

Haptic shared control: smoothly shifting control authority?

David A. Abbink · Mark Mulder · Erwin R. Boer

Received: 19 July 2011 / Accepted: 16 September 2011 / Published online: 5 November 2011
© The Author(s) 2011. This article is published with open access at Springerlink.com

Abstract Literature points to persistent issues in human-automation interaction, which are caused either when the human does not understand the automation or when the automation does not understand the human. Design guidelines for human-automation interaction aim to avoid such issues and commonly agree that the human should have continuous interaction and communication with the automation system and its authority level and should retain final authority. This paper argues that haptic shared control is a promising approach to meet the commonly voiced design guidelines for human-automation interaction, especially for automotive applications. The goal of the paper is to provide evidence for this statement, by discussing several realizations of haptic shared control found in literature. We show that literature provides ample experimental evidence that haptic shared control can lead to short-term performance benefits (e.g., faster and more accurate vehicle control; lower levels of control effort; reduced demand for visual attention). We conclude that although the continuous intuitive physical interaction inherent in haptic shared control is expected to reduce long-term issues with human-automation interaction, little experimental evidence for this is provided. Therefore, future research on haptic shared

control should focus more on issues related to long-term use such as trust, overreliance, dependency on the system, and retention of skills.

Keywords Automation · Human-machine interface · Haptic guidance · Shared control · Neuromuscular identification · Levels of automation · Force feedback

1 Introduction

Over the last decades, pushed and pulled by technological progress, our society has created more and increasingly complex machines in order to increase comfort, production, and safety for ourselves. Automation, i.e., giving partial or full authority to machines, has gone hand in hand with this trend, relieving us of the workload associated with controlling these machines. With the current level of technology, much is possible, but as Wiener and Curry (1980) already recognized, “...the question is no longer whether one or another function can be automated, but, rather, whether it should be.”

Our society agrees that some functions clearly should be automated: automation is widely accepted in many well-structured and predictable areas of our life. For example, nobody talks about the risks or human factors issues of washing machines or automatic assembly lines in the food industry. In these cases, the role of the human requires no more supervision than turning the machines on or off and sporadically monitoring whether the machine still works properly. Additionally, the worst-case scenario of a system failure is very unlikely to be life threatening, but would merely cause discomfort.

On the other hand, a high level of automation is not universally applicable: it can also have undesirable effects,

D. A. Abbink (✉) · M. Mulder
Department of BioMechanical Engineering, Faculty of 3mE,
Delft University of Technology, Mekelweg 2, 2628 CD Delft,
The Netherlands
e-mail: d.a.abbink@tudelft.nl

M. Mulder
Department of Control and Simulation, Faculty of Aerospace
Engineering, Delft University of Technology, Kluyverweg 1,
2629 HS Delft, The Netherlands

E. R. Boer
Entropy Control, San Diego, CA 92122, USA

especially in the control of more safety-critical dynamic processes in unpredictable environments. Literature on human factors in automation has widely reported on the disadvantages of inappropriate automation (Bainbridge 1983; Sheridan 2002; Sheridan and Parasuraman 2006). One of the more influential human factors' papers (Parasuraman and Riley 1997) argues that it is helpful to think of automation in terms of use (humans using automation to perform tasks otherwise performed manually), misuse (humans using automation to perform tasks it is not designed to handle), disuse (humans not using the automation where it could have been helpful), and abuse (when automation is implemented without sufficiently considering the effect on human operators). But as Lee (2008) states, a less broadly recognized theme of Parasuraman and Riley's paper (1997) cuts across these issues: a vicious cycle created by misuse and disuse of automation. Lee (2008) argues that when wrongly implemented automation (i.e., abuse) increases misuse and disuse of automation, it might in turn lead overzealous managers or designers to use more (or higher level) automation. For example, when pilots make errors due to loss of skills and overreliance on a non-perfect autopilot, regulations might be implemented to take away even more tasks, which would induce further loss of skills. This vicious cycle, a disbalance in control between the designer, manager, and the operator, to this day, is not yet broken, and many of the human-automation interaction issues described decades ago still persist (Parasuraman and Wickens 2008). How do we break this cycle? What design guidelines are available that should be followed more often in order improve the design of human-automation interaction, and thereby reduce automation abuse?

1.1 Design guidelines for human-automation interaction

Over two decades ago, Norman (1990) already put forward that appropriate human-automation interaction "...should assume the existence of error, it should continually provide feedback, it should continually interact with operators in an effective manner, and it should allow for the worst of situations." These guidelines appear difficult to meet in practice. Engineers tend to downplay the existence of automation errors. They find it hard to realize continuous feedback and continuous interaction because of annoyance or increased workload, and handing over control in the worst of situations is usually solved by an alert that tells the human operator that they have just been handed back full control authority over their vehicle or device.

Another important guideline is the concept of human-centered automation (Billings 1997). In broad terms, the concept echoes Norman's guidelines and states that the human must always be in control, must be actively involved

and adequately informed, and that humans and automation must understand each other's intent in complex systems.

If we rephrase the literature, the design guidelines associated with automation are twofold.

