Driving automation: Learning from aviation about design philosophies

Mark S. Young[1], Neville A. Stanton[1] & Don Harris[2]

[1]School of Engineering & Design, Brunel University, Uxbridge, Middlesex, UB8 3PH, UK

[2]Department of Human Factors, School of Engineering, Cranfield University, Cranfield, Bedford, MK43 0AL

Abstract: Whether we like it or not, automation is gradually taking over the driver's role. Full vehicle automation is predicted to be on British roads by 2030 (Walker et al., 2001). However, experience in aviation gives us some cause for concern for the 'drive-by-wire' car (Stanton & Marsden, 1996). Problems of workload and situation awareness have been attributed as the root cause of some aircraft accidents (Beaty, 1995). These problems may be approached in two different ways, depending on the level of automation used. These two philosophies are exemplified by the positions of the two major aircraft manufacturers. Airbus use a 'hard protection' system in their A320 and A340 series, employing automation to prevent error, and hence it can override the pilot. Boeing, on the other hand, opted for 'soft protection' in their 777 aircraft, using automation as a tool to aid pilots, and not giving it the authority to override pilot control (Hughes & Dornheim, 1995). This paper speculates whether hard or soft automation provides the best solution for road vehicles, and considers an alternative design philosophy in vehicles of the future based on coordination and cooperation.

Biographical Notes: Mark Young is a Research Lecturer in the School of Engineering and Design at Brunel University, UK. He received a B.Sc. and Ph.D. from the University of Southampton, both in the Department of Psychology. His doctoral and postdoctoral research were focused on the cognitive ergonomics of transport systems, particularly with advanced vehicle technologies. He spent two years working in industry for the Rail Safety and Standards Board, applying his experience to research and regulations for the UK rail network. Before joining Brunel in October 2004, Mark was a Visiting Fellow in the Department of Aviation, University of New South Wales, Australia, working on a Royal Society fellowship. Mark is a Registered Member of The Ergonomics Society, and sits on the vehicle design working group for the Parliamentary Advisory Council for Transport Safety (PACTS).

1 Driving automation

During the course of the 20[th] century, the automobile has grown from a humble horseless carriage into one of the most technologically advanced mass market commodities available. In the first few years of the 21[st] century, we are starting to see a new generation of technology enter the automobile, in the form of driving automation.

Of course, vehicle automation has been around for some time, with the first automatic gearboxes becoming widely available in the 1940s. Indeed, the 'auto-mobile' itself is a holistic example of automation. However, where traditional automatic systems have sought to assume the lower-level, operational components of vehicle control, new technologies are taking over more tactical and even strategic aspects of driving (cf. Ranney, 1994). Thus our choice of terminology above – <u>driving</u> automation vs. <u>vehicle</u> automation – was quite deliberate, as we distinguish between low-level vehicle control and higher-level cognitive driving tasks.

An automatic gearbox, then, largely carries out a mechanical, vehicle control task (though we accept that there is also some cognitive element regarding the decision to change gear). Similarly, conventional cruise control (CC) simply adjusts the throttle to maintain speed. On the other hand, adaptive cruise control (ACC) removes a cognitive task from the driver – perceiving the speed of a lead vehicle, deciding whether to adjust speed in response, and taking appropriate action. Collision warning and avoidance systems (CWS/CAS) take this a step further, by making a potentially stressful decision about whether to take emergency action. Even automatic steering systems (AS), which might appear to be an example of vehicle automation, relieve the driver of a significant cognitive workload (Young & Stanton, 2002), owing to the fact that steering is a second-order tracking task (Wickens et al., 1998). With more of these advanced systems being offered in mass production vehicles (e.g., Richardson et al., 1997), it seems that the automation is increasingly capable of taking over driving as well as vehicle control tasks. At present, the technologies are working independently of each other, resembling the architecture of early generation jet transports (Applegate & Graeber, 2005). Nevertheless, the degree of automation integration in road vehicles is increasing all the time, and full integration of these systems will make an autonomous vehicle a commercial reality. Futurologists and ergonomics researchers alike are predicting that by 2030 fully automated vehicles will be on our roads (e.g., Walker et al., 2001).

Whilst the engineering of these vehicles seems on track, the understanding of the interaction between the computing, the vehicle's behaviour and the driver's reactions seems much less clear. Where many drivers may be happy to relinquish control to automatic systems, others will still want the satisfaction of controlling their vehicles manually. While we still retain a human in the loop, we need to consider how the design of the automation fits around the driver. Fundamental human factors research in automation and human supervisory control spans over two decades (e.g., Bainbridge, 1982; Hollnagel, 1993; Norman, 1990; Reason, 1990), but it is only more recently that we are catching up with research in vehicle automation (e.g., Stanton & Young, 2000; Young & Stanton, 1997).

Much of the ideology behind driving automation has been around for some time in aviation systems, and there is considerable ergonomics knowledge about the promises and problems that can be extrapolated from the aviation domain (Billings, 1993; Stanton & Marsden, 1996; Wiener & Curry, 1980). Some of that knowledge might give us cause for concern. Indeed, Stanton & Marsden (1996) noted that automation has been

implicated in a number of fatal aviation accidents. Root cause analysis by accident investigators has identified psychological factors such as boredom and inattention under conditions of low workload, cognitive strain under conditions of very high workload, failure of automated systems to meet pilots' expectations, and over-reliance on the technology. At the psychological level, none of these human factors problems is domain-specific, and we can surely anticipate them with vehicle or driving automation. Mental workload can be too high with the proliferation of in-vehicle systems, or too low as more driving tasks are assumed by the automation (Young & Stanton, 2002). Equipment reliability with systems like ACC or CAS is reminiscent of similar problems in aviation, with direct implications for driver trust (e.g., Kazi et al., 2005; Lee & See, 2004). Skills will fade if driving automation is relied upon to substitute manual operation. In a critical situation, all of these problems could interact to cause safety consequences of the like we have already seen in aviation. The challenge for ergonomics is to ensure that driving automation systems are designed to avoid these problems while maximising the benefits.

