2023 62nd IEEE Conference on Decision and Control (CDC)
December 13-15, 2023. Marina Bay Sands, Singapore

MDP Abstractions from Data: Large-Scale Stochastic Networks

Abolfazl Lavaei, *Senior Member, IEEE*

***Abstract*— This work proposes a compositional data-driven technique for the construction of finite Markov decision processes (MDPs) for large-scale stochastic networks with *unknown* mathematical models. Our proposed framework leverages dissipativity properties of subsystems and their finite MDPs using a notion of *stochastic storage functions* (SStF). In our data-driven scheme, we first build an SStF between each unknown subsystem and its data-driven finite MDP with a certified probabilistic confidence. We then derive dissipativity-type compositional conditions to construct a *stochastic bisimulation function* (SBF) between an interconnected network and its finite MDP using data-driven SStF of subsystems. Accordingly, we formally quantify the probabilistic distance between trajectories of an unknown large-scale stochastic network and those of its finite MDP with a guaranteed confidence. We illustrate the efficacy of our data-driven results over a room temperature network composing 100 rooms with unknown models.**

I. Introduction

Providing a formal analysis framework for large-scale stochastic networks to fulfill complex logic properties is generally very challenging. This is particularly due to (i) dealing with uncountable state/input sets with large dimensions, (ii) stochastic nature of dynamics, (iii) complex logic requirements, and (iv) lack of closed-form mathematical models in many real-world applications. To mitigate the aforesaid difficulties, one rewarding solution is to approximate the original (concrete) system by a finite MDP as a finite-state model. By establishing a similarity relation between each concrete system and its finite MDP using a notion of *stochastic simulation functions*, the probabilistic mismatch between two systems can be quantified within a guaranteed error bound.

There have been numerous studies, conducted in the past two decades, on the abstraction-based analysis of stochastic systems. Existing results encompass construction of (in)finite abstractions for *stochastic* dynamical systems with continuous state sets [1]–[4]. However, the main bottleneck of those techniques is *curse of dimensionality* problem due to discretizing state and input sets. *Compositional techniques* for constructing finite abstractions have then been proposed to alleviate the underlying state-explosion problem: one can build a finite abstraction for a large-dimensional system using finite abstractions of smaller subsystems [5]–[8].

Although the above-mentioned studies on constructing finite abstractions are comprehensive, unfortunately, they require knowing the mathematical model of the system. Accordingly, one cannot leverage those techniques for many practical scenarios with unknown models. Although *identification techniques* have been proposed to learn approximate models of unknown systems, obtaining a precise model is computationally very burdensome [9, and references herein]). In addition, even if a model can be identified using system identification techniques, the relation between the identified model and its finite abstraction should be still constructed. Consequently, the computational complexity exists in two levels of identifying the model and establishing the similarity relation. In this work, we develop a *direct* data-driven scheme, without performing any system identification, and construct finite abstractions together with their associated similarity relations by directly gathering data from trajectories of unknown concrete systems.

The original contribution here is to propose a compositional data-driven technique for constructing finite MDPs for large-scale stochastic control networks with unknown mathematical models. We leverage dissipativity properties of subsystems and their finite MDPs using a notion of *stochastic storage functions* (SStF). In our data-driven scheme, we recast conditions of SStF as a robust optimization program (ROP). By gathering samples from trajectories of each unknown subsystem, we then provide a scenario optimization program (SOP) for the original ROP. By quantifying the closeness between the optimal values of SOP and ROP, we build an SStF between each unknown subsystem and its data-driven finite MDP with a guaranteed probabilistic confidence. We then derive a dissipativity-type compositional condition to construct *stochastic bisimulation functions* (SBF), between an interconnected network and its finite MDP, using data-driven SStF of subsystems. Eventually, we quantify the probabilistic closeness between trajectories of an unknown interconnected network and its finite MDP with a guaranteed confidence level. We demonstrate the efficacy of our proposed data-driven results over a room temperature network composing 100 rooms with unknown models.

There has been a limited number of work on data-driven construction of symbolic models (in deterministic setting) and finite MDPs (in stochastic setting). Existing results include: data-driven abstraction of monotone systems with disturbances [10]; data-driven construction of symbolic abstractions via a probably approximately correct (PAC) approach [11]; data-driven construction of finite abstractions for verification of unknown systems [12]; data-driven construction of symbolic models for incrementally input-to-state stable systems [13]; and data-driven construction of finite MDPs for incrementally input-to-state stable systems [14]. In comparison, we propose here a compositional data-driven framework using dissipativity approach for constructing finite MDPs for large-scale interconnected networks, whereas the results in [10]–[14] are all tailored to monolithic systems. As a result, the approaches in [10]–[14] suffer from the

The author is with the School of Computing, Newcastle University, United Kingdom. Email: abolfazl.lavaei@newcastle.ac.uk.

979-8-3503-0123-6/23/$31.00 ©2023 IEEE

sample complexity problem and are not useful in practice when dealing with high-dimensional systems. In addition, the works [10]–[13] construct *symbolic* abstractions from unknown *deterministic* systems, whereas we develop here a data-driven technique for building *finite MDPs* for *stochastic* systems which is more challenging given the stochastic nature of unknown dynamics. Due to space limitations, we refer to the arXiv version [15] for the proofs of all statements.

II. Discrete-Time Stochastic Control Systems

A. Notation and Preliminaries

In this work, $\mathbb{R}, \mathbb{R}^+$, and $\mathbb{R}_0^+$, represent sets of real, positive, and non-negative real numbers, respectively. Symbols $\mathbb{N} := \{0, 1, 2, ...\}$ and $\mathbb{N}^+ = \{1, 2, ...\}$ denote, respectively, sets of non-negative and positive integers. A column vector, given N vectors $x_i \in \mathbb{R}^{n_i}$, is represented by $x = [x_1; \ldots; x_N]$. Given a set X, its power set is denoted by 2^X. We denote the minimum and maximum eigenvalues of a symmetric matrix P, respectively, by $\lambda_{\min}(P)$ and $\lambda_{\max}(P)$. Given any scalar $a \in \mathbb{R}$ and vector $x \in \mathbb{R}^n$, $|a|$ and $\|x\|$ represent, respectively, the absolute value and the Euclidean norm. For a matrix $P \in \mathbb{R}^{m\times n}$, $\|P\| := \sqrt{\lambda_{\max}(P^\top P)}$. We denote the supremum of a function $f : \mathbb{N} \to \mathbb{R}^n$ by $\|f\|_\infty :=$ (ess)sup$\{\|f(k)\|, k \geq 0\}$. Given a system Λ and a property φ, $\Lambda \vDash \varphi$ denotes that Λ fulfills φ.

