Interrupted by my car? Implications of interruption and interleaving research for automated vehicles

Christian P. Janssen[a,*], Shamsi T. Iqbal[b], Andrew L. Kun[c], Stella F. Donker[a]

[a] Utrecht University, Experimental Psychology and Helmholtz Institute, Heidelberglaan 1, 3584 CS Utrecht, the Netherlands
[b] Microsoft Research, One Microsoft Way, Redmond, WA 98052, United States
[c] University of New Hampshire, Electrical and Computer Engineering, Kingsbury Hall, Durham, NH 03824, United States

ARTICLE INFO

ABSTRACT

As vehicles of the future take on more of the driving responsibility and the role of the driver transitions into more of a monitoring capacity, the traditional notions of interruption and attention management needs to be reconsidered for automated vehicles. We argue that the transfer of control between the automated vehicle and the human driver can be considered as an interruption handling process, and that this process goes through a series of ten explicit stages. Each stage has its own characteristics and implications for practice and future research. Therefore, in this paper we identify for each stage what is known from theory, together with important implications for safety, design, and future research, especially for human-machine interaction. More generally, the framework makes explicit that it is not appropriate to think of transfer of control as a single event or even small set of events. The framework also highlights that it might not be realistic to expect human drivers to immediately respond correctly to a system initiated request to transfer control, given that humans interleave their attention between non-driving and driving tasks, and given that a transition constitutes of multiple stages. These nuances are accounted for in the framework.

1. Introduction

In manual driving, the human operator is responsible for the moment-to-moment control of all driving-related functions. Nevertheless, drivers interleave driving with non-driving related tasks (Dingus et al., 2016; Klauer et al., 2014). Although there might be occasional benefits (Atchley and Chan, 2011), more generally, performing other tasks, such as making a phone call, can distract from the driving task. Visual-manual tasks in particular seem to lead to poorer driving performance (e.g., see meta-reviews in Caird et al., 2008, 2014, 2018; Horrey et al., 2009), increase the likelihood of a crash, especially for novice drivers (Klauer et al., 2014), and seem to precede a majority of traffic incidents (Dingus et al., 2016).

Driver distraction has also been explained using theoretical models. One prominent class of models is that of task interruptions (e.g., Altmann and Trafton, 2002; Boehm-Davis and Remington, 2009; Borst et al., 2015; Couffe and Michael, 2017; Salvucci and Taatgen, 2008; 2011; Sanderson and Grundgeiger, 2015). These models have been applied to many domains, given the wide spread of multitasking and interruptions in our daily lives (Janssen et al., 2015). These models typically distinguish two tasks, commonly labelled as (1) an original,

primary task, and (2) an interrupting, or secondary task that interrupts the primary task temporarily. In the case of driver distraction, driving is traditionally considered to be the original, primary task and other tasks the interrupting, secondary task (e.g., making a phone call, listening to the news, interacting with an in-vehicle system, or having a conversation).

As cars become more and more automated, this perspective needs to change. As the automated system's capabilities improve (i.e., as the level of automation increases, SAE International, 2014), the frequency with which human supervision and intervention is needed will be reduced. In effect, human drivers might start focusing on other things to do, only to be interrupted occasionally by their car to assist in the drive. Indeed, one of the motivations for automation is that having a reliable automated vehicle will allow humans to reclaim their time for work and play (Kun et al., 2016). In one extreme view, driving might even be considered the 'distraction' that interrupts the human from doing other tasks (see also Hancock, 2013). To further this point, a meta-review by De Winter et al. (2014) shows that as in-car automation increases, drivers (a) distract themselves more, (b) have less awareness of the traffic situation around them, and, in effect, (c) have increased delays in responding to critical incidents.
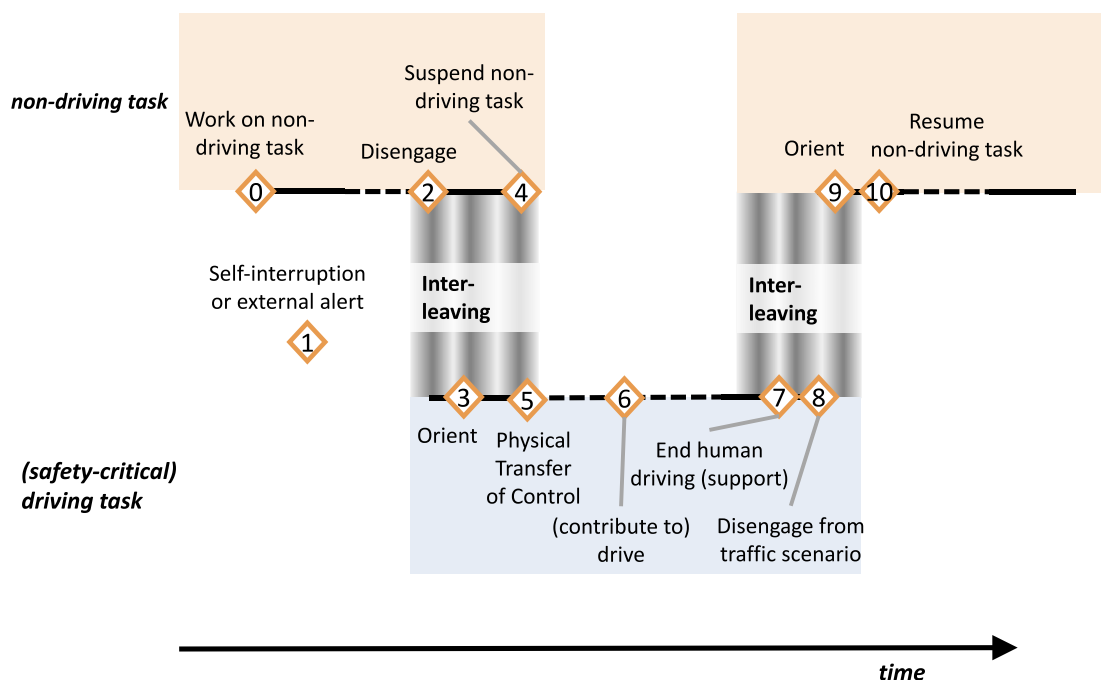
Fig. 1. The stages of a transition of control in an automated driving context, as seen from an interruption perspective. Instead of thinking of a transition of control as a single, or small set of stages, based on the interruption literature multiple stages can be distinguished. The figure is a modification of Fig. 1 in Boehm-Davis and Remington (2009). We added more explicit labels for driving (as the interrupting task) and also explicitly highlighted stages for interleaving. Note that stages 2–4 are not the mirror image of stages 7–10. We comment on the order of stages 8 and 9 in their relevant sections, as these stages might also occur in the inverse order (i.e., there might be interleaving).

Given the likelihood that drivers work on non-driving tasks in (semi-) automated vehicles, and the serious nature of traffic incidents, it is important that there is also a theoretical understanding of interruptions and task switching in the context of (semi-) automated vehicles. Based on theory, predictions can be made for future scenarios about distraction. These predictions can then be used to improve the design of automated vehicles, and to inform policy on automation and traffic safety. Fortunately, such predictions can be made by reviewing the existing literature on task interruptions and task interleaving, and applying those frameworks to (semi-) automated vehicles.

In this paper, we modify the framework of task interruptions to apply to automated driving settings that involve human interaction. This modification, discussed in the next section, makes explicit that it is not appropriate to think of transfer of control as a single event or even a small set of events, but as a series of stages instead. Such explicit consideration of multiple stages (instead of a single event) allows for further research, design and interventions at each stage, and therefore opens up opportunities for further research and insight.

Another contribution of our paper is that it highlights that it might not be realistic to expect human drivers to immediately respond correctly to a system-initiated request to transfer control, given that humans interleave their attention between non-driving and driving tasks, and given that a transition proceeds through multiple stages. This contrasts with the current literature which typically focuses on fast transitions of control and measurement of the minimum time needed to transition from the non-driving task to driving (e.g., Gold et al., 2013).

2. Framework: the stages of interruption applied to transition of control

Interruptions and multitasking research has a long history in human factors and general psychology (e.g., Telford, 1931) and in driver distraction specifically (e.g., Senders et al., 1967). It also is a recurring theme in human-computer interaction (e.g., Brumby et al., 2019; Couffe and Michael, 2017; Gould et al., 2012; Janssen et al., 2015; Li et al.,

2011; McFarlane, 2002; McFarlane and Latorella, 2002; Naujoks et al., 2017; Rivera-Rodriguez and Karsh, 2010; Salvucci et al., 2009). Detailed theoretical models and frameworks have been developed to describe and predict the cognitive processes involved in an interruption (e.g., Altmann and Trafton, 2002; Boehm-Davis and Remington, 2009; Borst et al., 2015; Couffe and Michael, 2017; Salvucci and Taatgen, 2008, 2011; Sanderson and Grundgeiger, 2015). Although these frameworks vary in some details, they all describe the interruption process as a sequence of *stages*. At each stage, the context (or experimental manipulations) can influence how the interruption is handled, and how well the original task is performed.

We focus on one framework in particular, by Boehm-Davis and Remington (2009), as this framework is the most detailed about the entire interruption process. We have adapted it to fit to the scenario of interruptions in an automated vehicle. Fig. 1 shows the result. The starting point of our version of the framework is that the human driver in a semi-automated vehicle is engaged in a non-driving (or original) task as the car is in an automated driving mode. From this starting point onwards, the transition to a driving task, and the later return to the non-driving task, can be characterized in ten stages: (1) the driver self-interrupts or receives an external warning that their input is needed for the driving task, (2) they disengage for the first time from their original task to start a period of interleaving attention between the original task and the driving task, (3) they orient towards the traffic environment and the car, (4) they suspend their original task, (5) there is a physical transition of (part of the) control of the vehicle or some input from the human driver is needed, (6) the human driver drives or contributes crucial input to the car to drive, which is followed by another interleaving period during which (7) the human no longer needs to provide input to the car, (8) they disengage from driving, (9) orient to their original non-driving task, and (10) resume suspended activities on their original task. Note that the "ramp down" stages after the interruption (i.e., stages 7 to 10) are not the direct mirror opposites of the "ramp up" stages leading up to the interruption (1 to 4), due to the safety-critical nature of the driving task in between: this requires sufficient

preparation before take-over, and might invoke some monitoring after the task has been transferred back to the driving system again.

