

Sequential Experiment Design for Parameter Estimation of Nonlinear Systems using a Neural Network Approximator

Raksha Ramakrishna*, Yuqi Shao[†], György Dán*, Nicole Kringos[‡]

* Division of Network and Systems Engineering, KTH Royal Institute of Technology, Stockholm, Sweden

[†] KTH Royal Institute of Technology, Stockholm, Sweden

[‡] Department of Civil and Architectural Engineering, KTH Royal Institute of Technology, Stockholm, Sweden

Abstract—We consider the problem of sequential parameter estimation of a nonlinear function under the Bayesian setting. The designer can choose inputs for a sequence of experiments to obtain an accurate estimate of the system parameters based on observed outputs, while complying with a constraint on the expected outputs of the system. We quantify the accuracy of the obtained estimate in terms of the ℓ_2 norm. We propose to solve the problem by casting it as the problem of minimizing the Bayesian Mean Square Error (BMSE) of the parameter estimate subject to a constraint on the expected deviation of the output from the desired target value. We develop a greedy policy to solve the problem in the sequential setting, and we characterize the solution structure based on analytical results for the Gaussian case. For a computationally tractable update of the posterior, we propose the use of a surrogate model combined with approximate Bayesian computation. We evaluate the proposed approach on the use case of smart road compaction, where the goal is to estimate asphalt parameters while reaching the desired compaction level, by choosing the value of the loading pressure. Simulation results on a synthetic road compaction dataset show the efficacy of the proposed solution scheme in both parameter estimation and effective compaction of the road.

I. INTRODUCTION

An accurate estimate of the model parameters of equipment and materials has a variety of applications in chemistry, manufacturing and construction [1], [2]. In the lack of non-destructive testing methodologies that would be applicable post production, parameter estimates often have to be obtained during production in real-time, i.e., under time constraint. This is often the case, for example, during road compaction and metalworking. Importantly, the estimation of the parameters during the production phase should not be detrimental to production quality and is limited to interaction with the material or equipment for the purpose of production, by providing a sequence of inputs and observing the outputs.

The problem of designing inputs to a system with the purpose of estimating its parameters has been studied in the framework of optimal experimental design (OED) [3]. In OED, an optimal set of inputs is designed for a system described by a parameterized model, so as to obtain an accurate estimate of the model parameters. A commonly used measure of accuracy is the variance of the parameter estimate, and the task is to reach a target variance through as few experiments as possible, because experiments are considered costly [4], [5]. In many

applications, however, even evaluating a model of the system can be costly, e.g., if it involves finite element (FE) method simulations of material behavior [2].

The computational cost of evaluating the system model becomes of crucial importance for sequential OED, i.e., when the observations from previous experiments are to be used for the design of future experiments [6]. As the problem is essentially a sequential decision making problem under uncertainty, it can be cast as a dynamic programming problem, but existing solution approaches rely on Monte Carlo simulation methods, which are computationally prohibitive if the evaluation of the system model is computationally intensive [6], [7]. In this case even greedy approaches are computationally challenging due to the need for updating the posterior after each experiment, which typically involves many evaluations of the system model for different parameters [8], [9].

In this paper we formulate the problem of estimating model parameters through sequential experiments subject to the constraint of reaching a target output value by sequentially designing inputs to the system. We develop a greedy approach for average optimal experiment design, i.e., minimizing the variance of the estimate in a Bayesian setting, and propose a computationally lightweight approach for updating the posterior through a surrogate system model based on a neural network approximator used for posterior update via approximate Bayesian computation. A key novelty of our approach is its applicability to problems where the evaluation of the system model is computationally demanding. In addition, the use of the Bayesian mean-squared error as a cost is novel in the sequential OED setting. We show the potential of the proposed scheme through estimating the material parameters during road compaction, while maintaining the error between target and actual compaction bounded in expectation.

The rest of the paper is organized as follows. In Section II we review related work, in Section III we introduce the sequential parameter estimation problem in a Bayesian setting and in Section IV we formulate the input design problem as Bayesian-mean squared error (BMSE) minimization. In Section V we design a computationally lightweight greedy algorithm for sequential ED, and we show its efficacy on the use case of asphalt compaction in Section VI-D. Section VII concludes the paper.

II. RELATED WORK

As mentioned before, our work is closely related to the field of optimal experiment design in statistics. In system identification problems, it is also known as optimal input design [10] wherein the parameters of the system model to be estimated are treated as deterministic quantities and the input to the system model is shaped or designed in a way that the estimation error of the parameters is minimized [11]. In linear models, the inverse of error covariance matrix of the parameters is related to the spectrum of the input which can be parameterized and therefore used for shaping the input to minimize the trace (A-optimality) or determinant (D-optimality) of error covariance matrix. Sequential version of the problem is also studied with applications in trajectory planning. More recently, theoretical guarantees for non-linear system identification have also been provided in [12].

The Bayesian version of OED is reviewed in [13]. The most commonly used utility function in the Bayesian version is the KL divergence between the prior and posterior parameter distributions averaged over all the observations [13]. Since this utility is intractable in general, a variational approach to approximate the posterior distribution is undertaken in [14]. As another approach, authors in [15] introduce a lower bound on the utility and maximize it with respect to design criteria parameterized by a neural network. For linear models, i.e. Gaussian priors and posteriors in the Bayesian model, it is possible to obtain closed form expressions for optimal experiment design as shown in [16], which leads to the well-known criterion of D-optimality in experiments. For Bayesian non-linear state-space models authors in [17] minimize the posterior Cramer-Rao lower bound (PCRB) in lieu of maximizing a utility function based on Bayesian mean-squared error (BMSE).

The sequential setting of OED was studied in [18], where a single experiment is performed at a time, the posterior of the parameters is updated and the process continues until the allowed number of experiments are performed. The utility considered is the KL divergence between the prior and posterior of the parameter distribution. The overall problem is formulated as a dynamic program and accounts for a dynamic environment. Authors in [6] consider the same problem and address computational and numerical issues through approximate dynamic programming. More recent papers utilize deep learning for input design, e.g., a neural network was trained to represent the input design policy in a Bayesian setting in [19].

III. PROBLEM STATEMENT

We consider a sequential parameter estimation problem, with outputs denoted by $\mathbf{y}_t \in \mathcal{Y} \subset \mathbb{R}^N$ and inputs denoted by $\mathbf{x}_t \in \mathcal{X} \subset \mathbb{R}^M$. The generative model for the output at time instant t is of the form

$$\mathbf{y}_t = \phi(\mathbf{x}_t; \boldsymbol{\theta}) + \boldsymbol{\eta}_t, \quad (1)$$

where $\phi(\cdot; \boldsymbol{\theta})$ is a function with parameter $\boldsymbol{\theta} \in \mathbb{R}^K$. The term $\boldsymbol{\eta}_t$ is used to model both model uncertainty due to

approximation of the generative model and measurement noise due to environmental factors. It is assumed that $\boldsymbol{\eta}_t$ follows a zero-mean Gaussian distribution with known variance σ^2 , $\boldsymbol{\eta}_t \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbb{I})$. We consider that the family of functions that $\phi(\cdot)$ belongs to is known, but the true value $\boldsymbol{\theta}^*$ of the parameter is unknown. In addition there is a desired target output $\bar{\mathbf{y}}$ that the designer intends to meet with the appropriate value of the input \mathbf{x} and parameter $\boldsymbol{\theta}$.

