

# Path Tracking of Highly Dynamic Autonomous Vehicle Trajectories via Iterative Learning Control

Nitin R. Kapania<sup>1,2</sup> and J. Christian Gerdes<sup>1</sup>

**Abstract**—Autonomous vehicles will benefit from the ability to perform aggressive driving maneuvers in safety-critical situations where the full use of available tire-road friction is required. Unfortunately, vehicle steering dynamics become highly nonlinear and difficult to model near the limits of tire adhesion, making accurate control of these maneuvers difficult. One promising approach is to use iterative learning control (ILC) as a method of gradually determining the proper steering input for a transient driving maneuver by repeating the maneuver several times and using information from previous iterations to improve the reference tracking performance. This paper explores the viability of this concept by applying learning algorithms for multiple-lap path tracking of an autonomous race vehicle. Racing is an ideal scenario for iterative learning control because race cars drive the same sequence of turns while operating near the physical limits of tire-road friction. This creates a repeatable set of nonlinear vehicle dynamics and road conditions from lap to lap. Simulation results are used to design and test convergence of both a proportional-derivative (PD) and quadratically optimal (Q-ILC) iterative learning controller, and experimental results are presented on an Audi TTS race vehicle driving several laps around Thunderhill Raceway in Willows, CA at combined vehicle accelerations of up to  $8 \text{ m/s}^2$ . Both control algorithms are able to correct transient path tracking errors and significantly improve the performance provided by a reference feedforward controller.

## I. INTRODUCTION

Recent advances in perception, drive-by-wire, and control technology will enable autonomous vehicles to improve passenger safety by intervening in situations where quick, highly dynamic maneuvers beyond the capability of the typical human driver are required. Examples of these situations include obstacle avoidance maneuvers [1] and scenarios where the tire-road friction limit has been temporarily exceeded due to ambient driving conditions or human error [2]. A major difficulty in performing aggressive driving maneuvers in safety-critical situations is that at the limits of tire adhesion, the steering dynamics associated with lateral tracking of the desired path become highly nonlinear. Furthermore, difficult-to-measure disturbances such as bank, grade and local friction variation of the road surface have a large effect on the transient dynamics of the vehicle [3].

Iterative learning control (ILC) is a promising method for determining a steering control input for path tracking of highly dynamic vehicle trajectories. Iterative learning has traditionally been used in controlled environments to improve the tracking performance of systems that perform the same task in repeatable conditions. By using error information

from prior attempts, iterative learning controllers gradually determine control inputs that cause the system output to accurately track the repeated reference trajectory. Recent research has expanded the use of ILC for autonomous ground and aerial vehicle applications. Chen and Moore [4] proposed a simple ILC scheme in 2006 to improve path-following of a ground vehicle with omni-directional wheels. In 2013, Sun et al. [5] proposed an ILC approach for overspeed protection of high-speed trains. Additionally, Purwin and D'Andrea [6] synthesized an iterative learning controller using a least-squares approach to perform aggressive transient maneuvers of an unmanned aerial vehicle (UAV).

This paper applies iterative steering control algorithms for the specific case of an autonomous race vehicle. Racing provides an ideal scenario for initial automotive learning control studies because race cars are driven multiple laps around the same sequence of turns on a closed track, while operating near the physical limits of tire-road friction to minimize lap times. Because the reference road curvature is unchanging from lap-to-lap, the unknown transient disturbances and vehicle dynamics tend to be repeatable and can therefore be accounted for via iterative learning control. A notable exception occurs when the vehicle has significantly understeered due to lack of front tire force availability. In this case, additional steering will have no impact on path tracking, and the desired vehicle velocity and/or cornering radius must be made less aggressive.

The paper is divided as follows. Section II introduces a linear model for the planar dynamics of a race car following a fixed reference path. Because the transfer function between the steering wheel input and the vehicle's path deviation is open-loop unstable, a stabilizing lanekeeping controller is added to the steering system and the closed loop dynamics are represented in the commonly used "lifted domain". Section III presents a PD-type iterative learning controller with a low-pass filter used to speed up convergence. Gain tuning and stability at low lateral accelerations are shown using lifted domain techniques, while nonlinear simulations are used to predict a desirable tracking response in the presence of high lateral acceleration. Section IV presents a quadratically optimal (Q-ILC) iterative learning controller, which has the benefit of explicitly accounting for changes in vehicle speed along the race track. Section V presents experimental data of both controllers implemented on an Audi TTS race vehicle at combined lateral/longitudinal accelerations of up to  $8 \text{ m/s}^2$ .

