

# ACCELERATING STOCHASTIC SIMULATION WITH INTERACTIVE NEURAL PROCESSES

**Anonymous authors**

Paper under double-blind review

## ABSTRACT

Stochastic simulations such as large-scale, spatiotemporal, age-structured epidemic models are computationally expensive at fine-grained resolution. We propose Interactive Neural Process (INP), a Bayesian active learning framework to proactively learn a deep learning surrogate model and accelerate simulation. Our framework is based on the novel integration of neural process, deep sequence model and active learning. In particular, we develop a novel spatiotemporal neural process model to mimic the simulator dynamics. Our model automatically infers the latent process which describes the intrinsic uncertainty of the simulator. This also gives rise to a new acquisition function based on the latent information gain. We design Bayesian active learning algorithms to iteratively query the simulator, gather more data, and continuously improve the model. We perform theoretical analysis and demonstrate that our approach reduces sample complexity compared with random sampling in high dimension. Empirically, we demonstrate our framework can faithfully imitate the behavior of a complex infectious disease simulator with a small number of examples, enabling rapid simulation and scenario exploration.

## 1 INTRODUCTION

Computational modeling is now more than ever at the forefront of infectious disease research due to the COVID-19 pandemic. Stochastic simulations play a critical role in understanding and forecasting infectious disease dynamics, creating what-if scenarios, and informing public health policy making (Cramer et al., 2021). More broadly, stochastic simulation (Ripley, 2009; Asmussen & Glynn, 2007) produces forecasts about complex interactions among people, environment, space, and time given a set of parameters. They provide the numerical tools to simulate probabilistic models and stochastic processes in finance (Lamberton & Lapeyre, 2007), chemistry (Gillespie, 2007) and many scientific disciplines.

Unfortunately, stochastic simulations at fine-grained spatial and temporal resolution are extremely computationally expensive. For example, epidemic models whose aim is to simulate realistic diffusion dynamics via in-silico experiments require the exploration of a large parameter space (e.g. characteristics of a virus, policy interventions, people’s behavior). Therefore, hundreds of thousands of simulations are required to explore and calibrate the epidemic model with observed epidemic surveillance data. This process significantly hinders the adaptive capability of existing stochastic simulators, especially in “war time” emergencies, due to the lead time needed to execute new simulations and produce actionable insights that could help guide decision makers.

Learning deep surrogate models to speed up complex simulation has been explored in climate modeling and fluid dynamics for *deterministic* dynamics (Sanchez-Gonzalez et al., 2020; Wang et al., 2020; Holl et al., 2019; Rasp et al., 2018), but not for *stochastic* simulations. Furthermore, these surrogate models can only approximate specific system dynamics and fail to generalize under different parametrization. Especially for pandemic scenario planning, we desire models that can predict futuristic scenarios under different conditions. Furthermore, the majority of the surrogate models are trained *passively* using a simulation data set. This requires a large number of simulations to cover different parameter regimes of the simulator and ensure generalization.

We propose Interactive Neural Process (INP), a Bayesian active learning framework to speed up the stochastic simulation. Given parameters such as disease reproduction number, incubation and infectious periods, mobility dynamics, and initial state, current mechanistic simulators generate

future outbreak states with time-consuming numerical integration. Our framework accelerates the simulation by learning the input-output map between parameters and future states, hence bypassing numerical integration. We develop a novel neural process (NP) (Garnelo et al., 2018) model to mimic the mechanistic simulator with uncertainty. We extend NP to spatiotemporal setting: we introduce a time-evolving latent process for temporal dynamics and integrate graph convolution for spatial modeling.

Instead of learning passively, we design *active learning* algorithms to interact with the simulator and update our model in “real-time”. We derive a new acquisition function, Latent Information Gain (LIG), based on our unique model design. Our algorithm selects the parameters with the highest LIG, queries the simulator to generate new simulation data, and continuously updates our model. Compared with Gaussian processes (GP), INP is more expressive, more accurate, and scales easily to high-dimensional sequence data. We provide theoretical guarantees for the sample efficiency of this procedure over random sampling. We also demonstrate the efficacy of our method on large-scale spatiotemporal epidemic models. In summary, our contributions include:

- Interactive Neural Process: a deep Bayesian active learning framework for accelerating large-scale stochastic simulation
- A novel spatiotemporal neural process model (STNP) for high-dimensional time series data with temporal latent process and spatial graph convolution.
- New acquisition function, Latent Information Gain (LIG), based on the inferred temporal latent process to quantify uncertainty with theoretical guarantees
- Real-world application to speed up stochastic, spatiotemporal epidemic models.

## 2 RELATED WORK

**Bayesian Active Learning and Experimental Design.** Bayesian active learning, or experimental design is well-studied in statistics and machine learning (Chaloner & Verdinelli, 1995; Cohn et al., 1996). Gaussian Processes (GPs) are popular for posterior estimation e.g. Houlby et al. (2011) and (Zimmer et al., 2018), but often struggle in high dimension. Deep neural networks provide scalable solutions for active learning. Deep active learning has been applied to discrete problems such as image classification (Gal et al., 2017) and sequence labeling (Siddhant & Lipton, 2018) whereas our task is continuous time series. Our problem can also be viewed as sequential experimental design where we design simulation parameters to obtain the desired outcome (imitating the simulator). Foster et al. (2021) propose deep design networks for Bayesian experiment design but they require an explicit likelihood model and conditional independence in experiments. [Kleinegesse & Gutmann \(2020\)](#) consider implicit models where the likelihood function is intractable, but computing the Jacobian through sampling path can be expensive and their experiments are mostly limited to low ( $\leq 10$ ) dimensional design. In contrast, our design space is of much higher-dimension and we do not have access to an explicit likelihood model for the simulator.

**Neural Processes.** Neural processes (NP) (Garnelo et al., 2018) model distributions over functions and imbue neural networks with the ability of GPs to estimate uncertainty. NP has many extensions such as attentive NP (Kim et al., 2019) and functional NP (Louizos et al., 2019). However, NP implicitly assumes permutation invariance in the latent variables and can be limiting in modeling temporal dynamics. Singh et al. (2019) proposes sequential NP by incorporating a temporal transition model into NP. Still, sequential NP assumes the latent variables are independent conditioned on the hidden states. We propose *spatiotemporal* NP with temporal latent process and spatial graph convolution, which is well-suited for modeling the spatiotemporal dynamics of infectious disease. We apply our model to real-world large-scale Bayesian active learning. Note that even though Garnelo et al. (2018) has demonstrated NP for Bayesian optimization, it is only for toy 1-D functions.

**Stochastic Simulation and Dynamics Modeling.** Stochastic simulations are fundamental to many scientific fields (Ripley, 2009), especially epidemic modeling. Data-driven models of infectious diseases are increasingly used to forecast the evolution of an ongoing outbreak (Arik et al., 2020; Cramer et al., 2021; Lourenco et al., 2020). However, very few models can mimic the internal mechanism of a stochastic simulator and answer “what-if questions”. [GPs are commonly used as surrogate models for expensive simulators \(Meeds & Welling, 2014; Gutmann et al., 2016; Qian et al., 2020\)](#), but GPs do not scale well to high-dimensional data. [Likelihood-free inference Lueckmann](#)

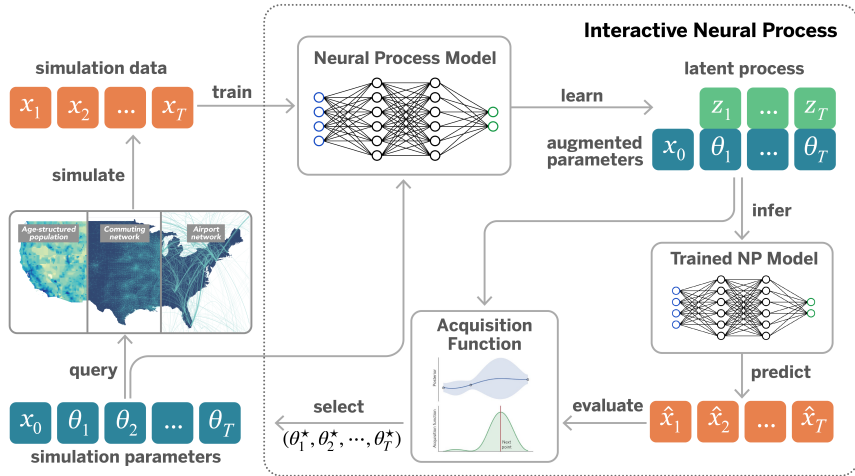


Figure 1: Illustration of the interactive neural process (INP). Given simulation parameters and data, INP trains a Neural Process model to infer the latent process. The inferred latent process allows prediction and uncertainty quantification. The uncertainty is used to select the next set of parameters to query the simulator, and simulate more data.

et al. (2019); Papamakarios et al. (2019); Munk et al. (2019); Wood et al. (2020) aim to learn the posterior of the parameters given observed data. They use neural density estimators to learn an emulator, but require a lot of simulations. We propose a novel neural process model to learn the likelihood. For active learning, instead of relying on Monte Carlo sampling, we directly compute the information gain in the latent process. For epidemic modeling, Qian et al. (2020) only consider the simple SEIR model. In contrast, we demonstrate the use of deep sequence model as a prior distribution in Bayesian active learning for real-world, large-scale, spatiotemporal dynamics. Our framework is also compatible with other deep sequence models for time series, e.g. Deep State Space (Rangapuram et al., 2018), Neural ODE (Chen et al., 2018) and etc.

### 3 METHODOLOGY

Consider a stochastic process  $\{X_1, \dots, X_T\}$ , governed by time-varying parameters  $\theta_t \in \mathbb{R}^K$ , and the initial state  $x_0 \in \mathbb{R}^D$ . In epidemic modeling,  $\theta_t$  can represent the effective reproduction number of the virus at a given time, the effective contact rates between individuals belonging to different age groups, the people’s degree of short- or long-range mobility, or the effects of time varying policy interventions (e.g. non-pharmaceutical interventions). The state  $x_t \in \mathbb{R}^D$  includes both the daily prevalence and daily incidence for each compartment of the epidemic model (e.g. number of people that are infectious and number of new infected individuals at time  $t$ ).

