Experimental Study of an Optimal-Control-Based Framework for Trajectory Planning, Threat Assessment, and Semi-Autonomous Control of Passenger Vehicles in Hazard Avoidance Scenarios

Sterling J. Anderson, Steven C. Peters, Tom E. Pilutti, and Karl Iagnemma*

Abstract. This paper describes the design of an optimal-control-based active safety framework that performs trajectory planning, threat assessment, and semi-autonomous control of passenger vehicles in hazard avoidance scenarios. The vehicle navigation problem is formulated as a constrained optimal control problem with constraints bounding a navigable region of the road surface. A model predictive controller iteratively plans an optimal vehicle trajectory through the constrained corridor. Metrics from this "best-case" scenario establish the minimum threat posed to the vehicle given its current state. Based on this threat assessment, the level of controller intervention required to prevent departure from the navigable corridor is calculated and driver/controller inputs are scaled accordingly. This approach minimizes controller intervention while ensuring that the vehicle does not depart from a navigable corridor of travel. It also allows for multiple actuation modes, diverse trajectory-planning objectives, and varying levels of autonomy. Experimental results are presented here to demonstrate the framework's semi-autonomous performance in hazard avoidance scenarios.

1 Introduction

Recent traffic safety reports from the National Highway Traffic and Safety Administration show that in 2007 alone, over 41,000 people were killed and another 2.5 million injured in motor vehicle accidents in the United States [1]. The

Sterling J. Anderson, Steven C. Peters, and Karl Iagnemma
Department of Mechanical Engineering, Massachusetts Institute of Technology,
77 Massachusetts Ave. Cambridge, MA 02139, USA
e-mail: {ster,scpeters,kdi}@mit.edu

Tom E. Pilutti
Ford Research Laboratories, Dearborn, MI 48124, USA
e-mail: tpilutti@ford.com

* Corresponding author.

A. Howard et al. (Eds.): Field and Service Robotics 7, STAR 62, pp. 59–68.
springerlink.com

longstanding presence of passive safety systems in motor vehicles, combined with the ever-increasing influence of active systems, has contributed to a decline in these numbers from previous years. Still, the need for improved collision avoidance technologies remains significant.

Recent developments in onboard sensing, lane detection, obstacle recognition, and drive-by-wire capabilities have facilitated active safety systems that share steering and/or braking control with the driver [2,3]. Among existing proposals for semi-autonomous vehicle navigation, lane-keeping systems using audible warnings [4], haptic alerts [5], steering torque overlays [6], and various combinations of these have been developed [7].

Many of the navigation systems developed in previous work address only one piece of the active safety problem. While some use planning algorithms such as rapidly-exploring random trees [3], evolutionary programming [8] or potential fields analysis [9] to plan a safe vehicle path, others simply begin with this path presumed [10]. The threat posed by a particular path is seldom assessed by the controller itself and is often only estimated by a simple threat metric such as lateral vehicle acceleration required to track the path [11]. Finally, hazard avoidance is commonly performed using one or more actuation methods without explicitly accounting for the effect of driver inputs on the vehicle trajectory. Such controllers selectively replace (rather than assist) the driver in performing the driving task. Yu addressed this problem in mobility aids for the elderly by designing an adaptive shared controller which allocates control authority between the human user and a controller in proportion to the user's current and past performance [12]. While sufficient to control low-speed mobility aids, this reactive approach to semi-autonomy is not well suited for higher-speed applications with significant inertia effects and no pre-planned trajectory.

In this paper, a framework for passenger vehicle active safety is developed that performs vehicle trajectory planning, threat assessment, and hazard avoidance in a unified manner. This framework leverages the predictive and constraint-handling capabilities of Model Predictive Control (MPC) to plan trajectories through a preselected corridor, assess the threat this path poses to the vehicle, and regulate driver and controller inputs to maintain that threat below a given threshold. The next section describes the semi-autonomous control framework and its associated trajectory prediction, control law, threat assessment, and intervention law. Experimental setup and results are then presented, and the paper closes with general conclusions and recommendations.

