Creating Predictive Haptic Feedback For Obstacle Avoidance Using a Model Predictive Control (MPC) Framework

Avinash Balachandran, Matthew Brown, Stephen M. Erlien and J. Christian Gerdes

***Abstract*— New sensing technologies allow modern vehicles to perceive the environment around them even when human visual perception is limited due to poor lighting or fog. Steer-by-wire technology enables active steering capability in which the driver's command to the roadwheels is augmented for maintaining safety. Predictive controllers can leverage both of these technologies to create shared control safety systems that work with the driver to ensure a safe and collision-free vehicle trajectory. The earlier the system intervenes, the smoother the intervention but the more it interferes with the driver's control authority. Ideally, predictive controllers should still intervene late but also indicate upcoming environmental threats to the driver as early as possible. Haptic feedback provides a good means of communicating information to the driver early. Together with a controller still providing a late intervention fallback, this regime provides an ideal framework for predictive shared control systems. This paper presents a novel technique for creating haptic steering feedback, based on future differences between the predictive controller and the driver. This feedback mirrors the tension between the sometimes competing controller objectives of following the driver and maintaining a feasible path. The paper uses simulation and experiment to investigate the inherent trade-offs of predictive haptic feedback and qualitatively discuss its impact.**

I. INTRODUCTION

New perception technologies are equipping modern vehicles with knowledge of the environment around them [1]. Sensing systems like radar, cameras and laser systems provide rich information about the environment in real-time even when human visual perception is limited due to poor lighting or fog. Steer-by-wire systems also empower vehicles by decoupling the handwheel from the roadwheels allowing the driver's steering command to the roadwheel to be augmented, a feature known as active steering. Ackermann et al. showed that active steering can aid vehicle control in various situations like yaw control and rollover avoidance [2].

Leveraging both the enhanced environmental perception and steer-by-wire actuation technology, predictive controllers can share control with the driver and enhance the safety of modern vehicles. Kawabe et al. presented a receding horizon control framework that uses information about the surrounding environment to generate optimal paths to guide a human driver [3]. Erlien et al. proposed a shared control framework, particularly for obstacle avoidance, that uses Model Predictive Control (MPC) for maintaining a safe trajectory using information about the environment [4].

Using predictive controllers to share control with the driver results in a trade-off between the sometimes competing objectives of following the driver and maintaining a feasible path. Biasing the controller heavily to the latter objective results in early and frequent interventions that interfere with the driving task and reduce the driver's control authority. Heavily biasing the controller to follow the driver results in later interventions that are more aggressive and uncomfortable but leave control authority with the driver for longer. Haptic feedback balances the need to warn the driver of upcoming environmental threats early while still having a late intervention fallback. Ideally, haptic feedback will guide a driver to modify his trajectory early so that the late intervention is minimal or not necessary.

Switkes et al. performed a simulator study that used the current lane-keeping potential field force acting on a vehicle as haptic steering feedback [5]. Steele and Gillespie generated haptic feedback on a simulator using the difference between the driver's current steering angle and the desired steering angle as determined by a autonomous path-tracking controller and a pre-defined static path [6]. Brandt et al. used a similar approach to generate haptic feedback on a simulator except, instead of static paths, a potential field technique was used to generate dynamic paths [7]. Predictive controllers enable haptic feedback based on the future tension between the objectives of following the driver or maintaining a feasible path. This feedback could guide the driver to modify his actions early.

This paper proposes a novel method of generating haptic steering feedback based on the future tension between the competing objectives of following the driver and maintaining a feasible path. The paper begins by describing the vehicle model that is used for simulation and, when linearized, is used for the predictive controller. The next section briefly describes the Model Predictive Controller used here which was developed by Erlien et al. for obstacle avoidance [4]. The paper then describes, using simulation, how to obtain predictive haptic feedback from the MPC solution so that the future tension between both objectives is captured. Trade-offs with this approach are also discussed. Lastly, the paper describes preliminary experimental driver data to qualitatively discuss the impact of this haptic feedback on the obstacle avoidance task.