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

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

## II. VEHICLE DYNAMICS AND PROBLEM OVERVIEW

For this paper, the objective of piloting an autonomous vehicle along a fixed race track in minimum time is divided into separate lateral and longitudinal vehicle control problems. The lateral controller uses the steering wheel input to track a desired “racing-line”, shown in Fig. 1a. The racing line is frequently represented by a path curvature function  $\kappa$  parameterized by distance along the track (Fig. 1b). The longitudinal controller tracks a desired speed profile that keeps the vehicle at a specified lateral-longitudinal acceleration magnitude, typically near the limits of tire-road friction (Fig. 1c). For the Thunderhill Raceway circuit shown in Fig. 1a, minimum-time velocity and curvature profiles are computed using methods published in [7]. Additionally, tracking of the velocity profile is handled by a longitudinal controller [8] that is not affected by the iterative learning controller.

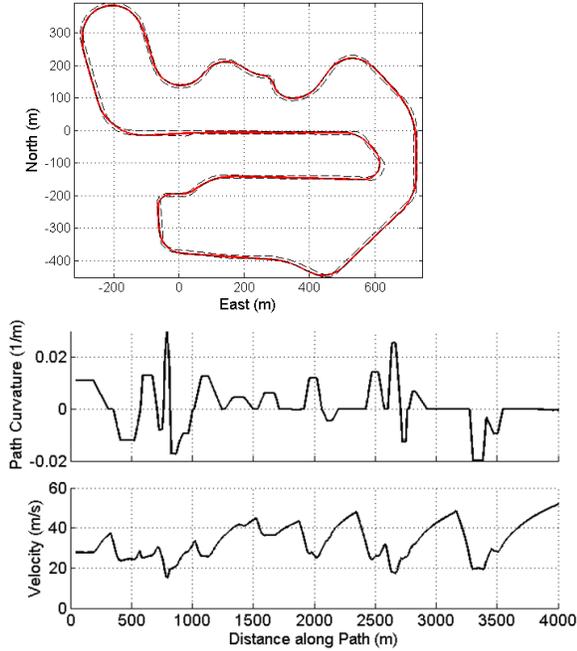


Fig. 1. (a) Overhead plot of racing line for Thunderhill Raceway, Willows CA, USA. (b) Racing line represented by curvature along distance traveled. (c) Velocity profile to keep vehicle at combined lateral/longitudinal acceleration of  $8 \text{ m/s}^2$ .

### A. Lateral Vehicle Dynamics

Fig. 2a shows a schematic of a vehicle following a path with curvature profile  $\kappa$ . The lateral deviation of the vehicle from the desired path  $e$  is the measured output for the ILC algorithms. Since iterative learning controllers determine appropriate *feedforward* control inputs over several iterations, it is necessary for the open-loop dynamics between the control input and control output to be asymptotically stable. For the case of active steering control, the transfer function between the vehicle steer angle  $\delta$  and the lateral path error  $e$  is characterized by two poles at the origin, requiring the addition of a stabilizing feedback controller.

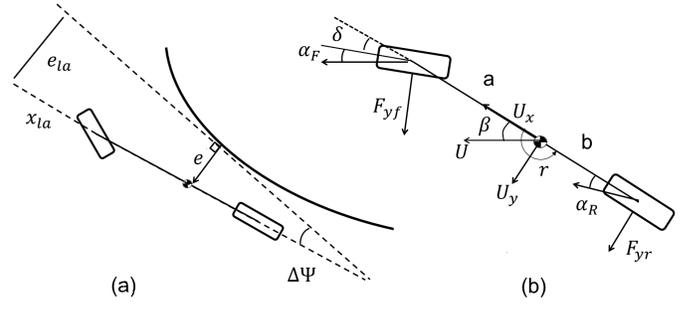


Fig. 2. (a) Diagram showing tracking error  $e$ , heading error  $\Delta\Psi$ , lookahead distance  $x_{LA}$ , and lookahead error  $e_{LA}$ . (b) Schematic of bicycle model.

