

GaPT: Gaussian Process Toolkit for Online Regression with Application to Learning Quadrotor Dynamics

Francesco Crocetti^{1,2*}, Jeffrey Mao^{1*}, Alessandro Saviolo¹, Gabriele Costante², and Giuseppe Loianno¹

Abstract— Gaussian Processes (GPs) are expressive models for capturing signal statistics and expressing prediction uncertainty. As a result, the robotics community has gathered interest in leveraging these methods for inference, planning, and control. Unfortunately, despite providing a closed-form inference solution, GPs are non-parametric models that typically scale cubically with the dataset size, hence making them difficult to be used especially on onboard Size, Weight, and Power (SWaP) constrained aerial robots. In addition, the integration of popular libraries with GPs for different kernels is not trivial. In this paper, we propose GaPT, a novel toolkit that converts GPs to their state space form and performs regression in linear time. GaPT is designed to be highly compatible with several optimizers popular in robotics. We thoroughly validate the proposed approach for learning quadrotor dynamics on both single and multiple input GP settings. GaPT accurately captures the system behavior in multiple flight regimes and operating conditions, including those producing highly nonlinear effects such as aerodynamic forces and rotor interactions. Moreover, the results demonstrate the superior computational performance of GaPT compared to a classical GP inference approach on both single and multi-input settings especially when considering large number of data points, enabling real-time regression speed on embedded platforms used on SWaP-constrained aerial robots.

SUPPLEMENTARY MATERIAL

Code: <https://github.com/arplaboratory/GaPT>

Video: <https://youtu.be/Bi5x2sQcW3s>

I. INTRODUCTION

Gaussian Processes (GPs) have become popular models in a wide range of robotics inference [1], [2], [3], planning [4], and control problems [5], [6], [7]. However, their inference time scales cubically with respect to the dataset size, thus making these models unsuitable for many real-time robotics problems, especially on Size, Weight, and Power (SWaP) constrained aerial robots. This motivates the need to develop novel approximation techniques that can perform inference at a much higher rate without sacrificing the accuracy of the original model. In this paper, we present GaPT, a GP

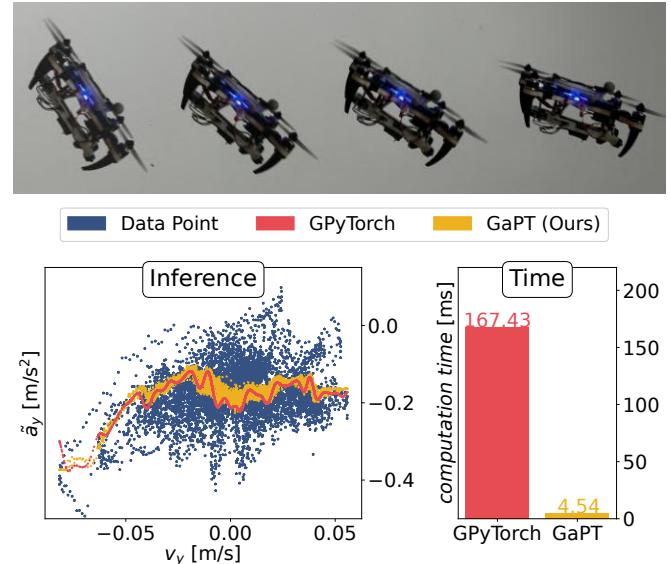


Fig. 1: Comparison between GaPT and a classical GP inference approach (GPyTorch) for learning quadrotor residual dynamics during agile flight on 6000 test data points.

Toolkit that converts a GP into a linear state space form. This form processes information sequentially to learn the data trend allowing the computation to scale linearly with the dataset instead. GaPT can easily instantiate any kernel or combination of them in C++ with minor or no user effort. The approach does not perform the expensive matrix inversions or multiplications of a typical GP and is easily portable to SWaP-constrained robots.

Our main contributions are the following. First, we design and present GaPT, a GP Toolkit that can perform inference at a much faster rate than the standard inference model used in popular libraries such as [8] by translating GPs into a linear state space form. The linear state space form or State Space Gaussian Process (SSGP) scales linearly with respect to the dataset size. This allows quicker data processing and overcomes one of the major hurdles related to real-time robotics applications, especially when employing large datasets on SWaP constrained aerial robots. In fact, the histogram in Fig. 1 shows the superior computational advantages of GaPT which requires $\times 37$ less time compared to the GP baseline, while also producing similar inference results when estimating external disturbances. Second, GaPT includes the state space realization for Single Input Single Output (SISO) as well as for Multiple Input Single Output (MISO) robotics systems. By expanding the input space to multi-dimensions, GaPT better captures the high-

*These authors contributed equally and listed alphabetic order.

¹The authors are with the New York University, Tandon School of Engineering, Brooklyn, NY 11201, USA. email:{fc2353, jm7752, as16054, loiannog}@nyu.edu.

²The authors are with the Department of Engineering, University of Perugia, 06125 Perugia, Italy. email:{francesco.crocetti, gabriele.costante}@unipg.it.

This work was supported in part by the NSF CAREER Award 2145277, the DARPA YFA Grant D22AP00156-00, the Technology Innovation Institute, Qualcomm Research, Nokia, and NYU Wireless. Giuseppe Loianno serves as consultant for the Technology Innovation Institute. This arrangement has been reviewed and approved by the New York University in accordance with its policy on objectivity in research.

frequency characteristic of the true signal, outperforming single input GPs on unfiltered real-world data, as experimentally validated in our results. Third, we provide a variety of ready-to-use kernel implementations such as RBF, Periodic, and Matern kernels. Fourth, we validate our approach on real-world learning quadrotor dynamics inference problems to capture the system behavior in multiple flight regimes and operating conditions, including those producing highly nonlinear effects such as aerodynamic forces and torques, and rotor interactions. Finally, our method includes a quick training routine and can easily export C++ ready functions that can be integrated into several optimization libraries (e.g., [9], [10], [11]) often used on small-scale aerial [1] and ground [2] robots. Our approach and its open source implementation will contribute to reducing the entry barrier to the adoption of these approaches avoiding the complex integration and computation time required by several libraries, especially for SWaP robots.

