Liability in Robotics: An International Perspective on Robots as Animals

Richard Kelley[1], Enrique Schaerer[2], Micaela Gomez[1], and Monica Nicolescu[1]

[1] *University of Nevada, 1664 N. Virginia St., Reno, NV 89557, {rkelley | mgomez | monica}@cse.unr.edu*

and

[2] *United States Court of Appeals for the Ninth Circuit, San Francisco, CA 94109, enrique_schaerer@ca9.uscourts.gov*

Abstract

As service robots become increasingly common in society, so too will accidents involving service robots. Current law functions effectively to adjudicate the disputes that arise from such accidents, but as technology improves and robot autonomy grows, it will become much harder to apply currently-existing laws. Instead, new legal frameworks will have to be developed to address questions of liability in human-robot interaction. We have already proposed the framework "Robots As Animals," in which robots are analogized to domesticated animals for legal purposes in disputes about liability. In our initial presentation, though, we focused exclusively on the common law in the United States Federal Government. In this paper, we examine the laws concerning domesticated animals in countries in Europe, Asia, and North America. We apply the lessons learned from our analysis to build an expanded framework that better reflects the established norms of several nations and more explicitly balances the competing interests of producers and consumers of robot technology. We also provide examples of ways in which our new framework may be applied.

keywords: human-robot interaction, roboethics, robot legal studies, liability, robots as animals

1 INTRODUCTION

As robots become increasingly common in daily life, there will be a corresponding growth in the rate of accidents and injuries involving robots. Although the law in many countries is adequate to address the present needs of society, it is unlikely that this state of affairs will continue as technological advances push robot autonomy to ever more impressive levels. With this in mind, there is a growing recognition of the need to develop laws governing many facets of human-robot interaction. In particular, there is a growing need to develop a legal framework that deals with the problem of *liability* in human-robot

interaction: namely, when a robot and a human are involved in an accident, who bears responsibility for the accident?

Any solution to this problem must balance a number of competing concerns. In the first place, the framework must address the fundamental novelty of robots in our society. Given this novelty, it seems unlikely that currently-existing laws can be naïvely extended to cover all the relevant concerns in human-robot interactions. At the same time, it would be ideal if the laws developed to regulate liability in human-robot interactions naturally evolved from presently-existing legal standards; such a natural evolution would allow the new standards for robots to rely on the strength of precedent established in other areas of law over the course of decades or even centuries. Similarly, the framework must strike a balance between the need to protect consumers of robot technology and the need for laws to allow producers of robot technology to innovate. A framework that provides too little protection exposes consumers to unnecessary amounts of risk. However, if the legal frameworks developed to govern human-robot interaction are too restrictive on or unreasonable for robot manufacturers and stifle innovation, then it is unlikely that society will be able to benefit to the greatest extent possible from developments in robotics. Clearly, balancing these and other possible interests will challenge lawmakers and the robotics community, which must have input in the development of the relevant legal frameworks if there is any to be hope for those frameworks to function effectively.

In previous work, we proposed the liability framework "Robots As Animals," which establishes a standard set of principles for assigning liability among the manufacturer, the owner, and the victim of a robot involved in an accident. In essence, this standard says that when a robot is involved in an accident, if it is not defective, then courts should treat the robot as if it were a domesticated animal when assigning liability to the victim and to the owner of the robot involved in an accident. Our initial analysis was based largely on the common law in the United States, in which the obligations of domesticated animal owners are enforced through civil (as opposed to criminal) law. However, a careful analysis of other nations' laws shows that the approach in the United States on the matter of owners' responsibilities for their domesticated animals is hardly universal. In fact, approaches to this issue by European and Asian countries are very different in many important respects. This raises the question: if we wish to develop a standard for liability in human-robot interaction that is widely applicable and consistent with the currently-existing legal norms in many nations, how can we extend Robots As Animals to take account of international variations in the laws governing domesticated animals? In this paper, we propose an answer to this question, extending Robots As Animals based on an analysis of the relevant laws of several nations in Europe and Asia. We contend that the extensions offered here improve the framework substantially, allowing it to reflect well-accepted norms in many nations. Moreover, we also contend that the framework offered here performs the necessary balancing described above, in particular balancing the needs of consumers and producers of robotics technology.

