

Behavior Synthesis via Contact-Aware Fisher Information Maximization

Hrishikesh Sathyanarayan and Ian Abraham
Yale University, New Haven, CT, USA
Email: {hrishi.sathyanarayan, ian.abraham}@yale.edu

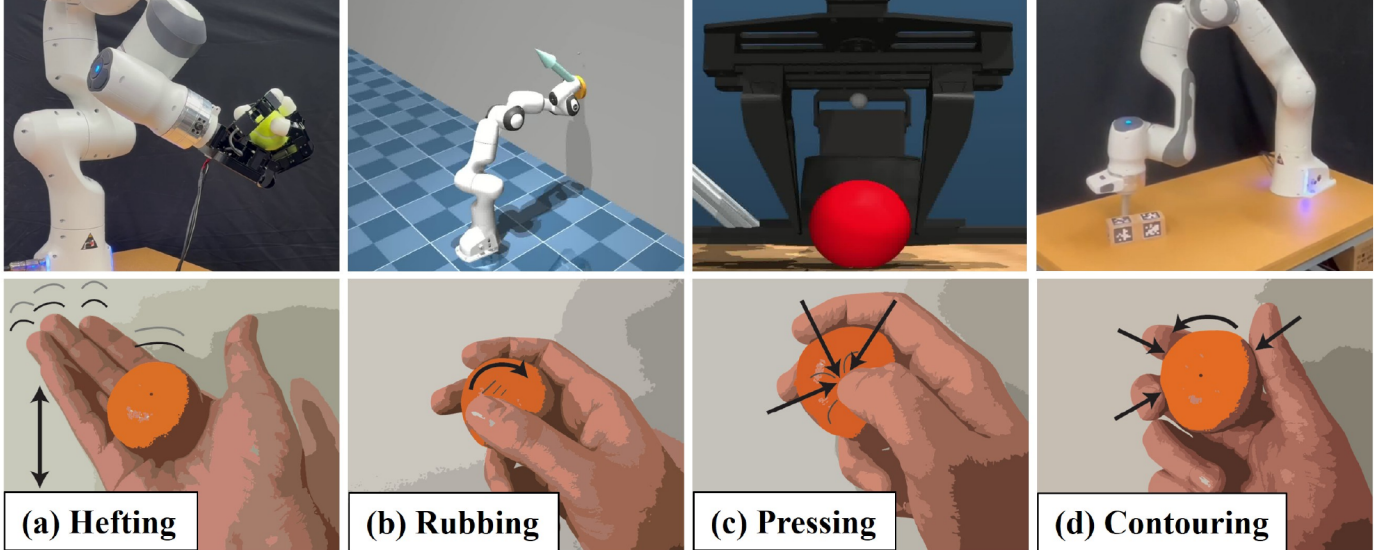


Figure 1: **Emergent Contact-Based Learning Behaviors.** Here, we show emergent tactile behaviors resulting from the proposed contact-aware Fisher information maximization that results in human-like tactile behaviors for learning (a) mass and weight, (b) friction and textures, (c) stiffness, and (d) shape [20]. Additional media and code is provided in https://github.com/ialab-yale/contact_aware_active_learning.git.

Abstract—Contact dynamics hold immense amounts of information that can improve a robot’s ability to characterize and learn about objects in their environment through interactions. However, collecting information-rich contact data is challenging due to its inherent sparsity and non-smooth nature, requiring an active approach to maximize the utility of contacts for learning. In this work, we investigate an optimal experimental design approach to synthesize robot behaviors that produce contact-rich data for learning. Our approach derives a contact-aware Fisher information measure that characterizes information-rich contact behaviors that improve parameter learning. We observe emergent robot behaviors that are able to excite contact interactions that efficiently learn object parameters across a range of parameter learning examples. Last, we demonstrate the utility of contact-awareness for learning parameters through contact-seeking behaviors on several robotic experiments.

I. INTRODUCTION

Contact dynamics are commonly used in robotics to manipulate the robot itself, e.g., through locomotion, or manipulate objects in its environment. However, the utility of contacts goes beyond just manipulation, and instead, contact can be seen as a medium to transmit information that can help a robot learn about its environment. In fact, prior work has demonstrated the information-richness of contact as a means

to improve parameter estimation problems [8, 21, 27]. The underlying challenge is enabling robot behaviors that can actively acquire contact data for learning.

Most prior work that leverages contact for learning acquire contact data as a by-product of implicit interactions [8, 21]. However, information “rich” contact interactions are scarcely found through unintentional interactions, necessitating the need for robots to *actively* explore contacts. Past work has shown the advantages of active exploration as a means to improve data collection for learning [44, 39], but more work is needed to improve the quality of contact data collected in exploration. In particular, some physical parameters of objects, e.g., shape and inertial parameters, can only be detected by continuously breaking and making contact [32, 30, 45, 28, 14, 3]. Thus, the goal of this work is to develop the theoretical foundations that optimize for robot behaviors that produce contact-rich data for learning.

To achieve this goal, we first formulate the maximum likelihood formulation for parameter learning subject to contact-based physics constraints. In this formulation, we connect the robot’s dynamics, and contact constraints, with the generation of sensor data for learning. We employ techniques from opti-

mal experimental design and derive a utility measure based on the Fisher information [9, 2] that explicitly takes into account information-rich sensor data as a function of contact. By maximizing this contact-aware Fisher information measure, we observe emergent robot behaviors that reasons about contact interactions that significantly improves parameter learning.

We focus on a class of example problems where the robot is tasked to interact with and learn object properties through interaction (see Fig. 1). These types of problems are interesting because many physical parameters can only be estimated from physical interaction [8, 21]. In this work, we show that a contact-aware optimal experimental design approach can produce *meaningful* contact interactions for robots to obtain information-rich data that facilitates improvement in parameter learning. In summary, the contributions of this paper are the following:

- 1) Formulation of a contact-aware Fisher information measure to identify information-rich contacts.
- 2) Development of an optimal control approach that plans contact sequences that improves parameter learning.
- 3) Demonstration of our approach generating emergent contact behaviors on a variety of robot parameter learning scenarios.

This paper is structured in the following manner: Section II overviews prior related work. Section III overviews preliminary material on Parameter Estimation, Fisher information, and Experimental Design. Section IV details the methodologies used to construct our approach. Section VI discusses our results and insights from using our approach. Section VII overviews limitations and future works. We conclude the paper in Section VIII.

II. RELATED WORK

A. Parameter Estimation

Parameter estimation has been widely explored where unknown physical parameters about the robot need to be estimated more precisely through measurements [21, 11, 16, 23]. In general, parameter estimation involves utilizing sensor measurements from a system in order to fit unknown parameters of interests to the sensor readings based on a model of choice that captures the dynamical behaviors of the system. In the context of learning parameters from contact, prior work remarks on the observability and identification of parameters [8], stating that not all interactions a system makes can be useful for estimating parameters. For example, if a system makes no effort to interact with the environment, and vice versa, parameters such as mass and inertia are indistinguishable due to the homogeneity of the equations of motion, and can only retrieve a ratio of the parameters.

Parameter estimation algorithms generally consist of fitting parameters comprised in a model structure with sensor readings from the robot. Depending on the type of sensor used, certain classes of physical parameters can be reasonably identifiable [4, 22]. For instance, estimating inertial parameters becomes intractable if the only available sensor for parameter

estimation is an accelerometer, and perceived contact forces have no reference value to account for inertia. Therefore, choosing the *type* of sensors available for parameter estimation is crucial before deducing what unknown parameters should be estimated.

Common methods for parameter estimation are Bayesian-inference based approaches [29, 34] and least squares estimation [8]. Bayesian approaches for parameter estimation bake parameter uncertainty into the problem and can generalize well for complex, nonlinear problems, at the cost of computational efficiency for large data. Least squares estimation works well for simple, linear regression problems, but struggles to generalize to complex, nonlinear problems. Other methods are presented in [38], where parameter estimation is solved via a factor graph-based approach. However, factor graph-based methods are computationally heavy to compute, and requires careful engineering to construct. While there are many approaches to parameter estimation, this paper formulates the parameter estimation as a maximum likelihood problem [12], which allows us to consider the dependence of data in parameter estimation. Specifically, the maximum likelihood formulation is amenable to techniques in optimal experimental design that provides meaningful requirements for measurement data.

B. Optimal Experimental Design and Fisher Information

The study of optimal experimental design [7, 9] provides an insight towards how to best sample data for estimating parameters in a maximum likelihood problem. Experimental design involves optimizing over an information metric in order to solve for expected measurements in an experiment that yield the most *information-rich* data about unknown parameters of interest. Fisher information is often used in experimental design, as it directly captures how sensitive the sensor measurements are to slight perturbations of the parameter estimates. By maximizing Fisher information [2, 7], we take an outer product of the gradients over the maximum likelihood reward function with respect to parameters of interest and solve for regions in this new reward landscape where the function is the *most sensitive* to slight perturbations of the parameters. This optimization problem is known as experimental learning. The Fisher information metric is often used as it captures information-richness in datasets and is directly related to the variance of parameter estimates of interest [33]. By optimizing directly over the Fisher information landscape, we yield experiments that lead to desirable, information-rich data [43, 44].

In the context of robotics, experimental learning provides us a means to acquire information-rich data that describes unknown parameters of interest accurately [43, 44]. The learning algorithm will output a control sequence that is applied to the robot which generates a motion that collects information-rich data through equipped sensors.

This paper takes a conventional experimental design problem and primes the Fisher information with contact-awareness, allowing the robot to optimize for *information-rich* contact interactions to improve parameter learning. In doing so, we pose an experimental design problem where contact is an implicit variable for optimization as a function of robot trajectory. Contact-implicit trajectory optimization [24, 26, 40] synthesizes trajectories wherein contact forces, robot state, and control inputs are posed as decision variables in the optimization problem. They are generally solved by optimizing for robot trajectory and contacts over an objective, with constraints being robot dynamics and a contact model of choice [32, 8, 30, 40]. Essentially, these methods can generate motions without explicitly specifying contact mode sequences.