II. RELATED WORKS

In literature, GPs have been applied to learning residual or full system dynamics [1], [2], [3], reward functions [12], planning [4], adaptive control [13], or learning safety regions for repeated tasks such as in reinforcement learning [5], [6]. These methods are often implemented by hard coding the full GP model in C++ such as in [1], [2]. In the realm of Bayesian Machine learning, Python libraries, such as GPytorch [8] or GPyflow [14], are useful for training hyperparameters and inference. These libraries leverage PyTorch or TensorFlow systems respectively for back propagation and GPU integration allowing the easy implementation of new kernels or more complex GP models [15], [16], but are complicated to interface with robotics optimization solvers [9], [10], [11].

Being non-parametric, the inference computation time of GPs generally scales with $O(n^3)$ with n the dataset size even when using the aforementioned libraries although the corresponding processing time is internally mitigated by leveraging GPU parallelism. The process of accelerating GPs is a major field of study. Data dimensionality reduction techniques [17], [18], [19] attempt to reduce the dataset size to a fraction of the points while keeping similar predictive performance. However, by performing data selection, a GP can never retain its full expressiveness especially if a large amount of data has to be discarded especially if converting inference time from a cubic order to a linear order a correspondingly large amount of data must be discarded. Our current method sidesteps this issue by converting the GP to an approximate state space form which only requires $O(n)$ time complexity allowing us to perform rapid inference without discarding data. This approach is inspired by recent signal processing results [20], [21]. However, [20] derives the state space form but only for a SISO system and is tailored for signal processing problems. In addition, it is very difficult to obtain a step-by-step formulation and portable solution tailored for small-scale robots. For further computation time increases among state space models, [21] also proposes a framework for SSGPs, but the code relies heavily on pre-

computing the Kalman filter predictions and updates steps to achieve these computation time gains making it less suitable for online learning. The aforementioned SSGP works are also not formulated for processing multiple inputs in the GPs and only consider single outputs with the exception of [22]. However, this approach is difficult to derive, implement, and practically use for dimensions beyond two. GaPT is easy to derive, and implement and has been experimentally validated. Finally, our framework seeks to promote the use of additional kernels by implementing several variants and provides a clear approach to enable any user to derive with minor effort additional ones based on the research needs.

III. GAUSSIAN PROCESS REGRESSION

Considering a multi dimensional input $\mathbf{x} \in \mathbb{R}^n$, a Gaussian Process (GP) is stochastic process $f(\mathbf{x})$ where any collection of random variables has a joint Gaussian Distribution $[f(\mathbf{x}_1), f(\mathbf{x}_2), \dots, f(\mathbf{x}_n)] \sim \mathcal{N}(\boldsymbol{\mu}, \mathbf{K})$. A GP is commonly defined as introducing the kernel formalism [23]

$$f(\mathbf{x}) \sim GP(\boldsymbol{\mu}(\mathbf{x}), k(\mathbf{x}, \mathbf{x}')), \quad (1)$$

where $\boldsymbol{\mu}(\mathbf{x}) : \mathbb{R}^n \rightarrow \mathbb{R}$ and $k(\mathbf{x}, \mathbf{x}') : \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}$ represent the mean and the covariance functions respectively. For a general regression problem, given a set of training data $\mathcal{D} = \{\mathbf{X}, \mathbf{y}\}$, with $\mathbf{X} = \{\mathbf{x}_i \in \mathbb{R}^d\}_{i=1}^n$ and $\mathbf{y} = \{y_i \in \mathbb{R}\}_{i=1}^n$, the interest is finding a function $\mathbf{f} : \mathbb{R}^n \rightarrow \mathbb{R}$ so that given an unknown input $\mathbf{x}_* \in \mathbb{R}^d$ the predicted value $y_* \in \mathbb{R}$ is reliable. In GP regression, $f(\mathbf{x})$ are assumed sampled from a zero mean Gaussian process (see eq. (1)) and the observations are assumed corrupted by additive Gaussian noise [22], [24]

$$\text{GP Prior} \quad f(\mathbf{x}_*) \sim GP(0, k(\mathbf{x}, \mathbf{x}'; \boldsymbol{\theta})) \quad (2)$$

$$\text{Measurements} \quad y_i = f(\mathbf{x}_i) + \epsilon, \quad \epsilon \sim \mathcal{N}(0, \sigma_{\text{noise}}^2) \quad (3)$$

where $\boldsymbol{\theta}$ represents the kernels hyper-parameters (length-scales, magnitude, period length, etc) which could be inferred from the marginal likelihood $\mathcal{N}(\mathbf{y} | \mathbf{0}, \mathbf{K}_{xx})$ maximizing the log-marginal likelihood

$$p(\mathbf{y} | \mathbf{X}, \boldsymbol{\theta}) = -\frac{1}{2} \log |\mathbf{K}| - \frac{1}{2} \mathbf{y}^\top \mathbf{K}^{-1} - \frac{n}{2} \log 2\pi, \quad (4)$$

