

Machine Ethics: The Design and Governance of Ethical AI and Autonomous Systems

By **ALAN F. WINFIELD** 

Guest Editor

KATINA MICHAEL 

Guest Editor

JEREMY PITT

Guest Editor

VANESSA EVERS

Guest Editor

I. INTRODUCTION

The so-called fourth industrial revolution and its economic and societal implications are no longer solely an academic concern, but a matter for political as well as public debate. Characterized as the convergence of robotics, AI, autonomous systems and information technology – or cyber-physical systems – the fourth industrial revolution was the focus of the World Economic Forum, at Davos, in 2016 [1]. Also in 2016 the US White House initiated a series of public workshops on artificial intelligence (AI) and the creation of an interagency working group, and the European Parliament Committee for Legal Affairs published a draft report with recommendations to the Commission on Civil Law Rules on Robotics.

Since 2016 there has been a proliferation of ethical principles in AI; seven sets of principles were published in 2017 alone – including the influential Future of Life Institute’s Asilomar principles for beneficial AI. Nor has there been a shortage of industry initiatives, perhaps the most notable is the Partnership in AI; founded in 2016 by representatives of Apple, Amazon, DeepMind and Google, Facebook, IBM, and Microsoft, the partnership now has more than 50 members. Most significant however has been the number of national AI initiatives announced [2]. Since Canada published its national AI strategy early

The papers in this special issue focus on various aspects of machine ethics, including the central question of how autonomous systems can be imbued with ethical values.

in 2017 more than 20 other countries (and groupings including the EU) have announced similar initiatives. The largest and most comprehensive is undoubtedly China’s Next Generation AI development plan, launched in July 2017.

Notably all of these initiatives express the need for serious consideration of the ethical and societal implications of robotics and artificial intelligence. Robot and AI ethics has been transformed from a niche area of concern of a few engineers, philosophers and law academics, to an international debate. For these reasons we believe this special issue – focused on the ethics of intelligent autonomous systems – is not only timely but necessary.

The primary focus of this special issue is machine ethics, that is the question of how autonomous systems can be imbued with ethical values. Ethical autonomous systems are needed because, inevitably, near future systems are moral agents; consider driverless cars, or medical diagnosis AIs, both of which will need to make choices with ethical consequences. This special

Digital Object Identifier 10.1109/JPROC.2019.2900622

issue includes papers that describe both implicit ethical agents, that is machines designed to avoid unethical outcomes, and explicit ethical agents: machines which either encode or learn ethics and determine actions based on those ethics. Of course ethical machines are socio-technical systems thus, as a secondary focus, this issue includes papers that explore the societal and regulatory implications of machine ethics, including the question of ethical governance. Ethical governance is needed in order to develop standards and processes that allow us to transparently and robustly assure the safety of ethical autonomous systems and hence build public trust and confidence.

II. THE LANDSCAPE OF ROBOT AND AI ETHICS

The field of robot and AI ethics is broadly divided into two main branches. By far the largest of these is concerned with the vitally important question of how human developers, manufacturers and operators should behave in order to minimize the ethical harms that can arise from robots and AIs in society, either because of poor design, inappropriate application, or misuse. This branch is concerned with the ethical application of robots and AIs and is generally referred to as either AI ethics or robot ethics, and has already led to the development of ethical principles [3], [4], standards [5], and proposals for good practice [6], [7].

The smaller branch of AI ethics is concerned with the question of how robots and AIs can themselves behave ethically. Referred to as ethical AI, ethical robots or – more generally – *machine ethics*, the field spans both philosophy and engineering. Philosophers are concerned with questions such as “could a machine be ethical, and if so which ethics should determine its behavior?” alongside larger questions such as “should society delegate moral responsibility to its machines?” while engineers are interested in solving the (significant) technical problem of how to build

an ethical machine. The two disciplines are not disconnected; philosophers are interested in the outcomes of practical machine ethics not least because – if successful – they lend urgency to the moral questions around ethical machines in society. Equally, engineers engaged in designing ethical machines need philosophers to advise on the definition of appropriate ethical rules and values for these machines.

Of course the idea that robots should not only be safe but also actively capable of preventing humans from coming to harm has a long history in science fiction. In his short story *Runaround*, Asimov [8] expressed such a principle in his now well-known Laws of Robotics. Although no-one has seriously proposed that real-world robots should be “three-laws safe,” work in machine ethics has advanced the proposition that future robots should be more than just safe. In their influential book *Moral Machines – Teaching Robots Right From Wrong*, Wallach and Allen [9] set out the philosophical foundations of machine ethics, coining the term Artificial Moral Agent (AMA); Wallach and Allen write:

“If multipurpose machines are to be trusted, operating untethered from their designers or owners and programmed to respond flexibly in real or virtual world environments, there must be confidence that their behaviour satisfies appropriate norms. This goes beyond traditional product safety ... if an autonomous system is to minimise harm, *it must also be ‘cognisant’ of possible harmful consequences of its actions, and it must select its actions in the light of this ‘knowledge’*, even if such terms are only metaphorically applied to machines” (italics added).

Since Čapek’s 1920 play *Rossum’s Universal Robots*, and the stories of Asimov, Frayn [10] and many others,

science fiction narratives have played a key role in the exploration of artificial morality [11]. In the study of machine ethics we can, for the first time, begin to investigate artificial morality in fact.

III. THE STATE OF THE ART IN MACHINE ETHICS

Machine ethics is a new field. Although the antecedents of machine ethics can be found in computer ethics [12], the earliest works on machine ethics were published less than 20 years ago. Those works are, not surprisingly, theoretical and philosophical. The field of machine ethics was *de facto* established by Allen *et al.* [13], [14], Asaro [15], Moor [16], Powers [17], Anderson and Anderson [18], [19], and Wallach and Allen [9].

