

Minimal Parametric Sensitivity Trajectories for Nonlinear Systems

Alex Ansari and Todd Murphey

Abstract—Trajectory optimization techniques can be used to minimize error norms on both state and controls with respect to a reference. However, without sufficient control authority, resulting controllers can be sensitive to variations in model parameters. This paper presents a method to incorporate parametric sensitivity in optimal control calculations to develop optimally insensitive trajectories which can be better tracked under open loop and limited gain feedback conditions. As it builds on existing nonlinear optimal control theory, the approach can be easily implemented and applies to a variety of system types. The effectiveness of the technique is demonstrated using a simplified vehicle model. Controllers developed are capable of tracking highly aggressive and dynamically infeasible paths while minimizing sensitivity to variations in modeled friction.

I. INTRODUCTION

For linearly controllable systems, sufficiently high gain feedback can be used to track arbitrary system trajectories. However, it is not always feasible to implement such controllers due to limitations on control authority. While well designed feedforward terms can improve tracking of reference trajectories and reduce feedback requirements, these strategies perform poorly when real world conditions vary from those modeled. This paper introduces a sensitivity optimization technique which addresses these concerns by producing optimal trajectories which make minor sacrifices in tracking a desired path in order to minimize a norm on parametric sensitivity. By incorporating these terms into optimal control calculations, the approach results in alternative trajectories which are less sensitive to fluctuations in model parameters when tracked using feedforward control.¹As a result, better tracking can be achieved using significantly less feedback.

The research presented in this paper builds on existing work in sensitivity optimization and optimal control. In [2], ensemble control is utilized to develop approximate control strategies which steer a unicycle subject to bounded model perturbations. [13] demonstrates the potential for improved feedforward control to enhance disturbance rejection for controllers with both feedback and feedforward terms. [3] and [9] introduce methods similar to those in this paper to optimize trajectory sensitivity for linear time invariant systems. [4] and [5] derive H_∞ methods for controller sensitivity minimization for linear and time discretized systems. While additional examples can be cited, the existing sensitivity optimization techniques identified rely on specific

topological or structural system properties, linear dynamics, or are in other ways limited in generality.

A primary benefit of the parametric sensitivity optimization approach presented in this paper is that it remains applicable to a broad class of linear and nonlinear, time varying systems. It applies to continuous time state and control signals i.e. the method does not require *a priori* time discretization. Additionally, it builds on existing nonlinear optimal control techniques and theory by modifying the form of the state space vector but otherwise leaves the existing procedure largely unchanged. As such, the methods discussed can be easily implemented and generally applied.

As an application of the sensitivity optimization approach, a simple model of a vehicle is used. We show that the techniques presented are capable of producing optimally insensitive vehicle driving styles and associated control laws with improved tracking performance under limited feedback conditions. Following, Section II provides a description of this vehicle model. Section III presents the iterative, nonlinear optimal control technique leveraged to provide feasible solutions for trajectories and controllers at each iteration. Section IV describes how these methods can be adapted to minimize a norm on parametric sensitivity. Finally, Sections V and VI discuss results obtained applying sensitivity optimization to the vehicle example under conditions of varying friction and available feedback.

To illustrate the difference between an optimal trajectory derived using standard trajectory optimization techniques and one derived using the sensitivity optimization methods discussed, Figures 1a and 1b are presented below. In each case, the dotted blue curves reflect optimal trajectories for the model vehicle's center of geometry as it tracks the dynamically infeasible desired trajectory represented by the dotted gray curve. Surrounding curves represent simulated trajectories which would result if the optimal feedforward control laws associated with each method were applied under levels of friction varying from 0.55 to 1.75 times the nominal value of 0.7. As reflected in Figure 1b, less sensitive trajectories result in tighter groupings of all these curves because they minimize an L_2 error norm between the optimal trajectory produced under nominal conditions and those due to varying model conditions. Further details and a discussion of these figures are presented in Sections V and VI.

