

Autonomous Drifting with 3 Minutes of Data via Learned Tire Models

Franck Djeumou¹, Jonathan Y.M. Goh², Ufuk Topcu¹, and Avinash Balachandran²

Abstract—Near the limits of adhesion, the forces generated by a tire are nonlinear and intricately coupled. Efficient and accurate modelling in this region could improve safety, especially in emergency situations where high forces are required. To this end, we propose a novel family of tire force models based on neural ordinary differential equations and a neural-ExpTanh parameterization. These models are designed to satisfy physically insightful assumptions while also having sufficient fidelity to capture higher-order effects directly from vehicle state measurements. They are used as drop-in replacements for an analytical brush tire model in an existing nonlinear model predictive control framework. Experiments with a customized Toyota Supra show that scarce amounts of driving data – less than three minutes – is sufficient to achieve high-performance autonomous drifting on various trajectories with speeds up to 45mph. Comparisons with the benchmark model show a 4× improvement in tracking performance, smoother control inputs, and faster and more consistent computation time.

I. INTRODUCTION

Maximizing tire force usage is critical to safely negotiating highly dynamic situations, e.g., emergency obstacle avoidance. Yet, accurately predicting the effective force generated by the four tires on a car is a difficult challenge. Firstly, the tire in isolation has many complex nonlinear phenomenon, including force saturation, camber thrust, and nonlinear load dependence. Indeed, significant effort has gone into developing analytical and empirical models for a *single* tire [1]–[7], including the Magic Formula [1] which is frequently used in industry. Despite its popularity, fitting the many parameters of the Magic Formula is difficult and often requires specialized testing and facilities [2], [3].

When attached to a vehicle, the complexity compounds, as every input to these models is coupled into suspension dynamics, weight transfer, and other effects. Many control approaches in the literature thus resort to using a single-track assumption [8]–[13], where these effects are ‘lumped’ into a single tire model at the front and rear axles, and empirically fit to measured vehicle data. This includes the Fiala brush model [14], which has been experimentally demonstrated in autonomous vehicle control scenarios at the limits of handling, including emergency obstacle avoidance, drifting, and racing [15]–[20]. Although the simplicity aids control development, this single tire lumping often fails to accurately capture the intricate coupling created by higher-order effects.

Neural networks, which have universal approximation properties in the limit, could offer a solution. Black-box



Fig. 1. A photocomposite showing stills from an overhead drone video of a fully autonomous experiment superimposed at 1s intervals.

and Magic Formula-based neural network models [4]–[7], [21]–[23] have been explored in the literature. However, they do not retain physics-based guarantees, and none has been tested on a full-size car operating near or at the limits of handling. In general, their complexity has to be balanced against overfitting and computational efficiency, especially when reliable, physically insightful extrapolation is required for real-time control.

Our first contribution is to combine the physical insights of tire models with the modelling power of neural networks. We propose a novel family of tire force models based on neural ordinary differential equations (NODE) [24], [25] and neural-ExpTanh, a novel parameterization which uses curves generated by the $\exp(\cdot)\tanh(\cdot)$ function. These are designed to have high fitting fidelity while also incorporating fundamental tire modelling insights [1]–[3], including the friction ellipse constraint and ‘S-shaped’ saturation trend. The NODE model defines a differential equation whose family of solutions includes established models such as the Magic Formula [1] and Fiala brush model [14]. Through optimization-based techniques [24], [26], the model is trained to fit vehicle state measurements. To address the computational complexity of training and evaluating NODE models, which requires integrating a differential equation, we also introduce neural-ExpTanh, a subset of the NODE model’s solutions. Neural-ExpTanh can be trained efficiently and targets real-time control precisely due to its cheap function and gradient computation time.

Our second contribution provides an extensive experimental evaluation of these NODE and ExpTanh models on a full-size, heavily-modified Toyota Supra. We first compare our models to the Magic Formula and Fiala models on a dataset from the vehicle. The results show that NODE and ExpTanh satisfy the tire fundamentals while being up to 2× denser than the baselines around zero-mean prediction error.

This material is based on work supported by Toyota Research Institute.

¹F. Djeumou and U. Topcu are with the University of Texas at Austin, Austin, TX, USA. Email: {fdjeumou, utopcu}@utexas.edu

²J. Goh and A. Balachandran are with Toyota Research Institute, Los Altos, CA, USA. Email: {jonathan.goh}@tri.global

We then use these learned models as drop-in replacements for an analytical brush model in an existing nonlinear model predictive control framework [16], [27]. These are compared in autonomous drifting experiments on two different trajectories which consistently excite the nonlinear regime. Compared to the baseline, the results show improved tracking, fewer steering oscillations, and lower computation time.

