

Embodied Design Improvisation for Autonomous Vehicles

David Sirkin, Sonia Baltodano, Brian Mok, Dirk Rothenbücher,
Nikhil Gowda, Jamy Li, Nikolas Martelaro, David Miller, Srinath Sibi,
and Wendy Ju

Abstract We have developed a generative, improvisational and experimental approach to the design of expressive everyday objects, such as mechanical ottomans, emotive dresser drawers and roving trash barrels. We have found that the *embodied design improvisation* methodology—which includes storyboarding, improvisation, video prototyping, Wizard-of-Oz lab studies and field experiments—has also been effective in designing the behaviors and interfaces of another kind of robot: the autonomous vehicle. This chapter describes our application of this design approach in developing and deploying three studies of autonomous vehicle interfaces and behaviors. The first, *WoZ*, focuses on the conceptual phase of the design process, using a talk-aloud protocol, improvisation with experts, and rapid prototyping to develop an interface that drivers can trust and hold in esteem. The second, the *Real Road Autonomous Driving Simulator*, explores people’s naturalistic reactions to prototypes, through an autonomous driving interface that communicates impending action through haptic precues. The third, *Ghost Driver*, follows the public deployment of a prototype built upon frugal materials and stagecraft, in a field study of how pedestrians negotiate intersections with autonomous vehicles where no driver is visible. Each study suggests design principles to guide further development.

1 Introduction

Car: “So, I’m noticing that one of my tires is low on air. Is there something about yourself that you want to talk about?” Driver, after two seconds, laughs out loud: “That’s the most bizarre question I’ve ever heard from a car.”

D. Sirkin (✉) • S. Baltodano • B. Mok • D. Rothenbücher • N. Gowda • J. Li • N. Martelaro •
D. Miller • S. Sibi • W. Ju
Center for Design Research, Stanford University, 424 Panama Mall, Stanford, CA 94305, USA
e-mail: sirkin@stanford.edu

Driver: “Tell me about yourself, car.” Car: “What would you like to know?” Driver: “Where are you from?” Car: “I’m from Japan. Where are you from?” Driver, after five seconds: “This is like a ‘Her’ moment.”

-Conversations Between Drivers and Autonomous System
WoZ Study

We have previously written about our method of *embodied design improvisation* to design machines and robots where (a) physical motions, gestures or patterns are employed, (b) the design space of possible actions, mechanisms and dimensions is vast, and (c) the cost in time, money and effort to build fully functioning systems is high (Sirkin and Ju 2015). This method exemplifies the challenges of design thinking, where the understandings that designers draw upon are often tacit, where an exhaustive search of the design space is not possible, and where the costs of full-scale solutions precludes iterative empirical testing.

The context for these prior articles was the form and movement of *expressive everyday objects*: non-anthropomorphic household furnishings that can initiate, conduct, and conclude interactions with people in a meaningful, improvisational way. In this chapter, we illustrate how we have applied the same approach to the design of interactions in another domain: autonomous vehicle interfaces.

Each of the following three sections summarizes a behavioral experiment in this new context, and focuses on aspects of the *embodied design improvisation* process: *WoZ* relates to early-stage exploration of design ideas within the vehicle through improvisation and enacted scenarios, the *Real Road Autonomous Driving Simulator* highlights rapid prototyping of physical interfaces applied to a field experiment, and *Ghost Driver* underscores the iterative redesign of study procedures to observe and understand pedestrian interactions outside of the vehicle.

2 Case Studies in Design Improvisation

These case studies center on the design of interactions between people and *near future* autonomous vehicles: in particular, how vehicles that operate with varying degrees of autonomy—from basic driver assistance to fully self-driving systems—can, and should, interact with drivers and passengers inside of the vehicle, and with pedestrians and bicyclists outside of it. Over the course of these projects, our goals have been (a) to observe and understand how people respond to vehicles that exhibit intrinsic agency during everyday driving, and (b) to develop and explore vehicle interfaces and behaviors that express that agency, including their features and limitations.

Each study draws upon *Wizard of Oz* techniques, which are frequently employed in interaction design, and which we further motivate and describe in the next section. Experimenters operate as *stand-ins* for future technologies that would otherwise perform extended, contingent interactions with people. The name comes from the novels of L. Frank Baum, wherein a Wizard is believed by all of the denizens of the Land of Oz to be a magical being, where in fact, he is an

ordinary man employing a variety of tricks to project an illusory reality (Baum 1900).

2.1 *Wizard of Oz (WoZ)*

2.1.1 Introduction

The WoZ system (Mok et al. 2015) is designed (a) to explore how drivers and their partially or fully autonomous vehicles can share and exchange driving tasks, and (b) to understand the thoughts and feelings that drivers experience during these transitions. Driving tasks include controlling the vehicle as it rolls down the road, or navigating through town to some destination. A transition is where the driver indicates that the car should assume some responsibility that he or she currently holds and the car accepts that role, or the reverse: where the car indicates and the driver accepts. Taking this perspective, driver and car act as a team (Inagaki 2009), collaborating to reach their destination in a safe, legal, timely and comfortable way. It thus becomes important that they learn each other's abilities, and communicate and understand each other's intentions and actions. By integrating WoZ into a driving simulator (see Fig. 1), we can alter the car's abilities and actions at a moment's notice, or update the surroundings (including pedestrians, other cars, and their behaviors) from one drive to the next.