Alex Ansari and Todd Murphey are with the Department of Mechanical Engineering, McCormick School of Engineering and Applied Science, Northwestern University, 2145 Sheridan Road, Evanston, IL 60208, USA, alexanderansari2011@u.northwestern.edu, t-murphey@northwestern.edu

¹Unless otherwise noted the optimal control techniques discussed in this paper are local methods. Therefore, optimal trajectories and controllers discussed are locally optimal with respect to the cost functionals used.





Fig. 1: Standard state based trajectory optimization (1a) and sensitivity (1b) optimization techniques tracking a desired trajectory using feedforward control. Orange, green, purple, red, brown, and black reflect trajectories which would result if actual friction were 55%, 70%, 85%, 125%, 150%, and 175% of the nominal modeled value of 0.7.

II. DYNAMIC MODELING

The vehicle model used to demonstrate the optimization and control techniques discussed in this paper is based on the robotic variable inertial vehicle (VIV) presented in [11]. The simplified version of this robotic vehicle used in this paper is parallel steered with four wheels modeled in point contact with the ground subject to viscous friction. The model includes a controllable mass which can be accelerated between its front and back end along the vehicle's central axis. Adjusting the placement of this mass modifies the vehicle's overall inertia and the normal forces at the tire contacts as it drives.

The vehicle controls are defined as the forward/backward thrust developed by the tires, $F(t)$, a kinematic control for the steering angle, $\Psi(t)$, and the acceleration of the adjustable mass in one dimension between its front and rear, $\nu(t)$. These are represented by the first three elements of the control vector $U(t)$ respectively. The state space for the model, reflected in Figure 2, consists of its $(x(t), y(t))$ location and heading angle $\theta(t)$ provided relative to a fixed world frame, the displacement of the mass center relative to the vehicle's geometry center, $M_c(t)$, and the derivatives of these states. The state space vector, $X(t)$, and control vector, $U(t)$, are

$$X(t) = [x(t), y(t), \theta(t), M_c(t), \dot{x}(t), \dot{y}(t), \dot{\theta}(t), \dot{M}_c(t)]^T$$

and

$$U(t) = [F(t), \Psi(t), \nu(t)]^T.$$

The dynamic equations of motion are derived using a Lagrangian formulation and solving the forced Euler Lagrange (EL) equations to incorporate the viscous friction forces. The equations are too unwieldy to be included but can be calculated using standard symbolic processing software. In




Fig. 2: Dynamic vehicle model with configuration variables

the shorthand version of the dynamics vector copied below $EL_{\ddot{x}(t)}$, $EL_{\ddot{y}(t)}$, and $EL_{\ddot{\theta}(t)}$ represent the solutions for $\ddot{x}(t)$, $\ddot{y}(t)$, and $\ddot{\theta}(t)$ derived from the forced EL equations.

$$\dot{X}(t) = [\dot{x}(t), \dot{y}(t), \dot{\theta}(t), \dot{M}_c(t), EL_{\ddot{x}(t)}, EL_{\ddot{y}(t)}, EL_{\ddot{\theta}(t)}, \dot{M}_c(t)]^T$$

III. THE OPTIMAL CONTROL ALGORITHM

The following section formulates a method to develop locally optimal controllers capable of driving a dynamical system along a desired trajectory while minimizing a norm on the state error, and applied controls. To these ends a projection based nonlinear optimal control technique detailed in [6], [7], and [12] is utilized to perform iterative optimization

of a cost functional of the form

$$J = \int_{T_0}^{T_f} l(X(t), U(t)) dt + m(X(T_f)) \quad (1)$$

with respect to $\xi(t) = (X(t), U(t))$, constrained by the dynamics

$$\dot{X}(t) = f(X(t), U(t)). \quad (2)$$

At each iteration of optimization a local quadratic model subject to a locally linear constraint is used to approximate the cost functional. Where $\zeta(t) = (z(t), v(t))$ represents perturbations to the state and control elements of the trajectory $\xi(t)$, local approximations are derived using the quadratic model given by²

$$g(\zeta) = DJ(\xi) \cdot \zeta + \frac{1}{2} \int_{T_0}^{T_f} \|\zeta\|^2 dt \quad (3)$$

subject to the constraint that the perturbations locally obey the linearized dynamics

$$\dot{z}(t) = A(t)z(t) + B(t)v(t). \quad (4)$$

Minimizing the quadratic model above with respect to its argument, ζ , provides the direction of steepest descent in the cost functional. This minimization can be calculated as the solution to an LQR problem for which there are well known methods of optimization (see [1]). As such, obtaining the steepest descent direction at each iteration is straightforward.

