

# Shared Control Up To The Limits of Vehicle Handling

John Talbot<sup>1,2</sup>, Matthew Brown<sup>3</sup>, and J. Christian Gerdes<sup>1</sup>

**Abstract**—Despite significant advances in vehicle safety, roadway accidents remain a substantial danger. To bring about safer vehicles, researchers and manufacturers continue to develop new systems to assist drivers in dangerous situations. Current approaches often implement several independent systems, each of which can assist in only specific situations, leaving open critical paths for accidents to occur. We propose a general approach to driver assistance based on nonlinear model predictive control (NMPC). This system can intervene in both lateral and longitudinal commands to keep the vehicle safely within the boundaries of the road, but allows the driver freedom to maneuver the vehicle as they wish when they act safely. This work builds on previous MPC approaches by incorporating a notion of safe operating speed. Properly modulating speed ensures the vehicle's limits are never exceeded, and not reached unless necessary. Experimental results on a full-size steer-, brake-, and throttle-by-wire vehicle, validate the performance of this system. These experiments show that the controller effectively matches the driver's commands when possible but will deviate from those commands when necessary for safe operation.

**Index Terms**—Driver Assistance, Model Predictive Control, Shared Control, Vehicle Safety

## I. INTRODUCTION

EVERY year, thousands of Americans lose their lives to motor vehicle crashes [1]. Of the many factors that can contribute to a crash, a study conducted by NHTSA showed that human error accounts for 94% of the critical reasons for pre-crash events [2].

To address driver shortcomings, manufacturers develop and deploy many safety systems on consumer vehicles to reduce traffic accidents. Anti-lock Braking Systems (ABS) were introduced in passenger vehicles throughout the 1980's and early 1990's. The system modulates brake pressures to prevent wheel lockup while maintaining the maximum possible brake force. Preventing front wheel lockup also aids the driver by maintaining steering authority of the vehicle. Despite success on the test track, Kahane and Dang found through a long term study that the system had a negligible reduction in fatal roadway crash risk [3]. While not pointing definitively to a reason for the inefficacy of ABS, the authors suggest that while ABS reduced fatalities on the road, it increased the risk of the vehicle leaving the roadway and a fatality occurring.

Though not directly effective at reducing fatalities, ABS technology became a fundamental component of Electronic



Fig. 1. X1, the drive-by-wire, student build research vehicle. The drive-by-wire capabilities are well suited to driver assistance system testing, allowing the system to override driver commands when necessary for safety.

Stability Control (ESC) which has successfully improved vehicle safety. Van Zanten *et al.* proposed an ESC system that modulates brake pressures not just to prevent wheel lockup, but to ensure the vehicle maintains stable operation in challenging maneuvers [4]. The system infers the driver's intended trajectory from steering wheel and pedal measurements. It then stabilizes the vehicle around this intended trajectory by modulating brake pressure to each wheel as needed to induce a supplementary yaw moment. Adding a yaw moment through braking also improves the steering authority of the driver in an understeering condition. By including a concept of driver intent, ESC performed much more effectively than previous ABS concepts. Webb showed that ESC successfully prevented over 7000 roadway fatalities in the United States between 2011 and 2015 [5].

Modern advanced driver assistance systems (ADASs) significantly expand the ability of a vehicle to keep a driver safe by including a notion of the vehicle's environment. Collision imminent braking systems apply brake pressure without input from the driver to mitigate the effects of a front or rear-end collision [6]. Lane keeping assist systems helps drivers avoid drifting into an adjacent lane [7]. Taking the vehicle's environment into account marks a great improvement in safety; however, each system utilizes only a specific concept of the environment. This limited scope of the environment allows the system to intervene only in specific situations, leaving open the possibility that no system would have the authority or knowledge to intervene appropriately in a complex emergency scenario.

Ideally, an assistance system could reason more generally about the safety of the vehicle alongside the wishes of the driver. It is important that the vehicle always remains in a safe state, so the system must be able to make significant corrective actions when in an emergency. It must also allow the driver to operate the vehicle unimpeded so long as their actions result in safe operation. These objectives come into tension when the

Toyota Research Institute provided funds to support this work.

<sup>1</sup>Department of Mechanical Engineering, Stanford University, Stanford, CA 94305 USA

<sup>2</sup>Corresponding author: jat0021@stanford.edu

<sup>3</sup>Woven Planet North America, 900 Arastradero Road, Palo Alto, CA 94304 USA

driver acts incorrectly or negligently, so the system must be able to trade-off between them, overriding the driver's actions when necessary.

Many approaches have been proposed for such a system, and one technique, Model Predictive Control (MPC), is popular among them. Borrelli *et al.* describe MPC as an optimal control technique where an open-loop, finite time, constrained optimization problem is repeatedly solved given the system's sampled states [8]. The first inputs of the resulting optimal trajectory are executed, and after some finite time, the system's states are resampled and the process repeats. MPC typically refers to optimization problems that constrain the problem's decision variables using only linear combinations of those decision variables. See Rawlings' tutorial of MPC for an accessible reference on the topic [9].

MPC offers a convenient way to encode trade-offs between competing objectives like driver matching and safe operation. Beal and Gerdes developed a linear convex MPC to enforce vehicle stability while directly incorporating driver steering commands into the optimization [10]. It effectively stabilized the vehicle, but no environmental information was included. Erlien and Gerdes extended Beal's approach by adding a safe driving corridor constraint to the problem [11]. This constraint enforced that the vehicle's trajectory could not go off the roadway or through obstacles. Guo *et al.* went further, dynamically varying the incentive to match the driver's steering input based on a metric of hazard [12]. Anderson *et al.* developed a technique to blend the driver's commands and the output of a MPC path tracking controller based on a heuristic metric of vehicle threat [13].

Importantly, these architectures only control steering angle and have no way of influencing vehicle speed. While steering control alone works well in many circumstances, effective speed control becomes increasingly important as the vehicle reaches its handling limits. Inagaki, et al., showed that increasing speed through a maneuver decreases the size of the stable convergence region of the vehicle's yaw-rate and lateral velocity states [14]. For this reason, production ESC systems limit a driver's ability to accelerate when active and often provide braking specifically to reduce vehicle speed in addition to providing a yaw moment [4].

Because MPC encodes a notion of time through its prediction horizon, it naturally allows for temporal trade-offs such as reducing speed now to avoid saturating cornering forces later. This behavior offers a great advantage in automotive applications; however, the nonlinear interdependence of longitudinal and lateral dynamics creates an obstacle to using a classic linear MPC. Nonlinear Model Predictive Control (NMPC) shares the fundamental idea of MPC, but generalizes the underlying optimization problem to accept nonlinear constraints and objectives. It can therefore incorporate coupled lateral and longitudinal vehicle dynamics directly. Brown and Gerdes showed that NMPC could be used to navigate a challenging double lane change to avoid two obstacles leaving only a small navigable path. [15]. In that work, coordinating longitudinal and lateral forces throughout the maneuver was critical to maintaining vehicle stability while avoiding the obstacles as the tires on each axle were operating at their limits.




Fig. 2. Single-Track vehicle model

The fullest treatment of applying NMPC for driver assistance comes from Schwarting *et al.*, who developed a framework directly incorporating a driver's steering and acceleration commands, while maintaining a safe trajectory [16]. They proposed a method of dynamically weighting the two main objectives of the system, matching driver commands and tracking a safe reference path, over the NMPC prediction horizon which we adopt in this work. In their formulation, the authors use a maximum speed objective through the NMPC horizon as an approximate progress objective. This contrasts with ESC and other proposed safety systems which look to decrease vehicle velocity in emergency scenarios. While showing promising results in simulation, pushing vehicles nearer their limits can reduce the robustness of the system in some situations and should be avoided when possible.