Many of these performance issues can be addressed by appropriate design of the automation. Previous work in the general area of automation has often concentrated on how much control the automation should have over the task as a way of managing interaction with the human operator (e.g., Kaber & Endsley, 1997; Labiale, 1997). Taxonomic approaches to automation (e.g., Parasuraman et al., 2000) typically classify the technology along several levels from full operational control to simple decision support or task assistance. By using strategic allocation of function decisions and mixing levels of control between the human and the machine, it is theoretically possible to manage mental workload, trust, and keep the driver in the loop to maximise the benefits that automation can offer.

In this paper we suggest a slightly different emphasis to the solution. Underpinning all of these levels of control, though, is a more fundamental question concerning the level of authority assigned to the automation. If it is the 'strong, silent' type, it could cause problems for driver performance as it is unclear who is in control at any particular time, and therefore what appropriate actions should be taken. At the macro level, this question represents our guiding philosophy when implementing automation – who has the ultimate power of veto, human or technology? As we shall see in the next section of this paper, the answer to that question can have a profound influence on safety and performance.

2   Learning from aviation

In aviation, the two major aircraft manufacturers, Airbus and Boeing, have developed different philosophies regarding the authority of automation: hard automation (the Airbus philosophy) and soft automation (the Boeing philosophy).

Airbus, in its fly-by-wire aircraft (A320/A330/A340 and the new A380 series) uses a 'hard protection' system for its flight control systems, believing that the automation technology exists to prevent the pilot from inadvertently exceeding safety limits. For instance, Airbus systems have hard speed envelope protection features that will prevent the pilot from stalling the aircraft and from pulling more than 2.5g even in an emergency (Hughes & Dornheim, 1995). The rationale behind hard protection is largely to protect the airframe – if the pilot should inadvertently take the aircraft beyond its performance envelope, automatic interventions will prevent damage and maintain flight dynamics. Hard automation, then, employs the technology to prevent error; as such it has ultimate authority and can override the human operator's inputs.

Boeing, on the other hand, opted for 'soft protection' in their 777 aircraft, using automation as a tool to aid pilots, and giving pilots full authority to override the automated systems if they want (or need) to. The pilots therefore have access to the full performance envelope and will not be overridden by the automated systems. There are still automated advisories in this soft protection scheme – if the pilot wishes to exceed set limits (such as exceeding 3.5 degree of bank, or pulling the yoke back as the aircraft decelerates below the minimum manoeuvre speed), s/he is required to apply more force than normal on the yoke (Hughes & Dornheim, 1995).

Hard and soft automation therefore use a similar set of sensors and control devices, but to different ends. Hard automation takes the pilot's input, determines whether they are appropriate for the (assumed) situation, and if necessary takes its own action before passing the instructions on to the control surfaces. This can be advantageous for some relatively simple tasks, like flight envelope protection. A good example is if the pilot has received a collision warning and, in a panic reaction, pulls hard back to gain altitude. Without an associated increase in thrust, the aircraft would soon stall. In that situation, the Airbus will itself apply all available thrust to climb without stalling. However, there are circumstances in which this level of computer authority can cause problems rather than resolving them – problems which usually result from the context of operation. The crash of an A320 at an air show near Paris in 1988 was caused because the automation had made an incorrect interpretation of the pilot's inputs (Beaty, 1995). Making a low-level fly-by with the undercarriage lowered, the computer assumed the pilot wanted to land and so throttled back the engines. To make matters worse, as the aircraft descended through 50 feet, many of the flight envelope protections were also removed when it automatically went into landing mode. When the pilot attempted to pull clear, the necessary thrust was not available, and the aircraft plunged into woodland at the end of the runway (Investigation Commission of Ministry of Transport, 1989). Clearly, the automation was unaware of the context of the flight (i.e., a low pass along the runway at an air show as opposed to a landing at an airport) and the intent of the pilot (i.e., not to land). The pilot's intent and the automation's intent were not congruent.

Soft automation makes a similar assessment of pilot inputs, but will only give feedback if the control requests appear to represent a safety risk. If the pilot persists, the soft automation will then pass the inputs directly to the control surfaces without intervention. Again, there are certain situations in which the pilot may legitimately wish to take the airframe beyond its performance limits. An incident involving an engine failure on a China Airlines Boeing 747 in 1985 was only recovered after the aircraft had lost more than 30,000 feet in an uncontrollable dive (see Norman, 1990). Needless to say, the airframe was significantly stressed during both the descent and the recovery, and substantial damage was caused (though only a few injuries were sustained on board). Interestingly, though, if that aircraft had been fitted with a hard protection system, the pilot would not have been able to recover control, as the necessary control inputs exceeded the performance envelope (although it could be argued that an Airbus flight control system would not have allowed the aircraft to get into such a situation in the first place). Both philosophies have advantages, then, and naturally, both Airbus and Boeing think that their scheme is best (Hughes & Dornheim, 1995).

In a coarse attempt to determine the relative merits of hard vs. soft protection in aviation automation, we have analysed all of the major (i.e., hull-loss) accidents

involving the modern automated[1] type aircraft of both Airbus and Boeing over the last 20 years. These aircraft types are associated with the automated systems under discussion here, and thus provide a reasonable comparison between the two philosophies. Table 1 shows all those accidents which can be classified as automation related – i.e., as having a direct cause attributable to some mismatch between human and automated activities. In table 2, we summarise those major accidents which are not automation related.