First, the human operator should be able to understand the automation system, fully and intuitively. To ensure this, continuous feedback about, and interaction with, the automation system is vital (Norman 1990; Billings 1997). If this guideline is not met, overreliance, complacency, loss of situation awareness, and/or confusion due to automation-induced surprises will occur (i.e., misuse of the systems), which will ultimately lead to distrust and disuse. Second, the automation system should include knowledge of the human operator. If this guideline is not met, the automation system may not match the goals, capabilities, and limitations of the human operator. This will result in automation abuse: wrong "engineering assumptions" about the human operators that will use the system. Better understanding of the human based on measurements and modeling (Parasuraman and Wickens 2008; Abbink and Mulder 2010) is expected to lead to less disagreement between automation system and human operator. But perhaps more important, in order to allow for the human to disagree, or to take over in the "worst of situations", the automation system should have the appropriate level of automation (LoA) (Sheridan 1992).

In order to satisfy the discussed automation design guidelines, the appropriate LoA needs to be determined, which is not a trivial task. The LoA might be designed to be constant (Endsley and Kiris 1995), but several studies have shown that the detection of automation failures is substantially degraded in automation, where the LoA remains fixed over time (e.g., Parasuraman and Riley 1997). Alternatively, the LoA could be variable, as was first proposed more than 30 years ago (Rouse 1976). This constitutes the concept of adaptive automation (for an overview, see: Inagaki 2003), which initially focused on system-driven adaptation (e.g., Scerbo 2001; Kaber and Endsley 2004). In other words, the automation system decides when to change its LoA. In order to remain aware, knowledge about the LoA should then also be continuously communicated (or at least be available) to the human operator. Such changes are most easily realized by binary switches of authority, from the point of view of the engineer. In other words, either the human is in control or the automation is in control of a particular task. But this may not be the most natural way of shifting or communicating the LoA (Flemisch et al. 2008). Indeed, "...what is needed is a soft, compliant technology, not a rigid, formal one" (Norman 1990). Still, adaptive automation introduces complexity during task allocation that can result in new issues with awareness of system functionality and automation failure detection (more automation abuse).

In contrast to automation-initiated changes in LoA, with adaptable automation, the human is the one to initiate changes in level of automation (Opperman 1994; Scerbo 2001). It is thought that “...adaptable automation can lead to benefits similar to those of adaptive automation while avoiding many of its pitfalls” (Miller and Parasuraman 2007). In particular, adaptable automation is thought to better involve the human operator in the task, thereby improving situation awareness and reducing skill degradation compared to adaptive automation or automation with fixed LoA. As a result, benefits in performance and workload are expected. Nevertheless, with adaptable automation, misuse and disuse are still likely to happen.

1.2 Problem statement and goal

In short, there is an ample body of literature that provides arguments and evidence for the need for human-automation interaction. Based on the discussed literature, we propose a set of four design guidelines for human-automation interaction: the human operator should

1. always remains in control, but can experience or initiate smooth shifts between levels of automation;
2. receive continuous feedback about the automation boundaries and functionality;
3. continuously interact with the automation; and
4. benefit from increased performance and/or reduced workload.

We believe that all four design guidelines mentioned above can be met by sharing control between the human and the automation on a physical level (i.e., through forces): haptic shared control. Abbink and Mulder (2010) define haptic shared control as a method of human-automation interaction that “...allows both the human and the [automation] to exert forces on a control interface, of which its output (its position) remains the direct input to the controlled system.” This implies that—depending on the direction and magnitude of the force that either human and automation exerts on the control interface—there can be a rich two-way interaction between human and automation, which will be elaborated upon in Sect. 2 below.

Haptic shared control has been investigated in areas such as vehicle control (e.g., Griffiths and Gillespie 2005; Forsyth and Maclean 2006; Goodrich et al. 2006; Abbink 2006; de Stigter et al. 2007; Mulder et al. 2008; Mulder et al. 2008; Abbink and Mulder 2009; Lam et al. 2009; Abbink and Mulder 2010; Mulder et al. 2010; Alaimo et al. 2010), robotic control (Rosenberg 1993; Marayong and Okamura 2004), and for learning and skill transfer (O’Malley et al. 2006). However, this approach has not (yet) received substantial attention in the automation and human factors domain. Also, in the haptics domain, only

few papers discuss the design and evaluation of haptic shared control in light of the known human factors issues with automation.

The goal of this paper is to bring the haptic world and the human-automation world closer together, and we will attempt to do this in two ways: on the one hand, by providing arguments and evidence that haptic shared control (alternatively dubbed continuous force feedback or haptic guidance) is a very promising human–machine interface for automation systems in the area of robotics and vehicular control; and on the other hand, by discussing how design choices in haptic shared control may prevent some of the issues with automation commonly reported in literature.

To reach these goals, Sect. 2 will provide the reader with a solid background on the benefits and limitations of different realizations of haptic shared control as a human–machine interface for dealing with automation. In Sect. 3, we will discuss some personal lessons regarding the design of haptic shared control based on two automotive case studies from our laboratory at the Delft University of Technology. The work will be discussed in light of other relevant publications in Sect. 4, which will also provide recommendations for future research.

2 Haptics for human-automation interaction

Essentially, sharing control through haptics implies that the human operator experiences additional forces on the control interface (e.g., joystick or steering wheel) that is used for controlling a system (e.g., robotic device or vehicle). One of the approaches is to generate additional forces as virtual fixtures (Rosenberg 1993), which constitute repulsive feedback forces used to protect forbidden regions in, for example, robotic surgery (Marayong and Okamura 2004) or UAV control (Lam et al. 2009). Essentially, this approach defines the boundaries within which operators can maneuver their system, and the closer the operators get to these boundaries, the higher the repulsive forces become. In other words, virtual fixtures *push* human operators away from pre-defined operational boundaries.

