

Scalable control synthesis for stochastic systems via structural IMDP abstractions

Mathiesen, Frederik Baymller; Haesaert, S.; Laurenti, L.

DOI

[10.1145/3716863.3718031](https://doi.org/10.1145/3716863.3718031)

Licence

CC BY

Publication date

2025

Document Version

Final published version

Published in

HSCC '25: Proceedings of the 28th ACM International Conference on Hybrid Systems: Computation and Control

Citation (APA)

Mathiesen, F. B., Haesaert, S., & Laurenti, L. (2025). Scalable control synthesis for stochastic systems via structural IMDP abstractions. In *HSCC '25: Proceedings of the 28th ACM International Conference on Hybrid Systems: Computation and Control* Article 14 ACM. <https://doi.org/10.1145/3716863.3718031>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.



Scalable control synthesis for stochastic systems via structural IMDP abstractions

Frederik Baymler Mathiesen
 f.b.mathiesen@tudelft.nl
 Delft University of Technology
 Delft, The Netherlands

Sofie Haesaert
 s.haesaert@tue.nl
 Eindhoven University of Technology
 Eindhoven, The Netherlands

Luca Laurenti*
 l.laurenti@tudelft.nl
 Delft University of Technology
 Delft, The Netherlands

Abstract

This paper introduces a novel abstraction-based framework for controller synthesis of nonlinear discrete-time stochastic systems. The focus is on probabilistic reach-avoid specifications. The framework is based on abstracting a stochastic system into a new class of robust Markov models, called orthogonally decoupled Interval Markov Decision Processes (odIMDPs). Specifically, an odIMDPs is a class of robust Markov processes, where the transition probabilities between each pair of states are uncertain and have the product form. We show that such a specific form in the transition probabilities allows one to build compositional abstractions of stochastic systems that, for each state, are only required to store the marginal probability bounds of the original system. This leads to improved memory complexity for our approach compared to commonly employed abstraction-based approaches. Furthermore, we show that an optimal control strategy for a odIMDPs can be computed by solving a set of linear problems. When the resulting strategy is mapped back to the original system, it is guaranteed to lead to reduced conservatism compared to existing approaches. To test our theoretical framework, we perform an extensive empirical comparison of our methods against Interval Markov Decision Process- and Markov Decision Process-based approaches on various benchmarks including 7D systems. Our empirical analysis shows that our approach substantially outperforms state-of-the-art approaches in terms of both memory requirements and the conservatism of the results.

CCS Concepts

- Mathematics of computing → Markov processes; • Theory of computation → Stochastic control and optimization; • Computer systems organization → Embedded systems.

Keywords

Markov Decision Processes, Robust Value Iteration, Reachability, Control Synthesis, Verification.

ACM Reference Format:

Frederik Baymler Mathiesen, Sofie Haesaert, and Luca Laurenti. 2025. Scalable control synthesis for stochastic systems via structural IMDP abstractions. In *28th ACM International Conference on Hybrid Systems: Computation and Control (HSCC '25), May 6–9, 2025, Irvine, CA, USA*. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3716863.3718031>

*L.L. is partially supported by the NWO (Grant OCENW.M.22.056).



This work is licensed under a Creative Commons Attribution 4.0 International License.
HSCC '25, Irvine, CA, USA
 © 2025 Copyright held by the owner/author(s).
 ACM ISBN 979-8-4007-1504-4/2025/05
<https://doi.org/10.1145/3716863.3718031>

1 Introduction

Modern cyber-physical systems, such as autonomous vehicles or robotic systems, are often modeled as stochastic dynamical systems [10, 32], where nonlinear dynamics transform stochastic quantities over time. Consequently, to deploy cyber-physical systems in safety-critical applications, it has become of paramount importance to develop techniques for controller synthesis and verification of nonlinear stochastic dynamical systems with guarantees of performance and correctness [13]. However, unfortunately, existing techniques still lack scalability and are commonly limited to low-dimensional systems [26, 28].

Abstraction-based methods arguably represent the state-of-the-art for formal analysis of nonlinear stochastic systems [11, 28, 31, 49]. Abstraction-based methods approximate a given complex system (the “concrete” system) with a simpler one (the “abstraction”) for the purpose of performing analysis while maintaining a (bi)simulation relation between the concrete and abstract systems [11, 31, 50, 51]. Then, by relying on the simulation relation, controller strategies or verification results can be mapped from the abstraction to the concrete system with provable correctness guarantees. In the context of stochastic dynamical systems, the abstraction model is usually taken to be a variant of a Markov model with finite state space. In particular, abstractions to robust Markov models and, especially, Interval Markov Decision Processes (IMDPs) [17] have recently shown great potential to achieve state-of-the-art results [5, 11, 12, 22]. However, these approaches are still limited both in terms of memory and computational requirements [11, 26, 35].

In this paper, we introduce a novel framework for formal analysis of stochastic dynamical systems based on abstraction to robust Markov models. Our framework uses compositional reasoning and relies on the structural properties of the system, namely independence of the stochasticity between dimensions, to improve the tightness and computational efficiency of abstraction-based methods that use robust finite-state Markov models. In particular, we introduce a new class of robust Markov models, called orthogonally decoupled Interval Markov Decision Processes (odIMDPs), where the transition probabilities between each pair of states lie within a product of intervals. We show that this form of transition probability is particularly suitable to abstract stochastic systems, where each interval in the product can encode the marginal probabilities of the system. This approach has two main advantages: (i) for each pair of state actions, odIMDPs are only required to store the marginal distributions of the system instead of the transition probabilities to all the other states in the abstraction. This leads to substantial improvements in the memory requirement of the method, which is arguably the main bottleneck of existing abstraction-based methods,

(ii) by encoding constraints on the marginals, the resulting abstraction model exploits the structural properties of the concrete system, leading to improved tightness of the resulting error bounds. To guarantee computational efficiency of our framework, we propose a novel approach to synthesize robust strategies for an odIMDP against a probabilistic reach-avoid property based on a divide-and-conquer algorithm that recursively solves a linear programming problem and returns upper and lower bounds on the probability that the odIMDP satisfies the property. Finally, we show that the resulting strategy can be mapped back on the concrete system and that the returned bounds are guaranteed to contain the probability that the concrete system satisfies the property.

To empirically evaluate the framework, we tested it on benchmarks from the literature [1, 43, 49, 53], including stochastic affine and nonlinear systems, neural network dynamic models, Gaussian Processes (GPs), and Gaussian mixture models. We compared our framework with existing state-of-the-art abstraction-based tools for verification of stochastic systems based on abstractions to Markov Decision Processes (MDPs) and IMDPs [49, 53]. As expected, the experiments show that our approach outperforms the state-of-the-art both in terms of memory requirement and conservativity of the results, allowing one to synthesize optimal strategies for systems that current methodologies cannot handle. For example, on both 6D and 7D systems, while existing approaches returned trivial results, our approach would successfully synthesize optimal policies with a guaranteed lower bound on the satisfaction probability > 0.95 . In summary, the main contributions of the paper are:

- We introduce a new class of robust Markov models called orthogonally decoupled Interval Markov Decision Processes (odIMDPs) and show how robust value iteration for odIMDPs can be performed by relying on linear programming.
- We show that a large class of stochastic dynamical systems can be successfully abstracted as odIMDPs and derive efficient and scalable methods to construct the abstraction.
- We performed a large-scale empirical evaluation and compared our approach with state-of-the-art abstraction-based verification methods for stochastic systems.

Related Works. Abstraction-based approaches for stochastic systems generally rely on abstracting the concrete system to either MDPs or IMDPs [28]. odIMDPs shares many features with IMDPs [17] and abstraction of stochastic systems to IMDPs has been extensively studied, including stochastic hybrid systems [11], neural networks [3], and Gaussian Processes [21, 43], and using data-driven construction methods [5, 20]. Furthermore, many tools support value iteration over IMDPs, including: PRISM [23], bmdp-tool [24], Storm [19], IntervalMDP.jl [35], and IMPaCT¹ [53]. However, a common limitation of these approaches is in the memory requirements in building the abstraction, which is generally quadratic in $|S|$, the size of the IMDP, which is, in turn, exponential in n , the dimension of the state space of the concrete system. In contrast, our approach based on odIMDPs guarantees an improved memory complexity (see Subsection 4) of a factor $\sqrt[n]{|S|}$. Furthermore, we also show how our approach generally returns tighter bounds compared to IMDP-based approaches. To the best of our knowledge,

the only paper that has considered compositional reasoning for IMDPs [18], which, however, focuses on the parallel composition of given IMDPs, while in this paper, we focus on compositional construction of the abstraction. Approaches based on abstractions to (non-robust) MDPs have also been widely employed in the literature [2, 28, 47]. To scale to larger-dimensional systems, various works [9, 16, 29, 40, 49] have devised frameworks that precede a gridding approach, such as the one introduced in our paper, with model order reductions or other modular approaches. However, in these approaches, gridding is still the limiting factor. Consequently, our approach is complementary to these order-reducing approaches and should lead, when combined with these methods, to abstractions of even higher-dimensional systems.

In parallel to abstraction-based methods, various works have focused on abstraction-free methods based on Stochastic Barrier Functions (SBFs) [42, 45]. The idea is to construct a function that allows one to upper-bound the probability that the system exits a given safe set without directly studying the evolution of the system over time. The synthesis of SBFs has been formulated as sum-of-squares optimization [38, 45], with neural networks [33], and as linear programming [37]. Furthermore, data-driven synthesis techniques have also been developed [36, 44]. However, compared to abstraction-based methods, SBFs tend to be conservative [26].

Structure of the paper. In Section 2, we formally state our problem of verifying probabilistic reach-avoid for stochastic systems and outline our approach. Then in Section 3, we briefly review IMDPs and how stochastic systems are abstracted to IMDPs. We also review weaknesses of this approach that leads us to introduce our new Markov model, orthogonally decoupled Interval Markov Decision Processes (odIMDPs), which is introduced in Section 4. Section 4 also discusses the space complexity required to store an odIMDPs and analyse their ambiguity set and compare them to those of state-of-the-art abstraction methods. In Section 5, we show how a stochastic system can be successfully abstracted to odIMDPs. Then, in Section 6 we derive algorithms for robust value iteration of odIMDPs based on linear programming. Finally, in Section 7, we empirically benchmark the proposed methodology.

