

Real-Time Online Learning for Model Predictive Control using a Spatio-Temporal Gaussian Process Approximation

Lars Bartels, Amon Lahr, Andrea Carron and Melanie N. Zeilinger

Abstract—Learning-based model predictive control (MPC) can enhance control performance by correcting for model inaccuracies, enabling more precise state trajectory predictions than traditional MPC. A common approach is to model unknown residual dynamics as a Gaussian process (GP), which leverages data and also provides an estimate of the associated uncertainty. However, the high computational cost of online learning poses a major challenge for real-time GP-MPC applications.

This work presents an efficient implementation of an approximate spatio-temporal GP model, offering online learning at constant computational complexity. It is optimized for GP-MPC, where it enables improved control performance by learning more accurate system dynamics online in real-time, even for time-varying systems. The performance of the proposed method is demonstrated by simulations and hardware experiments in the exemplary application of autonomous miniature racing.

Video: <https://youtu.be/x4qo66R2-Ds>

I. INTRODUCTION

Model predictive control (MPC) is an advanced control strategy that optimizes control actions by predicting future system behavior based on a dynamic model [1]. By solving an optimization problem at each time step, MPC can handle constraints and adapt to changing conditions, making it highly effective for complex and safety-critical applications. However, the control performance of traditional MPC relies on the predictive capabilities of the underlying model of the system dynamics. For many nonlinear systems, deriving such an accurate model analytically can be impractical. This complication motivates the incorporation of data-driven models into the MPC framework [2] to enhance both performance and constraint satisfaction.

Gaussian processes (GPs) are a popular choice to model the mismatch between the nominal model and the true system dynamics in learning-based control, as they offer data-efficient model adaptation and provide formal uncertainty quantification. GP-based MPC (GP-MPC) leverages these properties to maintain high performance and safety under model mismatch and has already been successfully applied across various domains (cf. Tab. 2 in [3]), including autonomous racing [4], [5]. However, the cubic computational scaling of exact GP inference [6] remains a critical bottleneck for real-time GP-MPC. This issue is particularly acute in online learning scenarios, where the continuous accumulation

The authors are with the Institute for Dynamic Systems and Control, ETH Zürich, CH-8092 Zürich, Switzerland. This work is supported by the European Union's Horizon 2020 research and innovation programme, Marie Skłodowska-Curie grant agreement No. 953348, ELO-X.




Fig. 1. High-speed miniature racing despite a time-varying steering perturbation: Online learning using spatio-temporal Gaussian process approximations enables real-time adaptation to time-varying disturbances. The car aims to race along the track centerline (blue) while staying within track boundaries (red) and therefore plans its future trajectory (yellow) including uncertainty estimates based on the GP model of the residual dynamics.

of training data eventually leads to computational latencies that exceed the real-time constraints.

A. Approximate GP Inference for MPC

To maintain real-time feasible computations, [7], [8], [4] use a *subset of data* (SoD) approximation by conditioning the GP model on a *data dictionary* containing a fixed number of data points selected from the full training data set. A key challenge with this approach is to decide on which data points to keep in the dictionary based on how much relevant information they contribute. By replacing old data points in the dictionary with newer ones, [9], [10] allow the GP model to *evolve* [11], making it suitable for dealing with time-varying disturbances.

Inducing-point GP approximations employ a smaller set of strategically placed inducing points to summarize the training data. This reduces computational complexity while preserving essential information for accurate predictions [12], [13], [14]. In the online-learning setting, [8], [4] similarly employ a data dictionary to limit the computational complexity. The finite-dimensional nature of inducing-point GPs also allows for constant-time online updates by updating the inducing-point distribution [15], [16]. However, for time-varying models, the incurred limited model fidelity necessitates to update also the inducing-point locations to cover the extending input domain. Variational inducing-point approximations perform

this update by re-training on mini-batches of the updated training data set [17], [18], or by performing gradient steps on the “collapsed” evidence lower bound [19], [20], which adds to the computation time for online inference. Still, to the best of the authors’ knowledge, these methods have seen limited application in the domain of predictive control. In this regard, existing GP-MPC applications [21], [4] have resorted to heuristics, dynamically placing inducing points along the previously predicted trajectory, since these locations represent the most relevant region in time and space at the current MPC iteration. The dual GP approach proposed by [22] instead employs a time-invariant inducing-point GP for long-term learning, complemented by an additional short-term model adapting to variations over time online by recursively phasing out old data.

Spatio-temporal GPs as described in [23], [24], [25], [26] are naturally well-suited for application in time-varying settings. They exploit Markovian temporal covariance structures to rewrite the GP prior as a state-space model defined through a stochastic differential equation [27] according to their kernel, allowing to use Kalman filtering for learning at a constant computational cost per added data point. In the context of MPC, a similar formulation is used by [28]; however, it is limited to purely temporal data. The authors of [29] additionally include a discrete non-temporal *operating condition*, but they also do not allow for data in a continuous space as it naturally arises from a dynamic system.

To benefit from the complementary strengths of both inducing-point GPs in space and spatio-temporal GPs in time, the two approaches have been combined by [30], [31] into an approximate spatio-temporal GP model by constructing a Markov chain to track the evolution of a set of spatial inducing points over time.

B. Contribution

In this work, we develop an approximate spatio-temporal GP model based on [31], tailored to the application in a GP-MPC setting. The proposed approach enables efficient real-time inference and life-long online learning by recursively adapting to new data over time at a constant computational cost. Unlike the recursive short-term GP of [22], it supports more intricate temporal covariances, and in contrast to other work on spatio-temporal GP-MPC [28], [29], it allows to operate over a continuous spatial input space.

We provide an open-source implementation of this model to integrate into the `L4acados` GP-MPC framework [32], expanding its feature set by a GP model that is particularly suitable to online learning of a time-varying mismatch in the nominal dynamics. The effectiveness of the proposed method is demonstrated through experiments in the exemplary real-world application of autonomous miniature racing.

