

# Streaming Generated Gaussian Process Experts for Online Learning and Control

Zewen Yang<sup>1</sup>, Dongfa Zhang<sup>1</sup>, Xiaobing Dai<sup>2</sup>, Fengyi Yu<sup>1</sup>, Chi Zhang<sup>3</sup>, Bingkun Huang<sup>1</sup>,  
Hamid Sadeghian<sup>1</sup>, Sami Haddadin<sup>4</sup>

<sup>1</sup>Chair of Robotics and Systems Intelligence, Munich Institute of Robotics and Machine Intelligence, Technical University of Munich

<sup>2</sup>School of Computation, Information and Technology, Technical University of Munich

<sup>3</sup>Department of Informatics, School of Computation, Information and Technology, Technical University of Munich

<sup>4</sup>Mohamed bin Zayed University of Artificial Intelligence

zewen.yang@tum.de

## Abstract

Gaussian Processes (GPs), as a nonparametric learning method, offer flexible modeling capabilities and calibrated uncertainty quantification for function approximations. Additionally, GPs support online learning by efficiently incorporating new data with polynomial-time computation, making them well-suited for safety-critical dynamical systems that require rapid adaptation. However, the inference and online updates of exact GPs, when processing streaming data, incur cubic computation time and quadratic storage memory complexity, limiting their scalability to large datasets in real-time settings. In this paper, we propose a streaming kernel-induced progressively generated expert framework of Gaussian processes (SkyGP) that addresses both computational and memory constraints by maintaining a bounded set of experts, while inheriting the learning performance guarantees from exact Gaussian processes. Furthermore, two SkyGP variants are introduced, each tailored to a specific objective, either maximizing prediction accuracy (SkyGP-Dense) or improving computational efficiency (SkyGP-Fast). The effectiveness of SkyGP is validated through extensive benchmarks and real-time control experiments demonstrating its superior performance compared to state-of-the-art approaches.

**Code** — <https://github.com/Zewen-Yang/SkyGP>

**Extended version** — <https://arxiv.org/abs/2508.03679>

## 1 Introduction

Real-time learning has become essential for modeling the dynamics of physical systems, where machine learning models must be continuously updated to adapt to changing environments while satisfying the requirements of responsiveness and safety. This capability is especially critical for autonomous systems, such as underwater vehicles, aerial drones, and healthcare robots, that operate in complex and safety-critical environments (Yang et al. 2025a). In such scenarios, the ability to learn accurate models online and seamlessly integrate them into control loops is key to ensuring robust and efficient operation (Dai et al. 2023).

Gaussian Processes (GPs) provide a powerful nonparametric framework for modeling dynamical systems, particularly in safety-critical applications, due to their ability

to quantify uncertainty (Wachi et al. 2018). However, despite their modeling flexibility, standard GP inference suffers from severe scalability limitations. It scales cubically with the number of data points, requiring  $\mathcal{O}(N^3)$  time and  $\mathcal{O}(N^2)$  memory for  $N$  training data. This scalability bottleneck makes conventional GPs impractical for online applications involving continuously streaming data or long-term deployments.

To mitigate these limitations, a variety of scalable approximation techniques have been developed. Among global approximations, sparse Gaussian Processes introduce a set of  $M \ll N$  inducing points to efficiently summarize the training data, reducing the complexity to  $\mathcal{O}(NM^2)$  for training and  $\mathcal{O}(M^2)$  for prediction. Notable methods include fully independent training conditional approximation (Snelson and Ghahramani 2005), variational free energy (Naishguzman and Holden 2007), and latent projection techniques (Reeb et al. 2018). While streaming variants such as incremental sparse spectrum GP (ISSGP) (Gijssberts and Metta 2013) and streaming sparse GP (SSGP) (Bui, Nguyen, and Turner 2017) allow for incremental online updates, these frequently require computationally expensive optimization at each update, which can undermine their practical use in latency-sensitive settings. Moreover, these methods typically forgo guarantees on prediction error bounds, sacrificing reliability in critical applications.

Distributed Gaussian Processes (DGPs) provide complementary scalability by partitioning data across multiple processors or agents for parallel processing. Within this paradigm, mixture-of-experts (MoE) approaches (Tresp 2000b; Yuan and Neubauer 2008; Trapp et al. 2020) aggregate predictions from independently trained GP experts using fixed or learned weights. Extensions such as product-of-experts (PoE) (Cohen et al. 2020), generalized PoE (gPoE) (Cao and Fleet 2015), and correlated PoE with sparse GPs (Schürch et al. 2023) have been proposed to better utilize uncertainty in aggregation. The Bayesian committee machine (BCM) (Tresp 2000a) and its robust variant (rBCM) (Deisenroth and Ng 2015; Liu et al. 2018) explicitly integrate the GP prior to mitigate overconfidence. To address real-time learning problems within distributed frameworks, Lederer et al. (2021) proposed LoG-GP, which incrementally constructs a tree-structured ensemble of GP experts of-

fering a scalable solution for DGPs. However, the reliance on partitioning along a single dimension limits effectiveness in high-dimensional spaces. The structure of DGPs enables parallel and distributed deployment, with each GP model hosted on a separate computational node or agent. This is widely employed in multi-agent systems, where agents can independently operate local GPs and collaborate through a communication network for cooperative learning (Yang et al. 2021, 2024a,b; Lederer et al. 2023). Nevertheless, these methods overlook online learning requirements, particularly the need to efficiently incorporate new observations and adapt to nonstationary environments (Yuan and Zhu 2024; Dai et al. 2024, 2025). While Yang, Dai, and Hirche (2025) tackle asynchronous communication, their computation-aware random splitting strategy fails to exploit spatial or temporal correlations in the streaming data.

These limitations motivate our streaming kernel-induced progressively generated expert framework for Gaussian processes (SkyGP), which handles non-stationary, streaming data by dynamically allocating GP experts based on kernel similarity and temporal recency. Our main contributions:

- We propose a progressive expert generation strategy that leverages kernel-induced centers to determine whether incoming data should be incorporated into an existing expert model or used to initialize a new one. SkyGP enables dynamic partitioning of streaming data and addresses the limitations of the state-of-the-art (SOTA) approaches.
- We develop a time-aware and configurable expert aggregation framework that incorporates temporal weighting to effectively manage the generated GP experts. The framework adapts to system constraints, such as memory and computational budgets, while ensuring bounded complexity during both training and inference.
- We provide a learning-based policy for dynamical system control tasks with a rigorous theoretical analysis. The main theorem reveals the relationship between model uncertainty and control performance.
- We validate the proposed approach on real-world benchmark datasets and real-time control tasks. Extensive experiments demonstrate that the Sky-GP outperforms SOTA methods in terms of prediction accuracy, computational efficiency, and closed-loop control performance.

## 2 Problem Statement and Preliminaries

This paper investigates the problem of online function approximation in a streaming data environment for learning and control tasks. The objective is to infer an unknown target function  $f(\cdot) : \mathbb{R}^m \rightarrow \mathbb{R}$  using a estimated function  $\hat{f}(\cdot)$  based on given an sequential stream of input-output data pairs  $(\mathbf{x}^s, y^s)$  indexed by  $s \rightarrow \infty$ , where  $m, s \in \mathbb{N}_{>0}$ . Each input  $\mathbf{x}^s \in \mathbb{R}^m$  is paired with an output  $y^s = f(\mathbf{x}^s) + \varepsilon \in \mathbb{R}$ , where the noise  $\varepsilon$  follows a normal distribution with zero mean and variance  $\sigma_n^2$ , and  $\sigma_n \in \mathbb{R}_{>0}$ .