2.1.2 Prototype Systems

The WoZ station allows researchers to communicate with participant drivers in the simulator vehicle through a speech interface, as well as to initiate or respond to transfers of control and operate the vehicle during periods of autonomous driving (see Fig. 2). In fact, the vehicle does not drive autonomously: rather, researchers control the car's actions using a Wizard of Oz protocol (Dahlbäck et al. 1993; Cross 1977), where drivers are told that they are interacting with an autonomous system, but their interactions are mediated by a human operator. As noted by Hoffman and Ju (2014), this approach allows us to explore a wide range of features and functions without first building a fully operational system.

Due to the challenge of attending to several simultaneous tasks, WoZ has dual control stations. The Interaction Wizard observes and interacts with the participant using video cameras and a text-to-speech interface, making the car appear to be able to detect, and respond, to the drivers' movements, facial expressions and utterances. The Driving Wizard controls the car's autonomous driving, triggers events in the simulated environment—such as pedestrians crossing the road or surrounding cars braking quickly, and updates elements of the car's visual interface—such as instruments panels.



Fig. 1 The simulator used for experiments in autonomous vehicle driving and control. Participants sit in a fixed-base car surrounded by a 270° screen depicting the study environment



Fig. 2 The Wizard of Oz control station. Wizard 1 (the Interaction Wizard) interacts with the participant, while Wizard 2 (the Driving Wizard) initiates transfers of control and operates the vehicle during autonomous driving mode

2.1.3 Designed Behaviors

From the driver's perspective, the car presents as a physical, social agent, able to (a) operate autonomously, with varying degrees of competence (which, unbeknownst to the driver, may be high or low, depending on the researchers' agenda), (b) carry on spoken dialog, in a male or female voice, about driving topics such as the roadway and navigation, or notably, non-driving topics such as the driver's preferences, expectations, experiences or emotional state, and by combining these (c) initiate and respond to (for example) requests to change speed if running late, detour for lunch, or transfer control if the driver feels sleepy.

2.1.4 Improvisation Sessions

We invited 12 interaction and interface design experts to act as participants in individual design improvisation sessions of about 30 min each. Given the exploratory nature of the study, we provided participants with little information about the car or instruction for interacting with it, and encouraged them to actively improvise to discover its abilities.

Participants traveled through a course comprising four sections: (a) a brief practice to introduce the simulator and its environment, which includes straight roads, intersections and roundabouts, (b) a stretch of forests and hills, (c) a city with densely placed buildings, pedestrian intersections and crossing traffic, and (d) a highway. Notably, in each of the latter three sections, several potentially dangerous events occurred: car cutoffs, pedestrian incursions and cars pulling onto the road. During the drive, including these events, the car offered to assume control at certain times, and requested that participants resume control at other times. The car (at the behest of its Wizards) did not follow a controlled response protocol: rather, it freely offered explanations for its driving (*pushing* information to drivers), or responding to queries about its behavior (*pulling* information from drivers). We recorded the entire exchange, and interviewed participants after the session about their preferences, experiences and thoughts about the events and car's responses.

Desire for Shared Control

On several occasions, the car intentionally drove *imperfectly*, drifting laterally within its lane, crossing into the sidewalk during a turn, or closely approaching people or cars directly ahead. Although the car's performance was flawed, for the most part, participants still held the system in esteem, and preferred to make gentle corrections to the pedals or steering wheel, and remain in autonomous mode.

This form of shared control is not considered under National Highway Traffic Safety Administration's current Levels of Automation model (National Highway Traffic Safety Administration 2013). The design challenge therefore becomes how shared control *should* function. Participants expected that the car would modify its behavior based on their guidance. But how can the car know when that guidance had ended? Also, some participants kept their hands on the steering wheel as a way to monitor the car's actions, or to lessen their feelings of unease, rather than as a way to prime the car's future behavior, making it difficult to interpret their intent. One approach may be to monitor speech, which for the current study, included comments such as "*you're drifting to the right*" or "*you're too close to the car ahead.*" Such explicit signals may be unstructured, but at least they carry clear intent and can be interpreted.

Trust in Autonomy

Participants reported that two behaviors significantly improved their trust in the car's autonomous system: (a) successfully traversing challenging sections of road, including a traffic circle and an s-curve, and (b) calling out features or events in the environment *that a human would have found noteworthy*. Regarding the latter, drivers felt that noting every possible event in the environment would become tiresome, but that highlighting events related to safety, or which the driver might have wanted to observe but had missed, made the car seem more like a peer. At one intersection, the car commented that one pedestrian in a group at the crosswalk was lagging behind the others, and almost all participants interpreted that observation as the car perceiving the world the same way that they (as humans) do.

