Published in final edited form as:

IEEE Trans Robot. 2014 December 5; 30(6): 1358–1370. doi:10.1109/TRO.2014.2345918.

Trajectory Synthesis for Fisher Information Maximization

Andrew D. Wilson, IEEE [Student Member], Jarvis A. Schultz, IEEE [Student Member], and
Todd D. Murphey, Member [IEEE]

Department of Mechanical Engineering, Northwestern University, Evanston, IL 60208 USA

Andrew D. Wilson: awilson@u.northwestern.edu; Jarvis A. Schultz: jschultz@u.northwestern.edu; Todd D. Murphey: t-murphey@northwestern.edu

Abstract

Estimation of model parameters in a dynamic system can be significantly improved with the choice of experimental trajectory. For general nonlinear dynamic systems, finding globally “best” trajectories is typically not feasible; however, given an initial estimate of the model parameters and an initial trajectory, we present a continuous-time optimization method that produces a locally optimal trajectory for parameter estimation in the presence of measurement noise. The optimization algorithm is formulated to find system trajectories that improve a norm on the Fisher information matrix (FIM). A double-pendulum cart apparatus is used to numerically and experimentally validate this technique. In simulation, the optimized trajectory increases the minimum eigenvalue of the FIM by three orders of magnitude, compared with the initial trajectory. Experimental results show that this optimized trajectory translates to an order-of-magnitude improvement in the parameter estimate error in practice.

Keywords

Maximum likelihood estimation; optimal control; parameter estimation

I. Introduction

THE design of trajectories for experimental identification of parameters in dynamic systems is an important problem in a variety of fields ranging from robotics to biology to chemistry and beyond. Improving model accuracy facilitates the design of quality controllers that can greatly increase system performance. When attempting to estimate a set of system parameters from observable data, the choice of experimental trajectory can have a significant effect on the precision of the parameter estimation algorithm; however, for nonlinear dynamic systems, the trajectory is constrained to nonlinear equations of motion, which leads to challenges in parameter estimation and trajectory optimization. Additionally, since the trajectories evolve on a continuous-time domain, it is important to ensure that the dynamics are satisfied throughout the entire time domain of any synthesized trajectory.

A variety of estimation techniques are used in practice including Kalman filtering, maximum likelihood estimators, and Monte–Carlo estimators [1]–[3]. This paper will focus solely on the problem of estimating static model parameters in nonlinear dynamic systems. A widely used method for estimating static model parameters is a maximum-likelihood estimator known as the batch least-squares estimator [3]. The batch least-squares method compares a set of measurements taken along the evolution of a trajectory with predicted observations of the system using the model equations and estimates of the parameters. This comparison is made using the sum of the least-squares error along the trajectory from which a new update to the parameter estimate is then calculated.

A fundamental quantity that affects how well the batch least-squares estimator performs with respect to parameter estimation is Fisher information. Finding a trajectory that maximizes the Fisher information of the model parameters will produce the best experimental estimate of the parameters, as indicated by the Cramer–Rao bound [4]. With these notions in mind, we will introduce a continuous-time optimization method that locally maximizes the Fisher information for a given nonlinear dynamic system.

A. Related Work

Since the design of an experimental trajectory has a wide range of potential uses, there have been a number of contributions to the area from different fields. A large amount of literature on optimal experimental design exists in the fields of biology [5]–[7], chemistry [8], and systems [9]–[12]. Many of these results focus on particular applications to experiments specific to their respective fields; however, the underlying principles of information theory remain the same.

A common metric used in these areas of experimental design—also the key metric in this paper—is the Fisher information matrix (FIM) computed from observations of the system trajectory [13]. Metrics on the Fisher information are used as a cost function in many optimization problems including work by Swevers on “exciting” trajectories [14]. This work, as well as related works [15], [16], synthesizes trajectories for nonlinear systems that can be recast as linear systems with respect to the parameters.

Further research has resulted in optimal design methods for general nonlinear systems. In work by Emery and Nenarokomov [17], least-squares and maximum-likelihood estimation techniques are combined with Fisher information to optimize the experimental trajectories. In this case and a number of others, the dynamics are solved as a discretized constrained optimization problem [18], [19].

Discretization of the dynamics *a priori* has several problems. First and foremost, an arbitrary choice about the discretization needs to be made. Second, adaptive time-stepping methods cannot be used, and a discretization appropriate for the initial trajectory cannot be expected to be appropriate for the final trajectory. Finally, discretization can lead to high-dimensional constrained optimizations (dimensions of 10^7 to 10^{12} are common in practical problems) that are impractical to solve numerically. To avoid discretization of continuous dynamics, a class of methods has been developed that relies on sets of basis functions to synthesize optimal controls for the system [20]–[22]. These methods allow the full trajectory to be

optimized on a continuous-time domain; however, the optimization problem is still subject to a finite set of basis function coefficients. One example of the basis function set includes the Fourier basis, which is used to create a class of trajectories over which the optimization can be performed [23].

In the robotics field, trajectory optimization and parameter identification algorithms have been developed for special classes of robotic systems. For serial robot arms and similarly connected systems, chain-based techniques and linear separation of parameters can be used [24], [25]. Techniques have also been adapted for parallel robots and manipulators [26]–[28]. While these techniques perform well for the intended class of robots, we seek an algorithm that has the ability to work on general nonlinear systems, only requiring differentiability of the dynamics and some form of control authority.

B. Contribution

The main contribution of this paper is the formulation of a continuous-time trajectory optimization algorithm that maximizes the information obtained from discrete observations during execution of a trajectory. The algorithmic approach we use is based on projection-based optimal trajectory tracking [29], [30]; however, this paper extends the algorithm to include a non-Bolza cost function—maximizing the information gained by observations of the dynamic trajectory. The extension of this algorithm is related to prior work by the authors on optimizing trajectories to improve estimation conditioning and convergence rates [31] with this paper presenting a derivation and results using Fisher information optimization to improve the accuracy of estimation. A preliminary version of this contribution without experimental validation was presented at the 2014 American Control Conference [32].

The benefit of using a variational technique is that a perturbation to the trajectory can be formulated in the continuous-time domain without the use of basis functions to define the trajectory. This allows for adaptive timestepping when numerically solving the differential equations governing the system dynamics. As an example, for the first iteration of the cart–pendulum example presented in Section V the average timestep used to compute the trajectory perturbation is 0.0051 s with a standard deviation of 0.0029 s; however, the minimum timestep used over the trajectory is 3.6×10^{-7} s. An algorithm discretizing the dynamics at a fixed timestep equal to this minimum timestep produces an optimization problem with $\sim n \times 10^7$ dimensions, making the calculation impractical. Subsequent iterations have varying timesteps with the minimum timestep of the final iteration at 7.4×10^{-6} s. Hence, optimizing control in a finite-dimensional fixed-timestep setting may yield intractable optimization problems compared with adaptive timestep techniques.

This paper is organized as follows. Section II introduces the problem formulation and estimation concepts required to compute the FIM for the system; Section III derives the required cost function and optimization routine for the optimization over the Fisher information; Section IV presents an overview of the experimental system that is used to validate our optimization routine; and Section V presents the experimental results. Detailed equations for the calculation of the descent direction are provided in the Appendix.

