The trouble with autopilots:
Assisted and autonomous driving on the social road

Barry Brown
Mobile Life @ Stockholm University
University of Stockholm, Sweden
Barry@mobilelifecentre.org

Eric Laurier
School of Geosciences
University of Edinburgh, Scotland
Eric.laurier@ed.ac.uk

ABSTRACT

As self-driving cars have grown in sophistication and ability, they have been deployed on the road in both localised tests and as regular private vehicles. In this paper we draw upon publicly available videos of autonomous and assisted driving (specifically the Tesla autopilot and Google self-driving car) to explore how their drivers and the drivers of other cars interact with, and make sense of, the actions of these cars. Our findings provide an early perspective on human interaction with new forms of driving involving assisted-car drivers, autonomous vehicles and other road users. The focus is on social interaction on the road, and how drivers communicate through, and interpret, the movement of cars. We provide suggestions toward increasing the transparency of autopilots' actions for both their driver and other drivers.

Author Keywords

Autonomous cars; self-driving; video analysis; interaction; social road; automobile interfaces; human-robot-interaction

INTRODUCTION

Commercially available cars such as the Tesla Model S and the Volvo XC90 now feature advanced driving assistance and semi-autonomous driving functions, with tens of thousands of these cars being driven on roads worldwide. Along with the controlled tests of fully autonomous cars by organisations like Tesla, Delphi, and Google, we are in the midst of a large scale field-test of an unfamiliar new technology. While fully autonomous test cars have driven several million miles on American roads [28], this is dwarfed by the 140 million miles driven worldwide by Tesla's autopilot [53].

While there is growing commercial activity in this area, there is little published research available on how these new cars are interacting with their own drivers, other road users, and the road environment. In this paper we use interaction

analysis to study recordings of self-driving cars, obtained from the YouTube video sharing website. These 'in the wild' videos, made by drivers and passengers, offer a distinctive early perspective on how assisted and autonomous driving is developing on the road and being responded to by other road users. We focus here on recordings of Tesla's 'Autopilot' alongside a smaller set of recordings of Google's self-driving car. The videos offer insights into how they mainly match but also at times disturb driver expectations by manoeuvring rudely or at risk.

We pursue the notion of 'co-driving' where cars take over some functions of driving but rely upon monitoring work done by their human drivers, who intervene should the autopilot, for instance, mis-recognise interpret a road situation. All assisted and autonomous car systems are currently at this level–requiring human intervention in particular road situations. Consequently, we are interested in the interaction between driver and system, but also how these systems interact with *other* drivers. The road is a social space. For human and drivers, each movement of a car is not only functional but communicative [29,39,40,55]. Small changes (or absence of changes) in speed, direction and so on, communicate not only driver intent but also their mood and character (e.g., aggressive, hesitant, unpredictable, selfish).

Recordings made by both drivers of autonomous cars and from the perspective of other road users, pick out what they have identified as "autonomous troubles" [10]. While poorly fashioned movements by self-driving cars make visible the limitations of current driving systems, they also reveal aspects of co-driving that call for specific new design responses. This paper outlines the 'envelope of interaction' in and around self-driving cars: how interactions between cars, drivers and new systems play out, and the opportunities to make those interactions work better. Co-driving presents new challenges and opportunities for design research to shape how cars of the future will be encountered by, drive and be driven by humans.

BACKGROUND

While limited assisted and autonomous driving functions (such as cruise control [6]), have been commercially available in cars since at least the early 1970s, research at Carnegie Mellon in the 1980s pioneered autonomous driving in the form of the Navlab research vehicles [71]–

vehicles which could manoeuvre and reason about their environment as they drove around it. In recent years, interest in autonomous driving has exploded as commercial system capabilities have increased, most notably with advances in computer vision [25], and the success of research prototypes such as the Google self-driving car [26,53]. This has resulted in massive industrial investment into autonomous car research [7].

One clear impetus for industry was the DARPA challenge – a series of DARPA run self-driving car competitions. Fletcher et al [24], dissected a number of "crash" events occurring during that challenge. More recent work has engaged with problems such as challenges in image recognition [23], speed of execution, the role of machine learning in understanding road terrain [72], as well as discussions of the socio-economic impacts and challenges these systems provoke [45]. It is important to note that while there has been impressive progress [44], significant challenges remain on basic tasks–such as using vision to locate a car [11]. Moreover, 'difficult' driving conditions such as police intervention, bad weather conditions, sun glare, even turning left through traffic [43] present significant unsolved problems.

Autonomous vs. assisted driving

The NHTSA classification differentiates different levels of autonomous driving [47])–with levels 1-3 listing different forms of assisted driving (where a human must still monitor the driving and intervene if need be). At Level 4 cars are classified as 'fully autonomous' and can operate unmanned.

Much of the data we examine here involves drivers making use of the Tesla autopilot. There are three main Tesla autopilot functions. First, an adaptive cruise control that relies on distance to the car in front to accelerate and decelerate the car, an auto-steer function that uses road markings and if that fails, uses the car in front, to stay in its lane. In addition, there is an assisted lane-change function, where the driver can use the indicator stalk and, if it is safe, the autopilot will move the car into a parallel lane. The Tesla displays 'what the system has seen' in a display in the centre of the instrument screen, showing vehicles on the road in-front of the car and road markings detected by the camera. These functions sit at level 2 of the NHTSA classification as driver *assistance*. Indeed, the Tesla frequently disengages and gives control back to the driver. In a challenging test of the Tesla, "Car and driver" magazine found that the Tesla required manual intervention 29 times over a 50-mile route in mixed highway, city and rural driving [66] (it should be added that this was far superior to the other cars tested).

The Google Car more squarely fits with NHTSA description of an autonomous car (level 3), in that it is designed to travel and manoeuver autonomously through the environment, making decisions about where to go, and how to navigate with other drivers on the road. The Google Car drives around areas that Google has carefully mapped (some of it using the same Google Car). However, the Google car is still level 3 because human monitoring is required. An indication of the reliability of the software is given by reports by Google to the Californian Department of Motor Vehicles concerning "disengagements" (where either the human takes over or the car itself hands over to the driver). In 2015, the Google Car generated roughly one disengagement per 3,000 miles of autonomous driving. By comparison the human accident rate is around one crash per half a million miles [8]. While very few of these disengagements have led to accidents, a pessimistic reading is that even the Google Car, in a mapped and relatively simple road environment, is not yet equal to a human driver. This suggestion is supported by reports from Google, for example, that there are still challenges in fundamentals, such as detecting traffic lights [23,61].

Human robot interaction

Given the move toward autonomy in movement, one resource for thinking about human interaction with self-driving cars is the human-robot-interaction community [27]. Human responses to the design and movement of different robots has been extensively investigated [33,42,51,62]. For example, Alač's [3] work on children interacting with educational robots, used video analysis to understand how children differentiate between robots, toys and other humans. Alač discusses how robots move from being seen as static objects, to being perceived as 'agents'– objects that have a reactive logic to their movements. Yet there is the potential of robots going further and being perceived as social, engaging in the world of humans, with some limited access to the situation [69, 46].

Transport safety

A second academic field with direct interest in autonomous functioning in transport systems is that of transport safety–a field that has shaped (and perhaps dominated) much of our understanding of the driving task and opportunities for automation (e.g. [2,67]). Drawing on experience with aircraft, autopilot researchers have become increasingly concerned by potential safety shortcomings that increased autonomy and de-skilling of drivers could create [12]. In aircraft, the autopilot has been associated with lowered awareness of the situation, over-dependence on automated features and simple lack of practice in manoeuvring the vehicle [20]. This research highlights that automating certain system functions may not necessarily result in safer vehicles. Within HCI, a number of workshops have discussed [48,49] issues around autonomous cars, alongside research on topics such as the transition between driver and system [41] and the 'ironies of automation' [5] in that automation may lead to less effective systems overall as automatic systems often do not make problems apparent until too late [47].

