

Trajectory-oriented optimization of stochastic epidemiological models

Arindam Fadikar

Decision and Infrastructure Sciences
Argonne National Laboratory

Mickaël Binois

Acumes project-team
Inria centre at Université Côte d’Azur

Nicholson Collier

Decision and Infrastructure Sciences
Argonne National Laboratory

Abby Stevens

Decision and Infrastructure Sciences
Argonne National Laboratory

Kok Ben Toh

Department of Preventive Medicine
Institute of Global Health
Northwestern University

Jonathan Ozik

Decision and Infrastructure Sciences
Argonne National Laboratory

ABSTRACT

Epidemiological models must be calibrated to ground truth for downstream tasks such as producing forward projections or running what-if scenarios. The meaning of calibration changes in case of a stochastic model since output from such a model is generally described via an ensemble or a distribution. Each member of the ensemble is usually mapped to a random number seed (explicitly or implicitly). With the goal of finding not only the input parameter settings but also the random seeds that are consistent with the ground truth, we propose a class of Gaussian process (GP) surrogates along with an optimization strategy based on Thompson sampling. This *Trajectory Oriented Optimization* (TOO) approach produces actual trajectories close to the empirical observations instead of a set of parameter settings where only the mean simulation behavior matches with the ground truth.

1 INTRODUCTION

Stochastic epidemiological models have been shown to be a critical tool to support public health decision-making in times of crisis (see, e.g., [26]). These models are parameterized by calibrating their output trajectories against observed data, a crucial step for their use in policy-making. Calibration, however, can be computationally infeasible, particularly for complicated simulators and large populations. Rather than requiring simulations to be re-run with all possible parameter combinations, many have proposed using surrogate models for calibration, which replaces the expensive simulation with a simpler model such as a Gaussian process (GP) [15, 7, 17, 20].

A typical GP surrogate models the *mean* behavior of a stochastic simulator. Given some initialization, repeated runs of the stochastic simulator will result in an ensemble of random realizations of output trajec-

tories from a distribution (as opposed to a fixed response in the case of deterministic models). In the case of a noisy epidemic model, the stochasticity in the output can safely be represented by a mean and variance (assuming white noise) and a GP surrogate can be produced to predict the mean behavior of this process for each simulation input that, once calibrated, is consistent with the ground truth. In other modeling settings, however, each replicate may represent something more structured than white noise. For example, in the case of an agent-based epidemiological model, each replicate might define a different population mixing, and averaging outputs from all replicates may not correspond to a feasible mixing pattern.

In such cases, the random seed governing the stochasticity of each replicate can be taken as an additional input to the model, and the calibration procedure is searching not for the input parameters that most closely match the empirical data on average, but rather for the best *trajectory* (parameter-seed pair). This trajectory-oriented optimization (TOO) approach seeks to capture specific mixing patterns in stochastic models, enabling more empirically informed scenario modeling and forward projections. In this paper, we propose a class of trajectory-oriented GP surrogates with an associated optimization strategy. We leverage the common random number framework detailed by, e.g., [8] and account for non-stationarities with a local formulation of the model [11]. We develop a flexible Bayesian optimization technique tailored for this specific case using a Thompson sampling approach that can be generalized to multiple targets. The main contribution of this paper is the combined surrogate model and optimization strategy framework for trajectory like response. We apply our method to an Susceptible-Exposed-Infectious-Recovered (SEIR) epidemiological model of COVID-19 infections in Chicago, IL, U.S.A., and demonstrate its benefits over standard GP surrogates.

2 METHOD

Let us denote the stochastic computer simulation as $f : \mathbb{R}^d \rightarrow \mathbb{R}$ that maps a d dimensional input $\mathbf{x} \in \mathcal{X}^d \subseteq \mathbb{R}^d$ to a vector output. For stochastic models, f is not one-to-one and the output $f(\mathbf{x})$ can be indexed via replicate indices if repeated runs are available at each input. For the purpose of this paper and the trajectory-oriented approach, we augment the d -dimensional input $\mathbf{x}$ with a trajectory identifier r that can also be associated with the random seed used in the simulation. Later, this additional input dimension is used as a categorical input and is exploited in our proposed GP surrogate. In summary, an output from the stochastic computer simulation f would be characterized by the tuple $(\mathbf{x}, r)$. At this level of specification, f is treated as a deterministic mapping between the tuple $(\mathbf{x}, r)$ and its output. We also slightly abuse notation in adding the time variable t in $\mathbf{x}$ to denote time-indexed scalar output from f .

