

Particle-based Gaussian process optimization for input design in nonlinear dynamical models

Patricio E. Valenzuela, Johan Dahlin, Cristian R. Rojas and Thomas B. Schön

Abstract—We propose a novel approach to input design for identification of nonlinear state space models. The optimal input sequence is obtained by maximizing a scalar cost function of the Fisher information matrix. Since the Fisher information matrix is unavailable in closed form, it is estimated using particle methods. In addition, we make use of Gaussian process optimization to find the optimal input and to mitigate the problem of a large computational cost incurred by the particle method, as the method reduces the number of functional evaluations. Numerical examples are provided to illustrate the performance of the resulting algorithm.

Index Terms—System identification, input design, Gaussian process optimization.

I. INTRODUCTION

Input design concerns the maximization of the information retrieved from an experiment. Some of the first contributions in this area have been introduced in [1], [2]. Since then, several approaches to experiment design have been developed (see e.g. [3] and the references therein).

Recently, the problem of input design for the identification of nonlinear dynamical models has gained interest. One of the main difficulties in this case is that a closed form expression for the Fisher information matrix is typically not available. In addition, the frequency domain techniques employed in the linear case [4] are no longer valid, which implies that other formulations are required. Contributions in this field consider nonlinear FIR models [5], multilevel excitation [6], [7], [8], and nonlinear state space models [9], among others.

As the Fisher information matrix is unavailable in closed form, we need to rely on estimates. However, such estimates are always subject to uncertainty, which results in difficulties when implementing traditional optimization methods.

In this work, we explore the reduction of the computational complexity when calculating the objective function used in input design for identification of nonlinear dynamical models. To this end, a Gaussian process optimization (GPO) based algorithm is presented. By assuming that the scalar function of the Fisher information matrix is a realization from a Gaussian process (GP), we can compute its predictive posterior distribution given a set of samples over the feasible set. The predictive posterior distribution acts as a surrogate

of the intractable objective function, and is employed to compute the next sample over the feasible set by using an acquisition rule. This technique recursively explores the feasible set to determine the element maximizing a surrogate function. The advantage of this approach when compared with existing techniques is that it can handle uncertainty in the estimates of the objective function, and it drives the exploration of the input space towards those regions where an improvement of the objective function is expected.

As with most approaches in experiment design, we rely on prior information about the system for computing an optimal design. This assumption can be overcome by implementing an adaptive scheme [10], or by using a robust input design scheme on top of it [11]. However, this is beyond the scope of this paper.

II. PROBLEM FORMULATION

Consider the discrete time, nonlinear state space model (SSM) defined for all $t \geq 1$ by

$$x_t|x_{t-1} \sim f_\theta(x_t|x_{t-1}, u_{t-1}), \quad (1a)$$

$$y_t|x_t \sim g_\theta(y_t|x_t, u_t), \quad (1b)$$

$$x_0 \sim \mu_\theta(x_0), \quad (1c)$$

where f_θ , g_θ , and μ_θ are known probability density functions (pdf) parameterized by the unknown parameter $\theta \in \Theta \subset \mathbb{R}^{n_\theta}$. Here, $u_t \in \mathcal{C} \subseteq \mathbb{R}^{n_u}$ denotes the input signal, $x_t \in \mathbb{R}^{n_x}$ are the (unobserved/latent) internal states, and $y_t \in \mathbb{R}^{n_y}$ are the measured outputs. We assume that there exists a $\theta_0 \in \Theta$ such that the pdfs in (1) describe the true pdfs of the system when $\theta = \theta_0$, i.e., there is no undermodelling [3].

The objective is to design $u_{1:T} := (u_1, \dots, u_T) \in \mathcal{C}^T$, such that the parameter θ in the model (1) can be identified with maximum accuracy as defined by a scalar function of the Fisher information matrix $\mathcal{I}_F^{\theta_0}$ [2], given by

$$\mathcal{I}_F^{\theta_0}(u_{1:T}) := \mathbf{E} \{ \mathcal{S}(\theta_0) \mathcal{S}^\top(\theta_0) | u_{1:T} \}, \quad (2)$$

with $\mathcal{S}(\theta_0)$ denoting the score function, i.e.,

$$\mathcal{S}(\theta_0) := \nabla \ell_\theta(y_{1:T})|_{\theta=\theta_0}, \quad (3a)$$

$$\ell_\theta(y_{1:T}) := \log p_\theta(y_{1:T} | u_{1:T}). \quad (3b)$$

We note that the expected value in (2) is with respect to the stochastic processes in (1).

In the following, we consider $u_{1:T}$ as a realization of a stationary process. Hence, we will be interested in the *per-sample Fisher information matrix*, given by

$$\mathcal{I}_F^{\theta_0, \text{av}}(u_{1:T}) := \frac{1}{T} \mathbf{E}_u \{ \mathcal{I}_F^{\theta_0}(u_{1:T}) \}. \quad (4)$$

This work was supported by the Swedish Research Council under contracts 621-2013-5524 and 621-2009-4017.

Patricio E. Valenzuela and Cristian R. Rojas are with the Department of Automatic Control and ACCESS Linnaeus Center, School of Electrical Engineering, KTH Royal Institute of Technology, SE-100 44 Stockholm, Sweden (e-mail: {pva, crro}@kth.se).

Johan Dahlin and Thomas B. Schön are with the Department of Information Technology, Uppsala University, SE-751 05 Uppsala, Sweden (e-mail: uu@johandahlin.com, thomas.schon@it.uu.se).