Notation. \mathbb{R} and \mathbb{N} represent the set of real and natural numbers with $\mathbb{N}_0 = \mathbb{N} \cup \{0\}$. The set \mathbb{S}_{++}^n is the set of symmetric positive-definite matrices of size $n \in \mathbb{N}$. For a finite set S , $|S|$ is its cardinality. A probability distribution γ over a finite set S is a function $\gamma : S \rightarrow [0, 1]$ satisfying $\sum_{s \in S} \gamma(s) = 1$. $\mathcal{D}(S)$ is the set of all probability distributions over S . For $\underline{p}, \bar{p} : S \rightarrow [0, 1]$ such that $\underline{p}(s) \leq \bar{p}(s)$ for each $s \in S$ and $\sum_{s \in S} \underline{p}(s) \leq 1 \leq \sum_{s \in S} \bar{p}(s)$, an interval ambiguity set $\Gamma \subset \mathcal{D}(S)$ is the set of distributions such that $\Gamma = \{\gamma \in \mathcal{D}(S) : \underline{p}(s) \leq \gamma(s) \leq \bar{p}(s) \text{ for each } s \in S\}$. \underline{p}, \bar{p} are often referred to as the interval bounds of the interval ambiguity set. The set of all interval ambiguity sets over S is denoted by $\text{int amb}(S)$. For n finite sets S_1, \dots, S_n we denote by $S_1 \times \dots \times S_n$ their Cartesian product. Given $S = S_1 \times \dots \times S_n$ and n ambiguity sets $\Gamma_i \in \mathcal{D}(S_i)$, $i = 1, \dots, n$, the product ambiguity set $\Gamma \subseteq \mathcal{D}(S)$ is defined as $\Gamma = \{\gamma \in \mathcal{D}(S) : \gamma(s) = \prod_{i=1}^n \gamma^i(s^i), \gamma^i \in \Gamma_i\}$, where $s = (s_1, \dots, s_n) \in S$. In what follows, with a slight abuse of notation, we denote the product ambiguity set as $\Gamma = \bigotimes_{i=1}^n \Gamma_i$. Each Γ_i is called a marginal

¹IMPaCT includes functionality for the abstraction of stochastic systems to IMDPs.

or component ambiguity set. $1_X(x)$ denotes the indicator function for the set X , that is, $1_X(x) = 1$ if $x \in X$ and 0 otherwise.

2 Problem Statement

Consider a stochastic process described by the following stochastic difference equation

$$x_{k+1} = f(x_k, u_k, v_k) \quad (1)$$

where $x_k \in \mathbb{R}^n$ and $u_k \in U \subseteq \mathbb{R}^m$ are respectively state of the system and input at time k . We assume that $|U|$ is finite. $v_k \in V \subseteq \mathbb{R}^{n_v}$ is an independent and identically distributed (i.i.d.) random variable for all time $k \in \mathbb{N}_0$ that represents the noise that affects the system at each time step. The vector field $f : \mathbb{R}^n \times U \times V \rightarrow \mathbb{R}^n$ represents the controlled stochastic dynamics of the system.

To define a probability measure for System (1), we introduce the one-step stochastic kernel $T(X | x, u)$ that defines the probability that System (1) from state x will transition to region $X \subseteq \mathbb{R}^n$ in one time step under action u . For this work, we assume that the transition kernel is a mixture of K Gaussians for $K > 0$, each with diagonal covariance. That is,

$$T(X | x, u) := \sum_{r=1}^K \alpha_r(x, u) \int_X N(x' | \mu^r(x, u), \Sigma^r(x, u)) dx', \quad (2)$$

where $\mu^r(x, u) \in \mathbb{R}^n$ and $\Sigma^r(x, u) \in \mathbb{S}_{++}^n$ are action and state-dependent mean and covariance functions of the r -th Gaussian in the mixture. $\alpha_r(x, u)$ is a weighting function such that $\alpha_r(x, u) \geq 0$ and $\sum_{r=1}^K \alpha_r(x, u) = 1$ for each source-action pair (x, u) . We assume μ^r, Σ^r , and α_r are continuous in x for each u . When $K = 1$, we omit r from the weighting probability function, mean, and covariance.

REMARK 2.1. Note that, while the assumption of diagonal covariance for each distribution in Eqn. (2) introduces independence of the stochasticity at different dimensions, the resulting model still encompasses a large class of models of practical interest. In fact, Eqn. (2) includes linear and nonlinear systems with additive Gaussian noise [3, 4, 11] and systems learned via Gaussian process regression [52]. Furthermore, we should stress that mixtures of Gaussians with diagonal covariance can arbitrarily well approximate any distribution [39, Proposition 4.1].

Given an initial condition $x_0 \in \mathbb{R}^n$ and a control policy² $\pi_x : \mathbb{R}^n \times \mathbb{N}_0 \rightarrow U$, the stochastic kernel T induces a unique and well-defined probability measure P^{x_0, π_x} [7, Prop. 7.45], such that for sets $X_0, X_{k+1} \subseteq X$ it holds that

$$P^{x_0, \pi_x}[x_0 \in X_0] = 1_{X_0}(x_0),$$

$$P^{x_0, \pi_x}[x_{k+1} \in X_{k+1} | x_k = x, u_k = \pi_x(k, x)] = T(X_{k+1} | x, \pi_x(x, k)).$$

With the definition of P^{x_0, π_x} , we can now make probabilistic statements on the trajectories of System (1).

Definition 2.2 (Probabilistic reach-avoid). Consider a compact region of interest $X \subseteq \mathbb{R}^n$. Let $R \subset X$ and $O \subset X$ be closed sets with $R \cap O = \emptyset$, $x_0 \in X$ be the initial state, and $H \in \mathbb{N}$ be the

²For the class of systems and problem considered in this paper time-dependent Markov policies suffices for optimality [26].

time horizon. Then, for a given control strategy π_x , *probabilistic reach-avoid* is defined as

$$P_{\text{ra}}(R, O, x_0, \pi_x, H) = P^{x_0, \pi_x}[\exists k \in \{0, \dots, H\}, x_k \in R, \forall k' \in \{0, \dots, k\}, x_{k'} \notin O \wedge x_{k'} \in X].$$

Note that the assumption of a given region of interest is standard [11, 49] and allows one to only focus on the behavior of the system in a bounded set. The problem we consider in this work is then formally defined as follows.

Problem 2.3 (Verification and controller synthesis). Consider a compact region of interest $X \subseteq \mathbb{R}^n$. Then, for sets $R, O \subset X$, initial state $x_0 \in X$, and time horizon H , synthesize a control strategy π_x^* such that $\pi^* \in \arg \max_{\pi_x} P_{\text{ra}}(R, O, x_0, \pi_x, H)$.

Problem 2.3 requires one to compute the strategy π_x that maximizes the probability that a trajectory of System (1) starting from x_0 enters R within H time steps and avoids O until R is reached. Maximizing the probability is without loss of generality as, as we will show in Section 6, the case min follows similarly. Note that focusing on probabilistic reach-avoid in Problem 2.3 is not limiting. In fact, certifying reach-avoid properties enables verification of more complex temporal logic specifications [6, Alg. 11].

Approach. To solve Problem 2.3, in Section 4, we introduce a class of robust Markov models, called orthogonally decoupled Interval Markov Decision Processes (odIMDPs), where transition probabilities lie in some product ambiguity set. To abstract System (1) into an odIMDP, we first partition the region of interest into a grid of discrete regions. Then, in Section 5, we show that one can take advantage of the product form of an odIMDP and, for each region, simply store the marginal probabilities of the original system. This allows us to construct the abstraction for System (1) in a compositional manner, with improved memory complexity and guaranteeing tighter error bounds compared to existing methods. We then show how an optimal strategy for an odIMDP can be found via linear programming and mapped back to the original system, guaranteeing efficiency. Before formally introducing odIMDPs, in the next Section, we start by reviewing IMDPs and the standard approach to abstracting stochastic systems to IMDPs.

3 Preliminaries on abstractions to Interval Markov Decision Processes

Interval Markov Decision Processes (IMDPs) (also called bounded-parameter MDPs) are a generalization of MDPs in which the transition probability distributions between states lie within some independent intervals [17].

Definition 3.1. An Interval Markov Decision Process (IMDP) is a tuple $M = (S, A, \Gamma)$ where

- S is a finite set of states,
- A is a finite set of actions assumed to be available at each state, and
- $\Gamma = \{\Gamma_{s,a}\}_{s \in S, a \in A}$ are sets of feasible transition probability distributions with $\Gamma_{s,a} \in \text{int amb}(S)$.

IMDPs are commonly used as the abstraction model for complex stochastic systems [5, 11, 12, 22, 25, 46]. In particular, the standard approach to abstract systems of the form of Eqn. (1) to an IMDPs is

first to assign to A the actions in U . Then, the region of interest X is partitioned into regions $\{s_1, \dots, s_{|S|}\}$ and each region is associated with a state in S . In what follows, for the sake of simpler notation, with an abuse, we will denote by s and t , that is, for a transition $s \xrightarrow{a} t$, both the abstract states in S and the corresponding regions of X . Then, for a state $s \in S$ and action $a \in A$, $\Gamma_{s,a}$ is defined as

$$\begin{aligned} \Gamma_{s,a} &= \{\gamma_{s,a} \in \mathcal{D}(S) : \forall t \in S, \\ &\min_{x \in s} T(t \mid x, a) \leq \gamma_{s,a}(t) \leq \max_{x \in s} T(t \mid x, a)\}. \end{aligned} \quad (3)$$

Once the abstraction is built, then one can rely on existing polynomial time algorithms for IMDPs to compute reach-avoid probabilities and optimal strategies on the abstraction IMDP, which can then be mapped back to the concrete system [11, 25]. However, the abstraction approach outlined above, while general and sound, has the following drawbacks:

- Memory consumption is often the main bottleneck. In fact, to perform value iteration on an IMDP M , for each pair of regions $s, t \in S$ and $a \in A$ we need to store $\underline{p}_{s,a}(t) = \min_{x \in s} T(t \mid x, a)$ and $\bar{p}_{s,a}(t) = \max_{x \in s} T(t \mid x, a)$, leading to a memory complexity of $O(|S|^2|A|)$.
- For any $s, t \in S$, the computation of $\min_{x \in s} T(t \mid x, a)$ and $\max_{x \in s} T(t \mid x, a)$ generally requires approximations, which lead to conservatism. For instance, in the case of Gaussian additive noise $x_{k+1} = f(x_k, u_k) + v_k$, where $v_k \sim \mathcal{N}(0, \Sigma)$ with diagonal covariance matrix $\Sigma \in \mathbb{S}_{++}^n$, under the assumption that to each $s \in S$ is associated a hyperrectangular region, i.e., $s = [\underline{s}_1, \bar{s}_1] \times \dots \times [\underline{s}_n, \bar{s}_n]$, the state-of-the-art approach [3, 11, 27] is to rely on the following under-approximation for the min case (similar reasoning holds for the max case):

$$\begin{aligned} \min_{x \in s} T(t \mid x, a) &= \min_{x \in s} \int_t \mathcal{N}(y \mid f(x, a), \Sigma) dy \geq \\ &\prod_{i=1}^n \min_{x \in s} \int_{\underline{L}_i}^{\bar{L}_i} \mathcal{N}(y_i \mid f(x, a)_i, \Sigma_{ii}) dy_i, \end{aligned} \quad (4)$$

which introduces conservatism by allowing each marginal to be optimized independently.