2 Framework Description

The navigation framework operates as follows. First, an objective function is established to capture desirable performance characteristics of a safe/"optimal" vehicle path. Boundaries tracing the edges of the drivable road surface are assumed to have been derived from forward-looking sensor data and a higher-level corridor planner. These boundaries establish constraints on the vehicle's projected position and are used together with a model of the vehicle dynamics to calculate an optimal sequence of inputs and the associated vehicle trajectory. The predicted trajectory

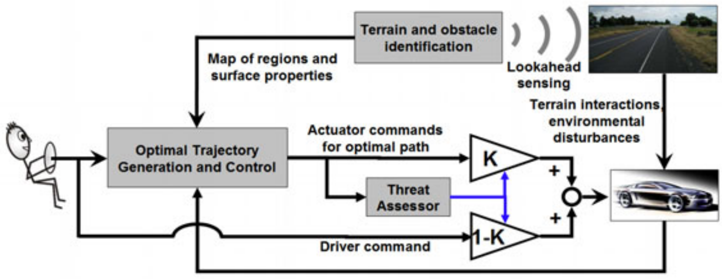

Fig. 1 Diagram of an active safety system.

is assumed to be a "best-case" scenario that poses the minimum threat to the vehicle given its current state. This threat is then used to calculate the intervention required to prevent departure from the navigable corridor and driver/controller inputs are scaled accordingly. Fig. 1 shows a block diagram of this system.

In this paper it is assumed that road lane data is available and that road hazards have been detected, located, and mapped to form the boundaries of a 2-dimensional corridor of travel. Radar, LIDAR, and vision-based lane-recognition systems [3,13], along with various sensor fusion approaches [14] have been proposed to provide the lane, position, and environmental information needed by this framework. Additionally, where multiple corridor options exist, it is assumed that a high-level path planner has selected a single corridor through which the vehicle should travel.

2.1 Vehicle Path Planning

The best-case (or baseline) path through a given region of the state space is established by a model predictive controller. Model Predictive Control is a finite-horizon optimal control scheme that iteratively minimizes a performance objective defined for a forward-simulated plant model subject to performance and input constraints. At each time step, t, the current plant state is sampled and a cost-minimizing control sequence spanning from time t to the end of a control horizon of n sampling intervals, $t+n\Delta t$, is computed subject to inequality constraints. The first element in this input sequence is implemented at the current time and the process is repeated at subsequent time steps. The basic MPC problem setup is described in [15].

The vehicle model used in this paper considers the kinematics of a 4-wheeled vehicle, along with its lateral and yaw dynamics. Vehicle states include the position of its center of gravity $[x, y]$, its yaw angle ψ, yaw rate $\dot{\psi}$, and sideslip angle β, as illustrated in Fig. 2. Table 1 defines and quantifies this model's parameters.

Tire compliance is included in the model by approximating lateral tire force (F_y) as the product of wheel cornering stiffness (C) and wheel sideslip (α or β for front or rear wheels respectively) as shown in (1).

$$F_y = C\alpha \tag{1}$$

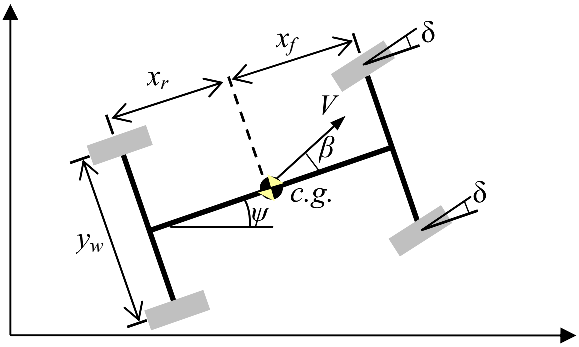

Fig. 2 Vehicle model used in MPC controller.

Table 1 Vehicle model parameters.

Symbol	Description	Value [units]
m	Total vehicle mass	2050 [kg]
I_{zz}	Yaw moment of inertia	3344 [kg·m^2]
x_f	C.g. distance to front wheels	1.43 [m]
x_r	C.g. distance to rear wheels	1.47 [m]
y_w	Track width	1.44 [m]
C_f	Front cornering stiffness	1433 [N/deg]
C_r	Rear cornering stiffness	1433 [N/deg]
μ	Surface friction coefficient	1

Linearized about a constant speed and assuming small slip angles, the equations of motion for this model are (where δ represents the steering angle input)

$$\dot{x} = V \tag{2}$$

$$\dot{y} = V(\psi + \beta) \tag{3}$$

$$\dot{\beta} = \frac{-(C_r + C_f)}{mV}\beta + \left(\frac{(C_r x_r - C_f x_f)}{mV^2} - 1\right)\dot{\psi} + \frac{C_f}{mV}\delta \tag{4}$$

$$\ddot{\psi} = \frac{(C_r x_r - C_f x_f)}{I_{zz}}\beta - \frac{(C_r {x_r}^2 + C_f {x_f}^2)}{I_{zz}V}\dot{\psi} + \frac{C_f x_f}{I_{zz}}\delta \tag{5}$$

where C_f and C_r represent the cornering stiffness of the lumped front wheels and the lumped rear wheels, and x_f and x_r are the longitudinal distances from the c.g. of the front and rear wheels, respectively.

