

Universal Learning of Stochastic Dynamics for Exact Belief Propagation Using Bernstein Normalizing Flows

Peter Amorese and Morteza Lahijanian

Department of Aerospace Engineering Sciences, University of Colorado Boulder, USA
 Peter.Amorese@colorado.edu, Morteza.Lahijanian@colorado.edu

Abstract

Predicting the distribution of future states in a stochastic system, known as *belief propagation*, is fundamental to reasoning under uncertainty. However, nonlinear dynamics often make analytical belief propagation intractable, requiring approximate methods. When the system model is unknown and must be learned from data, a key question arises: *can we learn a model that (i) universally approximates general nonlinear stochastic dynamics, and (ii) supports analytical belief propagation?* This paper establishes the theoretical foundations for a class of models that satisfy both properties. The proposed approach combines the expressiveness of *normalizing flows* for density estimation with the analytical tractability of *Bernstein polynomials*. Empirical results show the efficacy of our learned model over state-of-the-art data-driven methods for belief propagation, especially for highly non-linear systems with non-additive, non-Gaussian noise.

Code — <https://github.com/peteramorese/BernsteinFlow/tree/anonymous>

[//github.com/peteramorese/BernsteinFlow/tree/anonymous](https://github.com/peteramorese/BernsteinFlow/tree/anonymous)-zip

Extended version — <https://arxiv.org/abs/2509.15533>

1 Introduction

At the heart of intelligent reasoning under uncertainty is the ability to predict future random outcomes. Prediction accuracy has profound implications for the safety and effectiveness of real-world systems. Achieving accurate and informative predictions requires both a representative *model* of the underlying stochastic process and a principled method for *reasoning* with that model. When the model is known but nonlinear (or non-Gaussian), reasoning, specifically predicting the future state distribution, known as *belief propagation*, becomes analytically intractable and typically necessitates approximation. When the underlying dynamics are unknown, machine learning provides a data-driven means to model them. However, the choice of learning model should not only aim for expressive power but also support the type of reasoning required. In the case of belief propagation, this leads to a central question: *Can we learn a model that (i) universally approximates general nonlinear stochastic dynamics, and (ii) supports analytical belief propagation?*

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

This work focuses on the above question and lays the theoretical groundwork for a class of general nonlinear models capable of learning complex stochastic systems while supporting *analytical* belief propagation. The key insight is that, under a Bayesian (Markov chain) framework, belief propagation reduces to two fundamental operations, *multiplication* and *integration*, over two functions: the *prior* and the *conditional* probability density functions (PDFs). If these functions are represented as polynomials, these operations remain exact, since polynomials are closed under both multiplication and integration. This eliminates the need for approximation during propagation. Building on this observation, we develop a learning framework in which polynomials are trained to soundly represent *true* PDFs, i.e., non-negative functions that integrate to one over their domain. Specifically, in our approach, dubbed *Bernstein Normalizing Flows* (BNFs), the underlying PDFs are modeled using Bernstein polynomials, which offer both favorable analytical properties and universal approximation capabilities. However, polynomials alone are ill-suited for modeling PDFs over unbounded state spaces. To overcome this limitation, we leverage insights from normalizing flows (Papamakarios et al. 2021) to enable proper universal density estimation with polynomials. Hence, by embedding Bernstein polynomials in the normalizing flow architecture, we obtain models that can approximate arbitrary nonlinear stochastic systems while supporting efficient and exact belief propagation.

To the best of our knowledge, the proposed BNF framework is the first class of general nonlinear Markov chain models that simultaneously support universal approximation of stochastic dynamics and exact analytical belief propagation. Beyond the introduction of BNFs, this work makes three key contributions: (i) an explicit-constraint-free training procedure for learning valid PDFs, (ii) a method for enhancing model expressiveness without increasing the number of parameters, and (iii) empirical validation demonstrating the effectiveness of BNFs for belief propagation in comparison with state-of-the-art data-driven approaches.

Related Work Belief propagation in nonlinear stochastic systems is a fundamental challenge in probabilistic reasoning and control. A variety of approximate methods have been developed to handle this task, particularly in the contexts of filtering, density estimation, and probabilistic modeling.

These approaches vary in their trade-offs between accuracy, tractability, and the ability to provide formal guarantees.

Approximate belief propagation has been extensively studied in the context of nonlinear filtering (Schei 1997; Julier and Uhlmann 2004). Gaussian Mixture Models (GMMs) are commonly used to represent multimodal beliefs (Alspach and Sorenson 2003) while preserving some analytical tractability, such as efficient sampling and region-based integration, features valuable for planning. However, linearization-based methods like the Extended Kalman Filter (EKF) (Jazwinski 2013) introduce significant error under nonlinear dynamics. Component splitting techniques (Kulik and LeGrand 2024) mitigate this by recursively subdividing GMM components, but the resulting exponential growth in components renders long-horizon prediction intractable. These filtering approaches also typically assume additive Gaussian noise, making Gaussian Process (GP) regression a common choice for learning system dynamics (Deisenroth et al. 2011).

Monte Carlo methods (Djuric et al. 2003) offer a flexible alternative by representing beliefs with particle sets. Propagation is straightforward if state transitions can be sampled, but obtaining a usable PDF requires post-hoc density estimation. These methods are also prone to inaccuracies in low-sample regions and lack formal error bounds. In contrast to the approximate methods, our approach performs *analytical* belief propagation. Similar to GP regression with GMM propagation, our model can evaluate probabilistic future events, and easily draw samples from the current belief.

Several works have aimed to provide *formal* error bounds on belief propagation. For instance, Polymenakos et al. (2020) derive bounds for GP models, while grid-based techniques generalize formal GMM propagation to broader nonlinear systems (Figueiredo et al. 2024), albeit under restrictive assumptions (e.g., diagonal-covariance GMM approximations). Moment-based propagation methods also exist (Jasour, Wang, and Williams 2021), but moments alone are often insufficient to fully recover the underlying distribution (Akhiezer 2020). Unlike these methods, our approach does not assume a given model; instead, we focus on learning a model that supports exact belief propagation, eliminating the need for bounding the approximation error entirely.