On the other hand, after experiencing more egregious bouts of imperfect driving—such as maintaining uncomfortably short headway distance—participants tended to disengage automation and disapprove requests (by the car) to resume autonomous driving. Over time, their trust in the system could be repaired, but only after ongoing, non-driving-related conversation, such as the quotes at the start of the chapter (Sirkin et al. 2016), or several small trust-building tests, such as providing a warning prior to a car cutoff. For example, after the car announced “*There are people here,*” the driver responded “*Do you know how many people, car?*”

Driving Mode Transitions

Participants often felt that the timing of transitions was unclear. In particular, even short phrases such as “*I have control now,*” when spoken by the car, are expressed over time. The resulting uncertainty over when the transition occurred, or whether it was safe to relinquish control, led participants to ask “*Are you driving the car now?*” or “*Can I let go of the steering wheel now?*” We found that adding a chime, with sharp attack, provided a better transition demarcation (Fig. 3). While participants felt that a “3-2-1” spoken countdown prior to the chime provided even greater advance notice, it also extended the transition time, and became a nuisance after several times.

We also tested visual indicators of mode changes, including an instrument cluster graphic which changed color from gray to green, and wording from “*Autonomy Off*” to “*Autonomy On,*” during transitions to autonomous control. Participants found such cues helpful in determining when transitions had occurred, and suggested that a haptic indicator—such as a vibrating steering wheel, or tightening seatbelt, or moving seat—might provide even more effective notification.

Addressing Requests

Participants often instructed the car to perform certain tasks, such as “*pass that slow vehicle in front of us*” or “*tell me about today's news headlines*” or “*play music on*



Fig. 3 At the *top*, a spoken phrase like “*I have control now*” takes time, making the moment of transition unclear. At the *bottom*, a chime communicates the exact moment more clearly

my play list,” some of which might be inadvisable, and others unavailable. Drivers were typically satisfied if the car provided a technical rationale for not performing the task, such as not having access to certain audio files, but they were less willing to relent when they knew that the car was capable of performing the request, but refused to do so for some non-technical reason, such as exceeding the speed limit. In this case, participants continued to ask the car to speed up, with several eventually choosing to disengage autonomy and drive over the speed limit.

Response Latency

The car typically responded to participants’ requests, commands and conversation within about 10 s, limited by Wizard 1’s ability to type the message quickly. This delay was particularly noticeable for longer answers, which created an extended, uncomfortable silence, causing participants to question whether the car had heard, or interpreted, what they said. One way to ameliorate the problem is to provide a short acknowledgment, allowing for a detailed follow-up. For example, saying “*let me find out*” signals that the car has received and understood the message, and is working on a response. An alternative, suggested by one participant, is to play audio tones (such as soft beeps) that suggest that the computer is processing the information.

2.2 Real Road Autonomous Driving Simulator (RRADS)

2.2.1 Introduction

RRADS (Baltodano et al. 2015) is an on-road vehicle platform and set of study procedures to help researchers design and test autonomous vehicle interfaces. We developed the system specifically to explore attitudes and concerns that people may have in real-world, rather than fixed-base simulator, autonomous vehicles.

There are currently few platforms available to support such research. Virtual lab-based simulations excel at creating highly structured and controlled events (Talone et al. 2013), however, they are difficult to acquire and maintain, and

struggle to replicate the rich sensory stimuli, inertial forces, changing lighting and weather, and unpredictable traffic patterns experienced in on-road settings. For these reasons, we were motivated to developing a low-cost, safe and reliable real-world testbed for human-autonomous vehicle interactions.

2.2.2 Prototype Systems

RRADS involves two Wizards (Kelley 1983; Dahlbäck et al. 1993) in a single vehicle: a *Driving Wizard*, who controls the vehicle from the usual driving position, and an *Interaction Wizard*, who operates the interfaces being developed from the rear seat. Three GoPro cameras record road events, participants' reactions, and the actions of the Wizards. A partition made of stiff, opaque material obfuscates participants' view of the Driving Wizard (Fig. 4).

The partition plays a dual role during experiments: preventing the participant from seeing the Driving Wizard, while not compromising the Wizard's ability to use driving controls and safety features, including steering wheel, shift lever, pedals and mirrors. It is constructed of stiff, 2 cm thick foam core board, affixed to the vehicle interior using gaffer's tape. Figure 5a, b show the partition installed in an Infiniti M45, one of our two test vehicles (the other being a Jeep Compass).

Figure 5b also shows a (non-functional) steering wheel mounted to the dashboard directly ahead of the participant. We found that even this small gesture suggests the participant's role as being more than just passenger. It also supports the participant's suspension of disbelief, and in turn, increases the effectiveness of the simulation.

Three GoPro cameras, oriented as shown in Fig. 4, record events during the experiment, with the camera shown in Fig. 5b focusing on the participant's hand motions and facial expressions. Through these, the Interaction Wizard observes the participant's reactions, allowing in-the-moment, improvisational responses.

2.2.3 Designed Behaviors

The RRADS protocol has three main sections, each of which is designed to support participants' suspension of disbelief.

