

Grounding Robot Navigation in Self-Defense Law

James Zhu¹, Anoushka Shrivastava², and Aaron M. Johnson¹

Abstract—Robots operating in close proximity to humans rely heavily on human trust in them to successfully complete their tasks. But what are the real outcomes when this trust is violated? Self-defense law provides a framework for analyzing tangible failure scenarios that can inform the design of robots and their algorithms. Studying self-defense is particularly important for ground robots since they operate within public human environments, where they can pose a legitimate threat to human safety. Moreover, even if ground robots can guarantee human safety, the perception of a threat is enough to warrant human self-defense against robots. In this paper, we synthesize works in law, engineering, and the social sciences to present four actionable recommendations for how the robotics community can craft robots to mitigate the likelihood of self-defense situations arising. We establish how current U.S. self-defense law can justify a human protecting themselves against a robot, discuss the current literature on human attitudes toward robots, and analyze methods that have been produced to allow robots to operate close to humans. Finally, we present hypothetical scenarios that underscore how current robot navigation methods can fail to sufficiently consider self-defense concerns and the need for the recommendations to guide improvements in the field.

I. INTRODUCTION

Recent work in the robotics field has focused on deploying robots in public environments around humans for applications such as delivery [1], security [2], and personal assistance [3]. Delivery robotics companies such as Starship, Kiwi, and Serve have been rapidly growing the number of robots operating in public spaces like university campuses and sidewalks. In these environments, robots experience many close encounters with a diversity of humans. It is unavoidable that a number of these encounters will result in some kind of danger to the robot and/or human, such as a robot combusting due to a battery malfunction [4], college students vandalizing a robot [5], or a robot crashing into a car [6]. Though uncommon, these edge cases underscore that humans and robots can, and will, come into direct physical conflict with each other.

While the examples above were a result of some combination of accident, negligence, and wanton violence, there are cases where a human using physical force against a robot could be justified under fear of threat to one’s own well-being. Consider the following scenario, which is discussed in further detail in Section V:

- A person is tired after a long day of work and slowly walks home at night. A robot, mindful of keeping a safe distance from the person, quietly treads behind them at a constant distance. The person, on alert due to walking home alone at night, becomes fearful of something stalking them from behind. Eventually, the robot decides there is enough room on the sidewalk to accelerate and pass the person on their left.

In this example, both the human and robot behave reasonably, but the context and environment in which this interaction occurs could lead the human to feel threatened and damage the robot in self-defense. Our research is focused on analyzing such situations where use of force in human-robot interaction may be justified under self-defense law.

We propose categorizing self-defense scenarios as the tuple: (threat, protector, protectee). In standard self-defense where a human protects themselves from another human, this self-defense tuple is (human, human, self). There are also many other legally established tuples such as (human, human, other), (human, human, property), and (property, human, self). Note that even though not all of these tuples represent a human defending themselves, we use the term self-defense in this work due to its broad familiarity. This paper analyzes in detail the (robot, human, self) tuple, a special case of (property, human, self), while our future work will entail characterizing other self-defense situations involving robots.

The objective of this paper is to first extract elements from prior works that are key to the formation of a (robot, human, self) self-defense scenario, and secondly present four actionable recommendations for roboticists to design systems that mitigate the likelihood and severity of these self-defense situations. We aim to provide evidence-based guidance for roboticists across the industry so that the field can make a positive, equitable impact on our communities.

This paper draws upon the fields of self-defense law, human-robot interaction, and robot path planning to present an unaddressed topic that lies at their intersection. These research areas are synthesized to establish how the current robotics state-of-the-art overlooks the possibility of non-expert, non-user humans acting in justified self-defense against a robot. The rest of the paper is organized as follows:

- Sec. II reviews the guiding principles that inform U.S. self-defense law and argues for the importance of considering (robot, human, self) scenarios.
- Sec. III connects these principles to findings from human-robot interaction literature on human attitudes toward robots and identifies aspects of human attitudes that require additional study.

This material is based upon work partially supported by the U.S. National Science Foundation under grant #CMMI-1943900.

¹Department of Mechanical Engineering, Carnegie Mellon University, Pittsburgh, PA, USA, {jameszhu, amj1}@andrew.cmu.edu

²School of Computer Science, Carnegie Mellon University, Pittsburgh, PA, USA, anoushk2@andrew.cmu.edu

- Sec. IV discusses planning strategies that have been developed for robots to operate around humans and some of their limitations on handling self-defense situations.
- Sec. V synthesizes the previous sections to present hypothetical scenarios where a human would be justified in taking self-defense action against a robot.
- Sec. VI summarizes this work and discusses future work that the authors will address next to bolster the connection between robot design and self-defense law.