In Textbox 1 we provide three example scenarios to illustrate how the framework and stages map to specific settings. As further reference, in Textbox 2 we provide a glossary in which the most important key terms and theoretical constructs are defined.

Before delving into the details of each stage, we highlight three modifications that we made to the original framework (Boehm-Davis and Remington, 2009) to make it appropriate for automated driving. First, we explicitly labeled stage 0 and stage 6, and we also included the option of self-interruption given that self-interruptions are prevalent, at least in office settings (Dabbish et al., 2011).

Second, we explicitly labeled two sequences of stages as interleaving stages. This contrasts with the classical framework of interruptions (Boehm-Davis and Remington, 2009), in which the stages are mostly treated sequentially (i.e., with each stage occurring only once, and in order). Instead, for an automated driving context we suspect that some stages can be repeated over time, making an interleaving perspective (i.e., going back-and-forth between the driving and non-driving task) more appropriate. This will be explained in more detail for the relevant stages (and see also the example scenarios in Textbox 1).

Third, in contrast to the original literature, we avoid the terms "primary" and "secondary" task. These terms suggest some priority, with driving traditionally as "primary" and "original" task. However, such a distinction might get blurred in automated systems and drivers might not always adhere to it. Specifically, as cars become more automated, less time might be spent on driving (cf., de Winter et al., 2014). Instead of using the terms "primary" or "secondary" task, we recommend using terms that explicitly describe the nature or characteristics of the task, such as "time-sensitive", "safety-critical", and "in focus".

We will now go through the stages step-by-step. For each stage, we discuss the current knowledge from interruption theories and related work, as it applies to driver distraction settings. In addition, we discuss the implications for safety, design, and further theoretical study.

Textbox 1
: Example future scenarios

The following three futuristic scenarios illustrate how the framework can be applied to describe the stages of a transition of control, how consideration of the stages can help understand successful (scenario 1 and 3) and unsuccessful interruption handling (scenario 2). For each scenario we describe what type of measurement can be used to detect the onset or offset of the stage.

Scenario 1: Successfully handling e-mail before transitioning control to navigate roadworks in response to a pre-alert

A driver is in a car that drives at SAE level 3 on a highway, and is answering e-mails on a tablet (stage 0). During the drive, the car receives information from an online map system that roadworks are coming up in 3 kms and the car provides a pre-alert to the human driver that action is required in 30 s at the latest (stage 1, measured through external alert trigger). The driver briefly disengages from the e-mail task (stage 2, measured through eye-gaze and lack of tablet interaction) to orient (stage 3, measured through eye-gaze) to the driving situation to evaluate when to take over control. In this scenario, the driver notices that there is a merge before the roadworks start, but that it is still over 1 km away. Being close to finishing an e-mail, the driver quickly writes the last two sentences and sends off the e-mail (i.e., reaches a natural breakpoint) to be fully disengaged from the email task (stage 4, measured through eye-gaze). In the meantime, the driver occasionally looks up at the road to make sure that nothing has changed and no urgent action is needed (interleaving, measured through eye-gaze and tablet interaction). Eventually, they press a button in the car to take over control (stage 5, measured through button press and other sensors on steering wheel), before

the roadworks start - to gain a stable position and to then navigate through the lane merge and the road works (stage 6, measured through steering angle, vehicle position).

Once the roadworks are done, the driver continues to drive for a bit and then presses a button to hand back control to the car (stage 7, measured through button press). However, to be absolutely sure that the hand-over is successfully completed, they monitor the car for a while before disengaging fully (stage 8; measured through eye-gaze and lack of tablet interaction). In the meantime, the driver simultaneously prepares to re-engage with the non-driving task by placing the tablet on the lap and launching the e-mail application, while looking at the road in between (interleaving). Once the human driver is confident in the car's ability to drive safely by itself without human intervention, the e-mail task is being resumed (stage 9, measured through eye-gaze and tablet interaction). The driver can resume answering e-mails quickly, as the previous e-mail was already sent and now a completely new e-mail can be started (stage 10, eye-gaze and tablet interaction). To summarize, although in this scenario some time was taken to finish the e-mail before acting on traffic, this did not come at the cost of road safety (successful handling of critical event) or e-mail handling (quick resumption).

Scenario 2: Rushed and unsafe transition of control after a late (last-minute) alert

Now let's imagine scenario 1 in a case where the alert arrives only 10 s before the roadworks start. The transition through stages and their measurement in an experiment is similar. However, this time, the human driver only has 10 s to cycle through stages 1 (when the alert arrives) to 6 (when the road work starts), instead of the 30 s available in scenario 1. The driver disengages (stage 2) and while orienting (stage 3) sees that roadworks are imminent, drops the tablet (stage 4), and presses the button to take control (stage 5). Unfortunately, 10 s is a very short time to abandon the email task and safely start the driving task. For this reason the driver has difficulty in keeping the car within the lane markers (stage 6) and needs to slow the vehicle down by pressing the brakes rapidly. Additionally, as the orientation phase (stage 3) was short, the driver failed to notice that there was a merge between two lanes shortly before the roadworks. Fortunately, another driver halted, but there was a near crash.

After the roadworks, the driver wants to return to the urgent e-mail which they had to interrupt abruptly because of the need to take over driving control immediately. Given that the driver also wanted to respond to the e-mail urgently, they take the minimum amount of time to drive themselves, quickly hand back control to the car (stage 7) and disengage with traffic quickly (stage 8). As in this scenario the e-mail was interrupted in the middle of writing it, a longer time is needed to reorient to the e-mail task: the driver starts re-reading the e-mail and the reply from the top to get back into the flow of writing (stage 9) before resuming the task by writing the last two sentences (stage 10). In other words, the quick reacting that was required for driving after stage 1 both negatively affected the driving task (near crash) and the non-driving task (longer task resumption, due to need to build context on what the e-mail was about).

Scenario 3: Quick self-interruption to handle heavy rain conditions

In a third scenario, the human driver is answering e-mails while the car drives on a highway at SAE level 3. This time, there is a slight drizzle. Over time, the driver notices that the rain starts pouring faster and heavier, so they decide to self-interrupt (stage 1; in a lab experiment rain can be controlled, and time interval that passes before one looks up as a function of rain intensity can be determined) and look up at the road. As in scenario 1, the driver takes the time to finish the e-mail as there is no immediate need to take over control, then puts the tablet aside (stage 4) to only observe the car and traffic (orient, stage 3). When the driver notices that the car is leaving little headway to the car in front, they adjust the settings of the adaptive cruise control to increase the distance – as might be more appropriate during rainy conditions (stage 5, button press). In this case, the driver progresses

directly from stage 3 to stage 5, as they have already suspended the non-driving task (stage 4). The driver observes how the car handles itself, and becomes confident that there is no need to contribute further (stage 6), and quickly ends contributing (stage 7). Subsequently, the driver resumes the e-mail task (stage 9), but occasionally looks at the road and the car's dashboard to monitor if all systems are functioning properly (interleaving between stage 7 and 10). That is, a quick self-interruption at an appropriate time leads to fast adjustment of the vehicle's functioning and fast resumption of the original task. Of course, in such a scenario, a driver might also decide to continue to monitor the traffic situation more actively, thereby delaying the stage of working on the original task uninterrupted.

Textbox 2
: Glossary

• **Attention** Process related to bringing or having something (i.e., a stimulus or a task) in one's focus. Going back to at least William James (1890), attention has had many different definitions, each placing slightly different accents (e.g., Kahneman, 1973; Norman and Shallice, 1986; Van der Stigchel, 2019; Wickens and McCarley, 2008). Our definition is made as broad as possible, while acknowledging that one's attention is typically limited. In the context of interleaving and interruptions, attention is typically directed more to one task than the other, and a relevant question is how attention is "divided" between tasks.
     • **Automated vehicle** Transportation machine in which some aspect of the driving task is (partially) handled by a machine. There are multiple levels of automation, as defined by the SAE (SAE International, 2014).
     • **Disengage** Stage within the interruption process at which a task is temporarily abandoned to orient to another task. Note however, that after initial disengagement of a task a user might still return to it before fully suspending it. In other words, an initial disengagement might start an interleaving process.
     • **Forewarned interruption:** Situation where someone is notified by a so-called *forewarning* of an upcoming interruption or event, but no immediate action is required. An example in driving is the use of pre-alerts. Forewarned interruptions allow a user to go through the interruption process more carefully, for example by allowing them to suspend their original task (stage 4) at a natural breakpoint (rather than immediately) and to take time to orient on driving (stage 3) before the physical transfer of control (stage 5).
     • **Immediate interruption:** Situation where handling an interruption can not be postponed; immediate action is needed. An example is when a colleague rushes into an office to talk without allowing you to finish what you were working on. An example in driving is the use of *last-minute alerts/warnings*. The current Tesla model S cars, for instance, give an auditory warning when the driver has not touched the steering wheel for a specific time interval, and the driver needs to act on it immediately by grabbing the wheel. In experimental settings (especially outside of the driving domain), participants might even be "locked out" of completing their original task, thereby forcing a participant to work on the interrupting task (e.g., Li et al., 2008).
     • **Interleaving** Going back and forth between two (or more) tasks. In an interleaving process the assumption is that at each moment in time only one task has the main focus, but over time different tasks are being worked on. For example, going back and forth between steering a vehicle and typing on a phone (Janssen et al., 2012). Interleaving is a subclass of multitasking, in which there can also be parallel processing of tasks and stimuli. In our framework there are two interleaving phases: when starting to move attention from the original non-driving task to the driving task (stage 2–5) and when returning attention from the driving-task back to the original task (stage 7–10). More detailed definitions of interleaving can be found in Payne et al. (2007) and Janssen et al. (2019).
     • **Interruption** An event that temporarily shifts attention