Table 1: Automation related major accidents involving modern automated aircraft

Date	Aircraft	Fatalities	Location	Description
26/6/88	A320	3	Habsheim	Paris air show (see text)
14/2/90	A320	92	Bangalore	Engines in idle descent mode on approach; fell short of runway
20/1/92	A320	87	Strasbourg	Confused descent mode resulting in controlled flight into terrain
14/9/93	A320	2	Warsaw	Runway overrun as windshear on landing affected automatic braking systems
7/2/01	A320	0	Bilbao	Heavy landing following turbulence on approach; crew attempted go-around but automatic protection envelope prevented it
30/6/94	A330	7	Toulouse	Ground impact following test flight take-off due to misunderstanding of autopilot mode and overconfidence in aircraft abilities
20/12/95	B757	160	Cali	Hit mountain after confusion over directional beacon in flight management system
6/2/96	B757	189	Dominican Republic	Faulty airspeed indicator caused confusion with autopilot
15/4/02	B767	129	S. Korea	Struck mountain on circling approach after captain had taken over from autopilot and lost situation awareness

With twice as many major automation related accidents being in Airbus aircraft, these data would seem to imply that the hard automation philosophy leads to more problems of human performance than the soft protection approach. Even the Boeing accidents are only tenuously related to automation, as they were primarily problems of situation awareness following some fault on the flightdeck. Also clear is an almost reverse trend on non-automation related accidents (Boeing: 10; Airbus: 6), and the fact that despite all of them being hull-loss accidents, many did not result in fatalities. As we stated above, though, this is a coarse analysis, and does not control for absolute numbers of aircraft of each type, nor their range (which can affect accident statistics, as short-range aircraft go through more 'cycles' and are thus more prone to failure). Nonetheless, as a baseline comparison this analysis appears to reinforce the conclusion that poorly designed automation can decrease safety and performance in operation.

---

[1] For the purposes of this analysis, the aircraft types are Airbus A319/320/321, A330 and A340, and Boeing 737-600, -700, -800 and -900, 747-400, 757, 767, and 777. Source: www.aviation-safety.net

Table 2: Non-automation related major accidents involving modern automated aircraft
(N.B. this table does not include the September 11, 2001 terrorist attacks)

Date	Aircraft	Fatalities	Location	Description
10/3/97	A320	0	Abu Dhabi	Take-off difficulties led to runway overrun
22/3/98	A320	0	Philippines	Thrust left forward on no. 1 on landing
23/8/00	A320	143	Bahrain	Crashed during go-around; autopilot disconnected on visual approach
28/8/02	A320	0	Phoenix	Poor reverse thrust control on landing
21/3/03	A321	0	Taiwan	Landed on utility vehicle
2/8/05	A340	0	Toronto	Runway overrun in poor weather
4/11/93	B747	0	Hong Kong	Runway overrun
5/8/98	B747	0	Seoul	Runway overrun
31/10/00	B747	83	Taiwan	Took off on wrong runway and hit construction vehicle
2/10/90	B757	46	Guangzhou	Hit by crashing 737
2/10/96	B757	70	nr Lima	Faulty instruments confused flightcrew
14/9/99	B757	0	Costa Brava	Stormy conditions and loss of visual references destabilised approach
1/7/02	B757	2	Überlingen	Mid-air collision after conflicting instructions from ATC and TCAS
26/5/91	B767	223	Thailand	Reverse thrust isolator failed and deployed during flight
6/4/93	B767	0	Guatemala	Runway overrun
26/6/02	B767	0	Japan	Tail strike during training touch & go

The aviation industry has learned that high levels of automation autonomy, coupled with a great deal of automation authority could be problematical for the pilot (e.g., Sarter & Woods, 1994a). In a complex, automated system the combination of autonomy and authority may combine to produce the illusion of perceived animacy. Put another way, the automatic systems appear to develop a mind of their own. A good example of the role of high levels of authority and autonomy in a complex system can be found in the China Airlines Airbus A300 accident at Nagoya, Japan in 1994 (Aircraft Accident Investigation Commission - Ministry of Transport Japan, 1996). The aircraft was being flown manually on the approach (under flight director guidance) with the autothrottles engaged. At some point during the descent, go-around power was accidentally applied by the First Officer. This put the aircraft into an automatic (hard protection) go-around mode – it aborted the landing automatically – which caused a thrust increase and made the aircraft climb. The autopilot was subsequently engaged with the go-around mode still active. The pilots continued to push on the control in an attempt to re-initiate the descent into the airport, while the automatic systems in the aircraft wanted to climb. In other words, there was a mismatch in intent, and the flightcrew ended up fighting the automation. As the pilots applied down elevator, the aircraft compensated and tried to continue the go-around profile by applying opposite stabiliser trim. At this point the alpha floor function (an anti-stall protection) activated, applying full thrust from both engines and thus increasing pitch angle. The pilots elected to hold the throttles shut manually to re-initiate the descent. Nevertheless, the aircraft began to climb steeply, the

nose passed through about 40° pitch angle and airspeed finally fell to about 70 knots. Consequently, the aircraft stalled and crashed, killing 264 of the 271 people on board.

The Nagoya accident was one which began to raise concerns about the implementation of aircraft automation in the US Federal Aviation Administration, and in 1996 it produced a landmark report on the difficulties that flightcrew have when interacting with flight deck automation (FAA, 1996). The report made several recommendations at a systemic level, but primary amongst its concerns were the vulnerabilities in flightcrew management of automation, and situation awareness. The differing implementations of automation can exacerbate problems of mode awareness – perhaps the most significant safety concern associated with the use of automation in commercial aviation (see Sarter & Woods, 1992, 1994b, 1995, 2000). A lack of mode awareness has been linked with many aviation incidents and accidents and 'surprising' aircraft behaviour has been reported by many pilots when what were assumed to be 'hard' protections have not intervened in the control of the aircraft (Sarter & Woods, 1997). Typically, an automated procedure (such as flight level change) could be overridden by a later flightcrew action (increasing the rate of climb). This will cause an 'invisible' mode transition, and the flightcrew might not notice that the original mode has disengaged until an error occurs (in this case, an altitude bust).

