Introduction

A. Michael Froomkin

Like the Internet before it, robotics is a socially and economically transformative technology. The chapters that follow explore how the increasing sophistication of robots and their widespread deployment everywhere from the home to hospitals, public spaces, and the battlefield requires rethinking a wide variety of philosophical and public policy issues, interacts uneasily with existing legal regimes, and thus may counsel changes in policy and in law.

These discussions are necessary and timely. The state of thinking about legal and policy issues relating to robotics today is analogous to how scholars and policymakers approached the Internet before the World Wide Web. Imagine if in 1980 (more than ten years after the first Internet standards document, RFC 1), just as the Internet was starting to grow, or even in 1992, just as the first website went online, an interdisciplinary group of scholars attempted to engage with the policy and legal issues raised by the Internet. Undoubtedly the participants would have failed to foresee all the consequences of the Internet, good and bad, that we enjoy today. It is probable, however, that they would have identified key difficulties relating to domain names and trademarks, information security, access for the disabled, and privacy, to name only a few. By the time these issues were, in fact, recognized as significant, the installed base was sufficiently large to make changes to the relevant protocols and practices controversial (e.g., IPSEC) and in some cases (e.g., domain names) highly impractical.

As the chapters that follow demonstrate, although we are still at an early stage in thinking about the social and legal effects of robotics we already can identify a host of issues raised by the widespread deployment of robots in society. Robots raise issues spanning a very wide disciplinary focus and likely to impact not just all walks of daily and commercial life but also war. On the other hand, that very process of issue identification reveals just how much definitional and applied work remains to be done. Addressing these questions in an organized manner is a substantial undertaking, one that requires an extended conversation that must cross disciplines if it is to be informed and useful.

A measure of how early a stage we are in is that there is not yet a consensus regarding what should count as a “robot.” Most people, and undoubtedly all the contributors to this volume, would agree that a man-made object capable of responding to external stimuli and acting on the world without requiring direct – some might say constant – human control was a robot, although some might well argue for a considerably broader definition. The three key elements of this relatively narrow, likely under-inclusive, working definition are: (1) some sort of sensor or input mechanism, without which there can be no stimulus to react to; (2) some controlling algorithm or other system that will govern the responses to the sensed data; and (3) some ability to respond in a way that affects or at least is noticeable by the world outside the robot itself. (An even narrower version of that vision produces my favorite definition, prized for its pungency more than its accuracy, which says a robot is “an iPhone with a chainsaw attached.”)

This working definition permits a great deal of variation. Robots could be any size, ranging from nanobots to behemoths, but with the exception of the special issues potentially raised by nanomachines, size is not what usually makes a robot interesting and potentially problematic.

There is even greater variation – with more significance – relating to physical instantiation and independence. Robots can be corporeal, or they might be software-only creations so long as they take input from the outside world and they are able to respond in a way that has effects on the world beyond the machine that hosts the program; software is, of course, controlled by algorithms. Robots can have a very narrow range of function or have many; they can have a very narrowly circumscribed range of freedom of movement and of decision-making, or it can be very broad, potentially even limitless. At one extreme are the automata one finds on “robotic assembly lines.” In some cases these machines simply do the same thing over and over and do not count as robots. But in other cases even machines capable of only a narrow range of motion or discernment may technically fit the working definition set out above: a robot arm, for example, may be designed to sense the varying locations of a particular part as it flows down a conveyor belt, and may be designed to grab it and attach it in a set manner. That may be enough to qualify it as a robot in our terms, albeit not a very interesting one.

At the other extreme both in physicality and in inscrutability is the robot that features in the Millar and Kerr chapter, a medical diagnosis robot that is a program running on a supercomputer. It may collect information from sensors attached to the patient, or it could have no sensors at all and just respond to data input by diagnosticians. In either case, its responses are just words and numbers – it doesn’t do anything other than

advise, so its direct effect on the world is limited, although extant. But its independence is so great, and its method of finding correlations in giant data sets so opaque, that we may never be able to understand why it makes the choices it does. Instead we can only measure its success rate, compare it with human doctors, and begin to think seriously about what it means if – or rather when – the robot doctor starts compiling a superior batting average.