away from the task that one was originally working on. *External interruptions* are triggered by an external cue, something outside of the person that is interrupted, such as the ring of a phone call, an incoming e-mail sound, or a person walking into an office. *Self- (or internal) interruptions* are triggered without the presence of an external cue, for example, when one recalls an important task. See also Miyata and Norman (1986) and González and Mark, 2004.
     • **Multitasking** Process where two or more tasks or activities are performed, or seemingly performed, in parallel. Different types of multitasking can occur ranging from task interleaving (see above) to (seemingly) parallel processing in the brain. If there is no negative effect of the processing of one task or activity on the other (in terms of e.g. speed and accuracy) there is "perfect" time-sharing, and one could speak of parallel multitasking or parallel information processing. Perfect time-sharing can for example occur if the required mental resources (Wickens, 2008) and processes (Salvucci and Taatgen, 2011) of two tasks do not overlap. There is extensive debate whether the brain processes all information parallel, or whether there are "serial bottlenecks" that prevent parallel processes from occurring (Howes et al., 2009). In typical interleaving scenarios, at least some sequential (non-parallel) processing is assumed, thereby not allowing for perfect time-sharing (Janssen and Brumby, 2010; Wickens et al., 2015).
     • **Natural breakpoint** A position higher up in the task hierarchy, typically after completion of a subtask. Interleaving at such natural breakpoints compared to other positions in the task hierarchy typically is beneficial, as it is associated with less stress (Adamczyk and Bailey, 2004), reduced memory load (Borst et al., 2015) and mental workload (Bailey and Iqbal, 2008), and can provide beneficial speed-accuracy trade-offs (Janssen et al., 2012). It is therefore a relatively "natural" point to suspend one task in favor of another.
     • **Orienting** The process of building context about a task other than the task a person was focusing on so far. In our interruption framework there are two instances at which orienting takes place: at stage 3 (orienting to the interrupting task) and at stage 9 (orienting to the original task after the interrupting task is completed). Orienting can be thought of as starting to build relevant context, or situational awareness (Endsley and Garland, 2000) to act timely and accurately on the (new, interrupting) task.
     • **Original task** Task or activity that one is mostly working on before an interruption. In the context of automated driving this might be a non-driving related task (e.g., making a phone call, checking e-mail).
     • **Pre-alert (or forewarning)** Advanced warning to signal an upcoming critical event. In the context of automated driving, for example, pre-alerts can be used to foreshadow that a transition of control is coming up before the eventual last-minute warning is provided (Van der Heiden et al., 2017).
     • **Resumption** The process of picking up where one left (i.e., suspended) a task before temporarily disabandoning it (for example, due to an interruption) (see e.g., Altmann and Trafton, 2002; Borst et al., 2015).
     • **Suspension** Temporarily abandoning a task, for example to handle an interruption. In contrast to the disengage stage, after the suspension stage a user is assumed to first handle an interrupting task before returning to their original task that they suspend.
     • **Task** Goal directed activity (Card et al., 1983), which typically involves completing various substeps (i.e., subtasks).
     • **Task hierarchy** Arrangement of subcomponents of a task such that an end goal (i.e., task or goal directed activity) is achieved by first completing subcomponents (subtasks), that themselves can consist of smaller units. Goal-directed tasks typically have such a hierarchical structure (Card et al., 1983). For example, the task of calling someone might consist of the subtasks of opening a phone application, looking up their name, and selecting the right number (e.g., office or work). Each of those subtasks might themselves consist of different subtasks. For example, opening the phone application might consist of turning on

your phone and of scrolling to the window that contains the app. Each of those might again contain smaller steps.

• **Transition of control** The process during which there is a change in who operates the driving task, or a specific part or component of the driving task. Transitions can both be from human to machine (i.e., between stage 1 and 5 in Fig. 1), and from machine to human (i.e., at stage 7 in Fig. 1). Transition of control is sometimes also referred to as *"take-over"* (typically, when referring to an active act, such as when the human takes over control from the machine) or "hand-over" (typically, when referring to a more passive act, such as when the human gives control back to the machine). See Mirnig et al. (2017) for a more detailed discussion.

2.1. Stage 0: work on non-driving task

2.1.1. Theory

Drivers of currently conventional, non-automated vehicles perform other tasks while driving (Dingus et al., 2016; Klauer et al., 2014). While the negative effects of multitasking while driving a manual car are well documented in the literature, the research is more sparse on the effects of drivers engaging in non-driving tasks in a vehicle operating in (higher levels of) an automation mode. Based on a meta-review, it is expected that non-driving tasks are performed more frequently as the automation level of the car increases (de Winter et al., 2014). The types of tasks that have already been observed in naturalistic driving studies (Dingus et al., 2016; Klauer et al., 2014), or that people anticipate doing in automated vehicles (Pfleging et al., 2016) range from reaching to or interacting with objects in the car, to personal hygiene tasks, to dancing, to using phones, to talking and interacting with other passengers.

Although people engage in non-driving tasks, decades of research have shown that these tasks can distract from driving itself. Research has mostly studied the negative impact of visual-manual tasks and (cell-phone) conversations in the context of manual driving (e.g., Caird et al., 2008, 2014; Horrey et al., 2009). Activities that do not require visual-manual interaction can also distract. For example, this is the case for some conversations (Caird et al., 2008; Horrey et al., 2009). In particular, the more complex a conversation or linguistic task is (e.g., generating creative text versus simply repeating text), the more it affects driving performance (Iqbal et al., 2010; Kunar et al., 2008; Strayer and Johnston, 2001). There are even hypotheses that thinking itself can distract (e.g., Engström et al., 2017; Salvucci and Beltowska, 2008).

When people need to divide their time between two tasks, the time-on-task is affected by their priorities (Janssen and Brumby, 2010; Janssen et al., 2012). In the context of an automated vehicle, the non-driving task might be prioritized over the driving task, as the need for interaction with the car while it is in self-driving mode might be diminished. Therefore, even less attention might be given to monitoring the car and the driving environment.

Although our emphasis in the above review has mostly been on the distracting nature of non-driving tasks, performing non-driving activities might in some cases be beneficial for driving, as it might help in staying vigilant and preventing underload (Atchley and Chan, 2011; Young and Stanton, 2002). This is also relevant for automated driving conditions, where the reduced need to contribute to the driving task might be reclaimed for (non-driving related) work and play (Kun et al., 2016).

2.1.2. Implications for safety, design, and future research

No assumption can be made that people pay attention to the traffic environment while they are also performing other tasks. Tasks that require visual-manual interaction in particular are distracting, but tasks that do not require a driver to take their eyes off the road or hands off the wheel can also distract. Non-driving related tasks might occasionally be beneficial to avoid underload. Further insight is needed on the neuro-cognitive processes that underlie these tasks to understand why, how, and when these tasks are distracting in automated driving settings.

2.2. Stage 1: self-interruption or external alert

2.2.1. Theory

Interruptions can be initiated by one-self (self-interruption), or an external alert. It is unknown how frequently each type occurs in driving settings, but data is available from other domains. For example, self-interruptions make up approximately 50% of the observed interruptions in the office (González and Mark, 2004). Self-interruptions might be needed for lower levels of automation, where the human driver needs to actively monitor the car (i.e., including SAE level 3 vehicles that only drive automated in a limited set of driving contexts). As the automation level increase, we anticipate that the human can more and more rely on external alerts to trigger their assistance. If self-interruptions do occur, we assume that the driver has some awareness of the traffic situation, and therefore might go through the other stages of the interruption process, as presented in Fig. 1, faster.

Self-interruptions are also referred to as discretionary task interleaving, as the task switch is due to one's own choice (i.e., discretion), and the person might go back and forth between tasks (i.e., interleaves). Factors that make people stay on a task include the task's engagement and immersion level, and people's inherent desire to maintain a balance in the activities they are engaged in (Lewin, 1943), but also factors related to the task design. Factors that make people leave a task include a need for a break, for emotional homeostasis, or for rejuvenation (Mark et al., 2015). For goal-directed tasks (e.g., writing an e-mail, or filling out a spreadsheet) one of the important factors is having "natural breakpoints" in the task, which are formed by clusters of substeps in the task hierarchy (Janssen et al., 2012). Many dual-task studies observed switches at such natural breakpoints (e.g., Bailey and Iqbal, 2008; Bogunovich and Salvucci, 2010; Brumby et al., 2009; Iqbal et al., 2005; Janssen et al., 2012; Janssen and Brumby, 2010; Kun et al., 2013; Miyata and Norman, 1986; Payne et al., 2007; Salvucci, 2005; Yang et al., 2011). Interleaving at natural breakpoints can reduce stress (Adamczyk and Bailey, 2004), reduce the need to keep task-relevant information in memory (Borst et al., 2015), reduce mental workload (Bailey and Iqbal, 2008), and create beneficial speed-accuracy trade-offs for dual-task performance (Janssen et al., 2012). Moreover, interleaving at a natural breakpoint makes it easier to later resume the task (i.e., stage 10), as one can start with a new "sub-task" instead of resuming in the middle of a sub-task.

External interruptions are triggered by something outside of the individual. The interruption can come from immediate interruptions and forewarned interruptions. For immediate interruptions, immediate action might be required. An example is when a colleague rushes into an office to talk without allowing you to finish what you did. In experimental settings, a participant might even be "locked out" of working on their original task, thereby forcing the participant to work on the interrupting task (e.g., Li et al., 2008). Immediate interruptions are elicited in current vehicles by last-minute alerts. However, they are less suitable for vehicles with a high level of automation: cars should rather gradually transition the driver from a non-driving task to driving, allowing them to gain driving context during the transition. Studies have shown that a *minimum* warning time of 5 to 8 s is needed for drivers to safely take control from the automation in the case of last-minute alerts (e.g., Gold et al., 2013; Mok et al., 2017).

In contrast to immediate interruptions, forewarned interruptions give a person a notification about an upcoming potential interruption, but with the option to defer their response or to not act immediately. For automated driving, an example are pre-alerts that warn a driver twenty seconds before a transition of control needs to be acted on (Van der Heiden et al., 2017; Borojeni et al., 2018).

Independent of whether a warning is for an immediate interruption or a forewarned interruption, at least four factors influence the effectiveness of the associated notifications: presentation modality, timing of the alert, reliability of the alert, and required cognitive processing. For the presentation modality, a variety of options have been tested (see also Baldwin and Lewis, 2014; Petermeijer et al., 2017). Visual alerts and auditory alerts tend to dominate the automotive industry. Neither is bullet-proof however, as visual alerts can be overlooked, and auditory alerts might not be noticed when other sound sources are playing loud (e.g., radio), if the ears are obstructed (e.g., earphones), or if the driver has hearing impairments. Other modalities, such as haptics, are used more experimentally for in-car tests, but have more variable results in effectiveness (for a review, see Meng and Spence, 2015).

