12-5-2023

Issues in robot law and policy

A. Michael Froomkin

25. Issues in robot law and policy

A. Michael Froomkin[1]

1. INTRODUCTION

Robots are, or soon will be, ubiquitous. Already robots can be found on the road, in the air, on the sea, in the home or office, in the warehouse, in the hospital, and on the battlefield. As robotics improve, we can only expect to see more robots, doing more complicated things. It follows that robots will intersect more with people, and with the law, at almost every level and subject. The physical instantiation of most robots, the fact that all but pure software robots have a body, means that robot legal issues and robot regulation has salience at every level of government from the most local to the most international. At least in the United States, most robot law either adopts existing law or consists of frequently unanswered questions: vast tracts of law are waiting to be decided and written.

It is common to speak of robots as mechanizing the "OODA" (Observe, Orient, Decide, Act) Loop, and for current purposes, I will treat as a "robot" any device that is capable of affecting the world outside itself in response to some sensed trigger. This is a very broad definition, as it encompasses everything from simple devices like a programmable room thermometer that changes heating or cooling in response to a temperature measurement, to automated hammers that stamp down every time a sensor detects a nail on the assembly line, to complex software-only program trading algorithms that buy or sell shares in response to market movements, to complex military defense systems designed to shoot down incoming missiles.

Many simple robots raise no unusual legal issues. A thermostat may be governed by product liability law; the assembly line hammer governed by workplace safety regulations and standard tort or workers' compensation rules, but they each have those rules in common with other devices in the home, or on the shop floor. While even simple robots may occasionally raise complex issues of fact, normally there are no issues of theory unique to these robots as opposed to other neighboring products and devices.

In contrast, as discussed further below, more complicated and interesting robots—not least those capable of emergent behavior—do raise difficult and often unsettled legal issues of responsibility and liability, of regulatory competence, subsidiarity, and jurisdiction, and a host of related ethical issues as to who should be responsible for robot harms, what indeed counts as a harm, and whether and when certain types of robots should be restricted or prohibited. Many of these complex questions overlap with, or form part of, "AI law" since complex and interesting robots commonly are controlled by what we call an artificial intelligence. That controlling AI may reside onboard the robot, or it may control the robot remotely. The internal/external location distinction can be hugely important in robotic design, e.g., for battlefield robots designed to operate in the face of enemy attempts to degrade communications, but

[1] Thanks to Caroline Bradley, Benjamin Froomkin, David Froomkin, and to the chapter's reviewers for helpful suggestions.

commonly the location of the controller is not legally significant—although the location may make a legal difference in cases where the external controller is in a different jurisdiction from the physical part, or under the control of different persons.

2. GENERAL LEGAL STATUS OF ROBOTS

2.1 Questions of "Personality" and "Robot Rights"

As the law currently stands in the United States and, as far as I know, nearly everywhere else, the law treats all robots of every type as chattel. (The sole possible exception is Saudi Arabia, which gave "citizenship" to a humanoid robot, Sophia, in 2017. It is hard to see this as anything more than a publicity stunt, both because female citizenship in Saudi Arabia comes with restrictions that do not seem to apply to Sophia, and because "her" "life" consists of … marketing for her Hong-Kong-based creators (Reynolds, 2018).) That is, in the words of Neil Richards and William Smart,

> Robots are, and for many years will remain, tools. They are sophisticated tools that use complex software, to be sure, but no different in essence than a hammer, a power drill, a word processor, a web browser, or the braking system in your car. (Richards & Smart, 2016)

It follows that robot personhood (or AI personhood) under law remains a remote prospect, and that some lesser form of increased legal protections for robots, beyond those normally accorded to chattels in order to protect their owners' rights, also remain quite unlikely. Indeed, barring some game-changing breakthrough in neural networks or some other unforeseen technology, there seems little prospect that in the next decades machines of any sort will achieve the sort of self-awareness and sentience that we commonly associate with a legitimate claim to the bundle of rights and respect we organize under the rubric of personhood, although the possibility of machine rights has motivated both interesting thought experiments (e.g., Boyle, 2011) and vivid denunciations (e.g., Birhane & van Dijk, 2022).

There are, however, two different scenarios in which society or policymakers might choose to bestow some sort of rights or protections on robots beyond those normally given to chattels. The first is that we discover some social utility in the legal fiction that a robot is a person. No one, after all, seriously believes that a corporation is an actual person, or indeed that a corporation is alive or sentient (although Stross, 2018, suggests we think of corporations as "Slow AIs"), yet we accept the legal fiction of corporate personhood because it serves interests, such as the ability to transact in its own name, and limitation of actual humans' liability, that society—or parts of it—find useful. Although nothing at present suggests similar social gains from the legal recognition of robotic personhood (indeed issues of liability and responsibility for robot harms need more clarity, not less accountability), conceivably policymakers might come to see things differently. In the meantime, it is likely that any need for, say, giving robots the power to transact can be achieved through ordinary uses of the corporate form, in which a firm might for example be the legal owner of a robot. This has not, however, stopped speculation about how a robot or AI might own itself (Bayern, 2021, 2019; LoPuki, 2018).

Early cases further suggest that US courts are not willing to assign a copyright or a patent to a robot or an AI even when it generated the work or design at issue. Here, however, the primary justification has been straightforward statutory construction, holdings that the relevant

US laws only allow intellectual property rights to be granted to persons, and that the legislature did not intend to include machines within that definition.[2] Rules around the world may differ. For example, in *Commissioner of Patents v. Thaler* (2022), an Australian federal court ordered an AI's patent to be recognized by IP Australia. Similarly, a Chinese court found that an AI-produced text was deserving of copyright protection under Chinese law (Sawers, 2020). The North American literature on these topics is already vast; good starting points are Ryan Abbott's arguments in favor of AI patents (Abbott, 2016) and Annemarie Bridy's early survey of the arguments relating to copyrights (Bridy, 2012); the latter recently triggered a particularly elegant response (Carig & Kerr, 2021). There is also a voluminous EU-based literature (e.g., Hugenholtz & Quintais, 2021).

A more plausible scenario for some sort of robot rights begins with the observation that human beings tend to anthropomorphize robots. As Kate Darling observes, "Our well-documented inclination to anthropomorphically relate to animals translates remarkably well to robots" (Darling, 2016, p. 223), and ever so much more so to lifelike social robots designed to elicit that reaction—even when people know that they are really dealing with a machine (Nass & Moon, 2000). Similarly, studies suggest that many people are wired not only to feel more empathy toward lifelike robots than to other objects, but that as a result, harm to robots feels wrong (Darling, 2016, p. 223). Thus, we might choose to ban the "abuse" of robots (beating, torturing) either because it offends people, or because we fear that some persons who abuse robots may develop habits of thought or behavior that will carry over into their relationships with live people or animals, abuse of which is commonly prohibited. Were we to find empirical support for the hypothesis that abuse of lifelike, or perhaps humanlike, robots make abusive behavior toward people more likely, that would provide strong grounds for banning some types of harms to robots—a correlative to giving robots certain rights against humans (Darling, 2016, pp. 226–231). In Hohfeldian terms, if persons have a duty not to harm a robot, then, correlatively, the robot has the right not to be harmed by those persons (see Schlag, 2014, pp. 200–203).

2.2 Legal Issues in Human-Robot Interaction

The human tendency to anthropomorphize robots—and, conversely, the tendency among some to place undue faith in technology—can allow robots to become a means to manipulate people who interact with robots. The hypotheticals—and most are, for now, just hypotheticals—are legion. Perhaps if people become attached to their robot pets or companions, the firms making or running the robots could use this emotional bond to extract payments, e.g., for upgrades, from users. Or maybe the robots could be used to sneak ads into the home, in the guise of ordinary conversation—a particular worry for robots designed for children. Perhaps they will ask questions designed to reveal personal information (Darling, 2016, p. 221). Less hypothetical is the observation that people not only tend to anthropomorphize robots, but also tend to ascribe gendered characteristics to them based on their functions, as well as to apply common social behaviors such as politeness and reciprocity (Nass & Moon, 2000).

