^{xv} A classic legal definition of a profession was provided in *Hospital Computer Systems, Inc. v. Staten Island Hosp.*, 788 F.Supp. 1351, (D.NJ 1992): “A profession is not a business. It is distinguished by the requirements of extensive formal training and learning, admission to practice by a qualifying licensure, a code of ethics imposing standards qualitatively and extensively beyond those that prevail or are tolerated in the marketplace, a system for discipline of its members for violation of the code of ethics, a duty to subordinate financial reward to social responsibility, and, notably, an obligation on its members, even in non-professional matters, to conduct themselves as members of a learned, disciplined, and honorable occupation...Professionals may be sued for malpractice because of the higher standards of care imposed on them by their profession and by state licensing requirements engenders trust in them by clients that is not the norm of the marketplace. When no such higher code of ethics binds a person, such trust is unwarranted.” Other judgements from the United States that have not recognised software engineering and related vocations as a profession include: *Chatlos Systems, Inc. v. National Cash Register Corp*, *Triangle Underwriters, Inc. v. Honeywell, Inc.*, *Hospital Computer Systems, Inc. v. Staten Island Hospital*, *RKB Enterprise, Inc. v. Ernst and Young*. However, some progress towards formal professionalisation occurred in *Data Processing Services, Inc. v. L.H. Smith Oil Corp.* and *Diversified Graphics, Ltd. v. Groves*.

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 worth 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, government, and civil society?

1. 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: Table 2 for key questions). Self-regulatory frameworks must be embedded and highly visible in organisational culture to be effective.²⁴ Binding accountability structures as well as clear implementation and review processes are needed at a sectoral and organisational level. While select AI companies have created (and disbanded) ethics committees,⁶⁸ their remit, independence, and decision-making power is rarely clear. Publicising decisions of internal ethics committees could, for example, clarify the impact of principles on AI development and governance.¹⁷ Professional norms can be established by defining clear requirements for inclusive design, documentation of models and datasets, and independent auditing for bias, discrimination, and other ethical concerns.

Table 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?

2. Support 'bottom-up' AI Ethics

A 'top-down' approach to AI Ethics is uniquely difficult due to the diversity of technologies described as 'AI'. Inevitably, principles created to govern such a broad category of technologies are vague (see: Section 2.3). In such a diverse field, a bottom-up, case-based approach to ethics may be more effective. Professional ethics have historically developed in medicine and engineering following precisely this path; local practices and lessons emerge and spread across the field and industry from which principles and precedents can then be derived.¹⁴ Novel cases reveal new challenges for AI Ethics, which are desperately needed to move the field beyond well-worn cases^{11,69} and develop sector- and case-specific guidelines, technical solutions, and an empirical knowledge base. To complement the considerable top-down work already undertaken, increased support and access to development settings should be made available to support multi-disciplinary bottom-up research and development in AI Ethics. The recent rapid growth of multi-disciplinary research networks addressing ethical, social, and legal implications of AI (e.g. FAT-ML) gives cause for optimism.

3. 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.^{xvi} Doctors, lawyers, and other professions in the public service require a license to practice. It is a regulatory oddity that we license these 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. Licensing initiatives could initially target developers of systems with elevated risk or built for the public sector, such as facial recognition systems designed for policing.

4. Shift from professional ethics to business ethics

The outputs of many AI Ethics initiatives resemble professional codes of ethics that address design requirements and the behaviours and values of individual professions.¹ The legitimacy of particular applications and their underlying business interests remain largely unquestioned.^{1,71} This approach conveniently steers debate towards the transgressions of unethical individuals, and away from the collective failure of unethical businesses and business models.⁷² Developers will always be constrained by the institutions that employ them. To be truly effective, the ethical challenges of AI cannot be conceptualised as individual failures. Going forward, AI Ethics must become an ethics of AI businesses as well.

5. 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 (e.g. privacy, fairness),¹ but rarely propose technical definitions or explanations.^{5xvii} The rationale seems to be as follows: insufficient consideration of ethics leads to poor design decisions which create systems that harm users. Framing ethical challenges in terms of design flaws ensures they remain “fundamentally technical, shielded from democratic intervention” or regulation.¹ The need to translate high-level principles and essentially contested concepts into practical requirements, which can be difficult and time-consuming in practice, is seemingly avoided.

