UNIVERSITY OF LEEDS

This is a repository copy of *Models of Human Decision-Making as Tools for Estimating and Optimizing Impacts of Vehicle Automation*.

White Rose Research Online URL for this paper:
http://eprints.whiterose.ac.uk/127539/

Version: Accepted Version

Article:
Markkula, G orcid.org/0000-0003-0244-1582, Romano, R orcid.org/0000-0002-2132-4077, Madigan, R orcid.org/0000-0002-9737-8012 et al. (3 more authors) (2018) Models of Human Decision-Making as Tools for Estimating and Optimizing Impacts of Vehicle Automation. Transportation Research Record, 2672 (37). pp. 153-163. ISSN 0361-1981

https://doi.org/10.1177/0361198118792131

(c) National Academy of Sciences: Transportation Research Board 2018. Markkula, G , Romano, R , Madigan, R et al. (3 more authors) (2018) Models of Human Decision-Making as Tools for Estimating and Optimizing Impacts of Vehicle Automation. Transportation Research Record, 2672 (37). pp. 153-163. ISSN 0361-1981. Reprinted by permission of SAGE Publications

Reuse
See Attached

Takedown
If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

eprints@whiterose.ac.uk
https://eprints.whiterose.ac.uk/

MODELS OF HUMAN DECISION-MAKING AS TOOLS FOR ESTIMATING AND OPTIMISING IMPACTS OF VEHICLE AUTOMATION

Gustav Markkula
(corresponding author)
Institute for Transport Studies, University of Leeds
LS2 9JT Leeds, United Kingdom
Tel: +44 113 3439832; Email: g.markkula@leeds.ac.uk

Richard Romano
Institute for Transport Studies, University of Leeds
LS2 9JT Leeds, United Kingdom
Tel: +44 113 3438466; Email: r.romano@leeds.ac.uk

Ruth Madigan
Institute for Transport Studies, University of Leeds
LS2 9JT Leeds, United Kingdom
Tel: +44 113 3432071; Email: r.madigan@leeds.ac.uk

Charles W. Fox
Institute for Transport Studies, University of Leeds
LS2 9JT Leeds, United Kingdom
Tel: +44 113 3431794; Email: c.w.fox@leeds.ac.uk

Oscar T. Giles
Institute for Transport Studies, University of Leeds
LS2 9JT Leeds, United Kingdom
Email: o.t.giles@leeds.ac.uk

Natasha Merat
Institute for Transport Studies, University of Leeds
LS2 9JT Leeds, United Kingdom
Tel: +44 113 3436614; Email: n.merat@leeds.ac.uk

Word count: 5,736 words text + 7 figures x 250 words (each) = 7,486 words

Submission date: 1/08/2017
Revision date: 15/11/2017

ABSTRACT

With the development of increasingly automated vehicles (AVs) comes the increasingly difficult challenge of comprehensively validating these for acceptable, and ideally beneficial, impacts on the transport system. There is a growing consensus that virtual testing, where simulated AVs are deployed in simulated traffic, will be key for cost-effective testing and optimisation. The least mature model components in such simulations are those generating the behaviour of human agents in or around the AVs. In this paper, human models and virtual testing applications are presented for two example scenarios: (i) a human pedestrian deciding whether to cross a street in front of an approaching automated vehicle, with or without external human-machine interface elements, and (ii) an AV handing over control to a human driver in a critical rear-end situation. These scenarios have received much recent research attention, yet simulation-ready human behaviour models are lacking. They are discussed here in the context of existing models of perceptual decision-making, situational awareness, and traffic interactions. It is argued that the human behaviour in question might be usefully conceptualised as a number of interrelated decision processes, not all of which are necessarily directly associated with externally observable behaviour. The results show that models based on this type of framework can reproduce qualitative patterns of behaviour reported in the literature for the two addressed scenarios, and it is demonstrated how computer simulations based on the models, once these have been properly validated, could allow prediction and optimisation of AV impacts on traffic flow and traffic safety.

Keywords: automated vehicles, virtual testing, human behavior models

INTRODUCTION

In recent years, rapid technological progress has been made on automated vehicles (AVs); vehicles capable of taking over increasing shares of the driving task, with large hoped-for benefits in terms of increased mobility, improved traffic safety, reduced environmental footprint, and economic growth (*1*). As with any technology, it is important to subject AVs to sufficient testing and validation to ensure safe and effective operation. This is a daunting task, however, considering the essentially infinite variations of possible traffic situations that an AV needs to handle (*2*).

It is increasingly assumed that one important approach in the toolkit for AV validation will be *virtual testing* methods, where a computer simulation of an AV, integrating mathematical models of vehicle and sensors with the actual self-driving algorithms, can encounter a rich variety of simulated traffic situations (*3*, *4*). These traffic situations can for some purposes be replayed from large static databases of actual real-traffic recordings, but if the objectives of a virtual test depends on some form of *interaction* between the AV and surrounding road users—e.g., by means of external human-machine interfaces (*5*)—or interaction between the AV and its onboard operator, the involved human agents also need to be simulated.

Human road user models do exist, and are currently being actively developed and applied to AV testing (*3*, *6*), but so far mainly on the relatively coarse-grained level of *traffic microsimulation*, where the trajectory of each road user (car, truck, pedestrian, etc) is directly generated from equations of motion taking into account positions and velocities of surrounding road users. While this level of granularity might be adequate for some purposes, it seems likely that in many cases, the outcome of interactions between humans and AVs will hinge on finer details and time dynamics of human perception, scene interpretation, and decision-making (*7*, *8*).

Capturing such phenomena accurately in virtual simulation is of course very challenging, but recent literature suggest some possibly fruitful directions: There is by now ample neuroscientific support for the idea that perceptual decision-making in typical laboratory tasks is underpinned by noisy neural *evidence accumulation,* to a decision threshold at which the overt response is initiated (*9*). Interestingly, such models have also been proven useful for explaining human driver reaction times in responding to stimuli in traffic; discrete stimuli such as brake lights (*10*) but importantly also graded, dynamic stimuli like the *visual looming* of approaching road users or collision threats (*11*, *12*).

Here, this line of modelling will be taken one step further, and applied to the more complex decision-making situations that tend to arise in the AV context, where the emphasis of modelling naturally shifts away from low-level decisions like “more or less braking?” towards a higher-level assessment of the situation at hand; what is often referred to as *situational awareness* (*8*, *13*). The basic modelling idea to be explored here is that situational awareness can be modelled as a number of interrelated perceptual decisions about the world, where each decision is modelled as an accumulation process, but where these processes are also interconnected to influence each other in excitatory or inhibitory fashion. This is reminiscent of connectionist and activation dynamics type models of cognition (*14*, *15*). However, to our knowledge, the present paper is the first time several decision-making processes of an evidence accumulator type have been connected in a network arrangement as a means of modelling complex, multi-stage decision processes.

