

Modeling and Control for Dynamic Drifting Trajectories

Trey P. Weber  and J. Christian Gerdes , *Member, IEEE*

Abstract—Drifting, or cornering with rear tires that exceed slip limits, represents a trade-off of stability for controllability while operating at the limits of friction. Recent work has demonstrated exceptional performance by autonomous systems of stabilization and path tracking a vehicle around an unstable drifting equilibrium. However, safely navigating unexpected or challenging road conditions that require an autonomous vehicle to operate at the limits of friction is likely to require dynamic, non-equilibrium maneuvers. These trajectories activate underlying dynamics, such as weight transfer and wheelspeed, which significantly affect the forces acting on the vehicle. In this paper, we present a modeling and control framework for dynamic drifting trajectories. First, a novel vehicle model is proposed that strikes an appropriate balance of fidelity and complexity. Then, this vehicle model is embedded into a Nonlinear Model Predictive Control policy that can maintain stability and path tracking while performing dynamic drifting maneuvers. This work is validated experimentally using “Takumi”, an autonomous Toyota Supra, that demonstrates root mean squared path tracking error of 13 centimeters and a peak error of just 47 cm. Finally, a simulation study suggests parameter uncertainty, rather than additional model fidelity, is the primary limitation of further increasing controller performance.

Index Terms—Autonomous vehicles, control and optimization, drifting, nonlinear model predictive control, vehicle dynamics.

I. INTRODUCTION

PASSENGER vehicles are conventionally designed to understeer at the limit of handling. While cornering with maximum lateral acceleration, the front tires tend to reach the friction limit before the rear. This means the vehicle’s rotational dynamics remain predictable and stable - but at the expense of controllability. When the front tires are friction saturated, additional steering input can no longer alter the vehicle’s trajectory [1]. In contrast, limit oversteer - cornering with rear tires at the friction limit - offers the opposite trade-off. The vehicle maintains controllability despite operating at the saturation limit of the rear tires, but the rotational dynamics become unstable [2]. However, in this region of the state space unstable equilibria exist that the vehicle can operate around [3]. This is the foundation

of the sport of drifting - a cornering technique that involves deliberately oversteering and then countersteering so that the vehicle maintains a large sideslip angle. A skilled driver can exploit the stability-controllability tradeoff and perform agile maneuvers despite operating at the friction limit. Gray et al. demonstrated that drifting maneuvers, particularly in lower friction environments, have the potential to increase autonomous vehicle safety [4]. The authors implemented a hierarchical control framework with a motion primitive path planner for obstacle avoidance. In some scenarios, an obstacle could not be avoided without performing a drifting (limit oversteering) maneuver. Zhao et al. took this idea even further by justifying the use of emergency drifting maneuvers for autonomous vehicles [5]. By considering backward reachability of collision objects, the authors developed an architecture that determines when it’s absolutely necessary to switch from a baseline controller to a “beyond-the-limit” controller to avoid obstacles.

Recently, stabilization of an autonomous vehicle (AV) around a drifting equilibrium has been explored extensively. The potential to maintain stability despite rear tire force saturation - and even track a desired path while doing so - has important implications for safety systems of AVs. Velenis et al. stabilized a drifting vehicle using a linear controller that coordinates steering and drive inputs, and validated it in high fidelity simulation [6]. Hindiyeh and Gerdes achieved sustained, robust drifts of a full scale test vehicle using a sliding surface controller [7]. Several approaches have even managed to track a path in addition to maintaining stability. Goh and Gerdes followed a circular path while stabilizing a drift equilibrium in a full scale electric test vehicle [8]. The controller achieved sideslip angles as large as 45° and tracked a desired path within 0.4 meters. Peterson et al. achieved even smaller path tracking error using a linear quadratic regulator and constant speed assumptions [9]. But limiting autonomous drifting to states around *equilibria* significantly restricts the potential applications to safety systems for an AV. Controlling an AV around dynamic, *non-equilibrium* drifting trajectories could allow it to adeptly handle an emergency scenario such as maneuvering around obstacles or quickly changing directions, despite operating against the friction limit.

Dynamic drifting trajectories, characterized by rear tire force saturation, large sideslip angles, and non-zero (often large) state derivatives, are fundamentally more difficult for the controller to track. This is because they require the vehicle to operate against actuation limits and excite dynamics that would otherwise be dormant and neglected near an equilibrium. However, some approaches have been successful in following such trajectories

Manuscript received 9 October 2023; revised 13 November 2023; accepted 23 November 2023. Date of publication 8 December 2023; date of current version 29 April 2024. This work was supported in part by Toyota Research Institute, in part by the National Science Foundation, and in part by Stanford Graduate Research Fellowship. (Corresponding author: Trey P. Weber.)

The authors are with the Department of Mechanical Engineering, Stanford University, Stanford, CA 94025 USA (e-mail: tpweber@stanford.edu; gerdes@stanford.edu).

Color versions of one or more figures in this article are available at <https://doi.org/10.1109/TIV.2023.3340918>.

Digital Object Identifier 10.1109/TIV.2023.3340918

with an autonomous vehicle. Goh developed a control policy that regulates the rotation rate of the vehicle's velocity vector for path tracking and yaw acceleration for sideslip stabilization [10]. This controller executed dynamic "Figure 8" trajectories with about 1.5 meters of path tracking error in a full-scale electric test vehicle. Goh utilized online parameter estimation for the front cornering stiffness and rear tire friction, which varied by up to 44% and 19%, respectively, from their nominal values. This parameter variation potentially reflects the excitation of unmodeled dynamics such as weight transfer. Goel also performed "Figure 8" trajectories using the same electric test vehicle but using optimal control techniques [11]. This approach achieved a path tracking error of about 0.7 meters by using front brakes as an additional actuator. However, the author specifically notes the lack of a weight transfer model as a limitation of their approach and a primary direction for future work.

To harness the full potential of dynamic, non-equilibrium drifting trajectories, these modeling elements need to be addressed. Weight transfer is a well-studied phenomenon in more orthodox vehicle applications, such as while trail-braking during a cornering maneuver [12], but has not been extensively studied for drifting. Furthermore, it is well documented that tire parameters such as cornering stiffness and friction coefficient are a function of the normal load of the tire, creating a complex coupling between tire forces and weight transfer [13], [14]. Wheelspeed dynamics also play an important role in determining the forces acting on the driven axle while drifting. Considering the velocity vector of the rear tire contact patch using wheelspeed is critical in achieving accurate desired forces while drifting [10]. Wheelspeed dynamics encode the physical limitations on how quickly the rear tire forces can be changed, which is a limiting factor in performing transient drifting maneuvers. Accounting for these dynamics crucially allows us to select a control policy that *anticipates* the vehicle's behavior during a transient maneuver, rather than *reacting* to parameter variation or model mismatch.