II. PROBLEM STATEMENT

We consider a system with the general nonlinear discrete-time dynamics

$$x(k+1) = f(x(k), u(k)) + Bg(x(k), u(k), t_k) + w(k), \quad (1)$$

where $x(k) \in \mathbb{R}^{n_x}$ denotes the system state and $u(k) \in \mathbb{R}^{n_u}$ is the control input at time t_k . The system dynamics are composed of the time-invariant nominal dynamics $f: \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \rightarrow \mathbb{R}^{n_x}$, the unknown and potentially time-varying residual dynamics $g: \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \times \mathbb{R}_+ \rightarrow \mathbb{R}^{n_g}$ with linear mapping $B \in \mathbb{R}^{n_x \times n_g}$, and random process noise w assumed as i.i.d. Gaussian with zero mean and covariance Σ_w .

The residual dynamics are modeled as a Gaussian process such that

$$g(x, u, t) \sim \mathcal{N}(\mu_g(x, u, t), \Sigma_g(x, u, t)), \quad (2)$$

where the mean μ_g and covariance matrix Σ_g can be evaluated at the given values for (x, u, t) .

A. Stochastic OCP Formulation

Following the usual MPC approach, we aim to control the system with dynamics (1) by optimizing state trajectories over a receding horizon of length T . Therefore, consider a stochastic OCP of the general form

$$\begin{aligned} \min_{x, u} \quad & \mathbb{E} \left[l_T(x_T) + \sum_{i=0}^{T-1} l_i(x_i, u_i) \right] \\ \text{s.t.} \quad & x_{i+1} = f(x_i, u_i) + Bg(x_i, u_i, t_i) + w_i, \\ & \mathbb{P}(h_j(x_i, u_i) \leq 0) \geq p_j \text{ for } j = 1, \dots, n_h \\ & x_0 = x(k). \end{aligned} \quad (3)$$

Here, the control task is to minimize an expected cost function defined by the sum of stage costs $l(x_i, u_i)$ and a terminal cost $l_T(x_T)$. At the same time, the system is subject to n_h distinct chance constraints, imposing that the individual probability of a constraint $h_j(x_i, u_i) \leq 0$ being satisfied is at least p_j .

The interplay of the nonlinear nominal dynamics and the stochastic residual dynamics results in future state distributions whose exact computation is generally intractable [33]. Hence, directly solving the stochastic OCP (3) is challenging without further simplification.

B. Deterministic OCP Formulation

In order to obtain a tractable OCP, we approximate the system state distribution as Gaussian, parametrized by its mean μ^x and covariance Σ^x as in [5]. Therefore, consider a first-order Taylor approximation around the mean μ^x of the uncertainty propagation

$$\Sigma_{k+1}^x = \Psi(\mu_k^x, \Sigma_k^x, u_k, t_k), \quad (4)$$

$$\Psi(\mu^x, \Sigma^x, u, t) = \tilde{A}\Sigma^x\tilde{A}^\top + B\Sigma_g(\mu^x, u, t)B^\top + \Sigma_w, \quad (5)$$

where $\tilde{A} = \frac{\partial}{\partial x}[f(x, u) + B\mu_g(x, u, t)]_{x=\mu^x}$ denotes the Jacobian of the nominal and GP mean dynamics.

Using the Gaussian approximation of the state distribution, the stochastic chance constraints can be reformulated by deterministic tightened constraints with the tightenings given according to the uncertainty in the predictions as

$$h_j^\Sigma(\mu^x, u, \Sigma^x) = \alpha_j \sqrt{C_j(\mu^x, u)\Sigma^x C_j^\top(\mu^x, u)}, \quad (6)$$

where $C_j(x, u) = \frac{\partial}{\partial x} h_j(x, u)$ and the tightening factor is given by the inverse cumulative density function of a standard Gaussian distribution as $\alpha_j = \Phi^{-1}(p_j)$.

In a common approximation to the expected cost [5], we consider only the mean state in the cost function, i.e., $\mathbb{E}[l_i(x_i, u_i)] \approx l_i(\mu_i^x, u_i)$. In combination with the linearized covariance propagation (5) and constraint tightenings (6), this results in a nonlinear OCP whose optimization variables (μ^x, Σ^x, u) are all defined deterministically.

C. Zero-Order Optimization

Although a deterministic OCP can be formulated with the approximations discussed above, actually solving it remains challenging due to the high number of optimization variables and computationally expensive GP evaluations. Thus, these approximations alone are often insufficient to enable a real-time capable controller.

The zero-order SQP algorithm presented in [34], [35], [36] is an approach to greatly reduce computational cost and still find a feasible, albeit suboptimal, solution to the OCP. This algorithm assumes the covariance propagation (5) to have zero gradient with respect to μ^x and u during optimization. Following the recent implementations [37], [32], we further neglect the gradients of the constraint tightenings (6). The zero-order approximation thereby allows the elimination of the state covariances Σ^x as optimization variables. Instead, Σ^x and the resulting values of h^Σ can be precomputed based on the previous iterate $(\hat{\mu}^x, \hat{u})$ of the optimizer. This results in the reduced-size OCP

$$\begin{aligned} \min_{\mu^x, u} \quad & l_T(\mu_T^x) + \sum_{i=0}^{T-1} l_i(\mu_i^x, u_i) \\ \text{s.t.} \quad & \mu_{i+1}^x = f(\mu_i^x, u_i) + B\mu_g(\mu_i^x, u_i, t_i), \\ & h_j(\mu_i^x, u_i) + \hat{\beta}_j^{(i)} \leq 0 \text{ for } j = 1, \dots, n_h, \\ & \mu_0^x = x(k), \end{aligned} \quad (7)$$

where $\hat{\beta}_j^{(i)} = h_j^\Sigma(\hat{\mu}_i^x, \hat{u}_i, \hat{\Sigma}_i^x)$ with $\hat{\Sigma}^x$ being computed according to (5) based on $(\hat{\mu}^x, \hat{u})$. The algorithm [34] then proposes to solve the OCP (7) iteratively with SQP, where the fixed constraint tightenings $\hat{\beta}$ are updated between SQP iterations.

III. REAL-TIME SPATIO-TEMPORAL GP-MPC

For enabling a real-time capable GP-MPC strategy with online learning, we combine the zero-order GP-MPC algorithm [36] with an approximate spatio-temporal GP model presented by [31]. On this foundation, we contribute an efficient spatio-temporal GP algorithm by specifically tailoring and optimizing the model for application in an MPC setting.

