

Applied Professional Ethics for the Reluctant Robotist

John P. Sullins
Sonoma State University
Department of Philosophy
1801 East Cotati Avenue
Rohnert Park CA
01-707-664-2277
John.sullins@sonoma.edu

ABSTRACT

Since robots are designed to interact with humans, robotics applications always have ethical implications. Until recent decades, most robotics applications were of an industrial nature and therefore far removed from close interactions with people for safety reasons. In that milieu the ethical impacts of robots on human relations were more subtle and could be mostly ignored by robotists who might be reluctant to dive into the world of applied ethics. This situation has rapidly changed and with robots poised to enter homes, roads, hospitals, and battlefields worldwide, ethics has now become something that the designers of robots must take into careful consideration at some point during each project. In this paper we will see that robot ethics and computer ethics are related but not exactly the same. The difference between ethical and legal considerations will be addressed; we will introduce the basic ethical theories that are of value for considering in the design of robots, and conclude with a look at some of the main open challenges for the design of ethical robots will be listed. By the end of the paper we will have shown that ethics is not something to be avoided by robotists but it can actually present intriguing problems for researchers to solve and it also provides an opportunity for robotists to actively shape the future of human robot ethical relations.

Categories and Subject Descriptors

Social and professional topics: Professional topics– *Computing profession, Codes of ethics.*

General Terms

Human Factors, Theory, Legal Aspects, Verification.

Keywords

Robot ethics, Applied Ethics,

1. INTRODUCTION

Robotics is a massively multidisciplinary endeavor. Robotics projects have to bridge hardware and software design, which includes at the very least mechanical, electrical and systems engineering. Added to this list are the various subfields of computer science such as computer systems, networks, applied computing, security, to name just a few. When we are working

with biologically inspired robotic designs, then we need to add to this might various aspects of biology and/or applied chemistry. In the world of social robotics, skills in human behavior, sociology, and anthropology have been needed. This reality requires that robotists be the kinds of persons that are adept at thinking beyond traditional disciplinary boundaries. It is safe to say that those who do innovative work in robotics are already adept at working between paradigms and modes of thought.

Over the past few decades robotists have started to deliver the fruits of this adventurous and somewhat ad hoc style of problem solving with an explosion of robotics applications both big and small, as well as both industrial and commercial. This success has resulted in robot applications that are finding their way into even our personal and family lives. This very success is necessitating that robotists must acquire an even greater array of multidisciplinary skills in order to do their job well. Some of those skills include legal reasoning and still others require ethical reasoning. It is the former that we will be focusing on in this paper. It is argued here that even though robotists may not now see themselves as experts in the field of ethics, this natural reluctance must be overcome in order to solve some of the problems they will face in the application of the autonomous machines they are trying to create.

As an example let's look at the evolution of mobile robots. In the sixties, seventies, and eighties, robotists had enough to do just to get a robot to move across a lab in an efficient and accurate manner. A research robot like Shakey¹ had only distant moral and ethical implications. Some of the people working on this machine did engage in thinking about those implications, for instance see [1], [2], [3], but these thoughts were not necessary to advance the understanding of mobile robots. Fast forward to today and this line of research is breaking out of the labs and robotic cars are already logging many miles on public roads. At this point along with the various technical problems left to solve, there have also come large legal and ethical problems.

In many ways these more social problems are now more likely to keep robot cars from advancing more than the technical issues. Because of this, the research groups working on advancing this technology need to spend some time thinking about the ethics and legal problems that arise as unintended consequences of deploying these machines. It is not enough to make some vague claims about robot cars being safer than human driven cars and that they are necessary to save lives [4], it turns out that the issue is much more complex than it appears on the surface [5], [6], and

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage, and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s). Copyright is held by the author/owner(s).

The Emerging Policy and Ethics of Human Robot Interaction, Mar 2, 2015, Portland, OR, USA.

¹ Here is a nice history of Shakey the Robot:
<http://www.sri.com/work/timeline-innovation/timeline.php?timeline=computing-digital#!&innovation=shakey-the-robot>

talk of safety or the superiority of robots to human drivers is not enough to convince a skeptical public. This is not limited to robot cars but also to the use of drones for delivery in civil airspace where we have seen the ambitious plans for this technology grounded by the FAA [7]. We are likely to see this trend continue unless the robotics community can get ahead of these legal and ethical questions and have good counter arguments to rise when they get the inevitable pushback from the media and policy makers.

In this paper we will see that robot ethics and computer ethics are related but not exactly the same. The difference between ethical and legal considerations will be addressed; we will introduce the basic ethical theories that are of value for considering in the design of robots, and conclude with a look at some of the main open challenges for the design of ethical robots will be listed. By the end of the paper we will have shown that ethics is not something to be avoided by roboticists but it can actually present intriguing problems for researchers to solve and it also provides an opportunity for roboticists to actively shape the future of human robot ethical relations.