Such formulations of optimization problems have been shown to have impressive results, shown in [26, 40, 32]. These problems are formulated by a choice of contact model and optimization methods. Particular to the algorithm presented in [32], the resulting motion is generated through a time stepping method with linear complementarity constraints posed by [36, 1] in order to resolve rigid-body contact modes implicitly. In this work, we address a parameter learning problem formulated as an *contact-implicit optimization*, using a contact model of choice, where the parameters are embedded within the contact model. We aim to refine parameter estimates by optimizing contact modes such that they align closely with given contact data.

D. Belief Space Planning and Control

Belief space planning is commonly used in robotics to plan under parameter uncertainty by considering the robot's belief about the parameter, and gathering information to converge on parameter uncertainty. Consequently, the robot infers unknown properties of objects around it and plans for interactions accordingly. Early works in [31, 19, 15] present belief space control using a Partially Observable Markov Decision Process (POMDP), where the robot explores for relevant and rich information about parameters which is used to solve a maximum likelihood problem.

Recent approaches [42, 25] leverage the belief space formulation for improving contact-based robot tasks. Our approach is similar in that, we leverage the parameter uncertainty based on the sensitivity of the measurement model to improve robot task performance. The main difference is 1) the explicit consideration of contact mechanics in computing the Fisher information; and 2) that we can solve the contact-aware Fisher information maximization problem via sample-based control for real-time performance. The advantage is that we can generalize across a larger range of contact-based robotics tasks, and that can readily solve the information maximization problem under real-time control setting without needing to resort to computationally expensive reinforcement learning methods.

A. Parameter Learning via Maximum a-posteriori Estimate

In this work, we cast parameter learning as a maximum a-posteriori estimation over parameters $\theta \in \Theta$. Let $p(\theta) : \Theta \rightarrow \mathbb{R}$ be the prior distribution over θ and $\mathcal{D} = \{(x_i, y_i)\}_{i=1}^N$ be a collection of N input-output sensor data where $x_i \in \mathcal{X}$ is an input, and $y_i \in \mathcal{Y}$ is noisy sensor data. Given a likelihood model¹ $p(\mathcal{D}|\theta)$, the posterior belief $p(\theta|\mathcal{D})$ is given by the Bayesian update

$$p(\theta|\mathcal{D}) = \frac{p(\mathcal{D}|\theta)p(\theta)}{\int_{\Theta} p(\mathcal{D}|v)p(v)dv}. \quad (1)$$

It is often difficult or intractable to directly solve for the posterior distribution given data. Instead, it is easier to just optimize for the mode $\hat{\theta}$ of the posterior distribution.

Definition 1. Maximum a-posteriori Estimation. Given a prior $p(\theta) : \Theta \rightarrow \mathbb{R}$ over parameters $\theta \in \Theta$, a data set of N independent and identically distributed (i.i.d.) measurements $\mathcal{D} = \{(x_i, y_i)\}_{i=1}^N$, and a likelihood model $p(\mathcal{D}|\theta)$, the point estimate of the posterior distribution is given as the maximum a-posteriori (MAP) estimate $\hat{\theta}$ which is the solution to the optimization

$$\hat{\theta} = \arg \max_{\theta} p(\mathcal{D}|\theta)p(\theta). \quad (2)$$

In practice it is common to optimize for the log a-posteriori $\hat{\theta} = \arg \max_{\theta} \log(p(\mathcal{D}|\theta)p(\theta))$ which often has better numerical conditioning.

Remark 1. Not including the prior $p(\theta)$ in the maximization in (2) yields the simpler, but effective maximum likelihood problem. While the formulation remains useful for accounting for collected data, it is unable to account for prior information.

In general, solving the maximum a-posteriori estimation assumes data has been provided. We are interested in producing a method that provides the best data \mathcal{D} that improves the conditioning of (2).

B. Fisher Information

One way to establish better conditioning on maximum a-posteriori problems is to analyze the statistical manifold of the problem. In particular, we are interested in analyzing the sensitivity of the posterior distribution with respect to the parameter. The insight is that certain measurements yield more *information* about parameters, i.e., reduction in uncertainty, when the gradient of the likelihood model is largest. A common measure of information with respect to parameter sensitivity is the *Fisher information matrix*.

Definition 2. Fisher Information Matrix. Let $\mathcal{L}(\mathcal{D}|\theta)$ be the log likelihood of an unbiased estimator that is continuous and differentiable with respect to elements in \mathcal{D} and $\theta \in \Theta \subseteq \mathbb{R}^d$.

¹In our work, we call this the sensor model which provides the likelihood that the collected data \mathcal{D} is explained by parameter θ .

Then, the Fisher Information Matrix $\mathcal{F}(\mathcal{D}|\theta) \in \mathbb{R}^{d \times d}$ of θ is defined by

$$\mathcal{F}(\mathcal{D}|\theta) = \mathbb{E} \left[\frac{\partial}{\partial \theta} \mathcal{L}(\mathcal{D}|\theta) \frac{\partial}{\partial \theta} \mathcal{L}(\mathcal{D}|\theta)^\top \right] \quad (3)$$

and

$$\mathcal{F}(\mathcal{D}|\theta) = -\mathbb{E} [\nabla_\theta^2 \mathcal{L}(\mathcal{D}|\theta)] \quad (4)$$

if $\mathcal{L} \in \mathcal{C}^2$. Here, the expectation is with respect to the random variable \mathcal{D} and ∇_θ^2 is the Hessian [6].

Note that the Fisher Interestingly, the Fisher information matrix also has a unique interpretation as a lower-bound on the uncertainty of the posterior in maximum likelihood problem.

Definition 3. Cramér-Rao Lower Bound [33]. Let \mathcal{D} be a random variable drawn from a distribution $p(\mathcal{D}|\theta)$, and $\hat{\theta}$ be the unbiased estimator for θ , then

$$\mathcal{F}(\mathcal{D}|\theta) \geq \text{cov}(\hat{\theta})^{-1}. \quad (5)$$

For the interested reader, the proof of this lower bound can be found in [33]. The bound shows that maximizing the Fisher information matrix directly impacts the variance of the parameter estimate which we use to establish an optimization problem over *planned datasets*.

C. Optimal Experimental Design

We can maximize the Fisher information in order to choose measurements $\mathcal{D} = \{(x_i, y_i)\}_{i=1}^N$ that improves the precision of parameter estimation. The idea is to establish an optimization on the input design of an experiment, i.e., the independent variable x_i , such that the dependent variable y_i that has not been collected has more utility. To establish this optimization, we first assume maximum likelihood and normally distributed measurements[46] where $\mathcal{L}(\mathcal{D}|\theta) = \mathcal{L}(X|\theta)$ and $X = \{x_i\}_{i=1}^N$. Under a finite sampling of independent variables x_i we obtain the empirical Fisher information

$$\begin{aligned} \mathcal{F}(\mathcal{D}|\theta) &\approx -\frac{1}{N} \sum_{i=1}^N \frac{\partial^2 \mathcal{L}(X_i|\theta)}{\partial \theta^2} \\ &= \frac{1}{N} \sum_{i=1}^N \frac{\partial \mathcal{L}(X_i|\theta)}{\partial \theta} \frac{\partial \mathcal{L}(X_i|\theta)^\top}{\partial \theta}. \end{aligned} \quad (6)$$

Since X is now an independent control variable, the following *optimal design* is constructed

$$\max_X \psi(\mathcal{F}(X|\theta)) \quad (7)$$

where the Fisher information matrix is overloaded with argument X [33, 9], and $\psi(\cdot)$ is a matrix reduction operator over the Fisher information that returns a scalar measure of the matrix, e.g., the determinant (D-optimality), or the trace (T-optimality). Note that the Fisher information is symmetric and positive semi-definite [17, 5]. The inducing operator ψ is then a metric on the information landscape that alter the experiment design, e.g., D-optimality will maximize the total volume of the Fisher information matrix, whereas T-optimality maximizes the largest enclosing sphere. Our goal

is to formulate a *contact-aware* formulation of experimental design to facilitate with contact-based robot tasks that seeks information-rich contact for learning.

IV. CONTACT-AWARE EXPERIMENTAL DESIGN

To enable robots to actively seek out contact, we need to derive a notion of “information” related to parameter estimation subject to contact. Thus, we first need to establish the explicit dependence of contact in parameter estimation.

A. Contact-Aware Max. A-Posteriori Estimation (CA-MAP)

Let us define a robot with the state variable $x_t \in \mathcal{X}$ with control input $u_t \in \mathcal{U}$ at some discrete time $t \in \mathbb{Z}^+$. Additionally, let us define the dynamics constraints of the robot as $x_{t+1} = f_\theta(x_t, u_t, \lambda_t)$ and the sensor measurement model $y_t = g_\theta(x_t)$ parameterized by $\theta \in \Theta$ where $y_t \in \mathcal{Y}$ is a sensor measurement, and $\lambda_t \in \mathcal{C}_\theta(x_t)$ is the contact constraint. Here, $\mathcal{C}_\theta(x)$ denotes the contact conditions as a function of state, e.g., friction, non-penetration, and complementarity. We define an “experiment” as a set of T measurements $\mathcal{D} = \{(y_t, u_t) \forall t \in [0, T-1]\}$ where the goal is to maximize the a-posteriori estimate θ given a prior $p(\theta)$ and experiment \mathcal{D} subject to the dynamic, contact, and sensor model constraints.

Definition 4. Contact-Aware MAP (CA-MAP) Let $\tau = \{(x_t, \lambda_t) \forall t \in [0, T-1]\}$ be a trajectory of states x_t and contact forces λ_t for some time horizon $T \in \mathbb{Z}^+$. Given the likelihood model $p(\mathcal{D}|\tau, \theta)$ (i.e., a sensor model) that is normally distributed with mean $g_\theta(x)$ and variance Σ , the contact-aware maximum a-posteriori estimation (CA-MAP) problem is defined as

$$\begin{aligned} \hat{\theta} &= \arg \max_{\theta, \tau} \log p(\mathcal{D}|\tau, \theta) p(\theta) \\ \text{s.t.} \quad &\begin{cases} x_0 \text{ given} & (\text{initial cond.}) \\ x_{t+1} = f_\theta(x_t, u_t, \lambda_t) & (\text{dynamics model}) \\ \lambda_t \in \mathcal{C}_\theta(x_t) & (\text{contact constraints}) \\ \theta \in \Theta & (\text{feasible parameter}) \end{cases} \end{aligned} \quad (8)$$

where the subscript θ denotes the functional dependence of the model on the parameter.