After a first section below describing the two AV-human interaction scenarios addressed here, three consecutive sections will pursue the three main objectives of the paper: (i) To suggest models of human decision-making in the targeted scenarios, as examples of the modeling approach sketched above. (ii) To illustrate that these models can qualitatively reproduce non-trivial behavior patterns reported in the literature. These prior findings thus serve as a first high level data set, to which model parameters are manually fitted. (iii) To illustrate how models at this level of detail

can be useful in virtual testing of AVs. It is not an aim here to suggest final model formulations, nor to perform full model validation and tuning; these will be matters for future work.

TARGETED AV-HUMAN INTERACTION SCENARIOS

Figure 1 shows the two traffic scenarios considered here. The pedestrian interaction scenario in panel (a) is based on the test track study by Schneemann and Gohl (16). They considered the question of how an AV should behave at a zebra crossing, by studying the interactive behaviour of human drivers and pedestrians, as a function of initial car speed v_0 and time to collision (TTC; or alternatively in this scenario time to arrival) at the moment when the pedestrian was instructed to step up to the crossing, prompting decelerations from the (uninstructed) drivers. Here, what will be modelled is the pedestrian's decision to cross or yield as a function of car movements, as well as of more direct indications from the car (from either the driver or the AV, as the case may be) about having seen the pedestrian, and about its possible intentions to yield.

The safety-critical take-over scenario in panel (b) is based on the simulator study by Gold et al. (*17*), where drivers either manually drove a vehicle or were driven by an AV at v_0 = 120 km/h (75 mph), when a lead vehicle changed lane to reveal a stationary vehicle, at a TTC of either 5 s or 7 s, in the AV case prompting a *take-over request* to the driver. Similar types of take-over scenarios have recently been subject to intense research efforts; e.g. (*7*, *18*). What will be modelled here is the human's decisions, after the obstacle appears, on where to look and how to manoeuvre the car.

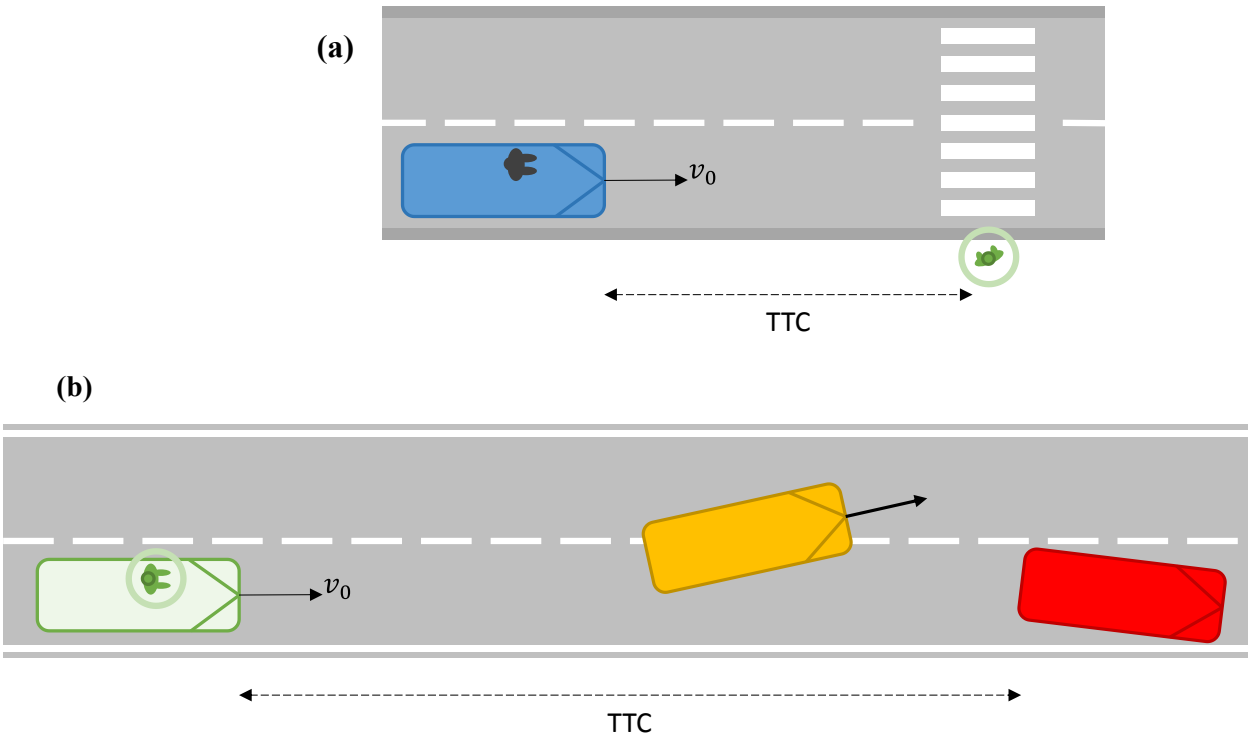

FIGURE 1 Schematic illustrations of the scenarios considered in this paper. The modeled human is indicated with a circle. TTC is short for time to collision. (a) Pedestrian interaction scenario. (b) Safety-critical take-over scenario.

MODELS

Model of pedestrian crossing decisions

The basic hypothesis behind the pedestrian model, illustrated in Figure 2, is that the pedestrian makes the action decision to cross based on either or both of two perceptual decisions: (i) A perceptual decision that it is possible to pass the road before the car arrives, made by observing the visual looming quantity $\tau = \theta/\dot{\theta}$, with θ the projected angle of the car on the pedestrian's retina. τ is a close approximation of TTC if assuming constant speeds (19), and is assumed to be compared to a threshold τ_{pass}. (ii) A perceptual decision that the car (driver or AV, but here "car" will be used for short) intends to yield, but not if the actual car movements clearly suggest that it is still unsafe to cross ($\tau < \tau_{pass}$). To make the decision about whether or not the car intends to yield, the pedestrian is assumed to monitor the quantity $\dot{\tau} = \mathrm{d}\tau/\mathrm{d}t$, which is ≥ -0.5 if the car stops at or before the zebra crossing (19), but also explicit communication acts (e.g., headlight flashes). However, if the pedestrian decides that the car has not seen him/her (e.g., based on driver head orientation), this is taken as evidence that the car is not intending to yield.

