



Delft University of Technology

Uncertainty Propagation in Stochastic Systems via Mixture Models with Error Quantification

Figueiredo, Eduardo; Patane, Andrea; Lahijanian, Morteza; Laurenti, Luca

DOI

[10.1109/CDC56724.2024.10886416](https://doi.org/10.1109/CDC56724.2024.10886416)

Publication date

2025

Document Version

Final published version

Published in

Proceedings of the IEEE 63rd Conference on Decision and Control, CDC 2024

Citation (APA)

Figueiredo, E., Patane, A., Lahijanian, M., & Laurenti, L. (2025). Uncertainty Propagation in Stochastic Systems via Mixture Models with Error Quantification. In *Proceedings of the IEEE 63rd Conference on Decision and Control, CDC 2024* (pp. 328-335). (Proceedings of the IEEE Conference on Decision and Control). IEEE. <https://doi.org/10.1109/CDC56724.2024.10886416>

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.

Green Open Access added to TU Delft Institutional Repository

'You share, we take care!' - Taverne project

<https://www.openaccess.nl/en/you-share-we-take-care>

Otherwise as indicated in the copyright section: the publisher is the copyright holder of this work and the author uses the Dutch legislation to make this work public.

Uncertainty Propagation in Stochastic Systems via Mixture Models with Error Quantification

Eduardo Figueiredo, Andrea Patane, Morteza Lahijanian, and Luca Laurenti

Abstract—Uncertainty propagation in non-linear dynamical systems has become a key problem in various fields including control theory and machine learning. In this work, we focus on discrete-time non-linear stochastic dynamical systems. We present a novel approach to approximate the distribution of the system over a given finite time horizon with a mixture of distributions. The key novelty of our approach is that it not only provides tractable approximations for the distribution of a non-linear stochastic system but also comes with formal guarantees of correctness. In particular, we consider the Total Variation (TV) distance to quantify the distance between two distributions and derive an upper bound on the TV between the distribution of the original system and the approximating mixture distribution derived from our framework. We show that in various cases of interest, including in the case of Gaussian noise, the resulting bound can be efficiently computed in closed form. This allows us to quantify the correctness of the approximation and to optimize the parameters of the resulting mixture distribution to minimize such distance. The effectiveness of our approach is illustrated on several benchmarks from the control community.

I. INTRODUCTION

Modern autonomous systems are becoming increasingly complex, commonly including non-linearities, data-driven components, and uncertain dynamics. The uncertainty affecting these systems can come from various sources including inherent probabilistic dynamics, lack of knowledge on parts of the model [1], and interactions with an uncertain environment [2]. Consequently, when autonomous systems need to be employed for safety-critical applications, where failure may translate into costly or deadly accidents, it becomes paramount to consider models that include non-linear stochastic dynamics. Unfortunately, for non-linear stochastic dynamical models, the distribution of the system at a given time cannot be computed in closed-form [3]. This leads to the main research question of this paper: Can we develop an efficient framework for uncertainty propagation in non-linear stochastic dynamical systems with formal guarantees of correctness?

Various methods have been proposed to propagate uncertainty in non-linear stochastic systems [4]. These either consider analytical approximation methods, such as Taylor expansions or Stirling's interpolation [5], [6], to approximate the one-step dynamics of the systems, or consider numerical approximation of the various integrals [7], [8] and/or only

propagate discrete samples from the system distribution [9], [10]. However, none of these methods comes with error bounds on the approximation error. An exception is [11], which is however only limited to Gaussian processes with additive noise and can only guarantee that with high probabilities the system trajectories will stay close to the mean. An alternative approach proposed in [12] is to rely on moment matching to propagate the moments of non-linear stochastic dynamical systems with non-linearities that are either polynomial or trigonometric functions. However, in practice, it is often required to propagate the full distribution, and not just the moments, e.g., for chance constraint satisfaction in planning problems [2]. Additionally, we should stress that, because of their universal approximation properties, various papers already attempted to use mixture distributions, and Gaussian Mixture Models (GMMs) in particular, to approximate the distribution of a stochastic dynamical system [13], [14]. However, these works are mostly based on heuristics on how to build the mixture and come with no guarantees of correctness.

In this work, given an error threshold $\delta > 0$ and a time horizon T , we develop a framework that for each time step $t \leq T$ builds a mixture distribution approximation of a given non-linear stochastic dynamical system, which is guaranteed to be δ -close to the distribution of the system at that time. To quantify the closeness between two distributions we consider the Total Variation (TV) distance, defined as the largest absolute difference between the probabilities that two probability distributions assign to an event [15]. Interestingly, we show that by appropriately picking the components of the approximating mixture distribution, an upper bound for its TV from the unknown distribution of the system at time t can be computed in closed form. Using this bound, we then propose a refinement algorithm that iteratively builds mixture models by efficiently placing distributions in regions of the state space by minimizing the TV bound. We empirically illustrate the efficacy of our framework in various benchmarks.

In summary, the main contributions of this work are: i) a novel framework for computing a mixture distribution approximating the distribution of a stochastic dynamical system at a given time, ii) formal bounds of correctness in terms of the TV distance and a proof of convergence of the approximating error to 0, iii) an adaptive grid refinement algorithm for bound optimization, and iv) empirical evaluation of the efficacy of our framework on four case studies including a Dubin's car model and an application on a planning problem.

E. Figueiredo and L. Laurenti are with the Delft Center for Systems and Control, TU Delft. Corresponding author's email: e.figueiredo@tudelft.nl

M. Lahijanian is with the Dept. of Aerospace Eng. Sciences and Computer Science, CU Boulder

A. Patane is with the School of Computer Science and Statistics, Trinity College Dublin

II. PROBLEM FORMULATION

A. Notation

Given a Borel measurable space $\mathcal{X} \subseteq \mathbb{R}^q$ we denote by $\mathcal{P}(\mathcal{X})$ the set of probability distributions on $\mathcal{X}$. For a random variable x_t taking values in $\mathcal{X}$, we denote by $\mathbb{P}_{x_t} \in \mathcal{P}(\mathcal{X})$ the probability measure associated to x_t , and by p_{x_t} the probability density function associated to $\mathbb{P}_{x_t}$. For all the probability distributions considered in this paper, we always assume that they have a density. Additionally, given a set of K probability distributions $\mathbb{P}_{\theta_1}, \dots, \mathbb{P}_{\theta_K} \in \mathcal{P}(\mathcal{X})$ and a set of weights $\omega_1, \dots, \omega_K \in [0, 1]$ such that $\sum_{k=1}^K \omega_k = 1$, we define a *mixture distribution* of size $K > 0$ as the distribution $\hat{\mathbb{P}} = \sum_{k=1}^K \omega_k \mathbb{P}_{\theta_k}$. If each $\mathbb{P}_{\theta_k}$ is a Gaussian distribution, then $\hat{\mathbb{P}}$ is called a *Gaussian mixture model* (GMM) or *Gaussian mixture distribution*. Finally, for a probability distribution $\mathbb{P} \in \mathcal{P}(\mathcal{X})$ and a measurable function $g : \mathcal{X} \rightarrow \mathcal{Y}$, where $\mathcal{Y} \subseteq \mathbb{R}^d$, we denote the push-forward measure of $\mathbb{P}$ by g as $g\#\mathbb{P}$ such that for all $A \subset \mathcal{Y}$, $(g\#\mathbb{P})(A) := \mathbb{P}(g^{-1}(A))$. We note that $g\#\mathbb{P}$ is still a probability distribution such that $g\#\mathbb{P} \in \mathcal{P}(\mathcal{Y})$.

