



HAL
open science

Automation Effects on Driver's Behaviour When Integrating a PADAS and a Distraction Classifier

Fabio Tango, Lucas Minin, Raghav Aras, Olivier Pietquin

► **To cite this version:**

Fabio Tango, Lucas Minin, Raghav Aras, Olivier Pietquin. Automation Effects on Driver's Behaviour When Integrating a PADAS and a Distraction Classifier. HCI 2011, Jul 2011, Orlando, Florida, United States. pp.503-512, 10.1007/978-3-642-21799-9_56 . hal-00618282

HAL Id: hal-00618282

<https://hal-centralesupelec.archives-ouvertes.fr/hal-00618282>

Submitted on 1 Sep 2011

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Automation Effects on Driver's Behaviour when integrating a PADAS and a Distraction Classifier

Fabio Tango, Luca Minin, Raghav Aras, Olivier Pietquin

Centro Ricerche Fiat – Strada Torino, 50 10043 Orbassano (Italy)
University of Modena & Reggio Emilia – via Amendola, 02 42122 Reggio Emilia (Italy)
IMS Supélec, Metz, France
fabio.tango@crf.it
luca.minin@unimore.it
raghav.aras@gmail.com
olivier.pietquin@supelec.fr

Abstract. The FP7 EU project ISi-PADAS aims at conceiving an intelligent system, called PADAS, to support drivers, which intervenes continuously from warning up to automatic braking in the whole longitudinal control of the vehicle. However, such supporting systems can have some unwanted side-effect: due to the presence of automation in the driving task, less attention and reaction are needed by the drivers to intervene in the longitudinal control of the vehicle. Such a paper aims at investigating the effects of the level of automation on drivers, in particular on their Situation Awareness, when the user is supported by a specific PADAS application, integrated with a driver's distraction classifier.

Introduction

The FP7 EU project ISi-PADAS (Integrated Human Modelling and Simulation to support Human Error Risk Analysis of Partially Autonomous Driver Assistance Systems) aims at conceiving an intelligent system called PADAS (Partially Autonomous Driver Assistance System) which aids human users to drive safely, by providing them with pertinent and accurate information in real time about the external conditions and by acting as a co-pilot in emergency situations. The system interacts with the driver through a Human-Machine Interface (HMI) installed on the vehicle using an adequate Warning and Intervention Strategy (WIS). Such a system intervenes continuously from warning up to automatic braking in the whole longitudinal control of the vehicle [1]. Recent data have identified inattention – and in particular distraction – as the primary cause of accidents [3]. The ISi-PADAS project has developed a model, able to detect and classify driver's distraction, which has then been integrated into the PADAS, in order to avoid or mitigate the negative effects and make the system “smarter” and more adaptive [2].

This paper addresses the problem of evaluating the effects of the automation on driver's Situation Awareness when using a PADAS integrating a Distraction (DIS) classifier.

The Context and the Problem Addressed

Systems like the Adaptive Cruise Control (ACC) can (partially) automate the longitudinal driving task and thus support users in this perspective, but – on the other way – they can cause also less attention and reduced reaction time from drivers, especially in case of unexpected events. Ward found that ACC appeared to improve driving safety by reducing instances of unsafe headway distance in following tasks [5]. However, he also highlighted a relevant impairment of *situation awareness* (SA) because of a reduced attention dedicated to positioning and slower response times to unexpected events. In addition, other studies have proved the influence of ACC automation on performance, workload and attention allocation, e.g. [6]. Despite of the indubitable advantages, drivers may direct their attention away from the driving task when using ACC, creating an unsafe situation due to a *loss of situation awareness*. The study of Rudin-Brown confirms such results, with an unexpected increase in accidents from driver distraction due to the interaction with the secondary task [7].

The PADAS Concept

The choice to focus the attention on longitudinal driving control task derived by a preliminary in-depth accident analysis, conducted to derive hypotheses about causes of driver's errors responsible for crashes [8]. In this context, many studies have shown the benefits of supporting systems in reducing accidents [2]. Therefore, the project has developed a PADAS called "*Advanced Adaptive Cruise Control*" (ACC+) with the following functionalities. In normal condition, the system can use the "traditional" ACC. Then when the driver acts on the brake pedal, thus indicating the will to brake, the AB function is able to modulate the braking action automatically. Finally, if the driver ignored warning and AB did not intervene on the brakes, EB acts autonomously in order to minimise the effects in case this is not avoidable anymore.

