



Human driver modelling and simulation into a virtual road environment

Thierry Bellet, Pierre Mayenobe, Jean-Charles Bornard, Jean-Christophe Paris, Dominique Gruyer, Bernard Claverie

► To cite this version:

Thierry Bellet, Pierre Mayenobe, Jean-Charles Bornard, Jean-Christophe Paris, Dominique Gruyer, et al.. Human driver modelling and simulation into a virtual road environment. Human Modelling in Assisted Transportation: Models, Tools and Risk Methods, Milan, Springer, pp.251-262, 2011. hal-00668899

HAL Id: hal-00668899

<https://hal.science/hal-00668899>

Submitted on 10 Feb 2012

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.

Bellet T., Mayenobe P., Bornard J.C., Paris J.C., Gruyer D., Claverie B. (2011).

“Human driver modelling and simulation into a virtual road environment”

In Cacciabue, P. C.; Hjalmdahl, M.; Lüdtkke A.; Riccioli, C. (Eds.): *Human Modelling in Assisted Transportation : Models, Tools and Risk Methods*.

Milan, Springer, pp. 251-262.

Human driver modelling and simulation into a virtual road environment

Thierry Bellet¹, Pierre Mayenobe¹, Jean-Charles Bornard^{1,3}, Jean-Christophe Paris¹, Dominique Gruyer², Bernard Claverie³

¹ INRETS-LESCOT, 25, Av. François Mitterrand, 69675 Bron, France, thierry.bellet@inrets.fr

² INRETS-LIVIC, 14 Route de la minière, 78000 Versailles, France, dominique.gruyer@inrets.fr

³ ENSC, 146, rue Léo Saignat 33076 Bordeaux, France, bernard.claverie@ensc.fr

Abstract: This paper presents the driver model developed by INRETS in the ISi-PADAS project, with the aim to dynamically simulate driver's mental activities carried out while driving. The methodology supporting this model development is based on empirical data collection on driving simulator, in the frame of a car-following task. After presenting the theoretical foundations of the modelling approach and the empirical data analysis, the functional architecture of our COgnitive Simulation Model of the DRIVER (COSMODRIVE) will be described, and the type of results liable to be obtained through simulations on a virtual Vehicle-Environment platform (SiVIC) will be presented. Then, the conclusion will briefly examine the perspectives of the model applications for driving aids virtual design.

Introduction

The general objective of this research is to design, develop and implement a cognitive simulation model of the car driver, able to virtually simulate the human drivers' mental activity carried out while driving, through an iterative "*Perception-Cognition-Action*" regulation process. This model is more particularly focused on three following functions: (1) *Perception* of the road environment and of the other road users behaviours, (2) *Cognition*, integrating elaboration of *mental representations* of the road scene (corresponding to driver's situational awareness) and *decision-making* processes (based on these mental models of the driving situation and on anticipations assessed from dynamic mental simulations) and (3) *Action*, corresponding to behavioural performances as decided at the cognitive level and then effectively implemented via actions on vehicle commands, in order to dynamically progress into the road environment. Moreover, the aim of the model is not only to simulate these perceptive, cognitive and executive functions in an optimal way, but also to generate human errors in terms of non-perception of event, erroneous situation awareness, or inadequate driving performances.