The designer is allowed to perform at most T experiments, by providing inputs and observing the outputs. Let us denote the inputs and outputs until experiment t as

$$\mathbf{X}_t \triangleq [\mathbf{x}_1 \quad \mathbf{x}_2 \quad \dots \quad \mathbf{x}_t], \quad \mathbf{Y}_t \triangleq [\mathbf{y}_1 \quad \mathbf{y}_2 \quad \dots \quad \mathbf{y}_t] \quad (2)$$

We adopt a Bayesian approach, and hence assume that the true parameter is a realization of a random variable $\boldsymbol{\theta}$. The belief about the distribution of $\boldsymbol{\theta}$ right before experiment t is referred to as the prior and is denoted by $p_t(\boldsymbol{\theta})$, while the belief about the distribution of $\boldsymbol{\theta}$ after experiment t is referred to as the posterior distribution and is denoted by $p(\boldsymbol{\theta} | \mathbf{Y}_t; \mathbf{X}_t)$. Given the posterior distribution, it is well-known that the estimator of $\boldsymbol{\theta}$ that minimizes the mean-squared error is the mean of the posterior distribution, $\mathbb{E}[\boldsymbol{\theta} | \mathbf{Y}_t; \mathbf{X}_t]$ and the designer wants this mean to be as close to the true value $\boldsymbol{\theta}^*$ as possible [20] after T experiments. Furthermore, the designer also wants the expected deviation from target output $\bar{\mathbf{y}}$ to be below a certain threshold δ . Formally, the designer wants to solve the optimization problem

$$\min_{\mathbf{X}_T} \|\mathbb{E}[\boldsymbol{\theta} | \mathbf{Y}_T; \mathbf{X}_T] - \boldsymbol{\theta}^*\|_2^2 \quad (3)$$

subject to

$$\mathbb{E}_{\boldsymbol{\theta} \sim p_t(\boldsymbol{\theta})} \left[\|\phi(\mathbf{x}_t; \boldsymbol{\theta}) - \bar{\mathbf{y}}\|_2^2 \right] \leq \delta, \quad (4)$$

$$\mathbf{x}_t \in \mathcal{X} \quad 1 \leq t \leq T$$

In constraint (4), the expectation is taken with respect to $\boldsymbol{\theta}$ drawn from the prior distribution $p_t(\boldsymbol{\theta})$. In an offline setup, where \mathbf{x}_T is decided for all experiments at once, the prior distribution is not updated based on the observed outputs, i.e., $p_t(\boldsymbol{\theta}) = p_0(\boldsymbol{\theta}) \forall t$. In a sequential setup, where \mathbf{x}_t for experiment t is decided based on past observations \mathbf{Y}_{t-1} and inputs \mathbf{X}_{t-1} , the prior distribution becomes the most recent posterior, i.e., $p_t(\boldsymbol{\theta}) = p(\boldsymbol{\theta} | \mathbf{Y}_{t-1}; \mathbf{X}_{t-1})$. The motivation behind the formulation stems from the goal of minimizing the mean squared error. Note that the objective (3) is expressed as the norm of the difference between the best estimate $\mathbb{E}[\boldsymbol{\theta} | \mathbf{Y}_T; \mathbf{X}_T]$ and the true parameter $\boldsymbol{\theta}^*$. Nonetheless, the true parameter is a realization of the random variable $\boldsymbol{\theta}$, and is not observable.

IV. BAYESIAN MEAN-SQUARED ERROR FORMULATION

In order to cope with the unknown parameter, we propose to adopt a Bayesian approach. Following the Bayesian approach, the estimate $\mathbb{E}[\boldsymbol{\theta} | \mathbf{Y}_T; \mathbf{X}_T]$ of the parameter is based on the observed outputs \mathbf{Y}_T , and minimizes the mean-squared error in the Bayesian sense, i.e., the Bayesian MSE (BMSE). This formulation is similar to other works in OED where the

goal is to minimize a metric of parameter uncertainty. Here, that metric is the BMSE, i.e., the variance of the posterior distribution.

A. BMSE Problem Formulation

In order to express the BMSE with respect to an estimator $\hat{\theta}$ considering the average error over the observations, let us denote by \mathbf{Y} the random vectors corresponding to the observations and use $p(\mathbf{Y}; \mathbf{X}_T)$ to denote the distribution of \mathbf{Y} parameterized by inputs \mathbf{X}_T . We can then express the BMSE of $\hat{\theta}$ by marginalizing with respect to the observations \mathbf{Y} and the parameter θ by considering the following factorization of their joint distribution,

$$p(\theta, \mathbf{Y}; \mathbf{X}_T) = p(\theta|\mathbf{Y}; \mathbf{X}_T)p(\mathbf{Y}; \mathbf{X}_T) \quad (5)$$

Then, the BMSE for estimator $\hat{\theta}$ is given as

$$\text{BMSE} \triangleq \int \int \left\| \hat{\theta} - \theta \right\|_2^2 p(\theta|\mathbf{Y}; \mathbf{X}_T)p(\mathbf{Y}; \mathbf{X}_T)d\theta d\mathbf{Y}_T. \quad (6)$$

Now, we can define the minimum BMSE with respect to the estimator $\hat{\theta}$ as $C(\mathbf{X}_T) \triangleq \min_{\hat{\theta}} \text{BMSE}$,

$$\begin{aligned} C(\mathbf{X}_T) &= \iint \left\| \mathbb{E}[\theta|\mathbf{Y}; \mathbf{X}_T] - \theta \right\|_2^2 p(\theta|\mathbf{Y}; \mathbf{X}_T)p(\mathbf{Y}; \mathbf{X}_T)d\theta d\mathbf{Y} \\ &= \int \text{Tr}(C_{\theta|\mathbf{Y}})p(\mathbf{Y}; \mathbf{X}_T)d\mathbf{Y} \end{aligned} \quad (7)$$

where $\text{Tr}(C_{\theta|\mathbf{Y}})$ refers to the trace of the conditional covariance matrix of θ given \mathbf{Y} . Note that the minimum MSE is expressed with respect to the deterministic parameters of the distribution of \mathbf{Y} and parameters of the prior distribution of θ . In the case of Gaussian noise as in (1), the minimum MSE is defined by the input \mathbf{X}_T and noise variance σ^2 and parameters of $p_t(\theta)$.