The earliest proposal for a practical *ethical governor* for robots – a mechanism for moderating or inhibiting a robot’s behavior to prevent it from acting unethically – was from Arkin [20],¹ although not tested on real robots. At the time of writing the number of experimental demonstrations of ethical robots remains very small indeed; to the best of our knowledge there have been only five such demonstrations to date: (i) the GenEth system of Anderson and Anderson [22], (ii) the Asimovian ethical robots of Winfield *et al.* [23] and Vanderelst and Winfield [24], (iii) Bringsjord *et al.*’s Akratic robot [25], (iv) the “sorry I can’t do that” robot of Briggs and Scheutz [26], and (v) the Intervening robot mediator in healthcare of Shim, Arkin and Pettinatti [27]. Papers which review and update the approaches of (i) and (ii) are included in this issue; we now briefly review (iii - v).

Based on earlier work proposing a general “logistic” method for engineering ethically correct robots [28], Bringsjord’s Akratic² robot is tested

¹The notion of a governor in autonomous systems dates back to the birth of cybernetics (from *Kyvernitis* - to govern), see [21].

²From the Greek *akrasia* referring to when a person acts in contradiction to their better judgment.

in scenarios in which it is charged with guarding a prisoner of war and must choose between retaliating with violence to an attack (thus satisfying a self-defense goal) or refraining from retaliation [25]. The robot's "ethical substrate model" incorporates logic that specifies which actions are forbidden under certain circumstances, or which are permitted or obligatory; importantly that logic, a deontic cognitive event calculus (DCEC), can be formally verified.

Briggs and Scheutz [26] demonstrate a robot capable of declining to carry out an instruction, if it reasons that to do so would harm itself. The robot may more properly be described as implicitly ethical since it is designed only to avoid unethical outcomes to itself, but the cognitive machinery it uses to avoid such outcomes is, like Bringsjord *et al.*'s robot [25], based on reasoning about obligation and permissibility.

In the work of Shim *et al.* [27] a robot equipped with an ethical governor monitors the interaction between a patient and healthcare advisor. The robot is able to sense (i) the loudness of the patient's speech, (ii) when the patient stands up to leave, and (iii) when the advisor asks the patient to take some medicine. In four test scenarios: "yelling," "quiet," "stay-in-the-room," and "safety-first," the robot senses the patient advisor interaction and uses both speech and gesture to intervene. In the safety-first scenario the robot might alert the advisor with "the previous records say he had a reaction (to this medicine). I think it's not safe."

Allen *et al.* [14] identified three approaches to machine ethics: top-down, bottom-up and a hybrid of top-down and bottom-up. The top-down approach requires training the machine to be able to recognize and correctly respond to morally challenging situations. The bottom up approach instead constrains the actions of the machine in accordance with pre-defined rules or norms. Of the five experimental demonstrations listed here only the first: Anderson and Anderson's GenEth system,

adopts a top-down approach. The other four – although very different in detail – are all examples of bottom-up constraint-based approaches.

All five trials are of robots with very limited ethics in constrained laboratory settings; they each demonstrate proof of concept but are far from practical application in real-world settings. Nevertheless we can be confident that it is possible, at least in principle, to build minimally ethical robots. Later in this article we shall consider the question of the degree of ethical agency of these robots.

All of the work above has, not surprisingly, focused on the core processes – the cognitive architectures – for making ethical choices. But many technical challenges remain: how, for instance, do we design a robot or AI that is "cognisant of possible harmful consequences of its actions?" This requires that a machine can reliably recognize when and how its actions have ethical salience; for a robot interacting with a human that means that it needs to perceive that the human is at risk – and the nature of that risk. The accurate perception of both risk and context is non trivial even in controlled environments – and even more challenging in real-world settings. Context might also change the risk calculation; an elderly person who is frail is clearly at greater risk than one who is physically fit, and the presence of other humans (a nurse for instance) clearly also changes the context for a care robot.

If and when ethical machines are ready for real-world application we would need to be sure that their ethical decision-making processes are both guaranteed and safeguarded against misuse. We would also need robust frameworks for ethical governance, including technical standards for ethical machines alongside processes of accident investigation. And bigger societal questions would need to be addressed around the extent to which we are prepared to delegate moral responsibility to our machines.

IV. DEFINING MACHINE ETHICS

Although not providing a singular definition of machine ethics Moor's influential 2006 paper: "The nature, importance, and difficulty of machine ethics" [16], defines the field by articulating four categories of ethical agency. These are:

- 1) *Ethical Impact Agents*: Any machine that can be evaluated for its ethical consequences.
- 2) *Implicit Ethical Agents*: Machines that are designed to avoid unethical outcomes.
- 3) *Explicit Ethical Agents*: Machines that can reason about ethics.
- 4) *Full Ethical Agents*: Machines that can make explicit moral judgments and justify them.

In this special issue we follow Anderson and Anderson's 2006 definition [18]. To paraphrase: *an ethical machine is guided by an ethical rule, or set of rules, in deciding how to act in a given situation.* It follows that we are concerned here with *autonomous* machines: either software AIs, or their physically embodied counterpart, robots, which determine how to respond to input without direct human control. Although labeled as autonomous such systems are generally subject to human supervision or monitoring (and intervention if necessary) in a way that is more properly described as "supervised autonomy"; but what is important is that low-level decisions are made by the system rather than a supervising human. In this special issue we are concerned with ethical systems that fall into both Moor's second and third categories of implicit and explicit ethical agents.

All of the five demonstrations of ethical robots referenced in the previous section are certainly implicit and, arguably, explicit ethical agents, since they all have either learned or defined ethical rules and the cognitive machinery to take those rules into account when deciding how to act in a given situation and – if necessary – proactively acting or intervening to prevent harm; a process we can arguably describe as ethical reasoning

(albeit of a limited kind). It is important to note that Moor's definitions of ethical agency, and in particular the distinction between implicit and explicit agency, remain controversial [29], as does the question of which systems properly qualify as explicit ethical machines [30].