Applying a standard iterative descent technique, at each iteration the previous best trajectory is perturbed in the descent direction ζ . However, in doing so the resulting trajectories diverge from the manifold of feasible trajectories, \mathcal{T} . By using cost functional (1) one can set up and then solve a second LQR problem in order to calculate an optimal feedback control law $K(t)$ (see [6]). This control law can then be applied in the projection given by equations (2) and

$$U(t) = \mu(t) + K(t)[\alpha(t) - X(t)] \quad (5)$$

to project each infeasible perturbed trajectory, $(\alpha(t), \mu(t))$, back onto \mathcal{T} .

This approach to nonlinear optimal control returns trajectories $\xi(t)$ on the manifold of feasible trajectories at the end of each iteration rather than requiring the optimization terminate before a feasible solution is returned. Thus, the optimization process can be halted at any time to return the current best feasible solution, $(X(t), U(t))$,³necessary for open loop control. Additionally, it provides each of the corresponding optimal feedback control laws $K(t)$. For more detail on these methods refer to [6] and [7].

² $DJ(\xi)$ in equation (3) refers to the slot derivative of $J(\xi)$ with respect to its argument. More generally $D_n F(arg_1, arg_2, \dots, arg_n)$ refers to the slot derivative of function F with respect to its n^{th} argument.

³After the optimal control calculations converge, $U(t)$ represents the locally optimal feedforward control law associated with locally optimal state trajectory $X(t)$.

Optimal Control Algorithm:

- Approximate cost functional using a local quadratic model
- Minimize the local model by solving an LQR problem to obtain $\zeta(t)$
- Minimize a 2^{nd} quadratic model to obtain optimal feedback gain $K(t)$
- Perturb $\xi(t)$ by $\zeta(t)$
- Use $K(t)$ to project perturbed $\xi(t)$ onto \mathcal{T}
- Apply Armijo line search to scale $\zeta(t)$ until the perturbed and projected $\xi(t)$ provides sufficient decrease
- Return $\xi(t)$ as current best solution
- Repeat until convergence

A. State Based Trajectory Optimization

These optimal control methods can be applied to the vehicle model to produce locally optimal trajectories which minimize applied controls and tracking errors developed in following a desired path. To accomplish this, it is necessary to define appropriate $l(X(t), U(t))$ and $m(X(T_f))$ terms in cost functional (1).

As formulated, the term $l(X(t), U(t))$ is an incremental cost term designed to be integrated over the time horizon. For this example, it can represent a combined norm on the state tracking error and applied controls as illustrated in

$$l(X(t), U(t)) = \frac{1}{2} (X(t) - X_d(t))^T \cdot Q \cdot (X(t) - X_d(t)) + \frac{1}{2} (U(t) - U_d(t))^T \cdot R \cdot (U(t) - U_d(t)). \quad (6)$$

The term $m(X(T_f))$ incorporates a terminal cost into equation (1). It can provide a norm on the state tracking error at the terminal time in the form

$$m(X(T_f)) = \frac{1}{2} (X(T_f) - X_d(T_f))^T \cdot P \cdot (X(T_f) - X_d(T_f)). \quad (7)$$

Plugging these in and optimizing, one can produce optimal controllers capable of driving the model along highly aggressive and dynamically infeasible trajectories.