II. VEHICLE MODEL

The vehicle model used in this work is a constant speed, planar bicycle model. The vehicle's motion is described by two states: sideslip β and yawrate r. These are defined in

The authors are with the Dynamic Design Lab, Department of Mechanical Engineering, Stanford University, Stanford, CA, 94305 USA email:avinashb@stanford.edu

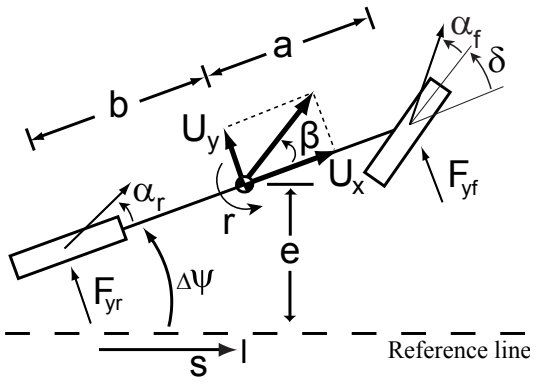

Fig. 1. Bicycle model schematic

Figure 1 and have the following equations of motion:

$$\dot{\beta} = \frac{F_{\mathrm{yf}} + F_{\mathrm{yr}}}{mU_{\mathrm{x}}} - r \qquad \dot{r} = \frac{aF_{\mathrm{yf}} - bF_{\mathrm{yr}}}{I_{\mathrm{zz}}} \tag{1}$$

where $F_{\mathrm{y[f,r]}}$ is the lateral tire force of the [front, rear] axle, m is the vehicle mass, U_{x} is the longitudinal velocity in the body fixed frame, I_{zz} is the yaw inertia, and a and b are the distances from the center of gravity to the front and rear axles, respectively.

The simplified Fiala tire model introduced by Pacejka in [8] gives a useful approximation of the non-linear relationship between $F_{\mathrm{y[f,r]}}$ and the tire slip angles $\alpha_{\mathrm{[f,r]}}$:

$$\begin{aligned} F_{\mathrm{y}} &= \begin{cases} -C_\alpha \tan\alpha + \frac{C_\alpha^2}{3\mu F_{\mathrm{z}}} |\tan\alpha| \tan\alpha \ldots \\ \qquad -\frac{C_\alpha^3}{27\mu^2 F_{\mathrm{z}}^2} \tan^3\alpha, & |\alpha| < \tan^{-1}\left(\frac{3\mu F_{\mathrm{z}}}{C_\alpha}\right) \\ -\mu F_{\mathrm{z}} \mathrm{sgn}\ \alpha, & \text{otherwise} \end{cases} \quad (2) \\ &= f_{\mathrm{tire}}(\alpha) \quad (3) \end{aligned}$$

where μ is the surface coefficient of friction, F_{z} is the normal load, and C_α is the tire cornering stiffness.

The tire slip angles, α_{f} and α_{r}, can be derived from the kinematics of the vehicle as:

$$\alpha_{\mathrm{f}} = \tan^{-1}\left(\beta + \frac{ar}{U_{\mathrm{x}}}\right) - \delta_{\mathrm{rw}} \tag{4}$$

$$\alpha_{\mathrm{r}} = \tan^{-1}\left(\beta - \frac{br}{U_{\mathrm{x}}}\right) \tag{5}$$

where δ_{rw} is the roadwheel angle.

The vehicle's position is specified relative to a reference line using three states: heading deviation $\Delta\psi$, lateral deviation e, and distance along the path s as defined in Figure 1.

The equations of motion of these states are then:

$$\dot{\Delta\psi} = r - U_x K \tag{6}$$

$$\dot{e} = U_{\mathrm{x}} \sin(\Delta\psi) + U_{\mathrm{y}} \cos(\Delta\psi) \tag{7}$$

$$\dot{s} = U_{\mathrm{x}} \cos(\Delta\psi) - U_{\mathrm{y}} \sin(\Delta\psi) \tag{8}$$

where K is the curvature of the reference line at position s.

III. MODEL PREDICTIVE OBSTACLE AVOIDANCE CONTROLLER

Generating predictive haptic feedback requires a technique for projecting a vehicle's trajectory into the future. MPC lends itself well to this purpose and the control scheme described in this section is based on the MPC-based obstacle avoidance controller developed by Erlien et al. [4].

This controller is based on a shared control scheme. Vehicle handling limits define the stable handling envelope while obstacles and road boundaries in the environment define the environmental envelope. An MPC scheme continually balances between the objectives of following the driver versus maintaining a feasible trajectory, intervening only as a last resort. This controller is briefly described in this section.