- The abstraction process does not exploit any structural property of the system to reduce the size of the resulting ambiguity set Γ . For instance, if the noise is independent across the various dimensions (as is the case for System (1)), one may want to account for this information to reduce the size of the resulting feasible set of distributions in the abstraction.

In Section 4, we will introduce a new subclass of robust MDPs, called odIMDP, where the feasible set of probabilities has a product form and requires only storing the marginal ambiguity sets, i.e. the right-hand side of Eqn. (4) without the product operator. We will show that by abstracting System (1) to an odIMDP one can alleviate the issues identified above and obtain state-of-the-art performance.

4 Orthogonally Decoupled Interval Markov Decision Processes

Merging ideas from IMDPs [17] and compositional analysis of MDPs [40], we propose a new subclass of robust MDPs, which we call *orthogonally decoupled Interval Markov Decision Processes* (*odIMDPs*). Intuitively, in an odIMDPs for each transition, the ambiguity set

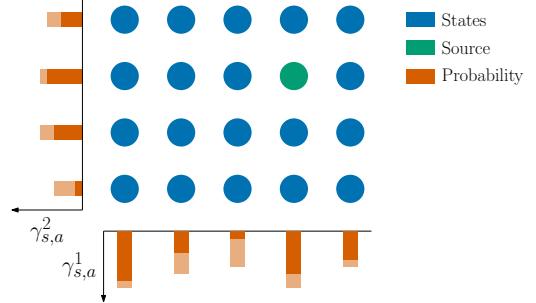


Figure 1: An example of an odIMDP where, given a source-action pair (s, a) (the green state), the ambiguity set for the transition probability can be decomposed into the product between two independent interval ambiguity sets, $\Gamma_{s,a}^1, \Gamma_{s,a}^2$. In this example, it is only necessary to store 9 transitions (per source-action pair) in contrast to 20 transitions for a traditional IMDP.

of the transition probabilities is a product of marginal interval ambiguity sets. In what follows, in this section, we first formally introduce odIMDPs. Then, we show how the memory requirements for this class of models are generally orders of magnitude lower compared to those of IMDPs. Furthermore, we also prove that they can produce tighter ambiguity sets compared to IMDPs. These properties will be employed in Section 5 to efficiently abstract System (1) into an odIMDPs and solve Problem 2.3.

Definition 4.1. An orthogonally decoupled Interval Markov Decision Process (odIMDP) with n marginals, also called components, is a tuple $M = (S, A, \Gamma)$ where

- $S = S_1 \times \dots \times S_n$ is a finite set of joint states with S_i being the set of states in marginal i ,
- A is a finite set of actions assumed to be available in each state, and
- $\Gamma = \{\Gamma_{s,a}\}_{s \in S, a \in A}$ are sets of feasible transition probability distributions where $\Gamma_{s,a} = \bigotimes_{i=1}^n \Gamma_{s,a}^i$ with $\Gamma_{s,a}^i \in \text{int amb}(S_i)$.

Similarly to IMDPs [25], a *path* of an odIMDP is a sequence of states and actions $\omega = (s_0, a_0), (s_1, a_1), \dots$, where $(s_k, a_k) \in S \times A$. We denote by $\omega(k) = s_k$ the state of the path at time $k \in \mathbb{N}_0$ and by Ω the set of all paths. A *strategy* or *policy* for an odIMDP is a function $\pi : S \times \mathbb{N}_0 \rightarrow A$ that assigns an action to a given state of an odIMDP. We note that since the ambiguity set is independent for each source-action pair, also called (s, a) -rectangularity, any optimal policy is time-varying and deterministic [26, 48]. An adversary is a function that assigns a feasible distribution $\gamma \in \Gamma$ to a given state [17]. Given a strategy and an adversary, an odIMDP collapses to a finite Markov chain, with the transition probability matrix specified by marginal distributions.

Analysis of space complexity. Before proceeding further, we now discuss the memory requirements to store an odIMDP. To simplify the presentation, we assume without loss of generality that the number of states in each marginal is equal, that is, $|S_i| = |S_{i+1}|$ for all $i = 1, \dots, n-1$. Hence, we have that the number of states along each marginal is $|S_i| = \sqrt[n]{|S|}$. To store an odIMDP, we need to store

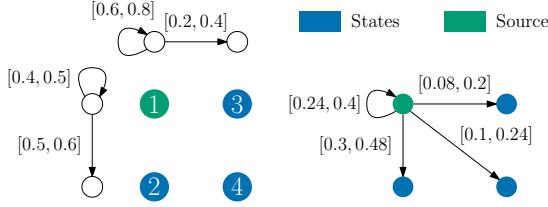


Figure 2: On the left, for the green state, we report two marginal interval ambiguity sets, i.e. an interval ambiguity set for each marginal of a product ambiguity set, with outgoing transitions to all other states. On the right, an IMDP is constructed by multiplying the interval bounds of the marginal ambiguity sets. By this multiplication of bounds, joint distributions are introduced that cannot be represented as a product of distributions from the marginal ambiguity sets.

the bounds on the transition probability. Then, for each source-action pair, we must store bounds for the ambiguity set along each marginal. That is, the upper and lower bound for $|S_i|$ states for each marginal. In total, this requires storing $2|S||A|n\sqrt[4]{|S|}$ scalar values. This is a crucial difference compared to traditional IMDP, where, as discussed in Section 3, for each state-action pair, we need to store $2|S|^2|A|$ scalar values. An example to clarify the different memory requirements between odIMDP and IMDP is reported in Fig. 1.

4.1 Comparison of ambiguity sets

At first glance, the ambiguity set of an odIMDPs may appear equivalent to that of an IMDP obtained by multiplying the interval bounds of the interval ambiguity sets of each marginal for each source-action pair, as, for instance, is done in Eqn. (3) using the approximation step in Eqn. (4). Remarkably, this is, however, a fallacy. In fact, multiplying the interval bounds of the marginal ambiguity sets introduces distributions that are not part of the product ambiguity set, as they may not satisfy the additional constraints on the marginals present in odIMDPs. To better understand why this theoretical quirk occurs, consider the following example.

Example 4.2. Fig. 2 shows an odIMDP represented by an IMDP for each marginal. The states are numbered to index into the joint probability distribution as $[0, 1]^4$. Furthermore, on the right, an interval ambiguity set is constructed from the odIMDP by multiplying the marginal bounds together corresponding to the joint transition. If we consider a distribution $[0.4, 0.3, 0.08, 0.22]^\top$, we see that it is clearly contained in the joint interval ambiguity set on the right of Fig. 2. However, since the value 0.4 in the first entry is equal to the upper bound for that entry in the ambiguity set, it forces $[0.5, \cdot]$ and $[0.8, \cdot]$ in the distribution for the vertical and horizontal marginals, respectively (because $0.4 = 0.5 \cdot 0.8$). For the vertical marginal to sum to one, we must have $[0.5, 0.5]$. The marginal decomposition of the last element 0.22 thus needs to be $0.5p = 0.22$, i.e. $p = 0.22/0.5 = 0.44$, which is not contained in the horizontal marginal ambiguity set (upper bound 0.4). Hence, the distribution contained in the joint interval ambiguity set cannot be factored into two distributions from the marginal ambiguity sets.

The intuition provided in Example 4.2 is formalized in Theorem 4.3 below, where we show that the ambiguity set of an odIMDP is contained in that of an IMDP obtained by multiplying the interval bounds of the interval ambiguity sets of each marginal of the odIMDPs. This result will allow us to build abstractions for System (1) that not only require less memory to be stored compared to those described in Section 3, but are also guaranteed to provide tighter error bounds for the same partition size of the state space.

