

# EPIC2016

## Case Studies 2 – Emergent, Underserved & New Markets

### Developing Socially Acceptable Autonomous Vehicles

ERIK VINKHUYZEN

*Nissan Research Center – Silicon Valley*

MELISSA CEFKIN

*Nissan Research Center – Silicon Valley*

*Recognizing that the movement of cars on the road involves inherently social action, Nissan hired a team of social scientists to lead research for the development of autonomous vehicles (AVs) that engage with pedestrians, bicyclists, and other cars in a socially acceptable manner. We are expected to provide results that can be implemented into algorithms, resulting in a challenge to our social science perspective: How do we translate what are observably social practices into implementable algorithms when road use practices are so often contingent on the particulars of a situation, and these situations defy easy categorization and generalization? This case study explores how our cross-disciplinary engagements have proceeded. A particular challenge for our efforts is the limitations of the technology in making observational distinctions that socially acceptable driving necessitates. We also illustrate some of the significant successes we've already achieved, including the identification of road use practices that are translatable into AV software and the development of a concept, called the Intention Indicator, for how the AV might communicate with other road users. We continue to investigate road use to uncover and describe the ways in which the social interpretation of the world can enhance the design and behavior of AVs.*

#### BACKGROUND

Nissan, just like many automotive companies (OEMs), is developing autonomous vehicles (AV) and like many other OEMs has invested in a Silicon Valley-based research center where key aspects of the AV's software systems are being developed. The particular focus for the lab is on autonomous driving for the city. From its establishment in 2013, the director of the research center, Maarten Sierhuis, maintained that a central challenge for autonomous vehicles would be effective interaction with other road users. The demands of city driving require this; urban contexts are interaction rich. A goal of Nissan's AV development, therefore, is to ensure "socially acceptable" autonomous driving. To this end, he contracted Gitti Jordan, a world-renowned anthropologist, and later built a small group of social scientists to help work out what socially acceptable driving might mean in practice.

In Maarten's original vision socially acceptable autonomous driving would be driving in which AVs, when interacting with other road users, operate smoothly and in a manner appropriate to the specific interactional context; AVs that behave neither too aggressively nor yield incessantly to other road users, neither impede the normal flow of traffic nor cause undue notice. In short, socially acceptable autonomous driving would mean that Nissan AVs would smoothly integrate into the flow of traffic and handle roadway interactions without disrupting other road users moving down the road. Determining what it would take for AVs to operate in this manner has been one of the key areas of focus for our group.

There could be other interpretations of socially acceptable autonomous driving, of course. One could aim to understand what consumers in different markets think they want from autonomous driving. Indeed surveys suggest the general public is rather divided about whether they truly want and trust AVs. Another interpretation of social acceptability would be that autonomous driving would bring a wholesale shift to some socially desirable outcome, such as new mobility options for the disabled persons. This certainly is an often touted and highly anticipated benefit of AVs by both OEMs and representatives of disability communities. Our work has yet to focus directly on either of these dimensions. Instead it is positioned in upstream research, at the very inception of the core software architectures and programming necessary for an AV to move about autonomously, ahead of the stage of engineering in which the ideas conceptualized and tested in research are “hardened” or firmed up so that vehicles can go into production and marketing strategies can be formulated. Our work thus precedes the kind of research ethnographers may more familiarly be brought into associated with the developing new car models, business strategies, or marketing plans.

Here we tell the story about how our group has undertaken research to ground the notion of socially acceptable autonomous driving in empirical investigation with a social scientific lens. We thus add to the literature on the impact of social scientific research in high-tech industries (the EPIC archives are filled with instances of this, as are numerous other publications over the proceeding two decades), exploring both some of the successes and challenges we’ve had integrating research findings with technology development.

## **APPROACH**

Excited about the possibility to impact the design of the AV and thus the mobility of the future, we embarked on our quest to help define socially acceptable autonomous driving. To have impact, we felt we needed to have a certain amount of independence in order to avoid being constrained by the starting assumptions of the engineering-driven effort. We could make a contribution to the notion of autonomous driving by doing ethnographically-informed research that we could analyze on our own terms, staying honest to our own discipline before attempting to adapt our findings to the terms deemed most relevant to the engineering team. Doing more than paying lip service to the idea that AVs must learn to behave in socially appropriate ways demands understanding what happens on the roads. And what happens on the road is undeniably social, as a broad mixture of people find their way to various destinations using a variety of transportation options, ultimately, and nearly unavoidably, by means of interacting with other road users to establish a self-organized order of traffic. And yet, whether regarding traffic as a particular form of public social life can be made useful and compelling to those actually building Nissan’s AV technologies remains a tremendous challenge, as we discuss below.

