To appear in IEEE Conference on Decision and Control (CDC), Dec. 2020.

Safety Verification of Unknown Dynamical Systems via Gaussian Process Regression

John Jackson[1], Luca Laurenti[2], Eric Frew[1], and Morteza Lahijanian[1]

***Abstract*— The deployment of autonomous systems that operate in unstructured environments necessitates algorithms to verify their safety. This can be challenging due to, e.g., black-box components in the control software, or undermodelled dynamics that prevent model-based verification. We present a novel verification framework for an unknown dynamical system from a given set of noisy observations of the dynamics. Using Gaussian processes trained on this data set, the framework abstracts the system as an uncertain Markov process with discrete states defined over the safe set. The transition bounds of the uncertain process are derived from the probabilistic error bounds between the regression and underlying system. An existing approach for verifying safety properties over uncertain Markov processes then generates safety guarantees. We demonstrate the versatility of the framework on several examples, including switched and nonlinear systems.**

I. INTRODUCTION

The ability to provide formal guarantees is essential for *safety-critical* systems. Without assurances, innovations such as self-driving cars, medical robotics, and autonomous aerial vehicles will remain bounded to narrow domains. To address this need, formal verification offers powerful frameworks with rigorous analysis techniques [1], [2]. They provide formal guarantees with respect to the system model. In many applications, however, an accurate model of an autonomous system is either unavailable due to, e.g., the use of a black-box controller, or if available, it is not in a closed form that can be used for formal verification. This work aims to develop a verification method that can provide safety guarantees for systems with unknown dynamics.

Formal verification of control systems has been widely studied, e.g., [3]–[9]. These methods are typically based on model checking algorithms [1], [2] which take a simple discrete, finite model and return a *yes* or *no* as to whether the model satisfies a given specification. To bridge the gap between continuous and discrete domains, those works construct an *abstraction*, a finite representation of the control system with a simulation relation [10]. This abstraction is in the form of a finite graph if the underlying system is deterministic or a finite Markov process if the underlying system is stochastic. Even though they admit strong guarantees, these methods are model-based and require full knowledge of the system model. Hence, they cannot be employed for analysis of systems with unknown dynamics.

Recent work has focused on data-driven analysis of dynamical systems, e.g., [11]–[14], which assume partial knowledge about the system and provide some performance assurances. The work in [12] uses techniques based on Bayesian inference to compute the confidence over a property of interest for partly unknown linear systems. Work [13] introduces an algorithm based on chance-constrained optimization to provide probabilistic stability guarantees for an unknown switched linear system from a finite number of observations of trajectories. Despite their strengths, those data-driven methods assume the unknown model is linear. Work [14] relaxes this assumption and considers safety assessment of a dynamical system whose model is fully unknown. Their method is based on approximation of the dynamics using a piecewise-polynomial function and safety assessment through barrier certificates. This safety analysis is sound with respect to the polynomial function but cannot be extended to the underlying system in a straightforward manner.

A powerful approach to approximate an unknown function is *Gaussian process* (GP) regression [15]. GP regression is a Bayesian machine-learning framework, which has been receiving special attention in safety-critical applications due to its ability to capture the uncertainty in the learning process [16], [17]. Recent works [18]–[21] successfully derive theoretical bounds on the distance between the regressed GP and the underlying (unknown) system. These results have led to the increased use of GPs in safe learning frameworks, e.g., [22]–[25]. In [22], [23], the proposed algorithms learn the unknown dynamics as a GP model, which is then used within a reinforcement learning algorithm to learn a reachability policy under safety constraints. Similarly, [24] introduces a method of learning a policy safely based on GP modeling with stability guarantees. Nevertheless, it is unclear whether those algorithms designed for learning policies can be employed for formal verification purposes.

In this work, we focus on the safety verification of control systems with unknown dynamics via GP regression. We introduce an algorithm that, given a set of noisy data, generates formal probabilistic guarantees for the unknown system to remain in a given safe set for every initial state. These guarantees are with respect to an arbitrary control policy. The algorithm uses a discretization of the safe set and GP regression to construct a finite abstraction with probabilistic bounds. This abstraction is in the form of an uncertain Markov model that captures all possible behaviors of the unknown

This work was supported in part by the University of Colorado Boulder Autonomous Systems Interdisciplinary Research Theme, and the NSF Center for Unmanned Aircraft Systems, and the National Science Foundation under award IIP - 1650468.

[1]J. Jackson, E. Frew, and M. Lahijanian are with the Department of Smead Aerospace Engineering Sciences, University of Colorado Boulder, CO, USA. {john.m.jackson, eric.frew, morteza.lahijanian}@colorado.edu

[2]L. Laurenti is with the Department of Computer Science, University of Oxford, UK. luca.laurenti@cs.ox.ac.uk

system through a derivation for the error bounds between the regression and underlying system. Then, the algorithm determines the safety probability bounds for the unknown system by performing safety verification on the abstraction.

The main contribution of this work is a framework for formal verification of unknown dynamical systems. This is the *first* abstraction-based verification technique that does not assume known dynamics to the best of our knowledge. This work lays the theoretical foundation for formal reasoning about unknown systems against complex specifications given, e.g., as *temporal logic* formulas [2]. Another contribution of the paper is a derivation of probability bounds on the transition from a point to a region for the unknown dynamics. These bounds are general and hence can be applied to systems with various levels of knowledge about their dynamics.

II. PROBLEM FORMULATION

Consider a controlled dynamical system with noisy observation (measurement) in the form of

$$\begin{aligned} \mathbf{x}(k+1) &= f(\mathbf{x}(k), \mathbf{u}(k)), \\ \mathbf{y}(k) &= \mathbf{x}(k) + \mathbf{v}(k, \mathbf{u}(k-1)), \end{aligned} \tag{1}$$

where

$$\mathbf{x}(k) \in \mathbb{R}^n, \mathbf{u}(k) \in \mathcal{U}, \mathbf{y}(k) \in \mathbb{R}^n, \mathbf{v}(k, \mathbf{u}(k-1)) \sim \mathcal{D}_{\mathbf{u}(k-1)},$$

$f : \mathbb{R}^n \times U \to \mathbb{R}^n$ is a possibly non-linear and unknown function that represents the dynamics of the system, $\mathcal{U} = \{a_1, \ldots, a_{|\mathcal{U}|}\}$ is a finite set of actions or control laws, and for each $a \in \mathcal{U}$, $\mathbf{v}(k, a)$ is a noise term sampled from distribution $\mathcal{D}_a$. We assume the noise $\mathbf{v}$ is an arbitrary zero-mean martingale difference sequence, i.e., for each $k > 0$ and $k' < k$

$$\mathbb{E}\big[\mathbf{v}(k, \mathbf{u}(k)) \mid \mathbf{v}(k', \mathbf{u}(k'))\big] = 0.$$

We further assume that $\|\mathbf{v}\| < \sigma$ almost surely for some $\sigma > 0$ at each step k and that the noise on the various components of $\mathbf{x}$ is independent, i.e., component $\mathbf{v}_i$ is independent of $\mathbf{v}_j$ for $i, j \in \{1, ..., n\}$.

