

Manifold-constrained Gaussian process inference for time-varying parameters in dynamic systems

Yan Sun

Georgia Institute of Technology

Shihao Yang (✉ shihao.yang@isye.gatech.edu)

Georgia Institute of Technology

Research Article

Keywords: ordinary differential equations, inverse problem, time-varying parameter estimation, Gaussian process, Bayesian inference

Posted Date: September 27th, 2022

DOI: <https://doi.org/10.21203/rs.3.rs-2102691/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Manifold-constrained Gaussian process inference for time-varying parameters in dynamic systems

Yan Sun¹ and Shihao Yang^{2*}

¹Department of Industrial and Systems Engineering, Georgia Institute of Technology, North Ave NW, Atlanta, 30332, Georgia, USA.

^{2*}Department of Industrial and Systems Engineering, Georgia Institute of Technology, North Ave NW, Atlanta, 30332, Georgia, USA.

*Corresponding author(s). E-mail(s): shihao.yang@isye.gatech.edu;
Contributing authors: ysun614@gatech.edu;

Abstract

Identification of parameters in ordinary differential equations (ODEs) is an important and challenging task when modeling dynamic systems in biomedical research and other scientific areas, especially with the presence of time-varying parameters. This article proposes a fast and accurate method, TVMAGI (Time-Varying MANifold-constrained Gaussian process Inference), to estimate both time-constant and time-varying parameters in the ODE using noisy and sparse observation data. TVMAGI imposes a Gaussian process model over the time series of system components as well as time-varying parameters, and restricts the derivative process to satisfy ODE conditions. Consequently, TVMAGI completely bypasses numerical integration and achieves substantial savings in computation time. By incorporating the ODE structures through manifold constraints, TVMAGI enjoys a principled statistical construct under the Bayesian paradigm, which further enables it to handle systems with missing data or unobserved components. The Gaussian process prior also alleviates the identifiability issue often associated with the time-varying parameters in ODE. Unlike existing approaches, TVMAGI can be applied to general nonlinear systems without specific structural assumptions. Three simulation examples, including an infectious disease compartmental model,

are provided to illustrate the robustness and efficiency of our method compared with numerical integration and Bayesian filtering methods.

Keywords: ordinary differential equations, inverse problem, time-varying parameter estimation, Gaussian process, Bayesian inference

1 Introduction

Ordinary Differential Equations (ODEs) are often used to analyze the behavior of dynamic systems, such as the spread of infectious diseases [1], interactions between species [2], and viral dynamics [3]. This paper studies a general formulation of ODE equations, where some of the parameters are allowed to be time-varying:

$$\dot{\mathbf{x}}(t) \equiv \frac{d\mathbf{x}(t)}{dt} = \mathbf{f}(\mathbf{x}(t), \boldsymbol{\theta}(t), \boldsymbol{\psi}, t), t \in [0, T] \quad (1)$$

Here, $\mathbf{x}(t)$ is the series of system outputs from time 0 to T , $\boldsymbol{\psi}$ denotes time-constant parameters, $\boldsymbol{\theta}(t)$ denotes time-varying parameters, and $\mathbf{f}$ is a set of general functions that characterize the derivative process. When $\mathbf{f}$ is non-linear, the system outputs $\mathbf{x}(t)$ typically do not have analytic solutions. To solve $\mathbf{x}(t)$ given initial conditions $\mathbf{x}(0)$ and parameters $\boldsymbol{\theta}(t)$ and $\boldsymbol{\psi}$, numerical integration methods are often required, such as Euler's Method or Runge-Kutta Method [4].

This paper focuses on the inverse problem that, given the observations, how to efficiently draw inference on the ODE parameters. Our goal is to estimate time-constant parameters $\boldsymbol{\psi}$ and time-varying parameters $\boldsymbol{\theta}(t)$ inside the ODE from data. In real world, the observation data of system components $\mathbf{x}$ are often obtained at discrete time points and are subject to measurement errors. We thus assume that we observe $\mathbf{y}(\boldsymbol{\tau}) = \mathbf{x}(\boldsymbol{\tau}) + \boldsymbol{\epsilon}(\boldsymbol{\tau})$, where $\boldsymbol{\tau}$ denotes the observation time points while error $\boldsymbol{\epsilon}(\boldsymbol{\tau})$ denotes Gaussian noise. We focus on the inference of $\boldsymbol{\theta}(t)$ and $\boldsymbol{\psi}$ given $\mathbf{y}(\boldsymbol{\tau})$, with emphasis on nonlinear structure $\mathbf{f}$.

The time-varying parameter $\boldsymbol{\theta}(t)$ in the ODE is often important yet challenging to recover from real-world data. For example, during a pandemic, the time-varying disease reproduction number is critical for public health policy decisions. However, its estimation can still be crude despite the best effort [5]. The time-varying $\boldsymbol{\theta}(t)$ provides too much degree of freedom to the ODE system, and two different $\boldsymbol{\theta}(t)$ can both give $\mathbf{x}(t)$ that fits the observation data, resulting in identifiability issues [6]. Such high degree of freedom in the time-varying parameters also gives overfitting issues for the usual gold standard numerical integration method, which tends to produce $\boldsymbol{\theta}(t)$ that fits the observation data exactly.

2 Review of related literature

Most existing methods for ODE inference only accommodate time-constant parameters [7–9]. For time-varying parameters inference in the ODE system, existing methods all have their deficiencies [10]. For example, [11] relied on time-consuming numerical integration; [12] proposed a Bayesian parametric approach to model time-varying coefficients in the HIV-1 dynamic model, sacrificing some flexibility; [13] developed an efficient two-stage local polynomial estimation method that circumvents numerical integration for a non-parametric time-varying parameters, but required ODE system to have linear dependency on the time-varying parameters (see Eq.(S9.2) in Supplementary Material). Bayesian filtering methods are also explored in the time-varying ODE parameter inferences, although lacking some statistical rigor. For example, [14] and [15] applied Ensemble Adjustment Kalman Filter (EAKF) algorithm to estimate parameters in a metapopulation SEIR model. [16] proposed an extended Kalman Filter approach based on Gauss-Markov process that can infer time-varying parameter but cannot accommodate time-constant parameters any more. To the best of our knowledge, there is no existing Bayesian inference method that eliminates numerical integration for a general ODE system with both time-constant and time-varying parameters.

3 Our contribution

We propose a fast and statistically principled method to infer time-varying $\theta(t)$ and time-constant ψ from noisy observations of ODE. The key idea is to use Bayesian approach and place Gaussian process (GP) prior on $\mathbf{x}(t)$ and time-varying parameters $\theta(t)$, thus the identifiability issue is mitigated using the informative prior that favors smoother parameter curves. Our method is built upon the prior work of MANifold-constrained Gaussian process Inference [8] where the Gaussian process $\mathbf{x}(t)$ is restricted on a manifold that satisfies the ODE system. Placing a Gaussian process on $\mathbf{x}(t)$ facilitates a fast inference on $\theta(t)$, as it completely bypasses numerical integration. Our approach also adheres to the classical Bayesian paradigm with principled posterior derivation. Through a Gaussian process model on $\theta(t)$, we are able to generalize the MANifold-constrained Gaussian process Inference to the situation where time-varying and time-constant parameters co-exist. We name our method TVMAGI (Time-Varying MANifold-constrained Gaussian process Inference), emphasizing its capability in handling time-varying parameters. We demonstrate the effectiveness of TVMAGI through three realistic simulation examples, where TVMAGI works well even when some of the system components $\mathbf{x}(t)$ are partially observed. Through these simulation examples, we also show that TVMAGI can outperform benchmark methods including a numerical integration approach, a Bayesian filtering approach, and a two-stage approach. Thanks to the computational savings of bypassing numerical integration, TVMAGI has great potential to be generalized in high-dimensional and large-scale systems. TVMAGI has a distinct contribution from the previously-proposed

time-constant parameter inference methods ([8, 9]) by investigating a much more complicated problem with functional estimate of time-varying parameters. The change from time-constant parameter to time-varying parameter also creates a notable difference in the scientific context, as parameters to be inferred in most real-world phenomena are non-stationary or changing over time.

4 Method of TVMAGI

The prior

Following standard Bayesian notation, the D -dimensional dynamic system $\mathbf{x}(t)$ is a realization of stochastic process $\mathbf{X}(t) = (X_1(t), \dots, X_D(t))$, and the P -dimensional time-varying parameters $\boldsymbol{\theta}(t)$ is a realization of stochastic process $\boldsymbol{\Theta}(t) = (\Theta_1(t), \dots, \Theta_P(t))$. We assume that $\boldsymbol{\Theta}(t)$ is continuous and differentiable in t during time period $[0, T]$, which helps to prevent overfit and to alleviate identifiability issue, but can be relaxed later. The prior distribution of $\mathbf{X}$ and $\boldsymbol{\Theta}$ in each dimension is independent Gaussian process. That is,

$$\Theta_p(t) \sim \mathcal{GP}(\mu_p^\Theta, \mathcal{K}_p^\Theta), t \in [0, T], p \in \{1, \dots, P\} \quad (2)$$

$$X_d(t) \sim \mathcal{GP}(\mu_d^X, \mathcal{K}_d^X), t \in [0, T], d \in \{1, \dots, D\} \quad (3)$$

where $\mathcal{K}_d^X$ and $\mathcal{K}_p^\Theta: \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$ are positive definite covariance kernels for GP, while μ_d^X and $\mu_p^\Theta: \mathbb{R} \rightarrow \mathbb{R}$ denote mean functions.

The likelihood

The observations are denoted as $\mathbf{y}(\boldsymbol{\tau}) = (y_1(\boldsymbol{\tau}_1), \dots, y_D(\boldsymbol{\tau}_D))$, where $\boldsymbol{\tau} = (\boldsymbol{\tau}_1, \boldsymbol{\tau}_2, \dots, \boldsymbol{\tau}_D)$ is the collection of observation time points across all components. Each component $X_d(t)$ can have its own set of observation time points $\boldsymbol{\tau}_d = (\tau_{d,1}, \dots, \tau_{d,N_d})$, where N_d is the number of observations of the d -th component. If the d -th component is not observed, then $N_d = 0$, and $\boldsymbol{\tau}_d = \emptyset$. The observation likelihood is thus assumed to be

$$Y_d(\boldsymbol{\tau}_d) = X_d(\boldsymbol{\tau}_d) + \epsilon(\boldsymbol{\tau}_d), \quad \epsilon(\boldsymbol{\tau}_d) \sim \mathcal{N}(0, \sigma_d^2) \quad (4)$$

In this paper, notation t shall refer to time generically, and $\boldsymbol{\tau}$ shall denote specifically the observation time points.