II. Problem Formulation

This paper considers estimation of a set of parameters in a system assumed to have noisy measurements but negligible process noise. This assumption will allow the expression for Fisher information to be simplified later in the paper into a closed form. Output variables may be a function of the states, controls, and parameter set. The noise from measurements is assumed to be independent and normally distributed. Thus, the model of the system is defined as

$$\dot{x}(t)=f(x(t), u(t), \theta) \quad y(t)=g(x(t), u(t), \theta)+w_y \quad (1)$$

where $x \in \mathbb{R}^n$ defines the system states, $y \in \mathbb{R}^h$ defines the measured outputs, $u \in \mathbb{R}^m$ defines the inputs to the system, $\theta \in \mathbb{R}^p$ defines the set of model parameters to be estimated, and w_y is the additive output noise, where $p(w_y) = N(0, \Sigma)$.

A. Least-Squares Parameter Estimation

Since normally distributed measurement noise is assumed on $y(t)$, a least-squares estimator is used to estimate the system parameters. This method is equivalent to maximum likelihood estimation due to the assumption of Gaussian noise [3]. Using a set of measurements, nonlinear batch least-squares estimation can be performed using either a gradient descent or Newton–Raphson search method.

The least-squares estimator can be written as

$$\hat{\theta} = \underset{\theta}{\operatorname{argmin}} \beta(\theta) \quad (2)$$

$$\beta(\theta) = \frac{1}{2} \sum_i^h (\tilde{y}(t_i) - y(t_i))^T \cdot \Sigma^{-1} \cdot (\tilde{y}(t_i) - y(t_i)). \quad (3)$$

$\tilde{y}(t_i)$ is the observed state at the i th index of h measurements, $\Sigma \in \mathbb{R}^{h \times h}$ is the covariance matrix associated with the sensor measurement error, and $\hat{\theta}$ is the least-squares estimate of the parameter set.

Given this estimator, we will use Newton's method with a backtracking linesearch to find optimal parameter values by minimizing the least-squares error in (3). To perform this optimization, the first and second derivatives of $\beta(\theta)$ w.r.t. θ must be calculated.

The first derivative can be computed by the following equation¹:

$$D_{\theta} \beta(\theta) = \sum_i^h (\tilde{y}(t_i) - y(t_i))^T \cdot \Sigma^{-1} \cdot \Gamma_{\theta}(t_i) \quad (4)$$

where

¹The notation $D_{\kappa} \alpha(\kappa)$ represents the partial derivative of α w.r.t. κ .

$$\Gamma_{\theta}(t_i) = D_x g(x(t_i), u(t_i), \theta) \cdot D_{\theta} x(x(t_i), u(t_i), \theta) + D_{\theta} g(x(t_i), u(t_i), \theta).$$

This equation requires the evaluation of $D_{\theta} x(x(t_i), u(t_i), \theta)$ of (1), which is computed by the following ordinary differential equation (ODE):

$$\dot{\psi}(t) = D_x f(x(t), u(t), \theta) \cdot \psi(t) + D_{\theta} f(x(t), u(t), \theta) \quad (5)$$

where

$$\psi(t) = D_{\theta} x(x(t), u(t), \theta) \in \mathbb{R}^{n \times p} \text{ and } \psi(0) = \{0\}^{n \times p}.$$

The Hessian of the least-squares estimator is also required to use Newton's method. Differentiating (4) w.r.t. θ yields

$$D_{\theta}^2 \beta(\theta) = \sum_i^h \Gamma_{\theta}(t_i)^T \cdot \Sigma^{-1} \cdot \Gamma_{\theta}(t_i) + (\tilde{y}(t_i) - y(t_i))^T \cdot \Sigma^{-1} \cdot [D_x g(\cdot) \cdot D_{\theta}^2 x(\cdot) + D_x^2 g(\cdot) \cdot \psi(t_i) + D_{\theta}^2 g(\cdot)]. \quad (6)$$

For compactness, the arguments $(x(t), u(t), \theta)$ are replaced by $(\cdot)$. Finally, the formulation for the Hessian requires $D_{\theta}^2 x(\cdot)$, which is given by the following ODE²:

$$\dot{\Omega}(t) = \left[D_x^2 f(\cdot) \cdot \psi(t) + D_{\theta} D_x f(\cdot) \right]^{T_{(1,3,2)}} \cdot \psi(t) + D_x f(\cdot) \cdot \Omega(t) + D_x D_{\theta} f(\cdot) \cdot \psi(t) + D_{\theta}^2 f(\cdot)$$

where

$$\Omega(t) = D_{\theta}^2 x(\cdot) \in \mathbb{R}^{n \times p \times p} \text{ and } \Omega(0) = \{0\}^{n \times p \times p}.$$

B. Estimation Algorithm

Using (4) and (6), parameter estimation is performed using either gradient descent or Newton steps based on the conditioning of the optimization. Typically, when far from an optimal solution, a few gradient descent steps are necessary before Newton steps will be effective.

Additionally, a backtracking linesearch method [33] is used to ensure that a sufficient decrease condition is satisfied for each iteration. For a well-conditioned problem, Newton steps will generally provide quadratic convergence rates with the stepsize $\gamma_j = 1$; however, in practice, ill-conditioned problems can cause convergence rates to degrade.

The estimation algorithm is outlined in Algorithm 1. Iterations continue until the gradient of the cost function $\beta(\theta)$ drops to a specified tolerance level.

²Since the equation for $\Omega(t)$ involves tensors, we will use the notation $[\cdot]^{T_{(1,3,2)}}$ to indicate a transpose of the second and third tensor dimensions.

Algorithm 1**Parameter Estimation**

Choose initial $\theta_0 \in \mathbb{R}^p$, tolerance ε , $j = 0$

while $D_0\beta(\theta_j) > \varepsilon$ **do**

 Calculate descent using gradient descent or Newton:

 Gradient descent: $d_j = -D_0\beta(\theta_j)$

 Newton: $d_j = -[D_\theta^2\beta(\theta_j)]^{-1} \cdot D_\theta\beta(\theta_j)$

 Compute γ_j with Armijo backtracking search

$\theta_{j+1} = \theta_j + \gamma_j d_j$

$j = j + 1$

end while

C. Fisher Information

Fisher information quantifies the amount of information a set of observations contains about a set of unknown parameters [13]. For dynamic systems, a set of measurements taken during the execution of a trajectory provides information about model parameters. The effectiveness of the parameter estimation technique outlined in the previous section is related to the FIM.

Assuming that the measurement noise of the system is normally distributed with zero process noise, the FIM for the system is given by

$$I(\theta) = \sum_i^h \Gamma_\theta(t_i)^T \cdot \Sigma^{-1} \cdot \Gamma_\theta(t_i). \quad (7)$$

Note that this expression has been simplified from the general form of Fisher information given the assumptions stated above. Details of the simplification can be found in [3].

In practice, increasing the Fisher information of the system tends to increase the maximum precision that can be obtained by an estimation algorithm; however, this is not explicitly guaranteed. This relationship is quantified by the Cramer–Rao bound [4] given by

$$\text{cov}_\theta(\hat{\theta}) \geq I(\theta)^{-1} \quad (8)$$

where $\hat{\theta}$ is the least-squares estimator defined in (2). While the Cramer–Rao bound provides the minimum possible covariance of the parameter estimate, it does not bound how large the covariance may be for a given trial.