Theoretical foundations

From a theoretical point of view, this research is based on the COSMODRIVE model (i.e. *COgnitive Simulation MOdel of the DRIVER*; Bellet et al. 1999; 2007) dedicated to driver's mental activities modelling. Basically, driving a car requires (i) to select relevant information in the environment, (ii) to understand the current situation and to anticipate its change in the more or less long term, (iii) to take decisions in order to dynamically interact with the road environment and the other road users, (iv) and to manage own resources (physical, perceptive and cognitive) in order to satisfy the time constraints of the task, inherent to the dynamic nature of the driving situation. The selective dimension of information collection is especially important as drivers cannot take in and process all the information available in the road environment. This information selection mechanism is the result of a complex process whose keystone are the driver's mental representations of the driving situation. Mental representations correspond to the driver's *Situation Awareness*, according to Endsley (1995) definition of this concept. These mental models are dynamically produced through a matching process between pre-existing *operative* knowledge and the perceptive information extracted in the road scene. At the *tactical level* (Michon, 1985), these mental representations provide ego-centred and goal-oriented visual-spatial models of the driving situation, which are dynamically produced and continually updated, as and when the drivers carry out their activity. One of their core-functions is to support cognitive anticipations, through mental simulations, providing expectations of future situational states. Moreover, driver modelling requires to consider two different levels of activity control: an automatic and implicit mode *versus* an attentional and explicit mode. This dichotomy is well established in the literature with the distinction put forward by Schneider et al. (1977) between *controlled processes*, requiring cognitive resources and which are only performed sequentially, and *automatic processes*, which can be performed in parallel without any attentional effort. In the same way, Rasmussen (1986) distinguishes 3 levels of activity control according to whether the behaviours implemented rely on (i) highly integrated abilities (*Skill-based behaviours*), (ii) decision rules for managing familiar situations (*Rule-based behaviours*), or (iii) more generic knowledge that is activated in new situations for which the driver have not any prior experience (*Knowledge-based behaviours*). Regarding the functioning of the human cognitive system while driving, a large share of the driver's activity relies on heavily integrated and automated empirical know-how partly escaping to conscious control, but nonetheless relying on an implicit form of awareness and activity monitoring, to guarantee that the goals explicitly defined are reached. Thus, these two levels of control support themselves, and are embedded in each other (Bellet et al., 2009). By considering this theoretical background, the computational version of COSMODRIVE implemented on the SIVIC platform during the ISIPADAS project is composed of three functional modules (i.e. *Perception*, *Cognition*, and *Action* modules) and is able to drive a virtual car on a virtual road from 2 synchronized "*Perception-Cognition-Action*" regulation loops: (i) an attentional

control loop, based on COSMODRIVE *Driving Schemas*, and (ii) an automatic control loop, simulated through the *Envelope Zones* strategy and the *Pure-Pursuit Point* method.

Modelling the tactical and explicit cognition: the Driving Schemas

Based on both the Piaget's concept of operative *scheme* and the Minsky (1975) *frames theory*, driving schema is a computational formalism defined at INRETS for driving knowledge modelling at the tactical level (Bellet et al, 1999; 2007). They correspond to prototypical situations, actions and events, learnt by drivers from their practical experience. From a formal point of view, a Driving Schema is composed of (i) a functional model of road *Infrastructure*, (ii) a *Tactical Goal* (e.g. turn left), (iii) a sequence of *States* and (iv) a set of *Zones*. Two types of zone are distinguished: *Driving Zones*, corresponding to the driving path of the vehicle as it progresses on the road, and the *Perceptive Exploration Zones*, in which the driver seeks information (e.g. potential events liable to occur). Each driving zone is linked with *Actions* to be implemented (e.g. braking or accelerating, in view to reach a given state), with a set of *Conditions* for performing these actions, and with perceptive zones that permitting to check these conditions. A *State* is defined by a vehicle *Position* and *Speed*. The different sequences of the driving zones make up the *Driving Paths* that progress from the initial to the final state (i.e. achievement of the tactical goal). Once activated in the working memory and instantiated with the characteristics of road environment, the active driving schema becomes the *tactical mental representation* of the driver, which will be continuously updated as and when s/he progresses into the road infrastructure. This representation corresponds to the driver's explicit awareness of the driving situation and provides a mental model of the road, functionally structured according to the tactical goal followed by the driver in the current context (e.g. turn left).

Modelling the operational skills and the implicit cognition:

At the operational level, corresponding to the automatic control loop, the COSMODRIVE model regulation strategy is jointly based on the *envelope zones* and the *pure pursuit point* approaches. The concept of *envelope zones* recalls two classical theories in psychology (Bellet et al., 2007): the notion of *body image* proposed by Schilder (1950), and the theory of *proxemics* defined by Hall (1966), relating to the distance keeping in social interactions with other humans. Regarding car driving, envelope zones refer to *safety margins* (Gibson et al., 1938). At this level, our driver model is based on Kontaratos' work (1974) who distinguished a *safety zone*, a *threat zone*, and a *danger zone*. Envelope zones correspond to the part of the path of driving schemas to be occupied by the car in the near future. As an "hidden dimension" of the social cognition, as suggested by Hall's theory (1966), these proxemics zones are also mentally projected to other road users, and are then used to dynamically interact with them, as well as to anticipate and manage collision risks. This virtual skin is permanently active while driving, as an implicit awareness of the expected allocated space for moving. As with Schilder's body schemas, it belongs to a highly integrated cognitive level (i.e. implicit regulation

loop), but however favours the emergence of critical events in the driver's explicit awareness. Therefore, the envelope zones play a central role in the regulation of "social" as well as "physical" interactions with other road users under normal driving conditions (e.g. inter-vehicle distance keeping), and in risk assessment if a critical situation occurs (commitment of emergency reactions).