II. U.S. SELF-DEFENSE LAW

In this section, we outline the key tenets of self-defense law in the United States. These laws vary appreciably by jurisdiction in the extent and environments in which self-defense is justified. For example, most U.S. states oblige no duty to retreat from a threat, while others impose a limited duty to retreat in public spaces [7]. Because of the great variance that exists in this domain, practitioners should use discretion with the specific self-defense statutes in their jurisdiction. However, we believe the principles outlined in this section are broadly applicable to the vast majority of self-defense codes in both the U.S. and other countries, and the conclusions drawn in this work are largely independent of the specific intricacies of individual statutes. The legal discourse regarding self-defense is lengthy and the analysis presented here is quite brief. For a deeper discussion on the theory of self-defense law, see [8,9].

The two primary principles that underpin American self-defense law are:

- a reasonable belief of imminent physical harm
- a proportional response to the threat

While terse, these aspects of self-defense carry centuries of legal nuance that must be interpreted by the courts on a case-by-case basis. Below we analyze some important details of these two principles.

A. A Reasonable Belief of Imminent Physical Harm

This phrase can be broken down into two parts. Firstly, a “reasonable belief” means that the protector in the (threat, protector, protectee) tuple does not need to have definitive proof that the threat of imminent harm is true, but only that it would be reasonable for a person to view a situation as threatening [10]. For instance, if the protector is acting with limited information about the threat, there can exist a reasonable belief of harm even if the protector is ultimately found to have been incorrect about the nature of the threat. Similarly, the reasonableness of the belief is subject to prior experiences of the protector. So even if an external third party would not necessarily view the belief as reasonable, context such as prior history of physical abuse [11] or legitimate verbal threats of violence that occurred directly prior [12] can provide justification for acting in self-defense.

The second aspect of this definition is “imminent physical harm” [13]. Imminence indicates that the harm to the protectee must be actively happening or about to happen. This means that force used either proactively or after the harm has subsided is invalid to justify self-defense. There has

been debate centered around the soundness of imminence as an indicator of the necessity of self-defense, but in practice, this interpretation has been upheld [13]. Additionally, the law tends to interpret the harm that is incurred as needing to be physical. While some have argued that other types of harm such as invasion of privacy may be tantamount to a self-defense justification [14], American common law has yet to grant this argument [15].

B. A Proportional Response to Threat

Once the protector reasonably believes that self-defense is warranted, they must act in a manner appropriate to the threat level [16]. For self-defense between humans, this means that relatively minor force, such as a punch or kick, can not be responded to with lethal force. However, certain situations can justify a lethal response to a less than lethal threat, such as when the protector is in their own home [17].

In a (robot, human, self) scenario, it is key to understand that the robot is property and has no intrinsic right to act in self-preservation in the way that humans do. Because of this, a self-defense act that destroys the robot can be justified even with a lesser threat to the human. This concept of proportional response emphasizes the need for engineers to carefully design robots to avoid self-defense situations, because any perception of threat could lead to justified destruction of the robot.

C. Self-Defense Against Robots

Ground robots exhibit characteristics that can uniquely give rise to self-defense justifications. While previous work has discussed self-defense against aerial drones [14], the argument justifying self-defense against these technologies was weakened by the typically large distance between the drone and any given person. Ground robots, on the other hand, are expected to come into close, immediate contact with humans during normal operation, so designing for self-defense situations is critical.

Since we have established that the proportionality criteria can be broadly satisfied in (robot, human, self), the primary challenge to determining when people may be justified to defend themselves against a robot is defining what a reasonable belief is. Because robots are still such a novel and unfamiliar technology to most people, the standard of reasonableness may be lowered to take into account the misconceptions and misunderstandings non-experts tend to carry as they interact with and react to robots.

To begin codifying self-defense law as a serious consideration for robot engineers and to assuage the public’s fears about coming into contact with potentially dangerous robots, we make the following recommendation:

Recommendation 1. *Robotics companies and research organizations should publicly advertise that in situations where a human and robot are in direct physical conflict, the human’s physical well-being is always valued more greatly than that of the robot, even if the result is damage or destruction of the robot.*

This recommendation reinforces the justification of the proportionality criteria for self-defense against robots, which will empower people with the understanding that they will have at least some immediate recourse if they feel physically under threat by a robot. In turn, this transparency from the robotics field on the rights that people have when interacting with robots can engender a greater sense of trust and openness to the deployment of robots in public spaces. As discussed in Section III, the attitudes that humans have towards robots has a large impact on how a human-robot interactions play out.

III. REASONABLE PERSPECTIVES TOWARD ROBOTS

In this section, we take a step toward understanding standards of reasonableness in human-robot interaction. We start by analyzing how behavioral norms can dictate (human, human, self) self-defense scenarios. We next contrast the role of behavioral norms in human-human and human-robot interactions and present results from prior literature on human attitudes toward robots, which contribute to an understanding of reasonable behavior. There lacks consistently interpretable robot behaviors that could be used to establish human-robot norms, and instead suggest considering diverse attitudes toward robots to inform what constitutes reasonable behavior.