where $\mathbf{K} = k(\mathbf{X}, \mathbf{X}, \boldsymbol{\theta}) + \sigma_{\text{noise}}^2 \mathbf{I}$ is the covariance matrix for the targets \mathbf{y} . The main advantage of employing GPs as regression models is the analytically tractable and close form solution of the predictive distribution [23] which is a Gaussian distribution with mean and variance

$$\begin{aligned} \mathbf{k}_* &= [k(\mathbf{x}_*, \mathbf{x}_0) \ k(\mathbf{x}_*, \mathbf{x}_1) \ \dots \ k(\mathbf{x}_*, \mathbf{x}_n)]^\top, \\ \mu_i(f(\mathbf{x}_*)) &= \mathbf{k}_*^\top (\mathbf{K} + \sigma_{\text{noise}}^2 \mathbf{I})^{-1} \mathbf{y}, \\ \mathbb{V}[f(\mathbf{x}_*)] &= k(\mathbf{x}_*, \mathbf{x}_*) - \mathbf{k}_*^\top (\mathbf{K} + \sigma_{\text{noise}}^2 \mathbf{I})^{-1} \mathbf{k}_*. \end{aligned} \quad (5)$$

Consequently, one of the main drawbacks of using Gaussian processes is the computational complexity, which can be limiting in applications where a large dataset is required to learn a complex function. The mean and variance evaluations (see Eq. (5)) are ruled by matrix inversions. Even given a pre-computed inverse it still takes $O(n^3)$ due to the multiplication which limits the amount of data usable in time-sensitive tasks such as robot control.

IV. METHODOLOGY

We describe a hierarchical approach to converting stationary kernel into its state space form for MISO systems and present a clear derivation of the state space form for different kernels useful in robotics.

A. Dual GP representation for stationary kernels

Performing a proper statistical inference often requires a large number of samples. This clashes with the high computational complexity required by GP-based inference. However, for stationary kernels, the Gaussian Process can be thought of as a solution to an approximate m -th order Stochastic Differential Equation (SDE) [25]. The Linear State Space Model (LSSM) representation of the system allows us to consider the statistical inference problem as a state estimation task that can be solved with an equivalent Kalman filter. The computational cost for a single evaluation depends on the m -order of the system $O(m^3)$, and if the number of the observations $n \gg m$, the influence of m^3 can be neglected obtaining $O(nm^3) \sim O(n)$. m in this case is a user-controlled hyperparameter maps to a Taylor Series approximation order of the Kernel. The approach represents an advantageous alternative to the classical GP inference method (with $O(n^3)$ cost). Therefore, this framework takes into account some covariance functions of a stationary process (Radial Basis kernel, Matern Kernel, Periodic Kernel), where the covariance functions are input invariant to translations (i.e., $k(\mathbf{x}, \mathbf{x}') = k(\Delta)$ where $\Delta = \mathbf{x} - \mathbf{x}'$). All stationary kernels can be directly derived by their spectral density [26], [27]. Because all stationary kernel functions are positive definite and stationary by the GP definition, we can apply Bochner's theorem as below to form a Fourier dual with the covariance function and its power spectrum.

Theorem 4.1 (Bochner): A continuous stationary function is positive definite if and only if k is the Fourier transform of a finite positive measure

$$k(\Delta) = \int \mathbf{S}(\omega) e^{2\pi i \omega \Delta} d\omega, \quad (6)$$

where \mathbf{S} is a positive finite measure.

If $k(\Delta)$ has a density $\mathbf{S}(\omega)$ then \mathbf{S} is known as the spectral density (or power spectrum) corresponding to k . In this case, the covariance function and spectral density are Fourier duals [28]

$$\mathbf{S}(\omega) = \int k(\Delta) e^{-2\pi i \omega \Delta} d\Delta. \quad (7)$$

B. State Space Gaussian Process (SSGP)

A Gaussian Process can be defined by two major components described in eqs. (2) and (3). Unfortunately, the GP Prior equation, eq. (2) is in a large matrix form making it computationally complex. We turn our attention to an alternative formulation of GP Prior described solely as a differential equation driven by a white noise process (zero-mean Gaussian process with a Dirichlet kernel function of magnitude q_c) [22]. We start by considering a single input

Algorithm 1 State Space Gaussian Process Conversion

- 1: $\mathcal{F}[k(\tau)] = \mathbf{S}(\omega) = F(j\omega)F(-j\omega)$ {Bochner}
- 2: $\frac{d^2}{d\omega^2} \mathbf{S}(\omega) = s(j\omega)s(-j\omega) \approx \sum_{i=0}^m c_i(\omega^2)^i$ {Taylor Series}
- 3: $\sum_{i=0}^m c_i(\omega^2)^i = \prod_{i=1}^m (r_i^2 + \omega^2)$ {Root Factor}
- 4: $\sum_{i=0}^m c_i(\omega^2)^i = \prod_{i=1}^m (r_i - j\omega)(r_i + j\omega)$ {Decomp.}
- 5: $s(j\omega) = \prod_{i=1}^m (r_i + j\omega)$ {Positive conjugate,}
- 6: $s(j\omega) = \sum_{i=0}^m a_i(j\omega)^i$ {Defactorization,}
- 7: $\mathcal{F}^{-1}[s(j\omega)] = \sum_{i=0}^m a_i \frac{\partial^i f(x)}{\partial x^i}$ {Inverse Fourier}

single output formulation in terms of a single input variable $x \in \mathbb{R}$ as

$$Q(x) = a_l \frac{\partial^l f(x)}{\partial x^l} + \cdots + a_1 \frac{\partial f(x)}{\partial x} + a_0 f(x), \quad (8)$$