A second "hidden dimension" of the implicit cognition as implemented at the operational level of COSMODRIVE concerns the executive functions of lateral and longitudinal controls of the car, to be carried out in order to dynamically progress along the driving path of the driving schema. This automatic control loop is based on the *Pure Pursuit Point* method. The *Pure-Pursuit Point* method was initially introduced by Amidi (1990) for modelling the lateral and the longitudinal controls of automatic cars along a trajectory, and has been adapted by Sukthankar (1997), and then Mayenobe (2004), for driver's situational awareness modelling. Mathematically, the pure-pursuit point is defined as the intersection of the desired vehicle path and a circle of radius centred at the vehicle's rear axle midpoint (assuming front wheel steer). Intuitively, this point describes the steering curvature that would bring the vehicle to the desired lateral offset after travelling a distance of approximately l . Thus the position of the pure-pursuit point maps directly onto a recommended steering curvature: $k = -2x/l$, where k is the curvature (reciprocal of steering radius), x is the relative lateral offset to the pure-pursuit point in vehicle coordinates, and l is a parameter known as the look-ahead distance. According to this definition, the vehicle-control abilities of COSMODRIVE for driving a virtual car are implemented as a dynamic regulation loop that permanently keeps the Pursuit Point on the driving path of the current driving schema, to a given speed assigned with each segment of the current tactical driving schema, as instantiated in the mental representation.

Methodology for model design and data collected

The methodological specificity of the driver modelling approach implemented in this research was to use the SiVIC virtual Platform (Gruyer et al., 2006) as (i) a driving simulator for empirical data collection with real drivers, and then, as (ii) a virtual road environment to be interfaced with the driver model for future virtual simulations. According to this approach, human drivers' behaviour and driver model performances will be observed for the same driving scenarios of car following, in the same virtual road environment. From these similarities, it is expected to facilitate the model validation and to increase its validity.

Participants

Twenty experienced drivers of middle-age (from 23 to 56 years old) have participated to the experiment. All the drivers have a minimum of 5 years of driving experience and they drive a minimum of 5000 kilometres per year.

Driving scenarios and visual Secondary Task (ST)

The participants' driving task was to follow a lead car in different driving conditions. Four main sources of variation have been more particularly investigated: (1) the driving context (i.e. motorway, rural road and urban area), and consequently the vehicle speed required (respectively 130, 90, and 50 km/h), (2) the nature of the car-following task (i.e. free *vs* imposed car following distance at a given Inter-Vehicular Time [IVT] of 0.6 sec.), (3) the lead car behaviour (having a steady *vs* irregular speed), and (4) the necessity to perform a Secondary Task (ST) while driving. Concerning more specifically visual distraction, the ST to be performed by the participants was the following: a set of 3 visual pictograms, accompanied with an auditory beep, were displayed on an additional screen (situated on the right side, near the usual position of the radio). Some seconds later (from 3 to 4 sec.), 1 of this 3 pictograms appeared under the first set, and the driver had to use a 3-buttons command for indicating which pictogram was replicated.

Main results

The results presented here only concern the negative impact of a visual ST on the drivers' performances, more particularly by considering the driving behaviour modifications in normal conditions (e.g. inadequate following distance), and the accident risk increasing for critical scenarios (i.e. when the lead car brakes).

Table 1: the percentage of collision with the lead for critical scenarios

<i>Context</i>	<i>Driving scenario</i>	<i>Without ST</i>	<i>ST-Visual</i>
Highway	Free steady lead car following	55 %	50 %
	Free unsteady lead car following*	35 %	50 %
	Constrained steady lead car following	65 %	70 %
	Constrained unsteady lead car following	70 %	70 %
Rural	Free unsteady lead car following	60 %	60 %
	Constrained unsteady lead car following*	55 %	80 %
Urban	Free steady car lead following*	20 %	30 %
	Free unsteady lead car following	30 %	30 %
	Constrained steady lead car following	30 %	30 %
	Constrained unsteady lead car following*	25 %	90 %

(***Bold values** indicate significant differences between without-ST *vs* ST conditions; T-test, $p < 0.05$)

In *normal driving conditions*, two main differences due to visual distraction have been observed: (i) a significant reduction (T-test, $p < 0.001$) of the safety margins in *free following conditions* (without ST, mean value of IVT is of 3 s. without ST *vs* 2.65 s. with ST) and (ii) a significant degradation ($p < 0.05$) of the following performance in constrained following conditions (in these scenarios, drivers have to follow the lead car at an imposed IVT of 0.6 s., and the percentage of time when this value is performed is of 57% without ST, *vs* 44 % with ST). These results show a negative effect of visual ST for short following distance keeping.