Polynomials have long been used for density estimation (Yu and Loskot 2023) due to their normalization convenience. Bernstein polynomials, in particular, have been applied to smooth approximations of empirical distributions and in non-parametric settings (Babu, Cauty, and Chaubey 2002; Belalia, Bouezmarni, and Leblanc 2017), though often via kernel-based methods. Sampling from general multivariate polynomial densities remains a challenge. Our work addresses this by embedding Bernstein polynomials within a normalizing flow architecture. Univariate Bernstein polynomials have been used in normalizing flows in conjunction with neural networks for robustness and interpretability (Ramasinghe et al. 2022; Arpogaus et al. 2025). However, due to the neural network, the modeled density is not a polynomial. In contrast, our proposed model is entirely polynomial, enabling efficient sampling, maintaining model sparsity, and supporting exact belief propagation.

Notation Throughout the paper, bold symbols denote vectors, e.g., $\mathbf{x} = [x_1, \dots, x_n] \in \mathbb{R}^n$, and polynomial functions are denoted as $\pi(\mathbf{x})$. Given a random variable \mathbf{x} , its probability density function (PDF), which we also refer to as *density*, is denoted $p(\mathbf{x})$, and the probability of \mathbf{x} being in region $R \subseteq \mathbb{R}^n$ is denoted $P(\mathbf{x} \in R)$. Multivariate polynomials (e.g., in the monomial-basis) $\pi(\mathbf{x}) = \sum_{j_1=0}^{d_1} \dots \sum_{j_n=0}^{d_n} c_{j_1, \dots, j_n} x_1^{j_1} \dots x_n^{j_n}$ are denoted with multi-indices $\pi(\mathbf{x}) = \sum_{\mathbf{j}=0}^{\mathbf{d}} c_{\mathbf{j}} \mathbf{x}^{\mathbf{j}}$, and the set of corresponding coefficients is denoted in bold as \mathbf{c} . The set of all degree- \mathbf{d} polynomials is denoted $\Pi^{\mathbf{d}}[\mathbf{x}]$. For a vector-valued function $f(\mathbf{x})$, the i -th component is denoted $f_i(\mathbf{x})$. If f depends only on the first $l \leq n$ components of \mathbf{x} , we denote this by writing $f(\mathbf{x}_{\leq l})$. If f depends only on the i -th component of \mathbf{x} , we use $f(\mathbf{x}_i)$. Finally, the *image* of region R via mapping f is denoted $f(R) := \{f(\mathbf{x}) \mid \mathbf{x} \in R\}$.

2 Problem Formulation

Consider a discrete-time, non-linear stochastic dynamical system of the form

$$\mathbf{x}_{k+1} = f(\mathbf{x}_k, \mathbf{v}_k) \quad (1)$$

where $\mathbf{x}_k \in X \subseteq \mathbb{R}^n$ is the state, and $\mathbf{v}_k \sim p(\mathbf{v})$ is a stochastic process noise variable. We assume function f is continuous and differentiable, and the initial state $\mathbf{x}_0 \sim p(\mathbf{x}_0)$ is random. This work considers a state space X that is either unbounded, i.e., $X = \mathbb{R}^n$, or bounded hyperrectangular, i.e., $X = \times_{i=1}^n I_i$ for some closed intervals $I_i \subset \mathbb{R}$.

The dynamics of System (1) can be equivalently represented by the state-transition distribution

$$p(\mathbf{x}_k \mid \mathbf{x}_{k-1}) = f(\mathbf{x}_{k-1}, \cdot) \# p(\mathbf{v}) \quad (2)$$

where $\#$ denotes the push-forward operator. This representation, which implicitly captures the non-linear dependence of f on \mathbf{x}_k and \mathbf{v}_k , may be very complex for highly non-linear systems. For ease of presentation, we denote this distribution as $p(\mathbf{x}' \mid \mathbf{x})$. Hence, the stochastic evolution of the system (1) from the initial distribution $p(\mathbf{x}_0)$ is captured by a continuous-state Markov chain $\mathcal{M}_f = (X, p(\mathbf{x}_0), p(\mathbf{x}' \mid \mathbf{x}))$.

We consider a scenario where both the dynamics f and the initial state distribution (belief) $p(\mathbf{x}_0)$ are *unknown* and must be learned from data. Specifically, we assume data is provided in the form of state input-output pairs $\mathcal{D}_{\mathbf{x}'} = \{(\hat{\mathbf{x}}, \hat{\mathbf{x}}')_i\}_{i=1}^N$, where $\hat{\mathbf{x}} \in X$ is a point and $\hat{\mathbf{x}}' = f(\hat{\mathbf{x}}, \mathbf{v})$ for a realization of \mathbf{v} , and initial states $\mathcal{D}_{\mathbf{x}_0} = \{(\hat{\mathbf{x}}_0)_i\}_{i=1}^{N_0}$ for $N, N_0 \in \mathbb{N}$.

We are interested in belief propagation, i.e., representing the (marginal) state distribution $p(\mathbf{x}_K)$, also referred to as the *belief* of the state at some future time $K \in \mathbb{N}$, and probabilistic reachability of a region of interest $R \subset X$ at time step K , i.e., $P(\mathbf{x}_K \in R)$, subject to starting at the initial belief $p(\mathbf{x}_0)$. Using the Markov chain \mathcal{M}_f , the belief can be propagated recursively for $k = 1, \dots, K$, using:

$$p(\mathbf{x}_k) = \int_{\mathbf{x}_{k-1}} p(\mathbf{x}_k \mid \mathbf{x}_{k-1}) p(\mathbf{x}_{k-1}) d\mathbf{x}_{k-1}. \quad (3)$$

Once $p(\mathbf{x}_K)$ is constructed, then, the probabilistic reachability is given by

$$P(\mathbf{x}_K \in R) = \int_R p(\mathbf{x}_K) d\mathbf{x}_K. \quad (4)$$

We refer to (4) as belief *evaluation*.

Challenges The analytical feasibility of the integral operations described by (3) and (4) is critically dependent on the coupling between the functional form of both distributions $p(\mathbf{x}_k)$ and $p(\mathbf{x}' | \mathbf{x})$. In fact, for most known systems that have non-linear dynamics, or even linear f but non-Gaussian $p(\mathbf{v})$, propagation via (3) is analytically infeasible, necessitating approximation.

