

Gaussian Processes with State-Dependent Noise for Stochastic Control

Menner, Marcel; Berntorp, Karl

TR2023-100 August 18, 2023

Abstract

This paper considers a stochastic control framework, in which the residual model uncertainty of a dynamical system is learned using a Gaussian Process (GP). In the proposed formulation, the residual model uncertainty consists of a nonlinear function and state-dependent noise. The proposed formulation uses a posterior-GP to approximate the residual model uncertainty and a prior-GP to account for state-dependent noise. The two GPs are interdependent and are thus learned jointly using an iterative algorithm. Theoretical properties of the iterative algorithm are established. Advantages of the state-dependent formulation include (i) faster convergence of the GP estimate to the unknown function as the GP learns which data samples are more trustworthy and (ii) an accurate estimate of state-dependent noise, which can, e.g., be useful for a controller or decision-maker to determine the uncertainty of an action. Simulation studies highlight these two advantages.

IEEE Conference on Control Technology and Applications (CCTA) 2023

Gaussian Processes with State-Dependent Noise for Stochastic Control

Marcel Menner and Karl Berntorp

Abstract—This paper considers a stochastic control framework, in which the residual model uncertainty of a dynamical system is learned using a Gaussian Process (GP). In the proposed formulation, the residual model uncertainty consists of a nonlinear function and state-dependent noise. The proposed formulation uses a *posterior-GP* to approximate the residual model uncertainty and a *prior-GP* to account for state-dependent noise. The two GPs are interdependent and are thus learned jointly using an iterative algorithm. Theoretical properties of the iterative algorithm are established. Advantages of the state-dependent formulation include (i) faster convergence of the GP estimate to the unknown function as the GP learns which data samples are more trustworthy and (ii) an accurate estimate of state-dependent noise, which can, e.g., be useful for a controller or decision-maker to determine the uncertainty of an action. Simulation studies highlight these two advantages.

I. INTRODUCTION

Recent progress in computational resources, data science, robotics, and control has sparked an interest in considering uncertainty bounds within a controller or decision-maker [1]. Gaussian-process (GP) regression is a popular paradigm for estimating a function as it provides not only a mean estimate of the desired quantity, but also a confidence bound [2]. Using a probabilistic interpretation, a GP can accurately reflect the distribution of the data for a given and constant prior on the measurement noise. This paper proposes a GP formulation for learning an unknown function and its state-dependent noise distribution for stochastic control.

A. Motivating Example

GP regression uses the posterior distribution given data points to fit a function mapping from a value x to a noisy output y . Fig. 1 illustrates a comparison between a classical GP and the proposed GP with state-dependent noise. It shows that the classical solution fails to reduce its uncertainty bounds based on the proximity of the output measurements, y . This behavior becomes apparent by inspecting the mean and covariance equations for a classical GP solution, in which the mean is a function of both x and the associated measurement y , i.e., $\mu(x)$ depends on the data x^i, y^i , whereas the covariance $\Sigma(x)$ depends only on the input to the unknown function x^i , and not on y^i . In contrast, the proposed GP formulation accurately reflects the noise in the data for both regions of higher noise and lower noise.

B. Contributions and Algorithmic Realization

This paper considers a stochastic control framework, where the residual model uncertainty is learned using GP regression. We propose a formulation, where the residual

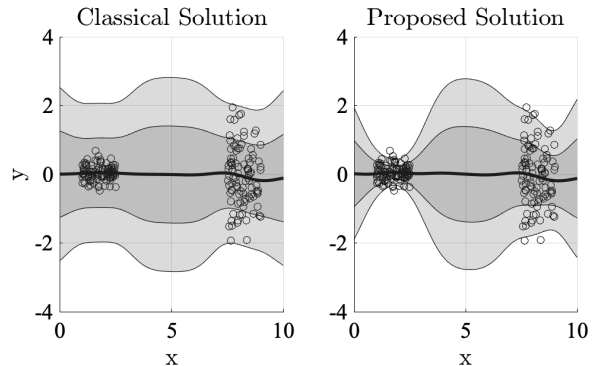


Fig. 1. Motivating example for GP regression with state-dependent noise. The displayed scenario uses 100 samples between $x = [1, 2.5]$ with lower noise and between $x = [7.5, 9]$ with higher noise. Left: Classical solution without state-dependence. Right: State-dependent solution. The example that the GP on the left evenly reduces uncertainty based on the state, x . The GP on the right reduces uncertainty based on the state, x , but also based on the value, y . Hence, if the measurement values y are “closer together” (see left part of the two plots), the uncertainty is reduced appropriately.

model uncertainty consists of a nonlinear function and state-dependent noise. The proposed GP formulation learns both the nonlinear function and the noise distribution simultaneously. It uses one GP to approximate the posterior distribution, which we refer to as the *posterior-GP*. Different from a classical GP, the proposed GP uses a data-based prior estimate to adjust the prior to match the noise distribution in the training data. In this paper, this data-based prior estimate is provided by a second GP, which we refer to as *prior-GP*. The proposed posterior-GP uses the training data directly, and the prior-GP uses the variance of the training data. However, the variance requires a mean estimate—provided by the posterior-GP. On the other hand, the posterior-GP requires an estimate of the prior—provided by prior-GP. We propose an iterative algorithm to address this interdependence and learn both GPs simultaneously. We present theoretical properties of the iterative algorithm such as boundedness and guaranteed convergence under some simplifying assumptions. The main advantages of the proposed solution are (i) less data are needed to accurately estimate the unknown function due to the GP learning which data samples are more trustworthy and (ii) a decision-maker can accurately determine the uncertainty of an action. Both advantages are illustrated using simulation examples.

