

Dynamically Deployed Support as a Potential Solution to Negative Behavioral Adaptation

Rita Kovordányi, Kjell Ohlsson, and Torbjörn Alm

Abstract—Advanced driver assistance systems are designed to make driving easier that is, to alleviate the driver’s workload, and to increase traffic safety. However, traffic safety is affected by negative behavioral adaptation, meaning that drivers tend to increase speed and pay less attention to driving when supported by an advanced assistance system. We relate behavioral adaptation to reinforcement learning at a subconscious level, and propose that driver assistance is dynamically varied within predetermined safety limits. The aim of employing a dynamic assistance policy is to prevent the driver from noticing a constant improvement in vehicle handling. We conclude by describing ongoing work for empirically evaluating an improved lane departure warning system that uses a dynamic assistance policy.

Index Terms—advanced driver assistance systems, lane departure warning systems, lane keeping assistance systems, negative behavioral adaptation, reinforcement learning, dynamic assistance policy

I. INTRODUCTION

The evolution of a new generation of advanced driver assistance systems is to a large extent propelled by advances in sensor technology, steadily increasing computing power and fast algorithms for analyzing in real-time multiple sources of sensor data [1]-[3]. Whereas much emphasis is put on the safety and reliability of the technical system, there is little concern of the human driver. However, the driver is of central importance for traffic safety. Empirical studies indicate that drivers have a tendency to adapt their driving style and misuse the increased safety margins created by advanced driver assistance systems, a phenomenon called *negative behavioral adaptation* [4]-[7]. Drivers may, for example, increase the driving speed and pay less attention to the driving task, to such an extent that the safety margins created by the driver assistance system are cancelled out [8], [9]. While this problem is widely acknowledged, to date no technically-

based, engineering solutions have been proposed that could mitigate the adverse safety effects of negative behavioral adaptation.

II. TECHNICAL SYSTEM EFFICIENCY

In today’s serial production cars, drivers are offered assistance with lateral vehicle control. The assistance offered ranges from simple warning to active contribution to steering.

At one end of the scale, *Lane Departure Warning Systems (LDWS)* alert the driver when *Time To Line Crossing (TTL)*—the time it would take the wheels of the car to cross the lane boundary if the car would continue to move along its present trajectory—falls below a predetermined threshold. Whereas some LDWS use directional auditory warnings in these critical situations, other systems employ a more discreet warning, for example, by vibrating the driver’s seat (such as the Citroën C5). Yet other systems take into account the stimulus-response compatibility [9] of warnings, making sure that the warning produced by the system is naturally associated with the right kind of reaction. For example, vibration of the steering wheel activates the driver’s hands, which in turn can shorten the driver’s reaction time for executing a corrective steering maneuver.

At the other end of the scale, *Lane Keeping Assistance Systems (LKAS)* can actively contribute to steering. These systems use sophisticated algorithms that take into account the projected path of the vehicle to determine the steering wheel torque that is required for bringing back the car into the lane [2], [3]. These systems are designed to blend seamlessly with the driver’s steering action, and provide up to a predetermined maximum portion (e.g., 80%) of the required torque. Hence, the system’s output is combined with the driver’s steering maneuvers to produce the final steering output. This solution contributes actively to steering and at the same time keeps the driver in the loop [11].

These systems are designed to make driving easier that is, to alleviate the driver’s workload and to increase traffic safety by improving lateral vehicle control. However, these top-of-the-line systems are exposed to negative behavioral adaptation, since drivers have a demonstrated tendency to adapt a more reckless driving style (i.e., increase driving speed and pay less attention to driving) during assisted driving.

Manuscript received December 15, 2004.

Rita Kovordányi, Ph.D. is assistant professor in cognitive technology at the Department of Computer and Information Science, Linköpings universitet, Linköping, SE-581 83 Sweden (phone: +46 13 281430; fax: +46 13 142231; e-mail: ritko@ida.liu.se).