A strategy to determine the input \mathbf{X}_T could be to minimize the quantity $C(\mathbf{X}_T)$ with respect to \mathbf{X}_T subject to the constraint pertaining to the target output \bar{y} in (4). Overall, the problem can be formulated from (7) as

$$\mathbf{X}_T^* = \arg \min_{\mathbf{X}_T} C(\mathbf{X}_T) \quad (8)$$

$$\begin{aligned} \text{subject to } & \mathbb{E}_{\theta \sim p_t(\theta)} \left[\|\phi(x_t; \theta) - \bar{y}\|_2^2 \right] \leq \delta, \quad (9) \\ & x_t \in \mathcal{X} \quad 1 \leq t \leq T \end{aligned}$$

B. Analytical Example

As an illustration, let us consider a linear measurement model with scalar parameters,

$$y = x\theta + \eta, \quad \eta \sim \mathcal{N}(0, \sigma^2), \quad (10)$$

a single experiment $T = 1$, and a Gaussian prior for parameter $\theta \sim \mathcal{N}(\mu_0, \sigma_0^2)$. Under these assumption, it is well known that the posterior distribution of θ is Gaussian with mean and variance given by [20]

$$\mathbb{E}[\theta|y; x] = \mu_{\theta|y} = \mu_0 + \frac{C_{\theta|y}}{\sigma^2} x(y - x\mu_0) \quad (11)$$

$$C_{\theta|y} = \left(\frac{1}{\sigma_0^2} + \frac{x^2}{\sigma^2} \right)^{-1}. \quad (12)$$




Fig. 1: Error curve $\mathbb{E}_{\theta} [(x\theta - \bar{y})^2]$ with $\mu_{\theta} = 1, \sigma_{\theta} = 0.1, \bar{y} = 0.25$, and threshold $\delta = 0.05$. The feasible set is $\mathcal{X}_{\delta} = [x_1, x_2] = [0.0266, 0.4280]$ and the optimal value is $x^* = 0.4280$ for the experiment.

The cost (7) can thus be expressed as

$$C(x) = \int C_{\theta|y}p(y; x)dy = C_{\theta|y} = \left(\frac{1}{\sigma_0^2} + \frac{x^2}{\sigma^2} \right)^{-1}, \quad (13)$$

and it does not depend on y . Also, the constraint pertaining to the target output (4) can be written as

$$\mathbb{E}_{\theta} [(x\theta - \bar{y})^2] = x^2\mathbb{E}_{\theta}[\theta^2] - 2x\bar{y}\mathbb{E}_{\theta}[\theta] + \bar{y}^2 \leq \delta \quad (14)$$

$$\implies x^2(\sigma_{\theta}^2 + \mu_{\theta}^2) - 2x\bar{y}\mu_{\theta} + \bar{y}^2 \leq \delta. \quad (15)$$

Thus, the input x has to lie in the set

$$\mathcal{X}_{\delta} = \{x | x^2(\sigma_{\theta}^2 + \mu_{\theta}^2) - 2x\bar{y}\mu_{\theta} + \bar{y}^2 \leq \delta\} \quad (16)$$

Note that the expectation is taken with respect to θ using the prior distribution. The overall problem (8) thus becomes

$$\min_x \left(\frac{1}{\sigma_0^2} + \frac{x^2}{\sigma^2} \right)^{-1} \quad (17)$$

$$\text{subject to } x^2(\sigma_{\theta}^2 + \mu_{\theta}^2) - 2x\bar{y}\mu_{\theta} + \bar{y}^2 \leq \delta. \quad (18)$$

Observe that the cost for minimization translates to maximizing x^2 subject to a quadratic inequality (18). If the discriminant of the quadratic is positive or zero, then we have a solution to (17). If the discriminant is negative, then the parabola in x lies entirely above the x -axis thereby rendering no solution to the problem in (17) subject to the constraint. For the discriminant to be positive, the following inequality must hold

$$4\bar{y}^2\mu_{\theta}^2 - 4(\sigma_{\theta}^2 + \mu_{\theta}^2)(\bar{y}^2 - \delta) \geq 0 \implies \delta \geq \frac{\bar{y}^2\sigma_{\theta}^2}{(\sigma_{\theta}^2 + \mu_{\theta}^2)}. \quad (19)$$

If the condition above is satisfied then the solution is

$$x^* = \left| \frac{\bar{y}\mu_{\theta}}{(\sigma_{\theta}^2 + \mu_{\theta}^2)} \right| + \frac{\sqrt{\delta(\sigma_{\theta}^2 + \mu_{\theta}^2) - \bar{y}^2\sigma_{\theta}^2}}{(\sigma_{\theta}^2 + \mu_{\theta}^2)}. \quad (20)$$

Fig. 1 shows the error curve, the threshold and the feasible set of inputs for a numerical example.

Remark 1: Observe that the optimal solution mainly depends on the parameters of the prior distribution of θ , the target output and the accuracy budget δ . Also note that the optimal solution x^* is at the boundary points of the feasible set, which we will make use of in the next section.

V. OPTIMAL INPUT DESIGN

Minimizing the BMSE with respect to input \mathbf{x} provides the optimal solution, but it is not always possible to express the cost in closed form as a function of \mathbf{x} . In what follows, we discuss alternative solution approaches that do not assume that the BMSE can be expressed in closed form, but still respect the constraints on the input at every experiment.

Generalizing (16), let us define the set of feasible values of \mathbf{x} given the prior distribution $p_t(\boldsymbol{\theta})$ as

$$\mathcal{X}_\delta^t = \{\mathbf{x} | \mathbb{E}_{\boldsymbol{\theta} \sim p_t(\boldsymbol{\theta})} [\|\phi(\mathbf{x}; \boldsymbol{\theta}) - \bar{\mathbf{y}}\|_2^2] \leq \delta, \mathbf{x} \in \mathcal{X}\}. \quad (21)$$

In order to satisfy the constraint, we have to ensure that $\mathbf{x} \in \mathcal{X}_\delta^t$ upon every experiment. Nonetheless, \mathcal{X}_δ^t depends on the most recent prior $p_t(\boldsymbol{\theta})$, hence in what follows we distinguish between offline and sequential input design.

A. Offline Input Design

We start with considering the offline case, i.e., observations become available only after all experiments are done, and thus the prior cannot be updated, i.e., $p_t(\boldsymbol{\theta}) = p(\boldsymbol{\theta})$, and hence the decision regarding all inputs can be made at once. As stated before, the cost in (8) is in general not analytically tractable, and as we will show in Section V-A2, the solution may be to use the same input T times. Hence we explore computationally efficient alternatives.