B. System Description

We consider a discrete-time stochastic process described by the following stochastic difference equation

$$x_{t+1} = g(x_t, \varepsilon_t), \quad x_0 \sim \mathbb{P}_{x_0}, \quad \varepsilon_t \sim \mathbb{P}_\varepsilon, \quad (1)$$

where $g : \mathbb{R}^d \times \mathbb{R}^q \rightarrow \mathbb{R}^d$ is a possibly non-linear continuous function representing the one-step dynamics of System (1), x_0 is a random initial condition with distribution $\mathbb{P}_{x_0} \in \mathcal{P}(\mathbb{R}^d)$, and ε_t is a noise term distributed according to a absolutely continuous distribution independent and identically distributed at each time step¹ $\mathbb{P}_\varepsilon \in \mathcal{P}(\mathbb{R}^q)$. Intuitively, System (1) represents a general model of a d -dimensional discrete-time stochastic system with possible non-linearities both in the state and noise.

In order to study the time evolution of System (1), we first introduce $\mathbb{T}_x$, the one-step transition kernel of System (1), which for any time t represents the probability distribution of x_t given that $x_{t-1} = x$ and is formally defined as:

$$\mathbb{T}_x = g(x, \cdot)\#\mathbb{P}_\varepsilon.$$

For simplicity, with an abuse of notation, we denote by $p(y | x)$ the probability density associated to $\mathbb{T}_x$.

Example 1. In the case where $g(x, \varepsilon) = f(x) + \varepsilon$ and the noise distribution is $\mathbb{P}_\varepsilon = \mathcal{N}(0, \Sigma_\varepsilon)$, that is, a zero mean Gaussian with covariance matrix Σ_ε , we have that $\mathbb{T}_x = \mathcal{N}(f(x), \Sigma_\varepsilon)$, a Gaussian distribution with mean $f(x)$ and covariance matrix Σ_ε .

The probability distribution of System (1) at time t , $\mathbb{P}_{x_t}$, can be described by its density p_{x_t} , which can be com-

puted iteratively over time by using the following Chapman-Kolmogorov equation [3], [13]:

$$p_{x_t}(y) = \int_{\mathbb{R}^d} p(y | x) p_{x_{t-1}}(x) dx. \quad (2)$$

That is, p_{x_t} is obtained by marginalizing the one-step transition kernel w.r.t. $p_{x_{t-1}}$. Unfortunately, even though $p(y | x)$ can be expressed in closed form in many cases of interest, Eqn (2) is in general not tractable. For instance, this is the case for the system in Example 1 where $p(y | x)$ is Gaussian, but p_{x_t} is intractable and non-Gaussian if f is non-linear. Consequently, approximation schemes are required.

C. Problem Statement

Our goal in this paper is to design an approximating mixture distribution $\hat{\mathbb{P}}_{x_t}$, which is close to the actual probability distribution $\mathbb{P}_{x_t}$, but that is also tractable and can be computed analytically in closed form. To do that, we need to introduce a metric to quantify the distance between probability distributions. While many metrics have been introduced for this task [15], in this paper, we focus on the Total Variation distance.

Definition 1 (Total Variation). Let $(\mathcal{X}, \mathcal{B}(\mathcal{X}))$ be a measurable space, and $\mathbb{P}, \mathbb{Q} \in \mathcal{P}(\mathcal{X})$ be probability distributions. Then, the Total Variation distance between $\mathbb{P}$ and $\mathbb{Q}$ is defined as:

$$\text{TV}(\mathbb{P}, \mathbb{Q}) := \sup_{A \in \mathcal{B}(\mathcal{X})} |\mathbb{P}(A) - \mathbb{Q}(A)|.$$

Intuitively, the TV is the largest difference that the probability of the same event can have according to two probability distributions. Consequently, because of its intuitive and worst-case nature, the TV represents a particularly convenient choice to quantify the distance between the distribution of stochastic dynamical systems. We should also stress that, differently from the TV case, closeness in other measures commonly employed, such as Wasserstein distance, does not necessarily imply closeness in the probability of events [15].

We are finally ready to formulate our problem statement.

Problem 1. Consider a finite time horizon $\{1, \dots, T\} \subset \mathbb{N}$, and a threshold $\delta > 0$. Find a set of mixture distributions of finite size $\hat{\mathbb{P}}_{x_1}, \dots, \hat{\mathbb{P}}_{x_T}$ such that for all $t \in \{1, \dots, T\}$ it holds that:

$$\text{TV}(\mathbb{P}_{x_t}, \hat{\mathbb{P}}_{x_t}) \leq \delta. \quad (3)$$

Note that Problem 1 is stated for a general mixture distribution $\hat{\mathbb{P}}_{x_t}$. This is because depending on the one-step transition kernel of System (1), one may want to focus on different classes of mixture distributions. Nevertheless, because of their modeling flexibility and because of the existence of TV bounds between two Gaussian distributions [16], in this paper we will give particular attention to GMMs.

We should also note that solving Problem 1 is particularly challenging not only because Eqn (2) cannot be solved in closed form, but also because it requires propagating over time the uncertainty introduced by each approximation step. In fact, according to Eqn (2), to compute $\mathbb{P}_{x_t}$, one needs to

¹We remark that the assumption of time invariance of the noise distribution is only introduced to simplify notation and our methods naturally extend to time-varying noise distributions, as we will emphasize in Remark 2 in Section III.

marginalize the one step transition kernel wrt $\mathbb{P}_{x_{t-1}}$, which is unknown for any $t > 0$.

a) Approach: To solve Problem 1, we propose a stochastic approximation scheme consisting of an iterative process based on the approximation of Eqn (2) with a mixture distribution. Intuitively, we partition $\mathbb{R}^d$ in various subspaces $\mathcal{X}_k$ and select representative points $x^{(k)} \in \mathcal{X}_k$ such that $\int_{\mathbb{R}^d} p(y | x) p_{x_t}(x) dx$ is close to $\sum_{k=1}^K p(y | x^{(k)}) \int_{\mathcal{X}_k} p_{x_t}(x) dx$. In Section III, we detail our stochastic approximation scheme, and we state proof of its correctness and convergence to the distribution of System (1). In Section IV, we present the resulting algorithm used to build the mixture approximations. Finally, in Section V, we present numerical experiments.