The last set of experiments demonstrates data efficiency and generalization of our models. We switch to a different set of tires, collect 3 minutes of manual driving data, train an ExpTanh model in a few seconds, and then perform figure-8 autonomous drifting experiments, shown in Figure 1. The learned model shows similarly good closed-loop performance, while the performance of the baseline model drops.

II. FUNDAMENTALS

Vehicle Dynamics. The dynamics are described using a planar single-track model [8], [10]–[13], shown in Figure 2.

They are expressed in a curvilinear coordinate system [15], [18], [19], where the vehicle position is relative to a reference trajectory. The position coordinate is described by the distance s along the path, the relative heading $\Delta\phi$ with respect to a planned course ϕ_{ref} , and the lateral deviation e from the path. The motion of the state $x = [r, V, \beta, \omega_{f,r}, s, e, \Delta\phi]$ is given by $\dot{x} = M(x, u)[F_{xf}, F_{yf}, F_{xr}, F_{yr}]$, where the matrix M can be found in [15]. The states r, V, β , and $\omega_{f,r}$ denote, respectively, the yaw rate, velocity, sideslip angle, and front/rear axle wheelspeed. The components of the control $u = [\delta, \tau_{f,r}]$ represent the steering angle and the torque exerted on the front/rear axle. Further, F_{xf}, F_{yf}, F_{xr} , and F_{yr} define, respectively, the front longitudinal, front lateral, rear longitudinal, and rear lateral tire forces. In what follows, we use F_x or F_y to refer to F_{xf}, F_{xr} or F_{yf}, F_{yr} , where it is clear from the context.

Tire Force Fundamentals. Modeling the nonlinear tire forces F_x, F_y has been explored extensively. Many models consider these forces to be generated by relative ‘slip’ velocity between the tire contact patch and the road. The slip angle α , longitudinal slip ratio σ , and total slip κ as given by (1) are often used as model inputs for the tire forces:

$$\alpha_f = \arctan \frac{V \sin \beta + ar}{V \cos \beta} - \delta, \quad \alpha_r = \arctan \frac{V \sin \beta - br}{V \cos \beta}$$

$$\sigma_{f,r} = \frac{r_w \omega_{f,r} - V_{xf,xr}}{V_{xf,xr}}, \quad \kappa_{f,r} = \sqrt{\tan^2 \alpha_{f,r}^2 + \sigma_{f,r}^2} \quad (1)$$

where $V_{xf} = V \cos(\delta - \beta) - ar \sin \delta$, $V_{xr} = V \cos \beta$, r_w is the effective radius of the wheel, and a, b are the distances from the center of gravity to the front and rear axles. We follow the sign convention of $F_y \leq 0$ for nonnegative α and $F_x \geq 0$ for nonnegative σ .




Fig. 2. Single-track model of a vehicle on a reference path.




Fig. 3. The left figure shows the inflection points $\alpha_1, \alpha_0, \alpha_1$ and the changes in the convexity/concavity of F_y in the pure slip regime. The right figure shows the inflection point κ_1 of F_{tot} in the combined slip regime.

Measurements of Tire Forces. In this work, we learn tire models from the state measurements and estimates of the effective lumped axle forces. While there are different strategies for estimating these forces, for simplicity in this paper, we consider conditions where we can assume $F_{xf} = 0$, e.g., no torque on the front wheels. Then, we compute $\dot{r}, \dot{V}, \dot{\beta}$ from measured states and invert through the matrix M to obtain estimated forces, indicated by $\bar{F}$ [15].

III. PHYSICS-INFORMED LEARNED TIRE FORCES

In this section, we describe our physics-based, neural ordinary differential equation (NODE) model and the derived ExpTanh parameterization. In keeping with tire modelling convention, we divide the discussion into pure slip and combined slip regimes. In pure slip, the tire is only creating force along one axis ($\sigma = 0$ or $\alpha = 0$); in combined slip, the tire is creating both longitudinal and lateral force ($\sigma \neq 0$ and $\alpha \neq 0$). From tire fundamentals [1]–[3], we expect the following generalized behavior, summarized in Figure 3.

Characteristic ‘S-shape’ curve. As the absolute value of the input slip increases, the tire force magnitude also increases until a peak force is attained and the tire contact patch starts to slide. Beyond this point, the force decreases, following an ‘S-shape’ curve. In the pure slip regime, the input slip is σ or α . In the combined slip regime, the input slip is some combination of σ and α , and the output is the total force magnitude $F_{\text{tot}} = \sqrt{F_x^2 + F_y^2}$.