Thus we see that, whilst hard automation may ostensibly cause more human factors problems, there are also issues with soft automation systems (and indeed, this classification may be more of a continuum than a dichotomy). Even the FAA's (1996) report stated that automation vulnerabilities were noted in varying degrees across aircraft type, manufacturer, and operator. If anything, there is even more variability in the driving environment, vehicles and drivers than in aviation (Harris & Harris, 2004). The aviation context is actually relatively simple. Control of the aircraft in all three axes can be reduced to simple, mathematical equations. There are many external aids for navigation and control of aircraft and interactions between aircraft are carefully controlled. The road environment is far more complex. Control of the vehicle per se is only one small part of the problem. The relatively uncontrolled environment in which drivers operate (with impromptu and ad hoc interactions with other drivers and unexpected behaviour of other road users and pedestrians) makes this a far more complex environment into which to insert appropriate automated assistance. So we may wonder which philosophy is best for driving automation.

3   Hard or soft driving?

In anticipating the degree of vehicle automation that might become standard issue in 2030, it is wise to consider the question of whether hard or soft automation provides the best solution for road vehicles, such as buses, lorries and cars. As with the examples from Airbus and Boeing, hard automation will overrule the driver if he or she exceeds the vehicle protection envelope whereas soft automation will allow the driver to override it, and have access to the full operating limits of the vehicle.

Before looking forward, though, let us review existing systems from this perspective. The automatic gearboxes already discussed are primarily categorised as hard automation – whilst the driver may usually make limited gear selections (e.g., the use of 'kickdown' or rudimentary gear lever settings), in the main the choice of gear is decided by the automation. Anti-lock braking systems (ABS), traction control, and electronic stability programs (ESP) are similar – leaving aside the possibility to arm or disarm the system,

interventions are made purely on an assessment of the vehicle's dynamics. Their operation is almost transparent to the driver. Conventional cruise control, on the other hand, can be classified as soft automation – the driver decides how and when to set the system, and can resume control at any time. Previously, we identified these kinds of systems as constituting 'vehicle automation' (as opposed to 'driving automation'). On the basis of existing systems, then, there is a trend (though by no means exclusive) towards vehicle automation being synonymous with hard protection.

Moving on to the novel technologies available at present and in the near future, ACC and AS represent examples of soft automation systems, in that they are fully selectable by the driver and any manual control inputs will override them. Parking aids also provide visual and/or auditory assistance in a task that is still subject to manual control. Similarly, collision warning systems (CWS) offer 'soft' information and advice to the driver without necessarily assuming control. It is proposed that in the near future these systems may be linked to automatic steering and braking systems for 'hard' intervention, producing what is effectively a collision avoidance system (CAS). This is analogous with Traffic Alert and Collision Avoidance System (TCAS), which has been mandatory in commercial aircraft for over a decade (see Harris & Smith, 1997, or Harris & Harris, 2004, for a list of parallel aviation and automotive technologies). Whilst TCAS has undoubtedly saved many lives, some of its early implementations actually provided resolution advisories that put the aircraft into greater danger. Incidents have been cited in which TCAS advised a descent in an area of high terrain (e.g., Mellone, 1993) – revealing its lack of awareness, and again demonstrating the importance of context for hard automation.

Likewise, intelligent speed adaptation (ISA) also represents hard automation. ISA automatically restricts speed by linking GPS position monitoring and map databases of speed limits with the engine management system and/or the vehicle's brakes. It has been claimed (Carsten & Tate, 2005) that such a system could save up to 37% of all injury accidents. However, serious concerns remain about its use. Consider the scenarios in which a driver is overtaking a slow-moving or stationary obstacle, and part way through the manoeuvre the driver is confronted by an oncoming car. Most drivers' natural reaction (and in many cases the safer response) is to apply more throttle and elect to temporarily exceed the speed limit in order to avoid a collision. Intervention of a speed limiter at this time could actually increase the risk of a crash. The engineering solution – to incorporate an override switch – is unacceptable as this will inevitably increase reaction times in the intensely time-critical road environment (cf. Harris & Harris, 2004). An alternative 'soft' system using a haptic throttle (i.e., a throttle pedal that gets more stiff the greater the excursion from the speed limit – analogous to the Boeing 777 flight control system protection) was rejected as it ability to be overridden (by pressing harder on the accelerator) as this was thought to be a safety shortcoming.

As with CAS, then, ISA illustrates that such hard protection systems are at best only partially context aware – a dangerous combination, as has been found on occasion in the aviation industry. Furthermore, the ISA system may mis-interpret the intent of the driver when attempting to accelerate in a built up area. The system may perceive the driver's intent to be one of exceeding the speed limit (thereby representing a potentially unsafe behaviour) rather than avoiding an accident by taking evasive action (a safety promoting behaviour). Whilst cost-benefit analyses (e.g., Carsten & Tate, 2005) seem to paint a positive picture, then, it seems that they are missing any potential dis-benefits of the system where the automated system may actually cause accidents. It can be argued that

'hard' protections are only appropriate where their application is not sensitive to the context of use (e.g., ESP or ABS) – in other words, what we have classified as vehicle automation.

