Minimal Effort, Maximum Impact: Fighting Humanoids and Interface Design

Joel Wester
joel.wester@di.ku.dk
University of Copenhagen
Copenhagen, Denmark

Eike Schneiders
eike.schneiders@soton.ac.uk
University of Southampton
Southampton, United Kingdom

Jacy Reese Anthis
anthis@uchicago.edu
University of Chicago
Chicago, USA

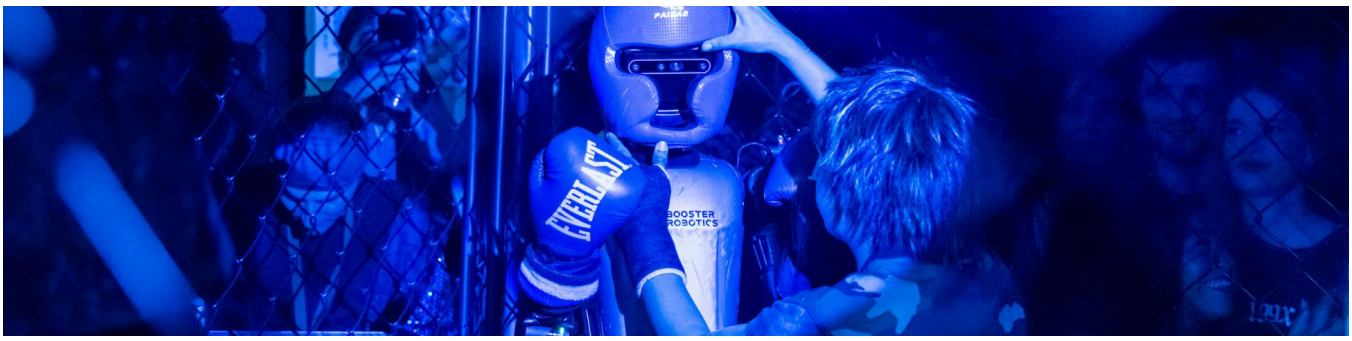

Figure 1: Screenshot from the Ultimate Fighting Bots website (https://ufb.gg). The platform allows users to watch humanoid fights or participate in proxy-based fighting: "*Ultimate Fighting Bots is the world's first robot combat league where humans pilot robots in real-time battles, live from underground arenas in San Francisco or from anywhere on the planet.*"

Abstract

From televised *Robot Wars* competitions to modern warfare with teleoperated drones, robot combat has long been of interest in entertainment and practical application. These systems are often anthropomorphised, such as being given humanlike names, and they are increasingly capable of real-world action, often operating at least semi-autonomously. These capabilities have given humans, sometimes everyday people, the ability to enact their will through simplified interfaces in which complex and consequential actions are triggered by a single button press. The disconnect between user input and robotic action raises pressing questions for interface design, in robot combat but also other contexts such as autonomous coding agents in which humans can enact significant real-world consequences with trivial effort. Prior HCI research has addressed some of these questions, such as work on drones that surfaced ethical challenges related to proxy-based control and responsibility. In this paper, we explore these dynamics by reflecting on fighting humanoid robots, a case that epitomizes the intersection of anthropomorphism, robot autonomy, human responsibility, and interface design. Drawing on these reflections, we foreground interactional friction as a theoretical lens to consider how doing more for less is not always better. By developing HCI theory in this domain, we can support human wellbeing as human-like artificial intelligence and robots move out of the boxing ring and into everyday, real-world contexts.

Keywords

Robot, Humanoids, Fight, Combat, Design

ACM Reference Format:
Joel Wester, Eike Schneiders, and Jacy Reese Anthis. 2026. Minimal Effort, Maximum Impact: Fighting Humanoids and Interface Design . In *Proceedings of the CHI 2026 Workshop: Ethics at the Front-End: Responsible User-Facing Design for AI Systems (CHI '26), April 13–17, 2026, Barcelona, Spain.* ACM, New York, NY, USA, 4 pages. https://doi.org/xx.xxxx/xxxxxxx.xxxxxxx

1 Backdrop

Robots have long captured public attention, and this fascination is especially pronounced when robots take on humanoid form. Human-like embodiment shapes how people perceive robots, often leading to heightened expectations about their cognitive, physical, or social capabilities [14, 23]. When a robot looks human, people intuitively ascribe to it a human-like range of motion, dexterity, intelligence, and competence—even when these capabilities are only partially realised [31]. The power of human-likeness deserves careful consideration as youth interact with increasingly human-like systems, and powerful artificial intelligence (AI) systems such as language models are already endowed with human-like voices [e.g., 6] and could soon be endowed with human-like virtual avatars or humanoid robotic bodies.