In this paper, we present the first combination of dynamic model and controller capable of leveraging drifting to execute highly dynamic maneuvers at a level of precision necessary for collision avoidance without relying on additional actuation. This involves extending previous models to explicitly include wheel-speed and dynamic longitudinal load transfer. Furthermore, we present some evidence that the residual tracking error with this approach lies within that expected from reasonable parameter variation, suggesting that this level of modeling complexity is not only necessary but potentially sufficient for control.

This paper is structured as follows: in Section II we describe the vehicle model developed for this work. Section III details the NMPC formulation that allows us to embed the vehicle model directly in the high-level control policy. Then, Section IV introduces Takumi, the full scale vehicle testing platform used in this work. Section V presents simulations and real world experiments to validate the vehicle model and controller before presenting results for dynamic drifting trajectories in Section VI. Finally, in Section VII we discuss the convergence of the NMPC optimization, compare our approach with the state-of-the-art, and present a simulation study on tire parameter uncertainty before concluding remarks in Section VIII.




Fig. 1. Single track vehicle model with reference path. Variable definitions are included in Table I.

TABLE I
VARIABLE DEFINITIONS & PARAMETER VALUES

Variable	Symbol	Units
Yaw Rate	r	rad/s
Velocity	V	m/s
Body Sideslip Angle	β	rad
Longitudinal Weight Transfer	ΔF_z	N
Rear Axle Speed	ω_R	rad/s
Distance Along Path	s	m
Lateral Path Error	e	m
Course Error	$\Delta\phi$	rad
Front Axle Steering Angle	δ	rad
Rear Axle Torque	τ	Nm
Front Axle Slip Angle	α_F	rad
Front Axle Lateral Force	F_{yF}	N
Rear Axle Slip Angle	α_R	rad
Rear Axle Slip Ratio	κ_R	<i>none</i>
Rear Axle Lateral Force	F_{yR}	N
Rear Axle Longitudinal Force	F_{xR}	N
Path Curvature	κ_{ref}	$1/m$
Velocity Vector Rotation Rate	ϕ	rad/s
Parameter	Symbol	Value
Mass	m	1496 kg
Yaw moment of inertia	I_z	2241 kgm^2
Distance from CG to front axle	a	1.217 m
Distance from CG to rear axle	b	1.231 m
Total Wheelbase	L	2.448 m
Height of CG	h_{cg}	0.45 m
Front axle affine C_α parameters	c_0	82845 N/rad
	c_1	-32.5 $1/rad$
Rear axle lateral stiffness	C_y	185 kN/rad
Rear axle longitudinal stiffness	C_x	213 kN/rad
Front axle friction coefficient	μ_F	1.00
Rear axle affine friction parameters	μ_0	0.814
	μ_1	$8.2 * 10^{-6} 1/N$
ΔF_z first-order dynamics gain	K	5 $1/s$
Drivetrain inertia	J	15 kgm^2
Wheel radius	R_e	0.32 m

II. VEHICLE MODELING

A. Single Track Vehicle Model

The vehicle is modeled as a rigid body, with the tires on each axle lumped together as seen in Fig. 1. Variable definitions and parameter values for the vehicle model are listed in Table I. The rate of change of the vehicle's yaw rate, velocity, and sideslip can then be derived using Newton's Laws:

$$\dot{r} = \frac{1}{I_{zz}} [a(F_{yF} \cos(\delta)) - bF_{yR}] \quad (1)$$

$$\dot{V} = \frac{1}{m} [F_{xR} \cos(\beta) - F_{yF} \sin(\delta - \beta) + F_{yR} \sin(\beta)] \quad (2)$$




Fig. 2. Velocity and acceleration configurations for two different sideslip angles.

$$\dot{\beta} = \frac{1}{mV} [F_{yF} \cos(\delta - \beta) - F_{xR} \sin(\beta) + F_{yR} \cos(\beta)] - r \quad (3)$$

To place the vehicle's center of mass relative to the reference path, the following path coordinates are used: $[s, e, \Delta\phi]$. The rate of change of these path coordinates can be written as follows:

$$\dot{s} = \frac{V \cos \Delta\phi}{1 - e\kappa_{\text{ref}}} \quad (4)$$

$$\dot{e} = V \sin \Delta\phi \quad (5)$$

$$\Delta\dot{\phi} = \dot{\phi} - \dot{\phi}_{\text{ref}} \quad (6)$$

$$\dot{\phi} = \dot{\beta} + r \quad (7)$$

B. Weight Transfer Dynamics

Dynamic weight transfer has a large influence on vehicle design. For example, performance oriented vehicles are generally rear wheel drive to take advantage of the rearward shift in normal load under acceleration. And for most vehicles, brakes are larger on the front axle to take advantage of the additional weight borne by those tires under braking.

Fig. 2 illustrates a drifting vehicle in two configurations: -55° and -25° of sideslip angle, which represents the sideslip range demonstrated later in this paper. Because drifting is a cornering technique, the vehicle has a large centripetal acceleration even when speed is constant. Due to larger sideslip angles than while conventionally cornering, this acceleration vector is pointed towards the front of the vehicle, causing a significant amount of weight to shift to the rear axle.

Longitudinal weight transfer has an important effect on the vehicle dynamics at the limit of friction. The yaw dynamics are particularly sensitive. Any longitudinal load transfer changes the total force capacity of each axle and their ability to balance the total yaw moment. In contrast, *lateral* weight transfer has a minimal effect on the total axle forces (i.e. F_{yF} , F_{yR} , or F_{xR} in Fig. 1) and the net moment due to a difference in coaxial forces is relatively small. Capturing the effect of lateral weight transfer presents a significant increase in model complexity. Considering four tires independently requires more vehicle states, more complicated dynamics, and twice the number of tire parameters. As demonstrated later in Section VII, dynamic

drifting maneuvers have a sensitivity to tire parameter variation. Using a model that requires more parameters could amplify this effect. Finally, the controller presented in Section III relies on computing a nonlinear optimization problem online, which motivates developing a vehicle model that strikes an appropriate balance of complexity and fidelity. For these reasons, we only model weight transfer along the longitudinal axis of the vehicle.

We model the normal load at each axle using first order dynamics similar to Subosits [15], such that the load transfer states asymptotically approach their steady-state values due to the forces acting at each axle.

$$\Delta\dot{F}_z = -K \left[\Delta F_z - \frac{h_{cg}}{L} (F_{xR} - F_{yF} \sin(\delta)) \right] \quad (8)$$

The normal force acting at each axle is then:

$$F_{zF} = \frac{bmg}{L} - \Delta F_z \quad (9)$$

$$F_{zR} = \frac{amg}{L} + \Delta F_z \quad (10)$$

C. Rear Wheelspeed Dynamics

While drifting, the vehicle's rear axle operates at the friction limit which creates a coupling between lateral and longitudinal tire forces. Increasing F_{xR} reduces F_{yR} , and vice versa, because the tire contact patch can only generate a total force of μF_{zR} . This means that opening the throttle, which for conventional driving only affects the longitudinal dynamics of the vehicle, also has a significant effect on the yaw rate and sideslip angle.