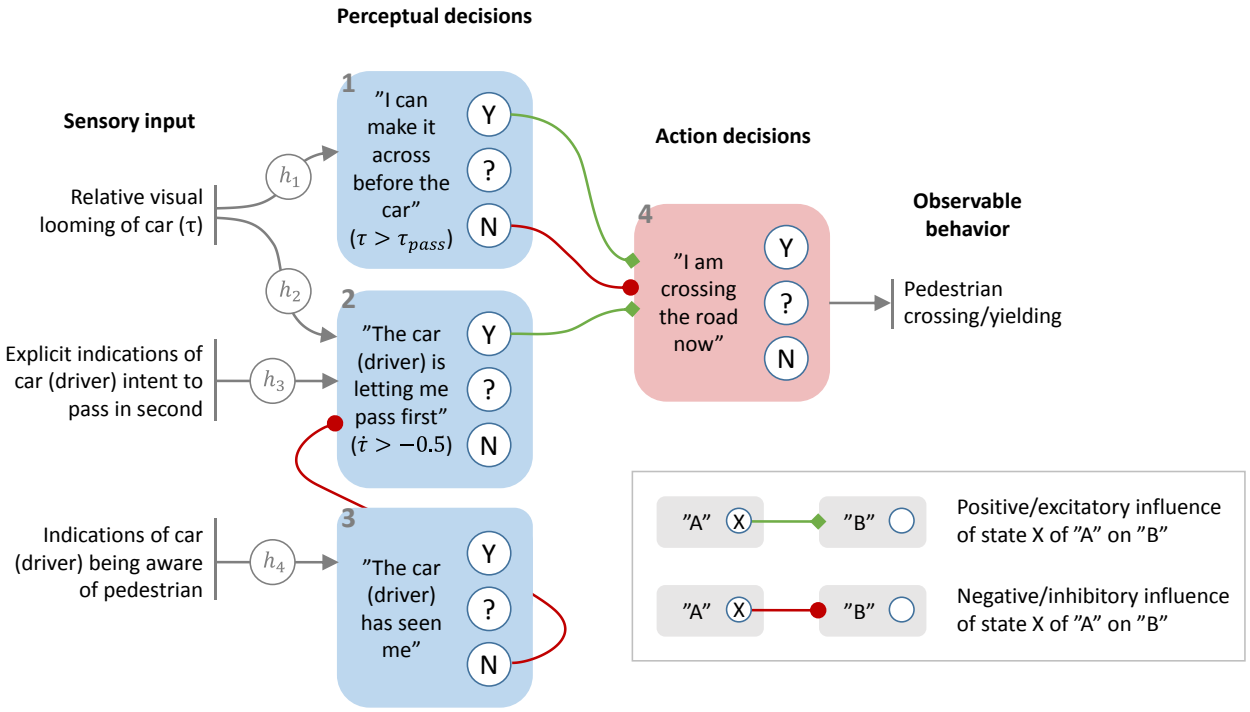

FIGURE 2 Schematic representation of the model for the pedestrian interaction scenario.

The remainder of this section provides a full mathematical specification of the model, but is not crucial for understanding the rest of this paper: Figure 2 shows four interconnected *decision units*, and at any point in time, the state of the model is described by the vector of unit *activations*, $\boldsymbol{A}(t) = [A_1(t) \ldots A_4(t)]^T$, where each A_i is a two-sided accumulator, bounded at plus and minus one, implying "yes" and "no" decisions, respectively, and zero implying maximum uncertainty. This type of accumulator is often referred to as a "drift-diffusion model for two-choice decisions" or similar (*20*), although interconnecting several of these or simulating them beyond a threshold being reached is possibly novel. At each new simulation time step, the A_i are updated as follows:

$$\frac{\mathrm{d}}{\mathrm{d}t}\boldsymbol{A}(t) = -\frac{1}{T}\boldsymbol{A}(t) + \boldsymbol{W}_Y\sigma_Y\big(\boldsymbol{A}(t)\big) + \boldsymbol{W}_?\sigma_?\big(\boldsymbol{A}(t)\big) + \boldsymbol{W}_N\sigma_N\big(\boldsymbol{A}(t)\big) + \boldsymbol{F}(t) + \boldsymbol{\nu}(t) \quad (1)$$

where T is a decay time constant, the $\boldsymbol{W}_\bullet$ are matrices of connections between decision units, the $\sigma_\bullet$ are activation functions, $\boldsymbol{F}(t)$ is a vector of external inputs, and $\boldsymbol{\nu}(t)$ is a vector of noise, accounting for between-trial behavioral variability due to for example variations in brain activity, modelled here as Gaussian noise with standard deviation $\Sigma\sqrt{\Delta t}$, where Σ is a parameter and Δt is the simulation time step (9). In this first model implementation, to simplify the manual tuning process, all elements in $\boldsymbol{W}_\bullet$ were kept either at zero (no connection) or ±1 (excitatory and inhibitory connections), rather than being freely tuned. Also for simplicity, the activation functions were kept piecewise linear: $\sigma_Y(A) = 0$ and 1 for $A \leq 0$ and $A \geq 1$, respectively, and linearly increasing in between, and $\sigma_N(A) = \sigma_Y(-A)$, such that the two functions respond maximally when A indicates "yes" and "no", respectively. Further, $\sigma_?(A) = 1 - \sigma_Y(A) - \sigma_N(A)$, i.e., zero at $A = \pm 1$ and linearly increasing from both sides to a maximum of one at $A = 0$. The external inputs to the decision units were:

- $F_1(t) = h_1\big(\tau(t)\big) = k_1\big(\tau(t) - \tau_{pass}\big)$, corresponding to the perceptual decision about whether the time margin for crossing is currently above $\tau_{pass} = 3$ s.
- $F_2(t) = h_2(\tau(t)) + h_3(t)$, where $h_2(\tau(t)) = k_2\big(\dot{\tau}(t) - \dot{\tau}_{stop}\big)$, corresponding to the perceptual decision about whether the car is decelerating to stop before the zebra crossing. Here $\dot{\tau}_{stop} = -0.55$ was used, to accumulate positive evidence also in the limit case $\dot{\tau} = -0.5$. The input $h_3(t)$ was set to one to model the case where the car is providing some communicative act interpreted by the pedestrian as indicating an intention to yield, e.g., flashing the headlights, zero otherwise.
- $F_3(t) = h_4(t)$ was set to one to model the car or its driver doing something that the pedestrian takes as evidence of the car seeing the pedestrian, e.g., eye contact or related head pose in the case of a human driver, or some external HMI in the case of an AV, or minus one if the car is doing something that the pedestrian takes as evidence of the opposite, e.g., a car driver is looking somewhere else.
- $F_4(t) = 0$.

Manual tuning indicated satisfactory model behavior for $T = 2$ s; $\Sigma = 0.5$; $k_1 = 0.5$; $k_2 = 4$.

Model of driver behavior during safety-critical take-over

The model for the take-over scenario, illustrated in Figure 3, is slightly more complex. It does not terminate at the first action decision like the pedestrian model; instead each new action affects how the scenario continues playing out over time. For reasons of space, all connections shown in Figure 3 will not be discussed individually here, but the main hypotheses behind the model are as follows: Driver braking is assumed to function as proposed in (*12*), as discrete increases of brake pedal position applied each time a discrepancy between visual looming, quantified as τ^{-1}, and predicted visual looming τ_p^{-1}, updated after each new brake adjustment, accumulates to threshold. The decision to change lane can come about in the model in two ways: an early detection of catching up with the lead vehicle (which does not in itself prompt braking) or

a later realisation that attempted braking is not solving the conflict, and both of these are also modulated by visual looming directly. However, the lane changing decision is in all cases inhibited—as in made less probable—if the driver does not know that it is safe to change lane. Gaze sampling is assumed driven to a large extent by the need to resolve uncertainties in the perceptual decisions (*21*), and the current gaze target determines what sensory input the model receives (the "switches" in Figure 3). In this first version of the model, there is no direct account of the process of actual control take-over; as a first approximation it will be investigated to what extent the effects of taking over from the AV can be accounted for as a simple time delay, combined with lower initial awareness of the adjacent lane.