A. Approximate Spatio-Temporal GP Model

The residual dynamics (2) are modeled as a GP with the spatial input features $z = [x^\top \ u^\top]^\top$ and the temporal input feature t . We assume this GP to have zero prior mean for simplicity and its kernel to be separable as $k((z, t), (z', t')) = k^s(z, z')k^t(\tau)$ into a spatial kernel k^s and a stationary temporal kernel k^t with $\tau = t - t'$. This separability allows

us to treat the spatial and temporal component of the GP independently according to the approach outlined in [31].

The spatial component of the kernel is approximated using a set of inducing points V placed at M fixed locations in space. Let the inducing-point values $v_k \in \mathbb{R}^M$ at any fixed time t_k follow a Gaussian distribution, whose mean $\mu_v^{(k)}$ and covariance $\Sigma_v^{(k)}$ may change over time. Conditioned on v_k , the distribution of an observation y_k at spatial inputs Z_k and time t_k only depends on the spatial kernel component k^s by the GP conditional

$$y_k | v_k \sim \mathcal{N}(K_{Z_k V} K_{V V}^{-1} v_k, K_{Z_k Z_k} - Q_{Z_k Z_k} + \Sigma_\epsilon), \quad (8)$$

where K_{**} is a spatial covariance matrix such that $[K_{AB}]_{i,j} = k^s(A_i, B_j)$ for two sets (A, B) of spatial input points. Here, we define $Q_{Z_k Z_k} = K_{Z_k V} K_{V V}^{-1} K_{V Z_k}$ and denote by Σ_ϵ the measurement noise variance. In accordance with the standard inducing-point GP inference equations [12], the laws of total expectation and total variance yield the distribution of the function value $g(z, t) \sim \mathcal{N}(\mu_g(z, t), \Sigma_g(z, t))$ with

$$\begin{aligned} \mu_g(z, t) &= K_{Z V} K_{V V}^{-1} \mu_v^{(t)}, \\ \Sigma_g(z, t) &= K_{Z Z} - K_{Z V} K_{V V}^{-1} (K_{V V} - \Sigma_v^{(t)}) K_{V V}^{-1} K_{V Z}. \end{aligned} \quad (9)$$

In the temporal component, the procedure presented in detail in [25], [27] is applied. This approach uses the Fourier transform of the stationary temporal kernel by (approximately) letting $\mathcal{F}[k^t(\tau)](\omega) = W(i\omega)W(-i\omega)$ such that $W(s)$ is a rational transfer function representing a linear time-invariant state-space model with dynamics matrix F , input matrix G , and output matrix H . If k^t is a member of the half-integer Matérn kernels, an explicit formula to compute (F, G, H) is given in appendix A. The evolution over time of the inducing points is then described by the discrete-time linear Gaussian state-space model

$$\begin{aligned} \bar{v}_0 &\sim \mathcal{N}(0, K_{V V} \otimes P_\infty), \\ \bar{v}_{k+1} | \bar{v}_k &\sim \mathcal{N}((\mathbb{I}_M \otimes A_k) \bar{v}_k, K_{V V} \otimes Q_k), \\ v_k &= (\mathbb{I}_M \otimes H) \bar{v}_k, \end{aligned} \quad (10)$$

whose state transition matrices are defined as Kronecker products (denoted by \otimes) of $A_k = \exp(\Delta t_k F)$ and $Q_k = P_\infty - A_k P_\infty A_k^\top$, where $\Delta t_k = t_{k+1} - t_k$, $\exp(*)$ denotes a matrix exponential, and the stationary covariance P_∞ is given as the solution to the continuous-time Lyapunov equation $F P_\infty + P_\infty F^\top = -G G^\top$.

Since the stochastic state-space model (10) and the observation conditional (8) are both fully linear and Gaussian, the standard Kalman filtering technique [38] can be applied to recursively learn the inducing mean $\mu_v^{(k)}$ and covariance $\Sigma_v^{(k)}$. For simplicity of notation, define $\bar{A}_k = \mathbb{I}_M \otimes A_k$, $\bar{Q}_k = K_{V V} \otimes Q_k$, and $\bar{C}_k = K_{Z_k V} K_{V V}^{-1} (\mathbb{I}_M \otimes H)$. The recursive estimation scheme first predicts the new mean and covariance by applying the Markovian state transition model (10) to the previous estimate

$$\begin{aligned} \hat{\mu}_v^{(k)} &= \bar{A}_{k-1} \mu_v^{(k-1)}, \\ \hat{\Sigma}_v^{(k)} &= \bar{A}_{k-1} \Sigma_v^{(k-1)} \bar{A}_{k-1}^\top + \bar{Q}_{k-1}. \end{aligned} \quad (11)$$

Then, the Kalman gain

$$K_k = \hat{\Sigma}_{\bar{v}}^{(k)} \bar{C}_k^\top \left(\bar{C}_k \hat{\Sigma}_{\bar{v}}^{(k)} \bar{C}_k^\top + R_k \right)^{-1}, \quad (12)$$

where $R_k = K_{Z_k Z_k} - Q_{Z_k Z_k} + \Sigma_{\epsilon}$, allows to update the current estimate based on the most recent measurements y_k :

$$\begin{aligned} \mu_{\bar{v}}^{(k)} &= \hat{\mu}_{\bar{v}}^{(k)} + K_k \left(y_k - \bar{C}_k \hat{\mu}_{\bar{v}}^{(k)} \right), \\ \Sigma_{\bar{v}}^{(k)} &= \left(\mathbb{I}_{Md} - K_k \bar{C}_k \right) \hat{\Sigma}_{\bar{v}}^{(k)}. \end{aligned} \quad (13)$$

B. Spatio-Temporal GP Inference for MPC

For applicability of the spatio-temporal GP model presented above in an MPC setting, ensuring computational efficiency is of paramount importance in the algorithm design, as GP-MPC imposes strict real-time constraints on computations for GP learning and inference at each time step.