In order to perform dynamic drifting maneuvers, rear wheel speed is used to quickly rotate the rear tire force vector between the lateral and longitudinal directions. To capture the physical limitations on how fast this vector can be rotated, the rate of change of the rear axle wheelspeed is modeled as a function of applied torque. These dynamics are found by performing a moment balance on the rear axle:

$$J\omega_R = \tau - R_e F_{xR} \quad (11)$$

D. Tire Models

For the front tire, the force-slip relationship is represented using a brush tire model that assumes purely lateral slip [16]. The front slip angle is calculated using the single track vehicle kinematics as follows:

$$\alpha_F = \tan^{-1} \left(\frac{V \sin(\beta) + ar}{V \cos(\beta)} \right) - \delta \quad (12)$$

The lateral front tire force is then defined using:

$$F_y(\alpha) = \begin{cases} -C_\alpha \tan(\alpha) + \frac{C_\alpha^2}{3F_{y,\max}} |\tan(\alpha)| \tan(\alpha) \\ -\frac{C_\alpha^3}{27F_{y,\max}^2} \tan^3(\alpha), & |\alpha| \leq \alpha_{\text{slide}} \\ -F_{y,\max} \text{sign}(\alpha), & |\alpha| > \alpha_{\text{slide}} \end{cases} \quad (13)$$

where C_α is the cornering stiffness, $\alpha_{\text{slide}} = \tan^{-1}(\frac{3F_{y,\max}}{C_\alpha})$, and $F_{y,\max} = \mu F_z$, because the total force available is related only to friction and normal load when there is no longitudinal braking force applied.

The combined effects of large lateral and longitudinal rear tire slip are captured with a coupled-slip tire model [14]. The rear slip angle and slip ratio are computed as follows:

$$\alpha_R = \tan^{-1} \left(\frac{V \sin(\beta) - br}{V \cos(\beta)} \right) \quad (14)$$

$$\kappa_R = \frac{R_e \omega_R - V \cos(\beta)}{V \cos(\beta)} \quad (15)$$

Then, the rear tire forces are defined as:

$$f = \sqrt{C_x^2 \left(\frac{\kappa_R}{\kappa_R + 1} \right)^2 + C_y^2 \left(\frac{\tan(\alpha_R)}{\kappa_R + 1} \right)^2} \quad (16)$$

$$F = \begin{cases} f - \frac{f^2}{3\mu_R F_{zR}} + \frac{f^3}{27(\mu_R F_{zR})^2}, & f \leq 3\mu_R F_{zR} \\ \mu_R F_{zR}, & f > 3\mu_R F_{zR} \end{cases} \quad (17)$$

$$F_{xR} = \frac{FC_x \kappa_R}{f(\kappa_R + 1)} \quad (18)$$

$$F_{yR} = -\frac{FC_y \tan(\alpha_R)}{f(\kappa_R + 1)} \quad (19)$$

E. Modified Tire Model Parameters

Below the limits of friction, cornering stiffness determines the sensitivity of lateral tire force to slip angle. Normal load has a strong effect on cornering stiffness. This effect is approximated as affine in a region around nominal operating conditions similar to Pacejka [14]:

$$C_{\alpha F}(\Delta F_z) = c_0 + c_1 \Delta F_z \quad (20)$$

We use a similar approach for the rear friction coefficient:

$$\mu_R(\Delta F_z) = \mu_0 + \mu_1 \Delta F_z \quad (21)$$

III. NONLINEAR MODEL PREDICTIVE CONTROL

Performing dynamic drifting maneuvers requires the vehicle to operate in an unstable region of the state space, close to actuator limitations, and balance multiple objectives. This problem is well suited for Nonlinear Model Predictive Control (NMPC). NMPC can consider nonlinear dynamics, reason about actuator constraints, and can be used for trajectory following. Due to the receding horizon nature of this control strategy and the ability to directly embed our vehicle model, NMPC enables us to consider how control inputs will affect state evolution over future time-steps and track complex trajectories.

In the following section, we detail the setup of the NMPC problem. First we discretize our dynamic vehicle model from Section II into equality constraints, ensuring trajectories computed by NMPC are dynamically feasible with respect to our model. Then we describe the cost function which embeds the state tracking objective and the input constraints which describe the actuator limitations of the vehicle.

The NMPC algorithm is implemented using CasADi - an open source tool for nonlinear optimization and algorithmic differentiation [17]. We utilize an interior point optimizer (IPOPT) as the solver and use Robot Operating System (ROS) to manage timing




Fig. 3. Block diagram for the overall NMPC structure. The reference trajectory and measured state from the vehicle are inputs to the NMPC problem. The previously computed control horizon is used to simulate forward the measured state for delay compensation. Steering and torque commands are sent to the vehicle at every time step from the most recently computed control horizon.

between function calls to the NMPC algorithm and sending actuator commands to the vehicle [18], [19]. A block diagram for the overall NMPC structure can be seen in Fig. 3.

A. Dynamics Constraints

Gathering the dynamic equations from Section II, our state vector is $x = [r, V, \beta, s, e, \Delta\phi, \omega_R, \Delta F_z, \delta, \tau]^T$ and input vector is $u = [\dot{\delta}, \dot{\tau}]^T$. Modeling inputs as the rate of change of our actuators allows us to constrain and cost the actuator slew rates directly. For the dynamic constraints that link each stage of our time-horizon, we convert to a spatial coordinate system of our path distance variable s . This is a convenient way to formulate the drifting problem, as it allows us to cost our deviation from any reference states with respect to distance along the desired path, rather than time.

$$\frac{dx}{ds} = \frac{dx}{dt} \frac{dt}{ds} = \frac{\dot{x}}{\dot{s}} \quad (22)$$

The state horizon is integrated trapezoidally:

$$\begin{aligned} \int_{s_k}^{s_{k+1}} \frac{\dot{x}}{\dot{s}} ds &= \int_{s_k}^{s_{k+1}} dx = x_{k+1} - x_k \\ &\approx \frac{1}{2}(s_{k+1} - s_k) \left(\left. \frac{dx}{ds} \right|_{s_{k+1}} + \left. \frac{dx}{ds} \right|_{s_k} \right) \end{aligned} \quad (23)$$

Finally, for a uniformly spaced horizon of N stages, the dynamic equality constraints for the optimization problem can be written as:

$$x_{k+1} = x_k + \frac{1}{2} \Delta s \left(\left. \frac{dx}{ds} \right|_{s_{k+1}} + \left. \frac{dx}{ds} \right|_{s_k} \right) \quad (24)$$

$$\forall k \in \{1, \dots, N-1\} \quad (25)$$

$$x_1 = x_{\text{initial}}$$

B. Cost Function and Actuator Constraints

Before formulating the cost function, reference states are subtracted and the quantity is normalized with a maximum desired value. Reference states are computed by taking the current measured s position and interpolating from a predefined desired path. This serves two purposes: it makes cost function weight tuning more intuitive and when gradients are scaled appropriately, the optimization algorithm (IPOPT) will converge

faster.