Meet and Greet

At the start of a session, a researcher greets and guides the participant to the vehicle, approaching it from the passenger side (Fig. 6). The vehicle is parked along the curbside with the Driving Wizard inside, but not visible through the windows, and the Interaction Wizard waiting by the rear passenger door. The Interaction Wizard is introduced as monitoring the autonomous system, and the participant is seated.

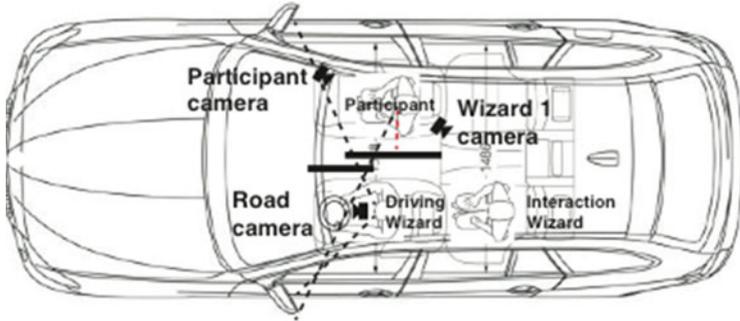


Fig. 4 The RRADS platform vehicle, noting placement of the Driving and Interaction Wizards, study participant, recording cameras, and opaque partition

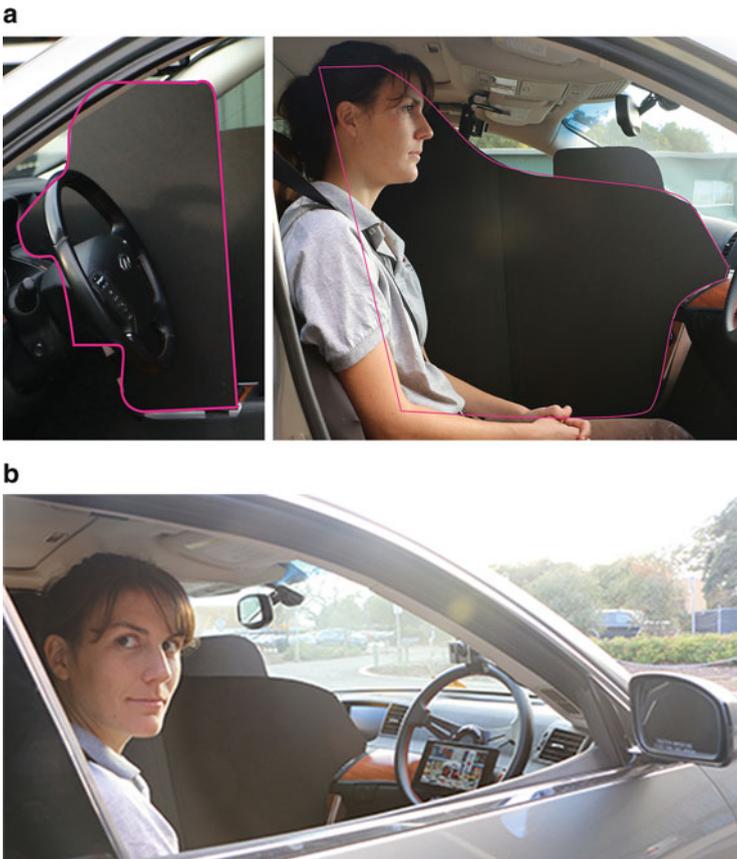


Fig. 5 (a) The Driving Wizard partition, as viewed from the driver's side (on the left), and passenger's side (on the right). (b) The participant's portion of the vehicle interior includes non-functional steering wheel, tablet interface, and video camera for recording gestural and facial reactions

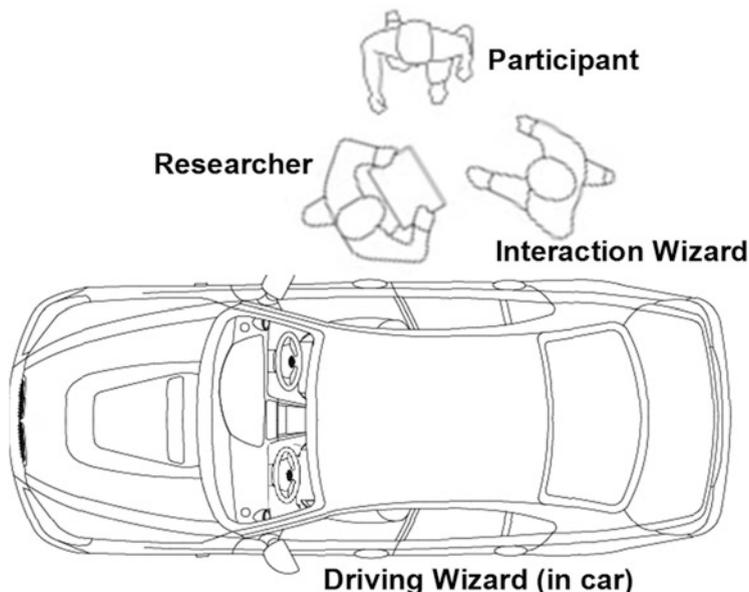


Fig. 6 Staging a participant’s introduction to the RRADS vehicle. A researcher guides the participant into position; the Driving Wizard is already concealed behind the partition

On the Road

The vehicle follows a typically pre-selected course, which is predictable and safe, and which the Driving Wizard knows. Pedestrians, traffic lights, speed limit changes and high-density traffic can all be sources of opportunity or complication to study design. When the vehicle returns, the researcher opens the participant’s door and engages him or her in light conversation, to allow the Driving Wizard the opportunity to drive away unseen.