Both vehicle and driving automation systems can therefore be designed under soft or hard protection philosophies, and indeed there is a semblance of a pattern emerging in terms of automation strategies – hard protection for vehicle automation, and soft protection for driving automation. Table 3 summarises the relevant technologies and classifications we have covered here.

Table 3: Matrix of hard and soft automation categories against vehicle and driving automation types

	Hard automation	Soft automation
Vehicle automation	Automatic transmission ABS Traction control ESP	CC
Driving automation	CAS ISA	ACC AS CWS Parking aids

On the basis of this classification and the assumptions from aviation, we might think that driving automation is heading in the right direction – if we take the view that soft protection is better for human performance. We now need to evaluate whether we truly can extrapolate from the problems in aviation, as Stanton & Marsden (1996) predicted. Whilst automobile accidents are rarely investigated in the same level of detail as aviation incidents, there is a notable history of accidents involving vehicle automation, in the phenomenon of unintended acceleration (Schmidt, 1993). Typically associated with automatic gearboxes (though also observed with cruise control), the problem occurs when drivers unwittingly hold their foot on the accelerator pedal rather than the brake when selecting gear. Consequently, the car speeds off, and the driver – thinking their foot is already on the brake – gets into a state of cognitive lockup (Moray & Rotenberg, 1989), presses even harder, and the vicious circle only ends when the car crashes into an obstacle. Furthermore, ABS, traction control, and ESP systems could all potentially change the driver's style, encouraging them to 'push the envelope' as they know there is a safety net available. As with aviation again, then, what was the last line of defence becomes the first point of control (cf. Billings, 1993).

More recently, ergonomics concerns with automation have centred around the cognitive and performance implications. Various researchers have looked at issues of mental workload, situation awareness, trust, and mental models (see Stanton & Young, 2000, for a summary). Generally, these concerns are focused on the higher-level driving automation systems, rather than vehicle automation. So we may seek to classify these cognitive problems according to whether the system falls into the hard or soft automation category.

Issues of mental workload have been identified with some soft driving automation systems. Young & Stanton (2001, 2002) found that AS significantly reduced driver mental workload, and the consequent underload led to performance problems when the driver needed to reclaim control. Nilsson (1995) also found that many drivers failed to

reclaim control from ACC in certain emergency scenarios. Similarly, Landau (2002) suggested that a proliferation of driver support systems could overload the driver, thus nullifying any stress or satisfaction benefits of each individual system.

Hard driving automation, on the other hand, is largely associated with problems of trust, situation awareness, and mental models. If the system is designed to assume control with little input from or feedback to the driver, then the driver may have difficulty in developing an appropriate mental model of its operation in a given scenario. Without knowing exactly how it might behave, the driver could become distrustful of the system (i.e., lack of trust) or even develop misplaced trust (i.e., overtrust or complacency; cf. Parasuraman & Riley, 1997). Lee & See (2004) argue that it is critical for a human operator's level of trust to be calibrated according to the genuine capabilities of the technology to avoid misuse and disuse of automation.

Although we have again been somewhat coarse here, this analysis indicates that the range of cognitive and performance problems would seem to be more severe when implementing driving automation in the hard protection category. However, the cognitive factors involved when interacting with automation are all interdependent (Stanton & Young, 2000), and so it is difficult to make generalisations on this basis without detailed further research. The distinction between the relative merits of soft and hard protection is more blurred than in aviation, as significant problems can occur whether hard or soft, vehicle or driving automation. Unintended acceleration (hard vehicle automation) has also been observed with cruise control (soft vehicle automation), while Young (2004) postulated that ACC could cause problems of uncommanded acceleration. Furthermore, behavioural adaptation to ABS or ESP (hard vehicle automation) can also occur with ACC (soft driving automation). Rudin-Brown & Parker (2004) found that drivers actively relied on the ACC system to keep a safe distance from the vehicle in front, rather than using it as the 'comfort and convenience' system for which it was designed (cf. Richardson et al., 1997). Finally, the cognitive issues of driving automation seem to straddle both hard and soft systems, with potential safety consequences in both cases.

Thus, whilst soft automation is generally favourable, we may need to develop the philosophy one step further for automobile automation. Rather than revisiting existing levels of automation taxonomies (e.g., Parasuraman et al., 2000), we would like to suggest a whole new design philosophy to automation. Once again, we borrow from the aviation domain as inspiration for this approach. The concluding section of this paper summarises our thoughts and outlines this proposed design approach.

4   A third philosophy

Young & Stanton (1997) noted that ergonomics researchers are not going to stop the tide of vehicle technology – nor, indeed, would we probably want to. However, rather than asking whether we should, we can have an impact on how such technology is implemented. This paper has reviewed two philosophical stances to automation and has applied them to current and near-future vehicle technology systems. We have derived two classifications of such technology – vehicle automation, in which low-level vehicle control aspects are automated; and driving automation, in which the driver is relieved of higher-level tactical or strategic tasks.

Whilst we have noted that the distinctions are blurred between hard and soft automation on the one hand, and vehicle and driving automation on the other, when

focusing within driving automation it does seem that hard automation presents more problems than promises. Different levels of automation offer different levels of protection in different contexts. The driver needs to be aware of this, especially if mode changes are automatic. For instance, Honda's new Advanced Driver Assistance System (ADAS) links ACC with a Lane Keeping Assistance System (LKAS, similar to the AS discussed earlier). This combination keeps the vehicle centred in the roadway following a vehicle in front at a set distance without driver input to the steering wheel or throttle – effectively analogous to a flight management system (cf. Harris & Harris, 2004). In the ADAS system, three levels of automation are available to the driver: simple cruise control (automation has authority over the throttle); advanced cruise control, with automatic distance keeping to the vehicle in front (automation has authority over the throttle and brakes) and ADAS (where the vehicle has authority over throttle, brakes and steering). However, it becomes obvious that not all automatic protections are available in all modes and as in aviation, the automated systems only work within well-defined parameters even when selected – which arguably caused the A320 crash at Habsheim (as the protection system did not operate below 50 feet altitude). The ADAS system automatically transitions between modes, hence protections that the driver may think are in place are actually not available. Automatic mode transitions were one of the principle concerns identified in the implementation of automation on the flight deck (FAA, 1996).