We would argue that *all* non-trivial examples of real-world AIs and robots are ethical impact agents as defined by Moor. The ethical impact of an online store's AI in offering suggestions for purchases may be slight, but there can be no doubt that search engines and AIs that determine which social media posts and advertisements appear on your home page can have significant societal and political impact during, for instance, election campaigns [31].

Consider AIs that recommend loans, determine welfare payments, or suggest prison sentences. We only need to reflect on the impact of an incorrect decision to see that these are systems with ethical impact. Indeed bias in AIs trained with uncurated data sets is now a well-known problem [32]. Even chatbots, which may appear inconsequential, can fall victim to gaming by users and quickly become deeply offensive [33], [34]. Connected AI-toys designed to interact and converse with children are even more worrying. More obvious examples of systems with clear ethical impact include medical diagnosis AIs, assisted-living (care) or companion robots, and driverless cars. All of these systems have the potential for negative ethical impact, either as a result of poor design or simply a lack of forethought about how the system might be used [35]. And few systems, regardless of how well designed, are immune to malicious use [36].

We propose that *all* robots and AIs of the kinds mentioned here should be designed to avoid negative ethical impacts; in other words they should be designed to be implicit ethical agents within Moor's schema. Frameworks for the design and test of implicit ethical agents are now emerging; the IEEE Standards Association document *Ethically Aligned*

Design [6] and associated standards currently in draft such as IEEE P7000 *Model Process for Addressing Ethical Concerns During System Design*, together provide a methodology for imbuing ethical values into intelligent systems. Those values are then expressed implicitly in such a way that the system meets the highest standards of, for instance, transparency and explainability (P7001), privacy (P7002) and is – as far as possible – free of bias (P7003). Adamson *et al.* outline these and other IEEE ethics initiatives in this special issue [47].

V. HOW ETHICAL ARE CURRENT EXPLICITLY ETHICAL MACHINES?

One objection frequently leveled at the ethical machines demonstrated thus far is that they simply extend the envelope of safe behavior, and that they should not be described as ethical, but simply as “safety plus” machines.

To counter this objection let us consider a simple thought experiment: you are walking on the street and notice a child who is not looking where she's going (perhaps engrossed in her smart phone); you see that she is in imminent danger of walking into a large hole in the pavement. Suppose you act to prevent her falling into the hole. Most bystanders would regard your action as that of a good person – in more extreme circumstances you might be lauded a hero. Of course your action *does* keep the child safe, but it is also a moral act. Your behavior is consequentialist ethics in action because you have (a) anticipated the consequences of her inattention and (b) acted to prevent a calamity. Your act is ethical because it is an expression of care for another human's safety and well-being. To claim that ethical robots merely exhibit safety plus behavior is to miss the key point that it is the *intention* behind an act that makes it ethical. Of course, robots and AIs don't have intentions but their designers do, and an ethical machine is an instantiation of those good intentions.

Moral philosophers are also, and perhaps not surprisingly, doubtful over claims that any machines can be described as ethical. Some argue that morality is the exclusive preserve of humans. Of course humans are not the only animals that demonstrate altruism, but we are almost certainly the only species capable of both consciously reflecting upon and justifying the morality of our actions.

One moral philosopher [38] offered the following opinion on the minimally ethical robots described in [23]:

“On the assumption that consequences have at least some significance in our moral deliberations, you can claim that your robot is capable of attending to one kind of moral consideration, even if you don't make the much stronger claim that it is capable of choosing the right action all things considered.”

It does therefore appear to be reasonable to make the limited claim that the ethical robots demonstrated in the laboratory are at least “capable of attending to one kind of moral consideration.”

We need to make the distinction between what might in principle be achievable in the near future and the far future. For machines to be as good as humans at moral reasoning they would need, to use Moor's terminology, to be full ethical agents [16]. The best we have demonstrated to date is a handful of proof-of-concept explicit ethical agents and even those, as we have commented, are only minimally ethical. In no sense are such agents better at moral reasoning than humans.

Does the fact that we have arguably reached the third category in Moor's scheme (explicit ethical agents) mean that full ethical agents are on the horizon? The answer again must be “no.” The scale of Moor's scheme is not linear. It is a relatively small step from ethical impact agents to implicit ethical agents, then a very much bigger and more difficult step to explicit ethical agents, which we

are only just beginning to witness. But then there is a huge gulf to full ethical agents, since they would almost certainly need something approaching human equivalent AI; full explicit ethical agents, by Moor's definition [16], require consciousness, intentionality, and free will. Indeed, we think it is appropriate to think of explicit ethical machines as examples of narrow AI; full ethical agents would almost certainly require artificial general intelligence (AGI).

A. The Ethical Turing Test

Several have proposed a test for ethical machines somewhat akin to the Turing Test. Allen *et al.* [13] for instance outline a “comparative Moral Turing Test” (cMTT) in which a human interrogator is presented with pairs of examples of morally significant actions of a human and an ethical machine. The interrogator is asked to judge which of each pair of actions is less moral, and if the less moral actions are, more often, human decisions, then the machine is judged to pass the test. As Allen *et al.* point out, there are several problems with this test; one is that human behavior is often less than perfectly ethical, so the cMTT may be setting the moral threshold too low – they note that “decisions that result in harm to others are likely to be much less tolerated in a machine than in another human being.”

Anderson and Anderson [22] propose an alternative “Ethical Turing Test” (ETT) in which a panel of ethicists are presented with the machine's ethical decisions across a range of application domains. Each ethicist is asked whether they agree or disagree with those decisions – if a significant number are in agreement (i.e., the ethicist would have made the same choice in the same situation) – then the machine is judged to pass the test. In a run of this test with a panel of 5 ethicists, described in [22], the level of agreement ranges from 75% to 100%. Anderson and Anderson note that the most contested domain of “treatment reconsideration” – “should the health care

worker (or robot) accept the patient's decision regarding treatment as final, or attempt to change their mind?” – was also the most ethically sensitive for humans. A finding that supports our view that ethical decisions difficult for humans will be equally challenging for machines and their designers.

