Adaptive Artificial Co-pilot as Enabler for Autonomous Vehicles and Intelligent Transportation Systems

Elvio G. Amparore, Marco Beccuti, Marco Botta, Susanna Donatelli, Fabio Tango
{amparore, beccuti, botta, susi}@di.unito.it, fabio.tango@crf.it

Abstract

This paper illustrates the concept of "co-pilot" as an enabling technology for autonomous driving. A co-pilot system mixes the features of commercial Advanced Driver Assistance Systems (like blind spot, forward-collision warning, lane change assistant, overtaking assistant, and others) with human factors like driver distraction and intention. The co-pilot can provide a "suggested action" to the human driver through a dedicated Human-Machine Interface (a set of screens on the dashboard) or, alternatively, can be the enabling technology to build effective and user-friendly future intelligent transportation systems (i.e. Autonomous Driving Functions). We illustrate the results achieved by the European projects HoliDes and the next steps foreseen in the EU project AutoMate.

1 Introduction

A number of intelligent agents is entering our lives and supporting us in a wider variety of tasks; in particular, this is definitely true for automotive domain, where automation in passenger cars is constantly increasing. In fact, current roadmaps of car-manufacturers and suppliers predict automated vehicles on highways by 2020. The reasons for that are threefold:

- *Zero Emission*, with the reduction of fuel consumption and CO_2 emission, as well as traffic flow optimization.
- *Demographic Change*, including the support to unconfident drivers and the enhancement of mobility for elderly people.
- *Vision zero*, which is the potential for more driver support by avoiding human driving errors (on which our paper is more focused).

Research in Intelligent Transportation Systems (ITS) began in the late '80s with the PATH program in the US, the PROMETHEUS project in the EU and the ASV projects in Japan [Bishop, 2005]. Today, the development of highly automated driving is the research focus of many OEMs and research institutes, addressing the specific principles of smart collaboration between humans and systems (such as the studies of [Flemisch, 2003] and [Inagaki, 2008] and [Da Lio et al., 2015]), which may include full automation as one extreme point of the interaction spectrum. An overview can be found in [Li et al., 2012].

Open questions regarding highly automated vehicles include the strengthening of driver's sensing ability; the information in case of errors; and the reduction of the driving effort as well as an increased usability. Indeed human drivers are limited in recognizing, interpreting, understanding and operating in critical situations; moreover, they are prone to misbehaviors, drowsiness and distraction [HAVEit, 2013]. Nowadays, there are already on the market several ADAS (*Advanced Driving Assistance System*) applications (e.g. blind spot, lane departure, emergency braking, semi-automatic parking, etc.) that are designed either to automate specific tasks or to provide additional information to the driver.

This research presents an artificial agent, named *co-pilot*, which provides a unique adaptive framework for supporting the human driver during critical situations or, alternatively, it can be regarded as an enabling technology for the *Autonomous Driving Functions* (ADFs). Indeed, the co-pilot is the core of such ADFs, by computing a (sub) optimal maneuver that takes into account both the lateral and longitudinal tasks under a common view. In addition the co-pilot is adaptive, namely the decision accounts for critical human factors, i.e. an estimate of the driver status (visual distraction) and intention. This is crucial to make the system response closer to human needs: for example, if the system detects that the driver is distracted, then it avoids to suggest more demanding maneuvers (such as a take-over).

The paper is organized as follows. Section 2 describes the technology we have considered and we are still using, inside the EU aforementioned projects. Section 3 describes the system architecture we adopted. Section 4 gives an overview of the current preliminary results. Finally, section 5 ends the paper, by providing the conclusions and illustrating the next steps of our research.

2 Enabling Technologies

This section describes the methods and techniques that are used to realize the co-pilot, with a specific focus on the relevant human factors that enable the adaptation.

2.1 Models for driver intention recognition

Driver intention recognition is mainly concerned with the recognition of maneuver (e.g. lane change) intentions, and a comparative review of works on maneuver intention estimation can be found in [Börger, 2013], [Doshi and Trivedi, 2011], [Kobiela, 2011], [Lefèvre et al., 2014]. The modeling landscape is mostly based on *Dynamic Bayesian Networks* (DBNs), *Hidden Markov Models* (HMMs) and their variants, or probabilistic and non-probabilistic discriminative models.

Intention recognition based on DBNs is usually organized as follows: for each addressed maneuver, a distinct DBN is learnt from a sequence of observations (training). The DBN therefore models the dynamic evolution of the vehicle state and/or position for the specific maneuver. Given a new sequence of observations, the actual maneuver intention is estimated by comparing the likelihood of observations for each DBN (for details, see for example [Oliver and Pentland, 2000], [Kumagai and Akamatsu, 2006], [Liebner et al., 2012], [Tay, 2009]). For instance, in [Oliver and Pentland, 2000] seven distinct HMMs are used to recognize seven driving maneuvers, evaluating four different combinations of feature vectors (i.e. vehicle data, lane position, and driver gaze information). On average, the resulting models were able to recognize the addressed maneuvers one second before “any significant (20% deviation) change in the car or contextual signals” took place.