The rest of the paper is structured as follows: first, we discuss some of the previous work done by roboticists on questions of roboethics and robot legal studies. We then review Robots As Animals before moving on to an analysis of the laws of several nations regarding domesticated animals such as dogs.

We then examine how lessons learned from the international analysis can be applied to the framework, and describe the necessary extensions. Before concluding the paper, we provide a number of examples showing how the framework may be applied in real-world situations.

2 RELATED WORK

The potential significance of legal standards for analyzing questions in roboethics has already been recognized [1] [2]. However, few researchers have looked at the social problems of robotics from a strictly legal perspective. In particular, there has been little discussion in the robotics community about specific laws that should be passed to regulate aspects of human-robot interaction.

Similarly, while a number of researchers have considered the question of responsibility in human-robot interaction[3][4], none to date have considered how our understanding how responsibility attribution should be applied to create a practical legal framework. In contrast, the present work seeks not to further our understanding of how people view the issue of responsibility in HRI, but to apply our understanding to create a legal framework that may be applied to real-world situations that will arise in the near future as robotics technology becomes even more pervasive throughout society.

The framework described here was originally proposed in [5]. As mentioned above, the framework as described in that paper was based largely on the laws of the United States. The present work is a direct extension of the original framework, and is designed to refine the framework based on a more expansive analysis of liability in other countries.

3 ROBOTS AS ANIMALS: A REVIEW

Before we extend our framework beyond American law, and before we refine our framework through robot taxonomy, we offer a brief review of the initial Robots as Animals framework. For further discussion of the ethical and legal considerations that underlie this framework, along with specific examples of the frameworks application, we refer the reader to our original work on this topic [5]. In short, this framework requires a strict products liability analysis, followed by a negligence analysis if the criteria of strict products liability do not apply.

3.1 Strict Products Liability for Manufacturers

Strict liability is the starting point in our initial Animals as Robots framework. Strictly liability imposes liability without fault: If strict liability applies, a plaintiff may hold a defendant liable for damages resulting from an accident regardless of the defendants fault. The defendant is held strictly liable because no level of care by the defendant can avoid liability.

Strict liability applies in the Animals as Robots framework as follows. If a robot is defective in manufacture or design, or has an inadequate warning label, a plaintiff injured by that robot may hold the robots manufacturer strictly liable for damages. The criteria for strict liability in this context are

that the robot manufacturer be a commercial supplier, and that the robot be defective at the time of sale. If these criteria are met, the robot manufacturer is liable regardless of the robot owners own negligence in bringing about the plaintiffs injury at the hands of the robot. This provides robot manufacturers with an incentive to avoid defects in the manufacture and design of their robots, and to place adequate warning labels on their robots.

3.2 Negligence for Robot Owners

If the criteria of strict liability are met, the robot manufacturer is held liable and the analysis ends. Otherwise, the Animals as Robots framework calls for a negligence analysis to determine whether the robot owner should be held liable to a plaintiff injured by the owners robot. Negligence imposes liability only if the defendant was at fault.

To establish that a robot owner was negligent within the Robots as Animals framework, a plaintiff must show the robot owner owed the plaintiff a duty of care, the owner breached that duty because the owners conduct fell below that level of care, and this breach caused the plaintiffs damages. We justify this negligence analysis by analogizing robots to domesticated animals, whose owners are as a general rule subject to negligence liability because such animals are generally predictable, rather than wild animals, whose owners are as a general rule subject to strict liability because such animals are generally unpredictable.