Mobility

A third body of literature of use in understanding the consequences of assisted and autonomous driving is social scientific work on road mobility [15,65]. The key issue

raised by this work is that the road is a setting which is not only socially organised [29] but, is sometimes sociable [35,36], and almost constantly risks becoming an anti-social environment. Jack Katz's [37] study of road rage underlined how most drivers are highly sensitive to the actions of other drivers. Driving together in traffic is, then, a collective practice, particularly susceptible to disruption by actions that appear not to follow, maintain or repair its moral order [16,39,50,52].

METHODS

Online collections of third party video provide a rarely tapped source of data on interaction and technology-use across an array of settings, particularly around the use of new technologies (e.g. unboxing new products, tutorials, trials etc.). YouTube now forms the world's largest repository of third party video, though care should be exercised over how its videos are searched for, organised and re-used [38]. Several researchers in CHI have already used YouTube data [9,32]. In the social sciences, YouTube data fits well with established conventions of using secondary data and connects with recent initiatives in large scale qualitative datasets.

Third party recordings are a rapid method for providing first person perspectives of early adopters and user reports of general issues which can then invoke more detailed experimental or ethnographic work. In some cases, third parties provide access to events which planned data collection would find almost impossible to capture (in our case, certain drop-outs and near-misses on the road). Third party video collected from online sources has been used previously to study policing practices [34], intergenerational communication [30], argumentative communication [60] and mobile device use [9,32]. Using third party data does mean a loss of control compared to methods such as equipping drivers with cameras (e.g. [10]) but does have advantages in terms of scope, access and availability.

For this study, we collected 10.5 hours of video from YouTube. There are 69 clips come from 63 individual up-loaders from UK, USA, Germany, France, Sweden, Hong Kong, Iceland and Canada. The clips are, on average, 9 minutes long, some are very much longer–7 of our clips were over 30 minutes long. The clips were selected for their depiction of uncut footage of driving under the control of an assisted or autonomous driving system. We used a range of search terms to collect videos of the Tesla 'autopilot', the Google Car, and other videos of driving assistance in use. We discarded news reports, or commercially produced material and sought 'naturalistic' videos of cars being driven. These car journeys were recorded by drivers using in-car dash-cams and similar set-ups to those used by professional researchers. In one case an uploader recorded trips using no less than 8 cameras to capture different perspectives in and around his car (Figure 4). In contrast, about one third (26) of the clips were recorded using a single mobile phone. Together these videos contain a mix of commentaries, comic moments (e.g., friends being shocked by the Tesla acceleration), reviews, travelogues and long stretches of silent driving. While the shorter clips are dominated by performances for the camera, this is not the case in the longer clips. The 'performed' driving was itself useful given much of it has equivalences with common data collection procedures used by CHI researchers such as 'think-alouds' and interviewee responses. The majority of the videos contained Tesla self-driving functionality; although we also found two clips of the Volvo XC90, and one of a Honda Civic which both have similar 'pilot assist' functionality. In terms of the Google Car we found a smaller corpus of 11 minutes from 9 clips, along with clips from a Google SXSW presentation of the Google Car, notable for including six video clips of 'interesting incidents' involving the Google Car.

For our analysis, clips were downloaded broken into smaller fragments which were tagged with keywords and transcribed. The fragments were analysed in group data sessions. The analysis focused on looking for examples of other drivers reacting to autonymous functions, on road 'troubles', driver interventions and interactions between the car and other drivers. In further individual and group data analysis sessions we surveyed the clips, editing the video down to 93 short (around one minute) edited clips for focused analysis, of which we then selected 34 clips for detailed analysis. Our analytic approach builds on earlier analyses of how drivers co-ordinate and manage their driving with each other [10,29,40,50] and on Watson's [73] studies of specialized forms of driving.

ETHICS

While nearly all of our clips make use of dash-cameras, around a fifth of the recordings appear to have been made by drivers holding their phone while they drive. While a 'new' technology (like the Tesla) is interesting to many YouTube viewers, the desire to record can encourage drivers to film in less than safe conditions. While we did not gather the data, we have concerns over drivers driving unsafely to produce recordings. A related concern is that the clips that we have analysed often feature driving outside the guidelines issued by Tesla about the use of their function, such as using the Autopilot on non-divided roads. As researchers, however, we would argue that examining and presenting this data alongside 'safe' practice is important in order to understand driving as it is actually occurring [10].

FINDINGS

As mentioned above, our goal in looking at these videos is to begin outlining the new 'envelope of interaction' surrounding self-driving cars, and thereby to understand how the interactions of both drivers and other road users are developing as autonomous cars become more prevalent.

Our results span two aspects of autonomous driving. First we discuss the role that assisted driving plays in reshaping driving practice, and in particular the development of 'co-

driving' which is a form of involvement of both driver and autopilot. We examine how drivers adapt their driving to take advantage of the features and limitations of their car's new systems. Second, we look at the relationship between the autopilot and other road users and a selection of the problems that arise through the lack of accountability and understanding that autonomous and assisted cars communicate.

Driving with the autopilot
An initial striking feature of many of the videos is that they show cars being manoeuvred with no hands on the steering wheel. Indeed, some of our drivers start their clips with a slow passing of the hand around the wheel, reminiscent of a magician's 'look no hands' gesture. However, assisted driving is certainly not completely 'hands free'. Drivers monitor their car because they have experienced it making unpredictable and potentially dangerous movements, as well as dropping control in more predictable ways:



```
"The first time it went through this road it would
lurch. I had to grab the steering wheel and it was
like, every one, I had to be ready. Eventually I
could keep auto-steer on and I would hold the
steering wheel. Sometimes it would deactivate
because it really wanted to get off and sometimes
it wouldn't. […] I am still ready to take over.
Here comes another "Do I take it no I don't take
it". I'm not sure that the exact voice autopilot
thinks in but it's probably close. Ehm. See it
learnt [images] Oh there. It was a pretty big
lurch. That wasn't me at all. So it lurched over
but then it corrected itself, so apparently it
doesn't know that one very well"
```

Figure 1: Lurching (arrow indicates brief shift in direction)
(source Tesla EVangelist)

Monitoring 'lurches' and lines on the dashboard display
Figure 1 offers an example of the Tesla autopilot driving on the road. Here a driver is reporting at some length on how the autopilot 'thinks' – describing 'lurches' it makes when it drives past different exit lanes it is unsure about. Since the autopilot has in the past moved toward exiting, the driver is now prepared to intervene, having learnt to detect the autopilot's exit-move at its first initiation, and thus being ready to force the car to stay in the correct lane. This first clip demonstrates how the autopilot requires a level of monitoring about what will happen next. Indeed, in this clip the driver also goes on to discuss the 'higher level' functions of the autopilot where the system 'learns' over time, so as to not take unwanted highway exits. The discussion shows not only that the driver is monitoring his

car as it drives, but also that he is understanding it through its movements on the road system. Through those moves he is learning about how it learns, and how it makes some of its decisions in relation to the road's features.
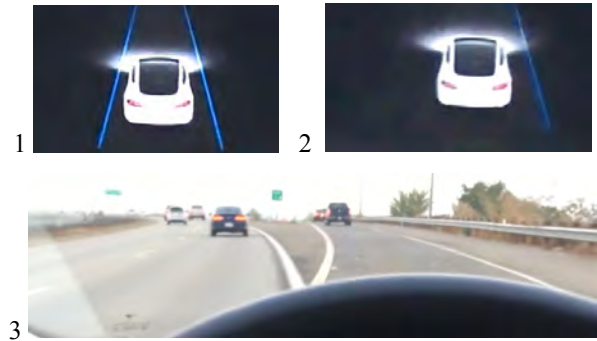

```
"Oh. it exited me I didn't want to exit. […] It
made a mistake"
```

Figure 2: Disappearing lane markings & exiting (source tuan)

Part of co-driving with the autopilot is anticipating or noticing situations where it will make mistakes. In Figure 2 the autopilot followed the edge-markings of an exit lane as a continuation of the current lane and the car exits the highway incorrectly. The edge-markings it is following are displayed on the dashboard and disappear when it cannot track them (see figure 2). The driver is left to exclaim that the car 'exited me', and 'I didn't want to exit'. Failing to notice and correct the autopilot's misrecognitionmis-interpretation, visible on the dashboard, then leads to taking the wrong exit, with all the potential navigational troubles this involves.