Kjell Ohlsson, Ph.D. is professor in human-machine interaction at the division of Industrial Ergonomics, Department of Mechanical Engineering, Linköpings universitet, Linköping, SE-581 83 Sweden (e-mail: kjeoh@ikp.liu.se).

Torbjörn Alm is lecturer at the division of Industrial Ergonomics, Department of Mechanical Engineering, Linköpings universitet, Linköping, SE-581 83 Sweden (e-mail: toral@ikp.liu.se).

III. TRAFFIC SAFETY IS AFFECTED BY THE DRIVER'S BEHAVIOR

Traffic safety is affected not only by the assistance system's technical efficiency but also by the way the driver adapts to the new system. Behavioral adaptation refers to a basic ability to modify behavior dynamically that is, on the fly, to meet the demands of changing circumstances. Normally, behavioral adaptation enhances an individual's chances of survival in a dynamically changing world. When it comes to traffic safety, however, behavioral adaptation can in addition have negative consequences. A driver may, for example, become more reckless and exhibit poorer lateral control over the vehicle after he/she learns to rely on the support provided by advanced driver assistance systems [8], [14].

This change in behavior is aimed at keeping constant a subjectively perceived risk level, according to risk homeostasis theory [15]. The zero risk theory elaborates this view, and states that drivers strive to minimize the subjective risk level by behaving so that the subjectively perceived risk level remains at zero [16]. The subjective risk level is in turn determined by a dynamic balance between a number of positive and negative factors, such as how the car handles, whether the driver is concentrated on driving or not, whether the driver is self confident, et cetera. These factors are in turn affected by the driver's recent actions and experiences. For example, if the driver has recently tried to make an overtaking maneuver, and was close to colliding with a meeting car, the driver's subjectively perceived risk level would be adjusted upwards. These minor chock-effects can stay on and raise the perceived risk level for extended periods after the incident [17]. According to [17], whenever the perceived risk level is above zero, the driver will take precautionary measures, such as decreasing speed and/or allocating more attention to driving to bring back the perceived risk level to zero. Likewise, when the risk level drops below zero, the driver would often use this extra margin to increase speed and/or allocate less attention to driving.

Other theories focus on the relationship between the driver's perceptual and motor processes and the driver's assessment of available time to execute maneuvers [18]-[21]. Available time could be reflected by TTC (Time To Collision) or TTLC (Time To Line Crossing, [22]). The driver feels that he/she is in control as long as available time matches the time needed for maneuvering the car.

A. Behavioral adaptation can occur at multiple levels

Behavioral adaptation has been observed for lane departure warnings systems that warn the driver when he/she is too close to the edge of the road, also indicating in which direction the steering wheel should be turned [4], [5], [7], [23]). Behavioral adaptation has also been observed in other domains using relatively simple experimental settings, such as a simulated process control task in a chemical plant [24]. There is therefore reason to believe that behavioral adaptation can be studied in driving simulators. This is important, as for safety reasons it is awkward to conduct experiments in real traffic when drivers' reaction to critical situation is studied.

The breadth of empirical studies on behavioral adaptation [6], [24], [25] suggests that this phenomenon can occur at multiple scales, both at a low, perceptual-motor level and at a higher cognitive level (cf. [17]). Of these two levels, low-level behavioral adaptation seems to be more dangerous, because contrary to processes that occur at a higher cognitive level, lower-level motor-perceptual processes are out of reach for the driver's conscious control. For this reason, we focus on low-level behavioral adaptation in the present paper.

Based on the above theories, we adopt the working hypothesis that drivers employ an individually chosen risk margin. This margin is continuously compared with the difference between available time and the time needed for maneuvering the car. This comparison forms the basis of the driver's choice of speed, allocation of attention and other higher-level tactical and strategic decisions. For example, precautionary drivers would want to keep the time needed to maneuver well within subjectively perceived available time. When this is not case, the driver will tend to decrease speed and/or attend more to driving than before.