A. Behavioral Norms in Human Interactions

In human-human interactions, behavioral norms between humans play a crucial role in determining reasonable behavior. When evaluating (human, human, self) cases, courts judge the reasonableness of a person's behavior based on broad, implicit understandings of how humans typically behave. For example, in the case of *Rowe v. United States*, it was found that even though Rowe kicked a man, which prompted the man to attack him with a knife, Rowe stepping back after the kick revived his right to self-defense [18].

Following the literature [19], we define a norm as a widely adhered to and understood action that helps coordinate behavior. This stepping back is a norm that was found to be a widely understood indication that Rowe was no longer a threat and withdrawing from the fight, which protected him from being retaliated against. Such implicit understandings between humans can be considered a basis of self-defense law because they create standards of reasonableness, allow for generalizations across cases, and ultimately promote fairness in decision-making.

Norms for human-robot interaction, on the other hand, are not well established in the sense that actions taken by the robot are not widely understood and do not facilitate coordinated behavior between the human and robot, which can make robot behavior unpredictable to a non-expert. Even though human behavior is not completely predictable, years of lived experience interacting with other humans allows for a deeply developed understanding of typical human behavior. Due to the novelty of robots, this understanding between humans and robots is inadequate. Because of the lack of established norms for robots, the standard for justified self-

defense in (robot, human, self) should be lower than in (human, human, self) cases.

Frameworks such as COMPANION have attempted to address this gap by encoding human social norms (such as maintaining personal space and moving to the right to avoid colliding with people approaching from the opposite direction) within robot behavior [20]. However, more research is needed to determine whether humans actually expect robots to behave with the same norms as humans. In fact, some studies have indicated otherwise. For example, [21] found that robots were considered more trustworthy when approaching a person quickly (possibly since faster robots were more noticeable than slower robots), whereas humans were more trustworthy when approaching slowly. The study also found that humans performed more corrective reactions (such as stepping back or adjusting eye contact) when a robot invaded their personal space compared to a human. This indicates that humans react differently to violations depending on if the offending party is a human or robot, and simply having robots adopt human norms does not guarantee self-defense situations will be avoided.

Even if robot behavioral norms become standardized among the industry, there is no guarantee that humans will, in their split-second decision making, have enough trust to assume that a robot can reliably follow certain norms. It is important for robots to conform to the preferences and expectations that humans have for their behavior, and we recommend further study of these topics:

Recommendation 2. *Because of the differences in human preferences and expectations when interacting with robots compared to other humans, researchers should explore whether there are robot behaviors that humans react consistently to and if these behaviors can be encoded into a standardized framework of human-robot norms.*

Results that establish norms for even a subset of ground robots (such as wheeled robots, humanoids, or quadrupeds) could begin to establish a more refined definition of reasonable behavior around robots. Possibilities of behavioral norms for robots could be exhibiting body language or digital facial expressions [22], which could in turn improve the legibility of robots in public environments (as discussed in Sec. IV-B). Establishing norms may be difficult due to the variance in attitudes that humans exhibit toward robots. Even if consistent norms cannot be identified for many aspects of robot behavior, it is still important to characterize how attitudes vary among humans and under what circumstances human attitudes can be well modeled.

B. Human Attitudes Toward Robots

Human attitudes towards robots tend to vary based on several factors, including a person's familiarity with robots and how well a robot's behavior aligns with the human's expectations and preferences. Studies suggest that the more familiar a person is with a robot and the more their expectations align with the robot's behaviors, the more positive their attitude towards that robot [23]. Conversely, when there

are gaps and discrepancies in these areas, attitudes tend to shift negatively. The real consequences of negative attitudes toward ground robots and a violation of expectations during their deployment emerge as justified self-defense scenarios. Alleviating the public’s negative attitudes and aligning robot design with expectations is essential for safe human-robot interactions. This reinforces the necessity of Recommendation 1, which can help inform the public’s expectations of how they can interact with robots.

One highly-documented example of this is the largely negative attitudes communities have expressed toward the deployment of robots by police departments in several U.S. cities [24]. One issue that arose was that communities affected by the police’s usage of ground robots were not involved in the development process and expressed frustration over the expensive and possibly dangerous technology [24,25], calling the robot “another danger for Black & Latino residents” [25] and expressing fear toward the futuristic appearance of the robots [26]. Requests to police departments for more information on the purpose of the robots were not always met [24,27].

When impacted communities are not involved in the development of robots, the deployed products can be misaligned with the community’s expectations of how these new tools should be used. Research indicates that the level of familiarity people have with robots and the preconceptions they hold influence their attitudes, such as how the fear of sentient robots correlates with negative attitudes [23]. Another study suggests that people may be more likely to support robots doing jobs that require less experience and communication [28]. Therefore, attitudes toward delivery robots, which satisfy both of these conditions, could be more positive than those toward police robots, though further research should test this theory.