Through the choice of Q, R, and P matrices, this cost functional also provides a straightforward means for designers to influence trade-offs made during optimization. Interestingly, the process allows for the selection of locally optimal trajectories and controllers which follow distinctly different driving styles. For the vehicle model, two main driving styles emerge. In the first, controllers maintain higher speed and forward orientation around sharp hair-pin turns. In the second, controllers slow the vehicle down and back into sharp turns to better approximate corners along the desired path.

IV. INCORPORATING PARAMETRIC SENSITIVITY

The cost functional and optimization approach previously described provide control laws with both feedforward and full state feedback terms capable of optimally driving the vehicle along a desired trajectory under modeled conditions.

However, model parameters such as the coefficient of friction, tire diameter, etc. are only approximations. There is no guarantee that controllers calculated in this fashion will be capable of driving a vehicle in a real environment where significant variations in model parameters are common and where there are limitations on the availability and magnitude of feedback which can be applied.

To address these concerns this section derives a strategy to incorporate and minimize a norm on parametric sensitivity in the optimization process that can be applied to optimize both linear and nonlinear dynamical systems. The approach selects for optimally insensitive trajectories which can be closely followed in spite of parametric fluctuations using only open loop control. As these trajectories lend themselves to more accurate feedforward tracking, significantly less feedback is required to achieve the same tracking performance as optimal controllers using feedforward and feedback terms to follow optimal trajectories derived using the standard method previously discussed.

This approach to sensitivity optimization requires an appropriate measure of sensitivity to minimize. The most natural measure of sensitivity to use in the proposed optimization is given by $\frac{dX(t)}{dp}$, which is the first order change in the system's state with respect to a change in the parameter of interest, $p \in \mathbb{R}$. Because $X(t)$ is calculated as the solution to the ODE given by (2) subject to initial condition $X(T_0) = X_0$, one means to calculate this derivative is to set up and solve a separate differential equation. Due to the property that mixed partial derivatives commute, the expression $\frac{d}{dt} \frac{dX(t)}{dp}$ is equivalent to $\frac{d\dot{X}(t)}{dp}$. Therefore, $\frac{dX(t)}{dp}$ can be calculated by taking the derivative of (2) with respect to the parameter and solving the resulting ODE, given by

$$\frac{d\dot{X}(t)}{dp} = D_1 f(X(t), U(t), p) \cdot \frac{dX(t)}{dp} + D_3 f(X(t), U(t), p), \quad (8)$$

subject to the initial condition $\frac{dX(t)}{dp} \Big|_{t=T_0} = \mathbf{0}$.

In the optimal control formulation discussed in Section III, the incremental cost functional (6) only requires solutions to its arguments, $X(t)$ and $U(t)$, to evaluate. As derived above, calculation of the sensitivity term depends on the forward integration of a differential equation of the dynamics and thus cannot be added as an additional term to the functional in a straightforward manner. Instead, the parametric sensitivity term can be appended to the state vector, $X(t)$, and the number of columns and rows of the Q matrix doubled to accommodate and provide the desired norm on the appended state vector. The appended state vector is denoted $\bar{X}(t) = [X(t), \frac{dX(t)}{dp}]^T$.

Appending to the state in this manner modifies the quadratic models needed to locally approximate the trajectory manifold, \mathcal{T} . Therefore, these changes also impact the solutions to the LQR problems required to obtain the feedback control law $K(t)$ and the descent direction $\zeta(t)$ during optimization. Only minor changes are required to include the appended state vector and modified Q matrix

when calculating $DJ(\xi) \cdot \zeta$ in (3) as the form of the cost functionals (6) and (7) remains unchanged. However, the linearizations of the dynamics required to locally constrain the solutions to the LQR problems require reformulation and the calculation of two new differential equations at each iteration.