IV. SUGGESTED UNDERLYING MECHANISMS

We propose that negative behavioral adaptation reflects basic properties of reinforcement learning that is naturally employed by humans [26]. Biologically based reinforcement learning is an unsupervised form of learning based on an internal system for rewarding successful behavior. For example, a need for change in vehicle state (e.g., if the car is drifting off the road) makes the driver execute an appropriate corrective maneuver (turn the steering wheel). The timing of this action and the steer angle is determined on the basis of the driver's previous experiences on what kind of outcome a particular timing and steer angle would result in. Hence, the driver expects that the chosen torque will achieve the change in vehicle state that was desired (cf. Fig. 1).

After execution, the driver observes the outcome (Δ_2 in Fig. 1) of the maneuver and compares the outcome with previous expectations (Δ_1 in Fig. 1). If the actual outcome matches with previous expectations, an internal reward will be issued that strengthens those processes that were responsible for executing the correct action, and in this way reinforces the behavior for future use (1). If the outcome does not match with expectations, a negative reward is issued, which makes the driver adjust his/her expectations according to the actual outcome. The size of this adjustment depends on the reward r (which ranges from -1 to 1), the expected reward $V(t)$, and the actual reward issued internally $V(t+1)$ (cf. 1). Here, γ is a weighting factor that determines to what extent the possibility of future rewards should be considered when adjusting current behavior.

$$\delta(t) = (r(t) + \gamma \hat{V}(t+1)) - \hat{V}(t) \quad (1)$$

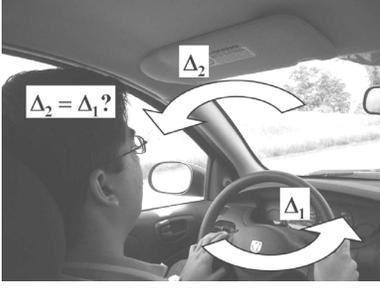


Fig. 1. The motor-perceptual loop forms a basis for executing actions and comparing the outcome of maneuvers with predicted outcome. When the outcome is better than expected, the driver can use this extra margin to increase speed and/or pay less attention to driving.

In cases when the outcome is persistently better than expected, the driver will consider this as an improved margin for maneuvers. This extra margin can be used or misused to allow for more slack in the timing of maneuvers, and since correcting maneuvers do not have to be executed as frequently and as precisely as before, to increase driving speed and to pay less attention to the driving task.

Hence, reinforcement learning can explain how drivers can keep constant a subjectively perceived risk level: They do not actually calculate the risk; they simply adjust those factors (e.g., driving speed) that affect the available time to execute maneuvers. Drivers can do this by observing and comparing the outcome of maneuvers. On the basis of the above discussion, it is also reasonable to assume, as was originally proposed by [16], that risk level is subjectively perceived to be at zero.

A. The effect of adaptive driver support systems

We propose that adaptive support systems interfere with the motor-perceptual loop that is at the foundation of risk perception and behavioral adaptation. At the outset the driver's expectations are based on manually executed previously unsupported maneuvers. When an assistance system is introduced, the driver's maneuvers (or even the lack of maneuvers) will turn out to yield a much better outcome than expected because the driver assistance system executes correcting maneuvers (this would be true for systems like adaptive cruise controls and LKAS that directly help with the driving task). The driver experiences through perceptual feedback that the outcome of maneuvers is persistently better than expected. This will cause the driver's motor-perceptual loop to become biased so that it now reflects a skewed relationship between the driver's maneuvers and outcomes that incorporates the contribution of the support system. When the driver notices the extra margins that arise, he/she will adjust the driving style accordingly (sloppier, less attentive driving).

