Towards Data-driven Verification of Unknown Dynamical Systems

John Jackson*, Luca Laurenti†, Eric Frew*, Morteza Lahijanian*
*Smead Aerospace Engineering Sciences, University of Colorado Boulder
{john.m.jackson, eric.frew, morteza.lahijanian}@colorado.edu
†Department of Computer Science, University of Oxford, UK
luca.laurenti@cs.ox.ac.uk

***Abstract*—This paper summarizes an abstraction-based formal verification framework for unknown dynamical systems. At its core, it uses Gaussian process regression to learn unmodelled dynamics from a given data set of noisy system measurements. While this paper envisions such a framework, details can be found in [8].**

I. Introduction

As the capabilities and complexities of autonomous systems grow, the ability to provide guarantees for safety-critical systems becomes more challenging. Conventional methods for the formal verification of control systems generate guarantees with respect to a given system model [5, 2]. In many applications, a system may have unmodelled dynamics in order to reduce complexity (e.g. linearization) or because the resources required to obtain a high-fidelity model are unavailable. Both of these factors may apply to a black-box system, such as a closed-source autopilot or an off-the-shelf component. Factors that were unaccounted for, including control actuator changes and environmental impacts, may render initial safety guarantees inapplicable.

Data-driven verification approaches can overcome these limitations if we can account for the modelling error. Haesaert et al. [7] and Kenanian et al. [9] focus on learning linear systems from data and computing a measure of safety, though their method is limited to linear systems. Similarly, Ahmadi et al. [1] use a piece-wise polynomial approximation to fit a potentially nonlinear model for safety assessment. While they produce safety guarantees using barrier certificates, the safety result is not applicable to the underlying system.

Recently, Gaussian process (GP) regression has become popular for learning dynamics from data, as the procedure can produce probabilistic error bounds between the regression and unknown system [13, 4]. GP regression has been used in the context of reachability and policy learning with stability guarantees [14, 3]. This highlights the opportunity for formal verification approaches that use a data-driven model *and* incorporate the modelling error to make safety guarantees with respect to the unknown system.

Our work combines GP regression with an extension of model-based formal verification [10] to generate safety guarantees for systems learned from data. For example, one may ask if the state safety property "Do not enter the unsafe region with a minimum probability of 99% over the next 1 hour" is satisfied from a subset of possible initial states. We accomplish this by using GP regression to learn the unknown dynamics, then build an uncertain Markov decision process whose transitions are defined using the probabilistic error bounds from the regression, and finally employ a worst-case and best-case solver to find a lower-bound probability of satisfying the safety property over a (possibly unbounded) time horizon. For a more in-depth discussion of the presented concepts, we refer interested readers to the full version of our paper [8].

II. Problem and Approach

We consider a discrete-time control system

$$\mathbf{x}(k+1) = f(\mathbf{x}(k), \mathbf{u}(k)),$$
$$\mathbf{y}(k) = \mathbf{x}(k) + \mathbf{v}(k, \mathbf{u}(k-1)),$$

where

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

$f : \mathbb{R}^n \times \mathcal{U} \to \mathbb{R}^n$ is the unknown dynamics function, $\mathcal{U} = \{a_1, \ldots, a_{|\mathcal{U}|}\}$ 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 $\mathcal{D}_a$.

The problem is summarized as follows. Let $\mathcal{X}_{\text{safe}} \subset \mathbb{R}^n$ be a compact safe set. The safety probability $P_{\text{safe}}(x) \in [0, 1]$ measures the chance that a trajectory beginning at an initial state $x \in \mathcal{X}_{\text{safe}}$ remains in the safe set over a time horizon $T \in \mathbb{N}_+ \cup \{\infty\}$. Given a data set D of state-action-observation tuples, find $p_{\min}(x)$ and $p_{\max}(x)$ such that $P_{\text{safe}}(x) \in [p_{\min}(x), p_{\max}(x)]$ for every initial state $x \in \mathcal{X}_{\text{safe}}$ under any choice of control $\mathbf{u}(k)$ for all $k \leq T$.