We propose a NMPC approach to driver assistance that directly incorporates a driver's steering and longitudinal force commands while ensuring the vehicle maintains a safe trajectory through its environment. We contribute a novel progress objective matching a safe reference speed, allowing the system to operate up to the limits of handling when necessary but giving it no incentive to do so normally. We contribute a method to mitigate delays between the driver's commands and the vehicle's actions using their command rates. We contribute an improved objective function for driver matching improving upon similar approaches by eliminating small unnecessary interventions while allowing significant interventions when necessary.

This paper is structured as follows. Section II details the dynamic model used to describe the vehicle's behavior. Section III explains the setup of the NMPC problem, in particular a novel formulation for the driver matching and path progress objectives. Section IV provides two sets of simulation results verifying the selected objective parameters and the systems response to different driving behaviors. Experimental results in Section V demonstrate the effectiveness of this controller on a full size vehicle on a closed test course.

## II. MODELING

To adequately capture vehicle behavior up to the limits of handling, we employ the planar single-track vehicle model shown in Fig. 2. Three velocity states describe the vehicle's motion with respect to the vehicle body frame: the longitudinal velocity  $u_x$ , the lateral velocity  $u_y$ , and the yaw rate  $r$ . Three

position states locate the vehicle in the curvilinear frame of the reference path: progress along the path  $s$ , lateral offset from the path  $e$ , and heading offset  $\Delta\Psi$ .  $\kappa$  represents the instantaneous curvature of the path. The steering angle  $\delta$  and the longitudinal forces at each axle,  $F_{xf}$  and  $F_{xr}$  act as inputs to the model.  $F_{yf}$  and  $F_{yr}$  are lateral tire forces as described in Section II-B. The model is parameterized using the vehicle's mass  $m$ , the distances from the vehicle's center of mass to its front and rear axles  $a$  and  $b$ , respectively, and its yaw axis moment of inertia  $I_{zz}$ . The following six equations describe the behavior of the vehicle:

$$\dot{u}_x = \frac{1}{m} (F_{xf} \cos \delta - F_{yf} \sin \delta + F_{xr}) + ru_y \quad (1)$$

$$\dot{u}_y = \frac{1}{m} (F_{yf} \cos \delta + F_{xf} \sin \delta + F_{yr}) - ru_x \quad (2)$$

$$\dot{r} = \frac{1}{I_{zz}} (aF_{yf} \cos \delta + aF_{xf} \sin \delta - bF_{yr}) \quad (3)$$

$$\dot{s} = \frac{1}{1 - \kappa e} (u_x \cos \Delta\Psi - u_y \sin \Delta\Psi) \quad (4)$$

$$\dot{e} = u_x \sin \Delta\Psi + u_y \cos \Delta\Psi \quad (5)$$

$$\Delta\dot{\Psi} = r - \kappa\dot{s} \quad (6)$$

#### A. Longitudinal Weight Transfer

During heavy braking and acceleration events, the weight transferred between the front and rear axle significantly changes the normal load at each axle. The change in normal load then affects the force producing capability of that axle's tires as described in Section II-B. A static moment balance sufficiently captures this phenomenon while maintaining model simplicity. The instantaneous normal load on each axle is given by

$$F_{zf} = \frac{1}{L} (mgb - h_{cm}F_x) \quad (7)$$

$$F_{zr} = \frac{1}{L} (mga + h_{cm}F_x) \quad (8)$$

where  $g$  is acceleration due to gravity,  $h_{cm}$  is the height of the center of mass measured from the ground, and  $L = a + b$  is the total wheelbase of the vehicle.

#### B. Tire Model

The lateral tire forces,  $F_{yf}$  and  $F_{yr}$ , arise from the slip angles,  $\alpha_f$  and  $\alpha_r$ , respectively. The angle from a tire's longitudinal heading vector to its velocity vector defines its slip angle. The front and rear slip angles are given by (9) and (10), respectively.

$$\alpha_f = \arctan\left(\frac{u_y + ar}{u_x}\right) - \delta \quad (9)$$

$$\alpha_r = \arctan\left(\frac{u_y - br}{u_x}\right) \quad (10)$$

The lateral tire forces relate to the tire slip angles using a modified Fiala brush tire model derived from models proposed by Pacejka [17]. This brush tire model describes the lateral

force at a given slip angle using two parameters: the road-tire friction coefficient,  $\mu$ , and the cornering stiffness of the tire  $C_\alpha$ . The sliding slip angle,  $\alpha^{sl}$  characterizes the transition from the unsaturated to the saturated operating region of the tire model.  $F_y^{max}$  is the maximum lateral force the tire can produce.

$$F_y = \begin{cases} -C_\alpha \tan \alpha \dots \\ \quad + \frac{C_\alpha^2}{3F_y^{max}} |\tan \alpha| \tan \alpha \dots \\ \quad - \frac{C_\alpha^3}{27(F_y^{max})^2} \tan^3 \alpha, & \text{if } |\alpha| < \alpha^{sl} \\ -F_y^{max} \operatorname{sgn} \alpha, & \text{otherwise} \end{cases} \quad (11a)$$

$$\alpha^{sl} = \arctan\left(\frac{3\mu F_z}{C_\alpha}\right) \quad (12)$$

When operating at the limits, the maximum lateral and longitudinal forces of the tire are coupled, limited by the total force generation capacity of the tire. We describe this coupling by derating the maximum lateral force capability of the tire by the commanded longitudinal force as shown by Hindiyeh and Gerdes [18]

$$F_y^{max} = \sqrt{(\mu F_z)^2 - F_x^2} \quad (13)$$

### III. NMPC PROBLEM FORMULATION

We formulate the proposed ADAS as a NMPC with  $N = 60$  stages of  $\Delta t = 50$  milliseconds each for a total horizon length of  $N\Delta t = 3$  seconds. Converting the continuous vehicle model from Section II to a discrete model using trapezoidal integration over the stage duration yields a discrete temporal mapping from current vehicle states and inputs to the vehicle states at the next stage. Two variable changes are implemented between the vehicle model proposed in Section II and the NMPC formulation. First, the front and rear longitudinal force are combined,  $F_x = F_{xf} + F_{xr}$ , to avoid competing front and rear forces. Secondly, the inputs to the vehicle model  $\delta$  and  $F_x$  are treated as states in the NMPC formulation, and their rates of change  $\dot{\delta}$  and  $\dot{F}_x$  are treated as inputs. This choice of inputs allows the simple inclusion of actuator slew rate constraints and penalties to ensure command feasibility and smoothness. The state vector at stage  $k$  of the NMPC horizon is  $z^{(k)} = [u_x^{(k)}, u_y^{(k)}, r^{(k)}, s^{(k)}, e^{(k)}, \Delta\Psi^{(k)}, \delta^{(k)}, F_x^{(k)}]^T$ , and the input vector is  $u^{(k)} = [\dot{\delta}^{(k)}, \dot{F}_x^{(k)}]^T$ .

The complete NMPC formulation is given by (14). As is typical in receding horizon controllers, the initial state vector is constrained to the actual measured state of the vehicle and actuators (14b). Each subsequent state is constrained by the discretized dynamics described in Section II (14c). At each time step the actuator position limits (14d) (14e) and slew rate limits (14f) (14g) are constrained by their physical limits. The total longitudinal force is mapped into a front and rear longitudinal force using brake and drive biases at each axle. The biases,  $\chi_f$  for front axle bias, and  $\chi_r$  for rear axle bias are computed using the smooth approximation presented by Laurence, et al., adapted for a rear wheel drive vehicle [19].