The manifold constraint

We introduce a variable W to quantify the difference in the derivative process $\dot{\mathbf{X}}(t)$ between Gaussian process and ODE:

$$W = \sup_{t \in [0, T], d \in \{1, \dots, D\}} |\dot{X}_d(t) - \mathbf{f}(\mathbf{X}(t), \boldsymbol{\Theta}(t), \boldsymbol{\Psi}, t)_d| \quad (5)$$

Intuitively, W is the L_∞ norm of derivative difference, and $W = 0$ if and only if $\dot{\mathbf{X}}(t)$ strictly satisfies the ODE structure, which is equivalent to constraining $\mathbf{X}(t)$ on the manifold of the ODE solutions. The advantage of

L_∞ norm is further discussed in Supplementary Material Section S1. In the ideal situation where $W \equiv 0$, the posterior distribution of $\Theta(t)$, Ψ , and $\mathbf{X}(t)$ shall be formulated as $P_{\Theta(t), \Psi, \mathbf{X}(t)|W, \mathbf{Y}(\tau)}(\theta(t), \psi, \mathbf{x}(t)|W = 0, \mathbf{Y}(\tau) = \mathbf{y}(\tau))$. However, such ideal posterior is not computable in practice. Therefore, we approximate W by finite discretization on time points $\mathbf{I} = \{t_1, t_2, \dots, t_n\}$, such that $\tau \subset \mathbf{I} \subset [0, T]$. We similarly define $W_{\mathbf{I}}$ on the discretization set $\mathbf{I}$ as the L_∞ distance of the derivative from GP and that from ODE:

$$W_{\mathbf{I}} = \sup_{t \in \mathbf{I}, d \in \{1, \dots, D\}} |\dot{X}_d(t) - \mathbf{f}(\mathbf{X}(t), \Theta(t), \Psi, t)_d| \quad (6)$$

Here $W_{\mathbf{I}}$ is the maximum on a finite set, and $W_{\mathbf{I}} \rightarrow W$ monotonically as $\mathbf{I}$ becomes dense. The associated computable Bayesian probability of the discretized manifold constraint $W_{\mathbf{I}} = 0$ is

$$\begin{aligned} & P(W_{\mathbf{I}} = 0 | \mathbf{X}(\mathbf{I}) = \mathbf{x}(\mathbf{I}), \Theta(\mathbf{I}) = \theta(\mathbf{I}), \Psi = \psi) \\ &= P(\dot{\mathbf{X}}(\mathbf{I}) - \mathbf{f}(\mathbf{X}(\mathbf{I}), \Theta(\mathbf{I}), \Psi, t_{\mathbf{I}}) = \mathbf{0} | \mathbf{X}(\mathbf{I}) = \mathbf{x}(\mathbf{I}), \Theta(\mathbf{I}) = \theta(\mathbf{I}), \Psi = \psi) \\ &= P(\dot{\mathbf{X}}(\mathbf{I}) = \mathbf{f}(\mathbf{x}(\mathbf{I}), \theta(\mathbf{I}), \psi, t_{\mathbf{I}}) | \mathbf{X}(\mathbf{I}) = \mathbf{x}(\mathbf{I})) \end{aligned} \quad (7)$$

which is a multivariate Gaussian distribution since the time derivative $\dot{X}_d(t)$ of GP is also a GP with specific mean and covariance kernel.

Supplementary Material Section S2 presents additional intuition regarding the manifold constraint $W_{\mathbf{I}} = 0$.

The posterior

Therefore, a computable discretized posterior for TVMAGI inference of $\mathbf{X}(t)$, $\Theta(t)$, and Ψ is:

$$P_{\Theta(\mathbf{I}), \Psi, \mathbf{X}(\mathbf{I})|W_{\mathbf{I}}, \mathbf{Y}(\tau)}(\theta(\mathbf{I}), \psi, \mathbf{x}(\mathbf{I}) | W_{\mathbf{I}} = 0, \mathbf{Y}(\tau) = \mathbf{y}(\tau)) \quad (8)$$

Equation (8) is the computable discretized posterior of TVMAGI inference. In this paper, we consider the Maximum A Posteriori (MAP) as the fast point estimate from TVMAGI, while the Posterior Mean and the Posterior Interval are the formal Bayesian inference results that further quantify the uncertainty.

The closed-form derivation

The posterior distribution of $\mathbf{X}(t)$, $\Theta(t)$, and Ψ in Eq.(8) can be further derived as

$$\begin{aligned} & P_{\Theta(\mathbf{I}), \Psi, \mathbf{X}(\mathbf{I})|W_{\mathbf{I}}, \mathbf{Y}(\tau)}(\theta(\mathbf{I}), \psi, \mathbf{x}(\mathbf{I}) | W_{\mathbf{I}} = 0, \mathbf{Y}(\tau) = \mathbf{y}(\tau)) \\ & \propto P(\Theta(\mathbf{I}) = \theta(\mathbf{I}), \Psi = \psi, \mathbf{X}(\mathbf{I}) = \mathbf{x}(\mathbf{I}) | W_{\mathbf{I}} = 0, \mathbf{Y}(\tau) = \mathbf{y}(\tau)) \end{aligned} \quad (9)$$

$$\begin{aligned} & \propto \pi_{\Psi}(\psi) \times \underbrace{P(\Theta(\mathbf{I}) = \theta(\mathbf{I}) | \Psi = \psi)}_{\text{1st Part, which is Eq.(2)}} \\ & \quad \times \underbrace{P(\mathbf{X}(\mathbf{I}) = \mathbf{x}(\mathbf{I}) | \Theta(\mathbf{I}) = \theta(\mathbf{I}), \Psi = \psi)}_{\text{2nd Part, which is Eq.(3)}} \\ & \quad \times \underbrace{P(\mathbf{Y}(\tau) = \mathbf{y}(\tau) | \mathbf{X}(\mathbf{I}) = \mathbf{x}(\mathbf{I}), \Theta(\mathbf{I}) = \theta(\mathbf{I}), \Psi = \psi)}_{\text{3rd Part, which is Eq.(4)}} \\ & \quad \times \underbrace{P(W_{\mathbf{I}} = 0 | \mathbf{Y}(\tau) = \mathbf{y}(\tau), \mathbf{X}(\mathbf{I}) = \mathbf{x}(\mathbf{I}), \Theta(\mathbf{I}) = \theta(\mathbf{I}), \Psi = \psi)}_{\text{4th Part, which is Eq.(7)}} \end{aligned} \quad (10)$$

$$\begin{aligned}
&= \pi_{\Psi}(\psi) \times \exp \left\{ - \underbrace{\frac{1}{2} \left(\sum_{p=1}^P \left[|\mathbf{I}| \log(2\pi) + \log |\mathcal{K}_p^{\Theta}(\mathbf{I})| + \|\theta_p(\mathbf{I}) - \mu_p^{\Theta}(\mathbf{I})\|_{\mathcal{K}_p^{\Theta}(\mathbf{I})^{-1}}^2 \right] \right)}_{\text{1st Part, which is Eq.(2)}} \right. \\
&+ \sum_{d=1}^D \left[\underbrace{|\mathbf{I}| \log(2\pi) + \log |\mathcal{K}_d^X(\mathbf{I})| + \|x_d(\mathbf{I}) - \mu_d^X(\mathbf{I})\|_{\mathcal{K}_d^X(\mathbf{I})^{-1}}^2}_{\text{2nd Part, which is Eq.(3)}} \right. \\
&+ \underbrace{N_d \log(2\pi\sigma_d^2) + \|x_d(\boldsymbol{\tau}_d) - y_d(\boldsymbol{\tau}_d)\|_{\sigma_d^2}^2}_{\text{3rd Part, which is Eq.(4)}} \\
&\left. \left. + \underbrace{|\mathbf{I}| \log(2\pi) + \log |C_d| + \|\mathbf{f}_{d,\mathbf{I}}^{\mathbf{x},\boldsymbol{\theta},\psi} - \dot{\mu}_d^X(\mathbf{I}) - {}'\mathcal{K}_d^X(\mathbf{I})\mathcal{K}_d^X(\mathbf{I})^{-1}\{x_d(\mathbf{I}) - \mu_d^X(\mathbf{I})\}\|_{C_d^{-1}}^2}_{\text{4th Part, which is Eq.(7)}} \right] \right\} \quad (11)
\end{aligned}$$

where $\|v\|_A^2 = v^T A v$, $|\mathbf{I}|$ is the cardinality of $\mathbf{I}$, and $\mathbf{f}_{d,\mathbf{I}}^{\mathbf{x},\boldsymbol{\theta},\psi}$ is short for the d -th component of $\mathbf{f}(\mathbf{x}(\mathbf{I}), \boldsymbol{\theta}(\mathbf{I}), \psi, t_{\mathbf{I}})$, and $C_d = \dot{\mathcal{K}}_d^X(\mathbf{I})\mathcal{K}_d^X(\mathbf{I})^{-1}\mathcal{K}_d^X(\mathbf{I})$ is the conditional covariance matrix of $\dot{X}_d(\mathbf{I})$ given $X_d(\mathbf{I})$.

A deeper look into the above equation reveals that Eq.(9) is the joint probability in Bayesian statistics, and Eq.(10) further decomposes it into parts. The 1st Part (which is Eq.(2)) corresponds to independent GP prior distribution of $\boldsymbol{\Theta}(\mathbf{I})$, as the prior of $\boldsymbol{\Theta}(t)$ and Ψ are independent. The 2nd Part (which is Eq.(3)) is the prior of GP on $\mathbf{X}(\mathbf{I})$, because the prior of $\mathbf{X}(\mathbf{I})$ is independent from $\boldsymbol{\Theta}(t)$ and Ψ . The 3rd Part (which is Eq.(4)) is the level of observation noise, and given the value of underlying true components $\mathbf{X}(\boldsymbol{\tau})$, the additive Gaussian observation noise $\epsilon(\boldsymbol{\tau})$ is independent from everything else. The 4th Part (which is Eq.(7)) can be simplified to be the conditional probability of $\dot{\mathbf{X}}(\mathbf{I})$ given $\mathbf{X}(\mathbf{I})$ evaluated at $\mathbf{f}(\mathbf{x}(\mathbf{I}), \boldsymbol{\theta}(\mathbf{I}), \psi, t_{\mathbf{I}})$. All four parts are multivariate Gaussian distributed. Especially, The 4th Part (which is Eq.(7)) is Gaussian because conditional $\dot{\mathbf{X}}(\mathbf{I})$ given $\mathbf{X}(\mathbf{I})$ has a multivariate Gaussian distribution, provided that the GP kernel $\mathcal{K}^X$ is twice differentiable.

