Formal Abstraction of General Stochastic Systems via Noise Partitioning

John Skovbekk[1], *Student Member, IEEE*, Luca Laurenti[2], Eric Frew[1], *Member, IEEE*, and Morteza Lahijanian[1], *Member, IEEE*

***Abstract*— Verifying the performance of safety-critical, stochastic systems with complex noise distributions is difficult. We introduce a general procedure for the finite abstraction of nonlinear stochastic systems with non-standard (e.g., non-affine, non-symmetric, non-unimodal) noise distributions for verification purposes. The method uses a finite partitioning of the noise domain to construct an interval Markov chain (IMC) abstraction of the system via transition probability intervals. Noise partitioning allows for a general class of distributions and structures, including multiplicative and mixture models, and admits both known and data-driven systems. The partitions required for optimal transition bounds are specified for systems that are monotonic with respect to the noise, and explicit partitions are provided for affine and multiplicative structures. By the soundness of the abstraction procedure, verification on the IMC provides guarantees on the stochastic system against a temporal logic specification. In addition, we present a novel refinement-free algorithm that improves the verification results. Case studies on linear and nonlinear systems with non-Gaussian noise, including a data-driven example, demonstrate the generality and effectiveness of the method without introducing excessive conservatism.**

***Index Terms*— Autonomous systems, Markov processes, stochastic systems**

I. INTRODUCTION

THE deployment of autonomous systems for safety-critical applications, such as medical robotics and self-driving vehicles, requires diligent verification of their behavior. Such systems are inherently stochastic due to uncertainty in physical components (e.g., noise in sensors and actuators) or black-box software components. Formal methods provides rigorous techniques for verifying stochastic systems subject to temporal logic specifications [1], [2]. In particular, powerful model checking algorithms exist for finite-state Markov processes that can scale to large systems [1]. However, to apply them to continuous-space systems, finite abstractions with correctness guarantees are required [2], [3], which is difficult in both accuracy and scalability. For this reason, most existing work focuses on specific classes of stochastic systems often with strong assumptions on the dynamics or noise models [4]–[7], which we aim to relax in this work.

This work was supported in part by NSF grant 2039062.

[1]John Skovbekk, Eric Frew, and Morteza Lahijanian are with the Smead Department of Aerospace Engineering Sciences, CU Boulder, Boulder, CO, USA {firstname.lastname}@colorado.edu

[2]Luca Laurenti is with the Center for Systems and Control at TU Delft, Delft, Netherlands L.Laurenti@tudelft.nl

Uncertain Markov models, namely interval Markov chains (IMCs [8]) have proven to be effective abstraction models for stochastic systems [4]–[6], [9], [10]. Beyond capturing stochasticity, they also provide a means to incorporate other sources of uncertainty (e.g., discretization error), thereby facilitating correctness. Yet, the difficulty remains for general stochastic models due to the need to correctly compute stochastic transition kernels. Existing techniques rely on standard (unimodal, symmetric and zero-mean) or affine noise distributions [6], [10]–[12], linear systems [7], [9]. Additionally, stochastic systems may possess multiple sources of uncertainty, such as data-driven settings [5], [7], [10], [13], [14]. Thus, IMC abstraction approaches for nonlinear systems that admit a wider class of distributions and structures are necessary to lift these limitations.

Another difficulty facing abstraction is the state-explosion dilemma in higher dimensions. Common approaches to this problem are focused on parallelizing computation [15] and adaptive refinement [4], [16]. Despite these efforts, the state-explosion problem remains, and new ideas are needed for further mitigation. Specifically, using the continuous system in tandem with the abstraction to improve the verification without refinement is largely unexplored.

Contributions: We present an abstraction method for nonlinear stochastic systems with non-affine, non-standard noise that admits known and data-driven systems. Our method generalizes an approach for systems learned from data with affine, sub-Gaussian noise [5]. It is based on partitioning the noise domain to bound the transition kernel of the IMC, sidestepping the need to evaluate it. We show optimality criteria for the noise partitions for systems with noise monotonicity, and provide explicit partitions for affine and multiplicative structures. To help address the state-explosion problem, we also propose a refinement-free method to improve the verification results of an abstraction by using the continuous process. Finally, we demonstrate the efficacy of the method by verifying linear, nonlinear and data-driven systems without introducing excessive conservatism.

In summary, our contributions are (1) a procedure for constructing abstractions via noise partitioning (Theorem 1); (2) optimal noise partition sizes and values for a general class of distributions (Theorem 2); (3) a procedure to improve the verification of the abstraction without refinement (Algorithm 1), and (4) evaluations and applications to nonlinear systems with non-standard and multiplicative noise (Section VI).

II. PROBLEM FORMULATION

We first introduce the stochastic process and its finite abstraction, and then formulate two main problems.