The input $u_{1:T}$ optimizes a scalar function of (4). We define this scalar function as $h: \mathbb{R}^{m \times m} \rightarrow \mathbb{R}$, assumed to be a matrix nondecreasing function [12, p. 108].

The problem presented here can be summarized as
Problem 1: Find an input signal $u_{1:T}^{\text{opt}} \in \mathcal{C}^T$ as

$$u_{1:T}^{\text{opt}} := \arg \max_{u_{1:T} \in \mathcal{C}^T} h(\mathcal{I}_F^{\theta_0, \text{av}}(u_{1:T})), \quad (5)$$

where $h: \mathbb{R}^{m \times m} \rightarrow \mathbb{R}$ is a matrix nondecreasing function, and $\mathcal{I}_F^{\theta_0, \text{av}}(u_{1:T})$ is given in (4). ■

III. GAUSSIAN PROCESS OPTIMIZATION IN INPUT DESIGN

Problem 1 is difficult to solve. One of the main challenges is the characterization of $h(\mathcal{I}_F^{\theta_0, \text{av}}(u_{1:T}))$ for all $u_{1:T} \in \mathcal{C}^T$. Indeed, a closed form expression for $h(\mathcal{I}_F^{\theta_0, \text{av}}(u_{1:T}))$ is only available under specific assumptions on the model structure (1) and the input properties, implying that we need to rely on approximations in the general case. Moreover, even if an estimate of $h(\mathcal{I}_F^{\theta_0, \text{av}}(u_{1:T}))$ is available, the estimate is subject to uncertainty, and this uncertainty cannot be handled by part of the available optimization methods.

Instead, we employ the iterative procedure discussed in [13] to solve Problem 1. The procedure generates a sequence of iterates $\{u_{1:T}^{(k)}\}_{k \geq 0}$ for the input excitation. Each iteration consists of three steps:

- (i) Given $u_{1:T}^{(k)}$, compute an estimate of the objective function $h(\mathcal{I}_F^{\theta_0, \text{av}}(u_{1:T}^{(k)}))$, denoted by \hat{h}_k .
- (ii) Given the collection of tuples $\{u_{1:T}^{(j)}, \hat{h}_j\}_{j=0}^k$, create a model of the (unavailable) objective function $h(\mathcal{I}_F^{\theta_0, \text{av}}(u_{1:T}))$.
- (iii) Use the model as a surrogate for $h(\mathcal{I}_F^{\theta_0, \text{av}}(u_{1:T}))$ to generate a new iterate $u_{1:T}^{(k+1)}$.

The procedure only requires one estimate of $h(\mathcal{I}_F^{\theta_0, \text{av}}(u_{1:T}))$ at each iteration, hence keeping the number of estimates as low as possible. Moreover, it requires fewer iterations than a random search, since it focuses on regions of \mathcal{C}^T where an improvement is expected.

For step (i), we employ particle methods to estimate $h(\mathcal{I}_F^{\theta_0, \text{av}}(u_{1:T}^{(k)}))$. This is discussed in Section III-A.

For steps (ii) and (iii) we use the GPO framework [14], [15]. We first compute a surrogate of the objective function by modelling it as a Gaussian process, and computing the predictive posterior distribution based on $\{u_{1:T}^{(j)}, \hat{h}_j\}_{j=0}^k$. This is discussed in Section III-B.

Then we make use of a heuristic, referred to as the *acquisition rule* (presented in Section III-C), to compute $u_{1:T}^{(k+1)}$ based on the GP model. The acquisition rule favours values of $u_{1:T}$ for which the model predicts a large value of the objective function and/or where there is high uncertainty. This establishes a trade-off between exploration and exploitation of the input set. Finally, to employ the GPO framework in input design, we need tractable parameterizations of \mathcal{C}^T , which are discussed in Subsection III-D.

A. Estimating the Fisher information matrix

Given $u_{1:T}^{(k)} \in \mathcal{C}^T$, we need to approximate (4). To this end, we consider the estimator in [16], which is based on

one estimate of $\mathcal{S}(\theta_0)$ (provided a sufficiently large T) to approximate (4) by [17]

$$\hat{\mathcal{I}}_F^{\theta_0, \text{av}} := \frac{1}{T} \left[\sum_{t=1}^T \hat{\mathcal{S}}_t(\theta_0) (\hat{\mathcal{S}}_t(\theta_0))^\top - \frac{1}{T} \hat{\mathcal{S}}(\theta_0) (\hat{\mathcal{S}}(\theta_0))^\top \right], \quad (6)$$

where the Fisher identity [18] can be used to write¹

$$\mathcal{S}(\theta') = \sum_{t=1}^T \mathcal{S}_t(\theta'), \quad (7)$$

$$\mathcal{S}_t(\theta') := \int \nabla \xi_\theta(x_{t-1:t})|_{\theta=\theta'} p_{\theta'}(x_{t-1:t} | \mathbf{y}, \mathbf{u}) dx_{t-1:t},$$

$$\xi_\theta(x_{t-1:t}) := \log f_\theta(x_t | x_{t-1}, u_{t-1}) + \log g_\theta(y_t | x_t, u_t),$$

where $x_{t-1:t} := \{x_{t-1}, x_t\}$. From (6), we note that an estimate for (7) is required, which can be obtained from particle methods [19].