The MDP Approach

Consider a system that assists the driver by applying corrective controls to the vehicle (giving the driver warning signals and intervening in the braking) so that collisions may be avoided. When a collision is imminent, there is no doubt about what controls to apply: a clear collision warning must be given to the driver, and brakes must be applied as hard as possible to bring the host vehicle to a halt as quickly as possible. When there is a safe distance between the two vehicles, it is not clear what the optimal controls are. Should no warnings, no braking interventions occur? It could be, for example, that not giving a warning when the vehicle is at a certain "safe distance" leads (eventually) to a situation where a collision is unavoidable whereas a warning leads to an avoidance of the eventual collision. Deciding what corrective controls to exercise in order to avoid an eventual collision is essentially a problem of "credit assignment": suppose an outcome is a consequence of a sequence of decisions. How do we decide what part each of decisions plays in the outcome? In other words, the credit assignment problem calls for a system for associating decisions to their

long-term outcomes. One of the most important theories for formulating and solving credit assignment in sequential decision-making problems is the Markov decision process (MDP) theory ([2] for details). In modelling a problem as an MDP, we contemplate a decision-maker who is required to take decisions over a sequence of discrete time periods. In each period, the problem occupies one of N possible states and the decision maker is called upon to choose from K possible alternatives. Such a decision and the state of the problem together determine, according to fixed probabilities P , the state occupied by the problem in the next period. Each choice is evaluated by an immediate *cost* c which is a function only of the state in which the choice is made (*Markovian Property*). The immediate cost tells the decision maker how good or bad a choice is in the short turn (i.e., its immediate consequence). It can use the immediate costs to determine the long-term desirability of a sequence of decisions. The objective of the decision maker to make such decisions that *minimize* his long-term *expected* cost. The solution to a sequential decision making problem is of course a *sequential decision making policy*. A policy is a function which assigns to each state an alternative. If the decision maker is following the optimal policy, it is the case that in any period, whatever be the state occupied by the problem, the decision the policy dictates will lead to long-term expected cost that is as small as or smaller than any other decision he could have made [15]. Thus the MDP approach applied to conceiving an optimal PADAS consists of two elements: modelling the problem as an MDP; solving the (possibly unknown) MDP.

Modelling a PADAS Target Application as an MDP

First, the state of the MDP in the collision avoidance problem has to be defined. In this problem, the decision maker is the warning and intervention system (of the PADAS). The relevant factors that determine if a collision with the leading vehicle can occur are: time to collision; headway; velocities of the host and leading vehicles; accelerations of the host and leading vehicles; brake and gas pedals position; driver's distraction level. So, the state in the proposed MDP can be considered as a 10-vector. However, the number of states of such an MDP would be too large for solving the MDP. If each element of the state is allowed to take only 3 possible values (a very conservative estimate, indeed most elements would take *continuous* values), we would end up with 3^{10} states!

For the sake of computational tractability, we must determine which of these elements contains *most* information relevant to collision avoidance. With this in mind, the time to collision¹ and driver's distraction variables are two good candidates.

Distraction Model as an MDP State to improve the PADAS

A distraction classifier has been developed in the ISI-PADAS project and integrated into the optimal WIS for the PADAS, with the aforementioned MDP

¹ defined as the time it would take for a collision to occur if the leading and follower vehicles maintain their current velocities, that is the distance rated by their relative speed.

approach. The distraction model has been implemented following also in this case a Machine Learning approach, which constitutes a quite common method to detect human visual distraction in a non-intrusive way, based on the vehicle dynamic data ([9], [10]). Concerning the distraction classifier developed by the project, a broad explanation is outside the scope of this paper, but some details are provided as following (the interested reader can see [11]). In particular, a Support Vector Machines based model (SVM, well-known data-mining methods) has been selected, since achieving the best performances [11]. So, to sum up, the PADAS is conceived as an MDP, including an SVM-distraction classifier in the states defining the MDP.