2. Ethics vs. Compliance

An important conceptual difficulty must be addressed before we can get much further in this discussion. It can be tempting for those reluctant to get involved in discussing the ethical impacts of ones research to just maintain that they already comply to the law and have legal counsel working on identifying the laws and contract obligations they may have. This is thought to absolve them of having to worry about ethics, since the law is designed to make sure their products and the use of their products are done in an ethical and socially acceptable way.

It is vitally important to remember that the domains of law and ethics should ideally overlap, but in reality they do not. We know that historically it might have been legal at certain times to own slaves but it was never a moral or ethical situation. More close to home, it is often the case that laws are written to benefit one group more than another in ethically dubious ways. And directly on point is the case that new products in information technology and robotics enter an initially unregulated marketplace in what the philosopher Jim Moor calls a “policy vacuum” where no one is quite sure how the new product should be used and what legal limits can be place on them [8]. Lawrence Lessig makes a similar point when he argues that the rate of change in information technology can easily outstrip the ability of law to set policies for the use of new technologies leaving computer programmers as de facto policy makers who change and disrupt society beyond effective legal control [9].

The literature on law and robotics is just now ramping up, and one nice example is the work of Bryant Walker Smith [10]. Walker Smith reminds us that when it comes to the use of a robotic system we have to remember that there will be the intended use of the designer and this may come into conflict with the actual use of the product along with demands of legal use and reasonable use [10]. Everything is fine if all of these kinds of use happen to co-occur, but look how this is playing out in the use of personal aerial drones. This technology is quite a fad right now and gives users unparalleled abilities to take video and photos from stunning angles. But as recent events such as the drone crash on the White House lawn [11], or the drone hooligans of Paris [12] show, actual use of robots can be well outside of intended, legal and reasonable use. But even if we can somehow assure that one’s

particular robotic system is used only in a legal, reasonable, and intended way, there is still a potential problem and that is the use may still fall outside of what is ethical or moral. This would be something like what the roboticist Illah Nourbakhsh calls “robot smog” [13]. This smog results from systems that are all operating as intended and in a legal manner but nonetheless result in uncomfortable or immoral situations. For instance, going to a public beach where a dozen or more kites are flying might actually enhance the experience, but going to a public beach with a dozen or more multirotor drones buzzing overhead is not a pleasant experience.

There is also a more pragmatic reason to add a concern to ethics and morality to the design process of sociable robotics. Customers and other users will expect that the machines they take into their homes and lives will have basic competencies in human robot interactions that involve ethically sensitive situations. Failure to provide this will result in product recalls and loss of customer confidence in your product.

We can now see that beyond compliance to the law we have to add another area of use, and that is ethical and moral use. And along with the design of intended use that roboticists already engage in, they must factor in ethical design practices as well. Let’s now try to move forward with those issues.

3. Ethical Theories Relevant to Robotics

Moral and ethical philosophy is a vast field of study that specialists devote their lives to. It is understandable that roboticists would be reluctant to add this vast field of study to the long list of topics they must already master to produce worthy research in their own field. Luckily earning an advanced degree in philosophical ethics is not required in order to do good values centered design. Just as the roboticist does not have to become a lawyer to take into account the legal implications of her designs, she also only needs to only have sensitivity to the ethical implications in her work. If deep problems arise, there are professionals to consult, just as there is in law.

Another issue that can be worrisome is that anyone who has read even a little in the field of philosophical ethics will be aware of the fact that no single theory has rose to the top of the field. While the field of ethical and moral philosophy is not settled, and in fact may be impossible to settle, there are still some useful ideas in the field that can be applied to help roboticists think through the ethical implications of their projects. In the same way that mathematics does not have to be a settled field of study in order for roboticists to apply certain ideas from the field in their programming and engineering, neither does ethics. It is best to think of the various theories found in philosophical ethics as potential tools that may or may not be appropriate for a given robotics application. Roboticists are free to ignore the doctrinal battles and turf wars that might exist for professional philosophers and just go straight to the good ideas no matter what school of thought they may come from.

It is also the case that most roboticists will have encountered many of the best theories in ethics at some point during their education or upbringing, so this material will not be entirely unknown to them. In the last few millennia there has been quite a bit of work written about ethics and morality. Some of it is more useful than others when it comes to specifically applying these ideas to projects in robotics. In this section, those theories that have specific applicability will be discussed.

Socrates is famous for saying “the unexamined life is not worth living,” and I would like to extend that idea by saying that the unexamined technology is not worth building. By this I mean that it is better for a technology to not be built, if its designers have not spent any time thinking about why they are building it in the first place and what this technology stands to add or detract from society.

3.1 A Brief look at Ethical Theories of Note

This is a quick list of the ethical theories that have been shown to be relevant to the design of information technology and robotics as well as many other technologies.

