
Safe Active Learning for Time-Series Modeling with Gaussian Processes

Christoph Zimmer Mona Meister Duy Nguyen-Tuong
Bosch Center for Artificial Intelligence, Renningen, Germany
{christoph.zimmer,mona.meister,duy.nguyen-tuong}@de.bosch.com

Abstract

Learning time-series models is useful for many applications, such as simulation and forecasting. In this study, we consider the problem of actively learning time-series models while taking given safety constraints into account. For time-series modeling we employ a Gaussian process with a nonlinear exogenous input structure. The proposed approach generates data appropriate for time series model learning, i.e. input and output trajectories, by dynamically exploring the input space. The approach parametrizes the input trajectory as consecutive trajectory sections, which are determined stepwise given safety requirements and past observations. We analyze the proposed algorithm and evaluate it empirically on a technical application. The results show the effectiveness of our approach in a realistic technical use case.

1 Introduction

Active model learning deals with the problem of sequential data labeling for learning an unknown function. Data points are sequentially selected for labeling such that the information required for approximating the unknown function is maximized, according to some measures. The overall goal is to create an accurate model without having to supply more data than necessary and, thereby reducing the annotation effort and measurement costs. Active learning has been well studied for classification tasks, e.g. for image labeling [12], but in the field of regression, the active learning approach, related to the optimal experimental design problem [8], is not yet widespread.

For actively learning time-series models representing physical systems, the data has to be generated such that the relevant dynamics can be captured. In practice, the physical system needs to be excited by dynamically moving around in the input space using *input trajectories*, such that the collected data, i.e. input and output trajectories, contain as much information about the dynamics as possible. Commonly used input trajectories include sinusoidal functions, ramps and step functions, white noise, etc. [13, 17]. When employing input excitation on physical systems, however, additional aspects of safety need to be considered. The excitation must not damage the physical system while dynamically exploring the input space, making it crucial to identify safe regions where dynamic excitation can be performed.

In this paper we consider the problem of safe exploration for active learning of time-series models. The goal is to generate input trajectories and output measurements which are informative for learning time-series models. To do so, our input trajectories are parametrized in consecutive sections, e.g. as consecutive piecewise ramps or splines. These consecutive sections of the input trajectory are determined *stepwise* in an explorative approach. Given observations, the next trajectory sections are determined by maximizing an information gain criterion with respect to the model. In this paper, we employ a Gaussian process with a *nonlinear exogenous* structure as the time-series model for which an appropriate exploration criterion is desired. An additional Gaussian process model is simultaneously used for predicting safe input regions, given safety requirements. The sections of

the input trajectory are determined by solving a constraint optimization problem, taking the safety prediction into account. The main contributions of the paper can be summarized as:

- We formulate an active learning setting for learning time-series models with dynamic exploration, in the context of the Gaussian process framework.
- We incorporate the safety aspect into the exploration mechanism and derive a criterion appropriate for the dynamic exploration of the input space with trajectories.
- We provide a theoretical analysis of the algorithm, and empirically evaluate the proposed approach on a realistic technical use case.

The remainder of the paper is organized as follows. In Section 2, we provide an overview on related work. In Section 3, we introduce the algorithm for safe active learning of time-series models. Section 4 provides a theoretical analysis, and in Section 5, we highlight our empirical evaluations in learning time-series model in several settings. The Appendix contains the proofs of the theoretical analysis section and some more experimental investigations.

2 Related Work

Most existing work for safe exploration in unknown environments is in the reinforcement learning setting [16, 10, 9]. For example, the safe exploration in finite MDP relies on the restriction of suitable policies, ensuring ergodicity at a user-defined safety level [16]. In [10], the ergodic assumption for the MDPs is dropped by introducing fatal absorbing states. In [9], the authors consider the use of a multi-armed, risk-aware bandit setting to prevent hazards when exploring different tasks. Strategies for exploring unknown environments have also been reflected in the framework of global optimization with Gaussian processes [1, 23, 11]. For example, [11] propose an efficient submodular exploration criterion for near-optimal sensor placements, i.e. for discrete input spaces. In [1], a framework is presented which yields a compromise between exploration and exploitation through confidence bounds. In [23], the authors show that under reasonable assumptions, strong exploration guarantees can be given for Bayesian optimization with Gaussian processes.

Safe exploration using Gaussian processes (GP) has also been considered in the past, such as for safe active learning [20] and safe Bayesian optimization [2, 24]. In [24], for example, a two-steps process is proposed for a safe exploration and efficient exploitation of identified safe areas. In safe active learning, [20] proposes a method for safe exploration based on the GP variance for stationary, i.e. pointwise, measurements. In contrast to the work by [20], we consider the setting of safe dynamical exploration, i.e. using trajectory-wise measurements. This setting is especially useful when actively learning time-series models.