III. STOCHASTIC APPROXIMATION SCHEME

To solve Problem 1, we introduce a stochastic approximation scheme based on iteratively approximating the distribution of System (1) by mixtures of the one-step transition kernel. In particular, for each time step t we consider a partition of $\mathbb{R}^d$ in K_t non-overlapping regions $\mathcal{X}_t := \{\mathcal{X}_t^{(1)}, \dots, \mathcal{X}_t^{(K_t)}\}$ and for each $\mathcal{X}_t^{(k)} \in \mathcal{X}_t$ we consider a representative point $x_t^{(k)} \in \mathcal{X}_t^{(k)}$. Then, we define the approximating mixture distribution $\hat{\mathbb{P}}_{x_t}$ iteratively as follows:

$$\begin{aligned} \hat{\mathbb{P}}_{x_0} &:= \mathbb{P}_{x_0} \\ \hat{\mathbb{P}}_{x_t} &:= \sum_{k=1}^{K_{t-1}} \omega_t^{(k)} \mathbb{T}_{x_{t-1}^{(k)}} \end{aligned} \quad (4)$$

with $\omega_t^{(k)} := \hat{\mathbb{P}}_{x_{t-1}}(\mathcal{X}_{t-1}^{(k)})$. Intuitively, at each time t , Eqn (4) builds a mixture approximation of $\mathbb{P}_{x_t}$ by averaging the one-step transition kernel of System (1) with weights that depend on the approximating mixture distribution at the previous time step.

Example 2. For the system dynamics of Example 1, Eqn (4) leads to a mixture distribution approximation at time step t given by the GMM $\hat{\mathbb{P}}_{x_t} = \sum_{k=1}^{K_{t-1}} \omega_t^{(k)} \mathcal{N}(f(x_{t-1}^{(k)}), \Sigma_\varepsilon)$, where $\omega_t^{(k)} = \hat{\mathbb{P}}_{x_{t-1}}(\mathcal{X}_{t-1}^{(k)})$ is the probability mass of the GMM $\hat{\mathbb{P}}_{x_{t-1}} = \sum_{k=1}^{K_{t-2}} \omega_{t-1}^{(k)} \mathcal{N}(f(x_{t-2}^{(k)}), \Sigma_\varepsilon)$ (built in the previous step) within the region $\mathcal{X}_{t-1}^{(k)}$.

Remark 1. Note that the mixture $\hat{\mathbb{P}}_{x_t}$ as defined in Eqn (4) is not necessarily a Gaussian mixture. If for the particular application it would be preferable to obtain a GMM, then this can always be obtained by approximating each $\mathbb{T}_{x_{t-1}^{(k)}}$ in Eqn (4) with a Gaussian (mixture) distribution. As the total variation is a distance that satisfies the triangle inequality, then the bounds on the approximation error developed in the rest of this section would hold also for this setting, with a correction term depending on the closeness of each $\mathbb{T}_{x_{t-1}^{(k)}}$ to the respective Gaussian or Gaussian mixture approximation.

It is obvious that the quality of the approximation obtained from Eqn (4) depends on the choice of the weights and representative points. How to perform this choice is the object of Section IV. In the remainder of this section, we focus on quantifying $\text{TV}(\hat{\mathbb{P}}_{x_t}, \mathbb{P}_{x_t})$.

A. Total variation bounds

We start our analysis with Theorem 1, where we bound $\text{TV}(\hat{\mathbb{P}}_{x_t}, \mathbb{P}_{x_t})$. Proofs for Theorems and Lemmas of this Section can be found in the Section VI.

Theorem 1. Let function $s(x, x^{(k)})$ be defined as

$$s(x, x^{(k)}) := \frac{1}{2} \int_{\mathbb{R}^d} |p(y | x) - p(y | x^{(k)})| dy \quad (5)$$

Then, for any $t > 0$ it holds that

$$\begin{aligned} \text{TV}(\mathbb{P}_{x_t}, \hat{\mathbb{P}}_{x_t}) &\leq \text{TV}(\mathbb{P}_{x_{t-1}}, \hat{\mathbb{P}}_{x_{t-1}}) + \\ &\sum_{k=1}^{K_{t-1}} \max_{x \in \mathcal{X}_{t-1}^{(k)}} s(x, x_{t-1}^{(k)}) \hat{\mathbb{P}}_{x_{t-1}}(\mathcal{X}_{t-1}^{(k)}), \end{aligned} \quad (6)$$

where $\text{TV}(\mathbb{P}_{x_0}, \hat{\mathbb{P}}_{x_0}) = 0$.

Theorem 1 guarantees that the approximation error at each time step propagates at most linearly over time. Furthermore, because of term $\hat{\mathbb{P}}_{x_{t-1}}(\mathcal{X}_{t-1}^{(k)})$, only regions with non-negligible probability mass w.r.t. $\hat{\mathbb{P}}_{x_{t-1}}$ influence the TV bound. The key component of Theorem 1 is the function $s(x, x_{t-1}^{(k)})$, which is equivalent to $\text{TV}(\mathbb{T}_x, \mathbb{T}_{x_{t-1}^{(k)}})$, i.e., the total variation between the one-step transition kernel at x and $x_{t-1}^{(k)}$. Consequently, the computation of $s(x, x_{t-1}^{(k)})$ depends on the specific noise distribution in System (1).

Luckily, closed-form expressions for $s(x, x_{t-1}^{(k)})$ or for upper bounds exist for many cases of interest, including when $\mathbb{T}_x$ follows a Gaussian, Gamma, or uniform distribution. Because of its widespread use, the Gaussian case is explicitly considered in detail in the next paragraph. For the Gamma distribution, one can use the fact that s can be bounded by the KL divergence using Pinsker's inequality and then rely on closed forms for the KL divergence between Gamma distributions [17]. For the uniform distribution, computing s can be performed by separating the integral in regions based on the overlapping of the supports (which can be easily done when the dimensions are independent, and hence the supports are hyper-rectangles).

a) Total Variation Bounds for the Gaussian Case: The following result is a corollary of Theorem 1 in the case of additive Gaussian noise.

Corollary 1. Assume that $g(x_t, \varepsilon_t) = f(x_t) + \varepsilon_t$ and $\varepsilon_t \sim \mathcal{N}(0, \Sigma_\varepsilon)$. Then, for any $t > 0$ it holds that

$$\begin{aligned} \text{TV}(\mathbb{P}_{x_t}, \hat{\mathbb{P}}_{x_t}) &\leq \text{TV}(\mathbb{P}_{x_{t-1}}, \hat{\mathbb{P}}_{x_{t-1}}) + \\ &\sum_{k=1}^{K_{t-1}} \text{erf} \left(\max_{x \in \mathcal{X}_{t-1}^{(k)}} h(x, x_{t-1}^{(k)}) \right) \hat{\mathbb{P}}_{x_{t-1}}(\mathcal{X}_{t-1}^{(k)}), \end{aligned} \quad (7)$$

where $\text{TV}(\mathbb{P}_{x_0}, \hat{\mathbb{P}}_{x_0}) = 0$ and

$$h(x, x_{t-1}^{(k)}) := \frac{1}{2\sqrt{2}} \sqrt{(f(x) - f(x_{t-1}^{(k)}))^T \Sigma_\varepsilon^{-1} (f(x) - f(x_{t-1}^{(k)}))}.$$

Eqn (7) follows from Theorem 1 using the fact that when the one-step transition kernel is Gaussian, then $s(x, x_{t-1}^{(k)})$

can be computed in a closed form in terms of the function $h(x, x_{t-1}^{(k)})$. Note also that for any $x \in \mathcal{X}_{t-1}^{(k)}$ the function $h(x, x_{t-1}^{(k)})$, and consequently the tightness of the bounds, depends on the following factors: (i) $f(x) - f(x_{t-1}^{(k)})$, which quantifies how much f varies locally, (ii) Σ_{ϵ}^{-1} , which suggests that lower noise variances may lead to worse bounds. This should not appear surprising when considering the definition of TV in Definition 1. In fact, from the definition of TV it follows that the TV between two delta Dirac distributions is one unless they have exactly the same mean. Consequently, the smaller the noise variance, the more local variation of $f(x)$ from the mean of each component of the GMM will affect the bound.