Finally, the magnitude of the front and rear maximum longitudinal tire force are constrained by their respective friction limits (14j) (14k). As the lateral forces are constrained within the derated brush tire model itself, this completely encodes the friction limits of each of the tires.

$$\min_{z, u} J \quad (14a)$$

$$\text{s.t.} \quad z^{(1)} = z_{\text{measured}} \quad (14b)$$

$$z^{(k+1)} = f_{\text{dyn}}(z^{(k)}, u^{(k)}) \quad \forall k = 1, \dots, N-1 \quad (14c)$$

$$\underline{\delta} \leq \delta^{(k)} \leq \bar{\delta} \quad \forall k = 1, \dots, N \quad (14d)$$

$$\underline{F}_x \leq F_x^{(k)} \leq \bar{F}_x \quad \vdots \quad (14e)$$

$$\underline{\dot{\delta}} \leq \dot{\delta}^{(k)} \leq \bar{\dot{\delta}} \quad (14f)$$

$$\underline{\dot{F}}_x \leq \dot{F}_x^{(k)} \leq \bar{\dot{F}}_x \quad (14g)$$

$$F_{xf}^{(k)} = \chi_f^{(k)} F_x^{(k)} \quad (14h)$$

$$F_{xr}^{(k)} = \chi_r^{(k)} F_x^{(k)} \quad (14i)$$

$$|F_{xf}^{(k)}| \leq \mu \cos(\alpha_f) F_{zf}^{(k)} \quad (14j)$$

$$|F_{xr}^{(k)}| \leq \mu \cos(\alpha_r) F_{zr}^{(k)} \quad (14k)$$

The total objective  $J$  combines  $J_{\text{driver}}(k)$ , the objective to match the driver's commands (18),  $J_{\text{path}}(k)$ , the objective to tracking the reference path (20),  $J_{\text{actuator}}(k)$ , the objective to minimize actuator slew rates (21),  $J_{\text{edge}}$ , the objective to stay within the road boundaries (16), and  $J_{\text{term}}$ , the terminal objective of the NMPC problem (22).  $\lambda(k)$  is a dynamic weighting function between matching the driver and tracking a reference path (17). The following sections detail the individual components of the objective function.

$$J = \sum_k \lambda(k) J_{\text{driver}}(k) + \sum_k (1 - \lambda(k)) J_{\text{path}}(k) + \sum_k J_{\text{actuator}}(k) + \sum_k J_{\text{edge}}(k) + J_{\text{term}} \quad (15)$$

### A. Safe Operation

Approximating the vehicle body as a set of circles as presented by Brown and Gerdes [15] yields a computationally inexpensive formulation to find the signed distance from the vehicle to the road boundaries. A set of two circles, one at the front axle and one at the rear approximates the vehicle body in this work. Within the road boundaries, the signed distance is positive and crossing a road boundary yields a negative distance. We define the safe operating region for the vehicle as the region where the distance to both road edges are positive. At each stage,  $k$ , while within the limits of the safe operating region, the vehicle incurs no cost. Violating these boundaries incurs a high quadratic cost given by (16) where  $d_{\text{edge}}^{(k)}$  is the signed distance to the road boundary at stage  $k$ , and  $W_{\text{edge}}$  is the cost on crossing the road boundary.

$$J_{\text{edge}}(k) = \begin{cases} 0 & d_{\text{edge}}^{(k)} > 0 \\ W_{\text{edge}} \left(d_{\text{edge}}^{(k)}\right)^2 & \text{otherwise} \end{cases} \quad (16a)$$

$$(16b)$$

### B. Time Variable Secondary Objective

While the system must always keep the vehicle safe, to be accepted, it must match the driver's intent as closely as possible. Because the driver may continuously vary their driving inputs, predictions of their behavior lose validity as they are made further into the future. To mitigate this Schwarting proposed trading between two different secondary objectives over the prediction horizon through an exponential weighting function [16]. In the near term, predictions of the driver's steering and longitudinal commands are more certain, making matching them directly more reasonable. Later in the horizon as the uncertainty in those predictions grows, tracking a predetermined safe path becomes more appropriate. We employ the same strategy in this work, decreasing the weight on matching the driver's commands exponentially as a function of time through the horizon according to (17) where  $\gamma$  is a tuning parameter:

$$\lambda(k) = \exp(-\gamma k \Delta t) \quad (17)$$

The relative contributions of driver matching and path tracking over the prediction horizon are shown in Fig. 3.

### C. Driver Intent

The system obtains the driver's intended steering commands by measuring their hand wheel angle directly. The driver's intended longitudinal commands cannot be measured directly, so the system infers them through position measurements of the throttle and brake pedals. These pedal positions map linearly onto the force generating capabilities of the motor and brakes, respectively. We assume that the current driver steering rate and pedal position remain constant throughout the NMPC horizon to illustrate this system performs well without a complex model of the driver's behavior. This is discussed in greater detail in Section V-B.

Ideally, small, unnecessary deviations from the driver's commands would be eliminated from an ADAS while still allowing for significant intervention when needed. A 1-norm

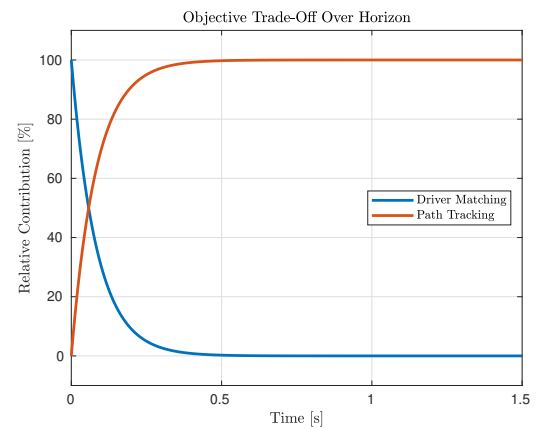


Fig. 3. Relative contribution of objective function terms. Further into the prediction horizon, the projected driver intent becomes less valid. To mitigate this, the objective of the NMPC formulation transitions from driver matching to path tracking.

or absolute value cost function eliminates small deviations while allowing large deviations when needed, but is not smooth which can make it difficult for numeric solvers to find a feasible, optimal solution. We use a log cosh cost function as shown in (18) where  $\delta_{\text{driver}}$  is the driver's steering command,  $F_{\text{driver}}$  is the driver's longitudinal force command,  $W_{\delta_{\text{driver}}}$  is the cost on deviations from the driver's steering commands, and  $W_{F_{\text{driver}}}$  is the cost on deviations from the driver's longitudinal force commands. This function approximates the 1-norm cost function behavior well, while also being twice-differentiable and smooth. A scaling factor  $\eta_{\delta}$  is used to normalize the relatively small value of steering angle in radians and  $\eta_{F_x}$  scales the relatively large values of longitudinal force in newtons to kilonewtons.

$$J_{\text{driver}}(k) = W_{\delta_{\text{driver}}} \log \cosh \left( \eta_{\delta} (\delta^{(k)} - \delta_{\text{driver}}) \right) + W_{F_{\text{driver}}} \log \cosh \left( \eta_{F_x} (F_x^{(k)} - F_{\text{driver}}) \right) \quad (18)$$