The problem of active learning for time-series models has not yet been considered extensively in the machine learning literature. Work on related topics is mostly in the field of online design of experiments, e.g. [6]. In [6], the authors employ a parametric model for learning dynamical processes, in which the data for model learning is obtained by exploring using the Fisher information matrix. In contrast to their work, we explore unknown environments by employing a criterion defined for the non-parametric GP model, while also taking into account safety requirements. Furthermore, our proposed exploration scheme is rigorously analyzed providing further algorithmic insights.

3 Safe Active Learning for Time-Series Modeling

Our goal is to approximate an unknown function $f : X \subset \mathbb{R}^d \rightarrow Y \subset \mathbb{R}$. In the case of time-series models, e.g. the well-established *nonlinear exogenous* (NX) model, the input space consists of discretized values of the so-called manipulated variables [3]. Thus, \mathbf{x}_k at time k can be given as

$$\mathbf{x}_k = (\mathbf{u}_k, \mathbf{u}_{k-1}, \dots, \mathbf{u}_{k-d_2+1}),$$

where $(\mathbf{u}_k)_k, \mathbf{u}_k \in \mathbb{R}^{d_2}$, represents the discretized manipulated trajectory. Here, d_1 is the dimension of the system's input space, d_2 the dimension of the NX structure, and $d = d_2 \cdot d_1$. In practice, the elements \mathbf{u}_k are measured from physical systems and need not be equidistant, however, for notational convenience we assume equidistance in this setting. In general, the manipulated trajectories are continuous signals and can be explicitly controlled. In the model learning setting, we observe data in

the form of n consecutive piecewise trajectories $\mathcal{D}_n^f = \{\tau_i, \rho_i\}_{i=1}^n$, where the input trajectory τ_i is a matrix and consists of m input points of dimension d , i.e. $\tau_i = (\mathbf{x}_1^i, \dots, \mathbf{x}_m^i) \in \mathbb{R}^{d \times m}$. The output trajectory ρ_i contains m corresponding output measurements, i.e. $\rho_i = (y_1^i, \dots, y_m^i) \in \mathbb{R}^m$.

The considered problem is to determine the next piecewise trajectory τ_{n+1} as input excitation to the physical system such that the *information gain* of \mathcal{D}_{n+1}^f – with respect to modeling f – is increased. At the same time, τ_{n+1} should be determined subject to given safety constraints. In this section we elaborate on the setting and describe the algorithm. The definition of the considered information gain and corresponding analysis are provided in Section 4.

3.1 Modeling Trajectories with Gaussian Processes

We employ a Gaussian Process (GP) model to approximate the function f (see [19] for more details). A GP is specified by its mean function $\mu(\mathbf{x})$ and covariance function $k(\mathbf{x}_i, \mathbf{x}_j)$, i.e. $f(\mathbf{x}_i) \sim \mathcal{GP}(\mu(\mathbf{x}_i), k(\mathbf{x}_i, \mathbf{x}_j))$. Given noisy observations of input and output trajectories, the joint distribution according to the GP prior is given as

$$p(\mathbf{P}_n | \mathbf{T}_n) = \mathcal{N}(\mathbf{P}_n | \mathbf{0}, \mathbf{K}_n + \sigma^2 \mathbf{I}),$$

where $\mathbf{P}_n \in \mathbb{R}^{n \cdot m}$ is a vector concatenating output trajectories and $\mathbf{T}_n \in \mathbb{R}^{n \cdot m \times d}$ a matrix containing input trajectories. The covariance matrix is represented by $\mathbf{K}_n \in \mathbb{R}^{n \cdot m \times n \cdot m}$. In this paper, we employ the Gaussian kernel as the covariance function, i.e. $k(\mathbf{x}_i, \mathbf{x}_j) = \sigma_f^2 \exp(-\frac{1}{2}(\mathbf{x}_i - \mathbf{x}_j)^T \boldsymbol{\Lambda}_f^2 (\mathbf{x}_i - \mathbf{x}_j))$, which is parametrized by $\boldsymbol{\theta}_f = (\sigma_f^2, \boldsymbol{\Lambda}_f^2)$. Furthermore, we have a zero vector $\mathbf{0} \in \mathbb{R}^{n \cdot m}$ as mean, an $n \cdot m$ -dimensional identity matrix \mathbf{I} , and σ^2 as output noise variance (see [19]). Given the joint distribution, the predictive distribution $p(\boldsymbol{\rho}^* | \boldsymbol{\tau}^*, \mathcal{D}_n^f)$ for a new piecewise trajectory $\boldsymbol{\tau}^*$ can be expressed as

$$p(\boldsymbol{\rho}^* | \boldsymbol{\tau}^*, \mathcal{D}_n^f) = \mathcal{N}(\boldsymbol{\rho}^* | \boldsymbol{\mu}(\boldsymbol{\tau}^*), \boldsymbol{\Sigma}(\boldsymbol{\tau}^*)), \quad (1)$$