where $Q(x)$ represents a white noise process and $f(x)$ represents the function we would like to fit over our data. Conceptually, we can see the reformulation of the prior from matrix form to a differential equation as the difference in taking the averaged sum of correlations from all points in the dataset eq. (5) as opposed to the trend of motion described by the previous points in the dataset in eq. (8). The additional advantage of seeing the GP in terms of the trend of the previous input points as opposed to the sum of all correlations to points is the ability to process information sequentially. We can derive the transfer function of eq. (8) as

$$F(j\omega) = \frac{q_c}{\underbrace{a_n(j\omega)^l + \cdots + a_1(j\omega) + a_0}_{s(j\omega)}}, \quad (9)$$

where $q_c \in \mathbb{R}$ is a constant describing the Fourier transform of $Q(x)$ of the process. The q_c parameter can be obtained through either optimization in hyperparameter training or tuned by a user. An important facet to emphasize is that a stationary kernel function $k(\Delta)$ is described completely and solely by its spectral density and eq. (9) also has a spectral density. We derive the steps needed to convert an arbitrary stationary kernel function $k(\Delta)$ into a corresponding differential equation. The key component is to obtain a differential equation in the form of eq. (8) from $s(j\omega)$ components.

The full method to construct this relationship is described in Algorithm 1. First, in step (1), the left and right side equivalences are due to the Bochner theorem and the definition of spectral density respectively. Very few kernels can be represented perfectly as a linear differential equation. This necessitates an approximation with a high enough m order Taylor series polynomial shown in step (2) of our algorithm. The Taylor series expansion is in terms of ω^2 to allow easy spectral decomposition in positive and negative conjugate components as in step (4). Only the positive component from step (5) is required. From step (5), we can de-factorize the positive conjugate in step (6) to obtain the coefficients of the polynomial. Finally, step (7) clearly shows that we obtain the form of eq. (8) after performing the inverse Fourier transform. Overall, in this process, we matched the spectral density of eq. (9) with the spectral

density described by eq. (7). If the two stationary GP match spectral densities, then they are the same GP. Therefore, given the linear differential equation in eq. (8) and $\mathbf{f}(x) = \left[f(x) \quad \frac{\partial f(x)}{\partial x} \quad \dots \quad \frac{\partial^m f(x)}{\partial x^m} \right]^\top$, we can create an equivalent state space form with the process model as

$$\frac{\partial \mathbf{f}(x)}{\partial x} = \underbrace{\begin{bmatrix} 0 & 1 & & \\ & \ddots & \ddots & \\ & & 0 & 1 \\ -a_0 & -a_1 & \dots & -a_{m-1} \end{bmatrix}}_{\mathbf{A}} \mathbf{f}(x) + \underbrace{\begin{bmatrix} 0 \\ \vdots \\ 0 \\ 1 \end{bmatrix}}_{\mathbf{L}} Q(x),$$

$$y = \underbrace{\begin{bmatrix} 1 & 0 & \dots & 0 \end{bmatrix}}_{\mathbf{H}} \mathbf{f}(x) + \varepsilon, \quad (10)$$

For the output equation recall that eq. (3) is already in linear form therefore no modification is needed. Once this state space equation is derived, a Kalman Filter or other similar technique can be employed to linearly iterate across the dataset as a sequence to make a prediction. To further elaborate, given a random input x_* , and an initial state $\mathbf{f}(x)$, the system will iteratively move through the dataset \mathcal{D} performing Kalman predictions and updates till it updates on $x_i \in \mathcal{D}$ which is the closest data point to x_* . Consequently, the derivatives such as $\frac{\partial \mathbf{f}(x)}{\partial x}$ and higher are updated during this motion allowing the system to observe a data trend. This facilitates a final prediction between x_i and x_* based on the previously observed trend. Consequently, our system is no longer deterministic as the system can either approach from the right or left of x_* . Extending the formulation to multiple input single output (MISO), can be obtained by stacking the eq. (10) diagonally for the \mathbf{A} matrix, horizontally \mathbf{H} , and vertically for \mathbf{L}

$$\begin{bmatrix} \frac{\partial \mathbf{f}_0(x_0)}{\partial t} \\ \vdots \\ \frac{\partial \mathbf{f}_d(x_d)}{\partial t} \end{bmatrix} = \begin{bmatrix} \mathbf{A}_0 & \dots & \mathbf{0} \\ & \ddots & \\ \mathbf{0} & \dots & \mathbf{A}_d \end{bmatrix} \begin{bmatrix} \mathbf{f}_0(x_0) \\ \vdots \\ \mathbf{f}_d(x_d) \end{bmatrix} + \begin{bmatrix} \mathbf{L}_0 \\ \vdots \\ \mathbf{L}_d \end{bmatrix} Q(\mathbf{x}),$$

$$y = \begin{bmatrix} \mathbf{H}_0 & \dots & \mathbf{H}_d \end{bmatrix} \begin{bmatrix} \mathbf{f}_0(x_0) \\ \vdots \\ \mathbf{f}_d(x_d) \end{bmatrix} + \varepsilon. \quad (11)$$

Consequently, the conversion of kernel functions into linear state space forms expressed as matrices allows easy exportation to other frameworks. This is because there is no need to write additional functions simply swapping transition matrix parameters is sufficient, and consequently dropping any library dependencies needed to solve for the transition matrices. Our toolbox has this functionality which allows exporting python models directly to C++.

C. Kernel Formulation

We give example derivations for our toolbox's kernel using Algorithm 1.