In *critical driving conditions*, the two main negative impacts of the visual ST on drivers' performances are (i) an increasing of Reaction time for braking (the differences are only significant for the constrained following task: 0.89 s. vs 1.1 s.; $p < 0.05$), and (ii) a risk of crash increasing: Table 1 presents the percentage of collision with the lead car/total number of required emergency braking, for the different driving conditions investigated. It appears that the risk of collision due to a ST is significantly increased for 4 of the 10 driving scenarios requiring an emergency braking. The highest negative impacts of visual ST are observed for the *constrained unsteady car following* scenarios, in both urban and rural areas.

Driver Model description

The functional architecture of the version of the COSMODRIVE model implemented into SiVIC is composed of three main modules (Figure 5): A Perception Module, a Cognition Module, and an Action Module.

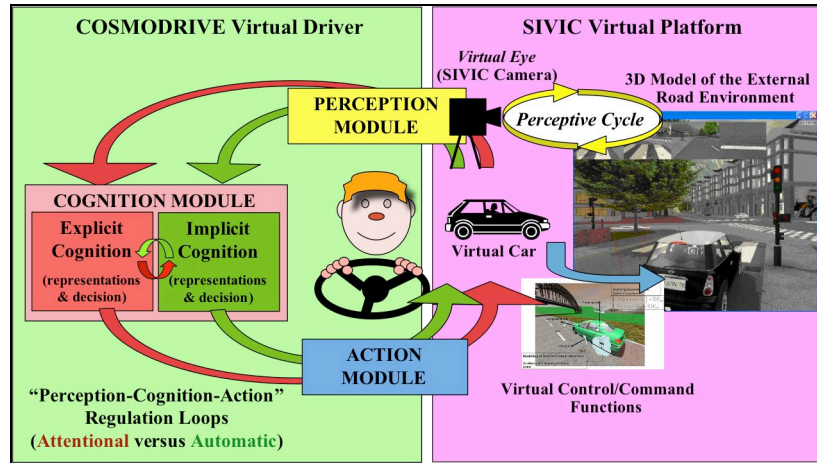


Figure 1: COSMODRIVE model interfaced with the SiVIC Virtual Environment

By implementing COSMODRIVE into the SiVIC Platform, it becomes possible to generate dynamic simulations of the driver model interacting with a virtual road environment, through actions on a virtual car.

Perception Module

The Perception module is based on a *virtual eye*, designed as a new type of SiVIC virtual sensor, adapted from the virtual camera model pre-existing in this platform. This virtual eye includes three visual field zones (fig. 1): the *foveal vision* (solid angle of 2.5° centered on the fixation point) with a high visual acuity, *para-foveal vision* (from 2.5° to 9°), and *peripheral vision* (from 9° to 150°). Visual strategies implemented by the driver model are simulated through a dynamic visual scanning of the road scene by the virtual eye. The visual strategies, modelled as a sequence of fixation points, are implemented by progressively considering *perceptive queries*

received by the Perception module from the Cognition module. Each query requires to focus the virtual eye on a specific area of the road scene. Perceived data is then integrated into the implicit and the explicit mental representations of the Cognition Module.




Figure 2: the COSMODRIVE virtual eye, as implemented on SiVIC

However, two complementary perceptive processes have been implemented. The first one is the *perceptive integration* (that is a data-driven process, i.e. bottom-up integration of perceptive data), allowing the cognitive integration of the *perceptible* data in the mental representations of the Cognition Module. The second process is the *perceptive exploration* (that is a knowledge-driven process) in charge to move the virtual eye in the road scene, from a point of fixation to another one, according to the perceptive queries received by the Perception module.