Given a probability space $(\Omega, \mathcal{F}_\Omega, \mathbb{P}_\Omega)$, with Ω being a sample space, $\mathcal{F}_\Omega$ a sigma-algebra on Ω, and $\mathbb{P}_\Omega$ a probability measure, N-Cartesian product set of Ω and its associated product measure are denoted, respectively, by Ω^N and $\mathbb{P}^N$. A set X is Borel, denoted by $\mathcal{B}(X)$, if it is homeomorphic to a Borel subset of a Polish space, *i.e.*, a separable and metrizable space.

B. Discrete-Time Stochastic Control Systems

Here, we first formally define discrete-time stochastic control systems as the following.

Definition 2.1: A discrete-time stochastic control system (dt-SCS) is characterized by

$$\Lambda = (X, U, D, \varsigma, f), \tag{1}$$

where:

- $X \subseteq \mathbb{R}^n$ is a Borel state set;
- $U = \{\nu_1, \nu_2, \ldots, \nu_m\}$, with $\nu_i \in \mathbb{R}^{\bar m}, i \in \{1, \ldots, m\}$, is a discrete input set;
- $D \subseteq \mathbb{R}^p$ is a Borel disturbance set;
- ς is a sequence of independent-and-identically distributed (i.i.d.) random variables from the sample space Ω to a set $\mathcal{H}_\varsigma$, *i.e.* $\varsigma := \{\varsigma(k) : \Omega \to \mathcal{H}_\varsigma,\ k \in \mathbb{N}\}$;
- $f : X \times U \times D \times \mathcal{H}_\varsigma \to X$ is a transition map, which is assumed to be unknown.

The evolution of dt-SCS can be described by

$$\Lambda\colon x(k+1) = f(x(k), \nu(k), d(k), \varsigma(k)), \quad k \in \mathbb{N}, \tag{2}$$

for any $x \in X$, $\nu(\cdot) : \Omega \to U$, and $d(\cdot) : \Omega \to D$. The *state trajectory* of Λ under $\nu(\cdot), d(\cdot)$ starting from $x(0) = x_0$ is denoted by $x_{x_0\nu d}\colon \Omega \times \mathbb{N} \to X$.

Since the ultimate objective is to construct a finite MDP for an *interconnected* dt-SCS, we consider dt-SCS in (2) as a *subsystem* and present another definition for *interconnected* dt-SCS without disturbances d as a composition of individual dt-SCS with disturbances d.

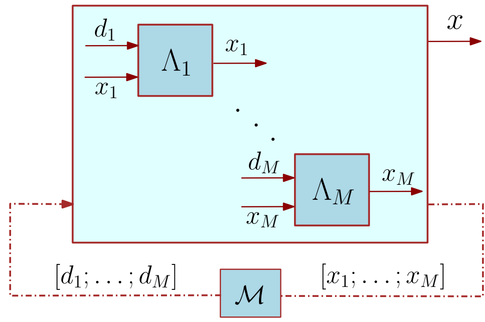

Fig. 1. Interconnected dt-SCS $\mathcal{I}(\Lambda_1, \ldots, \Lambda_M)$.

Definition 2.2: Consider $M \in \mathbb{N}^+$ dt-SCS $\Lambda_i = (X_i, U_i, D_i, f_i, \varsigma_i)$, $i \in \{1, \ldots, M\}$, with a matrix $\mathcal{M}$ as a coupling among them. An interconnection of Λ_i is characterized as $\Lambda = (X, U, f, \varsigma)$, represented by $\mathcal{I}(\Lambda_1, \ldots, \Lambda_M)$, where $X := \prod_{i=1}^M X_i$, $U := \prod_{i=1}^M U_i$, $f := [f_1; \ldots; f_M]$, and $\varsigma := [\varsigma_1; \ldots; \varsigma_M]$, such that:

$$\Big[d_1; \cdots; d_M\Big] = \mathcal{M}\Big[x_1; \cdots; x_M\Big]. \tag{3}$$

Such an interconnected dt-SCS is described by

$$\Lambda\colon x(k+1) = f(x(k), \nu(k), \varsigma(k)), \text{with } f : X \times U \times \mathcal{H}_\varsigma \to X. \tag{4}$$

An interconnected dt-SCS Λ is schematically depicted in Fig. 1.

C. Finite Markov Decision Processes

Here, we construct finite MDPs as finite-state approximations of dt-SCS. To this end, we first partition state and disturbance sets as $X = \cup_i \mathsf{X}_i$ and $D = \cup_i \mathsf{D}_i$, and then pick representative points $\hat x_i \in \mathsf{X}_i$ and $\hat d_i \in \mathsf{D}_i$ within those partitions sets as finite states and disturbances.

The dt-SCS in (1) can be *equivalently* considered as a continuous-space MDP $\Lambda = (X, U, D, \mathsf{T}_\mathsf{x})$ [16], with $\mathsf{T}_\mathsf{x} : \mathcal{B}(X) \times X \times U \times D \to [0, 1]$ being a conditional stochastic kernel that assigns to any $x \in X, \nu \in U, d \in D$, a probability measure $\mathsf{T}_\mathsf{x}(\cdot \,|\, x, \nu, d)$ such that for any set $\mathcal{X} \in \mathcal{B}(X)$:

$$\mathbb{P}\big\{x(k+1) \in \mathcal{X} \,\big|\, x(k), \nu(k), d(k)\big\} = \int_{\mathcal{X}} \mathsf{T}_\mathsf{x}(\mathsf{d}x(k+1) \,|\, x(k), \nu(k), d(k)).$$

One can *uniquely* determine the conditional stochastic kernel T_x using (ς, f) [16]. In the next definition, we formalize the construction of finite MDPs.