Taking f as completely unknown may lead to an ill-posed problem. We employ the following standard assumption [18], which guarantees f is a well-behaved function that can be approximated using GP regression.

Assumption 1. *For a compact set $\mathcal{K} \subset \mathbb{R}^n$, let $\kappa : \mathbb{R}^n \times \mathbb{R}^n \to \mathbb{R}_{>0}$ be a given kernel and $\mathcal{H}_\kappa(\mathcal{K})$ the reproducing kernel Hilbert space (RKHS) of functions over $\mathcal{K}$ corresponding to κ with norm $\|\cdot\|_\kappa$ [18]. Then, for each $a \in \mathcal{U}$ and $i \in \{1, ..., n\}$, $f_i(\cdot, a) \in \mathcal{H}_\kappa(\mathcal{K})$ and for a constant $B_i > 0$, $\|f_i(\cdot, a)\|_\kappa \leq B_i$, where f_i is the i-th component of f.*

Assumption 1 limits the class of functions that can be considered in Process (1) [18]. The class of functions strictly depends on the kernel under consideration. A universal kernel, such as the widely-used squared exponential kernel, has the property that $\mathcal{H}_\kappa(\mathcal{K})$ is a set which is dense in $\mathcal{C}(\mathcal{K})$ – the set of continuous functions over $\mathcal{K}$. That is, every continuous function over $\mathcal{K}$ can be approximated arbitrarily well by members of $\mathcal{H}_\kappa(\mathcal{K})$ [26].

Let $\omega_{\mathbf{x}}(k) = x_0 \xrightarrow{u_0} x_1 \xrightarrow{u_1} \ldots \xrightarrow{u_{k-1}} x_k$, where $x_1, \ldots, x_k \in \mathbb{R}^n$, be a trajectory of Process (1) up to time k with the observation (measurement) trajectory $\omega_{\mathbf{y}}(k) = y_0 y_1 \ldots y_k$. Then, a *control strategy* $\pi_{\mathbf{x}}$ is a measurable function that selects an action (control law) at time k for the system given the observation trajectory up to that time, i.e., $\pi_{\mathbf{x}}(\omega_{\mathbf{y}}(k)) \in \mathcal{U}$. Note that $\mathbf{v}$ is a stochastic process. As a consequence, $\pi_{\mathbf{x}}$ and $\mathbf{x}$ are stochastic processes.

A. Problem

The focus of this paper is the safety analysis of Process (1) from a set of samples, each in the form of (x, u, y), where y is an observation of Process (1) with state x and input u. Note that this analysis needs to be probabilistic due to the reasons stated above and the partial knowledge (finite noisy samples) of Process (1). The focus is specifically on the verification problem, where the goal is to check if a given safety probability threshold is guaranteed. Therefore, the problem is centered on computing the probability range that $\mathbf{x}(k)$ remains safe for a given (possibly unbounded) time horizon under all possible strategies. This problem is formally defined below.

Problem 1. *Let $\mathrm{D} = \{(x_i, u_i, y_i) \mid i \in \{1, ..., n_{\mathrm{D}}\}\}$ be a set of n_{D} samples of Process (1). Then, for a compact safe set $\mathcal{X}_{\text{safe}} \subset \mathbb{R}^n$, a time-horizon $T \in \mathbb{N} \cup \{\infty\}$, and every $x \in \mathcal{X}_{\text{safe}}$, compute the bounds of safety probability $P_{\text{safe}}(x)$ defined by*

$$\begin{aligned} p_{\min}(x) &\leq \min_{\pi_{\mathbf{x}}} \Pr(\forall k \in [0, T], \omega_{\mathbf{x}}(k) \in \mathcal{X}_{\text{safe}} \mid \\ &\qquad \mathbf{x}(0) = x, \pi_{\mathbf{x}}, \mathrm{D}), \\ p_{\max}(x) &\geq \max_{\pi_{\mathbf{x}}} \Pr(\forall k \in [0, T], \omega_{\mathbf{x}}(k) \in \mathcal{X}_{\text{safe}} \mid \\ &\qquad \mathbf{x}(0) = x, \pi_{\mathbf{x}}, \mathrm{D}). \end{aligned}$$

That is, $P_{\text{safe}}(x) \in [p_{\min}(x), p_{\max}(x)]$ for all possible strategies.

B. Approach

Our approach to Problem 1 is through a discrete abstraction of Process (1) in a form of an uncertain Markov decision process. A crucial part of the construction of this abstraction is the derivation of the uncertainty bounds for the transition probability of $\mathbf{x}(k)$ to region $q' \subset \mathbb{R}^n$ given that $\mathbf{x}(k-1) \in q \subset \mathbb{R}^n$. Section IV shows how these bounds can be computed by incorporating the uncertainty from the GP learning process. Intuitively, the regressed GP may not accurately approximate the posterior of Process (1) since the observation noise $\mathbf{v}$ is not Gaussian, i.e., $\mathbf{v}$ is bounded and the fact that only a finite amount of data is available. A correction term that captures this discrepancy is required. Section V proves the correctness of the proposed method.

III. PRELIMINARIES

Our approach is based on GP regression and Markov processes, which are formally defined in this section.

A. Gaussian Process Regression

Gaussian Process (GP) regression is a non-parametric Bayesian machine learning method [15]. For an unknown function $\mathrm{f} : \mathbb{R}^n \to \mathbb{R}$, the basic assumption of GP regression is that f is a sample from a GP with zero mean[1] and covariance $\kappa : \mathbb{R}^n \times \mathbb{R}^n \to \mathbb{R}_{>0}$. GP regression is often used when only noisy observations of f are available in the form $\mathrm{y} = \mathrm{f}(\mathrm{x}) + \mathrm{v}$, where v is assumed to be normally distributed with variance σ^2. Note that here $\mathrm{y}, \mathrm{v} \in \mathbb{R}$ are different from $y, v \in \mathbb{R}^n$.