Remark 2. Corollary 1 can be straightforwardly generalized to the case where the covariance matrix varies with time, as long as it is independent of x . If heteroscedasticity is required in the noise model, one possible strategy is to adapt the proof using the fact that KL divergence can be used to upper bound the TV distance and use Prop. 2.1 of [16] to bound term s .

For both Theorem 1 and Corollary 1, at each time step the TV bound is obtained by a sum over all components in the mixture. An interesting question is whether the bound decreases with a mixture of a larger size. We answer positively to this question in the following subsection, where we prove the convergence of our method: the $\text{TV}(\mathbb{P}_{x_t}, \hat{\mathbb{P}}_{x_t})$ goes to zero for approximating mixture models of arbitrarily large size.

B. Convergence Analysis

In Theorem 2 below we show that the approximation error introduced by Eqn (4) can be made arbitrarily small by increasing the size of the mixture.

Theorem 2. For $\bar{x} \in \mathbb{R}^d$ assume that $\mathbb{T}_{\bar{x}}$ -almost surely (a.s.) for every $y \in \mathbb{R}^d$ it holds that

$$\|x - \bar{x}\|_{\infty} \rightarrow 0 \implies p(y | x) \rightarrow p(y | \bar{x}),$$

where $\|x - \bar{x}\|_{\infty}$ is the infinity norm of vector $x - \bar{x}$. Then, for any $\epsilon > 0$ and any time $t < \infty$, there exists a sequence of partitions $\{\mathcal{X}_0, \dots, \mathcal{X}_{t-1}\}$ such that

$$\text{TV}(\mathbb{P}_t, \hat{\mathbb{P}}_t) \leq \epsilon. \quad (8)$$

Note that the assumption of almost surely continuity of $p(y | x)$ in Theorem 2 is not particularly restrictive. For instance, it holds for distributions with continuous density, e.g., Gaussian, or even with bounded support, where the boundary of the support has measure zero.

Remark 3. We should remark that the convergence in Theorem 2 is proved assuming that each $\mathcal{X}_i$ is a uniform partition of a compact set containing $1 - \frac{\epsilon}{2t}$ probability mass of $\hat{\mathbb{P}}_{x_i}$. While sufficient for showing convergence, this choice of partition is obviously sub-optimal, as in general, we would like to place more distributions from the mixture in regions where there is more probability mass. This is explored in our algorithmic framework detailed in the next section.

IV. ALGORITHM

We summarize our overall procedure to solve Problem 1 in Algorithm 1. For every time t , in Lines 2-5 we compute the approximating mixture at the current time step following Eqn (4). Then, in Line 6 we identify a hyper-rectangle $\mathcal{S}_t$ containing at least $1 - \epsilon$ probability mass of $\hat{\mathbb{P}}_{x_t}$, where ϵ is a given parameter, which in our implementation, for time step t is taken as $\epsilon \ll \frac{\delta}{T}$, such that the impact of not partitioning region $\mathbb{R}^d \setminus \mathcal{S}_t$ on the bound is small compared to our desired precision. Such a region can either be identified analytically as $\hat{\mathbb{P}}_{x_t}$ is known, e.g., in the case $\hat{\mathbb{P}}_{x_t}$ is a GMM, or using samples. $\mathcal{S}_t$ is then partitioned in an initial set of K_t hyper-rectangles in Line 7 such that all regions contain (approximately) probability p_{thr} (a parameter usually set to 1% in our experiments) by iteratively cutting the axis in half and checking the inside probability until the threshold is met. This allows us to segment initial regions by probability mass (see left grid in Figure 1), and not by size, as is the case with uniform/equidistant grids. We encode the resulting set of representative points corresponding to partition $\mathcal{X}_t := \{\mathcal{X}_t^{(1)}, \dots, \mathcal{X}_t^{(K_t)}\}$ by $\mathbf{x}_t^{\text{repres}} := \{x_t^{(1)}, \dots, x_t^{(K_t)}\}$ and for each $\mathcal{X}_t^{(k)}$ we select $x_t^{(k)}$ as its center point. The resulting grid is then refined in Lines 10-12 until the TV bound is smaller than a given threshold, $t \frac{\delta}{T}$, with a refinement procedure, which is detailed in subsection IV-A. Note that, because at each time step the TV bound is propagated at the next time step (see Theorem 1), this guarantees that for each $t \in \{0, \dots, T\}$ we have that $\text{TV}(\mathbb{P}_{x_t}, \hat{\mathbb{P}}_{x_t})$ will be smaller than δ . We should also stress that, in practice, the computational complexity of Algorithm 1 is dominated by Step 8, which for each region requires to compute the probability mass of each component of the mixture, thus requiring a total number of integral evaluations quadratic in the size of the mixture (as in general the number of components and regions are close across consecutive iterations).

Algorithm 1: Mixture approx. with TV bounds for $T < \infty$ time steps given an error threshold δ

Input: $\epsilon, p_{\text{thr}}, \delta$

Output: Mixture approximations $\{\hat{\mathbb{P}}_{x_t}\}_{t=1}^T$, and TV bounds $\{\text{TV}_t\}_{t=1}^T$

```

1 for  $t \in \{0, \dots, T-1\}$  do
2   if  $t = 0$  then
3      $\hat{\mathbb{P}}_{x_0} \leftarrow \mathbb{P}_{x_0}$ ,  $\text{TV}_0 \leftarrow 0$ 
4   else
5      $\hat{\mathbb{P}}_{x_t} \leftarrow \sum_{k=1}^{K_{t-1}} \omega_t^{(k)} \mathbb{T}_{x_{t-1}^{(k)}}$ 
6      $\mathcal{S}_t = \text{IdentifyHighProbRegion}(\hat{\mathbb{P}}_{x_t}, \epsilon)$ 
7      $\mathcal{X}_t, \mathbf{x}_t^{\text{repres}} \leftarrow \text{BuildGrid}(\mathcal{S}_t, p_{\text{thr}})$ 
8      $\omega_t \leftarrow \{\hat{\mathbb{P}}_{x_t}(\mathcal{X}_t^{(1)}), \dots, \hat{\mathbb{P}}_{x_t}(\mathcal{X}_t^{(K_t)})\}$ 
9      $\text{TV}_{t+1} \leftarrow \text{TV}(\mathbb{P}_{x_{t+1}}, \hat{\mathbb{P}}_{x_{t+1}})$  using Thm 1
10    while  $\text{TV}_{t+1} \geq t \frac{\delta}{T}$  do
11       $\mathcal{X}_t, \mathbf{x}_t^{\text{repres}}, \omega_t \leftarrow \text{Refine}(\mathcal{X}_t, \mathbf{x}_t^{\text{repres}}, \omega_t)$ 
12       $\text{TV}_{t+1} \leftarrow \text{TV}(\mathbb{P}_{x_{t+1}}, \hat{\mathbb{P}}_{x_{t+1}})$  using Thm 1
13 return  $\{\hat{\mathbb{P}}_{x_1}, \dots, \hat{\mathbb{P}}_{x_T}\}, \{\text{TV}_1, \dots, \text{TV}_T\}$ 

```




Fig. 1: Example of a partition used to propagate the distribution of the System presented in Example 1 with $\varepsilon \sim \mathcal{N}(0, 0.1I)$. The initial grid (left) is refined (right) with a contribution threshold of $\gamma = 1 \times 10^{-3}$. We see that the algorithm decided to cut all regions except for the corners.