Definition 2.3: Consider a continuous-space MDP $\Lambda = (X, U, D, \mathsf{T}_\mathsf{x})$. The *finite MDP*, constructed from Λ, is characterized by $\hat\Lambda = (\hat X, U, \hat D, \hat{\mathsf{T}}_\mathsf{x})$, with $\hat X$ and $\hat D$ being discrete state and disturbance sets of $\hat\Lambda$ and

$$\hat{\mathsf{T}}_\mathsf{x}(x' \big| x, \nu, d) = \mathsf{T}_\mathsf{x}(\Xi(x') \big| x, \nu, d), \forall x, x' \in \hat X, \forall \nu \in U, \forall d \in D,$$

where $\Xi : X \rightarrow 2^X$. Equivalently, given a dt-SCS $\Lambda = (X, U, D, \varsigma, f)$, its constructed *finite MDP* can be characterized as [16]

$$\hat{\Lambda} = (\hat{X}, U, \hat{D}, \varsigma, \hat{f}),$$

where $\hat{f} : \hat{X} \times U \times \hat{D} \times \mathcal{H}_\varsigma \rightarrow \hat{X}$ is a transition function defined as

$$\hat{f}(\hat{x}, \nu, \hat{d}, \varsigma) = \mathcal{P}_x(f(\hat{x}, \nu, \hat{d}, \varsigma)), \tag{5}$$

and $\mathcal{P}_x : X \rightarrow \hat{X}$ is a quantization map with a *state discretization parameter* ρ fulfilling the following inequality:

$$\|\mathcal{P}_x(x) - x\| \leq \rho, \quad \forall x \in X. \tag{6}$$

III. Stochastic Storage and Bisimulation Functions

In this section, we aim at quantifying the probabilistic mismatch between trajectories of an interconnected dt-SCS and its finite MDP using a notion of *stochastic bisimulation functions*, as defined next.

Definition 3.1: Given an interconnected dt-SCS $\Lambda = (X, U, \varsigma, f)$ and its finite MDP $\hat{\Lambda} = (\hat{X}, U, \varsigma, \hat{f})$, a function $\mathcal{V} \colon X \times \hat{X} \rightarrow \mathbb{R}_0^+$ is a stochastic bisimulation function (SBF) between $\hat{\Lambda}$ and Λ, represented by $\hat{\Lambda} \cong_{\mathcal{V}} \Lambda$, if

$$\forall x \in X, \forall \hat{x} \in \hat{X}: \qquad \gamma\|x - \hat{x}\|^2 \leq \mathcal{V}(x, \hat{x}), \tag{7a}$$

$$\forall x \in X, \forall \hat{x} \in \hat{X}, \forall \nu \in U:$$

$$\mathbb{E}\Big[\mathcal{V}(f(x, \nu, \varsigma), \hat{f}(\hat{x}, \nu, \varsigma)) \,\big|\, x, \hat{x}, \nu\Big] \leq \alpha \mathcal{V}(x, \hat{x}) + \varpi, \tag{7b}$$

for some $\gamma \in \mathbb{R}^+$, $0 < \alpha < 1$, and $\varpi \in \mathbb{R}_0^+$, where $\mathbb{E}$ is the expected value associated to ς.

We now leverage SBF $\mathcal{V}$ and quantify the probabilistic mismatch between trajectories of an interconnected system and its finite MDP, as in the next theorem [8].

Theorem 3.2: Given an interconnected dt-SCS Λ and its finite MDP $\hat{\Lambda}$, let $\mathcal{V}$ be an SBF between $\hat{\Lambda}$ and Λ. Then the probabilistic closeness between state trajectories of dt-SCS (*i.e.* $x_{x_0\nu}(\cdot)$) and its finite MDP (*i.e.* $\hat{x}_{\hat{x}_0\nu}(\cdot)$) within a time horizon $\mathcal{T} \in \mathbb{N}$ can be quantified as

$$\mathbb{P}\left\{\sup_{0 \leq k \leq \mathcal{T}} \|x_{x_0\nu}(k) - \hat{x}_{\hat{x}_0\nu}(k)\| \geq \varepsilon \,\big|\, x_0, \hat{x}_0\right\} \leq \delta, \tag{8}$$

where

$$\delta := \begin{cases} 1-(1-\frac{\mathcal{V}(x_0,\hat{x}_0)}{\gamma\varepsilon^2})(1-\frac{\varpi}{\gamma\varepsilon^2})^{\mathcal{T}}, & \text{if } \gamma\varepsilon^2 \geq \frac{\varpi}{1-\alpha}, \\ (\frac{\mathcal{V}(x_0,\hat{x}_0)}{\gamma\varepsilon^2})\alpha^{\mathcal{T}} + (\frac{\varpi}{(1-\alpha)\gamma\varepsilon^2})(1-\alpha^{\mathcal{T}}), & \text{if } \gamma\varepsilon^2 < \frac{\varpi}{1-\alpha}, \end{cases}$$

with $\varepsilon \in \mathbb{R}^+$ being an arbitrary threshold. If $\varpi = 0$ in (7b), the closeness guarantee in (8) can be generalized to infinite horizons as

$$\mathbb{P}\left\{\sup_{0 \leq k < \infty} \|x_{x_0\nu}(k) - \hat{x}_{\hat{x}_0\nu}(k)\| \geq \varepsilon \,\big|\, x_0, \hat{x}_0\right\} \leq \frac{\mathcal{V}(x_0, \hat{x}_0)}{\gamma\varepsilon^2}.$$

In general, constructing SBF for large-scale interconnected networks is very expensive (if it is not impossible). To tackle this computational difficulty, we present a notion of *stochastic storage functions* for individual subsystems and propose, in Section VI, some compositional dissipativity conditions to construct an SBF for an interconnected network using SStF of subsystems.

Definition 3.3: Given a dt-SCS $\Lambda = (X, U, D, \varsigma, f)$ and its finite MDP $\hat{\Lambda} = (\hat{X}, U, \hat{D}, \varsigma, \hat{f})$, a function $\mathcal{S} \colon X \times \hat{X} \rightarrow \mathbb{R}_0^+$ is a stochastic storage function (SStF) between $\hat{\Lambda}$ and Λ, represented by $\hat{\Lambda} \cong_{\mathcal{S}} \Lambda$, if

$$\forall x \in X, \forall \hat{x} \in \hat{X}: \qquad \gamma\|x - \hat{x}\|^2 \leq \mathcal{S}(x, \hat{x}), \tag{9a}$$

$$\forall x \in X, \forall \hat{x} \in \hat{X}, \forall \nu \in U, \forall d \in D, \forall \hat{d} \in \hat{D}:$$

$$\mathbb{E}\Big[\mathcal{S}(f(x, \nu, d, \varsigma), \hat{f}(\hat{x}, \nu, \hat{d}, \varsigma)) \,\big|\, x, \hat{x}, \nu, d, \hat{d}\Big]$$

$$\leq \alpha\mathcal{S}(x, \hat{x}) + \varpi + \begin{bmatrix} d - \hat{d} \\ x - \hat{x} \end{bmatrix}^\top \underbrace{\begin{bmatrix} \mathcal{Z}^{11} & \mathcal{Z}^{12} \\ \mathcal{Z}^{21} & \mathcal{Z}^{22} \end{bmatrix}}_{\mathcal{Z}} \begin{bmatrix} d - \hat{d} \\ x - \hat{x} \end{bmatrix}, \tag{9b}$$

for some $\gamma \in \mathbb{R}^+$, $0 < \alpha < 1$, $\varpi \in \mathbb{R}_0^+$, and a symmetric matrix $\mathcal{Z}$ with partitions $\mathcal{Z}^{jj'}$, $j, j' \in \{1, 2\}$.