To estimate the score function in (7), we require the two-step smoothing distribution $p_\theta(x_{t-1:t} | \mathbf{y}, \mathbf{u})$, which is not available analytically for a general SSM. Instead, we approximate it using an empirical distribution

$$\hat{p}_\theta(dx_{t-1:t} | \mathbf{y}, \mathbf{u}) := \sum_{i=1}^N w_t^{(i)} \delta_{x_{t-1:t}^{(i)}}(dx_{t-1:t}), \quad (8)$$

where $x_t^{(i)}$ and $w_t^{(i)}$ denote particle i and its normalized weight at time t . Here, $\{x_t^{(i)}, w_t^{(i)}\}_{t=1}^T$ denotes the *particle system* generated by a particle filter and $\delta_{x'}$ denotes the Dirac measure located at $x = x'$.

Following [16], here we use the bootstrap particle filter (bPF), see Algorithm 1 [21]. However, the estimator (8) based only on the bPF often suffers from poor accuracy due to particle degeneracy, see e.g. [19]. To mitigate this problem, we use a particle smoother that introduces a backward sweep after the forward run of the bPF. Here, we use the forward-filtering backwards simulator (FFBSi) with rejection sampling and early stopping [20].

Algorithm 2 presents the pseudo-code for the FFBSi. Here, $\text{Multi}(\{p^{(i)}\}_{i=1}^N)$ and $\text{Uniform}([a, b])$ denote the multinomial distribution over N elements, with $p^{(i)}$ being the probability of choosing the i -th element, and the uniform distribution with support $[a, b]$, respectively. We note that the parameter ρ required by Algorithm 2 is chosen such that $f_\theta(x_t | x_{t-1}, u_{t-1}) \leq \rho$ for all $t \in \{1, \dots, T\}$. The computational complexity of FFBSi is of order $\mathcal{O}(NMT)$, where N and M denote the number of filter and smoother particles, respectively. We refer to [20] for more details on the effects of N , M and T in the accuracy of the estimator and to [19] for more details on Algorithms 1-2.

B. Modelling the objective function

We explore the use of a GP to model the objective function $h(\mathcal{I}_F^{\theta_0, \text{av}}(u_{1:T}))$ [22]. GPs can be understood as a generalization of the multivariate Gaussian distribution and are commonly used as priors over functions [23]. In this perspective, the posterior obtained by conditioning on the

¹For conciseness, we write $\mathbf{v} := v_{1:T}$ for any vector $v_{1:T}$. In addition, we remove the dependence on k of the input, state, and measurements.

Algorithm 1 Bootstrap particle filter (bPF)

 INPUTS: An SSM (1), \mathbf{y} (observations), \mathbf{u} (inputs), $N \in \mathbb{N}$ (no. particles).

 OUTPUT: $\{x_t^{(i)}, w_t^{(i)}\}_{i=1}^N$, $t = 1, \dots, T$.

```

1: Sample  $x_0^{(i)} \sim \mu_\theta(x_0)$  and set  $w_0^{(i)} = 1/N$ .
2: for  $t = 1$  to  $T$  do
3:   for  $i, j = 1$  to  $N$  do
4:     (Resampling) Sample  $a_t^{(i)}$  from a multinomial distribution with  $\mathbf{P}(a_t^{(i)} = j) = w_{t-1}^{(j)}$ .
5:     (Propagation) Sample  $x_t^{(i)} \sim f_\theta(x_t^{(i)} | x_{t-1}^{(i)}, u_t)$ .
6:     Set  $x_{0:t}^{(i)} = \{x_{0:t-1}^{(i)}, x_t^{(i)}\}$ .
7:     (Weighting) Calculate  $\tilde{w}_t^{(i)} = g_\theta(y_t | x_t^{(i)}, u_t)$ .
8:     Normalize  $\tilde{w}_t^{(i)}$  (over  $i$ ) to obtain  $w_t^{(i)}$ .
9:   end for
10: end for
  
```

Algorithm 2 Fast forward-filtering backward-simulator with early stopping (FFBSi-ES)

 INPUTS: Inputs to Algorithm 1, $M \in \mathbb{N}$ (no. backward trajectories), $N_{\text{limit}} \in \mathbb{N}$ (limit for when to stop using rejection sampling), $\rho > 0$.
 OUTPUT: $\hat{\mathcal{I}}_F^{\theta_0, \text{av}}(\mathbf{u})$ (estimate of the Fisher information matrix).

```

1: Run Algorithm 1 to obtain the particle system  $\{x_t^{(i)}, w_t^{(i)}\}_{i=1}^N$  for  $t = 1, \dots, T$ .
2: Sample  $\{b_T(j)\}_{j=1}^M \sim \text{Multi}(\{w_T^{(i)}\}_{i=1}^N)$ .
3: Set  $\tilde{x}_T^{(j)} = x_T^{b_T(j)}$  for  $j = 1, \dots, M$ .
4: for  $t = T - 1$  to 1 do
5:    $L \leftarrow 1, \dots, M$ .
6:   {Rejection sampling until  $N_{\text{limit}}$  trajectories remain.}
7:   while  $|L| \geq N_{\text{limit}}$  do
8:      $n \leftarrow \text{Multi}(\{1/|L|\}_{i=1}^{|L|})$ .
9:      $\delta \leftarrow \emptyset$ .
10:    Sample  $\{I(k)\}_{k=1}^n \sim \text{Multi}(\{w_t^{(i)}\}_{i=1}^N)$ .
11:    Sample  $\{U(k)\}_{k=1}^n \sim \text{Uniform}([0, 1])$ .
12:    for  $k = 1$  to  $n$  do
13:      if  $U(k) \leq f(\tilde{x}_{t+1}^{I(k)} | x_t^{I(k)})/\rho$  then
14:         $b_t(L(k)) \leftarrow I(k)$ .
15:         $\delta \leftarrow \delta \cup \{L(k)\}$ .
16:      end if
17:    end for
18:     $L \leftarrow L \setminus \delta$ .
19:  end while
20:  {Use standard FFBSi for the remaining trajectories [20].}
21:  for  $j \in L$  do
22:    Compute  $\tilde{w}_{t|T}^{(i,j)} \propto w_t^{(i)} f(\tilde{x}_{t+1}^{(j)} | x_t^{(i)})$  for  $i = 1, \dots, N$ .
23:    Normalize the smoothing weights  $\{\tilde{w}_{t|T}^{(i,j)}\}_{i=1}^N$ .
24:    Draw  $b_t(j) \sim \text{Multi}(\{\tilde{w}_{t|T}^{(i,j)}\}_{i=1}^N)$ .
25:  end for
26:  Set  $\tilde{x}_{t:T}^{(j)} = \{x_t^{b_t(j)}, \tilde{x}_{t+1:T}^{(j)}\}$  for  $j = 1, \dots, M$ .
27:  Calculate
      
$$\hat{\mathcal{S}}_t^{(k)}(\theta) = \frac{1}{M} \sum_{j=1}^M \nabla \xi_\theta(\tilde{x}_{t:t+1}^{(j)}).$$

28: end for
29: Compute  $\hat{\mathcal{I}}_F^{\theta_0, \text{av}}(u_{1:T})$  using (6).
  