The rationale for holding the keepers of wild animals strictly liable is that the behavior of such animals is erratic. According to the popular refrain, You can take the animal out of the wild, but you cant take the wild out of the animal. The law therefore forces those who take on this extraordinary risk to bear the consequences when their wild animals injury people or property. By contrast, domesticated animals are far more predictable, such that the law holds the owners of such animals to a less stringent standard of negligence. The law likewise should hold robot owners to the same standard of negligence because, unlike a wild animal, a robot is programmed to act within a predictable range of behavior; a robot makes certain decisions on its own, albeit in general accordance with a prewritten program. The behavior of a robot therefore does not rise to the level of unpredictability one would expect from a wild animal. Instead, its behavior is more like that of a well-schooled canine, which typically does as he is trainedbut not always.

Thus, assuming robots are expected to interact with people and property to some extent, robot owners should be held liable for negligence with respect to their robots, much like dog owners are held liable for negligence with respect to their dogs. Of course, negligence liability for robots, as for dogs, is a default rule subject to stricter standards of liability depending on the circumstances. For example, the owner of a dangerous attack dog may be held strictly liable for not warning trespassers of the danger posed by the attack dog, or if the attack dog escapes and injures people or property. By the same token, the owner of a dangerous security robot may be held to a more stringent standard of liability. The law may even impose segregation policies on that security robot to limit the robots interaction with people

and property, just as many countries have imposed similar policies on attack dogs, as set out below.

In sum, the initial Animals as Robots framework forces robot manufacturers to bear the cost of defective robots, providing an incentive for manufacturers to produce non-defective robots without destroying their incentive to enter the robotics industry. For non-defective robots, the framework holds robot owners liable for their own negligence rather than absolving them of all liability, providing them with an incentive to exercise care in their use of robots. The framework therefore strikes a proper balance between robot manufacturers and robot owners.

4 AN INTERNATIONAL PERSPECTIVE ON OWNER RESPONSIBILITY FOR DOMESTICATED ANIMALS

Given the framework described above as a starting point, we should look beyond the United States to see how other countries deal with liability for domesticated animals. In this section, we examine the laws of European and Asian countries concerning, in particular, dangerous dogs. We also look at Canadian law and take a second look at the United States to see how the military and certain states have implemented regulations similar to those in Europe and Asia.

We focus on dogs for two reasons. First, the division of dogs into breeds is naturally analogous to robots, which can be divided into categories based on functionality (service robots, medical robots, military robots, and so on). Also, and perhaps more importantly, there are far more laws and regulations on dangerous dogs than on, say, "dangerous hamsters" or other domesticated animals. Thus the body of law regulating dangerous dogs provides a large and varied source of inspiration for the development of law governing robots.

4.1 Europe

European countries have implemented a number of restrictions on animals considered dangerous to the public. A number of countries (including Denmark, Germany, Portugal, and Great Britain) have implemented bans on specific dog breeds that are considered dangerous. These bans are difficult if not impossible to circumvent by legal means. Still other European nations (including Spain, Poland, Ireland, and France) have implemented substantial restrictions on breeds considered dangerous. Although the restrictions differ in each country, they include: microchip implantation, requirements for expanded liability insurance in case the dog in question attacks a person, requirements that owners of a dangerous dog display signs warning of the presence of a dangerous dog, requirements for reinforced enclosures for the dangerous dog, requirements that the dog in question be muzzled and on a leash while in public, requirements that an adult must be in control of the dog while in public, special collars for identification, and additional notification requirements in the event that the dog is lost, dies, moves out of the area, or moves within the same jurisdiction. Moreover, in many nations, there are substantial fines and other penalties for owners of illegal dogs or restricted dogs that attack a person. In cases of attack, the dog

may be (and often is) confiscated and destroyed.

4.2 Asia

Countries in Asia and Oceania that have regulated certain dog breeds tend to have similar regulations to those listed for Europe. In particular, both Singapore and Australia have implemented regulations similar to those in use in Europe. Additionally, individual Australian states have implemented their own regulations, again similar to those listed above.