with

$$\begin{aligned} \boldsymbol{\mu}(\boldsymbol{\tau}^*) &= \mathbf{k}(\mathbf{T}_n, \boldsymbol{\tau}^*)^T (\mathbf{K}_n + \sigma^2 \mathbf{I})^{-1} \mathbf{P}_n, \\ \boldsymbol{\Sigma}(\boldsymbol{\tau}^*) &= \mathbf{k}^{**}(\boldsymbol{\tau}^*, \boldsymbol{\tau}^*) - \mathbf{k}(\mathbf{T}_n, \boldsymbol{\tau}^*)^T (\mathbf{K}_n + \sigma^2 \mathbf{I})^{-1} \mathbf{k}(\mathbf{T}_n, \boldsymbol{\tau}^*), \end{aligned} \quad (2)$$

where $\mathbf{k}^{**} \in \mathbb{R}^{m \times m}$ is a matrix with $k_{ij}^{**} = k(\mathbf{x}_i, \mathbf{x}_j)$. The matrix $\mathbf{k} \in \mathbb{R}^{n \cdot m \times m}$ contains kernel evaluations relating $\boldsymbol{\tau}^*$ to the previous n input trajectories. As the covariance matrix \mathbf{K}_n is fully occupied, the input points \mathbf{x} are fully correlated within a piecewise trajectory, as well as across different trajectories; this enables the exploitation of high capacity correlations. However, due to the potentially large dimension $n \cdot m$, inverting the matrix $\mathbf{K}_n + \sigma^2 \mathbf{I}$ can be infeasible. GP approximation techniques can be employed, e.g. using sparse inducing inputs or variational approaches [18, 22, 26].

3.2 Modeling the Safety Condition

The safety status of the system is described by an unknown function $g : X \subset \mathbb{R}^d \rightarrow Z \subset \mathbb{R}$, mapping an input point \mathbf{x} to a safety value z , which acts as a safety indicator. The values z are computed using information from the system, and are designed such that all values equal or greater than zero are considered safe for the corresponding input \mathbf{x} . Example 1 shows a construction for computing z values. More examples can be found in the evaluation in Section 5.

Example 1 (A safety indicator for a high-pressure fluid system). *In a high-pressure fluid system, we can measure the pressure ψ for a given input state \mathbf{x} . Additionally, we know the value of the maximal pressure ψ_{max} which can act on the physical system. Given the current pressure ψ , the safety values z can be computed as*

$$z(\psi) = 1 - \exp((\psi - \psi_{max})/\lambda_p), \quad (3)$$

where λ_p describes the decline, when ψ increases towards ψ_{max} .

Note that z is continuous and, intuitively, indicates the distance of a given point \mathbf{x} from the *unknown* safety boundary in the input space. Thus, given the function g – or an estimate of it – we can evaluate the level of safety for a trajectory $\boldsymbol{\tau}$. We consider a trajectory as safe, if the probability that its safety values z are greater than zero is sufficiently large, i.e.

$$\int_{z_1, \dots, z_m \geq 0} p(z_1, \dots, z_m | \boldsymbol{\tau}) dz_1, \dots, z_m > 1 - \alpha,$$

with $\alpha \in (0, 1]$ representing the threshold for considering τ unsafe. Given data $\mathcal{D}_n^g = \{\tau_i, \zeta_i\}_{i=1}^n$, with $\zeta_i = (z_1^i, \dots, z_m^i) \in \mathbb{R}^m$, we employ a GP to approximate the function g . The predictive distribution $p(\zeta^* | \tau^*, \mathcal{D}_n^g)$ given a piecewise trajectory τ^* is then computed as

$$p(\zeta^* | \tau^*, \mathcal{D}_n^g) = \mathcal{N}(\zeta^* | \mu_g(\tau^*), \Sigma_g(\tau^*)), \quad (4)$$

with $\mu_g(\tau^*)$ and $\Sigma_g(\tau^*)$ being the corresponding mean and covariance. The quantities μ_g and Σ_g are computed as shown in Eq. (2), then with $\mathbf{Z}_n \in \mathbb{R}^{n \cdot m}$ as the target vector concatenating all ζ_i . By employing a GP for approximating g , the safety condition $\xi(\tau)$ for a trajectory τ can be computed as

$$\xi(\tau) = \int_{z_1, \dots, z_m \geq 0} \mathcal{N}(\zeta | \mu_g(\tau), \Sigma_g(\tau)) dz_1, \dots, z_m > 1 - \alpha. \quad (5)$$

In general, the computation of $\xi(\tau)$ is analytically intractable, and thus needs to rely on some approximation, such as Monte-Carlo sampling or expectation propagation [15].

3.3 The Algorithm

In the previous sections, we elaborated on the modeling of the predictive distribution and the safety condition for a given piecewise trajectory τ in the input space. For efficiently choosing an optimal τ , the trajectory needs to be appropriately parametrized. The most straightforward possibility is to parametrize in the input space. We illustrate the trajectory parametrization in the following Example 2, using ramp parameterization.