The experimental Set-up Description

The scope of this experiment was twofold. On one hand, to assess the effect of automation on driver's situation awareness; on the other hand, to compare the effects of the ACC+ solutions in reducing the risk of collision and on system acceptability from a driver perspective. Therefore the following phases have been considered:

- ACC+ with a starting (initial) MDP policy p_0 including a model for drivers distraction;
- ACC+ with a MDP running p^* policy (that is, the optimal policy developed by solving the MDP problem), including a model for driver distraction detection
- ACC+ with a policy based on deceleration ($pDec$ policy, described in [25]);
- Without supporting systems (reference line).

Concerning the reference policy $pDec$, it provides the driver with the same kind of information of the p^* policy, but the approach is different. The main idea of $pDec$ is to give the needed acceleration for the host vehicle, in order to reach the same speed of the lead vehicle within a certain distance. Then, if the value is over certain thresholds, a corresponding warning or action is provided [25]. This policy represents a traditional algorithm for the implementation of a longitudinal distance control/warning system.

For each of the above configurations, ten participants in the age between 21 and 45 years ($M = 31.0$, $SD = 9.0$) were asked to complete 12 driving sessions: each condition consisted in a car-following task on a highway road where unexpected braking events of the leading car have been reproduced. Hence, each of the 10 subjects performed 36 driving sessions, divided as follows:

- 12 with the ACC+ working with the p^* policy;
- 12 with the ACC+ working with the $pDec$ policy;
- 12 without any supporting system

Two different levels of road visibility (low, LV; high, HV) have been reproduced, combined with the activation of the secondary task and the level of automation of the vehicle longitudinal control. LV (200m) has been introduced with fog effect to force low situation awareness situations. LV condition has been compared with the HV condition (7000m).

The secondary task adopted is the SuRT (Surrugate visual Research Task [24]). Drivers were asked to interact with SuRT every time the task was presented on the lateral (right hand side) touch screen installed in the simulator cabin. A high level of SuRT difficulty has been set: high number of distractors with a reduced time to task completion (i.e. the SuRT screen was updated each 3-6 seconds). The secondary task has been activated in 6 conditions out of 12 for each driver.

SA has been measured at the end of each session by means of a self reporting questionnaire called SART (Situation Awareness Rating Technique [23]) submitted to participants: we expected a decrease of situation awareness when scenarios are critical (i.e. low visibility and secondary task on) and the level of automation is low.

The number of occurred collisions has been also calculated, to evaluate the effects of highly automated longitudinal control in particular while dealing with low situation awareness conditions. Finally, in order to verify the influence of the different ACC+ system automation on performance, workload and attention allocation ([8] and [9]), additional investigations have been conducted on specific indicators: NASA-Tlx and Steering Entropy (SE).

The experiment has been conducted on the static driving simulator of the University of Modena and Reggio Emilia (see www.scaner2.com).

Data Processing and Results

During the experiment, several data on drivers' behaviour were recorded from the simulator network at a sample rate of 20Hz (1 data-point each 0.05s).

Indicators and Indexes of Performance

SA has been measured by means of the SART questionnaire. It is a quick and easy self-rating measurement technique [23]. As a result of the interviews provided by the questionnaire, 10 dimensions were derived and used to measure SA. For each basic driving condition (4 in total) and for each driving session (3 in total, 2 with ACC+ and 1 without) the following measures have been analysed:

- *Number of accidents occurred*, allowing the evaluation of drivers' longitudinal performances with reference to the level of longitudinal automation and the scenario complexity
- *Steering Entropy* [28] has been computed to evaluate drivers' impairment in lateral control for each of the above mentioned situations. SE is a measure of randomness in a driver's steering control: it is higher when drivers make larger erratic steering movements, indicating potentially unsafe behaviour.
- *Nasa-Tlx* [29] has been used to analyse, in order to have a general understanding of the level of cognitive and physical effort required by drivers while coping with concurrent tasks (car-following and secondary task interaction).

Finally, participants were asked to fill up the subjective questionnaire defined by the ISI-PADAS consortium to assess both ACC+ running p^* and ACC+ running $pDec$ policies.