Stochastic simulation uses a mechanistic model  $F(\theta; \xi)$  to simulate the process where the random variable  $\xi$  represents the randomness in the simulator. Let  $\theta := (x_0, \theta_1, \dots, \theta_T)$  represent the initial state and all the parameters over time. For each  $\theta$ , we obtain a different set of simulation data  $\{(x_1, \dots, x_T)_m\}_{m=1}^M$ . However, realistic large-scale stochastic simulations require the exploration of a large parameter space and are extremely computationally intensive. In the following section, we describe a pipeline called Interactive Neural Process (INP) to proactively query the stochastic simulator, generate simulation data, in order to learn a fast surrogate model for rapid simulation.

#### 3.1 INTERACTIVE NEURAL PROCESS

Our high-level goal is to train a deep surrogate model to mimic the stochastic simulator. As show in Figure 1, given parameters  $\theta$ , we query the simulator, i.e., the mechanistic model to obtain a set of simulations  $\{(x_1, \dots, x_T)_m\}_{m=1}^M$ . We train a neural process (NP) model to learn the probabilistic map from parameters to future states. Our NP model is spatiotemporal and it captures the unique disease dynamics of the epidemic simulator. During inference, the model needs to generate predictions  $(\hat{x}_1, \dots, \hat{x}_T)$  at the target parameters  $\theta$  corresponding to different scenarios.

Instead of simulating at a wide range of parameter regimes, we take a Bayesian active learning approach to proactively query the simulator and update the model incrementally. Using NP, we can infer the latent temporal process  $(z_1, \dots, z_T)$  that encodes the uncertainty of the current surrogate




Figure 2: Graphical model comparison: Neural Process, Sequential Neural Process and our Spatiotemporal Neural Process. The key difference lies in the inference of the latent variable  $z_t$ . The distributions of  $z_t$  for NP, SNP and STNP are  $q(z_t|x_{t-1}, x_t, \theta_t, A)$ ,  $q(z_t|x_{t-1}, x_t, \theta_t, z_{1:t-1}, A)$ , and  $q(z_t|x_{0:t}, \theta_{1:t}, A)$  respectively.  $A$  is the adjacency matrix for the spatial graph, which is omitted in the figure for simplicity.  $h_t$  is the encoder hidden state and  $H_t$  is the decoder hidden state.

model. Then we propose a new reward/acquisition function, latent information gain (LIG), to select the  $\theta^*$  with the highest reward. We use  $\theta^*$  to query the simulator, and in turn generate new simulation to further improve the model. Next, we describe each of the components in detail.

### 3.2 SPATIOTEMPORAL NEURAL PROCESS

Neural process (NP) (Garnelo et al., 2018) is a type of deep generative model that represents distributions over functions. It approximates a stochastic process by combining Gaussian process with deep learning. NP introduces a global latent variable  $z$  to capture the stochasticity and learns the conditional distribution  $p(x_{1:T}|\theta)$  by optimizing the evidence lower bound (ELBO):

$$\log p(x_{1:T}|\theta) \geq \mathbb{E}_{q(z|x_{1:T}, \theta)} [\log p(x_{1:T}|z, \theta)] - \text{KL}(q(z|x_{1:T}, \theta)||p(z)) \quad (1)$$

Here  $p(z)$  is the prior distribution for the latent variable. We use  $x_{1:T}$  as a shorthand for  $(x_1, \dots, x_T)$ . The prior distribution  $p(z)$  is conditioned on a set of context points  $\theta^c, x_{1:T}^c$  as  $p(z|x_{1:T}^c, \theta^c)$ . The set of context points is randomly sampled from the training data. Its complement is the target set.

However, the global latent variable  $z$  in NP can be limiting for non-stationary, spatiotemporal dynamics in the epidemics. We propose *spatiotemporal* neural process (STNP) with two extensions. First, we introduce a temporal latent process  $(z_1, \dots, z_T)$  to represent the unknown dynamics. The latent process provides an expressive description of the internal mechanism of the stochastic simulator. Each latent variable  $z_t$  is sampled conditioning on the past history. Second, we explicitly model the spatial dependency in  $x_t \in \mathbb{R}^D$ . Rather than treating the dimensions in  $x_t$  as independent features, we capture their correlations with a graph. For instance, the travel graph between locations can be represented as an adjacency matrix  $A \in \mathbb{R}^{D \times D}$ .

Given parameters  $\{\theta\}$ , simulation data  $\{x_{1:T}\}$ , and the spatial graph  $A$  as inputs, STNP models the conditional distribution  $p(x_{1:T}|\theta, A)$  by optimizing the following ELBO objective:

$$\log p(x_{1:T}|\theta, A) \geq \mathbb{E}_{q(z_{1:T}|x_{1:T}, \theta, A)} \log p(x_{1:T}|z_{1:T}, \theta, A) - \text{KL}(q(z_{1:T}|x_{1:T}, \theta, A)||p(z_{1:T})) \quad (2)$$

where the distributions  $q(z_{1:T}|x_{1:T}, \theta, A)$  and  $p(x_{1:T}|z_{1:T}, \theta, A)$  are parameterized with neural networks. The prior distribution  $p(z_{1:T})$  is conditioned on a set of contextual sequences  $p(z_{1:T}|x_{1:T}^c, \theta^c, A)$ . Figure 2 visualizes the graphical models of our STNP, the vanilla NP (Garnelo et al., 2018), and Sequential NP (Singh et al., 2019).

We implement STNP following an encoder-decoder architecture. The encoder parametrizes the mean and standard deviation of the variational posterior  $q(z_{1:T}|x_{1:T}, \theta, A)$  and the decoder approximates the predictive distribution  $p(x_{1:T}|z_{1:T}, \theta, A)$ . To incorporate the spatial graph information, we use a Diffusion Convolutional Gated Recurrent Unit (DCGRU) layer (Li et al., 2017) which integrates graph convolution in a GRU cell. We use multi-layer GRUs to obtain hidden states from the inputs. Using re-parametrization (Kingma & Welling, 2013), we sample  $z_t$  from the encoder and then decode  $x_t$  conditioned on  $z_t$  in an auto-regressive fashion. To ensure fair comparisons, we adapt NP and SNP to graph-based settings and use the same architecture as STNP to generate the hidden states shown in Figure 6. See Appendix B.3 for architecture details.

### 3.3 BAYESIAN ACTIVE LEARNING

Algorithm 1 details a Bayesian active learning algorithm, based on Bayesian optimization (Shahriari et al., 2015; Frazier, 2018). We train a NP model to interact with the simulator and improve

**Algorithm 1:** Interactive Neural Process

---

**Input:** Initial simulation dataset  $\mathcal{S}_1$

- 1 Train the model  $\text{NP}^{(1)}(\mathcal{S}_1)$  ;
- 2 **for**  $i = 1, 2, \dots$  **do**
- 3     Learn  $(z_1, z_2, \dots, z_T) \sim q^{(i)}(z_{1:T}|x_{1:T}, \theta, \mathcal{S}_i)$  ;
- 4     Predict  $(\hat{x}_1, \hat{x}_2, \dots, \hat{x}_T) \sim p^{(i)}(x_{1:T}|z_{1:T}, \theta, \mathcal{S}_i)$  ;
- 5     Select a batch  $\{\theta^{(i+1)}\} \leftarrow \arg \max_{\theta} \mathbb{E}_{p(x_{1:T}|z_{1:T}, \theta)} [r(\hat{x}_{1:T}|z_{1:T}, \theta)]$  ;
- 6     Simulate  $\{x_{1:t}^{(i+1)}\} \leftarrow \text{Query the simulator } F(\theta^{(i+1)}; \xi)$  ;
- 7     Augment training set  $\mathcal{S}_{i+1} \leftarrow \mathcal{S}_i \cup \{\theta^{(i+1)}, x_{1:T}^{(i+1)}\}$  ;
- 8     Update the model  $\text{NP}^{(i+1)}(\mathcal{S}_{i+1})$  ;
- 9 **end**

---

learning. Let the superscript  $(i)$  denote the  $i$ -th interaction. We start with an initial data set  $\mathcal{S}_1 = \{\theta^{(1)}, x_{1:T}^{(1)}\}$  and use it to train our NP model and learn the latent process. During inference, given the augmented parameters  $\theta$  that are candidate scenarios, we use the trained NP model to predict the states  $(\hat{x}_1, \dots, \hat{x}_T)$ . We evaluate the models' predictions with an acquisition function  $r(\hat{x}_{1:T}, z_{1:T}, \theta)$  and select the set of parameters  $\{\theta^{(i+1)}\}$  with the highest reward. We query the simulator with  $\{\theta^{(i+1)}\}$  to augment the training data set  $\mathcal{S}_{i+1}$  and update the NP model for the next iteration.

The choice of the reward (acquisition) function  $r$  depends on the goal of the active learning task. For example, to find the model that best fits the data, the reward function can be the log-likelihood  $r = \log p(\hat{x}_{1:T}|\theta, A)$ . To collect data and reduce model uncertainty in Bayesian experimental design, the reward function can be the mutual information. In what follows, we discuss different strategies to design the reward/acquisition function. We also propose a novel acquisition function based on information gain in the latent space tailored to our STNP model.

### 3.4 REWARD/ACQUISITION FUNCTIONS

For regression tasks, standard acquisition functions for active learning include Maximum Mean Standard Deviation (Mean STD), Maximum Entropy, Bayesian Active Learning by Disagreement (BALD), and random sampling (Gal et al., 2017). With the distributions represented by NP, we explore various acquisition functions and their approximations. The temporal latent process in STNP encodes the internal mechanism of the stochastic simulator. We introduce a new acquisition function based on our unique NP design called Latent Information Gain (LIG).

**Maximum Mean STD.** Mean STD (Gal & Ghahramani, 2016) is a heuristic used to estimate the model uncertainty. For each augmented parameter  $\theta$ , we sample multiple  $z_{1:T}$  and generate a set of predictions  $\{\hat{x}_{1:T}\}$ . For a length  $T$  sequence with dimension  $D$ , we compute the standard deviation  $\sigma_{t,d}$  for time step  $t$  and feature  $d$ . Mean STD computes  $\bar{\sigma} = \frac{1}{TD} \sum_{t=1}^T \sum_{d=1}^D \sigma_{t,d}$  for each parameter  $\theta$ . We select the  $\theta$  with the maximum  $\bar{\sigma}$ . Empirically, we found that Mean STD often becomes over-conservative and tends to explore less.