1) *Radial Basis Function Kernel*: The Radial Basis Function (RBF) kernel is also known as the squared exponential

kernel. No closed-form solution for any arbitrary approximation order m exists. Fortunately, following algorithm 1, an appropriate approximation can always be derived for some order m and hyperparameter $z \in \mathbb{R}_{>0}$ known as the lengthscale. For future reference $\mathbb{R}_{>0}$ is the set of all positive real numbers excluding zero

$$k(\Delta) = e^{-0.5z^2\Delta^2},$$

$$s(j\omega)s(-j\omega) = 1 + (z\omega)^2 + \dots \frac{1}{m!}(z\omega)^{2m}. \quad (12)$$

Considering $m = 2$ and $z = 1$ we obtain

$$s(j\omega)s(-j\omega) = 1 + (\omega)^2 + 0.5(\omega)^4,$$

$$= 0.5(\omega^2 + (1+j))(\omega^2 + (1-j)),$$

$$s(j\omega) = \frac{1}{\sqrt{2}}(j\omega + \sqrt{1+j})(j\omega + \sqrt{1-j}), \quad (13)$$

$$s(j\omega) \approx \frac{1}{\sqrt{2}}(j\omega)^2 + \frac{2.2}{\sqrt{2}}(j\omega) + 1.$$

2) *Matern Kernel*: Matern is a more complex and generalizable kernel compared to the RBF Kernel and approaches the RBF Kernel as $\nu \rightarrow \infty$. This kernel is one of the few kernels where there is a closed-form solution for any arbitrary set of hyperparameters

$$k(\Delta) = \sigma^2 \frac{2^{1-\nu}}{\Gamma(\nu)} \left(\sqrt{2\nu} \frac{\Delta}{z} \right)^\nu K_\nu \left(\sqrt{2\nu} \frac{\Delta}{z} \right),$$

$$s(j\omega)s(-j\omega) = (\lambda^2 + \omega^2)^{(\nu+1/2)},$$

$$s(j\omega) = (\lambda + j\omega)^{(\nu+1/2)}, \quad (14)$$

where $\lambda = \frac{\sqrt{2\nu}}{z}$. $\sigma, \nu, z \in \mathbb{R}_{>0}$ are hyperparameters of the kernel. Γ and K_ν are the gamma function and modified Bessel function of the second kind respectively.

3) *Periodic Kernel*: The periodic kernel has good expressiveness with oscillating features

$$k(\Delta) = e^{-2z^2 \sin^2(0.5w_0\Delta^2)},$$

$$s(j\omega)s(-j\omega) = \sum_{j=0}^{\infty} q_j^2 (\delta(\omega - j\omega_0) + \delta(\omega + j\omega_0)). \quad (15)$$

Unfortunately, the power spectrum is a summation of impulse functions, δ . It is possible to obtain a slightly modified state space approximating the series to the limit $J \rightarrow \infty$ [29]

$$A_j^k = \begin{bmatrix} 0 & -j\omega_0 \\ j\omega_0 & 0 \end{bmatrix}. \quad (16)$$

The measurement matrix \mathbf{H} block row vector is a modified $\mathbf{H}_j^k = [1 \ 0]$, while the noise diffusion is the standard $\mathbf{L}_j^k = [0 \ 1]^\top$. To acquire a better approximation, the system k matrices are stacked similarly to the MISO system k times.

V. LEARNING QUADROTOR DYNAMICS

We validate our approach by considering a learning residual quadrotor dynamics problem. In the following section, we describe the formulation of our problem for both a Single input single output (SISO) GP implementation and MISO GP

TABLE I
AVERAGE INFERENCE TIME PER DATA POINT [ms]

Points	SISO SSGP			SISO GPT			MISO SSGP			MISO GPT		
	PER	RBF	MAT	PER	RBF	MAT	PER	RBF	MAT	PER	RBF	MAT
10	5.77	3.57	3.36	2.01	0.96	0.99	8.11	6.09	5.79	1.68	0.97	1.01
50	5.74	3.51	3.34	1.73	1.01	1.10	30.18	27.85	25.96	1.71	0.97	1.06
200	5.83	3.53	3.37	2.09	1.18	1.33	32.35	30.40	28.29	2.29	1.21	1.33
1000	6.61	4.43	4.27	16.04	3.23	4.16	43.23	41.21	39.44	46.45	3.16	4.95
2500	5.71	3.61	3.48	93.64	34.24	57.71	57.69	55.76	53.77	237.86	33.98	57.32
6000	6.91	4.50	4.43	402.47	167.29	255.39	88.19	86.26	84.26	1040.41	154.40	254.38

formulation. For our SISO GP, we follow a similar modeling approach to [1]. We consider the following dynamics

$$m\mathbf{a} = \mathbf{R}_B^W \tau \mathbf{e}_3 - mge_3, \quad (17)$$

where $\mathbf{e}_3 = [0 \ 0 \ 1]^\top$, \mathbf{a} is the acceleration in the world frame of the quadrotor, g the gravitational value, and \mathbf{R}_B^W the rotation converting points from the quadrotor's body frame to the world frame. The vehicle's velocity in the body frame, \mathbf{v}_b is used to estimate the acceleration error on the body of the vehicle δ_a . This is based on standard fluid dynamic approximations of the air resistance for a moving body when evolving at different speeds. In our dataset, we calculate the term of δ_a as

$$\delta_a = \mathbf{R}_B^W \left(\frac{\mathbf{v}_b - \hat{\mathbf{v}}_b}{\delta t_k} \right), \quad (18)$$

where $\hat{\mathbf{v}}_b$ is the predicted velocity from integrating eq. (17) from the previous time step, and δt_k is the length of that time step. We can then construct our SISO GP using \mathbf{v}_b as the input and δ_a as the output, acceleration error, we would like to predict. To handle multiple outputs, we consider each body axis independent and apply an independent GP individually to each axis component. This makes our augmented GP dynamic formulation of eq. (17) as

$$\mathbf{a} = \frac{1}{m} \mathbf{R}_B^W (\tau \mathbf{e}_3 + f_{GP}(\mathbf{v}_b)) + g \mathbf{e}_3. \quad (19)$$