Results for Automation Effects on Driver's Behaviour

SART questionnaire has been submitted to each subject four times for each of the three automation conditions. The aim of the evaluation conducted on collected data was to compare the average values of SART indicator of the aggregated SART scales *between subjects* for each level of automation and basic driving conditions. According to [23] the 10-scales of the SART have been aggregated: *Demands* on attentional resources (D), *Supply* of attentional resources (S) and *Understanding* of the situation (U).

SART indicator is computed as: $U-(D-S)$, where U, D, and S are the sum of the answers to the SART scales belonging to these aggregations:




Fig. 1. Average values of SART computed on the whole sample

In the figure, it is shown the quantitative variation of situation awareness among the driving (x-axis) and the automation conditions. Variations among automation conditions are relevant between automated (ACC+) and not-automated situations and there is a general positive increment of SA for the former conditions. Moreover, it is possible to see the effect of the driving conditions (SuRT and visibility) on the perceived SA: the absence of the secondary task and fog clearly increases the level of perceived situation awareness. A *t-Test* analysis of SART values was conducted for each driving condition and (in pairs) between the three automation configurations, in order to evaluate whether there are significant differences between the correspondent SART levels. Significant differences ($p < 0.05$) between average values of SART *with* and *without* automation have been found for all driving conditions but the “SuRT + Visibility low” one. All in all, no significant differences were found comparing the average values of SART computed for the ACC+ p^* and ACC+ $pDec$.

Concerning the lateral driving behaviour, the average values of steering entropy (SE) were computed for all subjects for the two automation types: ACC+ with p^* and ACC+ with $pDec$. As anticipated, SE revealed steering impairments due to factors influencing the vehicle lateral control. In this case, these factors are the interaction with SuRT and the low visibility. The *t-Test* analysis conducted on SE values showed that there is a significant statistical difference ($p < 0.05$) in SE performances between

driving conditions (x-axis of the next figure) for each separated automated condition. SE significantly increased with the presence of the secondary task, if compared to the conditions without the SuRT (see the following figure):




Fig. 2. SE variations in the 4 driving conditions, w.r.t. the two kinds of automation

For each driving condition, again a *t-Test* has been also conducted to compare the average values of SE in the two automated conditions to highlight potential statistical differences. No statistical differences between these values, however, were found in each of the driving conditions.

Finally, the automation effects on driver's workload are analyzed. The NASA-Tlx questionnaire has been submitted to each participants after each of the last four driving conditions for each driving session (ACC+ *p**, ACC+ *pDec* and without supporting systems). The following figure shows the average values of NASA-Tlx computed for each driving condition (x-axis) and for each automation type:




Fig. 3. NASA-Tlx average values

In general, the level of workload perceived by drivers was higher in the conditions where SuRT was active, lower when no secondary task was present and visibility was high. Significant differences among workload levels measured for each automation type in the four driving conditions were not found, except for the "SuRT OFF + Visibility high" where participants stated that driving without longitudinal assistance induced high workload. In this case, the NASA scales showing this pattern are physical effort (because drivers have to modulate the speed autonomously) and effort (physical and mental demand).

Moreover, the level of NASA-Tlx was higher in the ACC+ *pStar* conditions if compared to ACC+ *pDec*: the NASA scale contributing to the increase of the level of workload perceived by drivers is "Frustration". Drivers stated they were annoyed by a continuous warning signal provided by the ACC+ *pStar* also in conditions where they

did not perceive a risk of collision and where they were able to modulate the speed by themselves in case that more deceleration was needed.

Results for PADAS Performances

Some preliminary results on the performances of PADAS applications, integrating the Distraction classifier are pointed out in this section. Figure 4 (top) shows the effects of the type of longitudinal automation (ACC+ with p^* , ACC+ with $pDec$) compared to the control condition (no automation, No-PADAS) in the four basic driving conditions:




Fig. 4. Total collision using different PADAS configurations: cumulative collisions for all subjects in the top figure; collisions occurred for each subject in every driving session

The benefits of the ACC+ running the p^* policy are evident: besides some user-interaction issues discussed in the previous sections, the WIS, developed following an MDP approach, allows drivers to intervene in time on the brakes and avoid collision with the lead vehicle when it suddenly stopped. All in all, drivers found ACC+ with p^* policy more cautious, in the sense that the system provided warning signals also in conditions where the lead vehicle braking was smoother than the sudden one and where they did not perceive an imminent risk of collisions. This “learning effect” of the system induced drivers to be prepared to brake before the sudden stop event. Figure 4 (bottom) also reveals drivers impairment, for each driving condition, in coping not only with longitudinal control but also with secondary tasks activation and road visibility (i.e. conditions where SA was critical, in particular without the support of longitudinal automation).