For intention recognition based on discriminative models, commonly used techniques are *Support Vector Machines* (SVMs), Multi-Layer Perceptrons or Logistic Regressions (see [P. Kumar et al., 2013], [Garcia-Ortiz et al., 2011]). To the best of our knowledge, the most sophisticated model implemented in a real vehicle up to date is the discriminative model described by [Morris et al., 2011]. They used Relevance Vector Machines for learning a model for online recognition of lane-change intentions, which can be seen as a Bayesian alternative to SVMs, in that they provide a probabilistic classification. The resulting model is able to predict lane change intentions of human drivers up to approx. three seconds prior to the actual crossing of the lane.

Due to the sound foundation of machine-learning methods and the direct interpretability of their structure and parameters, we use DBNs for modeling driver intention recognition used in the co-pilot. In contrast to the aforementioned approach, we refrain from modeling the dynamic evolution of the vehicles state and position for different maneuvers in favor of a more direct representation of the statistical relationships between driver intentions and the complex traffic situation. We believe that this approach will provide for an earlier recognition of intentions solely based on the current environment (traffic, speed, car position) without the need for a lane change to have already started (i.e. without relying on the light indicators).

2.2 Machine learning for distraction recognition

Driver distraction is a critical human factor with significant safety concerns [Regan et al., 2011]. Deriving knowledge on the human operator status can be very valuable for the operative system conditions. In this work we consider the following definition for driver’s distraction: “*the diversion of attention toward a competing activity, which may result in insufficient or no attention at all to activities critical for safe driving*”. Such a definition is quite general, and at the same time it allows us to capture the key elements of distraction, together with the important notion of insufficient or no attention being given to activities that are critical for safe driving. In practice, distraction can be split into visual and cognitive aspects. In this work we mainly consider the visual distraction, which is the diversion of attention toward a competing activity that requires the driver to look at a secondary target inside the vehicle instead of looking at the road.

In the literature several studies proved that visual distraction can be successfully inferred using Machine Learning (ML) approaches, that usually outperform other analytical methods (see [Liang et al., 2007] for more details). We investigated different ML techniques and, in particular we used neural networks. As for DBNs, they are learned from observations and used to classify new observations. *Single Layer Feed-forward Neural Networks* (SLFN) are the most common class of neural networks, where neurons are organized in stratified layers (input→hidden→output), and connections are weighted. SLFN training typically involves iterative algorithms, which perform some learning step aimed at minimizing the error function, over the space of network parameters. The *Extreme Learning Machine* (ELM) algorithm introduced in [Huang et al., 2006] works by training a neural network in a single step without using an iterative procedure. This notably reduces the computational cost while preserving a good generalization. With ELM, the output connection weights are determined by the Moore-Penrose generalized inverse (or pseudo-inverse) of the hidden layer output matrix. In particular, SLFN networks have been chosen because of their tradeoff between the implementation simplicity and their capacity to satisfy hard real-time constraints for the evaluation.

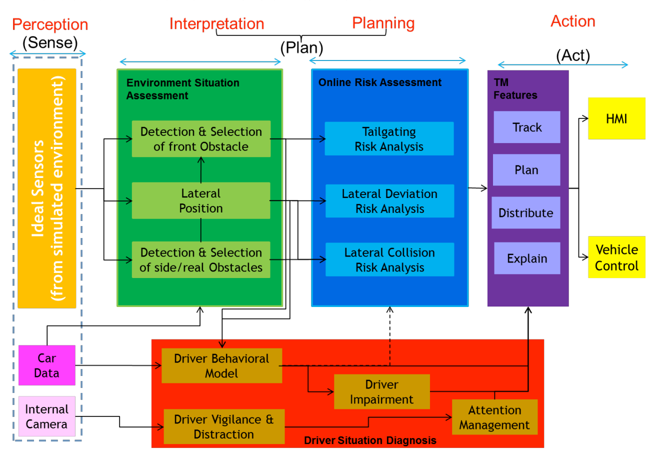

Figure 1: prototype architecture implemented in the demonstrator.

3 Implementation of the co-pilot

A critical aspect needed to design adaptive autonomous systems is the decision making task, which has to weight several possibly conflicting data sources in order to decide a safe driving plan.