THEOREM 4.3. For $S = S_1 \times \cdots \times S_n$, consider the interval ambiguity sets $\Gamma^1 \in \text{int amb}(S_1), \dots, \Gamma^n \in \text{int amb}(S_n)$, where $\Gamma^i = \{\gamma \in \mathcal{D}(S_i) : \underline{p}^i(s^i) \leq \gamma(s^i) \leq \bar{p}^i(s^i) \text{ for each } s^i \in S_i\}$. Call $\Gamma = \bigotimes_{i=1}^n \Gamma^i$. Then, it holds that $\Gamma \subseteq \bar{\Gamma}$, where

$$\bar{\Gamma} = \left\{ \gamma \in \mathcal{D}(S) : \prod_{i=1}^n \underline{p}^i(s^i) \leq \gamma(s) \leq \prod_{i=1}^n \bar{p}^i(s^i) \text{ for each } s \in S \right\}.$$

PROOF. It is sufficient to show that any distribution $\gamma \in \Gamma$ is also contained in $\bar{\Gamma}$, that is, it satisfies the interval bounds $\prod_{i=1}^n \underline{p}^i(s^i) \leq \gamma(s) \leq \prod_{i=1}^n \bar{p}^i(s^i)$ for each $s \in S$. By multiplying the (nonnegative) inequalities $\underline{p}^i(s^i) \leq \gamma^i(s^i) \leq \bar{p}^i(s^i)$ for each component $i = 1, \dots, n$, it holds for each $s \in S$ that $\prod_{i=1}^n \underline{p}^i(s^i) \leq \prod_{i=1}^n \gamma^i(s^i) \leq \prod_{i=1}^n \bar{p}^i(s^i)$. Therefore, we can conclude that the product distribution $\gamma = \bigotimes_{i=1}^n \gamma^i$ is contained in $\bar{\Gamma}$. \square

5 Abstraction

In order to describe the abstraction process of System 1 into an odIMDP, we first start from the case where the transition kernel is Gaussian and then move to the more general case. In what follows, we assume that the region of interest $X \subset \mathbb{R}^n$ is hyperrectangular³, that is, $X = [x_1, \bar{x}_1] \times \cdots \times [x_n, \bar{x}_n]$. To abstract a system to an odIMDP, we start by partitioning X into a grid $\tilde{S} = \tilde{S}_1 \times \cdots \times \tilde{S}_{n_x}$ and associate with each (hyperrectangular) region $s = [s_1, \bar{s}_1] \times \cdots \times [s_n, \bar{s}_n] \in \tilde{S}$ an abstract state in \tilde{S} . Similarly to Section 3, we abuse the notation and denote by s both the abstract state in \tilde{S} and the corresponding region in X . The decomposed abstract state is denoted by $s = (s^1, \dots, s^n) \in \tilde{S}$. In addition, we add one sink state \tilde{s}^i for each axis i to represent exiting the set $[x_i, \bar{x}_i]$ and denote $S_i = \{\tilde{s}^i\} \cup \tilde{S}_i$. The reach and avoid states of the abstract model, R and O ⁴ are respectively the set of states satisfying $s \subseteq R$ and $s \cap O \neq \emptyset$ [11, Section 4]. All sink states are classified as avoid states. Finally, we assign an abstract action $a \in A$ to each $u \in U$.

In the following subsections, we focus on computing bounds on the transition probability starting with the Gaussian case.

5.1 Gaussian case

We start by considering the case of a Gaussian transition kernel:

$$T(X \mid x, u) = \int_X \mathcal{N}(x' \mid \mu(x, u), \Sigma(x, u)) dx'. \quad (5)$$

To define the marginal interval ambiguity sets $\Gamma_{s,a}^i$, we first need to introduce some notation. In particular, for each (s, a) we define

³Note that if the region of interest is not hyperrectangular, but bounded, one can always over-approximate it with a hyperrectangular region.

⁴We abuse notation slightly by denoting both the concrete reach/avoid regions and the abstract reach/avoid states by R and O respectively.

intervals $[\underline{\mu}_{s,a}^i, \bar{\mu}_{s,a}^i]$ and $[\underline{\Sigma}_{s,a}^i, \bar{\Sigma}_{s,a}^i]$ such that for all $x \in s$:

$$\underline{\mu}_{s,a}^i \leq \mu(x, a)_i \leq \bar{\mu}_{s,a}^i \quad \text{and} \quad \underline{\Sigma}_{s,a}^i \leq \Sigma(x, a)_{ii} \leq \bar{\Sigma}_{s,a}^i. \quad (6)$$

That is, $[\underline{\mu}_{s,a}^i, \bar{\mu}_{s,a}^i]$ and $[\underline{\Sigma}_{s,a}^i, \bar{\Sigma}_{s,a}^i]$ represent interval bounds for the mean and variance of the i -th marginal of System (1) starting from region s . Then, computation of $\Gamma_{s,a}^i$ reduces to computing for each $t^i \in \tilde{S}_i$ the following transition probabilities

$$\underline{p}_{s,a}^i(t^i) = \min_{\mu^i \in [\underline{\mu}_{s,a}^i, \bar{\mu}_{s,a}^i], \Sigma^i \in \{\underline{\Sigma}_{s,a}^i, \bar{\Sigma}_{s,a}^i\}} \int_{t_i}^{\bar{t}_i} \mathcal{N}(y_i \mid \mu^i, \Sigma^i) dy_i, \quad (7)$$

$$\bar{p}_{s,a}^i(t^i) = \max_{\mu^i \in [\underline{\mu}_{s,a}^i, \bar{\mu}_{s,a}^i], \Sigma^i \in \{\underline{\Sigma}_{s,a}^i, \bar{\Sigma}_{s,a}^i\}} \int_{t_i}^{\bar{t}_i} \mathcal{N}(y_i \mid \mu^i, \Sigma^i) dy_i. \quad (8)$$

REMARK 5.1. Note that analytical solutions for Eqn. (7) and (8) exist [3, 11]. Furthermore, we stress that the computation of the bounds on mean and variance in Eqn. (6) is a well-studied problem for which there exist tools based on convex optimization [3, 11, 30, 51] that can also be applied in the context of Gaussian process regression [41].

For the transition to the sink state, we observe that this is equivalent to the complement of transitioning to inside X . Hence, we may compute the interval bounds via the complement probability.

$$\underline{p}_{s,a}^i(\check{t}^i) = 1 - \max_{\mu^i \in [\underline{\mu}_{s,a}^i, \bar{\mu}_{s,a}^i], \Sigma^i \in \{\underline{\Sigma}_{s,a}^i, \bar{\Sigma}_{s,a}^i\}} \int_{\underline{x}_i}^{\bar{x}_i} \mathcal{N}(y_i \mid \mu^i, \Sigma^i) dy_i, \quad (9)$$

$$\bar{p}_{s,a}^i(\check{t}^i) = 1 - \min_{\mu^i \in [\underline{\mu}_{s,a}^i, \bar{\mu}_{s,a}^i], \Sigma^i \in \{\underline{\Sigma}_{s,a}^i, \bar{\Sigma}_{s,a}^i\}} \int_{\underline{x}_i}^{\bar{x}_i} \mathcal{N}(y_i \mid \mu^i, \Sigma^i) dy_i. \quad (10)$$

Finally, the sink states are absorbing, that is, for all states $s = (s^1, \dots, s^{i-1}, \check{s}^i, s^{i+1}, \dots, s^n)$, we define $\underline{p}_{s,a}^i(\check{s}^i) = \bar{p}_{s,a}^i(\check{s}^i) = 1$ and $\underline{p}_{s,a}^i(t^i) = \bar{p}_{s,a}^i(t^i) = 0$ for any $t^i \in \tilde{S}_i$. From the computation of all interval bounds, we can define Γ as the product of marginal interval ambiguity sets $\Gamma_{s,a}^i$ for each source-action pair (s, a) . Then the odIMDP is $M = (S, A, \Gamma)$.

5.2 Mixture Models

Eqn. (2) is a generalization of the Gaussian case with $K > 1$. However, a key additional technical difficulty is that even if the covariance matrix of each Gaussian in the mixture is diagonal, then the covariance matrix for the mixture distribution is not necessarily diagonal. Specifically, the joint covariance of the mixture distribution in Eqn. (2), omitting the dependence on (x, u) , is

$$\Sigma = \sum_{r=1}^K \alpha_r \Sigma^r + \sum_{r=1}^K \alpha_r (\mu^r - \bar{\mu})(\mu^r - \bar{\mu})^\top \quad (11)$$

Due to the second term of Eqn. (11), the joint covariance Σ is generally not diagonal. Consequently, we cannot directly abstract the model to odIMDPs as it does not have the required product form. To solve this issue, in what follows, we build an odIMDPs for each Gaussian in the mixture following the approach in Section 5.1 and then abstract System (1) into a mixture of odIMDPs.

Definition 5.2. A mixture of K odIMDPs with n marginals is a tuple $M = (S, A, \Gamma, \Gamma^\alpha)$ where

- $S = S_1 \times \dots \times S_n$ is a finite set of joint states with S_i being the set of states in marginal i ,
- A is a finite set of actions,
- $\Gamma = \{\Gamma_{r,s,a}\}_{r \in K, s \in S, a \in A}$ are a set of K product ambiguity sets, where $\Gamma_{r,s,a} = \bigotimes_{i=1}^n \Gamma_{r,s,a}^i$ with $\Gamma_{r,s,a}^i \in \text{int amb}(S_i)$, and
- $\Gamma^\alpha = \{\Gamma_{s,a}^\alpha\}_{s \in S, a \in A}$ are sets of feasible weightings distributions, where $\Gamma_{s,a}^\alpha \in \text{int amb}(K)$. A feasible weighting distribution for a source-action pair (s, a) is denoted by $\alpha_{s,a} \in \Gamma_{s,a}^\alpha$.

$M = (S, A, \Gamma, \Gamma^\alpha)$ can be interpreted as a set of K odIMDPs sharing the same (decomposed) states S and the same set of actions A , but that can have different individual ambiguity sets $\Gamma_{r,s,a}$ for each source-action pair (s, a) . The weighting $\alpha_{s,a} \in \Gamma_{s,a}^\alpha$ defines the probability of selecting a component from the mixture. Thus, a feasible distribution for M is any distribution $\gamma_{s,a} = \sum_{r \in K} \alpha_{s,a}(r) \gamma_{r,s,a}^r$, where $\alpha_{s,a} \in \Gamma_{s,a}^\alpha$ and $\gamma_{r,s,a}^r \in \Gamma_{r,s,a}$ for each $r \in K$.

The process of abstracting System (1) to a mixture of odIMDPs, is as follows:

- (1) S and A are computed similarly to Section 5.1 and are shared among all Gaussians in the mixture,
- (2) for the r -th Gaussian in the mixture follow the approach in Section 5.1 to compute $\Gamma_{r,s,a}$ for any $s \in S$ and $a \in A$,
- (3) for each source-action pair (s, a) , the interval ambiguity set $\Gamma_{s,a}^\alpha \in \text{int amb}(K)$ is constructed such that for all $x \in s$ there exists $\alpha_{s,a} \in \Gamma_{s,a}^\alpha$ with $\alpha^r(x, a) = \alpha_{s,a}(r)$, for all $r \in K$. Similar to the marginal interval ambiguity sets case, we do that by constructing $\Gamma_{s,a}^\alpha$ such that for each $r \in K$ for all $\alpha_{s,a} \in \Gamma_{s,a}^\alpha$ it holds that $\min_{x \in s} \alpha^r(x, a) \leq \alpha_{s,a}(r) \leq \max_{x \in s} \alpha^r(x, a)$.