A common and powerful measure to improve efficiency is caching invariant and reoccurring expressions preemptively in order to avoid expensive repeating computations. For instance, the product $K_{ZV} K_{VV}^{-1}$ appears multiple times in the inference equations (9) and does not need to be recomputed.

In MPC, it usually holds that the step size Δt between consecutive time steps is constant. This allows for additional computational savings since the discrete-time state transition matrices A, Q from (10) in the GP model consequently also become invariant and can thus be cached. Also note that by choosing $\Delta t = 0$, we can model a time-invariant GP.

Beyond computational efficiency, prevention of numerical instabilities is another critical challenge. Recursive algorithms can be susceptible to small errors amplifying over time and culminating in inaccurate results. In Kalman filtering, such errors can cause the estimated covariance matrix to become indefinite, rendering the estimation algorithm invalid. To prevent this, the proposed algorithm takes inspiration from the square-root Kalman filter [39] and always computes the covariance matrix in its Cholesky decomposition. While this introduces some additional computational cost, it ensures that the covariance matrix remains positive semi-definite after each update, preserving the integrity of the Kalman filter.

The proposed algorithm of an approximate spatio-temporal GP model comprises three methods, which are presented in Alg. 1, 2, 3 respectively. The model is first initialized by computing the state-space representation of the temporal kernel, caching invariant terms, and obtaining the prior inducing-point distribution. Then, the model can be updated to transition to the next time step and incorporate new data if it is available. Finally, the GP can be evaluated at given spatial inputs to provide an uncertainty-aware residual dynamics estimate as required for stochastic MPC.

C. Python Implementation & Integration into *L4acados*

The proposed implementation is designed for compatibility with `GPYTORCH` [40], leveraging its core functionalities for evaluating GP kernel functions and Gaussian distributions. This integration allows us to rely on the `PYTORCH` [41] automatic differentiation feature for computing the Jacobian of GP predictions required by the optimization algorithm.

Algorithm 1 Initialize Approximate Spatio-Temporal GP

Require:

- GP with zero mean, spatial kernel k^s , temporal half-integer Matérn kernel k^t , noise σ_ϵ^2
- spatial inducing-point locations V
- time step size Δt

▷ state-space representation of temporal kernel

- 1: compute (F, G, H) from k^t by using (15)
- 2: obtain P_∞ by solving $FP_\infty + P_\infty F^\top = -GG^\top$

▷ cache invariant expressions

- 3: compute K_{VV} by evaluating k^s
- 4: $K_{VV}^{1/2} = \text{chol}(K_{VV})$
- 5: $\bar{H} = \mathbb{I}_M \otimes H$
- 6: $\bar{A} = \mathbb{I}_M \otimes A$ with $A = \exp(\Delta t \cdot F)$
- 7: $\bar{Q}^{1/2} = K_{VV}^{1/2} \otimes \text{chol}(Q)$ with $Q = P_\infty - AP_\infty A^\top$
- 8: $K_{VV}^{-1/2} = (K_{VV}^{1/2})^{-1}$
- 9: $L = K_{VV}^{-1/2} \bar{H}$

▷ initialize inducing-point distribution

- 10: $\mu_{\bar{v}} \leftarrow 0_{Md}$
 - 11: $\Sigma_{\bar{v}}^{1/2} \leftarrow K_{VV}^{1/2} \otimes \text{chol}(P_\infty)$
-

Algorithm 2 Update Approximate Spatio-Temporal GP

- 12: **procedure** UPDATEGP($Z \in \mathbb{R}^{n_z}, y \in \mathbb{R}^{n_y}$)

- 13: $\mu_{\bar{v}} \leftarrow \bar{A} \mu_{\bar{v}}$
 - 14: $\Sigma_{\bar{v}}^{\text{root}} \leftarrow \left[\bar{A} \Sigma_{\bar{v}}^{1/2} \quad \bar{Q}^{1/2} \right]$
 - 15: **if** new data (Z, y) is available **then**
 - 16: compute K_{ZZ} and K_{ZV} by evaluating k^s
 - 17: $Q_{ZZ}^{\text{root}} = K_{ZV} K_{VV}^{-1/2}$
 - 18: $R = K_{ZZ} - Q_{ZZ}^{\text{root}} (Q_{ZZ}^{\text{root}})^\top + \sigma_\epsilon^2 \mathbb{I}_{|Z|}$
 - 19: $\bar{C} = Q_{ZZ}^{\text{root}} L$
 - 20: $P = \bar{C} \Sigma_{\bar{v}}^{\text{root}}$
 - 21: $K = \hat{\Sigma}_{\bar{v}}^{\text{root}} P^\top (PP^\top + R)^{-1}$ ▷ Kalman gain
 - 22: $\mu_{\bar{v}} \leftarrow \mu_{\bar{v}} + K (y - \bar{C} \mu_{\bar{v}})$
 - 23: $\Sigma_{\bar{v}}^{\text{root}} \leftarrow \left[\Sigma_{\bar{v}}^{\text{root}} - KP \quad K \text{chol}(R) \right]$
 - 24: **end if**
 - 25: $\Sigma_{\bar{v}}^{1/2} \leftarrow \text{chol}(\Sigma_{\bar{v}}^{\text{root}} (\Sigma_{\bar{v}}^{\text{root}})^\top)$
 - 26: **end procedure**
-