3.1 System Architecture

Errore. L'origine riferimento non è stata trovata. shows the main building blocks of the car architecture, where the co-pilot manages the automated functions according to the situation and the driver needs, also taking into account the environment constraints. The central point for any automated systems is the ability to assess perception and decision performance under a given condition in a certain situation. With reference to the perception-cognition-decision process, as defined in [Stiller et al., 2007], input is received from sensors (considering several aspects and sources, e.g. internal camera for gestures and eye movements, from maps, from the environment and so on) over several processing steps via a geometrical-symbolic representation of the current traffic environment to the generation and control of suitable behavior. In this context, robustness is essential: one successful method to obtain it is to consider data-fusion from several sensors. This may happen on a sub-symbolic or symbolic level, in order to generate more robust hypothesis. Thereby, it is crucial to not only propagate knowledge through the cognition scheme but to augment this knowledge with confidence measures, which are consistently processed at each step of the cognition chain, considering the confidence of previous processing steps along with additional noise introduced by sensors and the uncertainty introduced by the individual algorithms. Given that, the co-pilot plans the safe maneuvers considering all these factors and then distribute the shared maneuver execution to driver and automation, including handing-over tasks to the driver or accepting/rejecting tasks assigned by the driver to the automation. In order to maintain the common frame of reference (see the "meta-cooperation" in Hoc's framework [Hoc et al., 2009]), the system has always to "explain" maneuvers, situation and task distribution to the driver. The following three subsections describe the three main components of the prototype architecture in *Figure 1*.

3.2 The co-pilot module for decision-making

The co-pilot module is designated to support the maneuver decisions of the human driver, using a *Bounded Markov Decision Process* (BMDP) for the decision process. [Givan et al., 2000]. *Figure 2* depicts the logical flow of the module. It starts by building the initial BMDP state s_0 using the sensor data (world representation).

The set of actions *Act* considered in the prototype are:

- *Keep Your Lane* (KYL): the EV (ego vehicle) continues following the current lane at the current speed.
- *Brake* (Brk): the EV will try to brake (considering a span of possible decelerations).
- *Change Left Lane* (CLL): the EV moves to the next lane on the left (considering a span of possible lateral accelerations) to start an overtake maneuver.
- *Change Right Lane* (CRL): like before for the right lane, usually to conclude an overtake maneuver.
- *Slowdown* (Slw): the EV decelerates following the current lane.

The projection function $F_{\updownarrow}(s, act)$ produces a new BMDP state s' starting from state s and simulating the consequence of action *act*. This function follows the formulas of **Errore. L'origine riferimento non è stata trovata.**, and involves vehicle dynamics and object kinetics. The projection function propagates uncertainty of the state parameters.

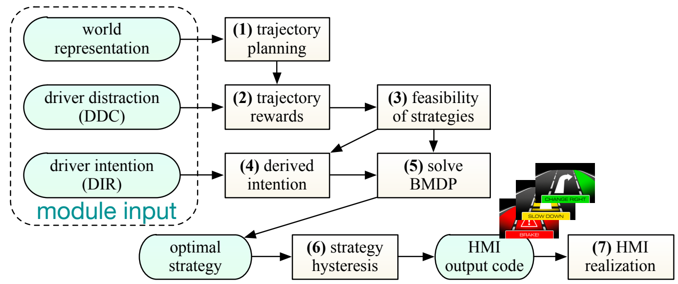

Figure 2: logic-flow of the co-pilot module for decision-making.

The set of states is generated using a variation of an *Online Sparse Sampling Algorithm* (OSSA) for BMDP solution: Starting from s_0, a tree of possible states is generated using $F_{\updownarrow}$, evaluating all the actions up to time T. Each path in the tree is a *trajectory*. A sequence of actions (i.e. a *policy*) is mapped into a set of multiple trajectories, due to the uncertainty encoded by $F_{\updownarrow}$. Step (2) assigns a reward to each trajectory, taking into account a safety measure of the state based on the standard time-to-collision measure:

$$R\updownarrow(s, act) = \mathrm{ramp}(ttc, r_{\min}, r_{\max})$$

Where ramp(*value*, *min*, *max*) is a linear ramp function, and the reaction times used in the prototype are $r_{\min}=3$, $r_{\max}=4$ if the driver is not distracted, $r_{\min}=4$, $r_{\max}=6$ otherwise (hence distraction raises the human reaction time). In step (3) trajectories that exceed a safety threshold value are considered *not safe*. Step (4) considers the feasibility of the estimated driver intention, which could result in two outcomes:

- Intended action is safe: this generates a (positive) suggestion in the HMI of doing that action, like: "you may change left".
- Intended action is not safe: this generates a (negative) indication in the HMI that the action is dangerous, like: "slowdown" because there is a vehicle ahead, or "do not change right" since the other lane is occupied.