**Maximum Entropy.** Maximum entropy computes the maximum predictive entropy as  $H(\hat{x}) = -\mathbb{E}[\log p(\hat{x}_{1:T})]$ . In general, entropy is intractable for continuous output. Our NP model implicitly assumes the predictions follow a multivariate Gaussian, which allows us to compute the differential entropy (Jaynes, 1957). We follow the same procedure as Mean STD to estimate the empirical covariance  $\Sigma \in \mathbb{R}^{TD \times TD}$  and compute the differential entropy for each parameter as  $H = \frac{1}{2} \ln |\Sigma| + \frac{TD}{2} (1 + \ln 2\pi)$ . We select the parameter  $\theta$  with the maximum entropy.

**BALD and EIG.** BALD (Houlsby et al., 2011) quantifies the mutual information between the prediction and model posterior  $H(\hat{x}_{1:T}|\theta) - H(\hat{x}_{1:T}|z_{1:T}, \theta)$ , which is equivalent to the expected information gain (EIG). Computing the EIG/BALD is challenging since  $p(\hat{x}_{1:T}|z_{1:T}, \theta)$  cannot be found in closed form in general. The integrand is intractable and conventional MC methods are not applicable (Foster et al., 2019). One way to get around this is to employ a nested MC estimator with quadratic computational cost for sampling (Myung et al., 2013; Vincent & Rainforth, 2017). We prove the equivalence between the expected information gain for selecting the parameter  $\theta$  and the

expected KL divergence in the latent processes, as illustrated by the following proposition. In this way, conventional MC method becomes applicable, which helps reduce the quadratic computational cost to linear.

**Proposition 1.** *The expected information gain (EIG) for neural process is equivalent to the KL divergence between the likelihood and posterior in the latent process, that is*

$$\text{EIG}(\hat{x}_{1:T}, \theta) := \mathbb{E}[H(\hat{x}_{1:T}) - H(\hat{x}_{1:T}|z_{1:T}, \theta)] = \mathbb{E}_{p(\hat{x}_{1:T}|\theta)} [\text{KL}(p(z_{1:T}|\hat{x}_{1:T}, \theta)||p(z_{1:T}))] \quad (3)$$

See proof in the appendix A.1. Inspired by this fact, we propose a novel acquisition function based on the information gain in the latent process.

**Latent Information Gain.** The trained NP model produces a variational posterior given the current dataset  $\mathcal{S}$  as  $p(z_{1:T}|\mathcal{S})$ . For every parameters  $\theta$  remained in the search space, we can predict  $\hat{x}_{1:T}$  with the decoder. We use  $\hat{x}_{1:T}$  and  $\theta$  as input to the encoder to re-evaluate the posterior  $p(z_{1:T}|\hat{x}_{1:T}, \theta, \mathcal{S})$ . LIG computes the distributional difference with respect to the latent process  $z_{1:T}$  as  $\mathbb{E}_{p(\hat{x}_{1:T}|\theta)} [\text{KL}(p(z_{1:T}|\hat{x}_{1:T}, \theta, \mathcal{S})||p(z_{1:T}|\mathcal{S}))]$  where  $\text{KL}(\cdot||\cdot)$  denotes the KL-divergence between two distributions. We select a batch of  $\theta$  with the highest latent information gain.

### 3.5 THEORETICAL ANALYSIS

We shed light onto the intuition behind choosing adaptive sample selection over random sampling via analyzing a simplifying situation. Assume that at a certain stage we have learned a feature map  $\Psi$  which maps the input  $\theta$  of the neural network to the last layer. Then the output  $X$  can be modeled as  $X = \langle \Psi(\theta), z^* \rangle + \epsilon$ , where  $z^*$  is the true hidden variable,  $\epsilon$  is the random noise.

Our goal is to generate an estimate  $\hat{z}$ , and use it to make predictions  $\langle \Psi(\theta), \hat{z} \rangle$ . A good estimate shall achieve small error in terms of  $\|\hat{z}_t - z^*\|_2$  with high probability. In the following theorem, we prove that greedily maximizing the variance of the prediction to choose  $\theta$  will lead to an error of order  $\mathcal{O}(d)$  less than that of random exploration in the space of  $\theta$ , which is significant in high dimension.

**Theorem 1.** *For random feature map  $\Psi(\cdot)$ , greedily optimizing the KL divergence,  $\text{KL}(p(z|\hat{x}, \theta)||p(z))$ , or equivalently the variance of the posterior predictive distribution  $\mathbb{E} \left[ (\langle \Psi(\theta), \hat{z} \rangle - \mathbb{E} \langle \Psi(\theta), \hat{z} \rangle)^2 \right]$  in search of  $\theta$  will lead to an error  $\|\hat{z}_t - z^*\|_2$  of order  $\mathcal{O}(\sigma d / \sqrt{t})$  with high probability.*

*On the other hand, random sampling of  $\theta$  will lead to an error of order  $\mathcal{O}(\sigma d^2 / \sqrt{t})$  with high probability.*

See proofs in the Appendix A.2.

## 4 EXPERIMENTS

We start with a simple susceptible-exposed-infectious-removed (SEIR) compartmental model. We compare our method with Gaussian process for validation. We then apply our INP framework to mimic the dynamics of a real-world large-scale Local Epidemic and Mobility model (LEAM-US) simulator and report the performance using different acquisition functions. We implement GP using gpytorch (Gardner et al., 2018) with Matern kernel. All NP models are implemented using pytorch (Paszke et al., 2019).

### 4.1 SEIR COMPARTMENTAL MODEL

**SEIR Simulation.** To show that INP indeed provides a more flexible and faster alternative to GP, we begin with a simple stochastic, discrete, chain-binomial SEIR compartmental model as our stochastic simulator. In this model, susceptible individuals ( $S$ ) become exposed ( $E$ ) through interactions with infectious individuals ( $I$ ) and are eventually removed ( $R$ ), see Appendix B.1 for details.

We set the total population  $N = S + E + I + R$  as 100,000, the initial number of exposed individuals as  $E_0 = 2,000$ , and the initial number of infectious individuals as  $I_0 = 2,000$ . The latent period varies between approximately 1.5 to 4 days, i.e. latent individuals move to the infectious stage at a rate  $\varepsilon \in [0.25, 0.65]$  (step 0.05) that is inversely proportional to the latent period, the infectious period






Figure 3: Neural Processes vs. Gaussian Processes for Bayesian active learning in SEIR compartmental model. Left: uncertainty quantification comparison with the truth. Right: MAE loss versus the percentage of data for both NP and GP using different acquisition functions.

$\mu^{-1}$  is set to be equal to 1 day, and we let the basic reproduction number  $R_0$  (which in this case coincides with the transmissibility rate  $\beta$ ) vary between 1.1 and 4.0 (step 0.1). Here, each  $(\beta, \varepsilon)$  pair corresponds to a specific scenario, which determines the 2-dimensional parameters  $\theta$ . We simulate the first 100 days of the epidemic for the infectious population (output dimension 100) with a total of 300 scenarios. We generate 30 samples for each scenario.

**Model Implementation.** We compare INP with GP for Bayesian active learning. In this toy experiment, we only predict the number of individuals in the infectious compartment. The input is  $(\beta, \varepsilon)$  pair and the output is the 100 days’ infection prediction. As the simulator is simple, we use the vanilla NP model with the global latent variable  $z$ . For each epoch, we randomly select 10% of the samples as context. Implementation details are deferred to Appendix B.5.

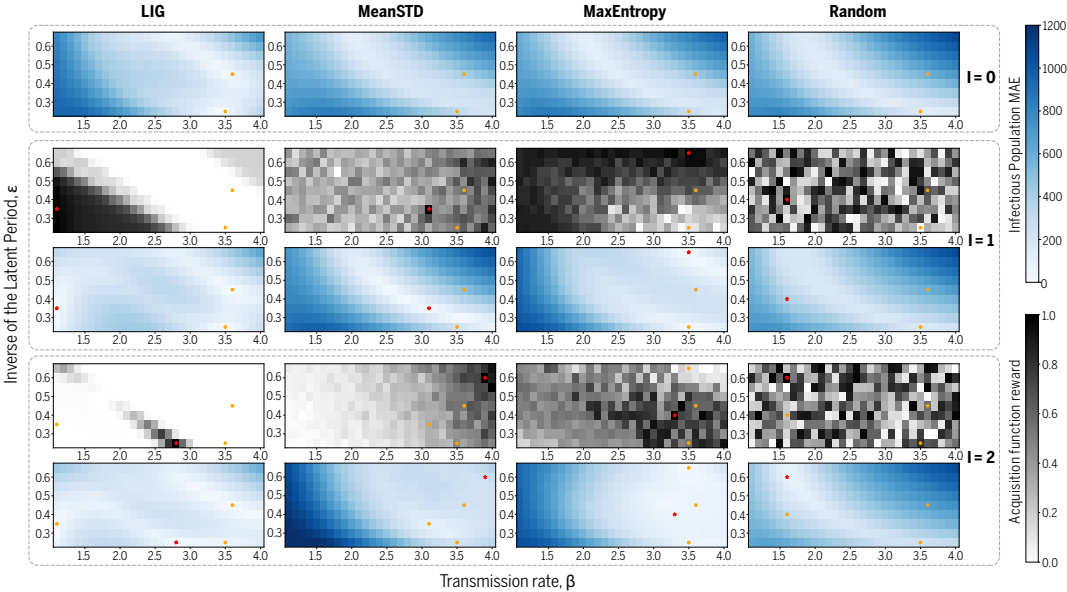


Figure 4: NP Acquisition functions behavior visualization in 2D. For each iteration, top row represents the current MAE mesh in infectious population for all  $(\beta, \varepsilon)$  candidates. Bottom row is the score evaluated by the acquisition function. Yellow dots are existing parameters. Red stars are the newly selected parameters.

**Uncertainty Quantification.** Figure 3 compares the NP and GP performance on one scenario in the held-out test set. It shows the ground truth and the predicted number of infectious population for the first 50 days. We also include the confidence intervals (CI) with 5 standard deviations for ground truth and NP predictions and 1 standard deviation for GP predictions. We observe that NP fits the simulation dynamics better than GP for mean prediction. Moreover, NP has closer CIs to the truth, reflecting the simulator’s intrinsic uncertainty. GP shows larger CIs which represent the model’s own uncertainty. Note that NP is much more flexible than GP and can scale easily to high-dimensional data.