Attitudes of marginalized groups toward robots are especially important for developers to consider, since they have been disproportionately affected by harmful uses of novel technologies [29]. Police robots have often been deployed to patrol low-income, Black neighborhoods [24,30], while women have repeatedly been targets of unwelcome surveillance by drones [31,32]. Additionally, [33] found that women tended to be less receptive to the concept of patrolling police robots than men. To combat this inequity, [24] suggests involving marginalized community members in the technology design process. Factoring in the preferences of the stakeholders who interact most closely with robots will help developers align robot design with expectations, reduce negative attitudes toward robots, and promote equity by working for marginalized communities instead of against them. We recommend further investigation into how attitudinal differences manifest among disadvantaged groups:

Recommendation 3. *Due to the variance in human attitudes toward robots and the disparate effects technologies have had, researchers should examine and catalog the attitudinal differences among different groups of people, especially from those that have historically been marginalized.*

Reasonable human behavior varies greatly due to differences in background and past experiences. These differences can be measured in the attitudes, perspectives, and reactions people exhibit toward robots. Ultimately, this variability suggests that even severe human behaviors towards robots can be justified and considered reasonable, at least until the establishment of robot behavioral norms that are broadly understood by people of many backgrounds. Instead of basing robot behavior on unestablished norms that people must adhere to, it is essential to consider people’s diverse attitudes and expectations regarding robots and design robots with this context in mind.

IV. HUMAN-AWARE PLANNING

While self-defense has so far gone unconsidered in the design and implementation of robots, there has been ample related work in planning robot motions in human environments. Generally, robot path planning is performed by sampling many possible paths a robot can take and selecting the most optimal choice, often based on the shortest path [34,35]. Algorithms are also able to obey specified constraints such as avoiding obstacles. Recent research has adapted these path planning algorithms to predict and react to human obstacles, and to minimize risk of collision with people [36]. Other work has developed robotic behaviors to satisfy desired outcomes such as visibility [37], active communication [38], and following social norms [20,39]. In this section, we examine two primary research thrusts in human-aware planning that have seen significant attention: explainability and legibility. We analyze not only what work has been done, but also the reasons stated in the literature for why these aspects of human-aware planning are important. While aspects of explainability and legibility are useful in mitigating the potential for self-defense situations, current implementations lack the capability to address all environments in which a self-defense scenario may arise.

A. Explainability

Drawing from [40], we define explainability as the ability of an autonomous agent to produce records of the decisions it has made and understandable reasoning for why those decisions were made. This definition is compatible with how explainability is discussed in prior works, such as generating contrastive explanations (i.e. why A and not B?) [41] and explanations that satisfy user-defined preferences [42]. Post-hoc explanations are designed to be generated after a robot decision has been executed and in response to some kind of questioning, while some work has examined generating concurrent explanations for behaviors as they are happening [43].

Drawing from the existing discourse on explainability in the field of autonomous vehicles (AVs), [44] discusses post-hoc and concurrent explanations for AVs by analyzing a scenario in which an AV fails to recognize a pedestrian crossing in front of it. Post-hoc explanations to characterize why the AV failed could be useful in a post-accident investigation and for regulators to hold manufacturers accountable. But

these post-hoc explanations would not be able to prevent accidents from occurring. Concurrent explanations, such as communication to a passenger that the car will continue through a crosswalk because no pedestrian has been detected, could allow passengers to take emergency actions when they recognize the vehicle has made an error. In this instance, a passenger could activate an emergency brake that stops the car before it collides with the pedestrian.

In the context of ground robots, post-hoc explanations assume that the people who desire explanations for robot behaviors have access to the robot afterward. These explanations could be useful to operators who could recognize errors in their usage of the robot, developers who could better understand errors and implement fixes, and members of the judiciary who could use explanations to assign liability after an accident. However, post-hoc explanations generally exclude members of the general public who interact with the robot for just a fleeting moment, such as passing each other on the sidewalk. Considering that the majority of people interacting with a robot in a public environment will likely not have access to that robot afterward, this exclusion is significant. Concurrent explanations, on the other hand, are able to actively communicate to people in the robot's immediate surroundings. However, concurrent explanations may be difficult to convey to certain people in real-world environments. Explanations announced verbally may not be heard by people on the phone or listening to music, or may be drowned out in loud environments such as construction. Similarly, explanations presented visually may not be suitable for people with visual impairments or in night-time environments. A robot must also consider that some people may not speak the robot's default language.

In a survey of 62 papers on explainability, [45] found that the most commonly stated motivation for the work was transparency (i.e. allowing people to better understand the inner workings of the robot), followed by trust and collaboration. These motivations go hand in hand, as increased transparency would naturally lead people to trust being around the robot and working with it. Of these surveyed papers, many framed trust around the relationship between robots and their operators or teammates, and the faith these people had that their robots would work reliably [46,47]. With the European Union's recent General Data Protection Regulation (GDPR) outlining a person's legal "right to explanation" when encountering autonomous agents [48], explainability has come to even greater relevance.