To use the FIM as a metric for trajectory optimization, $I(\theta) \in \mathbb{R}^{p \times p}$ must be mapped to a scalar value. There is a significant amount of literature on different types of mapping choices [5], [17], [18]; however, this paper will restrict itself to the design choice of E-optimality. The choice of E-optimal design is used to improve the worst-case variances of the parameter set by maximizing the minimum eigenvalue of the FIM. Thus, the algorithm

creates a variation on the trajectory to proportionally improve the information acquired. Other experimental design approaches such as A- and D-optimality conditions, which provide different goals for the desired information acquisition, can be realized with modifications to the objective function. The following section presents the derivation of the E-optimal objective function and presents the iterative algorithm.

III. Nonlinear Trajectory Optimization

The trajectory optimization algorithm used in the results to follow is a projection-based optimal trajectory tracking algorithm, which has been extended to include Fisher information as an objective [29], [30]. This algorithm is formulated in continuous time to allow the time discretization of the dynamics to be easily decoupled from the measurement time points as well as allow for the use of any numerical integration method. Therefore, the objective including any costs on the Fisher information must be formulated in continuous time.

A. Objective Function

To define an objective function dependent on the FIM on a continuous-time domain, the maximization of the information matrix needs to be cast from a finite set of discrete measurements into an appropriate continuous analog. To satisfy this condition, the information equation (7) will be written as

$$\tilde{I}(\theta) = \int_{t_0}^{t_f} \Gamma_{\theta}(t)^T \cdot \Sigma^{-1} \cdot \Gamma_{\theta}(t) dt. \quad (9)$$

Assuming that observations are taken regularly along the entire trajectory, the optimal trajectory $x^*(t)$ maximizing the eigenvalues of the continuous $\tilde{I}(\theta)$ will approximately optimize the eigenvalues of the sampled $I(\theta)$. As the sampling rate increases, the values of $\tilde{I}(\theta)$ and $I(\theta)$ will converge.

If measurement only occurs along certain portions of the trajectory with predetermined times, a weighting function can be added to $\tilde{I}(\theta)$ to ensure that sensitivity is maximized specifically in the sampled areas of the trajectory. However, if the observation time is not predetermined, the sensitivity along the entire trajectory will be maximized using $\tilde{I}(\theta)$.

The optimization objective function will, therefore, be given by

$$J = \frac{Q_p}{\lambda_{\min}} + \frac{1}{2} \int_{t_0}^{t_f} \left[(x(t) - x_d(t))^T \cdot Q_{\tau} \cdot (x(t) - x_d(t)) + u(t)^T \cdot R_{\tau} \cdot u(t) \right] dt \quad (10)$$

where $\lambda_{\min}$ is the minimum eigenvalue of $\tilde{I}(\theta)$, Q_p is the information weight, $x_d(t)$ is a reference trajectory, Q_{τ} is a trajectory tracking weighting matrix, and R_{τ} is a control effort weighting matrix. The weights must be chosen such that $Q_p \geq 0$, Q_{τ} is positive semi-definite, and R_{τ} is positive definite.

The various weights allow for design choices in the optimal trajectory that is obtained. The requirements of positive definiteness and positive semidefiniteness of the weighting matrices are necessary to maintain a locally convex optimization problem including the fact that $\lambda \geq 0$ [30]. Increasing the control weight will result in less-aggressive trajectories, generally decreasing the obtained information. Using a reference trajectory allows for an optimal solution that remains in the neighborhood of a known trajectory.

While the cost function is generally well defined, it does become ill-conditioned and singular in the case of a zero eigenvalue. This will occur if the initially chosen trajectory yields no information about at least one unknown parameter. For a dynamic system, this may occur if a stationary trajectory is chosen, in which case, perturbing the initial trajectory of the system using a heuristic method will likely provide a sufficient initial trajectory for the algorithm. If a perturbation fails to provide any information, the system geometry or sensor configuration may prohibit the experiment from resolving the parameters simultaneously for any trajectory. In this case, sensors may need to be added, or unknown parameter removed to create a well-posed experiment.

B. Extended Dynamic Constraints

The optimal control algorithm was previously formulated for trajectory tracking problems where the objective function is explicitly a function of the system states [30]. However, given (10), the cost also depends on $\psi(t) = D_{\theta}x(t)$. Since the objective is to minimize a norm that includes $\psi(t)$, which depends nonlinearly on $x(t)$, $\psi(t)$ is treated as an additional state. Appending $\psi(t)$ to the state vector as an additional dynamic constraint allows for variations in $\psi(t)$ in the optimization algorithm. For convenience, the extended state will be defined by $\bar{x}(t) = (x(t), \psi(t))$, and $\eta(t) = (\bar{x}(t), u(t))$ defines a curve which satisfies the nonlinear system dynamics.

C. Projection Operator

The minimization of (10) is subject to the dynamic constraints given by (1) and (5). The minimization of this unconstrained optimization involves iteratively calculating a descent direction followed by a projection that maps the unconstrained trajectory, formed by the sum of the current iterate and the descent direction, onto the dynamic constraints as detailed in [29]. The projection operator uses a stabilizing feedback law to take a feasible or infeasible trajectory, which is defined by $\xi(t) = (\bar{\alpha}(t), \mu(t))$, and maps it to a feasible trajectory $\eta(t) = (\bar{x}(t), u(t))$.

The projection operator used in this paper is given by

$$P(\xi(t)): \begin{cases} u(t) = \mu(t) + K(t)(\bar{\alpha}(t) - \bar{x}(t)) \\ \dot{x}(t) = f(x(t), u(t)) \\ \dot{\psi}(t) = D_x f(x, u, \theta)^T \psi(t) + D_{\theta} f(x, u, \theta)^T. \end{cases}$$

The feedback gain $K(t)$ can be optimized as well by solving an additional linear quadratic regulation problem. Details of the optimal gain problem can be found in [29], but any stabilizing feedback may be used.

With the addition of the projection operator, the problem can be reformulated from an optimization over the constrained trajectory $\eta(t)$ to an optimization over the unconstrained trajectory $\xi(t)$. This relation is given by

$$\operatorname{argmin}_{\eta(t)} J(\eta(t)) \leftrightarrow \operatorname{argmin}_{\xi(t)} J(P(\xi(t))).$$

The unconstrained formulation allows variations of the trajectory to be calculated free of the constraint of maintaining feasible dynamics; however, the solution is projected to a feasible trajectory at each iteration of the optimization algorithm.