The reward of a policy is the minimum of the rewards of that policy trajectories. The BMDP solution of step (5) consists in selecting the policy with the maximum reward function. The first action of the optimal policy is passed to the HMI module, after a hysteresis step (6).

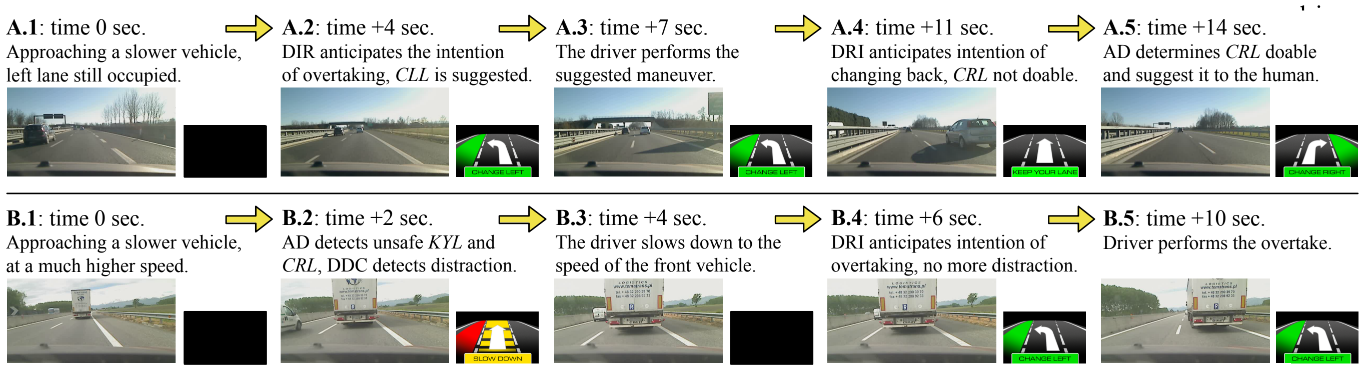

Figure 3: Two examples of the system on the prototype vehicle, showing the external camera view and the HMI.

3.3 Driver Intention Recognition (DIR)

In order to take an optimal decision, as illustrated in *Figure 2*, two blocks are also taken into account: the *Driver Intention Recognition* (DIR) and the *Driver Distraction Classification* (DDC).

The DIR module is the software component designed to provide context-dependent estimates of the hidden lane-change maneuver intentions of the human driver. The hope is to be able to detect the intention prior to the turn indicators. This is an important aspect for the co-pilot, because – for example – it can learn the driver preferences during the "normal" driving and then adopt this during the automated driving task.

At runtime, the module receives input from the vehicle sensors (actuator states, current velocity, position data provided by a lane detection camera, surrounding obstacles as seen by the LIDARs). Within the DIR module, the available sensor information is first synchronized. An internal world-model based on particle filters [Koller and Friedmann, 2009] then augments the available information with better estimates of the environment and vehicle position and classifies surrounding vehicles according to predefined roles (e.g., the lead vehicle, the vehicle behind on the left lane, etc.). For actual intention recognition, the DIR module utilizes a DBN that describes the statistical relations among the intentions, the behaviors, and the information of the vehicles state and traffic situations: the DIR model. The details of the DIR module can be found in [Eilers et al., 2016] and in [Yan et al 2016].

3.4 Driver Distraction Classification (DDC)

The purpose of the DDC module is to classify the visual driver distraction based on vehicle dynamics data and internal camera, using the machine learning techniques described as following. This module provides information about the operator's degree of distraction.

DDC consists of two components: the first component learns the classifier from a stream of sensory data. The second component uses the classifier to make prediction on the distraction status of the driver. The classifier can be trained either offline or online during its use. In this prototype, the classifier has been trained offline from system dynamics data collected from the prototype vehicle during 30 sessions. The 30 subjects drove for about 1 hour on normal and highway roads and they had to perform a SURT (surrogate task) while driving, in order to induce distraction.

The vehicle dynamic variables considered in this study are:

Speed [m/s] Time To Lane Crossing [s] Time To Collision [s] Position of accelerator pedal [%] Heading Angle [deg] Position of the brake pedal [%] Steering Angle [deg] Turn indicators [on/off]	Lateral Position [m] X,Y coordinates of front car (if any) Lane Width [m] Speed of car in front (if any) Road Curvature [%] Output of the monitoring system (head position and eyes tracking)

These values are directly available from the vehicle sensor data, or can be derived from those (e.g., time to collision is computed using the LIDAR data). The frequency of data collection is 20 Hz (1 data-point each 0.05s). Each of the continuous input variables above generates five input channels, namely the average, minimum, maximum, standard deviation and first derivative in a sliding window of given width. Discrete variables enter directly as input channels.