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. It is foolish to assume that very old and complex normative questions can be solved with technical fixes or ‘good’ design alone.^{xviii} The risk is that complex, difficult ethical debates will be oversimplified to make the concepts at hand computable and implementable in a straightforward but

^{xvi} Formal public service obligations could also be recognised for commercial AI companies. A related proposal in the United States involves re-conceptualising data-driven companies as “information fiduciaries” with formal fiduciary duties towards data subjects.⁷⁰

^{xvii} Initiatives such as the IEEE ‘Ethically-Aligned Design’ programme, which has gone to great lengths to begin translating high-level principles into practical requirements and technical standards for AI, are a rare exception to this trend.⁶

^{xviii} Current challenges faced in the field support this position: even for concepts such as ‘fairness’ that have seen significant research and development, consensus has not been reached and implementation outside research environments remains limited.⁷³ This situation is particularly worrying because ‘fairness’ is perhaps the most tractable principle currently addressed in AI Ethics; fairness concerns the distribution of treatments and outcomes, which allows it to be easily (but perhaps over simplistically) quantified and measured. It is the essentially contested concept perhaps most amenable to ‘technical fixes’, and yet it remains underdeveloped and under implemented.

conceptually shallow manner.¹¹ Concepts are like physical tools; they wear out and must be philosophically reconstructed in different contexts and in response to new technologies.⁴¹

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.

References

1. Greene, D., Hoffmann, A. L. & Stark, L. Better, Nicer, Clearer, Fairer: A Critical Assessment of the Movement for Ethical Artificial Intelligence and Machine Learning. *10* (2019).
2. Whittaker, M. *et al.* *AI Now Report 2018*. (AI Now Institute, 2018).
3. AI Ethics Guidelines Global Inventory. *AlgorithmWatch*
4. Nemitz, P. Constitutional democracy and technology in the age of artificial intelligence. *Philos. Trans. R. Soc. Math. Phys. Eng. Sci.* **376**, 20180089 (2018).
5. Hagendorff, T. The Ethics of AI Ethics -- An Evaluation of Guidelines. *ArXiv190303425 Cs Stat* (2019).
6. The IEEE Global Initiative on Ethics of Autonomous and Intelligent Systems. *Ethically Aligned Design*. (IEEE, 2019). doi:10.1007/978-3-030-12524-0_2
7. Calo, R. Artificial Intelligence Policy: A Primer and Roadmap Symposium - Future-Proofing Law: From RDNA to Robots (Part 2). *UC Davis Law Rev.* **51**, 399–436 (2017).
8. Whittlestone, J., Nyrup, 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).
9. Floridi, L. *et al.* AI4People—An Ethical Framework for a Good AI Society: Opportunities, Risks, Principles, and Recommendations. *Minds Mach.* **28**, 689–707 (2018).
10. OECD. Forty-two countries adopt new OECD Principles on Artificial Intelligence - OECD. *Forty-two countries adopt new OECD Principles on Artificial Intelligence* (2019). Available at: <http://www.oecd.org/science/forty-two-countries-adopt-new-oecd-principles-on-artificial-intelligence.htm>. (Accessed: 22nd May 2019)
11. European Group on Ethics in Science and New Technologies. *Statement on Artificial Intelligence, Robotics and ‘Autonomous’ Systems*. (European Commission, 2018).
12. European Group on Ethics in Science and New Technologies. Open letter to the European Commission. (2019).
13. High Level Expert Group on Artificial Intelligence. *Ethics Guidelines for Trustworthy AI*. (European Commission, 2019).
14. Filipović, A., Koska, C. & Paganini, C. *Developing a Professional Ethics for Algorithmists: Learning from the Examples of Established Ethics*. (Bertelsmann Stiftung, 2018).
15. Beauchamp, T. L. & Childress, J. F. *Principles of biomedical ethics*. (Oxford University Press, 2009).
16. Marshall, T. H. The recent history of professionalism in relation to social structure and social policy. *Can. J. Econ. Polit. Sci. Can. Econ. Sci. Polit.* **5**, 325–340 (1939).
17. Frankel, M. S. Professional codes: Why, how, and with what impact? *J. Bus. Ethics* **8**, 109–115 (1989).