C. Related Works

GPs for stochastic control: GPs are often used for learning unknown parts of system dynamics [3]–[7] or cost functions [8], [9]. In [3], [4], a GP is used to learn a residual dynamics model and its uncertainty bound is subsequently

used within a chance-constrained MPC. Further, [4] presents approximations for efficient computation. In [5], a control approach based on learning GPs online is presented and applied to a race car. The work in [6] uses a GP to model tire friction for improving the performance of a stochastic MPC. In [7], a GP-based stochastic MPC is used for cooperative adaptive cruise control, in which the uncertainty is used for predictive behavioral models. In [8], [9], GPs are utilized to learn a cost function that can be used for control.

Constraints for GP regression: Recent work studies constrained GP regression [10]–[16], which can, e.g., be used to enforce physics-based information. E.g., the approach in [10] uses a linear operators to enforce equality constraints, [11] uses virtual measurements to enforce inequality constraints, [12] uses tractability constraint to learn a reference trajectory from data for model predictive control, and [13] uses a basis-function approach to enforce constraints on the GP model.

GPs with input-dependent noise: Input-dependent noise is considered in [17]–[21], e.g., for robotic perception and localization. In [17], the posterior distribution of the noise distribution is sampled using Markov chain Monte Carlo methods. In contrast, we use an algorithm that does not rely on sampling and iteratively refines both the posterior-GP and the prior-GP. The work in [18] also considers an iterative estimation algorithm, which is not guaranteed to converge and may lead to oscillations. Compared to [18], we state theoretical properties of the proposed iterative algorithm such as boundedness and convergence under some mild assumptions. Further, [19] considers input-dependent noise for optimal design, and [20] presents a robot localization task using GPs with input-dependent noise. Different from [17]–[21], we present GPs with state-dependent noise with application to stochastic control and theoretical properties of the algorithm.

D. Preliminaries and Notation

We use $[x]^i$ to indicate the i -th element of a vector x and $[X]^{ii}$ to indicate the element in the i -th row and i -th column of a matrix X . We use the Hadamard product (element-wise product) \circ , i.e., for $x, w \in \mathbf{R}^{n_x}$, the i -th element of $z = x \circ w \in \mathbf{R}^{n_x}$ is $[z]^i = [x]^i [w]^i$. To ease exposition, we define $x^{\circ 2} := x \circ x$. We use I for the identity matrix and $\mathbf{1}$ and $\mathbf{0}$ for the one and zero vector of appropriate dimension. For a vector x , $x \sim \mathcal{N}(\mu, \Sigma)$ indicates that $x \in \mathbf{R}^{n_x}$ is Gaussian distributed with mean μ and covariance Σ .

Throughout, we use a common choice of kernel $k(x, w)$ with the radial basis function,

$$k(x, w) = \exp(-l(x - w)^T(x - w)), \quad (1)$$

where l is the length scale defining the spread of the basis function. Note that $k(x, w) = 1$ if $x = w$ and $k(x, w) \rightarrow 0$ for x far away from w . For brevity, $k_{x,w} := k(x, w)$,

$$k_{x,w} = [k_{x,w^1} \quad k_{x,w^2} \quad \dots \quad k_{x,w^M}],$$

the collection of vectors

$$\begin{aligned} X &= [x^1 \quad x^2 \quad \dots \quad x^N] \in \mathbf{R}^{n_x \times N} \\ W &= [w^1 \quad w^2 \quad \dots \quad w^M] \in \mathbf{R}^{n_x \times M}, \end{aligned}$$

and

$$K_{X,W} = \begin{bmatrix} k_{x^1,w^1} & k_{x^1,w^2} & \dots & k_{x^1,w^M} \\ k_{x^2,w^1} & k_{x^2,w^2} & \dots & k_{x^2,w^M} \\ \vdots & \vdots & \ddots & \vdots \\ k_{x^N,w^1} & k_{x^N,w^2} & \dots & k_{x^N,w^M} \end{bmatrix}.$$

For a function $\lambda(x)$ with $\lambda : \mathbf{R}^{n_x} \rightarrow \mathbf{R}$, we use bold symbol $\lambda(X)$ with $\lambda : \mathbf{R}^{n_x \times N} \rightarrow \mathbf{R}^N$ as shorthand notation to indicate that the i -th element of $[\lambda(X)]^i = \lambda(x^i)$. Similarly, we use the diagonal matrix $\bar{\lambda}(X)$ with $\bar{\lambda} : \mathbf{R}^{n_x \times N} \rightarrow \mathbf{R}^{N \times N}$, whose diagonal elements $[\bar{\lambda}(X)]^{ii} = \lambda(x^i)$. Hence, $\bar{\lambda}(X) = \text{diag}(\lambda(X))$.