3.1.1 Consequentialist or Utilitarian Ethics

One of the most famous and in some cases infamous codes of ethics has been those based on judging the ethics of an action based on the results of an action. Under this theory, an action is ethical if the consequences of that action are good. How you define “good” matters a lot. One school of thought is Utilitarianism which argues that measurable happiness or beneficial outcomes is what matters and we should work to ensure that the greatest number of people enjoy the highest beneficial outcomes possible. For a more complete explanation of this theory see [14]. An example of how Utilitarianism might be employed in the design of a robot can be seen here [15], where the author describes a machine called the Utilibot that is designed with:

The core ethical decision-making capacity of the Utilibot consists of two dynamic Bayesian networks that model human and environmental health, a dynamic decision network that accounts for decisions and utilities, and a Markov decision process (MDP) that decomposes the planning problem to solve for the optimal course of action to maximize human safety and well-being [14].

3.1.2 Deontological or Moral and Ethical Duties

It doesn't take long to see that if we push hard on consequentialism we can imagine situations where perhaps the majority of people can be made happy or a beneficial result can occur but at the cost of the happiness or interests of some minority. A deontologist would be unwilling to allow this to occur.

For them we have certain moral duties, rights and responsibilities that must be respected regardless of the results, certain things are right or wrong intrinsically and this cannot be changed by outside concerns or expediencies. This line of reasoning is most famously championed by the philosopher Kant.

There is nothing simple about his philosophy but one key thought is that an ethical agent should, “act only in accordance with that maxim through which you can at the same time will that it become a universal law” [16]. This means that the moral agent is bound to a universal law that they can access through the power of reason. A sketch of how Kant's ethics might be applied to a robot can be found in the work of Thomas Powers [17].

The Three Laws of Robotics as imagined by the great science fiction writer Isaac Asimov might also fit in our discussion here since they are designed to be unquestionably obeyed.² There is a

growing literature on this from both philosophers and roboticists. This is probably due to the fact that it is the first known attempt to codify an ethical system for robots even if it occurred in works of fiction. Those works also inspired many roboticists to enter the field so it is understandably one of the first things thought of when it comes to robot ethics. But it is also the case that that these laws have some very dangerous conceptual flaws as evidenced by the dramatic events in the stories where they are used as a plot device. They also require a machine that has quite a complex ability to understand broad concepts like “harm,” which can be hard enough for a human mind to fully reason through. A taste of this debate can be found here [18], [19], [20].

3.1.3 Virtue Ethics

Virtue Ethics has roots that go back to the work of Aristotle and is one of the three most influential ethical systems today most notably in the rise of the feminist ethics of care [21]. This system focuses on the individual character of the moral agent through the development of beneficial traits and habits of interaction. A famous example is given by Aristotle who noticed that real courage was to be found between the vices of cowardliness and foolhardiness. Even Aristotle did not give a complete list of virtues and vices and different ideas on what counts as a virtues or a vice have changed and evolved over time and in various cultures. The main benefit of this way of thinking about ethics is that it avoids some of the problems we have already seen with consequentialism and deontology where different duties or goals come into conflict and resolves them by saying an act or decision is ethical if that act or decision was brought about by a virtuous person. So we no longer focus on the specific act or goal but on the agent bringing them about.

There are two main ways that this ethical theory can be used in robotics research. One is in the development of the character of the researchers themselves. As we have already seen above, most of what is being done in high technology research occurs in a policy vacuum where there is little outside legal, let alone ethical guidance available. In these cases we can only rely on the moral character of the researchers themselves so it is in everyone's best interest to develop that sensibility during the education and mentoring of these researchers. An example of what that kind of training might look like can be found here [22].

The second location virtue ethics can inform design in robotics is found in the machines themselves. We have to be very careful here in that it is very easy to anthropomorphize robots and assume levels of character they do not have. Movies are full of plucky little virtuous robots that are always there to give a helping hand right when the hero needs them. It would be wonderful if reality could get close to that vision but we are a long ways away from that. Instead what I am talking about here is that just as a telerobot might need human operator to take over when in needs help with a tricky bit of navigation, so too might machines have to be teleoperated by humans when they are attempting to navigate tricky human relations. This could also be partially automated by allowing the machine to learn from the actions of human exemplars. Certainly much of etiquette and custom could be approximated that way. Much can also be done in making sure the machines at least operate as if they were a virtuous agent. For instance we want robots in the home and at work to behave as if

² It could be argued that the Three Laws of Robotics are more properly thought of as an example of Rule Utilitarianism since the laws are only justified on the fact that they are good for the

human owners of the robots. But this is philosophical hairsplitting and I promised I would refrain from that.

they were a friendly agent with your best interests in mind [23]. At the very least, we can define some low level mechanistic virtues such as data security, integrity, and accessibility that we want in robots as well as in any information technology [24]. Others have argued that robots should be programmed with a kind of “slave” virtue give that they do not have human rights or moral status [25]. Thus, at whatever level a particular robot system might be located at, there are ways to get some value out of applying virtue ethics to the design of the system.