The intuition behind this problem formulation is that we are optimizing for the parameter estimation *and* the trajectory τ that best explains the data \mathcal{D} . The role of the contact model for λ is to explicitly consider contact constraints as part of the parameter estimation problem. In many scenarios, contact forces are not explicitly measured, but do appear in the physical constraints of the system and are critical for parameter estimation. Thus, incorporating the contact model enables us to explicitly optimize contact interactions for experimental design.

B. CA-MAP Fisher Information Matrix

Given the contact-aware a-posteriori estimation problem, our goal is to derive an equivalent notion of the Fisher information that explicitly considers contact. To achieve this,

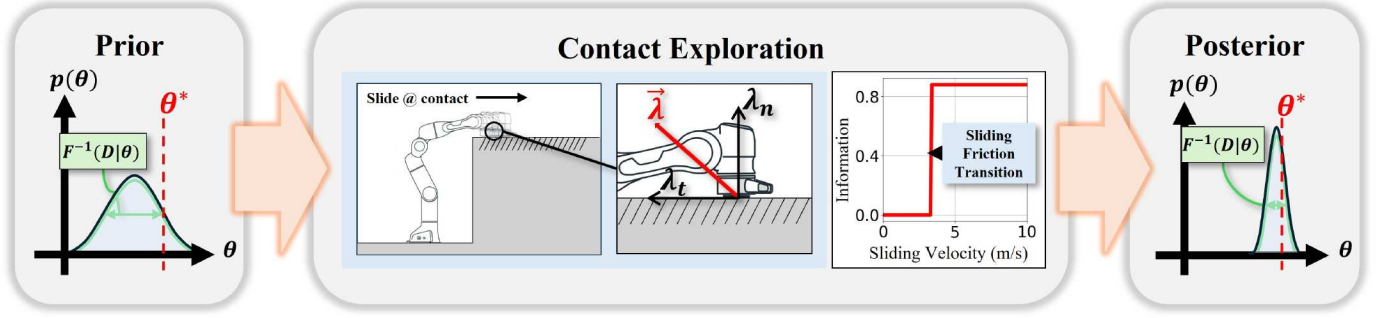


Figure 2: **Collapsing Parameter Uncertainty Through Contact.** In this example, the robot seeks to estimate the friction coefficient, μ , of a nearby table. Given an initial prior of the parameter estimate, the robot seeks to update its belief of the friction coefficient by interacting with the environment through contact. In doing so, the robot collects contact sensor measurements that hold information, \mathcal{F} (green curve), about the parameter, which is then used to update the parameter estimate and variance. We can see that in order to reduce parameter variance and converge to a more precise parameter estimate, the robot exhibits a *sliding* behavior that excites terms in the Fisher information, which is a *lower bound* on parameter uncertainty [33].

we first define the Lagrangian of the contact-aware parameter estimation problem as

$$\mathcal{L}(\mathcal{D}|\tau, \theta, \alpha) = \ell(\mathcal{D}|\tau, \theta) + \alpha d(\tau, \theta) \quad (9)$$

where α is the Lagrange multiplier for each constraint $d(\tau, \theta)$ and $\ell(\mathcal{D}|\tau, \theta) = \log p(\mathcal{D}|\tau, \theta)p(\theta)$.

Theorem 1. Let $\mathcal{L}(\mathcal{D}|\tau, \theta, \alpha)$ be the Lagrangian of the contact-aware estimation problem Def. 4 where $\mathcal{L} \in \mathcal{C}^2$. The contact-aware Fisher information matrix $\mathcal{F}(\mathcal{D}|\tau, \theta, \alpha) \in \mathbb{R}^{d \times d}$ for $\theta \in \Theta \subseteq \mathbb{R}^d$ is given by

$$\mathcal{F}(\mathcal{D}|\tau, \theta, \alpha) = -\nabla_{\theta}^2 \mathcal{L}(\mathcal{D}|\tau, \theta, \alpha) \quad (10)$$

defined by the steepest ascent problem

$$\delta\theta^* = \arg \max_{\delta\theta} \nabla_{\theta} \mathcal{L}(\mathcal{D}|\tau, \theta)^{\top} \delta\theta + \frac{1}{2} \delta\theta^{\top} \nabla_{\theta}^2 \mathcal{L}(\mathcal{D}|\tau, \theta) \delta\theta \quad (11)$$

where $\mathcal{F}(\mathcal{D}|\tau, \theta) = \nabla_{\theta}^2 \mathcal{L}(\mathcal{D}|\tau, \theta)$ and the dependence of the Lagrange multipliers are dropped and assumed to be constants.

Proof: See Appendix A for the complete proof. ■

Theorem 1 connects a variation of the Fisher information matrix to a Newton’s method conditioning for the steepest descent problem. Note that in this case, we define the empirical Fisher information (where the summation over data is implicit in the augmented log-likelihood model). Furthermore, consider that the Fisher information has as an argument the trajectory $\tau = \{(x_t, \lambda_t)\}_{t=0}^T$, and the measured data $\mathcal{D} = \{(y_i, u_i)\}_{i=0}^T$. We can use this fact to *solve* for robot behaviors that maximize the Fisher information, ultimately improving the conditioning on the steepest descent for the CA-MAP problem in the *data* space where contact can be explicitly sought to excite information about physical parameters.

V. CONTACT-AWARE OPTIMAL EXPERIMENTAL DESIGN

We now derive the contact-aware optimal experimental design approach to synthesize robot behaviors for learning. Our approach integrates the contact-aware Fisher information into a closed-loop optimal design procedure to facilitate robot parameter learning from information-rich contact data.

A. Contact-Seeking Behavior Synthesis

We first establish the optimization formulation that maximizes the contact-aware Fisher information matrix.

Definition 5. (Contact-Aware Optimal Experimental Design (CA-OED)). Let $\psi : \mathbb{R}^{d \times d} \rightarrow \mathbb{R}$ define a continuous and differentiable metric and $\mathcal{J}(\mathcal{D}, \tau)$ denote a cost term on the trajectory. Additionally, assume the contact-aware Fisher information matrix $\mathcal{F}(\mathcal{D}|\tau, \theta)$ is symmetric and positive semi-definite [35, 37]. Then, for a given estimate θ , the contact-aware optimal experimental design problem is defined as

$$\hat{\mathcal{D}}, \hat{\tau} = \arg \max_{\mathcal{D}, \tau} \psi(\mathcal{F}(\mathcal{D}|\tau, \theta)) - \mathcal{J}(\mathcal{D}, \tau) \quad (12)$$

$$\text{s.t.} \quad \begin{cases} x_0 \text{ (given)} \\ x_{t+1} = f_{\theta}(x_t, u_t, \lambda_t), \\ \lambda_t \in \mathcal{C}_{\theta}(x_t) \\ y_t = g_{\theta}(x_t) \end{cases}$$

where $\hat{\mathcal{D}} = \{(y_t, u_t)\}_{t=0}^T$ is the predicted dataset and $\hat{\tau} = \{(x_t, \lambda_t)\}_{t=0}^T$ is a reference trajectory.

We can solve this optimization problem in (12) using any number of optimization techniques. We opted to solve this problem using predictive sampling [13] where contact constraints are implicitly solved through the dynamics model [32, 8]. In this formulation of optimal experimental design, it is important to recognize that the measurements y_t that are optimized are only predictions. The purpose of the prediction

is to specify the Fisher information matrix which explicitly considers the parameter sensitivity to the measurement. As such, we can use the sensor predictions as an interpretation of what “information-rich” contact behaviors produce.

In this work, we use the trace for ψ as the metric on the Fisher information. By maximizing the trace Fisher information matrix, we improve the condition number of the matrix used for Newton’s step in the CA-MAP problem thus allowing empirically faster convergence and numerical stability. An example of this phenomena is illustrated in Fig. 2.

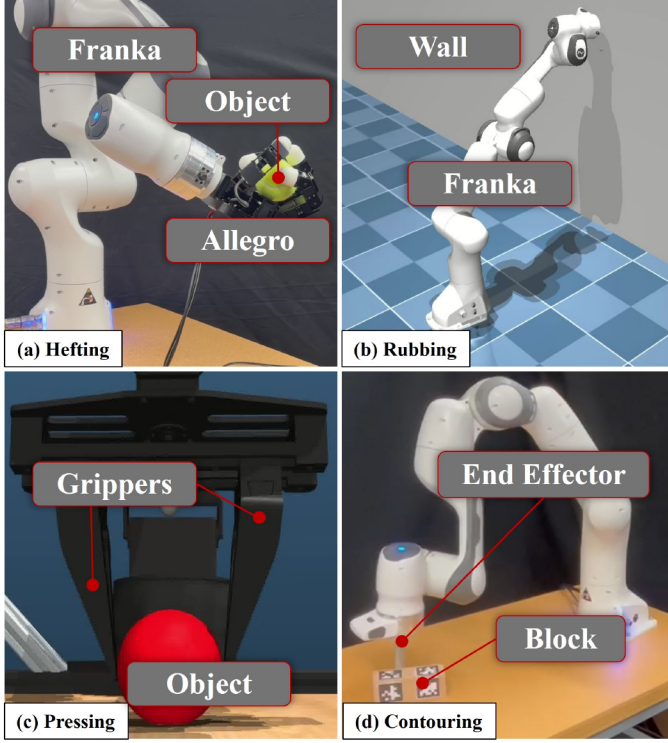


Figure 3: **Experimental Setup.** The experimental setup for (a) hefting experiments for mass estimation, (b) rubbing experiments for friction coefficient estimation, (c) pressing experiments for material stiffness and dampening estimation, and (d) contouring experiments for shape estimation.