18. *Black's law dictionary*. (West, 2009).
19. MacIntyre, A. *After Virtue: A Study in Moral Theory*. (Gerald Duckworth & Co Ltd, 2007).
20. Van den Bergh, J., Deschoolmeester, D., Deschoolmeester, D. & Vlerick Leuven Gent Management School, Gent, Belgium. Ethical Decision Making in ICT: Discussing the Impact of an Ethical Code of Conduct. *Commun. IBIMA* 1–11 (2010). doi:10.5171/2010.127497
21. Manders-Huits, N. & Zimmer, M. Values and Pragmatic Action: The Challenges of Introducing Ethical Intelligence in Technical Design Communities. **10**, 8 (2009).
22. Pellegrino, E. D. & Thomasma, D. C. *The virtues in medical practice*. (Oxford University Press, 1993).
23. Iacovino, L. Ethical Principles and Information Professionals: Theory, Practice and Education. *Aust. Acad. Res. Libr.* **33**, 57–74 (2002).
24. Boddington, P. *Towards a code of ethics for artificial intelligence research*. (Springer Berlin Heidelberg, 2017).
25. McDowell, B. *Ethical Conduct and the Professional's Dilemma: Choosing Between Service and Success*. (Quorum Books, 1991).
26. 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).
27. Wakabayashi, D. & Shane, S. Google Will Not Renew Pentagon Contract That Upset Employees. *The New York Times* (2018).
28. Conger, K. & Wakabayashi, D. Google Employees Protest Secret Work on Censored Search Engine for China. *The New York Times* (2018).
29. Wong, J. C. Demoted and sidelined: Google walkout organizers say company retaliated. *The Guardian* (2019).
30. Parker, D. B. *Ethical Conflicts in Computer Science and Technology*. (AFIPS Press, 1981).
31. Brotherton, S., Kao, A. & Crigger, B. J. Professing the Values of Medicine: The Modernized AMA Code of Medical Ethics. *JAMA* **316**, 1041–1042 (2016).
32. Panensky, S. A. & Jones, R. Does IT Go Without Saying? *Prof. Times Summer 2018*, 30–33 (2018).
33. Perlman, D. T. Who Pays the Price of Computer Software Failure Notes and Comments. *Rutgers Comput. Technol. Law J.* **24**, 383–416 (1998).
34. Mittelstadt, B., Allo, P., Taddeo, M., Wachter, S. & Floridi, L. The ethics of algorithms: Mapping the debate. *Big Data Soc.* **3**, (2016).
35. Jones, T. M. Ethical Decision Making by Individuals in Organizations: An Issue-Contingent Model. *Acad. Manage. Rev.* **16**, 366–395 (1991).
36. Metcalf, J. & Crawford, K. Where are human subjects in Big Data research? The emerging ethics divide. *Big Data Soc.* **3**, 2053951716650211 (2016).
37. Burrell, J. How the Machine 'Thinks:' Understanding Opacity in Machine Learning Algorithms. *Big Data Soc.* (2016). doi:10.1177/2053951715622512
38. Floridi, L. Faultless responsibility: on the nature and allocation of moral responsibility for distributed moral actions. *Philos. Trans. R. Soc. Math. Phys. Eng. Sci.* **374**, 20160112 (2016).
39. Awad, E. *et al.* The Moral Machine experiment. *Nature* **563**, 59 (2018).
40. Ess, C. Ethical pluralism and global information ethics. *Ethics Inf. Technol.* **8**, 215–226 (2006).
41. van den Hoven, J. Computer Ethics and Moral Methodology. *Metaphilosophy* **28**, 234–248 (1997).