Consider a data set of noisy samples $\mathrm{D} = \{(\mathrm{x}_i, \mathrm{y}_i), i \in \{1, \ldots, n_\mathrm{D}\}\}$. Let X and Y be ordered vectors with all points in D such that $X_i = \mathrm{x}_i$ and $Y_i = \mathrm{y}_i$. Further, call $K(X, X)$ the matrix with $K_{i,j}(X_i, X_j) = \kappa(\mathrm{x}_i, \mathrm{x}_j)$, $K(\mathrm{x}, X)$ the vector such that $K_i(\mathrm{x}, X) = \kappa(\mathrm{x}, X_i)$, and $K(X, \mathrm{x})$ defined accordingly. Assuming the noise is i.i.d., the predictive distribution of f at a test point x is given by the conditional distribution of f, which is Gaussian and with mean μ_D and variance σ^2_D given by

$$\begin{aligned}\mu_\mathrm{D}(\mathrm{x}) &= K(\mathrm{x}, X)\big(K(X, X) + \lambda I_{n_\mathrm{D}}\big)^{-1} Y \\ \sigma^2_\mathrm{D}(\mathrm{x}) &= \kappa(\mathrm{x}, \mathrm{x}) - \\ &\qquad K(\mathrm{x}, X)\big(K(X, X) + \lambda I_{n_\mathrm{D}}\big)^{-1} K(X, \mathrm{x}),\end{aligned}$$

where I_{n_D} is the identity matrix of size $n_\mathrm{D} \times n_\mathrm{D}$ and λ is a free parameter (often taken to be σ^2 when f is distributed according to the posterior).

In our setting, we do not assume that f is sampled from a GP and noise v is not Gaussian, so the assumptions for GP regression are not satisfied. Nevertheless, Assumption 1 permits using GP regression even in our scenario. In particular, the following Lemma provides a bound on the distance between μ_D and f so long Assumption 1 holds. This is an important result for safety verification, where the distance between the regression and the true system needs to be considered.

Lemma 1 ([20], Theorem 2)**.** *Let $\mathcal{K}$ be a compact set, $\delta \in (0, 1)$, α_D the maximum information gain parameter associated with κ and data set* D *training points, and $B > 0$ such that $\|\mathrm{f}\|_\kappa \leq B$. Assume that $|\mathrm{v}| < \sigma$ almost surely and μ_D and σ_D are found with $\lambda = 1 + 2/n_\mathrm{D}$. Define $\beta = (\sigma/\sqrt{\lambda})(B + \sigma\sqrt{2(\alpha_\mathrm{D} + 1 + \log 1/\delta)})$. Then, it holds that*

$$\Pr\big(\forall x \in \mathcal{K}, |\mu_D(\mathrm{x}) - \mathrm{f}(\mathrm{x})| \leq \beta\sigma_D(\mathrm{x})\big) \geq 1 - \delta.$$

Lemma 1 computes a probabilistic bound between the regressed GP and the underlying unknown function and takes into account the modelling errors in running GP regression with observation noise with the parameter λ and scaling factor $(\sigma/\sqrt{\lambda})$. The constraint on $\|\mathrm{f}\|_\kappa$ implies f is L_f-Lipschitz continuous with $L^2_\mathrm{f} \propto B$ [24]. The information gain term α_κ can be upper bounded for certain kernels as shown in [18].

B. Uncertain Markov Processes

A *bounded-parameter* [27] or *interval* [28] Markov decision processes (IMDP) is a generalization of the Markov decision process that allows interval-valued transition probabilities. An IMDP is defined as follows.

Definition 1 (IMDP)**.** *An interval Markov decision process (IMDP) is a tuple $\mathcal{I} = (Q, A, \check{P}, \hat{P})$, where*

- *Q is a finite set of states,*
- *A is a finite set of actions where $A(q)$ denotes the set of actions available at state $q \in Q$,*
- *$\check{P} : Q \times A \times Q \to [0, 1]$ is a function, where $\check{P}(q, a, q')$ defines the lower bound of the transition probability from state q to state q' under action $a \in A(q)$,*
- *$\hat{P} : Q \times A \times Q \to [0, 1]$ is a function, where $\hat{P}(q, a, q')$ defines the upper bound of the transition probability from state q to state q' under action $a \in A(q)$.*

For all $q, q' \in Q$ and $a \in A(q)$, it holds that $\check{P}(q, a, q') \leq \hat{P}(q, a, q')$ and

$$\sum_{q' \in Q} \check{P}(q, a, q') \leq 1 \leq \sum_{q' \in Q} \hat{P}(q, a, q').$$

A path ω of an IMDP is a sequence of states $\omega = q_0 \xrightarrow{a_0} q_1 \xrightarrow{a_1} q_2 \xrightarrow{a_2} \ldots$ such that $a_i \in A(q_i)$ and $\hat{P}(q_i, a_i, q_{i+1}) > 0$ for all $i \in \mathbb{N}$. We denote the last state of a finite path ω^{fin} by $last(\omega^{\mathrm{fin}})$ and the set of all finite and infinite paths by $Paths^{\mathrm{fin}}$ and $Paths$, respectively.

Let $\mathcal{D}(Q)$ denote the set of discrete probability distributions over Q. Given $q \in Q$ and $a \in A(q)$, we call $\gamma^a_q \in \mathcal{D}(Q)$ a *feasible distribution* reachable from q by a if

$$\check{P}(q, a, q') \leq \gamma^a_q(q') \leq \hat{P}(q, a, q')$$

for each state $q' \in Q$. We denote the set of all feasible distributions for state q and action a by Γ^a_q.

A *strategy* defines a choice of action at each state of the IMDP, and an *adversary* is a function that chooses a feasible distribution.