A lookahead controller provides the stabilizing feedback command for this paper. The “lookahead error” is defined by:

$$e_{LA} = e + x_{LA}\Delta\Psi \quad (1)$$

where  $\Delta\Psi$  is the vehicle heading error and  $x_{LA}$  is a lookahead distance projected in front of the vehicle, typically 5-20 meters for autonomous driving (Fig. 2a). The resulting feedback control law is

$$\delta_{FB} = -k_P e_{LA} \quad (2)$$

with proportional gain  $k_P$ . The control law (2) is a natural extension of potential field lanekeeping, as described by Rossetter et al., which also provides heuristics for selecting  $k_P$  and  $x_{LA}$  [9]. Desirable stability properties over significant tire saturation levels are demonstrated in [10].

With the feedback controller added, closed loop dynamics of the lateral path deviation are dependent on three other states: vehicle sideslip  $\beta$ , yaw rate  $r$  and heading error  $\Delta\Psi$ . For controller development and testing, these dynamics are given by the planar bicycle model:

$$\dot{\beta} = \frac{F_{yf} + F_{yr}}{mU_x} - r \quad \dot{r} = \frac{aF_{yf} - bF_{yr}}{I_z} \quad (3a)$$

$$\dot{e} = U_x(\beta + \Delta\Psi) \quad \dot{\Delta\Psi} = r - U_x\kappa \quad (3b)$$

Where  $U_x$  is the vehicle forward velocity and  $F_{yf}$  and  $F_{yr}$  the front and rear lateral tire forces. The vehicle mass and yaw inertia are denoted by  $m$  and  $I_z$ , and the geometric parameters  $a$  and  $b$  are shown in Fig. 2b.

As automotive racing frequently occurs near the limits of tire force saturation, lateral tire force is modeled using the nonlinear Fiala brush tire model, assuming a single coefficient of friction and a parabolic force distribution [11]. The front and rear lateral tire forces  $F_{yf}$  and  $F_{yr}$  are functions of the front and rear tire slip angles  $\alpha_f$  and  $\alpha_r$  respectively, as shown in (4).

$$C = [1 \ 0 \ 0 \ 0] \quad (10)$$

$$F_{y\star} = \begin{cases} -C_\star \tan \alpha_\star + \frac{C_\star^2}{3\mu F_{z\star}} |\tan \alpha_\star| \tan \alpha_\star \\ -\frac{C_\star^3}{27\mu^2 F_{z\star}^2} \tan^3 \alpha_\star, & |\alpha_\star| < \arctan\left(\frac{3\mu F_{z\star}}{C_\star}\right) \\ -\mu F_{z\star} \operatorname{sgn} \alpha_\star, & \text{otherwise} \end{cases} \quad (4)$$

The disturbance  $d(t)$  is assumed constant from lap to lap and is given by

$$d(t) = [0 \ -\kappa U_x(t) \ 0 \ 0]^T \quad (11)$$

where  $\mu$  is the surface coefficient of friction, the symbol  $\star \in [f, r]$  denotes the lumped front or rear tire,  $F_{z\star}$  is the corresponding tire normal load, and  $C_\star$  is the cornering stiffness for the corresponding tire. The linearized tire slip angles are given by

$$\alpha_f = \beta + \frac{ar}{U_x} - \delta \quad (5a)$$

$$\alpha_r = \beta - \frac{br}{U_x} \quad (5b)$$

### B. Linear Time Varying Model in the Lifted Domain

While the nonlinear tire model presented in (4) captures the effect of tire saturation at the limits of handling, the two iterative learning controllers considered in this study require a linear system description. With the assumption of low lateral acceleration, a simple linear tire model is given by,

$$F_{y\star} = -C_\star \alpha_\star \quad (6)$$

The lateral error dynamics of the vehicle in response to a *feedforward* steer angle input (recall that there is also a feedback steer angle  $\delta_{FB}$  in the loop) can be represented by the continuous, linear time varying (LTV) state space model:

$$\dot{x}(t) = A(t)x + B(t)\delta_L + d(t) \quad (7a)$$

$$e = Cx(t) \quad (7b)$$

$$x = [e \ \Delta\Psi \ r \ \beta]^T \quad (7c)$$

Where  $\delta_L$  is the learned steering input. The LTV framework is chosen to account for the variation in vehicle velocity  $U_x$  along the track (Fig. 1). While the longitudinal dynamics are not explicitly accounted for in (3a), allowing for the state matrices to vary with measured values of  $U_x$  enables more accurate vehicle control than assuming a constant velocity around the track. The time-varying matrices  $A$ ,  $B$ , and  $C$  are given by

