An Exploration of Mock Scenarios as a Prospective Method for Informing Law and Policy Thinking on Service Robotics

Ayan Robinson
robinsay@oregonstate.edu
Oregon State University
Corvallis, OR, USA

Cindy Grimm
grimmc@oregonstate.edu
Oregon State University
Corvallis, OR, USA

Ruth West
ruth.west@unt.edu
University of North Texas
Denton, TX, USA

Naomi T. Fitter
naomi.fitter@oregonstate.edu
Oregon State University
Corvallis, OR, USA

ABSTRACT

Modern service robots (such as delivery robots) are becoming more commonplace, and it is important to bridge the gap between legal and engineering thinking to be able to propose and implement effective policy. As a step in this direction, our presented work aims to inform lawyers and policy-makers on useful policy thinking surrounding service robots. Specifically, we use a mock robotic technology ordinance and accompanying mock scenarios to assess whether different types of shared context and information seem to influence robot law and policy experts' analysis of ordinances related to robotic technology. The presented work provides preliminary insights on how perceptions and interpretations of robot policy might change across different levels of background grounding and robotics understanding. Our findings can contribute to the ongoing discourse on responsible development and regulation of service robotic systems.

CCS CONCEPTS

• **Computer systems organization** → **Robotics**; • **Applied computing** → **Education**; • **Human-centered computing** → **Human computer interaction (HCI)**; • **Security and privacy** → *Human and societal aspects of security and privacy*.

KEYWORDS

robot law and policy, education via mock scenarios, service robotics

ACM Reference Format:
Ayan Robinson, Cindy Grimm, Ruth West, and Naomi T. Fitter. 2024. An Exploration of Mock Scenarios as a Prospective Method for Informing Law and Policy Thinking on Service Robotics. In *Companion of the 2024 ACM/IEEE International Conference on Human-Robot Interaction (HRI '24 Companion), March 11–14, 2024, Boulder, CO, USA.* ACM, New York, NY, USA, 4 pages. https://doi.org/10.1145/3610978.3640566

HRI '24 Companion, March 11–14, 2024, Boulder, CO, USA

ACM ISBN 979-8-4007-0323-2/24/03.
https://doi.org/10.1145/3610978.3640566

1 INTRODUCTION

Robots — such as the modern wave of autonomous delivery robots — are entering day-to-day human-populated settings, with relevant policy often trailing one step behind and sometimes struggling to connect to the realities of robotic systems. Bridging the gap between legal and technical perspectives is challenging but crucial for effective communication and policy creation [2, 3]. As one step toward this productive interchange, we considered the prospective learning intervention of mock ordinances and cases relevant to service robotics. These were loosely drawn from real ordinances and events that are already happening with delivery robots.

The broader context of this work has to do with artificial intelligence (AI) literacy, which refers to an individual's ability to understand, interpret, and engage with AI concepts, technologies, and applications. Past work has focused on developing and evaluating AI literacy programs for university education based on a comprehensive conceptual framework [8]. The rich domain of explainable AI (XAI) also represents a type of AI literacy, focusing on creating AI systems that can provide clear, human-interpretable explanations for decisions [1, 5, 6]. A related body of work shows that framing information is essential for thinking about service robot policy-relevant matters from privacy [7] to trust calibration [9]. Efforts in the domain of teaching law and policy experts about the essential technical intricacies needed for shaping reasonable frameworks around service robots (e.g., a past keynote presentation encouraging the bridging across these domains [4]), however, are a less common, yet essential, topic of study. We focus on this developing space in the work presented here.

Our central research question in this paper is whether the presentation of different mock service robot scenarios, along with a common base mock ordinance, influence thinking when it comes to evaluating the base robotic technology-relevant ordinance. We conducted a pilot between-subjects investigation with these scenarios, the full text of which is included in the paper's Appendices. The scenarios specifically focus on autonomous service robots falling under the category umbrella of food delivery or last-mile delivery, as these types of systems are quickly beginning to proliferate day-to-day life in many regions. The main contributions of this work include the scenario text itself, as well as early insights for hypothesis generation as to the future application and impact of this type of intervention.

2 METHODS

In this pilot between-subject investigation, participants in the 2023 We Robot conference (including individuals with a range of experience from law, ethics, and policy to technical robotics) responded to questions about our mock ordinances and scenarios, as further described below. This work was conducted with approval from our university ethics board.

2.1 Pilot Evaluation Design

Participants in our pilot investigation saw one of three written cases, as described below. The scenarios were written to exploit ambiguities in the base ordinance.