It is important to note that the idea of an ethical Turing Test remains controversial. Arnold and Scheutz [39] for instance argue against such a test. One of their objections is, essentially, that the ethical Turing Test treats the ethical agent as a black box by considering only its decisions and not the reasoning behind those decisions. Instead Arnold and Scheutz advocate verification, which seeks “predictable, transparent, and justifiable decision-making and action” – a position with which we strongly agree.

B. The Problem of Learned Ethical Principles

Consider the distinction between a machine that takes decisions according to predetermined principles of ethics and one that learns those principles from observed decisions. In fact, the latter is closer to the “scientific” study of (rather than the application of) ethics: the objective is to identify principles which serve to explain judgments of morality (distinctions between good and bad, or right and wrong) in terms of the reasoning used to justify them. This, rather than training, is the real motivation behind studying abstract situations such as the trolley problem [40]. These problems provide data (in the form of a set of moral judgments) from which a general principle of ethics can be derived. As with any good theory, this principle should have predictive leverage: so, if the problem or its parameters are changed, it should be possible to test whether the principle still holds (or not).

It is one thing to propose principles of ethics, implement them in a machine, and use them in a restricted context (the same context from which the principle was derived, perhaps). It is another thing to apply Machine

Learning algorithms to “learn” the principle(s) from the data; in which case, we need to be absolutely certain that the dataset has not been biased and that the explanatory principle so learned really does have predictive leverage when applied to a different context. For sure, if the machine learns a “wrong” or “inadequate” principle, or even just a “simple” principle, then there will be problems if we try to apply it in other situations.

Intuitively, the situation appears to be analogous to the problem of distributive justice, where in any given situation there are a number of possible ways of distributing rewards or punishments. These are called legitimate claims [41]. Equally, in any situation requiring a moral judgment, there may well be a number of principles that have stronger or weaker relative relevance. The requirement is to work out which principles apply in any situation, how to accommodate them in case of plurality, and how to reconcile them in case of conflict [42].

VI. THE MORALITY OF BUILDING ETHICAL MACHINES

Consider the question: is it ethical to delegate moral responsibility to machines? We routinely delegate task-level responsibilities to our machines; from simple automata such as washing machines to advanced safety-critical systems such as airplane autopilots, we trust a wide range of machines to undertake tasks that used to be exclusively performed by humans. To extend this kind of delegated responsibility to encompass ethics may therefore appear uncontroversial, but – we contend – it is a step that should only be taken with great care.

In considering the morality of building ethical machines we need to differentiate between implicit ethical machines – those designed to avoid unethical outcomes – and explicit ethical machines – those that reason about ethics – because the calculus of risk is quite different in each case.

Consider first the category of implicit ethical machines. We have already argued that *all* intelligent autonomous systems that have the potential to cause harm should be classed as implicit ethical machines, and designed using processes of ethically aligned design. The risks of not doing so are already apparent as outlined in Section IV. Building implicit ethical machines would seem to be an ethical course of action.

Conversely, explicit ethical machines do bring risk. In this special issue Cave *et al.* identify four broad categories of risk: “a) the risk that ethically aligned machines could fail, or be turned into unethical ones; b) the risk that ethically aligned machines might marginalize alternative value systems; c) the risk of creating artificial moral patients; and d) the risk that our use of moral machines will diminish our own human moral agency” [51]. It is by no means clear that we should build ethical machines, even if we can.

Of course real-world explicit ethical machines are not inevitable. There are, to the best of our knowledge, no explicit ethical agents in real-world use. As we have outlined here the only explicit ethical agents that exist are a handful of proof-of-concept laboratory prototypes. These are minimally ethical machines with limited functionality designed only to test hypotheses about how to build such machines.

A. The Ubiquitous Trolley Problem

In the public discourse around the ethics of AI, there is a common assumption that ethical machines – and in particular driverless car autopilots – should be able to resolve ethical dilemmas and choose between two equally undesirable outcomes. This assumption is fuelled by the Trolley Problem [41], of which there are many variations. For driverless cars the problem is often posed as the car having to choose whether to kill one human or several. A recent high profile study asked respondents for their preferences in a range of driverless

car scenarios and, perhaps not surprisingly, the study revealed significant cultural variations [43].

From an engineering perspective building a machine that can both reliably perceive and then choose between two unethical outcomes in any real-world environment – let alone in fast moving dynamic situations – is far beyond the state of the art. Yet in the public discourse this is rarely made clear, giving rise to unrealistic expectations of the decision-making capabilities of near future driverless cars. The trolley problem is a thought experiment; it undoubtedly has value both as a philosophical tool (as outlined in Section V) and at a statistical level – as suggested by Bonnefin *et al.* in this special issue [49] – but should not influence designs for driverless cars or any other autonomous systems.

Even if the technical problem of machines able to resolve real-world ethical dilemmas were solved society-wide debate would then be needed to discuss and agree on the rules and protocols for such machines, not least because society as a whole needs to take responsibility for the human causalities of accidents caused by such machines. It is notable that the federal government of Germany has ruled that “in the event of unavoidable accident situations, any distinction based on personal features (age, gender, physical or mental constitution) is strictly prohibited. It is also prohibited to offset victims against one another” [44].