Alternatively, forces can be designed to guide the human along some sort of optimal trajectory, which is an approach often taken in automotive applications (Griffiths and Gillespie 2005; Forsyth and MacLean 2006; Mulder et al. 2008; Flemisch et al. 2008) and in aviation applications (Goodrich et al. 2008). With this approach, operators are not experiencing forces that push them away from pre-defined boundaries; instead, the forces they experience try to *pull* them back to the optimal trajectory when they deviate from it.

For either approach, additional forces are presented on top of any inherent control interface forces, such as

friction, spring stiffness, or breakout forces. The additional forces from the haptic shared controller essentially communicate that the current human control input (i.e., the force needed to deflect the control interface such that it will reach a certain position) will yield a different control interface position than what the automation system deems optimal. Essentially, this enables continuous feedback about—and continuous interaction with—the automation system (design guidelines #2 and #3).

Haptic shared control also allows for an interesting approach to designing shifts in control authority that can be realized through physical interaction; for the same error (e.g., the difference between current steering angle and the automation's proposed optimal steering angle), the additional guiding forces can be very large (yielding the equivalent of full automation) or entirely absent (yielding full manual control). The power of haptic shared control lies in the fact that *any level* in between these two extremes may be realized. This will be elaborated upon in the next section.

In response to these forces, the human operator can choose to intuitively react not only cognitively but also on a neuromuscular level (Abbink and Mulder 2010). Human operators can greatly vary their response to guiding forces (Abbink 2006; Abbink et al. 2011), for example, by ignoring them (staying relaxed), by resisting them (through co-contraction and reflexive activity), or by amplifying them (through actively giving way and reflexive activity). When the additional forces are designed to not exceed the maximal forces humans can generate, the human operator will always remain in control (design requirement #1).

So, theoretically, haptic shared control seems to be able to meet the first three design guidelines for human-automation interaction that were presented in the introduction. But do realizations of haptic shared control systems actually provide the human operator with benefits (guideline #4) while avoiding creating new pitfalls in automation? Several examples from literature are discussed here in an attempt to answer this question.

2.1 Haptic shared control with fixed authority

Rosenberg's work on virtual fixtures (Rosenberg 1993) might very well be the first instance of haptic shared control. The first realization he proposed was to simulate "*rigid planar surfaces*" that fully prevent the operator from venturing beyond the fixture—essentially full automation to restrict movements. But Rosenberg noted that it is also possible to "...consider modeling compliant surfaces [...] or even attractive or repulsive fields." Rosenberg tested several realizations with fixed levels of automation and concluded that an impedance surface increased operator performance by up to 70%.

Griffiths and Gillespie (2005) extended the concept of "...virtual fixtures [that] are usually fixed in the shared workspace..." by making dynamic virtual fixtures that are "...animated by the automation system." They described the general working of such haptic shared control, in their case for steering an automobile: "...a steering wheel can be given a 'home' position that is itself animated according to sensed vehicle position within a lane. The automatic controller can create virtual springs that attach the steering wheel to a moving home angular position that corresponds to the vehicle direction recommended by the automation."

This idea was also explored in other studies (for an overview: Abbink and Mulder 2010), but all faced the same, important design question: what is the correct level for the forces delivered by such virtual springs? In this paper that level is called level of haptic authority (LoHA), which is related to, but notably different from, Sheridan's level of automation (LoA). The LoHA constitutes how forceful the human-automation interface connects human inputs to automation inputs and mainly addresses the provided support on a skill-based level through a single control interface (most likely to control vehicles or robots). Sheridan's levels of automation are more widely applicable (for example, also to more complex systems like power plants).

The optimal choice of LoHA depends on many factors, such as the quality of the automation system, traditional human factors issues, but also on the task at hand and on the properties of the human operator, both cognitive and neuromuscular. Concerning choosing the correct LoHA, Griffiths and Gillespie recognized that "...the stiffness of the shared controller must be tuned to balance two conflicting goals [...]: if the virtual spring is too stiff, the driver may find it difficult to overpower the controller's actions, but if the spring is very weak, disturbances would cause excessive error..." In other words, for high stiffness—a high LoHA—the performance is expected to be improved—of course, only as long as the human operator agrees with the automation and there are no automation failures. For low stiffness—a low LoHA—there is not much support from the system, so little performance increase is expected, but the system will be easy to overrule. This hypothesized behavior was experimentally verified in our laboratory (Abbink and Mulder 2009).

Regrettably, none of the user-studies with a haptic shared control system with fixed level of automation mention the rationale based on which the forces—and thus the LoHA—were determined. Trial-and-error tuning of the forces seems the most used approach. Still, the trial-and-error tuning often yielded beneficial results: e.g., Griffiths and Gillespie (2005) showed that their haptic assist system improved drivers' lane-following performance with at least 30%, while reducing visual demand by some 29%.

A more scientific method is to experimentally determine the trade-off in performance and control effort associated with different LoHA (Marayong and Okamura 2004; Passenberg et al. 2011) or to base it on validated human models (Abbink and Mulder 2010), for which experimental techniques have recently been made available (Abbink et al. 2011).