```

observations corresponds to the functions that could have generated the observations.

In the following, we model the function $h(\mathcal{I}_F^{\theta_0, \text{av}}(\cdot))$ as




Fig. 1. Left: Histogram of ν and plot of the scaled pdf of an $\mathcal{N}(0, 1^2)$ distribution (continuous line), Example 1. Right: Quantile-quantile plot of the samples of ν and the $\mathcal{N}(0, 1^2)$ distribution, Example 1.

being a priori distributed according to a GP. That is

$$h(\mathcal{I}_F^{\theta_0, \text{av}}(\cdot)) \sim \mathcal{GP}(m(\cdot), \kappa(\cdot, \cdot)), \quad (9)$$

where the process is fully described by the mean function $m(\cdot)$ and the covariance function $\kappa(\cdot, \cdot)$. Examples of these functions are a constant for m and a Matérn $s/2$ function for κ [22, p.84].

To simplify the discussion, we will focus on a specific iteration k of the proposed procedure. Let $\mathcal{D}_k := \{\mathbf{u}_{1:T}^{(k)}, \hat{\mathbf{h}}_k\}$ denote a set of iterates, where $\mathbf{u}_{1:T}^{(k)}$ and $\hat{\mathbf{h}}_k$ denote matrices obtained by stacking input realizations and estimates of the objective function up to iteration k , respectively. In addition, we will assume that

$$\hat{\mathbf{h}}_k = h(\mathcal{I}_F^{\theta_0, \text{av}}(u_{1:T}^{(k)})) + z, \quad (10)$$

where $z \sim \mathcal{N}(0, \sigma_z^2)$, and $\sigma_z > 0$. We note that σ_z is unknown a priori, and it needs to be estimated using \mathcal{D}_k . The assumption (10) seems strict, but the continuous mapping theorem [24, Theorem 2.7] shows that the central limit theorem also applies to the estimate $\hat{\mathbf{h}}_k$, as it is satisfied by (3b) asymptotically in the number of particles.

Example 1: Consider

$$x_{t+1} | x_t \sim \mathcal{N}(\phi x_t + u_t, 0.1^2), \quad (11a)$$

$$y_t | x_t \sim \mathcal{N}(\alpha x_t, 0.1^2), \quad (11b)$$

where the parameters are $\theta = \{\phi, \alpha\}$. We generate $T = 10^3$ observations from (11) with $\theta_0 = \{0.8, 1\}$.

We are interested in estimating $h(\mathcal{I}_F^{\theta_0, \text{av}}(u_{1:T})) = \log \det(\mathcal{I}_F^{\theta_0, \text{av}}(u_{1:T}))$, where $u_{1:T}$ is a binary white noise process with values $\{-1, 1\}$.

The estimate of the Fisher information matrix is obtained using Algorithms 1-2, with $N = 2.5 \cdot 10^3$ particles, $M = 100$ backward trajectories and $N_{\text{limit}} = \lfloor \sqrt{N} \rfloor$ in the FFBSi smoother. Figure 1 shows the histogram based on 10^3 realizations of the random variable

$$\nu := \frac{\sqrt{M}(\hat{h} - \bar{h})}{\sigma \sqrt{M\hat{h}}}, \quad (12)$$

where $\hat{h} := h(\hat{\mathcal{I}}_F^{\theta_0, \text{av}}(u_{1:T}))$, and \bar{h} , $\sigma^2_{\sqrt{M\hat{h}}}$ are the sample mean of \hat{h} and variance of $\sqrt{M}\hat{h}$, respectively. As a comparison, we also present the scaled pdf of an $\mathcal{N}(0, 1^2)$ distribution. We can see that the histogram follows the shape of the pdf of a $\mathcal{N}(0, 1^2)$ distribution. This is also confirmed by the quantile-quantile (QQ) plot in Figure 1, where the

quantiles of ν coincides with those given by an $\mathcal{N}(0, 1^2)$ distribution. ■