**Acquisition Functions.** We compare different acquisition functions in Bayesian active learning. We select 2 scenarios as the initial training dataset. Then we continue augmenting the training set using

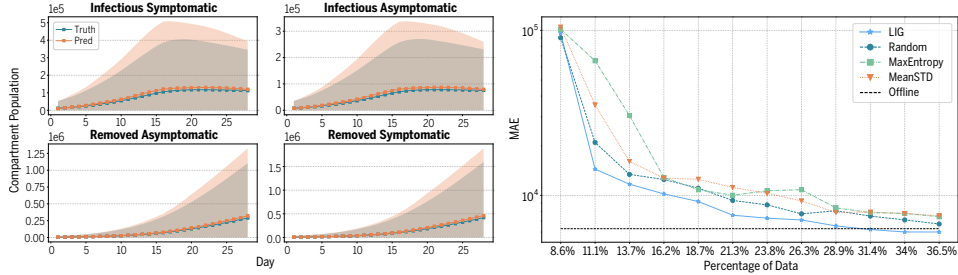


Figure 5: Left: INP predictions for the number of individuals in Infectious and Removed compartments in LEAM-US model. Right: acquisition function comparisons on the LEAM-US simulator. MAE loss versus the percentage of samples for INP during Bayesian active learning.

the remaining scenarios until the validation loss converges. GP is not a latent variable model, hence LIG does not apply. We report the average test mean absolute error (MAE) over three random runs. The candidate set for training includes 270 scenarios. We set aside 15 scenarios each for validation and test. The scenarios are evenly distributed in the search space.

Figure 3 shows the testing MAE versus the percentage of data used for training for NP and GP with different acquisition functions. None of the GP methods converge after selecting 4.07% of the data for training while NP methods converge much faster. Our proposed acquisition function LIG is more sample efficient in NP. It takes only 4.07% of the data to converge and reach the NP offline performance, which uses the entire training set for training.

**Exploration Exploitation Trade-off.** To understand the large performance gap for LIG vs. baselines on NP, we visualize the values of test MAE and the reward from acquisition functions for each Bayesian active learning iteration, shown in Figure 4. For Mean STD and Maximum Entropy, both tend to exploit the region with large transmission rate for the first 2 iterations. Including these scenarios makes the training set unbalanced. The MAE in the region with small transmission rate become worse after 2 iterations. Meanwhile, Random is doing pure exploration. The improvement of MAE performance is not apparent after 2 iterations. Our proposed LIG is able to reach a balance by exploiting the uncertainty in the latent process and encouraging exploration. Hence, with a small number of iterations ( $I = 2$ ), it has already selected “informative scenarios” in the search space.

#### 4.2 LOCAL EPIDEMIC AND MOBILITY MODEL (LEAM-US)

As a real-world experiment, we use a large-scale, spatiotemporal, age-structured epidemic model, LEAM-US, as our stochastic simulator. Our goal is to use INP to mimic the internal mechanism of this simulator to accelerate disease modeling and scenario creation.

**LEAM-US Simulator.** The Local Epidemic and Mobility model (LEAM-US) is a stochastic, spatial, age-structured epidemic model based on a metapopulation approach which divides the US in more than 3,100 subpopulations, each one corresponding to a each US county or statistically equivalent entity. Population size and county-specific age distributions reflect Census’ annual resident population estimates for year 2019. We consider individuals divided into 10 age groups. Contact mixing patterns are age-dependent and state specific and modeled considering contact matrices that describe the interaction of individuals in different social settings (Mistry et al., 2021). LEAM-US integrates a human mobility layer, represented as a network, using both short-range (i.e., commuting) and long-range (i.e., flights) mobility data, see more details in Appendix B.2.

**Dataset.** We separate data in California monthly to predict the 28 days’ sequence from the 2nd to the 29th day of each month from March to December. Each  $\theta$  includes the county-level parameters of LEAM-US and state level incidence and prevalence compartments. The total number of dimensions in  $\theta$  is 16,264, see details in Appendix B.2. Overall, there are 315 scenarios in the search space, corresponding to 315 different  $\theta$  with total 16,254 samples. We split 78% of the data as the candidate train set, and 11% for validation and test. We use 8.6% data for initial training. We first perform offline learning to show that the STNP model can accurately learn the LEAM-US dynamics. For Bayesian active learning, we use the candidate train set as the search space.

**Model Implementation.** We instantiate a STNP model to mimic an epidemic simulator that has  $\theta$  at both county and state level and  $x_t$  at the state level. We use county-level parameter  $\theta$  together



with a county-to-county mobility graph  $A$  in California as input. We use the DCGRU layer (Li et al., 2017) to encode the mobility graph in a GRU. We use a linear layer to map the county-level output to hidden features at the state level. For both the state-level encoder and decoder, we use multi-layer GRUs. For each epoch, we randomly select 20% samples as context sequence.

**STNP Predictions.** Figure 5 left visualize the STNP predictions in four key compartments of a typical scenario with  $R_0 = 3.1$  from March 2nd to March 29th. The confidence interval is plotted with 2 standard deviations. We can see that both the mean and confidence interval of STNP predictions match the truth well. These results demonstrate the promise that the generative STNP model can serve as a deep surrogate model for the LEAM-US simulator.

**Acquisition Functions.** We compare different acquisition functions with STNP. All methods start with the same initial train set. Then we continue adding 8 scenarios to the training set until the validation loss converges. We measure the average performance over three random runs and report the test MAE. Figure 5 right shows the log scale MAE versus the percentage of samples included for training, which scales linearly in computation. It shows our proposed LIG always has the best MAE performance until the convergence. Moreover, after using 31.4% of data from the candidate train set, LIG already outperforms the offline model.

**Ablation Study on NP Models.** We perform ablation study on our proposed NP model, STNP. Compared with vanilla NP (Garnelo et al., 2018) and SNP (Singh et al., 2019), the key innovation of model is temporal latent process, modeled by state level encoder and decoder as in Figure 6. To ensure fair comparison, we modified NP by adding the convolutional layers with diffusion convolution (Li et al., 2018) to embed the graphs. Similarly, we modified SNP by replacing the convolutional layers with diffusion convolution. Figure 2 shows the corresponding graphical models. Table 1 shows the testing MAE of different NP models trained in an offline fashion.

Our STNP significantly improves the performance and can accurately learn the simulator dynamics.

**Batch Active Learning with LIG.** In the case where we add more than one scenario of data per iteration, LIG allows us to calculate the corresponding reward for a batch of scenarios due to the mean aggregation layer in the NP encoder. In our LEAM-US experiment, we add 8 scenarios per iteration (batch size 8). As an ablation study, we compared 4 different batch sizes, see Appendix Figure 8. We observe that larger batch leads to better performance. This strategy combines LIG and random search, where it first randomly selects a set of groups and then uses LIG to measure the reward for the data within each group. Random search encourages exploration whereas LIG exploits the highest reward.

## 5 CONCLUSION & LIMITATION

We present a new way to seamlessly interact with existing stochastic simulators and accelerate simulation. Our Neural Interactive Process (INP) includes a novel spatiotemporal neural process model to approximate the underlying simulation dynamics. It infers the latent process which describes the intrinsic uncertainty of the simulator. We exploit this uncertainty and propose LIG as an acquisition function in Bayesian active learning. We perform theoretical analysis and demonstrate that our approach reduces sample complexity compared with random sampling in high dimension. Using a complex epidemic simulator, we demonstrate our method provides a faster and more flexible alternative to Gaussian process for policy optimization and scenario creation.

One limitation of our current work is the optimization step. Currently, we use a brute force search mechanism to select the parameters according to the value of an acquisition function. However, this search will become expensive in high dimension. We plan to leverage Monte Carlo sampling and Bayesian optimization techniques to parameterize the reward function. Then we would also be able to directly optimize for the target parameters with auto-differentiation.

Table 1: NP model comparison in LEAM simulator, population divided by 1000.

Model	MAE
NP (Garnelo et al., 2018)	$24.231 \pm 5.884$
SNP (Singh et al., 2019)	$21.781 \pm 0.825$
STNP (Ours)	$6.291 \pm 0.848$

## REPRODUCIBILITY STATEMENT

The implementation code is included in the supplementary material. The readme file includes the corresponding instructions. The full proof of the Theorem 1 can be found in Appendix A.2.

## REFERENCES

- Sercan Arik, Chun-Liang Li, Jinsung Yoon, Rajarishi Sinha, Arkady Epshteyn, Long Le, Vikas Menon, Shashank Singh, Leyou Zhang, Martin Nikoltchev, et al. Interpretable sequence learning for covid-19 forecasting. *Advances in Neural Information Processing Systems*, 33, 2020.
- Søren Asmussen and Peter W Glynn. *Stochastic simulation: algorithms and analysis*, volume 57. Springer Science & Business Media, 2007.
- Duygu Balcan, Vittoria Colizza, Bruno Gonçalves, Hao Hu, José J Ramasco, and Alessandro Vespignani. Multiscale mobility networks and the spatial spreading of infectious diseases. *Proceedings of the National Academy of Sciences*, 106(51):21484–21489, 2009.
- Duygu Balcan, Bruno Gonçalves, Hao Hu, José J Ramasco, Vittoria Colizza, and Alessandro Vespignani. Modeling the spatial spread of infectious diseases: The global epidemic and mobility computational model. *Journal of computational science*, 1(3):132–145, 2010.
- Kathryn Chaloner and Isabella Verdinelli. Bayesian experimental design: A review. *Statistical Science*, pp. 273–304, 1995.
- Ricky TQ Chen, Yulia Rubanova, Jesse Bettencourt, and David Duvenaud. Neural ordinary differential equations. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, pp. 6572–6583, 2018.
- Zizhong Chen and Jack J. Dongarra. Condition numbers of gaussian random matrices. *SIAM Journal on Matrix Analysis and Applications*, 27(3):603–620, 2005.
- Matteo Chinazzi, Jessica T Davis, Marco Ajelli, Corrado Gioannini, Maria Litvinova, Stefano Merler, Ana Pastore y Piontti, Kunpeng Mu, Luca Rossi, Kaiyuan Sun, et al. The effect of travel restrictions on the spread of the 2019 novel coronavirus (covid-19) outbreak. *Science*, 2020.
- David A Cohn, Zoubin Ghahramani, and Michael I Jordan. Active learning with statistical models. *Journal of artificial intelligence research*, 4:129–145, 1996.
- Estee Y Cramer, Velma K Lopez, Jarad Niemi, Glover E George, Jeffrey C Cegan, Ian D Dettwiller, William P England, Matthew W Farthing, Robert H Hunter, Brandon Lafferty, et al. Evaluation of individual and ensemble probabilistic forecasts of covid-19 mortality in the us. *medRxiv*, 2021.
- Jessica T Davis, Matteo Chinazzi, Nicola Perra, Kunpeng Mu, Ana Pastore y Piontti, Marco Ajelli, Natalie E Dean, Corrado Gioannini, Maria Litvinova, Stefano Merler, Luca Rossi, Kaiyuan Sun, Xinyue Xiong, M. Elizabeth Halloran, Ira M Longini, Cécile Viboud, and Alessandro Vespignani. Estimating the establishment of local transmission and the cryptic phase of the covid-19 pandemic in the usa. *medRxiv*, 2020.
- Simon Du, Jason Lee, Haochuan Li, Liwei Wang, and Xiyu Zhai. Gradient descent finds global minima of deep neural networks. In *Proceedings of the 36th International Conference on Machine Learning (ICML)*, pp. 1675–1685, 2019.
- Conor Durkan, Iain Murray, and George Papamakarios. On contrastive learning for likelihood-free inference. In *International Conference on Machine Learning*, pp. 2771–2781. PMLR, 2020.
- A Foster, M Jankowiak, E Bingham, P Horsfall, YW Tee, T Rainforth, and N Goodman. Variational bayesian optimal experimental design. Conference on Neural Information Processing Systems, 2019.
- Adam Foster, Desi R Ivanova, Ilyas Malik, and Tom Rainforth. Deep adaptive design: Amortizing sequential bayesian experimental design. *Proceedings of the 38th International Conference on Machine Learning (ICML)*, 2021.
- Peter I Frazier. A tutorial on bayesian optimization. *arXiv preprint arXiv:1807.02811*, 2018.
- Yarin Gal and Zoubin Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *international conference on machine learning*, pp. 1050–1059. PMLR, 2016.