Drivers learn to co-drive and intervene in order to prevent common autopilot mistakes, e.g. grabbing the steering wheel when the car lurches as in and/or follows a line or vehicle in error (as in fig. 1 and 2). Co-driving with an autopilot thus includes monitoring the movement of the car and watching the dashboard for instances when the autopilot has 'lost its way'

Dealing with 'troubles'
Drivers learn how much monitoring the autopilot needs, and adopt expectations for situations that lead to autopilot error. In our earlier work [10] we discussed how GPS does not make drivers 'docile', instead drivers incorporate new technology into their driving: watching for mistakes and/or trying to decode the 'troubles' that an imperfect technology creates in a real world context. Likewise, drivers of the Tesla adapt their analyses of the road situation. They develop an understanding of how the autopilot is sensing the road, and, correspondingly, the autopilot's capacities and limitations. In the videos, drivers discuss how they have become aware of the features of the road that lead to problems for the autopilot (such as sunshine, blind summits, blind corners). Similar to other forms of specialised driving (e.g. driving in forests, rally driving) autopilot drivers acquire a sensitivity to what features of the road are

'dangerous' in relation to the autopilot, even though they would cause minimal challenges to a typical human driver.

This awareness of autopilot capacities and incapacities produces distinct 'co-driving' practices (cf. [58]) that can also be seen in how drivers use the controls available to the driver while in autopilot. In particular, the autopilot allows drivers to set a 'target speed' which it will attempt to accelerate toward. While it is capped at the speed limit plus 10 mph, we observed drivers changing the target speed according to an assessment of the local road conditions. So, for example in Figure 3 a driver slows down the car as they approach a corner. Drivers also adjusted their speed to control overtaking behaviour such as when approaching hazards (e.g. tunnels & bridges) or tight curves in the road.

"Around this corner I'm going to slow down the cruise control cause it's a blind corner"

Figure 3: adjusting speed on corner (source Mark Jackson)

Figure 4: overtaking manoeuvre delayed until motorcycle passes (source Furious S8)

Drivers monitored the road and traffic to time their overtaking manoeuvres. In the Tesla, the driver can initiate a semi-automatic lane change by selecting the car's indicator lever. In Figure 4, for example, one driver overtakes two trucks and, just as he is about to return to the right lane, a motorcycle appears and overtakes the car on the inside. The driver holds their hand at the indicator, only flicking it, once the motorcycle passed the car. Waiting on the motorbike demonstrates a form of driving that is sensitive to the abilities of the autopilot, as well as remaining engaged with the environment.

Appropriating the autopilot
Co-driving is reminiscent of Suchman's description of 'human machine reconfigurations' [68] where both the driver, and the car jointly produce driving – a joint action that is accentuated rather than eliminated by the autopilot. Drivers discussed different expectations about driving, about what the autopilot can 'learn', and what situations are suitable or not for using the self-driving features. One driver discussed how the car can conduct the stop-go driving required in a slow-moving traffic queue:

"An interesting use for autopilot that I've found is if you're waiting at a long wait for the four-way stop sign […] autopilot will just keep an eye on the car for you and do stop and go until it's your turn. Just remember when you finally approach the stop sign the car will just continue to follow the one behind it. At this current version of software 7.1 the car will not stop or see stop signs […] please, please, exercise caution. [...] I will demonstrate. The car in front of me goes, now my car will continue to go, but since there's nobody in front of me, I will take control."

By selecting the autopilot, he had found that it could undertake part of the repetitive tasks of queuing. Yet, as this user is also aware, the autopilot edges the car forward by piggybacking on the car-in-front's stopping and starting (which, of course, is also piggybacking on the next car-in-front). Consequently, he stresses the need to take over from the autopilot when it reaches the intersection.

Tesla advises drivers to only use autopilot when a car is moving on a highway, with the autopilot maintaining speed. Yet, here, the driver demonstrates how he has discovered another use, deploying it at queued stop-intersections. While the autopilot can "stop and go", the driver makes further sense of the queue with their greater awareness of the situation [68], taking over the car when it becomes necessary, "Since there's nobody in front of me I will take control." Other drivers in our collection remarked upon (and demonstrated) the autopilot queuing at controlled stops (such as traffic lights), by simply following the movements of the car in front. In many cases, then, autopilot can undertake otherwise dull driving tasks–with supervision.

Driving safely
As one of the first cars on the road with advanced autonomous features, the safety of the Tesla has generated considerable media attention. In particular, a dispute

between Mobileye (who make image recognition software for assisted driving) and Tesla arose following a fatal accident involving a Tesla that failed to recognise a car with a trailer crossing its path. Tesla claimed that the autopilot was safer than human highway driving, while Mobileye commented on the danger of "consumer and regulatory confusion [resulting in] mistrust that puts in jeopardy technological advances that can save lives[1]." Clearly there is dispute even amongst manufacturers about the appropriate situations for safe autopilot use.

The Tesla car has extra safety features, such as auto-brake, that can engage during ordinary non-autopilot supported driving to prevent accidents. One clip in our corpus involves a Tesla driving at night, where a car drives suddenly across the lane in front of the Tesla, with the car automatically and suddenly braking safely. Yet, these features are by no means perfect. In a contrasting clip, the Tesla creeps along a highway at a slow speed approaching a stationary truck and trailer which are partially positioned on the hard shoulder. The stationary truck is not detected either by the autopilot (which should slow down), or the emergency braking system. The Tesla collides with the trailer, fortunately at a low speed. The problem appears to be one of differentiating between objects in, and not, in the car's path. Trailers seem to cause particular problems for the Tesla's visual system at this point.

Indeed, many of the safety-related videos in our corpus feature 'drop outs' from autopilot control where the car occupants are put in danger. These videos are given dramatic titles; "auto pilot auto steer almost crash", "Tesla Auto Pilot fail!", "Tesla Autopilot tried to kill me!" (Figure 5).

Image 1. *Autopilot on, sun glare on camera*
Image 2. *Autopilot alarm starts*
Image 3. *Car veers suddenly left into lane of approaching car.*
Image 4. *Driver grabs wheel and turns quickly right, "Whooo-a"*

Figure 5: "Tesla Autopilot tried to kill me!" (source RockTreeStar)

In Figure 5, a Tesla driver documented a near miss with oncoming vehicles when the autopilot drops out due to sun glare. The seriousness of the failures underlines the problem of drivers establishing and understanding the situated dangers of autopilot use and becoming attuned to those settings where the autopilot will struggle. As we have noted earlier in the paper, many drivers are learning how to co-drive with autopilots and which elements of the labour of driving they can delegate to the autopilot.