Example 2 (Consecutive ramps as piecewise trajectory). *A ramp can be parametrized with its start and end point. As the start point is the last point of the previous trajectory, the end point η is the only free quantity, and therefore a ramp can be parametrized as*

$$\begin{aligned} \tau(\eta) &= (\mathbf{x}_1(\eta), \dots, \mathbf{x}_m(\eta)) \\ \text{with for } 1 \leq k \leq m : \mathbf{x}_k(\eta) &= \left(\mathbf{u}_0 + \frac{k}{m}(\eta - \mathbf{u}_0), \dots, \mathbf{u}_0 + \frac{k - d_2 + 1}{m}(\eta - \mathbf{u}_0) \right) \end{aligned} \quad (6)$$

where \mathbf{u}_0 is the start point of the ramp. For $k - i \geq 0$, the manipulated input variable is on the currently planned trajectory, and for $k - i < 0$ it can be read from the list of already executed trajectories.

Given a trajectory parametrization with its predictive distribution in Eq. (1) and safety condition in Eq. (5), the next piecewise trajectory $\tau_{n+1}(\eta^*)$ can be obtained by solving the following constrained optimization problem

$$\eta^* = \operatorname{argmax}_{\eta \in \Pi} \mathcal{I}(\Sigma(\eta)) \quad (7)$$

$$\text{s.t. } \xi(\eta) > 1 - \alpha, \quad (8)$$

where η represents our trajectory parametrization, Π is domain of the manipulated variable, and \mathcal{I} an optimality criterion we will discuss later. As shown in Eq. (7), we employ the predictive variance Σ from Eq. (1) for the exploration, which is common in the active learning setting, especially in combination with a GP model [20, 14]. In contrast to previous work, due to the nature of the considered trajectory, we have a *covariance matrix* Σ instead of the variance value usually employed in the active learning and Bayesian optimization setting [23]. The covariance matrix is mapped by an optimality criterion \mathcal{I} to a real number, as indicated by Eq. (7). Various optimality criteria can be used for \mathcal{I} , as discussed in the system identification literature [8]. For example, \mathcal{I} can be the *determinant*, equivalent to maximizing the volume of the predictive confidence ellipsoid of the multi-normal distribution, the *trace*, equivalent to maximizing the average predictive variance, or the *maximal eigenvalue*, equivalent to maximizing the largest axis of the predictive confidence ellipsoid [8].

The constraint in Eq. (8) represents a probabilistic safety criterion, motivated by our probabilistic modeling approach for the safety. The probabilistic approach flexibly allows us to control the trade-off between exploration speed and safety consideration. For example, a 100% safe exploration would keep the algorithm from leaving the initial safe area and, hence, would not lead to an exploration of new areas. On the other hand, a 0% safe exploration will explore without any safe considerations which will result in many safety violations. This trade-off provides the users an additional degree of freedom, depending on how much they “trust” the behavior of their physical systems.

Algorithm 1 Safe Active Learning for Time-Series Modeling

- 1: Input: Safety threshold $0 \leq \alpha \leq 1$
 - 2: Initialization: Collect n_0 safe trajectories, i.e. $\mathcal{D}_0^{f,g} = \{\tau_i, \rho_i, \zeta_i\}_{i=1}^{n_0}$ with $n = n_0$.
 - 3: **for** $k = 1$ **to** N **do**
 - 4: Update regression model approximating f using $\mathcal{D}_{k-1}^f = \{\tau_i, \rho_i\}_{i=1}^n$, according to Eq. (1)
 - 5: Update safety model approximating g using $\mathcal{D}_{k-1}^g = \{\tau_i, \zeta_i\}_{i=1}^n$, according to Eq. (4)
 - 6: Determine new piecewise trajectory τ_{n+1} , by optimizing η according to Eq. (7 and 8)
 - 7: Execute τ_{n+1} on the physical system, while measuring ρ_{n+1} and ζ_{n+1}
 - 8: Include new trajectories into \mathcal{D}_{k-1}^f and \mathcal{D}_{k-1}^g with $n = n + 1$.
 - 9: **end for**
 - 10: Update and return regression model and safety model
-

Algorithm 1 summarizes the basic steps of the proposed algorithm, which needs to be initialized by n_0 safe trajectories. In practice, the initial trajectories are located in a small, safe region chosen beforehand using prior knowledge. The incremental updates of the GP models for new data, i.e. steps 4 and 5 in Algorithm 1, can be efficiently performed, e.g. through rank-one updates [21]. The optimization problem in Eq. (7) can be solved using gradient-based optimization approaches, e.g. [4, 5]. In this paper, we employ the NX-structure in combination with the GP model for time-series modeling. However, this approach can also be extended to the general nonlinear auto-regressive exogenous case [3], i.e. a GP with NARX input structure $\mathbf{x}_k = (y_k, y_{k-1}, \dots, y_{k-q}, u_k, u_{k-1}, \dots, u_{k-d})$. In this case, for optimization and planning of the next piecewise trajectories, one can use the predictive mean of $p(\rho|\tau, \mathcal{D}_n^f)$ as surrogate for y_k . Note that the input excitation is still performed through the manipulated variable u_k in the case of NARX.