D. Optimization Routine

Algorithm 2 defines the iterative method using a gradient descent approach to solve the optimization problem. Each iteration requires a descent direction $\zeta_k(t) = (z, v)$ to be computed from the following equation [30]:

$$\zeta_k(t) = \operatorname{argmin}_{\zeta_k(t)} DJ(P(\xi_k(t))) \circ \zeta_k(t) + \frac{1}{2} \langle \zeta_k(t), \zeta_k(t) \rangle \quad (11)$$

such that

$$\dot{\bar{z}} = A\bar{z} + Bv$$

where $\zeta_k(t) \in T_{\eta_k} \mathcal{T}$ i.e., the descent direction for each iteration lies in the tangent space of the trajectory manifold at η_k . The components of the descent direction $\zeta_k = (z(t), v(t))$ are defined by $z(t)$, i.e., the perturbation to the extended state, and $v(t)$, i.e., the perturbation to the control. The quantity $\langle \zeta_k(t), \zeta_k(t) \rangle$ is a local quadratic model computed as an inner product of $\zeta_k(t)$. Matrices A and B are the linearizations of the extended state dynamics. Since (11) is a quadratic function of ζ_k with linear constraints, the descent direction can be computed using LQR techniques, which depends on the linearization of the cost function $DJ(P(\xi_k(t)))$ and the local quadratic model $\langle \zeta_k(t), \zeta_k(t) \rangle$. The detailed calculation of the descent direction is provided in the Appendix.

Given the descent direction ζ_k , a backtracking linesearch of the projection, i.e., $P(\eta_k(t) + \gamma_k \zeta_k)$, provides a feasible trajectory assuming the step size γ_k satisfies the Armijo sufficient decrease condition [33]. Iterations upon the feasible trajectories continue until a given termination criteria is achieved.

Algorithm 2

Trajectory Optimization

Initialize $\eta_0 \in \mathcal{T}$, tolerance ε , $k = 0$
while $DJ(\eta_k(t)) \circ \zeta_k > \varepsilon$ **do**
 Calculate descent, ζ_k

$$\zeta_k = \arg \min_{\zeta_k(t)} DJ(P(\xi_k(t))) \circ \zeta_k + \frac{1}{2} \langle \zeta_k, \zeta_k \rangle$$

Compute γ_k with Armijo backtracking search

Calculate the infeasible step:

$$\xi_k(t) = \eta_k(t) + \gamma_k \zeta_k$$

Project trajectory onto dynamics constraints:

$$\eta_{k+1}(t) = P(\xi_k(t))$$

$k = k + 1$

end while

IV. Experimental Setup

To demonstrate the optimization of Fisher information for a dynamic system, a simulation and experimental test of a two-link cart–pendulum system is considered. The system has three configuration variables, $q = (x(t), \phi_1(t), \phi_2(t))$, where $x(t)$ is the horizontal displacement of the cart, and $\phi_i(t)$ is the rotational angle of each link as seen in Fig. 1.

The control input to the system is the acceleration of the cart, which is given by u . The cart can accelerate in either direction with positive acceleration to the right. Rotational friction is modeled at each pendulum joint, but the joints remain unactuated. The goal of the optimization algorithm is to accurately estimate the mass of the top pendulum link and the damping coefficient of the joints.

A. Experimental Testbed Setup

1) Robotic System—The experimental setup shown in Fig. 2 consists of a differential drive mobile robot with magnetic wheels moving in a plane. The pivot point for the first link of the double pendulum is provided by mounts on the robot. A plastic membrane provides the robot’s driving surface. The membrane is tensioned in all directions within an aluminum frame mounted parallel to the ground at a height of approximately 2.5 m. An unpowered magnetic idler mechanism is placed on the top side of the plastic to provide adherence for the robot. The idler has the same geometric footprint as the robot, but its magnetic wheels have opposite polarity.

Two 24-V DC motors drive the robot’s magnetic wheels. PID loops running at 1500 Hz close the loop around motor velocity using optical encoders for feedback. Two 12-V lithium iron phosphate batteries provide all required power. Desired velocity commands are sent to the robot wirelessly using Digi XBee modules at 50 Hz, and a 32-bit Microchip PIC microprocessor handles all on-board processing, motor control, and communication. The motors are powerful relative to the inertias of the robot, the double-pendulum payload, and the wheels. Thus, they are capable of accurately tracking aggressive trajectories up to a maximum rotational velocity. Since the input to the two-link cart–pendulum system is the direct acceleration of the cart, the PID loops on the motor velocity allow the robot to accurately track a given velocity profile. The Robot Operating System (ROS) is used for

visualization, interpreting, and transmitting desired trajectories and for processing and recording all experimental data [34].

2) Tracking System—A Microsoft Kinect is employed to track the system and obtain experimental measurements. The Kinect provides uncolored point clouds at 30 Hz as seen in Fig. 3, and the Point Cloud Library (PCL) is used for data processing [35]. Each raw point cloud is first downsampled, and then, pass-through filters eliminate points that lie outside of a predetermined bounding box of expected system configurations. A Euclidean cluster extraction algorithm is then used to extract three separate clouds—one for the robot and one for each of the links. The pendulum links are made from an acrylic that is transparent to the Kinect; opaque markers adhered to each of the links provide a visible surface. The markers have a gap between them, ensuring that the software uniquely detects the three clusters. The coefficients of the line equations along the axis of each link are then extracted using a sample consensus segmentation filter on the clusters representing the links. Once the coefficients for these lines are known, simple trigonometry allows calculation of the desired link angles, ϕ_1 and ϕ_2 .

Since the Kinect will provide measurements of the absolute rotation angle of each pendulum link, it is important to estimate the variance of the measurements for use in the optimization algorithm. To measure the covariance of the angles computed from the point cloud, measurements were recorded for 30 s at sensor's fixed frequency of 30 Hz with the double pendulum hanging in a stationary position. Although the variance of the measurements will have some dependence on the link velocity and configuration, we assume independence for the purpose of this experiment. The angles calculated from the point cloud are assumed to be independent for the two links; therefore, the off-diagonal elements are set to zero, yielding a covariance matrix given by

$$\Sigma = \begin{bmatrix} 1.12 \times 10^{-4} & 0 \\ 0 & 4.79 \times 10^{-4} \end{bmatrix} \text{rad}^2. \quad (12)$$

B. System Model

A mathematical model of the system is created by deriving standard rigid body dynamics equations. A Lagrangian $L = T - V$ is constructed where the kinetic energy of the system T is

given by $T = \frac{1}{2} \sum_{i=1}^2 m_i v_i(x, \phi_1, \phi_2)^2 + \mathcal{I}_i \omega_i(x, \phi_1, \phi_2)^2$, where v_i is the translational velocity, and ω_i is the rotational velocity of each link. The potential energy V is a function of

the height z_i of each link's center of mass, i.e., $V = \sum_{i=1}^2 m_i g z_i(x, \phi_1, \phi_2)$, given the gravitational constant g . For each link, the moment of inertia $\mathcal{I}_i$ is defined as a function of the mass and measured center of mass location. The center of mass on the link is also used to define the potential energy of the link in the Lagrangian.

Using the Euler–Lagrange equation

$$0 = \frac{\partial L}{\partial q} - \frac{d}{dt} \left(\frac{\partial L}{\partial \dot{q}} \right)$$

the full equations of motion can be solved in terms of $q(i)$ (keeping in mind that $\ddot{x} = u$ by assumption).