$$A(t) = \begin{bmatrix} 0 & U_x(t) & 0 & U_x(t) \\ 0 & 0 & 1 & 0 \\ \frac{-ak_P C_f}{I_z} & \frac{-ak_P x_L A C_f}{I_z} & \frac{-(a^2 C_f + b^2 C_r)}{U_x(t) I_z} & \frac{bC_r - aC_f}{I_z} \\ \frac{-k_P C_f}{mU_x(t)} & \frac{-k_P x_L A C_f}{mU_x(t)} & \frac{bC_r - aC_f}{mU_x(t)^2} - 1 & \frac{-(C_f + C_r)}{mU_x(t)} \end{bmatrix} \quad (8)$$

$$B(t) = [0 \ 0 \ \frac{aC_f}{I_z} \ \frac{C_f}{mU_x(t)}]^T \quad (9)$$

The next step is to discretize (7) by the controller sample time  $T_s$ , resulting in the discrete time system

$$x_j(k+1) = A(k)x_j(k) + B(k)\delta_j^L(k) \quad (12)$$

$$e_j(k) = Cx_j(k) \quad (13)$$

where  $k = 1, \dots, N$  is the time sample index, and  $j = 1, \dots, M$  is the number of iterations (i.e. the number of laps around the track). Development and analysis of the iterative learning controllers in the next section will be made easier by representing the system dynamics in the ‘‘lifted-domain’’, where the inputs and outputs are stacked into arrays and related by matrix multiplication, as follows:

$$e_j = \mathbf{P}\delta_j^L + \mathbf{d}_j \quad (14a)$$

$$\delta_j^L = [\delta_j^L(0) \ \dots \ \delta_j^L(N-1)]^T \quad (14b)$$

$$e_j = [e_j(1) \ \dots \ e_j(N)]^T \quad (14c)$$

Where the elements of the  $N \times N$  matrix  $\mathbf{P}$  are given by:

$$p_{lk} = \begin{cases} 0 & \text{if } l < k \\ CB(k) & \text{if } l = k \\ CA(l)A(l-1) \dots A(k)B(k) & \text{if } l > k \end{cases} \quad (15)$$

Note that for the case where  $U_x$  is constant, we have a linear time invariant (LTI) system given by only  $N$  elements, with  $p(k) = CA^{k-1}B$ , and the resulting lifted-domain matrix  $\mathbf{P}$  is Toeplitz.

### III. CONTROLLER DESIGN

With the representation of the steering dynamics for a given lap represented by (14), the next step is to design algorithms that determine the learned steering input  $\delta_{j+1}^L$  for the next lap, given the error response from the completed lap  $e_j$ . A common framework for iterative learning algorithms is to choose  $\delta_{j+1}^L$  as follows [12],

$$\delta_{j+1}^L = \mathbf{Q}(\delta_j^L - \mathbf{L}e_j) \quad (16)$$

where  $\mathbf{Q}$  is the  $N \times N$  filter matrix, and  $\mathbf{L}$  is the  $N \times N$  learning matrix. In following sections, the matrices  $\mathbf{Q}$  and  $\mathbf{L}$  will be obtained by designing a PD type iterative learning controller as well as a quadratically optimal (Q-ILC) learning controller.

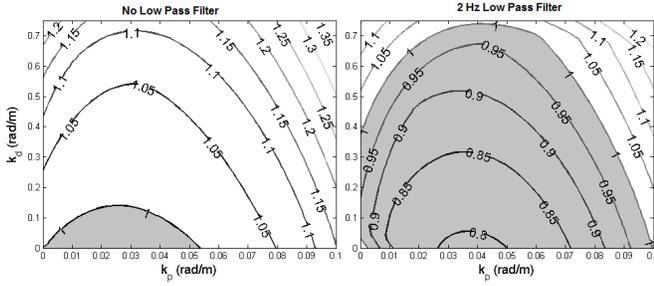


Fig. 3. Values of convergence bound  $\gamma$  vs.  $k_p$  and  $k_d$  for PD iterative learning controller with (left) no filtering and (right) with a 2 Hz low-pass filter. Lower values of  $\gamma$  correspond to faster convergence. Shaded regions correspond to systems with monotonic stability.