A. Stochastic Process Model

Consider the following discrete-time stochastic process

$$\mathbf{x}(k+1) = f(\mathbf{x}(k), \mathbf{w}(k)), \tag{1}$$

where $\mathbf{x} \in \mathbb{R}^n$, $\mathbf{w} \in W \subseteq \mathbb{R}^{n_\mathbf{w}}$ is i.i.d. process noise sampled from distribution $p(\mathbf{w})$ with possibly bounded support, and $f : \mathbb{R}^n \times W \to \mathbb{R}^n$ is a possibly nonlinear function. Distribution $p(\mathbf{w})$ is allowed to be non-standard, i.e., non-uniform and non-symmetric. Let $X \subset \mathbb{R}^n$ be a Borel measurable set. The one-step transition kernel, which defines the probability of $\mathbf{x}(k+1) \in X$ given $\mathbf{x}(k) = x_k$ is

$$T(X \mid x_k) = \int_X f(x_k, \mathbf{w}(k)) p(\mathbf{w}) d\mathbf{w}. \tag{2}$$

The transition kernel T is the basis for probability measures of paths of System (1) [17], i.e., given an initial condition $\mathbf{x}(0) = x_0$, $\Pr\big(\mathbf{x}(0) \in X \mid x_0\big) = \mathbf{1}(x_0 \in X)$ and $\Pr\big(\mathbf{x}(k+1) \in X \mid x_k\big) = T(X \mid x_k)$, where $\mathbf{1}(\cdot)$ is the indicator function that returns 1 if the argument is true and 0 otherwise.

B. Interval Markov Chains

A finite abstraction of System (1) is often an interval Markov chain [2], which defines a space of Markov chains.

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

- Q *is a finite set of states;*
- $\check{P} : Q \times Q \to [0,1]$ *is the transition interval lower-bound function, where,* $\forall q,q' \in Q$, $\check{P}(q,q') \leq \Pr(q,q')$*;*
- $\hat{P} : Q \times Q \to [0,1]$ *is the transition interval upper-bound function, where* $\forall q,q' \in Q$, $\hat{P}(q,q') \geq \Pr(q,q')$*;*

It holds that, for every $q \in Q$, $\sum_{q' \in Q} \check{P}(q,q') \leq 1 \leq \sum_{q' \in Q} \hat{P}(q,q')$. Define the *adversary* $\gamma : Q \times Q \to [0,1]$ as a true transition probability function such that, for all $q, q' \in Q$, $\gamma(q,q') \in [\check{P}(q,q'), \hat{P}(q,q')]$ and $\sum_{q' \in Q} \gamma(q,q') = 1$. The set of all adversaries is denoted by Γ. Under adversary γ, the IMC reduces to a Markov chain with a well-defined probability measure over its paths.

Consider a path property ϕ. The probability that all paths initiated at $q \in Q$ satisfy ϕ is denoted by $\Pr(q \models \phi)$. When ϕ is expressed in probabilistic computation tree logic (PCTL) or linear temporal logic (LTL) [1], $\Pr(q \models \phi)$ is equivalent to the reachability probability on an IMC that composes $\mathcal{I}$ with ϕ. W.L.O.G., let $Q_\phi \subseteq Q$ be the set of states, reaching which satisfies ϕ. While the exact value of $\Pr(q \models \phi)$ cannot be computed, it can be bounded, i.e., $\Pr(q \models \phi) \in [\check{p}(q), \hat{p}(q)]$, using dynamic programming [4]. For the lower bound,

$$\check{p}^0(q) = \mathbf{1}(q \in Q_\phi),\ \check{p}^k(q) = \min_{\gamma \in \Gamma} \sum_{q' \in Q} \gamma(q,q') \check{p}^{k-1}(q'). \tag{3}$$

The upper bound $\hat{p}$ is computed by replacing the $\min$ with $\max$ operator and $\check{p}$ with $\hat{p}$. The computation of the satisfaction bounds $\check{p}(q)$ and $\hat{p}(q)$ for all $q \in Q$ is called the IMC verification procedure.

C. Problem Statements

Verifying System (1) against ϕ can be performed by discretizing the state space of (1) to build an IMC abstraction $\mathcal{I}$ that soundly models (1), and then verifying $\mathcal{I}$ against ϕ. Given a compact set $X \subset \mathbb{R}^n$, we let Q_X denote a finite partitioning of X, with no preference on its inception. With an abuse of notation, $q \in Q_X$ is both one of these partitions and an IMC state. The verification results can be extended to (1), i.e., for every $x \in q$, $\Pr(x \models \phi) \in [\check{p}(q), \hat{p}(q)]$, if the abstraction satisfies the soundness definition below as shown in [5, Theorem 2].

Definition 2 (Abstraction Soundness)**.** *An IMC abstraction* $\mathcal{I}$ *is sound with respect to System* (1) *if, for all* $x \in q$, $\check{P}(q,q') \leq T(q' \mid x) \leq \hat{P}(q,q')$ *holds for all* $q, q' \in Q_X$*.*

To satisfy this definition, we assume that one of the requirements of ϕ is to remain within a bounded (safe) set $X \subset \mathbb{R}^n$ and refer to $\mathbb{R}^n \setminus X$ as an unsafe set.

Existing methods for IMC abstraction of stochastic systems are largely limited to simple dynamics – affine in noise with unimodal or symmetric distributions, or linear dynamics. The first problem considered here aims to establish a method that jointly addresses these limitations.

Problem 1 (Abstraction Construction)**.** *Construct a sound IMC abstraction for System* (1) *with a nonlinear* f *and non-affine and non-standard* $p(\mathbf{w})$*.*

In Section III, we propose a method that partitions the domain of $p(\mathbf{w})$ to construct the transition bounds of the IMC which are valid for arbitrary distributions. Solving this problem allows the application of IMC abstractions to a wider class of systems, including data-driven systems.