The Kinect measures the absolute angle y_i of each link i , which provides two measurement outputs, $y(t) = [y_1(t), y_2(t)]$. Since the system states define the angle of link 2 relative to link 1, the output function is

$$g(x(t), u(t), \theta) = \begin{bmatrix} \phi_1(t) \\ \phi_1(t) + \phi_2(t) \end{bmatrix}.$$

C. System Parameter Selection

We have not yet addressed the choice of which parameters to identify for a given system. Choosing the set of parameters will depend on the context of the estimation needs as well as concerns with identifiability. In certain situations, some parameters may not be identifiable, regardless of trajectory selection. For example, given the two-link cart-pendulum system with no damping in the joints, estimating both link masses independently is not possible. To illustrate this, Fig. 4 shows a contour plot of β , i.e., the parameter estimator cost, versus changes in each mass value. The straight isolines in the plot indicate that the parameters are coupled through a constant scaling factor resulting in no unique set of mass parameters that minimize the cost. Rather, the estimator cost can be minimized with an infinite linear combination of the mass parameters; therefore, the specific parameter set is nonidentifiable by the estimator.

The nonidentifiability condition can be tested for a given trajectory by examining the eigenvalues of the Fisher information matrix. If and only if a zero eigenvalue exists, neither parameter can be uniquely identified. If a zero eigenvalue exists for the trajectory, another trajectory should be tested for identifiability. If no trajectory can be found, the algorithm cannot be initialized, and a different set of parameters should be chosen. The experiment that follows in this paper is performed under the assumption that the two unknown system parameters are the mass of the top link m_1 and the damping coefficient of the top link joint and bottom link joint c . Since both links use the same bearings, we will assume that the damping coefficients of the two joints are equal. The remaining system parameters have been measured and are presented in Table I. Additionally, Table II lists the eigenvalues of the FIM for the initial trajectory. Since both eigenvalues are nonzero, the parameters can be simultaneously estimated to some precision, which will be discussed in Section VI-C.

Given that m_1 and c are the uncertain parameters, the following values will be used as initial estimates of the parameters:

$$m_1 = 0.085 \text{ kg}, c = 0.50 \text{ g/s.} \quad (13)$$

Note that the choice of units weights the desired precision between estimated parameters. In this case, the damping coefficient units are scaled to provide more relative precision in the mass estimate. The remaining system parameters that appear in Table I will be considered the actual known model parameters.

V. Results

A. Optimization Results

The optimization algorithm was run until a convergence criterion of $|DJ(\xi(t)) \circ \zeta| < 10^{-1}$ was satisfied, starting from an initial descent magnitude around 8×10^3 . The comparison of initial and optimized trajectories can be seen in Fig. 5.

Table II shows the initial and optimized eigenvalues of the FIM, $\tilde{I}(\theta)$, and the initial and optimized cost J . The results show that the minimum eigenvalue λ_2 increases by over a factor of 10^3 . Additionally, the other eigenvalue increases, although not included directly in the cost function.

Examining the plots of the optimized trajectory, it is clear that more information is gained by oscillating the pendulum back and forth. In particular, more information about the damping parameter c is gained by increasing the oscillation. This observation leads to a hypothesis that the cost function and optimization may be driven strongly by information concerning the damping parameter, although the result increases the overall amount of information for both parameters.

Using this optimized trajectory, simulation and experimental results of the system, as well as the experimental testbed, are presented in the following sections to demonstrate the improvement in the estimation of system parameters.

B. Monte–Carlo Simulation Analysis

Given the system model and parameters, results are first obtained in the simulation to assess the effectiveness of the algorithm given an exact model and pure Gaussian noise distributions. An experimental measurement of the trajectory is simulated by generating the trajectory of the deterministic system using our estimate of the parameters and sampling the trajectory at a discrete number of points adding Gaussian additive noise at each sample. This results in a discrete set of noisy measurements that will be used for the subsequent parameter optimization routine. For the simulation example, a 5-s trajectory is simulated at the Kinect's sampling rate of 30 Hz. The sampling rate is set by the frequency of sensors being used. The uncertainty of each state measurement is normally distributed with zero mean and variance equivalent to that measured by the Kinect during the testbed setup given by (12).

A Monte–Carlo simulation is performed using 300 trials, re-sampling the trajectory with new additive Gaussian noise samples for each trial. For the scope of the Monte–Carlo simulation analysis, the parameter estimates from (13) are treated as the actual system parameter set used to generate the simulated output, and the initial parameter estimate are uniformly varied in the range of $\pm 100\%$ of each parameter value in (13). This provides simulation results that are independent of the initial estimate.

After running the simulation, histogram plots in Fig. 6 show the distribution of parameter estimates using the optimized trajectory. The histogram of the optimized trajectory is shown by the red bins with a Gaussian fit to the mean and standard deviation shown by the solid red line. The dashed blue line indicates the Gaussian fit to the histogram results of the initial trajectory; however, to simplify the plot, the initial histogram is not shown. The averaged statistics of the simulation trials can be seen in Table III. Since the expected information has been significantly increased with the optimized trajectory, the results confirm that the covariance of the estimated parameters decreases dramatically. The precision of both parameter estimates with respect to the Lebesgue measure is improved by a factor of 10^3 .

C. Cramer–Rao Lower Bound

The estimator's covariance result should also be compared with the theoretical Cramer–Rao bound. As discussed in Section II the Cramer–Rao bound places an absolute lower bound on the variance of the parameter estimate that can be obtained using the batch least-squares estimator or other unbiased estimator. The bound is given in (8).

Table III lists the Cramer–Rao bounds for the initial and optimized trajectories. The covariance of the initial trajectory is clearly subject to a higher bound than that of the optimized trajectory. Due to round-off and other numerical errors in the algorithms and Monte–Carlo simulations, the covariance of the Monte–Carlo estimates is higher than the lower bound; however, overall remains within a factor of 3 of the predicted best-case variance estimates according to the Cramer–Rao bound.

D. Experimental Results

After validating the optimization results in simulation, the trajectories were tested on the experimental testbed to determine their experimental effectiveness. Each trajectory was run on the system with the observed angles and position of the robot recorded by the Kinect tracking system. The recorded data can be seen in Figs. 7 and 8.

Using the collected data as the reference for the least-squares parameter optimization, the parameter set estimate from (13) is used as the initial estimate for the parameter optimization routine. Using the batch-least-squares estimation method, the best estimates of the parameters were found based on the data collected. The results can be seen in Table IV. Actual values of the parameters for the experimental system were obtained by disassembling the pendulum system. The mass of each link was determined by individually weighing each link, and the damping coefficient was obtained from a batch-least-squares estimate of a single link in a free swinging trajectory.

Since the noise present in the initial trajectory is so high relative to the observed angles, the parameter estimation algorithm is not able to make any progress from the initial estimate of the parameters. This results in a 22.7% error in the mass and 316.7% error in the damping coefficient from the baseline values. The optimized trajectory provides parameter estimates with only 2.7% error in the mass and 17.2% error in the damping coefficient values.

A contour plot of the parameter estimator cost β is shown in Fig. 9. Using the experimental trajectory data, the cost was computed for a portion of the parameter space. The figure

shows a well-defined optimization basin for the optimized trajectory where the basin of the initial experimental trajectory is far less convex. This optimized contour plot illustrates why the data from the optimized trajectory provides the better estimate of the parameters of interest given the observed measurements. The E-optimality cost function, however, does not explicitly condition the optimized basin as seen in Fig. 9(b). In some cases, it may be advantageous to use an additional approach to improve the condition number of the optimization problem as well [31].