Gaussian Processes represent a Bayesian non-parametric approach commonly employed for function regression and approximation tasks. It assumes that the target function  $f(\cdot)$  is sampled from a GP prior, denoted as  $f(\cdot) \sim$

$\mathcal{GP}(m(\cdot), \kappa(\cdot, \cdot))$ , where  $m(\cdot)$  is the mean function (often set to zero when no prior knowledge is assumed) (Rasmussen and Williams 2006) and  $\kappa(\cdot, \cdot)$  is the covariance kernel (e.g., the squared exponential kernel). With the streaming observations accumulated up to time step  $t_k \in \mathbb{R}_{\geq 0}$  forming the dataset  $\mathbb{D}(t_k) = \{(\mathbf{x}^s, y^s)\}_{s=1,2,\dots,N(t_k)}$ , where  $N(t_k) = |\mathbb{D}(t_k)|$ , the posterior mean and variance functions of the GP can be derived as follows

$$\mu(\mathbf{x}) = \kappa(\mathbf{x}, \mathbf{X})\boldsymbol{\alpha}, \quad \sigma^2(\mathbf{x}) = \kappa(\mathbf{x}, \mathbf{x}) - \mathbf{v}^T \mathbf{v}, \quad (1)$$

where  $\kappa(\mathbf{x}, \mathbf{X}) = [\kappa(\mathbf{x}, \mathbf{x}^s)]_{s=1,\dots,N(t_k)}$  is the covariance vector between the test input  $\mathbf{x}$  and training inputs  $\mathbf{X} = [\mathbf{x}^1, \dots, \mathbf{x}^{N(t_k)}]$ , and  $\boldsymbol{\alpha}, \mathbf{v}$  are computed using the Cholesky decomposition  $\mathbf{L} = \text{cholesky}(\mathbf{K} + \sigma_n^2 \mathbf{I})$  as

$$\boldsymbol{\alpha} = \mathbf{L}^\top \backslash (\mathbf{L}^\top \backslash \mathbf{y}), \quad \mathbf{v} = \mathbf{L} \backslash \kappa(\mathbf{x}, \mathbf{X}) \quad (2)$$

with the training output  $\mathbf{y} = [y^1, \dots, y^{N(t_k)}]^\top$ . The kernel matrix  $\mathbf{K}$  is defined by its elements  $K_{ij} = \kappa(\mathbf{x}^i, \mathbf{x}^j)$  (Rasmussen and Williams 2006).

In online learning scenarios, the predictive mean and variance computations in (1) demonstrate  $\mathcal{O}(N(t_k))$  and  $\mathcal{O}(N^2(t_k))$  complexities respectively, indicating these operations computationally intractable as data grows without bound. To address these scalability challenges, a widely adopted strategy involves partitioning the data into multiple subsets and training separate GP experts with their means  $\mu_i(\cdot)$  and variances  $\sigma_i^2(\cdot)$ , where  $i \in \mathbb{N}$  denotes the index of the GP expert. Various aggregation techniques have been proposed for DGPs, and the structure of such aggregation methods can be formulated in a general framework as

$$\tilde{\mu}(\mathbf{x}) = \sum_{i \in \mathcal{N}} \omega_i(\mathbf{x}) \mu_i(\mathbf{x}), \quad (3a)$$

$$\tilde{\sigma}^2(\mathbf{x}) = \sum_{i \in \mathcal{N}} \varpi_i(\mathbf{x}) \sigma_i^2(\mathbf{x}), \quad (3b)$$

where  $\mathcal{N} \subset \mathbb{N}$  is the set of all GP experts, the functions  $\omega_i(\cdot) : \mathbb{R}^n \rightarrow \mathbb{R}_{\geq 0}$  and  $\varpi_i(\cdot) : \mathbb{R}^n \rightarrow \mathbb{R}_{\geq 0}$  represent the aggregation weights for the mean and variance, respectively.

## 3 Generated Experts of Gaussian Processes

In this section, we introduce an efficient online distributed learning framework designed to adaptively process streaming data in real time. At the core of SkyGP is a fully online and adaptive architecture that maintains a bounded, scalable set of GP experts. To enable dynamic expert allocation and effective aggregation, we first present the kernel-based center mechanism in Section 3.1. Building on this, we describe the progressive generation strategy for updating and predicting with the GP experts in Section 3.2, followed by the aggregation strategy in Section 3.3. Finally, we analyze the bounded computational complexity of SkyGP in Section 3.4.

### 3.1 Kernel-Induced Distance

To enable partitioning and localization for newly generated GP experts, we adopt a center-based representation of the training inputs, where the center of the  $i$ -th expert  $\mathcal{GP}_i$  is defined as  $\mathbf{c}_i = \frac{1}{N_i} \sum_{k=1}^{N_i} \mathbf{x}_i^k$ , following (Nguyen-tuong, Peters, and Seeger 2008). Each expert maintains its own representative center to support fast and scalable online allocation. To facilitate online computation, the center of each

---

Algorithm 1: Expert Localization with an Adaptive Window

---

**Require:**  $\mathbf{x}^k, \nu_{\text{prev}}, \{\vartheta_i^{k-1}\}_{i \in \mathcal{N}(t_{k-1})}$

**Step I: Compute adaptive window size**

- 1: **if**  $k > 1$  **then**
- 2:   Compute distance:  $d_{\text{temp}} \leftarrow 1/\kappa(\mathbf{x}^{k-1}, \mathbf{x}^k)$
- 3:   Window size:  $W \leftarrow \min(\bar{W}, \lfloor \exp(d_{\text{temp}}/\rho) \rfloor)$
- 4: **else**  $W \leftarrow 0$  **end if**

**Step II: Locate nearest expert**

- 5: **if**  $W \neq 0$  **then**
- 6:    $\mathcal{I} \leftarrow \{\nu_{\text{prev}} - W, \dots, \nu_{\text{prev}} + W\}$  from  $\mathcal{T}(t_{k-1})$
- 7:    $\mathcal{I} \leftarrow \{i \in \mathcal{I} \cap \mathcal{N}(t_{k-1}) \mid \vartheta_i^{k-1} > \bar{\vartheta}\}$
- 8: **else**  $\mathcal{I} \leftarrow \nu_{\text{prev}}$  **end if**
- 9:  $\nu_{\text{nr}} \leftarrow \arg \max_{i \in \mathcal{I}} \kappa(\mathbf{x}^k, \mathbf{c}_i)$ ;  $\nu_{\text{prev}} \leftarrow \nu_{\text{nr}}$
- 10:  $\mathcal{GP}_{\text{nr}}(t_k) \leftarrow$  Choose GP expert corresponding to  $\nu_{\text{nr}}$
- 11: **return**  $\mathcal{I}$  and  $\nu_{\text{prev}}$

---

expert is updated incrementally as new data points are assigned. Specifically, when a new point  $\mathbf{x}^k$  arriving at time step  $t_k$  is allocated to the selected expert, the expert center is updated as

$$\mathbf{c}_i^k = (k-1)\mathbf{c}_i^{k-1}/k + \mathbf{x}^k/k. \quad (4)$$

To evaluate the distance between the expert center  $\mathbf{c}_i^k$  and the newly arrived input  $\mathbf{x}^k$  in the feature space induced by the kernel, we define the kernel-based distance as follows

$$d_i^k(\mathbf{c}_i^k, \mathbf{x}^k) = 1/\kappa(\mathbf{c}_i^k, \mathbf{x}^k). \quad (5)$$

### 3.2 Progressive Expert Generation Strategy

**Expert Localization and Generation** In conventional online GP approximation methods, new data points are typically assigned to the first available expert that has not yet reached its capacity (Nguyen-tuong, Peters, and Seeger 2008; Lederer et al. 2021; Yang, Dai, and Hirche 2025). Once the expert becomes full, it is either no longer updated or directly split into sub-experts, often without explicitly accounting for the underlying distributional properties and space features of the data. In contrast, our SkyGP framework dynamically allocates each incoming data point to the most appropriate expert based on kernel proximity and time-aware factor. Specifically, from the prediction phase, we have already calculated the kernel distances between the new arriving point and the experts within an adaptive window  $W \in \mathbb{I}_{[1, \bar{W}]}$  and found the nearest experts with in the expert set  $\mathcal{N}(t_k)$ , where  $\mathcal{N}(t_0) = 1$  (see Algorithm 1). After locating the previous index  $\nu_{\text{prev}}$  of the nearest expert denoted by  $\nu_{\text{nr}}$ , we are able to process the streaming data based on the index list  $\mathcal{I}$  filtered by the time-aware factor  $\vartheta \in (0, 1]$  initialized by 1, which encodes the usage history of the experts. Moreover, we define the upper bound of  $\vartheta$  as  $\bar{\vartheta}$  to filter out the outdated experts. If the expert has not reached its capacity, the data point is added to it (see Figure 1). Otherwise, the two SkyGP variants handle this situation in two different ways outlined in Algorithm 2.