In recent years, humanoid robots have increasingly entered competitive and sport-like contexts, including the world's first Humanoid Robot Games, hosted by China in 2025 [32]. This context has two emerging paradigms. First, autonomous robots compete

CHI '26, Barcelona, Spain

ACM ISBN xxx-x-xxxx-xxxx-x/xx/xx
https://doi.org/xx.xxxx/xxxxxxx.xxxxxxx

against one another as proxies for their human development teams, directly translating software and hardware achievements into physical performance. Second, remotely controlled robots, such as racing drones or fighting humanoids, are proxies for live human–human competition between the operators, where outcomes are mediated through front-end interfaces and robotic bodies.

These developments raise pressing questions for interaction and interface design. When complex and consequential actions are accessible through simple physical controls or language-based commands, design choices can determine how users perceive those consequences, how much responsibility users feel for consequences, and how users are affected by dark patterns [10] and other drivers of user action. This concern resonates with broader debates in HCI about design friction, particularly in the context of AI and robotics [e.g., 17], debates that have become increasingly urgent as large language models and conversational interfaces transform how people interact with human-like AI systems. Against this backdrop, fighting humanoids provide a timely and provocative context for reflecting on how interface design shapes agency and responsibility in human-AI interaction.

In this paper, we first anchor our discussion in popular conceptualisations of robots as shaped by movies, social media, and science fiction narratives. We then turn to the related context of drones, in which prior HCI work has characterised proxy-based control and interface design and its consequences. We conceptualise these empirical phenomena primarily in theories of design friction, distilling its relevance for the interface design of fighting humanoids. Finally, we articulate a provocation to stimulate discussion around ethical interface design and potential directions for future work on fighting humanoids and other domains in which advanced technology connects trivial user actions to real-world action.

2 Violence, Robot Narratives, and Drones

Popular culture has long been fascinated with robots engaged in combat. Robot combat framed as entertainment includes televised competitions such as *Robot Wars*[1] and *BattleBots*[2], where machines with a variety of weapons (e.g., chainsaws and flamethrowers) are purposefully designed to damage or destroy one another, and large-scale spectacles such as the Giant Robot Duel (marketed as the world's first human-piloted, giant-robot fight) [5].

Robot violence has also been a recurring part of science fiction and media narratives. Films such as *Real Steel* depict humanoid robots with individual human operators, fighting as proxies for the rivalry between human individuals, while social media has popularised aggressive robot performances, including robot dogs framed as intimidating or militaristic [37]. Robot dogs are a particularly interesting example because, in addition to biological dogs as companion animals, they have been a popular and beloved toy, with the AIBO robot dog even being honoured in funerals by Japanese users [7]. HCI research has shown that people readily attribute mental states to robot dogs [16] and show moral concern for their welfare [30]. Across these representations, robots are repeatedly positioned as what could be called 'vessels for violence', agents that enact human will in the physical domain, enabling users to inflict physical harm on robots—and potentially on others—with incredible ease.

Taken together, these portrayals—be it in movies or at displays such as the CES [33]—suggest a form of vessels [21], in which robots become stand-ins for individual human aggression and competition. Importantly, the popularity of fighting robots does not remain confined to entertainment contexts; rather, it risks shaping broader expectations of robotic systems across domains, influencing how robots are designed and perceived in non-combat settings, as evidenced by AI spillover effects, the tendency of attitudes towards one system to spill over to attitudes towards others [27] and shifting perceptions of AI consciousness across different physical and digital systems [24]. This makes fighting robots a critical site for exploring the role of proxy interface design and effortfulness in contemporary human–robot interaction.