Rather than an overarching philosophy of soft or hard automation for driving (as has been seen in aviation), a blend throughout the driving subtasks may prove most efficient. So, now we are thinking in terms of shared authority, rather than either human or technological authority. The FAA's (1996) recommendations emphasised the need for flightcrews to be informed about the relevance and authority of automation in different circumstances. In other words, humans are being trained to compensate for the automation's failure to act as a team player. Dekker (2004) makes a similar point, noting that it is not the quantity of automation which causes the problem, but the quality. Instead of designing automation on a 'who does what' basis, successful automation depends on designers answering the question of 'how do we get along?' On the flight deck, such teamwork comes under the umbrella of Crew Resource Management (CRM; Wiener et al., 1993). Rather than thinking of CRM purely in terms of liveware-to-liveware or liveware-to-software (cf. Fitzgerald, 1997), though, why not also invoke it for software-to-liveware scenarios? After all, as Jensen (1997) points out, CRM is all very well, but the first line of defence should be the design of the system, not the flight crew.

In many ways, we are back to Norman's (1990) problem of feedback, and the automation not acting as a human co-pilot (or co-driver) would. If we are expecting the automation to behave as a team member – coordinating and cooperating with the driver – then we should apply notions of team performance to automation design. That means effective communications, group processes, team decision making, leadership, shared situation awareness, conflict resolution, and recognition of others' behavioural styles (Jensen, 1997). Christoffersen & Woods (2000) suggest that the more powerful automated systems become (i.e., high in autonomy and authority), the more feedback they need to supply to make their behaviour observable. From the human's perspective, if the capabilities and activities of automation are more transparent (as would be expected of a human colleague), then the problems of mental models, workload and trust should be ameliorated. Indeed, models of trust in automation have been built upon models of interpersonal trust in humans (Muir, 1994), so it seems logical to apply another aspect of human-human cooperation to human-machine cooperation.

This broadly agrees with the general opinion towards technological support systems rather than automated replacement of the driver. Stanton & Young (2005) illustrate this with ACC, suggesting that instead of simply assuming the task of longitudinal control, the same radar technology could instead be used to provide enhanced information about time-to-contact – typically an area of variability and error in human judgement. In this way, we would be using the capabilities of the technology to exploit the human's strengths while compensating for their weaknesses – what Schutte (1999) refers to as 'complemation'. Moreover, this soft philosophy is much more aligned to a socio-technical systems perspective, whereby humans and technology cooperate as a team to achieve an overall goal. Of course, we have seen that the reality is not as simple as that, with soft automation causing problems of mental workload which can be equally detrimental to performance. In all likelihood, then, we will have to match different elements of the driving task with different philosophies for optimum performance. In a sense, we have already begun, with traditional vehicle automation mostly falling into the hard automation category, whilst more driving automation systems seem to fit the bill of soft protection.

Taking this one step further, if the automation can adapt its style to its human colleagues, the principle of complemation may be maintained. For instance, the impact of CRM training is partly determined by the personalities of participants (Helmreich & Wilhelm, 1991). One factor may be locus of control – whether people rely on external determinants of performance, or their own behaviour. Rudin-Brown & Parker (2004) demonstrated that drivers with an external locus of control were more likely to trust the ACC system, and this slowed their responses to a simulated failure. A human co-pilot (or driver) would be sensitive to such social dynamics and alter their behaviour accordingly, thus maintaining balance amongst the team. The automated co-pilot should be designed to adapt in the same way. In practice, this may mean revisiting some ideas of dynamic allocation of function or adaptive automation, which have their own set of problems (e.g. Scallen et al., 1995). Nevertheless, treating the automation as a team member may put a different perspective on the attitude to design, but is perfectly in keeping with the systems perspective of ergonomics – both human and machine are working in harmony towards a common goal.

Therefore it seems that on the face of it, soft automation fits much better with ergonomics principles and research than hard automation. In closing, let us revisit the problems of automation identified by Stanton & Marsden (1996), to see whether a 'soft' CRM-designed system would hypothetically solve them. The first problem of automatic systems is the shortfalls in expected benefits when introduced into the operational arena. In terms of vehicular automation, this could mean that they turn out to be less reliable (e.g., a collision avoidance system fails to detect an approaching object) or have an adverse impact upon human performance (e.g., automation seems to make the easy tasks boring and the difficult tasks even more difficult). An openly communicative automated system, designed to support drivers under a soft protection philosophy rather than replace them, would certainly help here by optimising situation awareness and mental workload. Stanton & Young's (2005) alternative perspective on ACC technology is a prime example here, and would even help drivers as a collision warning system with no further adjustments.

Next, automatic systems can have problems related to equipment reliability. In terms of vehicular automation, this could mean that drivers lose their trust in the automated systems (e.g., the driver prefers to choose the manual alternative), intermittent faults

could go undetected until the context becomes critical (e.g., the failure reveals itself immediately prior to the vehicle impacting at high speed into another vehicle) and the driver becomes so dependent upon the automated systems that they operate them beyond design limits (e.g., invoking ACC in non-motorway situations). Whilst technical reliability is improving all the time, it would be unrealistic to expect a complemation system to be reliable 100% of the time. Nonetheless, consistent communications would help the human to detect and intervene in a more timely fashion if and when the system fails.