The autopilot and other drivers on the road
So far, we have examined the relationship between autopilot and the car's driver. In this section we move on to the relationship the autopilot has with *other* drivers. It is here that the 'communicative' aspects of car movement become all the more significant given that the autopilot is analysing the movement of vehicles which communicate the intended-actions of those other drivers, and those other drivers are also analysing the movement of cars controlled by autopilot as communicating their intended-actions.

Offers and rejections
Much of the Tesla's autopilot driving passes without incident, with the car remaining in lane correctly, and with cars moving around it. In our corpus there are nevertheless occasions when the Tesla causes 'troubles' in traffic. One aspect of watching the autopilot make its way through traffic is that, like many AI systems before it, it lacks equal access to other drivers' intended-actions. For example, in Figure 6, the autopilot changes lane causing minor troubles for the car behind. The move is initiated by driver using the indicator lever to issue the command to move to the left lane (which also activates the flashing indicator). The driver's action itself is in response to his noticing a white truck trying to merge onto the highway from a slip-road. The episode, then, starts with a demonstration of what "good" driving is - the graciousness of letting drivers in from a slip road.

In the moments that these first actions take to unfold (and as is typical for heavy traffic) a silver car approaches from behind, in the desired lane to the Tesla's left. The accelerating car, however, on seeing the indicator flashing on the Tesla, stops accelerating – leaving a space on the road and thereby making an 'offer' of a gap to the Tesla (Figure 6, Image 1). From the silver car's perspective, the Tesla appears not to be taking the offer (it makes no change in trajectory). The absence of response leads the silver car to remove its offer by starting to close the gap. At the last minute the Tesla moves across into the lane and 'cuts the car up' (Image 2). In the clip then while the driver has been considerate toward one driver, the autopilot has potentially upset another in the next moment. The action produced by the autopilot's robotic coordination might not 'bother' another autopilot, but is recognisable by a human driver as rude. The driver of the Tesla also notices this, '*well* we pulled in ahead of that guy […] ', his "well" preface marking the trouble he's about to report.

---

[1] http://www.mobileye.com/press-room/mobileye-responds-to-false-allegations/

*"Well… we pulled in ahead of that guy and from what I saw it wasn't something he was exactly encouraging"*

Figure 6: Ignoring the offer (Source: Garth Woodworth)

With its limited access to the traffic situation, the autopilot could not, of course, be expected to recognise the offer as considerate and its delay as marking that it declined an offer and its subsequent abrupt lane shift potentially upsetting the silver car. Our point rather is that the 'rudeness' of the autopilot demonstrates that driving involves norms of mobile conduct with others. Autopilots have been introduced into a social situation, and while at times they can produce safe and reliable driving, they also at times can fail to recognise and display normative features of driving in traffic, e.g. making and accepting offers. This is not to say that human drivers are consistently aware of other drivers but, at their best, they produce movements that are not merely agented–responsive to conditions to satisfy goals–they are *social*, responsive to a background knowledge of what manoeuvres are thoughtless, polite, safe in traffic.

Analysing what just this gap could mean
Related to autopilots not being able to access what fellow members of traffic are up to, they then have limited access to how their own movements on the road are being taken up by other drivers. For an example of failing to understand what a particular spatial arrangement potentially means, we can pass over to the Google Car–and a clip extracted from a talk by the principle engineer of the Google Car (Chris Urmson). In an overview talk he outlines how the car must cope not only with the routine but also the exceptional acts of other drivers on the road.**Error! Reference source not found.**

Figure 7 shows both a representation of the autopilot's model of the situation and the view from the rear of two Google cars travelling together. Another car comes out of an exit on the left, crosses across two lanes and wedges itself orthogonally to the traffic flow, in a space between two self-driving cars. The narrator uses the driver as an example of eccentric driving practices–"people who do, I do not know what". Yet the narrator is forgetting that other drivers do not see two Google cars travelling in a convoy. They see a second car make a noticeably slow approach to the car in front, projecting. Moreover, while approaching slowly, the second car is also leaving a large gap between it and the car in-front. A gap that is straight ahead of the car that is stuck in a yard. It appears to then be potentially offering a way out for the stuck car. In fact, the Google Car does deal with the car's acceptance appropriately, by stopping leaving sufficient space for sideways car.
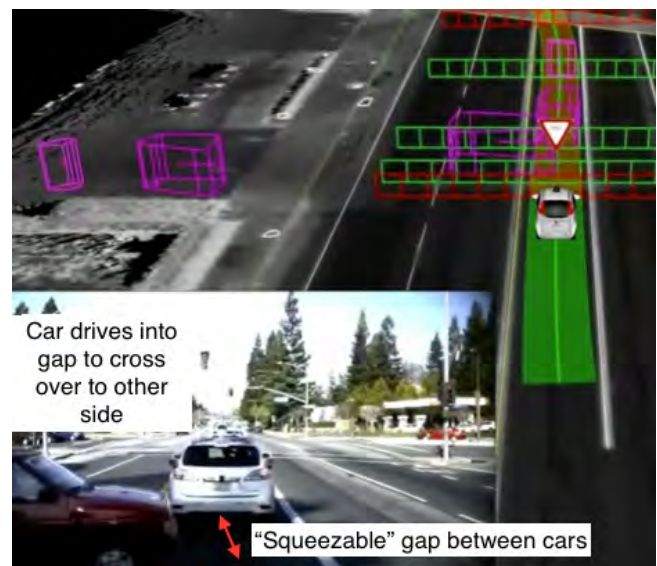
"And of course we have people who do, I do not know what sometimes on the road. Like this guy [image] pulling out between two self-driving cars. (laughter) You have to ask what are you thinking"

Figure 7: Seeing the gap (source: SXSW)

Gaps between cars (such as that between the two Google Cars) are not just 'safe/unsafe' distances on the road, but are used in displaying attitudes, maintaining norms and so on. Moving quickly to form a minimal gap with a car in front can be seen by other drivers as 'denying' any requests to merge or move through the gap. In this case, by contrast, slowly driving forward and leaving a gap, the second Google Car appears to have been recruited to help the stuck car. While the clip's narrator formulates the other driver's move as obscure and perhaps dangerous, to us the situation is recognisable as the give and take encountered on busy roads. Drivers frequently push into gaps where they can find them, and 'bend' traffic laws (such as turning left and crossing the highway margin). Congested traffic in complex

road layouts rely on driver generosity, negotiation, compromises, fiat, bullying and more.

The 'trouble' that the Google Car's deceleration and inter-car vehicle spacing occasionally produces also occurs with Tesla's autopilot, in that it also leaves 'gap' which are not attuned to local social circumstances. In our collection the 'gap' is frequently taken by other drivers, resulting in the autopilot slowing down and pulling back to create the same space, into which other cars then drive. Local judgements are required because while good driving mandates that one leaves a sufficient space in front of one's car (for braking distance), in congested traffic situations drivers typically leave less inter-vehicle space, and, as we have seen, depending on the timing and spacing of its appearance, other drivers make inferences about the gaps as offers, opportunities, refusals or avoidance.

For example, in one clip from our collection, an autopilot approached an entry ramp on the highway, and 'allowed' a car onto the highway in front. Of course, the car was not 'allowing', but simply because of its slower speed it left a 'gap' which was visible to the merging driver as an opportunity. The gap itself was tight and the merging driver hesitated slightly before moving to take it. The merging driver quickly moving into the next lane as soon as it was clear. Maintaining a one car gap with the car in front 'offers' a space to other drivers, but since the gap is not an offer to the autopilot (just a gap to be maintained) the 'offer' is not then adjusted by the autopilot on recognizing it as being 'taken'.