Discussion and Conclusions

This paper has presented the evaluation of the impact on driver’s behaviour, caused by supporting systems with high level of automation. In addition, also the system performances have been illustrated. We have focused on the following aspect: situation awareness, lateral behaviour and workload. Concerning the automatic system, we have developed the PADAS application following a MDP approach. Experiments have been carried out in a static driving simulator.

Concerning SA, the SART questionnaires showed significant benefits between automated and not automated driving conditions. In the former, higher values of

situation awareness reveal a relevant impact of the ACC+ (indifferently p^* or $pDec$) in providing the driver with a better understanding of the situation, particularly while dealing with concurrent tasks like monitoring the driving environment and executing secondary task. Effects of the ACC+ p^* and $pDec$ on lateral vehicle control have been also assessed using the SE indicator, revealing that the two automation solutions did not generate different lateral control impairments while significant different levels of lateral impairments have been measured among driving conditions because of the presence of the secondary task and low visibility conditions.

Drivers' feedbacks on the different NASA scales revealed an overall workload increase in conditions where the secondary task was active and the visibility was obstructed, as expected. Within these conditions, we did not found relevant workload differences among the two ACC+ solutions, but for most of these conditions driver workload was higher in the conditions without automation than with ACC+. This can be linked to the difficulty of performing secondary task like SURT in a harsh environment (low visibility conditions). So, we can conclude that high-automated systems, such as ACC+, did not reduce neither the impairment of drivers in controlling the vehicle, nor their situation awareness; moreover, driver's workload did not enhance, due to system use.

Another very interesting aspect was the evaluation of PADAS performances, in particular considering its efficacy in avoiding extremely dangerous situations. From preliminary results presented here, considering the number of collisions occurred when using ACC+ with p^* or $pDec$, compared with the situation when no supporting system is used, results pointed in favour of the ACC+ p^* system. In fact, this solution proved to be the most effective in avoiding accidents, if compared to the numbers measured for the other automation solution and for the driving condition without automation. The major benefit introduced by the ACC+ p^* was the cautious warning information and assistance strategy that allowed the driver to anticipate a potential sudden brake of the lead car.

Although these good results, from internal acceptance questionnaires, drivers were annoyed by this cautious strategy providing AB signals (hap-tic + acoustic in particular) frequently along the test session.

This research represents a feasibility study for an innovative approach to PADAS applications development, based on MDP and integrating a SVM driver's distraction classifier. Of course, all these methods, taken separately, are not new in literature, but at the best of authors' knowledge, their use in this kind of domain and with this type of integrated approach is quite innovative. With respect other traditional algorithms used for implementing the WIS of ADAS applications (i.e. $pDec$), our approach is based on learning the "behaviour" directly from the interaction between driver and system (vehicle + automation) in the specific environment; therefore we regard it as more appropriate to provide the right information in the right time (this has been proved by the drastic reduction in the number of collisions). On the other hand, such a system needs to be refined, since users assessed it as too conservative; however, due to that direct interaction between human and system, we are strongly confident to tune appropriately this kind of PADAS. So, next steps will involve the refinement of cost function, a more precise definition of the parameters constituting the state of the MDP, new possible technique for the MDP problem solution and – last but not least –

a deeper investigation and integration of the distraction model into the WIS of the system.

The authors would like to specially thank the ISi-PADAS consortium that has supported the development of this research.