2.1.1 Constraint Setup and Objective Function Description

As mentioned above, this framework assumes that the environment has been delineated previously. The boundaries of the navigable road surface at each timestep are then described by the constraint vectors

$$\begin{aligned}\mathbf{y}^{y}{}_{\max}(k) &= \left[y^{y}{}_{\max}(k+1) \quad \cdots \quad y^{y}{}_{\max}(k+p)\right]^{T} \\ \mathbf{y}^{y}{}_{\min}(k) &= \left[y^{y}{}_{\min}(k+1) \quad \cdots \quad y^{y}{}_{\min}(k+p)\right]^{T}\end{aligned} \tag{6}$$

By enforcing vehicle position constraints at the boundaries of the navigable region of the road surface (i.e. the lane edges on an unobstructed road), the controller forces the MPC-generated path to remain within the constraint-bounded corridor whenever dynamically feasible. Coupling this lateral position constraint with input constraints $\mathbf{u}_{\text{min/max}}$, input rate constraints $\Delta\mathbf{u}_{\text{min/max}}$, and vehicle dynamic considerations, the navigable operating corridor delineated by $\mathbf{y}^{y}{}_{\max}$ and $\mathbf{y}^{y}{}_{\min}$ translates to a safe operating region within the state space.

The controller's projected path through the constraint-imposed tube is shaped by the performance objectives established in the MPC cost function. While many options exist for characterizing desirable vehicle trajectories, here, the total sideslip angle at the front wheels (α) was chosen as the trajectory characteristic to be minimized in the objective function. This choice was motivated by the strong influence front wheel sideslip has on the controllability of front-wheel-steered vehicles since cornering friction begins to decrease above critical slip angles. In [16] it is shown that limiting tire slip angle to avoid this strongly nonlinear (and possibly unstable) region of the tire force curve can significantly enhance vehicle stability and performance. Further, the linearized tire compliance model described here does not account for this decrease, motivating the suppression of front wheel slip angles to reduce controller-plant model mismatch.

The MPC objective function with weighting matrices $R_{(\cdot)}$ then takes the form

$$J_k = \sum_{i=k+1}^{k+p} \frac{1}{2}\alpha_i^{\mathbf{T}} R_\alpha \alpha_i + \sum_{i=k}^{k+p-1} \frac{1}{2}\delta_i^{\mathbf{T}} R_\delta \delta_i + \sum_{i=k}^{k+p-1} \frac{1}{2}\Delta\delta_i^{\mathbf{T}} R_{\Delta\delta} \Delta\delta_i + \frac{1}{2}\rho_\varepsilon \varepsilon^2 \tag{7}$$

where ε represents constraint violation and was included to soften select position constraints as $\mathbf{y}^{j}{}_{\min} - \varepsilon\mathbf{V}^{j}{}_{\min} \le \mathbf{y}^{j} \le \mathbf{y}^{j}{}_{\max} + \varepsilon\mathbf{V}^{j}{}_{\max}$.

2.2 *Threat Assessment*

The vehicle path calculated by the MPC controller is assumed to be the best-case or safest path through the environment. As such, key metrics from this prediction are used to assess instantaneous threat posed to the vehicle. By setting constraint violation weights (ρ_ε) significantly higher than the competing minimization weight (R_α) on front wheel sideslip, optimal solutions satisfy corridor constraints before minimizing front wheel sideslip. When constraints are not active, front wheel sideslip – and the corresponding controllability threat – is minimized. When the solution is constrained, predicted front wheel sideslip increases with the severity of the maneuver required to remain within the navigable corridor.

Various approaches are available to reduce the predicted front wheel sideslip vector $\boldsymbol{\alpha}$ to a scalar threat metric Φ. In this paper,

$$\Phi(k) = \max\left(\left|\alpha_{k+1} \quad \alpha_{k+2} \quad \cdots \quad \alpha_{k+p}\right|\right)^T \tag{8}$$

was chosen for its good empirical performance when used to regulate controller intervention (described in the next section).