A standard GP [30, 15] prior characterized by a mean function $\mu(\cdot)$ and a covariance function $k(\cdot, \cdot | \Phi)$ on a set of output of size N , $\mathbf{Y} = (y_1, \dots, y_N)^T$ at input locations $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_N)^T$ provides a likelihood given by the multivariate Gaussian distribution. The covariance kernel $k : \mathbb{R}^{d+1} \times \mathbb{R}^{d+1} \rightarrow \mathbb{R}$ is positive definite and is often parameterized by additional parameter set Φ that controls various aspects of the input dependency. Inference on this additional unknown parameter is typically done via likelihood maximization or MCMC. A common practice in the computer experiments literature is to replace the actual simulation f by a fitted GP on N experiments for downstream tasks such as optimization, calibration or sensitivity analysis (see e.g., [22, 15]). When the observations $\mathbf{Y}$ from a computer experiment are noisy, white noise with fixed variance τ^2 is added to the covariance matrix such that $\mathbf{Y} \sim MVN(\mu, \mathbf{K}_N + \Sigma_N)$, where $\mathbf{K}_N$ is the $N \times N$ matrix with ij coordinates $k(\mathbf{x}_i, \mathbf{x}_j)$ and $\Sigma_N = \tau^2 \mathbb{I}$. Prediction at a new input $\mathbf{x}$ conditioning on the observations, i.e., $\mathbf{Y}(\mathbf{x}) | \mathbf{Y}$ also follows a Gaussian distribution with mean and variance given by:

$$\begin{aligned} \mu(\mathbf{x}) &= \mathbf{k}(\mathbf{x})^T (\mathbf{K}_N + \Sigma_N)^{-1} \mathbf{Y}, \text{ where } \mathbf{k}(\mathbf{x}) = (k(\mathbf{x}, \mathbf{x}_1), \dots, k(\mathbf{x}, \mathbf{x}_N))^T; \\ \sigma^2(\mathbf{x}) &= k(\mathbf{x}, \mathbf{x}) + \tau^2 - \mathbf{k}(\mathbf{x})^T (\mathbf{K}_N + \Sigma_N)^{-1} \mathbf{k}(\mathbf{x}), \end{aligned} \tag{1}$$

Without loss of generality, we assume $\mu = 0$ for the rest of the description.

When replicates (repeated experimental runs at the same design point) are available, the above GP formulation remains valid. However, we can achieve computational gain without sacrificing statistical properties by aggregating replicates across inputs and modeling their mean with an updated covariance matrix to account for non-constant variance across inputs. This is often called the stochastic kriging model [1], while other options to model input-dependent noise can be found in [14, 21, 6, 12].

In this work, however, we are explicitly interested in modeling individual *trajectories* rather than the mean behavior, a desirable task when the random seed encodes some essential structural information into the simulation. During calibration, rather than finding the parameter $\mathbf{x}$ that best matches observational data in expectation, we search for the optimal parameter-seed pair $(\mathbf{x}, r)$.

2.1 Common Random Number Gaussian Process (CRNGP)

We now define a new class of GP surrogates for stochastic computer simulations that models each replicate by exploiting the common random numbers [8, 27, 28]. Referring back to the augmented input tuple $(\mathbf{x}, r)$ for stochastic f , we formally define f to be a mapping $\mathbb{R}^d \times \mathcal{S} \rightarrow \mathbb{R}$, where $\mathcal{S}$ is the set of random seeds. In general, $\mathcal{S}$ can be arbitrary, but the exact specification is not necessary for its use. Without loss of generality, we assume $\mathcal{S}$ to be countable. Now that f is deterministic, we fall back to the standard GP prior with a carefully drafted covariance kernel k to deal with this augmented input. We assume a two-component separable covariance structure here such that, $k((\mathbf{x}_i, r), (\mathbf{x}_j, r')) = k(\mathbf{x}_i, \mathbf{x}_j) \times k_s(r, r')$. The novelty is in the kernel k_s where we put our assumption(s) on replicate similarity while the kernel k for $\mathbf{x}$ can be set to a common family of covariance kernels such as Gaussian or Matérn.