The conventional approach to improving the satisfaction intervals of an IMC is to refine the discretization Q_X, which contributes to the state-explosion dilemma. The next problem aims to improve the intervals on the same discretization Q_X by leveraging the model of (1).

Problem 2 (Verification Improvement)**.** *Given abstraction* $\mathcal{I}$ *of System* (1)*, reduce the verification error* $\hat{p}(q) - \check{p}(q)$ *for all* $q \in Q_X$ *without refining* Q_X*.*

In Section V, we propose an approach based on clustering states in Q_X that uses the structure of the transition bounds and (1) to reduce the gap between $\hat{p}$ and $\check{p}$.

Remark 1. *While we focus on IMC abstractions, the results are trivially applied to interval Markov decision process (IMDP) abstraction methods via concatenation of IMCs.*

III. ABSTRACTION VIA NOISE PARTITIONS

The IMC abstraction for System (1) involves discretizing the continuous state-space and computing transition probability bounds between the resulting states.

A. State Discretization

Constructing a finite-state abstraction for System (1) requires a bounded subset $X \subset \mathbb{R}^n$. The abstraction is sound on

X, but not the entire state-space as discussed in Definition 2. X is partitioned into a finite set of bounded and convex regions Q_X, which implies, for every $q, q' \in Q$, $q \cap q'$ has zero measure and $\bigcup_{q \in Q_X} q = X$. Let $q_{\neg X}$ represent the remainder (i.e., unsafe set) $\mathbb{R}^n \setminus X$. Then, the complete state set of the IMC is $Q = Q_X \cup \{q_{\neg X}\}$. The next IMC abstraction step computes the transition bounds between states.

B. Transition Bounds with Noise Partitioning

The definition and computation of the transition bound functions $\check{P}, \hat{P}$ begins with states in Q_X. The transitions to $q_{\neg X}$ is a modified case. The connection between System (1) and the abstraction arises from the transition kernel T in (2) over IMC states. From a given $x \in q$, the transition kernel to q' is $T(q' \mid x)$. Finding bounds on the kernel amounts to searching over all $x \in q$, i.e., $\min_{x \in q} T(q' \mid x)$, and $\max_{x \in q} T(q' \mid x)$. To satisfy Definition 2, $\check{P}(q, q')$ and $\hat{P}(q, q')$ must bound these extrema. For tractable evaluation of T in (2) with non-standard distributions, the probability measure of $\mathbf{w}$ is evaluated over partitions of its domain W.

Definition 3 (Noise Partition)**.** *A noise partition set C is a measure-preserving discretization of W, i.e., $\bigcup_{c \in C} c = W$ and $\forall c \in C$, $\sum_{c \in C} \int_c p(\mathbf{w}) d\mathbf{w} = \int p(\mathbf{w}) d\mathbf{w} = 1$.*

For brevity, c is used in place of $\mathbf{w}(k) \in c$, and its probability is $\Pr(c) = \int_c p(\mathbf{w}) d\mathbf{w}$. For a given $c \in C$, the posterior of region q is $Post(q, c) = \{f(x, w) \mid x \in q, w \in c\}$. The following theorem bounds the transition kernel.

Theorem 1. *Let $q, q' \in Q_X$ and C be a partition of W according to Definition 3. Then, the transition kernel is lower- and upper-bounded, respectively, by*

$$\min_{x_k \in q} T(q' \mid x_k) \geq \sum_{c \in C} \mathbf{1}(Post(q, c) \subseteq q') \Pr(c) \tag{4a}$$

$$\max_{x_k \in q} T(q' \mid x_k) \leq \sum_{c \in C} \mathbf{1}(Post(q, c) \cap q' = \emptyset) \Pr(c) \tag{4b}$$

Proof. We begin with finding the upper bound. Using T and finding the maximizing point,

$$\max_{x_k \in q} T(q' \mid x_k) = \max_{x_k \in q} \int \mathbf{1}(\mathbf{x}(k+1) \in q' \mid x_k, w_k) p(\mathbf{w}) d\mathbf{w} \tag{5}$$

The integral is split according to the partitions in C,

$$(5) = \max_{x_k \in q} \sum_{c \in C} \int_c \mathbf{1}(\mathbf{x}(k+1) \in q' \mid x_k\, w_k) p(\mathbf{w}) d\mathbf{w}, \tag{6}$$

which maintains equality due to the linearity of the integral. The indicator function is upper-bounded by the existence of a point in the intersection of $Post(q, c)$ with q',

$$(6) \leq \sum_{c \in C} \mathbf{1}(Post(q, c) \cap q' \neq \emptyset) \Pr(c),$$

where the max operator is dropped, as x_k is subsumed by q. The lower-bound is similar, instead doing under-approximation by checking if $Post(q, c) \subseteq q'$. □

The transition bounds found using Theorem 1 require two components: $Post(q, c)$ and $\Pr(c)$. Note that for the bounds in (4a)-(4b), an over-approximation of $Post(q, c)$ can be used, which can be obtained for nonlinear systems using local linear bounds of $f(\mathbf{x}(k), \mathbf{w}(k))$ [18], [19], discretization with Taylor model flowpipes [20], or mixed-monotone maps [21] depending on the knowledge of System (1). $\Pr(c)$ can be computed analytically for distribution-dependent soundness guarantees, or statistically for sampling-based guarantees [7]. The next section discusses how partitions are selected to optimize the bounds in Theorem 1.