However, as research is scarce on interface design for steering robots—instead typically focusing on human-robot interaction, such as their relative autonomy [15, 20], we draw from HCI research on drones as a closely related domain. Drone systems are increasingly deployed in complex physical environments (e.g., delivery tasks [26]) and are frequently operated through proxy-based control. While early work on drones focused primarily on navigation and autonomy, interface design has gained increasing attention in recent years. For example, Bahodi and colleagues have explored the role of interface design in coordinating drone fleets using battery visualisation [3], supporting surveillance tasks through autonomy design [2], and enabling supervisory control with different types of warnings appearing in the interface to reduce cognitive load [4]. Related, Cauchard et al. [12] explored interaction techniques for drones, highlighting that people used metaphorical gestures like putting up a hand to signal to the drone to stop—related to recent work by La Delfa et al. [22] on drones and metaphors. Boonyard et al. [9] instead investigated interface design for firefighting drones operating in high-risk environments. Ahlskog et al. [1] explored how interface design shapes trust in drone systems, highlighting the complexity around the combination of monitoring and manually controlling drones in more critical situations. One could easily imagine how principles from current drone practices could be transferred to the domain of fighting humanoids. For instance, a drone crash might be considered inappropriate drone flying, resulting in a penalty in the form of either enforced down-time or time to repair, and flying a drone until the battery runs out might result in cooldown time, enforced by a rule or battery recharge, prior to re-start. Similarly, fighting robots could—like drones—make use of eligibility criteria, such as age, experience, robot licenses, or previous drone-flying behaviour before participation. These limitations to how one might engage with fighting robots can be implemented in the form of deliberate design frictions—both soft and hard, which we outline in Table 1.

Despite this growing body of work, significant gaps remain in understanding how interface design mediates human responsibility and ethical engagement in proxy-controlled robots. Brown [11] highlight several contemporary challenges in 'military HCI', including issues of accountability, distancing, and moral disengagement—issues that emerge outside the military context with fighting humanoids.

[1] https://www.bbc.co.uk/programmes/p03h7r7n
[2] https://battlebots.com

Table 1: Frictions in front-end command design across phases of proxy fighting humanoid interaction.

Friction	Before fights	During fights	After fights
Soft frictions			
● wait	➤ Force queueing	➤ Delay high-impact commands	➤ Enforce pauses
● explain	➤ Potential consequences of commands	➤ Cues to illustrate command impact	➤ Impact summary
● warn	➤ Warnings prior to enabling controls	➤ Warnings at escalation thresholds	➤ Warnings at repeated patterns
● reflect	➤ Communicate responsibilities	➤ Check-ins at command escalation	➤ Reflection on decisions
Hard frictions			
■ gate	➤ Eligibility, age, and credentials	➤ Permission for advanced moves	➤ Gated re-entry after violations
■ limit	➤ Limit match and intensity levels	➤ Cap frequency or intensity of actions	➤ Block rapid matches
■ refuse	➤ Refuse access to certain modes	➤ Deny unsafe commands outright	➤ Refuse re-entry under conditions
■ defer	➤ Delay until conditions are met	➤ Postpone actions pending system state	➤ Cooldown before re-entry

3 Designing Interface Friction

Across product and service design, various mechanisms are used to introduce friction and safeguard users: cigarettes carry warning labels, gambling platforms implement stop or cooling-off mechanisms, firearms are subject to restrictions and waiting periods, and prescription medications require consultation with a medical professional. In HCI, research has explored friction as a design resource to positively influence interactions with technology. This includes slowing down interactions to positively influence mindfulness [28], time delays to decrease smartphone overuse [18], negative interface cues to impact harassment perceptions [25], lock-out mechanisms to support meaningful system usage [19], or emotion regulation on social media [13]. While users can be supported by frictions, frictions can leave negative impressions [8]. For examples, users can feel as though their autonomy is not being respected as companies or governments appear to be deliberately impairing their ability to seek out resources and entertain themselves.

Recently, there has been a surge of interest in how large language models can be configured to handle inappropriate or irrelevant user requests; refusal of a user instruction could be considered absolute design friction, though the model could also offer alternative, lower-friction interaction possibilities. This work includes technical approaches such as abstention [34] and refusal mechanisms [29, 36, 38], as well as more HCI-focused approaches that examine how denials should be designed and communicated to users to optimise the user experience [35]. Research on design friction in interface design is increasing. However, AI-powered robots and proxy fighting humanoids have not yet been included in such research.