As described in section III.A, behavioral adaptation can occur at multiple levels. A driver may, for example, decide to increase driving speed when under time pressure, or choose to ignore visible signs for adverse road conditions for the same reason. As these decisions are made on a conscious level, they are also subject for external influences, for example, through legislations and speed limits.

At another level, the driver constantly adapts his/her behavior on the basis of a motor-perceptual loop where the outcome of maneuvers is compared to expectations. The processes that are at work at this level cannot be controlled consciously, which aggravates the problem. The driver not only risks driving unsafely, he/she is also largely unaware of the problem.

We propose two ways of mitigating negative behavioral adaptation. One way is to present the driver with explicit feedback on inappropriate behavior, in this way bringing the result of the behavior to the user's conscious attention. A driver could, for example be warned if he/she was driving with higher speed, or was paying less than usual attention to the road scene. We note that recognition of deviant driving behavior has to be performed on an individual basis, which requires advanced user modeling.

A second way to improve today's driver assistance systems is to change or eliminate those system characteristics that trigger negative behavioral adaptation in the unsuspecting driver. In particular, based on our previous analyses of the role of reinforcement learning in behavioral adaptation, the assistance provided by a driver assistance system should be adjusted according to a dynamic assistance policy. For lane departure warning systems, the threshold for activating a warning can be randomly varied within a predetermined range. Let $TTLC$ (Time To Line Crossing) denote the time it would take the wheels of the car to cross the lane boundary. Normally, a warning is activated when $TTLC$ falls below a statically determined threshold $\tau_{cautious}$, τ_{danger} :

$$\begin{cases} TTLC < \tau_{cautious} & \rightarrow \text{cautious warning} \\ TTLC < \tau_{danger} & \rightarrow \text{danger warning} \end{cases} \quad (2)$$

In order to mitigate behavioral adaptation, a stochastically adjusted dynamic threshold τ_{dyn} for activating cautious warnings could be used instead of the static threshold (Fig 2):

$$\begin{aligned} TTLC < \tau_{dyn}, \quad \tau_{dyn} = rnd(\tau_{min}, \tau_{max}), \\ \text{where } \tau_{min} = \tau_{danger} + \epsilon, \quad \tau_{max} = \tau_{cautious} \end{aligned} \quad (3)$$

We note that that the time-span between τ_{min} and τ_{max} must be kept sufficiently large—approximately on the order of 2 s—for drivers to perceive the variation in assistance deployment. For this reason, it may also be necessary to vary ϵ depending on the driving speed, so that sufficient time span is

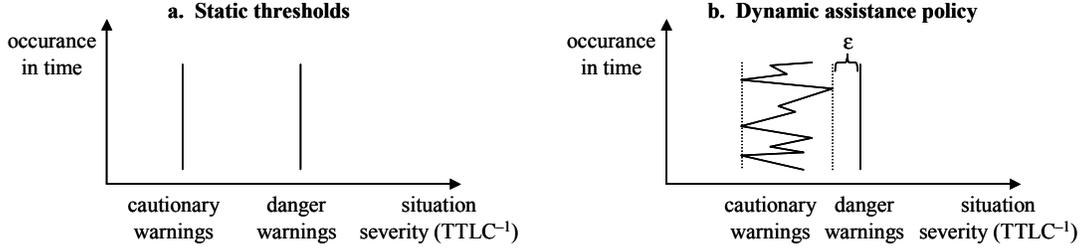


Fig. 2. a. Traditional LDWS using static thresholds. b. LDWS where the threshold for cautionary warnings are varied across occasions according to a dynamic assistance policy.

ensured also during highway driving.