Fundamentally, the GP term f_{GP} in eq. (19) infers a first order propeller drag term [1], [30]. However, the vehicle is generally subject to additional types of disturbances. These disturbances are a function of motor thrust which is a function of ω_i of the motor's angular velocity through a k_f coefficient, obtaining the total thrust as

$$\tau = T_0 + T_1 + T_2 + T_3, \quad T_i = \omega_i^2 * k_f, \quad (20)$$

where T_i is the motor thrust generated by the motor i . However, the thrust mapping is only an approximation and does not account for other external factors or changes in the propeller over time [31]. Our system augments this term by using a MISO GP to calculate the disturbances related to motor thrust from both the fluid resistance from high body rate changes and inaccuracies in the thrust mapping. Therefore, our full system augmented model of eq. (17) for MISO becomes

$$\mathbf{a} = \frac{1}{m} \mathbf{R}_B^W (\tau \mathbf{e}_3 + f_{GP}(\mathbf{v}_b, \omega_0, \omega_1, \omega_2, \omega_3)) + g \mathbf{e}_3. \quad (21)$$

VI. RESULTS

We design our evaluation procedure to address the following questions. (i) What is the computational speed gain from converting to state space versus a traditional matrix inversion eq. (5)? (ii) Can the SSGP perform robust inference to the real-world data comparable to the baseline matrix inversion implementation of eq. (5)? (iii) Can extending the input space to multiple inputs greatly improve the inference? We perform several experiments to address these questions. First, we compare our computational time for solving a full dataset between our implementation and a traditional method. Next, we compared the inference RMSE and fitting of our SSGP approach with a traditional inference approach of matrix inversion for both SISO and MISO. For our baseline implementation of eq. (5) we leverage GPytorch (GPT), [8] a library that performs robust and efficient matrix operations and optimizations by integrating with Pytorch's ecosystem. This allows GPT to perform both efficient optimization of hyperparameters for the objective eq. (4) and inference with eq. (5). We evaluated our approach on a real-world dataset of quadrotor dynamics and flight, [32]. This dataset contains real-world unfiltered data resulting in much noisier signals and stressing our models' ability to perform inference. We chose three different flight trajectories: Circle (training), Parabola (validation), and Lemniscate (validation). To validate inference time, we tested our implementation on a 11th Gen Intel i7 for both SSGP and GPT in Table I with pure CPU computation. The results were repeated with a Xavier NX achieving similar timing trends. This table records the average amount of time taken to complete one inferences given a dataset of size n . Our time comparison consists of the following kernels: RBF, Matern, and periodic given different approximation orders, RBF $m = 6$, Matern $m = 3$, and periodic $m = 18$ and compared with both GPT in a SISO format and a MISO (5 independent inputs). Our model for SISO or $m \leq 10$ can perform inference much faster than a comparable baseline implementation, especially for datasets > 1000 . However, the SSGP formulation slows down in the MISO formulation. This is because computationally MISO for the SSGP of order $m = 6$ and 5 inputs is equivalent to a $m = 30$ approximation. This gives even slower computations than the GPT version until 2500+ points are used. Our inference task on the dataset is to model the residual accelerations defined in Section V with body velocity as the input to prediction acceleration error for SISO, and body velocity and motor speeds for MISO. For our RMSE evaluation in

