

Detecting Emerging Challenges in Social Sidewalk Navigation

Vroon, Jered; Zhou, Y.; Rusak, Z.

Publication date

2020

Document Version

Final published version

Citation (APA)

Vroon, J., Zhou, Y., & Rusak, Z. (2020). *Detecting Emerging Challenges in Social Sidewalk Navigation*. Paper presented at Urban AI: Formulating an agenda for the interdisciplinary research of artificial intelligence in cities, Eindhoven, Netherlands.

Important note

To cite this publication, please use the final published version (if applicable). Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights. We will remove access to the work immediately and investigate your claim.

Detecting Emerging Challenges in Social Sidewalk Navigation

Jered Vroon

Knowledge & Intelligence Design
Industrial Design Engineering
Delft University of Technology
the Netherlands
j.h.vroon@tudelft.nl

Yunzhong Zhou

Knowledge & Intelligence Design
Industrial Design Engineering
Delft University of Technology
the Netherlands
y.zhou-13@tudelft.nl

Zoltán Rusák

Knowledge & Intelligence Design
Industrial Design Engineering
Delft University of Technology
the Netherlands
z.rusak@tudelft.nl

Abstract

When mobile urban robots will share the sidewalk with people, the resulting interactions can cause unexpected undesirable outcomes to emerge – from people running away scared to people deliberately teasing and harassing such systems. How can we design such AI systems to aptly handle the unexpected? Directly anticipating and/or detecting these kinds of situations will inherently be unreliable; they are unexpected, after all. And yet, there exists a very clear signal for social slip-ups: the emotional response of people. We thus argue that such systems need to be imbued with a capacity to interpret the socio-emotional reactions to their own behavior.

Author Keywords

Urban Robotics; Social Navigation; Social Signal Processing; Emergent Behavior

CCS Concepts

•**Computing methodologies** → *Cognitive robotics*; •**Human-centered computing** → **Interaction design theory, concepts and paradigms**;

Introduction

Social Sidewalk Navigation will be fundamental to the success of the upcoming field of Urban Robotics. Companies and municipalities are more and more putting robots in ur-

ban environments, to fulfil functions ranging from last-mile delivery to garbage collection and from handing out flyers to guiding people. For the most, people seem quite willing to accept such robots and share the sidewalk with them: “once we put one of the robots out there onto the sidewalk, [...] the vast majority of the public didn’t pay any attention whatsoever to the robots, even those seeing it for the first time” [7].

And yet, introducing such systems into the ‘wild’ also causes wildly unexpected reactions from people to emerge, which can have undesirable and unanticipated consequences (**emerging challenges**). Consider, for example, wheelchair users getting stuck on the street because robots occupy the curb¹. Or what to think of the robot handing out flyers being harassed and assaulted by children [3] and other cases of robots being bullied in the wild [4]?

This raises the question: *How can we design the social AI for such urban robotics to ably handle such emerging challenges?*

We can’t reliably avoid these challenges, because they are unknown upfront. While it will often be possible to patch them after they become apparent, such a patch will necessarily be applied only *after* the problem has arisen the first time. In addition, this may require an unfeasible amount of patches, given the complex and ever-changing dynamics of our social interaction; patches may not even be compatible with each other.

Even detecting when a system runs into such a challenge will be difficult, because of their unanticipated nature – how can we train a system to detect something unexpected?

¹An account of one person, Twitter user Emily E. Ackerman, experiencing this, can be found on twitter.com/EmilyEAckerman/status/1186363305851576321

We argue that the only way to detect when such social problems emerge will thus be from the socio-emotional reactions of the people involved. It is impossible to reliably detect or predict when an urban robot will make mistakes, but it may be doable to detect when people are upset with you (**detecting emerging challenges**).

Patching in Socially Aware Navigation

There is an extensive body of work on Socially Aware Navigation, which has been successfully deployed in urban robotics and beyond; many approaches aim to navigate such that the system avoids getting closer to people than certain set distances (derived from notions of personal space, such as Hall’s proxemics [11] and others [14]).