B. Closing the Loop

Given the solution to Eq. (12), the following step is to execute the behaviors on the robot, record the measured data and repeat the process until a desired parameter confidence defined by $p(\theta|\mathcal{D})$. Since our problem formulation only accounts for the posterior parameter mode $\hat{\theta}$, we use the approximate update equation on the posterior variance

$$\Sigma_{\theta}^{+} = \left(\mathcal{F}(\mathcal{D}|\tau, \hat{\theta}) + \Sigma_{\theta}^{-1} \right)^{-1} \quad (13)$$

where the Fisher information is calculated from the collected experiment data, Σ_{θ} is the covariance of the prior, Σ_{θ}^{+} is the posterior covariance, and $\hat{\theta}$ is solved initially using Eq. 8. Note that this equation over-approximates the uncertainty, which we found in practice assists with parameter convergence.

Remark 2. Because we are estimating continuous parameters θ , and not contact forces or potentially discontinuous state transitions, we can assume belief propagation is also continuous. Thus, we do not require techniques such as saltation matrices [18] to estimate parameter belief.

We present an algorithm that comprises of 1) the contact-aware experimental design optimization that searches for information-rich contact, and 2) the contact-aware maximum a-posteriori estimation problem that estimates parameters from collected data in Algorithm 1. We summarize the contact-aware experimental design procedure in Fig. 4.

Algorithm 1 Contact-Aware Behavior Synthesis

Require: initial state x_0 , prior parameter estimate $p(\theta)$, parameterized dynamics, sensor, and contact models, total number of experiments k_{\max} , and experiment duration T .

- 1: $k = 0$
- 2: **while** $k < k_{\max}$ **do**
- 3: $\hat{\mathcal{D}}, \hat{\tau} \leftarrow$ CA-OED Eq. (12) ▷ calc. contact behavior
- 4: $\mathcal{D} \leftarrow$ Execute experiment $\hat{\mathcal{D}}, \hat{\tau}$ on robot
- 5: $\hat{\theta}, \Sigma_{\theta}^{+} \leftarrow$ CA-MAP Belief update Eq. (8), Eq. (13)
- 6: $k \leftarrow k + 1$
- 7: **end while**

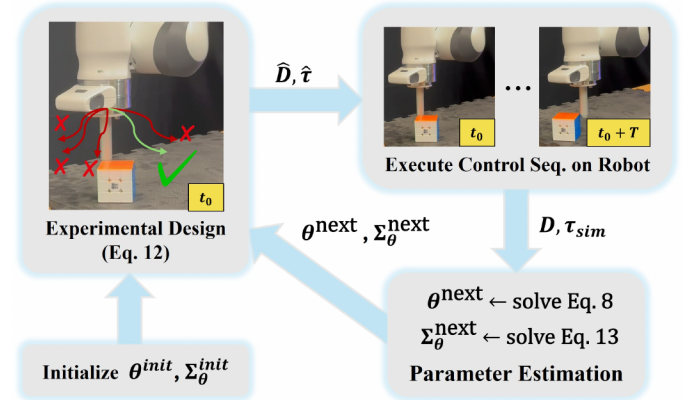


Figure 4: **Experimental Design Pipeline.** Here we outline our approach using Alg. 1. Given a prior on the parameter, we plan trajectories that excite information-rich contact modes. We then apply the planned trajectories to the robot and gather sensor measurements. These measurements are used to update the parameters and the variance of the prior to obtain the posterior. We repeat this process until the robot estimates parameters at a parameter certainty tolerance of choice.

VI. RESULTS AND DISCUSSION

This section details simulated and experimental results of our contact-aware experimental design approach in several scenarios. Our approach is evaluated using contact-implicit trajectory optimization methods where we optimize for contact modes that maximize information about objects with unknown parameters of interest, thus improving parameter

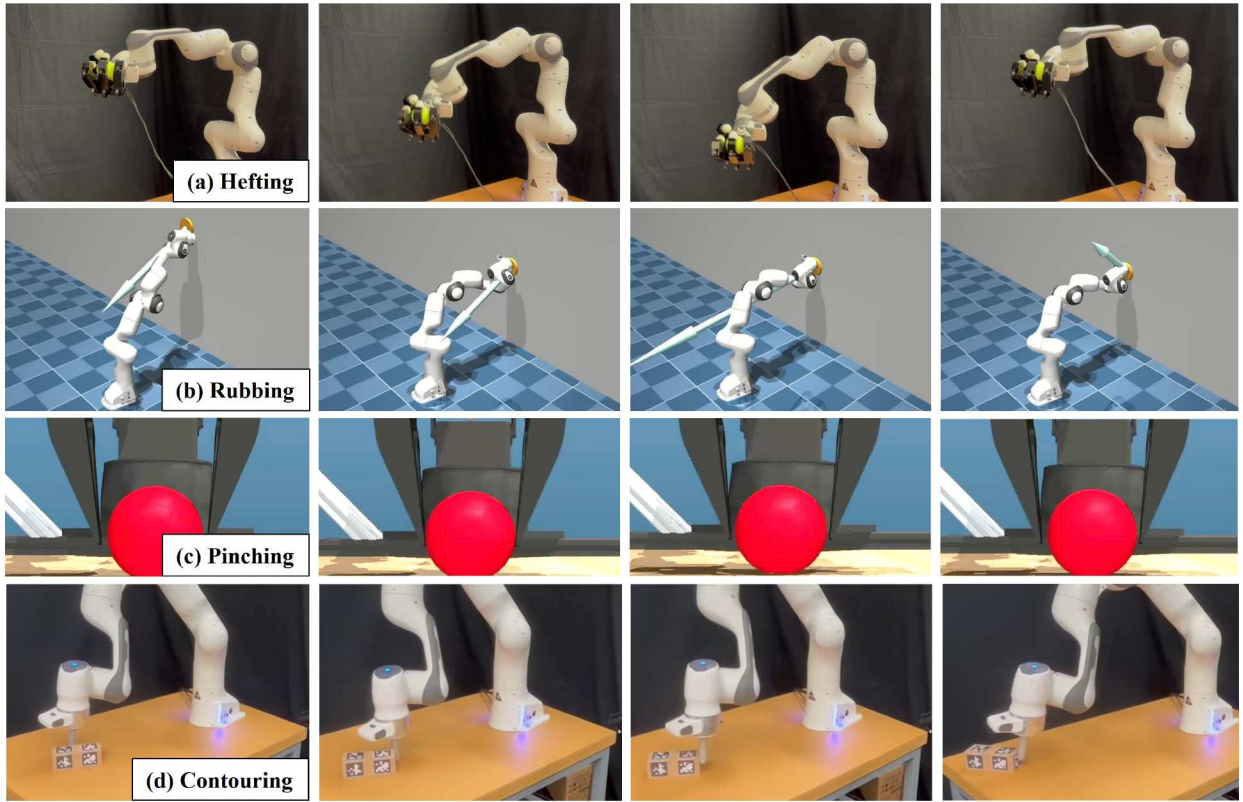


Figure 5: Emergent Behaviors from Experimental Design. Time series of experiments conducted over all scenarios. By maximizing contact-aware Fisher information, the robot seeks information-rich contacts for efficient identification of (a) mass, (b) friction coefficient, (c) material properties, and (d) shape. Simulation results are visualized on MuJoCo [41]. The robots were given a planned trajectory to execute, where we assume perfect robot state tracking upon executing the trajectory. We see the emergent behaviors in (a) a hefting motion that perturbs contact force on the object, (b) the end-effector rubbing against the wall with a nonzero normal force and tangent contact velocity, (c) the gripper pinching the object to excite contact modes, and (d) an object contouring behavior with both end effectors to excite contact modes across different regions of the object.

learning performance. We compare our approach to a similar approach [25], which executes trajectories based on a Fisher information maximization approach where contact is not considered explicitly. Particularly, we are interested in observing emergent behaviors from our approach that answer the following questions:

- Q1) What classes of parameter learning problems require more *compounding* and *constant* contact interactions as the resulting optimal emergent behavior?
- Q2) What classes of parameter learning problems require *making* and *breaking* contact as its resulting optimal emergent behavior?
- Q3) What convergence do we see for parameter learning through our optimized contact behaviors?
- Q4) How does the information loss landscape change for different parameter estimation problems?

We answer the questions aforementioned in this section sequentially with corresponding answers ‘A(-)’.

Robot Scenarios

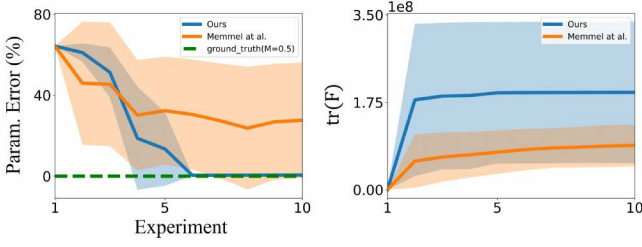
The results in this paper answers these questions by implementing our Fisher information maximization approach on the

following robot scenarios shown in Fig. 3:

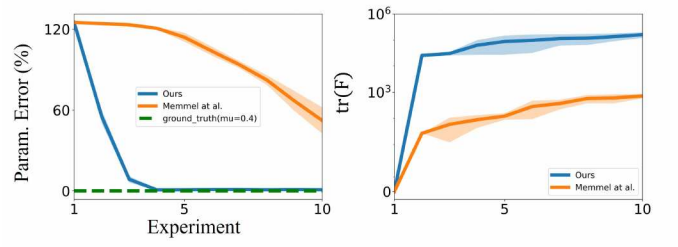
- 1) **Weight Estimation**, where an Allegro-Franka system estimates the mass of a grasped object through hefting behaviors.
- 2) **Friction Estimation**, where a Franka robot interacts with a wall to estimate friction coefficient.
- 3) **Stiffness Estimation**, where a robot gripper learns material stiffness and dampening properties of a uniform object through its grasping interactions.
- 4) **Shape Estimation**, where a Franka robot estimates the shape of a box (length and width) through non-prehensile contact interactions.

Experiment details are shown in Appendix B. Note that we implement these robot scenarios in the presence of *high contact force sensor noise*, which makes small perturbations of contact ambiguous information to identify parameters. We demonstrate the robustness of our approach to high sensor noise by enabling robots to reason about exploring for information-rich contacts *outside* some contact uncertainty envelope.