In this work, we question if explanations are the key to inspiring public trust in robots and allowing for a seamless deployment of robots in public spaces. Explanations can be valuable, but are ineffective in fully mitigating the possibility of self-defense scenarios due to the difficulties of access and communication with typical bystanders. What robots need is the capability to generate implicit methods of communication that can foster an improved understanding of robot behavior. The human-aware planning concept of legibility may be a more suitable method to accomplish this.

B. Legibility

While explainability focuses on the producing reasons for why a robot behaved a certain way, legibility characterizes how a robot communicates what it is doing or intends to do. Based on the work from [49], we define legibility as the ability of a human to understand a robot's intentions based on observation. For instance, cars have turn signals to indicate to others what action they are about to perform (e.g. turning right). The turn signal does not explain why the car is turning right, but allows others to understand what it is about to do and react accordingly. Explanations can help robots become more legible, but there are many other factors that can improve legibility such as providing cues [50], mimicking human behavior [51], moving quickly toward goals [52], and staying in people's field of view [37].

Legibility is a somewhat vague concept that is difficult to define and measure experimentally. To do this, some authors have evaluated legibility by asking people to predict a robot's future behavior based on past observed behavior [53,54]. Others evaluated human performance of an unrelated task while a robot navigated around them [55]. Others still used questionnaires to gauge how well subjects felt they understood a robot's intentions [50], while [52] proposed a numerical measure of a trajectory's legibility.

Studies have found that legibility is correlated with increased feelings of safety, comfort, and acceptance [55–57]. These goals align very closely with the stated purposes of explainability, but may differ in the groups of people these methods are designed for. While explainability is often framed around expert users, significant work has focused on robot legibility to non-experts, as the lesser amount of information needed may make it easier for non-experts to comprehend. However, as with generating explanations, interacting with a diverse group of humans may make producing legible behaviors much more difficult. In these situations, it may not be possible to stay in everybody's field of view or to expect all people to notice the robot's gestures. The complicating factors discussed for explainability can make communication difficult here as well, as visual or verbal signals can breakdown in certain cases. Engineers must also consider how obvious any given message is to non-expert non-users. This relates closely to the concept of human-robot norms, discussed in Sec. III. Norms are not yet established for robots, so developing robots that can be legible to people across diverse backgrounds is an enormous open problem.

C. Human-Aware Planning for Non-Expert, Non-Users

Based on this section, we conclude that current human-aware planning algorithms have not sufficiently addressed how robots should operate in dense human environments where communication is impeded. In particular, non-expert, non-users have received little attention as to how robots should interact with them. For instance, the IEEE Standard for Transparency of Autonomous Systems lays out guidelines for what information robots should be able to communicate to people [58]. The transparency standard for users

is grounded primarily in providing explanations, but the transparency standard for non-users among the general public is focused exclusively on data privacy and lacks considerations for how robots must communicate with non-users. We recommend these standards be amended to acknowledge how non-expert, non-users interact with robots:

Recommendation 4. *Standards, guidelines, and regulations from influential organizations such as IEEE should detail how robots should interact with non-expert, non-users to cultivate perceived safety and trust amongst the public. By drawing from concepts such as legibility, these standards can provide clear direction for how robots should be developed to minimize self-defense occurrences.*

While current standards for transparent and legible robots are unsatisfactory, the literature provides guidance to suggest possible standardization of certain aspects of robot behavior, such as establishing a standard mapping of light color to indicate behavior [14,31] and requiring the deployment of noticeable robots instead of silent, stealthy ones [21]. Another regulation that could be considered is how robots should operate at night. The Federal Aviation Administration currently restricts flying drones at night [14,59] and encouraging similar regulations for ground robots could be productive. However, unlike drones, there is not one government agency that can dictate regulations on ground robots. State and local governments control the use of their roads and sidewalks, so adoption of a consistent set of regulations is unlikely. Organizations like IEEE and leading companies like Boston Dynamics and Agility Robotics may be the vanguards of establishing industry standards for operating robots. Regardless, there is still a long way to go in establishing standards and regulations that fully address self-defense against robots.

V. JUSTIFIED HUMAN SELF-DEFENSE AGAINST ROBOTS

Even when explainability and legibility are incorporated into robot planning, there is still potential for self-defense situations to arise between robots and non-expert, non-user humans. Practitioners must understand the dependence a human-robot interaction has on the exact person and environment the interaction takes place in. In the Introduction, we suggested a hypothetical case of a lone person walking home at night, who is followed from behind and then passed by a robot. Consider two additional hypothetical cases where a robot that behaves according to conventional human-aware planning principles makes a human feel uncomfortable and even threatened:

- A robot equipped with gaze tracking technology attempts to stay as close as possible to the center of a person's field of view as they wait for a bus. This person, however, is attempting to view the numbers of the buses that are passing by, and is unsettled that the robot appears to be blocking them from finding the right bus to leave on. As the robot approaches, they become afraid of the robot's single-minded focus on them.