4 Theoretical Results

In this section, we provide some results on the theoretical analysis of the proposed approach. First, we investigate the safety aspect of the algorithm. In Section 4.2, we provide a bound on the decay rate of the predictive variances for the case when the criterion \mathcal{I} is a determinant, i.e. $\mathcal{I}(\Sigma(\eta)) = \det(\Sigma(\eta))$. The proofs can be found in the Appendix.

4.1 Safe Exploration

To satisfy the safety requirements, it is necessary to bound the probability of failures during exploration. Theorem 1 provides an upper bound on the probabilities for unsafe trajectories.

Theorem 1. *Let us assume that we have recorded n_0 initial safe trajectories, and that their observations are enough to model g well, in the sense that our GP quantifies the uncertainty of predictions for g correctly, i.e. $P(\mu_g - \nu\sigma_g \leq z \leq \mu_g + \nu\sigma_g) = \text{Erf}(\nu/\sqrt{2})$ for all $\nu \geq 0$. Let $\delta \in [0, 1]$ be the desired failure probability when determining the next N consecutive piecewise trajectories. Set $\alpha = \delta/N$ and let this α be the probability bound for a trajectory being unsafe (as in Eq. (5)). Then, the iterative exploration for the next N trajectories is unsafe with probability at most δ , i.e.*

$$P\left(\bigcup_{i=n_0+1}^{n_0+N} \{g(\mathbf{x}_j^i) < 0 \text{ for some } 1 \leq j \leq m | \xi(\tau_i) > 1 - \alpha\}\right) \leq \delta.$$

Theorem 1 supplies us with a useful rule of thumb to select α for sequentially determining the next N trajectories.

4.2 Decay of Predictive Variance

The remainder of the analysis is to show that the proposed exploration scheme makes the predictive uncertainty Σ decrease as n increases. In this paper we use the *determinant* of Σ as an exploration criterion, which has been shown to have a close relationship to the information gain [14, 23], defined as the mutual information I .

First, we point out that this relationship still holds true in case of trajectories as observations. Subsequently, we introduce the maximum information gain as an upper bound, which can further be

used to show the decrease of the predictive uncertainty. Lemma 1 clarifies the relationship between determinant and mutual information. Let us denote the predictive variance after recording $i - 1$ trajectories as $\Sigma_{i-1}(\tau_i)$, $\Sigma_0(\tau_1) = \mathbf{k}^{**}(\tau_1, \tau_1)$, and set $\tilde{\rho}_i = (f(\mathbf{x}_1^i), \dots, f(\mathbf{x}_m^i))$.

Lemma 1. *The mutual information $I(\{\rho_i\}_{i=1}^n; \{\tilde{\rho}_i\}_{i=1}^n)$ can be related to the predictive co-variances $\Sigma_{i-1}(\tau_i)$ as follows*

$$I(\{\rho_i\}_{i=1}^n; \{\tilde{\rho}_i\}_{i=1}^n) = 1/2 \sum_{i=1}^n \log |\mathbf{I}_m + \sigma^{-2} \Sigma_{i-1}(\tau_i)|$$

Next, we introduce the maximum information gain after observing n trajectories as $\gamma_n := \max_{\{\tau_i\}_{i=1}^n \subset X^m} I(\{\rho_i\}_{i=1}^n; \{\tilde{\rho}_i\}_{i=1}^n)$, (see Srinivas et al. [23] for more details). The maximum information gain is the information which could be gathered when exploring the system in a non-iterative way, by optimally designing all trajectories simultaneously (which is in practice hard as it would require a solution of a high dimensional optimization problem, and would not allow us to incorporate safety information from observations during the experiment). According to Srinivas et al. ([23], Theorem 5) the maximum information gain satisfies $\gamma_n = O(\log(n)^{d+1})$, i.e. the maximum information grows slower than the number of additional trajectories. This will be crucial later on, but first we investigate the relation between the determinant of the covariance and γ_n . Using Lemma 1 the determinant of the covariance can be bounded, as given in Lemma 2.

Lemma 2. *After observing n trajectories $\{\tau_i\}_{i=1}^n$ (according to Eq. (7)), the determinant of the covariance is upper bounded by*

$$\frac{1}{n} \sum_{i=1}^n |\Sigma_{i-1}(\tau_i)| \leq C \frac{\gamma_n}{n},$$

where Σ_{i-1} is the predictive variance computed using the previous $i - 1$ trajectories, γ_n is the maximum information gain, and $C = 2\sigma_f^{2m} / \log(1 + \sigma^{-2m} \sigma_f^{2m})$ is a constant.

The first step in proving Lemma 2 is to upper bound the predictive variance using the mutual information via Lemma 1. Subsequently, the mutual information is upper bounded by γ_n . Using Lemma 2 and Theorem 5 in [23], we can provide a decay rate on the average determinant of predictive variances.