Definition 2 (Strategy)**.** *A strategy π of an* IMDP *model $\mathcal{I}$ is a function $\pi : Paths^{\mathrm{fin}} \to A$ that maps a finite path ω^{fin} of $\mathcal{I}$ onto an action in A.*

Definition 3 (Adversary)**.** *Given an* IMDP *$\mathcal{I}$, an adversary is a function $\gamma : Paths^{\mathrm{fin}} \times A \to \mathcal{D}(Q)$ that, for each finite path $\omega^{\mathrm{fin}} \in Paths^{\mathrm{fin}}$ and action $a \in A(last(\omega^{\mathrm{fin}}))$, assigns a feasible distribution $\gamma(\omega^{\mathrm{fin}}, a) \in \Gamma^a_{last(\omega^{\mathrm{fin}})}$.*

Given a strategy π and an adversary γ, a Markov chain is composed from an IMDP. This Markov chain defines a probability measure over the paths of the IMDP [8].

IV. ABSTRACTION

In order to solve Problem 1, we abstract Process (1) as an IMDP $\mathcal{I} = (Q, A, \check{P}, \hat{P})$ as detailed below.

A. States & Actions

First, we partition the compact safe set $\mathcal{X}_\mathrm{safe}$ into a set of cells (regions) that are non-overlapping. Let $Q_\mathrm{safe} = \{q_1, ..., q_{|Q_\mathrm{safe}|}\}$ be the resulting set of cells. Then, $\cup_{q \in Q_\mathrm{safe}} q = \mathcal{X}_\mathrm{safe}$, and

$$q \cap q' = \emptyset, \quad \forall q, q' \in Q_\mathrm{safe}, \text{ and } q \neq q'.$$

[1]Extensions with non-zero mean are a trivial generalization [15]

Each region is associated to a state of IMDP $\mathcal{I}$. With an abuse of notation, q denotes both the region, i.e., $q \subset \mathcal{X}_{\text{safe}}$, as well as its corresponding IMDP state, i.e, $q \in Q$. From the context, the correct interpretation of q should be clear. Furthermore, let q_{u} denote the unsafe set $\mathbb{R}^n \setminus \mathcal{X}_{\text{safe}}$. Then, the set of states of $\mathcal{I}$ is defined as

$$Q = Q_{\text{safe}} \cup \{q_{\text{u}}\}.$$

The set of actions A of $\mathcal{I}$ is given by the set of actions in $\mathcal{U}$, and all actions are allowed to be available at each state of $\mathcal{I}$, i.e., $A(q) = A$ for all $q \in Q$.

B. Transition Probability Bounds

In order to account for the probabilistic behavior of Process (1), we define the following conditions for the transition probability bounds of $\mathcal{I}$:

$$\check{P}(q, a, q') \leq \min_{x \in q} \Pr(\mathbf{x}(k) \in q' \mid \mathbf{x}(k-1) = x, \mathbf{u}(k-1) = a, \mathrm{D}), \quad (2)$$

$$\hat{P}(q, a, q') \geq \max_{x \in q} \Pr(\mathbf{x}(k) \in q' \mid \mathbf{x}(k-1) = x, \mathbf{u}(k-1) = a, \mathrm{D}), \quad (3)$$

for all $q, q' \in Q$. Notice that even though the action is fixed in (2) and (3), a probabilistic statement is necessary because f is unknown and the samples in D are noisy. Conditions (2) and (3) guarantee that the full probabilistic behavior of Process (1) is accounted for in the abstraction as shown in Section V. In order to compute the bounds that satisfy these conditions, we partition the set of samples D according to actions $a \in \mathcal{U}$, i.e., $\mathrm{D} = \cup_{a \in \mathcal{U}} \mathrm{D}_a$, where

$$\mathrm{D}_a = \{(x_i, a, y_i) \mid (x_i, a, y_i) \in \mathrm{D}\}.$$

GP regression on D_a for each a results in a Gaussian posterior distribution characterized by mean μ^a_{D} and diagonal covariance matrix Σ^a_{D}. Recall that, even though f is unknown and $\mathbf{v}$ is not Gaussian, Lemma 1 allows one to characterize the distance between the posterior mean μ^a_{D} and $f(\cdot, a)$.

1) Transitions to Safe States: For all the safe states $q, q' \in Q_{\text{safe}}$, the transition probability bounds in (2) and (3) are given by Proposition 1. In order to state this result, we introduce the notions of reduction and enlargement of a compact set.

For a scalar $\epsilon > 0$ and a compact set $q \subset \mathbb{R}^n$, let $\underline{q} \subset q$ be a subset of q such that the distance between each of its points to the boundary of q is at least ϵ. Moreover, let $\overline{q}$ be such that $q \subset \overline{q}$ and $\overline{q}$ contains all the points that are within a ϵ margin from the boundary of q. Sets $\underline{q}$ and $\overline{q}$ are the *ϵ-reduced* and *ϵ-enlarged* versions of q, respectively. We are now ready to state the following result:

Proposition 1. *Let $q, q' \subset \mathbb{R}^n$ be compact sets. For $\epsilon > 0$, define the enlarged and reduced sets*

$$\overline{q}' = \{x \in \mathbb{R}^n \mid \exists x' \in q' \ \ s.t. \ \|x - x'\|_\infty \leq \epsilon\}$$

and

$$\underline{q}' = \{x \in q' \mid \forall x' \in \partial q', \ \|x - x'\|_\infty > \epsilon\},$$

where $\partial q'$ is the boundary of q'. Then, for a given action $a \in \mathcal{U}$, it holds that

$$\min_{x \in q} \Pr(\mathbf{x}(k) \in q' \mid \mathbf{x}(k-1) = x, \mathbf{u}(k-1) = a, \mathrm{D}) \geq \min_{x \in q} \Big(\mathbf{1}_{\underline{q}'}(\mu^a_{\mathrm{D}}(x)) \prod_{i=1}^{n} \Pr(\forall x' \in q, |f_i(x', a) - \mu^a_{i,\mathrm{D}}(x')| \leq \epsilon \mid \mathrm{D}) \Big),$$

and

$$\max_{x \in q} \Pr(\mathbf{x}(k) \in q' \mid \mathbf{x}(k-1) = x, \mathbf{u}(k-1) = a, \mathrm{D}) \leq \max_{x \in q} \Big(1 - \prod_{i=1}^{n} \Pr\big(\forall x' \in q, \ |f_i(x', a) - \mu^a_{i,\mathrm{D}}(x')| \leq \epsilon \mid \mathrm{D}\big) \big(1 - \mathbf{1}_{\overline{q}'}(\mu^a_{\mathrm{D}}(x))\big) \Big),$$

where $\mu^a_{i,\mathrm{D}}$ is the i-th component of vector μ^a_{D}, and $\mathbf{1}_H(h)$ is the indicator function which is 1 if $h \in H$ and 0 otherwise.