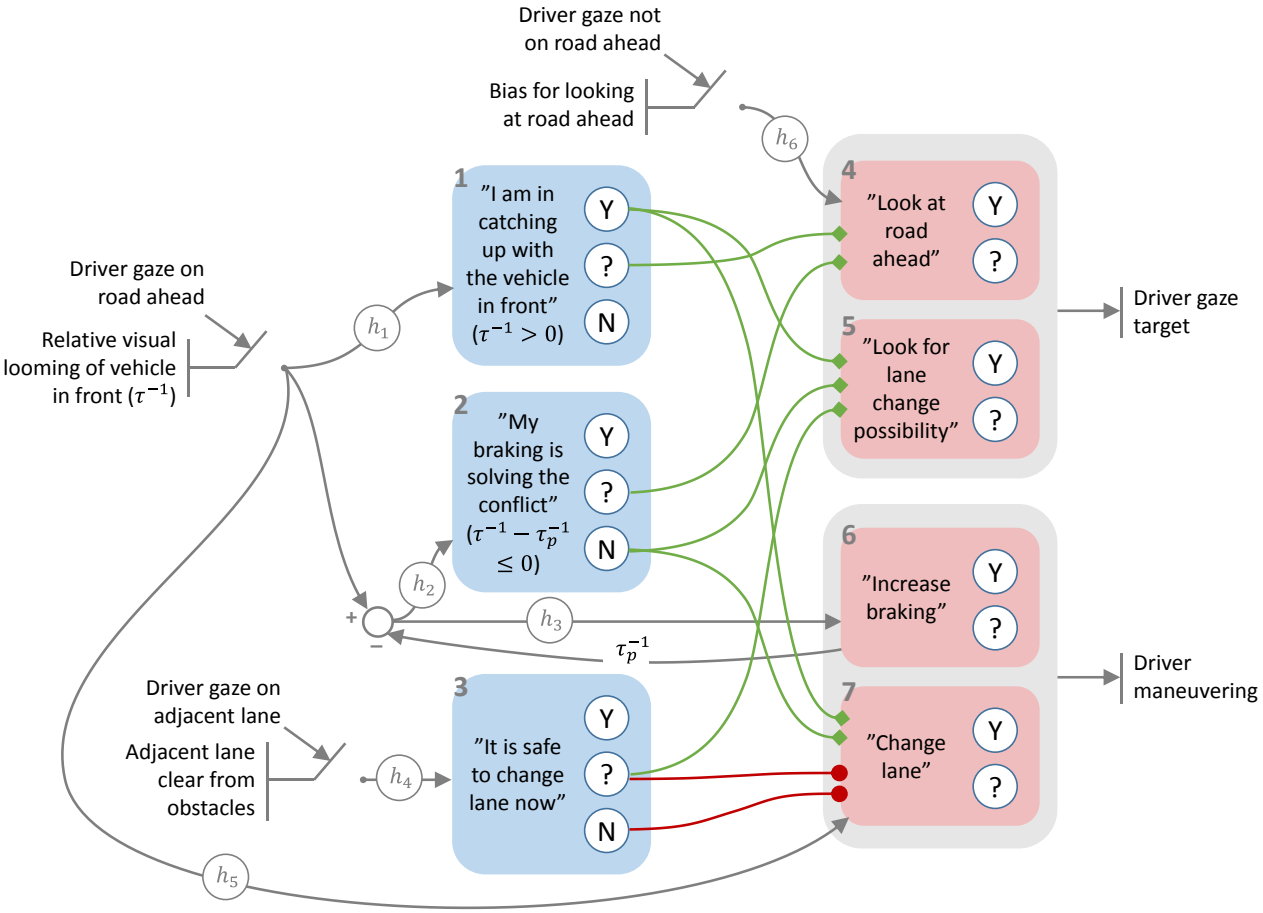

FIGURE 3 Schematic representation of the model for the safety-critical take-over scenario.

Again, the rest of this section gives mathematical model details: Just as in the pedestrian model, Equation (1) governs activations in the network. However, decisions between competing actions about where to look (units 4 and 5) or how to maneuver (units 6 and 7), are here modelled as "races" (*20*), bounded downward by zero instead of minus one, and each time an action threshold $A_i = 1$ is reached, the race in question is restarted by resetting the competing accumulators to zero. The exception is the "increase braking" action, which has its accumulator reset to 0.7, to reflect an increased brake readiness after the first brake application (*12*). Also, before the first brake application, decision unit 2 is inactive. The looming-related inputs to the network, all set to zero if the driver's eyes are off the road, were defined as:

- $F_1(t) = k_1\tau^{-1}(t)$, modeling the perceptual decision of whether the vehicle ahead is coming closer.
- $F_2(t) = -k_2(\tau^{-1}(t) - \tau_p^{-1}(t))$, modeling the perceptual decision of whether braking is solving the conflict.
- $F_6(t) = k_3(\tau^{-1}(t) - \tau_p^{-1}(t) - \tau_B^{-1})$, modeling the decision to increase braking (12).
- $F_7(t) = k_4(\tau^{-1}(t) - \tau_S^{-1}) - \min(0, k_5(\tau^{-1}(t) - \tau_{NS}^{-1}))$, modeling a positive looming contribution to the decision to change lane and a negative contribution to the same decision coming into play once $\tau^{-1} > \tau_{NS}^{-1}$, where $\tau_{NS}^{-1} > \tau_S^{-1}$.

The remaining inputs were defined as:

- $F_3(t) = k_6$ and $F_4(t) = k_7$, both if the driver is directing gaze towards the adjacent lane (whether by mirror checks, shoulder checks, etc, is not defined in the model), zero otherwise.
- $F_5(t) = 0$.

Manual tuning indicated satisfactory model behavior for $T = 2$ s; $\Sigma = 0.3$; $k_1 = k_2 = 20$; $k_3 = k_5 = 6$; $k_4 = 3$; $k_6 = 1.5$; $k_7 = 1.2$; $\tau_B^{-1} = \tau_S^{-1} = 0.1$ s^{-1}; $\tau_{NS}^{-1} = 0.5$ s^{-1}.