Exit Interview

A qualitative exit interview provides an opportunity to uncover the salient points to the passenger’s experience. Qualitative pilot phrases, such as “*How did the drive go, in general terms?*” can yield in-depth narrative responses which can be mined afterward.

2.2.4 Improvisation Sessions

The goal of our improvisation was to understand if pre-advising occupants of an autonomous vehicle’s upcoming maneuvers—its starts, stops, lane changes or turns—might influence their trust in the system: effectively saying “*we’re about*

to turn” just before turning (or starting or stopping). We drove 35 participants along a course which included haptic (relating to touch) cues—expressed by the vibration or movement of the seatback or floorboard—to indicate these upcoming maneuvers.

The first prototype was a pneumatic base in the participant’s foot well, just below his or her feet, that tilted in the direction the car was about to move. The second was an array of vibration motors, embedded in the seat back, which expressed patterns suggesting vehicle movement—such as a *cascade* from top to bottom to indicate an upcoming stop. The third was a pneumatic bladder placed behind the participants’ shoulders which, when inflated on one side or the other, indicated that the vehicle was about to turn the opposite direction. We found that participants’ guesses before the event were earliest in the case of the pneumatic floorboard, and latest in the case of the vibration array, suggesting that the floorboard provided the most effective early warning mechanism. We also found that environmental cues from the road, and from the vehicle’s motion were always present, and played a significant role in participants’ guesses.

The Hello Effect

At the start of the study, the vehicle greeted the participant using one of the prototype devices, or in the control case, with a revving of the engine. This was built into the study for practical reasons, to verify that the prototype was functioning properly, however, the interaction had an unintended effect on participants’ experiences, with several cited the greeting as a source of comfort and a sign of amicability: *“I was surprised how much I trusted it. Even from the beginning, when it said, ‘Hello,’ it had enough of a personality. That one thing gave it enough personality for me to trust it,”* or *“When the car said, ‘turn on,’ or something, and then there was the air, it just kind of shot up, and I was like, ‘okay, that’s kind of interesting,’ but it’s like its way of communicating with you, rather than a voice thing.”*

Trust Through Driving Style

The driving style was, in itself, a form of improvisation. The Driving Wizard had to assume the role of an autonomous car and drive like one consistently in an unpredictable world. Road hazards, school and construction zones, and emergency vehicles required more attention. At these times, it was important that the Driving Wizard stay in character and keep the study going. The Wizard also had to prepare for changing lighting conditions, wear the right clothes to reduce noise while driving, and even hold in a sneeze every now and then.

Trust was high through all conditions. In fact, we did not find any significant statistical differences between conditions despite several of the pre-cuing systems (especially the floor boards) being very effective pre-cuing devices. We expect that this lack of difference is due to the driver’s consistent and conservative driving style, so that the strength of this signal may have overwhelmed those from all other inputs.

Smooth Driving Is Safe Driving

During interviews, participants often referenced the car's safe driving style when questioned about trust. 30 % of responses to "*Did you trust the vehicle?*" strongly related to descriptions of smooth driving: "*The main thing was that it drove very smoothly.*" This relevance also emerged when we asked participants to elaborate on why they trusted the vehicle: "*Because it was smooth and it wasn't too fast or jerky,*" or "*It wasn't anything sudden, or things that would normally make me go, oh my God, this is scary, stop.*" The descriptions of smooth driving also related to descriptions of vehicle planning and awareness: "*Definitely smooth starts and stops. It sort of made it feel like the vehicle was planning what it was doing,*" or "*Something's a bit smoother, you realize that the person or the car kind of knows what's going on.*"

Trust and Belief

The improvisation protocol did not employ deception. The partition separating the Wizard Driver from the participant was intended merely to help facilitate the illusion of an autonomous vehicle. And yet, about 25 % of participants believed that the system was fully autonomous. Another large portion of the participants believed the vehicle was partially autonomous and remotely controlled by the Interaction Wizard.