VII. HOW SHOULD WE GOVERN ETHICAL MACHINES?

To recap: *all* intelligent autonomous systems should be regarded as ethical impact agents. Of course some may have no ethical impact whatsoever, but it would be foolhardy to assume this. Ideally all systems should be first, subjected to an ethical risk assessment, of the kind set out in standard BS 8611 [5], and second, redesigned to reduce the impact of any ethical risks exposed

by that risk assessment. Those systems that are shown to have some ethical impact (and we would wager that this would be almost all systems), which are then redesigned to avoid or minimize that impact, move into the category of implicit ethical machines.

It follows that we regard ethical risk assessment as the cornerstone of ethical governance. However, ethical risk assessment on its own is not enough, especially for implicit ethical machines. One of the general ethical principles set out in *Ethically Aligned Design* [6] concerns transparency; the principle asserts that it should always be possible to find out how and why an autonomous system made a particular decision. This kind of transparency is not a property of autonomous systems by default. It needs to be designed in, alongside sub-systems for securely logging system inputs, outputs and decisions – the robot/AI equivalent of an aircraft flight data recorder [45]. Without transparency discovering the causes of, for instance, a driverless car accident, or misdiagnosis by a medical diagnosis AI, becomes all but impossible. That process of discovery is vital if the faults that caused the accident are to be fixed, and accountability established [46].

Consider now the governance of explicit ethical machines. We already have very high expectations for the safety and reliability of our machines – especially those that have the potential to cause serious harm if they go wrong – but if and when real-world robots and AIs are explicitly ethical those expectations will be even higher. Because of this additional burden of expectation their governance will need to be especially robust.

Explicit ethical machines require two governance considerations over and above those needed for implicit ethical machines. The first is in the choice of ethical rules or – if those rules are learned – the choice of training cases. Those ethical rules or training cases need to be carefully scrutinized, ideally by a panel of

users, ethicists, lawyers and representatives of civil society. The second concerns the transparency of the ethical decision making process – it is critical that each ethical decision should be logged for later analysis and that it should be possible to understand why the system made those decisions. This is not only in the event that a system fails or causes harm. We believe that explicit ethical machines should be subject to a “probationary period,” during which *all* ethical decisions are reviewed (perhaps by the same panel that scrutinizes the system’s ethical rules). Only after this period is successfully concluded would monitoring revert to the baseline of fully investigating actual or near miss accidents. Given that explicit ethical machines are neither inevitable nor imminent, these considerations do not have the same urgency as those suggested for implicit ethical machines.

It is clear that *all* ethical machines need to be developed and used responsibly [37], and – for those that are ethically-critical – subject to standards and strong regulatory frameworks (many of which do not yet exist). In particular transparency should extend beyond the ethical machines themselves, to encompass the processes of design and operation, within frameworks of ethical governance [7]. Without such frameworks it is hard to see how ethical machines will be trusted.

VIII. THE PAPERS IN THIS SPECIAL ISSUE

Adamson *et al.* contribute a paper titled: “Designing a values-driven future for ethical autonomous and intelligent systems” [47], in which they argue that human values must drive our future autonomous systems in a way that both protects and benefits humanity. The paper describes IEEE’s work in this area and includes a table of approximately 50

activities within the IEEE related to ethics.

Anderson *et al.* contribute a paper titled “A value-driven eldercare robot: Virtual and physical instantiations of a case-supported principle-based behavior paradigm” [48]. The paper describes both simulated and real-robot implementations of an eldercare robot in which ethical principles are learned, via inductive logical programming, from a set of training examples provided by a project ethicist using GenEth: a general ethical dilemma analyzer.

Bonnefon *et al.* provide an insightful Point of View article titled: “The trolley, the bull bar, and why engineers should care about the ethics of autonomous cars” [49]. This paper, which brings attention to what the authors call “the statistical trolley dilemma,” is the only paper of the special issue focussed on autonomous vehicles

Bremner *et al.* contribute a paper titled: “On proactive, transparent and verifiable ethical reasoning for robots” [50], in which they review and update an approach to the design of ethical robots based on a simulation-based internal model. This model allows the robot to anticipate when another robot – acting as a proxy human – might be at risk of harm and intervene if necessary; the ethical robot’s reasoning is both transparent and verifiable.

Cave *et al.* present a paper aptly titled: “Motivations and risks of machine ethics” [51]. In this paper the authors clarify various philosophical issues surrounding the concept of an ethical machine and the aims of machine ethics. The authors argue that while there are good *prima facie* reasons for pursuing machine ethics, there are also potential risks that must be considered and managed.

Ema *et al.*, an 11 person research team based in Japan, present the

paper “Clarifying privacy, property, and power: Case study on value conflict between communities” [52]. Based around a controversial case study on the “flaming” of fan fiction, which was complicated by the ambiguous legal position of such fan fiction content in Japan, The paper aims to clarify notions of privacy and draw lessons for the ethical governance of (ethical) AI in the presence of value conflicts, through interdisciplinary collaboration.

Robertson *et al.* contribute a paper titled: “Engineering based ethical design methodology for embedding ethics in autonomous robots” [53]. The paper explores the process of robotics and autonomous systems development using a co-design approach to reduce end-user risk with respect to an endoscopic capsule for diagnosis and drug delivery. The contribution of the paper is a method for embedding ethics into the design of a machine in a socio-technical application.

“Inside the organization: Why privacy and security engineering is a challenge for engineers” is a paper contributed by Spiekermann *et al.* [54]. The study surveys 124 engineers in order to understand the issues facing ethical systems development with respect to privacy and security engineering. Their findings indicate that while many engineers regard security and privacy as important, they do not enjoy working on them, and struggle with their organizational environment.