In particular, a data replacement strategy is employed in SkyGP-Dense. When the new data pair  $(\mathbf{x}^k, y^k)$  is confirmed to add to  $\mathbb{D}_{\text{nr}}(t_k)$ , and the data point within  $\mathcal{GP}_{\text{nr}}$  that




Figure 1: Data allocation process

is furthest from its center is cast off and moved to a separate dropped dataset  $\mathbb{D}_{\text{nr}}^{\text{off}}(t_k)$ , whose center  $\mathbf{c}_{\text{nr}}^{\text{off}, k}$  is then updated in the same method in (4). Since such a replacement requires recomputing the expert’s kernel matrix and Cholesky factorization, we introduce an event-triggered mechanism to limit updates only to critical cases. Let  $\mathbf{c}_{\text{nr}}$  and  $\mathbf{c}_a^{\text{off}}$  be the center and the dropped center of expert  $\mathcal{GP}_{\text{nr}}$  with  $N_{\text{nr}}(t_k)$  data points  $\{(\mathbf{x}_{\text{nr}}^s, y_{\text{nr}}^s)\}_{s=1}^{N_{\text{nr}}(t_k)}$ . The event-triggered data forwarding strategy is designed as

$$\mathbb{D}_{\text{nr}}(t_k) = \begin{cases} \mathbb{D}_{\text{nr}}(t_{k-1}) \cup (\mathbf{x}^k, y^k) \setminus (\mathbf{x}^{\text{off}}, y^{\text{off}}), & \text{if } \Delta < 0 \\ \mathbb{D}_{\text{nr}}(t_{k-1}), & \text{otherwise} \end{cases}, \quad (6)$$

where  $\Delta \in \mathbb{R}$  is defined as  $\Delta = \max_{s=1, \dots, N_{\text{nr}}(t_k)} \Delta(s)$  and

$$\Delta(s) = \kappa(\mathbf{x}^s, \mathbf{c}_{\text{nr}}) - \kappa(\mathbf{x}^s, \mathbf{c}_{\text{nr}}^{\text{off}}) - \kappa(\mathbf{x}^k, \mathbf{c}_{\text{nr}}) + \kappa(\mathbf{x}^k, \mathbf{c}_{\text{nr}}^{\text{off}}).$$

This kernel-based selection strategy encourages experts to retain points that align with the current local distribution while avoiding repeated inclusion of previously discarded regions. Although computationally expensive, this step is infrequent and remains tractable under the predefined data threshold  $\bar{N}$ . If no expert accepts  $\mathbf{x}^k$ , a new expert will be created and the center will be initialized with  $\mathbf{x}^k$ . In SkyGP-Fast, no replacement is performed. This design ensures all update GP expert is a rank-one update to the Cholesky factor, which incurs a computational complexity of  $\mathcal{O}(N_{\text{nr}}^2)$ . SkyGP-Fast bypasses recomputations by appending new experts instead of replacing data, favoring low-latency updates and scalability. In contrast, SkyGP-Dense limits memory usage by reusing expert slots, making it more suitable for resource-constrained settings. Together, these two modes offer a flexible trade-off between computational efficiency and memory control across diverse online learning scenarios.

**Dynamic Expert List Creation** The expert search is implemented over a center-indexed list that maintains a globally ordered list of expert centers based on their insertion positions. The index of this list is defined as  $\nu_i \in \mathbb{N}$  with corresponding  $\mathcal{GP}_{\nu_i}$  expert. The newly instantiated expert with center  $\mathbf{x}^k$  will be inserted adjacent to the nearest existing expert identified during search. To update the table when new GP expert generated, we compare the kernel-based distance between  $\mathbf{x}^k$  and the left/right neighbors of the selected experts with their corresponding center  $\mathbf{c}_{\text{nr}}^-$  and  $\mathbf{c}_{\text{nr}}^+$ :

$$d_{\text{left}} = 1/\kappa(\mathbf{x}^k, \mathbf{c}_{\text{nr}}^-), \quad d_{\text{right}} = 1/\kappa(\mathbf{x}^k - \mathbf{c}_{\text{nr}}^+) \quad (7)$$

---

**Algorithm 2: Expert Update Strategy**


---

**Require:**  $\mathbf{x}^k, \nu_{\text{prev}}, \mathcal{I}, \{\vartheta_i^{k-1}\}_{i \in \mathcal{N}(t_{k-1})}$

**Step I: Aggregated Prediction**

- 1:  $\mathcal{I}_{\text{temp}} \leftarrow \{\nu \in \mathcal{I} \mid \vartheta_\nu^{k-1} > \bar{\vartheta}\}$
- 2:  $\mathcal{I}_{\text{agg}} \leftarrow$  indices of the largest  $\bar{N}$  values in  $\mathcal{I}_{\text{temp}}$
- 3:  $\tilde{\mu}(\mathbf{x}^k), \tilde{\sigma}(\mathbf{x}^k) \leftarrow$  (3) with MOE/POE/BCM (9) to (11)

**Step II: GP Expert Model Update**

- 4:  $\{\vartheta_i^k\}_{i \in \mathcal{N}} \leftarrow \{\vartheta_i^{k-1}\}_{i \in \mathcal{N}}$
- 5: **for all**  $\nu \in \mathcal{I}_{\text{agg}}$  **do**
- 6:    $\vartheta_\nu^k \leftarrow 1, \mathbb{D}_{\text{temp}} \leftarrow \{\mathbf{x} \in \mathbb{D}_\nu(t_{k-1}) \mid \kappa(\mathbf{x}, \mathbf{c}_\nu) > \kappa(\mathbf{x}^k, \mathbf{c}_\nu)\}$
- 7:   **if**  $|\mathbb{D}_\nu(t_{k-1})| < \bar{N}$  **then**
- 8:      $\mathbb{D}_\nu(t_k) \leftarrow \mathbb{D}_\nu(t_{k-1}) \cup (\mathbf{x}^k, y^k)$
- 9:     Update  $\mathcal{GP}_\nu(t_k) \leftarrow$  Cholesky rank-one update
- 10:    **return**  $\{\vartheta_i^k\}_{i \in \mathcal{N}(t_k)}$
- 11:   **else if** SkyGP-Dense and  $\mathbb{D}_{\text{temp}} = \emptyset$  **then**
- 12:     Update  $\mathbb{D}_\nu(t_k) \leftarrow$  (6),  $\mathcal{GP}_\nu(t_k) \leftarrow$  (2)
- 13:     Update  $\text{off } \mathbf{c}_\nu \leftarrow$  (4) with dropped data
- 14:    **return**  $\{\vartheta_i^k\}_{i \in \mathcal{N}(t_k)}$
- 15:   **end if**
- 16: **end for**
- 17: Create  $\mathcal{GP}_{|\mathcal{N}(t_{k-1})|+1}$  with dataset  $\mathbb{D}(t_k) = (x_t, y_t)$
- 18: Update expert list  $\mathcal{T}(t_k) \leftarrow \mathcal{T}(t_{k-1})$  and (8)
- 19: Update  $\mathcal{N}(t_k) \leftarrow \{\mathcal{N}(t_{k-1}), |\mathcal{N}(t_{k-1})| + 1\}$

**Step III: Update Time-Aware Factor**

- 20: **for**  $i = \mathcal{N} \setminus \mathcal{I}_{\text{agg}}$  **do**  $\vartheta_i^k \leftarrow \rho \vartheta_i^{k-1}$  **end for**
- 21: **return**  $\{\vartheta_i^k\}_{i \in \mathcal{N}(t_k)}$

---

According to the previous  $\nu_{\text{nr}}$  in Algorithm 1, the new index of  $\mathcal{GP}_{\text{new}}$  is

$$\nu_{\text{new}} = \begin{cases} \min(\nu_{\text{nr}} |\mathcal{N}|), & \text{if } d_{\text{right}} < d_{\text{left}} \\ \max(\nu_{\text{nr}} - 1, 0) & \text{otherwise} \end{cases}. \quad (8)$$