A. MPC Vehicle Model

The vehicle model used in the online MPC controller is a linearization of the vehicle model presented in Section II. The input to this vehicle model is F_{yf}, allowing the controller to directly consider tire non-linearity by computing an optimal front tire force which is then mapped to δ_{rw} using (3) and (4):

$$\delta_{\mathrm{rw}} = \beta + \frac{ar}{U_{\mathrm{x}}} - f_{\mathrm{tire}}^{-1}(F_{\mathrm{yf}}) \tag{9}$$

where f_{tire}^{-1} is computed numerically.

At each execution of the predictive controller, the measured speed is assumed to be constant over the prediction horizon. By linearizing around a given path and nominal α_r, a discrete, time-varying vehicle model can be expressed as:

$$x^{(k+1)} = A^{(k)}_{\bar{\alpha}_{\mathrm{r}}, t_{\mathrm{s}}} x^{(k)} + B^{(k)}_{\bar{\alpha}_r, t_s} {F_{\mathrm{yf}}}^{(k)} + d^{(k)}_{\bar{\alpha}_r, t_s} \tag{10}$$

where $x = [\beta \quad r \quad \Delta\psi \quad s \quad e]^T$, k is a time step index, subscript $\bar{\alpha}_{\mathrm{r}}$ denotes linearization of the rear tire model around rear slip angle $\bar{\alpha}_{\mathrm{r}}$, and subscript t_{s} denotes discretization of the vehicle model using time step length t_{s}.

B. Environmental Envelope

The environmental envelope consists of a set of obstacle-free points in s and e, as in Figure 2 (left). To avoid collision with the environment, the vehicle's trajectory must be fully contained within this set. The environmental envelope defines time-varying constraints on the lateral deviation of the vehicle from the reference line:

$$e^{(k)} \leq e^{(k)}_{\max} - \frac{1}{2}d - d_{\mathrm{buffer}} \tag{11}$$

$$e^{(k)} \geq e^{(k)}_{\min} + \frac{1}{2}d + d_{\mathrm{buffer}} \tag{12}$$

where $e^{(k)}_{\max}$ and $e^{(k)}_{\min}$ indicate the lateral deviation bounds for time step k, d is the vehicle's width, and d_{buffer} specifies a preferred minimum distance between obstacles and the vehicle to ensure driver comfort.

Inequalities (11) and (12) can be compactly expressed as:

$$H_{\mathrm{env}} x^{(k)} \leq G^{(k)}_{\mathrm{env}} \tag{13}$$

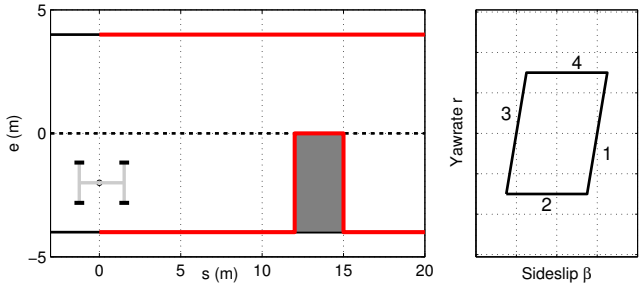

Fig. 2. The environmental envelope (left) provides a provides safe bounds on e for every point in s. The stable handling envelope (right) represents a safe region in the sideslip-yawrate space

where subscript env indicates the environmental envelope. A vehicle trajectory is collision free throughout the prediction horizon if and only if it satisfies inequality (13) for all k.

C. Stable Handling Envelope

The stable handling envelope defines limits on the vehicle states yaw rate and sideslip as illustrated in Figure 2 (right). Originally presented by Beal and Gerdes, this envelope reflects the maximum capabilities of the vehicle's tires such that at any point within this envelope, there exists a steering command to safely remain inside, ensuring stability [9].

The maximum steady-state yawrate defines bounds (2) and (4) in Figure 2 (right), while the rear slip angle at which lateral force saturates serves as the basis for bounds (1) and (3). The desire to restrain the vehicle states within the stable handling envelope is compactly expressed as the inequality:

$$\left| H_{\mathrm{sh}} x^{(k)} \right| \leq G_{\mathrm{sh}} \tag{14}$$

where subscript sh indicates the stable handling envelope.