Given the broad remit for our work, we quickly initiated a program to collect a variety of data to study social practices of road use. We collected video data from stationary cameras at different city intersections that captured the interactions between road users. We supplemented this data with intercept interviews to gain participants’ sense of the nature of their interactions. We also captured first person perspective of road users, by conducting “travel-alongs”, interviewing participants on their own transportation experiences, traveling along with them as they drove or walked through local urban contexts, and inviting them to

review the video recordings with us and reflect on their experiences. (This phase of research was ably advanced by Logan McLaughlin, a Master's student at the University of North Texas and who joined us as an intern for the summer of 2015; McLaughlin 2016). We furthered our first person perspective, and diversified our focus through a brief immersive study in Sao Paulo, Brazil. We also collected data recorded by bicyclists riding through Sao Paulo and in Amsterdam.

Methods of interaction analysis have been key to our analytical practices. To date, video analysis has provided an anchor to our analysis, which has also included attention to how people talk about and reflect on their own road use practices. Beyond these more typical ethnographic modes of inquiry to develop a basis for our understanding, we have also designed some concepts and performed preliminary testing of those to advance the design and refine our empirical understanding, which we will describe below.

Naturally, we continue to advance our knowledge and thinking through review of the literature, from social and cultural histories of transportation and mobility, to walking and pedestrian life, to in-car behavior, to the development of autonomous systems and more. We have collaborated with others in aspects of this work, including members of the DesignLab at UC San Diego under the direction of Don Norman. And we have also enjoyed the opportunity to work with classes at both the University of North Texas (see Jordan and Wasson's 2015 EPIC paper for a description of this work) and San Jose State University with Jan English-Lueck.

## **DUALING SCIENCES IN THE ENGINEERING OF AUTONOMOUS VEHICLES**

What does it mean to bring an understanding of the profoundly social nature of driving and road use into the very foundations of technical development?

That driving is not only a technical but also a social skill is obvious from the moment one progresses from driving in a parking lot to driving on the public roads. Indeed for anyone to manage their movement through time and space, whether in a car, on a bike, or as a pedestrian, is a social act that involves the interpretation of cultural signs and signals, to interacting with others. Even aspects of mobility one might deem purely technical at first—accelerating, slowing down, keeping one's balance on a bicycle—must be considered profoundly social skills on second viewing.

Take the problem of locating where the AV is. Despite advanced GPS capabilities, this remains a non-trivial technical challenge (see Brown et al. [CITE] for a demonstration that driving with a GPS is also a non-trivial cognitive task). GPS is needed to help determine the precise location of the vehicle, exactly where it is on the road in relation to lane markings and curbs, for instance, but it must also integrate such information with the maps in order to determine precisely where it is. This gets especially tricky when the GPS signal isn't completely reliable (for instance, next to tall buildings) and the system needs to decide whether to trust the GPS location or the information from its more proximate sensors that track the markings on the road.

And aside from the technical challenges, there are the social considerations of a location. For a car driving down the street, for instance, an area where children are playing by the side of the road changes the sense of location for a driver significantly. We might expect that an AV take such social considerations into account when it drives down the road, yet that necessitates that the AV has a concept of what "playing children" are—that it is able to

recognize not only children, but also their behavior as playing—and that it could adjust its driving style dynamically. (The same street without playing children does not require a similar level of caution).

Or take the rules of the road. Some may seem simple and could easily be implemented algorithmically. For instance, an AV can be designed to adhere strictly to the speed limit. Yet we all know that driving the speed limit can be too fast in some cases (the aforementioned street with playing children, e.g.), whereas in others the socially acceptable way of driving is to exceed the speed limit. Similarly, an AV can be programmed to never cross a double yellow line, yet drivers often break this rule to skirt around a turning vehicle (Picture 1) or around a bicyclist on a narrow road, and refraining from doing so could engender frustration for traffic behind the AV.