The issue of (over) trust gains salience from the reality that some robots are designed to collect information about users, such as for medical monitoring, although consent is required by the user or their guardian. Others record information about their environment, such as cleaning robots assembling floor plans, and then share that information with the company that

[2] See *Thaler v. Vidal*, 43 F.4th 1207 (Fed. Cir. 2022).

makes the cleaners (Kaminsky, 2015, discusses the risks with a focus on consent and disclosure issues). We know that some devices that rely on voice recognition have had conversations monitored by remote humans without notice to the humans. Indeed, potential robot fakery has multiple modes ranging from robots pretending to be people in screen chats to people pretending to be robots in order to make their product seem more sophisticated, or their advice appear scientifically based (Brennan-Marquez, Levy & Susser, 2019). All of these scenarios excite some appetite for regulation (see for example Hartzog, 2015), although little has come to pass as yet. In some special cases, such as in treatment or fiduciary relationships, these deceptions are either illegal or a violation of professional ethics and duties; but in other contexts, while immoral they may become actionable fraud only when someone can show injury—and the US Supreme Court has suggested in recent cases that privacy harms without some tangible monetizable consequences will not support a damages action, even if there is a law providing statutory damages.[3]

Regardless of their underlying legal status, some robots are capable of far more rapid reactions to stimuli than people. That can be extremely valuable, whether adjusting the trim of an aircraft, changing cooling settings in a nuclear power plant, or targeting incoming ordinance. But as sensors and programming are never perfect, rapid reactions also raise the specter of rapid errors, sometimes catastrophic, leading to suggestions that when a robot has the power to do significant harm, the law should require a "human in the loop" to reduce the risk of unwanted outcomes (Jones, 2015). One risk, however, is that inserting a human into the OODA loop will in the most critical cases mean that the system (robot + human) responds to danger more slowly, perhaps in some cases too slowly for safety or effectiveness.

Madeline Elish describes another family of risks which she calls "moral crumple zones," in which the human is assigned responsibility, and thus legal liability, without the practical means to shoulder it. For example, if the overall system is poorly designed, when a crisis occurs the human purportedly in the loop may never have enough time to evaluate the facts and make a good decision. A related failure mode becomes likely if the human's monitoring job is ordinarily routine and boring, leading attention to flag, or if the human's monitoring role involves overseeing so many robot tasks simultaneously that real oversight becomes practically impossible, as in for example a security guard tasked with monitoring a large number of security robots (Elish, 2019).

If robots are not legal persons, it follows that a robot cannot be legally responsible for harms or crimes that it causes. Ordinarily one would expect robot harms and crimes to be charged to someone seen as responsible for the robot's action. If, however, there is no "human in the loop" then who should that be? For civil damages, such as tort, commonly we seek to find the person whose actions were the proximate cause, or whose negligence permitted the harm to occur. Sometimes that may be clear: someone set the robot's actions in motion, or someone failed to anticipate a foreseeable contingency in the robot's construction or operation, or someone with a heightened duty of care negligently failed to monitor the robot's actions. Other situations, however, will be considerably less clear, and then one may need to resort to more elaborate concepts of (sometimes gross) negligence.

Crimes committed by (or is it "via") robots present additional complexities. It is not hard to say that if a person were to program a robot to steal, that person bears criminal liability for the theft, just as they would with any other instrumentality. But what if the person owning

---

[3] See *Spokeo, Inc. v. Robins*, 578 US 330 (2016).

or operating the robot lacks *mens rea*, as in the actual case of the robot programmed to randomly order things online and have them shipped to an art exhibition—which then ordered some Ecstasy tablets (Power, 2014). Is *mens rea* present when the person who set the robot in motion neither intended nor foresaw the crime of illegal purchase of narcotics? Sometimes, as in the case of the shopping robot, it may be sufficient to say that the person operating the robot should have foreseen the action and guarded against it. Sometimes, however, that judgment may be quite problematic, especially if the crime is remote in time or probability.

2.3 Legal Issues from Emergent Behavior

The most interesting, but also most legally difficult, robots may be those designed to learn from experience, ideally with relatively little human supervision. Suitably primed robots have taught themselves to walk (Wu et al., 2022). As robots interact increasingly with people, they will be called upon to take on, and in some cases to learn, increasingly complex tasks that cannot be easily specified in advance—or perhaps at all. Increasingly, therefore, designers will program robots to learn by doing. While this is flexible and often efficient, it also means that robots will inevitably learn to do things in unpredictable ways, and indeed learn to do unpredictable things altogether. This "emergent behavior," that is, "behavior which is useful but cannot be anticipated in advance by operators" (Calo, 2014, p. 5) is very much a feature, not a bug,

Both the complexity and the interactivity of robot systems capable of emergent behavior create opportunities for injury. In ordinary cases where poor design, poor manufacture, or a programming error, causes an automated tool to injure someone or something, it is not difficult (at least theoretically) to assign moral and legal responsibility. Similarly, in the absence of designer/manufacturer/programmer error, if a tool's operator misuses it in some fashion, there too the liability seems clear. And, indeed, modern tort law also has a variety of means to share out liability in scenarios in which designer, operator, victim, and even bystander may all be partly responsible for a harm.

But what is to be done when the emergent behavior is the result of multiple interactions with different people over the course of the robot's existence? Is it right to blame the designer for including a feature that is normally not just benign but may be necessary for the robot to learn to undertake complex tasks? Perhaps we should blame the robot's operator for failure to supervise it properly, but the nature of an emergent behavior is that it could manifest without warning. We do not at present have any consensus, and indeed little experience, as to how to assign liability in the case of injury caused by a robot's emergent, unexpected, behavior. The problem is, however, certain to emerge.

3. AUTONOMOUS VEHICLES

3.1 Robot Cars

Conversely, although we do have extensive experience as to how to assign liability for car accidents, and indeed how to regulate cars for safety, this experience turns out not to make the project of assigning liability for accidents involving robot cars (aka "self-driving cars" or "autonomous vehicles") as easy as one might hope. Part of the problem is that in the United States, the regulation of motor vehicles is shared between national, state, and local entities.

The federal National Highway Safety Administration (NHTSA) regulates many aspects of motor vehicle safety. NHTSA mandates minimal safety standards, e.g., for crashworthiness. It initiates recalls when it finds defects in the car manufacturing process. And it regulates or coordinates a number of traffic safety rules. States issue drivers' licenses, subject to a number of federal rules, e.g., about acceptable forms of identification. Most traffic rules, such as speed limits (subject to a national cap set by Congress) and parking, are the domain of state governments, and often localities. Enforcement of traffic laws varies by state, but local police and traffic enforcement departments usually do most of it. Robot cars and their users are subject to all of these regulatory regimes.

There is also an international aspect to robotic vehicle regulation. The World Forum for Harmonization of Vehicle Regulations is an intergovernmental platform, hosted by the United Nations Economic Commission for Europe (UNECE), that creates regulatory frameworks for the performance and safety of vehicles which it invites member states to adopt. For example, it recently extended the maximum recommended speed for land-based vehicles with automated lane-keeping systems from 60 km/h to 130 km/hr.[4] Transnational issues will also arise when users seek to take a self-driving vehicle across national borders (see Smith, 2020).

Many US states have enacted laws designed to encourage the testing and ultimate deployment of self-driving cars by creating exceptions to rules requiring drivers. Several firms have deployed small fleets of cars, some with monitors in the front seat, but also some without anyone but the passenger. For example, Cruise introduced a driverless taxi service in parts of San Francisco—but it then recalled all its robotaxis in the United States following a crash (Wessling, 2022).