Based on (10), it follows that the predictive posterior distribution is

$$h(\mathcal{I}_F^{\theta_0, \text{av}}(u_{1:T}) | \mathcal{D}_k) \sim \mathcal{N}(\mu(u_{1:T} | \mathcal{D}_k), \sigma^2(u_{1:T} | \mathcal{D}_k) + \sigma_z^2), \quad (13)$$

where $\mu(u_{1:T} | \mathcal{D}_k)$ and $\sigma^2(u_{1:T} | \mathcal{D}_k)$ denote the posterior mean and variance given \mathcal{D}_k . From standard results for the Gaussian distribution, we have

$$\begin{aligned} \mu(u_{1:T} | \mathcal{D}_k) &= m(u_{1:T}) \\ &\quad + \kappa(u_{1:T}, \mathbf{u}_{1:T}^{(k)}) \Gamma^{-1} \left\{ \widehat{\mathbf{h}}_k - m(u_{1:T}) \right\}, \end{aligned} \quad (14a)$$

$$\begin{aligned} \sigma^2(u_{1:T} | \mathcal{D}_k) &= \kappa(u_{1:T}, u_{1:T}) \\ &\quad - \kappa(u_{1:T}, \mathbf{u}_{1:T}^{(k)}) \Gamma^{-1} \kappa(\mathbf{u}_{1:T}^{(k)}, u_{1:T}), \end{aligned} \quad (14b)$$

with $\Gamma := \kappa(\mathbf{u}_{1:T}^{(k)}, \mathbf{u}_{1:T}^{(k)}) + \sigma_z^2 \mathbf{I}_k$, where \mathbf{I}_k denotes the $k \times k$ -identity matrix.

In the GP model introduced here, we use mean and covariance functions that possibly depend on some unknown hyperparameters. In addition, we also need to estimate σ_z characterizing the random variable z in (10). To estimate these quantities, we adopt the empirical Bayes procedure, where the marginal likelihood of the data is numerically optimized with respect to the hyperparameters [25].

C. Acquisition rules

To implement step (iii), we need to generate $u_{1:T}^{(k+1)} \in \mathcal{C}^T$. One option is to perform a random walk over \mathcal{C}^T , which works well provided that the parameterization of $u_{1:T}$ is of small dimension. However, this approach is inefficient as the dimension of the parameterization for $u_{1:T}$ increases.

Instead, we make use of acquisition rules that balance exploration and exploitation of the parameter space and employ the posterior distribution obtained from the GP. Here, we use the expected improvement (EI) technique [26].

Consider the *predicted improvement*

$$I(u_{1:T}) := \max \left\{ 0, h(\mathcal{I}_F^{\theta_0, \text{av}}(u_{1:T})) - \mu_{\max} - \xi \right\}, \quad (15)$$

where $\xi \in \mathbb{R}$ is a user defined coefficient balancing exploration and exploitation², and

$$\mu_{\max} := \max_{u_{1:T} \in \mathbf{u}_{1:T}^{(k)}} \mu(u_{1:T} | \mathcal{D}_k), \quad (16)$$

the expected peak of $h(\mathcal{I}_F^{\theta_0, \text{av}}(u_{1:T}))$ at iteration k .

By using the posterior distribution obtained from the GP, we define the EI as³

$$\mathbf{E} \{ I(u_{1:T}) \} = \sigma(u_{1:T}) \left\{ Z(u_{1:T}) \Phi(Z(u_{1:T})) - \phi(Z(u_{1:T})) \right\}, \quad (17a)$$

$$Z(u_{1:T}) := \sigma^{-1}(u_{1:T}) \left\{ \mu(u_{1:T}) - \mu_{\max} - \xi \right\}, \quad (17b)$$

² $|\xi|$ should be a fraction of $|\mu_{\max}|$ to promote exploration ($\xi < 0$) or exploitation ($\xi > 0$) in the method.

³For simplicity, the dependence on \mathcal{D}_k is dropped from the notation.

with Φ and ϕ denoting the cumulative distribution function and the pdf of the standard Gaussian distribution, respectively. Then, an acquisition rule is

$$u_{1:T}^{(k+1)} = \arg \max_{u_{1:T} \in \mathcal{C}^T} \mathbf{E} \{ I(u_{1:T}) | \mathcal{D}_k \} + e_{1:T}, \quad (18)$$

where $e_{1:T}$ is a random vector⁴ with a user defined distribution such that $\mathbf{E}\{e_{1:T}\} = 0$. Hence, the input in the next iteration is a realization of a random variable promoting exploration around the element maximizing the EI. From (17) we see that the EI assigns a large value when both the variance $\sigma(u_{1:T})$ and the mean difference $\mu(u_{1:T}) - \mu_{\max}$ are large, in line with the desired behavior of an acquisition function, as it is explained at the beginning of Section III.

D. Parameterizing the input

To implement the GPO for solving the input design problem, we need a parameterization of \mathcal{C}^T . Here we briefly explain two options:

1) *Stationary Markov processes*: If we restrict \mathcal{C} to be finite and $u_{1:T}$ to be a realization from an n -dimensional stationary Markov process of a given order, then the parameterization employed in [8] can be used. This parametrization is of interest when amplitude constraints on the input must be satisfied during the experiment.