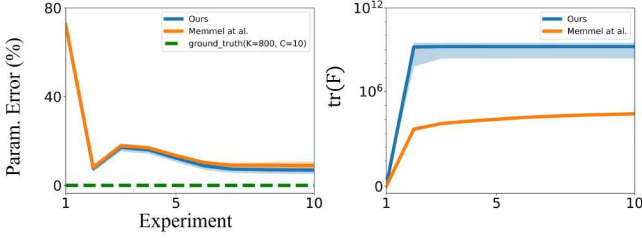
Our experimental pipeline is shown in Fig. 4. We compare



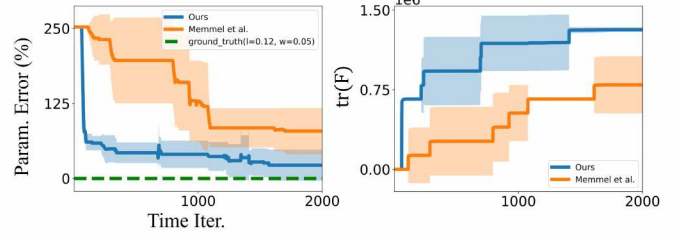
(a) **Mass Estimation.** The robot seeks to estimate the mass of a grasped object.



(b) **Friction Estimation.** The robot seeks contacts with the wall to estimate the friction coefficient.



(c) **Material Property Estimation.** The robot gripper interacts with an object to learn its material properties (stiffness and dampening coefficients).



(d) **Shape Estimation.** The robot that interacts with a box to learn its shape (length and width). Estimation is solved in-the-loop over discrete time.

Figure 6: Contact Excitations Improves Parameter Learning. We compare our experimental design approach with baseline similar to works in [25] for all robot scenarios. The baseline implements a belief-space planning controller that reasons about contact naïvely. We evaluate each robot scenario over a finite number of experiments. Here, we model the parameter error as $\%err = 100(\hat{\theta} - \theta^*)/\theta^*$, where $\hat{\theta}$ is the current estimate and θ^* is the true parameter. By maximizing the contact-aware Fisher information, we see that the robot seeks contacts that are information-rich about the unknown parameters of interest, and thus converging closer to the true parameter values θ^* , with resulting higher Fisher information gains.

our approach to a recent belief space planning baseline in [25], which explores for information-rich data about parameters without the explicit reasoning over contact.

A1. Compounding Contact Behavior

As shown in results in Fig. 5, the hefting experiment (a), where the robot pregrasps an object and learns the mass of that object through perturbing its constant normal contact force, as well as the rubbing experiment (b), where the robot interacts with a nearby wall to estimate the friction coefficient, both maintain nonzero contact in order to estimate parameters. When estimating friction coefficient, we can see in the contact sensor model in Appendix B that this term is identifiable when exciting both a high normal force as well as a non-zero tangent velocity. Therefore, maximizing over contact-aware Fisher information yields a high excitation of tangent velocities at the contact frame along with a high excitation of normal contact force. Therefore, the resulting emergent behavior results in compounding frictional contact with the wall.

In the case of the hefting experiments (a), it is possible to estimate mass parameters *even if* the object *momentarily* leaves the hand of the robot as long as the object returns to the hand. However, this adds new sparsities in the contact forces, and thus the sensor readings during the moments where the object leaves the hand would not contribute to

significant updates of the parameter prior. Therefore, in order for the robot to *effectively* learn the mass parameter of the object, it is desirable for the object to not leave the hand, and rather experience *perturbations* in normal contact forces. Such perturbations would excite the contact sensor model output written in Appendix B such that the gradient of the sensor model can extract mass information from the signed distance function ϕ_n between the object and the hand, as well as the normal velocity at the contact frame v_n . As a result, the resulting behavior for this experiment using our experimental design approach yields a hefting behavior where the robot moves its hand in an up-and-down motion in order to perturb normal force readings between the object and the robot *optimally*.

A2. Making and Breaking Contact Behavior

As shown in Fig. 5, the pressing experiment (c), where a robot gripper interacts with an object by grasping it, as well as the contouring experiment (d), where the robot interacts with an object its end effector, both make and break contact in order to significantly perturb the contact forces. Note that for experiment (d), the robot identifies the parameters online, where we execute Alg. 1 in the loop during robot operation. In these experiments, the robot is tasked with reasoning about material properties stiffness and dampening (c), as well as shape (d). To do so, our experimental design approach gives us

a planned trajectory that excites contact modes with the object by making and breaking contact in order to significantly perturb the contact forces. Intuitively, learning material properties naturally requires excitations in the normal forces between the robot and the object, and in doing so we can excite gradients in the contact sensor model as written in Appendix B.

A3. Parameter Convergence Over Experiments

We aim to demonstrate the convergence of parameter learning by comparing our experimental design approach for all robot scenarios shown in Fig. 5 to a baseline similarly implemented in [25]. The baseline reasons about information-rich trajectories to improve learning, only that it does not reason over contact explicitly by simply approximating the gradients of the sensor model using a finite differencing method. The goal of this section is to show that by reasoning over contact explicitly, we can converge better on parameter uncertainty.

The results are shown in Fig. 6. We can see that upon evaluating over experimental runs, we can observe that our experimental design approach is able to converge to a more precise parameter estimate with higher Fisher information gains. The emergent behaviors that result from contact-aware Fisher information maximization actively search for contact modes that are information-rich, resulting in naturally emerging behaviors that allow for precise identification of parameters. The reason that our approach improves the works in [25] is due to our direct consideration of contact interactions in the optimization as well as our approach in structuring the Fisher information with contact. The baseline approaches contact more naively, where Fisher information is constructed by approximating the gradients of the trajectory measurement model through a finite differencing method for a small perturbation in parameter space. By structuring the Fisher information based on a structured contact measurement model, the robot is able make deliberate decisions through its contact seeking behaviors.

In the case of the pinching experiments in Fig. 5, we can observe that our contact-aware Fisher information maximization approach only marginally improves parameter learning compared to the baseline. This due to the constrained nature of the problem, where the robot motion is constrained by a bounded 1 dimensional motion, assuming that the gripper is at a pre-grasped state. Because of high contact sensor noise, the robot is forced to perturb the contact force between its grippers and the object such that the resulting contact readings are outside of the contact uncertainty envelope in order to properly identify parameters. Because of the highly constrained nature of contact in this robot scenario, directly reasoning over information-rich contact will result in similar emergent behaviors as the baseline approach. Similar behaviors for the hefting experiments can be seen if the contact sensor is not as noisy. However, we can observe that our approach can outperform the baseline in more complex and unconstrained scenarios.

In Fig.7, we demonstrate the robustness of our approach as a function of parameter initialization. We observe that our approach is highly robust to initial priors within $\sim 30\%$

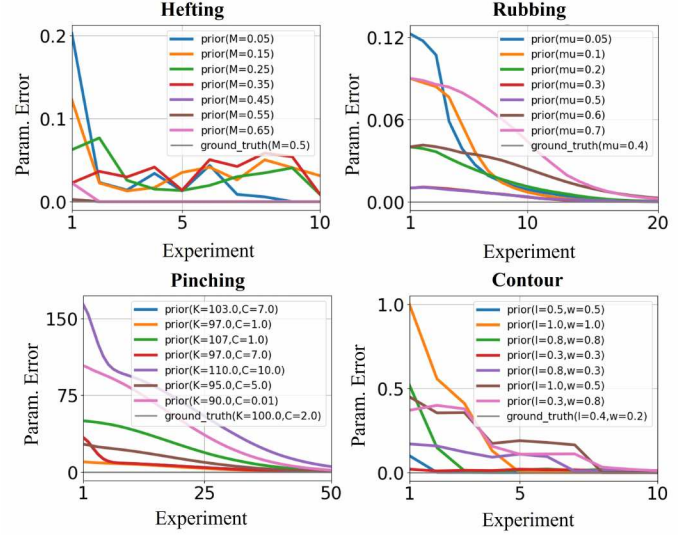


Figure 7: **Robustness to Initialization.** We study robustness of our approach is to 7 initial parameter priors for all scenarios. Note here that we model the parameter error as a distance $\delta = |\hat{\theta} - \theta^*|$. We observe that our contact-aware Fisher information maximization approach is resilient across a diverse initialization of priors over all experiments.

for hefting, and even more robust for other scenarios. We also expand on the remarks in [8] that estimating multiple parameters that are *co-dependent* to one another makes the estimation unidentifiable, unless we make assumptions of the contact interactions (e.g. bounding contact impulses), provided within the problem definition in Eq. 8 in the paper.

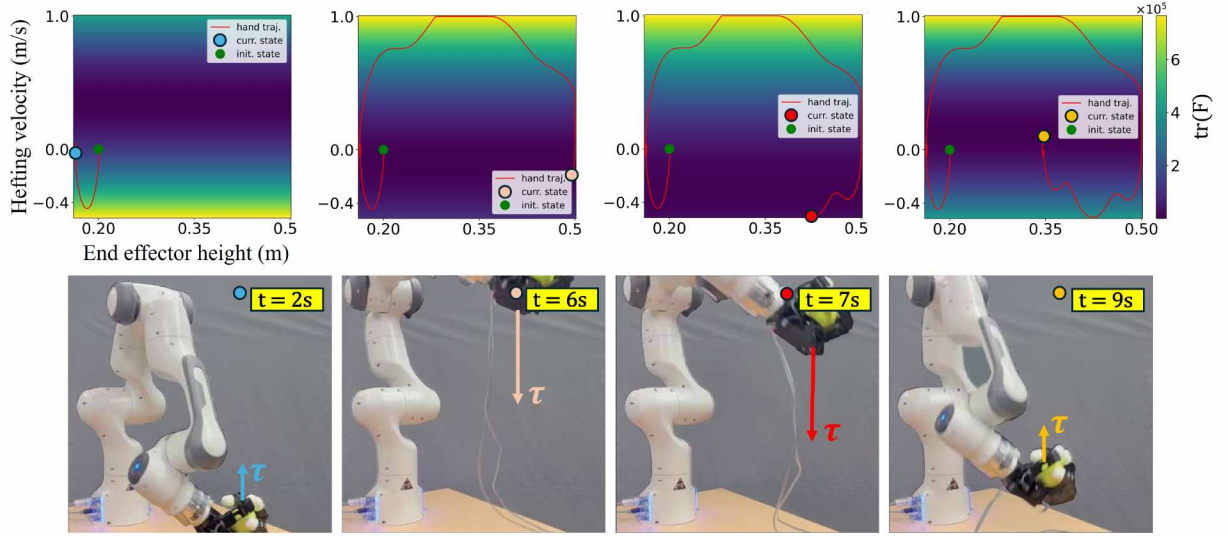
A4. Information Loss Landscape

We aim to provide a visual intuition of the information landscape that we optimize over in our contact-aware Fisher information maximization problem in Eq. 12. The information landscape is dictated by the physics models chosen for each robot scenario as well as the unknown parameters to estimate. Given that the robot has engaged contact with the environment, we show in Fig. 8 the changes in the information landscape with respect to different kinematic parameters that impact the landscape the most.