MODEL BEHAVIOR – REPRODUCING FINDINGS FROM LITERATURE

Reproducing pedestrian crossing behavior

Figure 4 shows example simulations with the pedestrian model, all at $v_0 = 50$ km/h (31 mph) and with TTC = 4 s at the moment when the pedestrian steps up to the zebra crossing, at which point the simulated car also initiates a deceleration to stop exactly at the zebra crossing. Panel (a) shows example simulations both with and without noise, whereas panels (b) and (c) show distributions for 10,000 simulations with noise.

In panel (a), note in the simulation without noise how the pedestrian comes close to a quick decision to cross while τ is above the threshold $\tau_{pass} = 3$ s, but since the margin $\tau - \tau_{pass}$ is small, the evidence accumulation is so slow that τ falls below τ_{pass} before a complete crossing decision has formed. From that point, despite picking up some signs of car deceleration, the pedestrian still judges from τ that crossing might be risky. Thus, as further illustrated by the delay time distribution in panel (b), one qualitative human behavior result that the model reproduces is the finding reported in (*16*), where 20 % of pedestrians waited until the car came to a complete stop before crossing. In the model simulations, the remaining 80 % are spread out in time with a wide peak in the time region while τ is still above τ_{pass}, where the model can make the decision to cross even if it is unsure of the car's intentions. This aligns well with another finding in (*16*), where 25 % of pedestrians stated that they made the decision to cross without having seen the car's deceleration. Exactly what activation A of the deceleration-detecting perceptual decision would yield a positive response to such a question is not defined in the model, but as a rough comparison, 25 % of the light brown dots in panel (b) lie below $A = 0.3$.

The scenario in Figure 4(b) assumes that the pedestrian gets indications of being seen by the driver (e.g., eye contact) but no explicit communication about yielding intentions, whereas panel (c) shows variations in this interactive behaviour. In the best case, where the pedestrian gets indications both of the car being aware of him/her, and of the car's intention to stop, the model pedestrian almost never waits until the car is stationary (dotted black line), whereas more than half of the model pedestrians wait that long in the worst case where neither of those indications are present (solid gray line).

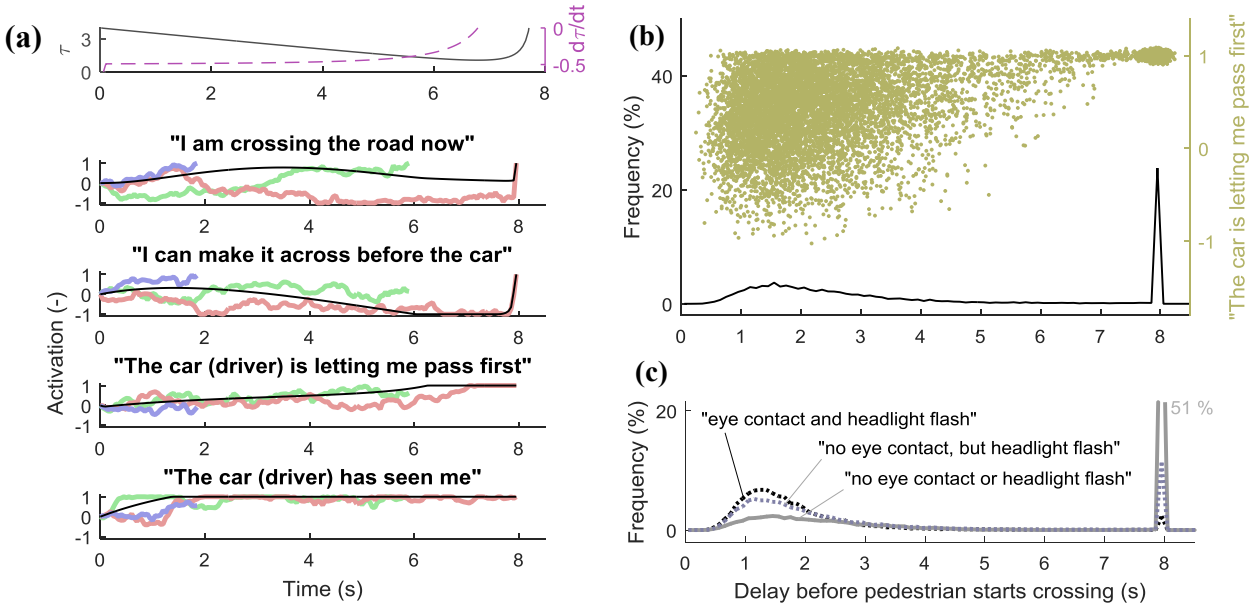

FIGURE 4 Example simulations of the pedestrian model. (a) Visual looming input (top panel) and resulting model activations. The black activation traces show a model simulation without noise, and the lighter colored activation traces show three example simulations with noise. (b) The black line shows the probability distribution of delay times for the model pedestrian's decision to cross the road. The light brown dots show the activation of the "the car is letting me pass first" unit at the time of crossing decision, as a function of the decision delay time. (c) Distribution of delay times for variations of the same scenario.

Reproducing safety-critical take-over behavior

Figure 5(a) shows four example time histories generated by the take-over driver model, all with TTC = 4 s and intermediate ($A = 0.5$) initial awareness of the adjacent lane being empty. These simulations differ only in the initial seed of the random noise in the decision processes, but it is notable that the model can nevertheless produce a range of qualitatively different sequences of behavior, both with and without one or more long or short glances away from the road ahead, and both with and without some degree of braking before a lane change is initiated (terminating the simulation).

Figure 5(b) shows the model-predicted frequency of different avoidance maneuvers, as a function of TTC at obstacle appearance, for a case with intermediate ($A = 0.5$) initial awareness of the adjacent lane being empty, as a possible emulation of a typical manual driving case. The obtained model distributions, and specifically the frequencies marked by vertical lines in the figure, align well with the pattern of driver responses reported in (*17*) for the manual driving condition: All human drivers (8/8; 100 %) responded with steering in the TTC = 7 s condition, whereas in the TTC = 5 s condition, also some combined braking and steering was observed (1/5 drivers; 20 %).

Figure 5(c) shows that without any initial awareness of the adjacent lane ($A = 0$), the probability of braking increases. This was also reported in (*17*) for the AV condition. Specifically, the model frequencies at the initial TTCs marked with vertical lines (4 s and 3.5 s) align qualitatively with the human responses in (*17*) for the experimental conditions with take-over

request at 7 s and 5 s TTC: A near-50/50 division between steering only and combined steering/braking was observed in the 7 s condition, whereas in the 5 s condition combined steering/braking rose to more clearly being the most common maneuver, followed by steering only, and some observations of pure braking. In other words, overall, the model simulations suggest the possibility that the take-over process in this scenario can be accounted for as a combination of reduced situational awareness and an urgency-dependent time delay, longer for the less critical scenario ($7\ \text{s} - 4\ \text{s} = 3\ \text{s}$, compared to $5\ \text{s} - 3.5\ \text{s} = 1.5\ \text{s}$) . This aligns with the report in (*17*) that hands-on-wheel and eyes-on-road delay times were longer in the less critical 7 s take-over scenario.