II. PROBLEM STATEMENT

A. Stochastic Control Formulation

We consider an uncertain dynamical system with the measurable state $x_k \in \mathbf{R}^{n_x}$ in discrete time,

$$x_{k+1} = f(x_k, u_k) + B(h(x_k, u_k) + \epsilon_k), \quad (2)$$

where $u_k \in \mathbf{R}^{n_u}$ is the input. The dynamical system model in (2) is composed of a known nominal function f , the known B , and an initially unknown function $h : \mathbf{R}^{n_x} \times \mathbf{R}^{n_u} \rightarrow \mathbf{R}^{n_h}$, which is to be learned using data. Finally, $\epsilon_k \in \mathbf{R}^{n_\epsilon}$ denotes noise or uncertainty arising from imperfect sensors, perception, model uncertainty, etc. We assume that the elements in ϵ_k are uncorrelated, i.e., ϵ_k is a realization of a distribution with a diagonal covariance matrix, however, the diagonal entries are allowed to be state-dependent,

$$\epsilon_k \sim \mathcal{N}(\mathbf{0}, \Sigma(x_k, u_k)) \quad (3)$$

with the diagonal matrix $\Sigma(x_k, u_k)$.

The dynamical system is subject to state and input constraints, which are enforced as chance constraints,

$$\Pr(x_k \in X) \geq p_x \quad (4a)$$

$$\Pr(u_k \in U) \geq p_u, \quad (4b)$$

where X and U are the state and input constraint sets that are to be satisfied with p_x and p_u probability, respectively.

The goal is thus to control the dynamical system in (2) in the presence of the noise/disturbances in (3) such that the chance constraints in (4) are not violated. In the following, we drop the dependence of h and Σ on u_k to ease exposition. In order to learn the unknown function h and unknown covariance matrix Σ , we use (2) to form

$$y = B^+(x_{k+1} - f(x_k, u_k)) = h(x_k) + \epsilon$$

with the Moore-Penrose pseudo inverse B^+ .

B. Gaussian-Process Regression

As the covariance in (3) is diagonal, we present the estimation problem for the one-dimensional case $n_h = 1$ for simplicity. However, $n_h \geq 2$ follows immediately by using n_h GPs, i.e., one GP for each of the n_h dimensions. Hence,

$$y^i = h(x^i) + \epsilon^i, \quad (5a)$$

where i indicates the i^{th} data point and

$$\epsilon^i \sim \mathcal{N}(0, g(\mathbf{x}^i)), \quad (5b)$$

with state-dependent variance $g(\mathbf{x}^i) > 0$ of the noise ϵ^i .

This paper shows how to simultaneously learn a surrogate function of h with state-dependent variance g using GPs with the training data \mathcal{D}_λ with N data samples

$$\mathcal{D}_\lambda = \{\mathbf{x}^i, y^i\}_{i=1}^N \quad (6a)$$

$$\mathbf{X} = [\mathbf{x}^1 \quad \mathbf{x}^2 \quad \dots \quad \mathbf{x}^N] \in \mathbf{R}^{n_x \times N} \quad (6b)$$

$$\mathbf{Y} = [y^1 \quad y^2 \quad \dots \quad y^N]^T \in \mathbf{R}^N. \quad (6c)$$

III. GPs WITH STATE-DEPENDENT NOISE

This section presents the proposed learning algorithm, in which one GP is used to approximate the unknown function h in (5a), which we refer to as *posterior-GP*. The posterior-GP uses a prior estimate of the noise distribution. This prior estimate is provided by a second GP, which adjusts the prior to reflect the noise distribution and is thus referred to as *prior-GP*. Both the posterior-GP and the prior-GP are determined using the data in (6), where the posterior-GP uses the training data in (6c) directly, and the prior-GP uses the variance of the training data in (6c). However, computing the variance requires a mean estimate—provided by the posterior-GP. In turn, the mean estimate depends on the prior—provided by the prior-GP. Hence, the posterior-GP and the prior-GP are interdependent, which we address next.

A. Posterior-GP

In order to derive the posterior distribution of the unknown function in (5a) given the measurement data, we use

$$\begin{bmatrix} \mathbf{Y} \\ y \end{bmatrix} \sim \mathcal{N}\left(\mathbf{0}, \begin{bmatrix} \mathbf{K}_{\mathbf{X}, \mathbf{X}} + \sigma_0^2 \mathbf{I} + \bar{\gamma}(\mathbf{X}) & \mathbf{k}_{\mathbf{X}, \mathbf{x}} \\ \mathbf{k}_{\mathbf{x}, \mathbf{X}} & k_{\mathbf{x}, \mathbf{x}} + \sigma_0^2 + \gamma(\mathbf{x}) \end{bmatrix}\right)$$

where $\gamma(\mathbf{x})$ is an initially unknown function, which assigns a prior to the measurement data. Then, the conditional distribution $p(y|\mathbf{Y}, \mathbf{X}, \mathbf{x})$ is a normal distribution with

$$y \sim \mathcal{N}(\lambda(\mathbf{x}), \Lambda(\mathbf{x})),$$

where the mean and the posterior variance are given by

$$\lambda(\mathbf{x}) = \mathbf{k}_{\mathbf{x}, \mathbf{X}} (\mathbf{K}_{\mathbf{X}, \mathbf{X}} + \sigma_0^2 \mathbf{I} + \bar{\gamma}(\mathbf{X}))^{-1} \mathbf{Y} \quad (7a)$$

$$\begin{aligned} \Lambda(\mathbf{x}) &= k_{\mathbf{x}, \mathbf{x}} + \sigma_0^2 + \gamma(\mathbf{x}) \\ &\quad - \mathbf{k}_{\mathbf{x}, \mathbf{X}} (\mathbf{K}_{\mathbf{X}, \mathbf{X}} + \sigma_0^2 \mathbf{I} + \bar{\gamma}(\mathbf{X}))^{-1} \mathbf{k}_{\mathbf{X}, \mathbf{x}}. \end{aligned} \quad (7b)$$

Remark 1 (Connection to classical GPs): In classical GP regression, the prior is assumed to be known and $\bar{\gamma}(\mathbf{x}) = 0$. We will use this as baseline comparison in Section V.