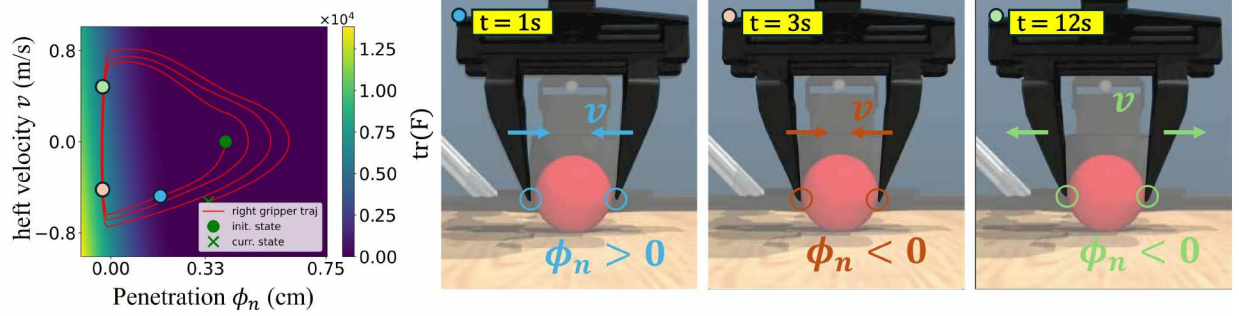
Based on the results in Fig. 8, we choose to focus on observations from the information landscape for the weight, friction, and material stiffness estimation scenarios.

Weight Estimation

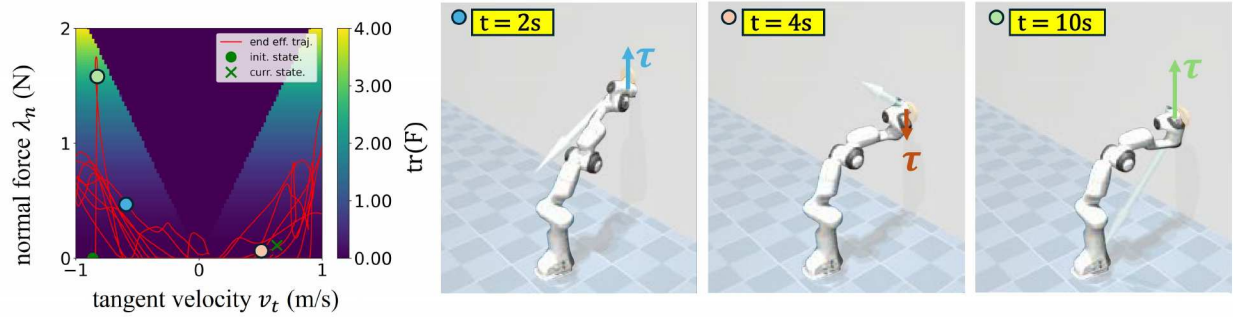
The emergent hefting behaviors shown in Fig. 5 optimizes for perturbations in normal contact between the hand and the object, with assumed soft body interactions. As a result, a natural consequence from our Fisher information maximization approach is exciting the ‘perturbation’ depth ϕ_n between the hand and the object upon contact and the normal velocity v_n at the contact frame. Fig. 8a shows the smooth information landscapes resulting from the current configuration of the hand-object system. The Fisher information landscape is linear with respect to the kinematic inputs, and results in the system



(a) **Hefting Behavior Information Landscape.** The robot excites penetration deformation ϕ_n by varying its hand positions and non-zero contact normal velocity v_n in order to excite information-rich contact about mass, resulting in an up-and-down hefting behavior shown in Fig. 5



(b) **Pinching Behavior Information Landscape.** Here, Fisher information is maximized by exciting contacts by activating the penetration distance into the object ϕ_n as well as the deformation rate $\frac{\partial \phi_n}{\partial t}$ through pinching. This yields natural behaviors of squeezing and contracting the object in order to identify material properties optimally.



(c) **Rubbing Behavior Information Landscape.** Here, information is maximized by exciting normal contact velocity v_n and tangent contact velocity v_t in order to collect information-rich contact about the friction coefficient, resulting in a rubbing behavior shown in Fig. 5

Figure 8: Optimizing Over Contact-Seeking Information Landscapes. We show the contact-aware Fisher information landscape that our experimental design method optimizes over for all scenarios.

being ‘pulled’ to the velocity boundaries, resulting in the robot exhibiting a hefting behavior.

Friction Estimation

From our approach, we observe in Fig. 5 that the robot optimally learns the friction coefficient between its end-effector and the wall by maintaining non-zero tangent contact forces

through rubbing. Intuitively, the friction coefficient parameter appears only in the tangent forces in the sensor model used for the experiment (See Appendix B), so the robot requires consistent and compounding excitations of tangent forces by pressing against the wall with a normal contact velocity v_n and exerting a non-zero tangent velocity v_t at the contact frame.

As shown in Fig. 8c, we visualize the information landscape that the robot optimizes over to obtain information-rich data to estimate the friction coefficient. The robot maximizes information about the friction coefficient during contact by sliding velocity v_t and pressing velocity v_n against the wall simultaneously, thus exhibiting a rubbing behavior. Note in the figure that the robot seeks to choose velocities at contact that are approach the peaks of the information landscape, but are bounded by physical constraints. This yields the robot in exhibiting a rubbing behavior with the wall.

Stiffness Estimation

The results in Fig. 5 show that the robot optimally learns stiffness and dampening parameters of a ball by interacting with its gripper. The gripper excites perturbations in the normal forces between the gripper fingers and the ball by making and breaking contact, and pressing into the ball during contact engagement. The information landscape shown in Fig. 8b indicates that there is a locally optimal trade-off between the penetration ϕ_n length of the gripper finger and the ball surface and the penetration rate $\frac{\partial \phi_n}{\partial t}$ (equivalent to normal contact velocity v_n). We can observe that such normal force excitations can lead to excitations of the gradients of the sensor model detailed in Appendix B. The robot excites its grippers to pinch and release contact with the object by seeking trajectories that maximize penetration and pinching velocity as shown by the information landscape. This naturally results in the robot exhibiting a ‘squeezing and contracting’ behavior with the object.

VII. LIMITATIONS

The experimental design method in this paper has several limitations worth discussing below.

Deterministic Dynamics in a Dynamically Changing World. Our approach assumes that the dynamics of the system are deterministic, and we fit stochastic data with a sensor model that assumes zero mean Gaussian noise. Such assumptions can be especially tricky to deal with since we do not capture any uncertainties in our dynamics model, which may lead to under or over approximations of contact behaviors and limit the performance of our method in practice. Additionally, our dynamics model captures contact modes at certain body frames of the robot system; if the robot interacts with the world outside the body frame of interest, the desired synthesized behavior may not be accurately exhibited on the robot.

Nonlinear Optimization Landscape. In general, the complexity of the sensor model and dynamics of the robot create a challenging optimization landscape in the CA-MAP problem in Eq. 8 that is nonlinear and non-convex in nature. As a result, optimizing contact behaviors for parameter learning are likely

to be sub-optimal solutions. Consequently, finding effective behaviors in real-time settings can be difficult to plan for.

Contact Model Parameter Tuning. The experiments presented in this paper requires a great deal of parameter tuning in the contact model. For example, the hefting experiments demonstrated in this paper assumes that we have a priori knowledge of the material properties of the object of interest, which is required for the soft contact model implemented. This required careful tuning of the contact model parameters in order to obtain reasonably realistic contact model behaviors.

Limitations from Fisher Information. Optimizing over Fisher information, in accordance with Definition 3, reduces the variance of the parameter estimates by pushing the tight lower bound down. Future work will investigate stronger notions of variance reduction by means of reducing an *upper bound* of parameter variance. Additionally, Fisher information only provides notions of parameter *precision*, and no guarantees on accuracy, which therefore may lead to systematic error in parameter estimates over time.

VIII. CONCLUSION AND FUTURE WORK

This paper presents a contact-aware Fisher information maximization approach that results in emergent learning behaviors that precisely identifies parameters of interest. We develop the theoretical foundations of our problem formulation, and why maximizing over Fisher information yields the steepest descent in parameter learning, thus improves convergence of parameters through select contact behaviors. Last, we demonstrate our approaches ability to estimate physics parameters on a variety of robot settings, resulting in different emergent learning behaviors depending on the parameters being estimated.

Future work will explore how our Fisher information maximization approach will impact downstream policy learning and trajectory optimization. A limitation of the proposed work is that our approach structures learning using deterministic dynamics. Future work will explore how process noise impacts our contact-aware information maximization approach. Furthermore, future directions will also explore contact-aware information maximization approaches to more complex, higher dimensional robotic systems.

IX. ACKNOWLEDGEMENTS

This work is supported by the National Science Foundation under award NSF FRR 2238066. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

The authors extend their gratitude to Ethan K. Gordon and Michael Posa from the GRASP Lab at the University of Pennsylvania for all the insightful discussions and constructive feedback for this work.