D. MPC Formulation

The controller's objectives can be expressed as an optimization problem to be evaluated over a finite prediction horizon:

$$\text{minimize} \quad \sum_{k=0}^{10} \left| F_{\mathrm{yf,driver}} - F_{\mathrm{yf,opt}}^{(k)} \right| \tag{15a}$$

$$+ \sum_{k} \gamma^{(k)} \left| \left| F_{\mathrm{yf,opt}}^{(k)} - F_{\mathrm{yf,opt}}^{(k-1)} \right| \right|_2 \tag{15b}$$

$$\text{subject to} \quad x^{(k+1)} = A_{\bar{\alpha}_{\mathrm{r}},t_{\mathrm{s}}}^{(k)} x^{(k)} + B_{\bar{\alpha}_{\mathrm{r}},t_{\mathrm{s}}}^{(k)} F_{\mathrm{yf,opt}}^{(k)} + d_{\bar{\alpha}_{\mathrm{r}},t_{\mathrm{s}}}^{(k)} \tag{15c}$$

$$\left| F_{\mathrm{yf,opt}}^{(k)} \right| \leq F_{\mathrm{yf,max}} \tag{15d}$$

$$\left| H_{\mathrm{sh}} x^{(k+1)} \right| \leq G_{\mathrm{sh}} \tag{15e}$$

$$H_{\mathrm{env}} x^{(k+1)} \leq G_{\mathrm{env}}^{(k+1)} \tag{15f}$$

$$\left| F_{\mathrm{yf,opt}}^{(k)} - F_{\mathrm{yf,opt}}^{(k-1)} \right| \leq F_{\mathrm{yf,max\ slew}}^{(k)} \tag{15g}$$

where $k = 0 \ldots 29$ and the variables to be optimized are the optimal input trajectory ($F_{\mathrm{yf,opt}}$). As is typical in receding horizon control, only the calculated optimal force at the first time step, $F_{\mathrm{yf,opt}}^{(0)}$, is applied to the vehicle, and the optimization problem (15) is re-solved at the next time step.

A tunable weight in this optimization is γ, which trades off between generating smooth steering inputs and matching driver intent. Cost term (15a) uses the l_1 norm as a convex approximation to the objective of *identically matching* the driver's command where $F_{\mathrm{yf,driver}}$ is the front tire force corresponding to the driver's hand wheel , δ_{hw}. The previously described tire model (3) maps from δ_{hw} to $F_{\mathrm{yf,driver}}$:

$$F_{\mathrm{yf,driver}} \quad = \quad f_{\mathrm{tire}} \left(\beta + \frac{ar}{U_x} - \delta_{\mathrm{hw}} \right) \tag{16}$$

Constraint (15d) reflects the maximum force capabilities of the front tires and (15g) reflects the slew rate limits of a physical steering system. Constraints (15e) and (15f) enforce the stable handling and environmental envelopes, respectively. In practice, slack variables are used on these constraints to ensure (15) is always feasible.

The prediction horizon used in optimization (15) uses time steps of different length in the near and long terms of the horizon to consider obstacles in the long term without compromising the prediction of vehicle states in the near term. CVXGEN [10] is used to leverage the significantly sparse structure of convex optimization problem (15) to produce an efficient solver for real-time implementation [11]. For the interested reader, Erlien et al. provides a more detailed description of this controller [4].

IV. Generating Predictive Haptic Feedback

This section discusses the generation of haptic feedback based on predictive information obtained from the obstacle avoidance MPC controller described in Section III. The haptic feedback given is proportional to the difference between the driver's steering input and the controller's desired steering action at a predetermined future prediction index. The haptic torque is generated as follows:

$$\tau_{\mathrm{haptic}} = K_{\mathrm{haptic}} \left(\delta_{\mathrm{opt}}^{(k)} - \delta_{\mathrm{driver}} \right) \tag{17}$$

where τ_{haptic} is the predictive haptic feedback torque, K_{haptic} is the gain for the haptic feedback, $\delta_{\mathrm{opt}}^{(k)}$ is the optimal steering angle at a pre-determined prediction index k, obtained from the optimization, and δ_{driver} is the driver's steering input.