B. Prior-GP

To infer an accurate uncertainty quantification or noise estimate from data, this paper proposes to utilize a prior-GP. The prior-GP is trained using the measurements, y^i , and the mean estimate of the posterior-GP, $\lambda(\mathbf{x})$ in (7a). Hence, the prior-GP uses training data computed as the second moment of the measurements, y^i , i.e.,

$$\mathcal{D}_\gamma = \{\mathbf{x}^i, z^i\}_{i=1}^N, \quad (8a)$$

with

$$z^i = (y^i - \lambda(\mathbf{x}^i))^2 - \sigma_0^2, \quad (8b)$$

$$\mathbf{Z} = [z^1 \quad z^2 \quad \dots \quad z^N]^T \in \mathbf{R}^N. \quad (8c)$$

Then, the prior-GP uses the data in (8) and

$$\begin{bmatrix} \mathbf{Z} \\ z \end{bmatrix} \sim \mathcal{N}\left(\mathbf{0}, \begin{bmatrix} \mathbf{K}_{\mathbf{Z}, \mathbf{Z}} + \sigma_0^2 \mathbf{I} & \mathbf{k}_{\mathbf{Z}, \mathbf{x}} \\ \mathbf{k}_{\mathbf{x}, \mathbf{Z}} & k_{\mathbf{x}, \mathbf{x}} + \sigma_0^2 \end{bmatrix}\right).$$

Using the conditional distribution $p(z|\mathbf{Z}, \mathbf{X}, \mathbf{x})$, the prior-GP is given by

$$z \sim \mathcal{N}(\gamma(\mathbf{x}), \Gamma(\mathbf{x})),$$

with

$$\gamma(\mathbf{x}) = \mathbf{k}_{\mathbf{x}, \mathbf{Z}} (\mathbf{K}_{\mathbf{Z}, \mathbf{Z}} + \sigma_0^2 \mathbf{I})^{-1} \mathbf{Z}, \quad (9a)$$

$$\Gamma(\mathbf{x}) = k_{\mathbf{x}, \mathbf{x}} + \sigma_0^2 - \mathbf{k}_{\mathbf{x}, \mathbf{Z}} (\mathbf{K}_{\mathbf{Z}, \mathbf{Z}} + \sigma_0^2 \mathbf{I})^{-1} \mathbf{k}_{\mathbf{Z}, \mathbf{x}}. \quad (9b)$$

On the one hand, the training data for the prior-GP in (8) require a mean estimate of the unknown function h . On the other hand, the posterior-GP requires a prior estimate of the noise distribution. Hence, the posterior-GP in (7) and the prior-GP in (9) need to be estimated jointly, which we address using an iterative algorithm in Section III-C.

Remark 2 (Exploitation-exploration): GP models are often used as a surrogate function for black-box optimization, e.g., in Bayesian optimization. For this application, $\lambda(\mathbf{x})$ in (7a) quantifies the mean estimate and $\Gamma(\mathbf{x})$ in (9b) quantifies the uncertainty, which can be used for exploration, whereas $\Lambda(\mathbf{x})$ in (7b) quantifies the noise in the data.

C. Iterative algorithm

We jointly estimate the unknown function in (5a) and the noise distribution of its data samples in (5b) by iteratively refining the posterior estimate and the prior estimate. Algorithm 1 summarizes the procedure. At iteration j , in Line 4, the iterative algorithm updates the posterior-GP using the mean estimate of the prior-GP at the previous iteration,

$$\lambda_{j+1}(\mathbf{X}) = \mathbf{K}_{\mathbf{X}, \mathbf{X}} (\mathbf{K}_{\mathbf{X}, \mathbf{X}} + \sigma_0^2 \mathbf{I} + \bar{\gamma}_j(\mathbf{X}))^{-1} \mathbf{Y}. \quad (10a)$$

Then, in Line 6, it uses the updated posterior-GP estimate in (10a) to update the training data for the prior-GP,

$$\mathbf{Z}_{j+1} = \left((\mathbf{Y} - \lambda_{j+1}(\mathbf{X}))^2 - \sigma_0^2 \mathbf{1} \right), \quad (10b)$$

Next, in Line 8, it uses the training data (10b) to update the prior GP,

$$\gamma_{j+1}(\mathbf{X}) = \mathbf{K}_{\mathbf{Z}, \mathbf{Z}} (\mathbf{K}_{\mathbf{Z}, \mathbf{Z}} + \sigma_0^2 \mathbf{I})^{-1} \mathbf{Z}_{j+1}. \quad (10c)$$

We use a stopping criterion for Algorithm 1 based on the basis-function interpretation of GPs, i.e.,

$$\alpha_{j+1} = (\mathbf{K}_{\mathbf{X}, \mathbf{X}} + \sigma_0^2 \mathbf{I} + \bar{\gamma}_{j+1}(\mathbf{X}))^{-1} \mathbf{Y},$$

where α can be thought of as weights for the basis functions, i.e., $\lambda(\mathbf{x}) = \mathbf{k}_{\mathbf{x}, \mathbf{X}} \alpha$. This choice is motivated by the impact that the prior-GP has on the posterior-GP. In the following, we show that this iterative procedure converges to its optimum under some assumptions.