These and other experiments raise questions of both social policy and liability law. How law and policy will react will depend on how safe self-driving cars prove to be, and even more so on *how* they are unsafe and to whom they are dangerous. Society will benefit if self-driving cars prove to be, on average, safer than legacy cars with human drivers, and many have suggested that they should not be allowed to proliferate until this is proven to be the case. But even once robot cars clear this bar, we will need to know if the harms from their accidents replicate existing patterns, or if they instead have a tendency to cause a different sort of accident. For example, self-driving cars might be more likely than legacy cars to injure other vehicles, or swipe bystanders, instead of harming the passengers. Or if, for example, self-driving cars were generally more safe but more likely to run over children, that ought to give regulators pause (cf. The Dawn Project, 2022). More generally, the issue of how to tune the safety tradeoff between the passengers and others remains to be confronted. Some philosophers have suggested that cars should quiz passengers as to their driving preferences, e.g., whether to prioritize the safety of people in the car or outside it, or whether to drive quickly or carefully, in order to replicate human driving behavior (see, e.g., Millar, 2017; Millar & Kerr, 2016), but this seems unlikely to catch on—which is probably just as well.

The safety tradeoff issue captured the popular imagination with Judith Jarvis Thompson's "Trolley Problem" which explores the moral consequences of various actions and inactions by positing accidents with different sorts of casualties (Thompson, 1976). Aspects of the Trolley Problem were gamified in a website called the "Moral Machine" (Awad et al., 2018) that asked

---

[4] See *World Forum for Harmonization of Vehicle Regulations, Amendment to UN Regulation*. United Nations: Economic and Social Council. (adopted June 22, 2022). Retrieved from https://unece.org/fileadmin/DAM/trans/main/wp29/wp29resolutions/ECE-TRANS-WP29-1140e.pdf

participants to give their views about "moral decisions made by machine intelligence, such as self-driving cars." The site attracted millions of participants from around the world, prompting a formidable philosophical critique that the "Trolley Problem is precisely the wrong tool" for thinking about how automated vehicles should make life or death decisions while driving: "[T]he Trolley Problem frames the question as if all we need to do is figure out what an individual ought to do while driving, and then make that the rule for autonomous vehicles" (Jacques, 2019). Crowdsourcing moral decisions to generate rules, Abby Everett Jacques argued, leads to repugnant results based on asking the wrong questions and providing the wrong, or vastly incomplete, information for decisions—not to mention that often the information provided on the website (e.g., the character or occupations of potential victims) is not information that would be available to a driver, whether human or mechanical. Even offering gross and visible information (baby in stroller vs. elderly victim) confuses an individual choice with a morally sensible policy or algorithm. Were we to rule that we must program our cars to deprioritize elderly jaywalkers over younger people, for example, that would in effect amount to legislating an increased chance of the death penalty for older jaywalkers (Jacques, 2019, pp. 7–8).

Even if self-driving cars prove safer in every way than legacy cars, for the foreseeable future they will still have accidents. (It may also bear mention that the issues of liability and compensation for car and other accidents is particularly important in nations like the United States that do not have a functioning national health system and have high medical costs, since this can greatly increase the amount at stake in any subsequent dispute as to liability.) In the meantime, we have driver-assistance, and partially self-driving cars which return control to the operator when the onboard guidance system encounters something it cannot recognize or deal with. Drivers with "Level 3" conditional driver automation, as defined in the SAE International Taxonomy,[5] are already notoriously inattentive when the guidance system is in control. As the level of automation increases to "Level 4," high driving automation, and hand-offs back to passenger control become less frequent, we can reasonably expect that drivers will be even less prepared for them. In short, absent judicial or regulatory intervention, drivers who purchase cars that are equipped with anything less than fully autonomous guidance systems are at risk of becoming the legal if not moral crumple zone (Goldenfein et al., 2020).

In contrast, if the passengers never have control of a vehicle that causes an accident (nor any say in how aggressively it drives), then it would seem absurd to make them in effect the insurer of the robot car's truly responsible party or parties. In the absence of an intentional harm, we commonly seek to put liability on the "least cost avoider" for an accident—the party who might have prevented it at the lowest social cost. Doing so, economic theory teaches, is most likely to align incentives to prevent accidents without creating too great a disincentive to undertake activities that might cause harms. Who, then, is the least cost avoider for an accident caused by a fully autonomous vehicle? Most likely it is its maker: the maker made the hardware and either created or chose the software running the vehicle, and thus is best positioned to choose how much to invest in testing and safety features. One thus would expect that the maker of a vehicle, or perhaps the provider of a taxi service (since they are free to choose what kind of robot taxi to acquire and provide), should become the holder of liability.

---

[5] See *J3016C Taxonomy and Definitions for Terms Related to Driving Automation Systems for On-Road Motor Vehicles*. SAE International. (April 30, 2021).

Yet, this is not clearly the answer to the liability question, at least not in the United States. In the absence of legislation changing the tort law defaults, how the United States determines liability for self-driving cars—and probably many other sorts of robots also—will likely turn initially on whether the robot presents as a product or a service.

If the robot presents as a service, e.g., a taxi or other on-demand arrangement, then ordinary rules of negligence will apply to the consumer's liability since the consumer is not the purchaser of the product; the service provider may be subject to product liability law as described below. That said, even for consumers, complexities abound as to whose negligence is at issue and against what standard of care it would be measured. For example, a robocar might be considered a product when sold. But if the software operating it were sold sufficiently separately, and continually updated, that software running the robocar would arguably be a separate service; on the other hand, if there is only one software choice for a given brand or model of robocar, then there would also be a good argument that the two are sufficiently inextricably bound, and the hardware predominates in the overall transaction, and thus the software would be lumped in as a product too.

If end-users buy (or long-term lease) their vehicle, software included, as the majority of US drivers currently do, then US law presumably will apply product liability law to all of it. If there is a manufacturing defect, where due to some factory error the product diverged from the intended design, then the states are uniform in following the Restatement (Second) Of Torts § 402A (1965),[6] and placing joint and several liability on the manufacturer and all the parties in the chain of sale to the end-user. If, however, the manufacturer produced the autonomous vehicle according to plan but there is an underlying design defect, the states are far from uniform as to whether a negligence or a strict-liability-like standard applies, and thus also not uniform as to the showings a victim, be it the end-user or a third party, must make in order to recover damages from the manufacturer.

Under the Restatement 2d of Torts, as amended in the 1960s and then developed in a series of court decisions during the next decades, the majority rule for design defects became the "consumer expectations test" which makes the seller of a product liable if the product is in a defective condition unreasonably dangerous to the consumer.[7] The trier of fact can infer the existence of a design defect if the product fails to meet the reasonable expectations of consumers. Thus, even where there is no evidence, direct or circumstantial, available to prove exactly what sort of manufacturing flaw existed, a plaintiff may establish her right to recover by proving that the product did not perform in keeping with the reasonable expectations of the user; for this reason and others, the consumer expectations test engenders critique (see Owen & Davis, 2020, at § 5:16).

A product falls beneath consumer expectations when the product fails under conditions concerning which an average consumer of that product could have fairly definite expectations, which makes it in effect a strict liability rule once that line is crossed. Importantly, however, the consumer expectations test does not apply where technical and mechanical defects are

---

6 See Restatement (Second) Of Torts § 402A (1965).

7 Section § 402A states that anyone "who sells any product in a defective condition unreasonably dangerous to the user or consumer" is strictly liable for the damages. *Id.* Comment *i* defined "unreasonably dangerous" as being dangerous "beyond that which would be contemplated by the ordinary consumer who purchases it." *Id.* at cmt. *i*, The modern test is sometimes traced to a deeply influential article, John Wade, *On the Nature of Strict Tort Liability for Products*, 44 Miss. L.J. 825, 837–838 (1973) (sixth factor of multi-factor test).

alleged which require an understanding of precise behaviors of obscure components of products under complex circumstances of a particular accident.[8] Exactly when robots move from machines with technical and obscure parts to devices as commonplace as cars, microwaves, or refrigerators, is an evolving and as yet largely unexplored question. Roombas are ubiquitous and surely fall into the ordinary consumer product category, as do perhaps hobby drones and robotic pets—but would the control software for a fully self-driving car, as opposed to a Tesla with Level 3 driver assistance? Or is the answer that Tessla might be estopped from arguing a lack of consumer expectations due to advertising the excellence of the software?