Where the appended dynamics vector, $\bar{f}(X(t), U(t), p)$, is given by

$$\bar{f}(X(t), U(t), p) = \begin{pmatrix} f(X(t), U(t), p) \\ \frac{d\dot{X}(t)}{dp} \end{pmatrix}, \quad (9)$$

the linearization of the appended dynamics with respect to the appended state vector, $\bar{X}(t)$, and the control vector, $U(t)$, are given by

$$\bar{A}(t) = \begin{pmatrix} D_1 f(X(t), U(t), p) & \mathbf{0} \\ \frac{\partial}{\partial X(t)} \frac{d\dot{X}(t)}{dp} & D_1 f(X(t), U(t), p) \end{pmatrix} \quad (10)$$

and

$$\bar{B}(t) = \begin{pmatrix} D_2 f(X(t), U(t), p) \\ \frac{\partial}{\partial U(t)} \frac{d\dot{X}(t)}{dp} \end{pmatrix}, \quad (11)$$

respectively. The two new equations which emerge from the linearizations, $\bar{A}(t)$ and $\bar{B}(t)$, are obtained by taking the appropriate partial derivatives of (8). These required partial derivatives are reflected in

$$\frac{\partial}{\partial X(t)} \frac{d\dot{X}(t)}{dp} = D_1^2 f(X(t), U(t), p) \cdot \frac{dX(t)}{dp} + D_1 D_3 f(X(t), U(t), p) \quad (12)$$

and

$$\frac{\partial}{\partial U(t)} \frac{d\dot{X}(t)}{dp} = D_2 D_1 f(X(t), U(t), p) \cdot \frac{dX(t)}{dp} + D_2 D_3 f(X(t), U(t), p). \quad (13)$$

Plugging equations (12) and (13) into (10) and (11) provides linearizations of the appended dynamical system required to constrain the solutions to the LQR problems through each iteration of optimization. By appending the state vector with sensitivity terms in this manner, the existing projection based optimization approach can therefore be applied to sensitivity optimization without modification of the algorithm. Mentioned previously, the optimization routine is capable of computing trajectories which can be tracked in a manner that simultaneously minimizes a weighted L_2 norm on the state error, applied controls, and state sensitivity for general linear and nonlinear dynamical systems.

V. SIMULATION RESULTS

In the following section we compare the performance of the standard state based and appended state sensitivity optimization approaches (see Sections III-A and IV) in developing trajectories insensitive to parametric variations for the vehicle example introduced in Section II. The two methods will be referred to as the standard approach and the sensitivity optimization approach, respectively. Results demonstrate




Fig. 3: L_2 trajectory error versus friction as a percent of nominal design friction for controllers with feedforward and no (left), medium-gain (middle), and high-gain (right) feedback terms. Blue curves reflect standard state based optimization results and red curves reflect appended state sensitivity optimization results.

how well trajectories produced using each method can be tracked by their respectively derived control laws as both friction and the degree of available feedback vary.

All simulations and figures presented are calculated using the matrices

$$Q = \text{diag}[15, 15, 100, 45, \frac{1}{2}, \frac{1}{2}, \frac{1}{10}, 1], \quad (14)$$

$$R = \text{diag}[\frac{1}{4000}, \frac{900}{\pi}, 1], \quad (15)$$

and

$$P = \text{diag}[\frac{5}{2}, \frac{5}{2}, 50, 25, \frac{5}{2}, \frac{5}{2}, \frac{1}{4}, 2] \quad (16)$$

to provide norms on the state error and controls in the cost function in equation (1) for the standard optimization. The sensitivity optimization approach uses the same R matrix but appends the Q matrix with an 8×8 identity matrix to provide the norm on the 8 additional sensitivity terms in the state vector. For this example, the P matrix above is also appended with an 8×8 zero matrix to ignore the sensitivity cost at the final time.

Referring to Figures 1a and 1b in Section I, under modeled conditions each optimally derived controller tracks its respective optimal trajectory (dotted blue curves). As friction is varied from the nominal value of 0.7, controllers which follow less sensitive driving styles produce trajectories which vary less than their more sensitive counterparts. These controllers minimize an L_2 error norm between trajectories resulting from variations in friction and the optimal trajectory.