IV. Data-Driven Construction of SStF

In our data-driven framework, we consider SStF in the form of $\mathcal{S}(\kappa, x, \hat{x}) = \sum_{j=1}^{z} \kappa_j h_j(x, \hat{x})$ with basis functions $h_j(x, \hat{x})$ and unknown variables $\kappa = [\kappa_1; \ldots; \kappa_z] \in \mathbb{R}^z$. We now cast conditions (9a)-(9b) of SStF as a robust optimization program (ROP):

$$\text{ROP}\colon \begin{cases} \min\limits_{[\mathcal{Q};\psi]} & \psi, \\ \text{s.t.} & \max_j \Big\{\Upsilon_j(x, \hat{x}, \nu, d, \hat{d}, \mathcal{Q})\Big\} \leq \psi, \ j \in \{1, 2\}, \\ & \forall x \in X, \forall \hat{x} \in \hat{X}, \forall \nu \in U, \forall d \in D, \forall \hat{d} \in \hat{D}, \\ & \mathcal{Q} = [\gamma; \alpha; \varpi; \mathcal{Z}^{11}; \mathcal{Z}^{12}; \mathcal{Z}^{22}; \kappa_1; \ldots; \kappa_z], \\ & \gamma \in \mathbb{R}^+, \alpha \in (0, 1), \varpi \in \mathbb{R}_0^+, \mathcal{Z}^{jj'}, \psi \in \mathbb{R}, \end{cases} \tag{10}$$

where:

$$\Upsilon_1 = \gamma\|x - \hat{x}\|^2 - \mathcal{S}(\kappa, x, \hat{x}),$$

$$\Upsilon_2 = \mathbb{E}\left[\mathcal{S}(\kappa, f(x, \nu, d, \varsigma), \hat{f}(\hat{x}, \nu, \hat{d}, \varsigma)) \,\big|\, x, \hat{x}, \nu, d, \hat{d}\right]$$

$$- \alpha\mathcal{S}(\kappa, x, \hat{x}) - \varpi - \begin{bmatrix} d - \hat{d} \\ x - \hat{x} \end{bmatrix}^\top \begin{bmatrix} \mathcal{Z}^{11} & \mathcal{Z}^{12} \\ \mathcal{Z}^{21} & \mathcal{Z}^{22} \end{bmatrix} \begin{bmatrix} d - \hat{d} \\ x - \hat{x} \end{bmatrix}. \tag{11}$$

When ψ_R^*, the optimal value of ROP, is less than or equal to zero, it is straightforward to confirm that conditions (9a)-(9b) are met.

The ROP in (10) is not solvable due to unknown maps f, $\hat{f}$ appearing in Υ_2. To resolve this difficulty, we collect N i.i.d. data within $X \times D$, denoted by $(\bar{x}_i, \bar{d}_i)_{i=1}^N$. We now propose a scenario optimization program (SOP), with an optimal value ψ_N^*, associated to the original ROP:

$$\text{SOP}_N\colon \begin{cases} \min\limits_{[\mathcal{Q};\psi]} & \psi, \\ \text{s.t.} & \max_j \Big\{\Upsilon_j(\bar{x}_i, \hat{x}, \nu, \bar{d}_i, \hat{d}, \mathcal{Q})\Big\} \leq \psi, \ j \in \{1, 2\}, \\ & \forall \bar{x}_i \in X, \forall \bar{d}_i \in D, \forall i \in \{1, \ldots, N\}, \\ & \forall \hat{x} \in \hat{X}, \forall \hat{d} \in \hat{D}, \forall \nu \in U, \\ & \mathcal{Q} = [\gamma; \alpha; \varpi; \mathcal{Z}^{11}; \mathcal{Z}^{12}; \mathcal{Z}^{22}; \kappa_1; \ldots; \kappa_z], \\ & \gamma \in \mathbb{R}^+, \alpha \in (0, 1), \varpi \in \mathbb{R}_0^+, \mathcal{Z}^{jj'}, \psi \in \mathbb{R}. \end{cases} \tag{12a}$$

We can now replace the unknown function $f(\bar{x}_i, \nu, \bar{d}_i, \varsigma)$ in Υ_2 by observing the one-step transition of dt-SCS starting from $\bar{x}_i$ under ν and $\bar{d}_i$. Regarding $\hat{f}(\hat{x},\nu,\hat{d},\varsigma)$ in Υ_2, we begin by initializing the unknown model at $\hat{x}$ under ν and $\hat{d}$ to compute $f(\hat{x},\nu,\hat{d},\varsigma)$. With a state discretization parameter ρ in place, we then compute $\hat{f}(\hat{x},\nu,\hat{d},\varsigma)$ as the point nearest to $f(\hat{x},\nu,\hat{d},\varsigma)$, where condition (6) is satisfied.

By proposing SOP (12a), the problem of unknown maps $f,\hat{f}$ in ROP (10) got solved. However, the proposed SOP in (12a) is not still tractable since there is no closed-form solution for computing the expected value in Υ_2. To resolve this issue, we propose another version of SOP, denoted by SOP_ς, by computing the expected value using its empirical approximation:

$$\text{SOP}_\varsigma\colon\begin{cases}\min\limits_{[\mathcal{Q};\psi]} \psi,\\ \text{s.t.}\quad \max\Big\{\Upsilon_1(\bar{x}_i,\hat{x},\nu,\bar{d}_i,\hat{d},\mathcal{Q}),\bar{\Upsilon}_2(\bar{x}_i,\hat{x},\nu,\bar{d}_i,\hat{d},\mathcal{Q})\Big\}\leq\psi,\\ \qquad \forall \bar{x}_i\in X,\forall\bar{d}_i\in D,\forall i\in\{1,\dots,N\},\\ \qquad \forall\hat{x}\in\hat{X},\forall\hat{d}\in\hat{D},\forall\nu\in U,\\ \qquad \mathcal{Q}=[\gamma;\alpha;\varpi;\mathcal{Z}^{11};\mathcal{Z}^{12};\mathcal{Z}^{22};\kappa_1;\dots;\kappa_z],\\ \qquad \gamma\in\mathbb{R}^+,\alpha\in(0,1),\varpi\in\mathbb{R}_0^+,\mathcal{Z}^{jj'},\kappa_i,\psi\in\mathbb{R},\end{cases}\tag{12b}$$