For many years, critics of the "consumer expectation" test objected that plaintiffs could prevail without showing flaws in the design, and had no obligation to present a reasonable alternative design. Instead, if the defense wants to argue there was no better practical design, it bears the burden of production on this issue once the plaintiff has made a prima facie case on what consumers would expect. In response, the more recent Restatement 3d of Torts: Product Liability § 2 takes a very different view of how design defect claims and defenses should work. It abolishes the "consumer expectation" test and replaces it with a "risk-utility" test (see Twerski & Henderson, 2009). Under this test, to prevail the plaintiff must show that the risk of the design exceeded its value; commonly the only way to do this is to proffer a safer alternate design that is not (substantially) more expensive to produce and which would not have caused, or greatly reduced the chance of, the harm.

Since the defendant has all the information about how the robot was designed and made, it will be hard for most plaintiffs to meet this burden. In most cases, it will require expensive experts, which makes it much harder to bring cases, and may require juries to hear very technical evidence. Critics of the 3d Restatement, including a substantial number of state supreme courts that have decided not to adopt this provision of the Restatement 3d, agree that it moves the liability standard away from strict liability by focusing on the foreseeability of the risk of harm, including a cost-benefit analysis. "Rather than focusing on the design of the product, it focuses on the conduct of the manufacturer" (see generally Owen & Davis, 2020, at §§ 5.6, 5.7).[9]

Thus, at present, liability in the United States for many robot-caused injuries presents up to three levels of ambiguity: 1) whether the robot presents as a product or service; 2) whether the type of robot or the relevant part of it is an ordinary consumer product about which consumers have grounded expectations; 3) whether the state whose law applies uses Restatement 2d strict liability for design defects or Restatement 3d more negligence-like principles. Liability may fall out differently in civil law regimes. For example, one analysis suggests that despite the EU Directive establishing strict liability for makers of defective products ((EU) Council Directive No. 85/374/EEC of July 25, 1985), in Germany the primary liability in practice will fall on the "keeper" of the vehicle due to the difficulty of proving fault against manufacturers (Ebers, 2022).

As states are likely to be anything but uniform in their answers to these questions, the robot accident liability issue seems ripe for a uniform national solution, ideally one that would reduce the substantial costs of litigating each of these issues in every state. This might involve legislation, or even a national robot regulator (Calo, 2014) but in some industries such as self-driving cars might also have a partial de facto solution if manufacturers choose to bundle

[8] See *Soule v. Gen. Motors Corp.*, 882 P.2d 298, 305 (Cal. 1994).
[9] See *Aubin v. Union Carbide Corp.*, 177 So. 3d 489, 506 (Fla. 2015).

car insurance for passengers into every sale in order to trumpet the safety of their product. Examples of the extensive writing on insurance issues include Templeton (2020), Geistefeld (2017), and Lior (2022).

In this vein, Kenneth Abraham and Robert Ravin argue powerfully that autonomous vehicles present a chance to move away from old tort paradigms for car accidents, and move to an administrative compensation scheme (Abraham & Rabin, 2019). It would be interesting to see how far this might be generalized to other robot-related accidents.

3.2 Drones (UAVs)

Unmanned aerial vehicles (UAVs, or more colloquially, drones) are either remote-controlled or autonomous vehicles, and thus potentially candidates for a similar liability regime to terrestrial vehicles, but before taking to the air they must navigate a substantially different regulatory regime. First, the national regulator, the Federal Aviation Administration, has sole regulatory powers over "navigable airspace" (Dolan & Thompson II, 2013, at 2), although at times the FAA has suggested it has regulatory authority over all airspace.[10] Second, the FAA has some concurrent authority with states for lower altitude flights, and it has not been shy about using it. Consequently, there is already a detailed set of national rules that drone manufacturers must follow to sell their products, and also significant rules for both commercial and hobbyist use of drones. The FAA's primary concern has been safety, which it defines to mean that drones should not cause physical injury, should not interfere with flight operations, should not intrude into airspace defined as security areas (e.g., most of Washington, DC), and should carry identifying marks such that if a drone is misused, not least for terrorism, it should be possible to identify the owner. By mid-2021, users had registered well over a million recreational drones with the FAA (NCSL, 2021). The FAA has also been very cautious about allowing commercial drone companies to operate delivery drones outside the line of sight of the operator, although trials of more relaxed rules are in progress.

Because they carry cameras into spaces people are accustomed to think of as private, ranging from remote wilderness to right outside windows on high floors of apartment buildings, drones create new privacy threats. Thomasen, 2018, suggests that drone regulation should be considered via a feminist lens. Indeed, there is extensive anecdotal evidence to support the claim that men use drones to spy on women much more than the reverse.

Despite academic calls to consider privacy harms (e.g., Froomkin & Colangelo, 2020; Froomkin, Arencibia & Colangelo-Trenner, 2022), the FAA does not see its mission as extending to the ways in which drones can intrude on the personal privacy of those overflown.[11] Drone trespassing and drone-enabled spying, voyeurism, and stalking are thus primarily an issue of state tort and, increasingly, statutory law. The trespass issue is complicated by uncertainty in most states as to the extent of the "vertical curtilage" and to what extent if any airspace above private property below the 400-foot line may be open to drone flights. As to spying, although torts such as intrusion upon seclusion are actionable nationwide, they remain

---

[10] See FAA. (2018, July 20). *FAA Statement-Federal vs Local Drone Authority*. United States Department of Transportation. Retrieved from https://www.faa.gov/newsroom/faa-statement-federal-vs-local-drone-authority.

[11] See FAA. (2019, December 31). Remote Identification of Unmanned Aircraft Systems. Federal Aviation Administration.

exotic and rare claims, and damages can be very hard to monetize—without which there is no claim. Several states have legislated special protections against overflights for favored industries, sporting events, or police stations (Skorup, 2022); several more have enacted various limits on surveillance via drone by state and local law enforcement, although similar limits usually do not apply to manned vehicles such as helicopters and airplanes (McNeal, 2014). A few states have also legislated bans or limits on overflights of private property (Skorup, 2022, pp. 163–164). The extent to which landowners may engage in self-defense against drone overflights also remains largely untested (see Froomkin & Colangelo, 2015).

Just as robots intrude into formerly private spaces, they also challenge our conception of what is public space and how one should regulate it. May terrestrial delivery robots use the sidewalk? May often small and slow delivery robots go on the road where there is no sidewalk? If they are UAVs, may they overfly streets, or private property? More generally, the combination of enhanced delivery services with on-demand robotic transport likely will enable and perhaps necessitate a rethinking of many zoning and urban planning rules. Discussions of the complexities include Marks, 2019, Thomasen, 2020, Woo, Whittington and Arkin, 2020, and Gilbert and Dobbe, 2021.

3.3 Autonomous Vehicles at Sea

At present, the place where autonomous vehicles may have the freest rein is likely the ocean. Although weather and sea present challenges absent on paved roads or in airspace on a clear day, generally speaking there are fewer obstacles to maneuver around, and, at least away from ports and coastlines, the chance of running into another vehicle or a person is relatively low. Numerous trials of both civilian and military, and more or less autonomous, surface and subsurface robotic vessels are underway.

One difficulty, however, is that the current international regime governing ships at sea assumes that "vessels" are ships with persons aboard. Vessels that fly the flag of their nation of registration enjoy many legal protections; indeed, to seize a manned vessel against the will of its crew is one of the oldest and universal international crimes—piracy. A self-propelled, self-guided, fully-functioning robotic ship with no one aboard does not fit either the classical definition of a "vessel" nor that of the United Nations Convention of the Law of the Sea Article 94, which presumes that a ship carries a "master and officers" (Brett, 2022). Indeed, some might suggest that it more closely fits the definition of "salvage"—property abandoned at sea and thus in theory available for any finder to acquire and keep. Undersea drones raise additional complexities, including potential privacy issues if they surveil passing ships or undersea installations (Brett, 2019). While unmanned vehicles may have relatively clear sailing on the seas, the legal regime governing sea-based robotic ships and drones will remain needlessly turbulent until maritime law is updated by international agreement (Brett, 2022).