A higher LoHA can also be reached by using automation and haptics in a different manner than described above. For example, a haptic flight director was proposed for aviation purposes (de Stigter et al. 2007), where a stick input of zero caused the automatic controller to deliver its optimal control input to the aircraft. The haptic flight director made the control actions of the automatic controller tangible as disturbance forces, which pilots could mitigate by keeping the stick centered. On top of that, own steering inputs could be added to those of the automatic controller by moving the stick away from the center. This approach is sometimes called “mixed-input shared control” (Abbink and Mulder 2010) or “indirect haptic feedback” (Alaimo et al. 2010). Note that this approach presents the haptic control information as disturbances that need to be resisted, which is quite different from the guidance approach, in which the system and the human act together to realize the required torques. As a result, it will be difficult for the human to understand when and how to disagree with the automation, since it is unclear what control inputs the automation gives to the vehicle. Additionally, “mixed-input shared control” essentially alters the controlled dynamics of the vehicle over time, which will likely degrade the internal model of the pilot about the aircraft dynamics, and therefore decrease situational awareness and skill in case of automation failure. Such long-term effects have not been studied in haptic shared control literature, but need to be taken into account when assessing the human factors impact of different design alternatives for haptic shared control.

2.2 Haptic shared control with variable authority

It is increasingly well understood that haptic feedback can not only be implemented with a static level of control authority but also that it lends itself very well to adaptive or adaptable automation. Our own first experiences with variable authority come from developing a driver support system for car following, as an alternative to the autonomous adaptive cruise control (ACC) system. We designed a haptic gas pedal with force feedback and variable stiffness (Abbink 2006; Mulder et al. 2009; Mulder et al. 2010; Mulder et al. 2011). The forces on the gas pedal were designed to actively communicate changes in the car-following situation. The stiffness could be increased to increase the authority with which the forces were applied to the gas pedal. The closer a lead vehicle was to the own

vehicle, the higher the stiffness of the gas pedal, and hence the authority with which the control system forces were applied to the gas pedal. With the unidirectional stiffness of the gas pedal, greatly increasing stiffness would not lead to an automatic controller, because accelerating a vehicle is achieved by depressing the gas pedal, which is actually made more difficult by increasing the stiffness. When the driver cannot override the stiffness of the haptic gas pedal, there is, therefore, no control possible at all. In contrast, for steering—which has bidirectional stiffness (e.g., Griffiths and Gillespie 2005)—haptic shared control enables the steering wheel to center around the steering angle determined by the control system. Increasing the stiffness in this case leads to an increased enforcement of the steering angle coming from the control system (Abbink and Mulder 2009). Hence, when the stiffness is large enough to prevent the driver from overriding the steering angle imposed by the control system, the control system effectively has become an automatic control system. Note that traditional methods to intervene with such automatic system could still be possible, like switching it off.

We learned that the support system yielded better results if the continuous haptic guidance was based on a solid experimentally grounded understanding of how drivers naturally respond to traffic situations (Mulder et al. 2011), which holds equally true for autonomous systems (Goodrich and Boer 2003), and therefore for haptic shared control as well.

Another interesting implementation of variable-autonomy haptic shared control systems was proposed by Goodrich et al. (2008), who opted for two fixed levels of automation based on the metaphor of horse riding, the so-called “H-mode”. When riding a horse, the rider can opt for loose-rein control, where the horse has most authority, and the rider loosely feels on the reins what the horse is doing (high LoHA). Alternatively, if the rider wants to enforce a certain path, he can grip the reins tighter and enforce his will (low LoHA of the automatic controller). These two modes are implemented in a flight simulator, with a “loose-rein” mode, where the system has a high LoHA (i.e., a high stiffness around the optimal steering angle) and where the pilot can then feel the system’s actions. The second mode is the “tight-rein” mode, where the stiffness is lower and the pilot assumes most of the control authority, hence the system has a low LoHA. They propose a binary user-initiated change in autonomy, or adaptable automation: “...if the pilot sustains a firm input while in loose reins, a transition into tight reins occurs. [And] while operating in tight reins, the automation [may] offer or initiate the transition to appropriate loose reins behavior.” The H-mode concept was also extended to automotive applications (Flemisch et al. 2008), but not much experimental evaluation has been published yet.

A more continuous shifting of autonomy can be realized by smoothly adapting the LoHA (i.e., control device stiffness around the optimal steering position) as a function of a dynamic trade-off between performance and control effort (Passenberg et al. 2011), or as a function of criticality (Abbink and Mulder 2009). These methods will result in a kind of “haptically *adaptive* automation”: system-initiated changes in LoHA. For user-initiated changes in LoHA (i.e., “haptically *adaptable* automation”), a signal from the driver is needed to extract when he or she is both willing and able to change LoHA.

In short, both continuously adaptive and adaptable automation interfaces are possible with haptic shared control. This will be argued in Sect. 3, and two case studies of adaptive automation are presented from our laboratory.

3 Case studies on automotive haptic shared control

Since 2002, the authors of this paper have collaborated on designing haptic shared control for automotive applications. Our research is strongly driven by the conviction that we need to understand drivers in order to support them best. In our efforts to do so, we have developed several support systems (Mulder et al. 2011; Mulder et al. 2010; Mulder et al. 2009; Abbink and Mulder 2009) and several experimental techniques to quantify driving behavior (Abbink et al. 2011; Mulder et al. 2008) and even to describe skill-based driving (i.e., at the operational level: steering and longitudinal control) using computational models (Abbink and Mulder 2010). In particular, we have focused on understanding the neuromuscular response to forces, both as perturbations or as guidance forces (Abbink 2006; Abbink et al. 2011).

Our approach has led us to the following design philosophy for haptic shared control: to design force feedback based on human capabilities and essentially “mirror the human”. If the human can adapt his/her impedance around his/her desired trajectory, so should the haptic shared controller. Figure 1 illustrates a model of this design philosophy: properties of human are mirrored in the controller, and together, they determine the steering input to the controlled system (e.g., a vehicle).