Likewise, for lane keeping assistance systems, if the steer angle required for bringing back the car to the lane is δ_{req} , the amount of steer angle, δ_{sys} , that normally would be provided by the system can be described as:

$$\delta_{sys} = q\delta_{req}, \text{ where } q \in [0, q_{max}] \quad (4)$$

Here q is the portion of steer angle that is provided by the LKAS; the remaining portion of δ_{req} is provided by the driver. A useful value for q_{max} has proven to be 0.8 [11]. In today's LKAS q is adaptively set to smaller values when the driver turns the steering wheel forcefully [11], in order to make it easy for the driver to override the system. After q has been calculated, steering assistance can be dynamically attenuated (decreased) in order to avoid triggering behavioral adaptation in the driver, according to the following scheme:

$$\delta_{sys} = q_{dyn}\delta_{req}, \text{ where } q_{dyn} = rnd(c_{min}q, q), \text{ } c_{min} \in [0, 1] \quad (5)$$

The aim of employing a dynamic assistance policy is to prevent the driver from experiencing a constantly available, readily predictable support that could induce the driver to adjust the internal mapping that the driver normally maintains between maneuvers and their expected outcome.

In essence, a dynamic assistance policy utilizes the fact that reinforcement learning requires a consistency in the outcome of actions. Somewhat simplified, if the outcome of actions has not been consistently changed, there is no informational basis for deciding whether to change behavior in this way or that way. Therefore, reinforcement learning cannot take place, if the driver does not experience a *systematic* shift in the outcome of maneuvers towards the better.

At the other side of the coin, by employing a dynamic assistance policy, the system provides sub-optimal assistance. How drivers would actually react to an assistance system that employed a dynamic assistance policy must therefore be empirically studied.

VI. ONGOING DRIVING SIMULATOR EXPERIMENTS

We are preparing a simulator study with the aim to evaluate the traffic safety of an LDWS [27]. Mainly, we want to see if behavioral adaptation is effectively mitigated by a dynamic assistance policy.



Fig. 3. The fixed-base driving simulator used in our study is software-oriented and allows for rapid prototyping of new driver assistance systems and HMI-features.

A. The LDWS used in the simulator study

The LDWS that will be used in our experiments is currently being evaluated by Scania, a Swedish truck manufacturer, for future deployment in serial production vehicles. The system employs directional auditory warnings, based on a static threshold for TTLC (Time To Line Crossing). As the wheels of big trucks are always close to the lane boundaries, there is a desire to decrease the threshold for warnings in order to decrease the number of false alarms, without also decreasing traffic safety.

At the same time, a *dynamically* decreased threshold for cautionary warnings, where the threshold is allowed to vary between two predetermined limits, would mitigate drivers' tendency for negative behavioral adaptation. We therefore hope to achieve the double aim of decreasing the number of false alarms in the system, and increasing traffic safety by mitigating negative behavioral adaptation in the driver.

B. Experimental design

20 participants will drive through a 45 km test-route in a fixed-base driving simulator (Fig. 3). Each participant will be tested in three conditions in a balanced design:

1. Unassisted driving
2. Traditional lane departure warning system using static thresholds
3. Improved system using a dynamic assistance policy

Using these three conditions, we intend to test the traffic safety of the improved system by comparing driver performance with the traditional system. In addition, we intend to relate driver performance with the improved system to base level performance during unassisted driving.

In order to force the need for manual intervention the car's lateral position will be artificially perturbed in a number of critical scenarios.

We want to measure participants' performance in these scenarios with respect to:

- Average driving speed
- Standard deviation of lateral position
- Mental workload (as measured by NASA-TLX)
- TTLC at the moment when the participant reacts to the car's lateral position being artificially perturbed (start of correcting steering maneuver)
- Steer angle immediately following participants' first reaction (how abruptly participants correct the car's position)
- Minimum TTLC (with what margin participants manage to avoid running off the road)

In addition, participants will be asked to fill out a questionnaire for the NASA-TLX scale [28], the Driving Internality Externality Scales [29] and the Sensation Seeking Scale [30]; the last two because drivers who tend to relinquish control to the technical system (experience external locus of control) and/or are liable to sensation seeking seem to be especially prone to negative behavioral adaptation [6].