4. BATTLEFIELD ROBOTS (LAWS)

There is another, even more consequential, area where the governing law is at least as much international as national: the battlefield. On the one hand, the prospect of having robots fighting promises to reduce military casualties for those deploying the robots. On the other hand, reliance on so-called "killer robots" might make decision-makers more willing to engage in hostilities as reliance on robots would reduce risks to human troops. Either way, warbots raise difficult practical and legal issues relating to dangers to civilians, with opinions divided on whether they would be safer (Arkin, 2010 and, very optimistically, Lewis, 2020) or more dangerous (Jenks & Liivoja, 2018) for bystanders.

Whether and how there can be "human in the loop" control is a critical question for military robots, as robots with autonomous capabilities, also known as Lethal Autonomous Weapons Systems (LAWS), create a tension with the doctrine of command responsibility that is fundamental to the modern laws of armed combat (LOAC) (Schwarz, 2018). For almost 70 years, if not more, it has been a cornerstone of both US and international law that an officer's, and indeed, every combatant's, responsibility for war crimes cannot be disclaimed by saying "I was just following orders." The duty to not commit war crimes, e.g., by causing militarily disproportionate harm to civilians, is uniquely personal and may indeed require violating orders in extreme cases. On the other hand, a commander is usually not held responsible for things s/he could not control except perhaps if it was highly foreseeable and could have been prevented.

LAWS implicate command responsibility because they are so fast, and because they have "black box" elements that the people using them—often soldiers, sailors, airmen, or marines who may not have high rank—may not understand. Indeed officers may not understand or be able to control them either, and this undermines a duty that military officers commonly see as moral as well as legal.

Some scholars argue that LAWS are per se illegal under international law, either under Article 36 of the Additional Protocol I (1977) of the 1949 Geneva Convention (see Press, 2017, pp. 1345–1347 for a summary) or, even more controversially, under the Martens Clause in the 1899 Hague Convention (II) (Asaro, 2016b); others have proposed various national and international civil remedies for harms from autonomous weapons (e.g., Crootof, 2016, 2018). Several international and national NGOs banded together in 2013 to form an international "Campaign to Stop Killer Robots,"[12] a plea which has been endorsed by more than 20 countries, although by none of the major world military powers. The campaign has not stopped investment in robotic and autonomous military technology, although it should be noted that the US Department of Defense has also produced high-minded, if perhaps not completely constraining, "AI Principles: Recommendations on the Ethical Use of Artificial Intelligence" (US Dept. of Defense, 2019).

5. POLICE ROBOTS ("ROBOCOPS")

"What develops first in the military often finds its way to domestic policing" (Joh, 2016, p. 528). Advances in military technology regularly filter back to the civilian economy, either directly or via laws that allow the US Defense Department to transfer military hardware to local law enforcement agencies (Dansby, 2020). Robotic policing ("robocops") could have either appealing or horrifying implementations. In the appealing vision, mechanizing some policing functions would remove undesirable features of current policing practices. In theory, robocops could be programmed to treat like cases alike, to act consistently with everyone regardless of race and regardless of which robot is doing the policing. Human police will be safer; for example, since robots are not people, they could be sent into dangerous

[12] The group's homepage is https://www.stopkillerrobots.org/.

situations, such as hostage-takings, or active shooters in schools and public places, in which they could collect information and engage in negotiations—and perhaps attempt to subdue suspects—without risk to police officers. Although robocops might be built to carry firearms or less-lethal tasers, they also could be limited to non-lethal force, although if the robot were designed to subdue armed suspects that would create a very substantial design, and if at all autonomous, programming challenge. And since robots can be set to record everything they do, audits will be easier and, in principle implementation of rule changes will be easier (just change the programming!) and instantly ubiquitous.

But the reality does not, and in the foreseeable future is unlikely to, even come close to this idyllic vision: as legal scholars have demonstrated, police departments are buying expensive robot and AI technologies that do not work as advertised, rely on databases that replicate or exacerbate existing biases due to reliance on "dirty data" (Richardson, Schultz & Crawford, 2019), impose sometimes invisible surveillance on citizens—usually against communities that already suffer from over-policing (Joh, 2016, 2022)—and can lead to what Elizabeth Joh has called "unexpected consequences," distorting the practice of policing (Joh, 2022). Furthermore, to the extent that robot selection and deployment decisions are made via procurement, they avoid structures designed to give citizens the right to seek judicial review of administrative regulations, including, often, systemic policing policies. These issues tend to be explored in the context of AI (see, e.g., Huq, 2020; Mulligan & Bamberger, 2019; Conglianese & Lehr, 2017), but they are equally applicable to robots, and not just because any but the simplest robotic policing system will rely on an AI to function.

As Peter Asaro notes, policing involves many discrete activities. Just a simple stop-and-frisk can be divided into: 1) profiling, detecting signs of possible illegal activity and choosing who to investigate or stop-and-frisk); 2) implementation, doing the stop-and-frisk including deciding if there is a legal violation worthy of an arrest or citation; 3) "justice," that is review, auditing, charging decisions, and in some cases trials. Each of these stages requires an algorithm, relevant data, and a feedback loop, which may involve a human in the loop during the stop-and-frisk and/or subsequent incorporation of new data from the encounter into the underlying database (Asaro, 2016c).

In fact, none of these activities are easy to mechanize except perhaps the recording/surveillance/detection function. First, there are physical and mechanical issues with designing a robot capable of safely frisking the wide variety of sizes and shapes of persons ranging from children, to very large people, to people with medical conditions requiring casts, wheelchairs, or connected medical devices—and capable of restraining them with proportionate force if needed. Second, and even more difficult, is the programming problem involved in enforcing even a very simple statute.

As an interdisciplinary group of engineers and lawyers demonstrated with a simple experiment, coding enforcement of even a simple speed limit is a surprisingly difficult problem that requires extensive interpretation of the legal rules by the coders, including judgments as to whether their goal is to enforce the law on the books, what they take to be the intent of the law on the books, or the law in action as they understand it (Shay et al., 2016b). The problems only get worse if one tries to automate more complicated aspects of law enforcement, especially areas where subjective judgments regarding large numbers of factors (what exactly is "suspicious activity") may be needed (Shay et al., 2016a).

Further, existing US law likely—but it is too early to say "certainly"—constrains robotic law enforcement in various ways. Some sense-enhanced detection technologies that one might wish to deploy on a robot may amount to an unconstitutional warrantless search if the technology is not in common use (see Kyllo v. United States, 533 US 27 (2001))—unless the fact that the detector is applied to targets in public places allows the courts to find that the contraband or other matter detected was effectively as exposed as if in plain sight. Similarly, at present, US legal doctrine justifies a number of exceptions to the ban on warrantless searches of persons and of the interior of vehicles on the grounds that police need to assure themselves that a person being stopped for questioning has no weapons that could be used to harm the officer. That safety justification for a large number of pat-downs and searches might vanish if the search is conducted by a machine rather than a person. In addition, if robotic cars meticulously follow traffic laws, the minor-violation justification for most stops and searches of cars may also vanish. Many other issues that are reasonably settled for human law enforcement, such as what constitutes reasonable force, will also be open to re-examination when the actor is a robot (Simmons, 2020).

Early returns on robotic policing are not encouraging. In 1996, Dallas police strapped a bomb to a robot and used it to blow up a person who had barricaded himself second floor of a building after killing five police officers and wounding seven others (Peterson, 2016). More recently, in 2021, the New York police department leased a robotic dog from Boston Dynamics, and touted Digidog's ability to "save lives." However, deployment in a public housing building spurred claims that it symbolized overly aggressive policing of poor communities, and the NYPD terminated its contract (Zaveri, 2021)—only to welcome robotic dogs back early in 2023 (Rubenstein, 2023).