This mirrors the tension between the competing objectives of following the driver versus maintaining a feasible path. If there is no environmental danger detected, there is no tension as the driver's input will result in a feasible path. Therefore, the driver's input and the controller's desired steering action will be the same and there will be no feedback. However, as this tension grows so will the difference between the driver's input and controller, resulting in a feedback torque.

Figure 3 shows a simulation demonstrating the generation of this predictive feedback torque. It also investigates the impact of using different prediction indices, k. The scenario used for generating this simulation involves a vehicle starting on the right lane of a two-lane road and traveling at a

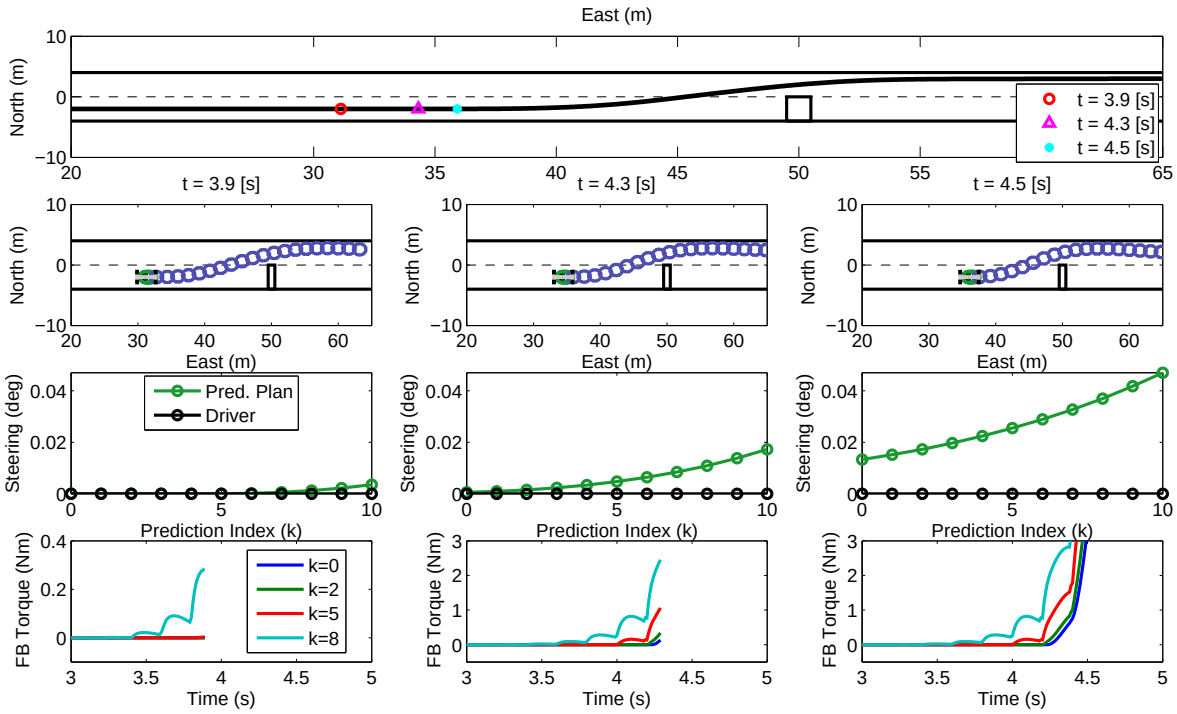

Fig. 3. Simulation results

constant speed towards an obstacle blocking the right lane. The driver is modeled as holding a constant zero steering angle throughout the simulation. The uppermost plot of Figure 3 represents the trajectory of the vehicle over the simulation. The columns of the 3 vertical plots provide a visualization of the MPC solution and the haptic feedback torque at simulation execution time steps $t = 3.9[s]$, $t = 4.3[s]$ and $t = 4.5[s]$. The first row of these columns shows, spatially, the trajectory solution obtained from the MPC solution. The second row of these columns shows the driver's steering input (in this case always zero) and the MPC solution's optimal desired steering action for the first 10 prediction indices, k. The last row of these columns shows the haptic feedback torque calculated by using Equation (17) and different prediction indices, up to and including that particular column's execution time step.