The timing of alerts has been investigated extensively (e.g., Borojeni et al., 2018; Dogan et al., 2017; Gold et al., 2013; Mok et al., 2015, 2017; van der Heiden et al., 2017; Walch et al., 2015). While a minimum warning time of 5– 8 s might work in some scenarios as a last-minute alert (e.g., Gold et al., 2013; Mok et al., 2017), recent studies argued that even earlier warnings (i.e., forewarnings) might be necessary and more effective (Borojeni et al., 2018; van der Heiden et al., 2017). Such forewarned cases allow the driver to finish their task at a natural breakpoint and allow more time to gain situational awareness. Such additional time is particularly relevant for automated driving cases where human input is requested infrequently.

The reliability of an alert is affected by the frequency of false alarms. For example, nuisance alerts might lead to the "cry wolf effect" (Breznitz, 1983; Sorkin, 1989), in which people ignore alerts altogether (though see Wickens et al., 2009 for conditions under which this might not be avoided).

Recent neuroscience studies suggest that the brain is less susceptible to unexpected auditory signals when people are driving (Wester et al., 2008), or when one is being driven by an automated vehicle (van der Heiden et al., 2018). Therefore, it should not be assumed that an alert that is provided is also (fully) processed. Moreover, an unexpected alert might also startle a driver and thereby negatively affect the ability to effectively take over control of the driving (Bliss and Acton, 2003).

2.2.2. Implications for safety, design, and future research

In-car alerts can alert the driver to a transition of control of the car. As uni-modal alerts might be missed, and as the brain's susceptibility to unexpected alerts is reduced while driving or being driven, multi-modal alerts are essential for critical notifications. Forewarned alerts provide an additional layer of security. The exact timing of such warnings is not yet determined, but intervals between 20 and 40 s seem to be current guesses (Borojeni et al., 2018; Merat et al., 2014; van der Heiden et al., 2017). As these forewarned alerts are not yet common practice, there is also an opportunity for training and regulation: should we test drivers' ability to take over control appropriately (see also Inners and Kun, 2017)?

2.3. Stage 2: disengage from original non-driving task

2.3.1. Theory

Gracefully disengaging from an ongoing task to switch to another task is a critical component in multitasking, as appropriate disengagement can help with easier resumption later. People have a tendency to disengage from tasks at natural breakpoints (e.g., Bailey and Iqbal, 2008; Bogunovich and Salvucci, 2010; Brumby et al., 2009; Iqbal et al., 2005; Janssen et al., 2012; Janssen and Brumby, 2010; Kun et al., 2013; Miyata and Norman, 1986; Payne et al., 2007; Salvucci, 2005; Yang et al., 2011). Therefore, immediate disengagement might not be expected, but delayed until such a breakpoint is reached. Moreover, cases outside of the driving domain have been reported in which participants tend to stick with their original task, even though they need to switch to another more urgent and time-critical task. This is also referred to as "cognitive lock-up" (e.g., Neerincx, 2003; Schreuder and Mioch, 2011).

Taken together, although stage 2 is the first step towards disengagement from the original task, it should not be assumed that this task is given up immediately. Instead, during stages 2 to 4 there might be occasional *interleaving* of attention between the non-driving task and orienting towards the driving task. Allowing for such adequate disengagement can delay when the driver takes control of the vehicle. However, as we will see for later stages (e.g., stage 5), it might prevent or reduce the driver from having active thoughts about, or continuing to work on, other non-driving tasks at a time where their full attention should be on driving.

In the classical interruption framework of Boehm-Davis and Remington (2009) there is not an explicit interleaving phase, as all stages are mostly treated sequentially. In automated driving we anticipate that there will be interleaving between driving and non-driving activities. Although there will be more back and forth between such tasks, the start and end of each stage can still be detected using for example eye-tracking (where do people watch) and interaction data (where do people 'act': on their phones or on the car's interface?). This can help to detect the first moment of disengagement (stage 2, i.e., the first moment after an alert where someone looks away from their original task), the first moment where they orient to the road (stage 3), the last moment they work on the original task (stage 4: suspension) before starting to contribute to the driving task (stage 5). In between these stages, there might be multiple glances at each of the tasks. See also our example scenarios in Textbox 1.

2.3.2. Implications for safety, design, and future research

People need some time to disengage from the tasks they were working on. This disengagement time should be considered when evaluating the safety of a system; instantaneous disengagement should not be assumed. There are ways to design for faster disengagement, for example by including occasional natural breakpoints in the task structure. However, not every task has such natural breakpoints – especially tasks that are not goal-directed. As drivers can bring many tasks with them to the car through their phone, there is a benefit to explore whether there are general solutions that can aid interleaving and resumptions, and what factors contribute to success without breaking a user's "flow" when they are mono-tasking.

2.4. Stage 3: orient to driving task

2.4.1. Theory

When drivers orient themselves towards the driving task, they need time to get a reliable understanding of the environment and the system state to react appropriately. The amount of time that is needed depends on (1) the characteristics and complexity of the traffic situation (for example: is the car driving on a clear highway without other traffic?), and (2) whether the driver has been keeping track of the driving environment before they needed to orient themselves to it. If the driver was triggered by an external interruption in stage 1, it is likely that more time is needed to orient to driving, as the driver so far had not noticed themselves that their assistance was needed.

Theory on situational awareness (Endsley and Garland, 2000), including how it applies to traffic (e.g., Gugerty, 1997; Kass et al., 2007) can guide these efforts. However, more work is needed for automated driving specifically. In the traditional interruption models (e.g., Boehm-Davis and Remington, 2009) suspension of the original task (stage 4) comes after the orientation stage, not before it. Immediate orientation and full attention can therefore not be assumed. Moreover, as we argued in stage 2, there might be interleaving of attention between driving and non-driving task at this point.

2.4.2. Implications for safety, design, and future research

Similar to the preceding stages, drivers cannot be assumed to respond immediately to a transition of control request in the orientation phase. Instead, this depends on the degree of interleaving that preceded

the orientation phase. The required take-over time also depends on the complexity of the transfer of control (e.g., is a "yes/no" response sufficient, or is there some critical steering action needed?), the driver's engagement in the original interrupted task, and the required time for disengagement. Finally, there should be clear guidance on what information is important for the current request to transition control from the car to the human driver. Prioritizing relevant information acquisition and providing support to quickly acquire that information should contribute towards improving safer task orientation.

An open theoretical question is what affects the length of the period during which situational awareness is built. This can be studied in various ways. One way that we suspect might be valuable is to look back at, and apply, models building on the seminal work of Senders and colleagues (Senders et al., 1967) on how systematic deprivation of stimulus input affected driving. In these studies, drivers' view of the road was occluded systematically through a visor, for which drivers had to push a button to open it. Senders and colleagues measured how quickly drivers pushed the button after the visor had closed to see the surrounding traffic again, as a function of road demands. This work has recently been tied to theoretical models of attention (e.g., Kujala et al., 2016; Chen and Milgram, 2013), but has so far only considered regular driving and not yet distracted driving. In addition, the role of other modalities such as audio (e.g., traffic sounds, alerts), and motion stimuli (e.g., sense of speed), including deprivation of such stimuli, can be investigated further.

2.5. Stage 4: suspend original non-driving task

2.5.1. Theory

After some period of interleaving between the non-driving task and orienting to the driving task, the driver might finally suspend their original, non-driving task. Intuitively, the suspension of the non-driving task might suggest that this task is no longer placing demands on the human driver. Recent studies, however, suggest that distracting effects can last until after a task is finished. For example, Strayer and colleagues (Strayer et al., 2015) measured drivers' response to an operation span (OSPAN) task while driving, including during and after in-car interacting with a mobile device. The reaction time to the OSPAN task was found to still be significantly longer than during baseline performance for up to 27 s after the in-car distraction ended. What is interesting about the Strayer study is that the in-car distracting tasks mostly relied on visual-manual interaction, not on memorizing information or planning. Therefore, despite the fact that no explicit thought process was needed after the interaction stopped (e.g., there was no need for memory or planning), there were still some cognitive remnants of distraction for a prolonged time.

2.5.2. Implications for safety, design, and future research

An original, non-driving task can create remnant distraction after the driver has finished the task, even if the original task did not heavily rely on cognitive processes (e.g., memory, planning). As this observation has been made recently, there is a need for replication. Moreover, knowledge is needed on what factors influence the duration and strength of this interference in various contexts. This includes a need for knowledge about the cognitive mechanisms that cause the distraction.

Given that engagement with driving might not be in full gear immediately (cf. Merat et al., 2014; see also stage 6), a critical decision for the human driver is to decide whether and when to disengage from their non-driving task. In systems that have a lower level of automation (i.e. SAE level 3), the default should be that the human responds immediately, and systems should be designed to support such quick responses. In a higher level of automation, like SAE level 4, there is more discretion for the human driver whether and when to respond. Design efforts can therefore explore how to best support the human driver in this decision (e.g., Larsson et al., 2015).

2.6. Stage 5: physical transfer of control

2.6.1. Theory

Eventually, the driver takes over control of the vehicle. Although all stages of the interruption process are part of the transition of control process, we see stage 5 as the point where the human driver takes *physical* control of some part of driving for the first time. This physical control can be done through different modalities, including manual and vocal (cf. Mirnig et al., 2017). This will also depend on the tasks that the human driver needs to take over (i.e., basic vehicle control might require visual-manual interaction; decisions on route alternatives might be done vocally).

From an interruptions perspective, the duration of the interruption (in this case: controlling the car) affects the later resumption (step 10 in our diagram). The assumption in typical interruption research is that while a person is working on an interrupting task, they need to keep information on the original task in memory. These memories are thought to decay at an exponential rate (Altmann and Trafton, 2002; Borst et al., 2015). Shorter interruptions are therefore less harmful for later task resumption (as there is less memory decay) and, therefore, to be preferred. A hidden assumption behind the memory theories is that the person is no longer actively thinking about their original task. However, there might be cases where the task, or active thoughts about it, continue, and continue to distract from driving. Adequate disengagement in stages 2 and 4 may help to prevent or reduce such situations.

2.6.2. Implications for safety, design, and future research

Taking the theoretical knowledge from stage 4 and stage 5 together, the implication is that short interruptions are preferred for actions that do not require a full construction of situational awareness. An example could be an interruption by the car to ask the human whether the car should take the fast or the scenic route for a drive. By keeping such interaction short, the original task can later be resumed more quickly as its content have not decayed in memory.

In cases where humans need to act directly in the environment (e.g., take over lateral or longitudinal control of the vehicle), such short interactions are not desired for. For these types of action, the driver needs to gain full situational awareness, which can be gathered if the stage leading into the physical transition is longer and if the in-car interface supports such situation awareness gathering activities.