with

$$\bar{\Upsilon}_2=\frac{1}{L}\sum_{q=1}^{L}\mathcal{S}(\kappa,f(\bar{x}_i,\nu,\bar{d}_i,\varsigma_q),\hat{f}(\hat{x},\nu,\hat{d},\varsigma_q))-\alpha\mathcal{S}(\kappa,x_i,\hat{x})$$
$$-\varpi+\mu-\begin{bmatrix}d-\hat{d}\\x-\hat{x}\end{bmatrix}^\top\begin{bmatrix}\mathcal{Z}^{11}&\mathcal{Z}^{12}\\\mathcal{Z}^{21}&\mathcal{Z}^{22}\end{bmatrix}\begin{bmatrix}d-\hat{d}\\x-\hat{x}\end{bmatrix},$$

where $\mu\in\mathbb{R}_0^+$ and $L\in\mathbb{N}_0^+$ are the approximation error and required number of realizations, respectively. We denote the optimal value of SOP_ς by ψ_ς^*.

We now leverage Chebyshev's inequality [17] to construct a relation between solutions of SOP_ς and SOP_N with a guaranteed confidence level $\beta_1\in(0,1]$.

Lemma 4.1: Let $\mathcal{S}$ be a feasible solution for SOP_ς in (12b). For a desired confidence level $\beta_1\in(0,1]$ and an approximation error $\mu\in\mathbb{R}_0^+$, one has

$$\mathbb{P}\Big\{\mathcal{S}(\kappa,x,\hat{x})\models\text{SOP}_N\Big\}\geq 1-\beta_1,$$

provided that $L \geq \frac{\mathcal{C}}{\beta_1\mu^2}$, where $\text{Var}\big[\mathcal{S}(\kappa,f(x,\nu,d,\varsigma),\hat{f}(\hat{x},\nu,\hat{d},\varsigma))\big]\leq\mathcal{C}$, $\forall x\in X,\forall\hat{x}\in\hat{X},\forall\nu\in U,\forall d\in D,\forall\hat{d}\in\hat{D}$.

Remark 4.2: As it can be observed, there is a bilinearity between unknown variables κ and α in Υ_2. To resolve it, we consider α in a discrete set as $\alpha\in\{\alpha_1,\dots,\alpha_l\}$. The cardinality l is then taken into account when determining the necessary amount of data to solve SOP, as shown in (13).

V. Data-Driven Guarantee for SBF Construction

Here, we aim at constructing an SStF between each unknown subsystem and its finite MDP with a certified confidence level by establishing a probabilistic relation between optimal values of SOP_ς and ROP [18].

Theorem 5.1: Consider unknown dt-SCS Λ in (1). Let Υ_1 and Υ_2 be Lipschitz continuous, with respect to x and (d,x) with Lipschitz constants, respectively, $\mathscr{L}_1,\mathscr{L}_{2_t}$, for given α_t where $t\in\{1,\dots,l\}$, and any $\nu\in U$. Consider the SOP_ς in (12b) with ψ_ς^*, $\mathcal{Q}^*=[\gamma^*;\varpi^*;\mathcal{Z}^{11*};\mathcal{Z}^{12*};\mathcal{Z}^{22*};\kappa_1^*;\dots;\kappa_z^*]$, and

$$N(\varepsilon_{2_t},\beta_2){:=}\min\Big\{N\in\mathbb{N}\,\Big|\sum_{t=1}^{l}\sum_{i=0}^{c-1}\binom{N}{i}\varepsilon_{2_t}^i(1-\varepsilon_{2_t})^{N-i}\leq\beta_2\Big\},\tag{13}$$

where $\beta_2,\varepsilon_{2_t}\in[0,1]$ for any $t\in\{1,\dots,l\}$, with c,l being, respectively, number of unknown variables in SOP_ς, and cardinality of finite set of α. If

$$\psi_\varsigma^*+\max_t\mathscr{L}_{\Upsilon_t}\eta^{-1}(\varepsilon_{2_t})\leq 0,\tag{14}$$

with $\mathscr{L}_{\Upsilon_t}:=\max\big\{\mathscr{L}_1,\mathscr{L}_{2_t}\big\}$ for any $t\in\{1,\dots,l\}$, and $\eta(r):\mathbb{R}_{\geq 0}\to[0,1]$, which depends on the geometry of $X\times D$ and the sampling distribution, then the data-driven $\mathcal{S}$ is an SStF between $\hat{\Lambda}$ and Λ, with a confidence of $1-\beta$ with $\beta=\beta_1+\beta_2$, *i.e.,*

$$\mathbb{P}^N\big\{\hat{\Lambda}\cong_{\mathcal{S}}\Lambda\big\}\geq 1-\beta_1-\beta_2,$$

where $\beta_1\in(0,1]$ is as in Lemma 4.1.

In the next lemma, we compute the function η which is required for checking condition (14).

Lemma 5.2: The function η in (14) fulfills the following condition [18, Proposition 3.8]:

$$\eta(r)\leq\mathbb{P}\big[\mathbb{B}_r(x,d)\big],\qquad\forall r\in\mathbb{R}_{\geq 0},\forall(x,d)\in X\times D,\tag{15}$$

with $\mathbb{B}_r(c)\subset X\times D$ being an open ball with center c and radius r. By gathering data from an $(n+p)$-dimensional *hyper-rectangle* uncertainty set $X\times D$ with a *uniform* distribution, the function η in (15) is then quantified as

$$\eta(r)=\frac{\text{Vol}(\mathbb{B}_r(x,d))}{2^{n+p}\text{Vol}(X\times D)}=\frac{\frac{\pi^{\frac{n+p}{2}}}{\Gamma(\frac{n+p}{2}+1)}r^{n+p}}{2^{n+p}\text{Vol}(X\times D)}$$
$$=\frac{\pi^{\frac{n+p}{2}}r^{n+p}}{2^{n+p}\Gamma(\frac{n+p}{2}+1)\text{Vol}(X\times D)},\tag{16}$$

with $\text{Vol}(\cdot)$ and Γ being volume set and Gamma function, respectively.