### 3.3 Aggregated Prediction

Leveraging DGP framework in (3), our proposed SkyGP can be seamlessly integrated with existing aggregation strategies to enhance predictive performance and uncertainty estimation across decentralized data streams. These weights are typically chosen to satisfy certain properties, such as non-negativity and normalization, ensuring that the aggregated predictions are coherent and interpretable. For example, the aggregated weight functions of MoE are

$$\omega_i(\mathbf{x}) = w_i, \quad (9a)$$

$$\varpi_i(\mathbf{x}) = w_i(\sigma_i^2(\mathbf{x}) + \mu_i^2(\mathbf{x})) - \tilde{\mu}^2(\mathbf{x})/w_i, \quad (9b)$$

satisfying  $\sum_{i \in \mathcal{I}_{\text{agg}}} w_i = 1$ , which is also utilized in the mixture of explicitly localized experts. Considering the posterior variance in GP experts, the PoE family methods employ the following aggregation structure

$$\omega_i(\mathbf{x}) = w_i \sigma_i^2(\mathbf{x}) / \varpi_i(\mathbf{x}), \quad (10a)$$

$$\varpi_i(\mathbf{x}) = 1 / \sum_{s \in \mathcal{I}_{\text{agg}}} w_s \sigma_s^{-2}(\mathbf{x}). \quad (10b)$$

Furthermore, these aggregation schemes can be extended to incorporate the prior variance of the unknown function, denoted as  $\sigma_*$ , to further refine the combined predictive uncertainty in the BCM family approaches. In this context, the

weight function for the mean remains consistent with that of the PoE in (10a), while the corresponding weight function for the aggregated posterior variance is given by

$$\varpi_i(\mathbf{x}) = 1 / \left( \sum_{s \in \mathcal{I}_{\text{agg}}} w_s \sigma_s^{-2}(\mathbf{x}) + (1 - \sum_{j \in \mathcal{I}_{\text{agg}}} w_j) \sigma_*^{-2} \right). \quad (11)$$

This extension ensures that the aggregated predictions remain well-calibrated, especially in distributed or federated learning scenarios where each expert may have access to different subsets of the data.

According to Algorithm 1, given a new input  $\mathbf{x}^k$ , SkyGP performs prediction by first identifying the definable  $\bar{N}$  number of experts within the search window. Based on the returned aggregation expert list  $\mathcal{I}_{\text{agg}}$ , each expert predicts the posterior mean  $\mu_i(\mathbf{x}^k)$  and variance  $\sigma_i^2(\mathbf{x}^k)$ . These predictions then are fused using a principled weighting strategy, e.g., MoE, PoE or BCM, yielding the final prediction in (3).

### 3.4 Bounded Complexity Analysis

In this section, we formalize the computational efficiency of the proposed SkyGP framework, which is designed to maintain bounded per-step complexity during both model update and prediction. The primary sources of computational cost include nearest expert localization, aggregated prediction, and expert model update. Each component is explicitly designed to ensure low complexity and scalability in streaming scenarios, while maintaining high prediction accuracy.

**Model Update** For expert localization, as described in Algorithm 1, the computational cost from evaluating kernel similarities between the current input and the expert centers to obtaining index set  $\mathcal{I}$  (line 1-8) is  $\mathcal{O}(1)$ . The subsequent selection of the nearest expert from the candidate set  $\mathcal{I}$  in line 9, even when incorporating the time-aware factor, requires only a linear search over  $\mathcal{I}$  with the same order of complexity, i.e.,  $\mathcal{O}(W)$ . Consequently, the overall computational complexity of Algorithm 1 is  $\mathcal{O}(W)$ .

The overall computational cost of the Algorithm 2 is dominated by three components. First, the selection of the aggregated expert set in line 2 requires filtering the candidate index set  $\mathcal{I}$  and selecting the top  $\bar{N}$  elements, which can be performed in  $\mathcal{O}(|\mathcal{I}| \log \bar{N})$  time. Second, the update phase iterates over the  $\bar{N}$  aggregated experts. For each, the construction of the temporary data subset incurs  $\mathcal{O}(|\mathbb{D}_\nu(t_{k-1})|)$  cost, which is bounded by  $\mathcal{O}(\bar{N})$  considering  $|\mathbb{D}_\nu(t_k)| \leq \bar{N}$ . Moreover, when a GP model is updated via a rank-one Cholesky modification for SkyGP-Fast, an additional  $\mathcal{O}(\bar{N}^2)$  cost per expert is incurred. On the other hand, SkyGP-Dense requires full Cholesky factor recomputation when performing replacement operations, incurring a higher cost of  $\mathcal{O}(\bar{N}^3)$ . Third, the time-aware weighting update over the remaining experts requires  $\mathcal{O}(\bar{N})$  operations. Hence, the total worst-case complexity can be expressed as  $\mathcal{O}(W \log \bar{N} + \bar{N} \bar{N}^2)$  for SkyGP-Fast, and similarly  $\mathcal{O}(W \log \bar{N} + \bar{N} \bar{N}^3)$  for SkyGP-Dense.

**Prediction** Aggregated mean and variance in SkyGP are computed using the bounded top  $\bar{N}$  nearest experts. As each local GP performs inference with  $\mathcal{O}(N^2(t_k))$  complexity, the overall prediction complexity per step is  $\mathcal{O}(\bar{N} N^2(t_k))$ .

## 4 Safe Learning-based Control Policy

In this section, we present a learning-based control policy that integrates the proposed SkyGP framework into nonlinear systems with unknown dynamics. The key idea is to leverage real-time GP predictions to estimate model uncertainties and design a feedback controller that ensures safe and stable trajectory tracking. Following the introduction of the prediction error bound of SkyGP in Section 4.1, we analyze the control performance for general dynamical systems and design a tracking control policy specifically for Euler–Lagrange (EL) systems in Section 4.2.

### 4.1 Learning Performance of SkyGP

We consider a function  $f$  belonging to a reproducing kernel Hilbert space (RKHS)  $\mathcal{H}_\kappa$  associated with a positive definite kernel  $\kappa(\cdot, \cdot)$ , such that  $\|f\|_\kappa \leq \Gamma$ . Let  $\mu_i(\mathbf{x})$  and  $\sigma_i(\mathbf{x})$  denote the posterior mean and standard deviation from the  $i$ -th GP expert trained on dataset  $\mathbb{D}_i$  for  $i = 1, \dots, \mathcal{N}$ . Then the following probabilistic prediction error bound holds.

**Lemma 1.** *Consider the regression task in a compact domain  $\mathbb{X}$ , and suppose the kernel function  $\kappa(\cdot, \cdot)$  is Lipschitz with Lipschitz constant  $L_\kappa \in \mathbb{R}_{0,+}$ . Choose  $\delta \in (0, 1/\mathcal{N})$  and  $\tau \in \mathbb{R}_+$ , then the prediction error satisfies*

$$|f(\mathbf{x}) - \tilde{\mu}(\mathbf{x} | \mathbb{D})| \leq \beta\sigma(\mathbf{x}) + \gamma(\mathbf{x}), \quad \forall \mathbf{x} \in \mathbb{X} \quad (12)$$

with a probability of at least  $1 - \mathcal{N}\delta$ , where  $\gamma(\mathbf{x}) = \sum_{i \in \mathcal{N}} \omega_i(\mathbf{x})\gamma_i$ ,  $\sigma(\mathbf{x}) = \sum_{i \in \mathcal{N}} \omega_i(\mathbf{x})\sigma_i(\mathbf{x})$ , and

$$\beta = 2 \left( 2 \sum_{j=1}^n \log \left[ \frac{\sqrt{n}}{2\tau} (\bar{x}_j - \underline{x}_j) \right] - 2 \log \frac{\delta}{n} \right)^{1/2}, \quad (13)$$

$$\gamma_i = (\sqrt{\beta_\delta} L_{\sigma,i} + \Gamma \sqrt{2L_\kappa + L_{\mu,i}}) \tau, \quad (14)$$

with  $\bar{x}_j = \max_{x \in \mathbb{X}} x_j$ ,  $\underline{x}_j = \min_{x \in \mathbb{X}} x_j$  and  $x_j$  as the  $j$ -th dimension of  $\mathbf{x}$ . The positive constants  $L_{\mu,i} \in \mathbb{R}_{0,+}$  and  $L_{\sigma,i} \in \mathbb{R}_{0,+}$  are the Lipschitz constant for posterior mean  $\mu_i(\cdot)$  and variance  $\sigma_i(\cdot)$ , respectively for  $i \in \mathcal{N}$ .