Cognition Module

The Cognition Module is implemented through two regulation processes: an attentional control process, based on an explicit awareness of the driving situation requiring cognitive resources for sequential reasoning, and an automatic process, based on an implicit situational awareness and cognitive skills liable to run in parallel. Moreover, two main cognitive functions are implemented in this module: mental representation elaboration and decision-making. Concerning mental representations elaboration, this process is based on *driving schemas* instantiation with the external environment characteristics. As visual-spatial models of the environment, mental representations modelling required to use several instances of the SiVIC 3D graphical engine (i.e. *representation* of current the driving situation, and *anticipated representations* corresponding to the driver's expectations on future situational states). These internal models of the external environment are continually fed by the *perceptive integration* and the *perceptive exploration* processes implemented in the *Perception module*. It is thus possible to simulate human errors in terms of inadequate mental representations (e.g. non-integration of perceptive data or event false-updating due to distraction). Concerning Decision-Making, this process is dually implemented in the Cognition Module. At the attentional level (i.e. *explicit* decisions), this process is based on State-Transition rules integrated into the *driving schemas*. At the automatic level (i.e. *implicit* decision-making), the decisional process is implemented via the *envelop-zones* regulation mechanism. Moreover, in order to support decision based on cognitive anticipations, a process of *mental deployment* (Bellet et al., 2009) of the current driving schemas has been implemented, by using a third specific instance of SiVIC.

Action Module

The Action Module is in charge to perform vehicle-control skills, according to the driving actions decided, anticipated and then planned at the representational level by the Cognition module. The two core regulation mechanisms implemented in the Action Module are the (i) *Pure-Pursuit Point* method and (ii) the *Envelope-Zones* regulation process. These vehicle-control abilities have been implemented on the SiVIC platform as a new type of the pre-existing SiVIC models of *vehicle controls* (Gruyer et al., 2006).

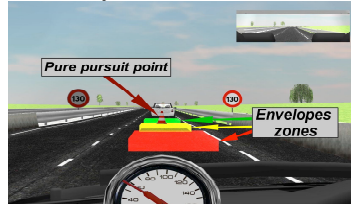


Figure 3: Pursuit Point and Envelope Zones on SiVIC platform

Indeed, a new class of “*COSMO-CAR*” has been defined, integrating the *pursuit point* and the *envelope zones*. Figure 6 illustrates such a regulation strategy in a car-following task: the pursuit point determines the car to be followed by the ego-car, and the envelope zones are used for keeping the IVT distance with the lead-car.

Model Results

In its current status, the COSMODRIVE model implemented on the SiVIC platform is able to observe, mentally analyse, decide and dynamically progress into a virtual road, through continuous actions on a virtual car. Indeed, model results take the form of dynamic simulations of the driver’s activity at four levels.

At the visual level (i.e. Perception module), by dynamic simulation of a sequence of visual fixation points, corresponding to the areas of interest successively explored by the driver while progressing on the road, according to its own tactical intentions, or as influenced by a visual secondary task while driving, requiring to alternate the road scene scanning and an on-board screen observation.

At the cognitive level (i.e. Cognition module), by dynamic elaboration of mental representations (i.e. situational awareness simulated through 3-Dimensional mental models of the road scene, integrating *driving schemas*, *envelope-zones* and *pure pursuit points* abilities), and decision-making processes simulations in charge to determine which relevant action should be implemented in the current driving context, as perceived, understood and anticipated by the driver model.

At the behavioural level (i.e. Action module), corresponding to the driver’s action actually performed on the virtual car commands (e.g. presented in Fig.9 through the curve describing the brake pedal status) for dynamically progressing into the virtual road environment and interacting with the other road users.

At the performance level as a whole, corresponding to the consequences of a dual “Perception-Cognition-Action” loop of regulation, continuously implemented by the driver model (e.g. respective speeds and positions of the vehicle and thus, Inter-Vehicular distances keeping), and which is dynamically simulated through the actual effects on driver’s action on the current driving situation, as virtually modelling into the SiVIC environment.

This last level of global performance, including potential critical consequences of human errors, is more particularly connected with the practical objectives of the Risk-Based Design methodology of driving aids to be implemented in the Isi-PADAS project, and that is focused on the human reliability issues. However, by considering the respective underlying simulations implemented by the *Perception*, the *Cognition*, and the *Action* modules, it becomes possible to investigate in detail human errors and thus to open the door for an “in-depth” understanding and analysis of the human driver’s reliability *versus* unreliability issues.