In their Point of View “Toward the agile and comprehensive international governance of AI and robotics” [55], Wallach and Marchant propose an agile ethical/legal model for the international and national governance of AI and robotics, building on their recommendations on the formation of Governance Coordinating Committees (GCCs) for coordinated oversight of emerging technologies. ■

REFERENCES

- [1] K. Schwab, *Fourth Industrial Revolution*. New York, NY, USA: Portfolio Penguin, 2017.
- [2] T. Dutton. (2018). *An Overview of National AI Strategies*. Medium. [Online]. Available: <https://medium.com/politics-ai/an-overview-of-national-ai-strategies-2a70ec6edfd>
- [3] M. Boden *et al.*, “Principles of robotics: Regulating robots in the real world,” *Connection Sci.*, vol. 29, no. 2, pp. 124–129, Apr. 2017.
- [4] P. Boddington, *Towards a Code of Ethics for Artificial Intelligence*. Cham, Switzerland: Springer, 2017.

- [5] *Robots and Robotic Devices: Guide to the Ethical Design and Application of Robots and Robotic Systems*. British Standard BS8611, 2016.
- [6] "Ethically aligned design: A vision for prioritizing human well-being with autonomous and intelligent systems (A/IS)," IEEE Standards P7000, version 2, IEEE Standards Association, 2017.
- [7] A. F. Winfield and M. Jirotko, "Ethical governance is essential to building trust in robotics and artificial intelligence systems," *Phil. Trans. Royal Soc. A*, vol. 376, Oct. 2018, Art. no. 20180085.
- [8] I. Asimov, *I, Robot*. New York, NY, USA: Gnome Press, 1950.
- [9] W. Wallach and C. Allen, *Moral Machines: Teaching Robots Right from Wrong*. New York, NY, USA: Oxford Univ. Press, 2009.
- [10] M. Frayn, *The Tin Men*. Collins, 1965.
- [11] S. Cave and K. Dihal, "Hopes and fears for intelligent machines in fiction and reality," *Nature Mach. Intell.*, vol. 1, pp. 74–78, Feb. 2018.
- [12] T. Forester and P. Morrison, *Computer Ethics: Cautionary Tales and Ethical Dilemmas in Computing*. Cambridge, MA, USA: MIT Press, 2018.
- [13] C. Allen, G. Varner, J. and Zinser, "Prolegomena to any future artificial moral agent," *J. Exp. Theor. Artif. Intell.*, vol. 12, no. 3, pp. 251–261, 2000.
- [14] C. Allen, I. Smit, and W. Wallach, "Artificial morality: Top-down, bottom-up, and hybrid approaches," *Ethics Inf. Technol.*, vol. 7, pp. 149–155, Sep. 2005.
- [15] P. M. Asaro, "What should we want from a robot ethic?" *Int. Rev. Inf. Ethics*, vol. 6, no. 12, pp. 9–16, 2006.
- [16] J. H. Moor, "The nature, importance, and difficulty of machine ethics," *IEEE Intell. Syst.*, vol. 21, no. 4, pp. 18–21, Jul./Aug. 2006.
- [17] T. M. Powers, "Prospects for a Kantian machine," *IEEE Intell. Syst.*, vol. 21, no. 4, pp. 46–51, Jul./Aug. 2006. doi: [10.1109/MIS.2006.77](https://doi.org/10.1109/MIS.2006.77).
- [18] M. Anderson and S. L. Anderson, "Special issue on machine ethics," *IEEE Intell. Syst.*, vol. 21, no. 4, Jul./Aug. 2006.
- [19] M. Anderson and S. L. Anderson, Eds., *Machine Ethics*. Cambridge, U.K.: Cambridge Univ. Press, 2011.
- [20] R. Arkin, *Governing Lethal Behavior in Autonomous Systems*. Evanston, IL, USA: Routledge, 2009.
- [21] G. Adamson, R. R. Kline, K. Michael, and M. G. Michael, "Wiener's cybernetics legacy and the growing need for the interdisciplinary approach [Scanning Our Past]," *Proc. IEEE*, vol. 103, no. 11, pp. 2208–2214, Nov. 2015.
- [22] M. Anderson and S. L. Anderson, "GenEth: A general ethical dilemma analyzer," in *Proc. 28th AAAI Conf. Artif. Intell.*, 2014, pp. 253–261.
- [23] A. F. T. Winfield, C. Blum, and W. Liu, "Towards an ethical robot: Internal models, consequences and ethical action selection," in *Lecture Notes in Computer Science*, vol. 8717, Berlin, Germany: Springer, 2014, pp. 85–96.
- [24] D. Vanderelst and A. Winfield, "An architecture for ethical robots inspired by the simulation theory of cognition," *Cogn. Syst. Res.*, vol. 48, pp. 56–66, May 2018.
- [25] S. Bringsjord, N. Sundar, D. Thero, and M. Si, "Akratic robots and the computational logic thereof," in *Proc. IEEE Int. Symp. Ethics Eng., Sci., Technol.*, Piscataway, NJ, USA, 2014, pp. 7:1–7:8.
- [26] G. M. Briggs and M. Scheutz, "Sorry, I can't do that: Developing mechanisms to appropriately reject directives in human-robot interactions," in *Proc AAAI Fall Symp. Ser.*, 2015, pp. 1–5.
- [27] J. Shim, R. Arkin, and M. Pettinatti, "An Intervening Ethical Governor for a robot mediator in patient-caregiver relationship: Implementation and Evaluation," in *Proc. IEEE Int. Conf. Robot. Automat. (ICRA)*, Singapore, 2017, pp. 2936–2942.
- [28] S. Bringsjord, K. Arkoudas, and P. Bello, "Toward a general logicist methodology for engineering ethically correct robots," *IEEE Intell. Syst.*, vol. 21, no. 4, pp. 38–44, Jul./Aug. 2006.
- [29] S. Dyrkolbotn, T. Pedersen, and M. Slavkovic, "On the distinction between implicit and explicit ethical agency," in *Proc AAAI/ACM Conf. AI, Ethics, Soc.*, 2018, pp. 74–80.
- [30] A. Sharkey, "Can we program or train robots to be good?" in *Ethics and Information Technology*, 2017. doi: [10.1007/s10676-017-9425-5](https://doi.org/10.1007/s10676-017-9425-5).