It is important to note that the error in the experimental case is far greater than the predicted Cramer–Rao bound in Table III. This is due to additional disturbances and unmodeled effects such as out-of-plane swinging motion, error in the control signal, and sensor nonlinearities. These effects will cause bias error in the estimator; however, even with these sources of bias, the optimized trajectory provides a much better estimate of model parameters than the initial system trajectory.

VI. Conclusion

This paper has presented a method to optimize the trajectory of a nonlinear system, maximizing the Fisher information with respect to a set of model parameters. The simulation results of the cart–pendulum simulation show that the optimization algorithm results in an increase of the minimum eigenvalue of the Fisher information as well as a decrease in the estimation variance from Monte–Carlo simulation. Additionally, experimental results confirm a substantial improvement in parameter estimation accuracy when using the optimized trajectory.

While the formulation of the algorithm allows for general nonlinear dynamics, limitations on tractable noise models remain a formidable challenge. Since process noise is assumed to be negligible for the purpose of this algorithm, results may only be useful on systems with accurate input models. Although the algorithm should improve prediction on most systems, predicted precision improvements in simulation may be overestimated compared with the experimental results due to unmodeled bias and system modeling errors. Still, with appropriate selection of weighting matrices on controls and tracking error, improvement in estimator performance is expected compared with arbitrary choices of the experimental trajectory.

Acknowledgment

This work was supported by the National Institute of Health under NIH R01 EB011615 and by the National Science Foundation under Grant CMMI 1334609. 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.

Appendix

Computing the Descent Direction: $\zeta_k(t)$

To find a descent direction for the optimal control algorithm, (11) must be solved. As shown in (11), the descent direction depends on the linearization of the cost function, i.e.,

$DJ(P(\xi_k(t)))$, and the local quadratic model, i.e., $\langle \zeta_k(t), \zeta_k(t) \rangle$. Using a quadratic model and expanding the linearizations of the cost function, (11) is rewritten as

$$\underset{\zeta_k(t)}{\operatorname{argmin}} = \int_{t_0}^{t_f} a(t)^T \bar{z}(t) + b(t)^T v(t) + \frac{1}{2} \bar{z}(t)^T Q_n \bar{z}(t) + \frac{1}{2} v(t)^T R_n v(t) dt \quad (14)$$

such that

$$\dot{\bar{z}} = A\bar{z} + Bv$$

where $a(t)$ and $b(t)$ are the linearizations of the cost function with respect to x and u , and Q_n and R_n are weighting matrices for the local quadratic model approximation. Design of these weighting matrices can lead to faster convergence of the optimal control algorithm depending on the specific problem. The derivations of the cost function linearizations and the dynamics linearizations are now presented.

A. Cost Function Linearization

The linearization of the cost function $DJ(P(\xi_k(t)))$ is found by taking the directional derivative of (10) with respect to the extended states $x(t)$ and the control vector $u(t)$.

The derivative of (10) with respect to x yields

$$\frac{\partial J}{\partial \bar{x}} = -\frac{Q_p}{\lambda_{\min}^2} \frac{\partial \lambda_{\min}}{\partial \bar{x}} + \int_{t_0}^{t_f} [(x(t) - x_d(t))^T \cdot Q_\tau] dt. \quad (15)$$

Since the cost function involves eigenvalues of $\tilde{f}(\theta)$, this linearization requires the calculation of the derivatives of eigenvalues. A process for this calculation was formalized by Nelson [36]. Given an eigensystem of the form

$$AX = X\Lambda$$

where Λ is a diagonal matrix of distinct eigenvalues $(\lambda_1, \lambda_2, \dots, \lambda_n)$, and X is the associated matrix of eigenvectors, the derivative of an eigenvalue λ_m is given by

$$D_x \lambda_m = \omega_m^T \cdot D_x A \cdot v_m$$

where ω_m is the left eigenvector, and v_m is the right eigenvector associated with λ_m .

If eigenvalues are not distinct, i.e., multiplicity greater than 1, the Nelson's method does not hold since the choice of eigenvectors is not unique. However, in the case of repeated eigenvalues and repeated eigenvalue derivatives, a set of eigenvectors can be determined up to a scalar multiplier, as detailed in [37]. Once the eigenvectors are calculated, each

eigenvalue and eigenvector pair can be used to compute the direction of steepest descent for the objective.

Using these methods to compute the eigenvalue derivative, $\frac{\partial \lambda_{\min}}{\partial \bar{x}}$ from (15) can be calculated. Taking the derivative of the eigenvalue of $\tilde{J}(\theta)$ from (9) with respect to the extended state yields

$$\frac{\partial \lambda_s}{\partial \bar{x}} = \omega_s^T \frac{\partial}{\partial \bar{x}} \left(\int_{t_0}^{t_f} \Gamma_\theta(t)^T \cdot \Sigma^{-1} \cdot \Gamma_\theta(t) \right) \nu_s$$

where s denotes the index of the minimum eigenvalue and eigenvector. Since the partial derivative and eigenvectors are evaluated only at the final time, the equation can be rewritten to a running cost formulation given by

$$\frac{\partial \lambda_s}{\partial \bar{x}} = \int_{t_0}^{t_f} \omega_s^T \frac{\partial}{\partial \bar{x}} \left(\Gamma_\theta(t)^T \cdot \Sigma^{-1} \cdot \Gamma_\theta(t) \right) \nu_s dt. \quad (16)$$

Finally, differentiating the inner product of the gradients yields

$$\frac{\partial}{\partial \bar{x}} \left(\Gamma_\theta(t)^T \cdot \Sigma^{-1} \cdot \Gamma_\theta(t) \right) = \begin{bmatrix} 2\Gamma_\theta(t)^T \cdot \Sigma^{-1} \cdot (D_x^2 g(\cdot) \cdot \psi(\cdot) + D_x D_\theta g(\cdot)) \\ 2\Gamma_\theta(t)^T \cdot \Sigma^{-1} \cdot D_x g(\cdot) \cdot E \end{bmatrix} \quad (17)$$

where E is a tensor of the form

$$E_{ijkl} = \delta_{ik} \delta_{jl}$$

with δ as the Kronecker delta function.

Combining (15)–(17), $a(t)$ in (14) is given by

$$a(t) = \begin{bmatrix} (x(t) - x_d(t))^T Q_\tau \\ \{0\}^{1 \times n \times p} \end{bmatrix} - \frac{Q_p}{\lambda_{\min}^2} \omega_s^T \begin{bmatrix} 2\Gamma_\theta(t)^T \Sigma^{-1} (D_x^2 g(\cdot) \psi(\cdot) + D_x D_\theta g(\cdot)) \\ 2\Gamma_\theta(t)^T \Sigma^{-1} D_x g(\cdot) E \end{bmatrix} \nu_s.$$

The linearization $b(t)$ from (14), defining the derivative of the cost function with respect to the controls $u(t)$, is given by

$$b(t) = u(t)^T \cdot R_\tau$$

where R_τ is the weighting matrix from (10).