Without any prior information on how replicates are related in a given computer simulation, we put a simplified assumption that the marginal covariance between responses at two random seeds r and r' is constant, $k_s(r, r') = \rho$, $0 < \rho < 1$. With this new specification, we write the computer simulation output vector $\mathbf{Y} = (y_1^{(1)}, \dots, y_1^{(r_1)}, \dots, y_n^{(1)}, \dots, y_n^{(r_n)})$ and the unique input locations $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_n)$, where, r_i denotes the number of replicates at $\mathbf{x}_i$ with $\sum_{i=1}^n r_i = N$. Without loss of generality, we also denote the random seeds by the natural numbers since they are treated as categorical in the GP formulation. The elements in covariance matrix $\mathbf{K}_N$ is expressed as:

$$k((\mathbf{x}_i, r), (\mathbf{x}_j, r')) = \begin{cases} k(\mathbf{x}_i, \mathbf{x}_j) & \text{if } r = r' \\ k(\mathbf{x}_i, \mathbf{x}_j) \times \rho & \text{otherwise,} \end{cases} \quad (2)$$

where ρ can be interpreted as an adjustment to the similarity between responses at $\mathbf{x}_i$ and $\mathbf{x}_j$ if different seeds are used at each $\mathbf{x}$. Moreover, this specification of ρ results in a valid covariance kernel. A constant nugget τ^2 (typically a small positive number) is added to the diagonal of the matrix $\mathbf{K}_N$, i.e., $\mathbf{\Sigma}_N = \tau^2 \mathbf{I}$. With this formulation, the kriging formula in Equation (1) still holds true with a little modification of the new predicting location $\mathbf{x}$ which now actually means prediction of a particular replicate $(\mathbf{x}, r)$. Obtaining the mean behavior (ignoring the seed information) is equivalent to making a prediction at a new seed (i.e., for which no observation has been made) [28]. Figure 1 shows realizations from CRNGP for three different values of ρ where k is set to a Gaussian covariance kernel with length-scale and nugget at 0.01 and 0.001 respectively. The notable observation from this figure is the role of ρ in controlling the similarities among three replicated trajectories. As ρ increases, the trajectories tend to behave similarly at each $\mathbf{x}$. The bottom-right panel shows a modified version of the proposed CRNGP (hereafter local CRNGP) where the value of ρ depends on the input. We describe this new family of models in detail below.

Figure 2 provides a glimpse of the fundamental difference between CRNGP and other standard GP approaches, we choose heteroskedastic GP [6] here for demonstration. When the random seed information is available, CRNGP is perfectly able to interpolate the response surface at the replicate level (for $x = 0$ and $x = 1$), whereas, a random seed agnostic GP simply tracks the mean behavior of the response at each (x, T) .

The estimated value of ρ in CRNGP is 0.56. An optimization algorithm based on just the mean behavior may result in a trajectory that would be otherwise infeasible at any $\mathbf{x}$.




Figure 1: Realizations from CRNGP for different values of ρ (keeping the other GP hyperparameters fixed) except for the bottom-right panel which shows realization from a local CRNGP model. Each line corresponds to a random realization from the CRNGP which is vector output (trajectory) indexed by variable T (scaled to $[0, 1]$). In practice, T can be thought of as discrete time indices where the simulations are evaluated. Solid and dashed lines correspond to realizations at $x = 1$ and $x = 0$ respectively. The local CRNGP consists of a mixture of two CRNGPs with different values of ρ s on the partition on T -axis $\{(0, 0.5), (0.5, 1)\}$.

Locally approximated CRNGP Note that the modified covariance kernel in CRNGP still maintains the stationarity and hence can be restrictive in real applications. Following common approaches for dealing with non-stationarity via local approximations ([16], [9]), we propose a similar recipe to extend the stationary CRNGP to support non-stationary realizations. To relax the assumption that realizations at any $\mathbf{x}$ has a marginal covariance of ρ that does not depend on the actual value of $\mathbf{x}$, we start with L partition of the input space $\mathcal{X}^d = \bigcup_{l=1}^L \mathcal{X}_l^d$ with $\mathcal{X}_l^d \cap \mathcal{X}_{l'}^d = \emptyset$ for any $l \neq l'$. An independent CRNGP is fitted to each subset of the data $(\mathbf{X}_l, \mathbf{Y}_l)$ with its own hyperparameters according to its membership into one of the L bins given the partition. Prediction at a new $\mathbf{x}^* \in \mathcal{X}_{l^*}^d$ is followed from the model in l^* partition. While this strategy provides an easy and cheap alternative to a full input-dependent GP hyperparameter estimation, a realization from this locally approximated CRNGP is not smooth and can show unrealistic behaviors at the partition boundaries. Moreover, prediction at a boundary should be influenced by all partitions sharing the boundary than just one partition. For the downstream tasks (such as optimization using Thompson sampling, discussed in the next section), a non-smooth realization on the whole input space may give rise to inaccuracies.