Algorithm 1: GP with state-dependent noise

```
1 Initialize  $\alpha_0 = \mathbf{0}$ ,  $\mathbf{Z}_0 = \mathbf{0}$ ,  $\gamma_0(\mathbf{X}) = \mathbf{0}$ ,  $j = 0$  ;
2 do
3   %% update posterior-GP estimate;
4    $\lambda_{j+1}(\mathbf{X}) = \mathbf{K}_{\mathbf{X},\mathbf{X}} (\mathbf{K}_{\mathbf{X},\mathbf{X}} + \sigma_0^2 \mathbf{I} + \tilde{\gamma}_j(\mathbf{X}))^{-1} \mathbf{Y}$ ;
5   %% update training data for prior-GP using
   posterior-GP;
6    $\mathbf{Z}_{j+1} = ((\mathbf{Y} - \lambda_{j+1}(\mathbf{X}))^{\circ 2} - \sigma_0^2 \mathbf{1})$  ;
7   %% update prior-GP;
8    $\gamma_{j+1}(\mathbf{X}) = \mathbf{K}_{\mathbf{X},\mathbf{X}} (\mathbf{K}_{\mathbf{X},\mathbf{X}} + \sigma_0^2 \mathbf{I})^{-1} \mathbf{Z}_{j+1}$  ;
9   %% check stopping criterion;
10   $\alpha_{j+1} = (\mathbf{K}_{\mathbf{X},\mathbf{X}} + \sigma_0^2 \mathbf{I} + \tilde{\gamma}_{j+1}(\mathbf{X}))^{-1} \mathbf{Y}$  ;
11   $j \leftarrow j + 1$  ;
12 while  $\|\alpha_{j+1} - \alpha_j\|_2 \leq \delta$ ;
```

IV. THEORETICAL PROPERTIES OF ALGORITHM

In order to analyze the convergence characteristics of the algorithm, we make a simplifying assumption, i.e., we study the algorithm in the presence of a small length scale l in (1). Using a small length scale $l \rightarrow 0$, we can simplify $\mathbf{K}_{\mathbf{X},\mathbf{X}} \rightarrow \mathbf{I}$, which allows for formulating analytical expressions of the matrix inverses in (10). Note that the assumption of a small length scale is only needed to establish a series of (nontight) bounds for the optimizer in Algorithm 1. Moreover, the assumption is utilized to show how to choose σ_0^2 appropriately. Therefore, in practice, this assumption is not needed.

The main steps of the proof are outlined in Lemma 1, Lemma 2, and Theorem 1. Lemma 1 shows that if the prior-GP converges, then the posterior-GP converges as well; Lemma 2 shows that the estimated prior-GP remains bounded for all iterations; and Theorem 1 proves that the prior GP is contracting, i.e., the prior GP converges to its optimum.

Lemma 1: If the prior-GP converges, then the posterior-GP converges as well, i.e., if $\gamma_j(\mathbf{X}) \rightarrow \gamma_*(\mathbf{X})$, then $\lambda_j(\mathbf{X}) \rightarrow \lambda_*(\mathbf{X})$.

Proof: This follows immediately from (10) with $\lambda_*(\mathbf{X}) = \mathbf{K}_{\mathbf{X},\mathbf{X}} (\mathbf{K}_{\mathbf{X},\mathbf{X}} + \sigma_0^2 \mathbf{I} + \tilde{\gamma}_*(\mathbf{X}))^{-1} \mathbf{Y}$. ■

Lemma 2: Let the prior-GP be initialized as in Line 1 in Algorithm 1. For a small length scale $l \rightarrow 0$, the prior estimate $\gamma_j(\mathbf{X})$ is bounded below and above with

$$-\frac{\sigma_0^2}{1 + \sigma_0^2} \mathbf{1} \leq \gamma_j(\mathbf{X}) \leq \frac{1}{1 + \sigma_0^2} (\mathbf{Y}^{\circ 2} - \sigma_0^2 \mathbf{1})$$

for all iterations j .

Proof: [Sketch] This can be proved using induction arguments showing that the iterations j of Algorithm 1 remain bounded. ■

Theorem 1: Let σ_0^2 be chosen such that $\sigma_0^2 \mathbf{1} \geq \mathbf{Y}^{\circ 2}$ and let $\gamma_*(\mathbf{X})$ denote the optimum of the prior-GP. Then, using the assumption of a small length scale, we can show that Algorithm 1 converges to the optimum, $\gamma_j(\mathbf{X}) \rightarrow \gamma_*(\mathbf{X})$ for $j \rightarrow \infty$.

Proof: [Sketch] This can be proved using the bounds established in Lemma 2 to show that the prior-GP is contracting with

$$\gamma_{j+1}(\mathbf{X}) - \gamma_*(\mathbf{X}) = \mathbf{A}_j (\gamma_j(\mathbf{X}) - \gamma_*(\mathbf{X})),$$

i.e., the eigenvalues of \mathbf{A}_j are inside the unit circle for every iteration j of Algorithm 1. ■

Remark 3: The design choice and condition for convergence $\sigma_0^2 \mathbf{1} \geq \mathbf{Y}^{\circ 2}$ makes sense intuitively, because it means that the prior σ_0^2 should be chosen to cover the square root of the data points y^i . In other words, the prior should have at least the support of the posterior.