4 Looking Ahead

Although current humanoid fighting robots may appear farcical in their appearance and movement, there is little reason to assume they will remain so. Given the rapid pace of technological development in recent years, it is easy to envision a future in which fighting humanoids and other proxy robotic systems become significantly more capable, both in terms of physical movement and the range of actions available to users via commands. Proxy fighting may extend beyond simple button-based interfaces to incorporate alternative forms of control, such as conversational commands, augmented reality, brain-computer interfaces, or embodied controllers.

Correspondingly, we conclude with Table 1, which presents a speculative and extensive list of potential soft and hard frictions, some of which we highlighted in Section 2—that we hope will stimulate discussion at the workshop. Future research on these frictions could include co-designing policies and practices with users and other stakeholders, identifying the pros and cons of the various frictions, and assessing which frictions have historical precedent of benefit or harm in other technologies.

References

[1] Johanna Ahlskog, Maria-Theresa Bahodi, Artur Lugmayr, and Timothy Merritt. 2024. Fostering Trust Through User Interface Design in Multi-Drone Search and Rescue. In *Proceedings of the Second International Symposium on Trustworthy Autonomous Systems* (Austin, TX, USA) *(TAS '24)*. Article 2, 11 pages. doi:10.1145/3686038.3686052

[2] Maria-Theresa Bahodi, Maja Hornbæk Kristensen, Niels van Berkel, Mikael Skov, Nicolai Brodersen Hansen, and Timothy Merritt. 2026. Eyes from Above: Co-designing a Multi-drone System for Enhanced Surveillance of Critical Infrastructure. In *Social Robotics + AI*. 447–474.

[3] Maria-Theresa Bahodi, Nathan Lau, Niels van Berkel, Kasper Andreas Rømer Grøntved, Mikael B. Skov, and Timothy Merritt. 2026. Before It Falls: Supporting Drone Fleet Management Through Battery Visualizations. In *Human-Computer Interaction – INTERACT 2025*. 347–370.

[4] Maria-Theresa Bahodi, Niels van Berkel, Mikael Skov, and Timothy Merritt. 2024. Show Me What's Wrong: Impact of Explicit Alerts on Novice Supervisors of a Multi-Robot Monitoring System. In *Proceedings of the Second International Symposium on Trustworthy Autonomous Systems* (Austin, TX, USA) *(TAS '24)*. Article 4, 17 pages. doi:10.1145/3686038.3686069

[5] Ollie Barder. 2017. The Giant Robot Duel Between Japan And The U.S. Turned Out To Be Rather Disappointing. *Forbes* (18 Oct 2017). https://www.forbes.com/sites/olliebarder/2017/10/18/the-giant-robot-duel-between-japan-and-the-u-s-turned-out-to-be-rather-disappointing/ Accessed: 2026-01-13.

[6] Sarah Barrington, Emily A. Cooper, and Hany Farid. 2025. People Are Poorly Equipped to Detect AI-powered Voice Clones. *Scientific Reports* 15, 1 (March 2025), 11004. doi:10.1038/s41598-025-94170-3

[7] James Birch. 2018. AIBO Robot Dogs Given Buddhist Funeral in Japan. https://www.nationalgeographic.com/travel/article/in-japan–a-buddhist-funeral-service-for-robot-dogs.

[8] Kerstin Bongard-Blanchy, Arianna Rossi, Salvador Rivas, Sophie Doublet, Vincent Koenig, and Gabriele Lenzini. 2021. "I am Definitely Manipulated, Even When I am Aware of it. It's Ridiculous!" - Dark Patterns from the End-User Perspective. In *Proceedings of the 2021 ACM Designing Interactive Systems Conference* (Virtual Event, USA) *(DIS '21)*. 763–776. doi:10.1145/3461778.3462086

[9] Chandhawat Boonyard, Christophe Jouffrais, Jessica R. Cauchard, and Anke M. Brock. 2025. Firefighting with Drone Assistance: User Needs and Design Considerations for Thailand. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Article 432, 18 pages. doi:10.1145/3706598.3714172

[10] Harry Brignull. 2010. Dark Patterns: Dirty Tricks Designers Use to Make People Do Stuff. https://90percentofeverything.com/2010/07/08/dark-patterns-dirty-tricks-designers-use-to-make-people-do-stuff/index.html.