B. Dynamics Linearization

The two other quantities needed to compute the descent direction are $A(t)$ and $B(t)$ —the linearizations of the dynamics. The descent direction ζ_k will satisfy the linear constraint ODE given by

$$\dot{\bar{z}}_k(t) = A(t)\bar{z}_k(t) + B(t)v_k(t)$$

where $A(t)$ is the linearization of the nonlinear dynamics given by (1) and (5) with respect to $x(\bar{t})$, and $B(t)$ is the linearization with respect to $u(t)$. The linearization $A(t)$ of the dynamics with respect to the extended state $x(\bar{t})$ is given by

$$A(t) = \begin{bmatrix} \frac{\partial \dot{x}}{\partial x} & \frac{\partial \dot{x}}{\partial \psi} \\ \frac{\partial \dot{\psi}}{\partial x} & \frac{\partial \dot{\psi}}{\partial \psi} \end{bmatrix} = \begin{bmatrix} D_x f(\cdot) & \{0\}^{n \times n \times p} \\ D_x^2 f(\cdot) \cdot \psi(t) + D_x D_\theta f(\cdot) & D_x f(\cdot) \cdot E \end{bmatrix}.$$

In addition to the state linearizations, the control linearizations are required. This linearization matrix $B(t)$ is given by

$$B(t) = \begin{bmatrix} \frac{\partial \dot{x}}{\partial u} \\ \frac{\partial \dot{\psi}}{\partial u} \end{bmatrix} = \begin{bmatrix} D_u f(\cdot) \\ D_u D_x f(\cdot) \cdot \psi(t) + D_u D_\theta f(\cdot) \end{bmatrix}.$$

Given the linearizations $a(t)$, $b(t)$, $A(t)$, and $B(t)$, (14) can be used to compute the descent direction $\zeta_k(t)$. At each iteration of the optimization algorithm, the perturbed trajectory $\eta_k(t) + \gamma_k \zeta_k$ must be projected to satisfy the dynamics. As shown in Algorithm 2, the process is repeated until a termination criteria is satisfied.

References

1. Kalman RE, Bucy RS. New results in linear filtering and prediction theory. J. Basic Eng. 1961; 83(1):95–108.
2. Ljung, L. System Identification: Theory for the User. 2nd ed.. Englewood Cliffs, NJ, USA: Prentice-Hall; 1999.
3. Bar-Shalom, Y.; Li, XR.; Kirubarajan, T. Estimation With Applications to Tracking and Navigation: Theory Algorithms and Software. New York, NY, USA: Wiley; 2004.
4. Rao CR, Wishart J. Minimum variance and the estimation of several parameters. Math. Proc. Cambridge Philos. Soc. 1947 Oct.43:280–283.
5. Faller D, Klingmuller U, Timmer J. Simulation methods for optimal experimental design in systems biology. Simulation. 2003 Dec.79:717–725.
6. Baltes M, Schneider R, Sturm C, Reuss M. Optimal experimental design for parameter estimation in unstructured growth models. Biotechnol. Progr. 1994 Sep.10:480–488.
7. Lindner PFO, Hitzmann B. Experimental design for optimal parameter estimation of an enzyme kinetic process based on the analysis of the Fisher information matrix. J. Theor. Biol. 2006; 238(1): 111–123. [PubMed: 16039672]
8. Franceschini G, Macchietto S. Model-based design of experiments for parameter precision: State of the art. Chem. Eng. Sci. 2008 Oct.63:4846–4872.
9. Dietrich, F.; Raatz, A.; Hesselbach, J. A-priori Fisher information of nonlinear state space models for experiment design; Proc. IEEE Int. Conf. Robot. Autom; 2010 May. p. 3698-3702.
10. Mårtensson J, Rojas CR, Hjalmarsson H. Conditions when minimum variance control is the optimal experiment for identifying a minimum variance controller. Automatica. 2011 Mar.47:578–583.
11. Joshi S, Boyd S. Sensor selection via convex optimization. IEEE Trans. Signal Process. 2009 Feb. 57(2):451–462.

12. Gevers M, Bombois X, Hildebrand R, Solari G. Optimal experiment design for open and closed-loop system identification. *Commun. Inform. Syst.* 2011; 11(3):197–224.
13. Lehmann, EL.; Casella, G. *Theory of Point Estimation*. New York, NY, USA: Springer; 1998.
14. Swevers J, Ganseman C, Tukul D, de Schutter J, Van Brussel H. Optimal robot excitation and identification. *IEEE Trans. Robot. Autom.* 1997 Oct.13(5):730–740.
15. Armstrong B. On finding exciting trajectories for identification experiments involving systems with nonlinear dynamics. *Int. J. Robot. Res.* 1989 Dec.8:28–48.
16. Gautier M, Khalil W. Exciting trajectories for the identification of base inertial parameters of robots. *Int. J. Robot. Res.* 1992 Aug.11:362–375.
17. Emery AF, Nenarokomov AV. Optimal experiment design. *Meas. Sci. Technol.* 1998 Jun.9:864–876.
18. Mehra R. Optimal input signals for parameter estimation in dynamic systems—Survey and new results. *IEEE Trans. Automat. Control.* 1974 Dec.AC-19(6):753–768.
19. Vincent TL, Novara C, Hsu K, Poolla K. Input design for structured nonlinear system identification. *Automatica.* 2010 Jun.46:990–998.
20. Wahlberg, B.; Hjalmarsson, H.; Annergren, M. On optimal input design in system identification for control; *Proc. 49th IEEE Conf. Decision Control*; 2010 Dec.. p. 5548–5553.
21. Rackl, W.; Lampariello, R.; Hirzinger, G. Robot excitation trajectories for dynamic parameter estimation using optimized B-splines; *Proc. IEEE Int. Conf. Robot. Autom.*; 2012 May. p. 2042–2047.
22. Wu W, Zhu S, Wang X, Liu H. Closed-loop dynamic parameter identification of robot manipulators using modified Fourier series. *Int. J. Adv. Robot. Syst.* 2012 May.9:1–9.
23. Park K-J. Fourier-based optimal excitation trajectories for the dynamic identification of robots. *Robotica.* 2006 Mar.24:625–633.
24. Gautier M, Khalil W. Direct calculation of minimum set of inertial parameters of serial robots. *IEEE Trans. Robot. Autom.* 1990 Jun.6(3):368–373.
25. Radkhah, K.; Kulic, D.; Croft, E. Dynamic parameter identification for the CRS A460 robot; *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*; 2007 Oct.. p. 3842–3847.
26. Vivas A, Poignet P, Marquet F, Pierrot F, Gautier M. Experimental dynamic identification of a fully parallel robot. *Proc. IEEE Int. Conf. Robot. Autom.* 2003; 3:3278–3283.
27. Kopacek P, Lutz P, Pashkevich A, Wu J, Wang J, Wang L, Li T. Dynamic formulation of redundant and nonredundant parallel manipulators for dynamic parameter identification. *Mechatronics.* 2009; 19(4):586–590.
28. Farhat N, Mata V, Page A, Valero F. Identification of dynamic parameters of a 3-DOF RPS parallel manipulator. *Mech. Mach. Theory.* 2008; 43(1):1–17.
29. Hauser J. A projection operator approach to the optimization of trajectory functionals. *Proc. World Congr.* 2002 Jul.15:310.
30. Hauser J, Meyer D. The trajectory manifold of a nonlinear control system. *Proc. 37th IEEE Conf. Decision Control.* 1998; 1:1034–1039.
31. Wilson AD, Schultz JA, Murphey TD. Trajectory optimization for well-conditioned parameter estimation. *IEEE Trans. Autom. Sci. Eng.* 2015 Jan.12(1)
32. Wilson, AD.; Murphey, TD. Local e-optimality conditions for trajectory design to estimate parameters in nonlinear systems; *Proc. Amer. Control Conf.*; 2014 Jun.. p. 443–450.
33. Armijo L. Minimization of functions having Lipschitz continuous first partial derivatives. *Pacific J. Math.* 1966; 16(1):1–3.
34. Quigley M, Conley K, Gerkey B, Faust J, Foote T, Leibs J, Wheeler R, Ng AY. ROS: An open-source Robot Operating System. *Proc. ICRA Workshop Open Source Softw.* 2009; 3:1–6.
35. Rusu, RB.; Cousins, S. 3D is here: Point Cloud Library (PCL); *Proc. IEEE Int. Conf. Robot. Autom.*; 2011 May. p. 1–4.
36. Nelson RB. Simplified calculation of eigenvector derivatives. *AIAA J.* 1976 Sep.14:1201–1205.
37. Friswell MI. The derivatives of repeated eigenvalues and their associated eigenvectors. *J. Vib. Acoust.* 1996 Jul.118:390–397.