Theorem 2. *Let $\{\tilde{\tau}_i\}_{i=1}^n$ be n arbitrary trajectories within a compact and convex domain X , and k be a kernel function such that $k(\cdot, \cdot) \leq 1$. If Σ_{i-1} come from our exploration scheme Eq. (7) (i.e. without safety considerations), then we have*

$$\frac{1}{n} \sum_{i=1}^n |\Sigma_{i-1}(\tilde{\tau}_i)| = \mathcal{O}\left(\frac{\log(n)^{d+1}}{n}\right).$$

We sketch the proof here: as our algorithm (without safety considerations) always chooses the trajectory with the highest determinant (D criterion), the average determinant of an actively learned scheme is always higher than or equal to the average determinant of an arbitrary scheme. Therefore, $\frac{1}{n} \sum_{i=1}^n |\Sigma_{i-1}(\tilde{\tau}_i)| \leq \frac{1}{n} \sum_{i=1}^n |\Sigma_{i-1}(\tau_i)|$, which is $\mathcal{O}(\log(n)^{d+1}/n)$ when employing Lemma 2 and the Theorem 5 of [23].

By Theorem 2, for any sequence of trajectories, the average of the determinants of their predictive covariances tends to zero. As the determinant corresponds to the volume of the confidence ellipsoid, we can conclude that the average volume of confidence ellipsoids tends to zero as well, indicating that on average, our predictions become precise. However, as the safety constraint in Eq. (8) changes at every iteration, we extend the statement of Theorem 2:

Theorem 3. *Let us assume that there exists a compact and convex domain X , and a kernel function k such that $k(\cdot, \cdot) \leq 1$, that covers the whole area which is explorable (independent of whether it is safe or not). Then, the statement of Theorem 2 still holds for our Algorithm 1 with iteration-dependent safe areas S_i .*

Theorem 3 guarantees the decay of averaged determinants of covariances during safe exploration.




Figure 1: The columns show the progress of the approximation of f (inlay) and the identified safety region (main figure) at different iterations. Each iteration corresponds to a consecutive planning of a new piecewise trajectory (here: 2D ramp). As shown by the results, the current estimation of the safe region (green area) gradually covers the actual safe area (red line), and the approximation error gradually decreases (as shown in the subfigures). An illustrative video showing all iterations can be found in the Appendix.

5 Evaluations

In section 5.1 we illustrate the proposed approach using synthetic models, comparing our safe active learning approach (SAL-NX) with random selection using safety constraints. Subsequently, we employ the approach to learn a dynamics model of a physical, high-pressure fluid system in Section 5.2. For simplicity we employ ramps for the piecewise trajectory parametrization, but other curve parameterizations could also be used instead, e.g. spline parameterization. The form of the input trajectory has an impact on the excitation of the system, as comprehensively studied in the field of system identification [17].

5.1 Simulated Experiments

Experiment 1 In this experiment, a toy example is employed to illustrate the concept of input space exploration with piecewise trajectories and safe region detection. A function $f: \mathbb{R}^2 \rightarrow \mathbb{R}$, $f(\mathbf{x}) = (x^{(1)} - 2)^2 + (x^{(1)} - 2)(x^{(2)} - 2) + (x^{(2)} - 2)^2$ with $\mathbf{x} = (x^{(1)}, x^{(2)})$ is used as the ground-truth. An observation is given by $y = f(\mathbf{x}) + \epsilon$ with $\epsilon \sim \mathcal{N}(0, 1)$. The safe region is characterized by $g: \mathbb{R}^2 \rightarrow \mathbb{R}$ with $g(\mathbf{x}) = (x^{(1)} - 5)^2 + (x^{(1)} - 5)(x^{(2)} - 5) + (x^{(2)} - 5)^2$. The safety indicator z is given as $z = -0.005 \cdot g(\mathbf{x}) + 1 + \varsigma$ with $\varsigma \sim \mathcal{N}(0, 1)$. It is considered to be safe for $z > 0$, otherwise unsafe.

We proceed as shown in Algorithm 1, where the piecewise trajectories are parametrized as 2D-ramps with 5 discretization points (i.e. $m = 5$, see Example 2). We start with 10 initial safe trajectories and consecutively determine new piecewise trajectories in the input space X , while also collecting outputs y and computing safety indicator values z . As the exploration progresses, the approximation of f and g becomes more and more accurate, as shown in Fig. 1. The current estimation of the safe region (green area) gradually covers the actual safe area, and the approximation error gradually decreases (as shown in the subfigures). An illustrative video showing all iterations can be found in the Appendix.

Experiment 2 In this experiment, we learn a time-series model given as a GP with NX-structure. We have two manipulated variables $u_k^{(1)}$ and $u_k^{(2)}$ at time k . The NX-structure is determined to be $\mathbf{x}_k = (u_k^{(1)}, u_k^{(2)}, u_{k-1}^{(1)}, u_{k-1}^{(2)})$, an input space with $d = 4$. The ground-truth models of f and g are provided in the Appendix. The piecewise trajectory is again parametrized as 4D-ramps with $m = 5$. We initialize the models using 10 collected piecewise ramps in a safe area, and start exploring in the input space. For a fair comparison, we benchmark the proposed algorithm against a *random selection with safe constraints* of next piecewise trajectories. Instead of optimizing the ramp parameter η as shown in Eq. (7) and (8), we randomly select η and pick the first one which fulfills the safety constraint $\xi(\eta) > 1 - \alpha$. Fig. 2 shows the results of the comparison of the proposed approach (SAL-NX) with random selection.