1) *Feasible Set Sampling*: Feasible set sampling considered here focuses on the feasibility of the solutions in (4). A feasible set of solutions \mathbf{X}_T can be obtained by sampling from the set \mathcal{X}_δ , but it is unclear how to best sample from \mathcal{X}_δ . A straightforward approach would be to sample uniformly at random from \mathcal{X}_δ . To obtain a better sampling strategy, recall that the solution in the Gaussian example shown in Section IV-B was to choose the input from the boundary of the feasible set. We thus consider sampling from a distribution that assigns a larger value at the boundary of the feasible set. Thus, we use the distribution

$$p_s(\mathbf{x}) = \begin{cases} \frac{1}{C_{\mathcal{X}_\delta}} \exp\left\{\frac{1}{\delta} \mathbb{E}_{\boldsymbol{\theta} \sim p(\boldsymbol{\theta})} [\|\phi(\mathbf{x}; \boldsymbol{\theta}) - \bar{\mathbf{y}}\|_2^2]\right\}, & \mathbf{x} \in \mathcal{X}_\delta \\ 0, & \text{otherwise} \end{cases}$$

$$C_{\mathcal{X}_\delta} = \int_{\mathbf{x} \in \mathcal{X}_\delta} \exp\left\{\frac{1}{\delta} \mathbb{E}_{\boldsymbol{\theta} \sim p(\boldsymbol{\theta})} [\|\phi(\mathbf{x}; \boldsymbol{\theta}) - \bar{\mathbf{y}}\|_2^2]\right\} d\mathbf{x} \quad (22)$$

This strategy reduces the entropy of the posterior distribution since, on average, conditioning reduces entropy. We use this offline strategy as a baseline in our evaluation.

2) *Analytical Example*: We continue with the example from Section IV-B, focusing now on the case with T experiments. We can concatenate the observations, $\mathbf{y}^\top \triangleq [y_1 \ y_2 \ \dots \ y_T]$, to arrive at the observation model

$$\mathbf{y} = \mathbf{x}\boldsymbol{\theta}, \quad (23)$$

where the inputs $\mathbf{x}^\top \triangleq [x_1 \ x_2 \ \dots \ x_T]$ need to be determined in an offline manner, i.e., all the inputs are decided at

once before making any observations. As before, the posterior distribution of θ is Gaussian with mean and variance as

$$\mathbb{E}[\theta | \mathbf{y}; \mathbf{x}] = \mu_{\theta | \mathbf{y}} = \mu_\theta + \frac{C_{\theta | \mathbf{y}}}{\sigma^2} \mathbf{x}^\top (\mathbf{y} - \mathbf{x}\mu_\theta) \quad (24)$$

$$C_{\theta | \mathbf{y}} = \left(\frac{1}{\sigma_\theta^2} + \frac{\|\mathbf{x}\|_2^2}{\sigma^2} \right)^{-1}. \quad (25)$$

The cost for minimization is $C(\mathbf{x}) = C_{\theta | \mathbf{y}}$. The constraints are decoupled with respect to each observation y_i , thus we have the overall optimization problem, which as before translates to maximizing with respect to \mathbf{x} ,

$$\max_{\mathbf{x}} \left(\frac{1}{\sigma_\theta^2} + \frac{\|\mathbf{x}\|_2^2}{\sigma^2} \right) \quad (26)$$

$$\text{subject to } x_t^2(\sigma_\theta^2 + \mu_\theta^2) - 2x_t\bar{y}\mu_\theta + \bar{y}^2 \leq \delta, \quad \forall t. \quad (27)$$

Note that the objective (26) can be decomposed with respect to each input x_i . Due to the absence of constraints that couple the inputs of different experiments, we get the same solution $x_t = x^* \quad \forall t$, where x^* is given in (20). Thus, a sufficiently diverse set of outputs cannot be obtained to change the posterior distribution significantly. In addition, observe that x^* , as in Section IV-B for $T = 1$, lies at the boundary of the feasible set.

Feasible Set Sampling: The expected error is given in the left-hand side (LHS) of (15) and the feasible set of x in (16) can be rewritten as

$$\mathcal{X}_\delta = \left[\frac{\bar{y}\mu_\theta}{(\sigma_\theta^2 + \mu_\theta^2)} + \frac{\sqrt{\delta(\sigma_\theta^2 + \mu_\theta^2) - \bar{y}^2\sigma_\theta^2}}{(\sigma_\theta^2 + \mu_\theta^2)}, \frac{\bar{y}\mu_\theta}{(\sigma_\theta^2 + \mu_\theta^2)} - \frac{\sqrt{\delta(\sigma_\theta^2 + \mu_\theta^2) - \bar{y}^2\sigma_\theta^2}}{(\sigma_\theta^2 + \mu_\theta^2)} \right], \quad (28)$$

with the upper and lower limits switched if $\bar{y}\mu_\theta < 0$. The sampling distribution can be written as

$$p_s(x) = \begin{cases} \frac{\exp\left\{\frac{1}{\delta}(x^2(\sigma_\theta^2 + \mu_\theta^2) - 2x\bar{y}\mu_\theta + \bar{y}^2)\right\}}{C_{\mathcal{X}_\delta}}, & x \in \mathcal{X}_\delta \\ 0, & \text{otherwise} \end{cases} \quad (29)$$

with $C_{\mathcal{X}_\delta}$ evaluated as in (22).

B. Sequential Input Design

We now turn to the sequential case, where the prior distribution of the parameters can be updated based on the outcome of previous experiments, before the next input is chosen, i.e., the posterior parameter distribution obtained after $t - 1$ experiments serves as the prior distribution for choosing x_t ,

$$p_t(\boldsymbol{\theta}) = p(\boldsymbol{\theta} | \mathbf{Y}_{t-1}; \mathbf{X}_{t-1}). \quad (30)$$

Ideally, one would solve (8) in the sequential setting by formulating it as a dynamic programming problem [18]. This involves solving the problem by backward induction to obtain a policy to determine the set of inputs \mathbf{X}_T^* . The following result follows from basic principles of dynamic programming.

Proposition 1: Consider experiment T with prior $p_T(\boldsymbol{\theta}) = p(\boldsymbol{\theta}|\mathbf{Y}_{T-1}; \mathbf{X}_{T-1})$. The optimal input \mathbf{x}_T^* is the solution to

$$\min_{\mathbf{x}_T} C(\mathbf{X}_T) \text{ s.t. } \mathbb{E}_{\boldsymbol{\theta} \sim p_T(\boldsymbol{\theta})} \left[\|\boldsymbol{\phi}(\mathbf{x}_T; \boldsymbol{\theta}) - \bar{\mathbf{y}}\|_2^2 \right] \leq \delta \quad (31)$$

Solving with backward induction, at experiment $t < T$, with prior $p(\boldsymbol{\theta}|\mathbf{Y}_{t-1}, \mathbf{X}_{t-1})$ and optimal inputs $\mathbf{x}_{t+1}^*, \dots, \mathbf{x}_T^*$, the cost $C(\mathbf{X}_T)$ has to be evaluated by computing multi-dimensional integrals with respect to \mathbf{Y}_T . Therefore, for addressing the multi-stage problem, in what follows we propose to adopt a greedy approach.