3.1.4 Fairness and Social Justice

This theory grows largely out of Anglo-American social and political philosophy. It argues that all human beings deserve to be treated equally or if they are treated unequally there must be a firm justification for that mistreatment. Here we might argue that all humans have certain inalienable rights that must be honored by any legal apparatus and by fellow citizens. For instance all humans have a right to freedom that they cannot have taken from them, nor that they themselves can barter away. However, a person’s infringement of the rights of others may temporarily cause them to lose their freedom by being placed in prison or jail and this is ethical as long as the system is fair and justified.

Roboticians need to consider this theory in the cases where their proposed systems might be unfairly impacting some set of human or legal rights. For instance we can imagine many robotics applications that might have a negative impact on the right to privacy [26], and this could only be ethical if that infringement was both fair and justified. Allowing an aerial drone to fly over a public square might be an unnecessary infringement of privacy most of the time but if that flight is part of a rescue operation, then it is both fair and justified to infringe on the privacy rights of the people in the square.

3.1.5 Common Goods

This is a more community based ethical theory. Many ethical theories base their reasoning solely on the individual moral agent. The idea is that even if the agent was alone in the universe, she through reason alone, could conceive of the ethical theory in question. This theory is more communitarian in that it preconceives of the existence of the community and places various ethical constraints on the individual based on this fact. So here the ethical duty is at the society level towards the individual, and demands that the common welfare and economic wellbeing of every individual is maintained to a certain standard. We can see that this theory has particular relevance to many European and Asian countries where social goods are distributed through various political means.

Roboticians working in or with these societies will need to pay attention that their systems fit into this way of living.

3.1.6 Religious Ethics

Religion plays an undeniable role in the development of thoughts and attitudes towards ethics and morality in every human society. Most of the ethical beliefs of various researchers in robotics will be based on or in reaction to their upbringing within a religious tradition or having been in close proximity to those who have. Since most of this research is done in pluralistic societies, bringing one particular religious tradition to the programming of a robot could lead to disagreements and trouble. Since most religious ethics that matter can be justified using the secular theories listed above, it is likely that we can make the most

number of users comfortable without explicit appeal to any one religious tradition.

It should be noted that both Buddhist ethics [27], and Christian divine command theory [28], have been proposed as viable avenues for solving issues in robot ethics.

3.1.7 Information Ethics

Information ethics is a very recent addition to the corpus of ethical study. Here we try to solve some of the thorny ethical paradoxes that occur in other theories by focusing on the relationship between the creation, organization, dissemination, and use of information at an abstract level and at the practical level of what is the best policies and procedures and ethical codes for governing the use of information. Luciano Floridi is perhaps the most prolific writer in this field of study, [29], [30],[31]. Floridi describes information ethics as an ethics that takes into account the ethics of information as a resource, a product, an environment as well as a metaethical framework [30], [32]. I might want to add to that that looking at information in this way can allow us to see certain robots as informational agents and by extension moral agents [33]. Illah Nourbakash defines a robot as the interface of the digital world with the physical world [13], and if we look at it that way we can see that robots are helping us more effectively interact with the cyber world and the distinction between reality and information is becoming less important. This means that information ethics is one of the most important ethical theories for a robotician to become aware of and it is having a very large impact on our moral values as well [34].

3.1.8 Hybrid Approaches

Even in this brief look at some of the highlights of ethical theory it should be very apparent by now that one theory does not stand out as more effective at application than any other. Because of this fact, we have to be a bit pragmatic about all of this and the robotician should see these theories as various potential tools that can be used to do necessary work in their designs. Just as no one tool is the best for building robots, no one ethical tool is appropriate to solve all ethical problems that result in the design of a robot. The best designs will find ways of building a hybrid approach to ethics within their design. Wendell Wallach and Colin Allen argue persuasively for this point [35]. They claim that certain ethical theories are “top-down,” in that they are highly rational and rule based, whereas others are “bottom-up” growing and evolving from the specifics of a situation [35]. Any working artificial moral agent like a robot programmed to behave ethically, will require designs that are sensitive to both approaches and the machine will appeal to both top-down and bottom up ethics in its operation [35].

3.2 Institutional Review Boards and Ethics Officers for Robotics

Many roboticians are familiar with institutional review boards (IRBs) which help give them some ethical guidance at least when it comes to issues of experimentation where human subjects are used. This is a good thing to have but it is by no means universal. Not all university systems have a robust system with experts that know how to apply ethics to robotics. Many times the IRB has a large caseload and can only give cursory attention to a project, thus missing subtle ethical concerns. IRBs are not common outside academic research though some companies have ethics officers or managers. But these entities are often seen in private industry as something that might slowdown or discourage

research so they are something to be avoided or only tolerated if they are required by law to have them [36], [41],[42]. I would like to argue that it is possible to design ethical review systems that are beneficial to the design process and not an unnecessary encumbrance. The best way forward here is not to have only an external review board, but we should also have ethicists involved in the design team. This way ethics is not seen as imposed from outside the team, but as an integral part of the design. The external review must still be in place to make sure the imbedded ethicists work is appropriate and that they are being properly utilized by the design team. What we need to avoid is the feeling that ethics is only about proscribing correct behavior like some kind of morality police, but instead that they help the team solve problems in HRI in new and effective ways. Robotics is an interdisciplinary endeavor, as the systems become more robust, one of those disciplines is ethics and this expertise must be added to the design process at multiple levels.