A. Refinement Procedure

The refinement procedure is described in Algorithm 2. The underlying idea of the refinement procedure is to rely on Theorem 1, where we can observe that for any t , $\text{TV}(\mathbb{P}_{x_t}, \hat{\mathbb{P}}_{x_t})$ depends on a sum over regions in $\mathcal{X}_{t-1}$, where each $\mathcal{X}_{t-1}^{(k)} \in \mathcal{X}_{t-1}$ contributes to the TV bound with the quantity $\max_{x \in \mathcal{X}_{t-1}^{(k)}} s(x, x_{t-1}^{(k)}) \hat{\mathbb{P}}_{x_{t-1}}(\mathcal{X}_{t-1}^{(k)}) := c^{(k)}$. The refinement procedure boils down to partitioning regions whose contribution to the TV bound is high in relative terms. More specifically, if the contribution $c^{(k)}$ of the hypercube $\mathcal{X}^{(k)}$ is higher than a threshold $\gamma > 0$, we cut the region in 2^d subregions by dividing each axis by 2 in its midpoint. An example of a refinement step is illustrated in Figure 1, where we can see that the refinement step decided to cut all inner regions except for the corners.

Algorithm 2: Refining the grid

```

1 function Refine( $\mathcal{X}, x^{\text{repres}}, \omega$ ):
2   for  $\mathcal{X}^{(k)}, x^{(k)}$  in  $\mathcal{X}, x^{\text{repres}}$  do
3      $c^{(k)} \leftarrow \max_{x \in \mathcal{X}^{(k)}} s(x, x^{(k)}) \hat{\mathbb{P}}_{x_{t-1}}(\mathcal{X}^{(k)})$ 
4     if  $c^{(k)} > \gamma$  then
5       Cut  $\mathcal{X}^{(k)}$  in  $2^d$  hypercubes
6    $\omega \leftarrow \{\hat{\mathbb{P}}_{x_t}(\mathcal{X}_{\text{new}}^{(1)}), \dots, \hat{\mathbb{P}}_{x_t}(\mathcal{X}_{\text{new}}^{(K^*)})\}$ 
7   return  $\mathcal{X}, x^{\text{repres}}, \omega$ 

```

V. NUMERICAL EXPERIMENTAL RESULTS

For our experiments², we consider the following benchmarks:

- A 2-D linear system with bimodal initial distribution with dynamics given by $g(x_t, \varepsilon_t) = Ax_t + \varepsilon_t$, where $A = \begin{bmatrix} 0.84 & 0.10 \\ 0.05 & 0.72 \end{bmatrix}$ and $\varepsilon_t \sim \mathcal{N}([0, 0]^T, 0.03I)$. The initial distribution is a mixture of two Gaussian distributions, that is, $x_0 \sim 0.5 \times \mathcal{N}([6, 10]^T, 0.005I) + 0.5 \times \mathcal{N}([8, 10]^T, 0.005I)$.

²The code can be found at <https://github.com/EduardoFMDCosta/DUQviaTotalVariation>

- A 2-D linear system with dynamics equal to the previous one, but with noise and initial condition following a uniform distribution. Specifically, we set $\varepsilon_t \sim \mathcal{U}([-0.3, 0.3]) \times \mathcal{U}([-0.3, 0.3])$, that is, a uniform distribution for each component of the noise between $[-0.3, 0.3]$. For the initial distribution we assume $x_0 \sim \mathcal{U}([-0.1, 0.1]) \times \mathcal{U}([-0.1, 0.1])$.

- An unstable polynomial system $g(x_t, \varepsilon_t) = f(x_t) + \varepsilon_t$ where $f(x) = [f_1(x), f_2(x)]^T$ with

$$f_1(x) = x_1 + 1.25hx_2$$

$$f_2(x) = 1.4x_2 + 0.3h(0.25x_1^2 - 0.4x_1x_2 + 0.25x_2^2)$$

for $h = 0.05$, $\varepsilon_t \sim \mathcal{N}([0, 0]^T, \sigma^2 I)$, where $\sigma \in \{1, 0.1, 0.001, 0.0001\}$ and $x_0 \sim \mathcal{N}([1, 1]^T, 0.002I)$.

- A 3D Dubins car model with constant velocity from [18], with $v = 5$, $h = 0.3$, and $u = 1/0.5$, with $x_0 \sim \mathcal{N}([0, 0, 0]^T, [0.005, 0.005, 0.001])$, and $\varepsilon_t \sim \mathcal{N}([0, 0, 0]^T, [0.06, 0.06, 0.01])$.

In Algorithm 1, we set $p_{\text{thr}} = 0.01$ (0.001 for Dubins) to build the initial grid, and for the refinement in Algorithm 2 we choose γ ranging from 10^{-6} to 10^{-7} depending on the case. All experiments were run on an Intel Core i7-1365U CPU with 16GB of RAM. Step 8 of Algorithm 1 was parallelized on 10 cores (given that this step is the bottleneck in terms of computational complexity).

A. Tightness of the Bounds

We start our analysis from Table I, where we study the efficacy of Algorithm 1 to build a GMM approximation for the polynomial system benchmark. The table shows that unless the variance of the noise, and consequently of the one-step transition kernel, becomes very small (below 10^{-2}) even a GMM of relatively few mixtures (order of 100) can obtain relatively tight TV bounds and be employed to solve Problem 1. However, when the variance of the noise becomes small and the noise distribution becomes closer to a delta Dirac, as already discussed right below Corollary 1, obtaining tight TV bounds become harder. Nevertheless, as illustrated in Figure 2, we should stress that, even if obtaining tight TV bounds in this setting becomes more challenging, the distribution obtained with our framework is still empirically very close to the true one even with mixtures of relatively small size.

In Table II we report an empirical evaluation of our approach for all four benchmarks. For all cases, we consider 5 refinement steps and compare the mixture model obtained from Algorithm 1 with the standard approach in [4] (called *equidistant grid*, which we will sometimes refer to the uniform grid)³, where each distribution in the mixture is uniformly spaced in a region containing most of the probability mass of the system and obtained via samples. We consider a time horizon of $T = 10$ for the bimodal case, $T = 7$ for the polynomial, and $T = 5$ for Dubins car and uniform settings. In all cases we can see how Algorithm 2

³Other works have proposed adaptive (i.e. non-uniform) grids for similar problems, e.g. [19], but ours is designed with guarantees of improvement in the TV distance.