Figure 8: Near miss as two lanes merge (source: Garth Woodworth)

Accountability: Driving as if other people are watching
A concept that we draw on to understand vehicle movement in traffic is *accountability*. In common use accountability is often equated with responsibility, in ethnomethodology accountability is taken to be deeply engrained in interaction: human actions are produced with an orientation toward their observable-reportable character. We assume

that each party can witness one another's action as those actions and call them to account for them. Accountable activities become important in road traffic because vehicles use shared mobile methods which make their actions observable to other vehicles, and their accountability is reflexively tied to each emerging traffic situation. Failing to recognise what other drivers are doing can produce not only rude, but also dangerous driving, especially when other drivers expect their movements to be understood [36]. In Figure 8, a Tesla on autopilot approaches two lanes that are about to merge into one. Before the lanes merge a truck quickly overtakes in the 'fast' lane (Image 1) while the autopilot holds its speed at 90 km/h (the speed limit+10km). During the section where the lanes are merging, a second car accelerates up in the fast lane, also attempting to overtake in front of the Tesla (Image 2). Yet, the autopilot maintains its speed–is slower than the overtaking car (allowing it to overtake), but it does not fall back in the rapidly narrowing road (Image 3). The other drivers expect the Tesla, because it is in the lane being merged to defer to the cars continuing lane. The autopilot mechanically maintains the set speed and trajectory, only narrowly avoiding clipping the overtaking car. The Tesla fails to recognise where it ought to decelerate in order to defer. Our point, again, is not the failings of the autopilot as such (the operations are outside of its designed specifications), but that vehicles make local judgements on which car should progress and which should defer. In this case the car is making an accountable dash to go in front of the Tesla before the road narrows to a single lane. Yet the projected action is not taken into account by the autopilot, and moreover, the autopilot cannot grasp that its own failure to close the gap behind the truck is indicating that it is going to 'let the other car in'.

As a driver, one's movements in traffic on the road are *accountable* to others – they are monitored by others. They are treated as first, having a local history, second, in their course and, third, as indications of future intent. They are treated as (mostly) responsive to the ongoing actions of fellow members of traffic. Autopilots do not share the same sense of accountability and that is part of the troubles they raise for other members of traffic.

Hesitation and creeping
Last, let us look at a traffic situation where precise timing and small movements are used by multiple drivers–driving through a four-way stop-intersection. In Figure 9, a driver with a dash-cam (from which they regularly upload to YouTube) spots the Google Car and follows it. The Google Car stops at a four-way junction in a manner which looks appropriate given that cars are approaching from the sides. The Google Car begins to edge forward to take its turn but then brakes suddenly while still at the start of its manoeuvre. The car behind (from which the driver is recording the Google Car) also starts to move, but then also has to brake suddenly, when the Google Car breaks. Accidents at junctions like this are commonplace with

novice or beginner drivers, because their hesitation at a junction is not anticipated by drivers behind. In this case, two cars go in front of the 'slot' the Google Car should have taken at the intersection, with the second 'late to junction car' driving onto the intersection after the Google Car is moving, causing it to suddenly brake when it sees that its path forward is blocked.
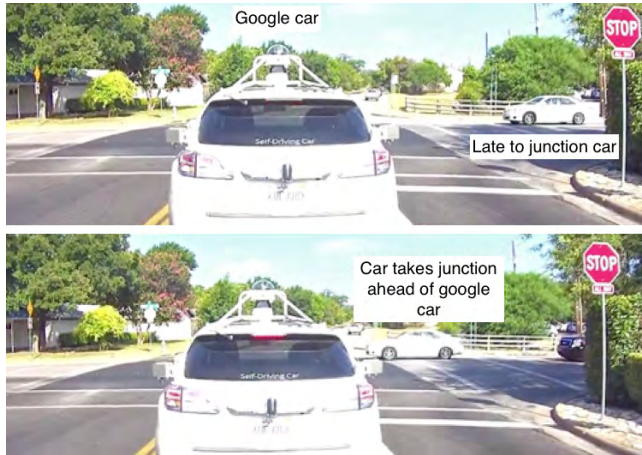
**Figure 9: Google Car hesitates at a four-way junction
(Source: The Dashcam Store)**

Research on four-way stop-intersections has underlined the importance of 'creeping' into the intersection to reserve a turn [13,14]. The creeping acts to show that a driver is attentive to the situation, and that they are ready to precisely take their 'slot' on the four-way intersection. By *not* creeping the Google Car finds itself 'cut up' by a car that slows but then accelerates to cross the intersection. In this case, it is the *lack* of motion by the Google Car, at just that moment of the other car slowing, that is seen as hesitation, which then leaves the 'slot' to be taken by that other driver. Naturally programmed to drive safely the Google Car frequently appears to other drivers as a hesitant and slow driver. In other clips, we saw similar cut-ups and drivers tailgate the Google Car to 'push' it through junctions (e.g., in Figure 8).

DISCUSSION
Our goal here has not been to review the current generation of autonomous vehicles and driver assistance systems, they are changing rapidly and require careful consideration by researchers with the relevant technical expertise. Rather, taking inspiration from earlier reflections on AI [18,19], we have examined what happens when AI meets the complex social, embodied world of driving in traffic, and hinted at the challenges of designing autonomous vehicles for this world. In our discussion we focus on three issues. The first concern how drivers will have to remain engaged with autopilots, though in new ways that will reconfigure how driving is accomplished. The second engages with the transparency of autopilot actions and designing for the accountability of its activities to drivers inside and outside the car. The third turns to the importance for sense-making

of pre-actions and considers the troubles raised by total transparency.

Engaging with assistance
In an earlier paper [10] we documented the troubles raised by GPS navigation systems for drivers. Given that driving assistance systems rely, to greater or lesser degrees, on navigation systems, our earlier arguments have renewed relevance. 'Natural normal troubles' of GPS use arose that needed to be dealt with as part of wayfinding. For the drivers we observed, successful navigation involved 'active driving'–listening and judging the GPS's instructions. Similarly, the co-driving we have described here depended upon the active engagement of a driver. The assisted driving systems often made minor mistakes, suffered drop-outs, or upset other drivers. Successful co-driving requires retained levels of engagement that vary with the situation. Emerging situations thus needed ongoing assessment by the driver as to whether the driver is likely to be called upon to intervene in an executive role.

Debates around the safety of assisted and autonomous driving seem only likely to intensify as the technology becomes both more common and more advanced [75]. Although piloting planes happens within a heavily regulated transport system we can still learn from the longer history of safety in that sector [1,56]. While autopilot use has made plane travel much safer, they heave created new troubles for commercial passenger plane pilots. For instance, when autopilot systems fail there is the danger of the 'startle' effect, where pilots that have not been involved in the situation lose valuable time in their initial panic and absence of understanding of what has led to an alarm. Designing for co-driving will require attention to how to keep drivers engaged, so that they are ready to return to driving when necessary. Designing for situated driver-assistance is far safer than designing for near-autonomy.

For other drivers on the road, the autopilot creates trouble in part because when assumed to be just another driver it behaves unexpectedly, leading to confusion and frustration. As we have described above, troubles in interaction with other drivers, such as unexpected lane changes, hesitations and losing a turn in traffic, emerge not from incompetent or faulty autonomous driving but rather vehicle movements that inadvertently communicate intended courses of action which are then not taken or a different action emerges than expected. While these troubles did not lead to accidents, they do lead to unnecessary disruption to traffic. One simple solution, which the current iteration of the Google Car has adopted by becoming a 'bubble car', is to be visible at-a-glance as a 'self-driving' vehicle. Obvious external markers shift other drivers' expectations of the predictably and meaning of a vehicle's actions.