$$\bar{x} = \left[\frac{r-r_{\text{ref}}}{e_{r,\text{max}}}, V, \frac{\beta-\beta_{\text{ref}}}{e_{\beta,\text{max}}}, s, \frac{e}{e_{\text{max}}}, \frac{\Delta\phi}{e_{\Delta\phi,\text{max}}}, \frac{\omega_R-\omega_{R,\text{ref}}}{e_{\omega_R,\text{max}}}, \Delta F_z, \delta, \tau \right] \quad (26)$$

$$\bar{u} = \left[\frac{\dot{\delta}-\dot{\delta}_{\text{ref}}}{e_{\dot{\delta},\text{max}}}, \frac{\dot{\tau}-\dot{\tau}_{\text{ref}}}{e_{\dot{\tau},\text{max}}} \right] \quad (27)$$

where $e_{*,\text{max}}$ is the maximum desired deviation from reference for that state variable. The objective function is formulated as quadratic in state and input vectors and summed along the horizon:

$$J = \sum_{j=1}^N (\bar{x}_j^T Q \bar{x}_j + \bar{u}_j^T R \bar{u}_j) + \bar{x}_N^T Q_{\text{term}} \bar{x}_N \quad (28)$$

and the Q and R matrices are diagonal:

$$Q = \text{diag}([k_r \ 0 \ k_\beta \ 0 \ k_e \ k_{\Delta\phi} \ k_{\omega_R} \ 0 \ \dots]) \quad (29)$$

$$R = \text{diag}([k_{\dot{\delta}} \ k_{\dot{\tau}}]) \quad (30)$$

$$Q_{\text{term}} = k_{\text{term}} Q \quad (31)$$

Relatively large costs are put on tracking the desired path (k_e) and desired sideslip angle (k_β). A small cost on rear wheelspeed (k_{ω_R}) encourages the controller to operate close to the equilibrium wheelspeed and use steering for any small corrections, while $k_{\Delta\phi}$ improves damping in the lateral path tracking state. Finally, a very small cost on yaw rate (k_r) aids in convergence time of the nonlinear optimization, while not overly restricting the controller's ability to use yaw rate to generate sideslip angle. Costs on actuator slew rates ($k_{\dot{\delta}}$ and $k_{\dot{\tau}}$) help ensure a smooth closed-loop response.

The terminal cost matrix Q_{term} incentivizes the NMPC to plan trajectories that are continuously feasible. In addition to the dynamic state constraints in Section III-A, actuator limitations are imposed on steering, steering rate of change, wheel torque, and wheel torque rate of change.

The resulting optimization problem is:

$$\underset{x,u}{\text{minimize}} \quad J \quad (32)$$

$$\text{subject to } x_1 = x_{\text{initial}} \quad (33)$$

$$x_{k+1} = x_k + \frac{1}{2} \Delta s \left(\frac{dx_k}{ds} + \frac{dx_{k+1}}{ds} \right) \quad (34)$$

$$\forall k \in [1, N-1]$$

$$\delta_{\min} \leq \delta_k \leq \delta_{\max} \quad (35)$$

$$\dot{\delta}_{\min} \leq \dot{\delta}_k \leq \dot{\delta}_{\max} \quad (36)$$

$$\tau_{\min} \leq \tau_k \leq \tau_{\max} \quad (37)$$

$$\dot{\tau}_{\min} \leq \dot{\tau}_k \leq \dot{\tau}_{\max} \quad (38)$$

where the actuator limitations in (35)–(38) come from the physical specifications of the vehicle, and x_{initial} comes from a delay compensation scheme that projects the measured state forward



Fig. 4. “Takumi”: A purpose-built Toyota Supra drifting platform. Photo Credit: Alex Green.

in time by the nominal NMPC solve time, similar to Brown and Gerdes [20].

The controller runs at a nominal rate of 50 Hz and is implemented on the drifting platform described in the following section.

C. Horizon Selection

The length of the NMPC horizon (30 meters) was selected so that the NMPC can plan through the entire dynamic drifting maneuver performed in Section VI (~20 meters).

Previous work has shown 50 ms to be an adequate integration time for vehicle dynamics [21]. Specifically, for NMPC at the friction limits, Brown and Gerdes compared a variety of integration methods and sampling times [20]. A 50 ms sampling time using second order Runge-Kutta integration was found to perform very well and with reasonable computation time. For this work, Δs was selected as 0.5 meters to correspond to a sampling time of approximately 50 ms for the trajectories performed (velocity of ~10 m/s).

IV. TAKUMI: A PURPOSE-BUILT TOYOTA SUPRA DRIFTING PLATFORM

Takumi is a fifth-generation Toyota Supra prototype vehicle that has been heavily modified to become an autonomous drifting platform (Fig. 4). Autonomous experiments are performed with a driver behind the wheel for safety. However, data in this work is recorded while the safety driver is not in control of any actuators.

A. Powertrain

Takumi is powered by a BMW B58 inline 6 cylinder engine with an aftermarket BorgWarner turbocharger. A 6-speed Samsonas sequential transmission and OS Giken limited slip differential deliver power to the rear wheels. Takumi produces 526 HP and 536 ft-lbs of torque @ 6450 rpm.

B. Actuation

Takumi's steering system has been modified with a Nanotech drive motor, Sendix F3668 optical CAN encoder, and Wisefab

WFA90 front steering angle suspension kit. Takumi can command ± 42 degrees of steering at the road wheel. A proprietary brake-by-wire system allows for independent braking pressure to be commanded at any of the four wheels. Finally, a Motec M142 ECU provides throttle-by-wire capability. Appendix B details the implementation of the steering and engine low-level control laws.

C. Sensing and Computation

Takumi receives position, heading, velocity, and acceleration measurements from an Oxford Technical Systems RT4003 GNSS/INS with dual GPS antennae. A dSpace MicroAutoBox II real time computer runs a Simulink model of the vehicle. This manages the timing and communication between the various low-level systems on the car - such as the Motec ECU, an auxiliary computer running the NMPC framework, and various sensors.

A custom-built RAVE Linux computer runs the ROS program that interfaces the NMPC framework and the MicroAutoBox via UDP packets sent over ethernet. This computer is equipped with an Intel Xeon 8 core, 3.3 GHz CPU and NVIDIA GeForce RTX 3080 Ti GPU.

V. MODEL VALIDATION

Before testing the vehicle model and controller proposed in Sections II and III for dynamic drifting trajectories, it's important to first assess the validity of this approach. The following section describes simulations and experiments conducted to demonstrate the importance of the additional modeling of weight transfer and wheelspeed dynamics for drifting.

An experiment was designed to excite weight transfer and wheelspeed dynamics and observe the controller's tracking performance. While drifting along a constant radius path of 10 meters, the reference sideslip angle sweeps between -55° and -25° , as depicted in Fig. 2. This sideslip angle range spans the actuation space of Takumi, requiring steering angles as large as -38° . Additionally, it corresponds to a ΔF_z range of $\{916, 1705\}$ N, or about 12% to 23% change of the static normal load of either axle, and a large rear wheelspeed range of $\{33, 64\}$ radians/second.