Algorithm 3 Evaluate Approximate Spatio-Temporal GP

- 27: **function** EVALUATEGP($Z \in \mathbb{R}^{N \times n_z}$)
 - 28: compute K_{ZZ} and K_{ZV} by evaluating k^s
 - 29: $\bar{C} = K_{ZV} K_{VV}^{1/2} L$
 - 30: $\hat{\mu}_{\bar{v}}^{(0)} = \mu_{\bar{v}}$ and $\hat{\Sigma}_{\bar{v}}^{(0)} = \hat{\Sigma}_{\bar{v}}^{1/2} (\hat{\Sigma}_{\bar{v}}^{1/2})^\top$
 - 31: **for** $k = 1, \dots, N$ **do**
 - 32: $\hat{\mu}_{\bar{v}}^{(k)} = \bar{A} \hat{\mu}_{\bar{v}}^{(k-1)}$
 - 33: $\hat{\Sigma}_{\bar{v}}^{(k)} = \bar{A} \hat{\Sigma}_{\bar{v}}^{(k-1)} \bar{A}^\top + \bar{Q}$
 - 34: $m^{[k]} = \bar{C}_k \hat{\mu}_{\bar{v}}^{(k)}$ ▷ $\bar{C}_k = \bar{C}^{[k:]}$
 - 35: $S^{[k]} = \bar{C}_k \hat{\Sigma}_{\bar{v}}^{(k)} \bar{C}_k^\top$
 - 36: **end for**
 - 37: $\mu = [m^{[1]} \quad \dots \quad m^{[N]}]^\top$
 - 38: $\Sigma = K_{ZZ} - Q_{ZZ}^{\text{root}} (Q_{ZZ}^{\text{root}})^\top + \text{diag}(S^{[1]}, \dots, S^{[N]})$
 - 39: **return** $\mathcal{N}(\mu, \Sigma)$
 - 40: **end function**
-

Our Python implementation of an approximate spatio-temporal GP model can be accessed as a contribution to L4acados¹ [32], which provides a framework for learning-based MPC, including an efficient GP-MPC implementation utilizing the zero-order SQP optimization algorithm [36]. Within L4acados, the proposed model can be used interchangeably with the existing standard GPyTorch models to capture residual dynamics as a GP. Thereby, we expand the feature set of L4acados by a GP model that enables GP-MPC with online learning at a constant computational cost per time step without neglecting any data.

IV. AUTONOMOUS MINIATURE RACING

A. Control Task

To demonstrate the presented implementation, we apply a GP-MPC controller utilizing the proposed approximate spatio-temporal GP model in autonomous miniature racing simulations and on hardware². The environment for these experiments is provided by the CRS platform [42].

The car is controlled by a model predictive contouring control (MPCC) formulation as it has been presented in [43]. It aims to find the optimal torque a and steering angle δ to maximize progress along the track parametrized by θ while respecting the track boundaries. For our experiments, we use an MPCC controller with a horizon length of $T = 40$ that is called at a frequency of 30Hz. Accordingly, the MPC model dynamics are discretized with a $\frac{1}{30}$ s time step size.

The nominal car dynamics are adopted from [42] as a dynamic bicycle model with a simplified Pacejka tire friction model [44]. It has a six dimensional state containing the car's position (x_p, y_p) and orientation φ in the global reference frame as well as the associated velocities (v_x, v_y, ω) in the car's reference frame. Together with the control inputs (a, δ, θ) , which are added as an additional integrator state, this results in the system state $x = [x_p \ y_p \ \varphi \ v_x \ v_y \ \omega \ a \ \delta \ \theta]^\top \in \mathbb{R}^9$ and control input $u = [u_T \ u_\delta \ u_\theta]^\top \in \mathbb{R}^3$.

A time-varying model perturbation is introduced by dynamically altering the neutral steering offset δ_0 , as depicted in Fig. 2. This setup serves as a controllable and reproducible proxy for real disturbances, which allows to isolate and evaluate the controller's online learning capability. Note that this perturbation does not physically limit the vehicle's maneuverability. Since the full range of steering angles remains accessible, the theoretically optimal racing performance is preserved, provided the learning-based controller successfully identifies and compensates for the mismatch. Each run begins with δ_0 at its nominal value of zero, such that only inherent disturbances are present in the system. After 15s, the value of δ_0 starts to vary with an amplitude of 0.15rad.

The perturbation directly affects the true dynamics of the linear and angular velocities (v_x, v_y, ω) . Hence, the residual dynamics of these three states are modeled as




Fig. 2. Steering perturbation mapping parametrized by the neutral steering offset δ_0 (red dot). The shaded area illustrates the set of all employed perturbation mappings, with magenta indicating positive values and green indicating negative values of δ_0 .

a GP. The covariance function of this GP is defined as the product of a spatial RBF kernel with input features $(v_x, v_y, \omega, a, \delta)$ and a one-time-differentiable temporal Matérn kernel (exponent $\nu = 1.5$) with time t as its only input feature. For the learning-based MPCC, we then apply a zero-order GP-MPC strategy [32] with real-time iteration (RTI) [45], where the residual dynamics are modeled by the specified GP. GP hyper-parameters are commonly obtained via minimization of the negative marginal log-likelihood or the evidence lower bound [13], determining the parameters that best explain the training data. In our experiments, we have additionally found that choosing a “smoother” GP posterior mean and covariance improves convergence of the optimizer, rendering larger lengthscale hyperparameters desirable. Consequently, we select the GP hyperparameters via gradient-based optimization of the marginal log-likelihood, incorporating lower bounds on the lengthscale values. Hyperparameter values are the same across all GPs.

A controller baseline is provided by a nominal MPCC controller, as used for example in [43], [42], without any awareness on the mismatch between nominal model and real system. Depending on the perturbation, open-loop predictions by the MPCC may be highly inaccurate and the car must purely rely on controller feedback to prevent crashing. Besides the GP-MPC strategy utilizing our approximate spatio-temporal GP model implementation, we also apply an exact GP model with a spatio-temporal kernel and a conventional, purely spatial inducing-point GP model for reference. In order to maintain real-time capable online learning, the latter two rely on a subset-of-data approximation following the GP-MPC example from [32]. Regardless of which model is applied, the inclusion of a GP provides the MPC controller with an estimate on the prediction uncertainty. After the first lap has been completed, online learning is activated and the GP model successively incorporates newly gathered data points, enabling it to adapt its mean and variance estimates

¹Code: <https://github.com/IntelligentControlSystems/l4acados>

²Experiment data and video material: doi:10.3929/ethz-c-000796944, <https://gitlab.ethz.ch/ics/spatio-temporal-gp-mpc>




Fig. 3. Experimental GP-MPC solve times using different GP models with thick lines representing a rolling average of the last 100 solves. Left: Exact GP with a SoD approximation using the 400 most recent data points (blue), compared to simulation results using the full dataset (light blue). Right: Proposed approximate spatio-temporal GP model with 80 spatial inducing points based on hardware experiments (green).

of the residual dynamics to the real system.