A significant difference between traditional interruption research, and interruptions in the semi-automated vehicle, is that interruptions in the semi-automated vehicle are typically time-critical. That is, when the car alerts for assistance, some action by the user is needed and perhaps hard to defer. In contrast, in the typical experiment the 'secondary task' is often meant as a manipulation to induce load or distraction on the driver. Given that these serve different purposes, more research is needed that investigates situations in which an original task is interrupted by a safety-critical task, as there might be differences.

2.7. Stage 6: (contribute to) drive

2.7.1. Theory

Even when drivers contribute to the drive, thoughts on tasks they were working on before might still linger in the back of their mind and distract from driving (e.g., Engström et al., 2017; Salvucci and Beltowska, 2008; Strayer et al., 2015). Moreover, given that drivers of non-automated vehicles distract themselves with other tasks (Dingus et al., 2016; Klauer et al., 2014), the assumption that drivers are fully focusing on the road after a transition of control to them might again be incorrect. This lack of full focus (visual and/or cognitive) might negatively affect various measures of driving performance (Kun, 2018).

Nonetheless, it would be ideal if drivers fully focus on driving during stage 6. This contrasts with all the preceding stages of the interruption process. In stage 0, which is also a single-tasking stage

(focusing on a non-driving task), we suggest that it is desirable for the driver to occasionally engage with or monitor the driving task in anticipation of a possible future request to transfer control. Compared to stage 1–4, where drivers might divide their attention between the driving and the non-driving task, in stage 6 drivers ideally should not have another task to deal with. Moreover, one assumption might be that the car has invoked human assistance because the traffic scenario is complex, and requires human attention. This differs from a regular "single-task driving" scenario in which there might occasionally be periods that are monotonous, and where distraction away from the driving might help in staying vigilant and preventing underload (Atchley and Chan, 2011).

Even in cases where a driver is fully focusing on driving after the transition of control, a study by Merat and colleagues (Merat et al., 2014) has shown that driving after the transition of control is not immediately similar to baseline driving without automation. Specifically, in their study, drivers had to take over control of a simulated vehicle after a predictable situation that was system-initiated, or after the driver had been looking away from the road for a predefined interval (i.e., when they were distracted). Although human gaze patterns and lateral control over the vehicle were better in the system-initiated predictable condition, in both conditions it took up to 35–40 s before lateral control of the vehicle was fully stable, and comparable to baseline non-automated driving. This research suggests that even after a driver has actively taken over control of the vehicle, it takes a while before they are in stable control of the vehicle.

2.7.2. Implications for safety, design, and future research

Research suggests that during driving there might be remnant distractions (e.g., Engström et al., 2017; Salvucci and Beltowska, 2008; Strayet et al., 2015), and drivers might not immediately have full control over the vehicle (Merat et al., 2014). These insights from theory again reinforce one of the main implications of the interruption framework: that a transfer of control should be signaled (stage 1) well in advance when possible. This will allow the driver to gain stable control before they need to handle a critical incident. The challenge is that such an advanced warning conflicts with the suggestion for stage 5 that encourages only short driving engagement. A possible solution lies in allowing prolonged periods of interleaving during stages 1 to 4. The driver then knows in advance that their input is needed, can finish tasks they were engaged in at a natural breakpoint (instead of at other points), can take their time to orient to the driving task, and in effect act more effectively (and when possible, in a shorter time interval) on the driving requirements when needed. That is, lengthy preparation might aid safe, fast, and accurate action later.

2.8. Stage 7: end human driving (support) action

2.8.1. Theory

In the typical interruption experiment (outside of driving), an interrupting task has a forced endpoint, which arrives either after a specific time interval or after a specific task step has been completed (e.g., Borst et al., 2015). Such an approach is motivated by the desire of classical interruption studies to study, in a controlled way, how the content and duration of an interruption impacts later resumption. However, this is different for an automated driving setting, where the "interrupting" driving (support) task can be ended when either the human, the system, or both human and system deem the environment safe to transition control back to the automated system (see also Mirnig et al., 2017).

A situation where the car actively takes over control aligns best with the forced interruption and resumption scenarios from classical interruption studies. Although such systems are, to the best of our

knowledge, currently not commercially available, they might be designed in the future to take-over in cases where the system detects that the human is not acting appropriately.

The analog to voluntary transition (i.e., where the human hands over control to the car) aligns better with the classical interleaving literature than with the interruptions literature. In interleaving studies, humans can change tasks at their own discretion. Moreover, within our framework we acknowledge that from stage 7 onward there is again an interleaving *phase*: the input from the human driver is no longer needed (stage 7), but they might still pay some attention to the road occasionally, while also starting to resume their original tasks (stage 9). That is, a full disengagement (stage 8) might also happen after stage 9.

2.8.2. Implications for safety, design, and future research

The end of human driving (support) actions can be initiated by different actors: the human, the car, or both human and car. These different situations put different demands on the human and the human-system interaction (see also Mirnig et al., 2017). As most interruption literature focuses on system-initiated endings of interruptions, more studies are needed on human-initiated endings. In addition, more knowledge is needed on effective communication of a transition of control from human to system: what is the best way to initiate it (human or system?), how is this communicated effectively (e.g., what modalities?), and how is the change of roles controlled and checked by the system? While the 'negotiated' and the 'mediated' interruption management processes from McFarlane's parlance (McFarlane, 2002) are typically considered in traditional interruption management scenarios, it may be interesting to see how the same theory applies in the context of transfer of control from a human back to the automated vehicle.

2.9. Stage 8: disengage from traffic scenario

2.9.1. Theory

Even after the human stopped contributing to the drive, it might take a while before they fully disengage from the traffic scene. How do they know that the automation is in safe control of the car again? This seems to be a particular issue for system-initiated transitions. If the human did not yet feel comfortable with this 'forced' transition, they might continue to monitor the system for a longer time. Although this can have safety benefits (e.g., maintaining situational awareness), it also comes with challenges. For one, a longer interruption makes later resumption of the original non-driving task harder due to decay of memory (Borst et al., 2015). Similarly, mode confusion might arise if the human missed the communication of the car that control transitioned to the car again (e.g., they might falsely believe they are still in control; see also Janssen et al., 2019).

We already noted for stage 7 that system-initiated transitions are the most well studied scenario in the interruption literature. Unfortunately, despite this wealth of knowledge, there are still gaps in our understanding, as in the typical interruption experiment there is no chance for the human driver to *return* to the interrupting task once it has been closed. This contrasts with the car situation, in which the human driver's input might not be required for the car, but nonetheless the driving environment still provides input (e.g., the car is still driving, and there is still visual input from the road). As a result, there might be interleaving of attention.

In a case where the human handed over control back to the car, it is less clear what happens during the disengagement stage. A likely factor to contribute to how quickly the human disengages from the traffic scenario is the belief and trust that the human places in the system's reliability. The study of trust is a field in itself, and there is not yet clear convergence on how to communicate the appropriate level of trust to

the human driver (e.g., Noah et al., 2017).

In our overview of interruption stages (Fig. 1), we placed the disengagement from traffic scenario stage (stage 8) before the stage to orient to the original task (stage 9). However, the timing of these stages might also be reversed if there is interleaving of attention between the road and some other task. That is, before the human driver fully disengages from traffic (stage 8), they might already orient towards the original task (stage 9). This differs from more classical interruption studies, where an original task could only be resumed after an interrupting task has been ended.

2.9.2. Implications for safety, design, and future research

The degree of disengagement from a driving system will depend on the initiator of the transition of control (system, human, or both) and the degree of belief and trust that the human driver has in the automation. From a safety perspective, if a human places too much reliance in a system, this can be particularly harmful. For example, a situation of overtrust (Lee and See, 2004) might be when, after driving through a snowy road 'by hand', a human driver transfers control to the car on a road that is still slippery and with occasional objects marking/blocking the road, that the car cannot handle.

The main implication for design is that a transition of control from the human to the vehicle should be clearly communicated by the system, independent of who initiated this transition, to avoid mode confusion. However, it is less clear what the ideal ways of communication are. Both overtrust and undertrust in the system need to be minimized, so the human driver knows when they can safely work on non-driving activities.

To further understand trust and handling of uncertainty in automated systems, insights are needed on system-design and human attention. Engineering can provide insights on the capabilities of the car, and express the uncertainty in the actions. Human attention studies can gain insight in relevant questions such as: What is the human driver's understanding of the system's capabilities? What type of information makes them believe the system can take over control or not? What is the most effective and least disruptive way to communicate this to them?

We highlighted that it is unclear whether the disengagement stage (stage 8) always comes before the stage of orienting to the original task (stage 9; Boehm-Davis and Remington, 2009) for automated driving scenarios. For automated driving in particular, there is an interesting conundrum in that the human might no longer be in control of the vehicle (stage 7), but still attend the road (between stage 7 and 8). How does this ability to observe, but not act on the car's course affect human behavior? Might there be cases where a human driver has handed back control to the car, to quickly determine (after stage 7, but before stage 10) to take-over control *again*? And how would that affect performance? These new questions now require answers.

2.10. Stage 9: orient to original non-driving task

2.10.1. Theory

When a driver wants to resume their original non-driving task, they might first orient themselves to the original task to see where they can resume. Cues from the task environment can aid resumption, particularly if the interruption has been too long to rely on memories from previous task steps (cf. Altmann and Trafton, 2002; Borst et al., 2015). Cues can act as place-keepers (Gray, 2000) and draw the person's attention to the relevant context. The context can create associations with past memories (Anderson, 1983), and thereby further aid reconstruction of the task context (see also Salvucci and Taatgen, 2011, Chapter 4; Borst et al., 2013).

Although shorter resumption times are associated with fewer errors in general, this perspective has recently been challenged (Brumby et al.,

2013). Brumby and colleagues found that some forced delay in task resumption (i.e., a slower resumption) *reduced* the number of errors, instead of finding the theoretically expected increase in number of errors. One explanation might be that the delay forced participants to better remember where they were in the task before moving to an interrupting task (cf. Ballard et al., 1997; Gray et al., 2006; Gray and Fu, 2004, see also Janssen and Gray, 2012), another might be that this leaves more time for task reconstruction (cf. Salvucci and Taatgen, 2011). That said, forced delays in task resumption might also encourage people to forego the original task altogether and move to another (third) task (Gould et al., 2015).