The calculation of guidance torques can be separated in two distinct mappings. The first mapping is that from system states in the environment to desired steering angle. Based on system states in the perceived environment, both human and haptic shared control system have reference trajectories they want to achieve (ref_{human} and ref_{sys}), which results in an optimal steering angle for each (x_{des} and x_{opt} , respectively). This part essentially is the automatic controller, and its output could be used to directly control the vehicle.

However, the essence of haptic shared control lies in the second mapping, that from desired steering angle to guiding force, that interacts with the human force. Note that if the human wants to impose his desired steering angle x_{des} , he will regard other forces that do not steer toward x_{des} as perturbations and become more stiff to resist them. Literature on human motion control has identified the neuromuscular mechanisms behind this stiffness adaptation and has offered methods to quantify the total physical interaction dynamics H_{pi} and include them into computational models (Abbink and Mulder 2010).

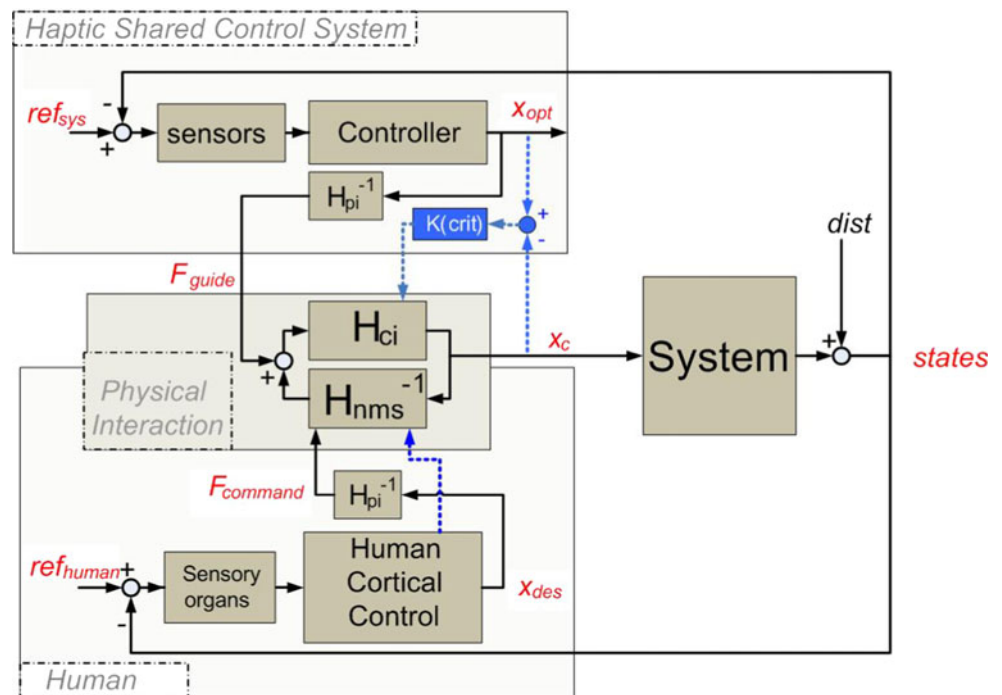
The proposed architecture in Fig. 1 allows the steering wheel system to respond likewise: when a driver does not respond adequately to a critical situation, the impedance of the steering wheel around x_{opt} (i.e., the LoHA) can be smoothly and temporarily increased by the function $K(crit)$, thereby guiding the driver toward the automation’s optimal steering angle x_{opt} . We initially explored the interactions between force feedback and changing stiffness feedback in an abstract steering task (Abbink and Mulder 2009), showing that a high LoHA generally yields benefits in performance and control effort (but increased workload in case the automation controller and human disagreed).

We then investigated haptic shared control for automotive steering, first for simple cases where drivers needed to stay within a lane on a curving road (a single optimal trajectory), before progressing toward supporting lane changes and bifurcation choices in evasive maneuvers (multiple optimal trajectories).

3.1 Case study 1: Automotive steering guidance for lane keeping and curve negotiation

The lane-keeping support system we developed (Mulder et al. 2008) in our laboratory was based on an automation system that used a single look-ahead point to generate trajectories that resembled human curve negotiation (cutting curves). The lateral error at that look-ahead time (0.7 s) was translated to forces, in order to communicate the optimal steering angle x_{opt} . As an innovation with respect to other literature, we used the concept of smooth system-initiated changes in LoHA (i.e., impedance around x_{opt}), in order to continuously communicate the criticality that the automation perceives. In our initial versions of the system, we related changes in LoHA to lateral error: in order to guide drivers more in case they were getting closer to the lane boundaries. Performance benefits were found in lane keeping (e.g., lateral error, time-to-lane crossing) as well as reduced control activity to realize those benefits. However, sometimes, subjects still fought the guiding forces, indicating the automation trajectories did not optimally match the driver’s trajectories.

Fig. 1 A schematic, symmetric representation of haptic shared control. Both the human and the shared controller have sensors to perceive changes in system states (possibly perturbed by *dist*), each having a goal (ref_{human} and ref_{sys} , respectively). During haptic shared control, both human and system can act with forces on the control interface (with $F_{command}$ and F_{guide} , respectively). Through physical interaction, the control interface (H_{ci}) exchanges force and position with the human limb (H_{nms}). Other variables are discussed in the text (*replicated from Abbink and Mulder 2010*)



3.2 Case study 2: Automotive steering guidance for lane changes and evasive maneuvers

All haptic guidance literature for automotive steering simplifies driver support to the previously discussed lane-keeping support; in other words, the automation follows only one reference trajectory. But in reality, drivers may also want to switch lanes, or avoid objects, such as braking vehicles or sudden hazards.