Nevertheless, many things that people commonly call “robots” are outside the working definition set out above, and also outside the scope of this project. Excluded from this working definition are many of the automata that figure heavily in popular consciousness. Machines fully controlled by remote human beings operating by telepresence – waldos – including, as far as we know, all military drones currently deployed in the field (as distinct from those on the drawing board and perhaps those being tested), are not robots in this sense because they lack any independent ability to respond to stimuli. They are no more robots than is a conventional motor vehicle, one that does not have a robot driver. They are not at the very least the sort of robots that generate the interesting problems we are concerned with here. Conversely, so-called cyborgs, that is, persons who have had their bodies enhanced by the use or implantation of some technology, are not in themselves “robots” because they are still people, for the same kinds of reasons that driving a conventional motor vehicle fails to make the driver a robot.

That said, some of those enhancements, notably robotic prosthetic limbs, raise issues much like their more independent cousins even if the devices might not be considered robots in the narrowest sense. Consider, for example, a robot leg engineered to respond to – and perhaps even anticipate – commands sent along human nerves connected to the machinery. Ideally, from the user’s point of view, the conscious intentionality needed to walk would be no different for the mechanical limb than for the biological one. One need only imagine, however, what happens if the prosthesis is not perfectly tuned: the limb’s owner is trapped at a cocktail party, listening to a particularly boring and unpleasant speaker and imagines someone giving him a good swift kick. Unconscious impulses travel down the nerves, meet the robotic controller, and suddenly the party gets more interesting.

Joking aside, this image of this semi-robotic kicker underlines the most significant way in which the legal and social issues raised by (at least non-software) robots with any degree of autonomy differ from those raised by the Internet: because they exist in real space, robots can cause direct physical effects on the world. They can cause physical damage, hurt – even kill – people either by accident or, in the case of military robots, by design.

Robots have great potential to take on unpleasant, difficult, or boring tasks, but they also present real risks that require careful planning by designers and by policy-makers.

Work on the social, political, ethical, and legal effects of new technologies tends to suffer from unusual barriers to entry when the technology is in its infancy. Industrial scientists deploying cutting-edge technologies tend to be focused on making them work, on getting them to market, and on making them profitable. They tend to view the social and policy, not to mention legal, issues as obstacles to be ignored, if possible, and overcome otherwise. Academic scientists are more likely to consider the broader questions, but they do not always have training in the relevant disciplines, whether social science, philosophy, or law. Conversely, the academics with training in those areas frequently lack the necessary background in the relevant technology and are often unfamiliar with the relevant literatures; often they may find them impenetrable. These problems tend to be solved as a technology is deployed: as the technology enters the mainstream it attracts popularizers. More importantly, as deployment takes place the new issues raised by the new technology become manifest – but by then it can be too late to alter the technology to design around the problem. Similarly, the presence of potential winners and losers may make the implementation of either policy or technology fixes more difficult than if they had been put in place earlier.

This volume collects the efforts of a diverse group of scholars who each, in their way, has worked to overcome these early barriers to entry and thus sought to grapple with the larger consequences of the increasingly discernible future of robotics. Most of the contributions in this volume are updated versions of papers first presented at “We Robot,” an annual interdisciplinary conference that began in 2012 and has been held at the University of Miami (twice), Stanford University, and the University of Washington. Like the authors of the chapters that follow, the attendees at We Robot are united by a belief that we have a lot of work to do if we are to build robots that play well with all the environments they are colonizing. We face even larger tasks if we are to craft a legal environment that, on the one hand, correctly balances the needs to encourage experimentation and investment in technology with the goal of protecting public safety while, on the other hand, balancing the needs of competing economic players in the robotics industry – not to mention considering the interests of those involved in the industries that robots promise to enhance – or replace.