A. Simulation Results

We conduct three simulations, using parameters listed in Tables I & II. First, the NMPC controller is simulated using all of the dynamics described in Section II in both the simulator and dynamical constraints of the controller (34). Then, two more simulations are conducted, one without the weight transfer dynamics (8) and one without wheelspeed dynamics (11) in the controller. In Fig. 5 the baseline case (including the full vehicle model in the NMPC controller) tracks the desired reference states closely. Then, omitting the weight transfer model results in increased tracking errors in both the low and high sideslip regions. This matches intuition, as the tire parameters for this simulation were selected using nominal normal loads for a sideslip of -40° . Finally, the simulation omitting wheelspeed

TABLE II
CONTROLLER PARAMETERS

Parameter	Symbol	Value
Number of NMPC horizon stages	N	60
NMPC horizon spacing	Δs	0.5 m
Maximum steering angle	δ_{max}	43°
Maximum steering angle rate	$\dot{\delta}_{max}$	$90^\circ/s$
Minimum steering angle rate	$\dot{\delta}_{min}$	$-90^\circ/s$
Maximum torque	τ_{max}	3.5 kNm
Minimum torque	τ_{min}	-1 kNm
Maximum torque rate	$\dot{\tau}_{max}$	2 kNm/s
Minimum torque rate	$\dot{\tau}_{min}$	-4 kNm/s
Maximum desired yaw rate error	$e_{r,max}$	45°
Maximum desired sideslip angle error	$e_{\beta,max}$	1°
Maximum desired course error	$e_{\Delta\phi,max}$	1°
Maximum desired path tracking error	$e_{e,max}$	0.05 m
Maximum desired wheelspeed error	$e_{\omega_R,max}$	1 rad/s
Maximum desired steering slew rate error	$e_{\dot{\delta},max}$	$4^\circ/s$
Maximum desired torque slew rate error	$e_{\dot{\tau},max}$	4 kNm/s
Cost on yaw rate error	k_r	1.0
Cost on sideslip angle error	k_β	0.8
Cost on course error	$k_{\Delta\phi}$	1.0
Cost on path tracking error	k_e	1.2
Cost on wheelspeed error	k_ω	1.0
Cost on steering slew rate	$k_{\dot{\delta}}$	1.0
Cost on torque slew rate	$k_{\dot{\tau}}$	0.5
Cost on terminal states	k_{term}	25




Fig. 5. Simulations of a sideslip variation trajectory to excite weight transfer and wheelspeed dynamics. First a baseline simulation with the full vehicle model in the NMPC constraints is conducted. Then, weight transfer and wheelspeed dynamics are removed.

dynamics results in large oscillations in all states as the controller begins to attempt to track the changing sideslip angle reference. This is due to the controller's inherent assumption that the rear tire forces can be changed according to the $\dot{\tau}$ limits in (38) and there is no drivetrain inertia. The simulation results in oscillations that don't decay until the reference sideslip angle no longer changes.

B. Experimental Results

To validate the conclusions drawn from simulation, they are repeated as experiments. These tests were performed using Takumi, the vehicle platform described in Section IV, on a




Fig. 6. Vehicle states during the sideslip variation experiment. Data is overlaid from two experiments, one with and one without incorporating longitudinal weight transfer in the vehicle model. From (14), Takumi's rear slip angle can be estimated to be approximately 50° , well past the peak slip angle of the tires.




Fig. 7. Estimated longitudinal weight transfer from Takumi's onboard accelerometer overlaid with the predicted state from NMPC.

closed course at Thunderhill Raceway Park in Willows, CA. The sideslip variation experiment is conducted first with weight transfer and wheelspeed dynamics in the vehicle model of the NMPC controller (baseline case), and then without weight transfer dynamics. The simulation omitting wheelspeed dynamics was considered unsafe to test on a real vehicle.

Takumi's tracking performance for velocity, sideslip angle, and lateral path error can be seen in Fig. 6. Takumi achieved RMS errors of 0.1 m/s, 1.7 degrees, and 6.7 cm, respectively, in the baseline (full vehicle model) case. When weight transfer is removed from the vehicle model, there is similar tracking performance in velocity and sideslip angle, but the controller is unable to track the path with the same accuracy (RMS errors of 0.12 m/s, 1.6 degrees, and 25.8 cm, respectively). These two experimental results closely resemble the corresponding simulations in Fig. 5, with respect to trend and magnitude of the state errors. This underlines the significant improvement in controller performance that the vehicle model developed in Section II provides.

By modeling weight transfer and wheelspeed dynamics in the high level tracking controller, Takumi maintains consistent performance despite the redistribution of almost 1 kN of normal force throughout the experiment, verified by an onboard accelerometer in Fig. 7. This weight transfer has a peculiar




Fig. 8. Observed front lateral tire force (see Appendix A) normalized by maximum lateral tire force and plotted against slip angle. Increasing ΔF_z corresponds to decreasing normal load on the front axle.




Fig. 9. Modeled forces with and without weight transfer and observed forces for sideslip variation experiment.

effect on the front tire curve in Fig. 8. There is a significant reduction in F_{yF} of $\sim 20\%$ as slip angle magnitude increases from -8° to -13° . This effect comes from the sensitivity of front cornering stiffness to normal load, as modeled by (20). While front slip angle increases in magnitude, the normal load decreases due to the change in sideslip angle which results in reduction of front cornering stiffness. The net effect is less lateral force. In Fig. 9, tire forces are plotted throughout the experiment. Observed tire forces (computed using the force observer described in Appendix A) are overlaid with the modeled forces with and without weight transfer. While the front cornering stiffness sensitivity causes a large model mismatch in F_{yF} , the rear tire forces (F_{xR} , F_{yR}) are less sensitive to the changing normal load. This is because the rear friction sensitivity to weight transfer is small and any change in normal load is distributed between the longitudinal and lateral components of the rear tire force. Additionally, by modeling the time delays of wheelspeed associated with drivetrain inertia, we do not observe any of the oscillations seen in the simulated case of omitting wheelspeed dynamics.




Fig. 10. Overhead illustration of dynamic transition maneuver (half of “Figure 8”).

By incorporating longitudinal weight transfer and tire parameter sensitivity into a model predictive control policy, changes in normal load and the subsequent effect on the dynamics are anticipated. Longitudinal weight transfer has a dominating effect on the front tire model due to cornering stiffness sensitivity. A direct comparison of a drifting experiment with/without weight transfer dynamics demonstrates an improvement of path tracking RMS error from 25.8 to 6.7 cm (and peak error from 50 cm to 10 cm). Additionally, incorporating wheelspeed dynamics is essential to prevent state oscillations due to rear tire force delays caused by drivetrain inertia.