Combined slip regime. For combined slip, the components of F_{tot} are distributed according to some ratio of the slip angle and longitudinal slip vs. the combined slip. An example is schematically shown in Figure 3, for fixed σ and $|\alpha|$ increasing from 0: The proportion of longitudinal force decreases while the lateral force increases until saturation.

Friction limits. In both regimes, the peak force is constrained by the maximum available tire/road adhesion capability, μF_z , with μ the friction coefficient and F_z the normal load on the tires. μF_z is difficult to know precisely as it depends on the surface, tire orientation, and the normal load – which in turn vary with the vehicle’s state due to weight transfer/suspension dynamics. Yet, this notion of a maximum force greatly eases analysis for control and safety.

Throughout this paper, we assume a given set of measurements $\mathcal{D} = \{(\alpha_{f,r}, \sigma_{f,r}, r, V, \beta, \omega_{f,r}, \bar{F}_x, \bar{F}_y, \mu \bar{F}_z)\}_{i=1}^N$, where $\mu \bar{F}_z = \bar{\mu} mg$ is a rough estimate of the nominal load and $\bar{\mu}$ encodes any available approximate knowledge on μ .

A. Physics-Informed NODE for Tire Force Modeling

We seek a generalizable model that satisfies these physical insights. First, for the characteristic ‘S-shape’, instead of intuiting a curve, e.g., Magic Formula, that satisfies the ‘S-shape’, we characterize the family of physically-feasible curves using notions of convexity, concavity, and inflection points. Then, we optimize for the function in this family that best fits the data via stochastic gradient descent.

Specifically, at critical inflection points (Figure 3), the curve changes convexity or concavity. In the pure slip regime, F_y contains *three* inflection points $\alpha_{-1}, \alpha_0, \alpha_1$. We seek a family of curves such that $F_y(\alpha)$ is convex for all $\alpha \leq \alpha_{-1}$ and $\alpha \in [\alpha_0, \alpha_1]$, and $F_y(\alpha)$ is concave otherwise, for some $\alpha_{-1}, \alpha_0, \alpha_1$. We have the same properties for $-F_x(\sigma)$, for some $\sigma_{-1}, \sigma_0, \sigma_1$. In the combined slip regime, the family of curves for F_{tot} contains an inflection point κ_1 such that $\forall \kappa \leq \kappa_1$, $F_{\text{tot}}(\kappa)$ is concave, and convex otherwise. Convexity and concavity correspond to nonnegative and nonpositive second-order derivatives, respectively. Thus, the main idea is to learn the inflection points and the second derivative of the tire forces with respect to the corresponding slips while enforcing the desired convexity/concavity properties; the forces are then obtained by integration. Further, we enforce soft constraints on the peak force as required by the friction limits.

Pure Slip NODE Model. The lateral force F_y is a solution of the second-order differential equation given by

$$\begin{aligned} \dot{F}_y^\theta &= G_y^\theta, \quad z = [\alpha, F_y^\theta, G_y^\theta, \alpha_{-1}^\theta, \alpha_0^\theta, \alpha_1^\theta] \\ \dot{G}_y^\theta &= \begin{cases} \exp\{\text{NN}_1^\theta(z, \text{feat})\} & \text{if } \alpha \leq \alpha_{-1}^\theta \text{ or } \alpha \in [\alpha_0^\theta, \alpha_1^\theta] \\ -\exp\{\text{NN}_2^\theta(z, \text{feat})\} & \text{otherwise} \end{cases} \end{aligned} \quad (2)$$

where the derivative here is taken with respect to α . The set of features used for learning are $\text{feat} = [r, V, \beta, \mu\bar{F}_z]$ for the front axle and $\text{feat} = [r, V, \mu\bar{F}_z]$ for the rear. We select the feature set feat such that for fixed feat , α given in (1) is not uniquely defined. NN_x^θ denotes a neural network, where θ is the set of all parameters for the model. The inflection points are parameterized as $[\alpha_{-1}^\theta, \alpha_0^\theta, \alpha_1^\theta, F_0^\theta, G_0^\theta] = \text{NN}_3^\theta(\text{feat})$ with F_0^θ, G_0^θ being the initial states to use when integrating the differential equation. Note that choosing the exp function in $\dot{G}_y^\theta$ enforces the nonnegative and nonpositive second-order derivatives constraints. We then compute the parameters θ by solving the following optimization problem

$$\begin{aligned} \min_{\theta} \frac{1}{N} \sum_{\substack{\alpha, r, V, \beta, \\ \bar{F}_y, \mu\bar{F}_z \in \mathcal{D}}} \left(\underbrace{\text{ode}((2), [\alpha_0^\theta, \alpha], [F_0^\theta, G_0^\theta])}_{[F_y^\theta, G_y^\theta]} - \bar{F}_y \right)^2 \\ + \lambda(\min\{\mu\bar{F}_z - |F_y^\theta|, 0\})^2 \end{aligned} \quad (3)$$