To complete the abstraction, transitions to the unsafe state $q_{\neg X}$ are defined using the following corollary.

Corollary 1 (Unsafe State Transitions)**.** *For every state $q \in Q_X$, the transition bounds to $q_{\neg X}$ are $\check{P}(q, q_{\neg X}) = 1 - \max_{x_k \in q} T(X \mid x_k)$ and $\hat{P}(q, q_{\neg X}) = 1 - \min_{x_k \in q} T(X \mid x_k)$. Additionally, the transition bounds between $q_{\neg X}$ and itself are $\check{P}(q_{\neg X}, q_{\neg X}) = \hat{P}(q_{\neg X}, q_{\neg X}) = 1$.*

Remark 2. *Theorem 1 can be applied to general (non-probabilistic) uncertainty sets by interpreting $\Pr(c)$ as a deterministic indicator function. For example, for the bounded uncertainty set W, choose $c = W$ so $\Pr(c) = 1$, and $\Pr(c') = 0$ for every other $c' \in C$. Effectively, using Theorem 1 in this case results in a non-deterministic transition system.*

IV. Optimal Partitions

The transition bounds in Theorem 1 return valid bounds for any choice of partition, and C can differ between (4a) and (4b). However, haphazard partitions can result in the trivial transition probability interval $[0, 1]$. The optimal noise partitions minimize the distance between the transition bounds, i.e., given $q, q' \in Q_X$,

$$\underline{C}^* = \arg\max_{C} \sum_{c \in C} \mathbf{1}(Post(q, c) \subseteq q') \Pr(c),$$

$$\overline{C}^* = \arg\min_{C} \sum_{c \in C} \mathbf{1}(Post(q, c) \cap q' = \emptyset) \Pr(c), \tag{7a}$$

Hence, noise partitions can be selected to optimize the transition bounds for each pair (q, q') independently. To begin the analysis on these partitions, we assume component-wise noise as defined below.

Definition 4 (Component-wise Noise)**.** *For $i \in \{1, \ldots, n\}$, let $M^i \in \{0, 1\}^{n \times n}$ be a matrix whose i, i element is one and all the other elements are zeros. Then, noise $\mathbf{w}(k) \in W \subset \mathbb{R}^n$ is called* component-wise *if $M^i f(\mathbf{x}(k), \mathbf{w}(k)) = f(\mathbf{x}(k), M^i \mathbf{w}(k))$.*

In other words, the noise vector shares the size of $\mathbf{x}(k)$, and each component $\mathbf{w}^i(k)$ only affects $\mathbf{x}^i(k+1)$, which admits (but is not limited to) affine and multiplicative noise (see Example 1 below). Definition 4 does not preclude the noise from being correlated. Assuming the noise satisfies Definition 4, the next step is to explore how $Post(q, c)$ changes with variations in c. If increasing $\mathbf{w}^i_k$ consistently increases (or decreases) $\mathbf{x}^i(k+1)$, it can lead to partitions C that induce non-empty intersections between $Post(q, c)$ and q'. This occurs if the system is monotone with respect to the noise, which is defined below.

Definition 5 (Noise Monotonicity)**.** *Monotonicity is the condition that, for all scalars $a, b \in \mathbb{R}$, $a > b$ implies $f^i(\cdot, a) - f^i(\cdot, b)$ has the same sign. System (1) is* monotonic *with respect to $\mathbf{w}(k)$ if each f^i is monotonic with respect to $\mathbf{w}^i(k)$.*

Example 1. *Consider the system $\mathbf{x}(k+1) = f(x(k)) \odot \mathbf{w}(k)$, where each $\mathbf{w}^i(k) \geq 0$ and $\odot$ is the element-wise product. Then the noise acts component-wise, and the system is monotonic with respect to $\mathbf{w}(k)$.*

Hitherto, we have made no assumptions about the convexity of $Post(q, c)$. Let the i component of a set in $\mathbb{R}^n$ refer to its projection on the i-th unit axis. The following theorem discusses non-convexity in terms of discontinuities (or holes) in each component of $Post(q, c)$. The theorem bounds the sizes of $\underline{C}^*$ and $\overline{C}^*$ for a system with monotonic noise.

Theorem 2 (Partition Size)**.** *Let $q, q' \subset \mathbb{R}^n$ be bounded and convex, and d be the largest number of discontinuities in each component of $Post(q, c)$. If System (1) is monotonic with respect to component-wise, uncorrelated noise $\mathbf{w}(k)$, then $|\overline{C}^*|$ and $|\underline{C}^*|$ are at most $(3 + 2d)n$.*

Proof. The proof is provided for the upper-bound partition (7a). The lower-bound is the same but uses $Post(q, c) \subseteq q'$.

$\overline{C}^*$ is found my choosing the constraint set that satisfies to (7a). As the noise is uncorrelated, it is sufficient to minimize the area intersection of $Post(q, c)$ with q' to minimize (7a). Let q, q' be convex and bounded, let $Post(q, c)$ contain at most d discontinuities for any choice of c, and let System (1) be monotonic with respect to $\mathbf{w}(k)$.