Fig. 2: Upper row considers the Dubins car setting with low noise variance (i.e. 10^{-3} for the positions x_1 and x_2 , and 10^{-4} for the steering angle x_3), while the lower row considers a higher variance noise (10^{-2} , and 2×10^{-3} , respectively). The mixtures are generated with only one refinement (see Algorithm 2) with $\gamma = 10^{-6}$, resulting in mixtures of size 2500-3000. *True distribution* means Monte Carlo sampling from the system, while the plots on the right are generated by sampling our mixture distributions.

substantially outperforms the approach mentioned in [4] by obtaining much better bounds for an equivalent number of regions. In the Dubins car, for instance, we were not able to obtain informative bounds (e.g. $\ll 1$) for a reasonable number of regions, while Algorithm 2 allowed for bounds between 3% (first-time step) to 20% (last).

Noise variance σ^2	1		10^{-1}		10^{-2}		10^{-3}	
# refinements	TV	GMM size	TV	GMM size	TV	GMM size	TV	GMM size
Initial grid	0.020	121	0.061	124	0.205	115	0.471	127
1	0.011	485	0.031	497	0.107	461	0.279	509
2	0.006	1,937	0.016	1,982	0.054	1,841	0.151	2,033
3	0.003	7,667	0.008	7,868	0.027	7,331	0.078	8,108
4	0.002	26,288	0.005	30,824	0.014	28,997	0.039	32,213
5	0.002	31,070	0.003	59,660	0.007	112,994	0.020	127,220

TABLE I: TV bounds for a 1-time step propagation of the polynomial system for the same initial distribution x_0 for various values of variance σ^2 . In each row, we show an iteration of the refinement algorithm and report both the resulting TV bound and the size of the approximating GMM.

B. Application to chance constrained certification

To highlight the usefulness of our approach, we show how it can be employed to quantify the probability that, at each time step, System (1) satisfies probabilistic constraints,

Grid method	Bimodal		Polynomial		Uniform		Dubins	
	Alg 2	Equid	Alg 2	Equid	Alg 2	Equid	Alg 2	Equid
TV ₁	0.004	0.007	0.004	0.006	0.004	0.003	0.028	0.99
TV _T	0.061	0.092	0.099	0.179	0.041	0.048	0.198	1.00
Avg TV	0.033	0.049	0.039	0.070	0.022	0.022	0.101	1.00

TABLE II: TV bounds for the first and last propagation steps (as well as for the average) for each system. The bounds are computed from grids generated by Algorithm 2 (Alg 2) with 5 refinements and by the uniform/equidistant grids (Equid) of equivalent size.

a setting commonly occurring in planning problems for stochastic systems [20]–[22]. In particular, for the bimodal system, we consider the setting in Figure 3, where the goal is to quantify the probability that at each time step the system will enter the unsafe region in red (which we denote by p_{hit}). In Table III, we report both the empirical value obtained by simulating the specific system 10^4 times and the upper bound of the hitting probability at each time step according to our framework (computed as the probability that the GMM approximation hits the obstacle + the TV bound at the particular time step computed according to Theorem 1). We observe that the upper bound for p_{hit} is fairly close to the actual hitting probability of the system, due to two factors: i) the hitting probability estimated by our mixture is very close to the true value (analogous to what is shown e.g. in Figure 2), and ii) the TV bounds are reasonably close to zero (see also Table II).

Bimodal		
Time step	Actual hitting probability p_{hit}	Upper bound p_{hit}
1	0.0	0.4
2	0.0	0.9
3	0.0	1.6
4	0.0	2.4
5	18.4	21.5
6	38.0	42.1
7	20.9	25.6
8	3.5	8.5
9	0.0	5.8
10	0.0	6.2

TABLE III: Actual p_{hit} is the unsafe set hitting probability computed via Monte Carlo simulation of the actual system (10^4 samples), while the upper bound is given by the probability mass of the GMM inside the unsafe set + TV bound.


VI. PROOFS

A. Proof of Theorem 1

For the proof we introduce a new auxiliary distribution $\tilde{\mathbb{P}}_{x_t}$, whose density is defined as

$$\tilde{p}_{x_t}(y) := \int_{\mathbb{R}^d} p(y | x) \hat{p}_{x_{t-1}}(x) dx.$$

That is, $\tilde{\mathbb{P}}_{x_t}$ is the probability distribution given by the propagation of mixture $\hat{\mathbb{P}}_{x_{t-1}}$ through the dynamics of the



(a) True distribution (b) Mixture approximation

Fig. 3: Bimodal setting for which we present the hitting probability study. The left image shows a Monte Carlo simulation of the system, while the right displays samples from our mixture approximation. p_{hit} is computed w.r.t. the red set.

system. Then, the proof uses the fact that by the triangle inequality

$$\text{TV}(\mathbb{P}_{x_t}, \hat{\mathbb{P}}_{x_t}) \leq \text{TV}(\mathbb{P}_{x_t}, \tilde{\mathbb{P}}_{x_t}) + \text{TV}(\tilde{\mathbb{P}}_{x_t}, \hat{\mathbb{P}}_{x_t}).$$

Consequently, to bound $\text{TV}(\mathbb{P}_{x_t}, \hat{\mathbb{P}}_{x_t})$ it is enough to bound $\text{TV}(\mathbb{P}_{x_t}, \tilde{\mathbb{P}}_{x_t})$ and $\text{TV}(\tilde{\mathbb{P}}_{x_t}, \hat{\mathbb{P}}_{x_t})$. We start from $\text{TV}(\tilde{\mathbb{P}}_{x_t}, \hat{\mathbb{P}}_{x_t})$. To do that we first note that for every $y \in \mathbb{R}^d$, the following holds:

$$\begin{aligned} |\tilde{p}_{x_t}(y) - \hat{p}_{x_t}(y)| &= \left| \int_{\mathbb{R}^d} p(y | x) \hat{p}_{x_{t-1}}(x) dx - \sum_{k=1}^{K_{t-1}} \omega_t^{(k)} p(y | x_{t-1}^{(k)}) \right| \\ &= \left| \sum_{k=1}^{K_{t-1}} \int_{\mathcal{X}_{t-1}^{(k)}} (p(y | x) - p(y | x_{t-1}^{(k)})) \hat{p}_{x_{t-1}}(x) dx \right| \\ &\leq \sum_{k=1}^{K_{t-1}} \int_{\mathcal{X}_{t-1}^{(k)}} |p(y | x) - p(y | x_{t-1}^{(k)})| \hat{p}_{x_{t-1}}(x) dx, \end{aligned}$$

where the second line uses the definition of $\omega_t^{(k)}$ in Eqn (4), and the third the triangle inequality. From this, it follows that