The DDC module employs a SLFN network with 63 input neurons (one for each input channel), 100 hidden neurons and 2 output neurons. Weights are determined offline using ELM algorithm, and are loaded by the DDC module at runtime. The two output neurons generate the distracted/non-distracted probability distribution, which is then discretized to obtain the distraction classification used as input for the decision module.

3.5 The system in action

Figure 3 shows how the system works in practice on the prototype vehicle with two small examples. These examples are extracted from a test drive done with the prototype vehicle on a highway near Torino, Italy. The purpose is to illustrate the adaptation on the driver intention/distraction in the decision process of the whole system. In example **A** the driver is fast approaching a slower car on the right lane. The DIR module infers that the most probable intention of the driver is to overtake that car, and the DDC considers the driver to be attentive. With this setting, the co-pilot module verifies that the overtake is safe (no obstacles), and suggests the driver with 7 seconds in advance that (s)he may overtake, supporting all the maneuvers. The suggestion does not depend on the turn indicators of the car. At stage A.4 in *Figure 3* the DIR module infers that the driver should return to the right lane. The co-pilot module first shows a "keep your lane" enforcement signal until the right lane is occupied (A.4), and finally, shows the CRL message to support the reentrance. Note that a forward-collision-warning and a blindspot will behave differently. They would both signal the longitudinal and lateral dangers (FCW at A.3 and blindspot at A.5), because they are not adaptive to the overtake intention.

Example **B** shows a similar scenario where the driver is again fast approaching a slower truck, but the left lane is occupied by another car, and the DDC module considers the driver to be inattentive. In this case the co-pilot module determines that the safest (highest reward) policy is a slowdown action, and it emits a "slowdown/do not change left" warning to the driver. If the driver does not respond quickly, an emergency brake signal would appear. After having adjusted the speed to follow the truck, the DDC module determines that the driver is attentive again (unclear whether attention is deduced because the speed has been reduced or because the module says so and in this scenario we consider the case in which the driver is attentive again at this point in time, and a CRL message is shown to suggest overtaking.

4 Experimental Phase

A prototype has been developed to test the feasibility of the co-pilot. The prototype addresses the driving in a highway scenario, which is adequate since the focus of the project is the adaptation to the human behavior more than the adaptation to the environment. The prototype has been realized in two forms: a vehicle and a simulator. Both runs the same software stack built on the RT-Maps 4 framework[1].

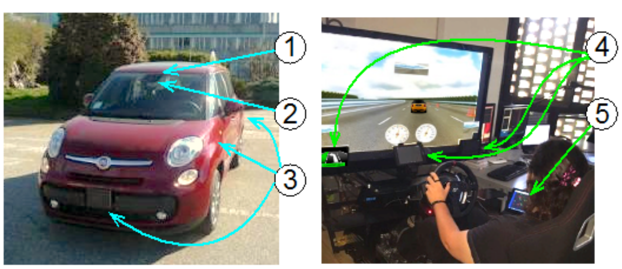

Figure 4: vehicle and driving simulator prototypes used for evaluation phase.

4.1 Overview of the experimental prototype

In order to test the feasibility of the system, including the co-pilot, DIR and DDC modules, we have considered one prototype (as illustrated in *Figure 4*), which is adequate for the adaptation to the human behavior and to the environment. The prototype has been realized in two forms: a vehicle and a simulator. *Figure 4*(left) shows the sensors used for the data collection: 1) an external camera used by the lane-detection algorithm (to build the road model); 2) an internal camera used to scan the human driver face; 3) four LIDAR sensors that detect environment obstacles. *Figure 4*(right) shows the setup of the simulator environment, which includes 4) the distributed co-pilot HMI; 5) the SURT used to distract the users. The simulator runs the SCANeR II software.

4.2 Evaluation of the Classification modules

A large number of experiments was carried out to test the DDC module, by varying the learning algorithm parameters, such as the number of neurons, learning rates, number of training instances, etc. Moreover, the collected data have been averaged over a period of time that varies, in the various experiments, between 1s and 2s. In order to be consistent with the target variable (*distracted* or *not-distracted*), data have been labelled *distracted* when the driver was not looking at the road for the whole considered period (using the internal camera of the vehicle), *not distracted* otherwise. Table 1 shows the main results:

Accuracy Learning Set	*Accuracy Test Set*	*False Positive Rate*	*False Negative Rate*
98.99+/-0.04	87.52+/-1.37	0.07+/-0.02	0.18+/-0.02

Table 1: classifier performances on the test-set.

The table reports the average classification rates (accuracy, false negative rate and false positive rate) obtained by training the DDC module using a leave-one-out strategy. In detail, we selected 25 drivers out of the 30 ones for which we collected data (5 were discarded due to a small number of distraction cases). In turn, one of these dataset has been used for testing while the other 24 have been combined together to form the learning set, and the process repeated for every dataset (25 runs in total).