- Yarin Gal, Riashat Islam, and Zoubin Ghahramani. Deep bayesian active learning with image data. In *International Conference on Machine Learning*, pp. 1183–1192. PMLR, 2017.
- Jacob R Gardner, Geoff Pleiss, David Bindel, Kilian Q Weinberger, and Andrew Gordon Wilson. Gpytorch: Blackbox matrix-matrix gaussian process inference with gpu acceleration. *arXiv preprint arXiv:1809.11165*, 2018.
- Marta Garnelo, Jonathan Schwarz, Dan Rosenbaum, Fabio Viola, Danilo J Rezende, SM Eslami, and Yee Whye Teh. Neural processes. *arXiv preprint arXiv:1807.01622*, 2018.
- Daniel T Gillespie. Stochastic simulation of chemical kinetics. *Annu. Rev. Phys. Chem.*, 58:35–55, 2007.
- Michael U Gutmann, Jukka Corander, et al. Bayesian optimization for likelihood-free inference of simulator-based statistical models. *Journal of Machine Learning Research*, 2016.
- Friedrich Götze and Alexander Tikhomirov. Rate of convergence in probability to the Marchenko-Pastur law. *Bernoulli*, 10(3):503 – 548, 2004.
- Philipp Holl, Nils Thuerey, and Vladlen Koltun. Learning to control pdes with differentiable physics. In *International Conference on Learning Representations*, 2019.
- Neil Houlsby, Ferenc Huszár, Zoubin Ghahramani, and Máté Lengyel. Bayesian active learning for classification and preference learning. *arXiv preprint arXiv:1112.5745*, 2011.
- IATA, International Air Transport Association, 2021. URL <https://www.iata.org/>. <https://www.iata.org/>.
- Edwin T Jaynes. Information theory and statistical mechanics. *Physical review*, 106(4):620, 1957.
- Hyunjik Kim, Andriy Mnih, Jonathan Schwarz, Marta Garnelo, Ali Eslami, Dan Rosenbaum, Oriol Vinyals, and Yee Whye Teh. Attentive neural processes. *International Conference on Learning Representation*, 2019.
- Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- Steven Kleinegese and Michael U Gutmann. Bayesian experimental design for implicit models by mutual information neural estimation. In *International Conference on Machine Learning*, pp. 5316–5326. PMLR, 2020.
- Damien Lamberton and Bernard Lapeyre. *Introduction to stochastic calculus applied to finance*. CRC press, 2007.
- Yaguang Li, Rose Yu, Cyrus Shahabi, and Yan Liu. Diffusion convolutional recurrent neural network: Data-driven traffic forecasting. *arXiv preprint arXiv:1707.01926*, 2017.
- Yaguang Li, Rose Yu, Cyrus Shahabi, and Yan Liu. Diffusion convolutional recurrent neural network: Data-driven traffic forecasting. In *International Conference on Learning Representations (ICLR)*, 2018.
- Christos Louizos, Xiahao Shi, Klamer Schutte, and Max Welling. The functional neural process. *Advances in Neural Information Processing Systems*, 2019.
- Jose Lourenco, Robert Paton, Mahan Ghafari, Moritz Kraemer, Craig Thompson, Peter Simmonds, Paul Klenerman, and Sunetra Gupta. Fundamental principles of epidemic spread highlight the immediate need for large-scale serological surveys to assess the stage of the sars-cov-2 epidemic. *MedRxiv*, 2020.
- Jan-Matthis Lueckmann, Giacomo Bassetto, Theofanis Karaletsos, and Jakob H Macke. Likelihood-free inference with emulator networks. In *Symposium on Advances in Approximate Bayesian Inference*, pp. 32–53. PMLR, 2019.

- Edward Meeds and Max Welling. Gps-abc: Gaussian process surrogate approximate bayesian computation. In *Proceedings of the Thirtieth Conference on Uncertainty in Artificial Intelligence*, pp. 593–602, 2014.
- Song Mei and Andrea Montanari. The generalization error of random features regression: Precise asymptotics and double descent curve. arXiv: 1908.05355, 2019.
- Dina Mistry, Maria Litvinova, Ana Pastore y Piontti, Matteo Chinazzi, Laura Fumanelli, Marcelo FC Gomes, Syed A Haque, Quan-Hui Liu, Kunpeng Mu, Xinyue Xiong, et al. Inferring high-resolution human mixing patterns for disease modeling. *Nature communications*, 12(1):1–12, 2021.
- Andreas Munk, Adam Ścibior, Atılım Güneş Baydin, Andrew Stewart, Goran Fernlund, Anoush Poursartip, and Frank Wood. Deep probabilistic surrogate networks for universal simulator approximation. *arXiv preprint arXiv:1910.11950*, 2019.
- Jay I Myung, Daniel R Cavagnaro, and Mark A Pitt. A tutorial on adaptive design optimization. *Journal of mathematical psychology*, 57(3-4):53–67, 2013.
- OAG, Aviation Worlwide Limited, 2021. URL <http://www.oag.com/>. <http://www.oag.com/>.
- George Papamakarios, David Sterratt, and Iain Murray. Sequential neural likelihood: Fast likelihood-free inference with autoregressive flows. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pp. 837–848. PMLR, 2019.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. In *Advances in neural information processing systems*, pp. 8026–8037, 2019.
- Zhaozhi Qian, Ahmed M Alaa, and Mihaela van der Schaar. When and how to lift the lockdown? global covid-19 scenario analysis and policy assessment using compartmental gaussian processes. *Advances in Neural Information Processing Systems*, 33, 2020.
- Syama Sundar Rangapuram, Matthias W Seeger, Jan Gasthaus, Lorenzo Stella, Yuyang Wang, and Tim Januschowski. Deep state space models for time series forecasting. *Advances in neural information processing systems*, 31:7785–7794, 2018.
- Stephan Rasp, Michael S Pritchard, and Pierre Gentine. Deep learning to represent subgrid processes in climate models. *Proceedings of the National Academy of Sciences*, 115(39):9684–9689, 2018.
- Brian D Ripley. *Stochastic simulation*, volume 316. John Wiley & Sons, 2009.
- Alvaro Sanchez-Gonzalez, Jonathan Godwin, Tobias Pfaff, Rex Ying, Jure Leskovec, and Peter Battaglia. Learning to simulate complex physics with graph networks. In *International Conference on Machine Learning*, pp. 8459–8468. PMLR, 2020.
- Bobak Shahriari, Kevin Swersky, Ziyu Wang, Ryan P Adams, and Nando De Freitas. Taking the human out of the loop: A review of bayesian optimization. *Proceedings of the IEEE*, 104(1): 148–175, 2015.
- Aditya Siddhant and Zachary C Lipton. Deep bayesian active learning for natural language processing: Results of a large-scale empirical study. *arXiv preprint arXiv:1808.05697*, 2018.
- Gautam Singh, Jaesik Yoon, Youngsung Son, and Sungjin Ahn. Sequential neural processes. *Advances in Neural Information Processing Systems*, 32:10254–10264, 2019.
- Michele Tizzoni, Paolo Bajardi, Chiara Poletto, José J Ramasco, Duygu Balcan, Bruno Gonçalves, Nicola Perra, Vittoria Colizza, and Alessandro Vespignani. Real-time numerical forecast of global epidemic spreading: case study of 2009 a/h1n1pdm. *BMC medicine*, 10(1):165, 2012.
- Benjamin T Vincent and Tom Rainforth. The darc toolbox: automated, flexible, and efficient delayed and risky choice experiments using bayesian adaptive design. *PsyArXiv. October*, 20, 2017.

- Rui Wang, Karthik Kashinath, Mustafa Mustafa, Adrian Albert, and Rose Yu. Towards physics-informed deep learning for turbulent flow prediction. In *Proceedings of the 26th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 2020, 2020.
- Frank Wood, Andrew Warrington, Saeid Naderiparizi, Christian Weilbach, Vaden Masrani, William Harvey, Adam Scibior, Boyan Beronov, John Grefenstette, Duncan Campbell, et al. Planning as inference in epidemiological models. *arXiv preprint arXiv:2003.13221*, 2020.
- Qian Zhang, Kaiyuan Sun, Matteo Chinazzi, Ana Pastore y Piontti, Natalie E Dean, Diana Patricia Rojas, Stefano Merler, Dina Mistry, Piero Poletti, Luca Rossi, et al. Spread of Zika virus in the Americas. *Proceedings of the National Academy of Sciences*, 114(22):E4334–E4343, 2017.
- Christoph Zimmer, Mona Meister, and Duy Nguyen-Tuong. Safe active learning for time-series modeling with gaussian processes. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, pp. 2735–2744, 2018.



