Towards Efficient Data Refinement for Data-driven Abstraction and Verification

John Skovbekk[1], **Luca Laurenti**[2], **Eric Frew**[1], **Morteza Lahijanian**[1]

[1]University of Colorado, Boulder, {firstname.lastname}@colorado.edu
[2]TU Delft, L.Laurenti@tudelft.nl

Introduction

Recent advances in Artificial Intelligence (AI) have propelled its integration into autonomous systems (Winfield et al. 2019). However, applying AI to *safety-critical* systems poses a significant challenge, demanding rigorous verification of decision-making properties for safety assurance. This challenge has given rise to an active research area known as data-driven verification (or control synthesis), focusing on verifying properties of AI-enabled systems against logical specifications using data (Fan et al. 2017; Jackson et al. 2020; Badings et al. 2023; Reed, Laurenti, and Lahijanian 2023; Martin, Schön, and Allgöwer 2023). Current approaches include statistical- and Bayesian-based methods, which respectively provide confidence-based and hard probabilistic results (Fan et al. 2017; Jackson et al. 2020). While Bayesian methods provide harder guarantees, their effectiveness relies heavily on data *quality*. That is, there is a trade-off between amount of data and computational cost. Hence, in situations with limited data, a critical question emerges: *where should one strategically collect data to optimize the verification outcome, a concept known as* data-refinement*?*

In this talk, we present our progress in addressing the above question, with a focus on efficient data-refinement for Interval Markov Decision Process (IMDP) abstraction and verification. Our previous work laid the foundation for constructing IMDP abstraction models from data using Gaussian process (GP) regression (Jackson et al. 2020; Skovbekk et al. 2023), and investigated online data collection (Jackson et al. 2021). Similarly, work (Jiang, Zhao, and Coogan 2022) performs safe exploration on IMDPs via sub-optimal random sampling. Given the trade-off between size of dataset and computational cost, an important aspect of data-refinement must be data-efficiency, especially in online settings. Our ongoing work aims to advance data-refinement strategies for IMDPs that can efficiently provide consistent improvements in the abstraction and verification results.

Data-driven IMDP Abstraction & Verification

IMDPs generalize MDPs with uncertain transition probabilities between each pair of states, which lay in independent intervals (Givan, Leach, and Dean 2000). They are powerful models for abstracting systems where exact transition probabilities cannot be computed. IMDPs are particularly effective for abstracting nonlinear systems, including those learned via deep kernels (Reed, Laurenti, and Lahijanian 2023) and GP regression (Skovbekk et al. 2023; Jiang, Zhao, and Coogan 2022). The procedure first discretizes the state space and then computes lower- and upper-bounds for the transition probabilities between the discretized states to account for the multiple sources of uncertainty.

Consider stochastic system $x^+ = f(x, a) + w$, where $x \in X \subset \mathbb{R}^n$, a is a general decision or control input, w is a sub-Gaussian random variable, and f is unknown. Instead an estimate $\hat{f}$ is available (via, e.g., GP regression or deep kernels) with learning error e_L. When performing IMDP abstraction, there are three sources of uncertainty: 1) the discretization of the set X, 2) the disturbance w, and 3) the learning error e_L. Each of these can be accounted for in the transition probability intervals as described in our previous work (Jackson et al. 2020; Skovbekk et al. 2023).

The IMDP abstraction can then be verified against a temporal logic specification ϕ, including probabilistic computational-tree logic (Lahijanian, Andersson, and Belta 2015) and linear temporal logic over finite traces (Wells et al. 2020) specifications. The verification result is a probability interval for each state that contains the true probability of satisfying ϕ. As the abstraction is sound (uncertainty is correctly handled), the satisfaction interval also contains the probability that the original stochastic system satisfies ϕ (Jackson et al. 2020).

