

Mitigating undesirable emergent behavior arising between driver and semi-automated vehicle

Timo Melman[†]
Cognitive Robotics
Delft University of Technology
Email: T.Melman@tudelft.nl
[†] Co-first author

Niek Beckers[†]
Cognitive Robotics
Delft University of Technology
Email: N.W.M.Beckers@tudelft.nl
[†] Co-first author

David Abbink
Cognitive Robotics
Delft University of Technology
Email: D.A.Abbink@tudelft.nl
Website: www.delfthapticslab.nl

Abstract—Emergent behavior arising in a joint human-robot system cannot be fully predicted based on an understanding of the individual agents. Typically, robot behavior is governed by algorithms that optimize a reward function that should quantitatively capture the joint system’s goal. Although reward functions can be updated to better match human needs, this is no guarantee that no misalignment with the complex and variable human needs will occur. Algorithms may learn undesirable behavior when interacting with the human and the intrinsically unpredictable human-inhabited world, thereby producing further misalignment with human users or bystanders. As a result, humans might behave differently than anticipated, causing robots to learn differently and undesirable behavior to emerge. With this short paper, we state that to design for Human-Robot Interaction that mitigates such undesirable emergent behavior, we need to complement advancements in human-robot interaction algorithms with human factors knowledge and expertise. More specifically, we advocate a three-pronged approach that we illustrate using a particularly challenging example of safety-critical human-robot interaction: a driver interacting with a semi-automated vehicle. Undesirable emergent behavior should be mitigated by a combination of 1) including driver behavioral mechanisms in the vehicle’s algorithms and reward functions, 2) model-based approaches that account for interaction-induced driver behavioral adaptations and 3) driver-centered interaction design that promotes driver engagement with the semi-automated vehicle, and the transparent communication of each agent’s actions that allows mutual support and adaptation. We provide examples from recent empirical work in our group, in the hope this proves to be fruitful for discussing emergent human-robot interaction.

I. BACKGROUND ON AUTOMOTIVE USE CASE

Emergent behavior arises from the interaction between multiple agents in an environment, and are often unexpected and undesired (see Fig. 1). This short paper focuses on avoiding undesirable emergent behavior arising from the interaction between two embodied agents in an automotive scenario: a *driver* interacting with the *algorithms* of a semi-automated vehicle (sAv) in a real-world environment. Predicting undesired emergent behavior is an extremely challenging task as it cannot result from understanding the vehicle in isolation of its driver, or vice versa. As such, the occurrence of emergent behavior depends not only on the driver but also on the automation algorithms, the interaction design, and the environment itself.

Interestingly, vehicle automation algorithms are often based on data gathered in isolation, e.g. by mimicking naturalistic

driving behavior. These algorithms may fail to generalize their learned behavior to scenarios that were not encountered before. The algorithms might even learn undesirable behavior and do not account for potential emergent behavior. This may cause misalignment that is annoying or even dangerous, hence the industry standard is to leave final responsibility with the driver. The driver may also behave in undesirable or even unexpected ways when interacting with the sAV. For example, drivers may not accept the automation due to misalignment or annoyance (disuse), or they may over-rely on the automation (misuse), which can lead to inappropriate responses (e.g., nodding off while driving in auto-pilot mode) or deskilling [8]. In human factors literature, disuse and misuse are referred to as *behavioral adaptation*: undesirable adaptations in human behavior arising from the interaction between human and automation that mitigate the very benefits the automation aimed to realize. Disuse and misuse depend on the driver’s attitude, the interaction design, the automation algorithms, and the dynamic environment. We, therefore, view behavioral adaptation as an important source of emergent behavior, arising from the interactive nature of driver and sAV.

II. MITIGATING UNDESIRABLE EMERGENT BEHAVIOR

This automotive example illustrates the point that to avoid undesirable emergent behavior in human-robot interactions, we need a three-pronged approach as illustrated in Fig. 1. First, we should leverage human factors knowledge on driver’s underlying behavioral mechanisms and incorporate these in the sAV’s algorithm to mitigate misalignment at the source. Second, human factors knowledge of driver adaptations *due to the interaction with the sAV* needs to be included in the interaction design. This should avoid undesirable behavioral adaptations to the sAV. Third, all this should be grounded in an interaction design that limits the impact of misalignment *when* it occurs, and allows for transparent communication of each agent’s underlying intentions. In our group, we integrate these three approaches in an attempt to understand, observe, and mitigate undesirable emergent behavior.

A. Including behavioral mechanisms in the sAV algorithms

For automated driving, the sAV’s actions are generally selected by optimizing for a reward function that is defined




Fig. 1. A schematic illustration of how undesirable emergent behavior may arise from the dynamic interaction between driver, self-learning algorithms, and the sAV in its environment (left panel). We propose a cybernetic interaction-based approach to mitigate this (right panel), combining three approaches: 1) leveraging human factors knowledge on driver’s behavioral mechanisms and incorporate these in the sAV’s algorithm, 2) including models of driver behavioral adaptation *in interaction* with sAV’s, and 3) resolving remaining misalignment and undesirable emergence through a driver-centered interaction design.

in the design phase. To account for misalignment between driver preferences and sAV, algorithms can be used to infer the driver’s (initially unknown) preferences, even while the sAV is already deployed (e.g., [3, 5]). Literature provides several approaches to estimate and model naturalistic driving behavior, typically using *rational* reward functions or inverse reinforcement learning techniques [10]. In our lab we adhere to the perspective of *bounded rationality* in driving, which states that drivers tend to keep performance within *acceptable, not optimal* states [2]. We have recently proposed a bounded-rationality risk-based reward function to propose a driver model from which risk-motivated adaptations in speed and trajectory naturally emerge [4]. This risk-based driver model was able to replicate a wide variety of such adaptations from literature, for example in response to changes in road width or oncoming traffic. It was also implemented in a test vehicle for a recent on-road study, providing human-like trajectory adaptations to a parked vehicle outside of the lane.