- Navigating around a blind corner, a legged robot unexpectedly bumps into a person turning the corner in the opposing direction. In this situation, the robot is unable to satisfy the personal space constraint it has been programmed with and reverts to a "safe mode", which is to sit down on the ground. Already flustered by the sudden encounter with the robot, the person finds this behavior particularly unexpected and feels unsafe due to this unpredictability.

In each of these cases, the robots embody some aspects of current human-aware planning methods, which in many circumstances may be appropriate and increase the transparency, trust, and perceived safety that nearby people feel. However, these cases highlight ways that naive implementation of these methods can cause unintended negative effects. The robot in the first case from the Introduction takes care to maintain a safe distance and pass according to typical social norms, but fails to account for the context and environment that causes the person to be fearful of any nearby entity approaching quietly from behind. In the second case, the robot attempts to maximize its visibility, but without the understanding that the human would prefer to not have the robot so central in their field of vision and feels uncomfortable with the intense attention the robot is paying them. Finally, the third case highlights a robot's attempt to embody a norm that indicates a non-threatening disposition. However, this norm is not obvious enough to a person that must make a split-second decision on whether the robot could harm them.

Even though the robots in these examples may not pose an actual threat to the humans, the behaviors of these robots coupled with the people's backgrounds and the unique environment they are in can lead to a perception of threat. This perceived threat could manifest into the humans acting in self-defense against the robots once they are sufficiently close to each other. Self-defense in these (robot-human-self) cases would be justified because there exists a reasonable belief of imminent physical harm. These scenarios could result in damage or destruction of the robot, which would be an appropriate proportional response to the threat.

VI. DISCUSSION AND CONCLUSION

As roboticists work to rapidly ramp up deployment of their robots in public environments, it is crucial to understand the genuine physical harm these robots could cause human bystanders. Developers must design robots not only to guarantee human safety, but also to maximize the perceived safety of nearby humans. However, as robots are still largely unfamiliar to most of the general population and are often viewed with negative preconceptions, it is likely that some humans will see robots as threats to their physical safety and act in self-defense. In this work, we discuss how self-defense law applies to human encounters with ground robots, the human norms and attitudes that dictate the outcome of human-robot interactions, and the need to expand explainability and legibility to address self-defense cases. Synthesizing these three concepts, we identify scenarios where human

self-defense against robots could be justified, even under reasonable robot behavior.

These considerations inform four recommendations to roboticists that aim to reduce the likelihood of justified human self-defense against robots. Recommendation 1 addresses robotics companies and research institutions to provide open communication to the public on the rights that they have when interacting with robots they perceive as dangerous, which will promote public trust and improve human attitudes toward robots. Recommendation 2 suggests researchers examine if any implicit robot behaviors are widely interpretable to humans and if a framework of human-robot norms can begin to be established. Recommendation 3 calls for a more detailed exploration of the attitudes that marginalized groups such as Black communities and women hold toward robot deployment. This will work toward considering previously-excluded people in the development of novel technologies and reinforcing the rights of these marginalized populations. Finally, Recommendation 4 advocates for an overhaul in robot standards, guidelines, and regulations to address legible robot behavior to non-expert, non-users.

We argue that contextualizing robot navigation in self-defense law establishes tangible, relevant outcomes that developers can use to evaluate their algorithms on. We hope that this work will contribute to keeping people of all backgrounds safe and secure as robots are increasingly deployed around them.

Based on Recommendations 2 and 3, one line of future work we would like to explore in more detail is the reasonableness criteria for self-defense against robots. Specifically, we would like to experimentally establish the aspects of robot locomotion that people would be more or less likely to perceive as threatening. The gaits of legged robots in particular are an unexplored aspect of human-robot interaction we aim to investigate. We would also like to expand the (threat, protector, protectee) framework by exploring scenarios where robots could act as protectors. The legal justification for such a scenario may be much more narrow than humans protecting themselves, but some work has begun to look at robots defending themselves [60] or their human users [61].