Proposition 1 guarantees that upper and lower bounds of $\Pr(\mathbf{x}(k) \in q' \mid \mathbf{x}(k-1) \in q, \mathbf{u}(k-1) = a, \mathrm{D})$ can be derived by checking if the posterior mean is within a reduced (or enlarged) version of q' and computing a uniform error bound on the distance between the posterior mean of the GP learnt via GP regression and f (the underlying dynamics in Process (1)). Such a bound can be computed by employing Lemma 1. Proposition 1 and Lemma 1 are combined in the following theorem.

Theorem 1. *Let $q, q', \overline{q}', \underline{q}'$ be as defined in Proposition 1. For $i \in \{1, ..., n\}$, consider $B_i > 0$ such that for a given $a \in \mathcal{U}$, $\|f(\cdot, a)\|_\kappa \leq B_i$. Define β_i as in Lemma 1 and for $\delta \in (0, 1)$ select $\epsilon = \max_{i \in \{1,...,n\}} \beta_i^{\frac{1}{2}} (\Sigma^{a,(i,i)}_{\mathrm{D}})^{\frac{1}{2}}$, where $\Sigma^{a,(i,i)}_{\mathrm{D}}$ is the i-th element of the diagonal of Σ^a_{D}. Then, it holds that*

$$\min_{x \in q} \Pr(\mathbf{x}(k) \in q' \mid \mathbf{x}(k-1) = x, \mathbf{u}(k-1) = a, \mathrm{D}) \geq \min_{x \in q} \mathbf{1}_{\underline{q}'}(\mu^a_{\mathrm{D}}(x))(1 - \delta)^n,$$

and

$$\max_{x \in q} \Pr(\mathbf{x}(k) \in q' \mid \mathbf{x}(k-1) = x, \mathbf{u}(k-1) = a, \mathrm{D}) \leq \max_{x \in q} \Big(1 - (1 - \delta)^n (1 - \mathbf{1}_{\overline{q}'}(\mu^a_{\mathrm{D}}(x))) \Big).$$

Note that Theorem 1 holds for every choice of constant δ, and hence of ϵ. As discussed in Section VI, this constant should be selected in order to maximize the tightness of the bound. In fact, ϵ controls both the tightness of the bound between the posterior mean and the underlying system and how much q' is reduced and enlarged.

Remark 1. *Proposition 1 is general; it does not make use of Assumption 1. It just assumes the existence of a bound between the posterior mean and the unknown function. Therefore, it can be applied to other settings where Assumption 1 is not satisfied. For instance, if function f is a sample from a*

GP prior with Gaussian observation noise, our framework can still be used in combination with existing error bounds developed for this scenario, such as those in [21].

2) Transitions to Unsafe State: We obtain upper and lower bounds for the transitions to the unsafe region q_{u} as a corollary of Theorem 1. That is, for every $q \in Q_{\mathrm{safe}}$,

$$\check{P}(q, a, q_u) = 1 - \max_{x \in q} \Pr(\mathbf{x}(k) \in \mathcal{X}_{\mathrm{safe}} \mid \mathbf{x}(k-1) = x, \mathbf{u}(k-1) = a, \mathrm{D}),$$

$$\hat{P}(q, a, q_u) = 1 - \min_{x \in q} \Pr(\mathbf{x}(k) \in \mathcal{X}_{\mathrm{safe}} \mid \mathbf{x}(k-1) = x, \mathbf{u}(k-1) = a, \mathrm{D}).$$

Both of these terms can be computed by employing Theorem 1. To complete the construction of abstraction $\mathcal{I}$, we make q_{u} absorbing, i.e., $\check{P}(q_{\mathrm{u}}, a, q_{\mathrm{u}}) = \hat{P}(q_{\mathrm{u}}, a, q_{\mathrm{u}}) = 1$ for all $a \in A$, to ensure that $\mathcal{I}$ does not count the transitions to $\mathcal{X}_{\mathrm{safe}}$ from q_{u} of Process (1) as a safe behavior.

V. VERIFICATION

Given the IMDP abstraction $\mathcal{I}$, we are interested in computing the probabilities of remaining in Q_{safe} for $T \in \mathbb{N} \cup \{\infty\}$ time steps from every $q \in Q_{\mathrm{safe}}$. Note that, under strategy π, the safety probability is a range due to the transition probability intervals of $\mathcal{I}$. The values in this range correspond to the feasible transition probabilities $\gamma_q^a \in \Gamma_q^a$ at every state $q \in Q$ and action $a \in A$ chosen by π. The choice of this feasible transition probability is made by adversary γ. Therefore, the minimum safety probability is achieved when both strategy π and adversary γ are minimizing the safety probability. Similarly, the maximum safety probability is given when both π and γ are maximizing.

This optimization problem can be formulated through the Bellman equation as detailed in [8]. Let $\check{p}^k(q)$ and $\hat{p}^k(q)$ denote the minimum and maximum probability of remaining safe in k time steps starting from state $q \in Q$, respectively. Then, the safety probability bounds for a finite time duration T can be computed by T recursive evaluations of

$$\check{p}^k(q) = \min_{a \in A(q)} \min_{\gamma_q^a \in \Gamma_q^a} \sum_{q' \in Q} \gamma_q^a(q')\, \check{p}^{k-1}(q') \tag{4}$$

$$\hat{p}^k(q) = \max_{a \in A(q)} \max_{\gamma_q^a \in \Gamma_q^a} \sum_{q' \in Q} \gamma_q^a(q')\, \hat{p}^{k-1}(q') \tag{5}$$

with initial values of $\check{p}^0(q) = 1$ for $q \in Q_{\mathrm{safe}}$ and $\check{p}^0(q_{\mathrm{u}}) = 0$. In the case of an infinite time horizon T, recursive evaluations of (4) and (5) need to continue until convergence, which is guaranteed to occur in finite time [8].

This method of evaluation is similar to value iteration. The additional step involves first optimizing over the adversaries, which itself can be performed iteratively via an ordering of the states in Q according to their values [8]. Once the optimal adversaries are obtained for all $a \in A(q)$, an optimization over the actions is performed to complete the computation for one time step in (4) and (5). This algorithm computes the safety probability bounds $\check{p}^T(q)$ and $\hat{p}^T(q)$ for each $q \in Q$. The complexity of the algorithm is polynomial in the size of the IMDP $\mathcal{I}$ [8].