Conclusion: model use for virtual design

The research presented in this paper takes place in the frame of a Human Centred Design approach, aiming at setting up a virtual simulation platform to design and evaluate in-vehicle systems interest and potential impact on road safety (Bellet et al. 2010a). In this objective, it was proposed as to implement a cognitive simulation model of the driver on a Vehicle-Environment platform, in order to provide a simulation platform liable to support virtual design of vehicle automation technologies. This driver model implemented on the SiVIC platform aims to simulate human drivers’ perception, cognition and behaviour in order to dynamically progress in, and interact with, a virtual road environment. One objective during the second phase of the ISI-PADAS project will be to contribute to the Risk Based Design methodology defined by Cassini and Cacciabue (2010), and requiring human errors simulation results. Indeed, like a human driver, this model is not only able to simulate the driving performance in an optimal way, but is also able to generate *human errors* (e.g. non-perception of events, erroneous situational awareness and thus, decision-making, or inadequate behavioural performance in terms of safety margin keeping or reaction time), liable to occur in particular driving conditions. According to these functionalities, the model interest for driving assistances design will more particularly concern the initial design phases corresponding to the driving aid concept definition. At this earlier stage, our driver model could thus be used for virtual simulations allowing the designers to estimate human drivers’ performances in case of unassisted driving, in order to identify and specify the most critical driving scenarios for which the target-system to be developed should provide a palliative assistance. These critical scenarios will correspond to driving situations where the human driver reliability - as assessed from our driver model performance - seems not sufficient to adequately manage the risk. Through these scenario simulations, it could be thus possible to provide ergonomics specifications of drivers’ needs in terms of assistance. Then, during the driving aid testing phases, coming later in the design process, it could be therefore possible to evaluate the assistance effectiveness for the specific sub-set of most critical scenarios, as selected through the model simulations, in order to test the efficiency of this device (and, therefore, its interest for the drivers) in these particular driving conditions. These issues will be investigated in the ISI-PADAS project.

Acknowledgment

This research has received funding from the European Commission Seventh Framework Programme (FP7/2007-2013) under grant agreement n°218552 Project ISi-PADAS.

References

- Amidi O (1990), *Integrated mobile robot control*, Technical Report CMU-RI-TR-90-17, Carnegie Mellon University Robotics Institute.
- Bellet, T., Bailly-Asuni, B., Mayenobe, P., Banet, A. (2009). A theoretical and methodological framework for studying and modelling drivers' mental representations, *Safety Science*, 47, 1205–1221.
- Bellet, T., Bailly, B., Mayenobe, P., Georgeon, O. (2007). Cognitive modelling and computational simulation of drivers mental activities. In: P.C. Cacciabue (Ed.), *Modelling Driver Behaviour in Automotive Environment: Critical Issues in Driver Interactions with Intelligent Transport Systems*, 315-343, Springer Verlag.
- Bellet, T., Tattegrain-Veste, H. (1999). A framework for Representing Driving Knowledge. *International Journal of Cognitive Ergonomics*, 3 (1), 37-49.
- Cassani M., Cacciabue C. (2010). *Methodology for an improved Human Error Risk Analysis in design processes*, ISI-PADAS Deliverable D.5.2, 60 p.
- Endsley M.R. (1995). Toward a theory of situation awareness in dynamic systems. *Human Factors*, 37 (1), 32-64.
- Gibson J.J., Crooks L.E. (1938). A theoretical field-analysis of automobile driving. *American Journal of Psychology*, 51, pp. 453-471.
- Gruyer D., Royère C., du Lac N., Michel G., Blosseville J.M. (2006). SiVIC and RT-MAPS Interconnected platforms for the conception and the evaluation of driving assistance systems. *Proceedings of the ITS World Congress*, London, UK, october.
- Hall E.T. (1966). *The hidden dimension*. Doubleday, New York.
- Mayenobe P. (2004). *Perception de l'environnement pour une gestion contextualisée de la coopération Homme-Machine*. PhD thesis, University Blaise Pascal de Clermont-Ferrand..
- Michon J.A. (1985). A critical view of driver behavior models: what do we know, what should we do?. In Evans, L., Schwing, R.C. (Eds), *Human behavior and traffic safety*, pp. 485-520, Plenum Press, New York.
- Minsky M. (1975). A Framework for Representing Knowledge. In Winston, P.H. (Ed.), *The Psychology of Computer Vision*, 211-277, Mc Graw-Hill, New York.
- Neisser U. (1976). *Cognition and reality: principles and implications of cognitive psychology*. W. H. Freeman, San Francisco.
- Rasmussen J. (1986). *Information processing and human-machine interaction : an approach to cognitive engineering*, Amsterdam, North Holland.
- Scheinder W., Shiffrin R.M. (1977). Controlled and automatic human information processing: Detection, search and attention. *Psychological Review*, 84, 1-88.
- Sukthankar R. (1997). *Situation Awareness for Tactical Driving*, Phd thesis, Carnegie Mellon University, Pittsburgh, PA, United States of America.