The greedy approach for the sequential input design consists of finding the optimal solution for experiment t given prior $p_t(\boldsymbol{\theta})$ and ignoring the terms in the cost comprising of future observations and inputs. In particular, at experiment t we consider the cost

$$C(\mathbf{x}_t) = \int \text{Tr}(C_{\boldsymbol{\theta}|\mathbf{Y}_t}) p(\mathbf{Y}_t; \mathbf{X}_t) d\mathbf{Y}_t. \quad (32)$$

From (30), note that the prior distribution $p_t(\boldsymbol{\theta})$ for experiment t was obtained based on the previous observations and inputs. This means that the posterior $p(\boldsymbol{\theta}|\mathbf{Y}_t; \mathbf{X}_t)$ and its mean $\mathbb{E}[\boldsymbol{\theta}|\mathbf{Y}_t; \mathbf{X}_t]$ have previous observations \mathbf{Y}_{t-1} and inputs \mathbf{X}_{t-1} as constants. Then $p(\mathbf{Y}_t; \mathbf{X}_t) = p(\mathbf{y}_t; \mathbf{x}_t) \prod_{i=1}^{t-1} \delta(\mathbf{y}_i - \tilde{\mathbf{y}}_i)$ where $\tilde{\mathbf{y}}_i$ are the observed values and

$$p(\mathbf{y}_t; \mathbf{x}_t) = \int p(\mathbf{y}_t|\boldsymbol{\theta}; \mathbf{x}_t) p_t(\boldsymbol{\theta}) d\boldsymbol{\theta} \quad (33)$$

$$= \int p(\mathbf{y}_t|\boldsymbol{\theta}; \mathbf{x}_t) p(\boldsymbol{\theta}|\mathbf{Y}_{t-1}; \mathbf{X}_{t-1}) d\boldsymbol{\theta}. \quad (34)$$

Also, since observations are independent of past observations given the parameter $\boldsymbol{\theta}$, the cost $C(\mathbf{x}_t)$ only depends on the input at experiment t , i.e.,

$$C(\mathbf{x}_t) = \int \text{Tr}(C_{\boldsymbol{\theta}|\mathbf{y}_t}) p(\mathbf{y}_t; \mathbf{x}_t) d\mathbf{y}_t. \quad (35)$$

Thus, following the greedy approach the input \mathbf{x}_t^G at experiment t is the solution to

$$\mathbf{x}_t^* = \arg \min_{\mathbf{x}_t} C(\mathbf{x}_t) \quad (36)$$

$$\text{subject to } \mathbb{E}_{\boldsymbol{\theta} \sim p_t(\boldsymbol{\theta})} \left[\|\boldsymbol{\phi}(\mathbf{x}; \boldsymbol{\theta}) - \bar{\mathbf{y}}\|_2^2 \right] \leq \delta. \quad (37)$$

1) *Analytical Example:* Continuing with the analytical example in the sequential setting, observe that the initial prior is Gaussian and the observation y and θ are jointly Gaussian. Thus, at experiment t , the prior $p_t(\boldsymbol{\theta})$ is Gaussian and has mean $\mu_{\theta, t-1}$ and variance $\sigma_{\theta, t-1}^2$ given by

$$\mu_{\theta, t-1} = \mu_{\theta, t-2} + \frac{\sigma_{\theta, t-1}^2}{\sigma^2} x_{t-1} (y_{t-1} - x_{t-1} \mu_{\theta, t-2}) \quad (38)$$

$$\sigma_{\theta, t-1}^2 = \left(\frac{1}{\sigma_{\theta, t-2}^2} + \frac{x_{t-1}^2}{\sigma^2} \right)^{-1}. \quad (39)$$

The cost for experiment t is the same as in the single experiment case (eqn. (13)) and does not depend on observation

\mathbf{y}_t . Thus, the optimal input \mathbf{x}_t^* is given by solving (17) with $\mu_{\theta} = \mu_{\theta, t-1}$ and $\sigma_{\theta}^2 = \sigma_{\theta, t-1}^2$ and is given by

$$\mathbf{x}_t^* = \frac{|\bar{y} \mu_{\theta, t-1}|}{(\sigma_{\theta, t-1}^2 + \mu_{\theta, t-1}^2)} + \frac{\sqrt{\delta(\sigma_{\theta, t-1}^2 + \mu_{\theta, t-1}^2) - \bar{y}^2 \sigma_{\theta, t-1}^2}}{(\sigma_{\theta, t-1}^2 + \mu_{\theta, t-1}^2)}. \quad (40)$$

For experiment $t+1$, the posterior of $\boldsymbol{\theta}$ is updated using the recursive formulae for mean in (38) and variance (39) and the optimal input is also obtained recursively using (40).

C. Training of Surrogate Model for Posterior Update

To efficiently update the posterior distribution after observing the outputs, a closed form expression of the prior distribution, evaluation of the distribution $p(\mathbf{y}_t|\mathbf{x}_t, \boldsymbol{\theta})$ and computation of the function $\boldsymbol{\phi}(\mathbf{x}; \boldsymbol{\theta})$ are needed. In lack of these, the computation of the posterior distribution has to be based on a numerical method. To address this problem, we propose to use a neural network approximator of the function $\boldsymbol{\phi}(\mathbf{x}; \boldsymbol{\theta})$ as a surrogate model, combined with approximate Bayesian computation [21] to update the posterior distribution. For illustration, consider the sequential setting, where the update of the posterior distribution after observing the output \mathbf{y}_t of the experiment t is

$$p(\boldsymbol{\theta}|\mathbf{Y}_t, \mathbf{X}_t) = \frac{p(\mathbf{y}_t|\mathbf{x}_t, \boldsymbol{\theta}) p(\boldsymbol{\theta}|\mathbf{Y}_{t-1}, \mathbf{X}_{t-1})}{\int p(\mathbf{y}_t|\mathbf{x}_t, \boldsymbol{\theta}) p(\boldsymbol{\theta}|\mathbf{Y}_{t-1}, \mathbf{X}_{t-1}) d\boldsymbol{\theta}} \quad (41)$$

In many applications the prior $p(\boldsymbol{\theta}|\mathbf{Y}_{t-1}, \mathbf{X}_{t-1})$ is not given in closed form and the distribution $p(\mathbf{y}_t|\mathbf{x}_t, \boldsymbol{\theta})$ is difficult to evaluate, especially if the function $\boldsymbol{\phi}(\mathbf{x}; \boldsymbol{\theta})$ takes a long time to compute.

The surrogate model is denoted by $\boldsymbol{\phi}_{\omega}(\mathbf{x}; \boldsymbol{\theta})$, where ω are the parameters of the neural network. The surrogate model is trained utilizing data generated by the original model $\boldsymbol{\phi}(\mathbf{x}; \boldsymbol{\theta})$, and is done offline, prior to performing the experiments. The surrogate $\boldsymbol{\phi}_{\omega}(\mathbf{x}; \boldsymbol{\theta})$ can be used for posterior update in both offline and sequential input design scenarios.