A. Correctness

The following theorem guarantees that the safety probability ranges computed by our framework are sound, i.e., they give lower and upper bounds for $p_{\min}(x)$ and $p_{\max}(x)$ as defined in Problem 1.

Theorem 2. *Let $x \in \mathcal{X}_{\mathrm{safe}}$ and $q \in Q$ such that $x \in q$. Then, it holds that $[p_{\min}(x), p_{\max}(x)] \subseteq [\check{p}^T(q), \hat{p}^T(q)]$.*

VI. CASE STUDIES

We evaluate the performance of our framework in three case studies. The first case study involves three single-action linear systems and shows the effect of various choices for parameter ϵ. The second case uses two of the linear systems to define a switched system with two actions. The final case considers the safety a nonlinear system. Additional results can be found in our extended manuscript [29].

In all three case studies, $\mathcal{X}_{\mathrm{safe}}$ is a two-dimensional square defined by the region $[-4, 4] \times [-4, 4]$. We performed a regression of each dynamical system using a pair of Gaussian processes, one for each output component. The GP prior used the zero mean and squared-exponential functions. The training process used one thousand i.i.d. training points with noise parameter $\sigma = 0.01$ to optimize the hyperparameters of the prior functions and train the Gaussian processes using the GaussianProcesses.jl Julia package [30]. We modified an existing tool to perform the verification over the resulting IMDP [8].

A. Single-Action Linear Systems

We performed the verification procedure on three linear systems $f(\mathbf{x}(k)) = A_i\, \mathbf{x}(k)$, $i \in \{\text{rotation}, \text{upper}, \text{lower}\}$,

$$A_{\text{rotation}} = \begin{bmatrix} 0.9 & -0.4 \\ 0.4 & 0.5 \end{bmatrix},$$

$$A_{\text{upper}} = \begin{bmatrix} 0.8 & 0.5 \\ 0 & 0.5 \end{bmatrix}, \qquad A_{\text{lower}} = \begin{bmatrix} 0.5 & 0 \\ -0.5 & 0.8 \end{bmatrix}.$$

We discretized the safe set $\mathcal{X}_{\mathrm{safe}}$ into squares with side length 0.25. Figure 1 shows the 10-step safety probability ($T = 10$) for each cell using $\epsilon = 0.12$. The legend above the figures maps the intensity of the shade of each cell to a probability value between zero and one. The white cells in Figures 1(a), (c) and (e) correspond to a minimum safety probability of one. If the system is initialized within one of these cells, then it is certain to remain in the safe set.

Figures 1(a), (c) and (e) also include cells where the minimum probability of safety is zero due to flow that leaves the $\mathcal{X}_{\mathrm{safe}}$ before returning. This does not mean that it is impossible to stay in the safe set starting at one of these cells, because the maximum probability of safety may be greater than zero. The maximum probability is shown in Figures 1(b), (d) and (f). These results indicate that if the system were initialized in the cells with a maximum probability of safety near zero (e.g. the black corners in (b)), it is certain to leave $\mathcal{X}_{\mathrm{safe}}$. Cells with a safety probability minimum of zero

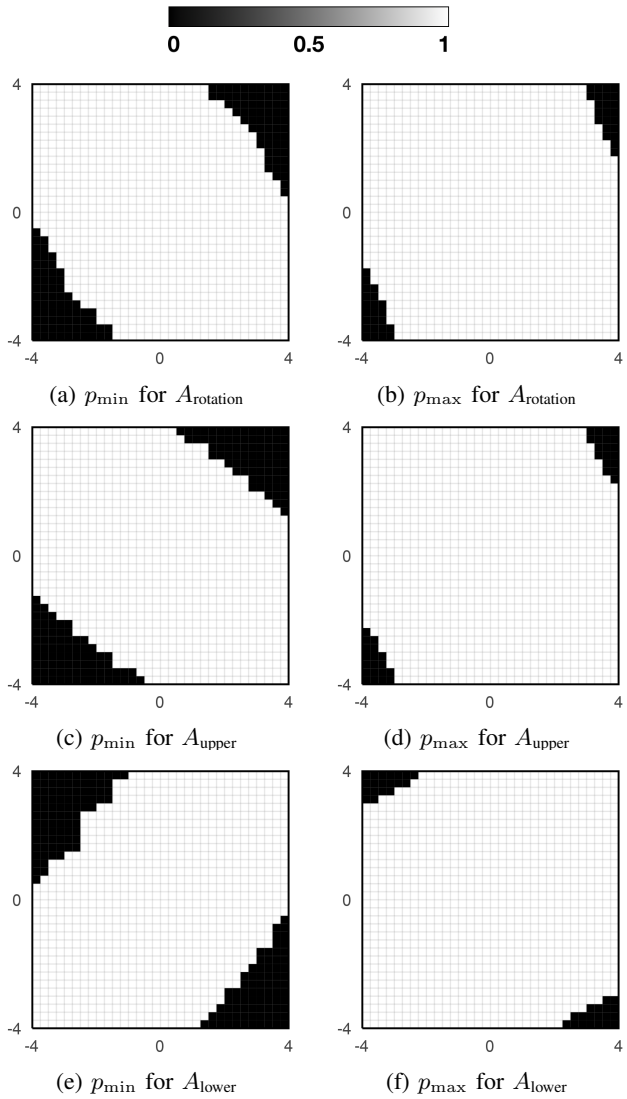

(a) $p_{\min}$ for A_{rotation} (b) $p_{\max}$ for A_{rotation}

(c) $p_{\min}$ for A_{upper} (d) $p_{\max}$ for A_{upper}

(e) $p_{\min}$ for A_{lower} (f) $p_{\max}$ for A_{lower}

Fig. 1: Minimum and maximum 10-step ($T = 10$) safety probability for the linear systems with $\epsilon = 0.12$.