Tension between the objectives of following the driver and maintaining a feasible trajectory builds as the steering action required to maintain a feasible trajectory starts differing greatly from the driver's input. This tension begins far out in the prediction horizon and propagates to smaller prediction indices if the driver continues to not modify his steering input. This can been seen in the second row of steering plots. The difference between the driver's input and the controllers desired steering action is the largest at the end of the horizon and propagates to smaller prediction indices as the execution time step increases (i.e. going from left to right). In the steering plot for the execution time step, $t = 4.5[s]$, the driver's input and the controller's desired steering angle no longer agree on the first prediction index, $k = 0$, indicating that an active steering intervention is occurring to maintain a collision-free trajectory. Intuitively, this shows the tension growing as the driver continues without changing his steering input and travels towards the obstacle. By correlating these steering plots with the trajectory plots in the row above, notice that this tension grows as the required trajectory to avoid the obstacle becomes more aggressive.

The last row shows the feedback based on using Equation (17) and different prediction indices, k. The larger the prediction index used to generate the feedback, the earlier this feedback is felt by the driver. This makes sense as the tension propagates down from further out in the prediction horizon to smaller prediction indices over time. Therefore, using a larger prediction index to generate the feedback results in the feedback kicking in earlier and indicating to the driver that the tension between the competing objectives is growing. In this way, predictive feedback can be given to the driver to guide him in modifying his actions early to reduce this tension. Since these objectives vary with different situations, the tension between balancing them is also highly situation dependent. Therefore, unlike other obstacle avoidance haptic warnings, the haptic feedback using this technique is not based on a fixed time to collision (TTC).

However, there is a trade-off in using the larger prediction indices. As with any prediction, the longer into the prediction horizon that one looks at, the more variability there will be. The variability in the tension when looking far into the prediction horizon finds its way into the haptic feedback

torque generated if large prediction indices are used. This can been seen in the feedback torque plot of execution time step $t = 4.5[s]$ where the larger the prediction index, k, used to generate the feedback, the more variability in that feedback torque. Therefore, as the prediction index, k, increases the feedback kicks in earlier but has more variability. Much of this variability can be damped out by the inherent dissipative dynamics of the steering system and the driver's arms. The authors have found that using prediction indices between 4 and 10 result in early feedback with reduced variability.

V. EXPERIMENTAL RESULTS

Experiments run on the steer-by-wire vehicle, X1 (illustrated in Figure 4) were used to verify the impact of this predictive haptic feedback on the obstacle avoidance task. X1 uses an integrated Global Positioning System (GPS) and Inertial Measurement System (INS) to obtain vehicle states and has a Force Feedback (FFB) steering system that generates user-defined torque on the handwheel. A dSPACE MicroAutoBox II (DS1401) interfaces with the subsystems of the vehicle at a rate of 500 [Hz] with the model predictive controller implemented on a single core of an Intel i7 processor utilizing MATLAB's real-time xPC toolbox at a rate of 100 [Hz]. To ensure that the participants feel comfortable driving the test vehicle, a nominal steering feel, similar to that of a production vehicle, needs to be created. The supportive steering feel model based on tire moments and steering system dynamics proposed by Balachandran et al. [12] was used to create this feel. The haptic feedback is then overlaid on this steering feel. This section presents a qualitative discussion of a single user using the predictive haptic feedback system.

A. Test Scenario

To test the efficacy of predictive haptic feedback, the visual indication of an upcoming obstacle must be decoupled from the predictive haptic feedback. This ensures that recorded driver reactions are due to the predictive haptic feedback rather than the visual cue of an upcoming obstacle. This decoupling models scenarios in which vehicles with enhanced perception technologies have better information about the environment than the driver particularly during poor lighting or foggy conditions. Jensen et al. used a similar approach of decoupling haptic and visual cues when studying, using a simulator, the efficacy of using sinusoidal haptic feedback to alert the driver to upcoming obstacles [13].