#### D. Path Tracking and Safe Progress Objective

As the influence of the driver's current commands diminishes throughout the horizon, the main objective changes to maintaining a safe and continuously feasible trajectory on the roadway. In this region of the NMPC horizon, a progress objective throughout the horizon ensures the horizon continuously recedes, rather than simply stopping the vehicle whenever approaching an environmental boundary. Schwarting's approach relies on a maximum speed objective as a proxy for maximum progress [16]. While such an objective performs adequately in many situations, it can unnecessarily incentivize the system to operate the vehicle near its operating limits, thus reducing the robustness of the system. In contrast, the proposed formulation incentivizes matching a predetermined safe speed profile as a progress objective. In this implementation, the maximum speed a point mass could navigate a given corner with the modeled friction limits and road curvature given by (19) determines the safe reference speed.

$$v_{\text{safe}} = \sqrt{\mu g / \kappa} \quad (19)$$

To accomplish path tracking, the NMPC formulation includes an objective to follow a predetermined reference path. A deviation from the reference lateral position, heading, or speed incurs a quadratic cost given in (20) where  $W_e$  is the cost on lateral path tracking error,  $W_{\Delta\Psi}$  is the cost on heading tracking error, and  $W_{u_x}$  is the cost on speed tracking error.

$$J_{\text{path}}(k) = W_e \left( e^{(k)} \right)^2 + W_{\Delta\Psi} \left( \Delta\Psi^{(k)} \right)^2 + W_{u_x} \left( v_{\text{safe}}^{(k)} - u_x^{(k)} \right)^2 \quad (20)$$

#### E. Actuator Rates

A penalty on actuator slew rates helps minimize sharp changes in steering and acceleration and smooth the final trajectory. The slew rate objective is given in (21) where  $W_{\dot{\delta}}$  is the cost on steering slew rate, and  $W_{\dot{F}_x}$  is the cost on

longitudinal force slew rate. These costs, in addition to the cost on matching driver commands, influence the aggressiveness of an intervention maneuver. Increasing the penalty on large actuator slew rates incentivizes the controller to deviate from the driver's commands sooner when an unsafe condition is predicted to avoid a sudden large sudden deviation. The actuator slew rate costs thus provide a method to tune the invasiveness of the system to fit the skill and comfort level of an individual driver. A less skilled driver or one that desires an earlier but smoother intervention can increase the slew rate costs and achieve that behavior from this system. A more skilled driver or one that desires less intervention can decrease these costs, yielding fewer but more aggressive and abrupt interventions.

$$J_{\text{actuator}}(k) = W_{\dot{\delta}} \left( \eta_{\delta} \dot{\delta}^{(k)} \right)^2 + W_{\dot{F}_x} \left( \eta_{F_x} \dot{F}_x^{(k)} \right)^2 \quad (21)$$

#### F. Terminal State

The NMPC formulation includes a strong incentive to keep the vehicle in a known safe state at the end of the prediction horizon by including a large terminal cost on the path tracking terms at the terminal stage of the controller. This promotes the feasibility of future solves. The terminal objective is given in (22) where  $k = N$  is the final stage of the NMPC prediction horizon,  $W_{e,N}$  is the cost on lateral path tracking error,  $W_{\Delta\Psi,N}$  is the cost on heading tracking error, and  $W_{u_x,N}$  is the cost on speed tracking error, all at the final stage of the NMPC prediction horizon.

$$J_{\text{term}} = W_{e,N} \left( e^{(N)} \right)^2 + W_{\Delta\Psi,N} \left( \Delta\Psi^{(N)} \right)^2 + W_{u_x,N} \left( v_{\text{safe}}^{(N)} - u_x^{(N)} \right)^2 \quad (22)$$

## IV. SIMULATION RESULTS

In the following simulations, we show firstly that our choice of optimization parameters is appropriate for the stated goal. Secondly, we compare progress objective formulations to show the advantages of matching a safe reference speed rather than maximizing speed.

#### A. Validating Optimization Parameters

A capable ADAS must respond appropriately to the full range of possible driver commands while maintaining the safety of the vehicle. In this section, we show that the proposed system effectively responds to different driver steer angles given the parameters listed in Table I. All simulations and experiments presented in this paper make use of these parameters. The following simulations show the system's response to five equally spaced driver steering commands from the left-most to the right-most steering limits. In each simulation, the vehicle travels down a straight road near the right road boundary at 8 m/s. Because the vehicle is nearer the right road boundary, the system will need to intervene significantly when the driver steers right, but very little when they steer

TABLE I  
OPTIMIZATION PARAMETERS

Parameter	Value	Description
$\gamma$	8	decay rate of driver matching
$W_{edge}$	1000	path overrun cost
$W_{\delta_{driver}}$	$1/(\pi/180)$	driver steering angle cost
$W_{F_{driver}}$	$1/0.2$	driver longitudinal force cost
$W_{u_x}$	$1/30^2$	speed tracking cost
$W_e$	$1/16^2$	path tracking error cost
$W_{\Delta\Psi}$	$1/(300\pi/180)^2$	path heading error cost
$W_{u_x,N}$	1	terminal speed tracking cost
$W_{e,N}$	$1/0.5^2$	terminal path tracking error cost
$W_{\Delta\Psi,N}$	$1/(0.5\pi/180)^2$	terminal path heading error cost
$W_{\dot{\delta}}$	$1/(500\pi/180)^2$	steering slew rate cost
$W_{F_{\dot{x}}}$	$1/10^2$	longitudinal force slew rate cost
$\eta_{\delta}$	10	steering angle normalization factor
$\eta_{F_x}$	1/1000	force normalization factor

left. The initial steer angle state is set to match the driver's input, and the driver's steering rate is set to zero.

Figure 4 shows the optimal NMPC horizon of steer angles for each of the driver's commanded steering angles. The driver may freely steer the vehicle toward the left boundary, as there is no danger of the vehicle leaving the roadway in that direction. As the driver steers right, however, the system takes significant action to keep the vehicle on the road. In each case, the system responds appropriately to the driver's commands. The incurred cost of the driver steer angle matching term in (18) provides a metric for the severity of the intervention.

### B. Maximum Speed vs. Safe Speed Matching

Increasing the speed of a vehicle in a challenging maneuver increases the risk of dangerous outcomes including loss of vehicle control. As illustrated by Brown and Gerdes, naively

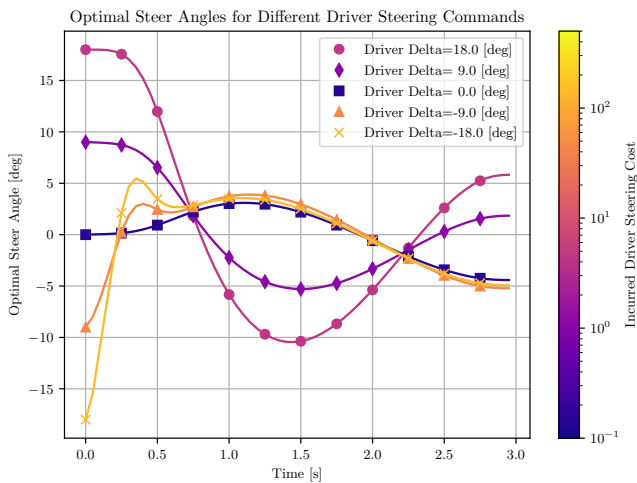


Fig. 4. Optimal NMPC horizons of steering angle for several different driver steering inputs. The driver may steer left relatively unimpeded as even the left-most steer angle would not cause the vehicle to leave the road. The system intervenes increasingly aggressively as the driver steers farther right near the right road boundary. Note the incurred costs are colored based on a logarithmic scale.