The results in Fig. 2 show that SAL-NX continuously improves the model approximation (shown as RMSE) and provides fast coverage of safe regions. The models for f and g are updated after every iteration by including new sample points. The hyperparameters can be estimated beforehand




Figure 2: The **first two** pictures from the left show the comparison of the SAL-NX (red line) with random selection (blue line). SAL-NX yields faster convergence in model approximation (left picture) and coverage of safe regions (right picture), while having less variance and outliers (indicated as small circles). The **last two** pictures show the impact of the safety threshold α . The left picture shows the RMSE of SAL-NX for 4 different values of α . The right picture shows the model approximation error as RMSE (red line) and percentage of *unsafe* trajectories (blue line) as a function of α . All pictures show boxplots over 5 repetitions.

or updated after every iteration. For the required number of initial trajectories, we refer to the lower bound as given in [20]. For computing the safety condition $\xi(\tau)$ from Eq. (5), we employ Monte-Carlo sampling. Our experiments are performed on a desktop computer. The algorithm is sufficiently fast for real-time applications.

In this experiment, we also compare our exploration approach with the one proposed in [6], however, *without* safety requirements in order to cope with the setting from [6]. We adapt their criterion based on the Fisher information for our GP model by employing the GP mean function. Additionally, we also compare the decrease in RMSE of the Fisher information based criterion to the decrease in RMSE of our algorithm. The results can be found in the Appendix 7.4 and show a competitive performance of our approach.

5.2 Learning a Surrogate Model of the High-Pressure Fluid System

The Use Case As a realistic technical use case, we employ the approach to actively learn a surrogate model of a high-pressure fluid injection system, as shown in Figure 3. Such systems are widely used in industry, e.g. in the automotive domain for injection of fuel into the combustion engine [25].




Figure 3: High-pressure fluid injection system with controllable inputs v_k, n_k and measured output ψ_k (picture taken from [25]).

The physical injection system is controlled by an actuation signal v_k and the speed of an external engine n_k , for every time step k . The goal is to obtain a surrogate model predicting the rail pressure ψ_k , which determines the amount of fluid coming out of the outlets. Due to the nature of the fluid and the mechanical components, the dynamics of the whole system are nonlinear, and thus model learning is an appropriate alternative compared to analytical models. However, generating the data for learning a time series surrogate model by varying the actuation signal and engine speed is not simple, as an inappropriate combination of them would result in hazardously high rail pressures, damaging the physical system.

Learning Time-Series Surrogate Models Due to the safety requirements and the fact that the safety boundary is not known beforehand, our safe active learning approach is very appropriate for approximating the dynamics model. The employed NX-structure is chosen to be $\mathbf{x}_k = (n_k, n_{k-1}, n_{k-2}, n_{k-3}, v_k, v_{k-1}, v_{k-3})$. We again parameterize with piecewise ramps in this 7D input space. The safety indicator value z is computed as shown in Example 1, with $\psi_{max} = 18$ MPa being the maximally allowed rail pressure. It should be noted that here z is computed as a function of the target output ψ , in contrast to the experiments in Section 5.1, where z is a function of the input. We initialize the model with 25 trajectories sampled around a safe point chosen by a domain expert. Subsequently, we start exploring the input space dynamically, considering both the safety constraint




Figure 4: The **first two** pictures from the left show the comparison of the SAL-NX (red line) with random selection with safe constraints (blue line), with respect to model approximation and coverage of safe regions. Here, $\alpha = 0.5$ and 250 trajectories are planned. The **last two** pictures show the impact of the safety threshold α on the approximation error, and failures during exploration. The results are displayed as a boxplot over 5 repetitions.

and the model information gain, while measuring the actuation and speed signals as input and the rail pressure as output.

Figure 4 shows the results after exploring the input space with 250 consecutive ramp trajectories, each consisting of $m = 5$ discretization points. We update the hyperparameters after every iteration. We compare our SAL-NX approach with the random selection, as described in the previous experiment. The figure also shows the impact of varying the threshold value α on both the model approximation error and the percentage of selected unsafe trajectories. In practice, the execution of the trajectories on the physical system is interrupted, when the system notices a violation of the maximal pressure ψ_{max} . The selected piecewise trajectory is then indicated as unsafe. For the evaluation of the coverage (second picture from the left), the “ground-truth” safe region is estimated beforehand with an extensive procedure.

6 Conclusions

In this paper we present an approach for active learning of a time-series model, given as a GP model with NX-structure. In this setting, the exploration is performed while taking safety requirements into account. For the successful application of the algorithm, it is crucial that the system can be actively controlled by a set of inputs and a safety signal can be observed during the system’s operation. The proposed approach is evaluated on toy examples, as well as on a realistic technical use case. The results show that this approach is appropriate for real-world applications, especially, in the industrial setting, where safety is a key requirement during operation.