We choose Matern kernel with degree of freedom $\nu = 2.01$ for both $\boldsymbol{\Theta}(t)$ and $\mathbf{X}(t)$ to guarantee a differentiable GP that allows more flexible patterns:

$$\mathcal{K}_{\nu}(l) = \phi_1^2 \frac{2^{1-\nu}}{\Gamma(\nu)} (\sqrt{2\nu} \frac{l}{\phi_2})^{\nu} K_{\nu}(\sqrt{2\nu} \frac{l}{\phi_2}), \quad l = |s - t| \quad (12)$$

where K_{ν} denotes the modified Bessel function of the second kind. In this case, $'\mathcal{K} = \frac{\partial}{\partial s} \mathcal{K}(s, t)$, $\mathcal{K}' = \frac{\partial}{\partial t} \mathcal{K}(s, t)$, and $\mathcal{K}'' = \frac{\partial^2}{\partial s \partial t} \mathcal{K}(s, t)$ are all well-defined.

5 Algorithm

This section provides a detailed computational scheme of TVMAGI, including the hyper-parameter settings. The implementation is available on GitHub. Overall, the Maximum A Posteriori (MAP) of $\mathbf{X}(\mathbf{I})$, $\boldsymbol{\Theta}(\mathbf{I})$, and Ψ is obtained

by optimization, while the posterior mean/interval is obtained by Hamiltonian Monte Carlo. To set the hyper-parameters and initiate the optimizer, we introduce a multi-stage approach in the algorithm. The advantages of the multi-stage mechanism are discussed in Supplementary Material Section S10.

5.1 Initialization and inference of the mean

At the first stage, we impose a GP only on $\mathbf{X}(t)$ and substitute the time-varying $\boldsymbol{\theta}(t)$ with its unknown mean $\boldsymbol{\mu}^\Theta$ in the entire model. This formulation ignores the time-varying property of $\boldsymbol{\theta}(t)$ and treats it as time-constant, which fits in the time-constant parameters inference framework of [8]. As such, we can use MAGI package ([8]) to obtain point estimates for the parameters and system components, denoted as $\hat{\boldsymbol{\mu}}^\Theta$, $\boldsymbol{\psi}^{(0)}$, and $\mathbf{x}(\mathbf{I})^{(0)}$. The $\hat{\boldsymbol{\mu}}^\Theta$ is subsequently used as the prior mean value for the time-varying $\boldsymbol{\theta}(t)$ in an empirical Bayes fashion, and will be plugged in Eq.(11). The $\boldsymbol{\psi}^{(0)}$ and $\mathbf{x}(\mathbf{I})^{(0)}$ will be used as the initial values for $\boldsymbol{\psi}$ and $\mathbf{x}(\mathbf{I})$ in the later MAP optimization.

The hyper-parameters $(\phi_{1,d}^X, \phi_{2,d}^X)$ for kernel $\mathcal{K}_d^X$ in Eq.(12) and the noise level σ for each system component $\dot{X}_d$, $d = 1, \dots, D$, are also estimated in MAGI package, using the Gaussian process smoothing marginal likelihood ([8]). The MAGI estimated noise level $\sigma^{(0)}$ will serve as initial value for later joint MAP optimization.

5.2 Point-wise inference of the time-varying parameters

At the second stage of TVMAGI, we obtain an initial estimate for time-varying $\boldsymbol{\theta}(\mathbf{I})$ by removing the smoothing GP prior. That is, we maximize the partial posterior Eq.(13) conditioning on $\boldsymbol{\Theta}(\mathbf{I})$, without considering Part 1 Eq.(2):

$$\begin{aligned} & \tilde{\mathbf{x}}(\mathbf{I}), \tilde{\boldsymbol{\theta}}(\mathbf{I}), \tilde{\boldsymbol{\psi}}, \tilde{\sigma} \\ &= \arg \max_{\mathbf{x}, \boldsymbol{\theta}, \boldsymbol{\psi}, \sigma} p_{\boldsymbol{\Psi}, \mathbf{X}(\mathbf{I})|W_I, \mathbf{Y}(\tau), \boldsymbol{\Theta}(\mathbf{I})}(\boldsymbol{\psi}, \mathbf{x}(\mathbf{I})|W_I = 0, \mathbf{Y}(\tau) = \mathbf{y}(\tau), \boldsymbol{\Theta}(\mathbf{I}) = \boldsymbol{\theta}(\mathbf{I})) \\ & \propto \pi_{\boldsymbol{\Psi}}(\boldsymbol{\psi}) \times \underbrace{P(\mathbf{X}(\mathbf{I}) = \mathbf{x}(\mathbf{I}))}_{\text{2nd Part Eq.(3)}} \end{aligned} \quad (13)$$

$$\begin{aligned} & \times \underbrace{P(\mathbf{Y}(\tau) = \mathbf{y}(\tau)|\mathbf{X}(\mathbf{I}) = \mathbf{x}(\mathbf{I}))}_{\text{3rd Part Eq.(4)}} \times \\ & \underbrace{P(\dot{\mathbf{X}}(\mathbf{I}) = \mathbf{f}(\mathbf{x}(\mathbf{I}), \boldsymbol{\theta}(\mathbf{I}), \boldsymbol{\psi}, t_I)|\mathbf{X}(\mathbf{I}) = \mathbf{x}(\mathbf{I}))}_{\text{4th Part Eq.(7)}} \end{aligned} \quad (14)$$

The optimization is initialized at $\mathbf{x}(\mathbf{I})^{(0)}$, $\boldsymbol{\psi}^{(0)}$, $\sigma^{(0)}$, and the $\boldsymbol{\theta}(\mathbf{I})$ is initialized at $\hat{\boldsymbol{\mu}}^\Theta$. We denote the optimized $\boldsymbol{\theta}(\mathbf{I})$ as $\tilde{\boldsymbol{\theta}}(\mathbf{I})$, and $\mathbf{x}(\mathbf{I})^{(0)}$, $\boldsymbol{\psi}^{(0)}$, $\sigma^{(0)}$ are updated to a new optimum $\tilde{\mathbf{x}}(\mathbf{I})$, $\tilde{\boldsymbol{\psi}}$ and $\tilde{\sigma}$. We call $\tilde{\boldsymbol{\theta}}(\mathbf{I})$ the point-wise estimate since there is no requirement on the smoothness or continuity of $\tilde{\boldsymbol{\theta}}(t)$ on $\mathbf{I}$. Although wiggling and possibly overfitting the data, the point-wise estimate $\tilde{\boldsymbol{\theta}}(\mathbf{I})$ captures the trend of parameter changes, which provides information to set the hyper-parameters of GP kernels $\mathcal{K}_p^\Theta$ for $\boldsymbol{\Theta}(t)$.

5.3 GP hyper-parameters for time-varying ODE parameters

Length scale parameter ϕ^Θ controls how fast $\theta(t)$ could change. Provided the point-wise estimate $\tilde{\theta}(\mathbf{I})$, we use Gaussian Process smoothing method to set the hyper-parameters $\phi_{1,p}^\Theta, \phi_{2,p}^\Theta$ of GP kernels $\mathcal{K}_p^\Theta$ in Eq.(12). We shall treat $\tilde{\theta}(\mathbf{I})$ as observations of $\Theta(\mathbf{I})$, and operate on each dimension of time-varying ODE parameters separately.

Recall the prior $\Theta_p(\mathbf{I}) \sim \mathcal{GP}(\hat{\mu}_p^\Theta, \mathcal{K}_p^\Theta(\mathbf{I}, \mathbf{I}))$, where the mean $\hat{\mu}_p^\Theta$ is obtained in Section 5.1. We use the empirical Bayes approach again to set $\phi_{1,p}^\Theta, \phi_{2,p}^\Theta$ by maximizing its posterior density at $\tilde{\theta}_p(\mathbf{I})$:

$$\hat{\phi}_{1,p}^\Theta, \hat{\phi}_{2,p}^\Theta = \arg \max_{\phi_1, \phi_2} \pi_{\Phi_p}(\phi) P(\tilde{\theta}_p(\mathbf{I}) | \phi) \quad (15)$$

where $\tilde{\theta}_p(\mathbf{I}) | \phi \sim \mathcal{N}(\hat{\mu}_p^\Theta, \mathcal{K}(\phi) + \text{diag}(\delta^2))$, the δ is the nuisance parameter governing the induced noise in point-wise estimate $\tilde{\theta}_p(\mathbf{I})$, and the $\pi_{\Phi_p}(\cdot)$ is the hyper-prior. In practice, the hyper-prior $\pi_{\Phi_p}(\cdot)$ is often set to be uniform on a reasonable interval depending on the context to ensure desired level of smoothness for the time-varying ODE parameter $\theta_p(t)$.

5.4 Maximum A Posteriori (MAP) optimization

All hyper-parameters are now set and will be held as constant when optimizing Eq.(11) to get the MAP, with initial values $\theta(\mathbf{I})^{(0)} = \tilde{\theta}(\mathbf{I})$, $\mathbf{x}(\mathbf{I})^{(0)} = \tilde{\mathbf{x}}(\mathbf{I})$, $\psi^{(0)} = \tilde{\psi}$ and $\sigma^{(0)} = \tilde{\sigma}$, all from Section 5.2. The joint posterior function Eq.(11) of $\mathbf{x}(\mathbf{I})$, $\theta(\mathbf{I})$, ψ and σ is optimized with Adam optimizer [17] in PyTorch to get the MAP estimation of TVMAGI. Finally, to mitigate the potential issue of Adam optimizer converging to local optimum, we suggest trying multiple initial values, including starting $\mathbf{x}(\mathbf{I})$ at linear interpolations from the observations $\mathbf{y}(\tau)$.

5.5 Interval estimation of parameters

In addition to the MAP point estimate, we also quantify of the parameter uncertainty in TVMAGI using posterior samples. In particular, we sample the posterior function Eq(11) using Hamilton Monte Carlo (HMC), while holding all the hyper-parameters at the same constant value as in Section 5.4. Details about the HMC algorithm can be found in Supplementary Material Section S3. The interested reader may refer to [18] for more thorough introduction to HMC. Specifically in all illustration examples of this paper, we set step size $\epsilon = 10^{-5}$, number of leap-frog steps $L = 100$, sample size 8000, burn-in ratio 0.5, and the HMC is initialized at the MAP estimate.