- *Base Ordinance (Control Condition):* Participants were presented with the base ordinance without an additional overlying scenario. This served to establish a baseline understanding of the ordinance alone. See Appendix A for the complete ordinance text.
- *MoonBot (Exploratory Condition 1):* Participants received the base ordinance, followed by a hypothetical description of a robot that a company, MoonBot, is preparing to deploy. See Appendix B for the complete scenario text.
- *Scooter (Exploratory Condition 2):* Participants were presented with the base ordinance, followed by a mock scenario in which a motorized scooter collided with a delivery robot. See Appendix C for the complete scenario text.

2.2 Procedure

Participants were recruited at the We Robot conference, an interdisciplinary conference on robot law and policy, as representative parties with interest in shaping better policies around real-world robots. Prospective participants were approached during short breaks between conference sessions and invited to participate. Willing participants received one of the three above-described cases to read. After reading the provided text, participants completed a set of questions about the case, as further described below. The scenario and questions were presented on paper to facilitate rapid deployment and data collection. Aside from one question about robotics experience, no demographic information was collected due to the brief and public nature of the study.

2.3 Measurement

The question set following the cases consisted of the following:

- (For MoonBot and Scooter Conditions) *Opening Questions about Whether the Ordinance Was Violated*: Participants responded yes/no/maybe, and completed a free-response question about their underlying reasoning.
- (For All Conditions) *Likert Scale Questions*: Participants rated the likelihood of the ordinance accomplishing each of the following on a 5-pt Likert-type scale from "Not Likely" (1) to "Very Likely" (5):
  - Improve the safety performance of autonomous delivery robots.
  - Preserve the privacy of residents.
  - Reduce the likelihood that a company will deploy an autonomous delivery system.
- (For All Conditions) *Closing Free-Response Question*: Participants shared feedback on what (if anything) they would change about the ordinance.
- (For All Conditions) *Robotics Experience Self-Report*: Participants reported their familiarity with robotic technology on a 5-pt Likert-type scale.

2.4 Participants

A total of 41 attendees of the 2023 We Robot conference completed the pilot. Demographics-wise, a substantial number of participants had a relatively high level of experience with robotic technology ($M = 4.0$ out of 5, $SD = 1.2$).

2.5 Analysis

For scale-wise questions, we used descriptive statistics to assess trending between different condition experiences, as well as analysis of variance (ANOVA) tests with an $\alpha = 0.05$ significance level to assess statistically significant differences across conditions. We used thematic analysis to analyze the free-response data.

3 RESULTS

All 41 participants completed the full questionnaire. Of the 41 participants, 11 experienced the base ordinance condition, 14 the MoonBot condition, and 16 the scooter condition. A synopsis of the categorical, scale-wise, and free-response results follows.

3.1 Ordinance Violation Verdict Results

For the MoonBot and scooter conditions, participants determined whether the ordinance had been violated. In the MoonBot case, no participants thought the ordinance had been violated, three responded maybe, and 11 responded no. For the scooter condition, perceptions were more mixed. Four participants stated that the ordinance was violated, eight responded maybe, and four responded no. Reasoning for these responses ranged widely, from "no uploads to the cloud allowed" as a reason for a clear violation for many MoonBot participants to "the ordinance isn't clear enough on what the contact info needs to include" as a reason behind a response of maybe for the scooter case.

3.2 Scale-Wise Question Results

Trends in the descriptive statistics appear in Table 1. Overall, ratings for the base ordinance tended to be more optimistic than ratings for one or both conditions when it came to preserving safety and privacy of citizens and reducing levels of robot deployment (the latter of which could be seen as good or bad, depending on the context). Once an accompanying scenario was introduced, ratings in all areas tended to decrease, aside from the one case of MoonBot safety ratings (which increased relative to the base ordinance condition

Table 1: User responses to scale-wise questions across condition. Results are reported as mean ± standard deviation. The highest mean value in each category is bolded for emphasis.

	Base Ordinance	MoonBot	Scooter
Safety	2.72 ± 1.00	**2.92 ± 0.91**	2.37 ± 0.95
Privacy	**3.27 ± 1.10**	2.57 ± 1.08	2.68 ± 1.25
Reduce Deployment	**3.27 ± 1.48**	3.07 ± 1.43	2.37 ± 1.14

ratings). Compared to the scooter condition, the MoonBot condition tended to yield higher perceptions of safety and likelihood of the ordinance to reduce deployment. The scooter condition tended to yield higher ratings for privacy preservation.