Our framework uses GP regression and probabilistic error analysis to learn the latent function f from noisy observations [12]. We assume that f is a member of the reproducing kernel Hilbert space (RKHS) associated with the prior kernel function used for the GP regression. Let f_i denote the ith component of the dynamics, and let μ_i and σ_i denote the posterior mean and covariance functions, respectively, derived from measurements of f_i. We use a result from Chowdhury and Gopalan [4] that (informally) states with probability at least $1 - \delta$,

$$|f_i(x) - \mu_i(x)| \leq \Big(B_i + R\sqrt{2(1 + \gamma_D + \log(1/\delta))}\Big)\sigma_i(x),$$

where B_i is an upper bound on the RKHS norm of f_i, R is a parameter related to the noise distribution, and γ_D is the maximal information gain term. We note that the proposed framework is also applicable to the case of f being drawn directly from the GP using different probabilistic error bounds [11].

We create an interval (uncertain) Markov decision process (IMDP) [6] abstraction of the system over a discretization of the state space. The decisions correspond to the control actions afforded to the system, and the discrete state labels are consistent with the continuous space. The transition probability bounds between two IMDP states q, q' are computed by checking the intersection between the image of q under the posterior mean function and q'. We incorporate the probabilistic model error into the transition bounds to tie the IMDP to the underlying system.

Finally, we find bounds for $p_{\min}(x)$ and $p_{\max}(x)$ by using an existing method for solving IMDPs that results in best- and worst-case analyses. The verification procedure finds a lower and upper bound of the probability of satisfaction by finding those transition probabilities and control actions of the IMDP that are minimizing and maximizing, respectively. These double minimization and maximization problems can be formulated as Bellman equations that are evaluated at every step [10]. This method can be used for both finite- and infinite-horizon problem statements, and the complexity of the algorithm is polynomial in the size of the IMDP.

III. Examples

We demonstrate the framework on a two-mode switched linear system and a nonlinear system. The safe set is defined as the square $[-4, 4] \times [-4, 4]$, and the verification task is staying inside the safe set over a time horizon. One-thousand random data points with additive noise were used for the regression.

A. Switched Linear System

The switched system $\mathbf{x}(k+1) = A_i \mathbf{x}(k)$ where $i \in \{\text{upper}, \text{lower}\}$ uses the matrices in Figure 2(a) and is permitted to switch between the two modes at each time step. With two actions available to the system, the worst-case result occurs if one action can drive the system to an "unsafe" region of the other mode. Figure 1 shows the minimum probability of safety after one and 1000 steps. Notably, the system is guaranteed to remain in the safe set after 1000 steps *regardless of the applied control* so long it starts in a cell with a minimum safety probability of one.

B. Nonlinear System

The nonlinear system is given by

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

The vector field for the true system is shown in Figure 2(b). Many vectors flow away and out of $\mathcal{X}_{\text{safe}}$ near parts of the border, while the field slowly spirals away from the origin. 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 2(c).

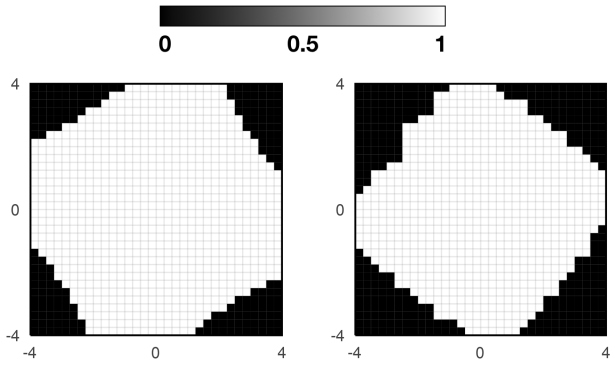

(a) $T = 1$ step (b) $T = 1000$ steps

Fig. 1: Minimum probability of safety for a switched system comprised of the upper and lower triangular system modes.