In case of lane changes, an inherent haptic shared control design problem arises “... from the fact that lane keeping and lane changing are two opposite objectives, which can not be met at the same time. It occurs when there exists a mismatch between the goal of the support system (i.e., lane keeping) and the goal of the driver (i.e., changing lanes)” (Tsoi et al. 2010). Griffiths and Gillespie (2005) chose to mitigate this inherent negative effect by tuning the LoHA to be low, so that lane changes were still possible, in spite of the counteracting from the haptic shared controller. In an experiment, they used objects in the center of the desired trajectory to enforce lane changes and observed higher control effort and a larger number of collisions with these objects when receiving haptic guidance, compared to manual control. The issue could be avoided altogether by temporarily switching off the system (e.g., by linking it to the turn indicator). Since drivers sometimes forget to use the turn indicator, such switching might be confusing and

annoying. And afterwards, the benefits of lane keeping will be absent until the driver activates the system again.

We explored a more continuous solution by avoiding conflicts between the automation and the human, in particular, by incorporating algorithms to switch the desired trajectories of the haptic shared controller (Tsoi et al. 2010). The switching was initiated by the automation, based on environmental constraints and the forces the driver exerted on the steering wheel. This felt to the driver as slightly pushing the vehicle over a “hill”, and by using the turn indicator, this “hill” was less steep. Experimental evidence showed that all the benefits of lane-keeping support persisted, while comfortably allowing lane changes, with low levels of control effort.

Subsequently, we investigated how we could use automation in more time-critical situations, in particular where multiple trajectories need to be considered? For example, how should we deal with an object suddenly appearing in the reference trajectory of the automation system? In such cases, even if the sensors would accurately detect the object, an engineering problem remains to determine which way to avoid the object. We believe it would be best if we could safely leave this choice to the driver, who may have better insight into the best solutions to avoid the situation. Additionally, automotive companies may not want to be responsible for difficult choices (should I hit the suddenly appearing pedestrian or should I escape into a ditch to

avoid him?) and vice versa; drivers are not likely to want to leave this up to an autonomous system.

Instead of suddenly switching off automation or haptic shared control in such complicated cases, would it not be feasible to expand the existing framework of continuous guidance? A novel concept was proposed (Della Penna et al. 2010) that is based on temporarily reducing the steering wheel stiffness around the steering angle that will steer straight into the object, thereby making it easier to steer away to avoid the crash. In very time-critical situations, the steering wheel stiffness might even become negative, forcing the driver to make active effort to add his own stiffness (through co-contraction) to keep on steering toward the object. This leaves the authority of choice completely up to the driver, but once the choice is made, the driver is assisted to steer fast in the direction he chooses, where he will be guided on the redirected trajectory. Note that as long as the driver does not make a choice, he will feel that the automation system communicates increasing criticality and wants the driver to make a choice.

Experimental results in a driving simulator showed that when avoiding objects in highly critical situations, the designed extended haptic shared control system—compared to unsupported, manual control—, helped to reduce the number of crashes, with reduced control effort (forces) and control activity (steering actions) (Della Penna et al. 2010). Essentially, the haptic shared control system allowed drivers to choose their preferred trajectory around the object and helped to quickly execute this choice without deteriorating overshoot that might result from excessive steering. A decreased response time of at least 100 ms when using the haptic shared control system compared to manual control without any haptic guidance indicated that subjects could now respond on a neuromuscular level to the situation, something also found for the haptic gas pedal (Abbink 2006; Abbink et al. 2011).

In short, haptic shared control is not only useful to guide drivers along a single trajectory but also when drivers are changing lanes or making choices how to maneuver around unexpected obstacles. In other words, one concept of continuous human-automation interaction is possible that leaves the driver in final authority and results in performance benefits at reduced levels of control activity. One can imagine how the concept of calculating one or more desired trajectories can be extended to include other vehicles as well. As the supported traffic situations become more complex and more likely to suffer from automation-induced surprises or mismatches, it will be even more important to keep the driver continuously involved, who should be able to feel when the steering wheel is not moving according to the driver's wishes, and then to easily overrule it.

4 Discussion and conclusion

Although this article has presented many arguments in favor of haptic shared control as a human-automation interface, still, the systems are not optimal yet. Although subjective reports were very positive, some subjects complained that they were sometimes fighting the system, which is shown in most experimental studies by sporadic increases in interaction forces (e.g., Griffiths and Gillespie 2005; Forsyth and MacLean 2006; Tsoi et al. 2010).

In some cases, the human will accept that someone/something else takes over his tasks and might even prefer this. In other cases, the human will want to feel free and in control. But how do we ensure as engineers that the human operator feels free to act? Sociologist Zygmunt Bauman (2000) argues that "...to feel free means to experience no [...] resistance to the moves [that are either] intended, or conceivable to be desired." This can be applied directly to our engineering guidelines: if we can make human-automation interaction in such a way that the human agrees with the actions of the automation, the human is likely to feel free and in control.