6 Synthesis of Optimal Policies for odIMDPs

Once an abstraction of System (1) into an odIMDPs or a mixture of odIMDPs is built, what is left to do is to synthesize an optimal strategy, which can then be mapped back on the original system. In what follows, we consider the odIMDPs case; the mixture case follows similarly as we will illustrate in Eqn. (16). For an odIMDPs $M = (S, A, \Gamma)$ synthesizing an optimal strategy for $P_{\text{Ra}}(R, O, x_0, \pi, H)$, reduces to the following optimization problem:

$$\max_{\pi} \min_{\gamma} \mathbb{P}_{\pi, \gamma} [\omega \in \Omega \mid \exists k \in \{0, \dots, K\}, \omega(k) \in R, \forall k' \in \{0, \dots, k\}, \omega(k') \notin O \wedge \omega(k') \in \tilde{S}], \quad (12)$$

where $\mathbb{P}_{\pi, \gamma}$ is the probability of the Markov chain induced by strategy π and distribution γ . In particular, similar to the IMDPs case [17, 25], Eqn. (12) can be computed by solving the following pessimistic robust value iteration:

$$\begin{aligned} V_0(s) &= \mathbf{1}_R(s) \\ V_k(s) &= \max_{a \in A} g \left(s, \min_{\gamma_{s,a} \in \Gamma_{s,a}} \sum_{t \in S} V_{k-1}(t) \gamma_{s,a}(t) \right) \\ &= \max_{a \in A} g \left(s, \min_{\gamma_{s,a}^i \in \Gamma_{s,a}^i, i=1, \dots, n} \sum_{t \in S} V_{k-1}(t) \prod_{i=1}^n \gamma_{s,a}^i(t^i) \right), \end{aligned} \quad (13)$$

where $g(s, v) = \mathbf{1}_R(s) + \mathbf{1}_{S \setminus T}(s)v$ with $T = R \cup O$ and the last equality is due to the fact that for odIMDPs $\Gamma_{s,a} = \bigotimes_{i=1}^n \Gamma_{s,a}^i$. That is,

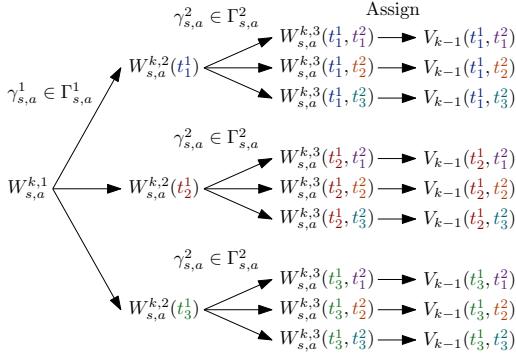


Figure 3: An example of the recursive structure for the proposed algorithm for value iteration over odIMDPs. The example is for an odIMDP with two marginals of three states each. Notice going right-to-left that we optimize over $\gamma_{s,a}^2$ three times, as for each subproblem, we assume t^1 is given.

the ambiguity set is the product of the marginal ambiguity sets. If a strategy is given, we denote the value function by $V_k^\pi(s)$ and if furthermore the inner minimization is replaced by a maximization, we use the notation $\hat{V}_k^\pi(s)$. As a consequence of the product structure, solving Eqn. (13) reduces to iteratively solving the following optimization problem:

$$\min_{\gamma_{s,a}^i \in \Gamma_{s,a}^i, i=1,\dots,n} \sum_{t \in S} V_{k-1}(t) \prod_{i=1}^n \gamma_{s,a}^i(t^i) \quad (14)$$

Because of its multilinear structure, Eqn. (14) cannot be solved directly using standard linear programming approaches for IMDPs, e.g. O-maximization [17]. To solve this issue, we develop a divide-and-conquer approach in which we recursively compute a lower bound on Eqn. (14) by decomposing the problem into a set of linear programming. Specifically, we compute an under-approximation of Eqn. (14), $W_{s,a}^k$ defined recursively as follows:

$$\begin{aligned} W_{s,a}^{k,n}(t^1, \dots, t^n) &= V_{k-1}(t) \\ W_{s,a}^{k,i-1}(t^1, \dots, t^{i-1}) &= \min_{\gamma_{s,a}^i \in \Gamma_{s,a}^i} \sum_{t^i \in S_i} W_{s,a}^{k,i}(t^1, \dots, t^i) \gamma_{s,a}^i(t^i), \\ &\text{for } i = 2, \dots, n \end{aligned} \quad (15)$$

$$W_{s,a}^k := W_{s,a}^{k,0} = \min_{\gamma_{s,a}^i \in \Gamma_{s,a}^i} \sum_{t^i \in S_i} W_{s,a}^{k,1}(t^1) \gamma_{s,a}^i(t^1).$$

Fig. 3 gives an example of Eqn. (15) for a 2-marginals odIMDP with three states in each marginal. The intuition is that we derive a bound for Eqn. (14) by solving a set of simpler optimization problems, one marginal at a time. In particular, $W_{s,a}^{k,i-1}(t^1, \dots, t^{i-1})$ optimizes the expectation of $W_{s,a}^{k,i}$, which is itself a bound of V_{k-1} , for the marginal i , while keeping the values of the other marginals fixed to t^1, \dots, t^{i-1} . Note that, as illustrated in Fig. 3, the resulting algorithm is exponential in n . However, as each of the optimization problems is a particularly simple linear problem that can be solved efficiently using, e.g., the O-maximization algorithm [17, 25] and as we will show in Section 6, the resulting computational time complexity will still be lower compared to the approach described in Section 3

based on IMDPs. First, in the rest of this Section, in Theorem 6.1 and Proposition 6.2 we show, respectively, that the approach in Eqn. (15) is sound and guaranteed to improve the tightness of the results compared to the approach described in Section 3.

THEOREM 6.1. *For any source-action pair $(s, a) \in S \times A$, it holds that*

$$W_{s,a}^{k,0} \leq \min_{\gamma_{s,a}^i \in \Gamma_{s,a}^i, i=1,\dots,n} \sum_{t \in S} V_{k-1}(t) \prod_{i=1}^n \gamma_{s,a}^i(t^i).$$

PROOF.

$$\begin{aligned} W_{s,a}^{k,0} &= \min_{\gamma_{s,a}^1 \in \Gamma_{s,a}^1} \sum_{t^1 \in S_1} W_{s,a}^{k,1}(t^1) \gamma_{s,a}^1(t^1) = \\ &\min_{\gamma_{s,a}^1 \in \Gamma_{s,a}^1} \sum_{t^1 \in S_1} \min_{\gamma_{s,a}^2 \in \Gamma_{s,a}^2} \sum_{t^2 \in S_2} W_{s,a}^{k,2}(t^1, t^2) \gamma_{s,a}^1(t^1) \gamma_{s,a}^2(t^2) \leq \\ &\min_{\gamma_{s,a}^1 \in \Gamma_{s,a}^1} \min_{\gamma_{s,a}^2 \in \Gamma_{s,a}^2} \sum_{t^1 \in S_1} \sum_{t^2 \in S_2} W_{s,a}^{k,2}(t^1, t^2) \gamma_{s,a}^1(t^1) \gamma_{s,a}^2(t^2) \leq \\ &\min_{\gamma_{s,a}^1 \in \Gamma_{s,a}^1, i=1,\dots,n} \sum_{t^1 \in S_1} \dots \sum_{t^n \in S_n} W_{s,a}^{k,n}(t^1, \dots, t^n) \prod_{i=1}^n \gamma_{s,a}^i(t^i) = \\ &\min_{\gamma_{s,a}^i \in \Gamma_{s,a}^i, i=1,\dots,n} \sum_{t \in S} V_{k-1}(t) \prod_{i=1}^n \gamma_{s,a}^i(t^i). \end{aligned}$$

□

Note that an alternative option to the divide-and-conquer algorithm considered here would be to simply multiply the interval bounds of the marginal ambiguity sets for each source-action pair, obtaining an ambiguity set identical to the one in Eqn. (3). As already mentioned, we do not do that because it would increase the conservativism of the resulting approach. This is proved in the following proposition that guarantees that our approach in Eqn. (15) is guaranteed not to be more conservative than the one described in Section 3. In practice, in the experimental section, we show how our approach consistently returns substantially tighter bounds.

PROPOSITION 6.2. *For an odIMDPs $M = (S, A, \Gamma)$ consider the function $V_{k-1} : S \rightarrow \mathbb{R}$ and define the interval ambiguity set associated to M as $\bar{\Gamma}_{s,a} = \left\{ \gamma \in \mathcal{D}(S) : \underline{p}_{s,a}(t) \leq \gamma_{s,a}(t) \leq \bar{p}_{s,a}(t) \text{ for each } t \in S \right\}$, where $\underline{p}_{s,a}(t) = \prod_{i=1}^n \underline{p}_{s,a}^i(t^i)$ and $\bar{p}_{s,a}(t) = \prod_{i=1}^n \bar{p}_{s,a}^i(t^i)$. Then, for any $(s, a) \in S \times A$ it holds that $W_{s,a}^k \geq \min_{\gamma_{s,a} \in \bar{\Gamma}_{s,a}} \sum_{t \in S} V_{k-1}(t) \gamma_{s,a}(t^i)$.*

PROOF. To conclude the proof, it suffices to show that for any source-action pair $(s, a) \in S \times A$, the joint distribution $\gamma_{s,a}(t)$ obtained by optimizing Eqn. (15) is contained in $\bar{\Gamma}_{s,a}$. By definition of sets $\Gamma_{s,a}^i$ in Eqn. (15) it holds that the i -th marginal of $\gamma_{s,a}(t)$ must be contained within $[\underline{p}_{s,a}^i(t^i), \bar{p}_{s,a}^i(t^i)]$. This implies that $\prod_{i=1}^n \underline{p}_{s,a}^i(t^i) \leq \gamma_{s,a}(t) \leq \prod_{i=1}^n \bar{p}_{s,a}^i(t^i)$. □

The proposed divide-and-conquer algorithm extends readily to mixtures of odIMDPs by executing the algorithm in Eqn. (15) for each element of the mixture and optimizing for the worst mixture distribution. That is, by solving:

$$W_{s,a}^k := \min_{\alpha \in \Gamma_{s,a}^\alpha} \sum_{r \in K} \alpha(r) W_{r,s,a}^{k,0} \quad (16)$$

Table 1: A summary of the benchmarks and their characteristics. Note that the number of states for the different experiments is the specified amount unless otherwise stated.