Since both $p(\mathbf{x}_0)$ and $p(\mathbf{x}' | \mathbf{x})$ are unknown and must be learned, the following question arises: *Is there a functional form for $p(\mathbf{x}_k)$ and $p(\mathbf{x}_k | \mathbf{x}_{k-1})$ that can model general non-linear stochastic dynamics such that (3) and (4) become analytically tractable?* To rigorously differentiate such desirable models from linear-Gaussian models, or even non-linear models that assume *additive-Gaussian* noise, we define a universal distribution approximator.

Definition 1 (Universal Distribution Approximator). Let \mathcal{P} be the set of all continuous PDFs supported on a compact bounded support $\mathcal{X} \subset \mathbb{R}^n$, i.e., $p(\mathbf{x}) > 0$ for all $p \in \mathcal{P}$, $\mathbf{x} \in \mathcal{X}$. Let \mathcal{P}_θ be a family of distributions parameterized by a set of parameters θ where $|\theta| < \infty$. Then, under any probability divergence $\mathcal{D}(\cdot | \cdot)$, $p_\theta \in \mathcal{P}_\theta$ is said to be *universal* if for any $p \in \mathcal{P}$ and $\epsilon > 0$, if there exists a θ such that $\mathcal{D}(p | p_\theta) < \epsilon$.

Definition 1 can be easily extended to describe universal *conditional* PDF approximators by asserting that $p_\theta(\mathbf{x} | \mathbf{y})$ is universal for all \mathbf{y} in a compact bounded region \mathcal{Y} .

This work tackles the challenge of choosing a functional form for learning an arbitrary Markov chain \mathcal{M}_f subject to:

- C1. $p_\theta(\mathbf{x}_0)$ and $p_\theta(\mathbf{x}_k | \mathbf{x}_{k-1})$ are *universal*,
- C2. $p_\theta(\mathbf{x}_k)$ can be computed exactly via (3), and
- C3. $P(\mathbf{x}_k \in R)$ can be computed exactly via (4) for any hyper-rectangle $R \subset X$.

Hence, the belief propagation and evaluation problem that we consider is as follows.

Problem 1. Given datasets $\mathcal{D}_{\mathbf{x}_0}$ and $\mathcal{D}_{\mathbf{x}'}$ generated from the stochastic System (1) and a time horizon $K \in \mathbb{N}$, learn a Markov chain \mathcal{M}_f that adheres to conditions C1-C3, and compute the distribution $p(\mathbf{x}_K)$.

Approach Overview Ensuring that all three constraints hold poses a challenging functional representation problem. Recall that the recursive operation in (3) involves *multiplication* and *marginalization*, while (4) requires *definite integration*. Among common function classes, *polynomials* are one of the few that are closed under these operations and are also known for their strong approximation capabilities. We leverage these properties to address Problem 1.

In the remainder of the paper, all probability densities $p_\theta(\cdot)$ assumed to be parameterized models. For notational simplicity, we omit the subscript θ . All proofs are provided in the extended version.

3 Preliminaries

Our approach builds on two existing concepts, normalizing flows and Bernstein polynomials, which we review here.

Normalizing Flow Normalizing flows (Papamakarios et al. 2021) are a powerful and expressive parametric form for distribution modeling. In essence, a normalizing flow consists of two parts: i) a diffeomorphism (invertible and differential map) $g : X \rightarrow Z$ which maps from a “feature” space X to a “latent” space Z that is homeomorphic to X , and ii) a simple distribution $p_z(\mathbf{z})$ (often the standard Gaussian) over the latent space Z . In essence, g^{-1} transforms $p_z(\mathbf{z})$ into an arbitrary density $p_x(\mathbf{x})$ in the feature space X , and g “reverses” the transformation, allowing $p_x(\mathbf{x})$ to be expressed using the differential change in volume, i.e.,

$$p_x(\mathbf{x}) = p_z(g(\mathbf{x})) |\det(J_g(\mathbf{x}))|, \quad (5)$$

where $\det(J_g)$ is the determinant of the Jacobian of g .

For density estimation or generative modeling, g is typically parameterized with a universal *invertible* function approximator. In order to make g invertible *and* have a tractable Jacobian determinant, triangular maps are often used, i.e., $g_i(\mathbf{x}) = g_i(\mathbf{x}_{\leq i})$, simplifying (5) to

$$p_x(\mathbf{x}) = p_z(g(\mathbf{x})) \left| \prod_{i=1}^n \frac{\partial g_i}{\partial x_i}(\mathbf{x}_{\leq i}) \right|. \quad (6)$$

To ensure that the triangular map is invertible, g_i must be monotonic along the i -th dimension, i.e., $\frac{\partial g_i}{\partial x_i}(\mathbf{x}_{\leq i}) > 0$ for all $\mathbf{x}_{\leq i} \in \mathbb{R}^i$. Importantly, under mild conditions (Papamakarios et al. 2021), (6) can approximate any arbitrary feature distribution p_x as long as g_i can approximate any monotonic function along i .

Bernstein Polynomials Any multivariate polynomial $\pi(\mathbf{x})$ of degree \mathbf{d} can be equivalently expressed in the Bernstein basis as $\pi(\mathbf{x}) = \sum_{\mathbf{j}=0}^{\mathbf{d}} b_{\mathbf{j}} \phi_{\mathbf{j}}^{\mathbf{d}}(\mathbf{x})$, where $\phi_{\mathbf{j}}^{\mathbf{d}}(\mathbf{x}) = \prod_{i=1}^n \binom{d_i}{j_i} x_i^{j_i} (1 - x_i)^{d_i - j_i}$ are the Bernstein basis polynomials, and $b_{\mathbf{j}}$ are the coefficients. A coefficient set in the Bernstein basis is denoted with \mathbf{b} . The Bernstein basis allows one to easily bound the polynomial using its coefficients. We take advantage of this property in order to construct valid polynomial distribution models.

4 Polynomial Distribution Modeling

In this section, we present our modeling framework for $p(\mathbf{x}_0)$ and $p(\mathbf{x}_k | \mathbf{x}_{k-1})$. We aim to express both density functions as multivariate polynomials to preserve the feasibility of (3) and satisfy conditions C2 and C3. However, polynomials inherently cannot represent PDFs supported over unbounded domains such as \mathbb{R}^n , since a valid PDF must be a non-negative function that integrates to one over its domain, whereas any non-negative polynomial $\pi(\mathbf{x})$ yields $\int_{\mathbb{R}^n} \pi(\mathbf{x}) d\mathbf{x} = \infty$. To address this, we first map the unbounded support \mathbb{R}^n to a bounded domain and then define the polynomial PDF over this transformed space.