### A. Proportional-Derivative Controller

The proportional-derivative iterative learning controller computes the steering addition  $\delta^L$  for the current lap  $j$  based on the error  $e_{j-1}$  and error derivative at the same time index  $k$  from the previous lap:

$$\delta_j^L(k) = \delta_{j-1}^L(k) - k_p e_{j-1}(k) - k_d (e_{j-1}(k) - e_{j-1}(k-1)) \quad (17)$$

where  $k_p$  and  $k_d$  are the proportional and derivative gains. In the lifted domain representation from (16), the resulting learning matrix  $\mathbf{L}$  is given by

$$\mathbf{L} = \begin{bmatrix} -(k_p + k_d) & & 0 \\ & \ddots & \\ k_d & & \\ 0 & k_d & -(k_p + k_d) \end{bmatrix} \quad (18)$$

An important design consideration in choosing the gains  $k_p$  and  $k_d$  is avoiding a poor lap-to-lap “transient” response, where the path tracking error increases rapidly over the first several laps before eventually decreasing to a converged error response  $e_\infty$ . This is a commonly encountered design requirement for ILC systems, and can be solved by ensuring the following *monotonic convergence* condition is met [12]:

$$\gamma \triangleq \bar{\sigma}(\mathbf{P}\mathbf{Q}(I - \mathbf{L}\mathbf{P})\mathbf{P}^{-1}) < 1 \quad (19)$$

where  $\bar{\sigma}$  is the maximum singular value. In this case, the value of  $\gamma$  provides an upper bound on the change in the tracking error norm from lap to lap, i.e.

$$\|e_\infty - e_{j+1}\|_2 \leq \gamma \|e_\infty - e_j\|_2 \quad (20)$$

Fig. 3 shows values of  $\gamma$  for both an unfiltered PD controller ( $\mathbf{Q} = I$ ), and for a PD controller with a 2 Hz low pass filter. The  $\gamma$  values are plotted as a contour map against the controller gains  $k_p$  and  $k_d$ . Addition of the low-pass filter assists with monotonic stability by removing oscillations in the control input generated when trying to remove small reference tracking errors after several iterations. Since the filtering occurs when generating a control signal for the next lap, the filter  $\mathbf{Q}$  can be zero-phase.

However, testing for linear stability is insufficient for controller design given that racing frequently occurs near the

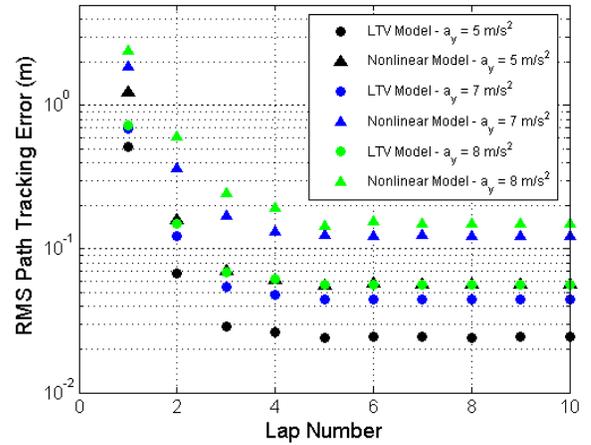


Fig. 4. Simulated results for root-mean-square path tracking error for PD iterative learning control at several values of vehicle lateral acceleration. Controller gains are  $k_p = .05$ ,  $k_d = .05$ . Results from a nonlinear vehicle dynamics simulation are compared to results from the linear vehicle model.

limits of vehicle handling, when the vehicle dynamics are described by the nonlinear equations of motion presented in Section II. To test the PD controller feasibility, the vehicle tracking performance over multiple laps is simulated using the path curvature and speed profile shown in Fig. 1. Simulated results for the root-mean-square (RMS) tracking error are shown in Fig. 4 for both the linear state dynamics described by (13) and the nonlinear dynamics model given by (3) and (4). The results indicate that as the vehicle corners closer to the limits of handling, the tracking performance of the ILC degrades relative to the expected performance given by the linear model, but can still be expected to converge over relatively few iterations.