$$\begin{aligned} \text{TV}(\tilde{\mathbb{P}}_{x_t}, \hat{\mathbb{P}}_{x_t}) &= \frac{1}{2} \int_{\mathbb{R}^d} |\tilde{p}_{x_t}(y) - \hat{p}_{x_t}(y)| dy \\ &\leq \frac{1}{2} \int_{\mathbb{R}^d} \left[\sum_{k=1}^{K_{t-1}} \int_{\mathcal{X}_{t-1}^{(k)}} |p(y | x) - p(y | x_{t-1}^{(k)})| \hat{p}_{x_{t-1}}(x) dx \right] dy \\ &= \sum_{k=1}^{K_{t-1}} \int_{\mathcal{X}_{t-1}^{(k)}} \underbrace{\left[\frac{1}{2} \int_{\mathbb{R}^d} |p(y | x) - p(y | x_{t-1}^{(k)})| dy \right]}_{s(x, x_{t-1}^{(k)})} \hat{p}_{x_{t-1}}(x) dx \\ &\leq \sum_{k=1}^{K_{t-1}} \max_{x \in \mathcal{X}_{t-1}^{(k)}} s(x, x_{t-1}^{(k)}) \int_{\mathcal{X}_{t-1}^{(k)}} \hat{p}_{x_{t-1}}(x) dx \\ &= \sum_{k=1}^{K_{t-1}} \max_{x \in \mathcal{X}_{t-1}^{(k)}} s(x, x_{t-1}^{(k)}) \hat{\mathbb{P}}_{x_{t-1}}(\mathcal{X}_{t-1}^{(k)}) \end{aligned}$$

where the change order of integration in the third line is possible due to the application of Fubini's theorem after the

observation that the inner function is integrable on $\mathbb{R}^d \times \mathbb{R}^d$.

Now, to conclude, for the TV distance between $\mathbb{P}_{x_t}$ and $\tilde{\mathbb{P}}_{x_t}$ we have that

$$\begin{aligned} \text{TV}(\mathbb{P}_{x_t}, \tilde{\mathbb{P}}_{x_t}) &= \frac{1}{2} \int_{\mathbb{R}^d} |p_{x_t}(y) - \tilde{p}_{x_t}(y)| dy \\ &= \frac{1}{2} \int_{\mathbb{R}^d} \left| \int_{\mathbb{R}^d} p(y | x) p_{x_{t-1}}(x) dx - \int_{\mathbb{R}^d} p(y | x) \hat{p}_{x_{t-1}}(x) dx \right| dy \\ &= \frac{1}{2} \int_{\mathbb{R}^d} \left| \int_{\mathbb{R}^d} p(y | x) (p_{x_{t-1}}(x) - \hat{p}_{x_{t-1}}(x)) dx \right| dy \\ &\leq \frac{1}{2} \int_{\mathbb{R}^d} \int_{\mathbb{R}^d} p(y | x) |p_{x_{t-1}}(x) - \hat{p}_{x_{t-1}}(x)| dx dy \\ &= \frac{1}{2} \int_{\mathbb{R}^d} \underbrace{\left[\int_{\mathbb{R}^d} p(y | x) dy \right]}_{=1} |p_{x_{t-1}}(x) - \hat{p}_{x_{t-1}}(x)| dx \\ &= \frac{1}{2} \int_{\mathbb{R}^d} |p_{x_{t-1}}(x) - \hat{p}_{x_{t-1}}(x)| dx \\ &= \text{TV}(\mathbb{P}_{x_{t-1}}, \hat{\mathbb{P}}_{x_{t-1}}) \end{aligned}$$

where the fourth line uses the triangle inequality, and the fifth consists of a change in the integration order applying Fubini's theorem.

B. Proof of Corollary 1

From the definition of total variation, we have that

$$s(x, x_{t-1}^{(k)}) = \text{TV}(\mathcal{N}(f(x), \Sigma_\varepsilon), \mathcal{N}(f(x_{t-1}^{(k)}), \Sigma_\varepsilon)).$$

It then follows from [16] (or Theorem 1 in [23]), that

$$\begin{aligned} s(x, x_{t-1}^{(k)}) &= \mathbb{P} \left(\mathcal{N}(0, 1) \in \left[-\sqrt{2}h(x, x_{t-1}^{(k)}), \sqrt{2}h(x, x_{t-1}^{(k)}) \right] \right) \\ &= \Phi \left(\sqrt{2}h(x, x_{t-1}^{(k)}) \right) - \Phi \left(-\sqrt{2}h(x, x_{t-1}^{(k)}) \right) \\ &= 2\Phi \left(\sqrt{2}h(x, x_{t-1}^{(k)}) \right) - 1 = \text{erf} \left(h(x, x_{t-1}^{(k)}) \right) \end{aligned}$$

C. Proof of Theorem 2

Proof. From Theorem 1 we have that

$$\text{TV}(\mathbb{P}_s, \hat{\mathbb{P}}_s) \leq \sum_{t=1}^s \sum_{k=1}^{K_{t-1}} \max_{x \in \mathcal{X}_{t-1}^{(k)}} s(x, x_{t-1}^{(k)}) \hat{\mathbb{P}}_{x_{t-1}}(\mathcal{X}_{t-1}^{(k)}).$$

Consequently, it is enough to prove that for any $\bar{\varepsilon} > 0$, for any t , there exists a partition of $\mathbb{R}^d = \{\mathcal{X}^{(1)}, \dots, \mathcal{X}^{(K+1)}\}$ in $K+1$ sets such that

$$\sum_{k=1}^{K+1} \max_{x \in \mathcal{X}^{(k)}} s(x, x_{t-1}^{(k)}) \hat{\mathbb{P}}_{x_{t-1}}(\mathcal{X}^{(k)}) \leq \bar{\varepsilon}. \quad (9)$$

In fact, we can then conclude by taking $\bar{\varepsilon} < \frac{\varepsilon}{t}$. In order to do that, we first consider a hypercube $\mathcal{S}$ of diameter r w.r.t. the infinity norm containing at least $1 - \frac{\bar{\varepsilon}}{2}$ probability mass of $\hat{\mathbb{P}}_{x_{t-1}}$, which is guaranteed to exist as $\hat{\mathbb{P}}_{x_{t-1}}$ is almost surely bounded by assumption. We then select $\mathcal{X}^{(K+1)} = \mathbb{R}^d \setminus \mathcal{S}$

and let $\{\mathcal{X}^{(1)}, \dots, \mathcal{X}^{(K)}\}$ be a uniform partition of $\mathcal{S}$ in K hypercubic regions each of diameter $\frac{r}{\sqrt{K}}$. What is left to show is that there exists $\bar{K}$ such that for any $K \geq \bar{K}$

$$\sum_{k=1}^K \max_{x \in \mathcal{X}^{(k)}} s(x, x^{(k)}) \hat{\mathbb{P}}_{x_{t-1}}(\mathcal{X}^{(k)}) \leq \frac{\bar{\epsilon}}{2}. \quad (10)$$

Now, note that:

$$\begin{aligned} & \sum_{k=1}^K \max_{x \in \mathcal{X}^{(k)}} s(x, x^{(k)}) \hat{\mathbb{P}}_{x_{t-1}}(\mathcal{X}^{(k)}) \\ & \leq \max_{x \in \mathcal{X}^{(k)}} \max_{\bar{x}: \|x - \bar{x}\|_\infty \leq \frac{r}{\sqrt{K}}} s(x, \bar{x}) \sum_{k=1}^K \hat{\mathbb{P}}_{x_{t-1}}(\mathcal{X}^{(k)}) \\ & \leq \max_{x \in \mathcal{X}^{(k)}} \max_{\bar{x}: \|x - \bar{x}\|_\infty \leq \frac{r}{\sqrt{K}}} s(x, \bar{x}) \cdot 1. \end{aligned}$$