3.3 Codes of Ethics for Robotics

One important level where ethics can enter into robotics research is through written codes of ethics. Professional codes of ethics seem to have begun with medicine and the Hippocratic Oath but by now nearly every profession has a set of fairly detailed codes that prescribe how members of their profession should behave under certain ethical constraints. Engineering and Computer science are vital to robotics and both of them have well thought out codes of ethics; ABET: Accreditation Board for Engineering and Technology code of ethics for engineers [37], and the Association of Computing Machinery (ACM) Code of Ethics and Professional Conduct [38]. These cover a lot of ground but are silent on some of the pressing ethical issues that occur in robotics but not in the other two professions. Robotics as a discipline has yet to formally accept a code of ethics but there have been some proposals such as [39], which do a better job giving guidance to roboticists in working through the ethical implications of their work including considerations for Human dignity, design ethics, legal ethics, and social duties [39].

One problem with codes of ethics is that they must be updated, discussed and taught. They do no good if they are just stuck to a wall in a lab where no one looks at them. The good news is that there is some evidence that shows that actively using codes of ethics, case studies and review can be worked into mentoring and other instruction and that this does increase ethical outcomes in research processes [40].

4. Open Ethical Challenges in Robotics

There are certain ethical challenges that are somewhat unique to robotics. The following section is not meant to be completely exhaustive but it is included in order to get the reader thinking about these implications and draw connections to their own work. I have organized these challenges into the categories; near term, on the horizon, and future. In each one of those categories we have; personal and home robotics, military, and commercial industrial. I have written about this issue before [43], but this is a good time to update those thoughts and bring in ideas from others [39], [42].

4.1 Near Term Ethical Challenges

As these technologies become more ubiquitous, ethical challenges will grow from novel curiosity to a crisis level if not properly dealt with.

4.1.1 Personal and Home Robotics

Personal drones are the first really popular personal robotics technology. Used mostly for taking photographs from the cameras mounted on the drones, these machines are unfortunately noisy and privacy can be a concern. As more people use them in an unregulated or semi regulated way, we get an example of robot smog [13].

Robotic toys continue to grow in popularity. While most are not that robust in their capabilities, there are some potential ethical issues. Privacy is a concern in any robot that keeps data on its user. But most important is affective manipulation, which occurs when a toy uses faux affective attachment to motivate its own user. In this case it becomes hard to tell if the user is playing with the toy or if it is the other way around. Affective attachments can be formed with very little in the way of computational power since much of it can be bootstrapped just on the form the robotic doll takes [23].

4.1.2 Military robotics

The last few decades has seen a considerable growth in robotics applications and this category contains the most pressing ethical concerns. There is no doubt of the efficiency and usefulness of robotic weaponry but there remains serious ethical concerns about how these systems might color our ability to reason ethically when it comes to decisions to use lethal force.

These weapons might also decrease the friction against going to war in that they lower the cost both economically and in blood that is spent by the aggressor.

Privacy and environmental concerns are important and occur when people have to live for many years under the constant fear of attack from aerial robot systems

While the majority of these machines are teleoperated, autonomy is increasing and it is more a matter of ethics that is holding this development back than the technology itself.

There has been reaction to these machines that has resulted in a movement to ban killer robots with the guiding principle that we should have the human right that a decision to kill a human should not be made by a machine [44].

Every robotics researcher that enters into military robotics research needs to make a through ethical accounting of their decision to take money and resources for the development of systems the use of which they might not entirely know or perhaps support.

4.1.3 Commercial and Industrial

Medical robotics brings up the majority of the concerns in this area followed closely by robotic cars [39], [42], [43].

Surgical robots are a reality. Most of the systems people think about are teleoperated but some of them, such as bone milling machines or image taking machines are more autonomous. There are a number of easy to understand ethical issues here such as a patient's right to know the risks involved and for both patients and doctors to not assume too much in the skills of the robot. But there are other harder to see issues like the inability to export robotic surgery to the developing world and the lack of development of traditional surgery skills that are more easily applied in developing nations among others [45], [46].

Autonomous robotic cars are still in the experimental stage but there might start to be more applications of this technology taking to the roads soon. This is an interesting case where the

technology is being marketed as a way to stop the mass carnage that exists in the amount of automobile deaths that happen each year. This unmitigated good has begun to be challenged and ethicists like Patrick Lin are getting involved with research groups and asking them to think through the nuances of the robotic car proposal [5], [6]. Autonomous cars will be faced with no win situations and their choice of action will have to be programed in beforehand.