Transforming the State Space Let $\mathbb{U}^n := (0, 1)^n$ be the n -dimensional open unit box. To enable the use of polynomial density functions, we employ a mapping $\Omega : X \rightarrow \mathbb{U}^n$ to transform the original state space X into the unit box \mathbb{U}^n . As long as Ω is a diffeomorphism and diagonal, i.e., $\Omega_i(\mathbf{x}) =$

$\Omega_i(\mathbf{x}_i)$, densities over X can be equivalently expressed over \mathbb{U}^n using (5) as follows.

$$p_{\mathbf{x}}(\mathbf{x}) = p_{\mathbf{u}}(\Omega(\mathbf{x})) |\det J_{\Omega}(\mathbf{x})| \quad (7a)$$

$$= p_{\mathbf{u}}(\Omega(\mathbf{x})) \prod_{i=1}^n \left(\frac{d\Omega_i}{dx_i}(x_i) \right). \quad (7b)$$

In fact, a valid choice for each component Ω_i of Ω is a univariate cumulative distribution function (CDF) of a continuous distribution supported on \mathbb{R} , e.g., the Gaussian CDF, since the CDFs are monotonic and continuous, i.e., they are diffeomorphisms. The properties of the proposed models do not depend on the specific choice of Ω ; however, certain choices of Ω may allow the model to learn more accurate representations of the underlying Markov chain (see the extended version for details on how to choose Ω).

The properties of Ω preserve the integration property

$$\int_{I_i} p_{\mathbf{x}}(\mathbf{x}) dx_i = \int_{\Omega(I_i)} p_{\mathbf{u}}(\mathbf{u}) du_i, \quad (8)$$

over any interval I_i of an axis of X . By extension, marginalizing or integrating $p_{\mathbf{x}}(\mathbf{x})$ over hyper-rectangular regions $R \subseteq X$ is equivalent to marginalizing or integrating over its image $\Omega(R)$ in \mathbb{U}^n . Note that $\Omega(R)$ is guaranteed to also be hyper-rectangular since Ω is diagonal and monotone.

This allows us to express a polynomial density $p_{\mathbf{u}}(\mathbf{u}) = \pi(\mathbf{u})$, such that marginalizing over u_i reduces to simply integrating the polynomial $\pi(\mathbf{u})$ over the interval $[0, 1]$. Additionally, it is easy to verify that the product of any two \mathbb{U}^n -polynomial densities of different random variables, e.g., $p(\mathbf{u}_k | \mathbf{u}_{k-1}) p(\mathbf{u}_{k-1})$, as in (3), results in a polynomial and remains integrable. For the remainder of the paper, we drop the subscript when denoting densities, e.g., $p_{\mathbf{u}}(\mathbf{u}) = p(\mathbf{u})$.

Since both *multiplication* and *marginalization* operations in (3) (including integration in (4)) become simple polynomial operations of the equivalent \mathbb{U}^n -polynomial densities, we can learn the Markov chain \mathcal{M}_f directly in \mathbb{U}^n without imposing any restrictions on the generality of the underlying systems or state distributions. Therefore, all that remains is to specify the polynomial model parametric form for learning $p(\mathbf{u}_0)$ and $p(\mathbf{u}_k | \mathbf{u}_{k-1})$.

Polynomial Density Estimation Density estimation is a functional optimization problem with two constraints over the support: i) the integral of the function over the support must be 1, and ii) the function must be non-negative. Regarding (i), the class of all polynomials of a fixed degree \mathbf{d} that sum to 1 over \mathbb{U}^n , can be easily characterized as a linear constraint over coefficients (Yu and Loskot 2023).

Enforcing constraint (ii), however, is more challenging since even checking if a polynomial is non-negative over \mathbb{U}^n is known to be **NP-hard** (Murty and Kabadi 1985). To address this, we employ the Bernstein polynomial basis, which provides a *relaxation*, i.e., a sufficient (but not necessary) condition for bounding a polynomial (and its derivatives) over \mathbb{U}^n . We embed this relaxation within a normalizing flow architecture to ensure the necessary invertibility conditions for synthesizing valid polynomial densities.

Bernstein Normalizing Flow

We propose the Bernstein Normalizing Flow (BNF) for density estimation, particularly for learning $p(\mathbf{u}_0)$. Recall, triangular-map normalizing flows are an expressive parametric form for universal density estimation, which can be easily extended to universal *conditional* density estimation.

We aim to formulate a polynomial distribution over \mathbb{U}^n via the diffeomorphism $g : \mathbb{U}^n \rightarrow \mathbb{U}^n$ where the latent space is also the unit box. In order to ensure that g is a diffeomorphism, it suffices to show that g_i is monotonic in u_i and spans the full range, i.e., for all $\mathbf{u} = (u_1, \dots, u_n) \in \mathbb{U}^n$, $g_i(\mathbf{u}) = 0$ if $u_i = 0$, and $g_i(\mathbf{u}) = 1$ if $u_i = 1$.

Suppose each component g_i is a multivariate Bernstein polynomial $g_i(\mathbf{u}) = \pi_g^i(\mathbf{u}_{\leq i})$ of degree \mathbf{d} . By making the latent density uniform over \mathbb{U}^n , the normalizing flow density (5) simplifies to

$$p(\mathbf{u}) = \prod_{i=1}^n \frac{\partial \pi_g^i}{\partial u_i}(\mathbf{u}_{\leq i}), \quad (9)$$

which is simply a product of multivariate polynomials.