Picture 1. A car crosses a double yellow line to skirt around a car waiting to turn right. The car performed the same maneuver; it is a common practice to break the rules of the road to keep traffic flowing.

Moreover there are many rules of the road that explicitly refer to a driver's judgment: just consider yield signs and four-way-stop intersections; you must yield when others are present but not when nobody is, and at a stop you must of course stop and go in the order of arrival, but what counts as one's arrival when cars stop (or keep rolling) at different distances from their lines? As AVs are driven by means of a computer program, they don't have the capability to use this required "judgment". How to act appropriately must be pre-specified by its engineers who thus have the Herculean task of defining all possible situations.<sup>1</sup> This is something that is easily done in case of games like Chess or Go where all possible legal moves are well-defined, but nigh impossible when one deals with a real-world environment such as human behavior on public roads.

And this points to one of the key differences between social science research and AI research. While there are vast differences in approach among social scientists, our

anthropological and ethnomethodological backgrounds hold in common that we regard human behavior as massively contingent on the situation, in some regards the very opposite of algorithmic. Even a relatively circumscribed form of human behavior, like movement in traffic, is dependent on a host of external circumstances. These include aspects of the physical environment—the road signs, the markings on the street—temporal factors—the time of day, the season, the weather—and social factors—perceptions and meanings of the environment and specific locations (a “neighborhood”), the presence of other road users including bicyclists and pedestrians, their ages and physical ability, the cause or reasons for people being on the road in the first place (for work? for pleasure?) and special events such as a parade or a farmers market. We also recognize that how people move about in traffic is but a small aspect of people’s life, that their presence on the road and movement from place to place is but a fleeting moment (and often a not a very notable one at that) in an on-going set of personal experiences. While such considerations of the contingencies of human behavior are a starting point for our research, for AV software engineers these considerations just aren’t very helpful, as they try to get on with the programming which is dependent on defining just what the AV should do in what situation.

## **PREDICTING ROAD USER TRAJECTORIES**

One of the things that the autonomous vehicle’s designers (AI researchers, researchers in the fields of robotics, machine learning, agent modeling, human perception, etc.), need are rules and models that give them a way to predict what other road users will do. Since we are considered the experts in human behavior within the laboratory, it seems only natural for them to ask us to provide them with such models of social actions. We have been asked to provide, for instance, state transition diagrams to specify in detail the intentional states of road users and how they change their state as a result of changing circumstances. As an example, a particularly relevant thing for an AV to notice about pedestrians would be the change from “waiting” to “crossing the street,” or for a bicyclist the transition from “going straight” to “turning left.” Detailed observations of how these transitions occur could allow the AI researchers to construct models that would specify how the AV should act, given its perception of the environment.

While we acknowledge that for the engineers there is a necessity to break down the world into agents’ behavioral states and the transitions between states, the challenge for us as social scientists is that the reduction of road user behavior into transition diagrams is not only extremely difficult, but that even when such an attempt is made, something essential about the ‘social’ seems to be lost in the process. For instance, take the case of this pedestrian waiting to cross at an intersection in San Francisco.



Pictures 2a & 2b. A pedestrian is waiting (2a) and follows a man who crosses before the light has turned to walk (2b).

Depicted in picture 2a & 2b is someone who is waiting among a group of people for the light to turn before crossing the street. The fact that he is among a group of people is relevant, since it makes it more likely that he might join the others in crossing a few seconds before the light turns to walk. One can try to break down a pedestrian's behavior into its constituent parts, consisting of, for example, body posture, gaze direction, positioning on the road or sidewalk, but that break down seems to miss something essentially social and may not help you the next time you encounter a similar, but differently executed, social action.



Pictures 3a, 3b, 3c, 3d. Two women arrive at the intersection (3a). They step into the street to cross (3b) when the one on the left slows down and looks left. The woman on the right looks where her partner on the left is pointing (3c) as they halt at the edge of the road. The car that was waiting for them now crosses the intersection (3d).