Instead, we define the predictive distribution at any $\mathbf{x}^*$ to be a weighted prediction across all L local models where the weights are proportional to the inverse of the distance between $\mathbf{x}^*$ and $\mathcal{X}_l^d$ for all $l = 1, \dots, L$. Aggregating local models is also the idea behind product of experts or Bayesian committee machines, which come in many flavors, see, e.g., [31, 11] and references therein. For the purpose of generating realizations from this construction, one can bypass the exact analytic form of the mean and variance of the predictive




Figure 2: Comparison between fitted CRNGP and Heteroskedastic GP (HetGP) on vector-valued response at 4 different values of input x . Each color corresponds to a random seed value except the purple which represents the predicted mean surface from the HetGP model. The top panel shows the prediction from CRNGP which occurs at trajectory level and the bottom panel shows one mean prediction at each x by a heterogeneous GP.

distribution at $\mathbf{x}^*$ and directly apply the averaging on realizations from individual models. Let $\mathbf{Y}^{(l)}(\mathbf{x}^*, s)$ be the realization from the predictive distribution of $\mathbf{Y} | (\mathbf{X}_l, \mathbf{Y}_l)$, then the combined realization is given by,

$$\mathbf{Y}(\mathbf{x}^*, s) = \sum_{l=1}^L \alpha(\mathbf{x}^*, \mathbf{X}_l) \mathbf{Y}^{(l)}(\mathbf{x}^*, s), \quad (3)$$

where the α function depicts the distance between the prediction location and the training data in each partition.

2.2 Optimization using CRNGP surrogate

Gaussian processes are the workhorse of Bayesian optimization [13], also called efficient global optimization [18]. The principle is to rely on the prediction and uncertainty quantification capabilities of a GP surrogate to balance between exploitation of pertinent areas where the mean is promising and exploration of unknown regions of the parameter space where the variance is large. This task is handled by the acquisition function or infill criterion, such as the expected improvement [24] for global optimization. We refer to [13] for a more detailed presentation of BO.

In Bayesian optimization, scalar outputs are the most typical use case, with many possible variations. Related to our TOO approach, [2] solve similar trajectory problems by exploiting the composite aspect of the cost function to optimize: $g(\mathbf{y}(\mathbf{x})) = \sum_j (y_j(\mathbf{x}) - y_j^{obs})^2$ where y_j^{obs} are the actual observations that the simulator trajectory should match. But they solely focus on the mean behavior and not on the actual trajectories. This is also the case of [27, 28] where they exploit common random numbers with an enriched CRNGP model.

Our goal here is to identify

$$\left\{ (\mathbf{x}^{(1)}, r^{(1)}), \dots, (\mathbf{x}^{(p)}, r^{(p)}) \text{ s.t. } g(\mathbf{y}(\mathbf{x}^{(i)}, r^{(i)})) \leq g(\mathbf{y}(\mathbf{x}, r^{(j)})) \forall \mathbf{x} \in \mathcal{X} \forall 1 \leq i \leq p, 1 \leq j \leq |\mathcal{S}| \right\}$$

where p is the number of trajectories needed. To do so, we look at the following optimization problem as a proxy:

$$\arg \min_{\mathbf{x} \in \mathcal{X}, r \in \mathcal{S}} \sum_j \left(f(\mathbf{x}, t_j, r) - y_j^{(obs)} \right)^2 \quad (4)$$

where the set $\mathcal{S}$ is countable (in practice, we set $\mathcal{S}$ to be a large finite set of random seeds). Compared to the more usual goal of solving $\arg \min_{\mathbf{x} \in \mathcal{X}} \mathbb{E}_r \left[\sum_j \left(f(\mathbf{x}, t_j, r) - y_j^{(obs)} \right)^2 \right] = g(\mathbf{y}(\mathbf{x}))$, we are not focusing on finding the parameters $\mathbf{x}$ that perform best on average but on finding good trajectories even if the mean behavior does not correspond to the data. This would allow us to provide a more diverse set of inputs than when focusing solely on the mean. Note that this is also a different approach than focusing on risk-averse (or robust) optimization.