V. SIMULATION RESULTS

This section presents validation studies. The first is a non-physical system meant to illustrate the efficacy of the proposed method. The second example shows how to leverage the state-dependent variance estimates for stochastic control.

A. Illustrative Estimation Example

We generate data according to the ground-truth function

$$y^i = 0.5(\cos x^i + 1) + \epsilon^i \quad (11a)$$

$$\epsilon^i \sim \mathcal{N}(0, 0.19x^i + 0.1), \quad (11b)$$

where the data points x^i are sampled from a uniform distribution $x^i \sim \mathcal{U}(0, 10)$. Next, we study the number of data samples required to accurately approximate both the mean and the noise distribution in (11).

1) *Qualitative results:* Fig. 2 illustrates a comparison of a classical implementation of a GP and the proposed solution for different numbers of data points. First, consider $N = 1$. As the one collected data point is far from zero (around the $2\sigma_0$ bound), the proposed solution does not trust the data point as much as the classical solution. Hence, in such an instance, the proposed solution is more conservative. Next, consider $N = 5$. Here, it can be seen that the ‘‘closeness’’ of the four measurements between $x \in [0, 5]$ gives the proposed GP confidence in its estimate. Consequently, the uncertainty is decreased more than for the classical GP. For $N = 20$, the proposed algorithm increases the uncertainty between $x \in [5, 10]$ as measurements are more spread out. Here, it can be seen that the classical GP underestimates the noise for $x \in [5, 10]$ and overestimates the noise for $x \in [0, 5]$. In contrast, the proposed GP correctly predicts the trend of the uncertainty in the data. For $N = 100$, the proposed solution accurately estimates both the mean function and its noise distribution, whereas the classical GP solution can only approximate the mean.

2) *Quantitative results:* Fig. 3 shows the convergence rate of the GP to the ground truth function, i.e., $\|\lambda(x) - h(x)\|_2$ and $\|\gamma(x) - g(x)\|_2$. It shows that the proposed GP with state-dependent noise is able to learn both the mean and the variance of the ground-truth function. First, the convergence rate of the mean estimate in the proposed method is faster than the standard GP implementation with a constant prior. The reason for the faster convergence rate is that the proposed GP implicitly learns that it can trust the measurements more if they are closer together, and also that it cannot trust the measurements as much if they are further apart. As $N \rightarrow \infty$, the mean estimates of both the classical GP and the proposed GP formulations converge to the ground truth. However, the variance estimates of the constant-prior implementation will inaccurately reflect the true uncertainty, which can be detrimental when used in combination with a feedback controller.

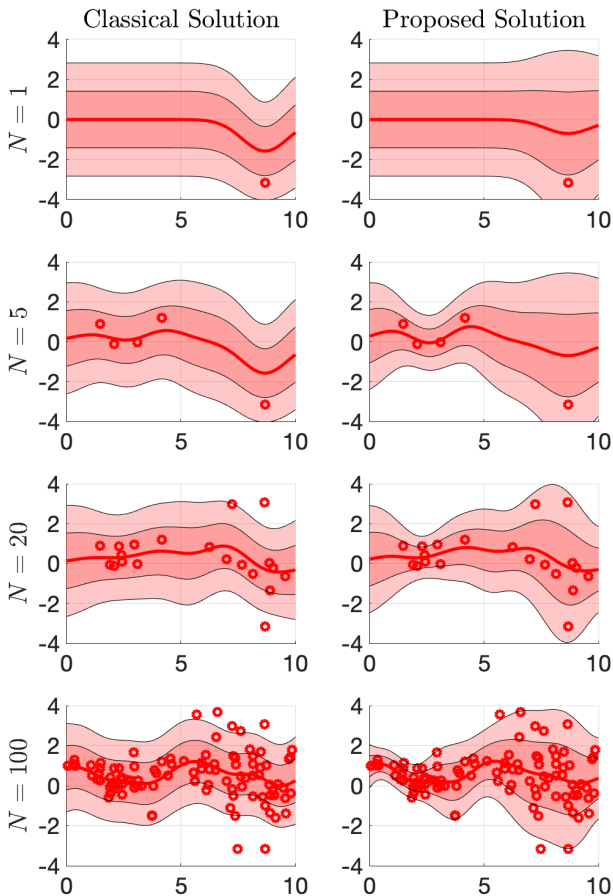


Fig. 2. Comparison of classical GP regression and the proposed GP regression with state-dependent noise. Left column: Classical GP regression. Right column: Proposed GP regression. The plot illustrates the mean in red as well as the 1σ bound and 2σ bound in different shades of red. The training data points are illustrated in and are the same for both classical GP regression and the proposed method.