where ode , an integration scheme, solves (2) from α_0^θ to the measured α with the initial condition given by $[F_0^\theta, G_0^\theta]$. The term $\lambda(\min\{\mu\bar{F}_z - |F_y^\theta|, 0\})^2$ enforces the friction limits knowledge as a soft constraint by penalizing values that exceed the estimated nominal load $\mu\bar{F}_z$. The hyperparameter λ specifies the confidence in $\mu\bar{F}_z$: Low values enable the peak force to be adjusted according to the data while high values constrain the peak force to be less than $\mu\bar{F}_z$.

Combined Slip NODE Model. The total force F_{tot} is a solution of the second-order differential equation given by

$$\begin{aligned} \dot{F}_{\text{tot}}^\theta &= G_{\text{tot}}^\theta, \quad z = [\kappa, F_{\text{tot}}^\theta, G_{\text{tot}}^\theta, \kappa_1^\theta] \\ \dot{G}_{\text{tot}}^\theta &= \begin{cases} -\exp\{\text{NN}_1^\theta(z, \text{feat})\} & \text{if } \kappa \leq \kappa_1^\theta \\ \exp\{\text{NN}_2^\theta(z, \text{feat})\} & \text{otherwise} \end{cases} \end{aligned} \quad (4)$$

where the derivative is with respect to the combined slip κ , $[\kappa_1^\theta, F_0^\theta, G_0^\theta] = \text{NN}_3^\theta(\text{feat})$, and the features feat are again picked such that α and σ are not uniquely defined. Then, to learn the component distribution of this total force, we define $[s_1^\theta, s_2^\theta] = \text{NN}_4^\theta(\alpha, \sigma)$ and estimate F_y^θ and F_x^θ by scaling F_{tot}^θ as follows

$$F_y^\theta = \frac{s_1^\theta F_{\text{tot}}^\theta}{\sqrt{(s_1^\theta)^2 + (s_2^\theta)^2}}, \quad F_x^\theta = \frac{s_2^\theta F_{\text{tot}}^\theta}{\sqrt{(s_1^\theta)^2 + (s_2^\theta)^2}} \quad (5)$$

Thus, the parameters θ are obtained by solving the problem

$$\begin{aligned} \min_{\theta} \frac{1}{N} \sum_{\substack{\alpha, \sigma, r, V, \beta, \bar{F}_y, \\ \bar{F}_x, \mu\bar{F}_z \in \mathcal{D}}} \left(\underbrace{\text{ode}((4), [\kappa_0^\theta, \kappa], [F_0^\theta, G_0^\theta])}_{[F_{\text{tot}}^\theta, G_{\text{tot}}^\theta]} - \bar{F}_{\text{tot}} \right)^2 \\ + (F_y^\theta - \bar{F}_y)^2 + (F_x^\theta - \bar{F}_x)^2 + \lambda(\min\{\mu\bar{F}_z - F_{\text{tot}}^\theta, 0\})^2 \end{aligned}$$

where the measured total force is $\bar{F}_{\text{tot}} = \sqrt{\bar{F}_x^2 + \bar{F}_y^2}$.

Remark 1: In the pure slip regime, learning $-F_x^\theta$ follows exactly the description of F_y^θ with α replaced by σ . Despite the rich class of functions encoded by the NODE formulation, solving (2) and (4) to estimate the forces slows down training and hinders the direct application of the formulation for control. In practice, we address this issue by first learning the parameters θ , then training a new neural network to mimic the solutions of (2) and (4) via overfitting. Thus, evaluating the obtained neural network and its Jacobian becomes computationally cheap for real-time control.

B. ExpTanh: A New Family of Tire Models

We restrict the NODE model’s set of solutions to a family of functions, namely ExpTanh , satisfying the second-order derivative condition without the need to integrate a differential equation. ExpTanh curves are given by

$$\text{ExpTanh}^\theta(z) = a_0^\theta + (a_1^\theta + a_2^\theta e^{-a_3^\theta |z|}) \tanh(a_4^\theta(z - a_5^\theta))$$

where a_k^θ are constants or neural network functions such that $a_1^\theta, a_2^\theta, a_3^\theta \geq 0$. Importantly, the maximum/minimum values z_+^θ, z_-^θ can be found analytically:

$$z_\pm^\theta = a_5^\theta \pm a_4^\theta \text{atanh}\left(\frac{\sqrt{(a_2^\theta a_3^\theta)^2 + 4(a_4^\theta)^2} - a_2^\theta a_3^\theta}{2a_4^\theta}\right) \quad (6)$$