VI. DYNAMIC TRAJECTORY VALIDATION

The vehicle model and controller are now put to the test of tracking a dynamic “Figure 8” trajectory. This experiment requires Takumi to rapidly transition between -40° and 40° of sideslip angle in just 2 seconds (an overhead view of the path can be seen in Fig. 10). Two such maneuvers are linked together to create a full “Figure 8” trajectory. During each transition maneuver, the trajectory requires Takumi to rapidly vary wheelspeed between 44 and 34 rd/s. The resulting acceleration causes ΔF_z to vary between 1400 and 450 N over the same duration. The reference trajectory was computed offline in a fashion similar to Goh [10].

A. Experimental Results

Takumi navigates the “Figure 8” trajectory with ease, successfully stabilizing the vehicle despite rapid changes in sideslip angle and wheelspeed. Throughout the experiment, Takumi achieves RMS errors of 0.24 m/s, 2.4° , and 13 cm in velocity, sideslip angle, and lateral path error (Fig. 11). Peak lateral errors for each dynamic transition maneuver of the trajectory are just 47 and 37 centimeters.

In Fig. 13, modeled forces are overlaid with observed forces during the “Figure 8” experiment. The tire models in Section II perform well, closely matching the observed forces throughout the trajectory. There are small disturbances in F_{yF} and F_{xR} , particularly at $s = 260$ and 380 meters. These disturbances coincide with peaks in path tracking error in Fig. 11. Once the vehicle departs from the path, the controller encounters a predicament: the velocity vector must be rotated such that the vehicle reduces path tracking error but also maintains the desired




Fig. 11. Takumi successfully performs a “Figure 8” experiment, maintaining sideslip stabilization and tracking despite the excitation of wheelspeed and longitudinal weight transfer dynamics.



Fig. 12. Overhead composite photo of Takumi during the “Figure 8” maneuver - each frame taken 0.5 seconds apart. Photo Credit: John Talbot.




Fig. 13. Modeled and observed forces during “Figure 8” experiment.

sideslip angle. These competing demands can normally be resolved by assigning a higher cost to tracking the path in the cost function. However, this relationship is not as straightforward during a dynamic drifting trajectory as the vehicle operates near actuation limits and with large yaw rates. Returning to the path too aggressively incurs large costs in other states and risks exceeding stability limits. This behavior was observed empirically. Further increasing the path tracking cost (k_e) begins to induce oscillations in yaw rate, sideslip angle, and velocity. Eventually, Takumi becomes unstable during the transition of the “Figure 8” when the path tracking cost is too large. This suggests that for this control strategy and vehicle model, path tracking performance cannot be further increased without significant trade-offs with stability.

In summary, force disturbances during a dynamic drifting trajectory present complex trade-offs for the controller to navigate. The result in this paper demonstrate the robustness of the NMPC, as disturbances of several thousand Newtons incur less than half a meter of path tracking error, before the controller quickly returns to the path while maintaining stability. In the following section, we discuss the convergence of the NMPC, compare these results with prior work, and suggest limitations on further increasing controller performance.

VII. DISCUSSION

A. Remarks on NMPC Convergence

Due to the nonlinearity of the vehicle model from Section II, the optimization has no guarantee of convergence. Empirically, the optimization converged reliably at a mean solution time of 18.91 ms during the experiment in Section VI. A possible explanation for this efficiency of the NMPC lies in an analysis performed by Peterson et al. regarding the underlying linear structure of the single track vehicle model [9]. Simulations of a linearized vehicle model while drifting matched the full nonlinear dynamics well. The eigenstructure was shown to vary smoothly and consistently even as the linearization point changes.

This characteristic of the single track vehicle model at large sideslip angles could explain why the nonlinear solver IPOPT reliably solves our NMPC problem. While IPOPT relies on a complicated line-search and barrier method, the underlying solver uses Newton’s method for each step [18]. Applying Newton’s method on a linearly-constrained quadratic program has favorable convergence properties.

B. Comparison With Prior Art

The experimental results in Section VI demonstrate that our vehicle model and controller can perform highly dynamic drifting trajectories that exceed the tracking performance of prior work. Goh first achieved dynamic drifting maneuvers using a controller that imposes stable path tracking and sideslip dynamics by inverting a single track vehicle model [10]. While similar to the vehicle model described in this work, it does not incorporate wheelspeed dynamics or weight transfer. Goh used online parameter estimators for the front axle cornering stiffness and rear axle friction coefficient, which varied as much as 44%




Fig. 14. “Figure 8” transition is simulated with parameter uncertainty and compared to experimental results from Takumi.

and 19%. Goh’s work is a significant achievement - the first demonstration of such maneuvers by an autonomous vehicle - but had about 1.5 meters of path tracking error. This may not be sufficient for the safety system of an autonomous vehicle. Our work builds on Goh’s result by modeling the dynamics causing this parameter variation and using a predictive controller that can anticipate these effects, rather than react to them.

Goel demonstrated impressive results for a similar “Figure 8” maneuver by leveraging the same vehicle model as Goh in a Nonlinear Model Predictive Controller and incorporating front brakes as an additional actuator [11]. This approach achieved a path tracking error of 0.7 meters for the maneuver. Goel highlighted that the use of front brakes was a significant factor in managing model mismatch, particularly from neglecting longitudinal weight transfer. Additionally, Goel incorporated hard constraints on the magnitude of front lateral force, front braking force, and rear drive force, which may limit the possible range of maneuvers or tracking performance. Goel’s call to attention for the need to model longitudinal weight transfer was an important inspiration for our work. To directly quantify this difference, Fig. 6 compares our model with one omitting longitudinal weight transfer. Finally, when applying our model to a dynamic drifting maneuver, the controller achieves a peak tracking error of 47 cm without the use of additional actuation (front brakes).

C. Sensitivity to Parameter Variation

How much further can path tracking error be improved? Is there a control strategy, or additional modeling that can further reduce the tracking error for dynamic drifting trajectories? Potentially, but tracking precision is inherently limited by the certainty of our model parameters, particularly the tires.

Tire friction changes both dynamically and spatially. The large wheelspeeds associated with drifting result in a wide range of rear tire temperatures [22]. Due to large slip ratios, tire wear is a factor in determining friction as well. And even in a controlled testing environment, road surface variation occurs. For example, in Fig. 12 Takumi can be observed drifting on an asphalt surface with residual tire rubber from previous experiments. In some

areas these markings are much thicker, causing the tire friction coefficient to vary. And while we have assumed a first order longitudinal weight transfer model to address this, other effects such as pitch, roll, suspension dynamics, and lateral weight transfer are neglected. Cornering stiffness, as demonstrated by this work, is sensitive to changes in normal load. It is also a function of temperature and inflation pressure [23].

To understand how parameter uncertainty affects path tracking of dynamic drifting trajectories, we present a simulation study using the “Figure 8” trajectory. In Fig. 14, the NMPC is simulated with mismatches in tire parameters and the results are overlaid with the experimental data from Fig. 11. A tire parameter uncertainty of just 10% yields similar tracking errors to that observed empirically with Takumi. This sensitivity imposes an inherent constraint on tracking performance.