Transparency and accountability
In discussions of algorithmic systems [21,59] more broadly, there have been calls for systems to be more *transparent* in showing their current state and ongoing processes. Notable

here is Pasquale's critique of the lack of 'algorithmic transparency' in AI systems [57]. Recent work has, however, questioned whether systems, such as artificial neural networks, could ever be transparent in a form that humans could easily make sense of or, equally, that attempting to achieve transparency will make unrealistic demands on regulators or users [4]. Transparency is a problem for both the driver within a vehicle and other drivers outside of that vehicle. For the latter, a potential solution to the problem of transparency is to build on how car drivers already make their actions transparent to other drivers. Researchers at Nissan have explored how autopilots could show recognition of the acts of other road users through extensions on the car indicator [14] thereby building on existing signalling system for communicating with other drivers.

From the findings in this paper our suggestion is that for both the driver inside the car and for other cars, the movements of the vehicle made by the autopilot should be designed with an understanding of how those movements will be understood by both the driver and by other vehicles involved in emergent situations. Cars would then make movements not just as a by-product of travelling towards their destination, but as a central part of how they communicate. Small movements can show that they have, noticed another vehicle, are ready to take their slot, or, are denying a slot to another driver.

What would drawing on existing embodied communicative practices by drivers mean in practice? Rather than attempting to offer visualizations of its state, the autopilot would–through its motions and related actions–demonstrate what it is doing to others. So, for example, when merging into traffic an autopilot will turn the car slightly to direct it towards a parallel car in order to indicate the intention to enter in front of that car. This movement is followed by a wait for a response from the other driver, and if one is not forthcoming an autopilot then decelerates to enter behind that car (see figure 8). Developing these movement protocols for autopilots requires an explicit focus, then, on current communicative practices of vehicles in traffic. In other words, building for transparency asks us to look back to early CSCW discussions where the notion of accountability prioritised designing systems which produce accounts that are useful as resources for users from without (driving in traffic) rather than within (AI states & processes) [17].

The social road
Work on multi-agent models for autonomous vehicles has recently engaged with systems that discover features of different drivers [74,54], as well as understanding the effects that drivers' movements have on other road users. As Sadigh *et al* write, *"other drivers do not operate in isolation: an autonomous car's actions will actually have effects on what other drivers will do"[63]*. In this paper we have attempted to move these arguments further – to explore how thinking of the road as inherently social, as well as logistical, should influence the technologies we build. Here, we have only started to outline the ways in which drivers make sense of and make use of the actions of others.

One valuable resource is interactional work which has documented how numerous actions (such as offers, invitations and rejections etc.) are signalled before they are underway with "pre-actions" [64, chapter 4]. For example, when collaboration involves tools, the project-ability of a request is relied upon to support close collaboration [70]. Parties collaborating in a situation need to see in common what others will be doing next, so (for example) picking up a microphone can be seen as a "pre" to using that microphone to make an announcement [31, chapter 4]. In driving in traffic there are many actions that other drivers make that can be seen as 'pre' moves, moves that are available to other drivers. For example, in our discussion above of the Google Car at a stop intersection, 'creeping' acts as 'pre' to crossing the intersection.

An added social complexity, however, is that while driving is often cooperative it can also be competitive and conflictual. Consequently, making a driver's intentions transparent to others, on congested roads, will, at times, act against the interests of a driver. Katz [37] quotes a driver who explained that they never indicated when driving in LA, as that would just be 'giving information away to the enemy'. Accordingly, a completely transparent autopilot might end up being open to exploitation by ruthless and/or harried human drivers. So, for example, a self-driving car that avoids collisions by always slowing down could be safely 'cut up' by a human driver, because they can assume that the autonomous car will always stop. With all safety technologies, there is also the challenge of 'risk compensation' [22], where improved safety is adopted by drivers to then drive faster, or more aggressively. In terms of self-driving cars it may be that this 'risk compensation' will occur not only with those in the car being controlled, but also with other drivers who might take advantage of the cautiousness of the autonomous driving system.

CONCLUSION
In this paper we have documented the challenges that drivers with autopilots experience on real world roads. Outside of test environments autonomous cars are not things of wonder, but are just another car progressing through the social order of traffic. Using videos of autonomous and assisted cars selected from third party recordings we have engaged with both how HCI can help design better autopilots and the dangers that badly designed autopilots might create. Rather than assuming human drivers will disappear with the rise of autonomous vehicles we have shown instead how it is the nature of driving that is reconfigured through allowing the autopilot to undertake simple tasks as a co-pilot, with situated supervision.

REFERENCES

1. Airbus. Flight Operations Briefing Note: Standard Operating Procedures Optimum Use of Automation: FLT_OPS – SOP – SEQ 02 – REV 03 – JUL. 2006.

2. Motoyuki Akamatsu, Paul Green, and Klaus Bengler. 2013. Automotive Technology and Human Factors Research: Past, Present, and Future. *International Journal of Vehicular Technology* 2013: e526180. https://doi.org/10.1155/2013/526180

3. Morana Alač. 2015. Social robots: Things or agents? *AI & SOCIETY*: 31, 4, 519-535. https://doi:10.1007/s00146-015-0631-6

4. Mike Ananny and Kate Crawford. 2016. Seeing without knowing: Limitations of the transparency ideal and its application to algorithmic accountability. *New Media & Society* online first https://doi.org/10.1177/1461444816676645

5. Lisanne Bainbridge. 1983. Ironies of automation. *Automatica* 19, 6: 775–779.

6. K. Bengler, K. Dietmayer, B. Farber, M. Maurer, C. Stiller, and H. Winner. 2014. Three Decades of Driver Assistance Systems: Review and Future Perspectives. *IEEE Intelligent Transportation Systems Magazine* 6, 4: 6–22. https://doi.org/10.1109/MITS.2014.2336271

7. Mark Bergen. 2015. Meet the Companies Building Self-Driving Cars for Google and Tesla (And Maybe Apple). *Recode*. Retrieved September 20, 2016 from http://www.recode.net/2015/10/27/11620026/meet-the-companies-building-self-driving-cars-for-google-and-tesla

8. Kavi Bhalla and Dinesh Mohan. 2016. Understanding the road safety performance of OECD countries. In *Transport planning and traffic safety: making cities, roads, and vehicles safer*, Geetam Tiwari and Dinesh Mohan (eds.) CRC Press, UK, 1–15.

9. Mark Blythe and Paul Cairns. 2009. Critical Methods and User Generated Content: The iPhone on YouTube. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '09), 1467–1476. https://doi.org/10.1145/1518701.1518923

10. Barry Brown and Eric Laurier. 2012. The Normal Natural Troubles of Driving with GPS. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '12), 1621–1630. https://doi.org/10.1145/2207676.2208285

11. Cesar Cadena, Luca Carlone, Henry Carrillo, Yasir Latif, Davide Scaramuzza, Jose Neira, Ian D. Reid, and John J. Leonard. 2016. Simultaneous Localization And Mapping: Present, Future, and the Robust-Perception Age. *arXiv preprint arXiv:1606.05830*. Retrieved September 20, 2016 from http://arxiv.org/abs/1606.05830

12. Stephen M. Casner, Edwin L. Hutchins, and Don Norman. 2016. The challenges of partially automated driving. *Communications of the ACM* 59, 5: 70–77.