Amazon has famously proposed using drone technology for home deliveries but this has been put on hold at least in the US [7], but some sort of drone delivery seems inevitable and this will require that the FAA will need to spend some time regulating the airspace under four hundred feet where these machines tend to fly. More obviously ethical uses of this technology exist as drones could be used as an inexpensive way to deliver medications needed by remote communities in locations where traditional air travel is difficult or impossible.

4.2 Ethical Concerns on the Horizon

These are technologies that have had proof of concept in the lab and are likely to come to market at some point in the next decade.

4.2.1 Personal and Home Robotics

We can expect many more aerial drone applications and a much more crowded urban and even interior airspace exacerbating all the issues raised above. Just as aerial drones can track their user now, personal assistant robots might start appearing in shopping malls and stores. Again the primary ethical issues are privacy and the safety of others as machines interact with not only persons other than their users, but the other robots that also interact in this space. Since there is unlikely to be standardization, these interactions could lead to misunderstandings and accidents.

Toys will gain in capability, with our advanced robots of today becoming child's toys of tomorrow such as kiddie autonomous riding vehicles. Dolls will likely continue to use our affections against us, perhaps tricking us to make in app purchases against our better judgment.

Growing out of the toy market, we could see more capable sex robots arrive on the market, the ultimate in HRI. Right now these are a laughable curiosity but a good deal of product research is going into this technology and it has been promised for a number of years now [47], when it takes off, there will be many strange ethical issues to resolve [48].

Personal autonomous cars might become available, but certainly our cars will be assisting us in ever increasing ways making them at least semi-autonomous at the very least. What will the roads be like when there are all these various levels of vehicles on the road, from fully driver operated, thorough semi-autonomous to fully autonomous.

4.2.2 Military Robots

Military necessity and competition between nations will make autonomy necessary. Unless there is some arms control placed on these technologies we are very likely to see a rapid arms race as each country tries to make more and more capable systems.

Robotic prosthesis technology could make losing a limb hardly a reason to end a soldier's deployment. This will start to blur the distinction between human and machine in ethically challenging ways.

Assistive robots will not be limited to shopping or drone delivery at home; these technologies will be brought to the battlefield as well perhaps finally delivering on the promise of allowing the military to deploy less troops who can nonetheless do more given all this assistive technology.

4.2.3 Commercial and Industrial

In the medical area care and medical assistant robots will begin to be deployed at the rate we see robot surgery systems entering the hospital today. The main ethical concern here is that we do not lose the human touch and human care that is emblematic of good healthcare and replace it with cheaper robotic systems.

Many jobs will disappear as they become automated given that even small companies will be able to afford light industrial robots. This may just mean that other jobs open up and new opportunities for employment might emerge, but it is certain that many of today's jobs will disappear and we need to prepare now for ways to ethically transition this workforce.

Most of our economic interactions with companies and services will be through automated conversations with various robotic systems. This will be a big change in the social contract that will stretch our already overburdened traditional ethical systems since it will be hard to know when your are dealing with a human or a machine, humans may start to converse as if they were a machine, and we will have to negotiate when it is right to use a robot agent to do our business and when our full human attention is needed [39].

4.3 Future Robotic Ethical Concerns

It is hard to say anything here that might not seem wrong or even silly upon later reading. So I will dispense with any predictions and instead look at the big ethical issues looming over the horizon.

4.3.1 Moral praise and blame

All of the traditional ethical systems presuppose that the only agent worthy of ethical praise or blame is other human beings. Recently, environmental ethics has extended our concern to environments and animals but as we build more and more robust robotics systems, eventually we will have to consider the best of them as something like a moral agent [33].

4.3.2 Artificial Moral Agency

Once we have made that conceptual leap, we have to learn how to create machines that will reason ethically and hopefully treat us in a just way.

4.3.3 Robotic Environmentalism

We will have to make sure these systems work well with our environments. Right now they are vastly different from the biological world. We did a terrible job on meshing our technology with the environment last century; we will not survive another century if we get this wrong.

4.3.4 Finding a New Humanism

Our race for automation in every last aspect of our lives will make a world that is easier to survive in, but we also must make sure it is a world worth living in. To make that happen we will have to closely examine each and every step we plan to take and chose only those that will take us to a better, more meaningful existence as humans.

5. Exciting Challenges

When we think about the future, it can bring thoughts of trepidations but it is fully possible to mitigate those fears. Instead we should look with excitement at the possibilities and have a sober plan on how to achieve them. Robot ethics is not an attempt to throttle innovation in HRI. Robot ethics is a way to focus our efforts in HRI in ways that will make real differences in people's lives through a focused series of thought experiments and reasoned applications of useful theories from the vast territory of ethical theories. We are reengineering these ideas and creating a technology of ethics with the ultimate goal of creating machines that are capable of engaging in philosophical thought with us.