6 Benchmark methods and evaluation metrics

6.1 Benchmark methods

We compare our method with two common approaches for time-varying parameter inference in ODE: numerical integration methods, represented by Runge-Kutta method [4], and Bayesian filtering methods, represented by Ensemble Adjustment Kalman Filter (EAKF)[15], which has been used in estimating the influenza disease spread SIRS model parameter [15] and studying time-varying fatality rate In COVID-19 disease spread modeling [19]. Supplementary Material Section S6 provides the review of two approaches and some additional theoretical discussion about the limitations and the statistical rigor of the benchmark methods for ODE inference when time-varying parameters and time-constant parameters co-exist.

6.2 Evaluation Metrics

To assess the quality of the parameter estimates and the system recovery, we consider two metrics based on root mean squared error (RMSE). First, we examine the accuracy of the parameter estimates, using *parameter RMSE*. For the time-constant parameters, we directly calculate the RMSE of the parameter estimates to the true parameter value across simulations. For the time-varying parameters, we additionally average over discretization set $\mathbf{I}$ for the RMSE. Second, we examine the system recovery, using *trajectory RMSE*. Due to the potential identifiability issue that different parameters can give similar system observations, we measure how well the system components are recovered as another independent evaluation. To calculate the trajectory RMSE, we use numerical integration to reconstruct the trajectory based on the TVMAGI inferred parameters and initial conditions. The RMSE of the reconstructed trajectory to the true system is then calculated at observation time points.

We emphasize that the numerical integration is only used for evaluation purpose, and throughout our TVMAGI approach, no numerical integration is ever needed. For better distinction, we refer to the MAP of $\mathbf{x}(\mathbf{I})$ directly from TVMAGI as the *inferred trajectory*, and refer to the numerically integrated $\mathbf{x}(t)$ based on the TVMAGI inferred parameters and initial conditions as the *reconstructed trajectory*.

To assess the quality of the interval estimates, we consider the Frequentist coverage of our posterior intervals. For the time-constant parameters, we directly calculate the proportion of repeated simulations where our posterior interval covers the truth. For the time-varying parameters, we additionally average over discretization set $\mathbf{I}$ for the coverage. The coverage of the inferred trajectory can be similarly calculated, averaging over discretization set $\mathbf{I}$. We do not compute the coverage of the reconstructed trajectory as it will require numerical solver for each posterior sample of the parameters and initial conditions.

7 Results

We illustrate the accuracy and efficiency of TVMAGI through three realistic simulation studies of ODE models in epidemiology, ecology, and system biology. We begin with a disease compartmental model that demonstrates the effectiveness of TVMAGI for problems with partially observed system component(s). We then use an ecology example to show how TVMAGI can mitigate the identifiability issue through the informative GP prior that favors smoother time-varying parameters. Lastly, we apply TVMAGI on a system biology example with non-stationary rapid-changing time-varying parameters, and presents TVMAGI's competitive performance with one additional tailor-made benchmark method for such ODE.

7.1 SEIRD model

Consider a COVID-19 cases/deaths modeling using an infectious disease Susceptible-Exposed-Infectious-Recovered-Deceased (SEIRD) compartmental ODE model [20, 21], where the entire population is classified into S , E , I , R , D components, and any transitions from one state to another state (i.e., the disease spreading dynamics) are modeled as ODE:

$$\frac{dS}{dt} = -\frac{\beta IS}{N}, \quad \frac{dE}{dt} = \frac{\beta IS}{N} - v^e E, \quad \frac{dI}{dt} = v^e E - v^i I, \quad \frac{dD}{dt} = v^i I \cdot p^d \quad (16)$$

N is the total population, and the cumulative recovered population is $R = N - S - E - I - D$. The S, E, I and D denote the susceptible, exposed, infected population and cumulative death respectively. The 4 parameters of interest are investigated: rate of contact by an infectious individual (β), rate of transferring from state of exposed to infectious (v^e), rate of leaving infectious period (v^i) and fatality rate (p^d). During a pandemic, parameters in the SEIRD model can evolve over time due to pharmaceutical and non-pharmaceutical interventions. We assume that β is time-varying due to the mutation of disease and policy interventions during a specific time; p^d is time-varying depending on the sufficiency of medical treatments; v^e is time-varying due to the different levels of public awareness or complacency, and v^i is assumed to be unknown time-constant parameter to avoid identifiability issues.

In the experiment we set $v^i = 0.1$, $\beta_t = 1.8 - \cos(\pi t/8)$, $v_t^e = 0.1 - 0.02 \cos(\pi t/8)$, $p_t^d = 0.05 + 0.025 \cos(\pi t/8)$, and focus on a time horizon of 32 days. The initial values of four components are set as (100000, 100, 50, 50) for (S, E, I, D) . We assume S, I, D are observed on daily frequency with log-normal multiplicative observation noise at 3% level. The exposed population E is assumed to be only sparsely observable at 3% noise level, with one observation per two days, due to the high cost of data acquisition from sampling test. Such pandemic settings [22, 23] capture the periodic fluctuation of parameters often observed in the real world.

We apply TVMAGI on a log-transformed system (by taking the log of populations in each of the S, E, I, D state) over 100 simulation datasets, with




Fig. 1 Results of parameter inference (upper) and reconstructed trajectory (lower) of TVMAGI in 100 simulated datasets for SEIRD model. The mean and the 95% interval here refer to the point estimates across 100 simulated datasets. One sample simulation dataset is also presented to visualize the noise level and observation schedule.

Table 1 Accuracy comparison for the SEIRD model based on 100 simulation datasets. The mean of RMSE is reported first with the standard deviation across 100 replications followed after $\pm$ for the parameters and the reconstructed trajectories. The last column is the coverage of interval estimates for the parameters and the inferred trajectories. The last row shows the computing time (in seconds) needed to obtain point estimates from all methods.

	Point Estimate				TVMAGI Posterior Samples	
	RMSE	TVMAGI-MAP	Runge-Kutta	EAKF	Posterior Mean RMSE	Interval Coverage
Parameter	β	0.114$\pm$ 0.039	0.178 $\pm$ 0.094	0.706 $\pm$ 0.010	0.110 $\pm$ 0.043	98.2%
	v^e	0.009$\pm$ 0.010	0.051 $\pm$ 0.030	0.057 $\pm$ 0.003	0.007 $\pm$ 0.005	97.4%
	v^i	0.005 $\pm$ 0.003	0.004$\pm$ 0.003	0.151 $\pm$ 0.008	0.007 $\pm$ 0.004	91.0%
	p^d	0.019$\pm$ 0.029	0.083 $\pm$ 0.073	0.039 $\pm$ 0.003	0.011 $\pm$ 0.008	98.0%
Trajectory	S	581.7$\pm$ 272.1	1084.8 $\pm$ 195.3	3868.3 $\pm$ 132.2	615.3 $\pm$ 294.5	98.8%
	E	704.7$\pm$ 218.3	951.7 $\pm$ 142.3	5376.1 $\pm$ 496.6	660.7 $\pm$ 202.3	96.6%
	I	439.0$\pm$ 140.4	556.2 $\pm$ 90.6	3167.0 $\pm$ 404.9	415.0 $\pm$ 158.7	96.4%
	D	38.3 $\pm$ 4.9	33.3$\pm$ 5.0	907.3 $\pm$ 48.6	14.0 $\pm$ 5.2	94.2%
Computing Time (s)		1006.7 $\pm$ 115.54	2904.4 $\pm$ 195.5	7.3 $\pm$ 0.4	-	-

2 discretizations per day. Figure 1 shows the results of parameter inference and the TVMAGI reconstructed trajectory $\mathbf{X}(\mathbf{I})$ of the ODE system. The parameter RMSE and trajectory RMSE introduced in Section 6.2 are presented in Table 1, where TVMAGI is shown to be more accurate than Runge-Kutta or EAKF.

For point estimates, Figure 1 and Table 1 show that, even when the exposed population is sparsely observed, TVMAGI is still capable of providing good results of inference. As the most important parameter when assessing the spread of disease, β_t can be accurately and robustly inferred. v^i can also be accurately inferred as constant. p_t^d has larger variability at the start, as initial deaths are too few to provide enough information. In comparison, the variability of v_t^e inference increases at the end of the period, because susceptible population has decreased to nearly zero while infectious population reaches plateau. Despite variations in the inferred parameters, the inferred system trajectories are all very close to the truth, confirming the intuition that the system is possibly less sensitive to p^d in earlier state and v^e in later stage. Supplementary Material Section S5

has the visual illustration for Runge-Kutta or EAKF, and their accuracy is far from satisfactory: Runge-Kutta method will overfit the observation noise, and EAKF reconstructed trajectory completely misses the truth.

For Interval estimates, Figure 10 in Supplementary Material Section S7 gives a visual illustration for 10 sample datasets. The coverage of Posterior Interval across 100 simulated datasets is included in Table 1. The empirical coverage of the interval is reasonable around the 95% nominal value. The intervals are wider for p^d at the starting time, and wider for v^e at the ending time, which are consistent with the intuition about their sensitivity discussed above. More interval estimation results are available in Supplementary Material Section S7.

On the computational cost, Table 1 also shows that TVMAGI is much faster than the Runge-Kutta numerical integration methods. EAKF is fast, but gives unreliable results (see Supplementary Material Section S6 for more discussion on the reliability of EAKF).

7.2 Lotka-Volterra Model

Lotka-Volterra (LV) model (a.k.a. predator-prey model) is widely used to describe population fluctuation of predators and preys and their interactions in the ecosystem [24]. With the introduction of time-varying parameters, the system becomes weakly identifiable during certain time range, which creates a challenge in the inference. Specifically, the ODE system is characterized as:

$$\frac{dx}{dt} = \alpha_t x - \beta xy, \quad \frac{dy}{dt} = \delta xy - \gamma_t y \quad (17)$$

where x and y denote the population of preys and predators. α_t indicates the birth rate of the prey and γ_t denotes the death rate of the predator, both of which are assumed to fluctuate according to seasonality. β and δ describe the interaction relationships between predators and preys, and are assumed constant. We set the parameters $\beta = 0.75$, $\delta = 1$, $\alpha_t = 0.6 + 0.3 \cos(\pi t/5)$, and $\gamma_t = 1 + 0.1 \sin(\pi t/5)$. The time is measured on a yearly basis, and data for 20 years are generated with monthly observations contaminated by 3% multiplicative log-normal noise. The initial values of predators and preys are 1 and 3, as an ideal ratio in real ecology systems [25].

Fig. 2 Comparison of inferred $\theta(t)$ and reconstructed $\mathbf{X}(t)$ of LV model. The mean and the 95% interval here refer to the point estimates across 100 simulated datasets. One sample simulation dataset is also plotted to visualize the noise level.




Figure 2 shows the estimated time-varying parameters and the reconstructed trajectory $\mathbf{X}(I)$, with parameter RMSE and trajectory RMSE presented in

Table 2 Accuracy comparison of estimated parameters and reconstructed trajectory in the LV model based on 100 simulation datasets. See legend of Table 1 for detailed description.