Thirdly, automatic systems seem to lead to problems related to training and skills maintenance. In terms of vehicular automation, this could mean that driving skills could be stripped away through lack of practice by automation being in control. This is likely to make the driver even more dependent upon the automated systems. If drivers are not performing a function, how can they be expected to take it over adequately when the automated systems fail to cope? This is where a CRM system comes into its own, as it plays to the strengths of each member of the team (whether human or machine), and can actually serve to develop the skills of the human by fostering an error-tolerant environment in which to learn (cf. Ivancic & Hesketh, 2000).

Finally, automatic systems seem to induce errors in users. In terms of vehicular automation, this could mean that design flaws lead to driver errors when interacting with the automated systems, for example specifying the wrong target speed and distance with ACC. Of particular concern is the possible introduction of mode errors, which are most likely when controls have more than one function and the system's mode is not transparent. The A320 crash at the Paris air show was a prime example, and one which could have been prevented if both human and machine were more aware of each other's intentions. Once again, then, communication is the key, as we move from 'strong, silent' technology towards a cooperative, supportive system.

From our analysis of automation in aviation, the problem of lack of coordination has emerged as a central issue (cf. Dekker, 2004). We propose that allocation of function and automation design needs to explicitly examine coordination and cooperation between human and automated subsystems if the problems of automation in aviation are not to be replicated through automation in automobiles.

5   References

Aircraft Accident Investigation Commission - Ministry of Transport Japan (1996). China Airlines Airbus Industrie A300B4-622R, B1816 Nagoya Airport April 26, 1994. Tokyo: Ministry of Transport.

Applegate, J.D. & Graeber, R.C. (2005). Integrated System Safety Design and Human Factors Considerations for Jet Transport Aeroplanes. In, D. Harris and H.C. Muir (Eds.), Contemporary Issues in Human Factors and Aviation Safety (pp. 3- 24). Aldershot: Ashgate.

Bainbridge, L. (1982). Ironies of automation. In G. Johannsen et al. (Eds.), Analysis, design and evaluation of man-machine systems (pp. 151-157). Pergamon.

Beaty, D. (1995). The naked pilot: The human factor in aircraft accidents. Shrewsbury: Airlife Publishing Ltd.

Billings, C. (1993). Aviation Automation. New Jersey: Lawrence Erlbaum.

Carsten, O. M. J. & Tate, F. N. (2005). Intelligent Speed Adaptation: Accident Savings and Cost Benefit Analysis. Accident Analysis and Prevention, Vol. 37, 407-416.

Christoffersen, K., & Woods, D.D. (2000). How to make automated systems team players. Columbus, OH: Institute for Ergonomics, The Ohio State University.

Dekker, S. (2004). On The Other Side of a Promise: What Should We Automate Today? In D. Harris (Ed.), Human Factors for Civil Flight Deck Design (pp. 183-198). Aldershot: Ashgate.

Fitzgerald, R. E. (1997). Call to action: We need a new safety engineering discipline. Professional Safety, Vol. 42(6), 41-44.

Harris, D. & Harris, F.J. (2004). Predicting the successful transfer of technology between application areas; a critical evaluation of the human component in the system. Technology in Society, Vol. 26(4), 551-565.

Harris, D. & Smith, F.J. (1997). What can be done versus what should be done: a critical evaluation of the transfer of human engineering solutions between application domains. In D. Harris (Ed.), Engineering Psychology and Cognitive Ergonomics Volume One - Transport Systems (pp. 339-346). Aldershot: Ashgate.

Helmreich, R. L. & Wilhelm, J. A. (1991). Outcomes of Crew Resource Management Training. International Journal of Aviation Psychology, Vol. 1(4), 287-300.

Hollnagel, E. (1993). Context and Control. New York: Academic Press.

Hughes, D., & Dornheim, M. A. (1995, January 30). Accidents direct focus on cockpit automation. Aviation Week and Space Technology, 52-54.

Investigation Commission of Ministry of Transport (1989). Final report concerning the accident which occurred on June 26th 1988 at Mulhouse-Habsheim (68) to the Airbus A 320, registered F-GFKC. Paris: Ministry of Planning, Housing, Transport and Maritime Affairs.

Ivancic, K. & Hesketh, B. (2000). Learning from errors in a driving simulation: effects on driving skill and self-confidence. Ergonomics, Vol. 43(12), 1966-1984.

Jensen, R. S. (1997). The boundaries of aviation psychology, human factors, aeronautical decision making, situation awareness, and crew resource management. International Journal of Aviation Psychology, Vol. 7(4), 259-267

Kaber, D. B., & Endsley, M. R. (1997). Out-of-the-loop performance problems and the use of intermediate levels of automation for improved control system functioning and safety. Process Safety Progress, Vol. 16(3), 126-131.

Kazi, T. A., Stanton, N. A., Young, M. S., & Harrison, D. A. (2005). Assessing drivers' level of trust in Adaptive Cruise Control and their conceptual models of the system: implications for system design. In L. Dorn (Ed.), Driver Behaviour and Training Volume 2. Aldershot: Ashgate.

Labiale, G. (1997). Cognitive ergonomics and intelligent systems in the automobile. In Y. I. Noy (Ed.), Ergonomics and safety of intelligent driver interfaces. (pp. 169-184). Mahwah, NJ: Lawrence Erlbaum Associates.

Landau, K. (2002). The development of driver assistance systems following usability criteria. Behaviour & Information Technology, Vol. 21(5), 341-344.

Lee, J. D. & See, K. A. (2004). Trust in automation: Designing for appropriate reliance. Human Factors, Vol. 46(1), 50-80.