- [31] D. Helbing, "Will democracy survive big data and artificial intelligence?" in *Towards Digital Enlightenment*, D. Helbing, Ed. Springer, 2019.
- [32] A. Caliskan, J. J. Bryson, and A. Narayanan, "Semantics derived automatically from language corpora contain human-like biases," *Science* vol. 356, no. 1334, pp. 183–186, 2017.
- [33] G. Neff and P. Nagy, "Automation, algorithms, and politics | talking to bots: Symbiotic agency and the case of tay," *Int. J. Commun.*, vol. 10, p. 17, Oct. 2016.
- [34] W. J. Wolf, K. W. Miller, and F. S. Grodzinsky, "Why we should have seen that coming: Comments on Microsoft's 'tay' experiment," and wider implications," *ACM SIGCAS Comput. Soc.*, vol. 47, no. 3, pp. 54–64, 2017. doi: [10.1145/3144592.3144598](https://doi.org/10.1145/3144592.3144598).
- [35] V. Charisi, A. Habibovic, J. Andersson, J. Li, and V. Evers, "Children's views on identification and intention communication of self-driving vehicles," in *Proc. ACM Conf. Interact. Design Children*, 2017, pp. 399–404.
- [36] M. Brundage et al. (2018). "The malicious use of artificial intelligence: Forecasting, prevention, and mitigation." [Online]. Available: <https://arxiv.org/abs/1802.07228>
- [37] M. Jirotko, B. Grimpe, B. Stahl, G. Eden, and M. Hartswood, "Responsible research and innovation in the digital age," *Commun. ACM*, vol. 60, no. 5, pp. 62–68, 2017.
- [38] R. Reilly-Cooper, personal communication, 2015.
- [39] T. Arnold and M. Scheutz, "Against the moral Turing test: Accountable design and the moral reasoning of autonomous systems," *Ethics Inf. Technol.*, vol. 18, no. 2, p. 103, 2016. doi: [10.1007/s10676-016-9389-x](https://doi.org/10.1007/s10676-016-9389-x).
- [40] J. J. Thomson, "The trolley problem," *Yale Law J.*, vol. 94, p. 1395, 1985.
- [41] N. Rescher, *Distributive Justice*. Indianapolis, IN, USA: Bobbs-Merrill, 1966.
- [42] J. Pitt, D. Busquets, and S. Macbeth, "Distributive justice for self-organised common-pool resource management," *ACM Trans. Auton. Adapt. Syst.*, vol. 9, no. 3, pp. 14:1–14:39, 2014.
- [43] E. Awad et al., "The moral machine experiment," *Nature*, vol. 563, pp. 59–64, Oct. 2018.
- [44] *Ethics Commission: Automated and Connected Driving*, German Federal Ministry of Transport, BMVI, Jun. 2017.
- [45] A. F. Winfield and M. Jirotko "The case for an ethical black box," in *Towards Autonomous Robotic Systems (Lecture Notes in Computer Science)*, vol. 10454. Springer, 2017, pp. 262–273.
- [46] S. Wachter, B. Mittelstadt, and L. Floridi, "Transparent, explainable, and accountable AI for robotics," *Sci. Robot.*, vol. 2, no. 6, 2017, Art. no. eaa6080.
- [47] G. Adamson, J. C. Havens, and R. Chatila, "Designing a value-driven future for ethical autonomous and intelligent systems," *Proc. IEEE*, vol. 107, no. 3, Mar. 2019. doi: [10.1109/JPROC.2018.2884923](https://doi.org/10.1109/JPROC.2018.2884923).
- [48] M. Anderson, S. L. Anderson, and V. Berenz, "A value-driven eldercare robot: Virtual and physical instantiations of a case-supported principle-based behavior paradigm," *Proc. IEEE*, vol. 107, no. 3, Mar. 2019. doi: [10.1109/JPROC.2018.2840045](https://doi.org/10.1109/JPROC.2018.2840045).
- [49] J.-F. Bonnefon, A. Shariff, and I. Rahwan, "The trolley, the bull bar, and why engineers should care about the ethics of autonomous cars," *Proc. IEEE*, vol. 107, no. 3, Mar. 2019. doi: [10.1109/JPROC.2019.2897447](https://doi.org/10.1109/JPROC.2019.2897447).
- [50] P. Bremner, L. A. Dennis, M. Fisher, and A. F. Winfield, "On proactive, transparent and verifiable ethical reasoning for robots," *Proc. IEEE*, vol. 107, no. 3, Mar. 2019. doi: [10.1109/JPROC.2019.2898267](https://doi.org/10.1109/JPROC.2019.2898267).
- [51] S. Cave, R. Nyrupe, K. Vold, and A. Weller, "Motivations and risks of machine ethics," *Proc. IEEE*, vol. 107, no. 3, Mar. 2019. doi: [10.1109/JPROC.2018.2865996](https://doi.org/10.1109/JPROC.2018.2865996).
- [52] A. Ema, "Clarifying privacy, property, and power: Case study on value conflict between communities," *Proc. IEEE*, vol. 107, no. 3, Mar. 2019. doi: [10.1109/JPROC.2018.2837045](https://doi.org/10.1109/JPROC.2018.2837045).
- [53] L. J. Robertson, R. Abbas, G. Alici, A. Munoz, and K. Michael, "Engineering-based design methodology for embedding ethics in autonomous robots," *Proc. IEEE*, vol. 107, no. 3, Mar. 2019. doi: [10.1109/JPROC.2018.2889678](https://doi.org/10.1109/JPROC.2018.2889678).
- [54] S. Spiekermann, J. Korunovska, and M. Langheinrich, "Inside the organization: Why privacy and security engineering is a challenge for engineers," *Proc. IEEE*, vol. 107, no. 3, Mar. 2019. doi: [10.1109/JPROC.2018.2866769](https://doi.org/10.1109/JPROC.2018.2866769).
- [55] W. Wallach and G. Marchant, "Toward the agile and comprehensive international governance of AI and robotics," *Proc. IEEE*, vol. 107, no. 3, Mar. 2019. doi: [10.1109/JPROC.2019.2899422](https://doi.org/10.1109/JPROC.2019.2899422).