In details, the first two columns report the average accuracy on the training sets and test set, respectively. The accuracy on the test cases shows good generalization ability towards new drivers. For what concerns the false negative rate, it should be noted that in practical situations there are many more cases of driver not distracted than there are cases of distraction, and the latter are more difficult to identify, in general, because good drivers tend to drive safe even when partially distracted.

The evaluation of the DIR module has been the subject of a separate analysis. Interested readers can find it in [Eilers et al., 2016] and in [Yan et al 2016].

4.3 Evaluation results of the system prototype

The aim of the system experimental assessment was to objectively measure its performance against a state-of-the-art baseline system. The assessment was done on the driving simulator, for safety reasons. The baseline system includes both a blindspot and FCW systems for lateral and longitudinal warnings.

[1] [1] RT-Maps framework. https://intempora.com/.

The test involved 30 subjects (15 men, 15 women; average age: 39 years; average years of driving license: 20; average km/year: 15.660) on five tracks, without anticipating the kind of support system that they would have seen while driving. After a warm-up track (to get confidence with the simulator), the test is performed two times on two tracks (5 minutes each), with the baseline and with the prototype system. Half of the participants see the system after the baseline, while the other half see it before (to randomize and avoid bias).

The two test tracks are two-lane highway scenarios designed to create specific dangerous situations to the driver, like: a preceding vehicle that brakes unexpectedly, a slow preceding vehicle with an incoming car on the left lane, … Details on the tested situations can be found in [HoliDes, 2016] (D9.9). Drivers were requested to perform a SURT task that activated at random every 30/45 seconds on a secondary screen that is located on the side of the simulator screen (to trigger a visual distraction).

The following Performance Indicators where considered:

PI1. Number of accidents occurred during the test.

PI2. Percentage of driving time where the TTC of the preceding vehicle is less than 2 seconds.

PI3. Number of times the driver presses the brake strongly, achieving a sudden hard braking (deceleration of more than -8m/s2).

PI4. Average distance to the preceding vehicle when the user performs the lane change. The purpose is to measure if the new system increases safety, avoiding the Peltzman effect [Peltzman, 1975] (inducing confidence in making risky maneuvers).

PI5. Average TTC when the driver starts pressing the brake, to measure the impact of slowdown/brake suggestions.

Table II shows the average PI values obtained from the user tests, and the statistical significance. The values show favorable performance of the new system over the baseline.

	PI1	PI2	PI3	PI4	PI5
Baseline	0.1724	0.0126	1.3103	73.5552	2.7134
co-pilot	0.0862	0.0069	1.3793	82.1026	3.3742
Signif.	<0.0001	<0.0001	0.0011	0.5461	0.0085

Table II: PI results of the evaluation.

For the number of accidents, they are strongly reduced: in the baseline, there are 10 accidents in the test cases, while in the new system there are only 5 accidents. This means the co-pilot was able to halve the number. A similar improvement has been achieved also for PI2, where the time spent by the drivers in potentially critical situations has been reduced by around 50%.

The higher value of the system on the indicator PI5 means that, when the driver starts braking, the TTC is greater in the new system than it is in the baseline, meaning that the co-pilot increases the awareness of the driver to dangerous situations. The number of hard brakings is almost equal, which could be caused by the fact that the scenarios make these braking necessary due to the critical situations that happens.

PI4 takes into account the effect of the recognition of driver intention: when the DIR module infers the intention of overtaking, the system suggests the maneuver; suggesting the lane change maneuver in advance seems to modify the attitude of the users in changing the lane in advance, increasing the safety distance.

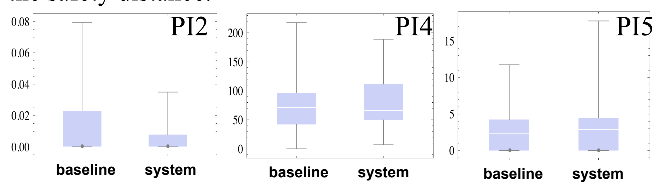

Figure 5: Box-plots of the continuous PIs.

Figure 5 shows the box-plots of the three continuous PIs (left is baseline, right is new system). The statistical test shows that all PIs have enough samples to assess for a statistical difference, except for PI4 (even if it is showing slightly different distributions), which requires more samples.

5 Discussion and Conclusions

In the European co-funded project HoliDes [HoliDes, 2016], a first version of the co-pilot has been implemented and used in a limited form to produce a comprehensive preventive safety system capable of giving information and warnings only. However, this concept of the co-pilots are potentially suited to more sophisticated applications, such as the ones developed in the EU project AutoMate (http://www.automate-project.eu/). In fact, if necessary, the driver could “loosen” control and let the system autonomously navigate, or can “tighten” control and reclaim authority. On the other way around, if necessary, the system may be programmed to completely take over from the driver in certain conditions.