The prompt "*Did you trust the vehicle?*" was particularly helpful in uncovering how immersed participants became: "*I guess the computer was pretty cautious, which was pretty awesome. It was a much better driver than most humans that I know,*" or "*It made me feel like even though it wasn't a human, it wasn't of malicious intent.*" A few participants who believed the system was autonomous revealed reservations about the technology: "*I just don't fully trust that car to drive on its own. Even though I had no bad experiences with this car, it just seems strange to me still and foreign to me that a car can drive itself.*"

During more complicated maneuvers, some participants ascribed agency to the Interaction Wizard: "*There was a construction site. The guy was waving for me to move and I was like, I don't know what to do, so I was like, 'I really hope the car does something smart.' The car backed up and then the guy made more hand signals. I don't know if [the Interaction Wizard] or if the car did it.*"

2.3 Ghost Driver

2.3.1 Introduction

With increasing capability of self-driving cars (Lari et al. 2014), each of a vehicle's occupants, including its operator, may become mere passengers, with no visible

human driver. Long-established practices of communication between drivers and road users outside of the vehicle—such as making eye contact, nodding one’s head or giving hand signs—may no longer be possible. Beside the arising concerns around safety, how comfortable might people feel walking or bicycling in front of autonomous cars if they do not receive acknowledgment that they have been noticed?

Prior work in social science, psychology and civil engineering has shown that a driver’s gaze first goes to the face of a bicyclist (Walker and Brosnan 2007), and that pedestrians who stared at approaching drivers elicited greater stopping (Guéguen et al. 2015). Luoma and Peltola (2013) found that high speed was a signal from drivers that they did not intend to give way to pedestrians, and a similar field observation (Velde et al. 2005) showed that over half of pedestrians do not look for vehicles after arriving at a curb, but that all of them look at oncoming vehicles while crossing.

2.3.2 Prototype Systems

To explore these questions, we needed to evoke the impression that a car was driving autonomously, and in turn, deprive pedestrians and bicyclists of any chance to interact with a human (anywhere) in the car. While the California Department of Motor Vehicles issues licenses for testing autonomous technology, regulations require a human operator to occupy the driver seat at all times: to take over in an emergency, or for driving when autonomy is turned off (California 2015). Thus, a self-driving car would *have* to have a visible operator, and even though that operator might not be driving the car at the moment of an interaction, participants could wrongly interpret him or her as doing so.

We therefore developed a system (Rothenbücher et al. 2016) that would disguise both vehicle and operator, to make them appear to be autonomous to the outside world. For the car, we attached (that is, we did not connect and utilize) components from functioning autonomous cars—including a laser-based detector on the roof, radar units on the front corners, and cameras on the roof and dashboard—to an ordinary VW eGolf, and added vinyl stickers on the hood and doors saying *Stanford Autonomous Car* (Fig. 7).

For the driver, inspired by an invisible driver prank published on YouTube (Hossain 2013), we designed a *car seat costume*, to make him or her invisible to anyone outside of the vehicle (Fig. 8). The basic shape of the original seat was formed in wire mesh, stabilized with paper-mâché, and covered with a regular seat cover. To give the driver peripheral vision, we covered the wire mesh around the headrest (only) with sheer, see-through black fabric. The driver was dressed in black, including hands. Using only the bottom of the steering wheel, the driver could maneuver the car without being seen.



Fig. 7 The Ghost Driver car features autonomous vehicle props like a laser-based detector, radar units, interior and exterior cameras and decals on the hood and doors



Fig. 8 The seat cover costume includes two arm outlets so that the concealed driver could steer the car using the bottom of the steering wheel

2.3.3 Designed Behaviors

Implicit interaction theory (Ju 2015) suggests that pedestrian-autonomous vehicle interaction patterns at intersections might resemble the following: (a) a pedestrian approaches an intersection, (b) a car approaches the same intersection, (c) the pedestrian makes eye contact, (d) the driver makes eye contact, (e) the driver indicates not giving way, (f) the pedestrian waits, (g) the driver moves through the crosswalk, and finally (h) the pedestrian crosses—or alternatively: (e′) the driver indicates giving way, (f′) the driver stops and waits, (g′) the pedestrian crosses, and (h′) the driver moves through the crosswalk.

While this is our expected behavior, an autonomous car with no visible driver would likely break down the pattern in steps (c) and (d), critical points in which driver and pedestrian intent are communicated. This lack of nonverbal signals from a driver—or more correctly, the lack of a driver—is atypical, and serves to heighten awareness of the car, thereby making it appear to be more proactive. The transfer of agency from driver to car may lead to attempts to *repair* the interaction through improvised forms of iteration: repeatedly searching for a driver, staring at the location where the driver would be, trying to verbally engage a driver.

2.3.4 Improvisation Sessions

We developed a *breaching* experiment (Garfinkel 1991; Weiss et al. 2008), wherein we placed our mock autonomous car in a natural setting (Zhuang and Wu 2011), to observe how participants would respond. We video recorded interactions from multiple perspectives (on top of the car, on the street) (Millen 2000; Crabtree 2004), analyzed the footage, looking for behavior patterns and responses, and asked participants open-ended questions about their reactions and if they believed that the car was driving autonomously.

We held sessions in two locations: a parking lot with a pedestrian crosswalk at its exit leading onto a street, and a traffic circle with a large ratio of bicyclists. Both locations are highly frequented, especially between lectures and during lunch hour, so we ran sessions over 3 days from 11 am to 2 pm. At the crosswalk, the car waited in the parking lot, facing the exit and barely visible from the sidewalk. As soon as a pedestrian approached intersection, the invisible driver accelerated to arrive at the moment that the pedestrian was about to cross the street. We varied driving style from conservative on the first day to more aggressive at the second day, with the car approaching at a higher speed and stopping later. To create ambiguity, the car also briefly lurched forward after it had come to a full stop just as the pedestrian was about to cross. On the third day (a week later), we moved to the traffic circle, where the car entered and did a few circles before exiting again.