ExpTanh Pure Slip. We model F_y as $F_y^\theta(\alpha, \text{feat}) = \text{ExpTanh}^\theta(\alpha)$, where $(a_i^\theta)_{i=0}^5 = \text{NN}^\theta(\text{feat})$, feat are the same features as in the NODE version, and θ is the set of all parameters. In practice, we pass $a_1^\theta, a_2^\theta, a_3^\theta$ through an exponential function to enforce nonnegative values. The optimum parameters θ are given by

$$\min_{\theta} \frac{1}{N} \sum_{\substack{\alpha, r, V, \beta, \\ \bar{F}_y, \mu\bar{F}_z \in \mathcal{D}}} \left(F_y^\theta - \bar{F}_y \right)^2 + \lambda(\mu\bar{F}_z - |F_y^\theta(\alpha_\pm^\theta, \text{feat})|)^2$$

where the second term is a similar soft penalty on exceeding the estimated maximum friction force.

ExpTanh Combined Slip. We model the total force as $F_{\text{tot}}^\theta(\kappa, \text{feat}) = \text{ExpTanh}^\theta(\kappa)$, where $(a_i^\theta)_{i=0}^5 = \text{NN}^\theta(\text{feat})$ and feat represents the same set of features as in the combined slip NODE model. The forces F_y^θ and F_x^θ depend on F_{tot}^θ as given by (5), where the functions s_1^θ and s_2^θ are to be learned. Specifically, we compute θ by solving

$$\min_{\theta} \frac{1}{N} \sum_{\substack{\alpha, \sigma, r, V, \beta, \bar{F}_y, \\ \bar{F}_x, \mu \bar{F}_z \in \mathcal{D}}} \left(F_{\text{tot}}^\theta(\kappa, \text{feat}) - \bar{F}_{\text{tot}} \right)^2 + (F_y^\theta - \bar{F}_y)^2 + (F_x^\theta - \bar{F}_x)^2 + \lambda (\mu \bar{F}_z - F_{\text{tot}}^\theta(\kappa_+, \text{feat}))^2$$

Remark 2: Firstly, by incorporating selected subsets of the measured states, feat , in addition to the slip values, the proposed models are able to capture the intricate coupling between the effective lumped tire force curves and vehicle motion. While we made one choice for feat , other selections are likely suitable, depending on the vehicle. Secondly, for fixed parameters θ , ExpTanh requires only two evaluations of the function $\exp$, which is computationally cheap compared to the Magic Formula requiring three evaluations of $\arctan$; the gradient is also easier to compute.

IV. EXPERIMENTS

We demonstrate the data efficiency, prediction accuracy, and computational efficiency of our tire models through several experiments: Comparisons to the Magic Formula and Fiala brush on the testbed vehicle dataset, and autonomous drifting on slalom and figure-8 trajectories.

The experiments in this section were performed on the Toyota Supra described in [16], [27] and heavily modified for high-performance autonomous driving. Vehicle state measurements are obtained from a commercial RTK GPS-INS unit at a rate of 250Hz. As a rear-wheel drive vehicle, the front tires operate in the pure slip regime with $F_{xf} = 0$, and the rear tires operate in the combined slip regime. We assume standard units for all quantities when they are not specified.

A. Evaluation of the Learned Tire Models

To compare the tire models, we used a dataset $\mathcal{D}$ (Figure 4) of manual and autonomous driving/drifting. The dataset contains 306887 state measurements at 100 Hz, totaling ~ 1 hour accumulated over the span of three months on the same surfaces under similar summer weather conditions.

For the NODE model, NN_1^θ , NN_2^θ , and NN_4^θ have 2 hidden layers with 16 nodes per layer, while NN_3^θ have 4 nodes per layer. For the ExpTanh model, NN^θ and NN_4^θ have 2 hidden layers with 3 nodes per layer. All neurons used $\tanh$ as activation function. We used $\lambda = 0.01$ to express low confidence in the estimated $\mu \bar{F}_z = 7000$. We trained the models via Adam optimizer [28], where the learning rate is set to decay exponentially with a rate of 0.01 and an initial value of 0.001. On a laptop with GeForce RTX 2060, training both the pure slip F_{yf} and the combined slip F_{xr} and F_{yr} models took ~ 27 minutes for NODE, and only ~ 4 minutes for ExpTanh .