First, consider $d = 0$, so $Post(q, c)$ is convex for a given c in all components. Then, $Post(q, c) \cap q'$ must be convex. As q' is bounded, $Post(q, c) \cap q'$ is also bounded. For each i, due to the monotonicity of System (1), at most 3 partitions of W^i are needed to induce the minimum intersection with q'^i due to the convexity and boundedness of q'^i. Thus, $\overline{C}^*$ consists of $3n$ partitions at most when $d = 0$.

Next, consider $d > 0$ for a component of $Post(q, c)$ and begin with $\overline{C}^*$ as found above. The intersection $Post(q, c) \cap q'$ is possibly non-convex due the projection of discontinuities of $Post^i(q, c)$. For each discontinuity, only two additional partitions are needed to induce the minimum intersection with q'^i due to the monotonicity of System (1). This is repeated for each component in $[1, n]$ for the result $|\overline{C}^*| \leq (3 + 2d)n$. Repeating this procedure for the lower bound yields the same number of partitions in $\underline{C}^*$. □

Theorem 2 shows that the sizes of $\underline{C}^*$ and $\overline{C}^*$ are bounded, but it leaves them unspecified. The following corollaries specify the partitions for affine and multiplicative noise in the case $Post(q, c)$ is convex (hence $d = 0$). To facilitate this, let $Post_f(q) = \{f(x) \mid x \in q\}$ be the f-dependent posterior. The corollaries partition W^i into three intervals, $\overline{C}^* = \{[-\infty, \epsilon_1], [\epsilon_1, \epsilon_2], [\epsilon_2, \infty]\}$ and $\underline{C}^*\{[-\infty, \epsilon_3], [\epsilon_3, \epsilon_4], [\epsilon_4, \infty]\}$, according to Theorem 2.

Corollary 2 (Partitioning for Affine Noise)**.** *Assume System (1) satisfies the requirements of Theorem 2 and has affine noise, i.e. $f(\mathbf{x}(k)) + \mathbf{w}(k)$. Let the i-th component endpoints of the target region q' be $a_{q'}, b_{q'}$, and a_q, b_q for the f-dependent posterior, and let $l = b_q - a_q$. Let* ***C1*** *be $l > b_{q'} - a_{q'}$,* ***C2*** *be $b_q > b_{q'}$, and* ***C3*** *be $a_q < a_{q'}$. Then $\epsilon_1 = a_{q'} - b_q$, $\epsilon_2 = b_{q'} - a_q$ in all cases. For the lower-bound,*

$$\epsilon_3 = \begin{cases} 0 & \text{if } \mathbf{C1} \\ a_{q'} - b_q + l & \text{if } \mathbf{C2} \\ a_{q'} - a_q & \text{o.w.} \end{cases}, \epsilon_4 = \begin{cases} 0 & \text{if } \mathbf{C1} \\ b_{q'} - l - a_q & \text{if } \mathbf{C3} \\ b_{q'} - b_q & \text{o.w.} \end{cases}$$

Proof. The proof uses the relative positions in Figure 1 to find partitions that minimize and maximize the intersection between the posterior and target region. For the upper bound (4b), $\mathbf{w}(k) < a_{q'} - b_q$ or $\mathbf{w}(k) > b_{q'} - a_q$ is the largest interval that ensures no intersection can occur with q'. For the lower bound (4a), **C1** occurs when the posterior is larger than the target, so no partition of noise can cause intersection. For **C2**, $a_{q'} + l - b_q < \mathbf{w}(k) \leq b_{q'} - b_q$ is the largest interval that ensures an intersection. **C3** is the same as **C2** with mirrored positions. When no cases are true, $Post_f(q) \subseteq q'$, and $a_{q'} - a_q \leq \mathbf{w}(k) \leq b_{q'} - b_q$ ensures intersection. □

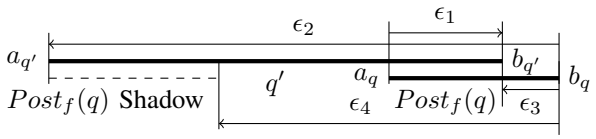

Fig. 1: Endpoints $a_{q'}, b_{q'}, a_q, b_q$ and distances between a component of q' and $Post_f(q)$.

Corollary 3 (Partitioning for Multiplicative Noise)**.** *Assume System (1) satisfies the requirements of Theorem 2 and has (w.l.o.g.) positive multiplicative noise, i.e., Example 1. Let the component endpoints be the same as Corollary 2, and let them all be positive. Then, $\epsilon_1 = a_{q'}/b_q$, $\epsilon_2 = b_{q'}/a_q$, $\epsilon_3 = a_{q'}/a_q$, and $\epsilon_4 = b_{q'}/b_q$.*

Proof. The proof is similar to that of Corollary 2. The lower bound is maximized when both $b_q \mathbf{w}(k) < b_{q'}$ and $a_q \mathbf{w}(k) > a_{q'}$, so $a_{q'}/a_q < \mathbf{w}(k) \leq b_{q'}/b_q$. Note that if $a_{q'}/a_q > b_{q'}/b_q$, then the CDF evaluates to zero, so the partition set is trivial. Likewise, the upper bound is minimized when $a_q \mathbf{w}(k) > b_{q'}$ or $b_q \mathbf{w}(k) < a_{q'}$, so $\mathbf{w}(k) < a_{q'}/b_q$ or $\mathbf{w}(k) > b_{q'}/a_q$. □

These corollaries can be extended in the non-convex case ($d > 0$) if the f-dependent posterior is readily available. The case studies in Section VI show that the partitioning approach finds accurate abstractions for systems with non-standard noise when compared to a specialized method.