B. Controller Performance Comparison

The following results were obtained from real-world racing experiments, conducted as described above.

1) *Computational Performance:* Due to the real-time constraint, we aim for the total GP-MPC computation time to remain below 30ms. With the zero-order SQP RTI algorithm, only a single quadratic subproblem of (7) is solved per time step. These computations are relatively fast (below 7ms), which leaves GP learning and inference as the main contributor to the total computational cost.

The solve times plotted in Fig. 3 compare GP-MPC using our approximate spatio-temporal GP implementation with standard exact GP-MPC. It is clear that continuously incorporating more data into an exact GP through online learning is not computationally feasible in the long term. For our experiments, we therefore used a SoD approximation, conditioning the GP only on the most recent 400 data points and discarding the rest, which allows to mostly satisfy the computational sub-30ms constraint. The approximate spatio-temporal GP is not limited through scaling with the amount of data points and it can be observed that solve times remain roughly constant at all times, allowing to run GP-MPC with online learning indefinitely without neglecting any data.

2) *Predictive Performance:* The efficacy of the GP’s online adaptation may be assessed by the one-step-prediction error between the car’s actual state and the MPC predicted state from the previous step. The plots in Fig. 4 show this error for the nominal MPC and the proposed spatio-temporal GP-MPC approaches. Depending on the steering perturbation currently present in the system, it can be seen how relying on the nominal model results in a significant model mismatch in the lateral velocity and angular velocity dynamics. Learning this nominal mismatch online with the proposed approximate spatio-temporal GP model implementation considerably reduces the MPC prediction error such that the remaining residual appears independently distributed.

While not shown in the plot, the exact GP model with spatio-temporal kernel achieves a similar prediction accu-




Fig. 4. One-step MPC prediction error over experiment runtime without a GP (red) and with online learning by an approximate spatio-temporal GP (blue) with a $2\text{-}\sigma$ confidence interval to model the residual dynamics. The bar up top illustrates the time-varying steering perturbation present in the system with the neutral steering offset δ_0 transitioning from zero through positive (green) and negative (magenta) values and back to zero.

racy despite the SoD truncation, whereas the purely spatial inducing-point GP produces inconsistent prediction errors since its kernel lacks the temporal dimension necessary to distinguish between varying disturbance states at the same spatial coordinates.

3) *Racing Performance:* How the reduction in prediction error affects racing performance in terms of lap times is shown in Fig. 5 for all four controller variants under the same steering perturbation.

For the first 15s, there is no steering perturbation present and all controllers result in a similar racing performance.

As soon as the perturbation commences, it is obvious that the nominal MPC is unable to maintain this performance level. In contrast, with the proposed approximate spatio-temporal GP model adapting to the changing conditions, the controller’s performance level remains largely unaffected by the perturbation. It can also be noted that the SoD exact GP model performs similarly well, indicating that its buffer of 400 data points is sufficient to carry all temporally

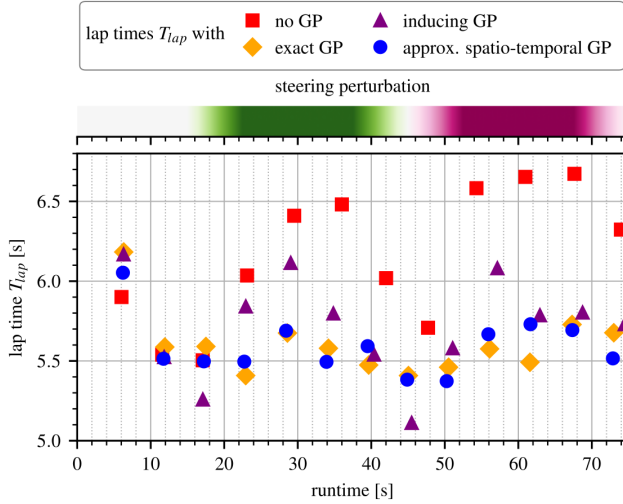


Fig. 5. Evolution of the most recent lap times over the experiment runtime under a time-varying steering perturbation for nominal MPC (red squares) and GP-MPC with an exact GP model (SoD) (orange diamonds), a spatial conventional inducing-point GP model (SoD) (purple triangles), and an approximate spatio-temporal GP model (blue circles). The exact GP and spatial inducing-point GP employ a subset of data approximation and are conditioned on the most recent 400 data points. Lap times are recorded and plotted at the moment the car has completed a lap. The bar up top illustrates the time-varying steering perturbation present in the system with green corresponding to a positive neutral steering offset δ_0 and magenta conversely indicating a negative value.

relevant information in this scenario. For the complete data persistence of our approach to offer a significant performance benefit here, the residual dynamics might have to evolve more slowly relative to the data accumulation rate, causing the buffer to discard relevant data points before the system could revisit those spatial locations.

More context to the differing racing performance is given by the trajectories of the car on track over the full experiment run for nominal MPC and approximate spatio-temporal GP-MPC shown in Fig. 6. It can be observed that the steering calibration perturbation clearly hinders the nominal controller to drive the car on a stable racing line. In particular, the car follows an undulating trajectory on the straights and the safety margin to the track boundary is violated in some corners. Meanwhile, the GP-MPC controller is able to adapt and compensate for the changing perturbation, which allows the car to maintain much more consistent racing lines throughout the experiment runtime.