REFERENCES

- [1] D. Lee, G. Kang, B. Kim, and D. H. Shim, "Assistive delivery robot application for real-world postal services," *IEEE Access*, 2021.
- [2] A. Lopez, R. Paredes, D. Quiroz *et al.*, "Robotman: A security robot for human-robot interaction," *International Conference on Advanced Robotics*, 2017.
- [3] A. Vega-Magro, L. V. Calderita, P. Bustos, and P. Núñez, "Human-aware robot navigation based on time-dependent social interaction spaces: a use case for assistive robotics," *IEEE International Conference on Autonomous Robot Systems and Competitions*, 2020.
- [4] S. Brown-Heidenreich, "KiwiBot catches fire outside Martin Luther King Jr. student union," *The Daily Californian*, 2018.
- [5] L. O'Haver, "UT students arrested after damaging food delivery robot," *WATE*, 2022.
- [6] E. Parks and D. Zoga, "NBC 5 responds: What happens if you're in a crash with a robot?" *NBCDFW*, 2020.
- [7] R. M. Brown, *No duty to retreat: Violence and values in American history and society*. Oxford University Press, 1991.
- [8] R. A. Epstein, "The theory and practice of self-help," *Journal of Law, Economics, and Policy*, 2005.
- [9] D. I. Brandon, M. L. Cooper, J. H. Greshin *et al.*, "Self-help: Extrajudicial rights, privileges and remedies in contemporary American society," *Vanderbilt Law Review*, 1984.
- [10] K. W. Simons, "Self-defense: Reasonable beliefs or reasonable self-control?" *New Criminal Law Review*, 2008.
- [11] L. S. Scheff, "People v. Humphrey: Justice for battered women or a licence to kill," *University of San Francisco Law Review*, 1997.
- [12] H. P. Thompson, "The State v. Sutton," *Supreme Court of Georgia*, 2015.
- [13] R. A. Rosen, "On self-defense, imminence, and women who kill their batterers," *North Carolina Law Review*, 1993.
- [14] A. M. Froomkin and Z. Colangelo, "Self-defense against robots and drones," *Connecticut Law Review*, vol. 48, no. 1, 2015.
- [15] D. Schneider, "Open season on drones?" *IEEE Spectrum*, 2014.
- [16] S. Uniacke, "Proportionality and self-defense," *Law and Philosophy*, 2011.
- [17] B. Levin, "A defensible defense?: Reexamining castle doctrine statutes," *Harvard Journal on Legislation*, 2010.
- [18] D. B. Kopel, "How the United States Supreme Court confronted a hanging judge in the nineteenth century and taught some lessons for jurisprudence in the twenty-first," *American Journal of Criminal Law*, 2000.
- [19] J. E. Anderson and D. Dunning, "Behavioral norms: Variants and their identification," *Social and Personality Psychology Compass*, 2014.
- [20] R. Kirby, R. Simmons, and J. Forlizzi, "COMPANION: A constraint-optimizing method for person-acceptable navigation," *IEEE International Symposium on Robot and Human Interactive Communication*, 2009.
- [21] A. Sardar, M. Joosse, A. Weiss, and V. Evers, "Don't stand so close to me: Users' attitudinal and behavioral responses to personal space invasion by robots," *ACM/IEEE International Conference on Human-Robot Interaction*, 2012.
- [22] I. Cohen, R. Looije, and M. A. Neerinx, "Child's recognition of emotions in robot's face and body," *ACM/IEEE International Conference on Human-Robot Interaction*, 2011.
- [23] M. K. Strait, C. Aguilon, V. Contreras, and N. Garcia, "The public's perception of humanlike robots: Online social commentary reflects an appearance-based uncanny valley, a general fear of a "technology takeover", and the unabashed sexualization of female-gendered robots," *IEEE International Symposium on Robot and Human Interactive Communication*, 2017.
- [24] A. Yunus and S. A. Doore, "Responsible use of agile robots in public spaces," *IEEE International Symposium on Ethics in Engineering, Science and Technology*, 2021.
- [25] M. Zaveri, "N.Y.P.D.'s robot dog returns to work, touching off a backlash," *The New York Times*, 2021.
- [26] S. A. Periwal, "Boston Dynamics' robot dog strikes fear online as fans recall Black Mirror episode 'Metalhead,'" *Sportskeeda*, 2020.
- [27] A. Jarmanning, "Mass. state police tested out Boston Dynamics' Spot the robot dog, civil liberties advocates want to know more," *WBUR*, 2019.
- [28] L. Takayama, W. Ju, and C. Nass, "Beyond dirty, dangerous and dull: What everyday people think robots should do," *ACM/IEEE International Conference on Human-Robot Interaction*, 2008.
- [29] J. Buolamwini and T. Gebru, "Gender shades: Intersectional accuracy disparities in commercial gender classification," *Conference on Fairness, Accountability and Transparency*, 2018.
- [30] C. Gidaris, "The rise of the robots: Technocapitalism and the policing of race," *Fast Capitalism*, 2021.
- [31] K. M. Thomasen, "Beyond airspace safety: A feminist perspective on drone privacy regulation," *Canadian Journal of Law and Technology*, 2018.
- [32] D. Tourjée, "Several women in this rural town say they were stalked by drones," *Vice Media Group*, 2017.
- [33] A. Gurinskaya, "Support for robocops: Measuring effects of attitudes towards police and policing technologies," *Digital Transformation and Global Society*, 2022.
- [34] P. E. Hart, N. J. Nilsson, and B. Raphael, "A formal basis for the heuristic determination of minimum cost paths," *IEEE Transactions on Systems Science and Cybernetics*, 1968.
- [35] S. M. LaValle, "Rapidly-exploring random trees: A new tool for path planning," *Tech. Rep. TR 98-11*, 1998.
- [36] J. Chu, F. Zhao, S. Bakshi *et al.*, "Risk-aware path planning with uncertain human interactions," *American Control Conference*, 2021.