We compare our models with the Magic Formula and Fiala model. The parameters of the Magic Formula (Chapter 4, Section 4.3.2 of [1]) were obtained by optimizing a mean-square-error loss over the dataset. The Fiala model parameterization was empirically tuned by usage in the existing autonomous drifting NMPC framework [15], [16]. Figure 4 summarizes our findings: Our tire models significantly improve prediction accuracy over the Magic Formula and Fiala while satisfying the tire fundamentals. The NODE model provides the best prediction accuracy while taking significantly more time to train and evaluate. In contrast, ExpTanh achieves slightly lower prediction accuracy compared to the NODE model while being easy to train and evaluate.

Figure 5 shows how the learned ExpTanh model efficiently captures the coupling with the vehicle states r, V, β . This suggests that our models not only fit the tires but can also incorporate complex chassis interactions (e.g., weight transfer and suspension dynamics). In the pure slip model, r shifts the center of the curve. This could be due to significant static and dynamic camber from the test vehicle's aggressive




Fig. 4. Comparison of the different tire models trained and tested on a real-world driving dataset. The first row shows the density distribution of the prediction error, and the second row shows the forces as a function of the slip values for a fixed state $r, V, \beta = 0.7, 20, 0.1$, where $F_{yr}(\alpha_r)$ is obtained for fixed $\sigma_r = 0.4$ and $F_{xr}(\sigma_r)$ is obtained for fixed $\alpha_r = 0.02$. In the density plot, NODE and ExpTanh are at least $1.5\times$ denser around zero-mean error than expert-designed Fiala and Magic Formula (MF). The second row validates that the learned NODE and ExpTanh enforce the tire fundamentals.

drift-specific front suspension setup. Low speed values tend to flatten the curve, while the slip angle corresponding to the peak force decreases with increasing sideslip angle β . For the combined slip model, the dependency on r, V is most significant in the nonlinear transitional region at low longitudinal slip. As expected from the tire fundamentals, the magnitude of F_{xr} decreases as the slip angle increases for fixed σ_r .




Fig. 5. Impact of the states r, V, β on the learned ExpTanh model. For the front tire, the blue curve corresponds to fixed $r = 1.8, \beta = 0.9$, and V ranging from 5 to 20 with lower values represented by lighter colors. The green curve follows the blue curve but with $r = 0$. The orange curve uses fixed $r = -1.8, V = 12$, and β varying from -0.9 to 0.9 . For the rear tire, V ranges from 5 to 20 with fixed $r = -1.8$ for the blue curve, r ranges from -1.8 to 1.8 with fixed $V = 12$ for the green curve, and both V and r vary for the orange curve.

B. Autonomous Drifting with Learned Tire Models

To evaluate their practical closed-loop performance, we use our learned models in Figure 4 as direct drop-in replacements for a Fiala model in an existing closed-loop NMPC framework for autonomous drifting [16], [27]. The reference trajectories were pre-computed via nonlinear optimization with the benchmark Fiala model. The NMPC cost function primarily penalizes the lateral error e , the deviation from

the reference sideslip angle $e_\beta = \beta - \beta_{ref}$, and the relative deviation $\Delta\phi$. The sideslip β_{ref} enforces the drifting profile.

1) *Drifting on Slalom Trajectory:* For the first experiment, we compare the closed-loop performance of the benchmark Fiala, NODE, and ExpTanh models on a transient slalom trajectory (Figure 6). The integrated NODE formulation was approximated with a neural network trained on its output. The slalom trajectory has corners with reference sideslip angle of up to 43° and velocity between 31mph and 45mph. Figure 6 demonstrates improved tracking performance: In terms of root mean squared error, ExpTanh tracks the path (e and $\Delta\phi$) up to $3.5\times$ better than the Fiala model while achieving up to $1.5\times$ better sideslip tracking performance. The NODE model achieves slightly lower performance than ExpTanh, possibly due to the loss of accuracy from the approximation procedure. Importantly, we also note fewer steering oscillations δ when ExpTanh and NODE are used, as compared to the baseline Fiala model.

2) *Autonomous Drifting with 3 Minutes of Data:* This set of experiments investigates the generalizability of the ExpTanh model. First, we perform experiments on a figure-8 trajectory with both the benchmark Fiala model and the ExpTanh model. We then changed the rear tires from Toyo Proxes Sport 275/35R18, which we used for all previous tests, to Bridgestone Potenza Sport 275/35R18. A safety driver then manually drove the car on the skidpad, with unstructured grip and drift maneuvers, for ~ 3 minutes. This data was then used to train an ExpTanh model; this took $< 15s$ on a laptop with a GeForce RTX 2060. This freshly fitted model was then again compared to the benchmark model on the same figure-8 trajectory.