This requires an even better understanding of the human, captured in quantitative models, as mentioned by Parasuraman and Wickens (2008): "Importantly, designers in fields such as aerospace systems are beginning to call for computational models of human-automation interaction, which can both predict the effectiveness of automation systems and yield guidelines for designers to follow." Such models have been argued for automation systems in previous work (Goodrich and Boer 2003), but they also apply to shared control. Therefore, we believe that more research is needed to ground the design of haptic shared control on two experimental model mappings:

- a human-centered automation system, i.e., the mapping from error states to optimal steering angle; and
- a human-centered force generation system, i.e., the mapping from optimal steering angle to forces that also allow scaling of the haptic LoHA (i.e., impedance around optimal steering angle).

It will be interesting to explore adaptive mappings that learn the preferences of individual drivers, and so better match the needs and goals of the individual driver.

Human factors literature points to persistent issues in human-automation interaction, even after decades of research. According to this paper, these issues can be summarized as follows: human-automation interaction issues are likely to arise

- if the human does not understand the automation (in terms of capabilities, boundaries of operation, current functionality, goals, and LoA); or

- if the automation (and the engineer that made it) does not understand the human, (in terms of capabilities, goals, and inputs to the vehicle).

Based on literature, four design guidelines for human-automation interaction were chosen that state that the human should

1. always remain in control, but should be able to experience or initiate smooth shifts between levels of authority;
2. receive continuous feedback about the automation boundaries and functionality;
3. continuously interact with the automation; and
4. should benefit from increased performance and/or reduced workload.

Haptic shared control is a promising approach to meet these four guidelines, but the worlds of haptic interfacing and human-automation interaction design seem to be somewhat isolated. This paper has attempted to bridge the gap, by discussing several realizations of haptic shared control in light of these four design guidelines. We have discussed haptic shared control literature that indicates that all four guidelines can be met. We have also shown that experimental evidence exists that haptic shared control can lead to short-term performance benefits, such as faster and more accurate control of vehicles and robotic devices, often at lower levels of control effort. A few studies have shown reduced demand for visual attention.

However, such benefits are usually discussed in terms of benefits over manual control, and not compared to automation. Moreover, the automation system, on which the continuous forces are based, is often not separately evaluated and simply assumed to be perfect. Therefore, many important human-automation interaction issues remain unaddressed for haptic shared control (de Winter and Dodou 2011). For example, we hypothesize that it is easier to maintain skills and situation awareness and to catch an automation error when the continuous actions of an automatic controller is presented haptically, rather than through binary alerts, but this has not been proven yet. These and other human factors issues have received little experimental attention as of yet, in particular: issues with long-term use, trust, overreliance, dependency on the system, and retention of skills. This should be addressed in future research on haptic shared control.

In short, we believe that haptics can be key to designing better human-automation interaction (especially for vehicular control), but that more experimental evidence is needed to understand long-term human factors issues. Therefore, it is important to join the expertise in the fields of human-automation interfacing and haptics, with the common goal of making human-automation

interaction better: more intuitive, avoiding misuse, disuse, and abuse.

Acknowledgments The contribution of D.A. Abbink was supported by VENI Grant 10650 from the Netherlands Organization for Scientific Research (NWO).

Open Access This article is distributed under the terms of the Creative Commons Attribution Noncommercial License which permits any noncommercial use, distribution, and reproduction in any medium, provided the original author(s) and source are credited.

References

- Abbink DA (2006) Neuromuscular analysis of haptic gas pedal feedback during car following. PhD thesis, Delft University of Technology
- Abbink DA, Mulder M (2009) Exploring the dimensions of haptic feedback support in manual control. *J Comput Inform Sci Eng* 9(1):011006–1–011006–9
- Abbink DA, Mulder M (2010) Neuromuscular analysis as a guideline in designing shared control. *Adv Haptics* 109:499–516
- Abbink DA, Mulder M, Van Der Helm FCT, Mulder M, Boer ER (2011) Measuring neuromuscular control dynamics during car following with continuous haptic feedback. *IEEE Trans Syst Man Cybern, Part B*, 41:1239–1249
- Alaimo SMC, Pollini L, Bresciani JP, Bühlhoff HH (2010) A comparison of direct and indirect haptic aiding for remotely piloted vehicles. 19th IEEE International Symposium on Robot and Human Interactive Communication Principe di Piemonte—Viareggio, Italy, 12–15 Sept 2010, pp 541–547
- Bainbridge L (1983) Ironies of automation. *Automatica* 19(6): 775–779
- Bauman Z (2000) *Liquid modernity*. Polity Press, Malden
- Billings CE (1997) *Aviation automation: the search for a human-centered approach*. Erlbaum, Mahwah
- de Stigter S, Mulder M, van Paassen MM (2007) Design and evaluation of a haptic flight director. *J Guid Control Dyn* 30(1):35–46
- de Winter JCF, Dodou D (2011) Preparing drivers for dangerous situations: A critical reflection on continuous shared control. In: *IEEE proceedings of the IEEE SMC conference*, Anchorage, Alaska, USA
- Della Penna M, Paassen MM van, Abbink DA, Mulder M, Mulder M (2010) Reducing steering wheel stiffness is beneficial in supporting evasive maneuvers. In: *IEEE proceedings of the IEEE SMC conference*, pp 1628–1635, Istanbul, Turkey
- Endsley M, Kiris E (1995) The out-of-the-loop performance problem and level of control in automation. *Hum Factors* 37:381–394
- Flemisch F, Kelsch J, Löper C, Schieben A, Schindler J (2008) Automation spectrum, inner/outer compatibility and other potentially useful human factors concepts for assistance and automation. In: de Waard D, Flemisch FO, Lorenz B, Oberheid H, Brookhuis KA (eds) *Human factors for assistance and automation*. Shaker Publishing, Maastricht, pp 1–16
- Forsyth BAC, MacLean KE (2006) Predictive haptic guidance: intelligent user assistance for the control of dynamic tasks. *IEEE Trans Vis Comput Gr* 12(1):103–113
- Goodrich MA, Boer ER (2003) Model-based human-centered task automation: a case study in ACC system design. *IEEE Trans Syst Man Cybern-Part A: Syst Hum* 33(3):325–336
- Goodrich K, Schutte P, Flemisch F, Williams R (2006) Application of the H-mode, a design and interaction concept for highly