We have a lot to learn, reluctance must be put aside and the investigation of the ethical commitments inherent in HRI begun with enthusiasm.

6. ACKNOWLEDGMENTS

I would like to thank the organizers of The Emerging Policy and Ethics of Human-Robot Interaction workshop at HRI 2015, for inviting me to present this work.

7. REFERENCES

- [1] Moravec, Hans. 1990. *Mind Children: The Future of Robot and Human Intelligence*. Harvard University Press, Cambridge, MA.
- [2] Moravec, Hans. 2000. *Robot: Mere Machine to Transcendent Mind*. Oxford University Press, Oxford, UK.
- [3] Brooks, Rodney. 2002. *Flesh and Machines: How Robots Will Change Us*. Pantheon Books, New York, NY.
- [4] Simonite, Tom. 2013. Data Shows Google's Robot Cars are Smoother, Safer Drivers Than You or I. *Technology Review*. October 25. DOI=<http://www.technologyreview.com/news/520746/data-shows-googles-robot-cars-are-smoother-safer-drivers-than-you-or-i/>
- [5] Lin, Patrick. 2013. The Ethics of Saving Lives with Autonomous Cars Is Far Murkier Than You Think. *Wired*. July 30. DOI= <http://www.wired.com/2013/07/the-surprising-ethics-of-robot-cars/>
- [6] Lin Patrick. 2014. Here's a Terrible Idea: Robot Cars With Adjustable Ethics Settings. *Wired*. August 18. DOI=<http://www.wired.com/2014/08/heres-a-terrible-idea-robot-cars-with-adjustable-ethics-settings/>
- [7] Morris, Chris. 2015. FAA Plans Ground Amazon's Drone Program. *Fortune*. February 16. DOI=<http://fortune.com/2015/02/16/faa-amazon-drones/>
- [8] Moor, James H. 1985. What is Computer Ethics?. *Metaphilosophy*, Volume 16 issue 4 (30 September), Pgs. 266-275.
- [9] Lessig, Lawrence. 1999. *Code and Other Laws of Cyberspace*. Basic Books, New York.
- [10] Bryant Walker Smith, Forthcoming 2015. Lawyers and Engineers Should Speak the Same Robot Language, in *Robot Law*, DOI=<http://newlypossible.org>
- [11] Acosta, Jim and Dimond, Jeremy. 2015. U.S. Intel Worker Blamed for White House Drone Crash. *CNN*, January 27. DOI=<http://www.cnn.com/2015/01/26/politics/white-house-device-secret-service/index.html>
- [12] BBC. 2015. Paris Night Drone Mystery Deepens. February 25. DOI=<http://www.bbc.com/news/world-europe-31619099>
- [13] Nourbakhsh, Illah. 2013. *Robot Futures*, MIT Press, Cambridge MA.
- [14] Sinnott-Armstrong, Walter, "Consequentialism", The Stanford Encyclopedia of Philosophy (Spring 2014 Edition), Edward N. Zalta (ed.), DOI = <http://plato.stanford.edu/archives/spr2014/entries/consequentialism/>
- [15] Cloos, Christopher. 2005. The Utilibot Project: An Autonomous Mobile Robot Based on Utilitarianism. *AAAI Fall Symposia*. DOI=<http://www.aaai.org/Papers/Symposia/Fall/2005/FS-05-06/FS05-06-006.pdf>
- [16] Johnson, Robert, "Kant's Moral Philosophy", The Stanford Encyclopedia of Philosophy (Summer 2014 Edition), Edward N. Zalta (ed.), DOI = <http://plato.stanford.edu/archives/sum2014/entries/kant-moral/>
- [17] Powers, Thomas M., 2011. Prospects for a Kantian Machine. In *Machine Ethics*, Anderson and Anderson (eds.), pgs. 464-294.
- [18] Clarke, Roger, 2011. Asimov's Laws of Robotics: Implications for Information technology. In *Machine Ethics*, Anderson and Anderson (eds.), pgs. 254-285.
- [19] Anderson, Susan Leigh, The Unacceptability of Asimov's Three Laws of Robotics as a Basis for Machine Ethics. In *Machine Ethics*, Anderson and Anderson (eds.), pgs. 285-297.
- [20] Hall, J. Storrs, 2007. *Beyond AI: Creating the Conscience of the Machine*. Prometheus Books.
- [21] Hursthouse, Rosalind, "Virtue Ethics", The Stanford Encyclopedia of Philosophy (Fall 2013 Edition), Edward N. Zalta (ed.), DOI = <http://plato.stanford.edu/archives/fall2013/entries/ethics-virtue/>
- [22] Sullins, John P. 2014. A Case Study in Malware Research Ethics Education: When teaching bad is good. *Cyber-Security Research Ethics Dialog & Strategy (CREDS II)*, May 17, San Jose, CA. DOI=http://www.cs.sonoma.edu/ledin/malware/pdf/Sullins_Creds2014.pdf
- [23] Sullins, John P. 2008. Friends by Design: A Dersign Philosophy for Personal Robotics Technology. *Philosophy and Design: From Engineering to Architecture*. Springer, pgs 143-157. DOI=http://link.springer.com/chapter/10.1007/978-1-4020-6591-0_11#
- [24] Sullins, John P. 2014. Deception and Virtue in Robotics and Cyberwarfare. *The Ethics of Information Warfare: Law, Governance and Technology Series*, Volume 14, Pgs. 187-201. DOI=http://link.springer.com/chapter/10.1007/978-3-319-04135-3_12
- [25] Petersen, Stephen, 2006. The Ethics of Robot Servitude. DOI=<http://stevepetersen.net/professional/petersen-robot-servitude.pdf>