References

- [1] J. Alsen et al. (2010). Integrated Human Modelling and Simulation to support Human Error Risk Analysis of Partially Autonomous Driver Assistance Systems: the ISi-PADAS Project. Transport Research Arena Europe, Brussels 2010
- [2] Tango F., Aras R. and Pietquin O. (2010, June-July, 30-02). Learning optimal control strategies from interactions for a Partially Autonomous Driver Assistance System. In the Proceeding of I Workshop on Human Modelling in Assisted Transportation (HMAT). [to be published by Springer] Belgirate, Italy
- [3] Beirness D.J. et al.: The Road Safety Monitor: Driver Distraction. Traffic Injury Research Foundation. Ontario, Canada (2002)
- [4] Wang J., Knipling R. R., and M. J. Goodman, "The role of driver inattention in crashes: New statistics from the 1995 crashworthiness data system (CDS)," in *Proc. 40th Annu.: Assoc. Advancement Automotive Med.*, 1996, pp. 377–392
- [5] Ward, N.J., (2000). Automation of task processed: an example of intelligent transportation systems. *Human Factors and Ergonomics in Manufacturing* 10 (4), 395–408
- [6] Parker, H.A., Malisia, A.R. & Rudin-Brown, C.M. (2003). Adaptive cruise control and driver workload. In: *Proceedings of XVth Triennial Congress of the International Ergonomics Association Conference (CD-ROM)*, Seoul, Korea, August 24–29, 2003.
- [7] Rudin-Brown, C.M., Parker, H.A. & Malisia, A.R. (2003). Behavioral adaptation to adaptive cruise control. In: *Proceedings of the 47th Annual Meeting of the Human Factors and Ergonomics Society. Human Factors and Ergonomics Society, Santa Monica, CA*
- [8] S. Briest and M. Vollrath; "In welchen Situationen machen Fahrer welche Fehler? Ableitung von Anforderungen an Fahrerassistenzsysteme durch"; In-Depth- Unfallanalysen. In VDI (Ed.), *Integrierte Sicherheit und Fahrerassistenzsysteme* (pp. 449 - 463). Wolfsburg: VDI (2006) – in German language
- [9] Liang Y., Reyes M.L. and Lee J.D.: Real-time Detection of Driver Cognitive Distraction Using Support Vector Machines. In *Proc. IEEE Transactions on Intelligent Transportation Systems*, Vol. 8, No. 2 – June 2007
- [10] Salvucci, D. D. (2005). Modeling tools for predicting driver distraction. In *Proceedings of the Human Factors and Ergonomics Society 49th Annual Meeting*. Santa Monica, CA: Human Factors and Ergonomics Society
- [11] Tango, F., Minin L., Montanari, R., & Botta, M. (2010, August, 16-19). Non-intrusive Detection of Driver Distraction using Machine Learning Algorithms. In the proceeding of the XIX European Conference on Artificial Intelligence (ECAI). Lisbon, Portugal
- [12] Taylor, R.M. (1989). Situational awareness rating technique (SART): The development of a tool for aircrew systems design. *Proceedings of the AGARD AMP Symposium on Situational Awareness in Aerospace Operations, CP478*. Seuilly-sur Seine: NATO AGARD
- [13] Mattes, S. (2003) The lane change task as a tool for driver distraction evaluation. In H. Strasser, H. Rausch & H. Bubb (Eds.), *Quality of work and products in enterprises of the future* (pp. 57-60). Stuttgart: Ergonomia Verlag
- [14] Saroldi, D. Bertolino & Sidoti, C. (1997). Driving in the Fog with a Collision Warning System: a Driving Simulator Experiment ITS'97 - 4th World Congress on Intelligent Transport Systems, Berlin, 21-24 October 1997
- [15] Pietquin O., Tango F. and Aras R. (2011). Batch Reinforcement Learning for Optimizing Longitudinal Driving Assistance Strategies. *IEEE SSCI 2011 Symposium Series on Computational Intelligence*, Paris (France), 11-15 April 2011.
- [16] Trochim, William M. (October 20, 2006). "Likert Scaling". *Research Methods Knowledge Base*, 2nd Edition.. Retrieved April 30, 2009
- [17] Nakayama, O., Futami, T., Nakamura, T., & Boer, E. R. (1999). Development of a steering entropy method for evaluating driver workload. Paper presented at SAE International Congress and Exposition, Detroit, Michigan, USA
- [18] Hart, S.G. & Staveland, L.E. (1988) Development of NASA-TLX (Task Load Index): Results of Empirical and Theoretical Research. In Hancock, P.A., & Meshkati, N. (Eds.), *Human Mental Workload* (pp. 239-250). Amsterdam: North Holland Press