Hence, the two invertibility conditions of g , i.e., monotonicity and coverage of the full range, can be expressed in differential form as

$$\frac{\partial g_i}{\partial u_i}(\mathbf{u}_{\leq i}) > 0 \quad (10a)$$

$$\int_0^1 \frac{\partial g_i}{\partial u_i}(\mathbf{u}_{\leq i}) du_i = 1 \quad (10b)$$

for all $\mathbf{u}_{\leq i} \in \mathbb{U}^i$. To enforce these conditions on the Bernstein polynomial $\pi_g^i(\mathbf{u}_{\leq i})$, we can directly parameterize the partial derivative $\pi_{\partial g}^i(\mathbf{u}_{\leq i}) = (\partial \pi_g^i / \partial u_i)(\mathbf{u}_{\leq i})$ and enforce constraints (10) on the coefficients via the following lemmas:

Lemma 1 (Bernstein Relaxation (Lorentz 2012)). *A Bernstein polynomial $\pi(\mathbf{u})$ with coefficients $\mathbf{b}_{\mathbf{j}}$ is bounded on \mathbb{U}^n by the extrema of its coefficients, i.e., $\pi(\mathbf{u}) > \min_{\mathbf{j}} \mathbf{b}_{\mathbf{j}}$ and $\pi(\mathbf{u}) < \max_{\mathbf{j}} \mathbf{b}_{\mathbf{j}}$ for all $\mathbf{u} \in \mathbb{U}^n$.*

Lemma 2. *Let $\pi(\mathbf{u})$ be a Bernstein polynomial of degree $\mathbf{d} = (d_1, \dots, d_n)$ with coefficients $\mathbf{b}_{\mathbf{j}} = (b_{j_1}, \dots, b_{j_n})$. Then,*

$$\int_0^1 \pi(\mathbf{u}) du_i = 1 \iff \sum_{j_i=0}^{d_i} b_{j_i} = d_i + 1 \quad \forall j_i. \quad (11)$$

Let $\tilde{\mathbf{b}}^i$ be the coefficients of $\pi_{\partial g}^i$, which has degree $\tilde{\mathbf{d}}^i = (d_1, \dots, d_i - 1, \dots, d_n)$. Then, condition (10a) can be easily enforced by constraining $\tilde{\mathbf{b}}_{\mathbf{j}}^i \geq 0$ for all $\mathbf{0} \leq \mathbf{j} \leq \tilde{\mathbf{d}}^i$, and condition (10b) can be enforced by ensuring that all $\tilde{\mathbf{b}}_{\mathbf{j}}^i$ sum to d_i along the i -th axis. With both constraints, polynomials π_g^i correctly parameterize a diffeomorphism, and therefore $p(\mathbf{u})$ is a valid density.

Remark 1. *To achieve more expressivity, multiple diffeomorphism “layers” g can be composed. Doing so, however, raises the degree multiplicatively of the fully-expanded polynomial (which is necessary for belief propagation and evaluation).*

Conditional Density Estimation

To parameterize a conditional distribution $p(\mathbf{u} \mid \mathbf{w})$ for $\mathbf{w} \in \mathbb{U}^n$, we can construct a flow function $h : \mathbb{U}^n \times \mathbb{U}^n \rightarrow \mathbb{U}^n$, such that given any \mathbf{w} , $h(\mathbf{u}, \mathbf{w})$ is a diffeomorphism in \mathbf{u} . The differential constraints (10a) and (10b) can be modified accordingly by adding \mathbf{w} as a parameter. Making each component h_i of h a polynomial $h_i(\mathbf{u}, \mathbf{w}) = \pi_h^i(\mathbf{u}_{\leq i}, \mathbf{w})$ of degree $\mathbf{D} = (d_1, \dots, d_n, d_1, \dots, d_n)$, yields a polynomial conditional distribution of the form

$$p(\mathbf{u} \mid \mathbf{w}) = \prod_{i=1}^n \frac{\partial \pi_h^i}{\partial \mathbf{u}_i}(\mathbf{u}_{\leq i}, \mathbf{w}). \quad (12)$$

By ensuring that the coefficients of each π_h^i are all non-negative, and sum to $d_i - 1$ along the i -th axes, (12) is a valid conditional distribution.

Theorem 1. *The conditional Bernstein normalizing flow is a universal conditional distribution approximator.*

Belief Propagation

By modeling $p(\mathbf{u}_0)$ as a BNF and $p(\mathbf{u}' \mid \mathbf{u})$ as a conditional BNF, the operations in (3) and (4) can be carried out exactly using tensor operations on the coefficients of each model. Additionally, both multiplication and integration operations can be carried out in the Bernstein basis, improving the numerical stability of each prediction (Farouki and Rajan 1987). Since belief propagation is exact with respect to the model, all prediction error can be attributed to modeling/learning error (and floating-point numerical inaccuracies).

The following theorem states another crucial consequence of using BNF to model \mathcal{M}_f for belief propagation.

Theorem 2. *Given a state-transition conditional-BNF $p(\mathbf{u}' \mid \mathbf{u})$ that is of degree \mathbf{d}' in \mathbf{u}' , starting from any valid polynomial belief $p(\mathbf{u}_0)$, $p(\mathbf{u}_k) \in \Pi^{\mathbf{d}'}[\mathbf{u}]$ for every $k \in \mathbb{N}$.*

Theorem 2 has computational implications for long-horizon prediction. Namely, in each propagation recursion (3), the amount of memory needed to store each belief $p(\mathbf{u}_k)$ does not grow in time, unlike methods such as GMM-splitting (Kulik and LeGrand 2024). Additionally, the computation needed for $p(\mathbf{u}_k)$ is $O(k)$. Note that number of parameters in BNF scales with the number of coefficients in a degree- \mathbf{d} $2n$ -variate (for \mathbf{u} and \mathbf{u}') Bernstein polynomial, i.e. $O(d^{2n})$, which is memory intensive for large n . For scalability purposes, we recognize sparse models as valuable future work.

The proposed BNF methods are not limited to just belief propagation and can be trivially extended to other applications as discussed in the following remark.

Remark 2 (Bayesian Belief Update Analog). *The proposed methods can be employed for exact Bayesian belief updating with learned likelihood and prior models, i.e., $p(\alpha \mid \beta) = p_\theta(\beta \mid \alpha)p_\theta(\alpha)/p(\beta)$. Computing the normalization constant $p(\beta)$ needed to calculate the posterior belief $p(\alpha \mid \beta)$ often results in an infeasible integral for complex likelihood models. Nonetheless, the multiplication-normalization operation sequence closely parallels the propagation operation described in (3), and thus the models presented in this paper can be readily applied to exact Bayesian Belief updating.*

Remark 3. *The normalizing flow structure allows one to easily sample from both of the learned distributions. A latent state $\hat{\mathbf{z}}$ is sampled uniformly over \mathbb{U}^n , then mapped through the inverse of g or h , i.e., $\hat{\mathbf{u}} = g^{-1}(\hat{\mathbf{z}})$ or $\hat{\mathbf{u}} = g^{-1}(\hat{\mathbf{z}}; \mathbf{w})$.*

5 Training Procedure

In this section, we outline the procedure for training each piece of the Markov chain. Specifically, we show how the Bernstein relaxation constraints described in Sec. 4 can be enforced during training. Additionally, we describe a procedure that can be implemented during training to tighten the relaxations and increase the expressivity of the the model without increasing the number of parameters.