PART I: STARTING POINTS

Our starting point in this volume, appropriately, is a collaboration between a lawyer and a roboticist: Neil M. Richards and William D. Smart's "How Should the Law Think about Robots?" Faced with the coming challenge of larger and larger numbers of robots in public spaces, homes, offices, and hospitals, the legal system will need to decide how to treat them. Currently our laws tend to treat machines that do what they are designed to do as legal extensions of the people who set them into motion. When machines malfunction we try to assign the blame to the appropriate manufacturer, designer, maintainer, or modifier. Neither of these approaches transfers easily to the complex world of robots, especially robots with any degree of autonomy. Traditional legal approaches to fault and product liability also wobble when confronted with robots capable of emergent behavior. The complex interrelation among the manufacturer, the programmers, the owner, the user, and other parties who may have intentionally or otherwise trained the robot creates the potential for both evidentiary and theoretical thickets. Professors Richards and Smart argue that how we conceptualize robots – the metaphors we use to understand and relate to them – will shape how engineers design robots and how society, and especially lawyers, will regulate them. The authors argue that the issue is one of imagination – the imagination of designers and regulators – and how best to blend those two very different ways of thinking through problems. Thinking about robots in different ways – whether as tools, pets, servants, or family members – will affect how they are built and how the law recognizes them.

PART II: RESPONSIBILITY

Part II picks up the theme of responsibility for what robots do. It begins with two lawyers and their largely opposing perspectives on the ability of the tort system to deal effectively with the legal and social perturbations likely to be introduced by robots. On the one hand, F. Patrick Hubbard's "Allocating the Risk of Physical Injury from 'Sophisticated Robots'" offers a somewhat sanguine assessment of the common law's ability to cope with technological change. He begins by showing how current doctrines might apply to various robot use cases, such as robot cars, and suggests that, in the main, outcomes under current doctrine will not stray too far off the mark from where we would want it to. Furthermore, Prof. Hubbard reminds us that today's tort law is itself the product of a continual evolution, and he suggests that the common law system thus

contains the internal mechanisms necessary to make the relatively small changes that a roboticized world will require. He therefore casts a generally skeptical eye on most contemporary law reform proposals aimed at fostering innovation in sophisticated robots. We should, he concludes, be very wary of proposals to make big changes in a system that is already a careful balance of competing interests between, on the one hand, entrepreneurs and proponents of new devices and, on the other, ordinary people who are entitled to redress from injury.

Judge Curtis E.A. Karnow offers a contrasting vision in “The Application of Traditional Tort Theory to Embodied Machine Intelligence,” in part because he focuses on the hardest case – genuinely autonomous robots. In Judge Karnow’s view, traditional doctrines of negligence and strict liability will not prove adequate to address their challenges. Judge Karnow begins by proposing a working definition of what it means for a robot to be “autonomous” – by no means the least important feature of this chapter. Armed with this definition, and U.S. tort law’s definition of “foreseeability,” Judge Karnow finds inadequacies in a legal system that presumes linear causation, but will be confronted with ever more complex machines interacting with the larger environment. He believes that the robots’ actions will be nonlinear, effectively unpredictable, and that a tort system designed to place most liability on those thought to deserve it because they had a reasonable opportunity to control or foresee the risk will not be up to the task. The tort system will need to change, our expectations will need to change, and – in a conclusion that evokes a theme introduced in the first chapter – robot design will need to find a way to build robots that have “common sense” so that their behavior does not veer too far from our expectations.

Bryant Walker Smith’s chapter, “Lawyers and Engineers Should Speak the Same Robot Language,” also addresses conceptual themes about the law’s encounter with robotics, this time with illustrations drawn from nascent attempts to create rules and standards for driverless cars. The conceptual portion argues that it is a mistake for either engineers or lawyers to think of a robot in isolation from the larger system of which it is a part. Many robots are or will be designed to interact with other machines, the cloud, and – critically – humans. Designs and rules must take this embeddedness into account. Yet the recognition that a robot is part of a system, in turn, creates a new problem, that of defining the relevant system boundaries. Having suitably complicated the question, Prof. Smith then takes us on a tour of actual industry standards and proposed or actual regulatory documents relating to autonomous motor vehicles and finds serious conceptual and terminological inconsistencies, especially when viewed through the lens of systems analysis. He concludes that the best

way to avoid this confusion, and any concomitant physical risk to motorists and third parties, is for engineers and lawyers to talk more so that they can build a common language about robots.