For the trajectory-oriented setup, we propose to use Thompson sampling (TS) [35], a simple and versatile approach still balancing exploration and exploitation. TS amounts to sampling one realization from the GP posterior distribution $\tilde{\mathbf{Y}} \sim \mathcal{N}(\mu(\tilde{\mathbf{X}}, S), \sigma^2(\tilde{\mathbf{X}}, S))$ over discrete sets $\tilde{\mathbf{X}}$ and S , then computing $g(\tilde{\mathbf{Y}})$ before selecting $(\mathbf{x}^*, r^*) \in \arg \min_{(\mathbf{x} \in \tilde{\mathbf{X}}, r \in S)} g(\tilde{\mathbf{Y}})$ as next input to evaluate. Extension to a continuous input space can be

performed by using varying random discrete sets or optimizing a continuous approximation of the realization [25]. In practice, as in [28], the set S only needs to contain seed values that have already been evaluated, plus one additional seed value accounting for the mean behavior over all unobserved seeds. This way, a new seed is evaluated only if it does better than all existing ones while promising seeds can be kept to improve over $\mathbf{x}$ (line 4 in the Algorithm 1). The general pseudo-code of the approach is illustrated in Algorithm 1. The batch extension of this Thompson sampling scheme is performed by sampling several posterior samples and using their respective best solutions as batch candidates [19]. The choice of the number of initial simulations— N_0 and total simulations— $N_{\max}$ would generally be dictated by the total computational budget available.

Algorithm 1 Pseudo-code for CRNGP based trajectory search

Require: N_0 (initial budget), $N_{\max}$ (total budget)

- 1: Create initial design of experiments, $(\mathbf{X}_n, S_n)$ of size $n = N_0$ and evaluate it $\mathbf{Y}_n = f(\mathbf{X}_n, S_n)$
 - 2: Train the GP model on $(\mathbf{X}_n, S_n, \mathbf{Y}_n)$
 - 3: **while** $n \leq N_{\max}$ **do**
 - 4: $S_n \leftarrow \{S_n, (\max S_n) + 1\}$
 - 5: Sample from the GP posterior: $\tilde{\mathbf{Y}} \sim \mathcal{N}(\mu(\tilde{\mathbf{X}}, S_n), \sigma^2(\tilde{\mathbf{X}}, S_n))$
 - 6: Choose $(\mathbf{x}_{n+1}, s_{n+1}) \in \arg \max_{\mathbf{x} \in \tilde{\mathbf{X}}, s \in S_n} g(\tilde{\mathbf{Y}})$
 - 7: Update the GP model by conditioning on $(\mathbf{x}_{n+1}, s_{n+1})$.
 - 8: $n \leftarrow n + 1$
 - 9: **end while**
-

The calibration task may have $p \geq 2$ targets, corresponding to various observables matching outputs of the simulators. In such a case, there is in general no solution perfectly matching all outputs simultaneously, and the solution is defined based on compromise. That is, a solution is said to be Pareto non-dominated if there is no other one that is at least as good on all objectives and strictly better in at least one. The set of non-dominated solutions is called the Pareto set in the input space, its image in the output space being the Pareto front. In addition to fitting one GP to each output and sampling from their respective posteriors, the sequential selection procedure must be modified as well. For instance, the hypervolume contribution, e.g., [4] of each element of the GP posterior, i.e., $(g(\tilde{\mathbf{Y}}_1^{(i)}), \dots, g(\tilde{\mathbf{Y}}_p^{(i)}))$, is the volume added to the current Pareto formed by the $(g(y_1^{(i)}), \dots, g(y_p^{(i)}))_{1 \leq i \leq n}$.




Figure 3: Thompson sampling (TS) based single objective optimization using CRNGP surrogate. Figures in the left show the predicted mean surface by CRNGP conditioning on the observations (represented by dots) and the figures in the right show random realizations from the fitted CRNGP based on which the new evaluations are obtained in the TS algorithm. Replicates are represented by different colors. Rows correspond to different iterations in the TS. In each iteration a new x for next evaluation is set at the arg min of a realization (right column). The global minimum is found in 11 iterations.

3 Empirical evaluation

3.1 Stochastic Epidemic Model

We used a published stochastic SEIR compartmental model [32] to simulate SARS-CoV-2 transmission and COVID-19 disease states. The model included multiple symptom statuses (asymptomatic, presymptomatic, mild, and severe), and severe disease outcomes (requiring hospitalization, critical illness requiring intensive care unit ICU admission, and deaths).

At the start of the epidemic, everyone is susceptible (S). The susceptible population is exposed and infected by the virus at a rate controlled by the transmission rate parameter, the number of infectious and susceptible people. After a latent period of a few days, exposed individuals (E) become infectious without symptoms: they are either completely asymptomatic (As) or presymptomatic (P). At the end of the incubation period, presymptomatic people proceed to develop either mild (Sm) or severe (Ss) symptoms. Asymptomatic or people with mild symptoms recovered (R) eventually, while those with severe symptoms would require hospital care (H) after a few days in the Ss compartment. Whether or not they eventually sought care, people who need hospital care either recover without complications or develop critical illness

(C) requiring ICU care. The latter could eventually recover from the illness by first moving back to the post-ICU hospitalization state (H_p) or die (D) from the illness. All exposed people are undetected. When a person enters the asymptomatic, presymptomatic, symptomatic mild, and symptomatic severe state, a fraction of them would be detected after a few days, which leads to isolation and reduced infectiousness. A schematic representation of the model is provided in Figure 4.