2.3 Hazard Avoidance

Given a best-case vehicle path through the environment and a corresponding threat, desired inputs from the driver and controller are blended and applied to the vehicle. This blending is performed based on the threat assessment: a low predicted threat causes more of the driver's input and less of the controller's input to be applied to the vehicle, while high threat allows controller input to dominate that of the driver. This "scaled intervention" may thereby allow for a smooth transition in control authority from driver to controller as threat increases.

Denoting the current driver input by u_{dr} and the current controller input by u_{MPC}, the blended input seen by the vehicle, u_v, is defined as

$$u_v = K(\Phi)u_{MPC} + (1 - K(\Phi))u_{dr} \tag{9}$$

The intervention function K is used to translate predicted vehicle threat Φ (obtained from the MPC trajectory plan) into a scalar blending gain. This function is bounded by 0 and 1 and may be linear, piecewise-linear, or nonlinear. Linear and piecewise-linear forms of this function may be described by

$$K = f(\Phi) = \begin{cases} 0 & 0 \le \Phi \le \Phi_{eng} \\ \dfrac{\Phi_{aut} - \Phi}{\Phi_{aut} - \Phi_{eng}} & \Phi_{eng} \le \Phi \le \Phi_{aut} \\ 1 & \Phi \ge \Phi_{aut} \end{cases} \tag{10}$$

where Φ_{eng} and Φ_{aut} represent the threat level at which the controller engages and the level at which it is given full control authority and effectively acts as an autonomous controller.

Using predicted threat (Φ) as calculated in (8) with an appropriate cost function formulation of the form (7) ensures that 1) the threat metric regulating controller intervention is minimized in the path plan (and associated control calculation) and 2) the controller maintains full control authority when constraints are binding. Increasing Φ_{eng} widens the "low threat" band in which the driver's inputs are unaffected by the controller. Increasing the value of Φ_{aut}, on the other hand, delays complete controller intervention until more severe maneuvers are predicted. The friction-limited bounds on the linear region of the tire force curve (1) suggest a natural upper limit of $\Phi_{aut} \le 5$ degrees in order to ensure that by the time the predicted maneuver required to remain within the safe region of the state space reaches this level of severity, the controller has full control authority and can – unless unforeseen constraints dictate otherwise – guide the vehicle to safety.

3 Experimental Setup

Experimental testing was performed at 14 m/s using a test vehicle and three human drivers. Driver and actuator steering inputs were coupled via an Active Front Steer (AFS) system. An inertial and GPS navigation system was used to measure vehicle position, sideslip, yaw angle, and yaw rate while a 1 GHz dSPACE processor ran controller code and interfaced with steering actuators.

Three common scenarios were used to analyze system performance. In each scenario, obstacles, hazards, and driver targets were represented to the driver by cones and lane markings and to the controller by a constrained corridor (with on-board sensing and constraint mapping assumed to have been performed previously by "virtual sensors" and high-level planners respectively). Only results from multiple-hazard-avoidance tests are shown below. In these tests (illustrated in Fig. 3), both lanes of travel were blocked at different locations, forcing the vehicle to change lanes to avoid the first hazard, then change lanes again to avoid the second.

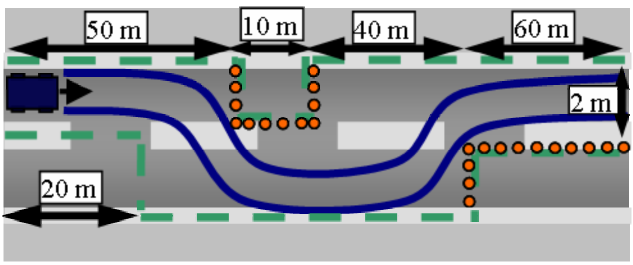

Fig. 3 Multiple hazard avoidance test setup showing hazard cone placement (circles) and lane boundaries (dashed).

Table 2 Controller parameters.