Or take these two women in picture 3, who look like they are about to cross based on the way they approach the intersection (3a). Yet after they step into the street, they stop (3b), look around and point (3c), clearly engaged in an interaction to figure out where they should go. The analysis that they are “not crossing” would be highly relevant to an AV—and note that the car's driver has easily figured this out and crosses (3d)—but depends on seeing that these women are not two individuals, but that they are together (the idea that walking down

the sidewalk is dependent on an ability to see that certain people are together was the subject of early work in ethnomethodology [Lincoln Ryave & Schenkein, 1974]).

Or, consider this picture 4 taken from the front of a moving car in southern California.



Picture 4. A woman steps out into the street behind a parked car.

The woman steps into the street with a large step behind a car. Considered simply as a moving object, as AV algorithms tend to do, she may represent a potential conflict for the car from which the picture was taken. Instead of a potential hazard, however, what we see is that she is the driver of the car parked on the side of the road and she is walking around it to get into it the left front door, where we know the driver sits in vehicles (this is in California). The woman standing on the sidewalk behind her will presumably ride shotgun adding to the Gestalt of a woman stepping out into the street to get into her vehicle.

To us as social scientists these examples demonstrate that social understanding imbues our understanding of street life, of people's behavior in traffic, and of how we perceive the world in general. This social lens is paramount, and is both more than and qualitatively different from the raw sensory input (a 'bit cloud' from a Lidar sensor or a sequence of video images taken from a camera, for instance) from which the engineers must build up the AV's interpretation of a scene. While it is certainly easy enough to discuss compelling examples with the engineers, it remains much more challenging to define the concrete implications of these social observations for their technical work.

## **MAKING PROGRESS**

We have, nonetheless, had successes, not just in generating conversations across our disciplinary bounds, but also in tangible input to the development of the AV. An example of a successful result was our analysis of a road user practice we called "piggybacking." This is a practice observed with some frequency at stop intersections. As stated, the traffic rules for four-way stops prescribe that cars may cross the intersection in the order in which they

arrive and that when there is conflict the car from the right should go first. Piggybacking is a practice that systematically breaks this rule, but in a socially acceptable manner. Drivers who piggyback travel through the intersection by jumping ahead in the order, taking advantage of the priority established by a car ahead of them, often due to their place in the queue being blocked from going by another road user. The picture below helps to illustrate.



Picture 5. Piggybacking. The Blue car arrived earlier at the intersection than the yellow car, but the yellow car takes advantage of the fact that the red car is blocking the blue car from going.

Pedestrians too were seen to piggyback on the priority established by other pedestrians. A car may also piggyback on the right of way established by a slower moving bike. In other words, there are many ways that piggybacking occurs at four-way stops, resulting in subtle yet socially recognizable and socially acceptable ways in which the order in which road users cross may be altered.

Piggybacking was readily accepted by the technical team as the kind of social behavior they were eager to program into the perception and possible action of the AV. We believe this was because the behavior was recognizable and describable at the level of decision logic that was relatively easy for the AV engineers to implement; the AV had a concept of the queuing order, and thus was able to perceive and categorize the right objects for piggybacking (i.e., the order in which cars arrive at an intersection, and what constitutes a clear path through the intersection). Since it is a practice that breaks a driving rule in order to achieve better overall traffic flow, would it not constitute quintessential evidence that the AV was driving in a socially acceptable manner if it could piggyback? Moreover, while the public understands that AVs would drive conservatively and therefore safely, there is a great concern that their robotic driving style will impede traffic flow; piggybacking is a tangible example of how such concerns might be addressed.

## **MAKING FURTHER PROGRESS, THE INTENTION INDICATOR**

Piggybacking wasn't the only practice that stood out for us at stop intersections. We also noted that when there was apparent conflict about the order in which road users can cross, people often negotiate about who should go first, in particular between pedestrians and drivers. Pedestrians have the right of way in a crosswalk, but it is not clear whether they do



when they wait at the curbside or only when they step into the road. Regardless, most pedestrians will not blithely step out into the crosswalk in front of an oncoming car. Pedestrians may seek eye contact with a car, or at least take a clear accounting of a car's behavior before stepping out in front of it. We observed pedestrians that waved cars on, and drivers that waved on pedestrians. The pictures 6a, 6b, 6c, below illustrate one of these moments.