We employed a team of five researchers: a driver, a coordinator giving instructions to the driver and shooing away participants who came too close or lingered too long, and three interviewers, who asked participants questions just following their

interactions. We recorded 67 interactions (49 at the crosswalk and 18 at the traffic circle) and interviewed 30 participants.

Pedestrians' Interaction Behaviors

Most people who interacted with the car noticed the missing driver (80%) and believed that it was driving on its own (87%). Also, bystanders were excited to see a car that appeared to be self-driving: taking photos or videos, talking about it with friends, even tweeting about it. These observations, and self-reported impressions on our questionnaire, show that our Wizard-of-Oz approach worked well and achieved its purpose.

The props and decals drew attention to the car, so that most participants stared a while to focus on it. Part of this focus was to look at the driver, which explains why so many saw that there was none, making it more difficult for them to predict what would happen next.

And yet, the video revealed that people were not overly shy about walking in front of the car. Although they assumed that it was self-driving, their behavior proceeded in the usual way. Of the 49 crosswalk interactions, only two people clearly tried to avoid passing in front of the car by walking around it, both on the second day when the car restarted after having come to a full stop. One said later *"I waited for a while to see what it's going to do, then tried to cross, but then while I was trying to cross, it attempted to start, so I stopped and waited."* More often, we saw moments of hesitation, like stopping short or slowing walking pace. Sometimes, we saw people walk with greater deliberation and expressive motion, as if to explicitly signal that *a person is walking here* to the car (Fig. 9). So, particularly when the car stopped and restarted, established crossing procedures broke down, and participants' demand for interaction increased to resolve the ambiguity (Ju 2015)—only in this case, there was no driver to communicate and align with.

Forgiving the Newbie

Although the car sometimes misbehaved, nearly all people seemed to be tolerant and forgiving. We saw just one person on the second day who seemed to be upset by the car creeping into the sidewalk. Another person said *"I guess, if it were a person I'd have a really negative reaction towards them, but then, the autonomous car is a really interesting concept, so it was less negatively impacted."*

Overall, people seem to have lower expectations than they would for human drivers, grant that the car is still learning, and admit that mistakes are part of that learning process. But at the same time, their expectations seem to be higher, since the technology is meant to eliminate human error. One participant, upon walking in front of the car, said *"the risk I took by crossing the intersection was higher than I realized, because nobody is behind the wheel of the car. At the same time, there are no human errors, there are just car sensors."* Even though participants reported



Fig. 9 Participants walk in front of Ghost Driver, looking for a driver (on the *left*) and exaggerating walking behavior (on the *right*)

liking and trusting the car, some mentioned a certain unease, where they “*didn’t feel very comfortable*,” “*wanted to make sure that it wasn’t going to hit me*” or “*kept an eye out while crossing*.” But what seems contradictory might just be a dual concept of trust: one which is spontaneous, in-the-moment and action-driven, and another which is more conceptual, and derived from mental models built over time about technology.

Compliance and Acknowledgment

There are two distinguishable elements of the interaction between pedestrians or bicyclists with a driver: one is to evoke *compliance* by connecting with the driver, and the other is to get *acknowledgment* that one was noticed. The first becomes irrelevant when the driver is a robot, in that people do not assume that the robot responds in a moody, impulsive or brash way. In other words, we do not need to give robots *the look*. Acknowledgment, however, remains relevant. People expect recognition that they have been seen, and if there is no driver, they look for it from the car, through its movement and behavior. This provides a design opportunity: to leverage people’s knowledge of car behavior, such as expressive movement (like when the car’s front *dives* as it stops at an intersection, or its rear *squats* as it accelerates away), lighting, or even sound, as explicit cues of the car’s intentions.

3 Conclusion

In this paper, we have illustrated how the design thinking research techniques of *embodied design improvisation* have been applied to discover how interactions with autonomous vehicles should be designed.

In the WoZ sessions, participants wanted to share control with the car without assuming full control, and wanted to know exactly when the mode switch occurred. The car's delays in response and unperformed requests were acceptable, as long as it provided technical explanations. Each of these influenced participants' sense of trust, and the acceptability of the car as a conversation partner.

The RRADS platform pointed to influences that might be greater levers in trust than pre-cueing. The *Hello Effect* seems to indicate that an autonomous vehicle's perceived personality and driving style may be incredibly strong and salient factors in building users' trust in the system. In addition, participants seemed to be actively evaluating the vehicle's competence when it encountered complex situations, and its apparent ability to interact with construction vehicles or bicyclists seemed to reassure those who were initially skeptical of autonomous driving technology.

For Ghost Driver, further research will focus on the question of how important eye contact is for safety, once compliance is no longer an issue. In other words, is signaling acknowledgment a safety, or rather a convenience, feature of the car? We assume that this question is crucial for the design of such a signaling system because it defines the real needs of pedestrians and other human road users.