Model parameters such as incubation period, disease symptomatocity and severity, and hospital and ICU lengths of stay, are informed by literature and observed data in Chicago. The transmission rate, probability of detection, and the disease recovery rate are varied over time to better fit the daily ICU census, hospital census, and reported deaths in Chicago. A detailed description of the model and the parameters along with the source code can be found at [this github repository](#) [33].




Figure 4: **SEIR model structure to simulate SARS-CoV-2 transmission.** S=Susceptible, E=Exposed, A=Asymptomatic, P=Presymptomatic, Sm=Mild symptomatic, Ss=Severe symptomatic, H=Requires hospitalization, C=Critically ill (ICU), D=Death, R=Recovered. Subscripts u and d denote if the infection is undetected or detected. Asymptomatic and presymptomatic infections may also be detected, which would result in less infectiousness.

3.2 Single and Bi-objective Optimization

To illustrate our method, we use a three parameter version of the stochastic compartmental model described above. The three input parameters are the initial disease transmissibility rate and the rate of exposed individuals becoming infectious for both asymptomatic and symptomatic cases. Other disease parameters in the model are kept fixed at values informed by literature and the observed data in Chicago. Each simulation is run for the first 100 days of the disease propagation and output is obtained at 5 discretized time values, essentially resulting in a vector-valued output of size 5. For the output of interest, we choose cumulative hospital admissions in the single objective case along with cumulative deaths in the bi-objective case. A trajectory corresponding to a random (input, random seed) pair is treated as the ground truth and we set our goal to discover all trajectories that are consistent with this ground truth. In each case we set our initial simulation budget $N_0 = 50$, which is uniformly distributed on the three dimensional input space according to a Latin hypercube design [23] and maximum budget of $N_{\max} = 200$.

We also use the local version of the CRNGP surrogate, based on a fixed partition on the time component of the input. Essentially our CRNGP surrogate for vector valued output operates on five dimensional input space – three disease model parameters, discretized time indices, and the random seed. We choose only the time dimension for constituting a partition on which local surrogates would be more appropriate. This is further evident from Figure 5 and Figure 6 as to how trajectories become more diverse with time. The other aspect of a useful partition for local surrogates has to do with the number of actual simulations in each

partition. A surrogate trained on only a few simulations in each partition in a high input dimension setting may defeat the goal of creating an accurate surrogate.

Figure 5 shows the 50 best trajectories that match with the ground truth (denoted by solid black dots) after the total simulation budget is exhausted. Searching for optimal input parameter values using HetGP surrogate in TS results in settings where the mean behavior of the stochastic simulation is expected to match the ground truth, whereas, the TOO approach facilitates finding not only the optimum input parameter values but also the replicates, i.e., an actual trajectory that matches with the ground truth. When calibrating to two objectives – hospitalizations and deaths, Figure 6 shows the best trajectories with respect to both quantity of interest.




Figure 5: Best 50 trajectories of cumulative hospitalizations that are consistent with the observations denoted by black dots when using CRNGP and HetGP surrogates in Algorithm 1




Figure 6: Best trajectories with respect to two objectives—cumulative hospitalizations and deaths simultaneously. Solid and broken lines represent CRNGP and HetGP surrogates respectively.

4 CONCLUSION

In the expensive simulation context, see, e.g., [3], where simulation time is not negligible, the use of surrogate models such as Gaussian processes is required to efficiently obtain a solution of an inverse problem within a limited computational budget. For a simulation that produces a stochastic response, often characterized by a random seed, we have proposed the CRNGP surrogate that can map each individual simulation response to a unique (input, random seed) pair. Furthermore, we have devised an optimization scheme that when paired with CRNGP can facilitate finding actual simulation outputs that are consistent with observations. The proposed method has been developed for handling vector-valued output, hence we use the term *trajectory-oriented optimization*.