There seems to be some tension between the literature that focuses on the human tendency to anthropomorphize and trust robots, and the strong reaction against some kinds of robot policing. On the other hand, we have not yet encountered much in the way of robots as criminal instrumentalities. Were robots to become the tool of choice for, say, bank robbery, one might reasonably expect that the public might accept, or even demand, anti-robot robotic police.

6. ROBOTS AND EMPLOYMENT

If robots do not kill us, perhaps they will just take our jobs. The plethora of tasks that robots (and AIs) seem likely to be capable of—ranging from construction and warehouse work to driving trucks to highly skilled jobs in the financial, legal, and medical worlds—inevitably raises concerns about the possibility that robots will displace workers on a large scale. Indeed, there is evidence from around the world that firms seek to replace workers—not just industrial and retail workers but also professionals when possible—with robots since they can work 24/7, can all be taught new behaviors simultaneously, are resistant to pandemics (although vulnerable to their own kind of virus), and do not go on strike.

Acemoglu and Restropo (2020) estimated that in the US one more reprogrammable industrial "robot per thousand workers reduces the aggregate employment-to-population ratio by about 0.2 percentage points and wages by about 0.42%" nationally, but in the area where the robot is deployed the employment-to-population ratio declines 0.77% and wages decline 0.77%. Every robot added to a commuting zone (a geographic areas used for economic analysis) reduces employment by six workers in that area. The study's authors speculate that future effects will be larger as industry uses more robots. Meanwhile, Acemoglu, LeLarge and Restrepo (2020) used French data to suggest that "firms adopting automation technologies reduce their costs and may expand at the expense of their competitors." Will there be a race to the bottom as firms vie not to be eliminated for being the last to replace their workers with robots?

There is at present no consensus as to whether the increasing use of robots (and AI) will cause substantial permanent unemployment, or just a significant but temporary effect due to skill mismatch. Even a "temporary" effect might be generational and severe, not just among industrial workers but also among other potentially replaceable lower-skilled workers, such as (in the United States) the approximately 1 million truckers who have few transferrable skills and the circa 10 million cashiers, retail salespersons, and first-line retail supervisors (Anderson, 2020). Many of the workers who are not displaced, including many in white-collar professions, may find themselves subjected to robot-administered monitoring and supervision where every action is tallied in search of maximum productivity (Kantor & Sundaram, 2022; Dzieza, 2020; Harwell, 2021).

Robots require capital expenditures, changing the ratio of capital to labor; with fewer workers, the productivity rates of those who remain employed should increase. But capital cost is sensitive to macroeconomic factors such as interest rates, and to legal factors; among the legal factors are changes in worker protection (increases make robots seem more attractive), and tax law. There has been substantial scholarly attention to tax issues as they affect the decision to invest in robots, and also as to whether tax policy might be used to protect some jobs from robot displacement. For example, because wages are not deductible expenses, but capital investment in machinery is either a deductible cost or creates a depreciable asset, researchers suggest that the tax system creates an incentive to replace labor with capital, as does the legal obligation to make unemployment insurance and other payments for human workers—but not for robots (e.g., Abbot & Bogenschneider, 2018; Kovacev, 2020). And some, including some members of the European Parliament, have proposed a robot tax in order both to discourage that substitution and to create a fund that would help displaced workers, although the Parliament did not take up the proposal (Reuters, 2017). So far, the idea has foundered on the fear of discouraging innovation and the difficulty of measuring or defining what number or fraction of jobs a given machine has displaced. Proposals to tax robots generally also will run into issues of defining what would be a taxable robot, and how to decide what level of taxation per robot would be appropriate.

The public has begun to take note of the employment threat that robots may present, although so far robots tend to be more unpopular in nations with higher inequality (Shoss & Cirlante, 2022). The academic debate, however, is overshadowed by an awareness that claims about the job-destroying effects of automation were prevalent in the 1960s—and proved utterly unfounded (Jaffe & Froomkin, 1968). Is this time different?

7. CONCLUSION

One thing that is clearly not different this time is that every major new technology—be it electricity, railroads, the Internet, or robots—creates new sets of ethical, legal, and social problems. Currently we are only in the early days of roboticization: the quality, quantity, and variety of robots are each poised for rapid growth. This creates important and necessary work for legislators, regulatory authorities, judges, standards-makers, and academics who seek to maximize the benefits of robots while minimizing the harms. This brief survey, itself only a snapshot in time, leaves out much, but even so it demonstrates that the tasks before us are deeply consequential.

New technologies also sometimes create opportunities for change in fundamental social arrangements. The roboticization of work, for example, offers the hope of freedom from dangers and drudgery, but could in some scenarios open the door to mass unemployment and immiseration. LAWS tantalize with the prospect of reducing human military casualties, but also present clear dangers of attenuation of command responsibility, civilian carnage, and that the prospect of war where aggressors suffer few casualties may prove too tempting to some. The odds that robots will become our masters in the foreseeable future seem negligible, but it is anything but inevitable that all robots-inspired changes will be benign, or that we will use robots to humanity's best advantage. The challenges for lawyers, policymakers, ethicists, and indeed everyone, are upon us.

REFERENCES

Abbott, R. (2016). I Think, Therefore I Invent: Creative Computers and the Future of Patent Law. *Boston College Law Review, 57*, 1079–1126.

Abbott, R. & Bogenschneider, B. (2018). Should Robots Pay Taxes? Tax Policy in the Age of Automation. *Harvard Law & Policy Review, 12*, 145–175.

Abraham, K. & Rabin, R. (2019). Automated Vehicles and Manufacturer Responsibility for Accidents: A New Legal Regime for a New Era. *Virginia Law Review, 105*, 127–171.

Acemoglu, D., LeLarge, C. & Restrepo, P. (2020). Competing with Robots: Firm-Level Evidence from France. *American Economic Association*. Retrieved from https://doi.org/10.1257/pandp.20201003

Acemoglu, D. & Restrepo, P. (2020). *Robots and Jobs: Evidence from US Labor Markets*. Chicago, IL: The University of Chicago Press.

AI Principles: Recommendations on the Ethical Use of Artificial Intelligence (2019, October 31). Department of Defense, Defense Innovation Board. Retrieved from https://media.defense.gov/2019/Oct/31/2002204458/-1/-1/0/DIB_AI_PRINCIPLES_PRIMARY_DOCUMENT.PDF

Anderson, D. (2020, September 8). *Retail Jobs Among the Most Common Occupations*. U.S. Census. Retrieved from https://www.census.gov/library/stories/2020/09/profile-of-the-retail-workforce.html

Arkin, R. (2010). *The Case for Ethical Autonomy in Unmanned Systems*. London: Journal of Military Ethics.

Asaro, P. (2016a). "Hands Up, Don't Shoot!": HRI and the Automation of Police Use of Force. *Journal of Human-Robot Interaction, 5*, 55–69.

Asaro, P. (2016b). *Robotic Weapons and the Martens Clause*. Cheltenham: Edward Elgar.

Asaro, P. (2016c). *Will #BlackLivesMatter to RoboCop?* We Robot. Retrieved from https://robots.law.miami.edu/2016/wp-content/uploads/2015/07/Asaro_Will-BlackLivesMatter-to-Robocop_Revised_DRAFT.pdf

Awad, E., Dsouza, S., Kim, R., Schulz, J., Henrich, J., Shariff, A., Bonnefon, J. & Rahwan, I. (2018). *The Moral Machine Experiment*. Berlin: Nature.

Bayern, S. (2021). *Autonomous Organizations*. Cambridge: Cambridge University Press.

Bayern, S. (2019). Are Autonomous Entities Possible? *Northwestern University Law Review, 114*, 23–47.