## A THEORETICAL ANALYSIS

### A.1 LATENT INFORMATION GAIN

**Proposition 1.** *The expected information gain (EIG) for neural process is equivalent to the KL divergence between the prior and posterior in the latent process, that is*

$$\text{EIG}(\hat{x}, \theta) := \mathbb{E}[H(\hat{x}) - H(\hat{x}|z, \theta)] = \mathbb{E}_{p(\hat{x}|\theta)} [\text{KL}(p(z|\hat{x}, \theta)||p(z))] \quad (4)$$

*Proof of Proposition 1.* The information gained in the latent process  $z$ , by selecting the parameter  $\theta$  and generate  $\hat{x}$  is the reduction in entropy from the prior to the posterior  $\text{IG}(\theta) = H(\hat{x}) - H(\hat{x}|z, \theta)$ . Take the expectation of  $\text{IG}(\hat{x}, \theta)$  under the marginal distribution, we obtain from the conditional independence of  $z$  and  $\theta$  that

$$\begin{aligned} \mathbb{E}_{p(\hat{x}|\theta)} [\text{KL}(p(z|\hat{x}, \theta)||p(z))] &= \mathbb{E}_{p(\hat{x}, z|\theta)} \left[ \log \frac{p(z|\hat{x}, \theta)}{p(z)} \right] \\ &= \mathbb{E}_{p(\hat{x}, z|\theta)} \left[ \log \frac{p(z|\hat{x}, \theta)}{p(z|\theta)} \right] \\ &= \mathbb{E}_{p(\hat{x}, z|\theta)} [\log p(z, \hat{x}, \theta) - \log p(\hat{x}, \theta) - \log p(z, \theta) + \log p(\theta)] \\ &= \mathbb{E}_{p(\hat{x}, z|\theta)} [\log p(\hat{x}|z, \theta) - \log p(\hat{x}|\theta)] \\ &= \mathbb{E}_{p(z)} \left[ \mathbb{E}_{p(\hat{x}|z, \theta)} [\log p(\hat{x}|z, \theta)] - \mathbb{E}_{p(\hat{x}|\theta)} [\log p(\hat{x}|\theta)] \right] \\ &= \mathbb{E}_{p(z)} [H(\hat{x}|\theta) - H(\hat{x}|z, \theta)] \\ &= \text{EIG}(\hat{x}, \theta). \end{aligned}$$

□

### A.2 SAMPLE EFFICIENCY OF ACTIVE LEARNING

From the main text we know that in each round, the output random variable

$$X = \langle \Psi(\theta), z^* \rangle + \epsilon. \quad (5)$$

We further assume that the random noise  $\epsilon$  is mean zero and  $\sigma$ -subGaussian.

Using this information, we treat  $z$  as an unknown parameter and define a likelihood function so that  $p(X|z; \theta)$  has good coverage over the observations:

$$p(X_k|z; \theta_k) \propto \exp \left( -\frac{1}{2\sigma^2} (X_k - \langle \Psi(\theta_k), z \rangle)^2 \right).$$

Let the prior distribution over  $z$  be  $p(z|\theta_k) = p(z) \propto \exp \left( -\frac{m}{2\sigma^2} \|z\|^2 \right)$ . Here we use  $k$  instead of  $(i)$  in the Algorithm 1 to represent the number of iterations. We can form a posterior over  $z$  in the  $k$ -th round:

$$p(z|X_1, \theta_1, \dots, X_k, \theta_k) \propto \exp \left( -\frac{m}{2\sigma^2} \|z\|^2 - \frac{1}{2\sigma^2} \sum_{s=1}^k (X_s - \langle \Psi(\theta_s), z \rangle)^2 \right).$$

Focusing on the random variable  $z \sim p(\cdot|X_1, \theta_1, \dots, X_k, \theta_k)$ , the estimate of the hidden variable, we can express it at  $k$ -th round as:

$$z_k = \hat{z}_k + \sigma V_k^{-1} \eta_k, \quad (6)$$

where  $\hat{z}_k = V_k^{-1} \sum_{s=1}^k X_s \Psi(\theta_s)$ ,  $V_k = m\mathbf{I} + \sum_{s=1}^k \Psi(\theta_s) \Psi(\theta_s)^\top$ , and  $\eta_k$  is a standard normal random variable.

We can either choose action  $\theta$  randomly or greedily. A random choice of  $\theta$  corresponds to taking

$$\theta_k \sim \mathcal{N}(0, \mathbf{I}), \quad (7)$$

A greedy procedure is to choose action  $\theta_k$  in the  $k$ -th round to optimize  $\text{KL}(p(z|\hat{x}, \theta) \| p(z)) = \mathbb{E}_{p(z|\hat{x}, \theta)} \left( \log \frac{p(z|\hat{x}, \theta)}{p(z)} \right)$ , where we denote the estimated output variable  $\hat{x}$  given  $\theta$  and  $z$  as  $\hat{x} = \langle \Psi(\theta), z \rangle$ . This optimization procedure is equivalent to maximizing the variance of the prediction:

$$\theta_k = \arg \max_{\theta \in \mathbb{R}^d} \mathbb{E}_{z \sim p(\cdot | X_1, \theta_1, \dots, X_{k-1}, \theta_{k-1})} \left[ \left( \langle \Psi(\theta), z \rangle - \mathbb{E}_{z \sim p(\cdot | X_1, \theta_1, \dots, X_{k-1}, \theta_{k-1})} \langle \Psi(\theta), z \rangle \right)^2 \right]. \quad (8)$$

For both approaches, we assume that the features  $\Psi(\theta)$  are normalized.

We compare the statistical risk of this approach with the random sampling approach.

Assume that the features are normalized, so that for all  $\theta \in \mathbb{R}^d$ ,  $\Psi(\theta) \in \mathbb{S}^{d-1}$ . Define a matrix  $\mathbf{A}_k \in \mathbb{R}^{d \times k}$  containing all the column vectors  $\{\Psi(\theta_1), \dots, \Psi(\theta_k)\}$ . We can then express the estimation error in the following lemma.

**Lemma 1.** *The estimation error  $\|\hat{z}_k - z^*\|_2$  can be bounded as follow.*

$$\begin{aligned} \|\hat{z}_k - z^*\|_2 &\leq m \left( m + \sigma_{\min}(\mathbf{A}_k \mathbf{A}_k^T) \right)^{-1} \cdot \|z^*\|_2 \\ &\quad + \min \left\{ 1/(2\sqrt{m}), 1/\left( \sqrt{\sigma_{\min}(\mathbf{A}_k \mathbf{A}_k^T)} + \frac{m}{\sqrt{\sigma_{\min}(\mathbf{A}_k \mathbf{A}_k^T)}} \right) \right\} \cdot \sigma\sqrt{d}. \end{aligned}$$

We now analyze random sampling of  $\theta$  versus greedy search for  $\theta$ .

If the feature map  $\Psi(\cdot) = \text{id}$ , then from random matrix theory, we know that for  $\theta$  randomly sampled from a normal distribution and normalized to  $\|\theta\| = 1$ ,  $\sigma_{\min}(\frac{1}{k}\mathbf{A}_k \mathbf{A}_k^T)$  will converge to  $\left( \sqrt{1/k} - \sqrt{1/d} \right)^2$  for large  $k$ , which is of order  $\Omega(1/d)$ . This will lead to an appealing risk bound for  $\|\hat{z}_k - z^*\|_2$  on the order of  $\mathcal{O}(d/\sqrt{k})$ .

However, in high dimension, this feature map is often far from identity. In the proof of Theorem 1 below, we demonstrate that even when  $\Psi(\cdot)$  is simply a linear random feature map, with i.i.d. normal entries, random exploration in  $\theta$  can lead to deteriorated error bound. This setting is motivated by the analyses in wide neural networks, where the features learned from gradient descent are close to those generated from random initialization Du et al. (2019); Mei & Montanari (2019).

**Theorem 1** (Formal statement). *Assume that the noise  $\epsilon$  in equation 5 is  $\sigma$ -subGaussian.*

*For a normalized linear random feature map  $\Psi(\cdot)$ , greedily optimizing the KL divergence,  $\text{KL}(p(z|\hat{x}, \theta) \| p(z))$  (or equivalently the variance of the posterior predictive distribution defined in equation equation 8) in search of  $\theta$  will lead to an error  $\|\hat{z}_k - z^*\|_2 = \mathcal{O}(\sigma d/\sqrt{k})$  with high probability.*

*On the other hand, random sampling of  $\theta$  following equation 7 will lead to  $\|\hat{z}_k - z^*\|_2 = \mathcal{O}(\sigma d^2/\sqrt{k})$  with high probability.*

*Proof of Theorem 1.* For a linear random feature map, we can express  $\Psi(\theta) = \Psi\theta$ , where entries in  $\Psi \in \mathbb{R}^{d \times d}$  are i.i.d. normal. The entries of  $\Psi\theta$  are then normalized.

- For random exploration of  $\theta$ , the matrix containing the feature vectors becomes  $\mathbf{A}_k = \Psi\Theta_k$ , where matrix  $\Theta_k \in \mathbb{R}^{d \times k}$  collects all the  $k$  column vectors of  $\{\theta_1, \dots, \theta_k\}$ . Then  $\mathbf{A}_k \mathbf{A}_k^T = \Psi\Theta_k \Theta_k^T \Psi^T$ . From random matrix theory, we know that the condition number of  $\Psi$  is equal to  $d$  with high probability Chen & Dongarra (2005). Hence for normalized  $\Psi$  and  $\theta$ ,  $\sigma_{\min}(\Psi\Theta_k \Theta_k^T \Psi^T) \geq \sigma_{\min}^2(\Psi) \sigma_{\min}(\Theta_k \Theta_k^T) = \frac{1}{d^2} \sigma_{\min}(\Theta_k \Theta_k^T)$ . The inequality holds because the smallest singular value is the inverse of the norm of the inverse matrix.