13. Melisa Cefkin and Erik Vinkhuyzen. 2016. Developing Socially Acceptable Autonomous Vehicles. *Ethnographic Praxis in Industry Conference Proceedings*, p. 522–534 https://www.epicpeople.org

14. Melissa Cefkin. 2016. Human-Machine Interactions and the Coming Age of Autonomy | Platypus. *Platypus: the castac blog*. Retrieved September 20, 2016 from http://blog.castac.org/2016/01/age-of-autonomy/

15. Tim Dant. 2004. The driver-car. *Theory, Culture & Society* 21, 4–5: 61–79.

16. Elwys De Stefani. 2015. Formulating instructions. First actions in driving lessons. Retrieved September 20, 2016 from https://lirias.kuleuven.be/handle/123456789/533703

17. Paul Dourish. 1997. Accounting for system behaviour: Representation, reflection and resourceful action. *Computers and Design in Context, MIT Press, Cambridge, MA, USA*: 145–170.

18. Hubert Dreyfus, Stuart E. Dreyfus, and Tom Athanasiou. 2000. *Mind over machine*. Simon and Schuster. Retrieved September 20, 2016 from https://www.google.com/books?hl=en&lr=&id=e9W9m_4q4pYC&oi=fnd&pg=PR9&dq=mind+over+machine&ots=FiOEE8W6O2&sig=LkUriFSZzKIYLpgAn5cqkva2NJs

19. Hubert L. Dreyfus. 1992. *What computers still can't do: a critique of artificial reason*. MIT press. Retrieved September 20, 2016 from http://books.google.com/books?hl=en&lr=&id=7vS2y-mQmpAC&oi=fnd&pg=PR9&dq=info:tDOnc5P0VbAJ:scholar.google.com&ots=1U2dRHuvTz&sig=8YrnhoWEpclD-PdrCp3CzQtIfNY

20. Mica R. Endsley and Esin O. Kiris. 1995. The out-of-the-loop performance problem and level of control in automation. *Human Factors: The Journal of the Human Factors and Ergonomics Society* 37, 2: 381–394.

21. Motahhare Eslami, Aimee Rickman, Kristen Vaccaro, Amirhossein Aleyasen, Andy Vuong, Karrie Karahalios, Kevin Hamilton, and Christian Sandvig. 2015. I always assumed that I wasn't really that close to [her]: Reasoning about Invisible Algorithms in News Feeds. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, 153–162. Retrieved December 21, 2016 from http://dl.acm.org/citation.cfm?id=2702556

22. William N. Evans and John D. Graham. 1991. Risk reduction or risk compensation? The case of mandatory safety-belt use laws. *Journal of Risk and Uncertainty* 4, 1: 61–73.

23. Nathaniel Fairfield and Chris Urmson. 2011. Traffic light mapping and detection. In *Robotics and Automation (ICRA), 2011 IEEE International Conference on*, 5421–5426. Retrieved September 20, 2016 from http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=5980164

24. Luke Fletcher, Seth Teller, Edwin Olson, David Moore, Yoshiaki Kuwata, Jonathan How, John Leonard, Isaac Miller, Mark Campbell, Dan Huttenlocher, and others. 2008. The MIT–Cornell collision and why it happened. *Journal of Field Robotics* 25, 10: 775–807.

25. A. Geiger, P. Lenz, and R. Urtasun. 2012. Are we ready for autonomous driving? The KITTI vision benchmark suite. In *2012 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 3354–3361. https://doi.org/10.1109/CVPR.2012.6248074

26. Lee Gomes. 2014. Google's Self-Driving Cars Still Face Many Obstacles. *MIT Technology Review*. Retrieved September 20, 2016 from https://www.technologyreview.com/s/530276/hidden-obstacles-for-googles-self-driving-cars/

27. Michael A. Goodrich and Alan C. Schultz. 2007. Human-robot interaction: a survey. *Foundations and trends in human-computer interaction* 1, 3: 203–275.

28. Google. 2016. *Google Self-Driving Car Project Monthly Report August 2016*. Retrieved from https://static.googleusercontent.com/media/www.googl e.com/en//selfdrivingcar/files/reports/report-0816.pdf

29. Pentti Haddington. 2012. Movement in action: Initiating social navigation in cars. Retrieved September 20, 2016 from http://www.degruyter.com/view/j/sem.2012.2012.issue -191/sem-2012-0059/sem-2012-0059.xml

30. Dave Harley and Geraldine Fitzpatrick. 2009. YouTube and intergenerational communication: the case of Geriatric1927. *Universal access in the information society* 8, 1: 5–20.

31. Christian Heath and Paul Luff. 2000. *Technology in Action*. Cambridge University Press

32. Juan Pablo Hourcade, Sarah L. Mascher, David Wu, and Luiza Pantoja. 2015. Look, My Baby Is Using an iPad! An Analysis of YouTube Videos of Infants and Toddlers Using Tablets. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (CHI '15), 1915–1924. https://doi.org/10.1145/2702123.2702266

33. Chien-Ming Huang and Bilge Mutlu. 2012. Robot behavior toolkit: generating effective social behaviors for robots. In *Proceedings of the seventh annual ACM/IEEE international conference on Human-Robot Interaction*, 25–32. Retrieved September 20, 2016 from http://dl.acm.org/citation.cfm?id=2157694

34. Nikki Jones and Geoffrey Raymond. 2012. "The Camera Rolls" Using Third-Party Video in Field Research. *The Annals of the American Academy of Political and Social Science* 642, 1: 109–123.

35. Oskar Juhlin. 2001. Traffic behaviour as social interaction - Implications for the design of artificial drivers. In *The Social Production of Technology : On the everyday life of things*, Hans Glimell and Oskar Juhlin (eds.). BAS, Goteburg, Sweden, 19–38.

36. Oskar Juhlin. 2010. *Social Media on the Road*. Springer

37. Jack Katz. 1999. *How Emotions Work*. Chicago University Press.

38. Eric Laurier. 2016. YouTube: fragments of a video-tropic atlas. *Area* 48, 4: 488-495

39. Eric Laurier, Barry Brown, and Hayden Lorimer. 2012. What it means to change lanes: actions, emotions and wayfinding in the family car. *Semiotica* 191, 1: 117-135.

40. Eric Laurier, Hayden Lorimer, Barry Brown, Owen Jones, Oskar Juhlin, Alison Noble, Mark Perry, David Pica, Paul Sormani, and Ignatz Strebel. 2008. Driving and "passengering": notes on the ordinary organization of car travel. *Mobilities* 3, 1: 1–23.

41. Key Jung Lee, Yeon Kyoung Joo, and Clifford Nass. 2014. Partially Intelligent Automobiles and Driving Experience at the Moment of System Transition. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '14), 3631–3634. https://doi.org/10.1145/2556288.2557370

42. Iolanda Leite, Carlos Martinho, and Ana Paiva. 2013. Social robots for long-term interaction: a survey. *International Journal of Social Robotics* 5, 2: 291–308.

43. John Leonard. 2015. Mapping, Localization, and Self-Driving Vehicles. Retrieved from http://ilp.mit.edu/images/conferences/2016/ict/presenta tions/Leonard.2016.ICT.pdf

44. S. Lowry, N. Sünderhauf, P. Newman, J. J. Leonard, D. Cox, P. Corke, and M. J. Milford. 2016. Visual Place Recognition: A Survey. *IEEE Transactions on Robotics* 32, 1: 1–19. https://doi.org/10.1109/TRO.2015.2496823

45. Markus Maurer, J. Christian Gerdes, Barbara Lenz, and Hermann Winner. 2016. Autonomous driving: technical, legal and social aspects. Retrieved September 20, 2016 from http://dl.acm.org/citation.cfm?id=2965082

46. Ross Mead and Maja Mataric. 2016. Robots Have Needs Too: How and Why People Adapt Their Proxemic Behavior to Improve Robot Social Signal Understanding. *Journal of Human-Robot Interaction* 5, 2: 48–68.