Point Estimate					TVMAGI Posterior Samples	
RMSE		TVMAGI	Runge-Kutta	EAKF	Posterior mean RMSE	Interval Coverage
Parameter	$\alpha(t)$	0.0395±0.0302	0.1155± 0.0920	0.2330±0.1589	0.0450 ± 0.0346	40.5%
	β	0.0154±0.0098	0.0480± 0.0069	0.1041±0.0621	0.0187 ± 0.0122	62.0%
	δ	0.0156±0.0111	0.0166± 0.0058	0.1605±0.0831	0.0160 ± 0.0119	58.0%
	$\gamma(t)$	0.0304±0.0226	0.0863± 0.0756	0.0974±0.1847	0.0341 ± 0.0201	46.0%
Traj.	x (prey)	0.0606± 0.0550	0.0314±0.0085	0.4701±0.0903	0.0838 ± 0.0568	69.1%
	y (predator)	0.0813± 0.0622	0.0384±0.0131	0.2773±0.0646	0.0989 ± 0.0714	62.2%
Computing Time (s)		910.8± 113.7	2042.1± 85.2	22.7± 1.1	-	

Table 2. Our recovered system components x and y are very close to the truth, despite the weak identifiability of the parameter α_t and γ_t when the x and y are at peak (year=12). Most notably, α_t could deviate from the truth in the attempt to best fit the observed noisy data at the peak of x_t , resulting in a biased inference of the time-varying parameters at the weakly identifiable time points, although all deviations are still within the range of smoothness constraints on α_t . Nevertheless, both TVMAGI inferred system components x and y are still accurate.

Comparing with benchmark models in Table 2, TVMAGI gives the most accurate parameter inference thanks to the GP smoothing prior that mitigates the identifiability issue. The numerical method of Runge-Kutta gives better trajectory inference, but it cannot handle the identifiability issue in the parameters. The coverage from TVMAGI is not ideal, possibly due to the bias in $\alpha(t)$ estimate and the variance in $\gamma(t)$ estimate – if the GP smoothing prior is too strong, the point estimates will be biased, and if the GP smoothing prior is too weak, the point estimates will have large variance (see Supplementary Material Figure 10). The comparison on computational cost again demonstrates the expected advantage of TVMAGI over Runge-Kutta, while EAKF is the fastest method with the worst accuracy.

This example illustrates the performance of TVMAGI in the presence of weak identifiability – the inferred time-varying parameters at the weakly identified time points could subject to deviation from the truth, although the parameters are smooth and still fit the observed data well.

7.3 HIV model

In this example, we compare TVMAGI with a state-of-the-art two-stage Efficient Local Estimation (ELE) method proposed by [26] in an HIV dynamic model that they studied. This is a challenging case for GP modeling as the true time-varying parameter has non-periodic non-stationary trends with rapid changes [3, 27]. To use the ELE method of [26], the ODE system must fit in the linear form of Eq.(18):


$$X'(t) = \sum_{i=1}^d a_i(t) Z_i(t) - cX(t) \quad (18)$$

Table 3 Accuracy comparison of estimated parameters and trajectories in the HIV model based on 100 simulation datasets. See legend of Table 1 for detailed description.

Point Estimate					TVMAGI Posterior Samples	
RMSE	TVMAGI	Runge-Kutta	EAKF	ELE	Posterior mean RMSE	Interval Coverage
$a(t)$	281.9 ± 71.8	695.7 ± 50.9	818.6 ± 54.9	291.5 ± 48.4	359.0 ± 127.8	74.1%
$x(t)$	0.057 ± 0.002	0.038 ± 0.003	0.181 ± 0.004	0.075 ± 0.003	0.069 ± 0.004	65.9%
Computing Time (s)	897.6 ± 72.1	1940.2 ± 79.7	5.4 ± 0.3	10.7 ± 0.1	-	-

where $Z_i(t)$ is the known covariate, and $a_i(t)$ is the unknown time-varying coefficient. For benchmark comparison, we treat $a_1(t)$ and $a_2(t)$ in Eq.(18) as unknown time-varying parameters for TVMAGI. Detailed illustration of HIV model formulation is provided in Supplementary Material S9.

Figure 3 shows the TVMAGI inferred parameter $a(t) = \sum_{i=1}^d a_i(t)Z_i(t)$ and the reconstructed trajectory $X(t)$. The parameter/trajectory RMSEs of TVMAGI and the benchmark methods are reported in Table 3. TVMAGI has a small advantage over the state-of-the-art method on HIV model inference of $X(t)$. Further visual comparison to benchmark methods (Supplementary Material Figure 9) shows that TVMAGI is slightly more accurate at the beginning phase of the system, which is in fact the most challenging phase for HIV inference as viral load drops sharply due to the drug effect. TVMAGI also achieves competitive inference result on $a(t)$, which is of clinical importance for the generation rate of HIV virus [13]. The TVMAGI posterior interval coverage is less ideal because of the decreased accuracy in $a(t)$ towards the ending period (Supplementary Material Figure 12). Most importantly, while the benchmark method requires a highly restricted form of ODE formulation, TVMAGI assumes no specific form of ODE equations, and is thus applicable for general ODE systems, albeit with longer computing time.

Fig. 3 TVMAGI inferred $a(t)$ and reconstructed $\mathbf{X}(t)$ of HIV model.

Overall in this example, we compare TVMAGI with an additional benchmark method that can only be applied to the ODEs with a specific form, where TVMAGI is shown to provide competitive inference accuracy while having much more general applicability. The application in HIV model also illustrates that TVMAGI could work well with non-stationary trends in the time-varying parameters, where the time-varying parameters is not periodic and have rapid changes in part of the time horizon.

8 Sensitivity analysis

We conduct three sensitivity analysis to show the robustness of our approach: the number of discretization, the selection of GP kernel, and the mis-specified time-varying parameters. Detailed discussion is provided in Supplementary Material Section S8, along with tables and visualizations of SEIRD model results.

9 Discussion

In this paper, we introduce a Bayesian approach, TVMAGI, for time-varying parameters inference in ODE dynamic systems. TVMAGI models time-varying parameters and system components as Gaussian process, and is constrained to have the derivative processes satisfy the ODE dynamics. We show that TVMAGI is statistically principled and illustrate its general applicability through three simulation examples. Results have shown that TVMAGI yields accurate and robust parameter inference from noisy observations, with reasonable interval estimates as well. Moreover, TVMAGI can mitigate the identifiability issue and the over-fitting issue in the time-varying parameters using the informative GP smoothing prior. TVMAGI is also generally applicable in the presence of missing observations.

TVMAGI is more accurate than the benchmark methods because TVMAGI addresses the challenges of the numerical integration method and the Bayesian filtering method for ODE time-constant and time-varying parameter inference. Numerical integration methods are the gold standard for the ODE parameter inference when all parameters are time-constant. However, with the presence of time-varying parameters, due to the lack of smoothness structure on $\theta(t)$, the inferred time-varying parameters from Runge-Kutta will overfit the noisy observation data, resulting in volatile $\theta(t)$ with little information about the true trends. The Bayesian filtering approach, on the other hand, cannot infer time-constant parameter ψ because the update in ψ is not permissible in a state-space model fashion. We can nevertheless enforce an update on the time-constant parameter, but there will be no guarantee on the accuracy of the reconstructed trajectory. The ELE two-stage approach relies on a regression technique that can only be used if the ODE has linear dependency on the time-varying parameters. Therefore, TVMAGI is the only approach that is theoretically sound, practically accurate, and generally applicable for the ODE inference problem when time-constant and time-varying parameters co-exist.

On the computational time comparison, TVMAGI has notable advantage of reduced computation cost compared to numerical integration method, while the inference is more accurate compared to the fast-yet-unreliable Bayesian filtering methods. Even for the three small-sized problems in this paper, TVMAGI is more than twice as fast than the numerical integration method of Runge-Kutta with better accuracy. When dealing with large-scale system, the gain in computational time is likely to be even larger, as TVMAGI computational time would scale linearly as the dimension of system components grow,

while inference with numerical integration method typically scales super linearly. Therefore, TVMAGI has strong potential in large-scale systems, where numerical integration is expensive.

There are two settings that may require tuning in TVMAGI. First, the number of discretization can affect the inference results. When observed components are sparse, the number of discretization should increase until the results are stabilized. However, over-densed discretization will lead to higher computation cost. For example, in SEIRD model, we set discretization as 2 data points per day for optimized performance, as further increasing the discretization will not improve the result accuracy. Second, the inference results on TVMAGI can be affected by hyper-parameter settings of the GP kernel for $\theta(t)$. To achieve the desired variability level of time-varying parameters, we find it helpful to use informative hyper-prior that specifies the range of length-scale (a.k.a. bandwidth) parameter of the GP kernel for $\theta(t)$ to prevent obvious over-smoothing or over-fitting.

The Gaussian process modeling of $\theta(t)$ with Matern kernel $\nu = 2.01$ ensures continuously differentiable time-varying parameters, which prevents overfitting the parameter to the observation noise. The variability in time-varying parameter $\theta(t)$ can be further controlled by the length-scale GP hyperparameter ϕ_2 through its hyper-prior. The Matern kernel together with the hyper-prior on the length-scale hyperparameter ensures the smoothness and the degree of variability in $\theta(t)$, which in turn prevents over-fitting and mitigates identifiability issues. If a more flexible $\theta(t)$ is desired, Matern kernel $\nu = 1.5$ with hyper-prior favoring smaller GP hyperparameter ϕ_2 can be used to allow rapid non-differentiable changes in $\theta(t)$.

One limitation of TVMAGI is its inherent bias. Just like any other Bayesian approaches, TVMAGI could be biased towards smoother curves due to the GP prior. The results of inference would be less accurate when the true time-varying parameters have rapid changes, and the posterior interval coverage could suffer. But as shown in the examples, the magnitude of such bias is small in practice, and our accuracy is still comparable with state-of-the-art approaches while TVMAGI having much better universal applicability. TVMAGI is also not suitable if the underlying time-varying parameter is a jump process. In this case, methods in change point detection literature might be more applicable [28]. Alternatively, we can place the prior of continuous-time Markov chain or Poisson process on $\theta(t)$, instead of Gaussian process, to model the jump process.

There are also many other interesting future directions for TVMAGI. We currently focus on empirical performance of TVMAGI through simulation examples. More theoretical study on the convergence property, identifiability issue, and asymptotic behavior of the time-varying parameter estimate are all natural directions of future research. It would be of future interest to extend TVMAGI for partial differential equation of spatial-temporal dynamics [29], or stochastic differential equation of inherent noise modeling [30].