The parameterization of the input is given by the stationary distribution of the Markov process, which is constrained to

$$\begin{aligned} \mathcal{P}_{\mathcal{C}} := \left\{ p_u : \mathcal{C}^n \rightarrow \mathbb{R} \mid p_u(\mathbf{x}) \geq 0, \forall \mathbf{x} \in \mathcal{C}^n; \right. \\ \left. \sum_{\mathbf{x} \in \mathcal{C}^n} p_u(\mathbf{x}) = 1; \right. \\ \left. \sum_{v \in \mathcal{C}} p_u(v, \mathbf{z}) = \sum_{v \in \mathcal{C}} p_u(\mathbf{z}, v), \forall \mathbf{z} \in \mathcal{C}^{n-1} \right\}. \end{aligned} \quad (19)$$

Following [8], we parameterize (19) as the convex hull of its extreme points, which are computed using graph theoretical techniques. Therefore, the decision variable in this case corresponds to the weighting vector of the extreme points describing an element in $\mathcal{P}_{\mathcal{C}}$. Assuming that $\mathcal{P}_{\mathcal{C}}$ has $n_{\mathcal{V}}$ extreme points, then the weighting vector $\alpha := [\alpha_1 \ \dots \ \alpha_{n_{\mathcal{V}}}]^T \in \mathbb{R}^{n_{\mathcal{V}}}$ is used to compute $p \in \mathcal{P}_{\mathcal{C}}$ as

$$p = \sum_{i=1}^{n_{\mathcal{V}}} \alpha_i p^{(i)}, \quad (20)$$

with α satisfying

$$\alpha_i \geq 0, \text{ for all } i \in \{1, \dots, n_{\mathcal{V}}\}, \quad (21a)$$

$$\sum_{i=1}^{n_{\mathcal{V}}} \alpha_i = 1. \quad (21b)$$

In (20), $\{p^{(i)}\}_{i=1}^{n_{\mathcal{V}}}$ corresponds to the probability mass functions (pmf) that are the extreme points of $\mathcal{P}_{\mathcal{C}}$.

⁴The implementation of the random vector is over the parameter space characterizing \mathcal{C}^T .

Algorithm 3 GPO for input design

 INPUTS: Algorithm 2, K (no. iterations) and $u_{1:T}^{(0)} \in \mathcal{C}^T$ (initial excitation).

 OUTPUT: $\{x_t^{(i)}, w_t^{(i)}\}_{i=1}^N, t = 1, \dots, T$.

- 1: Sample $u_{1:T}^{(0)} \in \mathcal{C}^T$.
 - 2: **for** $k = 0$ to K **do**
 - 3: Use Algorithm 2 to compute $\hat{h}_k := h(\hat{\mathcal{I}}_F^{\theta_0, \text{av}}(u_{1:T}^{(k)}))$.
 - 4: Compute (13)-(14) to obtain $h(\mathcal{I}_F^{\theta_0, \text{av}}(u_{1:T}))|_{\mathcal{D}_k}$.
 - 5: Compute (16) to obtain $\mu_{\max}^{(k+1)}$.
 - 6: Compute (18) to obtain $u_{1:T}^{(k+1)}$.
 - 7: **end for**
 - 8: Compute the maximizer of $\mu(u_{1:T}|\mathcal{D}_K)$ to obtain $u_{1:T}^{\text{opt}}$.
-

Once a new sample $\alpha \in \mathbb{R}^{n_\nu}$ satisfying⁵ (21) is generated, we compute the associated pmf $p \in \mathcal{P}_\mathcal{C}$ by (20), and we generate $u_{1:T}$ by running a Markov chain with stationary distribution p .

2) *Stationary AR processes*: If variance constraints are imposed on the input, we can restrict $u_{1:T}$ to be a filtered white noise process, as it is proposed in [27]. In this case, the decision variables are the filter coefficients, and the properties of the white noise. For example, we can assume that $u_{1:T}$ is a realization from a stationary AR process

$$A(q)u_t = e_t, \quad (22)$$

where $\{e_t\}$ is Gaussian white noise, with variance σ_e^2 , and

$$A(q) := \sum_{i=0}^{n_a} a_i q^{-i}, \quad (23)$$

with $n_a > 0$ given, $a_i \in \mathbb{R}$ for all $i \in \{1, \dots, n_a\}$, and $a_0 = 1$. For this example, the decision variables are $\sigma_e > 0$, and $\{a_i\}_{i=1}^{n_a}$, such that $A(q)$ has all its zeros strictly inside the complex unit disc⁶.

E. The final procedure

Algorithm 3 presents the resulting procedure for input design using Gaussian process optimization. We note that only one functional evaluation is required per iteration, reducing the computational effort when optimizing over \mathcal{C}^T .

IV. NUMERICAL EXAMPLES

Example 2: Consider the linear Gaussian SSM in Example 1. We are interested in maximizing $h(\mathcal{I}_F^{\theta_0, \text{av}}(u_{1:T})) = \log \det(\mathcal{I}_F^{\theta_0, \text{av}}(u_{1:T}))$, where $u_{1:T}$ ($T = 10^3$) is a realization of a stationary Markov process (see Section III-D), with $n_m = 1$ and $\mathcal{C} = \{-1, 1\}$.

For Algorithm 3, we use $K = 500$, $\xi = 0.01$, and $u_{1:T}^{(k+1)}$ given as a random walk centered around the parametrization of the argument maximizing the EI at iteration k , uniformly distributed on $[-0.01, 0.01]$. The estimate of the Fisher information matrix is obtained using Algorithms 1-2, which are implemented as in Example 1. For the prior distribution of

⁵This can be achieved by sampling α satisfying (21a), and then normalizing the entries of α to satisfy (21b).