Name	Dimension	Dynamics type	Property	# inputs	# states
Car parking [49]	2	Additive linear	Reach-avoid	9	1600
Robot reachability [53]	2	Additive linear (non-linear control)	Reachability	121	400
Robot reach-avoid [53]	2	Additive linear (non-linear control)	Reach-avoid	441	1600
Building automation system [53]	4	Additive affine	Safety	4	1225
Van der Pol [49]	2	Additive polynomial	Reachability	11	2500
NNDM Cartpole	4	Additive Neural Network Dynamic Model (NNDM)	Safety	2	192000
6D linear model	6	Additive linear	Safety	0	262144
7D linear model	7	Additive linear	Safety	0	2097152
Dubin's car GP [43]	3	Gaussian process with deep kernel learning	Reach-avoid	7	25600
Stochastic switched linear	2	Gaussian mixture, each linear	Reach-avoid	0	1600

where $W_{r,s,a}^{k,0}$ is computed as in Eqn. (15). As $\Gamma_{s,a}^\alpha$ is an interval ambiguity set, the above is a linear program that can be efficiently solved using O-maximization [17].

Analysis of time complexity. Similarly to the analysis of the space complexity (see Subsection 4), we assume that the number of states in each marginal is equal, i.e. $|S_i| = \sqrt[n]{|S|}$. Since value iteration is performed independently for each source-action pair, to compute the time complexity of the algorithm, we can analyze each pair separately. To do this, we start by observing the number of (branching) internal nodes in the recursion tree required to solve Eqn. (15) (see Fig. 3) is a geometric series with base $|S_i|$ up to exponent $n-1$. The complexity of the geometric series is $O\left((|S_i|)^{n-1}\right)$ [14]. At each internal node of the tree, the O-maximization algorithm [25] is executed once on a set of size $|S_i|$. O-maximization is limited by the required sorting, which is of complexity $O(|S_i| \lg |S_i|)$. Now, since $|S_i| = \sqrt[n]{|S|} = |S|^{\frac{1}{n}}$, the complexity for each source-action pair is $O(|S| \lg |S_i|)$, and the full complexity is thus $O(|S|^2 |A| \lg \sqrt[n]{|S|})$. In contrast, value iteration for a regular IMDP [17] has a time complexity of $O(|S|^2 |A| \lg |S|)$.

6.1 Correctness

In this subsection, we map the strategy π to a switching strategy $\pi_x : \mathbb{R}^n \times \mathbb{N}_0 \rightarrow U$ for System (1), ensuring the correctness of our approach. To this end, define the function $J : \mathbb{R}^n \rightarrow S$ that maps each concrete state to an abstract state and its corresponding region. That is, $J(x) = s$ if and only if $x \in s$. Then the switching strategy for System (1) is defined as $\pi_x(x, k) = \pi(J(x), k)$ for all $x \in \mathbb{R}^n$. The following theorem ensures the correctness of π_x .

THEOREM 6.3. *Let M be an odIMDP abstraction of System (1). Then, for any strategy π , it holds that for any $x_0 \in X$*

$$P_{\text{ra}}(R, O, x_0, \pi_x, H) \in [V_H^\pi(J(x_0)), \hat{V}_H^\pi(J(x_0))],$$

where $V_H^\pi(s), \hat{V}_H^\pi(s)$ are the lower and upper bounds on the satisfaction probability for M , as introduced in Section 6.

The proof follows similarly to [26, Theorem 4] by induction on Eqn. (13) and relying on Theorem 6.1 to guarantee the soundness of our approach at any time step. Note that if the regions R and O do not align with the partitioning, then extra care is needed for computing the upper bound on satisfaction [11, Section 7].

7 Experiments

To show the efficacy of our approach, we conducted a wide range of empirical studies with benchmarks ranging from linear 2D systems to nonlinear systems, linear 7D systems, Neural Network Dynamic Models (NNDMs), and Gaussian mixture models. A summary of the benchmarks can be found in Table 1 and extended descriptions are provided in Appendix A of [34]. All benchmarks are verified for a time horizon $H = 10$ unless otherwise specified. We compare our results against state-of-the-art tools for abstractions to IMDPs and MDPs. In particular, for IMDPs, we compare with the method in [3, Theorem 1], as well as IMPaCT [53], which uses nonlinear optimization and Monte Carlo integration to compute the transition probabilities needed to compute the abstraction as illustrated in Section 3. Furthermore, we also compare with SySCoRe [49], which is based on model reduction and abstractions to MDPs. To run value iteration for both the proposed method and the method of [3], we use IntervalMDP.jl [35], which is our package for parallelized value iteration over IMDPs and odIMDPs written in Julia [8]. The experiments are all run on the TU Delft supercomputer, DelftBlue [15], on one memory partition node with an Intel Xeon E5-6248R CPU (12 cores allocated) and 200 GB memory. We stress that despite recent advances in performing value iteration on GPUs [35, 53], the experiments are all conducted on CPUs.

7.1 Comparison with IMDP-based approaches

To compare abstraction approaches based on odIMDPs and IMDPs, for each method, we report abstraction and certification time and memory usage. Furthermore, to quantify the conservatism of the results, we consider the following metrics: (i) the mean lower bound $\text{mean}(\{V_H^\pi(s) : s \in S \setminus T\})$ and (ii) the mean error, $\epsilon = \text{mean}(\{\hat{V}_H^\pi(s) - V_H^\pi(s) : s \in S \setminus T\})$ where T is the set of terminal states and π is the optimal strategy for Eqn. (15). \hat{V}_H^π and V_H^π are the lower and upper bounds on the satisfaction probability returned by our approach. That is, the solution of Eqn. (12) using the approach in Eqn. (15). Furthermore, we report the (aggregated) difference in V between the methods as $\text{agg}(\{V_H^{\text{odIMDP}}(s) - V_H^{\text{other}}(s) : s \in S \setminus T\})$ where the aggregation functions are $\text{agg} \in \{\min, \max, \text{mean}\}$ and the methods are $\text{other} \in \{\text{IMDP}, \text{IMPaCT}\}$ (IMDP refers to the method in [3, Theorem 1]). We consider the same discretization size in the abstraction process for all methods, reported in Table 1.

The results of the comparison with IMDP-based approaches can be found in Table 2 and Table 3 for computation performance and satisfaction probability, respectively. Comparing the computation time of [3, Theorem 1] and IMPaCT [53] with that of our method based on odIMDPs, the abstraction time of odIMDPs is at least 5x faster, up to 80x faster, for 2D benchmarks, and at least two orders of magnitude faster for >3D benchmarks. The certification time is comparable to or faster than both [3, Theorem 1] and IMPaCT. As expected, the memory requirements are considerably lower for odIMDPs: Compared to IMPaCT, odIMDPs use at least an order of magnitude less memory. Most remarkably, for the 6D and 7D linear models, odIMDPs uses 4682x and 30476x less memory, respectively, compared to IMPaCT. If we compare the satisfaction probability results in Table 3, the proposed method is at least as good as the method [3, Theorem 1] for all states in all examples. This ranges from almost the same satisfaction probability for the 2D robot with a reach-avoid specification to 17.33 percentage points better on average for the 4D building automation system. A similar analysis holds for IMPaCT, with the caveat that IMPaCT uses non-linear optimization to compute Eqn. (3), thus its results are generally tighter compared to those of [3, Theorem 1].

7.2 Comparison with SySCoRe (MDP-based approach)

We now also compare with SySCoRe⁵ [49], which represents a state-of-the-art tool for control and verification of stochastic systems based on abstractions to MDPs. Note that, as SySCoRe only supports gridding of 2D systems, we cannot compare on systems of higher dimension. In particular, in Fig. 4, we consider the car parking benchmark introduced in Table 1, and for various sizes of the abstraction, we report the error ϵ , as introduced in the previous subsection. From the figure, we have three key observations:

- SySCoRe requires a relatively large number of regions in the partition before its mean error decreases with larger time horizons.
- odIMDPs tends to obtain substantially tighter bounds for the same abstraction size. For example, the mean error at time 100 of SySCoRe with an MDP of size $90k$ is larger than that of our approach with an odIMDPs of $2.5k$ states.
- The mean error for odIMDPs first increases with time horizons to a level slightly above SySCoRe with the same number of regions and then decreases for longer horizons.

We should, however, stress that SySCoRe requires significantly less memory compared to our approach for the same partition size. The memory usage compared to the mean error at convergence for both SySCoRe and odIMDPs is reported in Table 4. Nevertheless, especially for longer time horizon, because of the fewer regions per dimension required by our approach to achieve a similar level of error, we expect our approach based on odIMDPs to be substantially more scalable for higher dimensions. In fact, this has already been observed in [11] for IMDP-based approaches compared to

⁵We only compare with the finite-state abstraction of SySCoRe, i.e. without model order reductions, to assess the impact of the model choice. As mentioned in Section 1, model order reductions are also interesting in the context of IMDP-based abstractions and is subject to future work.

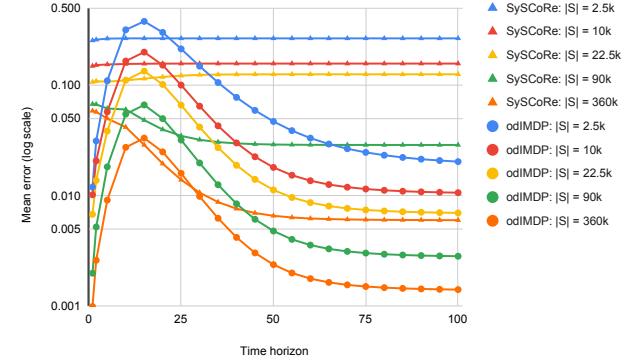


Figure 4: Comparing ϵ for a varying number of regions on the abstraction (assuming a uniform partition of the state space) and time horizons for both SySCoRe [49] and odIMDPs.