47. Alexander Meschtscherjakov, Rod McCall, Nicolas Louveton, Thomas Engel, Manfred Tscheligi, and Vincent Koenig. 2015. MaDSAV: maintaining driving skills in semi-autonomous vehicles. In *Adjunct Proceedings of the 7th International Conference on Automotive User Interfaces and Interactive Vehicular Applications*, 136–139. Retrieved September 20, 2016 from http://dl.acm.org/citation.cfm?id=2809732

48. Alexander Meschtscherjakov, Manfred Tscheligi, Dalila Szostak, Sven Krome, Bastian Pfleging, Rabindra Ratan, Ioannis Politis, Sonia Baltodano, Dave Miller, and Wendy Ju. 2016. HCI and Autonomous Vehicles: Contextual Experience Informs Design. In *Proceedings of the 2016 CHI Conference*

Extended Abstracts on Human Factors in Computing Systems (CHI EA '16), 3542–3549. https://doi.org/10.1145/2851581.2856489

49. Alexander Meschtscherjakov, Manfred Tscheligi, Dalila Szostak, Rabindra Ratan, Roderick McCall, Ioannis Politis, and Sven Krome. 2015. Experiencing Autonomous Vehicles: Crossing the Boundaries Between a Drive and a Ride. In *Proceedings of the 33rd Annual ACM Conference Extended Abstracts on Human Factors in Computing Systems* (CHI EA '15), 2413–2416. https://doi.org/10.1145/2702613.2702661

50. Lorenza Mondada. 2012. Talking and driving: multi-activity in the car. *Semiotica* 2012, 191: 223–256.

51. Bilge Mutlu, Takayuki Kanda, Jodi Forlizzi, Jessica Hodgins, and Hiroshi Ishiguro. 2012. Conversational gaze mechanisms for humanlike robots. *ACM Transactions on Interactive Intelligent Systems (TiiS)* 1, 2: 12.

52. Maurice Nevile. 2012. Interaction as distraction in driving: A body of evidence. *Semiotica*, 191: 169–196.

53. Nick Statt. 2016. Tesla customers have driven 100 million miles with Autopilot active. *The Verge*. Retrieved from http://www.theverge.com/2016/5/24/11761098/tesla-autopilot-self-driving-cars-100-million-miles

54. J. Nilsson, J. Silvlin, M. Brannstrom, E. Coelingh, and J. Fredriksson. 2016. If, When, and How to Perform Lane Change Maneuvers on Highways. *IEEE Intelligent Transportation Systems Magazine* 8, 4: 68–78. https://doi.org/10.1109/MITS.2016.2565718

55. Don Norman. 1992. *Turn Signals Are the Facial Expressions of Automobiles*. Addison Wesley.

56. Raja Parasuraman and Victor Riley. 1997. Humans and automation: Use, misuse, disuse, abuse. *Human Factors: The Journal of the Human Factors and Ergonomics Society* 39, 2: 230–253.

57. Frank Pasquale. 2015. *The black box society: The secret algorithms that control money and information*. Harvard University Press.

58. E. Pollard, P. Morignot, and F. Nashashibi. 2013. An ontology-based model to determine the automation level of an automated vehicle for co-driving. In *2013 16th International Conference on Information Fusion (FUSION)*, 596–603.

59. Gary Pritchard, John Vines, Pam Briggs, Lisa Thomas, and Patrick Olivier. 2014. Digitally driven: how location based services impact the work practices of London bus drivers. *In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '14)*. ACM, New York, NY, USA, 3617-3626. http://dx.doi.org/10.1145/2556288.2557156.

60. Edward Reynolds. 2011. Enticing a challengeable in arguments: sequence, epistemics and preference organisation. *Pragmatics* 21, 3: 411–430.

61. Daniel Rosenband. 2016. Are we there yet? Silicon in self-driving cars. Retrieved from http://www.hotchips.org/video-for-keynote-2-are-we-there-yet-silicon-in-self-driving-cars-by-daniel-rosenband-of-google-available-to-all/

62. Selma Sabanovic, Marek P. Michalowski, and Reid Simmons. 2006. Robots in the wild: Observing human-robot social interaction outside the lab. In *9th IEEE International Workshop on Advanced Motion Control, 2006.*, 596–601. Retrieved September 20, 2016 from http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=1631758

63. Dorsa Sadigh, Shankar Sastry, Sanjit A. Seshia, and Anca D. Dragan. 2016. Planning for autonomous cars that leverages effects on human actions. In *Proceedings of the Robotics: Science and Systems Conference (RSS)*. Retrieved December 21, 2016 from https://pdfs.semanticscholar.org/e3e1/a23e0096f8ca51b783812310d44673a3c89a.pdf

64. Emanuel A. Schegloff. 2007. *Sequence organization in interaction: Volume 1: A primer in conversation analysis*. Cambridge University Press.

65. Mimi Sheller. 2007. Bodies, cybercars and the mundane incorporation of automated mobilities. *Social & Cultural Geography* 8, 2: 175–197.

66. Don Sherman. 2016. Semi-Autonomous Cars Compared! Tesla Model S vs. BMW 750i, Infiniti Q50S, and Mercedes-Benz S65 AMG - Feature. *Car and driver*. Retrieved September 20, 2016 from http://www.caranddriver.com/features/semi-autonomous-cars-compared-tesla-vs-bmw-mercedes-and-infiniti-feature

67. Neville A. Stanton and Philip Marsden. 1996. From fly-by-wire to drive-by-wire: Safety implications of automation in vehicles. *Safety Science* 24, 1: 35–49. https://doi.org/10.1016/S0925-7535(96)00067-7

68. Lucille Alice Suchman. 2007. *Human-machine reconfigurations: plans and situated actions*. Cambridge University Press.

69. Lucy Suchman. 1987. *Plans and situated actions: The problem of human-machine communication*. Cambridge University Press.

70. Marcus Sanchez Svensson, Christian Heath, and Paul Luff. 2007. Instrumental action: the timely exchange of implements during surgical operations. In *ECSCW 2007*. Springer, 41–60. Retrieved September 20, 2016 from http://link.springer.com/10.1007/978-1-84800-031-5_3

71. C. Thorpe, M. Herbert, T. Kanade, and S. Shafer. 1991. Toward autonomous driving: the CMU Navlab. I. Perception. *IEEE Expert* 6, 4: 31–42. https://doi.org/10.1109/64.85919

72. Sebastian Thrun, Michael Montemerlo, and Andrei Aron. Probabilistic Terrain Analysis For High-Speed Desert Driving. Retrieved September 20, 2016 from http://robots.stanford.edu/papers/thrun.mapping-Stanley.pdf

73. Rod Watson. 1999. Driving in forests and mountains: A pure and applied ethnography. *Ethnographic studies* 3: 50–60. https://doi.org/10.5449/idslu-001103654

74. Junqing Wei, John M. Dolan, and Bakhtiar Litkouhi. 2013. Autonomous vehicle social behavior for highway entrance ramp management. In *Intelligent Vehicles Symposium (IV), 2013 IEEE*, 201–207. Retrieved December 21, 2016 from http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=6629471

75. Junko Yoshida. 2016. Autopilot: Flying vs. Driving. Can lessons learned from aviation apply to cars? *EE Times*. Retrieved December 10, 2016 from http://www.eetimes.com/document.asp?doc_id=1329271