In all cases, the important aspect is that the co-pilot will be adaptive and cooperative, thus the driven vehicle should appear to be driven by a human, being easily interpretable to other human road users. In this context, the co-pilot is the enabling technology for implementing ADFs, with the capacity to improve the human–vehicle interactions, by considering and exchanging intentions between co-driver-equipped vehicles, as well as by taking into account the driver state (i.e. visual distraction).

In this paper, we have attempted to set out a viable roadmap for producing the co-pilot enabling technology, making significant use of recent developments in cognitive systems, in order to address the adaptation and continuous support to the human driver. This system is realized using a decision process that balances multiple action outcomes with the inferred human status (intention, distraction). This produces contextualized strategies that are shown as graphical messages through a dedicated HMI on the vehicle dashboard, or alternatively, can be regarded as a “virtual driver”, able to take the vehicle control and thus implementing the ADFs.

Moreover, other achievements will need further research at the intersection between cognitive sciences and intelligent vehicles. In particular, we plan to investigate the impact of introducing a human-like behavior of the co-pilot, sharing experience and roles with human drivers. In this context, the system could use its emulation capacity (in a way similar to human rehearsing of possible experiences) to discover and learn higher-level behaviors that might prove more effective. This enables the co-pilot to become an expert driver without directly needing training examples from expert humans. In this sense, reinforcement-learning techniques would probably also allow to better tailor such a system to each human driver.

Acknowledgments

This work was supported by the EU Artemis Joint Undertaking research project HoliDes, grant no. 332933.

References

[Bishop, 2005] R. Bishop, "Intelligent vehicle R&D: A review and contrast of programs worldwide and emerging trends," Ann. Telecommun., vol. 60, no. 3/4, pp. 228–263, Apr. 2005.

[Inagaki, 2008] T. Inagaki, "Smart collaboration between humans and machines based on mutual understanding," Annu. Rev. Control, vol. 32, no. 2, pp. 253–261, Dec. 2008.

[Flemisch, 2003] F. O. Flemisch et al., "The H-metaphor as a guideline for vehicle automation and interaction," NASA Langley Res. Center, Hampton, VA, USA, NASA/TM-2003-212672, Dec. 2003.

[Da Lio et al., 2015] Da Lio, M. and Biral, F. and Bertolazzi, E. and Galvani, M. and Bosetti, P. and Windridge, D. and Saroldi, A. and Tango, F. "Artificial co-drivers as a universal enabling technology for future intelligent vehicles and transportation systems," IEEE Transactions on Intelligent Transportation Systems, 16 (1), pp. 244-263, 2015.

[Li et al., 2012] L. Li, D. Wen, N.-N. Zheng, and L.-C. Shen, "Cognitive cars: A new frontier for ADAS r search," IEEE Trans. Intell. Transp. Syst., vol. 13, no. 1, pp. 395–407, Mar. 2012.

[HAVEit, 2013] HAVEit-Highly Automated Vehicles for Intelligent Transport, [Accessed: 15-Jun-2017]. [Online]. Available: http://www.haveit-eu.org/.

[HoliDes, 2016] project: http://www.holides.eu/public-deliverables

[Börger, 2013] J. Börger, "Fahrerintentionserkennung und Kursprädiktion mit erweiterten Maschinellen Lernverfahren", dissertation, Universität Ulm, 2013.

[Doshi and Trivedi, 2011] A. Doshi and M. M. Trivedi, "Tactical Driver Behavior Prediction and Intent Inference: A Review", in Proceedings of the 14th International IEEE Conference on Intelligent Transportation Systems, 2011, pp. 1892-1897.

[Kobiela, 2011] F. Kobiela, "Fahrerintentionserkennung für autonome Notbremssysteme", Springer, 2011.

[Lefèvre et al., 2014] S. Lefèvre, D. Vasquez, and Ch. Laugier, "A Survey on Motion Prediction and Risk Assessment for Intelligent Vehicles". Robomech Journal, 1, 1, 2014, pp. 1-14.

[Oliver and Pentland, 2000] N. Oliver and A. P. Pentland "Driver Behavior Recognition and Prediction in a SmartCar", in Proceedings of the 2000 IEEE Intelligent Vehicles Symposium, 2000, pp. 7-12.

[Kumagai and Akamatsu, 2006] T. Kumagai and M. Akamatsu, "Prediction of Human Driving Behavior Using Dynamic Bayesian Networks", in IEICE Transactions on Information and Systems, Vol. E89-D, No. 2, 2006, pp. 857-860.