4.3 North America

4.3.1 Canada

The Canadian government has no regulations at the federal level concerning dangerous dogs. However, a number of provinces have implemented bans and restrictions on dangerous dogs.

4.3.2 The United States

The Military Although the the federal government has not implemented any bans or restrictions on dangerous dogs, both the United States Army and the Marine Corps have banned or restricted ownership of animals considered to be dangerous.

The States Similarly, one state (Ohio) and several local governments have implemented restrictions or bans on dangerous dogs.

5 EXTENDING ROBOTS AS ANIMALS

5.1 General Lessons from the International Analysis

In general, we find that the laws concerning so-called "dangerous dogs" in European and Asian countries are stricter than most of the laws in the United States. Whereas liability for domesticated animals is mostly a civil concern, in Europe and Asia, the regulations include criminal *and* civil penalties. The laws of these countries divide dogs into two broad categories: so-called "dangerous dogs" and all other dogs. This distinction is drawn at the breed level, although in some jurisdictions individual dogs that would otherwise be considered safe can be placed on a list of dangerous dogs at the discretion of the courts. The general sense of this approach seems useful for our purposes. In particular, the emphasis on breeds or *types* of animal has a natural analogy to robots, where the "type" of a robot can characterized by that robot's function. For example, we can view a robot such as iRobot's vacuum cleaner the Roomba as a service robot that is largely harmless, whereas a U.A.V. operated by the military or law enforcement would, especially if heavily armed, be considered "dangerous." In general, the idea of classifying robots according to their capacity to do harm seems potentially useful.

The extent of restrictions varies substantially by country. In the strictest cases, ownership of a dangerous dog is completely prohibited; dangerous dogs found in these jurisdictions are confiscated and destroyed. In other cases where dangerous dogs are regulated, owners are required to take special precautions to ensure that their dogs are sufficiently controlled to prevent harm to others. Again, it is easy to see how this approach can be applied to Robots As Animals. One can easily imagine types of robots that should be banned from ownership in the general public (e.g. armed military robots), as well as robots that might be available for widespread ownership but subject to regulation and restrictions (e.g. private security robots armed with non-lethal weapons). Additionally, one can further imagine robots that are, practically speaking, considered "safe" within the range of their intended uses. This could include small service robots or "toy" robot companions.

Finally, many of the restrictions employed to control dangerous dogs can be easily adapted to deal with robots. For example, all of the following can be applied to robots with almost no change at all: microchip tracking, liability insurance requirements, owner-displayed warning signs, adult control while in public, special identification while in public ("robot collars"), and expanded notification requirements.

5.2 Additions to the Basic Framework

With the above lessons in mind, we therefore propose a modified framework for liability in human-robot interactions. In the new framework, we distinguish between robots that are dangerous and robots that are considered largely safe. For robots that are considered generally dangerous, nations may choose to implement bans or restrictions such as those discussed above. For instance, it probably makes sense to require companies employing armed security guard robots to register those robots with the government at some level, even if those robots use only non-lethal weapons. Similarly, one can easily envision types of robots for which expanded insurance requirements, warning signs, and special identification would be appropriate. On the other hand, robots that are not considered dangerous should probably not be subject to these kinds of restrictions unless a special case warrants the classification.

In the event of an accident involving a robot, the new framework first considers the type of the dog involved in the accident. If the robot was classified as being sufficiently dangerous to be banned, the owner may face substantial criminal penalties. Even if the robot was restricted instead of banned, it may be that the owner of the robot in the accident failed to follow the laws concerning control over the robot. In such cases the robot's owner may still face penalties specified under criminal law. In addition to these criminal penalties, we also apply our original civil framework, employing the ideas of strict liability for manufacturers and negligence for owners.