### B. Quadratically Optimal Controller

An alternate approach to determining the learned steering input  $\delta_j^L$  is to minimize a quadratic cost function for the next lap:

$$J_{j+1} = e_{j+1}^T T e_{j+1} + \delta_{j+1}^{L T} R \delta_{j+1}^L + \Delta_{j+1}^T S \Delta_{j+1} \quad (21)$$

where  $\Delta_{j+1} = \delta_{j+1}^L - \delta_j^L$  and the  $N \times N$  matrices  $T$ ,  $R$ , and  $S$  are weighting matrices, each given by a scalar multiplied by the identity matrix for simplicity. This formulation allows the control designer to weight the competing objectives of minimizing tracking error, control effort, and change in the control signal from lap to lap. While constraints can be added to the optimization problem, the unconstrained problem in (21) can be solved analytically [13] to obtain desired controller and filter matrices:

$$\mathbf{Q} = (\mathbf{P}^T T \mathbf{P} + R + S)^{-1} (\mathbf{P}^T T \mathbf{P} + S) \quad (22a)$$

$$\mathbf{L} = (\mathbf{P}^T T \mathbf{P} + S)^{-1} \mathbf{P}^T T \mathbf{P} (\mathbf{T}^{1/2} \mathbf{P})^{-1} \mathbf{T}^{1/2} \quad (22b)$$

An advantage of the quadratically optimal control design over the simple PD controller is that the controller matrices  $\mathbf{Q}$  and  $\mathbf{L}$  take the linear time-varying dynamics  $\mathbf{P}$  into

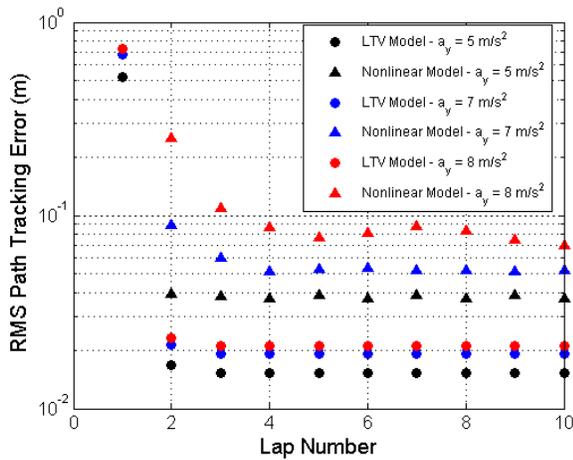


Fig. 5. Simulated results for root-mean-square path tracking error for Q-ILC at several values of vehicle lateral acceleration, with  $T = R = I$  and  $S = 100I$ . Results from a nonlinear vehicle dynamics simulation are compared to results from the linear vehicle model.



Fig. 6. Autonomous Audi TTS with electronic power steering, brake booster, and throttle by wire.

account. This allows the iterative learning algorithm to take into account changes in the steering dynamics due to changes in vehicle velocity. However, a disadvantage is that computing  $\delta^L$  in (16) requires matrix multiplications with the typically dense matrices  $\mathbf{Q}$  and  $\mathbf{L}$  for every lap, which can be computationally expensive for fast sampling rates.

Fig. 5 shows simulated results for the quadratically optimal ILC controller, with  $T = R = I$  and  $S = 100I$ . For simplicity, the weighting matrices were all chosen to be a scalar multiplied by the identity matrix, with specific values chosen coarsely based on trial and error. The results are similar to those in Fig. 4 in that the predicted controller performance is worse when the simulation accounts for nonlinear tire dynamics. However, the simulation still shows a rapid decrease in path tracking error over the first ten laps, with RMS errors on the order of 8-9 cm at lateral accelerations of 0.8  $g$ .

#### IV. EXPERIMENTAL RESULTS

Experimental data for both controllers were collected over multiple laps at Thunderhill Raceway, a three-mile paved racetrack in Willows, CA, with track boundaries shown in Fig. 1a. The experimental testbed is an autonomous Audi TTS equipped with an electronic power steering motor, active brake booster, and throttle by wire (Fig. 6). Vehicle and controller parameters are shown in Table 1.

An integrated Differential Global Positioning System (DGPS) and Inertial Measurement Unit (IMU) is used to obtain vehicle state information, and a localization algorithm determines the lateral path tracking error  $e$ , heading error  $\Delta\Psi$ , and distance along the desired racing line. The steering controller updates at 200 Hz, and the iterative steering corrections are calculated after each lap from data downsampled to 10 Hz. The corrections are then applied as a function of distance along the track using an interpolated lookup table. For safety reasons, a steady-state feedforward steering algorithm [3] is also applied to keep the tracking error on the first lap below 1 m.