Birhane, A. & van Dijk, J. (2022, February 7). *Robot Rights? Let's Talk About Human Welfare Instead*. New York: Association for Computing Machinery.

Boyle, J. (2011, March 9), *Endowed by Their Creator? The Future of Constitutional Personhood*. Brookings. Retrieved from https://www.brookings.edu/articles/endowed-by-their-creator-the-future-of-constitutional-personhood/

Brennan-Marquez, K., Levy, K. & Susser, D. (2019). Strange Loops: Apparent versus Actual Human Involvement in Automated Decision Making. *Berkely Technology Law Journal, 34*, 745–771.

Brett, A. (2019). Secrets of the Deep: Defining Privacy Underwater. *Missouri Law Review, 84*, 47–92.

Brett, A. (2022). Regulating the Autonomous Ocean. *Brooklyn Law Review, 88*, 1–55.
Bridy, A. (2012). Coding Creativity: Copyright and the Artificially Intelligent Author. *Stanford Technology Law Review, 5*, 1–28.
Calo, R. (2010). People Can Be So Fake. *Pennsylvania State Law Review, 114*, 809–855.
Calo, R. (2014, September 15). *The Case for a Federal Robotics Commission*. Brookings. Retrieved from https://www.brookings.edu/research/the-case-for-a-federal-robotics-commission/
Citron, D. & Solove, D. (2022). Privacy Harms. *Boston University Law Review, 102*, 793–863.
Conglianese, C. & Lehr, D. (2017). Regulating by Robot: Administrative Decision Making in the Machine-Learning Era. *The Georgetown Law Journal, 105*, 1147–1223.
Craig, C. & Kerr, I. (2021). The Death of the AI Author. *Ottawa Law Review, 51(1)*, 31–86.
Crootof, R. (2016). War Torts: Accountability for Autonomous Weapons. *University of Pennsylvania Law Review, 164*, 1347–1402.
Crootof, R. (2018). International Cybertorts: Expanding State Accountability in Cyberspace. *Cornell Law Review, 103*, 565–644.
Dansby, J. (2020). Hammers and Nails: 1033 Program Reforms to Halt Police Militarization. *Civil Rights Law Journal, 31*, 109–134.
Darling, K. (2016). Extending Legal Protection to Social Robots. In R. Calo, A. Froomkin & I. Kerr (Eds.). *Robot Law*. Cheltenham: Edward Elgar.
Dolan, A. & Thompson II, R. (2013). *Integration of Drones into Domestic Airspace: Selected Legal Issues 2*. Congressional Research Service. Retrieved from https://fas.org/sgp/crs/natsec/R42940.pdf
Dzieza, J. (2020, February 27). *How Hard Will Robots Make Us Work?* The Verge. Retrieved from https://www.theverge.com/2020/2/27/21155254/automation-robots-unemployment-jobs-vs-human-google-amazon
Ebers, M. (2022, February 5). *Civil Liability for Autonomous Vehicles in Germany*. SSRN. Retrieved from https://ssrn.com/abstract=4027594 or http://dx.doi.org/10.2139/ssrn.4027594
Elish, M. (2019). Moral Crumple Zones: Cautionary Tales in Human Robot Interaction, *Engaging Science, Technology, and Society 5*, 40–60.
Froomkin, A. & Colangelo, Z. (2015). Self-Defense Against Robots and Drones. *Connecticut Law Review, 48*, 1–69.
Froomkin, A. & Colangelo, Z. (2020). Privacy as Safety. *Washington Law Review, 95*, 141–203.
Froomkin, A., Arencibia, P. & Colangelo-Trenner, P. (2022). Safety as Privacy. *Arizona Law Review, 64*, 921–987.
Geistfeld, M. (2017). A Roadmap for Autonomous Vehicles: State Tort Liability, Automobile Insurance, and Federal Safety Regulation. *California Law Review, 105*, 1611–1694.
Gilbert, T. & Dobbe, R. (We Robot Conference Draft, 2021). *Autonomous Vehicle Fleets as Public Infrastructure*). We Robot 2021. Retrieved from https://werobot2021.com/wp-content/uploads/2021/09/Gilbert__et_al_Autonomous-Vehicle-Fleets.pdf
Goldenfein, J., Mulligan, D., Nissenbaum, H. & Ju, W. (2020). Through the Handoff Lens: Competing Visions of Autonomous Futures. *Berkeley Technology Law Journal*, 835–910.
Hartzog, W. (2015). Unfair and Deceptive Robots. *Maryland Law Review, 74*, 785–829.
Harwell, D. (2021, November 11). Contract Lawyers Face a Growing Invasion of Surveillance Programs that Monitor Their Work. *The Washington Post*. Retrieved from https://www.washingtonpost.com/technology/2021/11/11/lawyer-facial-recognition-monitoring/
Hugenholtz, P. & Quintais, J. (2021). Copyright and Artificial Creation: Does EU Copyright Law Protect AI-Assisted Output? *International Review of Intellectual Property and Competition Law, 52*, 1190–1216.
Huq, A. (2020). Constitutional Rights in the Machine-Learning State. *Cornell Law Review, 105*, 1875–1954.
Jacques, A. (2019). *Why the Moral Machine is a Monster*. We Robot Draft. Retrieved from https://robots.law.miami.edu/2019/wp-content/uploads/2019/03/MoralMachineMonster.pdf
Jaffe, A.J. & Froomkin, J. (1968). *Technology and Jobs: Automation in Perspective*. New York: Praeger.
Jenks, C. & Liivoja, R. (2018, December 11). H*umanitarian Law & Policy, Machine Autonomy and the Constant Care Obligation*. Humanitarian Law & Policy. Retrieved from http://blogs.icrc.org/law-and-policy/2018/12/11/machine-autonomy-constant-care-obligation/
Joh, E. (2016). Policing Police Robots. *UCLA Law Review Discourse, 64*, 516–543.

Joh, E. (2022, January 15). *Reckless Automation in Policing*. Berkeley Technology Law Journal Online. Retrieved from https://btlj.org/2022/07/reckless-automation-in-policing/

Jones, M. (2015, January 13). Regulating the Loop: Ironies of Automation Law: Tying Policy Knots with Fair Automation Practices Principles. We Robot, *Vanderbilt Journal of Entertainment & Technology Law, 77.*

Kaminsky, M. (2015). Robots in the Home What Will We Have Agreed To? *Idaho Law Review, 51*, 661–677.

Kantor, J., Sundaram, A., Aufrichtig, A. & Taylor, R. (2022, August 15). The Rise of the Worker Productivity Score. *The New York Times*. Retrieved from https://www.nytimes.com/interactive/2022/08/14/business/worker-productivity-tracking.html

Kerr, I. (2004). Bots, Babes, and the Californication of Commerce. *University of Ottawa Law and Technology Journal, 1*, 285–324.

Kovacev, R. (2020). A Taxing Dilemma: Robot Taxes and the Challenges of Effective Taxation of AI, Automation and Robotics In the Fourth Industrial Revolution. *The Contemporary Tax Journal, 9*, 23–49.

Kyllo v. United States, 533 US 27 (2001).

Lewis, L. (2020, January 10). Killer Robots Reconsidered: Could AI Weapons Actually Cut Collateral Damage? *Bulletin of the Atomic Scientists*. Retrieved from https://thebulletin.org/2020/01/killer-robots-reconsidered-could-ai-weapons-actually-cut-collateral-damage/

Lior, A. (2022). Insuring AI: The Role of Insurance in Artificial Intelligence Regulation. *Harvard Law & Technology, 35*, 469–530.

LoPuki, L. (2018). Algorithmic Entities. *Washington University Law Review, 95*, 887–953.