REFERENCES

- [1] Mihai Anitescu and Florian Potra. Formulating dynamic multi-rigid-body contact problems with friction

- as solvable linear complementarity problems. *Nonlinear Dynamics*, 14, 03 1997. doi: 10.1023/A:1008292328909.
- [2] Nikolay Atanasov, Jerome Le Ny, Kostas Daniilidis, and George J. Pappas. Information acquisition with sensing robots: Algorithms and error bounds, 2013.
- [3] Alp Aydinoglu and Michael Posa. Real-time multi-contact model predictive control via admm. In *2022 International Conference on Robotics and Automation (ICRA)*, page 3414–3421. IEEE Press, 2022. doi: 10.1109/ICRA46639.2022.9811957. URL <https://doi.org/10.1109/ICRA46639.2022.9811957>.
- [4] R. Bellman and K.J. Åström. On structural identifiability. *Mathematical Biosciences*, 7(3):329–339, 1970. ISSN 0025-5564. doi: [https://doi.org/10.1016/0025-5564\(70\)90132-X](https://doi.org/10.1016/0025-5564(70)90132-X). URL <https://www.sciencedirect.com/science/article/pii/002555647090132X>.
- [5] Gregory S Chirikjian. *Stochastic Models, Information Theory, and Lie Groups, Volume 1: Classical Results and Geometric Methods*. Applied and Numerical Harmonic Analysis. Birkhäuser Boston, Boston, 1 edition, 2009. ISBN 081764802X.
- [6] Bradley Efron and David V. Hinkley. Assessing the accuracy of the maximum likelihood estimator: Observed versus expected fisher information. *Biometrika*, 65(3): 457–482, 1978. ISSN 00063444, 14643510. URL <http://www.jstor.org/stable/2335893>.
- [7] A F Emery and Aleksey V Nenarokomov. Optimal experiment design. *Measurement Science and Technology*, 9 (6):864, jun 1998. doi: 10.1088/0957-0233/9/6/003. URL <https://dx.doi.org/10.1088/0957-0233/9/6/003>.
- [8] Nima Fazeli, Roman Kolbert, Russ Tedrake, and Alberto Rodriguez. Parameter and contact force estimation of planar rigid-bodies undergoing frictional contact. *The International Journal of Robotics Research*, 36(13-14): 1437–1454, 2017. doi: 10.1177/0278364917698749. URL <https://doi.org/10.1177/0278364917698749>.
- [9] Valerii Fedorov. Optimal experimental design. *Wiley Interdisciplinary Reviews: Computational Statistics*, 2(5): 581–589, 2010.
- [10] T.I. Fenner and G. Loizou. Matrix bounds on the spectral condition number. *Linear Algebra and its Applications*, 8(2):157–178, 1974. ISSN 0024-3795. doi: [https://doi.org/10.1016/0024-3795\(74\)90053-6](https://doi.org/10.1016/0024-3795(74)90053-6). URL <https://www.sciencedirect.com/science/article/pii/0024379574900536>.
- [11] M. Gautier and W. Khalil. On the identification of the inertial parameters of robots. In *Proceedings of the 27th IEEE Conference on Decision and Control*, pages 2264–2269 vol.3, 1988. doi: 10.1109/CDC.1988.194738.
- [12] Jiawei Ge, Shange Tang, Jianqing Fan, Cong Ma, and Chi Jin. Maximum likelihood estimation is all you need for well-specified covariate shift, 2023. URL <https://arxiv.org/abs/2311.15961>.
- [13] Taylor Howell, Nimrod Gileadi, Saran Tunyasuvunakool, Kevin Zakka, Tom Erez, and Yuval Tassa. Predictive sampling: Real-time behaviour synthesis with mujoco, 2022. URL <https://arxiv.org/abs/2212.00541>.
- [14] Wanxin Jin, Alp Aydinoglu, Mathew Halm, and Michael Posa. Learning linear complementarity systems. In Roya Firoozi, Negar Mehr, Esen Yel, Rika Antonova, Jeannette Bohg, Mac Schwager, and Mykel Kochenderfer, editors, *Proceedings of The 4th Annual Learning for Dynamics and Control Conference*, volume 168 of *Proceedings of Machine Learning Research*, pages 1137–1149. PMLR, 23–24 Jun 2022. URL <https://proceedings.mlr.press/v168/jin22a.html>.
- [15] Leslie Pack Kaelbling, Michael L. Littman, and Anthony R. Cassandra. Planning and acting in partially observable stochastic domains. *Artificial Intelligence*, 101(1):99–134, 1998. ISSN 0004-3702. doi: [https://doi.org/10.1016/S0004-3702\(98\)00023-X](https://doi.org/10.1016/S0004-3702(98)00023-X). URL <https://www.sciencedirect.com/science/article/pii/S000437029800023X>.
- [16] Pradeep K. Khosla and Takeo Kanade. Parameter identification of robot dynamics. In *1985 24th IEEE Conference on Decision and Control*, pages 1754–1760, 1985. doi: 10.1109/CDC.1985.268838.
- [17] J. Kiefer. Optimum experimental designs. *Journal of the Royal Statistical Society: Series B (Methodological)*, 21(2):272–304, 1959. doi: <https://doi.org/10.1111/j.2517-6161.1959.tb00338.x>. URL <https://rss.onlinelibrary.wiley.com/doi/abs/10.1111/j.2517-6161.1959.tb00338.x>.
- [18] Nathan J. Kong, J. Joe Payne, James Zhu, and Aaron M. Johnson. Saltation matrices: The essential tool for linearizing hybrid dynamical systems. *Proceedings of the IEEE*, 112(6):585–608, June 2024. ISSN 1558-2256. doi: 10.1109/jproc.2024.3440211. URL <http://dx.doi.org/10.1109/JPROC.2024.3440211>.
- [19] Mikko Lauri and Risto Ritala. Planning for robotic exploration based on forward simulation. *Robotics and Autonomous Systems*, 83:15–31, September 2016. ISSN 0921-8890. doi: 10.1016/j.robot.2016.06.008. URL <http://dx.doi.org/10.1016/j.robot.2016.06.008>.
- [20] Susan J. Lederman and Roberta L. Klatzky. Extracting object properties through haptic exploration. *Acta Psychologica*, 84(1):29–40, 1993. ISSN 0001-6918. doi: [https://doi.org/10.1016/0001-6918\(93\)90070-8](https://doi.org/10.1016/0001-6918(93)90070-8). URL <https://www.sciencedirect.com/science/article/pii/0001691893900708>. Tactile Pattern Recognition.
- [21] Taeyoon Lee, Jaewoon Kwon, Patrick M. Wensing, and Frank C. Park. Robot model identification and learning: A modern perspective. *Annual Review of Control, Robotics, and Autonomous Systems*, 2023. doi: <https://doi.org/10.1146/annurev-control-061523-102310>. URL <https://www.annualreviews.org/content/journals/10.1146/annurev-control-061523-102310>.
- [22] Jonas Linder, Martin Enqvist, Fredrik Gustafsson, and Johan Sjöberg. Identifiability of physical parameters in systems with limited sensors. *IFAC Proceedings Volumes*, 47(3):6454–6459, 2014. ISSN 1474-6670. doi: <https://doi.org/10.3182/20140824-6-ZA-1003>.

02272. URL <https://www.sciencedirect.com/science/article/pii/S147466701642625X>. 19th IFAC World Congress.
- [23] Lennart Ljung. *System identification (2nd ed.): theory for the user*. Prentice Hall PTR, USA, 1999. ISBN 0136566952.
- [24] Zac Manchester and Scott Kuindersma. Variational contact-implicit trajectory optimization. In *Proceedings of International Symposium on Robotics Research (ISRR '17)*, pages 985 – 1000, December 2017.
- [25] Marius Memmel, Andrew Wagenmaker, Chuning Zhu, Patrick Yin, Dieter Fox, and Abhishek Gupta. Asid: Active exploration for system identification in robotic manipulation, 2024. URL <https://arxiv.org/abs/2404.12308>.
- [26] Igor Mordatch, Emanuel Todorov, and Zoran Popović. Discovery of complex behaviors through contact-invariant optimization. *ACM Trans. Graph.*, 31(4), July 2012. ISSN 0730-0301. doi: 10.1145/2185520.2185539. URL <https://doi.org/10.1145/2185520.2185539>.
- [27] Yusuke Ogawa, Gentiane Venture, and Christian Ott. Dynamic parameters identification of a humanoid robot using joint torque sensors and/or contact forces. 2015: 457–462, 02 2015. doi: 10.1109/HUMANOIDS.2014.7041401.
- [28] Mihir Parmar, Mathew Halm, and Michael Posa. Fundamental challenges in deep learning for stiff contact dynamics, 2021.
- [29] Chiara Pepi, Massimiliano Gioffrè, and Mircea Grigoriu. Bayesian inference for parameters estimation using experimental data. *Probabilistic Engineering Mechanics*, 60:103025, 2020. ISSN 0266-8920. doi: <https://doi.org/10.1016/j.probengmech.2020.103025>. URL <https://www.sciencedirect.com/science/article/pii/S0266892020300102>.
- [30] Samuel Pfrommer, Mathew Halm, and Michael Posa. Contactnets: Learning of discontinuous contact dynamics with smooth, implicit representations. In *Conference on Robot Learning*, 2020. URL <https://api.semanticscholar.org/CorpusID:221857185>.
- [31] Robert Platt, Russ Tedrake, Leslie Pack Kaelbling, and Tomas Lozano-Perez. Belief space planning assuming maximum likelihood observations. In *Robotics: Science and Systems*, 2010. URL <https://api.semanticscholar.org/CorpusID:2693863>.
- [32] Michael Posa, Cecilia Cantu, and Russ Tedrake. A direct method for trajectory optimization of rigid bodies through contact. *The International Journal of Robotics Research*, 33(1):69–81, 2014. doi: 10.1177/0278364913506757. URL <https://doi.org/10.1177/0278364913506757>.
- [33] C. Radhakrishna Rao. Minimum variance and the estimation of several parameters. *Mathematical Proceedings of the Cambridge Philosophical Society*, 43(2):280–283, 1947. doi: 10.1017/S0305004100023471.
- [34] Jonas Rothfuss, Bhavya Sukhija, Lenart Treven, Florian Dörfler, Stelian Coros, and Andreas Krause. Bridging the sim-to-real gap with bayesian inference, 2024. URL <https://arxiv.org/abs/2403.16644>.
- [35] Alexander Soen and Ke Sun. On the variance of the fisher information for deep learning, 2021. URL <https://arxiv.org/abs/2107.04205>.
- [36] David E. Stewart and Jeffrey C. Trinkle. An implicit time-stepping scheme for rigid body dynamics with coulomb friction. *Proceedings 2000 ICRA. Millennium Conference. IEEE International Conference on Robotics and Automation. Symposia Proceedings (Cat. No.00CH37065)*, 1:162–169 vol.1, 2000. URL <https://api.semanticscholar.org/CorpusID:10920958>.
- [37] Ke Sun and Frank Nielsen. Relative Fisher information and natural gradient for learning large modular models. In Doina Precup and Yee Whye Teh, editors, *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 3289–3298. PMLR, 06–11 Aug 2017. URL <https://proceedings.mlr.press/v70/sun17b.html>.
- [38] Balakumar Sundaralingam and Tucker Hermans. In-hand object-dynamics inference using tactile fingertips, 2021. URL <https://arxiv.org/abs/2003.13165>.
- [39] J. Swevers, C. Ganseman, D.B. Tukul, J. de Schutter, and H. Van Brussel. Optimal robot excitation and identification. *IEEE Transactions on Robotics and Automation*, 13(5):730–740, 1997. doi: 10.1109/70.631234.
- [40] Yuval Tassa, Tom Erez, and Emanuel Todorov. Synthesis and stabilization of complex behaviors through online trajectory optimization. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 4906–4913, 2012. doi: 10.1109/IROS.2012.6386025.
- [41] Emanuel Todorov, Tom Erez, and Yuval Tassa. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 5026–5033. IEEE, 2012. doi: 10.1109/IROS.2012.6386109.
- [42] Jean-François Tremblay, David Meger, Francois Hogan, and Gregory Dudek. Learning active tactile perception through belief-space control, 2023. URL <https://arxiv.org/abs/2312.00215>.
- [43] Andrew D. Wilson, Jarvis A. Schultz, and Todd D. Murphey. Trajectory synthesis for fisher information maximization. *IEEE Transactions on Robotics*, 30(6): 1358–1370, December 2014. ISSN 1941-0468. doi: 10.1109/tro.2014.2345918. URL <http://dx.doi.org/10.1109/TRO.2014.2345918>.
- [44] Andrew D. Wilson, Jarvis A. Schultz, Alex R. Ansari, and Todd D. Murphey. Real-time trajectory synthesis for information maximization using sequential action control and least-squares estimation. In *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4935–4940, 2015. doi: 10.1109/IROS.2015.7354071.
- [45] William Yang and Michael Posa. Impact-invariant control: Maximizing control authority during impacts, 2023.
- [46] Dongliang Zheng, Jack Ridderhof, Panagiotis Tsiotras,