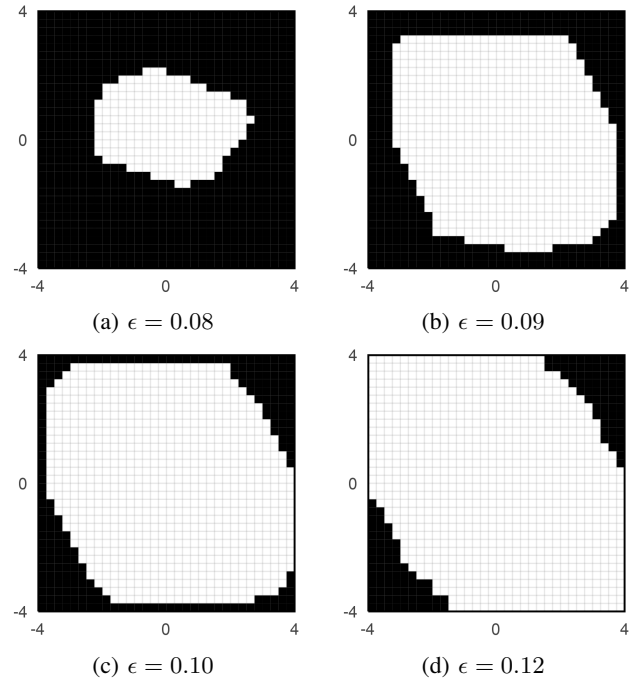

(a) $\epsilon = 0.08$ (b) $\epsilon = 0.09$

(c) $\epsilon = 0.10$ (d) $\epsilon = 0.12$

Fig. 2: Minimum 2-step safety probability for A_{rotation} for different values of ϵ.

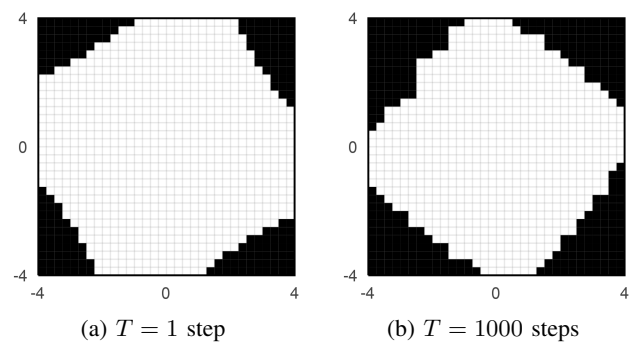

(a) $T = 1$ step (b) $T = 1000$ steps

Fig. 3: Minimum probability of safety for a switched system comprised of the A_{upper} and A_{lower} systems.

and maximum of one essentially indicate a nondeterministic transition to a safe or unsafe cell after T steps.

The effect on the verification results of reducing ϵ is shown in Figure 2 for A_{rotation}. The initial surely-safe areas diminish as ϵ decreases until the minimum safety probability becomes zero nearly everywhere. This highlights a trade-off when choosing ϵ. Recall that Proposition 1 depends on enlarging and reducing the target set. Small ϵ shrinks and enlarges the sets less, which can be beneficial when calculating the transition probabilities. Small ϵ also tightens the bound on the distance between the system and the process. Too small, and the resulting probabilities become trivial everywhere, i.e., a minimum of zero and maximum of one.

B. Switched Linear System

The switched system uses the A_{upper} and A_{lower} systems from the previous section, and enables switching between the two at each time step. The verification used the previous discretization of $\mathcal{X}_{\text{safe}}$ and $\epsilon = 0.12$. Recall that the verification problem aims to check if this system remains in the safe set for all possible strategies. With two actions available to the system, the worst-case result occurs if one action could drive the system to an "unsafe" region of the other action. Figure 3 shows the minimum probability of safety after one and 1000 steps. Due to the tight results of the component systems, the verification output of the switched system happens to be the superposition of the individual verification outputs. Notably, the system is guaranteed to remain in the safe set after 1000 steps *regardless of the underlying strategy* so long it starts in a cell with a minimum safety probability of one.

C. Nonlinear System

We demonstrate the verification on a nonlinear system given by

$$f(\mathbf{x}(k)) = [\,\mathbf{x}_1(k) - 0.05\,\mathbf{x}_2(k),\ \mathbf{x}_2(k) + 0.1\sin(\mathbf{x}_1(k))\,]^T$$

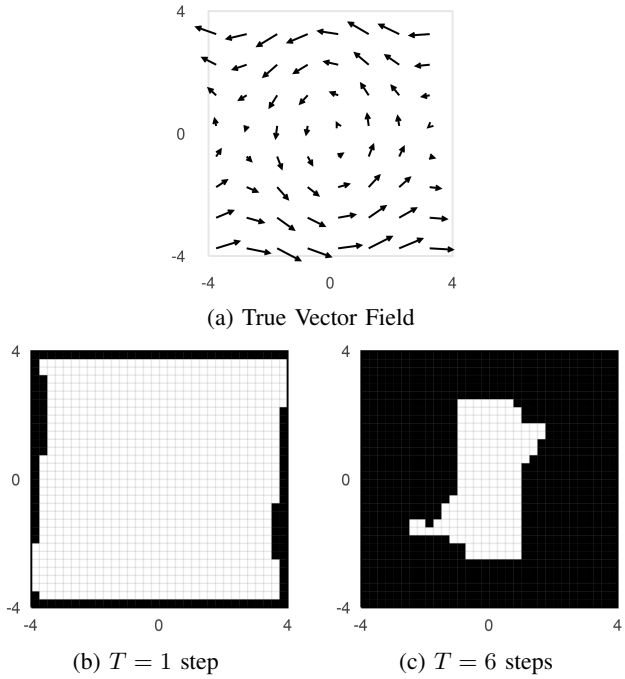

(a) True Vector Field

(b) $T = 1$ step

(c) $T = 6$ steps

Fig. 4: Vector fields and the minimum safety probability for multiple steps of the nonlinear system.

over a discretization of $\mathcal{X}_{\text{safe}}$ with squares of side length 0.25. The vector field for the true system is shown in Figure 4(a). After 1 step, the minimum probability of safety is zero around parts of the field that flow out of $\mathcal{X}_{\text{safe}}$ shown in Figure 4(b). However, the non-zero maximum probability of transitioning to parts of the field that flows out of $\mathcal{X}_{\text{safe}}$ cause the initially-large set to shrink after successive steps. After 6 steps, safety is guaranteed if the system starts in regions around the origin.

VII. CONCLUSION

We introduced a novel verification framework that generates safety guarantees for unknown dynamical systems. The approach is based on GP regression and an uncertain abstraction that incorporates probabilistic error bounds between the model learned from data and the underlying system. This versatile framework paves the way for generating guarantees for safety-critical systems that have black-box components.