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

$$A_{\text{upper}} = \begin{bmatrix} 0.8 & 0.5 \\ 0 & 0.5 \end{bmatrix}$$

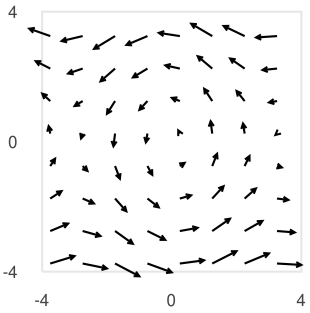

(a) Switched-system matrices (b) Nonlinear vector field

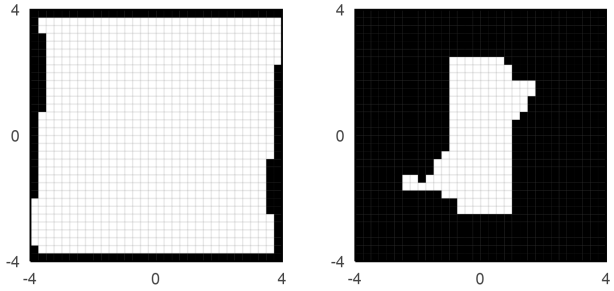

(c) $T = 1$ step (d) $T = 6$ steps

Fig. 2: Matrices for the linear switched system, nonlinear vector field and the minimum safety probability for multiple steps of the nonlinear system.

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 can only be guaranteed if the system starts in regions around the origin.

IV. Conclusion

In this short paper, we summarize our framework for the verification of unknown dynamical systems from noisy measurements.

REFERENCES

[1] Mohamadreza Ahmadi, Arie Israel, and Ufuk Topcu. Safety assessemt based on physically-viable data-driven models. In *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*, pages 6409–6414. IEEE, 2017.

[2] Christel Baier and Joost-Pieter Katoen. *Principles of Model Checking*. The MIT Press, Cambridge, MA, 2008.

[3] Felix Berkenkamp, Matteo Turchetta, Angela Schoellig, and Andreas Krause. Safe model-based reinforcement learning with stability guarantees. In *Advances in neural information processing systems*, pages 908–918, 2017.

[4] Sayak Ray Chowdhury and Aditya Gopalan. On kernelized multi-armed bandits. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 844–853. JMLR. org, 2017.

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

[6] Robert Givan, Sonia Leach, and Thomas Dean. Bounded parameter markov decision processes. In *European Conference on Planning*, pages 234–246. Springer, 1997.

[7] Sofie Haesaert, Paul MJ Van den Hof, and Alessandro Abate. Data-driven and model-based verification via bayesian identification and reachability analysis. *Automatica*, 79:115–126, 2017.

[8] John Jackson, Luca Laurenti, Eric Frew, and Morteza Lahijanian. Safety verification of unknown dynamical systems via gaussian process regression. *arXiv preprint arXiv:2004.01821*, 2020.

[9] Joris Kenanian, Ayca Balkan, Raphael M Jungers, and Paulo Tabuada. Data driven stability analysis of black-box switched linear systems. *Automatica*, 109:108533, 2019.

[10] Morteza Lahijanian, Sean B. Andersson, and Calin Belta. Formal verification and synthesis for discrete-time stochastic systems. *IEEE Transactions on Automatic Control*, 60(8):2031–2045, Aug. 2015.

[11] Armin Lederer, Jonas Umlauft, and Sandra Hirche. Uniform error bounds for gaussian process regression with application to safe control. In *Advances in Neural Information Processing Systems*, pages 657–667, 2019.

[12] Carl Edward Rasmussen. Gaussian processes in machine learning. In *Summer School on Machine Learning*, pages 63–71. Springer, 2003.

[13] Niranjan Srinivas, Andreas Krause, Sham M Kakade, and Matthias W Seeger. Information-theoretic regret bounds for gaussian process optimization in the bandit setting. *IEEE Transactions on Information Theory*, 58(5):3250–3265, 2012.

[14] Yanan Sui, Alkis Gotovos, Joel W Burdick, and Andreas Krause. Safe exploration for optimization with gaussian processes. *Proceedings of Machine Learning Research*, 37:997–1005, 2015.