Symbol	Description	Value [units]
p	Prediction horizon	{35, 40}
n	Control horizon	{18, 20}
$R_y^{(\ )}$	Weight on front wheel slip	0.2657
R_u	Weight on steering input	0.01
$R_{\ u}$	Weight on steering input rate (per t)	0.01
$u_{min/max}$	Steering input constraints	± 10 [deg]
$\ u_{min/max}$	steering input rate (per t) constraints	± .75 [deg] (15 deg/s)
$y^{y}_{\ min/max}$	Lateral position constraints	Scenario-dependent
	Weight on constraint violation	1 x 10^5
$[\ _{eng}\ \ _{aut}]$	Thresholds for controller intervention	{[0 3], [1 3]} deg
V	Variable constraint relaxation on vehicle position	[1.25, ···, 1.25, 0.01]

Two types of human driver inputs were tested. Drowsy, inattentive, or otherwise impaired drivers were represented by a constant driver steer input of zero degrees. In these tests, the unassisted driver's path formed a straight line directly through the obstacle(s). To represent active driver steer inputs, the drivers were asked to steer either around or into obstacles.

Controller parameters are described and quantified in Table 2.

4 Experimental Results

The semi-autonomous framework proved capable of keeping the vehicle within the navigable corridor for each of the maneuvers, using various system/controller configurations, and with three different human drivers. Results from multiple hazard avoidance experiments are shown below.

Fig. 4 compares a semi-autonomous multi-hazard-avoidance maneuver to an autonomous maneuver (K=1).

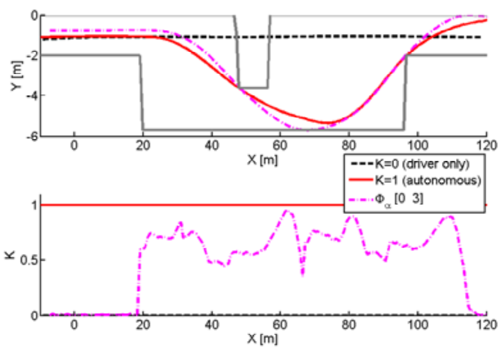

Fig. 4 Multiple hazard avoidance tests showing the similarity between semi-autonomous (dash-dot) and autonomous (solid) vehicle trajectories.

Notice that the semi-autonomous controller delayed intervention until the driver's inputs put the vehicle at risk of leaving the navigable road surface. When the framework did intervene, it allocated enough control authority to the controller to avert corridor departure or loss of control. Also notice that even with average controller intervention K_{ave}=0.44, the vehicle trajectory obtained using the semi-autonomous controller very closely resembles the "best case" trajectory taken by the autonomous controller. This results from the selective nature of the semi-autonomous system – it intervenes only when necessary, then relinquishes control to the driver once threat to the vehicle has been reduced.

Fig. 5 shows experiments in which the driver was instructed to swerve at the last minute to avoid hazards.

Notice that intervention by the semi-autonomous controller slightly preceded an otherwise-late driver reaction. The combined effect of both inputs was then sufficient to avoid both road hazards.

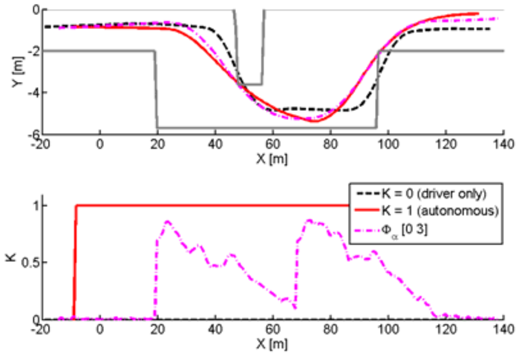

Fig. 5 Multiple hazard avoidance tests showing the vehicle trajectory with an unassisted driver input (dashed) and autonomous controller (solid), and semi-autonomous controller (dash-dot).

Finally, in each of the above experimental results, this shared-adaptive controller behaves as a stable closed-loop system. While this was also true of all of the other simulated and experimental results conducted to date, no rigorous stability proof is presented in this paper.

5 Conclusions

This paper presented an optimal-control-based framework that performs trajectory planning, threat assessment, and semi-autonomous control of passenger vehicles in hazard avoidance. This framework has been proven experimentally capable of satisfying position, input, and dynamic vehicle constraints using multiple threat metrics and intervention laws. Additionally, this framework has been shown to provide significant autonomy to a human driver, intervening only as necessary to keep the vehicle under control and within the navigable roadway corridor. Experimental results have also shown this control framework to be stable even in the presence of system-inherent time delays, though a rigorous stability proof is a topic of current investigation.

Finally, while human factors have not been studied in depth here, it is expected that with additional investigation, a best-case, or average driver-preferred intervention law may be described and intervention settings tuned accordingly. Further work is needed before this research is road-ready.