and Ali akbar Agha-mohammadi. Belief space planning: A covariance steering approach, 2021. URL <https://arxiv.org/abs/2105.11092>.

APPENDIX A PROOFS

Fisher Information Steepest Descent Proof

Proof: We first approximate $\mathcal{L}(\mathcal{D}|\tau, \theta + \delta\theta)$ in Eq.11 by a second order Taylor expansion,

$$\begin{aligned} \mathcal{L}(\mathcal{D}|\tau, \theta + \delta\theta) &\approx \mathcal{L}(\mathcal{D}|\tau, \theta) + \nabla_{\theta} \mathcal{L}(\mathcal{D}|\tau, \theta)^T \delta\theta \\ &\quad + \frac{1}{2} \delta\theta^T \nabla_{\theta}^2 \mathcal{L}(\mathcal{D}|\tau, \theta) \delta\theta \end{aligned} \quad (14)$$

Solving for $\delta\theta$ can then be done with the Taylor Expansion as

$$\delta\theta^* = \arg \max_{\delta\theta} \nabla_{\theta} \mathcal{L}(\mathcal{D}|\tau, \theta)^T \delta\theta + \frac{1}{2} \delta\theta^T \nabla_{\theta}^2 \mathcal{L}(\mathcal{D}|\tau, \theta) \delta\theta \quad (15)$$

which has the following closed form solution

$$\delta\theta^* = \mathcal{F}(\mathcal{D}|\tau, \theta)^{-1} \nabla_{\theta} \mathcal{L}(\mathcal{D}|\tau, \theta). \quad (16)$$

APPENDIX B IMPLEMENTATION DETAILS

Relationship between Condition Number and Trace

To improve the condition number of the Fisher information matrix, $\kappa(\mathcal{F})$, we optimize over the trace of the Fisher information matrix (T-optimality) as condition numbers bounded by the trace of a nonsingular matrix [10], as

$$\begin{aligned} (||\mathcal{F}||^2 ||\mathcal{F}^{-1}||^2 - (n-2))^2 &\geq (\kappa(\mathcal{F}) + \kappa^{-1}(\mathcal{F}))^2 \\ &\geq \begin{cases} n \text{ even, } \frac{4}{n^2} ||\mathcal{F}||^2 ||\mathcal{F}^{-1}||^2 \\ n \text{ odd, } \frac{4}{n^2-1} (||\mathcal{F}||^2 ||\mathcal{F}^{-1}||^2 - 1) \end{cases} \end{aligned} \quad (17)$$

where $\mathcal{F} \in \mathbb{R}^{n \times n}$, and the Schur matrix norm $||\mathcal{F}|| = \text{tr}(\mathcal{F}^T \mathcal{F})$, for a nonsingular matrix \mathcal{F} . By maximizing over the trace of the Fisher information matrix in Eq. 7, we thus maximize the bounds of the condition number.

Contact Model

The contact sensor model implemented in this paper utilizes the following soft contact model,

$$\begin{bmatrix} \lambda_n \\ \lambda_t \end{bmatrix} = \begin{bmatrix} \max(0, -K\phi_n - C|v_n|) \\ \hat{v}_t \max(-\mu\lambda_n, -R|v_t|) \end{bmatrix} \quad (18)$$

where K , C , μ , and R are the material stiffness, material dampening, friction coefficient, and friction resistance parameter, respectively, ϕ_n is the signed distance function between the system body and object, and $v_c = [v_n, v_t]^T = \mathbf{J}_c(q)\dot{q}$ are the normal and tangent velocities at the contact frame with unit vectors \hat{v}_n and \hat{v}_t , respectively, and where $\mathbf{J}_c(q)$ is the contact Jacobian that maps joint velocities to the contact frame.

Experiment	K [N/m]	C [Ns/m]	μ	R [Ns/m]
Hefting	500	0	N/A	N/A
Rubbing	100	1	0.4	2.0
Pinching	800	10	N/A	N/A
Contouring	500	0	0	0

Table I: **Contact Model Parameters.** Soft contact model parameters for all experiment scenarios. **Bold** parameters the unknown true parameters.

Experiment Details

Here we detail the experiment settings discussed in this paper.

Predictive Sampling Parameters. All experimental setups were conducted in discrete time settings with a short time horizon $t_H = 10$, and an initial nominal plan $\Pi_{\text{init}} = \text{zeros}(t_h, n_a)$, where n_a is the number of controllable joints in the system. The $N = 10$ sampled control splines were constructed using a Cubic Hermite spline, as demonstrated by [13]. For all experiments we fix the sampling variance as $\Sigma = 1.0$. To prevent the control sequences from violating joint constraints during forward rollouts and execution on the robot, we clamp the control splines to only sample within a domain that satisfy joint constraints.

Mass Estimation Experiments. The Franka end effector is free to move within 1 dimensional domain $\mathcal{W} = [0, 0.6]$ m in the z direction with respect to world frame coordinates (at the Franka root), with bounded joint velocity constraints $v_{ee} \in [-1.0, 1.0]$ m/s. The robot starts at initial end effector position $p_{ee,0} = 0.2$ m with zero initial velocity. The Allegro hand attached to the Franka estimates the mass of a grasped tennis ball of mass 0.05 kg (the results in this paper scale this value by a factor of 100). We initialize the variance of the mass parameter as $\Sigma_{\theta} = 10 \text{ kg}^2$. The contact model parameters in this scenario is shown in Table I.

Friction Estimation Experiments. The Franka end effector is free to move within 2 dimensional domain $\mathcal{W} = [0, 1.0] \times [0.5, 1.0]$ m with respect to a local wall frame coordinate system at the base of the wall, with bounded joint velocity constraints $v_{ee} \in [-1.0, 1.0]$ m/s in both x and y directions. The robot starts at initial end effector position $p_{ee,0} = [0.5, 0.8]$ m in wall frame coordinates (oriented 0.8 m in the y direction from the root of the Franka, with its x axis pointing away from the wall) with zero initial velocity. The robot estimates the coefficient of friction between its end effector and the wall, with true parameter $\mu = 0.4$. We initialize the variance of the mass parameter as $\Sigma_{\theta} = 1$. The contact model parameters in this scenario is shown in Table I.

Material Property Estimation Experiments. A robot gripper enters a pregrasped state, where the gripper surface is positioned 0.4 cm away from the surface of the ball object with stiffness coefficient $K = 800 \text{ N/m}$, and dampening coefficient $C = 10 \text{ Ns/m}$. The gripper is free to move within gripper velocity constraints and prismatic joint constraints as per an Aloha 2 system. We initialize the variance of the mass parameter as $\Sigma_{\theta} = [100.0 \text{ (N/m)}^2, 10.0 \text{ (Ns/m)}^2]$.

Shape Estimation Experiments. A Franka is tasked with estimating the length $l = 0.126$ m and width $w = 0.05$ m block. on a table surface. The robot is constrained in the world frame z axis such that its motion is planar about the surface of the table. The robot is initialized at end effector position $p_{ee} = [0.5, -0.1, 0.08]$ meters with respect to the world frame, and is constrained to a planar domain $\mathcal{W} = [0.3, -0.15] \times [0.8, 0.15]$ m, with velocity constraints set to $v_{ee} \in [-0.5, 0.5]$ m/s for all axes. We initialize the variance of the mass parameter as $\Sigma_\theta = 10 \text{ cm}^2$. The contact model parameters in this scenario is shown in Table I.