To assess condition (14), it is necessary to determine $\mathscr{L}_{\Upsilon_t}$. The following lemmas present computations of $\mathscr{L}_{\Upsilon_t}$ for both linear and nonlinear stochastic systems

Lemma 5.3: Given a linear dt-SCS $x(k+1)=Ax(k)+B\nu(k)+Ed(k)+\varsigma(k)$, let $(x-\hat{x})^\top P(x-\hat{x})$ be an SStF with a positive-definite matrix $P\in\mathbb{R}^{n\times n}$. Then $\mathscr{L}_{\Upsilon_t}$ is computed as $\mathscr{L}_{\Upsilon_t}=\max\big\{\mathscr{L}_1,\mathscr{L}_{2_t}\big\}$, with

$$\mathscr{L}_1=4s_1(\lambda_{\min}(P)+\lambda_{\max}(P)),$$
$$\mathscr{L}_{2_t}=2\lambda_{\max}(P)(2\mathcal{Y}_1^2s_1+2\mathcal{Y}_1\mathcal{Y}_2s_2+2\mathcal{Y}_1\mathcal{Y}_3s_3+\mathcal{Y}_1\rho+\mathcal{Y}_3\rho$$
$$+2\mathcal{Y}_3^2s_3+2\mathcal{Y}_3\mathcal{Y}_2s_2+2\mathcal{Y}_3\mathcal{Y}_1s_1+2s_1\alpha_t)+2s_4s_5,$$

where $\|A\|\leq\mathcal{Y}_1$, $\|B\|\leq\mathcal{Y}_2$, $\|E\|\leq\mathcal{Y}_3$, $\|x\|\leq s_1$ for any $x\in X$, $\|\nu\|\leq s_2$ for any $\nu\in U$, $\|d\|\leq s_3$ for any $d\in D$, $\|[d-\hat{d};x-\hat{x}]\|\leq s_4$ for any $x\in X$, $\hat{x}\in\hat{X}$, $d\in D$, $\hat{d}\in\hat{D}$, and $\|\mathcal{Z}\|=s_5$.

We now compute $\mathscr{L}_{\Upsilon_t}$ for *nonlinear* stochastic systems.

Lemma 5.4: Given a nonlinear dt-SCS $x(k+1)=f(x(k),\nu(k),d(k))+\varsigma(k)$, let $(x-\hat{x})^\top P(x-\hat{x})$ be an SStF with a positive-definite matrix $P \in \mathbb{R}^{n\times n}$. Then $\mathscr{L}_{\Upsilon_t}$ is computed as $\mathscr{L}_{\Upsilon_t} = \max\{\mathscr{L}_1, \mathscr{L}_{2_t}\}$, with

$$\begin{aligned}\mathscr{L}_1 &= 4s_1(\lambda_{\min}(P) + \lambda_{\max}(P)),\\ \mathscr{L}_{2_t} &= 2\lambda_{\max}(P)(2\mathcal{Y}_f\mathcal{Y}_x + \mathcal{Y}_x\rho + 2\mathcal{Y}_f\mathcal{Y}_d + \mathcal{Y}_d\rho + 2s_1\alpha_t)\\ &\quad + 2s_4s_5,\end{aligned}$$

where $\|f(x,\nu,d)\| \leq \mathcal{Y}_f$, $\|\partial_x f(x,\nu,d)\| \leq \mathcal{Y}_x$, $\|\partial_d f(x,\nu,d)\| \leq \mathcal{Y}_d$, $\|x\| \leq s_1$ for any $x \in X$, $\|[d-\hat{d}; x-\hat{x}]\| \leq s_4$ for any $x \in X$, $\hat{x} \in \hat{X}$, $d \in D$, $\hat{d} \in \hat{D}$, and $\|\mathcal{Z}\| = s_5$.

A. Data-Driven Finite MDPs via Maximum Likelihood Estimation

Here, we construct finite MDPs from data by estimating parameters of the probability distribution via maximum likelihood estimation (MLE) [19]. If the underlying stochasticity has a Gaussian distribution, its mean and standard deviation can be estimated via MLE as

$$\hat{\mu}_{\hat{N}} = \frac{1}{\hat{N}}\sum_{j=1}^{\hat{N}}\tilde{x}_j, \quad \hat{\sigma}^2_{\hat{N}} = \frac{1}{\hat{N}-1}\sum_{j=1}^{\hat{N}}(\tilde{x}_j - \hat{\mu}_{\hat{N}})^2,$$

where $\hat{\mu}_{\hat{N}}, \hat{\sigma}_{\hat{N}}$ are the *empirical* mean and standard deviation given $\hat{N}$ sampled data. Additionally, MLE approach can be used to estimate parameters of *any arbitrary* probability distributions. We then use the estimated parameters from MLE method and construct a finite MDP via the results of Section II-C. Although it is possible to provide an asymptotic confidence bound for MLE using Fisher information [20], we leave it to a future work for the sake of an easier presentation.

VI. Compositional Construction of SBF for Interconnected dt-SCS

Here, we propose a compositional dissipativity approach to build an SBF for an interconnected network using SStF of individual subsystems. The constructed SBF is then utilized to compute the probabilistic mismatch between trajectories of the interconnected system Λ and its finite MDP $\hat{\Lambda}$, as presented in Theorem 3.2.

Theorem 6.1: Consider an interconnected dt-SCS $\Lambda = \mathcal{I}(\Lambda_1,\ldots,\Lambda_M)$ composed of $M \in \mathbb{N}_0^+$ subsystems Λ_i. Let there exist an SStF between each subsystem Λ_i and its finite MDP $\hat{\Lambda}_i$ with a confidence of $1-\beta$, with $\beta = \beta_1 + \beta_2$, as in Theorem 5.1. Then

$$\mathcal{V}(\kappa, x, \hat{x}) := \sum_{i=1}^{M}\mathcal{S}_i(\kappa_i, x_i, \hat{x}_i), \tag{17}$$

is an SBF between $\hat{\Lambda} = \mathcal{I}(\hat{\Lambda}_1,\ldots,\hat{\Lambda}_M)$ and $\Lambda = \mathcal{I}(\Lambda_1,\ldots,\Lambda_M)$ with a confidence of $1-\sum_{i=1}^{M}\beta_i$, where $\beta_i = \beta_{1_i} - \beta_{2_i}$, if

$$\begin{bmatrix}\mathcal{M}\\ \mathbb{I}\end{bmatrix}^\top \mathcal{Z}_{cmp}\begin{bmatrix}\mathcal{M}\\ \mathbb{I}\end{bmatrix} \preceq 0, \tag{18}$$

$$\text{with}\quad \mathcal{Z}_{cmp} := \begin{bmatrix}\mathcal{Z}_1^{11} & & & \mathcal{Z}_1^{12} & & \\ & \ddots & & & \ddots & \\ & & \mathcal{Z}_M^{11} & & & \mathcal{Z}_M^{12}\\ \mathcal{Z}_1^{21} & & & \mathcal{Z}_1^{22} & & \\ & \ddots & & & \ddots & \\ & & \mathcal{Z}_M^{21} & & & \mathcal{Z}_M^{22}\end{bmatrix}.$$