Our discussion of issues of responsibility concludes with Jason Millar and Ian Kerr's "Delegation, Relinquishment, and Responsibility: The Prospect of Expert Robots." Here the question is when it is right and reasonable to rely on a robotic expert. The authors begin by focusing on IBM's Watson (the "deep question answering" robot who defeated two all-time *Jeopardy* champions), a system now being tested as a robotic diagnostician. What happens if studies show that these robots' method of basing diagnoses on correlations drawn from Big Data perform better than a logic-based human faced with the same patient exhibiting the same symptoms? The robots may become the preferred experts, forcing us to question whether humans should have a duty to delegate tasks such as medical diagnosis and driving to expert robots. Exploring notions of expertise and evidence-based decision-making, a view that believes that an action that produces the most favorable outcome is the most justifiable, the authors canvas various legal and ethical approaches to determining responsibility when these advanced systems make mistakes and for when people fail to rely on robots as much as they should have.

PART III: SOCIAL AND ETHICAL MEANING

The discussion of whether it is ethically proper, or even obligatory, to rely on robots when they are safer than people transitions us to broader ethical questions involving robots. "The Open Roboethics Initiative and the Elevator-Riding Robot" by AJung Moon et al. describes an open-source-like initiative to accelerate the bottom-up design of roboethics principles suitable for implementation. The proposal responds to both the fragmentation of current roboethics initiatives and the tendency of existing approaches to be run by experts; that is, to be top-down in their orientation. A bottom-up approach allows for greater public participation, something the authors argue is particularly appropriate to the development of ethics-based rules that must by their nature be very context-dependent. Illustrating just how context-dependent these rules might be, the authors provide a proof-of-concept using the example of an "elevator-riding robot" that for safety and other reasons must be the sole passenger on an elevator in the building where it delivers things. How should it deal with elevators that arrive containing people? Should it ask them to step out? What about when the robot is waiting for an elevator with people who are also intending to use it? Whatever the answers, the authors suggest that

their Open Roboethics initiative (ORi) would produce a result that while not necessarily the most ethical in the sense that philosophers use the term, and in no way guaranteed to apply across cultures, likely would reflect the popular sentiment and expectation as to behavior in any given culture. This distillation of common understandings is what designers need in the short run and could, in the longer term, serve as a starting point for discussions of what are the morally correct actions. Indeed, in the case of the elevator-riding robot, the authors designed a survey that mapped expectations given variables such as the urgency of the robot's tasks, the person's location (in-elevator or waiting) and the person's characteristics (e.g., in a wheelchair, carrying heavy objects). The very small proof-of-concept survey ($n = 8$) allowed the construction of a decision matrix from which the robot could yield, insist, or dialogue with the other passenger, and the authors further discuss how such a matrix could be implemented in a robot – and how it could be designed to allow easy modification as needed, an outcome that the authors suggest demonstrates the utility of their Open Roboethics model.

A different sort of openness is the subject of “The Application of a ‘Sufficiently and Selectively Open License’ to Limit Liability and Ethical Concerns Associated with Open Robotics” by Diana Marina Cooper. Open licensing promotes the rapid innovation that the robotics industry needs if it is to achieve the goal of a “robot in every home,” but that same openness carries with it risks of legal liability in various scenarios. Open licensing also means that one's inventions might find ethically dubious uses downstream. The challenge, therefore, is to find a license that is neither too open to these dangers nor so closed as to hamper innovation – a “sufficiently open license.” This chapter examines the obstacles to the adoption of a “sufficiently open” model then discusses means of overcoming those obstacles. It concludes with a concrete proposal: a first draft of the Ethical Robot License (ERL).