TABLE I  
VEHICLE PARAMETERS

Parameter	Symbol	Value	Units
Vehicle mass	$m$	1500	kg
Yaw Inertia	$I_z$	2250	$\text{kg} \cdot \text{m}^2$
Front axle to CG	$a$	1.04	m
Rear axle to CG	$b$	1.42	m
Front cornering stiffness	$C_F$	160	$\text{kN} \cdot \text{rad}^{-1}$
Rear cornering stiffness	$C_R$	180	$\text{kN} \cdot \text{rad}^{-1}$
Lookahead Distance	$x_{LA}$	15.2	m
Lanekeeping Gain	$k_{LK}$	0.053	$\text{rad m}^{-1}$
Lanekeeping Sample Time	$t_s$	0.005	s
ILC Sample Time	$T_s$	0.1	s
PD Gains	$k_p$ and $k_d$	0.02 & .4	$\text{rad m}^{-1}$
Q-ILC Matrix	$T$ and $R$	$I$	-
Q-ILC Matrix	$S$	$100I$	-

Fig. 7 shows the applied iterative learning signals and resulting path tracking error over four laps using the PD learning algorithm. The car is driven aggressively at peak lateral/longitudinal accelerations of 8  $\text{m/s}^2$ . On the first lap, despite the incorporation of a feedforward-feedback controller operating at a high sampling rate, several transient spikes in tracking error are visible due to the underdamped tire dynamics near the limits of handling. However, the iterative learning algorithm is able to significantly attenuate these transient spikes after just two or three laps. Similar qualitative results occur for the quadratically optimal ILC.

In Fig. 8, results are shown for versions of the velocity profile in Fig. 1 that are scaled to achieve different vehicle accelerations. The results show that at low vehicle accelerations, when vehicle dynamics are accurately described by a linear tire model, the feedback-feedforward controller is able to keep the vehicle close to the desired path, leaving little room for improvement through iterative learning control. However, as the speed profile becomes more aggressive, the path tracking degrades in the presence of highly transient tire dynamics, and iterative learning control can be effectively used to obtain tight path tracking over two or three laps of racing. In practice, the performance of both the PD algorithm and quadratically optimal algorithms are similar, and an important observation is that the RMS tracking error increases slightly from lap-to-lap at the end of some tests. While not predicted in simulation, this behavior is not unreasonable given unmodelled sensor noise and disturbances that vary from lap to lap. More refined tuning of the gain matrices may be able to prevent this RMS error increase, or the ILC

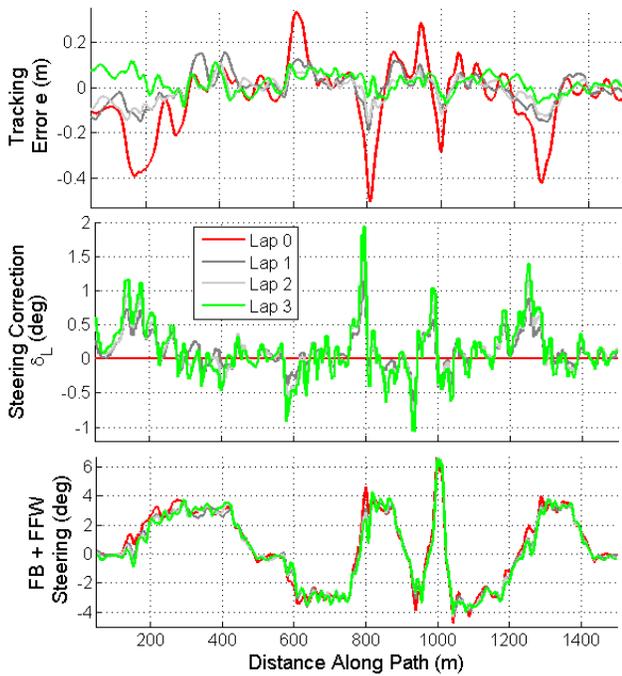


Fig. 7. Experimental results for path tracking error with PD learning controller, at peak lateral accelerations of  $8 \text{ m/s}^2$ .