VII. Case Study: Room Temperature Network

We showcase our data-driven results using a room temperature network consisting of 100 rooms, each with unknown models, interconnected in a circular topology, and equipped with cooling systems. The temperature dynamics, denoted as $x(\cdot)$, can be described through the following interconnected network [21]:

$$\Lambda\colon x(k+1) = Ax(k) + \theta T_c\nu(k) + \digamma T_E + \varsigma(k),$$

where the matrix A has diagonal entries $a_{ii} = 1 - 2\aleph - \digamma - \theta\nu_i(k)$, $i \in \{1,\ldots,M\}$, off-diagonal entries $a_{i,i+1} = a_{i+1,i} = a_{1,M} = a_{M,1} = \aleph$, $i \in \{1,\ldots,M-1\}$, and other entries being zero. Symbols $\aleph$, $\digamma$, and θ are thermal factors between rooms $i \pm 1$ and i, the outside environment and the room i, and the cooler and the room i, respectively. In addition, $x(k) = [x_1(k);\ldots;x_M(k)]$, $x(k) = [\varsigma_1(k);\ldots;\varsigma_M(k)]$, $T_E = [T_{e_1};\ldots;T_{e_M}]$, with $T_{e_i} = -1\,{}^\circ C$, $\forall i \in \{1,\ldots,M\}$, being the outside temperatures. The cooler temperature is $T_c = 5\,{}^\circ C$ and the control input is $\nu \in \{0, 0.05, 0.1, 0.15, 0.2\}$. Now by characterizing each individual room as

$$\begin{aligned}\Lambda_i\colon x_i(k+1) &= a_{ii}x_i(k) + \aleph(d_{i-1}(k) + d_{i+1}(k))\\ &\quad + \theta T_c\nu_i(k) + \digamma T_{e_i} + \varsigma_i(k),\end{aligned} \tag{19}$$

where $d_0 = d_M, d_{M+1} = d_1$, one has $\Lambda = \mathcal{I}(\Lambda_1,\ldots,\Lambda_M)$, with a coupling matrix $\mathcal{M}$ as $\bar{m}_{i,i+1} = \bar{m}_{i+1,i} = \bar{m}_{1,M} = \bar{m}_{M,1} = 1$, $i \in \{1,\ldots,M-1\}$, and other entries being zero. We assume the model of each room is unknown to us. The main target is to compositionally construct a finite MDP as well as a data-driven SBF via solving SOP (12b). Accordingly, we utilize the data-driven finite MDP and synthesize controllers regulating the temperature of each room in a safe set $X_i = [-0.5, 0.5]$ with a guaranteed probabilistic confidence.

We consider our SStF as $\mathcal{S}_i(\kappa_i, x_i, \hat{x}_i) = \kappa_{1_i}(x_i - \hat{x}_i)^4 + \kappa_{2_i}(x_i - \hat{x}_i)^2 + \kappa_{3_i}$. We also fix $\varepsilon_{t_i} = 0.025$, $\beta_{2_i} = 10^{-4}$, and $\rho_i = 0.05$, a-priori. According to (13), we compute $N_i = 911$ required for solving SOP in (12b). We also fix $\mu_i = 0.1$, $\beta_{1_i} = 10^{-4}$ and compute $L_i = 643$ according to Lemma 4.1. By solving SOP (12b) with N_i, L_i, we obtain the corresponding decision variables as

$$\begin{aligned}&\mathcal{S}_i(\kappa_i, x_i, \hat{x}_i) = 0.11(x_i - \hat{x}_i)^4 + 0.14(x_i - \hat{x}_i)^2 + 143,\\ &\mathcal{Z}_i^{11} = 0.001, \mathcal{Z}_i^{22} = -0.01, \mathcal{Z}_i^{12} = \mathcal{Z}_i^{21} = 0,\\ &\gamma_i^* = 141, \varpi_i^* = 0.42, \psi_{\varsigma_i}^* = -0.3019,\end{aligned} \tag{20}$$

with a fixed $\alpha_i = 0.99$. We now compute $\mathscr{L}_{\Upsilon_{t_i}} = 0.8$ according to Lemma 5.4. We also compute $\eta^{-1}(\varepsilon_{2_{t_i}})$ according to Lemma 5.2 as $\eta^{-1}(\varepsilon_{2_{t_i}}) = 0.362$. Since $\psi_{\varsigma_i}^* + \max_t \mathscr{L}_{\mathcal{H}_{t_i}}\eta^{-1}(\varepsilon_{2_{t_i}}) = -11 \times 10^{-3} \leq 0$, the constructed data-driven $\mathcal{S}_i$ is an SStF between each unknown room Λ_i and its finite MDP $\hat{\Lambda}_i$, with a confidence of at least $1 - \beta_{1_i} - \beta_{2_i} = 1 - 2\times 10^{-4}$.

We now construct an SBF for the interconnected rooms via SStF of individual rooms, constructed from data. By leveraging $\mathcal{Z}_i$ as in (20), the matrix $\mathcal{Z}_{cmp}$ is reduced to

$$\mathcal{Z}_{cmp} = \begin{bmatrix} 0.001\mathbb{I}_{100} & 0 \\ 0 & -0.01\mathbb{I}_{100} \end{bmatrix},$$

and compositionality condition (18) is reduced to

$$\begin{bmatrix} \mathcal{M} \\ \mathbb{I}_{100} \end{bmatrix}^\top \mathcal{Z}_{cmp} \begin{bmatrix} \mathcal{M} \\ \mathbb{I}_{100} \end{bmatrix} = 0.001\mathbb{I}_{100}\mathcal{M}^\top\mathcal{M} - 0.01\mathbb{I}_{100} \preceq 0.$$

Hence, one can certify that $\mathcal{V}(\kappa, x, \hat{x}) = \sum_{i=1}^{100}\{\mathcal{S}_i(\kappa_i, x_i, \hat{x}_i)\} = \sum_{i=1}^{100}\{0.11(x_i - \hat{x}_i)^4 + 0.14(x_i - \hat{x}_i)^2 + 143\}$ is an SBF between the interconnected rooms Λ and its finite MDP $\hat{\Lambda}$ with $\gamma = 14100, \alpha = 0.99, \varpi = 42$, and a confidence of $1 - \sum_{i=1}^{100}\beta_{1_i} - \sum_{i=1}^{100}\beta_{2_i} = 98\%$. Hence, by employing the results of Theorems 3.2 and 6.1, we guarantee that the mismatch between state trajectories of Λ and $\hat{\Lambda}$ remains within $\varepsilon = 0.5$ during $\mathcal{T} = 5$ (45 minutes) with a probability of 95% and a confidence of 98%.