Acknowledgments. The authors would like to thank Eric Tseng, Matt Rupp, Mitch McConnell, Len Johnson, Steve Hermann, Kyle Carey, Reid Steiger, Tim Zwicky, Roger Trombley, Chris Wallis, and Jeff Rupp, all of Ford Motor Co. for their assistance in conducting the abovementioned experiments.

References

1. National Highway Traffic Safety Administration (NHTSA), 2007 Traffic Safety Annual Assessment - Highlights, NHTSA National Center for Statistics and Analysis (2008)
2. Weilkes, M., Burkle, L., Rentschler, T., Scherl, M.: Future vehicle guidance assistance - combined longitudinal and lateral control. Automatisierungstechnik 42, 4–10 (2005)
3. Leonard, J., How, J., Teller, S., et al.: A perception-driven autonomous urban vehicle. Journal of Field Robotics 25, 727–774 (2008)
4. Jansson, J.: Collision avoidance theory with application to automotive collision mitigation. Doctoral Dissertation, Linkoping University (2005)
5. Pohl, J., Birk, W., Westervall, L.: A driver-distraction-based lane-keeping assistance system. Proceedings of the Institution of Mechanical Engineers. Part I: Journal of Systems and Control Engineering 221, 541–552 (2007)
6. Mobus, R., Zomotor, Z.: Constrained optimal control for lateral vehicle guidance. In: Proceedings of the 2005 IEEE Intelligent Vehicles Symposium, June 6-8, pp. 429–434. IEEE, Piscataway (2005)
7. Netto, M., Blosseville, J., Lusetti, B., Mammar, S.: A new robust control system with optimized use of the lane detection data for vehicle full lateral control under strong curvatures. In: ITSC 2006: 2006 IEEE Intelligent Transportation Systems Conference, Piscataway, NJ, September 17-20, pp. 1382–1387. Institute of Electrical and Electronics Engineers Inc., United States (2006)
8. Vaidyanathan, R., Hocaoglu, C., Prince, T.S., Quinn, R.D.: Evolutionary path planning for autonomous air vehicles using multiresolution path representation. In: 2001 IEEE/RSJ International Conference on Intelligent Robots and Systems, October 29-November 3, pp. 69–76. Institute of Electrical and Electronics Engineers Inc. (2001)
9. Rossetter, E.J., Christian Gardes, J.: Lyapunov based performance guarantees for the potential field lane-keeping assistance system. Journal of Dynamic Systems, Measurement and Control, Transactions of the ASME 128, 510–522 (2006)
10. Falcone, P., Tufo, M., Borrelli, F., Asgari, J., Tseng, H.: A linear time varying model predictive control approach to the integrated vehicle dynamics control problem in autonomous systems. In: 46th IEEE Conference on Decision and Control 2007, CDC, Piscataway, NJ, December 12-14, 2007, pp. 2980–2985. Institute of Electrical and Electronics Engineers Inc, United States (2008)
11. Engelman, G., Ekmark, J., Tellis, L., Tarabishy, M.N., Joh, G.M., Trombley, R.A., Williams, R.E.: Threat level identification and quantifying system, U.S. Patent US 7034668 B2, April 25 (2006)
12. Yu, H., Spenko, M., Dubowsky, S.: An adaptive shared control system for an intelligent mobility aid for the elderly. Autonomous Robots 15, 53–66 (2003)
13. McBride, J.R., Ivan, J.C., Rhode, D.S., Rupp, J.D., Rupp, M.Y., Higgins, J.D., Turner, D.D., Eustice, R.M.: A perspective on emerging automotive safety applications, derived from lessons learned through participation in the DARPA Grand Challenges. Journal of Field Robotics 25, 808–840 (2008)
14. Zomotor, Z., Franke, U.: Sensor fusion for improved vision based lane recognition and object tracking with range-finders. In: Proceedings of the 1997 IEEE Conference on Intelligent Transportation Systems, ITSC, November 9-12, pp. 595–600. IEEE, Piscataway (1997)
15. Garcia, C., Prett, D., Morari, M.: Model predictive control: theory and practice-a survey. Automatica 25, 335–348 (1989)
16. Falcone, P., Borrelli, F., Asgari, J., Tseng, H.E., Hrovat, D.: Predictive active steering control for autonomous vehicle systems. IEEE Transactions on Control Systems Technology 15, 566–580 (2007)