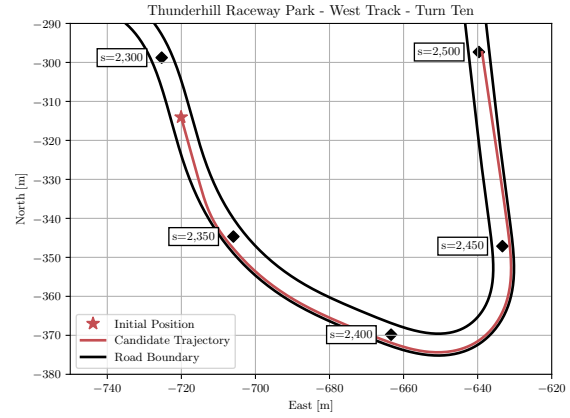


Fig. 5. Turn ten of the west track at Thunderhill Raceway Park. This hairpin corner will be used in the simulations detailed below as well as the experiments in Section V.

operating near the vehicle's limits can reduce robustness to unmodeled disturbances in MPC implementations [20]. Maintaining a buffer between the vehicle's operating states and its modeled limits improves the robustness of the controller. To do this, it is important to reduce the vehicle's speed before entering a maneuver that would challenge its limits.

Potentially in conflict with this objective of lowering speed during challenging maneuvers is the progress objective of the NMPC formulation. As described in Section III, this formulation incentivizes the vehicle to match a safe reference speed rather than maximizing speed as a proxy for progress along the reference path. We show the significance of how to influence the vehicle's speed with two sets of simulations implementing both NMPC formulations through a challenging hair-pin turn maneuver, turn ten of Thunderhill Raceway Park's west track. The simulations assume icy conditions with a nominal road-tire friction value of  $\mu = 0.3$ . Each simulation also demonstrates the performance of two conventional lane keeping assistants as a baseline. The first implements the potential field lane keeping assistant proposed by Rossetter, *et al.* [21], which is a reasonable proxy for a production lane keeping assist. The second adds a simple speed limit monitor to the potential field lane keeping assistant. This monitor augments the driver's force command with a linear feedback term between the vehicle's current speed and the speed limit. The monitor will only become active when the vehicle's speed exceeds the speed limit, and then may only use up to twenty percent of the vehicle's maximum braking capability to reduce its speed. The speed limit for this controller is identical to the safe reference speed used in the NMPC formulation presented in this work. These two systems differ fundamentally from the NMPC formulations as they actively nudge the driver toward the center of the lane.

The vehicle begins each simulation aligned with the road at its lateral center near the entry of the corner. The initial position of the vehicle and one candidate trajectory is show in Fig. 5. The safe reference speed of this corner is determined as 7 meters per second using (19). The vehicle's initial speed is 9.2 meters per second (5 miles per hour greater than the




Fig. 6. Comparison of different driver assistance formulations maneuvering through the hairpin corner shown in Fig. 5. Both NMPC formulations and the lane keeping assistant implementing a speed limit perform well with perfect knowledge of road-tire friction. The classic lane keeping assistant cannot maintain a safe trajectory through the corner. The lane keeping assistant with speed limit monitoring is much more invasive, restricting the driver's speed more than the other systems.

safe speed), and the driver commands zero steer angle and 250 newtons of longitudinal force throughout each simulations representing a very light acceleration. In the maximum speed objective formulation, the stage-wise cost on tracking the safe reference speed as shown in (20) becomes a maximum speed objective. Additionally, the terminal cost in (22) on matching the safe reference speed becomes a maximum speed objective. All other parameters remain equal for both NMPC based systems.

In the first set of simulations, each system has perfect knowledge of the friction between the vehicle's tires and the roadway. As shown in Fig. 6, the difference in behaviors of two NMPC formulations with different progress objectives is small. The classic lane keeping assistant that cannot influence speed is unable to keep the vehicle safely on the roadway. The lane keeping assistant with speed limit monitoring does successfully keep the vehicle near the center of the lane, but as shown in Fig. 6, it enforces a very conservative speed which can be frustrating to a driver.

In real-world operation, there are often slight differences in the modeled and true behavior. One example, described by Martin *et al.*, is winter conditions where friction between vehicle tires and the road varies with several conditions [22]. To highlight the sensitivity to unmodeled disturbances, we introduce a relatively small but typical mismatch in the friction coefficient used in the NMPC formulation and in the simulator. The modeled friction coefficient each of the systems at both tires is  $\mu = 0.3$ ; however, the friction actually available to the front tire is  $\mu_{actual} = 0.28$ .

Figure 7 illustrates the improved robustness of the safe reference speed objective over a maximum progress objective. Like the previous simulation results, the conventional lane keeping assistant cannot keep the vehicle safely on the road

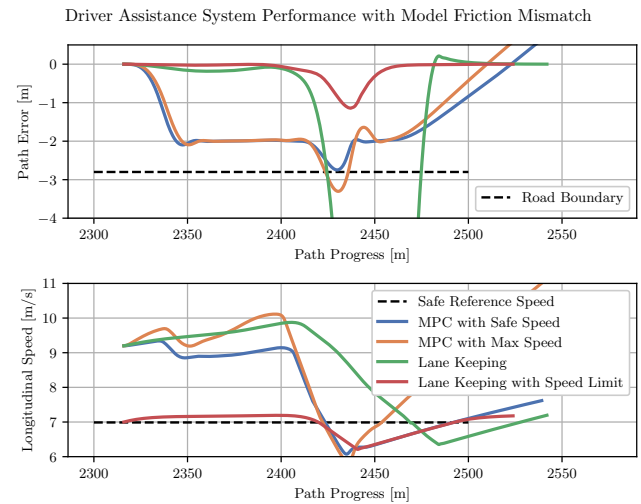


Fig. 7. Comparison of different driver assistance formulations maneuvering through the hairpin corner shown in Fig. 5 with a mismatch in modeled and experienced friction. The safe reference speed progress objective is more resilient to this model mismatch and keeps the vehicle safely on the road, while the maximum speed progress objective exits the roadway. The classic lane keeping assistant again cannot maintain a safe trajectory through the corner, and the lane keeping assistant with speed limit monitoring does keep the vehicle safely on the road, but is much more restrictive than the other systems.

because it cannot influence the vehicle's speed. The lane keeping assistant with speed limit monitoring again does well to keep the vehicle safe, but it is much more invasive in restricting the driver's desired speed. The simulated vehicle implementing the NMPC formulation with the maximum progress objective exited the roadway, while the simulated vehicle with the safe reference speed objective remained safely on the road. This improvement in robustness comes without introducing conservatism because the safe reference speed is based on the modeled friction value, so the NMPC will allow the driver to take the vehicle up to the modeled limits. The proposed controller has similar performance to the maximum speed objective controller when the vehicle is operating near the safe reference speed with no model mismatch; however, under even a small amount of friction variation, the proposed controller has much better performance, marking one novel contribution of this work.

## V. EXPERIMENTAL RESULTS

In this section, real-world experiments performed on a full-sized test vehicle will show firstly that the vehicle responds naturally to driver commands other than a simulated constant input. Secondly, the experiments will validate our choice of dynamic model, and the system architecture performs adequately amid real-world conditions on a test track. Keeping an instrumented vehicle safe up to its handling limits adds confidence that this system can perform its safety tasks in a dynamic and challenging environment.