Constrained Log-Likelihood Optimization

Maximum Likelihood Estimation (MLE) is a well established objective for density estimation (Pan and Fang 2002). We aim to maximize the likelihood BNF given state space data $D_{\mathbf{x}'}$ and $D_{\mathbf{x}_0}$. MLE of a \mathbf{x} -density with \mathbf{x} -data is equivalent to MLE of a \mathbf{u} -density against the data mapped to \mathbb{U}^n via Ω . The constrained MLE optimization problem over the parameters of each BNF model for learning $p(\mathbf{u}_0)$ is:

$$\arg \max_{\mathbf{b} \in \mathbf{B}} \mathbb{E}_{\mathbf{u}_0 \sim p^*(\mathbf{u}_0)} [\log p(\mathbf{u}_0)] \quad (13a)$$

$$\text{subject to } \tilde{\mathbf{b}}^i \geq 0 \quad (13b)$$

$$\sum_{j_i=0}^{d_i-1} \tilde{\mathbf{b}}^i = d_i \quad \forall 1 \leq i \leq n \quad (13c)$$

where $p^*(\cdot)$ denotes the true distribution, and $\tilde{\mathbf{B}} \ni \mathbf{b}$ denotes the parameter space. Expectations with respect to p^* are empirically evaluated over the data in $\mathcal{D}_{\mathbf{x}_0}$. A similar optimization problem can be formulated for training the transition distribution by using the empirical expectation of $\log p(\mathbf{u}' \mid \mathbf{u})$ over the data $\mathbf{u}', \mathbf{u} \sim p^*(\mathbf{u}', \mathbf{u})$ where $p^*(\mathbf{u}', \mathbf{u})$ denotes the true joint distribution that generated $\mathcal{D}_{\mathbf{x}'}$. Note that, even though the constraints (13b) and (13c) are linear, the objective (13a) is non-linear and non-convex. Thus, we leverage stochastic gradient descent (SGD) for optimization.

In order to avoid explicitly constraining the parameter space, we can define a differentiable function $\Psi : \Theta \rightarrow \tilde{\mathbf{B}}_{\text{feas}}$ where Θ represents an “unconstrained” parameter space, and $\tilde{\mathbf{B}}_{\text{feas}} \subset \tilde{\mathbf{B}}$ is the set of feasible parameters, i.e., the set of all $\tilde{\mathbf{b}}$ that satisfy (13b) and (13c). With Ψ , SGD can be performed in Θ instead of \mathbf{B} , enforcing the constraints implicitly. We can define Ψ for a given coefficient tensor \mathbf{b}^i as the composition of two operations: 1) map each vector to be positive, and 2) normalize the coefficients to sum to d_i along dimension i . Formally, we can construct Ψ as $\tilde{\mathbf{b}}^i = \Psi_i(\theta^i) = (\sigma \circ \delta)(\theta^i)$, where δ is a positive-range element-wise function (e.g., soft-plus) and σ is the normalizing function

$$\sigma(\theta^i) = \theta_i / \sum_{j_i=0}^{d_i-1} \theta_j^i. \quad (14)$$

Each operation is differentiable; thus, models parameterized by $\theta_0, \dots, \theta_{d^r-1}$ can be trained with standard SGD are guaranteed to satisfy the constraints.

Tightening the Relaxation

We have thus far ignored an important aspect inherent to the BNF that can significantly affect the models’ expressiveness: the tightness of the Bernstein relaxation in Lemma 1. Constraining coefficients $\mathbf{b} > 0$ is only a *sufficient* (but not necessary) condition for $\pi_{\partial g}^i \geq 0$. Thus, $\tilde{\mathbf{B}}_{\text{feas}}$ inner-approximates the set of all degree- \mathbf{d} polynomial diffeomorphisms, and limits the expressiveness of the model. Increasing the degree of each π_g^i improves the expressiveness, at the cost of adding many more parameters to the model. We propose an alternative method to tighten the relaxation during the training procedure *without* adding any parameters to the model.

The proposed procedure lifts a given Bernstein polynomial to a higher degree to reason over feasibility. Any Bernstein polynomial of degree \mathbf{d} can be equivalently expressed in a Bernstein basis of a higher degree $\mathbf{d}^+ \geq \mathbf{d}$ (Lorentz 2012). For simplicity of presentation, we describe the following assuming a one-dimensional Bernstein polynomial; however, all operations readily extend to the multivariate case. A Bernstein basis polynomial $\phi_j^{\mathbf{d}}(x)$ can be written as a linear combination of higher-degree basis polynomials:

$$\phi_j^{\mathbf{d}}(x) = \frac{\mathbf{d} + 1 - j}{\mathbf{d} + 1} \phi_j^{\mathbf{d}^+}(x) + \frac{j + 1}{\mathbf{d} + 1} \phi_{j+1}^{\mathbf{d}^+}(x). \quad (15)$$

One can apply (15) recursively $\mathbf{d}^+ - \mathbf{d}$ times to build a linear transformation of the original coefficients, in the form $\mathbf{b}^+ = \mathbf{M}_d^{\mathbf{d}^+} \mathbf{b}$ where \mathbf{b}^+ is the degree- \mathbf{d}^+ coefficient vector of the same polynomial represented by \mathbf{b} , and $\mathbf{M}_d^{\mathbf{d}^+}$ is a $(\mathbf{d}^+ + 1) \times (\mathbf{d} + 1)$ tall matrix.