- [26] Calo, M. Ryan, 2012. Robots and Privacy. In *Robot Ethics*, Lin, Abney and Bekey (eds.), pgs. 187-203.
- [27] Hughes, James, 2012. Compassionate AI and Selfless Robots: A Buddhist Approach. In *Robot Ethics*, Lin, Abney and Bekey (eds.), pgs. 69-85.
- [28] Bringsjord, Selmer and Taylor, Joshua, (2012). The Divine-Command Approach to Robot Ethics. In *Robot Ethics*, Lin, Abney and Bekey (eds.), pgs. 69-85.
- [29] Floridi, Luciano, 2010. *Information: a Very Short Introduction*. Oxford University Press.
- [30] Floridi, Luciano, 2010. *The Philosophy of Information*. Oxford University Press.
- [31] Floridi, Luciano, 2013. *The Ethics of Information*, Oxford University Press.
- [32] Floridi, Luciano, 2014. *The Fourth Revolution*, Oxford University Press.
- [33] Sullins, John P. 2006. When is a Robot a Moral Agent? *International Review of Information Ethics*, 6(1): 24-30. DOI= http://www.i-r-i-e.net/inhalt/006/006_Sullins.pdf
- [34] Sullins, John P. 2014. Information Technology and Moral Values, *The Stanford Encyclopedia of Philosophy*. Edward N. Zalta (ed.). DOI= <http://plato.stanford.edu/archives/spr2014/entries/it-moral-values/>
- [35] Wallach, Wendell and Allen, Colin, 2009. *Moral Machines: Teaching Robots Right from Wrong*. Oxford University Press.
- [36] Clark, Hannah, 2006. Chief Ethics Officers: Who Needs Them? *Forbes*, October 23. DOI= http://www.forbes.com/2006/10/23/leadership-ethics-hp-lead-govern-cx_hc_1023ethics.html
- [37] ABET: Accreditation Board for Engineering and Technology code of ethics. DOI= <https://www.iienet2.org/details.aspx?id=299>
- [38] ACM: Task Force for the Revision of the ACM Code of Ethics and Professional Conduct, ACM Code of Ethics and Professional Conduct. DOI= <https://www.acm.org/about/code-of-ethics>
- [39] Riek, Laurel D. and Howard, Don, 2014. A Code of Ethics for the Human-Robot Interaction Profession. In *Proceedings of We Robot 2014*. DOI= <http://robots.law.miami.edu/2014/wp-content/uploads/2014/03/a-code-of-ethics-for-the-human-robot-interaction-profession-riek-howard.pdf>
- [40] Tolich, Martin, Tumilty, Emma, 2014. Making ethics review a learning institution: The Ethics Application Repository proof of concept – *tear.otago.ac.nz*. *Qualitative Research*, Mar 24, 2001-212.
- [41] Calo, Ryan, 2014. America Needs a Federal Robotics Agency. *Time Magazine*, Sept 22.
- [42] Calo, Ryan, 2014. The Case for a Federal Robotics Commission. *Brookings*, September. DOI= <http://www.brookings.edu/research/reports2/2014/09/case-for-federal-robotics-commission>
- [43] Sullins, John P. 2011. Introduction: Open Questions in Roboethics. *Philosophy of Technology*, Springer-Verlag. 24:233–238.
- [44] Campaign to Stop Killer Robots. DOI= <http://www.stopkillerrobots.org/>
- [45] Sullins, John P. 2014. “Ethical Trust in the Context of Robot-Assisted Surgery,” *APA Newsletter on Philosophy and Computers*, Volume 14, 1, Fall.
- [46] Van Wynsberghe, A., and C. Gastmans. 2008. Telesurgery: An Ethical Appraisal. *Journal of Medical Ethics* 34, no. 10: e22.
- [47] Levy, David. 2008. *Love and Sex with Robots: The Evolution of Human-Robot Relationships*, Harper Perennial.
- [48] Sullins, John P. 2012. Robots, Love, and Sex: The Ethics of Building a Love Machine. *IEEE Transactions on Affective Computing*, Fourth Quarter 2012 (vol. 3 no. 4) Pgs. 398-409.