V. CONCLUSION

The proposed spatio-temporal GP-MPC strategy offers a principled and computationally feasible approach to predictive control with online learning. The underlying approximate spatio-temporal GP model preserves the flexibility and probabilistic rigor of conventional GPs, but is naturally much better suited to continuous online learning since it can be conditioned on an arbitrary number of data points without negatively affecting the computational cost of inference. In contrast to other GP-based methods, this enables consistent

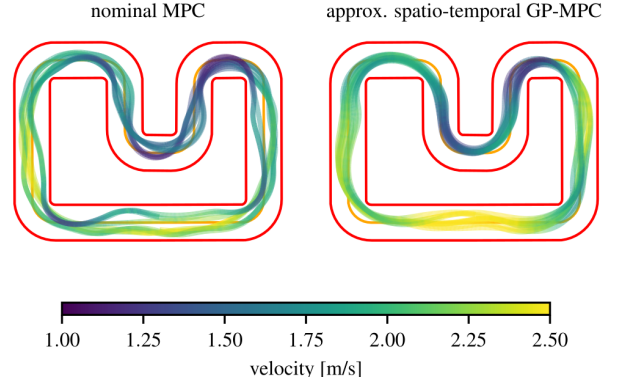


Fig. 6. Trajectory of racing car on track under a time-varying steering perturbation controlled by nominal MPC without GP (left) and approximate spatio-temporal GP-MPC with online learning (right).

real-time capability of closed-loop control without necessitating concessions such as discarding parts of the collected data. The resulting potential for performance benefits of the presented real-time implementation has been demonstrated in autonomous miniature racing experiments. For future work, it would be interesting to investigate how hyper-parameters can be similarly updated online, and theoretical closed-loop guarantees can be established.

APPENDIX

A. State-Space Representation of Matérn Kernel

For the stationary temporal kernel k^t of the spatio-temporal GP, we focus on the family of half-integer Matérn covariance functions as they permit an exact state-space representation. The covariance between two points in time separated by τ is then given as

$$k^t(\tau) = \frac{2^{1-\nu}}{\Gamma(\nu)} \left(\sqrt{2\nu} \frac{\tau}{\sigma_t} \right)^\nu K_\nu \left(\sqrt{2\nu} \frac{\tau}{\sigma_t} \right), \quad (14)$$

where Γ denotes the gamma function, K_ν denotes the modified Bessel function of the second kind, σ_t is the temporal length scale, and $\nu = D - 1/2$ for a positive integer D is a smoothness parameter such that the resulting GP is $D - 1$ times differentiable.

In accordance with [27] the state-space system matrices for such a Markovian kernel can be directly computed as

$$F = \begin{bmatrix} 0 & 1 & & 0 \\ \vdots & & \ddots & \\ 0 & 0 & & 1 \\ -a_1\gamma^D & -a_2\gamma^{D-1} & \cdots & -a_D\gamma \end{bmatrix}, \quad G = \begin{bmatrix} 0 \\ \vdots \\ 0 \\ 1 \end{bmatrix}, \quad (15)$$

$$H = [q \ 0 \ \cdots \ 0],$$

where $\gamma = \sqrt{2\nu}/\sigma_t$, $a_i = \binom{D}{i-1}$ denote binomial coefficients, and $q = \frac{((D-1)!)^2}{(2D-2)!} (2\gamma)^{2D-1}$ defines the diffusion coefficient.