Fig. 4. X1 Steer-by-wire vehicle

Figure 5 illustrates the two-lane test scenario use for this experiment. To decouple the visual and haptic cues of the obstacle, 'popup' obstacles were used. The vehicle had information about the upcoming obstacle before deployment so the driver would receive predictive haptic feedback before the deployment of the obstacle. Before the deployment of the obstacle, drivers started by driving on the right lane as illustrated in Figure 5(a). Cruise control was used to maintain a constant speed (7 [mps]) so that drivers avoided the obstacle through only the use of steering. During each trial, either one or none of the 3 'popup' obstacles (indicated in Figure 5 as the red dotted circles) would randomly deploy. If an obstacle deployed, it would block the rightmost lane as indicated in Figure 5(b) causing drivers to perform a lane change to the left lane. Drivers drove with both treatments, with and without predictive haptic feedback, where the future index used to obtain the feedback was 10 (i.e. $k = 10$). Since the experimental conditions were repeatable for each participant, this setup resulted in the haptic feedback occurring at about a TTC of 3.5 [s] for all participants.

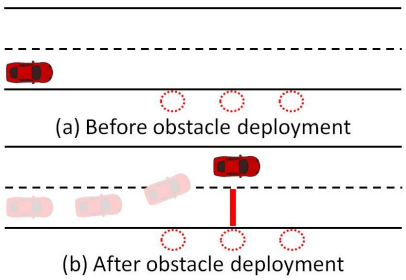

Fig. 5. Experimental setup

B. Discussion

Figure 6 illustrates experimental data from a single driver. A single trial with predictive feedback (red) and one without (blue) is shown. The large plot on the left shows the trajectory of the vehicle and indicates that with the predictive feedback, the driver modifies his trajectory to avoid the obstacle earlier. This is clearly seen in the handwheel angle plot on the top right where the red line indicates an earlier turn than the blue line. The peak steering input required is also smaller. The handwheel angle derivative (second plot from the top) also indicates that fewer steering corrections are required with the feedback. The bottom two plots indicate vehicle yawrate and lateral acceleration respectively and show that these quantities are generally smaller as well, indicating a smoother motion of the vehicle. Overall, the predictive feedback results in the driver reacting and turning earlier in response to an upcoming obstacle. This in turn results in a smaller required steering input, less required steering corrections and a smoother vehicle motion.

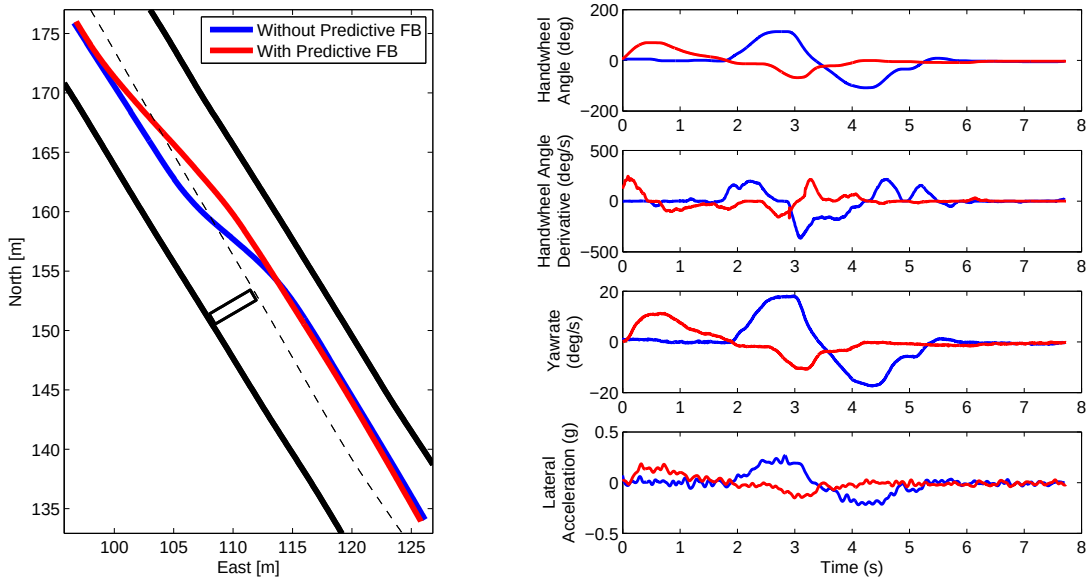

Fig. 6. Experimental driver data

VI. CONCLUSION

This paper proposes a novel method of generating haptic steering feedback for a shared control MPC-based controller. The feedback is based on the future tension between the competing objectives of following the driver and maintaining a feasible path. The controller is tuned to intervene late to maintain driver authority for as long as possible. The haptic feedback generated mirrors the future tension between the two objectives with larger feedback indicating greater future tension. The larger the prediction index used to generate the haptic feedback, the earlier the feedback is felt by the driver. However, simulation shows that using information from far in the prediction horizon injects the variability from these predictions into the feedback. Lastly, preliminary experimental driver data is used to qualitatively verify the efficacy of the predictive haptic feedback. Future work will focus on the completion of a full driver study with multiple participants to quantitatively show the efficacy of the predictive haptic feedback.