We then use the fact from random matrix theory that for normalized random  $\theta$ , the asymptotic distribution of the eigenvalues of  $\frac{1}{k}\Theta_k\Theta_k^T$  follow the (scaled) Marchenko–Pastur distribution, which is supported on  $\lambda \in \left[ \left( \sqrt{1/k} - \sqrt{1/d} \right)^2, \left( \sqrt{1/k} + \sqrt{1/d} \right)^2 \right]$ , where the  $1/d$  scaling comes from the fact that  $\theta$  is normalized Götze & Tikhomirov (2004). Hence for large  $k$ ,  $\sigma_{\min}(\Theta_k\Theta_k^T) \geq \left(1 - \sqrt{k/d}\right)^2$  with high probability. This combined with the previous paragraph yields that for the random feature model,

$$\sigma_{\min}(A_k A_k^T) = \Omega\left(\frac{1}{d^2} \left(1 - \sqrt{k/d}\right)^2\right)$$

with high probability. Plugging this result into Lemma 1, we obtain that the error  $\|\hat{z}_k - z^*\|_2$  for random exploration in the space of  $\theta$  is of order  $\mathcal{O}(d^2/\sqrt{k})$ .

- We then analyze the error associated with greedy maximization of the posterior predictive variance. We first note that the variance of the posterior predictive distribution in equation 8 can be expressed as follows using equation 6:

$$\mathbb{E} \left[ \left( \langle \Psi(\theta), z \rangle - \mathbb{E} \langle \Psi(\theta), z \rangle \right)^2 \right] = \sigma^2 \mathbb{E} \left[ \left( \langle \Psi(\theta), V_{k-1}^{-1} \eta_k \rangle \right)^2 \right] = \sigma^2 \Psi(\theta)^T V_{k-1}^{-2} \Psi(\theta), \quad (9)$$

where the expectations are with respect to  $z \sim p(\cdot | X_1, \theta_1, \dots, X_{k-1}, \theta_{k-1})$ .

We perform a singular value decomposition  $A_k = U_k \Lambda_k W_k$ . Then  $\sum_{s=1}^k \Psi(\theta_s) \Psi(\theta_s)^T = A_k A_k^T = U_k \Lambda_k \Lambda_k^T U_k^T$ , and that  $V_{k-1}^{-2} = (m\mathbf{I} + A_{k-1} A_{k-1}^T)^{-2} = U_{k-1} (m\mathbf{I} + \Lambda_{k-1} \Lambda_{k-1}^T)^{-2} U_{k-1}^T$ . Via this formulation, we see that maximizing  $\Psi(\theta)^T V_{k-1}^{-2} \Psi(\theta)$  in equation 9 to choose  $\theta_k$  is equivalent to choosing  $\Psi(\theta_k) = (U_{k-1})_{(c,l)}^T$ , where  $l = \arg \min_{i \in \{1, \dots, d\}} (\Lambda_{k-1} \Lambda_{k-1}^T)_{(i,i)}$ . In words, when we use greedy method and maximize the variance of the prediction, it corresponds to taking  $\Psi(\theta_k)$  in the direction of the smallest eigenvector of  $V_{k-1}$ .

Since every  $\Psi(\theta)$  is normalized and we initialize uniformly:  $V_0 = m\mathbf{I}$ , the process is equivalent to scanning the orthogonal spaces of normalized vectors in  $\mathbb{R}^d$  for  $\lfloor k/d \rfloor$  times. For large  $k$ , entries in  $\Lambda_k \Lambda_k^T$  are approximately uniform and are all larger than or equal to  $\lfloor k/d \rfloor$ . Then  $\sigma_{\min}(A_k A_k^T) = \Omega(k/d)$ . Plugging into the bound of Lemma 1, we obtain that

$$\|\hat{z}_k - z^*\|_2 = \mathcal{O}\left(\frac{\sigma d}{\sqrt{k}}\right).$$

□

*Proof of Lemma 1.* We first express the estimate  $\hat{z}_k$  as follows.

$$\hat{z}_k = V_k^{-1} \sum_{s=1}^k X_s \Psi(\theta_s) = V_k^{-1} \sum_{s=1}^k \Psi(\theta_s) \Psi(\theta_s)^T z^* + V_k^{-1} \sum_{s=1}^k \epsilon_s \Psi(\theta_s).$$

Then

$$\begin{aligned} \|\hat{z}_k - z^*\|_2 &= \left\| \left( V_k^{-1} \sum_{s=1}^k \Psi(\theta_s) \Psi(\theta_s)^T - \mathbf{I} \right) z^* + V_k^{-1} \sum_{s=1}^k \epsilon_s \Psi(\theta_s) \right\|_2 \\ &\leq \underbrace{\left\| \left( V_k^{-1} \sum_{s=1}^k \Psi(\theta_s) \Psi(\theta_s)^T - \mathbf{I} \right) z^* \right\|_2}_{T_1} + \underbrace{\left\| V_k^{-1} \sum_{s=1}^k \epsilon_s \Psi(\theta_s) \right\|_2}_{T_2}. \end{aligned}$$

Define a matrix  $A_k \in \mathbb{R}^{d \times k}$  containing all the column vectors  $\{\Psi(\theta_1), \dots, \Psi(\theta_k)\}$  and perform a singular value decomposition  $A_k = U_k \Lambda_k W_k$ . Then  $\sum_{s=1}^k \Psi(\theta_s) \Psi(\theta_s)^T = A_k A_k^T = U_k \Lambda_k \Lambda_k^T U_k^T$ ,

and  $V_k = m\mathbf{I} + \Lambda_k \Lambda_k^T$ , We further define vector  $e_k \in \mathbb{R}^s$  where  $(e_k)_s = \epsilon_s$ . We use this definition to simplify the two terms further.

For term  $T_1$ ,

$$\begin{aligned} \left\| \left( V_k^{-1} \sum_{s=1}^k \Psi(\theta_s) \Psi(\theta_s)^T - \mathbf{I} \right) z^* \right\|_2 &= m \|V_k^{-1} z^*\|_2 \\ &\leq m \|V_k^{-1}\|_2 \cdot \|z^*\|_2 \\ &= m (m + \sigma_{\min}(\Lambda_k \Lambda_k^T))^{-1} \cdot \|z^*\|_2. \end{aligned}$$

For term  $T_2$ , we define a diagonal matrix  $\bar{\Lambda}_k \in \mathbb{R}^{k \times k}$  which satisfies  $(\bar{\Lambda}_k)_{i,i} = 1$  if  $i \leq d$  and  $(\bar{\Lambda}_k)_{i,i} = 0$  if  $i > d$ , when  $k > d$ . The following bound on  $T_2$  can be achieved.

$$\begin{aligned} \left\| V_k^{-1} \sum_{s=1}^k \epsilon_s \Psi(\theta_s) \right\|_2 &= \|V_k^{-1} \Lambda_k e_k\|_2 \\ &= \left\| U_k (\Lambda_k \Lambda_k^T + m\mathbf{I})^{-1} U_k^T U_k \Lambda_k \bar{\Lambda}_k W_k e_k \right\|_2 \\ &\leq \left\| U_k (\Lambda_k \Lambda_k^T + m\mathbf{I})^{-1} \Lambda_k \right\|_2 \cdot \|\bar{\Lambda}_k W_k e_k\|_2 \\ &= \left\| (\Lambda_k \Lambda_k^T + m\mathbf{I})^{-1} \Lambda_k \right\|_2 \cdot \|\bar{\Lambda}_k W_k e_k\|_2 \\ &\leq \min \left\{ 1/(2\sqrt{m}), 1/\left( \sqrt{\sigma_{\min}(\Lambda_k \Lambda_k^T)} + \frac{m}{\sqrt{\sigma_{\min}(\Lambda_k \Lambda_k^T)}} \right) \right\} \cdot \|\bar{\Lambda}_k W_k e_k\|_2. \end{aligned}$$

Assuming that noise  $\epsilon_s$  is  $\sigma$ -subGaussian, then so is  $W_k e_k$  since  $W_k$  is a unitary matrix. Multiplied by the diagonal matrix  $\bar{\Lambda}_k$  which has zero,  $\|\bar{\Lambda}_k W_k e_k\|_2 \leq \sigma\sqrt{d}$ . Therefore,

$$\left\| V_k^{-1} \sum_{s=1}^k \epsilon_s \Psi(\theta_s) \right\|_2 \leq \min \left\{ 1/(2\sqrt{m}), 1/\left( \sqrt{\sigma_{\min}(\Lambda_k \Lambda_k^T)} + \frac{m}{\sqrt{\sigma_{\min}(\Lambda_k \Lambda_k^T)}} \right) \right\} \cdot \sigma\sqrt{d}.$$

□

## B EXPERIMENT DETAILS

### B.1 SEIR SIMULATOR

Our SEIR simulator is a simple stochastic, discrete, chain-binomial compartmental model. In this model, susceptible individuals ( $S$ ) become exposed ( $E$ ) through interactions with infectious individuals ( $I$ ). Exposed individuals which are infected but not yet infectious transition to infectious compartment at a rate  $\varepsilon$  that is inversely proportional to the latent period of the disease. Lastly, infectious individuals transition to the removed compartment at a rate  $\mu$  which is inversely proportional to the infectious period. Removed individuals ( $R$ ) are assumed to be no longer infectious and they are to be considered either recovered or dead. All transitions are simulated by randomly drawn from a binomial distribution.

### B.2 LEAM-US MODEL

LEAM-US integrates a human mobility layer, represented as a network, using both short-range (i.e., commuting) and long-range (i.e., flights) mobility data. Commuting flows between counties are obtained from the 2011-2015 5-Year ACS Commuting Flows survey and properly adjusted to account for differences in population totals since the creation of the dataset. Instead, long-range air traveling flows are quantified using origin-destination daily passenger flows between airport pairs as reported by the Official Aviation Guide (OAG) and IATA databases (updated in 2021) (OAG,

Aviation Worldwide Limited, 2021; IATA, International Air Transport Association, 2021). In addition, flight probabilities are age and country specific.

The model is initialized using a multi-scale modeling approach that utilizes GLEAM, the Global and Epidemic Mobility model (Balcan et al., 2009; 2010; Tizzoni et al., 2012; Zhang et al., 2017; Chinazzi et al., 2020; Davis et al., 2020), to simulate a set of 500 different initial conditions for LEAM-US starting on February 16th, 2020. The disease dynamics are modeled using a classic SEIR-like model and initial conditions are determined using the Global and Epidemic Mobility model (Balcan et al., 2009; 2010; Tizzoni et al., 2012; Zhang et al., 2017) calibrated to realistically represent the evolution of the COVID-19 pandemic (Chinazzi et al., 2020; Davis et al., 2020). Lastly, travel restrictions, mobility reductions, and government interventions are explicitly modeled to mimic the real timeline of interventions of the events that occurred during the COVID-19 pandemic.