Figure 7 summarizes our findings. For the original Toyo tires, performance was significantly better with the ExpTanh model than the baseline Fiala model. Importantly, the closed-loop behavior with the ExpTanh model was similar after switching to the new Bridgestone tires and retraining the network. In contrast, the performance with




Fig. 6. Drifting on a slalom figure. Our approaches show better accuracy at trajectory tracking and fewer steering oscillations than Fiala model.




Fig. 7. Drifting on a figure-8 trajectory with 3 minutes of data. ExpTanh shows better tracking performance with both tires, especially in transitional regions. The red line indicates where the *Bridgestone* + Fiala test was ended due to the safety driver feeling uncomfortable with the tracking error.

the unchanged Fiala model significantly degraded, showing that there was indeed a notable difference in the behavior of the tires that the ExpTanh model successfully adapted to with sparse data. This is also reflected in the root mean squared e and e_β values: compared to the baseline, we see a $> 4\times$ improvement in lateral error, and $> 2\times$ improvement in sideslip tracking for both tires.

RMS	Toyo		Bridgestone	
	Fiala	ExpTanh	Fiala	ExpTanh
e (m)	1.77	0.40	2.02	0.27
e_β (rad)	0.013	0.006	0.011	0.004

With ExpTanh, the controller tracks the sideslip β reference with less overshoot and less steering oscillation compared to the Fiala model. This difference is particularly pronounced at the end of each transition, as shown in the zoomed section on the sideslip evolution, where we expect to see more complex interactions between the vehicle states due to transient load transfer and high yaw rates. This suggests not only that the ExpTanh model is able to capture these effects but also that the controller benefits by exploiting this in closed loop. In contrast, the controller with the Fiala model tends to overshoot severely during these transitions.

Another region with complex coupling is the slow transition from drifting to grip driving at the end of the experiment (e.g. $s \in [600, 660]$). Here, the baseline *Toyo* + Fiala model combination exhibits steering and sideslip oscillations. In contrast, with both tires, ExpTanh smoothly tracks sideslip, and has better lateral error performance.

Figure 8 compares the observed optimal control problem computation time during the *Toyo* experiments. Due to its simplicity, forward evaluation of the Fiala model is likely faster than ExpTanh. However, in NMPC, the fidelity of the tire model and the smoothness of its Jacobian are important. Figure 8 shows that NMPC with Fiala often needs more gradient iterations than ExpTanh to converge to a solution.

This is exacerbated in the transitional regime, $s \geq 600$, where the number of iterations triples with the Fiala model, but remains similar with ExpTanh. This suggests that NMPC with ExpTanh can be both faster and more consistent.




Fig. 8. Compute time and number of gradient iterations from the controller.

V. CONCLUSION

We propose a family of tire force models based on neural ordinary differential equations (NODE) and ExpTanh. These models combine physics-based tire modelling fundamentals with the ability to directly learn, using onboard sensor measurements, higher-order effects from the interaction between the tires, suspension, and road. Autonomous drifting experiments, which subject the model to extreme conditions, demonstrate improved tracking performance, optimized computation time, and unprecedented data efficiency: Learning with only 3 minutes of driving. Finally, our rapid training time (usually a few seconds) suggests that future work could explore using these models in a life-long learning setting, where the tire curves are updated online during driving.

Acknowledgements. The authors would like to thank our colleagues at Toyota and UT Austin's Autonomous Systems Group for their generous support. We thank ONR, AFOSR, and NSF for their support on our prior works on NODE.