⁶This can be guaranteed by factorizing $A(q)$ into first and second order polynomials in q , and imposing the constraint on each of these factors.

$h(\mathcal{I}_F^{\theta_0, \text{av}}(u_{1:T}))$, we consider a constant mean function, and a covariance function composed of a Matérn $s/2$ structure and a constant, where the value of s and the constant are estimated using empirical Bayes. The Matérn $s/2$ structure is chosen in this example as it contains information about the smoothness of $h(\mathcal{I}_F^{\theta_0, \text{av}}(u_{1:T}))$. Other choices for the covariance function are possible and we refer to [22, Chapter 4] for more details.

Algorithm 3 is implemented in Matlab using the `fmincon` command for (18) and the GPML toolbox [28] to infer the hyperparameters and estimate the predictive posterior distribution of $h(\mathcal{I}_F^{\theta_0, \text{av}}(u_{1:T}))$.

The solution given by Algorithm 3 is $u_t = 1$ for all $t \geq 0$. In this example, a nonzero constant input introduces a nonzero offset in the measurements, which helps to estimate θ in the presence of process disturbance and measurement noise. As a reference, we draw $u_{1:T}$ as a realization from a binary white noise process with values $\{-1, 1\}$. The results are $h(\mathcal{I}_F^{\theta_0, \text{av}}(u_{1:T}^{\text{opt}})) = 14.57$ for the optimal input and $h(\mathcal{I}_F^{\theta_0, \text{av}}(u_{1:T})) = 10.18$ for the binary white noise. ■

Example 3: Consider the system

$$x_{t+1}|x_t \sim \mathcal{N}\left(\frac{1}{\gamma + x_t^2} + u_t, 0.1^2\right), \quad (24a)$$

$$y_t|x_t \sim \mathcal{N}\left(\beta x_t^2, 1^2\right), \quad (24b)$$

where the parameters are $\theta = \{\gamma, \beta\}$. We generate $T = 10^3$ observations from the model with $\theta_0 = \{2, 0.8\}$. We note that estimating γ in (24) is inherently difficult, since two different values of x_t can explain y_t equally well.

We consider the same setting and function h as in Example 2, but we consider three cases for \mathcal{C} :

- Case 1: $\mathcal{C} = \{-1, 1\}$.
- Case 2: $\mathcal{C} = \{-1, 0, 1\}$.
- Case 3: $\mathcal{C} = \{-1, -1/3, 1/3, 1\}$.

Table I presents the value of $h^{\text{opt}} := h(\hat{\mathcal{I}}_F^{\theta_0, \text{av}}(u_{1:T}^{\text{opt}}))$ for each case, where $u_{1:T}^{\text{opt}}$ corresponds to the optimal input obtained from Algorithm 3. As comparison, we also present $h(\hat{\mathcal{I}}_F^{\theta_0, \text{av}}(u_{1:T}))$ when $\{u_t\}$ is a first order stationary AR process given by (22)-(23), where $a_1 = -0.9$ and $\sigma_e^2 = 1 - a_1^2$, and a binary distributed white noise with values $\{-1, 1\}$ (AR Gaussian and Binary in Table I, respectively). We see that the stationary AR process gives less information than the binary white noise process, and that the binary white noise seems to be optimal when $\mathcal{C} = \{-1, 1\}$, as it is confirmed by the value of h^{opt} for Case 1. We also note that adding intermediate values to the input alphabet increases the amount of information in the data, as h^{opt} is greater in Cases 2 and 3 than in Case 1.




Fig. 2. Optimal input $u_{1:T}^{\text{opt}}$ for Case 3 in Example 3.




Fig. 3. Value of \hat{h}_k and μ_{\max} at iteration k for Case 3 in Example 3.

TABLE I
 h^{opt} FOR DIFFERENT INPUT REALIZATIONS, EXAMPLE 3.

Input h^{opt}	AR Gaussian	Binary	Case 1	Case 2	Case 3
	2.07	4.11	4.11	4.15	4.44

Figure 2 presents the optimal input obtained for Case 3. We note that the optimal input includes a nonzero offset to improve the accuracy of the parameter estimates.

To illustrate the evolution of \hat{h}_k , Figure 3 presents the samples $\{\hat{h}_k\}_{k=1}^{100}$, together with the value of μ_{\max} at every iteration. The first 20 samples are drawn at random from \mathcal{C}^T to provide an initial estimate of the hyperparameters in the GP prior. We note that some of the samples in $\{\hat{h}_k\}_{k=1}^{20}$ are not close to the optimal cost, which is expected due to random sampling. However, once Algorithm 3 is executed from iteration 21 onwards, we observe that the samples are close to μ_{\max} , which implies that the space \mathcal{C}^T is explored only in those regions where h can only increase with respect to the current estimates. Hence, the proposed technique drives the parameter search towards those regions where an improvement in the objective function is expected. ■

V. CONCLUSIONS

A Gaussian process optimization algorithm for input design for the identification of nonlinear dynamical models has been introduced. The method maximizes a scalar cost function of the Fisher information matrix over the parameter set for the input sequence. Since the objective function is unavailable in closed form, a Gaussian process approach is employed to compute a surrogate function. Numerical examples show that the algorithm can provide a good alternative to solve the input design problem.

Future work on this subject will consider a better estimator of the Fisher information matrix with a better particle smoother, alternative parameterizations of $\{u_t\}$ and convergence analysis of the method.