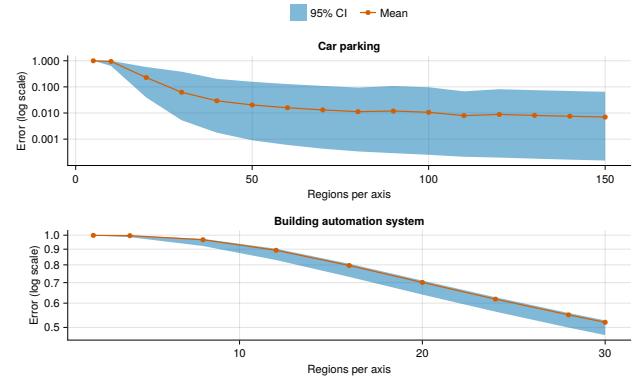


Figure 5: Mean error and 95% confidence interval, with respect to a uniform distribution of initial conditions, for various sizes of partitions for the car parking and building automation system benchmarks.

MDP-based approaches and, as illustrated in Table 2, our approach substantially outperforms that of [11] in terms of scalability.

7.3 Empirical convergence analysis

As hinted in Fig. 4, the error of our approach is reduced with finer partitioning of the region of interest X . To analyze this behavior in more detail, in Fig. 5 we report the mean error and 95% confidence interval as a function of the number of regions per axis for the car parking and building automation system benchmarks. In both cases, the mean error decreases with more regions for both systems. The slight increase at various points for the car parking benchmark is attributed to a misalignment of the reach and avoid regions with the partitioning of X (see Section 6.1 for more details). The confidence interval for the car parking benchmark suggests that while a larger number of regions quickly reach low errors, some hardly converge. Future work will study the convergence from a theoretical perspective.

Table 2: Computation time and memory requirements for IMDP-based approaches across different benchmarks. Time is measured in seconds, and memory is measured in MB. OOM denotes out of memory. To compare the memory requirements for the benchmarks with OOM, we have calculated the amount of memory required to store the transition interval bounds (as a dense matrix) using the formula for IMDPs in Subsection 4.

Benchmark	Our method			Adams et al. [3]			IMPaCT [53]		
	Abs. time	Cert. time	Mem.	Abs. time	Cert. time	Mem.	Abs. time	Cert. time	Mem.
Car parking	0.150	0.146	19.2	13.497	0.255	138.1	19.570	0.846	304.9
Robot reachability	0.665	0.168	21.6	12.659	0.129	88.0	20.611	0.765	306.7
Robot reach-avoid	14.259	7.641	547.2	1136.720	6.739	4143.8	918.856	37.526	16388.0
Building automation system	0.023	0.237	3.1	7.273	0.285	105.1	17.564	0.318	96.0
Van der Pol	1.658	0.257	45.5	27.255	0.827	353.6	113.235	3.233	1093.4
NNMD Cartpole	236.154	550.747	610.8	42326.400	3.751	1590.9	Incompatible dynamics		
6D linear model	8.959	359.887	237.3	OOM			OOM ($\approx 1.1TB$)		
7D linear model	66.231	13903.540	2310.8	Timeout (24h)			OOM ($\approx 70.4TB$)		
Dubin's car GP	13.816	19.562	352.2	336.940	31.333	14265.8	Incompatible dynamics		
Stochastic switched linear	0.033	0.045	4.5	3.391	0.038	41.0	NLopt failure		

Table 3: Satisfaction probabilities for IMDP-based approaches. V denotes the lower bound satisfaction probability and ϵ the mean error. For Adams et al. [3] and IMPaCT [53], we also report the (minimum, maximum, and mean over regions) difference δ in V to our method, where positive means that our method yields a higher satisfaction probability and vice versa. The unreliable result for IMPaCT with the Van der Pol benchmark is due to it returning NaN for every state.

Benchmark	Our method		Adams et al. [3]					IMPaCT [53]					
	Mean V	ϵ	Mean V	ϵ	Min δ	Max δ	Mean δ	Mean V	ϵ	Min δ	Max δ	Mean δ	
Car parking	0.269	0.3885	0.213	0.5315	0.0040	0.1428	0.0560	0.213	0.5183	0.0040	0.1422	0.0556	
Robot reachability	0.889	0.1108	0.881	0.1186	0.0060	0.0119	0.0079	0.890	0.1098	-0.0058	0.0022	-0.0010	
Robot reach-avoid	0.980	0.0199	0.979	0.0208	0.0005	0.0027	0.0008	0.980	0.0202	-0.0001	0.0018	0.0003	
Building automation system	0.263	0.7336	0.090	0.9076	0.0510	0.2297	0.1733	0.174	0.8237	0.0304	0.1131	0.0897	
Van der Pol	0.069	0.3367	0.051	0.4178	0.0000	0.0529	0.0177	Unreliable results					
NNMD Cartpole	0.004	0.7634	0.000	0.7184	0.0000	0.4101	0.0037	Incompatible dynamics					
6D linear model	0.958	0.0419	OOM			OOM			OOM				
7D linear model	0.952	0.0483	Timeout			OOM			OOM				
Dubin's car GP	0.362	0.3461	0.216	0.5046	0.0000	0.8383	0.1458	Incompatible dynamics					
Stochastic switched linear	0.411	0.2828	0.366	0.3605	0.0000	0.0979	0.0456	NLopt failure					

Table 4: The memory usage and the mean error ϵ at 100 time steps for SySCoRe [49] and odIMDPs on the car parking benchmark (see Table 1).

# regions	SySCoRe		odIMDPs	
	Mem. (MB)	ϵ	Mem. (MB)	ϵ
2500	0.80	0.2650	35.55	0.0204
10000	3.20	0.1568	279.51	0.0107
22500	7.20	0.1254	937.86	0.0070
90000	28.81	0.0288	7459.02	0.0028
360000	115.21	0.0060	59499.08	0.0014

8 Conclusion

We presented a novel approach for control synthesis of non-linear stochastic systems against probabilistic reach-avoid specifications. Our approach is based on abstractions to orthogonally decoupled Interval Markov Decision Process (odIMDP), which is a new class of robust Markov models with uncertain transition probabilities

having the product form. We showed how such a structure on the transition probabilities allows one to abstract a large class of stochastic systems by obtaining substantial improvements both in terms of memory requirements and tightness of the results compared to state-of-the-art. The theoretical findings are supported by experimental results, showing how our new approach can successfully outperform existing competing methods in various tasks.

Naturally, our proposed approach is not without limitations. First, our framework only applies to systems with independence in the stochasticity across dimensions. As discussed, this includes various systems commonly used in practice. However, if there is dependence across the marginals, then the transition probabilities do not have the product form. Consequently, our framework cannot be applied. How to extend our approach to this setting is an important future direction. Secondly, exactly and efficiently solving the multi-linear problem of Eqn. (14) is still an open problem for which we employed relaxations. Finally, our approach assumes on a grid-partitioning of the region of interest, although the literature [3, 43] on refinement has shown that heterogeneous abstractions can be beneficial.