Given a probability threshold for the satisfaction of the specification, the states of the IMDP (and consequently the original system) can be classified as satisfying, violating, or *possibly* satisfying. The latter is undesirable in verification, and the goal is to reduce these states. The root cause of existence of these states is typically high error values in the abstraction due to, e.g., the space discretization being too coarse or the learning uncertainty being too high. *Discretization-refinement* improves the abstraction (reduces its error) by making the discretized states finer, reducing the size of transition probability intervals. While this can have a positive effect on the verification results, discretization-refinement is computationally expensive and can lead to the state-explosion problem (Valmari 1996). Alternatively, *data-refinement* directly reduces the learning error e_L by col-

AAAI 2024 Spring Symposium on User-Aligned Assessment of Adaptive AI Systems. Stanford University, Stanford, CA, USA.

lecting additional data to improve the abstraction without changing the space discretization. Our first motivating question is *when is it more beneficial to perform data-refinement over state-refinement?* Currently, heuristics are used to guide data-collection on an IMDP abstraction at each step after deployment of the system, such as choosing the action that provides the most progression towards satisfying ϕ (Jackson et al. 2021). However, those heuristics do not provide guarantees on the quality of the final abstraction or verification result. Our second motivating question is *which data-refinement strategies lead to verification improvement?*

Space vs. Data Refinement

To determine the impact of the learning error in the verification result, we run the abstraction procedure with $\hat{f}$ alone and neglect the learning error e_L, essentially codifying $\hat{f}$ as the true model (similar to the concept of certainty-equivalence in stochastic control (Whittle 1981)). The uncertainty embedded in this *pseudo-abstraction* is due to the space discretization and stochastic disturbance alone. The pseudo-abstraction is a potential "best-case" for the current discretization if the learning error were driven to zero.

Example. To illustrate, consider a 2D linear system $x^+ = Ax + w$ where A is an unknown, stable matrix, and $w \sim \mathcal{N}(\mathbf{0}, 0.01I)$. Initially, 50 datapoints are sampled from the system to build an IMDP abstraction. Then, the pseudo-abstraction is constructed by using the learned model as the true model and treating e_L as zero everywhere. Another abstraction is made using 600 randomly-sampled points (large data) to approximate the zero-error verification results. All are verified against a specification that the system eventually reaches the region with label "G" while avoiding "O" (shown in Figure 1) with a probability greater than 90%. The results for pseudo- and large-data abstractions in Figures 1(b) and 1(c) are similar, which suggests 1) the pseudo-abstraction is a good zero-error approximation, and 2) some states need discretization-refinement to be definitively classified, and others can be improved using data-refinement. Establishing formal relationships using the pseudo-abstraction is ongoing work.

Data Refinement Strategies

Efficiently collecting data to realize the results of the pseudo-abstraction is an open question. As uniform sampling to drive e_L to zero everywhere is infeasible, we study several data-refinement strategies for efficient collection. Currently, we assume that we can collect data from any state. In the future, we will generalize to the online setting where system trajectories are considered. We present preliminary results for the following data-collection strategies:

- Random - sample from a random target state;
- Max Sigma - sample the state with the max σ;
- Trans. Width & σ - sort states according to the product of transition interval improvements using the pseudo-abstraction and max σ;
- Ver. Int. Width & σ - sample the state with the largest product of satisfaction interval width and max σ.

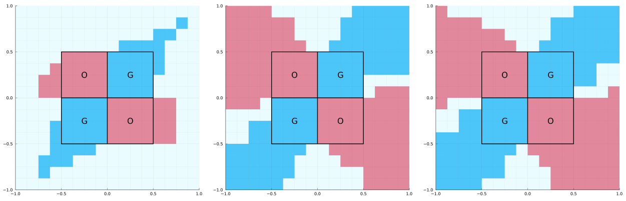

(a) Original (b) Pseudo (c) Large data

Figure 1: Verification results for the original (small-dataset), pseudo (error-free), and large-dataset abstractions. 🟦 are satisfying, 🟥 are violating, and □ are possibly satisfying.

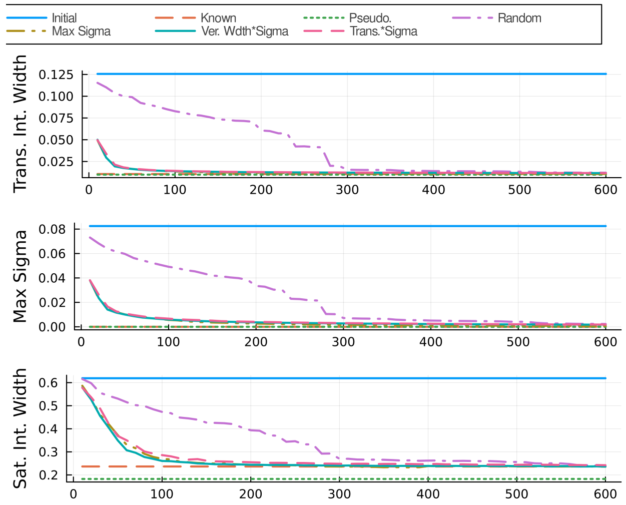

Figure 2: Comparison of metrics from employing different data-collection strategies.