Lemma 1 shows the probabilistic theoretical error bound of the aggregated prediction  $\tilde{\mu}(\cdot)$ , which also holds for the proposed Sky-GP<sup>1</sup>. Note that the boundness of the prediction error does not rely on the choice of aggregating strategy reflected by the weighting function  $\omega_i(\cdot)$ . For notational simplicity, denote the prediction error bound as  $\eta(\mathbf{x}) = \beta\sigma(\mathbf{x}) + \gamma(\mathbf{x})$ , which is used in Section 4.2.

**Remark 1.** Compared with the other probabilistic (Srinivas et al. 2012; Scharnhorst et al. 2022) or deterministic bounds (Maddalena, Scharnhorst, and Jones 2021; Hashimoto et al. 2022), the error bound chosen in this paper follows (Lederer, Umlauf, and Hirche 2019), which results in an amplitude constant  $\beta$  of posterior variance  $\sigma(\cdot)$ . Moreover, it is shown that there always exists a sufficiently small grid factor  $\tau \in \mathbb{R}_+$  such that  $\gamma(\mathbf{x}) \ll \beta\sigma(\mathbf{x})$  for all  $\mathbf{x} \in \mathbb{X}$ , i.e., the posterior variance dominates the error bound  $\eta(\mathbf{x})$ . Therefore, the prediction error bound could also be written as  $\eta(\mathbf{x}) \leq 2\beta\sigma(\mathbf{x})$  (Lederer et al. 2024).

<sup>1</sup>The proofs of all lemmas and theorems are provided in Appendix A in (Yang et al. 2025b).

## 4.2 Safe Critical Control Application

**A General Control System** The objective is to design a GP-based control law that drives an unknown dynamical system to exhibit a desired behavior. Specifically, we consider a class of control applications focused on output stabilization, where the goal is to ensure that the system output converges to the equilibrium point asymptotically. The continuous-time nonlinear system is described as follows

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}, \mathbf{u}) + \mathbf{g}(\mathbf{x}, \mathbf{u}), \quad \mathbf{z} = \mathbf{h}(t, \mathbf{x}) \quad (15)$$

where  $\mathbf{x} \in \mathbb{X} \subset \mathbb{R}^m$ ,  $\mathbf{u} \in \mathbb{U} \subset \mathbb{R}^{m_u}$  and  $\mathbf{z} \in \mathbb{R}^{m_z}$  denote the system state, control input and system output, respectively. The transition function  $\mathbf{f}(\cdot, \cdot) = [f_1(\cdot, \cdot), \dots, f_m(\cdot, \cdot)]^\top : \mathbb{X} \times \mathbb{U} \rightarrow \mathbb{R}^m$  is unknown, where each scalar  $f_j(\cdot, \cdot)$  belongs to a RKHS  $\mathcal{H}_{\kappa,j}$  corresponding to a kernel function  $\kappa_j(\cdot, \cdot) : \mathbb{X}_e \times \mathbb{X}_e \rightarrow \mathbb{R}_{0,+}$  with  $\mathbb{X}_e = \mathbb{X} \times \mathbb{U} \subset \mathbb{R}^{m+m_u}$  for all  $j = 1, \dots, m$ . The measurement function  $\mathbf{h}(\cdot, \cdot) : \mathbb{R}_{\geq 0} \times \mathbb{X} \rightarrow \mathbb{R}^{m_z}$  is known. To achieve system output stabilization, i.e.,  $\lim_{t \rightarrow \infty} \mathbf{z}(t) = \mathbf{0}_{m_z \times 1}$ , a control policy  $\pi(\cdot, \cdot, \cdot) : \mathbb{R}_{\geq 0} \times \mathbb{X} \times \mathcal{H}_\kappa \rightarrow \mathbb{R}^{m_u}$  is designed satisfying the following assumption.

**Assumption 1.** Suppose there exists a differentiable function  $V(\cdot, \cdot) : \mathbb{R}_{\geq 0} \times \mathbb{R}^{m_y} \rightarrow \mathbb{R}_{\geq 0}$  satisfying

$$\underline{\alpha}(\|\mathbf{z}\|) \leq V(t, \mathbf{z}) \leq \bar{\alpha}(\|\mathbf{z}\|) \quad (16)$$

with class- $\mathcal{K}$  functions  $\underline{\alpha}(\cdot), \bar{\alpha}(\cdot) : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{0,+}$ . The control law  $\pi(\cdot, \cdot, \cdot)$  is designed such that

$$\begin{aligned} & \nabla_{\mathbf{x}} V(t, \mathbf{x})(\boldsymbol{\mu}(\mathbf{x}, \pi(t, \mathbf{x}, \boldsymbol{\mu})) + \mathbf{g}(\mathbf{x}, \pi(t, \mathbf{x}, \boldsymbol{\mu}))) \\ & + \nabla_t V(t, \mathbf{x}) \leq -\alpha(V(t, \mathbf{h}(t, \mathbf{x}))) - \|\nabla_{\mathbf{x}} V(t, \mathbf{x})\|^2 / (4\varepsilon), \end{aligned} \quad (17)$$

with a positive constant  $\varepsilon \in \mathbb{R}_{>0}$ , where the function  $\alpha(\cdot) : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$  belongs to class- $\mathcal{K}$  and

$$\begin{aligned} \nabla_t V(t, \mathbf{x}) &= \frac{\partial V(t, \mathbf{h}(t, \mathbf{x}))}{\partial t} + \frac{\partial V(t, \mathbf{h}(t, \mathbf{x}))}{\partial \mathbf{h}(t, \mathbf{x})} \frac{\partial \mathbf{h}(t, \mathbf{x})}{\partial t}, \\ \nabla_{\mathbf{x}} V(t, \mathbf{x}) &= \frac{\partial V(t, \mathbf{h}(t, \mathbf{x}))}{\partial \mathbf{h}(t, \mathbf{x})} \frac{\partial \mathbf{h}(t, \mathbf{x})}{\partial \mathbf{x}}, \end{aligned} \quad (18)$$

for any  $\boldsymbol{\mu}(\cdot) = [\mu_1(\cdot), \dots, \mu_m(\cdot)]^\top$  satisfying  $\mu_i(\cdot) \in \mathcal{H}_{\kappa,j}$  for all  $j = 1, \dots, m$ .

This assumption indicates the asymptotic output stability of the equivalent system  $\dot{\mathbf{x}} = \boldsymbol{\mu}(\mathbf{x}, \mathbf{u}) + \mathbf{g}(\mathbf{x}, \mathbf{u})$  using the control law  $\pi(t, \mathbf{x}, \boldsymbol{\mu})$ . Compared to the condition in conventional asymptotic stability (Khalil 2002), the additional term  $\|\nabla_{\mathbf{x}} V(t, \mathbf{x})\|^2 / (4\varepsilon)$  introduced in (17) is used following input-to-state control Lyapunov functions to partially compensate for the potential effect from unknown disturbances. Later, an example is provided to demonstrate how to design a model-based controller that satisfies this assumption. Notably, the true dynamical system is unknown  $\mathbf{f}(\cdot)$  instead of  $\boldsymbol{\mu}(\cdot)$ , whose performance is shown as follows.

**Theorem 1.** *If there exists a control law  $\pi(\cdot, \cdot, \cdot)$  satisfying Assumption 1. Choose  $\delta \in (0, 1/m)$ , then the output  $\mathbf{z}(t)$  of the true system (15) is ultimately bounded by*

$$\lim_{t \rightarrow \infty} \|\mathbf{z}(t)\| \leq \underline{\alpha}^{-1}(\alpha^{-1}(\varepsilon \bar{\eta}^2)) \quad (19)$$

with a probability of at least  $1 - m\delta$ . The positive constant  $\bar{\eta} = \max_{t \in \mathbb{R}_{0,+}, \mathbf{x} \in \mathbb{X}} \|\boldsymbol{\eta}(\boldsymbol{\xi}(t, \mathbf{x}, \boldsymbol{\mu}))\|$  denotes an upper bound on the model error between  $\mathbf{f}(\cdot)$  and its GP approximation  $\boldsymbol{\mu}$ , where  $\boldsymbol{\xi}(t, \mathbf{x}, \boldsymbol{\mu}) = [\mathbf{x}^T, \boldsymbol{\pi}^T(t, \mathbf{x}, \boldsymbol{\mu})]^T$  and  $\boldsymbol{\eta}(\cdot) = [\eta_1(\cdot), \dots, \eta_{d_x}(\cdot)]^T$  and  $\eta_i(\cdot)$  obtained using prediction error bound.