- [37] A. V. Taylor, E. Mamantov, and H. Admoni, "Observer-aware legibility for social navigation," *IEEE International Conference on Robot and Human Interactive Communication*, 2022.
- [38] E. Sanoubari, B. David, C. Kew *et al.*, "From message to expression: Exploring non-verbal communication for appearance-constrained robots," *IEEE International Conference on Robot and Human Interactive Communication*, 2022.
- [39] Y. F. Chen, M. Everett, M. Liu, and J. P. How, "Socially aware motion planning with deep reinforcement learning," *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2017.
- [40] P. Langley, B. Meadows, M. Sridharan, and D. Choi, "Explainable agency for intelligent autonomous systems," *AAAI Conference on Artificial Intelligence*, 2017.
- [41] M. Cashmore, A. Collins, B. Krarup *et al.*, "Towards explainable ai planning as a service," *ICAPS Workshop on Explainable Planning*, 2019.
- [42] S. Sohrabi, J. A. Baier, and S. A. McIlraith, "Preferred explanations: Theory and generation via planning," *AAAI Conference on Artificial Intelligence*, 2011.
- [43] L. Zhu and T. Williams, "Effects of proactive explanations by robots on human-robot trust," *Social Robotics*, 2020.
- [44] S. Atakishiyev, M. Salameh, H. Yao, and R. Goebel, "Explainable artificial intelligence for autonomous driving: An overview and guide for future research directions," *arXiv preprint arXiv:2112.11561*, 2021.
- [45] S. Anjomshoar, A. Najjar, D. Calvaresi, and K. Fr  mling, "Explainable agents and robots: Results from a systematic literature review," *International Conference on Autonomous Agents and Multiagent Systems*, 2019.
- [46] M. W. Boyce, J. Y. Chen, A. R. Selkowitz, and S. G. Lakhmani, "Effects of agent transparency on operator trust," *ACM/IEEE International Conference on Human-Robot Interaction Extended Abstracts*, 2015.
- [47] N. Wang, D. V. Pynadath, and S. G. Hill, "The impact of POMDP-generated explanations on trust and performance in human-robot teams," *International Conference on Autonomous Agents & Multiagent Systems*, 2016.
- [48] General Data Protection Regulation, "Regulation EU 2016/679 of the European parliament and of the council of 27 April 2016," *Official Journal of the European Union*, 2016.
- [49] E. A. Sisbot, L. F. Marin-Urias, R. Alami, and T. Simeon, "A human aware mobile robot motion planner," *IEEE Transactions on Robotics*, 2007.
- [50] N. J. Hetherington, E. A. Croft, and H. M. Van der Loos, "Hey robot, which way are you going? Nonverbal motion legibility cues for human-robot spatial interaction," *IEEE Robotics and Automation Letters*, 2021.
- [51] J. Guzzi, A. Giusti, L. M. Gambardella *et al.*, "Human-friendly robot navigation in dynamic environments," *IEEE International Conference on Robotics and Automation*, 2013.
- [52] A. D. Dragan, K. C. Lee, and S. S. Srinivasa, "Legibility and predictability of robot motion," *ACM/IEEE International Conference on Human-Robot Interaction*, 2013.
- [53] B. Busch, J. Grizou, M. Lopes, and F. Stulp, "Learning legible motion from human-robot interactions," *International Journal of Social Robotics*, 2017.
- [54] S. Wallk  tter, M. Chetouani, and G. Castellano, "A new approach to evaluating legibility: Comparing legibility frameworks using framework-independent robot motion trajectories," *arXiv preprint arXiv:2201.05765*, 2022.
- [55] D. Bortot, M. Born, and K. Bengler, "Directly or on detours? How should industrial robots approximate humans?" *ACM/IEEE International Conference on Human-Robot Interaction*, 2013.
- [56] F. Eyssel, D. Kuchenbrandt, and S. Bobinger, "Effects of anticipated human-robot interaction and predictability of robot behavior on perceptions of anthropomorphism," *ACM/IEEE International Conference on Human-Robot Interaction*, 2011.
- [57] F. Dehais, E. A. Sisbot, R. Alami, and M. Causse, "Physiological and subjective evaluation of a human-robot object hand-over task," *Applied Ergonomics*, 2011.
- [58] "IEEE standard for transparency of autonomous systems," *IEEE Std 7001-2021*, 2022.
- [59] Federal Aviation Administration, "Sec. 107.29 operation at night," *Code of Federal Regulations*, 2021.
- [60] J. Zhu and A. M. Johnson, "By air or by land: How locomotion methods dictate drone ethics," *ICRA Workshop on Addressing Ethical and Technical Challenges in the Development, Use and Governance of Lethal Autonomous Weapons Systems*, 2022.
- [61] E. K. Duarte, M. Shiomi, A. Vinel, and M. Cooney, "Robot self-defense: Robots can use force on human attackers to defend victims," *IEEE International Conference on Robot and Human Interactive Communication*, 2022.