Marks, M. (We Robot Conference draft, 2019). *Robots in Space: Sharing Our World with Autonomous Delivery Vehicles*. We Robot 2019. Retrieved from https://robots.law.miami.edu/2019/wp-content/uploads/2019/04/Marks_Robots-in-Space.pdf

McFarland, M. (2022, September 1). *Cruise Recalls Its Robotaxis After Passenger Injured in Crash*. CNN Business. Retrieved from https://www.cnn.com/2022/09/01/business/cruise-robotaxi-recall/index.html

McNeal, G. (2014). *Drones and Aerial Surveillance: Considerations for Legislatures*. Brookings. Retrieved from https://www.brookings.edu/research/drones-and-aerial-surveillance-considerations-for-legislatures/

Millar, J. (2017). Ethics Settings for Autonomous Vehicles. In P. Lin, K. Abney & R. Jenkins (Eds.). *Robot Ethics 2.0: From Autonomous Cars to Artificial Intelligence*. Oxford: Oxford University Press.

Millar J. & Kerr, I. (2016). Delegation, Relinquishment, and Responsibility: The Prospect of Expert Robots. In R. Calo, A. Froomkin & I. Kerr (Eds.). *Robot Law*. Cheltenham: Edward Elgar.

Mulligan, D. & Bamberger, K. (2019). Procurement As Policy: Administrative Process for Machine Learning. *Berkely Technology Law Journal*, *34*, 773–851.

Nass, C. & Moon, Y. (2000). Machines and Mindlessness: Social Responses to Computers. *Journal of Social Issues, 56*, 81–103.

Owen, D. & Davis, M., *Products Liability*. Eagan: Thompson Reuters.

Power, M. (2014, December). What Happens When a Software Bot Goes on a Darknet Shopping Spree? *The Guardian*. Retrieved from https://www.theguardian.com/technology/2014/dec/05/software-bot-darknet-shopping-spree-random-shopper

Peterson, A. (2016, July 8). In An Apparent First, Dallas Police Used a Robot to Deliver Bomb That Killed Shooting Suspect. *The Washington Post*. Retrieved from https://www.washingtonpost.com/news/the-switch/wp/2016/07/08/dallas-police-used-a-robot-to-deliver-bomb-that-killed-shooting-suspect/

Press, M. (2017). Of Robots and Rules: Autonomous Weapon Systems in the Law of Armed Conflict. *Georgetown Journal of International Law, 48*, 1337–1366.

Reuters (2017, February 16), *European Parliament Calls for Robot Law, Rejects Robot Tax*. Retrieved from https://www.reuters.com/article/us-europe-robots-lawmaking/european-parliament-calls-for-robot-law-rejects-robot-tax-idUSKBN15V2KM

Reynolds, E. (2018, January 6). *The Agony of Sophia, the World's First Robot Citizen Condemned to a Lifeless Career in Marketing*. WIRED. Retrieved from https://www.wired.co.uk/article/sophia-robot-citizen-womens-rights-detriot-become-human-hanson-robotics

Richards, N. & Smart, W. (2016). How Should the Law Think About Robots? In R. Calo, A. Froomkin & I. Kerr (Eds.). *Robot Law*. Cheltenham: Edward Elgar.

Richardson, R., Schultz, J. & Crawford, K. (2019). Dirty Data, Bad Predictions: How Civil Rights Violations Impact Police Data, Predictive Policing Systems, and Justice. *New York University Law Review, 94*, 192–233.
Rubenstein, D. (2023, April 11), Security Robots. DigiDog. GPS Launchers. Welcome to New York. *New York Times*. Retrieved from https://www.nytimes.com/2023/04/11/nyregion/nypd-digidog-robot-crime.html
Sawers, P. (2020, January 10). *Chinese Court Rules AI-written Article Is Protected by Copyright*. VentureBeat. Retrieved from https://rai2022.umlaw.net/wp-content/uploads/2022/02/16_Chinese-court-rules-AI-written-article-is-protected-by-copyright.pdf
Schlag P. (2014). How to Do Things with Hohfeld. *Law & Contemporary Problems, 78*, 185-234.
Schwarz, E. (2018, August 29). The (Im)Possibility of Meaningful Human Control for Lethal Autonomous Weapon Systems. *Humanitarian Law & Policy*. Retrieved from https://blogs.icrc.org/law-and-policy/2018/08/29/im-possibility-meaningful-human-control-lethal-autonomous-weapon-systems/
Shay, L., Hartzog, W., Nelson, J., Larkin, D. & Conti, G. (2016a). Confronting Automated Law Enforcement. In R. Calo, A. Froomkin & I. Kerr (Eds.). *Robot Law*. Cheltenham: Edward Elgar.
Shay, L., Hartzog, W., Nelson, J. & Conti, G. (2016b). Do Robots Dream of Electric Laws? In R. Calo, A. Froomkin & I. Kerr (Eds.). *Robot Law*. Cheltenham: Edward Elgar.
Shoss, M. & Ciarlante, K. (Summer 2022). Are Robots/AI Viewed as More of a Workforce Threat in Unequal Societies? Evidence from the Eurobarometer Survey. *Technology, Mind, and Behavior*. Retrieved from https://tmb.apaopen.org/pub/rv1x9zq4/release/2#:~:text=Utilizing%20the%20Eurobarometer%2087.1%20data,threats%20of%20general%20job%20loss
Simmons, R. (2020). Terry in the Age of Automated Police Officers. *Seton Hall Law Review, 50*, 909–953.
Skorup, B. (2022). Drones, Airspace Design, and Aerial Law in States and Cities. *Akron Law Review, 55*, 157–186.
Smith, B. (2020). *New Technologies and Old Treaties*. University of South Carolina School of Law: Faculty Publications.
Stross, C. (2018, January 2). *Dude, You Broke the Future!* Charlie's Diary. Retrieved from https://www.antipope.org/charlie/blog-static/2018/01/dude-you-broke-the-future.html
Templeton, B. (2020, September 21). *What Happens To Car Insurance Rates After Self-Driving Cars?* Forbes. Retrieved from https://www.forbes.com/sites/bradtempleton/2020/09/21/what-happens-to-car-insurance-rates-after-self-driving-cars/?sh=2c9c4e9a5b97
The Dawn Project: In Scientific Test, Tesla "Full Self-Driving Technology Consistently Strikes Child-Sized Mannequins (July 13, 2022). Dawn Project. Retrieved from https://rai2022.umlaw.net/wp-content/uploads/2022/08/The_Dawn_Project___Tesla_FSD_Test__8_.pdf
Thomasen, K. (2018). Beyond Airspace Safety: A Feminist Perspective on Drone Privacy Regulation. *Canadian Journal of Law and Technology, 16*, 308–338.
Thomasen, K. (2020). Robots, Regulation, and the Changing Nature of Public Space. *Ottowa Law Review, 51*, 275–312.
Thomson, J. (1976). Killing, Letting Die, and the Trolley Problem. *The Monist, 59*, 204–217.
Twerski, A. & Henderson, J. (2009). Manufacturers' Liability for Defective Product Designs: The Triumph of Risk-Utility. *Brooklyn Law Review, 74*, 1062–1108.
Wade, J.W. (1974). On the nature of strict tort liability for products. *Insurance Law Journal, 1974*(3), 141–162.
Wu, P., Escontrela, A. & Hafner, D. (2022, June). *DayDreamer: World Models for Physical Robot Learning*. Cornell University: arXiv. Retrieved from https://arxiv.org/pdf/2206.14176.pdf
Wessling, B. (2022, June 24). Cruise Hits Milestone By Charging For Robotaxis Rides in SF. *The Robot Report*. Retrieved from https://www.therobotreport.com/cruise-begins-charging-the-public-for-robotaxis-rides/
Woo, J., Whittington, J. & Arkin, R. (2020). Urban Robotics: Achieving Autonomy in Design and Regulation of Robots and Cities. *Connecticut Law Review, 52*, 324–410.
Zaveri, M. (2021, May 11). N.Y.P.D. Robot Dog's Run Is Cut Short After Fierce Backlash. *The New York Times*. Retrieved from https://www.nytimes.com/2021/04/28/nyregion/nypd-robot-dog-backlash.html