These three case studies exemplify how design thinking can help us to understand how people will respond to technologies that do not yet exist, and thereby help us to steer the direction of future technologies towards interactions that are safe and desirable.

References

- Baum L (1900) *The wonderful wizard of Oz* (Books of Wonder). HarperCollins, Richmond, TX
- Baldodano S, Sibi S, Martelaro N, Gowda N, Ju W (2015) The RRADS platform: a real road autonomous driving simulator. In: *Proceedings of the International Conference on Automotive User Interfaces and Interactive Vehicular Applications*, ACM, pp 281–288
- California Department of Motor Vehicles (2015) Autonomous vehicles in California. <http://www.dmv.ca.gov/portal/dmv/detail/vr/autonomous/testing/>. Accessed Mar 2016
- Crabtree A (2004) Design in the absence of practice: breaching experiments. In: *Proceedings of Conference on Designing Interactive Systems*, ACM, pp 59–68
- Cross N (1977) *The automated architect*. Viking Penguin, New York
- Dahlbäck N, Jönsson A, Ahrenberg L (1993) Wizard of Oz studies—why and how. *Knowl-Based Syst* 6(4):258–266
- Garfinkel H (1991) *Studies in ethnomethodology*. Polity Press, Cambridge, p 36
- Guéguen N, Meineri S, Eyssartier C (2015) A pedestrian's stare and drivers' stopping behavior: a field experiment at the pedestrian crossing. *Saf Sci* 75:87–89
- Hoffman G, Ju W (2014) Designing robots with movement in mind. *J Hum-Robot Interact* 3(1):89–122
- Hossain R (2013) Drive thru invisible driver prank. <http://youtube.com/watch?v=xVrJ8DxECbg>. Accessed Mar 2016
- Inagaki T (2009) Human-machine collaboration for safety and comfort. In: *Proceedings of the ENRI International Workshop on ATM/CNS*, Tokyo, Japan
- Ju W (2015) *The design of implicit interactions*. Morgan & Claypool, San Rafael, CA
- Kelley J (1983) An empirical methodology for writing user-friendly natural language computer applications. In: *Proceedings of the ACM Conference on Human Factors in Computing Systems*, ACM, pp 193–196

- Lari A, Douma F, Onyiah I (2014) Self-driving vehicles: current status of autonomous vehicle development and Minnesota policy implications. Preliminary white paper, University of Minnesota, pp 1–20
- Luoma J, Peltola H (2013) Does facing traffic improve pedestrian safety? *Accid Anal Prev* 50:1207–1210
- Millen D (2000) Rapid ethnography: time deepening strategies for HCI field research. In: Proceedings of Conference on Designing Interactive Systems, Processes, Practices, Methods and Techniques, ACM, pp 280–286
- Mok B, Sirkin D, Sibi S, Miller D, Ju W (2015) Understanding driver-automated vehicle interactions through Wizard of Oz design improvisation. In: Proceedings of the International Driving Symposium on Human Factors in Driver Assessment, Training and Vehicle Design, pp 386–392
- National Highway Traffic Safety Administration (2013) Preliminary statement of policy concerning automated vehicles. US Department of Transportation, Washington, DC
- Rothenbücher D, Li J, Sirkin D, Mok B, Ju W (2016) A field study investigating the interaction between pedestrians and driverless vehicles. In: Proceedings of the International Symposium on Robot and Human Interactive Communication, IEEE
- Sirkin D, Ju W (2015) Using embodied design research as a design research tool. In: Proceedings of the International Conference on Human Behavior in Design, Design Society, pp 1–7
- Sirkin D, Fischer K, Jensen L, Ju W (2016) Eliciting conversation in robot vehicle interactions. In: Proceedings of the Association for the Advancement of Artificial Intelligence Spring Symposium Series: Enabling Computing Research in Socially Intelligent Human Robot Interaction, pp 164–171
- Talone A, Fincannon T, Schuster D, Jentsch F, Hudson I. (2013) Comparing physical and virtual simulation use in UGV research: lessons learned from HRI research with two test beds. In: Proceedings of the Human Factors and Ergonomics Society Annual Meeting 57(1), 2017–2021
- Velde A, van der Kamp J, Balera J, Savelsbergh G (2005) Visual timing and adaptive behavior in a road-crossing simulation study. *Accid Anal Prev* 37(3):399–406
- Walker I, Brosnan M (2007) Drivers' gaze fixations during judgements about a bicyclist's intentions. *Transport Res F: Traffic Psychol Behav* 10(2):90–98
- Weiss A, Bernhaupt R, Tscheligi M, Wollherr D, Kühnlenz K, Buss M (2008) A methodological variation for acceptance evaluation of human-robot interaction in public places. In: Proceedings of the IEEE International Symposium on Robot and Human Interactive Communication, IEEE, pp 713–718
- Zhuang X, Wu C (2011) Pedestrians' crossing behaviors and safety at unmarked roadway in China. *Accid Anal Prev* 43(6):1927–1936