Theorem 3 ((Garloff 1985)). *Given a Bernstein polynomial $\pi(\mathbf{u})$ of degree \mathbf{d} , let $\underline{\pi} = \inf_{\mathbf{u} \in \mathbb{U}^n} \pi(\mathbf{u})$ and $\underline{b} = \min_j \mathbf{b}$. Then, for $\mathbf{d}^+ \geq \mathbf{d}$, $\underline{b} \leq \underline{\pi} \leq \underline{b} + O(\mathbf{d}^{-1})$ and \underline{b} converges to $\underline{\pi}$ monotonically.*

Theorem 3 illustrates the convergence of the Bernstein relaxation as the raised-degree is increased.

Constraining degree- \mathbf{d} polynomials by the degree \mathbf{d}^+ representation inherently captures *more* non-negative degree- \mathbf{d} polynomials, and as $\mathbf{d}^+ \rightarrow \infty$, the raised degree constraint captures all non-negative polynomials. In other words, there may exist a non-negative degree- \mathbf{d} polynomial with coefficients \mathbf{b} that minimizes (13a), but violates constraint (13b) due to the conservatism of the relaxation.

Employing degree-raising during training is not as straightforward as using implicit constraints, as discussed in the previous section. To illustrate the challenge, consider the following sequence of operations within a single SGD iteration: 1) Raise the degree of the current (unconstrained) parameters via $\mathbf{b}^+ = \mathbf{M}_d^{\mathbf{d}^+} \mathbf{b}$. 2) Replace all negative entries with zero. 3) Project the rectified parameters back down to the degree- \mathbf{d} representation. 4) Normalize the parameters using (14). Unfortunately, even though the rectified $\mathbf{b}^+ \geq 0$ after step 2, the projection does not necessarily preserve this property.

To address this issue, we train the model with a soft-constraint violation penalty. Then, using an iterative projection, the parameters are moved into the feasible region. The constraint violation penalty loss can be formulated as

$\sum_j \max(0, \mathbf{b}^+)$, which is easily calculated with matrix multiplication. Then, to apply a hard constraint to the resulting parameters \mathbf{b} , we iteratively repeat the projection steps 1-3 until $\mathbf{b}^+ \geq 0$. This training procedure avoids cumbersome feasible-set projections during optimization, while still yielding a feasible solution.

6 Evaluations

This section evaluates the efficacy of learned BNF Markov chain models for uncertainty propagation and the effect of the choice of degree on the models’ expressiveness. Full details on system parameters, experimental setup, and additional visualizations are provided in the extended version.

Comparison Methods We compare BNF to three uncertainty propagation methods based on GMM belief representations: (i) first-order linearization around component means (EKF-style) (Jazwinski 2013), (ii) Whitened Spherical Average Second-Order Stretching (WSASOS) (Kulik and LeGrand 2024), and (iii) a grid-based approach from (Figueiredo et al. 2024). EKF and WSASOS assume nonlinear dynamics with additive Gaussian noise. The grid-based method assumes a diagonal-covariance GMM for $p(\mathbf{x}'|\mathbf{x})$, but lacks a procedure for constructing a formal GMM approximation from general nonlinear models, so we evaluate it assuming a single-component GMM.

All three methods rely on Gaussian assumptions, so we used GP regression to learn $p(\mathbf{x}' | \mathbf{x})$ and expectation maximization on a GMM to learn $p(\mathbf{x}_0)$. EKF serves as a widely used baseline, WSASOS as a state-of-the-art component-splitting method, and the grid-based approach as a formal method with guaranteed error bounds.

Experimental Setup Experiments were performed on two highly non-linear latent stochastic systems: Van der Pol with additive Gaussian noise and a stable oscillator with multiplicative non-Gaussian noise. The latent initial state distribution is a Gaussian distribution. For each GMM method, $p(\mathbf{x}_0)$ is learned with 10 mixture components to avoid overfitting to the initial state data. The resolution of the grid method is 20-by-20. Each comparison uses a Multitask Gaussian Process Regression with a radial basis function kernel. For BNF, we used Bernstein polynomials with degrees 10, 20, and 30 for each component g_i . Each method is trained with 1K initial state data points 10K state-transition data points. Prediction accuracy is evaluated using average log-likelihood (13a) on a distinct large test data set, consisting of Monte-Carlo samples from the true system. Each experiment was performed 10 times with random training seeds, and all log-likelihood results were found to have a variance below 10^{-3} .

Fig. 1 shows a visual comparison of the computed beliefs at time step $k = 9$ for the Van der Pol system (Fig. 1a–d) and the stable oscillator (Fig. 1e–h). Figs. 2–3 present the average log-likelihood results. Overall, we observe that BNF performs comparably to the grid-based method in the additive Gaussian case and significantly outperforms all baselines under non-Gaussian noise. Additionally, performance improves with increasing polynomial degree.




Figure 1: Visual Comparison for the final belief propagated $k = 9$ timesteps into the future for the (a)-(d) Van der Pol and (e)-(h) stable oscillator systems.




Figure 2: **Additive-Gaussian Noise** - Belief accuracy

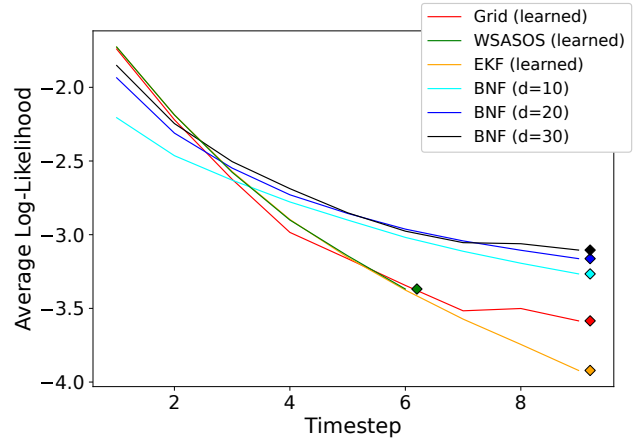


Figure 3: **Multipli. Non-Gaussian Noise** - Belief accuracy

Additive Gaussian Noise System As can be seen in Fig. 2, BNF generates a more accurate prediction than EKF, however, the grid method performs the best. Since the system is additive Gaussian, GP regression learns a *very* accurate model. The grid method produces the least propagation error among the GMM methods, and is able to achieve a very accurate prediction. WSASOS (learned) times out after 6 time steps since the number of components in the GMM explodes exponentially.