- automated vehicles, to aircraft. IEEE/AIAA 25th digital avionics systems conference, pp 1–13
- Goodrich K, Schutte PC, Williams RA (2008) Piloted evaluation of the H-mode, a variable autonomy control system, in motion-based simulation. In: AIAA Proceedings, pp 1–15
- Griffiths PG, Gillespie RB (2005) Sharing control between humans and automation using haptic interface: primary and secondary task performance benefits. *Hum Factors* 47(3):574–590
- Inagaki T (2003) Adaptive automation: sharing and trading of control. In: Hollnagel E (ed) *Handbook of cognitive task design*. Erlbaum, Mahwah, pp 147–169
- Kaber DB, Endsley MR (2004) The effects of level of automation and adaptive automation on human performance, situation awareness and workload in a dynamic control task. *Theor Issues Ergon Sci* 5(2):113–153
- Lam TM, Boschloo HW, Mulder M, van Paassen MM (2009) Artificial force field for haptic feedback in UAV teleoperation. *IEEE Trans Syst Man Cybern—Part A: Syst Hum* 39(6):1316–1330
- Lee JD (2008) Review of a pivotal human factors article: humans and automation: use, misuse, disuse, abuse'. *Hum Factors: J Hum Factors Ergon Soc* 50:404–410
- Marayong P, Okamura AM (2004) Speed-accuracy characteristics of human-machine cooperative manipulation using virtual fixtures with variable admittance. *Hum Factors: J Hum Factors Ergon Soc* 46:518–532
- Miller C, Parasuraman R (2007) Designing for flexible interaction between humans and automation: delegation interfaces for supervisory control. *Hum Factors* 49(1):57–75
- Mulder M, Abbink DA, Boer ER (2008) The effect of haptic guidance on curve negotiation behavior of young, experienced drivers. *IEEE International Conference on systems, man and cybernetics*, pp 804–809
- Mulder M, Mulder M, van Paassen MM, Abbink DA (2009) Haptic gas pedal feedback. *Ergonomics* 51(11):1710–1720
- Mulder M, Pauwelussen JJA, Van Paassen MM, Mulder M, Abbink DA (2010) Active deceleration support in car following. *IEEE Trans Syst Man Cybern—Part A: Syst Hum* 40(6):1271–1284
- Mulder M, Abbink DA, van Paassen MM, Mulder M (2011) Design of a haptic gas pedal for active car-following support. *IEEE Trans Intell Transp Syst* 12(1):268–279
- Norman DA (1990) The “problem” with automation: inappropriate feedback and interaction, not “over-automation”. *Philos Trans R Soc B-Biol Sci* 327(1241):585–593
- O’Malley MK, Gupta A, Gen M, Li Y (2006) Shared control in haptic systems for performance enhancement and training. *J Dyn Syst Meas Control* 128(1):75
- Opperman R (1994) *Adaptive user support*. Erlbaum, Hillsdale
- Parasuraman R, Riley V (1997) Humans and automation: use, misuse, disuse, abuse. *Hum Factors* 39:230–253
- Parasuraman R, Wickens CD (2008) Humans: still vital after all these years of automation. *Hum Factors: J Hum Factors Ergon Soc* 50(3):511–520
- Passenberg C, Groten R, Peer A, Buss M (2011) Towards real-time haptic assistance adaptation optimizing task performance and human effort. *World Haptics Conference, Istanbul*
- Rosenberg LB (1993) Virtual fixtures: perceptual tools for telerobotic manipulation. *Proceedings of IEEE annual virtual reality international symposium*, pp 76–82
- Rouse W (1976) Adaptive allocation of decision making responsibility between supervisor and computer. In: Sheridan TB, Johansson G (eds) *Monitoring behavior and supervisory control*. Plenum Press, New York, pp 221–230
- Scerbo M (2001) Adaptive automation. In: Karwowski W (ed) *International encyclopedia of ergonomics and human factors*. Taylor & Francis, London, pp 1077–1079
- Sheridan TB (1992) *Telerobotics, automation, and human supervisory control*. MIT Press, Cambridge
- Sheridan TB (2002) *Humans and automation: system design and research issues*. HFES issues in human factors and ergonomics series, vol 3. Wiley, New York, ISBN 0-471-23428-1
- Sheridan TB, Parasuraman R (2006) Human-automation interaction. In: Nickerson R (ed) *Reviews of human factors and ergonomics*, vol 1. Human Factors and Ergonomics Society, Santa Monica, pp 89–129
- Tsoi KK, Mulder M, Abbink DA (2010) Balancing safety and support: changing lanes with a haptic lane-keeping support system. *Proceedings of the IEEE SMC conference 2010, Istanbul*
- Wiener EL, Curry RE (1980) Flight-deck automation: promises and problems. *Ergonomics* 23:995–1011