Mellone, V.J. (1993). Genie out of the bottle? ASRS Direct Line, Issue 4 (June). http://asrs.arc.nasa.gov/directline_issues/dl4_tcas.htm (Accessed 15 May 2006).

Moray, N. & Rotenberg, I. (1989). Fault management in process control: eye movements and action. Ergonomics, Vol. 32(11), 1319-1342.

Muir, B. M. (1994). Trust in automation: Part I. Theoretical issues in the study of trust and human intervention in automated systems. Ergonomics, Vol. 37(11), 1905-1922.

Nilsson, L. (1995). Safety effects of adaptive cruise control in critical traffic situations. Proceedings of the second world congress on intelligent transport systems: Vol. 3 (pp. 1254-1259).

Norman, D. A. (1990). The 'problem' with automation: inappropriate feedback and interaction, not 'over-automation'. Phil. Trans. R. Soc. London B, Vol. 327, 585-593.

Parasuraman, R. & Riley, V. (1997). Humans and automation: use, misuse, disuse, abuse. Human Factors, Vol. 39(2), 230-253.

Parasuraman, R., Sheridan, T. B. & Wickens, C. D. (2000). A model for types and levels of human interaction with automation. IEEE Transactions on Systems, Man, and Cybernetics – Part A: Systems and Humans, Vol. 30(3), 286-297.

Ranney, T. A. (1994). Models of driving behavior: a review of their evolution. Accident analysis and prevention, Vol. 26(6), 733-750.

Reason, J. T. (1990). Human Error. Cambridge: Cambridge University Press.

Richardson, M., Barber, P., King, P., Hoare, E., & Cooper, D. (1997). Longitudinal driver support systems. Proceedings of Autotech '97 (pp. 87-97). London: IMechE.

Rudin-Brown, C. M. & Parker, H. A. (2004). Behavioural adaptation to adaptive cruise control (ACC): Implications for preventive strategies. Transportation Research Part F: Traffic Psychology and Behaviour, Vol. 7(2), 59-76

Sarter, N.B. & Woods, D.D. (1992). Pilot interaction with cockpit automation: Operational experiences with the flight management system. International Journal of Aviation Psychology, Vol. 2(4), 303-321.

Sarter, N. B., & Woods, D. D. (1994a). Decomposing Automation: Autonomy, Authority, Observability and Perceived Animacy. In, R. Parasuraman and M. Mouloua. (Eds.), Proceedings of the 1st Automation technology and human performance conference (Washington D.C.) Hillsdale, NJ: Erlbaum Associates.

Sarter, N.B. & Woods, D.D. (1994b). Pilot interaction with cockpit automation: II. An experimental study of pilots' model and awareness of the flight management system. International Journal of Aviation Psychology, Vol. 4(1), 1-28.

Sarter, N.B. & Woods, D.W. (1995). How in the world did we ever get into that mode? Mode error and awareness in supervisory control. Human Factors, Vol. 37(1), 5-19.

Sarter, N.B. & Woods, D.D. (1997). Team play with a powerful and independent agent: Operational experiences and automation surprises on the Airbus A-320. Human Factors, Vol. 39(4), 553-569.

Sarter, N.B. & Woods, D.D. (2000). Team play with a powerful and independent agent: A full-mission simulation study. Human Factors, Vol. 42(3), 390-402.

Scallen, S. F., Hancock, P. A., & Duley, J. A. (1995). Pilot performance and preference for short cycles of automation in adaptive function allocation. Applied Ergonomics, Vol. 26(6), 397-403.

Schmidt, R. A. (1993). Unintended acceleration: Human performance considerations. In B. Peacock & W. Karwowski (Eds.), Automotive Ergonomics (pp. 431-451). London: Taylor & Francis

Schutte, P. (1999). Complemation: An alternative to automation. Journal of Information Technology Impact, Vol. 1(3), 113-118.

Stanton, N. A., & Marsden, P. (1996). From fly-by-wire to drive-by-wire: safety implications of automation in vehicles. Safety Science, Vol. 24(1), 35-49.

Stanton, N. A. & Young, M. S. (2000). A proposed psychological model of driving automation. Theoretical Issues in Ergonomics Science, Vol. 1(4), 315-331.

Stanton, N. A. & Young, M. S. (2005). Driver behaviour with Adaptive Cruise Control. Ergonomics, Vol. 48(10), 1294-1313.

Walker, G.H., Stanton, N.A., & Young, M.S. (2001). Where is computing driving cars? International Journal of Human Computer Interaction, Vol. 13(2), 203-229.

Wickens, C. D., Gordon, S. E. & Liu, Y. (1998). An introduction to human factors engineering. New York: Longman.

Wiener, E. L., & Curry, R. E. (1980). Flight-deck automation: promises and problems. Ergonomics, Vol. 23(10), 995-1011.

Wiener, E. L., Kanki, B. G. & Helmreich, R. L. (1993). Cockpit resource management. San Diego: Academic Press.

Young, M. S. (2004). I thought you were driving! A story about vehicle automation. Ergonomics Australia, Vol. 18(3), 16-19.

Young, M. S. & Stanton, N. A. (1997). Automotive automation: Investigating the impact on drivers' mental workload. International Journal of Cognitive Ergonomics, Vol. 1(4), 325-336.

Young, M. S. & Stanton, N. A. (2001). Size matters. The role of attentional capacity in explaining the effects of mental underload on performance. In D. Harris (Ed.), Engineering Psychology and Cognitive Ergonomics: Vol. 5 – Aerospace and Transportation Systems (pp. 357-364). Aldershot: Ashgate.

Young, M. S. & Stanton, N. A. (2002). Malleable Attentional Resources Theory: A new explanation for the effects of mental underload on performance. Human Factors, Vol. 44(3), 365-375.