This theorem guarantees that the system output ultimately remains within a bounded neighborhood of the desired equilibrium, despite the presence of model uncertainties. To instantiate this theorem and derive a practically applicable controller, we next consider a representative Euler–Lagrange (EL) system and demonstrate the design and performance of a corresponding learning-based policy.

**Euler-Lagrange System Control** Consider an EL system

$$M(\mathbf{q})\ddot{\mathbf{q}} + \mathbf{C}(\mathbf{q}, \dot{\mathbf{q}})\dot{\mathbf{q}} + \mathbf{g}(\mathbf{q}, \dot{\mathbf{q}}) = \mathbf{u} + \mathbf{d}(\mathbf{q}, \dot{\mathbf{q}}), \quad (20)$$

where  $\mathbf{q} \in \mathbb{Q} \subset \mathbb{R}^{m_q}$  and  $\mathbf{u} \in \mathbb{U} \subset \mathbb{R}^{m_u}$  denote the generalized coordinate and control input with  $m_q = m/2$ . Define  $\mathbf{x} = [\mathbf{q}^T, \dot{\mathbf{q}}^T]^T \in \mathbb{X}$ , then the dynamics of (20) is rewritten as

$$\dot{\mathbf{x}} = \underbrace{\begin{bmatrix} \dot{\mathbf{q}} \\ M^{-1}(\mathbf{q})(\mathbf{u} - \mathbf{C}(\mathbf{q}, \dot{\mathbf{q}})\dot{\mathbf{q}} - \mathbf{g}(\mathbf{x})) \end{bmatrix}}_{\mathbf{g}(\mathbf{x}, \mathbf{u})} + \underbrace{\begin{bmatrix} \mathbf{0}_{m_q \times 1} \\ M^{-1}(\mathbf{q})\mathbf{d}(\mathbf{x}) \end{bmatrix}}_{\mathbf{f}(\mathbf{x}, \mathbf{u})}$$

The control task is to track a pre-defined trajectory  $\mathbf{x}_d(\cdot) = [\mathbf{q}_d^T(\cdot), \dot{\mathbf{q}}_d^T(\cdot)]^T : \mathbb{R}_{\geq 0} \rightarrow \mathbb{X}$ , such that the output is defined as tracking error as  $\mathbf{z}(t) = \mathbf{x}(t) - \mathbf{x}_d(t)$ . To achieve this control objective, a learning-based control law inspired by computed torque control is proposed as

$$\boldsymbol{\pi}(t, \mathbf{x}, \boldsymbol{\mu}) = \mathbf{C}(\mathbf{x})\dot{\mathbf{q}} + \mathbf{g}(\mathbf{x}) - \boldsymbol{\mu}(\mathbf{x}) + M(\mathbf{q})\ddot{\mathbf{q}}_d(t) \quad (21)$$

$$+ M(\mathbf{q})(\mathbf{K}_p(\mathbf{q} - \mathbf{q}_d(t)) + \mathbf{K}_d(\dot{\mathbf{q}} - \dot{\mathbf{q}}_d(t))) - \mathbf{B}^T(\mathbf{x})\mathbf{P}\mathbf{z}/(2\varepsilon),$$

where  $\mathbf{B} = [\mathbf{0}_{d_q \times d_q}, M^{-T}(\mathbf{q})]^T$ . The control gains  $\mathbf{K}_p$  and

$\mathbf{K}_d$  are chosen such that  $\mathbf{A} = \begin{bmatrix} \mathbf{0}_{m_q \times m_q} & \mathbf{I}_{m_q} \\ \mathbf{K}_p & \mathbf{K}_d \end{bmatrix}$  is Hurwitz.

The matrix  $\mathbf{P}$  is the solution to the Lyapunov equation associated with the quadratic Lyapunov function  $V(t, \mathbf{z}) = \mathbf{z}^T \mathbf{P} \mathbf{z}$ , where  $\mathbf{P} \succ 0$  satisfies  $\mathbf{A}^T \mathbf{P} + \mathbf{P} \mathbf{A} = -\mathbf{Q}$  with  $\mathbf{Q} \succ 0$  being a positive definite matrix. The existence and uniqueness of the solution  $\mathbf{P}$  is guaranteed due to Hurwitz  $\mathbf{A}$  (Khalil 2002). Then, the control performance in terms of tracking error bound is shown as follows.

**Theorem 2.** Consider the EL system (20) driven by the control law (21) using the proposed Sky-GP satisfying all assumptions in Lemma 1. Choose  $\delta \in (0, 1/m_q) \subset \mathbb{R}$ , then the tracking error  $\mathbf{z}$  is ultimately bounded by

$$\lim_{t \rightarrow \infty} \|\mathbf{z}(t)\| \leq \varepsilon \bar{\lambda}(\mathbf{P})(\lambda(\mathbf{Q})\lambda(\mathbf{P}))^{-1} \bar{\eta}^2 \quad (22)$$

with a probability of at least  $1 - m_q \delta$ .

Theorem 2 shows the tracking error bound obtained by using the controller (21) with the proposed Sky-GP, which is relevant to the worst-case learning performance  $\bar{\eta}$  and the desired convergence strength reflected by  $\mathbf{Q}$ . Specifically, smaller prediction error  $\bar{\eta}$  and larger eigenvalues of  $\mathbf{Q}$  induce smaller tracking error. Note that increasing the control gains with larger eigenvalues of  $\mathbf{Q}$  also leads to high sensitivity of the measurement noise, which may deteriorate the control performance. Therefore, an accurate prediction for a smaller  $\bar{\eta}$  is essential to achieve high tracking precision.

## 5 Experimental Evaluation

We conduct comprehensive experiments to evaluate the proposed SkyGP framework from both regression and control perspectives. In Section 5.1, we assess the regression performance on multiple real-world benchmark datasets. Then, a closed-loop control performance comparison is evaluated against SOTA approaches on a nonlinear dynamical system in Section 5.2. Finally, we present an ablation study analyzing different algorithmic design choices<sup>2</sup>.

### 5.1 Regression Performance Evaluation

We evaluate the proposed SkyGP framework on four real-world benchmark datasets. The SARCOS dataset contains 44,484 samples with 21 input features. The PUMA-DYN32NM (PUMA) dataset consists of 7,168 samples with 32 input features. The KIN40K dataset includes 10,000 samples with 8-dimensional inputs. The ELECTRIC dataset records household electricity consumption, comprising over 2 million samples with 11-dimensional inputs. For simplicity, we use the first 20,000 samples of ELECTRIC dataset. To assess the performance of SkyGP, we compare our two variants against several strong baselines: LoG-GP (Lederer et al. 2021) and Local GPs (Nguyen-tuong, Peters, and Seeger 2008) with each expert holding at most 50 data points, SSGP with 50 inducing points (Bui, Nguyen, and Turner 2017), and ISSGP with 200 random features (Gijbbergs and Metta 2013). All methods use an ARD kernel with hyperparameters pre-trained with the first 1000 samples in each dataset. Each model is evaluated in a sequential setting, where predictions and updates are performed on-the-fly as each new data point arrives. We evaluated the average prediction and update times, as well as the standardized mean squared error (SMSE) and the mean standardized log loss (MSLL) in a sequential interpretation (Lederer et al. 2021). The maximal number of experts  $\bar{N} \in \{1, 2, 4\}$  are used in all benchmark. Each expert with SkyGP holds at most  $\bar{N} = 50$  data points and the maximal search window size  $\bar{W} = 40$  and uses the rBCM aggregation method. Set the time-aware decay weighting factor  $\rho = 0.995$  and  $\bar{\vartheta} = 10^{-3}$ . SkyGP-Fast with a maximum of one expert is denoted by SkyGP-F-1, and SkyGP-Dense with one expert by SkyGP-D-1. Other variants follow the same naming convention.