ABOUT THE AUTHORS

Alan F. Winfield (Member, IEEE) received the Ph.D. degree in digital communications from the University of Hull, Kingston upon Hull, U.K., in 1984. Dr. Winfield is Professor of Robot Ethics at the University of the West of England (UWE), Bristol, U.K., Visiting Professor with the University of York, York, U.K., and Associate Fellow of the Leverhulme Centre for the Future of Intelligence, Cambridge, U.K. In 1984 he co-founded and led APD Communications Ltd, until taking up appointment at UWE in 1992. He co-founded the Bristol Robotics Laboratory, UWE, where his research is focused on cognitive robotics.



Dr. Winfield is passionate about communicating research and ideas in science, engineering and technology; he led UK-wide public engagement project Walking with Robots, awarded the 2010 Royal Academy of Engineering Rooke medal for public promotion of engineering. Until recently he was director of UWE's Science Communication Unit. He has published over 240 works, including *Robotics: A Very Short Introduction* (Oxford University Press, 2012).

Dr. Winfield is an advocate for robot ethics; he is a member of the British Standards Institute working group that drafted BS 8611: *Guide to the Ethical Design of Robots and Robotic Systems*. He sits on the executive of the IEEE Standards Association Global Initiative on Ethics of Autonomous and Intelligent Systems, and chairs IEEE Working Group P7001, drafting a new standard on Transparency of Autonomous Systems.

Katina Michael (Senior Member, IEEE) received the B.S. degree in information technology from the School of Mathematical and Computing Science, University of Technology, Sydney, NSW, Australia, in 1996, the Doctor of Philosophy degree in information and communication technology from the Faculty of Informatics, University of Wollongong, Wollongong, NSW, Australia, in 2003, and the Master of Transnational Crime Prevention degree (distinction) from the Faculty of Law, University of Wollongong in 2009.



Dr. Michael has held visiting academic appointments at Nanjing University, Nanjing, China, and the University of Southampton, Southampton, U.K. From 1996 to 2001, she was a Senior Network Engineer with Nortel Networks, Wollongong, NSW, Australia. She was also a Systems Analyst with Andersen Consulting, North Sydney, NSW, Australia, and OTIS Elevator Company, Minto, NSW, Australia. She is currently a Professor with the School for the Future of Innovation in Society and School of Computing, Informatics and Decision Systems Engineering, Arizona State University, Tempe, AZ, USA, where she is also the Director of the Center for Engineering, Policy and Society. She has authored or co-authored over seven edited books, 500 page reference volume: *Innovative Automatic Identification and Location-Based Services: From Bar Codes to Chip Implants* (Hershey, PA, USA: IGI, 2009), and 150 peer-reviewed papers. Her current research interests include emerging technologies and technologies used for national security and their corresponding social implications.

Ms. Michael was a recipient of the Brian M. O'Connell Distinguished Service Award in the Society for the Social Implications of Technology. She is the Guest Editor of 15 special issues including in the PROCEEDINGS OF THE IEEE, *Computer*, *IEEE Robotics and Automation Magazine*, *IEEE Potentials*, *Journal of Location Based Services*, *Computer Communications*, *Electronic Commerce Research*, and *Prometheus*. She was the Editor-in-Chief of the *IEEE Technology and Society Magazine* from 2012 to 2017 and has been a Senior Editor of the *IEEE Consumer Electronics Magazine* since 2015. She is the founding Editor-in-Chief of the new IEEE Transactions on Technology and Society that will begin publication in 2020.

Jeremy Pitt received the B.Sc. in Computer Science from Manchester University in 1984 and was awarded the Ph.D. in Computing from Imperial College London in 1991.



Dr. Pitt is Professor of Intelligent and Self-Organising Systems in the Department of Electrical & Electronic Engineering at Imperial College London. His inter-disciplinary research focuses on the use of mathematical and computational logic to provide the foundations for the design and development of self-organising socio-technical systems, with applications in sustainability, community energy systems and transportation. In all of these applications, there are issues involving collective action, polycentric governance, knowledge management, non-monetary economies, and computational justice; and his research programme aims to incorporate solutions to these issues in intelligent systems which use information for socially productive purposes that complements human ingenuity, creativity and innovation. His most recent work in the intersection of computer science and political science has generated the idea of "Democracy by Design," establishing design principles for socio-technical systems which ensure that democratic processes, practices and principles are considered as a primary requirement in systems development. He is a Fellow of IET, a Fellow of the BCS, and Editor-in-Chief of *IEEE Technology and Society Magazine*, reflecting his concern for the impact of technology on society and the social, legal, ethical and environmental implications of new technology.

Vanessa Evers received a PhD degree in Computer Science/Educational Technology from the Open University, Milton Keynes, U.K. in 2001.



Dr. Evers is full Professor of Human Media Interaction at the University of Twente, vice-dean of research and scientific director of the DesignLab. Her research focuses on the design and development of Socially Intelligent Agents. This concerns human interaction with autonomous agents such as robots or machine learning systems and cultural aspects of Human Computer Interaction. She is best known for her work on social robotics such as the FROG robot (fun robotic outdoor guide), SPENCER (The airport service robot) and DE-ENIGMA (robot for autism education) that can interpret human behavior automatically and respond to people in a socially acceptable way. She is very active organizing scientific conferences and as editor of academic journals, she is a speaker on AI and Robotics at international events such as the World Economic Forum and is a frequent contributor to the media in newspapers and tv-shows.