2.10.2. Implications for safety, design, and future research

Safety is hindered if task resumption is not a smooth process, as it might encourage drivers to keep paying attention to the non-driving task while they are supposed to be focusing on driving. Cues can be beneficial in two ways to make resumption more smooth. First, they provide a concrete starting point from which (memory) associations about task-relevant factors can start. Second, if the person knows that they can rely on cues for task resumption, they might feel less need to continue working on it or to keep information about the task in memory. This can reduce workload (Bailey and Iqbal, 2008) and stress (Adamczyk and Bailey, 2004). Future design work can consider under what circumstances cues are provided and when, so as to aid, but not distract, the user.

Out of all the stages of the interruption process, the task resumption stage is best understood. One open area is what causes the apparent speed-accuracy trade-off in task resumption: why is sometimes a fast, and sometimes a slow resumption beneficial in terms of error reduction? If a smooth resumption is helped by some forced delay in the resumption, how can this be done in a way that avoids switching to other tasks (Gould et al., 2015), and without taxing workload during the drive (cf. Bailey and Iqbal, 2008)?

2.11. Stage 10: resume original non-driving task

Once a human driver has resumed their original task, it is a similar situation to stage 0. The driver might again focus mostly on non-driving tasks, and it can not be assumed that they are focusing on the traffic scenario.

3. General discussion

We adapted a general framework of interruptions (modified from Boehm-Davis and Remington, 2009) to capture the process of transfer of control in automated driving. Contrary to an implicit common assumption in current automated driving research, our framework makes explicit that transfer of control is not a single step, but instead a series of multiple stages, as illustrated in Fig. 1. For each stage, we highlighted relevant theory and important directions for safety and design. Moreover, the identification of the stages allowed the identification of important new directions of research for each stage. Consistent with the interruption literature, we only looked at situations in which two tasks are being completed. However, there might be cases where more tasks are being juggled (e.g., taking over driving while also holding a conversation and monitoring the navigation system).

3.1. Claims and points of reflection

From our analysis, five overarching claims and points of reflection emerge that go beyond the more detailed descriptions of each stage:

1 Transfer of control in automated driving can be considered as going

through a series of multiple stages (see Fig. 1). It is not accomplished in a single stage, or a small set of stages, even though this is the focus of most research to date (e.g., measurements of the speed of transfer of control at a single moment). Although we have explicitly separated all the stages here, in some situations, stages might co-occur. For example, if a driver has been paying attention to the road before an external alert (i.e., before stage 1), because they were working on an 'easy' task (e.g., listening to music), the disengagement and suspension stages (stage 2 and 4) might co-occur (or happen rapidly after each other), and the physical handover (stage 5) might occur without the need of much additional orientation (stage 3).

2 Consideration of transfer of control through the lens of interruptions is necessary, given expectations in the field that a transfer of control request might become less frequent and less urgent as vehicle automation increases. Humans might, as a result, perform other, non-driving related activities more frequently and for longer periods of time.

3 Labeling driving and other activities as primary/secondary might not be appropriate. Instead, we encourage the use of more neutral and factual terms such as "driving", "non-driving", "time-sensitive", and "safety-critical" tasks and activities. Note that in some cases non-driving related activities might also contain time-sensitive characteristics.

4 People's attention division should not be seen as being dedicated fully to one or another task (e.g., to either driving or non-driving activities). That is, unlike in the classical interruption framework (where people are thought to work on one task at a time), within settings of automated driving there might be multiple periods in which attention is *interleaved* between multiple activities, including driving and non-driving activities.

5 Expectations that human drivers immediately take over control when requested by the vehicle might be unrealistic for at least three reasons. First, research shows that in early stages of the transfer (i.e., stages 2 to 5), drivers might not immediately direct their attention to the drive (i.e., they might interleave) and they might not have sufficient awareness of their environment to act appropriately. Second, research has shown that earlier tasks might negatively impact later tasks (such as driving) even if those earlier tasks were discontinued (stage 6). Third, there is no empirical evidence that people indeed fully disengage consistently from other activities when taking over control of the vehicle, whereas there is evidence that they engage with other non-driving activities under regular driving conditions.

Although the aim of some levels of automation (e.g., SAE level 5) is to completely surpass human involvement in the driving, human involvement is expected for all the other levels and for systems that are commercially available now and in the years to come. The interruption framework that we presented is particularly relevant for SAE levels 3 and 4, where the car is assumed to take over control of the driving for prolonged periods of time in specific operational design domains (e.g., a highway under regular traffic circumstances), and where the human might work on other tasks in the meantime. The framework might also be used in some cases of the highest level of automation, where some human input might occasionally be requested (e.g., to confirm a change of route), and interrupt a user from other non-driving activities they were doing. To make automated driving systems successful, more research is needed about each of the stages of the interruption framework in the context of automated vehicles. We therefore identified multiple important directions for future work.

There are many other settings in which humans interact with automation. In particular, there is an increase in research on automated systems that are used in safety critical settings (e.g., medicine, aviation) or that are used by non-professional users (e.g., use of automation on phones and devices), see review in Janssen et al. (2019). Such settings might also benefit from application of our framework for attention management. Future research can look into what aspects and what processes are consistent or different across domains.

3.2. Are all the stages needed?

Description of all 10 stages of the interruption process in the context of transfer of control allows for an accurate, detailed description of the transition of control process. However, it is an open question whether all stages and constructs are involved in all situations of transfer of control. We already highlighted that stage 8 might for example not preceed stage 9, as drivers might already orient to their previously original task before disengaging from the driving task.

Similarly, there might be situations where stages are completed rapidly or even skipped. For example, in cases where a driver self-interrupts to take over control of the driving task, they might disengage from their original task and immediately suspend the original task (i.e., collapsing stages 2 and 4). In such a scenario, the driver will have had reasons to act and might need little or no time to orient to the driving task (i.e., dropping or briefly going through stage 3, see also scenario 3 in Textbox 1), and quickly taking physical control (stage 5). Further research is needed to understand under what conditions specific phases might be shortened or omitted. With the current insights from the literature, we adhered to using the extensive model, as the utility of this framework is that it is more consistent with the original interruptions literature, and allows for explicit discussion of which stages were left out in any specific application of this framework to a concrete setting. Thereby, situations where phases are dropped might prove to be the exception rather than the default.

A next question is then how to detect transitions through the stages. In our example scenarios (Textbox 1), we provided a couple of examples of how this can be done in experiments. For example, eye-tracking might be used to detect whether one is looking at a particular task (e.g., looking at the road versus looking at one's phone) and interaction data might be used to detect whether one is actively acting on that task (e.g., steering wheel movements, button presses on a phone). Multiple measures might be needed to detect the transition between stages. Moreover, during the two interleaving phases, a driver is going back and forth between the two tasks, and researchers might want to distinguish the original moment at which a stage started (e.g., the first time one orients to the driving task, stage 3) and moment where one (after interleaving) continues such a stage (e.g., continues to orient to the driving task to build further context). Our example scenarios (Textbox 1) provide various examples.

4. Conclusion

We present a framework that considers transitions of control by automation through the lens of interruption and interleaving processes. The framework makes explicit that it is not appropriate to think of transfer of control as a single event or even as a small set of events. It also highlights that it might not be realistic to expect human drivers to immediately respond correctly to a system initiated request to transfer control, given that humans interleave their attention between non-driving and driving tasks, and given that a transition constitutes of multiple stages. These nuances are accounted for in the framework. The

framework opens up multiple directions in the design and evaluation of the process of transfers of control.

Declarations of interest

None.

Acknowledgements

References

Adamczyk, P.D., Bailey, B.P., 2004. If not now, when? In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. New York, NY. ACM Press, pp. 271–278.

Altmann, E., Trafton, J.G., 2002. Memory for goals: an activation-based model. Cogn. Sci. 26 (1), 39–83.

Anderson, J.R., 1983. A spreading activation theory of memory. J. Verbal Learn. Verbal Behav. 22 (3), 261–295.

Atchley, P., Chan, M., 2011. Potential benefits and costs of concurrent task engagement to maintain vigilance: a driving simulator investigation. Hum. Factors 53 (1), 3–12. http://doi.org/10.1177/0018720810391215.

Bailey, B.P., Iqbal, S.T., 2008. Understanding changes in mental workload during execution of goal-directed tasks and its application for interruption management. ACM Trans. Comput.-Hum. Interact. 14 (4), 1–28.

Baldwin, C.L., Lewis, B.A., 2014. Perceived urgency mapping across modalities within a driving context. Appl. Ergon. 45 (5), 1270–1277.

Ballard, D.H., Hayhoe, M.M., Pook, P., Rao, R., 1997. Deictic codes for the embodiment of cognition. Behav. Brain Sci. 20 (04), 723–742.

Bliss, J.P., Acton, S.A., 2003. Alarm mistrust in automobiles: how collision alarm reliability affects driving. Appl. Ergon. 34 (6), 499–509. http://doi.org/10.1016/j.apergo.2003.07.003.

Boehm-Davis, D.A., Remington, R.W., 2009. Reducing the disruptive effects of interruption: a cognitive framework for analysing the costs and benefits of intervention strategies. Accid. Anal. Prevent. 41 (5), 1124–1129. http://doi.org/10.1016/j.aap.2009.06.029.

Bogunovich, P., Salvucci, D.D., 2010. Inferring multitasking breakpoints from single-task data. In: Proceedings of the 32nd Annual Meeting of the Cognitive Science Society, pp. 1732–1737.

Borojeni, S.S., Weber, L., Heuten, W., Boll, S., 2018. From reading to driving – priming mobile users for take-over situations in highly automated. In: Proceedings of the 20th International Conference on Human-Computer Interaction with Mobile Devices and Services.

Borst, J.P., Buwalda, T.A., van Rijn, H., Taatgen, N.A., 2013. Avoiding the problem state bottleneck by strategic use of the environment. Acta Psychol. 144 (2), 373–379.

Borst, J.P., Taatgen, N.A., van Rijn, H., 2015. What makes interruptions disruptive?: a process-model account of the effects of the problem state bottleneck on task interruption and resumption. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. New York, NY. ACM Press, pp. 2971–2980.

Breznitz, S., 1983. Crywolf: The psychology of False Alarms. Lawrence Erlbaum, Hillsdale, NJ.

Brumby, D.P., Cox, A.L., Back, J., Gould, S.J., 2013. Recovering from an interruption: investigating speed − accuracy trade-offs in task resumption behavior. J. Exp. Psychol. 19 (2), 95.