Declarations

The authors have no financial or non-financial interests that are directly or indirectly related to this paper.

Appendix A Supplementary materials

A.1 Advantage of the infinity norm in Eq.(4.5)

In this section we illustrate how L_∞ norm in Eq.(4.5) of main text facilitates theoretical construction, compared to L_2 norm. First, with L_∞ norm in W , it is clear that on the discretization subset I , the corresponding W_I will simply be the maximum over I . However, with the L_2 -norm of $\int_0^T (X(t) - f(X(t), \Theta(t), \Psi, t))^2 dt$, the formulation of the corresponding W_I is not as clear. Second, using L_∞ makes the theoretical justification easier. To mathematically study the properties of TVMAGI while avoiding Borel paradox, one can use the fact that $\{W_I < \epsilon\} \equiv \cap_{i \in I} \{W_i < \epsilon\}$, thanks to W_I being the L_∞ norm over the set I . Third, the L_∞ norm in Eq.(4.5) and Eq.(4.6) automatically transforms into L_2 loss for likelihood calculation in Eq.(4.7) and Eq.(4.11) through a simple mathematical derivation, which facilitates computation while maintaining the theoretical rigor. This is because when a Gaussian distributed vector is constrained to have zero deviation with some fixed value (i.e., vector L_∞ distance to the fixed value is zero), the fixed value will be plugged into the Gaussian probability density function, inducing an L_2 loss in the target function Eq.(4.11).

A.2 Intuition of the manifold constraint on Gaussian process

In this section we illustrate the intuition behind the manifold constraint $W = 0$ (or $W_I = 0$) for the Gaussian process.

We consider the following simple bi-variate Gaussian example:

$$(Z_1, Z_2) \sim \mathcal{N}_2(0, \begin{pmatrix} 1 & 0.5 \\ 0.5 & 1 \end{pmatrix})$$


Let $W = Z_1 - 0.5 \times (Z_2^2 + 1)$. Then conditioning on $W = 0$, the (Z_1, Z_2) is distributed on the parabola as in the left panel of Figure A1. The sampling is possible as the blue dots in the right panel of Figure A1, where the density is proportional to the original bi-variate Gaussian but only on the parabola curve.

The Gaussian process X_I in the main text and the constraint $W_I = 0$ apply the same intuition on $|I|$ -dimensional Gaussian vector, thus having a manifold constraint induced by $W_I = 0$.

A.3 HMC algorithm

We outline the HMC procedure for sampling from a target probability distribution. Algorithm 1 provides the details of our HMC implementation.

Fig. A1 Illustration of the manifold constraint on bi-variate Gaussian. Left panel: blue dots are the samples from joint density of (Z_1, Z_2) without the manifold constraint $W = 0$. The red curve is the manifold constraint $W = 0$ to be imposed. Right panel: blue dots are the samples from $(Z_1, Z_2)|W = 0$ where all points lie on the parabola. The density is proportional to the original bi-variate Gaussian but only on the parabola curve.



Algorithm 1 HMC sampling in TVMAGI

Input:

U : Log likelihood function in Eq.(4.11)

ϵ : step size of HMC

L : number of leaf frog steps

N : number of samples

Initialize: $x(I), \theta(I), \psi, \sigma$

```

1: for  $i$  in 1:N do
2:    $q_{\text{current}} = \text{vector}(x(I), \theta(I), \psi, \sigma)$ 
3:    $q = q_{\text{current}}$ 
4:    $p = \text{rnorm}(\text{length}(q), 0, 1)$ 
5:    $p_{\text{current}} = p$ 
6:    $p = p - \epsilon \nabla U(q)/2$ 
7:   for  $j$  in 1:L do
8:      $q = q + \epsilon p$ 
9:      $p = p - \epsilon \nabla U(q)$ 
10:  end for
11:   $p = p - \epsilon \nabla U(q)/2$ 
12:  if  $\text{runif}(1) < \exp(U(q_{\text{current}}) - U(q) + \text{sum}(p_{\text{current}}^2 - p^2)/2)$  then
13:    return  $q$                                 ▷ (Accept)
14:  else
15:    return  $q_{\text{current}}$                         ▷ (Reject)
16:  end if
17: end for

```

A.4 Additional benchmark methods of Bayesian filtering approaches

Compared with Ensemble Adjustment Kalman Filter (EAKF), other Bayesian filtering methods, such as Extended Kalman Filter (EKF), Unscented Kalman Filter (UKF) and Ensemble Kalman Filter (EnKF) are less discussed in ODE parameter inference applications. In this section we also include a few more Bayesian filtering benchmark methods of Extended Kalman Filter (EKF), Unscented Kalman Filter (UKF), and EnKF. All the Bayesian filtering methods have the inherent limitation that all parameters must be assumed time-varying, and thus cannot accommodate time-constant parameters. To further illustrate the Bayesian filtering approaches, we also provide inference results using the other three methods. Fig.A2 and Fig.A3 illustrate the results of SEIRD model. Fig.A4 shows the results of LV model and Fig.A5 shows the results of HIV model.




Fig. A2 Inferred $\theta(t)$ of EKF, UKF and EnKF approach for SEIRD model. The inferred parameters cannot capture the ODE structure, thus yielding inaccurate reconstructed trajectory.




Fig. A3 Reconstructed trajectory of EKF, UKF and EnKF approaches for SEIRD model. The accuracy is far from satisfactory.




Fig. A4 Inferred parameter and reconstructed trajectory of EKF, UKF and EnKF approaches for LV model. The accuracy is far from satisfactory.




Fig. A5 Inferred parameter and reconstructed trajectory trajectory of EKF, UKF and EnKF approaches for HIV model. The accuracy is worse than TVMAGI, Runge-Kutta, or ELE.

A.5 Additional results of the main benchmark methods

Due to the limited space of the main text, we include visualizations of the main benchmark methods of Runge-Kutta method and EAKF method for the three examples in the main text here. Fig.A6 is for SEIRD model parameter inference, and Fig.A7 is for SEIRD model reconstructed trajectory. Fig.A8 is for LV model. Fig.A9 is for HIV model.




Fig. A6 Results of parameter inference in 100 simulated datasets using Runge-Kutta and EAKF for SEIRD model.




Fig. A7 Reconstructed trajectory using inferred parameters of Runge-Kutta and EAKF methods for SEIRD model. One sample simulation dataset is also presented to visualize the noise level and observation schedule.




Fig. A8 Comparison of inferred $\theta(t)$ and reconstructed $X(t)$ of LV model. We also plot one sample simulation dataset to visualize the noise level.




Fig. A9 Comparison of inferred $\theta(t)$ and reconstructed $X(t)$ of HIV model. We also plot one sample simulation dataset to visualize the noise level.

A.6 Review on benchmark methods

A.6.1 Runge-Kutta method

Runge-Kutta method is a brute-force way for parameter inference in ODE systems. As a non-linear least square method, Runge-Kutta method minimizes the MSE of observations and reconstructed trajectory using numerical integration from the proposed initial conditions and parameters. The objective function is given by:

$$\min_{\mathbf{x}_0, \boldsymbol{\theta}(\mathbf{I}), \boldsymbol{\psi}} \sum_{\tau \in \tau_d} \sum_{d=1}^D (y_d(\tau) - \mathbb{X}_{\tau}^{RK4}(\mathbf{x}_0, \boldsymbol{\theta}(t), \boldsymbol{\psi})_d)^2 \quad (\text{A1})$$

where $\mathbb{X}^{RK4}$ denotes the reconstructed trajectory using the 4th Order Runge-Kutta method.

A.6.2 Ensemble Adjustment Kalman Filter

Ensemble Adjustment Kalman Filter (EAKF) is a variation of Kalman Filter that is popular for parameter calibration of ODE systems in practice. It is a specially designed fast Bayesian filtering method. As a data assimilation technique, EAKF represents filtered distribution using Monte Carlo samples, and replaces the covariance matrix with sample covariance. The Kalman update assumes all probability distributions involved are Gaussian. As a major difference with Ensemble Kalman Filter (EnKF), EAKF uses a deterministic update instead of stochastic update.

A.7 Limitations of Runge-Kutta method

Numerical integration methods are the gold standard for the ODE parameter inference when all parameters are time-constant. However, with the presence of time-varying parameters, there are several inherent disadvantages of numerical integration methods. First and foremost, without any structure on the time-varying parameters, the numerical integration method will give time-varying parameter estimate that perfectly fits the observation data, resulting in overfitting issues in the time-varying parameter. Second, with the increase of time points and size of the system, the objective function becomes expensive to evaluate, resulting in high computation cost. Third, the numerical methods are sensitive to the initial value of the optimization, while searching for a good initial point can be challenging in the high-dimensional scenarios, as optimization of objective functions using algorithms such as Adam can be easily stuck at local minimum. We point out that using random initial values in our examples can lead to high level of error, making numerical integration methods completely fail. In this case, all the optimization for Runge-Kutta method are initialized at the TVMAGI initial points from Section 5.1 in the main text examples.

A.8 Limitations of Bayesian filtering methods

Although Bayesian filtering method is the fastest, examples have shown that applying Bayesian filtering methods to ODE parameter inference problems has failed to provide satisfactory results. Even though we included the Bayesian filtering methods as baseline comparison methods, we emphasize that state-space model (Bayesian filtering) and ODE parameter inference (TVMAGI) are fundamentally different problems. To illustrate the difference in a simplified framework from a theoretical perspective, consider the following example of time-constant parameter inference where all parameters are denoted as θ .

The fundamental difference is the lack of randomness in the state transition given the model parameters, and thus the Bayesian update given the observation will have zero effect. With the ODE structure, there is no randomness in $\mathbf{x}_t|\mathbf{x}_{t-1}, \theta$. The state transition distribution $p(\mathbf{x}_t|\mathbf{x}_{t-1}, \theta)$ essentially has shifted Dirac delta distribution. Therefore, regardless of emission probability $p(\mathbf{y}_t|\mathbf{x}_t)$, the hidden state $\mathbf{x}_t$ will not depend on $\mathbf{y}_t$. As such, all Bayesian updates will have zero effect to shift distribution of $p(\mathbf{x}_t|\mathbf{x}_{t-1}, \theta)$, which is still shifted Dirac delta distribution. The exact Bayesian filtering results will simply be the solution of ODE dynamics given the initial sample of $\mathbf{x}_0$ and the parameter θ . In this case, the parameter estimation in the exact Bayesian filtering reduces to using numerical solver to generate the entire ODE curve given $\mathbf{x}_0, \theta$, and then using a least square approach to compare the solved curve and the observations to find the best $\mathbf{x}_0, \theta$. The exact Bayesian filtering in this case degenerates to a numerical integration method. From another particle filter perspective, each particle of sampled $\mathbf{x}_0, \theta$ will evolve in time according to ODE without any randomness, and the parameter estimation becomes a brute-force search of the particle of sampled $\mathbf{x}_0, \theta$ that provides the smallest mean squared error to the observation. In light of this, the Bayesian filtering/smoothing is more suitable for inference of Stochastic Differential Equations (SDEs) parameters.