Given the surrogate, we perform the update of the posterior distribution when a closed form expression does not exist using approximate Bayesian computing [21], which is a rejection sampling approach. Algorithm 1 details the steps to produce samples of $\boldsymbol{\theta}$ from posterior distribution $p(\boldsymbol{\theta}|\mathbf{Y}_t; \mathbf{X}_t)$, given prior $p_t(\boldsymbol{\theta})$, input \mathbf{x}_t , and corresponding observation \mathbf{y}_t . The approach involves repeatedly generating a sample $\boldsymbol{\theta}'$ from the prior distribution $p_t(\boldsymbol{\theta})$ and using the sampled parameter to generate a sample \mathbf{z}_t using the surrogate model, $\mathbf{z}_t = \boldsymbol{\phi}_{\omega}(\mathbf{x}_t; \boldsymbol{\theta}')$, in lieu of simulating from the actual function $\boldsymbol{\phi}(\boldsymbol{\theta}'; \mathbf{x}_t)$. Then, the generated sample \mathbf{z}_t is compared with the observation \mathbf{y}_t using $\rho(\mathbf{y}_t, \mathbf{z}_t) = \|\mathbf{y}_t - \mathbf{z}_t\|_2^2$. If $\rho(\mathbf{y}_t, \mathbf{z}_t) \leq \epsilon$, the sample $\boldsymbol{\theta}'$ is accepted and added to the set of samples, Θ , that will constitute the posterior. The procedure is repeated until there are enough samples (N_s) to form an empirical posterior distribution.

For the offline input design case, the prior is the same for all experiments, i.e. $p_t(\boldsymbol{\theta}) = p(\boldsymbol{\theta}) \quad \forall t$. Therefore, given observations \mathbf{Y}_T and inputs \mathbf{X}_T , samples $\mathbf{z}_1, \dots, \mathbf{z}_T$

Algorithm 1: Approximate Bayesian Computing

Data: Observation \mathbf{y}_t , Input \mathbf{x}_t , tolerance level ϵ , number of samples N_s

Result: Posterior sample set

$$\Theta = \{\theta | \theta \sim p(\theta | \mathbf{Y}_t; \mathbf{X}_t)\}$$

```

1  $\Theta = \emptyset$ ;
2 while  $|\Theta| < N_s$  do
3   Sample  $\theta' \sim p_t(\theta)$ ;
4   Generate  $\mathbf{z}_t = \phi_\omega(\theta'; \mathbf{x}_t)$ ;
5   if  $\rho(\mathbf{y}_t, \mathbf{z}_t) \leq \epsilon$  then
6      $\Theta = \Theta \cup \theta'$ ;
7   end
8 end
  
```

are generated as before using $\theta' \sim p(\theta)$. Then, we use $\rho(\mathbf{Y}_T, \mathbf{Z}_T) = \sum_{t=1}^T \|\mathbf{y}_t - \mathbf{z}_t\|_2^2$ to accept or reject θ' using a new tolerance value $\tilde{\epsilon}$ which we set to be $\tilde{\epsilon} = \epsilon T$ to have the same tolerance per experiment as in the sequential design case.

VI. APPLICATION TO INTELLIGENT ROAD COMPACTION

In what follows we show how the considered problem applies to the use case of intelligent road compaction [22]. During road compaction, compaction actions (speed, vibration) are applied on asphalt, modeled as an elastoplastic material with uncertain parameters, so as to achieve a desired deformation, called permanent strain. Considering a contiguous stretch of asphalt of homogeneous material as a sequence of T sections, compaction can be modeled as T subsequent experiments. Learning the parameters of the compacted asphalt, which is the aim of the experiment, in turn allows to optimize predictive maintenance, thereby reducing maintenance costs.

A. Mathematical model of compaction simulation

The compaction models of asphalt in the literature [2] are quite complex and require the knowledge of several parameters in order to simulate the effect of compaction. In this paper, we use a bilinear elastoplastic material model along with site geometry and boundary conditions. The process of compaction can be modeled as a loading pressure (stress) x that leads to strain ε [23]. Due to the plastic nature of the material there is a plastic strain, which leads to permanent deformation. For simplicity, we consider the input in experimental design to be the stress x . For a bilinear elastoplastic material, the relationship between stress x (in MPa) and permanent deformation (in m) can be described using three parameters: yield point σ_{yield} , linear elastic modulus E_e and the isotropic tangent modulus E_{Tiso} . Thus, $\theta = [\sigma_{\text{yield}} \ E_e \ E_{\text{Tiso}}]^\top$. In the elastic region, there is no deformation since the removal of load reverses the strain. However, after the stress reaches the yield point σ_{yield} , the rate of increase in strain changes according to E_{Tiso} .

In one dimension, we can write the relationship between the input stress x_t and the expected resulting output, i.e., plastic




Fig. 2: Compaction simulation plastic displacement field generated by COMSOL [24] with parameters $\sigma_{\text{yield}} = 0.2$, $E_e = 3$, $E_{\text{Tiso}} = 2$ and loading pressure $x = 2$. (all values in MPa).




Fig. 3: Geometry and Load by Compactor

deformation as

$$\phi(x_t; \theta) = \frac{x_t - \sigma_{\text{yield}}}{y_0 \left(\frac{1}{E_{\text{Tiso}}} - \frac{1}{E_e} \right)}. \quad (42)$$

where y_0 is the original depth of the material before compaction.

In two dimensions, the displacement field in the material can only be obtained by using finite element (FE) methods, which require one to numerically solve a system of differential equations that govern the spatial variation in stress, and compute the strain, which leads to displacement or deformation in 2D space. An example of displacement field generated using COMSOL Multiphysics [24] is shown in Fig. 2, assuming that the measured permanent deformation corresponds to the displacement at the mid-point of the compactor in the direction perpendicular to the compactor wheel.

B. Generation of Synthetic Data

We simulate the transverse direction of the road that the compactor is passing along. Assuming the contact length is $2m$ as a typical width of a roller, and the contact width is $0.1m$ in the longitudinal direction. We simulated a cross section of a road by 2D solid mechanics with plane strain setting. For simplicity, there are only 2 layers: stiff underlying soil layer and compaction material layer. The left and bottom boundaries are fixed, and the right boundary is a symmetric boundary which constrains the movement in X -axis. The geometry is shown in Fig. 3. For simplicity, we assume that the stiff soil layer is linear elastic and has a Young Modulus $E = 300$ MPa and Poisson's ratio $v = 0.35$. The compaction material layer is bilinear elastoplastic with a constant Poisson's ratio of $v = 0.15$, and the parameters to be estimated are $\theta = [\sigma_{\text{yield}} \ E_e \ E_{\text{Tiso}}]^\top$. Since the constitutive model we used is not dependent on loading time, we focus on varying

loading pressure magnitude alone. We used COMSOL Multiphysics [24] for the FE method computations, the response contains the plastic deformation of the compaction material top surface. Treating the result of the FE method as the function ϕ , for given loading pressure x and parameters θ we obtain the observed deformation as

$$y_t = \phi(x_t; \theta) + \eta_t, \quad \eta_t \sim \mathcal{N}(0, \gamma^2). \quad (43)$$