References

- [1] P. Auer. Using Confidence Bounds for Exploitation-Exploration Trade-Offs. *Journal of Machine Learning Research*, 2002.
- [2] F. Berkenkamp, A. Krause, and A. P. Schoellig. Bayesian Optimization with Safety Constraints: Safe and Automatic Parameter Tuning in Robotics. Technical report, arXiv, February 2016.
- [3] S. Billings. *Nonlinear System Identification: Narmax Methods in the Time, Frequency, and Spatio-Temporal Domains*. John Wiley & Sons, 2013.
- [4] R. H. Byrd, J. C. Gilbert, and J. Nocedal. A trust region method based on interior point techniques for nonlinear programming. *Mathematical Programming*, 89(1):149–185, 2000.
- [5] T. F. Coleman and Y. Li. On the convergence of reflective newton methods for large-scale nonlinear minimization subject to bounds. Technical report, Ithaca, NY, USA, 1992.
- [6] M. Deflorian, F. Kloepper, and J. Rueckert. Online dynamic black box modelling and adaptive experiment design in combustion engine calibration. In *6th IFAC Symposium Advances in Automotive Control*. Elsevier, 2010.

- [7] M. Deflorian, F. Kloepper, and J. Rueckert. Design of experiments for nonlinear dynamic system identification. In *Proceedings of the 18th World Congress, The International Federation of Automatic Control*. 2011 IFAC, 2011.
- [8] V. Fedorov and P. Hackl. *Model-Oriented Design of Experiments*. Lecture Notes in Statistics. Springer New York, 2012.
- [9] N. Galichet, M. Sebag, and O. Teytaud. Exploration vs Exploitation vs Safety: Risk-Aware Multi-Armed Bandits. In *Proceedings of the 5th Asian Conference on Machine Learning*, 2013.
- [10] P. Geibel. Reinforcement Learning with Bounded Risk. In C. E. Brodley and A. P. Danyluk, editors, *Proceedings of the 18th International Conference on Machine Learning*, pages 162–169, 2001.
- [11] C. Guestrin, A. Krause, and A. Singh. Near-Optimal Sensor Placements in Gaussian Processes. In *Proceedings of the 22nd International Conference on Machine Learning*, 2005.
- [12] A. J. Joshi, F. Porikli, and N. Papanikolopoulos. Multiclass active learning for image classification. In *IEEE Conf. on Computer Vision and Pattern Recognition*, pages 2372–2379, 2009.
- [13] L. Ljung and T. Söderström. *Theory and Practice of Recursive Identification*. MIT Press series in signal processing, optimization, and control. MIT Press, 1985.
- [14] D. J. C. MacKay. Information-Based Objective Functions for Active Data Selection. *Neural Computation*, 4(4):590–604, 1992.
- [15] T. P. Minka. Expectation Propagation for Approximate Bayesian Inference. In *Uncertainty in Artificial Intelligence*. Morgan Kaufmann, 2001.
- [16] T. M. Moldovan and P. Abbeel. Safe Exploration in Markov Decision Processes. In *Proceedings of the 29th International Conference on Machine Learning*, 2012.
- [17] R. Pintelon and J. Schoukens. *System Identification: A Frequency Domain Approach*. Wiley, 2012.
- [18] J. Quiñero-Candela and C. E. Rasmussen. A Unifying View of Sparse Approximate Gaussian Process Regression. In *Journal of Machine Learning Research*, 2005.
- [19] C. E. Rasmussen and C. K. I. Williams. *Gaussian Processes for Machine Learning*. The MIT Press, 2006.
- [20] J. Schreiter, D. Nguyen-Tuong, M. Eberts, B. Bischoff, H. Markert, and M. Toussaint. Safe Exploration for Active Learning with Gaussian Processes. In *ECML/PKDD*, volume 9286, 2015.
- [21] M. Seeger. Low rank updates for the Cholesky decomposition, 2007.
- [22] E. L. Snelson and Z. Ghahramani. Sparse Gaussian Processes using Pseudo-inputs. In *Advances in Neural Information Processing Systems*, 2006.
- [23] N. Srinivas, A. Krause, S. M. Kakade, and M. W. Seeger. Information-Theoretic Regret Bounds for Gaussian Process Optimization in the Bandit Setting. *Transactions on Information Theory*, 2012.
- [24] Y. Sui, V. Zhuang, J. Burdick, and Y. Yue. Stagewise safe bayesian optimization with gaussian processes. In *35th International Conference on Machine Learning*, 2018.
- [25] N. Tietze, U. Konigorski, C. Fleck, and D. Nguyen-Tuong. Model-based calibration of engine controller using automated transient design of experiment. In *14. Internationales Stuttgarter Symposium*. Springer Fachmedien Wiesbaden, 2014.
- [26] M. K. Titsias. Variational Learning of Inducing Variables in Sparse Gaussian Processes. In *Proceedings of the Twelfth International Conference on Artificial Intelligence and Statistics*, 2009.