To conclude it is enough to note that, for any $x \in \mathcal{X}^{(k)}$, if $p(y | \bar{x}) \rightarrow p(y | x)$ when $K \rightarrow \infty$ for every y a.s., then it holds that

$$\lim_{K \rightarrow \infty} \left(\max_{\bar{x}: \|x - \bar{x}\|_\infty \leq \frac{r}{\sqrt{K}}} s(x, \bar{x}) \right) = 0,$$

since $\lim_{K \rightarrow \infty} s(x, \bar{x}) = 0$ due to Scheffe's Lemma [24] (Lemma 13.3). $\square$

VII. CONCLUSION

We introduced a framework to approximate the distribution of a non-linear stochastic dynamical system over time with a mixture of distributions. By deriving bounds for the Total Variation distance, we showed that for each time step, the resulting approximating mixture distribution is guaranteed to be δ -close to the distribution of the system. On a set of experiments, we showed the efficacy of our framework. In our view, there are several opportunities for improvement of the current framework, such as i) mixture compression (reducing the number of components in the approximating mixture with formal guarantees) to decrease the computational complexity, ii) bound tightening (in Theorem 1, we are working on ways to improve the bound without resorting to the maximization), and iii) extension of the formal guarantees to other metrics on probability distributions, such as the Wasserstein distance.

VIII. ACKNOWLEDGEMENTS

L.L. and E.F. are partially supported by the NWO (grant OCENW.M.22.056).

REFERENCES

- [1] M. Alami, "Learning against uncertainty in control engineering," *Annual Reviews in Control*, vol. 53, pp. 19–29, 2022.
- [2] È. Pairet, J. D. Hernández, M. Carreras, Y. Petillot, and M. Lahijanian, "Online mapping and motion planning under uncertainty for safe navigation in unknown environments," *IEEE Transactions on Automation Science and Engineering*, vol. 19, no. 4, pp. 3356–3378, 2021.
- [3] A. Girard, C. Rasmussen, J. Q. Candela, and R. Murray-Smith, "Gaussian process priors with uncertain inputs application to multiple-step ahead time series forecasting," *Advances in neural information processing systems*, vol. 15, 2002.
- [4] D. Landgraf, A. Völz, F. Berkel, K. Schmidt, T. Specker, and K. Graichen, "Probabilistic prediction methods for nonlinear systems with application to stochastic model predictive control," *Annual Reviews in Control*, vol. 56, p. 100905, 2023.
- [5] G. Einicke, *Smoothing, filtering and prediction: Estimating the past, present and future*. BoD-Books on Demand, 2012.
- [6] T. S. Schei, "A finite-difference method for linearization in nonlinear estimation algorithms," *Automatica*, vol. 33, no. 11, pp. 2053–2058, 1997.
- [7] J. Dunik, S. K. Biswas, A. G. Dempster, T. Pany, and P. Closas, "State estimation methods in navigation: Overview and application," *IEEE Aerospace and Electronic Systems Magazine*, vol. 35, no. 12, pp. 16–31, 2020.
- [8] K. Ito and K. Xiong, "Gaussian filters for nonlinear filtering problems," *IEEE transactions on automatic control*, vol. 45, no. 5, pp. 910–927, 2000.
- [9] J. S. Lui and R. Chen, "Sequential monte carlo methods for dynamic systems," *Journal of the American Statistical Association*, vol. 93, no. 443, p. 1032, 1998.
- [10] M. S. Arulampalam, S. Maskell, N. Gordon, and T. Clapp, "A tutorial on particle filters for online nonlinear/non-gaussian bayesian tracking," *IEEE Transactions on signal processing*, vol. 50, no. 2, pp. 174–188, 2002.
- [11] K. Polymenakos, L. Laurenti, A. Patane, J.-P. Calliess, L. Cardelli, M. Kwiatkowska, A. Abate, and S. Roberts, "Safety guarantees for iterative predictions with gaussian processes," in *2020 59th IEEE Conference on Decision and Control (CDC)*. IEEE, 2020, pp. 3187–3193.
- [12] A. Jasour, A. Wang, and B. C. Williams, "Moment-based exact uncertainty propagation through nonlinear stochastic autonomous systems," *arXiv preprint arXiv:2101.12490*, 2021.
- [13] F. Weissel, M. F. Huber, and U. D. Hanebeck, "Stochastic nonlinear model predictive control based on gaussian mixture approximations," in *Informatics in Control, Automation and Robotics: Selected Papers from the International Conference on Informatics in Control, Automation and Robotics 2007*. Springer, 2009, pp. 239–252.
- [14] G. Terejanu, P. Singla, T. Singh, and P. D. Scott, "Uncertainty propagation for nonlinear dynamic systems using gaussian mixture models," *Journal of guidance, control, and dynamics*, vol. 31, no. 6, pp. 1623–1633, 2008.
- [15] A. L. Gibbs and F. E. Su, "On choosing and bounding probability metrics," *International statistical review*, vol. 70, no. 3, pp. 419–435, 2002.
- [16] L. Devroye, A. Mehrabian, and T. Reddad, "The total variation distance between high-dimensional gaussians with the same mean," *arXiv preprint arXiv:1810.08693*, 2018.
- [17] A. B. Tsybakov, "Introduction to nonparametric estimation," *Springer Series in Statistics*, 2009.
- [18] F. B. Mathiesen, S. C. Calvert, and L. Laurenti, "Safety certification for stochastic systems via neural barrier functions," *IEEE Control Systems Letters*, vol. 7, pp. 973–978, 2022.
- [19] M. Šimandl, J. Královec, and T. Söderström, "Advanced point-mass method for nonlinear state estimation," *Automatica*, vol. 42, no. 7, pp. 1133–1145, 2006.
- [20] S. Adams, M. Lahijanian, and L. Laurenti, "Formal control synthesis for stochastic neural network dynamic models," *IEEE Control Systems Letters*, vol. 6, pp. 2858–2863, 2022.
- [21] T. Brüdigam, M. Olbrich, D. Wollherr, and M. Leibold, "Stochastic model predictive control with a safety guarantee for automated driving," *IEEE Transactions on Intelligent Vehicles*, vol. 8, no. 1, pp. 22–36, 2021.
- [22] Q. H. Ho, Z. N. Sunberg, and M. Lahijanian, "Gaussian belief trees for chance constrained asymptotically optimal motion planning," in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 11 029–11 035.
- [23] S. Barsov and V. V. Ulyanov, "Estimates of the proximity of gaussian measures," in *Sov. Math., Dokl.*, vol. 34, 1987, pp. 462–466.
- [24] B. K. Driver, "Math 280 (probability theory) lecture notes," *Lecture notes, University of California, San Diego, Department of Mathematics*, vol. 112, pp. 92 093–0112, 2007.