Laurense showed that accurate friction parameters are essential to operate at the true limit of handling for autonomous racing applications, and even a 10% friction estimation error is difficult to achieve [24]. When drifting, because the rear tires are constantly at the friction limit, online friction estimation may be feasible. However, this underlines the fundamental limitation of reacting to model variation instead of anticipating it. Although NMPC is effective in planning dynamic drifting maneuvers over a receding horizon, it is still vulnerable to discrepancies in the predicted dynamics due to tire parameter variation.

VIII. CONCLUSION

Using a Nonlinear Model Predictive Control policy, dynamic drifting trajectories are performed using a full scale Toyota Supra test vehicle. By considering weight transfer and wheelspeed evolution over a receding horizon, the controller navigates the trajectory with velocity, sideslip angle, and path tracking RMS errors of 0.24 m/s, 2.4°, and 13 cm respectively and a peak path tracking error of just 47 cm.

Experiments performed with and without weight transfer dynamics in the vehicle model demonstrated a significant difference in performance. A simulation study suggests that parameter uncertainty is the primary limitation to further decreasing path tracking error for transient drifting maneuvers. Advantages of further increasing model fidelity, such as a double track model with lateral weight transfer, therefore could provide diminishing returns at the expense of computational complexity.

This work presents an important step towards safely controlling autonomous vehicles in the presence of rear-tire saturation. Despite operating at the limits of handling and actuation, we achieve unprecedented tracking performance for dynamic drifting trajectories by explicitly considering wheelspeed and weight transfer dynamics. This opens up new possibilities for autonomous vehicles in emergency scenarios, particularly on low friction surfaces. Driving at the rear friction limits to avoid an obstacle or maintain stability does not need to come at the expense of motion planning fidelity.

APPENDIX A

TIRE FORCE OBSERVER

The tire force estimation problem is formulated as a linear unknown input observer where forces (F_{yF} , F_{xR} , F_{yR}) are

thought of as unknown inputs to the vehicle. The primary benefit of this method is obtaining observed tire forces while drifting without having to select tire models or tire parameters. Additionally, there's no need to take numerical derivatives of noisy measurement data, which can be problematic. First we convert our vehicle states $[r, V, \beta]$ to longitudinal and lateral velocity $[r, U_x = V \cos \beta, U_y = V \sin \beta]$. Now our dynamics are:

$$\dot{x} = f(x) + A\hat{F} \quad (39)$$

$$\begin{bmatrix} \dot{r} \\ \dot{U}_x \\ \dot{U}_y \end{bmatrix} = \begin{bmatrix} 0 \\ -rU_y \\ rU_x \end{bmatrix} + \begin{bmatrix} \frac{a \cos \delta}{I_z} & 0 & \frac{-b}{I_z} \\ \frac{-\sin \delta}{m} & \frac{1}{m} & 0 \\ \frac{\cos \delta}{m} & 0 & \frac{1}{m} \end{bmatrix} \begin{bmatrix} F_{yF} \\ F_{xR} \\ F_{yR} \end{bmatrix} \quad (40)$$

Since $\hat{F}$ is unknown, let $\hat{x}$ be the estimated dynamics. We then write $\hat{x}$ as the sum of our measured dynamics and a feedback gain (λ) on the difference between our measured and estimated states:

$$\dot{\hat{x}} = f(\hat{x}) + \lambda(x - \hat{x}) \quad (41)$$

$$\hat{s} = \hat{x} - f(\hat{x}) = \lambda(x - \hat{x}) \quad (42)$$

$$\hat{F} = A^L \hat{s} \quad (43)$$

where A^L denotes a left inverse of A .

APPENDIX B

LOW-LEVEL ACTUATOR CONTROL

The autonomous steer and throttle systems rely on low-level control laws to achieve desired actuator commands. Desired steering angle is converted to a current command for the steering motor using a feedback/feedforward policy as follows:

$$I_{\delta, \text{command}} = I_{\delta, \text{fb}} + I_{\delta, \text{ff}} \quad (44)$$

$$I_{\delta, \text{fb}} = K_p(\delta_{\text{des}} - \delta_{\text{meas}}) + K_d(\dot{\delta}_{\text{des}} - \dot{\delta}_{\text{meas}}) \quad (45)$$

$$I_{\delta, \text{ff}} = K_{\text{ff}} F_{yF, \text{des}} \quad (46)$$

where K_p and K_d are proportional and derivative gains, respectively, and K_{ff} is a constant factor that maps desired front lateral force to feedforward steering motor current.

To achieve desired engine torque, we rely on an intake manifold pressure control law. First, we compute the mass of air in the intake manifold and convert it to a desired manifold pressure and a desired air mass using the ideal gas law and dynamometer data:

$$m = \frac{P_{\text{man}} V_{\text{man}}}{RT_{\text{man}}} \quad (47)$$

$$P_{\text{man, des}} = \Theta(\tau_{\text{engine}}) \quad (48)$$

$$m_{\text{des}} = \frac{P_{\text{man, des}} V_{\text{man}}}{RT_{\text{man}}} \quad (49)$$

where R is the ideal gas constant, V_{man} is the volume of the intake manifold, T_{man} is the air temperature, and Θ is a lookup table that maps desired engine torque to desired intake manifold pressure. A feedback control policy is used to control the mass air flow into the intake manifold:

$$\dot{m}_{\text{in, des}} = \dot{m}_{\text{out}} - K_{\text{engine}}(m - m_{\text{des}}) \quad (50)$$

$$\dot{m}_{\text{out}} = \rho_{\text{air}} \dot{V}_{\text{engine}} = \frac{P_{\text{man}}}{RT_{\text{man}}} \frac{V_{\text{cyl}} \omega_{\text{rpm}}}{120} \quad (51)$$

Assuming temperature does not change, the throttle flow coefficient is written as follows [25]:

$$\frac{\dot{m}_{\text{in}}}{A_{\text{eff}}(\alpha_{\text{throttle}})} = \frac{P_{\text{boost}}}{RT_{\text{man}}} \Pi^{\frac{1}{\gamma}} \sqrt{\frac{2\gamma}{\gamma-1}} \left(1 - \Pi^{\frac{\gamma-1}{\gamma}}\right) \quad (52)$$

$$\Pi = \max\left(\frac{P_{\text{man}}}{P_{\text{boost}}}, \left(\frac{2}{\gamma+1}\right)^{\frac{\gamma}{\gamma-1}}\right) \quad (53)$$

$$C_{\text{flow}} = C_d A_{\text{eff}}(\alpha_{\text{throttle}}) = \frac{\dot{m}_{\text{in,des}}}{\dot{m}_{\text{in}}} \quad (54)$$

where P_{boost} is the turbocharger boost pressure upstream of the throttle, C_d is the throttle discharge coefficient, A_{eff} is the effective area of the throttle, and γ is the specific heat of air. Finally, we convert C_{flow} to a commanded throttle angle using a lookup table ϕ_{throttle} :

$$\alpha_{\text{throttle}} = \phi_{\text{throttle}}(C_{\text{flow}}) \quad (55)$$

ACKNOWLEDGMENT

The authors would like to thank the Bridgestone Corporation for their generous support. The authors would also like to thank their colleagues in the Dynamic Design Lab for their support and assistance in experimental testing at Thunderhill Raceway Park. A special thanks to Joe Farasati of Bay Area Drifting for his expertise in the art of drifting.