The plots in Figure 3 demonstrate the impact of varying friction on L_2 error for three different full state feedback scenarios. The leftmost plot reflects the scenario illustrated in Figure 1 in which controllers apply no feedback, using only the optimally derived feedforward control terms to drive the vehicle under various levels of friction. In the middle and rightmost plots similar results are presented for controllers using these same feedforward terms along with medium-gain and high-gain feedback terms respectively. Viewed from left to right, these plots demonstrate the impact of increasing feedback on tracking error for each method.




Fig. 4: L_2 trajectory error versus friction as a percent of nominal design friction for controllers with feedforward and medium-gain partial state feedback terms. The blue curve reflect standard state based optimization results and the red curve reflects sensitivity optimization results.

To simulate the different feedback conditions applied, feedback control laws are derived by minimizing local quadratic models of a cost functional similar to (1) in the manner discussed in Section III. By modifying terms in the R matrix used in this cost functional, the magnitude of allowable feedback can be influenced to select for controllers of desired gain. Applying Q and R matrices (14) and (15) produces the feedback controllers used in the high-gain case. Medium-gain conditions are simulated by scaling the R matrix by a factor of 50. To control for the impact of applied feedback gains on tracking performances between the two methods, optimal feedback controllers were calculated around the standard optimization trajectory only rather than deriving separate controllers for the standard and sensitivity derived optimal trajectories separately. The resulting controllers provide feedback for both optimization approaches in the various gain scenarios presented in Figure 3.⁴

Figure 4 provides L_2 error norms versus friction for the two methods when the optimal feedforward controllers are

⁴The L_2 error norms for sensitivity optimized trajectories were not significantly influenced when using feedback controllers derived specifically for these trajectories.

combined with partial state, medium-gain feedback terms. The figure presents tracking performance results under varying friction in a common scenario in which the robotic vehicle has access to data for position and velocity but no means to accurately measure heading angle for feedback control. To simulate these conditions, the medium-gain feedback law was modified to zero out columns associated with errors in both the $\theta(t)$ and $\dot{\theta}(t)$ state variables.

VI. DISCUSSION

The locally optimal trajectories reflected in Figures 1a and 1b result from feedforward controllers providing control inputs similar in peak magnitude. These figures and the L_2 errors in the leftmost plot of Figure 3 show that while the sensitivity optimized trajectory is slightly worse at tracking the desired trajectory, the L_2 errors between this optimal trajectory and those produced by variations in friction are significantly reduced. Under open loop conditions, these findings confirm that sensitivity optimized trajectories result in improved tracking for similar degrees of control authority.

Trends in L_2 errors also indicate that increasing feedback reduces controller tracking sensitivity in all cases. The middle plot of Figure 3 demonstrates that when medium-gain feedback is incorporated into the optimal feedforward control laws, the sensitivity optimized driving style continues to outperform the one derived using the standard approach, but less so than in the open loop case. As the level of available feedback is further increased, the error norms between the two approaches begin to converge. The rightmost plot shows that when feedback of sufficient magnitude can be applied, the difference in tracking performance between the two methods becomes insignificant.

These findings reflect the fact that when limited in gain, feedback can only partially compensate for trajectory errors. Driving styles produced by the sensitivity optimization approach use feedforward control to ensure that trajectories associated with significant variations in friction remain closer to the optimal. As such, less feedback is required to track these trajectories. When allowable feedback is increased further, feedforward terms become less significant until feedback control dominates and dictates tracking performance. Thus, with similar high-gain state feedback controllers, one can expect similar tracking performance between the two methods. In these cases sensitivity optimization is of little to no benefit as it results in slightly decreased tracking performance when following the gray desired trajectory and can increase computational complexity when deriving control laws.

The medium-gain, partial state feedback conditions simulated to produce Figure 4 demonstrate the benefits of the sensitivity optimization approach when full state feedback is not an option. Under all friction conditions simulated, sensitivity optimization trajectories resulted in superior tracking performance when compared to counterparts derived using the standard approach. These results are attributed to the fact that sensitivity optimized trajectories can be tracked more closely by their feedforward control terms. As

such, state errors associated with variables not included in feedback remain small enough for the partial state feedback controllers to be effective. In the example presented, errors in heading, $\theta(t)$, and angular velocity, $\dot{\theta}(t)$, are thus more significant when tracked by the feedforward laws associated with optimal trajectories derived by the standard approach. Because no feedback is available to correct for these errors in the partial feedback scenario, tracking performance is more significantly impacted.