Figure 5(b) and (c) also show simulated crash frequencies. The model's braking alone was typically not enough to avoid collision, whereas in (*17*) seemingly no collisions occurred. This may indicate some problem with the tuning of the braking model. If the model applied steering avoidance, avoidance was considered successful as long as the time margin before collision, given speed and acceleration at time of steering initiation, was above 1 s.

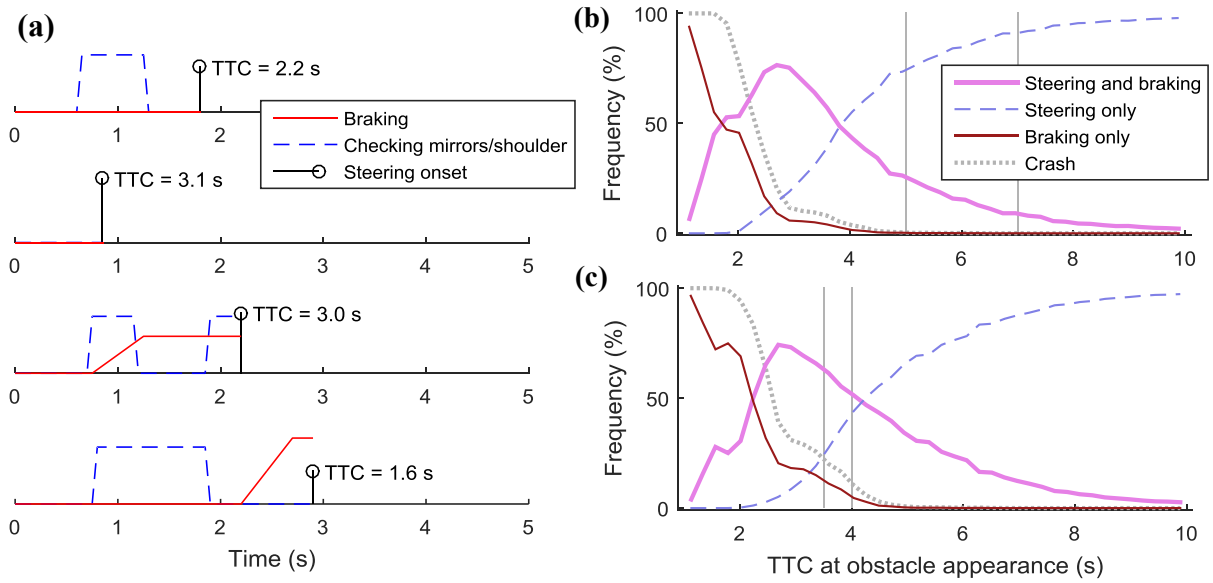

FIGURE 5 Example simulations of the hand-over model. (a) Typical model behavior time histories, all for the same initial condition. (b) Frequency of avoidance maneuvers performed by the model when run in a way so as to emulate manual driving (intermediate initial awareness of the adjacent lane), and resulting crash frequency, as a function of TTC at the time of obstacle appearance. The vertical lines indicate points of correspondence with (*17*), discussed in the text. (c) As panel b, but emulating a transition out of automated driving (with zero initial awareness of the adjacent lane).

MODEL APPLICATIONS IN VIRTUAL TESTING AND TUNING

This section will give examples of how the proposed models could be put to applied use in virtual, testing of AVs. It should be emphasized that since the models have yet to be fully validated and parameterized, the actual results given below are preliminary at best; especially for the transition scenario where the model is both more complex and showed signs of possible limitations in its braking behavior. The main aim here is to illustrate the *type* of results that can be obtained.

Optimizing AV traffic flow at a pedestrian crossing

If an AV intends to yield to a pedestrian, how should it behave so as to help the pedestrian make the crossing decision as quickly as possible? To answer this question, it was assumed that once the pedestrian starts walking, the AV adapts acceleration to pass behind the pedestrian, and then accelerates at 1 m/s^2 back up to the initial 50 km/h (31 mph). The time lost in the interaction was then calculated by comparing the distance travelled before regaining the initial speed, to the distance that would have been travelled in the same amount of time at constant speed, i.e., as if there had been no pedestrian.

Figure 6 shows, across a range of pedestrian appearance TTCs, (i) the model's predicted effect on time loss, of applying more than the deceleration needed to stop exactly at the zebra crossing, i.e., rising above the curved edge of the gray shaded area in each panel, and (ii) the effect on time loss of external AV-human communication. Panel (a) shows a "minimal communication" case where the AV does not give any explicit external signs of being aware of the pedestrian (i.e., nothing to replace eye contact) nor any signs, besides deceleration, of its intentions to yield, and panel (b) is a case where the AV does show such explicit external signs.

Perhaps counter to what one might expect, the model predicts that increased deceleration can save considerable time. This is especially true for the AV without external communication means, where for some TTCs (e.g. TTC = 4.5 s) an extra 0.5 m/s^2 of deceleration can change a 12 s time loss to near zero. The reason this happens in the model is that the pedestrian's judgment of whether the car will stop before the crossing (i.e., of whether $\dot{\tau} > -0.5$) is difficult when the margin is narrow (i.e., when $\dot{\tau} \approx -0.5$), and the exaggerated deceleration makes this perceptual decision much easier and quicker.

The model predicts that time savings can also be achieved by letting the AV communicate its situational awareness and intentions (assuming that this can be achieved by some external HMI that is understandable to the pedestrian). The benefits of exaggerated decelerations are still present, but less pronounced, except for the lowest TTCs, where it starts being questionable whether the AV should at all yield to a waiting pedestrian, considering the comfort and safety of AV occupants and any following vehicles.

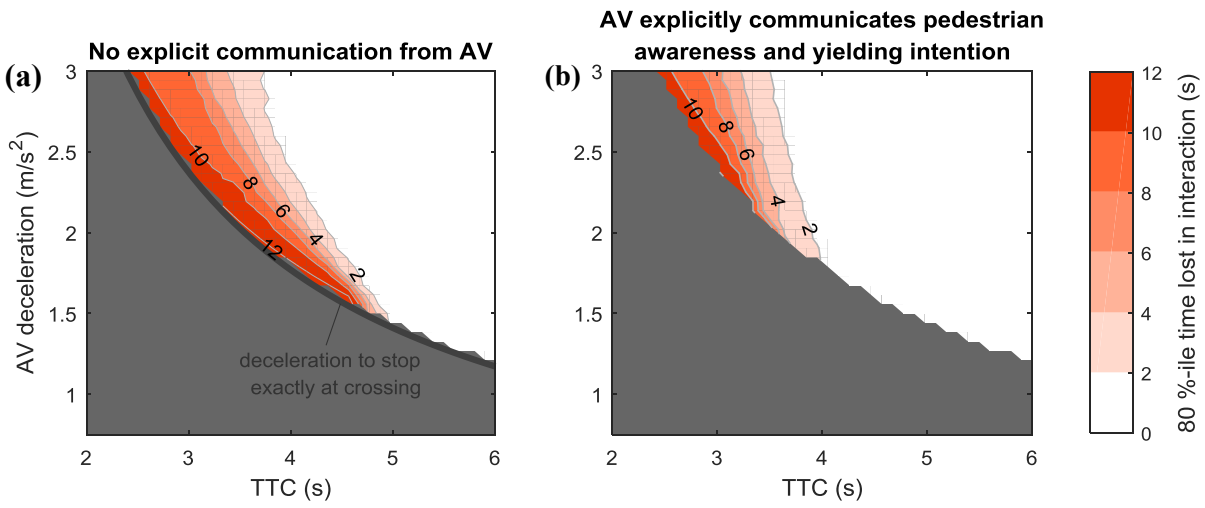

FIGURE 6 Model estimation of time lost for the AV in the pedestrian interaction, as a function of TTC at time of pedestrian presentation, and AV deceleration applied at the same moment.