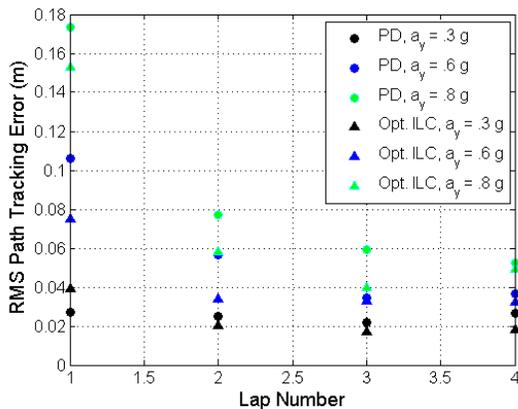


Fig. 8. Experimental results for both controller types over a variety of lateral accelerations.

algorithm can be stopped after several iterations once the tracking performance is acceptable.

## V. CONCLUSION

This paper demonstrates the application of iterative learning control methods to achieve accurate steering control of an autonomous race car over multiple laps. Two different algorithms, proportional-derivative and quadratically optimal learning control are tested in simulation and then used to experimentally eliminate path tracking errors caused by the highly transient nature of the vehicle lateral dynamics near the limits of tire-road friction. Both learning algorithms provide comparable lap-to-lap tracking performance, although the proportional-derivative method is computationally fast

enough to run in real time. While the proposed learning algorithms perform well for an autonomous race car driving the same trajectory every lap, future work will consider the challenge of applying iterative learning control for more general autonomous driving applications. One promising avenue for future research is to adjust a feedforward vehicle model based on information learned from one autonomous driving maneuver, enabling the vehicle to perform a wide variety of similar maneuvers with high accuracy.

## ACKNOWLEDGMENT

This research is supported by the National Science Foundation Graduate Research Fellowship Program (GRFP). The authors would like to thank the members of the Dynamic Design Lab at Stanford University and the Audi Electronics Research Lab.

## REFERENCES

- [1] J. Funke and J. C. Gerdes, "Simple clothoid paths for autonomous vehicle lane changes at the limits of handling," in *ASME 2013 Dynamic Systems and Control Conference*. American Society of Mechanical Engineers, 2013, pp. V003T47A003–V003T47A003.
- [2] V. Turri, A. Carvalho, H. Tseng, K. H. Johansson, and F. Borrelli, "Linear model predictive control for lane keeping and obstacle avoidance on low curvature roads," 2013.
- [3] N. R. Kapania and J. C. Gerdes, "An autonomous lanekeeping system for vehicle path tracking and stability at the limits of handling," in *12th International Symposium on Advanced Vehicle Control*, 2014.
- [4] Y. Q. Chen and K. L. Moore, "A practical iterative learning path-following control of an omni-directional vehicle," *Asian Journal of Control*, vol. 4, no. 1, pp. 90–98, 2002.
- [5] H. Sun, Z. Hou, and D. Li, "Coordinated iterative learning control schemes for train trajectory tracking with overspeed protection," *IEEE Trans. on Automation Science and Engineering*, vol. 10, no. 2, pp. 323–333, 2013.
- [6] O. Purwin and R. DAndrea, "Performing and extending aggressive maneuvers using iterative learning control," *Robotics and Autonomous Systems*, vol. 59, no. 1, pp. 1–11, 2011.
- [7] P. A. Theodosios and J. C. Gerdes, "Generating a racing line for an autonomous racecar using professional driving techniques," in *Dynamic Systems and Control Conference*, 2011, pp. 853–860.
- [8] K. Kritayakirana, "Autonomous vehicle control at the limits of handling," Ph.D. dissertation, Stanford University, 2012.
- [9] E. J. Rossetter and J. C. Gerdes, "A study of lateral vehicle control under a virtual force framework," in *Proc. International Symposium on Advanced Vehicle Control, Hiroshima, Japan*, 2002.
- [10] K. L. Talvala and J. C. Gerdes, "Lanekeeping at the limits of handling: Stability via lyapunov functions and a comparison with stability control," in *Dynamic Systems and Control Conference*, 2008, pp. 361–368.
- [11] H. B. Pacejka, *Tire and Vehicle Dynamics*, 3rd ed. Butterworth-Heinemann, 2012.
- [12] D. A. Bristow, M. Tharayil, and A. G. Alleyne, "A survey of iterative learning control," *Control Systems, IEEE*, vol. 26, no. 3, pp. 96–114, 2006.
- [13] D. A. Bristow and B. Hencney, "A q,l factorization of norm-optimal iterative learning control," in *47th IEEE Conference on Decision and Control*, 2008. IEEE, 2008, pp. 2380–2384.