V. STATE CLUSTERING

Improving the satisfaction intervals of the IMC directly impacts the guarantees on System (1), but relying solely on refining the space discretization can lead to an explosion in the number of states. Here, we propose a novel method based on clustering the states of the IMC to improve the satisfaction intervals without refinement to solve Problem 2.

Consider state $q \in Q_X$ and its set of successor states Q'. By the structure of (4a), $\check{P}(q, q')$ increases with the size of q' as it depends on $Post(q, c) \subset q'$ being true. Algorithm 1 is based on this principle. The sorting of Q_X on Line 1 makes the algorithm start with states with large $\check{p}$, as its successor states have larger $\check{p}$, making improvement more likely. Then, a subset of Q' is chosen to cluster into a

Algorithm 1 Clustering-based IMC improvement

Require: IMC $\mathcal{I}$, verification results $\check{p}$, $\hat{p}$

$Q_X \leftarrow$ sort by $\check{p}(q)$ in descending order
for each $q \in Q_X$ **do**
  $\tilde{q} \leftarrow$ cluster $\tilde{Q} \subset Q'$ into a single state
  $\check{p}(\tilde{q}) \leftarrow \min_{q' \in \tilde{Q}} \check{p}(q')$
  $\check{P}(q,\tilde{q}), \hat{P}(q,\tilde{q}) \leftarrow$ Theorem 1
  $\check{p}_{new}(q) \leftarrow \min_{\gamma \in \Gamma} \sum_{q' \in Q' \setminus \tilde{Q}} \gamma(q,q')\check{p}(q') + \gamma(q,\tilde{q})\check{p}(\tilde{q})$
  $\hat{p}_{new}(q) \leftarrow$ similarly with (3)
  if $\check{p}_{new}(q) > \check{p}(q)$ or $\hat{p}_{new}(q) < \hat{p}(q)$ **then**
    Save these values to $\check{p}$, $\hat{p}$
return Improved intervals $\check{p}$, $\hat{p}$

single state $\tilde{q}$. The transition interval to $\tilde{q}$ is computed using Theorem 1, and the satisfaction intervals are recalculated. Clustering any states from Q' improves the transition interval, but it must be balanced with $\check{p}(q')$ for effective enhancement. We leave the question of the optimal clustering choice for future work. However, our application of Algorithm 1 to a data-driven example reveals that even sub-optimal clustering yields improvements to the verification results.

VI. Evaluations

We evaluate the proposed methods on linear, non-linear, and data-driven systems using the PCTL specification ϕ that states "the probability of reaching goal G within k steps (infinity unless otherwise noted) while avoiding obstacles O is ≥ 0.9." Figures show states satisfying ($\models \phi$) or violating ($\not\models \phi$) the specification, and possibly either ($?\phi$).

1) Linear System Comparison: First, we compare our set-based transition interval calculations against the direct point-search method for specific noise proposed in [21] over the same state discretization. The system is linear with additive truncated Gaussian noise bounded on $[-0.4, 0.4]$ from [21]. As shown in Figure 2, the classification results are nearly identical with average and maximum differences of 8×10^{-4} and 0.02, respectively, in the lower-bound satisfaction probability. Computation time for discretization, verification, and refinement is approximately 90 seconds. In this example, the set-based criteria of Theorem 1 with the optimality of Corollary 2 are sufficient to provide an accurate abstraction. In addition, we highlight the shortcomings of using Theorem 1 without Corollary 2. The hatching on Figure 2a indicates the verification results using fixed noise partitions $[-0.1, 0.1]$ to construct C for both components, which is a shadow compared to using the optimal partitions.

2) Multiplicative Noise: Next, we consider a system with multiplicative noise, which existing abstraction approaches cannot explicitly handle, to the best of our knowledge. The dynamics are $\mathbf{x}(k+1) = A\mathbf{x}(k) \odot \mathbf{w}(k)$, where the 1st and 2nd rows of A are $(0.7, 0.1)$ and $(0.1, 0.8)$. Each noise component is from a truncated Normal distribution with support $[0.9, 1.1]$, mean 1, and variance 0.1. Figure 3a shows verification results using Corollary 3, which took 2 min. to compute. Figure 3b shows the mean and 1-sigma bounds of 1000 sampled paths. These validations are consistent with the classifications, as

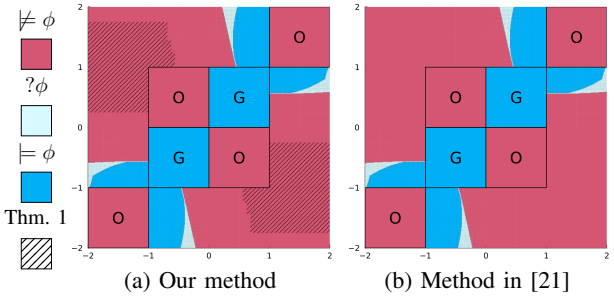

(a) Our method (b) Method in [21]

Fig. 2: Comparison of verification results.