If the robot involved in an accident is not classified as dangerous by the law, then in general we do not apply criminal penalties. In these cases, we simply apply the civil framework established in our previous work to assign liability among the robot's manufacturer, owner, and the victim in the accident.

5.2.1 The Robot Taxonomy

In order for the modified framework to be useful in practice, the first step in using the framework must be to carefully design the robot taxonomy to be used. In general, this taxonomy should consider how dangerous a robot is, as we do here, but could also examine other issues, such as the extent to which the robot is responsible for humans' safety or well-being.

5.2.2 Owner Responsibility for Other Robots

Although we have considered so-called "dangerous robots" here, it is clear that there are other kinds of robots that do not naturally fall into either the "safe" or the "dangerous" category. For example, assistive robots may have a substantial role in caring for the elderly, and may cause substantial harm if they are poorly designed, manufactured, or just don't work well. However, these robots seem dangerous in a very different sense from that of the robots mentioned above. With this in mind, a reasonable taxonomy may include more factors than simply a robot's capacity to do harm.

5.2.3 Special Cases and the Courts

In the case of dangerous dogs, there are several nations in which the courts have substantial discretion in the way they apply existing law. We should allow for similar discretion in the case of robots. This can work in two ways. First, a dangerous robot that is involved in an accident may not face the full penalties under the law in some cases – it may simply be unnecessary to destroy a robot, even if it is considered dangerous by the state. At the same time, it may be that a robot that is considered "safe" turns out, by its actions, to be more similar to a dangerous robot. In such cases, it would be appropriate for the court to treat the robot as if it were dangerous, placing restrictions on the robot in the future.

This flexibility points to an important aspect of our framework. Although much of the framework is implemented via laws passed by legislatures, the application of the framework is carried out by courts that have some measure of flexibility in how they apply the law. We contend that this approach is much better than one that simply tries to solve the problem of liability in HRI via legislative action alone, since the courts are much better able to adapt to changing circumstances and advancing technologies than are legislatures, which may be slow to act or (just as likely) may overreact and stifle innovation by passing laws that place an onerous burden on robot developers.

Extending the Taxonomy: Guiding Principles The initial Robots as Animals framework is a good starting point. Robots are products and, therefore, should be subject to strict products liability if they are defective. Otherwise, robots should be subject to negligence liability because robots are more like domesticated animals than wild animals based on degree of predictability. But predictability has limited usefulness for robot taxonomy, which is a refinement of the initial Robots as Animals framework. Unlike an animal, a robot may be most unpredictable when it is somehow defective in its manufacture or design (i.e., hold the robot manufacturer strictly liable), or when it is subject to wear and tear

(i.e., hold the robot owner liable for failing to perform regular maintenance). Thus, the better guiding principle for our robot taxonomy is the degree of danger the robot poses. A robots dangerousness is based on environmental factors (e.g., where the robot operates and how it is expected to interface with its surroundings, especially people but also property) and the characteristics of the robot (e.g., hardware and software, size and function, and so on).

6 THE NEW FRAMEWORK IN ACTION

To show how the framework described above may be used in practice, we consider the following examples.

6.1 The Friendly Robot Companion

We start with a simple example. Suppose that Owen owns a friendly robot companion, similar to the Paro robotic seal. The companion makes sounds, responds to touch, and can make small motions, but cannot transport itself in any way. Our framework would classify such a robot as decidedly safe, would not require special restrictions of any kind, and in the (admittedly unlikely) event of an accident involving this robot, we would simply apply our original framework without modification.

Although this example may seem uninteresting for legal purposes, it demonstrates the important fact that for many types of service robots, our original framework can function perfectly well to govern human-robot interactions. It is only when greater danger or responsibility come into play that our new standard becomes relevant.