Biographies

Andrew D. Wilson (S'10) received the B.S. degree in mechanical engineering from Penn State University, University Park, PA, USA. He is currently working toward the Ph.D. degree in mechanical engineering with the Neuroscience and Robotics Laboratory, Northwestern University, Evanston, IL, USA.

In 2010, he received the National Defense Science and Engineering Graduate Fellowship from the U.S. Department of Defense. His research interests include modeling and control of dynamic systems using informative learning methods.

Jarvis A. Schultz (S'11) received the B.S. and M.S. degrees in mechanical engineering from the University of Wyoming, Laramie, WY, USA, and the Ph.D. degree in mechanical engineering from Northwestern University, Evanston, IL, USA.

He is currently the Assistant Director of the Master of Science in Robotics Program with Northwestern University, with research interests focused on embedded control for complex nonlinear systems.

Todd D. Murphey (S'02–M'03) received the B.S. degree in mathematics from the University of Arizona, Phoenix, AZ, USA, and the Ph.D. degree in control and dynamical systems from the California Institute of Technology, Pasadena, CA, USA.

He is an Associate Professor of mechanical engineering with Northwestern University, Evanston, IL, USA. His laboratory is part of the Neuroscience and Robotics Laboratory, and his research interests include robotics, control, computational methods for biomechanical systems, and computational neuroscience.

Dr. Murphey received the National Science Foundation CAREER Award in 2006, membership in the 2014–2015 DARPA/IDA Defense Science Study Group, and

Northwestern's Professorship of Teaching Excellence. He is a Senior Editor of the IEEE TRANSACTIONS ON ROBOTICS.

Fig. 1.
Cart-pendulum system.

Fig. 2.
Experimental setup with annotations of the important components provided.

Fig. 3.

Image processing performed on the point cloud provided by the Kinect. The top image shows the raw point cloud, and the bottom shows the data after all processing; different colors are assigned to each cluster, and white lines are displayed to show the lines produced by the sample consensus filter.

Fig. 4.

Contour plot of the unidentifiable system. The deterministic value of the parameter is shown by the green dot in the contour plot with isolines indicating the estimator cost $\beta(\theta)$. The straight lines indicate a poorly conditioned estimation problem, which results in nonidentifiable parameters.

Fig. 5.

Comparisons of the trajectory before and after Fisher information optimization. (a) Cart position $x(t)$. (b) Link 1 angle $\phi_1(t)$. (c) Link 2 angle $\phi_2(t)$. (d) Control input $u(t)$.

Fig. 6. Monte-Carlo histogram approximations of the PDF with respect to the Lebesgue measure of the optimized trajectory with the optimized variance fit in red and the initial trajectory variance in blue. (a) PDF histograms of the top link mass, m_1 . (b) PDF histograms of the damping coefficient c .

Fig. 7.

Comparisons of the measured trajectories and estimated model trajectory for the initial choice of experimental trajectory. (a) Cart position $x(t)$. (b) Link 1 angle $\phi_1(t)$. (c) Link 2 angle $\phi_2(t)$.

Fig. 8. Comparisons of the measured trajectories and estimated model for the optimized trajectory based on the Fisher information metric. (a) Cart position $x(t)$. (b) Link 1 angle $\phi_1(t)$. (c) Link 2 angle $\phi_2(t)$.

Fig. 9.

Comparisons of parameter optimization cost contours with isolines indicating the estimator cost $\beta(\theta)$. The deterministic estimate of the parameter is shown by the green dot in the contour plot. The parametric map of the optimized trajectory is significantly more convex than the initial trajectory. (a) Initial trajectory. (b) Optimized trajectory.

TABLE I

Measured System Parameters

<i>Parameter</i>	Link 1 Value	Link 2 Value
Mass, m	Estimate	0.0847 kg
Link Length, L	0.305 m	0.305 m
Link Width, w	0.0445 m	0.0381 m
Bearing offset from edge	0.0127 m	0.0127 m
Center of mass position, s	0.146 m	0.125 m
Damping coefficient, c	Estimate*	Estimate*

* The values for the damping coefficient of each joint will be assumed to be equal.

TABLE II

Optimization Results

	λ_1	λ_2	J
Initial:	123.3	0.0315	15.9
Optimized:	1.43×10^5	32.6	0.618
Fisher Information Matrix			
Initial:	$\begin{bmatrix} 3.72 \times 10^3 & 5.14 \\ -5.14 & 9.57 \times 10^{-1} \end{bmatrix}$		
Optimized:	$\begin{bmatrix} 4.31 \times 10^6 & -3.77 \times 10^3 \\ -3.77 \times 10^3 & 9.91 \times 10^2 \end{bmatrix}$		

TABLE III

Monte-Carlo Simulation Results

Monte-Carlo Covariance	
Initial:	$\begin{bmatrix} 3.16 \times 10^{-4} & 3.45 \times 10^{-3} \\ 3.45 \times 10^{-3} & 1.06 \end{bmatrix}$
Optimized:	$\begin{bmatrix} 2.37 \times 10^{-7} & 1.63 \times 10^{-6} \\ 1.63 \times 10^{-6} & 1.19 \times 10^{-3} \end{bmatrix}$
Cramer-Rao Lower Bound	
Initial:	$\begin{bmatrix} 2.71 \times 10^{-4} & 1.46 \times 10^{-3} \\ 1.46 \times 10^{-3} & 1.05 \end{bmatrix}$
Optimized:	$\begin{bmatrix} 2.32 \times 10^{-7} & 8.83 \times 10^{-7} \\ 8.83 \times 10^{-7} & 1.01 \times 10^{-3} \end{bmatrix}$

TABLE IV

Experimental Results

	m_1 (kg)	c (g/s)
Initial Estimate:	0.085	0.500
Measured Actual Value:	0.110	0.180
Initial Trajectory Estimate:	0.085	0.500
% Error from Actual Value:	22.7%	316.7%
Optimal Trajectory Estimate:	0.107	0.211
% Error from Actual Value:	2.7%	17.2%