B. Application to Stochastic Control

Next, we apply the proposed GP regression model to an illustrative stochastic control example. We employ the solver in [22]. However, any stochastic control method, see e.g., [4], can be utilized as long as the control formulation considers uncertainty bounds. We consider the dynamical system

$$\begin{bmatrix} p_{k+1} \\ v_{k+1} \end{bmatrix} = \begin{bmatrix} 1 & T_s \\ 0 & 1 \end{bmatrix} \begin{bmatrix} p_k \\ v_k \end{bmatrix} - \begin{bmatrix} 0 \\ T_s F(v_k) \end{bmatrix} + \begin{bmatrix} 0 \\ T_s \end{bmatrix} u_k,$$

with the a-priori unknown friction term $F(v_k)$ given by

$$F(v) \sim \mathcal{N}(h(v), g(v)), \quad (12a)$$

$$h(v) = v + \sin(3v), g(v) = 0.8 \max(0, v) + 0.2. \quad (12b)$$

Hence, the dynamics model studied in this section exhibits an uncertain friction term (12a), which we learn using the proposed algorithm. The GP model is subsequently used within a stochastic controller as state constraints, which are enforced as chance constraints,

$$\Pr(v_k \in X) \geq p_x, \quad (13)$$

with $X = [-1, 1]$ m/s and confidence interval $p_x = 95.44\%$, i.e., $\pm 2\sigma$ for a normal distribution. The results in the following use GP models learned using $N = 100$ data points,

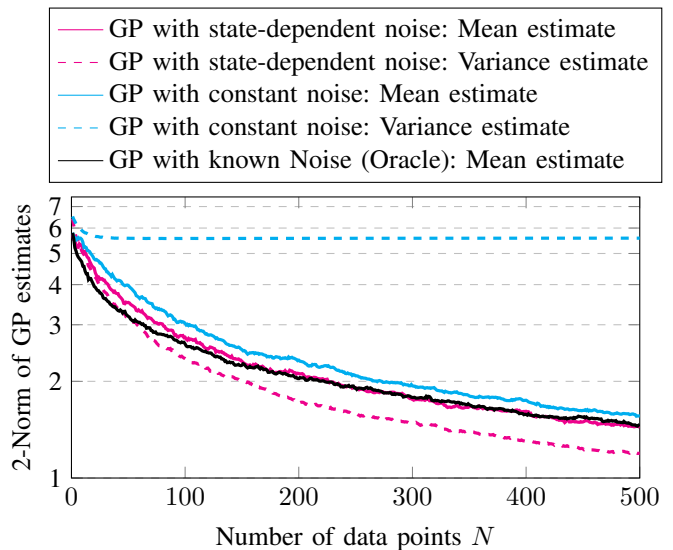


Fig. 3. Statistics of GP estimate accuracy. The plot illustrates the median of the 2-norm of 500 Monte-Carlo trials. The plot shows how quickly the mean estimate can decrease if the noise distribution were known (in black), which provides a bound on the convergence rate. Further, the plot shows two GP models: a GP with constant noise (in cyan) and the proposed GP with state-dependent noise (in magenta).

which are randomly generated according to $v^i \sim \mathcal{U}(-1, 1)$ for each Monte-Carlo trial. The optimal controller optimizes the cost function $\sum_{k=0}^N (10(p_k - p_{\text{ref}})^2 + u_k^2)$, where $p_{\text{ref}} = 1$ for $0-2s$ and $p_{\text{ref}} = 0$ after $2s$.

Fig. 4 shows statistics of 500 Monte-Carlo trials of the stochastic controller using three different GP models. It shows the stochastic controller using the proposed GP model and two GPs with constant priors as a comparison, one that uses a prior to remain cautious and one that uses a prior to be aggressive. The aggressive GP implementation uses a prior $\sigma_0^2 = 0.2$, which is the smallest noise in the data (12). The cautious GP uses $\sigma_0^2 = 1$, which is the highest noise level inside the chance constraints in (13). Fig. 4 illustrates the median of the position and the velocity, as well as the 2σ -bound of the velocity spread. It shows that the proposed GP implementation is cautious between $0-1s$, which is needed to satisfy the chance constraints in (13). During this time interval, this behavior resembles the cautious GP very closely (see middle plot). Next, Fig. 4 shows that the proposed GP implementation is aggressive during the time interval $2-3s$, which is caused by the low noise characteristics for negative velocities. During this time interval, this behavior resembles the aggressive GP very closely (see bottom plot). On the other hand, the aggressive GP is too aggressive during $0-1s$ and the cautious GP is not aggressive enough during $2-3s$. Overall, the proposed GP is cautious when its needed to satisfy (13), but is aggressive whenever possible.

Table I summarizes the closed-loop cost and percentages of constraint violations of the three GP implementations. As expected, the aggressive GP violates the chance constraints in (13) with an average 22% of constraints violations during intervals $0-1s$ and $2-3s$. Next, the cautious GP is overly conservative, which can be seen by its higher closed-loop




Fig. 4. Qualitative performance results of different GP models, i.e., three stochastic controllers. Top: Position achieved by the three different stochastic controllers. Middle: Velocity of controller using proposed GP with state-dependent noise inference (in black) and using GP with cautious choice of constant prior (in yellow). Bottom: Velocity of controller using proposed GP with state-dependent noise inference (in black) and using GP with aggressive choice of constant prior (in purple). The figure shows the mean of 500 Monte-Carlo runs as well as the 2σ bound.

TABLE I
COST AND CONSTRAINT VIOLATIONS

Approach	Closed-loop Cost			Constraint Violations during 0–1s and 2–3s
	Min.	Median	Max.	
Proposed	82.2	86.8	92.1	2%
Aggressive	77.1	81.3	86.6	22%
Cautious	89.0	94.1	100.2	1%

operating cost. The proposed GP’s constraint violations matches the 2σ -bound according to the requirements in the chance constraints in (13). Consequently, its operating cost are lower than the cautious GP.