As with any set of rules, these approaches always run into exceptional situations not sufficiently covered by the rules. For example, in earlier work, a failure to properly predict human motion could result in erratic paths being executed [9] and the aforementioned case of a robot blocking someone in a wheelchair on the street. To further complicate this, people will often actively respond to the navigation behavior of such systems (e.g. [13, 10, 12]), which can make it even harder to anticipate what interactions will emerge.

To some extent, these exceptional situations can of course be ‘patched’ by updating the software after they arise and are picked up by the responsible engineers and designers. This is what happened to resolve most of the problems mentioned above: for example, the robot bullied by children was patched to seek the safety of their parents [3] and the robot with the erratic paths was patched to slow down when needed [9]. Of course there is a limit to how many such band-aids can be used; some patches will cause new problems, and some situations will have conflicting needs in terms of the patches that apply. In other words, a more fun-

damental solution is needed, ideally one that would allow for a solution to be found on the spot.

Detecting Socio-Emotional Reactions to Behavior

Human emotion is expressed through – and can thus be detected from – a wide range of social signals and cues [15], including facial expressions, verbal cues, body language [6], and gait/trajectories [5]. These latter two are most applicable to the context of social sidewalk navigation, because of their relative robustness against occlusions, effectiveness at a longer range, and because they require less privacy-sensitive information to function.

Recent work has attempted to take such socio-emotional reactions into account in social navigation, e.g. by assuming that negative emotions are caused by a robot violating peoples' comfort zones [2, 1], or that robots should give pedestrians more personal space when negative emotions are detected [8].

Though these approaches work well in their context, they still leave a lot of the expressivity of such socio-emotional reactions untapped. Other aspects of robot behavior, such as movement speed or obstructing a pedestrian's path, could also trigger negative emotions result in unexpected human responses. To tackle emerging challenges, a mobile robot should not only detect whether the negative emotion is caused by its own behavior [16] but also distinguish which aspect of its behavior is the cause, be it its speed, its distance to the pedestrian, its approaching angle, or another aspect.

Towards Responsive Social Sidewalk Navigation

Mobile robots navigating in urban environments will have to cope with emerging interactions and the unexpected

challenges that can arise from them. This means that, as argued above, in order to respond appropriately, such systems will need to detect the socio-emotional reactions of pedestrians to their own behavior. We briefly discussed recent advances in social signal processing that suggest that such detections may soon be within reach.

Though these kinds of detections will thus be a necessary starting point, even with them it will still be a (design) challenge to handle emerging challenges appropriately. One approach would be to include a human operator in the loop, who can assume manual control when the system detects it is causing a negative reaction in nearby pedestrians. An alternative would be to have the system try to resolve the socio-emotional reaction autonomously, e.g. by using it as feedback for an online reinforcement learning system that tries to find/learn the appropriate behaviour.

Beyond urban robotics, these arguments may well also apply to the broader topic of urban AI. While the physicality of robotics gives a more direct urgency to responding to people's socio-emotional reactions, other systems may well benefit from a capacity to properly handle socio-emotional reactions.

And then, maybe, one day, if we are offended or otherwise inconvenienced by that robot in our city, it can detect that socio-emotional reaction and thus respond in a more understanding way.

REFERENCES

- [1] Aniket Bera, Tanmay Randhavane, and Dinesh Manocha. 2019a. Improving Socially-aware Multi-channel Human Emotion Prediction for Robot Navigation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*. 21–27.

About the Authors

Jered Vroon is a post-doc who investigates the dynamic interactions between people and intelligent artefacts. His background is in cognitive AI and social robotics.

Yunzhong Zhou is a PhD student exploring the detection of social feedback cues. His background is in mechanical engineering and optimization.