The issue of ethically dubious uses of robots occupies center stage in Sinziana Gutiu's “The Robotization of Consent,” which considers the likely consequences of realistic sexbots on how humans interact with each other. Gutiu first argues that market-driven design choices will result in female robots that, even if they are not sexbots as such, reproduce unrealistic, stereotypical, female characteristics. By depicting female robots as passive, obedient, and sexualized, these devices will entrench false conceptions about what women should look and act like, thus harming real women by damaging their image and self-worth, and indeed harming men also by distorting their ability to form relations with women. Furthermore, the creation of increasingly realistic-seeming female sexbots will dehumanize sex and intimacy in male–female relationships. Most critically, unlike

real people, consent to sexual relations will be irrelevant for robots, further promoting misconceptions about women and eroticizing sexual slavery, which risks distorting men's understanding of how consent operates in intimate relations. While this robotic depiction of women's subordination raises ethical questions and fits Catherine MacKinnon's definition of pornography, it may not, Gutiu says, conform to the definition of pornography under either Canadian or U.S. law. We should not, she concludes, give in to the creation of sexbots – even though early models already exist, but should consider how to regulate sexbots while remaining mindful of the competing interests of freedom of expression.

The challenge of regulating misuses of social robots is taken up directly by Kate Darling in "Extending Legal Protection to Social Robots: The Effects of Anthropomorphism, Empathy, and Violent Behavior Towards Robotic Objects." She begins by noting that the practice of assigning rights to nonhuman entities is not new. We create penalties for harming certain animals (or, if you prefer, give those animals certain rights) because of our inherent desire to protect things we care about and with which we feel sympathy. Research shows that just as we project our emotions onto animals, so too we apply our inherent tendency to anthropomorphism – to attribute human qualities, such as cognition and emotions – onto robots, and perhaps especially "social robots," defined as robots that "communicate through social cues, display adaptive learning behavior, and mimic various emotional states." This emotional projection justifies regulating human interaction with social robots – be they toy animals or realistic human models – differently from interaction with a toaster, for many of the same reasons that we prohibit animal abuse. Even though robots are not alive, there are, Darling argues, other compelling psychological and philosophical parallels in how humans relate to the robotic victim of abuse that are relevant to animal abuse regulation. Among these are the risk that, as social robots become increasingly lifelike, we may subconsciously equate them with living things and then risk transposing behaviors applied from robots to the living. Similarly, if robots are lifelike, the same concern for the feelings of observers that justifies some prohibitions on animal abuse might apply to robots, even though the robot itself feels no pain. The argument also has a Kantian resonance, for just as our actions to nonhuman animals reflects our morality, so, too, may inhumanity against social, realistic robots lead to reinforcement of inhuman behavior more generally. Banning cruelty to robots – including but not limited to sexual abuse – may thus be both necessary and legitimate.

PART IV: LAW ENFORCEMENT

Legitimacy is a significant concern in Part IV, which addresses different aspects of robots in law enforcement. We begin with “Confronting Automated Law Enforcement” by Lisa Shay et al. The authors note that robots and associated automation systems such as sensor systems with law-enforcement algorithms and punishment feedback loops are an attractive tool for law enforcement because they seem to promise evenhanded, tireless, law enforcement while reducing manpower requirements. But of course it is not quite that simple: deployment of such systems may meet some of these goals, but they also risk unintended consequences. The five authors set out a taxonomy of automation capabilities, divided into categories of multiple types of surveillance, then analysis, aggregation, and punishment. They propose a detailed analytic framework for analyzing the costs and benefits of law-enforcement automation, including a lengthy checklist of factors – notably issues of freedom, autonomy, due process, and privacy – that need to be considered in order to fully account for the full potential costs as well as benefits. As the authors note, these calculations are essential before a program is implemented, as once adopted automation schemes tend to take on a bureaucratic life of their own.