C. Expected results

Based on our account for behavioral adaptation (see section IV), we expect improved traffic safety effects for LDWS using a dynamic assistance policy. In particular, we expect that drivers using an LDWS with dynamic support policy will exhibit:

- Lower average driving speed
- Higher mental workload (which would indicate that the driver is more focused on driving)
- Larger TTLC (i.e., faster reaction from the driver when the car's lateral position is artificially perturbed)
- Softer steering (i.e., less panic reactions)
- Increased minimum TTLC (i.e., participants are able to bring the car back to the lane earlier when using the improved LDWS)

VII. SUMMARY

The intended safety effect of advanced driver assistance systems, such as Lane Departure Warning Systems (LDWS), can be eliminated through negative behavioral adaptation in drivers. As an example, drivers seem to increase driving speed and pay less attention to driving when using an LDWS [4], [5].

We propose that negative behavioral adaptation arises in the driver's motor-perceptual loop. This loop controls the execution of maneuvers and mediates perceptual feedback to the driver on the outcome of these maneuvers. When a driver assistance system is introduced, car handling will improve which will cause the driver to subconsciously adjust to the consistent discrepancy between expected and actual outcome. Simply expressed, drivers learn to expect better vehicle control even when executing less precise maneuvers less frequently.

In order to mitigate negative behavioral adaptation, we propose that advanced driver assistance systems employ a dynamic assistance policy that prevents the driver from anticipating improved car handling.

REFERENCES

- [1] S. Tokoro, K. Morizumi, T. Kawasaki, T. Nagao, K. Abe, and K. Fujita, "Sensor fusion system for pre-crash safety system", in *Proc. 2004 IEEE Intelligent Vehicles Symposium*, Parma, 2004, pp. 945-950.
- [2] H. Fritz, A. Gern, H. Schiemenz, and C. Bonnet, "A driver assistance system for commercial vehicles based on fusion of advanced ACC and lane keeping", in *Proc. 2004 IEEE Intelligent Vehicles Symposium*, Parma, 2004, pp. 495-500.
- [3] G. Meier, G. Roppenecker, and C. Wurmthaler, "Automatic lateral vehicle guidance using tracking control", in *Proc. 2004 IEEE Intelligent Vehicles Symposium*, Parma, 2004, pp. 145-149.
- [4] C. M. Brown, "The concept of behavioural adaptation: Does it occur in response to lane departure warnings?", in *Proc. Intl. Conf. Traffic and Transport Psychology*, pp. 4-7, 2000.
- [5] C. M. Rudin-Brown and Y. I. Noy, "Investigation of behavioral adaptation to lane departure warnings", *Transportation Research record*, vol. 1803, pp. 30-37, 2002.
- [6] C. M. Rudin-Brown and H. A. Parker, "Behavioral adaptation to adaptive cruise control (ACC): implications for preventive strategies", *Transportation Research*, part F, vol. 7, 2004, pp. 59-76.
- [7] P. C. Burns, "Behavioural Adaptation to an Advanced Driver Support System", Volvo Technology Corporation, Internal Report, 2001.
- [8] Z. Jęfć, J. Engström, and P. Piamonte, "Potential safety benefits of lane departure warning systems on Swedish roads: Pre-study", Swedish National Road Agency, Publication 2003:14, ISBN: 1401-9612, 2003.
- [9] H. Summala, "Behavioural adaptation and drivers' task control", in *Human factors for highway engineers*, R. Fuller and J. A. Santos, Eds. Oxford: Pergamon Press, 2002.
- [10] M. A. Vidulich and C. D. Wickens, "Stimulus-central processing-response compatibility: Guidelines for the optimal use of speech technology", *Behavior Research Methods, Instruments and Computers*, vol. 17, no. 2, 1985, pp. 243-249.
- [11] S. Ishida and J. E. Gayko, "Development, evaluation and introduction of a lane keeping assistance system", in *Proc. 2004 IEEE Intelligent Vehicles Symposium*, Parma, 2004, pp. 943-944.
- [12] R. Kovordányi, "When will advanced driver support systems be user-adaptive? The case of adaptive cruise control", in *Proc. Of AAAI Spring Symposium Series, Challenges to Decision Support in a Changing World*, to be published.
- [13] R. Kovordányi and S. Ohlsson, "Toward adaptive support: Modelling drivers' allocation of attention", in *Proc. of the 34th Annual Congress of the Nordic Ergonomic Society*, 2002, pp. 493-498.