Building and maintaining a research vehicle is no small feat and would not have been possible without the efforts of Larry Cathey, John Talbot, Tushar Goel, and Marsie T. Peterson. Thanks to GReddy Performance Products, Negishi Competition Systems, and Jojo Callos for their expertise and guidance in building Takumi.

REFERENCES

- [1] R. Y. Hindiyeh and J. C. Gerdes, "Equilibrium analysis of drifting vehicles for control design," in *Proc. Dyn. Syst. Control Conf.*, 2009, pp. 181–188.
- [2] R. Y. Hindiyeh, "Dynamics and control of drifting in automobiles," Ph.D. dissertation, Stanford Univ., Stanford, CA, USA, Mar. 2013.
- [3] E. Ono, S. Hosoe, H. D. Tuan, and S. Doi, "Bifurcation in vehicle dynamics and robust front wheel steering control," *IEEE Trans. Control Syst. Technol.*, vol. 6, no. 3, pp. 412–420, May 1998.
- [4] A. Gray, Y. Gao, T. Lin, J. K. Hedrick, H. E. Tseng, and F. Borrelli, "Predictive control for agile semi-autonomous ground vehicles using motion primitives," in *Proc. IEEE Amer. Control Conf.*, 2012, pp. 4239–4244.
- [5] T. Zhao, E. Yurtsever, and G. Rizzoni, "Justifying emergency drift control for automated vehicles," in *Proc. 10th IFAC Symp. Adv. Automot. Control*, 2022, pp. 141–148.
- [6] E. Velenis, D. Katzourakis, E. Frazzoli, P. Tsotras, and R. Happee, "Steady-state drifting stabilization of RWD vehicles," *Control Eng. Pract.*, vol. 19, no. 11, pp. 1363–1376, 2011.
- [7] R. Y. Hindiyeh and J. C. Gerdes, "A controller framework for autonomous drifting: Design, stability, and experimental validation," *J. Dyn. Syst. Meas. Control*, vol. 136, no. 5, 2014, Art. no. 051015.
- [8] J. Y. Goh and J. C. Gerdes, "Simultaneous stabilization and tracking of basic automobile drifting trajectories," in *Proc. IEEE Intell. Veh. Symp.*, 2016, pp. 597–602.
- [9] M. T. Peterson, T. Goel, and J. C. Gerdes, "Exploiting linear structure for precision control of highly nonlinear vehicle dynamics," *IEEE Trans. Intell. Veh.*, vol. 8, no. 2, pp. 1852–1862, Feb. 2023.
- [10] J. Goh, "Automated vehicle control beyond the stability limits," Ph.D. dissertation, Stanford Univ., Stanford, CA, USA, Jun. 2019.
- [11] T. Goel, "In complete control; simultaneous path, speed, and sideslip angle control of a drifting automobile," Ph.D. dissertation, Stanford Univ., Stanford, CA, USA, Dec. 2022.
- [12] E. Velenis and P. Tsotras, "Minimum time vs maximum exit velocity path optimization during cornering," in *Proc. IEEE Int. Symp. Ind. Electron.*, 2005, pp. 355–360.
- [13] J. C. Huston and D. B. Johnson, "Effect of the normal force dependence of cornering stiffness on the lateral stability of recreational vehicles," *SAE Trans.*, vol. 89, pp. 999–1005, 1980.
- [14] H. Pacejka, *Tyre and Vehicle Dynamics*. Oxford, U.K.: Butterworth-Heinemann, 2002.
- [15] J. K. Subosits and J. C. Gerdes, "Impacts of model fidelity on trajectory optimization for autonomous vehicles in extreme maneuvers," *IEEE Trans. Intell. Veh.*, vol. 6, no. 3, pp. 546–558, Sep. 2021.
- [16] F. Fiala, "Seitenkrafte am rollenden luftreifen," *Z. VDI*, vol. 29, pp. 973–979, 1954.
- [17] J. A. E. Andersson, J. Gillis, G. Horn, J. B. Rawlings, and M. Diehl, "CasADi: A software framework for nonlinear optimization and optimal control," *Math. Program. Computation*, vol. 11, no. 1, pp. 1–36, 2019.
- [18] A. Wächter and L. T. Biegler, "On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming," *Math. Program.*, vol. 106, no. 1, pp. 25–57, 2006.
- [19] M. Quigley et al., "ROS: An open-source robot operating system," in *Proc. ICRA Workshop Open Source Softw.*, 2009, Art. no. 5.
- [20] M. Brown and J. C. Gerdes, "Coordinating tire forces to avoid obstacles using nonlinear model predictive control," *IEEE Trans. Intell. Veh.*, vol. 5, no. 1, pp. 21–31, Mar. 2020.
- [21] P. Falcone, M. Tufo, F. Borrelli, J. Asgari, and H. E. Tseng, "A linear time varying model predictive control approach to the integrated vehicle dynamics control problem in autonomous systems," in *Proc. IEEE 46th Conf. Decis. Control*, 2007, pp. 2980–2985.
- [22] M. Mizuno, H. Sakai, K. Oyama, and Y. Isomura, "Development of a tyre force model incorporating the influence of the tyre surface temperature," *Veh. Syst. Dyn.*, vol. 43, no. sup1, pp. 395–402, 2005.
- [23] G. Vorošević, B. Rakicevic, S. Mitić, and D. Stamenković, "Determination of cornering stiffness through integration of a mathematical model and real vehicle exploitation parameters," *FME Trans.*, vol. 41, pp. 66–71, Jan. 2013.
- [24] V. Laurence, "Integrated motion planning and control for automated vehicles up to the limits of handling," Ph.D. dissertation, Stanford Univ., Stanford, CA, USA, 2019.
- [25] J. E. Hadeef, G. Colin, Y. Chamaillard, S. Orlar, P. Rodriguez-Ayerbe, and V. Talon, "Explicit-ready nonlinear model predictive control for turbocharged spark-ignited engines," in *Proc. 7th IFAC Symp. Adv. Automot. Control*, 2013, pp. 189–194.



Trey P. Weber received the B.S. degree in mechanical engineering from the University of California, Los Angeles, Los Angeles, CA, USA in 2019, and the M.S. degree in mechanical engineering in 2021 from Stanford University, Stanford, CA, where he is currently working toward the Ph.D. degree in mechanical engineering. He works with the Dynamic Design Lab, where his research focuses on vehicle dynamics, motion planning, and control at the limits of friction.



J. Christian Gerdes (Member, IEEE) received the Ph.D. degree from the University of California at Berkeley, Berkeley, CA, USA, in 1996. He is currently a Professor in mechanical engineering with Stanford University, Stanford, CA, and the Director of the Center for Automotive Research with Stanford, Stanford University. 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. When not teaching on campus, he can often be found at the racetrack with students trying out their latest prototypes for the future.