Pictures 6a, 6b, 6c. Negotiation between a pedestrian and a driver at a four-way stop in town. In (a) the woman holds her partner back and waves the driver on, in response (b) the driver waves the pedestrians on. The pedestrians cross (c)

How could an AV interact and communicate with pedestrians, bicyclists and other drivers about the order of traffic? How could an AV express that it was letting a pedestrian go ahead, or, by contrast, that it was planning to go and that a pedestrian had better not step into the street? The AV should be able to express its intentions, but also take account of the specific situation and expectations of other road users in that setting, and adjust its behavior accordingly. We started to explore the hardware and software solutions that would allow the AV to communicate with other road users and developed a concept that we continue to adapt and refine.

The concept we developed involves a signaling system that would be added to the vehicle and could change modes according the interactional context when potential conflicts arise with other road users. With the system, AVs could engage in roadway negotiations based on the AV's perception of what vehicles and other people on the road are doing and planning to do, and adapt its behavior in order to be a good 'social partner.' The concept involves a light strip viewable from the front and side of the car with programmable signals that communicate the cars intention to other road users in the vicinity.

One way we have begun to test the concept was to try it out in our facility on a remotely controlled toy car. The objective of these tests was to get feedback on the specific instantiations of the concepts – there were many ways to execute the details and we wanted to sort out how the different options would work. Which were the most easily understood, what unintended reactions did they provoke, and so on, focusing on whether the concept would have any meaningful bearing on interactions on the road (and ultimately traffic flow). We ran the experiment twice, in one case creating a simulation of an intersection on the road

and asking people to take specific actions such as: cross the road as you normally would, linger before crossing, jaywalk, and so on. In the second test we did not direct people at all, but simply observed what people who moved around the lab experienced when they encountered the toy car.<sup>2</sup>



Picture 7. Employee interacting with the prototype remotely controlled car.

Although we had initially envisioned that the light strip would negotiate and interact with other road users in a way that would full-fill the function of hand waving, eye gaze, and the reading of ‘signals’ about other road users’ intent that we observed among pedestrians and drivers, at this point an actual interaction requires both a level of continuous perception and an ability to adapt the actions of the AV in real time that is far beyond the current system’s capabilities. Rather, our solution is for the light strip to display the AV’s intention. It has come to be called the Intention Indicator.

Despite these limitations the Intention Indicator concept was an immediate hit within the organization. Indeed it was taken up by the designers of the Nissan’s autonomous concept car, the IDS (Intelligent Driving System), and paired with other communication devices, including a technology that could detect pedestrians and bicyclists and could confirm that they had been ‘seen’ and an LED text display to reinforce messages to near-by road users. The IDS concept car was revealed at the Tokyo Motor Show in October 2015 (Picture 8).



Picture 8. The Nissan IDS concept car with the intention indicator represented by the blue light strip running along the sides and front of the car, as well as a display in the window.

There is much research to be done to develop and harden the Intention Indicator concept. Assuming the Intention Indicator could be added to the vehicle (US vehicle code regulation does not allow for any colored signaling lights to the front of the vehicle to avoid confusion with emergency vehicles, for instance), it remains to be seen if inclusion of an Intention Indicator will have meaningful impact on the road, both in the early days of introducing autonomous vehicles to the road, and over time. After all, car drivers communicate their intention largely through movement, and it is an open question whether having a light strip that displays the AV's 'intention' is actually a helpful addition to its 'natural' communication through its movement.<sup>3</sup>

In a company where real impact is measured by actual technologies that find their way into physical cars, the Intention Indicator's inclusion on the Nissan IDS concept car was a great success for our newly founded group. (Elsewhere Cefkin began to explore [2016] some of the interesting dimensions of this in terms of cultural and social views of vehicle interactions.) And it gave us a real boost in status within the organization. Indeed it is one way in which we are using social scientific research to build a new field of automotive development, External HMI.

## CONCLUSIONS

This paper has been a case study of the challenges and successes of a small social science research group in an engineering laboratory dedicated to the development of autonomous

car. We have shown how we are attempting to make an impact on the design of the automobile of the future by considering the social organization of road use, with the ultimate goal of helping create socially acceptable autonomous vehicles.