In these experiments, information about the environment, namely the position of road boundaries and the safe reference speed of the track are determined a priori and embedded into the NMPC formulation. For simplicity, in these experiments,




Fig. 8. Example of delay between driver commands and NMPC output. In this example, the driver's current steering inputs are held constant over the prediction horizon. The driver's commands are effectively delayed by roughly 100 milliseconds.

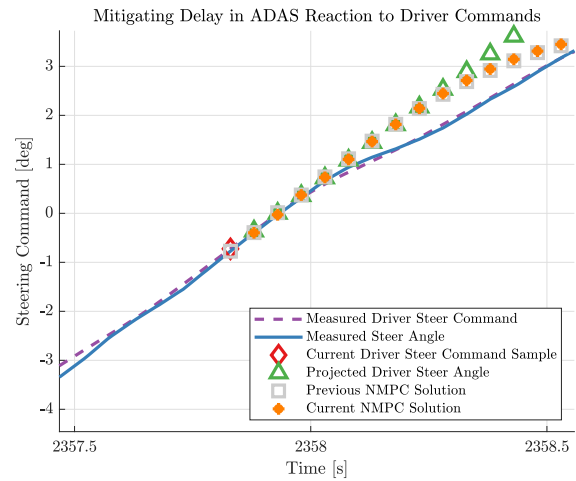


Fig. 9. Delay mitigated using first-order-hold. Assuming the driver's input rates remain constant over the horizon mitigates the delaying effects and improves responsiveness to the driver.

the nominal trajectory is the centerline of the racetrack, and the road bounds were set at 2.8 m to either side of the center line for a total navigable width of 5.6 m. These values enforce a physical safety buffer when experimenting on track but are easily adjustable to values appropriate for the scenario at hand, e.g. the lane width of a road. For production implementations this information could come from surveyed maps of the roadways being used along with their posted speed limits. An online perception stack could augment this information improving reliability and adding adaptability for a changing environment as shown by Wille, *et al.* [23]

### A. Experimental Setup

The driver assistance system in this paper was experimentally validated using X1, the student designed and built electric test platform shown in Fig. 1. X1 relies on steer-, brake-, and throttle-by-wire systems to control the on-board actuators. Sensors measure the position and speed of the steering handwheel as well as the position of the throttle pedal. A dSpace MicroAutoBoxII provides low-level actuator control and communications between all measurement and actuation systems and the PC described below. An OxTS RT4003 dual antenna RTK GNSS unit produces precision estimates of the vehicle's position and velocity states in real time. The nonlinear program (14) was implemented using the open-source automatic differentiation package Casadi [24], and solved using IPOPT, an open source interior point method solver [25]. The generated solver was run on a conventional PC with an Intel i7 CPU running Ubuntu 20.04 and the Robotic Operating System, (ROS) [26].

### B. Improving Responsiveness to Driver Commands

Assumptions on the driver's future behavior become important when employing this framework on an experimental vehicle where it must run in real-time. Average solve times during the experiments performed range between 50 to 100

milliseconds. Solves of this length introduce a delay between when current values are sampled and the solution which results from them. To mitigate this delay for vehicle stability and control, we adopt the replanning strategy proposed by Brown and Gerdes [15]. This strategy reduces the effect of delays caused by the relatively large solve time in two ways. First, it linearly interpolates the previously solved trajectory to supply commands to low-level controllers operating at a faster rate. Second, it uses the previously solved trajectory to project the current state of the vehicle forward in time by 50 milliseconds to become the initial conditions of the next solved trajectory.

Such solve times also introduce a delay between the driver's commands and those executed on the vehicle. This delay leads to confusion and frustration and may ultimately lead to the driver's rejection of the safety system. To mitigate these delays, we assume the driver's steering rate remains constant throughout the horizon, rather than assuming their current

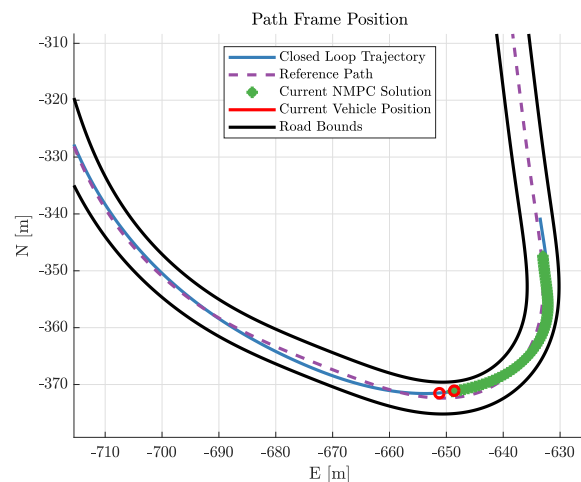


Fig. 10. Vehicle position for skilled driving behavior. The driver remains near the center of the road throughout, so no intervention is necessary.



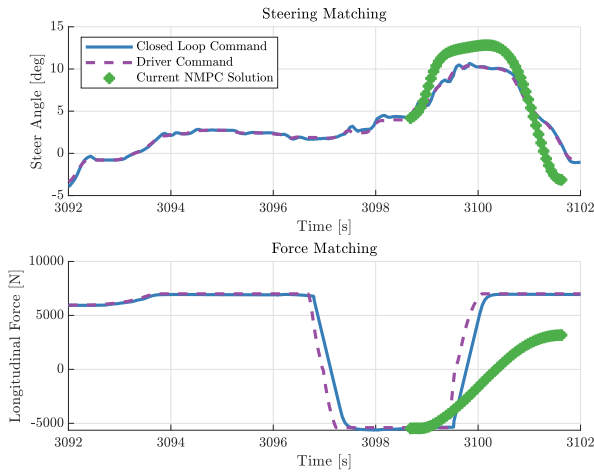


Fig. 11. Driver inputs vs ADAS commands for skilled driving behavior. Notice the commands output by the ADAS system follow the driver's commands closely while the driver operates the vehicle safely.

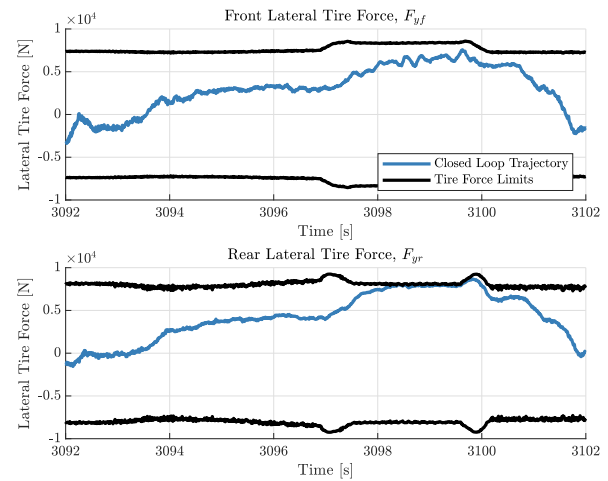


Fig. 12. Lateral tire forces for skilled driving behavior. Even though the driver approaches the modeled tire limits, the system does not intervene, because their commands result in a safe trajectory for the vehicle.

command will remain constant.