C. Training the Surrogate

In order to obtain a surrogate model $\phi_\omega(x_t; \theta)$, we created a data set using uninformative priors for the parameters, $p(\theta)$, i.e., uniform and independent distributions, as discussed next. The prior of the yield point σ_{yield} is uniform, $\sigma_{\text{yield}} \sim \mathcal{U}[\underline{\sigma}, \bar{\sigma}]$, and is independent of the other parameters. The parameters $E_e \in [\underline{E}_e, \bar{E}_e]$ and $E_{\text{Tiso}} \in [\underline{E}_{\text{Tiso}}, \bar{E}_{\text{Tiso}}]$ has to satisfy $E_{\text{Tiso}} < E_e$. This is incorporated into the prior distribution by having

$$p(E_e, E_{\text{Tiso}}) = \begin{cases} 1/C, & (E_{\text{Tiso}}, E_e) \in \mathcal{E} \\ 0, & \text{else} \end{cases} \quad (44)$$

where C is the area of the region \mathcal{E} where $E_{\text{Tiso}} < E_e$

$$C = \int_{E_e=\underline{E}_e}^{\bar{E}_e} \int_{E_{\text{Tiso}}=\underline{E}_{\text{Tiso}}}^{\bar{E}_{\text{Tiso}}} \mathbb{I}\{E_{\text{Tiso}} < E_e\} dE_e dE_{\text{Tiso}} \quad (45)$$

Overall the prior distribution is

$$p(\theta) = p(\sigma_{\text{yield}})p(E_e, E_{\text{Tiso}}) = \frac{1}{(\bar{\sigma} - \underline{\sigma})C}, \quad (46)$$

where $\sigma_{\text{yield}} \in [\underline{\sigma}, \bar{\sigma}]$, $(E_{\text{Tiso}}, E_e) \in \mathcal{E}$. We used the data set for training a feed-forward neural network with three hidden layers of 10 neurons each. The activation function used is the rectified linear unit (ReLU). The parameters ω of the surrogate model ϕ_ω were learned using mini-batch stochastic gradient with RMSprop [25], i.e., the learning rate was divided by the running average of the recent gradient magnitude. The learning rate used was 0.001, batch size was 50 and 1500 epochs. The test accuracy (RMSE) was 10^{-3} .

D. Numerical Results

For the evaluation we used the same prior parameter distribution as for creating the training data set. A histogram of samples drawn from the prior distribution is shown in Fig. 4. The input x_t (in MPa) can taken values in \mathcal{X} , $x_t \in \mathcal{X} = \{1, 1.05, \dots, 2.05\}$. The observation y_t is obtained by adding Gaussian measurement noise with standard deviation 1×10^{-3} to the output of the COMSOL model. The target deformation is $\bar{y} = -0.03$ where the sign indicates compaction of material. For the constraint, we used $\delta = 5 \times 10^{-4}$, which is lower than the test accuracy of the surrogate model. For each evaluation, the true value of the parameter θ^* was chosen from the prior distribution. We run the algorithm 10 times with different set of true parameters. For each set of true parameters, $T = 20$ compaction experiments are performed. In the offline setting, the value of input x_t is chosen at random from the feasible set \mathcal{X}_δ . We consider three different sampling strategies for x_t . First, we sample uniformly at random from the feasible set.

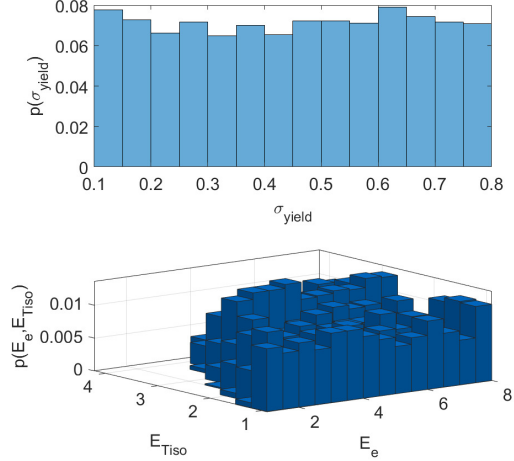


Fig. 4: Histogram of prior distribution. $\sigma_{\text{yield}} \sim U[0.1, 0.8]$. Parameters E_e, E_{Tiso} are jointly distributed according to (46) with $E_e \in [1, 8], E_{\text{Tiso}} = [1, 4]$.

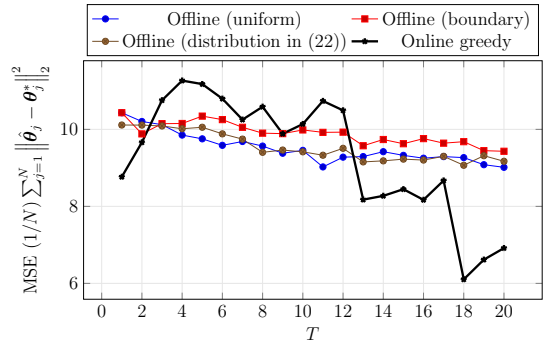


Fig. 5: Estimation performance using online and offline solution approaches. The estimator is posterior mean $\mathbb{E}[\theta | \mathbf{Y}_T; \mathbf{X}_T]$ after T experiments.

Secondly, we sample from the set of points that lie on the boundary of the feasible set. In this problem, it translates to sampling from a set with two values where the average (with respect to prior $p(\theta)$ deviation from the target \bar{y} is maximum. Thirdly, we sample the input x_t from the distribution in (22). In the sequential setting, the cost in (37) is evaluated numerically for an input x_t using Monte-Carlo integration method [26] with 2000 evaluation points due to the absence of closed form expressions. In both the sequential and offline methods, approximate Bayesian computation (ABC) is used for posterior update. For both types of methods, the estimator is the posterior mean, $\hat{\theta} = \mathbb{E}[\theta | \mathbf{Y}_T; \mathbf{X}_T]$.

Fig. 5 shows the estimation performance for the offline and the sequential greedy input design methods. The mean-squared error, $(1/N) \sum_{j=1}^N \|\hat{\theta}_j - \theta_j^*\|_2^2$, is plotted for different values of T . As expected, sequential input design outperforms all the offline methods due to the feedback in the form of sequential observations to improve the estimator. The behaviour of the MSE of the online greedy method seems erratic due to limited number of trials.

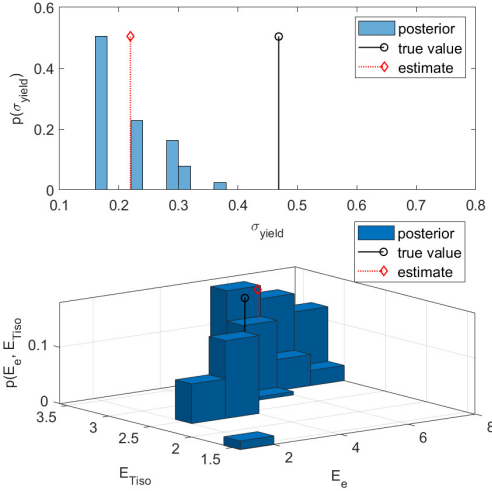


Fig. 6: Posterior distribution, actual and estimated parameters after $T = 20$ experiments where inputs $\{x_t\}_{t=1}^T$ are solutions to the online greedy problem (37).