[Liebner et al., 2012] M. Liebner, M. Baumann, F. Klanner, and Ch. Stiller, "Driver Intent Inference at Urban Intersections using the Intelligent Driver Model", in Proceedings of the 2012 IEEE Intelligent Vehicles Symposium, 2012, p. 1162-1167.

[Tay, 2009] Ch. Tay, "Analysis of Dynamic Scenes: Application to Driving Assistance" Ph.D. dissertation, Institut National Polytechnique de Grenoble – INPG, 2009.

[Kumar et al., 2013] P. Kumar, M. Perrollaz, S. Lefèvre, and C. Laugier, "Learning-based approach for online lane change intention prediction", in Proceedings of the 2013 IEEE Intelligent Vehicles Symposium, 2013, pp. 797-802;

[Garcia-Ortiz et al., 2011] Garcia-Ortiz M, Fritsch J, Kummert F, Gepperth "A behavior prediction at multiple time-scales in inner-city scenarios" In: Proc. IEEE intelligent vehicles symposium, 2011, pp.1068–1073.

[Morris et al., 2011] B. T. Morris, A. Doshi, and M. M. Trivedi, "Lane-Change Intent Prediction for Driver Assistance: On-Road Design and Evaluation", in Proceedings of the 2011 IEEE Intelligent Vehicles Symposium, 2011, pp. 895-901.

[Regan et al., 2011] M. A. Regan and C. Hallett and C. P. Gordon, "Driver distraction and driver inattention: Definition, relationship and taxonomy", Accident Analysis & Prevention, vol. 43, no. 5, 2011.

[Liang et al., 2007] Liang Y. and 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.

[Huang et al., 2006] Huang, G.B. and Zhu, Q.Y. and Siew, C.-K "Extreme Learning Machine: A New Learning Scheme of Feedforward Neural Networks". In: Extreme Learning Machine, vol.70, pp.489-501, 2006.

[Stiller et al., 2007] C. Stiller, G. Farber, S. Kammel. “Cooperative Cognitive Automobiles”. Proceedings of the 2007 IEEE Intelligent Vehicles Symposium. Istanbul, Turkey, June 13-15, 2007.

[Hoc et al., 2009] J.M. Hoc, M. S. Young, and J.-M. Blo seville, “Cooperation between drivers and automation: implications for safety,” Theor. Issues Ergon. Sci., vol. 10, no. 2, pp. 135–160, Mar. 2009.

[Givan et al., 2000] R. Givan, and S.M. Leach, and T.L. Dean. Bounded-parameter Markov Decision Processes. J. Artificial Intelligence, volume 122(1-2), pages 71-109. 2000.

[Tango and Pietquin, 2012] Tango F. and Pietquin O. “A Reinforcement Learning Approach to Optimize the longitudinal Behavior of a Partial Autonomous Driving Assistance System”. 20th European Conference on Artificial Intelligence (ECAI) and 7th Conference on Prestigious Applications of Intelligent Systems (PAIS). Montpellier (France). 2012 (Best Paper PAIS Section).

[Puterman, 2005] M.L. Puterman. “Markov Decision Processes. Discrete Stochastic Dynamic Programming”. Wiley, Chichester, 2005.

[Wei, 2011] J. Wei, J. M. Dolan, J. M. Snider, B. Litkouhi, A point-based MDP for robust single-lane autonomous driving behavior under uncertainties, in: Robotics and Automation (ICRA), 2011 IEEE International Conference on, 2011, pp. 2586–2592.

[Kearns et al., 2002] M. Kearns, and Y. Mansour and Ng Y.A.A. Sparse sampling algorithm for near-optimal planning in large Markov Decision Processes. J. Machine Learning. volume 49(2), pages 193-208. 2002.

[Glaser et al., 2010] Glaser et al. “Maneuver-Based Trajectory Planning for Highly Autonomous Vehicles on Real Road With Traffic and Driver Interaction”, ITS, vol. 11, issue 3, Sept. 2010.

[Koller and Friedmann, 2009] D. Koller and N. Friedmann, „Probabilistic Graphical Models: Principles and Techniques”, The MIT Press, Cambridge, Massachusetts, London, England, 2009.

[Eilers et al., 2016] Fei Yan, Mark Eilers, Andreas Lüdtke, Martin Baumann, “Developing a model of driver's uncertainty in lane change situations for trustworthy lane change decision aid systems”. Intelligent Vehicles Symposium 2016: 406-411

[Yan et al 2016] Fei Yan, Mark Eilers, Andreas Lüdtke, Martin Baumann, “Development of a Lane Change Assistance System Adapting to Driver's Uncertainty During Decision-Making”. AutomotiveUI (adjunct) 2016: 93-98

[Peltzman 1975] “The Effects of Automobile Safety Regulation.” Journal of Political Economy, vol. 83, no. 4, 1975, pp. 677–725.