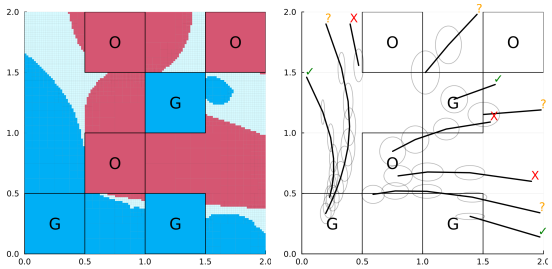

(a) Verification (b) Sampled Trajectories

Fig. 3: Verification of the system with multiplicative noise.

some pass too close to O to make a conclusion either way, and others consistently reach G or O.

3) Data-driven Verification: The linear system from the comparison example with $\mathbf{w}^i(k) \sim \mathcal{N}(0, 0.1^2)$ on each component is learned via Gaussian process (GP) regression with 200 data points. The transition bounds account for both the resulting uncertainty from the learning procedure and the inherent system noise, as the resulting noise distribution is a sum of normals. The efficacy of Algorithm 1 is demonstrated on the initial discretization by clustering the states that intersect with the f-dependent posterior of q, which is found similar to [5]. Figures 4a and 4b show additional satisfying and violating states are identified without refinement. The lower-bound satisfaction was improved in 8 states, with a 10% (absolute) average increase in the lower bound. Figures 4c and 4d show the similarity of classifications between the known and learned system after refining the abstraction. This shows the efficacy of the method in the data-driven setting.

4) Duffing Oscillator: The nonlinear Duffing oscillator has complex motion and chaotic behavior with continuous-time dynamics $\ddot{x} + \delta\dot{x} + \alpha x + \beta x^3 = \gamma\cos(\omega t)$, where, $\delta = 0.3$, $\alpha = -1.0$, $\beta = 1.0$, $\gamma = 0.37$, and $\omega = 1.2$. This system is discretized over the time-span $[0, 0.5]$, after which the forcing function is reset and noise $\mathbf{w}^i(k) \sim \mathcal{N}(0.1, 0.01^2)$ is drawn. Taylor models were used to over-approximate the $Post$ of each discrete region [22]. The abstraction and verification for $k = 10$, shown in Figure 5, took 4.5 hours to compute. The paths are the means of 1000 samples with 2-sigma confidence bounds at each point and the initial classification. The validating trajectories provide insight to why initial states bear their classification.

5) Dubin's Aircraft with Mixture: A 3D discrete-time Dubin's car model [23] with a constant right-turn control input is verified with additive mixture noise consisting of $\textsc{Uniform}(-0.05, -0.01)$ and $\textsc{Uniform}(0.0, 0.04)$, each with a 50% weighting, on each state component x, y, and heading

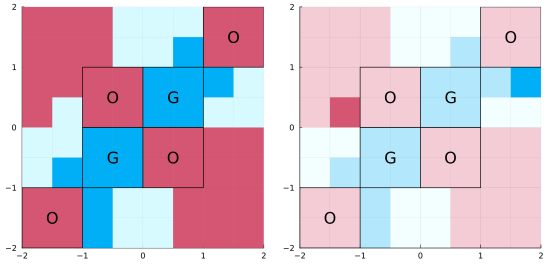

(a) Initial (b) Cluster (1 step)

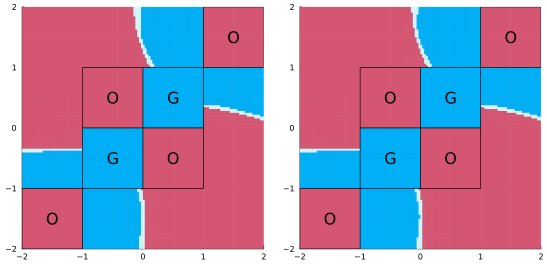

(c) Final (GP) (d) Final (Known)

Fig. 4: Verification of the learned linear system.

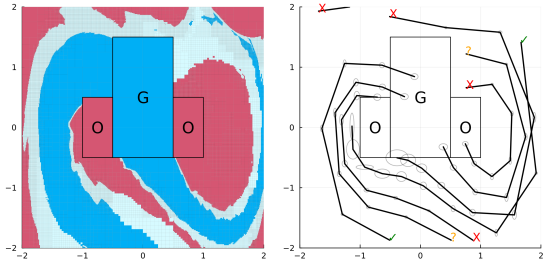

(a) Classifications (b) Sampled Trajectories

Fig. 5: Verification of the noisy Duffing oscillator.

θ. Figure 6 shows results in 3D and a 2D slice. Initial states are identified that are guaranteed to make the turn safely, or fail to meet the minimum safety threshold. This shows the efficacy of the method in the verification of autonomous system with non-standard distributions.

VII. CONCLUSIONS

We present an IMC abstraction method for nonlinear stochastic systems by partitioning the noise domain, and a refinement-free approach to improve IMC verification. This procedure admits systems with non-affine and non-standard noise distributions, and data-driven systems. Future work will add measurement models, generalize the optimal partitioning beyond component-wise noise and convex posteriors, and improve the efficacy of the clustering procedure.

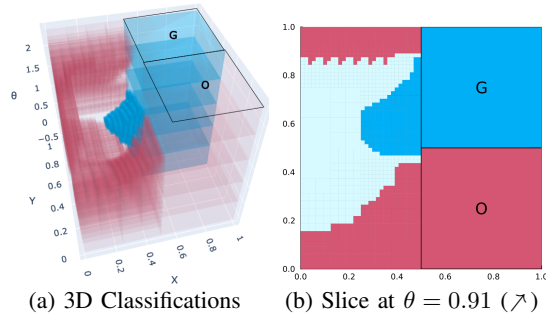

(a) 3D Classifications (b) Slice at $\theta = 0.91$ ($\nearrow$)

Fig. 6: Verification of the constant-turn Dubin's car system.