Brumby, D.P., Janssen, C.P., Mark, G., 2019. How do interruptions affect Productivity? In: Sadowski, C., Zimmermann, T. (Eds.), Rethinking Productivity in Software Engineering. Apress, pp. 85–110 Chapter 9.

Brumby, D.P., Salvucci, D.D., Howes, A., 2009. Focus on driving: how cognitive constraints shape the adaptation of strategy when dialing while driving. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. New York, NY. pp. 1629–1638.

Caird, J.K., Johnston, K.A., Willness, C.R., Asbridge, M., 2014. A meta-analysis of the effects of texting on driving. Accid. Anal. Prevent. 71, 311–318. http://doi.org/10.1016/j.aap.2014.06.005.

Caird, J.K., Simmons, S.M., Wiley, K., Johnston, K.A., Horrey, W.J., 2018. Does talking on a cell phone, with a passenger, or dialing affect driving performance? An updated systematic review and meta-analysis of experimental studies. Hum. Factors 60 (1), 101–133.

Caird, J.K., Willness, C.R., Steel, P., Scialfa, C., 2008. A meta-analysis of the effects of cell phones on driver performance. Accid. Anal. Prevent. 40 (4), 1282–1293. http://doi.

org/10.1016/j.aap.2008.01.009.

Card, S.K., Moran, T., Newell, A., 1983. The Psychology of Human-Computer Interaction. Lawrence Erlbaum Associates, Hillsdale, NJ.

Chen, H.Y.W., Milgram, P., 2013. A framework for modelling and analysing variability in visual occlusion experiments. In: Proceedings of the Human Factors and Ergonomics Society Annual Meeting. Los Angeles, CA. SAGE Publications, pp. 1884–1888.

Couffe, C.L., Michael, G.A., 2017. Failures due to interruptions or distractions: a review and a new framework. Am. J. Psychol. 130 (2), 163–181. http://doi.org/10.5406/amerjpsyc.130.2.0163.

Dabbish, L., Mark, G.J., González, V.M., 2011. Why do I keep interrupting myself?: Environment, habit and self-interruption. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. New York, NY. ACM Press, pp. 3127–3130. http://doi.org/10.1145/1978942.1979405.

de Winter, J., Happee, R., Martens, M.H., Stanton, N.A., 2014. Effects of adaptive cruise control and highly automated driving on workload and situation awareness: a review of the empirical evidence. Transp. Res. Part F 27, 196–217. http://doi.org/10.1016/j.trf.2014.06.016.

Dingus, T.A., Guo, F., Lee, S., Antin, J.F., Perez, M., Buchanan-King, M., Hankey, J., 2016. Driver crash risk factors and prevalence evaluation using naturalistic driving data. In: Proceedings of the National Academy of Sciences, , 201513271. http://doi.org/10.1073/pnas.1513271113.

Dogan, E., Rahal, M.C., Deborne, R., Delhomme, P., Kemeny, A., Perrin, J., 2017. Transition of control in a partially automated vehicle: effects of anticipation and non-driving-related task involvement. Transp. Res. Part F 46, 205–215.

Endsley, M.R., Garland, D.J. (Eds.), 2000. Situation Awareness Analysis and Measurement. Lawrence Erlbaum Associates, London.

Engström, J., Markkula, G., Victor, T.W., Merat, N., 2017. Effects of cognitive load on driving performance: the cognitive control hypothesis. Hum. Factors 59 (5), 734–764.

Gold, C., Damböck, D., Lorenz, L., Bengler, K., 2013. "Take over!" How long does it take to get the driver back into the loop? In: Proceedings of the Human Factors and Ergonomics Society Annual Meeting. 57. pp. 1938–1942. http://doi.org/10.1177/1541931213571433.

González, V.M., Mark, G.J., 2004. "Constant, constant, multi-tasking craziness": managing multiple working spheres. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. New York, NY. ACM Press, pp. 113–120.

Gould, S.J., Cox, A.L., Brumby, D.P., 2015. Task lockouts induce crowdworkers to switch to other activities. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. New York, NY. ACM Press, pp. 1785–1790.

Gould, S.J., Cox, A.L., Brumby, D.P., González, V.M., Salvucci, D.D., Taatgen, N.A., 2012. Multitasking and Interruptions: a SIG on bridging the gap between research on the micro and macro worlds. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems Extended Abstracts. New York, NY. ACM Press, pp. 1189–1192.

Gray, W.D., 2000. The nature and processing of errors in interactive behavior. Cogn. Sci. 24 (2), 205–248.

Gray, W.D., Fu, W., 2004. Soft constraints in interactive behavior: the case of ignoring perfect knowledge in-the-world for imperfect knowledge in-the-head. Cogn. Sci. 28 (3), 359–382.

Gray, W.D., Sims, C.R., Fu, W.-T., Schoelles, M.J., 2006. The soft constraints hypothesis: a rational analysis approach to resource allocation for interactive behavior. Psychol. Rev. 113 (3), 461–482. http://doi.org/10.1037/0033-295X.113.3.461.

Gugerty, L.J., 1997. Situation awareness during driving: explicit and implicit knowledge in dynamic spatial memory. J. Exp. Psychol. 3 (1), 42–66.

Hancock, P.A., 2013. Driven to distraction and back again. In: Regan, M.A., Lee, J.D., Victor, T.W. (Eds.), Driver Distraction and Inattention Advances in Research and Countermeasures. Ashgate Publishing limited, Surrey, UK, pp. 9–25.

Horrey, W.J., Lesch, M.F., Garabet, A., 2009. Dissociation between driving performance and drivers' subjective estimates of performance and workload in dual-task conditions. J. Safety Res. 40 (1), 7–12. http://doi.org/10.1016/j.jsr.2008.10.011.

Howes, A., Lewis, R.L., Vera, A., 2009. Rational adaptation under task and processing constraints: implications for testing theories of cognition and action. Psychol. Rev. 116 (4), 717–751.

Inners, M., Kun, A.L., 2017. Beyond liability: legal issues of human-machine interaction for automated vehicles. In: Proceedings of the 9th International Conference on Automotive User Interfaces and Interactive Vehicular Applications. New York, NY. ACM Press, pp. 245–253.

Iqbal, S.T., Adamczyk, P.D., Zheng, X.S., Bailey, B.P., 2005. Towards an index of opportunity: understanding changes in mental workload during task execution. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. New York, NY. ACM Press, pp. 311–320.

Iqbal, S.T., Ju, Y.-C., Horvitz, E., 2010. Cars, calls, and cognition: investigating driving and divided attention. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. New York, NY. ACM Press, pp. 1281–1290.

James, W, 1890. The Principles of Psychology. Holt, New York.

Janssen, C.P., Boyle, L., Kun, A., Ju, W., Chuang, L., 2019a. A hidden Markov framework to capture human-machine interaction in automated vehicles. Int. J. Hum. Comput. Interact. 35 (11), 947–955. https://doi.org/10.1080/10447318.2018.1561789.

Janssen, C.P., Brumby, D.P., 2010. Strategic adaptation to performance objectives in a dual-task setting. Cogn. Sci. 34 (8), 1548–1560. http://doi.org/10.1111/j.1551-6709.2010.01124.x.

Janssen, C.P., Brumby, D.P., Garnett, R., 2012. Natural break points the influence of

priorities and cognitive and motor cues on dual-task interleaving. J. Cogn. Eng. Decis. Mak. 6 (1), 5–29. http://doi.org/10.1177/1555343411432339.

Janssen, C.P., Donker, S.F., Brumby, D.P., Kun, A.L., 2019b. History and future of human-automation interaction. Int. J. Hum. Comput. Stud. http://doi.org/10.1016/j.ijhcs.2019.05.006.

Janssen, C.P., Everaert, E., Hendriksen, H.M.A., Mensing, G.L., Tigchelaar, L.J., Nunner, H., 2019c. The influence of rewards on (sub-)optimal interleaving. PLoS ONE 14 (3), e0214027. https://doi.org/10.1371/journal.pone.0214027.

Janssen, C.P., Gould, S.J., Li, S.Y.W., Brumby, D.P., Cox, A.L., 2015. Integrating knowledge of multitasking and interruptions across different perspectives and research methods. Int. J. Hum. Comput. Stud. 79, 1–5. http://doi.org/10.1016/j.ijhcs.2015.03.002.

Janssen, C.P., Gray, W.D., 2012. When, what, and how much to reward in reinforcement learning-based models of cognition. Cogn Sci 36 (2), 333–358. http://doi.org/10.1111/j.1551-6709.2011.01222.x.

Kahneman, D., 1973. Attention and Effort. Prentice-Hall, Englewood Cliffs, NJ.

Kass, S.J., Cole, K.S., Stanny, C.J., 2007. Effects of distraction and experience on situation awareness and simulated driving. Transp. Res. Part F 10 (4), 321–329. http://doi.org/10.1016/j.trf.2006.12.002.

Klauer, S.G., Guo, F., Simons-Morton, B.G., Ouimet, M.C., Lee, S.E., Dingus, T.A., 2014. Distracted driving and risk of road crashes among novice and experienced drivers. New Engl. J. Med. 370 (1), 54–59. http://doi.org/10.1056/NEJMsa1204142.

Kujala, T., Mäkelä, J., Kotilainen, I., Tokkonen, T., 2016. The attentional demand of automobile driving revisited: occlusion distance as a function of task-relevant event density in realistic driving scenarios. Hum. Factors 58 (1), 163–180. http://doi.org/10.1177/0018720815595901.

Kun, A.L., 2018. Human-machine interaction for vehicles: review and outlook. Found. Trends Hum.–Comput. Interact. 11 (4), 201–293. http://dx.doi.org/10.1561/1100000069.

Kun, A.L., Boll, S., Schmidt, A., 2016. Shifting gears: user interfaces in the age of autonomous vehicles. IEEE Pervasive Comput. 15 (1), 32–38.

Kun, A.L., Shyrokov, A., Heeman, P.A., 2013. Interactions between human–human multithreaded dialogues and driving. Pers. Ubiquitous Comput. 17 (5), 825–834.

Kunar, M.A., Carter, R., Cohen, M., Horowitz, T.S., 2008. Telephone conversation impairs sustained visual attention via a central bottleneck. Psychon. Bull. Rev. 15 (6), 1135–1140. http://doi.org/10.3758/PBR.15.6.1135.

Larsson, P., Johansson, E., Söderman, M., Thompson, D., 2015. Interaction design for communicating system state and capabilities during automated highway driving. Procedia Manuf. 3, 2784–2791.