Experiments performed on X1, implementing different assumptions on driver behavior, illustrate the effect of these assumptions on system performance. In these experiments, the driver manually inputs an oscillating steering input to the handwheel, and measurements of the handwheel and road wheel positions were captured. The first experiment assumes that the driver will hold their steer angle constant over the full prediction horizon. Figure 8 shows the delay in steering between the driver's handwheel position and the position of the road wheels. This delay between the driver and vehicle results from a few aspects of the NMPC formulation, namely:

- 1) Solve times taking at least 50 milliseconds
- 2) The first stage of the NMPC prediction horizon is constrained to the current measured actuator state
- 3) The driver's commands are assumed constant throughout

These factors effectively delay the driver's commands by approximately 50 to 100 milliseconds. Assuming the driver will hold a constant steering rate over the horizon yields a significant improvement in system responsiveness, making driving feel much more natural. Figure 9 shows the same experiment with the constant steering rate assumption. Importantly, though the predictions made using either assumption quickly lose accuracy compared to the driver's actual future inputs, the driver predictions only have a significant effect on the NMPC solution very early in the prediction horizon. As described in Section III, the objective transitions from driver matching to path tracking after the first few hundred milliseconds. It is also important to note that this assumption does not affect the safety objective of the formulation.

While this assumption tends to work quite well, it can cause a temporary deviation from the driver's intentions in some situations. Specifically, if the driver changes the direction of steering rate at the same time the solver experiences an especially long solve, the constant rate assumption will lead to a greater deviation from the driver's command than a constant position assumption.

Because this delay stems from the significant amount of time it takes to solve the optimization problem, it could also be mitigated with greater computing resources. We provide this solution as a means to implement this algorithm in a situation where computing resources may be constrained, which is often the case in production vehicles.

### C. Driver Matching and Safety

Two experiments demonstrate the proposed system's actions given two different driving behaviors, that of a skilled driver and of a worst-case unskilled driver. In both experiments, the vehicle enters turn ten of the Thunderhill Raceway Park west track at a high speed.

The first experiment illustrates the systems reaction to skilled driving behavior. At the entry of the corner, the driver

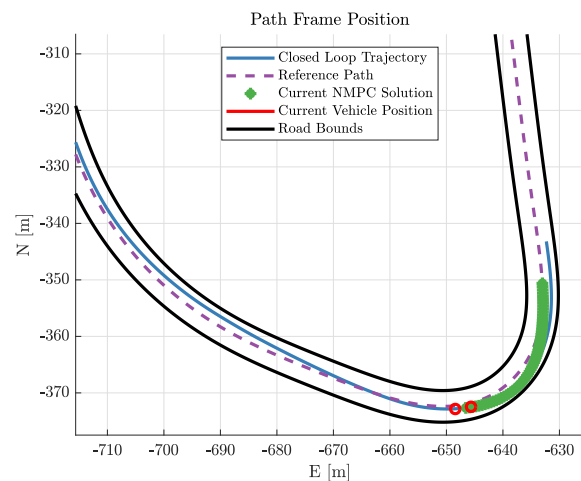


Fig. 13. Vehicle position for unskilled driving behavior. The vehicle nears and eventually touches the right road boundary. The driver does not take action to keep the vehicle on the road, so the ADAS system must make a significant intervention.

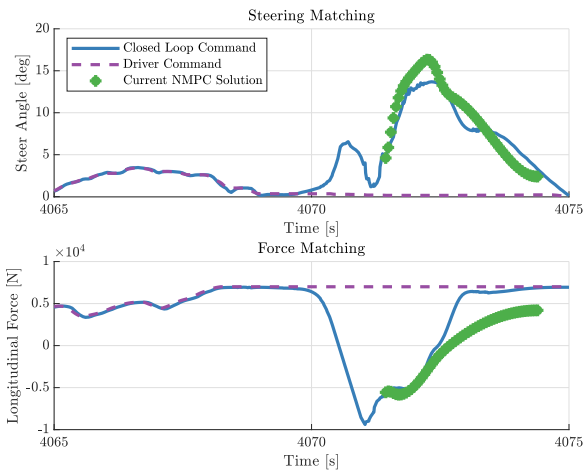


Fig. 14. Driver inputs vs. ADAS commands for unskilled driving behavior. Notice the ADAS must make a significant intervention in both longitudinal and lateral commands to keep the vehicle safely on the road.

releases the throttle, applies moderate brake pressure, and steers to navigate through the corner while stabilizing the vehicle. By braking before entering the corner, the driver can navigate safely without reaching the limits of the vehicle. While the controller does include a cost on tracking the path later in the horizon, the driver is able to control the vehicle as desired while they are driving safely and effectively. Figure 10 shows the driver moving freely within the road boundaries.

With the driver acting appropriately, both maintaining operation within the vehicle's limits and the road boundaries, the assistance system does not need to intervene. Indeed Fig. 11 shows that the driver's commands are effectively passed through the assistance system with no modification. Figure 12 shows the estimated limits of the tires and utilized tire forces according to the model in Section II. Even though the

driver operates the vehicle near the modeled friction limits, the vehicle remains safe and does not approach the road boundaries.

The next experiment illustrates the system's reaction to unskilled driving behavior in the same scenario. This provides an intuitive worst-case driver to benchmark the controller. Approaching the corner, the driver maintains full throttle and applies no steering to navigate through the corner. The vehicle would leave the roadway at a high speed if no intervention occurred, likely leading to a serious accident. In this case, the assistance system deviates significantly from the driver's commands as shown in Fig. 14. The assistance system both applies brake pressure to slow the vehicle down and turns left to compensate for the driver's lack of steering. Figure 13 shows that the assistance system succeeds in keeping the vehicle within the safe operating boundaries of the roadway.

Figure 15 shows the tire force utilization through the worst-case driver test. The controller successfully drives the vehicle at the limits of handling as required to safely navigate the turn at speed. This behavior represents the most important feature of this system: while it does not push the vehicle up to its limits during normal operation, the system effectively operates at the limits when necessary to keep the vehicle safe.

## VI. CONCLUSION

In this paper we have presented a novel approach to ADAS using NMPC. This controller tries to directly match driver commands until those commands would result in an unsafe situation for the vehicle. To do this effectively, the controller modulates the lateral and longitudinal inputs of the vehicle to ensure the tire limits are not exceeded and indeed not reached unless necessary. Simulations show the importance of choosing an appropriate progress objective that does not incentivize the vehicle to operate near its limits when unnecessary, as this reduces the robustness of the system. Two experiments on a test vehicle demonstrate the effectiveness of this system with two distinct driving behaviors: one of an experienced driver who can successfully navigate the vehicle on their own, and one of an unskilled driver who would run the vehicle off the road without intervention.

## ACKNOWLEDGMENTS

Toyota Research Institute provided funds to support this work. The authors also thank the students, staff, and alumni of the Dynamic Design Lab for their valuable discussions and assistance performing experiments.

## REFERENCES

- [1] National Center for Statistics and Analysis, "Overview of motor vehicle crashes in 2019," National Highway Traffic Safety Administration, Tech. Rep. DOT HS 813 060, December 2020.
- [2] S. Singh, "Critical reasons for crashes investigated in the National Motor Vehicle Crash Causation Survey," National Highway Traffic Safety Administration, Tech. Rep. DOT HS 812 115, February 2015.
- [3] C. J. Kahane and J. N. Dang, "The Long-Term Effect of ABS in Passenger Cars and LTVs," National Highway Traffic Safety Administration, Tech. Rep. DOT HS 811 182, August 2009.
- [4] A. T. Van Zanten, R. Erhardt, and G. Pfaff, "VDC, the vehicle dynamics control system of bosch," *SAE Technical Papers*, vol. 104, pp. 1419–1436, 1995.