In our specific epidemic application, the end goal of calibrating the simulation model is to collect simulated trajectories of disease outcomes that are consistent with empirical data trajectories. For a stochastic epidemic model (such as an ABM), each output of the simulation model may correspond to a specific instantiation of population mixing, which is often not controlled by the input parameters, but by the random seeds. Thus, in order to capture mixing patterns consistent with empirical data, one needs to find both optimal parameter settings and their corresponding random seeds. Specific mixing patterns are important to try to capture since they govern the pathways through which an epidemic will transmit, affecting how public health interventions may mitigate that transmission. Separately, in order to implement data assimilation for an ABM without rerunning the model from the very beginning of an epidemic, the state of the model, including the complete state of the agents in the model, needs to be captured at the end of each assimilating time period, so that it can be restarted from that time. The TOO approach explicitly seeks to find the best trajectories to keep in this case.

We have shown that using a CRNGP, we are not only able to find good trajectories, but we can do so within a smaller simulation budget. This fast time-to-solution aspect is particularly critical for producing actionable insights for decision makers and has been highlighted in [5]. Further improvement to time-to-solution is possible by considering the simulation time as an additional variable to control (see [34], [29]) and by developing open science platforms [10] to enable these analyses. The local CRNGP surrogate can also be extended to handle input-dependent replicate similarity by considering a more complex prior model for ρ .

5 ACKNOWLEDGEMENTS

This material is based upon work supported by the National Science Foundation under Grant No. 2200234, the National Institutes of Health under grants R01AI136056 and R01DA055502, the U.S. Department of Energy, Office of Science, under contract number DE-AC02-06CH11357, and the DOE Office of Science through the Bio-preparedness Research Virtual Environment (BRaVE) initiative. This research was completed with resources provided by the Laboratory Computing Resource Center at Argonne National Laboratory (Bebop cluster).

References

- [1] Bruce Ankenman, Barry L Nelson, and Jeremy Staum. Stochastic kriging for simulation metamodeling. *Operations research*, 58(2):371–382, 2010.
- [2] Raul Astudillo and Peter I Frazier. Thinking inside the box: A tutorial on grey-box bayesian optimization. In *2021 Winter Simulation Conference (WSC)*, pages 1–15. IEEE, 2021.
- [3] Evan Baker, Pierre Barbillon, Arindam Fadikar, Robert B Gramacy, Radu Herbei, David Higdon, Jiangeng Huang, Leah R Johnson, Pulong Ma, Anirban Mondal, et al. Analyzing stochastic computer models: A review with opportunities. *Statistical Science*, 37(1):64–89, 2022.