6.2 The Security Guard Robot

In contrast with the friendly companion described above, suppose that Owen is a store owner who has purchased a security robot to guard his store. This robot is able to autonomously patrol Owen's store, notify Owen of suspicious behavior by customers, and in some cases even use an electroshock stun gun to immobilize customers who are acting threatening. Suppose further that Vicky, a customer in Owen's shop, finds herself confronted by the security robot and, for whatever reason, ends up receiving a large electrical shock from the robot. If it turns out that the jurisdiction in which Owen's store is located requires registration for armed robots and Owen failed to maintain the proper registration, then he may be open to a criminal penalty (say, a fine) in addition to whatever penalties he faces under the original (civil) framework.

6.3 The Modified Robot Lawnmower

For a third example, suppose that Owen owns a robotic lawnmower. Although a lawnmower does have the potential to cause harm, it seems reasonable to say that the degree of harm that a lawnmower can cause during normal operation is smaller than an armed security robot. For this reason, it would probably be the case that the restrictions on such a service robot would be either weak or nonexistent

in general. Suppose, however, that Owen has decided to "augment" his robotic lawnmower with small flame throwers that fire periodically as the robot mows Owen's lawn. Suppose that this modified robot is mowing the lawn when crosses paths with Vicky, a passerby who stops to watch Owen's flame-throwing lawnmower. As the robot passes Vicky, it releases a burst of flame, causing Vicky severe burns. Although typically one would not expect a lawnmower to be subject to strict regulations, in this case it seems appropriate that the reaction of the government to Owen's modifications be fairly harsh. In particular, it is likely that the courts would decide to treat the lawnmower as a "dangerous robot" for legal purposes, perhaps even seizing the robot and destroying it. Beyond this, one would expect that Owen would face criminal penalties such as fines under our extended framework, and would additionally face civil penalties pursued by Vicky. In this case we see how a robot, considered harmless by default, can be deemed by the courts to be dangerous, so that the courts can apply the law with the greatest flexibility as needed.

7 CONCLUSION

In this paper, we have offered an extension of Robots As Animals that takes into account international norms concerning liability for domesticated animals. We have summarized the laws of several nations concerning dangerous pets, particularly dogs, and shown how these laws may be adapted for use in regulating human-robot interactions. We presented our expanded framework, identified some of the primary considerations in the design of such a framework for practical use, and showed how the framework could be used in several practical examples.

We are currently investigating a number of avenues for future research. First, we are engaging in a careful analysis of robot defects, aiming to identify a set of legal criteria to decide when a robot should be considered defective. We are also exploring other uses of the robot taxonomy, possibly outside of questions of law and liability. We are also examining legal questions related to the use of robots by government, for example by the police and military, to see how our framework might be adapted to those special cases.

REFERENCES

[1] P. M. Asaro, "Robots and Responsibility from a Legal Perspective," in *Proc. of the IEEE 2007 International Conference on Robotics and Automation (ICRA07)*, Rome, April 2007.

[2] Y. Weng, C. Chen, C. Sun, "Toward the Human-Robot Co-Existence Society: On Safety Intelligence for Next Generation Robots," *International Journal of Social Robotics*, 2009.

[3] T. Kim and P. Hinds, "Who Should I Blame? Effects of Autonomy and Transparency on Attributions in Human-Robot Interaction," in *Proc. of the International Symposium on Robot and Human Interactive Communication (RO-MAN06)*, September 2006, pp. 8085.

[4] P.A. Mudry, S. Degallier, A. Billard, "On the influence of symbols and myths in the responsibility ascription problem in roboethics - A roboticists perspective," in *Proc. of the International Symposium on Robot and Human Interactive Communication (RO-MAN08)*, August 2008.

[5] E. Schaerer, R. Kelley, M. Nicolescu, "Robots as Animals: A Framework for Liability and Responsibility in Human-Robot Interactions," in *Proc. of the International Symposium on Robot and Human Interactive Communication (RO-MAN '09)*, September 2009.

[6] W.L. Prosser, W.P. Keeton, *The Law of Torts*, 1984.