Still, no matter how well a risk-based reward function can capture naturalistic driving behavior, it will never do so perfectly for all drivers under all conditions. More importantly, a predefined reward function for the driver-sAV system does not necessarily distinguish between each agent’s contributions to the driving task, whether driver behavior or algorithm behavior in isolation, or some form of cooperative driver-sAV behavior. As such, the sAV’s algorithm does not account for undesirable emergent behavior such as when a driver - feeling supported - engages in more risk-seeking activities, resulting in an increase of the sAV’s support, thereby progressively replacing the driver. In safety-critical scenarios outside the operational design domain of the sAV, such emergent behavior may prove detrimental.

B. Including predictions of driver behavioral adaptations

A second approach is to include an empirical understanding of how driver behavior is shaped *through the interaction* with the sAV (See 2 in Fig. 1). Conceptual behavioral theories exist, called homeostasis theories, that aim to capture the underlying invariants (e.g., experienced task difficulty or risk) that govern driver adaptation to external changes. We recently

replicated the three main theories to quantitatively compare their predictive capabilities into behavioral adaptation models [7]. We also used this knowledge to explain why drivers started speeding when supported by a haptic lane-keeping assistance system, thereby negating the safety and comfort benefits present at a fixed speed. Subsequently, we reduced haptic assistance above the speed limit, which subconsciously stimulated drivers to settle for a dynamic equilibrium around the speed limit, where haptic assistance provided safety and comfort [6]. Such behavioral nudging may also occur *between* different vehicles: e.g., sAV’s can be designed such that they select their actions to elicit a desired behavioral adaptation from a human driver in another vehicle [9].

C. Driver-Centered Interaction Design

Because driver behavioral models and sAV reward functions will always be simplified representations of reality, approaches 1 and 2 will only reduce the occurrence and magnitude of misalignments and undesirable emergence, but not always prevent them. Therefore, a driver-centered interaction design is essential to cope with undesirable emergent behavior (from both sAV and driver) *when they occur* (for a recent overview concerning the interaction design space based on shared control, see [1]). As long as drivers remain responsible, we advocate haptic shared control which makes misalignments tangible and correctable, thereby establishing mutual support between driver and sAV and the opportunity for mutual learning due to awareness of their dynamic capabilities and preferences.

III. SYMBIOTIC DRIVING FOR BENEFICIAL EMERGENCE

When these three design approaches are combined, we speak of ‘symbiotic driving’. Recent test-track studies have demonstrated symbiotic driving outperforms sAV designs that have only improvements in human-like reward functions, or only improvements in interaction. Additionally, simulator studies show symbiotic driving avoids undesirable emergent behavior. In future work, we will explore symbiotic driving as a means to create the conditions for beneficial emergent behavior, also between multiple road users.

REFERENCES

- [1] D.A. Abbink, T Carlson, M Mulder, J.C.F. de Winter, F Aminravan, T.L. Gibo, and E.R. Boer. A Topology of Shared Control Systems—Finding Common Ground in Diversity. *IEEE Transactions on Human-Machine Systems*, 48(5):509–525, October 2018. doi: 10.1109/THMS.2018.2791570.
- [2] M.A. Goodrich, W.C. Stirling, and E.R. Boer. Satisficing revisited. *Minds and Machines*, 10(1):79–109, 2000. doi: 10.1023/A:1008325423033.
- [3] D. Hadfield-Menell, S. Milli, P. Abbeel, S. J Russell, and A. Dragan. Inverse Reward Design. *31st Conference on Neural Information Processing Systems*, page 10, 2017.
- [4] S.B. Kolekar, J.C.F. De Winter, and D.A. Abbink. Human-like driving behaviour emerges from a risk-based driver model. (in press).
- [5] M. Kuderer, S. Gulati, and W. Burgard. Learning driving styles for autonomous vehicles from demonstration. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2641–2646, Seattle, WA, USA, May 2015. IEEE. ISBN 978-1-4799-6923-4. doi: 10.1109/ICRA.2015.7139555.
- [6] T. Melman, J.C.F. de Winter, and D.A. Abbink. Does haptic steering guidance instigate speeding? A driving simulator study into causes and remedies. *Accident Analysis & Prevention*, 98:372–387, January 2017. doi: 10.1016/j.aap.2016.10.016.
- [7] T. Melman, D.A. Abbink, M.M. van Paassen, E. R. Boer, and J.C.F. de Winter. What determines drivers’ speed? A replication of three behavioural adaptation experiments in a single driving simulator study. *Ergonomics*, 61(7):966–987, July 2018. doi: 10.1080/00140139.2018.1426790.
- [8] R. Parasuraman and V. Riley. Humans and Automation: Use, Misuse, Disuse, Abuse. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 39(2):230–253, June 1997. doi: 10.1518/001872097778543886.
- [9] D. Sadigh, S.S. Sastry, and A.D. Seshia, S.A. and Dragan. Planning for Autonomous Cars that Leverage Effects on Human Actions. In *Robotics: Science and Systems XII*. Robotics: Science and Systems Foundation, 2016. ISBN 978-0-9923747-2-3. doi: 10.15607/RSS.2016.XII.029.
- [10] D. Sadigh, S.S. Sastry, S.A. Seshia, and A.D. Dragan. Information gathering actions over human internal state. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 66–73, Daejeon, South Korea, October 2016. IEEE. doi: 10.1109/IROS.2016.7759036.