- [14] OECD, "Behavioural adaptation to changes in the road transport systems", OECD, Paris, 1990.
- [15] G. J. S. Wilde, "The theory of risk homeostatis: Implications for traffic safety and health", *Risk Analysis*, vol. 2, 1982, pp. 209-225.
- [16] H. Summala and R. Näätänen, "The zero-risk theory and overtaking decisions," in *Road user behavior: Theory and research*, T. Rothengatter and R. de Bruin, Eds. Van Gorcum, Assen / Maastricht. 1988, pp. 82-92.
- [17] H. Summala, "Hierarchical model of behavioural adaptation and traffic accidents", in *Traffic and Transport Psychology: Theory and application*, T. Rothengatter and E. C. Vaya, Eds. Oxford: Pergamon Press, 1997, pp. 41-52.
- [18] D. N. Lee, "A theory of visual control of braking based on information about time to collision", *Perception*, vol. 5, 1976, pp. 437-459.
- [19] D. T. McRuer, R. W. Allen, D. H. Weir, and R. H. Klein, "New Results in Driver Steering Control Models". *Human Factors*, vol. 19, pp. 381-397, 1997.
- [20] E. Donges, "A Two Level Model of Driver Steering Behaviour", *Human Factors*, vol. 20, 1978, pp. 691-707.
- [21] P. C. Cacciabue, "Modelling and simulation of human behaviour in system control". London: Springer Verlag, 1998.
- [22] W. van Winsum, "From adaptive control to adaptive behaviour." PhD thesis, University of Groningen, Netherlands, 1996.
- [23] K. Suzuki and H. Jansson, "An analysis of drivers' steering behaviour during auditory and haptic warnings for the designing of lane departure warning systems." *Journal of the Society of Automotive Engineers of Japan Review*, vol. 24, no. 1, 65-70, 2003.
- [24] Johansson, B. and Rigas, G. (2004) Functional failures, time and control. In (Eds.) Thissen W, Wieringa P, Pantic M & Ludema M. (eds): *Proceedings of the 2004 IEEE International Conference on Systems, Man & Cybernetics*. Omnipress, The Netherlands.
- [25] L. Nilsson, "Safety effects of adaptive cruise controls in critical traffic situations", in *Proc. The Second World Congress on Intelligent Transport Systems*, 1995, pp. 1254-1259.
- [26] R. C. O'Reilly, "Six principles for biologically-based computational models of cortical processing" *Trends in Cognitive Sciences*, vol. 2, no. 11, 1998, pp. 455-462.
- [27] T. Alm, R. Woltjer, K. Ohlsson, and R. Kovordányi, "Advanced driver interface design and evaluation", in *Proc. Driving Assessment*, submitted for publication.
- [28] S. Hart and L. E. Staveland, "Development of NASA-TLX (Task Load Index): results of empirical and theoretical research", in P. A. Hancock and N. Meshkati, Eds. *Human Mental Workload*. Amsterdam: Elsevier Science Publishers, 1988.
- [29] I. Montag, A. L. Comrey, "Internality and externality as correlates of involvement in fatal driving accidents", *Journal of Applied Psychology*, vol. 72, 1987, pp. 339-343.
- [30] M. Zuckerman, "Behavioral expressions and biosocial bases of sensation seeking". Cambridge: University of Cambridge Press, 1994.