Lee, J.D., See, K.A., 2004. Trust in automation: designing for appropriate reliance. Hum. Factors 46 (1), 50–80. http://doi.org/10.1518/hfes.46.1.50_30392.

Lewin, K., 1943. Defining the "Field at a given time" Psychol. Rev. 50 (3), 292–310.

Li, S.Y.W., Blandford, A., Cairns, P., Young, R.M., 2008. The effect of interruptions on postcompletion and other procedural errors: an account based on the activation-based goal memory model. J. Exp. Psychol. 14 (4), 314–328.

Li, S.Y.W., Magrabi, F., Coiera, E., 2011. A systematic review of the psychological literature on interruption and its patient safety implications. J. Am. Med. Inf. Assoc. http://doi.org/10.1136/amiajnl-2010-000024.

Mark, G., Iqbal, S.T., Czerwinski, M., Johns, P., 2015. Focused, aroused, but so distractible: temporal perspectives on multitasking and communications. In: Proceedings of CSCW. New York, NY. ACM Press, pp. 903–916.

McFarlane, D.C., 2002. Comparison of four primary methods for coordinating the interruption of people in human-computer interaction. Hum.-Comput. Interact. 17 (1), 63–139.

McFarlane, D.C., Latorella, K.A., 2002. The scope and importance of human interruption in human-computer interaction design. Hum.-Comput. Interact. 17 (1), 1–61.

Meng, F., Spence, C., 2015. Tactile warning signals for in-vehicle systems. Accid. Anal. Prev. 75, 333–346.

Merat, N., Jamson, A.H., Lai, F., Daly, M., 2014. Transition to manual: driver behaviour when resuming control from a highly automated vehicle. Transp. Res. Part F 27, 274–282. http://doi.org/10.1016/j.trf.2014.09.005.

Mirnig, A.G., Gärtner, M., Laminger, A., Meschtscherjakov, A., Trösterer, S., Tscheligi, M., McCall, R., McGee, F., 2017. Control transition interfaces in semiautonomous vehicles: a categorization framework and literature analysis. In: Proceedings of the 9th International Conference on Automotive User Interfaces and Interactive Vehicular Applications. New York, NY. ACM Press, pp. 209–220.

Miyata, Y., Norman, D.A., 1986. Psychological issues in support of multiple activities. In: Norman, D.A., Draper, S. (Eds.), User Centered System design: New perspectives On Human-Computer Interaction. Erlbaum, Hillsdale, NJ, pp. 265–284.

Mok, B., Johns, M., Lee, K.J., Miller, D., Sirkin, D., Ive, P., Ju, W., 2015. Emergency, automation off: unstructured transition timing for distracted drivers of automated vehicles. In: Proceedings of the IEEE International Conference on Intelligent Transportation Systems. Gran Canaria. IEEE, pp. 2458–2464.

Mok, B., Johns, M., Miller, D., Ju, W., 2017. Drivers with active secondary tasks need more time to transition from automation. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. New York, NY. ACM Press, pp. 2840–2844. http://doi.org/10.1145/3025453.3025713.

Naujoks, F., Wiedemann, K., Schömig, N., 2017. The importance of interruption management for usefulness and acceptance of automated driving. In: Proceedings of the International Conference on Automotive User Interfaces and Interactive Vehicular Applications. New York, NY. ACM Press, pp. 254–263.

Neerincx, M., 2003. Cognitive modelling of pilot errors and error recovery in flight management tasks. In: Hollnagel, E. (Ed.), Handbook of Cognitive Task Design. CRC, pp. 283–306.

Noah, B.E., Wintersberger, P., Mirnig, A.G., Thakkar, S., Yan, F., Gable, T.M., Kraus, J., McCall, R., 2017. First workshop on trust in the age of automated driving. In: Proceedings of the 9th International Conference on Automotive User Interfaces and Interactive Vehicular Applications Adjunct. New York, NY. ACM Press, pp. 15–21.

Norman, D.A., Shallice, T., 1986. Attention to action. Consciousness and Self-Regulation. Springer, Boston, MA, pp. 1–18.

Payne, S.J., Duggan, G.B., Neth, H., 2007. Discretionary task interleaving: heuristics for time allocation in cognitive foraging. J. Exp. Psychol. 136 (3), 370–388. http://doi.org/10.1037/0096-3445.136.3.370.

Petermeijer, S., Bazilinskyy, P., Bengler, K., De Winter, J., 2017. Take-over again: investigating multimodal and directional TORs to get the driver back into the loop. Appl. Ergon. 62, 204–215.

Pfleging, B., Rang, M., Broy, N., 2016. Investigating user needs for non-driving-related activities during automated driving. In: Proceedings of the 15th International Conference on Mobile and Ubiquitous Multimedia. Rovaniemi, Finland. ACM, pp. 91–99. http://dx.doi.org/10.1145/3012709.3012735.

Rivera-Rodriguez, A.J., Karsh, B.T., 2010. Interruptions and distractions in healthcare: review and reappraisal. BMJ Qual. Saf. 19 (4), 304–312. http://doi.org/10.1136/qshc.2009.033282.

SAE International. (2014). *J3016: taxonomy and definitions for terms related to on-road motor vehicle automated driving systems*. Retrieved from http://standards.sae.org/j3016_201401/.

Salvucci, D.D., 2005. A multitasking general executive for compound continuous tasks. Cogn. Sci. 29 (3), 457–492. http://doi.org/10.1207/s15516709cog0000_19.

Salvucci, D.D., Beltowska, J., 2008. Effects of memory rehearsal on driver performance: experiment and theoretical account. Hum. Factors 50 (5), 834–844.

Salvucci, D.D., Taatgen, N.A., 2008. Threaded cognition: an integrated theory of concurrent multitasking. Psychol. Rev. 115 (1), 101–130. http://doi.org/10.1037/0033-295X.115.1.101.

Salvucci, D.D., Taatgen, N.A., 2011. The Multitasking Mind. Oxford University Press, New York, NY.

Salvucci, D.D., Taatgen, N.A., Borst, J.P., 2009. Toward a unified theory of the multitasking continuum: from concurrent performance to task switching, interruption, and resumption. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. New York, NY. ACM Press, pp. 1819–1828.

Sanderson, P.M., Grundgeiger, T., 2015. How do interruptions affect clinician performance? Negotiating fidelity, control, and potential generalizability in the search for answers. Int. J. Hum. Comput. Stud. 79, 85–96.

Schreuder, E.J., Mioch, T., 2011. The effect of time pressure and task completion on the occurrence of cognitive lockup. In: *Proceedings of the International Workshop on Human Centered Processes*. 2011. pp. 10–11.

Senders, J., Kristofferson, A., Levison, W., Dietricht, C., Ward, J., 1967. The attentional demand of automobile driving. Highway Res. Rec. 195, 15–33.

Sorkin, R.D., 1989. Why are people turning off our alarms? Hum. Factors Bull. 32 (4), 3–4.

Strayer, D.L., Johnston, W.A., 2001. Driven to distraction: dual-task studies of simulated driving and conversing on a cellular telephone. Psychol. Sci. 12 (6), 462–466.

Strayer, D.L., Cooper, J.M., Turrill, J., Coleman, J.R., Hopman, R.J., 2015. Measuring cognitive distraction in the automobile III: a comparison of ten 2015 in-vehicle information systems. AAA Found. Traffic Safety 1–46.

Telford, C.W., 1931. The refractory phase of voluntary and associative responses. J. Exp. Psychol. 36 (1), 1–36.

van der Heiden, R.M.A., Iqbal, S.T., Janssen, C.P., 2017. Priming drivers before handover in semi-autonomous cars. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. New York, NY. ACM Press, pp. 392–404.

van der Heiden, R.M.A., Janssen, C.P., Donker, S.F., Hardeman, L.E.S., Mans, K., Kenemans, J.L., 2018. Susceptibility to audio signals during autonomous driving. PLoS ONE 13 (8), e0201963. http://doi.org/10.1371/journal.pone.0201963.

van der Stigchel, 2019. How Attention Works. MIT Press, Boston, MA.

Walch, M., Lange, K., Baumann, M., Weber, M., 2015. Autonomous driving: investigating the feasibility of car-driver handover assistance. In: Proceedings of the ACM conference on Automotive User Interfaces. New York, NY. ACM, pp. 11–18. http://doi.org/10.1145/2799250.2799268.

Wester, A.E., Böcker, K., Volkerts, E.R., Verster, J.C., Kenemans, J.L., 2008. Event-related potentials and secondary task performance during simulated driving. Accid. Anal. Prevent. 40 (1), 1–7. http://doi.org/10.1016/j.aap.2007.02.014.

Wickens, C.D., 2008. Multiple resources and mental workload. Hum. Factors 50 (3), 449–455 http://doi.org/10.1518/001872008×288394.

Wickens, C.D., McCarley, J.S., 2008. Applied Attention Theory. CRC Press.

Wickens, C.D., Gutzwiller, R.S., Santamaria, A., 2015. Discrete task switching in overload: a meta-analyses and a model. Int. J. Hum. Comput. Stud. 79, 79–84. http://doi.org/10.1016/j.ijhcs.2015.01.002.

Wickens, C.D., Rice, S., Keller, D., Hutchins, S., Hughes, J., Clayton, K., 2009. False alerts in air traffic control conflict alerting system: is there a "cry wolf" effect? Hum. Factors 51 (4), 446–462.

Yang, F., Heeman, P.A., Kun, A.L., 2011. An investigation of interruptions and resumptions in multi-tasking dialogues. Comput. Linguist. 37 (1), 75–104.

Young, M.S., Stanton, N.A., 2002. Malleable attentional resources theory: a new explanation for the effects of mental underload on performance. Hum. Factors 44 (3), 365–375.

Christian P. Janssen is an assistant professor of experimental psychology at Utrecht University. He received his PhD in human-computer interaction from UCL (2012).

Andrew L. Kun is an associate professor of electrical and computer engineering at the University of New Hampshire. He received his PhD in Electrical Engineering from the University of New Hampshire (1997).

Shamsi T. Iqbal is a senior researcher at Microsoft Research AI. She obtained her PhD in human-computer interaction from the University of Illinois at Urbana-Champaign (2008).

Stella F. Donker is an associate professor of experimental psychology at Utrecht University. She received her PhD in movement sciences from the University of Groningen (2002).