Zoltán Rusák is an assistant professor who works on design methods for complex situations with emergent behaviors. His background is in cyber-physical systems, computer-aided design, and mechanical engineering.

- [2] Aniket Bera, Tanmay Randhavane, Rohan Prinja, Kyra Kapsaskis, Austin Wang, Kurt Gray, and Dinesh Manocha. 2019b. The emotionally intelligent robot: Improving social navigation in crowded environments. *arXiv preprint arXiv:1903.03217* (2019).
- [3] Dražen Brščić, Hiroyuki Kidokoro, Yoshitaka Suehiro, and Takayuki Kanda. 2015. Escaping from children’s abuse of social robots. In *Proceedings of the tenth annual acm/ieee international conference on human-robot interaction*. 59–66.
- [4] Merijn Bruijnes, Robby van Delden, and Jered Vroon. 2017. Telepresence Robots in the Wide Wild World. In *3rd Workshop on Ethical Encounters in HCI: Research in Sensitive Settings*.
- [5] Liqing Cui, Shun Li, and Tingshao Zhu. 2016. Emotion detection from natural walking. In *International Conference on Human Centered Computing*. Springer, 23–33.
- [6] Beatrice De Gelder. 2006. Towards the neurobiology of emotional body language. *Nature Reviews Neuroscience* 7, 3 (2006), 242–249.
- [7] Ahti Heinla. 2018. Hello, Robot! (Nov 2018). <https://medium.com/starshiptechnologies/hello-robot-28b9b73787bf>
- [8] Rui Jiang, Shuzhi Sam Ge, Nagacharan Teja Tangirala, and Tong Heng Lee. 2016. Interactive Navigation of Mobile Robots Based on Human’s Emotion. In *International Conference on Social Robotics*. Springer, 243–252.
- [9] Thibault Kruse, Alexandra Kirsch, Harmish Khambhaita, and Rachid Alami. 2014. Evaluating directional cost models in navigation. In *Proceedings of the 2014 ACM/IEEE international conference on Human-robot interaction*. ACM, 350–357.
- [10] Thibault Kruse, Alexandra Kirsch, Emrah Akin Sisbot, and Rachid Alami. 2010. Exploiting human cooperation in human-centered robot navigation. In *RO-MAN, 2010 IEEE*. IEEE, 192–197.
- [11] Thibault Kruse, Amit Kumar Pandey, Rachid Alami, and Alexandra Kirsch. 2013. Human-aware robot navigation: A survey. *Robotics and Autonomous Systems* 61, 12 (2013), 1726–1743.
- [12] Christina Lichtenthäler and Alexandra Kirsch. 2014. Goal-predictability vs. trajectory-predictability: which legibility factor counts. In *Proceedings of the 2014 acm/ieee international conference on human-robot interaction*. ACM, 228–229.
- [13] Ross Mead and Maja J Mataric. 2015. Robots Have Needs Too: People Adapt Their Proxemic Preferences to Improve Autonomous Robot Recognition of Human Social Signals. *New Frontiers in Human-Robot Interaction* (2015), 100.
- [14] Jorge Rios-Martinez, Anne Spalanzani, and Christian Laugier. 2015. From proxemics theory to socially-aware navigation: A survey. *International Journal of Social Robotics* 7, 2 (2015), 137–153.
- [15] Alessandro Vinciarelli, Maja Pantic, Dirk Heylen, Catherine Pelachaud, Isabella Poggi, Francesca D’Errico, and Marc Schroeder. 2011. Bridging the gap between social animal and unsocial machine: A survey of social signal processing. *IEEE Trans. on Affective Computing* 3, 1 (2011), 69–87.
- [16] Jered Vroon, Gwenn Englebienne, and Vanessa Evers. 2019. Detecting Perceived Appropriateness of a Robot’s Social Positioning Behavior from Non-Verbal Cues: ‘A robot study in scarlet’. In *Proc. conf. on Cognitive Machine Intelligence*. IEEE.