The next chapter, “Do Robots Dream of Electric Laws? An Experiment in the Law as Algorithm,” by four of the same five authors, asks a much more focused but still complex question: How will automated law enforcement actually work in practice? Laws on the books are not, the authors report, well-crafted to be automated. To illustrate this problem they ran an experiment in which they asked 52 programmers independently to code the enforcement of a set of traffic laws. They then applied the algorithms provided to data drawn from the actual movements of a car fitted with a sensor recording an hour’s driving in traffic. Although all the programmers were tasked with applying the same simple speed limits, the results varied widely. Giving a subset of the programmers extra guidance, such as varying the tolerance for speeding based on tiers of speed, reduced the variation but did not eliminate it. The experiment raises many avenues for future study, almost all of which require much greater cooperation between engineers and lawyers. One could, for example, work to make laws more automation-friendly by specifying them differently and in more detail; or one could work with engineers, perhaps guided by lawyers, to figure out coding strategies that are sensitive to the way law is interpreted and used in reality. If nothing else, the authors observe, the results demonstrate how sensitive any attempt to automate the enforcement of legal rules will be to the details of the specifications and instructions given to the people who code the enforcement system.

Different, but no less weighty, concerns animate Kristen Thomasen's chapter, "Examining the Constitutionality of Robot-Enhanced Interrogation." Here the issue is the legitimacy of state agencies combining human-computer interaction (HCI) technology with sensors that monitor human physiological responses in order to improve the extraction of truthful information from suspects during interrogations. These technologies are already being deployed in automated kiosks, where an individual interacts with an avatar interrogator. The HCI system uses a combination of visual, auditory, near-infrared, and other sensors to monitor a suspect's eye movements, voice, and various other qualities. The information is then aggregated and analyzed to determine deception. Thomasen argues that while these technologies do offer genuine benefits to law enforcement, they also pose serious risks to individual rights – such as privacy and the right to silence. Her study explores possible solutions and suggests that courts, HCI technology developers, and state agencies institute limits on how this emerging technology obtains and uses information.

PART V: WAR

Our final section looks at intentionally lethal robots. Like so many of the chapters preceding them, both chapters in this section grapple with the ability of some part of the legal system – in this case international human rights law – to address the challenges created by the deployment of autonomous robots in a sphere of life, war and preparations for war.

In "Asleep at the Switch? How Killer Robots Become a Force Multiplier of Military Necessity" Ian Kerr and Katie Szilagyi consider the effects of the – inevitable? – introduction of robotic weaponry on International Humanitarian Law (IHL). Robots look to offer enormous advantages on the battlefield: the advanced sensory capabilities and shorter reaction times of machines will someday outperform human soldiers while reducing the deploying side's human casualties. Robotic warfare proponents believe that entrusting kill/no-kill decisions to autonomous robotic weapons will save lives by more consistently comporting with IHL. While pessimistic about the ability of robots to discriminate enough to comply with IHL, the authors accept that robots likely will be effective weapons – which they see as the source of a problem: IHL justifies the use of military force with a necessity/proportionality calculus, which weighs a military operation's necessity against the harm resulting from carrying out that operation. Kerr and Szilagyi argue that, contrary to the claims of the proponents of killer robot deployment, lethal autonomous robots threaten to erode the IHL framework because the robotization of warfare will push

us to redefine IHL norms just as did the invention of the submarine. The laws of war, they suggest, are expressed in technologically neutral terms: they are framed as general laws focusing on the uses and effects of force, even when applied to nuclear bombs; the proportionality principle presumes that outright prohibition of particular technologies is inappropriate or perhaps impossible. As a result, IHL always accepts the possibility of the use, or at least deployment, of new weapons because it only focuses on misuse. The authors argue that instead of accepting the deployment of killer robots as potentially “necessary” under IHL (because the robots are successful force multipliers) – and thus appear inevitable – we should instead consider technology-specific approaches transcending current IHL paradigms that will address the transformative effects of robotic military technologies.