REFERENCES

- [1] H. B. Pacejka, *Tire and Vehicle Dynamics*. Butterworth-Heinemann, third edition ed., 2012.
- [2] J. Svendenius and M. Gäfvert, "Construction of novel semi-empirical tire models for combined braking and cornering," *Technical Reports; TFR*, vol. 7606, 2003.
- [3] J. Svendenius, *Tire modeling and friction estimation*. PhD thesis, Lund University, 2007.
- [4] M. Acosta and S. Kanarachos, "Tire lateral force estimation and grip potential identification using neural networks, extended kalman filter, and recursive least squares," *Neural Computing and Applications*, vol. 30, pp. 3445–3465, 2017.
- [5] P. Guarneri, G. Rocca, and M. Gobbi, "A neural-network-based model for the dynamic simulation of the tire/suspension system while traversing road irregularities," *IEEE Transactions on Neural Networks*, vol. 19, pp. 1549–1563, 2008.
- [6] N. Xu, H. Askari, Y. Huang, J. Zhou, and A. Khajepour, "Tire force estimation in intelligent tires using machine learning," *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, pp. 3565–3574, 2022.
- [7] J. Matuško, I. Petrović, and N. Perić, "Neural network based tire/road friction force estimation," *Eng. Appl. Artif. Intell.*, vol. 21, pp. 442–456, 2008.
- [8] B. Paden, M. Cáp, S. Z. Yong, D. S. Yershov, and E. Frazzoli, "A survey of motion planning and control techniques for self-driving urban vehicles," *IEEE Transactions on Intelligent Vehicles*, vol. 1, pp. 33–55, 2016.
- [9] C. You and P. Tsiotras, "Vehicle modeling and parameter estimation using adaptive limited memory joint-state ukf," *2017 American Control Conference (ACC)*, pp. 322–327, 2017.
- [10] R. Rajamani, *Vehicle dynamics and control*. Springer Science & Business Media, 2011.
- [11] P. Falcone, F. Borrelli, J. Asgari, H. E. Tseng, and D. Hrovat, "Predictive active steering control for autonomous vehicle systems," *IEEE Transactions on Control Systems Technology*, vol. 15, pp. 566–580, 2007.
- [12] J. Kong, M. Pfeiffer, G. Schilbach, and F. Borrelli, "Kinematic and dynamic vehicle models for autonomous driving control design," *2015 IEEE Intelligent Vehicles Symposium (IV)*, pp. 1094–1099, 2015.
- [13] P. Polack, F. Althé, B. d'Andréa Novel, and A. de La Fortelle, "The kinematic bicycle model: A consistent model for planning feasible trajectories for autonomous vehicles?," *2017 IEEE Intelligent Vehicles Symposium (IV)*, pp. 812–818, 2017.
- [14] E. Fiala, "Seitenkräften am rollenden luftreifen," *VdI*, vol. 96, pp. 973–979, 1954.
- [15] J. Y. Goh, T. Goel, and J. C. Gerdes, "Toward automated vehicle control beyond the stability limits: Drifting along a general path," *Journal of Dynamic Systems Measurement and Control-transactions of The Asme*, vol. 142, 2019.
- [16] A. Balachandran, T. L. Chen, J. Y. Goh, S. McGill, G. Rosman, S. Stent, and J. J. Leonard, "Human-centric intelligent driving: Collaborating with the driver to improve safety," in *Automated Road Transportation Symposium*, pp. 85–109, Springer, 2023.
- [17] R. Hindiyeh and J. C. Gerdes, "A controller framework for autonomous drifting: Design, stability, and experimental validation," *Journal of Dynamic Systems Measurement and Control-transactions of The Asme*, vol. 136, p. 051015, 2011.
- [18] J. K. Subosits and J. C. Gerdes, "Impacts of model fidelity on trajectory optimization for autonomous vehicles in extreme maneuvers," *IEEE Transactions on Intelligent Vehicles*, vol. 6, pp. 546–558, 2021.
- [19] J. Y. Goh and J. C. Gerdes, "Simultaneous stabilization and tracking of basic automobile drifting trajectories," *2016 IEEE Intelligent Vehicles Symposium (IV)*, pp. 597–602, 2016.
- [20] J. Y. M. Goh, *Automated vehicle control beyond the stability limits*. Stanford University, 2019.
- [21] S. Srinivasan, I. Sa, A. Zyner, V. Reijgwart, M. de la Iglesia Valls, and R. Y. Siegart, "End-to-end velocity estimation for autonomous racing," *IEEE Robotics and Automation Letters*, vol. 5, pp. 6869–6875, 2020.
- [22] N. A. Spielberg, M. Brown, N. R. Kapania, J. C. Kegelmann, and J. C. Gerdes, "Neural network vehicle models for high-performance automated driving," *Science Robotics*, vol. 4, 2019.
- [23] G. Williams, P. Drews, B. Goldfain, J. M. Rehg, and E. A. Theodorou, "Information-theoretic model predictive control: Theory and applications to autonomous driving," *IEEE Transactions on Robotics*, vol. 34, pp. 1603–1622, 2017.
- [24] T. Q. Chen, Y. Rubanova, J. Bettencourt, and D. K. Duvenaud, "Neural ordinary differential equations," in *NeurIPS*, 2018.
- [25] F. Djeumou, C. Neary, É. Goubault, S. Putot, and U. Topcu, "Neural networks with physics-informed architectures and constraints for dynamical systems modeling," in *LADC*, 2022.
- [26] F. Djeumou, C. Neary, É. Goubault, S. Putot, and U. Topcu, "Taylor-lagrange neural ordinary differential equations: Toward fast training and evaluation of neural odes," in *IJCAI*, 2022.
- [27] J. Y. Goh, M. Thompson, J. Dallas, and A. Balachandran, "Nonlinear model predictive control for highly transient autonomous drifting," 15th International Symposium on Advanced Vehicle Control, 2022.
- [28] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.