### B.3 SPATIOTEMPORAL NP MODEL

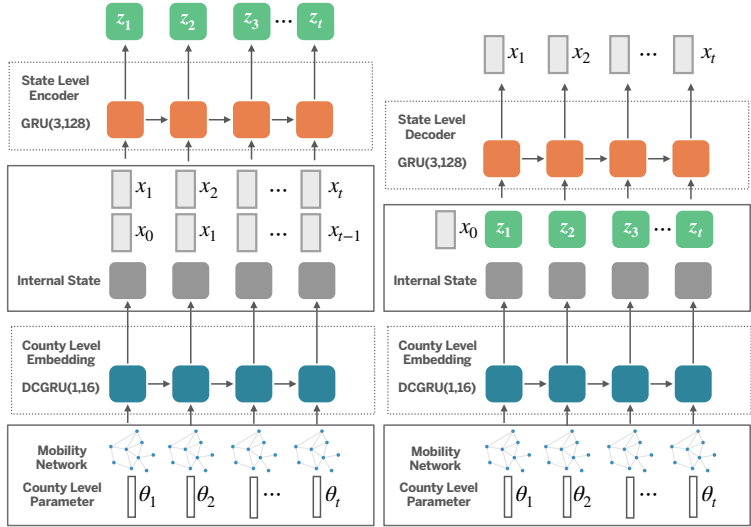


Figure 6: Visualization of the STNP model architecture. For both the encoder and the decoder, we use a diffusion convolutional GRU (DCGRU) Li et al. (2018) to capture spatiotemporal dependency.

As shown in Figure 6, our model has  $\theta$  at both county and state level and  $x_t$  at the state level. **The total number of dimensions of  $\theta$  is 16,264 (12 compartments  $\times$  2 (incidence + prevalence) + 28 days  $\times$  58 counties  $\times$  10 metadata parameters). The output dimension is 672 (12 compartments  $\times$  2 (incidence + prevalence)  $\times$  28 days)** We use county-level parameter  $\theta$  together with a county-to-county mobility graph  $A$  as input. We use the DCGRU layer (Li et al., 2017) to encode the graph in a GRU. We use a linear layer to map the county-level output to hidden features at the state level. For both the state-level encoder and decoder, we use multi-layer GRUs.

The input  $\theta_{1:t}$  is the county-level parameters for LEAM-US with a dimension of 10. The county level embedding uses 1 layer DCGRU with a width of 16. The internal state is at state level with a dimension of 16. The state level encoder and decoder use 3 layer GRUs with width of 128. The dimension of the latent process  $z_{1:t}$  is 32. The dimension of output  $x_{1:t}$  is 24, including the incidence and prevalence for 12 compartments. We trained STNP model for 500 steps with learning rate fixed at  $10^{-3}$  using Adam optimizer. We perform early stopping with 50 patience for both offline learning and Bayesian active learning.

### B.4 SEQUENTIAL NEURAL LIKELIHOOD (SNL) MODEL IMPLEMENTATION FOR SEIR SIMULATOR.

The input and output of SNL are the same as INP and GP. We use the likelihood-free inference code (Durkan et al., 2020) to implement SNL model. For each iteration, we perform Monte Carlo sampling from the trained likelihood model to report MAE. SNL takes 1 hour to complete training.

Table 2: Performance comparison of different acquisition functions in INP for SEIR simulator

Percentage of samples	LIG	Random	MeanSTD	MaxEntropy
1.11%	<b>365.87</b> $\pm$ 142.87	480.68 $\pm$ 5.24	480.22 $\pm$ 12.63	427.73 $\pm$ 61.36
1.85%	<b>236.9</b> $\pm$ 50.6	398.33 $\pm$ 131.05	314.75 $\pm$ 111.42	302.24 $\pm$ 119.84
2.96%	<b>119.26</b> $\pm$ 14.22	244.27 $\pm$ 148.89	158.94 $\pm$ 36.6	186.88 $\pm$ 57.48
4.07%	<b>96.73</b> $\pm$ 17.07	116.8 $\pm$ 9.1	127.36 $\pm$ 27.97	146.72 $\pm$ 26.06

Table 3: Performance comparison of different acquisition functions in GP for SEIR simulator

Percentage of samples	Random	MeanSTD	MaxEntropy
1.11%	663.76 $\pm$ 46.36	606.81 $\pm$ 6.89	<b>586.25</b> $\pm$ 58.44
1.85%	637.12 $\pm$ 13.45	<b>619.15</b> $\pm$ 36.42	628.54 $\pm$ 71.34
2.96%	597.3 $\pm$ 19.59	589.72 $\pm$ 24.9	<b>568.84</b> $\pm$ 19.05
4.07%	<b>519.98</b> $\pm$ 17.86	530.07 $\pm$ 32.95	578.34 $\pm$ 68.7

Table 4: Performance comparison of different acquisition functions in INP for LEAM-US simulator, population divided by 1000.

Percentage of samples	Random	MeanSTD	MaxEntropy	LIG
11.1%	20.961 $\pm$ 5.548	35.356 $\pm$ 28.706	65.498 $\pm$ 13.324	<b>14.447</b> $\pm$ 1.087
13.7%	13.418 $\pm$ 0.815	16.092 $\pm$ 3.11	30.496 $\pm$ 24.333	<b>11.704</b> $\pm$ 0.216
21.3%	9.332 $\pm$ 0.601	11.191 $\pm$ 0.184	10.028 $\pm$ 2.065	<b>7.593</b> $\pm$ 0.822
28.9%	8.077 $\pm$ 0.657	7.908 $\pm$ 0.536	8.417 $\pm$ 0.616	<b>6.539</b> $\pm$ 0.618
36.5%	6.719 $\pm$ 0.383	7.533 $\pm$ 0.861	7.431 $\pm$ 0.776	<b>6.008</b> $\pm$ 1.079

Table 5: Performance comparison between LIG in INP and SNL for SEIR simulator.

Percentage of samples	LIG	SNL
1.11%	<b>365.87</b> $\pm$ 142.87	<b>707.61</b> $\pm$ 44.42
1.85%	<b>236.9</b> $\pm$ 50.6	<b>669.03</b> $\pm$ 73.19
2.96%	<b>119.26</b> $\pm$ 14.22	<b>668.67</b> $\pm$ 72.42
4.07%	<b>96.73</b> $\pm$ 17.07	<b>685.28</b> $\pm$ 53.00

## B.5 IMPLEMENTATION DETAILS

For both GP and INP mimicking SEIR simulation, we ran experiments using CPU. No GPU accelerator is needed for this simple model. It takes 1.5 hours to complete training. For INP mimicking LEAM-US simulation, we ran experiments with GEFORCE RTX 2080. It takes 5 hours to complete training. For all experiments, we run with three different random seeds.

## B.6 BROADER IMPACT

The proposed approach has applications in the field of computational epidemiology as it provides a way to efficiently explore the parameter space of complex epidemiological model (such as LEAM-US) allowing to reduce the time-to-insights needed to provide actionable insights to policy makers. As an example, a single realization of the LEAM-US model requires about 1 hour to be completed (on a single thread of a AMD Ryzen Threadripper 3970x CPU, 3.7Ghz base clock) while in this work we show that using INP framework can achieve a performance comparable to the one of epidemic simulator while requiring less than one-third of the runs (as shown in Figure 5). In other words, this means that we can produce COVID-19 forecasts and projections three times faster than we would without the integration with the proposed framework. In addition, INP provides a general surrogate model that is not tailored to solve only for a specific target outcome (e.g. prediction of future deaths) but rather it retains the ability to describe the underlying *state of the system* as the original



mechanistic epidemic model. This feature preserves the ability of using the learned surrogate model as a situational awareness and outbreak analytics tool as it can also provide insights on unobservables (such as the fraction of immune population in a certain location/age group) which can be used by public health officials to assess the (potential) impact of past/future policies on the evolution of an epidemic.

## C ADDITIONAL RESULTS

### C.1 INP VERSUS GP AND SNL

Table 2 shows the average results together with the standard deviation of INP for SEIR simulator after running experiments three times. Table 3 shows the average results together with the standard deviation of GP for SEIR simulator. Figure 7 shows the average test MAE over three random runs versus the percentage of training data for SNL, NP and GP with different acquisition functions. Table 5 shows the comparison of average MAE together with the standard deviation for the SEIR simulator. We observe that LIG is always better than SNL until convergence. Table 4 shows the average MAE together with the standard deviation of INP for the LEAM-US simulator.

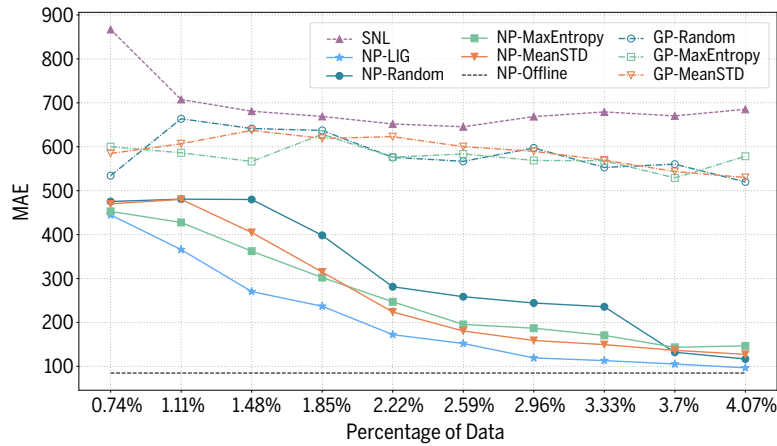


Figure 7: MAE loss versus the percentage of data for SNL, NP and GP using different acquisition functions.

### C.2 BATCH ACTIVE LEARNING WITH LIG.

Figure 8 compares 4 different setups: 8 batches (size 1), 4 batches (size 2), 2 batches (size 4), and 1 batch (size 8).




Figure 8: Batch size comparisons for LIG on the LEAM-US simulator. MAE loss versus the percentage of samples for INP during Bayesian active learning.