ACKNOWLEDGMENT

The authors would like to acknowledge Brian Canada and Phillip Garcia from Stanford University Parking & Transportation Services for their invaluable logistical support in conducting the experimental driver studies.

REFERENCES

[1] J. Leonard, J. How, S. Teller, M. Berger, S. Campbell, G. Fiore, L. Fletcher, E. Frazzoli, A. Huang, S. Karaman, O. Koch, Y. Kuwata, D. Moore, E. Olson, S. Peters, J. Teo, R. Truax, M. Walter, D. Barrett, A. Epstein, K. Maheloni, K. Moyer, T. Jones, R. Buckley, M. Antone, R. Galejs, S. Krishnamurthy, and J. Williams, "A perception-driven autonomous urban vehicle," *Journal of Field Robotics*, vol. 25, no. 10, pp. 727–774, Oct. 2008.

[2] J. Ackermann, D. Odenthal, and T. Bunte, "Advantages of active steering for vehicle dynamics control," *Proceedings of 32nd ISATA, Automotive Mechatronics Design and Engineering*, pp. 263–270, 1999.

[3] T. Kawabe, H. Nishira, and T. Ohtsuka, "An optimal path generator using a receding horizon control scheme for intelligent automobiles," in *Proceedings of the 2004 IEEE International Conference on Control Applications, 2004.*, vol. 2, no. 4. IEEE, 2004, pp. 1597–1602.

[4] S. M. Erlien, S. Fujita, and J. C. Gerdes, "Safe driving envelopes for shared control of ground vehicles," in *7th IFAC Symposium on Advances in Automotive Control*, Sep. 2013, pp. 831–836.

[5] J. P. Switkes, E. J. Rossetter, I. A. Coe, and J. C. Gerdes, "Handwheel force feedback for lanekeeping assistance: combined dynamics and stability," *ASME Journal of Dynamic Systems, Measurement, and Control*, vol. 128, no. 3, p. 532, 2006.

[6] M. Steele and R. B. Gillespie, "Shared control between human and machine: using a haptic steering wheel to aid in land vehicle guidance," *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, vol. 45, no. 23, pp. 1671–1675, Oct. 2001.

[7] T. Brandt, T. Sattel, and M. Bohm, "Combining haptic human-machine interaction with predictive path planning for lane-keeping and collision avoidance systems," in *2007 IEEE Intelligent Vehicles Symposium*. IEEE, Jun. 2007, pp. 582–587.

[8] H. Pacejka, *Tire and vehicle dynamics*, butterworth ed., 2012.

[9] C. E. Beal and J. C. Gerdes, "Model predictive control for vehicle stabilization at the limits of handling," *IEEE Transactions on Control Systems Technology*, vol. 21, no. 4, pp. 1258–1269, Jul. 2013.

[10] J. Mattingley and S. Boyd, "CVXGEN: a code generator for embedded convex optimization," *Optimization and Engineering*, vol. 13, no. 1, pp. 1–27, 2012.

[11] J. Mattingley, Y. Wang, and S. Boyd, "Code generation for receding horizon control," in *2010 IEEE International Symposium on Computer-Aided Control System Design (CACSD)*, 2010, pp. 985–992.

[12] A. Balachandran, S. M. Erlien, and J. C. Gerdes, "The virtual wheel concept for supportive steering feedback during active steering interventions," *ASME 2014 Dynamic Systems and Control Conference, San Antonio, TX*, 2014.

[13] M. J. Jensen, A. M. Tolbert, J. R. Wagner, F. S. Switzer, and J. W. Finn, "A customizable automotive steering system with a haptic feedback control strategy for obstacle avoidance notification," *IEEE Transactions on Vehicular Technology*, vol. 60, no. 9, pp. 4208–4216, Nov. 2011.