42. Gallie, W. B. Essentially Contested Concepts. *Proc. Aristot. Soc.* **56**, 167–198 (1955).
43. 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).
44. Papadakis, M. A. *et al.* Disciplinary Action by Medical Boards and Prior Behavior in Medical School. *N. Engl. J. Med.* **353**, 2673–2682 (2005).
45. Toulmin, S. How medicine saved the life of ethics. *Perspect. Biol. Med.* **25**, 736–750 (1982).
46. van de Poel, I. Translating Values into Design Requirements. in *Philosophy and Engineering: Reflections on Practice, Principles and Process* (eds. Michelfelder, D. P., McCarthy, N. & Goldberg, D. E.) **15**, 253–266 (Springer Netherlands, 2013).
47. Friedman, B., Hendry, D. G. & Borning, A. A Survey of Value Sensitive Design Methods. *Found. Trends® Human–Computer Interact.* **11**, 63–125 (2017).
48. Friedman, B. & Kahn, P. H. Human agency and responsible computing: Implications for computer system design. *J. Syst. Softw.* **17**, 7–14 (1992).
49. 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. *ArXiv190506876 Cs* (2019).
50. Shilton, K. Values Levers: Building Ethics into Design. *Sci. Technol. Hum. Values* **38**, 374–397 (2013).
51. 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).
52. Bynum, T. W. Ethical decision making and case analysis in computer ethics. in *Computer ethics and professional responsibility* (eds. Bynum, T. W. & Rogerson, S.) 60–87 (Blackwell, 2004).
53. Ladd, J. The quest for a code of professional ethics: an intellectual and moral confusion. in *Ethical issues in the use of computers* (eds. Johnson, D. G. & Snapper, J. W.) 8–13 (Wadsworth Publ. Co., 1985).
54. McNamara, A., Smith, J. & Murphy-Hill, E. Does ACM’s code of ethics change ethical decision making in software development? in *Proceedings of the 2018 26th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering - ESEC/FSE 2018* 729–733 (ACM Press, 2018). doi:10.1145/3236024.3264833
55. 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).
56. Helin, S. & Sandström, J. An Inquiry into the Study of Corporate Codes of Ethics. *J. Bus. Ethics* **75**, 253–271 (2007).
57. 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).
58. 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).
59. Shilton, K. “That’s Not An Architecture Problem!?”: Techniques and Challenges for Practicing Anticipatory Technology Ethics. 7 (2015).
60. Goertzel, K. M. Legal liability for bad software. *CrossTalk* **23**, (2016).
61. Laplante, P. A. Licensing professional software engineers: seize the opportunity. *Commun ACM* **57**, 38–40 (2014).

62. Pour, G., Griss, M. L. & Lutz, M. The push to make software engineering respectable. *Computer* **33**, 35–43 (2000).
63. Seidman, S. B. The Emergence of Software Engineering Professionalism. in *E-Government Ict Professionalism and Competences Service Science* (eds. Mazzeo, A., Bellini, R. & Motta, G.) **280**, 59–67 (Springer US, 2008).
64. Suchman, L. Corporate Accountability. *Robot Futures* (2018).
65. Abbott, A. Professional Ethics. *Am. J. Sociol.* **32** (1983).
66. O'Connor, J. E. Computer Professionals: The Need for State Licensing. *Jurimetr. J.* **18**, 256–267 (1978).
67. Gotterbarn, D. The Ethical Computer Practitioner—Licensing the Moral Community: A Proactive Approach. *SIGCSE Bull* **30**, 8–10 (1998).
68. Olson, P. Google Quietly Disbanded Another AI Review Board Following Disagreements. *Wall Street Journal* (2019).
69. Angwin, J., Larson, J. & Kirchner, L. Machine Bias. *ProPublica* (2016).
70. Balkin, J. M. Information Fiduciaries and the First Amendment. *UCDL Rev* **49**, 1183 (2015).
71. Benkler, Y. Don't let industry write the rules for AI. *Nature* **569**, 161–161 (2019).
72. 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**, (2019).
73. Holstein, K., Vaughan, J. W., Daumé III, H., Dudík, M. & Wallach, H. Improving fairness in machine learning systems: What do industry practitioners need? *ArXiv181205239 Cs* (2018). doi:10.1145/3290605.3300830