REFERENCES

[1] E. M. Clarke, O. Grumberg, and D. Peled, *Model Checking*. MIT Press, 1999.

[2] C. Baier and J.-P. Katoen, *Principles of Model Checking*. Cambridge, MA: The MIT Press, 2008.

[3] P. Tabuada, *Verification and control of hybrid systems: a symbolic approach*. Springer Science & Business Media, 2009.

[4] C. Belta, B. Yordanov, and E. A. Gol, *Formal methods for discrete-time dynamical systems*, vol. 89. Springer, 2017.

[5] L. Doyen, G. Frehse, G. J. Pappas, and A. Platzer, "Verification of hybrid systems," in *Handbook of Model Checking*, pp. 1047–1110, Springer, 2018.

[6] H. Kushner and P. G. Dupuis, *Numerical methods for stochastic control problems in continuous time*, vol. 24. Springer Science & Business Media, 2013.

[7] S. E. Z. Soudjani, C. Gevaerts, and A. Abate, "Faust2: Formal abstractions of uncountable-state stochastic processes," in *International Conference on Tools and Algorithms for the Construction and Analysis of Systems*, pp. 272–286, Springer, 2015.

[8] M. Lahijanian, S. B. Andersson, and C. Belta, "Formal verification and synthesis for discrete-time stochastic systems," *IEEE Transactions on Automatic Control*, vol. 60, pp. 2031–2045, Aug. 2015.

[9] L. Laurenti, M. Lahijanian, A. Abate, L. Cardelli, and M. Kwiatkowska, "Formal and efficient synthesis for continuous-time linear stochastic hybrid processes," *IEEE Transactions on Automatic Control*, 2020.

[10] A. Girard and G. J. Pappas, "Approximation metrics for discrete and continuous systems," *IEEE Transactions on Automatic Control*, vol. 52, no. 5, pp. 782–798, 2007.

[11] S. Dutta, S. Jha, S. Sankaranarayanan, and A. Tiwari, "Learning and verification of feedback control systems using feedforward neural networks," *IFAC-PapersOnLine*, vol. 51, no. 16, pp. 151–156, 2018.

[12] S. Haesaert, P. M. Van den Hof, and A. Abate, "Data-driven and model-based verification via bayesian identification and reachability analysis," *Automatica*, vol. 79, pp. 115–126, 2017.

[13] J. Kenanian, A. Balkan, R. M. Jungers, and P. Tabuada, "Data driven stability analysis of black-box switched linear systems," *Automatica*, vol. 109, p. 108533, 2019.

[14] M. Ahmadi, A. Israel, and U. Topcu, "Safety assessemt based on physically-viable data-driven models," in *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*, pp. 6409–6414, IEEE, 2017.

[15] C. E. Rasmussen, "Gaussian processes in machine learning," in *Summer School on Machine Learning*, pp. 63–71, Springer, 2003.

[16] L. Cardelli, M. Kwiatkowska, L. Laurenti, and A. Patane, "Robustness guarantees for bayesian inference with gaussian processes," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 33, pp. 7759–7768, 2019.

[17] F. Berkenkamp and A. P. Schoellig, "Safe and robust learning control with gaussian processes," in *2015 European Control Conference (ECC)*, pp. 2496–2501, IEEE, 2015.

[18] N. Srinivas, A. Krause, S. M. Kakade, and M. W. Seeger, "Information-theoretic regret bounds for gaussian process optimization in the bandit setting," *IEEE Transactions on Information Theory*, vol. 58, no. 5, pp. 3250–3265, 2012.

[19] P. Germain, F. Bach, A. Lacoste, and S. Lacoste-Julien, "Pac-bayesian theory meets bayesian inference," in *Advances in Neural Information Processing Systems*, pp. 1884–1892, 2016.

[20] S. R. Chowdhury and A. Gopalan, "On kernelized multi-armed bandits," in *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pp. 844–853, JMLR. org, 2017.

[21] A. Lederer, J. Umlauft, and S. Hirche, "Uniform error bounds for gaussian process regression with application to safe control," in *Advances in Neural Information Processing Systems*, pp. 657–667, 2019.

[22] A. K. Akametalu, S. Kaynama, J. F. Fisac, M. N. Zeilinger, J. H. Gillula, and C. J. Tomlin, "Reachability-based safe learning with Gaussian processes," in *IEEE 53rd Annual Conference on Decision and Control (CDC), 2014: 15-17 Dec. 2014, Los Angeles, California, USA*, pp. 1424–1431, IEEE, 2014.

[23] Y. Sui, A. Gotovos, J. W. Burdick, and A. Krause, "Safe exploration for optimization with gaussian processes," *Proceedings of Machine Learning Research*, vol. 37, pp. 997–1005, 2015.

[24] F. Berkenkamp, M. Turchetta, A. Schoellig, and A. Krause, "Safe model-based reinforcement learning with stability guarantees," in *Advances in neural information processing systems*, pp. 908–918, 2017.

[25] K. Polymenakos, L. Laurenti, A. Patane, J.-P. Calliess, L. Cardelli, M. Kwiatkowska, A. Abate, and S. Roberts, "Safety guarantees for planning based on iterative gaussian processes," *arXiv preprint arXiv:1912.00071*, 2019.

[26] I. Steinwart, "On the influence of the kernel on the consistency of support vector machines," *Journal of machine learning research*, vol. 2, no. Nov, pp. 67–93, 2001.

[27] R. Givan, S. Leach, and T. Dean, "Bounded-parameter Markov decision processes," *Artificial Intelligence*, vol. 122, no. 1-2, pp. 71–109, 2000.

[28] E. M. Hahn, V. Hashemi, H. Hermanns, M. Lahijanian, and A. Turrini, "Multi-objective robust strategy synthesis for interval Markov decision processes," in *Int. Conf. on Quantitative Evaluation of SysTems (QEST)*, (Berlin, Germany), pp. 207–223, Springer, Sep. 2017.

[29] J. Jackson, L. Laurenti, E. Frew, and M. Lahijanian, "Safety verification of unknown dynamical systems via gaussian process regression," *arXiv preprint:2004.01821*, 2020.

[30] J. Fairbrother, C. Nemeth, M. Rischard, J. Brea, and T. Pinder, "Gaussianprocesses. jl: A nonparametric bayes package for the julia language," *arXiv preprint:1812.09064*, 2018.