[11] Barry Brown. 2025. Designing for War? Ethical and Practical Challenges to a Military HCI. In *Proceedings of the Sixth Decennial Aarhus Conference: Computing X Crisis (AAR '25)*. 153–165. doi:10.1145/3744169.3744172

[12] Jessica R. Cauchard, Jane L. E, Kevin Y. Zhai, and James A. Landay. 2015. Drone & me: an exploration into natural human-drone interaction. In *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing* (Osaka, Japan) *(UbiComp '15)*. 361–365. doi:10.1145/2750858.2805823

[13] Katie Davis, Rotem Landesman, Jina Yoon, JaeWon Kim, Daniela E Munoz Lopez, Lucia Magis-Weinberg, and Alexis Hiniker. 2025. "You Go Through So Many Emotions Scrolling Through Instagram": How Teens Use Instagram To Regulate Their Emotions. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Article 1179, 16 pages. doi:10.1145/3706598.3713844

[14] Nathaniel Dennler, Changxiao Ruan, Jessica Hadiwijoyo, Brenna Chen, Stefanos Nikolaidis, and Maja Matarić. 2023. Design Metaphors for Understanding User Expectations of Socially Interactive Robot Embodiments. *J. Hum.-Robot Interact.* 12, 2, Article 21 (April 2023), 41 pages. doi:10.1145/3550489

[15] Cloe Z. Emnett, Terran Mott, and Tom Williams. 2024. Using Robot Social Agency Theory to Understand Robots' Linguistic Anthropomorphism. In *Companion of the 2024 ACM/IEEE International Conference on Human-Robot Interaction (HRI '24)*. Association for Computing Machinery, New York, NY, USA, 447–452. doi:10.1145/3610978.3640747

[16] Batya Friedman, Peter H. Kahn, and Jennifer Hagman. 2003. Hardware Companions?: What Online AIBO Discussion Forums Reveal about the Human-Robotic Relationship. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, Ft. Lauderdale Florida USA, 273–280. doi:10.1145/642611.642660

[17] Rachael Garrett, Patrick Brundell, Simon Castle-Green, Kat Hawkins, Paul Tennent, Feng Zhou, Airi Lampinen, Kristina Höök, and Steve Benford. 2025. Friction in Processual Ethics: Reconfiguring Ethical Relations in Interdisciplinary Research. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. ACM, Yokohama Japan, 1–15. doi:10.1145/3706598.3714123

[18] Luke Haliburton, David Joachim Grüning, Frederik Riedel, Albrecht Schmidt, and Nađa Terzimehić. 2024. A Longitudinal In-the-Wild Investigation of Design Frictions to Prevent Smartphone Overuse. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI '24)*. Article 243, 16 pages. doi:10.1145/3613904.3642370

[19] Jaejeung Kim, Hayoung Jung, Minsam Ko, and Uichin Lee. 2019. GoalKeeper: Exploring Interaction Lockout Mechanisms for Regulating Smartphone Use. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 3, 1, Article 16 (March 2019), 29 pages. doi:10.1145/3314403

[20] Stephanie Kim, Jacy Reese Anthis, and Sarah Sebo. 2024. A Taxonomy of Robot Autonomy for Human-Robot Interaction. In *Proceedings of the 2024 ACM/IEEE International Conference on Human-Robot Interaction*. ACM, Boulder CO USA, 381–393. doi:10.1145/3610977.3634993

[21] Andreas Krieg and Jean-Marc Rickli. 2018. Surrogate Warfare: The Art of War in the 21st Century? *Defence Studies* 18, 2 (April 2018), 113–130. doi:10.1080/14702436.2018.1429218

[22] Joseph La Delfa, Rachael Garrett, Airi Lampinen, and Kristina Höök. 2024. How to Train Your Drone: Exploring the umwelt as a design metaphor for human-drone interaction. In *Proceedings of the 2024 ACM Designing Interactive Systems Conference* (Copenhagen, Denmark) *(DIS '24)*. 2987–3001. doi:10.1145/3643834.3660737

[23] Ali Ladak, Jamie Harris, and Jacy Reese Anthis. 2024. Which Artificial Intelligences Do People Care About Most? A Conjoint Experiment on Moral Consideration. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. ACM, Honolulu HI USA, 1–11. doi:10.1145/3613904.3642403