References

- [1] Alessandro Abate, H.A.P. Blom, Nathalie Cauchi, Joanna Delicaris, Arnd Hartmanns, Mahmoud Khaled, Abolfazl Lavaei, Carina Pilch, Anne Remke, and More Authors. 2020. ARCH-COMP20 Category Report: Stochastic Models. In *EPiC Series in Computing*.
- [2] Alessandro Abate, Joost-Pieter Katoen, John Lygeros, and Maria Prandini. 2010. Approximate model checking of stochastic hybrid systems. *European Journal of Control* 16, 6 (2010), 624–641.
- [3] Steven Adams, Morteza Lahijanian, and Luca Laurenti. 2022. Formal control synthesis for stochastic neural network dynamic models. *IEEE Control Systems Letters* 6 (2022), 2858–2863.
- [4] Steven Adams, Andrea Patane, Morteza Lahijanian, and Luca Laurenti. 2024. Finite Neural Networks as Mixtures of Gaussian Processes: From Provable Error Bounds to Prior Selection. *arXiv e-prints* (2024), arXiv–2407.
- [5] Thom Badings, Licio Romao, Alessandro Abate, David Parker, Hasan A Poonawala, Marielle Stoelinga, and Nils Jansen. 2023. Robust control for dynamical systems with non-gaussian noise via formal abstractions. *Journal of Artificial Intelligence Research* 76 (2023), 341–391.
- [6] Christel Baier and Joost-Pieter Katoen. 2008. *Principles of Model Checking*. The MIT Press.
- [7] Dimitri P Bertsekas and Steven E Shreve. 2004. *Stochastic Optimal Control: the Discrete-time Case*. Athena Scientific.
- [8] Jeff Bezanson, Stefan Karpinski, Viral B Shah, and Alan Edelman. 2012. Julia: A fast dynamic language for technical computing. *arXiv preprint arXiv:1209.5145* (2012).
- [9] Craig Boutilier, Richard Dearden, and Moisés Goldszmidt. 2000. Stochastic dynamic programming with factored representations. *Artificial Intelligence* 121, 1 (2000), 49–107. doi:10.1016/S0004-3702(00)00033-3
- [10] Ashwin Carvalho, Yiqi Gao, Stéphanie Lefevre, and Francesco Borrelli. 2014. Stochastic predictive control of autonomous vehicles in uncertain environments. In *12th international symposium on advanced vehicle control*, Vol. 9.
- [11] Nathalie Cauchi, Luca Laurenti, Morteza Lahijanian, Alessandro Abate, Marta Kwiatkowska, and Luca Cardelli. 2019. Efficiency through uncertainty: Scalable formal synthesis for stochastic hybrid systems. In *Proceedings of the 22nd ACM international conference on hybrid systems: computation and control*. 240–251.
- [12] Krishnendu Chatterjee, Koushik Sen, and Thomas A Henzinger. 2008. Model-checking ω -regular properties of interval Markov chains. In *Foundations of Software Science and Computational Structures: 11th International Conference, FOSSACS 2008, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2008, Budapest, Hungary, March 29–April 6, 2008. Proceedings*. Springer, 302–317.
- [13] Steve A. Chien. 2022. Formal Methods for Trusted Space Autonomy: Boon or Bane?. In *NASA Formal Methods*, Jyotirmoy V. Deshmukh, Klaus Havelund, and Ivan Perez (Eds.). Springer International Publishing, Cham, 3–13.
- [14] Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, and Clifford Stein. 2001. *Introduction to Algorithms* (2 ed.). The MIT Press.
- [15] Delft High Performance Computing Centre. 2024. DelftBlue Supercomputer (Phase 2). <https://www.tudelft.nl/dhpc/ark/44463/DelftBluePhase2>.
- [16] Sadegh Esmaili Zadeh Soudjani, Alessandro Abate, and Rupak Majumdar. 2016. Dynamic Bayesian networks for formal verification of structured stochastic processes. *Acta Informatica* 54, 2 (Dec. 2016), 217–242. doi:10.1007/s00236-016-0287-9
- [17] Robert Givan, Sonja Leach, and Thomas Dean. 2000. Bounded-parameter Markov decision processes. *Artificial Intelligence* 122, 1-2 (2000), 71–109.
- [18] Vahid Hashemi, Andrea Turrini, Ernst Moritz Hahn, Holger Hermanns, and Khaled Elbassioni. 2017. Polynomial-Time Alternating Probabilistic Bisimulation for Interval MDPs. In *Dependable Software Engineering. Theories, Tools, and Applications*, Kim Guldstrand Larsen, Oleg Sokolsky, and Ji Wang (Eds.). Springer International Publishing, Cham, 25–41.
- [19] Christian Hensel, Sebastian Junges, Joost-Pieter Katoen, Tim Quatmann, and Matthias Volk. 2021. The probabilistic model checker Storm. *International Journal on Software Tools for Technology Transfer* 24, 4 (July 2021), 589–610. doi:10.1007/s10009-021-00633-z
- [20] John Jackson, Luca Laurenti, Eric Frew, and Morteza Lahijanian. 2020. Towards Data-driven Verification of Unknown Dynamical Systems. In *2nd RSS Workshop on Robust Autonomy: Tools for Safety in Real-World Uncertain Environments (RSS 2020)*.
- [21] John Jackson, Luca Laurenti, Eric Frew, and Morteza Lahijanian. 2021. Formal verification of unknown dynamical systems via Gaussian process regression. *arXiv preprint arXiv:2201.00655* (2021).
- [22] Jesse Jiang, Ye Zhao, and Samuel Coogan. 2022. Safe learning for uncertainty-aware planning via interval MDP abstraction. *IEEE Control Systems Letters* 6 (2022), 2641–2646.
- [23] Marta Kwiatkowska, Gethin Norman, and David Parker. 2011. PRISM 4.0: Verification of probabilistic real-time systems. In *Computer Aided Verification: 23rd International Conference*. Springer, 585–591.
- [24] Morteza Lahijanian. 2021. bmdp-tool. <https://github.com/aria-systems-group/bmdp-tool> Standalone software.
- [25] Morteza Lahijanian, Sean B Andersson, and Calin Belta. 2015. Formal verification and synthesis for discrete-time stochastic systems. *IEEE Trans. Automat. Control* 60, 8 (2015), 2031–2045.
- [26] Luca Laurenti and Morteza Lahijanian. 2023. Unifying Safety Approaches for Stochastic Systems: From Barrier Functions to Uncertain Abstractions via Dynamic Programming. *arXiv preprint arXiv:2310.01802* (2023).
- [27] Luca Laurenti, Morteza Lahijanian, Alessandro Abate, Luca Cardelli, and Marta Kwiatkowska. 2020. Formal and efficient synthesis for continuous-time linear stochastic hybrid processes. *IEEE Trans. Automat. Control* 66, 1 (2020), 17–32.
- [28] Abolfazl Lavaei, Sadegh Soudjani, Alessandro Abate, and Majid Zamani. 2022. Automated verification and synthesis of stochastic hybrid systems: A survey. *Automatica* 146 (2022), 110617.
- [29] Abolfazl Lavaei, Sadegh Soudjani, and Majid Zamani. 2021. Compositional abstraction-based synthesis of general MDPs via approximate probabilistic relations. *Nonlinear Analysis: Hybrid Systems* 39 (2021), 100991. doi:10.1016/j.nahs.2020.100991
- [30] Changliu Liu, Tomer Arnon, Christopher Lazarus, Christopher Strong, Clark Barrett, Mykel J Kochenderfer, et al. 2021. Algorithms for verifying deep neural networks. *Foundations and Trends in Optimization* (2021).
- [31] Jun Liu and Necmiye Ozay. 2014. Abstraction, discretization, and robustness in temporal logic control of dynamical systems. In *Proceedings of the 17th international conference on Hybrid systems: computation and control*.
- [32] Ryan Luna, Morteza Lahijanian, Mark Moll, and Lydia E. Kavraki. 2014. Fast stochastic motion planning with optimality guarantees using local policy re-configuration. In *2014 IEEE International Conference on Robotics and Automation (ICRA)*. 3013–3019. doi:10.1109/ICRA.2014.6907293
- [33] Frederik Baymler Mathiesen, Simeon C. Calvert, and Luca Laurenti. 2023. Safety Certification for Stochastic Systems via Neural Barrier Functions. *IEEE Control Systems Letters* 7 (2023), 973–978. doi:10.1109/LCSYS.2022.3229865
- [34] Frederik Baymler Mathiesen, Sofie Haesaert, and Luca Laurenti. 2025. Scalable control synthesis for stochastic systems via structural IMDP abstractions. *arXiv:2411.11803 [eess.SY]* <https://arxiv.org/abs/2411.11803>
- [35] Frederik Baymler Mathiesen, Morteza Lahijanian, and Luca Laurenti. 2024. IntervalMDP.jl: Accelerated Value Iteration for Interval Markov Decision Processes. *arXiv:2401.04068 [eess.SY]* <https://arxiv.org/abs/2401.04068>
- [36] Frederik Baymler Mathiesen, Licio Romao, Simeon C. Calvert, Alessandro Abate, and Luca Laurenti. 2023. Inner Approximations of Stochastic Programs for Data-Driven Stochastic Barrier Function Design. In *2023 62nd IEEE Conference on Decision and Control (CDC)*. 3073–3080. doi:10.1109/CDC49753.2023.10383306
- [37] Rayan Mazouz, Frederik Baymler Mathiesen, Luca Laurenti, and Morteza Lahijanian. 2024. Piecewise Stochastic Barrier Functions. *arXiv:2404.16986 [cs.RO]* <https://arxiv.org/abs/2404.16986>
- [38] Rayan Mazouz, Karan Muvvula, Akash Ratheesh Babu, Luca Laurenti, and Morteza Lahijanian. 2022. Safety guarantees for neural network dynamic systems via stochastic barrier functions. *Advances in Neural Information Processing Systems* 35 (2022), 9672–9686.
- [39] Vassili Nesteridis, Sebastian Schmutzhard, and Vangelis Stefanopoulos. 2011. Universal series induced by approximate identities and some relevant applications. *Journal of Approximation Theory* (2011).
- [40] Petter Nilsson, Sofie Haesaert, Rohan Thakker, Kyohei Otsu, Cristian Ioan Vasile, Ali-Akbar Agha-Mohammadi, Richard M Murray, and Aaron D Ames. 2018. Toward Specification-Guided Active Mars Exploration for Cooperative Robot Teams.. In *Robotics: Science and systems*, Vol. 14. 1–9.
- [41] Andrea Patane, Arne Blaas, Luca Laurenti, Luca Cardelli, Stephen Roberts, and Marta Kwiatkowska. 2022. Adversarial robustness guarantees for gaussian processes. *Journal of Machine Learning Research* 23, 146 (2022), 1–55.
- [42] Stephen Prajna, Ali Jadbabaie, and George J. Pappas. 2007. A Framework for Worst-Case and Stochastic Safety Verification Using Barrier Certificates. *IEEE Trans. Automat. Control* 52, 8 (2007), 1415–1428. doi:10.1109/TAC.2007.902736
- [43] Robert Reed, Luca Laurenti, and Morteza Lahijanian. 2023. Promises of deep kernel learning for control synthesis. *IEEE Control Systems Letters* (2023).
- [44] Ali Salamat, Abolfazl Lavaei, Sadegh Soudjani, and Majid Zamani. 2021. Data-Driven Safety Verification of Stochastic Systems via Barrier Certificates. *IFAC-PapersOnLine* 54, 5 (2021), 7–12. doi:10.1016/j.ifacol.2021.08.466 7th IFAC Conference on Analysis and Design of Hybrid Systems ADHS 2021.
- [45] Cesar Santoyo, Maxence Dutreix, and Samuel Coogan. 2021. A barrier function approach to finite-time stochastic system verification and control. *Automatica* 125 (2021), 109439. doi:10.1016/j.automatica.2020.109439
- [46] John Skovbekk, Luca Laurenti, Eric Frew, and Morteza Lahijanian. 2023. Formal Abstraction of General Stochastic Systems via Noise Partitioning. *IEEE Control Systems Letters* (2023).
- [47] Sadegh Esmaili Zadeh Soudjani and Alessandro Abate. 2014. Aggregation and control of populations of thermostatically controlled loads by formal abstractions. *IEEE Transactions on Control Systems Technology* 23, 3 (2014), 975–990.
- [48] Marnix Suilen, Thom Badings, Eline M Bovy, David Parker, and Nils Jansen. 2024. Robust Markov Decision Processes: A Place Where AI and Formal Methods

Meet. In *Principles of Verification: Cycling the Probabilistic Landscape: Essays Dedicated to Joost-Pieter Katoen on the Occasion of His 60th Birthday, Part III*. Springer, 126–154.

[49] Birgit Van Huijgevoort, Oliver Schön, Sadegh Soudjani, and Sofie Haesaert. 2023. SySCoRe: Synthesis via stochastic coupling relations. In *Proceedings of the 26th ACM international conference on hybrid systems: Computation and control*.

[50] Birgit C. van Huijgevoort and Sofie Haesaert. 2022. Similarity quantification for linear stochastic systems: A coupling compensator approach. *Automatica* 144 (2022), 110476. doi:10.1016/j.automatica.2022.110476

[51] B. C. van Huijgevoort, S. Weiland, and S. Haesaert. 2023. Temporal Logic Control of Nonlinear Stochastic Systems Using a Piecewise-Affine Abstraction. *IEEE Control Systems Letters* (2023). doi:10.1109/LCSYS.2022.3230765

[52] Christopher Williams and Carl Rasmussen. 1995. Gaussian processes for regression. *Advances in neural information processing systems* 8 (1995).

[53] Ben Wooding and Abolfazl Lavaei. 2024. IMPaCT: Interval MDP Parallel Construction for Controller Synthesis of Large-Scale STochastic Systems. In *Quantitative Evaluation of Systems and Formal Modeling and Analysis of Timed Systems*.