VI. CONCLUSIONS

This paper considered a stochastic control formulation, in which the residual model uncertainty of the dynamical system consists of a nonlinear function and state-dependent noise. It presented an iterative algorithm to learn the unknown nonlinear function along with the state-dependent noise distribution using GP regression. The iterative algorithm used a posterior-GP to approximate the residual model uncertainty and a prior-GP to account for state-dependent noise. The iterative algorithm was shown to converge under simplifying assumptions. Simulation results showed that the proposed GP formulation enables faster convergence of the GP as the algorithm leverages the trustworthiness of the data by means of the prior. Further, simulation results showed the

advantages of using the proposed GP formulation within a stochastic control framework.

REFERENCES

- [1] L. Hewing, K. P. Wabersich, M. Menner, and M. N. Zeilinger, “Learning-based model predictive control: Toward safe learning in control,” *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 3, pp. 269–296, 2020.
- [2] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning*. The MIT Press, 2005.
- [3] L. Hewing, A. Liniger, and M. N. Zeilinger, “Cautious nm-pc with gaussian process dynamics for autonomous miniature race cars,” in *2018 European Control Conference (ECC)*, pp. 1341–1348, 2018.
- [4] 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, 2020.
- [5] J. Kabzan, L. Hewing, A. Liniger, and M. N. Zeilinger, “Learning-based model predictive control for autonomous racing,” *IEEE Robotics and Automation Letters*, vol. 4, no. 4, pp. 3363–3370, 2019.
- [6] S. Vaskov, R. Quirynen, M. Menner, and K. Berntorp, “Friction-adaptive stochastic predictive control for trajectory tracking of autonomous vehicles,” in *2022 American Control Conference (ACC)*, pp. 1970–1975, 2022.
- [7] S. Mosharafian, M. Razzaghpour, Y. P. Fallah, and J. M. Velni, “Gaussian process based stochastic model predictive control for cooperative adaptive cruise control,” in *2021 IEEE Vehicular Networking Conference (VNC)*, pp. 17–23, 2021.
- [8] S. Levine, Z. Popovic, and V. Koltun, “Nonlinear inverse reinforcement learning with Gaussian processes,” *Advances in neural information processing systems*, vol. 24, 2011.
- [9] S. Levine and V. Koltun, “Continuous inverse optimal control with locally optimal examples,” in *29th International Conference on International Conference on Machine Learning*, pp. 475–482, 2012.
- [10] C. Jidling, N. Wahlström, A. Wills, and T. B. Schön, “Linearly constrained Gaussian processes,” *Advances in neural information processing systems*, vol. 30, 2017.
- [11] H. Maatouk and X. Bay, “Gaussian process emulators for computer experiments with inequality constraints,” *Mathematical Geosciences*, vol. 49, pp. 557–582, 2017.
- [12] J. Matschek, A. Himmel, K. Sundmacher, and R. Findeisen, “Constrained Gaussian process learning for model predictive control,” *IFAC-PapersOnLine*, vol. 53, no. 2, pp. 971–976, 2020.
- [13] K. Berntorp and M. Menner, “Online constrained Bayesian inference and learning of Gaussian-process state-space models,” in *2022 American Control Conference (ACC)*, pp. 940–945, 2022.
- [14] L. P. Swiler, M. Gulian, A. L. Frankel, C. Safta, and J. D. Jakeman, “A survey of constrained Gaussian process regression: Approaches and implementation challenges,” *Journal of Machine Learning for Modeling and Computing*, vol. 1, no. 2, 2020.
- [15] S. Da Veiga and A. Marrel, “Gaussian process modeling with inequality constraints,” in *Annales de la Faculté des sciences de Toulouse: Mathématiques*, vol. 21, pp. 529–555, 2012.
- [16] S. Da Veiga and A. Marrel, “Gaussian process regression with linear inequality constraints,” *Reliability Engineering & System Safety*, vol. 195, p. 106732, 2020.
- [17] P. Goldberg, C. Williams, and C. Bishop, “Regression with input-dependent noise: A Gaussian process treatment,” *Advances in neural information processing systems*, vol. 10, 1997.
- [18] K. Kersting, C. Plagemann, P. Pfaff, and W. Burgard, “Most likely heteroscedastic gaussian process regression,” in *24th International Conference on Machine Learning*, pp. 393–400, 2007.
- [19] A. Boukouvalas, D. Cornford, and M. Stehlík, “Optimal design for correlated processes with input-dependent noise,” *Computational Statistics & Data Analysis*, vol. 71, pp. 1088–1102, 2014.
- [20] R. Miyagusuku, A. Yamashita, and H. Asama, “Gaussian processes with input-dependent noise variance for wireless signal strength-based localization,” in *2015 IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR)*, pp. 1–6, 2015.
- [21] Q.-H. Zhang and Y.-Q. Ni, “Improved most likely heteroscedastic gaussian process regression via Bayesian residual moment estimator,” *IEEE Trans. on Signal Processing*, vol. 68, pp. 3450–3460, 2020.
- [22] M. Menner, A. Chakrabarty, K. Berntorp, and S. Di Cairano, “Learning optimization-based control policies directly from digital twin simulations,” in *2022 IEEE Conference on Control Technology and Applications (CCTA)*, pp. 895–900, 2022.