The focus of our lab is on the software for the AV to drive in city traffic. We have struggled to have a direct impact on this software development, in part because our observations of traffic often hinge on the recognition that road users make decisions based on their own social assessment of other road users and their intentions. While we have been successful to present these observations to the engineers in a compelling manner, it is far from obvious what the implications of our research should be given the limitations of the sensory systems of AVs that cannot detect the social world of traffic in all the subtle detail road users do. It is therefore no surprise that the engineering teams have a tendency to set aside our observations as compelling yet rather irrelevant, unless they can be translated into state-based diagrams, the onus for the production of which they put on us. This is challenging for several reasons, among them that the production of formal models has not been the focus of our education, that the formal models are difficult to produce without intimate knowledge of the relevant categories of the AVs software—what social actions can it recognize—and that formal models very quickly become only a poor, watered-down representation of the richness of the social world.

Nevertheless, our ethnographically inspired studies of traffic and road users have born some fruit. Not only have they resulted in actual hardware on Nissan's concept autonomous vehicle, but our research has also led to various observations that are becoming part of how the organization talks about autonomous vehicles, both within the company and to the public. We may indeed help redefine and challenge some of the fundamental categories that may seem natural from an engineering perspective, but have limited use when considered against a backdrop of the actual social reality of the methods people use to navigate the city streets. However, the interest and desire to integrate understandings of the social nature of road use into the design of autonomous vehicles remains high within the organization. That work is challenging, but it is a challenge we accept eagerly.

**Erik Vinkhuyzen** is a Senior Researcher at Nissan Research Center in Silicon Valley and specializes in video-based studies of people and technology. The focus of his current research is driving, walking and bicycling in city and suburban environments, with the goal of identifying social practices that can aid in the development of socially acceptable autonomous vehicles. Before joining Nissan Erik was a principal researcher at the Xerox Palo Alto Research Center (PARC). He also worked at NASA Ames Research Center, and the Institute for Research on Learning (IRL). He received his Ph.D. in cognitive psychology from the University of Zurich.

**Melissa Cefkin** is a Principal Scientist & Design Anthropologist at Nissan Research in Silicon Valley where she explores the potential of having autonomous vehicles as interactive agents in the world. Her work is at the intersection of ethnographic and anthropological research and the worlds of business, design, and technical system development. Melissa is the author of numerous publications including the *Ethnography and the Corporate Encounter* (editor, Berghahn Books 2009) and served in a wide range of the leadership roles, including president and conference co-chair, for the EPIC (Ethnographic Praxis in Industry Conference) organization. She worked previously at IBM Research, Sapient and the Institute for Research on Learning (IRL).

## NOTES

1. Machine learning techniques can be used, of course, making the engineering a little easier perhaps, but that is not the same as using judgment, as Button et al. (1995) have argued forcefully.
2. One of the side benefits of these experiments was that they engaged our colleagues throughout the lab in our work. Given that our facility hosts people from a number of groups and divisions across Nissan and Renault, many people had only been vaguely familiar with the nature and direction of our work until then. These testing activities provided them with a sense of our approach and the direction of our thinking, and allowed us to sense how people might react more generally.
3. Although it should be noted that when the engineers implement how the car moves, they don't do so in the first instance with an eye to what such movement communicates to other road users, a consideration that we have always stressed in our presentations to them about social acceptability of road users.

## REFERENCES CITED

- Brown, B., and Laurier E.  
2012) The Normal, Natural Troubles of Driving with GPS. Proceedings of CHI 2012.
- Button, G., Coulter, J., Lee, J., and Sharrock, W.  
1995 Computers, Minds and Conduct. Polity.
- Cefkin, M.  
2016 *Human-Machine Interactions and the Coming Age of Autonomy*. Platypus, the CASTAC Blog.  
<http://blog.castac.org/2016/01/age-of-autonomy/>
- Jordan, B. and C. Wasson,  
2015 Autonomous Vehicle Study Builds Bridges between Industry and Academia. EPIC Proceedings, pp. 24 – 35. <https://www.epicpeople.org/autonomous-vehicle-study-builds-bridges-between-industry-and-academia/>
- Lincoln Ryave A. and Schenkein, J.N.  
1974 *Notes on the art of walking*. In *Ethnomethodology: selected readings* (Roy Turner, ed.), Harmondsworth, Penguin, pp. 265-74.
- McLaughlin, L.  
2016 *Understanding Road Use And Road User Interaction: An Exploratory Ethnographic Study Toward The Design Of Autonomous Vehicles*. M.A Thesis, University of North Texas.