Statistical tests indicated that within the current sample, none of these trends corresponded to an associated significant difference. There were no significant differences in the privacy, safety, or deployment reduction ratings across condition (all $p > 0.191$).

3.3 Free-Response Question Results

For the closing question about changes that participants would make to the presented ordinance, key emergent themes were as follows: vehicle characteristics, data upload to the cloud, contact information requirements, privacy concerns, safety regulations, mapping and recording restrictions, and emergency stop mechanisms, as further shown in Table 2. One participant in the MoonBot condition and three in the scooter condition left this field blank; accordingly, we divide the values in the table by 11 base ordinance condition participants, 13 MoonBot participants, and 13 scooter participants when computing percentages.

Occurrence rates of different codes across condition shows some potential for the presented scenarios to guide thinking about robotics policy. For example, without the scenarios overlaid, participants thought more about vehicle characteristics (73% of cases) and mapping/recording restrictions (36% of cases), compared in other conditions. The MoonBot condition led to relatively more thinking about data upload (occurrence in 54% of responses) and privacy concerns (46% of responses) compared to the other conditions. On the other hand, the scooter condition seemed to elicit the most thinking about safety and emergency stop functions (both 38% occurrence), at least by a small margin compared to other conditions. Once a scenario was laid over the core ordinance, participants seemed to think more about contact information requirements (38% of cases for both MoonBot and scooter, compared to 18% in the base ordinance condition).

4 DISCUSSION

Our findings can be contextualized within the broader landscape of research on human-robot interaction, autonomous systems, and urban mobility. Although the pilot focused on a small set of beginning mock scenarios, future work in related domains can potentially use a similar approach to provide additional insights and contribute to a more comprehensive understanding of user perceptions of and preferences for service robots (and beyond). The ANOVA results showed no statistically significant differences across scenarios in terms of safety performance, preservation of privacy, and autonomous delivery system deployment likelihood, but trending in the data might inform future scenario design and evaluation. For example, having additional scenario context tended to raise concern levels for safety and privacy.

Several common themes emerged in the critique of the current ordinance, such as restrictions on vehicle characteristics, the importance of clear contact information, and concerns about data uploading and data privacy. These themes have significant design implications for the development of autonomous systems and related regulations. Our results hint that service robot designers should prioritize ensuring robust data privacy measures, providing explicit contact information to address user concerns and enhance overall system safety, and incorporating clear and accessible emergency stop mechanisms (as a few examples). The results also highlight the need for standardized language and clear definitions in city ordinances related to autonomous systems. We noticed while parsing the qualitative data that ambiguities in terminology, such as the interpretation of "the cloud," can lead to confusion and misinterpretation of regulations.

There are clear limitations to the work, both in the small sample size and the relative simplicity of the questions and responses. However, as indicated by Table 2, a concrete ordinance and scenario still can elicit specific and fitting areas for discussion and critique. This demonstrates the promise of even such relatively simple ordinances and scenarios to serve as a potential mechanism for initiating focused and apt discussions in the coming waves of important robot policymaking tasks. In future steps, we plan to generate ordinances and scenarios with closer ties to real-world laws and cases, in addition to collecting data from a larger sample with less background training in robotics.

Overall, this work offers valuable insights on perceptions and regulation of autonomous systems in human-populated settings such as dense urban environments. By examining responses to different scenarios, we shed light on the factors that influence public attitudes towards safety, privacy, and deployment surrounding autonomous delivery systems.

Table 2: Overview of key emergent themes in the qualitative analysis of changes respondents would make to the ordinance, with frequencies for each condition and example quotes. Themes are ordered from highest to lowest frequency of occurrence.

	Base Ordinance	MoonBot	Scooter	Illustrative Quote
Vehicle Characteristics	8	2	8	"Reduce size and mass to be less dangerous and obtrusive"
Data Upload to Cloud	4	7	2	"Uploading to cloud allowed with monitoring"
Contact Information Requirements	2	5	5	"Contact details provided [for a] human agent officially authorized to deal with safety concerns and accidents"
Privacy Concerns	3	6	3	"Data privacy regulations need extra clarification, different wording"
Safety Regulations	4	1	5	"[The ordinance needs] clearer safety regulations"
Mapping and Recording Restrictions	4	2	3	"Have stronger limits on data recording and data recording management"
Emergency Stop Mechanisms	3	1	5	"The vehicle needs a clearly marked & accessible emergency stop that a person physically next to [the robot] can trigger"

5 ACKNOWLEDGEMENTS

We thank Ariana Hill and Sydney Ayers for their data collection support. This work was funded in part by NSF grants NRI 2024872, 2024673, and 2024643.