REFERENCES

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

[2] A. Lavaei, S. Soudjani, A. Abate, and M. Zamani, "Automated verification and synthesis of stochastic hybrid systems: A survey," *Automatica*, vol. 146, p. 110617, 2022.

[3] R. Alur, T. Henzinger, G. Lafferriere, and G. Pappas, "Discrete abstractions of hybrid systems," *Proceedings of the IEEE*, vol. 88, no. 7, pp. 971–984, 2000.

[4] M. Lahijanian, S. B. Andersson, and C. Belta, "Formal verification and synthesis for discrete-time stochastic systems," *IEEE Transactions on Automatic Control*, vol. 60, no. 8, pp. 2031–2045, 2015.

[5] J. Jackson, L. Laurenti, E. Frew, and M. Lahijanian, "Strategy synthesis for partially-known switched stochastic systems," in *Int. Conf. on Hybrid Systems: Computation and Control*, 2021.

[6] M. Dutreix, J. Huh, and S. Coogan, "Abstraction-based synthesis for stochastic systems with omega-regular objectives," *Nonlinear Analysis: Hybrid Systems*, vol. 45, p. 101204, 2022.

[7] T. Badings, L. Romao, A. Abate, D. Parker, H. A. Poonawala, M. Stoelinga, and N. Jansen, "Robust control for dynamical systems with non-gaussian noise via formal abstractions," *Journal of Artificial Intelligence Research*, vol. 76, pp. 341–391, 2023.

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

[9] N. Cauchi, L. Laurenti, M. Lahijanian, A. Abate, M. Kwiatkowska, and L. Cardelli, "Efficiency through uncertainty: Scalable formal synthesis for stochastic hybrid systems," in *ACM Int. Conf. on hybrid systems: computation and control*, 2019, pp. 240–251.

[10] J. Jiang, Y. Zhao, and S. Coogan, "Safe learning for uncertainty-aware planning via interval mdp abstraction," *IEEE Control Systems Letters*, vol. 6, pp. 2641–2646, 2022.

[11] J. Jackson, L. Laurenti, E. Frew, and M. Lahijanian, "Formal verification of unknown dynamical systems via gaussian process regression," *arXiv preprint arXiv:2201.00655*, 2021.

[12] B. C. van Huijgevoort, S. Weiland, and S. Haesaert, "Temporal logic control of nonlinear stochastic systems using a piecewise-affine abstraction," *IEEE Control Sys. Letters*, vol. 7, pp. 1039–1044, 2023.

[13] K. Hashimoto, A. Saoud, M. Kishida, T. Ushio, and D. V. Dimarogonas, "Learning-based symbolic abstractions for nonlinear control systems," *Automatica*, vol. 146, p. 110646, 2022.

[14] S. Adams, M. Lahijanian, and L. Laurenti, "Formal control synthesis for stochastic neural network dynamic models," *IEEE Control Systems Letters*, vol. 6, pp. 2858–2863, 2022.

[15] S. E. Z. Soudjani, C. Gevaerts, and A. Abate, "Faust: F ormal a bstractions of u ncountable-st ate st ochastic processes," in *Tools and Alg. for the Const. and Analys. of Sys.* Springer, 2015, pp. 272–286.

[16] S. Esmaeil Zadeh Soudjani and A. Abate, "Adaptive and sequential gridding procedures for the abstraction and verification of stochastic processes," *SIAM Journal on Applied Dynamical Systems*, vol. 12, no. 2, pp. 921–956, 2013.

[17] A. Klenke, *Probability Measures on Product Spaces*. Springer Science & Business Media, 2013.

[18] Z. Jin, Q. Shen, and S. Z. Yong, "Mesh-based piecewise affine abstraction with polytopic partitions for nonlinear systems," *IEEE Control Systems Letters*, vol. 5, no. 5, pp. 1543–1548, 2021.

[19] F. B. Mathiesen, S. C. Calvert, and L. Laurenti, "Safety certification for stochastic systems via neural barrier functions," *IEEE Control Systems Letters*, vol. 7, pp. 973–978, 2022.

[20] X. Chen, E. Ábrahám, and S. Sankaranarayanan, "Flow*: An analyzer for non-linear hybrid systems," in *Computer Aided Verification*. Springer, 2013, pp. 258–263.

[21] M. Dutreix and S. Coogan, "Efficient verification for stochastic mixed monotone systems," in *2018 ACM/IEEE 9th International Conference on Cyber-Physical Systems (ICCPS)*. IEEE, 2018, pp. 150–161.

[22] S. Bogomolov, M. Forets, G. Frehse, K. Potomkin, and C. Schilling, "Juliareach: a toolbox for set-based reachability," in *ACM Int. Conf. on Hybrid Systems: Computation and Control*, 2019, pp. 39–44.

[23] M. Owen, R. W. Beard, and T. W. McLain, *Implementing Dubins Airplane Paths on Fixed-Wing UAVs**. Dordrecht: Springer Netherlands, 2015, pp. 1677–1701.