- [4] Nicola Beume, Boris Naujoks, and Michael Emmerich. Sms-emoa: Multiobjective selection based on dominated hypervolume. *European Journal of Operational Research*, 181(3):1653–1669, 2007.
- [5] Mickael Binois, Nicholson Collier, and Jonathan Ozik. A portfolio approach to massively parallel Bayesian optimization. *arXiv:2110.09334 [math, stat]*, 10 2021. arXiv: 2110.09334.
- [6] Mickaël Binois, Robert B. Gramacy, and Mike Ludkovski. Practical heteroscedastic Gaussian process modeling for large simulation experiments. *Journal of Computational and Graphical Statistics*, 27(4):808–821, 2018.
- [7] Mathieu Carmassi, Pierre Barbillon, Matthieu Chiodetti, Merlin Keller, and Eric Parent. Bayesian calibration of a numerical code for prediction. *Journal de la société française de statistique*, 160(1):1–30, 2019.
- [8] Xi Chen, Bruce E Ankenman, and Barry L Nelson. The effects of common random numbers on stochastic kriging metamodels. *ACM Transactions on Modeling and Computer Simulation (TOMACS)*, 22(2):1–20, 2012.
- [9] Hugh A Chipman, Edward I George, and Robert E McCulloch. Bart: Bayesian additive regression trees. 2010.
- [10] Nicholson Collier, Justin M. Wozniak, Abby Stevens, Yadu Babuji, Mickaël Binois, Ardindam Fadikar, Alexandra Würth, Kyle Chard, and Jonathan Ozik. Developing distributed high-performance computing capabilities of an open science platform for robust epidemic analysis. 2023. Publisher: arXiv Version Number: 1.
- [11] Adam M Edwards and Robert B Gramacy. Precision aggregated local models. *Statistical Analysis and Data Mining: The ASA Data Science Journal*, 14(6):676–697, 2021.
- [12] Arindam Fadikar, Dave Higdon, Jiangzhuo Chen, Bryan Lewis, Srinivasan Venkatramanan, and Madhav Marathe. Calibrating a stochastic, agent-based model using quantile-based emulation. *SIAM/ASA Journal on Uncertainty Quantification*, 6(4):1685–1706, 2018.
- [13] Roman Garnett. *Bayesian optimization*. Cambridge University Press, 2023.
- [14] Paul W Goldberg, Christopher KI Williams, and Christopher M Bishop. Regression with input-dependent noise: A gaussian process treatment. *Advances in neural information processing systems*, pages 493–499, 1998.
- [15] Robert B. Gramacy. *Surrogates: Gaussian Process Modeling, Design and Optimization for the Applied Sciences*. Chapman Hall/CRC, Boca Raton, Florida, 2020.
- [16] Robert B. Gramacy and Daniel W. Apley. Local Gaussian process approximation for large computer experiments. *Journal of Computational and Graphical Statistics*, 24(2):561–578, 2015.
- [17] Dave Higdon, Marc Kennedy, James C Cavendish, John A Cafeo, and Robert D Ryne. Combining field data and computer simulations for calibration and prediction. *SIAM Journal on Scientific Computing*, 26(2):448–466, 2004.
- [18] Donald R. Jones, Matthias Schonlau, and William J. Welch. Efficient global optimization of expensive black-box functions. *Journal of Global Optimization*, 13(4):455–492, 1998.
- [19] Kirthevasan Kandasamy, Akshay Krishnamurthy, Jeff Schneider, and Barnabás Póczos. Parallelised bayesian optimisation via thompson sampling. In *International Conference on Artificial Intelligence and Statistics*, pages 133–142. PMLR, 2018.
- [20] M. Kennedy and A. O’Hagan. Bayesian calibration of computer models (with discussion). *Journal of the Royal Statistical Society, Series B*, 63:425–464, 2001.
- [21] Kristian Kersting, Christian Plagemann, Patrick Pfaff, and Wolfram Burgard. Most likely heteroscedastic gaussian process regression. In *Proceedings of the 24th international conference on Machine learning*, pages 393–400. ACM, 2007.
- [22] Jack PC Kleijnen. *Design and analysis of simulation experiments*. Springer, 2018.
- [23] Michael D McKay, Richard J Beckman, and William J Conover. A comparison of three methods for selecting values of input variables in the analysis of output from a computer code. *Technometrics*, 42(1):55–61, 2000.
- [24] Jonas Mockus, Vytautas Tiešis, and Antanas Žilinskas. The application of Bayesian methods for seeking the extremum. *Towards Global Optimization*, 2(117-129):2, 1978.
- [25] Mojmir Mutny and Andreas Krause. Efficient high dimensional bayesian optimization with additivity and quadrature fourier features. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems 31*, pages 9005–9016. Curran Associates, Inc., 2018.
- [26] Jonathan Ozik, Justin M Wozniak, Nicholson Collier, Charles M Macal, and Mickaël Binois. A population data-driven workflow for covid-19 modeling and learning. *The International Journal of High Performance Computing Applications*, 35(5):483–499, 2021.
- [27] Michael Pearce, Matthias Poloczek, and Juergen Branke. Bayesian simulation optimization with common random numbers. In *2019 Winter Simulation Conference (WSC)*, pages 3492–3503. IEEE, 2019.
- [28] Michael Arthur Leopold Pearce, Matthias Poloczek, and Juergen Branke. Bayesian optimization allowing for common random numbers. *Operations Research*, 70(6):3457–3472, 2022.

- [29] Victor Picheny and David Ginsbourger. A nonstationary space-time gaussian process model for partially converged simulations. *SIAM/ASA Journal on Uncertainty Quantification*, 1(1):57–78, 2013.
- [30] Carl Edward Rasmussen and Christopher K. I. Williams. *Gaussian Processes for Machine Learning*. MIT Press, Cambridge, MA, 2005.
- [31] Didier Rulhière, Nicolas Durrande, François Bachoc, and Clément Chevalier. Nested kriging predictions for datasets with a large number of observations. *Statistics and Computing*, 28:849–867, 2018.
- [32] Manuela Runge, Reese AK Richardson, Patrick A Clay, Arielle Bell, Tobias M Holden, Manisha Singam, Natsumi Tsuboyama, Philip Arevalo, Jane Fornoff, Sarah Patrick, et al. Modeling robust covid-19 intensive care unit occupancy thresholds for imposing mitigation to prevent exceeding capacities. *PLOS Global Public Health*, 2(5):e0000308, 2022.
- [33] Manuela et. al. Runge. Covid-chicago. <https://github.com/numalariamodeling/covid-chicago>, 2022. [Online; accessed 8-July-2023].
- [34] Jasper Snoek, Hugo Larochelle, Jasper Snoek, and Ryan P. Adams. Practical Bayesian optimization of machine learning algorithms. In *Neural Information Processing Systems*, 2012.
- [35] William R Thompson. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25(3-4):285–294, 1933.