REFERENCES

- [1] D.R. Cox, *Planning of experiments*, New York: Wiley, 1958.
- [2] G.C. Goodwin and R.L. Payne, *Dynamic System Identification: Experiment Design and Data Analysis*, Academic Press, New York, 1977.
- [3] L. Ljung, *System Identification. Theory for the User, 2nd ed.*, Upper Saddle River, NJ: Prentice-Hall, 1999.
- [4] H. Jansson and H. Hjalmarsson, "Input design via LMIs admitting frequency-wise model specifications in confidence regions," *IEEE Transactions on Automatic Control*, vol. 50, no. 10, pp. 1534–1549, 2005.
- [5] C. Larsson, H. Hjalmarsson, and C.R. Rojas, "On optimal input design for nonlinear FIR-type systems," in *Proceedings of the 49th IEEE Conference on Decision and Control*, Atlanta, USA, December 2010, pp. 7220–7225.
- [6] M. Forgiione, X. Bombois, P.M.J. Van den Hof, and H. Hjalmarsson, "Experiment design for parameter estimation in nonlinear systems based on multilevel excitation," in *Proceedings of the 13th European Control Conference*, Strasbourg, France, June 2014.
- [7] A. De Cock, M. Gevers, and J. Schoukens, "A preliminary study on optimal input design for nonlinear systems," in *Proceedings of the IEEE Conference on Decision and Control*, Florence, Italy, December 2013, pp. 4931–4936.
- [8] P.E. Valenzuela, C.R. Rojas, and H. Hjalmarsson, "A graph theoretical approach to input design for identification of nonlinear dynamical models," *Automatica*, vol. 51, pp. 233–242, 2015.
- [9] P.E. Valenzuela, J. Dahlin, C.R. Rojas, and T.B. Schön, "A graph/particle-based method for experiment design in nonlinear systems," in *Proceedings of the 19th IFAC World Congress*, Cape Town, South Africa, August 2014.
- [10] L. Gerencsér, H. Hjalmarsson, and J. Mårtensson, "Identification of ARX systems with non-stationary inputs—asymptotic analysis with application to adaptive input design," *Automatica*, vol. 45, no. 3, pp. 623–633, 2009.
- [11] C.R. Rojas, J.S. Welsh, G.C. Goodwin, and A. Feuer, "Robust optimal experiment design for system identification," *Automatica*, vol. 43, no. 6, pp. 993–1008, June 2007.
- [12] S. Boyd and L. Vandenberghe, *Convex Optimization*, Cambridge University Press, 2004.
- [13] J. Dahlin and F. Lindsten, "Particle filter-based Gaussian process optimisation for parameter inference," in *Proceedings of the 19th IFAC World Congress*, Cape Town, South Africa, August 2014.
- [14] B. Shahriari, K. Swersky, Z. Wang, R.P. Adams, and N. de Freitas, "Taking the human out of the loop: A review of Bayesian optimization," *Proceedings of the IEEE*, vol. 104, no. 1, pp. 148–175, 2016.
- [15] M.A. Osborne, R. Garnett, and S.J. Roberts, "Gaussian processes for global optimization," in *3rd International Conference on Learning and Intelligent Optimization*, 2009, pp. 1–15.
- [16] P.E. Valenzuela, J. Dahlin, C.R. Rojas, and T.B. Schön, "On robust input design for nonlinear dynamical models," *Automatica*, 2016, (Provisionally accepted).
- [17] M. Segal and E. Weinstein, "A new method for evaluating the log-likelihood gradient, the Hessian, and the Fisher information matrix for linear dynamic systems," *IEEE Transactions on Information Theory*, vol. 35, no. 3, pp. 682–687, 1989.
- [18] O. Cappé, E. Moulines, and T. Rydén, *Inference in Hidden Markov Models*, Springer, 2005.
- [19] F. Lindsten and T.B. Schön, "Backward simulation methods for Monte Carlo statistical inference," *Foundations and Trends in Machine Learning*, vol. 6, no. 1, pp. 1–143, 2013.
- [20] R. Douc, A. Garivier, E. Moulines, and J. Olsson, "Sequential Monte Carlo smoothing for general state space hidden Markov models," *Annals of Applied Probability*, vol. 21, no. 6, pp. 2109–2145, 2011.
- [21] A. Doucet and A. Johansen, "A tutorial on particle filtering and smoothing: Fifteen years later," in *The Oxford Handbook of Nonlinear Filtering*, D. Crisan and B. Rozovsky, Eds. Oxford University Press, 2011.
- [22] C.E. Rasmussen and C.K.I. Williams, *Gaussian processes for Machine Learning*, MIT press, 2006.
- [23] P. Boyle, *Gaussian processes for regression and optimisation*, Ph.D. thesis, Victoria University Wellington, Wellington, New Zealand, 2007.
- [24] P. Billingsley, *Convergence of probability measures, 2nd ed.*, John Wiley & Sons, 1999.
- [25] B.P. Carlin and T.A. Louis, *Bayes and empirical Bayes methods for data analysis*, London: Chapman and Hall, 1996.
- [26] D.R. Jones, "A taxonomy of global optimization methods based on response surfaces," *Journal of Global Optimization*, vol. 21, no. 4, pp. 345–383, 2001.
- [27] R.B. Gopaluni, T.B. Schön, and A.G. Wills, "Input design for nonlinear stochastic dynamic systems - A particle filter approach," in *Proceedings of the 18th IFAC World Congress*, Milano, Italy, August 2011.
- [28] C.E. Rasmussen and H. Nickish, *Gaussian process regression and classification toolbox, version 3.6*, July 2015.