[24] Ali Ladak, Matti Wilks, Steve Loughnan, and Jacy Reese Anthis. 2025. Robots, Chatbots, Self-Driving Cars: Perceptions of Mind and Morality Across Artificial Intelligences. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Article 1101, 19 pages. doi:10.1145/3706598.3713130

[25] Song Mi Lee, Andrea K. Thomer, and Cliff Lampe. 2022. The Use of Negative Interface Cues to Change Perceptions of Online Retributive Harassment. *Proc. ACM Hum.-Comput. Interact.* 6, CSCW2, Article 334 (Nov. 2022), 23 pages. doi:10.1145/3555226

[26] Shiva Nischal Lingam, Rutger Verstegen, Sebastiaan Martinus Petermeijer, and Marieke Martens. 2024. Design of Delivery Drone and the Interaction: A Public User Perspective. In *Proceedings of the 12th International Conference on Human-Agent Interaction* (Swansea, United Kingdom) *(HAI '24)*. 356–357. doi:10.1145/3687272.3690880

[27] Aikaterina Manoli, Janet V. T. Pauketat, and Jacy Reese Anthis. 2025. The AI Double Standard: Humans Judge All AIs for the Actions of One. *Proc. ACM Hum.-Comput. Interact.* 9, 2, Article CSCW185 (May 2025), 24 pages. doi:10.1145/3711083

[28] Teale W. Masrani, Jack Jamieson, Naomi Yamashita, and Helen Ai He. 2023. Slowing it Down: Towards Facilitating Interpersonal Mindfulness in Online Polarizing Conversations Over Social Media. *Proc. ACM Hum.-Comput. Interact.* 7, CSCW1, Article 90 (April 2023), 27 pages. doi:10.1145/3579523

[29] John Mavi, Diana Teodora Găitan, and Sergio Coronado. 2025. From Rogue to Safe AI: The Role of Explicit Refusals in Aligning LLMs with International Humanitarian Law. arXiv:2506.06391 [cs.CY]

[30] Jon Mooallem. 2015. Mr. Know-It-All: Is It OK to Kick a Robot? *Wired* (2015).

[31] Eike Schneiders, Eleftherios Papachristos, and Niels van Berkel. 2022. The Effect of Embodied Anthropomorphism of Personal Assistants on User Perceptions. In *Proceedings of the 33rd Australian Conference on Human-Computer Interaction* (Melbourne, VIC, Australia) *(OzCHI '21)*. Association for Computing Machinery, New York, NY, USA, 231–241. doi:10.1145/3520495.3520503

[32] Sean Sinico. 2025. World's First Humanoid Robot Games Begin in China. https://www.dw.com/en/worlds-first-humanoid-robot-games-begin-in-china/a-73652714.

[33] Sky News. 2026. *Why Is a Sky Reporter Having a Punch-Up With a Robot at CES?* https://news.sky.com/video/why-is-a-sky-reporter-having-a-punch-up-with-a-robot-at-ces-13492685 Video report; accessed: 2026-01-22.

[34] Bingbing Wen, Jihan Yao, Shangbin Feng, Chenjun Xu, Yulia Tsvetkov, Bill Howe, and Lucy Lu Wang. 2025. Know Your Limits: A Survey of Abstention in Large Language Models. *Transactions of the Association for Computational Linguistics* 13 (06 2025), 529–556. doi:10.1162/tacl_a_00754

[35] Joel Wester, Tim Schrills, Henning Pohl, and Niels van Berkel. 2024. "As an AI language model, I cannot": Investigating LLM Denials of User Requests. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI '24)*. Article 979, 14 pages. doi:10.1145/3613904.3642135

[36] Michele Zappavigna. 2025. 'I'm sorry Dave, I'm afraid I can't do that': Moral regulation in refusals by LLM chatbots. *New Media & Society* 0, 0 (2025). doi:10.1177/14614448251356686

[37] Mihir Zaveri. 2021. N.Y.P.D. Robot Dog's Run Is Cut Short After Fierce Backlash. *The New York Times* (April 2021).

[38] Mingqian Zheng, Wenjia Hu, Patrick Zhao, Motahhare Eslami, Jena D. Hwang, Faeze Brahman, Carolyn Rose, and Maarten Sap. 2025. Let Them Down Easy! Contextual Effects of LLM Guardrails on User Perceptions and Preferences. arXiv:2506.00195 [cs.CL]