Table 1 shows the average SMSE and MSLL across three benchmark datasets, demonstrating that the proposed SkyGP variants consistently outperform baseline methods such as LoG-GP and LocalGPs in both predictive accuracy and uncertainty calibration. While ISSGP achieves the best SMSE and MSLL on the PUMA dataset, its prediction and update time is over 30 times slower, as detailed in Table 2, making it less suitable for real-time applications. Notably, SkyGP-D-4 achieves the best overall performance, with the lowest SMSE (0.017) and MSLL (-2.03) on the SARCOS dataset, and competitive results on PUMA and ELECTRIC. Table 2 summarizes the average prediction and update times across three datasets. SkyGP-F-1 achieves the lowest overall latency, with update times consistently at 0.04s and the fastest

<sup>2</sup>For additional results and ablation study analysis, refer to Appendix B in (Yang et al. 2025b).

Model	SARCOS		PUMA		ELECTRIC	
	$\epsilon_{\text{smse}}$	$\epsilon_{\text{msll}}$	$\epsilon_{\text{smse}}$	$\epsilon_{\text{msll}}$	$\epsilon_{\text{smse}}$	$\epsilon_{\text{msll}}$
LoG-GP	0.044	-1.83	0.20	-1.00	0.11	-3.08
SkyGP-F-1	0.037	-1.83	0.26	-0.92	0.08	-3.29
SkyGP-F-4	0.024	-1.91	0.09	-1.39	0.07	-3.36
SkyGP-D-1	0.031	-1.89	0.23	-1.10	0.08	-3.31
SkyGP-D-4	<b>0.017</b>	<b>-2.03</b>	<b>0.08</b>	-1.40	0.07	<b>-3.38</b>
LocalGPs-1	0.031	-1.90	0.18	-1.11	0.08	-3.26
LocalGPs-4	0.071	-1.38	0.18	-0.79	0.14	-2.20
ISSGP	0.023	-1.91	<b>0.08</b>	<b>-1.46</b>	<b>0.06</b>	-2.19
SSGP	0.068	-1.11	0.09	-1.05	-	-

Table 1. Average SMSE and MSLL on 3 datasets.

Model	SARCOS		PUMA		ELECTRIC	
	$t_{\text{pred}}$	$t_{\text{up}}$	$t_{\text{pred}}$	$t_{\text{up}}$	$t_{\text{pred}}$	$t_{\text{up}}$
LoG-GP	0.30	0.22	<b>0.22</b>	0.20	0.18	0.24
SkyGP-F-1	<b>0.16</b>	<b>0.04</b>	0.26	<b>0.04</b>	0.15	<b>0.04</b>
SkyGP-F-4	0.23	<b>0.04</b>	0.35	<b>0.04</b>	0.21	<b>0.04</b>
SkyGP-D-1	0.17	0.09	0.25	0.08	<b>0.14</b>	0.06
SkyGP-D-4	0.24	0.16	0.28	0.16	0.22	0.08
LocalGPs-1	1.17	0.15	1.07	0.08	2.78	0.06
LocalGPs-4	1.25	0.23	1.14	0.08	2.88	0.06
ISSGP	18	7	3	6	4	8
SSGP	5	4	2	6	-	-

Table 2. Average prediction and update time on 3 datasets.

prediction on SARCOS (0.16s), making it well-suited for real-time applications. SkyGP-D-1 has the best prediction time on ELECTRIC (0.14s) while maintaining competitive update performance. In contrast, traditional baselines such as ISSGP and SSGP exhibit significantly higher computational costs, with ISSGP requiring up to 18s for prediction and 7s for update on SARCOS, rendering them impractical for online deployment. Notably, SSGP fails to produce results on the ELECTRIC dataset within 20s, so we omit it.

## 5.2 Control Performance Evaluation

We evaluate the control performance of the system described as  $m\ddot{q} + 9.8 = u + f(\mathbf{x})$  with  $m = 1$ , where

$$f(\mathbf{x}) = 1 + x_1x_2/10 + \cos(x_2)/2 - 10\sin(5x_1) + (1 + \exp(-x_2/10))^{-1}/2. \quad (23)$$

Choose  $x_1 = q, x_2 = \dot{q}$ , then the dynamics is reformulated as  $\dot{x}_1 = x_2, \dot{x}_2 = g(u) + f(\mathbf{x})$  with  $g(u) = u - 9.8 = (u - c\dot{q} - g)/m$  similar as (21) with  $c = 0$  and  $g = 9.8$ .

$$\pi(t, \mathbf{x}, \tilde{\mu}) = 9.8 - \tilde{\mu}(\mathbf{x}) - a_r w_r^2 \sin(w_r t) + k_p(q - q_d(t)) + k_d(\dot{q} - \dot{q}_d(t)) - [0, 1] \mathbf{P}(\mathbf{x} - \mathbf{x}_d)/(2\varepsilon) \quad (24)$$

with control gains  $k_p = 5, k_d = 10$  and  $\varepsilon = 1$ . The matrix  $\mathbf{P}$  is obtained by solving the Lyapunov equation with  $\mathbf{Q} = \mathbf{I}_2$ . The two variants of SkyGP are with  $\tilde{N} = 50, \tilde{W} = 10, \Gamma = 1$  and the maximal data size set to 100.




Figure 2: Learning and tracking performance comparison of the control task from the 100 times Monte Carlo tests.

Other hyperparameters are set the same in Section 5.1. Furthermore, the desired reference is chosen with the form  $q_r(t) = a_r \sin(w_r t)$  with coefficients  $a_r = 1, w_r = 0.1$ .

The maximum norm values of tracking errors  $\mathbf{z} = \mathbf{x} - \mathbf{x}_d$  and prediction errors  $f(\mathbf{x}) - \tilde{\mu}(\mathbf{x})$  are shown in Figure 2. The box plot summarizes the distribution of the performance on prediction error and tracking error over 100 trials with a random initial state  $\mathbf{x}(0)$  uniformly distributed in  $[0, 1]^2$ . The proposed SkyGP variants demonstrate both lower average and median in prediction and tracking errors, indicating more accurate and consistent learning and control behavior compared to all baselines. Particularly, the prediction and tracking errors of SkyGP-Dense are lower than those of SkyGP-Fast with sufficiently fast computation time, which is consistent with the regression benchmarks in Section 5.1.

## 6 Conclusion and Discussion

In this paper, we introduced SkyGP for scalable and adaptive learning. Designed to address the computational and memory bottlenecks of exact GPs in real-time settings, SkyGP maintains a bounded set of experts while preserving strong predictive performance with uncertainty quantification. SkyGP-Fast, targeted for computational efficiency, and SkyGP-Dense, designed for high prediction accuracy. Our extensive empirical evaluation across multiple real-world regression datasets and real-time control tasks demonstrates the superiority of SkyGP over SOTA baselines.

A limitation of the center update rule in Equation (4) is that it may fail to accurately capture special data distributions, such as annular or multimodal patterns. We adopt the method from (Nguyen-tuong, Peters, and Seeger 2008) for its simplicity and computational efficiency. Another limitation lies in the use of the dynamic table list, which may struggle to handle data with sudden distributional shifts, potentially affecting the timely allocation or reuse of experts.

## Acknowledgements

The authors acknowledge the financial support by the Federal Ministry of Education and Research of Germany in the programme of “Souverän. Digital. Vernetzt.” under joint project 6G-life with project identification number: 16KISK002.