Minimising crash risk in time-critical AV hand-overs

It has been suggested that detriments in performance during AV-human hand-overs are the result of a loss of situational awareness (*8*), with a recommendation that designers need to give deep consideration as to how best to facilitiate the process of bringing drivers back "into the loop" during transitions from automation (*18*). Exactly how to do so is of course a non-trivial matter, but here it is examined what benefit it might have in the studied scenario *if* a solution could be devised ensuring that the driver quickly knows, without having to look away from the road ahead, that it is possible to change lane, emulated here by initializing the "it is safe to change lane now" decision at full certainty. It is also examined here whether any such benefit might interact, according to the model, with emergency braking applied by the AV itself.

As shown in Figure 7, the model predicts a range of initial TTCs between 1.5 s and 3.5 s (remember that factoring in the take-over process, this might correspond to a considerably higher obstacle appearance TTC) where benefits of increased situational awareness may exist. For example, even if the AV decelerates with 6 m/s^2 in a TTC = 2 s scenario, immediate driver awareness of the feasibility of steering avoidance changes the model-predicted crash risk from about 80 % to about 20 %. In fact, according to the model, this low crash risk remains nearly unchanged even if the AV does not apply any emergency braking at all, which seems beneficial from a perspective of minimising risk of being struck from behind.

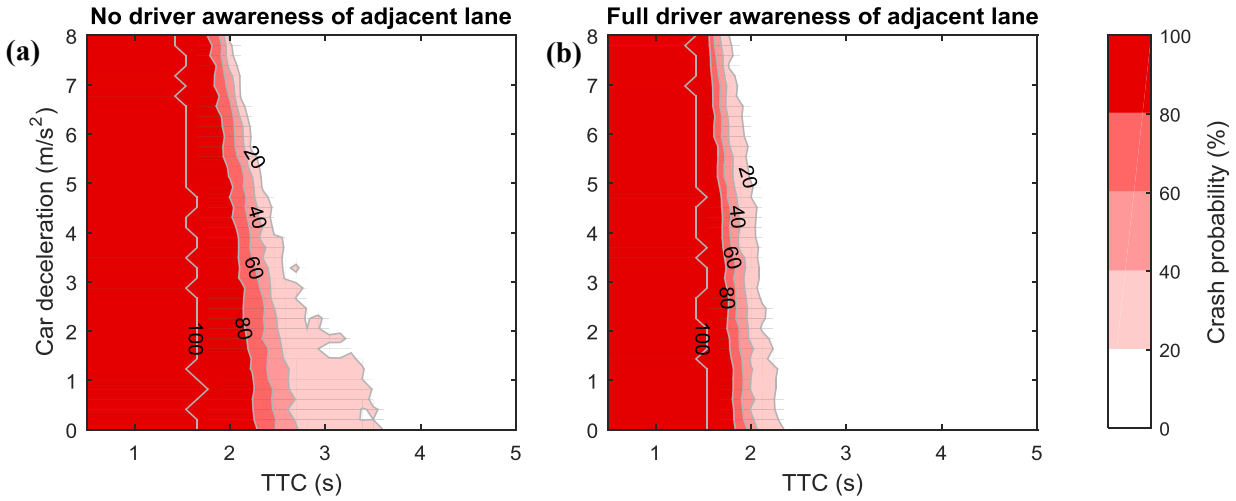

FIGURE 7 Model estimated effect of driver awareness of possibility to change lane, and AV emergency braking, on crash risk in the studied take-over scenario.

DISCUSSION AND CONCLUSIONS

This paper has presented first draft versions of novel models of human decision-making and behavior in two scenarios relevant to development and testing of AVs. Both models seem to offer useful developments beyond what is currently available in the literature for these scenarios.

Pedestrian road-crossing has been modeled here as an interrelated set of perceptual decisions about oncoming vehicle movement and intentions. The model has illustrated how a car that applies exactly the right amount of deceleration for stopping might make it difficult for a waiting pedestrian to decide on crossing before the car has come to a complete stop, thus explaining empirical observations of such late pedestrian crossing decisions (in (*16*), but anecdotally also in many people's everyday experience). In contrast, existing microsimulation pedestrian models tend to emphasise the question of *whether* a pedestrian accepts a certain time gap between passing cars (*22*), and less on *when* the pedestrian accepts it, as a function of exact movement and communicative acts of a vehicle, making the model proposed here a potentially useful addition to the toolkit for virtual testing and tuning of AV behavior.

With regards to the other scenario targeted here, focusing on human avoidance maneuvering close to a potential rear-end crash, existing driver models for such situations typically decide on braking and/or steering by drawing from fixed, situation-independent probability distributions (*23*), whereas it is well known that human avoidance responses vary significantly with not least urgency (*16*, *24*), in ways similar to what is exhibited by the model proposed here. The present work also provides some first steps towards modeling the actual process of taking over from an AV, suggesting that, at least in the presently studied scenario, the take-over could be thought of as an urgency-dependent time delay, combined with a low initial situational awareness.

Indeed, a striking feature of both models is that they provide concrete operationalizations of the concept of situational awareness in the given scenarios, in terms of neurobiologically plausible models of perceptual decisions. This may be a useful complement to existing computational models of situational awareness; e.g., (*13*). Overall, the results presented here seem promising for the general modeling approach of interconnecting several perception and action decision accumulators.

Related but different models have been developed from connectionist and dynamical systems perspectives (*14*, *25*), as well as in the field of naturalistic decision-making (*15*). One useful next step would be to benchmark the present modeling framework against these existing ones, and ideally also against more conceptually different alternatives based on cognitive architectures such as ACT-R (*26*), or from robotics (*27*).

An obvious next step will be to properly test and parameterize the models on detailed data sets of human behavior in the studied scenarios. If such validation tests are successful, the (so far very preliminary) virtual testing results presented here suggest that models at this level of detail, capable of predicting full AV-human interaction sequences unfolding in dynamic traffic scenarios, could be of great use when testing, validating, and optimising future automated transport systems.