VII. CONCLUSIONS

This paper introduces a method to incorporate and optimize a norm on parametric sensitivity in optimal control. Compared with existing techniques for sensitivity optimization, the method presented is more broadly applicable as it builds on an approach to optimal control which works for both linear and nonlinear systems and requires minimal assumptions. To demonstrate its effectiveness, the technique is applied to a dynamic vehicle model. Results show the approach successfully produces locally optimal control laws and driving styles that are less sensitive to fluctuations in modeled friction.

While high-gain feedback controllers may be robust with respect to parameter variations, a benefit of the sensitivity optimization techniques presented is that they result in locally optimal trajectories which can be reliably tracked under limited feedback conditions. Optimal trajectories developed using these techniques result in controllers applicable to a much wider assortment of physical systems and which remain effective in spite of the inevitable variations and constraints associated with real world conditions.

ACKNOWLEDGMENT

This material is based upon work supported by the National Science Foundation under Grant IIS 1018167. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

REFERENCES

- [1] Brian D. O. Anderson and John B. Moore. *Optimal control: linear quadratic methods*. Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 1990.
- [2] A. Becker and T. Bretl. "Approximate steering of a unicycle under bounded model perturbation using ensemble control". *IEEE Transactions on Robotics*, 28(3):580–591, June 2012.
- [3] P. Byrne and M. Burke. "Optimization with trajectory sensitivity considerations". *IEEE Transactions on Automatic Control*, 21(2):282–283, Apr. 1976.
- [4] Hong Chen, Miao-Miao Ma, Hu Wang, Zhi-Yuan Liu, and Zi-Xing Cai. "Moving horizon H_∞ tracking control of wheeled mobile robots with actuator saturation". *IEEE Transactions On Control Systems Technology*, 17(2):449–457, March 2009.
- [5] Caio Fernandes de Paula and Luís Henrique Carvalho de Ferreira. "An easy-to-use H_∞ /LTR control solution with mixed-sensitivity properties". *IEEE Transactions on Automatic Control*, 56(7):1709–1713, July 2011.
- [6] J. Hauser and D.G. Meyer. "The trajectory manifold of a nonlinear control system". In *Proceedings of the 37th IEEE Conference on Decision and Control*, volume 1, pages 1034–1039 vol.1, 1998.

- [7] John Hauser. "A projection operator approach to the optimization of trajectory functionals". In *IFAC World Congress*, Barcelona, Spain, Jul. 2002.
- [8] Chris T. Kelley. *Iterative Methods for Optimization*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 1999.
- [9] E. Kreindler. "Formulation of the minimum trajectory sensitivity problem". *IEEE Transactions on Automatic Control*, 14(2):206 – 207, Apr. 1969.
- [10] Richard M. Murray, Zexiang Li, and S. Shankar Sastry. *A Mathematical Introduction to Robotic Manipulation*. CRC Press, Boca Raton, FL, 1994.
- [11] Chenghui Nie, S.C. Van Dooren, J. Shah, and M. Spenko. "Execution of dynamic maneuvers for unmanned ground vehicles using variable internal inertial properties". In *Proceeding of IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4226 – 4231, Oct. 2009.
- [12] A. Rucco, G. Notarstefano, and J. Hauser. "Dynamics exploration of a single-track rigid car model with load transfer". In *Proceedings of the 49th IEEE Conference on Decision and Control (CDC)*, pages 4934 – 4939, Dec. 2010.
- [13] N. C. Singer and W. P. Seering. "Preshaping command inputs to reduce system vibration". *Journal of Dyanamic Systems, Measurement, and Control*, 112(2):76 – 82, 1990.