## References

- Bui, T. D.; Nguyen, C.; and Turner, R. E. 2017. Streaming Sparse Gaussian Process Approximations. In Guyon, I.; Luxburg, U. V.; Bengio, S.; Wallach, H.; Fergus, R.; Vishwanathan, S.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Cao, Y.; and Fleet, D. J. 2015. Generalized Product of Experts for Automatic and Principled Fusion of Gaussian Process Predictions. arXiv:1410.7827.
- Cohen, S.; Mubvha, R.; Marwala, T.; and Deisenroth, M. 2020. Healing Products of Gaussian Process Experts. In III, H. D.; and Singh, A., eds., *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, 2068–2077. PMLR.
- Dai, X.; Lederer, A.; Yang, Z.; and Hirche, S. 2023. Can Learning Deteriorate Control? Analyzing Computational Delays in Gaussian Process-Based Event-Triggered Online Learning. In Matni, N.; Morari, M.; and Pappas, G. J., eds., *Proceedings of The 5th Annual Learning for Dynamics and Control Conference*, volume 211 of *Proceedings of Machine Learning Research*, 445–457. PMLR.
- Dai, X.; Yang, Z.; Xu, M.; Zhang, S.; Liu, F.; Hattab, G.; and Hirche, S. 2024. Decentralized event-triggered online learning for safe consensus control of multi-agent systems with Gaussian process regression. *European Journal of Control*, 80: 101058. 2024 European Control Conference Special Issue.
- Dai, X.; Yang, Z.; Zhang, S.; Zhai, D.-H.; Xia, Y.; and Hirche, S. 2025. Cooperative Online Learning for Multi-agent System Control via Gaussian Processes With Event-Triggered Mechanism. *IEEE Transactions on Neural Networks and Learning Systems*, 36(7): 13304–13318.
- Deisenroth, M.; and Ng, J. W. 2015. Distributed Gaussian Processes. In Bach, F.; and Blei, D., eds., *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, 1481–1490. Lille, France: PMLR.
- Gijsberts, A.; and Metta, G. 2013. Real-time model learning using Incremental Sparse Spectrum Gaussian Process Regression. *Neural Networks*, 41: 59–69. Special Issue on Autonomous Learning.
- Hashimoto, K.; Saoud, A.; Kishida, M.; Ushio, T.; and Dimarogonas, D. V. 2022. Learning-based Symbolic Abstractions for Nonlinear Control Systems. *Automatica*, 146: 110646.
- Khalil, H. K. 2002. *Nonlinear Systems*. Prentice-Hall.
- Lederer, A.; Begzadić, A.; Hirche, S.; Cortés, J.; and Herbert, S. 2024. Safe barrier-constrained control of uncertain systems via event-triggered learning. *arXiv preprint arXiv:2408.16144*.
- Lederer, A.; Conejo, A. J. O.; Maier, K. A.; Xiao, W.; Umlauf, J.; and Hirche, S. 2021. Gaussian Process-Based Real-Time Learning for Safety Critical Applications. In Meila, M.; and Zhang, T., eds., *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, 6055–6064. PMLR.
- Lederer, A.; Umlauf, J.; and Hirche, S. 2019. Uniform error bounds for Gaussian process regression with application to safe control. *Advances in Neural Information Processing Systems*, 32.
- Lederer, A.; Yang, Z.; Jiao, J.; and Hirche, S. 2023. Cooperative Control of Uncertain Multiagent Systems via Distributed Gaussian Processes. *IEEE Transactions on Automatic Control*, 68(5): 3091–3098.
- Liu, H.; Cai, J.; Wang, Y.; and Ong, Y. S. 2018. Generalized Robust Bayesian Committee Machine for Large-scale Gaussian Process Regression. In Dy, J.; and Krause, A., eds., *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, 3131–3140. PMLR.
- Maddalena, E. T.; Scharnhorst, P.; and Jones, C. N. 2021. Deterministic Error Bounds for Kernel-based Learning Techniques under Bounded Noise. *Automatica*, 134: 109896.
- Naish-guzman, A.; and Holden, S. 2007. The Generalized FITC Approximation. In Platt, J.; Koller, D.; Singer, Y.; and Roweis, S., eds., *Advances in Neural Information Processing Systems*, volume 20. Curran Associates, Inc.
- Nguyen-tuong, D.; Peters, J.; and Seeger, M. 2008. Local Gaussian Process Regression for Real Time Online Model Learning and Control. In Koller, D.; Schuurmans, D.; Bengio, Y.; and Bottou, L., eds., *Advances in Neural Information Processing Systems*, volume 21. Curran Associates, Inc.
- Rasmussen, C. E.; and Williams, C. K. I. 2006. *Gaussian Processes for Machine Learning*. Adaptive Computation and Machine Learning. Cambridge, Mass: MIT Press. ISBN 978-0-262-18253-9.
- Reeb, D.; Doerr, A.; Gerwinn, S.; and Rakitsch, B. 2018. Learning Gaussian Processes by Minimizing PAC-Bayesian Generalization Bounds. In Bengio, S.; Wallach, H.; Larochelle, H.; Grauman, K.; Cesa-Bianchi, N.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc.
- Scharnhorst, P.; Maddalena, E. T.; Jiang, Y.; and Jones, C. N. 2022. Robust Uncertainty Bounds in Reproducing Kernel Hilbert Spaces: A Convex Optimization Approach. *IEEE Transactions on Automatic Control*, 68(5): 2848–2861.
- Schürch, M.; Azzimonti, D.; Benavoli, A.; and Zaffalon, M. 2023. Correlated product of experts for sparse Gaussian process regression. *Machine Learning*, 112(5): 1411–1432.
- Snelson, E.; and Ghahramani, Z. 2005. Sparse Gaussian Processes using Pseudo-inputs. In Weiss, Y.; Schölkopf, B.; and

- Platt, J., eds., *Advances in Neural Information Processing Systems*, volume 18. MIT Press.
- Srinivas, N.; Krause, A.; Kakade, S. M.; and Seeger, M. W. 2012. Information-theoretic Regret Bounds for Gaussian Process Optimization in the Bandit Setting. *IEEE Transactions on Information Theory*, 58(5): 3250–3265.
- Trapp, M.; Peharz, R.; Pernkopf, F.; and Rasmussen, C. E. 2020. Deep Structured Mixtures of Gaussian Processes. In Chiappa, S.; and Calandra, R., eds., *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, volume 108 of *Proceedings of Machine Learning Research*, 2251–2261. PMLR.
- Tresp, V. 2000a. A Bayesian committee machine. *Neural computation*, 12(11): 2719–2741.
- Tresp, V. 2000b. Mixtures of Gaussian Processes. In Leen, T.; Dietterich, T.; and Tresp, V., eds., *Advances in Neural Information Processing Systems*, volume 13. MIT Press.
- Wachi, A.; Sui, Y.; Yue, Y.; and Ono, M. 2018. Safe Exploration and Optimization of Constrained MDPs Using Gaussian Processes. *Proceedings of the AAAI Conference on Artificial Intelligence*, 32(1).
- Yang, Z.; Dai, X.; Dubey, A.; Hirche, S.; and Hattab, G. 2024a. Whom to Trust? Elective Learning for Distributed Gaussian Process Regression. In *Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems*, AAMAS '24, 2020–2028. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems. ISBN 9798400704864.
- Yang, Z.; Dai, X.; Fang, L.; Zhou, J.; and Yan, Z. 2025a. Safe event-triggered control of unmanned surface vehicles with Gaussian processes: Resilience in denial of service attacks and uncertain dynamics. *Engineering Applications of Artificial Intelligence*, 153: 110776.
- Yang, Z.; Dai, X.; and Hirche, S. 2025. Asynchronous Distributed Gaussian Process Regression. *Proceedings of the AAAI Conference on Artificial Intelligence*, 39(21): 22065–22073.
- Yang, Z.; Dong, S.; Lederer, A.; Dai, X.; Chen, S.; Sosnowski, S.; Hattab, G.; and Hirche, S. 2024b. Cooperative Learning with Gaussian Processes for Euler-Lagrange Systems Tracking Control Under Switching Topologies. In *2024 American Control Conference (ACC)*, 560–567.
- Yang, Z.; Sosnowski, S.; Liu, Q.; Jiao, J.; Lederer, A.; and Hirche, S. 2021. Distributed Learning Consensus Control for Unknown Nonlinear Multi-Agent Systems based on Gaussian Processes. In *2021 60th IEEE Conference on Decision and Control (CDC)*, 4406–4411.
- Yang, Z.; Zhang, D.; Dai, X.; Yu, F.; Zhang, C.; Huang, B.; Sadeghian, H.; and Haddadin, S. 2025b. Streaming Generated Gaussian Process Experts for Online Learning and Control: Extended Version. arXiv:2508.03679.
- Yuan, C.; and Neubauer, C. 2008. Variational Mixture of Gaussian Process Experts. In Koller, D.; Schuurmans, D.; Bengio, Y.; and Bottou, L., eds., *Advances in Neural Information Processing Systems*, volume 21. Curran Associates, Inc.
- Yuan, Z.; and Zhu, M. 2024. Lightweight Distributed Gaussian Process Regression for Online Machine Learning. *IEEE Transactions on Automatic Control*, 69(6): 3928–3943.