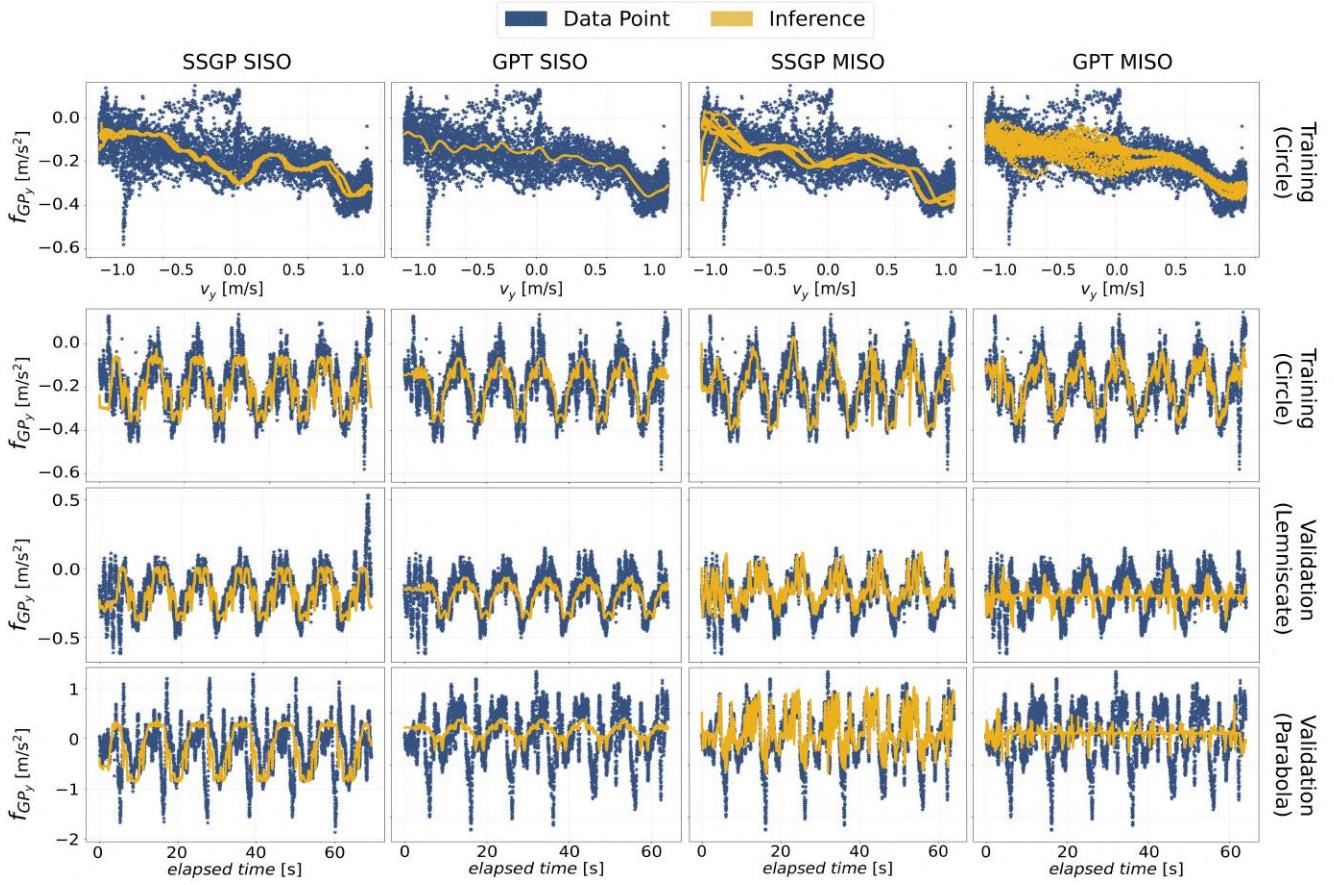


Fig. 2: Predictive performance of the propeller drag term on the y-axis of f_{GP} denoted as f_{GP_y} over training and validation datasets. Data trend between body velocity and residual acceleration error (first row). Visualization of the fitting on the residuals with respect to time (second to fourth rows).

TABLE II
PREDICTIVE PERFORMANCE RMSE [m s^{-2}]

Trajectory	SISO		MISO	
	SSGP	GPT	SSGP	GPT
Circle (training)	0.0938	0.0755	0.0852	0.0666
Transposed Parabola (val)	0.1330	0.1164	0.0889	0.1516
Extended Lemniscate (val)	0.6079	0.4561	0.3176	0.4698

Fig. (2), we focused on estimating the drag term on the y-axis, f_{GP_y} , choosing the RBF Kernel with an approximation order $m = 6$. The residual inference performance of SSGP is compared to the same kernel implemented in GPT using the same hyperparameters. As observed, our SSGP method despite performing similarly to GPT on the training set better generalizes on the testing dataset with lower RMSEs across all test sets (see Table II), and can perform faster inference. We believe the greater generalizability of our system depends on the prior induced by the SSGP structure where inference is achieved through observing the data trend rather than a correlation on the dataset like in GPT. Furthermore scaling our inputs to include the motor rates allows the MISO SSGP to better capture high-frequency features than a SISO system. Similar results are obtained for the other drag components f_{GP_x} , f_{GP_z} .

VII. CONCLUSION

In this work, we presented GaPT, a scalable and fast toolkit for state space GPs. This method shows superior computational performance compared to a classical GP inference approach on both single and multi-input settings especially for large datasets. GaPT enables easy exporting of GPs into C++ functions allowing direct integration in common robotics libraries enabling real-time inference speed even on embedded platforms used on SWaP-constrained aerial robots. GaPT incorporates both single or multiple inputs solutions. The latter further improves the inference capabilities of the proposed approach. Future works will involve adding additional kernels into the Toolkit such as spectral mixture or deep state-space models which can better capture complex patterns. This process will also be facilitated by the open source implementation that the robotics community can leverage to collaborate by adding additional functionalities that may be helpful for different purposes. Furthermore, we will show the integration of our method in various additional robotic use cases such as for control in aggressive maneuvers [33], calculating safety sets, or exploiting the full power of model predictive control with online refined learned dynamics.

REFERENCES

[1] G. Torrente, E. Kaufmann, P. Föhn, and D. Scaramuzza, “Data-driven mpc for quadrotors,” *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 3769–3776, 2021.

[2] L. Hewing, J. Kabzan, and M. N. Zeilinger, “Cautious model predictive control using gaussian process regression,” *IEEE Transactions on Control Systems Technology*, vol. 28, no. 6, pp. 2736–2743, 2019.

[3] M. P. Deisenroth, D. Fox, and C. E. Rasmussen, “Gaussian processes for data-efficient learning in robotics and control,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 37, no. 2, pp. 408–423, 2015.

[4] T. D. Barfoot, C. H. Tong, and S. Särkkä, “Batch continuous-time trajectory estimation as exactly sparse gaussian process regression,” in *Robotics: Science and Systems*, vol. 10. Citeseer, 2014, pp. 1–10.

[5] A. K. Akametalu, J. F. Fisac, J. H. Gillula, S. Kaynama, M. N. Zeilinger, and C. J. Tomlin, “Reachability-based safe learning with gaussian processes,” in *53rd IEEE Conference on Decision and Control*, 2014, pp. 1424–1431.

[6] L. Wang, E. A. Theodorou, and M. Egerstedt, “Safe learning of quadrotor dynamics using barrier certificates,” in *IEEE International Conference on Robotics and Automation (ICRA)*, 2018, pp. 2460–2465.

[7] A. Loquercio, A. Saviolo, and D. Scaramuzza, “Autotune: Controller tuning for high-speed flight,” *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 4432–4439, 2022.