Our final chapter suggests that this new law can be rooted in a venerable principle: the Martens Clause, which first appeared in the Preamble to the Hague Convention II on The Laws and Customs of War on Land in 1899, and which has been restated several times since. In “*Jus nascendi*, Robotic Weapons, and the Martens Clause,” Peter M. Asaro summarizes the Martens Clause as stating “that acts are not legal, or permissible, simply because they are not explicitly prohibited by the law.” This, he argues, provides a hook not just for traditional unwritten prohibitions but also for a recognition of evolving standards of morality. In other words, because it includes the Martens Clause, traditional IHL contains both the seeds for a prohibition on killer robots and a mechanism by which such a principle can be recognized by international tribunals even in the absence of a specific new treaty. If the goal of IHL is to protect an essential core of humanity, even among the great inhumanity of war, then one should consider the norms that will be challenged or violated by battlefield robotics and try to determine how best to protect those norms from being undermined. Prof. Asaro finds in current international discussions “something that looks very much like an emergent principle,” although he admits that at present “[i]t is difficult to call this a norm proper,” in part because the principle of human control has never needed to be stated before, as there were little other than booby traps and mines to challenge it. In the face of the introduction of killer robots, the principles of humanity and the dictates of public conscience argue for articulating and establishing new rules of IHL that address the threats these technologies pose to fundamental norms of responsibility and accountability, and to human rights and human dignity.

Looking at these essays as a whole, several common themes emerge. The three most notable are: (1) these are still very early days for issues of robots and regulation; (2) the technical issues are far more complex than

lawyers tend to imagine, and the legal, ethical, and philosophical issues more controversial (and also sometimes more complicated) than engineers tend to imagine; and (3) that the only way we are going to resolve these questions at all well is to widen and deepen our interdisciplinary efforts. Making the world ready for robots and robots ready for the world must be a team project – or else it likely will go poorly.

That these are the early days of Robot Law almost goes without saying. Yet, it needs to be said in order to emphasize that even if these are early days they are not in any way *too* early days. Indeed, as some of our authors note, path-dependence is a giant problem: once systems are in place they are hard to amend and even harder to dislodge. If we want to get the standards right, if we want to get the law right, we need to start now – if not yesterday. Here the example of the Internet, and especially of the domain name system, is highly instructive. Were we to design a domain name system on a blank slate, knowing what we now know about the controversies it has engendered, the DNS would look nothing like what it is today. We would design around the trademark issues, we would reduce or eliminate the single point of failure, we would build in a simpler mechanism to grow the top-level namespace. We have none of these options because of the legacy of design choices a group of young software engineers made decades ago – choices that were not technologically required but seemed inoffensive and convenient at the time and that, for understandable reasons (no one foresaw the growth of the Internet), seemed reasonable to make in isolation from any input by lawyers, ethicists, or political scientists. With robots, we know we can do better. And we also know that, like the early days of Internet policy-making, the choices we make will have long-term consequences.

The chapters in this volume demonstrate that there is a great need for interdisciplinary teamwork, but also that this teamwork can produce results. Interdisciplinary teamwork will allow robot designers to avoid some legal and ethical issues by designing around them. And when it is not possible or commercially feasible to design around these issues – or if someone chooses to ignore them – regulators need help from engineers and ethicists to help them craft rules that are well-targeted to grapple with the issues at hand but also likely to work for future iterations of the technology. The culture of Internet startups was not, in the main, friendly to lawyers except as facilitators of capital acquisition. Telling a software engineer that sometimes she ought to talk to lawyers before designing a product was about as fun as telling her she would need a root canal. Yet, sometimes fairly innocuous design changes can remove obstacles to robot deployment. Guidance as to the saliency of obstacles, and the identification and careful specification of issues, can guide development

of robot standards and design. Conversely the designers and makers of robots are necessary players if we are to achieve quality regulation, for they are best placed to explain what is coming, what is feasible, and what is commercially sensible.

Your editors hope that by identifying issues that threaten to shape or delay the deployment of robots, this very interdisciplinary set of chapters addressing the impacts of widespread robot deployment on society will influence plans and designs for future robots as well as the first iterations of Robot Law. Ideally, the conversations generated by these studies will help shape how the problems identified may be solved, mitigated, or avoided. We expect that the essays that follow will be of interest and of use to as wide a variety of groups as the authors, who include engineers, ethicists, lawyers, roboticists, philosophers, and serving military. We have many very interesting problems to solve together.