However, Bayesian filter methods still can be applied in ODE inference problem if we can forego some statistical rigor. We can treat **all** parameters as time-varying, and artificially introduce additional randomness in $\theta_t|\theta_{t-1}$ and $\psi_t|\psi_{t-1}$. In this case, the simultaneous estimation of system components $\mathbf{x}(t)$, time-constant parameter ψ and time-varying parameter $\theta(t)$ becomes an estimation of joint hidden state $(\mathbf{x}_t, \psi_t, \theta_t)$. Then Bayesian filtering methods such as EKF, UKF, EnKF and EAKF become applicable. However, this is not a statistically principled approach, because (1) time-constant parameter ψ is now changing with time, and (2) the inference on the distribution is not exact as Gaussian distributional approximation is used on system components $\mathbf{x}(t)$. Nevertheless, this method could work empirically, with the notable success of SIRS-EAKF model.

In our example, we found the Bayesian filter methods are highly sensitive to the initialization of the parameters. Randomized initialization often lead to insensible results. Therefore, all the parameters for Bayesian filter methods are initialized at the TVMAGI initial points from Section 5.1 in the main text examples, which is the same as Runge-Kutta method.

Even then, our numerical experiments suggest that Bayesian filtering methods are the fastest, but often yield unreliable inference results. Among EKF, UKF, EnKF and EAKF, we see that UKF, EnKF and EAKF yield similar results, all outperforming EKF in LV and HIV examples, while EKF has a slight advantage in SEIRD model inference compared with other Bayesian filtering methods, although not robust. All of them yield orders of magnitude worse trajectory RMSE compared to TVMAGI or Runge-Kutta. Two factors contribute to the unsatisfactory trajectory RMSE of Bayesian filtering methods. First, all filtering approaches fail to yield an accurate and robust parameter estimation, especially in SEIRD model when the observation points are limited. Second, the variance of time-varying parameters is large at weakly identifiable time points (Figure A4), and the estimates for the later time points are no longer accurate due to the cascading effect.

The failure of Bayesian filtering in our setting is not surprising. First, it violates the assumption of the model, as it is not principled to allow time-constant parameters ψ to change over time. Although we used the average $\bar{\psi}$ of the filtered parameter ψ_t to be the final estimate for the ψ , the approximation error can still be large. The idea of changing a time-constant parameter to be time-varying during inference and later plugging in the average of the inferred values is not theoretically sound. The trajectory RMSE precisely evaluates how accurate the estimated parameters can be used to reconstruct the entire system given the ODE structure, of which $\bar{\psi}$ would fail. Second, contrary to the typical setting of Bayesian filtering in machine learning where there is a long sequence of data, our experiments are designed to see how the method performs with short time series and sparse observations, as in most scientific experiment settings. The lack of long series of data poses a challenge to Bayesian filtering methods. Third, the ODE structure is no longer exactly followed in Bayesian filtering, which loosens the structure constraints and creates additional loss of information from the observation data that is already sparse.

A.9 Additional results of TVMAGI interval estimates

In this section we present the visualization for the interval estimation results. We see that for long time series of observations, the estimated intervals tend to be narrow and may not contain the true values, which is a limitation of TVMAGI.




Fig. A10 Illustration of interval estimation of SEIRD model of 10 sample datasets for β (upper), v^e (middle), p^d (lower). The shadow indicates the 95% posterior interval from HMC samples, the solid lines indicate the posterior mean, and the red surface indicates the true value.




Fig. A11 Median of upper & lower bound of β , v^e and p^d in SEIRD model. We randomly plot 10 replications.




Fig. A12 Estimated interval of $\alpha(t)$ and $\gamma(t)$ in LV model. We randomly plot 10 replications.


(a) Interval estimation of $a(t)$ (b) Median upper and lower bound of $a(t)$ **Fig. A13** Estimated interval of $a(t)$ in HIV model. We randomly plot 10 replications.

Table A1 Parameter and trajectory RMSE of TVMAGI in SEIRD model under different discretization level. The computing time (in seconds) is reported in the last row.

RMSE		Discretization Level		
		1	2	4
Trajectory/Parameter	β	0.156 ± 0.083	0.114 ± 0.039	0.102 ± 0.036
	v^e	0.007 ± 0.007	0.009 ± 0.010	0.010 ± 0.010
	v^i	0.006 ± 0.004	0.005 ± 0.003	0.005 ± 0.003
	p^d	0.018 ± 0.027	0.019 ± 0.029	0.021 ± 0.033
	S	1253.4 ± 332.0	581.7 ± 272.1	603.9 ± 355.2
	E	1010.6 ± 231.9	704.7 ± 218.3	679.5 ± 209.1
	I	584.6 ± 195.6	439.0 ± 140.4	401.0 ± 132.8
	D	44.7 ± 8.8	38.3 ± 4.9	39.0 ± 3.4
computing time (s)		622.1 ± 49.6	1013.3 ± 109.1	1773.9 ± 161.3

A.10 Additional results on TVMAGI sensitivity study

In this section we conduct three sensitivity analysis: the number of discretization, the selection of GP kernel, and the mis-specified time-varying parameter. For the number of discretization, our theoretical derivation ensures that the inference result will converge as the discretization increase, and we recommend gradually increasing the number of discretization points until the result is stabilized. Here we empirically demonstrate such convergence by presenting results with various discretization level. For the GP kernel selection, we relax the GP kernel of time-varying parameter to be Matern kernel with different degrees of freedom $\nu = 2.5$ and $\nu = 1.5$. When $\nu = 1.5$, $\theta(t)$ is only required to be continuous, but not necessarily differentiable. For the mis-specified time-varying parameter, we examine the TVMAGI time-varying estimate when one parameter is in fact time-constant.

A.11 Number of discretization

In this section we explore the sensitivity of TVMAGI to the number of discretization. In the paper we used discretization level of 2 in the SEIRD model, that is, with 32 observation points, we have a total of 64 discretization points (2 discretization per observation). For comparison, we use the same observation data with discretization level of 1 and level of 4, corresponding to 32 discretization points (1 discretization per observation) and 128 discretization points (4 discretization per observation), respectively. Table A1 shows that the inference accuracy indeed converges. However, the computational time scales up linearly with the number of discretization points. In practice, we recommend gradually increasing the number of discretization points until the result is stabilized, trying to balance the inference accuracy with the computing cost.

Table A2 Parameter and trajectory RMSE of TVMAGI in SEIRD using different kernels.

RMSE		ν for $\theta(t)$		
		1.5	2.5	2.01
TrajectoryParameter	β	0.085 ± 0.005	0.139 ± 0.021	0.114 ± 0.039
	v^e	0.012 ± 0.004	0.015 ± 0.005	0.009 ± 0.010
	v^i	0.004 ± 0.005	0.003 ± 0.004	0.005 ± 0.003
	p^d	0.058 ± 0.004	0.044 ± 0.012	0.019 ± 0.029
	S	423.0 ± 130.2	853.3 ± 288.5	581.7 ± 272.1
	E	412.3 ± 111.7	836.1 ± 186.5	704.7 ± 218.3
	I	356.3 ± 94.1	569.4 ± 171.8	439.0 ± 140.4
	D	69.1 ± 21.5	20.4 ± 8.6	38.3 ± 4.9

A.12 Selection of kernel


In this section we discuss how the kernel selection will affect the performance of TVMAGI. In the paper we recommend modeling $\theta(t)$ as Gaussian process with Matern kernel $\nu = 2.01$ to guarantee a continuous and differentiable time-varying parameters while maintaining high flexibility. We can also use other GP kernels or hyperparameters to control the smoothness. For example, Matern kernel with $\nu = 2.5$ can be used for even smoother GP with a simple closed-form kernel. The condition of differentiability can also be further relaxed if we substitute the kernel with $\nu = 1.5$, and then the parameters are only assumed with continuity without differentiability, allowing more flexible patterns for the time-varying parameters $\theta(t)$. Table A2 shows the result under both kernels, where the performance is similar to the recommended $\nu = 2.01$ in SEIRD model, indicating that TVMAGI is not sensitive to the choice of kernels.

A.13 Mis-specified time-varying parameters

In this section, we explore the TVMAGI estimation when a time-constant parameter is mis-specified to be time-varying, i.e., a time-constant ODE parameter is falsely recognized as a time-varying parameter. Ideally, the inferred parameter curve from TVMAGI should be a horizontal line which is close to the true constant value, while still maintaining smoothness. In this example, we alter the settings of SEIRD model by setting parameter v^e as constant $v^e = 0.1$, and treat it as time-varying parameter in the TVMAGI estimation. As shown in Figure A14, TVMAGI is capable of dealing with the mis-specified parameters, where the inferred v^e is approximately a horizontal line close to true parameter value (except the ending time when the v^e is difficult to estimate), while the inference of other parameters are not affected. The parameter and reconstructed RMSE results are presented in Supplementary Material


Table 1, and the accuracy is comparable to the Table 1 in main text where all time-varying parameters are correctly specified.

Fig. A14 Mis-specified parameter v^e in SEIRD model. Figure y-scale is the same as Figure 1 for better visualization.




A.14 Discretization and choice of kernels

In this section we give additional visualizations for TVMAGI sensitivity study when varying the discretization level, and Matern kernel degree of freedom ν .




(a) Discretization level = 1




(b) Discretization level = 4

Fig. A15 TVMAGI inferred parameters in SEIRD model with different discretization level.



(a) Discretization level = 1



(b) Discretization level = 4

Fig. A16 TVMAGI reconstructed trajectories in SEIRD model with different discretization level.**Table A3** Parameter and reconstructed RMSE table for mis-specified time-varying parameter v^e .