Let us now synthesize a controller for Λ via its data-driven finite MDP $\hat{\Lambda}$, constructed via the MLE approach with $\hat{N} = 10^5$, such that the controller regulates state of each room within $[-0.5, 0.5]$. To do so, we first synthesize a controller for each abstract room $\hat{\Lambda}_i$ via AMYTISS [22] and then refine it back over unknown original room Λ_i. Accordingly, the overall controller for the network would be a vector whose entries are controllers for individual rooms. Closed-loop trajectories of a representative room with several noise realizations are depicted in Fig. 2. As observed, all trajectories respect the safety specification.

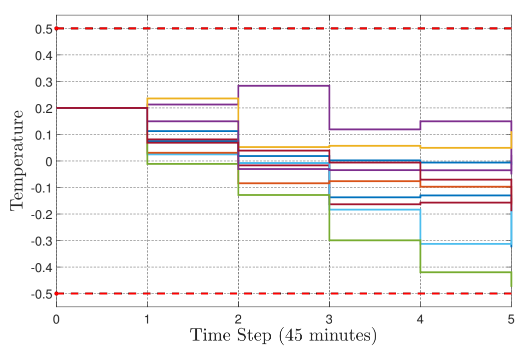

Fig. 2. Closed-loop state trajectories of an unknown representative room with several noise realizations.

VIII. Conclusion

In this work, we developed a compositional data-driven technique using dissipativity reasoning for constructing finite MDPs for large-scale stochastic networks with unknown mathematical models. The main goal was to leverage stochastic bisimulation functions (SBF) and quantify the closeness between an unknown original network and its data-driven finite MDP, while proposing a certified probabilistic confidence. In our proposed scheme, we first constructed a stochastic storage function between each unknown subsystem and its data-driven finite MDP with an a-priori confidence level. We then provided dissipativity-type compositional conditions to construct an SBF for an unknown interconnected network using its data-driven SStF of subsystems. We verified our results over a room temperature network composing 100 rooms with unknown dynamics.

References

[1] A. Abate, M. Prandini, J. Lygeros, and S. Sastry, "Probabilistic reachability and safety for controlled discrete-time stochastic hybrid systems," *Automatica*, vol. 44, no. 11, pp. 2724–2734, 2008.

[2] A. A. Julius and G. J. Pappas, "Approximations of stochastic hybrid systems," *IEEE Transactions on Automatic Control*, vol. 54, no. 6, pp. 1193–1203, 2009.

[3] M. Zamani, P. Mohajerin Esfahani, R. Majumdar, A. Abate, and J. Lygeros, "Symbolic control of stochastic systems via approximately bisimilar finite abstractions," *IEEE Transactions on Automatic Control*, vol. 59, no. 12, pp. 3135–3150, 2014.

[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] E. M. Hahn, A. Hartmanns, H. Hermanns, and J.-P. Katoen, "A compositional modelling and analysis framework for stochastic hybrid systems," *Formal Methods in System Design*, vol. 43, no. 2, pp. 191–232, 2013.

[6] A. Lavaei and E. Frazzoli, "Scalable synthesis of finite MDPs for large-scale stochastic switching systems," in *2022 IEEE 61st Conference on Decision and Control (CDC)*, 2022, pp. 7510–7515.

[7] A. Nejati and M. Zamani, "Compositional construction of finite MDPs for continuous-time stochastic systems: A dissipativity approach," *IFAC-PapersOnLine*, vol. 53, no. 2, pp. 1962–1967, 2020.

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

[9] Z. Hou and Z. Wang, "From model-based control to data-driven control: Survey, classification and perspective," *Information Sciences*, vol. 235, pp. 3–35, 2013.

[10] A. Makdesi, A. Girard, and L. Fribourg, "Efficient data-driven abstraction of monotone systems with disturbances," *IFAC-PapersOnLine*, vol. 54, no. 5, pp. 49–54, 2021.

[11] A. Devonport, A. Saoud, and M. Arcak, "Symbolic abstractions from data: A PAC learning approach," in *Proceedings of the 60th IEEE Conference on Decision and Control*, 2021, pp. 599–604.

[12] R. Coppola, A. Peruffo, and M. Mazo Jr, "Data-driven abstractions for verification of deterministic systems," *arXiv:2211.01793*, 2022.

[13] A. Lavaei and E. Frazzoli, "Data-driven synthesis of symbolic abstractions with guaranteed confidence," *IEEE Control Systems Letters*, vol. 7, pp. 253–258, 2022.

[14] A. Lavaei, S. Soudjani, E. Frazzoli, and M. Zamani, "Constructing MDP abstractions using data with formal guarantees," *IEEE Control Systems Letters*, vol. 7, pp. 460–465, 2022.

[15] A. Lavaei, "MDP abstractions from data: Large-scale stochastic networks," *arXiv:2309.07941*, 2023.

[16] O. Kallenberg, *Foundations of modern probability*. Springer-Verlag, New York, 1997.

[17] J. G. Saw, M. C. Yang, and T. C. Mo, "Chebyshev inequality with estimated mean and variance," *The American Statistician*, vol. 38, no. 2, pp. 130–132, 1984.

[18] P. Mohajerin Esfahani, T. Sutter, and J. Lygeros, "Performance bounds for the scenario approach and an extension to a class of non-convex programs," *IEEE Transactions on Automatic Control*, vol. 60, no. 1, pp. 46–58, 2014.

[19] I. J. Myung, "Tutorial on maximum likelihood estimation," *Journal of mathematical Psychology*, vol. 47, no. 1, pp. 90–100, 2003.

[20] L. Le Cam, *Asymptotic methods in statistical decision theory*. Springer Science & Business Media, 2012.

[21] P. J. Meyer, A. Girard, and E. Witrant, "Compositional abstraction and safety synthesis using overlapping symbolic models," *IEEE Transactions on Automatic Control*, 2017.

[22] A. Lavaei, M. Khaled, S. Soudjani, and M. Zamani, "AMYTISS: Parallelized automated controller synthesis for large-scale stochastic systems," in *Proceedings of the 32nd International Conference on Computer-Aided Verification (CAV), Lecture Notes in Computer Science 12225*, 2020, pp. 461–474.