Batches of 10 datapoints were sampled using each strategy, up to 600 additional datapoints. Figure 2 shows the results of three metrics: average transition interval width, average max σ in each state, and average satisfaction interval width. The strategy using the product of max verification width with uncertainty was the most effective in reducing the satisfaction interval widths. While the pseudo-abstraction estimates the best-possible results, using the transition interval improvements to guide data-refinement was not effective. These results encourage our further development of the pseudo-abstraction and efficient data-refinement strategies.

Conclusion

While IMDP abstractions are effective for verifying stochastic systems from data, when and where additional data is the most beneficial remains a challenge. Our initial work uses the error-free model called pseudo-abstraction to approximate the best-possible results and compares data-collection strategies that progress towards this outcome. Our preliminary results show that this method is effective. Ongoing work includes the formalization of the pseudo-abstraction and strategies, and adapting them for the online setting.

References

Badings, T.; Romao, L.; Abate, A.; and Jansen, N. 2023. Probabilities are not enough: Formal controller synthesis for stochastic dynamical models with epistemic uncertainty. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 14701–14710.

Fan, C.; Qi, B.; Mitra, S.; and Viswanathan, M. 2017. DryVR: Data-driven verification and compositional reasoning for automotive systems. In *International Conference on Computer Aided Verification*, 441–461. Springer.

Givan, R.; Leach, S.; and Dean, T. 2000. Bounded-parameter Markov decision processes. *Artificial Intelligence*, 122(1-2): 71–109.

Jackson, J.; Laurenti, L.; Frew, E.; and Lahijanian, M. 2020. Safety verification of unknown dynamical systems via gaussian process regression. In *2020 59th IEEE Conference on Decision and Control (CDC)*, 860–866. IEEE.

Jackson, J.; Laurenti, L.; Frew, E.; and Lahijanian, M. 2021. Synergistic offline-online control synthesis via local gaussian process regression. In *2021 60th IEEE Conference on Decision and Control (CDC)*, 2232–2239. IEEE.

Jiang, J.; Zhao, Y.; and Coogan, S. 2022. Safe Learning for Uncertainty-Aware Planning via Interval MDP Abstraction. *IEEE Control Systems Letters*, 6: 2641–2646.

Lahijanian, M.; Andersson, S. B.; and Belta, C. 2015. Formal verification and synthesis for discrete-time stochastic systems. *IEEE Transactions on Automatic Control*, 60(8): 2031–2045.

Martin, T.; Schön, T. B.; and Allgöwer, F. 2023. Guarantees for data-driven control of nonlinear systems using semidefinite programming: A survey. *Annual Reviews in Control*, 100911.

Reed, R.; Laurenti, L.; and Lahijanian, M. 2023. Promises of Deep Kernel Learning for Control Synthesis. *IEEE Control Systems Letters*.

Skovbekk, J.; Laurenti, L.; Frew, E.; and Lahijanian, M. 2023. Formal Abstraction of General Stochastic Systems via Noise Partitioning. *IEEE Control Systems Letters*.

Valmari, A. 1996. The state explosion problem. In *Advanced Course on Petri Nets*, 429–528. Springer.

Wells, A. M.; Lahijanian, M.; Kavraki, L. E.; and Vardi, M. Y. 2020. LTLf synthesis on probabilistic systems. *arXiv preprint arXiv:2009.10883*.

Whittle, P. 1981. Risk-sensitive linear/quadratic/Gaussian control. *Advances in Applied Probability*, 13(4): 764–777.

Winfield, A. F.; Michael, K.; Pitt, J.; and Evers, V. 2019. Machine ethics: The design and governance of ethical AI and autonomous systems [scanning the issue]. *Proceedings of the IEEE*, 107(3): 509–517.