RMSE				
Parameter		Reconstructed		
β	0.120 ± 0.046	S	513.13 ± 232.08	
v^e	0.008 ± 0.009	E	647.16 ± 232.81	
v^i	0.007 ± 0.004	I	364.94 ± 197.92	
p^d	0.022 ± 0.036	D	27.61 ± 6.52	

A.15 Structure of HIV model

The ODE model that characterizes the response of anti-viral regimens during HIV infection is given by:

$$\begin{cases} T'(t) \equiv \frac{d}{dt}T(t) = \lambda - \rho T(t) - k[1 - r(t)]T(t)X(t) \\ T^{*'}(t) \equiv \frac{d}{dt}T^*(t) = k[1 - r(t)]T(t)X(t) - \delta T^*(t) \\ X'(t) \equiv \frac{d}{dt}X(t) = N\delta T^*(t) - cX(t) \end{cases} \quad (\text{A2})$$

$T(t)$ denotes the concentration of uninfected CD4+ T cells, which can be accurately measured clinically; $T^*(t)$ denotes the unknown unobservable concentration of infected T cells; $X(t)$ is the HIV-1 viral load in plasma, and can

(a) Inferred parameters of $\nu = 1.5$ (b) Reconstructed trajectory of $\nu = 1.5$ **Fig. A.17** TVMAGI inferred parameters and reconstructed trajectories of SEIRD model with Matern kernel $\nu = 1.5$

be observed with noise. λ is the rate of new T cell generation; ρ is the death rate of T cells; k is the infection rate of T cells by HIV virus; δ is the death rate of infected cells; N is the total production of new virions by an infected T cell; c denotes the known constant rate of free virion clearance; $r(t)$ is the time-varying antiviral drug efficacy coefficient, which may decay through time due to drug resistance. Our simulation settings are based on as $\lambda = 36$, $\rho = 0.108$, $k = 5 \times 10^{-4}$, $\delta = 0.1$, $N = 1000$, $c = 3.5$, $X(0) = 1000$, $T(0) = 350$, $T^*(0) = 20$, and $r(t) = \cos(\pi t/500)$. Time horizon is set as 100 days, with observation noise level at 5%. Hulin Wu (2008) transformed the system Eq.(A2) into Eq.(7.17) by taking $d = 2$, $a_1(t) = -NT^*(t)$, $a_2(t) = Nk[1 - r(t)]X(t)$, $Z_1(t) = 1$ and $Z_2(t) = T(t)$, and then used their ELE method estimate time-varying coefficients $a_i(t)$.

A.16 Advantages of multi-stage algorithm

Compared with joint optimization of hyperparameters and parameters together, the multi-stage optimization method enjoys several advantages. First, the GP hyperparameters Φ^X for the system components are set at the first stage and held as constant in the rest of the optimization so that the inverse of kernel matrix only needs to be computed once. Second, GP hyperparameters Φ^Θ for the time-varying parameters could not be set without any information about $\Theta(I)$. Therefore, a multi-stage procedure is necessary, where a point-wise $\tilde{\theta}(I)$ is obtained in one stage without GP, and then GP hyperparameters Φ^Θ is estimated in the following stage based on $\tilde{\theta}(I)$. Lastly, The multi-stage optimization ensures that each step of the optimization starts with sensible




Fig. A18 TVMAGI inferred parameters and reconstructed trajectories of SEIRD model with Matern kernel $\nu = 2.5$

initial value obtained from previous modularized optimization, thus drastically decreasing the chance of Adam optimizer stuck in local mode. Experiments have shown that our carefully designed multi-stage optimization is faster and achieves better results than joint optimization with randomized starting values.

However, although we carefully designed the multi-stage optimization and sampling schedule, occasionally the Adam optimizer or the HMC sampler could still get stuck. Among the total of 300 simulated datasets across three examples, the algorithm got stuck in one particular dataset of the SEIRD model. In the stuck case, some manual tuning of the hyper-parameters or jittering of the sampled parameters might be needed. We will continue to improve the robustness of our proposed algorithm and our software implementation.

References

- [1] Li, M.Y., Muldowney, J.S.: Global stability for the seir model in epidemiology. *Mathematical Biosciences* **125**(2), 155–164 (1995). [https://doi.org/10.1016/0025-5564\(95\)92756-5](https://doi.org/10.1016/0025-5564(95)92756-5)
- [2] Takeuchi, Y., Du, N.H., Hieu, N.T., Sato, K.: Evolution of predator–prey systems described by a lotka–volterra equation under random environment. *Journal of Mathematical Analysis and Applications* **323**(2), 938–957 (2006). <https://doi.org/10.1016/j.jmaa.2005.11.009>
- [3] Perelson, A.S., Neumann, A.U., Markowitz, M., Leonard, J.M., Ho, D.D.: Hiv-1 dynamics in vivo: virion clearance rate, infected cell life-span, and

- viral generation time. *Science* **271**(5255), 1582–1586 (1996)
- [4] Lapidus, L., Seinfeld, J.H.: Numerical Solution of Ordinary Differential Equations. Academic press, ??? (1971)
 - [5] Abbott, S., Hellewell, J., Thompson, R.N., Sherratt, K., Gibbs, H.P., Bosse, N.I., Munday, J.D., Meakin, S., Doughty, E.L., Chun, J.Y., *et al.*: Estimating the time-varying reproduction number of sars-cov-2 using national and subnational case counts. *Wellcome Open Research* **5**(112), 112 (2020)
 - [6] Miao, H., Xia, X., Perelson, A.S., Wu, H.: On identifiability of nonlinear ode models and applications in viral dynamics. *SIAM review* **53**(1), 3–39 (2011)
 - [7] Dondelinger, F., Husmeier, D., Rogers, S., Filippone, M.: Ode parameter inference using adaptive gradient matching with gaussian processes. In: *Artificial Intelligence and Statistics*, pp. 216–228 (2013). PMLR
 - [8] Yang, S., Wong, S.W.K., Kou, S.C.: Inference of dynamic systems from noisy and sparse data via manifold-constrained gaussian processes. *Proceedings of the National Academy of Sciences* **118**(15) (2021) <https://www.pnas.org/content/118/15/e2020397118.full.pdf>. <https://doi.org/10.1073/pnas.2020397118>
 - [9] Wenk, P., Gotovos, A., Bauer, S., Gorbach, N.S., Krause, A., Buhmann, J.M.: Fast gaussian process based gradient matching for parameter identification in systems of nonlinear odes. In: Chaudhuri, K., Sugiyama, M. (eds.) *Proceedings of the Twenty-Second International Conference on Artificial Intelligence and Statistics. Proceedings of Machine Learning Research*, vol. 89, pp. 1351–1360. PMLR, ??? (2019)
 - [10] Wu, H.: Statistical methods for hiv dynamic studies in aids clinical trials. *Statistical methods in medical research* **14**(2), 118–134 (2005). <https://doi.org/10.1191/0962280205sm390oa>
 - [11] Li, L., Brown, M., Lee, K., Gupta, S.: Estimation and inference for a spline-enhanced population pharmacokinetic model. *Biometrics* **58**(3), 601–11 (2002). <https://doi.org/10.1111/j.0006-341x.2002.00601.x>
 - [12] Huang, Y., Liu, H. D.and Wu: Hierarchical bayesian methods for estimation of parameters in a longitudinal hiv dynamic system. *Biometrics* **62**(2), 413–423 (2006). <https://doi.org/10.1111/j.1541-0420.2005.00447.x>
 - [13] Cao, J., Huang, J.Z., Wu, H.: Penalized nonlinear least squares estimation of time-varying parameters in ordinary differential equations. *Journal of computational and graphical statistics* **21**(1), 42–56 (2012)

- [14] Pei, S., Shaman, J.: Initial simulation of sars-cov2 spread and intervention effects in the continental us. *MedRxiv* (2020)
- [15] Shaman, J., Karspeck, A.: Forecasting seasonal outbreaks of influenza. *Proceedings of the National Academy of Sciences* **109**(50), 20425–20430 (2012)
- [16] Schmidt, J., Krämer, N., Hennig, P.: A probabilistic state space model for joint inference from differential equations and data. *arXiv preprint arXiv:2103.10153* (2021)
- [17] Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014)
- [18] Neal, R.M.: MCMC Using Hamiltonian Dynamics. In: Brooks, S., Gelman, A., Jones, G., Meng, X.L. (eds.) *Handbook of Markov Chain Monte Carlo*. Chapman & Hall/CRC Handbooks of Modern Statistical Methods, pp. 113–162. CRC Press, ??? (2011). Chap. 5
- [19] Yang, W., Kandula, S., Huynh, M., Greene, S.K., Van Wye, G., Li, W., Chan, H.T., McGibbon, E., Yeung, A., Olson, D., et al.: Estimating the infection fatality risk of covid-19 in new york city, march 1-may 16, 2020. *MedRxiv* (2020)
- [20] Hethcote, H.W.: The mathematics of infectious diseases. *SIAM review* **42**(4), 599–653 (2000)
- [21] Hao, X., Cheng, S., Wu, D., Wu, T., Lin, X., Wang, C.: Reconstruction of the full transmission dynamics of covid-19 in wuhan. *Nature* **584**(7821), 420–424 (2020)
- [22] Dong, E., Du, H., Gardner, L.: An interactive web-based dashboard to track covid-19 in real time. *The Lancet infectious diseases* **20**(5), 533–534 (2020)
- [23] Mwalili, S., Kimathi, M., Ojiambo, V., Gathungu, D., Mbogo, R.: Seir model for covid-19 dynamics incorporating the environment and social distancing. *BMC Research Notes* **13**(1), 1–5 (2020)
- [24] Goel, N.S., Maitra, S.C., Montroll, E.W.: On the volterra and other nonlinear models of interacting populations. *Reviews of modern physics* **43**(2), 231 (1971)
- [25] Donald, D.B., Stewart Anderson, R.: Resistance of the prey-to-predator ratio to environmental gradients and to biomanipulations. *Ecology* **84**(9), 2387–2394 (2003)

- [26] Chen, J., Wu, H.: Efficient local estimation for time-varying coefficients in deterministic dynamic models with applications to HIV-1 dynamics. *Journal of the American Statistical Association*, 103:481, 369-384 (2008). <https://doi.org/10.1198/016214507000001382>
- [27] Huang, Y., Rosenkranz, S.L., Wu, H.: Modeling hiv dynamics and antiviral response with consideration of time-varying drug exposures, adherence and phenotypic sensitivity. *Mathematical biosciences* **184**(2), 165–186 (2003)
- [28] Cuenod, C.-A., Favetto, B., Genon-Catalot, V., Rozenholc, Y., Samson, A.: Parameter estimation and change-point detection from dynamic contrast enhanced mri data using stochastic differential equations. *Mathematical biosciences* **233**(1), 68–76 (2011)
- [29] Xun, X., Cao, J., Mallick, B., Maity, A., Carroll, R.J.: Parameter estimation of partial differential equation models. *Journal of the American Statistical Association* **108**(503), 1009–1020 (2013)
- [30] Kou, S.C., Xie, X.S.: Generalized langevin equation with fractional gaussian noise: subdiffusion within a single protein molecule. *Physical review letters* **93**(18), 180603 (2004)