Fig. 7: Input loading pressure x_t obtained by solving the greedy problem (37) in the online setting and corresponding observed plastic deformation y_t which is close to the desired target value \bar{y} .

For the sequential greedy input design, the posterior distribution after T experiments is plotted in Fig. 6. We see that the mean of the posterior distribution is close to the actual value for E_e, E_{Tiso} , the posterior distribution is sharper compared to the prior, and its weight is concentrated closer to the actual value, thereby effectively reducing uncertainty in parameter estimate. Fig. 7 shows the input load x_t obtained by solving the greedy problem (37) and the corresponding observed permanent deformation y_t . We see a close proximity to the desired target compaction \bar{y} , which shows the efficacy of sequential input design. The variation in the value of the input and the corresponding observations are a result of

competing objectives of reducing the uncertainty of parameters and constraining to meet the target value; the variation of the input values helps improve the parameter estimates. The observed permanent deformation can be made closer to the target value with a stricter constraint, i.e., a lower value of δ .

VII. CONCLUSIONS AND FUTURE WORK

We investigated the problem of sequential parameter estimation of a nonlinear function under the Bayesian setting while also meeting the system output constraints by optimally choosing the inputs to the system. We formulate this problem as a minimization of the Bayesian mean-squared error (BMSE) under constraints, which is a novelty of our paper, and we put forward a greedy solution for the sequential estimation problem. The proposed use of a neural network surrogate provides computational ease compared to other methods in the literature for the posterior update. In the road compaction case, it is found that the greedy approach is effective in accurate estimation of parameters as well as compaction of the road. Future work involves analytical development of input design methods using approximate dynamic programming, investigate deep learning based decision models for sequential input design under output constraints, and investigation of generalized surrogate development methods for unseen operating conditions.

REFERENCES

- [1] L. Fleitmann, J. Pyschik, L. Wolff, J. Schilling, and A. Bardow, "Optimal experimental design of physical property measurements for optimal chemical process simulations," *Fluid Phase Equilibria*, vol. 557, p. 113420, 2022.
- [2] F. Chen, R. Balieu, and N. Kringos, "Thermodynamics-based finite strain viscoelastic-viscoplastic model coupled with damage for asphalt material," *International Journal of Solids and Structures*, vol. 129, pp. 61–73, 2017.
- [3] A. F. Emery and A. V. Nenarokomov, "Optimal experiment design," *Measurement Science and Technology*, vol. 9, no. 6, p. 864, 1998.
- [4] X. Huan and Y. Marzouk, "Simulation-based optimal bayesian experimental design for nonlinear systems," *J. Comput. Phys.*, vol. 232, p. 288–317, 2013.
- [5] C. Agrell and K. R. Dahl, "Sequential bayesian optimal experimental design for structural reliability analysis," *Statistics and Computing*, vol. 31, no. 3, pp. 1–29, 2021.
- [6] X. Huan and Y. M. Marzouk, "Simulation-based optimal bayesian experimental design for nonlinear systems," *Journal of Computational Physics*, vol. 232, no. 1, pp. 288–317, 2013.
- [7] P. Carlin, J. B. Bradley, Kadane, and A. E. Gelfand, "Approaches for optimal sequential decision analysis in clinical trials," *Biometrics*, vol. 54, p. 964–975, 1998.
- [8] A. Solonen, H. Haario, and M. Laine, "Simulation-based optimal design using a response variance criterion," *J. Comput. Graph. Stat.*, vol. 21, pp. 234–252.
- [9] C. C. Drovandi, J. McGree, and A. N. Pettitt, "Sequential monte carlo for bayesian sequentially designed experiments for discrete data," *Comput. Stat. Data Anal.*, vol. 57, pp. 320–335, 2013.
- [10] B. Wahlberg, H. Hjalmarsson, and M. Annergren, "On optimal input design in system identification for control," in *49th IEEE Conference on Decision and Control (CDC)*. IEEE, 2010, pp. 5548–5553.
- [11] X. Bombois, G. Scorletti, M. Gevers, R. Hildebrand, and P. Van den Hof, "Cheapest open-loop identification for control," in *IEEE Conference on Decision and Control (CDC)*. IEEE, 2004, pp. 382–387.
- [12] H. Mania, M. I. Jordan, and B. Recht, "Active learning for nonlinear system identification with guarantees," *arXiv preprint arXiv:2006.10277*, 2020.
- [13] K. Chaloner and I. Verdinelli, "Bayesian experimental design: A review," *Statistical Science*, pp. 273–304, 1995.

- [14] A. Foster, M. Jankowiak, E. Bingham, P. Horsfall, Y. W. Teh, T. Rainforth, and N. Goodman, "Variational bayesian optimal experimental design," *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [15] A. Foster, D. R. Ivanova, I. Malik, and T. Rainforth, "Deep adaptive design: Amortizing sequential bayesian experimental design," in *International Conference on Machine Learning*. PMLR, 2021, pp. 3384–3395.
- [16] K. Chaloner, "Optimal Bayesian Experimental Design for Linear Models," *The Annals of Statistics*, vol. 12, no. 1, pp. 283 – 300, 1984.
- [17] A. Tulyan, S. R. Khare, B. Huang, R. B. Gopaluni, and J. F. Forbes, "Bayesian identification of non-linear state-space models: Part i-input design," *Proc. of IFAC*, vol. 46, no. 32, pp. 774–779, 2013.
- [18] X. Huan and Y. M. Marzouk, "Sequential bayesian optimal experimental design via approximate dynamic programming," *arXiv preprint arXiv:1604.08320*, 2016.
- [19] T. Blau, E. V. Bonilla, I. Chades, and A. Dezfouli, "Optimizing sequential experimental design with deep reinforcement learning," in *Proc. of International Conference on Machine Learning (ICML)*, 2022, pp. 2107–2128.
- [20] S. M. Kay, *Fundamentals of statistical signal processing: estimation theory*. Prentice-Hall, Inc., 1993.
- [21] J.-M. Marin, P. Pudlo, C. P. Robert, and R. J. Ryder, "Approximate bayesian computational methods," *Statistics and Computing*, vol. 22, no. 6, pp. 1167–1180, 2012.
- [22] R. Kuenzel, J. Teizer, M. Mueller, and A. Blicke, "Smartsite: Intelligent and autonomous environments, machinery, and processes to realize smart road construction projects," *Automation in Construction*, vol. 71, pp. 21–33, 2016.
- [23] E. Zenkov and V. Sobolev, "Numerical modeling of the dynamics of multi-storey buildings with elastoplastic seismic insulators," in *Journal of Physics: Conference Series*, 2018.
- [24] COMSOL AB, "COMSOL Multiphysics®." [Online]. Available: <https://www.comsol.com>
- [25] G. Hinton, N. Srivastava, and K. Swersky, "Neural networks for machine learning lecture 6a overview of mini-batch gradient descent," *Cited on*, vol. 14, no. 8, p. 2, 2012.
- [26] R. E. Caflisch, "Monte Carlo and quasi-Monte Carlo methods," *Acta numerica*, vol. 7, pp. 1–49, 1998.