A BASE ORDINANCE

All participants, across the three conditions, saw the following base ordinance first:

"This is a made-up example of a city ordinance written by a medium-sized city USA with an older, dense downtown surrounded by suburbs. Recently, a growing number of younger tech industry people have been moving into downtown. The (mostly older) established residents are concerned about new technology disrupting the 'look and feel' of their downtown.

"Ordinance:

(1) Any autonomous vehicle operating in downtown must be less than 100lbs, travel at less than 15 mph, and be narrower than 2'
(2) All autonomous vehicles need to have clear markings (phone number or website) that identity the operator of the vehicle
(3) No vehicle shall record or upload video directly "to the cloud"
(4) No autonomous vehicles are allowed to stop within 15' of a ramp or other accessible entrance to a building"

B MOONBOT CONDITION

For participants in the MoonBot condition, the further text read:

"Company MoonBot is getting ready to deploy its robotic fleet of 2'×2' 60lb delivery bots. In preparation for deployment, MoonBot uses teleoperation to drive robots through the city, recording sensor readings (including cameras) of the city. This video is uploaded to the cloud where, through a partially automated, partially manual process it is annotated to create a detailed abstract 3D map of the city that marks all of the sidewalk edges, curb cuts, driveways, sign posts, etc. This information is used by the robots both to tell where they are, avoid restricted areas, and to identify places where a human should step in to check the surroundings before proceeding (e.g., busy intersections).

"Technical note: The individual robots do not have enough memory storage or processing power to build the map on the robot itself, and then upload the map."

C SCOOTER CONDITION

For participants in the Scooter condition, the further text read:

"After the ordinance was put into effect the following happened:

"A motorized scooter hit a delivery robot when it veered off of the sidewalk and into their path. Both the motor scooter and robot were damaged. The owner of the motor scooter called the number on the side, but just got the automated menu option for food delivery. The robot was too heavy to move far (and they didn't want someone to think they were stealing it) but they did manage to drag it out of the driveway it was blocking and into the bicycle lane, where it sat for a day."

REFERENCES

[1] Hassan Ali, Sulaiman Riza, Abdulhak Mansoor, and Kahtan Hasan. 2021. Towards User-Centric Explanations for Explainable Models: A Review. *Journal of Information System and Technology Management* 6 (2021).
[2] Annie I. Antón, Travis D. Breaux, Dimitris Karagiannis, and John Mylopoulos. 2008. First International Workshop on Requirements Engineering and Law (RELAW). In *Proc. of the IEEE International Requirements Engineering Conference*. 1–4.
[3] Anna Bobkowska and Magdalena Kowalska. 2010. On efficient Collaboration between Lawyers and Software Engineers When Transforming Legal Regulations to Law-Related Requirements. In *Proc. of the IEEE International Conference on Information Technology (ICIT)*. 105–109.
[4] Ryan Calo. 2019. How We Talk About AI (and Why It Matters). In *Proc. of the AAAI/ACM Conference on AI, Ethics, and Society (AIES '19)*. 1.
[5] Hyesun Choung, Prabu David, and Arun Ross. 2023. Trust in AI and Its Role in the Acceptance of AI Technologies. *International Journal of Human-Computer Interaction* 39, 9 (2023), 1727–1739.
[6] Leilani H. Gilpin, David Bau, Ben Z. Yuan, Ayesha Bajwa, Michael Specter, and Lalana Kagal. 2018. Explaining Explanations: An Overview of Interpretability of Machine Learning. In *Proc. of the IEEE International Conference on Data Science and Advanced Analytics (DSAA)*. 80–89.
[7] Matthew Rueben, Frank J Bernieri, Cindy M Grimm, and William D Smart. 2017. Framing effects on privacy concerns about a home telepresence robot. In *Proc. of the ACM/IEEE International Conference on Human-Robot Interaction*. 435–444.
[8] Kong Siu-Cheung, Cheung William Man-Yin, and Zhang Guo. 2023. Evaluating an Artificial Intelligence Literacy Programme for Developing University Students' Conceptual Understanding, Literacy, Empowerment and Ethical Awareness. *Educational Technology & Society*, 16–30.
[9] Auriel Washburn, Akanimoh Adeleye, Thomas An, and Laurel D Riek. 2020. Robot errors in proximate HRI: How functionality framing affects perceived reliability and trust. *ACM Transactions on Human-Robot Interaction (THRI)* 9, 3 (2020), 1–21.