ACKNOWLEDGMENT

This work is part of the interACT project, funded by the European Union's Horizon 2020 research & innovation programme under grant agreement 723395. It reflects the authors' views and the European Commission is not responsible for any use that may be made of the information herein.

REFERENCES

1. Fagnant, D. J., and K. Kockelman. Preparing a nation for autonomous vehicles: opportunities , barriers and policy recommendations. *Transportation Research Part A*, Vol. 77, 2015, pp. 167–181. http://dx.doi.org/10.1016/j.tra.2015.04.003.
2. Kalra, N., and S. M. Paddock. Driving to safety: How many miles of driving would it take to demonstrate autonomous vehicle reliability? *Transportation Research Part A*, Vol. 94, 2016, pp. 182–193. http://dx.doi.org/10.1016/j.tra.2016.09.010.
3. Chen, B., D. Zhao, and H. Peng. Evaluation of Automated Vehicles Encountering Pedestrians at Unsignalized Crossings. https://arxiv.org/abs/1702.00785.
4. Waymo. *On the road to fully self driving: Waymo safety report*. https://waymo.com/safetyreport/, accessed 2017-10-16, 2017.
5. Clamann, M., M. Aubert, and M. L. Cummings. Evaluation of Vehicle-to-Pedestrian Communication Displays for Autonomous Vehicles. Presented at *Transportation Research Board 96th Annual Meeting*, 2017.
6. Kesting, A., M. Treiber, M. Scho, et al. Adaptive cruise control design for active congestion avoidance. *Transportation Research Part C*, Vol. 16, 2008, pp. 668–683. http://dx.doi.org/10.1016/j.trc.2007.12.004.
7. Zeeb, K., A. Buchner, and M. Schrauf. What determines the take-over time? An integrated model approach of driver take-over after automated driving. *Accident Analysis and Prevention*, Vol. 78, 2015, pp. 212–221. http://dx.doi.org/10.1016/j.aap.2015.02.023.
8. Kaber, D. B., and M. R. Endsley. The effects of level of automation and adaptive automation on human performance , situation awareness and workload in a dynamic control task. *Theoretical Issues in Ergonomics Science*, Vol. 5, No. 2, 2004, pp. 113–153. http://dx.doi.org/10.1080/1463922021000054335.
9. Purcell, B. A., R. P. Heitz, J. Y. Cohen, et al. Neurally Constrained Modeling of Perceptual Decision Making. *Psychological Review*, Vol. 117, No. 4, 2010, pp. 1113–1143.
10. Ratcliff, R., and D. Strayer. Modeling simple driving tasks with a one-boundary diffusion model. *Psychon Bull Rev*, 2013. http://dx.doi.org/10.3758/s13423-013-0541-x.
11. Markkula, G., J. Engström, J. Lodin, et al. A farewell to brake reaction times? Kinematics-dependent brake response in naturalistic rear-end emergencies. *Accident Analysis & Prevention*, Vol. 95, 2016, pp. 209–226.
12. Svärd, M., G. Markkula, J. Engström, et al. A quantitative driver model of pre-crash brake onset and

control. Presented at *Human Factors and Ergonomics Society 61st Annual Meeting*, 2017.
13. McCarley, J. S., C. D. Wickens, J. Goh, et al. A Computational Model of Attention/Situation Awareness. Presented at *Human Factors and Ergonomics Society 46th Annual Meeting*, 2002.
14. Roe, R. M., J. R. Busemeyer, and J. T. Townsend. Multialternative decision field theory: a dynamic connectionist model of decision making. *Psychological Review*, Vol. 108, No. 2, 2001, pp. 370–392.
15. Patterson, R., L. Fournier, B. Pierce, et al. Modeling the Dynamics of Recognition-Primed Decision Making. Presented at *The 9th International Conference on Naturalistic Decision Making*, 2009.
16. Schneemann, F., and I. Gohl. Analyzing Driver-Pedestrian Interaction at Crosswalk : A Contribution to Autonomous Driving in Urban Environments. Presented at *IEEE Intelligent Vehicles Symposium*, 2016.
17. Gold, C., D. Damböck, L. Lorenz, et al. "Take over!" How long does it take to get the driver back into the loop? Presented at *Human Factors and Ergonomics Society 57th Annual Meeting*, 2013.
18. Merat, N., A. H. Jamson, F. C. H. Lai, et al. Transition to manual: Driver behaviour when resuming control from a highly automated vehicle. *Transportation Research Part F: Psychology and Behaviour*, Vol. 27, 2014, pp. 274–282. http://dx.doi.org/10.1016/j.trf.2014.09.005.
19. Lee, D. N. A theory of visual control of braking based on information about time-to-collision. *Perception*, Vol. 5, No. 4, 1976, pp. 437–459.
20. Bogacz, R., E. Brown, J. Moehlis, et al. The Physics of Optimal Decision Making : A Formal Analysis of Models of Performance in Two-Alternative Forced-Choice Tasks. *Psychological Review*, Vol. 113, No. 4, 2006, pp. 700–765. http://dx.doi.org/10.1037/0033-295X.113.4.700.
21. Friston, K., R. A. Adams, L. Perrinet, et al. Perceptions as hypotheses : saccades as experiments. *Frontiers in Psychology*, Vol. 3, 2012. http://dx.doi.org/10.3389/fpsyg.2012.00151.
22. Lu, L., G. Ren, W. Wang, et al. A cellular automaton simulation model for pedestrian and vehicle interaction behaviors at unsignalized mid-block crosswalks. *Accident Analysis and Prevention*, Vol. 95, 2016, pp. 425–437.
23. van Auken, R. M., J. W. Zelner, D. P. Chiang, et al. *Advanced Crash Avoidance Technologies (ACAT) Program - Final Report of the Honda-DRI Team*. Publication DOT HS 811 454. USDOT/NHTSA, 2011.
24. Lee, S. E., E. Llaneras, S. Klauer, et al. *Analyses of rear-end crashes and near-crashes in the 100-car naturalistic driving study to support rear-signaling countermeasure development*. Publication DOT HS 810 846. USDOT/NHTSA, 2007.
25. Pellecchia, A., C. Igel, J. Edelbrunner, et al. Making Driver Modeling Attractive. *IEEE Intelligent Systems*, Vol. 20, No. 2, 2005, pp. 8–12. http://dx.doi.org/10.1109/MIS.2005.31.
26. Keller, J., C. Lebiere, R. Shay, et al. Cockpit system situational awareness modeling tool. In *Human Performance, Situation Awareness and Automation: Current Research and Trends: HPSAA II* (D. A. Vincenzi, M. Mouloua, and P. A. Hancock, eds.), Psychology Press, pp. 66–71.
27. Bai, H., S. Cai, N. Ye, et al. Intention-Aware Online POMDP Planning for Autonomous Driving in a Crowd. Presented at *IEEE International Conference on Robotics and Automation*, 2015.