[8] J. Gardner, G. Pleiss, K. Q. Weinberger, D. Bindel, and A. G. Wilson, “Gpytorch: Blackbox matrix-matrix gaussian process inference with gpu acceleration,” *Advances in neural information processing systems*, vol. 31, 2018.

[9] R. Verschueren, G. Frison, D. Kouzoupis, N. van Duijkeren, A. Zanelli, R. Quirynen, and M. Diehl, “Towards a modular software package for embedded optimization,” in *Proceedings of the IFAC Conference on Nonlinear Model Predictive Control (NMPc)*, 2018.

[10] A. Domahidi and J. Jerez, (2014–2019) Forces professional. Embotech AG. [Online]. Available: <https://embotech.com/FORCES-Pro>

[11] C. Mastalli, R. Budhiraja, W. Merkt, G. Saurel, B. Hammoud, M. Naveau, J. Carpenter, L. Righetti, S. Vijayakumar, and N. Mansard, “Crocoddyl: An Efficient and Versatile Framework for Multi-Contact Optimal Control,” in *IEEE International Conference on Robotics and Automation (ICRA)*, 2020.

[12] E. Biyik, N. Huynh, M. J. Kochenderfer, and D. Sadigh, “Active preference-based gaussian process regression for reward learning,” in *Robotics: Science and System XVI*, 2020.

[13] A. Gahlawat, P. Zhao, A. Patterson, N. Hovakimyan, and E. Theodorou, “L1-gp: L1 adaptive control with bayesian learning,” in *Learning for Dynamics and Control*. PMLR, 2020, pp. 826–837.

[14] A. G. d. G. Matthews, M. Van Der Wilk, T. Nickson, K. Fujii, A. Boukouvalas, P. León-Villagrá, Z. Ghahramani, and J. Hensman, “Gpflow: A gaussian process library using tensorflow.” *J. Mach. Learn. Res.*, vol. 18, no. 40, pp. 1–6, 2017.

[15] G. Parra and F. Tobar, “Spectral mixture kernels for multi-output gaussian processes,” in *Proceedings of the 31st International Conference on Neural Information Processing Systems*, ser. NIPS’17. Red Hook, NY, USA: Curran Associates Inc., 2017, p. 6684–6693.

[16] Z. Zhao, M. Emzir, and S. Särkkä, “Deep state-space gaussian processes,” *Statistics and Computing*, vol. 31, no. 6, pp. 1–26, 2021.

[17] S. Ambikasaran, D. Foreman-Mackey, L. Greengard, D. W. Hogg, and M. O’Neil, “Fast direct methods for gaussian processes,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 38, no. 2, pp. 252–265, 2016.

[18] J. Quiñonero-Candela and C. E. Rasmussen, “A unifying view of sparse approximate gaussian process regression,” *Journal of Machine Learning Research*, vol. 6, no. 65, pp. 1939–1959, 2005.

[19] “Fast gaussian process regression for big data,” *Big Data Research*, vol. 14, pp. 12–26, 2018.

[20] S. Särkkä, “Linear operators and stochastic partial differential equations in gaussian process regression,” in *International Conference on Artificial Neural Networks*. Springer, 2011, pp. 151–158.

[21] A. Corenflos, Z. Zhao, and S. Särkkä, “Temporal gaussian process regression in logarithmic time,” in *25th International Conference on Information Fusion (FUSION)*, 2022, pp. 1–5.

[22] S. Sarkka, A. Solin, and J. Hartikainen, “Spatiotemporal learning via infinite-dimensional bayesian filtering and smoothing: A look at gaussian process regression through kalman filtering,” *IEEE Signal Processing Magazine*, vol. 30, no. 4, pp. 51–61, 2013.

[23] C. E. Rasmussen and C. K. I. Williams, *Gaussian processes for machine learning.*, ser. Adaptive computation and machine learning. MIT Press, 2006.

[24] I. Murray and R. P. Adams, “Slice sampling covariance hyperparameters of latent gaussian models,” *Advances in neural information processing systems*, vol. 23, 2010.

[25] J. Hartikainen and S. Särkkä, “Kalman filtering and smoothing solutions to temporal gaussian process regression models,” in *IEEE international workshop on machine learning for signal processing*, 2010, pp. 379–384.

[26] S. Bochner *et al.*, *Lectures on Fourier integrals*. Princeton University Press, 1959, vol. 42.

[27] M. L. Stein, *Interpolation of spatial data: some theory for kriging*. Springer Science & Business Media, 1999.

[28] C. Chatfield, *The analysis of time series: an introduction*, 6th ed. CRC Press, 2004.

[29] A. Solin and S. Särkkä, “Explicit link between periodic covariance functions and state space models,” in *Artificial Intelligence and Statistics*. PMLR, 2014, pp. 904–912.

[30] J. Svacha, J. Paulos, G. Loianno, and V. Kumar, “Imu-based inertia estimation for a quadrotor using newton-euler dynamics,” *IEEE Robotics and Automation Letters*, vol. 5, no. 3, pp. 3861–3867, 2020.

[31] X. Dai, Q. Quan, J. Ren, and K.-Y. Cai, “An analytical design-optimization method for electric propulsion systems of multicopter uavs with desired hovering endurance,” *IEEE/ASME Transactions on Mechatronics*, vol. 24, no. 1, pp. 228–239, 2019.

[32] A. Saviolo, G. Li, and G. Loianno, “Physics-inspired temporal learning of quadrotor dynamics for accurate model predictive trajectory tracking,” *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 10256–10263, 2022.

[33] J. Mao, S. Nogar, C. M. Kroninger, and G. Loianno, “Robust active visual perching with quadrotors on inclined surfaces,” *IEEE Transactions on Robotics*, pp. 1–17, 2023.