REFERENCES

- [1] C. E. García, D. M. Prett, and M. Morari, "Model predictive control: Theory and practice—A survey," *Automatica*, vol. 25, no. 3, 1989.
- [2] L. Hewing, K. P. Wabersich, M. Menner, and M. N. Zeilinger, "Learning-Based Model Predictive Control: Toward Safe Learning in Control," *Annu. Rev. Control Robot. Auton. Syst.*, vol. 3, no. 1, 2020.
- [3] A. Scampicchio, E. Arcari, A. Lahr, and M. N. Zeilinger, "Gaussian processes for dynamics learning in model predictive control," 2025.
- [4] J. Kabzan, L. Hewing, A. Liniger, and M. N. Zeilinger, "Learning-Based Model Predictive Control for Autonomous Racing," *IEEE Robot. Autom. Lett.*, vol. 4, no. 4, 2019.
- [5] 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, 2020.
- [6] C. Williams and C. Rasmussen, "Gaussian Processes for Regression," *Advances in neural information processing systems*, vol. 8, 1995.
- [7] D. Nguyen-Tuong and J. Peters, "Incremental online sparsification for model learning in real-time robot control," *Neurocomputing*, vol. 74, no. 11, 2011.
- [8] A. Ranganathan, M.-H. Yang, and J. Ho, "Online Sparse Gaussian Process Regression and Its Applications," *IEEE Transactions on Image Processing*, vol. 20, no. 2, 2011.
- [9] M. Maiworm, D. Limon, and R. Findeisen, "Online learning-based model predictive control with Gaussian process models and stability guarantees," *International Journal of Robust and Nonlinear Control*, vol. 31, no. 18, 2021.
- [10] D. Bergmann, K. Harder, J. Niemeyer, and K. Graichen, "Nonlinear MPC of a Heavy-Duty Diesel Engine With Learning Gaussian Process Regression," *IEEE Trans. Contr. Syst. Technol.*, vol. 30, no. 1, 2022.
- [11] D. Petelin and J. Kocijan, "Evolving Gaussian process models for predicting chaotic time-series," in *2014 IEEE Conference on Evolving and Adaptive Intelligent Systems (EAIS)*, 2014.
- [12] J. Quinero-Candela and C. E. Rasmussen, "A unifying view of sparse approximate Gaussian process regression," *The Journal of Machine Learning Research*, vol. 6, 2005.
- [13] M. Titsias, "Variational Learning of Inducing Variables in Sparse Gaussian Processes," in *Proceedings of the Twelfth International Conference on Artificial Intelligence and Statistics*, ser. Proceedings of Machine Learning Research, vol. 5. Hilton Clearwater Beach Resort, Clearwater Beach, Florida USA: PMLR, 2009.
- [14] A. G. d. G. Matthews, J. Hensman, R. Turner, and Z. Ghahramani, "On Sparse Variational Methods and the Kullback-Leibler Divergence between Stochastic Processes," in *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics*, ser. Proceedings of Machine Learning Research, vol. 51. Cadiz, Spain: PMLR, 2016.
- [15] L. Csató and M. Opper, "Sparse On-Line Gaussian Processes," *Neural Computation*, vol. 14, no. 3, 2002.
- [16] H. Bijl, J.-W. van Wingerden, T. B. Schön, and M. Verhaegen, "Online sparse Gaussian process regression using FITC and PITC approximations," *IFAC-PapersOnLine*, vol. 48, no. 28, 2015.
- [17] J. Hensman, N. Fusi, and N. D. Lawrence, "Gaussian Processes for Big Data," 2013.
- [18] C.-A. Cheng and B. Boots, "Incremental variational sparse Gaussian process regression," *Advances in Neural Information Processing Systems*, vol. 29, 2016.
- [19] T. D. Bui, C. Nguyen, and R. E. Turner, "Streaming Sparse Gaussian Process Approximations," in *Advances in Neural Information Processing Systems*, vol. 30. Curran Associates, Inc., 2017.
- [20] W. J. Maddox, S. Stanton, and A. G. Wilson, "Conditioning Sparse Variational Gaussian Processes for Online Decision-making," in *Advances in Neural Information Processing Systems*, vol. 34. Curran Associates, Inc., 2021.
- [21] L. Hewing, A. Liniger, and M. N. Zeilinger, "Cautious NMPC with Gaussian Process Dynamics for Autonomous Miniature Race Cars," in *2018 European Control Conference (ECC)*, 2018.
- [22] Y. Liu, P. Wang, and R. Tóth, "Learning for predictive control: A Dual Gaussian Process approach," *Automatica*, vol. 177, p. 112316, Jul. 2025.
- [23] J. Hartikainen and S. Sarkka, "Kalman filtering and smoothing solutions to temporal Gaussian process regression models," in *2010 IEEE International Workshop on Machine Learning for Signal Processing*. Kittila, Finland: IEEE, 2010.
- [24] 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 Process. Mag.*, vol. 30, no. 4, 2013.
- [25] A. Carron, M. Todescato, R. Carli, L. Schenato, and G. Pillonetto, "Machine learning meets Kalman Filtering," in *2016 IEEE 55th Conference on Decision and Control (CDC)*. Las Vegas, NV, USA: IEEE, 2016.
- [26] M. Todescato, A. Carron, R. Carli, G. Pillonetto, and L. Schenato, "Efficient spatio-temporal Gaussian regression via Kalman filtering," *Automatica*, vol. 118, 2020.
- [27] S. Särkkä and A. Solin, *Applied Stochastic Differential Equations*, 1st ed. Cambridge University Press, 2019.
- [28] N. Schmid, J. Gruner, H. S. Abbas, and P. Rostalski, "A real-time GP based MPC for quadcopters with unknown disturbances," in *2022 American Control Conference (ACC)*, 2022.
- [29] Y. Li, R. Chen, and Y. Shi, "Spatiotemporal learning-based stochastic MPC with applications in aero-engine control," *Automatica*, vol. 153, 2023.
- [30] J. Hartikainen, J. Riihimäki, and S. Särkkä, "Sparse Spatio-temporal Gaussian Processes with General Likelihoods," in *Artificial Neural Networks and Machine Learning – ICANN 2011*. Berlin, Heidelberg: Springer, 2011.
- [31] W. Tebbutt, A. Solin, and R. E. Turner, "Combining pseudo-point and state space approximations for sum-separable Gaussian Processes," in *Proceedings of the Thirty-Seventh Conference on Uncertainty in Artificial Intelligence*. PMLR, 2021.
- [32] A. Lahr, J. Näf, K. P. Wabersich, J. Frey, P. Siehl, A. Carron, M. Diehl, and M. N. Zeilinger, "L4acadros: Learning-based models for acados, applied to Gaussian process-based predictive control," 2024.
- [33] A. Girard, C. Rasmussen, J. Q. Candela, and R. Murray-Smith, "Gaussian Process Priors with Uncertain Inputs Application to Multiple-Step Ahead Time Series Forecasting," in *Advances in Neural Information Processing Systems*, vol. 15. MIT Press, 2002.
- [34] A. Zanelli, J. Frey, F. Messerer, and M. Diehl, "Zero-Order Robust Nonlinear Model Predictive Control with Ellipsoidal Uncertainty Sets," *IFAC-PapersOnLine*, vol. 54, no. 6, 2021.
- [35] X. Feng, S. D. Cairano, and R. Quirynen, "Inexact Adjoint-based SQP Algorithm for Real-Time Stochastic Nonlinear MPC," *IFAC-PapersOnLine*, vol. 53, no. 2, 2020.
- [36] A. Lahr, A. Zanelli, A. Carron, and M. N. Zeilinger, "Zero-order optimization for Gaussian process-based model predictive control," *European Journal of Control*, vol. 74, 2023.
- [37] J. Frey, Y. Gao, F. Messerer, A. Lahr, M. Zeilinger, and M. Diehl, "Efficient Zero-Order Robust Optimization for Real-Time Model Predictive Control with acados," 2023.
- [38] R. E. Kalman, "A New Approach to Linear Filtering and Prediction Problems," *Journal of Basic Engineering*, vol. 82, no. 1, 1960.
- [39] P. Kaminski, A. Bryson, and S. Schmidt, "Discrete square root filtering: A survey of current techniques," *IEEE Transactions on Automatic Control*, vol. 16, no. 6, 1971.
- [40] J. R. Gardner, G. Pleiss, D. Bindel, K. Q. Weinberger, and A. G. Wilson, "GPyTorch: Blackbox Matrix-Matrix Gaussian Process Inference with GPU Acceleration," 2021.
- [41] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala, "PyTorch: An Imperative Style, High-Performance Deep Learning Library," in *Advances in Neural Information Processing Systems*, vol. 32. Curran Associates, Inc., 2019.
- [42] A. Carron, S. Bodmer, L. Vogel, R. Zurbrügg, D. Helm, R. Rickenbach, S. Muntwiler, J. Sieber, and M. N. Zeilinger, "Chronos and CRS: Design of a miniature car-like robot and a software framework for single and multi-agent robotics and control," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*, 2023.
- [43] A. Liniger, A. Domahidi, and M. Morari, "Optimization-based autonomous racing of 1:43 scale RC cars," *Optimal Control Applications and Methods*, vol. 36, no. 5, 2015.
- [44] R. Rajamani, *Vehicle Dynamics and Control*. Springer Science & Business Media, 2011.
- [45] M. Diehl, H. Bock, and J. Schlöder, "Real-Time Iterations for Non-linear Optimal Feedback Control," in *Proceedings of the 44th IEEE Conference on Decision and Control*, 2005.