To isolate *propagation error* from *learning error*, we compared BNF to baseline methods using the true underlying system dynamics, effectively isolating only the error accumulated via propagation. Alternatively, since BNF propagation is exact, the induced error is *only* learning error. As can be seen, BNF performs comparably despite starting from a more inaccurate initial distribution.

In BNF, the effect of the degree of the polynomial is most notably observed the estimate of the initial distribution. The initial distribution has relatively small covariance, thus, the polynomial must be able to attain a steep transition, which requires polynomials of higher degree.

Multiplicative Non-Gaussian Noise System Fig. 3 illustrates the universal modeling capability of BNFs. Since the true underlying system has highly non-linear non-Gaussian noise, the GP regression models become unable to truly capture the stochastic nature of the system, despite having ample data. Consequently, unlike the additive-Gaussian case, the GMM methods now incur significant propagation *and* learn-

ing error. Conversely, BNF is able to more accurately capture the complex stochasticity in the true system (visually seen in Fig. 1 (f-g)). In particular, the inaccuracy in the stochastic model for GP-based collapses the belief over the stable points of the system, whereas BNF maintains a better (more uncertain) representation of the belief. BNF exhibits a consistent trend with the previous experiment when varying the degree. This experiment demonstrates the power of shifting all prediction error to only learning error: prediction accuracy is solely a function of the representational capability of the model and the methods used for training.

These experiments illustrate the power of BNF in learning and propagating PDFs of complex, realistic systems. Unlike baseline methods, BNF maintains high predictive fidelity even under non-Gaussian noise. Given BNF’s exact propagation property, its sole error is rooted in learning, which can be mitigated using higher polynomial degrees.

7 Conclusion

In this work, we introduced a Bernstein Normalizing Flows (BNFs), a novel model class that enables universal PDF approximation and exact belief propagation by combining Bernstein polynomials with normalizing flows. BNFs eliminate approximation error while supporting efficient sampling and evaluation. We provided the theoretical foundations for this approach and empirically demonstrate that BNFs are particularly effective in nonlinear, non-Gaussian settings, where state-of-the-art methods suffer.

References

- Akhiezer, N. I. 2020. *The classical moment problem and some related questions in analysis*. SIAM.
- Alspach, D.; and Sorenson, H. 2003. Nonlinear Bayesian estimation using Gaussian sum approximations. *IEEE transactions on automatic control*, 17(4): 439–448.
- Arpogaus, M.; Kneib, T.; Nagler, T.; and Rügamer, D. 2025. Hybrid Bernstein Normalizing Flows for Flexible Multivariate Density Regression with Interpretable Marginals. *arXiv preprint arXiv:2505.14164*.
- Babu, G. J.; Canty, A. J.; and Chaubey, Y. P. 2002. Application of Bernstein polynomials for smooth estimation of a distribution and density function. *Journal of Statistical Planning and Inference*, 105(2): 377–392.
- Belalia, M.; Bouezmarni, T.; and Leblanc, A. 2017. Smooth conditional distribution estimators using Bernstein polynomials. *Computational Statistics & Data Analysis*, 111: 166–182.
- Deisenroth, M. P.; Turner, R. D.; Huber, M. F.; Hanebeck, U. D.; and Rasmussen, C. E. 2011. Robust filtering and smoothing with Gaussian processes. *IEEE Transactions on Automatic Control*, 57(7): 1865–1871.
- Djuric, P. M.; Kotecha, J. H.; Zhang, J.; Huang, Y.; Ghirmai, T.; Bugallo, M. F.; and Miguez, J. 2003. Particle filtering. *IEEE signal processing magazine*, 20(5): 19–38.
- Farouki, R. T.; and Rajan, V. 1987. On the numerical condition of polynomials in Bernstein form. *Computer Aided Geometric Design*, 4(3): 191–216.
- Figueiredo, E.; Patane, A.; Lahijanian, M.; and Laurenti, L. 2024. Uncertainty propagation in stochastic systems via mixture models with error quantification. In *2024 IEEE 63rd Conference on Decision and Control (CDC)*, 328–335. IEEE.
- Garloff, J. 1985. Convergent bounds for the range of multivariate polynomials. In *Int. Symp. on Interval Mathematics*, 37–56. Springer.
- Jasour, A.; Wang, A.; and Williams, B. C. 2021. Moment-based exact uncertainty propagation through nonlinear stochastic autonomous systems. *arXiv preprint arXiv:2101.12490*.
- Jazwinski, A. H. 2013. *Stochastic processes and filtering theory*. Courier Corporation.
- Julier, S. J.; and Uhlmann, J. K. 2004. Unscented filtering and nonlinear estimation. *Proceedings of the IEEE*, 92(3): 401–422.
- Kulik, J.; and LeGrand, K. A. 2024. Nonlinearity and uncertainty informed moment-matching Gaussian mixture splitting. *arXiv preprint arXiv:2412.00343*.
- Lorentz, G. G. 2012. *Bernstein polynomials*. American Mathematical Soc.
- Murty, K. G.; and Kabadi, S. N. 1985. Some NP-complete problems in quadratic and nonlinear programming. Technical report.
- Pan, J.-X.; and Fang, K.-T. 2002. Maximum likelihood estimation. In *Growth curve models and statistical diagnostics*, 77–158. Springer.
- Papamakarios, G.; Nalisnick, E.; Rezende, D. J.; Mohamed, S.; and Lakshminarayanan, B. 2021. Normalizing flows for probabilistic modeling and inference. *Journal of Machine Learning Research*, 22(57): 1–64.
- Polymenakos, K.; Laurenti, L.; Patane, A.; Calliess, J.-P.; Cardelli, L.; Kwiatkowska, M.; Abate, A.; and Roberts, S. 2020. Safety guarantees for iterative predictions with Gaussian processes. In *2020 59th IEEE Conference on Decision and Control (CDC)*, 3187–3193. IEEE.
- Ramasinghe, S.; Fernando, K.; Khan, S.; and Barnes, N. 2022. Robust normalizing flows using Bernstein-type polynomials. *33rd British Machine Vision Conference Proceedings*.
- Schei, T. S. 1997. A finite-difference method for linearization in nonlinear estimation algorithms. *Automatica*, 33(11): 2053–2058.
- Yu, Y.; and Loskot, P. 2023. Polynomial distributions and transformations. *Mathematics*, 11(4): 985.