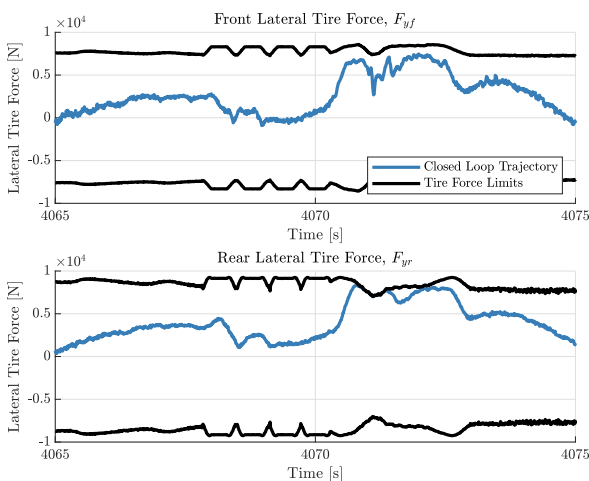


Fig. 15. Lateral tire forces for unskilled driving behavior. The system must operate at the modeled limits of tire friction for a short time to keep the vehicle on the road, but does not remain at the tire limits when it is not necessary to.

- [5] C N Webb, "Estimating Lives Saved by Electronic Stability Control, 2011–2015," National Highway Traffic Safety Administration, Tech. Rep. DOT HS 812 391, March 2015.
- [6] A. Doi, T. Butsuen, T. Niibe, T. Takagi, Y. Yamamoto, and H. Seni, "Development of a rear-end collision avoidance system with automatic brake control," *Jsaee Review*, vol. 15, no. 4, pp. 335–340, 1994.
- [7] S. Ishida and J. E. Gayko, "Development, evaluation and introduction of a lane keeping assistance system," in *IEEE Intelligent Vehicles Symposium, 2004*. IEEE, 2004, pp. 943–944.
- [8] F. Borrelli, A. Bemporad, and M. Morari, *Predictive control for linear and hybrid systems*. Cambridge University Press, 2017.
- [9] J. Rawlings, "Tutorial overview of model predictive control," *IEEE Control Systems Magazine*, vol. 20, no. 3, pp. 38–52, Jun. 2000, conference Name: IEEE Control Systems Magazine.
- [10] C. E. Beal and J. C. Gerdes, "Model Predictive Control for Vehicle Stabilization at the Limits of Handling," *IEEE Transactions on Control Systems Technology*, vol. 21, no. 4, pp. 1258–1269, July 2013.
- [11] S. M. Erlien, S. Fujita, and J. C. Gerdes, "Safe driving envelopes for shared control of ground vehicles," *IFAC Proceedings Volumes*, vol. 46, no. 21, pp. 831–836, 2013.
- [12] H. Guo, L. Song, J. Liu, F. Y. Wang, D. Cao, H. Chen, C. Lv, and P. C. K. Luk, "Hazard-evaluation-oriented moving horizon parallel steering control for driver-automation collaboration during automated driving," *IEEE/CAA Journal of Automatica Sinica*, vol. 5, no. 6, pp. 1062–1073, November 2018.
- [13] S. J. Anderson, J. M. Walker, and K. Iagnemma, "Experimental performance analysis of a homotopy-based shared autonomy framework," *IEEE Transactions on Human-Machine Systems*, vol. 44, no. 2, pp. 190–199, Apr. 2014.
- [14] S. Inagaki, I. Kshiro, and M. Yamamoto, "Analysis on Vehicle Stability in Critical Cornering Using Phase-Plane Method," *International Symposium on Advanced Vehicle Control*, p. 216, 1994.
- [15] M. Brown and J. C. Gerdes, "Coordinating Tire Forces to Avoid Obstacles Using Nonlinear Model Predictive Control," *IEEE Transactions on Intelligent Vehicles*, vol. 5, no. 1, pp. 21–31, 2020.
- [16] W. Schwarting, J. Alonso-Mora, L. Paull, S. Karaman, and D. Rus, "Safe Nonlinear Trajectory Generation for Parallel Autonomy With a Dynamic Vehicle Model," *IEEE Transactions on Intelligent Transportation Systems*, vol. 19, no. 9, pp. 2994–3008, September 2018.
- [17] H. B. Pacejka, *Tire and Vehicle Dynamics*. Burlington, MA: Elsevier, 2006.
- [18] R. Y. Hindiyeh and J. C. Gerdes, "Design of a dynamic surface controller for vehicle sideslip angle during autonomous drifting," *IFAC Proceedings Volumes*, vol. 43, no. 7, pp. 560–565, 2010.
- [19] V. A. Laurence and J. C. Gerdes, "Long-Horizon Vehicle Motion Planning and Control Through Serially Cascaded Model Complexity," *IEEE Transactions on Control Systems Technology*, pp. 1–14, 2021.
- [20] M. Brown and J. C. Gerdes, "Robust Stabilization and Collision Avoidance through Minimizing Open-Loop Velocity Uncertainty," *IEEE Intelligent Vehicles Symposium, Proceedings*, pp. 259–264, 2020.
- [21] E. Rossetter, J. Switkes, and J. Gerdes, "A gentle nudge towards safety: experimental validation of the potential field driver assistance system," in *Proceedings of the 2003 American Control Conference, 2003.*, vol. 5, Jun. 2003, pp. 3744–3749.
- [22] D. P. Martin and G. F. Schaefer, "Tire-Road Friction in Winter Conditions for Accident Reconstruction," *SAE Transactions*, vol. 105, pp. 732–750, 1996, publisher: SAE International.
- [23] J. M. Wille, F. Saust, and M. Maurer, "Stadtpilot: Driving autonomously on Braunschweig's inner ring road," in *2010 IEEE Intelligent Vehicles Symposium*, Jun. 2010, pp. 506–511.
- [24] J. A. E. Andersson, J. Gillis, G. Horn, J. B. Rawlings, and M. Diehl, "CasADi – A software framework for nonlinear optimization and optimal control," *Mathematical Programming Computation*, vol. 11, no. 1, pp. 1–36, 2019.
- [25] A. Wächter and L. T. Biegler, "On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming," *Mathematical Programming*, vol. 106, no. 1, pp. 25–57, May 2006.
- [26] M. Quigley, K. Conley, B. Gerkey, J. Faust, T. Foote, J. Leibs, R. Wheeler, and A. Y. Ng, "ROS: an open-source Robot Operating System," *ICRA workshop on open source software*, vol. 3, no. 3.2, p. 5, 2009.



**John Talbot** received the Bachelor of Mechanical Engineering degree from Auburn University, Auburn, AL, USA in 2016. He is currently pursuing the Ph.D. degree in mechanical engineering with Stanford University, Stanford, CA, USA. His current research interests include advanced driver assistance systems as well as vehicle motion planning and control.



**Matthew Brown** received the B.S. degree in mechanical engineering from the University of California at Berkeley, Berkeley, CA, USA in 2013 and the Ph.D. degree in mechanical engineering Stanford University, Stanford, CA, USA, in 2020. He is currently a Research Scientist at Woven by Toyota. He is interested in trajectory planning and control of autonomous vehicles, particularly the use of model predictive control to control vehicles in highly dynamic scenarios. He is also interested in robust control and methods for leveraging learned models.



**J. Christian Gerdes** received the Ph.D. degree from the University of California at Berkeley, Berkeley, CA, USA, in 1996. He is currently Professor Emeritus of mechanical engineering at Stanford University, Stanford, CA, USA. His laboratory studies how cars move, how humans drive cars, and how to design future cars that work cooperatively with the driver or drive themselves. He can often be found at the racetrack with students trying out their latest prototypes for the future.